A Network Based Simulation Approach to Cybersecurity Policy

Sean P. Gorman*, Rajendra G. Kulkarni,
Laurie A. Schintler, Ph.D., and Roger R. Stough, Ph.D.

School of Public Policy,
George Mason University
Fairfax, Virginia 22003, U.S.A.
*Corresponding author
e-mail: sgorman1@gmu.edu

ABSTRACT

Cybersecurity is an issue of increasing concern since the events of September 11th. Many questions have been raised concerning the security of the Internet and the rest of US's information infrastructure. This paper begins to examine the issue by analyzing the Internet's autonomous system (AS) map. Using the AS map, generic malicious infections are simulated and different defense strategies are considered in a cost effectiveness analysis framework. The results show that protecting the most connected nodes provides significant gains in security and that after the small minority of the most connected nodes are protected there are diminishing returns for further protection. Although if parts of the small minority of the most connected firm are not protected, such as non-US firms, protection levels are significantly decreased. A simple cost function is also proposed to prepare cost effectiveness analysis of different security strategies, and specific financial and US federal government policies will be simulated and compared.

Keywords: cybersecurity, policy, cost effectiveness analysis, network analysis

INTRODUCTION

The Internet is an amalgam of thousands of interconnected networks. Some of these networks are vast global networks like Worldcom (MCI) or Cable & Wireless while others are small local networks like a university. The individual networks that compose the Internet are commonly called autonomous systems (AS) and number roughly 12,000 active AS's with 22,000 assigned and roughly 65,000 AS's possible (based on a 16 bit number) (Gao 2001). The task of trying to provide a minimum level of security for all these networks is a daunting effort, but one that has been increasingly highlighted as an area of importance for national security (CIPB 2002, NRC 2002, NSTAC 2002). Innovative approaches are called for to tackle a problem of such a large scale and increasingly global nature. Recently, researchers in the many fields have begun work concerning the fundamental structure of complex interaction of the networks that comprise the Internet (Faloutsos et al 1999, Tangmunarunkit 2002, Lakhina et al 2002, Yook et al 2001). Much of the work has revolved around the finding that the Internet at the AS and router level form a scale free network (Faloutsos 1999, Albert and Barabasi 2002). An understanding of the mechanics underlying the growth and evolution of the Internet provides a new perspective for the role policy can play in helping foster a more secure Internet. A review of the literature pertaining to complex networks will be discussed with specific emphasis on implications for security in networks. The goal of this paper is to investigate possible least effort strategies to protect the network with a minimal level of intervention. The research will then be placed into the context of the current policy debate over cybersecurity.

Attack Effects and Internet Structure

On Saturday January 15th 2003 at 5:30 UTC the SQL Slammer worm emerged from somewhere in East Asia and propagated around the globe, doubling every 8.5 seconds and infecting 90% of vulnerable machines in under 10 minutes (Moore et al 2003). While the SQL Slammer did not carry a malicious payload, the sheer amount of traffic it produced swamped networks causing 13,000 Bank of America ATMs to become disconnected, cancelled airline flights, and disrupted elections and 911 services. The spread of SQL slammer worm was a warning of not only the speed and scope of malicious worms but the level of interdependency of the Internet with other critical infrastructures (banking and finance, transportation, medical, public safety and governance). The speed of the worm is all the more confounding when the spread and the complexity of infrastructure it traversed is considered.

The individual autonomous systems that compose the Internet broadly fall under three categories:

- **Stub AS** - It is connected to only one other AS. For routing purposes it is treated as part of the parent AS.
- **Multihomed AS** - It is connected to more than one other AS, but does not allow transit traffic. Internally generated traffic can be routed to any of the connected ASs. It is used in large corporate networks that have a number of Internet connections, but do not want to carry traffic for others.

- **Transit AS** - It is connected to more than one other AS and it can be used to carry transit traffic between other AS's. (Magoni and Pansiot 2001)

In addition to a basic typology AS's are often ranked into tiers, from 1-5. The tier 1 AS's are global networks, down to tier 5 networks consisting of local area networks for organizations and firms. The complexity of the Internet's infrastructure is daunting; these networks reside in numerous countries and fall under a wide variety of jurisdictions and most often are subject to little to no regulation, oversight or central control.

Error and Attack Tolerance of Complex Networks

Scale free networks have many implications, but a far-reaching consequence of their unique structure is they are very fault tolerant but also susceptible to attack (Albert et al 2000). Specifically, a scale free network model remains connected when up to 80% of nodes are randomly removed from the network, but when the most connected nodes are removed the average path length of the network increases rapidly, doubling its original value when the top 5% of nodes are removed (Albert et al 2000). In short, targeting the most connected nodes can cause significant damage to a scale free network, making them highly susceptible to a coordinated and targeted attack against them. Albert et al's work was complimented by the analysis of Callaway et al (2000) modeling network robustness and fragility as a percolation and by Cohen et al (2001) using related methodologies. Preliminary analyses of these models on spatial network data have shown similar results when cities are the nodes and fiber connections between them are the links. Utilizing a model of node connectivity and path availability Grubesic et al

(2003) find that the disconnection of a major hub city can cause the disconnection of peripheral cities from the network. Spatial analysis of network failure has also been done for airline networks finding similar results for the Indian airline network (Cliff et al 1979).

The Spread of Malicious Attacks in Complex Networks

The scale free structure of the Internet also has implications for how malicious attacks like worms and viruses are propagated throughout the network. Viruses and worms are not trivial computer nuisances, but high-cost problems:

> By the end of August, the cost of virus attacks in 2001 totaled nearly $10.7 billion, according to researchers at Computer Economics. In previous years, computer viruses have done quite a bit of financial damage, the group says. During 2000, virus attacks cost an estimated $17.1 billion, with the Love Bug and its 50 variants doing about $8.7 billion worth of harm. And in 1999, the estimated damage was reported to be $12.1 billion…Code Red accounted for $2.6 billion in damage -- $1.5 billion in lost productivity and $1.1 billion in clean-up costs (Lawyer 2003).

