Evolving cascading failure resilience in complex networks

David Newth[1], Jeff Ash[2]

[1]*CSIRO Centre for Complex Systems Science*
GPO Box 284 Canberra ACT 2601 AUSTRALIA

Email: david.newth@csiro.au

[2]*Environmental and Information Science, Charles Sturt University,*
PO Box 789, Albury NSW 2640 AUSTRALIA

Email: {dnewth, jash}@csu.edu.au

Abstract

Our modern society has come to depend on large scale infrastructure networks to deliver resources to our homes and businesses in an efficient manner. Over the past 10 years there have been numerous examples where a local disturbance has led the global failure of the system. In this paper we use an evolutionary algorithm to evolve complex networks that are resilient to cascading failure. We then analyse these networks for topological regularities that explain the source of such resilience. The analysis reveals that clustering, modularity and long path lengths all play an important part in the design of robust large scale infrastructure.

1. Introduction

Many manmade networks such as the World Wide Web, the Internet and power transmission grids, as well as social networks, financial networks and other organisational and economic infrastructures exhibit similar statistical characteristics (Albert & Barabási 2002). Scale-free networks carry with them a well-recognised strength—tolerance of random failures. But they are particularly susceptible to failure of specific nodes that are highly connected (hubs) and if such removals occur, the networks disintegrate rapidly (Albert *et al.* 2000). When the dynamics of the network are also examined, another weakness is revealed—susceptibility to cascading failures (Motter & Lai 2002).

Cascading failures in power grids are well documented. On August 10 1996 in Oregon a combination of hot weather and abnormally high electricity demand caused power lines to sag into trees and trigger a cascade failure of power stations, distribution substations, and assorted other infrastructure which affected power supplies to 11 states (CNN 1996). On August 14 2003 a similar train of events starting in Ohio triggered the largest blackout in North American history (U.S. Canada Power 2004). Australia and New Zealand have not been left untouched. In Auckland the failure of four major distribution cables began on January 22 1998, and when the last of these collapsed almost a month later on February 20 the city was left totally without power. 17 days later this city had still only managed to regain 40 percent of its capacity (Davis 2000). Where a network is carrying a flow of some particular resource (electricity, gas, data packets, information, etc) nodes individually experience a load, and in normal circumstances this load doesn't exceed the capacity of that node. Also, in normal circumstances when flows on the network change, or nodes are added or removed, dynamic adjustment of flows on individual nodes occurs automatically, keeping all nodes loaded below capacity.

Cascade failures are initiated when a heavily loaded node is lost for some reason, and the load on that node (i.e. the flow passing through it) must be redistributed to other nodes in the network. This redistribution may cause other nodes to exceed their capacity causing them also to fail. Even if an overloaded node doesn't actually fail, protection mechanisms designed into the network may cause it to shut down anyway, in an attempt to prevent node failure. Hence the number of failed or stressed nodes increases, propagating throughout the network. In particularly serious cases the entire network is affected.

In this paper we use an evolutionary algorithm to evolve complex networks that are robust to cascading failures. We then apply network statistics to identify topological structures that promote resilience to cascading failure. The following section outlines the model of cascading failures used in this paper. This model serves as our fitness function. Section 3 outlines the evolutionary algorithm used to evolve complex networks. Section 4 details the network statistics used. Section 5 outlines the major results. Finally section 6 provides a discussion of the results, and further areas of study.

2. A model of cascading failures

Crucitti *et al.* (2003) (and others e.g. Motter & Lai 2002) proposed a simple model for representing cascading failures on complex networks. They showed that the breakdown of a single node that causes load redistribution to other nodes is sufficient to cause global system failure. This basic model is adopted here as a simulation of cascading failure. Essentially each node is characterized by a given capacity to handle traffic or resource flow. Initially the network is in a stationary state in which the load at each node is smaller than its capacity. The network is then perturbed, by the "breaking down" of a node. The breakdown changes the balance of flows within the network, leading to the redistribution of load to other nodes. If the extra load cannot be handled by those nodes it is then redistributed to other adjacent nodes. Eventually the string of failures will result in the reduction in the performance of the entire network. This cascading failure is somewhat analogous to the cascading failures observed in electrical power grids and the Internet. In the remainder of this section we will introduce the cascading failure model.

