Analysis of Spam

Anselm Lambert

A dissertation submitted to the University of Dublin,
in partial fulfilment of the requirements for the degree of
Master of Science in Computer Science

Department of Computer Science,
University of Dublin, Trinity College

September 2003

Declaration

I declare that the work described in this dissertation is, except where otherwise stated, entirely my own work and has not been submitted as an exercise for a degree at this or any other university.

Signed: ____________________

Anselm Lambert

September 9th, 2003

Permission to lend and/or copy

I agree that Trinity College Library may lend or copy this dissertation upon request.

Signed: ____________________

Anselm Lambert

September 9th, 2003

Acknowledgments

I would like to thank Pádraig Cunningham for submitting this intriguing topic as a dissertation project. Thanks also to my classmates who strived to achieve the highest standards throughout the year while at the same time injecting humour and novelty into the learning process.

Abstract

Spam is a pervasive annoyance in the lives of the Internet user. It has exploded into all facets of communications from mobile phones to personal organisers, and it has become a topical subject of discussion due to recent media coverage. Spam has a tangible cost measured in lost productivity, bandwidth usage, administration, and invasion of privacy. As a result, an anti-spam industry has evolved in order to counter the spam attack with a focus on two spam-filtering categories: collaborative techniques and content-analysis techniques.

This research involved analysing a wide variety of e-mail in order to produce a profile of spam and, more importantly, develop a profile of the spammer. A number of fundamental questions are answered, for example: are current definitions of spam adequate and if so, are they globally applicable? There was also an investigation to examine the possibility of a spammer successfully targeting e-mail to an individual or group of individuals.

In this study, honeypot accounts were created and positioned to receive spam. The outcome of this research is a definitive guide to spam, which will provide researchers and regular Internet users alike with knowledge that will aid them in the fight against spam and facilitate the improvement of spam filtering techniques.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

As Vice President and Chief Information Officer (CIO) of a Fortune 1000 company in the United States, it was customary for me to come into the office early before most other people. One particular morning in 1999, I was surprised to find my e-mail inbox empty. Usually I could expect to see 20 to 30 new e-mails, so I decided to investigate further. Immediately upon entering the data-centre, I noticed a lot of activity on the console of the e-mail gateway. It was possible to see enormous amounts of e-mail coming in and almost instantaneously getting re-routed out again. By opening up one of these e-mails and finding a product advertisement, the source of the problem became apparent. A spammer was using my company's equipment to send out tens of thousands of spam messages to unsuspecting victims. Unfortunately we had left the SMTP (Simple Mail Transfer Protocol) relay open on the e-mail server, which provided a conduit for a spammer to send unsolicited bulk e-mail onto the Internet.

It was a cat-and-mouse game for the next four hours as we tried to stop the company's system being used in such a manner. At one stage, the e-mail gateway crashed as a result of all the activity. Eventually, all external e-mail had to be disabled and only internal e-mail was possible, which was routed using the mail servers located at the main offices throughout the country. As we were a sales organisation and e-mail was an integral part of our day-to-day business, it was vital to recover the ability to send and receive Internet e-mail. Closing the open relay solved the problem.

I received spam in the past, but only regarded it as a nuisance that disappeared by hitting the delete key. In this instance however, my company's equipment was hijacked and used in an illegal manner. It was at that moment that I realised the havoc that could be caused by spam and it has been a continuous battle since then.

1.1 Motivation

According to eMarketer, 76 Billion spam e-mails will be delivered in 2003 [1]. But this seems to be a low estimate as the anti-spam company Brightmail blocks more than 60 billion e-mails per month on behalf of Internet Service Providers (ISP) and corporate customers. MSN Hotmail stated in an e-mail sent to all it users in May 2003, that they were blocking more that 2 billion spam e-mails every day. America Online (AOL) reached the 1-billion-spam-per-day mark in March 2003, which meant that the company was blocking an average of more than 28 spam e-mails reaching every AOL account on a daily basis. Ferris Research reported that in 2003 the corporate user would waste 15 hours deleting e-mail, compared to 2.2 hours in the year 2000. Gartner Inc. reported that in 2002, 25% of all e-mails qualified as spam. Other estimates declare that spam accounts for approximately 40% of global e-mail traffic, and analysts predict that by September 2003 over half of all e-mail will be spam [2].

Even though there are discrepancies in the amount of spam e-mail generated each year, the estimated numbers are still staggering. Developing a suitable profile on spam and learning about the behaviour of the spammer can lead to more accurate spam filters and can also lead to other methodologies to reduce spam. This will result in substantial savings in both administrative and end-user resource costs. This is the motivation behind the research. The main goal of the project is to present a definitive guide to spam.

From a personal perspective, as the Information Technology (IT) leader in a large company I saw the effects of spam on the emotions of some recipients. There was one incident where a chain mail message was e-mailed to a user in the company. The contents of the e-mail dealt with a boy that was missing (presumed kidnapped) and included an attachment with a picture of the boy. The trusting recipient did not know that it was a hoax and forwarded the e-mail (and attached picture) to every employee throughout the United States, including the company’s President and all the senior executives. I reacted to the message immediately, confirming that it was indeed a hoax and sent e-mail to all users in the company explaining that the contents of the message was a hoax and to ignore it. I also briefly mentioned the nature of chain mail e-mails and the consequences of spam. Even though I sent the "ignore" e-mail within 10 minutes of receiving the chain mail, at least 8 users replied stating that they had already forwarded the e-mail to families and friends, with the majority of the messages going to their work e-mail addresses. Upon arriving at the desk of the person who sent out the chain e-mail, she described how upset she felt. She was tricked into sending this e-mail that appeared legitimate and "for a good cause". It made her feel humiliated and believed that her trust was violated. Other users who forwarded the chain mail e-mail expressed similar sentiments. I felt bad for my co-workers, as I knew that the spam incident had a negative impact on their trusting nature. Hitting the delete key can easily erase a spam message and it ends there, but as this story establishes, sometimes spam messages can have harmful consequences.

Spam lists are a cheap commodity and can be easily purchased on the Internet. If factors such as bandwidth, processing power, disk usage and administration costs are considered for a list of 1,000,000 e-mail addresses, then the time and effort required in delivering spam messages to the users on this list is astonishing. If 10% of the e-mail addresses on this list are invalid then it should be taken into account that an undeliverable message is sent to the ISP and the sender. Since the spammer's address is probably invalid in order to hide their identity, this increases the amount of messages going back and forth. It is usually the case that there is a number of attempts made by the mail server to deliver e-mail before finally giving up, so there could be 1,000,000 messages just to handle the invalid addresses on this one list. It is important to note that the spammer does not pay the cost of these messages.

1.2 Objectives

Spam is rising exponentially at a rate of 1000% per year [3]. In May 2003, the general manager of Microsoft's anti-spam technology and strategy group stated: "Spam has reached epic proportions and we are in a crisis situation". This was the view from a company that sends 20 million marketing e-mails each month. There are systematic efforts underway to analyse spam, as it is regarded not just as a nuisance but a serious threat to the very existence of the Internet. SpamArchive.org was launched in November 2002 with the aim to collect 1.5 million spam messages in its first year. It is a resource that provides a database of spam to be used for developing, testing and benchmarking anti-spam tools. There are other research centres and organisations that collect and analyse spam also, so is there a need for more spam analysis research?

In order to develop a spammer profile and carry out a worthwhile analysis on spam, it is necessary to understand the underlying factors on why spam exists in the first place. The objective of this thesis is to create a more rounded view of spam, including the types of spam currently on the Internet. The tools used to send spam, anti-spam techniques, spam targeting and current legislative responses to spam are crucial to this research as they provide the bigger picture in spam research. These are just some of the topics in this study. A more detailed breakdown of the core objectives is as follows.

1.2.1 Spam Analysis

A wide variety of spam will be analysed. Spam will be retrieved from honeypot (bait) addresses, which were created specifically for this research. In this analysis, information will be gleamed from present-day spam in an effort to examine the tricks and techniques used in the effort to bypass spam filters. Keeping current with the state of the art spamming methods will give us an insight on how spam is evolving. The ultimate goal for the anti-spam community is the complete eradication of spam but a more realistic view would be to develop systems that would keep spam at a controlled level. The spam analysis in this study will contribute to the existing body of knowledge in spam research.

1.2.2 Creation of a Spam Profile

A comprehensive spam profile will be created in order to define and categorise spam. The various definitions of spam will be examined with the aim to determine if any are globally acceptable. The economic and legal aspects of spam will be discussed, as will the role of government in addressing spam. Spam scams will also be investigated as they can have more serious consequences.

1.2.3 Creation of a Spammer Profile

Since most spam comes from a small percentage of spammers, an attempt will be made to develop a profile of a spammer. The focus will be on the various methods used by the spammer to send spam. During this research, I was able to download spamming tools from the Internet that could have enabled me to have a spamming operation up and running within a day, sending out hundreds of thousands of spam messages to victims worldwide. It was surprisingly easy to obtain the tools of the spam trade and it is not difficult for a computer-savvy novice to become a spammer. No programming skill is required and there are multiple resources on the Internet that provides guidance to spam newcomers. To become a professional spammer, companies sell software that has stealth features to hide the spammer's identity and mask routing information. They also provide filter avoidance tools. If the spammer deems the risk too great, they can become truly anonymous and use a bulletproof service that would send out their spam from places like China, where the long hand of western anti-spam legislation cannot reach. The spammers usually pay for these services using wire transfers, so tracking the source can be very difficult if not impossible. These bulletproof services guarantee over 99% uptime and most offer web site hosting where spammers can advertise their products and services without the worry of getting shut down by their Internet Service Provider.

The risks associated with spam include the alienation of the vast majority of recipients of the spam message, the loss of creditability in the spam message, violation of the ISP's acceptable usage policy that could result in loss of service, and civil/criminal liability. Do the risks outweigh the benefits and why would a spammer try this form of advertising if they stand to lose so much? This research will attempt to answer these questions and understand the logic behind the people who send spam.

1.2.4 Examination of Spam Targeting

As there is a lot of information and research material already collected on spam, I will take the research to a new level and investigate the concept of spam targeting. The fundamental question surrounding spam targeting will be answered i.e. can the spammer target certain individuals or groups of people, or it just pot luck who receives spam?

There is a famous cartoon depicting a dog at a computer console browsing the Internet with the caption “On the Internet, nobody knows that you are a dog” [4]. The creator of the cartoon was trying to portray the anonymity of the Internet user. This is in fact reality as users can hide behind any pseudonym and profile, but in this dissertation there will be an examination to find out if the spammer’s targeting mechanism is accurate. There are dire consequences if spam messages are directed to a wrong group of individuals due to inaccurate targeting, e.g. e-mail with adult content sent to children. Recipients of such messages should be “qualified” (fall into the correct category) in order to receive them.

1.3 Document Outline

This section outlines the contents of the chapters in this document

Chapter 1: The **Introduction** looks at the motivation behind this research and the desired objectives.

Chapter 2: **State of the Art** deals with current research on spam and will examine areas of interest such as economic and legislative factors in the spam industry.

Chapter 3: **Spammer Tactics and Tools** deals with the spammer, their tools of the trade and the services that support them. Spammer tactics, tricks and scams are also examined.

Chapter 4: **Anti-Spam Techniques** presents the methodologies used to fight spam. Ways to avoid spam are also discussed in this chapter.

Chapter 5: **Implementation** examines the methods that were used to create the honeypot e-mail addresses, which were then positioned to receive spam.

Chapter 6: **Evaluation** presents the results and findings in this research.

Chapter 7: **Spam Targeting** concentrates on the methods used by spammers in their attempt to target audiences for specific spam messages. There is also an examination of the accuracy of this targeting.

Chapter 8: **Conclusion** includes a summation and suggested future work on spam.

Chapter 2

State of the Art

2.1 Spam Overview

Spammers argue that spam is just a form of free-expression, but anti-spammers say that it stifles freedom of speech, as it keeps users away from newsgroups and other Internet-related discussion forums. This chapter deals with the current state of the art in spam-related topics.

2.2 Spam Profile

The reason why spam exists is because it works. E-mail marketing elicits between 0.1 and 1 percent follow-ups. Direct mail gets a 1 - 3 percent response, and SMS advertising attracts a staggering 10 – 20 percent response rate [5]. A profile of spam is introduced in this section.

2.3 Definition of Spam

“SPAM®” is a brand name of luncheon meat and is a registered trademark of Hormel Foods Incorporated. On the other hand, “spam” (all lowercase) is a term that was light-heartedly adopted by the Internet community after a famous Monty Python sketch, to label unsolicited mass postings on USENET newsgroups. The focus of this research will be on the latter meaning of the word.

The word “spam” is also used in Internet vernacular to describe the use of any words, scripting, HTML (Hypertext Markup Language) code, or programming on web pages that is not meant to benefit the end-user experience. “Spamdexing” (or simply spam) is the term coined to describe the taking of extreme or excessive measures to achieve top search engine positions. Other interesting definitions of spam exist, one of which is the following standard technical definition, which is used by mail-abuse.org [6] and the anti-spamming group at the Spamhaus Project:

> An electronic message is “spam” IF: (1) the recipients personal identity and context are irrelevant because the message is equally applicable to many other potential recipients; AND (2) the recipient has not verifiably granted deliberate, explicit, and still-revocable permission of it to be sent; AND (3) the transmission and reception of the message appears to the recipient to give a disproportionate benefit to the sender.

The word “spam” now appears in the English dictionary [7] and described as:

> Unsolicited e-mail, often of a commercial nature, sent indiscriminately to multiple mailing lists, individuals, or newsgroups; junk e-mail.

Another definition [8] that is used by companies that block sources of spam is:

> Internet Spam is one or more unsolicited messages, sent or posted as part of a larger collection of messages, all having substantially identical content.

Some definitions classify jokes and chain letters as spam messages. The Supreme Court of the State of Washington formally endorsed the following definition:

> The term spam refers broadly to unsolicited bulk e-mail (or junk e-mail), which can be either commercial (such as an advertisement) or non-commercial (such as a joke or chain letter).

The standard accepted definition of spam is unsolicited commercial e-mail (UCE) or unsolicited bulk e-mail (UBE). However, in my opinion an e-mail message should be considered spam if it is unsolicited AND has no benefit to the recipient.

I believe that it is very difficult to have a globally acceptable definition of spam. In a scenario where a person received an unsolicited commercial e-mail but found that the spam was useful e.g. it allowed the recipient to purchase an item of clothing at a discount, then was this e-mail really spam and if so can it be classified as “good” spam? In this situation, both the sender and recipient received benefit from the e-mail, so by using the technical definition, this e-mail cannot be regarded as spam. It follows then that if this e-mail is not spam, then filters should not block it. But it may be spam to other users who deem the e-mail to be of no benefit to them, which makes it spam by the technical definition. This is why I believe that the decision should trickle down to the individual user. The e-mail in the previous example should be classified as spam because for the majority of its recipients it is has a disproportionate benefit to the sender. However the e-mail should be delivered to the recipient’s “bulk” or “possible spam” folder, where the user can review it. The individual should have the final decision on whether it is spam or not. It is the very same concept as art. One person may view a painting as a masterpiece, while another person may view the same painting as abhorrent and demand that it should be censored. It is not safe to have the e-mail administrator, filter developer, or company executive decide what is spam. Filters should be in place to catch probable spam and allow the end-user to decide if it is or not. These self-learning filters could then become more accurate as the spam corpus (created by the individual) increases.

2.3.1 Characteristics of Spam

There is a fine line between legitimate e-mail marketing and spam. Generally, spam has some or all of the following characteristics:

- Return address is not valid. If you try replying to spam e-mail, you will get a delivery error.
- Forged headers. To hide the origin of the e-mail, routing headers are usually forged, which makes it very difficult to trace persistent spammers.
- The identity of the recipient is irrelevant because the e-mail is valid for many other recipients.
- Dictionary attack address. If the ‘To’ address line is examined, you may see different variants of the recipient’s e-mail addresses. For example, if your e-mail address happened to be JohnSmith@hotmail.com, you would see JohnSmith@yahoo.com, JohnSmith@aol.com, JohnnySmith@hotmail.com etc.
- Subject line has no bearing on the content of the e-mail. Bulk e-mailers use programs to randomly generate characters in the subject line in order to bypass spam filters.

- E-mail content is of a dubious nature. For example, some of the topics would cover get-rich schemes, body enhancements, etc. Others are well-known frauds such as various forms of the Nigerian e-mail scam.
- Unsubscribe does not work in spam e-mail. If you try to unsubscribe from spam e-mail, it is often the case that the link does not work, or opens up an advertisement's web site. What it does confirm however is the legitimacy of your e-mail address, which would be duly added to another spam list of verified e-mail addresses. It should be noted that a small percentage of unsubscribed links will work, but in the course of this research only one link appeared to work in 12 spam e-mails.
- May contain hidden scripts. If spam contains HTML, it may contain hidden JavaScript, which can open up web sites and activate advertisement popup windows.

2.3.2 Spam Breakdown

According to Brightmail, a leading anti-spam company, 88% of spam is non-porn related, which may surprise a lot of people [9]. One-third of all spam that users receive fall under the category of "products" e.g. cheap ink cartridges etc. 24% of spam deals with financial offers e.g. credit card offers, low cost mortgages etc. Spam scams e.g. the notorious Nigerian scam, accounts for only 5% of all spam. The remaining percentage can be classified as miscellaneous e.g. body enhancements.

2.4 Top Spam E-mails

Here is a sampling of the top spam e-mails currently roaming the Internet:

- "URGENT AND CONFIDENTIAL". This is the Nigerian spam asking e-mail recipients for bank account information in order to deposit large amounts of money from Nigeria.
- "GET A FREE PASS TO THOUSANDS OF XXX SITES!" Spam e-mail advertising pornographic web sites.

- “Protect Your Computer Against Viruses for $9.95”. An offer for a cheap anti-virus software program.
- “Verification Department”. This is a credit card scam e-mail that purports to be from a credit card company offering the e-mail recipient the opportunity to receive a credit card, irrespective of credit history or employment status.
- “Online Auction Marketing Secrets!” Another spam scam that claims to reveal secrets to online auctioning success.
- “Important news Kuira”. This is a prolific spam that has many variations with the sole purpose to sell septic tank systems.
- “Printer Cartridges – Save up to 80% - free shipping offer”. Spam e-mail to sell printer cartridges with enormous savings.
- “$100 FREE, Please Play Now!” This spam e-mail is an offer to join an online casino, where you will receive free credit of $100.

2.5 False Positives

False positives are legitimate e-mail messages that are classified as spam and the amount of false positives that a filter creates is crucial to the success of the filter. False negatives, on the other hand, are spam e-mails that get through the filtering system. Spam filters should allow more false negatives than false positives. Concerns about false positives are one of the reasons why the corporate user sees so much spam. The fear of losing a lucrative business deal as a result of a false positive has led to many IT managers letting the end-user decide what is spam or not. Spam filter vendors argue that false positives account for only 0.001% of filtered messages. However, for a large company that receives 50,000 e-mails on a daily basis, that small percentage can make a difference to the company’s bottom line. One of those false positive e-mails could be of vital importance to the company.

Another interesting scenario is where you have medical and pharmaceutical companies that receive a large number of e-mails containing words common in the lexicon of spam e-mail, such as sex, Viagra, drug etc. It is vital that these e-mails get through the filtering process. A common method in dealing with false positives is to forward all e-mails flagged as spam to a special folder e.g. bulk e-mail. Users can then examine these quarantined e-mails, which will be automatically deleted after an expiration date. A large company may have to store many gigabytes of this type of e-mail, but with the low cost of storage devices, and the high cost of losing a big contract due to a false positive, the choice of keeping the filtered messages is the obvious one.

2.6 Economics of Spam

Spam has become an industry and the economics of this fledgling industry are impressive even when the financial dealings of bulk e-mail companies, who operate within the confines of the law, are excluded.

2.6.1 Cost of Spam

Spam has become an industry, mainly due to the low cost involved in sending spam. The average spam costs the spammer 0.00032 cents, so sending bulk e-mail is very cheap. However, the reason why it is so cheap is because of the notion of cost shifting, which means the cost to the recipient is greater than the cost to the sender. A spammer can use a simple modem to send e-mail to thousands of Internet users. The cost to the recipient could include resources such as processor time, disk storage, bandwidth, anti-virus software, and spam filtering software. If you have a finite storage space with an ISP, spam e-mail may fill up your allocated space and block legitimate e-mail. The cost to the ISP is passed on to the consumer. A single piece of spam e-mail from source to destination could pass through a number of ISPs along the way, with the burden of delivery placed on the ISP (and not the sender).

On the other hand, these very same ISPs sign lucrative contracts with known spammers. Under the guise of not regulating content, some backbone providers are happy to sell 45Mbps DS3 circuits to bulk e-mailers. The price tag for one of these circuits from a Tier-1 provider could range anywhere from $30,000 – $40,000 per month. A well-known Louisiana-based spammer claims to use three of these circuits in order to send 84 Million bulk e-mails a day.

Spam does not only target e-mail inboxes. Mobile phones, Blackberry pagers, PDAs (Personal Digital Assistants) and other mobile devices are also targets of spam. If a company whose sales force communicates with each other and gain access to their corporate e-mail using Blackberry pagers, then the cost of spam is substantial. With these systems, the user is charged by the amount of bandwidth that is consumed, so each spam message has a varying monitory cost. Spam has also begun to invade other wireless platforms where the user is charged on a per-message basis.

2.6.2 The Spam Industry

The spam industry can be regarded as those who make a business out of sending spam and those who business revolve around anti-spam software and services. It is interesting to see how the spam industry mimics other marketing strategies. During the holiday seasons of Christmas and Thanksgiving in 2002, there was a marked increase in the amount of spam being detected [10], corresponding to the marketing blitz strategies of consumer industries in an attempt to encourage people to spend more during these lucrative periods.

There are organisations that support the rights of spammers to send bulk e-mail. In the United States, the Direct Marketing Association (DMA) is a well-funded lobbying group that maintains a pro-spam stance. The DMA believe that any non-porn, non-fraud bulk e-mail should not be classified as spam. The association represents approximately 5,000 companies that send both postal and electronic mail. At its conference in October 2002, it was reported that 2/3 of its members generated increased sales due to e-mail marketing. Some of the members stated that an e-mail campaign could result in up to 12 times the response rate of ordinary postal mail. With statistics like this, it is no wonder why advertisers find this medium to be very attractive.

Some spammers operate out in the open as registered companies, where they pay taxes and generally adhere to accepted business practices. One Ohio-based spammer in the United States has 60 employees and yearly sales of $12 million. They churn out 5 billion spam e-mails every month for everything from nutritional supplements to house improvements. The owner of the company asserts that they have hundreds of corporate clients, including publicly traded companies.

2.6.2.1 Reputable Bulk E-mailers

Is there such a thing as a reputable bulk e-mailer? Does good spam exist? For example, would you not like to hear about the price reduction of an electronic gadget, which you intended to buy when the price was right? Is it fair to label a perfectly legitimate e-mail marketing company as a spammer? Is it right to blacklist a company just because they send bulk e-mail? Can an online opinion polling company be regarded as a spammer?

