

Geographic Distribution of Disruptions in Weighted Complex Networks: An Agent-Based Model of the U.S. Air Transportation Network

David C. Earnest

Old Dominion University, Norfolk, Virginia, USA
dearnest@odu.edu

Abstract

International networks, although highly efficient, may produce surprising threshold effects that shift costs to geographically distant locations. International utility, transportation, and information networks facilitate the efficient flow of information, energy, goods and people. These networks exhibit a scale-free network structure with a few large “hubs”. Yet their efficiency belies their lack of robustness. Because such networks transcend national boundaries, furthermore, disruptions to the network in one geographic region may have profound economic and national security costs for countries in another region. To illustrate how complex networks may transmit costs among countries, this paper builds an agent-based model (ABM) of the international air transportation system. The ABM employs a genetic algorithm to identify “small” disruptions that produce cascading network failures. The study makes two contributions. First, it demonstrates how some complex networks evolve into network structures that trade off robustness for efficiency. Second, it illustrates how researchers can combine agent-based modeling, evolutionary computation, and network analysis to simulate differing failure modes for global networks. This convergence of simulation methodologies characterizes the emerging field of computational social science.

Cascading Failures in Complex Networks

Global trade and commerce depends upon efficient transportation, information and utility networks. As these complex networks have evolved to meet the demands of consumers, however, they have assumed structures that are efficient but not very robust—that is, when they experience disruptions, they may take some time to resume their efficient operation. While such networks move people, information and goods, furthermore, they also transmit the costs of disruptions to geographically distant locations. The Northeast Blackout of 2003, triggered by erroneous power

readings in Indiana, caused power outages not only in the Midwest but also the Northeast as well as Ontario and Quebec in Canada, affecting 50 million people. In March 2008, the government of Haiti increased inspections of cargo through its ports in part to fight corruption, and in part to interdict the transshipment of narcotics to the United States. The inspections led not only to rotting and putrid food shipments on docks in Port-Au-Prince, but also to a backlog of containers in the Port of Miami as shippers had nowhere to store goods in Haiti (Katz and Kay 2008). In another recent telling example, in September 2011 a utility worker at a power substation in Yuma, AZ removed a faulty piece of monitoring equipment. The resulting cascade of power outages blacked out electricity throughout San Diego and much of Baja Mexico, leaving ATMs, traffic lights, 911 call centers and the San Diego airport powerless (Watson 2011). These examples illustrate how global networks may generate threshold effects that shift disruptions to geographically distant locations.

To illustrate the dynamics of such cascading failures in global networks, this study uses an agent-based model of the United States air transportation network. Previous research has found that air transportation networks exhibit the properties of a small-world scale-free network (Amaral et al. 2000, Guimera et al. 2005). Because such networks tend to have a few large “hubs,” or nodes with a large number of links, they tend to be robust to random failures but sensitive to targeted disruptions such as a terrorist attack. To understand the failure modes of the U.S. air transportation network, this study uses a genetic algorithm to act as a “smart terrorist”—the GA learns which attack strategies produce the largest disruption in air transportation for the least amount of effort. The article proceeds as follows. First, it reviews prior research on network failures, differentiating between studies that examine static metrics of network structure and those that measure dynamic flows across networks. It argues that agent-based modeling offers a better simulation method for

studying complex dynamic networks. The paper then reviews the structure of the U.S. air transportation network, using data from the U.S. Department of Transportation's Bureau of Transportation Statistics (BTS). This data confirms prior research that characterizes air transportation network as a small-world scale-free network. The study uses the BTS data to seed an agent-based model of the dynamic flow of passengers through the air transportation network. Given the large number of possible failure modes arising from the combinations of disrupted nodes, the study uses a genetic algorithm to explore the parameter space. It illustrates how a few apparently minor airports (such as Santiago, Chile) can nevertheless produce surprising backlogs in the flow of passengers in the United States. These findings illustrate how researchers can combine network analysis, evolutionary computation, and agent-based modeling to study the dynamics of global networks. The paper concludes with a discussion of directions for future research.

1. Methods of Analyzing Network Failure

Prior research suggests that several properties of a network may affect its vulnerability to disruption. A "network" simply consists of a number n of nodes (vertices), each of which has k connections (edges) to some of the other nodes. Networks may be undirected (that is, the edges between nodes represent reciprocal relations) or directed (edges represent one-way relationships). The number of edges a given node has is its "degree". One can characterize a network by its density D (the ratio of number of edges to the total number of possible edges, or $D = \sum k / n(n-1)$); the distribution of k ; and many other measures. Two important measures are a network's clustering coefficient and its average path length (sometimes referred to as its geodesic distance). A "path length" between two nodes is the number of edges along the shortest route between them. The average path length thus measures the mean distance between nodes in a network. The clustering coefficient of a network is the probability that two nodes are connected given that both are connected to a common third node. In this respect, the clustering coefficient measures the tendency of nodes to cluster together (Watts and Strogatz 1998).

