Advanced Simulation for Analysis of Critical Infrastructure: Abstract Cascades, the Electric power grid, and Fedwire[1]

Robert J. Glass, Walt E. Beyeler, and Kevin L. Stamber

Advanced Modeling and Techniques Investigations (AMTI)

The National Infrastructure Simulation and Analysis Center (NISAC)
Sandia National Laboratories
Albuquerque, New Mexico 87185

[1] This work was supported by the National Infrastructure Simulation and Analysis Center (NISAC) with funding from the United States Department of Homeland Security. NISAC's core partners are Sandia National Laboratories and Los Alamos National Laboratory. Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000. Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405 ENG-36.

Executive Summary:

Critical Infrastructures are formed by a large number of components that interact within complex networks. As a rule, infrastructures contain strong feedbacks either explicitly through the action of hardware/software control, or implicitly through the action/reaction of people. Individual infrastructures influence others and grow, adapt, and thus evolve in response to their multifaceted physical, economic, cultural, and political environments. Simply put, critical infrastructures are complex adaptive systems. In the Advanced Modeling and Techniques Investigations (AMTI) subgroup of the National Infrastructure Simulation and Analysis Center (NISAC), we are studying infrastructures as complex adaptive systems.

In one of AMTI's efforts, we are focusing on *cascading failure* as can occur with devastating results within and between infrastructures. Over the past year we have synthesized and extended the large variety of abstract cascade models developed in the field of complexity science and have started to apply them to specific infrastructures that might experience cascading failure. In this report we introduce our comprehensive model, *Polynet*, which simulates cascading failure over a wide range of network topologies, interaction rules, and adaptive responses as well as multiple interacting and growing networks. We first demonstrate *Polynet* for the classical Bac, Tang, and Wiesenfeld or *BTW sand-pile* in several network topologies. We then apply *Polynet* to two very different critical infrastructures: the high voltage electric power transmission system which relays electricity from generators to groups of distribution-level consumers, and Fedwire which is a Federal Reserve service for sending large-value payments between banks and other large financial institutions. For these two applications, we tailor interaction rules to represent appropriate unit behavior and consider the influence of random transactions within two stylized networks: a regular homogeneous array and a heterogeneous scale-free (fractal) network.

For the stylized electric power grid, our initial simulations demonstrate that the addition of geographically unrestricted random transactions can eventually push a grid to cascading failure, thus supporting the hypothesis that actions of unrestrained power markets (without proper security coordination on market actions) can undermine large scale system stability. We also find that network topology greatly influences system robustness. Homogeneous networks that are "fish-net" like can withstand many more transaction perturbations before cascading than can scale-free networks. Interestingly, when the homogeneous network finally cascades, it tends to fail in its entirety, while the scale-free tends to compartmentalize failure and thus leads to smaller, more restricted outages. *In the case of stylized Fedwire*, initial simulations show that as banks adaptively set their individual reserves in response to random transactions, the ratio of the total volume of transactions to individual reserves, or "turnover ratio", increases with increasing volume. The removal of a bank from interaction within the network then creates a cascade, its speed of propagation increasing as the turnover ratio increases. We also find that propagation is accelerated by patterned transactions (as expected to occur within real markets) and in scale-free networks, by the "attack" of the most highly connected bank. These results suggest that the time scale for intervention by the Federal Reserve to divert a cascade in Fedwire may be quite short.

Ongoing work in our cascade analysis effort is building on both these specific stylized applications to enhance their fidelity as well as embracing new applications. We are implementing markets and additional network interactions (e.g., social, telecommunication, information gathering, and control) that can impose structured drives (perturbations) comparable to those seen in real systems. Understanding the interaction of multiple networks, their interdependencies, and in particular, the underlying mechanisms for their growth/evolution is paramount. With this understanding, appropriate public policy can be identified to guide the evolution of present infrastructures to withstand the demands and threats of the future.

1. Introduction

National and economic security and the quality of life in the U.S. depend on reliable operation of complex infrastructures. The National Infrastructure Simulation and Analysis Center, or NISAC, provides modeling and simulation capabilities for analyzing critical infrastructures. These capabilities will improve the robustness of our nation's critical infrastructures by informing policy analysis, investment and mitigation planning, education and training, and near real-time assistance to crisis response organizations. The Advanced Modeling and Techniques Investigations Task (AMTI) is one of NISAC's long-term investments in understanding critical infrastructures and their interdependencies (Glass et al., 2003). Our mandate is to identify theories, methods, and analytical tools from the study of general complex adaptive systems that are useful for understanding the structure, function, and evolution of complex interdependent critical infrastructures.

The complexity of the interdependent and ever changing systems that comprise critical infrastructures makes understanding and modeling them difficult. Fortunately, there has been a great deal of basic research over the past few years focused on understanding complex adaptive systems and developing theories to explain how they behave under stress. This fundamental research has begun to explain the evolution of generic complex network structures, and to identify their vulnerabilities and strengths. We are applying and extending the results of complexity theory to model critical infrastructures. Complexity theory allows us to better understand how general features, such as network connectivity and operational pressures, influence system robustness, determine operating margins, and control system behavior and evolution. This perspective may disclose strategies to make critical infrastructures more robust by strengthening a given set of components, or through the formulation of appropriate long range policy whereby the infrastructure evolves robustness over time. AMTI's initial work has focused on one of the hallmarks of complexity theory, *the cascade model*. In the past decade, cascade models have been used to represent many natural and social phenomena that undergo catastrophic response such as earthquakes, mass extinctions, major forest fires, landslides, epidemics, wars, and revolutions. They have also been proposed to represent a variety of critical infrastructure problems, such as the failure of electric power grids and telecommunication systems, traffic jams, financial market crashes, and the behavior of groups of people in crisis.

An example of a simple cascade model is the now-classical "BTW sand-pile" introduced in 1987 by Bak, Tang, and Wiesenfeld. This model has created its own cascade of activity with well over 2000 citations and applications in fields ranging across physics, biology, economics, and geology, as well as some spillover to critical infrastructure. BTW systems can develop to a state of *self-organized criticality*. For systems that exhibit self-organized criticality, events of all sizes may occur, and the frequency of events as a function of size is represented by a power-law over a range limited by the size of the system. Important repercussions of this behavior are the lack of predictability beyond a statistical representation and "heavy tails" in the cascade statistics due to the power-law. Thus, although they are rare, large system-spanning events are to be expected.

In the AMTI effort, we have first worked toward synthesis of the large variety of cascade models reported in the literature and their extension and application to critical infrastructure systems. As part of this effort we are building a comprehensive model, *Polynet*, which goes beyond the simple cascade models of the past. *Polynet* implements a wide range of network topologies, interaction rules, and adaptive responses as well as the ability to consider multiple interacting and growing networks. In addition to considering a variety of purely abstract problems, we are also applying *Polynet* to specific critical infrastructures. By grounding the abstract in the specific, as well as generalizing the specific in the abstract, our effort strives to enhance our fundamental understanding of the structure, function, and evolution of complex interdependent critical infrastructures. If we can understand "robustness" as a function of the critical characteristics of infrastructures and understand how infrastructures respond to evolutionary pressures, we can identify public policy that will foster the evolution of robust infrastructures.

