Shall We Stop All Unsolicited Email Messages?

Rui Dai
College of Management
Georgia Institute of Technology
rui.dai@mgt.gatech.edu

Kang Li
Department of Computer Science
University of Georgia
kangli@cs.uga.edu

Abstract. Spam is commonly defined as unsolicited email messages, and the goal of spam filtering is to eliminate these messages. In this paper, we argue that not all unsolicited email messages are useless to recipients, and dropping all unsolicited messages is not optimal for either the recipient or the sender. Increasing the sender's cost (with stamps, puzzles or tarpits) has been another approach aimed at discouraging senders from generating large volumes of spam. But the optimal level of cost to be imposed upon the senders, the most important factor in this approach, has yet to be determined. In this paper, we propose a combination of cost and filtering approaches: charging senders of massive unknown messages or unsolicited messages based on aggregated filter information.[1] We present both an economic model to calculate optimal surcharges under different conditions and the results of a simulation meant to study the consequences of different charging mechanisms. We use theory and simulation to show that our *differentiated surcharge* mechanism can improve social welfare compared with a stamp or a "perfect" filter. Our work can also shed light on how much extra cost we should impose on senders to regulate incoming email messages, which is a fundamental factor neglected in previous research.

1. Introduction

Unsolicited email messages, also known as spam, have become a serious problem for Internet users. They cause disruptions to many users every day [1]. Recent studies show that between one-seventh and one-half of email messages going into an Internet user's mailbox are spam [2]. Currently, anti-spam warriors and their email marketing opponents are in an arms race. Right now, the main anti-spam mechanism is the spam filter [3, 4, 5], the purpose of which is to prevent all unsolicited (marketing) messages from reaching recipients. Marketing senders (including spammers) who want to get unsolicited email to receivers disguise their messages as non-marketing messages to circumvent filters.

In this paper, we argue that this arms race is very hard to resolve and is leading to a complicated and inefficient email system for both sides. Ideally, a perfect filter would only let solicited email messages enter the inbox. Spam would then disappear. In practice, this filter is nowhere close to being built, and marketers are still circumventing filters. As a result, the arms race is still going on. Sophisticated spam filters, such as learning based Bayesian filters, have been built to identify spam messages. To counter attack that, marketers have started to use more textural obscurities, images, or even apply the Bayesian algorithm to infer the filter behaviors in order to bypass them. As a result, users are unhappy with the amount of low quality email messages in their mailbox, which seems to indicate the failure of spam filters. Marketers are also unhappy with the process they experience. They have to make a special effort to disguise messages, such as sending individual email messages rather than batching them together as bulk messages.

Instead of continuing to build more complicated filters, we propose an economic approach to control the sender and recipient behaviors so that we can have a simple yet still effective system. Using an economic approach to control spam is not a new idea, and we review a few of them in the related work section. Our solution differs from them by combining pricing (or cost in a general sense) with filters for massive bulk messages. In particular, we would like to use filter output as an indication of user preference and use that to determine the price charged to the sender. This is different from the stamp method since we do not impose a uniform surcharge [2, 9]. It depends on the quality of the email messages, which is a potential benefit. It uses ex ante preference information to decide the quality of an email, instead of ex post report [12], which can avoid a moral hazard problem. Our objective is to maximize recipient benefit (enterprise solution) or overall social welfare (government solution), which is different from Van Zandt's work [13]. A simple illustration of the system is like this: Unsolicited massive messages are charged based on the aggregated filter score. This only applies to massive bulk messages, and we assume that bulk

[1] Here charging is not limited to money transfer. It can be implemented in different forms.

mail sent in the form of individual email messages can be detected at the recipient side. The details of the detection are out of the scope of this discussion. When an email is solicited (e.g. identified by a filter, white-list etc.), it gets into the inbox. When an email is unsolicited but is not a bulk mail, then the processing is left to the individual user's filter.

