Game-theoretic Patrolling Strategies for Intrusion Detection in Collaborative Peer-to-Peer Networks

Pratik Narang, Kunal Mehta, Chittaranjan Hota
Department of Computer Science & Information Systems
Birla Institute of Technology and Science-Pilani, Hyderabad Campus
Hyderabad, A.P., India
Email: {p2011414, f2010425, hota}@hyderabad.bits-pilani.ac.in

***Abstract*—This work studies the problem of optimal positioning of Intrusion Detection Systems (IDSs) in a collaborative Peer-to-Peer (P2P) environment involving a number of peers and super-peers. This scenario applies to network architectures like that of Gnutella, Skype or Tor, which involve a huge number of leaf-peers and a selected number of super-peers who have higher responsibilities in the network.**

A malicious entity may become part of the P2P network by joining from any part of the network. It can attack a super-peer (through APTs, packets with malicious payloads etc.) and thus disrupt the P2P network. Peers may try to secure the network by running IDSs at strategic locations in the network. But a deterministic schedule of running and positioning the IDSs can be observed and thwarted by an adversary. In this paper, we explore the problem of optimally positioning IDSs in a P2P network with a randomized, non-deterministic, game-theoretic approach. Our approach distributes the responsibility of running the IDSs between the peers in a randomized fashion and minimizes the probability of a successful attack.

I. Introduction

Peer-to-Peer (P2P) networks are overlay networks consisting of distributed collection of autonomous end-system computing devices called 'peers'. The 'peers' form a set of interconnections to share and mobilize resources (such as content, storage, bandwidth, CPU cycles etc.) such that peers have symmetric roles in the overlay for both message routing and resource sharing [1]. The construction of P2P networks is on the top of IP layer, typically with a decentralized protocol which allows 'peers' to share resources. The ease of sharing resources – whether in the form of music and videos (BitTorrent) or computing resources (SETI@ home project) – has led to the runaway success of P2P applications. The P2P paradigm has also seen wide deployment for IPTV (LiveStation) and Voice-over-IP based services (Skype).

We consider the scenario of a P2P network such as that of Gnutella or Skype[1] which involves a super-peer (aka ultra-peers) architecture and super-peers hold higher responsibilities and/or privileges in the network. It is also applicable to networks like Tor which bear resemblance to P2P networks. Tor uses 'relay nodes' which can be equated to super-peers since they take the responsibility of routing and relaying the traffic in the Tor network.

A P2P network lacks any centralized authority, and is thus more vulnerable to security threats than the traditional client-server architectures. Although the decentralized and distributed nature of P2P network offers resilience towards network-breakdowns, the super-peer architecture is more sensitive in this regard since an adversary can disrupt (albeit not break-down) the P2P network by attacking a super-peer node. For example, a Denial of Service (DoS) attack targeted on certain relay nodes in Tor can lead to an increased latency and higher time-outs in the network. In traditional networks, such a scenario can be secured by use of Intrusion Detection Systems (IDS), which might be deployed at the backbone router of an enterprise (Network-based IDS or NIDS) or at each end-host (Host-based IDS or HIDS). Owing to its decentralized and distributed nature, a NIDS is not feasible in P2P networks. Furthermore, a HIDS-based solution will provide security only to the node running the IDS, and not to the network. Although a solution based on Distributed IDS can be explored, self-interested peers may not want to spend their resources in running an IDS.

Since the P2P architecture is inherently about peers coming together to share and mobilize resources, we utilize this strength of P2P architecture in its favor. In this work, we propose a novel approach of monitoring and securing the P2P network through Distributed IDSs installed on participating super-peers. Super-peers are required to run the IDS only for a small chunk of time in a given time-slice, thereby reducing the load on any single participating peer. This scheduling is done with a game-theoretic approach to minimize the probability of a successful attack (or equivalently, minimize the *payoff* of the attacker).

