Analysis of Cascading Failures in Power Networks

by

Irina V. Abarinov

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degrees of

Bachelor of Science in Electrical Science and Engineering

and

Master of Engineering in Electrical Engineering and Computer Science
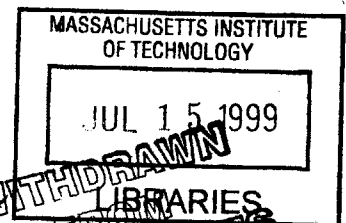
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1999

© 1999 Irina V. Abarinov. All rights reserved.

The author hereby grants to M.I.T. permission to reproduce and distribute publicly paper and electronic copies of this thesis and to grant others the right to do so.

Author . . . .
        Department of Electrical Engineering and Computer Science
                                              May 18, 1999

Certified by. . . . .
                                              George C. Verghese
                                    Professor of Electrical Engineering
                                              Thesis Supervisor

Accepted by . . . .
                                              Arthur C. Smith
        Chairman, Department Committee on Graduate Students

Analysis of Cascading Failures in Power Networks

by

Irina V. Abarinov

Abstract

This thesis studies the characteristics of cascading failures in a simulated linear swing-equation model for simple power networks. The model transitions from displaying only small average failure size to undergoing large average failure size as the loading of the system is increased. For some parameter values, the transition region becomes critical, and it exhibits failures of all sizes whose distribution is approximated by power-laws. Comparisons of the average failure size are made for networks with different generator/load configurations, transmission line densities, system sizes, frequency deviation thresholds, and angle thresholds. Connections with the literature on power-law distributions and criticality are discussed.

Thesis Supervisor: George C. Verghese
Title: Professor of Electrical Engineering

2

Acknowledgments

I would like to thank my thesis supervisor, George Verghese, for providing inspiration, guidance, and support throughout the project, as well as for giving me an opportunity to be a teaching assistant in his course. I am grateful to Bernard Lesieutre and Chalee Asavathiratham for their valuable comments on the thesis work during our weekly meetings. I would also like to thank my academic advisor, Gerald Wilson, for his advice and encouragement throughout my years at MIT.

I am particularly grateful to Mikhail Ulinich for providing me with computational resources for the project, as well as for the technical support that came with those resources. I would like to thank him for his patience throughout the project and his valuable comments and corrections on the draft of this thesis.

Contents

4

List of Figures

List of Tables

Chapter 1

Introduction

The U.S. power grid is one of the largest and most intricate systems in the world, and its continuous operation without interruptions, such as instabilities, network-wide power oscillations, or major cascading failures, is a strong necessity. The complexity of power systems, however, makes the task of planning, management, and analysis very difficult, especially as deregulation of the power industry is starting to take place. The development of broad interconnections in the power networks was motivated, among other things, by the desire for reliability, since loads in an interconnected system are not dependent on any one particular generator. Recent outages, however, such as the event that took place on the West Coast on August 10, 1996, which affected about 7.5 million customers and cost billions of dollars, demonstrate that global connectivity can lead to the multiplication of failures, making reliability a more subtle issue.

The graph in Figure 1-1 shows the distribution of the size of power outages measured by the number of affected customers in the U.S. electric grid for the years of 1984 through 1997, from the data in [4]. The log of the frequency per year of failures greater than N is plotted against $\log(N)$. The striking feature of this figure is that the distribution is not Gaussian, but can be characterized by a power law with "heavy tails". This means that the number of large, costly, cascading failures is not an anomaly, but a regular occurrence in the power network. The graph indicates that outages affecting 5 million people occur every 10 years, and failures affecting millions happen yearly. In addition, the deregulation of the industry has the potential
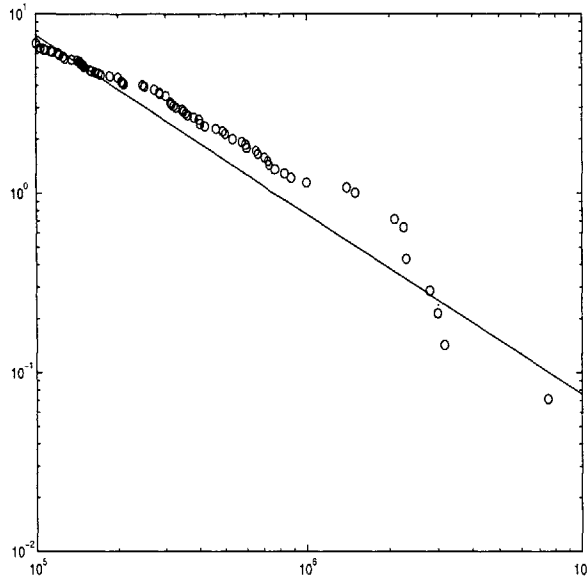
Figure 1-1: Frequency of outages vs. number of customers affected for U.S. power grid, 1984-1997.

of making the situation even worse, as power companies start to use more complicated management techniques to meet the increasing demand and looser cooperation requirements. Thus, there is a growing need for understanding the dynamics and characteristics of cascading fault events in power systems.

The presence of power law distributions is a common statistical feature of many complex interactive systems, such as species extinction, natural catastrophes, and traffic patterns [1, 2, 6]. In these systems, interconnectedness tends to increase the probability that many components of the system fail at once, leading to frequent occurence of catastrophic events. This contrasts sharply with a Gaussian or an exponential distribution, where the occurence of an event larger than some characteristic size is highly unlikely.

Power law distributions in connected systems have been studied widely in the context of percolating networks, where power laws arise at specific parameter values known as critical points [5, 12, 10]. Critical point behavior and power law distributions can also develop spontaneously, without any parameter tuning, in certain many-body systems through the mechanism known as *self-organized criticality* [6]. Recent work in this area, however, has suggested that power laws appear in complex designed systems

10

through a different process, termed *highly-optimized tolerance* [3], which is due to the tradeoffs associated with providing robust performance in uncertain environment.

The work presented in this thesis studies the dynamics and characteristics of cascading failures in power networks through computer-based simulations of simple models. The thesis also discusses the connections between the latest literature on power-law distributions in complex networks and the problem of propagating fault events in power systems.

The remainder of this thesis is organized into 5 chapters. Chapter 2 reviews the concepts of percolation, self-organized criticality, and highly optimized tolerance, and discusses their application to the power network setting. The model of power systems which is used to study cascading failures in the thesis is introduced in Chapter 3. Chapter 4 describes the methods and the computer program used to simulate the model of the network, as well as the typical behavior of the simulated systems. Observed cascading failures are discussed in Chapter 5. Finally, Chapter 6 summarizes the thesis and presents some suggestions for future work.

Chapter 2

Literature Review

Power law distributions in interconnected systems are associated with *criticality* - a behavior that arises when a system that undergoes a continuous transformation from one state to another passes through a point where events of all sizes are possible. In the critical state, all members of the system influence each other, so interactions stretch across the entire system.

This chapter reviews some prototypes of systems that exhibit critical behavior and power law distributions. The first section discusses the theory of percolation, which deals with critical behavior of equilibrium systems. In such systems the critical point occurs only at certain parameter values, and careful tuning of the system is necessary to obtain criticality. The second section deals with nonequilibrium systems and the notion of self-organized criticality (SOC), which claims that certain systems will drive themselves into critical states. SOC has been found useful in modeling a broad class of natural systems, such as sandpiles and earthquakes [6]. In the third section, another mechanism for generating power laws in interconnected systems, called highly-optimized tolerance (HOT), is presented. HOT applies to evolved or engineered systems that have to provide high performance in uncertain environmental conditions. Finally, the last section formulates the problem of cascading failures in power systems in the context that is used to study percolation, SOC, and HOT, and discusses similarities and differences among these systems.
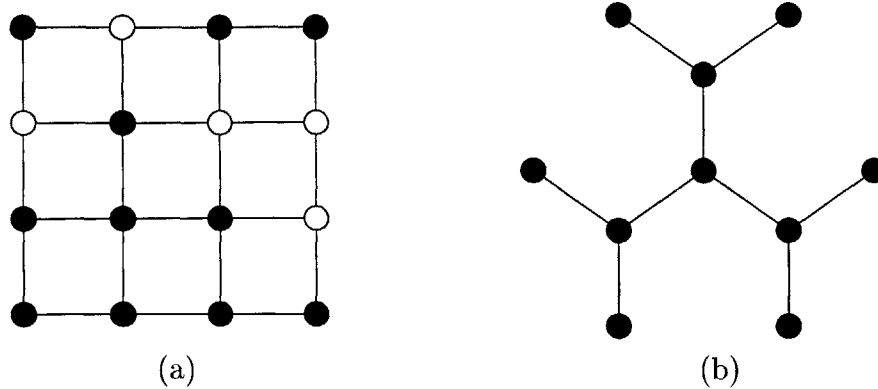
Figure 2-1: Percolation examples: a) square lattice and b) Bethe lattice (both extend indefinitely beyond the portions shown).

2.1 Percolation

The idea of percolation can be demonstrated on an infinite square lattice such as the one shown in Figure 2-1 (a). Each node of the lattice can be in one of two states: the node can be occupied with probability p or it can be vacant with probability $1 - p$. The probability that a node is occupied is independent of the state of other nodes. A group of nodes consisting of occupied neighbors (occupied nodes connected by an edge) forms a cluster, such as the cluster of size 8 in the bottom left corner of Figure 2-1 (a). Percolation theory deals with the number and properties of these clusters [12].

As the probability p of a square being occupied is increased, the clusters become larger in size on average. At some critical probability p_c, called the percolation threshold, a cluster of infinite size forms in the lattice. That is, for all $p > p_c$ there exists with unit probability an infinite cluster extending from one edge of the lattice to another, while for all $p < p_c$ the probability of such a cluster is zero [12]. This is an example of a criticality phenomenon.

The percolation threshold depends on the topography of the lattice and can be derived analytically for a few very special configurations. An example of a lattice, known as the Bethe lattice, for which p_c can be determined analytically is shown in Figure 2-1 (b). The Bethe lattice starts with a central node (origin) having z bonds

13

that lead to other nodes. The structure in Figure 2-1 (b), for example, has $z = 3$. These nodes (first generation) also have z bonds, one of them leading to the origin, and the other $z - 1$ bonds ending in new nodes. This branching process continues until the surface of the lattice is reached; the surface nodes have only one bond connecting them to the interior of the lattice.

To find the percolation threshold for the Bethe lattice, consider a structure with n generations. Assume that the origin is occupied, and check if there is a chance of finding a path that connects the central node to the surface by computing the average number of paths that lead there. Since there are no closed loops in the Bethe lattice, all paths to the surface exist with the same probability p^n, which is equal to the probability that every node in the path is occupied. Moreover, the surface contains $z(z-1)^{n-1}$ nodes, so there are that many paths leading to the surface. This gives us

$$E\{number\ of\ paths\} = z(z-1)^{n-1}p^n. \tag{2.1}$$

As n goes to infinity, Equation 2.1 goes to 0 if $(z-1)p$ is less than 1, so an infinite path does not exist. However, if $(z-1)p$ is greater than one, the average number of paths becomes infinite, which means that an infinite cluster exists with probability 1. Therefore, the percolation threshold p_c is

$$p_c = \frac{1}{1-z}.$$

For other lattice configurations, the percolation threshold can be estimated by simulating the system on a computer. The percolation threshold for the square lattice, for instance, was computed to be 0.592746 [12].

