Incentive Mechanism for Prompting ISPs to Implement Outbound Filtering of Unwanted Traffic

Yousuke Takahashi, Keisuke Ishibashi

Incentive Mechanism for Prompting ISPs to Implement Outbound Filtering of Unwanted Traffic

Yousuke Takahashi
NTT Service Integration Laboratories
NTT Corporation
Tokyo 180-8585, Japan
Email: takahashi.yousuke@lab.ntt.co.jp

Keisuke Ishibashi
NTT Service Integration Laboratories
NTT Corporation
Tokyo 180-8585, Japan
Email: ishibashi.keisuke@lab.ntt.co.jp

***Abstract*—Methods of filtering unwanted traffic in ISPs are categorized into inbound filtering at receiver ISPs and outbound filtering at sender ISPs. Outbound filtering is effective because it can directly utilize sender information such as dynamic IP address information, but not widely used in the current Internet. This is because outbound filtering does not benefit the ISP that installs it and pays the cost, but benefits ISPs that receive unwanted traffic from the installing ISP. The objective of our study is to solve this incentive problem. We propose an incentive mechanism where ISPs pay penalties for sending unwanted traffic to peering ISPs. This payment is combined with payments of transit fees so that the mechanism will be deployed through the current ISP relationship. The key question here is whether the mechanism will be deployed or not, and we confirmed through simulations that the mechanism is spontaneously adopted by a sufficient number of ISPs who behave in a selfish manner. Accordingly, outbound filtering was also widely installed. In our simulation scenario, there are some ISPs who send little unwanted traffic and should not install the outbound filtering due to its cost. Our simulation showed that with an appropriate penalty charge parameter setting, our proposed mechanism can achieve the state that almost all ISPs who should install the outbound filtering install the filtering, and otherwise not.**

I. Introduction

Unwanted traffic such as spam mail and distributed denial-of-service attacks seriously damages ISPs. For example, 89.1% of all e-mail is spam [1], and ISPs pay immense facility and operating costs in dealing with this unwanted traffic [2]. To reduce these costs, ISPs need to filter unwanted traffic. We focus here on spam mail as a typical example of unwanted traffic.

Methods of filtering unwanted traffic in ISPs are categorized into inbound filtering at receiver ISPs and outbound filtering at sender ISPs.

Receiver-side inbound filtering is important from the viewpoint of protecting network facilities and the ISP's own customers. Many ISPs have implemented some kind of receiver-side inbound filtering against unwanted traffic.

However, $80-90\%$ of spam mail has been generated by botnets in recent years [1], so the origins of spam mail are widely distributed. Consequently, it has been difficult for receiver ISPs to configure access control lists with a high degree of accuracy.

In this case, outbound filtering at sender ISPs is effective because sender ISPs can utilize information that receiver ISPs cannot access. For example, outbound port 25 blocking at a sender ISP, which filters spam mail using the dynamic IP address information of the sender ISP, is effective in blocking spam mail generated by botnets.

However, the adoption ratio of outbound filtering at sender ISPs is much lower than that of inbound filtering at receiver ISPs [3]. One of the reasons for this low ratio is reported to be the incentive problem [4], in that the ISP that pays the cost of implementing outbound filtering is different from the ISP that benefits from the filtering. Accordingly, total optimization cannot be achieved because each ISP acts to achieve an individual optimum. Therefore, low adoption ratio of outbound filtering is caused by the difference between the individual optimum and total optimization. In that case, it is necessary to introduce appropriate mechanism design so that each ISP will take actions to achieve total optimization. Mechanism design is to design a system or a rule to give each player incentives for achieving a total optimum. The objective of our study is to solve the incentive problem for sender ISPs to adopt outbound filtering by performing appropriate mechanism design on ISPs.

