

SPAM ON THE INTERNET:
Is it here to stay or can it be eradicated?

Sara de Freitas and Mark Levene

SPAM ON THE INTERNET:	1
ABSTRACT	3
INTRODUCTION:	3
PART ONE: THE PROGRESS OF SPAM	4
PART TWO: DEFINING THE DIFFERENT TYPES OF UNSOLICITED BULK E-MAIL	7
PART THREE: COMBATING SPAM	7
CONCLUSIONS: IS SPAM HERE TO STAY?	16
ACKNOWLEDGEMENTS.....	17
REFERENCES	17
APPENDIX A: CLASSIFYING SPAM	21

ABSTRACT

This report outlines the growing problem of spam (unsolicited bulk e-mail), which has become a pervasive problem for Internet activity and has important implications for further and higher education institutions. The report provides a brief history of the development of spam, an explanation of how to define the different types of spam and an overview of technological and social ways of combating spam. The report provides a starting point for understanding the scale of the problem and begins a consideration of what further and higher education institutions can do to readdress the pervasive problem of unsolicited bulk e-mail.

INTRODUCTION:

At a recent address to the World Economic Forum, Microsoft's Bill Gates promised that 'spam will soon be a thing of the past' (Weber, 2004). Although others do not share Bill Gates's optimism (Arthur, 2004), his hopes to completely eradicate all spam by 2006 reflect a growing impatience with the problem of escalating amounts of unsolicited bulk e-mail.

Unsolicited bulk e-mail – or spam as it is popularly called – currently accounts for 63 per cent of all received e-mail in March 2004 (Brightmail, 2004; Salem, 2004). Of the 70 million e-mails that Brightmail filtered in September 2003 alone 54 per cent was unsolicited mail and that percentage is increasing year on year. In addition, Shinya Akamine, chief executive of Postini Inc., a US spam-filtering company, told a recent US Congress hearing that she believes spam has grown from 78 per cent to 83 per cent of all e-mail traffic this year (Krim, 2004). But although Bill Gates's plan to use a combination of different ways of filtering e-mail may lead to a significant reduction of spam in the short term, many are concerned that spam will never be completely eradicated (Hypönnen, 2004; Linford, 2004).

Spam has increasingly become a problem for all sectors of industry and education since the development of the World Wide Web and the increased use of e-mail for business and education (Salem, 2004). A series of attempts, both technological and non-technological, have been made to try to combat the increasing problems of congested mailboxes and to counter the heavy weight of unwanted e-mail traffic, which will have a strong effect upon the overall performance of the Internet. This has obvious implications for further and higher education in terms of the priority of maintaining institution-wide systems that are being used to support administrative tasks, and are increasingly being used for the delivery of learning materials and to support online communities of learners (de Freitas and Roberts, 2004).

In order to more fully consider the possible solutions to the spam problem, this paper will provide: a brief overview of the development of spam from the earliest direct marketing of Charles Ponzi to the modern day spammers, and a consideration of the different types and examples of spam. We will also consider the scale of the problem

and provide a technological review of the current methods being used to filter, track and block spam. We will also consider some current and future non-technological solutions including legislation, financial penalties and collaborative systems. We will conclude with some observations about the possible future of spam.

While the authors note a paucity of academic literature, with a particular need for literature on spamming written from social scientific perspectives, the report therefore draws upon interviews with noted experts in the field as well as sourcing from a wide range of technical and journalistic reports and articles.

Notably there is an increasing issue of spam affecting instant messaging and texting services delivered to mobile devices (Syntegra, 2003). This paper primarily concerns the use of e-mail for spamming due to the large scale of the problem.

PART ONE: THE PROGRESS OF SPAM

This section will provide a brief overview of spam from the earliest direct marketing in 1919 until today when the increasing use of spam is creating potential problems that may even lead to a collapse of the e-mail system (Hypönnen, 2004).

In 1919, Charles Ponzi developed the direct marketing pyramid scheme allowing investors to 'double their money in 90 days', a claim that resulted in mass investment into the scheme, making it an overnight success. Ponzi had promised money back fast but these promises turned out to be unfounded - and within four months many people had lost a lot of money and Ponzi was imprisoned. Direct marketing using the postal system was established and its potential applications were just beginning. While Ponzi was the first to use the postal service for direct marketing, the uses of spamming have evolved since then but remain mainly commercial.


With the advent of the electronic mail system in the 1970s a new opportunity for direct marketing using unsolicited electronic mail became apparent. In 1978, Gary Thuerk compiled a list of those on the Arpanet and then sent out a huge mailout publicising Digital Equipment Corporation (DEC - now Compaq) systems. The reaction from the Defense Communications Agency (DCA) - who ran Arpanet - was very negative and it was this negative reaction that ensured that it was a long time before unsolicited bulk e-mail was used again (Templeton, 2003). As long as the U.S. Government controlled a major part of the backbone, most forms of commercial activity were forbidden (Hayes, 2003). However, in 1993 the Internet Network Information Center was privatised, and with no central government controls, spam, as it is now called, came into wider use.

The term 'spam' was taken from the Monty Python Flying Circus (a British comedy team) and their comedy skit that featured the ironic 'spam song'. The purpose of the sketch is to say that 'spam is something you get whether you order it or not, and

eventually the noise of 'spam' will drown out everything else' (Viatel, 2004, p. 3). Conversely, where 'spam' came to mean unsolicited e-mail, the term 'ham' has come to mean e-mail that *is* wanted. Brad Templeton, a UseNet Pioneer and chair of the Electronic Frontier Foundation, has traced the first usage of the term 'spam' back to MUDs (Multi User Dungeons) - or real time multi-person shared environment - and the MUD community. These groups introduced the term spam to the early chat rooms (Internet Relay Chats).