Hopefully, these questions will get the reader to consider the other side of the bulk e-mail business. These are the companies that stake their reputation on respecting opt-out responses from e-mail recipients. They operate openly and within the full confines of the law. Some companies argue that organisations like AOL want to keep advertising revenues to themselves and stopping bulk e-mail to the millions of AOL users is an unfair business practice. Forrester Research estimates that corporations will spend approximately $6 billion on e-mail marketing in 2005. Traditional companies such as Ford and Procter & Gamble are joining new age companies such as Amazon.com in the highly lucrative e-mail marketing strategy. E-mail facilitates access to the global marketplace and allows companies to achieve sales in far-reaching places of the world. It is simple for bulk e-mail recipients to click through to a company's web site and make a purchase.

2.6.3 The Anti-Spam Industry

The European Union estimates the cost of spam worldwide at $8 - $10 billion annually. An anti-spam industry has evolved with many companies hopeful to reap the rewards of the cash injected into the industry by eager venture capitalists. The industry can be categorised into two main areas of concentration: anti-spam software and anti-spam services. Some businesses are capable of offering both solutions, and are comparable to other companies in the software/services industry. Their offerings are geared to the large enterprises, service providers and educational institutions, right down to the home user.

Spam analysis can be very costly in terms of time and processor usage, and if it is performed on the contents of every e-mail message, then there is a noticeable delivery overhead for organisations with substantial numbers of users. Mail throughput is very important to many companies. This in turn leads to competition within the anti-spam industry to produce faster and more reliable filters. There were approximately 200 businesses in the anti-spam industry at the start of 2003 with more to follow, as the projected market is expected to grow to $181 million in 3 years [11].

The irony with the anti-spam industry is that some of the companies that offer anti-spam solutions have used spam to develop business. At least two of these companies acquired spam mail lists and sent e-mail to addresses on the lists offering their anti-spam services [12].

2.7 Legal Response

The core question when you talk about a legal response to spam is: Will it make a difference? If a spammer is intent on sending out bulk e-mails, they will find ways to circumvent the laws, or just ignore them altogether. Legislation will make a difference to bulk e-mail marketers, as it will put controls into place that will have to be adhered to. But new laws will make little difference to the core spammer groups, whose tactics are tantamount to fraud to begin with. The delicate element in passing anti-spam legislation in the United States is that spam is regarded as commercial speech, and as such, protected by the first amendment to the constitution.

The states of Washington and California have laws that make spamming legal if some basic rules are followed. They are attempting to stop spammers from hiding “underground” by using fake e-mail addresses and using open relays. These laws declare that if you want to send unsolicited e-mail, then you will have to provide your full name, address, phone number, valid reply e-mail address, and the option to be removed from further mailings. These laws also prohibit the use of open relays. Some organisations like CAUCE (Coalition Against Unsolicited Commercial E-mail), the Centre for Digital Democracy, and the SpamCon Foundation would like to outlaw spam completely. They believe that as spammers impose the cost of the e-mail on the recipient, the correct policy is to prohibit it. They would like a similar law to the Telephone Consumer Protection Act of 1991 (TCPA), which outlawed junk faxes. Prior to this act, companies were inundated with commercial faxes, which not only tied up fax machines, but also shifted the cost burden to the recipient. This law was very successful and after it was passed, companies very rarely received junk faxes. Anti-spam groups argue that the same results could be achieved with a similar junk e-mail law and in the United States they are actively lobbying to have it introduced to Congress.

2.7.1 Legal Action against Spammers

In the Supreme Court of California, the justices were divided on the issue of whether the sending of e-mail to a company could constitute trespass. This case stemmed from a spamming incidence where a former Intel employee sent 30,000 e-mails to current Intel employees disparaging the company. The ruling is expected in late 2003 and could have major implications for other spam cases. Two lower courts has already sided with Intel agreeing that the former employee’s messages constituted trespass to the company’s servers.

There is an organisation called the SpamCon Foundation [13], which has a law centre web site that tracks information about spam cases, texts of relevant laws, and legislative news regarding spam. Transcripts of actual court cases against spammers are available on their web site. It also provides stories of individuals who successfully sued spammers and received monitory awards.

In Buffalo City, New York, the man known as the “Buffalo Spammer” was arrested and arraigned in May 2003. He allegedly sent over 825 Million spam e-mails, but his arrest was not as a result of the spam he sent, as sending spam in New York is not illegal. The criminal offence was his use of forgery and identity theft to send out the spam. He stole the identities of two residents and opened up Internet access accounts with Earthlink, Inc. He then forged the e-mail headers sent from these accounts. He eventually opened more than 343 e-mail accounts with stolen identities, which cost Earthlink over $1 million.

Earthlink, the 3rd largest ISP in the United States, won a $16 million settlement against this spammer the week prior to his arrest. In April 2003, America Online sued five spammers accusing them of sending more than 1 billion spam e-mails to its subscribers. These lawsuits had the support of its subscribers as 8 million of them sent individual complaints reporting spam that they received using a “Report Spam” feature that the company introduced 6 months prior to the lawsuit. America Online sought damages of $10 million and an end to the spamming. Even though most of the defendants in this case are unknown, filing the lawsuits will give AOL the authority to subpoena the spammers’ service providers and reveal their identities.

To date, the largest legal action taken against spammers occurred on June 17th, 2003 when Microsoft filed 13 civil suits in the company’s home state of Washington against U.S. spammers and filed a further two suits in Great Britain. Microsoft claimed that the defendants in the 15 suits were responsible for sending more than 2 billion spam messages to the company’s customers. The lawsuits in Great Britain also accuse the spammers of illegally harvesting Microsoft e-mail addresses for use in building spam lists. Microsoft pointed out specific e-mail sending practices used by the spammers, including the use of deceptive and fraudulent e-mails, bogus virus warnings and unsolicited messages relating to pornography. Many victims of spam worldwide have received spam messages from the “Hotmail” service, which is now a Microsoft-owned company, but the company claims that spammers spoofed most of these addresses. Washington State is known to have the strongest anti-spam laws in the United States, while in Great Britain the lawsuits are based on that country’s Misuse of Computers Act.

There appears to be a cohesive attempt by some of the larger companies like Microsoft, AOL Time Warner and Earthlink to take their spam grievances to court in order to seek financial remedies. These suits also attempt to bring the spammers out into the open as the companies use the ability to subpoena documents that would identify the spammers.

It should be mentioned that irrespective of spoofing, ISPs claim that a substantial amount of spam is sent from Microsoft’s Hotmail servers, which are abused by spammers using deficiencies such as open relays and proxies. As a consequence, in May 2003 Microsoft imposed a limit of 100 outgoing e-mails per day on all Hotmail accounts in an effort to deter spammers. Another point of interest is that Microsoft sends out tens of millions of marketing e-mail messages every month to Internet users and it appears that these messages bridges the fine line between spam and legitimate e-mail. It seems that one company’s spam is another company’s marketing campaign.

2.7.2 Legal Action against Anti-Spammers

Once again, it should be noted that there is a fine line between a spammer and a legitimate bulk e-mailer. One such bulk e-mailer, Yesmail, won a temporary restraining order against MAPS (Mail Abuse Prevention System), prohibiting it from placing the reputable e-mail marketer on the black hole list of known spammers. The online opinion polling company Harris Interactive sued MAPS and the ISPs that blocked their e-mail messages from getting to end users. Putting Harris on a blacklist caused irreparable harm to the company’s image and reputation.

At the other end of the spectrum, you have actual spammers suing organisations that have blacklisted them. One such case is Case Number 03-80295 filed in the United States District Court, Southern District of Florida by EMarketersAmerica.org on behalf of anonymous senders of Unsolicited Bulk E-mail. The lawsuit was against the Spamhaus Project [14]. It appears that the same lawyer who filed the case formed EMarketersAmerica.org only 4 weeks before the lawsuit. The suit sought an injunction to stop Spamhaus from publishing the IP addresses of the anonymous entity that the lawyer represented. The Spamhaus Block List is a blacklist of IP addresses used by companies and Internet Service Providers to block e-mail from confirmed spam senders. It so happens that the lawyer in this case is not only the plaintiff, but also the personal lawyer of America’s top spammer.

You would expect the people at WorldJustice.com to be a group of altruistic individuals dedicated to the cause of stamping out world hunger or defeating injustice worldwide. These goals however are not included in the organisation’s charter. Their purpose is to file class action lawsuits against ISPs and organisations that block spam and who have injured and maligned online entrepreneurs due to “possible illegal interference”. They also claim that federal anti-trust laws may have been breached by these “self appointed police agencies”. They are obviously taking aim at ISPs that block spam and companies that provide blacklists such as mail-abuse.org and their MAPS RBL (Mail Abuse Prevention System, Real-time Blackhole List).

2.8 Anti-Spam Legislation

This anti-spam legislation section looks at the role of government in addressing spam and examines the approaches used by governments worldwide.

2.8.1 United States

In the United States, each state has its own laws and regulations that are independent of the other states. There are currently no national (federal) laws that regulate unsolicited e-mail, but 33 individual states have enacted anti-spam laws. Notable in the absence of laws is Florida, which is the number one location of spammers in the world.

In April 2003, Virginia enacted the toughest anti-spam legislation to date by imposing harsh criminal penalties for sending spam through deceptive means. Spammers who send unsolicited bulk e-mail from/to Virginia using fake return addresses or open relays face prosecution that could result in jail terms and sizeable fines. One should note the interesting fact that the Virginia governor signed the new law at the Dulles-based headquarters of AOL. Also in 2003, Senator Debra Bowen introduced an anti-spam law in California that is a replica of the successful junk fax law, but it has to get through the state’s legal process before coming into law.

In early 2003, there were a number of proposed spam laws before the U.S. Congress in response to the spam problem, but it appears that these bills attempt to regulate rather than outlaw spam. The Burns-Wyden CAN-SPAM Act requires the senders of unsolicited bulk e-mail to put “ADV” in the subject header of e-mails. This is hardly a deterrent to spammers. “The Reduction in Distribution of Spam Act”, sponsored by congressman

Billy Tauzin who is chairman of the House Committee on Energy and Commerce, requires recipients to respond to spam so as to “Opt-Out” from future mailings. The House Committee on Energy and Commerce is the committee that examines all anti-spam laws before passing them on to Congress and the Senate for approval. Any of the Acts passed by the U.S. Congress would supersede existing state laws, including the Virginian law that makes spamming a criminal offence. However, it does look like 2003 may be a key year in anti-spam legislation in the United States as there were 9 anti-spam bills introduced to Congress with some now at the Senate voting stage. These bills target different aspects of spam such as dictionary attacks, misleading headers, non-functioning return addresses, resale of e-mail addresses and other similar spam characteristics. Due to the recent outrage at the extent of the spam problem, many of these bills were enhanced with versions that increased spam penalties e.g. from $10 to $100 per e-mail sent with misleading header information. One of the bills, the Criminal Spam Act of 2003 would make it a crime to hack into mail servers to send spam. Of course, these laws would only be applicable to spammers in America and could easily be evaded by using offshore resources.

2.8.2 Europe

From October 2003, a European Union directive will make unsolicited bulk e-mails illegal across member states. This directive will require prior consent or an existing customer relationship before unsolicited e-mails can be sent to individuals for direct marketing purposes. The problem with this law is that it only deals with spam within the European Union and does not address spam originating from outside the member states.

2.8.3 Legislative Action Worldwide

The following countries (as of June 2003) have enacted opt-in laws, where permission has to be first received from an individual, before sending them commercial e-mail:
Austria, Denmark, Finland, Germany, Greece, Hungary, Italy, Norway, Spain, Poland and Slovenia.

OECD is the Organisation for Economic Co-operation and Development, which is based in Paris and has 30 member states that include the United States, Japan, Germany and Great Britain. They issued guidelines in June 2003 that suggested an international approach in the fight against unsolicited commercial e-mail, as they believed it was the medium that exacerbated the international fraud problem. Based on these guidelines, OECD wanted a cooperative approach to the spam problem: “They should build on existing projects to gather and share information, including consumer complaints and notifications of pending investigations and cases, through online tools and databases.” The Federal Trade Commission in the United States urged widespread adoption of the guidelines and singled out spam and the World Wide Web as the primary tools in fraudulent international marketing. The OECD guidelines should make it easier for member states to collect evidence from foreign agencies and to cooperate in the prosecution of spammers.

2.9 Success of the Legislative Approach

The question should then be asked about the success of the legislative approach. During my research on spammers, I was surprised to find one particular spamming group who removed the e-mail addresses of users from ISPs in the states of Virginia, Washington and California from over 15 million e-mail addresses that they had in their possession. This was in direct response to the strong anti-spam laws that are in place in these states. It does not mean that all spammers will respect these laws, but since the spammer community share information and techniques among their members, it may mean that spammers will not risk being sued by users in states and countries where there are strict anti-spam laws.

2.10 Opt-In versus Opt-Out Legislation

Over 90% of all spam hitting Europe comes from mostly Florida-based spammers [15], so how will the implementation of the European opt-in legislation due in October 2003 affect this onslaught? The opt-in rule requires prior permission before sending unsolicited commercial e-mail. Unfortunately, the European legislation only applies to European spam and a similar law would have to be implemented in the United States in order to see any reduction in the amount of spam. The United States, on the other hand, is taking the opposite approach in which the spam recipient would have to respond to a spam message requesting no further e-mails i.e. the opt-out approach. Some anti-spam activists believe that this in fact legalises the sending of spam. The "Reduction in Distribution of Spam Act", currently before the U.S Congress, would legalise the sending of spam and bulk e-mail as long as there was an opt-out option for the recipient. There are numerous small businesses in the United States who would like to send bulk e-mail to potential customers but fear the stigma of being labelled a spammer. This Act would legitimise the sending of unsolicited bulk e-mail and if you consider the large number of small businesses and the 92 million U.S. inboxes currently in use, then it certainly appears that this Act would just exasperate the spam problem.

Double opt-in is the process of confirming an initial list subscription request. As so many spammers claim that recipients have opted-in to receive their messages, the double opt-in option requires the list owner to receive confirmation from list members agreeing to the subscription.

2.11 Chapter Summary

2003 is a pivotal year in the legislative approach to the spam problem. With the new prior consent law in Europe and the strong American anti-spam laws currently before Congress, it appears that legislation will be a formidable component of the multi-pronged attack on spam. Companies are also taking the legal battle to the spammer with many lawsuits being filed in an attempt to identify spammers in order to recoup financial losses and to force the cessation of spamming from known sources. A profile of spam was introduced in this chapter where the definitions and characteristics of spam were discussed. The economic factors surrounding spam, which is core to the proliferation of this type of bulk e-mail was also presented.

Chapter 3

Spammer Tactics and Tools

3.1 Overview

There is a multitude of anti-spam techniques used in the fight against spam but it is apparent that spammers are not slowing down in their efforts to counteract these techniques. It is at the stage in the fight where spammers actively recruit anti-spammers to develop systems to defeat filters. There are spam services available that allow spam messages to be processed through filters prior to distribution in order to fine-tune the e-mail message and improve the chances of spam getting through commonly used filters.

In order to effectively fight spam, a suitable profile of the spammer needs to be developed. The spammer profile is introduced in this section. There is also a discussion on the tactics, tools and scams employed by spammers.

3.2 Spammer Profile

The names and addresses of the most prolific spammers are readily available on the Internet. Some anti-spam web sites and enthusiasts go into great detail describing the spammers' modus operandi such as the aliases that they use, P.O. box locations, business and home addresses, banks that they use, ISP names and IP addresses, domain names and e-mail addresses, and the web sites that are used to advertise spam products and services.

Spammers can be categorised by the following descriptions:

- **Casual Spammer**. These are the Internet users that forward chain letters and pyramid schemes and are generally regarded as just a nuisance. It is interesting to note that some providers, including Yahoo, regard this type of e-mail as a violation of their usage agreement and could result in account termination.

- **Small Scale Spammer**. The vast majority of spammers fall into this category and they are sometimes called “spam kiddies” who do not understand how open relays and filtering software work. They tend to use spamming toolkits and mail lists that were purchased on CDs or downloaded from the web.
- **Hacker Spammer**. This is the more sophisticated spammer who actively creates software to circumvent spam filters.
- **Large Scale Spammer**. This is the category of spammer that does the most damage and creates most of the spam on the Internet. They are usually well funded and run their operations as a business.

Boca Raton in Florida is regarded as the spam capital of the world [16], with 40 out of the world’s 200 most prolific spam operations located there. Spammers tend to know each other and there is great deal of sharing of information e.g. sharing lists of publicly known SMTP open relays and mailing lists. But not all spammers send e-mail advertising services or products. The election campaign of Jeb Bush, who happens to be the governor of Florida, was accused of spamming the state electorate with e-mail messages urging voters to re-elect their candidate. Well-known companies such as Microsoft have also been accused of spamming consumers, with the millions of marketing e-mails that they send out every month.

I was curious to find out about the pioneers in the spamming business and my research led to 2 Arizona-based lawyers. On April 12th 1994, they cross-posted their “Green Card Assistance Services” to 6,000 newsgroups, at the same time. When the thousands of USENET members logged onto their particular newsgroup, they found the Green Card advertisement placed by the lawyers. Many of these users were outraged and responded with e-mails that reflected their feelings. However this piece of spam is pre-dated by an e-mail sent on May 2nd, 1978 by a salesman at DEC (Digital Equipment Corporation), who advertised the new computer that his company was selling, to everyone that had an e-mail address on the Arpanet (the predecessor to the Internet). The recipients of this e-mail were not too happy about the incident and discussed censorship as the solution to the problem, but eventually decided against it.

3.2.1 Spammer Motivation

The motivation for most if not all spammers is money. Jupiter Communications predicts that commercial e-mail marketing will become a $7.3 billion business by 2005 [17]. Spammers receive commission for every referral that is successful from their spam message. They can also receive an additional commission on every sale. Companies who do not want to send spam offer “Affiliate Programs”, where spammers send their e-mails for them. A self-confessed spammer in Oregon willingly shared his secrets with a local newspaper in May 2003 [18] and claimed earnings of $1,000 per week sending out 10 million unsolicited e-mail advertisements on a daily basis for various companies. It is known that some of the bigger spammers have criminal backgrounds and have been involved in other dubious entrepreneurial ventures, but this particular spammer was a former police sergeant who got into the business after joining a spammer bulletin board.

The information that he provided is very useful as he gave an insight into the financial rewards in the spamming business. The response rate to spam is extremely low (less than one tenth of one percent) but when recipients do respond to spam it can be very lucrative for the spammer. Viagra distributors pay individual spammers based on the number of sales, as high as $60 per $150 order, and with such hefty commissions it is no wonder why we see so many Viagra-based spam e-mails. Financial companies pay for spam recipients who request more information - $12 for mortgage leads and $5 for insurance referrals.

3.2.2 How a Spammer gets your E-mail Address

This is not an exhaustive list but spammers primarily use the following methods to obtain e-mail addresses.

Social Engineering: This method is used to convince people into giving up their e-mail addresses, for example via chain letters, and free offers.

List Buying: If the spammer does not wish to invest in software that extracts addresses, spam lists containing millions of users can be bought quite easily, which can be downloaded online or purchased on a CD. Bulk-email-lists.com is a notorious purveyor of e-mail lists that sells a package of over 200 million e-mail addresses for $499. On their web site they explain that they built these lists by buying or exchanging lists, and the remaining e-mail addresses were collected from web sites and newsgroups. They also state that their list of 200 million addresses are mixed up from various parts of the world and can be used for adult oriented sites. Spammers advertise the fact that some lists are built using customer databases, client rosters and e-mail addresses obtained from Internet Service Providers.

USENET Postings: If you post to a newsgroup, your e-mail address is there for the whole world to see. Spammers use software to extract addresses from newsgroups.

Web Sites: Similar to the software that spammers use to extract addresses from newsgroups, spammers also employ software to harvest addresses from web sites.

Mailing Lists: Mail list servers can be tricked into sending lists to spammers or they could use other illegitimate methods to obtain lists of subscribers to mailing lists.

Chat Rooms: There are programs that specialise in retrieving the screen names of chat room users. The screen names usually represent the first part of an e-mail address. The domain part of the e-mail address is retrieved from the location of the chat room e.g. AOL or Yahoo.

Yellow and White Pages: These pages have e-mail directories that can be easily harvested by a spammer.

Dictionary Attack: This is a brute force attack where spammer programs guess e-mail addresses by using multiple combinations of common names. This form of attack is described in greater detail in section 3.5.1.

Web Browser: Some people configure web browsers with their e-mail addresses for ease of use (e.g. to FTP files). Web sites can trick the browser into sending these addresses to the spammer.

Guestbook and Forms: Filling in an online form or signing a guestbook on a web site can lead to your e-mail address being compromised, as dishonest site administrators can sell it.

Replying to Spam: Even though the majority of the usernames that send spam are bogus, some of them are valid. Replying to valid spam addresses will lead to more spam.

Unsubscribing: Clicking on a link to unsubscribe from future mailings only verifies the authenticity of your address, which results in more spam. Some spammers though will honour a removal request.

More detailed explanations of some of these methods are discussed in Section 3.5. There are other sophisticated methods used by spammers to obtain e-mail addresses such as hacking into web sites and using finger daemons, but spammers who carry out such attacks are in the minority.

3.3 Spamware

Spamware is the software used by people in the spam business. The creators of these software packages prefer to call them e-mail marketing packages and there are numerous packages available on the Internet. It should be mentioned that the creators of many of these applications also offer "Enterprise Editions" in order for the spammer to operate in a continuous, expeditious mode. Many spammers run fire-and-forget operations where they send out as much spam as possible to a great number of users, then they close down the operation and move to another ISP or open relay. Plug-ins and gateways are available (at a price) for many applications that provide extra functionality to spamware, e.g. gateways that automatically convert e-mail messages (SMTP) to different formats such as MAPI (Mail/Messaging Applications Programming Interface) and SMS (Short Message Service) text messages. There are some companies that sell spamware as well as conventional software like anti-virus packages and operating systems. What is remarkable is the fact that these companies proudly advertise the seedy features of their spamware products such as the ability to use open relays, forge the header message ID, add a fake sender in the header, and include bogus received from/by lines that are intended to mask the spammer's identity.

3.3.1 E-mail Harvester/Extractor

E-mail Harvester software is designed to extract e-mail addresses from Internet sources.

The extractor that I used in this research did not work with proxy servers, but there are extractors available that will go through proxies e.g. Atomic E-mail Hunter. It supports multi-thread page loading and filters can be used to limit the scanning depth i.e. the dept that the extractor will travel within a web site – usually measured in levels. Other extractors will allow e-mail extraction using keywords, which is used to target audiences. It works by entering the topic into a text box. The program then searches for e-mail addresses from the top ranked pages retrieved from popular search engines. This keyword search can be a powerful tool for spammers who want to target spam to specific groups of individuals.

