A Model for Opportunistic Network Exploits: The Case of P2P Worms*

Michael Collins
CERT/NetSA,
Carnegie Mellon University
mcollins@cert.org

Carrie Gates†
CA Labs,
Islandia, NY
carrie.gates@ca.com,

Gaurav Kataria
Heinz School,
Carnegie Mellon University
gauravk@andrew.cmu.edu

Abstract

We segregate attacks into two categories – targeted and opportunistic – based on whether the attacker compromises a specific target (targeted) or a number of intermediate targets to fulfill his end goal (opportunistic). We assume that opportunistic attackers consider targets indistinguishable except for their vulnerabilities, and are interested in acquiring as many targets as possible. We therefore hypothesize that opportunistic attackers will develop attacks involving services which have the largest number of potential targets. We test this hypothesis in a limited way by correlating worm releases on P2P file sharing networks with the number of users on the networks being targeted. Our results demonstrate that this relationship exists only for variants of worms and not for new worms. We further demonstrate that the results are service specific, and that there is no general model that represents the entire file sharing vector.

1 Introduction

As the Internet has grown, so too has the number of worms and viruses, from the single Morris worm in 1988 to over twenty thousand worms in 2005[1]. To that end, we need to understand current attacker practice to inform the design of detection and defense systems.

In this paper, we classify attackers into two categories: *targeted* and *opportunistic*. A targeted attacker has some *a priori* knowledge or interest in the target before initiating the attack. An example of a targeted attacker is the "jaeger", described by Stoll[17], who was attacking specific machines in an attempt to obtain American protected information for the Russian government. In contrast, an opportunistic attacker has no specific interest in a target outside of its exploitability. An opportunistic attacker gains by exploiting a large number of remotely controlled computers. For example, Mirkovic and Reiher[12] developed a taxonomy of DDoS attack tools and defenses that described an initial *recruitment* phase during which attackers acquire resources for DDoS; this recruitment phase is a specific form of opportunistic attack. Where a targeted attacker seeks maximum gain through exploiting specific vulnerabilities in a unique target, an opportunistic attacker seeks to control as many targets as possible. An opportunistic attacker is therefore likely to attack a large number of targets indiscriminately; given the large number of computers connected to the Internet, a small probability of success is sufficient to acquire a large number of systems.

Worms and viruses are examples of opportunistic attack mechanisms. Given that the goal of an opportunistic attack is to amass resources, we hypothesize that worms and viruses will target those services that have the largest market share. As numerous factors simultaneously affect an attacker's decision to exploit a service, it is difficult to disentangle the relative importance of each. Yet, by carefully constructing a fairly controlled dataset on release of worms and viruses on P2P file sharing networks over a three year period,

*This work was supported in part by grant no. CNS-0433540 from the National Science Foundation.

†This work was performed while with CERT/NetSA at Carnegie Mellon University.

we are able to identify some significant correlates of attacker choice. Specifically, we estimate the effect of market share on the release of worms and worm variants.

The ultimate goal of our research is to develop a predictive model of exploits that can inform network administrators when a newly introduced service is likely to be affected by worms and viruses. Such a model could be used by a network administrator to determine at what point in time he should implement a policy regarding usage of a service based on his assessment of risk. In this paper, we attempt to model the current exploitation of services, in specific P2P file sharing, recognizing that if we cannot explain current trends, then we are even less likely to be able to predict future trends.

We present an overview of related research in Section 2, followed by a more in-depth description of our problem domain in Section 3. We describe our modeling approach in Section 4 and discuss the results from our models in Section 5. Section 6 then provides some concluding remarks.

2 Related Work

While there are few studies of attackers, some taxonomies of attackers have been developed. The earliest taxonomy was developed by Landreth in 1985 [8], who had five divisions based loosely on the motivation of the attacker. The motivations ranged from the desire to learn, to damaging a system, to stealing from this system. By 1995, motivations such as national interest and financial gain were emerging and being included in attacker taxonomies [5, 14]. In 2000, Rogers [16] argued that a generally accepted taxonomy of attackers needed to be developed where the categories would allow psychological profiles to be assigned. He proposed a new taxonomy [15] that consisted of seven categories where he linked technical sophistication and motivation. The second least-sophisticated category was *cyber-punks*, which he described as having enough computer knowledge to sometimes write their own limited software. He further commented that they tend to engage in malicious activity, and that they were the group on which the media tends to focus. We expect that this is the group that aligns closely with our notion of opportunistic, or disinterested, attackers. However, given that a large number of denial-of-service attacks are motivated by financial goals [11] (*e.g.*, extortion of the victim based on threatening to attack) which would require the attacker to amass a large number of attack machines, the "professional criminal" category, which is the second most sophisticated form of attacker, will also apply.

