2010 Sixth International Conference on Mobile Ad-hoc and Sensor Networks

High Performance Firewalls in MANETs

Hang Zhao
Department of Computer Science
Columbia University
zhao@cs.columbia.edu

Steven M. Bellovin
Department of Computer Science
Columbia University
smb@cs.columbia.edu

***Abstract*—Doing route selection based in part on source addresses is a form of policy routing, which has started to receive increased amounts of attention. In this paper, we extend our previous work on ROLF (ROuting as the Firewall Layer) to achieve source prefix filtering. This permits easy definition of "inside" and "outside", even in MANET environment where there is no topological boundary. We present algorithms for route propagation and packet forwarding using ROFL; we measure its performance in a simulated environment with two different ad hoc routing protocols. Simulation results demonstrate that ROFL can significantly reduce unwanted packets without extra control traffic incurred, and thus improves overall system performance and preserves battery power of mobile nodes. ROFL is the first scheme to provide a concrete defense against some battery exhaustion attacks in MANETs. Moreover, it requires only minor changes to existing ad hoc network routing protocols, making it practical and feasible to be deployed in real world.**

Keywords: Firewall, Policy Routing, Security, MANETs

I. Introduction

Firewalls have long been a mainstay of network security. While their utility has diminished in recent years [4] due to increasingly rich topologies, they are still valuable. In recent work [19], [21], we extended firewalls to work with mobile ad hoc networks (MANETs), using routing protocols to implement the firewall layer. MANETs, which are self-configuring networks consisting of mobile devices connected by wireless links, are not amenable to protection by traditional firewalls. Nodes in MANETs are free to move, causing network topology to change rapidly and unpredictably. The lack of network infrastructure makes each node serve as a router that relays packets towards other nodes. There are no topological choke points to block unwanted traffic. Our solution, ROFL (Routing as the Firewall Layer), treats destination port number as part of the address during routing. Consequently, packets intended for forbidden services are dropped closer to source. This serves two important purposes: it not only helps implement a security policy, it causes unwanted packets to be dropped as early as possible, thus conserving battery power.

In common with most firewalls and routing protocols, ROFL makes its security decisions based on destination addresses and port numbers. For wired networks, this is quite proper, since it is not possible to trust source addresses come from beyond the firewall [9]. In MANETs, where connectivity patterns are constantly changing, the situation is subtly different. While it remains true that behavior of "untrusted" nodes (i.e., those not protected by firewall) cannot be relied upon for security purposes, adding source address constraints to firewall rules have a second, and equally important function: they define the boundaries of *policy region*. That is, source address rules define the boundary between the "inside" — the portion of the network protected by firewall — and the outside. In a traditional wired network, firewall itself is the boundary marker, as network topology determines the inside and outside.

In the MANET shown in Figure 1(a)(i), nodes I1 and I2 provide firewall functionality against outside node O1. In Figure 1(a)(ii), nodes I1 and O1 have moved, changing the connectivity patterns. I2 can no longer reach I1 except by going through O1; in addition, I4 has a direct link to O1 and must activate firewall mechanisms against traffic originating from it. We must now rely on O1's source address to make such decisions, rather than on a fixed topology. In ROFL, we use the routing system to implement firewall functionality. If we now use source addresses as part of our firewall mechanism, we must also use them in routing. This, then, *requires* that we implement a policy routing scheme.

From a quantitative perspective, ROFL has some interesting characteristics. Clearly, it will save power by dropping unwanted packets early; how great the savings are will depend on the rate of unwanted traffic. A tactical MANET is never a closed network; there is always the chance that some network node will be compromised and used to attack other nodes. Accordingly, very fine-grained permission mechanisms are highly desirable, to minimize the damage from any such incidents. In other words, in any given MANET, not all nodes will be able to access all services on all other nodes.

As the network's topology changes, this constraint can be problematic. In Figure 1(a)(i), I1, I2, I3, I4, and I5 are allowed to speak to each other for some service, while O1 is not. As the nodes move (Figure 1(a)(ii)), I3 should no longer even attempt to transmit to I1, since its only path goes through O1, a node untrusted. If I3 transmits nevertheless, that traffic is must be dropped at some point. Without ROFL, the drop point would be I2 or I4, since they are aware of the path. With ROFL, I3 itself would know not to transmit, thereby saving power. The tradeoffs are complex. We have saved power and preserved security by not sending via an untrusted node. On the other hand, a desired communication has been blocked, not because O1 has become evil but because we fear that it may happen.