A generic communication/transport network is represented as a weighted undirected graph **G**, with N nodes and K edges. **G** describes an $N \times N$ interaction matrix. If there is an edge between nodes i and j then g_{ij} is assigned a weight in the range $(0, 1]$. Otherwise $g_{ij} = 0$. The value of g_{ij} can be thought of as a measure of cost, efficiency or a measure of resistance for moving between node i and j via g_{ij} . Initially at time t =0, all nonzero g_{ij} values are set to 1, meaning that all paths are functioning equally. The model then iteratively applies a rule for the time evolution of **G** that mimics the flow redistribution following the failure of a node. The most efficient path—ie the shortest-between nodes i and j—is represented by E_{ij} . The average efficiency of a network is then calculated as:

$$E(\mathbf{G}) = \frac{1}{(N(N-1))} \sum_{i \neq j \in \mathbf{G}} \varepsilon_{ij} \qquad (1)$$

and is used to measure the efficiency or performance of **G** at a given time step. The load $L_i(t)$on node i at time t is defined as the total number of shortest paths passing through node i at time t. Each node is characterized by a capacity C_i. The capacity is the maximum load that a node can handle before it becomes overloaded and its performance is diminished. The capacity is the number of shortest paths passing through node i at time t =0, multiplied by a tolerance parameter$\propto$:

$$C_i = \alpha \cdot L_i(0) \quad i = 0, 1, 2, \ldots, N \qquad (2)$$

where $\propto \leq 1$. $\propto$ can be thought of as a measure of the *stress* that the network is under at time $t = 0$.$\propto = 1$ means that network is operating at maximum capacity. $\propto$ » 1, depicts a network carrying a light load.

The breakdown of a network is simulated by reducing the capacity of a given node i by some margin. For this simulation this margin is set to 75%, (ie C_i 0.75). The change in performance of a given node affects the performance of other nodes. Consequently, the redistribution of loads overloads other nodes. At each time step t the following iterative rule is applied to simulate the load redistribution:

$$\varepsilon_{ij}(t+1) = \begin{cases} \varepsilon_{ij}(t) \cdot \frac{C_i}{L_i} & \text{if } L_i(t) > C_i \\ \varepsilon_{ij}(0) & \text{otherwise} \end{cases} \qquad (3)$$

where j are all those nodes in the local neighbourhood of node i.

Figure 1 illustrates the behaviour of the cascading failure model. Figure 1A shows the time evolution of efficiency for three values of (1.05, 1.1, and 1.35). Figure 1B compares the performance of an Erdös-Rényi random graph (Erdös & Rényi 1959), and a network from this study, over tolerance parameters ($\propto$ values) between 12. Note that when $\propto > 1.35$, there is no difference in the performance of these networks.

