SAND2006-6371C

Network-layer Selective Security

Casey T. Deccio[*]
Sandia National Laboratories[†]
7011 East Ave, MS 9012
Livermore, CA 94550
Telephone: (925) 294-4784
Fax: (925) 294-2776
ctdecci@sandia.gov

Mark J. Clement
Brigham Young University
3361 TMCB PO Box 26576
Provo, UT 84602
Telephone: (801) 422-7608
Fax: (801) 422-0169
clement@cs.byu.edu

Quinn O. Snell
Brigham Young University
3361 TMCB PO Box 26576
Provo, UT 84602
Telephone: (801) 422-5098
Fax: (801) 422-0169
snell@cs.byu.edu

Spencer L. Cox
Brigham Young University
3361 TMCB PO Box 26576
Provo, UT 84602
Telephone: (801) 422-7290
Fax: (801) 422-0169
slc48@cs.byu.edu

Abstract

The Internet and other large computer networks have become an integral part of numerous daily processes. Security at the network layer is necessary to maintain infrastructure survivability in the case of cyber attacks aimed at routing protocols. In order to minimize undesired overhead associated with added security at this level, the notion of selective security is proposed. This research identifies elements in network topologies that are most important to the survivability of the network. The results show that the strategic placement of network security at critical elements will improve overall network survivability without the necessity of universal deployment.

1 Introduction

The Internet is the foundation of world-wide digital communication. Many critical applications depend on this enormous infrastructure for their functionality. The survivability of this and other large networks is vital to maintaining stability in many daily processes. This research aims to increase the dependability of a network by identifying and securing the most critical points within it.

At the heart of large network infrastructures, such as the Internet, are the network-layer protocols. At this layer routing mechanisms establish virtual links to fully connect entire networks, making global communication possible. These protocols were originally designed to operate in a trusted environment, without the threat of malicious nodes. This assumption has led to vulnerabilities in the network core. It has been commented that "abuse of the routing mechanisms and protocols is probably the simplest protocol-based attack available" [7]. The increased availability of tools allowing direct access of network resources to malicious users has made these attacks a reality [25]. Specific attacks targeting the routing infrastructure include routing table poisoning attacks, packet mistreatment attacks, and denial-of-service (DoS) attacks [9].

Security for higher-level (e.g., transport, application layers) protocols has been the focus of much recent research, but without security at the lower layers, computer networks are left vulnerable to attack. Security proposals for various routing protocols have surfaced in research, but the deployment rate of these mechanisms is low. Often this neglect is attributed to the performance cost, political logistics, and uncertainty associated with configuring something new onto all nodes in a stable network environment [19].

Rather than universal application of arbitrary security at all points within a given network, this re-

[*] Indicates contact author.

[†] Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

search suggests a selective deployment of network-layer security to protect *critical elements*—elements whose failure or attack would be most detrimental to network survivability. Overhead incurred by securing these critical elements is justified because of the risk associated with leaving them vulnerable. When the most critical elements of a network are secured against attack, the collective network graph remains more resilient to attackers.

This paper defines survivability in terms of the network routing infrastructure. Using metrics for measuring network performance, critical elements are identified, whose functionality determines in large part the survivability of the entire network. Related research and simulation results from empirical analysis are used to fortify the claim that securing critical elements will reduce the risk of a network catastrophe in the case of attack. The research and conclusions presented in this paper will provide a basis upon which a future selective security model might be designed and implemented.

2 Network-layer Security

Deployed network-layer protocols are some of the most vital mechanisms maintaining connectivity across local, national, and international boundaries. Though transparent to the end user, routing protocols are an integral part of every network system. This section gives an overview of the network layer and introduces the concept of selective security.

2.1 The Network Layer

The routing infrastructure maintains paths from all nodes to all other nodes within a network. Internet routing is hierarchical. *Autonomous systems* (AS) are networks managed by a central entity and utilize an *interior routing protocol* to manage routing within the network, such as the Open Shortest Path First (OSPF) protocol [23] or the Routing Information Protocol (RIP) [22]. The Internet is a complex network of AS that communicate using an *exterior routing protocol*, such as the Border Gateway Protocol (BGP) [28].

When successful protocol attacks are executed at the network layer, the effects are far-reaching. In 1997 routers at MAI Network Services, an Internet service provider (ISP) headquartered in Virginia, relayed bad router information from one of its customers onto Sprint's backbone. The bogus information propagated throughout Sprint's network, advertising MAI's network as the best route to get anywhere, and causing routers operated by Sprint and other ISPs to transmit all Internet traffic to MAI's network [29]. MAI's network was overwhelmed almost instantly by the extreme load, but routers nationally, and perhaps internationally, continued to forward data, creating a "black hole" scenario for several hours. During the outage, Sprint reported that most of its network was at 10% utilization, while the affected area was completely overloaded [31].

Although the cause of the mentioned routing incident was not a malicious attack, it demonstrates the ripples that can be felt throughout large networks, even when only one small part is compromised. Care should be taken to secure the network layer against undesired mishaps or attacks.

2.2 Selective Security

To protect network-layer protocols from attacks, research has produced security mechanisms. Often, however, development of these low-level security mechanisms does not reach a stable state, or they are simply not deployed. An example is the protocol for OSPF with Digital Signatures [24]. Although the draft for this was written in 1997, the status of the protocol is still "experimental" in 2006.

Why does routing security often fall short of deployment? Universal deployment of security mechanisms may seem unappealing for various reasons, which may include inter-organization logistics or politics, performance concerns, or concerns with complicating a stable network environment. The objective of selective security mechanisms is to *effectively* secure critical elements while striving to maintain a lower overhead (computational, political, or otherwise) than that accrued if all nodes were secured in a similar fashion.

An analysis of network-layer protocols will show that some network elements are more essential than others in maintaining a dependable network. The routing infrastructure exhibits a hierarchical characteristic, which means that elements have varying importance respective to the survivability of the overall infrastructure. Subsequent sections show that certain behaviors of network topologies place a higher reliance on particular elements in order to maintain survivability. Selective security techniques will apply the necessary measures to protect network elements of higher importance.

Two questions must be addressed in regard to the idea of selective network-layer security. First, what is the plausibility of applying various security mechanisms to nodes within the same infrastructure? The answer for each case depends on the specific protocol to which security is being applied. As an example, using authentication on only select routers in an OSPF network might require a non-trivial change to the specification. However, the OSPF with Digital Signatures specification currently allows for the possibility of non-authenticated *areas*—divisions within an OSPF network—working with authenticated areas [24]. A scheme with more variance from the original specification may have more trouble getting approved and deployed than one that closely compares with the original.