To the best of our knowledge, ROFL is the first scheme to provide a concrete defense against some battery exhaustion attacks in MANETs. Other contributions include: 1) we extend our previous work on ROFL to include source prefix filtering

978-0-7695-4315-4/10 $26.00 © 2010 IEEE
DOI 10.1109/MSN.2010.30

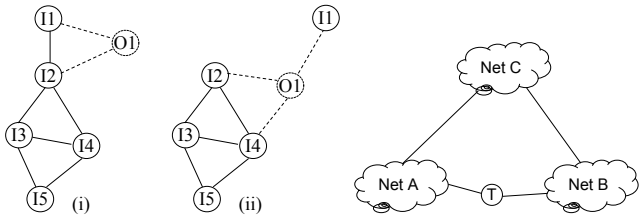

(a) Two different topologies in a MANET, after O_1 has moved.

(b) Two subnets connected by both a transit node and a transit net.

Fig. 1: MANET topologies.

to implement a policy routing scheme; 2) we provide algorithms for route propagation and packet forwarding to cope with existing routing protocols; 3) we demonstrate the high performance of ROFL as a firewall mechanism in MANETs using simulation. The results show that, for AODV routing protocol when nodes moving at a moderate speed of 10 m/s and a 10% bad traffic rate, ROFL reduced 8.4% of data packets and 11.5% of control packets transmitted; with 20% bad traffic, the reductions were 18.6% and 21.4%. The same trend was observed when the amount of malicious traffic growths. Similar results were also observed with OLSR routing protocol.

In Section II, we summarize the basic ROFL scheme proposed previously. Section III introduces the new policy routing mechanism. Section IV provides algorithms for route propagation and packet forwarding for the new scheme. Section V describes the protocol design for implementing ROFL with existing MANET routing protocols. In Section VI, we study ROFL's performance through network simulations. We finish with a discussion of related work and conclusions drawn.

II. Background

A. Basic ROFL Scheme

The ROFL scheme is based on a simple notion: services — that are, port numbers — should be treated as part of the IP address in the routing system. If a certain service is not advertised to a particular network, no host on that network can reach it; the routing system will not deliver the packets. We thus use every router along the path as a firewall. There are many benefits to this scheme, especially in MANETs where battery power is limited. If unwanted packets are dropped very early, a lot of power can be saved by not transmitting those packets. In ROFL, we extend the prefix field of a conventional routing advertisement to include a service s, and a ROFL announcement is

$$R = \{d : s/m,\ M\}$$

where d is an address prefix, s is a service port number, m is a prefix length, and M is a routing metric. To block a service, we advertise an infinity route with $M = \infty$. Virtually all useful protocols require replies from servers to the clients. This implies that there must be a route to client ports. Such kind of routes to client ports be handled similarly to service-specific routes. (Full details are given in [21].)

ROFL takes advantage of underlying routing protocol to disseminate service-specific route information. There is no need to implement a separate route propagation phase. Moreover, ROFL advertisements are handled just like any other routing advertisements: the source node and any intermediate routers do a longest-prefix match on the advertisement. If there is no matching route, the packet is dropped.

B. Extensions of ROFL

There are a number of extensions that may also prove to be useful. One is to route on IPsec SPIs instead of the port numbers, since port numbers are not visible to intermediate nodes for encrypted traffic. Routing on SPIs permits early drop of many fake packets. It is especially useful if fine-grained SPIs are used, as in distributed firewalls. Note that ROFL and cryptography are orthogonal issues. Cryptographic techniques provide confidentiality and integrity of communications; the goal of ROFL is not to replace cryptography, but rather to drop unwanted packets early to save transmission power.

A second extension is to separate port number from IP address, and instead pass a list of permitted or prohibited port numbers along with a prefix announcement. This would allow a single announcement (and hence FIB entry) to handle many services. More importantly, it would easily handle ROFL announcement for an entire network, such as $\{192.0.2.0/24, \{25, 80\}, M\}$. However, this would require an additional lookup step not part of destination prefix matching; additionally, it might conflict with the desire to route different services differently.

III. Policy Routing with ROFL

A. ROFL with Source Prefix Filtering

To implement policy routing, we augment each prefix announcement with a set of permissible source prefixes S:

$$R = \{d : s/m,\ S,\ M\}$$

such that data traffic coming from those address prefixes is allowed to access service s provided by destination address d. We define $S = \{p_1, p_2, ..., p_n\}$; each p_i $(0 \leq i \leq n)$ is a source address prefix. Thus no source prefix constraint is specified if $S = \phi$. (Alternately, $S = \{0/0\}$, i.e., all addresses are accepted; for clarity in this paper, we prefer to distinguish between no source prefix constraints and one that happens to have no effect.) In some situations, it is possible to implement S as a Bloom filter [5] on the set of source addresses or networks of a given prefix length. Bloom filters are a space-efficient data structure that can compress the representation of a set of members in a compact manner, albeit with some chance of false positives. Bloom filters are particularly useful and feasible in MANETs, where there is little topological structure and each allowed node is identified by a flat address.

