Epidemic Profiles and Defense of Scale-Free Networks

Linda Briesemeister, Patrick Lincoln, Phillip Porras
SRI International
333 Ravenswood Avenue, Menlo Park, CA 94025, U.S.A.
firstname.lastname@sri.com

ABSTRACT

In this paper, we study the defensibility of large scale-free networks against malicious rapidly self-propagating code such as worms and viruses. We develop a framework to investigate the profiles of such code as it infects a large network. Based on these profiles and large-scale network percolation studies, we investigate features of networks that render them more or less defensible against worms. However, we wish to preserve mission-relevant features of the network, such as basic connectivity and resilience to normal nonmalicious outages. We aim to develop methods to help design networks that preserve critical functionality and enable more effective defenses.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General—*security and protection*

General Terms

Algorithms, Security

Keywords

Scale-free networks, computer epidemics, self-propagating malicious code

1. INTRODUCTION

The escalating dependence in our nation on cyber infrastructure to control and transport valuable information has left many in precarious situations, overdependent on unreliable and nonsurvivable systems. The disturbingly frequent outbreak of malicious worms and viruses in the broader public Internet often penetrate into even well-protected enterprise networks, or cause major disruption through targeted or widespread denial of service attack. Thus we are motivated to study the problem of defending a large network infrastructure from rapidly propagating malicious code.

WORM'03, October 27, 2003, Washington, DC, USA.

We report here on preliminary research aimed at building a framework to help understand how to defend networks with certain properties from worms. The starting points of our research include three main points: (1) studies of worms and viruses and the strategies they employ to infect a given system, and more importantly, their strategies for propagation (how they target new systems to attack), (2) studies of percolation and epidemic spread in large networks exhibiting key properties observed in intranets such as scale-freeness, and (3) studies of preservation of mission critical functionality for a network.

Obviously, relative safety from malicious worm outbreaks is achievable through reduced network connectivity. In the extreme case of truly disconnected networks, worms cannot propagate from one partition of the network to another (never mind that many network partitions claimed to be "airgapped" suffered from worm and virus outbreaks originating from other partitions. Forensic evidence that the networks were connected in unauthorized ways suggest it may be impractical to rely on this method of defense.) If we presume networks are connected, defensibility considerations motivate extremely sparse connectivity between network partitions. However, random failures can, with unacceptably high probability, render such sparsely connected networks disconnected, potentially compromising the mission.

We presume the mission of a given network includes preservation of connectivity in some form (defined more formally later). We assume that random, uncorrelated network outages occur with a given probability P for all nodes. Finally, we assume that malicious rapidly propagating worms will be released into the network. Under these assumptions, we study network models and start to develop a framework to analyze defensibility of networks against these kinds of threats.

2. RELATED WORK

Moore et al. [1] study the spread of CodeRed and related worms through a large fraction of the actual machines infected by CodeRed, under various scenarios involving content blocking and address blacklisting. They find that under reasonable assumptions, no response time is fast enough to protect against widespread epidemic. Their model network topology is based on actual Internet route maps during the initial spread of CodeRed, but they eliminate redundant routes from end-user machines. Thus their results are conservative, and the actual threat we face remains dire.

Albert et al. [2] study error and attack tolerance in scale-

free networks. The metrics of interest concern the connectivity of networks (cluster size) after few nodes have been removed. The authors find scale-free networks robust against random error, but not against deliberate attack of highly connected nodes. The authors' approach has a dynamic component of studying the network after repeated removal of nodes, which could be seen as cascading failures.

Pastor-Satorras and Vespignani [3, 4, 5] and Eguíluz and Klemm [6] model spreading of infections under the susceptible-infected-susceptible (SIS) model in scale-free networks. While Pastor-Satorras and Vespignani use the BA-model (proposed by Barabási and Albert), Eguíluz and Klemm use their KE-model for describing scale-free networks. Due to different network models, the first author pair finds the absence of an epidemic threshold that determines prevalence, whereas the second author pair finds finite epidemic thresholds. In our work we study both the BA-model and KE-model of scale-free networks.

Eguíluz et al. [7] characterize fractal properties of complex networks. The authors derive analytical results, which agree with numerical results obtained from simulation of percolation in scale-free networks.

Dezsö and Barabási [8] study spreading of viruses under the SIS model in scale-free BA-model networks with randomly and deliberately (favor nodes with high degrees) distributed cures. While randomly placed cures are ineffective, policies protecting the hubs can restore the epidemic threshold.

Leveille [9] proposes a new epidemiological model PSIDR geared toward describing the behavior of computer worms. The author simulates PSDIR in different homogeneous and scale-free networks.

Newman et al. [10] employ a simple spreading algorithm with 100% efficiency on a graph model that they obtained from referencing email address books. The resulting graph is semi-directed and shows a strongly connected giant component. Using this real topology, the authors show that the outbreak size can be significantly reduced when removing up to 10% of nodes in decreasing order of their out-degree from the giant component.

3. EPIDEMIC PROFILES

We are concerned with network security management in defense of networks against fast moving malicious code epidemics. Success or failure of defense against malicious self-propagating code depends greatly on the availability of communication channels between susceptible nodes in the network.

Networks that are inherently defendable can, with relatively few alterations, prevent or significantly delay an infection from reaching its maximum saturation potential. Significant network segmentation, lack of communication channels among vulnerable nodes, and IP filtering to limit scanning all play a role in the rate at which various epidemics will find targets of opportunity.

