

Survivability of Large Networks in the Presence of Malicious Attacks

Casey T. Deccio, Spencer Cox, Matthew Smith, Jacob Wan,
Mark Clement, and Quinn Snell
Computer Science Department
Brigham Young University
Provo, UT 84602

Abstract

The Internet has become the foundation for world-wide digital communication. The survivability of this critical network infrastructure is crucial to businesses, universities, and government agencies. This survivability can be compromised by the failure of a small percentage of critical routers within the network. Additional security is justified on these critical routers, even if resources are not available to secure all routers in the network topology. This research identifies critical routers in a network by estimating the damage to the network if that router were the target of a Denial of Service (DoS) attack. Topological data from several real-world studies is used to validate this research.

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 networks is vital to maintaining stability in many daily processes.

At the heart of large network infrastructures, such as the Internet, is the network-layer protocols. At this layer routing mechanisms establish virtual links to fully connect entire networks, making possible global communication. 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” [5]. The increased availability of tools allowing

permissive network access to malicious users has made these attacks a reality. Specific attacks include Domain Name System (DNS) hacking attacks, routing table poisoning attacks, packet mistreatment attacks, and DoS attacks [6].

Securing higher-level (e.g., transport, application layers) security protocols have been the focus of much recent research, but without security at the lower layers, computer networks are left vulnerable to attack. Security features for routing protocols such as the Border Gateway Protocol (BGP) [15] and the Open Shortest Path First (OSPF) protocol [11] have been proposed [12–14, 16, 17]. However, these features are often not deployed because of their expense, their performance cost, and the risk associated with introducing something new into a stable network environment [9].

In this research we analyze the effects of different attacks and failures on computer networks so that custom levels of network-layer security can be applied where they will make the most difference. The functional loss of particular network nodes can have a varying effect on the stability of the graph as a whole. Some networks are greatly affected by the incidental failure of random nodes, while others remain virtually unchanged. Some topologies are fault tolerant but experience degraded performance in the presence of malicious attacks. An important contribution in this research is the identification of *critical nodes* in a network, where additional security can be justified.

The remainder of this paper is organized as follows: Section 2 describes some of the previous work in network-layer security; Section 3 discusses the characteristics of large network topologies relative to our research; Section 4 discusses the attack models used in our experiments; and Section 5 contains experimental results.

2 Network-layer security

Some of the seminal work in securing the network layer was published in the Ph.D. thesis of Radia Perlman [14]. Her work describes different levels of network robustness and outlines a public key system for identifying trusted nodes in a link-state network. Terms borrowed from this thesis are *simple failure*, in which a node or link becomes inoperative, and *Byzantine failure*, in which nodes or links continue to operate, but

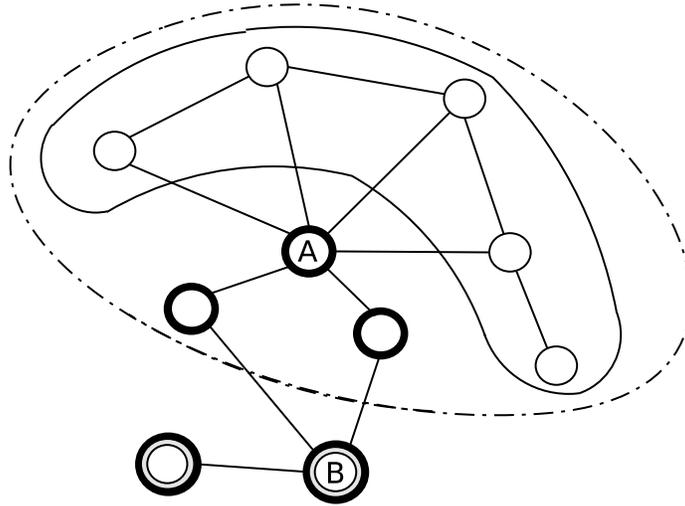


Figure 1: Impact of node failure on the relative diameter of the network. If node A is lost due to an attack, all of the nodes with black bold outline will be disconnected and the relative diameter will increase dramatically. If node B is lost, the nodes with grey outlines will be disconnected and the relative diameter will exhibit a minimal increase. Additional security should be implemented for node A .

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.

Other research has extended Perlman's work to secure specific protocols. Murphy, et al, [12, 13] have developed a digital signature scheme for protecting OSPF. In securing BGP, Smith, et al, [16, 17] have used encryption, digital signatures, and the use of predecessor information to verify autonomous system (AS) paths.

