



Simulation and Analysis of Cascading Failure in Critical Infrastructure

**Robert Glass, Walt Beyeler, Kevin Stamber, Laura Glass,
Randall LaViolette, Stephen Conrad, Nancy Brodsky,
Theresa Brown, Andy Scholand, and Mark Ehlen**

**Advanced Modeling and Techniques Investigations
National Infrastructure Simulation and Analysis Center (NISAC)
Sandia National Laboratories**



NISAC's core partners are Sandia National Laboratories and Los Alamos National Laboratory. Sandia is a multiprogram 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.

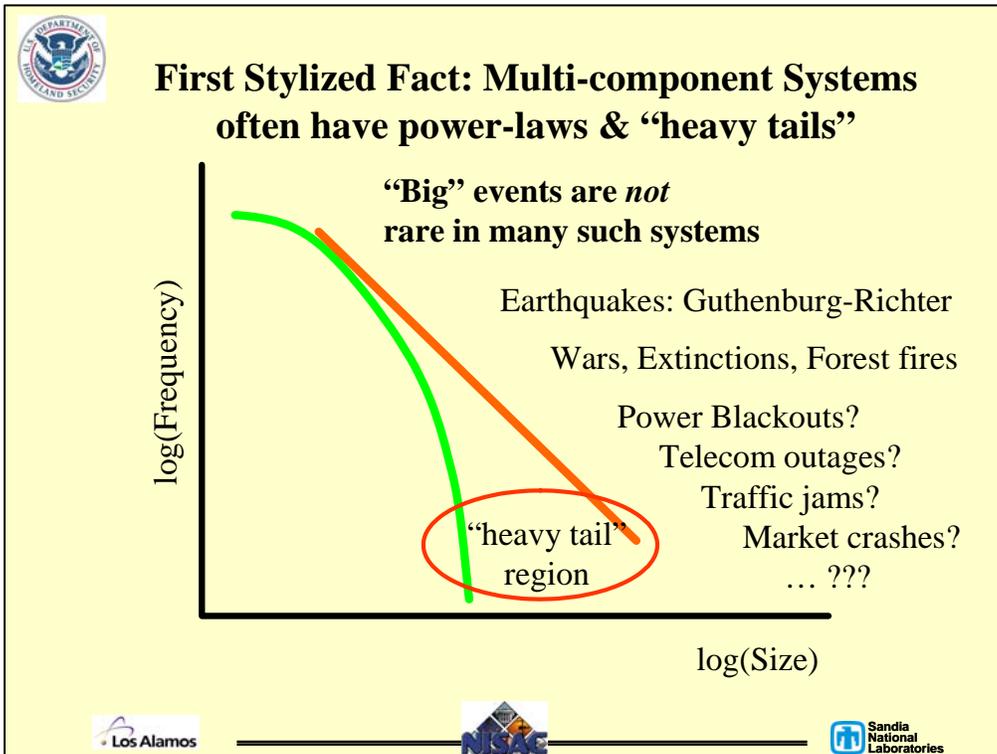
National and economic security and the quality of life in the U.S. depend on reliable operation of complex infrastructures. The **National Infrastructure Simulation and Analysis Center**, or **NISAC**, provides modeling and simulation capabilities for analyzing critical infrastructures. NISAC was founded by Congress in the late 90's as a joint effort between Sandia and Los Alamos National Laboratories. When DHS was formed NISAC moved over to DHS/IAIP where it has continued its efforts. NISAC analyses focus on the such things as projecting the consequences of disruptions in infrastructure services and changes in security policy (power outages, hurricanes, floods, terrorist attacks, security measures, etc.). NISAC combines simulation of the various infrastructures with perturbations (natural and anthropogenic) along with disease and economic models to evaluate consequences to public health, economics of the region, their distribution and duration. A major focus of NISAC is understanding **interdependencies**, quantifying their effects and identifying effective strategies for reducing the potential consequences. We are focused on how and when a perturbation spills over or **cascades** from one infrastructure to another. We use coupled network models, agent-based simulation tools and system dynamics models with feedbacks within and between infrastructures to try to model and understand this process, evaluate consequences, and ultimately suggest mitigation strategies that minimize the compounded effects. Of course, there's a lot of integration that you have to do to play this game.

What I will introduce you to today is some of our work being done within the **Advanced Modeling and Techniques Investigations Task** (AMTI), 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**.

In one of AMTI's efforts, we are focusing on **cascading failure** as can occur with devastating results within and between infrastructures.

References:

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).



First Stylized Fact

Infrastructures are **very large multi-component systems**.

Many multi-component systems exhibit “**heavy tails**” that can often be represented as a **power law** for event frequency as a function of event (or outage) size.

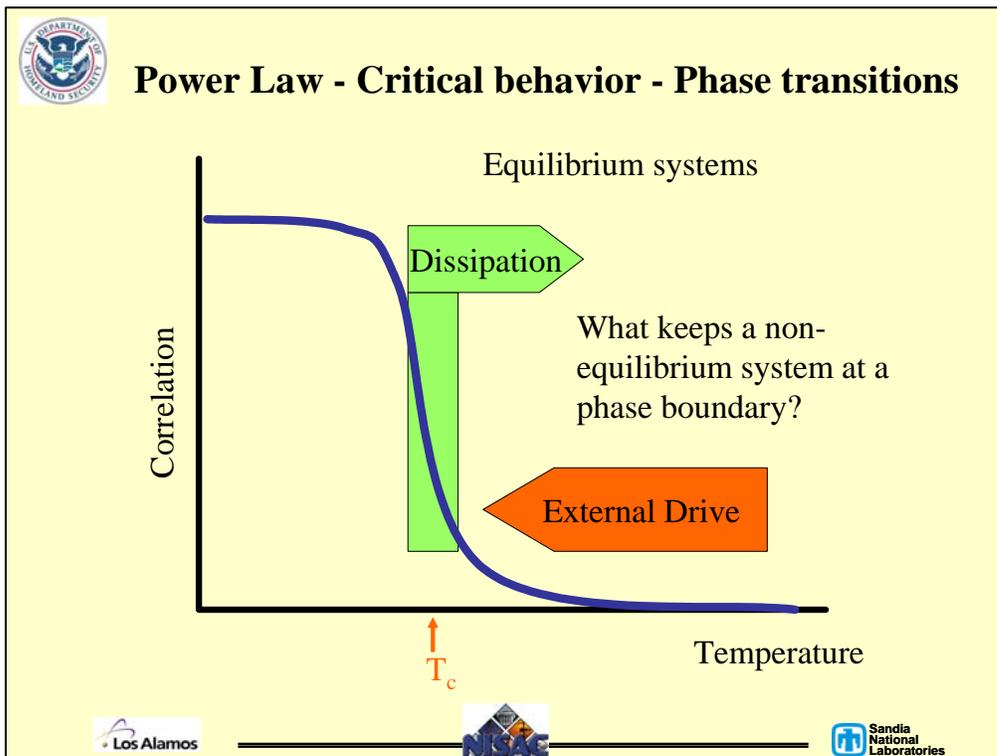
The green curve represents a standard normal distribution while the orange line is a power law. The “**heavy tail**” region of the power-law shows that big events are not rare in such systems.