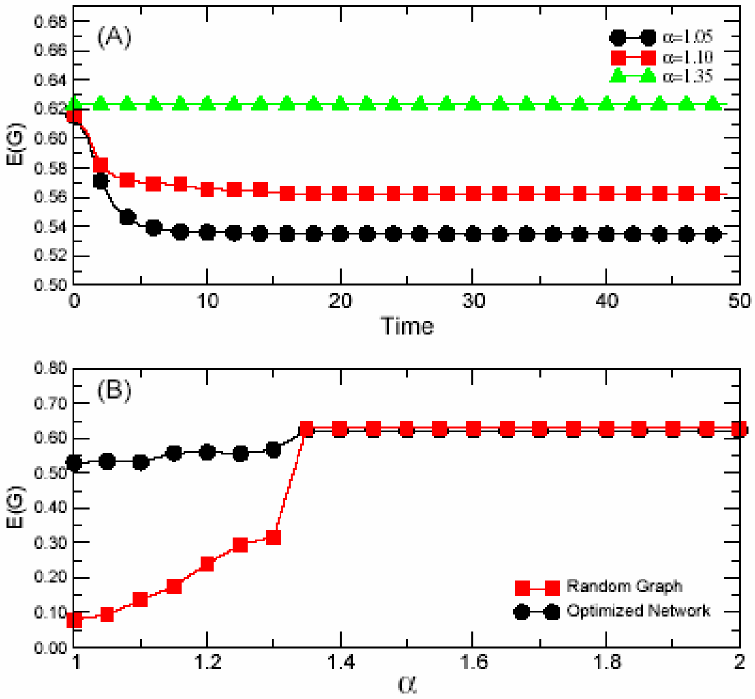

Figure 1. *Behaviour of the Cascading Failure Model. (A) Simulation of a cascading failure on a random network with values of 1.05, 1.10 and 1.35. (B) Comparison between an optimized network from this study and a random ensemble with same numbers of nodes and links, over values between 1-2. Over lower values of between 1-1.35 (highly stressed networks), the optimized network considerably outperforms the random ensembles. However above this threshold both networks display similar behaviour.*

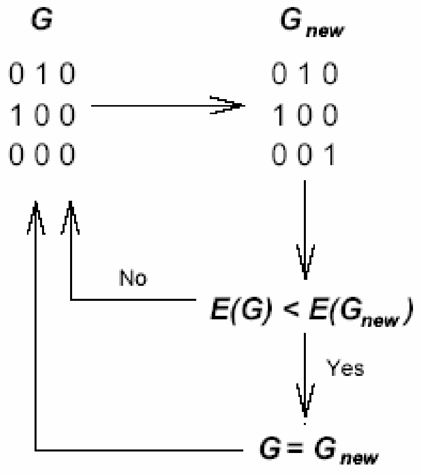

Figure 2. *Basic scheme of the evolutionary algorithm used in this study. Starting from a network/matrix* $\boldsymbol{G}$ *(here a* 3×3 *matrix), the algorithm mutates a given matrix entry. The objective function* $E(\boldsymbol{G})$ *is then evaluated (see §2), and the new network is accepted provided that a greater efficiency is achieved. Otherwise, we start again with the original network* $\boldsymbol{G}$*. At* $t = 0$*,* $\boldsymbol{G}$ *is initialized with a fixed density of ones.*

3. Evolving failure resilient networks

To search for failure resilient networks we have combined the recursive load redistribution algorithm and efficiency measures described in §2, with a variation of the Metropolis algorithm (Metropolis *et al.*1953) or a 1+1 evolutionary algorithm. The evolutionary algorithm is used to maximise $E(\mathbf{G})$. The value was fixed at 1.05, to simulate a highly stressed network. The metropolis type algorithm was selected over a population based approach, because of the computational complexity of evaluating $E(\mathbf{G})$. Figure 2 shows the basic configuration of evolutionary algorithm used here.

The networks used in this study had a total of 50 nodes. (While being quite small due to computational constraints, this does provide a sufficiently sized network to gain statistically significant attributes). Initially the network was seeded with a random network containing 100 edges. At each time step the child network was mutated in one of three ways: (1) a new link was added between two previously unconnected nodes; (2) a link connecting two nodes was deleted; or (3) a link connecting nodes i and j was reassigned to connect i and k. These mutation operations were selected without bias. The new network $\mathbf{G}_{new}$, was evaluated only if:

(1) $\mathbf{G}_{new}$ contained only one component (i.e. there were no disconnected nodes); and (2) the connectivity of the network was below a critical level $C_{critical}$. $C_{critical}$ was set at 150 edges to give a level of connectivity similar to that seen in many natural and physical systems (Albert & Barabási 2002). Also, constraining $C_{critical}$ prevented the algorithm from converging to a trivial completely connected network which was unrealistic. The algorithm was executed for 500 time steps. The topological properties of the final network, $\mathbf{G}_{f\,inal}$, were then analysed. The following section outlines the network properties that were analysed.

4. Complex networks

Recent studies have shown that many complex networks share common global statistical and structural features (see Albert & Barabási (2002) for an overview). These properties include scale free degree distribution, short path length and high clustering (the so called “small world” properties), modularity, and assortativeness. The remainder of this section briefly describes each of these properties.

