A framework for MAC protocol misbehavior detection in wireless networks *

Svetlana Radosavac, John S. Baras
Department of Electrical and Computer Engineering
and The Institute for Systems Research
College Park, MD 20742
{svetlana,baras}@isr.umd.edu

Iordanis Koutsopoulos
Department of Computer and Communications Engineering
University of Thessaly
Volos, Greece
jordan@uth.gr

ABSTRACT

The pervasiveness of wireless devices and the architectural organization of wireless networks in distributed communities, where no notion of trust can be assumed, are the main reasons for the growing interest in the issue of compliance to protocol rules. Reliable and timely detection of deviation from legitimate protocol operation is recognized as a prerequisite for ensuring efficient and fair use of network resources and minimizing performance losses. Nevertheless, the random nature of protocol operation together with the inherent difficulty of monitoring in the open and highly volatile wireless medium poses significant challenges. In this paper, we consider the fundamental problem of detection of node misbehavior at the MAC layer. Starting from a model where the behavior of a node is observable, we cast the problem within a minimax robust detection framework, with the objective to provide a detection rule of optimum performance for the worst-case attack. The performance is measured in terms of required number of observations in order to derive a decision. This framework is meaningful for studying misbehavior because it captures the presence of uncertainty of attacks and concentrates on the attacks that are most significant in terms of incurred performance losses. It also refers to the case of an intelligent attacker that can adapt its policy to avoid being detected. Although the basic model does not include interference, we show that our ideas can be extended to the case where observations are hindered by interference due to concurrent transmissions. We also present some hints for the problem of notifying the rest of the network about a misbehavior event. Our work provides interesting insights and performance bounds and serves as a prelude to a future study that would capture more composite instances of the problem.

*Research supported in part by the U.S. Army Research Office under CIP URI grant No DAAD19-01-1-0494

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
WiSe '05 Cologne, Germany
Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...$5.00.

Categories and Subject Descriptors

C.2.0 [**Computers-Communication Networks**]: General-Security and Protection

General Terms

Design, Security

Keywords

Wireless networks, MAC layer, minmax robust detection, protocol misbehavior

1. INTRODUCTION

Deviation from legitimate protocol operation in wireless networks has received considerable attention from the research community in recent years. The pervasive nature of wireless networks with devices that are gradually becoming essential components in our life-style justifies the rising interest on that issue. In addition, the architectural organization of wireless networks in distributed secluded user communities raises issues of compliance with protocol rules. More often than not, users are clustered in communities that are defined on the basis of proximity, common service or some other common interest. Since such communities are bound to operate without a central supervising entity, no notion of trust can be presupposed.

Furthermore, the increased level of sophistication in the design of protocol components, together with the requirement for flexible and readily reconfigurable protocols has led to the extreme where wireless network adapters and devices have become easily programmable. As a result, it is feasible for a network peer to tamper with software and firmware, modify its wireless interface and network parameters and ultimately abuse the protocol. This situation is referred to as protocol misbehavior. The goals of a misbehaving peer range from exploitation of available network resources for its own benefit up to network disruption. The solution to the problem is the timely and reliable detection of such misbehavior instances, which would eventually lead to network defense and response mechanisms and isolation of the misbehaving peer. However, two difficulties arise: the random nature of some protocols (such as the IEEE 802.11 medium access control one) and the nature of the wireless medium with its inherent volatility. Therefore, it is not easy to distinguish between a peer misbehavior and an occasional protocol malfunction due to a wireless link impairment.

Protocol misbehavior has been studied in various scenarios in different communication layers and under several mathematical frameworks. The authors in [13] focus on MAC layer misbehavior in wireless hot-spot communities. They propose a sequence of conditions on some available observations for testing the extent to which MAC protocol parameters have been manipulated. The advantage of the scheme is its simplicity and easiness of implementation, although in some cases the method can be deceived by cheating peers, as the authors point out. A different line of thought is followed by the authors in [11], where a modification to the IEEE 802.11 MAC protocol is proposed to facilitate the detection of selfish and misbehaving nodes. The approach presupposes a trustworthy receiver, since the latter assigns to the sender the back-off value to be used. The receiver can readily detect potential misbehavior of the sender and accordingly penalize it by providing less favorable access conditions through higher back-off values for subsequent transmissions. A decision about protocol deviation is reached if the observed number of idle slots of the sender is smaller than a prespecified fraction of the allocated back-off. The sender is labeled as misbehaving if it turns out to deviate continuously based on a cumulative metric over a sliding window. This work also presents techniques for handling potential false positives due to the hidden terminal problem and the different channel quality perceived by the sender and the receiver. The work in [5] attempts to prevent scenarios of colluding sender-receiver pairs by ensuring randomness in the course of MAC protocol.

A game-theoretic framework for the same problem at the MAC layer is provided in [4]. Using a dynamic game model, the authors derive the strategy that each node should follow in terms of controlling channel access probability by adjustment of contention window, so that the network reaches its equilibrium. They also provide conditions under which the Nash equilibrium of the network with several misbehaving nodes is Pareto optimal for each node as well. The underlying assumption is that all nodes are within wireless range of each other so as to avoid the hidden terminal problem.

Misbehavior detection has been studied at the network layer for routing protocols as well. The work in [12] presents the watchdog mechanism, which detects nodes that do not forward packets destined for other nodes. The pathrater mechanism evaluates the paths in terms of trustworthiness and helps in avoiding paths with untrusted nodes. The technique presented in [3] aims at detecting malicious nodes by means of neighborhood behavior monitoring and reporting from other nodes. A trust manager, a reputation manager and a path manager aid in information circulation throughput the network, evaluation of appropriateness of paths and establishment of routes that avoid misbehaving nodes. Detection, isolation and penalization of misbehaving nodes are also attained by the technique above.

Node misbehavior can be viewed as a special case of denial-of-service (DoS) attack or equivalently a DoS attack can be considered as an extreme instance of misbehavior. DoS attacks at the MAC layer are a significant threat to availability of network services. This threat is intensified in the presence of the open wireless medium. In [7], the authors study simple DoS attacks at the MAC layer, show their dependence on attacker traffic patterns and deduce that the use of MAC layer fairness can mitigate the effect of such attacks. In [1] the focus is also on DoS attacks against the 802.11 MAC protocol. They describe vulnerabilities of 802.11 and show ways of exploiting them by tampering with normal operation of device firmware.

The nature of wireless networks operation dictates that decisions about the occurrence or not of misbehavior should be taken on-line as observations are revealed and not in a fixed observation interval. This gives rise to the sequential detection problem. A sequential decision rule consists of a stopping time which indicates when to stop observing and a final decision rule that indicates which hypothesis (i.e, occurrence or not of misbehavior) should be selected. A sequential decision rule is efficient if it can provide reliable decision as fast as possible. It has been shown by Wald [15] that the decision rule that minimizes the expected number of required observations to reach a decision over all sequential and non-sequential decision rules is the sequential probability ratio test (SPRT).

