2004-09-30

Network-layer Selective Security

Casey T. Deccio
Brigham Young University - Provo

NETWORK-LAYER SELECTIVE SECURITY

by

Casey T. Deccio

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

Department of Computer Science

Brigham Young University

September 2004

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Casey T. Deccio

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory.

Date	Mark J. Clement, Committee Chairman
Date	Quinn O. Snell, Committee Member
Date	Bryan S. Morse, Committee Member

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Casey T. Deccio in its final form and have found that (1) its format, citations and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

_____
Date

_____
Mark J. Clement
Committee Chairman

Accepted for the Department

_____
David Embley
Graduate Coordinator

Accepted for the College

_____
G. Rex Bryce
Associate Dean,
College of Physical and Mathematical Sciences

ABSTRACT

NETWORK-LAYER SELECTIVE SECURITY

Casey T. Deccio

Department of Computer Science

Master of Science

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 thesis 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.

ACKNOWLEDGMENTS

Many thanks goes out to those individuals have helped make the completion of this thesis possible. I thank my advisor, Dr. Mark Clement, who helped me to finally narrow in on a relevant and challenging topic, and also helped me to develop critical writing skills throughout the duration of research. Spencer Cox developed a tool used in this research for the analysis of large networks under attack. Jeremy Goold helped me organize my thoughts on numerous occasions, when I had obviously spent too many hours in the research lab. Occasional ping-pong matches with Daniel Ledesma during research breaks allowed my mind time away to refocus my thoughts. I am especially grateful to my loving wife, Talia, who was patient and understanding during the difficult months that I spent working on this research. It was always wonderful to be greeted by her and my beautiful daughter, Zion, after a day spent at work. Finally, I am grateful to a merciful Father in Heaven, who gave me the faith, capacity, and endurance necessary to complete this thesis.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The Internet is the foundation of world-wide digital communication. Universities, businesses, and government agencies rely on the Internet for common and vital tasks. 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" [10]. The increased availability of tools allowing direct access of network resources to malicious users has made these attacks a reality [31]. Specific attacks targeting the routing infrastructure include routing table poisoning attacks, packet mistreatment attacks, and denial-of-service (DoS) attacks [12].

As technology advances, more critical applications are being handled by electronic systems, creating a target for cyber-terrorists [5]. A recent network outage spanning more than three days at Beth Israel Deaconess Medical Center exemplifies the vitality of stable computer networks [9]. Infections on desktop computers in the research wing slowed traffic throughout the entire network. The systems used by doctors and nurses to track patients were interrupted, causing the staff to revert to

paper methods while the problems were solved. Had the incident been more severe, however, patients' lives could have been endangered by the outage. Since the incident, the hospital has heightened network security by decrypting and inspecting all virtual private network (VPN) traffic before passing it through the hospital's firewall. [36].

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 [23].

Rather than universal application of arbitrary security at all points within a given network, *critical elements* should be identified in the infrastructure—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.

Thesis Statement

Critical elements can be identified in a computer network, whose selective security will make the network more survivable when under attack. This research defines network survivability in terms of the network routing infrastructure and identifies critical elements in network topologies, 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 thesis will provide a basis upon which a future selective security model might be designed and implemented. This research will not, however, directly outline a model for selective security.

Early chapters discuss previous research involving threats and security at the network layer and theoretical studies that provide necessary background for this topic. Contributions, including methods for identifying critical network elements based on network performance, are found in later chapters. Chapter 2 summarizes threats of specific network-layer attacks and introduces the notion of selective security.

A case study of routing mechanisms and applied security is found in Chapter 3. Chapter 4 discusses notable characteristics of different network topologies. Network survivability and metrics for measuring network performance are found in Chapter 5. Methods for identifying critical network elements and an analysis of simulation results are discussed in Chapter 6. The conclusions of this work are outlined in Chapter 7.

Chapter 2

Network-layer Attacks

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. When successful protocol attacks are executed at this level, 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 [37]. 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 [39].

Although the cause of the above 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. This chapter discusses several attacks against network-layer protocols [12] and introduces the notion of custom security.

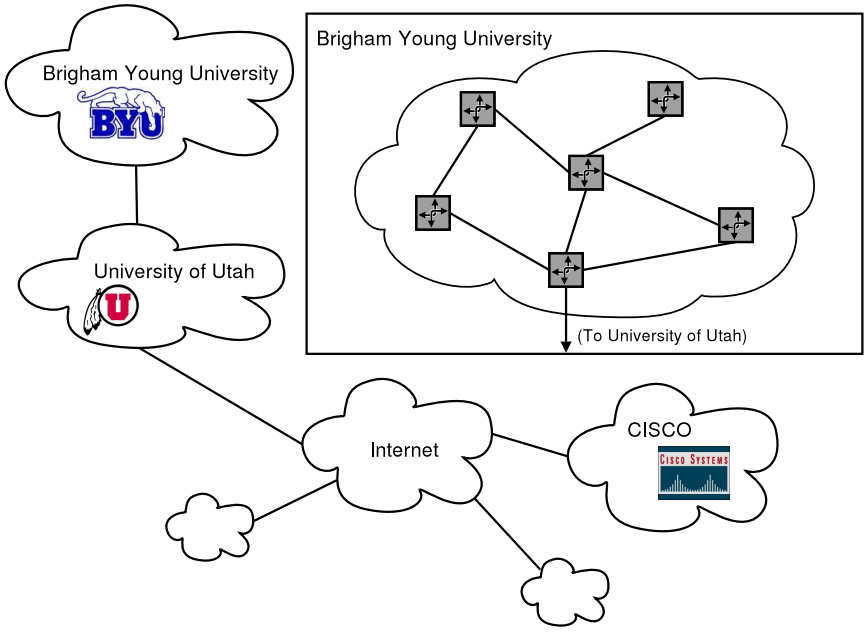

Figure 2.1: An example network hierarchy. Brigham Young University, the University of Utah, and CISCO use interior routing protocols to establish routes within their respective networks, but use exterior routing protocols to send traffic between themselves and other Internet hosts.

2.1 Protocol Attacks

The routing infrastructure maintains data 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 [28] or the Routing Information Protocol (RIP) [27]. The Internet is a complex network of AS that communicate using an *exterior routing protocol*, such as the Border Gateway Protocol (BGP) [35]. An example network hierarchy is shown in Figure 2.1. The AS (shown as clouds) transfer data within their organizations using an interior routing protocol. Data is transferred between AS using an exterior routing protocol.

Attacks at the routing level can fragment a network graph or place a large hole in the topology, crippling network performance. Less noticeable but equally compromising attacks could result in the exposure or manipulation of sensitive or classified material. Several routing attacks are described in this chapter.

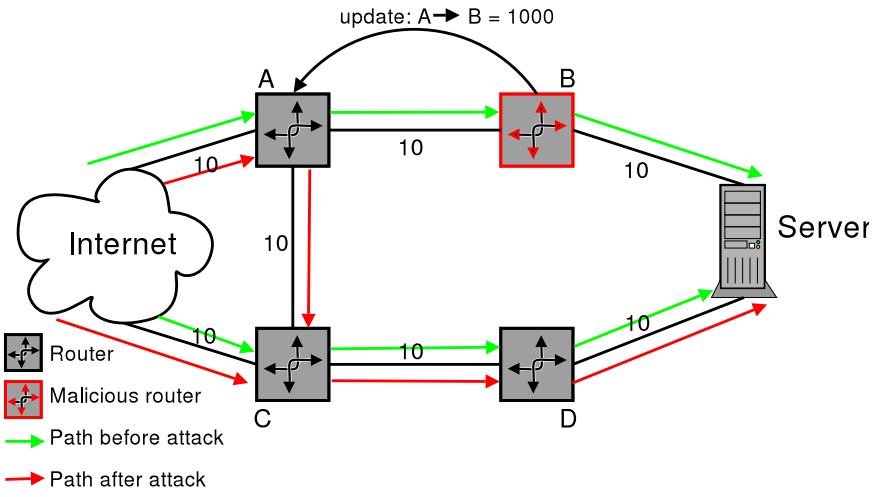

Figure 2.2: A routing protocol attack, caused by routing table poisoning. Router B advertises a $A \rightarrow B$ metric of 1000 to router A (much larger than the actual metric), and diverts traffic destined to the server through router C.

2.1.1 Routing Table Poisoning

The routing table is the data structure in a router which contains the information received by other routers, describing the topology of the network. This information is the basis of determining where to forward incoming data. *Poisoning* is caused by inaccurately updating the routing table of a functioning router. In Figure 2.1.1 router B sends an update to A claiming that the distance metric from A to B is 1000. The affected router, A, may even pass the poisoned information on to other routers, infecting them as well, until an accurate update remedies the situation. In accordance with shortest path routing, router A modifies its behavior to route traffic destined for the server through router C—a "shorter" path.

2.1.2 Denial of Service

DoS attacks seek to disable a particular network element by overwhelming it with large amounts of data, effectively overloading its resources. At the network layer, DoS attacks can overwhelm and disable vital network routers. This action would leave a hole in a network, or completely divide the network, until the affected routers are able to recover. One method for targeting a router for a DoS attack is accomplished by poisoning the routing table of selected routers, deliberately diverting traffic flows through a particular router, whose finite resources would quickly be consumed with

large quantities of data. Router C in Figure 2.1.1 might feel this effect. Another way to disable a router for a time would be to forward a large number of protocol packets to the router, from a single (DoS) or multiple (i.e., *distributed DoS* or DDoS attack) sources.

2.1.3 Man-in-the-middle Attack

In a *man-in-the-middle* attack, data is forwarded through an intruding or compromised router [41]. This malicious router then has the ability to eavesdrop for sensitive information, alter through data, or drop select data packets. The redirection of traffic through such a router might be the result of disabling links or nodes (e.g., through DoS attacks) in the network or colluding routers that poison routing tables in order to control traffic flow.

If routers B and C from Figure 2.1.1 are colluding, then the bogus update from B results in a man-in-the-middle attack, wherein sensitive information now passes through (the compromised) router C.

2.2 Selective Security

The routing infrastructure discussed in Section 2.1 is critical to the functionality of network processes across the Internet. To protect network-layer protocols from attacks, such as those mentioned previously, research has produced security mechanisms (some of these are discussed in more detail in Chapter 3). 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 [29]. Although the draft for this was written in 1997, the status of the protocol is still "experimental".