4.1 Scale-free degree distribution

One of the common structural properties found in many man made and natural complex networks is a degree distribution with a power-law tail $P(k) \approx k^{-\gamma}$, where exponent γ is in the range between 2 and 3 (Albert & Barabási 2002). The degree of a node is the number of links possessed by that node. Networks exhibiting these power-law degree distributions are known as scale free networks. Several mechanisms have been proposed for the formation of such topological features. Barabási and Albert (1999) showed that a preferential attachment mechanism leads to a degree distribution with a power-law tail. Ferrer i Cancho and Solé (2001) showed that the minimisation of path length and the number of links contained within a network also leads to scale free structures. In more recent studies Brede and Finnigan (2004) have suggested that scale free networks possess a number of topological properties that promote stability when confronted with perturbations. Combined, these mechanisms suggest that scale-free networks may be an efficient and stable configuration for many complex networks.

4.2 Small-world networks

Small world properties can be detected through two statistics, the average shortest path length and the clustering coefficient. The average shortest path length (*PL*) is defined as:

$$PL = \frac{1}{N(N-1)} \sum_{i=1}^{N} \sum_{j=1}^{N} PL_{min}(i,j) \tag{4}$$

where PL_{min} is the minimum distance between nodes *i* and *j*.

Clustering is a common feature found in many networks. The extent of clustering within a network can be quantified by the clustering coefficient (*CC*) (Watts & Strogatz 1998). Given a node N_i, with k_i neighbours, E_i is defined to be number of links between the k_i neighbours. The clustering coefficient is the ratio between the number of links that exist between neighbours (E_i) and the potential number of links $k_i(k_i - 1)/2$ between the neighbours. The average clustering coefficient is:

$$CC = \frac{1}{N} \sum_{i=1}^{N} \frac{2E_i}{K_i(k_i - 1)} \tag{5}$$

Erdös-Rényi (1959) random graphs have a binomial degree distribution that can be approximated by a Poisson distribution. A network is said to have small world properties if, compared to an Erdös-Rényi random graph, the following conditions hold: $PL \approx PL_{rand}$ and $CC >> CC_{rand}$.

4.3 Common neighbours and modularity

Many complex systems exploit modularity as a way of coping with complexity. Modules in complex networks are subgroups of nodes that are highly interconnected within, and loosely connected outside the group. Figure 3 depicts a social system that is composed of five tightly coupled communities, each of which is loosely connected to the other communities (Newman 2004). Pimm (1980) proposed that the ratio between the number of neighbours nodes *i* and *j* have in common divided by their total number of neighbours (S_{ij}) as a simple measure of system modularity or compartmentalization. The average modularization *M* is the average of S_{ij} across all nodes:

$$M = \frac{1}{N(N-1)} \sum_{i=1}^{N} \sum_{j=1}^{N} S_{ij} \tag{6}$$

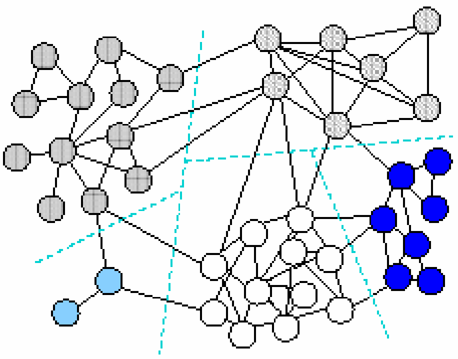

Figure 3. *A Complex Network With Five Distinct Modules. This figure shows a social network of comprising 38 nodes. The system is self-organised into a number of highly cohesive sub-groups (shading is used to depict members of the same group), that are highly coupled, with loose connections between groups (the loose couplings are those edges that cross the dashed lines).*

4.4 Assortativeness

A network displays assortative mixing if the nodes in the networks that have many connections tend to be connected to other nodes with many connections. A network is said to be disassortative if the highly connected nodes tend to be connected to nodes with few connections. The degree of assortativeness can be detected through the use of the Pearson's correlation coefficient. Such correlations can be defined as (Newman 2004):

$$\Gamma = \frac{c\sum_i j_i k_i - [c\sum_i \frac{1}{2}(j_i + k_i)]^2}{c\sum_i \frac{1}{2}(j_i^2 + k_i^2) - [c\sum_i \frac{1}{2}(j_i + k_i)]^2} \tag{7}$$

where j_i and k_i are the degrees of the vertices at the ends of the i^{th} edge. The constant c is defined as the reciprocal of m where m is the number of edges i.e. $c =$. A network displays assortative mixing when $\Gamma > 0$ and disassortative mixing when $\Gamma < 0$. Studies have found that social networks display assortative mixing, while systems with a power-law degree distribution are disassortatively mixed (Newman 2004).