Source prefix filtering (SPF) constraints in ROFL control route propagation and packet forwarding. A ROFL announcement is passed to a node only if that node is authorized to access the advertised service. During packet forwarding phase, a packet is dropped immediately if it is coming from a source address not specified in the SPF constraints of a matching route. Moreover, this new ROFL scheme is capable of implementing a complete set of filtering functionalities provided by traditional firewalls. (See our technical report [20] for a detailed discussion.) For simplicity, ROFL refers to the new scheme with SPF for the rest of the discussion.

Route selection is now based in part on the source address. Therefore the actual route announcement becomes a form of policy routing. There are two obvious approaches to inserting policy constraints into routing announcements. First, all relevant nodes along the route propagation paths could create or modify the policy statement, in accordance with some central policy. A better approach would be allowing only the node advertising the service — the route originator — to embed a source prefix constraint in routing announcement. Subsequent receivers of this route announcement do not alter the embedded policy statement. We suggest the second approach for a few reasons. First of all, route initiator has the best knowledge of authorized nodes for the service it provides. Secondly, allowing intermediate routers to modify policy statements requires trust relationship established amongst them. More importantly, the first approach might work for static nodes; but would not be able to cope with dynamically changing topologies, like in MANETs. Finally, if were to add some sort of routing security mechanism similar to S-BGP [12], [13], changes en route would be much more difficult.

B. The Transit Node Problem

Consider the network topology shown in Figure 1(b). A node in Net A is advertising a service that a node in Net B wishes to access; however, the transit node T is not an authorized source for this service. That is, the source address policy advertised for this service includes Nets B and C, but not T. What should happen? One option would be to permit traffic from B to T; after all, Net B is an authorized source. This in turn would require that the border node in Net A advertise the service to T, which would presumably pass it on to Net B. This option is insecure: when packets for the service arrive at Net A, it is not possible to tell whether they originated from Net B or were forged by transit node T. We thus adopt the following policies: 1) Routing advertisements are never propagated to a node not authorized for that service, according to the originator's policy; 2) Packets for a service are only accepted from nodes to which routing advertisements were sent. In a wired net, this is generally easy; in a MANET, it will likely require some form of cryptographic authentication of the neighboring node, since in a wireless environment it is generally very difficult to determine the precise source of a packet. These principles are not related to route aggregation. A shorter prefix may be transmitted to T if and only if such a prefix already existed and T was authorized for it.

IV. Algorithms

Routers are generally composed of two fundamental mechanisms, the control plane and the data plane. The control plane, sometimes known as route generation, produces a valid path from source to destination by exchanging routing information with other nodes. The data plane, or message forwarding, relays packets from node to node until they reach their final destination, following the selected route. In ROFL, we implement packet filtering by layering it on top of routing. ROFL is agnostic to the specific type of routing protocols used; only slight modifications are needed during the route propagation and the packet forwarding phases.

Because we do not change route calculations or prefix lookups, our new ROFL scheme can cope well with any distance vector or link state algorithms: route announcements in ROFL with source prefix filtering are handled the same way as conventional ones during this phase. There are some changes, however. Therefore, in this section we present our algorithms for route propagation and packet forwarding.

A. Route Propagation

Depending on the dissemination pattern of routing information provided by underlying routing protocol, we propose several variants for route propagation.

1) Route Propagation via Unicast: If point-to-point dissemination of routing announcements is possible, like in a wired network, a service provider has full control over the propagation of service specific routes.

```
    RoutePropucast(u)
1   if u provides s
2     R ← genRt(d, s, m, S, M);
3   else
4     S ← recvRt(R);
5     calculateFIB(R);
6   for each neighbor h ∈ Nb
7     if S == φ or h ∈ S
8       sendRt(R, h);
```

Service provider (i.e., the route initiator) generates a routing announcement R with appropriate routing information enclosed (Line 2), and then propagates R only to neighbors that are allowed to access that service (Line 6-8). Upon receipt of R, a node extracts source prefix constraint information S from R (Line 4), calculates its forwarding information base (FIB) to reflect the update from newly received route (Line 5), and continues forwarding R to its authorized neighbors.

2) Route Propagation via Broadcast: If routing information is disseminated into a network through broadcast by the underlying routing protocol, such as the flooding phase in many table-driven MANET routing protocols, it is the receiver's responsibility to make a processing and forwarding decision by consulting enclosed source prefix filtering information.