Percolation theory is relevant to the study of power law distributions, since near the percolation threshold quantities of interest are related to the occupancy probability p through power laws. For example, the normalized cluster number n_s, defined as the number of clusters consisting of s sites divided by the total number of nodes in the lattice, has been shown through numerical simulations to take the form:

14

$$n_s(p) = s^{-\tau} f[(p - p_c)s^{\sigma}] \qquad (p \to p_c, s \to \infty)$$

[12]. The coefficients τ and σ, as well as the precise form of the scaling function $f(z)$, have to be determined numerically for various lattice structures. However, it turns out that in almost all cases $f(z)$ approaches a constant value when $|z| \ll 1$ and decays very rapidly for $|z| \gg 1$, so $n_s(p)$ behaves according to a power law at the critical point.

Another example of a power law distribution is the mean cluster size S which scales as

$$S \propto |p - p_c|^{-\gamma}$$

as p approaches p_c. The exponent γ was calculated to be 1 in the Bethe lattice, and it is equal to 43/18 in the square lattice [12]. Another key feature of the critical state are the fractal shapes of the clusters at the percolation threshold. Away from the percolation threshold, statistical properties of the clusters scale exponentially for fixed p as $s \to \infty$:

$$\log n_s(p < p_c) \quad \propto \quad -s$$
$$\log n_s(p > P_c) \quad \propto \quad -s^{1-1/d}$$

For a lattice that has finite size, the appearance of a cluster which spans the whole lattice will not be as sudden as for infinite systems. Figure 2-2 shows the schematic graph of the probability Π that a lattice of linear size N contains a cluster extending from the top edge to the bottom edge. The transition region from $\Pi = 0$ to $\Pi = 1$ narrows as the size of the lattice is increased. For the finite size lattice, we can define an average occupation density p_{av} at which a percolation cluster appears for the first time:

$$p_{av} = \int p(\frac{d\Pi}{dp})dp$$

Numerical studies show that, as the size of the lattice increases, the effective perco-
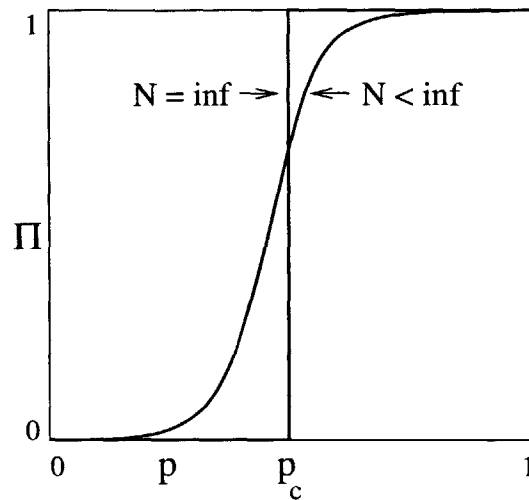
Figure 2-2: Variation of the probability Π that a cluster spans the whole system, for large and medium system sizes [12].

lation threshold p_{av} approaches the asymptotic value p_c as

$$p_{av} - p_c \propto N^{-1/\nu}$$

The parameter ν has been determined to be 4/3 for the square lattice and 1/2 for the Bethe lattice [12].

There is a large body of literature that deals with applications of percolation theory to different physical systems, such as oil fields, diffusion in disordered media, and conductance of random electrical networks [12]. The spread of forest fires, for example, can be examined in the context of percolation [10]. Imagine a forest that has been divided into sections to form a square lattice where each node is occupied by a tree with a certain occupation probability p. A randomly chosen tree is set on fire by an agent such as a lightning strike. If the fire can spread from a burning tree to its neighbors, how much of the forest will burn down? Whether the fire will be contained to a small group of trees or engulf the whole forest depends on the probability p, and the question can be answered by considering this as a percolation problem on a square lattice, as discussed above.

2.2 Power Laws and Self-Organized Criticality

Although power law relationships are present in percolating systems, the critical point which exhibits these power laws can only be reached through careful "tuning" of the occupation probability to equal the percolation threshold. Power laws, however, are observed in many natural and man-made phenomena without any such tuning taking place. Another theory is necessary to explain how power law behavior can develop spontaneously in real systems.

In 1987, Bak and coworkers introduced the concept of self-organized criticality (SOC), which shows how dynamic instabilities in systems drive them to a critical state [6]. The SOC claims that systems consisting of many interacting parts can organize themselves into a complex but general structure. At the critical point these systems exhibit behavior that is not described by any one event scale but has a variety of event sizes. This makes the dynamic response of SOC systems complex, but the statistical properties of the response follow simple power laws, a property that generalizes across many seemingly different systems. The main claim of SOC is that systems are capable of exhibiting this critical-point behavior, characteristic of percolating systems, but without any parameter tuning required.

To illustrate the concept of SOC, consider the model of a sandpile, which has become a standard paradigm for the discussion of criticality. The model consists of an N-by-N square lattice where every node can be occupied by a column of sand grains. The number of grains present at some node $h(i, j)$ determines the stability of that node. If $h(i, j)$ is smaller than some threshold value h_c, the location is stable. If, however, $h(i, j)$ equals or exceeds the threshold, the pile collapses and the sand grains spill on the neighboring nodes. The borders of the lattice are open, so the sand simply falls off the edges and out of the system at those locations. For example, for $h_c = 4$, the toppling rule can be stated as

If $h(i, j) \geq h_c$ then :

$$h(i, j) \rightarrow h(i, j) - 4$$
$$h(i + 1, j) \rightarrow h(i + 1, j) + 1 \text{ if } i \neq N$$

$$h(i-1,j) \rightarrow h(i-1,j) + 1 \text{ if } i \neq 0$$
$$h(i,j+1) \rightarrow h(i,j+1) + 1 \text{ if } j \neq N$$
$$h(i,j-1) \rightarrow h(i,j-1) + 1 \text{ if } j \neq 0$$

The pile of sand is generated by adding individual grains to randomly selected nodes according to the addition rule:

$$h(i,j) \rightarrow h(i,j) + 1$$

Starting from some stable initial conditions, the pile is built up by the addition rule until overcritical nodes with $h(i,j) \geq h_c$ are produced. Then the system is allowed to relax according to the toppling rule. When all nodes become stable once again, the process is continued.

Computer simulations of the sandpile model show that, starting from an arbitrary initial condition of the sand distribution, the average column height $\langle h(i,j) \rangle$ settles to some temporal mean, but continues to fluctuate slightly around that point [6]. At these sand densities, addition of more sand grains causes the system to exhibit critical behavior. The statistical measurements of the total number of nodes that are involved in an avalanche and the time it takes for an avalanche to die down reveal that they follow power law distributions [6]. In addition, the individual regions that exhibit avalanches are fractal in shape [3].

Numerical simulations have also been performed on a variety of similar models, which were usually based on some natural phenomena that exhibit power laws, such as the formation of river basins, eqarthquakes, and solar flares [12, 9, 1, 6]. These studies reveal some essential dynamical properties of SOC systems. The SOC behavior arises in systems that are dominated by the interactions among their parts, rather than the dynamics of the parts themselves. In the sandpile model, for instance, the behavior of individual sand columns causes avalanches, but transition of sand from one column to another generates complex behavior. The thresholds are also necessary ingredients of SOC systems, since they allow the systems in the critical state to exhibit a great variety of behavior. For example, an addition of a single sand grain to a column

18

will not cause any avalanches if the height of the sand column remains below the threshold, or, at the other extreme, an additional sand grain might start a system-wide avalanche. The existence of thresholds allows a variety of responses between these two extremes to exist. Another important requirement in SOC system is that the external drive of the system must be slow. In the sand pile context, this means that the system must be allowed to relax between additions of new sand grains, or the behavior of the system will be dominated by the strong drive instead of the dynamics of the sand pile.

2.3 Highly Optimized Tolerance

Although SOC provides a mechanism for generating power laws in complex networks, it does not account for the design and evolution that take place in the creation of engineered and biological systems. Moreover, the self-similar characteristics, which are a key feature of SOC systems, are not present in many designed or evolved networks. On the contrary, engineered and biological systems are self-dissimilar and hierarchical in nature. A satellite, for instance, consists of many smaller, interacting parts which are very different in function and form from the satellite as a whole.

In 1998, Carlson and Doyle proposed a new mechanism, called highly optimized tolerance (HOT), that produces power laws in interconnected systems which have the characteristics of engineered and biological systems. The HOT concept applies to systems that are optimized, either through natural selection or design, to perform in uncertain environments. The main claim of HOT is that the tradeoffs between performance and robustness involved in the design of the system give rise to the power law distributions.

To explore the differences between natural and designed networks, Carlson and Doyle performed numerical simulations on the forest fire model. The objective in the design of the forest is to maximize the yield Y, defined as the total number of trees left unburned after a lightning strike. Since the yield increases with the initial density of the trees, but decreases with the size of the fire, the final design is a tradeoff

between these two quantities.

In the undesigned system studied by Carlson and Doyle, the trees are placed on randomly selected nodes of a square N by N lattice until some specified density ρ is reached. This random system then is just a two-dimensional percolation problem with occupation probability $p = \rho$. In the random setting, all initial lattice configurations with density ρ are equally likely to occur. The initial conditions for the designed system, in contrast, are specially selected to optimize the yield.

A fire is initiated by a random lightning spark which has spatial distribution $L(i, j)$. If the lightning hits a non-vacant site, the fire spreads to its neighbors and eventually burns the entire cluster of connected trees. The yield is then computed to be $Y(\rho) = \rho - \langle f \rangle$, where $\langle f \rangle$ is the normalized number of burned trees averaged over different lattice configurations and many lightning strikes.

In the designed system, the initial configuration of trees is selected from many possibilities to maximize the yield of the system when the spatial distribution of the lightning $L(i, j)$ is known. The yield is improved by placing fire barriers (lines of unoccupied nodes) around locations with high distribution of the sparks, breaking the forest up into separated regions. The maximum size of the fire is thus constrained to the size of the largest isolated region. The amount of barriers available to protect against fire is limited by the desired tree density. Such a design process leads to highly structured configurations, with trees planted densely to maximize the timber yield, and fire breaks placed to minimize the spread of fire.

The results of the simulations indicate that the yield of the random system increases initially with the density ρ, but declines back to zero at high density. The transition, which occurs at the critical density ρ_c, becomes sharper as $N \to \infty$. The yield curve for the optimized system has similar structure, but its critical density is much higher than that of the random system and approaches 1 as $N \to \infty$. Thus, the designed systems can achieve much higher yields than the random systems at densities above ρ_c. The high yield systems, which are specifically optimized for maximum timber production, are systems that operate in the HOT state. At the critical point, HOT system exhibit power law distributions of fire events for a variety of light-

ning distribution functions. Moreover, the scaling exponents are independent of the particulars of $L(i, j)$.

