This paper was presented as part of the Mini-Conference at IEEE INFOCOM 2011

Collaborative Firewalling in Wireless Networks

Mahmoud Taghizadeh
Department of ECE
Michigan State University
taghizad@msu.edu

Amir R. Khakpour
Department of CSE
Michigan State University
khakpour@cse.msu.edu

Alex X. Liu
Department of CSE
Michigan State University
alexliu@cse.msu.edu

Subir Biswas
Department of ECE
Michigan State University
sbiswas@msu.edu

***Abstract*—Firewalls are one of the essential security elements to enforce access policies in computer networks. Open network architecture, shared wireless medium, stringent resource constraints, and highly dynamic network topology impose a new set of challenges on deploying firewalls in a mobile wireless environment. The current state-of-the-art demands for self protection by personal (*i.e.* local) firewalls for each node; however, this requires that all unwanted traffic travels all the way to the node before it is discarded at the destination. This wastes considerable bandwidth and power of all of the nodes in a network with multi-hop routing, specially if a node is under a denial of service (DoS) attack. In this paper, we develop a novel distributed firewalling scheme for wireless networks in which nodes collaboratively perform packet filtering to address resource squandering. The proposed scheme introduces techniques to distribute discarding rules based on both proactive and reactive routing protocols. It also proposes efficient rule placement mechanisms to maximize the number of packets discarded remotely before they reach the destination and minimize the number of unwanted packet forwardings. The scheme is evaluated through various simulation scenarios. The simulation results show that by distributing only 1% of the rules, about 42% of the unwanted traffic is discarded before it reaches the destination, which significantly saves the network resources. Saving about 30% of the wasted bandwidth can be crucial for the performance of a wireless network.**

I. Introduction

A. Problem Statement

Wireless networks, such as Mobile Ad Hoc Networks (MANETs) and wireless mesh networks, have become an integral part of the Internet infrastructure. On the Internet, firewalls are widely deployed on the border of private wired networks to stop unwanted traffic to and from the outside. However, unlike wired networks, it is difficult to deploy firewalls for wireless mobile networks because each wireless mobile node often manages itself and therefore a central firewall policy is often unrealistic to reach. Furthermore, due to the mobility and topology dynamism, wireless mobile networks often lack the concept of private networks and therefore have no clear line of defense. To defend against malicious attacks, each wireless node has to implement the firewall functionality by itself. However, discarding unwanted traffic at destination nodes in wireless networks leads to significant waste of scarce resources, such as bandwidth and power, used by intermediate node to forward unwanted traffic. This paper concerns the problem how to stop, at least reduce, unwanted traffic in wireless networks before reaching destination nodes.

B. Technical Challenges

To discard unwanted traffic before reaching destinations, for each wireless node, we need to distribute its firewall rules to other nodes. However, distributing firewall rules in wireless networks is a technically challenging problem. First, the topology in wireless mobile networks is dynamic and so are the forwarding paths. Thus, for the firewall rules that a node wants other node to enforce, it is difficult to identify which nodes these rules should be sent to. Second, the number of rules that a wireless node can handle is rather limited due to resource limitations. Thus, for the firewall rules that a node receives, it is difficult to decide which rules should be admitted and enforced given its resource constraints.

C. Our Approach

In this paper, we propose a distributed firewalling scheme for wireless mobile networks where nodes collaboratively discard unwanted packets for each other. We address the first challenge of topology and path dynamism by embedding firewall rules within routing messages and distributing a node's rules along the paths that the node receives unwanted traffic so that unwanted packets can be discarded before they reach the node. Coupling rule distribution with routing update messages allows us to find the paths that unwanted traffic are received and therefore distribute rules along them. For proactive routing protocols, where each node periodically sends routing messages to other neighbors, firewall rules are sent out along proactive routing messages so that the nodes receiving the routing messages can enforce these rules. For reactive routing protocols, when a node wants to send a packet, the header of the packet is included in the route request message. If the packet is unwanted, the destination node includes the rule for discarding this packet in the corresponding route reply messages and hereby notifies all intermediate nodes in the path from the packet source to the packet destination. We address the second challenge of resource limitations by each node enforcing portions of its received firewall rules based on rule admission policies and replacing obsolete rules by new rules based on a rule replacement policy. We propose a heuristic based rule admission and replacement algorithm to maximize the number of unwanted packets discarded before reaching their destinations and minimize the amount of unwanted packet forwarding in the network.