The first major UseNet (the world's largest online conferencing system) spam was sent in January 1994 - and was a religious posting: 'Global alert for all: Jesus is coming soon'. The term spam was more broadly popularised in April 1994 when two lawyers Canter and Siegel from Arizona posted their message advertising their information and legal services for immigrants applying for the US Green Card scheme. The message was posted to every newsgroup on UseNet, and after this incident the term spam became synonymous with junk - or unsolicited - e-mail. Spam spread quickly amongst the UseNet groups who were an easy target for spammers simply because the e-mail addresses of members were widely available (Templeton, 2003).

More recently spam has been spreading at an increasingly rapid rate, and while groups of spammers were relatively small in the past, the wide availability of 'spam kits' over the Internet (which include mailing lists and detailed instructions on how to set up a spam outfit) has spread the practice from the United States to China, Russia and South America (Thomson, 2003; Linford, 2004). Since 2000 the threat from ever-increasing volumes of spam, the spread of viruses through spamming and an increasing number of those spamming has contributed to significant increases of e-mail traffic and an increasing problem for IT Systems Groups everywhere. The scale of the problem is perhaps best highlighted when we consider the growth of spam since 2001, when the percentage of spam, according to Brightmail Inc., was 7 per cent of all received e-mail. By 2002 this had grown to 29 per cent, and by the end of 2003 the total stood at 54 per cent (Salem, 2004). In March 2004 the percentage had increased to 63 per cent and this is set to rise considerably higher. See Graph 1, below.



Graph 1: The escalation of spam worldwide, 2001-February 2004. Source: Brightmail.

This situation is obviously having a damaging impact on industry, e-commerce and education (Salem, 2004). For example, in tertiary education spamming is becoming a significant impediment, blocking up mailboxes and spreading viruses. A straw poll of the colleges of the University of London demonstrates wide disruption to tertiary education institutions caused by spamming, due to time spent by users dealing with the excessive load of unwanted e-mail and by systems teams having to provide software solutions to deal with the problem. In addition, network performance is affected by overloaded systems dealing with the e-mails. The scale of the problem is perhaps best demonstrated by example, from a survey of external e-mail received over seven days at Birkbeck, University of London. We estimate that identified spam picked up by the spam filter made up 30 per cent of e-mail (data collected from Central Computing Services Systems Team at Birkbeck, University of London, February 2004). While at University College London, systems staff estimate that spam accounts for around 40 per cent of all external e-mail (data collected from the Systems Manager at the University College, London).

The two biggest worries for IT managers concerning e-mail are viruses and spam. Until recently they were separate issues but during June 2003 a new sinister virus called SoBig was released (Stewart, 2003). This virus is used to install anonymous proxy servers between the network and the user on infected computers enabling spammers to send e-mail through these hidden servers without fear of their IP address being detected. This approach to using viruses to spread spam is extremely worrying and points to the importance of combining the wars against spam and viruses.

PART TWO: DEFINING THE DIFFERENT TYPES OF UNSOLICITED BULK E-MAIL

So who are the spammers? While notably spammers are invariably unknown to the recipients, the spammers can be divided into three main groups: direct marketers who want to make commercial gain from spamming; criminal groups who are using spamming to 'legitimise' their activities (Linford, 2004; Gleick 2004); and disaffected individuals who want to disrupt Internet services and who in many cases may have inside information about how the systems are structured.

So what kinds of spam are there and how can they be classified? We can look at the classification of spam in two ways, first, in terms of the intention of the spammers (Schwartz and Garfinkel, 1998). This outlines a classification which includes the following categories: unsolicited commercial e-mail, make money fast messages, reputation attacks, UseNet spam, fraudulent activity (scams) and excessive multipostings. The second way of identifying spam is in terms of subject matter. Classification based on subject matter seems to be a more effective way to identify actual spam, although there is some overlap between the two classification systems. This section therefore explores the various types of unsolicited bulk e-mail according to subject matter, giving examples of them so that they can be identified. While different classifications of spam are widely used, one helpful content-based classification system is used by Brightmail (2004) and includes the following groups of spam (see also Appendix A):

- Adult
- Financial
- Products
- Internet
- Spiritual
- Scams
- Leisure
- Health
- Other.

PART THREE: COMBATING SPAM

Some predict that the spam problem will get worse, at least in the short term. Hand in hand with the push for tighter legislation to tackle the problem, several technical solutions have been deployed and new ones are being proposed. Here we will present a review of the current technological efforts to combat spam, and also include an indication of how the present solutions including legislative ones are evolving and may further develop.

Before an e-mail arrives in your mailbox it passes through a mail server, which is either hosted within your organisation or through an Internet Service Provider (ISP). Filtering out spam at this early stage (pre-receipt) before the message arrives at your

machine is obviously desirable and many IT departments and ISPs have already installed anti-spam software on their servers. Tools also exist which are user-based and filter out e-mail that has already arrived at your mailbox (post-receipt). Due to the flood of spam that is relentlessly sent to us, for now, it is probably best to have filtering tools both at the server and the user ends.

Two problems, which need to be addressed by any spam filtering system, are the rates of *false positives* and *false negatives*. A false positive is a mail message that the filter tags as spam but is actually ham, while a false negative is a mail message that the filter tags as ham but is actually spam. Having no filter at all is the case of 0 per cent false positives and 100 per cent false negatives, and a filter that blocks everything is one with 100 per cent false positives and 0 per cent false negatives. Ideally we want 0 per cent false positive, i.e. all ham gets through the filter, and 0 per cent false negatives, i.e. all spam is blocked.