With this approach, we expect the receiver to be able to control the information quality entering the inbox by controlling the price. With this pricing control, senders are encouraged to behave efficiently (e.g. sending email messages to singular users separately rather than mass-lumping them all together in one email; targeting a more specific audience). We describe this economics-engineering approach and present a model to study the email system's behavior under our approach as well as other spam control approaches. We use the following performance metrics to evaluate these approaches: *average inbox quality* (the average score of all messages in the recipient's inbox), the *sender's benefit*, the *receiver's benefit*, and the *social welfare*. The paper is organized as following: Section 2 discusses the related work, Section 3 describes our spam control approach and presents an economic model, Section 4 presents our simulation results based on the model, Section 5 discusses the possible challenges such as deployment issues, and Section 6 concludes the paper.

2. Related Work

While spam filters, such as white-list, black-list and content-based filters [3,4,5,6], are the dominant anti-spam technique, we only review the economic-based approaches in this paper due to the space limit. Existing economic-based research [2, 7, 9, 10, 12, 13] has explored pricing mechanisms to control spam. Papers [2, 9, 10] focus on the stamp and authentication mechanisms. Before sending a message, a sender pays a "stamp" to let the recipient accept the message. By charging a small price to send a message, the pricing system shifts the task of screening messages from recipients, who do not priorly know the content of a message, to senders, who do. The stamp mechanism rewards senders for being selective in sending messages. Senders' information about the recipients' interests enables the senders to be more selective, increasing the chances of their messages being relevant to and read by the recipients. However, the researchers of this method are facing deployment issues. In order to use a "stamp" based email system, an authentication infrastructure for charging must be widely deployed to all senders and receivers. In addition, a flat pricing mechanism does not reflect the receiver's preference. Loder et al. [12] employ a combination of escrow services and charging mechanisms which rely on the recipient to decide if an email is spam and whether or not to charge. Theoretically, this approach can achieve socially optimal outcome. But it critically depends on subjective opinion, which can be biased because of conflicted interests. Van Zandt [13] addresses an information overload problem due to low transmission cost. This research focuses on how to increase sender benefit, which so far is not a major concern in anti-spam research.

3. Theoretical Model

In order to focus on the essential issues related to information overload and spam, we use the following assumptions to define the problem domain of this paper.

Assumption 1: We only consider mass email messages in our system. If an email is delivered, every recipient served by the mail server will receive a copy of the email. Otherwise no one will receive the email. [2]

Advanced filter technology (Bayesian) can be trained and used to estimate personal preference for specific email messages. However, the error of estimation can be quite significant for two reasons. First, the preference of each individual recipient is inherently ambiguous and changes over time. Second, email senders can derive individual preferences and manipulate the content of an email message. If we use feedback from an individual filter to take action – accept or block an email – a sender can manipulate the content of an email so that it has a better chance to be delivered in the future. Therefore, we only rely on aggregated information from all individual filters in our system. For example, the aggregated information of an incoming email can be "80% of the recipients will find this email useful." This can be viewed as a measure of the quality of the email. Then the mail server decides whether to deliver or deny the email. By applying aggregated information, we can reduce estimation errors and make it harder for a sender to detect a recipient's personal preferences. Since the information is aggregated, the actions taken by the mail server are also the same for all recipients.

[2] Here the recipient population can be thought of as employees in a company or an interest group.

Since it is possible for senders to use fake email titles to attract or mislead recipients into reading their email messages, recipients should not make decisions on whether to read an email based on its title. Therefore, we assume a recipient cannot precisely tell which email is more important simply by looking at email titles and thus has to really read an email to evaluate the potential value. Based on this assumption, we conclude that a recipient will either read all email messages in the mailbox or read none of them. This result is presented in Lemma 1.