Here, we first propose an incentive mechanism that gives sender ISPs incentives for adopting outbound filtering by extending transit contracts between ISPs. Under a transit contract, typically, a customer ISP pays transit fees for sending and receiving traffic to and from its provider ISP. This inter-ISP traffic includes both legitimate traffic and unwanted traffic. Our proposed incentive mechanism is a contract option regarding the unwanted traffic. Under this contract option, if a customer ISP sends unwanted traffic to a provider ISP, the provider ISP will raise transit fees depending on the received amount of unwanted traffic. Conversely, if the provider ISP sends unwanted traffic to the customer ISP, the provider ISP will discount the transit fee depending on the amount of unwanted traffic that is sent. Thus, the sender ISPs will have an economic incentive for adopting outbound filtering.

Though this simple contract option promotes the adoption of outbound filtering, it is uncertain whether this

contract option itself would be spontaneously adopted or not. The second contribution of this paper is that we examine through multi-agent simulation whether our proposed contract option can be diffused in a situation where each ISP can make decisions individually. The results of our simulation showed that when the proposed contract option was available, a certain number of ISPs adopted it, and consequently, about 80% of ISPs adopted outbound filtering for unwanted traffic.

This result, though not all ISPs adopted the filtering, is revealed that as a state close to the social optimum. This is because there are some ISPs who send little number of spam mails and whose cost of outbound filtering is higher than the benefit to them in reducing spam mail. With an appropriate spam-penalty parameter setting, our incentive mechanism achieves the state that almost all ISPs who should install the outbound filtering install the filtering, and otherwise not. The results suggest that our economic approach to incentivize individual players is superior to the institutional approach of forcing all players to obey, such as through legal requirements.

This paper is organized as follows. In Section II, previous studies on various mechanisms and their design are presented. In Section III, our proposed incentive mechanism is explained. In Section IV, we model an inter-ISP network and describe our multi-agent simulation results. Finally, Section V concludes this paper and briefly touches on future work.

II. Related Work

In Ref. [4], the authors discuss the incentive problem in installing outbound filtering for unwanted traffic. To solve this problem, they propose the creation of a new institution that monitors security levels. ISPs supervised by this institution must maintain network security by adopting outbound filtering. Consequently, safe communication can be achieved among these ISPs. Thus, this institution gives ISPs that want to participate in this network an incentive to adopt outbound filtering.

In Ref. [5], "Bonded Sender" was proposed, which is an accounting mechanism for ISPs according to the amount of spam mail sent. To participate in Bonded Sender, ISPs have to deposit a certain amount of money in the Bonded Sender organization. When the Bonded Sender organization accepts claims from users who have received spam mail, the organization checks the origin ISPs of the spam mail. If the origin ISP is part of the Bonded Sender program, the Bonded Sender organization donates the deposit from the origin ISP to a charity. Thus, Bonded Sender gives participant ISPs an economic incentive to adopt outbound filtering for spam mail.

CentMail [6] is an accounting mechanism to keep end users from sending enormous amounts of spam mail. Participants in CentMail can attach a CentMail stamp of 1 cent to e-mails they send. The CentMail stamp is an e-mail signature. Spammers usually send enormous amounts of spam mail to the general public, and thus, it is difficult for spammers to participate in CentMail economically. End users can distinguish legitimate emails from vast amounts of received emails easily by checking whether a CentMail stamp is attached.

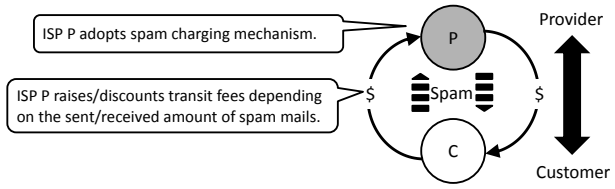

Fig. 1. Spam charging mechanism

These incentive mechanisms aim to decrease unwanted traffic by charging the player sending the unwanted traffic, as our proposed mechanism does. These mechanisms would work as expected if they were adopted by a sufficient number of players. To achieve such a charging mechanism, a charging player needs to conclude a contract with a charged player. There are an enormous number of players from diverse countries and regions on the Internet, so it is difficult to construct a common charging scheme for such players.

