A collaborative firewall for wireless ad-hoc social networks

Leonardo Maccari
DISI - University of Trento, Italy
leonardo.maccari@unitn.it

Keywords: collaborative firewall, wireless ad-hoc networks, security, privacy

Abstract: A collaborative firewall can be realized in a multi-hop distributed wireless network when all or some of the nodes in the network agree on a filtering policy and enforce it when routing a packet. Cooperative firewalling introduces many challenges, how to distribute the rules, how to enforce them, how to reduce the global rule-set in order to limit the impact on the network performance. This paper studies the performance of a collaborative firewall when only a subset of the nodes of the ad-hoc network filter the packets. In order to achieve higher performances the integration with OLSR protocol is proposed. Simulations on realistic scenarios are performed and the source code of the simulator is released.

1 Introduction

A firewall is a bastion host that divides two segments of a network with different trust levels. Packet filtering may be applied at network layer or at higher layers (i.e. application firewalls) in order to enforce finer policies for security and performance reasons. In a multi-hop wireless network the concept of bastion host and network segments fade. The network may access the Internet across a gateway which can be used to filter part of the contents, but there is no standard mean of applying network-wide filtering policies inside the multi-hop network.

This paper focuses on the integration of routing and packet filtering in order to implement a collaborative firewall, that is, a best effort network-wide filtering function in which nodes will participate in order to filter out as much of the unwanted traffic as possible. The three main contributions of this paper are:

- tests are performed on real devices to justify the need for optimization with large rule-sets. Our tests show that rule-set size has a serious impact on the average round trip time even with rule-sets smaller than the ones considered in literature.
- The integration with OLSR routing is proposed to limit the collaborative filtering to a subset of the MPR nodes.
- Simulations are performed on realistic mobility and path-loss models. The source code of the simulator is released to the public in order to stimulate the interest in this research topic.

2 Collaborative Firewalling

A collaborative firewall can be implemented in an ad-hoc network if the nodes of the network distribute their rule-sets, i.e. the rules that they want to be applied to the traffic directed to them, or generated by them. Each node in the network will collect the rule-sets of all the other nodes and build a global rule-set made of the union of the single rule-sets that will be used to filter the packets before forwarding them.

A collaborative firewall is useful in a wireless ad-hoc network for mainly three reasons: performance, security and privacy. From a performance point of view it avoids the waste of network resources. If node N_1 tries to connect to node N_2 on a service N_2 does not support, the TCP packets sent to N_2 are wasting resources along all the path from N_1 to N_2. If node N_2 acts as a gateway it may want to limit the services or the nodes that can access the Internet. A collaborative firewall can be used in all the situations in which it is necessary to limit the impact of unwanted traffic. For instance nodes may agree to filter file-sharing applications or other resource greedy applications. From a security point of view it can be a valuable instrument to limit the impact of attacks against the network. Port scans, floods, password brute force on open services can be stopped using specific rules. Note that a node may locally filter the traffic directed to himself, but this has only a limited effect against a flooding attack that can saturate the whole path from the attacker to the victim. Under a privacy point of view an user may want to limit the scope of the communications that are

generated by himself. For instance, the user may wish to send broadcast chat messages only to the nodes that are inside a certain geographical area (assuming the nodes support some localization protocol). Another example is to limit his visibility to a certain distance (in terms of hops) from his position in the topology. One way to achieve this is to modify the routing messages with firewall actions, removing the references to the node outside a certain range (consider that routing messages are not generated only by the user's node).

Some of these functions could be achieved with other means, for instance, group communications can be achieved with group-based cryptography and service discovery protocols can avoid opening unwanted connections. Some other instead are typical of firewalls, like application layer filtering and reactive filtering. A collaborative firewall is a valuable instrument not only because it will achieve the second set of goals but also because it will help solving the others without increased complexity, since firewalling is supported in many operative systems.

A distributed firewall can be imagined as an instrument made of three building blocks:

1. A rule-set generation tool. This can be completely user generated, derived by high-level policies or even generated in a reactive way.
2. A rule-set distribution protocol, which can be embedded in the routing protocol used.
3. A rule-set enforcement policy that will try to limit the number of effectively enforced rules per node in order to save energy and limit the search time.

