The Dangerous Economics of Spam Control

Alena Kimakova[*] and Reza Rajabiun[†]

2008 MIT Spam Conference

Keywords

Spam, mechanism design, filtering, technological change

Abstract

The adoption of a wide range of regulatory and technical measures against spam has not constrained its growth and sophistication. This paper provides a novel explanation for this puzzle by emphasizing the technological trade-offs between the accuracy and speed of filters facing network providers in the early to mid-2000s. Furthermore, the paper documents how antispam software developers have responded to the technological gap.

I. Introduction

While volatile on a daily basis, on bad days the volume of content defined as undesirable by end users, or their filters, can reach 90% of total messages [1]. Estimates by TrustedSource, from a network of sensors in 68 countries show that in early 2008, 120 out of 160 billion daily messages were spam. Why do we see more spam after the adoption of countermeasures aiming to increase the costs incurred by spammers (blacklists, reputation systems, civil and criminal penalties) and receiver side mechanisms (checksum and content filters)? An understanding of this question is required for the design of mechanisms likely to reduce the costs posed by this form of noise on end users and communication network infrastructure. Without a good picture of this history, we may be doomed to adopt countermeasures that simply motivate spammers to produce more, rather than less, of their advertisements.

To explain the growth and sophistication of spam, the analysis integrates perspectives from economics and computer science on the interactions between spammers, and antispam mechanisms. [2] We extend insights from game theoretical models that attribute the growth of spam to asymmetries in the quality of filters across network providers. Existing studies however do not explain why differences in filter quality exist across different networks. This paper contributes to ongoing debates by emphasizing the role of technological trade-offs between the accuracy and speed of filters.

II. Technological choice

Advances in content classification systems since the late 1990s have been impressive, and constrain the sensory threat spam poses to email. [3,4]. Nonetheless, the high ratio of noise to signal poses significant costs on owners and operators of messaging network, since they are forced to allocate more resources to processing a given level of legitimate traffic. Network costs of spam are of particular concern to developing countries due to the relative scarcity of bandwidth, processing capacity, and specialized administrative skills [5]. Existing research usually focuses only on the expected accuracy of filters, and ignores the implications of spam on network providers. End user and network costs of spam are however closely linked. The same technologies that lower the price of communications often also empower some individuals and groups to create more spam. Spam and antispam are hence likely to co-evolve over time.

Close substitutes to email such as telephone calls, conferences, or instant messaging typically impose higher bandwidth and machine processing capacity requirements than email. They are also inferior instruments

[*] York University, Toronto, Canada (akimakov@yorku.ca)
[†] COMDOM Software and York University, Toronto, Canada (reza@comdomsoft.com)

for archiving and processing old communications. Non-email forms of spam, such as printed material or telemarketing have non-negligible costs. In this context, it is easy to see that substitution away from email due to the costs of spam, or other malware that comes with it, is not likely to be an efficient response to the growth and sophistication of spam since the early 2000s.

To capture the main elements of our argument, consider the following problem faced by an ISP. Let the costs of spam (C) to a semi-autonomous network provider over a specific time interval be defined as a combination of end user (E) and network costs (N). End user costs are a function of the accuracy of a filter. The expected false negative rate (E1) of the filter has costs for end users in terms of their scarce attention, as in Loder et al. (2004). [6] False positives (E2) have opportunity costs in terms of reliability of email.

Any message sent by a spammer has the probability (z) of reaching the inbox of an end user, and a probability 1-z of getting filtered out. The relationship between this figure and the expected response rate from the population of end users can be estimated through trial and error by spammers, which in itself tends to increase the volume of spam.

In addition, network costs require an accounting for the bandwidth, hardware, and administrative resources that ISPs must allocate to spam control. If we ignore bandwidth costs, the hardware and administrative requirements can be estimated based on the number of servers (S), which is inversely related to the speed of the scanning and filtering systems in place.

Finally, both network and end user costs of spam are dependent on the architecture of the antispam system chosen by network administrators. A centralized system for the identification and classification of email into ham and spam exhibits economies of scale. However, a centralized administrator may not have an accurate picture of what constitutes ham/spam for different sub-groups of individuals using the network infrastructure. Decentralization of content filtering to reduces Type I and II errors, but unfortunately increases network costs to the carriers, downstream ISPs, and ultimately end users.

