Chapter 4
Modeling Interdependent Networks as Random Graphs: Connectivity and Systemic Risk

R. M. D'Souza, C. D. Brummitt and E. A. Leicht

Abstract Idealized models of interconnected networks can provide a laboratory for studying the consequences of interdependence in real-world networks, in particular those networks constituting society's critical infrastructure. Here we show how random graph models of connectivity between networks can provide insights into shifts in percolation properties and into systemic risk. Tradeoffs abound in many of our results. For instance, edges between networks confer global connectivity using relatively few edges, and that connectivity can be beneficial in situations like communication or supplying resources, but it can prove dangerous if epidemics were to spread on the network. For a specific model of cascades of load in the system (namely, the sandpile model), we find that each network minimizes its risk of undergoing a large cascade if it has an intermediate amount of connectivity to other networks. Thus, connections among networks confer benefits and costs that balance at optimal amounts. However, what is optimal for minimizing cascade risk in one network is suboptimal for minimizing risk in the collection of networks. This work provides tools for modeling interconnected networks (or single networks with mesoscopic structure), and it provides hypotheses on tradeoffs in interdependence and their implications for systemic risk.

R. M. D'Souza (✉) · C. D. Brummitt
University of California, Davis, CA 95616, USA
e-mail: raissa@cse.ucdavis.edu

C. D. Brummitt
e-mail: cbrummitt@math.ucdavis.edu

E. A. Leicht
CABDyN Complexity Center, University of Oxford, Oxford OX1 1HO, UK
e-mail: elizabeth.leicht@wolfson.ox.ac.uk

4.1 Introduction

Collections of networks occupy the core of modern society, spanning technological, biological, and social systems. Furthermore, many of these networks interact and depend on one another. Conclusions obtained about a network's structure and function when that network is viewed in isolation often change once the network is placed in the larger context of a network-of-networks or, equivalently, when viewed as a system composed of complex systems [13, 15]. Predicting and controlling these über-systems is an outstanding challenge of increasing importance because system interdependence is growing in time. For instance, the increasingly prominent "smart grid" is a tightly coupled cyber-physical system that relies on human operators and that is affected by the social networks of human users. Likewise, global financial markets are increasingly intertwined and implicitly dependent on power and communication networks. They are witnessing an escalation in high frequency trades executed by computer algorithms allowing for unanticipated and uncontrolled collective behavior like the "flash crash" of May 2010. Reinsurance companies uncannily forecast the increase of extreme events (in particular in the USA) just weeks before the onslaught of Superstorm Sandy [59] and stressed the urgent need for new scientific paradigms for quantifying extreme events, risk, and interdependence [54].

Critical infrastructure provides the substrate for modern society and consists of a collection of interdependent networks, such as electric power grids, transportation networks, telecommunications networks, and water distribution networks. The proper collective functioning of all these systems enables government operations, emergency response, supply chains, global economies, access to information and education, and a vast array of other functions. The practitioners and engineers who build and maintain critical infrastructure networks have long been cataloging and analyzing the interdependence between these distinct networks, with particular emphasis on failures cascading through coupled systems [19, 21, 29, 42, 51, 55, 56, 60, 61, 63].

These detailed, data driven models are extremely useful but not entirely practical due to the diversity within each infrastructure and due to difficulty in obtaining data. First, each critical infrastructure network is independently owned and operated, and each is built to satisfy distinct operating regimes and criteria. For instance, consider the distinct requirements and constraints of a municipal transportation system versus a region of an electric power grid. Even within a municipal transportation system there exist multiple networks and stakeholders, such as publicly funded road networks and private bus lines and train networks. Second, there are few incentives for distinct operators to share data with others, so obtaining a view of a collection of distinctly owned systems is difficult. Third, the couplings between the distinct types of infrastructure are often only revealed during extreme events; for instance, a natural gas outage in New Mexico in February 2011 caused rolling electric power blackouts in Texas [16]. Thus, even given the most detailed knowledge of individual critical infrastructure systems, it is still difficult to anticipate new types of failures mechanisms (i.e., some failure mechanisms are "unknown unknowns").

Idealized models for interdependent networks provide a laboratory for discovering unknown couplings and consequences and for developing intuition on the new emergent phenomena and failure mechanisms that arise through interactions between distinct types of systems. In fact, the idea of modeling critical infrastructure as a collection of "complex interactive networks" was introduced over a decade ago [3]. Yet idealized models are only starting to gain traction [58, 71], and they are largely based on techniques of random graphs, percolation and dynamical systems (with many tools drawn from statistical physics). Despite using similar techniques, these models can lead to contrasting conclusions. Some analytic formulations show that interdependence makes systems radically more vulnerable to cascading failures [15], while others show that interdependence can confer resilience to cascades [13].
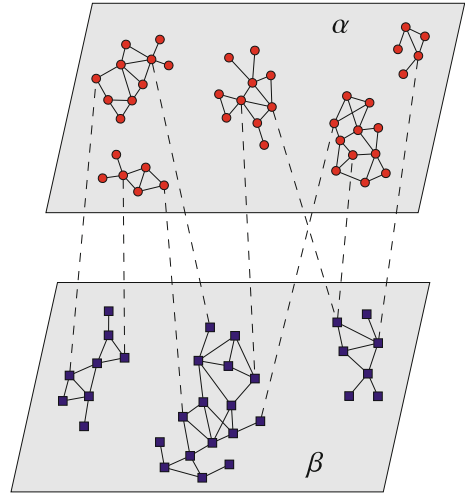
Given a specified set of network properties, such as a degree distribution for the nodes in the network, random graph models consider the ensemble of all graphs that can be enumerated consistent with those specified properties. One can use probability generating functions to calculate the average or typical properties of this ensemble of networks. In the limit of an infinitely large number of nodes, the generating functions describing structural and dynamic properties are often exactly solvable [52], which makes random graphs appealing models that are widely used as simple models of real networks. Of course there are some downsides to using the random graph approach, which will require further research to quantify fully. First, in the real-world we are typically interested in properties of individual instances of networks, not of ensemble properties. Second, percolation models on random graphs assume local, epidemic-like spreading of failures. Cascading failures in the real-world, such as cascading blackouts in electric power grids, often exhibit non-local jumps where a power line fails in one location and triggers a different power line hundreds of miles away to then fail (e.g., see Ref. [1]). This issue is discussed in more detail below in Sect. 4.3.4.1. Nonetheless, random graphs provide a useful starting point for analyzing the properties of systems of interdependent networks.

Here, in Sect. 4.2 we briefly review how random graphs can be used to model the structural connectivity properties between networks. Then, in Sect. 4.3 we show how, with the structural properties in place, one can then analyze dynamical process unfolding on interconnected networks with a focus on cascades of load shedding.

4.2 Random Graph Models for Interconnected Networks

Our model of "interconnected networks" consists of multiple networks (i.e., graphs) with edges introduced between them. Thus, the system contains multiple kinds of nodes, with one type of node for each network, and one type of edge. A simple illustration of a system of two interconnected networks is shown in Fig. 4.1. (A related class of graphs called multiplex networks considers just one type of node but multiple kinds of edges [49, 70].) This general framework can model different kinds of systems that have connections to one another, or it can capture mesoscopic structure in a single network, such as communities and core-periphery structure.