5. Results

The evolutionary algorithm outlined in §3 was executed 100 times. In Section 5.1 we analyse the topological properties of the resultant network. In the following section we examine how the statistical features of the networks change as they are optimized. Section 5.2 examines the degree distribution of the optimized networks. Finally section 5.3 compares the statistical characteristics of the optimized networks with those of two random null models.

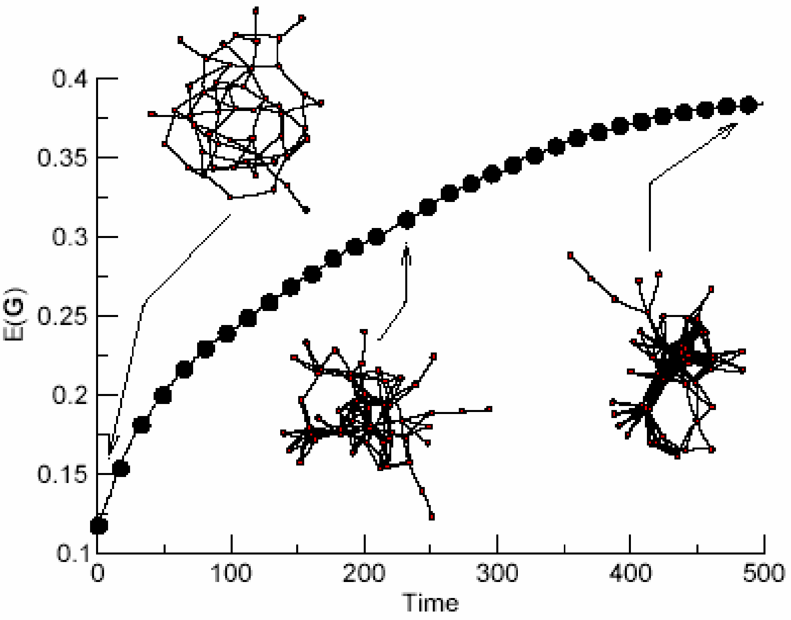

Figure 4. *Progression and Emergence of Structure. Initially at t=0 the network is seeded with a random graph. As the network becomes more robust structures such as hub, star-like substructures and highly connected sub-components begin to emerge (t=250). Finally at t=500 the network is highly optimized and displays some decisive structural characteristics, such as modules, and hubs.*

5.1 Network evolution

Figure 4, shows the improvement of the network efficiency ($E(\mathbf{G})$) over time, together with three snapshots of the network at time steps 0, 250, and 500. The figure illustrates that over time the network becomes more stable, and structure—emergence of hubs, clustering, bipartite sub structures, and modules—begins to occur. Figure 5, shows the evolution in five different network characteristics for a typical run.

5.2 Degree distribution

Extensive studies of the degree distribution of real-world networks has identified three main classes of networks: (1) scale-free networks, characterized by a vertex connectivity distribution that decays as a power law; (2) broad-scale networks, characterized by a connectivity distribution that has a power law regime followed by a sharp cutoff; and (3) single-scale networks, characterized by a connectivity distribution with a fast decaying tail (Anarak *et al*. 2000). Figure 6 shows the cumulative degree distribution for the networks evolved in this study. This suggests our evolved networks belong to the single-scale networks. This could be in part due to finite size effects, and to the fact that the random initial conditions (an Erdös-Rényi random graph), belong to the single scale networks.