Power-law behavior is also typical of what has been called “**1/f noise**” found in many natural and anthropogenic systems.

What about infrastructures?

Certainly **Power grid blackouts** have heavy tails, but also **Telecom outages**, **Traffic jams** and **Market crashes** as well.

Note that **roll off** in the power law at both ends occurs in all natural systems of finite size.



What is behind this power-law behavior?

In **equilibrium** systems, **power-laws** are correlated with **critical behavior** as often found at **phase transition boundaries**.

Phase transitions occur at specific **critical points**, T_c , and systems generally must be tuned to be there.

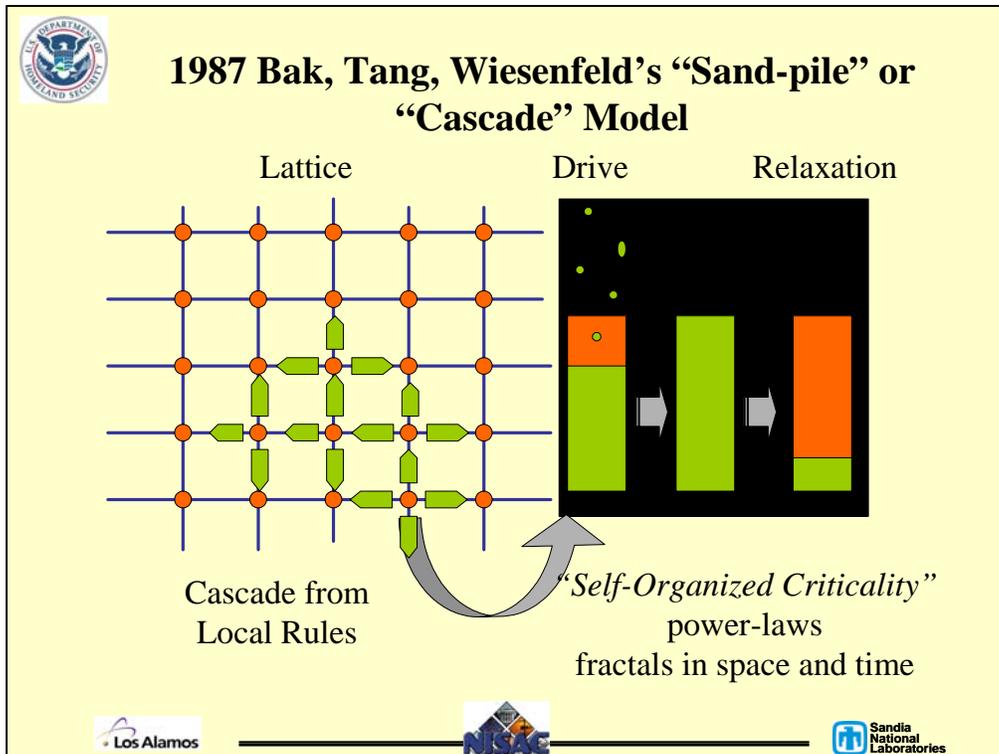
A magnet is a classic example where below the Curie point it behaves collectively (as a magnet) but above, does not.

Percolation theory has been developed to understand system behavior at the critical point where spatial-temporal fractals and power laws emerge.

What about **non-equilibrium** systems?

Many non-equilibrium systems (e.g. BTW sand-pile in the next slide) **maintain themselves in a critical state**. This can occur through the interaction of a driving process which pushes the system in one direction, and a dissipating process which only becomes effective because of properties that emerge (perhaps via long-range correlations) at the phase transition boundary.

For non-equilibrium systems to behave this way, they must be **placed and maintained within an energy gradient**.



BTW Sand-pile or Cascade model

In 1987, **Bak, Tang, and Wiesenfeld** formulated a very simple model that generates **cascades** with power-law distributions within a multi-component system from **simple local rules** operating on a square lattice within a slow random drive: 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.

Model relies on a **separation of time scales** such that the **drive** is very slow relative to the **relaxation** process. Thus cascades evolve to completion before additional sand is applied.

Dissipative system: for the original BTW sand-pile, dissipation occurs only at the boundaries where sand is lost. However, dissipation can occur within the local rule as well (i.e., friction).

This simple model based on local rules creates a state of **Self Organized Criticality** with power law distributions for cascades and fractals in space and time.

Since it's introduction, this simple model has been modified and applied in nearly every scientific field and the original paper has been referenced over 2000 times.