These methods are aimed at securing routing protocols against malicious intrusion and Byzantine failures. They all utilize some sort of cryptographic solution, which can add overhead and complexity to a router. These methodologies promise security and robustness when all routers in a large network adopt the prescribed security measures. In this research we identify *critical nodes* on which this security should be deployed in networks.

3 The topological nature of large networks

Determining critical nodes in a network can be accomplished by examining natural topological characteristics. This section examines some of the characteristics of the Internet and other large networks. Scaling and connectivity distribution are used to classify types of networks, including scale-free and homogeneous networks. Network diameter is a useful metric in determining the efficiency of a network, in the case of attack or failure.

3.1 Scale-free networks

The complexity of the Internet topology is attributed to the unmanaged and rapid growth that has occurred since its inception. The complex nature of the Internet makes it difficult to classify. Related research has categorized similar infrastructures for social and biological systems that occur in nature [3,18]. Researchers have observed that such large networks organize themselves into a *scale-free* state [4].

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: $P(k) \sim k^{-\gamma}$ [4]. Relatively few nodes are highly connected in a scale-free network; the majority of nodes have very few neighbors.

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 [8] merged with data also obtained in 1999 from the Internet Mapping Project at Lucent Bell Laboratories [10]¹ These studies produced a topology consisting of 284,805 Internet routers. 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 has 100 connections is $P(100) = 2.5 \times 10^{-5}$, while the probability that a node only neighbors one node 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

¹The Internet Mapping Project is now run by Lumeta Corporation. More information can be found at their Web site [7].

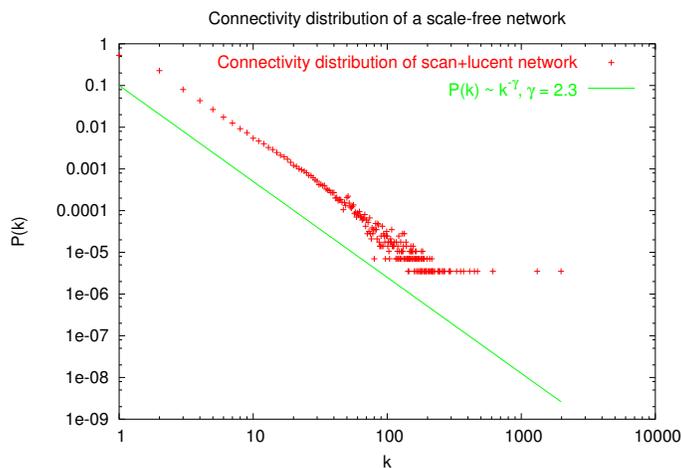


Figure 2: The normalized connectivity distribution of the scan+lucent network. The line below the plot points represents the connectivity distribution for the scale-free model, with $\gamma = 2.3$.

malicious attacks [2]. 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.

3.2 Homogeneous networks

While the graph of nodes and AS comprising the Internet tend to self-organize according to scale-free characteristics, nodes within an AS exhibit different behaviors. The connectivity of these nodes follows more of a normal distribution, indicating homogeneity in the graph. Data used to characterize homogeneous networks consisted of the topology of a major telecommunications vendor’s frame relay network, hereafter referred to as the “vendor” data. A normalized graph of the connectivity distribution of the 221 routers comprising this AS is shown in Figure 3. There are fewer critical nodes in this topology, and a large majority of the nodes have between eight and ten connections.

A homogeneous topology such as that shown in Figure 3 carries its own set of vulnerabilities. Any random node failure will likely cause equal damage in a topology with normally distributed connectivities. Such networks, while not particularly prone to malicious attack, can experience significant performance degradation in the presence

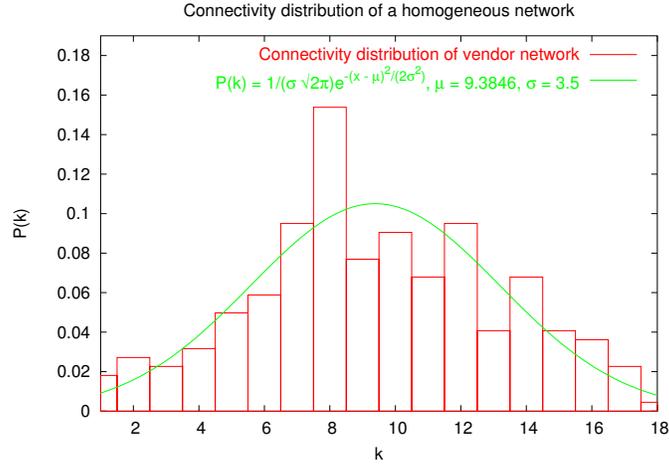


Figure 3: The normalized connectivity distribution of the vendor network.

of random node failures [2].