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 discussed in Section 2.1 exhibits a hierarchical characteristic, which means that not all elements have the same importance respective to the survivability of the overall infrastructure. Likewise, certain behaviors of network topologies place a higher reliance on particular elements in order to maintain survivability (as explained in subsequent chapters). 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 [29]. A scheme with more variance from the original specification may have more trouble getting approved 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 thesis does not directly discuss implementation of selective security, but rather helps identify critical points in the network that should be secured against attack.

Chapter 3

Network-layer Security: A Case Study

This chapter summarizes some of the efforts that have been applied towards securing the network layer. Some of the seminal work in securing network-layer routing was published in the 1988 Ph.D. thesis of Radia Perlman [33]. Her work describes different levels of network robustness and outlines a public key system for identifying trusted nodes in a link-state routing network. Perlman's thesis introduces the terms *simple failure*—a node or link becomes inoperative—and *Byzantine failure*—nodes or links continue to operate, but incorrectly. By the same notion *simple robustness* and *Byzantine robustness* are exhibited by networks that continue correct operation in spite of simple and Byzantine failures, respectively.

Perlman's thesis is a foundation for routing protocols and associated security that have been developed. As a case study, this chapter discusses one of those protocols—OSPF—and some vulnerabilities and securities of the protocol. OSPF is a widely-used *link-state* routing protocol, so categorized because each participating router is responsible for describing the state of its local neighborhood (i.e., metrics to neighboring networks, routers, and hosts) to all other nodes in the network. The protocol is "open", meaning that its specification is in the public domain.

3.1 OSPF Route Establishment

In the OSPF protocol, each router is responsible for maintaining an identical database which describes the topology of the AS, so it can determine the shortest routes for

packets to take to arrive at any particular destination. This *link-state database* is composed of *Link State Advertisements* (LSAs) received from other nodes in the AS, each containing information about the local neighborhood of the source node. To determine the shortest path to a destination, all routers run a shortest path algorithm (i.e., Dijkstra's Algorithm [15]) in parallel. Each uses its link-state database to construct a tree of shortest paths to all other destinations, with itself as the root. If there exist several equal-cost paths to a particular destination, then the traffic is distributed evenly among the routes [28].

3.1.1 Adjacencies

To build a description of the network topology each router must first establish *adjacencies* with its immediate neighbors. Routers first discover their neighbors by sending and receiving *Hello packets*. Using the *Hello protocol*, Hello packets are sent out periodically on each network interface by each router. Once neighboring routers have "met" via the Hello Protocol, the routers undergo a database exchange process to enforce database synchronization. This synchronization is repeated periodically to maintain data integrity.

3.1.2 Flooding

Once adjacencies have been established between neighboring nodes, a router organizes the information about its local neighborhood into a LSA to distribute to all other routers. The LSA reaches all other routers through reliable, intelligent *flooding*. This process is as follows: when a router receives an advertisement from a neighbor, it acknowledges receipt of the advertisement and, if the advertisement is new, forwards the advertisement to all other neighbors. Thus, all routers will have an identical topological database of LSAs from which they can derive shortest paths to all destinations.

3.1.3 OSPF Hierarchy

Routers can be grouped together in *areas* within an AS. Each area runs in parallel a copy of the basic link-state protocol. OSPF defines a two-level hierarchy among all areas in the AS: the top level is the *backbone*, and the next level consists of many areas connected to the backbone. A router which is part of two areas within the AS is

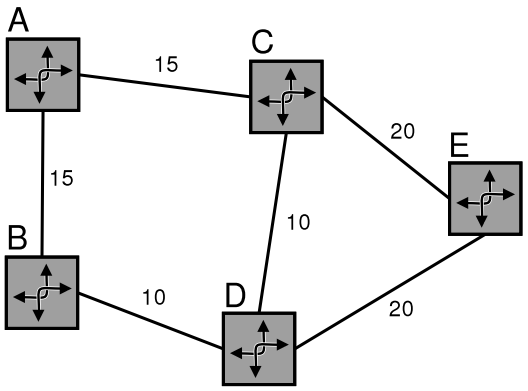

Figure 3.1: An area in an OSPF network.

known as an *Area Border Router* (ABR) and belongs to the backbone. Routers that are connected to points outside the AS are known as *Autonomous System Boundary Routers* (ASBRs). Routers within an area can determine the shortest path to a router outside the area using information from the ABRs. Likewise, routers within the AS can determine the shortest path to routers outside the AS via information from the ASBRs.

The remainder of this section will only discuss the basic OSPF protocol that is performed within a single area.

3.1.4 LSAs

LSAs are the building block for link-state databases. Each LSA describes the local neighborhood of the router from which it originated. This includes the nodes that are directly connected to the originating router and the metrics to each neighboring node. There several types of LSAs. Only the Type 1 LSA, in which a router sends a LSA to all the routers in its area to advertise its adjacencies, is related to this research. The LSA data from router C in Figure 3.1.3 would carry the LSA data shown in Table 3.1.4.

In addition to carrying the state of its neighboring nodes, all LSAs carry a header that contains identifying information for that LSA. This includes its age, special options, the type of LSA, the ID of the originating router, a sequence number, and a checksum. As a LSA is flooded, its age is incremented according to the time it has been circulating. This process is also performed by each router, as the LSA resides in its link-state database. When the age of a LSA reaches `MaxAge` (defined

		From				
		A	*B*	*C*	*D*	*E*
	A			10		
	B					
To	*C*					
	D			15		
	E			20		

Table 3.1: The Type 1 LSA data for router C in Figure 3.1.3.

in the OSPF specification), the router in which it resides flushes the LSA from its database and floods the LSA, so other routers flush the outdated LSA as well.

3.2 OSPF Vulnerabilities

The availability of the OSPF specification has allowed for review in the public domain. In recent research protocol vulnerabilities have surfaced, whose exploitation could threaten network security. Wang, et al. have categorized OSPF attacks as either insider attacks or outsider attacks [40]. *Outsider* attacks involve non-protocol participants (e.g., an attacker with access to links between routers), and *insider* attacks refer to valid routers that have been compromised or are showing Byzantine failures.

Attacks specific to the OSPF protocol include manipulation of LSAs of transit. Modifications of header fields— including the serial number and age—or the link-state data itself can violate integrity of the OSPF protocol.

3.3 OSPF Security

The OSPF Version 2 specification [28] outlines native security features for OSPF. These native features include the option for *authentication*, thus preventing an unauthorized router from posing as a valid router on the network. Every OSPF packet specifies an authentication type and designates a 64-bit field to be used for the authentication.

The simplest authentication technique is the use of a 64-bit password [28]. However, any intruder with access to the network could trivially obtain a password sent in the clear using a *passive attack* (i.e., learning the password by observing

network traffic) and thus authenticate itself.

Another authentication technique uses a *cryptographic* solution [28]. A shared secret key is configured for all routers within a subnet or network. For each OSPF packet, the key is used to generate and verify a *message digest* which is appended to the end of the packet. The digest is a one-way function of the secret key and the packet itself. The secret key is protected from passive attacks because it is never sent in the clear. Also, the increasing sequence number in each LSA prevents an attacker from replaying the packet and its corresponding digest.

3.4 OSPF with Digital Signatures

OSPF with Digital Signatures is not part of the OSPF specification; it is a proposed specification with experimental status at the time of this writing [30, 29]. This schema, outlined by Murphy, et al., uses a *public key infrastructure* in which each *authenticated router* in an internal network maintains a public and private key. When an authenticated router sends routing protocol packets, it signs the data using its private key—a signature that can only be verified using the corresponding public key. All authenticated routers in the AS maintain a database with the public keys of all other authenticated routers. Using the public key of the LSA's owner, any router receiving the LSA can verify the LSA's integrity. Also, a router generating faulty information can be identified by its signature on the faulty OSPF packet.

3.4.1 Authenticated Routers and LSAs

Authenticated OSPF routers are capable of performing all the same functions as standard OSPF routers. Additionally they generate signed routing information LSAs (*authenticated LSAs*), send *new key information LSAs*, manage key and signature algorithm verification, and verify signatures received. An authenticated LSA contains the same information as LSAs (Section 3.1.4). In addition it contains a signature of the LSA (all except the age field—so routers can increment the age while it is in transit), the router key used to sign the LSA, the ID of the trusted entity that produced the certificate, and the length of the signature.

In order to guard against a *premature aging* attack, the signature of an authenticated LSA includes the age if and only if the age is `MaxAge`. This prevents a malicious router from signaling other routers to flush another router's LSA from

their link-state databases.

3.4.2 Key Management and Distribution

Using the OSPF PKI, a special device acts as a *trusted entity* (TE) for authenticated routers. This TE maintains its own public and private key pair: its public key is known by all authenticated routers, and its private key is kept secret. In order for a router to be authenticated in this AS, the router must flood the network with a *Router Public Key LSA*, which has the following components: a regular LSA header (see Section 3.1.4) with LSA type 16; signature information, which will allow routers to interpret the router's signature; a certificate signed by the TE, which verifies the authenticity of the LSA header information; the signature of the trusted entity; and the router's signature of the LSA, excluding the age. Any authenticated router receiving this Router Public Key LSA will add the new authenticated router's id and public key to its database. Because the addition of new routers to a network is generally planned by network administrators, the necessary configurations for a new router to join an authenticated network may be performed offline, thus avoiding the risk of the trusted entity being compromised.

3.5 Security Analysis

The security measures in Sections 3.3 and 3.4 are effective in protecting a communications network against damaging protocol attacks. Message digests of OSPF protocol packets protect the integrity of protocol exchanges. Because the secret key used is undetectable and irreproducible by intruders, this security protects against outsider attacks. For protection against insider attacks, the OSPF with Digital Signatures allows routers to identify the source of malicious nodes or Byzantine failures.

While both of the above measures perform well in the security aspect, deploying these methods universally could affect other network performance. Overhead of any type at the network level should be minimized, so the protocols at that level are seemingly transparent. However, the RSA cryptography that is typically used in digital signature schemes is computationally expensive—about 1000 times slower that symmetric authentication schemes—even with low RSA exponents [40]. Further, OSPF with Digital Signatures adds protocol complexity to an already complicated OSPF protocol.