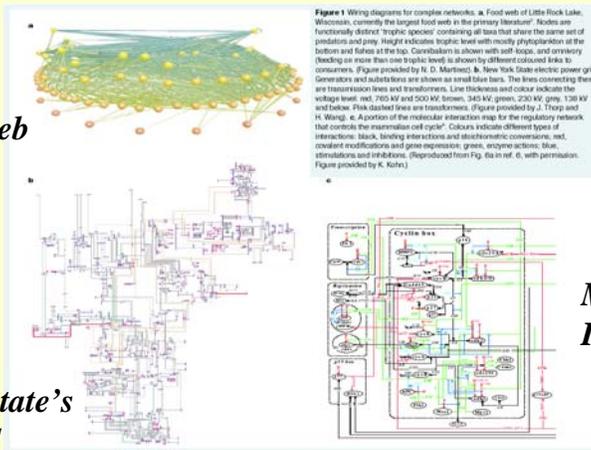
References:

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



Second Stylized Fact: Networks are Ubiquitous in Nature and Infrastructure

Food Web



Molecular Interaction

New York state's Power Grid

Illustrations of natural and constructed network systems from Strogatz [2001].



Second Stylized Fact

Another important feature of many natural and man made systems are that components are linked into **complex and often ramified networks**.

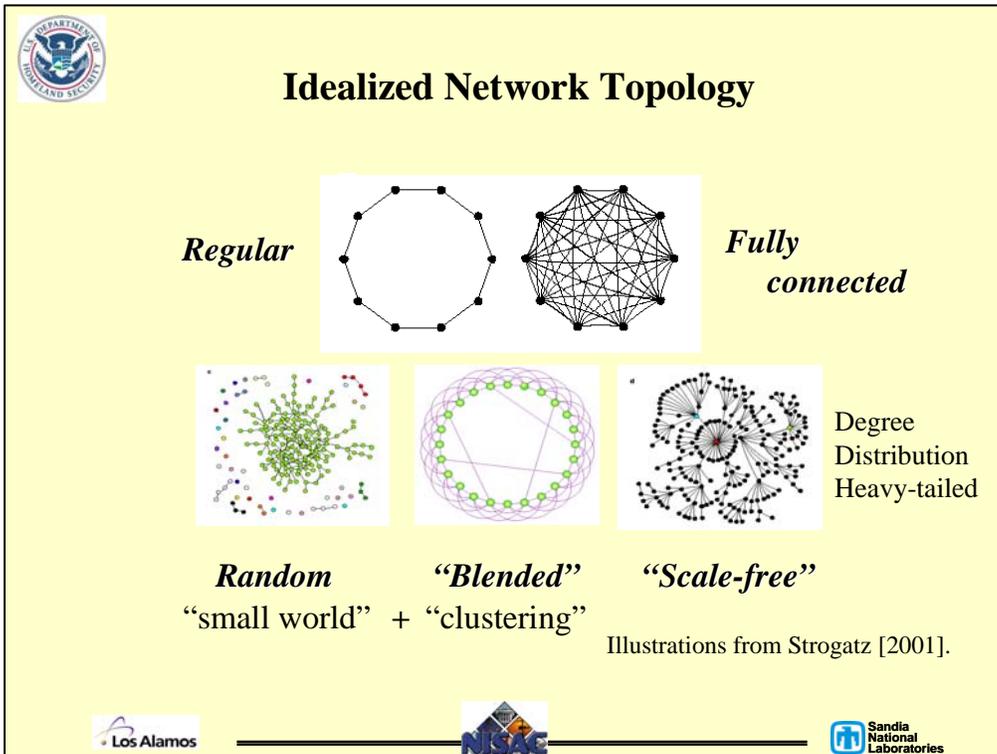
Designed by evolution or by man, **networks are ubiquitous**. Here are three examples from Strogatz (2001).

Nearly every system can be formulated and analyzed as a network!

We find: King pins, keystone species, critical nodes, critical reactions, rate determining steps...

References:

Strogatz, S.H., Exploring Complex networks, Nature, 410:268-276, 2001.



Idealized Network Topology

Graph theorists have generated and explored the properties of many idealized network topologies thus allowing us to identify and classify attributes.

At one end of the spectrum are **perfectly ordered, regular lattices**: crystals are an example. Regular lattices have the property of “**clustering**”, that is, your neighbors are often connected to each other.

At the other end of the spectrum are **disordered, random networks**, first studied by Erdos and Renyi (1959). Such networks are formed by joining two nodes at random and then repeating this selection and joining process over and over until a specified link density is achieved. Random networks have what is called the “**small world**” property, that is, it takes just a few steps within the network to go from one place to another. However, random networks are devoid of clustering.

Blending a Ring lattice with a Random network yields both the small world characteristic and clustering. This was first proposed by Watts and Strogatz (1998) as representative of many social networks.

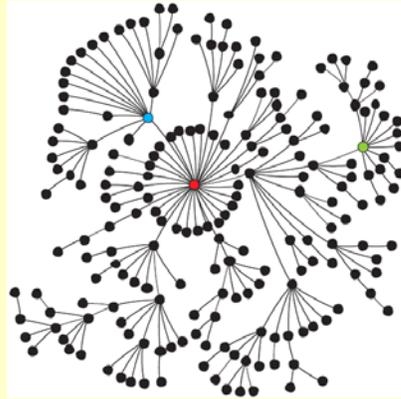
In many naturally occurring networks, one finds a power-law or near power-law for the **nodal degree distribution** such that a significant number of highly connected nodes exist (i.e., a heavy tail). Networks with this power-law distribution have fractal properties and are often called **scale-free** (Barabasi and Albert, 1999).

References:

- Erdos, P., and A. Renyi, On Random Graphs, I, Publicationes Mathematicae (Debrecen), 6:290-297, 1959.
- Watts, D.J., and S.H. Strogatz, Collective dynamics of ‘small world’ networks, Nature, 393:440-442, 1998.
- Barabasi, A.-L., and R. Albert, Emergence of scaling in random networks, Science, 286:5439:509-512, 1999.



Special properties of the “Scale-free” network



Example:
Barabasi and Albert, 1999

Preferential attachment
“rich get richer”

Hierarchical with
“King-pin” nodes

Properties:
tolerant to random
failure...
vulnerable to
informed attack



Special properties of the scale free network

Scale-free networks can be formed by many different processes or models.