The second question regarding selective security is how well selectively deployed security mechanisms will protect the network from protocol attacks. This issue should also be analyzed by (1) identifying the protocol to which security will be applied and (2) identifying specific attacks to the protocol.

Effective design of a selective security mechanism should involve an analysis of the protocol that is being protected. An example of an ineffective selective security method is an OSPF with Digital Signatures network in which critical routers sign their update packets, but few other routers check the signatures. This is analogous to requiring patrons to show a current drivers license at an airport where only few attendants are verifying this document. When designing a technique for selective security, the critical elements should be thoroughly secured in order to assure network survivability. For this to happen, it may be necessary to secure more than the selected important elements; perhaps the security of some superset of those elements is required.

Akin to designing a security model for any other environment, designing a selective security model for routing protocols involves consideration of possible attacks aimed at the protocol. For example, protection against a routing table poisoning attack can be applied using the hashing or digital signing of update packets. However, this implementation would not protect against a router DoS attack.

Selective security is an abstract term in itself, and specific implementations may vary. This research does not directly discuss implementation of selective security, but rather helps identify critical points in the network that should be secured against attack.

3 Network Topologies

Communications networks respond differently to applied instances of attack or failure. This section examines some of the characteristics of the Internet and other large networks in order to identify critical elements within them. Characteristics of scale-free and random networks are discussed in this research.

3.1 Scale-free Networks

The complexities of the Internet topology and World-wide Web are attributed to the unmanaged and rapid growth that has occurred since its inception. The complex nature of these networks makes them difficult to classify. Related research has categorized similar infrastructures for social and biological systems that occur in nature [4, 32]. Albert and Barabási, et al. have observed that such large networks organize themselves into a *scale-free* state, and the results of their research are used to identify critical elements in large networks [5, 2].

Scale-free networks are characterized by their connectivity distribution $P(k)$, the probability that a node in the network is connected to k other nodes. In scale-free networks $P(k)$ decays as a power-law: $P(k) \sim k^{-\gamma}$ [5]. Relatively few nodes are highly connected in a scale-free network; the majority of nodes have very few neighbors. This relationship places enormous significance on the nodes with the highest degree. Figure 3.1 shows the graph (a) and connectivity distribution (b) of a scale-free network generated and visualized using the Pajek Program for Large Network Analysis [6]. The network follows the Barabási-Albert extended model [1] and is comprised of 100 nodes connected by 400 directed links. Its connectivity distribution approximates the model $P(k) \sim k^{-1.6}$, which is also graphed in Figure 3.1b.

One set of Internet topology data used for analysis in this research consists of data from the SCAN project obtained in 1999 using the Mercator software [17] merged with data also obtained in 1999 from the Internet Mapping Project at Lucent Bell Laboratories [20][1] These studies produced a topology consisting of 284,805 connected Internet routers, with connectivity $P(k) \sim k^{-2.3}$. This data is hereafter referred to as the "scan+lucent" data. The connectivity distribution of the "scan+lucent" data is shown in Figure 2. The probability that a network node is connected to 100 others is $P(100) = 2.5 \times 10^{-5}$, while the probability that a node only has one neighbor is extremely high $P(1) = 0.53$.

The scale-free distribution carries with it properties of extreme robustness when nodes are disabled at random. However, it network functionality degrades rapidly when the most connected nodes are targeted [3]. Because the concentration of highly-connected nodes represents only a small percentage of the whole network, a loss of a small percentage of these critical nodes is extremely damaging to the functionality of the network. Section 5.1 discusses

[1] The Internet Mapping Project is now run by Lumeta Corporation. More information can be found at their Web site [13].

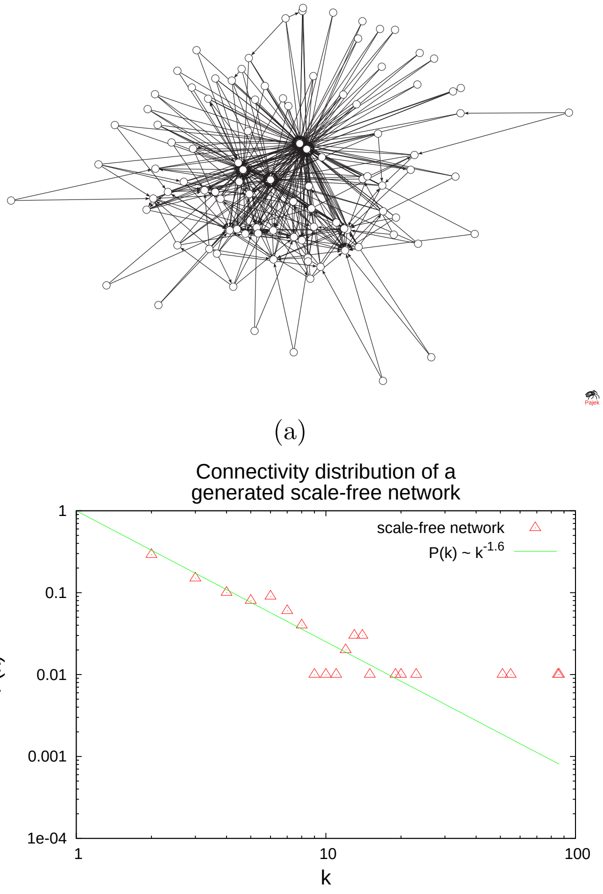

(a)

(b)

Figure 1: The graph (a) and connectivity distribution (b) of a scale-free network with 100 nodes connected by 400 directed links. The network was generated and visualized using the Pajek Program for Large Network Analysis [6] and follows the model $P(k) \sim k^{-1.6}$, which is plotted against its connectivity distribution.