5.3 Comparison

In previous studies (Barabási & Albert 1999; Barabási & Albert 2002; Milo *et al* 2002), it has been highlighted that many of the statistical properties of a network are derived directly from the degree distribution. In an attempt to determine how "unique" or "special" these evolved networks are we have compared their network statistics to those of two random null models. The first of these is an Erdös-Rényi (ER) random graph (1959) to determine those characteristics which can be accounted for by purely random interactions. The second model is the degree randomization model as described by Milo *et al.* (2002). This model assumes that the degree distribution is the source of the network properties. The model randomizes node connection (ie which node is connected to which other node), but preserves the individual node degree characteristics. Comparison between the evolved networks and the two null models, shows what is unique to the evolved networks, as well as what properties can be accounted for by random assemblage and the degree distribution.

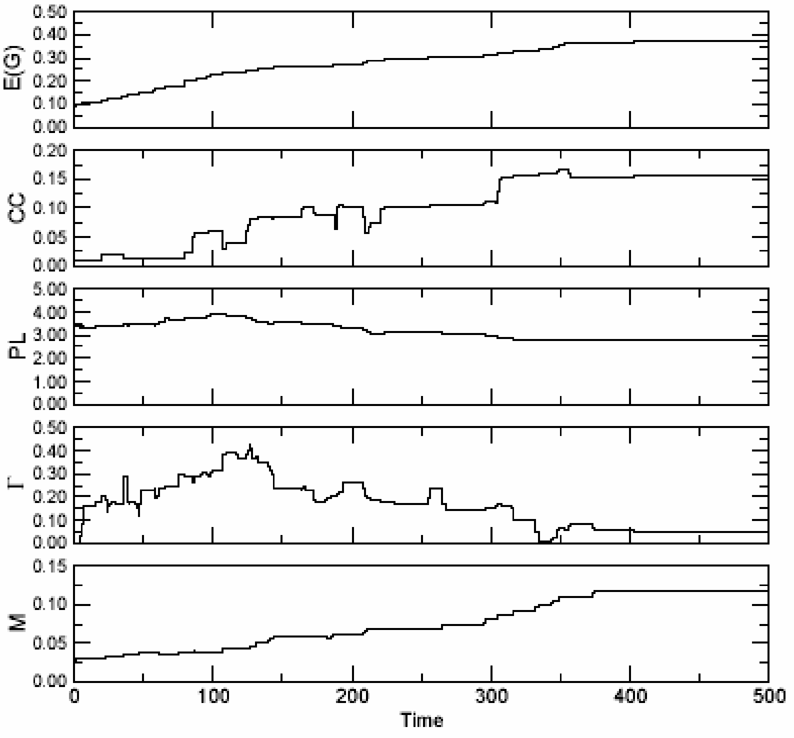

Figure 5. *Changes in Network Statistics during a Typical Run. As can be seen the efficiency* $E(G)$ *of the network when subjected to attack improves over time. The network also becomes more clustered and modular over time (*CC *and* M*), while the average shortest path length (*PL*) steadily decreases. Assortativeness (*Γ*) increases initially and then decreases.*

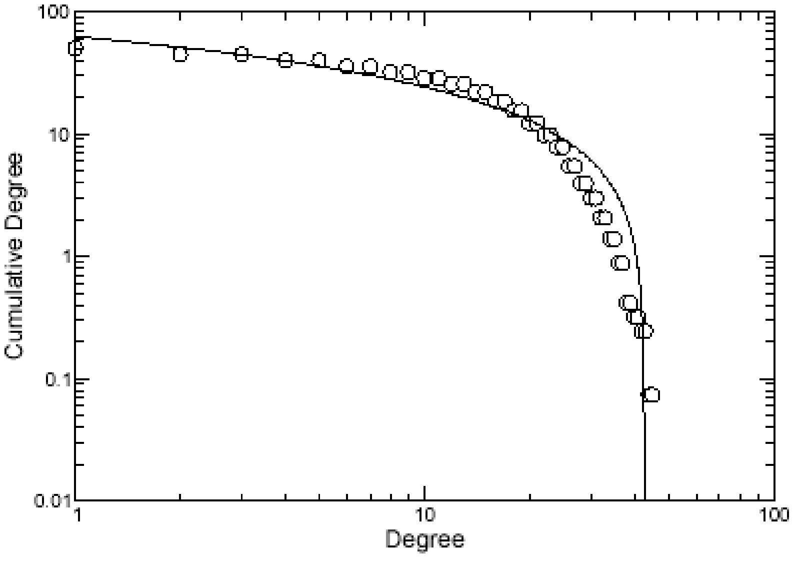

Figure 6. *Cumulative Degree Distribution for the Evolved Networks. The fast decaying tail is an indication of a single scale network.*