Moreover, whether or not the incentive mechanism can effectively reduce the volume of spam mail depends on whether it can involve spammer ISPs that send massive amounts of spam. This is because most spam mail is sent by several spammer ISPs [1]. However, spammer ISPs would not be keen to participate in these schemes because if they did, they would have to be pay high fees based on the amount of unwanted traffic they sent.

III. Incentive mechanism

A. Overview

As discussed in Sec.I, we propose a spam charging mechanism based on current transit contracts, whereas the previously mentioned mechanisms require new contract relationships (Fig. 1). When an ISP installs a spam charging mechanism, the sender-side ISP is charged for spam mail on the transit link between the ISP and its customer ISP, and the peering link between the ISP and its peer ISP. Spam mail charges are assessed at the same time as monthly traffic charges.

Our mechanism requires not only counting traffic volume but classifying traffic into unwanted or not. The actual methods of classification are out of scope of our paper, but recent researches on traffic classification can be applied for our mechanism [7]. Though those methods have still limitations in neither accuracy nor timeliness to classify all traffic between ISPs in real-time and drop the unwanted traffic, but may be enough to classify traffic off-line and charge according to the amount of unwanted traffic.

B. Adoption and diffusion of mechanism

We show scenarios in which our proposed charging mechanism is diffused among ISPs in two phases, how the

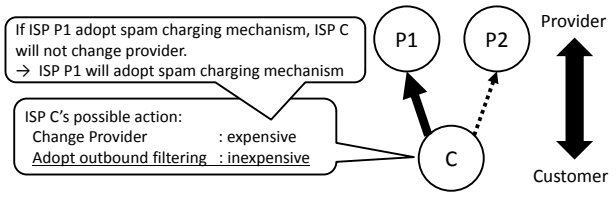

Fig. 2. Phase 1: first adopter

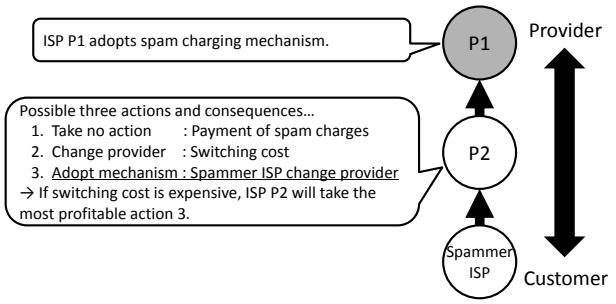

Fig. 3. Phase 2: diffusion

first ISP adopts the mechanism when there are no adopters, and how the mechanism is diffused after the first adopters.

1) Phase 1: first adopter: Here, we show how an ISP adopts the mechanism (Fig. 2). Let Customer ISP C sign a transit contract with provider ISP P1, because it offers a lower transit cost than another provider, ISP P2. In this case, when ISP P1 installs a spam charging mechanism, the profit of ISP C decreases due to payment of the spam charges. However, if the difference in the transit fee is more than the spam charges, ISP C will maintain a transit contract with ISP P1. Therefore, ISP P1 has an incentive to introduce a spam charging mechanism as a differentiation strategy. We think this differentiation strategy is plausible considering the recent takedown of McColo [8], which hosted multiple botnet controllers to send spam mail and was de-peered by its provider ISP. Our mechanism can be regarded as a graceful countermeasure against spammer ISPs. Another situation is when customer ISP C has no ISP to connect to except provider ISP P1, which is known as vendor lock-in. In this situation, the switching cost is too high for ISP C to change providers. Thus, our proposed incentive mechanism in the existing economic scheme has various triggers that stimulate ISPs to be first adopters.