In this report, we first introduce *Polynet* and demonstrate its behavior for the classical BTW sand-pile implemented on arbitrary network topologies (**Section 2**). We then apply *Polynet* to two example critical infrastructures, the electric power grid (**Section 3**) and Fedwire (**Section 4**). We choose these two applications because they exhibit very different interaction rules and contexts. For each we consider cascade behavior in context of random transaction perturbations in two stylized network topologies: homogeneous regular networks as well as complex scale-free networks that exhibit fractal qualities. Finally we conclude and sketch important directions of future study in the AMTI effort (**Section 5**).

2. *Polynet*

Polynet is an abstract automaton or simple agent based model that encompasses and extends previous algorithms published in the literature (e.g., Bak, Tang, and Wiesenfeld, 1987; Olami, Feder, and Christensen, 1992; Sakhtjen, Carrerras and Lynch, 2000; de Arcangelis and Herrmann, 2002; Goh, Lee, Kahng, and Kim, 2003). Generally, abstract nodes are linked together into a network of arbitrary topology. The state of each node is defined by a local state variable. External stresses can change a node's state. More importantly, one node's state can depend on the state of connected nodes through interaction rules that may range from local (nearest neighbor) to global (the state of each node depending on the state of all other nodes). When the value of a node's state variable exceeds a given threshold, the state of the node abruptly changes. The interaction rules propagate abrupt state changes to connected nodes and thus can create a cascade. This simple framework can produce rich and surprising behavior. It can also be extended to include additional processes found to be important in critical infrastructure, as well as other complex adaptive systems, that can act at a variety of nested time scales. Feedback between node activity and node behavior can drive nodes to adapt by adjusting their thresholds, their links to other nodes, as well as other critical parameters. Additionally, multiple networks with alternative interaction rules and topologies may be connected to consider inter-dependencies and inter-network cascades. A full description of *Polynet* is beyond the scope of this report and will be presented elsewhere.

The generality and flexibility of *Polynet* allows application to an extremely wide range of problems. As an introduction to the model, and before describing our two critical infrastructure applications, we will demonstrate *Polynet* with the simple abstract problem of the BTW sand-pile. In the BTW sand-pile, a grain of sand is added to a site chosen at random within a two dimensional square lattice. When the number of grains at a site exceeds 4, it distributes a grain of sand to each of its non-diagonal neighbors. If any of these sites are pushed over their thresholds, they too distribute their sand grains and thus contribute to the cascade. Sand is removed from the domain when it encounters the edge of the network. Application of the BTW sand-pile relies on a separation of time scale between a relatively slow process that adds sand to the domain and a fast process by which sand is redistributed within the domain. Thus the fast process fully relaxes the system and a cascade, when triggered, runs to its completion before the next addition of a sand grain to the domain; concurrent cascades are precluded. To generalize the BTW sand-pile and apply it to arbitrary network topologies, let us consider grains of sand to represent units of "energy", E, and specify a constant threshold value across all sites, E_c, at which a site changes state and distributes one unit of E to each of its neighboring sites. Let us also choose a small number of randomly distributed sites within the network to act as sinks that absorb all E distributed to them. These sites play the role of the edges of the original BTW sand-pile and allow closed networks to be considered. In this generalized form, we can now apply the BTW sand-pile to any network topology.

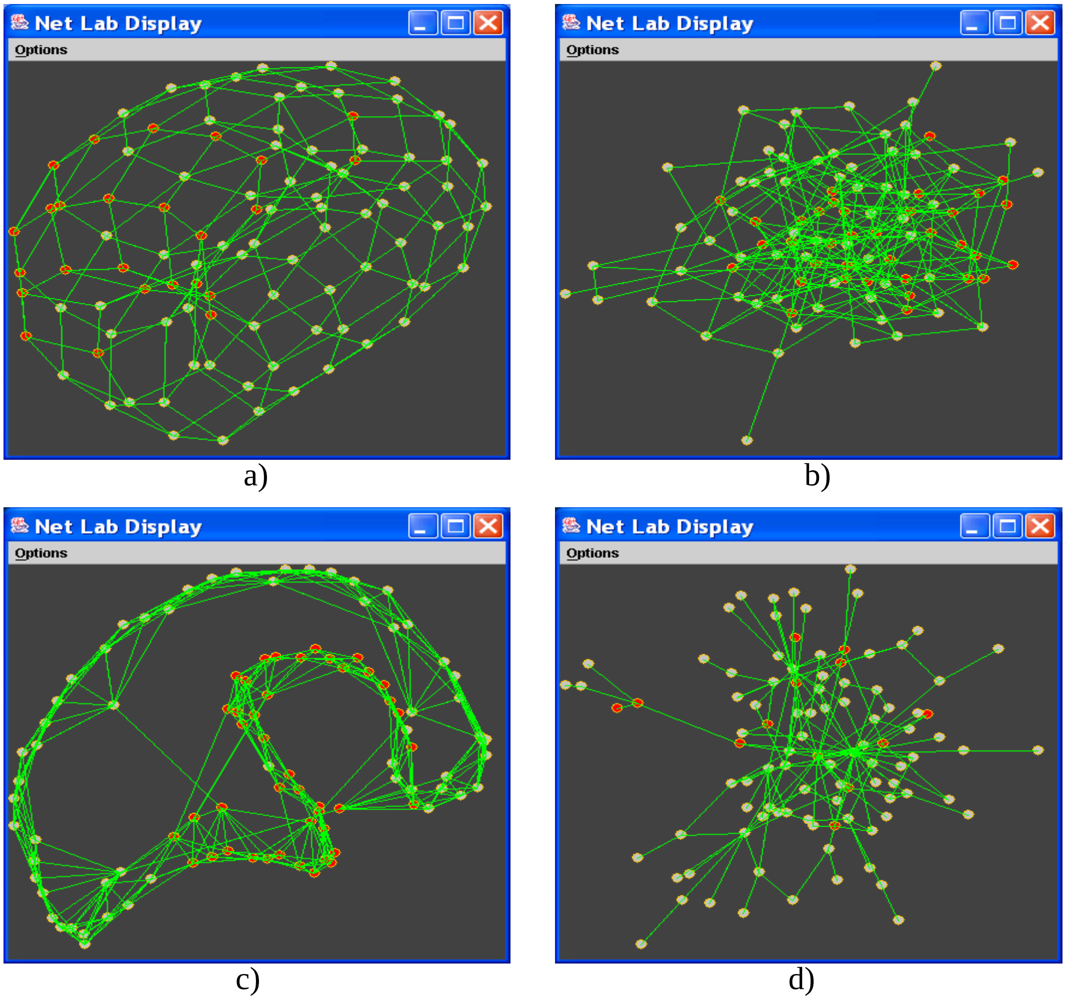

Figure 2.1: Example Stylized network topologies for a) Square Lattice "fish net" with wrapped edges to form a torus (or "donut"), b) Random network, c) Watts-Strogatz blended network, and d) Scale-Free network. Red nodes denote a BTW sand-pile cascade in progress.