Besides network performance concerns, there are political logistics involved with a stable network running without security. Any security added may require additional hardware and some down time. Also, the effort must be coordinated and supported equally by all parties involved.

The security developed for OSPF, and other protocols, may be more beneficial and readily deployable if a selective security scheme is adopted, rather than a solely universal scheme. This research identifies critical areas of network topologies where the cost of additional security is warranted.

Chapter 4

Network Topologies

Communications networks respond differently to applied instances of attack or failure. The functional loss of particular nodes or links within a network affects the stability of the entire graph. Some networks are greatly affected by the incidental failure of random nodes. Others exhibit fault tolerance but experience degraded performance in the presence of malicious attacks. This chapter examines some of the characteristics of the Internet and other large networks in order to identify critical elements within them.

The identification of critical nodes in a communications network requires an understanding of natural topological characteristics of different networks. Characteristics of scale-free and random networks are discussed in this research.

4.1 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 [6, 42]. 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 [7, 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: $P(k) \sim k^{-\gamma}$ [7]. 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 connectivities. Figure 4.1 shows the graph (a) and connectivity distribution (b) of a scale-free network generated and visualized using the Pajek Program for Large Network Analysis [8]. 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 4.1b.

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 [21] merged with data also obtained in 1999 from the Internet Mapping Project at Lucent Bell Laboratories [24][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 4.1. 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 amid even unrealistically high node failure rates. However, it also exhibits a vulnerability to malicious attacks, in which the highest-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 6.1 discusses the identification of critical elements in scale-free networks, and explains results of simulations involving attack and survivability of scale-free networks.

4.2 Random Networks

Relatively smaller networks, such as autonomous systems (AS) managed by a sole entity are often classified as *random networks*. There are many models describing random networks. The simplest model is denoted $G(n, M)$, $0 \leq M \leq \binom{n}{2}$, and defines n isolated nodes $1, 2, \ldots, n$, after which M edges are selected at random without replacement and added to the graph [11] [11]. Erdős and Rényi define a similar

[1]The Internet Mapping Project is now run by Lumeta Corporation. More information can be found at their Web site [16].

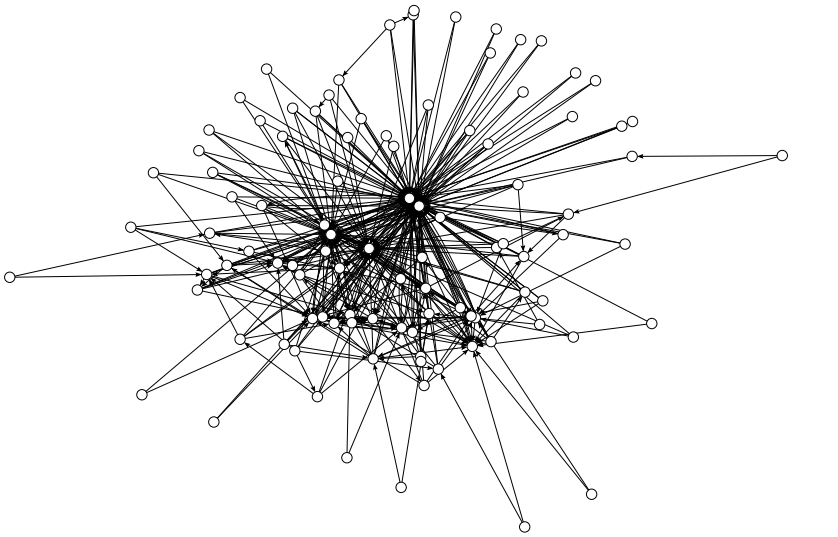

(a)

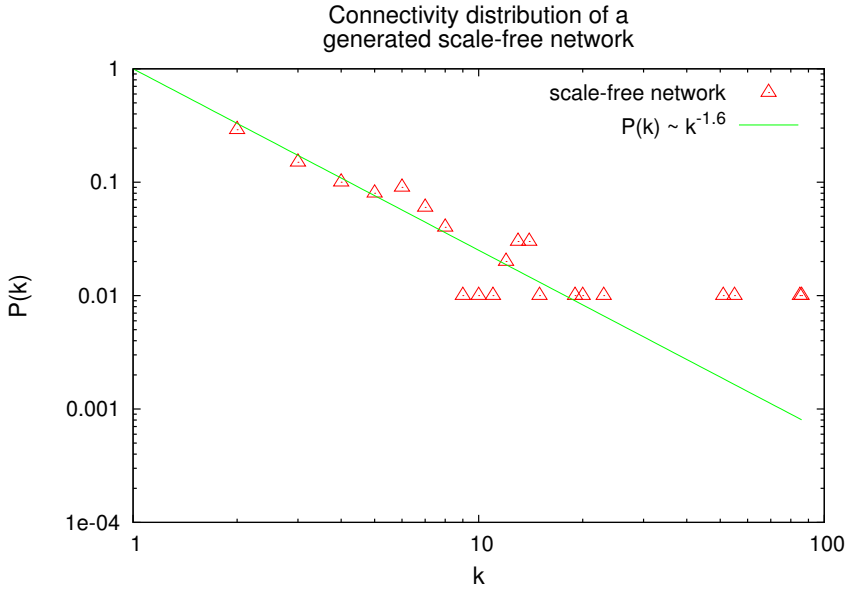

(b)

Figure 4.1: 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 [8] and follows the model $P(k) \sim k^{-1.6}$, which is plotted against its connectivity distribution.

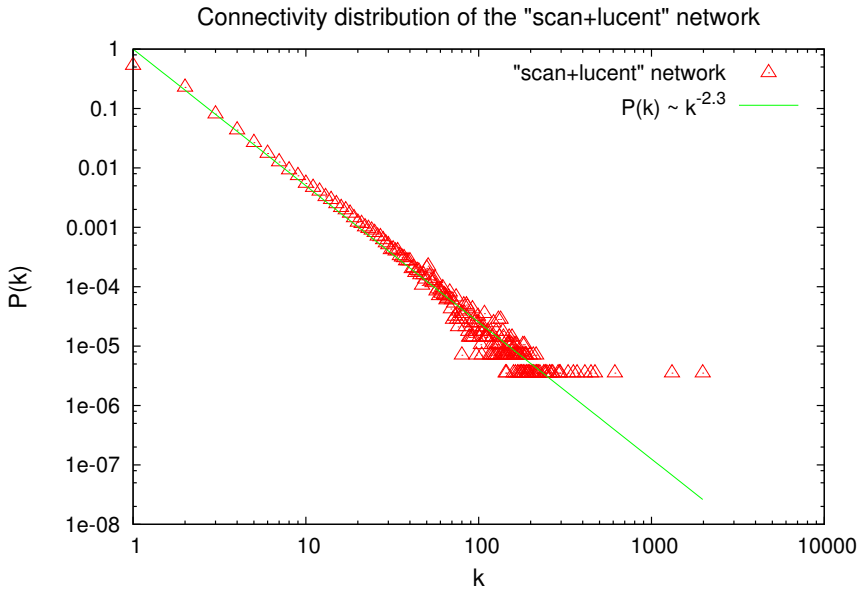

Figure 4.2: The connectivity distribution of the "scan+lucent" network. The line below the plot points represents the connectivity distribution for the scale-free model $P(k) \sim k^{-2.3}$.

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. At $p(n) = \frac{1}{2}$ any graph with n nodes is equiprobable. Figure 4.2 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 [8]. This network follows the Erdős-Rényi model [11] 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. This property means that 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, directed attacks at the highest-connected network nodes will not harm the network more than random node failures.

In order to show realistic results of attack and survivability, the topology of a major telecommunications vendor's frame relay network is used as a real-world graph for analysis. The connectivity distribution of the "vendor" network, comprised

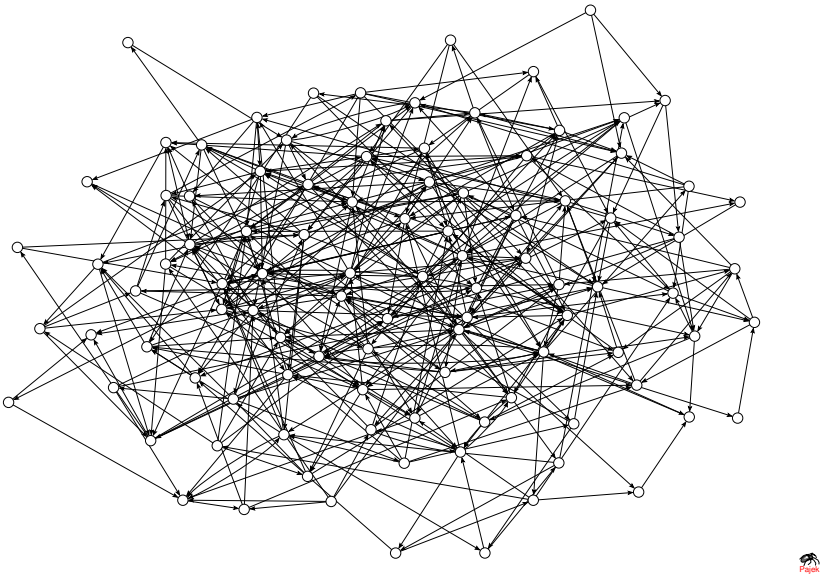

(a)

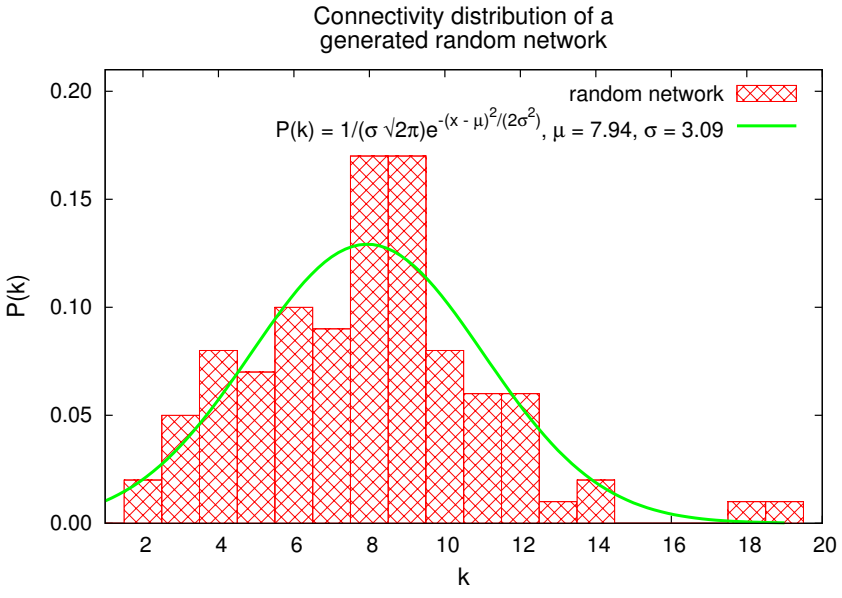

(b)

Figure 4.3: 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 [8] and approximates a normal distribution with $\mu = 7.94$ and $\sigma = 3.09$.