2) Phase 2: diffusion: Next, we show how our mechanism is diffused after the first adopter appears (Fig. 3). ISP P1 has installed a spam charging mechanism whereas ISP P2 has not. ISP P1 charges ISP P2 for spam mail sent by a spammer ISP through ISP P2. Because ISP P2 will have to pay spam charges if it takes no action, ISP P2 responds to ISP P1's action by adopting a spam charging mechanism or switching providers. Thus, if a provider ISP adopts a spam charging mechanism, it can then place pressure on its customer ISP to also adopt a spam charging mechanism.

IV. Simulation

In Section III, we conceptually showed the scenarios that our proposed charging mechanism could be voluntarily adopted by each ISP. However, there still remains a question whether the mechanism will be adopted by sufficient number of ISPs to decrease the total volume of spam mails.

In this section, we confirm whether the volume of spam was reduced (and the total benefit increased) when our mechanism was adopted through multi-agent simulation. We built our simulation model based on the existing agent-based model for the evolution of the Internet ecosystem [9]. It models a hierarchical structure of the Internet and simulates the dynamics of ISP peering/transit selection changes given an inter-autonomous system (AS) traffic matrix. We extended the model by adding a spam traffic matrix and selection options for ISPs to adopt the charging mechanism and the outbound filtering. Our goal is to reveal the diffusion process through practical simulation.

A. Simulation Model

Our simulation model is described here, focusing on the primary parts: ISP model, inter-ISP relationship model, inter-ISP traffic model, economic model, and ISP action model.

1) ISP model: In our simulation, we sampled 50 actual ASes from CAIDA AS data [10] as simulation agents. We also used the data on the number of customers in ASes [11].

ASes are categorized into three groups: large ISPs, small ISPs, and stub ISPs, each of which respectively corresponds to Tier-1 ASes, Tier-2 ASes, and edge ASes such as content providers or enterprise customers.

Moreover, as with the ITER model, we classify ISPs into six regions according to their location. An ISP can connect with ISPs that belong to the same region. Large ISPs belong to every region, small ISPs belong to two regions, and stub ISPs belong to one region. Stub ISPs select one small ISP that belongs to the same region as the provider ISP. Small ISPs buy a transit service from two large ISPs as provider ISPs. Small ISPs may choose peer ISPs from small ISPs that belong to the same region in order to cut transit costs. Meanwhile, all large ISPs always have peering links with each other to maintain reachability.

2) Inter-ISP relationship model: In this work, the inter-ISP relationship has two aspects: the connection relationship and the economic relationship (see Fig. 4). The connection relationship refers to an inter-ISP network connected with physical cables. The economic relationships are built on the connection relationships and categorized as a customer-provider relationship or peer-peer relationship.

The initial relationship is created by using CAIDA data [12] and changed by ISP action models described later in the simulation. Furthermore, inter-ISP traffic is forwarded through a shortest path in the relationships as well as by complying with the no-valley routing policy [13].

3) Inter-ISP traffic matrix model: We used two traffic matrices: legitimate traffic and spam traffic. As in [11], the legitimate traffic matrix was created using the gravity model, where the total traffic volume each ISP

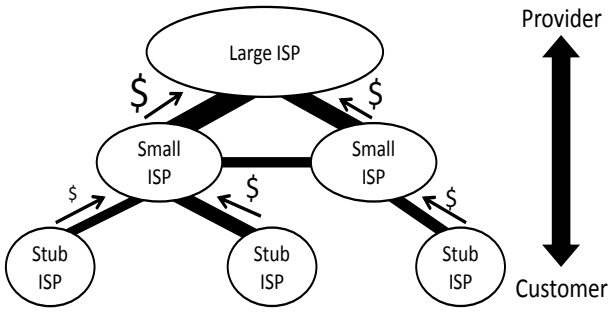

Fig. 4. Inter-ISP hierarchy

sends/receives was assumed to be linear to the number of its customers. The spam traffic matrix was also created with the gravity model, but we used spam data measured at our mail server for the number of spam mails each ISP sends. In the measured data, some ISPs sent a huge number of spam mails and appeared as spammer ISPs in our simulation. This is because a few particular ISPs are connected by a large number of bot-infected PCs [1]. In addition, network locality was considered so that the number of spam mails was doubled if the sender and receiver ISPs were located in the same region.