3.3.1.1 Extracting Addresses from Newsgroups

A variation on the e-mail extractor designed for web sites is the extractor that retrieves e-mail addresses of users who post on USENET newsgroups or series of groups. Not only will this software extract the addresses but it will also send the spam message. The software has the added feature of being able to extract the e-mail addresses from AOL chat room logs.

3.3.2 Desktop Server Software

The Desktop Server software package converts a home user's PC into an e-mail server with the primary purpose of sending bulk e-mail. This software package contains options that exports and imports large lists and has the ability to personalise spam, which means it can associate the recipient's actual name with the e-mail address. The spam message would then appear as if the e-mail was specifically written for the recipient with the aim of adding credibility to the message. It is known that people are more likely to read e-mail that begins with a reference to their name. Desktop Server 2000 was an application that I examined for this research that had the ability to cloak spam messages, with the stated goal of making it very difficult for "complainers" to find the spammer's Internet Service Provider. The application also had built-in functionality that automatically looked up DNS (Domain Naming System) numbers, instead of querying the ISP's DNS servers for this information. This software allows the spammer to be completely independent from the ISP. The only remaining item required to start a spamming operation is a pipe to the Internet, which can be provided by a simple modem connection.

3.3.3 E-mail List Verifier

The E-mail List Verifier application checks the validity of e-mail addresses. This is another powerful tool in the spammer's arsenal as it substantially cleans a mail list database that has been polluted by anti-spam devices, such as spider traps. See Section 4.2.6.1 for more information on spider traps. The verifier software checks the syntax of the e-mail addresses and flags incorrectly formatted addresses such as spambait!@aol.com. The software also checks the availability of the address domain, such as yahoo.com and hotmail.com, and flags non-existent domain. Some verifiers may go another step and check the validity of the e-mail addresses by establishing an SMTP connection with the mail server and testing the address there. This is of course very time consuming and the spammer has the option to skip this process. One has to remember though that every time that an e-mail address is verified at the source, it is using the resources (opening a connection etc.) at the host's mail server.

3.3.4 E-mail List Manager

The List Manager has features that remove duplicates, count e-mail addresses, extract e-mail addresses and names from lists, verify and validate lists, sort list files and filter lists based on specific words.

3.3.5 Targeting Software

There are programs developed for bulk e-mailers that are intended to target particular audiences. Targeting software allows spammers to collect e-mail addresses within a specific marketing area and send spam messages to that locality or group of individuals. There are obvious benefits for small businesses that would like to advertise their products or services locally, but in the hands of a spammer this software would still be indiscriminate, as it would target everyone in a location list irrespective of the content of the message. Prospect Finder 2002 is an application that searches the Internet by keyword, personal profile and by location, and then extracts e-mail addresses based on those results. There are e-mail address collectors available that are designed to target specific domains, such as Hotmail and Yahoo that contains millions of addresses. They can quickly gather large amounts of addresses from a specified domain, but they can also verify the authenticity of the address by connecting to the domain server. The creators of this type of software do not advertise the methodologies used, but I can only presume that some form of dictionary attack is being used. They claim that up to 6000 e-mail addresses per minute can be collected. Unfortunately, the load on the targeted domain mail server must be substantial if the verification process is used for thousands of e-mail addresses. The creators of one such application “Direct Email Collector” state that they open multiple connections to the mail server in order to verify the extracted e-mail addresses.

There are many other spamware products such as the Bulk SMS Sender, which is an application that allows SMS messages to be sent to mobile phones directly from the host computer. The IM (Instant Messenger) Mass Mailer gives the spammer the ability to send bulk IM messages. The CD E-mail Extractor is a program that searches the contents of a CD or DVD and extracts the e-mail addresses located on the disc.

3.4 Spammer Support Services

Spammers would not be able to do business in a cost effective manner if companies were not available to provide them with specialised services. There are also spammer newsletters and newsgroups that discuss the current state of their profession. These newsletters offer advice on how to buy and rent mailing lists, how to deliver e-mail to AOL, how to extract e-mail addresses from files and CDs, and how bulk e-mail marketing works. Since AOL is the largest online provider in the world, there are spammer tips available on the Internet that advise spammers how to format their messages so that it looks acceptable on AOL’s proprietary web browser.

3.4.1 Bulletproof Hosting

Most, if not all, reputable online providers have Terms of Service (TOS) or Acceptable Use Policies (AUP), which prohibit the sending of unsolicited commercial e-mail. Some spammers have web sites that promote their products and it is crucial for them to keep these sites running. In the event that these web sites are shut down due to violations of ISP agreements, it is a major ordeal for the spammer to get up and running again with a different ISP. Bulletproof hosting is a system where the spammer’s web site is guaranteed uptime, even if spam e-mail lists the web site in the message, which in normal circumstances would violate ISP agreement terms. In the United States and other countries, law enforcement agencies could still shut down these web sites, irrespective of the guaranteed uptime, but that depends on the hosting company residing in their jurisdictions. This is the reason why the majority of these bulletproof hosting companies are located in countries such as China. One bulletproof hosting company proudly proclaims on their web site that their servers are located in “some province in the highlands of China” [19]. A spammer can then set up their web site and advertise the product or service in the spam using the hosting company’s domain e.g. http://www.blackboxhosting.com/spamsite. Bulletproof hosting companies can be very discerning with their clients, where some may host pornographic web sites while others refuse to do so. The hosting service can be offered for as little as $299 per month.

ISPs in China not only host spammers’ web sites, but they are now also responsible for the transmission of spam that originates in the United States. The administrators of these ISPs are far beyond the legislative reaches of the U.S. and Europe. The problem with spam originating in China has become such a nuisance for some ISPs that they have blacklisted all of the IP address blocks assigned to China Telecom, which is the state-owned carrier of the backbone that serves the approximate 200 Chinese ISPs. This is a draconian measure that blocks legitimate e-mail along with spam. Some blacklist organisations have blocked only certain portions of China Telecom’s IP address range, with the hope that the governing body in China will eventually deal with the problem. Unfortunately, this is a measure that stifles freedom of speech in a country that sorely needs it.

3.5 Spammer Tactics

Spammer tactics are diverse so in this section there is a comprehensive examination of the various methods used by spammers to obtain e-mail addresses and send spam. There is also a very interesting section on why some spammers make every effort to detract the attention of those users and organisations that are more likely to react negatively to their spam.

3.5.1 Dictionary Attack

During a five-month period in late 2002 and early 2003, spammers conducted a massive dictionary attack to obtain the e-mail addresses of millions of users on the mail servers of Hotmail and MSN. In a dictionary attack, the spammer uses software to create every conceivable variation of a person’s name. They would use common first names and common last names and combine them. For example, they would try johndoe@msn.com, jonathandoe@msn.com, johnnydoe@msn.com etc. It is also called a brute force attack. The software can automatically generate millions of random addresses that are sent to providers and logs successful attempts so that the e-mail address can be added to a spam list. According to the Spamhuas Project, the source of the five-month dictionary attack on Hotmail’s system was servers in Beijing, China operated by American spammers.

3.5.2 Spambots

A spambot is a piece of software, usually written in C for speed and portability, with the primary purpose to crawl through the Internet looking for e-mail addresses. They are also called spiders, deviant web crawlers, harvesters or robots. The spambot needs a starting URL from which it collects e-mail addresses and other URLs. The addresses are then stored in a database for use in the creation of a spam list, and the URLs are stored as a source for other web links to crawl in order to locate more e-mail addresses.

3.5.3 Spoofing

When a spammer spoofs an e-mail message, they are making it appear that the e-mail is coming from a source other than the spammer, and they may even make the message appear as if it is coming from someone you know or a company that you trust. There are many reasons why a spammer would send a spoofed message. The primary reason is to trick the recipient into opening up the spam message, but another sinister reason is to retrieve sensitive information. If spam e-mail was spoofed to originate from a user’s ISP, then that user may be persuaded to send account and even credit card information. E-mail messages are also spoofed to hide the identity of the sender. It is relatively easy to spoof e-mail messages and the only way to defend against spoofing at the moment is by the use of digital signatures. By digitally signing an e-mail message, it provides an assurance that the message is from the actual sender. PGP (Pretty Good Privacy) has been used by businesses and home-users alike for many years to protect their e-mail messages, but digital certificates are also available. If an attempt is made to forge the contents, then the signature becomes invalid and the recipient is informed.

3.5.4 Bandwidth Theft

With the popularity of wireless networks and the prevalence of unsecured wireless access points, spammers are now targeting these networks to send out bulk e-mail. The modus operandi involves the spammer driving around hotspot locations (e.g. on or near a university campus) and when an unsecured access point is found, it is used to send out a deluge of spam. It is a stealth operation, with the spammer vacating the area before the system administrators finds out that the network is an unwitting source of spam. Another spamming technique involves a mobile spam command centre, where the spammer has a vehicle full of equipment to carry out spam operations. A worker at a small or medium sized ISP is then bribed for a certain (brief) number of hours on their network. Astonishingly, it is as easy as running a network cable out of a window or backdoor to the spammer’s vehicle.

Allowing spam through open relays has been a known vulnerability for some time, and many companies and ISPs have now blocked access to port 25 in order to close down the relaying of e-mail. However, spammers have turned to using insecure proxy servers instead, which are set up by many home users who have DSL (Digital Subscriber Line) or cable modem access to the Internet. Proxy servers allow the home user to share out Internet access to other users in the household. Unfortunately, it is often the case that the proxies are not configured correctly and can be used as a spam conduit without the owner even knowing about it. In July 2003, approximately 2000 PCs with high-speed Internet connections were hijacked by a rogue program, which used the connections to spam e-mail addresses with pornographic advertisements [20]. This Trojan program turned the victim's computer into a proxy server that fetched porn ads from unidentified servers and sent them, typically via e-mail, to the inboxes of many online users. Installing a reputable firewall and updating virus definition files is one way to block this type of stealth activity.

3.5.5 Bypassing Filters

Content filters are designed to look for occurrences of certain words in an e-mail message and the spammers' task is to get these words to the recipient without having the e-mail flagged as spam. Spammers bypass filters using one or more of the following methods. The image e-mail is an e-mail that contains nothing but an image or a link to an image. All of the words that a spam filter would generally catch are imbedded in the image. The filter could be redesigned to stop all e-mails with images, but this would also prevent legitimate e-mails with images. Spammers also use spaces and symbols between letters of a word in order to get the word pass the filter. For example, instead of using the word PORN, which would be caught by a filter, P O R N (with blank spaces between the letters) or P*O*R*N would be used. A variation on this tactic is to use a combination of letters and numbers e.g. VIDEO TAPE would be represented as V1DE0 T4PE.

HTML comments are also used to trick the filter. A user will not see HTML comments in their e-mail message. For example, a filter would read the word "mort<!-- rtvx934 --> gage" as a single token, but the comments are ignored in an HTML e-mail and the word "mortgage" is displayed. An HTML table is another tactic used to bypass spam filters. The words in a spam message are divided into the cells of a table, which are then rebuilt by the e-mail client (if it is capable of reading HTML e-mail). The filter just sees a complex string of code and text, and as there are no flagged words the message is not classified as spam.

3.5.6 Spam Lists

I was curious to find out the cost (to the spammer) of an e-mail address on a mailing list.

As of June 2003, the going rate was 0.0008 cents per address based on rates offered by millionsofemails.com. In addition to lists that can be purchased on CDs or downloaded from the Internet, there are clubs that specialise in e-mail lists and for a monthly fee members can get access to fresh addresses every week. Emaillistclub.com is one such club where members can also use the database of over 6 million addresses on the remove list. As part of the membership, spammers can fill out a form and receive targeted addresses based on their requirements. A verification service is provided to confirm the validity of the e-mail addresses. Of course instead of purchasing lists directly, an e-mail harvester can be used to extract e-mail addresses, which can be freely downloaded from the Internet.

3.5.7 Addresses that Spammers Avoid

It was interesting to find out that some spammers could be discriminating in who became victims of their spam. Even though spam lists contained millions of users, there was a concerted effort to avoid users that could be a source of consternation to the spammers. Users with GOV, MIL, US, EDU and ORG e-mail addresses were generally avoided. There was a substantial effort to remove the e-mail addresses of users from ISPs in the states that had very strong anti-spam legislation in place i.e. Virginia, Washington and California. The fear of a legal backlash was not worth sending spam to these states. Some spammers have filters that deliberately remove the e-mail addresses of known “complainers”. These complainers were users who requested to be removed from future mailings or users that clicked on an unsubscribe link. Filters are also used to bypass webmasters and administrators of mail servers and web sites.

3.6 Spammer Tricks

Spammers have invented some very innovative tricks in the pursuit of their trade. Some of these antics are discussed here, along with viable solutions to overcome this trickery. Encoding web pages in JavaScript is a method that they use to hide their web site’s source code. To counteract this encoding, JavaScript decoders are available on the Internet that allows the source code to be viewed. Or alternatively, users of Microsoft’s Internet Explorer browser can make use of the “Microsoft Web Developer Accessories” add-on to view the plain HTML that the spammer’s JavaScript sends to the browser. The reason why spammers make such an effort to hide their code is that they know that anti-spammers are able to retrieve important information about their operations from this source code.

A large number of Internet users who had the misfortune of opening up “loaded” spam messages have experienced some of the nastier tricks that spammers use. These tricks include barraging the user with multiple pop-ups, switching the browser to full-screen mode and disabling the user’s ability to close the program, or re-opening a web page each time a user closes it. Spammers also use JavaScript to disable the right-click option to view the source code. A practical solution to these problems is to put the web browser in restricted mode, which will disable JavaScript, ActiveX, cookies and other tools that the spammers employ. The downside in using the restricted mode is that the user may not have the full experience of legitimate web sites that also use these features.

3.7 Spam Scams

In this section, the most prolific spam scams are examined. Suggestions are offered to assist e-mail users in avoiding these spam traps. Some of these scams could be considered harmless, but others have proven to be deadly.

3.7.1 Modem Hijacking

Modem Hijackings occur when unauthorised software uses a computer to dial up to a website that bills the phone owner with exorbitant call charges. This software can be installed unknowingly by opening up a spam attachment or as the result of a rogue popup advertisement. If fallen prey to such a scam, the call charges should be disputed with the phone company. Software is also available that prevents modem hijacking.

3.7.2 Identity Theft

In early 2003, users of America Online received e-mails requesting their passwords and other personal information. Once this type of information is compromised, a person’s stolen identity could be used to make credit-card purchases and commit other fraudulent dealings. Many users of other online providers have experienced the same spam scam. It should be noted that credit card firms and ISPs would never ask their customers to reveal passwords or any other personal information in such a manner. These e-mails should not be responded to, but reported to the relevant authorities.

3.7.3 Scam Websites

Another spam scam going around at the moment are e-mails that suggest that the recipient could find out what gossip was posted about them on a particular web site. There is a web site called Word-of-Mouth.org, registered in the Philippines, which invites users to write gossip about other people whose e-mail address they have to include. The company then contacts the owner of this e-mail address with instructions on how to view the gossip, for a fee. Deleting the e-mail is the solution for this scam.

3.7.4 Computer Viruses

Unfortunately, there are numerous e-mails rooming the Internet that contain viruses with varying degrees of destructive powers. E-mail from unknown users should not be opened. However, there are viruses that can steal e-mail addresses in personal address books stored on computers, so it may appear that the e-mail is coming from someone that is known to the recipient. This is the trademark of viruses known as worms. The only solution to this problem is to make sure that the virus definition files are up-to-date and to confirm any attachment from a sender before opening it up. If in any doubt, do not open attachments and especially do not open files that end with the following extensions: VBS, PIF, SCR and BAT.

3.7.5 Credit/Debt Repair

This spam scam advertises debt or credit card repair. For a fee, companies like ClearCredit.com and GetCleanCredit.com will offer to rectify any credit problems that a person may be experiencing. After providing their credit card numbers, the consumer finds that these companies charged their cards for unauthorised services, often ending up in more debt than before receiving the spam message. Deleting the e-mail is the easy solution for this type of scam.

3.7.6 The Nigerian (419) E-mail Scam

The Nigerian e-mail scam is probably the most pervasive e-mail scam in existence. Also known as the 419 scam after the section in the Nigerian criminal code covering such activity, there are many variations but the theme is the same: a wealthy Nigerian needs assistance moving money out of the country. Victims of this scam are informed that they will receive a certain percentage of a multi-million dollar fortune if they provide their bank account as a temporary holding account. Upon receiving the pertinent information, the scam artists proceed to steal whatever money is in the account. Recently, spam has hit the Internet involving cashier’s cheques. Auction sellers, for example those found on eBay, are paid with a cashier’s cheque that is worth more than the agreed purchase price. The sellers are then asked to cash the cheque and wire the difference to a bank account in Nigeria. Usually, money deposited using cashier cheques is available faster than funds deposited using personal cheques. The victim, after depositing the cashier cheque, feels secure enough to wire the money to the requested account. Later when the cashier’s cheque is found to be phoney, the bank makes the victim pay the full amount, as they are liable for the cheque. In the Czech Republic, a victim received this spam and followed through with it and lost his life’s savings. In retaliation, it is alleged that he shot and killed a Nigerian diplomat. According to the U.S. State Department, over 25 murders of Americans have been directly linked to the Nigerian scam. This is the unfortunate situation where spam has a deadly consequence.

3.8 Chapter Summary

In this chapter the focus was on the spammer. A spammer profile was introduced which included the motivation behind their decision to get involved in the spamming business. Tools of the spamming trade were discussed, as were the ancillary services that keep them in operation. The sections that dealt with spammers’ tactics, tricks and scams gave an insight on what the spammer has to do to stay one step ahead of the anti-spamming community.

Chapter 4

Anti-Spam Techniques

4.1 Overview

This section deals with the technical remedies that are available to combat spam. The tools used in fighting spam and an in-depth analysis of the current state of the art in anti-spam techniques is presented. The effectiveness of these solutions is also evaluated.

4.2 Strategies

The strategies used to block spam are diverse and includes many promising techniques.

4.2.1 Closing Open Relays

Open relays were very common in the early days of the Internet and they occur when a mail server processes a message for a sender or recipient who is not a local user. It acts as a third party to relay e-mail between other mail servers. The simple concept can be seen in the following schematic.

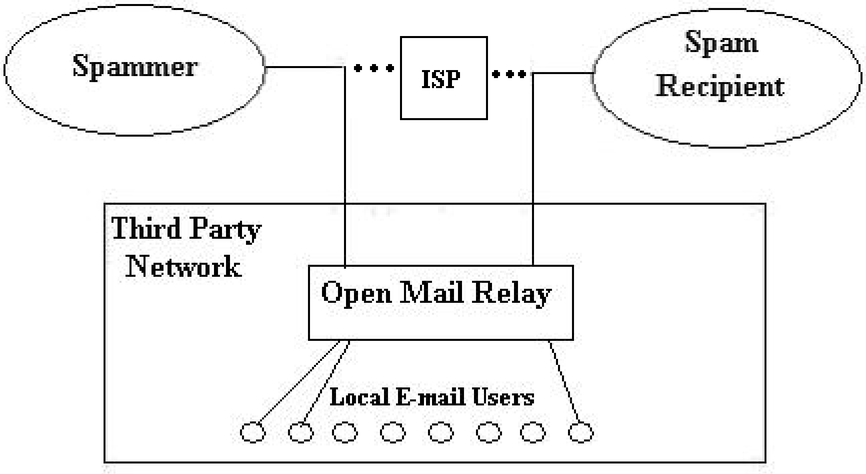

Figure 1: Open Relay E-mail System

The sender and recipient are outside the local domain (IP Network), but the spammer can still use the mail server to transfer the e-mail message. Today, there is absolutely no scenario that requires an open relay, but unfortunately there are many e-mail servers that act as open relays, and as such, provide a conduit for a spammer to inundate the Internet with spam. There are a number of reasons why a spammer would use an open relay. For example, there are dedicated spam operations flooding the Internet from static locations. It is quite easy for administrators to ban all connections from these locations. To avoid this type of blockade, spammers commandeer third-party e-mail servers to bypass the blocks that are in place. As administrators have not banned the open relay on these third-party networks, the spam will get through the blockade to its destination.

Another reason why spammers use open relays is to amplify the number of spam e-mails that they send. Many open relay mail servers are high-powered machines with fast Internet connections. Getting access to these servers allows spam to be pumped out at an extraordinary rate. Not only is the spammer getting access to the equipment to send out spam, but they are also using the victims' bandwidth. In no uncertain terms, this can only be classified as the theft of resources. The fact that a spammer can hide their identity behind an open relay is an added bonus. When the spammer relays spam through an open relay, the commandeered server will appear to be the source of this spam. This, coupled with forged e-mail headers, makes tracking the spammer very difficult. The consequences of an open relay being used by a spammer can be very expensive for the server's owner. The mail server may crash during a relaying episode, as a result of disk usage or processor overload. This would then prohibit legitimate e-mail activity, which is a serious consequence as it is an integral part of day-to-day operations for the typical modern business. After a mail relay hijacking, many unplanned person-hours are required to recover from the attack, which may involve restoring valid e-mail messages. There is also the cost of downtime and of course the tarnished image of being a source of spam. The domain that owns the open relay could be blacklisted as a result. Closing down open relays severely restricts the spammer by denying them a primary tool used in sending spam. RFC2505 describes in detail how to prevent unauthorized mail relaying [21].

4.2.2 Blacklists

A blacklist is a list of known IP (Internet Protocol) addresses that are used to send spam. Internet Service Providers and e-mail administrators can subscribe to these blacklist databases in order to filter out spam passing through their systems. A blacklist not only contains a database of spammer IP addresses, but also a database of known open relays, which facilitates the sending of spam. So much spam originates from open relays, that it was necessary to blacklist open relay servers.

ORDB.org is an organisation that stores the IP addresses of verified open SMTP (Simple Mail Transfer Protocol) relays. At the other extreme, the Internet Mail Relay Services Survey Project takes a more proactive approach and actively sweeps blocks of IP addresses worldwide looking for open mail relays. It stores the results in a database that can be used as a blacklist of open relays. While ORDB.org and similar groups focus on open relays, there are other Internet organisations that maintain lists of other sources of spam. The Spamhaus Project is an organisation that tracks down the Internet’s worst spammers and works with ISPs and law enforcement agencies to identify and shut down persistent culprits [22]. According to Spamhaus, 90% of all spam received by Internet users in North America and Europe originate from a group of approximately 180 individuals. The Register of Known Spam Operations (ROKSO) is a register of these known operations and is used to update blacklists. 50% of spam is sent from spammers directly, while the other half is sent (usually by the same spammers) with the help of open proxies and open relays. The Spamhaus Block List (SBL) is a real-time DNS-based list of IP addresses of verified spammers, spam gangs and spam services. It does not include open relays. ISPs and business networks alike use the SBL list. This list can be queried in real-time by mail systems to determine if the source of e-mail is a known spammer. As the list determines whether e-mail gets to the recipient, the integrity and accuracy of the list is crucial. All SBL entries are backed up with evidence that the IP address of the source e-mail is under the control of a spammer. The criteria to get listed in this blacklist is as follows:

- Spam Source. Static IP addresses where spam has originated.