II. Related Work

Super-peer based P2P networks such as Gnutella have been shown to be susceptible to query-flooding based DoS attacks [3]. Traffic analysis attacks on Skype Voice-over-IP (VoIP) calls, which compromise the privacy of Skype calls, have been explored using application-level features extracted from VoIP call traces [4]. Past research has also demonstrated attacks on Tor where anonymity in the Tor network can be compromised by traffic-analysis attacks by a global passive adversary [5] or by non-global adversaries with minimal resources [6].

Peer-to-Peer network haves been studied from the Game-theoretic perspective mainly with regard to incentives for sharing [7] and collaboration [8], managing trust [9], etc. Although security in P2P networks *per se* has not received much attention from a Game-theoretic perspective, the topic

[1] Skype has now moved to a cloud-based architecture [2]

of network security with Game theory has attracted a lot of attention. Work in [10] provides a good survey of the same.

Authors in [11], [12] consider optimal resource allocation by a defender in a network against potentially malicious packets by adopting a game-theoretic approach of inspecting only a fraction of all packets. The work of [11] was limited to a single source and single target, whereas authors in [12] considered multiple targets.

Janakiraman et al. in [13] presented a collaborative, Peer-to-Peer approach for building a distributed, scalable Intrusion Detection System amongst trusted peers. Locasto et al. [14] deploy a decentralized system for efficiently distributing alerts to collaborating peers by creating a distributed 'watchlist' from alert streams. Duma et al. [15] explore the challenge of collaborative P2P Intrusion Detection from the perspective of Insider Threat.

Our work deals with a different problem of Game-theoretic patrolling strategies in a P2P network. By running a single IDS device at any given point of time, we aim to conserve the resources of participating peers. And by positioning the IDS at strategic locations, we aim to minimize the gain of an attacker.

III. The 'Game' Environment

As mentioned before, we restrict our discussion to a P2P network architecture which focuses only on super-peer nodes. Leaf-peers are safely ignored in our model owing to the high 'churn-rate' (joining and leaving of peers) seen in them. Any security solution involving leaf-peers is bound to fail since a leaf-peer's lifespan in the network may be very short. Super-peers are selected based on a number of factors which include high uptime, higher network bandwidth, publicly visible IP address etc., and have lesser churn than leaf-peers.

We enlist certain assumptions which are required to clearly define the scope of our Game-theoretic approach:

- We attach a value to every super-peer in the P2P network. This value might be computed based on the node's uptime, the number of super-peers and/or leaf-peers attached to it, reputation etc. For the purpose of this work, the values are chosen arbitrarily.
- Our approach operates on a snapshot of the network topology, and thus requires the network topology to remain constant. Since we base our approach only on super-peers, high churn-rate in leaf-peers has no impact on our approach.
- We assume that an adversary can join the network in the form of a leaf-peer or infect an existing leaf-peer, and will try to disrupt the network by attacking certain super-peers which hold higher value in the network. The attack can be in the form of malicious packets which cause a buffer overflow, injecting malware in shared files, APTs etc.
- We assume that the attacker has full knowledge of the network topology, the location of super-peers, the paths leading to them, etc. This is true in real-life networks as well since an attacker is in a position to gain such information through a 'reconnaissance' involving port-scans etc.

TABLE I. Terms and Symbols used

T	The set of Target nodes
Q	All super-peer nodes in the network
T_j	The payoff associated with the target j.
A	The set of attacker paths chosen $A = A_1, A_2, A_3 \ldots$
D	The set of defender allocations $D = D_1, D_2, D_3 \ldots$
P_D	Defender's mixed strategy over D
P_A	Attacker's mixed strategy over A
$S_{k,j}$	Simple path from source k to target j
V	Value of the game

- We formulate the game with only a single IDS being active in the P2P network at any point of time.
- For the purpose of the Game-theoretic formulation, we limit ourselves to a zero-sum game (primarily because computational limits are reached even for small network graphs). This implies that in case of a successful attack, if the attacker gains a payoff of x, the payoff of the defender is $-x$. The payoff is zero for other cases. This is an important assumption because in zero-sum games, a minimax strategy strategy is equivalent to a Stackelberg strategy [16].
- Since Game theory primarily deals with rational players, we limit our discussion of the attacker or the defender(s) to rational, utility-maximizing players.