Infection strategy, meaning the method by which the epidemic seeks new targets, can be susceptible to variations in network infrastructure such as the addition of content filtering and address blacklisting. The size of the set of susceptible nodes relative to the entire target space is an important issue in understanding spread rate, and mapping infection criteria to the network node configurations is another important element in profiling an epidemic. Homogeneous network computing environments may be highly resilient to some forms of malicious contagions, while fatally susceptible to others. Beyond a basic SI (Susceptible-Infected) infection model, it is critical to understand the alignment of vulnerability dependencies and host configuration.

Underlying this study is the notion that a specific epidemic may be profiled sufficiently to allow an assessment of the probability of susceptibility in end nodes. Further, knowledge of the epidemic behavior may influence the estimation of infection cost per node. In the future we plan to bring these notions together to develop epidemic profiles, which assist in assessing and simulating the inherent vulnerability or resilience of a network to known or hypothesized epidemics before those epidemics are encountered in the wild. The present paper is a first step in that direction.

3.1 Infection Criteria

Worms and viruses have relied upon a number of infection methods, including network service buffer overflows, macro and script insertion, deception of binary code like time-of-check-to-time-of-use (TOCTTOU) vulnerabilities, and argument-driven subversion. The enumeration of vulnerability dependencies relevant to successful infection include features such as target operating system, enabled network services, patch revisions, configuration settings, hardware architecture, and resident applications. Although numerous infection techniques may be applied against a wide set of vulnerabilities, experience has shown that malicious applications have typically employed a limited set of exploit techniques, often producing outbreaks that are highly OS or application specific.

However, the emerging "blended threat" attack mode illustrates how a single contagion may employ numerous techniques to infect a large heterogeneous population of hosts, effectively producing a large infection criteria set. For example, Nimda [11] exploited five major infection methods, increasing both its potential to propagate across even highly segmented heterogeneous computing environments, and increasing the overall cost of defense.

3.2 Infection Strategy

To date, there have been a number of infection strategies documented, many of which have been experienced in the wild. It is important to understand infection strategies when evaluating security posture and formulating course of action in the presents of an initial infection. For example, highly segmented networks with strongly limited external to internal availability may provide significant channels to topological worms, while imposing few constraints on mail or contagion-based attacks. Further, whether a contagion-based worm is propagating through a network service that is critical or extraneous to the network mission will influence the cost deemed acceptable for responses that block the channel.

A dominant strategy in worm and virus propagation has been to employ a sequential process of scanning, in which an infected host searches for a target victim, then propagating, where the infected machine launches an infection attempt against the discovered target, and iterating. Various techniques have been used to explore the target space:

- Mail-based – Such as the Melissa virus [12], employ mail services and user address book information to propagate

- Topological – Such as the Morris worm [13], leverage substantial internal topological information on each compromised target to seek additional new targets
- Contagion – "Surreptitious" malicious code propagation, which embed contagions within normal communication channels
- Active Scanning – such as CodeRed [14], perform self-propagation via random scanning to identify potential contagion targets
- Coordinated Scanning – An optimization of active scanning, employ efficient segmentation of IP address space to accelerate scan coverage, resulting in so-called Warhol worms [15], famous for 15 minutes

Recently, an alternate strategy has been explored in which the malicious application performs an extended scanning phase without immediately attacking susceptible targets. Rather, the application constructs a list of candidate susceptible targets, and when this list is complete, may enter an aggressive infection phase, in which all candidate susceptible targets are attacked. Flash worms, hypothesized by Staniford et al. [16], might employ this strategy together with rough synchronization of infected hosts to achieve nearly immediate pandemics in susceptible populations.

Single stage worms have also been developed, which consolidate the scanning and infection processes into a single stage. For example, Sapphire [17] demonstrated that malicious code propagation does not require the two-stage approach of scan and infect. Sapphire introduce a single packet propagation strategy which merged the scan and infection phases into a single UDP packet.

The speed of these attacks motivate the search for highly automated defensive measures, and the study of network topologies that are more defensible.

3.3 Epidemic Subgraph Partitioning

An epidemic subgraph represents the subgraph of a network in which all end node systems possess the attributes that satisfy the infection criteria, plus the intermediate infrastructure nodes that are within the traversal path of vulnerable end nodes employing the infection strategy. An epidemic tree is a subtree of an epidemic subgraph representing a time series of first infection events for each node, and their interdependencies.

For example, the Sapphire worm's infection criteria includes only computers running Microsoft operating systems, and running Microsoft SQL server 2000 or Microsoft SQL server desktop MSDE as enabled services (most home machines running SQL server do not enable Internet access to this service, and thus do not satisfy this infection criteria; however many applications silently install MSDE 2000) not running service pack SP3. The infection strategy uses UDP port 1434 and attempts to connect to randomly generated IP addresses, including broadcast addresses such as x.y.z.0. The epidemic subgraph of Sapphire would include all hosts matching the infection criteria and intermediate nodes passing UDP 1434 traffic.

4. EXAMPLE EPIDEMIC PROFILE

Self-propagating malicious code may span a range of capabilities that increase the complexity of the infection criteria and may in fact employ multiple infection strategies. To motivate the problem and illustrate our approach, we consider a malicious code attack that spans multiple attack methods, and can operate across a heterogeneous network.