4) Economic model: The benefit for an ISP, B, is calculated using the following expression.

$$\begin{aligned} B &= T_i - T_p - C_f - S_d + \\ &\quad S_i * I_m - S_p * I_m - C_o * I_o \end{aligned}$$

T_i : Transit income from customer ISPs
T_p : Transit payment to provider ISPs
C_f : Facility cost
S_d : Damage of received spam mails
S_i : Spam charge income
S_p : Spam charge payment
C_o : Cost of installing outbound filtering

$$I_m = \begin{cases} 1 & \text{if the ISP or its provider adopt the mechanism} \\ 0 & \text{otherwise} \end{cases}$$

$$I_o = \begin{cases} 1 & \text{if the ISP adopts outbound filtering} \\ 0 & \text{otherwise} \end{cases}$$

The parameters T_i, T_p, and C_f are the same as in [9], while spam-related parameters S_d, S_i, S_p, C_o, are introduced as follows.

- S_d: damage of ISPs by receiving spam mails, such as ISP's mail server costs, help desk calls concerning spam, and inbound filtering costs[2].S_d is assumed to linear to the number of received spam mails.
- S_i (S_p): income (payment) of spam penalty for receiving (sending) spam mails. The spam penalty is also linear to the number of spam mails, we define spam penalty multiplier as penalty per one spam mail damage.
- C_o: cost of installing outbound filtering. In this paper, C_o is assumed to be linear to the number of ISP's users.

TABLE I
SIMULATION PARAMETERS

Number of ISPs	50
(Large / Small / Stub)	(5 / 20 / 25)
Damage per spam mail	10^{-5} [$]
Outbound filtering cost per user	0.5 [$]
Incentive mechanism cost per user	1.0 [$]
Total number of spam mails	$5.1 * 10^{13}$
Switching cost	10%
Number of regions	6

5) ISP action model: Each ISP selfishly takes the following actions to maximize its own benefit in turn. Each ISP can foresee the next action of its customer's ISPs after it acts. Note that each ISP only changes its state if the benefits are forecast to improve by more than 10%. Thus, an ISP will not take any action if it gains little benefit, which reflects the switching cost.

- Provider Selection (Small ISPs, Stub ISPs)
  Small ISPs and stub ISPs select provider ISPs to maintain reachability. Each ISP randomly determines its provider ISP out of economically near-optimal ISPs.
- Peer Selection (Small ISPs)
  To cut transit costs, small ISPs may choose a peering ISP with cost-benefit peering [9], where both ISPs must see a positive value from the peering link.
- Installing Outbound Filtering (All ISPs)
  All ISPs choose whether to install outbound filtering by comparing its cost as well as spam penalty payments. When an ISP installs outbound filtering, the ISP pay for the cost to install the filtering and the number of spam mails from that ISP decrease to zero.
- Installing Charging Mechanism (Large ISPs, Small ISPs)
  Large ISPs and small ISPs choose whether to install a spam charging mechanism considering the next action of its customer ISP. If the benefit of installing a charging mechanism is higher than the loss caused by its customer ISP switching to another provider ISP, a charging mechanism is installed.

Table I shows the simulation parameters. Note that the relative value of these parameters is important in our simulation.

B. Evaluation results

We ran the simulations by varying the spam penalty multiplier from 10^{-1} to 10^{7}. There were 50 ISPs, and 100 rotations (turns); thus, each simulation had 5,000 steps and was run 5 times. Each ISP selfishly chose actions to maximize its own benefits.

The results are shown from two aspects: macroscopic, such as the total benefit and adoption ratio, and microscopic such as the status of adopting individual ISPs. In addition, we show both transient state and steady state for the macroscopic behavior.