Stylized network topologies range from *regular* (as originally considered by BTW), to *random* (Erdos and Renyi, 1959), to *scale-free* (Barabasi and Albert, 1999), as well as others that blend any or all (e.g., Watts and Strogotz, 1998). Regular networks are classical constructs that can range from "fish-nets", to honeycombs, to rings (see **Figure 2.1a** where a fish-net with wrapped edges forms a torus or "donut"). In many regular lattices, nodes are connected to others who also have the same neighbors and thus they often demonstrate a high degree of *clustering*. While regular lattices can be found in nature, such as in crystals or bee hives, they are poor models for many real networks under study. Social networks are an example. There the *small-world* characteristic, i.e., a small number of steps connect any two nodes within the network, is ubiquitous but not found in regular lattices. A stylized network that has the small-world characteristic is the random network developed by Erdos and Renyi (1959). A random network is formed by the sequential random pairing of nodes (see **Figure 2.1b**). Unfortunately, random networks have very little clustering. To remedy this problem, Watts and Strogatz (1998) created a blended network that begins with a regular ring, where every node on the ring is connected not only to the next node but to a certain number ahead and behind on the ring, and then imposes random re-wiring to achieve both a high degree of clustering as well as the small-world characteristic (**Figure 2.1c**). In other naturally occurring networks, one finds a power-law or near power-law for the *nodal degree distribution* such that a significant number of highly connected nodes exist (i.e., a heavy tail). While such a distribution can be imposed on a random or Watts-Strogatz network through re-wiring, a intuitively pleasing approach is to form it naturally by growing the network such that nodes entering the network preferentially attach to ones of higher degree. Barabasi and Albert (1999) introduced such an approach to generate what is now called the *scale-free* network for its fractal properties (**Figure 2.1d**).

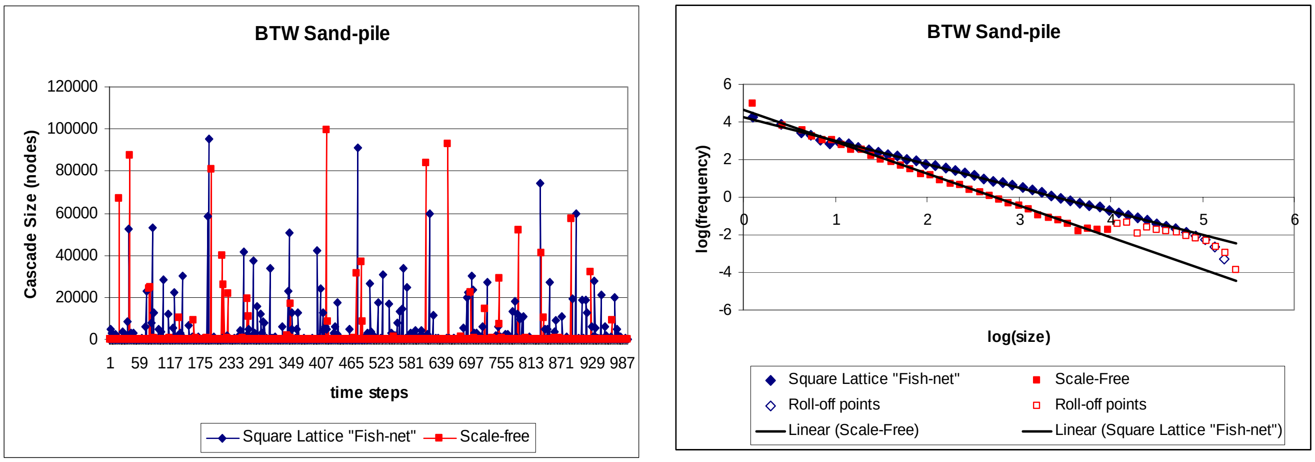

Figure 2.2: For the BTW sand-pile on fish-net and scale-free networks, example time series for cascade size (to the left) and cascade size distributions (on the right).

Example BTW sand-pile simulations for 10,000 node problems in each of the stylized network types presented above exhibit time series that are highly erratic (see examples for fish-net and scale-free in **Figure 2.2 left**). In the figure, cascade size (defined by the number of times nodes in the network are pushed about threshold and distribute E) is plotted in time defined by the number of unit additions of E to the network. The time between cascades appears to be random and the size of the cascade unpredictable. Cascade size distributions for each network type (**Figure 2.2 right)** exhibit the typical BTW sand-pile power-law with eventual exponential "roll-off" at large values. The power-law is indicative of self-organized criticality while the roll-off reflects the finite size of the simulation. The exponent of the power-law (slope of the line in **Figure 2.3 right**) is dependent on the network topology and conforms to the results obtained recently by others (de Arcangelis and Herrmann, 2002; Goh, Lee, Kahng, and Kim, 2003). Interestingly, but less noticeable in the figure, an additional influence of network topology is found in the fraction of cascades that involve only a single node. Nearly half of the cascades in the scale-free network are single node while only a tenth are single node in the other networks. This is due to the fact that in scale-free networks, nodes of degree one (singly connected) are always connected to nodes of higher degree (multiply connected). Because a node in the BTW sand-pile distributes a unit of energy to connecting node, an energy "well" equal to its degree appears for each node. Thus, cascades from lower degree nodes into higher degree nodes are buffered and a number of them must occur before the higher degree node is pushed above threshold. Combining this behavior with the fact that nodes of degree 1 make up roughly 40% of the scale-free network and thus receive that portion of the random perturbations, a large number of single node cascades result.

The BTW sand-pile considers simple local nearest neighbor interactions between nodes and models a transmission process within a network that is fast relative to the addition of accumulating perturbations. As it stands, such a model may have application to a variety of situations of importance in the analysis of critical infrastructures.[2] However, the constraints of the BTW sand-pile can be relaxed or replaced with others quite generally within *Polynet* and thus transform the model in many directions. In the next two sections, we explore such transformations in context of two applications: the electric power grid and Fedwire.

3. Stylized Electric power grid *Polynet* Application

Electric power grids have a propensity to cascade; historical examples include the summer of 1996 outages in the Pacific Northwest (see Kosterev et al, 1999 for details) and the recent August 14, 2003 outage in the Northeast United States and Canada. The first example of a cascade model applied to the electric power grid was presented by Sakhtjen, Carrerras and Lynch (2000). They conceived of the electric power grid as a closed network of transmission nodes and implemented nearest neighbor interaction rules to model a damage surge. Perturbations were imposed by random exchanges of voltage between one node and its neighbor. When a given node was pushed above its threshold, it distributed voltage to its neighbors similarly to the BTW sand-pile (e.g., **Section 2**). Taking a somewhat different approach, Motter and Lai (2002), and most recently Albert, Albert, and Nakarado (2004), considered the cascade of an electric power grid due to the redistribution of 'load' within a damaged network subsequent to the removal of a single transmission node. They implemented global interaction rules to calculate 'load' conceived of as the local congestion incurred within the network from the routing of electricity along the shortest paths between generator and consumer nodes. If the local load increased beyond a given threshold for a particular node, it too would fail and cause additional load redistribution.

While the application of a cascade model with simple interaction rules is a considerable simplification of the complex dynamics that transpire within a transient electric power grid as it fails, such application allows focus on a variety of situations that are currently difficult to assess by other means. The response of a cascade model to random surges such as considered by Sakhtjen, Carrerras and Lynch (2000) and node removal either under random or directed attach such as considered by both Motter and Lai (2002), and Albert, Albert, and Nakarado (2004) suggest that the electric power grid can be fragile under these types of perturbations. Another possible driver of perturbations is a simple, unregulated market, such as that which characterizes certain geographic elements of the present electric power market and has been hypothesized to promote large scale electric power grid instability. In our example application of *Polynet*, we ask whether an electric power grid is inherently prone to cascading failure through its operation in context of a simple market and whether network topology can play a role in either postponing cascades or modifying their extent once they occur. Thus, we consider a stylized electric power grid responding to simple random fluctuations in load within the network imposed by energy transactions between power generators and consumers.