A. The Game, Players and Payoffs

Stackelberg Security games are leader-follower games wherein the attacker or the defender takes the first move and the other follows sequentially. In our game formulation, the 'attacker' is a malicious peer who wants to disrupt the P2P network. The 'defenders' are benign super-peers in the network who want to collaboratively ensure smooth and secure operation of the P2P network. We restrict the scope of our research to the attacker-defender model, and do not consider the question of trust or effective collaboration amongst the peers in a P2P network (which are separate areas of study covered in past research [7], [8]).

As mentioned before, we attach values to super-peers which indicate their worth in the network. This value may be chosen based on their contribution to the network, the average number of super-peers/leaf-peers connected to them, their uptime etc. For this work, we choose this value arbitrarily. We consider the scenario where super-peers in the P2P network want to collaborate to protect those nodes which hold higher value in the network (and are thus lucrative targets for an attacker).

Although continuously running an IDS on all peers will provide the maximum security, such continuous monitoring is not viable because peers may not feel incentivized to do so. Although participating peers want successful operation of the P2P network, they are expected to be selfish towards contributing their resources. Moreover, even if multiple IDSs are deployed in the P2P network at certain locations, a deterministic scheduling of IDSs can be observed and thwarted by an active adversary. A randomized patrolling schedule of monitoring the network by running IDSs at strategic locations overcomes this shortcoming and allows peers to save resources at the same time. One of the super-peers generates a Game-theoretic, randomized patrolling schedule based on the 'value'

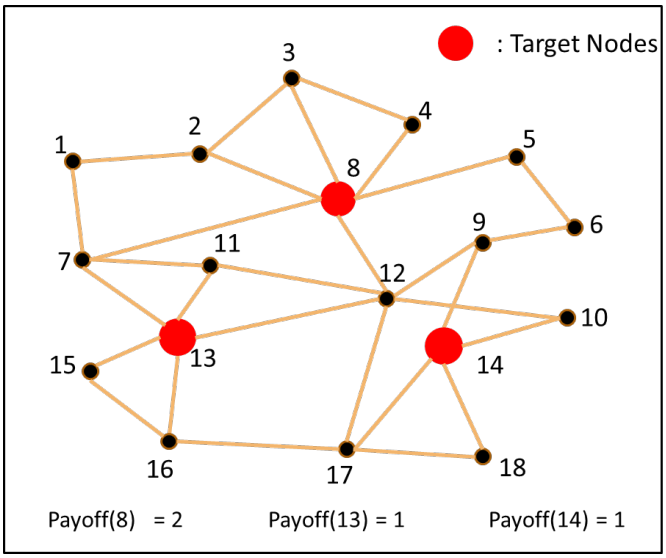

Fig. 1. A snapshot of the P2P network with only super-peers considered

of the 'target' super-peers and the possible paths which the attacker may choose to attack those nodes. The output is in the form of a probability distribution over certain participating super-peers. The probability value assigned to each super-peer indicates the percentage of time (in a given time-slice) for which it should run the IDS.

If the attacker successfully compromises a super-peer whose value is v, the attacker gains a payoff of v and the P2P network loses the value v (or gains $-v$). But the attacker can launch a successful attack only if there is no active IDS on the path between the attacker and its target node. If there is an active IDS, the attack is detected and thwarted. In this scenario, there is no payoff for the attacker or the defender(s). A non-zero-sum game scenario can consider negative payoff for the attacker launching an unsuccessful attack (which takes into account cost of launching the attack, punishment on getting caught etc.), or a cost for the defender(s) to monitor the network. But as we have already mentioned, non-zero-sum games are out-of-scope for this work.

