

Black-Hole Attacks in P2P Mobile Networks Discovered through Bayesian Filters*

Jorge Hortelano¹, Carlos T. Calafate¹, Juan Carlos Cano¹,
Massimiliano de Leoni², Pietro Manzoni¹, and Massimo Mecella³

¹ Departamento de Informática de Sistemas y Computadores
Universidad Politécnica de Valencia, Spain

² Department of Mathematics and Computer Science
Eindhoven University of Technology, The Netherlands

³ Dipartimento di Informatica e Sistemistica
SAPIENZA Università di Roma, Italy

Abstract. MANETs (Mobile Ad-hoc NETWORKs) are an example of Peer-to-Peer (P2P) mobile networks in which security attacks, as black-hole ones, may cause serious dangers to the whole system. The watchdog is a well-known sensor usually adopted for detecting black-holes in such networks, but typical watchdogs are characterized by a relatively high number of false positive and negative cases, which can affect the effectiveness and efficiency to deal with intrusions. This paper proposes a novel approach for detecting black-hole attacks and selfish nodes in mobile P2P networks by using a watchdog sensor and a bayesian filtering. We demonstrate the validity of the approach through testing.

1 Introduction

Peer-to-Peer (P2P) mobile networks, such as Mobile Ad hoc NETWORKs (MANETs), are distributed systems composed by wireless mobile nodes that can freely and dynamically self-organise into arbitrary and temporary topologies [8]. These networks have origins in military missions and recovery operations but, in the recent years, a wide range of possible civil applications emerged, e. g., vehicular networks (a.k.a. VANETs), a form of P2P mobile networks used for communication among vehicles and between vehicles and roadside equipment.

The main characteristic of such networks is that they allow different kinds of devices to easily interconnect in areas with no pre-existing communication infrastructure; there exist several protocol specifications, such as AODV [4], that aim to find routing paths between pairs of devices. These allow non-neighbouring nodes to communicate by using intermediate nodes as relays. But the majority of these protocols assume a friendly,

* The work of Hortelano, Calafate, Cano and Manzoni was partially supported by the Ministerio de Educación y Ciencia, Spain, under Grant TIN2008-06441-C02-01, and by the “Ayudas complementarias para proyectos de I+D para grupos de calidad de la Generalitat Valenciana (ACOMP/2010/005)”. The work of de Leoni – performed while he was at SAPIENZA Università di Roma – and Mecella was partly supported by SAPIENZA Università di Roma through the grants AST 2009 “METRO” and FARI 2008, and by the EU through the project SM4All.

reliable and cooperative environment. Therefore, a single malicious node can easily prevent a mobile network from working, and therefore the emerging need for research focused on the provision of practical proposals for securing them [5].

In this context, intrusion detection systems (IDSs) aim at monitoring the activity of the various nodes in the network in order to detect misbehaviours. A basic brick of some IDSs is the watchdog, a collective name for special sensors that can detect selfish nodes and black-hole attackers.

In few words, a black-hole is a type of attack to the network in which a node intends to disrupt the communication with its neighbourhood by attracting all traffic flows in the network and then dropping all packets received without forwarding them to their final destination.

A watchdog is continuously listening neighboring devices for verifying that they, when they are not the final expected recipients, forward packets/messages toward the final destinations. Indeed in MANETs every node is able to analyse the packet headers and learn whether neighbouring nodes are the actual receivers or, conversely, they should forward it to another node on the path to the destination. Devices that do not forward packets for which they are not recipient are considered as misbehaving.