- Spam Gangs. If registered in the ROKSO database, or listed as a known spammer and moved to a new IP subnet.
- Spam Services. Mail servers, web servers, DNS servers used in sending unsolicited bulk e-mail.
- Spam Support Services. Any support entity that is used in facilitating the sending of spam, such as bulletproof hosting, DNS resolution and related services.

Spamhaus asserts that as of February 2003, the SBL was protecting 110 million users worldwide. Blacklists have been available for years but according to a study by Network World [23], the MAPS RBL, which is one of best-known blacklists, only catches 24% of spam, with 34% false positives. False positives are legitimate e-mail messages that are classified as spam and such a high number is unacceptable to most companies. Another problem with blacklists is the possibility of legitimate organisations ending up on a list and suffering the consequences of having its e-mail rejected worldwide.

The U.S. Federal Trade Commission (FTC) is the federal agency in the United States who is primarily responsible for prosecuting spammers, and relies on e-mail as an important communication medium with the public. In 2002, the agency began using blacklists and the result was bounced e-mail from senders trying to complain about spam. If the sender’s ISP is on the blacklist used by the FTC, then all e-mail will be rejected from that provider’s domain. This problem raised a very interesting legal issue as it may have violated an Americans’ First Amendment right to petition the government. One of these senders, a technology policy analyst at a research institute had her e-mail to the director of the FTC’s bureau of consumer protection rejected. The irony in this case was that the sender was trying to send the director a proposal on how to limit spam.

It is worth mentioning that whitelists are also available, which is the opposite approach to blacklists. This is a list of e-mail addresses or domains from where the recipient organisation will accept e-mails, but these lists alone are hardly practical as a solution for large businesses. The other problem with whitelists is that it is not difficult to guess and spoof an “acceptable” domain name. It is interesting to note that the vast majority of IP addresses listed on blacklists are insecure servers in the United States, many of which are located in the country's universities.

4.2.3 Spam Filters

There are a several kinds of spam filters available today, each with debatable pros and cons. This section deals with these filters.

User-Defined filters are the filters that are included with many present-day e-mail clients. They work by forwarding e-mail to various mailboxes depending on the headers and content of the message. You could put e-mail from legitimate sources into a mailbox named after that source. Other filters could delete dubious e-mail. The end-user can be the person who sets up the filter by examining the spam that they receive and creating rules based on their spam. For example, rules could be created based on the appearance of "FF0000", which is HTML for bright red, which happens to be a good indicator of pornographic spam. It should be stressed that if the end-user is not in charge of the filtering process, then the entity in charge of delivering e-mail to the user's inbox can be legally responsible if objectionable material gets through.

Language filters are simple enough to understand – they filter out any e-mail that is not in your language of choice. Even though foreign language spam is not a major problem today, it is expected that it is only a matter of time that spammers, who are not particular in their methods, send bulk messages to e-mail lists in foreign countries. They send these messages irrespective of the ability of the recipient to be able to read them. Setting a filter to reject e-mail in other languages would eradicate this source of spam.

Header filters are somewhat more complicated than the filters discussed previously. Header filters examines the headers in the e-mails for forgery, which is a common trick used by spammers. Generally a user sees the sender, recipient and subject header fields on an e-mail message. It is easy for the user of most e-mail clients to view all of the e-mail headers, with for example, a "Show All Headers" button. There is a wealth of information in this expanded view. Deciphering e-mail headers is the first step in tracking down a spammer. Routing information is available in these headers, which tells the recipient the name and IP address of the different servers that the e-mail message passed through. It also shows what computer sent and received the mail and the exact date and time that it was sent. The most recent “received” line in the header is the server that delivered the e-mail to the user’s ISP. Since spammers do not want to be traced, they put bogus information in e-mail headers. Header filters detect forged headers in e-mail, which is almost a guarantee that the e-mail is spam. However, not all spam has spurious headers, so header filters should be used in conjunction with other filters.

Permission filters blocks all e-mail that does not come from a sanctioned source. Usually the first time e-mail is sent to a person using a permission filter, an automated response is sent back directing the sender to a web site where information is entered. When this is accepted, the sender’s e-mail address is sanctioned and all future e-mails will be accepted.

Content filters scan the contents of e-mail and uses a weighted mechanism to determine if the e-mail is spam or not. They are highly effective but can lead to the filtering of legitimate messages. Some experts believe it is possible to stop spam completely and that content-based filters are the solution. When you first look at a new e-mail message you know within a second or two, whether it is spam or not. All that is required now is the software that can make the same judgment. A Bayesian filter is a content-based filter that is proven to produce fewer false positives than blacklists and other filters. It works by assigning spam probabilities to individual words in an e-mail message. The filter can be taught to quarantine e-mail that use the spammers’ trick of jumbling words and using numbers instead of letters in order to avoid detection. There are many papers on text classification, which is a superset of spam filtering, but it was not until 1998 that Pantel and Lin [24] wrote the first paper on the Bayesian filter. Their filter caught 92% of spam, with 1.6% false positives, which was not sufficient to classify it as a successful filter.

Paul Graham achieved a major improvement with his Bayesian filter using statistical analysis [25]. With 0 false positives and missing less than 5 per 1000 spam e-mail (based on his own corpus), it is relevant to see how this powerful anti-spam technique was developed. One corpus of spam and one of non-spam e-mail is used. In Graham's early filter, he used 4000 messages of each type. The entire text of each message is then scanned, including the headers, embedded html and JavaScript. Alphanumeric characters, dollar signs, dashes and apostrophes were considered to be part of tokens, and everything else as a token separator. Then the number of times a token occurs in each corpus is counted, resulting in two large hash tables, one for each corpus, mapping tokens to the number of occurrences. A third hash table is then created mapping each token to the probability that an e-mail containing it as a spam message. Each corpus is considered to be a single long stream of text (in order to count occurrences), but the number of e-mails in each corpus is used as the divisor in calculating the spam probability, rather than the combined length. This was a deliberate bias in order to protect against false positives, which was a constant theme in Graham's work. He was well aware of the consequences of false positives being classified as spam. When new e-mail arrives, it is scanned into tokens and the fifteen tokens that are regarded as most interesting (i.e. how far their spam probability is from a neutral 0.5) is used to calculate the probability that the e-mail is highly likely to be spam. When a new word is found, there is a high probability that it is in a legitimate e-mail, as spam words tend to repeat.

According to Graham, the advantage of the statistical approach is the fact that a large number of spam e-mails do not have to be read, and additionally, it is known what is being measured. SpamAssassin™ (Section 4.2.9) is a feature recognising filter that assigns a spam score to e-mail. Graham believes that there are problems with the notion of a "score", because of the uncertainty of its meaning. The Bayesian approach assigns an actual probability, where specific rare spam words can even decrease the probability that e-mail is spam. The concept behind the Bayesian method is to filter each user's e-mail based on the spam and non-spam that they receive. A "delete as spam" button could be used to generate the spam corpus, and everything else would be applicable to the non-spam corpus. Many IT administrators believe that this is the main downside of Bayesian filters. The filters work for individuals and not groups, and also require constant maintenance to keep up with new spamming techniques. The other problem with word filters is that spammers can send e-mail that contains no words but consist of an image that loads from an external source. In order to defeat this type of spam, a filter would have to stop the HTML link from accessing the external image.

4.2.4 Machine Learning Techniques

The Artificial Intelligence (AI) community has taken a keen interest in the fight against spam. The Bayesian approach is only one approach to the problem and there are many papers that discuss other machine learning concepts. Support Vector Machines (SVM) is one possibility as a means of classifying e-mail as spam or non-spam [26]. The learning machine is trained based on sample e-mails that are either marked as spam or non-spam by a user. SVM are a collection of classification and regression algorithms where boosting using decision trees algorithms (such as C4.5) can be used to speed up the training and classification process. There are also research papers on ways to filter e-mail using k-nearest neighbour, centroid based approaches and other classification techniques. In fact, there has been so much AI work done on e-mail classification, that it could be the subject of a substantial thesis on the topic.

Artificial Intelligence solutions try to mimic the way our brains think when we receive spam. We can quickly tell with a quick scan if an e-mail message is spam or not. There are many variables that need to be examined in order to determine the legitimacy of an e-mail message. The problem with many filters at the moment is that just because an e-mail message contains phrases like “looks good”, “breast”, “toll-free number”, “feel satisfied”, it is a high probability that it would be flagged as a spam message. However, this could be an innocent e-mail discussing a chicken dinner recipe. Intelligent anti-spam solutions would look not just at the words, but other factors such as the time and date that the e-mail was sent, how often e-mail was received by the sending server, and other similar features. The recent focus on AI solutions to the spam problem is due to the increasing use of the self-learning filter.

4.2.5 Honeypot E-mail Addresses

Honeypot addresses are the most widespread proactive anti-spam technique in use today.

In general terms, a Honeypot or Tarpit is a system designed to look attractive to spammers and crackers. A mail server could be disguised as an Open Relay, but in fact it will not forward any spam. Honeypot Addresses are fake e-mail addresses designed as spider bait i.e. designed to get picked up by e-mail harvester programs. These fake addresses can then pollute the spammer's list database.

4.2.6 Deceiving E-mail Harvesters

Spammers use e-mail harvester programs as a common source of e-mail addresses. These harvesters troll the Internet searching for e-mail addresses on web sites. There are tools on the Internet that makes addresses visible to humans reading a web page but invisible to the web robots that is responsible for the address extraction. These "bots" (robots) specifically looks for the e-mail fingerprint associated with the HTML tag "<A HREF = mailto: ……>". These tags are very easy to find on a web page and they highlight the fact that an e-mail address exists next to the tag. There are JavaScript programs though, that can make these tags and associated e-mail addresses look like garbled text and therefore inaccessible by the web robots, but allows the web page and e-mail addresses to be viewed by a human. The only possible way for a harvester to retrieve JavaScript e-mail addresses is to include a JavaScript compiler with the harvesting software, which is a layer of complexity that most harvester developers do not want to deal with. However, there are rumours in some anti-spam newsgroup postings that new web robots are now being designed to defeat this defence mechanism.

Another trick that is used to fool the spammer's harvesting program is to use an HTML form instead of a "mailto" tag. Using a form, a CGI (Common Gateway Interface) script would send the contents of the form to an e-mail address, rather than publicly displaying that address. If it is a legal requirement that an e-mail address must appear on a web page, then a graphic image could be used that would display the address, but which a robot cannot retrieve. Some other methods to fool e-mail harvesters are discussed in the next section (Section 4.2.6.1).

4.2.6.1 Spider Traps

A spider trap is a web page that contains a large number of fake e-mail addresses with the intention of filling up spammer lists with worthless addresses. Some of these pages also contain URLs to other spider traps in order to propagate the uselessness of the harvested e-mail database. The spammer can never know to what extent their lists are polluted. Some of these spider traps are automatically generated using scripts, so when a spider accesses the page, it is consumed with an enormous amount of bogus data. A “poison” CGI script could be used to create a web page with bogus e-mail addresses that could include a link returning to the script page, which results in the spider getting stuck in an infinite loop harvesting unusable addresses.

Software has been developed to prevent spiders from accessing parts of web sites that are marked out-of-bounds. Robotcop is an example of an open source module for web servers and is used by many web site administrators. There are non-threatening spiders; for example a web crawling program that indexes web sites for search engines or spiders (web crawlers) that searches web sites for images and music files. These crawlers generally should have access to web pages, but it is the spider that harvests e-mail addresses that is targeted by this anti-harvesting software. Robotcop works by requesting all spiders to read a robots.txt file and expects these spiders to abide by the rules set out in the text file. If a spider breaks any of the rules then any further requests by the spider are intercepted and blocked. “Trap” directories that are marked off limits can be created to test the spiders, and if the spider falls into the trap then a number of countermeasures can be initiated. A misbehaving spider more than likely represents an e-mail harvester, so it could be redirected to a directory with bogus e-mail addresses in order to pollute the harvester’s database. More often than not though, the mischievous spiders are just intercepted and blocked with its IP address logged in order to reject any future requests.

4.2.7 Postage Stamps and Digital Signatures

Some anti-spam systems would like to use the concept of postage in an electronic, peer-to-peer model of the postage stamp. The idea behind this system is based on the electronic postage stamp, which is required for all e-mail coming into a user’s inbox. The currency of the postage stamp is processing power (CPU cycles), which is used to solve a puzzle and whose solution becomes the postage stamp. A single e-mail postage stamp would not be taxing on a user's system but if you were a spammer and would like to send spam to hundreds of thousands of individual e-mail addresses, then there is considerable overhead involved. Even though this anti-spam system would work well for single users, it would be expensive for online providers and business, which have to perform the puzzle solutions for their users. This is where the public keys concept enters into the equation. The provider or enterprise could collect the public keys of all messages passing through their mail systems in order to form a relationship, which in turn are used to validate future messages.

4.2.8 Mobile Agents

An innovative approach to fighting spam is with the use of mobile agents. Li Cheng and Wang Weinong [27] of Shanghai University introduced this concept at the 2002 Symposium on Applications and the Internet. In an SMTP system, when e-mail is communicated between clients, a transmission channel is setup between the target and sending MTA (Mail Transfer Agent). After this channel is established, the SMTP sender transmits a MAIL command. This is where the mobile agent enters into the system. The SMTP receiver responds with an AGENT reply instead of an OK reply and dispatches a mobile agent to the sender. The agent will then proceed to audit and filter the e-mail message, rejecting spam and forwarding acceptable mail to the recipient. The obvious upside to this system is that the filter processing is not carried out at the receiving end. If the e-mail is spam and destined to many thousands of recipients, then the spammer's mail server will have to bear the brunt of the processing used by these mobile agents. If the e-mail is legitimate, then the owners of the sending mail system should not be overly concerned by the taxing of their system with these mobile agents, as they know that the processing power used will be reciprocated by systems sending e-mail to their mail servers. This concept could be further enhanced to check e-mail for viruses, where a suspect attachment would never reach the interior of a network and allow itself to propagate. The major downside with the mobile agent approach is the main disadvantage associated with mobile agents in general, which is security and how to protect both the host and agent.

4.2.9 Other Anti-Spam Methodologies

Technical analysis is an anti-spam approach where a program performs an analysis of the incoming message and makes a determination on whether the message is spam or not based on specific criteria. For example, the program could run a forward DNS test to check the authenticity of the sending domain. It could also send a reply to the e-mail and if it does not receive a 'user-unknown' reply, which usually occurs within a very brief period of time, then that e-mail could be considered genuine. Unlike other common methods of spam filtering, technical analysis will not block e-mail from legitimate sources. Once e-mail has been successfully received, the sender can be assigned a trusted status, which will ensure that all future e-mails from this individual will be accepted. In this way, the performance overhead is reduced after receipt of the first e-mail.

Combination approaches are the most successful filters on the market at the moment. Tagged Message Delivery Agent (TMDA) is an OSI software application that combines a whitelist, a blacklist, and a cryptographically enhanced confirmation system. TMDA's whitelist strategy relies on keeping a list of trusted e-mail addresses, which can send e-mail unhindered. E-mail from unknown senders is held in a queue until a response is received to a one-time confirmation request. Once a response is received, the original message is considered legitimate and delivered. TMDA then adds this new e-mail address to the whitelist, where confirmations are not required.

Another combination approach is a system that uses a whitelist, an RBL (Real-time blacklist), a filter based on long distinctive phrases or sentences found in spam messages, and blocking all e-mail from countries that are notorious sources of spam. These countries are generally regarded to be China, Taiwan, South Korea, Brazil and Argentina. Even though you are blocking e-mail from whole countries, it is quite easy to add specific addresses from these countries to the whitelist for successful receipt of messages. Content filtering based on simple words or short phrases are considered by some to be error prone and a source of false positive e-mails. This is the reason why content filtering used by these systems are based on long, distinct phrases found in many spam e-mails.

SpamAssassin is a popular and sophisticated anti-spam program that checks e-mail for certain spam criteria. It uses text analysis and several real-time blacklists. Each criterion has a weight associated with it. If the combined weight of all the matching criteria exceeds a certain threshold, then the e-mail is considered spam and dealt with accordingly. Using its rule base, it utilises a wide range of heuristic tests on the e-mail headers and body text. One criterion is whether the “To” and “From” headers are the same. As this is common in spam, it is weighed high. Large fonts and multiple colours are also weighed. Other criteria checks whether the whole e-mail matches (or closely resembles) spam on the Razor database, which is a database of thousands of spam e-mails. SpamAssassin is one of the most widely used content-based anti-spam filtering packages.

Some online providers offer their users the opportunity of using a challenge-response option with their e-mail client. This technology obliges the sender of e-mail to verify their authenticity before the message is accepted. In May 2003, Earthlink, the third largest online provider in the United States, introduced the challenge-response option to its 5 million subscribers. However, many leading anti-spam experts fear it could backfire and render e-mail useless if this approach is widely embraced. Legitimate e-mail users now face an extra layer of message exchanges in order for their e-mail to be delivered to Earthlink users. Earthlink currently blocks 80% of spam, but it has increased by a factor of 6 in the last 18 months prior to May 2003. The fear of an increase in the number of false positive e-mails if a more selective filter was used caused the company to decide on the challenge-response option. America Online, the leading online provider in the United States, blocks 80% of all incoming traffic, or an astonishing 2 billion messages a day (May 2003). The burden on America Online’s system to send out over 2 billion challenges every day would be too demanding. With many USENET newsgroups having thousands of subscribers, this anti-spam technique could quickly become very user-unfriendly.

4.2.10 Success of Anti-Spam Techniques

There is constant debate on the success of these anti-spam techniques. A lot of research has gone into fighting spam, yet we still receive enormous quantities of it in our inboxes. For most users, even one spam message per day is very annoying. Spam works because some people respond to it. If there was a way to restrict the sales pitch of these spam e-mails, or make it cost prohibitive for spammers to send e-mail, then it is indeed possible to put them out of business. The spam response rate is very low (15 per million e-mails). If it was possible to reduce the response rate even further then the spammer would have to find another line of work. The solution lies in the constant effort to hinder the spammer with new filters and anti-spam techniques. In the future there will a time when there is a continuous flow of diverse anti-spam software solutions coming to the market on a regular basis, and that will be the time when spam loses its stranglehold on our inboxes.

4.3 Spam Revenge

Some anti-spam enthusiasts believe that spammers should be spammed back. There was a concerted effort to put the e-mail addresses of known spammers on mail lists and on web sites as bait for e-mail harvesters. This would lead to the paradox of spammers spamming other spammers. Other more vengeful victims of spam have put spammers on postal mail lists and subscribed them to multiple magazine subscriptions. One of the most ingenious ways to make the companies that manufacture and distribute spammer software pay a financial price for doing business with spammers can be found at http://www.spambattle.com. SpamBattle.com declares that: “it is the ultimate revenge and infuriates the spam software companies to no end”. This is how it works. Companies that create software helpful to spammers usually advertise on Pay-Per-Click search engines. Every time an Internet user clicks on one of these links, the spamware company is charged. A cost is assessed for a unique click in a 24-hour period, which means the advertiser will only be charged once per user during that time. Consequently, continuously clicking on a link will not cause the advertiser to be charged for each click. SpamBattle.com urges as many people as possible to search for topics like “bulk e-mail”, “email harvesters” etc. and then click on the results. The aim is to shift the cost burden of spam back to the companies that create the software.

4.4 How to Avoid Spam

Unsolicited bulk e-mail is the number one complaint on the Internet today. In order to receive spam, the spammer needs to know your e-mail address, so the best way to avoid spam is to keep your address off spam lists. There are a number of ways to achieve this.

4.4.1 Use E-mail Obfuscation

The spammers' number one source for e-mail addresses is from web sites. Web robots (or spiders) crawl websites for e-mail address and if your address is found, it will be added to a spam list and shared among spammers. The solution to this problem is not to put your e-mail address on a web site, but if it necessary for you to do this, then obfuscate the address using a number of known methods. An e-mail encoder could be used, or simply replacing the "@" symbol with the word "at" would also be sufficient. Another way to keep off spam lists is to avoid posting to USENET newsgroups. This is probably the second most successful method for spammers in obtaining e-mail address as your e-mail address is out there for everyone to see.

4.4.2 Pick an Unusual E-mail Address

In the analysis section of this thesis, it was found that by picking an unusual username also helped in spam avoidance. For example, john2000007f@hotmail.com is better to use than johnsmith@hotmail.com, as the former address is less susceptible to a dictionary attack. Expanding on this concept, there is also the notion of "address munging", which is a technique where changes are made to the domain portion of your e-mail address in order to trick e-mail harvesters. Instead of using johnsmith@hotmail.com, the address johnsmith@nospam.hotmail.com could be used. Mangling the username portion of your address will still allow spam to be sent to your ISP. However by mangling the domain portion will cause the spam message to be undeliverable (as the domain will not exist).

4.4.3 Do not List your E-mail Address

Listing an e-mail address in an online directory can also lead to the receipt of spam e-mail. When setting up an online account, in some instances there is an opportunity to check a box that will list you in the company directory. This box should be unchecked. One of the e-mail harvesters that I used in this research actually used the Yahoo online directory to retrieve e-mail addresses for spamming purposes.

4.4.4 Do not Post using your Private E-mail Address

Posting to public mailings lists could also be hazardous to your online privacy. Posting to a public forum verifies the authenticity of an e-mail address, even if users use throwaway addresses for this purpose. This leads to another solution to avoid spam – use disposable e-mail addresses.

4.4.5 Disposable Addresses

Disposable or throwaway e-mail addresses are practical for users who want to make use of public forums and newsgroups. The fear of receiving spam should not lead down the path of the Internet hermit. Many online users practice this spam avoidance technique by creating e-mail addresses specifically for these areas of the Internet. An e-mail address should be setup as "private" and provided to only family and friends, and e-mail addresses should be setup as "public" when you do not care if the account receives spam. When it comes to the juncture where the public address is deemed useless because of the amount of spam it receives, it can be discarded and replaced by a new disposable account.

4.4.6 Use Provided Filtering Systems

For the home user, most online providers provide filtering options with their accounts, which is a service that should be used in order to avoid spam. Hotmail is a provider that gives the user an opportunity to set the level of filtering available with their accounts. The default level catches obvious spam e-mail. The enhanced level is where most of the spam sent to the e-mail account is caught, and finally the exclusive level is where the user will only receive e-mail from a specified safe list. Hotmail also allows the user to decide what happens to the junk e-mail. The spam can be deleted immediately without being delivered to the junk folder, or it can be sent to the junk folder so that the user can review it, and it is then deleted after a certain period of time, usually 10 to 30 days. For the business user, there are two main tools used to avoid spam: anti-spam services and anti-spam software. Some businesses outsource their spam fight to a service that monitors incoming mail. They use a variety of methodologies to categorise the e-mail e.g. whitelists, spam matching, filtering etc. Anti-spam software is used by organisations that do not want to outsource this function. The software can either run on the mail server or on the local desktop, and it is the method of choice for privacy-concerned companies.

