

An analytic framework for modeling and detecting access layer misbehavior in wireless networks

Svetlana Radosavac¹

University of Maryland College Park

and

George Moustakides

University of Patras, Greece

and

John S. Baras

University of Maryland College Park

and

Iordanis Koutsopoulos

University of Thessaly, Greece

The widespread deployment of wireless networks and hot spots that employ the IEEE 802.11 technology has forced network designers to put emphasis on the importance of ensuring efficient and fair use of network resources. In this work we propose a novel framework for detection of intelligent adaptive adversaries in the IEEE 802.11 MAC by addressing the problem of detection of the worst-case scenario attacks. Utilizing the nature of this protocol we employ sequential detection methods for detecting greedy behavior and illustrate their performance for detection of least favorable attacks. By using robust statistics in our problem formulation, we attempt to utilize the precision given by parametric tests, while avoiding the specification of the adversarial distribution. This approach establishes the lowest performance bound of a given Intrusion Detection System (IDS) in terms of detection delay and is applicable in online detection systems where users who pay for their services want to obtain the information about the best and the worst case scenarios and performance bounds of the system. This framework is meaningful for studying misbehavior due to the fact that it does not focus on a specific adversarial strategies and therefore is applicable to a wide class of adversarial strategies.

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

General Terms: Design, Security

Additional Key Words and Phrases: Wireless networks, MAC layer, min-max robust detection, protocol misbehavior

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

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2008 ACM 0000-0000/2008/0000-0001 \$5.00

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 [Raya et al. 2004] 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 [Kysanur and Vaidya 2003], 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 pre-specified 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

¹Now with DoCoMo Communications Laboratories USA, Inc.

the sender and the receiver. The work in [Cárdenas et al. 2004] 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 [Čagalj et al. 2005]. 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 [Marti et al. 2000] 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 [Buechegger and Boudec 2002] 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 throughout 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 [Gupta et al. 2002], 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 [Bellardo and Savage 2003] 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 [Wald 1947] 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 min-max

formulation and identify the rule that optimizes worst-case performance over the class of allowed uncertainty conditions. The min-max 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 [Kassam and Poor 1985; Verdu and H.V.Poor 1984].

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 min-max 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 min-max robust detection problem. Our work contributes to the current literature by: (i) formulating the misbehavior problem at hand as a min-max 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. Although we do mention the impact of interference on the performance of the IDS and perform initial evaluation to illustrate its impact on detection delay, we do not perform extensive analysis and performance evaluation of the system in the presence of interference in the remainder of this article.

Compared to our preliminary work in [Radosavac et al. 2005], in this work we introduce a more sophisticated adversary model that captures the behavior of an adversary in the IEEE 802.11 MAC in a more precise manner. Namely, the work in [Radosavac et al. 2005] assumed that the back-off counters never freeze due to a perceived busy channel. However, according to the IEEE 802.11 MAC specification, each node freezes its back-off counter when the channel is busy. Consequently, the above assumption lead to a heavily approximated adversarial model. In this work, we improved the adversary model from [Radosavac et al. 2005] by assuming that the back-off counters freeze due to a perceived busy channel. More specifically, we used asymptotic theory to derive an expression of the attacker percentage of channel access and to define the attack classes of interest. We also extend this framework and generalize our treatment for multiple competing nodes and assess the performance of both the detection scheme and the adversary for such scenario. In particular, we incorporated terminology and derivations that apply to quickest change detection theory. The improved and precise version for the worst-case adversarial strategy

enabled us to provide rigorous proofs for the structure of the optimal attack and for the properties of the saddle point of the IDS-attacker game. In order to provide a more sophisticated and realistic simulation scenario, we implemented a misbehaving node model in the network simulator Opnet and evaluated the performance of both the adversary and the IDS in terms of trade-off between the expected time to False Alarm and the expected time to detection.

The rest of the paper is organized as follows. In Section 2.1, we discuss the issue of misbehavior in IEEE 802.11 MAC protocol. In Section 3 we introduce the min-max robust detection model with the underlying assumptions; and present our main results regarding misbehavior detection. Section 5 contains a number of numerical results; further issues are discussed in Section 4 and we conclude our study in Section 6.

In subsequent discussion, the terms “misbehavior” and “attack”, as well as, “misbehaving node” and “attacker” will be used interchangeably.

2. IEEE 802.11 MAC DCF: OVERVIEW OF THE PROTOCOL

The most frequently used MAC protocol for wireless networks is the IEEE 802.11 MAC protocol, which uses a distributed contention resolution mechanism for sharing the wireless channel. Its design attempts to ensure a relatively fair access to the medium for all participants of the protocol. In order to avoid collisions, the nodes follow a binary exponential back-off scheme that favors the last winner amongst the contending nodes.