Fig. 4.1 A stylized illustration of two interconnected networks, *a* and *b*. Nodes interact directly with other nodes in their immediate network, yet also with nodes in the second network
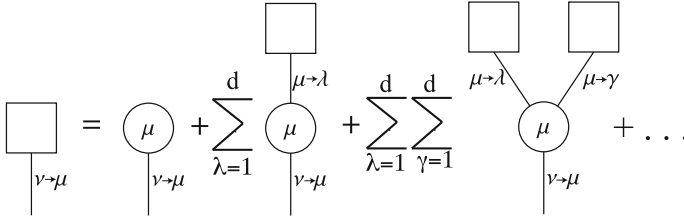

4.2.1 Mathematical Formulation

Here we briefly review the mathematics for calculating the structural properties of interconnected networks as discussed in Ref. [40]. In a system of $d \geq 2$ interacting networks, an individual network μ is characterized by a multi-degree distribution $\{p_{\mathbf{k}}^{\mu}\}$, where $\mathbf{k}$ is a d-tuple, $(k_1, \ldots, k_d)$, and $p_{\mathbf{k}}^{\mu}$ is the probability that a randomly chosen node in network μ has k_ν connections with nodes in network ν. A random graph approach considers the ensemble of all possible networks consistent with this multi-degree distribution. To realize a particular instance of such a network we take the "configuration model" approach [10, 47]. Starting from a collection of isolated nodes, each node independently draws a multi-degree vector from $\{p_{\mathbf{k}}^{\mu}\}$. Next, each node is given k_ν many "edge stubs" (or half-edges) of type ν. We create a graph from this collection of labeled nodes and labeled edge stubs by matching pairs of compatible edge stubs chosen uniformly at random. For instance, an edge stub of type ν belonging to a node in network μ is compatible only with edge stubs of type μ belonging to nodes in network ν. Generating functions allow us to calculate the properties of this ensemble.

The generating function for the $\{p_{\mathbf{k}}^{\mu}\}$ multi-degree distribution is

$$G_\mu(\mathbf{x}) = \sum_{k_1=0}^{\infty} \cdots \sum_{k_d=0}^{\infty} p_{\mathbf{k}}^{\mu} \prod_{\nu=1}^{d} x_\nu^{k_\nu}, \tag{4.1}$$

where $\mathbf{x}$ is the d-tuple, $\mathbf{x} = (x_1, \ldots, x_d)$. This is a generating function for a probability distribution already known to us (our multi-degree distribution for network μ), and thus not terribly informative on its own. However, we can derive additional generating functions for probability distributions of interest, such as the distribution of sizes of connected components in the system. However, we much first derive

Fig. 4.2 A diagramatical representation of the topological constraints placed on the generating function $H_{\mu\nu}(\mathbf{x})$ for the distribution of sizes of components reachable by following a randomly chosen ν-μ edge. The labels attached to each edge indicate type or *flavor* of the edge, and the sum runs over over all possible flavors

two intermediate generating function forms, one for the probability distribution of connectivity for a node at the end of a randomly chosen edge and a second for the probability distribution of component sizes found at the end of a random edge. Reference [52] contains a clear and thorough discussion of this approach for a single network, which we apply here to multiple networks.

First consider following an edge from a node in network ν to a node in network μ. The μ node is k_ν times more likely to have ν-degree k_ν than degree 1. Thus the probability $q_{\mathbf{k}}^{\mu\nu}$ of reaching a μ-node of ν-degree k_ν is proportional to $k_\nu p_{k_1 \cdots k_\nu \cdots k_d}^{\mu}$. Accounting for the fact that we have followed an edge from a node in ν to a node in μ, the properly normalized generating function for the distribution of additional edges from that μ-node is

$$G_{\mu\nu}(\mathbf{x}) = \sum_{k_1=0}^{\infty} \cdots \sum_{k_d=0}^{\infty} \frac{(k_\nu + 1) p_{k_1 \cdots (k_\nu+1) \cdots k_l}^{\mu}}{\overline{k}_{\mu\nu}} \prod_{\gamma=1}^{d} x_\gamma^{k_\gamma} = \frac{G_{\mu}^{'\nu}(\mathbf{x})}{G_{\mu}^{'\nu}(\mathbf{1})}. \quad (4.2)$$

Here $\overline{k}_{\mu\nu} = \sum_{k_1} \cdots \sum_{k_d} k_\nu p_{\mathbf{k}}^{\mu}$ is the normalization factor accounting for $G_{\mu\nu}(\mathbf{1}) = 1$ and $\overline{k}_{\mu\nu}$ is also the average ν-degree for a node in network μ. We use $G_{\mu}^{'\nu}(\mathbf{x})$ to denote the first derivative of $G_{\mu}(\mathbf{x})$ with respect to x_ν and thus $G_{\mu}^{'\nu}(\mathbf{1}) = \overline{k}_{\mu\nu}$. A system of d interacting networks has d^2 excess degree generating functions of the form shown in Eq. 4.2.

Now consider finding, not the connectivity of the μ-node, but the size of the connected component to which it belongs. This probability distribution for sizes of components can be generated by iterating the random-edge-following process described in Eq. 4.2, where we must consider all possible types of nodes that could be attached to that μ-node. For an illustration see Fig. 4.2. In other words, the μ-node could have no other connections; it might be connected to only one other node and that node could belong to any of the d networks; it might be connected to two other nodes that could each belong to any of the d networks; and so on. Iterating the

random-edge construction for each possibility leads to a generating function $H_{\mu\nu}$ for the sizes of components at the end of a randomly selected edge

$$H_{\mu\nu}(\mathbf{x}) = x_\mu q_{0\cdots0}^{\mu\nu} \qquad (4.3)$$

$$+ x_\mu \sum_{k_1\ldots k_d=0}^{1} \delta_{1,\sum_{\lambda=1}^{d} k_\lambda} q_{k_1\cdots k_d}^{\mu\nu} \prod_{\gamma=1}^{d} H_{\gamma\mu}(\mathbf{x})^{k_\gamma}$$

$$+ x_\mu \sum_{k_1,\ldots,k_d=0}^{2} \delta_{2,\sum_{\lambda=1}^{d} k_\lambda} q_{k_1\cdots k_d}^{\mu\nu} \prod_{\gamma=1}^{d} H_{\gamma\mu}(\mathbf{x})^{k_\gamma} + \cdots,$$

where δ_{ij} is the Kronecker delta. Reordering the terms, we find that $H_{\mu\nu}$ can be written as a function of $G_{\mu\nu}$ as follows:

$$H_{\mu\nu}(\mathbf{x}) = x_\mu \sum_{k_1=0}^{\infty} \cdots \sum_{k_d=0}^{\infty} q_{k_1\cdots k_d}^{\mu\nu} \prod_{\gamma=1}^{d} H_{\gamma\mu}(\mathbf{x})^{k_\gamma}$$

$$= x_\mu G_{\mu\nu}[H_{1\mu}(\mathbf{x}), \ldots, H_{d\mu}(\mathbf{x})]. \qquad (4.4)$$

Here again, for a system of d networks, there are d^2 self-consistent equations of the form shown in Eq. 4.4.

Now instead of selecting an edge uniformly at random, consider a node chosen uniformly at random. This node is either isolated or has edges leading to other nodes in some subset of the d networks in the system. The probability argument above allows us to write a self-consistency equation for the distribution in component sizes to which a randomly selected node belongs:

$$H_\mu(\mathbf{x}) = x_\mu G_\mu[H_{1\mu}(\mathbf{x}), \ldots, H_{d\mu}(\mathbf{x})]. \qquad (4.5)$$

With this relation for H_μ, we can now calculate the distribution of component sizes and the composition of the components in terms of nodes from various networks. However, our current interest is not in finding the exact probability distribution of the sizes of connected components, but in finding the emergence of large-scale connectivity in a system of interacting networks. To address this problem, we need only to examine the average component size to which a randomly chosen node belongs. For example, the average number of ν-nodes in the component of a randomly chosen μ-node is

$$\langle s_\mu \rangle_\nu = \frac{\partial}{\partial x_\nu} H_\mu(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{1}}$$

$$= \delta_{\mu\nu} G_\mu[H_{1\mu}(\mathbf{1}), \ldots, H_{d\mu}(\mathbf{1})]$$

$$+ \sum_{\lambda=1}^{d} G_\mu^{'\lambda}[H_{1\mu}(\mathbf{1}), \ldots, H_{d\mu}(\mathbf{1})] H_{\lambda\mu}^{'\nu}(\mathbf{1})$$

$$= \delta_{\mu\nu} + \sum_{\lambda=1}^{d} G_{\mu}^{'\lambda}(\mathbf{1}) H_{\lambda\mu}^{'\nu}(\mathbf{1}). \tag{4.6}$$

Table 4.1 shows the explicit algebraic expressions derived from Eq. 4.6 for a system of $d = 2$ networks with two different forms of internal degree distribution and types of coupling between networks. Where the algebraic expression for $\langle s_\mu \rangle_\nu$ diverges marks the percolation threshold for the onset of a giant component. For instance, the first case shown in Table 4.1 is for two networks, a and b, with internal Poisson distributions, coupled by a third Poisson distribution. For this situation, the percolation threshold is defined by the expression $(1 - \bar{k}_{aa})(1 - \bar{k}_{bb}) = \bar{k}_{ab}\bar{k}_{ba}$.