The high cost of virus and worm attacks on the Internet and connected businesses highlights the importance of understanding the nature of how these attacks spreads and what steps might be taken to mitigate them. The scale free and power law nature of the Internet illustrated by Barabasi and Albert (2002) and Faloutsos et al (1999) point to a

methodological framework for examining the issue. Analysis of epidemics in scale free networks first reported by Pastor-Satorras and Vespignani (2001), found that a wide range of scale free networks did not have an epidemic threshold[1]. The lack of an epidemic threshold meant that infections would persist and spread irrespective of the rate of the infections, however, the outcome is dependent on particular structure and topology of networks (21). This in theory could explain why viruses are rarely eradicated from the Internet and tend to spread quickly even when injected from peripheral places. Pastor-Satorras and Vespignani (2002) extended this work examining immunization of complex networks including an empirical test of the Internet at the AS level. In their test a SIS (susceptible-infected-susceptible) model was implemented, where half of the nodes in the network were infected and then nodes were immunized and the effect on infection rates were recorded. They found that targeted immunizations performed significantly better than uniform immunization strategies.

Dezos and Barabasi (2002) directly addresses the prospects of stopping such viruses, finding that traditional methods did not succeed is slowing spreading rates or eradicating viruses. The authors' instead found that selectively protecting the most connected nodes in the network could restore an epidemic threshold and "potentially eradicate a virus" (p.1). The study also points that a policy approach based on a "protect the hubs" strategy is cost effective, expending resources on only a few targeted nodes (Dezos and Barabasi 2002, p.3). The Dezsos and Barabasi (2002) study, based on theoretical models instead of empirical data, leaves some question of how effective their strategy would be with actual networks. A recent study by Newman et al (2002) studied

[1] The epidemic threshold is the point at which the percentage of unvaccinated people is high enough to risk an epidemic.

a 16,881-user email network to determine how viruses would spread across the network. While the structure of the network was not the power law distribution seen in the theoretical scale free models discussed above, the network's exponential distribution still reacted similarly to the predicted models. Protecting the most connected email users (in the form of anti-virus software or other measures) in the network had significantly better results than randomly protecting users across the network.

The collective work on the nature of complex networks and spread of worm/virus points to a possible fruitful approach for policies that could help provide greater cybersecurity. Questions, though, still remain as to: how a "protect the hubs" strategy would play out across the Internet as whole and what level of protection would be needed to gain the maximum level of security with the minimal level of investment. Is there a distinctive phase transition where protecting a certain percentage of nodes results in a big jump in overall network security? Further, considering the global nature of the Internet can any one country implement policies that would affect enough of the network to make an appreciable impact on global network security.

Public Policy and Cybersecurity

A negative externality of scale free network structure is security, the literature of complex networks points out that certain network formation, like scale free networks, are very vulnerable to targeted attacks. This has direct implications for critical infrastructure protection, which largely revolves around the protection of national networks. Hunker (2003) states that the policies required for critical infrastructure are intrinsically tied to understanding the complexity and interdependencies of the nation's networks. Part of

this understanding is the role of geography and what role it plays in the complexity of network interdependency and the policy implications (Rinaldi et. al. 2001). If the properties that make complex networks susceptible to attack also manifest themselves at a geographic level this will be an area of research interest with many policy implications. All ready cybersecurity and critical infrastructure have attracted considerable policy attention. Most notable have been the Federal Reserve, the Securities and Exchange Commission (SEC), the Office of the Comptroller of the Currency (OCC), and the State of New York Banking Dept. issuance of the "Draft Interagency White Paper on Sound Practices to Strengthen the Resilience of the U. S. Financial System" and the Whitehouse Critical Infrastructure Protection Board's February 2003 release of "The National Strategy to Secure Cyberspace". As the nation's strategy to develop policy addressing critical infrastructure an understanding of network

"The National Strategy to Secure Cyberspace" has received criticism from several venues specifically the lack of three features:

1. An assessment of the threat and the cost of inaction.
2. A link between the policies objectives and incentives
3. The strategy rejects regulation, government standards, and use of liability laws in addressing the cybersecurity issue (Berkowitz and Hahn p. 6 2002).

It is argued that the lack of these three features prevents any kind of cost benefit analysis from being done. Further, that without a cost benefit analysis there is no threshold for determining which of the activities and plans in the strategy are worth investing in or

pursuing (Berkowitz and Hahn p. 6 2002). Without an analysis of the range of policy options the best strategy cannot be determined. Hunker (2002) deals specifically with the range of public policy interventions that could deal with global information infrastructure security. Specifically he outlines:

1. Market forces
2. Regulation
3. Tort liability and contracts
4. Voluntary standards and best practices
5. Insurance
6. Public disclosure
7. Reputation/ratings
8. Procurement

In regards to these options Hunker also states that the large sunk costs in information infrastructure by network providers makes them unwilling to invest in implementing cybersecurity measures. Further, network providers security is dependent on the security of all other interconnected networks, resulting in a prisoner's dilemma situation (Kunreuther et al 2002). The National Strategy argues that policy interventions like regulation, standards, and liability would impede innovation and hamper international competitiveness. The plans strategy and counter veiling arguments remain to be empirically tested. What follows is a proposed methodology and preliminary results for evaluating the costs and benefits of different defense strategies and policy initiatives.

Methodology

In order to accomplish this task the problem will be simplified and examined at a macro level. The approach will be to examine the AS level topology of the Internet to determine what minimal level of protection will be required to protect the overall health of the network and prevent the wide scale spread of a generic malicious infections. Most malicious infections spread by email or random IP address probing and most infection simulations run at this level of analysis. The problem with this approach at a macro policy and economic level is that there are no economic entities that can be dealt with. The email end user or IP address holder are difficult to identify and even worse to coordinate for a protection effort. While the AS level is not technically accurate it does provide a set of economic actors that can be identified and allow for testing of the policy and economic implication of different protection strategies. The focus here is not on recreating realistic malicious attack scenarios but testing realistic policy strategies. Unfortunately we have not been able to locate data or techniques to fuse the two together, but we believe realistic worm simulations have been done (Moore et al 2003, Kephart and White 1991, Zou et al 2002), but not realistic policy analysis.

For this analysis, AS level nodes will be selected for protection (i.e. when a malicious infections encounters the node it will not become infected or pass along the infection to others), and the protection strategy will be tested to see how it affects the spread of an infection in the network. This approach looks at proactive measures, specifically where to invest limited resources, to stop malicious infections, instead of previous research that examines containing or eliminating existing worms (Pastor-Satorras and Vespignani 2002, Moore et al 2003). Considering the rapid spread of the SQL Slammer worm this could be a worthwhile path of investigation. The AS's selected

for protection will be first random and second based on their connectivity in the network. The threshold of AS's needed for protection will then be tested to determine at what point an acceptable level protection has been achieved. If the "protect the hubs" strategy proves a prudent strategy, further tests will be employed to determine what percentage of hubs is required for a least effort strategy to provide an adequate level of cybersecurity. Since it was not possible to acquire cost data for protection least effort is simply defined as the minimal number AS's that need to be protected. It is assumed there would be a wide variation in cost depending on the size of the AS. Further, how these AS's would be protected will not be endeavored, and the non-realistic assumption of 100% protection will first be assumed.