In reality users will tolerate a certain level of classification errors, although some would argue that the only acceptable level of false positives is zero. It is important in this respect that e-mail that has been marked as spam is available for user (or Systems Manager) inspection in a personal (system) spam folder. This way, if a user (Systems Manager) detects a false positive she can add it to her *safe list* (e.g. safe senders list on Microsoft Outlook) of e-mail addresses she has authenticated as valid to receive e-mail from, and the spam filter can then take this information into account.

Whatever the technical solution for filtering spam, it must take into account the fact that spammers will fight back and find new ways of fooling anti-spam software. The implications of this are twofold. On the one hand, technical solutions need to be *adaptive*, i.e. modifying their internal behaviour to tackle new types of spam messages. On the other hand, it is important to pursue the legal route in parallel to technical solutions, in order to stop known mass spammers. Recently, Bill Gates issued an open e-mail in which he stated that Microsoft is significantly stepping up their efforts to fight spam both on the technological and policy-making fronts. He emphasised the use of machine learning techniques in building anti-spam tools that are easy to use, precise and adaptable (Gates, 2003).

We will now discuss some of the technical solutions that are currently being used and developed to filter out spam.

BLOCK LISTING

Block lists contain Internet Protocol (IP) addresses of known sources of spam, and are used for blocking incoming e-mail from these addresses before reaching the user. According to Spamhaus (see: Spamhaus, 2004) 90 per cent of all the spam users receive in North America and Europe can be traced to a hard-core group of fewer

than 200 spam outfits, all of which are operating illegally. Spamhaus maintains a block list, the Spamhaus Block List (SBL), which contains a current list of verified IP addresses of spam sources. SBL can be queried, free of charge, by mail servers wishing to block e-mail from these sources. SBL's primary objective is to avoid false positives, and from SBL such mistakes are extremely rare. The UK's Education and Research Network Association (UKERNA) subscribes to the Realtime Blackhole List (RBL), another block list that is generated by the Mail Abuse Prevention System (MAPS) (See, Janet-Cert 2004).

PROTOCOL CHANGE

The Anti-Spam Research Group (ASRG), which is a subgroup of the Internet Research Task Force (IRTF, 2004), investigates tools and techniques to mitigate the effects of spam. Its main focus is on technical solutions and providing input to the standardisation efforts of the IRTF. One of the important issues the group is looking into is to propose improvements to the current standard for sending and receiving e-mail, Simple Mail Transfer Protocol (Postel, 1982). The problem with SMTP is that it has no safe-guards to prevent forging or 'spoofing' e-mail addresses.

One proposal is to modify the Domain Name System (DNS) in order to be able to identify the actual computers acting as mail servers rather than just the website the e-mail came from. Another proposal, to verify the sender of an e-mail, called domain keys, is to use public key cryptography to sign an e-mail before it is sent and then verify its source once it arrives. To enable this feature, backed by Yahoo, e-mail servers will have to install open-source software, causing debate about who is to take ownership of e-mail technology standards. Yet another proposal that is gaining momentum is called Sender Policy Framework (SPF) (see: Sender Policy Framework 2004), which is a safe listing system requiring domain owners to publish the IP addresses from where e-mail are sent. When an e-mail arrives at the server the IP address of the sender must match the published IP address for the domain mentioned in the e-mail, otherwise the e-mail is rejected before it arrives in the user's mailbox. These suggested patches to the SMTP protocol will not stop spam but will help anti-spam technologies to track its origin, forcing the offenders to move to new domains more frequently.

ECONOMIC SOLUTIONS

The underlying idea behind all economic solutions is to make spammers pay for each unsolicited bulk e-mail they send, deeming spam a financially unviable proposition. One straightforward idea is to allocate users a reasonable e-mail quota of e-mails they are allowed to send and to charge a fixed rate on all e-mails above the quota. One of the challenges in implementing such a system is the mechanism by which such micro-payments are collected. The profit from such a scheme could, for example, be donated to charity, but this is problematic as there is a non-trivial cost in managing and transferring these funds. One strong argument for an economic

solution to spam is that “money talks” providing the strongest deterrent for sending junk e-mail (Arrison, 2004).

A refinement of the micro-payment scheme is that the receiver of an e-mail charges the sender a fee for each e-mail (Fahlman, 2002). Users set the fee, individually, according to the worth they attach to sending them an e-mail. As with all other solutions the user keeps a safe list of people they will receive e-mail from at no charge. In the case that the sender is not on this list, the payment details must be agreed prior to the message going through. To implement this, system software for managing the payment will have to be plugged-in to our e-mail software, enabling the transfer of money to the recipient’s e-mail account. This could be done via e-stamps, which are digital tokens that represent the amount of money being transferred. The redemption of the e-stamp is optional, so that if the e-mail is not considered as being junk, the user will probably opt not to collect the fee.

COMPUTATIONAL SOLUTIONS

The basic idea behind computational solutions, in similarity to economic ones, is to make spammers pay for sending e-mail. Only this time, rather than a direct payment, the sender of an e-mail is required to perform a small calculation prior to sending an e-mail. As spammers send bulk e-mails regularly, it would be unfeasible for them to perform all the computations required by this proposal without heavy investment in hardware. This idea is being investigated by researchers at Microsoft’s Penny Black Project, which was inspired by the Penny Black stamp introduced in the UK in 1840 as the first prepaid stamp at the cost of one penny to the sender of a letter (Penny Black Project, 2004). The computation solution sent with an e-mail can be tagged to its header and verified when the e-mail arrives. It is important in this scheme that the verification of the computation can be performed at a tiny fraction of the computation itself, and this is where the intelligence of this mechanism lies. This idea can be combined with a safe list of trusted e-mails, so that if an e-mail arrives without the computational stamp, it can be checked against the safe list before it is rejected.