4.2.2 Consequences of Interactions

To quantify the consequences of interaction between distinct networks, we want to compare results obtained from the calculations above to a corresponding baseline model of a single, isolated network. Interesting differences already arise for the case of $d = 2$ interacting networks, which we focus on here. Consider two networks, a and b, with n_a and n_b nodes respectively. They have multi-degree distributions $p_{k_a k_b}^a$ and $p_{k_a k_b}^b$ respectively. The reference single network, $\mathcal{C}$, neglects the network membership of the nodes. It is of size $n_\mathcal{C} = n_a + n_b$ nodes, and has degree distribution

$$p_k = \left[f_a \sum_{k_a,k_b=0}^{k} \left(p_{k_a k_b}^a \delta_{k_a+k_b,k} \right) + f_b \sum_{k_a,k_b=0}^{k} \left(p_{k_a k_b}^b \delta_{k_a+k_b,k} \right) \right],$$

where $f_a = n_a/(n_a+n_b)$ and $f_b = n_b/(n_a+n_b)$. In other words, network $\mathcal{C}$ is a composite view that neglects whether a node belongs to network a or b. So a node that had degree $\{k_a, k_b\}$ in the interacting network view has degree $k = k_a + k_b$ in the composite, $\mathcal{C}$, view. We compare the properties of the ensemble of random graphs constructed from the interconnected networks multi-degree distribution, $\left\{ p_{k_a k_b}^a, p_{k_a k_b}^b \right\}$, to the properties of the ensemble constructed from the composite, p_k, degree distribution (Fig. 4.3).

In Ref. [39], we analyze the situation for two networks with distinct internal Poisson distributions coupled together via a third Poisson distribution. We show that large-scale connectivity can be achieved with fewer total edges if the network membership of the node is accounted for (i.e., the composite $\mathcal{C}$ view requires more edges to achieve a giant component).

Next we show that other effects are possible for different types of networks. For instance, the degree distributions that are a truncated power law describe many real-world networks, such as the connectivity between Autonomous Systems in the Internet and connectivity patterns in social contact networks [20]. Yet many critical infrastructure networks (such as adjacent buses in electric power grids) have very
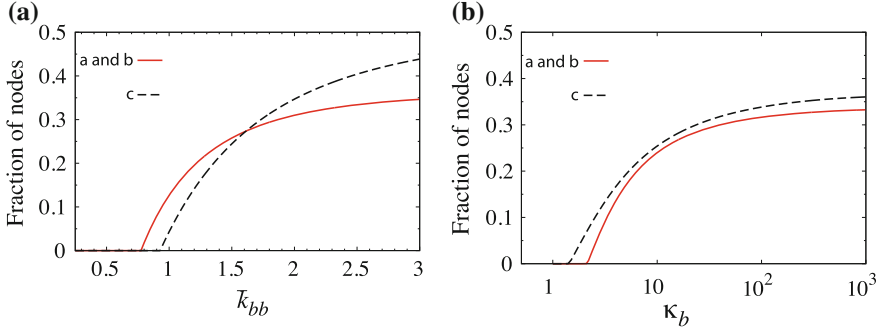
Table 4.1 Expressions for average component size by node type for three different interacting network topologies

Network topology				Average node count by type and initial network
a-a	*a-b*	*b-a*	*b-b*	
Distribution parameters				$\langle s_a\rangle_a$
Generating functions				$\langle s_a\rangle_b$
				$\langle s_b\rangle_a$
				$\langle s_b\rangle_b$
Poisson	Poisson	Poisson	Poisson	$1+\dfrac{\bar{k}_{aa}[1-\bar{k}_{bb}]+\bar{k}_{ab}\bar{k}_{ba}}{(1-\bar{k}_{aa})(1-\bar{k}_{bb})-\bar{k}_{ab}\bar{k}_{ba}}$
$\bar{k}_{aa}$	$\bar{k}_{ab}$	$\bar{k}_{ba}$	$\bar{k}_{bb}$	$\dfrac{\bar{k}_{ab}}{(1-\bar{k}_{aa})(1-\bar{k}_{bb})-\bar{k}_{ab}\bar{k}_{ba}}$
	$G_a(x_a,\mathbf{b})=e^{\bar{k}_{aa}(x_a-1)}e^{\bar{k}_{ab}(x_b-1)}$			$\dfrac{\bar{k}_{ba}}{(1-\bar{k}_{aa})(1-\bar{k}_{bb})-\bar{k}_{ab}\bar{k}_{ba}}$
	$G_b(x_a,\mathbf{b})=e^{\bar{k}_{ba}(x_a-1)}e^{\bar{k}_{bb}(x_b-1)}$			$1+\dfrac{\bar{k}_{bb}[1-\bar{k}_{aa}]+\bar{k}_{ab}\bar{k}_{ba}}{(1-\bar{k}_{aa})(1-\bar{k}_{bb})-\bar{k}_{ab}\bar{k}_{ba}}$
Power-law	Poisson	Poisson	Poisson	$1+\dfrac{\bar{k}_{aa}[1-\bar{k}_{bb}]+\bar{k}_{ab}\bar{k}_{ba}[1-G'^{\alpha}_{\alpha\alpha}(1)+G'^{\alpha}_{\alpha\beta}(1)]}{[1-G'^{\alpha}_{\alpha\alpha}(1)][1-\bar{k}_{bb}]-\bar{k}_{ab}\bar{k}_{ba}[1-G'^{\alpha}_{\alpha\alpha}(1)+G'^{\alpha}_{\alpha\beta}(1)]}$
τ_a,κ_a	$\bar{k}_{ab}$	$\bar{k}_{ba}$	$\bar{k}_{bb}$	$\dfrac{\bar{k}_{ab}[1-G'^{\alpha}_{\alpha\alpha}(1)+G'^{\alpha}_{\alpha\beta}(1)]}{[1-G'^{\alpha}_{\alpha\alpha}(1)][1-\bar{k}_{bb}]-\bar{k}_{ab}\bar{k}_{ba}[1-G'^{\alpha}_{\alpha\alpha}(1)+G'^{\alpha}_{\alpha\beta}(1)]}$
	$G_a(x_a,\mathbf{b})=\dfrac{\mathrm{Li}_{\tau_a}(x_a e^{-1/\kappa_a})}{\mathrm{Li}_{\tau_a}(e^{-1/\kappa_a})}e^{\bar{k}_{ba}(x_b-1)}$			$\dfrac{\bar{k}_{ba}[1+\bar{k}_{aa}-G'^{\alpha}_{\alpha\alpha}(1)]}{[1-G'^{\alpha}_{\alpha\alpha}(1)][1-\bar{k}_{bb}]-\bar{k}_{ab}\bar{k}_{ba}[1-G'^{\alpha}_{\alpha\alpha}(1)+G'^{\alpha}_{\alpha\beta}(1)]}$
	$G_b(x_a,\mathbf{b})=e^{\bar{k}_{ba}(x_a-1)}e^{\bar{k}_{bb}(x_b-1)}$			$1+\dfrac{\bar{k}_{bb}[1-G'^{\alpha}_{\alpha\alpha}(1)]+\bar{k}_{ab}\bar{k}_{ba}[1-G'^{\alpha}_{\alpha\alpha}(1)+G'^{\alpha}_{\alpha\beta}(1)]}{[1-G'^{\alpha}_{\alpha\alpha}(1)][1-\bar{k}_{bb}]-\bar{k}_{ab}\bar{k}_{ba}[1-G'^{\alpha}_{\alpha\alpha}(1)+G'^{\alpha}_{\alpha\beta}(1)]}$