Service provider u initiates a ROFL announcement R disseminated into the network via broadcast (lines 1-3). When a node u receives this announcement, it first retrieves the SPF constraints S enclosed (line 5), and then makes a processing and forwarding decision by consulting S (lines 6-10). If u is an authorized receiver, it shall update its routing table and forwarding information base (FIB) accordingly and then continue broadcasting R (lines 7-8); otherwise, R is silently discarded (line 10).

```
    RoutePropbcast(u)
1   if u provides s
2     R ← genRt(d, s, m, S, M);
3     bcastRT(R);
4   else
5     S ← recvRt(R);
6     if S == φ or u ∈ S
7       calculateFIB(R);
8       bcastRT(R);
9     else
10      discardRT(R);
```

3) Route Propagation via Multicast: Another alternative for route propagation can be performed via multicast. Suppose there is a separate multicast address for each policy, then ROFL announcements for that policy are sent to that multicast address; only nodes listening to that address will

normally receive it. This approach requires a pre-established multicast group for each policy implemented. Moreover, a secure multicast can be performed by assigning group key only to authorized receivers belonging to that group, using available key management mechanism for MANETs.

B. Packet Forwarding

Once a path is established between a source and destination pair by underlying routing protocol, every node along the path consults its local routing table T to make forwarding decision.

```
    PKTFORWARD(u)
1   p_s, p_d, s_d ← procPkt(K);
2   if u is the destination
3     S ← getSPF(s_d);
4     if S == φ or p_s ∈ S
5       sendToUpperLayer(K);
6     else
7       discardPkt(K);
8   else
9     if R ← findRt(T, p_d, s_d)
10      S ← extractRt(R);
11      if S == φ or p_s ∈ S
12        forwardPkt(K, R);
13      else
14        discardPkt(K);
15    else
16      discardPkt(K);
```

Upon receiving a data packet K, node u obtains source address p_s, destination address p_d and destination port number s_d from packet header (Line 1). If K arrives at its destination and indeed comes from an authorized source, the packet is passed to upper layer (Line 2-7); Otherwise K needs to be further routed. Node u consults its routing table T based on p_d and s_d (Line 9). If a matching route R is found, u extracts the source prefix constraint from R (Line 10). Then K is forwarded towards its destination if it's from an authorized source (Line 11-12). Otherwise, this packet is silently discarded. Line 3-7 implements a distributed firewall at destination host by checking packet source for service s_d. This would drop malicious traffic coming from direct neighbors, since there is no intermediate routers in between that can act as firewalls.

C. Discussion

Our *route propagation* and *packet forwarding* algorithms described above are similar to those dealing with conventional routing advertisements. With ROFL, though, there is an additional check against the source prefix constraint which constraints route propagation and packet forwarding decision at each node. For a detailed proof of correctness of those algorithms, please refer to our technical report [20].

V. PROTOCOL DESIGN

ROFL can be easily adopted by a large number of routing protocols with minor modifications, while being transparent to upper layers. Characteristics of underlying routing protocol determine applicable variant of route propagation algorithm discussed previously. Once the routing table and the forwarding information base (FIB) are successfully updated upon receipt of ROFL announcements, packet forwarding is straight forward following algorithm PKTFORWARD(). Thus, we focus on propagation of ROFL announcements in MANETs.

Classification of MANET routing protocols could be done in many ways [16]. One common taxonomy is to categorize them in terms of routing strategy: reactive routing (on demand), proactive routing (table-driven), and hybrid routing which combines previous two. Another popular classification is based on how routing paths are computing at each node using different routing information. For a comprehensive study, we carefully chose two routing protocols: *Ad hoc On-Demand Distance Vector* (AODV) protocol [15], a reactive routing protocol using distance vector information, and *Optimized Link Stat Routing* (OLSR) protocol [10], a proactive routing protocol using link state information.

A. ROFL with AODV

AODV [15] is one of the most popular on demand routing protocols for MANETs. It doesn't require routing paths to be established before the presence of active communication. If a route to the destination is not available, a node broadcasts a route request (RREQ) message using an expanding ring search technique. A route reply (RREP) message is then generated by either the destination node or an intermediate node that has a valid route to that destination.

To implement ROFL scheme with AODV, we replace conventional routing information embedded in a RREP message with service specific announcement that include a service port number and source prefix constraints. As RREP traverses towards the route requester, ROFL announcement is stored at each intermediate node following the ROUTEPROP$_{ucast}$() algorithm, because RREP is unicast back to the route initiator along the reverse path that RREQ traversed. We also piggyback client route information in RREQ messages initiated by the route requestor at the beginning. Therefore, our implementation doesn't require extra control messages to be transmitted comparing to its non-ROFL counterpart (i.e., pure AODV).

B. Optimized Link State Routing (OLSR)

OLSR [10] is a commonly deployed table-driven routing protocol for MANETs, that makes use of link state information to compute shortest paths to all destinations in the network. To reduce flooding overhead, only nodes, selected as Multipoint Relays (MPR), are responsible for forwarding control traffic disseminated into the entire network. Hello (HELLO) message contains the link information with the 2-hop neighborhood of a node, and Topology Control (TC) message declares link availability between a node to its MPR selectors. Without loss of generality, we assume that each node has only one interface participating in OLSR routing. Multiple Interface Declaration (MID) messages are irrelevant to route discovery.