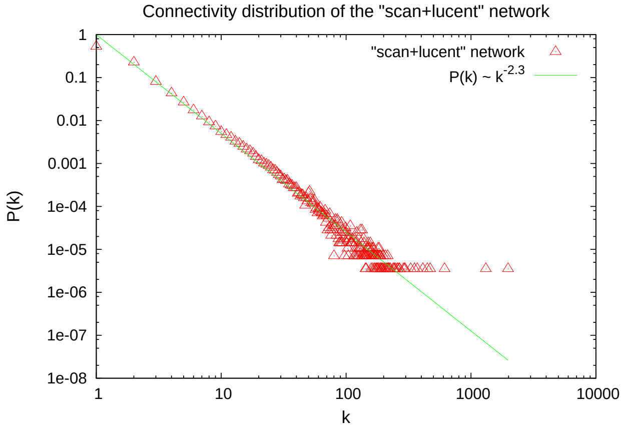

Figure 2: The connectivity distribution of the "scan+lucent" network. The line below the plot points represents the connectivity distribution for the scale-free model $P(k) \sim k^{-2.3}$.

the identification of critical elements in scale-free networks.

3.2 Random Networks

Relatively smaller networks, such as AS managed by a sole entity, are often classified as *random networks*. There are many models describing random networks. Erdős and Rényi define a model $G(n, P\{u, v\} = p(n))$, $0 \leq p(n) \leq 1$, in which each possible edge between two vertices u and v is added with probability $p(n)$ to the graph [8]. At $p(n) = \frac{1}{2}$ any graph with n nodes is equiprobable. Figure 3 shows the graph (a) and connectivity distribution (b) of a random network that was generated and visualized using the Pajek Program for Large Network Analysis [6]. This network follows the Erdős-Rényi model [8] and has 100 nodes connected by 398 directed links and approximates a normal distribution with $\mu = 7.94$ and $\sigma = 3.09$.

Networks following a random graph model exhibit characteristics different from those of scale-free networks. Most notably, random graphs generally follow a pattern of *homogeneity*; the connectivities of the nodes in this model are approximately the same, and each node in the network contributes equally to the stability of the entire network graph: if any network node is lost, the damage is approximately the same as if any other node were lost instead [3]. Thus, attacks directed at the most connected network nodes will not harm the network more than attacks at random nodes.

The results of simulations involving attack and survivability, as well as methods for identifying critical elements in random networks are discussed in Section 5.2.

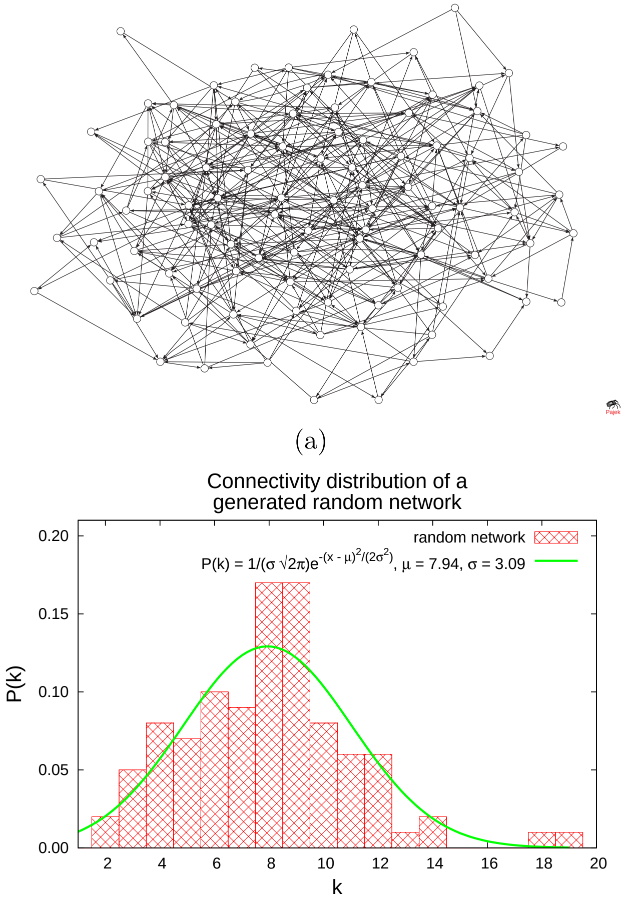

(a)

(b)

Figure 3: The graph (a) and connectivity distribution (b) of a random network with 100 nodes connected by 398 directed links. The network was generated and visualized using the Pajek Program for Large Network Analysis [6] and approximates a normal distribution with $\mu = 7.94$ and $\sigma = 3.09$.

4 Measuring Network Performance and Survivability

Defining survivability is a crucial step in identifying critical network elements. This section defines survivability and outlines several metrics for quantifying network performance in a communications network.

4.1 Network Survivability

A *survivable* system is able to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents [14]. Vital systems not exhibiting survivability may result in catastrophic consequences when undesired events are experienced—even the loss of human life.

An analysis of potential failures in real network environments demonstrates the notion of survivability in relation to the mission of an organization. AT&T—a large telecommunications company—maintains a standard service-level agreement guaranteeing 99.99% network availability to its customers. That is equivalent to 43 minutes of allowable down time per month. In 2001 the malfunction of a single switch on AT&T's ATM network overloaded 7% of all the network switches for about four hours, greatly exceeding the allowed down time in its service agreement [26].

Even the Internet—which urban legends claim could withstand the effects of a nuclear bomb [33]—is not without an "Achilles' Heel". If the functionality of just 5% of the Internet's most highly-connected nodes is lost through attack, the complete infrastructure becomes fragmented and unusable [3].

The first step in maintaining network survivability is to identify the network's "mission", so that execution efficiency of that mission can be evaluated in the presence of attack [14]. This research deals with general computer communications networks, which are expected to maintain a certain level of performance. It is therefore necessary to identify metrics for quantifying network performance in scale-free and random graphs. Metrics from queuing models are helpful for the monitoring, analysis, and quantifying of network behavior under a range of failures and attacks [18]. Topological characteristics desirable for good network performance are also outlined.

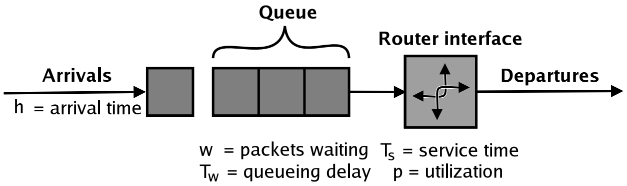

Figure 4: A queuing model for a network router interface. The parameters λ, T_w, T_s, and ρ affect network performance.