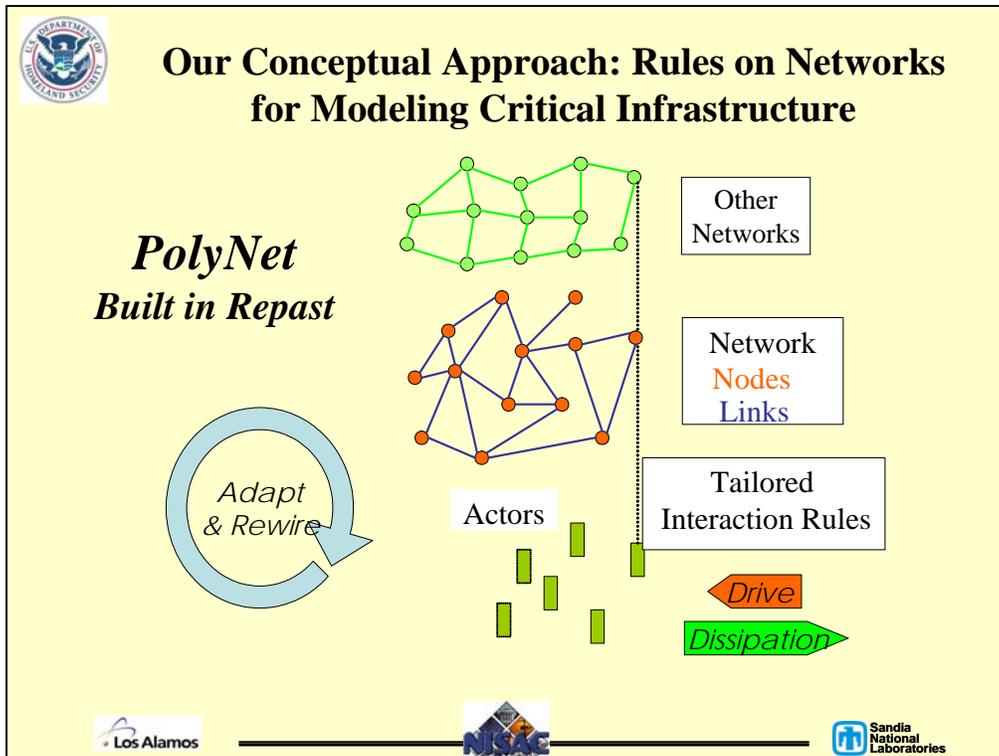
The preferential attachment algorithm of Barabasi and Albert (1999) was used to create the network shown in this slide.

Two additional features that one often finds in real and engineered systems are “**king-pin**” or “**key stone**” nodes that are critical to the operation of the entire system, and **hierarchies or “tree” structures** where some (or all but one) nodes are subservient to others. Both of these features are found in the Scale-free network.

Albert, Jeong and Barabasi (2000) demonstrated the critical properties of such a network: **tolerant to random failure** but **vulnerable to informed attack**. For example, if one chose a node at random to remove from the network in the slide, a degree one node would likely be selected, and its removal would do little to the connectivity at large. But if the red, highest degree node were selected, the network would fragment into many pieces, losing its large scale connectivity.

References:

Albert, R., H. Jeong, A.-L. Barabasi, Error and attack tolerance of complex networks, Nature, 406:378-382, 2000.



Our conceptual approach

We combine the 2 Stylized Facts to formulate our conceptual approach: **Bottom's up Simulation or Rules on Networks.**

Our general approach distills the system of interest to a network (or multiple networks) of nodes and connections with a set of tailored interaction rules (static to adaptive) for each. Combined with drives and dissipations we can evaluate how general features, such as network connectivity and interaction rules, can influence **cascading failure** and the choice of mitigation strategy once a cascade begins.

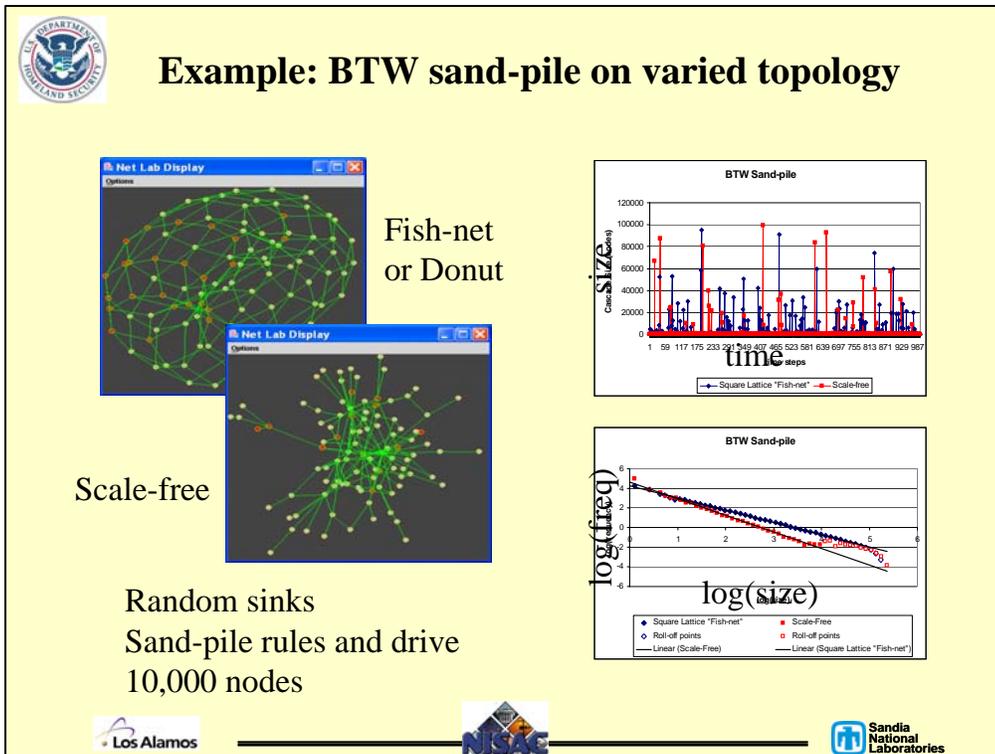
We have rolled all of this into an easily adapted code we call **Polynet**. Polynet is written in Java and inherits many of its classes from **Repast** written by researchers at Argonne National Laboratories and the University of Chicago.