Rather than performing a machine computation, a different kind of problem solving can be required from the sender of an e-mail such as solving a Captcha. A Captcha (Completely Automated Public Turing tests to tell Computers and Humans Apart) is a program that can generate and grade tests that most humans can pass but current computer programs cannot pass, such as recognising an image with distorted text (Captcha Project 2004). One problem with current Captchas is that they cause problems to people who have impaired vision, so that new types of Captchas, for example, using sound rather than vision, are being investigated.

E-MAIL ALIASING

The basic idea behind e-mail aliasing is to set up a variety of e-mail aliases (alternative addresses for a single user receiving e-mail) in such a way that each alias

can be restricted to a different group of users. The way it works is as follows: the user sets up a number of aliases to her e-mail with a set of manually configured attributes describing the acceptance criteria for each alias. Attributes include: how long the e-mail is valid, how many messages can be received until it is invalidated and who is allowed to use the alias for sending messages. In the vast majority of cases e-mail aliases can be created automatically by a system, called *Re-mailer* (Gburzynski and Maitan, 2004), which is a server-based system that users can subscribe to. The only difference between aliases set up by the user and those created automatically by the system, are that attributes of automatically created aliases are determined from pre-defined personalised templates. Personalisation in this context means that the same alias is used in all correspondence with the same sender. For every e-mail received for the first time, through a *master alias*, which is the e-mail address that is widely published, say name@bbk.ac.uk, *Re-mailer* creates an alias that is personalised to the sender, called a *quick alias*, say name@alias.bbk.ac.uk. The first e-mail received from anyone is bounced back to the sender using the quick alias and the sender has a certain amount of time, usually 48 hours, to validate the e-mail with a correct response to a challenge, such as solving a Captcha. If this happens the e-mail is sent through as validated and the sender's e-mail authenticated, otherwise it is rejected. Senders are burdened with one extra e-mail, which is the price they have to pay for this mechanism to work. This proposal involves an extension to e-mail servers but is compatible with existing e-mail infrastructure.

SENDER WARRANTED E-MAIL

This anti-spam method involves sending a special header in e-mails that certifies that the mail is ham. The method commercialised by Habeas (see: Habeas 2004) is to insert a Haiku, an ancient Japanese poetic form, into the headers of e-mails sent from companies licensing their method. An example of such as header is "winter into spring". Since copyright and trademark law protect the headers, their use by unlicensed spammers is illegal and they can be prosecuted. To supplement this, Habeas is building a safe list of all users who have licensed their method and a block list of those who have abused the method.

The advantage of this approach is its simplicity, as it requires no additional software, nor a change in the e-mail protocol. Its main disadvantage is that, although Haiku are protected, it is not clear that this will deter spammers if this method is widely adopted. Moreover, a patent for this method, if granted, will hinder wide adoption due to overly centralised control by the licensing company.

COLLABORATIVE FILTERING

SpamNet marketed by Cloudmark (see: Cloudmark, 2004) is a community-based tool that is an add-on to e-mail software. The idea is that the community collaborates in real-time to fight spam. A copy of the message is sent by the tool to a central spam database whenever any user in the community blocks a spam e-mail. All members

share the contents of this database so that if the *same message* appears in someone else's mailbox it is automatically blocked. To take care of the problem of false positives, blocked mail is moved into a spam folder rather than being removed. To reduce the number of false positives, Cloudmark also apply a trust system that checks the credentials of users when they notify the community of a spam e-mail. Cloudmark currently have a community of over 900,000 members and have reported the current success rate of their system to be 90 per cent. A particular problem with the current version of Cloudmark's software, at least for HE/FE, is that it is only compatible with Microsoft's Outlook.

One complication with a simple-minded implementation of this collaborative approach is that spammers tend to make random changes in e-mails as they are sent out, so detecting exact matches between e-mails as a blocking mechanism is not sufficiently robust to tackle the problem. A solution to this is to apply a similarity-checking program, which is insensitive to small differences between e-mails. Another recent strain of junk e-mail that needs to be dealt with in this context is *scramblespam*, where most of the message consists of random characters, thus confusing pattern-matching anti-spam algorithms.

One obvious drawback of the collaborative approach is scalability as the social network continues to grow. An alternative approach to maintaining a database of spam e-mail, without the need for a community, is implemented by Brightmail (see: Brightmail 2004). Their technique is to utilise what they call the 'probe network', which has over two million decoy e-mail accounts that attract about 15 million spam e-mails per day, that feeds into their database of known spam.

RULE-BASED SOLUTIONS

Rule-based filters maintain a collection of patterns that can be matched against an incoming e-mail to decide if it is spam. Each rule produces a score, and if the total score for the message exceeds a threshold value then it is classified as spam and blocked. (For example, a simple rule could state that if the subject of the e-mail contains the word "money" then the score assigned to the e-mail is increased by 1.) The most well-known rule-based filter is SpamAssassin, which is based on fuzzy logic rules to give a confidence on the accuracy of a rule. It is estimated to have over 30 million users and is claimed to be up to 95 per cent accurate (Sergeant, 2003). The specification of rules is handcrafted but a genetic algorithm, which is a flexible machine learning technique, does the assignment of scores to rules. It is easy to add new rules, and to customise the weights, i.e. relative scores, and thresholds of existing rules for identifying spam. SpamAssassin supports several rule categories including: header, body and message structure rules.

