Using Insurance to Increase Internet Security

Svetlana Radosavac, James Kempf and Ulaş C. Kozat
DoCoMo Communications Laboratories USA, Inc.
Palo Alto, CA 94304
{sradosavac,kempf,kozat}@docomolabs-usa.com

ABSTRACT

Managing security risks in the Internet has so far mostly involved methods to reduce the risks and the severity of the damages. Those methods reduce but do not eliminate risk, and the question remains on how to handle the residual risk. Current schemes applied by Internet Service Providers (ISPs) penalize the users, who suffer from the consequences. In this paper, we take a new approach to the problem of Internet security and advocate managing the residual risk by buying insurance against it and consequently re-arranging the incentive chain. We first analyze the current state of the Internet and investigate if it is possible to alleviate the existing problems by introducing insurance schemes. By performing detailed analysis we define an insurance policy that can survive in a competitive market. Following that, we analyze the impact of insurance-based ISPs on the rest of the network and attempt to answer whether using insurance can increase the overall security of the system and provide incentive to other ISPs to implement such policies.

Categories and Subject Descriptors

C.2.0 [**Computers-Communication Networks**]: General—*Security and Protection*

General Terms

Economics, Security

Keywords

DDoS, security, insurance, economy, incentive, correlated risk, risk transfer

1. INTRODUCTION

The Internet has become a fundamental part of life during the last decade and it has become of essential value to companies as well as to individual users to maintain stability of services that we rely upon on a daily bases. More than one billion people use the Internet and critical industries like banking heavily rely on it. However, the Internet was built under assumptions that don't hold any more: that all users of the network can be trusted and that the computers linked by the Internet are fixed objects. Hence, the Internet lacks inherent security architecture. Protections like firewalls and antispam software are add-ons and can be considered only as patches used until a real solution is found. The Internet has become just like real world: both good and malicious individuals have access to it. However, unlike the real world, it has become increasingly difficult to identify and trace Internet users. As a consequence, malicious individuals have strong incentives to shift their illegal activities to the Internet, where they can access more people in a shorter time period, while minimizing their chances of being discovered. We are now faced with the situation where the Internet's security problems are getting worse and at the same time society's dependence on it is deepening. Due to the current state of the Internet architecture, only the target (i.e. target of Distributed Denial of Service (DDoS) attack) bears the cost of the attack. Neither the infected users (that are actually responsible for the attack) nor the ISPs bear any of the cost and therefore do not have any short term incentive to invest in security measures. This situation results in the following paradox: it is widely accepted that decreasing the overall number and intensity of attacks will be beneficial to both individual users and e-businesses given the huge loss these attacks cause; on the other hand, organizations are reluctant to establish the defense given the costs they impose for their implementation.

One problem that heavily impacts Internet security is that ISPs focus only on inbound traffic control and such methods have limited effectiveness. Attack traffic has already traversed multiple domains and wastefully consumed network resources by the time it hits the targeted ISP's domain. Many of the attacks that originate from a single domain rapidly branch out toward many targets, making it much more difficult to control at the destination rather than at the source. So far, inbound traffic control has not been effective in prevention of large-scale attacks (such as DDoS attacks). The best example is handling of email spam. By applying new email filters, the amount of spam has significantly increased since spammers now have to increase the volume of junk email if they want to increase the probability of delivering their messages. To improve Internet security, it is essential that service providers control *both* outbound and inbound traffic. Outbound traffic control detects attacks at the source and thus minimizes the probability of spreading,

NetEcon'08, August 22, 2008, Seattle, Washington, USA.

without subjecting the network to congestion. Outbound control is especially effective when done by ISPs, which can leverage the direct relationship with their customers to hold them accountable and take punitive action against violators.

We can now see that one of the main problems of the current Internet is that the end users bear the complete cost of the attacks, while the ISPs or infected users do not bear any consequences. Consequently, this setting provides no incentive to the participating parties to invest into their own security. However, even if a specific user or a set of users that belong to an ISP invests into their own security, the security risks are not eliminated due to the fact that each user needs to interact with numerous users who may implement different security measures.

In this work, we build up on the initial risk analysis provided in [1] and propose a novel approach of handling the risk by re-arranging the economic incentives and transferring some part of the cost of attack to all involved parties (in the current system only the attack target bears all the cost). We analyze several different insurance scenarios for prevention of attacks. We attempt to answer three basic questions: (i) whether implementing insurance schemes in the current Internet architecture is feasible for loss compensation; (ii) how implementation of insurance schemes can help in mitigation/prevention of attacks and (iii) what is the effect of security measures implemented by one ISP (or a group of ISPs) on the security of other ISPs and the overall system security.