4.2 Queuing Model

Network delay is a critical performance metric which can be degraded in the presence of network attacks. Queuing theory can be used to estimate the incnreased delay caused by a network attack. A network router can be analyzed as a single-server queue (see Figure 4) [30]. Data packets arrive at the router at rate λ and must be serviced with average *service time* T_s. When λ is such that arriving packets cannot immediately be serviced, arriving packets must "wait" in line (queue) to be serviced behind other packets that arrived previously. When the router interface is available to service the next packet, a packet is selected from the w waiting packets according to some policy (e.g., a *first-in-first-out* or FIFO policy). The *link utilization* ρ is the fraction of time that the dispatching interface is "busy" servicing packets, measured over some period of time [30]:

$$\rho = \lambda T_s \tag{1}$$

When $\rho = 1.0$, the interface is saturated. Thus, the theoretical maximum input rate that can be handled by a router is [30]:

$$\lambda_{\max} = \frac{1}{T_s} \tag{2}$$

However, the finite buffer size of a router usually limits the maximum input rate to 70–90% of the theoretical maximum. *Queuing delay* T_w is the average time spent waiting to be serviced, and is calculated using Little's formula [30]:

$$T_w = \frac{w}{\lambda} = \frac{wT_s}{\rho} \tag{3}$$

As a link approaches capacity (i.e., $\rho \to 1.0$), delay becomes arbitrarily high [11].

The parameters of the above router queuing model will affect the overall flow of traffic through a network. The total data successfully transmitted across a network over a period of time is known as *aggregate throughput* H. Because seamless data transfer is the primary goal of a communications network, analysis of aggregate throughput amid varying conditions provides a measure of network efficiency, and higher throughput is an indicator of better performance.

The metrics described in this section will be used in Section 5 to evaluate network performance after networks have been targeted for attack. This will be a measure for how much elements affect overall network survivability.

4.3 Topological Characteristics

Topological characteristics can be used to indicate some measure of performance of the network. Several metrics have been defined to describe the *interconnectedness* of a network—a property describing how closely-linked the topology is. In a graph G, the *distance* $d(u, v)$ between two nodes u and v is defined as the length of the shortest path joining u and v. If $d(u, v) = \infty$ for any two network nodes, the network is *fragmented*—that is, there are isolated clusters of nodes in the network. *Diameter* $D(G)$ is defined as the maximum distance between any pair of nodes in G and corresponds to the delay of data passed through the network [10]. The *average distance* $\langle d \rangle$ over all pairs of nodes in a network is also helpful in determining network efficiency [2]. Small diameter and average distance are desirable characteristics for a communications network, and result in higher network performance [16].

If a network becomes fragmented as the result of the failure or attack of one or more nodes or links, remaining nodes are grouped into clusters according to which nodes or links have been disabled. As the network is fragmented into clusters, nodes have no reliable path for transmitting data to and from the nodes outside their cluster (see Figure 5), and the network's ability to successfully transfer data is diminished. In order to maintain reliable network communication security should guard against attacks that will fragment the network.

Analysis of network fragmentation suffered and increase in $\langle d \rangle$ incurred will be used in Section 5 to quantify network survivability when arbitrary network elements are protected against network-layer attacks.

5 Identifying Critical Network Elements

This section presents methods for quantifying the value of a network element within the network. The

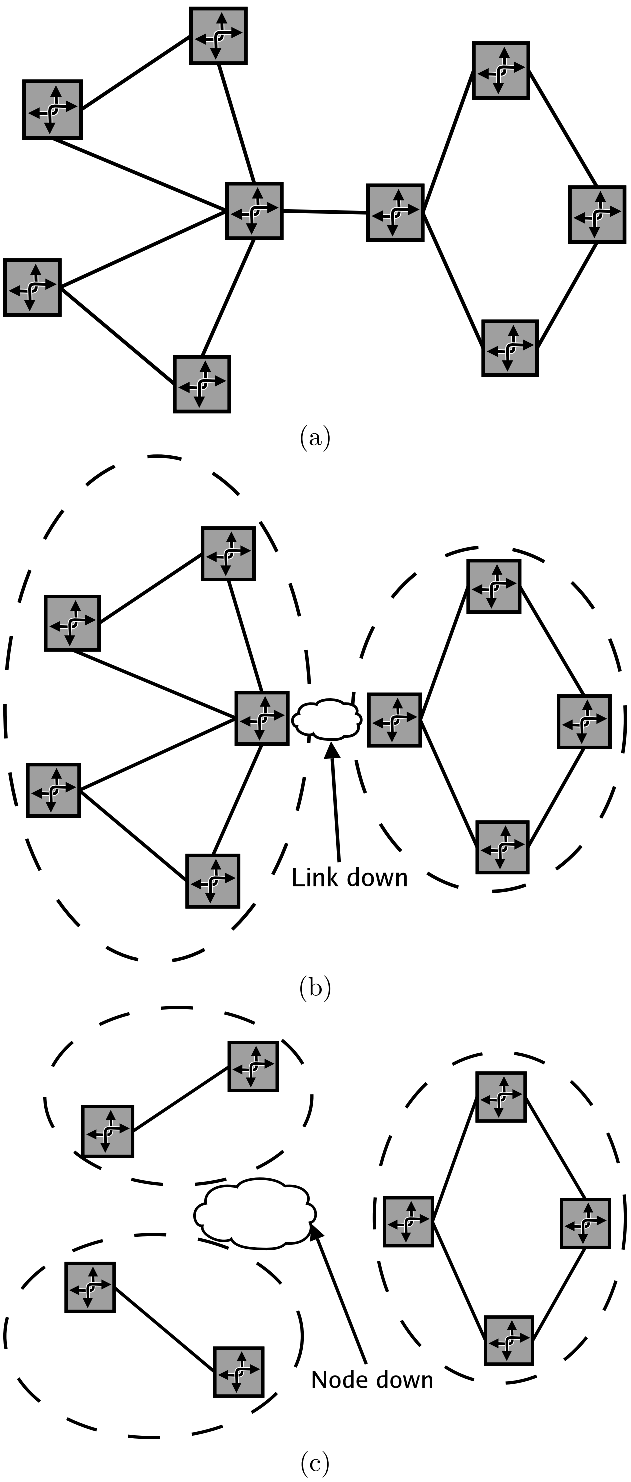

Figure 5: The formation of isolated clusters within a network, as the result of disabled link or node: (a) the original network; (b) clusters formed after a link is disabled; and (c) clusters formed after a node is disabled.

results are based on previous research as well as empirical data gathered from network simulation.