Based on the previously mentioned assumptions, we built a stylized model to formulate and compare different anti-spam schemes. Because it is difficult for an individual recipient to tell whether an email is spam, personal filters generate an estimated possibility that a recipient will find an email useful. We use a random variable, hit ratio $r_i \in [0,1]$, to denote how likely an email *i* is useful. The sender of an email knows the *expected hit ratio*, $h_i = \int_0^1 g_i(r)dr$, of email *i*, where g_i is the p.d.f of of r_i. We have argued previously that r_i can be very inaccurate. However, given a large number of recipients, the estimation of expected hit ratio, h_i, would be more stable and not easily biased by an individual recipient.

We use two parameters to characterize each email: *expected hit ratio* (h), which has been defined previously, and *payoff per hit (p)*, which is the compensation to the sender for each hit, i.e. a recipient finds an email valuable. Some email messages could have a very low expected hit ratio but a high payoff per hit, and vice versa. In order to capture the diversity of email messages, let f (h, p) denote the joint probability density function of all email messages' expected hit ratio and payoff per hit. By incorporating the sender's benefits, we can study the benefit on the the sender's side and how a different mechanism can change the sender's behavior.

Now, let us formulate the recipient's cost and benefit of reading email messages. Each recipient receives email messages at rate λ_i and can process email messages at rate μ. For ease of analysis, we assume that every time she reads a useful email, her utility will increase by *v*. However, processing each email will incur a constant positive cost, *c*, on the recipient. Following assumption 3, each recipient chooses process rate μ* to solve

$$\begin{aligned} &\underset{\mu}{Max}\ \mu v\bar{h} - c\mu, \\ &s.t.\ \ \mu \le \lambda_i \end{aligned} \qquad (1)$$

where $\bar{h}$ is the average expected hit ratio of all incoming email messages. It is easy to find the optimal solution to (1):

$$\mu^* = \begin{cases} \lambda_i & \text{if} \quad v\bar{h} \ge c \\ 0 & \text{if} \quad v\bar{h} < c. \end{cases} \qquad (2)$$

Lemma 1: A recipient will read either all email messages in her mailbox or none of them.

This is consistent with empirical evidence: when email messages in a mailbox are likely to be useful ones, users tend to read them all; otherwise, if a mailbox is full of trash email messages, users tend to stop using email systems even though there could be a few useful email messages. Suppose there are totally *m* identical potential email recipients served by the mail server. Although each sender could send a bunch of different email messages, whether to send out each email to all recipients can be deemed a sequence of independent decisions. Therefore, we focus on the decision problem of whether to send out an individual email to all *m* recipients given its characteristics (*h* and *p*) and hence the expected profits of sending this email. Without considering any extra marginal cost[3], from (2), the expected payoff to send an email with (*h'*, *p'*) to all recipients is

$$U(h', p') = \begin{cases} mh'p' & \text{if} \quad v\bar{h} \ge c \\ 0 & \text{if} \quad v\bar{h} < c. \end{cases} \qquad (3)$$

Regardless of the average expected hit ratio and payoff per hit, without any surcharge, every email will generate non-negative benefit for the sender. Therefore, in the absence of regulation, senders will always choose to send out all email messages without considering whether the recipients will find these email messages valuable or not. In this case, the expected hit ratio $\bar{h}$ can be expressed as

$$\bar{h} = \int_0^\infty \int_0^1 hf(h,p)dhdp\,. \qquad (4)$$

[3] Network connection cost and server cost can be viewed as a lump-sum cost.

As we know, most unsolicited email messages have a very low expected hit ratio, and a smaller $\bar{h}$ will lead to a lower expected benefit for reading an email, $v\bar{h}$. From (1) and (2), a recipient will eventually shut down her mailbox since there are too many junk email messages. On the senders' side, the total surplus will collapse to 0 according to (3).

We have used our framework to show that a free email system will lead to a "lose-lose" situation that is similar to the prisoners' dilemma. Several mechanisms have been proposed to solve the spam problem. In the following sections, we model and compare several proposed schemes and our charging mechanism using a unified framework. We start with an enterprise solution to maximize the recipient's benefit. Then we discuss a stamp system (a uniform surcharge approach). Finally, we show our charging scheme which maximizes overall benefits for all senders and recipients.