The basic feature of attack and misbehavior strategies is that they are entirely unpredictable. In the presence of such uncertainty, it is meaningful to seek models and decision rules that are robust, namely they perform well for a wide range of uncertainty conditions. One useful design philosophy is to apply a minimax formulation and identify the rule that optimizes worst-case performance over the class of allowed uncertainty conditions. The minimax design principle has been successfully applied in signal processing and control systems, where the goal is to design receiver filters of optimal performance with respect to a certain measure (e.g. signal-to-noise-ratio) in the presence of system modeling uncertainties and background noise [10, 14].

In a wireless network, information about the behavior of nodes can become readily available to immediate neighbors through direct observation measurements. If these measurements are compared with their counterparts for normal protocol operation, it is then contingent upon the detection rule to decide whether the protocol is normally executed or not. A minimax formulation translates to finding the detection rule with the minimum required number of observations to reach a decision for the worst instance of misbehavior. Clearly, such a scheme would guarantee a minimum level of performance which is the best minimum level possible over all classes of attacks. In this work, we address the problem of MAC protocol misbehavior detection at a fundamental level and cast it as a minimax robust detection problem. Our work contributes to the current literature by: (i) formulating the misbehavior problem at hand as a minimax robust sequential detection problem that essentially encompasses the case of a sophisticated attacker, (ii) quantifying performance losses incurred by an attack and defining an uncertainty class such that the focus is only on attacks that incur "large enough" performance losses, (iii) obtaining an analytical expression for the worst-case attack and the number of required observations, (iv) establishing an upper bound on number of required samples for detection of any of the attacks of interest, (v) extending the basic model to scenarios with interference due to concurrent transmissions. Our work constitutes a first step towards understanding the structure of the problem, obtaining bounds on achievable performance and characterizing the impact of different system parameters on it.

The rest of the paper is organized as follows. In section II, we discuss the issue of misbehavior in IEEE 802.11 MAC protocol. In section III we present the minimax robust de-

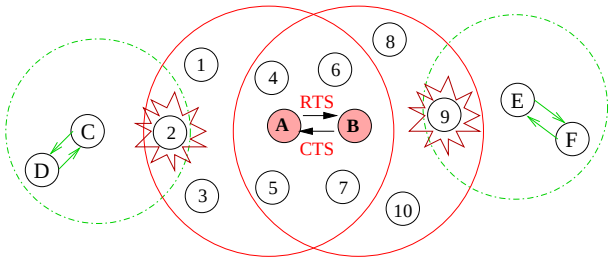

Figure 1: Observer nodes and effect of interference due to concurrent transmissions.

tection model and basic assumptions and demonstrate our approach. In section IV, we discuss some further issues and in section V we show some numerical results. Finally, section VI concludes our study. In subsequent sections, the terms "misbehavior" and "attack", "misbehaving node" and "attacker" will be used interchangeably with the same meaning.

2. MISBEHAVIOR IN THE IEEE 802.11 MAC PROTOCOL

In distributed coordinating function (DCF) of the IEEE 802.11 MAC protocol, coordination of channel access for contending nodes is achieved with carrier sense multiple access with collision avoidance (CSMA/CA) [9]. A node with a packet to transmit selects a random back-off value b uniformly from the set $\{0, 1, \ldots, W-1\}$, where W is the (fixed) size of the contention window. The back-off counter decreases by one at each time slot that is sensed to be idle and the node transmits after b idle slots. In case the channel is perceived to be busy in one slot, the back-off counter stops momentarily. After the back-off counter is decreased to zero, the transmitter can reserve the channel for the duration of data transfer. First, it sends a request-to-send (RTS) packet to the receiver, which responds with a clear-to-send (CTS) packet. Thus, the channel is reserved for the transmission. Both RTS and CTS messages contain the intended duration of data transmission in the duration field. Other hosts overhearing either the RTS or the CTS are required to adjust their network allocation vector (NAV) that indicates the duration for which they will defer transmission. This duration includes the SIFS intervals, data packets and acknowledgment frame following the transmitted data frame. An unsuccessful transmission instance due to collision or interference is denoted by lack of CTS or ACK for the data sent and causes the value of contention window to double. If the transmission is successful, the host resets its contention window to the minimum value W.

IEEE 802.11 DCF favors the node that selects the smallest back-off value among a set of contending nodes. Therefore, a malicious or selfish node may choose not to comply to protocol rules by selecting small back-off intervals, thereby gaining significant advantage in channel sharing over regularly behaving, honest nodes. Moreover, due to the exponential increase of the contention window after each unsuccessful transmission, non-malicious nodes are forced to select their future back-offs from larger intervals after every access failure. Therefore the chance of their accessing the channel becomes even smaller. Apart from intentional selection of small back-off values, a node can deviate from the MAC protocol in other ways as well. He can choose a smaller size of contention window or he may wait for shorter interval than DIFS, or reserve the channel for larger interval than the maximum allowed network allocation vector (NAV) duration. In this work, we will adhere to protocol deviations that occur due to manipulation of the back-off value.

The nodes that are instructed by the protocol to defer transmission are able to overhear transmissions from nodes whose transmission range they reside in. Therefore, silenced nodes can observe the behavior of transmitting nodes. The question that arises is whether there exists a way to take advantage of this observation capability and use it to identify potential misbehavior instances. If observations indicate a misbehavior event, the observer nodes should notify the rest of the network about this situation or could launch a response action in order to isolate the misbehaving nodes. Detecting misbehavior is not straightforward even in the simplest case, namely that of unobstructed observations. The difficulty stems primarily from the non-deterministic nature of the access protocol that does not lead to a straightforward way of distinguishing between a legitimate sender, that happens to select small back-offs, and a misbehaving node that maliciously selects small back-offs. The open wireless medium and the different perceived channel conditions at different locations add to the difficulty of the problem. Additional challenges arise in the presence of interference due to ongoing concurrent transmissions.