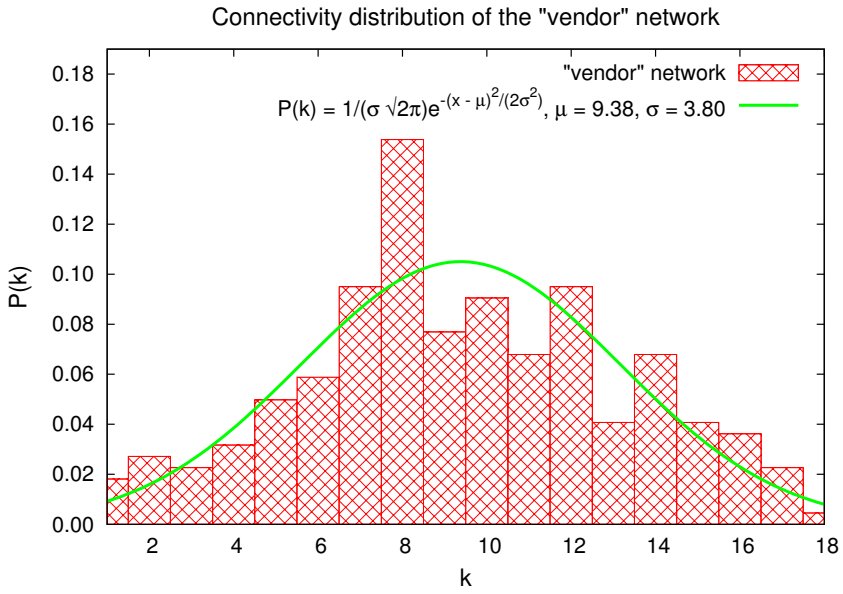

Figure 4.4: The connectivity distribution of the "vendor" network, plotted with the curve of a normal distribution with mean $\mu = 9.38$ and standard deviation $\sigma = 3.5$.

of 221 nodes, is plotted in Figure 4.2, along with the plot of a normal distribution with mean $\mu = 9.38$ and standard deviation $\sigma = 3.5$. The results of simulations involving attack and survivability, as well as methods for identifying critical elements in random networks are discussed in Section 6.2.

Chapter 5

Measuring Network Performance and Survivability

Network survivability is a hot topic in recent research. Defining survivability is a crucial step in identifying critical network elements. This chapter defines survivability and outlines several metrics for quantifying network performance in a communications network.

5.1 Network Survivability

With the increasing dependence on large, distributed architectures for operations on a daily basis, the need for critical systems analysis is invaluable. A *survivable* system is able to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents [17]. 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 [32].

Online systems at Beth Israel Deaconess Medical Center handle 40 terabytes of information daily. Doctors and nurses use these systems to track patient

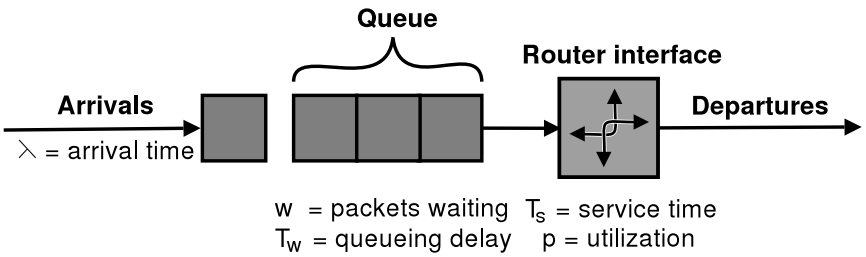

Figure 5.1: A queuing model for a network router interface. The parameters λ, T_w, T_s, and ρ affect network performance.

status and access clinical records, including problem lists, medications, allergies, and notes. In 2002 infections on desktop computers in the research wing of the hospital caused network slowdowns and interruptions throughout the entire network for more than three days [9]. The slowdowns caused the staff to revert to paper methods during this period, degrading the efficiency of their work.

Even the Internet—which urban legends claim could withstand the effects of a nuclear bomb [43]—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 research presented in this thesis is intended to make large networks more survivable against protocol-based attacks. Survivability is the property of an entire network, not simply that of one of its comprising nodes [25]. 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 [17]. 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 [22]. Topological characteristics of the networks are also analyzed in correlation with these metrics.

5.2 Queuing Model

An analysis of performance at the router level can help quantify the performance of the comprising network. A network router can be analyzed as a single-server queue by using a theoretical model (see Figure 5.2) [38]. Data packets arrive at the router

at rate λ and must be serviced with average *service time* T_s. Often it is required that arriving packets "wait" in line (queue) to be serviced behind other packets that arrived first. 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 [38]:

$$\rho = \lambda T_s \tag{5.1}$$

When $\rho = 1.0$, the interface is saturated. Thus, the theoretical maximum input rate that can be handled by a router is [38]:

$$\lambda_{\max} = \frac{1}{T_s} \tag{5.2}$$

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 [38]:

$$T_w = \frac{w}{\lambda} = \frac{wT_s}{\rho} \tag{5.3}$$

As a link approaches capacity (i.e., $\rho \to 1.0$), delay becomes arbitrarily high [14]. Network performance is higher if link utilization and queuing delay are minimized.

When queued data exceeds the buffer size of a router, the router will drop packets based on some predefined policy (e.g., drop the last packet that arrived). The *loss probability* L at a queue is the percentage of packets dropped by a particular router interface over a certain period of time. It is desirable for the loss probability to be as close to zero as possible.

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 Chapter 6 to evaluate network performance after networks have been targeted for attack. This will be a

measure for how much elements affect overall network survivability.

5.3 Topological Characteristics

Topological characteristics can be used to indicate some measure of performance of the network. Several metrics have been derived for describing 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 [13]. 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 [20].

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 5.3), and the network's ability to accomplish its mission of successful data transfer 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 D or $\langle d \rangle$ incurred will be used in Chapter 6 to quantify network survivability when arbitrary network elements are protected against network-layer attacks.

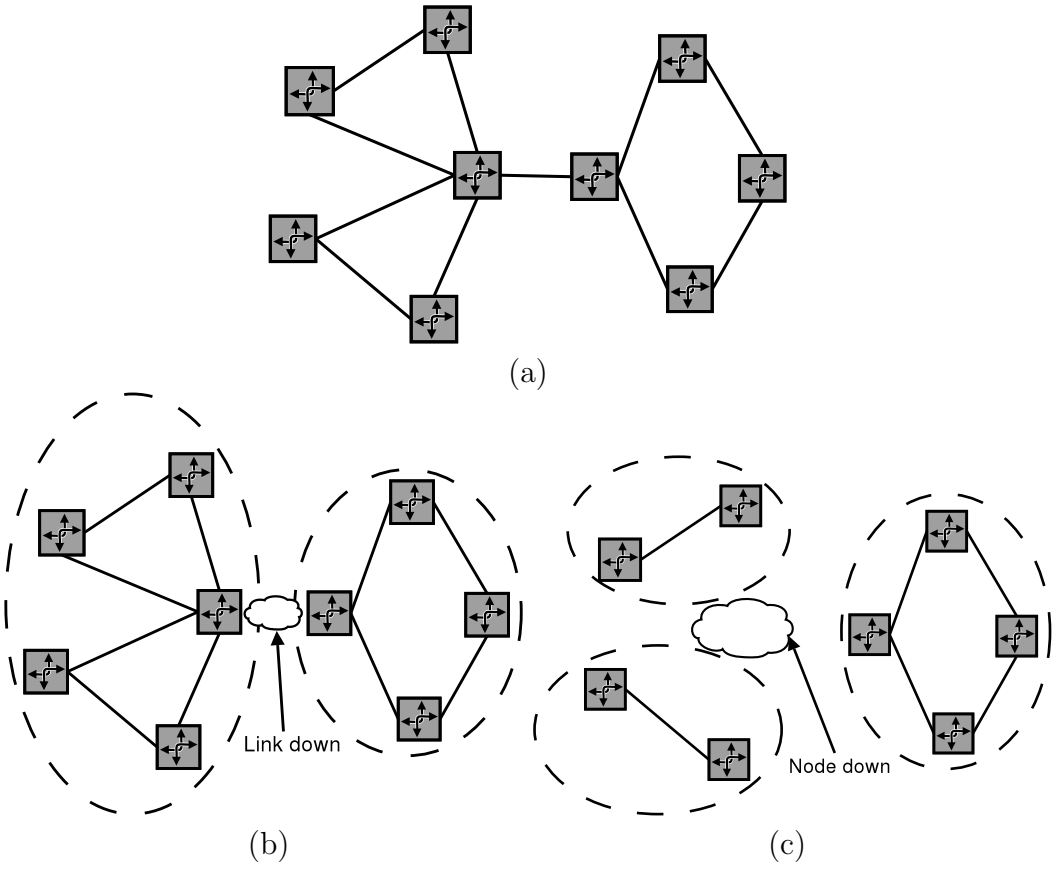

(a) (b) (c)

Figure 5.2: 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.

Chapter 6

Identifying Critical Network Elements

The contribution of this thesis is a resource allowing network analysts to identify critical elements in a communications network, whose failure is detrimental to the survivability of the entire network infrastructure. This chapter presents metrics for quantifying the value of a network element within the network. The results are based on previous research as well as empirical data gathered from network simulation.