Each node in the network analyzed will be an individual autonomous system connected to the Internet. The data for this analysis was obtained from the University of Michigan's Internet Topology Project[2] and is based on data extracted form Oregon Route views on September 30th, 2001 and consists of 11,955 individual autonomous systems. The AS data was then analyzed utilizing two different approaches a weak worm and a strong worm. Worm, in this case, is just a generic term for a malicious infection that affects the Internet at a network level as opposed to a virus, which typically is transmitted through email. The simulation is intended to look at how infections spread from one firm's network to another and not at the IP address level that worms have used to propagate in the past. The weak worm and strong worm will both be run with a "protect the hubs strategy" with the most connected node being protected first, the next most

[2] This project is supported in part by NSF Grant No. ANI-0082287, by ONR Grant No. N000140110617, and by AT&T Research.

protected second, and so on. For purposes of simplicity the protected nodes in these simulations will be referred to as the "core".

When the weak worm is run a node is randomly chosen and all of its neighbors are infected. Next one of those infected neighbors is randomly chosen and all of its neighbors are infected and the process is repeated until the infection refers back to the originating node. The worm takes a random walk across the network, infecting all the neighbors of each node in its walk. The strong worm on the other hand infects all of the neighbors instead of just selecting one node to follow. This allows the worm to infect all AS's in a rapid manner when no protection is in place. To manage the strong worm computationally a queue approach was used where the neighbors of the originally infected node are put into a queue and infected in turn. As the worm spreads each neighbor's neighbors are put into the queue and infected as well. This way the length of the queue, nodes to be infected, can be plotted along with the total number of infected nodes, total number of attempts to infect nodes per cycle, and the number of times the protected core is visited per cycle. A cycle is simply a single simulation run with *n* number of nodes protected. The output produced by the simulation takes the worst-case infection scenarios from 15 iterations of each cycle. The results from the weak worm strong worm, and random weak and strong worm are presented below as Figures 1 to 4.

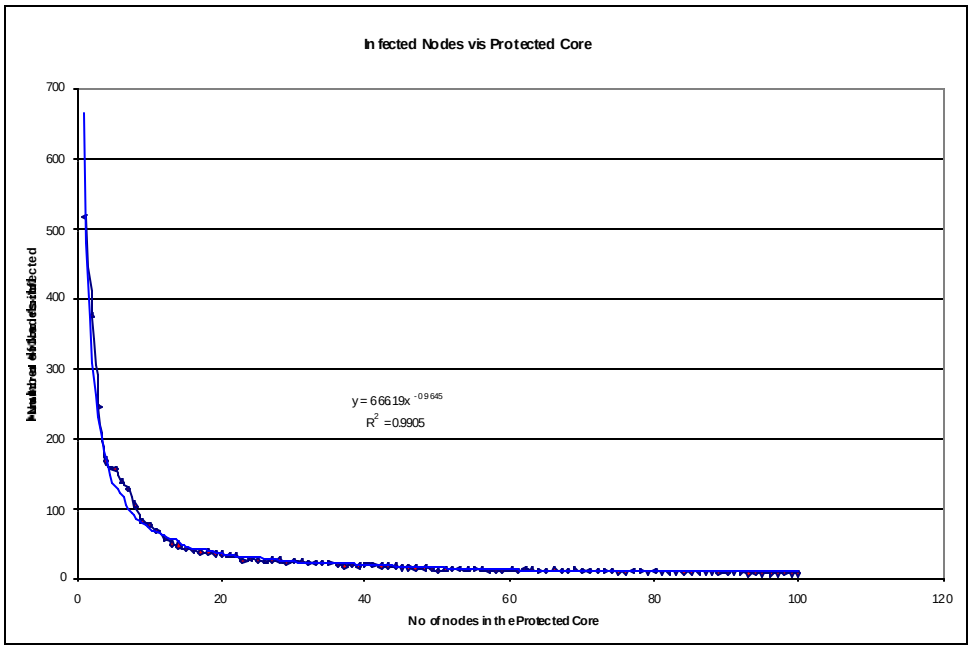

Figure 1: The number of nodes protected versus size of worst-case infected cluster.

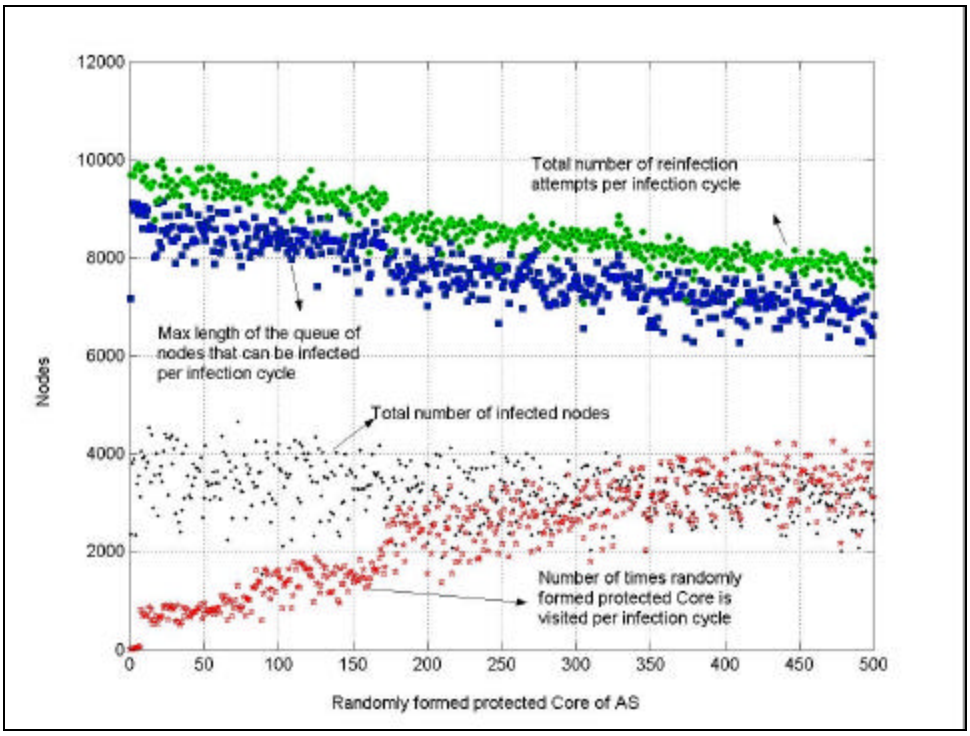

Figure 2: The number of nodes protected versus size of infected cluster with random protection strategy and weak worm