Studies of attack tools are more common in the literature than are studies of attackers. Early studies of attack tools were single-instance cases, such as histories of specific attacks or analyses of particular tools. Focusing in particular on worms, as they are a common example of an opportunistic attack method, studies of particular worms have been published, such as Moore *et al.*'s study of the Sapphire/Slammer worm [13] and Zou *et al.*'s study of Code Red [21]. Worms have also been aggregated into taxonomies, such as that of Weaver *et al.*'s [19], which includes a breakdown of the methods of worm propagation as well as a study on attacker motivation.

When worms have been modeled in the past, the focus has been on modeling their propagation. The first such analysis was performed in 1991 by Kephart and White [7], who modeled the spread of a virus as a function of time, as well as the impact of detecting and removing viruses quickly enough on the final number of infected hosts. More recently, Zou *et al.* [21] developed a propagation model that fit the spread of the Code Red worm. Wang *et al.* [18] describe propagation models in terms of graphs, removing some of the constraints of previous models where homogeneous connectivity between nodes is assumed. Propagation models of other non-homogeneous systems, such as wireless networks [10] where the issues include dynamically changing connectivity levels, have also been developed.

In addition to exploiting vulnerabilities in software (*e.g.*, CodeRed, SQL Slammer), worms have also been developed that use email for propagation. For example, Zou *et al.* [20] modeled the spread of email worms using factors such as email checking time and the probability of a user clicking on an attachment. Similarly, instant messaging worms have also developed, which have been analyzed by Mannan and van Oorschot [9].

Other models of attacks, which were not specifically worm based, have also been developed. For example, Arora *et al.* [3] developed and analyzed economic models investigating the affects of vulnerability disclosures on both the attacker community and the vendor community in terms of the number of attacks received as a function of time. Browne *et al.* [4] have studied the relationship between vulnerability announcement and exploit release. Arbaugh *et al.* [2] modeled the life cycle of a software vulnerability.

3 Problem Description

We divide attackers into two categories: targeted and opportunistic. A targeted attacker relies on skill or inside knowledge to exploit a vulnerability on a high value target. In contrast, an opportunistic attacker seeks to exploit common mode failures across multiple targets. Note that this distinction is based on the mode of attack and not behavior of attacker. The same individual may act as an opportunistic attacker to garner machines that can be used later to launch a targeted attack. In general, opportunistic attacks have a considerably higher rate of failure than targeted attacks; in order to ensure a reasonable rate of return, opportunistic attacks must reach a sufficiently large pool of targets.

Because opportunistic attacks must interact with thousands of targets, they are often automated. One method of automating these attacks is to use worms. For the purposes of this paper, we define a *worm* as an application that is installed on a computer without the informed consent of that computer's owners and which, upon installation and without external stimulus, attempts to repeat the installation upon other computers reachable from it. We note that our definition encompasses several conventional definitions of a virus.

In order to control a system, the attacker (a person or an automated tool) will use a *vector*. We define a vector as any method used by the worm to transmit itself to the target, such as a buffer overflow impacting the stack, (e.g. a remote attacker can send a malformed request to a vulnerable IIS web server to gain root access[1]), a list of frequently used passwords, or a deceptive message sent to a user. An initial list of vectors includes:

Email : A worm propagates through email by sending a deceptive message containing a copy of the worm as an attachment. The targeted user is then deceived into opening the attachment to continue the propagation cycle. Email vectors can be further delineated by how they acquire addresses and what methods they use to send mail (such as via Microsoft Outlook or through the use of an ad hoc SMTP engine).

Message : A message worm propagates like an email worm, but by using a common messaging protocol, such as AIM or MSN Messenger. Message vectors are differentiated by the messaging protocol used.

File Copy : A file copy vector transfers a file to a specific location on the target's filesystem, such as floppy disks or a mapped drive. File Copy vectors are differentiated by where on the filesystem they copy the file.

File Sharing : A file sharing vector copies the worm file to a directory known to be used by a particular file sharing application. The worm file is then given a deceptive name and thus may be pulled by other users on the same file sharing network. File Sharing worms can be differentiated by the file sharing application or network exploited by the worm.

