13th ICCRTS: C2 for Complex Endeavors

"GETS Survivability Analysis"

Topic(s):
Networks and Networking
Modeling and Simulation
Civil-Military Endeavors

Name of Authors:
Denise M. Masi, Ph.D., Noblis
Martin J. Fischer, Ph.D., Noblis
David A. Garbin, Noblis

Point of Contact:
Denise M. Masi
Noblis
3150 Fairview Park Drive South
Falls Church, VA 22042
703-610-1582
dmasi@noblis.org

GETS Survivability Analysis

Abstract

The ability to communicate during emergencies is essential for government personnel. The mission of the National Communications System (NCS) includes planning for and provisioning National Security/Emergency Preparedness (NS/EP) communications for the federal government under all circumstances, including crisis or emergency, attack, recovery, and reconstitution. In support of its mission, NCS runs several emergency telecommunications priority service programs for federal government users, including the Government Emergency Telecommunications Service (GETS) and is currently investigating the need to evolve the GETS program toward IP capability. This study is an analysis of the survivability objectives for Next Generation Network (NGN) GETS Voice Services. We used IP-SURVIV, a Noblis Survivability Analysis tool, to analyze throughput and connectivity of a network topology similar to major ISP backbone as network elements are disabled. Three failure scenarios were analyzed: random electronics failures, directed / terrorist attacks, and natural disasters. Throughput and connectivity were examined under each scenario, and worst case results will be incorporated into survivability objectives for the network service providers.

Keywords: modeling and simulation, survivability, network resilience

Introduction

The ability to communicate during emergencies is essential for government personnel. The mission of the National Communications System (NCS) includes planning for and provisioning National Security/Emergency Preparedness (NS/EP) communications for the federal government under all circumstances, including crisis or emergency, attack, recovery, and reconstitution. The NCS offers a wide range of communications services that support qualifying government personnel in performing their NS/EP missions, including the Government Emergency Telecommunications Service (GETS). GETS provides emergency access and priority processing in the local and long distance segments of the public switched wireline network, for use in an emergency or crisis situation.

Industry is moving from circuit switched to Internet Protocol (IP) technology for all telecommunications applications, including voice, and NCS is investigating the need to evolve toward IP capability to ensure continuity of priority traffic during emergencies. Modeling telecommunications for current and future programs providing emergency service is essential. We are conducting various modeling and simulation activities to support this effort.

The NCS is developing an NS/EP IMS Core Industry Requirements (IR), Voice Service Document, which is intended to be used to acquire service to ensure continuity of GETS and its wireless counterpart, the Wireless Priority Service program, as the infrastructure

migrates to packet-switched Next Generation Network (NGN). The document describes Quality of Service (QoS) and performance requirements for NGN GETS Voice service, in the areas of call establishment, quality of voice, call sustaining, reliability, availability, survivability. This paper describes the analysis that has been done in the area of survivability requirements for the NGN GETS Voice Industry Requirements document. The survivability requirements must be defendable, as shown by modeling results. Although our analysis is focused on a very specific application, it is helpful to first define what is meant by survivability for telecommunication networks, and review the literature in the area of the mechanics of complex networks.

The Alliance for Telecommunications Industry Solutions (ATIS) Telecom Glossary 2000 (ATIS, 2000) defines survivability as "a property of a system, subsystem, equipment, process, or procedure that provides a defined degree of assurance that the named entity will continue to function during and after a natural or man-made disturbance; *e.g.*, nuclear burst. *Note:* For a given application, survivability must be qualified by specifying the range of conditions over which the entity will survive, the minimum acceptable level or post-disturbance functionality, and the maximum acceptable outage duration."

Much research and analysis has been done in the area of network mechanics and their survivability. Two major approaches have been used in evaluating telecommunication network survivability: statistical approaches based on reliability calculations, and deterministic approaches based on graph theoretic calculations (Shake et al. 1999).

Albert and Barabási have done much research in both the theory and applications of modeling complex networks. Albert and Barabási (2002) is an extensive survey of modeling paradigms based on graph theoretic calculations for many types of complex networks, including the Internet, the World Wide Web, biological cells, networks of chemicals linked by chemical reactions, and others. They review random graphs whose node degree follows a Poisson distribution, small-world models which posit that in even large networks there is a short distance between any two nodes, and "scale-free" networks whose node degree follows a power-law tail. A power-law distribution is a distribution (such as the Pareto) of the form $P(k) \sim k^{-\gamma}$; in this case k is the node degree. This distribution for network degree implies that there is a non-trivial probability of having a very highly connected node. The models for scale-free networks focus on modeling the evolution of the network to result in the current topology. Results on network robustness due to random network errors and intentional attacks that cause network nodes or links to be disabled are also described. In the intentional attacks, the nodes which are most highly connected are removed, which they assume would be most damaging to a network. Their measures of network robustness or survivability under errors or attacks are the size of the largest connected cluster and the average path length of this cluster. These measures are examined for both random and scale-free networks having 10,000 nodes; the scale-free network they analyzed was generated using their network topology evolution model. One of the findings was that the network breaks down faster when nodes with the highest number of edges are targeted in an intentional attack versus when nodes are selected randomly.

Shake et al. 1999 focus on wide-area fiber-based telecommunication networks, and give as motivation ATIS statistics on the ease of some physical attacks, such as severing the fiber infrastructure. Their preferred measure of network vulnerability is the minimum number of nodes or links whose removal disconnects part of the network. They also propose assigning a vulnerability metric to each network link which could be based on a variety of factors. Vulnerability between pairs of nodes can then be computed based on the topologies and link vulnerabilities between the nodes. A network average vulnerability can then be computed. Their path vulnerability metric does not take into account traffic loadings, however.