In the remainder of the talk, I will present **one abstract example** (the BTW Sand-pile on arbitrary network topology) and **three very different applications**: 1) Cascading blackouts in the high voltage electric power transmission system which relays electricity from generators to groups of distribution-level consumers; 2) Cascading liquidity loss in payment systems, which are central bank services for sending large-value payments between banks and other large financial institutions; and 3) Cascading epidemics within a structured community of people when vaccine is low. In each, network topology and interaction rules are specifically tailored.

Note that **the Abstract Example and all Applications are highly "stylized"**.

References:

Glass, R.J., W.E. Beyeler, K.L. Stamber, Advanced simulation for analysis of critical infrastructure: Abstract cascades, the electric power grid, and Fedwire, 18 pages (SNL paper SAND 2004-4239).



Abstract Example: BTW sand-pile on varied topology

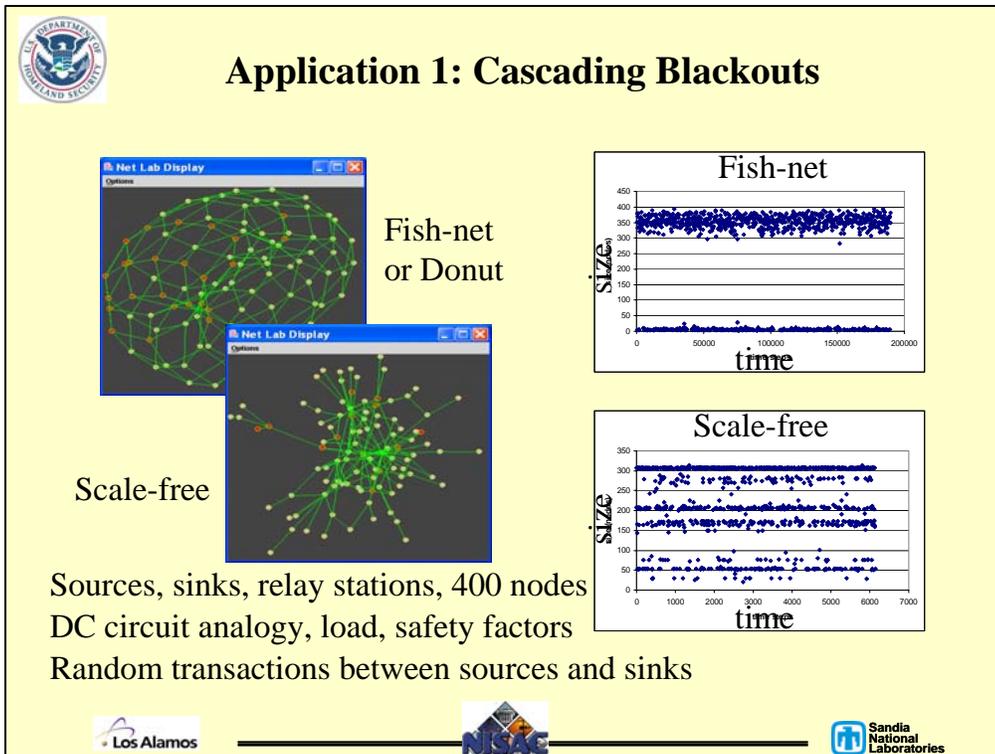
To generalize the BTW sand-pile and apply it to arbitrary network topologies, let us consider grains of sand to represent units of “**energy**”, E , and specify a constant threshold value across all sites, E_c , at which a site changes state and distributes one unit of E to each of its neighboring sites. 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.

Example BTW sand-pile simulations for **10,000 node** problems for fish-net and scale-free stylized network exhibit time series that are **highly erratic** (see **top right plot**). In the plot, cascade size (defined by the number of times nodes in the network are pushed about threshold and distribute E) is shown 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 (see **lower right plot**) 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. We see that the exponent of the power-law (slope of the line) is dependent on the network topology.

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. 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 remainder of the talk, we explore such transformations in context of three applications: the electric power grid, a payment system, and the spread of an infectious disease.

Reference:

Glass, R.J., W.E. Beyeler, K.L. Stamber, Advanced simulation for analysis of critical infrastructure: Abstract cascades, the electric power grid, and Fedwire, 18 pages (SNL paper SAND 2004-4239).



Application 1: Cascading Blackouts

A Stylized power grid is represented by ideal networks that “bracket” what we find in real systems: **regular fish-net lattice** and **scale-free**. Nodes represent **sources, sinks, and relays stations** for electricity. Sources and sinks are assigned **representative values for power grids**. **DC circuit analogy** is solved on the network to yield **loads** at each node and then nodes are given failure loads specified by a uniform **safety factor representative of grid design**. The system is driven by a **random, unregulated market** where pairs of sources and sinks are chosen at random to buy and sell electricity. After each transaction, load is recalculated within the network. This sequence continues until a node is pushed above failure threshold. The failed node is then removed, load is recalculated, nodes which are now pushed above threshold then fail and are removed, etc. The resulting **load based cascade** is followed to its completion. Following a cascade, the network is placed at its initial condition and random transactions are once again accumulated until the next cascade occurs, etc. Cascade size (number of nodes that fail) as a function of time (transactions) for two example networks each containing 400 nodes are shown in the plots on the right.