Vulnerability : A vulnerability worm scans IP space for machines that are vulnerable to a certain buffer overflow or other software vulnerability like race condition or environment error. The worm then launches an *exploit* to use that vulnerability to gain control of the target. Vulnerability vectors are differentiated by the platform and specific in-platform vulnerability exploited by the vector.

Password : A password worm uses common password lists or known weak passwords, such as the default passwords used with the Oracle or SQL Server databases, in order to gain access to a server. Password worms are differentiated by the password list used and the targeted service.

[1] CVE-2002-0364:Buffer overflow in the chunked encoding transfer mechanism in IIS 4.0 and 5.0 allows attackers to execute arbitrary code via the processing of HTR request sessions, aka "Heap Overrun in HTR Chunked Encoding Could Enable Web Server Compromise".

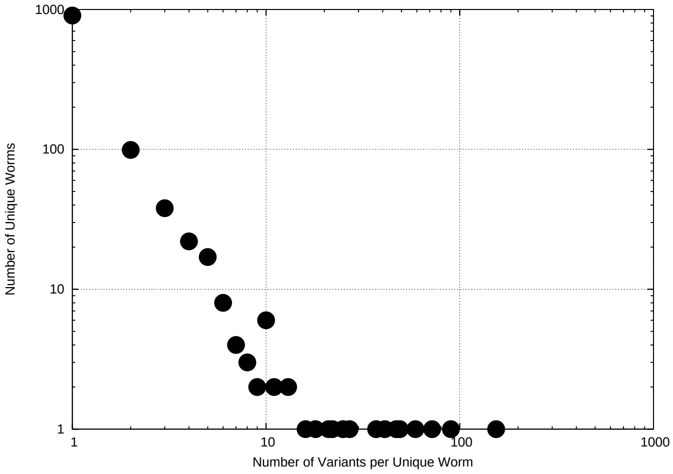

Figure 1: Count Of Variants Per Worm. The distribution fits a power law with an R^2 of 0.75.

We hypothesize that opportunistic attackers will choose to exploit those vectors that grant them access to the largest number of targets. Within a particular class of vectors, we therefore expect attackers to choose *services* that have the largest market share. We define services as the various programs within a particular class of vectors that can be exploited. For example, the vulnerability vector contains the DCOM and the LSASS services.

Interestingly, not all worms are written from scratch – attackers release new worms by modifying the source code of other worms that they encounter to suit their own needs. These new worms, which will consist of the original worm code with additional or reconfigured features, are called *variants*[2]. An example of this behavior was observed in case of modification of the Blaster worm as Blaster.B variant; the author of Blaster.B had no relationship to the author of Blaster, instead modifying the code he received when his machines were infected [6]. New variants of a worm can also be developed by friends/partners of the original worm writer, if he chooses to share his source code with them.

We use the Symantec threat database[3] to estimate the relationship between worms and their variants. Symantec names worms using a 3-part platform.name.variant format, such that W32.Bobax.C is the variant C of Bobax worm which affects Win32 (32 bit Windows) family of operating systems. As new variants of a worm are discovered different letters are used to represent them. Different vendors follow different nomenclature and classification schemes for worm reporting. In our analyses, we choose to follow Symantec's classification of worms. Our results are thus presented with the caveat that the worm variants detailed are the variants as described by the Symantec database.

Figure 1 shows the relationship between worms and their variants derived from the Symantec threat database. This distribution follows a power law: the majority of worms have only one variant, while a small number of worms are repeatedly modified. The worm that has the most variants is w32.mytob, with 154 variants.

Given this background, we have the following set of goals:

Goal 3.1 *Determine if the market share for a particular service is a significant predictor of the number of worms written for that service.*

Goal 3.2 *Determine if the market share for a particular service is a significant predictor of the number of variants written for that service.*

Goal 3.3 *Determine if specific services are a significant predictor of the number of worms written for that vector.*

[2] The original worm is considered the *first* variant.

[3] `http://www.symantec.com/avcenter`

Goal 3.4 *Determine if specific services are a significant predictor of the number of variants written for that vector.*

4 Estimation Approach

In this section we analyze the factors that could potentially influence an opportunistic attacker's decision to choose a particular attack vector. We then construct a sufficiently controlled dataset to isolate and test the impact of some key determinants.