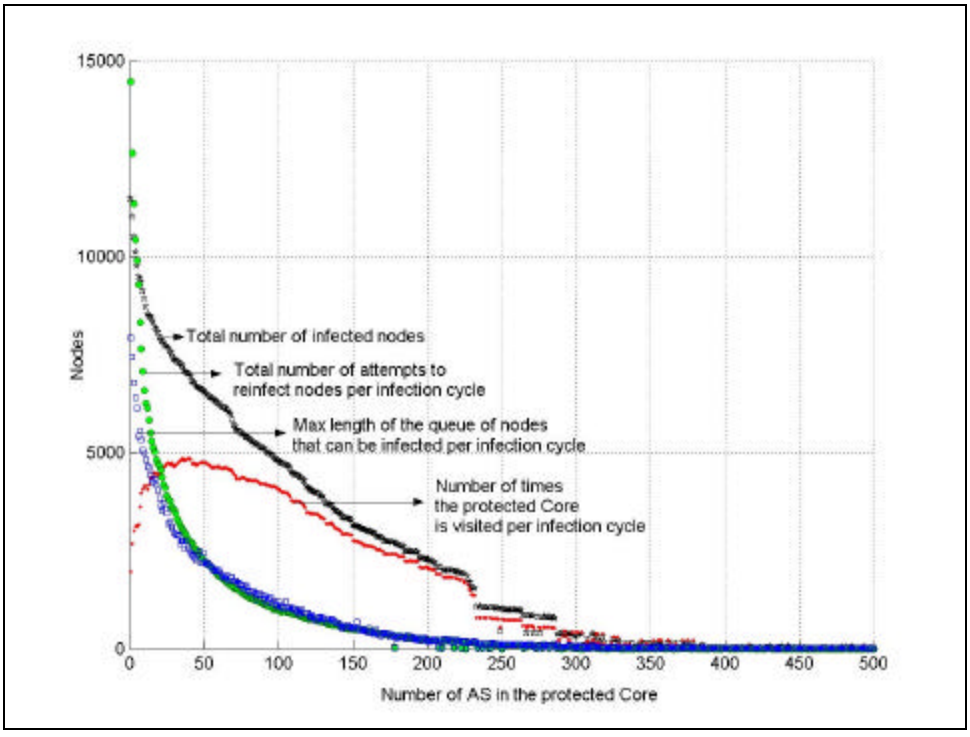

Figure 3: The number of nodes protected versus worst-case size of infected cluster over 15 iterations.

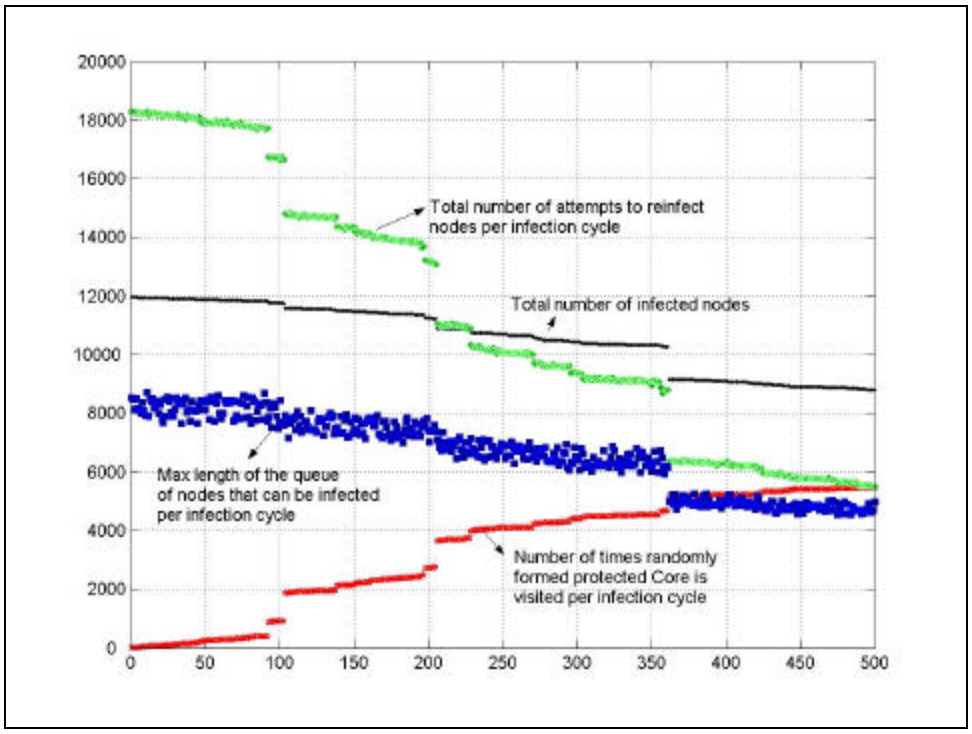

Figure 4: The number of nodes protected versus worst-case size of infected cluster with random protection strategy

This approach allows the testing of how increasing the size of the protected core affects the spread of virus/worm across the network of AS's. The results of the weak worm in figure 1 illustrate a precipitous drop in infection clusters with the first few nodes protected, and after about 20 nodes protected the change in infection cluster is relatively small. The sharp shift is indicative of the tight power law fit (R^2 = .9905) when the data is placed in a log-log format. The results seem to indicate a distinct point of inflection where there are diminishing returns for further investment in nodal protection. When this is compared to a random protection strategy with the weak worm the results are dramatic. The weak worm infections under the random strategy illustrate a random distribution with little noticeable decrease even after 500 nodes are protected.

The result of the strong worm (Figure 3) does not show the sharp power law decline seen in the weak worm, but there is still a definitive point where the infection drop off drastically (node 229) and then becomes largely ineffective (node 275). While the number of nodes requiring protection to contain the worm is larger than the weak worm, the total nodes protected are only 2.3% of the total network.