With ROFL, service specific routes with source prefix constraints are embedded in each HELLO message and TC message without extra control messages. Since those messages are broadcast, route propagation algorithm ROUTEPROP$_{bcast}$() is implemented for OLSR. Therefore, we implement ROFL scheme with AODV and OLSR respectively. Although the implementation is protocol-specific, ROFL announcements are treated just like any conventional routing advertisements in those protocols. Most of the implementation effort is dealing with routing table maintenance, and its cost has been analyzed quantitatively in our previous work [21].

VI. Performance Evaluation

We evaluate ROFL's performance with two routing protocols, namely AODV and OLSR, in a simulation environment, as there are not many MANETs actually deployed. We vary the amount of malicious traffic present in the network to analyze data traffic and control traffic incurred respectively in ROFL under a carefully chosen mobility model.

A. Simulation Setup

We implemented ROFL in GloMoSim [2], where 100 nodes, with transmission range of 250 m, are randomly deployed in a simulation area of 1500×1500 m^2. Our goal is to demonstrate the high performance of ROFL as an effective and efficient firewall mechanism especially for MANETs. Without presence of unwanted traffic, no firewall mechanism is able to prove its effectiveness. Hence, in our experiments, one key parameter is the amount of malicious traffic injected into the network. We generate 30 CBR flows, each consisting of data packets of 512 $Bytes$ transmitted between a pair of randomly chosen source and destination nodes at a rate of 4 $pkts/sec$. Thus we observe both intra-group and inter-group communications (w.r.t. RPGM). We vary the amount of malicious traffic from 0% to 100% with an increase of 10% each time. Situations with more than 70% malicious traffic, although unlikely to occur in reality, are used to verify the correctness of ROFL.

Individual random mobility patterns, like Random Waypoint (RWP) or Random Walk (RW), are found to have flaws or limitations [6]. In a MANET, such as in military deployment or disaster recovery, grouped motion behavior is very often. Thus we import Reference Point Group Mobility (RPGM) model generated using BonnMotion [1]. In RPGM [11], nodes are partitioned into groups based on their logical relationship. The movement of an individual node is captured by its group motion plus a random motion from its reference point. To evaluate ROFL's performance as a purely distributed firewall mechanism, we set the averaged number of nodes per group to 10, and vary the maximum speed of node movement from $[0, 5, 10, 15, 20]$ m/s. Notice that 20 m/s corresponds to 72 km/hr, which is a reasonable maximum speed of a vehicle in MANETs, though perhaps too fast for military models.

The performance of ROFL is evaluated against distributed firewalls [4], where firewall policies are centrally computed and pushed to each end host. Thus filtering of malicious traffic is at destinations only. We are particularly interested in the amount of data traffic and control traffic incurred both with and without ROFL; other metrics, such as throughputs, end-to-end delays, etc, are not our focus, as they are generally considered as performance metrics of underlying routing protocols. Each experiment runs for 900 sec with 30 sec pause period. Results collected from each experiment are averaged from 20 runs. (Due to the space limitation, we omit simulation results with ROFL using RWP mobility model.)

B. Performance of Routing Protocols with RPGM

Prior to conducting experiments with ROFL scheme, we evaluated the performance of AODV and OLSR across the set of RPGM models with various degrees of mobility in Figure 2(a). Clearly, both AODV and OLSR didn't achieve their best performance under RPGM. Previous studies on group mobility models [3], [14] demonstrated that compared with flat routing schemes, such as AODV or DSDV, hierarchical or cluster-based routing protocols exhibits much better scalability. Therefore it is not surprising to observe more fluctuation of these performance curves when parameters, like group size or movement pattern, are further tuned for this purpose. ROFL implements packet filtering by taking advantage of underlying routing mechanisms. Thus its performance cannot be completely isolated from the behavior of underlying routing protocols that are used. Results in Figure 2(a) serves as a baseline evaluation for subsequent experiments.

C. Scenario 1. ROFL with AODV

Figure 2(b) depicts the reduction of data packets received at destinations using ROFL comparing against the distributed firewall approach, given that the percentage of malicious traffic present in the network varies from 0% to 100%. The line with $V_{max} = 0$ m/s represents a static ad hoc network; whereas $V_{max} = 20$ m/s demonstrates a highly dynamic MANET. (Data packet loss occurs due to several reasons, such as node mobility, collision at link layer, etc.) Ideally, data packet reduction increases by 10% each time with the growth of malicious traffic. Obviously, with distributed firewalls, malicious packets are not dropped until reaching their final destinations; hence data packets observed at destinations are not affected by malicious traffic present.