In contrast to popular conception, only a small proportion of worms released in year 2005 used software vulnerabilities as attack vectors [1]. According to Symantec reports written for *new* worms in year 2005, the majority used email, with 57 write-ups, followed by P2P file sharing with 30 write-ups, out of a total of 126 worm write-ups. In comparison, 21 worms used the LSASS[4] vulnerability[5]. It may be hypothesized that opportunistic attackers prefer technically simple exploits to more complicated exploits like buffer overflows. At the same time, an equally compelling argument would be that attackers seek maximum exposure and hence their choice of attack vector. It is very difficult to isolate these two effects; for example, when comparing Windows to Linux it is not possible to pinpoint whether the large market share or the poor security and management of windows systems is responsible for its lion share of exploits, and to what extent. Furthermore, there are other contextual factors that may potentially influence an opportunistic attacker's choice of attack vector, such as attacker expertise, availability of toolkits (an indirect externality faced by developers of worm variants), or attacker preferences (liking or hatred toward a particular software or service). In practice, isolating and estimating the relative importance of each factor is hard. Aggregate statistics on worms and viruses serve only to highlight the current security risks. In order to be able to predict future risk associated with a particular software or service, under a given set of conditions, we set out by constructing a fairly controlled dataset on worms and variants affecting P2P file sharing.

Of the six attack vectors described in Section 3, we chose P2P file sharing because, by comparing the release of worms and variants on different P2P networks (*e.g.*, Gnutella, FastTrack, eDonkey) over time, we can identify the effect of change in user population (and market share) of different networks on the release of worms and variants affecting them. By comparing only within the file sharing (P2P) class of attack vectors, we try to ensure that the attackers decision is not affected by either difficulty of exploit, expertise, attack vector preference, or any other technical reason. Furthermore, as all P2P networks are quite similar to each other, there are limited sources of in-sample variation other than the difference between the number of users on each network, though we cannot rule out network specific fixed effects. Thus, this study design lets us estimate the impact of user market share on the count (and the corresponding proportion) of worms and variants released on various networks.

4.1 Data

Symantec Corporation maintains a security response database publicly available at `http://www.symantec.com/avcenter`. This database contains a list of Symantec anti-virus signatures, and writeups describing a significant subset of those signatures; for example, on March 9th, 2006, there were a total of 6,049 signatures and 2,981 writeups for worms under the "W" category.[6]

As mentioned earlier, we use Symantec's definition of worms and variants for our analyses. Figure 2 is a chart of the total number of malcode reports written by Symantec over a period of three years. As is evident from this figure, the number of reports written by Symantec appear to have a time variation, where some year end decline in reporting can be visibly noticed. We can not ascertain from the data available to us if this is due to a seasonal variation in Symantec's reporting or if it reflects actual attack trends. We are therefore unable to interpret any time trends in release of worms or variants.

[4] http://www.microsoft.com/technet/security/Bulletin/MS04-011.mspx

[5] Note that these values reflect the current trend of worms to incorporate multiple vulnerabilities. Thus 57 worms that used email and 30 that used file sharing might also have used other vectors, including both email and file sharing so that the worms

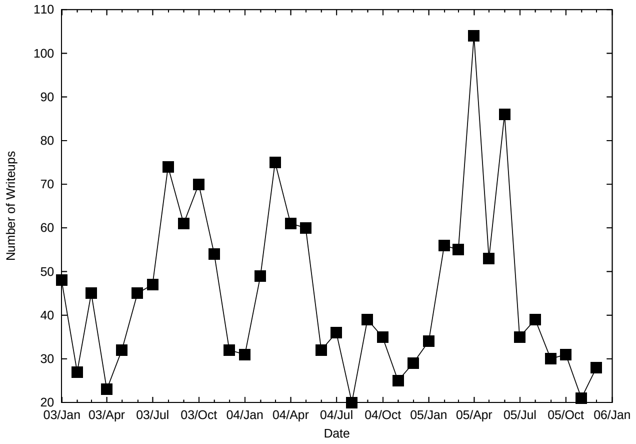

Figure 2: Total writeups produced by Symantec appear to have a possible time variation. We are unable to ascertain if this is due to reporting variation by Symantec or a reflection of actual attack trends.