Though networks differ widely along these and other measures, researchers have proposed three cardinal types of networks. A random network is one in which a random process creates edges among vertices. In random networks, the degree distribution approximates the Poisson distribution. The random attachment rule generally creates a network with very low clustering and relatively short path lengths. A "scale free" network is one for which the degree follows a power law distribution. This distribution arises because such networks grow through a process of

preferential attachment, whereby the probability of new edges incident to a node increases as the node's degree increases (Barabasi and Albert 1999). Scale-free networks are characterized by a few large "hubs," or vertices with a large number of incident edges, but most vertices have just a few edges (Barabasi and Albert 1999, Barabasi and Bonabeau 2003). For this reason, they tend to have lower clustering but higher average path lengths. By contrast, "small world" networks have higher clustering coefficients but the shorter average path lengths (Watts and Strogatz 1998; Watts 1999a, 1999b; Barrat et al. 2004).

The first studies of network failure tended to use static analysis. By measuring a network's degree distribution, density, and the size and number of its components (that is, a subset of the graph in which any vertex can reach another through some existing path), researchers can compare the structure of a network before and after the removal of a given vertex. In effect, this method allows researchers to compare how different networks "break apart." One such study compared random accidental failures of a node (as might occur in a power grid, for example) and attacks that targeted hub nodes. It found that in random networks, random failures tend to break the network into more, smaller components. By contrast, scale-free networks tend to be more robust to random failures but less so to attacks (Barabasi and Bonabeau 1999).

One problem with the static analysis of networks is that it tends to ignore several important features of networks, not the least of which is the volume of flows across edges. These flows may vary both across pairs of vertices as well as over time between any two given vertices. To capture these features, dynamic network analysis recently has examined how flows fluctuate through a network. Research has shown that the small-world network structure allows for efficient, parallel transmission, including mechanisms of disruptions to other nodes. For example, one study found that small world social networks are particularly efficient at transmitting diseases (Watts 1999a). Another method of analyzing dynamics is to create weighted network models, in which the edges between nodes are weighted by some measure of their traffic (Barrat et al. 2004; Wang and Chen 2008; Yang and Li 2011). Many networks exhibit such heterogeneity across vertices and edges; transportation networks, information systems, and power grids all have edges with varying flows. Using Monte Carlo simulation, the study of weighted networks can estimate the probability of failures. Several recent studies have used weighted network models to assess network robustness in general (Estrada 2006; Wang and Chen 2008); in transportation networks (Dall'Asta et al. 2006; Li and Mao 2006; Nagurney and Qiang 2007); and on the internet and in power grids (Yang et al. 2009).

It is useful to note that some complex global networks respond to geopolitical factors as well as technological and

economic ones. For example, one study found that the global air transportation network exhibits a surprising feature: the most connected cities in the network are not necessarily those that are the most “central” to the network (that is, that are on the shortest path length between other airports) (Guimera et al 2005). One consequence of this structure is that the smaller nodes may act as a bridge between different communities within the network, much as Anchorage is a critical bridge to the Alaskan air transportation subnetwork. These smaller bridge nodes may play a disproportionately large role in dynamic processes, whether the spread of infectious diseases or the transmissions of cascading failures.

Policy makers have expressed their concern about the robustness of global networks, not only because failures are costly but also because network vulnerabilities may lie elsewhere outside their jurisdiction. For example, title X of the Implementing Recommendations of the 9/11 Commission Act of 2007 (which became P.L. 110-53 with President Bush’s signature on August 3, 2007) calls for a national database on U.S. transportation assets whose loss “would have a negative or debilitating effect on the economic security, public health, or safety of the United States” (U.S. Public Law 110-53). Similarly, the U.S. National Infrastructure Protection Plan raises the prospect that transportation disruptions could generate cascading failures in the U.S. infrastructure (U.S. Department of Homeland Security 2006). Too often, however, policy makers have ignored the properties of networks that make these systems vulnerable. For example, a review of the Defense Critical Infrastructure Program Assessment Benchmarks for maritime transport focuses on physical security of specific ports, but does not address the networked characteristics that actually make the maritime shipping system vulnerable (US Department of Defense 2005).