Fig. 1 depicts a scenario where node A or B is malicious. At this stage, we assume that A is the only misbehaving node and that no other node in its vicinity transmits. We defer discussion about the collusion between nodes A and B for a subsequent section. We assume that nodes have clocks that are synchronized through the use of GPS devices. Additional issues arising from errors in clock synchronization will be investigated elsewhere. Node A accesses the channel by using a randomly selected back-off value within its contention window. When the back-off counter decreases to zero, A sends an RTS to B, which replies with a CTS. Node A's RTS message silences nodes 1 to 7, which are in A's transmission radius. Similarly, node B's CTS silences nodes 4 to 10. Following the RTS-CTS handshake, A sends a data segment to B. After the transmission is over, A attempts to access the channel anew by selecting a back-off value again and the procedure repeats. Nodes 1-10 can hear the transmissions of nodes A or B, or of both, depending on whose transmission radius they reside in. Consider the i-th transmission of node A. A node in its transmission range finds time point t_i of RTS packet reception from

$$t_i = T_{i-1} + T_{\text{DIFS}} + b_i, \ i > 1, \tag{1}$$

where T_{i-1} denotes the end time point of reception of the previous data segment and b_i is the random back-off value. Thus, the back-off values can be easily derived. Note that the back-off value before transmission of the first data segment cannot be found since there does not exist any previous reference point to compare it to. A node within transmission range of B can also compute the back-off used by A by using as a reference the time point of reception of the overheard ACK from node B for the previous data segment. Then, a node can measure time point t_i' of CTS packet reception and compute the back-off of node A by using

$$t_i' = T_{\text{ACK},i-1} + T_{\text{DIFS}} + b_i + T_{\text{RTS}} + T_{\text{SIFS}}, \ i > 1. \tag{2}$$

Similarly with the RTS, the first back-off value cannot be found. Clearly, the entire sequence of back-offs of node A is observable in this fashion. It should also be noted that the identity of the node who uses those back-offs (which could be potentially a misbehaving one) is revealed in the corresponding fields of RTS or CTS messages.

We now proceed to describe two scenarios in which observations of nodes 1-3 and 8-9 are hindered by interference and hence correctness of observations is influenced.

1. *Interference due to concurrent transmissions.* Assume that node C has obtained access to the channel and therefore node 2 is silenced. Node C is in the process of transmitting data packets to node D. If observer node 2 is within transmission range of C, C's transmission is overheard by node 2. Clearly, the ongoing transmission of C is experienced as interference at node 2 and obstructs node 2's observations. In case of significant interference level, node 2 may not be able to obtain the timing of received RTS of node A and find the back-off value. Additional ongoing transmissions increase the perceived interference level. Evidently, obstructed measurements due to interference create additional problems in detecting misbehavior, as will be seen in the sequel. The extent to which observations of node 2 are influenced by interference depends on the relative proximity of 2 to nodes A and to the interfering nodes, since the received signal strength of the RTS packet and the interference is a function of signal strength decay with distance.

2. *Interference due to simultaneous channel access.* Node 2 that is silenced by A's RTS observes the sequence of back-offs of node A. If node 2 is in the interference range of node C and C is out of the interference range of A, C may attempt to access the channel at the same time. If the RTS packets from nodes A and C overlap in time when received at node 2, node 2 receives a garbled packet and cannot distinguish neither the transmitter identity nor the packet reception time.

Interference from concurrent data transmissions and simultaneous channel access also affects measurements of nodes within the transmission range of node B. Both types of impairments lead to difficulties in misbehavior detection because they cause corruption of measurements. The probability of the second type of impairment is admittedly much lower than that of the first type, since it requires that nodes A and C access the channel almost at the same time. Although this problem is different from the first one, we will elaborate on obstruction of observations owing only to the first scenario.

A comment about the effect of misbehavior in a network-wide scale is in place here. Each node within transmission range of a malicious node increases its contention window exponentially after each unsuccessful transmission attempt. The same holds for nodes which are located out of the transmitter's range but are able to transmit to nodes that are silenced by the transmitter (in our case, nodes C and E). They may constantly attempt to communicate with silenced nodes and consequently increase their contention windows. In that respect, the effect of a malicious node spreads in an area much larger than their transmission range and may affect channel access of nodes throughout that area.

Another arising issue is the notification of the rest of the network about the misbehavior. Although all nodes within transmission range of nodes A and B above can deduce potential misbehavior, the nature of IEEE 802.11 MAC protocol prohibits them from obtaining access to the channel and transmitting notification information. In a subsequent section, we present a practical method to achieve this goal.

3. MINIMAX ROBUST MISBEHAVIOR DETECTION

In this section we present our approach for misbehavior detection when observations are not obstructed by interference. In section IV, analyze the scenario in the presence of interference due to ongoing concurrent transmissions.

3.1 Problem motivation and sequential detection

We focus on monitoring the behavior of node A for the single-hop communication with node B in figure 1. We assume that any node within the transmission range of A or B observes the same sequence of measurements of back-off values used by A. Since the sequence of observations is the same, the procedure that will be described in the sequel can take place in any of these observer nodes. Since the back-off measurements are enhanced by an additional sample each time A attempts to access the channel, an on-line sequential scheme is suitable for the nature of the problem. The basis of such a scheme is a sequential detection test that is implemented at an observer node. The objective of the detection test is to derive a decision as to whether or not a misbehavior occurs as fast as possible, namely with the least possible number of observation samples. Since the observation samples are random variables, the number of required samples for taking a decision is a random variable as well.

The probability of false alarm P_{FA} and the probability of missed detection P_M constitute inherent tradeoffs in a detection scheme, in the sense that a faster decision unavoidably leads to higher values of these probabilities while lower values are attained with the expense of detection delay. For given values of P_{FA} and P_M, the detection test that minimizes average number of required observations (and thus the average delay) to reach a decision among all sequential and non-sequential tests for which P_{FA} and P_M do not exceed the predefined values above is Wald's Sequential Probability Ratio Test (SPRT) [15]. When SPRT is used for sequential testing between two hypotheses concerning two probability distributions SPRT is optimal in that sense as well [6].

SPRT collects observations until significant evidence in favor of one of the two hypotheses is accumulated. After each observation at the k-th stage, we choose between the following options: accept one or the other hypothesis and stop collecting observations, or defer decision for the moment and obtain observation $k+1$. In SPRT, there exist two thresholds a and b of SPRT that aid the decision. The figure of merit at each step is the logarithm of the likelihood ratio of the accumulated sample vector until that stage. For the case of testing between hypotheses $\mathbf{H}_0$ and $\mathbf{H}_1$ that involve continuous probability density functions f_0 and f_1, the logarithm of likelihood ratio at stage k with accumulated samples $x_1, \ldots, x_k$ is

$$S_k = \ln \frac{f_1(x_1, \ldots, x_k)}{f_0(x_1, \ldots, x_k)}, \tag{3}$$

where $f_i(x_1, \ldots, x_k)$ is the joint probability density function of data $(x_1, \ldots, x_k)$ based on hypothesis $\mathbf{H}_i$, $i = 0, 1$. If the observation samples are statistically independent

$$S_k = \sum_{j=1}^{k} \Lambda_j = \sum_{j=1}^{k} \ln \frac{f_1(x_j)}{f_0(x_j)}, \tag{4}$$

with $f_i(\cdot)$ the probability density function of hypothesis $\mathbf{H}_i$, $i = 0, 1$. The decision is taken based on the criteria:

$$\begin{aligned} S_k \geq a &\Rightarrow \text{accept } \mathbf{H}_1, \\ S_k < b &\Rightarrow \text{accept } \mathbf{H}_0, \\ b \leq S_k < a &\Rightarrow \text{take another observation.} \end{aligned} \tag{5}$$

Thresholds a and b depend on the specified values of P_{FA} and P_M, as will be explained in the sequel.