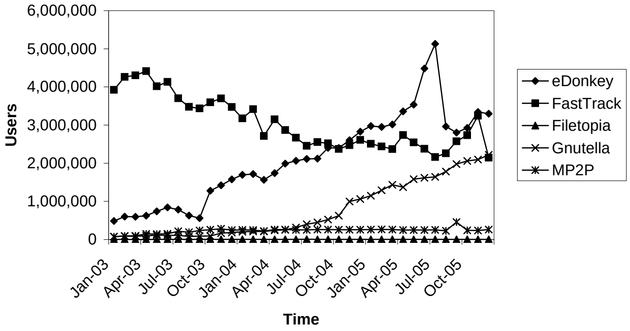

Figure 3: User populations for each of the file sharing services

Peer-to-peer users can use a variety of file sharing clients, but the majority of peer-to-peer activity takes place on three networks: FastTrack network (which includes Kazaa and Grokster clients), eDonkey network (which includes eDonkey and Overnet), and Gnutella network (which includes Gnutella, Gnucleus, Bearshare, Limewire etc.). The majority of clients use a single network, although there exist plugins that support multiple-network functionality and some clients (such as Shareaza) support multiple network communications.

A host on a peer-to-peer network shares files by copying them to a specific local folder and then publishing the contents of that folder to the network as a whole. The contents of that folder can then be searched by the rest of the file share network and copied by other users looking for that particular file. The file sharing worms we examine propagate by copying themselves to the shared directory under a deceptive name, such as a current video game, movie or list of passwords to pornography sites. Users who download those files in addition/instead copy the worm and thus continue to propagate it throughout the file sharing network.

File sharing clients provide statistics on the total size of the network, and a time series of these statistics are available at slyck.com[7] (Figure 3). We consider the total user population accessible to a worm as the sum of that of all the file sharing clients that the worm uses. For example, a worm that spreads using only KMD and Kazaa (both FastTrack clients) can reach the population of the FastTrack network, whereas, a worm that spreads using Gnutella client and Kazaa can spread across both Gnutella and FastTrack networks.

would appear in both categories.

[6] This category includes all worms in the Windows family, such as W32, W2K and W95.

[7] `http://www.slyck.com/stats.php`

Network	Worms	Variants
Multiple Networks	78	154
FastTrack only	62	122
eDonkey only	3	3
Gnutella only	1	3
FileTopia only	0	4
MP2P only	0	0

Table 1: Aggregate number of worms and variants released on various P2P networks between Jan 2003 and Dec 2005

Using this approach raises a danger of over counting as a user may be present on Kazaa and eDonkey at the same time. In addition, servers in peer to peer networks may artificially inflate their statistics. However, we assume that attackers knowledge about the potential network size may be no different than our/user knowledge of the network size.

We track release of worms and variants on five main P2P networks, over period of three years, Jan 2003 to Dec 2005. The aggregate number of worms and variants released in reported in Table 1. As noted, many worms were capable of affecting clients of multiple networks. Some of these worms search for all directories on the host computer with the name “share” and copy themselves to it, thus affecting multiple clients.

The FileTopia network is listed as having no worms, but four variants. This occurs because some worms were originally designed for one of the other networks, but were then modified to attack the FileTopia network and thus appear as a variant on FileTopia network. We include the MP2P network in our analysis, even though it does not have any worms or variants, because it is still a valid data point. In Table 1, “Multiple Networks” refers to those worms and variants that targeted two or more networks.

5 Results and Limitations

In order to estimate the effect of factors influencing release of worms and variants, we first ran a negative binomial regression for the worm count data. We used seven independent variables: the time that the data set covers (in months), the number of users affected (in millions), and five dummy variables representing each of the five networks. For those worms or variants that used multiple networks, the number of users was the sum of each of the attacked networks, and the corresponding dummy variables were set. We performed a negative binomial regression because the dependent variable (number of worms or number of variants released per month) fit well the negative binomial distribution. The results of the regression are reported in Table 2, column 2.

Only the coefficients on dummy for FastTrack and eDonkey networks turned out to be significant predictors of worm count[8]. Surprisingly, the coefficient on number of users was insignificant. A similar trend was observed for variant count data (Table 2, column 3).

An implicit assumption in the above estimation was that worm writers are focused only on P2P networks and respond directly to change in user population. However, that is not entirely true as worm writers may also compare and choose between various services (e.g. P2P, email etc.). Since looking at the actual count data can be misleading, we decided to focus on the proportion of worms released. In addition, we normalized the user population to estimate the effect of market share of network(s) on the share of worms it attracts. The independent variables included dummy for each of the five different networks. As the dependent variable in this case was proportion of worms, we used tobit regression censured on left at 0 and on right at 1. The results are reported in Table 3.