One drawback of the analysis of weighted networks is that such models typically assume that the weights assigned to vertices or edges are constant. While this is a useful simplifying assumption, it is unrealistic: the demand for electricity is greater on hot days than moderate ones; the interstates on Thanksgiving tend to be more crowded than a Tuesday in October. For this reason, weighted network models may suffer from a lack of external validity. Another challenge for dynamic analysis is interaction effects among vertices; which combinations of failed nodes produce a greater likelihood of network fragmentation? Agent-based modeling (ABM) offers a useful alternative method for analyzing these networks. ABM is an object-oriented modeling methodology that simulates interactions among autonomous actors. It is particularly useful for simulating systems characterized by a large number of actors who interact repeatedly over time, and have cause-effect relationships that exhibit nonlinear

<i>Measure</i>	<i>Value</i>
Nodes	797
Edges	12,745
Density	0.0201
Average Total Degree	36.88
Median Total Degree	12
Power Law Exponent	1.434
Average Path Length	3.093
Clustering Coefficient	0.652

Table 1: Summary statistics of the U.S. air route network.

relationships due to feedback, exponential growth, and interaction among parameters. ABM also is a useful method for studying rare phenomena (Lustick, Miodownik and Eidelson 2004). Complex weighted networks exhibit all these properties.

To illustrate how complex networks may transmit costs among countries, this study uses an ABM of the international air transportation system. Using data from the U.S. Department of Transportation’s Bureau of Transportation Statistics, the ABM simulates the movement of people on international flights to and from the United States. The ABM also employs a genetic algorithm to identify “small” disruptions that produce cascading network failures. Genetic algorithms are a particularly efficient method of searching large parameter spaces such as those characterizing network dynamics. The algorithm identifies conditions under which disruptions elsewhere in the international network produce large economic losses to the United States.

2. Data

The Bureau of Transportation Statistics (BTS) in the U.S. Department of Transportation collects monthly data on air transportation within the United States and between the U.S. and cities with direct service to American airports (U.S. Department of Transportation 2011). The data includes not only the names of the origin and destination airports of each route but also the airlines servicing the route; the available seats; the number of passengers flown; the amount of freight and mail flown; the distance and air time; and the load factor (the ratio of passenger miles to available seat miles). It is important to note that the BTS data records charter as well as regularly scheduled flights, so there likely is some small variation in the network structure from month to month. Nevertheless, the data is a reasonably valid representation of the contemporary U.S. air transport network. I used the data for February 2011 to construct a weighted network using the origin and destination airports as nodes and the passengers flown as weights for the edges. Table 1 summarizes the properties

of the network. These data suggest the U.S. air transport network approximates the structure of a small-world scale-free network. A number of airports in the dataset were “isolates,” or had no scheduled or charter flights in February 2011. Several other city-pairs were “pendants”—that is, they were components unconnected to the rest of the network. After removing these, the network consists of 797 airports (nodes) and 12,745 “city pairs” or routes (edges). The average number of in- and out-routes (i.e., total degree) is 36.88; the median total degree of 12 suggests a skewed distribution that typifies a scale-free network. Figure 2 shows the log-log plot of the out degree distribution and is visually similar to the distribution of scale-free networks. Researchers disagree about how best to determine whether empirical data scales according to a power law (Clauset, Shalizzi and Newman 2009). I used the estimation technique proposed by Clauset, Shalizzi and Newman (2009) and found the U.S. air transport network’s scaling exponent γ is approximately 1.43. This estimate approximates the scaling exponent that other studies have found for the global air transportation network ($\gamma \approx 1$)

(Guimera and Amaral 2004); China’s network ($\gamma \approx 1.7$) (Li and Cai 2004); and that of India ($\gamma \approx 2.2$) (Bagler 2008). Each of the studies concludes these networks all exhibit scale-free small-world properties. I likewise conclude the BTS data suggests the U.S. air transportation network, along with its international city-pairs, also is a scale-free small-world network.

3. The Model and Genetic Algorithm

Using the BTS data, I created an agent-based model in which agents represent airports in the network. Each airport agent has a set of out links to the destination airport agents recorded in the BTS dataset. Each link has a weight w equal to the daily number of passengers who transited that route. Because the BTS data is aggregated by month, the weight equals $1 / 28$ of the reported passengers on a given route in February 2011. The simulation gives each airport an initial endowment of passengers equal to the sum of the weights of its in-links. At each step in simulated time (the time step is equivalent to one day in the real-world network) airports “send” a number of