Different graph models, such as those discussed in Section 3, present different challenges for measuring survivability, using the metrics introduced in Section 4. However, strategies can be deployed to identify critical elements in networks having varying characteristics. In particular, methods of identification within scale-free and random networks will be discussed.

5.1 Critical Elements of Scale-free Networks

The scale-free networks discussed in Section 3.1 are distinguished by a small concentration of highly connected nodes. The work of Albert and Barabási, et al. has shown that as the nodes of a large network are disabled in order of decreasing connectivity, the network becomes fragmented and unusable when only 5% have been directly disabled [3]. Results of simulations produced in this research are comparable to the results published by Albert and Barabási, et al.; the most connected nodes in a scale-free network are most critical to the network's survivability.

5.1.1 Simulation Environment

In order to analyze survivability of a scale-free network, a software tool was created for simulating network attacks. The topological information from the "scan+lucent" network was imported into this simulator, and the simulator removed network nodes from the graph iteratively, without replacement. In order to reduce the computation time required to calculate essential network metrics on this large network, 200 of the original 284,805 nodes were disabled at each iteration for this simulation, and the resulting network after each iteration was the largest cluster of connected, functioning nodes. At each time step the resulting network was analyzed.

The simulation was performed once without any protection and once with the 10% most connected nodes secured against attack. In both runs the nodes with the highest degree were targeted. An additional run simulate attacks at random nodes. The attack model applied is construed as either a crippling of the node itself or the incapacitating of the set of links connecting it to other network nodes.

5.1.2 Simulation Analysis

In the attack model, wherein network nodes were disabled in order of decreasing connectivity, the fragmentation in the network severely crippled its abil-

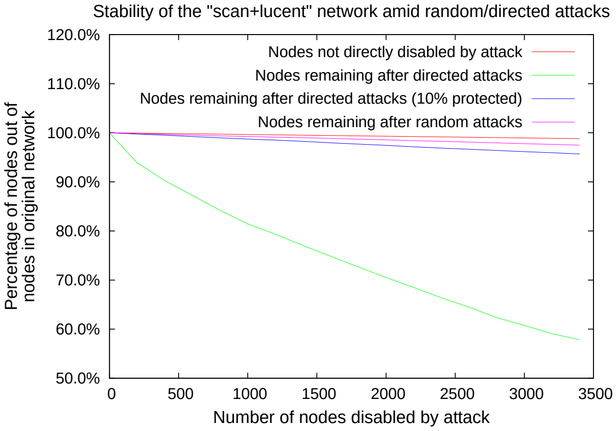

Figure 6: Network stability of the "scan+lucent" network under different attack strategies. When over 1% of the nodes are randomly removed from the network, the network retains over 97% of its original nodes, but when 1% of the most connected nodes are removed from the network, the size of the functional network drops below 60% of its original size. If only the 10% most connected nodes are protected, and the network suffers a directed attack,' the network remains resilient.

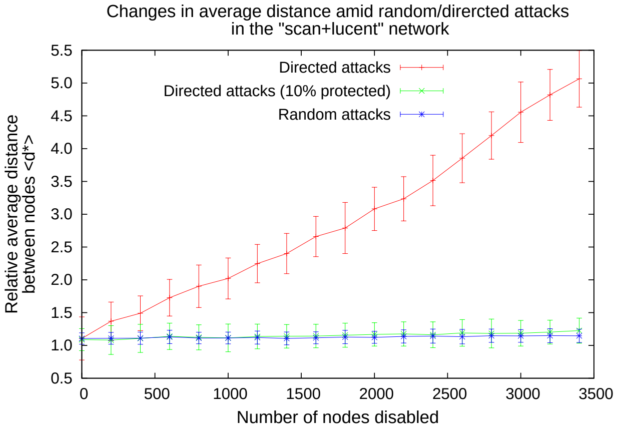

Figure 7: Changes in the relative average distance $\langle d^* \rangle$ between nodes of the "scan+lucent" network in the presence of different attack strategies, shown with 95% confidence intervals. The value $\langle d^* \rangle$ remains nearly unchanged even after 1% of the nodes are randomly removed from the network. However, $\langle d^* \rangle$ increases linearly as nodes are removed in order of decreasing connectivity. Protection of the 10% most connected nodes shows little change in $\langle d^* \rangle$, although the network suffers a directed attack.

ity to function (see Figure 5.1.2). When 1% of the most connected nodes were effectively disabled by attackers, the largest connected cluster remaining was comprised of less than 60% of the original nodes.

In contrast, Figure 5.1.2 also shows the result of losing large numbers of nodes due to random attacks. Although 1% of the nodes were directly attacked, the network retained over 97% of its original nodes. When the 10% most connected nodes were secured from attack (Figure 5.1.2) the network retained over 95% of the original nodes for operation.

The average distance $\langle d \rangle$ of the "scan+lucent" network was calculated at each iteration of attack and failure using a statistical sample of the entire remaining network. The average distance from each of 1000 randomly-selected nodes to all other nodes was calculated:

$$\langle d \rangle = \frac{\sum_{u \in S} \sum_{v \in G | v \neq u} d(u,v)}{|S|(|G|-1)} \tag{4}$$

where S is the set of randomly selected nodes, and G is the set of entire network nodes.

It should be noted that in general as the number of links increases in a network with a fixed number of nodes, $\langle d \rangle$ decreases [3]. This makes it difficult to compare $\langle d \rangle$ values for networks with different numbers of nodes or links. For this reason the term *relative average distance* $\langle d^* \rangle$ is introduced, which is a ratio of the calculated average distance to the number of unique source/destination pairs in G, multiplied by a constant c in order to bring the values being compared into a more reasonable range for comparison:

$$\langle d^* \rangle = \frac{\langle d \rangle}{|G|(|G|-1)} c \tag{5}$$

The simulation results with $c = 10^{10}$ are shown in Figure 7 with a 95% confidence interval. The value $\langle d^* \rangle$ increased linearly when the nodes were disabled in order of decreasing connectivity, but for the simulation involving simple failures $\langle d^* \rangle$ remained almost unchanged, despite the loss of over 1% of the network's nodes. When the 10% most connected nodes are protected against attack, $\langle d^* \rangle$ again shows little change, even when the target of a direct attack.