1) Adoption ratio and total benefit: First, we clarify whether or not our proposed mechanism will be diffused. Fig. 5 shows the changes in the adoption ratio of the

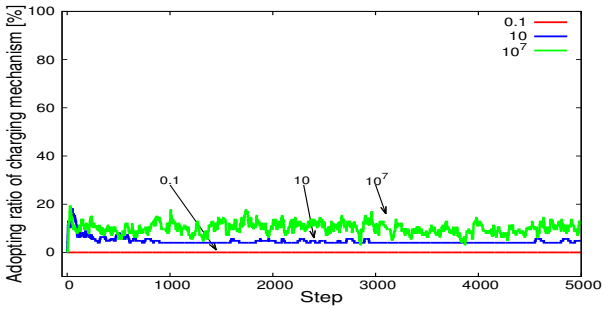

Fig. 5. Changes in adoption ratio of incentive mechanism

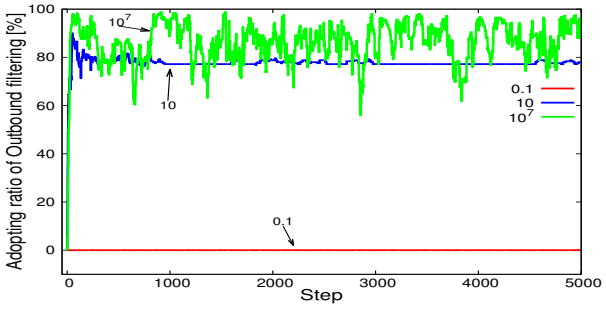

Fig. 6. Changes in adoption ratio of outbound filtering

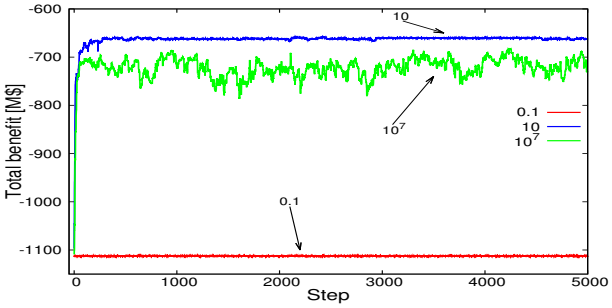

Fig. 7. Changes in total benefit

charging mechanism when the spam penalty multiplier is 0.1, 10, and 10^7. With a spam penalty of 0.1, the adoption ratio does not change at all. This is because the spam penalty is too small to give ISPs an economic incentive for adopting the charging mechanism. Conversely, with a spam penalty of 10^7, the adoption ratio increases to 10% at the end. However, the ratio oscillates quite a bit due to the very expensive spam penalty compared to the switching cost. With a spam penalty of 10, the adoption ratio increases rapidly within the first $1,000$ steps, and then converges at about 5%.

Moreover, we examine which ISP adopts charging mechanism at first. With the spam penalties of 10 or 10^7, first adopter is always LTP in our simulations. We investigate what motivates first adopter's decision to adopt charging mechanism through detailed data of each agent. ALL of first adopter LTP increases its transit income by attracting many STPs to make transit contract. When STPs make transit contract to first adopter LTP, they installed outbound filtering to decrease spam penalty to pay to its LTP. These STPs can increase their benefit by receiving spam penalty from first adopter LTP.