The drawback of the HOT configurations is that they are highly sensitive to design flaws and errors. For instance, a single tree growing on a fire barrier can connect two isolated areas of the forest, thereby increasing the size of the fire dramatically. By contrast, moving a tree from one node to another node in a random configuration does not alter the distribution of events. The HOT states are also highly dependent on the exact knowledge of the function $L(i, j)$. If the distribution of sparks is changed from Gaussian to uniform, for example, the distribution of fires in the designed system actually begins to increase with their size, leading to very low yields.

In summary, the HOT states have several features that distinguish them from random critical systems. They are optimized to have high yields in the presence of uncertainty that was included in the design process. The simulations show that the optimization process gives rise to power law distributions, which abound in real engineered and evolved systems. However, HOT systems are highly sensitive to design errors and unanticipated changes in the environment, unlike random systems. Another distinguishing trait of HOT systems is their highly structured and stylized configurations, in contrast to self-similar, fractal configurations of SOC and percolating systems.

2.4 Cascading Failures in the Power System Setting

The problem of cascading fault events in the power system setting can be readily formulated in the terms of percolation, SOC, or HOT scenarios. An external event, such as a break of a transmission line during a storm, puts additional stress on the network, which could result in increased power demand on some generators and increased power flows in some transmission lines. If the new operational regime causes the power flows or the generator frequency deviations to exceed their thresholds, other

components in the network will fail, and cascading fault events will occur.

Although the general description of cascading failures in power systems is similar to percolation, SOC, and HOT networks, the specific rules that govern the propagation of fault events are not as easily defined. One problem lies in determining the "neighbors" of the failed component which are going to experience changes in their states as a result of the failure. For example, a break of a transmission line could increase power flows in lines at spatially removed locations. To decide which transmission lines will be affected, equilibrium power flows must be determined for the whole modified network. Moreover, the dynamics of the power grid, which determine how the system moves from one equilibrium point to another, could cause some system components to fail during transition, making this problem even harder to analyze.

Chapter 3

Modeling

This chapter introduces power network models that can be used to study cascading failures through numerical simulations. The first section describes the swing-equation model, which is used in the experiments presented in this thesis. The section talks about the models for system components represented by the equations and discusses simplifying assumptions made to derive them. The next section presents a spring-mass system that is governed by the same equations as the power grid model. This mechanical analog is useful for visualizing the dynamics of the system and how system components interact with one another. In the third section, protective mechanisms that interfere with the dynamics of the swing equations and cause cascading fault events are added to complete the simulation model. The final section proposes a finite-state (FS) analog for systems of coupled oscillators, which coarsely aggregates the dynamics of such systems while still retaining some features of their complex behavior. This FS reduced model is closer in its structure to the models studied by percolation, SOC, and HOT, and could be used at some point to obtain theoretical insight into the issue of propagating failures.

3.1   Swing-Equation Model

In a power system, the defining variables are voltage (V), frequency ($\dot{\delta}$), and the flows of real and reactive power (P and Q). These variables interact through various means

and on different time scales. For an exploration of dynamic phenomena that cover large geographic territory, the natural starting point would be the swing equations, which relate real power flow to phase angle differences in the transmission network [7]. Swing dynamics occur on the time scale of seconds and can span the entire system. They also have the potential to trigger protection mechanisms of the system in a way that can cause cascading fault events.

In the swing-equation model, the behavior of each generator node or bus is described by a second order differential equation representing the movement of the rotor according to Newton's law in rotational form. We model the loads as synchronous motors, so their dynamics are also described by second-order differential equations. Under the assumption that all bus voltages are well regulated, the relationship between the mechanical power P_{mi} and the electrical power P_{ei} of the component at the ith node can be written as

$$P_{mi} - P_{ei} = M_i \ddot{\delta}_i(t) + D_i \dot{\delta}_i(t), \qquad (i = 1, \ldots, N) \qquad (3.1)$$

where δ_i is the angle of the rotor with respect to a synchronously rotating reference frame, M_i is the inertia, and N is the number of generators and loads in the system. The parameter D_i is the damping coefficient, which represents electromechanical damping, as well as effective damping produced by other components not directly modeled by the swing equations, such as power system stabilizers. If node i is a generator, the mechanical power P_{mi} is the power injected into the system and is a positive quantity. On the other hand, if the node i is a load, the parameter P_{mi} is the power taken out of the system and it is negative. The power supplied by the generators is adjusted to be equal to the power consumed by the loads through the action of governors, whose dynamics are not represented in this simple model.

The power flows in the swing-equation model are described by algebraic equations dependent on the load and network characteristics. The electrical power flowing between nodes i and j is given by

$$P_{ei} = Re[V_i^* I_i] = Re[V_i^* \sum_j (Y_{ij} V_j)] \qquad (i = 1, \ldots, N) \qquad (3.2)$$

where V_i is the voltage at node i, and $Y = G + jB$ is the matrix of network admittances. We assume that the transmission lines are lossless, i.e. $G = 0$, and that the magnitudes of the voltages remain fixed. Then, since $V_i = |V_i| e^{j\delta_i(t)}$, Equation 3.2 can be written as

$$P_{ei} = \sum_j |V_i||V_j| b_{ij} \sin(\delta_i(t) - \delta_j(t)) \qquad (i, j = 1, \ldots, N). \qquad (3.3)$$

The parameter b_{ij} is proportional to the susceptance of the line between nodes i and j, so it is 0 if buses i and j are not connected. Thus, the term in the summation in Equation 3.3 is only nonzero for buses that are connected to bus i. It is important to note that the matrix $B = [b_{ij}]$ directly represents the structure of the network; it is essentially an incidence matrix for the network graph. As the network topology changes due to failures, for example, the pattern of zero and nonzero entries of B evolves to reflect the differences in network structure.

The quantity $|V_i||V_j| b_{ij} \sin(\delta_i(t) - \delta_j(t))$ is the real power flow across the transmission line from bus i to bus j. This expression limits the maximum power that can be transmitted through the line to $|V_i||V_j| b_{ij}$. However, other factors, such as thermal dissipation constraints on the line, may restrict the maximum transferred power to values below this theoretical bound.

Combining Equation 3.2 with Equation 3.3 and rewriting, we obtain

$$M_i \ddot{\delta}_i(t) = -\sum_j b_{ij} \sin(\delta_i(t) - \delta_j(t)) - D_i \dot{\delta}_i(t) + P_{mi} \qquad (i, j = 1, \ldots, N). \qquad (3.4)$$

This equation presents an analytical difficulty, since a solution might not exist for all parameter values. For this reason, ant to get a simpler model, our analysis uses the linearization of Equation 3.4 about the "flat-start" nominal solution, where all variables are 0, which corresponds to the small-signal setting:
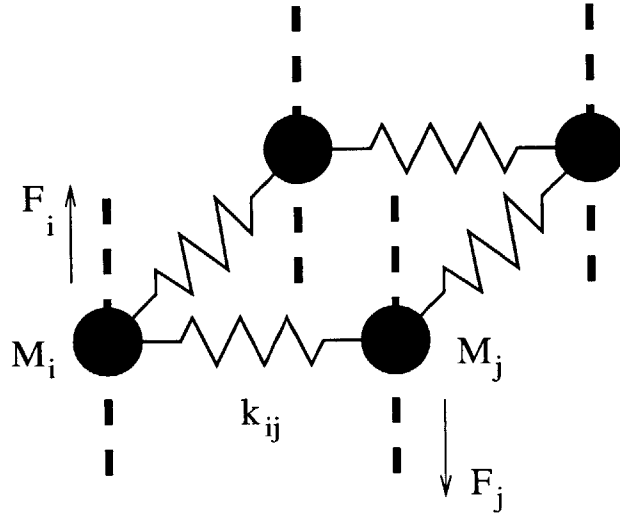
Figure 3-1: Spring-mass system — a mechanical analog of the swing-equation model.

$$M_i \ddot{\delta}_i(t) = -\sum_j b_{ij}(\delta_i(t) - \delta_j(t)) - D_i \dot{\delta}_i(t) + P_{mi} \qquad (i, j = 1, \ldots, N). \qquad (3.5)$$

3.2 A Mechanical Analog

The linearized swing-equations also represent the dynamics of a spring-mass system, shown in Figure 3-1. In this model, objects with mass M_i are arranged on a plane and connected together with springs. The masses slide on vertical poles, which restrict their movement to the direction normal to the plane and introduce frictional forces. Each mass also experiences an external force F_i perpendicular to the plane. This force points up for "generator" masses, and it is directed downward for "load" masses.

The motion of a mass i in an n-mass system is characterized by its vertical displacement δ_i from some frame of reference. Writing Newton's second law we get

$$M_i \ddot{\delta}_i(t) = \sum_j F_{sij} - D_i \dot{\delta}_i(t) + F_i \qquad (i, j = 1, \ldots, N) \qquad (3.6)$$

where F_{sij} is the force mass j exerts on mass i through the connecting spring, and D_i

is the coefficient of friction acting on mass i. The force F_{sij} is the vertical component of the force exerted by an extended spring with spring constant k_{ij}, so we have:

$$F_{sij} = k_{ij}\sqrt{1 + \delta_j - \delta_i} \times \frac{d}{\sqrt{1 + \delta_j - \delta_i}} = k_{ij}(\delta_j - \delta_i). \qquad (3.7)$$

This equation assumes that the rest length of the springs is equal to zero. Substituting Equation 3.7 into 3.6, we obtain the system of equations governing the motion of the masses:

$$M_i \ddot{\delta}_i(t) = -\sum_j k_{ij}(\delta_i(t) - \delta_j(t)) - D_i \dot{\delta}_i(t) + F_i \qquad (i, j = 1, \ldots, N)$$

This is exactly the linearized swing equation, Equation 3.5, with $k_{ij} = b_{ij}|V_i||V_j|$ and $F_i = P_{mi}$. The linear mechanical model can be augmented to represent nonlinear swing equations by introducing nonlinear spring characteristics given by the sinusoid term in Equation 3.4. This analog is useful in developing intuition for the dynamics of the power system, and it allows direct application of literature on the analysis and control of mechanical systems to the power network.

3.3 Failure Mechanisms

The model of the power network used in the simulations has several failure mechanisms, which represent the limitations and safeguards encountered in real networks. These protection mechanisms are capable of interacting with the dynamics of the swing equations to change the topological structure of the system and to produce a succession of fault events.

The first failure mechanism represents transmission line overloading. If the power transfer through a line exceeds the limit for that line, the line is removed from service by opening appropriate "breakers", so the connection between nodes i and j is disabled. In other words, if the difference in the rotor angles between two connected nodes, $|\delta_i - \delta_j|$, exceeds some quantity, called the angle threshold, the parameter b_{ij} is set to zero. Thus, line failures directly change the topology of the network.

27

The generators are also subject to failure events in the model. If the frequency deviation $\dot{\delta_i}$ of a generator exceeds some threshold, the generator is disconnected from the network. This represents the protective relay action that exists in actual generators, to prevent operation that is significantly outside normal speeds. In the model, the tripping of a generator is captured by setting P_{mi} and M_i to zero.