The results of the simulations performed on the scale-free "scan+lucent" network support the claim that critical elements can be identified in scale-free networks. When the most connected nodes are secured against attack in a scale-free network, major network fragmentation will be prevented. In addition, if the most connected nodes are secured from attack, $\langle d^* \rangle$ will not increase significantly, although other random nodes may have lost functionality. These attributes make the network more survivable.

5.2 Critical Elements of Random Networks

Because of the topological differences between scale-free and random networks, the analysis and conclusions drawn about scale-free networks in Section 5.1 do not necessarily apply to random networks. The research of Albert and Barabási, et al. [3] shows that nodes attacked in order of decreasing degree and at random disable the network in a similar fashion. In this section link analysis is used to identify critical elements in random networks.

5.2.1 Max-flow Min-cut

Communications networks support a finite flow of data through their systems. Each link $\{u, v\}$ has a limited *capacity* $C(u, v)$ that affects the overall behavior of traffic flow in the network. The *maximum flow* (max-flow) is the greatest rate at which data can be sent from a source s to a destination t without violating capacity constraints [12]. Data flows in the network are referred to as *commodities*, and each commodity has a demand $D(s, t)$ [21].

A *cut* $(U, \bar{U})$ of a graph G is a partition of G into U and $\bar{U} = G - U$. The *capacity* of this cut is the sum of the capacities linking U and $\bar{U}$:

$$C(U, \bar{U}) = \sum_{\{u,v\} | u \in U \wedge v \in \bar{U} \vee v \in U \wedge u \in \bar{U}} C(u, v) \tag{6}$$

The sum of demands whose source and sink are on opposite sides of the cut is the *demand* of the cut separating U and $\bar{U}$:

$$D(U, \bar{U}) = \sum_{\{s,t\} | s \in U \wedge t \in \bar{U} \vee t \in U \wedge s \in \bar{U}} D(s, t) \tag{7}$$

In a *uniform multicommodity flow problem*, it is assumed that there is a commodity for each unique node pair in the network, and each commodity has the same demand. The demand of such a cut is simply [12, 21]:

$$D(U, \bar{U}) = |U||\bar{U}| \tag{8}$$

In this paper uniform multicommodity flow problems are used as a case study for examination.

The *min-cut* of a network graph $\theta(G)$ is the cut with the lowest capacity-to-demand ratio [21]:

$$\theta(G) = \min_{U \subseteq V} \frac{C(U, \bar{U})}{D(U, \bar{U})} \tag{9}$$

The set of links comprising the min-cut might be characterized as a "bottleneck" in the network—vulnerable but vital strands which attach two network partitions. The vulnerability lies in the high utilization of those links spanning the cut. If one or more of that set are disabled, as the result of an attack, network congestion will likely increase. The load that was once distributed across several links will now rest on the remaining links, potentially overloading their already weighted load. If all of the links are successfully disabled, then the network becomes fragmented.

The min-cut problem suggests a solution to identifying critical elements in random networks. If the links comprising the min-cut of a network are attacked, the effects of fragmentation or congestion will be felt throughout the network. However, if these links are secured, the network is more survivable to attacks.

5.2.2 Link Valuability

The solution to the min-cut is a set of links within a network. Therefore, using only these metrics, it is difficult to quantify and compare the values of different links within the network. In order to effectively do this, link *valuability* ϑ of a link $\{u, v\}$ is defined for uniform multicommodity network here:

$$\vartheta(u, v) = \frac{\sum_{U | u \in U \wedge v \in \bar{U} \vee v \in U \wedge u \in \bar{U}} \frac{D(U,\bar{U})}{C(U,\bar{U})}}{|U \mid u \in U \wedge v \in \bar{U} \vee v \in U \wedge u \in \bar{U}|} c \tag{10}$$

where c is a constant used only to bring the values being compared into a more reasonable range for comparison. Link valuability is the average demand-to-capacity ratios of all network cuts of which it is a part. By definition as link valuability increases, the expected utilization of the link will increase, and the link's importance with respect to overall network survivability will increase.

5.2.3 Simulation Environment

In order to test how well the link valuability property holds, a simulation was designed to test the performance of a dumbbell-shaped graph that shares the homogeneous property of random networks (i.e., nearly all links have the same number of neighbors). This 4-graph, shown in Figure 8, is comprised of ten nodes, and was simulated using the network simulator *ns-2* [15]. Each network link had a capacity of 5.0Mbps. The network was designed to make links $\{3, 5\}$ and $\{4, 6\}$ the most critical to the infrastructure. Using $c = 10^7$, these links each had a valuability $\vartheta(3, 5) = \vartheta(4, 6) = 5.739$, and the other links had valuabilities ranging from 5.565 to 5.655. A link-state routing protocol was used in the network to establish routes.

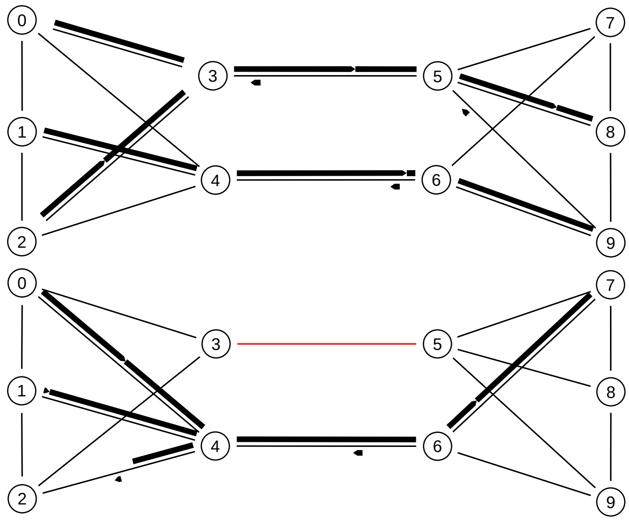

(a)

(b)

Figure 8: A dumbbell-shaped homogeneous network used for network simulation and analysis in which links $\{3, 5\}$ and $\{4, 6\}$ have higher valuability than the others: (a) FTP traffic flows from sources at nodes 0–2 to sinks at nodes 7–9 over highly utilized links $\{3, 5\}$ and $\{4, 6\}$; (b) link $\{3, 5\}$ has been disabled, and traffic originating at nodes 0, 1, and 2 is passed only through link $\{4, 6\}$ to destination nodes. Figures generated using the Network Animator (Nam) [15].