Different graph models, such as those discussed in Chapter 4, present different challenges for measuring survivability, using the metrics introduced in Chapter 5. 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.

6.1 Critical Elements of Scale-free Networks

Scale-free networks, as discussed in Section 4.1, are distinguished by their 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.

6.1.1 Simulation Environment

To produce results supporting the hypothesis that the most connected nodes are more critical to the survivability of a scale-free network, a software tool was created for simulating network attacks and simple failures (the loss of random nodes) and analysis of the resulting network. The simulator was written in the C++ programming language and does not utilize a scheduler for time-based and traffic performance, but rather only analyzes topological characteristics. 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 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 run once to simulate network attacks and once to simulate simple failures. The former involved disabling the 200 highest connected nodes in the network at each iteration; the latter involved disabling 200 randomly-selected network 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.

6.1.2 Simulation Analysis

In the attack model, wherein network nodes were disabled in order of decreasing connectivity, the fragmentation in the network severely crippled its ability to function (see Figure 6.1.2). When only 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 6.1.2 also shows the result of losing large numbers of nodes due to simple failures. Although 1% of the nodes were directly attacked, the network retained over 97% of its original nodes.

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

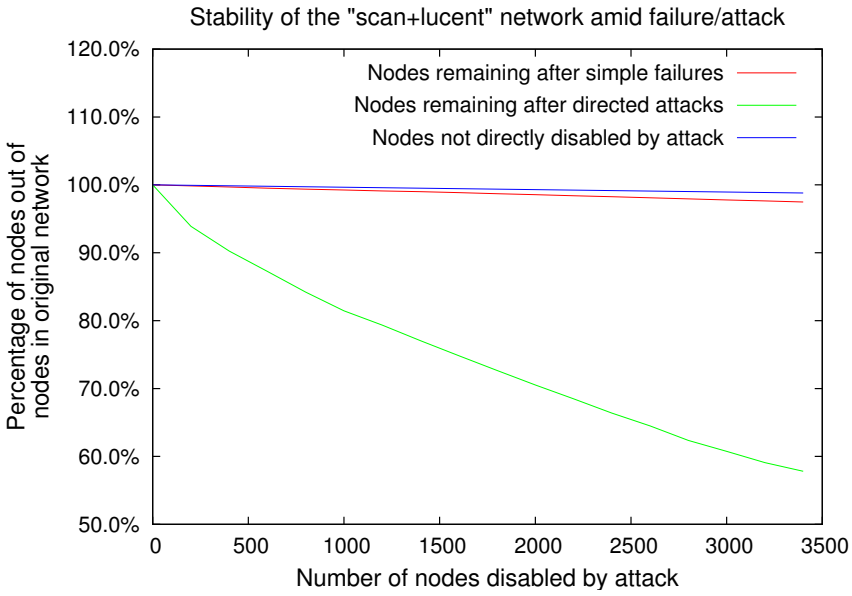

Figure 6.1: Network stability of the "scan+lucent" network under different attack strategies. When over 1% of the nodes are randomly removed from the network (simple failure), the network retains over 97% of its original nodes, but when 1% of the most connected nodes are removed from the network, the size of the functional network drops below 60% of its original size.

nodes was calculated:

$$\langle d \rangle = \frac{\sum_{u \in S} \sum_{v \in G | v \neq u} d(u, v)}{|S|(|G| - 1)} \tag{6.1}$$

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:

$$\langle d^* \rangle = \frac{\langle d \rangle}{|G|(|G| - 1)} c \tag{6.2}$$

The simulation results with $c = 10^{10}$ are shown in Figure 6.1.2 with a 95% confidence interval. The value $\langle d^* \rangle$ increased linearly when the nodes were disabled in

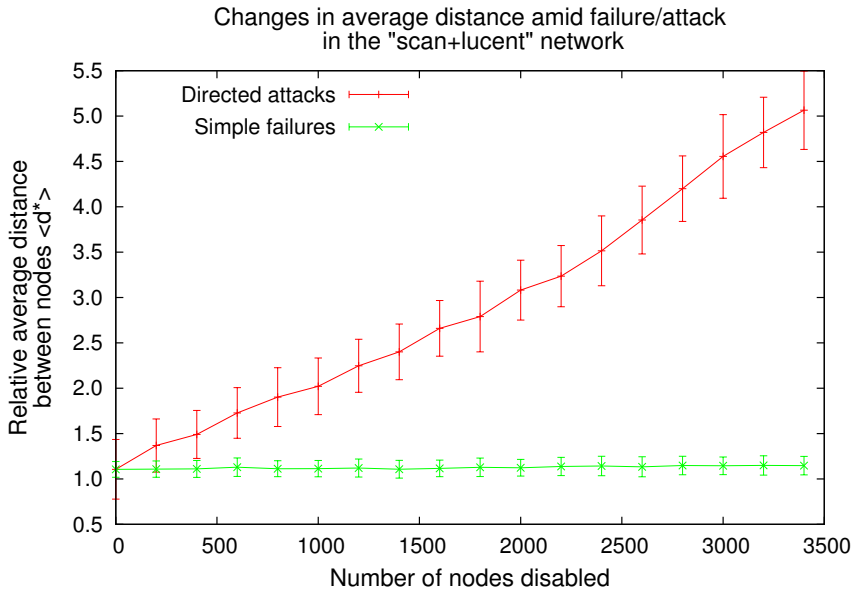

Figure 6.2: Changes in the relative average distance $\langle d^* \rangle$ between nodes of the "scan+lucent" network in the presence of different attack strategies, shown with 95% confidence intervals. The value $\langle d^* \rangle$ remains nearly unchanged even after 1% of the nodes are randomly removed from the network. However, $\langle d^* \rangle$ increases linearly as nodes are removed in order of decreasing connectivity.

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. Maintaining lower $\langle d \rangle$ values will help minimize delay and associated network congestion and data loss.

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 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.

6.2 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 6.1 do not necessarily apply to random networks. Simulations in this section explore different methods

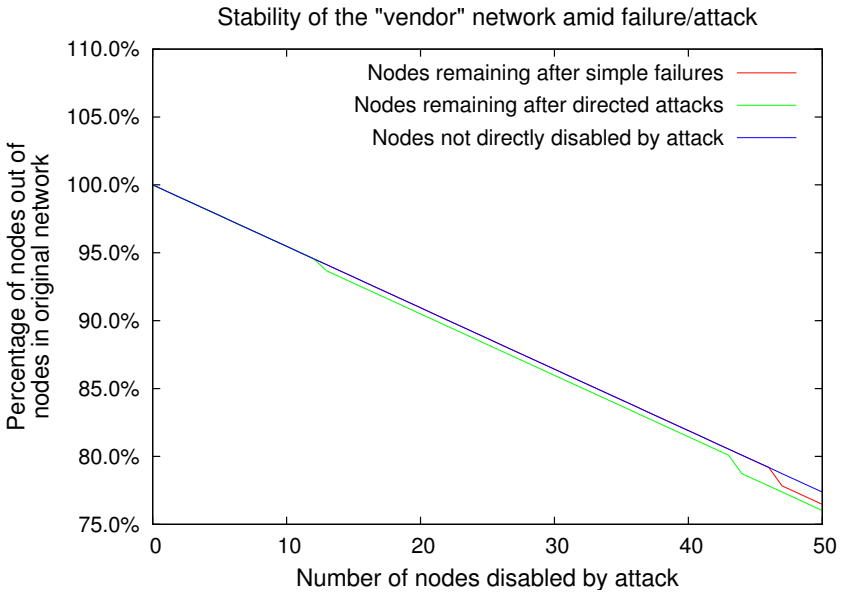

Figure 6.3: Network stability of the "vendor" network under different attack strategies. Nodes removed by both random selection and in order of decreasing connectivity exhibit similar characteristics. When nodes are removed by attack or by simple failure at random nodes, very few other nodes are made inaccessible.

for identifying critical elements in random networks.

6.2.1 Simulation of Node Attacks

Using the simulator described in Section 6.1.1, experiments similar to those performed on the "scan+lucent" network were run with random network data. Attack and failure scenarios were simulated on the "vendor" network by disabling one node per iteration at random or by decreasing connectivity, respectively.

Results from these simulations showed that network fragmentation was almost identical in the case of both attack and failure scenarios, as shown in Figure 6.2.1. This is due to the normal distribution of connectivities in random networks. Such homogeneity makes almost all nodes equal in importance with respect to the survivability of the network graph.

Because of the smaller size of the "vendor" network, the average distance $\langle d \rangle$ did not need to be approximated using a random sample, as in Equation 6.1; rather, it was calculated:

$$\langle d \rangle = \frac{\sum_{u \in G} \sum_{v \in G | v \neq u} d(u, v)}{|G|(|G| - 1)} \tag{6.3}$$

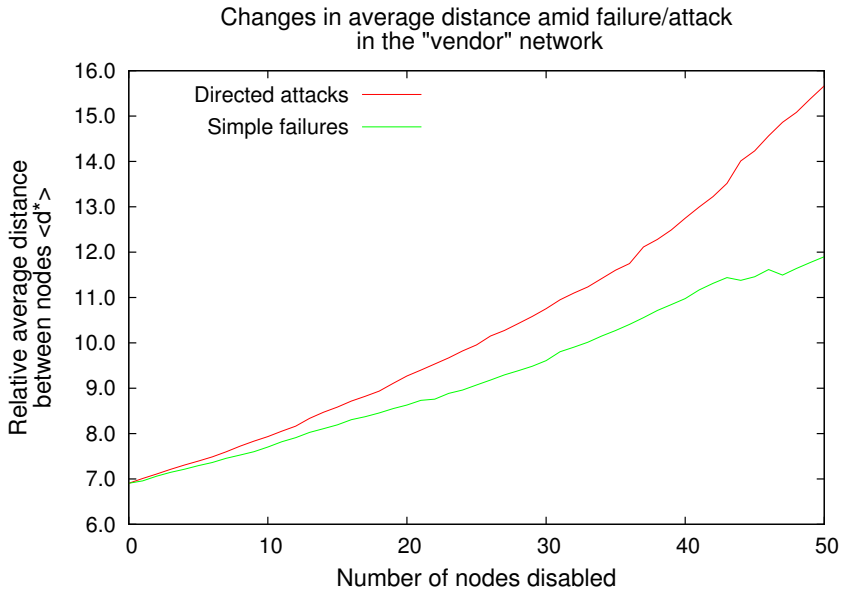

Figure 6.4: Changes in the relative average distance $\langle d^* \rangle$ of the "vendor" topology in the presence of different attack strategies. The value $\langle d^* \rangle$ increases nearly linearly both when nodes are removed in decreasing connectivity and when nodes are removed at random, although the rate of increase is greater for directed attacks.