The most common organization for rule-sets is the linear list. The price to pay for simplicity is the time of the search that is linear with the size of the rules. Many approaches exist to improve the look-up time of the search on large rule-sets using complex data structures. Most of these techniques are imagined for rule-sets where only network and transport layer fields are matched, so the regularity of the patterns to be matched can be used to organize the rule-sets. This is impossible with application-layer filtering, that is needed for many useful traffic filtering functions. Moreover, most of those techniques are aimed at slowly varying rule-sets since they move the complexity from the search operation to the data structure organization. In a mobile ad-hoc network nodes often enter and exit the network and may change their rule-set, so that the cost of keeping the data structure updated may be higher than the cost of a linear search. A review of the difficulties of using standard packet classification techniques in wireless distributed network can be found in (Fantacci et al., 2008).

In a wired network the most common firewall configuration is to set the default rule to drop all the packets and to explicitly permit the allowed services. This approach is hard to implement in ad-hoc networks, since the changes in topology and in rule-sets take some time to propagate in the network. At any time the global rule-set calculated by a node may be outdated and the firewall may take wrong decisions. It is tolerable to have false negatives (i.e. packets that are forwarded by a node even if a rule exists that would drop them) but not to drop valid packets. For this reason the default rule must be set to allow the packets and each rule will produce a drop. As a consequence, for each valid packet the whole rule-set must be scanned.

Cooperative firewalling is a known subject in literature, an initial model has been proposed by Bellovin et al. in (Ioannidis et al., 2000) where the firewall is moved from a bastion host to the endpoints of a still traditional centralized network. Recently, the subject has received more attention, Bellovin proposed a distributed policy enforcement platform (Zhao et al., 2008a; Zhao and Bellovin, 2009; Zhao and Bellovin, 2010; Zhao et al., 2011), as well as other authors (Alicherry et al., 2008). Other works focus on the application of hash functions to speed-up rule-matching (Fantacci et al., 2008; Neira et al., 2008) or on rule-set reduction techniques (Taghizadeh et al., 2011). Similarly, if higher security is required deny-by-default networks filter the communications between two nodes if a previous security handshake has been not performed (Zhang et al., 2008; Alicherry et al., 2009). This paper explores an alternative approach, the possibility of limiting the number of nodes that filter the packets in the network introducing integration with OLSR routing protocol to achieve more efficiency.

3 Considered Scenario

The network model that is the focus of this paper is a mobile ad-hoc social network where a group of people create a multi-hop network to support the typical applications of social networking, similarly to the model presented in (Li et al., 2012). The fragmentation of mobile software platforms and the lack of common protocols have slowed the penetration of ad-hoc networks out of academia but recently there is a concentration of available operative systems and peer-to-peer networking standards are emerging (such as Wi-Fi Direct, IEEE 802.11s). In general there seems to be more interest in this topic from market players

(for instance Qualcomm and Nokia [1]).

Such a network may be composed by a large number of nodes with pedestrian mobility that generate the typical traffic of social networks, such as chat or lightweight file-sharing (photo sharing, profile sharing, vCard sharing...). For this paper, the chosen routing protocol is OLSR since it better fits a situation in which the nodes need to know who is *on-line*. Rule-sets have been piggybacked to OLSR HELLO and TC messages. Instead of moving whole rule-sets only a numeric ID is used. The assumption is that the rule-sets are known in advance on the devices since the applications come with a set of pre-defined rule-sets. The users will choose one and sponsor just the ID in HELLO messages, that in turn will be included in TC messages. Alternatively, the users may periodically sponsor the whole rule-set at large intervals and just the ID on OLSR messages. OLSR can be hardened with cryptographic credentials in order to avoid well-known attacks against routing, as in (Zhao et al., 2008b). This will assure that the rule-set can not be spoofed or altered by some malicious users.

How big a rule-set can be? client hosts are normally equipped with few rules while border routers can be equipped with tens of thousands of rules. Application layer rules are a powerful instrument to avoid the spread of worms and spam, as well to shape the access to services. They can be used on clients in order to have a high security level irrespective of the security of the applications that run on the device. Empirically, a rule-set size between 10 and 50 rules can be considered realistic, but the more services the network offers, the more it can grow, especially for application layer rules. In the simulation scenarios of next section each node randomly chooses a 20 rules rule-set out of a set of ten. With a realistic network size (up to 100 nodes), the global rule-set that is the union of all the rule-sets can grow up to thousands. For simplicity only network and transport layers are used as matching parameters in the simulations.