Our approach is based on sequential detection. However, the main idea is that it places emphasis on the class of attacks that incur larger gain for the attacker (they result in higher chances of channel access). An attack in that class would have most devastating effects for the network, in the sense that it would deny channel access to the other nodes and would lead to unfair sharing of the channel. Besides, if we assume that the detection of an attack is followed by communication of the attack event further in the network so as to launch a network response, it would be rather inefficient for the algorithm to consider less significant (and potentially more frequent) attacks and initiate responses for them. Instead, it is meaningful for the detection system to focus on encountering the most significant attacks and at the same time not to consume resources of any kind (processor power, energy, time or bandwidth) for dealing with attacks whose effect on performance is rather marginal.

3.2 Minimax robust detection approach : Definition of uncertainty class

Previously, we stressed the sequential nature of our approach and the implicit need to consider most significant attacks. The approach should also cope with the encountered (statistically) uncertain operational environment of a wireless network, namely the random nature of protocols and the unpredictable misbehavior or attack instances. Hence, it is desirable to rely on robust detection rules that would perform well regardless of uncertain conditions. In this work, we adopt the minimax robust detection approach where the goal is to optimize performance for the worst-case instance of uncertainty. More specifically, the goal is to identify the least favorable operating point of a system in the presence of uncertainty and subsequently find the strategy the optimizes system performance when operating in that point. In our case, the least favorable operating point corresponds to the worst-case instance of an attack and the optimal strategy amounts to the optimal detection rule. System performance is measured in terms of number of required observation samples to derive a decision.

A basic notion in minimax approaches is that of saddle point. A strategy (detection rule) d^* and an operating point (attack) f^* in the uncertainty class form a saddle point if:

1. For the attack f^*, any detection rule d other than d^* has worse performance. Namely d^* is the optimal detection rule for attack f^* in terms of number of minimum number of required observations.
2. For the detection rule d^*, any attack f other than f^* gives better performance. Namely, detection rule d^* has its worst performance for attack f^*.

Implicit in the minimax approach is the assumption that the attacker has full knowledge of the employed detection rule. Thus, it can create a misbehavior strategy that maximizes the number of required samples for misbehavior detection delaying the detection as much as possible. Therefore, our approach refers to the case of an intelligent attacker that can adapt its misbehavior policy so as to avoid detection. One issue that needs to be clarified is the structure of this attack strategy. Subsequently, by deriving the detection rule and the performance for that case, we can obtain an upper bound on performance over all possible attacks.

According to the IEEE 802.11 MAC standard, the back-off for each legitimate node is selected from a set of values in a contention window interval based on uniform distribution. The length of contention window is $2^i W$ for the ith retransmission attempt, where W is the minimum contention window. In general, some back-off values will be selected uniformly from $[0, W]$ and others will be selected uniformly from intervals $[0, 2^i W]$, for $i = 1, \ldots, I_{\max}$ where $I_{\max}$ is the maximum number of re-transmission attempts. Without loss of generality, we can scale down a back-off value that is selected uniformly in $[0, 2^i W]$ by a factor of 2^i, so that all back-offs can be considered to be uniformly selected from $[0, W]$. This scaling property emerges from the linear cumulative distribution function of the uniform distribution. An attack strategy is mapped to a probability density function based on which the attacker selects the back-off value. Although the possible back-off values are discrete, without loss of generality we use continuous distributions to represent attacks in order to facilitate mathematical treatment and to demonstrate better the problem intuition. We consider continuously back-logged nodes that always have packets to send. Thus, the gain of the attacker is signified by the percentage of time in which it obtains access to the medium. This in turn depends directly on the relative values of back-offs used by the attacker and by the legitimate nodes. In particular, the attacker competes with the node that has selected the smallest back-off value out of all nodes.

Assume that a misbehaving and legitimate node intend to access the channel. In order to have a fair basis for comparison, assume that they start their back-off timers at the same time and that none of the counters freezes due to a perceived busy channel. Let the random variable Y stand for the back-off value of legitimate user, hence it is uniformly distributed in $[0, W]$. Also, let the random variable X stand for the misbehaving node (attacker), so that it has unknown pdf $f(x)$ with support $[0, W]$. The relative advantage of the attacker is quantified as the probability of accessing the channel, or equivalently the probability that its back-off is smaller than that of the legitimate node, $\Pr(X < Y)$. Starting from

$$\Pr(X < Y) = \int_0^W P(Y > X | X = x) f(x)\, dx \tag{6}$$

and using elementary probability, we obtain

$$\Pr(X < Y) = 1 - \frac{\mathbb{E}[X]}{W}, \tag{7}$$

where $\mathbb{E}[\,\cdot\,]$ denotes expectation of a random variable.

Suppose that both nodes were legitimate. If p is the access probability of each node, then the probability of successful

channel access is $p(1-p)$. This is maximized for $p^* = 1/2$ for each node. Now, if one node is the attacker, it receives gain from its attack if $\Pr(X < Y) > 1/2$ or equivalently $\mathbb{E}[X] < W/2$. This implies a shift of X to smaller back-offs compared to the uniform Y for which $\mathbb{E}[Y] = W/2$. The attack strategies of interest are the ones with large enough incurred benefit. In order to quantify this, let ϵ be a positive number and define class of attacks

$$\mathcal{F}_\epsilon = \left\{ f(x) : \int_0^W x f(x)\, dx \leq \frac{W}{2} - \epsilon \right\}. \tag{8}$$

This class includes attacks for which the incurred relative gain compared to legitimate operation exceeds a certain amount. The criterion above is equivalent to $\Pr(X < Y) > (1/2) + \epsilon/W$. The class $\mathcal{F}_\epsilon$ is the uncertainty class of the robust approach and the parameter ϵ is a tunable parameter. By defining the class $\mathcal{F}_\epsilon$, we imply that the detection scheme should focus on attacks with larger impact to system performance and not on small-scale or short-term attacks.

3.3 Minimax robust detection approach: Derivation of the worst-case attack

Hypothesis $\mathbf{H}_0$ concerns legitimate operation and thus the corresponding pdf f_0 is the uniform one. Hypothesis $\mathbf{H}_1$ corresponds to misbehavior with unknown pdf $f(\cdot)$.