Gorman et al. (2004) also studied resilience of telecommunications networks to physical attacks; their analysis allowed the modeling of both physical fiber networks and logical networks. They used a dataset consisting of aggregated Internet Service Provider (ISP) networks and studied the impact of removing nodes. Their primary measure of interest was the minimum number of hops required from the two farthest nodes in the network. They also did not study the impact of the node removal on the network traffic throughput.

Garbin and Shortle (2006) propose several resilience indices for telecommunications networks which is a promising concept, in that it can be compared across different network architectures. The behavior of these indices under real network conditions is the subject of ongoing research by Garbin and Shortle.

Other aspects of telecommunication network survivability research should also be mentioned. The impact of cyber attacks, rather than physical attacks, on the survivability of telecommunication networks was studied for example in Gorman et al. (2003). The design of survivable networks rather than their analysis has been the researched by, for example, by Medhi (1994), who presents algorithms to design a survivable network that considers the wide-area traffic network and the underlying transmission facility network simultaneously.

In the case of the NGN GETS, a single network provider, or perhaps a small number of providers, will be supplying service. Therefore, in this analysis we will be analyzing the survivability of an IP network representative of a single provider (rather than the large representations of the Internet studied by Albert and Barabási and others). In the following sections we will describe the methodology and numerical results for this survivability analysis.

Methodology

We will first describe our approach and scenarios. In this analysis we will use IP-SURVIV, a Noblis Survivability Analysis and Resilient Network Design Tool, to analyze throughput and connectivity of a network similar to a publicly available major ISP backbone as a function of outages in the network. The failure scenarios for analysis are

- Random electronics failures: affects fiber segments
- Directed/ terrorist attack: affects major cities (nodes)
- Natural disaster: affects both links and nodes in a region

The premise behind our analysis methodology is that logical backbone and edge links in a network are physically routed over a limited number of fiber optic transmission paths. Therefore, the disruption of a single fiber transmission segment can simultaneously fail several logical network links. Our analysis network is an IP network representative of a provider IP network. Given we are analyzing a single provider's network rather than the entire Internet as studied in Albert and Barabási (2002), the degree of our network nodes does not have a power-law tail distribution. Figure 1 shows the core provider nodes and links between them in this network. Figure 2 shows the provider edge routers and their connectivity to the core routers. The notional fiber connectivity for the network is shown in Figure 3; the limited number of fiber paths is evident.

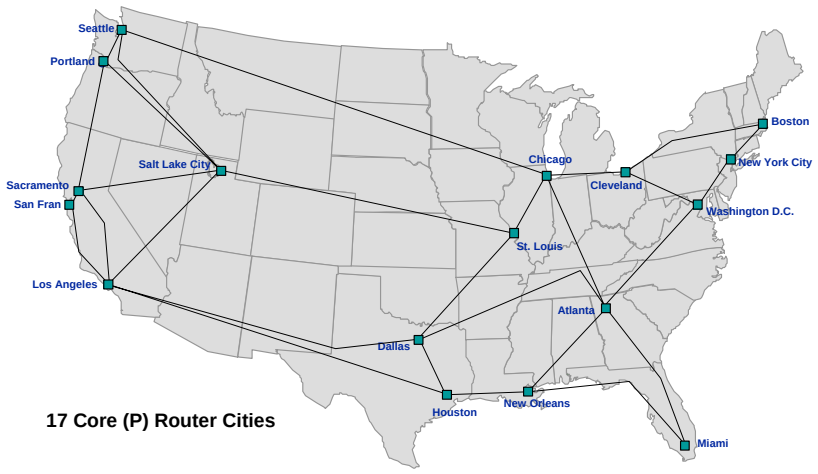

Figure 1. Network Core (Provider) Nodes and Links

Figure 4 shows an example of the shared risk groups in the network. The shared risk groups are the network fiber segments that correspond to multiple logical links. The network has a total of fifty fiber segments. Disruption of fiber segment 34, for example, causes outages on five separate network logical links (Dallas to Los Angeles, Houston to Los Angeles, Phoenix to Los Angeles, San Diego to Dallas, and Phoenix to San Diego).