To summarize, the network costs of spam facing our hypothetical network provider can be described by the following general functional form.

$$C = C\,(E\,(E1, E2), N\,(E1, E2, S))$$

There is little known about the nature of these relationships, and we should not assume that they are static. In practice, the parameters could be estimated for individual ISPs based on accounting information as well as the features of different antispam systems available at the time.

The optimal choice of filter for upstream ISPs may nevertheless differ from those of more specialized ISPs/ASPs. Upstream entities are likely to own and operate the network infrastructure, and hence are likely to be more sensitive to the speed of the antispam technology they chose than its accuracy. The popularity of centralized checksum and reputation based filters upstream reflects this factor. The fact that very accurate content filters are popular with end users that rely on email for business and personal communications also substantiates this hypothesis.

Within this framework, it is easy to see that improvements in accuracy are not sufficient to constrain the growth of spam. Upstream ISPs require countermeasures that are accurate, as well as fast. Given the presence of technological tradeoffs between filter accuracy and speed, choices by some sub-segment of networks may look suboptimal from the perspective of others downstream. Smaller ISPs and end users are then forced to buy and maintain their own filters, or outsource scanning and filtering of their communications to specialized firms.

III. Distribution of tastes

A theoretical analysis by Khong (2004) suggests that mechanisms that would allow spammers and receivers with a taste for spam

to connect, without causing problems for others, are superior to those that aim to impose costs on senders, or filters [7]. Filtering and blocking are viewed as second best solutions, since in the long run they encourage spammers to compensate for the losses in their response rates by producing more spam. For instance, opt-in mechanisms represent the first best instrument. While opt-in registries or lists, as well as digital gated communities, have become popular since 2004, the incentives of spammers to produce more email advertisements have become even more pronounced after the wide-spread adoption of such open channels.

Loder et al. (2004) offer a different interpretation of the Coasean approach to the analysis of decentralized economic conflicts. [6] Specifically, they propose an Attention Bond Mechanism (ABM) as the first best solution, relative to regulatory sanctions on senders or even a perfect Bayesian filter. Their mechanism aims to impose a decentralized price on spam messages. They argue that by asserting property rights, ABM is theoretically superior to regulation and filters because it accounts for a basic economic insight:

"In terms of individual and aggregate social welfare, a system that facilitates valuable exchange and side payments will generally dominate a system that grants only unilateral veto power to either party." (Loder et al., 2004, 3)

This view reflects one of the central assumptions of modern economic theory: The subjective theory of value. In the case of spam control, this perspective suggests that reasonable people are likely to disagree about what constitutes desirable and undesirable content. Most people find mass mailings about sexual enhancers wasteful or even offensive. Nonetheless, some subgroup of a population is open to that sort of advertisement, and apparently responds to it. Attempts to blocking spam by implementing a spam filter that excludes any messages containing "Viagra" is not necessarily efficient because it prevents exchange between buyers and sellers. Filters that do not account for the distribution of taste for underlying products or services, and instead block communications, motivate spammers to create noisy variations on the same theme as they search for their target audience. This problem exists because spammers do not necessarily know who wants the products ex ante, but can try to identify this audience through trial and error. The more sophisticated the filters in place, the higher the number of messages needed to achieve the same response rates.

For a more interesting application, consider the international dimensions of the pharmaceuticals trade. Because of variation in regimes for regulating the price of prescription drugs and intellectual property rights, there are significant differences between the prices of medicines in the United States and Canada. [8] Information about such price differentials was costly to obtain, particularly before the advent of the Internet. For the large sub-group of the U.S. population without good health insurance, the new communication technologies made it possible to see the differentials, generating a thriving Internet pharmacy business in Canada. While some potential buyers actively looked for alternative suppliers of their medicines by going to their home pages, active spamming campaigns have been an essential part of the arbitrage strategy of the Canadian suppliers. For people with good health insurance, these messages appear as spam, while for others they may be a lifeline to medicines they otherwise could not afford.