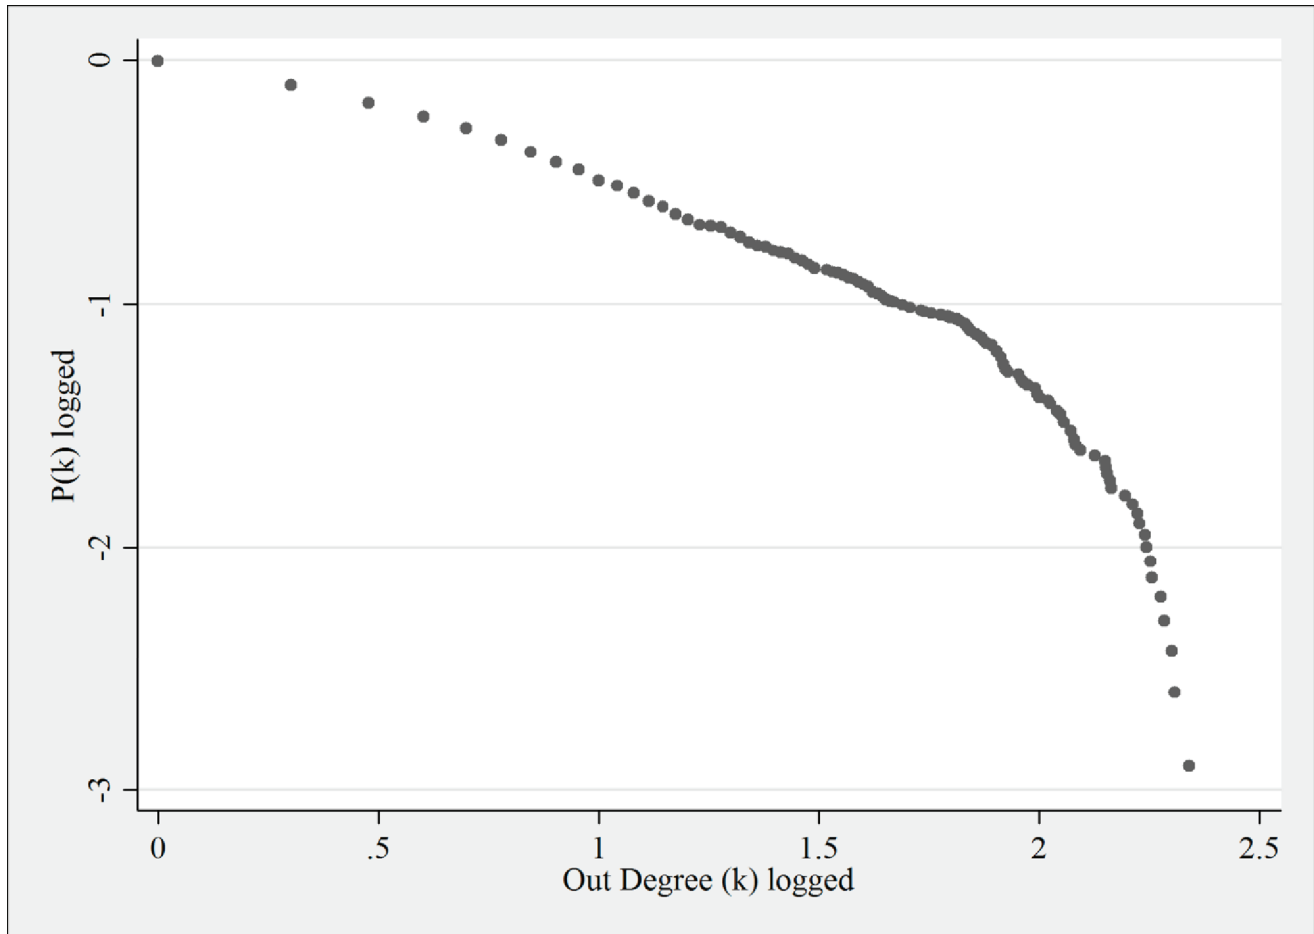


Figure 2: Log-log plot of the distribution of out degrees.

passengers to their network neighbors equal to the weight of each out link. Airport agents keep track of the stock of agents at each step in simulated time: the stock S is equal to the difference of the summed in-link weights and out-link weights: $S = \sum w_i - \sum w_o$. It is important to note that, in the simulation results presented here, the weights remain constant throughout the simulation. But there is no necessary reason for the ABM to have such a restriction. Indeed, one of the advantages of the model is that one can easily reprogram the simulation so that weights vary stochastically according to a schedule (such as fluctuations in passenger volume weekly or seasonally) or by drawing from a known statistical distribution.

Each airport agent has a throughput capacity equal to 1.25 times the sum of its in-link weights. This is equivalent to assuming each airport is operating at 80 percent of full capacity ($1 \div 0.8 = 1.25$). This capacity parameter gives airport agents some ability to send excess passengers to its network neighbors in the event of a backlog—that is, each out link from an airport has extra “seats” with which to move passengers if $\Delta S > 0$. The capacity parameter is comparable to the average load factor of routes in February 2011 (which was approximately 0.7). By assuming that airports are operating at less than capacity, the networks should exhibit some ability to recover from disruption as airport agents not directly affected by a disruption can use excess capacity to move a backlog of passengers. Conversely, as the capacity constraint grows toward 100 percent, one expects backlogs will build. Although it would be interesting to simulate the effect of variations in capacity on network backlogs, to focus on disruptions the simulation keeps the capacity constraint constant across airport agents and across experiments.