978-1-4244-9921-2/11/$26.00 ©2011 IEEE

D. Summary of Simulation Result

Our simulation results reveal that for each node by distributing only 1% of its firewall rules, about 36% of unwanted packets can be discarded in mobile networks and about 42% in static networks. Furthermore, using our rule admission and replacement policy, we can save the network bandwidth wasted by unwanted traffic up to 30% for different mobility patterns and speeds. The results also indicates that our collaborative firewalling scheme is effective in both low speed mobile networks such as MANETs and high speed mobile networks such as vehicular ad hoc networks (VANETs).

E. Key Contributions

We make five key contributions in this paper: (1) We propose an effective and efficient collaborative firewalling scheme for wireless mobile networks. (2) We introduce two performance metrics, Packet Discarding Ratio (PDR) and Forwarding Cost Ratio (FCR), for measuring the performance of our scheme. (3) We propose a heuristic based rule admission and replacement algorithm to maximize PDR and minimize FCR. (4) We implemented and evaluated our scheme on NS2 considering different mobile speeds.

II. Rule Distribution Framework

In this section, we present our rule distribution scheme addressing both rule exporting (*i.e.*, determining which rules to export to neighbors and how to export them) and rule importing (*i.e.*, determining which received rules to be enforced).

A. Overview

A mobile node is assumed to have some firewall policy that specifies what packets it does and does not want to receive from other nodes. The firewall policy is represented by an Access Control List (ACL) consisting of a sequence of rules. Each rule has a predicate over some packet header fields and a decision (*i.e.*, action) to be taken for the packets that match the predicate. The decision of a rule is typically `accept` (*i.e.*, permit) or `discard` (*i.e.*, deny). ACL rules are often overlapping and decisions are made based on the first match semantics (*i.e.*, the decision that an ACL makes for a packet is the decision of the first rule that the packet matches in the ACL). Thus, for rule distribution purposes, we need to convert each overlapping ACL to an equivalent non-overlapping ACL. Each node only exports discard rules (*i.e.*, the rules whose decision is `discard`) and they are stored in a table called Export-Policy Table (EPT). To ensure that the rules exported by a given node can only be matched by packets that destined to this node, the destination field of each rule in the node's EPT must be the address of the node. Accordingly, for each node, we store the rules that it receives and it wants to enforce in a table called Import-Policy Table (IPT). Before forwarding a packet, a node checks the packet header against IPT rules and discard the packet if it matches a rule in the IPT. Figure 1 shows a simple network where node D distributes the rules in its EPT, as shown in Table I, to other nodes. Note that for three rules of D, node A admits rule r_1 and r_2, B admits rule r_1, C admits rule r_2, and S admits nothing. Suppose node S sends a packet with destination port 80 to D, it is forwarded by B but discarded by A before reaching D because this packet matches rule r_2 in A's IPT.

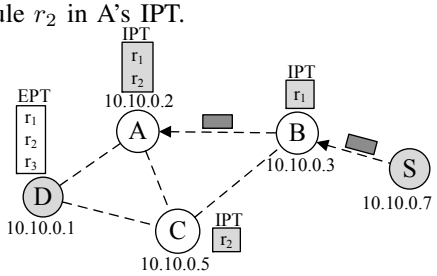

Fig. 1. Rule distribution on an example network

Rule	Src IP	Dest IP	Src Port	Dest Port	Protocol	Action
r_1	10.10.0.5	10.10.0.1	*	8080	TCP	discard
r_2	10.10.0.7	10.10.0.1	*	80	*	discard
r_3	10.10.0.3	10.10.0.1	*	*	*	discard

TABLE I
Export-Policy Table of node D

B. Rule Exporting

1) Constructing Export-Policy Table: To construct EPTs, we need to convert overlapping ACLs to equivalent non-overlapping ACLs. We perform this conversion using Firewall Decision Diagrams (FDD), a tree-based data structure for representing ACLs [1]. An FDD is a directed acyclic graph (DAG) with the following properties: (1) It has exactly one root node. (2) Each non-terminal node represents a packet field and each terminal node represents a decision. (3) A directed path from the root to a terminal node is called a *decision path*. The node labels on every decision path are unique. (4) Each edge is labeled with a non-empty set of integers within the domain of the field that labels the node where the edge is originated. (5) The sets of integers that label the outgoing edges of a node are non-overlapping. The union of these sets equals to the domain of the field that labels the node. After converting an overlapping ACL to an equivalent FDD, we can generate an equivalent non-overlapping ACL from the FDD by generating one rule per decision path. Deleting accept rules from the non-overlapping ACL yields the EPT.