Case 1: Maximize the Recipient's Surplus

We take the recipient's point of view. Suppose all recipients are within an organization. The organization tries to configure its mail server to maximize the total surplus for all recipients. Let λ_e denote the exogenous rate of email requests (from the clients of e-marketers that ask them to send out email advertisements) and $\theta(h, p)$ denote the probability that an e-marketer will send out an email of (p, h), $0 \le \theta(h, p) \le 1$. The objective function is:

$$\max_{\theta(p,h)} \; m\lambda_i(v\bar{h} - c) \qquad (5)$$
$$s.t. \quad v\bar{h} \ge c$$

where $\bar{h} = \dfrac{\iint h\theta(h, p) f(h, p)dhdp}{\iint \theta(h, p) f(h, p)dhdp}$ and $\lambda_i = \lambda_e \iint f(h, p)\theta(h, p)dhdp$.

By applying the Kuhn-Tucker theorem [16], we have

$$\theta^*(h, p) = \begin{cases} 1 & vh > c \\ \text{any value in } [0,1] & vh = c \\ 0 & vh < c \end{cases}$$

It is clear that, typically, the constraint in (5) is not binding.

Since we only consider the recipient's benefits, the expected hit ratio is the only criterion we need to consider. It must be larger than or equal to *v/c*, in order for an email message to be allowed into the system. We use Figure 1 to depict this result. Only email messages in the shadow area should be delivered to individual mailboxes.

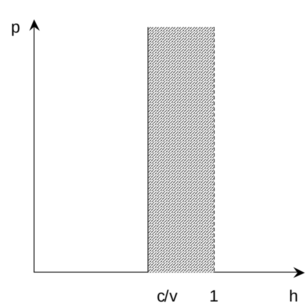

Figure 1. Optimal Enterprise Solution

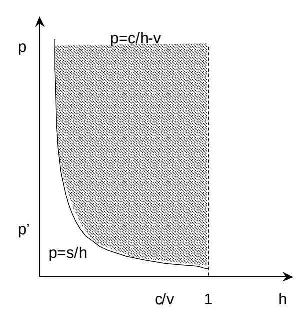

Figure 2. Fixed Surcharge

Case 2: Stamp – Fixed Surcharge

One proposal is to charge a fixed fee (or other kind of cost) for all email messages in order to reduce email volume. However, how much to charge for each stamp is left unanswered. In this section, we use our model to derive the efficient stamp price for an entire system. Suppose the stamp price is *s*. In order for a sender to send an email, the expected return must be larger than the stamp price, $hp \ge s$, given the assumption that the email will be read. Figure 2 shows the area where email messages will be sent (the shadow area) when the senders are faced with a fixed price *s*. We can tell from the figure that email messages with a very low expected hit ratio or a very low payoff will not be sent. We take the social planner's point of view and find the optimal s^* to maximize social welfare:

$$\underset{s}{Max}\ m\lambda_e \int_{s/h}^{\infty}\int_0^1 hpf(h,p)dhdp + m\lambda_i(v\bar{h}-c)$$

$$s.t.\quad v\bar{h} \geq c$$

Where $\bar{h} = \dfrac{\int_{s/h}^{\infty}\int_0^1 hf(h,p)dhdp}{\int_{s/h}^{\infty}\int_0^1 f(h,p)dhdp}$ and $\lambda_i = \lambda_e \int_{s/h}^{\infty}\int_0^1 f(h,p)dhdp$

The necessary condition for s^* to be socially optimal is:

$$-\int_0^1 s^* f(h,\frac{s^*}{h})dh - v\int_0^1 hf(h,\frac{s^*}{h})dh + c\int_0^1 f(h,\frac{s^*}{h})dh = 0. \qquad (6)$$

It is not difficult to tell from (6) that $s^* \leq c$. If we further assume $f_p'(\bullet,p) \geq 0$, according to the second order condition, there is a unique s* that satisfies (6).