SpamAssassin is widely used in higher education, as it is free, relatively easy to install, easy to configure and known to be successful in blocking a large percentage of spam. At the School of Computer Science and Information Systems in Birkbeck we have chosen to use SpamAssassin to filter e-mail before it arrives at the user's mailbox, and its rules currently block around 95 per cent of all incoming spam e-mail.

STATISTICAL SOLUTIONS

This type of solution is often implemented as a post-receipt system rather than a pre-receipt one, i.e. the spam filter only acts once the e-mail has arrived in the user's mailbox. It is not a deterrent as some other solutions are, in the sense that the spammer does not have to pay for sending junk, but if effective it will make spamming futile. The essence of the statistical method is to use Bayesian text classification to assign each e-mail message either to the spam category or the ham category. In order for this method to work it is necessary to have available a large corpus of spam e-mail, in order to build accurate statistical patterns for classification purposes.

The naïve Bayes approach, which is the one most commonly used due to its relative simplicity and effectiveness, simply counts the number of occurrences of all words in the body of the text so as to assign their probability of being present in a spam message. Assuming that the classification software is an add-on to the e-mail software on the user's machine, a statistical profile of the user's ham messages can also be computed from the messages in the user's inbox. When a new e-mail arrives in the user's inbox the Bayesian classifier will compute the probability of this message being spam or ham using the classifier's pre-computed probabilities. The classifier will then choose to label the e-mail with the category having the higher probability, and if this turns out to be spam then it can put it into a separate junk e-mail folder that the user can inspect, just in case it is a false positive. The strength of this approach is that the filter is adaptive, in the sense that it can re-compute the classifier's probabilities of spam and ham as new e-mails arrive and are classified. This is especially important when a user detects a false positive and moves the message out of the junk e-mail folder. Another advantage of this approach is that it can readily be refined to detect sub-classes of spam, such as adult and money categories, and also sub-classes of ham such as work and personal categories.

Microsoft has developed its own Bayesian filter, which is now bundled with the new version of Outlook that comes with Office 2003. It is based on a filter developed within Microsoft Research, which combines the statistical approach with a set of handcrafted rules (Sahami et al., 1998). In the current version each message is tested against more than half a million criteria that are used to score the message, and if the score is above a certain threshold it is considered to be spam. Microsoft has also recruited about 250,000 MSN members to manually classify their e-mail messages as

spam or ham, and this information is used to continuously train and improve the filter.

Apart from Microsoft's effort, several other Bayesian spam filters are being developed (see: Paul Graham, 2004) and we expect to see more in the near future. We mention Bogofilter (see: Bogofilter, 2004), which is a popular open-source Bayesian filter for non-windows platforms such as Linux (Altunergil, 2003). A recent introduction to machine learning methods used to filter spam, including Bayesian, neural networks and decision tree learning, which stresses how these can be integrated into freely available packages such as SpamAssassin, can be found in Massey et al. (2003). Using several techniques to filter spam has the advantage of being able to combine them to obtain greater accuracy at the expense of a more complex system.

Statistical methods are most effective when combined with other methods such as block listing. In order to fool Bayesian filters spammers are trying to make their messages look more like normal non-spam messages by adding 'innocent' text such as a topical news item or a paragraph from a book. Because statistical filters are adaptive, fooling them is extremely difficult, but if spammers can get feedback on which messages pass the filters' tests, then they could counter-adapt their messages to fool the filter (Graham-Cumming, 2003). In this context, one simple line of defence against spam is to insist that e-mails be written in plain text and not in HTML, which allows all forms trickery including the possibility of transmitting viruses.

LEGISLATIVE SOLUTIONS

Non-technical methods for preventing spam in the future include legislation and prosecution. Legislative methods have to date proved fairly ineffective due to the global extent of the activity, which has implications for how global law operates. One of the key impediments to this method of prevention therefore lies in the cases where spamming is being committed outside the legal jurisdiction. In other words, all the nations would have to agree to the same legislation in order for it to be enforceable. Another factor lies in how the law is enforced, which would need collaboration between the various countries as well as significant amounts of funding for enforcement.

Current social and political debates have centred upon the adoption of either using an opt-in or an opt-out approach. That is, should the public opt-in before any mail can be sent or should they opt-out whenever unsolicited bulk e-mail is received? So far legislation has largely included opt-out options – which can be likened to 'do not call' (Vaile, 2004; Hypönnen, 2004; Linford, 2004). The opt-out option aims to stop all fraudulent e-mail and unsolicited pornography but allows for legitimate business direct marketing e-mail. However, in the process, it does not make spam illegal, leaving a significant grey area in the legislation (Syntegra, 2003).

For example the US Can Spam Act of 2003 is an example of the opt-out mode of legislation which allows Americans to opt-out of receiving unsolicited bulk e-mail. However the law has been found to be weak as it does not ban junk e-mail outright (BBC News, 2004). Similarly, the recent UK Privacy and Electronic Communications Regulation legislation operates using an opt-out clause and similarly allows for direct e-mail marketing to businesses but not to consumers (BBC News, 2003). The only nation to adopt an opt-in approach to legislation is Italy and they make spamming a criminal activity (Linford, 2004). There are still many critics of the spamming legislation that say the opt-out clause significantly weakens the chance of any legislative effectiveness. In brief, legislative solutions seem to have significant impediments not least due to the global reach of the activity and the problems with enforcement.