B. Example

Since our approach deals only with super-peers, consider a random P2P network in a simplified form as given in Figure 1. All the nodes shown in the figure are super-peers which may have any number of leaf-peers connected to them. The leaf-peers and their connections with super-peers are neglected, and only super-peers and their connections with other super-peers are shown. For the sake of simplicity, let us consider that only three nodes in this network have a value attached to them: node 8 with value 2, and nodes 13 and 14 with value 1 each. All other nodes are taken to have zero value. Thus, nodes 8, 13 and 14 are the probable 'target nodes' for a malicious peer, and peers in the network would want to minimize the probability of attack on these targets.

An attacker can gain twice as much as payoff by attacking node 8 than by attacking node 13 or 14. He may connect himself to the P2P through any of the super-peers and make that node its 'source', and attempt to attack one of the target nodes through any of the possible paths from the source to the target. For example, the attacker may connect to node 2 and attack node 8 using the path `2->8` and obtain a payoff of 2. If any node on that path (that is, node 2 or 8) has an active IDS, the attacker will not be successful and he gains zero payoff. Then, the attacker may try to attack node 13 using the path `2->1->7->13` and obtain a payoff of 1, and so on. For the example taken by us in Figure 1, the maximum possible payoff for the attacker, or the maximum possible value of the game (V), is 2.

Our solution proposes a randomized strategy for keeping the IDS active on different nodes for different chunks of time (in any given time-slice). These chunks of time are determined using a Game-theoretic approach, as described in the next section.

IV. Proposed Solution

The network snapshot in Figure 1 represents an overlay network of trusted peers modeled by the graph $G(N, E)$. Each node (peer) has an Intrusion Detection System (IDS) installed on it. As mentioned before, we use a Distributed IDS (DIDS) approach. The IDS can monitor the network traffic on all the edges connected to that node. The resources spent by the peers in running an IDS can be conserved if the IDS is ran only on certain strategic locations in the network, and only a single IDS is in 'switched on' condition at any given point of time. Furthermore, this allocation must be done in such a way that the probability of attack on the 'target nodes' is kept to the minimum. The peers (super-peers, to be precise) play the role of defenders who want to defend the 'target nodes' from probable attacks, and collaborate to save their resources at the same time.

Each target node j has a payoff (value) associated with it, given by T_j (a real number). The value of T_j for each target node is determined by the relative importance of the node, which may be modeled as a function of its up-time, associativity etc.

A pure strategy D_i for the defender is to activate the IDS on node i. The set of all such allocations is D. A pure strategy A_k for the attacker is a path from any 'source' node to one of the target nodes in T. The set of all such paths is A. A source node is any super-peer from which an attack can be launched. As mentioned before, the game is a zero-sum game. The attacker gets a utility of T_j for successfully attacking the target j, and zero otherwise. Similarly, the defender gets a utility of $-T_j$ if the attacker successfully attacks target j. Success and failure are defined by the intersection of the Attacker and Defender allocations. If the defender has used node i in its allocation and an attacker path A_k passes through i, the attacker will get detected and his attack will fail. The value of the game is modeled as the utility derived by the attacker for playing the mixed strategy P_A over A.

The objective is to find a Minimax strategy P_D for the defender. Since it is a zero-sum game, the Minimax solution is also a Nash Equilibrium.

Below, we describe the computation of the solution for this attacker-defender game. This computation for the solution is done by the defender(s) by computing the best responses for the attacker as well as the defender in each situation.