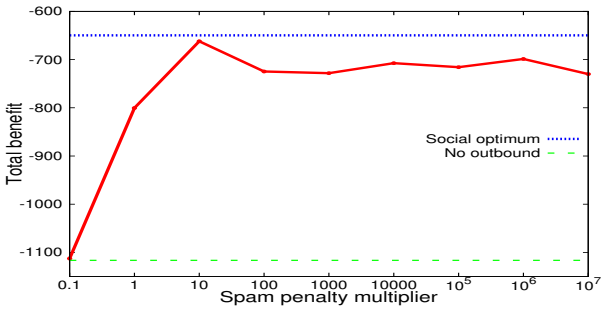

Fig. 8. Spam penalty multiplier vs. total benefit of all ISPs

Next we show how many ISPs adopt outbound filtering as the adoption ratio of our charging mechanism increases. Fig. 6 shows the results. The overall trend is the same with the adoption ratio of the incentive mechanism, but the ratios are high (around 80% to 90%) for spam penalties of 10 and 10^7 . This is because even if a provider ISP does not adopt the incentive mechanism, it will not contract with a customer ISP who does not implement outbound filtering if the upper provider ISP adopts a mechanism to avoid the spam penalty charge. Thus, a provider ISP implementing the mechanism indirectly forces its customer or customer's customer ISP to install the outbound filtering.

Finally, we show how much the total benefit increases with our mechanism (Fig. 7). Interestingly, with the spam penalty of 10^7, the total benefit is not the highest although the adoption ratio of outbound filtering is the highest. This result shows that the higher spam penalty and the higher adoption ratio of outbound filtering do not lead to a higher total benefit.

Fig. 8 plots the relationship between spam penalty and total benefit. Here, total benefit is calculated as the mean values of the last 50 steps, which can be regarded as steady states.

It is also observed that a higher spam penalty does not lead to a higher total benefit and that there is an optimal value of spam penalty. This is because in our simulation parameters, there were some ISPs whose own cost to install outbound filtering was more expensive than the damage caused by the ISPs to other ISPs by sending spam mail. From the viewpoint of total optimization, these ISPs should not install outbound filtering. If the spam penalty multiplier is too high, even these ISPs will install outbound filtering, and as a result, the total benefit for all ISPs will decrease. Here, the 'No Outbound' horizontal dashed line indicates the state where no ISPs have installed outbound filtering, and the 'Social Optimum' horizontal dotted line indicates the state where outbound filtering is only installed by ISPs whose outbound costs are lower than the damage caused to other ISPs by spam that they send. By appropriately setting the spam penalty, we can observe that the proposed mechanism can achieve a near optimal total benefit.

Our findings are summarized as follows: 1) our proposed mechanism is voluntarily diffused despite each ISP selfishly maximizing its own benefits; 2) our proposed mechanism increases the total benefit by stopping spam

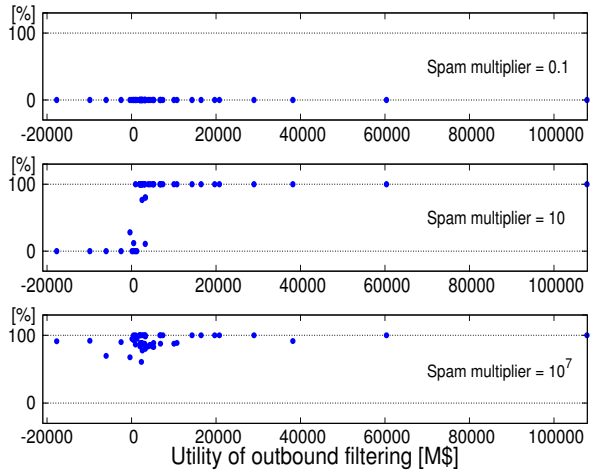

Fig. 9. Adoption ratio of outbound filtering with the utility of outbound filtering for each ISP

mail; and 3) a higher spam penalty does not lead to higher benefits and can result in unstable ISP states. Thus, an optimal spam penalty multiplier can be used to achieve a near total optimal benefit.

2) Utility of outbound filtering vs. adoption ratio of outbound filtering: We next show the adoption ratio of individual ISPs when the spam penalty multiplier is 0.1, 10, and 10^7 times the spam damage multiplier in steady state (Fig. 9). Here, adoption ratio of outbound filtering is calculated as the mean value of the last 50 steps, which can be regarded as steady states. The horizontal axis is 'the utility of outbound filtering', which is the difference between the spam damage reduced by adopting outbound filtering and the cost of the outbound filtering. The ISPs whose utility is positive should install outbound filtering because the total benefit increases, and vice versa.