3.3 Network efficiency

Because low latency is desired with data transmission through a network, the efficiency of a network with n nodes is related to the lengths of the data paths through the network. The average length of shortest paths (in hops) between any two nodes is referred to as the *diameter* d of a network [1]. A network's diameter is a measure of its connectivity and efficiency in transporting data.

A change in diameter is used to analyze the robustness of network efficiency amid malicious attack or failure, in which a is the number of attacked or failed nodes removed, and r represents the number of functioning nodes that remain connected to the network. If the diameter increases as the result of an attack, traffic takes longer to traverse the network, which results in increased delay, average queue size, and the possibility of congestion and data loss.

Relative diameter d' is used for comparison and analysis of the results, so that the characteristics of networks of varying sizes could more accurately be compared. Because diameter is defined as the shortest path between arbitrary nodes, we define the relative diameter as ratio of the actual diameter d to the number of paths in the graph r^2 , $d' = \frac{d}{r^2}$. As the diameter increases, this metric increases linearly. As nodes

are lost, the metric experiences more significant increases.

Figure 1 illustrates the varying impact of node failure on the connectivity and the relative diameter of a network. The relative diameter of the original network $d'_0 = 0.0201$. When node A is lost due to an attack, the relative diameter increases by a factor of four to $d'_a = 0.08$. When node B is the target of an attack the relative diameter of the network experiences a minor increase to $d'_b = 0.028$. The overhead of securing node A is justified by the significant increase in relative diameter that the network experiences when A is lost.

4 Attack models

In order to determine the impact of securing critical nodes attack models were applied to the scan+lucent data and the vendor data [2]:

- repeated simple failures of random network nodes
- repeated malicious attacks aimed at nodes of decreasing connectivity

4.1 Simple failures

Applying Perlman’s definition of simple failure [14] iteratively to random nodes, we compared the functionality loss of the scale-free scan+lucent network with that of the homogeneous vendor network. The stability of the scan+lucent and vendor networks at high failure rates are shown in Figures 4 and 5, respectively. When only approximately 1% of the nodes are removed from the scan+lucent network, over 97% of the network topology remains intact. This demonstrates the robustness of scale-free networks, even at extremely high failure rates. The large majority of nodes have only one neighboring node, so network losses remain low with increasing failures of random nodes. When 10% of the nodes are removed from the smaller vendor topology, the network maintains 90% of its nodes. This statistic reflects the nature of the homogeneous graph; the majority of the nodes have between eight and ten neighbors, so the loss of a single node is unlikely to bring down other nodes with it.

Figures 6 and 7 demonstrate the efficiency of the scan+lucent and vendor networks,

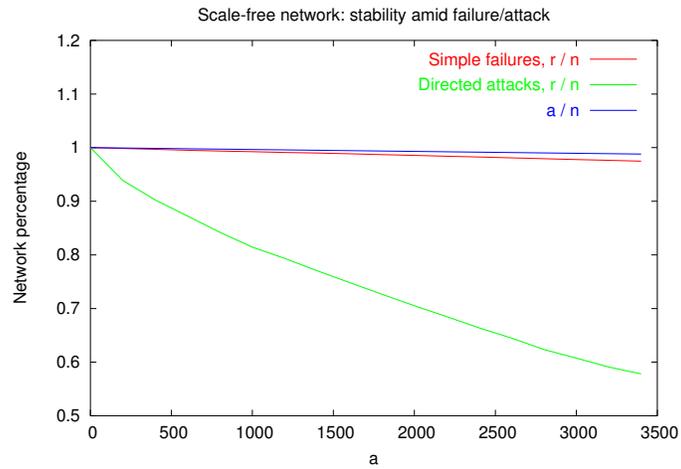


Figure 4: Network stability of the lucent+scan network. When over 1% of the nodes are randomly removed from the network, the network retains over 97% of its original nodes, but when the top 1% most connected nodes are removed from the network, the size of the functioning network drops below 60% of its original size.