A. Trivial Zero-sum Game

Since an attacker is targeting some selected 'target nodes' in the P2P network, a trivial solution is to let the IDS run on the target nodes themselves. In this case the attacker's strategy space can be filled with a simple path from any source to each of the target nodes: $A = A_1, A_2, A_3 \ldots$ where each A_k is an arbitrarily chosen simple path to target node k. The defender's strategy space will include all the target nodes: $D = D_1, D_2, D_3 \ldots$ where each D_k is a target node. It is presented in Algorithm 1:

Algorithm 1 Trivial Solution

```
procedure TRIVIAL SOLUTION
    for each node j ∈ T do:
        k = ChooseRandomSource(Q − T)
        A = A ∪ S_{k,j}
    end for
    D ← T
    (P_A, P_D) = MMSC(A, D)
end procedure
```

MMSC implies MiniMax Mixed Strategy Calculator. It calculates attacker's and defender's mixed strategies P_A and P_D respectively over A and D using the Von Neumann's MiniMax Theorem [17]. This works by formulating linear equations and solving them as follows: Both, the attacker and the defender, aim at making the other player indifferent towards choosing any of the available strategies. This is achieved by equating the average payoff that the opponent receives on choosing any of the available strategy. This is known as the Principle of Indifference [18].

If the game matrix between the attacker and defender is represented as M having $A = A_1, A_2, A_3 \ldots$ and $D = D_1, D_2, D_3 \ldots$, then

$$M = \begin{bmatrix} m_{11} \ldots & m_{1m} \\ m_{21} \ldots & m_{2m} \\ \cdots & \\ \cdots & \\ m_{n1} \ldots & m_{nm} \end{bmatrix}$$

The attacker is the row player and the defender is the column player. m_{ij} is the payoff to the attacker when attacker chooses A_i and defender chooses D_j. P_A is represented as n-tuple $(p_1, p_2, p_3 \ldots)^T$ such that $\sum p_i = 1$. Similarly, P_D is the m-tuple $(q_1, q_2, q_3 \ldots)^T$ such that $\sum q_i = 1$. If attacker chooses the mixed strategy P_A and defender chooses a pure strategy D_j, then the average payoff to attacker is $\sum_{i=1}^{n} p_i m_{ij}$. The defender (who is choosing only the strategy j) receives its exact negative. Similarly, if the defender chooses P_D and attacker chooses a pure strategy A_i then the average payoff to the defender is $-\sum_{j=1}^{m} q_i m_{ij}$. Alternatively, the pure strategy of selecting a single row or column can be represented as a unit vector e_i or e_j where all elements are zero except the i^{th}/j^{th} elements, which are 1.

The probability distribution (mixed strategy) for attacker is found by solving the equations

$$\sum_{i=1}^{n} p_i = 1 \tag{1}$$

TABLE II. TRIVIAL SOLUTION FOR GRAPH IN FIGURE 1

IDS positions:	8	14	13
Probability Distribution:	3/5	1/5	1/5

and

$$\forall j, k \in D, \sum_{i=1}^{n} p_i m_{ij} = \sum_{i=1}^{n} p_i m_{ik} \quad for \quad j \neq k \tag{2}$$

In a similar way, the mixed strategy for the defender is found by solving

$$\sum_{i=1}^{m} q_j = 1 \tag{3}$$

and

$$\forall i, l \in A, \sum_{j=1}^{m} q_j m_{ij} = \sum_{j=1}^{m} q_j m_{lj} \quad for \quad i \neq l \tag{4}$$

The value of the game is the average payoff to attacker when both play the above mixed strategy.

$$V = \sum_{i=1}^{n} \sum_{j=1}^{m} p_i m_{ij} q_j \tag{5}$$

The Solution: With the trivial solution, we obtain a probability distribution with which the IDS should be kept at 'switched on' position at those nodes. The trivial solution thus obtained for the graph in Figure 1 is given in Table II. The trivial solution states that IDS should run on node 8 with probability $\frac{3}{5}$, and on nodes 13 and 14 with probabilities $\frac{1}{5}$ each. The value of the game in this situation is 0.8.

In a real-life scenario, these probabilities may be converted to equivalent time-chunks in a time-slice. Nodes 8, 13 and 14 are the only nodes carrying some value in the network, and node 8 carries the highest value in the network. Thus, such a solution is 'trivial'. A better case will be explored in the non-trivial solution when we do not have the IDS running on the 'target' nodes.