The objective of a detection rule is to minimize the number of the required observation samples N so as to derive a decision regarding the existence or not of misbehavior. The performance of a detection scheme is quantified by the average number of samples $\mathbb{E}[N]$ needed until a decision is reached, where the average is taken with respect to the distribution of the observations. This number is a function of the adopted decision rule d and the attack p.d.f f, that is

$$\mathbb{E}[N] = \phi(d, f). \tag{9}$$

Let $\mathcal{D}$ denote the class of all (sequential and non-sequential) statistical hypothesis tests for which the false alarm and missed detection probabilities do not exceed some specified levels P_{FA} and P_M respectively. Generally, a hypothesis test consists of a decision function $g(\cdot)$ that acts on a set of k observations and takes values in the set of hypotheses, i.e, $g : \Omega^k \rightarrow \{\mathbf{H}_0, \mathbf{H}_1\}$. Let $\mathcal{G}$ be the space of all decision functions. A sequential test is a pair $(g_T(\cdot), T)$ where T is the stopping time and $g_T(\cdot)$ is the decision function that acts on observation samples collected up to time T. Thus, $\mathcal{D} = \mathcal{G} \bigcup (\mathcal{G} \times [0, \infty])$. In the context of the minimax robust detection framework, the problem is to optimize performance in the presence of worst-case attack, that is find

$$\mathbb{E}[N]^* = \min_{d \in \mathcal{D}} \max_{f \in \mathcal{F}_\epsilon} \phi(d, f)\,, \tag{10}$$

assuming that finite number of samples are needed (otherwise the "min-max" notation should change to "inf-sup"). We proceed to a formal definition of saddle point.

DEFINITION 1. *A pair (d^*, f^*) is called a saddle point of the function ϕ if*

$$\phi(d^*, f) \leq \phi(d^*, f^*) \leq \phi(d, f^*)\ \forall d \in \mathcal{D},\ \ \forall f \in \mathcal{F}_\epsilon. \tag{11}$$

A saddle point (d^*, f^*) of ϕ consists of a detection test d^* and an attack distribution f^*. Equation (11) is a formal statement of properties 1 and 2 that were mentioned in Sect. III-B. In order to facilitate solution of problem (10), we find the saddle point of ϕ. First, recall that the optimal detection test in the sense of minimizing expected number of samples needed for detection is SPRT. This means that SPRT is the test $d^* \in \mathcal{D}$, such that for a fixed (but unknown) attack f we have $\phi(d^*, f) \leq \phi(d, f)$ for all other tests $d \in \mathcal{D}$. The inequality above also holds for $f = f^*$, and hence the second inequality in (11) has been established.

We now prove the first inequality. Assuming that SPRT is used, we seek an attack distribution f^* such that $\phi(d^*, f^*) \geq \phi(d^*, f)$ for all other attacks $f \in \mathcal{F}_\epsilon$. In order to find f^*, we need an expression for the required average sample number (ASN) $\mathbb{E}[S_N]$ of SPRT. From Wald's identity [15]

$$\mathbb{E}[S_N] = \mathbb{E}[N] \times \mathbb{E}[\Lambda], \tag{12}$$

where $\mathbb{E}[\Lambda]$ is the expected value of the logarithm of likelihood ratio. By using a similar derivation as the one in [8, pp.339-340], we derive the following inequalities

$$1 - P_M \geq e^a P_{FA} \ \text{ and } \ P_M \leq e^b (1 - P_{FA}), \tag{13}$$

where a and b are the thresholds of SPRT. When the average number of required observations is very large, the increments Λ_j in the logarithm of likelihood ratio are also small. Therefore, when the test terminates with selection of hypothesis $\mathbf{H}_1$, S_N will be slightly larger than a, while when it terminates with selection of $\mathbf{H}_0$, S_N will be very close to b. Therefore, the above inequalities hold to a good approximation as equalities. Under this assumption, the decision levels a and b that are required for attaining performance (P_{FA}, P_M) are given by,

$$a = \ln \frac{1 - P_M}{P_{FA}} \ \text{ and } \ b = \ln \frac{P_M}{1 - P_{FA}}. \tag{14}$$

Following the derivations of [15, 8],

$$\mathbb{E}[S_N] = a P_D + b(1 - P_D) \tag{15}$$

where $P_D = 1 - P_M$ is the probability of detection of SPRT. Hence, the average number of samples is

$$\mathbb{E}[N] = \frac{\mathbb{E}[S_N]}{\mathbb{E}[\Lambda]} = \frac{C}{\mathbb{E}\left[\ln \frac{f(X)}{f_0(X)}\right]} \tag{16}$$

where $f_0(x) = 1/W$ denoting uniform distribution of normal operation and the expectation of denominator is with respect to the unknown attack distribution f. Since C is a constant, the problem of finding the attack that maximizes the required number of observations reduces to problem,

$$\min_f \int_0^W f(x) \ln \frac{f(x)}{f_0(x)}\, dx \tag{17}$$

subject to the constraints,

$$\int_0^W f(x) dx = 1 \text{ and } \int_0^W x f(x) dx \leq \frac{W}{2} - \epsilon. \tag{18}$$

The first constraint exists since f is a pdf and the second one is because $f \in \mathcal{F}_\epsilon$. By applying the Karush-Kuhn-Tucker (KKT) conditions, we find that the function f^* has the form

$$f^*(x) = e^{-\lambda - 1} e^{-\mu x},\ \mu > 0, \tag{19}$$

where λ and μ are the Lagrange multipliers that correspond to the constraints and are functions of W and ϵ only. These

can be obtained by the system of equations:

$$\frac{e^{-\lambda-1}}{\mu}\left(\frac{1-e^{-\mu W}}{\mu}-We^{-\mu W}\right) = \frac{W}{2}-\epsilon \quad (20)$$
$$e^{-\mu W} = 1-\mu e^{\lambda+1}$$

Interestingly, the result above shows that the worst-case attack distribution f^* in terms of maximizing number of required samples has exponential density. Since $\phi(d^*, f^*) \geq \phi(d^*, f)$ for all $f \in \mathcal{F}_\epsilon$, we proved the left inequality in (11). We have now shown that pair (d^*, f^*), where d^* is SPRT and $f^*(x)$ is the exponential density constitute a saddle point of ϕ. This means that the so-called minimax equality holds and we can interchange the order of min and sup in the optimization problem above [2]. Then, the problem

$$\max_{f\in\mathcal{F}_\epsilon}\min_{d\in\mathcal{D}}\phi(d,f) \quad (21)$$

has the same solution with (10). As a side remark, note that the derived exponential pdf has maximum differential entropy over all pdf's in class $\mathcal{F}_\epsilon$.

As was mentioned above, the minimax robust detection approach captures the case of an intelligent adaptive attacker. The SPRT algorithm is part of the intrusion detection system module that resides at an observer node. With the method outlined in Sect. II, an observer node monitors the behavior of another node with the objective to derive a decision as fast as possible. In other words the observer (and hence the system) attempts to minimize the number of required samples so as to improve its payoff in terms of improved chances for channel access. On the other hand, an intelligent attacker that knows the detection algorithm attempts to delay this decision as much as possible so as to increase his own benefit in terms of chances for channel access. The attacker aims at a strategy that causes performance degradation for other nodes by remaining undetected.