2) Rule Distribution: We distribute rules in EPTs along with routing messages. Below, we discuss our rule distribution scheme based on two types of routing protocols used in wireless networks: proactive protocols and reactive protocols.

a) Rule Distribution with Proactive Routing Protocols: In proactive routing protocols such as OLSR [2] and DSDV [3], each node periodically sends routing updates to its neighbors to keep the nodes' routing table consistent. Such protocols are suitable for stationary networks such as wireless mesh networks and the wireless networks that are communication intensive. Using such protocols, for each discard rule in the EPT of a given node, this node keeps track of the hit rate of the rule, *i.e.*, the number of received packets that match the rule per unit time; once the hit rate of a rule exceeds a threshold, the node sends out the rule piggybacked on its routing messages. Each exported rule r is associated to *hit rate* denoted h. To control the number of hops a rule traverses, each rule has a Time to Live (TTL) value. A rule is not forwarded if its TTL value is equal to 0. Figure 2 shows an example scenario that illustrates the above rule distribution process.

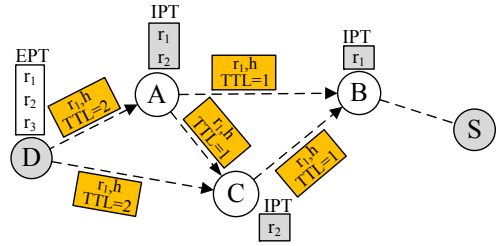

Fig. 2. Rule distribution with proactive routing protocols

b) Rule Distribution with Reactive Routing Protocols: Reactive routing protocols, such as AODV [4] and DSR [5], are suitable for wireless networks that are mobile and that are not communication intensive. In such protocols each node broadcasts a Route Request (RR) message to find a path to the destination node that it wants to send messages; upon receiving an RR message, a Route Reply (RRep) message is sent either by an intermediate node (if the node knows a path from itself to the destination) or the destination node. The RR messages include source and destination address to discover the route to the destination. In our rule distribution scheme, we modify RR messages so that they also include other packet header fields such as destination port, source port, and protocol type. Before an intermediate node replies to an RR message with an RRep message, it checks the packet header fields included in RR message against the rules in its IPT. The destination node, however, checks the packet header against its EPT. If the packet header does match any of the rules, the node includes the corresponding rule r with its hit rate h in the RRep message and sends it over to the source. Note that hit rate of rule r is updated each time a packet header matches rule r. When an RRep messages travels back to the source node, it informs all intermediate nodes in the path about rule r. Upon receiving the RRep message, the source node or the intermediate nodes can decide whether to import the rule based on its hit rate. Figure 3 shows an example scenario that illustrates the above rule distribution process. In this example, node S wants to send an HTTP request to node D over destination port 80. To find the path to node D, node S sends an RR message including source port (SP=1111), destination port (DP=80), and protocol (P=TCP). As rule r_2 in the EPT matches the packet header fields in the RR message, node D includes rule r_2 along with its h to the RRep message and sends it over to node S. Finally, node S realizes from the RRep message that its packet will be discarded by node D. Hence, node S stops sending any HTTP request.

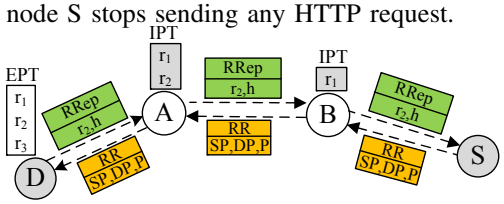

Fig. 3. Rule distribution with reactive routing protocols

As reactive routing protocols usually use route caching, a node may receive unwanted packets without receiving the corresponding RR message. Thus, only relying on reactive approach to distribute rules may not be sufficient. Therefore, we can use the proactive approach along with the reactive approach.

C. Rule Importing

When a node receives a rule, it decides whether to admit the rule to its IPT based on its admission and rule replacement policies. Rule admission policy is based on multiple factors such as the rule's originator and its hit rate. For instance, a node may ignore a received rule if its hit rate is less than a prescribed threshold. If a node decides to store a new rule while its IPT is full, it must evict another rule from its IPT to accommodate the new rule based on certain replacement policy. In section III, we present an admission policy to optimize rule placement based on our system performance metrics.