The mobility model is the one introduced in (Musolesi and Mascolo, 2006). It describes communities of people and shows statistical properties close to the ones measured on real traces. The path-loss model is a dual-slope model enriched with a ray-tracing algorithm (Sommer et al., 2011). A realistic mobility model and the presence of obstacles will generate topologies close to real ones.

To limit the overhead due to firewalling (for instance, in terms of delay introduced) rule-sets can be reduced as in (Taghizadeh et al., 2011), or the number of nodes that act as firewalls can be limited. In both cases false negatives will be generated. To evaluate the performance of the firewall two metrics are used:

M1: counts the false positives hop-by-hop, that is, each time a packet is forwarded from the sender to the destination IP (or to the point when it is dropped).

M2: counts the false positives end-to-end, that is, it counts the packets that arrive to the destination IP.

M1 measures the inefficiency of the firewall as a mean to reduce the traffic overhead for the whole network. M2 measures the inefficiency in filtering traffic directed against a certain host.

4 OLSR-based Collaborative Firewalling

To test the impact of large rule-sets on real devices the delay introduced by the firewall in mobile devices was measured. The tests were performed on two smartphones equipped with 400MHz and 600MHz ARM processors and a dedicated GPU. Even if the market offers faster CPUs for high-end devices, the processors used for the tests can be correctly considered as representing an average processor that can be found on smartphones, e-book readers, tablets.

The tests were performed using a PC connected over a wired USB connection with the mobile device in order to avoid the fluctuations of the Wi-Fi media. This does not influence the test since the aim is to have a measure of the delay due to the processing time of the CPU. The mobile device was equipped with a Linux kernel 2.6 and configured with rule-sets of increasing size. The PC sent 300 UDP ECHO request, the mobile devices answered with ECHO replies after having checked all the rules in the rule-set. For each test the average round trip time (RTT) was calculated and subtracted by the value measured without any rule-set present. The rule-sets were composed for 25% of layer 4 rules and for 75% of layer 7 rules (string matching on the first 50 bytes of the packet).

Figure 1 reports the results. It can be noticed that, as expected, the delay increases linearly with rule-set size and that the delay introduced is relevant even for few hundreds rules. In a network made of one hundred nodes with rule-sets made of 30 rules the global rule-set will be made of 3000 rules which introduces an average delay of 16 ms per hop. In the network used for the simulations an average path length of 3.2 hops has been measured, with standard deviation around 1.5, that accounts for an average RTT increment oscillating between 54ms and 150ms, to be added to the delay introduced by the other layers. In works like (Taghizadeh et al., 2011) very large rule-sets have been used (up to 10000). This test shows

[1] http://pervacy.eu/?p=171

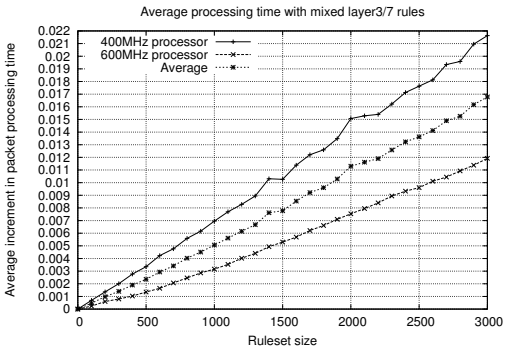

Figure 1: The delay introduced (seconds) changing rule-set sizes.

that actually the impact on network performance is present even with smaller sizes and that increased CPU load can discourage nodes to to participate in the collaborative firewall.

4.1 OLSR and firewall integration

The focus of this section is on a situation in which not all the nodes are willing to participate to the distributed firewall, in particular when only a subset of the nodes are computationally able to handle large rule-sets. In this situation the collaborative firewall is realized with a set of firewall nodes that is a subset of the total (from now on those nodes will be simply called *the firewalls*).