As in the previous set of regressions we find that the dummy variables for FastTrack and eDonkey are significant, for both worm proportion as well as proportion of variants released. In particular, we note that

[8] We refrain from interpreting month coefficient due to possible variation in Symantec’s reporting

	Dependent Variable	
Variables	Worm count	Variant count
Users (in Millions)	0.09 (0.14)	0.13 (0.11)
Month	-0.05^* (0.01)	-0.05^* (0.01)
FastTrack	2.90^* (0.68)	2.88^* (0.53)
eDonkey	-0.90^* (0.33)	-1.03^* (0.28)
Gnutella	0.21 (0.23)	0.28 (0.20)
FileTopia	-15.67 (1944.24)	0.32 (0.67)
MP2P	-15.68 (1933.14)	-13.40 (472.28)
Constant	-2.31 (0.56)	-1.76 (0.44)

Notes: * $p < 0.01$; Standard errors in parenthesis.

Table 2: Estimation using Negative Binomial Regression

	Dependent Variable	
Variable	Worm Proportion	Variant Proportion
Market share	0.39 (0.44)	0.63^* (0.29)
Month	-0.005 (0.006)	0.002 (0.003)
FastTrack	0.74^* (0.28)	0.33^* (0.15)
eDonkey	-0.46^* (0.17)	-0.44^* (0.10)
Gnutella	0.02 (0.11)	-0.009 (0.07)
FileTopia	-2.67 (-)	-0.07 (0.14)
MP2P	-2.68 (-)	-1.83 (-)
Constant	-0.75 (0.21)	-0.47 (0.12)

Notes: * $p < 0.01$; Standard errors in parenthesis.

Table 3: Estimation using Tobit Regression

the FastTrack coefficient is positive and the eDonkey coefficient is negative in both cases; this may be a result of network demographics, as FastTrack is more popular in U.S. whereas eDonkey is more popular elsewhere. We find that market share of networks is not a significant predictor of proportion of worms released, however, it is a significant predictor of proportion of variants released. Specifically, we can say that 1% increase in market share attracts 0.63% more variants.

Relating these results back to the four goals outlined in Section 3, our results indicate that the market share for a particular service is *not* a significant indicator of the number of worms written for that service. However, market share *is* a significant indicator of the number of variants written for that service. This implies that attackers who are truly opportunistic, modify only those worms that allow them to target more users. The power law relationship between number of variants and worms also seems to suggests that variant writers are disproportionately interested in a few worms.

Our other two goals were related to determining if the particular service (network) being exploited was a significant indicator of either the number of worms or variants written that month. In the case of both worms and variants we found that FastTrack and eDonkey are significant indicators. For variants, this implies that, while the relationship between market share and number of variants is significant, it is also influenced by the network being exploited.

Our results are limited by two main factors. The first is that we have demonstrated the relationship between market share and number of variants only in the context of file sharing. Thus we can only accept our hypothesis when dealing with file sharing applications. We have not demonstrated that this relationship can be generalized to other vectors, such as messaging or email. A second limitation is due to the number of categories available. For example, we identified only six vectors, within which we focused solely on file sharing networks of which there are only five. However, given that we have observed differences even between these five networks, we have demonstrated that no general model for this vector exists. A third limitation is that we have only demonstrated a relationship in existing data, and have not provided any indication of how well this will perform when making predictions about future behavior. Thus we do not know if, for example, our model overfits our data.

6 Conclusions And Future Work

In this paper we introduced two categories of attackers: targeted and opportunistic. We focused on opportunistic attackers, who have as a goal the acquisition of a large number of computers. One of the automated methods that can be used to acquire a large number of computers is through the use of worms.

We hypothesized that worm authors choose their target service based on the market share of that service. We further hypothesized that the number of variants for a worm would be similarly related to the market share of the attacked service. These hypotheses were based on the belief that an attacker will want to acquire as many machines as possible for as little effort as possible.

We employed an estimation strategy based on Tobit regression to find the effect of market share on the proportion of worms and variants released on various P2P networks. We found that the market share of a network is *not* a significant predictor of the proportion of new worms that are written for a particular network. Whereas, we found that the proportion of worm variants is significantly influenced by the market share. This suggests that attackers are likely to continue to modify and use worm code that attacks services with the greatest market share. To some extent this is evident in increasing number of variants on email network.