Consider a DoD wide area network comprised of several local area networks (LANs), all connected via a private leased network. Each LAN contains a heterogeneous mix of many Windows workstations and a few critical Unix servers that host a minimal set of network services, such as DNS and SMTP. LANs enforce strong filtering restrictions at their gateways, while internally allow open connectivity policies. In this scenario, the attacker seeks to maximize the saturation of a malicious application across the organization, which given the large distribution of Windows machines in each LAN, would make a Windows exploit an important part of the attack.

In this imagined network, between different autonomous systems or LANs, Windows machines do not directly communicate. Rather, the primary ingress method for propagating across LANs will require the exploitation of intra-LAN network communications, such as leveraging the SMTP channel via a technique like the Outlook MIME vulnerability (CVE-2001-0154) to auto-launch the worm when a user views an infected message, or an exploit against the DNS service (CVE-2002-0374). As the DNS service does not require human interaction via mail client, the attacker selects this method of propagation across LANs. The location of susceptible BIND servers may be discovered through typical propagation methods such as the random scan, contagion-based, or coordinated scanning techniques discussed previously.

Once inside a LAN, the attacker's objective is to saturate Windows hosts with copies of the worm, potentially enabling some payload to become active. Several methods may be employed to copy and execute the contagion on Windows machines via exploitation of drive shares, as discussed in (CAN 1999-[0518—0519—0520]). Once invoked on a Windows machine, the worm may continue to participate in replicating itself across drive shares until all available shares have been infected.

We presume the worm is constructed using metamorphic techniques which render usual antivirus signature checking useless. Research at Symantec and elsewhere on detection of polymorphic and metamorphic worms continues with success, but the costs of defending against metamorphic worms is high, and motivates the study of resource-constrained short-term responses to outbreaks.

5. COMPUTER NETWORK TOPOLOGIES

In order to study propagation of worms with a given epidemic profile, we study artificially generated network topologies. We endeavor to create model topologies close to actual network topologies in some key dimensions. We divide models of network topologies into two categories. The first category contains network models exhibiting a homogeneous degree distribution. Regular graph topologies generally fall into this category but also notably the prominent random graph model, which Erdös and Rényi (ER-model) proposed [18]. The second category consists of network models with degree distributions following a power law, which is commonly found in existing, large networks. Such models are also known as scale-free networks.

5.1 Scale-Free Networks

Many real networks being studied - despite their different natures - share some common features, namely scale-free distribution of degree (following a power law), high clustering, and short average path length. We consider two scale-free network models.

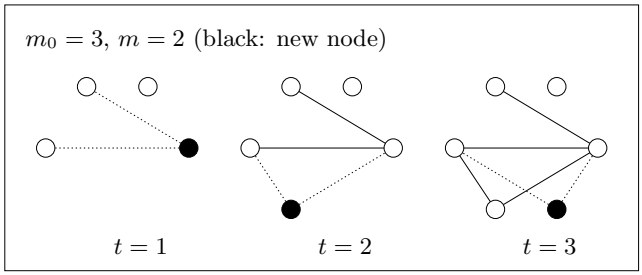

Figure 1: Example of generating BA-model

Barabási and Albert [19] gave a minimal model to generate scale-free networks. The BA-model takes three parameters: the number m_0 of initial nodes, the initial degree $m\,(\leq m_0)$ of every new node attached, and the number t of time steps. In every time step, one new node with m new edges is added to the graph. The new edges are connected to existing nodes according to the rule of preferential attachment. The probability Π to attach the new node to an existing node i depends on the degree k_i of this node, such that $\Pi(k_i) = k_i / \sum_j k_j$. Hence, new nodes prefer to attach to existing nodes with higher degrees. See Figure 1 for an example of generating a BA-model graph. The BA-model produces graphs with a power law degree distribution $P(k) = 2m^2 k^{-3}\,(k \geq m)$.

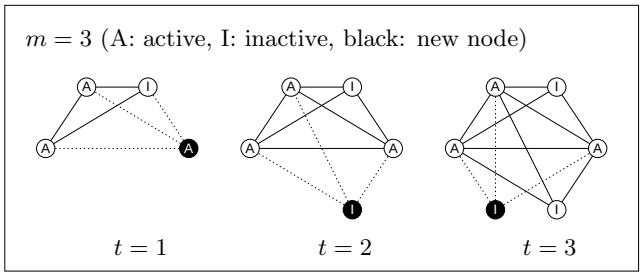

Figure 2: Example of generating KE-model

Klemm and Eguíluz [20] introduced the KE-model to model scale-free networks with a higher clustering coefficient than the BA-model. The authors also refer to their model as structured or highly clustered scale-free networks. The KE-model takes two parameters: the number of initial nodes m and the number t of time steps. Start with m fully connected, active nodes. In every time step, add one new node and attach it to all active nodes. Make the new node active as well. Then, choose one of the active nodes to be inactivated according to a probability Π inversely proportional to its current degree k_i using $\Pi(k_i) = ((\sum_j k_j^{-1})k_i)^{-1}$. See Figure 2 for an example of generating a KE-model graph. The KE-model generates networks with the same degree distribution as the BA-model, but exhibits other network topology features more similar to real computer networks.