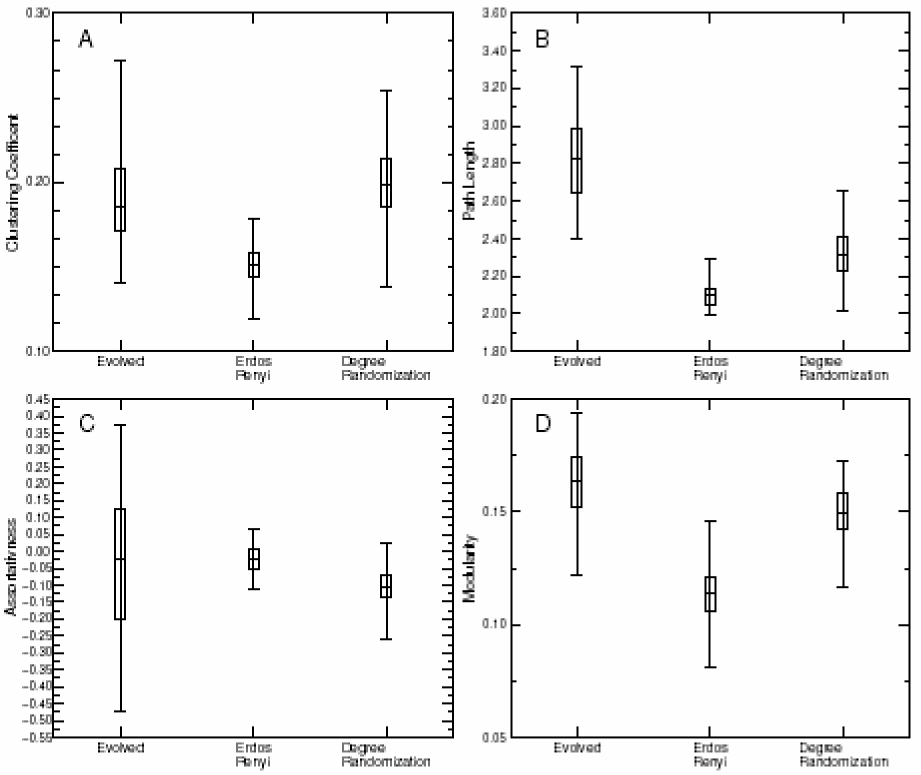

Figure 7. *Comparison between the Evolved Networks and the two Random Null Models. (A) Clustering Coefficient; (B) Path Length; (C) Assortativeness; and (D) Modularity.*

Figure 7 shows a comparison between the evolved networks and the two null models. 7A shows that the variation in clustering can be best accounted for by the degree distribution. The path length (7B) of the evolved networks is somewhat higher than that of the ER random graphs and the randomised degree networks. This characteristic can be considered to be unique to these networks. The wide variation in the assortativeness of the networks (7C) indicates that this characteristic is not an important one in the formation of cascade resilient networks. Finally the degree of modularity (7D) in the evolved networks is higher than that found in the random null models, although the degree distribution does account for some of the modularity.

6. Discussion

In this paper we have explored the concept of cascade failure resilient networks, and attempted to identify those characteristics that make these networks robust to cascading failure. Figure 4, illustrates that the optimized networks evolved by our model display a degree of structural regularity. This regularity emerges as the network organizes itself from its initial random configuration to an optimized, cascading failure-resistant configuration. Many of the emergent properties conform to the general principles of engineering design. Notions of modular design are common in many engineering projects. Self-contained components are often used as a means to isolate failures. We suggest that modularity in complex networks could also be serving a similar function. Local failures could be propagated locally and resolved, thus affecting only a small part of the network. Clustering also appears to be an important factor. High clustering provides a series of alternative pathways through which flows can pass, thus avoiding the failed component. The results here show that the degree distribution is an important factor in promoting modularity and clustering. That is, the way links are allocated to nodes drives the modular and clustered nature of the network.