We also found that eDonkey and FastTrack networks had significant effect both for the proportion of worms and the proportion of variants released. This highlights that due to some network specific fixed effects it may not be possible to write a general model that represents the relationship between market share and variants for the entire attack vector.

In a follow up of this study we would like to investigate the factors, in addition to market share, that influence variant writers.

References

[1] *Symantec Internet Security Threat Report, Volume IX.* Symantec Corporation, 20330 Stevens Creek Road, Cupertino CA, 2006.

[2] W. A. Arbaugh, W. L. Fithen, and J. McHugh. Windows of vulnerability: a case study analysis. *IEEE Computer*, 33:52–59, 2000.

[3] A. Arora, R. Krishnan, Anand Nandkumar, Rahul Telang, and Yubao Yang. Impact of vulnerability disclosure and patch availability- an empirical analysis. In *Workshop on the Economics of Information Security (WEIS)*, Minneapolis, MN, 2004.

[4] H. K. Browne, W. A. Arbaugh, J. McHugh, and W. L. Fithen. A trend analysis of exploitations. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, pages 214–231.

[5] Richard O. Hundley and Robert H. Anderson. Emerging challenge: security and safety in cyberspace. *IEEE Technology and Society Magazine*, 14(4):19 – 28, 1995.

[6] Gregg Keizer. FBI nabs teen suspect in MSBlaster variant. `http://www.techweb.com/wire/26801944`, August 2003.

[7] Jeffrey O. Kephart and Steve R. White. Directed-graph epidemiological models of computer viruses. In *Proceedings of the 1991 IEEE Computer Society Symposium on Research in Security and Privacy*, pages 343–359, May 1991.

[8] Bill Landreth. *Out of the Inner Circle: A Hacker's Guide to Computer Security.* Microsoft Press, 1985.

[9] Mohammad Mannan and Paul C. van Oorschot. On instant messaging worms, analysis and countermeasures. In *Proceedings of the 2005 ACM Workshop on Rapid Malcode*, pages 2–11, November 2005.

[10] James W. Mickens and Brian D. Noble. Modeling epidemic spreading in mobile environments. In *Proceedings of the 4th ACM Workshop on Wireless Security*, pages 77–86, September 2005.

[11] Jelena Mirkovic, Sven Dietrich, David Dittrich, and Peter Reiher. *Internet Denial of Service: Attack and Defense Mechanisms.* Prentice-Hall, 2005.

[12] Jelena Mirkovic and Peter Reiher. A taxonomy of ddos attacks and ddos defense mechanisms. *SIGCOMM Comput. Commun. Rev.*, 34(2):39–53, 2004.

[13] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. The spread of the sapphire/slammer worm. Technical report, CAIDA, ICSI, Silicon Defense, UC Berkeley EECS and UC San Diego CSE, 2003.

[14] Donn B. Parker. *Fighting Computer Crime.* Wiley Computer Publishing, New York, 1998.

[15] Marc Rogers. A new hacker taxonomy. http://www.escape.ca/˜mkr/hacker_doc.pdf. Last visited: 12 March 2002, 2000.

[16] Marc Rogers. Psychological theories of crime and hacking. http://www.escape.ca/˜mkr/crime_doc.pdf. Last visited: 12 March 2002, 2000.

[17] Clifford Stoll. *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage.* Doubleday, 1989.

[18] Yang Wang, Deepayan Chakrabarti, Chenxi Wang, and Christos Faloutsos. Epidemic spreading in real networks: An eigenvalue viewpoint. In *22nd International Symposium on Reliable Distributed Systems (SRDS'03)*, pages 25–34. IEEE Computer, October 2003.

[19] Nicholas Weaver, Vern Paxson, Stuart Staniford, and Robert Cunningham. A taxonomy of computer worms. In *First Workshop on Rapid Malcode (WORM)*, pages 11–18. Association for Computing Machinery, October 2003.

[20] Cliff C. Zou, Don Towsley, and Weibo Gong. Email worm modeling and defense. In *13th International Conference on Computer Communications and Networks*, pages 409–414, October 2004.

[21] Cliff Changchun Zou, Weibo Gong, and Don Towsley. Code red worm propagation modeling and analysis. In *Proceedings of the 9th ACM conference on Computer and communications security*, pages 138 – 147. ACM, November 2002.