Fig. 4.3 Comparing random graph models which account for interacting networks (*red line*) to random graph models with the identical degree distribution, but which neglect network membership (*dashed black line*). **a** The fraction of nodes in the largest connected component for two interconnected networks with Poisson degree distribution, as edges are added to network b. Accounting for network structure allows for a giant component to emerge with fewer edges. Here $n_a = 4n_b$. **b** The corresponding fractional size of the giant component for a network with a Poisson degree distribution coupled to a network with a truncated power law degree distribution as the power law regime is extended. Here was see the opposite effect to **a**, where large scale connectivity is delayed by accounting for network membership

narrow degree distributions, which we approximate here as Poisson. Thus, we are interested in the consequences of coupling together networks with these different types of distributions. Let network a have an internal distribution described by a truncated power law, $p_{k_a}^a \propto k_a^{-\tau_a} \exp(-k/\kappa_a)$, and network b have an internal Poisson distribution. Coupling these networks via a distinct Poisson distribution is described by the second case shown in Table 4.1. Here, the composite $\mathcal{C}$ view requires fewer edges to achieve a giant component, so large-scale connectivity requires more edges if the network membership of the nodes is accounted for. The effects in shifting the percolation transition can be amplified if the networks are of distinct size, $n_a \neq n_b$. For more details on these percolation properties of interconnected networks, see Refs. [39, 40]. Also, see Ref. [38] for a discussion of how correlations in *multiplex* networks can alter percolation properties.

4.3 Application: Sandpile Cascades on Interconnected Networks

Equipped with a random graph model of interconnected networks and an understanding of its percolation properties, we now use this framework to analyze systemic risk by studying a dynamical process occurring on such interconnected networks. Here we seek a model that captures risk of widespread failure in critical infrastructures.

4.3.1 The Sandpile Model as a Stylization of Cascading Failure in Infrastructure

A common feature of many infrastructures is that their elements hold load of some kind, and they can only hold a certain amount of it. For example, transmission lines of power grids can carry only so much electricity before they trip and no longer carry electricity [18]; banks can withstand only so much debt without defaulting [30]; hospitals can hold only so many patients; airports can accommodate only so many passengers per day. When a transmission line, bank, hospital or airport partially or completely fails, then some or all of its load (electricity, debt, patients or travelers) may burden another part of that network or a completely different kind of network. For instance, when a transmission line fails, electricity quickly reroutes throughout the power grid (the same network), whereas when an airport closes due to a catastrophe like a volcano eruption [31] travelers may overwhelm railway and other transportation networks.

In addition to loads and thresholds, another commonality among certain risks of failure in infrastructure are heavy-tailed probability distributions of event size. In electric power systems, for instance, the amount of energy unserved during 18 years of North American blackouts resembles a power law over four orders of magnitude, and similarly broad distributions are found in other measures of blackout size [18]. In financial markets, stock prices and trading volume show power law behavior, in some cases with exponents common to multiple markets [22, 26]. In interbank credit networks, most shocks to banks result in small repercussions, but the 2008 financial crisis demonstrates that large crises continue to occur. Similarly broad distributions of event sizes also occur in natural systems such as earthquakes [64], landslides [32] and forest fires [45, 65]. Some evidence suggests that engineered systems like electric power grids [18] and and financial markets [22], not to mention natural catastrophes like earthquakes [64], landslides [32] and forest fires [45, 65], all show heavy-tailed event size distributions because they self-organize to a critical point.

An archetypal model that captures these two features—of units with capacity for load and of heavy-tailed event size distributions—is the Bak-Tang-Wiesenfeld (BTW) sandpile model [5, 6]. This model considers a network of elements that hold load (grains of sand) and that shed their load to their neighbors when their load exceeds their capacity. Interestingly, one overloaded unit can cause a cascade (or avalanche) of load to be shed, and these cascades occur in sizes and durations distributed according to power laws. This deliberately simplified model ignores detailed features of real systems, but its simplicity allows comprehensive study that can in turn generate hypotheses to test in more realistic, detailed models, which we will discuss in Sect. 4.3.4.

4.3.2 Defining the Sandpile Model on Networks

First studied on a two-dimensional lattice [5, 6], the BTW model has recently been studied on random graphs [11, 13, 17, 27, 28, 35–37, 41, 53], in part because many critical infrastructure like power, water, transportation and finance have network structure. There are different ways to implement the sandpile model on a network, but these implementations only differ in specifics. Here we study the following natural formulation [27, 28, 36, 37].

Each node holds grains of sand, which we interpret as units of load. Nodes can hold only a certain number of grains. When the number of grains equals or exceeds the node's threshold, then the node topples and moves sand to its neighbors. A natural choice for the threshold of a node is its degree, so that when a node topples it sends one grain to each of its neighbors. Other thresholds have been studied [27, 36], but these other rules for the threshold require nodes to shed sand to (for example) a random subset of their neighbors.

The BTW sandpile model consists of a sequence of cascades (avalanches), defined as follows. First, drop a grain of sand on a node chosen uniformly at random. If the node's number of grains is greater than or equal to its threshold (i.e., its degree), then that node is considered overwhelmed or unstable, and that node sheds (moves) all its load to its neighbors by sending one grain to each neighbor. These neighbors may in turn exceed their thresholds and have to topple, and subsequently their neighbors may topple, and so on. Once no node exceeds its threshold, we record the number of nodes that toppled (the cascade size), and the process begins again by dropping a grain on a random node.

In order to prevent the system from becoming inundated with sand, grains of sand must somehow be removed. Following [28], we choose the following rule for dissipation of sand: whenever a grain of sand is sent from one node to another node, with some small, fixed probability that grain is removed from the system.

The quantities of interest are measured in the dynamical equilibrium state that the system reaches after many cascades have occurred, because the system self-organizes to a critical point. Specifically, if the network begins without any sand, then sand slowly builds up in the system. After a large number of cascades (e.g., an order of magnitude more than the number of nodes), the system appears to reach a dynamical equilibrium at which the amount of sand does not change significantly relative to the system size. On one hand, large cascades tend to destroy lots of sand because of the rule for dissipating sand described above. On the other hand, when the amount of sand in the system is low, then cascades are typically smaller, so the amount of sand tends to increase. These effects balance so that the system seeks a critical point at which the probability distribution of cascade size (and of other measures like cascade duration) show power law behavior [28, 53]. These power law–distributed cascades can therefore serve as a useful, stylized model for risk of large cascading failures in infrastructures.
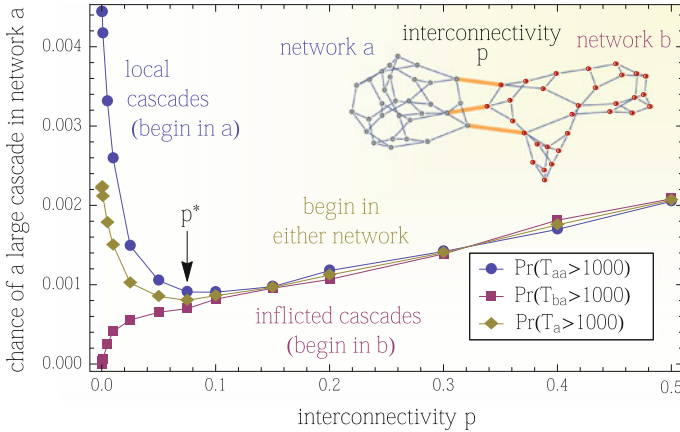
4.3.3 Results for the Sandpile Model on Interconnected Networks: Optimal Interconnectivity, the Yellowstone Effect, and Systemic Risk

In this subsection, we highlight three results from Ref. [13]. Next, in Sect. 4.3.4, we comment on current and future work to understand the sandpile model on isolated and interconnected networks, as well as on work to understand risk in interdependent infrastructures and other examples of optimal, intermediate amounts of connectivity.