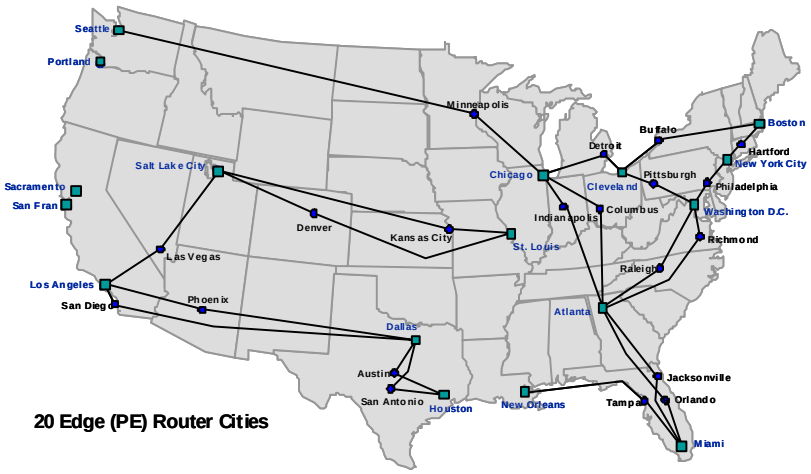

Figure 2. Network Provider Edge Nodes and Links

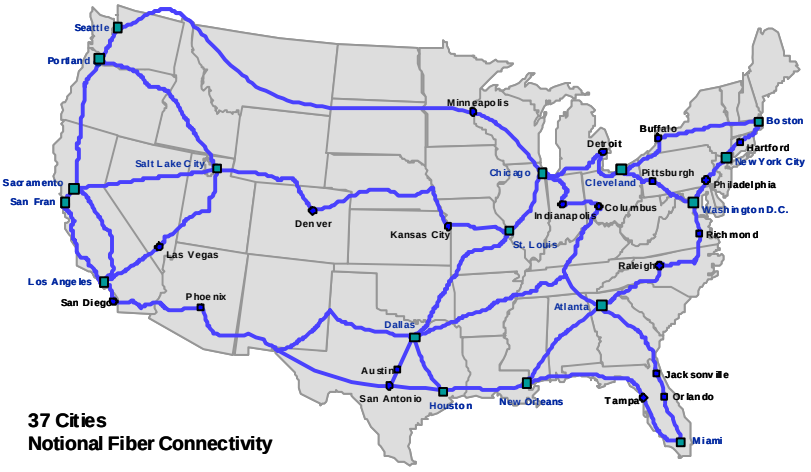

Figure 3. Network Notional Fiber Connectivity

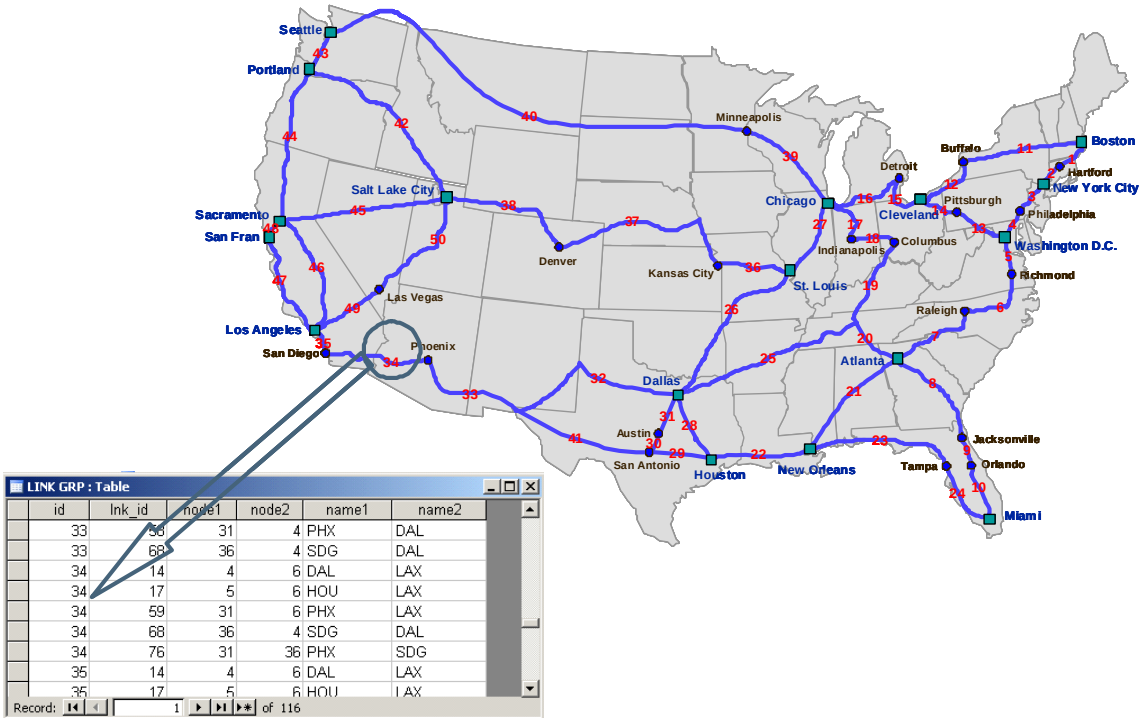

Figure 4. Network Shared Risk Groups

Our interest is in ensuring that the NS/EP traffic is carried during emergencies. Thus our metrics of interest will include not only connectivity (as was the primary focus in the literature that we reviewed), but also the traffic carried. Our metrics are:

- Percent of origin to destination (O-D) pairs communicating
- Throughput: the amount of data per time unit that has a route available to its final destination
- Maximum Round Trip Time (propagation delay)

We assume that traffic between all node pairs is population-based, and was developed from the 2005 Census data. The traffic is scaled so no packet loss in the baseline case with no failures. In addition, we assume that Open Shortest Path First (OSPF) routing is used, with the link distance used as the OSPF metric.

The impact of the three types of failures (random failures, directed attacks, natural disasters) on the above measures will be investigated. Random failures will be discussed first.

Random Failures Analysis

In the random failures analysis, the survivability of the network will be analyzed as increasing numbers of link groups in the network are randomly selected and disabled. Random link group outages are indicative of electronics failures occurring randomly through the network. The link groups are the sets of inter-switch connections lost from each of fifty segment outages, as shown by the inset in Figure 4.

Our IP-SURVIV tool is then configured to simulate disabling random selections of link groups. Segment outages are varied between zero and fifty simultaneous segments disabled. Replications (n=10,000) are performed for each number of disabled segments investigated.

Figure 5 shows both the mean throughput (as a percent of traffic) and mean percent of O-D pairs connected versus the percentage of segments disabled. Throughput is slightly lower than connectivity due to traffic weighting. However, the impact of the traffic assumptions here are minimal, as uniform traffic (percentage of O-D pairs connected) and non-uniform traffic (throughput) results are close.