To examine how disruptions affect flows on the simulated network, the model uses a genetic algorithm (GA) (Holland 1992, Miller 1998). Borrowing insights from natural selection, a GA is an evolutionary computation technique that efficiently explores very large parameter spaces. For simulations characterized by both large numbers of parameter combinations and interaction effects, factorial designs can be quite time-consuming. In the analysis of large complex networks, factorial designs can be prohibitively slow, particularly when one wishes to account for interaction effects among nodes. For example, when Chicago O’Hare Airport is snowed in, there likely will be a considerable backlog in the network; but when both Chicago and Atlanta Hartsfield Airport are closed, the backlogs may be exponentially larger. To generalize the example, a factorial design that wished to identify an optimal combination of two airport nodes to remove would have to test $797 \times 796 = 634,412$ combinations. To study a three-node combination, the number of experiments grows to 5×10^8 .

GAs can search more efficiently. The algorithm acts as an “optimal terrorist” of sorts, exploring the system to discover which disruption strategies produce the largest backlog in the system. In each experiment, the GA optimizes against one of two fitness criteria: the average number of passengers backlogged at U.S. airports, and the total number of backlogged passengers in U.S. airports divided by the number of out links disabled by the GA’s disruption. It measures these criteria for 90 steps (simulated days) after a disruption. The former fitness criterion measures macro-level effects across the entire network. The latter criterion by contrast encourages the GA to be efficient by finding the greatest backlog for the smallest attack—essentially a minimax strategy. In a sense, by penalizing the GA for picking the largest airports, this latter criterion is equivalent to looking to trigger for a network avalanche much like the shutdown of a power generation plant in suburban Cleveland triggered cascading failures in the Northeast power grid in 2003.

The GA starts with an initial set of 50 random strategies—a “strategy” is simply a list of airports to remove from the network, e.g. [Tegucigalpa, Tokyo, Toronto]. Because I am interested in how disruptions in geographically distant locations may affect the United States, the GA’s strategies consist only of airports outside of the United States (there are 207 such airports in the BTS data). The model runs the simulation once for each strategy, disabling the airports as well as their in- and out links. It then measures that strategy’s performance using one of the two the fitness criteria. After testing all 50 initial strategies (a “generation”) the GA uses a selection procedure to populate 40 strategies for the next generation. In half of the experiments, the GA uses a simple tournament selection that compares the fitness of two randomly chosen strategies. In the other half, the GA uses a fitness proportionate selection rule, in which the probability of a strategy surviving to the next round is higher for better performing strategies. As the simulation evolves, the GA creates novel strategies in three ways. First, after every tournament selection the winning strategies cross over with a probability of .75. Second, at the end of every generation, after the algorithm selects its fittest strategies the GA mutates each allele on a fit strategy with a probability of .005. Finally, the selection tournament provides only 40 fit strategies for subsequent generations. The remaining 10 strategies are randomly generated ones, assuring that in each generation fit strategies compete against 20 percent new strategies. Figure 3 is a box plot of fitness by generation for one of the experiments; the hollow circles represent the median fitness for each generation. By the 22nd generation, the GA has found an optimal strategy that survives and becomes the median strategy by the end of the experiment. The figure also clearly shows how the median value




Figure 3: Box plot of strategy fitness by GA generation.

increases and the interquartile range grows with each passing generation. This is the value of a GA: it simultaneously improves strategies while exploring a range of alternative strategies.

An experiment consists of the GA testing 50 generations of 50 strategies each, for a total of 2,500 simulations per experiment. To test for interaction effects, the GA ran experiments in which it selected a single node; two nodes; and three nodes for removal. I conducted twelve total experiments: three types of strategy (one, two or three airports attacked) \times two fitness criteria (total backlog versus backlog / disabled node) \times two selection rules (tournament versus fitness proportionate). For each experiment I recorded both the measures of network performance and the final generation of 50 strategies.

4. Findings

Table 4 reports the frequency with which airport nodes appeared in the final generation of the GA. Recall that each generation included 50 strategies, and that the twelve

experiments varied disruption strategies from one to a combination of three airports. For this reason, the GA identifies an average of 100 optimal airport nodes in each experiment's final generation, for a total of 1,200 selected nodes. For k successes in n trials with a probability of success of p , the binomial mass function is $p^k \times (1-p)^{(n-k)}$. Because the GA selects from only the 207 non-U.S. airports, $p = 1 \div 207 = .0048$. Thus, the probability the GA of randomly selects an airport 13 times in 1,200 trials is about .006. Table 4(a) reports the airports the GA selected with a frequency that is significantly greater than random selection at $p < .01$; only 11 airports appear with a frequency greater than 13. Table 4(b) reports the selected airports for the total backlog criterion, while 4(c) lists the airports when the GA sought to optimize the total backlog per disabled link.