---

[2] For instance, consider a simple abstraction of human mass action in response to crisis within say a financial market or a subway station. Let us interpret the nodes in our BTW sand-pile to be people connected within a social/financial network or simply adjacent to each other in space. E will be the "desire" an individual has and E_c the threshold desire level needed to make a given binary decision such as buy/sell or fight/flight. Many studies of animals and people have documented "herd" behavior whereby individuals are influenced by the action of their friends or neighbors. If we abstract this as the simple BTW sand-pile interaction rule and interpret random additions of unit E to be random events that increase the desire of a person to act, our analogy is complete.

In our stylized application, we place power generators and consumers within ideal network topologies and calculate the resulting flow of current within using a DC circuit analogy. At each node, load is defined by the sum of the currents conducted across the node. Threshold values for load (above which failure occurs) are given by a safety factor on the initial base loads imposed by the generator-consumer population. Both positive and negative perturbations to this initial state are imposed by generator-consumer pair transactions chosen at random. Transaction sizes are a fixed fraction of the consumer's present demand. After a transaction, the flow of current is recalculated within the network and the new loads are compared to their threshold values. If a node exceeds its threshold, it fails and is removed from the network. Loads within the network are then recalculated, compared to their threshold values, and failed nodes once again removed. The process of load redistribution and node removal proceeds until the cascade ends. If a perturbation does *not* result in the separation of consumers from generators, it is retained and a new perturbation is added. In this way, the system undergoes a random walk that varies both the total current and its distribution within the network. When a cascade occurs and consumers are separated from generators resulting in unmet demand, data from the cascade are recorded such as the time to cascade (number of perturbations added), its size (number of nodes that fail), duration (number of load redistribution cycles that take place), and fraction of demand that is unmet. The initial state is then restored and perturbations are once again imposed. In this way, we build cascade statistics over a large number of simulations for a given network configuration.

For our example problem, we consider a stylized high voltage power transmission grid with a set of parameters that are reasonably representative of real electric power grids. We take the ratio of generators to consumers as ~ 1:1 (where the term consumer is representative of a group of residential, commercial, and/or industrial consumers of electricity with a common connection point to the high voltage transmission system) and in combination, they make up ~ 2/5's of the nodes, the remaining 3/5's forming relay substations. While consumers have very similar demand (250 MW with a standard deviation of 20 MW), generators produce over a wide range (50-1200 MW with a mean of 300MW). We take the safety margin at a node as ~ 4 standard deviations, or ~ 30% above the base mean value. While the conductive properties of links in an electric power grid are heterogeneous across the network, we treat them as homogeneous in our example so that the variability of load within the network is entirely due to network topology and placement of consumers/generators. Finally, for convenience, we take the size of the transaction perturbation as ~ 4% of the particular consumer node's present demand.

The topology of realized electric power grid networks is under investigation (e.g., Albert, Albert, and Nakarado, 2004). For demonstration purposes, we considered two stylized network types that seem to somewhat bracket expected topology. The first is a scale-free network generated with the preferential attachment algorithm of Barabasi and Albert (1999) (e.g., see **Figure 2.1d**). For this stylized network, we designate all nodes of degree 1 to be either generators or consumers, with nodes of higher degree forming relay substations. This arrangement mimics the tendency of high voltage electric power grid networks to be sparsely connected at both the supply and demand points but be densely connected in between. Additionally, we consider the regular fish-net square lattice with wrapped edges so that the grid folds to be a torus (or "donut") (e.g., see **Figure 2.1a**). For this stylization, generators and consumers are distributed at random within the network.

We performed multiple 400 node realizations for each network type and in all cases, cascades occurred in our stylized electric power grids. Statistics were gathered for each realization in a period over which 1000 cascades occurred. In general, fish-net networks required ~ 30 times more perturbations than did the scale-free to complete the 1000 cascade period. We have chosen a typical simulation for each network type to depict in our figures. **Figure 3.1** shows cascade size over the course of a run for both networks. Across all realizations, the scale-free always contained multiple bands while the fish-net tended to have concentration at the high end with the occasional development of a small low end peak. We also find that the size of the cascade is directly related to the fraction of demand that is shed by the cascade (**Figure 3.2**) as well as to its duration (**Figure 3.3**). Interestingly, for the scale-free, both large and small events are of short duration while intermediate size events are much longer. The data for the scale-free is distributed tightly or "clumped" around a set of distinct loci. For the fish-net, larger events tend to be of longer duration however there is much greater scatter around the large event locus.

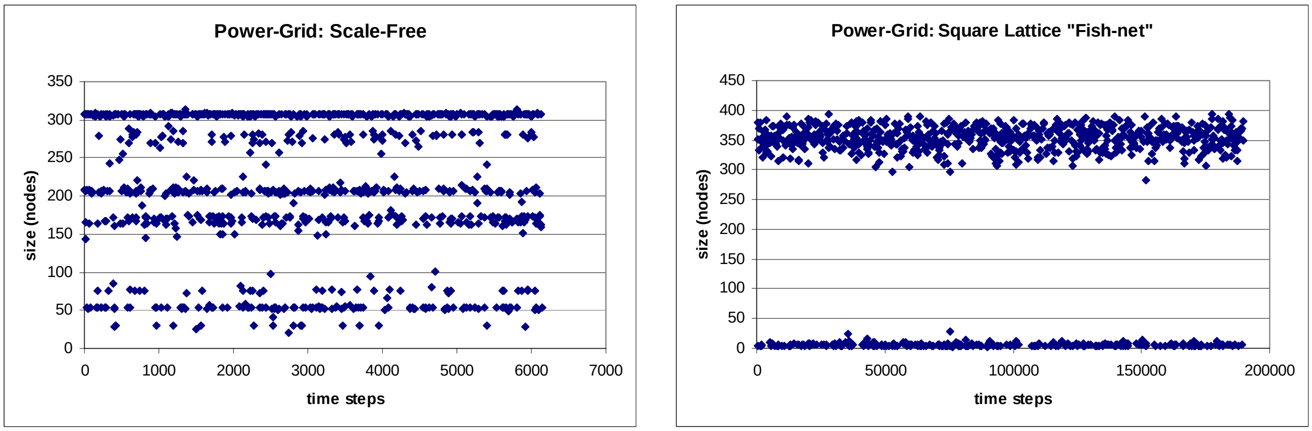

Figure 3.1: Example cascade time series for 1000 events on a 400 node scale-free (to the left) and fish-net (to the right) network. Note that the scale-free network shows a number of cascade size bands well below the size of the network while the fish-net exhibits only very small or very large cascade sizes.

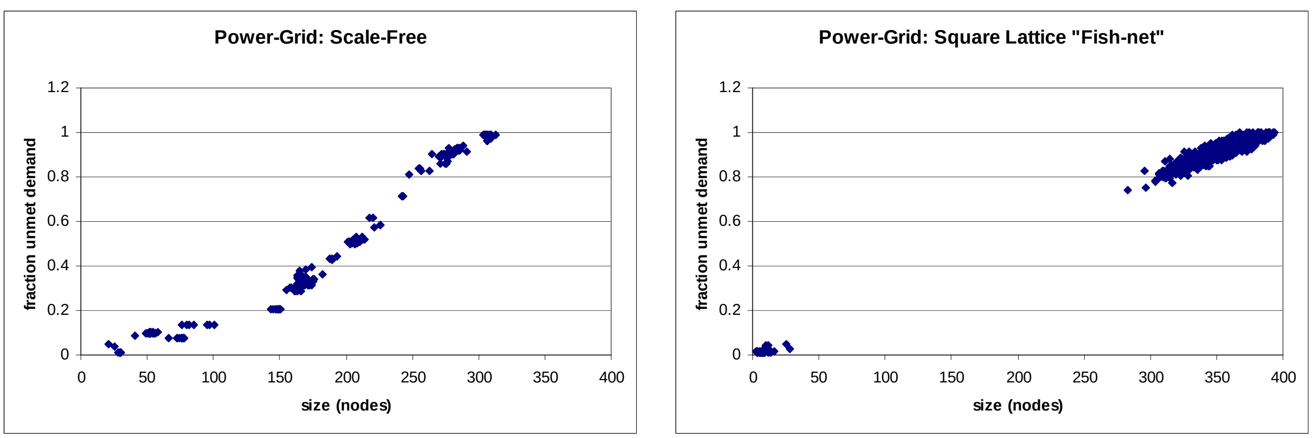

Figure 3.2: Example relation between fraction of unmet demand and cascade size for a 400 node scale-free (to the left) and fish-net (to the right) network.