Case 3: Differentiated Surcharge

In this section, we do not specify a particular pricing scheme at the very beginning. Instead, at first we develop an optimal centralized solution: suppose there is a social planner who decides which email messages should be sent to recipients. We want to know how the social planner makes decisions. Then we implement this solution in a decentralized system by imposing a surcharge on senders, so that they will make the same decisions as the social planner does.

Now let us characterize the social planner's decision problem. When the social planner is faced with an email of (h,p), the probability that she will send it to recipients is $\theta(h,p)$, $0 \leq \theta(h,p) \leq 1$. The social planner's objective function is:

$$\underset{\theta(p,h)}{Max}\ m\lambda_e \iint hp\theta(h,p)f(h,p)dhdp + m\lambda_i(v\bar{h}-c) \qquad (7)$$

$$s.t.\quad v\bar{h} \geq c$$

where $\bar{h} = \dfrac{\iint h\theta(h,p)f(h,p)dhdp}{\iint \theta(h,p)f(h,p)dhdp}$ and $\lambda_i = \lambda_e \iint \theta(h,p)f(h,p)dhdp$.

From the Kuhn-Tucker condition, if the constraint is not binding,[4] we have

$$\theta^*(h,p) = \begin{cases} 1 & (v+p)h > c \\ any\ value\ in\,[0,1] & (v+p)h = c \\ 0 & (v+p)h < c \end{cases} \qquad (8)$$

where $\theta^*(h,p)$ is the optimal solution to (7). This solution tells us that, in order to reach the socially optimal point, we should only allow email messages with either a relatively high payoff *p* or a relatively high expected hit ratio *h* to be sent to recipients. Figure 3 can help us understand this result better.

All the email messages that fall into the shadow area in Figure 3 should be forwarded to recipients. Email messages on the line of *p=c/h-v* can either be sent or not; email messages that fall in the area below the curve should be blocked.

Here we illustrate one possible implementation of this solution in a distributed system. Figure 4 shows how the system works. When the mail server receives an email from a sender, it gets the scores of this email from individual filters (either on the client or server side). Then the server calculates the expected hit ratio, *h*. If the sender wants the email delivered, she must accept a "surcharge"[5] of *max {(c/h-v), 0}*. It is not difficult to see that we can reach the same result as the centralized solution in figure 3: senders will only send email messages in the shadow area. Notice here that the surcharge depends only on the information available to the recipients. They do not need to know the

[4] For the binding case, please refer to a longer version of the paper that is available upon request (currently under-submission).

[5] This "surcharge" can be monetary or resource cost, such as bandwidth or time cost.

sender's payoff. The charging mechanism in this system could be similar to the stamp mechanism, which would require an authentication system deployed by a third party, or it could be implemented locally by charging for computational resources [9,14,15].

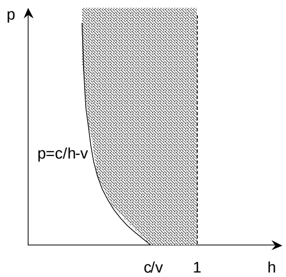

Figure 3. Optimal Differentiated Surcharge

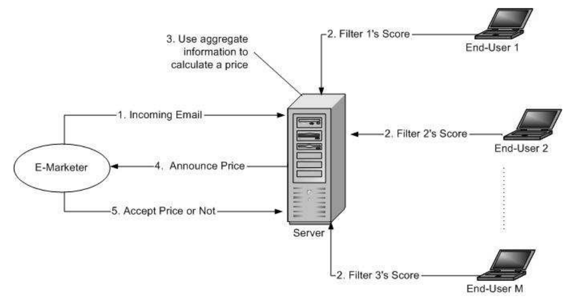

Figure 4. System Running Scenario

4. Experiments

To show the impact of various email control mechanisms, we built a simulation based on the theoretical model presented in Section 3. In these experiments, we took a mixture of solicited and unsolicited messages as the input to our simulation. For solicited email messages, we set the expected hit ratio to 1. For unsolicited messages, we used a spam filter, SpamAssassin [11], to estimate the expected hit ratio. SpamAssassin produces a spam likelihood score for each message. In this simulated environment, the generated score became the expected hit ratio (0 for spam and 1 for non-spam). The unsolicited email messages were obtained from spam samples that were archived in July 2003 [7].