The results illustrate how the GA found airports that can disrupt flows in the air transport network even though they are not central to the network. The airport with the highest betweenness centrality (that is, the probability that the

<i>Node</i>	<i>N</i>	<i>Percent</i>
<i>(a) All Experiments</i>		
Tokyo (Narita)	193	16.08
Santiago, Chile	108	9.00
Toronto (Pearson)	74	6.17
Brisbane	48	4.00
Seoul (Inchon)	43	3.58
Birmingham, UK	38	3.17
Montego Bay	36	3.00
Abu Dhabi	34	2.83
Cancun	28	2.33
Mumbai	19	1.58
Santa Marta, Colombia	16	1.33
Total	1200	100.00
<i>(b) Criterion = Total Passenger Backlog</i>		
Tokyo (Narita)	143	23.83
Toronto (Pearson)	74	12.33
Seoul (Inchon)	39	6.50
Montego Bay	36	6.00
Cancun	26	4.33
Subtotal	600	100.00
<i>(c) Criterion = Backlog per Disabled Link</i>		
Santiago, Chile	108	18.00
Tokyo (Narita)	50	8.33
Brisbane	48	8.00
Birmingham, UK	37	6.17
Abu Dhabi	34	5.67
Subtotal	600	100.00

Table 4: Results of the GA experiments.

airport lies on the shortest path between all other vertices) is Toronto at .0045. This is an expected result: the BTS data reports only traffic to and from U.S. airports but not, for example, between Toronto and Vancouver. By construction, then, all non-U.S. airports in the simulation have low betweenness centrality. Nonetheless, the results also show how relatively “small” these airports are in the network. Toronto has the greatest number of connections to the U.S. network with 72 out links; Cancun has 43 and Montego Bay 26. Santiago, Mumbai, Brisbane, Birmingham, Abu Dhabi and Santa Marta all have three or fewer out links to the United States. In terms of flows,

Toronto sent an average of about 12,000 passengers to the United States per day in February 2011; Tokyo sent about 10,000; and Cancun about 8,500. These are obviously rather small portions of the daily network flow of about 1.9 million passengers.

A comparison of tables 4(b) and 4(c) illustrates how the GA found different strategies when optimizing different criteria. To create the greatest total backlog of passengers, the GA identified large foreign airports with both lots of connections to the United States and relatively large passenger flows. Tokyo’s Narita Airport and Inchon Airport in Seoul are important gateways from Asia to North America. Likewise, Toronto serves as a bridge between the Canadian and American air transportation networks. Surprisingly, the GA selected no European airports to disrupt. Equally surprising is its selection of Montego Bay and Cancun. Because the model uses BTS data from February 2011, the GA might be capturing winter travel to these vacation destinations. Yet their inclusion may also reveal some of the structural properties of the U.S. network. As Caribbean destinations, Cancun and Montego Bay form a cluster in the network because numerous large hubs in the U.S. are connected to both, including Atlanta, Dallas-Fort Worth, Newark, both New York airports, Chicago O’Hare and Miami. Indeed, the two airports share 23 U.S. destinations. This suggests that, although individually the Cancun and Montego Bay are relatively small, the interaction effect of a simultaneous disruption creates congestion in major hub airports in the United States.

Table 4(c) shows the GA results when it optimized a minimax criterion: the most disruption for the least number of disabled links. The results illustrate that, although large airports can create sizeable disruptions to passenger flows, such disruptions are relatively “costly” in the sense that they require disabling many links. When measured on a per-link basis, smaller airports may have a greater impact. Santiago, Chile is connected to only three U.S. airports; Brisbane and Abu Dhabi each are connected to only two. Yet because of the scale-free nature of the air transportation network, the hub structure allows relatively small nodes like Santiago to introduce perturbations that the hub then transmits through the network.

Although table 4 presents the frequency with which the GA selects specific airports to disrupt, it does not summarize the frequency with which the GA selects specific strategies. In eight of the twelve experiments, the GA combined the disruption of two or three airports outside the United States. An examination of these strategies should indicate whether the GA identified interaction effects among airports. Table 5 reports the most frequently selected strategies, and reveals a few surprises. Although Tokyo and Montego Bay may be geographically distant, their passenger flows intersect at a number of hub

<i>Strategy</i>	<i>N</i>	<i>Percent</i>
Montego Bay, Tokyo	36	9.00
Brisbane, Santiago	34	8.50
Seoul, Tokyo, Toronto	32	8.00
Abu Dhabi, Birmingham, Santiago	29	7.25
Aguascalientes, San Salvador	10	2.50
Santa Marta, Toronto	10	2.50
Total	400	100.00

Table 5: Most frequently selected strategy sets.

airports including Atlanta, Chicago O’Hare, Dallas-Fort Worth, and Los Angeles. These hubs also connect Seoul, Tokyo and Toronto. More surprising is the strategy to disrupt both Aguascalientes, Mexico and San Salvador. Though quite small, both Latin American airports feed traffic through Atlanta and Dallas-Fort Worth. Similarly, Toronto and Santa Marta, Colombia are connected through Miami and JFK Airport in New York. All of these examples suggest that combinations of disruptions can produce nonlinear effects by pushing the passenger backlog of a U.S. hub airport above the capacity threshold.