Solution	Method	Benefits	Limitations
Block listing	Use of lists of IP addresses of known sources of spam (e.g., SBL and RBL)	Blocks a significant volume of spam	Cannot block all spam, and needs to be updated on a regular basis
Protocol change	To provide a method of tracking the source of an e-mail	Will help to identify spammers, and add spam addresses to block lists	Will not prevent spam as such
Economic solutions	Impose a fee for sending e-mail	Will deter spammers from sending large volumes of junk e-mail	Will be difficult and costly to implement a world-wide standard for collecting the fee
Computational solutions	Impose an indirect payment in the form of a machine computation prior to sending e-mail	It is a viable alternative to the economic solution, without needing the infrastructure to collect a fee	A protocol involving cryptographic techniques will need to be put in place, and software developed to implement the method
E-mail aliasing	Set up e-mail aliases for different groups of people with different acceptance criteria	Will reduce spam through an authentication process	This method involves an extension to current e-mail servers, and the management of e-mail aliases
Sender warranted e-mail	Use of a special header to certify the e-mail as valid	No need for additional software or e-mail protocol	Will probably not deter spammers if widely adopted, and wide licensing of the technology will be problematic

Collaborative filtering	Communities collaborate to fight spam using a collaborative tool that is an add-on to e-mail software	Possible eradication of large volumes of spam through collaborative reporting of spam	Still vulnerable to random changes in spam e-mail, and there are problems with scalability of this method
Rule-based solutions	These filters maintain a collection of patterns to be matched against incoming spam, as in SpamAssassin	It is easy to install and effective in blocking a large percentage of spam, and in the case of SpamAssassin is free	It needs a lot of tuning, and should be combined with other methods to filter out a larger volume of spam
Statistical solutions	Often deployed as a post-receipt spam filter using Bayesian text classification to tag e-mail as spam or ham	It is very effective and is also adaptive, so hard to fool	Most effective when used with other pre-receipt filter systems
Legislative solutions	National and global legislation to enforce anti-spam laws	Prosecution of individual spammers	Problems of enforcement, not least due to crossing of different jurisdictional boundaries

Table 1: Spam and methods of prevention

CONCLUSIONS: IS SPAM HERE TO STAY?

In the course of this study we have found that there was a dearth of guidance for further and higher education on the subject of how to cope with spam. Although spam has been around for several years, it is only recently that institutions have begun fighting back by dealing with it at an organisational level. However it is clear that more guidance and support for academic institutions would be helpful, particularly in terms of sharing anti-spamming good practice. This report could provide a starting point for a wider debate on spamming within the further and higher education sector.

Although in theory the best way to combat spam may be through a system of charging, it seems that in practice combined efforts that use block lists together with pre- and post-receipt filtering systems may be the most effective approach at present. We found that the institutions that we surveyed favoured this technical solution, while we await for other social, technological and legislative methods of blocking spam to become more established (see Table 1 for a summary of the methods of solutions for combating spam).

In addition, simple steps can be taken to reduce the amount of spam that we receive. It is already well-known that you should not publish your e-mail address on the Internet in a form that spammers can harvest easily; for example, on your web page replace `j.bloggs@uni.ac.uk` with `j(.)bloggs(at)uni(.)ac(.)uk`. The `mailto` function in HTML is another example where an e-mail is visible to spammers, since it is easily parsed and collected by software used by spammers. Enrique Salem (CEO of Brightmail) (2004) recommends that learners have two e-mail addresses one for

personal correspondence and one for shopping on the Internet. Certainly your e-mail address should be given out with some caution, particularly when deciding whether to opt in or out of marketing e-mail.

The question about whether spamming can be completely eradicated cannot be answered with any certainty at this time. Perhaps reassuringly, Bill Gates and Enrique Salem argue that spamming *will* be eradicated in the next few years through the solutions we detailed in Part Three. However, it seems unlikely that spam will completely disappear, although with the current force behind the anti-spam movement gaining momentum, we can expect to see less spam, but only with preventative measures such as those described in this report being put in place. In the near future however, the 'cat and mouse' game between spammers and anti-spammers is set to continue.

ACKNOWLEDGEMENTS

Thanks to the University College London Systems Team, the Birkbeck Central Computing Services Systems Team and to Rajesh Pampapathi, School of Computer Science and Information Systems, Birkbeck, for providing us with the screen shots and statistics.

REFERENCES

ALTUNERGIL, O., 2003. *Bayesian Filtering with bogofilter and sypheed claw*, January 2003.

See: www.linuxdevcenter.com/pub/a/linux/synd/2003/01/30/bogofilter.html. Last accessed 24th February 2004.

ARTHUR, C., 2004. *World Economic Forum: Gates aims to wipe out spam as UK broadband users unwittingly help the spammers*. The Independent, p. 22. January 26th.

ARRISON, S., 2004. Canning spam: An economic solution to unwanted e-mail. Pacific Research Institute.

See: www.pacificresearch.org/pub/sab/techno/2004/spam01-26-04.pdf. Last accessed 24th February 2004.

BBC NEWS., 2004. *US anti-spam law fails to bite*. 9th February 2004.

See: <http://news.bbc.co.uk/2/low/technology/3465307.stm>. Last accessed 25th February 2004.

BBC NEWS., 2003. *Top UK sites 'fail privacy test*. 11th December 2003.

See: <http://news.bbc.co.uk/2/hi/technology/3307705.stm>. Last accessed 25th February 2004.