In comparison, when a random protection strategy is implemented (figure 4) little to no protection is afforded even after 500 nodes are protected. Under the random protection strategy, protecting even 500 nodes results in 8789 infected nodes, 72.3% of the total network. The one large drop in the results is because AS 701 (Worldcom) the most connected node in the network was randomly chosen.

The results of this study and particular simulation approach illustrate a significant improvement in security from a "protect the hubs strategy", although the strategy becomes less effective as a worm is made more potent. A needed extension of this work is to investigate scenarios where protection is not 100% and some worms find their way through an AS's defense. This would provide an additional level of testing for the effectiveness of the strategy.

A Cost Effectiveness Analysis of Cybersecurity Strategies

One of the problems with the analysis proposed so far is that it considers the cost of protecting all autonomous systems as being equal. The cost of protecting a large global network like UUnet with 2451 connections and a local University with one connection are treated as the same. To remedy this problem a simple cost function will be introduced, where the cost of protecting a link connecting two autonomous systems

will be $1. Unfortunately better cost figures for Internet malicious attacks are not currently available. Several industry experts were consulted but approaches and cost figures varied widely. Chen (2003) calculated cost and benefits for protection from denial of service attacks, based of the frequency of attacks, number of subscribers, and AT&T's network charge to customers for security services. While the metrics are not appropriate to the analysis here, they did, interestingly, conclude that providing DOS protection services would be profitable strategy for ISP's (Chen 2003).

Under the $1 per link cost structure the total cost of protecting a network will then be the total number of links multiplied by $1. From the prior analysis of security strategies the most successful will be taken and the total cost of that strategy calculated. Suppose that the total budget for securing nodes/links is B $, and assume that it takes a fixed cost of C $ to secure each link connected to a set of nodes, then the total number of nodes/links that can be secured with budget of B $ is given by the following relation:

$$\sum_{i=1}^{n} L_i \times C \leq B \qquad (1)$$

Where, L_i , is the number of links for associated with node i, and n is the total number of nodes. Once the cost of the strategy has been determined the return on investment of that strategy will be compared to other strategies. Once a set of protected nodes has been defined for each strategy, the strong worm algorithm will be run for each. Once the algorithm has been run the following data will be collected:

- Number of infected nodes
- Number of non-infected links

- Cost effectiveness ratio
- Number of protected nodes
- Number of protected links
- The protected link multiplier = number of protected nodes / number of protected links
- The integrity of the network after node failures are accounted for

This data will allow an examination of the return on investment for different allocation of resources to protect the network. The money invested will be the same but the outcome will differ by the investment strategy. When the cost for the "protect the hubs" strategy was calculated with the top 276 nodes protected, using the equation outlined above, the total cost of protection was $19,544. The testing of the "protect the hub" strategy illustrated it as having maximum benefit with the least number of nodes protected. Now that the cost of the "protect the hubs" strategy has been calculated a comparison between cost and benefits can be made. First a comparison between the "protect the hubs" strategy and the random strategy will be done. The simulation will first protect the most connected node, then the second most connected node and so on until the allotted funds ($19,544) have been spent. Next nodes will be randomly selected and protected until the allotted funds have all been spent. Finally the malicious attack simulation will be run to see which strategy resulted in a better return on investment for the allotted funds (table 1). The malicious attack simulation was run thirty times and the worst-case scenario is presented in the results.

Strategy	Protected Nodes	Non-infected Nodes	CE[3] ratio	% Infected	Protected Paths	Path Multiplier	Network Integrity
Random	4637	7220	$2.71	40%	18790	4.1	No
Protect the Hubs	276	11186	$1.75	7%	41274	149.5	Yes

Table 1. Cost Effectiveness Analysis of Random and Hub Protection Strategies

Adding a cost component to the analysis illustrates that the "protect the hubs" strategy produces a better return on investment. For the same monetary investment in protection the "protect the hubs" strategy resulted in only 7% of the network being infected, and for the 276 nodes protected 41,274 paths to other nodes were defacto protected, a roughly 150x protection multiplier. This compares to the random strategy that allowed 40% of the network to be infected and only a 4x protection multiplier. Further the protection cost per node for the targeted strategy was only $1.77 compared to $2.71 for the random strategy. Also it is important to note that with 40% of the network compromised it becomes disconnected and large parts of network cannot communicate with each other, thus the network has lots its integrity.

While the "protect the hubs" strategy does better than a random strategy, are there other strategies that might outperform it? To test the comparative value of the "protect the hub" strategy three other strategies were devised for testing:

- Bottom tier – the least expensive node in the network to protect is chosen, followed by the next least expensive node, and so on until the total funds allotted for protection are consumed.

[3] Cost effectiveness ratio provides the cost per non-infected node in the network

- Middle tier – the first node not protected by "protect the hubs" strategy is selected for protection, followed by the next most expensive node to protect, and so on until the total funds allotted for protection are consumed.
- Mutli-tier – a node is randomly chosen from the "protect the hubs" strategy, then another is randomly chosen from the bottom tier, then another chosen randomly from the middle tier, and so on in that order until the total funds allotted for protection are consumed.

The same malicious attack simulation is run with all three new strategies, with thirty runs, selecting the worst-case scenarios for comparison. The simulation runs produced the results seen in table 2.

Strategy	Protected Nodes	Non-infected Nodes	CE Ratio	% Infected	Protected Paths	Path Multiplier	Network Integrity
Protect the Hubs	276	11186	$1.74	7%	41274	149.5	Yes
Bottom Tier	11028	11030	$1.77	8%	7058	0.6	No
Middle Tier	6861	8120	$2.41	32%	9808	1.4	No
Multi Tier	786	10331	$1.89	14%	37456	47.7	Yes

Table 2. Cost Effectiveness Analysis of Competing Cybersecurity Strategies

The "protect the hubs" strategy still performs the best of the group, but the bottom tier strategy also produces seemingly good results. The bottom tier strategy results in only 8% of the network being infected, just 1% higher than the "protect the hubs" strategy and a low $1.77 cost per node for protection. This "protect the small guy" strategy appears to give good results, but a closer inspection illustrates the shortcoming on this strategy.

While the bottom tier strategy directly protects a huge number of nodes, 11,028, it leaves the most connected nodes vulnerable. This is critical because the most connected nodes are what allow the least connected nodes to communicate with each other. The least connected nodes rarely link directly to each other. Thus, with the loss of the top most connected nodes the network is balkanized to an incredible degree, this is best seen in a before and after visualization of the network send in Figure 5.