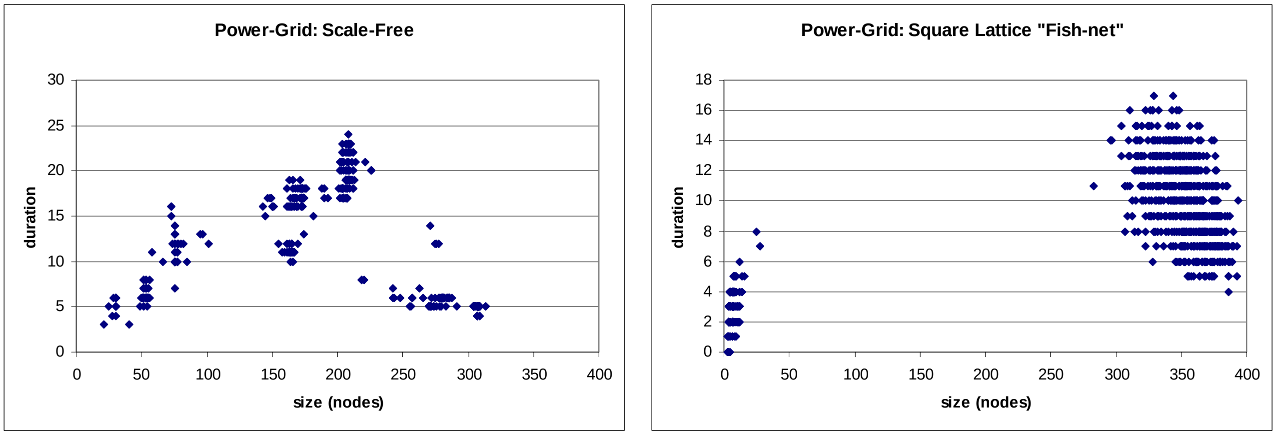

Figure 3.3: Example relation between cascade duration and cascade size on a 400 node scale-free (to the left) and fish-net (to the right) network. Note the clumped nature of the scale-free results.

Our results demonstrate that for the electric power grid cascade application, the distribution of event sizes is not a power-law as has been found in the classical BTW sand-pile (See **Section 2** above). Instead there is a tendency in fish-net networks and a possibility in the scale-free for cascade frequency to *increase* with size. While for the fish-net, most cascades are near the size of the system (and shed near 100% of the demand), in scale-free networks distinct sub-peaks occur, any of which may be the most common, with a high degree of variability from realization to realization. These differences between the BTW sand-pile and the electric power grid cascade results are hypothesized to be due to the combination of 1) our use of global versus local information in the calculation of node state and 2) network topology.

The global calculation of load does two things. First it allows the spreading of failure past nearest neighbors. Although there is only local connectivity among nodes, there is an implicit global interaction as the failure of one node instantaneously changes load at every other node in the network. Changes may be larger near the failed node, but there is the potential for failure to "jump" to any point in the network. Second, and likely more important, there is a tendency for cascades to accelerate because each additional node that fails pushes the loads of all nodes higher. Contrarily, for the BTW sand-pile, only local rules apply and acceleration is not possible. Acceleration is clearly implied in the fish-net results where cascades that reach a size of greater than ~ 10 nodes always run to span nearly the entire network (refer to **Figure 3.2**).

On top of the influence of global information on cascade dynamics, the scale-free network topology imposes preferred conduits for the flow of current. The scale-free topology is also naturally segmented or "hierarchical". This means that when a cascade begins, it can be isolated by the failure of a small number of critical nodes. If the initial allocation of generators and customers along with transaction perturbations have not built an over reliance of one region of the network on others, then sub-regions fail independently. However, if significant generation-consumption dependencies have developed across regions, cascades can spill over from one to other dependent sub-regions of the network. Such a behavior is likely responsible for the banded and clumped nature of the scale-free cascade response. Another influence of the scale-free network topology is that it makes the entire network more fragile to perturbations than the fish-net. On average, it takes ~ 1/30th of the transaction perturbations in the fish-net network to trigger a cascade in the scale-free. This is due to the fact that in the fish-net, there are many more possible paths between any two sources and sinks while in the scale-free, preferential attachment builds hierarchy instead of redundancy.

3.1 Summary and future extensions for the electric power grid application:

Our initial analyses show that small additive power transaction perturbations applied at random and with no geographical restrictions will drive our stylized electric power grids to failure in both fish-net and scale-free topologies. These results support the hypothesis that unrestricted power markets can undermine system stability however they do not yet speak to whether restricting markets will increase system robustness. We find that the fish-net topology is much more robust to additive transaction perturbations than the scale-free (roughly a factor of 30). However, when a cascade occurs, the scale-free naturally compartmentalizes failure while the fish-net does not and thus when the fish-net fails it does so near the scale of the entire network. This topological influence is hypothesized to be due to the fact that the scale-free consists of well-connected sub-regions with grouped power generators and consumers that are connected to other sub-regions through a small number of critical and more vulnerable nodes. Dependencies between sub-regions arise in our analysis through random power transactions that can span the entire network. It is these dependencies that likely cause failure in one sub-network to cascade into others.

Beyond testing the hypothesized influence of network topology, a number of future studies will improve the fidelity of our analyses. First, parametric studies that systematically vary network size, perturbation size, and the distribution of sources/sinks (size and location), would fully disclose present model behavior. Second, actual electric power grid topologies with appropriate node parameterization (resistances, etc.) can be rendered from available data to consider real rather than purely stylized network topologies. Third, we can introduce the added cascade of a "damage surge" that results from voltage spikes emanating from the source of the failure. Fourth, we can consider data-based event distributions for additional perturbation types, such as random human or component failure. Fifth, we can consider the adaptive setting of node thresholds such as to mimic the natural upgrading process that is imposed when a component fails. This study would begin to push beyond the interrogation of a static grid by incorporating active self-healing processes. Accordingly we must extend *Polynet* to include processes that drive the topology, composition, and operational demands of an infrastructure. For the electric power grid application, one such process is the dynamic power market interactively connected with additional networks (information, other markets, social connections, etc.) all of which are driven by a number of externalities such as weather and climate. With such an extended model, attack scenarios and policy related questions could be considered such as:

1) Will geographically restricting transactions decrease the probability of cascade occurrence or the size of the cascade when it does occur?
2) What might the influence of distributed power generation be on electric power grid stability?
3) Are there any general principles beyond those already designed and outlined within the North American Electric Reliability Council operating rules that may aid in the graceful fragmentation of the network either before or during a cascade?