The firewalls will have a larger impact on the network if the routing protocol is modified in order to make them forward more traffic. In OLSR this can be accomplished using a higher *willingness* parameter. The willingness parameter determines how much a node is willing to be elected MPR. Once elected MPR the node will generate TC messages, forward the broadcast traffic and will be preferentially selected as a next hop to any other destination in the network (as suggested in section 10, step 3.2 of RFC 3626). MPRs play a fundamental role in OLSR but since they generate control messages their number should be kept as low as possible. The willingness is expressed with a value that ranges from 0 to 7. With a 0 value the node will never be chosen as an MPR, while with 7 the node will always be chosen. For other values, when a node has two neighbors that can indifferently be elected MPR, the one with highest willingness will be preferred. The value default is 3.

In figures 2a and 2b are reported metrics M1 and M2 for a network of 60 nodes in which only an increasing fraction of randomly chosen nodes is a firewall. Details of the simulations are reported in table 1. Two curves are present in each figure, one refers

Dimension	600m x 600m, 60 nodes
Speed	uniform in [0.5, 2] mps
Rule-set size	20/1200 (single/global)
Total open ports	25
Obstacles	2 obstacles, 70m x 80m
Duration	2000 seconds
Traffic	avg. of 4800 UDP ECHO per run, random destination

Table 1: Simulation parameters

to the overall performance of the distributed firewall when willingness in set to 3 for every node (the curve marked as "Willingness off"), the second refers to the case of choosing 7 for nodes that are active firewall and 2 for nodes that are not firewalls. All the metrics are shown as a fraction of the unwanted packets that get delivered when there is no active firewall.

As a first remark it can be noted that the M2 metric decreases quite sharply even when only 20% of the nodes are firewalls, then it slows down for higher percentages. This is expected since the probability of reaching a random destination follows a binomial distribution with parameter p chosen as the fraction of firewalls against the total number of nodes. When the number of firewalls increases, only packets that travel along short routes will not be filtered. As a consequence, also M1 has the same trend since for each packet that is filtered, in average, a number of hops equal to half of the average route length is saved (if the firewalls are chosen at random).

As a second remark it can be noted that changing the willingness parameter produces a significant improvement. Even when only 10% of the nodes are firewalls, around 40% of the packets are dropped before destination (metric M2), which makes a generic communication very hard to complete. When 20% of the nodes are firewalls a denial of service attack against a node can be mitigated for the 65% of its effect against the target (metric M2) and 45% against the whole network (metric M1).

Figure 2c reports the distribution of the MPR selector set sizes for 4 runs with different parameters. The selector set size is the size of the set of nodes that elected the same node as MPR. This graph is shown in order to verify the impact on the choice of MPR nodes in the network. It can be seen that the impact on the network is minimal and mostly positive. In average there are 30.58 MPR when 30% of the nodes are firewalls against 34.16 when the willingness parameter is not changed. The distribution of the MPR selector set is more polarized, meaning that there are less nodes that have few selectors.