In Fig. 9, with a spam penalty of 0.1, almost none of the ISPs install outbound filtering. With a spam penalty of 10, the ISPs whose utility of outbound filtering is largely negative do not install outbound filtering, while those whose utility of outbound filtering is positive do install it. With a spam penalty of 10^7, even the ISP with the lowest utility level installs the outbound filtering, which leads to a decrease in the total benefit, as shown in (Fig. 8).

In this simulation, although all the ISPs adopt outbound filtering, the total benefit cannot reach the largest value, but when all the ISPs whose utility of outbound filtering is positive adopt outbound filtering and the other ISPs do not, the total benefit can reach the largest value. Even in such a situation, our incentive mechanism can nearly achieve an optimal total benefit. This suggests that an economic approach is more efficient than the institutional approach of forcing all players to obey, such as through legal means.

V. Conclusion

The objective of our study was to solve incentive problems in implementing outbound filtering of unwanted traffic. We proposed an incentive mechanism where ISPs pay a penalty for sending unwanted traffic to peering ISPs. This payment was combined with payments of transit fees so that the mechanism will be deployed through the current ISP relationship. We evaluated our proposed mechanism through simulations where individual ISPs behaved in a selfish manner. We showed that even with that behavior, many ISPs will spontaneously adopt our proposed mechanism. Accordingly, outbound filtering was also widely installed. Our proposed mechanism nearly achieved cost optimization. This is because it provides strong incentives to install outbound filtering, particularly to spammer ISPs. The results suggest that this approach to incentivizing an individual player is cheaper than the institutional approach of forcing all players to obey, such as through legal means.

In the future, we plan to further elaborate the simulation model. Specifically, in respect to simulation parameters, we will consider the case in which an inter-ISP spam traffic matrix and spam charging multiplier have not changed throughout the simulation.

References

[1] MessageLabs Intelligence: 2010 Annual Security Report. http://www.messagelabs.com.
[2] "Enisa 2009 spam survey," European Network and Information Security Agency(ENISA), Dec 2009.
[3] "Worldwide Infrastructure Security Report," Arbor Networks, vol.7, January 2009.
[4] M. Parameswaran and A.B. Whinston, "Incentive Mechanisms for Internet Security," *Information Assurance, Security and Privacy Services*, pp.101-142, 2009.
[5] Bonded sender program, http://www.bondedsender.com.
[6] S. Goel, J. Hofman, J. Langford, D. M. Pennock, and D. M. Reeves. Centmail: Rate Limiting via Certified Micro-Donations. In *Proc. of CEAS*, 2009.
[7] H. Kim, K. Claffy, M. Fomenkova, D. Barman, and M. Faloutsos. Internet Traffic Classification Demystified: The Myths, Caveats and Best Practices. In *Proc. of CoNEXT08, Madrid*, Dec. 2008.
[8] J. Armin, "McColo - Cyber Crime USA," HOSTEXPLOIT.COM. 2008. available at http://hostexploit.com.
[9] A. Dhamdhere and C. Dovrolis, "An Agent-based Model of the Internet Ecosystem," *Invited paper at COMSNETS*, Bangalore, India, January 2009.
[10] The CAIDA Autonomous System Taxonomy Repository. http://www.caida.org/data/active/as_taxonomy/
[11] H. Chang and S. Jamin, Z. Mao, and W. Willinger, "An Empirical Approach to Modeling Inter-AS Traffic Matrices," In *Proc. of Internet Measurement Conference*, Oct 2005.
[12] The CAIDA AS Relationships Dataset, 20100120. http://www.caida.org/data/active/as-relationships/
[13] L. Gao and F. Wang, "The extent of AS path inflation by routing policies," In *Proc. of IEEE Global Internet Symposium*, Nov 2002.