This example clearly shows the importance of the economic approach to the analysis of the spam problem. In particular, the taste for spam is not likely to be distributed normally across a population of end users. Instead, it seems more realistic to assume that taste for any particular class of advertisements can be characterized as one with a very long tail. Some value the specific content of a batch of spam very highly, but most of the population does not. The idea that decentralized mechanisms, like ABM or opt-ins, are more efficient than hierarchical systems for the identification and processing of email arises from the economic emphasis on the subjective theory of value.

It is pertinent to note that in natural sciences, long tailed distributions, as in the case of

consumer preferences for spam/ham, are associated with properties of scale free/complex networks. [9] This class of structures usually indicates the potential for phase transitions and the existence of multiple equilibria. In understanding spam, and designing robust mechanisms for mitigating its end user and network costs, this is an important point. It is possible that the current noise/signal ratio is not unique, and more effective combinations of countermeasures may be able to motivate a shift to a superior long term equilibrium level of spam.

In terms of microeconomic theory, the presence of long tailed distributions of information about the value of underlying assets is associated with a situation where markups are invariant to the number of sellers. [10] In the context of the spam problem, this insight means that the margin to spamming or expected response rates, are possibly invariant to the number of spammers at play. Consequently, the inherent distribution of preferences for spam explains why sender side cost mechanisms, like civil or criminal punishments, as well as advances in filtering, have not had the expected effects on the total volume of spam. Some risk averse spammers, and companies that paid them for their services, may have been driven out of the market. The reaction of the remaining firms has been to increase the quantities of spam they produce. In the language of game theory, this suggests the presence of strategic complementarities in the reactions of individual suppliers of spam, and countermeasures that make it more costly for them to reach their target audience. The presence of such complementarities is also associated with the possibility of multiple market equilibria in relation to prices and quantities. Wrong choices about filters by ISPs, or regulations by policy makers, may explain why we have now ended in such an inefficient state of affairs.

This inference may seem puzzling to proponents of countermeasures that assume imposing costs on spammers or receivers in terms of legal sanctions or IP reputation, for example, should thin out the market. [11,12] When the taste for spam does not follow a normal distribution across a population of end users, some spammers may stop production if they are faced with increasing costs. However, other spammers who remain in the market can chose to hide their identities and produce even more in response to the pricing scheme. Open channels and opt-in lists likely provide one element of the range of connections that will be necessary to reduce the incentives of businesses to advertise through spam.

IV. Strategic conflicts

Although spam is used to distribute malware, it primarily exists because it works as a cheap method for advertising goods and services. The most popular economic explanation for this phenomenon relies on the (in) famous tragedy of the commons. In this context, the overuse of a common resource is similar to the problems of depleting fisheries, or congesting public roads. In general this type of market failure arises because of the presence of too many legal rights to use a particular resource. In contrast to the monopoly problem, where too few usage rights lead to high prices and low quantities, market pressures in the commons result in excessively high quantities and low prices as some individuals and groups manage to externalize the costs of their actions. [13] The generic mechanism for solving the commons problem is to raise costs facing producers to account for the social costs of their actions.

Numerous regulatory and technical mechanisms have been proposed and adopted to raise the costs of sending spam in the late 1990s and early 2000s. These include ad hoc challenge response systems, blacklists, government regulations, as well as systems that try to identify and process spam based on the reputation of the physical domain or IP address of senders. [14] However, these countermeasures have not translated into lower volumes of spam. If anything, the tendency of spammers to send out large amounts of messages has been enhanced since the adoption of antispam regulations and other mechanisms that aim to reduce spam by increasing personal or computational costs of sending large volumes of untargeted advertisements. In the regulatory case, the

costs are raised by some expectation of criminal and civil sanctions. With blacklists/reputation based filters, the costs to spammers, or their robots, materialize in terms of future ability to send messages. An important policy question raised by the growth of spam relates to the desirability of adopting, or retaining, countermeasures that aim to increase costs of spam production.

The economic literature on crime and punishment, following Becker (1968, 1993) and Friedman (1984), provides an intuitive picture of the implication of this class of mechanisms. [15, 16, 17] This literature points out that the costs of enforcing any public norm, or tax, increases with the expected level of punishment. If expected punishments are too high, then economic agents will try to avoid them in more sophisticated and socially costly ways. As a result, the so-called "hang them all" enforcement strategies with heavy punishments are typically inefficient, and often counterproductive. The development and adoption of sophisticated cloaking techniques by spammers since the early 2000s seems to be a direct response to the increased expected punishments to senders of mass advertisements through email. To see this economic intuition, it is pertinent to recall that escalation of wars on drugs, prostitution, and other vices by imposing costs on private bargains often only exacerbate the social costs associated with the problem.