4.4.7 Do not Reply to Spam and do not use the Unsubscribe Link

Upon receiving spam, many recipients would like to send e-mail back to the sender with suggestions on what to do with the spam or with a request to unsubscribe from any future mailings. This is a very natural response as the user feels the spam intrusion as a violation and may consider it as an invasion of their privacy. My research concluded that when you do respond to spam, there is usually no one “at the other end”, that is to say, you will probably receive back an “unknown user” message. Spam recipients should never respond to spam, just in case the spammer’s e-mail address is valid. If you do respond, it is only confirmation that your e-mail address is genuine and you will probably end up on another spammer list. Continuing along this line of advice, it is highly recommended that the unsubscribe option in spam e-mail should not be used, as once again it only verifies the validity of your e-mail address. When this option was exercised by bait e-mail addresses created for this research, these “Opt-Out” links usually went nowhere or the link was redirected to a commercial site that was the source of the spam, and other opt-out links just opened up pop-up advertisements. Currently there are laws being proposed that would require unsubscribe links to be genuine and would assess a financial penalty if opt-out options were not respected. Until these laws are in place, it is imprudent to use these links.

4.4.8 Turn off Automatic Downloading

If your e-mail client has a feature that automatically downloads files, then it should be turned off for two main reasons. The first reason is the possibility of a virus or other harmful program being automatically downloaded to your PC. The second reason is that by automatically downloading a file, it registers the fact that your e-mail address is alive and genuine, which is comparable to a confirmation receipt for a spammer. It may be convenient to automatically download files e.g. HTML Graphics, but it is known that users who turn off this option receive less spam.

4.4.9 Check the Privacy Policy when Ordering Online

When purchasing merchandise online or when signing up for online subscriptions, read the privacy statement carefully to see what the seller proposes to do with your personal information. They may want to put you on a “future mailings” list or they may want to share your information with other partners. If possible, it is advisable to uncheck the option for future mailing. Even if a company respects a buyer’s decision not to share your information, there has been at least one instance where a bankrupt company has sold off their e-mail lists to the highest bidder. Lists of people who have purchased online in the past is highly sought after and regarded as a valuable asset.

4.5 Chapter Summary

Anti-spam techniques were presented in this chapter, which provided comprehensive explanations on how to counteract the methods used by spammers. There were stories on how some victims of spam sought retribution and finally, some advice was given on how to avoid spam and suggestions were made on how to stay off spam mailings lists. The irony of developing new successful anti-spam methods is that it in turn leads to more sophisticated spamming techniques. The positive spin with this fact is that spam is the only kind of e-mail that uses these types of tricks. If e-mail is masked in some way, then there is a strong possibility that it is spam, and the more complicated the masking technique, the greater the possibility that the message is spam. This information alone can be used to filter e-mail, without going through the time-consuming process of analysing the contents of the message.

Chapter 5

Implementation

5.1 Overview

The implementation in this research involved the creation of 59 unique e-mail accounts. These accounts are categorised as honeypot e-mail addresses, as their sole purpose was to attract spam e-mail.

5.2 Account Setup

Most of the accounts in this research were baited in a manner that would replicate normal e-mail use. In most cases I did not seek out spam, as it would provide bias in the spam corpus e.g. accounts were not setup on adult web sites due to the fact that the spam received from these sources would be biased towards pornographic spam messages. The accounts were created with obvious gender names e.g. John, Lisa, Michelle etc. The purpose of gender distinction was to find out if a spammer could target e-mail to a particular sex e.g. does a spammer care if e-mail created specifically for females is sent to male-sounding e-mail addresses. Also, the ages, occupation, and location were clearly defined in the profile section for each of these accounts, in order to establish if a spammer uses profile information for e-mail targeting. During the account setup, the checkbox was selected to list each account in the company's e-mail directory.

5.2.1 Honeypot Addresses

The full list of honeypot addresses and the baiting methods that were used on each account are listed in chronological order (by create date) in Table 1 (Section 6.2).

5.2.2 Baiting Methods

One of the most popular ways for spammers to obtain e-mail addresses is to use e-mail harvesters, which scour the web extracting e-mail addresses from web sites. Four e-mail addresses were put on a web page that allowed such indexing. Two other e-mail addresses were also included on the web page but were obfuscated using strategies that will be discussed later. The other most popular method used by spammers to obtain e-mail addresses is from USENET newsgroups. Eight e-mail addresses were listed in newsgroups, with each honeypot account joining at least two newsgroups. Some of these newsgroups were investor-oriented groups; the goal was to find out if spammers would target these e-mail addresses with spam specific to finance and investing.

Seven e-mail addresses responded to popup requests on web sites that I visited. Seven e-mail addresses were used to respond to recent spam. Four other e-mail addresses were used to send e-mail to known spammers. Thirteen e-mail addresses were used at various web site registrations. Three e-mail addresses were added to mail lists on web sites. Two e-mail addresses requested removal from a list, even though they were not originally listed. One e-mail address was used to register at a gambling web site, but specifically requested no mailings. One e-mail address was used to register with a chat service and another four addresses were registered with IwantSpam.com. Finally, three e-mail addresses were not baited in an attempt to find out if there were methodologies used by the spammers to collect addresses that I was not aware of. These three addresses were deemed less susceptible to a dictionary attach as the usernames were quite long and alphanumeric in nature e.g. john20000007f@hotmail.com.

Going into greater detail with the baiting methods, the IwantSpam.com bait addresses were easily configured. The requirement was to insert the e-mail address into a text box on a web site, checking the box “I want Spam”, and hitting the subscribe button. IwantSpam is a service that signs up the e-mail address to a variety of e-mail newsletters from various categories. Spamware.org is a web site where people can buy spam software, and is the organisation behind this free service. They issue a warning that if you submit your e-mail address then you will receive bulk e-mail. One has to wonder though about the process to verify the authenticity of the subscribed e-mail address as none of the bait addresses received a confirmation reply from IwantSpam.com. This would mean that an e-mail address could easily be entered for spam without the owner’s consent. I further realised that spammers themselves could use this service to enter the e-mail addresses that appear on their lists and claim that the recipients opted-in to receive their spam in this manner.

Newsgroups were deliberately selected to test if spammers were profiling e-mail accounts to be targeted with topical spam. These newsgroups were:
thestarwarsfanclub, the_star_wars_fan_club, fight-spam-by-spam, loversoftitanic, thegodfatherclub, theclubofjamesbond, dawnofthedeadclub, myhappyonline, canadianretirement, Russian_Adoption, GAARPadopt, askadviceman29, highyieldtradingprograms, ultimatetradersmomoclub, WorldGreatestTrader_com, aaahotstocks, sales-marketing-tips, and canslimstockpicks.
The wide range of newsgroups had various topics that included adoption, stock trading, retirement issues and newsgroups that were rated very popular by Yahoo. The goal was to see if a spammer would send e-mail based on the newsgroup topic e.g. financial-related spam to the financial-related newsgroup member. Each newsgroup bait address was configured to allow the profile to be viewed and at least one posting was made by each of the newsgroup accounts in order for the e-mail address to be publicly displayed. One bait address was registered with the notorious spam host newsgroup news.astraweb.com.

Two bait addresses requested removal from a list, even though the addresses had not received e-mail from the list host. The test in this case was to see if the list owner would use the fact that these were valid e-mail addresses and add them to a spam list. Some e-mail addresses were used to register at web sites that offered free offerings (e.g. online classes). In some instances, the checkbox to receive further mailings was unchecked in order to see if the host honoured the request. In other instances, the checkbox was checked to receive further mailings, with the objective to request removal from these lists at a future date and see if the request was honoured. One bait address registered with

Yahoo Chat and participated in a number of online chats in order to make the username and therefore the e-mail address publicly visible. There are numerous instances of chat room users receiving their first spam messages with minutes of signing into a chat room for the first time. E-mail accounts were also used to fill in popup requests i.e. windows that requested an e-mail address for further information or for a free prize, mainly in sites that are notorious for barraging users with popups.

E-mailing known spammers was another common baiting practice used in this research. This was done by either replying to existing spam that was received by other e-mail accounts, or by using lists of known spammers that are available on the Internet and sending simple one-line e-mail to these accounts. E-mail was sent to 155 spammers using these two methods.

5.3 E-mail Address Obfuscation

There are a number of ways to obfuscate or hide an e-mail address that appears on a web page. The easiest method is to simply use ‘at’ instead of using the ‘@’ symbol in the e-mail address. For example, on the bait web page in this research, maryjanedilbert at hotmail.com was used instead of maryjanedilbert@hotmail.com. It is easy for a human reader to know how to send an e-mail to the obfuscated address i.e. replace the ‘at’ with a ‘@’. However, when a web spider (robot) crawls the web page and finds maryjanedilbert at hotmail.com, it treats it as ordinary text and does not harvest it as a valid e-mail address. Since this is a popular method to hide e-mail addresses at the moment, it remains to be seen if the spammer has figured out a way to bypass this simple trick. It would be programmatically easy enough to harvest this address using a number of rules in a harvester, but the same rules could also lead to false hits. The only downside with this method of obfuscation is that it is not very user-friendly. Generally, a user would prefer to see the full e-mail address on the web site e.g. maryjanedilbert@hotmail.com and either just click on the underlined address to open up the native e-mail program and send an e-mail, or selecting and copying the underlined text and pasting it into an e-mail program. Instead the user would have to either retype the whole address, or just copy and paste the address, replacing the ‘at’ with the ‘@’ symbol, and deleting the spaces before and after the ‘at’. Another way to hide an e-mail address is to use an e-mail encoder. Encoding the full e-mail address through the use of character entities, which transforms the ASCII e-mail address into its equivalent decimal entity, obscures the address. For example, the HTML equivalent of john20000008@eudoramail.com is:

&#106;&#111;&#104;&#110;&#050;&#048;&#048;&#048;&#048;&#048;&#048;&#056;&#064;&#101;&#117;&#100;&#111;&#114;&#097;&#109;&#097;&#105;&#108;&#046;&#099;&#111;&#109;

There are numerous e-mail encoders on the Internet, where all you have to do is insert your e-mail address into a text box and the corresponding decimal encoded address is returned to you. You can then put this encoded address on a web page, and if an e-mail harvester indexes this page - the e-mail address will be bypassed. Taking a cynical view, if I were a spammer, then this would be an ideal way to collect valid e-mail addresses i.e. offer a free encoding service and just collect the e-mail addresses of users who use the service. However, I am not aware of spammers using this method (yet).

5.4 Account Monitoring

The 59 e-mail accounts were checked on a weekly basis to observe spam activity. The total number of e-mails for each account was counted and recorded. Even though both spam and legitimate e-mail were originally included in this count, it would still give an accurate depiction of incoming spam, as most of the e-mail in the inbox was spam related or e-mail from a subscribed service. The other non-spam sources in these accounts included e-mail from the mail provider and delivery notifications.

5.5 Chapter Summary

In this chapter, we looked at how honeypot e-mail addresses were setup and positioned to receive spam e-mail. A number of different methods were used to attract spam, most notably newsgroups, website registration, and listing the addresses on web sites. E-mailing known spammers, filling in popup ads, and responding to spam were some of the other methods. The process of setting up the 59 e-mail accounts was not very complicated but it was tedious and time-consuming. I initially considered developing a program to create the accounts but found that most of the providers issued a challenge during the registration process that was designed specifically to prevent this. For example, there would be a faded box with some random numbers and letters, and the user would be required to enter these characters into another text box. This feature was probably designed to defeat optical recognition systems that could be used in an automated registration program.

59 e-mail accounts were created because I wanted to include various profile characteristics such as gender, age, income, job description, geographical location, educational level (high school up to graduate school) etc. The painstaking effort to keep the profile in these accounts different was my attempt to monitor the types of spam that these users received and compare them to the profile. I wanted to find out if the spammer would target spam to the bait addresses based on the readily available profile information, which was highly unlikely for most spammers. I was also curious to find out if spammers had profile extractor programs to extract profile information in an automated fashion.

Chapter 6

Evaluation

6.1 Overview

There are numerous spam archives available on the Internet where researchers can download spam e-mail samples. Some of this spam goes back five years, so the natural evolution of spam messages can be viewed quite easily. The objective of the archive obtained in this research was to examine the current state of spam messages and to find out how spammer methodologies have changed in line with better filtering techniques.

6.2 Analysis of Honeypot E-mail

Not all of the e-mail that the honeypot addresses received can be classified as spam. Many e-mail messages were newsletters; other messages were from subscription-based lists. Even though these messages were spam-like in nature and had all hallmarks of spam, it would be inaccurate to label them as such for this research. As these messages were somewhat solicited, they therefore do not fall under the common description of spam i.e. unsolicited bulk e-mail. Other messages in the corpus were automatic notifications of unsuccessful attempts to deliver e-mail, for example when the bait addresses were used to send e-mail to known spammers or when these e-mail addresses replied to spam messages. Table 1 represents the total number of e-mails that each honeypot account received, the total number of spam messages for each account and the category of these spam messages. From a total number of 896 e-mails collected in a three-month period, 167 e-mails were deemed spam messages, which is 18.64% of the complete corpus. This is generally regarded as being in the low end of the spam scale, but proves the point that the amount of spam that an e-mail address receives depends on the usage (i.e. how a user propagates the address) and to a lesser degree, luck (i.e. unfortunate enough to be harvested).

E-mail Address	Baiting Method	Total Spam	Total E-mail	Spam Breakdown
mary2000003@yahoo.com	Newsgroup	5	17	N, 2RL, 2AU
mary2000004@yahoo.com	Newsgroup	3	23	N, RL, G
mary2000005@yahoo.com	Newsgroup	2	14	P, N
mary2000006@yahoo.com	Newsgroup	1	12	N
mary2000007@yahoo.com	Gambling.com Registration	3	11	P, N, AU
john20000003@yahoo.com	Popup Response	1	14	N
john20000004@yahoo.com	Newsgroup	1	13	N
john20000005@yahoo.com	Newsgroup	18	36	14F, 3RL, N
john20000006@yahoo.com	Newsgroup	1	13	N
john20000007@yahoo.com	Newsgroup	2	15	N, F
mary2000003@hotmail.com	Popup Response	14	171	2N, 8BE, 2WL, V, F
mary2000004@hotmail.com	Website Mail-list	16	23	2N,8BE,F,2P,2V,WL
mary2000005@hotmail.com	Website Mail-list	11	29	5BE, 2N, WL, 3P
mary2000006@hotmail.com	Popup Response	15	23	2N, 7BE, 2WL, F, 3V
mary2000007@hotmail.com	Responded to Spam	13	53	5V, 7BE, P
John20000003@hotmail.com	Website Mail-list	7	22	2N, 5X
John20000004@hotmail.com	Website Registration	5	18	2N, 3X
John20000005@hotmail.com	Website Registration	4	24	2N, 2X
John20000006@hotmail.com	Requested list removal	3	12	3X
John20000007@hotmail.com	Signed up with Iwantspam.com	21	29	7BE,2N,4X,2WL,4V,2P
John20000007f@hotmail.com	Not baited	0	7	N/A
mary2000003@campus.ie	Responded to Spam	0	13	N/A
mary2000004@campus.ie	Popup Response	0	4	N/A
mary2000005@campus.ie	Website Registration	0	4	N/A
mary2000006@campus.ie	Website Registration	0	7	N/A
mary2000007@campus.ie	Requested list removal	0	4	N/A
john20000003@campus.ie	Popup Response	0	5	N/A
john20000004@campus.ie	Website Registration	0	4	N/A
john20000005@campus.ie	Signed up with Iwantspam.com	0	5	N/A
john20000006@campus.ie	Website Registration	0	15	N/A
john20000007@campus.ie	Sent e-mail to known spammers	0	13	N/A
michelle2000003@mail.com	Sent e-mail to known spammers	0	20	N/A
michelle2000004@mail.com	Website Registration	0	5	N/A
michelle2000005@mail.com	Website Registration	0	29	N/A

E-mail Address	Baiting Method	Total Spam	Total E-mail	Spam Breakdown
michelle2000006@mail.com	Responded to Spam	0	6	N/A
michelle2000007@mail.com	Popup Response	0	31	N/A
john20000003@mail.com	Sent e-mail to known spammers	0	18	N/A
john20000004@mail.com	Website Registration	0	8	N/A
john20000005@mail.com	Signed up with Iwantspam.com	0	5	N/A
john20000006@mail.com	Not baited	0	5	N/A
john20000007@mail.com	Popup Response	0	32	N/A
lisa2000003@eudoramail.com	Sent e-mail to known spammers	2	7	2N
lisa2000004@eudoramail.com	Responded to Spam	2	4	2N
lisa2000005@eudoramail.com	Website Registration	2	6	2N
lisa2000006@eudoramail.com	Responded to Spam	2	4	2N
lisa2000007@eudoramail.com	Website Registration	2	19	2N
john20000003@eudoramail.com	Website Registration	2	5	2N
john20000004@eudoramail.com	Responded to Spam	2	4	2N
john20000005@eudoramail.com	Responded to Spam	2	5	2N
john20000006@eudoramail.com	Website Registration	3	13	2N, G
john20000007@eudoramail.com	Signed up with Iwantspam.com	2	5	2N
john20000008@eudoramail.com	Not baited	2	3	2N
johnpatrickdilbert@hotmail.com	Indexed	0	4	N/A
johnpatrickdilbert@campus.ie	Indexed	0	1	N/A
johnpatrickdilbert@mail.com	Indexed	0	3	N/A
johnpatrickdilbert@eudoramail.com	Indexed	0	1	N/A
maryjanedilbert@hotmail.com	Obscured Indexed	0	5	N/A
johnjosephdilbert@eudoramail.com	Obscured Indexed	0	1	N/A
John20000009@yahoo.com	Chat Room	0	1	N/A
	Totals	**167**	**896**	

Table 1: E-mail Corpus Breakdown and Categorisation

The spam breakdown abbreviations are as follows:

N = Nigerian, RL = Request to join List, AU = e-mail After Unsubscribing, G = Gambling, P = Product offer, F = Financial, BE = Body Enhancement, WL = Weight Loss, V = Viagra, X = adult entertainment, N/A = Not Applicable (No spam received).

6.3 Spam Corpus

This section deals with the e-mail messages that were received by the honeypot accounts and labelled as spam. In some instances, filters did not catch these messages as spam and could only be labelled as such by a visual inspection. Yahoo and Hotmail e-mail accounts received e-mail messages that should have been categorised as spam, but they were not caught by the ISPs' filters. Upon forwarding these messages to an account that used SpamAssassin, they were immediately caught as spam messages. SpamAssassin flagged them as spam based on the content of the messages only. Due to the fact that they were forwarded from a legitimate ISP, other information such as the sender and date headers became valid. So it appears that the type of filter that your ISP uses has a major impact on the decision on whether or not an e-mail message is categorised as spam. Some of the reasons why 28 honeypot accounts did not receive spam are:

- Obfuscation. The process of hiding addresses from web robots and harvesters worked.
- 10 of the accounts were registered with Campus.ie. It may be the case that this domain is too restrictive as it caters to students in Ireland.
- Activation period. All of the accounts were active for only 3 months. It is a strong possibility that given a longer time span, more accounts would have acquired spam.
- The majority of these accounts were deliberately setup to replicate normal use and did not actively seek spam. Among these, 3 accounts were not baited at all.
- There were 11 accounts baited by responding to spam or sending e-mail to known spammers. The spammer's address in each case was bogus.

6.3.1 Spam Diversity

The diversity in the spam received by the different honeypot addresses depended on the spammer who harvested the addresses. E-mail accounts on the same spam list will receive identical spam, but the spammer will use different headers in an attempt to bypass the filters. All of the spam messages received by the honeypot addresses could be classified into 10 categories. Even though, some of the categories had a broad scope (e.g. Product Offer), it was still surprising that the spam diversity that I expected did not exist.

The most prolific spam received was a Nigerian scam e-mail that crossed multiple ISPs and baiting methods. A noticeable finding in this research was the increase in the amount of the Nigerian spam and its different variations. 31 accounts received spam and only two of these accounts did not receive a Nigerian e-mail. As mentioned earlier in the spam profile section (Section 2.2.3), 5% of all spam is classified under the Nigerian category. However, based on the amount of Nigerian e-mails received by the honeypot accounts, the percentage has certainly increased. Other common spam e-mails received included Viagra, Weight Loss, and Body Enhancements. Some of these e-mails were farcical in nature, for example one message proposed to sell a product to lose weight while in the shower.

6.3.1.1 Reasons why Users Receive Different Spam

The question can be asked: What determines the kind of spam that a person receives? Based on this research, the type of spam that you receive is also based on the spammer who harvests your e-mail address. Some spammers do not want the headaches associated with pornographic spam and will not send them in any circumstance. On the other hand, if a spammer harvests your e-mail address and does not care about the content of the e-mail, then adult-related messages will eventually be delivered to your address. Unfortunately, due to the fact that spam lists are shared among spammers, it is only a matter of time that a harvested e-mail address receives the whole gamut of topics generally associated with spam messages.

It became apparent that a group of Hotmail honeypot accounts in this research were receiving spam e-mail from the same source. The spammer sent the messages in batches of two or three to the same accounts over the course of the research. The subject, date/time and sender headers were always different even though the content of the e-mails were identical. Due to the fact that these particular accounts were baited using different methods, it was obvious that the accounts were compromised by a dictionary attack. Another group of Hotmail accounts that had similar usernames also fell victim to a separate dictionary attack, as they were the only accounts in the research to receive identical pornographic spam messages even though they were also baited differently.

6.3.2 Spam Samples

In order to get a feel for the current breed of spam, here is a sampling of the spam messages that were sent to the honeypot addresses:

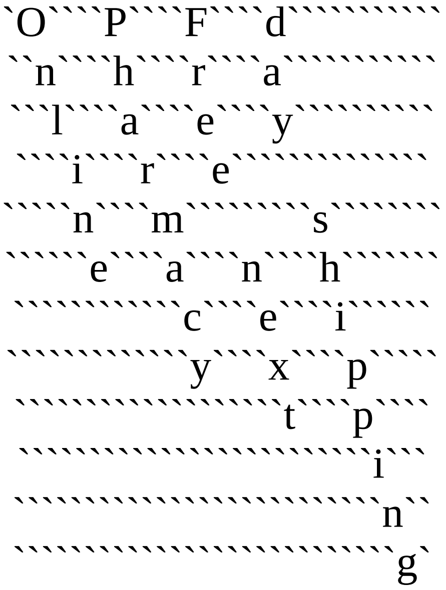

~~~
`O````P````F````d```````````
``n````h````r````a```````````
```l````a````e````y``````````
````i````r````e``````````````
`````n````m``````````s````````
``````e````a````n````h```````
````````````c````e````i``````
`````````````y````x````p`````
``````````````````t````p````
````````````````````````i```
`````````````````````````n``
``````````````````````````g`
~~~

Figure 2: Spam Sample I

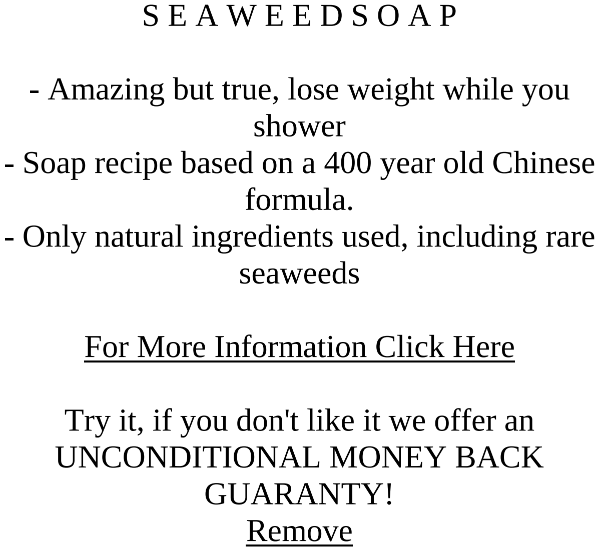

SEAWEEDSOAP

- Amazing but true, lose weight while you shower
- Soap recipe based on a 400 year old Chinese formula.
- Only natural ingredients used, including rare seaweeds

For More Information Click Here

Try it, if you don't like it we offer an UNCONDITIONAL MONEY BACK GUARANTY!
Remove

Figure 3: Spam Sample II

Figure 2 is a spam example that replicated other attempts that successfully fooled spam filters. The tactic used by these messages was the use of punctuation marks and HTML tags to separate the letters in the message.