Finally, it is interesting to note that although the combination of Brisbane and Santiago is the second most frequently selected strategy, they share no link neighbors. To fly from Brisbane to Santiago, a passenger would have to transit either LAX or JFK first, and then Miami, Atlanta, or Dallas-Fort Worth. The frequency with which the GA selected this strategy suggests the possibility of second-order interaction effects. By simultaneously disrupting Santiago and Brisbane, the GA may induce backlogs that build first in one U.S. hub airport and then in another. In this respect, hub airports can act as multipliers for disruptions, magnifying the cascades of backlogged passengers. Anyone who has faced a “weather” delayed flight on a sunny day is familiar with these second-order effects.

5. Future Research

Although these findings are interesting, the simplifying assumptions of the simulation limit their generality. Foremost is the assumption that the U.S. air transportation system is a discrete network. Of course, it is merely a subnetwork of the global air transportation system. As the 2010 eruption of the Eyjafjallokull volcano in Iceland demonstrated, delays in the European subnetwork can reverberate in the U.S. With data on both the structure of and traffic across the global air transportation network, the GA might identify other, more effective modes of disruption. Similarly, the simulation would benefit from finer-grained measures of the network’s dynamics. The

simulation presented here used daily passenger flows to affix constant weights to links in the network. Likewise, it assumes a constant capacity constraint across airports and across time. Although the BTS aggregates data by month, it may be possible to measure the variation in passenger flows among airports in the system. Such data would allow the model to simulate daily and seasonal variations in passenger traffic, and by extension the variation in capacity constraints at airports. With such a refinement, the GA could search not only for optimal disruptions but also for an optimal time at and sequence in which to disrupt the airports. It is likely that the sequence and timing of disruptions is just as important as the nodes the GA disrupts.

What are the financial costs of the disruptions identified by the GA? The results above do not quantify the backlog as a percentage of total throughput in the system, nor do they estimate the financial costs of such delays. It may be, for example, that although the GA has identified simultaneous disruptions of Brisbane and Santiago as an optimal disruption strategy, this may create backlogs of only a few hundred passengers per day. A more realistic simulation would measure the financial costs of disruptions. After all, airlines and regulators ultimately are more concerned about financial losses than the number of individuals who are inconvenienced. The costs may be considerable, furthermore. The Air Transport Association estimates that in 2009, a one-minute delay of a flight produces about \$61 in direct costs to airlines plus another \$0.62 in opportunity costs to passengers (Air Transportation Association 2011). To quantify this in terms of the simulation results presented above, one experiment in which the GA disabled Seoul, Tokyo and Toronto produced about 39,400 passenger delay days (i.e. one passenger delayed one day) or a daily average of about 438 passengers. The costs to passengers alone would be about \$391,000 per day. Using data like this, the GA could select among the most costly strategies rather than merely those that affect flows the most.

Finally, the simulation would benefit from “smart” airport agents. In the current implementation of the simulation, an airport agent simply moves its passenger backlog to all its network neighbors—in effect, it assumes passengers are homogenous when, in the real world, they differ in their destinations. Obviously, this implementation is unrealistic: Chicago O’Hare cannot reroute a Des Moines passenger through South Bend because that passenger probably will end up back in Chicago. One useful extension of the model would be to endow airports with evolutionary learning as well, so that they can dynamically evolve strategies for dispensing with passenger backlogs. In effect, airport agents would co-evolve strategies with the disruption strategies created by the optimal terrorist GA.

6. Conclusions

Weighted complex networks behave in surprising ways. When such networks span the borders of nation-states—and many utility, information and transportation networks do—they may produce unintended costs that governments cannot control. While the static analysis of the structural properties of such networks can reveal subgraphs, bridging nodes, and other critical features, it tends to overlook dynamic flows through the network. To understand cascading failures, researchers need to conduct dynamic analysis of flows. Because many global complex weighted networks have evolved in response to market demands, furthermore, they have developed scale-free properties that are highly efficient for moving information, people, and goods, but that may not be very robust in the face of disruptions. For this reason, researchers and policy makers alike need methods to analyze how complex weighted networks respond to disruptions. Using the U.S. air transportation system as an example, this study illustrates how researchers can combine agent-based modeling, evolutionary computation, and network analysis to simulate differing failure modes for global networks. By focusing on disruptions at non-U.S. airports, the study demonstrates how disruptions may interrupt flows at points in the network that are geographically distant. This is not only costly to individual, firms, and governments, but it also demonstrates that individual governments cannot manage network effects on their own. The United States air transportation network relies upon efficient networks in Europe and Asia, just as those regions depend upon safe and efficient transportation in the United States. In the absence of international coordination in the management and security of complex networks, nasty surprises will inevitably occur.