Answering such questions for a variety of network topologies as well as real electric power grids can aid informed policy analysis and allow the creation of new policy that can guide the electric power grid toward increased robustness.

4. Stylized Fedwire *Polynet* Application

Fedwire is a Federal Reserve service for sending large-value payments between banks and other large financial institutions. The average daily volume transmitted within Fedwire is ~ $1.6T, while the total of the reserve account balances supporting this flow is typically only ~ $10B. This extremely efficient use of capital (characterized by the *turnover ratio* of transaction volume to total balances) arises from and depends upon the close coordination of payments among banks. Failure of a participant to make timely payments, either through communications failure or liquidity shortfalls, can affect the ability of payees to fulfill their own obligations. The close coordination engendered by a reliable payment system creates a network of inter-bank dependencies, which is potentially subject to cascade failures in the absence of mitigating interventions.

McAndrews and Potter (2002) describe Fedwire's response during 9-11. While no banks failed, damage to property and communication systems perturbed the coordinated flow of funds and resulted in payment problems among banks. The Federal Reserve intervened by supplying abundant liquidity through discount window loans and open market operation that restored payment coordination. This episode emphasized both 1) concern that rising debt within the system might create a cascade that would bring the system to a standstill, and 2) the usefulness of the discount window and of intraday lending by the central bank for managing market demands for liquidity. To our knowledge, there has been no quantitative simulation of Fedwire that would allow evaluation of both normal and reactive behavior. In particular, understanding of the speed of the cascade process in Fedwire is important to construct intervention scenarios by the Federal Reserve in the event of acute infrastructure failure. The time in which the Federal Reserve must intervene to avert a liquidity crisis may be quite short given the high turnover ratio seen in the Fedwire system.

Here we apply *Polynet* to first simulate the normal behavior of banks conducting random transactions using the Fedwire system, and then consider the disruption and ensuing cascade when a bank is removed from the network. Nodes represent banks and links potential payment exchanges between banks within the network. Payments are assumed to have a uniform (unit) size. During normal operations, connected banks are randomly selected to transact a payment, which increases the debit count at the sending bank and decreases the debit count at the receiving bank. A large number of such transactions are simulated during a specified trading period. At the conclusion of this period each bank has a debit count reflecting its net position. If this debit count exceeds the bank's reserves, it is in an overnight overdraft position. Heavy penalties are imposed for overnight shortfalls. Banks therefore have an incentive to maintain adequate reserves to avoid overdraft penalties. Federal Reserve deposits do not earn interest, so banks also have a strong incentive to keep reserves low. Banks adjust their reserves at the end of each trading period and, over the course of many trading periods, balance overdraft risk against lost opportunity costs for holding reserve balances.

In our simulation, we associate a particular bank's reserves with its threshold and implement adaptive feedback to allow each node to set its threshold during a "training" period composed of a large number of trading periods. Once thresholds have adapted to the trading pattern across the network, we disrupt the payment system by disconnecting a single bank. Disconnection prevents this bank from making payments, while the payments due the disconnected bank are encumbered even though undeliverable. We note that this is simple nearest neighbor interaction within the network of potential payment exchanges. While the model is globally conservative in the normal (training) mode, disconnecting a bank results in a net injection of deficit into the system. The disruption propagates when bank's deficits exceed their thresholds (reserves) due to interaction with a disconnected bank, resulting in suspension of payments by that bank. Suspension injects additional deficit for any subsequent payments due by that bank, which induces suspensions in other banks and thus a deficit cascade.

We initially considered networks of 1000 nodes. Two stylized network topologies were used to explore the effects of heterogeneity in activity and interconnection: a scale-free network with nodes having degrees (number of connections) ranging from 1 to more than 100, and a regular ring network in which each node is connected to the 25 nodes on either side along the ring. For each network, a trading period consisted of 50,000 transactions between randomly selected pairs of adjacent nodes. Over the course of a training period composed of 3000 trading periods, each bank adjusted its threshold (which represents reserve account balance) so that the risk of an overnight overdraft is less than 0.001. The size of the training period was selected by observing the convergence behavior of the thresholds. Ten realizations of each of the two network types were examined.

The connection heterogeneity in the scale-free network naturally leads to concentration of transaction activity at highly connected nodes. To sharpen the comparison with the homogeneous ring network, we enhanced this concentration by making the probability of a bank interacting with a connected bank proportional to the degree of the connected bank. Following the training period, liquidity is not evenly distributed within the network. Banks with higher connectivity participate in more transactions, and therefore require larger reserves to achieve the same level of risk. While the ring yields a normal distribution for reserves (Average of 30.8, Standard deviation of 3.7), the scale-free has heavy tails. This distribution is a direct consequence of the distribution of node connectivity within the network (see **Figure 4.1 left**). Required reserves do not scale linearly with degree, even though the number of transactions increases with degree. This is because the standard deviation of the end-of-period balance, and therefore the level of reserves needed to satisfy a specified risk, scales with the square root of the number of transactions.

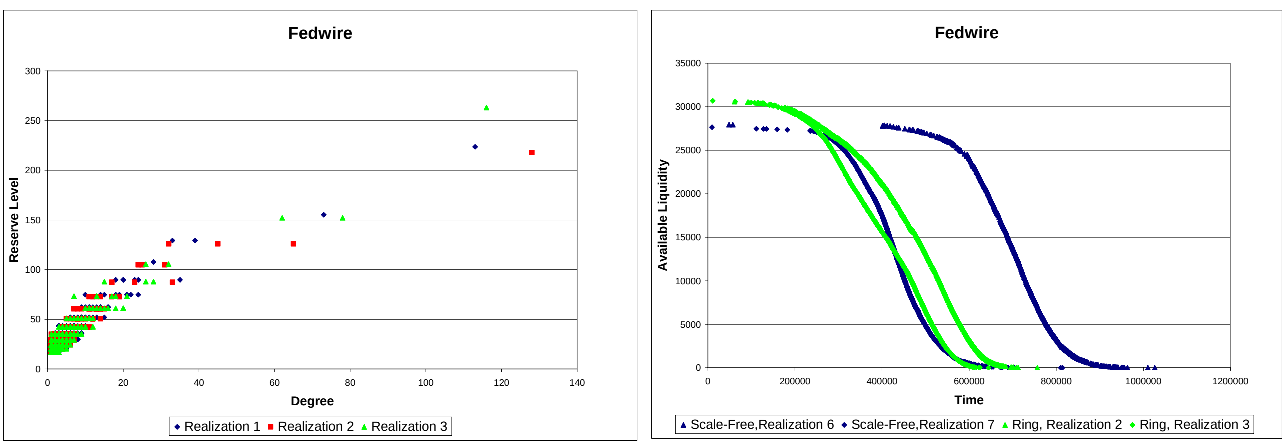

Figure 4.1: Reserve levels at nodes for three realizations of a 1000 node scale-free network are shown to the left. Most and least rapid liquidity decay during cascades instigated by the removal of a random node in 1000 node scale-free and ring networks are shown to the right.