In ROFL, unwanted packets are dropped further away from destinations depending on how far ROFL announcements can propagate in the network. We observe a significant reduction of data packets at destinations in Figure 2(b), that clearly increases with the growth of malicious traffic. When nodes start moving faster, we notice slight fluctuations caused by rapidly changing topology. We believe these fluctuations are closely related to the performance of underlying routing protocols rather than ROFL itself (see the baseline scenario in Figure 2(a)). With a maximum speed of 20 m/s, we still observe significant packet reduction following the same trend.

As energy conservation is crucial in MANETs, it is beneficial and desirable to save battery power by reducing unnecessary transmissions. Figure 2(d) and 2(e) illustrate the reduction of packet transmission for data traffic and control traffic respectively. Data packets transmitted in the network counts not only packets generated at their origins but relaying effort made by intermediate nodes. By dropping unwanted data packets as early as possible, we experience significant reduction of data packets with the growth of malicious traffic. In addition, ROFL triggers reduction of control traffic, as route requests (RREQs) from malicious (i.e., unauthorized) nodes are dropped silently by neighbors which have seen that ROFL announcement before. However, reduction of control traffic flattens when there are significant amount of malicious traffic but very little legitimate flows in the network (at $y = 80\%,\ 90\%,\ 100\%$ in Figure 2(e)). Recall that since AODV implements on-demand

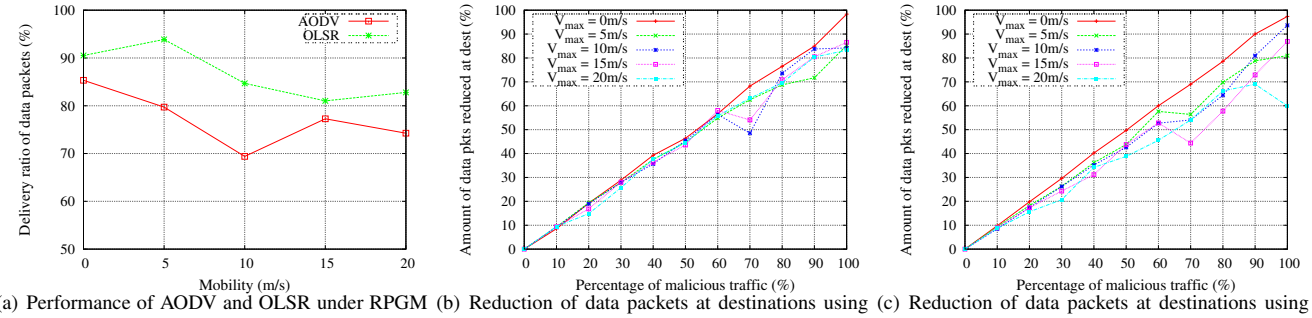

(a) Performance of AODV and OLSR under RPGM mobility model (no malicious traffic).

(b) Reduction of data packets at destinations using ROFL with AODV.

(c) Reduction of data packets at destinations using ROFL with OLSR.

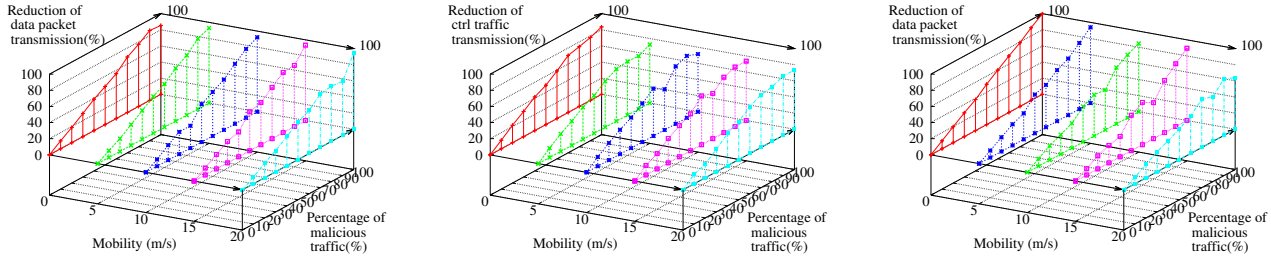

(d) Reduction of data packet transmission using ROFL with AODV.

(e) Reduction of control packet transmission using ROFL with AODV.

(f) Reduction of data packet transmissions using ROFL with OLSR.

Fig. 2: Performance evaluation of ROFL in a simulation environment.

routing strategy, a service-specific route announcement gets disseminated into the network when there exists a permissive routing request (RREQ). With very little legitimate data traffic, hence very few permissive route requests, we cannot benefit further from ROFL. Moreover, increases of nodes mobility only introduces slight degradation of ROFL's performance.