The paper is organized as follows. In Sect. 2 we present our system model and assumptions that will be adopted for the remainder of the paper. In Sect. 3 we first introduce basic definitions related to the notation adopted from economics. Following that, in Sect. 3.1 and Sect. 3.2, we offer a more detailed model of ISP's policies followed by a description of a model used for risk analysis in economics. In Sect. 3.3 we investigate how separating users based on their risk susceptibility and offering different policies to each type of users affects ISP's profits and user behavior. In the conclusion of that section we investigate whether the strategy of offering different policies to different classes of users can survive in a competitive market and examine the resulting scenario and possible outcomes. Guided by a model provided in [3], in Sect. 4 we investigate how security (or lack of security) of a specific ISP affects the overall security. Finally, Sect. 5 concludes our study.

2. SYSTEM MODEL

In this work we assume the existence of two entities: the ISP and users. The goal of both ISPs and users is to maximize their gain while minimizing their losses. We assume that the users are aware of the risks involved when they interact with other users and would like to insure themselves and minimize their own losses. On the other hand, the main goal of ISPs is to make a profit. The question is whether a policy that brings profit to the ISPs while protecting the users from risks exists. In this work we aim to provide a framework by using insurance mechanisms, where the ISPs offer certain types of insurance to the users in exchange for certain level of insurance premiums (insurance premiums will be offered by ISPs, not by insurance companies). We assume the existence of two types of users: *high* and *low* risk users, where the terms "high" and "low" define the probability that a certain user will seek a payment from the insurer. More specifically, a high risk user is more likely to ask for an insurance premium payout than a low risk user. In the subsequent sections, the terms "high risk" and "unsafe", "low risk" and "safe" will be used interchangeably with the same meaning.

We assume that each user has wealth w as a result of his Internet connectivity and activity. When this wealth is not insured, there exist two possible outcomes for the user. If he doesn't suffer any damage, his wealth will remain equal to w and his utility will be $U(w)$. On the other hand, if he suffers damage d, his wealth will be reduced to $w-d$ and the resulting utility will be $U(w-d)$. His expected wealth is then given by:

$$E(w) = p(w-d) + (1-p)w \tag{1}$$

and his expected utility is

$$EU(N) = pU(w-d) + (1-p)U(w) \tag{2}$$

where N in $U(N)$ stands for utility when no insurance is offered and p represents the probability of damage occurring.

Now consider the case with insurance offered. Assume that an individual purchases an insurance policy at price α_1. Hence, the initial wealth of a user is equal to $w-\alpha_1$. In the case of an attack, the insurance company pays out to the user the amount of money equal to α_2 and consequently the resulting wealth of an insured individual after the accident is equal to $w-\alpha_1-d+\alpha_2$. The user's expected utility in this case can be expressed as:

$$EU(I) = pU(w-\beta) + (1-p)U(w-\alpha_1) \tag{3}$$

where $\beta = \alpha_1 + d - \alpha_2$ and I in $U(I)$ stands for utility when insurance is offered. Looking at the above equation, we note that the payout insurance premium α_2 has to be a function of both the insurance premium α_1 and the probability of claiming insurance by individual users, p. Furthermore, the following notation is adopted for the remainder of the paper:

w_1 : final wealth of the user without attack

w_2 : final wealth of the user after the attack

We assume a user will have incentive to buy an insurance policy if the expected utility of being insured exceeds the expected utility of being uninsured. Our initial assumption in this work is that ISPs do not implement any kind of outbound traffic control. The only type of traffic control implemented is the standard inbound traffic control.

3. EFFECTS OF INSURANCE ON SECURITY

In order to proceed towards our analysis we first introduce some basic notation that will be used in the remainder of the paper.

Definition 1. **Insurance policy:** a contract of insurance, describing the term, coverage, premiums and deductibles.

More specifically, an insurance policy represents a set of payment and compensation rules enforced between the buyer and provider of the policy

Definition 2. An **insurance contract:** the set of rules under which the features of an insurance policy are enforced.

We also note here that vector $\alpha = (\alpha_1, \alpha_2)$, defined in Sect. 2, can be used to describe an insurance contract.