Malicious nodes' detections of current-day watchdogs are affected by several errors due to nodes' mobility and signal noises. This work aims at reducing the number of false positives and negatives by integrating watchdogs with bayesian filtering techniques. Bayesian filters can partly fade the problems by using historical information obtained by the watchdog in the previous time. The technique proposed is independent of the underlying routing protocols and, hence, is widely applicable in several different scenarios of P2P mobile networks. The general approach here proposed, even if demonstrated in the context of MANETs, is in general applicable to a wide spectrum of P2P networks, not only mobile, but also application overlay networks, etc. In a few words, the proposed approach can be summarized as follows:

1. Every node installs a watchdog, thus allowing for detecting misbehaviors (e.g., the number of packets that nodes should forward but that they do not do).
2. The percentage of packets that nodes do not correctly forward is used as input for the bayesian filters in order to predict the percentage of non-forwarded packets in the near future. If such a percentage is higher than a certain threshold, then the node is considered as malicious since it does not behave correctly. Please note that it is not possible to assume that "good" nodes forward correctly all packets because of radio noises, packet losses and other similar characteristics of the aerial medium, which cause some delivery attempts to fail.
3. Every node that detects this malicious behaviour enables consequently appropriate actions to avoid malicious nodes to influence the right network's functioning. Every device can take its own recovery actions, or, conversely, all nodes can reach a consensus on the collaborative actions to deal with the situation. However, this point is out of the scope of this paper: we focus on signalling malicious nodes, assuming another component to take care of mitigating the consequences of such attacks.

The rest of this paper is organised as follows. Section 2 summarises relevant work and the motivation of our work. Section 3 introduces bayesian filters, whereas Section 4 presents our adaptation of the bayesian filters for detecting black-hole attacks.

Section 5 complements our development proposal by explaining the different implementation trade-offs that should be taken into account. Section 6 shows the evaluation performed to validate our mechanism. Finally, Section 7 outlines possible future work.

2 Related Work

The concept of watchdog is not a novelty in the literature. Due to the effectiveness of this methodology and its relatively easy implementation, several proposals use it as the basis of their IDS solutions. Similarly to our approach, Obimbo et al. [3] implement a watchdog that listens neighboring nodes and checks whether they misbehave as they do not forward the packets they are supposed to. In the Pathrater approach [10], each node uses the information provided by watchdogs to rate neighbours. The Routeguard mechanism [6] combines the watchdog and Pathrater solutions to classify each neighbouring node as Fresh, Member, Unstable, Suspect or Malicious. Other approaches like Patwardhan [11] extend the detection capabilities provided by the watchdog with public key encryption and signatures. Marti et al. [12] uses watchdogs in order to prevent malicious nodes from breaking the routing protocols.

But as already pointed out in [10], the problem of all of these solutions is that the used watchdogs report a lot of false detections. Hence, they consider malicious nodes that really are not or, vice versa, actual malicious nodes are not detected as such. The approach we are proposing is more precise as it integrates techniques to mitigate the causes of erroneous detections, that are radio noises and the packet losses.

The appropriateness of bayesian filtering for our intent has been previously confirmed in several fields, such as to implement reputation systems [2] or to predict the nodes' disconnections [9].

3 Bayesian Filtering

Bayesian filters [1] probabilistically estimate a dynamic system's state from noisy observations. At time t , the state is estimated by a random variable θ , which is unknown and this uncertainty is modelled by assuming that θ itself is drawn according to a distribution that is updated as new observations become available. It is called *belief* or $Bel_t(\theta)$. To illustrate this, let's assume that there is a sequence of time-indexed observations z_1, z_2, \dots, z_n . The $Bel_t(\theta)$ is then defined by the posterior density over the random variable θ conditioned on all sensor data available at time t :

$$Bel_t(\theta) = p(\theta|z_1, z_2, \dots, z_t)$$

In our approach, the random variable θ belongs to $[0,1]$. Then we use for the *belief* the distribution $Beta(\alpha, \beta)$ that is suitable for this interval:

$$Bel_t(\theta) = Beta(\alpha_t, \beta_t, \theta)$$

where α and β represent the state of the system, and it is updated according to the following equations:

$$\begin{cases} \alpha_{t+1} = \alpha_t + z_t \\ \beta_{t+1} = \beta_t + z_t \end{cases}$$

The Beta function only needs two parameters that are continuously updated as observations are made or reported. In our approach, the observation z_t represents the information from the watchdog obtained in time interval $[t, t + 1]$ about the percentage of non-forwarded packets.