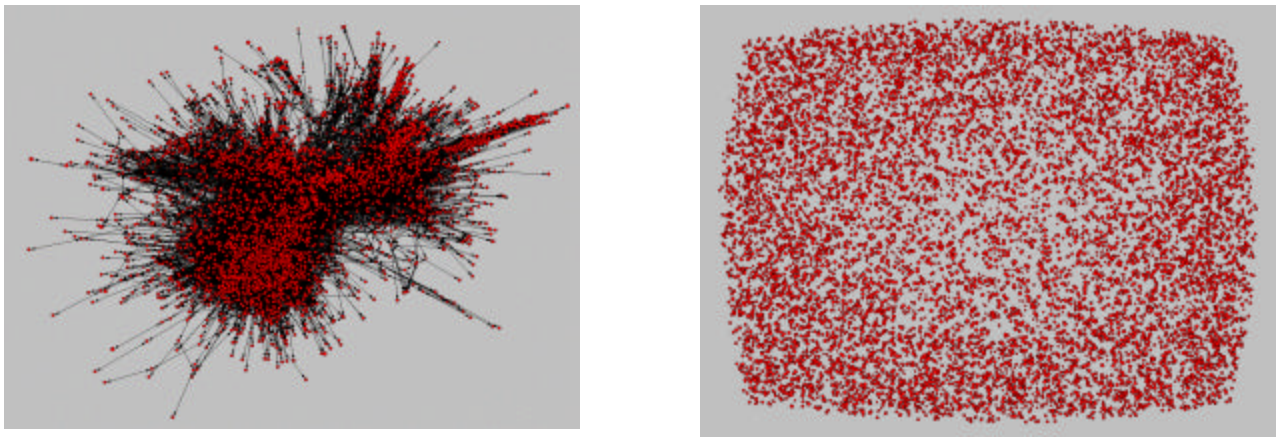

Before After

Figure 5. Before and After Network Visualization of the Bottom tier Strategy

Perhaps the most interesting result was the relatively good performance of the multi-tier strategy. Utilizing random selection from the other three strategies, the multi-tier approach resulted in only 14% of the network being infected, an almost 50x protection multiplier, relatively low $1.89 protection cost per node, and results in a network that does not loose its integrity after attack. This is important because the "protect the hubs" strategy depends on total cooperation of all the top nodes. The difficulty of this requirement will be illustrated in the next section.

The Impact of Policy on Cybersecurity Strategy Implementation

Their remains another important facet for this research assessing the impact of policy on the implementation of cybersecurity strategies. How many AS's in the network would fall under US jurisdiction and would a US policy affecting just US firms be enough to obtain a reasonable level of security? In order to begin examining this issue the top 350 AS's registered with US addresses were compiled. Next a protection scenario was run with the most connected US firm protected followed by the next most connected through the top 350 using the strong worm methodology outlined previously. The results of the strong worm algorithm with a US only protection strategy was then plotted over the same procedure using the top 350 global AS's – the result can be seen in Figure 6.

The results illustrate similar levels of protection with the US only strategy (blue) through the top 75 AS's, then the global AS strategy (red) begins to protect more AS's than the US only strategy. After the top 100 AS's the US only strategy flat lines with over 5,000 AS's still being infected, while the global strategy continues to decrease and largely contains infections around 275 AS's protected. The result makes a preliminary case that a US only policy strategy could be an inadequate measure. International cooperation or policies that can influence foreign firm's security policies appear to be needed.

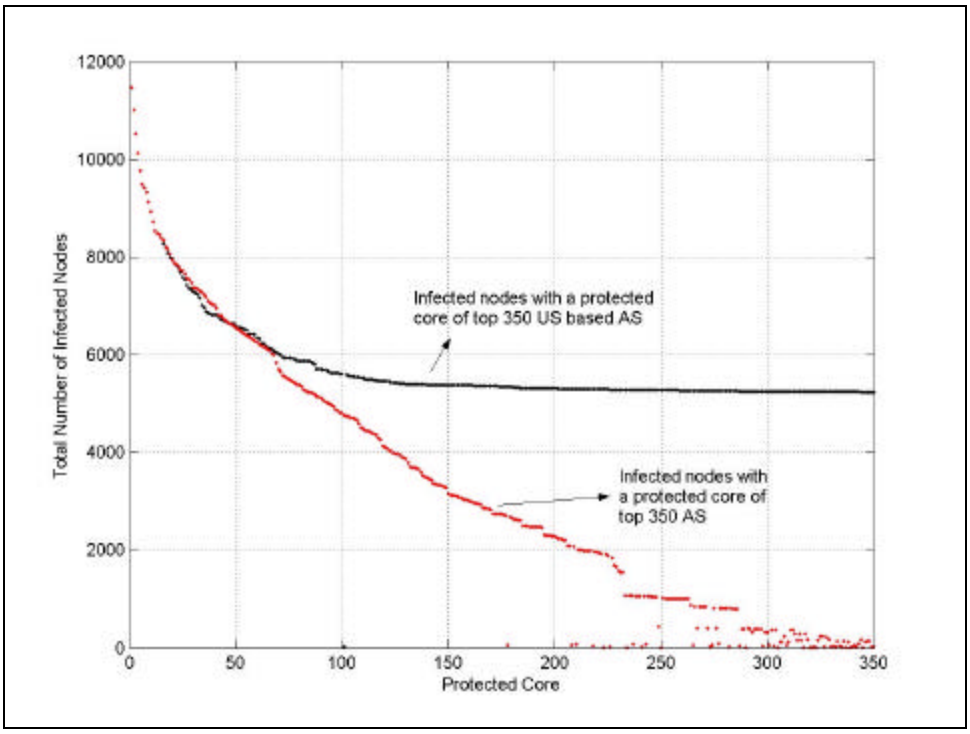

Figure 6: US vs. Global protection with the strong worm algorithm.

Policy Analysis and Implications

The result of the "protect the hubs" strategy has implications for policy in regards to best approaches to cybersecurity and critical infrastructure protection. Several studies have pointed out the fragility and vulnerability of the Internet to malicious attack (Albert et al 2000, Callaway et al 2000, Cohen et al 2001). There has been debate as to the best policy approach to the current security shortcomings of the US's information infrastructure. It has been offered that there are several options for dealing with the current situation ranging from regulation, market forces, contract law, standards/best practices, insurance, or government mandated procurement requirements (Hunker 2002). While a full discussion of all the possible interventions for security goes outside the

scope of this paper the results can shed some lights as to which general directions might bear the most fruit.