The mean results in Figure 5 do not give a full representation of the results, however. Figure 6 shows the mean throughput as well as the extremes (minimum and maximum) and .1 percentile and 99.9 percentile of throughput. There is much variability in the results for a given percent of segments disabled. When one fiber segment is disabled, all traffic can be re-routed and throughput is 100 percent. However, when three fiber segments (or six percent of network fiber segments) are disabled, not all traffic can be re-routed to its destination and throughput ranges between 48 percent and 100 percent.

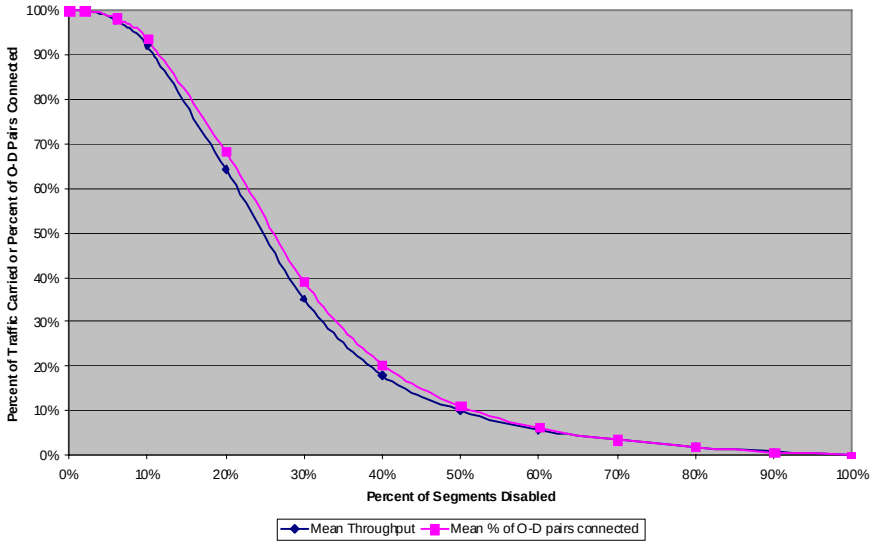

Figure 5. Mean Throughput and Connectivity versus Random Segment Damage

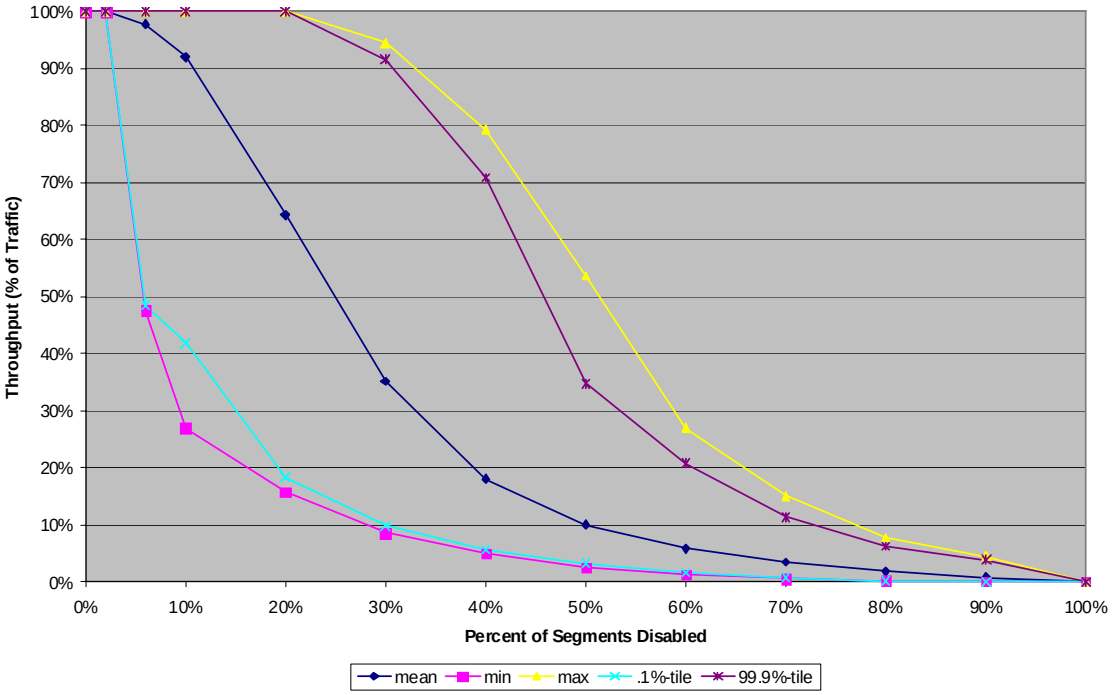

Figure 6. Means and Extremes of Throughput versus Random Segment Damage

The mean of the Maximum Round Trip Time versus the percentage of segments disabled is shown in Figure 7. The maximum round trip time is greatest when 20 percent of segments are disabled. When fewer then 20 percent of segments are disabled, much of the traffic is able to be rerouted to reach its final destination, but the routes are longer than the routes without segments disabled, resulting in a greater maximum round trip time. When greater than 20 percent of segments are disabled, connectivity is highly degraded and much of the traffic cannot reach its final destination, resulting in a lower maximum round trip time.