In the future, we will compare these model topologies with actual enterprise network topologies, as well as Internet-wide topology information such as connectivity advertised through BGP (border gateway protocol) to route Internet messages between Autonomous Systems.

5.2 Network Mission

We assume the network in question is built with some purpose. For example, the provision of reliable access to information. We model this network mission as the requirement that a majority of some set of C special clients be able to communicate with the majority of some other set of S clients. We assume that random failures or the malicious worm prevents this mission communication on a node if it has infected an S or C client, or if all communication paths from an S machine to a C machine pass through at least one down or infected node.

5.3 Fault Tolerance

KE networks, by their very nature, provide connectivity robust against random faults. In particular, all pairs of nodes in a KE network are connected through m completely disjoint paths. This pleasant property provides guarantees of fault tolerant connectivity.

LEMMA 1. *In a KE network with generation parameter m there are m disjoint paths between any node in the original set of m nodes and any other node.*

The proof of this lemma follows from results proved in [21, 22]. A sketch of the proof goes as follows. Assume counterexample KE networks exist. Choose the smallest network N with some node X disconnected from some original node Y by removing $m-1$ nodes. If X is an original node, then since the original m nodes are completely connected there can be no $m-1$ cutset. If X is not an original node, examine the m older "parent" nodes X was connected to when X was added to the network. Since at most $m-1$ nodes are removed from the network, one of these m parent nodes must not be removed. If X is disconnected from Y in the modified network, then so must its parent be. However, removing X and considering the parent provides a new counterexample violating our assumption of minimality.

THEOREM 1. *In a nontrivial KE network with generation parameter m there are m disjoint paths between any pairs of nodes.*

Theorem 1 can be proven by considering the smallest counterexample represented by two nodes $N1$ and $N2$ disconnected by removal of $m-1$ nodes. The $m-1$ cutset cannot disconnect the original set of m nodes, since they are completely connected. Thus the $m-1$ cutset must partition the network such that a subset of original nodes and $N1$ appear together, and $N2$ appears in another partition without any original nodes. But this $m-1$ cutset then disconnects $M2$ from at least one of the original m nodes, contradicting Lemma 1.

Theorem 1 can be used to bound the fault tolerance of a KE network. If the chance of random failure causing a node to be unavailable at a given time is .001 (that is 99.9% uptime), then for our 1000 node server network, the chance of random failures leading to no working paths existing between any two nodes in the KE network is less than 10^{-10}.

BA networks provide much weaker guarantees, only $m/2$ connectivity in the worst case. However, BA networks on average provide levels of fault tolerance equal to or greater than KE networks.

For missions which require network connectivity, both KE and BA networks provide adequate levels of average tolerance to random faults, under reasonable assumptions. Some other obvious server network architectures (hub-and-spokes, tree, ring, etc) provide much less resilience to random faults and thus demand much higher levels of individual node reliability or provide much lower levels of delivered uptime. Finally, guaranteed m-connectivity is also indicative of high service provision levels during normal (non attack) operations. In future work we hope to better understand the relationship between actual network architectures and these simplified models.

6. SIMULATION

Scale-free networks provide reasonable models of realistic network topologies. Topological properties of scale-free networks have been well-explored, though the dynamics of interaction in scale-free networks are just emerging as a subject of the latest research [6, 3]. Here we study the dynamics of malicious code propagation in various models of scale-free networks using simulation.

Each simulated topology has $N = 50,000$ nodes. We assume the network architecture and epidemic profile as described in Section 4: a large network containing $N_{WAN} = 1,000$ autonomous systems or LANs, where each LAN contains 50 nodes, and a blended threat capable of rapidly infecting topologically close client machines, or infecting a server which then randomly selects other target servers for infection. We consider both the BA- and the KE-models of scale-free networks for the wide-area network, and we consider only a simplified completely connected topology for LANs. Assuming general node uptime (freedom from random faults) exceeds 99% (achieved with even standard desktop configurations), provision of mission communications occurs with very high probability. We use generating parameters m_0, m, and t for the WAN models such that $m_0 = m$ and $m = 10$. Then, we perform $t = N_{WAN} - m$ steps to generate the topology.

Eguíluz and Klemm [6] pointed out the divergent behavior of epidemic spreading in the BA- and KE-model. These analytical results are based on the susceptible-infected-susceptible (SIS) spreading model.

The SIS model has one parameter λ of infection probability. Each individual of the population is either infected or susceptible at any point in time. If individual A is infected at time $t-1$, it is susceptible at time t. If, otherwise, individual A is susceptible and connected to at least one infected individual at time $t-1$, then with probability λ individual A is infected at time t.

Here, we employ susceptible-infected (SI) spreading, which behaves like SIS except that infected nodes stay infected and do not change back to susceptible. In the real-world scenario we imagine with rapidly spreading malcode, we assume that machines do not become uninfected (but still susceptible to the same contagion) as the SIS model assumes. Thus in our SI-based simulations the infected nodes continue spreading the disease.

As a variation of these spreading algorithms, we also include an individual susceptibility b_i of nodes, motivated by the above discussion on epidemic criteria and epidemic profiles. The probability $b_i \cdot \lambda$ determines the chance to contract the disease from an infected neighbor. We use $\lambda = 1$.