4 The Bayesian Watchdog

Our approach is based on the information of the incoming packets that devices have not forwarded, nonetheless they should have done so. Our bayesian watchdog relies on some basic assumptions:

1. Every device is equipped with a wireless card that allows for promiscuous mode: any device can listen the packets traversing its neighbourhood and, hence, monitor the activity of one-hop distant nodes.
2. Each node has an implementation of a watchdog sensor, let's indicate as i . The i -th watchdog of a given node monitors the incoming and outgoing traffic of every neighbouring node. In this way, analysing the packet headers, it is able to count the packets that nodes did not forward.
3. Each node has at least three neighbours. We assume a density of the network that makes different paths possible for reaching a destination, and each node is monitored by different neighbours.

The watchdog of device i is in charge of listening the packets' traffic in its neighborhood and verifying whether the fraction of packets that are not correctly forwarded by every neighboring device j . If a given j forwards less than a given fraction of packets than it should, the watchdog considers j as misbehaving. Device i does not know a priori such a fraction for each neighboring node j and, therefore, it defines a random variable $\theta_i(j)$ to estimate it for j . In fact, $\theta_i(j)$ is the viewpoint of device i for what concerns device j . It is worthy highlighting that taking only the last observation is not sufficiently reliable since this could be effected by noise. So the old observations should be considered.

Therefore our watchdog makes use of bayesian filtering, as described in Section 3. Variable $\theta_i(j)$ complies with the Beta distribution with parameters $(\alpha^{(i,j)}, \beta^{(i,j)})$. These parameters are continuously updated with new incoming observations of the fraction of non-forwarded packets. Node i makes periodical observations each t seconds (with t constant) of the behaviour of node j . Let s be the fraction of packets observed by i that are not forwarded by node j in this observation period. Parameters $\alpha^{(i,j)}$ and $\beta^{(i,j)}$ are updated as follows:

$$\begin{cases} \alpha^{(i,j)} := u \cdot \alpha^{(i,j)} + s \\ \beta^{(i,j)} := u \cdot \beta^{(i,j)} + (1 - s) \end{cases} \quad (1)$$

Values $\alpha^{(i,j)}$ and $\beta^{(i,j)}$ are initially set to 1.

The variable u is a fading mechanism for past experiences. This fading mechanism allows for redemption of a neighbour if its behaviour changes to a correct one along the time. This fading mechanism will be useful if there are false positives due to the environmental noise. Greater values for u corresponds to consider the old observations more significantly.

With the beta function defined previously we can define the reputation function of node j on node i $R_i(j)$ using the estimated distribution $Beta(\alpha_i(j), \beta_i(j))$ of variable $\theta_i(j)$:

$$R_i(j) := \begin{cases} 1 & P(\theta_i(j) < \gamma) \\ 0 & P(\theta_i(j) \geq \gamma) \end{cases} \quad (2)$$

where, according to the Probability Theory, the probability that $\theta_i(j)$ is less than γ is computable as the integration of the distribution of $\theta_i(j)$ between $-\infty$ and γ , i.e. in our case between 0 and γ :

$$P(\theta_i(j) < \gamma) = \int_0^\gamma Beta(\alpha_i(j), \beta_i(j))$$

If $R_i(j) = 0$, node i reputes j as malicious. This means that node j is malicious if the estimated fraction of packets that are not correctly forwarded is more than a given value γ , named *tolerance threshold*. This tolerance threshold may be depending on the environmental noise and must be defined for each scenario.

5 Tuning of the Bayesian Watchdog

The bayesian filtering used in our watchdog approach is based on some parameters that need in several cases to be tuned for the specific scenarios through a previous training procedure. In particular, the bayesian filtering depends on the following parameters:

Tolerance threshold γ . An higher value for the tolerance threshold requires more time for the watchdog to detect an attack, as an higher value of alpha is needed on function 2 to set the reputation as malicious; to achieve this higher value of alpha, the bayesian filter needs to perform more observations. But, on the other hand, the watchdog is more robust against environmental noise (less false positives). Conversely, if we set a low value for gamma, some nodes affected by noise would be declared as malicious ones and false positives would appear.