Definition 3. An **insurance premium:** the periodic payment made on an insurance policy (an amount of money a user pays to an insurance company regardless of whether he/she had an accident).

The questions we are interested in are if there exists a policy that encourages good behavior and whether or not it is possible to enforce such policy by regulatory dynamics.

In this paper we adopt the equilibrium definition from [4]. We assume that in equilibrium:

1. No insurance policy makes negative profits over all;
2. Insurers are assumed to have sufficient financial resources so that they are willing and able to sell a contract they expect to be profitable;
3. If there were a potential policy that could be offered that would be more profitable than the current policy offered in equilibrium, then the existing policy is suboptimal.

In the remainder of the paper we assume that both the users that access services through ISPs and the ISP have the goal of making a profit, while minimizing the risks involved. Also note that in the case of the user, the goal is to minimize the decrease in initial wealth w.

3.1 ISP insurance policies

It has already been mentioned in Sect. 1 that a user cannot eliminate the risk by only protecting himself. This is partially due to the fact that new threats appear and propagate with high speed and partially due to the fact that both ISPs and users interact with each other and thus they are highly dependant on each other's conditions. Therefore, it is not only important to find an optimal insurance policy that insures a specific class of users belonging to an ISP. It is essential to analyze how security of one ISP impacts security of ISPs it interacts with and vice versa.

To maintain simplicity, we consider only 2 types of users: high and low risk. A user is classified as either low or high risk depending on several factors, such as profitability of its business (more successful business is more likely to be a target), publicity of the user (e.g. a highly active political organization vs. local charity group), whether or not the user deals with sensitive and important data, etc. More specifically, the users are defined as:

High risk (H) : with probability of claiming insurance p_h

Low risk (L) : with probability of claiming insurance p_l

where $p_h > p_l$.

We have already stated that by introducing insurance, part of the risk is being transferred to the ISP. We assume that in the case when attacks happen, the ISP compensates for the damages of users who purchase insurance policies. The main goal of the ISP is to make a profit and in the remainder of this section we investigate insurance policies that can be offered by ISPs. Each insurance policy attracts a certain portion of low and high risk users. In the remainder of this section we investigate (i) if there exists a policy that is acceptable to users (brings satisfying level of compensation for an acceptable insurance premium) and ISPs (brings them profit) and (ii) if such policy exists, can it survive in the competitive market (i.e. is it stable).

For clarity purposes we now briefly present an overview of risk analysis method widely used in economics.

3.2 Risk modeling

In addition to the above information, we assume that all users know their own risk type p_i, but this information is not available to the insurance companies. Instead, the insurance company assumes that all users belong to the same risk category and will be claiming insurance with probability p, where $p_L < p < p_H$. This setup is more realistic because users in general know more about their risk type than the insurance companies. This claim is true even in the case of uneducated users. Namely, even though they do not know how insecure they are, they are aware that they are not using any security measures to protect themselves from becoming a victim.

In order to explain some basic notation that arises from economy and is used in the remainder of the paper, we analyze the setting presented in Fig. 1. The x-axis represents the wealth of the user before the attack and the y-axis represents the wealth of the user after the attack. Point E in Fig. 1 represents the *endowment* point. In case of no loss a user remains with wealth w (x-axis) and in case of attack a user remains with wealth $w - d$ (y-axis) at point E.

Curves U_H and U_L represent *indifference curves* for high and low risk users respectively. Namely, all points at U_H (U_L) yield the same utility for high (low) risk users and as a consequence a user is indifferent between the choices that lie on the same curve. The slope of indifference curves represents the MRS (Marginal Rate of Substitution, i.e. rate at which consumers are willing to substitute one good for the other). Both types of users have the same preferences but their indifference curves have different slopes at any point in the state space diagram since they face different probabilities of claiming the insurance premiums. The line MRS_L in Fig. 1 represents the market average fair odds line (or fair odds line) [2]. The market average fair odds are the odds that an insurer could offer to the average customer while breaking even on average as long as the contract was accepted by a random sample of both types of users. The optimal operating point (from the point of *expected* utility) is where the indifference curve is tangent to the fair odds line [2]. In the case of low risk users, the optimal operating point is point A in Fig. 1.