Androutsopoulos et al. (2005) describe the interactions between spammers and receivers within a two player adversary game, specifically in the case where all end user mailboxes have filters. [18] Based on this characterization, they argue that the spam game (almost) always has a single Nash equilibrium, and hence tends to settle in an infinitely repeated game. This conjecture implies that strategic interactions between spammers and receivers are likely to persist over time, unless there are changes to underlying technologies or the taste for spam.

Reshef and Solan (2006) provide another game theoretical model for the analysis of three classes of countermeasures. [11] These include filters of different qualities, sender side cost-raising mechanisms (authentication and reputation services, counter-attacks, and payments in monetary or computational terms), and a do-not-spam registry. They claim that when the cost of sending messages is not too high, the effect of improved filtering quality on the total volume of spam is ambiguous. On the other hand, they argue that when the costs of sending spam are high, improved filter quality reduces the total level of spam on the network. The key implication of their model is that mechanisms aiming to impose a cost on spammers behave as strategic complements to filters in fighting spam at the aggregate network level. In other words, the use of one class of instruments enhances the power of the other under their formulation, which is used to justify creating do-not-call registries as complements to filters. It should be also noted that their model does not capture the possibility of spamming innovation in reaction to mechanisms trying to impose costs on spammers.

Eaton et al. (2008) also develop a model that reflects a similar assumption about the complementarities between filters and sender side countermeasures as in Reshef and Solan (2006). [12] Their theoretical framework suggests that filtering alone, without a sender or receiver side payment scheme may be counterproductive. This is because spammers can respond to improved filter quality by increasing the total volume of messages they send out. Since countermeasures that try to fight spam by imposing sender side costs have been shown to be impractical due to the innovativeness of spammers in hiding their identities, Eaton et al. (2008) propose imposing costs on receivers for reading messages. Imposing costs on receivers would create a dangerous divide globally based on the ability to pay, and would thus threaten the positive economic and social spillovers that connection to the Internet represents.

In addition to their practical and normative limitations, the attempt to deal with the spam problem through either sender or receiver side pricing schemes shows the dangers posed by the economics of spam control. Attempts to increase the costs of undesirable behavior generate reactions by those who will have to

pay the suggested regulatory (civil or criminal sanctions) or market price (ABM or receiver pays mechanisms of Eaton et al.). Importantly, even royalty-free open source filters are a form of cost for receivers, since the software requires administrative and other network related expenditures. In a sense, there is already a receiver side payment in place, but not in explicit monetary terms.

An important omission in these game theoretical models is the lack of attention to the organizational features of the Internet, namely that email and other messaging systems are usually sold by Internet Service Providers (ISPs) as part of a larger bundle of substitutable products, such as voice services. Network providers compete on the price of these packages, not on a per piece basis. If separate receiver or sender prices are imposed, end users can switch to other technologies within the bundle. In general, email requires relatively less bandwidth and processing power. In this context, adopting bad filters may only push customers to less efficient modes of communications, and hence constrain network access on an aggregate basis. Ironically, the sender or receiver side countermeasures can be more costly than the problem they hope to solve.

Kearns (2005) points to another possible economic explanation for spam, specifically by focusing on the incentives of some ISPs to filter effectively. [2] Many large ISPs charge private and public sector users based on the volume of traffic to and from the particular customer, rather than an unlimited bundle. Hence, such sellers can view spam as a source of potential revenue, rather than a cost. Large incumbent carriers as a result may not have the right incentives to adopt the most effective countermeasures. When spammers are aware of this problem, but do not exactly know which networks use good or bad filters, they arguably have economic incentives to produce more spam rather than less. More spam functions as an instrument for evading filters, but also as a means to search for people with a taste for spam across heterogeneous populations of end users.

V. Speed versus accuracy

Existing literature suggests that even if they could read end user preferences about ham/spam accurately, providers of backbone infrastructure may not have sufficient financial incentives to adopt the right technological countermeasures.