We determine b_i from the degree d_i of node i. We implement a simple linear function that maps values between $\min(d_i)$ and $\max(d_i)$ to susceptibilities $0 < b_i < 1$. We use a descending function, so that nodes with small degrees are more susceptible than those with higher degrees. To avoid extreme susceptibilities of 0 and 1, we use the linear functions that maps $\min(d_i) - 1$ to 1 and $\max(d_i) + 1$ to 0.

To motivate our simulated immunization, we consider detection and mitigation techniques proposed in other research. First we assume immunization may be triggered by a recognition of malicious or anomalous content over network communications through which the infection is transmitted. Detection may occur via techniques that recognize malicious content embedded within known protocol traffic [23]. Alternatively, anomaly-based techniques may be employed to dynamically profile normal message content [24, 25] or to compare message content against abstract specifications of legal protocol content [26]. For this simulation we will consider immunization to represent the dynamic introduction of node-level blocking of message exchanges between those network applications or services found to be in an alerted state (e.g., detection may lead to the activation of a response device described in [27]). In the general case, we assume that while increased filtering of a network protocol that is currently in use as an infection vector will promote immunization, it may also impact other non-malicious network functions. Thus, topologies that promote immunization with minimal suppression of critical protocol communications could offer defensive advantage.

We compare unhindered epidemic spreading with increasingly aggressive node-level blocking, and we alter our assumptions of detection timeliness. In the latter setup, we monitor the prevalence ρ, which is the number of infected nodes divided by the number of nodes. If the prevalence exceeds a given response threshold ρ_{res} for the first time, we target the most connected nodes to be immunized (regardless of them being infected or not). We study two cases with 10 nodes (= 1% of all) and 100 nodes (= 10% of all) being immunized when the threshold ρ_{res} is met. We consider three different settings for this threshold: The response is launched when at least 20%, 5%, or 1% of all nodes are infected.

At the beginning of each simulation run, we select one node at random to be infected. A simulation run performs $T = 25$ time steps of epidemic spreading. For each set of parameters, we carry out 50 simulation runs with different seeds for the random number generator.

6.1 Simulation Results

Our simulation results are shown in Figures 3 and 4. The left diagrams in Figure 3 present simulation results for a 50,000 node network where the 1,000 node WAN is constructed from the BA-model. The right diagrams in Figure 3 and the diagrams in Figure 4 present results for a 50,000 node network where the 1,000 node WAN is constructed using the KE approach. Each row of diagrams from top to bottom denotes different settings for the response threshold ρ_{res} decreasing from 20% to 5% to 1% prevalence. The response threshold ρ_{res} corresponds to the level of infection in the network at which worm detection mechanisms are