In each experiment, we measured the total benefit for the sender, the receiver, and the whole system. The expected per-email benefit to the sender and the receiver depends on the expected hit ratio of the email message. In these experiments, we normalized the maximum per-email reading benefit (v) to 1. That is, the benefit for the recipient of reading a regular email is 1 and h for a marketing message with an expected hit ratio h. We set the per-email reading cost (c) to a constant, 0.7. The sender's potential benefit from a hit message is a random value generated by a truncated normal distribution (no negative values) with $\mu= 0.3$ and $\delta^2=1$.

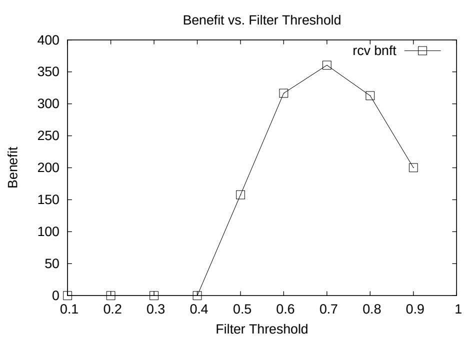

Figure 5: Receiver Benefit vs. Filter Threshold

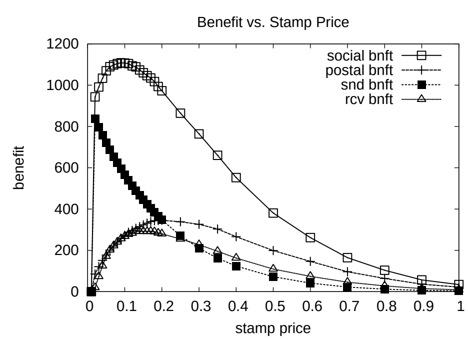

Figure 6: Benefit vs. Stamp Price

4.1 Receiver Benefit and Filter Threshold

We first study the filter behavior from the receiver's perspective, particularly the selection of a spam filter threshold and its impact on the receiver's benefit. Based on our theoretical model presented in section 3 and the values of the parameters set in our simulation, the optimal threshold should be c/v =0.7. To test this prediction, we used synthetic email messages as the input to the simulation, and the expected benefit of each email message and reading cost are calculated in the ways stated earlier. We chose a fixed input volume and measured the benefit of receivers with different filter thresholds.

The simulation results show that adjusting the filter threshold clearly affects the total benefit of the receiver. When the filter threshold is very low, the quality of email messages in the receiver's inbox is also very low, and the total reading cost becomes higher than the email message benefit. Under this condition, the best strategy for receivers is to stop reading email messages. When this happens, there is zero benefit for senders and receivers. With a “perfect” filter whose threshold is set to one, the quality of email messages in users’ inboxes is high, and reading email messages would thus increase user benefit because the expected value brought by solicited messages is higher than the reading cost (according to our simulation and reality), otherwise users would stop soliciting those messages. Getting only solicited messages would not maximize receiver benefit. Reading an unsolicited message with a high enough expected value (e.g. with threshold 0.7) can offset the reading cost as shown by the simulation results.

In the rest of the paper, when we compare the outcomes of different spam control approaches, we include two filter-based approaches. One uses the optimal filter threshold to control spam. The other one uses a “perfect” filter, in which only solicited messages are accepted.