On removal of a single randomly chosen bank, the cascade process begins. **Figure 4.1 right** shows the erosion of total available liquidity (sum of available reserves at all active banks) as a function of time for both the scale-free and ring networks. Simulations that bracket the results seen for the 10 realizations are shown for each network type. The variability among realizations for the scale-free is much larger than the variability seen in the ring case. The scale-free response is characterized by a highly variable delay between the initiating suspension and the first subsequent suspensions. However, after this *initiation period* the cascade response is similar for the sampled scale-free networks, and slightly more rapid than for the ring networks. The decline in liquidity seen in **Figure 4.1 right** is approximately described by a logistic curve. Such a response is seen in classical models of spreading infection where the rate of disease progression is proportional to the interaction rate between infected and uninfected populations. Rates are slow at the beginning and end of the response, when the population is almost entirely uninfected or infected, and are largest when the population sizes are equal. Between 8 and 14 trading periods elapse before half of the banks are suspended. In the scale-free case, the first 5 or 10 suspensions occupy most of this time. Once initiated, approximately 3 to 4 trading periods are required for the cascade to pervade the network.

In the above simulations, the liquidity to which the banks have tuned to control overdraft risk during the training period is ~ 27,000 for the scale-free and ~ 31,000 for the ring – just more than half of the period's trading volume of 50,000 units. This level of reserves indicates some re-use of funds within the system (slightly more in the scale-free due to its inherent heterogeneity), but does not approach the efficiencies seen in Fedwire use. The reserves retained by each bank buffer the propagation of payment shortfalls through the system. Normal Fedwire operations typically have reserves of less than 1/100 of daily volume. However, for the above simulations, reserves settle during the training period to ~ 3/5 of a period's trading volume. In our model, reserves are determined by the fluctuations in the ending balance at each bank, which is in turn determined by the sum of its trading period transactions. Because reserves required to meet a certain risk level are proportional to the standard deviation of this sum, the ratio of total reserves to total transaction volume, which we define as "turnover ratio", is proportional to $N^{-1/2}$, where N is the number of trading period transactions. Thus we can reduce the turnover ratio by increasing the number of transactions seen by each bank. As an illustration, we conducted simulations using a network of 100 rather than 1000 nodes, and an increased transaction frequency resulting in 500,000 rather than 50,000 transactions in a trading period. The reserve levels supporting this higher volume are approximately equal to the original levels, so that the reserve overhead has decreased from 3/5 to 3/50 of the trading period volume. The effect of increased turnover on the loss in liquidity in the system after a random bank failure is shown in **Figure 4.2 left**. We see that thinner reserves lead to much more rapid propagation of disruption: half of the banks are suspended after slightly more than half of a trading period.

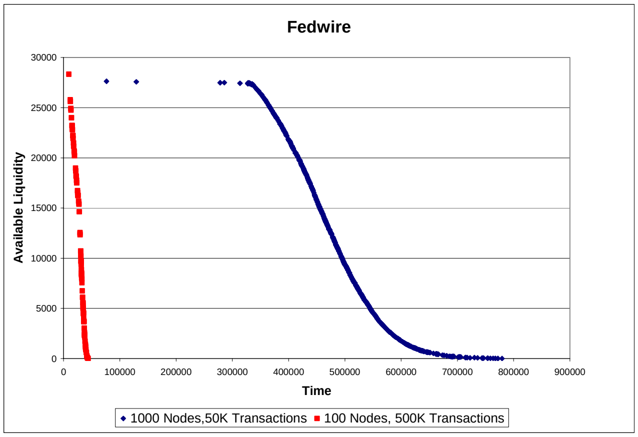

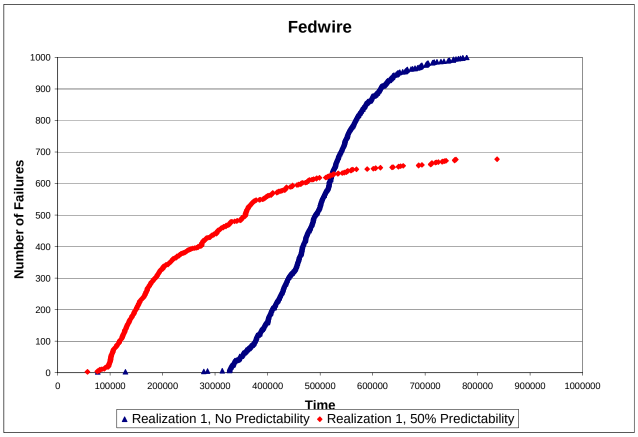

Figure 4.2: The effect of increasing transaction density (decreasing thickness) on liquidity decay in a scale-free network is shown to the left. The effect of adapting to predictable patterned transactions on bank failures in a scale-free network is shown to the right. Both decreasing thickness and patterned transactions accelerate the cascade.

Real wholesale payments also have patterns (i.e., temporal correlation for transactions with particular nodes), and it is reasonable to expect that banks exploit these patterns to manage their reserve accounts. Such patterns can be included explicitly in the model to consider their influence. We have begun exploring the effect of patterned payments by generating a fixed set of recurring transactions. Transactions seen by the system are a mixture of selections from this set of transactions and random transactions generated as before. **Figure 4.2 right** shows the effect on the number of failed banks in time when half of the daily transactions are selected from the recurring set in a 1000 node scale-free network. The time required until 50% of the banks are suspended has decreased from the completely random transaction result of 10 trading periods to slightly more than 7. The overall reserve balance is slightly higher at 29,000 units, but has the same distribution.

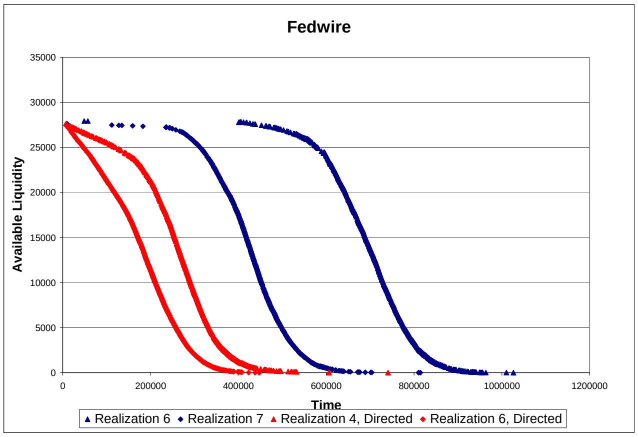

Figure 4.3: Most and least rapid liquidity decay during cascades instigated by the attack of the highest degree node (most connected, red) versus random removal (blue) on a scale-free network. Attack of the highest degree node entirely removes the initiation period for the cascade.

The cascades presented above were all instigated by the initial removal of a single randomly-selected bank. In other studies of cascade phenomena, scale-free networks often exhibit very different behavior depending on whether the initial disruption occurs at a randomly chosen node or at a well-connected node. To consider the susceptibility of the scale-free network to the attack of a highly connected node, we revisited the analysis presented in **Figure 4.1** and instead removed the maximally connected node to begin the cascade. **Figure 4.3** shows the resulting liquidity declines for two bounding network realizations. The corresponding curves for removal of a randomly chosen node from **Figure 4.2** are reproduced for comparison. We find that the long and variable initiation period found for random removal is entirely removed in each attack scenario.

4.1 Summary and extensions for the Fedwire application:

Our initial Fedwire analyses illustrate many expected responses. The simple adaptive response of banks to avoid both overdrafts and excess reserves adjusts individual levels such that the more transactions a bank has, the greater its level. The turnover ratio of the reserves, i.e., the total transaction level relative to the level of the reserves within the entire network, increases with the number of transactions. But with increasing turnover ratio, the speed of liquidity loss when a bank is removed from the network increases dramatically. Overall turnover ratio is a reflection of predictability, and greater predictability reduces the need for bank reserves. The high efficiency of reserve use seen during actual Fedwire operations suggests that payments are highly predictable or controllable. Our initial exploration of patterns in payments, and adaptations to those patterns, demonstrates an additional increase in the speed of liquidity loss on removal of a single bank. We also found that for the attack of the most highly connected bank, the initiation period before the system cascades is entirely elimiated. Thus, we expect that the potent combination of high transaction volume, imposed patterned transactions and high degree node attack may be result in an acute cascade failure for which the required intervention time by the Federal Reserve may be quite small.