An important consideration, which must be taken into account whenever a failure occurs, is the balance of powers flowing in and out of the system. A new equilibrium can be reached after a failure only if the total amount of power consumed by the loads is equal to the total amount of power supplied by the generators connected to those loads. Since failures change the topology of the network and/or supplied power, a new balance must be established after each failure. In the actual power network, this is accomplished through the use of governors that increase the mechanical power input to the generators if there is a drop in the local frequency. Since the governor dynamics are not included in our model, their action is represented by changing the mechanical power input of each generator equally to balance the power flows. Note that the power balance must be satisfied for every connected subnetwork in the system.

The complete model is a hybrid nonlinear system of coupled differential equations. The presence of failure mechanisms is what makes the model hybrid in nature. Relay actions have the effect of changing the topology of the system and some of its parameters, thus changing the differential equations. The complicated interactions between the dynamics of the swing equations and the overloading and generator protection mechanisms have the potential for generating complex fault events, as will be seen in the next chapter.

3.4 Finite State (FS) Model of Oscillating Systems

In addition to numerical simulations, some theoretical understanding of the problem of cascading failures needs to be developed. Existing methods for percolation, SOC, and HOT are constrained to static problems or cellular automata with a small number of states for each node. The dynamics of the swing equations, however, make the

28

(p_i, v_i)	State Number
$(-1, +1)$	-3
$(-1, 0)$	-2
$(0, -1)$	-1
$(0, 0)$	0
$(0, +1)$	1
$(+1, 0)$	2
$(+1, -1)$	3

Table 3.1: Numbering of states of the finite-state "mass-spring" model.

power system model very different from such networks and much more difficult to study. To aid in the development of methods for analyzing failures in dynamic coupled networks, we develop an a coarse finite-state model for oscillating systems, which bears closer resemblance to the models previously studied in the percolation literature.

In the model, the interacting objects or "masses" are arranged on a horizontal plane, and their movement is confined to the direction perpendicular to the plane. The "position" p_i of a mass i is its vertical displacement from the plane, and it can only assume one of three values: -1, 0, and $+1$. If the objects all have the same mass and if there are no frictional forces, the "velocity" v_i of an object determines how its position will change in the absence of outside forces. In this model, v_i is also restricted to take on only one of three possible values: -1, 0, and $+1$. A mass moves up if it has a positive velocity, it moves down if it has a negative velocity, and it remains in place if its velocity is zero. Thus, a mass in this model can only be in one of seven possible states, as depicted in Table 3.1. The states $(+1, +1)$ and $(-1, -1)$ are not allowed in the system, since the velocity in that state would cause the mass to move out of the allowed set of positions.

The masses in the system are connected together by "springs" which are at their rest length when the masses have the same position. If a mass is vertically displaced from its neighbors, its extended interconnection springs create a force on the surrounding masses that changes their positions and velocities. Therefore, a mass can make transitions between its states not only if it has non-zero velocity, but also if
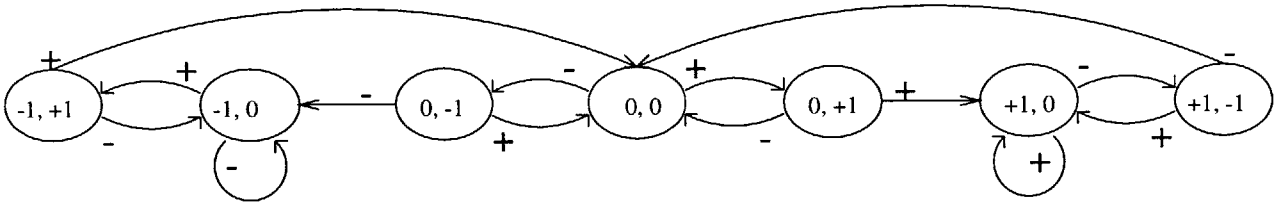
Figure 3-2: Transition graph for the finite-state "mass-spring" model.

there is a difference between the position of the mass and those of its neighbors.

The direction and the amount of transition x_i for mass i is determined at every time step in the simulation of the model by the equation

$$x_i = v_i + \sum_j (p_j - p_i)$$

where the summation is over the set of the masses attached to mass i. At every time interval the state transitions an amount $|x_i|$ in the direction of $\text{sign}(x_i)$ according to the transition graph in Figure 3-2.

Simulations of this system of "masses" and "springs" were done with the masses arranged on a two-dimensional square lattice, with each mass connected to its closest neighbors (diagonal neighbors are not connected). The simulation program allows the user to specify the number of masses in the system and the initial states for all masses. When the program is run, it continuously updates the states and redraws a picture of the system. The picture is a view of the masses from above the horizontal plane, as shown in Figure 3-3. The largest circles represent masses with position of +1, the intermediate-sized objects have position 0, and the smallest circles are masses that have −1 as their position. The number next to each mass corresponds to the state of that mass, as specified in Table 3.1.

The simulations performed with different initial conditions show a variety of oscillatory behavior, similar to that found in the continuous analog of the system. The reduction from an infinite array of possibilities for the state of each mass to only seven possible states makes this model much easier to track. At the same time, the
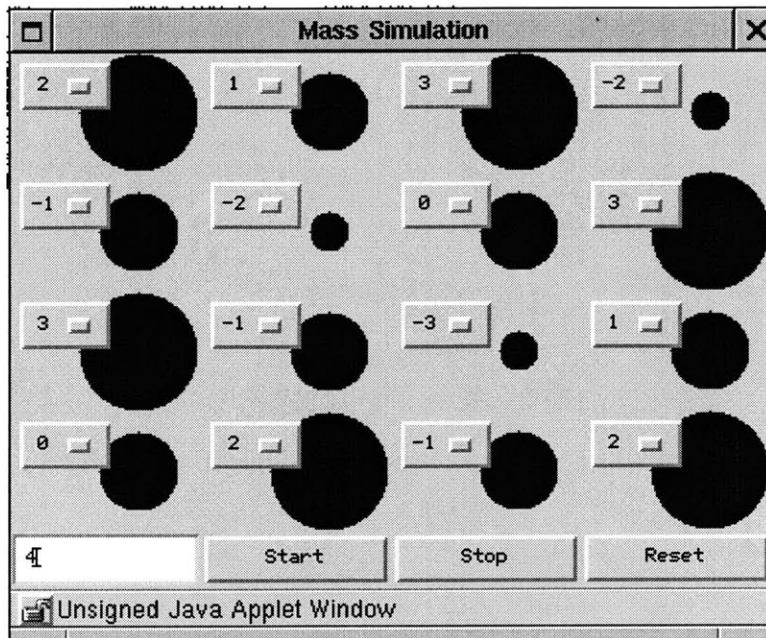
Figure 3-3: Simulation of the finite-state system.

complexity of the behavior of this simplified model indicates that it may be suitable for study of cascading failures from the percolation, SOC, or HOT viewpoint and can be used in future work in this area.

Chapter 4

Simulation Setup

In this chapter, we describe the setup of the simulation program used to perform numerical analysis of propagating failures in an interconnected, dynamic system. The program is based on the swing-equation model for power networks presented in the previous chapter. The first section explains the methods used to represent a continuous-time system through a discrete-time simulation program. The second section discusses the choice of values for system parameters and protection relay settings. The behavior of the simulation model and the cascading fault events produced by the model are described in the third section. The last section compares the response of the dynamic model used in the simulations with the response of an equivalent static system, where only equilibrium state is considered. The results show that the dynamics are important in generating cascading failures, and that accounting exclusively for steady-state behavior leads to different fault events.

4.1   Discretization

In order to simulate the dynamics of a continous system on a digital computer, the equations that describe the system need to be discretized. We use the backward Euler formula

$$\dot{\delta}(t) \approx \frac{\delta(t) - \delta(t - \Delta)}{\Delta}$$

to approximate the first derivative and we use the trapezoidal approximation

$$\ddot{\delta}(t) \approx \frac{\frac{\delta(t+\Delta)-\delta(t)}{\Delta} - \frac{\delta(t)-\delta(t-\Delta)}{\Delta}}{\Delta}$$

for the second derivative of $\delta(t)$. Substituting the above approximations in Equation 3.5, rearranging, and changing to discrete-time notation

$$\delta(t) = \delta[n]$$
$$\delta(t-\Delta) = \delta[n-1]$$
$$\delta(t-2\Delta) = \delta[n-2]$$

yields a set of discrete-time equations for angle deviations:

$$\delta_i[n] = -\frac{\Delta^2}{M_i}\sum_j b_{ij}(\delta_i[n-1] - \delta_j[n-1]) + \frac{\Delta^2}{M_i}P_{mi}[n-2] + $$
$$(2 - \frac{D_i\Delta}{M_i})\delta_i[n-1] + (\frac{D_i\Delta}{M_i} - 1)\delta_i[n-2] \qquad (i = 1,\ldots,N)$$

A combination of derivative approximation formulas was used to discretize the system of equations in order to achieve the best accuracy without increasing the complexity and computation time of the simulation program. The trapezoidal approximation produces a more accurate approximation of the oscillatory continuous time system [11]. However, when the trapezoidal formula is used for the first derivative, it destroys the state-space structure of the system. The discrete-time system created through the use of the trapezoidal approximation alone has the angle deviation of bus i at time n, $\delta_i[n]$, dependent on the angle deviations of other buses at time n. Thus, the equations have to be solved iteratively at every time step, increasing the computational complexity. This problem does not arise when the trapezoidal formula is applied to the second derivative only.

The backward Euler approximation formula, on the other hand, always preserves the state-space structure of the system. In the discretized equations, the angle deviation at time n of bus i, $\delta_i[n]$, depends only on past values of other angle deviations. The discrete-time equation is thus easy to simulate, but it is not accurate for all parameter values. The backward Euler formula gives precise approximations for small enough values of the step-size Δ only when the continuous-time system is stable [11]. The system described by the swing-equations is stable when the damping coefficient D_i is greater than zero. When D_i is equal to zero the continuous-time system is marginally stable and exhibits oscillatory behavior. The discrete-time equations derived from such a system through the use of the backward Euler formula, however, is unstable. By using different formulas for first and second derivatives we reduce approximation errors while still keeping the computation time and complexity of the simulation program small.

4.2 Model Parameter Values

The specific network models used in our simulations are square grids of various sizes. Each node of the square lattice contains a load or a generator at the beginning of the simulation. There is one generator for every eight loads, and the generators are uniformly distributed among the loads, as shown in Figure 4-1. The size N of the N by N grids used in simulations is restricted to be a multiple of three, so that the generator to load ratio and the structure of the network is not changed as the system is scaled.