We begin by studying one of the simplest interconnected networks, two random 3-regular graphs a and b with edges introduced between them. Specifically, each node in network a (b) has 3 neighbors in network a (b, respectively). Networks a and b have identical number of nodes. Next, a fraction p of nodes in a have one edge to a neighbor in the other network. (In the notation of the join degree distributions in Sect. 4.2.1, the degree distribution of network a is $p^a_{k_a,k_b} = \delta_{k_a,3} \left[p\delta_{k_b,1} + (1-p)\delta_{k_b,0} \right]$, and vice versa for network b.) This "interconnectivity" parameter $p \in [0, 1]$ measures the coupling between the two networks. The threshold of each node is its total degree.

One motivating example for this choice of interconnected networks are power grids. The degree of a typical node in the transmission grid (the part of a power grid that moves electric energy at high voltage) is approximately 3 [13], so we chose to study random 3-regular graphs. (Using 3-regular graphs rather than, say, Erdős-Rényi random graphs, simplifies the degree distribution to delta functions and hence simplifies branching process approximations of cascades [13].) Moreover, power grids have modular structure because they consist of "regions" or "control areas". Historically, each region was its own grid, and then these grids began connecting with one another, so that now one grid can span an entire continent. Each region of the grid is typically more densely connected within the region than with other regions. Furthermore, this modular structure is not static: grids continue to build new transmission lines between regions in order to, for example, accommodate wind power [34]. Increasing the interconnectivity p in our model vaguely captures the construction of new transmission lines between regions of a power grid.

Other infrastructures, from interbank credit networks [46] to transportation [48], exhibit modular structure at different scales. In some cases, these modules are becoming more interconnected over time, as lending, travel and trade become more global. Understanding how this increase in connectivity affects systemic risk is a problem that transcends disciplines. Though the sandpile model does not capture any one of these infrastructures accurately, it self-organizes to a critical point at which cascades occur in sizes described by power laws, and this behavior vaguely resembles large fluctuations in many engineered and natural systems. Thus, the sandpile model can be useful for generating hypotheses to test in more realistic models. Next we highlight three such hypotheses.

Fig. 4.4 The chance that a network a connected to another network b suffers a cascade larger than half its network [gold curve, $Pr(T_a > 1000)$] has a minimum at a critical amount of interconnectivity p^*. Networks that want to mitigate their largest cascades would prefer to build or demolish interconnections to operate at this critical point p^*. The *blue* (*red*) *curve* is the chance $Pr(T_{aa} > 1000)$ [$Pr(T_{ba} > 1000)$] that a cascade that begins in a (b) topples at least 1000 nodes in a. Increasing interconnectivity only exacerbates the cascades inflicted from b to a (*red*), but interestingly it initially suppresses the local cascades in a (For each p, we run a simulation on one realization of two random 3-regular graphs with 2000 nodes each; each node has a neighbor in the other network with probability p. The dissipation parameter is 0.01, the amount that makes the largest cascades slightly smaller than the size of the system. The *inset* depicts a small example with 30 nodes per network and $p = 0.1$.)

4.3.3.1 Optimal Interconnectivity

Suppose each network a, b is a region of a power grid and that each region is owned by a different utility. (To reiterate, the sandpile model misses crucial features of power grids, described below in Sect. 4.3.4.1, but we use the power grid example to facilitate interpretation of results.) If each network (think "each utility in the power grid") a, b wants to reduce the risk of cascading failure in its own network, then how many interconnections (edges between the networks) would they want?

Figure 4.4 shows the striking result that each network a, b would want to build some interconnections but not too many. Specifically, define a large cascade in a network as a cascade that topples at least half of the network. In Fig. 4.4, a has 2000 nodes, so a large cascade in a is one that causes at least 1000 toppling events in a. (The results are rather insensitive to changes in this cutoff for calling cascades large; see [13, Fig. 4]. Also, Sect. 4.3.3.2 explores the risk of small cascades.) The chance of a large cascade in a network is a measure of that network's risk. The gold curve of Fig. 4.4 shows that a network's risk decreases and then increases with the interconnectivity p, with the minimum occurring at an intermediate interconnectivity p^*. Thus, two initially isolated networks would want to build interconnections up to p^* in order to reduce their own risk of large cascades.
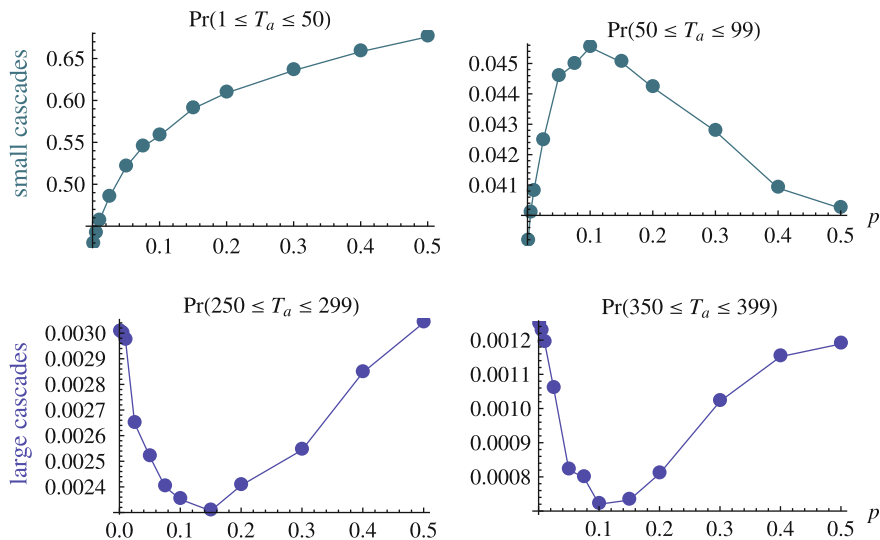
The 70 % drop in the risk of either network due to increasing interconnectivity p from 0.001 to $p^* = 0.075 \pm 0.01$ is significant. If these cascades were blackouts, then utility a (say) would experience 70 % fewer large blackouts. Why? By building $p^* N_a = 150 \pm 20$ edges (transmission lines) with its neighboring network b, the networks can collectively share their risk of large blackouts.

To further illustrate this "optimal interconnectivity" p^*, we distinguish cascades that begin in network a (the blue curve labeled "local cascades" in Fig. 4.4) from cascades that begin in network b (the red curve labeled "inflicted cascades"). As interconnectivity p increases, the chance of a large inflicted cascade increases monotonically, because building interconnections opens new avenues for large cascades to spread to the other network.

More interestingly, building some interconnections (but not too many) suppresses local cascades. That is, when interconnectivity p is small, the more edges a has with b, the lower the chance that a cascade begun in a topples a number of nodes in a greater than half the size of a. One reason for this suppression of local cascades is that nodes with an edge to the other network have larger threshold (because their degree is 4 rather than 3), so they topple less often when they receive sand. (However, the repercussions of toppling a degree-4 node are worse because they hold more sand.) Another reason that some interconnectivity suppresses local cascades is that more interconnections make the cascades less confined to one network and instead become more spread out among the two networks (see [13, Fig. S10]). This phenomenon of sharing risk resembles the tenet of diversification in investment portfolios in finance [2, 9].