Fading value u . This parameter indicates the weight of the old information obtained by the bayesian watchdog. When closer to 1, the old observations weight similarly to the new ones. In case of the change of the behavior of a given node, since the misbehavior detected in the latest observation periods is mostly as relevant as the good behavior observed in the past, the bayesian filters require more time to learn that the node has changed to a bad behavior. On the other hand, the effects of noise onto the filter become less relevant and they are mitigated by the past observations. Therefore, the percentage of false positives is lower. Clearly, for smaller u 's value, the opposite behavior should be observed.

Updating time. This is the period between two subsequent updates of parameters α and β according to the observations harvested about the packets that, wrongly, are not forwarded. Too frequent updates can cause problem to the bayesian filters if the noise is relatively high. If the the filter's parameter is updated frequently and the packet losses are high, it is likely that the number of packets received in the update period is nearly 0%, thus causing nodes to be wrongly considered as malicious. Conversely, if the observation time is too long, parameters are updated too infrequently and, hence, the time to detect a malicious node may become unacceptable.

6 Evaluation

We have performed several tests using the NS-2 simulator [13] in order to evaluate the approach and to tune some of the parameters of the bayesian filters used inside the watchdog. In our simulation we considered a network of 50 nodes in an area 870x870 meters wide.

We have performed our experiments both considering static and dynamic scenarios. The channel has been modelled in NS-2 as a noisy IEEE 802.11 wireless affected by the fading effect. In the scenarios where nodes have been moving, we have used random way point model provided by the NS-2. As a matter of fact, the mobility affects the accuracy of the watchdog for two main reasons: (i) routes used for the traffic flow need to be recalculated each time the topology changes, causing packet losses; and (ii) if the attacker is moving, there is a possibility that the malicious node moves outside the watchdog's signal range before it is detected. Both characteristics of these scenarios cause false negatives and false positives to be increased.

6.1 Static Scenario

We use random scenarios for validating the implementation of our bayesian watchdog. In the first place we perform several tests to evaluate the behaviour of the watchdog's module in a static scenario. Figure 1 shows how the bayesian watchdog detects the 100% of the attacks independently of the number of attackers that there are in the network, and therefore a 0% of false negatives. The absence of mobility makes the number of false positives negligible. But, that is not a realistic setting.

6.2 Dynamic Scenario

The experiments described below are targeted to find the best tuning of parameters in order to improve the effectiveness. The experiments have been conducted for different motion speeds of the MANET devices, thus verifying how the speed can affect the detection of malicious nodes.

Evaluation of the Tolerance Threshold γ . We perform different tests in scenarios with different mobility speeds and changing the tolerance threshold. Figure 2 shows the results measured. For both of diagrams, the x axis represents the various thresholds tested. In Figure 2(a) and Figure 2(b), the y axes measure respectively the percentage of actual attacks detected and false negatives. For the results' analysis, it seems any




Fig. 1. Actual detections and false positives in a static scenario




Fig. 2. Percentage of (a) actual attacks detected and (b) false positives for different tolerance threshold and for different devices' speed

threshold between 0.75 and 0.85 decreases the false positives while keeping a good rate detection.

A survey of the NS-2 trace shows that an higher value of gamma (closer to 1) causes the bayesian watchdog to be more strict when detecting an attack, decreasing the false positives but also decreasing the percentage of detection. This is caused by the fact that, as discussed in Section 5, the watchdog needs an higher value of alpha to decide if a neighboring node is malicious, and therefore, more time is needed to detect it.

Evaluation of the Fading Value u . The next step is to evaluate what is the influence of the fading value upon the accuracy of detection. Figure 3 shows the results obtained when varying the fading value of the bayesian watchdog. We use a gamma value of 0.85 for these tests, as it seemed the most suitable according to the results of the previous experiments for the tolerance threshold.

As shown in Figure 3, we can see how an high value of fading is more robust against false positives. But, when a node starts behaving maliciously, it takes longer to detect that. Therefore, such a longer time decreases the accuracy of detecting actual attacks. As a result, the optimal fading value may be depending on the needs of the network.