This result can be counter-intuitive, since more

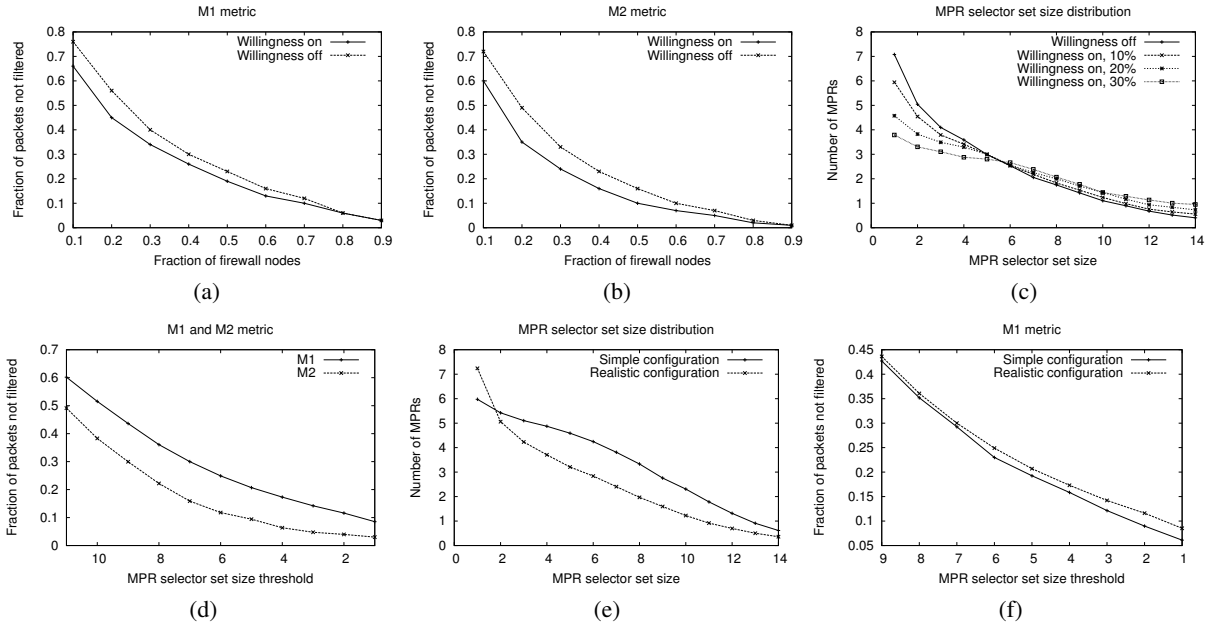

Figure 2: Metrics M1 (a) and M2 (b) increasing the fraction of firewall nodes. MPR selector set size distribution (c). M1 and M2 against the MPR selector set size threshold (d). Comparison between simple and realistic scenarios (e,f)

nodes are forced to be MPR than with standard OLSR. Actually, with OLSR the MPR choice is performed independently by each node which means that even if each node minimizes its own set, the intersection of all the MPR sets is not minimized. Having some nodes with higher willingness polarizes the choice thus reducing the intersection.

4.2 MPR-based filtering

In this section another configuration is considered in which all the nodes have equivalent resources. In order to reduce the overhead introduced in the RTT it is necessary to limit the set of firewalls trying to identify the nodes that are more *central* in the topology.

Centrality is a known concept in networking (Katsaros et al., 2010) that can be estimated with a few metrics. For the purpose of filtering the *shortest-path betweenness* seems the wisest choice, since it represents the fraction of shortest multi-hop routes that pass through a node. OLSR gives an approximated global knowledge of the full topology. It is theoretically possible for each node to calculate and keep up-to-date its betweenness but in practice it is computationally very costly. Betweenness can be approximated limiting the scope of the calculation to only the 2-hop neighbors. An MPR node is, by design, a node that maximizes the betweenness between the selector and its 2-hop neighborhood. Being an MPR is a local esteem of centrality, thus, the chosen strategy consists in ranking the MPRs using the size of their selector set and then implementing the firewall only in the nodes that have a rank over a certain threshold.

In figure 2d M2 and M1 metrics are reported when the threshold is varied from the 11 to 1 (the x axis is inverted in order to make the graph visually comparable with the ones in figure 2a and 2b). It is interesting to note that for a value of 9 and 7, which roughly corresponds to the x-value of 10% and 20% in figure 2b M2 is about half of the case when nodes are selected randomly, meaning that the choice of the MPR selector set size is an useful approximation of a centrality measure. Note that the overall distribution of the MPR selector sets size is known to any node (every TC carries the selector set of an MPR). Given this distribution some topological properties of the network may be derived. The possibility of dynamically choosing the MPR threshold to obtain a wanted value for M1 or M2 is considered as a future work.

4.3 Impact of simulation scenario

In order to stress the importance of using correct models, some simulations have been repeated with a random way-point mobility without obstacles. We refer to this configuration as *simple configuration* and to the model with obstacles, shadowing and Musolesi mobility as the *realistic configuration*. While the two

configurations are functionally comparable (the average size of the routing table differs for 1.5 destinations) the simple configuration generates shorter average paths (2.7 hops versus 3.2) and there is a big difference is in the distribution of the MPR selector set sizes as reported in figure 2e (for a total of 36 against 47). In figure 2e metric M1 is compared, the curves seem very close but the difference in the MPR selector set sizes makes the same threshold value correspond to a different number of firewalls in the network. If instead of comparing the MPR threshold the two curves are compared using the number of firewalls per simulation (roughly this can be done shifting the curve for the simple scenario two units left), then an average 10% difference can be measured. In conclusion this comparison suggests that the MPR selector size threshold used as an approximation of centrality seems to be independent from the topology at this network size and that the simple scenario overestimates the firewall performance compared to a realistic one due to the shorter average path.