Perhaps the most persuasive argument from the results is that universal regulation is most likely an excessive approach to the problem. At the same time an uncoordinated approach fostering random protection appears to be largely ineffective. In the case of telecommunications, industry wide regulation has been most often justified in the quest to provide universal service for a population (Dinc et al 1998, Tarjanne 1999). The results of this analysis illustrates that the universal protection theoretically offered by regulation of the Internet produces minimal returns in relation to the effort to protect all the networks connected to the Internet. In fact returns diminish significantly after the protection of the top 20 nodes in the network with a weak work and the top 275 nodes with a strong worm, which constitute only .17% and 2.3% respectively of total nodes in the network.

Further questions still remain, how many firms control the top AS's? The top 20 network providers control over 70 AS's, so it is likely there are not 275 separate firms to deal with. It was not possible to perform this analysis for this paper, but it is an important extension of this research that is under investigation.

The results presented in these simulations, in term of percentages, can be deceiving, while it only requires protection of 2.3% of total nodes to obtain a high level of security the cost of protecting nodes is not equal. The most connected nodes in the network are large global networks like MCI, Sprint and AT&T. The cost of securing global networks of this size are significant and dwarf the cost of securing smaller campus networks. Needless to say protection of 2.3% of nodes would not equal 2.3% of costs.

Further, the small number of firms represented by the top 20 or top 275 AS's would seem to indicate that public – private partnerships or selective regulation to address the problem would be beneficial. The difficult task is ensuring that as many of the top ASs' are protected as possible. Even with just the non-US networks removed the level of protection is significantly reduced. Also it remains to be seen if market forces or even public-private partnerships can provide adequate coverage of the top AS's. Selective regulation of the top AS's could ensure coverage but questions of equity and hampered competition and innovation could arise. Several of the alternative approaches delineated by Hunker (2002) could be answers to the dilemma. For instance, if the US government has contracts with a significant number of these networks basic security requirements built into RFP's could cover a large number of firms and also provide economic incentive for compliance in the form of increased service fees to cover upgrades.

Financial Sector versus the Federal Sector

Hunker's (2002) scenario can be tested with the framework to examine how effective the strategy would be. Utilizing the AS graph all federal government networks can be identified, then all their nearest neighbors calculated. The nearest neighbors are all networks that interconnect with federal government networks to provide connectivity services and would fall under the RFP regulation. Once this core has been identified the malicious attack scenario can be run and the cost effectiveness of the policy evaluated. In order to provide a comparison a financial sector policy will be simulated as well. In this scenario a policy will be simulated where the network providers to all US financial institutions would have to meet a minimal level of security in order to obtain contracts. The same approach as used with the federal RFP sector will be employed where all US

financial institutions in the US will be identified along with their nearest neighbors. The financial institutions and their nearest neighbors will constitute the protected core and the malicious attack simulation will be run. The results will be used to compare the cost effectiveness of the federal RFP policy to the financial institution policy. Theses results are presented in table 3.

Policy	Protected Nodes	Non-infected Nodes	Cost of Protection	CE Ratio	% Infected	Protected Paths	Path Multiplier	Network Integrity
Fed RFP	390	4,262	$ 9,120.00	$ 2.14	64.5%	16,940	43.44	Yes
US Financial	488	5,787	$ 7,480.00	$ 1.29	51.8%	22,086	45.26	Yes
World Financial	748	6,849	$ 7,966.00	$ 1.16	42.9%	24,308	32.50	Yes

Table 3: Cost Effectiveness Comparison of a Federal RFP versus a Financial Cybersecurity Policy

The results of the simulations demonstrate that the financial policy does a better job protecting the network and is more cost effective. The simulation results indicate that he financial policy results in 12.7% less of the network being infected in an attack than the federal policy. Further, the financial policy only costs $1.29 per node to implement while the federal policy costs $2.14 per node to implement. The likely reason for this is two fold; (1) the financial sector has more connectivity 121 connections to other network versus 79 for the federal policy, and (2) more of those connections are to foreign networks. The federal core connects to 18 (22.8%) foreign-based networks, including networks based in Canada, Netherlands, Japan, Singapore, Korea, China Taiwan, Jordan and Russia. The financial core, on the other hand, connected to 36 (30.0%) foreign-based networks, Australia, United Kingdom, Netherlands, Australia, South Africa, Italy,

Switzerland, Greece, Spain, Canada, Korea, Japan, India, Korea, Malaysia, Chile, Slovenia, Bulgaria, Russia, Romania, Philippines, and Panama. Also, the financial network connects to most OECD[4] countries, which the federal network does not. The third scenario in the table further tests the international hypothesis by including foreign financial institutions and their nearest neighbors in the financial network. The inclusion of foreign financial institutions not only decreases the level of infection by 9% it also decreases the per-unit cost of the protection strategy $1.16. This simulation is addition to the US versus non-US scenario make a strong argument for the value of international cooperation in cybersecurity efforts.

Recently there has been new theoretical support for why the federal and financial strategies could work well as cybersecurity policies. Cohen, Havlin, and ben-Avraham found that an acquaintance immunization strategy was effective method for protecting computer networks and populations with broad degree distributions. The acquaintance immunizations strategy randomly selects a group of nodes for immunization along with all their nearest neighbors (acquaintances). When the acquaintance strategy was tested it was found to produce a low immunization threshold for populations with a broad degree distributions. A low immunization threshold means that only a relatively few numbers of nodes in the network needed to be immunized to prevent the infection from spreading throughout the network. Broad degree distributions characterize a wide variety of real work networks including the Internet, the electric power grid, the World Wide Wed, friendship networks, email networks, and air transportation to name a few. The implications of the study are useful for many different applications.