In this context, the spam problem can be viewed as a coordination failure not among receivers and spammers (as in Khong, 2004), but among different classes of network owners and operators. Downstream entities may be better off with less incoming spam, since it lowers their infrastructure costs, but cannot force upstream entities to do the filtering for them. Hence, smaller sub-networks must buy their own filters, and install them on the perimeters of the connections with the outside. If upstream server side spam control is not effective, end users are forced to buy spam, as well as virus and other malware filters, for their desktops or downstream servers.

If network costs of spam did not matter, and only false negatives and false positives did, then a system for decentralized pricing negotiations between end users and spammers, like Loder et al. (2004)'s ABM, would generate the first best solution. Implementing pricing mechanisms requires reliable authentication techniques, which are not available today because of earlier advances by spammers. Statistical content filters that are able to learn individual end user preferences represent the second best solution. They mitigate the costs posed by spam on end users, and importantly, allow spammers and people for a taste for their products to connect with each other. Between these two extremes lie a number of other methods for identifying and processing of spam in place today, particularly by larger network operators. When network costs of spam are high, upstream ISPs must consider the throughput of a filter, as well as the expected error rates.

In practice, most antispam systems are a bundle of different types of filters. Basic filtering techniques used today in open source and commercial bundles have been around at least since the mid to late 1990s. These include:

1) Ad hoc feature selection models that look at the characteristics of past spam messages and judgments of administrators for classification, 2) Statistical (Bayesian) content filters that aim to read end user preferences and classify messages based on this knowledge, 3) Checksum/fingerprint systems that look at the prevalence of the same message across the network, 4) Blacklists, IP or domain reputation based countermeasures that focus on the network behavior of spammers. [19] Statistical content filters are the only one of this set that can be implemented in a decentralized manner, and are consequently more accurate as detailed by Cormack and Lynam (2007). [4] Checksum and reputation based measures that ISPs adopted were fast, but are easy to bypass through existing hash busting or IP hijacking techniques known by spammers. [20]

In the early 2000s, some large ISPs adopted centralized fingerprinting/checksum filters. This class of mechanisms was faster than the first generation of content filters available at the time (around 5 times), but was about 5% less accurate in its ability to detect spam.[1]. [21] Over time, spammers responded to this class of mechanisms by adopting algorithms that make each message unique, but essentially convey the same advertisement to the readers. As a result, reputation based systems became popular among large providers in the mid-2000s. Again, reputation based systems are fast, but are designed in a centralized manner, and hence have limited accuracy in terms of false negatives relative to decentralized content filters.(around 30%) [22]. Given the differences in accuracy, technological choices upstream may appear inefficient to downstream ISPs or corporate networks.

The technological tradeoff represented by the two classes of centralized filters adopted in this period provides a plausible explanation both for the growth, and persistence of spam. In the models reviewed in the last section, there is a broad consensus that high levels of spam we observe today would not exists if a perfect statistical filter existed, and was implemented by a good part of sub-networks. Intuitively, this means that when too many sub-networks have bad filters, for one reason or another, spammers are encouraged to send out more messages when aiming to maintain a constant response rate to their advertisements, or smooth their incomes.

An interesting illustration of the importance of network costs is found in the recent industry trend to outsource the processing of spam, as well as antivirus and other malware, to specialized firms. Providers of scanning and filtering services for ISPs allow network providers to externalize the risks of making mistakes about in-house technologies. Outsourcing to specialists essentially allows network providers to buy insurance against hard-to-define changes in spam and antispam technologies.

VI. The economic response

Development and adoption of relatively inaccurate, but fast centralized antispam systems in the early to mid 2000s helps explain why the level and sophistication of spam has continued to grow over time. An asymmetric distribution of good and bad filters motivates spammers to search for customers by sending out more messages, cloak their origins, and envelope their advertisement in images, pdf files, and other packages. These reactions invariably increase the network costs of spam. For instance, advertisement embedded in images take up more computational power and bandwidth than text messages.

This technological perspective on the levels, sophistication, and persistence of spam suggests a number of potential market responses. Given the costs of spam, network operators, end users, and software makers should have some incentives to organize and try to strengthen authentication protocols. The search for new standards like SPF and DKIM illustrates the presence of such economic incentives. [23, 24] Since there are a number of methods available for bypassing these mechanisms already, they are unlikely to provide a robust solution in the longer run.