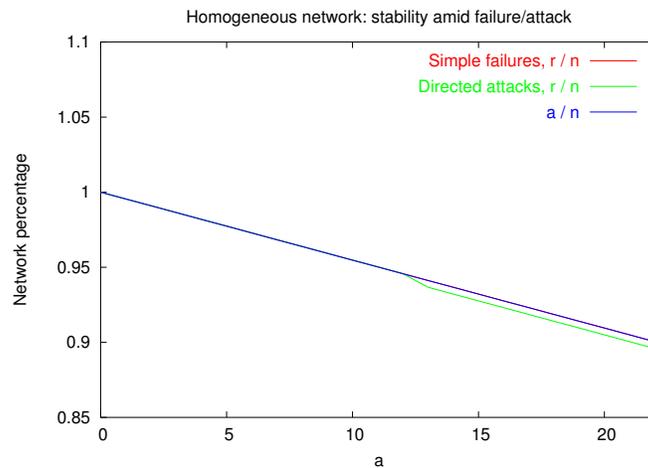


Figure 5: Network stability of the vendor network. Nodes removed by both random selection and in order of decreasing connectivity exhibit similar characteristics. When nodes are removed by attack or by simple failure at random nodes, very few other nodes are made inaccessible.

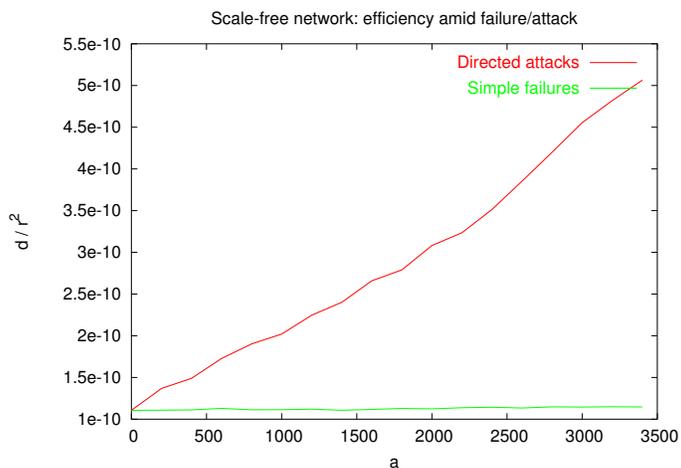


Figure 6: Network efficiency of the scan+lucent data in the presence of failure and attack. The relative diameter of the network remains nearly unchanged even after 1% of the nodes are removed from the network. The relative diameter increases linearly as nodes are removed in order of decreasing connectivity.

respectively, in the presence of simple failures. Although over 1% of the network nodes have been removed from the scan+lucent data, the relative network diameter remains unchanged, again displaying the inherent fault tolerant nature of scale-free networks. However, the relative diameter of the vendor network rises linearly, and is nearly 25% larger than the original when 10% of the nodes are removed. This is because the diameter of this graph depends not on the connectivity of a small concentration of nodes, but rather on the collective connectivity of the graph.

4.2 Directed attacks

Figures 4 and 5 display the results of repeating directed attacks on the scan+lucent and vendor topologies, respectively. Nodes at each iteration are removed in order of decreasing connectivity, as if they had been the target of a DoS attack. The scan+lucent graph, which stayed relatively stable amid simple failures, lost over 1/3 of its entire network when only 1% of its nodes were attacked. This is because the highest connectivity was concentrated in a small percentage of critical nodes, and when these nodes were attacked, the network quickly became disconnected. The vendor network, however, remained as tolerant to the attack as it did to the excessive simple failures: nearly 90%

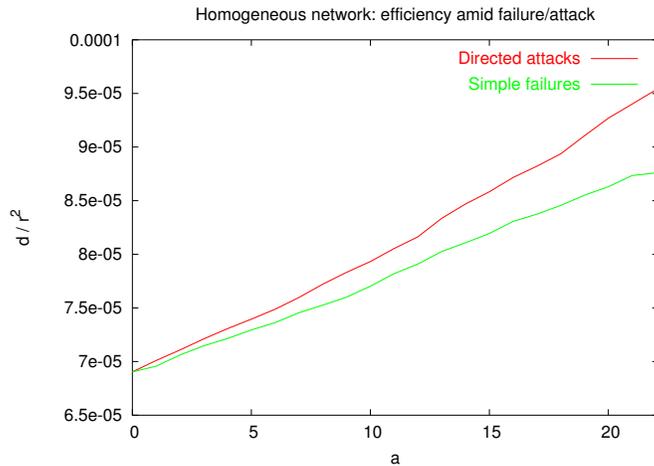


Figure 7: Network efficiency of the vendor topology in the presence of failure and attack. The relative diameter of the network increases linearly both when nodes are removed in decreasing connectivity and when nodes are removed at random, although the rate of increase is greater for directed attacks.

of the network remained in tact when 10% of the most connected nodes were removed. This is attributed to the normal distribution of connectivity in this network.