Before proceeding, we note a similarity between optimal interconnectivity and equilibria in economics. Just as rational agents seek more of something as long as the marginal benefits exceed the marginal costs, a network would seek more interconnectivity as long as the marginal benefits exceed the marginal costs. In the sandpile model, building interconnections confers more benefits than costs initially, where benefits are reduction in risk of large cascades. In a competitive market, consumers and firms converge on the optimal price p^* at which the marginal benefit of the last unit consumed equals the marginal cost. Analogously, two networks seeking to mitigate their risk of large cascades converge on the optimal interconnectivity p^* at which the marginal benefits of the last edge built equal the marginal cost. More realistic models of connections within and between networks would also incorporate the costs of building and maintaining a new link, and this cost would presumably change the optimal number of links p^*.

Perhaps many interconnected networks are what Nassim Taleb calls "antifragile", meaning that they become more robust against large-scale catastrophes [66] if they have some variability [67] from input from external networks (e.g., interconnectivity $p^* > 0$).
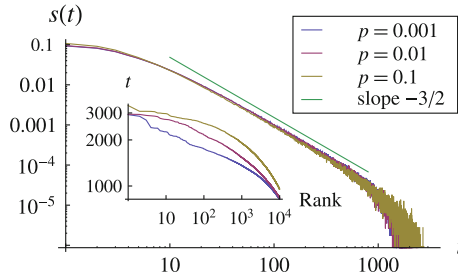
Fig. 4.5 a Networks mitigating the smallest cascades of size $T_a \in [1, 50]$ seek isolation $p = 0$, while **b** networks suppressing intermediate cascades $T_a \in [50, 100]$ seek isolation $p = 0$ or strong coupling $p = 1$, depending on the initial interconnectivity p in relation to the unstable critical point $p^* \approx 0.12 \pm 0.02$. But networks like power grids that mitigate large cascades **c, d** would seek interconnectivity at the stable equilibrium $p^* \approx 0.12 \pm 0.02$. The qualitative shape of the plots in the *bottom* figures and the location of p^* are robust to changes in the window $\ell \leq T_a \leq \ell + 50$ for all $200 \leq \ell \leq 800$ (Here we show results from simulations on two random 3-regular graphs with 1000 nodes each, which is half the network size as in Fig. 4.4, to show how p^* decreases with system size.)

4.3.3.2 Yellowstone Effect: Why Suppressing Small Cascades Can Increase the Risk of Large Ones

Rather than seeking to mitigate their risk of large cascades (and hence seeking interconnectivity p^*), what if the two networks a, b seek to mitigate their risk of small cascades? Figure 4.5 shows that the risk of small cascades increases monotonically with interconnectivity p. Thus, $p = 0$ minimizes the risk of small cascades.

However, by Fig. 4.4, $p = 0$ is a local maximum in the risk to each network. Thus, by seeking $p = 0$ to mitigate their own small cascades, networks a, b would increase their risks of large cascades. The same phenomenon is thought to occur in suppressing blackouts [18] and forest fires [45]. In fact, this phenomenon has been given the name the "Yellowstone effect" because suppressing small forest fires in Yellowstone National Park, WY, in the twentieth century densified forest vegetation and hence increased the risk of massive forest fires [45]. What Fig. 4.5 demonstrates is that interconnectivity is another mechanism that can cause the Yellowstone effect. This result suggests that we should look for similar phenomena in more realistic models.

Fig. 4.6 Increasing the interconnectivity p between two random 3-regular graphs extends the tail of the total avalanche size distribution $s(t)$, which does not distinguish whether toppled nodes are in network a or b. The *inset* shows a rank-size plot on log-log scales of the number of topplings t in the largest 10^4 avalanches (with 2×10^6 grains of sand dropped), showing that adding more edges between random 3-regular graphs enlarges the largest global cascades by an amount on the order of the additional number of interconnections. As expected theoretically [28], when a and b nodes are viewed as one network, $s(t) \sim t^{-3/2}$ for large t (*green line*)

4.3.3.3 Risk-Avoiding Networks Can Exacerbate System-Wide Risk

If two networks act in a greedy, rational manner to mitigate their own risk of large cascades, without regard to the risk of the other network, then by Sect. 4.3.3.1 each network would seek the optimal, intermediate amount of interconnectivity p^*. What is the effect of this self-interested behavior on the system as a whole?

Figure 4.6 shows that every increase in interconnectivity p tends to increase the size of the largest cascades in the whole system (where the size of the cascade no longer distinguishes types of nodes). The main plot of Fig. 4.6 is the cascade size distribution $s(t)$, which is the probability of observing a cascade with t many toppling events (in the equilibrium state of the system after many cascades have been run without collecting statistics), for interconnectivity $p = 10^{-3}, 10^{-2}, 10^{-1}$. (As expected [28], the avalanche size distribution shows power law behavior with exponent $-3/2$ over at least two orders of magnitude, and more detailed theoretical arguments confirm this conclusion [53].)

To illustrate the tail behavior of the cascade size distribution $s(t)$, the inset of Fig. 4.6 shows a rank-size plot of the largest cascades in the whole system. This plot shows that, as p increases, global cascades become larger by an amount on the order of the additional number of interconnections. Because each interconnection confers an addition to the threshold of nodes and hence to the capacity of the system to hold sand, the system holds more sand in the dynamic equilibrium state, so the largest cascades can involve more sand. Similar phenomena occur in transportation systems and electric power grids. Building a new bridge to Manhattan, for example, can paradoxically worsen traffic because people use the new capacity (the so-called Braess' Paradox [12]). Similarly, the benefits of upgrades to power grids are often fleeting, because operators seek to efficiently use their costly infrastructure [18].

4.3.4 Optimal Connectivity in Models Other Than the Sandpile Model

Because the sandpile model self-organizes to a critical point, it has relatively few parameters, so its behavior can be explored somewhat comprehensively. By contrast, detailed models of real systems can have so many parameters that they are difficult to comprehend, and many parameters are difficult to measure in the real world. Thus, the interesting behavior of the sandpile model—such as power law distributions of cascade size, optimal interconnectivity, the Yellowstone effect and system-wide risk—can serve as hypotheses to test and refine in more realistic models.

Next we review recent work in calculating optimal interdependence and optimal connectivity in power grids, bank networks and social systems.

4.3.4.1 Cascading Failures in Power Grids: Why Topological Models do not Suffice, and What Details are Needed

One of the promises of the burgeoning field of complex networks is to simplify the complicated behavior of real systems. Unfortunately, power grids are one example for which naïve, topological network approaches do not appear to suffice [14, 33].

Furthermore, most of these topological models, like the sandpile model, treat a cascading failure like an epidemic that spreads between adjacent nodes. By contrast, failures in the power grid spread non-locally: when a node in a power grid (such as a bus or a substation) fails or, more commonly, an edge (a transmission line) trips, the electric power re-routes almost instantly to all parallel paths, inversely proportionally to the impedances on lines. Thus, a failure can trigger other failures hundreds of kilometers away [24]. Models that lack this non-locality (or that have non-local failures but via completely different mechanisms) offer little insight into cascading failures in power grids [33].

What then is needed to offer insight for blackouts in power grids? A first step and essential ingredient are the linearized direct current (DC) equations, an approximation of the physics of electric power. These equations require data on the power injections at every node (which is positive for generators, negative for load buses) and the impedances and capacities of lines. Thus, the topological structure of a power grid is insufficient to run the DC power flow equations; these "thin" networks need to be augmented with data on the buses and transmission lines. (The MATLAB software MATPOWER [72] provides a useful starting point because it contains data on the Polish power grid. Another approach is to generate statistically correct power grid topologies [69].)

Equipped with a physics-based model of electric power flow in a grid, one must choose what failure mechanisms to model. Unfortunately, the number of failure mechanisms is large; examples include thermal overloads, relay failure, voltage collapse, dynamic instability and operator error [23]. The state-of-the-art is to model a subset of these failure mechanisms (see, e.g., [7, 23, 50]).