4. FURTHER ISSUES

4.1 Network notification

In the previous sections we discussed the issue of misbehavior detection of node A with the help of observer nodes that reside within transmission range of A or the receiver B (Fig. 1). During the misbehavior, the observer nodes are silenced due to exchange of RTS and CTS messages and are prevented from accessing the channel. A natural question that arises after misbehavior is detected concerns notifying the network about the attack. This is an essential step that needs to be accomplished so that the network learns about the attack and can initiate a response or isolate the attacker. It will also prevent further propagation of misbehavior in the network. The challenge lies in the fact that the nature of 802.11 MAC does not provide the observer nodes in transmission range of A or B an opportunity to communicate their messages unless a separate control channel is maintained, in which case the notification message can be transmitted in that channel.

In the absence of this control channel, an observer nodes that detect protocol deviations should get an opportunity to transmit a notification message as fast as possible. A practical solution to this problem would be for receiver B to adjust the transmit power level of the CTS message. Based on the received signal strength of received RTS packets and data from the attacker A, node B can estimate the distance of A. Subsequently, it can reduce the transmit power of the CTS to the minimum necessary level such that A remains in transmission range of B. Recall that the 802.11 MAC protocol has the capability of controlling transmit power.

The reduction of transmit power level releases some observer nodes that were previously silenced. These are now out of transmission range of CTS and can obtain access to the channel. Therefore, they can communicate information about the identity of the misbehaving node to their peers. We note that the released observer nodes also need to adjust their power levels so as not to cause interference to unreleased observers. Also, different released observers will have different observations since the latter are location-dependent (see subsection 4.3 below). This mechanism could be implemented as part of the IDS algorithm and works when node B and each of the observer nodes within B's range obtain the same (unobstructed) back-off observation samples. Further study is needed for the case where observations differ among different observer nodes in B's range due to different local perceived interference conditions.

This mechanism provides a paradigm of cross-layer coordination and information exchange. The physical-layer based transmit power adaptation can offer an additional degree of freedom to the MAC layer-based detection system and can enable network notification. Without the help from the physical layer, this task would be extremely difficult.

4.2 Colluding nodes

The problem treatment above assumed the existence of a single attacker and did not include the scenario of colluding nodes. In the communication scenario of figure 1, nodes A and B may collude if node B receives the RTS messages from attacker A and it intentionally delays the CTS message by some amount of time. This scenario exploits the nature of exponential backoff by choosing small backoff values and additionally breaks the protocol rules by waiting longer than SIFS between RTS and CTS signals. In this case, the observer nodes within transmission range of B perceive erroneous, higher back-off values from node A. As a result, they cannot detect potential misbehavior of A. They also cannot determine the maliciousness of receiver B. However, the remaining observes that can overhear both A and B can detect misbehavior with higher probability since it is not allowed to wait for periods that are longer than SIFS between RTS and CTS control signals.

In this fashion, a colluding node B decreases the number of observer nodes that can provide correct measurements. Misbehavior of node A can thus be observed only by nodes within transmission range of A. On the other hand, only observers residing within range of both A and B can monitor both A and B and therefore detect collusion of A and B by using a detection scheme similar to the one outlined in previous sections. The detection method can have two separate tests: one acting on the observed back-offs of A and one for measuring timing delays from the receiver in issuing CTS messages. The latter test should be a threshold rule, since normally the delay before issuing a CTS is deterministic. The decision about collusion is taken after combining results from both tests. However, note that in the event of collusion the mechanism of the previous subsection cannot help in network notification.

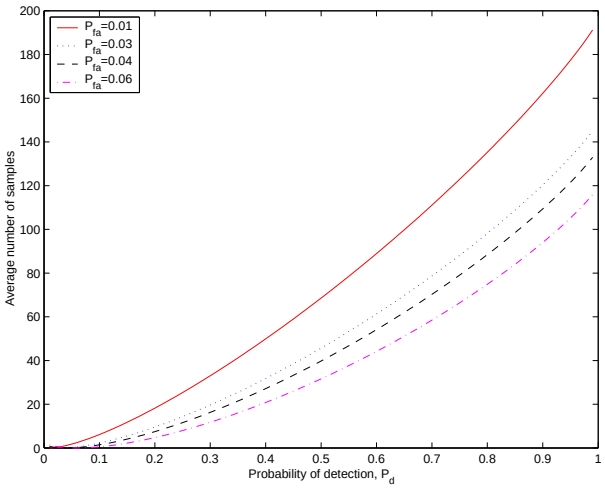

Figure 2: Average number of samples needed for misbehavior detection versus P_D in absence of interference for the worst-case attack.

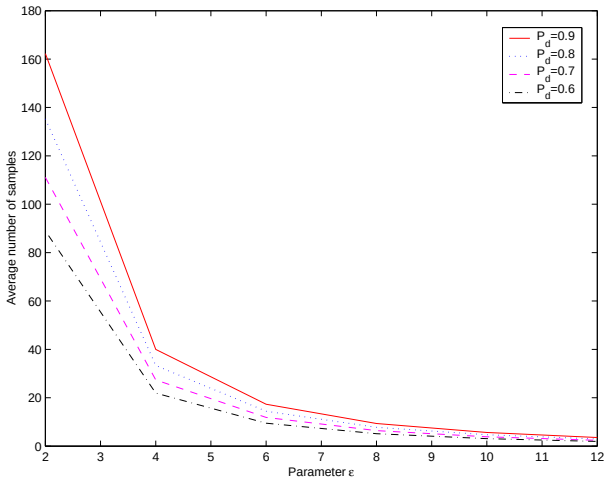

Figure 3: Average number of samples needed for misbehavior detection as a function of ϵ in absence of interference for the worst-case attack.

4.3 Inaccurate measurements due to interference

The underlying assumption of our approach was that the back-off value observations were collected in the absence of interference from ongoing concurrent transmissions. However, observations are affected by interference due to transmission of nodes that are located out of range of the attacker, but within range of an observer. For example, in figure 1, transmission of node C obstructs observations of 2. The presence of interference may corrupt some measurements and thus it is anticipated to increase the number of observation samples needed to derive a decision.

Since interference is caused due to ongoing data transmissions that are of much longer duration than that of an observed RTS or CTS packet, we can assume that the level of interference due to one such transmission remains constant for the duration of an observed RTS or CTS packet. Recall that RTS and CTS packets are sent with the lowest modulation level and coding rate. To enable analytical tractability, we consider an uncoded transmission and assume the use of BPSK (which is the lowest modulation level in 802.11a) in RTS/CTS transmission. The interference conditions during an RTS or CTS observed packet are captured by the signal-to-interference and noise ratio (SINR) γ. For fixed transmit power levels and certain variance of Gaussian noise at the receiver, this ratio depends on the relative proximity of the observer node to the transmitter of RTS or CTS message as well as to the interferers. The packet start point can be distinguished if the packet is received correctly. The bit error rate (BER) in the received RTS or CTS packet is given by $\text{BER} = Q(\sqrt{2\gamma})$ for BPSK modulation, where $Q(\cdot)$ denotes the Q-function. The probability of RTS or CTS packet error is the RTS-CTS packet error rate (PER) as

$$\text{PER} = 1 - (1 - \text{BER})^{8m} \tag{22}$$

where m is the number of bytes of the RTS and CTS packets and is 20 and 14 respectively. Since PER gives the percentage of observed packets received in error, the number of required observations to derive a decision is PER% higher than the corresponding number without interference. This PER value holds for uncoded transmission and thus it is an upper bound on PER when a coding scheme is used.