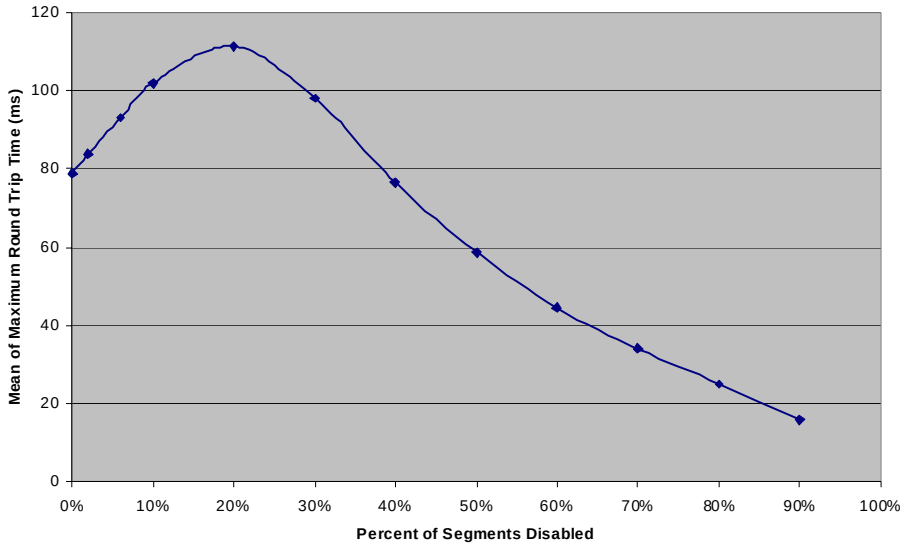

Figure 7. Mean of Maximum Round Trip Time versus Random Segment Damage

Directed Attack Analysis

In an intentional directed or terrorist attack, entire nodes would be targeted for disabling. In our analysis we remove a node by disabling all links connected to it. We investigate both single node outages and multiple node outages. Several methods of selecting nodes for simultaneous node outage are used:

- Population based selection: largest *m* nodes according to population are disabled simultaneously (e.g., if two nodes are attacked, they will be the two highest population nodes)
- Financial and Government Centers: New York, Los Angeles, Chicago, and Washington DC disabled simultaneously
- Random selection: *m* nodes are randomly selected for simultaneous disabling, with *n*=10,000 replications for each number of disabled nodes

Throughput, connectivity, and maximum round trip time are analyzed as network nodes are disabled. Figure 8 displays throughput in the case of single nodes disabled; the x-axis lists the nodes in order of decreasing traffic level. New York City, which has the heaviest traffic, results in the greatest degradation in throughput if it is disabled. Disabling centrally located nodes such as Chicago and St. Louis associated with links crossing the Mississippi River has the greatest impact on Maximum Round Trip Time. Albert and Barabási (2002) had conjectured that if nodes which are most highly connected are removed, that would be most damaging to a network; this is not the case for our network as seen by the results for disabling New York City. New York City has a degree of only two, yet has a larger impact on network throughput than disabling any of the highest degree nodes (Los Angeles, Atlanta, and Salt Lake City, which have degree five). This

difference is due to the Albert and Barabási focus on large scale-free networks, and their use of different measures such as the size of the largest network cluster that remains connected. Our interest is in throughput – ensuring that the NS/EP traffic is transmitted.