B. Non-trivial Zero-sum Game

The target nodes are super-peers which hold certain prime responsibilities or higher privileges in the P2P network. They already have high load on them. A good example would be of Tor's relay nodes. The trivial solution proposed running IDS on the target nodes themselves. Since the targets themselves become defenders, surely such a solution achieves lowest value of the game. However, at the cost of slightly higher value of the game, we explore the possibility of running the IDS on super-peers apart from the target nodes (and thus reduce the burden on the target nodes). This non-trivial case is explained in Algorithm 2. We follow the same notations as used in the trivial solution. The value of the game is calculated in the same way as mentioned previously. The functioning of **MMSC** also remains same. In the Algorithm, ABR and DBR stand for Attacker's Best Response and Defender's Best Response respectively. We discuss about them next.

Attacker's Best Response (A.B.R): Given the defender's strategy P_D, A.B.R determines the best response by the

Algorithm 2 Non-trivial Solution

procedure NON-TRIVIAL SOLUTION
  $D \leftarrow ChooseRandomDefender(Q - T)$
  $A \leftarrow ChooseRandomPath(G)$
  while V_{old}!=V_{new} **do**
    Calculate V_{old}
    $(P_D, P_A) = MMSC(D, A)$
    $a = A.B.R.(D, A)$
    $d = D.B.R.(D, A)$
    $D = D \cup d$
    $A = A \cup a$
    Calculate V_{new}
  end while
  return (P_D, P_A)
end procedure

attacker. The response in this case is addition of a new path to the attacker's strategy space which maximizes the value of game (or equivalently, the payoff to the attacker). For this, the attacker adds a simple path (p), from an arbitrary source node to a possible target node, to its strategy space A. It now calls MMSC with this modified strategy space (A') and defender's strategy space (D) to arrive at Minimax Mixed Probability Distributions. Using these new probability distributions, the value of the game (V') is calculated in the same way as outlined in the trivial solution case. This is repeated for all possible simple paths from all possible sources to all the target nodes. The path yielding maximum value (V') is selected as the best response. It is explained in Algorithm 3.

Algorithm 3 Attacker's Best Response

procedure A.B.R
  $V_{min} \leftarrow V_{current}$
  $q \leftarrow NULL$
  loop $\forall$ paths p from all source to all targets
    $A' = A \cup p$
    $(P'_D, P'_A) = MMSC(D, A')$
    Calculate V' using (P'_D, P'_A)
    if $V' > V_{min}$ **then**
      $V_{min} = V'$
      $q = p$
    end if
  end loop
  return q
end procedure

Defender's Best Response (D.B.R) Given Attacker's mixed strategy P_A, D.B.R calculates the best response by the defender. The defender aims to minimize the value of the game (thereby minimizing its own loss) and hence the best response will be the addition of such a node to the defender's strategy pool which reduces the current value of the game to the minimum possible. In order to find the best-response node, the defender adds a single node to its already existing strategy space (D) and uses MMSC with this new strategy space (D') and attacker's strategy space (A) to arrive at the Minimax Mixed Strategy Distributions. These new distributions are used to calculate the value of the game (V') at this point, in the same way as mentioned in the trivial solution case. This is done for all the (super) peers in the network and the one yielding minimum value (V') is selected as the best response. The same thing can be explained algorithmically in Algorithm 4. The Algorithm's run-time is $O(n)$ in number of nodes.

TABLE III. NONTRIVIAL SOLUTION FOR GRAPH IN FIGURE 1

IDS positions:	12	2	4	3	7	5	11
Probability Distribution:	1/6	1/6	1/6	1/6	1/6	1/6	0

Algorithm 4 Defender's Best Response

procedure D.B.R
  $V_{min} \leftarrow V_{current}$
  $r \leftarrow NULL$
  loop $\forall$ nodes $n \in \{Q - T\}$
    $D' = D \cup n$
    $(P'_D, P'_A) = MMSC(D', A)$
    Calculate V' using (P'_D, P'_A)
    if $V' < V_{min}$ **then**
      $V_{min} = V'$
      $r = p$
    end if
  end loop
  return q
end procedure

The Solution: With the non-trivial solution, we obtain a probability distribution for running the IDS on super-peers other than the target nodes. For the graph described in Figure 1, the non-trivial solution obtained is given Table III. Note that the probability value for the last node obtained is zero, which indicates the termination of the algorithm.