D. Scenario 2. ROFL with OLSR

Figure 2(c) shows the reduction of data packets received at destinations using ROFL with OLSR as the underlying routing protocol, with malicious traffic varying from 0% to 100%. Obviously, ROFL achieves an incremental reduction of received data packets at destinations along with the growth of malicious traffic, even under a highly dynamic network topology. With OLSR, service-specific route announcements originated from service providers are disseminated into the entire network via the flooding of *Topology Control (TC)* messages. Theoretically, with ROFL, each participant in the network has a complete knowledge of services, that they are allowed to access. Thus malicious packets are expected to be dropped very close to their origins. Figure 2(f) shows the reduction of data packet transmissions. Unlike with AODV, ROFL doesn't affect control traffic transmission in OLSR; because OLSR implements table-driven routing approach with routing information periodically flooded to the entire network.

E. Discussion

Experiments with two different MANET routing protocols demonstrate that ROFL outperforms distributed firewalls by effectively reducing the amount of data transmission with various amount of malicious traffic. Implementation overhead of ROFL includes both the extra cost of maintaining routing table with service specific entries and the cost of transmitting control traffic in the network. The former has been discussed analytically in our previous work [21]. With the growth of nodes mobility, we still observe the same trend. As a recap, we summarize the performance of ROFL with small amount of malicious traffic (as that is often the case in reality) under various degree of mobility in Tables I, II, and III.

	Node Mobility				
Malicious Traffic	0 m/s	5 m/s	10 m/s	15 m/s	20 m/s
10%	9.8%	8.5%	8.4%	8.7%	8.4%
20%	18.9%	15.9%	18.6%	17.9%	17.8%

TABLE I: Reduction of data packet transmission using ROFL with AODV, given small amount of malicious traffic.

10%	10.7%	11.6%	11.5%	11.7%	11.8%
20%	21.6%	21.8%	21.4%	19.0%	16.7%

TABLE II: Reduction of control packet transmission using ROFL with AODV, given small amount of malicious traffic.

10%	9.9%	9.4%	8.7%	7.7%	8.0%
20%	19.8%	18.0%	17.5%	16.2%	16.4%

TABLE III: Reduction of data packet transmission using ROFL with OLSR, given small amount of malicious traffic.

Clearly, ROFL achieves high performance in terms of the reduction of data traffic and control traffic transmission regardless of the network mobility. For instance, for AODV, with nodes moving at 10 m/s and a 10% bad traffic rate, ROFL reduced the number of data packet and control packet transmissions by 8.4% and 11.5% respectively; at 20% bad traffic, the reduction were 18.6% and 21.4%. For OLSR, with nodes moving at the same speed, ROFL reduced the amount of data packet transmission by 8.7% (at 10% bad traffic) and 17.5% (at 20% bad traffic) respectively.

Apart from above results, we also observe some interesting phenomena. In the first scenario with AODV, malicious nodes are induced to generate more routing requests (RREQs), as previous ones were silently dropped by neighbors who have seen the relevant ROFL announcement with embedded source

prefix filtering information before. Depending on how aggressive the attacker is, it may generate much more control traffic without being able to establish a valid route. Thus our scheme indeed helps to drain their battery power faster, which can disable the compromised nodes. Possibly, such an abnormal rate of route advertisements can be useful input to an intrusion detection system, though we have not investigated this.

One further aspect is worth mentioning. With group mobility models, if the groups move too far apart ROFL is more likely to result in network partitioning. We do not regard this as a flaw in ROFL; rather, it is a tradeoff between one form of availability (a connected network) and both a second form of availability (battery power) and confidentiality: traffic is not sent through nodes not trusted to handle it. A likely solution would involve some number of trusted, higher-powered transmitters coupled with a hierarchical routing scheme. We have not yet evaluated this alternative. Therefore, we conclude that ROFL is an effective and efficient firewall mechanism for ad hoc networks that are highly dynamic.

VII. Related Work

To the best of our knowledge, there has not been much work on implementing firewall techniques in MANETs, largely due to the rapidly changing network topology. As battery supply remains crucial in MANETs, filtering unwanted traffic as early as possible — which a firewall can do — is quite important.

A survey paper on power and security can be found in [17], where they categorized denial-of-sevice (DoS) attacks and defenses by protocol layers, especially for those attacks that target energy supply of a battery-powered device. However, like many other papers in this field they point out the severity of battery exhaustion attacks but do not provide a useful solution. SANE and Ethane [7], [8] are similar, in that they have many forwarding nodes within the network enforcing a security policy. However, the scheme is intended for centralized policy control, rather than source-specific policy decisions. As such, it requires the central policy controller to be in contact with all forwarding nodes whenever policy changes. Furthermore, it lacks the built-in universal enforcement guarantees provided by the routing mechanisms ROFL relies on; this in turn means that it could not cope well with a MANET environment where the topology frequently changes. There are previous works on en-route packet filtering in sensor networks [18], [22]. However, they took a different approach by focusing on implementing authentication schemes for sensor networks. Filtering decision are made by group checking on some kind of message endorsement generated by a number of sensor nodes, when the packet is en-route to its destination; whereas in ROFL, we implement packet filtering by using underlying routing mechanisms.