5. NUMERICAL RESULTS

The goal of the simulations is to assess the performance of our approach and identify the relative impact of different system parameters on it. The performance is measured in terms of the average required number of observation samples, $\mathbb{E}[N]$ in order to derive a decision, which essentially denotes the delay in detecting a misbehavior instance. In particular, we evaluate the performance with respect to the following parameters:

- Specified values of P_{FA} and P_M (or probability of detection, $P_D = 1 - P_M$).
- Perceived interference conditions, reflected in SINR γ.
- The tunable system parameter ϵ.

Parameter ϵ defines the class of attacks of interest since it specifies the incurred relative gain of the attacker in terms of the probability of channel access. In that sense, ϵ can be interpreted as a sensitivity parameter of the detection scheme with respect to attacks, which is determined according to the IDS requirements. IEEE 802.11 MAC is implemented and MATLAB is used to evaluate the performance of our scheme, taking into account the sequence of observed back-offs.

In order to obtain some intuition from the results, we consider the case of one attacker and the competing legitimate node (in this case $0 \leq \epsilon \leq 16$). First, we study the case with no interference from ongoing transmissions. In Fig. 2, we depict the average required number of observation samples as a function of probability of detection P_D for different values of P_{FA} and a fixed low value of ϵ, $\epsilon = 2$. This is the result of solving problem (10). It can be seen that for typical values of P_D and P_{FA} the average required number of samples for detecting the worst-case attack is 150 to 180. For

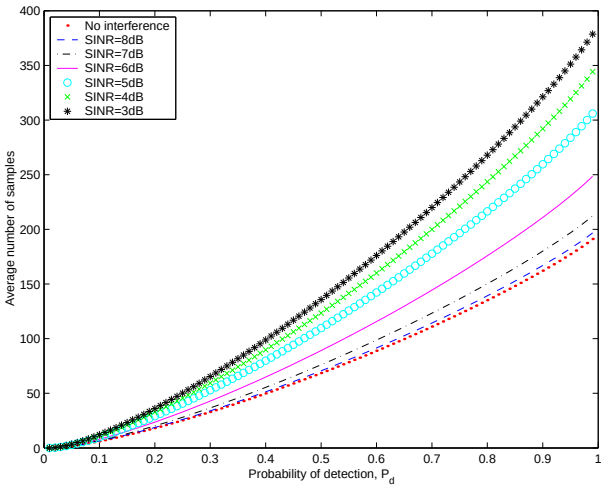

Figure 4: Average number of samples needed for misbehavior detection for P_{FA}=0.01 and different interference conditions. The observer node collects RTS packets.

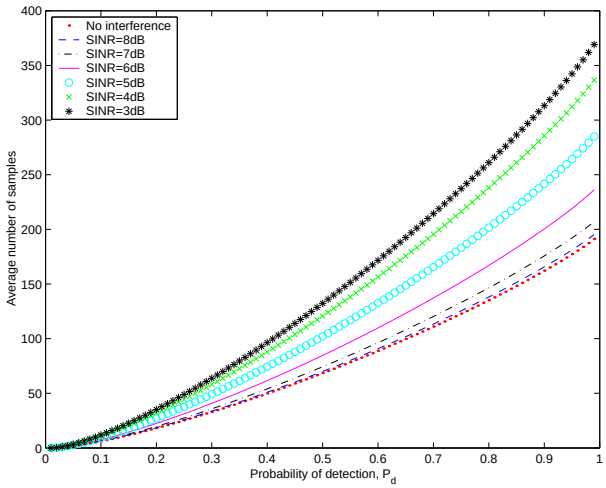

Figure 5: Average number of samples needed for misbehavior detection for P_{FA}=0.01 and different interference conditions. The observer node collects CTS packets.

more stringent specifications, namely larger P_D and lower P_{FA}, more samples are needed. In that sense, high values of P_D and P_{FA} may turn out to be beneficial for the attacker, since the latter intends increase the detection delay.

Fig. 3 illustrates the performance of the detection scheme as a function of the parameter ϵ for fixed $P_{FA} = 0.03$ and different values of P_D. The graph shows that low values of ϵ (e.g. up to 3) prolong the detection procedure, since in that case the attacker does not deviate significantly from the protocol. On the other hand, a large ϵ signifies a class of increasingly aggressive attacks for which the detection is achieved with very small delay. The results above provide useful insights about the response of the system with respect to the attack. For example, a value $\epsilon = 4$ denotes an attacker that obtains channel access 25% more times compared to the legitimate operation. In that case, 40 samples would suffice so as to detect the attack with $P_D = 90\%$. Therefore, a more aggressive attack policy (within the class $\mathcal{F}_\epsilon$ for large ϵ) brings significant benefits each time the attacker accesses the channel, but it allows limited number of channel uses before it is detected. On the other hand, a milder attack incurs lower benefit for each channel use but it enables the attacker to access the channel more times before it is detected.

Inspired by this observation, we can view DoS attacks as an extreme case of misbehavior. To illustrate this, we define the class of uniform pdf's $\{f_\delta(x) : f_\delta(x) = 1/\delta$ for $0 \leq x \leq \delta$ and 0 else$\}$. A DoS attack is described by pdf $f_\delta(x)$ with $\delta \ll W$ and $\delta \to 0$, since it chooses back-offs from a very small interval. By substituting $f_\delta(x)$ in (16), we obtain the expression for the expected number of samples needed for detection of DoS attacks as $C(\ln W - \ln \delta)^{-1}$.

We now proceed to quantifying the impact of interference on performance. Depending on interference conditions, a percentage of the back-off samples collected by the observer nodes are corrupted. In that case, the RTS or CTS PER indicates the amount of additional measurements required for reaching a decision, depending on whether the observer node resides within range of the attacker or the receiver of the attack. Fig. 4 shows the average required number of samples for achieving $P_{FA} = 0.01$ and P_D given by the x-axis for different intensity of interference, with $\epsilon = 2$. For large values of P_D it can be observed that intense interference conditions (reflected to SINR values of 3-4 dB) can increase the number of required samples by $85\% - 120\%$ compared to the case of no interference. In addition, for SINR$> 8dB$, the performance is not affected significantly by interference. Hence, interference can be viewed as providing additional benefit to the attacker in the sense that it prolongs detection. Similarly, for $P_{FA} = 0.04$, the corresponding percentage of additional samples is approximately $80 - 100\%$.