The value of the game obtained in this scenario was 1.667, which is greater than the value of the game obtained for the trivial case (0.8). As we had mentioned before, the value of the game will be lowest in the trivial case since the targets themselves are the defenders. Running the IDS on nodes other than the target nodes does result in a small de-merit in terms of value of the game (and equivalent negative payoff to the defender), but it saves the target nodes from the extra load of running an IDS.

V. DISCUSSION AND EVALUATION

For the example network considered by us in Figure 1, the maximum value of the game is 2. This value is obtained by the attacker when he successfully attacks node 8. With the trivial solution, the value of the game was reduced to 0.8, which is 40% of the maximum value. However, the trivial solution put the responsibility of running the IDS on the target super-peers. If such a scenario is acceptable, the trivial solution is the best possible solution and achieves the lowest possible value of the game. However, in case the (super) peers in the network do not want to burden the target super-peers with this additional responsibility, the non-trivial solution may be adopted. With the non-trivial solution, the value of the game came out to be 1.667, which is 83.35% of the maximum value.

At present, our approach considers the deployment of a single IDS at any given point of time. Compared to this, game-theoretic deployment of multiple IDSs will allow more paths to be monitored at any given point of time. This will certainly

lead to further reduction in the value of the game, and thus offer better security. However, the initial strategy space with a single IDS (in non-trivial solution) is $|Q - T|$. For a solution involving n IDSs, the Defender's strategy space will grow to ${}^{|Q-T|}C_n$. Computing the Defender's Best Response in this set will prove to be computationally prohibitive for large networks.

The output of our approach is in the form of a probability distribution over super-peer nodes. We interpreted this probability distribution in terms of 'chunks of time' for which these nodes will run the IDS in any given time-slice. In general, this time slice should not be very large. If the time-slice is very large, the IDS will remain fixed for a long time. An attacker may be able to learn the position of the IDS by, say, getting caught once, and then evade the IDS in the next turn by choosing a different path. The strength of our approach lies in its randomized, game-theoretic approach. A very high 'time-slice' value will make it ineffective. We envision a value of about one hour to be suitable for most practical purposes.

It must be noted that our approach operates on a snapshot of the network. An ardent reader may wonder on the practical value of this approach in a P2P environment where nodes join or leave the network at a high rate. But, churn rate is higher in leaf-peers than in super-peers. Rather, a node with short uptime is usually never chosen as a super-peer. Since our approach does not rely on leaf-peers, it is unaffected by new leaf-peers joining or leaving the network. As far as the issue of churn in super-peers is concerned, it will not only impact our approach, but also effect the attacker's knowledge of the network. The attacker will be forced to perform a fresh reconnaissance to gain knowledge of the network. Even the P2P network's routing and indexing does get impacted by ungraceful departure of an important node.

VI. Conclusions and Future Work

This work presented an evaluation of Game-theoretic patrolling strategies for Intrusion Detection in Collaborative P2P networks. Our example demonstrated that by adopting these Game-theoretic strategies, the defenders could bring down the payoff of the attacker (and thus minimize their own loss) to 40–83.35% of the maximum value.

In this work, the value of target nodes was chosen arbitrarily. In future work, we plan to explore more comprehensive ways of modeling the value of a target node. Also, this work considered the case of having a single IDS in 'switched on' mode at any point of time in the network. As mentioned before, we plan to explore the problem of deploying multiple IDSs in the network in a collaborative fashion. Since the problem is computationally prohibitive, heuristic-based solutions can be considered as a viable option.