Using this value for $\langle d \rangle$, relative average distance $\langle d^* \rangle$ was calculated using Equation 6.2 with $c = 10^5$. Figure 6.2.1 shows the changes in $\langle d^* \rangle$ after each iteration of attack or failure. The results show little difference between attack tolerance and error tolerance. In both cases, $\langle d^* \rangle$ rises approximately linearly, though in the case of directed attack, the slope is slightly higher.

The homogeneity of random networks necessitates a different approach for identifying critical elements to increase network survivability. Rather than using the connectivities of network nodes, an analysis of vital network links in random networks is performed in subsequent sections.

While the following mechanisms are outlined for identifying critical links in random networks, they will not be investigated for scale-free networks in this research. One reason for this is that critical links of scale-free networks have already been (indirectly) identified as the set of those linking the most connected nodes to their neighbors: by attacking this set of links, the most connected nodes are incapacitated. Additionally, the techniques used for simulation performed in this research allowed only smaller networks to be examined in reasonable time.

6.2.2 Max-flow Min-cut

This section summarizes some of the studies that have been performed with regard to data flow and network cuts, in order to identify network links critical to a survivable network. 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 [15]. Data flows in the network are referred to as *commodities*, and each commodity has a demand $D(s, t)$ [26].

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 $\bar{U}$:

$$C(U, \bar{U}) = \sum_{\{u,v\}|u\in U\wedge v\in\bar{U}\vee v\in U\wedge u\in\bar{U}} C(u, v) \tag{6.4}$$

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

$$D(U, \bar{U}) = \sum_{\{s,t\}|s\in U\wedge t\in\bar{U}\vee t\in U\wedge s\in\bar{U}} D(s, t) \tag{6.5}$$

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 [15, 26]:

$$D(U, \bar{U}) = |U||\bar{U}| \tag{6.6}$$

This thesis will use uniform multicommodity flow problems 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 [26]:

$$\theta(G) = \min_{U\subseteq V} \frac{C(U, \bar{U})}{D(U, \bar{U})} \tag{6.7}$$

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.

6.2.3 Network Bisection

Network *bisection* will only be discussed briefly here. It is a problem similar to the min-cut problem, but with cardinality constraints on the resulting sets—the number of nodes in each set must be equal: $|U| = |\bar{U}| = \frac{|G|}{2}$ [19]. An attacker who bisects a network successfully creates the largest possible division of network nodes by disabling the fewest number of links.

As in the case of the min-cut problem, securing the links comprising the minimum network bisection will patch a potential network vulnerability, increasing the network's survivability. Therefore, the minimum bisection also poses a solution to the identification of critical elements in a random network.

6.2.4 Link Valuability

The solution to both the min-cut and network bisection 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:

$$\vartheta(u,v) = \frac{\sum_{U|u\in U\wedge v\in\bar{U}\vee v\in U\wedge u\in\bar{U}} \frac{D(U,\bar{U})}{C(U,\bar{U})}}{|U \mid u \in U \wedge v \in \bar{U} \vee v \in U \wedge u \in \bar{U}|} c \tag{6.8}$$

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.

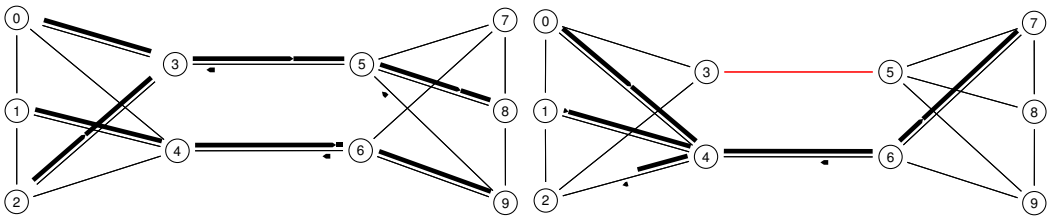

Figure 6.5: A dumbbell-shaped homogeneous network used for network simulation and analysis in which links $\{3, 5\}$ and $\{4, 6\}$ have higher valuability than the others: (a) FTP traffic flows from sources at nodes 0–2 to sinks at nodes 7–9 over highly utilized links $\{3, 5\}$ and $\{4, 6\}$; (b) 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) [18].

Dumbbell Network Simulation Environment

In order to test how well the link valuability property holds, a simulation model was designed to test the performance of a dumbbell-shaped graph that shares the homogeneous property of random network models (i.e., nearly all links have the same number of neighbors). This 4-graph, shown in Figure 6.2.4, is comprised of ten nodes, and was simulated using the network simulator *ns-2* [18]. 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 series of simulations was run for each traffic scenario. Each simulation series consisted of consecutive network simulations with identical traffic flows from sources at nodes 0–2 to destinations at nodes 7–9. In each run 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.

In the first simulation series, one flow of file transfer protocol (FTP) traffic was sent across the network from each of nodes 0–3 to respective nodes 7–9. The traffic load across links $\{3, 5\}$ and $\{4, 6\}$, before any link was disabled, utilized a respective 99.44% and 75.24% of link capacity. The results of this series of simulations are plotted on graphs, correlating network performance metrics with the valuability

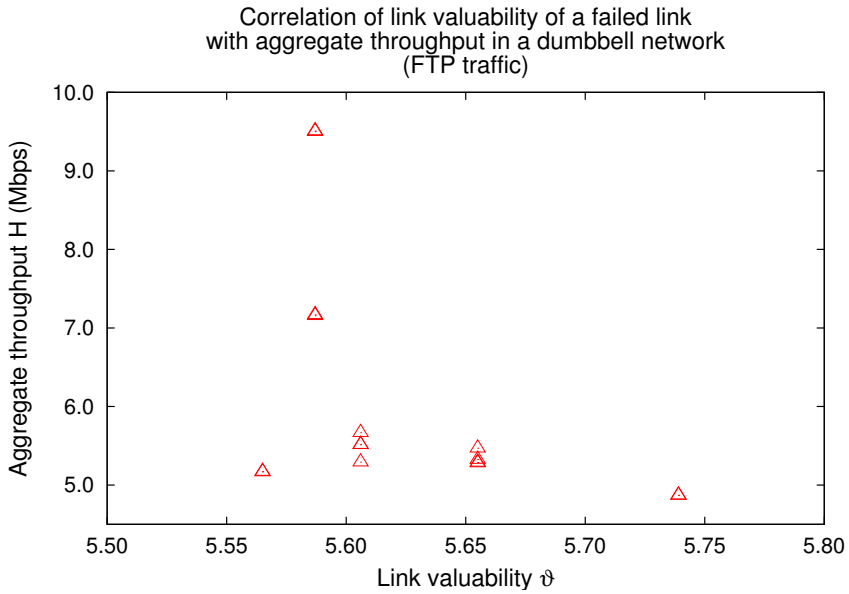

Figure 6.6: Correlation of aggregate throughput H from simulations involving FTP traffic in a dumbbell-shaped network with the link valuability $\vartheta(u, v)$ of a particular link $\{u, v\}$ dropped in each simulation. When the most valuable links are disabled, as a result of attack, the aggregate throughput is lowest. When less valuable links are disabled, the throughput remains relatively higher.

of each disabled link. Figure 6.2.4 shows the correlation of aggregate throughput with link valuability. The network produced the lowest aggregate throughput after the most valuable links $\{3, 5\}$ and $\{4, 6\}$ were disabled.

A decrease in aggregate throughput H experienced in a network could be attributed to several circumstances: decreased flow rate at the traffic source; or dropped packets at one or more intermediate routers. In the former case network link metrics might exhibit behaviors otherwise indicative of higher performance because less data is traversing the network. In the latter case, network link metrics would likely show signs of degraded performance because the links are the source of the degradation. A combination of these two circumstances may also result in smaller H.

For further analysis of network performance, the link utilization ρ of the simulations involving the FTP traffic was analyzed. Figure 6.2.4 shows the mean (a) and 85th percentile (b) link utilization graphed against the valuability of the failed link. The graphs show a decrease in ρ as the valuability of the failed link increases. The transmission control protocol (TCP), utilized by FTP, reduces the rate of traffic being sent across the network when congestion is detected (i.e., packets are dropped) [4]. This action reduces the packet arrival rate λ at intermediate routers,

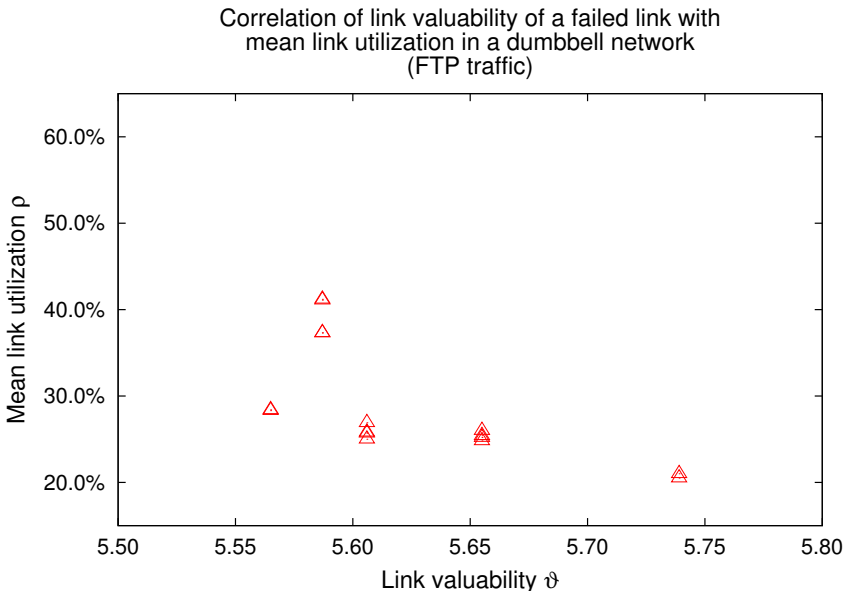

(a)