4.2 Stamp Price

In this section, we present simulation results of using a stamp to control unsolicited email messages. The sender decides whether to send a message or not based solely on the potential benefit of the message and the stamp price. The receiver takes all the incoming messages and reads them as long as the expected value is larger than the reading cost. As stated in section 3, a stamp directly affects the sender’s benefit and the social welfare since senders have to consider the additional cost introduced by the stamp and change their behavior accordingly. The stamp system can be implemented in different ways: it can be a pure money transfer between senders and recipients or a third party can serve as the gatekeeper (the government) whose money transfer is referred to as *postal benefit*.

The simulation output is presented in Figure 6. In this simulation, we used a fixed email volume (1000) and studied how stamp price affects the benefits. When the stamp price is too low (close to 0), the pricing mechanism is not effective in stopping senders, therefore recipients get low quality messages. As a result, recipients stop using email, and all parties have zero benefit. Once the stamp price increases to the level where message quality is high enough to bring the recipient’s attention back (because the expected benefit of reading an email message is larger than the reading cost), then the benefits of the receiver start to increase, as does the benefits of the sender, the benefits of the stamp issuer, and the general social welfare. As stamp price keeps increasing, senders start to send fewer email messages since the expected benefits cannot cover the cost. When the price is high enough, everyone starts to lose benefits. Again, just as predicted by our theoretical model in section 3, the optimal stamp price is around 0.1. In the later comparison of different spam control approaches, we use the optimal stamp price for the fixed surcharge mechanism.

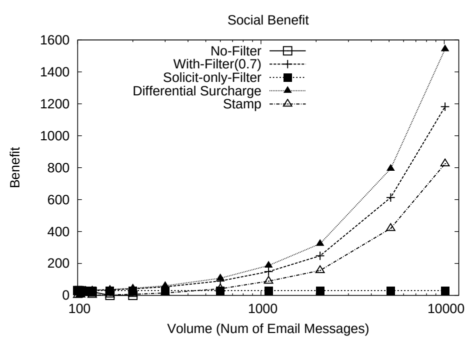

Figure 7: Social welfare vs. Email Volume

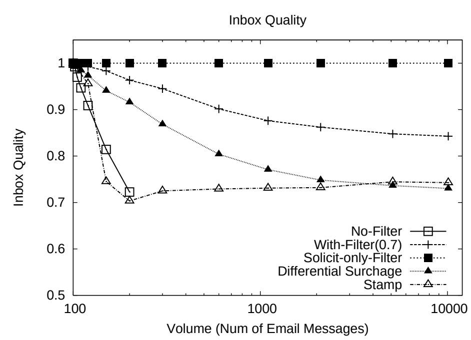

Figure 8: Inbox Quality vs. Email Volume

4.3 Email Message Quality and Social Welfare

In this section, we use simulation results to answer the question about the “future of the email system”, that is –as the volume of email messages increases (assuming senders are greedy and would increase the volume of email messages as long as the increment helps senders maximize their benefit), what would be the total benefit generated by each spam control mechanism, and what would be the inbox quality for the recipient side? Our simulation used synthetic input that combined a fixed amount (100) of solicited email messages, and a variable amount (from 0 to 10000) of unsolicited messages. With this input, we evaluated how the social welfare changes as the volume of email messages increases for different spam control mechanisms.

Figure 7 depicts the social welfare for each spam control mechanism with different email volumes. When no filter is applied, the recipient’s inbox is flooded with unsolicited messages, and the inbox quality is very low. As we have predicted, users would stop using email, and this would lead to zero benefit for senders as well as a low social welfare. When a “perfect” spam filter is used so that only solicited messages get into the inbox, receivers get the same set of messages no matter what volume of marked messages exist. Under this condition, the social welfare is a constant positive number.

All the other approaches (a filter with an optimal threshold, a differentiated surcharge approach, and a stamp approach) let some unsolicited email messages get into the recipients' mailbox but also regulate inbox quality so that receivers will still read email messages. As a result, the social welfare increases as email volume increases. Just as predicted in section 3, our differentiated surcharge scheme outperforms the other two schemes. The gap increases as the volume of messages gets larger.