Such a detailed, physics-based, data-driven model of cascading failures might find rather rich pictures of optimal interconnectivity between regions of a power grid. The model space would likely be much richer than that of the sandpile model. But solving this problem once is not enough because modern power grids are changing. For instance, rapid deployment of smart grid technologies enable greater control and measurement of the grid. Renewable energy will stress the grid in new ways, as generation becomes more intermittent and increasingly located in sunny, windy locations, thereby changing the import of power between regions. These changes to the grid make studies of optimal grid structure all the more timely and important.

4.3.4.2 Optimal Interconnectivity in Bank Networks, Coupled Transportation Networks and Social Systems

The notion of optimal, intermediate amounts of connectivity is not new. For example, Battiston et al. [9] found that a network of banks is most resilient to cascading default if banks have an intermediate amount of risk diversification. What made this result novel was its contrast with the common wisdom in the financial literature that more diversification is always better [2]. In another model of bank default, if banks lend to an intermediate number of other banks, then the banks can be the most fragile and still not suffer cascading default [8].

Optimal coupling has also been found in a model of transportation on coupled spatial networks [48]. If a transportation administrator wishes to minimize both the average travel time and the amount of congestion, then a nontrivial, optimal "coupling" between, say, a rail network and a road network can emerge. Like in the sandpile model on interconnected graphs [13], two competing forces (efficiency and congestion in the transportation model) can balance at optimal amounts of coupling.

Optimal connectivity has also been found in strategic games played on networks. For example, a social network playing the minority game is most efficient at an intermediate amount of connectivity [4, 44]. Optimal interconnectivity between two networks has been found in the public goods game, where the interconnectivity $p = 1/2$ maximizes cooperation [68].

These results in financial and social models suggest that optimal connectivity (and interconnectivity) may be common in networks. If the dynamics occurring on some network have opposing tradeoffs, then these tradeoffs may balance at critical points. Whether the corresponding real systems—such as power grids, bank networks or social networks—can sense and respond to these tradeoffs and hence operate at the optimal points remains an open question. The answers would likely be essential to any effort to control such systems [43].

4.4 Conclusion and Future Work

Why might two networks build connections between them? This chapter demonstrates two reasons: to efficiently provide global connectivity and to reduce the risk that either system suffers a large cascading failure.

This work belongs to a movement to study mesoscopic, intermediate-scale structure in networks, and not just global structure (like degree distributions) and microscopic structure (like clustering coefficients). Two prominent examples of mesoscopic structure in networks are community (or modular) structure and core-periphery. There exist many tools for finding community structure in networks (see the reviews [25, 57]) and comparatively fewer tools for finding core-periphery structure [62]. But we are only just beginning to learn about the effect of this mesoscopic structure on the system's percolation properties (Sect. 4.2, Ref. [40]) and on dynamics occurring on the network (Sect. 4.3, Ref. [13]).

Another challenge is to study the converse: how the dynamics on the network affect its mesoscopic structure. In the sandpile model on interconnected networks [13], large cascades in one network may convince it to build more interconnections and hence to change the mesoscopic structure. Similarly, in power grids, large blackouts can provoke upgrades to the grid, which can include new transmission lines that change the structure of the grid. Large financial crises alter web of financial interactions among banks [46]. Widespread defection in a social network may alter its social ties. Thus, rare, catastrophic events [66] may be a sign of a network in the throes of its path toward optimal connectivity, if one exists.

References

1. *Review of Selected 1996 Electric System Disturbances in North America.*North American Electric Reliability Council, 2002.
2. F. Allen and D. Gale.Financial contagion.*Journal of Political Economy*, 108(1):1–33, 2000.
3. M. Amin.National infrastructure as complex interactive networks.In T. Samad and J. Weyrauch, editors, *Automation, control and complexity: an integrated approach*, pages 263–286. John Wiley & Sons,Inc., 2000.
4. M. Anghel, Z. Toroczkai, K. Bassler, and G. Korniss. Competition-Driven Network Dynamics: Emergence of a Scale-Free Leadership Structure and Collective Efficiency. *Physical Review Letters*, 92(5):058701, Feb. 2004.
5. P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of $1/f$ noise. *Physical Review Letters*, 59(4):381–384, 1987.
6. P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Physical Review A*, 38(1):364–374, 1988.
7. R. Baldick, B. Chowdhury, I. Dobson, Z. Dong, B. Gou, D. Hawkins, H. Huang, M. Joung, D. Kirschen, F. Li, J. Li, Z. Li, C.-C. Liu, L. Mili, S. Miller, R. Podmore, K. Schneider, K. Sun, D. Wang, Z. Wu, P. Zhang, W. Zhang, and X. Zhang. Initial review of methods for cascading failure analysis in electric power transmission systems ieee pes cams task force on understanding, prediction, mitigation and restoration of cascading failures. In *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pages 1–8, July 2008.

8. S. Battiston, D. D. Gatti, M. Gallegati, B. Greenwald, and J. E. Stiglitz. Default cascades: When does risk diversification increase stability? *Journal of Financial Stability*, 8(3):138–149, Sept. 2012.

9. S. Battiston, D. D. Gatti, M. Gallegati, B. Greenwald, and J. E. Stiglitz. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control*, 36(8):1121–1141, Aug. 2012.

10. B. Bollobás. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European Journal of Combinatorics*, 1:311, 1980.

11. E. Bonabeau. Sandpile dynamics on random graphs. *Journal of the Physical Society of Japan*, 64(1):327–328, 1995.

12. D. Braess, A. Nagurney, and T. Wakolbinger. On a Paradox of Traffic Planning. *Transportation Science*, 39(4):446–450, Nov. 2005.

13. C. D. Brummitt, R. M. D'Souza, and E. A. Leicht. Suppressing cascades of load in interdependent networks. *Proc. Natl. Acad. Sci. U.S.A.*, 109(12):E680–E689, Feb. 2012.

14. C. D. Brummitt, P. D. H. Hines, I. Dobson, C. Moore, and R. M. D'Souza. A transdisciplinary science for 21st-century electric power grids. Forthcoming, 2013.

15. S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 464:1025–1028, 2010.

16. M. Chediak and L. M. Cold snap causes gas shortages across u.s. southwest. *Bloomberg News*, (http://www.bloomberg.com/news/2011-02-04/cold-snap-causes-gas-shortages-across-u-s-southwest.html), Feb, 2011

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

18. I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman. Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization. *Chaos*, 17(026103), 2007.

19. I. Dobson, B. A. Carreras, and D. E. Newman. A branching process approximation to cascading load-dependent system failure. In *Thirty-seventh Hawaii International Conference on System Sciences*, 2004.

20. R. M. D'Souza, C. Borgs, J. T. Chayes, N. Berger, and R. D. Kleinberg. Emergence of tempered preferential attachment from optimization. *Proc. Natn. Acad. Sci. USA*, 104(15):6112–6117, 2007.

21. L. Dueñas-Osorio and S. M. Vemuru. Cascading failures in complex infrastructures. *Structural Safety*, 31(2):157–167, 2009.

22. B. Dupoyet, H. R. Fiebig, and D. P. Musgrove. Replicating financial market dynamics with a simple self-organized critical lattice model. *Physica A*, 390(18–19):3120–3135, Sept. 2011.

23. M. J. Eppstein and P. D. H. Hines. A "Random Chemistry" Algorithm for Identifying Collections of Multiple Contingencies That Initiate Cascading Failure. *Power Systems, IEEE Transactions on*, 27(3):1698–1705, 2012.

24. Federal Energy Regulatory Commission. *Arizona-Southern California Outages on September 8, 2011*, Apr. 2012.

25. S. Fortunato. Community detection in graphs. *Physics Reports*, 486:75–174, 2010.

26. X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423(6937):267–270, 2003.

27. K. Goh, D. Lee, B. Kahng, and D. Kim. Cascading toppling dynamics on scale-free networks. *Physica A: Statistical Mechanics and its Applications*, 346(1–2):93–103, 2005.