BOGOFILTER. 2004. Bogofilter website.

See: <http://bogofilter.sourceforge.net/>. Last accessed 4th March 2004.

BRIGHTMAIL., 2004. The Brightmail website.

See: www.brightmail.com/spamstats.html. Last accessed 3rd February 2004.

CAPTCHA., 2004. The Captcha Project.

See: <http://www.captcha.net/>. Last accessed 24th February 2004.

CLOUDMARK., 2004. Cloudmark website.

see: www.cloudmark.com. Last accessed 4th March 2004.

DE FREITAS, S. AND ROBERTS, G., 2004. *Does distance e-learning work?* Association for Learning Technology Journal, 11(3), pp. 69-87 Cardiff: University of Wales Press.

FAHLMAN, S.E., 2002. *Selling interrupt rights: A way to control unwanted e-mail and telephone calls*, Technical Forum, IBM Systems Journal, Vol. 41, pp. 759-766.

GATES, B., 2003. *Towards a spam free future*, Bill Gates, 24 June 2003, Executive E-mail.

See: www.microsoft.com/mscorp/execmail/2003/06-24antispam-print.asp. Last accessed 24th February 2004.

GBURZYNSKI, P. AND MAITAN, J., 2004. *Fighting the spam wars; A re-mailer approach with restrictive aliasing*, in ACM Transactions on Internet Technology, Volume 4, pp. 1-30.

GLEICK, J., 2003. *Get out of my box*. In the Guardian Review, 2nd March 2003, pp.1-2.

GRAHAM-CUMMING, J., 2003. *Fooling and poisoning adaptive spam filters*, Sophos White Paper, November 2003.

HABEAS., 2004. Habeas website. See: www.habeas.com. Last accessed 4th March 2004.

HAYES, B., 2003. *Spam, spam, spam, lovely spam.* American Scientist, vol. 91, no. 3, pp. 200-204.

HYPÖNNEN, M., 2004. Interview with Mikko Hypönen (F-Secure) conducted on 27th February 2004.

INTERNET RESEARCH TASK FORCE, 2004. See: <http://www.irtf.org/charters/asrg.html>. Last accessed 26th April 2004.

JANET-CERT., 2004. Janet-Cert website.

See: www.ja.net/CERT/JANET-CERT/mail/mail-abuse/rbl-plus-guide.html. Last accessed 5th March 2004.

KRIM, J., 2004. *Senate hears mixed reviews of Anti-spam law*. Washington Post [online], 21st May. See: <http://www.washingtonpost.com/wp-dyn/articles/A43622-2004May20.html>, last accessed 19th June 2004.

LINFORD, S., 2004. Interview with Steve Linford (Spamhaus) conducted on 9th January 2004.

MASSEY, B., THOMURE, M. BUDREVICH, R. AND LONG, S., 2003. *Learning spam: Simple techniques for freely-available software*. USENIX Annual Technical Conference, pp. 63-76, San Antonio, Texas, June 2003.

PAUL GRAHAM., 2004. Paul Graham website.

See: www.paulgraham.com/filters.html. Last accessed 4th March 2004.

PENNY BLACK PROJECT., 2004. The Penny Black Project.

See: <http://research.microsoft.com/research/sv/PennyBlack/>. Last accessed 25th February 2004.

POSTEL, J., 1982. *Simple Mail Transfer Protocol*. Information Sciences Institute. University of Southern California. See: <http://www.ietf.org/rfc/rfc0821.txt>. Last accessed 26th April 2004.

SAHAMI, M., DUMAIS, S., HECKERMAN D. AND HORVITZ E., 1998. *A Bayesian approach to filtering junk e-mail*, AAAI Workshop on Learning for Text Categorization, Madison, Wisconsin, July 1998.

SALEM, E., 2004. Interview with Enrique Salem (CEO of Brightmail) conducted on 23rd February 2004.

SCHWARTZ, A. AND GARFINKEL, S., 1998. Stopping spam. Beijing. Cambridge. O'Reilly & Associates.

SENDER POLICY FRAMEWORK., 2004. Sender Policy Framework website.

See: <http://spf.pobox.com/>. Last accessed 5th March 2004.

SERGEANT, M., 2003. *Internet level spam detection and SpamAssassin 2.50*, Cambridge MA, January 2003.

See: <http://axkit.org/docs/presentations/spam/>. Last accessed 19th June 2004.

SPAMHAUS., 2004. Spamhaus website.

See: www.spamhaus.org/sbl/. Last accessed 4th March 2004.

STEWART, J., 2003. *Spam and SoBig: arm in arm*. Network Security, pp. 12-16, October 2003.

SYNTEGRA., 2003. Can spam kill the mobile messaging market? White Paper. See: <http://www.us.syntegra.com/acrobat/208950.pdf>. Last accessed 26th April 2004.

TEMPLETON, B., 2003. Brad Templeton's Home Page.

See: www.templetons.com/brad/spam_react.html. Last Accessed 29th October 2003.

THOMSON, I., 2003. *Mafia muscles in on spam and viruses*. Vnunet.com. See: <http://www.vnunet.com/News/1151421>. Last accessed 26th April 2004.

VAILE, D., 2004. *Spam canned*, Internet Law Bulletin, Vol. 6, no. 9. See: http://www.bakercyberlawcentre.org/Articles/vaile_spam_ilb_6.9.pdf. INTLB 113.

VIATEL., 2004. *Spam: now a corporate concern*. Viatel White Paper [online]. Egham, Surrey: Black Spider Technologies. See: http://www.viatel.com/uploads/VIATEL-097750_spam.pdf, last accessed 20th June 2004.