Figure 3 represents a plain text e-mail that was not flagged as spam.

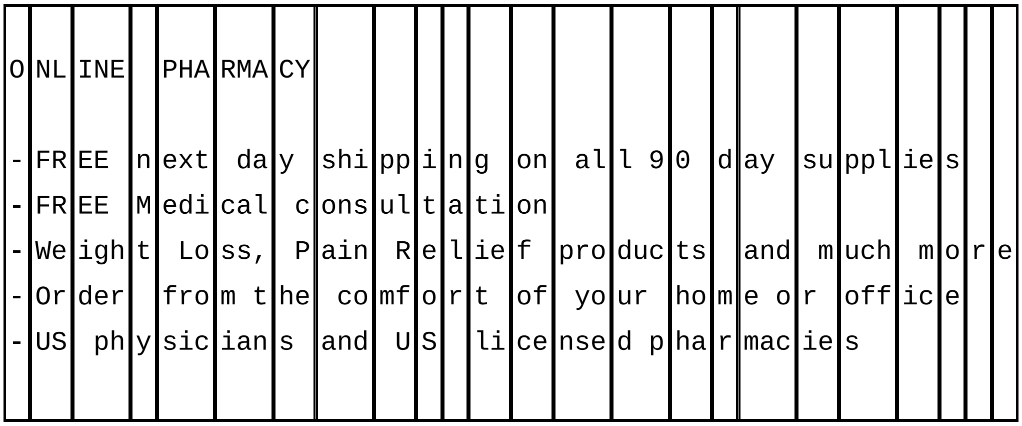

O	NL	INE		PHA	RMA	CY																	
-	FR	EE	n	ext	da	y	shi	pp	i	n	g	on	al	l 9	0	d	ay	su	ppl	ie	s		
-	FR	EE	M	edi	cal	c	ons	ul	t	a	ti	on											
-	We	igh	t	Lo	ss,	P	ain	R	e	l	ie	f	pro	duc	ts		and	m	uch	m	o	r	e
-	Or	der		fro	m t	he	co	mf	o	r	t	of	yo	ur	ho	m	e o	r	off	ic	e		
-	US	ph	y	sic	ian	s	and	U	S		li	ce	nse	d p	ha	r	mac	ie	s				

Figure 4: Spam Sample III with visible table

The preceding example in Figure 4 shows how HTML tables are used to confuse filters. In order to see the full effects of the table, it was necessary to highlight the borders produced by the table. Tables in spam messages are usually not visible to the recipient.

The following spam message is an attempt to personalise the spam message, but the name in this case is incorrect:

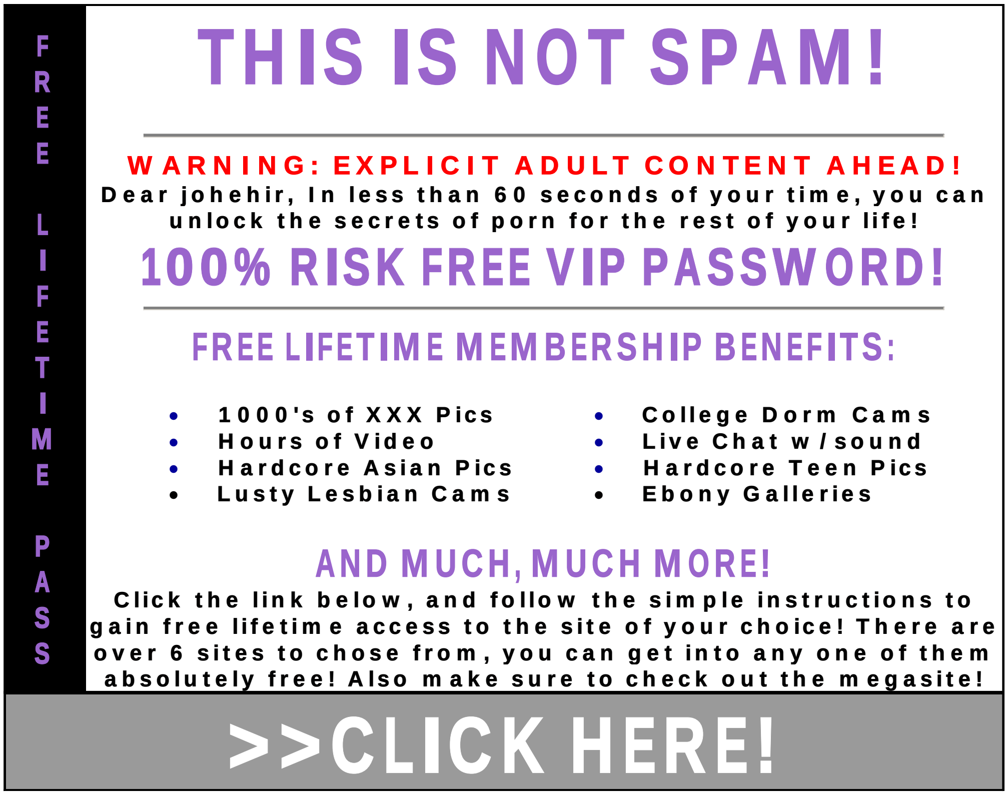

FREE LIFETIME PASS

THIS IS NOT SPAM!

WARNING: EXPLICIT ADULT CONTENT AHEAD!

Dear johehir, In less than 60 seconds of your time, you can unlock the secrets of porn for the rest of your life!

100% RISK FREE VIP PASSWORD!

FREE LIFETIME MEMBERSHIP BENEFITS:

- 1000's of XXX Pics
- Hours of Video
- Hardcore Asian Pics
- Lusty Lesbian Cams
- College Dorm Cams
- Live Chat w/sound
- Hardcore Teen Pics
- Ebony Galleries

AND MUCH, MUCH MORE!

Click the link below, and follow the simple instructions to gain free lifetime access to the site of your choice! There are over 6 sites to chose from, you can get into any one of them absolutely free! Also make sure to check out the megasite!

>>CLICK HERE!

Figure 5: Spam Sample IV

There is also an effort to reduce the scoring of the message by introducing infrequently used words. They are included in the HTML table but hidden from the recipient by having identical font and background colours. The hidden words were:

> fratricidal sheeting opiniastre interiorly heteroepic Europa featherer anounou commensalistic vistamente tomnoup registership rinch daggy anisilic protoreptilian multivious villously antikinase epicondylic possessingness sist decide toadlike hornwork tipsify flakily numerative embryographic homostyled guaba Cyclanthales dolichocercic brightness.

The next example is a Viagra spam e-mail, but in this case plain text in an HTML table is used.

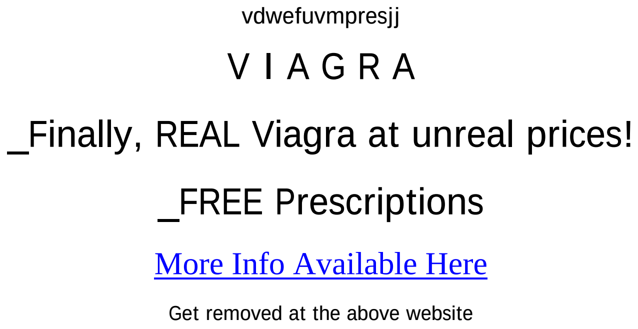

vdwefuvmpresjj

VIAGRA

_Finally, REAL Viagra at unreal prices!

_FREE Prescriptions

More Info Available Here

Get removed at the above website

Figure 6: Spam Sample V

The source code for the last line in this message is:

```
<font size =1>G<cwxhlf><fch><kohscub>e<fvuhpg>t<slboacz><vtrcty>
remov<szgnsou><nthe><ouvdoe><bwo>e<uzrdph><rzvtlgk><avjts><pzzc>d at the
<tnoq>ab<jmqjeo>ove<gyjgtpj> w<vyczlf>ebsit<jpiwizs>e<font>
<ivfm><kbavwji>
```

The ridiculous-looking HTML tags are just ignored by the browser. The only purpose of the tags is to trick the filter by breaking up the words in the text message.

The following representation is a sampling of other types of spam e-mail that were received in this research:

Date and Time	Subject	Sender	Spam Category
Thu, 29 May 2003 21:18:26	PLEASE ASSIST	real-dukemetu@zwallet.com	Nigerian
Sun, 13 Jul 2003 15:22:42	ddomIt"s.Bettter.Big, askk her!rp	vbemushirul@sfigbuwyfiegehen.com	Body Enhancement
Tue, 08 Jul 2003 04:12:45	dGGeta ilDiiploomaa fast	huakhom@aorpvjoztek.com	University Diploma
Sat, 12 Jul 2003 00:57:19	uduyj.EEmpoweer.your.. manhhooodd.-why wait62	whfnhon@esnogromanufa.com	Viagra
Thu, 10 Jul 2003: 00:29:29	cfbddkGGet skkiinny, easyy6648	wenjinamjlyo@stanislauswbchiak.com	Weight Loss
Tue, 12 Aug 2003 07:03:33	All inclusive VIP pass	briskish15133@excite.com	Porn Web Site

Table 2: Spam Samples

6.3.2.1 Classic Spam Example

The last spam example in the previous table (Table 2) contains almost all the classic signs of spam e-mail. First of all, the date was invalid as it was dated in the future (i.e. it was received on August 11th, but was marked sent on August 12th). When an attempt was made to reply to the sender, it was returned undeliverable with the message "User unknown in local recipient table" (i.e. the user did not exist), so the sender's e-mail address was spoofed. When the option to unsubscribe to future mailings was exercised by clicking on the unsubscribe link in the spam message, the link was not valid (i.e. the page was not available). The spam contents consisted of mostly HTML, which is a common spam masking technique used by contemporary spammers. Also, the 'To' header did not contain the recipient's e-mail address, which signifies the use of 'bcc' (blind carbon copy) to send the spam message. This is a snapshot of the actual spam e-mail:

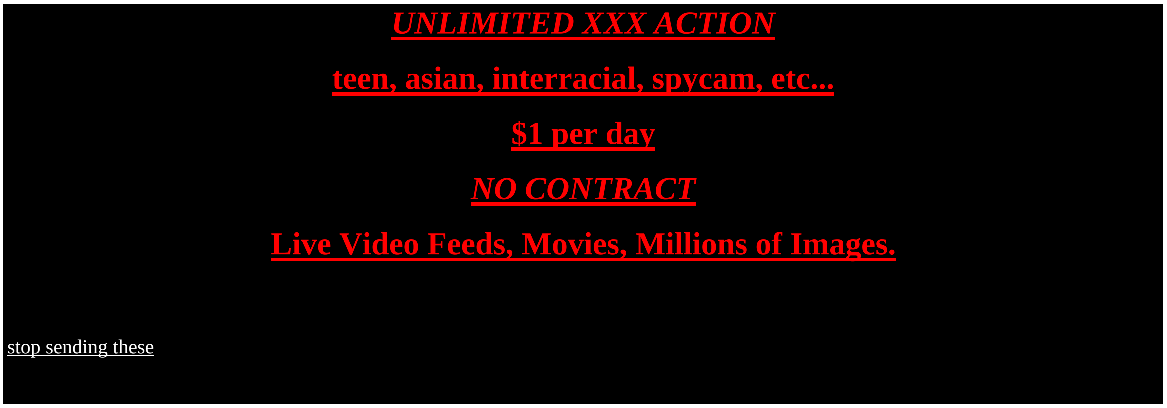

UNLIMITED XXX ACTION

teen, asian, interracial, spycam, etc...

$1 per day

NO CONTRACT

Live Video Feeds, Movies, Millions of Images.

stop sending these

Figure 7: Classic Spam Example

The spam filter (SpamAssassin) also scored the sender's e-mail address (i.e. ends with numbers) to indicate the probability that the e-mail was spam. One interesting additional feature in this message is the random text that appeared after the "stop sending these" unsubscribe link:

> theftdom bali Itelmes prevenance midwestward lymphaemia nonoccupation culpable Dicotylidae ovariostomy Mab venatorious heteropod urdee unselling yachtsman Lamiaceae unmoldable digestive amalgamative Halosaurus nictation episiotomy izle Neuroptera yellowtop fellation ultraexcessive hyperimmunization faradic languid sidelang Eloise Imeritian.

Even though the text is not visible in the message (as the font colour is black against a black background), it is used to trick the filter by posing as legitimate text. The fact that these words are unusual tells me that they are intended to defeat content filters, as rare words can decrease the probability that an e-mail is spam (See Bayesian filters in Section 4.2.3).

6.3.2.2 Spam Source Code

The following piece of code was extracted from the source code of a spam message received by one of the honeypot addresses:

```
W<n48c>e <rl52><xeuk>guar<fl29><h3r0>a<et40>n<v902>t<npz8><fn5f>ee
<r9k2><aqgj><ecjz>t<an10>h<v047>e<os5o><yt3i><mqj2><br><l33c><t97m>
<rr66>suc<klp0><uy09><i3p5>c<j481>es<n5r0><ys02><z69s><j92a><woju>s
of<vy06><ud19><aa7r><n545>
<o62j>o<q53v>u<rb54><g9tj><vkjd><oq22><h10b><r182>r<myvw>
p<tp3o>r<i7jd>o<zq9j><ek6n><bms5><kc5z>gram
o<j536><j42m>r<u3qu><h3aq><x943><m3j8><fki4><sh1x><n591><hia6> <wkza>we
<r5o0><m39o>w<ysi2><rb8m>i<pn83><o7wq>ll
r<m1os><jcj6><x742>e<n182><x721><li50><eo5z>f<hdzf>u<ba7d><hw65>n<n64r><s0fd><u
6p6><j28v><n6f8><l5l1>d
e<jdy9><z214><i9w6><rqen><v1s8><do3p><w9m9><wn11><rb22><f38d><t16c>very<br><kl
48><y2h3><u71u><zkqp>
<ug2d><i2qu>penn<dqg3><wvlm><b1xv><qd72>y.<mnf0> <r945><di18><s61m>
<urw7>Co<gcf1><w7o3>m<gn80>e<g749>
f<h951><yu24>i<gk1i>nd<w26s><x9a6><f9ka><on47> out<vga5>
<my79>w<ag55><qiqo><aat0><l0i2>hy<r952> <b3aa>m<fn1u><mq37><p32t>ore
<g670><zv6j>m<q83x><i5v0><oyx9>e<iz81>n <v2pv><fy00>A<zy92>ND<sa7p>
<w834>WO<rj68>MEN<br><vjr1><y8jn><v92a><o0q7><a312><u7rk><w4y1><d8vm>
<eo78><t231><zsc8><yzdx>c<dq70><q797><u165>o<yn39><x5p2>m<fd2j>e<yg96><ddx6>
<u13h>t<o9u9><hx48>o<vxn6> us<kvtv><h0o2><ji42><g0u5>
t<cb3z><eg4j>han<uvn8><q463> <l170><l2ds>any<lz3l>
<om2r><lev3>o<lp30>t<qe3s>he<ase9><yh10>r<mh74>
<ykxb>si<m360>te.<zi0w><br><o834>
<p37m><br>
```

Figure 8: Spam Source Code I

It can easily be seen how this code broke up the contents of the message by using HTML tags. All of this code was used to convey the following lines in the spam message:

“We guarantee the success of our program or we will refund every penny. Come find out why more men AND WOMEN come to us than any other site.”

This next example of spam source code was found embedded in a table for a Viagra spam message:

```
<vx2i><zc01>
<l254><l3vu><nbm2><j48d><font size = 5 face = "Courier New, Courier, mono"><b><kult>
_<ri0b><fib7><z14b>___<kry1>_<qn99>__<p9l6><k477><n1pn><mqo5>__<c1is>__<b40o><s
8id>__<t0rr>__<qr28><e7s6><r36y>_<h97j><k7uh>__<ilig><g371><e2rd>_<y557><j03q><xo
hd>__<fgy9>_<p23x><br><lip6>
_<d258>_R<s66v>_<e33d>_<b4f8>_V_<mp0q>___<s36p><r672><rc49>_<wbyq>_<d336><i2
2u>_<ai7p>U_<w08d>___<dq96>P__<rxfl><br>
<h5ui><i756>_<jz40>E__<o47r>_I___A_<l7i4>__<a08a>N_<lz5r>__<vgmx>_<c464>R__<br>
<iti1>
_<kb7i><expx><sb5j><xh7x><d67o><l3p3>_A<gt20>__<m139>_<ufzq>A___T<q57k>_<x48c
><dj7z><h022>__R<dvd0>_<i6z0>_<yi8x><hv84><i0ps>__I__<br><b24a>
__<oz4n><x4w9><f0h2><ae6k>L___G<h70v><w1l9>____<r51a>_<b379>_<p3rz>_E_<lt8l>_<
o2k6>__<l7l4>C<i95r>__<br><kh54>
______R___<k1y0>_<alqu>__<u811><mug9><iziz><n2n2><h8n3>_A<zn35><c6ps><esa4>_<
xo3c><pfs3>_<qv5x>__E<t9oq>__<b0j9><br><a7p0><naz6>
_<e885><i7l2>_<n0wq>_<oyl8>_<y2e6><x2l1><q5o1><ve6u><yj5r>__A__<ie7e><co12><vz6
7>_<s2v5>__<hz3w>__<x48b>L_<xg6t>_<f7e5><oib1>_<n37j><tuv5><y1ex>_S_<u056>_<vq
r5><br><br><cfht><t72i><o6br></b><o1b1></font>
```

Figure 9: Spam Source Code II

This code produced the following visually effective output:

```
__R___V_______U____P__
__E___I___A___N____R__
__A___A___T___R____I__
__L___G_______E____C__
______R_______A____E__
______A_______L____S__
```

Figure 10: Output of Spam Source Code II

6.3.3 How Spam Tricked the Filters

When the spam e-mail obtained in this research is examined closely, the methods used to avoid current state of the art spam filters became apparent. There was one particular persistence e-mail message that easily tricked the filters by using a GIF image as the main body content. This meant that there was no text in the body of the e-mail message for the filter to analyse. In order to defeat subject header analysis, the spammer changed the subject for each e-mail message that was sent. E-mail administrators and some filters might view an e-mail message with the same subject going to multiple recipients as a spam message, so the spammer has circumvented this spam flag by randomly changing the subject in the e-mail. They do attempt though to present some semblance of the nature of the e-mail in the subject of the e-mail. The following is an example of the different subject headers for the same spam e-mail that was received by a number of honeypot accounts:

ddomIt"s.Better.Big, ask her!rp
bbeqxBBettteer to be bigg uek
zzvBetttterr ttoo bee bigg eenn
gbvhh A mmaassivee "”one””f aawaits wdzz
ohggdIt's.Bettterr.Bigg,, ask herr!!qkkmnn
hjBetterr to be big kbb
rIt's.Bettter..Big, ask her!hh
maoaa A mmassive "'one"Unknowwnn Funnctioon awaits lg
vvttrIt's..Betterr.Big,aassk her!!w

All of these lines appeared on the subject header of the same spam message delivered to different honeypot accounts in this research. It can be observed that none of the subject lines are the same for these identical spam messages. Even though the text is a bit garbled, one would still have a fair idea of the nature of the e-mail from the subject line. The source of the spam was undoubtedly the same as they all appeared in the same time frame and the sender's e-mail address was spoofed with names like:
vbemushirul@sfigbuwyfiegehen.com and xthang@ycjdcmucplebanek.com.

The hyperlink on all of these particular spam messages pointed to the same web site, where you could purchase the product advertised in the e-mail. As I expected, the web site had special built-in features to hide the source code e.g. the right-click option to view the source was disabled, the alt key could not be used, the browser menus were hidden etc. In fact at one stage, this spam message was generating so much traffic to the advertised web site that it was temporarily shut down due to the web site exceeding its ISP's data quota. This fact has to be considered alarming as it proves that this type of spam message is capable of eliciting responses from spam recipients. Either that, or an anti-spammer was enacting revenge on the site.

Other spam messages in the corpus had similar features i.e. spoofed sender addresses and varying subject lines for the same message. I also noticed that the same spammer would send different types of spam to the same accounts. For example in one instance the spammer sent separate spam messages that advertised University Diplomas (with only text in the body of the e-mail), body enhancements (using a table with HTML tags in the body of the e-mail) and Viagra (text body). Even though the sender's e-mail address was obviously spoofed with non-functioning accounts, it was also obvious that the messages came from the one source as the same group of 5 recipients appeared in the "To" header of the e-mail, and the addresses were listed in same order as the other spam messages. Another interesting fact that was observed in the corpus was the spammer's use of animated GIF images in order to make their spam message more visually appealing to the reader.

6.3.4 Breakdown of Spam by Baiting Method

It is not accurate to assume that just because a greater percentage of spam was received by one baiting method in this research, then that is the method of choice for all spammers. There are many variables involved in the success of the baiting methods. The following chart is a breakdown of the baiting methods and the percentage of the overall spam that was received. Even though the chart lists the popup baiting method as the number one source of spam, it needs to be pointed out that 2 dictionary attacks skewed the actual results. This attack is discussed later.