A more promising economic response to the gap between accurate and fast filters has been further optimization of content filters. Opensource examples of this response can be

found in systems such as CRM114 discriminator, or Bogofilter, which extend the principles of content filtering to retain accuracy, but are much faster than the first generation of this class of countermeasures. [25, 26] To substantiate this point, the following panel illustrates the difference in throughput capacity between a first and a second generation statistical content classifier. For reference, recall that centralized checksum bundles popularized in the early to mid-2000s were approximately 5x faster than first generation content filters.

We focus on two Bayesian engines that highlight the extremes of the technological gap associated with the asymmetric distribution of filter quality across the network. The famous SpamAssassin is a useful benchmark for first generation systems since it continues to operate as the analytical core of a wide range of commercial software and appliance front ends, including a variety of commercial fingerprint and reputation based systems. [27]

COMDOM Antispam for servers represents a second generation statistical filter, released on a commercial basis in 2007. The tests show an improvement in terms of processing capacity of a magnitude of at least 30 times over a period less than 5 years in distributed statistical filters.

In terms of the mechanics of designing content scanning and filtering engines, the comparisons have a second implication. Processing emails takes time, which can be decomposed into two distinct components. The first part is for scanning the content, which depends on the size of each message. The second is the overhead for classification, which is fixed for a given technology. The throughput comparisons between first and second generation Bayesian content filters in the figures below reveal that the generational shift results primarily from reductions in variable costs involved in scanning and tokenization.

Throughput test - COMDOM® Antispam vs. Spam Assassin - msg / min:

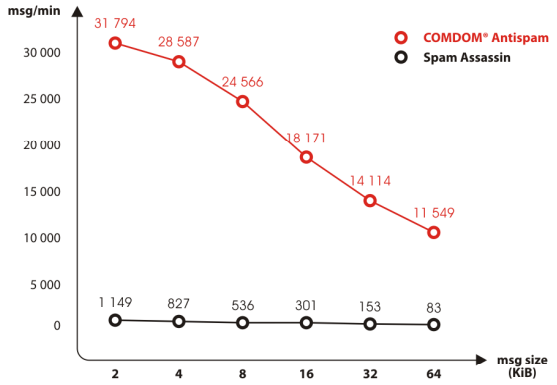

Throughput test - COMDOM® Antispam vs. Spam Assassin - KiB / min:

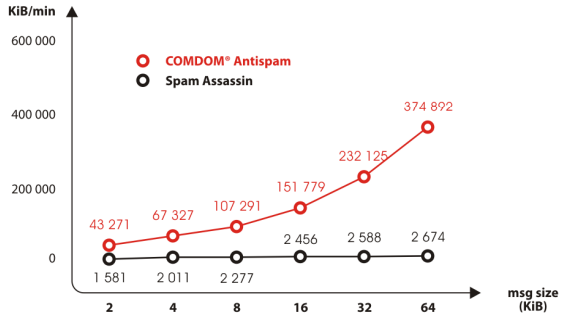

Test configuration:

SW:
COMDOM® Antispam version: 1.0
Spamassassin version: 3.1.0a
Spamassassin based SMTP/LMTP proxy daemon spampd version: 2.30
Benchmark test program: Postal (http://www.coker.com.au/postal/)

HW:
CPU: Intel® Pentium® M processor 1.70GHz with 2048 KB L2 Cache
MB: G5M100-N (Intel® 855GME chipset)
LAN1: Intel 82562EM, LAN2: Intel 82551QM
RAM: 2x1GB DDR400, HDD: 120GB ATA Disk Seagate ST3120026A

VII. Implications

The growth in levels and sophistication of spam has puzzled economists, computer scientists, and public policy makers. A wide range of regulatory and technical countermeasures have been proposed to fight spam by increasing costs on spammers, or

through pricing schemes for legitimate emails. Proponents of such mechanisms tend to ignore expected responses by spammers, as well as the value of email as an accessible and reliable channel for communications. Although explicit or implicit pricing schemes may seem as an obvious solution to some economists, implementing such mechanisms induces spammers to use even more costly techniques, and to send out a larger quantity of advertisements. Blacklists, reputation based filters, and checksum/fingerprinting systems invariably hope to fight spam by increasing the costs of mass mailings. This is because they provide centralized administrative structures to veto communications between advertisers and those with a taste for spam. As a result, spammers tend to produce larger volumes of undesirable content in searching for their audience.