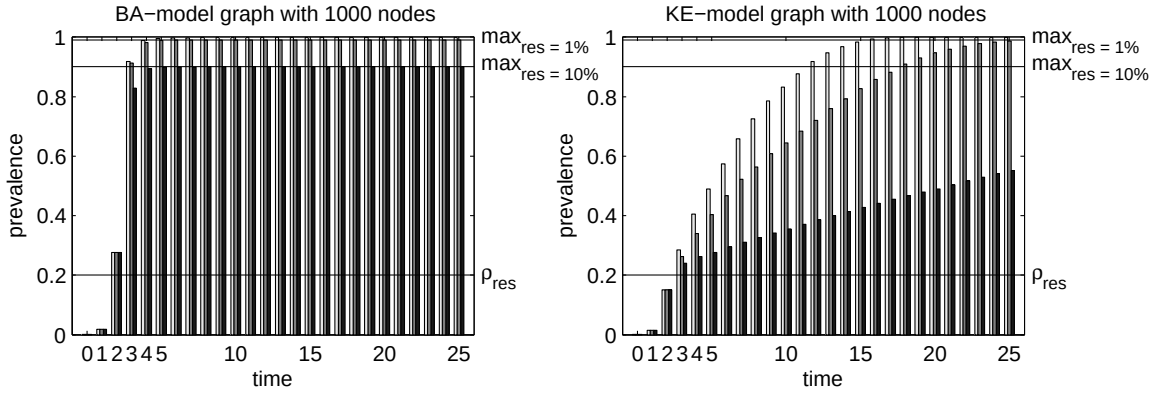

(a) Launching response when reaching threshold $\rho_{res} = 20\%$

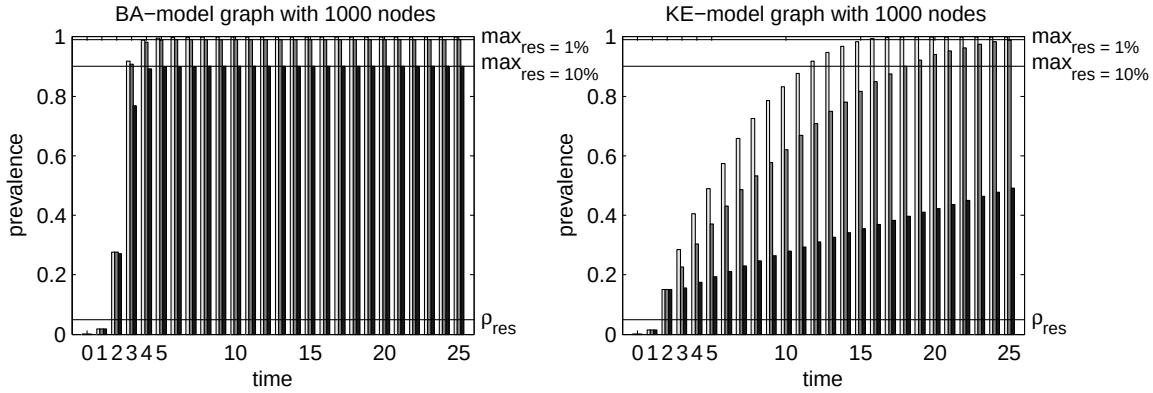

(b) Launching response when reaching threshold $\rho_{res} = 5\%$

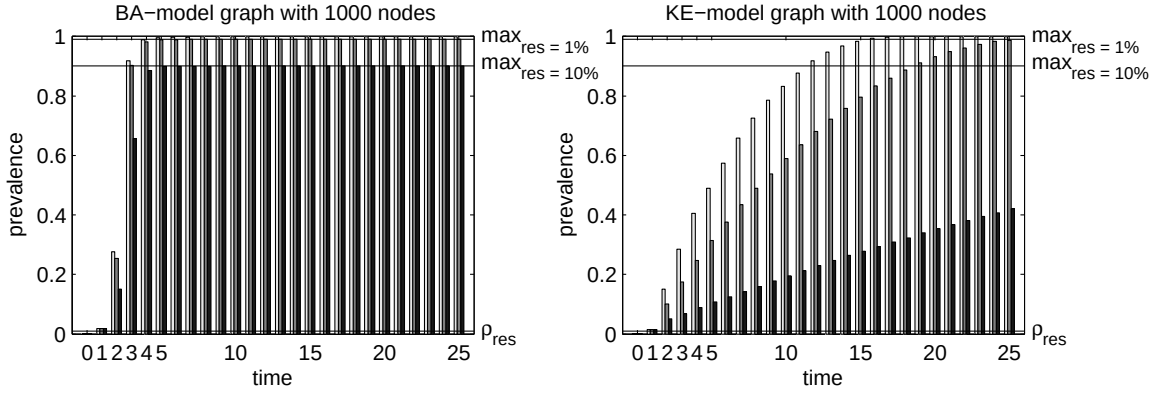

(c) Launching response when reaching threshold $\rho_{res} = 1\%$

Figure 3: Simulation results for BA- and KE-model with three bars for mean prevalences in each time step. Left bar (light gray) is no response to epidemic, middle bar (medium gray) is 1% nodes immunized, right bar (dark gray) is 10% nodes immunized when threshold ρ_{res} reached. Upper horizontal lines show maximum possible prevalence when responding, and lower horizontal lines represent the level of infection ρ_{res} at which a response is triggered.