If we assume that each \$1 of insurance costs q and that a units of insurance are bought, the insurer's goal is to choose the optimal level of a. Following the notation from Sect. 2 for the case when insurance is purchased, we obtain the following equations:

$$w_1 = w - aq \tag{4}$$

$$w_2 = w - aq - d + a \tag{5}$$

and the expected wealth is

$$\begin{aligned} E_i[W] &= (W - aq)(1 - p_i) + (W - aq - d + a)p_i \\ &= W - p_i d + a(p_i - q), \text{ where } i = \{L, H\}. \end{aligned} \tag{6}$$

We say that the price q of one unit of insurance is *fair* if it is equal to the expected cost of insurance, i.e. $q = p$. An individual is *completely insured* when $a = d$. When full insurance is purchased an individual's final wealth is $w - pd$, regardless of the occurrence of the loss, i.e. $w_1 = w_2$. This setting is represented with "full insurance line" in Fig. 1. Due to the fact that high risk users are aware of their risk category it is reasonable to assume that they will always seek to obtain a complete (full) insurance since this policy com-

pensates them for their losses at all times. On the other hand, full insurance does not create any incentive for high risk users to invest into their own security and the optimal policy for the ISP should avoid offering full insurance to high risk users in order to make a profit. We will see that this assumption holds after further analysis in Sect. 3.3. After solving Eq. 4 and Eq. 5 by eliminating parameter a and obtaining the relationship between w_1 and w_2 it is easy to show that the slope of the fair odds line depends on probabilities of claiming insurance p_h (p_l). It is known that the insurer is

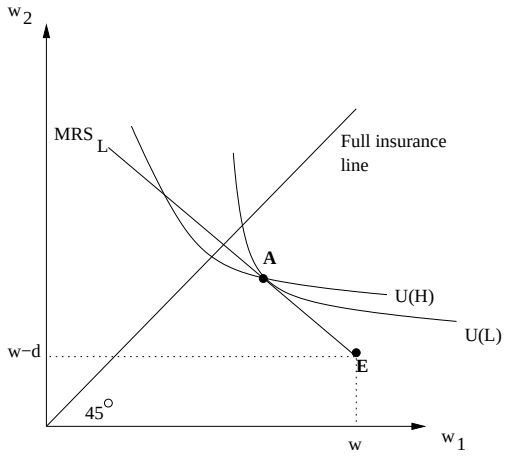

Figure 1: Graphical representation of risk modeling.

driven by market demand to offer the policy that optimizes the welfare of the low risk customers [2]. This policy is represented by point A in Fig. 1. Any contract below MRS would offer extra profits to the insurer if it could attract both types of customers. This kind of contract cannot be an equilibrium since competition would drive the contract to improve until it again reaches a point that lies on MRS. If an insurer offers a contract to the right of A along MRS, the contract could always be improved by another insurer offering a contract at A since both risky and non risky customers prefer that contract. Similarly, if an insurer offers a contract on the left side of A along the MRS, that contract could always be improved by offering a contract at point A (only the safe customers will prefer the contract A and thus it would attract all the safe types, with all the unsafe types remaining at the point at the left side of A). Therefore, a Nash equilibrium cannot exist in a setting where the same type of insurance is offered to all types of users. A more detailed analysis of this setting is presented in [2] and is beyond scope of this work.

Looking at this setup from a business point of view, the result makes sense. If a company loses money on one group of users and gains money on the other, there is a strong incentive to separate the two groups and charge different prices for insurance, which brings us to the notion of a separating equilibrium, where each risk type buys a different policy.

3.3 Each risk type buys a different policy

It has been shown in [2] that the setting when all the users are offered the same policy is infeasible as soon as an informed insurer enters the market, resulting in strict separation of low and high risk users. Now, if one or more ISPs decide on a policy where they offer fixed insurance premiums for all users, they are inevitably going to attract only the high risk users. This brings us to the following scenario presented in Fig.2. In this scenario, points A^L and

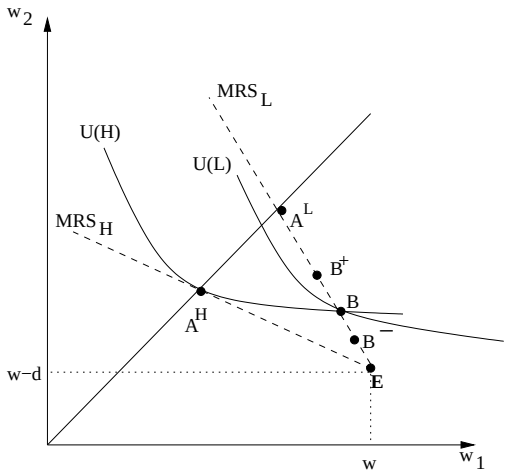

Figure 2: Different risk types are offered different policies.