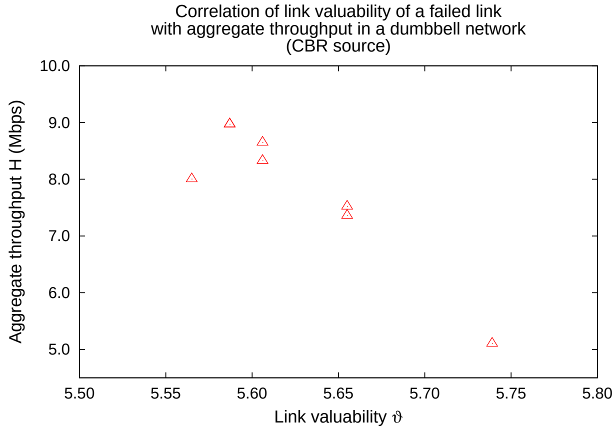

Figure 9: Correlation of aggregate throughput H from simulations involving CBR traffic in a dumbbell-shaped network with the link valuability $\vartheta(u, v)$ of a particular link $\{u, v\}$ dropped in each simulation. As the valuability of a failed link increases, the aggregate throughput in the network decreases.

A series of simulations was run in which 1.0Mbps constant bit rate (CBR) traffic was transmitted from each of nodes 0–2 to each of nodes 7–9 over the the user datagram protocol (UDP). In each simulation a particular link was disabled at 5.0 seconds of simulation time. The network traffic then continued to run for 15.0 additional seconds in order to monitor network performance following the link attack, and the metrics measured at each sampling interval were averaged over the entire 15.0 seconds. A simulation was run once for each link in the network, disabling the link on that run.

A graph correlating aggregate throughput H with link valuability ϑ in this series of simulations is shown in Figure 9. This graph displays a correlation between $\vartheta(u, v)$ and the H resulting from the attack of link $\{u, v\}$; as ϑ increases the resulting H decreases. When these results were analyzed using the statistical program R [27], it produced a high linear correlation value of 0.915.

An analysis of metrics at network routers shows how the loss of more valuable links further impacts network performance. Figure 10 maps the 85th percentile of ρ following a link failure to the corresponding valuabilities of disabled links. Utilization ρ increases as the valuability of the disabled link increases. R produces a correlation value of 0.453 for the 85th percentile ρ value in the dumbbell network with CBR traffic. Although the correlation is not as high as that of the aggregate throughput, the results show that failure of the most valuable links causes

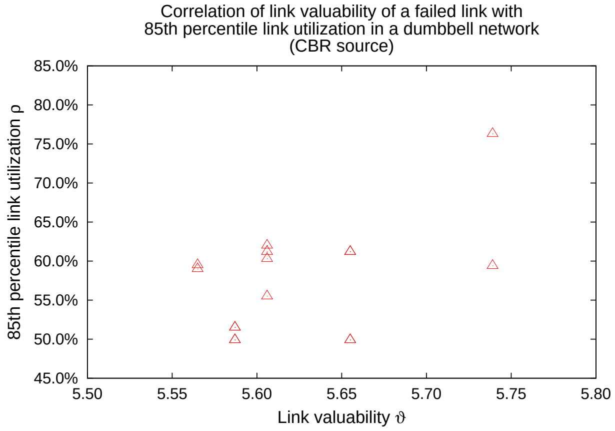

Figure 10: Correlation of 85th percentile link utilization ρ from simulations involving CBR traffic in a dumbbell-shaped network with the link valuability $\vartheta(u, v)$ of a particular link $\{u, v\}$ dropped in each simulation.

network utilization to increase the most.

Queuing delay T_w in networks involving link attack were also analyzed in the series of simulations using CBR traffic. The 85th percentile queuing delays, calculated using Little's formula [30], is shown in Figure 11. The latter graph shows that when the most valuable network links fail, queuing delay is the highest.

A network following the Erdős-Rényi model [8] was generated with the Pajek Program for Large Network Analysis [6], so that the results might be verified on another random network. This 7-graph consisted of 10 nodes connected by 24 10Mbps links. Traffic in the network was generated by 150 CBR sources distributed uniformly across the network, with corresponding destinations also uniformly distributed. A series of simulations was performed, as in the case of the dumbbell network. A mapping of $\vartheta(u, v)$ to the aggregate throughput of network data after the failure of corresponding link $\{u, v\}$ failed is shown in Figure 12. As in the case of the dumbbell network, as the valuability of a failed link increases, the aggregate throughput in the network decreases. R correlated this data with a value of 0.634.

The results of multiple simulations using a dumbbell network graph with several traffic models show that when links with higher valuabilities are disabled, network performance suffers more than if less valuable links are disabled. These results were verified on a generated random network.

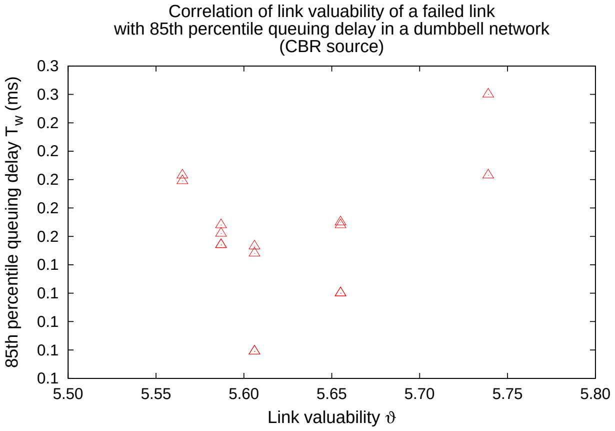

Figure 11: Correlation of 85th percentile queuing delay T_w from simulations involving CBR traffic in a dumbbell-shaped network with the link valuability $\vartheta(u, v)$ of a particular link $\{u, v\}$ dropped in each simulation.

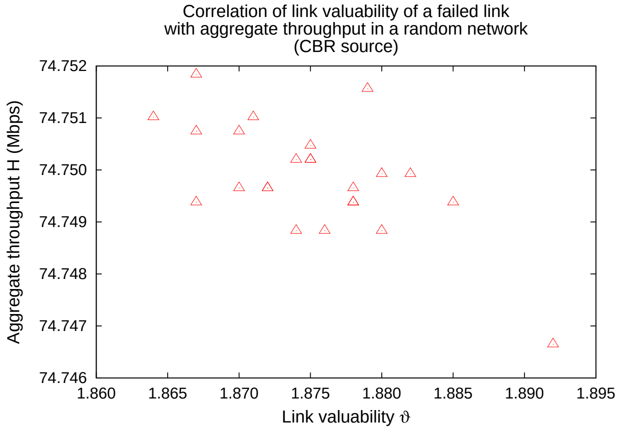

Figure 12: Correlation of aggregate throughput H from simulations involving CBR traffic in a random network with the link valuability $\vartheta(u, v)$ of a particular link $\{u, v\}$ dropped in each simulation. As the valuability of a failed link increases, the aggregate throughput in the network decreases.