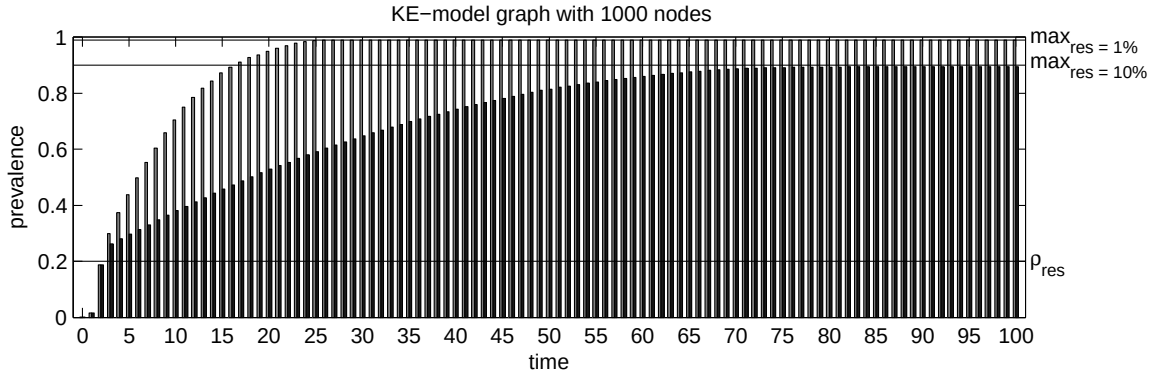

(a) Launching response when reaching threshold $\rho_{res} = 20\%$

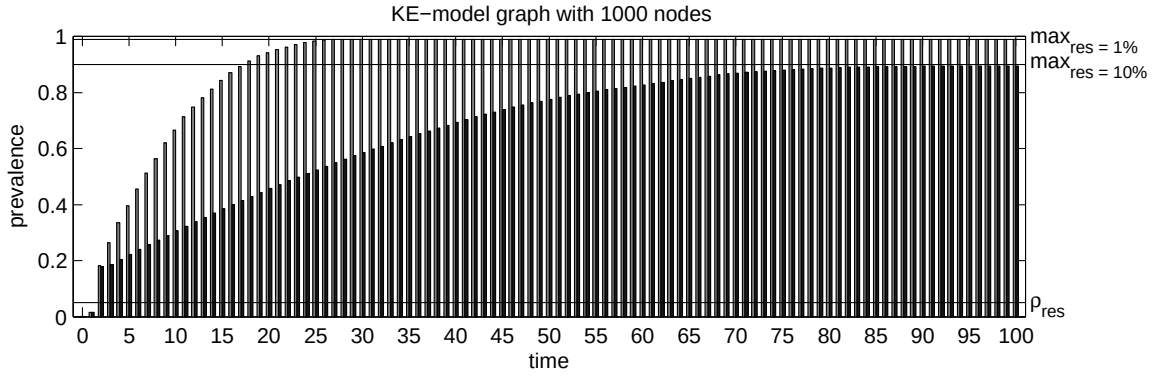

(b) Launching response when reaching threshold $\rho_{res} = 5\%$

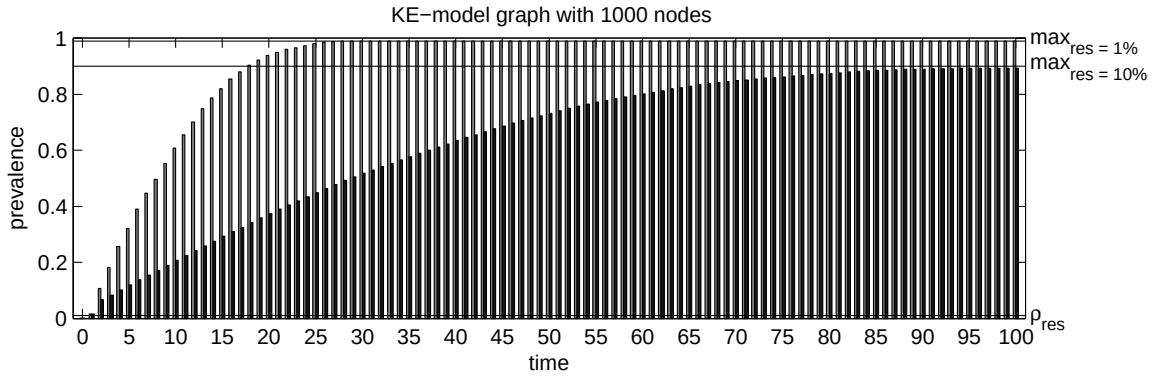

(c) Launching response when reaching threshold $\rho_{res} = 1\%$

Figure 4: Simulation results for asymptotic behavior of KE-model with 1% (lighter gray) and 10% (darker gray) nodes immunized when threshold ρ_{res} reached. Upper horizontal lines show maximum possible prevalence, and lower horizontal lines represent the level of infection ρ_{res} at which a response is triggered.

likely to detect the contagion and recognize it as a significant threat. The other simulation parameters are matched for these simulations, so they share average node connectivity, network diameter, and individual susceptibility rates. The graphs show prevalence of three methods as time progresses. Time units here are relative to the speed of worm propagation, for some worms time units would be measured in minutes, and for others it may be measured in seconds or even milliseconds.

In Figure 3, the first value (left bar in light gray) denotes epidemic spreading without any countermeasures, the second and third values (middle bar in medium gray and right bar in dark gray) correspond to responses taken with 1% and 10% of all nodes immunized, respectively. In Figure 4, we show the asymptotic behavior of the KE network when countermeasures are taken. Here, the left bars in lighter gray and the right bars in darker gray denote responding with immunizing 1% and 10% of the nodes, respectively. We kept all simulation parameters except performing 100 time steps for the results shown in Figure 4.

Note that the simulated worm spreads extremely rapidly in a BA network, even with defensive measured launched relatively early. In only a few time steps, the rapidly propagating simulated worm is able to infect the overwhelming majority of server and client machines, and compromises the mission-critical functions. In all of our simulations, BA networks rapidly suffer near-complete infection, while KE networks do so much more slowly. The simulation runs of the KE network extended to 100 time steps show saturation at a later point in time.

Importantly, we have found that network defenses put in place late in an attack (after $\rho_{res} = 20\%$ of susceptible nodes are already infected) can slow the spread of worms in certain network topologies. In particular, if successful defenses can be instituted on 10% of machines after $\rho_{res} = 20\%$ of machines are infected, the prevalence of the contagion only reaches 50% after 25 time steps and reaches maximum saturation of 90% after about 75 time steps. This might indicate that KE-like networks could slow rapid worm propagation enough to enable human or other responses to defend the remaining portions of the network, while BA-like networks do not provide the luxury of time for human response to rapidly spreading malcode.

7. TYING IT ALL TOGETHER

The simulation results described in Section 6.1 can be seen as an abstract model for the example epidemic profile described in Section 4. The epidemic subgraph of a DoD wide area network is modeled as the KE- or BA-style 1,000 node network, each node constituting the server for a 50 node local area network. The local area networks are thus heterogeneous, at least containing servers and client machines, and the entire network consists of 50,000 machines. We presume that in this network Windows client machines are prevented from sharing Windows drive shares across the WAN (as is usual), and thus one major path of infection between the Windows boxes is blocked. However, we model a blended threat worm comprising local Windows attacks, and also DNS exploits that can attack the server network. We presume the worm targets WAN victims from one server to another by random scan, but targets victim machines once inside a LAN through topology information. Since we presume open connectivity policies within each LAN, the worm is able to rapidly compromise all mission-critical clients in each LAN it infects. We also assume that simple firewalling, presumably deployed at every LAN entrypoint, is unable to detect or block the modeled worm, perhaps because the worm is polymorphic or metamorphic. The abstract modeling of network defensibility suggests WAN network architectural choices may be available that slow worm propagation but still provide strong guarantees of network connectivity.