D. Policy Table Consistency

When the ACL rules of a node are modified, probably by its administrator, the EPT of the node has to be recomputed. The new EPT rules may be different from the old EPT rules, which may have been exported and admitted by other nodes. To avoid such inconsistency, we propose two mechanisms that all nodes can employ to remove the admitted rules originated from a particular node: *rule revocation* and *rule expiration*. In the rule revocation mechanism, a node can broadcast a revocation message that contains all the rules that this node wants to revoke. In the rule expiration mechanism, each rule in an IPT has a lifetime and is deleted when it expires. These two mechanisms can be used together.

III. Rule Admission and Replacement Policy

We first define two performance metrics:(1) Packet Discard Ratio (PDR) that represents the percentage of unwanted packets discarded in the path before reaching the destination, and (2) Forwarding Cost Ratio (FCR) that represents the expected value of the portion of the path that an unwanted packet is forwarded before being discarded. We then present *Split Replacement Policy (SRP)* to maximize PDR and minimize FCR.

As mentioned before, due to limited size of IPTs, PDR and FCR are dependent on the number of rule duplications in the paths to the destination. Thus, we should carefully control the amount of rule duplication to tradeoff between the two potentially conflicting goals of minimizing FCR and maximizing PDR. To this end, we divide each IPT into two segments: the first segment of an IPT holds rules with highest hit rates with no constraints on number of duplications in the path. The second segment of an IPT stores the rules that are unique in the path. Keeping the rules with the highest hit rates in the first segment of an IPT helps decreasing FCR, and keeping the uniqueness of rules along the path to the destination in the second segment of IPT helps increasing PDR. By adjusting the size of the first and the second segments of IPTs, SRP can easily regulate the overall cost of unwanted packets in the network.

Using SRP, when a node receives a new rule, it imports the rule, if it has a free slot in its IPT; otherwise, it executes the following procedure: (1) If the rule has been stored in the path

and if there is a rule in the *first* segment of an IPT whose hit rate is smaller than the new rule's hit rate, the rule is replaced with the rule with the lowest hit rate; otherwise, the rule will not be imported in IPT. (2) If the rule has not been stored in the path and if there is a rule in the *second* segment of an IPT whose hit rate is smaller than the new rule's hit rate, the rule is replaced with the rule with the lowest hit rate; otherwise, the rule will not be imported in the IPT. In case a rule is added to an IPT, the rule admission distance is updated to node's distance to the destination and the rule will be forwarded to the next hop. SRP pseudocode is shown in Algorithm 1.

Algorithm 1: Split Replacement Policy

Input: rule r and rule hit rate h
node IPT :IPT
IPT size: c
Output: new IPT: IPT

if *($IPT.size \leqslant c$)* **then**
  $IPT.add(r)$
  return
else
  if *(the rule r has not been already stored in the path)* **then**
    $\hat{r} \leftarrow$ rule with lowest hit rate $\hat{h}$ in the second part of IPT
  else
    $\hat{r} \leftarrow$ rule with lowest hit rate $\hat{h}$ in the first part of IPT
  if *($h > \hat{h}$)* **then**
    replace $\hat{r}$ with r
forward the rule to the next node in the path

IV. Evaluation

We implemented our scheme on NS2 [6]. To evaluate the performance of our collaborative firewalling scheme, we measure PDR and FCR with different system parameters. We then used three different mobility profiles: (1) a static network with no mobility, (2) a low speed mobile network with average speed of 1m/s, and (3) a high speed mobile network with average speed of 10m/s. The mobility pattern is based on the popular random waypoint model for mobile network simulations [7].

A. Simulation Setup

To simulate all unwanted traffic in a network, we create N rules and N sets of unwanted packets where each rule exactly matches one unique set of unwanted packets. In this simulation, each set of unwanted packets is generated at a rate following a Zipf distribution (*i.e.*, power-law distribution).

When simulation starts, nodes begin generating unwanted packets. In our simulation, we only use the reactive approach for rule distribution because PDR and FCR are not dependent on the rule distribution framework; thus, presenting the results for reactive approach suffices. Thus, once the destination node receives an RR message for an unwanted packet, it sends the packet's corresponding rule to the source node. When a node in the path receives the rule, it may import the rule based on its admission and replacement policy. Table II summarizes the simulation parameters.