The parameter values of the simulation model are chosen arbitrarily, and are not meant to correspond precisely to the parameter values of an actual power system. Rather, they are selected to capture some of the qualitative features of real power systems. Table 4.1 summarizes the values of the different parameters used in the simulations. All of the transmission lines are assumed to have the same susceptance, so the parameter b_{ij} is equal to 10 for all lines. For simplicity, all of the generators are given the same mass of 1 and the same damping coefficient of 0.447. Similarly,

34

Parameter	Value	
	Generator Bus	*Load Bus*
M_i	1	0.05
D_i	0.447	0.1
P_{mi}	matched to equal power consummed by loads	uniformly distributed between -1 and 0

Table 4.1: Parameters used in simulations. In addition, $b_{ij} = 10$, $\Delta = 0.01$.

all the loads have a mass of 0.05 and a damping coefficient of 0.1. The masses are chosen so that a generator has 20 times the mass of a load. The damping coefficients are determined to make the system underdamped, and each angle deviation oscillates for about 10 cycles in response to a step change in the power flows before settling to a steady state. Such response is evocative of what is found in actual power systems [8]. A step size of $\Delta = 0.01$ is used throughout.

The power consumption at each load is chosen randomly at the beginning of each simulation, and it is distributed uniformly between 0 and 1. The total supplied mechanical power must balance the total consumed power in order for the system to have a stable equilibrium. Thus, the power supplied to the generators is set equal to the total load power, and it is divided equally among all generators in the system. As the simulation progresses and the network topology evolves, the program checks that the power balance requirement is satisfied at all times for all emerging subnetworks, and adjusts the generator power inputs if necessary.

The settings of the protective mechanisms are chosen so that the simulated system exhibits complex behavior. For instance, if the generator frequency relay is set to a very small value, any disturbance to the system will cause all of the generators to fail immediately and propagating failures will not occur. On the other hand, if the generator frequency relay is set too loosely, it will not have any effect on the system's behavior. Any cascading failure events that might happen will result from the actions of the transmission line relays only. Although such behavior is also interesting to study, it is not a typical scenario for cascading failures in real power networks [4].

Therefore, the protection relays are chosen so that the response of the simulated system to a disturbance typically activates both velocity and power flow relays to produce propagating failures.

4.3 An Example of Cascading Failure

Figure 4-1 shows a typical response of the simulated system to an initial disturbance, which is also described in Table 4.2. At the beginning of the simulation, the 36-bus system has four generators, represented by circles in the figure. A load is located at every remaining node of the square grid, and a transmission line is placed along every edge of the lattice. Initially, the system is in equilibrium and power flows are constant along all of the lines. The transmission line relay of each line is set to be 66% above the steady state power flow along that line. The generators are tripped when their frequencies exceed 0.05.

The failure events are initiated by removing a single transmission line chosen at random, as shown in the second frame of the figure. A short time later, this leads to a failure of one of the generators, as well as additional transmission line failures, as depicted in frame 3. Then, the transmission lines continue to break and one of the loads becomes disconnected from the generators and stops receiving power, as frame 4 shows. The faults continue to spread, as shown in frames 5 through 8, and the whole upper right quadrant of the system is left without power. Eventually, only a few transmission lines remain intact, but all of the generators are taken out of service, leading to a system-wide failure. This is shown in frames 9 through 12.

4.4 Comparison with a Static System

The dynamics of the swing equations play a crucial role in determining the extent of the damage to the system caused by cascading failures. Studying just the static or steady-state power flows through transmission lines also leads to propagating failure events, but it does not paint a complete picture of the system behavior. For exam-
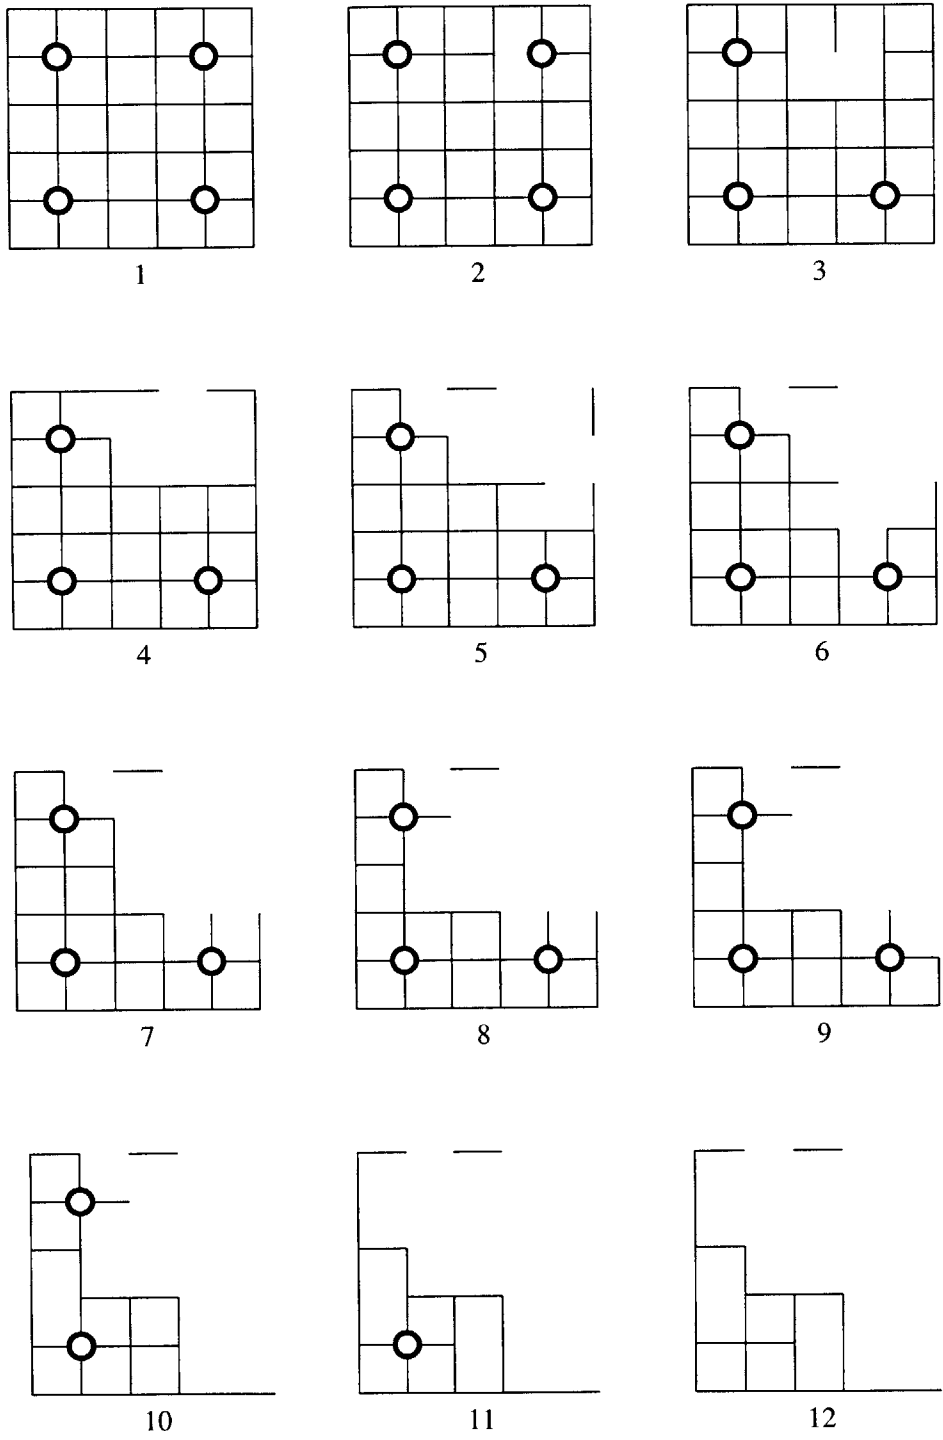
Figure 4-1: Typical scenario for cascading failures.

Time	Line Failures	Generator Failures
0.00	(10,11)	
0.05	(9,10), (10,16)	
0.10		11
0.11	(3,4), (4,10), (5,11), (11,12), (11,17)	
0.12	(4,5)	
0.13	(12,18)	
0.16	(17,18), (17,23)	
0.17	(2,3)	
0.19	(5,6)	
0.20	(16,17)	
0.21	(6,12)	
0.24	(16,22)	
0.25	(22,23)	
0.27	(18,24)	
0.28	(15,16), (23,24)	
0.30	(9,15)	
0.34	(15,21)	
0.36	(14,15)	
0.41	(24,30)	
0.48	(19,20)	
0.48		29
0.49	(23,29), (28,29), (29,30), (29,35), (30,36)	
0.50	(27,28)	
0.51		8
0.52	(2,8), (7,8), (8,9), (8,14)	
0.54		26

Table 4.2: An example of a cascading failure in a 6x6 system. The nodes are numbered from left to right, then top to bottom.
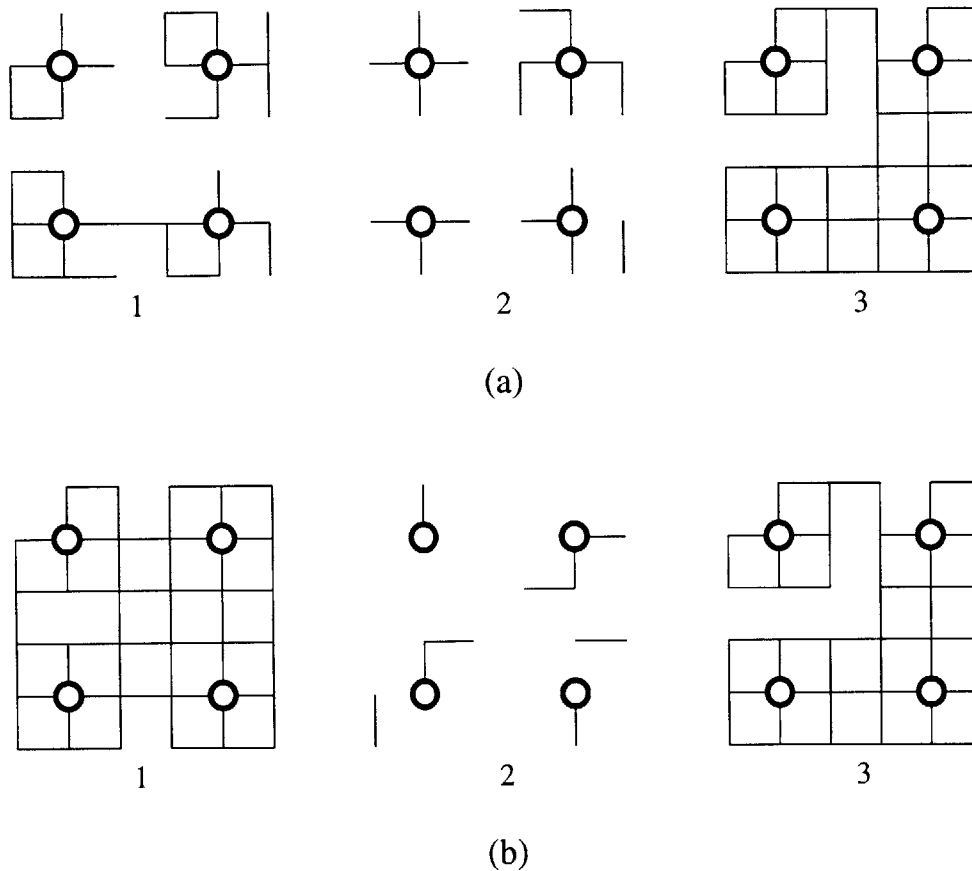
Figure 4-2: Final network configurations of three systems after an initiating failure event in (a) dynamic simulation and (b) static simulation.

ple, the generator velocity relays cannot be incorporated into static systems, since in steady state all velocities are equal to zero. But even if the generator velocity relays are ignored in the dynamic case, leaving only line relays, the simulations of cascading failures in a dynamic system yield very different results from the simulations of cascading failures in the same system without dynamics.