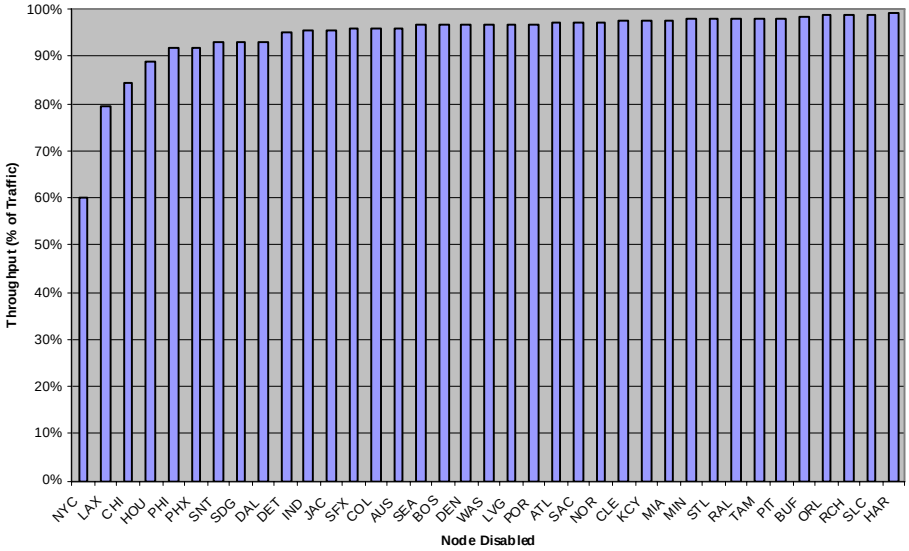

Figure 8. Throughput versus Single Node Disabled

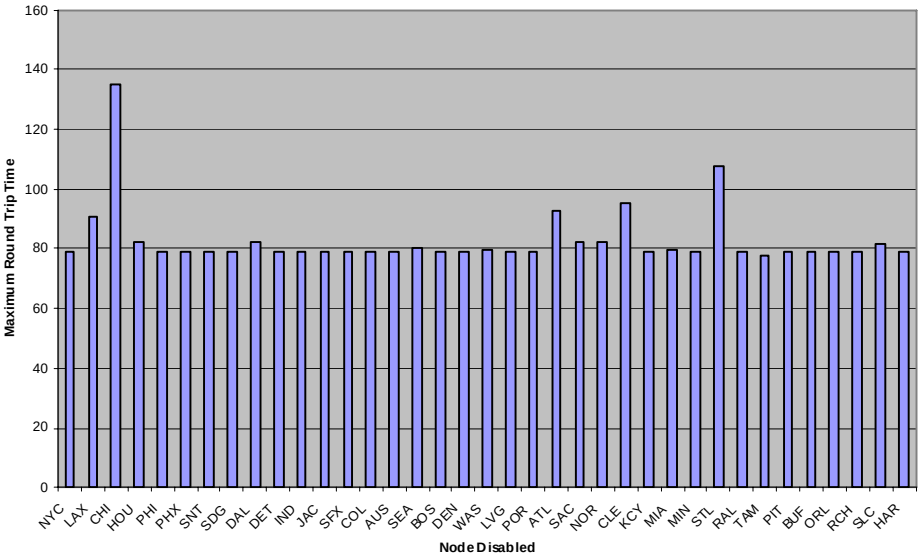

Figure 9. Maximum Round Trip Time versus Single Node Disabled

Increasing numbers of nodes were selected for simultaneous removal in order of decreasing node population (and thus originating traffic). These results are shown in

Figure 10. For example, when three nodes are disabled, they are New York City, Los Angeles, and Chicago. Throughput is shown as the blue curve, however, it is instructive in this case to examine loss (100% minus the throughput), which is the yellow curve. Traffic loss increases (or throughput decreases) at a higher rate than the amount of originating traffic or population that is disabled (shown as the pink curve), due to other traffic that cannot reach its destination due to the outages. Loss increases significantly when the nine largest nodes are disabled, and no path across the Mississippi River is available.

Maximum round trip time increases significantly when Chicago is disabled in addition to other nodes (see Figure 11). Also, maximum round trip time decreases for ≥ 24 percent of nodes disabled (the largest nine nodes), due to the inability for cross-country traffic to reach its destination.

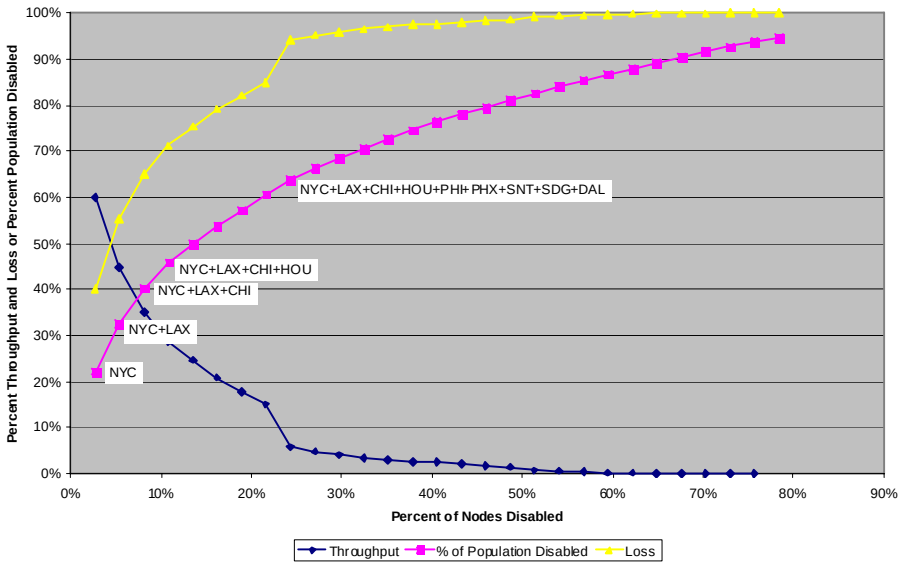

Figure 10. Throughput versus Percent of Nodes Disabled

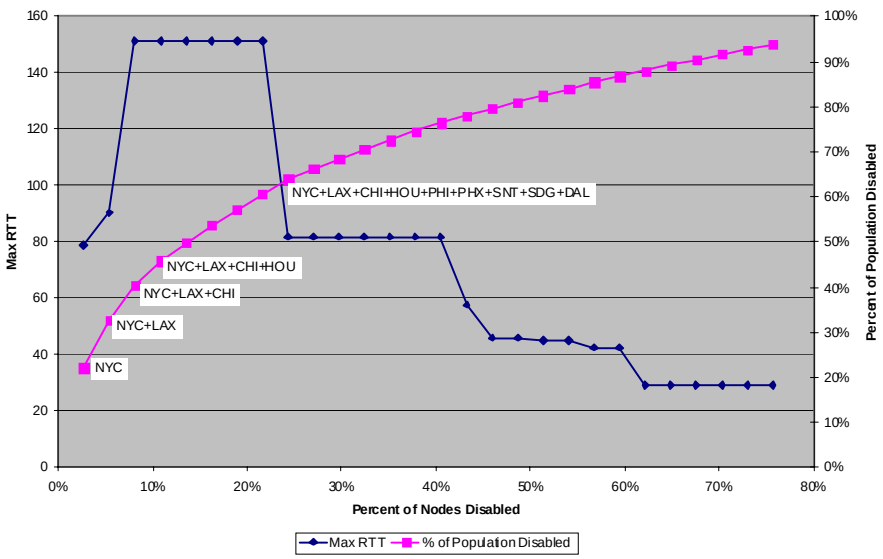

Figure 11. Maximum Round Trip Time versus Percent of Nodes Disabled

The impact of disabling financial and government centers on network survivability can also be examined. New York, Los Angeles, Chicago, and Washington DC are considered to be financial and government centers in the U.S., and we can compare the impact of disabling them on the network to disabling the four highest traffic and population cities (New York, Los Angeles, Chicago, and Houston) – see the first two rows of Table 1. Disabling New York, Los Angeles, Chicago, and Washington DC results in greater loss than if New York, Los Angeles, Chicago, and Houston are disabled, showing that survivability results are not easy to predict from traffic considerations alone.