[4] Organization for Economic Co-operation and Development

The authors focus specifically on human populations and computer networks, finding the biggest advantage of their strategy resulting from not needing to know global properties of the network. In order to effectively immunize the network one does not need to know who the most connected nodes are, or any other aspect of the networks structure for the strategy to be implemented. This facet has obvious benefits for human populations fighting epidemics like AIDS. It is very difficult to identify which members of the population have the greatest number of sexual partners for both pragmatic and privacy reasons. In computer networks the advantage is less obvious, the topology of most all computer networks is well known, or can be easily acquired. As this paper has demonstrated, though, getting all nodes in a global network to cooperate can be very difficult for a variety of economic and political reasons. In any heterogeneous network environment cooperation of all nodes can be difficult to achieve. The acquaintance strategy offers a method for overcoming such cooperation hurdles, and also explains why the financial and federal strategies in the previous section work far better than random strategies. The federal and financial strategies are in essence an implementation of the acquaintance strategy. Federal and financial AS's are randomly distributed across the AS graph, but are all likely to connect to hubs for transit across the network. When the federal or financial policy leverages its scope to include their big hubs in a protection strategy to overall resiliency of the network is dramatically increased.

To conclude it appears from these crude simulations that targeted strategies and efforts that emphasize international cooperation will result in more cost effective cybersecurity strategies. This approach could be refined to simulate a cooperate or defect

choice with different policy scenarios through an agent based model with the data presented in this study, and this is a future direction of research. Finally, sector specific cybersecurity policies could have significant effect in increasing overall resiliency of the Internet to malicious attack through the positive externalities of acquaintance immunization strategies.

REFERENCES

Albert R. and Barabási, A., 2002, Statistical mechanics of complex networks, Reviews of Modern Physics 74: 47-97.

Albert R., Jeong H., and Barabási A.L., 2000, Attack and error tolerance in complex networks, Nature 406: 378.

Callaway, D.S., Newman, M.E.J., Strogatz, S.H., and Watts, D.J., 2000, Network robustness and fragility: percolation on random graphs, Physical Review Letters 85:25.

Chen, L., 2003, Computational Models for Defenses against Internet-based Attacks Unpublished Dissertation, Carnegie Mellon University.

CIPB, 2002, The National Strategy to Secure Cyberspace Washington, DC: Whithouse Critical Infrastructure Protection Board - http://www.whitehouse.gov/pcipb/.

Cliff A., Haggett P., and Ord K.,1979, Graph theory and geography. In: Wilson R. and Beineke L. (Eds) Applications of graph theory, London: Academic Press.

Cohen, R., Erez, K., ben-Avraham, D., and Havlin, S, 2001 Breakdown of the Internet under intentional attack, Physical Review Letters 86:16.

Cohen, R., Erez, K., ben-Avraham, D., and Havlin, S, 2003 Efficient immunization strategies for computer networks and populations, Physical Review Letters 91:24.

Dezsos, Z., Barabasi, A.L., 2002, Halting viruses in scale-free networks, Physical Review E 65: 055103 (R).

Dinc M., Haynes K.E., Stough R.R., and Yilmaz S., 1998, Regional universal telecommunication service provisions in the US — Efficiency versus penetration, Telecommunications Policy 22 (6): 541-553.

Faloutsos C., Faloutsos P., and Faloutsos M., 1999, On power-law relationships of the Internet Topology, Computer Communication Review 29: 251.

Gao, L., 2001, On inferring autonomous system relationships in the Internet, IEEE/ACM Transactions on Networking 9 (6): 733.

Gorman, S.P. and Kulkarni, R., Forthcoming, Spatial small worlds: New geographic patterns for an information economy, Environment and Planning B.

Gorman, S.P., Schintler, L.A., Kulkarni, R.G., and Stough, R.R., forthcoming, The revenge of distance: Vulnerability analysis of critical information infrastructure, Journal of Crisis and Contingency Management.

Grubesic, T.H., O'Kelly, M.E., and Murray, A.T., (2003) A geographic perspective on telecommunication network survivability, Telematics and Informatics 20 (1): 51-69.

Hunker, J., 2002, Policy challenges in building dependability in global infrastructures, Computers & Security 21 (8): 705-711.

J. O. Kephart and S. R. White, 1991, Directed-graph Epidemiological Models of Computer Viruses, in Proc. of the 1991 IEEE Computer Society Symposium on Research in Security and Privacy 2, 343–359.

Lakhina, A., Byers, J.W., Crovella, M., and Matta, I., 2002, On the Geographic Locations of Internet Resources, http://www.cs.bu.edu/techreports/pdf/2002-015-internet-geography.pdf.

Lawyer, G., 2003, The battle of the bug: Government, industry move to protect Internet from cyber attacks, viruses, http://www.xchangemag.com/articles/1B1front4.html.

Magoni, D. and Pansiot, J.J., 2001, Analysis of the autonomous system network topology, Proceedings of ACM SIGCOMM'01.

Moore, D., Paxson, V., Savage, S., Colleen, S., Staniford, S., and Weaver, N, 2003, The spread of the Sapphire/Slammer worm, CAIDA - http://www.caida.org/outreach/papers/2003/sapphire/sapphire.html.

Moore, D., Shannon, C., Voelker, G.M., and Savage, S., 2003, Internet quarantine: Requirements for containing self-propagating code, INFOCOM 2003 http://www.caida.org/outreach/papers/2003/quarantine/.

Moreno, Y., Vazquez, A., 2003, Disease spreading in structured scale free networks, The European Physical Journal B 31: 265-271.

Newman, M.E.J., Forest, S., and Balthrop, J, 2002, Email networks and the spread of computer viruses, Physical Review E 66: 035101(R).

NRC, 2002, Cybersecurity Today and Tomorrow: Pay Now or Pay Later Washington, DC: National Academy Press.

NSTAC, 2002, Network Security/Vulnerability Assessments Task Force Report Washington, DC: The President's National Security Telecommunications Advisory Committee, http://www.ncs.gov/nstac/NSVATF-Report-(FINAL).htm.

Pastor-Satorras, R., and Vespignani, A., 2002, Immunization of complex networks, Physical Review E 65: 036104-1.

Pastor-Satorras, R., Vespignani, A., 2001, Epidemic dynamics and endemic states in complex networks, Physical Review E 63: 066117.

Tangmunarunkit H., Govindan R., Jamin S., Shenker S., and Willinger W., 2002, Network topologies, power laws, and hierarchy, Computer Communication Review 32 (1): 76-76.

Tarjanne, P., 1999, Preparing for the next revolution in telecommunications: implementing the WTO agreement, Telecommunications Policy 23 (1): 51-63.

Yook SH, Jeong H, Barabási AL, 2001, Modeling the Internet's Large-Scale Topology, http://xxx.lanl.gov/abs/cond-mat/0107417.

Zou, CC., Gong W., and D. Towsley, 2002, Code Red Worm Propagation Modeling and Analysis, in 9th ACM Conference on Computer and Communication Security 11, 121-138.