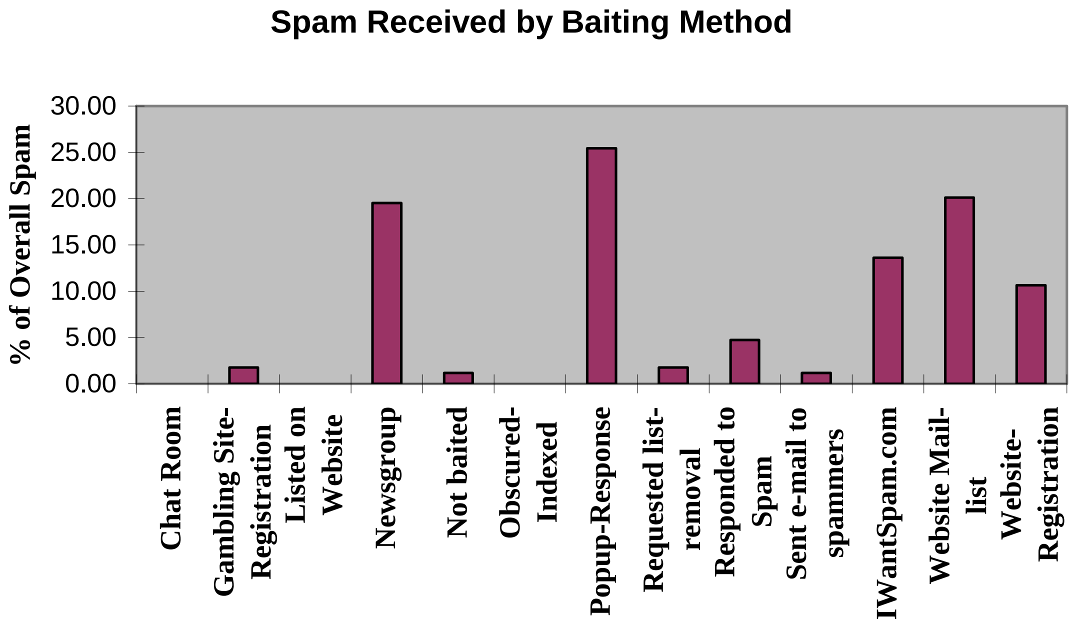

Figure 11: Breakdown of Spam by Baiting Method

It is widely known that the majority of spam is sent to lists that are created by harvesting e-mail addresses on web sites. There could be a number of reasons why the honeypot addresses listed on the web page did not receive as much spam as the other methods. One of the reasons could be the fact that spammers did not index the web page because it was not a popular web page, or the fact that the e-mail addresses appeared on the web page for only 3 months and the spammers' harvesters just did not get to the site yet. With regard to newsgroups, there is the possibility that the bait addresses listed in these newsgroups may have received more spam if I posted to the groups more often.

6.4 Spam Scoring

In order for e-mail messages to get classified as spam, a spam filter such as SpamAssassin uses a scoring mechanism where points are awarded based on certain criteria. Microsoft Outlook has a built-in feature that automatically filters e-mail using SpamAssassin and categorises any e-mail with a score greater than 5.0 as spam. Features of spam messages in this research that scored prominently included:

Invalid date: Sometimes spammers send messages with dates in the future.

From and To are the same: This feature scored high as it allowed the spammer to send a spam message to himself/herself, and then use the ‘bcc’ (blind carbon copy) option to send the same message to a large number of recipients. The recipient would not see the other users on the list.

From: ends in numbers.

Dear No-Name: As SpamAssassin reports: “How dear can you be if you don’t know my name?” However, spammers attempt to obtain at least the first names of e-mail address owners when they harvest or buy lists in order to personalise their message, and lower the overall spam score. Subject is all capitals, Subject has lots of exclamation marks, HTML-only mail with no text, are other features that increases the likelihood that the message is spam. It should be mentioned that spammers attempt to overcome these types of filters by creating messages that scores below 5.0.

It should also be noted that the current version of SpamAssassin has a staggering 921 similar tests that it performs in order to determine if a message is spam or not.

6.5 Related Findings

This section deals with some of the other interesting findings that were revealed during the course of the research.

6.5.1 Honeypot Anti-Spam Techniques are Flawed

There are many concerted efforts by anti-spam organisations and individuals to frustrate spammers by using honeypot techniques. Some of these efforts involve considerable effort and time and many involve sophisticated anti-spamming methods. For example, some sites automatically generate web pages with bogus addresses with the intent of both polluting the spammer’s database and slowing down the spider program. Not only does it slow down the web server that is hosting this site, but it also consumes the owner’s (victim’s) bandwidth. I have serious doubts if these techniques work at all. In this research, modern e-mail harvesters were able to detect these spider traps and exit out immediately and continue crawling to other web sites.

One particular web site went to the trouble of creating two web pages, each with 10,000 fake e-mail addresses. Unfortunately, most if not all of the current e-mail harvesters will just bypass these pages. If the spammer is using an old harvester, then it may try to extract these bogus addresses, however spammers are well aware of this type of spider trap. The web pages themselves are over 500 kilobytes in size and must have taken a considerable effort to create, even with a fake e-mail address generator. There are only 3½ lines of text at the start of the 2 pages that contains the 20,000 fake addresses, but there are still enough of keywords on these lines that signals a harvester that the web site is a trap. Here are the actual lines of text on the first web page:

> “Below are 10,000 bogus email addresses for the email spiders that harvest your email addresses for spam. If you are tired of spam, copy this list and post it on as many web servers as you can. If the spammers spend all their time sending junk to junk addresses then it may slow them down.”

Harvesters have spam detection utilities that searches for these words and upon finding even one of the words “bogus”, “spiders”, “harvest”, “spam”, “spammers”, “junk” (all of which are contained in this particular spider trap), the e-mail harvester will just exit that site, and proceed to the next URL on its list.

Most Internet users prefer to have an e-mail address with all letters e.g. Anselm.Lambert@tcd.ie is preferred over Anselm.Lambert01@tcd.ie. Spammers are aware of this fact and as part of the spider’s spambait detection mechanism many modern-day harvester will bypass e-mail addresses containing numbers. A great number of spider traps contain addresses with alphanumeric characters, which makes these traps useless against the current breed of e-mail harvester. On the other hand, this piece of information could also keep your e-mail address off spam lists i.e. picking an unusual username (with a numeric) improves your chances of not getting spam.

6.5.2 Spam Tools Findings

The best way to gain an insight into how the spammer collects e-mail addresses was to frequent web sites and newsgroups where they congregate. The spammer’s “tricks of the trade” are readily available on the web, and the following section describes some of the spammer applications used in this research.

6.5.3 E-mail Harvester Experiment

I found a powerful e-mail harvester at http://www.bulk-email-lists.com. Before I ran the setup executable to launch the harvester installation process, it was necessary to install security patches and updates to the Windows 2000 operating system on the research computer. It was also necessary to monitor network and Internet access using Sygate Personal Firewall, just in case there was rogue access to or from the system. The .NET Framework was required to run this program, so version 1.1 was downloaded from Microsoft and installed. The actual harvester program installation took only about 10 seconds to complete, and I was harvesting e-mail addresses within another 10 seconds. The only requirement to start the harvesting process was to include a starting URL in a very intuitive graphical interface. The following screenshot is the configuration screen for the harvester:

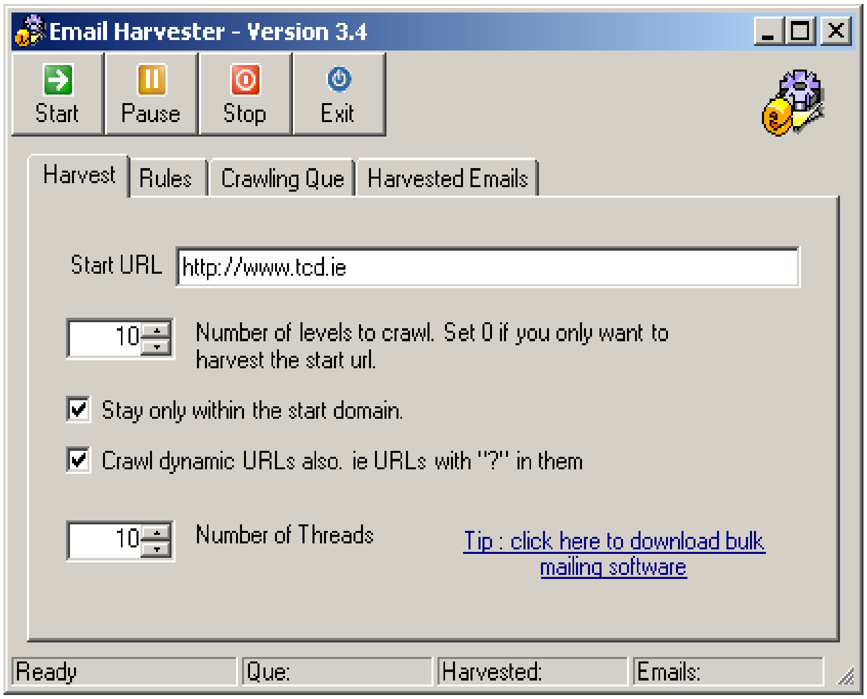

Start URL http://www.tcd.ie

10 Number of levels to crawl. Set 0 if you only want to harvest the start url.

☑ Stay only within the start domain.

☑ Crawl dynamic URLs also. ie URLs with "?" in them

10 Number of Threads

Tip : click here to download bulk mailing software

Ready | Que: | Harvested: | Emails:

Figure 12: Opening Screenshot of E-mail Harvester

6.5.3.1 E-mail Harvester Findings

The rules in this program consisted of rejecting pages that contained the words spam, bait, blacklist, robots and junk. Other words can also be included. One of the first astounding results in my research is that the vast majority of spambait web pages use at least one of these words and they also contain multiple bait e-mail addresses. This would indicate that these efforts are useless against the current breed of harvesting programs. For example, the creator of a spider trap would give an explanation of the e-mail addresses listed on the webpage. Using the smart spambait detection tool, the harvester would just bypass this page based on the appearance of any of the keywords. Another rule is the option to use the smart spambait detection utility that rejects pages that have more than a specified number of e-mail addresses. All of the spambait websites that I visited had more than 10 bait addresses listed on the bait page. 10 addresses is the default number that the harvester is configured with, which means it will skip pages with 10 e-mail addresses or more. It should be noted that there are many e-mail harvester programs available, but it does appear that the spammer has figured out a simple way to bypass honeypot addresses. The harvest results for the starting URL: http://www.tcd.ie can be seen in the following screenshot:

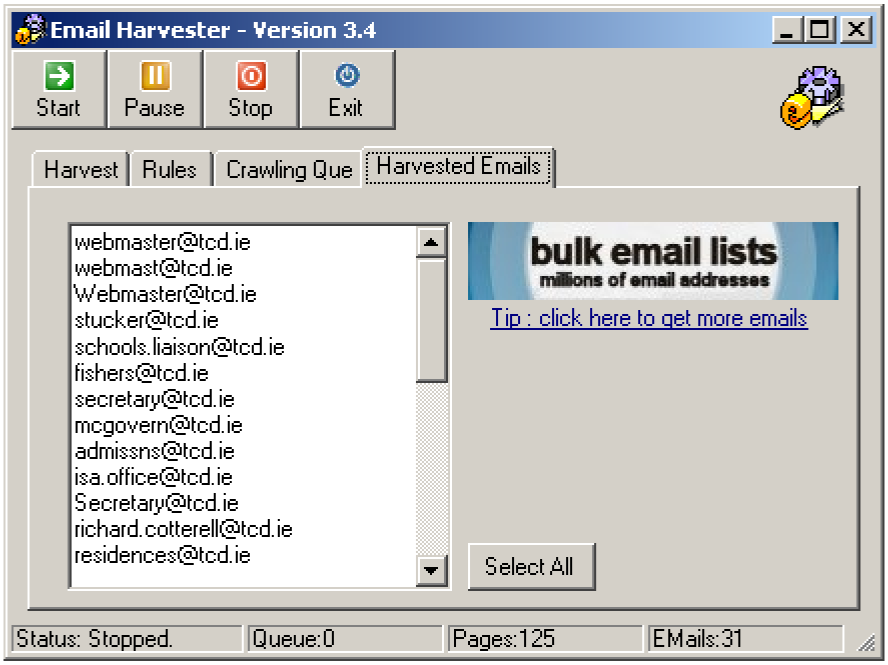
Email Harvester - Version 3.4
Start
Pause
Stop
Exit
Harvest
Rules
Crawling Que
Harvested Emails
bulk email lists
millions of email addresses
Tip : click here to get more emails
Select All

webmaster@tcd.ie
webmast@tcd.ie
Webmaster@tcd.ie
stucker@tcd.ie
schools.liaison@tcd.ie
fishers@tcd.ie
secretary@tcd.ie
mcgovern@tcd.ie
admissns@tcd.ie
isa.office@tcd.ie
Secretary@tcd.ie
richard.cotterell@tcd.ie
residences@tcd.ie

Status: Stopped. Queue:0 Pages:125 EMails:31

Figure 13: Partial Results of Program Harvesting the www.tcd.ie Website

The harvester was stopped after it crawled through 125 web pages and collected 31 e-mail addresses. It could have been configured to run indefinitely by specifying more than 10 levels. It also has the ability to crawl dynamic URLs i.e. URLs with "?" in them. By looking at the results screenshot, the next very interesting finding came to light. The harvester was bypassing e-mail addresses that included numbers. It was at this point, I came to the realisation that this was probably the reason why the honeypot addresses that I was using on my bait web page were not finding their way onto spam lists. All of my honeypot addresses were alphanumeric:
john20000007@hotmail.com, john20000007@eudoramail.com and so on.
New honeypot accounts had to be created and baited for indexing due to this finding.

When the starting URL on the harvester was given the web site with the honeypot addresses used in this research, it bypassed all of them. It did collect my real address at the bottom of the page i.e. Anselm.Lambert@cs.tcd.ie. This was an indicator of the suitable addresses that this program would collect. The indexed web page with the honeypot addresses was then edited to remove the numbers from the e-mail addresses (for example john20000007@eudoramail.com was replaced with john@eudoramail.com), and subsequently ran through the harvester again, which now collected all of the bait addresses. With regard to performance, there is an option on the harvester for increased effectiveness and speed. This is achieved by selecting the number of threads used to crawl the web looking for e-mail addresses.

A discovery that came to light at this juncture of the research was the ability of the spammer to harvest addresses from web sites that require a logon. Many site administrators believe that they may be protected from web robots extracting their addresses by using a username and password. However an internal browser that comes with the harvester is used to log into a site and as the pages are browsed manually, the e-mail addresses are harvested automatically as the web pages are loaded. The personal firewall that I was using on my system highlighted another interesting finding. One of the e-mail harvesters that I used in this research was designed to use the Yahoo online directory to retrieve e-mail addresses for targeted spamming purposes. This is discussed in more detail in Section 7.4.

6.5.3.2 Keyword Search Harvester Experiment

There is a type of harvesting program that allows spammers to harvest e-mail addresses based on a keyword search. I used one of these programs to examine the possibility of creating accurate target lists based on the topic in the keyword search. The program works by using the Google search engine to retrieve a list of web pages based on the keyword search criteria.

6.5.3.3 Keyword Search Harvester Findings

On one occasion, I entered “distributed systems” into the search box and using a 56Kbps dial-up connection, the program was allowed to run for 5 minutes. It crawled 334 links and retrieved 203 addresses using 5 open connections. The program gives the option to vary the number of connections so as to speed up the process of harvesting addresses, with the consequence of more processing power usage. Out of the 203 addresses, there were many duplicates but the program had a “clean addresses” buttons that removed duplicates and brought the number of harvested addresses down to 63. Out of these 63 e-mail addresses, there were many addresses that would be of no use to a spammer, for example: info@domain.com, webmaster@domain.com, pagemaster@domain.com etc.

In another keyword search, the keyword harvester retrieved 1618 addresses using the search criteria “anselm lambert trinity college”. Once again, a dial-up connection with 5 open connections and a depth level of 2 was chosen for this search. The depth level setting tells the program how deep into a web site it should crawl in order to retrieve addresses. After cleaning the initial list of 1618 addresses, it was reduced to 490 e-mail addresses. It should be noted that this harvester retrieved the new honeypot addresses used in this research located at http://www.cs.tcd.ie/Anselm.Lambert/, but the obfuscated addresses on the web page were skipped. This proves that e-mail address obfuscation does indeed work. None of the harvesters in this research were able to successfully retrieve the obfuscated e-mail addresses on the bait web site.

I did not find the keyword search harvester to be a reliable source of targeted e-mail addresses, as the program quickly lost its focus on the keyword topic and crawled many irrelevant web sites. This in turn harvested e-mail addresses that were inappropriate for the list based on the keyword search. Even though some addresses may be suitable for the list, the large number of inaccurate addresses made the list useless as a targeted list. However, I would imagine that this would not be of great concern to unscrupulous spammers and the fact that a small percentage of the retrieved addresses are accurate would be sufficient for them to spam the whole list. Another point of interest that I noted during the keyword search examination was that the harvester not only searched Google for the source of hyperlinks to crawl looking for addresses, but the program also searched Google's cache just in case the web page changed or was removed. Google's cache is a snapshot of the web page that it takes during the indexing process, so that it can be used to search web pages that are no longer available.

6.5.3.4 Beating the Spider Trap

In this research it was shown that the current state of the art harvesting program has the ability to realise if the program is stuck in a spider trap. Once this situation is detected, the harvester will not terminate but will get out of the trap and continue crawling the web for more e-mail addresses.

6.5.4 Opting out and Unsubscribing

Some of the honeypot accounts in this research were configured to receive newsletters and special offers for the purpose of examining the integrity of the unsubscribing or opting-out process. It was revealed in previous research that unsubscribing to such lists could be a tedious task, and could require a number of attempts before successful removal. But in the majority of cases it was eventually possible to be completely removed from these subscription-based lists [28]. After receiving a substantial number of spam-like e-mails by subscribing to various services and newsletters, I initiated the unsubscribe process on July 15th 2003 on six of the honeypot addresses. The purpose was to find out if the request to unsubscribe would be honoured and if so, how long it would take. If the process took too long, then any e-mail received from these sources after a reasonable amount of time (for example, 15 working days) should be judged as spam messages. Out of the six accounts, three addresses continued to receive e-mail even though three weeks had elapsed since the confirmed receipt of the unsubscribed action had been received from the ISP. The other three accounts stopped receiving newsletters almost immediately after the request. This was surprising as Mail.com honoured the request even though the company's accounts is notorious for being a source of spam and the honeypot accounts created with Mail.com were inundated with popup advertisements each time the accounts were checked for e-mail. On the other hand, Yahoo continued to send "special offers" to two honeypot addresses even though the company acknowledged receiving the opt-out requests. A second request to Yahoo was required to stop these mailings.

However, when an attempt was made to unsubscribe to future mailings by using the opt-out link in a spam message, only one link in 12 spam e-mails worked. The other attempts to unsubscribe led to the activation of popup advertisements when the link was clicked, or to web redirections to other advertisements. This confirms what is already known: Unsubscribing to spam from within the spam message itself or associated web site, does not work.

6.5.5 Other Findings

There is a very useful option with some e-mail clients that provides a confirmation if e-mail was successfully delivered and in some cases it can provide information on whether e-mail was opened and read. I thought that this would be a very useful feature for a spammer, as it would confirm the authenticity of an e-mail address. To test this feature, I used Microsoft Outlook to send messages to honeypot addresses at the 5 different e-mail providers used in this research i.e. Yahoo, Hotmail, Mail.com, Eudoramail and Campus.ie. Even though the Outlook client requested confirmation that the recipients opened the e-mail, none of the providers sent this confirmation when the e-mail was opened up and read by the accounts in this experiment. It makes perfect sense for the provider that offers free e-mail not to acknowledge successful receipt of e-mail as it not only denies the spammer confirmation of a legitimate address, but it also allows the provider to save on the costs of sending e-mail receipts.

6.5.6 Success of E-mail Address Obfuscation

The success of the e-mail obfuscation methods used in this research can be measured by the amount of spam that these accounts received. The e-mail addresses that were obfuscated in this research did not receive any spam.

6.6 Evaluation of Project Results

If an e-mail address appears on a web site then the degree of spam that the address receives depends on the extent of indexing of that web site. That is to say, a web site that is heavily indexed will be subject to a greater number of visits by e-mail harvesters. The e-mail addresses that were baited on the web page in this research would have received a greater number of spam messages if they were included on a more frequently visited web page. It is critical for serious spammers to obtain fresh e-mail addresses, so the harvesting of web sites and newsgroups is still the preferred choice to build spam lists. In this research, responding to popup windows was the baiting method that received the most spam. The e-mail received as a result of entering the bait addresses into the popup window was not counted as spam if the e-mail received by these addresses originated from the domain listed in the popup. Nevertheless, I believe that the spam delivered to these popup bait e-mail addresses was as a result of a dictionary attack (as discussed in Section 6.3.1.1) and had nothing to do with responding to a popup query. Two dictionary attacks actually resulted in the 10 compromised e-mail addresses receiving approximately two-thirds of the total spam received.

6.7 Chapter Summary

During the course of this research, I quickly realised that spammers will use all available existing technologies and services in order to ply their trade. For example, in order to create e-mail lists that are based on countries or states, the spammer can just use Yahoo’s e-mail directory. The work has already been done for them - the spammer just figured out how to retrieve the information from a legitimate source. The keyword search harvester, which is heralded by spamware creators as an excellent source for targeted lists, uses the services of the very respected Google search engine. I do not believe that the owners of Google and Yahoo would take great joy in the fact that spammers are abusing their services. In this chapter the results of the baiting methods was presented. Examples of the spam received by the honeypot addresses were shown to be consistent with the current breed of spam messages delivered to many e-mail users today. The spam received by the honeypot accounts was non-discriminating i.e. a user’s profile played no part in the type of spam that they received. Categorising e-mail as spam depended on the ISP and the level of filtering used. In some instances, the same type of e-mail would be flagged by one ISP and not by others. There was a number of cases when filters that were configured in default mode did not catch obvious spam messages, but with honeypot accounts that operated in enhanced mode, these same messages were automatically sent to the spam folder. The findings in this chapter are summarised in Section 8.2.

Chapter 7

Spam Targeting

7.1 Overview

One good lead is better than a thousand fruitless ones, so you would presume that targeting would be the method of choice for spammers. Unfortunately, many spammers send their bulk e-mail indiscriminately. I have seen reports from people who used bulk lists and did not receive a single positive response from a mass mailing (even with lists greater than one million e-mail addresses). The average response for a non-targeted e-mail campaign is 15 per million. It should be mentioned that almost every web site and company that deals with targeting lists declare that all of the e-mail recipients on their lists have opted-in to be on the list. With some of these lists reaching into the millions, it is highly doubtful that this is the case. Spammers, in general, claim as their first line of defence against a spamming accusation that the recipients on their lists elected to receive their spam messages by opting-in to the list. A point of interest that I noticed in my research of list sellers is the consistent spelling mistakes and grammatical errors on their web sites, which may reflect the accuracy and reliability of the content on these sites. Some spammers believe that if you target a wrong product to the wrong audience, then it should be classified as spam. Using that logic, they must also believe that targeting an audience with a product or service that the spammer feels is suited to the recipient is not spam. This logic is quite unreasonable.