8. CONCLUSION

We are working toward a framework that enables the analysis of real-world networks, helping network designers and administrators construct more defensible networks. We study epidemic profiles of worms, the strategies they use to infect a system and propagate across a network. We also study percolation or epidemic spread in artificially created scale-free network topologies. Finally, we bring to bear on these studies information regarding mission critical network requirements, such as reliable connectivity. Our initial research indicates that scale-free network topologies provide certain key advantages to provision of reliable network performance. Our initial simulation results suggest that some scale-free network topologies are inherently more defensible than others against rapidly spreading malicious worms. We hope to expand our studies to be able to draw crisp conclusions regarding organization of network topologies that are inherently defensible against malicious worm outbreaks, but which also provide mission-critical network services resilient to normal accidents and random outages.

9. REFERENCES

[1] David Moore, Colleen Shannon, Geoffrey M. Voelker, and Stefan Savage. Internet quarantine: Requirements for containing self-propagating code. In *Proc. of INFOCOM*, 2003.

[2] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance in complex networks. *Nature*, 406:387–482, 2000.

[3] Romualdo Pastor-Satorras and Alessandro Vespignani. *Epidemics and Immunization in Scale-Free Networks*, chapter Epidemics and immunization in scale-free networks. Wiley-VCH, Berlin, May 2002.

[4] Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic dynamics in finite size scale-free networks. *Physcal Review E*, 65, 2002. 035108.

[5] Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86:3200–3203, 2001.

[6] Victor M. Eguíluz and Konstantin Klemm. Epidemic threshold in structured scale-free networks. *Physical Review Letters*, 89(10), September 2002. 108701.

[7] Victor M. Eguíluz, Emilio Hernández-García, Oreste Piro, and Konstantin Klemm. Effective dimensions and percolation in hierarchically structured scale-free networks. submitted to Physical Review E, 2003.

[8] Zoltán Dezsö and Albert-László Barabási. Halting viruses in scale-free networks. *Physical Review E*, 65, 2002. 055103.

[9] Jasmin Leveille. Epidemic spreading in technological networks. Technical Report HPL-2002-287, HP Laboratories Bristol, October 2002.

[10] M. E. J. Newman, Stephanie Forrest, and Justin Balthrop. Email networks and the spread of computer viruses. *Physical Review E*, 66, 2002. 035101.

[11] A. Mackie, J. Roculan, R. Russel, and M.V. Velzen. Nimda worm analysis. *Security Focus*, Incident Analysis Report, Version 2, September 2002.

[12] C. C. Zou, D. Towsley, and W. Gong. Email virus propagation modeling and analysis. Technical Report TR-CSE-03-04, University of Massachusetts Amherst, Electrical and Computer Engineering Department, 2003.

[13] E. Spafford. An analysis of the internet worm. In *Proceedings of the European Software Engineering Conference*, volume LNCS 387, September 1989.

[14] Eeye Digital Security. ".ida code red' worm". Advisory AL20010717, July 2001.

[15] N. Weaver. Potential strategies for high speed active worms: A worst case analysis. Whitepaper, UC Berkeley, March 2002.

[16] S. Staniford, V. Paxon, and N. Weaver. How to 0wn the internet in your spare time. In *Proceedings of the 11th USENIX Security Symposium*, August 2002.

[17] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the slammer worm. *IEEE Security and Privacy*, 1(4):33–39, July 2003.

[18] Mark E. J. Newman. Random graphs as models of networks. In S. Bornholdt and H. G. Schuster, editors, *Handbook of Graphs and Networks*. Wiley-VCH, Berlin, 2002. To appear.

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

[20] Konstantin Klemm and Victor M. Eguíluz. Highly clustered scale-free networks. *Physical Review E*, 65, December 2002. 036123.

[21] Rebecca N. Wright, Patrick D. Lincoln, and Jonathan K. Millen. Efficient fault-tolerant certificate revocation. In *2000 ACM CCS*, Menlo Park, CA, jun 2000. SRI International.

[22] R.Wright, P. Lincoln, and J. Millen. Depender graphs: A method of fault-tolerant certificate distribution. *Journal of Computer Security*, 9(4):323–338, 2001.

[23] T. Toth and C. Krugel. Accurate buffer overflow detection via abstract payload execution. In *Proceedings of the 5th International Symposium on Recent Advanced in Intrusion Detection*, Zurich, Switzerland, October 2002.

[24] P.A. Porras and A. Valdes. Live traffic analysis of tcp/ip gateways. In *Proceedings of the 1998 ISOC Symposium on Network and Distributed Systems Security*, San Diego, California, March 1998. ISOC Press.

[25] A. Valdes and K. Skinner. Adaptive, model-based monitoring for cyber attack detection. In *Proceedings of the 3rd International Symposium on Recent Advanced in Intrusion Detection*, Toulouse, France, October 2000.

[26] R. Sekar, A. Gupta, J. Frulo, T. Shanbhad, A. Tiwari, Y. Yang, and S. Zhou. Specification-based anomaly detection: a new approach for detecting network intrusion. In *Proceedings of the ACM Conference on Computer and Communications Security*, Washington DC, November 2002.

[27] D. Nojiri, J. Rowe, and K. Levitt. Cooperative response strategies for large scale attack mitigation. In *Proceedings of the 2003 DARPA DISCEX Conferenece*, Washington DC, April 2003.