The effect of directed attacks against the scan+lucent and vendor topologies are shown in Figures 4 and 5, respectively. The relative diameter of the scan+lucent network rises linearly in the presence of an attack directed at nodes of decreasing diameter, rising to five times the relative diameter of the original network. This increase in diameter adds delay to traffic, and will likely cause congestion. When 10% of the highest connected nodes of the vendor data are removed, the relative diameter increases at a slightly higher rate than it did in the presence of overwhelming failure, indicating that the attack affects the homogeneous network, but only slightly more than normal failure would affect the network.

5 Conclusions

Network-layer security is vital to the stability and efficiency of digital communication. Secure implementations have been developed but not widely deployed because of costly resources and the risk involved. This research examines the nature of network topologies to identify critical nodes on which additional security can be justified because of

their influence on the network infrastructure. The effects of scale-free and homogeneous networks in the presence of simple failure and malicious attacks are examined. Experimental data shows that if attacks on scale-free networks are directed at nodes with decreasing connectivity, the size of the network is reduced significantly when relatively few nodes are removed, and the relative diameter rises significantly. Homogeneous networks exhibit similar changes in the presence of either simple failure or attack.

By securing the most vital portions of the network infrastructure, large scale-free networks can remain more resilient to attack and maintain their dependability.

References

- [1] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Diameter of the world-wide web. *Nature*, 401:130–131, Sep 1999.
- [2] Réka Albert, Hawoong Jeong, and Albert-László Barabási. The Internet’s Achilles’ Heel: Error and attack tolerance of complex networks. *Nature*, 406:378, 2000. Online at <http://citeseer.nj.nec.com/albert00internets.html>.
- [3] Jayanth R. Banavar, Amos Maritan, and Andrea Rinaldo. Size and form in efficient transportation networks. *Nature*, 399:130–132, May 1999.
- [4] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286:509–512, Oct 1999.
- [5] S. M. Bellovin. Security problems in the TCP/IP protocol suite. *Computer Communications Review*, 19(2):32–48, 1989. Online at <http://www.research.att.com/~smb/papers/ipext.pdf>.
- [6] Anirban Chakrabarti and G. Manimaran. Internet infrastructure security: a taxonomy. *IEEE Network*, 16(6):13–21, Nov/Dec 2002.
- [7] Lumeta Corporation. Internet mapping project. <http://www.lumeta.com/mapping>.

- [8] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for Internet map discovery. In *IEEE INFOCOM 2000*, pages 1371–1380, Tel Aviv, Israel, March 2000. IEEE.
- [9] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Efficient security mechanisms for routing protocols. In *Proceedings, Tenth Annual Network and Distributed System Security Symposium (NDSS 2003)*, Feb 2003.
- [10] USC Information Sciences Institute. Internet maps. <http://www.isi.edu/scan/mercator/maps.html>.
- [11] J. Moy. OSPF version 2. RFC 2328, Apr 1998.
- [12] S. Murphy, M. Badger, and B. Wellington. OSPF with digital signatures. RFC 2154, Jun 1997.
- [13] S. L. Murphy and M. R. Badger. Digital signature protection of the OSPF routing protocol. In *Proceedings, the Symposium on Network and Distributed System Security*, pages 93–102, Feb 1996. Online at <http://ieeexplore.ieee.org/iel3/3553/10644/00492416.pdf>.
- [14] Radia Perlman. *Network layer protocols with Byzantine robustness*. PhD thesis, Massachusetts Institute of Technology, 1988.
- [15] Y. Rekhter and T. Li. A border gateway protocol 4 (BGP-4), Jul 1994.
- [16] B. Smith and J. Garcia-Luna-Aceves. Securing the border gateway routing protocol. In *Proceedings, Global Internet '96*, London, UK, Nov 1996. Online at <http://citeseer.nj.nec.com/smith96securing.html>.
- [17] Bradley R. Smith, Shree Murphy, and J. J. Garcia-Luna-Aceves. Securing distance-vector routing protocols. In *Proceedings of the Symposium on Network and Distributed System Security*, pages 85–92, San Diego, CA, Feb 1997. Online at <http://citeseer.nj.nec.com/smith97securing.html>.
- [18] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, Jun 1998.