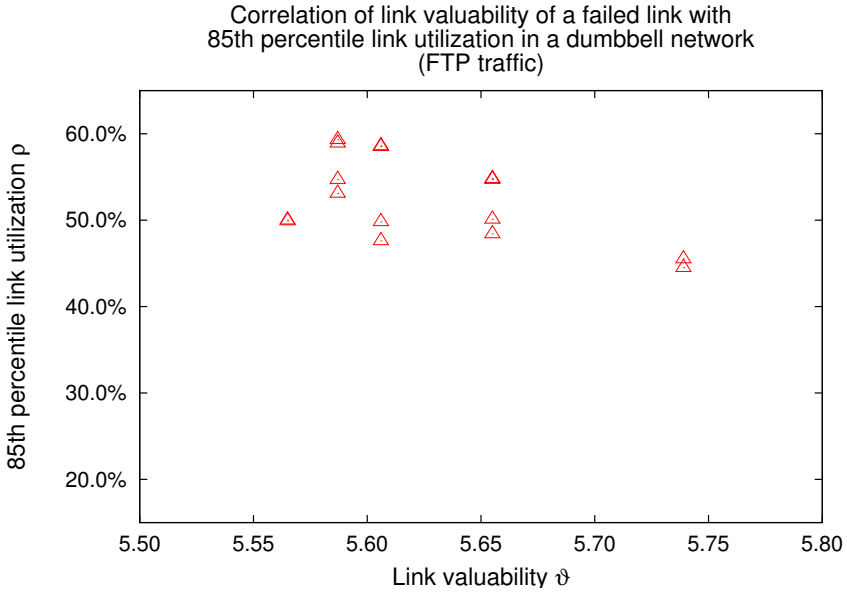

(b)

Figure 6.7: Correlation of mean (a) and 85th percentile (b) link utilization ρ from simulations involving FTP traffic in a dumbbell-shaped network with the link valuability $\vartheta(u, v)$ of a particular link $\{u, v\}$ dropped in each simulation.

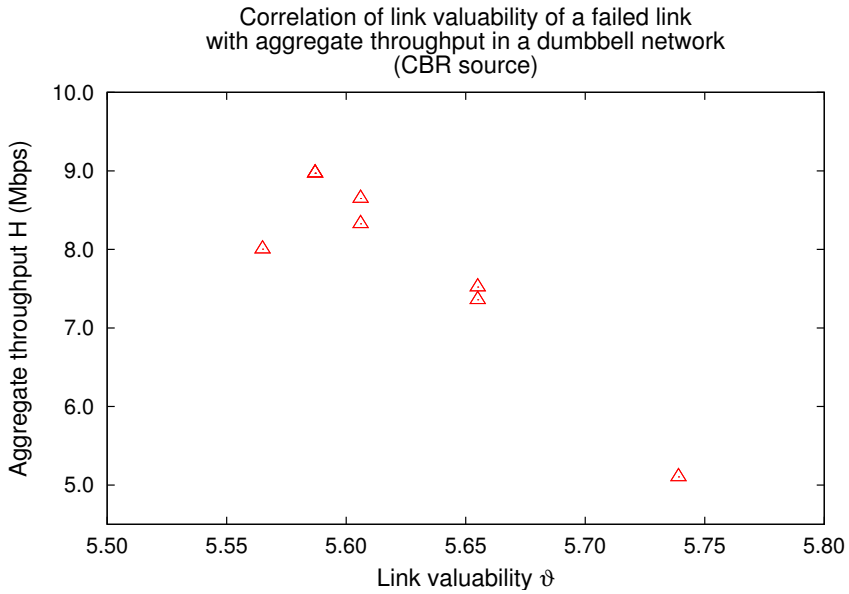

Figure 6.8: Correlation of aggregate throughput H from simulations involving CBR traffic in a dumbbell-shaped network with the link valuability $\vartheta(u, v)$ of a particular link $\{u, v\}$ dropped in each simulation. As the valuability of a failed link increases, the aggregate throughput in the network decreases.

resulting in a decreased ρ.

To analyze network performance without the effects of TCP in the simulations, the second series of network simulations utilized constant bit rate (CBR) traffic using the user datagram protocol (UDP), which does not adjust its rate based on network activity. In this series of simulations traffic sources at nodes 0–2 transmitted data at a rate of 1.0Mbps to each of nodes 7–9, utilizing 85–95% of the capacity of links $\{3, 5\}$ and $\{4, 6\}$.

A graph correlating aggregate throughput H with link valuability ϑ in this series of simulations is shown in Figure 6.2.4. 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 [34], 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 6.2.4 maps both the mean (a) and 85th percentile (b) of ρ following a link failure to the corresponding valuabilities of disabled links. The mean utilization, shown in Figure 6.2.4a, decreases as the valuability of failed links increases. However, for several reasons the mean value

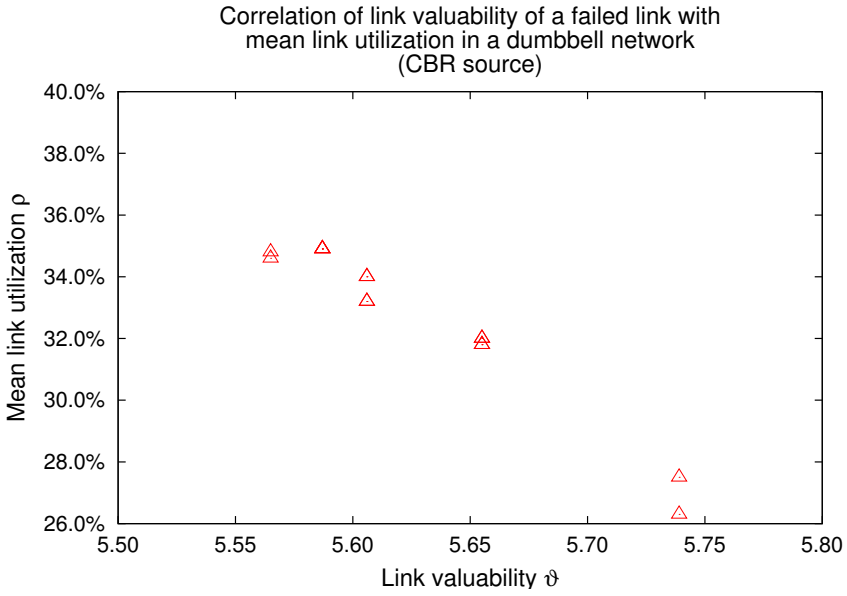

(a)

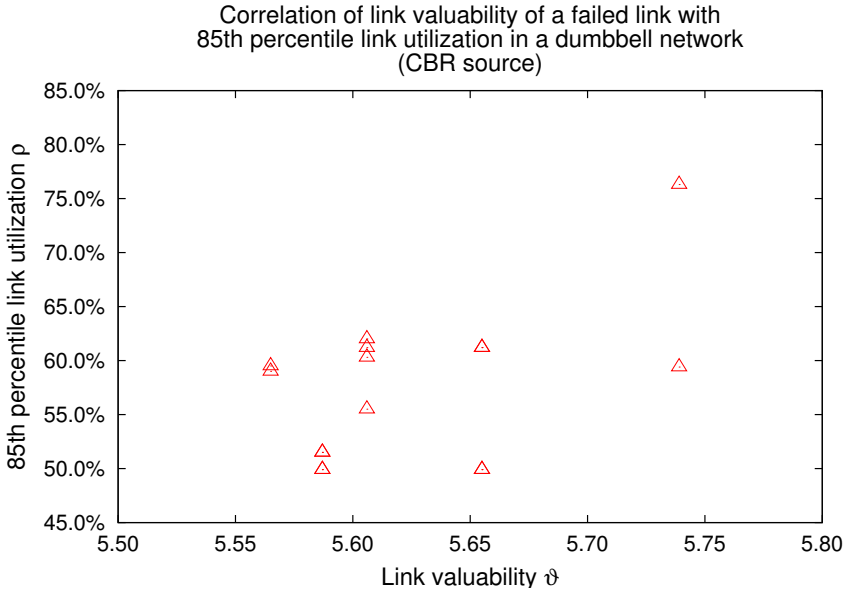

(b)

Figure 6.9: Correlation of mean (a) and 85th percentile (b) 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.

is deceiving when analyzing a performance metric such as link utilization. First, the available paths by which data can travel through the network from source to destination are limited after a particular link fails. Because of this, flows that before traveled separate paths may now overlap at some links, and other links will now be under-utilized. In this example flows which before traveled over links $\{3, 5\}$ and $\{4, 6\}$ (Figure 6.2.4a) use only link $\{4, 6\}$ after link $\{3, 5\}$ has been attacked (Figure 6.2.4b). In addition, if flows sharing a path cause saturation at some point in the path, and packets are dropped at that interface, then the utilization may appear more regular at subsequent hops in the path—the interface of node 4 leading to node 6, in this example. The mean value, therefore, will be calculated based on a smaller number of utilized links and a larger number of under-utilized links.

A better statistic for measuring network performance based on the metrics at comprising links is the *nth percentile*. At most $n\%$ of the utilization of remaining network links is less than this value, and at most $(100 - n)\%$ are greater. A network analyst using the nth percentile to analyze performance at router interfaces can better determine network health, because of an accounting for highly-utilized links that may not otherwise be considered. In this series of simulations, the 85th percentile of ρ was recorded, and is shown in Figure 6.2.4b. In this graph 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 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 mean (a) and 85th percentile (b) queuing delays, calculated using Little's formula [38], are shown in Figure 6.2.4. Again the mean T_w value, graphed in Figure 6.2.4a, is not indicative of network performance, so the 85th percentile T_w value is included in Figure 6.2.4b. This latter graph shows that when the most valuable network links fail, queuing delay is the highest.

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. This was most evident in the simulations using CBR traffic, in which TCP overhead was not felt in the network. These results support the claim that critical links can be identified in random networks, the preservation of which increases network survivability.