A^H are the full-insurance points for the two risk groups. Group L has higher wealth because its odds of experiencing a loss are lower. Note the point labeled B on the MRS_L, which represents the point where the indifference curve from the full-insurance point for the high risk group crosses the MRS_L. We first emphasize that B is the best policy that can be offered to low risk users that would not also attract high risk users. If an ISP offered another policy, say B^+, low risk users would strictly prefer it. However, the high risk users would also prefer this policy, bringing us to the scenario from Sect. 3.2, i.e. the non-sustainable scenario. If an ISP offered the policy B^-, high risk users would not select it, but low risk users would strictly prefer the original policy (B). Hence, *any* policy like B^- is dominated by B. So, B is the point that defines the separating constraint for low and high risk users. Any policy that is more attractive to high risk users would result in the scenario from Sect. 3.2.

We now have the scenario where the ISP offers two types of policies: A^H and B, where A^H is the best policy for high risk users and B is the best policy for low risk users. The ISP offers two types of policies, but does not enforce access control, i.e. a user is free to choose his own policy. Since each user is aware of his risk factor, high risk users are going to choose A^H, while low risk users choose B. Now, we see that with this insurance scenario high risk users are fully insured and low risk users are offered partial insurance (note that if a company offered a policy that fully insured the low risk users, it would also attract the high risk users). Hence, preferences of high risk users act as a constraint on the market. The insurance companies must maximize the well-being of low risk users subject to the constraint that they do not attract high risk users. The question that we need to answer is whether the proposed policies are in equilibrium. To answer that question we observe Fig. 3. We first consider the scenario where the market fair odds line, M_1, lies below the low risk user's indifference curve through C. In this case, any contract capable of attracting low risk users away from C would also attract high risk users from A and lie above the market average fair odds line, MRS_H thus introducing

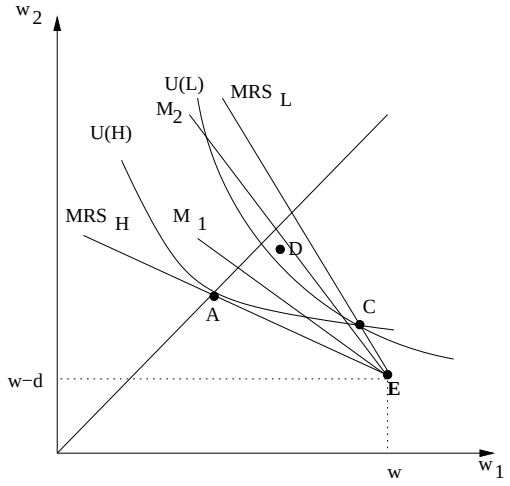

Figure 3: Existence of equilibrium when users are offered different policies.

a premium below the market average fair odds premium and producing expected losses for the insurer. An insurer faced with competitors offering the separating contracts could do no better than to offer those contracts and can find no other contract to offer which produces supernormal profits; the separating contract therefore represents Nash equilibrium.

The question is what would happen if the market fair odds line was at M_2? In this case the market fair odds line cuts the low risk user's indifference curve at point C. This scenario may arise in the case when there exists a higher proportion of low risk users in the market. If the indifference curve and market fair odds line cut in this way, it is always possible to find a new contract to offer that is capable of attracting both high and low risk customers away from the separating contract. We denote this contract as D in Fig. 3. Since D lies above the indifference curves for low and high risk users, the contract attracts both types of customers away from the separating contracts. Also, since D lies below M_2, the contract charges a premium higher than the market average fair odds premium, thus yielding positive expected profits to the insurer. An insurer faced with competitors offering the separating contracts will not maximize profit, given the actions of his competitors, by offering separating contracts, but will do better to offer the contract allowing customers to locate at point D. The separating contracts, therefore, do not produce a Nash equilibrium in this case. The contract located at point D is the same one as we have already analyzed in Sect. 3.2. We saw that no such contract ever produces a Nash equilibrium. It follows then that no Nash equilibrium exists in the latter case. We now offer an intuitive explanation of the above result. At the separating equilibrium, the low risk users are not fully insured and they are unhappy because of such a setting. A policy like D that requires just a little cross-subsidy to high risk users but offers more insurance is preferred by low risk users to policy C. Hence, if there are sufficiently few high risk users in the market, an ISP could profitably offer this policy and it will dominate the two separating policies. We now see that low risk users prefer more insurance at unfair price to less insurance at fair price. However, the market cannot tolerate this scenario (Sect. 3.2).