VIII. Conclusions

ROFL with source prefix filtering successfully solves the MANET firewall problem. First, it creates boundaries between regions that have different policies, even when the topology — and hence the boundary — changes. Furthermore, it does so in an appropriate fashion: the same policy specifications control both service routing advertisements and traffic acceptance. Second, by dropping unwanted packets very early, there is a large savings in battery power across the entire MANET, thus eliminating a major denial-of-service issue, in both of the major types of MANET routing.

As part of our ongoing work, we take a policy refinement approach by examining the transformation from high level security policies into low level rules that can be implemented using ROFL as the enforcement mechanism. There, we are particularly interesting in adaption to policy changes and how that would affect generation and dissemination of ROFL advertisements in consequence.

References

[1] "Bonnmotion, a mobility scenario generation and analysis tool."

[2] "GloMoSim: Global mobile information systems simulation library."

[3] F. Bai, N. Sadagopan, and A. Helmy, "IMPORTANT: A framework to systematically analyze the Impact of Mobility on Performance of RouTing protocols for Adhoc NeTworks," in *Proc. of Infocom*, 2003.

[4] S. M. Bellovin, "Distributed firewalls," *;login:*, pp. 39–47, November 1999.

[5] B. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Communications of ACM*, vol. 13, no. 7, pp. 422–426, July 1970.

[6] T. Camp, J. Boleng, and V. Davies, "A Survey of Mobility Models for Ad Hoc Network Research," *Wireless Communication and Mobile Computing: Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, vol. 2, no. 5, pp. 483–502, Sep. 2002.

[7] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker, "Ethane: taking control of the enterprise," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 1–12, 2007.

[8] M. Casado, T. Garfinkel, A. Akella, M. J. Freedman, D. Boneh, N. Mckeown, and S. Shenker, "Sane: A protection architecture for enterprise networks," in *Usenix Security Symposium*, 2006.

[9] W. R. Cheswick, S. M. Bellovin, and A. D. Rubin, *Firewalls and Internet Security; Repelling the Wily Hacker*, 2nd ed. Reading, MA: Addison-Wesley, 2003. [Online]. Available: http://www.wilyhacker.com/

[10] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)," Internet Engineering Task Force, RFC 3626, Oct. 2003.

[11] X. Hong, M. Gerla, G. Pei, and C.-C. Chiang, "A Group Mobility Model for Ad Hoc Wireless Networks," *Proc. of the ACM Int. Workshop on Modelling and Simulation of Wireless and Mobile Systems*, 1999.

[12] S. Kent, C. Lynn, J. Mikkelson, and K. Seo, "Secure border gateway protocol (S-BGP) – real world performance and deployment issues," in *Proceedings of the IEEE Network and Distributed System Security Symposium*, February 2000.

[13] S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol (Secure-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582–592, April 2000.

[14] G. Pei, M. Gerla, X. Hong, and C. chuan Chiang, "Wireless hierarchical routing protocol with group mobility," in *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 1999.

[15] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing," Internet Engineering Task Force, RFC 3561.

[16] C. Perkins, *Ad hoc Networking*. Addison-Wesley, 2001.

[17] D. Raymond and S. Midkiff, "Denial-of-service in wireless sensor networks: Attacks and defenses," *IEEE Pervasive Computing*, vol. 7, no. 1, pp. 74–81, Jan–Mar 2008.

[18] F. Ye, H. Luo, S. Lu, L. Zhang, and S. Member, "Statistical en-route filtering of injected false data in sensor networks," in *INFOCOM*, 2005.

[19] H. Zhao and S. M. Bellovin, "Policy algebras for hybrid firewalls," Department of Computer Science, Columbia University, Tech. Rep. CUCS-017-07, March 2007.

[20] ——, "Source prefix filtering in ROFL," Department of Computer Science, Columbia University, Tech. Rep. CUCS-33-09, July 2009.

[21] H. Zhao, C.-K. Chau, and S. M. Bellovin, "ROFL: Routing as the firewall layer," in *New Security Paradigms Workshop*, September 2008.

[22] S. Zhu, S. Setia, S. Jajodia, and P. Ning, "An interleaved hop-by-hop authentication scheme for filtering of injected false data in sensor networks," in *IEEE Symposium on Security and Privacy*, 2004.