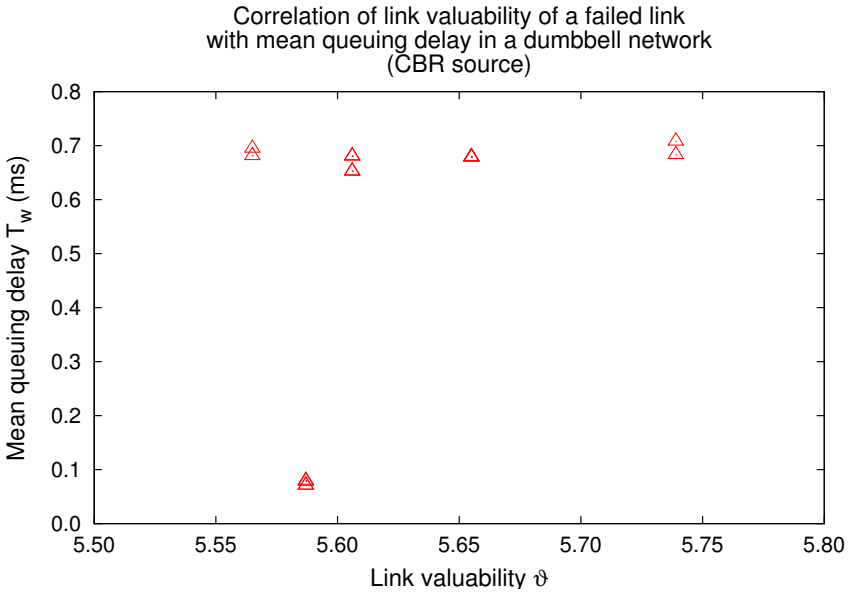

(a)

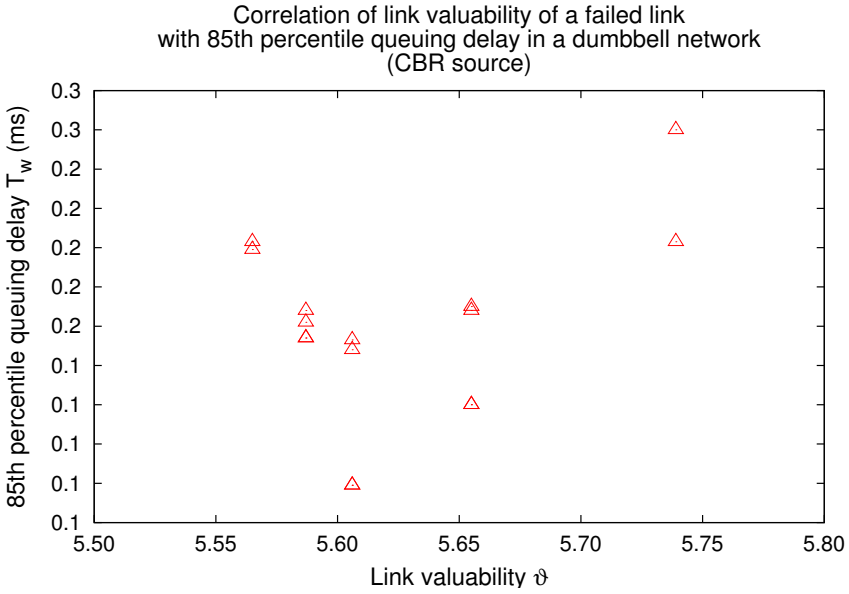

(b)

Figure 6.10: Correlation of mean (a) and 85th percentile (b) 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.

6.3 Time Complexity

The mechanisms described in this chapter have been shown by empirical and theoretical analysis to be effective in finding defining critical nodes. However, the feasibility of utilizing such mechanisms to identify critical network elements could require substantial resources, depending on the size of the network.

The number of nodes neighboring any particular node (i.e., connectivity) is fairly trivial to come by. A network administrator should know the topology of the network in question. If the network is running a link-state routing protocol (e.g., OSPF), this information is stored in the link-state database of every router. Determining the diameter and average distance of a network requires $O(n^3)$ for a n-node network [15]. However, a link-state protocol runs the shortest path algorithm in parallel (i.e., one instance per router), so the time complexity is $O(n^2)$ in such a network, making it more feasible.

The min-cut of a network can be determined in polynomial time, through the research of Tom Leighton, et al. [26]. Network bisection is a NP-hard problem, but an approximation can be found in $O(\sqrt{n}\log n)$ time [19].

The time complexity for finding the valuability of a link, as presented in this paper, is exponential. Although there is a large time complexity for this problem, multicommodity max-flow min-cut theorems have been used to approximate similar problems and may prove useful in solving this problem [26].

Chapter 7

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 thesis 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 summarized various network-layer protocol attacks, such as routing table poisoning attacks, DoS attacks, and man-in-middle-attacks. The notion of protecting against such attacks using a selective security model was discussed. As a case study, the OSPF routing protocol was examined, including some vulnerability studies, native security mechanisms, and the PKI implementation OSPF with Digital Signatures. The advantages to deploying a selective-security OSPF mechanism over a universal security scheme were discussed.

Network survivability was also discussed in this research, in order to measure accuracy in identifying critical network elements. The characteristics of scale-free and random network models were described, and data from real-world graphs following these models were obtained for survivability analysis. Survivability was defined in terms of communications networks, and performance metrics were outlined for measuring survivability in such networks. A queuing model was used to examine performance at the link level, and aggregate throughput was a measure of overall network

performance. Other desirable topological characteristics included a connected network graph (i.e., no fragmentation) and low diameter and average distance between nodes.

Attack models were simulated on a scale-free network topology, using a tool built in-house for topology analysis, to test its survivability in different kinds of attacks. Results showed that when the nodes with the highest connectivities were disabled by attack, the network suffered higher fragmentation, and the relative average distance—used for comparing networks of different sizes—increased more than when other (random) nodes were disabled in the network. The nodes with the highest connectivities in scale-free networks were identified as elements critical to network survivability.

The topology analysis tool was also used to examine attack models in random networks. The response in random networks differed from that of scale-free networks. Attacks that disabled nodes with high connectivities didn't damage the network significantly more than random failures did, in terms of fragmentation and increase in relative average distance. Different strategies were needed to identify critical elements in random networks.

In order to better identify critical elements in random networks, network data flow was examined, and network cuts were introduced. Previous research showed that the min-cut and minimum-bisection of a network graph are areas vulnerable to fragmentation and increased congestion, if attacked, and are thus critical to the survivability of random networks [26, 19].

The link valuability metric—a function of the network cuts of which a link is a part—was introduced in order to quantify and compare the critical nature of individual links in a random network. In order to test the validity of this metric, a sample network was created and simulated using the network simulator *ns-2* [18]. The network followed a pattern of homogeneity—a characteristic of random networks—but deliberately designated two links as those with the highest valuability by positioning half of the network nodes on each side of these links, in a dumbbell formation. In the simulations traffic was transmitted from one side of the network to the other, traversing the two most valuable links, and each link was removed in turn to test the damage to network performance and correlate that effect with the valuability of the link that was removed. Results from these simulations showed that as the valuability of the attacked link increased, aggregate throughput in the network decreased. Also,

the link utilization and queuing delay were highest after the most valuable links were disabled.

The results from simulations correlating link valuability with network performance showed that links with higher valuabilities are more critical to the survivability of random networks. These results support the claim that critical elements can be identified in random network topologies.

In this research, no specific mechanism was designed for implementing network-layer selective security, but a foundation was laid for future design of such a mechanism. The results from this thesis indicate that critical elements can be identified within scale-free and random networks. Security applied to these elements will increase survivability in such networks. The creation of a mechanism for the selective deployment of network security to elements critical to network survivability would ensure greater stability in networks of high importance.

Bibliography

[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] M. Allman, V. Paxson, and W. Stevens. TCP congestion control. RFC 2581, Apr 1999.

[5] Kevin Anderson. US 'fears al-Qaeda hack attack'. *BBC News Online*, Jun 2002. Online at http://news.bbc.co.uk/1/hi/sci/tech/2070706.stm.

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

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

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

[9] Ann Bednarz. Hospital sounds alarm after 3-day struggle. *Network World Fusion*, Nov 2002. Online at http://www.nwfusion.com/news/2002/1125bethisrael.html.

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

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

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

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

[14] 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.

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

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

[17] 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.

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

[19] Uriel Feige, Robert Krauthgamer, and Kobbi Nissim. Approximating the minimum bisection size. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 530–536, Portland, OR, May 1999.

[20] 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.

[21] 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.

[22] 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.

[23] 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.

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

[25] S. Jha, J. Wing, R. Linger, and T. Longstaff. Survivability analysis of network specifications. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN 2000)*, pages 613–622, New York, NY, Jun 2000.

[26] 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.

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

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

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

[30] S. L. Murphy and M. R. Badger. Digital signature protection of the OSPF routing protocol. In *Proceedings of the Symposium on Network and Distributed System Security*, pages 93–102, San Diego, CA, Feb 1996.

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

[32] Denise Pappalardo. Can one rogue switch buckle AT&T's network? *Network World Fusion*, Feb 2001. Online at http://www.nwfusion.com/news/2001/0223attupdate.html.

[33] Radia Perlman. *Network layer protocols with Byzantine robustness*. PhD thesis, Massachusetts Institute of Technology, 1988.

[34] 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.

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

[36] Paul Roberts. Slammer: One year later. *Network World Fusion*, Jan 2004. Online at http://www.nwfusion.com/news/2004/0120slammoney.html.

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

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

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

[40] B. Vetter, F. Wang, and S.F. Wu. An experimental study of insider attacks for OSPF routing protocol. In *Proceedings of the 1997 International Conference on Network Protocols*, pages 293–300, Atlanta, GA, Oct 1997.

[41] F. Wang and S. Wu. On the vulnerabilities and protection of OSPF routing protocol. In *Proceedings of the 7th International Conference on Computer Communications and Networks*, pages 148–152, Lafayette, LA, Oct 1998.

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

[43] David L. Wilson. The Internet vs. the bomb: Would the Internet survive the bomb? *CNN*. Online at http://edition.cnn.com/SPECIALS/cold.war/experience/technology/internet.bomb.