6 Conclusions

Recent technology has enhanced communication globally with the development of large communications networks, the largest of which is the Internet. In order to guard against network-layer protocol attacks, these infrastructures should be secured. Because the computational or logistical overhead of universally securing all nodes in a particular network can slow down or even prevent the application of secure routing mechanisms, this research shows that critical elements in computer networks can be identified, elements whose attack would be detrimental to the survivability of the collective network graph. Security applied to these elements will increase the survivability of entire networks.

This research discussed the notion of selective network-layer security to protect against cyber attacks. The characteristics of scale-free and random network models were described, and data from a real-world graph following these models was obtained for survivability analysis. Network survivability was defined, and link and topological metrics were used to quantify network performance.

Attack models simulated on a scale-free network topology showed that the nodes with the highest degree are most critical to network survivability. Simulated attacks on random networks showed that the links with the highest valuability are most critical to network survivability. The selective deployment of network security to critical elements can improve overall survivability in networks of high importance.

References

[1] Réka Albert and Albert-László Barabási. Topology of evolving networks: Local events and universality. *Physical Letter Reviews*, 85(24):5234–5237, May 2000.

[2] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Diameter of the World-Wide Web. *Nature*, 401:130–131, Sep 1999.

[3] Réka Albert, Hawoong Jeong, and Albert-László Barabási. The Internet's Achilles' Heel: Error and attack tolerance of complex networks. *Nature*, 406:378–482, Jul 2000.

[4] Jayanth R. Banavar, Amos Maritan, and Andrea Rinaldo. Size and form in efficient transportation networks. *Nature*, 399:130–132, May 1999.

[5] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286:509–512, Oct 1999.

[6] V. Batagelj and A. Mrvar. Pajek - program for large network analysis. *Connections*, 21(2):47–57, Spring 1998.

[7] S. M. Bellovin. Security problems in the TCP/IP protocol suite. *Computer Communications Review*, 19(2):32–48, Apr 1989.

[8] Béla Bollobás. *Random Graphs*. Cambridge University Press, Cambridge, UK, second edition, 2001.

[9] Anirban Chakrabarti and G. Manimaran. Internet infrastructure security: a taxonomy. *IEEE Network*, 16(6):13–21, Nov/Dec 2002.

[10] Fan R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, Providence, RI, 1991.

[11] Joe E. Cohen and Clark Jeffries. Congestion resulting from increased capacity in single-server queueing networks. *IEEE/ACM Transactions on Networking*, 2(5):305–310, Apr 1997.

[12] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. The MIT Press, Cambridge, Massachusetts, 1992.

[13] Lumeta Corporation. Internet mapping project. http://www.lumeta.com/mapping.

[14] B. Ellison, D. A. Fisher, R. C. Linger, H. F. Lipson, T. Longstaff, and N. R. Mead. Survivable network systems: An emerging discipline. Technical Report CMU/SEI-97-TR-013, Software Engineering Institute, Carnegie Mellon, 1997.

[15] Kevin Fall, Kannan Varadhan, and the VINT project. *The ns Manual*. The VINT Project. http://www.isi.edu/nsnam/ns/doc/ns_doc.pdf.

[16] Christos Gkantsidis, Milena Mihail, and Amin Saberi. Conductance and congestion in power law graphs. In *ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, volume 31, pages 148–159, San Diego, CA, Jun 2003.

[17] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for Internet map discovery. In *Proceedings of the Nineteenth Annual Joint*

Conference of the IEEE Computer and Communications Societies (IEEE INFOCOM 2000), volume 3, pages 1371–1380, Tel Aviv, Israel, Mar 2000. IEEE.

[18] Salim Hariri, Guangzhi Qu, Tushneem Dharmagadda, Modukuri Ramkishore, and Cauligi S. Raghavendra. Impact analysis of faults and attacks in large-scale networks. *IEEE Security & Privacy Magazine*, 1(5):49–54, Sep/Oct 2003.

[19] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Efficient security mechanisms for routing protocols. In *Proceedings of the Tenth Annual Network and Distributed System Security Symposium (NDSS 2003)*, San Diego, CA, Feb 2003.

[20] USC Information Sciences Institute. Internet maps. http://www.isi.edu/scan/mercator/maps.html.

[21] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, Nov 1999.

[22] G. Malkin. RIP version 2. RFC 2453, Nov 1998.

[23] J. Moy. OSPF version 2. RFC 2328, Apr 1998.

[24] S. Murphy, M. Badger, and B. Wellington. OSPF with digital signatures. RFC 2154, Jun 1997.

[25] P. Papadimitratos and Z. J. Haas. Securing the Internet routing infrastructure. *IEEE Communications Magazine*, 40(10):60–68, Oct 2002.

[26] Denise Pappalardo. Can one rogue switch buckle AT&T's network? *Network World Fusion*, Feb 2001. http://www.nwfusion.com/news/2004/0120slammoney.html.

[27] R Development Core Team. *R: A language and environment for statistical computing.* R Foundation for Statistical Computing, Vienna, Austria, 2004. http://www.R-project.org.

[28] Y. Rekhter and T. Li. A border gateway protocol 4 (BGP-4), Jul 1994.

[29] CNET News.com Staff. Router glitch cuts Net access. *CNET News.com*, Apr 1997. http://news.com.com/2100-1033-279235.html?legacy=cnet.

[30] William Stallings. Queueing analysis, 2000. http://www.williamstallings.com/StudentSupport.html.

[31] Michael Stutz. Net outage: The oops heard 'round the world. *Wired News*, Apr 1997. http://www.wired.com/news/technology/0,1282,3442,00.html.

[32] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, Jun 1998.

[33] David L. Wilson. The Internet vs. the bomb: Would the Internet survive the bomb? *CNN*. http://edition.cnn.com/SPECIALS/cold.war/experience/technol