Higher social welfare is achieved by letting the recipient read some unsolicited email messages. We measure the inbox quality to illustrate this tradeoff. Figure 8 portrays the inbox quality for each spam control approach as the input volume changes. When no spam control mechanism is used, the inbox quality decreases quickly as the total volume of email messages increases. With the "perfect" spam filter approach, we show that inbox quality can be kept high when the filter threshold is set to 1 so only solicited messages get in. When a filter threshold is used to optimize the receiver's benefit, the inbox quality degrades but is still kept above the filter threshold as the volume of input email messages increases. With our differentiated surcharge scheme, inbox quality also decreases as the email volume increases and stays lower than the case with the optimal filter threshold. The rationale is that we let some unsolicited email messages with a lower expected hit ratio be delivered since the expected return on the sender's side is high, and they would like to pay for them to get through. This can improve overall system efficiency. The stamp approach has even lower quality results. Because stamp price is fixed regardless of the message content, a sender can send email messages as long as the benefit remains higher than the stamp price. This could result in high sender benefit and low receiver benefit. The latter indicates that many lower quality messages could be delivered into the recipient's inbox.

5. Conclusion

In this paper, we argued that the goal of email filters should not be to eliminate all unsolicited email messages. Currently we have so many of them that people tend to believe all unsolicited messages are useless. In this paper, we showed that no matter whether from the viewpoint of the enterprise or social planner, getting a limited amount of unsolicited messages helps by increasing the receiver's benefit as well as social benefit. To keep the high benefit while preserving high email quality, we proposed an approach that combined economic incentive with filter information, in which receivers can express their preference in the email price. Similar to other pricing mechanisms, the differential surcharge mechanism gives incentives for marketers to refine their target audience. Our simulation results show that pricing with a differential surcharge can achieve high social and receiver benefits while still preserving high inbox quality, even when the email volume increases significantly.

Reference

[1] Nicholas Graham, "AOL blocks a billion spam emails in 24 hours", Mercury News Article, Mar 06, 2003.

[2] Sam Vaknin, "The Economics of Spam", published by *United Press International*. July 23, 2002.

[3] Xavier Carreras, and Llues Marquez, "Boosting Trees for Anti-Spam Email Filtering", in *Proceedings of 4th International Conference on Recent Advances in Natural Language Processing*, 2001.

[4] Paul Graham, "A Plan For Spam", in the first SPAM conference, http://www.paulgraham.com/spam.html.

[5] Distributed Spam Clearinghouse, available at http://www.rhyolite.com/anti-spam/dcc/

[6] Michelle Delio, "Not All Asian E-Mail Is Spam", Wired News article, Feb 19, 2002.

[7] Paul Q. Judge, "SPAM Archive". http://www.spamarchive.org/

[8] Ari Juels and John Brainard, "Client Puzzles: A Cryptographic Defense against Connection Depletion", in *Proceedings of NDSS-1999* (Networks and Distributed Security Systems), pages 151-165, 1999.

[9] The Penny Black Project. http://research.microsoft.com/research/sv/PennyBlack, Microsoft Research.

[10] C. Dwork and M. Naor, "Pricing via Processing or Combatting Junk Mail", in Proceeding of *CRYPTO'92*.

[11] SpamAssassin, http://au.spamassassin.org

[12] Thede Loder, Marshall Van Alystyne and Rick Wash,"*An Economic Answer to Unsolicted Communication", the* ACM Conference on Electronic Commerce (EC'04).

[13] Timothy P. Van Zandt,"Information Overload in a Network of Targeted Communication", INSEAD.

[14] Axel Zinser, et al, "Teergrubing" http://www.iksjena.de/mitarb/lutz/usenet/teergrube

[15] Marty Lamb, "Using Statistics to Cause Spammers Pain", http://www.martiansoftware.com/

[16] A. Dixit, "Optimization if Economic Theory", 2nd edition. New York, Oxford Univ. Press, 1990.