Since the mid 2000s, market responses to the technological asymmetries that accentuate the spam problem have materialized in two distinct ways. First, attempts have been made to improve the ability of receivers to authenticate the origins of incoming mail. Second, developers of decentralized and self-learning content filters have improved the capacity of their software to scan and classify messages.

The development of faster and more robust antispam technologies however does not necessarily mean that they will be adopted by a sufficient part of sub-networks. The speed of technological adjustment to self-learning content filters across will arguably condition future levels and patterns of spam. If a substantial proportion of semi-autonomous networks retain inaccurate centralized reputation based or checksum/fingerprint systems to lower infrastructure costs, spammers have corresponding incentives to produce more advertisement in their search for their target audience.

One way of dealing with the problem would be to improve the quality of information about technological tradeoffs facing network administrators. For instance, organizations providing security system certification services could add throughput tests to their usual assessment of false negatives and positives. [25] More accurate information about features of available filters should help the process of adjustment to more effective mechanisms for mitigating the costs of spam.

References

[1] http://www.trustedsource.org/

[2] Kearns (2005) Economics, Computer Science, and Policy. Issues in Science and Technology, Winter.

[3] Sahami, Dumais, Heckerman, and Horvitz (1998). A Bayesian Approch to Filteing Junk Email. AAAI Workshop on Learning for Text Categorization.

[4]Cormack and Lynam (2007) On-line Supervised Spam Filter Evaluation, ACM Transactions on Information Systems 25, 3.

[5] Rajabiun (2007) Spam and the Digital Divide, Virus Bulletin, December.

[6] Loder, Van Alstyne, and Wash (2004) Information Asymmetry and Thwarting Spam. http://web.mit.edu/marshall/www/home.html

[7] Khong (2004) An Economic Analysis of Spam Law. Erasmus Law and Economics Review, 1.

[8] http://www.ciparx.ca/

[9] http://www.nd.edu/~networks/

[10] Gabaix, Laibson, and Li (2005) Extreme Value Theory and the Effect of Competition on Profits. http://pages.stern.nyu.edu/~xgabaix/

[11] Reshef and Solan (2006) The Effects of Anti-spam Methods on Spam Mail. CEAS 2006.

[12] Eaton, MacDonald, and Meriluoto (2008). Spam – Solutions and their Problems. University of Calgary, Dept. of Economics Working Paper.

[13] Buchanan and Yoon (2000) Symmetric Tragedies: Commons and the Anticommons, Journal of Law and Economics, April.

[14] Alperovitch, Judge, and Krasser (2007) A Taxonomy of Email Reputation Systems, ICDCS Workshops.

[15] Becker (1968) Crime and Punishment: An Economic Approach. The Journal of Political Economy 76.

[16] Becker (1993) Noble Lecture: The Economic Way of Looking at Behavior, The Journal of Political Economy, 101.

[17] Friedman (1984) Efficient Institutions for the Private Enforcement of Law, Journal of Legal Studies, 13.

[18] Androutsopoulos, Magirou and Vassilakis (2005) A Game Theoretic Model of Spam E-Mailing. CEAS 2005.

[19] Ramachandran and Feamster. (2006) Understanding the Network-Level Behavior of Spammers, SIGCOMM 06, Pisa, Italy.

[20] http://www.jgc.org/tsc.html

[21] Antispam Technology Impact Assessment Guide: 2007. http://www.comdomsoft.com/en/antispam/white-papers/

[22] Results of Anti-Spam Solution Testing (2007), Opus One.

[23] http://www.openspf.org/

[24] http://www.dkim.org/

[25] http://bogofilter.sourceforge.net/

[26] http://crm114.sourceforge.net/

[27] http://spamassassin.apache.org/

[28] http://www.comdomsoft.com/

[25] The current certification protocols for antispam at two such organizations, West Coast Labs (http://www.westcoastlabs.org) and ICSA Labs (www.icsalabs.com) do not incorporate throughput testing.