Figure 4-2 shows the differences that arise between failure events in static and dynamic systems. The behavior of the three systems in Figure 4-2 (a) was simulated using the swing-equation model with only line power flow relays in effect. The systems were perturbed from steady state by the removal of a single transmission line, and a series of failure events occured. The figure shows the final topology of the three systems after they have reached a new equilibrium. Figure 4-2 (b) depicts the

same three systems after they have experienced the same initial failures, but with the dynamics of the systems ignored in the simulation. In the second simulation, new steady-state power flows were calculated after each failure event. Then only the steady-state power flows were compared with the transmission line relays, and violating lines were removed. This process also yielded cascading failures, but the final topology of the systems in equilibrium was different in most cases from that of the dynamically simulated systems.

For the first system in Figure 4-2, the dynamic simulation resulted in a much more extensive damage to the system. Six out of 32 load buses became disconnected from the generators, and 28 transmission lines were removed. In terms of the amount of power consumed, 20.3% of the loads were affected by the failure. In contrast, only one load became disconnected in the static simulation, and only 5 transmission lines failed. In the second system, the statically simulated system experienced the larger failure, with 82.5% of the loads affected. The dynamic simulation showed that only 58.7% of the loads actually stopped receiving power. In the dynamic case, some some intervening dynamic failure "redirected" the failure evolution to where the final damage was smaller. The dynamic simulation showed that only 58.7% of the loads actually stopped receiving power. The third system is an example of the case where both simulations lead to the same final network topology.

Clearly, the dynamics of the swing equations are very important in determining the extent of the failures, even when the generator velocity relays are not considered. Depending on the network configuration, static analysis can either underestimate or overestimate the failure size. It is interesting to note, however, that in some situations the static simulations yield the same results as the dynamic simulations, so static analysis might be a useful tool for studying cascading failures under certain conditions.

Chapter 5

Observed Cascading Failures

5.1 Failure Size Measurement

To study the behavior of cascading events in the network, we introduce a criticality parameter α. In the model setting, this parameter scales the amount of power P_{mi} flowing through the network, so the Equation 4.1 becomes

$$\delta_i[n] = -\frac{\Delta^2}{M_i}\sum_j b_{ij}(\delta_i[n-1] - \delta_j[n-1]) + \alpha\frac{\Delta^2}{M_i}P_{mi}[n-2] +$$
$$(2 - \frac{D_i\Delta}{M_i})\delta_i[n-1] + (\frac{D_i\Delta}{M_i} - 1)\delta_i[n-2] \qquad (i = 1, \ldots, N).$$

The parameter α is adjusts the total amount of power consumed by the loads, so it is referred to as the loading factor. The maximum power allowed to flow through the line is set some percentage above the steady-state flow for each line when α is equal to one.

The simulation is started from an equilibrium state with constant power flows through the transmission lines. A failure event is initiated by removing a single transmission line, chosen at random. Depending on the value of the parameter α in the simulation, the system might reach a new equilibrium without violating the power

flow and generator frequency deviation constraints, or it might undergo a succession of failure events such as shown in Figure 4-1. The simulation is run until the system reaches a new equilibrium or all the generators are tripped.

A typical measure for the size of the failure cascade in the power system setting is the number of customers affected by the outage [4]. Since in our model each load has different power requirements (P_{mi} for a load is chosen uniformly between -1 and 0), we measure the size of the failure as the total power requirement of loads which are no longer serviced as a result of the failure, normalized by the total power requirement of all loads. Thus, the failure size takes on values between 0 and 1. The number of affected loads depends on the specific network configuration and the specific triggering event, so the simulations are repeated for 500 different systems and initiating failures. We compute the average size of the failure, which is calculated as the mean value of the failure size over different runs of the simulation.

5.2 Transition Graphs of Average Failure Size

Figure 5-1 shows the average failure size plotted as a function of the loading factor α for several system sizes. For the data presented in this plot, the generator velocity relay is set to 0.05, and the transmission line relays are set to the steady-state flows with α equal to 1. The plot shows the existence of a critical value of the loading factor, which separates two regions of different qualitative behavior. As the size of the system is increased, the transition region becomes narrower.

Some of the characteristics of this failure plot are not surprising. For instance, it is reasonable that for small values of α, the average size of the failures will be small, since the network is not fully utilized. Also, it is expected that for values of α close to one, the size of the failures is large, because the system is overloaded. On the other hand, features such as the sharp transition between the two regions deserve closer attention. The relationship between the transition region width and the size of the system is similar to the dependence found in the percolation setting.
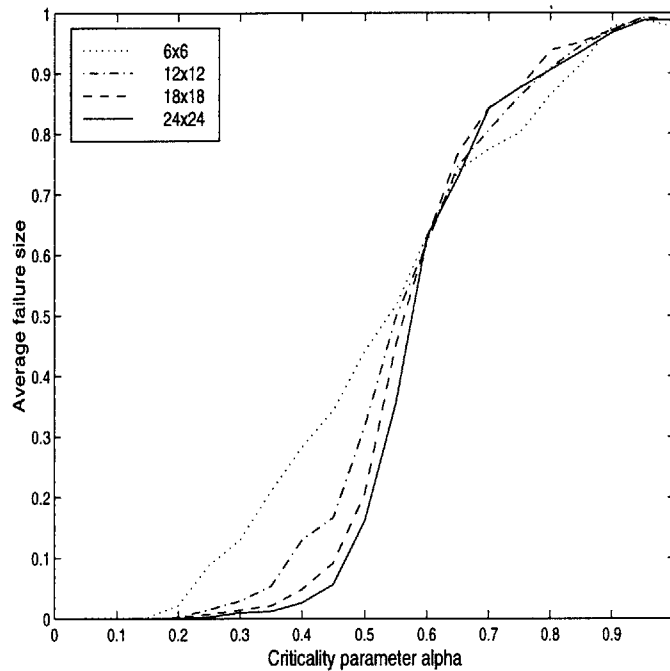
Figure 5-1: Average failure size vs. criticality parameter α for different system sizes.

5.3 Frequency Deviation Relay Settings and Cascading Failures

When the average failure size is plotted as a function of the loading factor for different values of the maximum frequency deviation, the features of the transition curve change. Figure 5-2 compares the transition graphs for a 12-by-12 system with different generator relays. If the maximum allowable frequency deviation is small, the system exhibits larger average failure for all values of the loading factor, and the entire curve shifts to the left. For such a relay setting, the system operates in an all-or-nothing regime for each value of α — either all of the generators are tripped and the failure size is 1, or none of the generators is tripped and the failure size is very close to zero. The type of failure experienced by the system depends on the network configuration and the initiating event. The curve for a relay setting of 0.01 in Figure 5-2 is the average of 500 such all-or-nothing failures. As the criticality parameter is increased, the number of complete failures grows, so the average failure size increases.
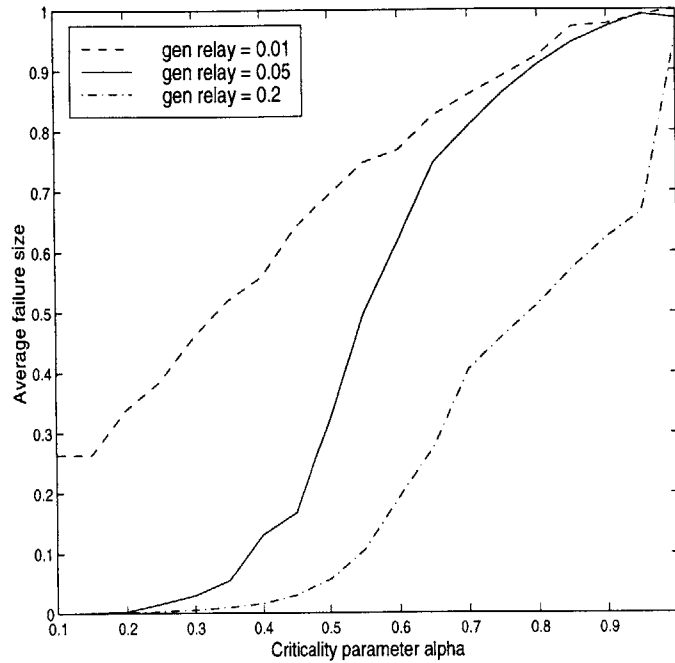
Figure 5-2: Average failure size vs. criticality parameter α for different frequency deviation relay settings for 12x12 systems.

In contrast to systems with tight frequency deviation relay, systems with larger relay setting exhibit cascading failure events of all sizes, a behavior that bears closer resemblance to actual power systems. For intermediate ranges of frequency deviation relay, the curve can be roughly divided into two distinct regions of behavior, as shown in Figure 5-2 for generator relay of 0.05. For values of the loading factor α below 0.4, systems tend to experience small failures, while for α greater than 0.7, the average failure size is close to 1. The transition between the regions of different qualitative behavior is sharp for intermediate values of the generator relay, but it becomes more gradual as the allowable frequency deviation is increased. As the generator relay becomes too loose, though, the graph is no longer broken into two distinct regions of behavior. Instead, the average failure size rises almost uniformly as α is increased. Also, the maximum failure size recorded in such systems increases with α, and a complete failure does not occur until the loading factor is 1. Thus, the critical value of the loading factor only exists for some values of the frequency deviation relay.

44

Note that changing the transmission line settings is equivalent to changing the values of the loading parameter α and the frequency deviation relays, since the system is linear. Moreover, the way the maximum allowable power flow is determined does not significantly change the transition graphs of the average failure size, as will be seen in Section 5.6.

5.4 Distribution of Failure Size

Figure 5-3 shows the cumulative distributions of the failure sizes at a critical value of α in the transition range for different system sizes. The failure size x is plotted on the horizontal axis, and the number of failures of size x or greater is plotted on the vertical axis. The plots are drawn on a logarithmic scale. The graphs show the presence of "heavy tails" in the distribution function — about one tenth of all initiating events result in system-wide outages for a 24-by-24 system, and the situation is even worse for smaller systems. The majority of the failures fall into the size range of 0.1 to 1. For such failure sizes, the distribution is well approximated by power laws, which are shown as straight lines on the logarithmic plots. For both very small and very large outages, the finite size effects play a significant role in determining the distribution graphs. This is due to the relatively small sizes of systems in this study when compared with percolation studies, where systems of size 10^{10} are usually used [12].