28. K.-I. Goh, D.-S. Lee, B. Kahng, and D. Kim. Sandpile on Scale-Free Networks. *Physical Review Letters*, 91(14):148701, Oct. 2003.

29. T. H. Grubesic and A. Murray. Vital nodes, interconnected infrastructures and the geographies of network survivability. *Annals of the Association of American Geographers*, 96(1):64–83, 2006.

30. A. G. Haldane and R. M. May. Systemic risk in banking ecosystems. *Nature*, 469:351–355, Jan 2011.

31. M. Hanlon. How we could all be victims of the volcano... and why we must hope for rain to get rid of the ash. *Daily Mail*, April 2010. http://www.dailymail.co.uk/news/article-1267111/

32. S. Hergarten. Landslides, sandpiles, and self-organized criticality. *Natural Hazards and Earth System Sciences*, 3:505–514, 2003.
33. P. Hines, E. Cotilla-Sanchez, and S. Blumsack. Do topological models provide good information about electricity infrastructure vulnerability? *Chaos*, 20(033122), Jan 2010.
34. C. Joyce. Building power lines creates a web of problems. *NPR*, April 2009. http://www.npr.org/templates/story/story.php?storyId=103537250
35. J. Lahtinen, J. Kertész, and K. Kaski. Sandpiles on Watts-Strogatz type small-worlds. *Physica A*, 349:535–547, 2005.
36. D. Lee, K. Goh, B. Kahng, and D. Kim. Sandpile avalanche dynamics on scale-free networks. *Physica A: Statistical Mechanics and its Applications*, 338(1–2):84–91, 2004.
37. K.-M. Lee, K.-I. Goh, and I. M. Kim. Sandpiles on multiplex networks. *Journal of the Korean Physical Society*, 60(4):641–647, Feb. 2012.
38. K.-M. Lee, J. Y. Kim, W.-K. Cho, K. Goh, and I. Kim. Correlated multiplexity and connectivity of multiplex random networks. *New Journal of Physics*, 14(3):033027, 2012.
39. E. A. Leicht and R. M. D'Souza. Random graph models of interconnected networks. *Forthcoming*.
40. E. A. Leicht and R. M. D'Souza. Percolation on interacting, networks. arXiv:0907.0894, July 2009.
41. S. Lise and M. Paczuski. Nonconservative earthquake model of self-organized criticality on a random graph. *Physical Review Letters*, 88:228301, 2002.
42. R. G. Little. Controlling cascading failure: Understanding the vulnerabilities of interconnected infrastructures. *Journal of Urban Technology*, 9(1):109–123, 2002.
43. Y.-Y. Liu, J.-J. Slotine, and A.-Á. Barabási. Controllability of complex networks. *Nature*, 473(7346):167–173, May 2011.
44. T. Lo, K. Chan, P. Hui, and N. Johnson. Theory of enhanced performance emerging in a sparsely connected competitive population. *Physical Review E*, 71(5):050101, May 2005.
45. B. D. Malamud, G. Morein, and D. L. Turcotte. Forest Fires: An Example of Self-Organized Critical Behavior. *Science*, 281(5384):1840–1842, sep 1998.
46. C. Minoiu and J. Reyes. A Network Analysis of Global Banking: 1978–2010. *Journal of Financial Stability*, 2013. in press.
47. M. Molloy and B. Reed. A critical point for random graphs with a given degree sequence. *Random Structures and Algorithms*, 6:161–180, 1995.
48. R. G. Morris and M. Barthelemy. Transport on Coupled Spatial Networks. *Physical Review Letters*, 109(12):128703, Sept. 2012.
49. P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela. Community structure in time-dependent, multiscale, and multiplex networks. *Science*, 328(5980):876–878, 2010.
50. D. P. Nedic, I. Dobson, D. S. Kirschen, B. A. Carreras, and V. E. Lynch. Criticality in a cascading failure blackout model. *International Journal of Electrical Power & Energy Systems*, 28(9):627–633, 2006.
51. D. E. Newman, B. Nkei, B. A. Carreras, I. Dobson, V. E. Lynch, and P. Gradney. Risk assessment in complex interacting infrastructure systems. In *Thirty-eight Hawaii International Conference on System Sciences*, 2005.
52. M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Phys. Rev. E*, 64(2):026118, 2001.
53. P.-A. Noël, C. D. Brummitt, and R. M. D'Souza. Controlling self-organizing dynamics using self-organizing models. Forthcoming, 2013.
54. W. of the ETH Risk Center. New Views on Extreme Events: Coupled Networks, Dragon Kings and Explosive Percolation. October 25–26th, 2012.
55. S. Panzieri and R. Setola. Failures propagation in critical interdependent infrastructures. *Int. J. Modelling, Identification and, Control*, 3(1):69–78, 2008.
56. P. Pederson, D. Dudenhoeffer, S. Hartley, and M. Permann. Critical infrastructure interdependency modeling: A survey of u.s. and international research. *Idaho National Laboratory, INL/EXT-06-11464*, 2006.

57. M. A. Porter, J.-P. Onnela, and P. J. Mucha. Communities in networks. *Notices of the American Mathematical Society*, 56(9):1082–1097, 2009.
58. E. Quill. When networks network. *Science News*, 182(6), 2012.
59. M. RE. North America most affected by increase in weather-related natural catastrophes. http://www.munichre.com/en/media_relations/press_releases/2012/2012_10_17_press_release.aspx, Oct. 17, 2012
60. S. Rinaldi, J. Peerenboom, and T. Kelly. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine*, December:11–25, 2001.
61. S. M. Rinaldi. Modeling and simulating critical infrastructures and their interdependencies. In *38th Hawaii International Conference on System Sciences*, pages 1–8, Big Island, Hawaii, 2004.
62. M. P. Rombach, M. A. Porter, J. H. Fowler, P. J. Mucha. Core-Periphery Structure in, Networks. arXiv:1202.2684, Feb. 2012.
63. V. Rosato, L. Issacharoff, F. Tiriticco, S. Meloni, S. D. Procellinis, and R. Setola. Modelling interdependent infrastructures using interacting dynamical models. *Int. J. Critical Infrastructures*, 4(1/2):63–79, 2008.
64. A. Saichev and D. Sornette. Anomalous power law distribution of total lifetimes of branching processes: Application to earthquake aftershock sequences. *Physical Review E*, 70(4):046123, Oct 2004.
65. P. Sinha-Ray and H. J. Jensen. Forest-fire models as a bridge between different paradigms in self-organized criticality. *Physical Review E*, 62(3):3216, sep 2000.
66. N. N. Taleb. *The Black Swan: The Impact of the Highly Improbable*. Random House Inc., New York, NY, 2007.
67. N. N. Taleb. *Antifragile: Things That Gain from Disorder*. Random House Inc., New York, NY, November 2012.
68. B. Wang, X. Chen, and L. Wang. Probabilistic interconnection between interdependent networks promotes cooperation in the public goods game. arXiv:1208.0468, Nov. 2012.
69. Z. Wang, A. Scaglione, and R. J. Thomas. Generating statistically correct random topologies for testing smart grid communication and control networks. *IEEE Transactions on Smart Grid*, 1:28–39, 2010.
70. S. Wasserman and K. Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge university press, 1994.
71. N. Wolchover. Treading softly in a connected world. *Simons Science News*, (http://www.wired.com/wiredscience/2013/03/math-prevent-network-failure/), 2013
72. R. Zimmerman, C. Murillo-Sánchez, and R. Thomas. Matpower: Steady-state operations, planning, and analysis tools for power systems research and education. *Power Systems, IEEE Transactions on*, 26(1):12–19, feb. 2011.