In figure 5, we show the case of obstructed observations based on received CTS packets for $\epsilon = 2$. It turns out that CTS observations are preferable to RTS ones in the sense that they result in fewer required number of samples to detect misbehavior. For example, for SINR values of 3-4 dB, $P_{FA} = 0.01$ and large P_D, we observe an increase of $85 - 100\%$ in the number of required samples compared to that with no interference. CTS observers become more efficient compared to RTS ones for more intense interference conditions. Therefore, when assigning observer roles to nodes, emphasis should be given to those nodes that are located within range of the receiver.

6. DISCUSSION

In this work, we presented a framework of study for the problem of MAC misbehavior detection. Our approach encompasses the case of an intelligent attacker that adapts its misbehavior strategy with the objective to remain undetected as long as possible. We cast the problem within a minimax robust detection framework, characterize the worst-case misbehavior strategy showing that the optimal detection rule is SPRT. Clearly, if the attacker is ignorant of the detection mechanism, the number of required observations to detect it under the same values of P_{FA} and P_D is lower than the corresponding value for the adaptive attacker. Our

results can thus shed light in the characterization of fundamental performance limits in terms of accuracy or detection delay for misbehavior detection. They can also serve as benchmarks for performance evaluation of other detection policies and can provide useful insights about the effect of interference on performance. Finally, we provided an instance of a case when cross-layer interaction offers a solution to the issue of notifying the network about the misbehavior.

Our work constitutes the first step towards building a theoretical framework for studying the structure of such misbehavior problems. The model can be extended to include obstruction of observations due to simultaneous channel access attempts. We now mention some issues for further study. A first issue concerns the exploitation of observations from several observers in order to improve performance. This amounts to the scenario where observers pass their measurements to a fusion center which then combines them appropriately and derives a decision as to the occurrence or not of attack. Due to different perceived channel conditions at different locations of observer nodes, the amount of interference at their receivers differs. If observers obtain the same sequence of measurements, different samples of the sequence are corrupted due to interference. The task of the fusion center is then simply to combine the received sequences of measurements in a fashion very similar to that of diversity combining. Given that there exists a certain cost (e.g. consumed energy) in passing measurements to a fusion center, an interesting issue pertains to the minimum number of observers that are necessary to achieve a certain level of performance in terms of detection delay or accuracy.

A far more challenging problem arises if each observer does not measure back-offs accurately but it obtains a sequence of distorted values. This situation may arise in case of occasional loss of synchronization between nodes or due to hardware (e.g. counter) malfunction. Another instance in which observers may have distorted back-off sequences is the following. At the $i > 1$ transmission, node A selects a back-off b and starts decrementing his counter. If the medium is sensed busy, the counter freezes (suppose for duration d) and restarts again when the medium is idle. When the counter reaches zero, the RTS message is sent. In that case, the observers perceive a back-off $\hat{b} = b + d$.

In our approach, we have assumed continuously backlogged nodes and have used channel access probability as a means of measuring the benefit of the attacker and corresponding performance loss of legitimate nodes. Implicitly, we assumed that fair sharing of the medium is reflected by this measure. However, fair sharing also involves the intention of a node to send a packet and therefore it is affected by packet arrivals from higher layers and backlogs at different nodes. This introduces the issue of throughput fairness and throughput benefit. The attacker causes more damage to the system if it prevents legitimate nodes from transmitting their payload.

The treatment of more than one attacker in the network is definitely worth investigating. It would be interesting to model and compare the case of attackers that act independently and that of attackers that co-operate. In the first case, the objectives of attackers may be conflicting in the sense that each of them attempts to maximize its own benefit. In the latter case, the optimal attack strategy, if it exists, can aid in quantifying the benefits of co-operation and its effects on performance degradation of legitimate nodes.

Finally, it would be very interesting to extend our approach and obtain results in the context of more sophisticated MAC protocols such as 802.11e with the special features regarding back-off control and differentiation in channel access opportunities that are incorporated in its enhanced DCF (EDCF) operation mode.

7. ACKNOWLEDGMENTS

The authors wish to thank George V. Moustakides from University of Thessaly, Greece for enlightening discussions. Finally, the constructive comments of anonymous reviewers are acknowledged.

8. REFERENCES

[1] J. Bellardo and S. Savage. 802.11 denial-of-service attacks: real vulnerabilities and practical solutions. In *Proc. of USENIX Security Symposium*, San Antonio, TX, June 2003.

[2] D. Bertsekas. *Convex analysis and optimization.* Athena Scientific, 2003.

[3] S. Buchegger and J.-Y. L. Boudec. Performance Analysis of the CONFIDANT Protocol. In *Proceedings of MobiHoc*, Lausanne, June 2002.

[4] M. Cagalj, S. Ganeriwal, I. Aad, and J.-P. Hubaux. On Cheating in CSMA/CA Ad Hoc Networks. Technical Report IC/2004/27, EPFL-DI-ICA, March 2004.

[5] A. A. Cardenas, S. Radosavac, and J. S. Baras. Detection and prevention of MAC layer misbehavior in ad hoc networks. In *Proceedings of SASN '04*, pages 17–22, 2004.

[6] V. Dragalin, A. Tartakovsky, and V. Veeravalli. Multihypothesis Sequential Probability Ratio Tests - Part I: Asymptotic optimality. *IEEE Trans. on Information Theory*, 45(7):2448 – 2461, Nov. 1999.

[7] V. Gupta, S. Krishnamurthy, and M. Faloutsos. Denial of service attacks at the MAC layer in wireless ad hoc networks. In *Proc. of MILCOM*, 2002.

[8] C. W. Helstrom. *Elements of signal detection and estimation.* Prentice-Hall, 1995.

[9] IEEE. IEEE wireless LAN medium access control (MAC) and physical layer (PHY) specifications, 1999.

[10] S. Kassam and H. Poor. Robust techniques for signal processing : a survey. *Proc. of the IEEE*, 73(3):433–481, March 1985.

[11] P. Kyasanur and N. Vaidya. Detection and handling of MAC layer misbehavior in wireless networks. In *Proc. of International Conference on Dependable Systems and Networks*, 2003.

[12] S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Mobile Computing and Networking*, pages 255–265, 2000.

[13] M. Raya, J.-P. Hubaux, and I. Aad. DOMINO: A system to detect greedy behavior in IEEE 802.11 Hotspots. In *Proceedings of MobiSys '04*, pages 84–97, 2004.

[14] S. Verdu and H.V.Poor. On minimax robustness: a general approach and applications. *IEEE Trans. on Information Theory*, 30(2):328–340, March 1984.

[15] A. Wald. *Sequential Analysis.* John Wiley and Sons, New York, 1947.