5 Conclusions

Cooperative firewalls can positively impact the performances of ad-hoc and mesh networks. They play a useful role in the enforcement of security policies, for reactive security when the network is under attack, to perform traffic shaping or to limit the diffusion of traffic without relying on the support of applications. Nevertheless, large rule-sets are hard to be handled by the limited computation power of common mobile devices. In this paper the impact of the computational overhead has been measured showing that thousands of rules introduce a large delay in RTT. It has been shown that limiting the number of firewalls in the network still allows to filter a high percentage of unwanted traffic and that without adding complexity, the performances can be improved using the information produced by OLSR. Since the behavior of a routing protocol is very scenario-dependent a network simulator with realistic mobility and path-loss models has been realized on top of Omnet++ platform. The source code is available on the project site: www.pervacy.eu.

ACKNOWLEDGEMENTS

Financed by Provincia di Trento under *The Trentino programme of research, training and mobility of post-doctoral researchers, incoming Post-docs 2010* CALL 1, PCOFUND-GA-2008-226070

REFERENCES

Alicherry, M., Keromytis, A., and Stavrou, A. (2008). Distributed firewall for manets. Technical report, Columbia University.

Alicherry, M., Keromytis, A. D., and Stavrou, A. (2009). Evaluating a collaborative defense architecture for manets. In *Conference on Internet multimedia services architecture and applications*, IMSAA.

Fantacci, R., Maccari, L., Ayuso, P., and Gasca, R. (2008). Efficient packet filtering in wireless ad hoc networks. *Communications Magazine, IEEE*, 46(2):104 –110.

Ioannidis, S., Keromytis, A. D., Bellovin, S. M., and Smith, J. M. (2000). Implementing a distributed firewall. In *ACM Conference on Computer and Communications Security*, Athens, Greece.

Katsaros, D., Dimokas, N., and Tassiulas, L. (2010). Social network analysis concepts in the design of wireless ad hoc network protocols. *Network, IEEE*, 24(6):23 –29.

Li, J., Wang, H., and Khan, S. U. (2012). A semantics-based approach to large-scale mobile social networking. *ACM/Springer Mobile Networks and Applications*, 17.

Musolesi, M. and Mascolo, C. (2006). A community based mobility model for ad hoc network research. In *International workshop on Multi-hop ad hoc networks: from theory to reality*, REALMAN '06.

Neira, P., Gasca, R., Maccari, L., and Lefevre, L. (2008). Stateful firewalling for wireless mesh networks. In *New Technologies, Mobility and Security, NTMS '08.*

Sommer, C., Eckhoff, D., German, R., and Dressler, F. (2011). A Computationally Inexpensive Empirical Model of IEEE 802.11p Radio Shadowing in Urban Environments. In *8th Conference on Wireless On demand Network Systems and Services (WONS 2011).*

Taghizadeh, M., Khakpour, A., Liu, A., and Biswas, S. (2011). Collaborative firewalling in wireless networks. In *Joint Conference of the IEEE Computer and Communications Societies, INFOCOM 2011.*

Zhang, H., DeCleene, B., Kurose, J., and Towsley, D. (2008). Bootstrapping deny-by-default access control for mobile ad-hoc networks. In *Military Communications Conference, 2008. MILCOM 2008. IEEE.*

Zhao, H. and Bellovin, S. M. (2009). Source prefix filtering in ROFL. Technical Report CUCS-033-09, Department of Computer Science, Columbia University.

Zhao, H. and Bellovin, S. M. (2010). High performance firewalls in MANETs. In *International Conference on Mobile Ad-hoc and Sensor Networks*.

Zhao, H., Chau, C.-K., and Bellovin, S. M. (2008a). ROFL: Routing as the firewall layer. In *New Security Paradigms Workshop*.

Zhao, H., Lobo, J., Roy, A., and Bellovin, S. M. (2011). Policy refinement of network services for MANETs. In *The 12th IFIP/IEEE International Symposium on Integrated Network Management (IM 2011).*

Zhao, S., Aggarwal, A., Liu, S., and Wu, H. (2008b). A secure routing protocol in proactive security approach for mobile ad-hoc networks. In *Wireless Communications and Networking Conference, WCNC 2008.*