α	Zipf parameter	$[0.5 \cdots 1]$
C	IPT size	$[10 \cdots 200]$
θ	SRP parameter	$[0 \cdots 1]$
v	Movement speed	0 m/s (static), 1m/s (slow) and 10m/s (fast)
N	Total number of rules	10000
n	Total number of nodes	100
A	Area	500×2500 meter
Simulation runs for 15000 seconds. Packets are generated every 2 seconds		

TABLE II
Simulation parameters

B. Simulation Results

We evaluate SRP performance by measuring PDR and FCR using different values of θ. Recall that we divide IPT tables into two segments where the first segment stores high hit rate rules and the second segment stores unique rules in the path that rules are distributed. In practice, the high hit rate rules are duplicated in the first segment of the IPT for all nodes in the path. We used parameter θ to indicate the percentage of IPT reserved in the first segment. Figure 4(a) shows the impacts of θ on PDR. The results correspond to $\alpha = 0.8$ and IPT size of 100, which is equal to 1% of the total number of rules in the destination node. For small θ values, nodes only store rules which have not been stored by other intermediates nodes in the path. Hence, the total number of different rules stored in the path increases which in turn causes more number of unwanted packets to be discarded in the path. This explains why PDR is high for small values of θ.

However, as θ increases, the amount of rule duplication along the path that the rule is distributed increases and therefore PDR reduces. In an extreme case, when $\theta = 1$ (*i.e.* no second segment in IPT), all nodes will store the same set of high hit rate rules, which in turn leads to the minimum PDR. The slope of decreasing trend of PDR is different for three mobility patterns. The more mobile a network is, the less significant the impact of θ is on PDR. The reason is when nodes are moving, rules in the second segment of IPT may not be unique along the path. Thus, the number of duplicated rules in a path stored in the second segment of IPT increases as nodes move faster. This *unnecessary* duplication of the rules in the second segment reduces the impact of θ on PDR.

Figure 4(b) shows the impact of θ on FCR. By increasing θ, greater number of high hit rate rules are stored in IPT. This leads to discarding more unwanted packets locally with minimum FCR of 0. However, as shown in Figure 4(a), by increasing θ, the number of packets that reaches the destination increases (*i.e.* PDR decreases). Thus, there are more number of unwanted packets with the maximum FCR of 1. The tradeoff between percentage of unwanted packets being discarded at the source and the percentage of unwanted packets that reach the destination brings up an optimal point for θ at which the average FCR is minimum. For instance, for a static network when $\theta = 0.62$ FCR is minimum. The optimal point for mobile networks shifts to the left as their mobility speed increases due to unnecessary rule duplications for the rules that are stored in the second segment of the IPT.

Figure 4(c) shows the impact of θ on the overall cost when $\gamma = 0.5$ (*i.e.* performance metrics PDR and FCR are equally important). In this case, the overall cost is minimum when

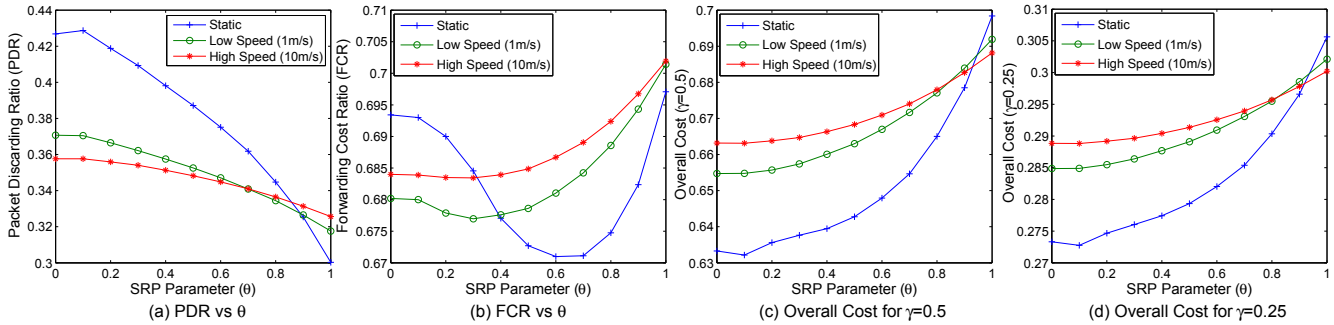

Fig. 4. SRP Packet Discarding Ratio and Forwarding Cost Ratio vs θ

$\theta = 0$ for mobile networks and $\theta = 0.1$ for static network. The results show in general the collaborative firewalling is indeed required to minimize the overall cost of unwanted packets. As in wireless mobile networks, FCR seems to be more important than PDR due to the power constraints on mobile nodes, Figure 4(d) shows the overall cost when $\gamma = 0.25$ where FCR is three times more important than PDR. The results show that the overall cost for SRP is down to 0.272 for static and 0.285 for mobile networks, respectively.