One of the more interesting findings is that the optimized networks have longer path lengths than the two comparative random null models. The existence of long path length means that a disturbance needs to pass through a greater number of intermediate steps (and can potentially be resolved at each step), before the entire network is exposed to the failure. The longer the delay the greater the chance the network will resolve the perturbation. The placement of hubs and highly connected components is crucial in maintaining resilience (as the balance between modularity, clustering and path length needs to be maintained).

Hubs play an important role in many large scale infrastructure networks. While the variation in assortative mixing (see figure 7C), and finite size effects make it difficult to determine the role and configuration of hubs that makes networks stable, we postulate that collections of interconnected hubs (hub ensembles) allow perturbations to be distributed and reabsorbed quickly. Visual inspection of the optimized networks shows that this is a common feature; however a systematic test needs to be developed to gain a full understanding of the interconnected nature of the hubs.

Finally, the work presented here opens a number of additional lines of study, and three deserve mention. (1) The networks studied here are all homogeneous. How does the system organize itself, where certain key components are more likely to have a major failure? (2) If capacity of a given node is measured in terms of a cost function, what configuration generates the most robust topology, while minimising cost? (3) Many natural systems display a high degree of resilience to cascading failures. How do these networks compare with the evolved networks, and to large scale infrastructure networks? What are the sources (reasons) for the variations? All these questions require further experimentation but can be explored in the context of the framework proposed here.

References

Albert, R. and Barabási, A.-L. (2002) Statistical mechanics of complex systems. *Rev. Mod. Phys.*, **74**. 47–97.

Albert, R., Jeong, H. and Barabási, A.-L. (2000) Error and attack tolerance of complex networks. *Nature*, 406–378.

Amaral, L. A. N., Scala, A., Barthlmy, M, and Stanley H. E. (2000) Classes of smallworld networks. *Proc. Nat. Acad. Sci.*, **97**(21). 11149–11152.

Barabási, A.-L. and Albert, R. (1999) Emergence of scaling in random networks. *Science*, **286** 509–512.

Brede, M. and Finnigan, J. (2004) Growing Networks with Enhanced Resilience to Perturbation. arXiv:condmat/0405076v1.

CNN Interactive (1996) Sagging power lines, hot weather blamed for blackout. `http://www.cnn.com/US/9608/11/power.outage/`

Cohen, R., Erez, K., ben-Avraham, D. and Havlin, S. (2001) Breakdown of the Internet under intentional attack. *Phys. Rev. Let.*, **86** 3682.

Crucitti, P., Latora, V. and Marchiori, M. (2003) A model for cascading failures in complex networks. arXiv:condmat/0309141.

Davis, P. (2000) Auckland: City of darkness. *Earth Island J.*, **15** 4 `http://www.earthisland.org/eijournal/`

Erdös, P. and Rènyi, A. (1959) On random graphs. *Publicationes Mathematicae*, **6**, 290–297.

Ferrer i Cancho, R. and Solé, R.V. (2001) Optimization in Complex Networks. Santa Fe Working Paper 01-11-068.

Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D. and Alon, U. (2002) Network motifs: simple building blocks of complex networks. *Science*, **298** 824–827.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. and Teller, E. (1953) Equations of state calculations by fast computing machines. *J. Chem. Phys.*, **21** 1087– 1091.

Molloy, M. and Reed, B. (1995) A critical point for random graphs with a given degree sequence. *Random Structures and Algorithms*, **6** 161–180.

Motter, A. and Lai, Y.C. (2002) Cascade-based attacks on complex networks. *Phys. Rev. E.* **66** 065102.

Newman, M. E. J. (2004) Detecting community structure in networks. *Eur. Phys. J. B.*, **38** 321– 330.

Pimm, S. (1980) Food Web design, and the effect of species deletion. *Oikos*, **35** 139–149.

U.S.-Canada Power System Outage Task Force (2004) Final report on the august 14th blackout in the United States and Canada `https://reports.energy.gov/BlackoutFinalWeb.pdf`

Watts, D., and Strogatz, S. (1998). Collective dynamics of "small-world" networks. *Nature*, **393** 440–442.