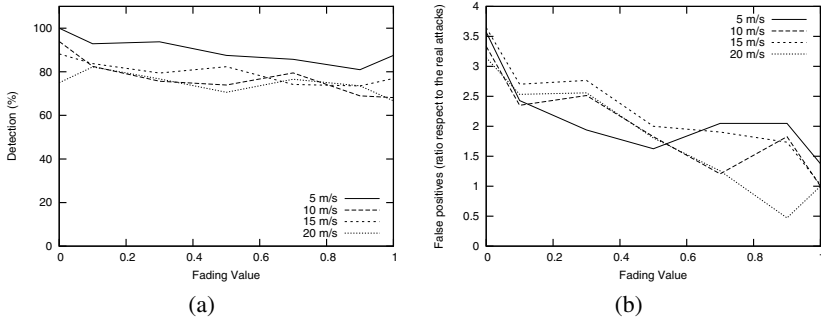


Fig. 3. Percentage of (a) actual attacks detected and (b) false positives for different fading values and different mobility speeds

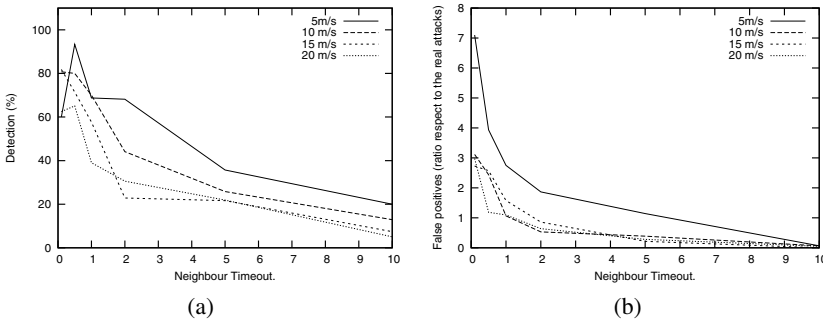


Fig. 4. Percentage of (a) actual attacks detected and (b) false positives when varying the updating time

E.g., if the routing protocol needs to recalculate new routes frequently due to a high value of the node’s speed, a higher value of fading is recommended. Or, if a malicious node performs intermittent attacks, a lower value of fading is needed.

Evaluation of the Update Time. Figure 4 confirms what stated in Section 5: shorter update time of the parameters of the bayesian filter increments the detection of false negatives but also decrements the accuracy for what concerns the false positives.

6.3 Comparing the Bayesian Watchdog with a Standard One

The section proposes a comparison between the standard watchdog and the bayesian watchdog in order to judge whether bayesian filters can really be supportive in the accuracy of the detection of malicious nodes. Specifically, we have used a standard watchdog that was previously implemented [7] (with tolerance threshold of 20%). As far as the bayesian watchdog, the tolerance has been set to 0.85, the fading u to 0.5 and the updating time to 2 seconds.

Figure 5 shows a comparison between the bayesian watchdog and the standard watchdog in a set of scenarios where the degree of mobility is varying. As far as the




Fig. 5. Comparison between both watchdogs with different degrees of mobility: (a) actual attacks detected and (b) false positives

percentage of actual attacks detected, Figure 5(a) shows that the bayesian watchdogs perform mostly at the same level as the standard. Some small improvements have been measured for scenarios where nodes were moving faster. In fact, the bayesian watchdog is less affected by the problem of the mobility as explained at the beginning of this section.

The best improvement has been measured for what concerns false positives that have been decreased of 20%. Indeed, the bayesian filter deals very well with noisy environment such as MANETs. A smaller influence of the noise in the measurement of the non-correctly forwarded packets has resulted in a lower number of false positives.

In addition to a significant decrease of false positives, the bayesian watchdog is also able to detect malicious behaviors more quickly than standard ones.