V. Related Work

Maccari *et al.*. proposed a firewalling scheme for mesh networks in [8]. In this scheme, all accepted packets for each node are represented by a bloom filter. In this scheme, they use a bloom filter to send the list of accepted packets to all nodes in the network. Thus, when a node wants to forward a packet, it queries the packet from all bloom filters it has received from other nodes. If it is found, the packet is forwarded; otherwise, it is discarded. Due to large number of accepted packets, the bloom filter is very big (it could be in order of GB or TB). To deal with this problem, they only consider packets with class C IP addresses and port numbers less than 1024, which is a considerable limitation of their work. The authors extend their work in [9] to support stateful firewalls and they use d-left counting bloom filter with handover support. Our scheme is different from this work in the following perspectives: (1) In Maccari *et al.*'s work, all firewall rules are distributed in form of sets of all accepted packets, whereas we incrementally send non-overlapping discard rules on-demand when they have hits. (2) Maccari *et al.*'s work is not scalable as the size of bloom filters is tightly dependent on the size of accepted packets, whereas using rules instead of bloom filters resolves the scalability issues. (3) In our scheme, we embed small rules in routing messages to minimize the communication overhead between nodes, whereas in Maccari *et al.*'s work, node's large bloom filters should be sent over to all other nodes in the network, which results in very high communication overhead.

Alicherry *et al.* presented a traffic authentication framework for MANETs in [10], [11]. In this framework, a set of trusted nodes is appointed as group controller that are responsible for distributing token policies. When a node wants to access an authorized service on a destination, the source node sends the corresponding token policy to the destination. The destination notifies the source node as well as all intermediate nodes along the path to the source about the amount of allocated bandwidth for the requested session. This work is fundamentally different from ours, as in our scheme a destination node exports its rules along the paths to reduce the number of unwanted packets it receives. Moreover, the scalability of this approach is questionable, as in a communication intensive network, each node needs to handle a large number of token policies that are not in compliance with the nodes that have limited storage and computation power capacity.

VI. Conclusion

In this paper, we propose a collaborative firewalling scheme for mobile networks. We introduce two performance metrics, namely Packet Discarding Ratio (PDR)and Forwarding Cost Ratio(FCR), to address the effectiveness of distributed firewalling in the network. We further propose a heuristic algorithm, namely the Split Replacement Policy to maximize PDR and minimize FCR. We finally evaluate the performance of the system by extensive simulation on mobile networks with different mobility profiles. Our results show that using the proposed collaborative firewalling scheme, a considerable portion of unwanted traffic can be discarded before reaching the destinations, which saves substantial amount of power and bandwidth.

References

[1] Mohamed G. Gouda and Alex X. Liu, "Structured firewall design", *Computer Networks Journal (Elsevier)*, vol. 51, no. 4, pp. 1106–1120, March 2007.

[2] T. Clausen and P. Jacquet, "Optimized link state routing protocol (OLSR)", *RFC 3626*, 2003.

[3] P. Bhagwat C. E. perkins, "Highly dynamic destination-sequenced distance vector (DSDV) for mobile computers, protocols and applications", *Proc. SIGCOMM*, 1994.

[4] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On-Demand Distance Vector (AODV) routing", *RFC3561*, 2003.

[5] D. Maltz D. Johnson, Y. Hu, "The Dynamic Source Routing Protocol (DSR) for mobile ad hoc networks for IPv4", *RFC4728*, 2007.

[6] "The network simulator-NS2", http://www.isi.edu/nsnam/ns/.

[7] Christian Bettstetter, Giovanni Resta, and Paolo Santi, "The node distribution of the random waypoint mobility model for wireless ad hoc networks", *IEEE Transactions on Mobile Computing*, 2003.

[8] Leonardo Maccari, Romano Fantacci, P. Neira, and R.M. Gasca, "Mesh network firewalling with bloom filters", in *Proc. ICC*, 2007.

[9] P. Neira, R.M. Gasca, Leonardo Maccari, and L. Lefevre, "Stateful firewalling for wireless mesh networks", in *IFIP International Conference on New Technologies Mobility and Security (NTMS)*, 2008.

[10] Mansoor Alicherry, Angelos D. Keromytis, and Angelos Stavrou., "Deny-by-default distributed security policy enforcement in mobile ad hoc networks", *Proc. Int. ICST Conference on Security and Privacy in Communication Networks*, 2009.

[11] Mansoor Alicherry and Angelos D. Keromytis, "DIPLOMA: Distributed Policy Enforcement Architecture for MANETs", in *Proc. Int. Conf. on Network and System Security (NSS)*, 2010.