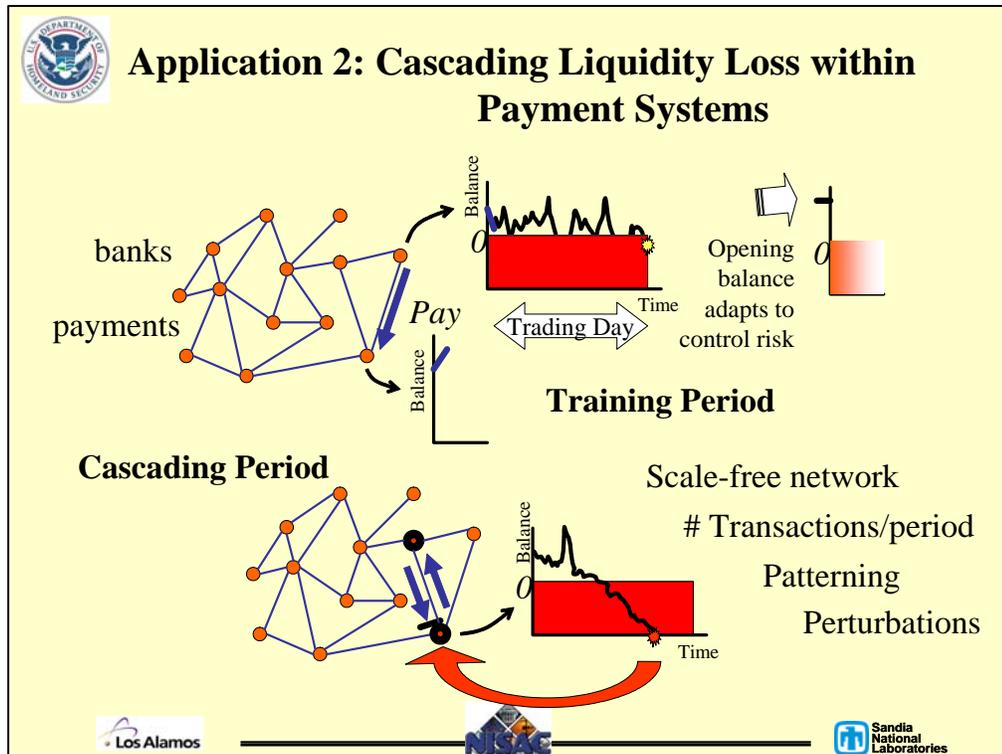
Fishnet: Cascades are either very small, or near the size of the system

Scale-free: sets of cascades occur that are specific to a given network realization and determined by the specifics of the network topology, natural breaks occur that fragment the system when cascades occur.

Also note that the **time scales for the two networks are over 2 orders of magnitude different** suggesting the **fish-net to be much more robust to market perturbations than the scale-free** (i.e., it can accumulate many more perturbations before cascading)

References:

Glass, R.J., W.E. Beyeler, K.L. Stamber, Advanced simulation for analysis of critical infrastructure: Abstract cascades, the electric power grid, and Fedwire, 18 pages (SNL paper SAND 2004-4239).



Application 2: Cascading Liquidity Loss within Payment Systems

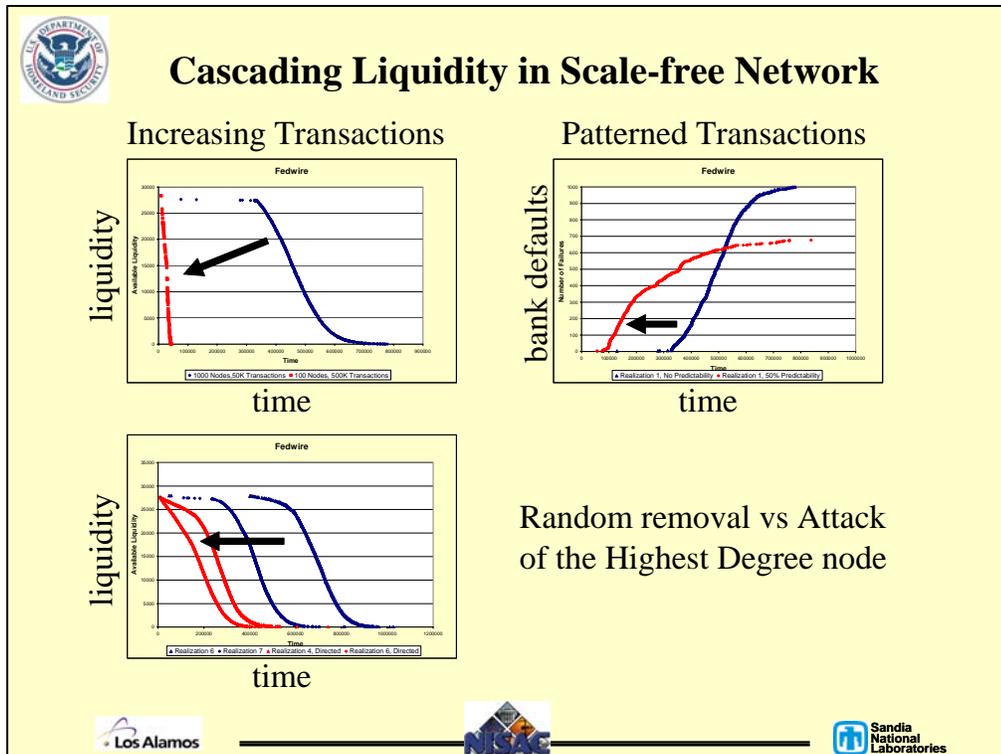
Payment systems allow banks to move money and securities between each other. An example of a payment system is **Fedwire**, the Federal Reserve’s service for sending large-value payments between banks and other large financial institutions. It is worth noting that nearly every central bank across the globe has adopted Fedwire-like payment systems to allow the smooth flow of funds. 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.

In our stylized example, **banks form nodes and transactions are links**. Through a training period (several thousands of trading days) banks adapt their balances to reduce the risk of borrowing or losing the use of funds at the end of the day. We then **trigger a cascade** by removing a bank. Payments to the bank are tied up while payments it should make do not arrive. Neighboring banks eventually default and pass the **cascade of liquidity loss** on to their neighbors contagiously. We consider a **scale-free network** of transactions between banks and study the influence of:

- 1) increasing the **number of transactions/period**
- 2) **patterning** as one might expect to occur due to processes that occur periodically between certain banks
- 3) removing a bank at **random versus selecting the highest degree bank**.