Nodes Disabled	Throughput	Loss	Max RTT	% O-D Pairs Communicating
NYC, LAX, CHI, WAS	21.9%	78.1%	126.2	51.1%
NYC, LAX, CHI, HOU	28.8%	71.2%	151.2	79.3%
Mean Results for Four Randomly Selected Nodes	76.8%	23.2%	92.6	76.7%

Table 1. Disabling Four Nodes Comparison

If node outages are selected in a random or less obvious manner, rather than based on financial, government, or population considerations, expected throughput could be much higher than in the figures and table above. For example, Table 1 (last row) shows that the average throughput when four nodes are randomly selected is much higher than for four large nodes (e.g., NYC, LAX, CHI, WAS). Obviously, there is much variability in the ability of the network to carry traffic to its destination when a given number nodes are

disabled – see Figure 12. This was also seen when link segments were disabled in the previous section. The combination of four nodes giving the lowest throughput (13 percent) in our simulation was Los Angeles, Chicago, Dallas, and Atlanta.

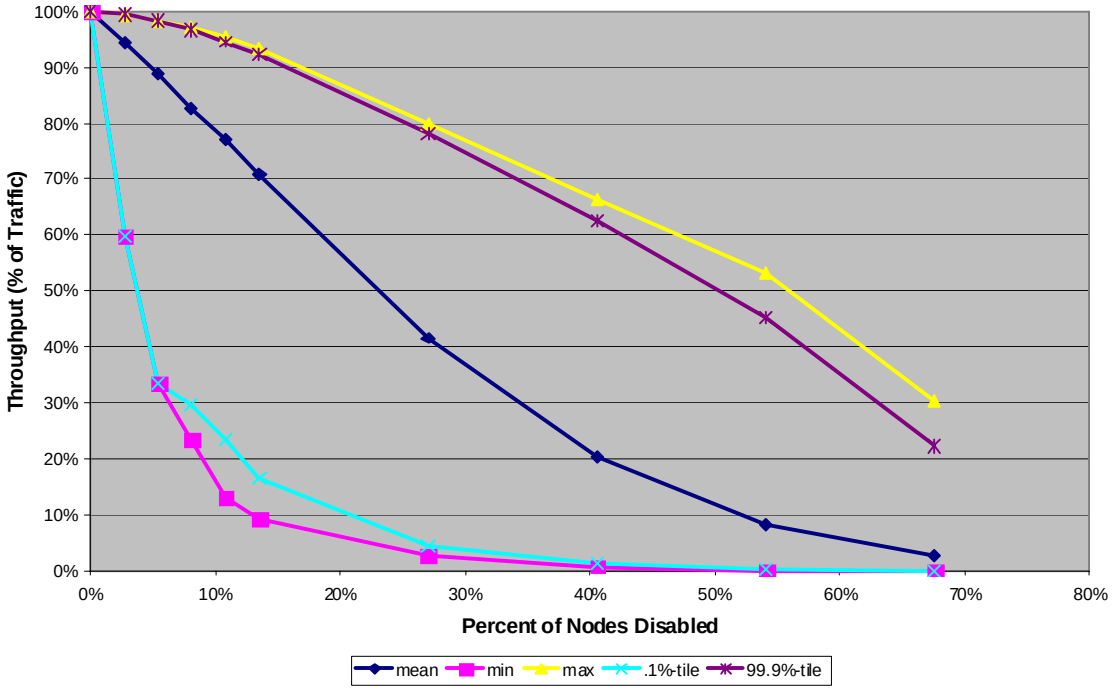

Figure 12. Means and Extremes of Throughput versus Random Node Damage

Natural Disaster Analysis

Network survivability to natural disasters was investigated by completely disabling regions of the U.S., and running IP-SURVIV. The U.S. Standard Federal Regions[1] were employed (Figure 13). In this analysis, outages are assumed to strike individual regions or neighboring groups of regions.

[1] The ten standard Federal Regions were established by OMB Circular A-105, *"Standard Federal Regions,"* in April, 1974. The regions are employed by a number of agencies including FEMA, HHS, GSA (GSA uses a slight deviation from the others).

Figure 13. U.S. Standard Federal Regions[2]

Figure 14 shows the throughput results as single regions or multiple neighboring regions are disabled. The single regions which most affect network throughput when disabled are Region 9, which contains six network nodes in California, Arizona, and Nevada, and Region 2, which contains the New York City node. When two regions are disabled, regions 5 and 6 and also regions 4 and 5 result in significantly reduced throughput (18 and 25 percent, respectively). This is because the disabling of either of these pairs of regions results in no routes for cross-Mississippi traffic.

Results for Maximum Round Trip Time are not shown. However, disabling Region 5 has a greater impact on Maximum Round Trip Time than any other single region, which is consistent with previous results of disabling the Chicago node alone. When two regions are disabled, maximum round trip time is most affected for Regions 6 and 7.