We can now see that even in the case when different insurance policies are offered to low and high risk users, a profitable business policy for an ISP may not exist. Additionally, due to the nature of Internet, different classes of users belonging to different ISPs frequently interact, which makes it more difficult for an ISP to properly assess user's security properties. In the remainder of the paper we present how interactions among users with different security properties belonging to different ISPs affect the overall network security and contribute to the failure of the existing ISP business model.

4. EFFECTS OF SECURITY MEASURES

It has been mentioned in the previous section that the overall security depends on the security of individual ISPs as well as on interactions among various ISPs (interactions among users belonging to different ISPs). In this section, applying the model presented in [3], we analyze the network scenario where a certain percentage of ISPs adopts insurance-based policies and the remaining percentage of ISPs employ no control over its users. We examine whether such a mixture of ISPs can exist, the impact of secure ISPs on non-secure ones and vice versa.

4.1 System model

We assume a setting with a network that consists of N ISPs. Each ISP chooses whether it wants to enforce insurance-based policies or not. As a consequence, the resulting network consists of N_1 ISPs that enforce insurance-based policies and N_2 ISPs that do not enforce such policies. We represent the network using a weighted directed graph $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$ that consists of nodes $\mathcal{N}$ and edges $\mathcal{E}$. Each node $\mathcal{N}$ represents an ISP and we assume that two nodes in the graph are connected with an edge if security of one ISP affects the security of another one. The set of nodes has N elements, where $N = N_1 + N_2$. There exists an edge e_{ij} (i,j=1,...,N) between ISP$_i$ and ISP$_j$ if the security of ISP$_i$ impacts the security of ISP$_j$. Taking into account the fact that levels of interaction between various ISPs are different, it is necessary to associate *weight* w_{ij} with each edge e_{ij}. In our setting, weight w_{ij} represents the impact the security of ISP$_i$ has on the security of ISP$_j$. The weight $w_{ij} < 0$ if i decreases overall security of j and $w_{ij} > 0$ if i increases security of j.

Depending on its own resources, each ISP$_i$ decides on the level of security investment x_i, where security investment can be outbound traffic control, insurance for the users, firewall implementation, etc. The investment levels of all ISPs can be represented in vector form as $\mathbf{x}$. More specifically, we can now say that if the security of ISP$_i$ affects the security of ISP$_j$ and ISP$_i$ invests x_i into its own security, ISP$_j$'s security is increased by $w_{ij}x_i$, where w_{ij} represents the impact level of i's security on j's security. Furthermore, a *weight matrix* $\mathcal{W}$ is defined as

$$\mathcal{W}_{ij} = \begin{cases} 1 & \text{if } i = j \\ w_{ij} & \text{if } \exists e_{ij} \\ 0 & \text{otherwise} \end{cases}$$

Now, $\mathbf{W}\cdot\mathbf{x}_i$, where $\mathbf{W} = \mathcal{W}^T$, represents the total security effects of all ISPs in the network on ISP$_i$. The total utility function of ISP$_i$ is now represented as:

$$U_i(\mathbf{x}) = V_i((\mathbf{W}\mathbf{x})_i) - c_i x_i \qquad (7)$$

for function $V_i(.)$ and $c_i > 0$. $c_i x_i$ is the cost of implementing security mechanisms by ISP_i [3]. Consequently, the utility $U_i(\mathbf{x})$ can be described as the total benefit by ISP_i resulting from (i) its own security investments and (ii) security investments of other ISPs. It is now easy to observe from Eq. 7 that $U_i(\mathbf{x})$ of ISP_i can increase (decrease) depending on interactions with other ISPs.

In our model, each ISP makes investments in its own security and optionally employs insurance-based policies for its users. Although both types of ISPs can have a certain portion of high risk users, we assume that insurance-based ISPs are safer since they take additional measures to detect and isolate non-secure users due to the fact that they have the incentive to marginalize the risk. Additionally, we assume that ISP_i's security is determined by:

1. Frequency of interactions between i and j, where $j \neq i$;
2. Frequency of interactions of ISP_j with ISP_k, $k \neq i, j$. If ISP_j has a tendency to interact with high risk ISP's, we assume this interaction impacts the security of ISP_j in negative way. Consequently, if ISP_i and ISP_j interact frequently, ISP_i's security is decreased.