Reference:

Glass, R.J., W.E. Beyeler, K.L. Stamber, Advanced simulation for analysis of critical infrastructure: Abstract cascades, the electric power grid, and Fedwire, 18 pages (SNL paper SAND 2004-4239).



Cascading Liquidity in Scale-free Network

The influence of **increasing the number of transactions** is shown in the upper left plot. Banks adjust balances such that 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, as we increase the number of transactions, we reduce the turnover ratio. **Cascades are greatly accelerated.**

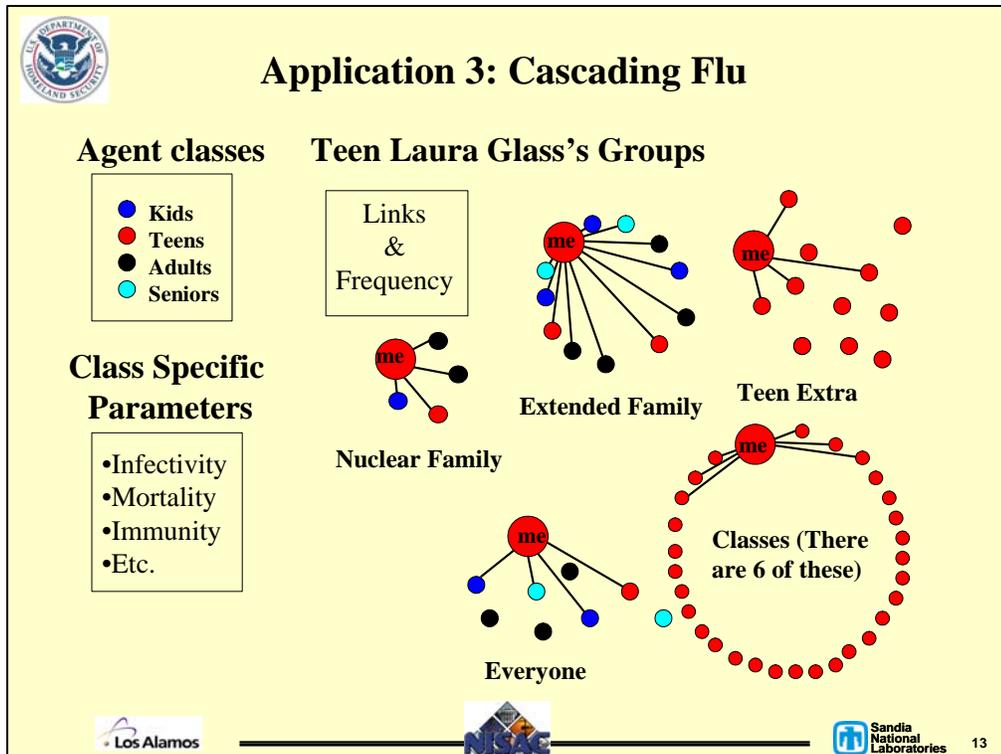
The influence of **patterned transactions** is shown in the upper right plot. Again, banks adjust to the patterning and **cascades are accelerated.**

Random removal vs Attack of the Highest Degree node is shown in the lower left. The two blue curves bracket the set of simulations where random nodes were removed, the red curves bracket the set that were attacked. **Again, cascades are accelerated by attack.**

Important observation: Payment systems should have a very high tendency to cascade, but in practice they don't. The reason for their robustness is due to **the role of the policies of the central bank** that catch and mitigate such cascading, often before it can even start.

Reference:

Glass, R.J., W.E. Beyeler, K.L. Stamber, Advanced simulation for analysis of critical infrastructure: Abstract cascades, the electric power grid, and Fedwire, 18 pages (SNL paper SAND 2004-4239).



Application 3: Cascading Flu

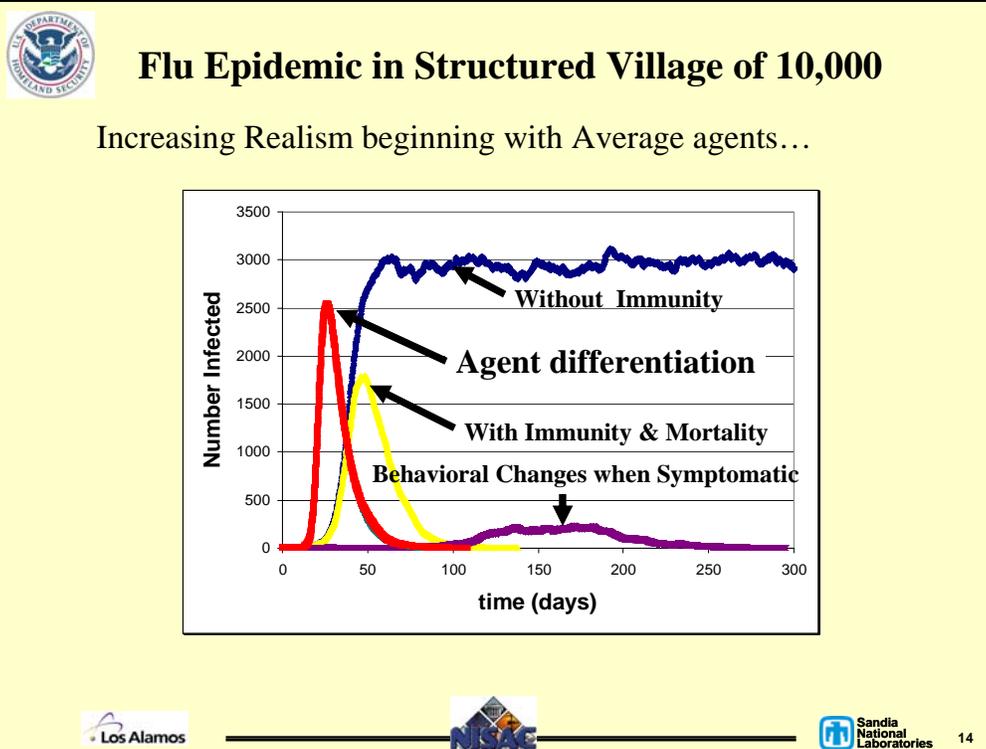
In Fall 2004, there was quite a bit of stir around the lack of usual levels of flu vaccine. The current policy for vaccination places those that have a high risk of death at higher priority for receiving vaccination, especially during a shortage. Does this policy make sense? We can apply **Polynet** to evaluate the **cascade of flu** on a **structured network of social contacts within a community**. This work was performed by Laura Glass under the guidance of Robert Glass and Walter Beyeler and will be presented at the 2005 Intel International Science and Engineering Fair in Phoenix, Arizona, the week of May 9, 2005.