WEBER, T., 2004. *Gates forecasts victory over spam*. BBC News website.

See: <http://news.bbc.co.uk/go/pr/fr/-/hi/business/3426367.stm>. Last accessed 3rd February 2004.

APPENDIX A: CLASSIFYING SPAM

Brightmail (2004) have developed a subject-classification system for use in their systems. Examples of these categories include the following:

Adult: contain services aimed at over 18s and include links to pornography, personal advertisements and relationship advice.

Subject: Hey! We have Heather Graham, Liz Hurley and many, many more featherweight
From: [Isaiah Tackett <hlyobkq5@excite.com.au>](mailto:Isaiah_Tackett<hlyobkq5@excite.com.au>)
Reply-To: [Isaiah Tackett <hlyobkq5@excite.com.au>](mailto:Isaiah_Tackett<hlyobkq5@excite.com.au>)
Date: 2004-02-21 12:12
To: [operator <operator@ccs.bbk.ac.uk>](mailto:operator@ccs.bbk.ac.uk)



Figure 1: Spam advertising adult material.

Financial: refer to money-related services such as investments and real estate loans.



Figure 2: Spam advertising financial services.

Products: offering or advertising commercial goods and services.



Figure 3: Spam advertising a product.

Internet: offering or advertising Internet or computer services including webhosting and spamware.

Spiritual: offering religious and spiritual services including astrology and psychic services.

Scams: fraudulent activity and include investment pyramid schemes and chain letters.

Subject: KIND ASSISTANCE
From: frdtaylor@survivormail.com
Reply-To: frdtaylor@survivormail.com
Date: 2003-09-26 16:51
To: cypheerpunks@minder.net

First may are solicit your confidentiality on this transaction,
by this virtue of its nature, I am Fredick taylor a cousin to Charles taylor the else president of Liberia,With the recent indictment of Charles Taylor by the international war crimes tribunal, he has mandated me to look for a reliable partner who will urgently assist in the collection of consignment of boxes (containing cash of \$25million dollars he kept in safe custody with a security management company in Accra Ghana and the security company is not aware of the contents of the consignment.The money was the proceed of Diamond Sales which the rebels group has been fighting to gain access which resulted in the Killing of innocent citizens and destruction of properties.Please if you can assist us in safeguarding and investing this money in any profitable business under your strict supervision, please send your reply to the above address and please also include your Telephone number for easy contact, I am presently living in Accra Ghana as a political asylum. We need !

Figure 3: Investment scam spam.

Leisure: offering prizes and other discounted activities and include vacation offers and online casinos.



TropicalTrips.net
Tropical Vacations for Couples or the Entire Family

Congratulations!
you have been selected to receive
Up TO 4 FREE

Tropical Florida Vacations!
Qualify and take one of our vacations and receive up to 4 FREE bonus vacations to several different destinations! Absolutely FREE!

If you qualify you will be contacted by telephone for address verification and complete vacation details.

Complete our Qualification Form

Destination 1: **5 Days/ 4 Nights** hotel accommodations in Orlando, FL home of **Disney, Epcot,** and the new **Universal Studios**

All fields are required Must be 18 to qualify

Email


First Name

Figure 4: Commercial spam: Selling vacations.

Health: health-related products and services and including pharmaceuticals and medical treatments.

From: Christian Read [LAJTTJQBKCW@level3.net]
To: nobody@ccs.bb.k.ac.uk; operator; p.eden@ccs.bb.k.ac.uk
Cc:
Subject: The Real Somc Deal on the Int.ernetmaden

Sent: Fri 27/02/2004 08:52



FREE Prescription
with **EVERY ORDER!**

- :: Safe, Secure Online Shopping
- :: No Long Waiting Lines
- :: Complete and Total Privacy
- :: US Licensed Pharmacies

WE MAKE IT EASIER AND FASTER THAN EVER!
Medications Prescribed and Shipped Overnight To You!

Welcome to the **LOWEST PRICES** Online!

AMBIEN New Low Price! **ORDER NOW**

FIORICET New Low Price! **ORDER NOW**

PHENTERMINE New Low Price! **ORDER NOW**

SOMA New Low Price! **ORDER NOW**

TRAMADOL New Low Price! **ORDER NOW**

VIAGRA New Low Price! **ORDER NOW**

Figure 5: Health related spam.

Other: e.g.: false virus alerts.

Subject: {Virus?} Information From: eveleva@ami.ru.acad.bg Date: 2004-02-25 06:14 To: advisory@ccs.bbk.ac.uk	Attachments:  VirusWarning.txt
---	---

Warning: This message has had one or more attachments removed
Warning: (about.zip).
Warning: Please read the "VirusWarning.txt" attachment(s) for more information.

Please see the attached file for details

This is a message from the MailScanner E-Mail Virus Protection Service

The original e-mail attachment "about.zip"
was believed to be infected by a virus and has been replaced by this warning
message.

If you wish to receive a copy of the *infected* attachment, please
e-mail helpdesk and include the whole of this message
in your request. Alternatively, you can call them, with
the contents of this message to hand when you call.

At Wed Feb 25 08:15:55 2004 the virus scanner said:
about.zip->about.exe Infection: W32/Mydoom.F@mm

Note to Help Desk: Look on the MailScanner post-office.ru.acad.bg in /var/spool/MailScanner/quarantine/20040225
(message i1P6Ftnk015374).

--
Postmaster
Mailscanner thanks transtec Computers for their support

Figure 6: False virus-alert spam.