7 Conclusions and Future Work

The work described in this paper is aimed at increasing the accuracy of the detection of malicious nodes' behaviour in P2P networks, as MANETs. One of the most significant problems of the standard watchdog are concerned with the influence of the noisy observation upon the accuracy. Here we have proposed a new class of watchdogs that rely on bayesian filters. Bayesian filters are broadly used in several scenarios due to their ability to reduce the influence of the noise on the measurements.

In the standard watching, most of the false positives and negatives are caused by the erroneous measurements of the packets that nodes should forward but actually they do not. The erroneous measurements are mostly caused by the unreliability of the wireless medium. Nothing can be done on reducing that. But bayesian filtering deal very well with preventing this node from influencing the judgement of the maliciousness of given devices.

We have devised a technique to integrate bayesian filtering techniques inside the watchdogs and we have conducted some experiments inside an ns-2 implementation to verify the approach. The integration of bayesian filtering inside the watchdogs has decreased the number of false positives detected while the percentage of the detection of the actual attacks has been kept quite high (or, even, slightly improved). As future

work, we intend to provide a concrete implementation of our bayesian watchdog and to perform a deeper experimental phase on the real devices. Moreover, we argue that the approach to detect malicious nodes can be applied also to other P2P networks (e.g., application overlay networks) by suitably modifying the concepts of what it is observed and what is the noisy (transmitted packets in MANETs, could be application messages in overlay networks, etc.)

References

1. Berger, J.O.: *Statistical Decision Theory and Bayesian Analysis*. Springer, Heidelberg (1993)
2. Buchegger, S., Boudec, J.Y.L.: *A robust reputation system for p2p and mobile ad-hoc networks* (2004)
3. Obimbo, C., Arboleda, L.M., Chen, C.Y.: *A Watchdog Enhancement to IDS in MANET*. In: IASTED conference on Wireless Networks (July 2006)
4. C. Perkins, E. Belding-Royer, S. Das: *Ad hoc on-demand distance vector (AODV) routing*. Request for Comments 3561, Network Working Group (July 2003), <http://www.ietf.org/rfc/rfc3561.txt> experimental
5. Yang, H.: *Security in mobile ad hoc networks: challenges and solutions*. *IEEE Wireless Communications* 11(1), 38–47 (2004)
6. Hasswa, A., Zulkernine, M., Hassanein, H.: *Routeguard: an intrusion detection and response system for mobile ad hoc networks*. In: *Wireless And Mobile Computing, Networking And Communications, 2005 (WiMob 2005)*, vol. 3, pp. 336–343. IEEE Computer Society Press, Los Alamitos (August 2005)
7. Hortelano, J., Ruiz, J.C., Manzoni, P.: *Evaluating the usefulness of watchdogs for intrusion detection in VANETs*. In: *ICC 2010 Workshop on Vehicular Networking & Applications, Cape Town, South Africa* (2010)
8. Chlamtac, I., Conti, M., Liu, J.J.: *Mobile ad hoc networking: imperatives and challenges*. *Ad Hoc Networks* 1(1), 13–64 (2003)
9. de Leoni, M., Humayoun, S.R., Mecella, M., Russo, R.: *A bayesian approach for disconnection management in mobile ad-hoc network*. In: *Ubiquitous Computing and Communication Journal* (2008)
10. Marti, S., Giuli, T.J., Lai, K., Baker, M.: *Mitigating routing misbehavior in mobile ad hoc networks*. In: *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking, MobiCom 2000*, pp. 255–265. ACM, New York (2000)
11. Patwardhan, A., Parker, J., Joshi, A., Iorga, M., Karygiannis, T.: *Secure Routing and Intrusion Detection in Ad Hoc Networks*. In: *Proceedings of the 3rd International Conference on Pervasive Computing and Communications*. IEEE Computer Society Press, Los Alamitos (March 2005), main Conference
12. Marti, S., Giuli, T.J., Lai, K., Baker, M.: *Mitigating routing misbehavior in mobile ad hoc networks*. In: *6th MobiCom, Boston, Massachusetts* (August 2000)
13. UC Berkeley, LBL, USC/ISI, and Xerox PARC researchers: *Network Simulator - ns (Version 2)* (1998), <http://www.isi.edu/nsnam/ns/>