For this application, **Nodes/agents become Kids, Teens, Adults and Seniors** in proportions specified by demographics. To build the social network in a structured community or village, **agents belong to multiple groups also specified by demographics**. Those groups for teen Laura Glass are shown in the slide as an example. Within each group, a **network is imposed of given ideal topology** (random and ring topologies are used in this example) and the **number of links and their frequency of activation specify the network of contacts in time**.

Parameters for the disease are chosen to be representative of the flu: incubation period, infectious period, symptomatic period, average infectivity, infectivity and mobility when symptomatic, mortality, and the inference of immunity.

Agent class specific parameters change the disease parameters in a relative fashion for: infectivity, mobility when symptomatic, mortality, and probability of vaccination.

Once the network of social connection is built, flu epidemics are **instigated with the arrival of 10 infected adults (travelers)** chosen at random from the population. The spread of the disease is followed in time: **number infected, number recovered, number of deaths**, etc.



Flu Epidemic in Structured Village of 10,000

Here we show the **influence of increasing realism** beginning with **average agents** that hold average values for all parameters but reside within a **structured Village of 10,000**.

Without immunity, the flu reaches an average value of 3000 people sick at any one time.

Adding Immunity after being sick and **Mortality**, the number infected rises to below 2000 and then falls to zero after ~100 days. These results would be represented very well by a standard SIR model (based on ordinary differential equations) with an empirical fit for its parameters.

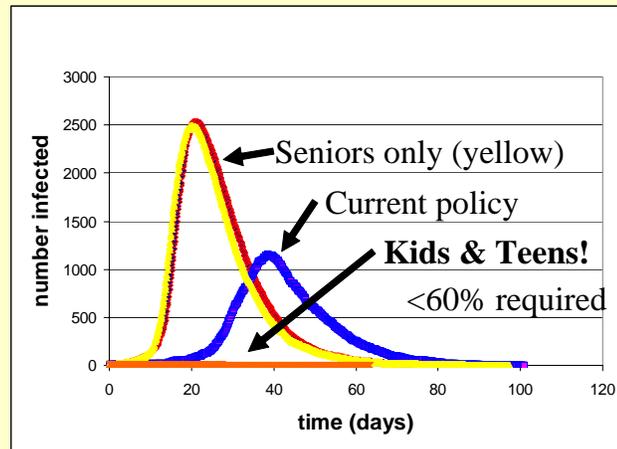
Adding behavioral changes when symptomatic that reduce the contact rate with those that are sick drops the peak infected to below 250 and shifts the peak to later time.

Now, differentiating agents with relative values for parameters such that kids and teens are more contagious and seniors less contagious than adults, the peak is pushed to ~25 days and rises to ~2500!

Having highly infectious Kids and Teens together in schools has a huge influence on the spread of the flu.



Flu Epidemic Mitigation: Vaccination Strategies



Flu Epidemic Mitigation: Vaccination Strategies

We now consider **vaccination strategies**. The red curve is the final, most realistic simulation from the previous slide. No one yet is vaccinated. **Current policy** is shown in blue (26% for kids and teens, 30% for adults, and 59% for seniors) and indeed decreases the speed of the spread of the flu and the number of people that get sick and die. If we had a vaccine shortage, we could choose to vaccinate only the highest risk group – seniors – resulting in the yellow curve. However, this does almost nothing to the spread of the flu, and others within the population die. If instead we vaccinate only the kids and teens, the flu doesn't even spark! In fact **only 60% of the kids and teens in our structured Village need to be vaccinated for the complete suppression of a flu epidemic.**

The projected shortfall of vaccine last fall would have been well within that needed to vaccinate all Kids and Teens. Ira Longini and coworkers at Emory have published similar results just in the past few months using a structured stochastic model that assumes full mixing within each group. **Keeping networks within groups as we do is useful beyond the constraints of this particular application.** In the case of highly virulent flu (likely avian) that breaks out without the pre-development of an effective vaccine, mitigation strategies including antiviral and **behavioral modification** such as quarantine can also be considered in light of such network modeling. Additionally, concepts such as “**shielding**” in context of bio attacks can also be informed through analysis and simulation that focuses on the underlying social network.

This application has shown that the combination of **topology** and **grouping of agents of like properties** is important. This combination is inherent in the concept of “**structure**”. Our results from this application have obvious implications for the prior two applications where stylized topologies were used to represent power grids and payment systems.

Concluding Remarks:

Interested?
rjglass@sandia.gov

Cascading behavior
 Mitigation strategies

Rules on Networks

Developmental directions

Detailed applications with
 Domain experts

Generalization/Abstraction

Critical Infrastructure
 with embedded humans

Complex Adaptive System

Los Alamos

Sandia National Laboratories

Concluding Remarks:

Results of analysis as presented in this talk allow us to better understand how general features, such as **network connectivity** and **interaction rules**, can influence system robustness and the choice of mitigation strategy. In all of these applications we are **working with domain experts** to better represent critical specifics of each application:

- Network topology**
- Refined interaction rules**
- Adaptive response**

Besides the examples shown here today, we are also working ones on information systems, bank systems and reaction in social nets, some of which involve a spatial domain as well. But our effort is also to **generalize and distil** what we learn from specific applications **to further our general understanding of how critical infrastructures can fail.**

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 evolve in response to their multifaceted physical, economic, cultural, and political environments. Simply put, **critical infrastructures are complex adaptive systems.**

Contact us if you are interested! rjglass@sandia.gov