From my research I found that spam targeting is not the norm and recent studies tend to agree with my findings. In June 2003, Applied Research [29] released the findings of a study that was commissioned by Symantec Corporation, which found that four out of every five children receive inappropriate spam e-mail. The nature of these e-mails ranged from get-rich-quick schemes to pornographic material. The majority of 1,000 children, aged between 7 and 18, interviewed for this survey said that they felt “uncomfortable and offended when seeing improper e-mail content.” It was found that one in five children opened up and read the spam. 47% of the children received spam with links to pornographic material. Whatever ideological feelings you may have regarding the rights of spammers to be able to send this type of material, the results of these findings are still alarming.

According to one creator of spam tools (InfinityMailer), targeting e-mail to prospective customers by demographic, geographic, or other specific interest will pay dividend in the long run. They believe it is possible to get as many leads with a targeted list of 15,000, as it is to get with an untargeted list of 150,000 (or more). This company provides anonymous bulk e-mail software and sells a tool that makes it possible to extract emails by geographic or specific interest groups.

7.2 The Art of Targeting

The methods used by spammers in targeting users with spam are not very sophisticated. For example, in order to send e-mail about a dietary supplement, spammers use e-mail harvesters to collect e-mail addresses from topical newsgroups such as diet, health and fitness groups. If a person posts a message on one of these groups, then there is a good chance that they would have an interest in a dietary supplement. Another common method used by the spammers is to use a keyword harvester, where the area of interest is included as a keyword search on a harvester. The harvester uses the keywords e.g. weight loss, to search the Google search engine. Then the program queues up the URLs retrieved from that search so that those web pages can be harvested for e-mail addresses. It is assumed (by the spammer) that the Google search results would yield a large group of e-mail addresses, whose users would be interested in the keyword topic. This is less discriminating than the newsgroup search as all the e-mail addresses in the keyword search are sent spam versus people who expressed interest in the topic on the newsgroups. Some targeted lists are created by selecting users who subscribed to relevant newsletters, or who work professionally in the specific area of interest, or from those who picked that particular topic as a hobby when filling out a questionnaire.

Another targeting method used is by country. List-sellers break down their list by country, which is quite easy to do by using the domain suffix. For example, all e-mail addresses ending in "ie" would be grouped into an Irish list. This method would miss all the Irish users who use other domain suffixes such as .com, .org etc. Similarly in the United States, the lists are broken down by state, city or zip code. State lists are very important for targeting recipients as some products and services may be illegal in some states, and some state laws may have actually ban spam and have criminal penalties for violations. "Geotargeting" is the term used by mass mailers to describe the process of targeting geographical locations.

There are some companies that sell e-mail targeting as a service. All that is required is the type of target audience and the e-mail message, from which these companies send out the spam e-mail. I contacted a number of companies that sell target lists in order to gleam any information on how they create these lists. Only one of them replied and they said that they got the e-mail addresses from their affiliates and brokers.

7.3 Targeting Individuals

Opt-in mailing lists have consistently shown to be very effective in starting and maintaining a relationship with a customer, more than other types of Internet advertising. It is proven to be 18% more effective than banner advertising, which has only a minuscule response rate of 0.65% [30]. Opt-in lists should be the method of choice to target individuals with specific e-mails, as the recipients have previously given approval to receive these messages, and more importantly they have already shown an initial interest in the topic of the e-mail. One of the bait addresses in this research opted-in to an e-mail list and received only e-mail specific to the indicated topic. The advertisers went an extra step to confirm the authenticity of the opt-in request by providing confirmation e-mail with the subscribed e-mail address, the IP address and name of the subscribing host, the date and time of the request, full contact information of the advertiser (list host) and finally instructions on how to opt-out of the list. This is the correct way to run an opt-in list and the companies that operate these lists understand the potential returns of responsibly targeting individuals in this manner. It should be pointed out that the list owners are responsible for privacy concerns, such as protecting subscriber information, and selling this information is unauthorised without the expressed consent of the subscriber.

E-mail personalisation is a technique that spammers use in order to improve the response rate of their message. By customising e-mail with the recipient's name, the spammers try to give the impression that the e-mail message is geared specifically to that user. Many list sellers will offer the first names (and sometimes the last names) of e-mail address owners as an added bonus to customers that are buying lists.

Well-known companies that use bulk e-mail marketing implement targeting as part of their business model. Some companies target individuals based on where they live and what they have bought in the past. Hewlett-Packard (HP) is an excellent example of a company who use targeting in their marketing strategy. They tracked customers who bought a printer and used a very effective bulk e-mail blitz to gain more sales. A month after an individual purchased a printer they would receive e-mail from HP that would give tips on how to clear a paper jam. A month later, another e-mail would follow explaining how to extend the life of the toner cartridge. Finally, another e-mail would be sent around the time that the toner cartridge is actually running out, guiding the user to a web site where they could purchase a replacement. One may say that this is targeting put to good use and it resulted in HP's case with a dramatic increase in sales.

7.4 Targeting Groups

CustomerBlast.com is a web site where spammers can purchase targeted lists of people that purportedly have opted-in to these lists. The lists are interesting as it gives an indication of the categories of people that spammers would like to target. CustomerBlast.com also offers to send out the spam to the target audience for a price of

$199 per million e-mails. Here is a sampling of the categories and lists that exists at the moment:

Generic Lists: By Domain Name, By USA Major City, By USA State, By USA Phone Area Code, By Foreign Surfers.

Webmaster Lists: InterNic, General Webmasters, Adult Webmasters, Foreign Webmasters, USA Webmasters.

Telecom Industry: Business Consumers, Residential Consumers, Sales Agents, Telecom Affiliates, Switchless Resellers, Interexchange Carriers, CLECs (Competitive Local Exchange Carrier) and Local Service, Wireless Companies, DSL and ISPs.

Student Database: High School, Technical Schools, Catholic Schools, Other Schools.

Medical Profession: Doctors Office, Nurse Care, EMS Workers, Pharmacy List, Hospital Lists, Other Medical.

Financial List: Stock Brokers, Finance Analysts, Accountants, Banks, Mortgage Brokers, Mortgage Consumer, Credit Repair Seek (presumably of people who sought credit repair online), Credit Card Seek, Other Finance.

Finally, there is a hot list that is regarded to be the target audience who will respond to spam messages in greater numbers than people on the other lists.

Hotlist Consumer: Casino Seekers, Medical Conditions, Fitness Seekers, Boat & Yacht, Golf & Tennis, Hunting Seekers, Smokers Lists, Antique Buyers.

U.S. state lists are a popular item for sale on these target list web sites and one may wonder what is stopping a person from buying a target list and reselling it for a greater price than what was paid for it. The answer is that there is nothing to stop this type of resale. It seems that there is no honour among spammers as some spam support sites have disclaimers that their lists are genuine and U.S. state lists offered on other web sites are just copies of these lists. They also assert in the disclaimer that the resale of these lists by other web sites is illegal. Target lists by individual states are deemed lucrative as it allows businesses to target individuals in their own locality. For example, if a business sold cheap car parts in shops in Los Angeles, then they would prefer to target bulk e-mail to potential customers in California rather than to users in New York.

As part of my research into the targeting of e-mail users from particular states in America, one of the harvesting programs that I used simply connected to email.people.yahoo.com and retrieved the e-mail addresses of users from that directory. When online users create accounts with Yahoo and other Internet providers, one of the requirements is to include their state of residence in the profile section. This harvesting program found a way to connect to the Yahoo directory and download e-mail addresses based on state addresses. With the Sygate Firewall software that was installed on the research computer, it was possible to view the connection that the program was using:

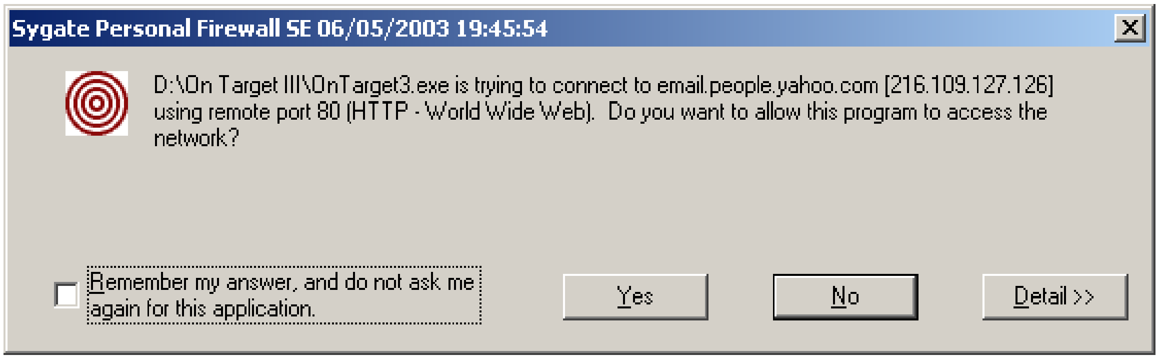

Sygate Personal Firewall SE 06/05/2003 19:45:54

D:\On Target III\OnTarget3.exe is trying to connect to email.people.yahoo.com [216.109.127.126] using remote port 80 (HTTP - World Wide Web). Do you want to allow this program to access the network?

□ Remember my answer, and do not ask me again for this application.

Figure 14: Screenshot of Harvesting Program Connecting to Yahoo

7.4.1 Targeting Experiment

In order to verify the veracity of targeted state lists created by spammer software, an experiment was carried using a software package called On Target III, which purportedly collects the e-mail addresses of users living in specified states in America. Not only could the user of the software program select the state for targeting, but also specific cities within that state could be selected. Using a 56Kbps dial-up connection and selecting the city of Little Rock in Arkansas, 339 addresses were collected in a 5-minute period. The program automatically sorted the e-mail addresses alphabetically as they were coming into the application. From this list, I randomly selected 10 e-mail addresses and sent e-mail to these users requesting verification of the fact that they lived in Little Rock. The result from the test was that 4 of the e-mails bounced back with delivery problems i.e. unknown user, and there was no reply from the remaining six addresses. Surprised by the fact that there was not even one response after a 2-week period, I ran the experiment again for the state of Ohio, and specifically for the city of Dublin. The spam targeting software collected 100 e-mail addresses, which according to the program were supposedly all from the state of Ohio. Again, I sent e-mail to 10 randomly selected addresses requesting confirmation that the recipient lived in that state. The domains were varied e.g. Yahoo, Hotmail and MSN, and the e-mail was also sent to one of my Hotmail accounts to confirm delivery and to make sure that the spam filter for that provider did not flag the e-mail as spam. In this new batch, there were two undeliverable messages but one recipient did respond stating that he did live near Dublin, Ohio at one time. In this experiment, 200 e-mail addresses were collected using the targeting software, and from this list 20 addresses were randomly selected to receive e-mail, with the final outcome of only one response from this group. It should be noted that the person who did respond no longer lived in the state in question, so it must be inferred that the targeting software used in this experiment was quite useless.

7.5 Chapter Summary

In this chapter we looked at the ways spammers attempt to send spam to selected audiences i.e. spam targeting. The accuracy of targeting can only be judged as haphazard at best. From research and the targeting experiment conducted in this research, e-mail targeting software generates lists of mostly old or defunct e-mail addresses. In order to obtain fresh addresses, harvesting of appropriate web sites and newsgroups would be required. I believe that the only way that a foolproof target list can be created is through a registration process. All users who wish to receive mailings would develop a profile, where the information they provide would be the basis of all target lists e.g. location, hobbies, gender, occupation etc. In order to be on these target lists, there would have to be the stipulation that the user has opted-in to receive bulk e-mail, which they may do quite willingly if they are confident that the bulk e-mail that they receive is relevant to their interests. The permission-based system would be beneficial to both the sender and the recipient, therefore this bulk e-mail would not be categorised as spam and allowed through spam filters. As there could be many “spam” words in the e-mail, I would envision some form of challenge-response mechanism followed by the whitelisting of the sender’s e-mail address. Topical e-mail newsletters would be an example of such e-mail.

Chapter 8

Conclusion

8.1 Objectives Fulfilled

This paper is a research-based analysis of spam. A definitive guide to spam was presented, fulfilling the objectives set out at the beginning of the study. Throughout the Internet and in many research papers there is a great deal of study and discussion on spam-specific topics, but I could not find a cohesive study that dealt with all the major aspects of spam. I thought that it was very important that if spam filters are discussed then the economic foundation of spam should also be discussed. This in turn, ties in with the current legislative approaches throughout the world. Anti-spam techniques and state of the art spamming techniques go hand-in-hand, as one greatly influences the other. Readers may not understand the full consequence of the spam problem without reading the extent that spam scams have affected people's lives. A solution is required for this problem but the banning of all bulk e-mail messages is a draconian measure and could even be used by repressive governments to prevent free speech.

8.2 Thesis Achievements and Findings

The honeypot addresses in this research could have been positioned to receive a greater number of spam messages, which is achieved quite easily by visiting adult and gambling web sites, and registering with other dubious services. However, the goal of the project was to try to replicate the usage of the typical e-mail user. Some of the honeypot accounts deliberately tried to attract spam, but most accounts tried to follow a regular Internet user's e-mail pattern e.g. posting to newsgroups, listing the e-mail address on a web site, signing up to receive newsletters etc.

The findings in this research can be summarised as follows:

1. Spammers continuously mutate their messages in order to defeat current spam filtering technologies.
2. The type of spam that a user receives depends on the spammer who obtains their e-mail address. For example, some spammers will not send pornographic spam to users on their lists. Unfortunately, spammers exchange lists, which is the reason why most e-mail users will eventually receive a wide range of diverse e-mails.
3. Many honeypot strategies that are being used at the moment are flawed, as the current breed of e-mail harvester have configurable spambait detection mechanisms. The substantial effort by anti-spammers to quell the tide of spam by using bait websites and honeypot addresses could be in serious jeopardy unless they stay current with state-of-the-art spamming tools.
4. It is possible to create e-mail addresses that are less susceptible to receive spam messages. Dictionary attacks may circumvent this strategy, which was the case for a group of Hotmail honeypot addresses in this study.
5. This research bore out the fact that e-mail address obfuscation works and it should be used if e-mail addresses are going to be listed on web sites.
6. The amount of spam that an e-mail address receives depends on the usage (i.e. how a user propagates the address) and to a lesser degree, luck (i.e. unfortunate enough to have the address listed in the wrong place at the wrong time and end up getting harvested by a spammer).
7. Two-thirds of the spam collected in this study was as a result of a dictionary attack on two groups of honeypot accounts.
8. The type of filter that your ISP uses has a major impact on the decision on whether or not an e-mail message is categorised as spam.
9. Some spammers will make a concerted effort to avoid sending spam to users that could be a source of consternation to them. They will try to avoid GOV, MIL, US, EDU and ORG e-mail addresses and the addresses of users who reside in states and countries with strict anti-spam legislation.

10. The Nigerian scam spam has certainly increased from its overall tally of 5% of all spam. It was the most prolific spam message received by the honeypot accounts in this research, crossing many domain (ISP) boundaries and baiting methods.
11. Opting-out does work if the user opts-out of a service that they had initiated e.g. subscribing to a newsletter or requesting special offers. In this research, three out of six requests to opt-out of future mailings were honoured quite promptly, while two required a second request before mailings stopped. One service did not honour the request and continued to send e-mail, which was then categorised as spam.
12. Opting-out does not work however, when the unsubscribe link is used in a spam message. Only one link in 12 spam e-mails actually appeared to work.
13. Generally, spammers do not use profiling to target audiences. Each of the honeypot accounts had accessible profiles, but these profiles were not used. For example, it did not matter whether the username and user details represented gender correctly, as spammers did not use this information.
14. The software used by many spammers to create their lists exploits the services of providers such as Google and Yahoo.
15. Off-the-self targeting software does not work. Even though all of the available programs in this category were not tested, the experiments in this research concluded that the lists created by these software packages created lists that contained a large proportion of defunct e-mail addresses.
16. Keyword search programs that are designed for targeting (i.e. you enter a topic and the program harvests addresses based on the results of the web search) can appear to be successful in creating targeted lists. But these programs can quickly lose focus on the desired topic by selecting irrelevant links to harvest, and as a consequence gather unrelated e-mail addresses.
17. Spammers have some work to do in order to fine-tune their spam targeting techniques. But due to the increased returns for targeting campaigns, there will be a greater effort to find strategies that work.

18. 2003 will be a pivotal year in the legislative response to the spam onslaught, with Europe enacting opt-in laws and America following close behind with stringent anti-spam laws before Congress.
19. There is no silver bullet to the spam problem. The solution is a combined approach of legislation, education and spam filtering technologies in order to contain the proliferation of spam e-mail.

8.3 Future Work

Spam mutates as the level of spam filter sophistication increases. It seems that better filters spawn better spamming techniques. How do you counteract this phenomenon? In the course of this research, I read about Apache James (Java Apache Mail Enterprise Server), which is a 100% pure Java (open source) SMTP and POP3 mail server. Due to the fact that it is an open source e-mail server, it is possible to modify and extend it to a great degree. There has been some discussion in the development community about building extensions to James that allows it to perform as a spam honeypot. In fact, I am one of the authors of a paper that uses James to perform this function [31]. A variation on this theme is to write an application (e.g. a Java API) that captures and records real-time statistics on incoming spam. One of the suggested ideas in a discussion group is to use the application to dynamically generate new spam filters as new sources of spam are discovered. I believe that there is tremendous potential with this concept, as it appears that anti-spam techniques are always one step behind current spamming trends. In my opinion, spam filters should be dynamic and mutate just like the spam that they are trying to stop. Of course, there would have to be some type of collaborative effort with this approach, as it is of no use if only one mail server has the information on a new form of spam message. There could be a peer-to-peer solution with e-mail servers sharing this information or it could be a centralised solution with a central repository of new spam sources.

8.4 Summation

Methods to stop spam from reaching your inbox were presented in this paper but currently there is no silver bullet solution to the problem. The ultimate solution will need to have 0.0 false positives and even though this is a lofty and some may say unattainable goal, it is of utmost importance that not one single legitimate e-mail messages is blocked by an anti-spam mechanism. It is preferred to receive false negatives in the filter tuning process, rather than miss any messages that should have been received.

Anti-spam software will soon become as popular as anti-virus software, with most of the PC manufacturers shipping it as a pre-installed option. In June 2003, Hewlett Packard (HP) started to ship anti-spam software bundled with its new desktop computers. Junk e-mail filters are becoming more intelligent with ISPs promising better protection for its member's inboxes. Their filters can create new rules based on what the user has previously marked as spam in the hope that the self-learning feature of filters is the key to solving the spam problem.

For this dissertation I researched enough information to write a book on spam, but it is obvious that more research is required to offset the tenacity of the spamming community. Due to the fact that e-mail marketing has a proven positive ROI (return on investment) and provides the added benefit of rapid time-to-market, bulk e-mailing will continue to be used as a marketing strategy. Whether or not this type of e-mail constitutes spam, should be the recipient's decision. A spam blocking system should be used to filter out "obvious" spam such as pornography and chain mail e-mails, but the remaining spam should be delivered to a bulk e-mail folder where the end-user can categorise the e-mail. It would be valuable if the filter was self-learning and used the recipient's categorisation of spam e-mail to block further spam messages.

I believe that the decisive solution to the spam problem involves a combined approach of many anti-spam technologies. Filters, blacklist, whitelists and other methodologies all have a role to play in winning the battle against the proliferation of spam. It needs to be said though that a multi-faceted approach should be solidified by legislative force. Weak laws will not suppress spam. However, it is crucial that the Internet retains the ability to be an open medium; therefore a legislative response should not stymie legitimate e-mail marketing but encourage the adherence to a set of unambiguous ground rules.

References

[1] Research firm eMarketer, http://www.eMarketer.com, 2003.

[2] Brightmail Inc., Press release, http://www.brightlight.com/pressreleases/070103_uk_spam_summit.html, July 2003.

[3] Joyce Graff and Maurene Grey, Gartner Research, http://www.itap.purdue.edu/presentations/2003/presforum04242003.pdf, 2003.

[4] P. Steiner, “New Yorker Magazine”, p.61, July 1993.

[5] Sam Vaknin, PhD. “The Economics of Spam”. United Press International, 2002.

[6] Definition of spam: http://www.mail-abuse.org/standard.html.

[7] The American Heritage Dictionary of the English Language, Fourth Edition, The Houghton Mifflin Company, 2000.

[8] Definition of spam: http://www.monkeys.com/spam-defined/.

[9] http://www.brightmail.com/spamstats.html, 2003.

[10] Brightmail Press release, http://www.brightmail.com/pressreleases/122302_holiday_spam_alert.html, December 2002.

[11] The Radicati Group, Inc., Anti-Virus, Anti-Spam and Content Filtering Market Trends 2002-2006, 2002.

[12] Newsgroup posting, http://www.google.com/groups?&as_umsgid=8cd9d315.0305032212.4aed5d80%40posting%2Egoogle%2Ecom, May 2003.

[13] SpamCon Foundation, http://www.spamcon.org/.

[14] The Spamhaus Project, London, UK, http://www.spamhaus.org.

[15] The Spamhaus Project, http://www.spamhaus.org/newsdog.lasso?article=117.

[16] Newspaper Article, http://www.miami.com/mld/miamiherald/news/columnists/fred_grimm/5964205.htm, May 2003.

[17] Jupitermedia Corporation, http://cyberatlas.internet.com/markets/advertising/print/0,,5941_356791,00.html, May 2000.

[18] Newspaper Article, http://www.oregonlive.com/business/oregonian/index.ssf?/base/business/105256787116000.xml, May 2003.

[19] Blackbox Hosting Website, http://www.blackboxhosting.com/box/faq.asp, 2003.

[20] CNN article, http://www.cnn.com/2003/TECH/internet/07/14/porn.backdoor.reut/index.html, July 2003.

[21] G. Linberg, Network Working Group. Request for comment: 2505. Best Current Practice, 1999.

[22] The Spamhaus Project, London, UK. http://www.spamhaus.org.

[23] Sharon Gaudin and Suzanne Gaspar. “The Spam Police”, http://www.nwfusion.com/research/2001/0910feat.html, 2001.

[24] Patrick Pantel and Dekang Lin. “SpamCop – A Spam Classification Organisation Program”, Proceedings of AAAI-98 Workshop on Learning for Text Categorisation, 1998.

[25] Paul Graham, “A Plan for Spam”, http://paulgraham.com/spam.html, August 2002.

[26] H. Drucker, D. Wu, N. Vapnik, “Support Vector Machines for Spam Categorizations”, *IEEE Transactions on Neural Networks*, Vol. 10, No. 5, September 1999.

[27] Li Cheng and Wang Weinong, “Internet Mail Transfer and Check System Based on Intelligence Mobile Agents”, Proceedings of the 2002 Symposium on Applications and the Internet, 2002.

[28] Center for Democracy and Technology, “Why am I getting all this spam?”, March 2003.

[29] Symantec Corporation, http://www.symantec.com/press/2003/n030609a.html, June 2003.

[30] Website, http://www.zvon.org/tmRFC/RFC3098/Output/chapter7.html, May 1999.

[31] J. Seigneur, A. Lambert, P. Argyroudis, C. Jensen, Y. Chen, E. Gray, “PR3 Email Honeypot”, June 2003.