By using the above rules, depending on the properties of the involved ISPs, we can now define the function $V_i(.)$ from Eq. 7.

The question that we attempt to answer is whether there exists a Nash equilibrium, i.e. state of the system where no user has an incentive to unilaterally deviate. In our case, we seek a scenario in which no ISP would be any better off in case it invests more in its security or decreases its security measures than it is in the current state. By following the results obtained in [3] we know that if $\mathbf{W}$ is strictly diagonally dominant (i.e. if $\mid \mathbf{W}_{ii} \mid > \sum_{j \neq i} \mid \mathbf{W}_{ij} \mid \forall i$), then the given game has a unique Nash equilibrium. Therefore, there is no unique answer to the question on whether insurance-based scheme can exist in a competitive market if such mechanism is not enforced by all ISPs. It is now reasonable to assume that the answer will depend on several factors, such as overall security of ISPs that do not offer insurance (which depends on the percentage of high and low risk users using such ISPs), interactions with other ISPs, etc. In general, we can assume the following outcomes:

1. An insurance-based ISP determines that it has to invest significantly more in security when interacting with a specific ISP that doesn't offer insurance and decides it would be better off blacklisting such ISPs; this results in increase of its own security and decrease of security investments. In this case the ISP that does not offer insurance policies can now decide whether it wants to employ such policies and re-establish its relationship with insurance-based ISPs or not.
2. Insurance-based ISP determines that its connections with ISPs that do not use insurance are too valuable (i.e. it has a lot of customers that interact with this type of ISPs) and decides to invest more in its own security since the overall benefit is larger than when no interactions exist. In this scenario, users belonging to the non-insurance based ISP "free ride" on the security measures employed by the insurance-based ISP.

We now note that if the insurance-based ISP decides to cut the connections with ISPs that do not use insurance, the security of the whole network decreases since the positive impact of the insurance-based ISP disappears. Consequently, the rest of insurance-based ISPs will have to either invest more in their own security or cut the connections with ISPs that do not use insurance. In this setting, two possible outcomes arise: (i) two separate networks are created: one with only insurance-based ISPs and the other one with non-insurance based ISPs or (ii) ISPs that do not use insurance decide to switch to insurance-based policies if they find the interactions with insurance-based ISPs to be too valuable for them.

5. CONCLUSIONS AND FUTURE WORK

In this work we present a preliminary analysis and results for regulating behavior of users and ISPs in the Internet by imposing insurance schemes and transferring part of cost of attacks to all parties involved in the exchange of traffic. We obtain an optimal insurance policy that can be offered to both low and high risk users and analyze its profitability in a competitive and dynamic market. Guided by an intuitive assumption that security of one ISP affects the security of other ISPs in either positive or negative manner, we model the interactions between various ISPs by using directed graphs and offer initial intuitive analysis of possible scenarios resulting from such interactions.

This work represents a preliminary look into a very complicated issue of fixing the current Internet architecture and proposes a novel, incentive based, method for prevention of attacks and increasing overall system security by using insurance. Two important conclusions arise from this work: (i) even when insurance is used for elimination of residual risk, it cannot be guaranteed that a profitable business model for an ISP will exist and (ii) security of an ISP does not depend only on the level of investments into its own security. It significantly depends on the level of interactions with other ISPs, which may be either secure or insecure. The first issue potentially leads towards proposing alternative architectures that enable complete elimination of residual risk under all conditions. The second issue leads to the conclusion that even in the case when risk can be transferred locally, there may not exist a global equilibrium if interactions with other ISPs significantly decrease the security of insurance-based ISP.

6. REFERENCES

[1] J. Bolot and M. Lelarge. A New Perspective on Internet Security using Insurance. In *Proc. of the IEEE Infocom*, pp. 1948-1956, April 2008.

[2] B. Hillier. *The Economics of Asymmetric Insurance.* Palgrave Macmillan, 1997.

[3] R. A. Miura-Ko, B. Yolken, J. Mitchell, and N. Bambos. Security decision-making among interdependent organizations. 2008.

[4] M. Rothschild and J. Stiglitz. Increasing risk: I. A Definition. *Journal of Economic Theory*, 2(3):225–243, 1970.