References

Air Transport Association. Annual and Per-Minute Cost of Delays to U.S. Airlines. Online resource available at <http://www.airlines.org/Economics/DataAnalysis/Pages/CostofDelays.aspx>. Last accessed 9 September 2011.

Amaral, L.A.N. et al. 2000. Classes of Small-World Networks. *Proceedings of the National Academy of Sciences* 97, 21: 11149-11152.

Bagler, G. 2008. Analysis of the airport network of India as a complex weighted network. *Physica A* 387, 12: 2972-2980.

Barabasi, A.L. and Albert, R. 1999. Emergence of Scaling in Random Networks. *Science* 286, 5439: 509-512

Barabasi, A.L. and Bonabeau, E. 2003. Scale-Free Networks. *Scientific American* 288, 5: 50-59.

Barrat, A. et al. 2004. The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences* 101, 11: 3747-3752

Clauset, A., Shalizzi, C. R. and Newman, M. E. J. 2009. Power-Law Distributions in Empirical Data. *SIAM Review* 51, 4: 661-703.

Dall'Asta, L. et al. 2006. Vulnerability of weighted networks. *Journal of Statistical Mechanics* April: 1-12.

Estrada, E. 2006. Network robustness to targeted attacks. The interplay of expansibility and degree distribution. *The European Physical Journal B* 52, 4: 563-74.

Guimera, R. and Amaral, L.A.N. 2004. Modeling the world-wide airport network. *European Physical Journal B* 38, 2: 381-385.

Guimera, R. et al. 2005. The worldwide air transportation network: Anomalous centrality, community structure, and cities' global role. *Proceedings of the National Academy of Sciences* 102, 22: 7794-7799.

Holland, J. H. 1992. Genetic Algorithms. *Scientific American* 267: 66-73 and 114-116.

Katz, J. M. and Kay, J. 2008. Red tape cutting off food. *The Virginian-Pilot* (March 7): A3.

Li, K., Gao, Z., and Mao, Z. 2006. A Weighted Network Model for Railway Traffic. *International Journal of Modern Physics C* 17, 9: 1339-1347.

Li, W. and Cai, X. 2004. Statistical analysis of airport network of China. *Physical Review E* 69, 4.

Lustick, I., Miodownik, D. and Eidelson, R. J. 2004. Secessionism in Multicultural States: Does Sharing Power Prevent or Encourage It? *American Political Science Review* 98, 2: 209-229.

Miller, J. H. 1998. Active Nonlinear Tests (ANTs) of Complex Simulation Models. *Management Science* 44, 6: 820-830

Nagurney, A. and Qiang, Q. 2007. A Transportation Network Efficiency Measure that Captures Flows, Behavior, and Costs with Applications to Network Component Importance Identification and Vulnerability. In Proceedings of the POMS 18th Annual Conference. Miami, FL: Production and Operations Management Society.

U.S. Department of Defense. Office of the Assistant Secretary of Defense for Homeland Defense. 2005. Defense Critical Infrastructure Program. Directive 3020.40. Washington, DC: U.S. GPO.

U.S. Department of Homeland Security. 2006. National Infrastructure Protection Plan. Washington, DC: U.S. GPO.

U.S. Department of Transportation. Bureau of Transportation Statistics. Air Carriers: T-100 International Market. Online resource available at <http://www.transtats.bts.gov>. Last accessed 9 September 2011

U.S. Public Law 110-53. 110th Cong., 1st sess., 2007. Implementing Recommendations of the 9/11 Commission Act of 2007.

Wang, W. and Chen, G. 2008. Universal robustness characteristics of weighted networks against cascading failures. *Physical Review E* 77, 2.

Watson, J. 2011. Up to 5 Million Lose Power in Calif., Ariz., and Mexico. *The Virginian-Pilot* (September 9): 10.

Watts, D. J. 1999a. Networks, Dynamics, and the Small World Phenomenon. *American Journal of Sociology* 105, 2: 493-527

Watts, D. J. 1999b. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton, NJ: Princeton University Press.

- Watts, D. J. and Strogatz, S. H. 1998. Collective Dynamics of 'Small World' Networks. *Nature* 393: 409-410.
- Yang, R. et al. 2009. Optimal weighting scheme for suppressing cascades and traffic congestion in complex networks. *Physical Review E* 79, 2.
- Yang, Y. and Li, W. 2011. Comparative Analysis on Weighted Network Structure of Air Passenger Flow of China and US. *Journal of Transportation Systems Engineering and Information Technology* 11, 3: 156-162.