[2] Figure from http://www.fema.gov/about/contact/regions.shtm

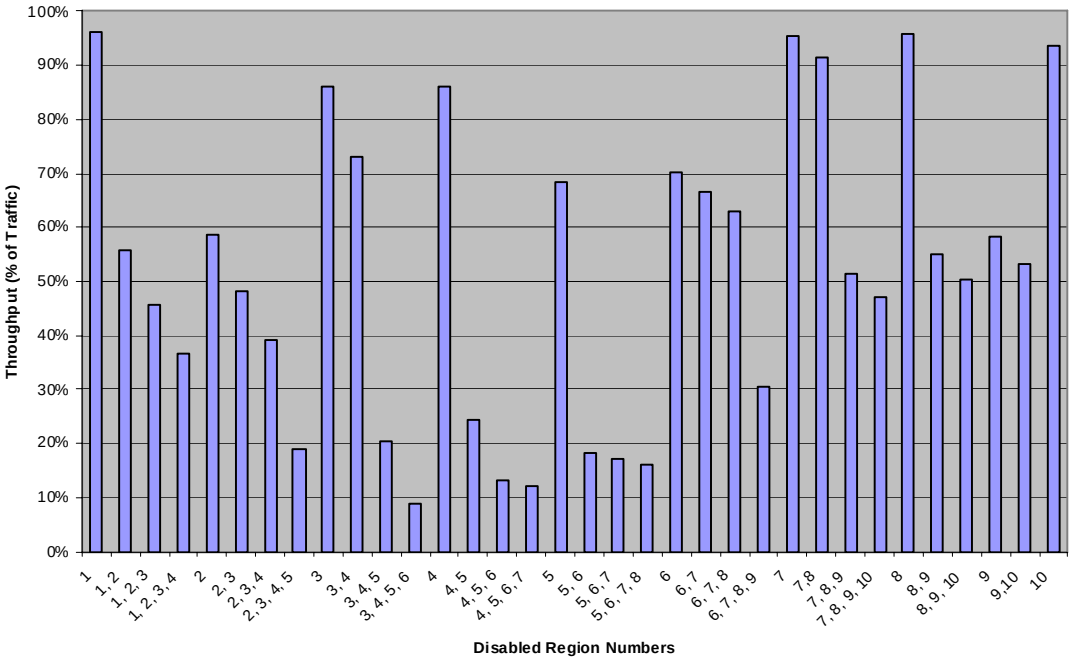

Figure 14. Throughput versus Region(s) Disabled

Summary and Conclusions

In summary, we have analyzed the impact of three types of failures on the survivability of an Internet Service Provider's backbone network. These failure types were random failures, directed attacks, and natural disasters. The measures we examined were throughput, percent of origin to destination pairs communicating, and maximum round trip time. These results will be used in specifying the survivability objective for the NGN GETS Voice Industry Requirements document. We found a large amount of variability in the results for a given percentage of the network (either nodes or link segments) disabled, and thus the survivability objective(s) will need to be well defined, to include perhaps the worst case results rather than mean results. For instance, based on our analysis, when 10 percent of network fiber segments are disabled in a random manner, network throughput could be as low as 27 percent or as high as 100 percent, depending on which segments were randomly selected. Our work differs from other studies in the literature which looked at different measures and sizes of networks. Most importantly, the high variability in the results point to the need for an automated tool like IP-SURVIV to analyze the impact of errors and attacks on the network.

We plan to conduct additional work in several areas. Both NGN and GETS survivability will depend on the ability of the call managers or processing elements to be reached by signaling. We have begun researching the impact of the number, placement, and connectivity of call managers on the network survivability. We also plan to extend the current survivability analysis to multiple domains, and investigate whether there is a marked difference in the ability to communicate depending on whether the GETS user is

limited to single carrier, or able to go between carriers. Lastly, we intend to develop models to analyze GETS survivability with wireless access.

Acknowledgments

This work was partially funded by the National Communications System Contract Number NBCH-D-02-0039 (Task Order Number D0200390095). Research funding was also provided by the Center for Network-Based Systems, a research collaboration of Noblis and George Mason University. We would like to thank Dr. Pat McGregor of Nyquetek for reviewing our assumptions and giving us feedback.

References

Alliance for Telecommunications Industry Solutions (ATIS), *Telecom Glossary 2000*, Document T1.523-2001, at http://www.atis.org/tg2k/

Albert, Réka and Albert-László Barabási. 2002. Statistical mechanics of complex networks, *Reviews of Modern Physics* 74: 47-97. At http://www.cs.unibo.it/~babaoglu/courses/cas/resources/tutorials/rmp.pdf

Garbin, D.A., Shortle, J.F. 2006. Measuring Resilience in Network-Based Infrastructures. *CIPP Working Paper 12-06*. Arlington, VA: George Mason University.

Gorman, S., Kulkarni, R., Schintler, L., and Stough R. 2003. Least Effort Strategies for Cybersecurity, *The Critical Infrastructure Protection Program: Workshop I Working Papers*. George Mason University, Fairfax, VA. At http://arxiv.org/ftp/cond-mat/papers/0306/0306002.pdf

Gorman, S., Schintler, L., Kulkarni, R., and Stough R. 2004. The Revenge of Distance: Vulnerability Analysis of Critical Information Infrastructure, *Journal of Contingencies and Crisis Management* 12: 48-63. At http://arxiv.org/ftp/cond-mat/papers/0310/0310427.pdf

D. Medhi. 1994. A Unified Approach to Network Survivability for Teletraffic Networks: Models, Algorithms and Analysis, *IEEE Transactions on Communications* 42: 534-548. At http://www.csee.umkc.edu/~dmedhi/papers/m_tcomm_94.pdf.

Shake, T.H., Hazzard, B. & D. Marquis. 1999. Assessing Network Infrastructure Vulnerabilities to Physical Layer Attacks. Paper prepared for the National Institute of Standards and Technology National Informations Systems Security Conference in Arlington, VA. At http://csrc.nist.gov/nissc/1999/proceeding/papers/p18.pdf