Our Fedwire *Polynet* application is in its initial stage and a number of future studies will improve the fidelity of present analyses. First, parametric studies that systematically vary network size and topology, transaction size, and transaction patterning will more fully disclose present model behavior. Second, it would be useful to apply the model to actual Fedwire transaction topologies with appropriate node parameterization (connectivities, reserves, etc.). Additionally, and as in the case of the electric power grid, transactions in Fedwire are made on behalf of their customers. These customers form a trading network, through markets and other relationships, that determines the demand for money transfers. Extending *Polynet* to explicitly include this underlying "economic" drive will constrain the payment activities in the Fedwire model, as well as provide another pathway for propagation of disruptions through the payment system. With such an extended model, policy related questions such as the following could be considered:

1) How fast must the Federal Reserve react after a bank removal is sensed? What are the bounds on injections required to stabilize the system?
2) How might the collapse of Fedwire influence and spread into other markets and infrastructures? Are there measures that could be taken to prevent such a spread?

Answering such questions can aid informed policy analysis and allow the creation of new policy that can guide Fedwire toward increased robustness.

5. Conclusion and Directions for the AMTI task

In this report, we have presented the beginnings of our AMTI effort to apply insights from the study of generic complex adaptive systems to understand the structure, function, and evolution of complex interdependent critical infrastructures. As our first step we have focused on cascading failure as can occur with devastating results within and between infrastructures. We have developed an abstract cascade model, *Polynet*, which encompasses a variety of complex system models. We demonstrated *Polynet* with one fully abstract problem, the BTW sand-pile in various network topologies, and for two critical infrastructure related applications, the electric power grid and Fedwire, both in scale-free and regular homogeneous topologies.

While our critical infrastructure applications are in their initial stage, they allow some conclusions to be drawn and hypotheses to be formulated. Both applications consider the effects of random transactions on idealized networks. However, they differ in the way cascades are initiated and the way failed nodes influence the subsequent cascade process. For the electric power grid, the addition of random transactions eventually pushes the grid to a cascading failure, while for Fedwire, the removal of a bank initiates the cascade. In the electric power grid, state changes are felt instantaneously throughout the network, while payment failures have only local influence on correspondent banks in Fedwire. For both cases, we observe a dependence on network geometry. Network geometry influences the electric power grid's robustness to transaction perturbations, as well as the cascade size disribution. On stylized scale-free electric power grids, cascades appear to be limited to network sub-regions, which we hypothesize are characterized by regional supply/demand imbalance. In contrast, cascades on the stylized fish-net most often occur at the scale of the entire network. For Fedwire, network topology significantly influences the cascade initiation period and for the scale-free network, attacking the most highly connected node can cause the cascade to begin immediately. In our stylized Fedwire application, we also considered the adaptive response of nodes to balance overdraft risk against lost opportunity costs for holding reserve balances. We found that the system naturally evolved to larger turnover ratios as we increased the transaction volume, thus leading to the acceleration of cascade failure. For the electric power grid one can imagine that the adaptive setting of load thresholds to mimic the strengthening of troublesome nodes might cause the scale-free network to evolve enhanced robustness.

As a whole, our studies emphasize the importance of capturing the relevant abstracted "physics", i.e., node interaction, network topology, and adaptive behavior, when considering cascade behavior in infrastructures. For each of the two critical infrastructure applications we have identified ways to better capture the governing dynamics. In both cases, these extensions include connections to other "networks" such as markets, telecommunication, and/or control (e.g., sensor networks that may be an array of widgets or an array of humans connected in a social network). Additional applications such as the spread of infectious diseases within agricultural systems, or the reaction of a crowd to crisis within a constrained physical environment (e.g., a network of mass transit terminals) will also require such extensions.

Networks form and operate in response to external drivers, and understanding the response of a single specific network cannot tell us whether the propensity to cascade is accidental or endemic. This distinction is essential for informing policy. Understanding the interaction of multiple networks, their interdependencies, and in particular, the underlying mechanisms for their growth/evolution is paramount. Because cycles of spontaneous failure and adaptive regeneration can endogenously re-shape the structure and function of networks, the problem cannot be fully considered with a static network as we have done in our initial work reported here. Growing/evolving network topology, and adapting interaction/behavioral rules ultimately determine the shape and performance of infrastructures. If we can understand which characteristics are critical for infrastructure behavior, and understand how infrastructures respond to evolutionary pressures, we can identify appropriate public policy to foster the evolution of robust infrastructures.

Acknowledgements

The authors thank the rest of the AMTI group members for contributing valuable insights over the course of development: George Barr, Mark Bastian, Nancy Brodsky, Steve Conrad, Mark Ehlen, Tina Jenkins, and Andy Scholand. We also appreciate the constructive reviews of this manuscript provided by Theresa Brown, Sharon Deland, and Ralph Keyser.

References

Albert, R., I. Albert, and G.L. Nakarado, Structural vulnerability of the North American power grid, arXiv:cond-mat/0401084v1, 7 Jan 2004.

Bak, P., C. Tang, and K. Wiesenfeld, Self-organized criticality: An explanation of 1/f noise, Physical Review Letters, 59:4:381-384, 1987.

Barabasi, A.-L., and R. Albert, Emergence of scaling in random networks, Science, 286:5439:509-512, 1999.

de Arcangelis, L., and H.J. Herrmann, Self-organized criticality on small world networks, Physica A, 308:545-549, 2002.

Erdos, P., and A. Renyi, On random graphs, I, Publicationes Mathematicae (Debrecen), 6:290-297, 1959.

Glass, R.J., W.E. Beyeler, S.H. Conrad, N.S. Brodsky, P.G. Kaplan, T.J. Brown, Defining research and development directions for modeling and simulation of complex, interdependent adaptive infrastructures, 32 pages (SNL paper SAND 2003-1778P).

Goh, K.-I., D.-S. Lee, B. Kahng, and D. Kim, Sand-pile on scale-free networks, Physical Review Letters, 91(14)148701, 2003.

Kosterev, D.N., C.W. Taylor, and W.A. Mittelstadt, Model validation for the August 10, 1996 WSCC system outage, IEEE Transactions on Power Systems v14, n3, pp. 967-979, 1999.

McAndrews, J.J., S.M. Potter, Liquidity effects of the events of September 11, 2001, FRBNY Economic Policy Review, pgs 59-79, November 2002.

Motter, A.E., and Y.C. Lai, Cascade-biased attacks on complex networks, Physical Review E, 66 065102(R), 2002.

Olami, Z., H.J.S. Feder, and K. Christensen, Self-organized criticality in a continuous, non-conservative cellular automaton modeling earthquakes, Physical Review Letters, 68:1244, 1992.

Sakhtjen, M.L., B.A. Carrerras and V.E. Lynch, Disturbances in a power transmission system, Physical Review E, 61:5:4877-4882, 2000.

Watts, D.J., and S.H. Strogatz, Collective dynamics of ‘small world’ networks, Nature, 393:440-442, 1998.