Acknowledgments

This work was supported by Grant number 12(13)/2012-ESD for scientific research under Cyber Security area from the Department of Information Technology, Govt. of India, New Delhi, India.
We would like to thank Professor Shambhu Upadhyaya for providing his insightful comments on this work.

References

[1] J. Buford, H. Yu, and E. K. Lua, *P2P Networking and Applications*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2008.

[2] "Skype's cloud-based architecture," http://blogs.skype.com/2012/07/26/what-does-skypes-architecture-do/, accessed on 3rd July 2014.

[3] N. Daswani and H. Garcia-Molina, "Query-flood dos attacks in gnutella," in *Proceedings of the 9th ACM Conference on Computer and Communications Security*, ser. CCS '02. New York, NY, USA: ACM, 2002, pp. 181–192.

[4] Y. Zhu and H. Fu, "Traffic analysis attacks on skype voip calls," *Computer Communications*, vol. 34, no. 10, pp. 1202–1212, 2011.

[5] S. J. Murdoch and G. Danezis, "Low-cost traffic analysis of tor," in *Proceedings of the 2005 IEEE Symposium on Security and Privacy*, ser. SP '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 183–195.

[6] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker, "Low-resource routing attacks against tor," in *Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society*, ser. WPES '07. New York, NY, USA: ACM, 2007, pp. 11–20.

[7] E. Anceaume, M. Gradinariu, and A. Ravoaja, "Incentives for p2p fair resource sharing," in *Peer-to-Peer Computing, 2005. P2P 2005. Fifth IEEE International Conference on*. IEEE, 2005, pp. 253–260.

[8] S. Ye, F. Makedon, and J. Ford, "Collaborative automated trust negotiation in peer-to-peer systems," in *Peer-to-Peer Computing, 2004. Proceedings. Proceedings. Fourth International Conference on*. IEEE, 2004, pp. 108–115.

[9] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The eigentrust algorithm for reputation management in p2p networks," in *Proceedings of the 12th international conference on World Wide Web*. ACM, 2003, pp. 640–651.

[10] M. H. Manshaei, Q. Zhu, T. Alpcan, T. Bacşar, and J.-P. Hubaux, "Game theory meets network security and privacy," *ACM Computing Surveys (CSUR)*, vol. 45, no. 3, p. 25, 2013.

[11] M. Kodialam and T. Lakshman, "Detecting network intrusions via sampling: a game theoretic approach," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, vol. 3. IEEE, 2003, pp. 1880–1889.

[12] O. Vaněk, Z. Yin, M. Jain, B. Bošanskỳ, M. Tambe, and M. Pěchouček, "Game-theoretic resource allocation for malicious packet detection in computer networks," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 2012, pp. 905–912.

[13] R. Janakiraman, M. Waldvogel, and Q. Zhang, "Indra: A peer-to-peer approach to network intrusion detection and prevention," in *Enabling Technologies: Infrastructure for Collaborative Enterprises, 2003. WET ICE 2003. Proceedings. Twelfth IEEE International Workshops on*. IEEE, 2003, pp. 226–231.

[14] M. E. Locasto, J. J. Parekh, A. D. Keromytis, and S. J. Stolfo, "Towards collaborative security and p2p intrusion detection," in *Information Assurance Workshop, 2005. IAW'05. Proceedings from the Sixth Annual IEEE SMC*. IEEE, 2005, pp. 333–339.

[15] C. Duma, M. Karresand, N. Shahmehri, and G. Caronni, "A trust-aware, p2p-based overlay for intrusion detection," in *Database and Expert Systems Applications, 2006. DEXA'06. 17th International Workshop on*. IEEE, 2006, pp. 692–697.

[16] Z. Yin, D. Korzhyk, C. Kiekintveld, V. Conitzer, and M. Tambe, "Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness," in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 1139–1146.

[17] R. Motwani and P. Raghavan, *Randomized algorithms*. Chapman & Hall/CRC, 2010.

[18] M. Dresher, *Games of Strategy: Theory and Applications*. Prentice-Hall, 1963.