For comparison, Figure 5-4 shows the cumulative distributions of the failure sizes at the same value of α for systems with a larger frequency deviation relay setting of 0.1. As was seen in Section 5.3, such systems have a gradual transmission curve that is not characteristic of critical behavior. The graphs in Figure 5-4 demonstrate that these systems have strikingly different distributions, which are well approximated by Gaussian functions. The systems near the critical point obtained with tighter frequency relay settings exhibit much heavier tails, and they experience larger sized events. Also note that larger non-critical systems have a maximum failure size of about 0.5, as opposed to the possible complete failures in critical systems.
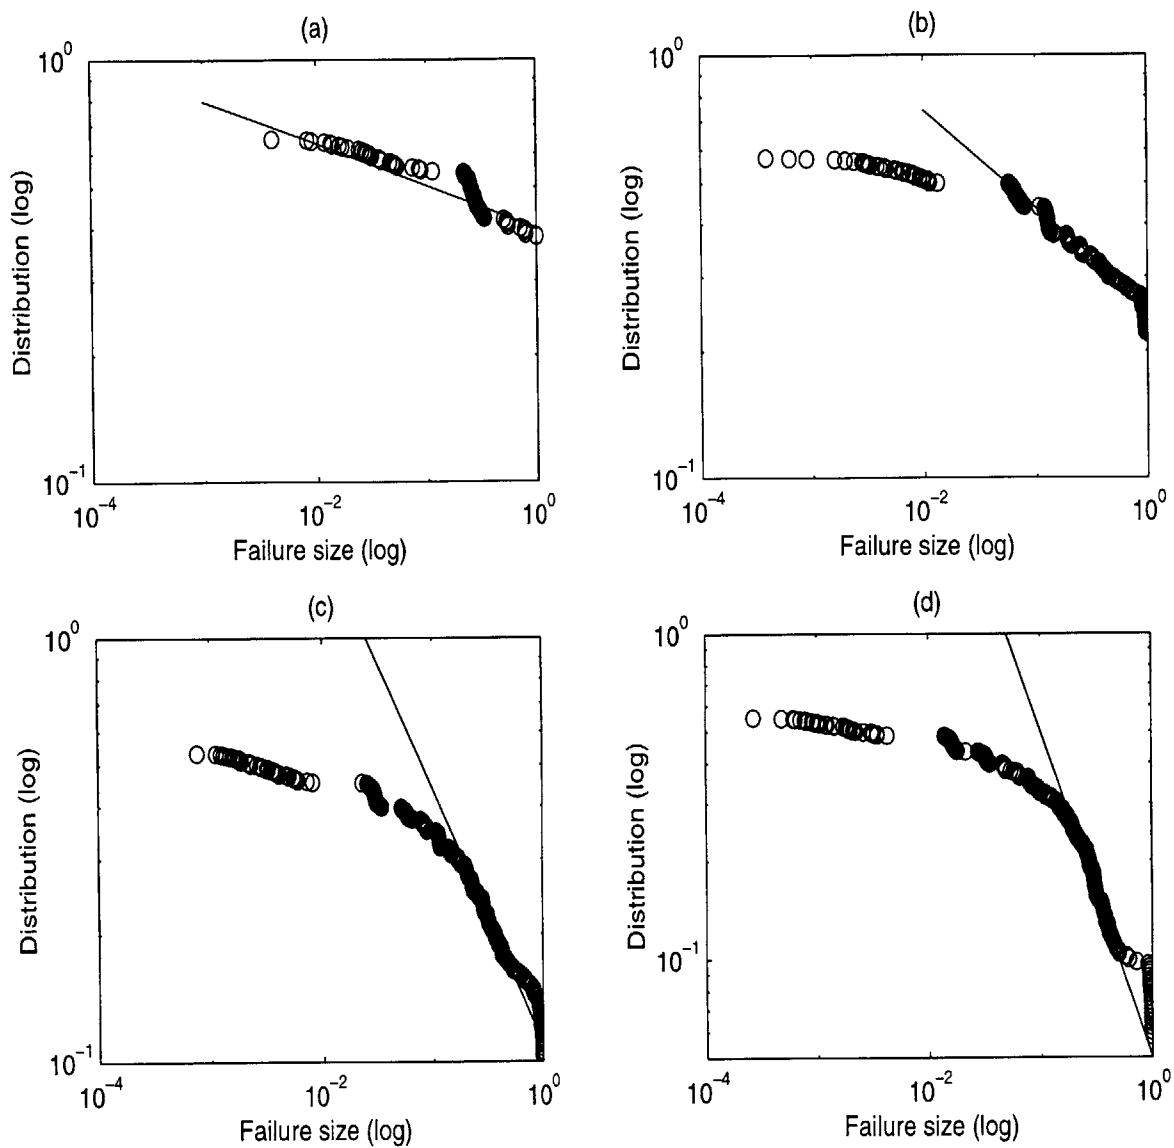
Figure 5-3: Distributions of failure sizes approximated by power laws at $\alpha = 0.5$ for critical (a) 6x6 system, (b) 12x12 system, (c) 18x18 system, and (d) 24x24 system (frequency deviation relay is 0.05 throughout).
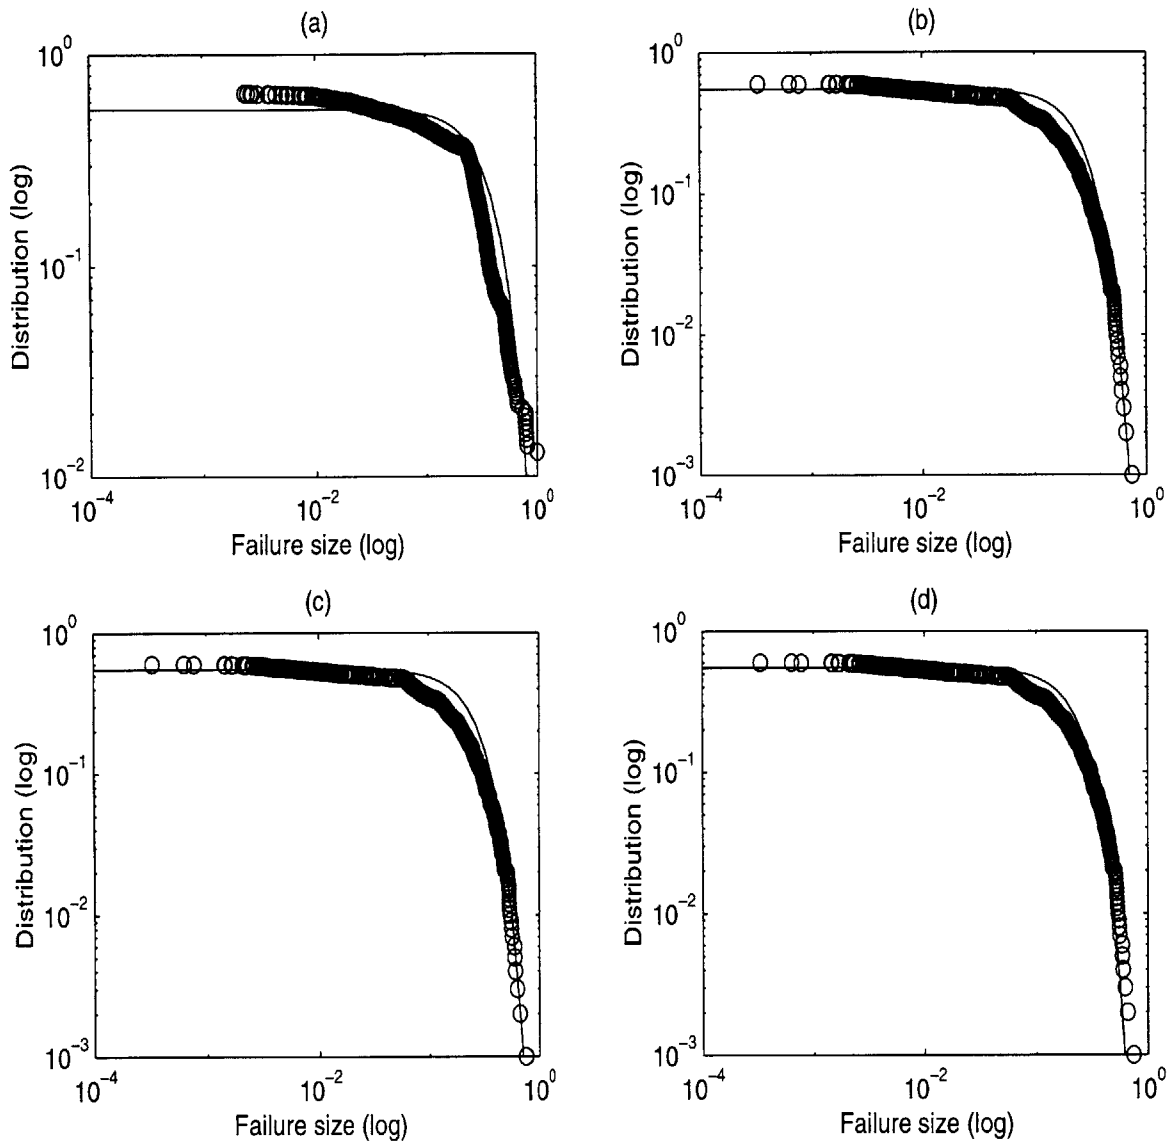
Figure 5-4: Distributions of failure sizes approximated by Gaussian relationships at $\alpha = 0.5$ for non-critical (a) 6x6 system, (b) 12x12 system, (c) 18x18 system, and (d) 24x24 system.

5.5 Sparse Networks

One of the main reasons for creating complex, interconnected power networks is their increased reliability. In highly connected systems, a load does not depend on one generator to supply its power; when a generator fails, other generators will be able to supply the load. But interconnected networks are also prone to cascading fault events, where a single failure is able to spread through the entire system, bringing down many transmission lines and generators. To study the role that high interconnectivity plays in creating or preventing cascading failures in this model, we altered the system configuration to have a lower transmission line density. In the modified model, a transmission line exists between two buses with probability r. If r is equal to 1, the system is as before, but if the value of r is decreased, the system becomes less connected.

Figure 5-5 plots the average failure size versus loading factor for different values of the probability r. The graph indicates that for large values of r, such as 0.9 and 0.95, less connected systems behave similarly to the fully connected systems. However, their average failure size is dramatically higher at low values of the criticality parameter. As the connection probability r is decreased further, the average failure size becomes the same for all values of α, and it is always less than 1. This happens because at lower probabilities the network becomes divided into disconnected subsystems, and the failure size can only be as large as the size of the largest subsystem. Thus, more sparse networks experience smaller average failure sizes under high loading conditions, but they are much more prone to failures at smaller values of the loading parameter.

5.6 Random Relay Settings

To explore the role design plays in the model, additional random elements are introduced into the system. In one experiment, the values of the transmission line relays are randomized. Instead of setting the maximum allowable power flow to some percentage above steady state line flows, the relays are determined according to a uniform
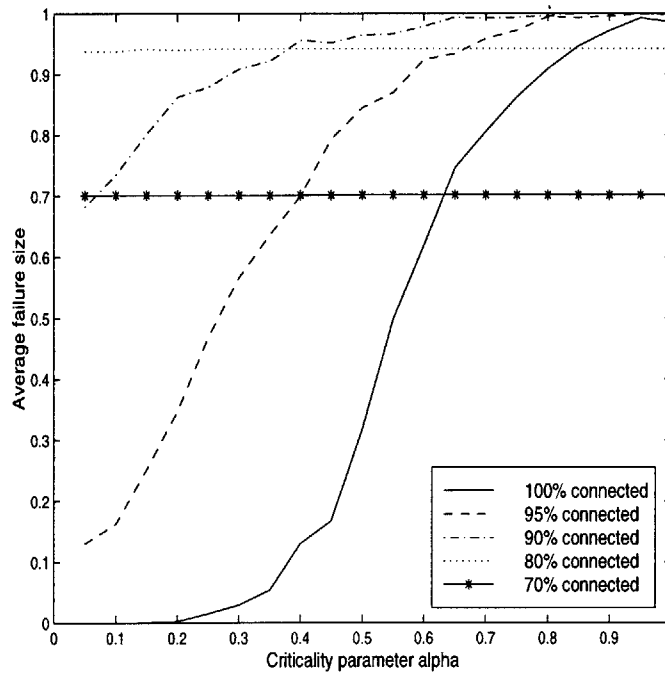
Figure 5-5: Average failure size vs. criticality parameter α for 12x12 systems with different transmission line densities.

probability distribution. To make the comparison with a designed system, the relays are picked so that their average value is equal to that of the "designed" system, and the fluctuations around that point are varied between different simulations.

Figure 5-6 (a) shows the results of the experiment for a 12x12 system. In the designed system, the transmission line relays are equal to 160% of the steady state power flows. In the random systems, the relays are uniformly distributed between 140% and 180%, 120% and 200%, and 100% and 220% for different runs. The figure shows that the performance, measured in terms of small average failure sizes, is not significantly decreased when the transmission line relays are picked randomly, as long as they are above the steady state power flows.

In another experiment, the configuration of the network was randomized. Instead of arranging the generators in a regular pattern and surrounding them by eight loads, the generator buses were chosen randomly from the nodes of the lattice. Figure 5-6 (b) shows the changes in the failure sizes between the designed and random systems.
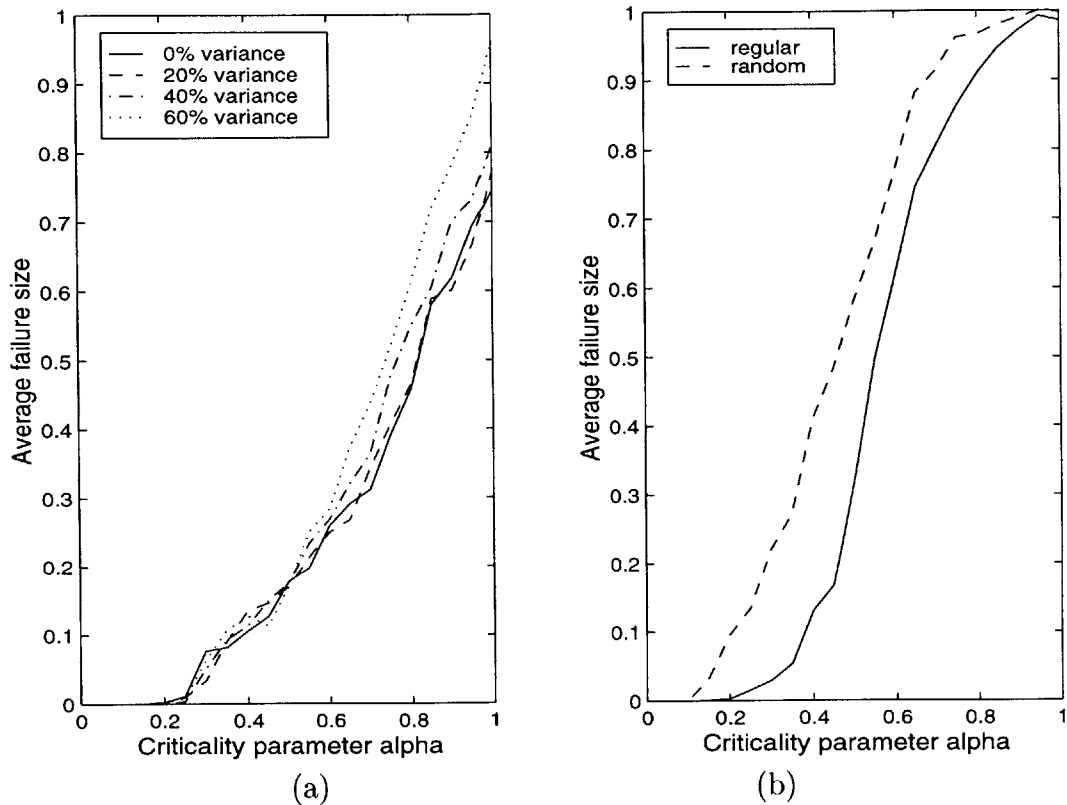
Figure 5-6: Average failure size vs. criticality parameter α for 12x12 systems with (a) random transmission line relay settings, and (b) random generator placement.

The graph demonstrates a clear performance advantage of the designed system.

5.7 Chapter Summary

The results of the simulations indicate that the swing-equation model produces cascading fault events. As the loading of the system is increased, the model undergoes a continuous transition from a region with small average failure size to a region where the average normalized failure size is close to one. The characteristics of the transition region depend on the setting of the frequency deviation threshold. If the threshold is set very tightly, propagating failures are not observed. The initiating event either results in no further failures, or causes all of the generators to be tripped. If, on the other hand, the generator threshold is set too loosely, it has no effect on the system,

and only the power flow relay is responsible for cascading failures. In that case the transition between the regions of different qualitative behavior is smooth, and the failure size has Gaussian distribution.

For intermediate ranges of the frequency deviation relay, the model has a sharp transition region, similar to the critical region found in percolating systems. The transition region becomes sharper as the size of the system is increased. At the values of the loading parameter corresponding to the critical region, the failure size distributions have heavy tails, and can be approximated by power laws. The small size of the simulated systems results in significant finite-size effects, which are apparent in the distribution graphs of the failure size.

Systems that have the same basic configuration but fewer connecting transmission lines exhibit larger average failure sizes. Such systems have better performance only at large loading levels, due to the fact that these systems are broken into independent subsections, so the initial failure can only affect the subsection in which it occurs. Systems that have generators distributed randomly throughout the network have similar failure size transition curves, but experience larger average failure sizes for all values of the loading factor. Finally, the way of choosing transmission line relays has small effects on the system's performance, as long as they are set above steady-state power flows.

Chapter 6

Conclusion

This thesis presented some investigations into the problem of cascading failures through simulations of a simple model of power systems. Chapter 2 discussed the existing literature on systems that exhibit power law distributions and criticality. The main concepts of percolation, self-organized criticality, and highly-optimized tolerance were presented, and their main ideas were illustrated on simple system examples of the square and Bethe lattices, the sand pile model, and the forest fire model. The chapter also described how the problem of cascading failures in the power grid could be viewed in this context. It was shown, however, that even the simplest power system models are much more complex than the systems normally studied by percolation, SOC, or HOT.

In Chapter 3, the models used to study cascading failures through simulations were presented. The nonlinear and linearized swing-equation models for power networks and their mechanical analog, the spring-mass system, were described. The chapter also discussed the protection mechanisms, which could lead to propagating fault events, and their representation in the models. In addition, an finite-state model of a coupled oscillating system was introduced in this chapter. This highly simplified model is easier to relate to the models used in percolation, SOC, and HOT studies.

Chapter 4 discussed the details of the simulation setup, such as the choice of the values for model parameters and protection thresholds. The kinds of propagating fault events that this model generates were also described in this chapter. The results

of simulating the dynamic swing-equation model were compared with the results generated through static analysis of steady-state power flows. It was concluded that, in general, these simulations will yield different results, although some cases were found where both simulations produced the same failure events.

The statistics of the observed cascading failures were presented in Chapter 5. The results indicate that the simulated model transitions from a condition where the average failure is restricted to a small section of the system, to a condition where the average failure encompasses the entire network. It was also shown that the behavior in the transition region depends on the value of the frequency deviation threshold, and that the region exhibits critical behavior for some values of the threshold. In the critical region, the failure size was found to be distributed according to a power law, with characteristic heavy tails.

The comparisons of average failure sizes for systems with different configurations were also presented in Chapter 5. It was shown that systems with fewer connecting transmission lines have larger average failure sizes. It was also demonstrated that systems with random generator placement experienced larger failure events than systems where generators were distributed evenly throughout the network. The choice of power flow relays was shown to have only small effects on average failure size.

6.1   Suggestions for Future Work

One direction for the continuation of the work on the analysis of cascading failures in power systems is to increase the complexity of the simulated model. The model can be extended to include nonlinear dynamics in the simulation, the size of the simulated model can be increased, and different generator/load ratios can be simulated. Alternative methods for setting angle and frequency deviation thresholds can be developed as well, in order to explore the role design plays in generating power law distributions.

Another direction for future work lies in simplifying the model, so that theoretical insight into the problem of cascading failures can be developed. The finite-state model of coupled oscillating systems, presented in this thesis, can be extended to

53

include spring and velocity thresholds similar to the angle and frequency deviation thresholds encountered in power systems. This change would make the model suitable for investigations of propagating fault events in a simplified system.

References

[1] P. Bak. *How Nature Works*. Springer-Verlag, 1996.

[2] J. M. Carlson and J. Doyle. Cascading failures in designed systems: A new mechanism for power laws. Technical report, California Institute of Technology, 1998.

[3] J. M. Carlson and J. Doyle. Highly optimized tolerance: Robustness and power laws in complex sysstems. Technical report, California Institute of Technology, 1998.

[4] Disturbance Analysis Working Group. North American Electric Reliability Council. Data from database at http://www.nerc.com/dawg/.

[5] G. Grimmett. *Percolation*. Springer-Verlag, 1989.

[6] H. J. Jensen. *Self-Organized Criticality*. Cambridge Lecture Notes in Physics. Cambridge University Press, 1998.

[7] P. Kundur. *Power System Stability and Control*. EPRI Power System Engineering Series. McGraw-Hill, 1994.

[8] G. N. Ramaswamy. *Modal Structures and Model Reduction, with Application to Power System Equivalencing*. PhD thesis, Massachusetts Institute of Technology, 1995.

[9] I. Rodriquez-Iturbe and A. Rinaldo. *Fractal River Basins: Chance and Self-Organization*. Cambridge University Press, 1997.

[10] M. Schroeder. *Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise.* W. H. Freeman, 1991.

[11] W. McC. Siebert. *Circuits, Signals, and Systems.* The MIT Press, 1986.

[12] D. Stauffer and A. Aharony. *Introduction to Percolation Theory.* Taylor and Francis, second edition, 1992.