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) [IEEE 1999]. A node with a packet to transmit selects a random back-off value b uniformly from the set $\{0, 1, \dots, W - 1\}$, where W is the (fixed) size of the contention window. The random back-off selected corresponds to the number of slots a station needs to wait in addition to the mandatory Distributed Interframe Space (DIFS) interval before attempting to transmit. 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 Short Interframe Space (SIFS) intervals, data packets and acknowledgment (ACK) 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 .

Fig. 1 illustrates the scenario of contending nodes using the protocol. In this




Fig. 1. Nodes A and C contend for accessing node B. In the first attempt A reserves the channel followed by successful access by node C.

specific scenario nodes A and C contend for accessing node B. In the first attempt, after waiting for the fixed interval DIFS, node A senses the channel idle and sends an RTS to node B. Consequently, node C overhears the RTS, which also contains the duration of the intended data exchange between A and B (d), and defers its transmission for the time interval equal to d (i.e. sets its NAV to d). After waiting for SIFS, node B senses the channel idle and responds with a CTS. After successfully exchanging the data, all participating stations wait for a fixed interval equal to DIFS followed by a back-off interval b , uniformly chosen within the interval $[0, W]$. In this scenario, node C chooses smaller back-off value and accesses the channel, forcing node A to defer its transmission. Typical parameter values for the MAC protocol depend on the physical layer that IEEE 802.11 uses. Table I shows the parameters used when the physical layer is using direct sequence spread spectrum (DSSS).

DIFS	$50\mu s$
SIFS	$10\mu s$
SlotTime	$20\mu s$
ACK	$112\text{bits} + \text{PHY_header} = 203\mu s$
RTS	$160\text{bits} + \text{PHY_header} = 207\mu s$
CTS	$112\text{bits} + \text{PHY_header} = 203\mu s$
DATA	$\text{MAC_header (30b)} + \text{DATA(0-2312b)} + \text{FCS(4b)}$
Timeouts	$300-350\mu s$
CW_{min}	32 time slots
CW_{max}	1024 time slots

Table I. Parameters for DSSS

2.1 Misbehavior in the IEEE 802.11 MAC protocol

The scenario provided in Fig. 1 illustrated the the IEEE 802.11 DCF favors the node that selects the smallest back-off value among a set of contending nodes. This opens space for misbehavior of protocol participants if no detection system is employed. More specifically, 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. 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




Fig. 2. Observer nodes and effect of interference due to concurrent transmissions.

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 Distributed Interframe Space (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. 2 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

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

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

$$t'_i = T_{\text{ACK},i-1} + T_{\text{DIFS}} + b_i + T_{\text{RTS}} + T_{\text{SIFS}}, \quad i > 1. \quad (2)$$

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. MIN-MAX ROBUST MISBEHAVIOR DETECTION

In this section we present our approach for misbehavior detection when observations are not obstructed by interference. In Sect. 4, we 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 Fig. 2. Our work assumes a stationary network where the node relative positions do not change with time. 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.

A sequential detection test is therefore a procedure which with every new information that arrives asks the question whether it should *stop* receiving more samples or continue sampling. If the answer to the first question is to stop (because sufficient information has been accumulated) then it proceeds to the phase of making a *decision* on the nature of the data. It is therefore clear that there are two quantities involved: a stopping time N which is a random variable taking positive integer values and denoting the time we decide to stop getting more data;

and a decision rule d_N which at the time of stopping N decides between the two hypotheses $\mathbf{H}_0, \mathbf{H}_1$, where \mathbf{H}_0 denotes legitimate behavior and \mathbf{H}_1 denotes selfish behavior, and therefore assumes the values 0,1. For simplicity let us denote with \mathcal{D} the combination $\mathcal{D} = (N, d_N)$ of the stopping time N and the decision rule d_N .

The probability of false alarm and the probability of missed detection constitute inherent tradeoffs in a detection scheme. Clearly we can obtain small values for both of these two decision error probabilities by accumulating more information, that is, at the expense of larger detection delay. A logical compromise would therefore be to prescribe some maximal allowable values for the two error probabilities, and attempt to *minimize* the expected detection delay. Expressing this problem under a more formal setting, we are interested in finding a sequential test $\mathcal{D} = (N, d_N)$ that solves the following constraint optimization problem

$$\inf_{N, d_N} \mathbb{E}_1[N], \quad \text{under the constraints } \mathbb{P}_0[d_N = 1] \leq \alpha; \mathbb{P}_1[d_N = 0] \leq \beta; \quad (3)$$

where $\mathbb{P}_i, \mathbb{E}_i$ denote probability and expectation under hypothesis $\mathbf{H}_i, i = 0, 1$, and $0 < \alpha, \beta < 1$ are the prescribed values for the probability of false alarm and miss respectively.

This interesting mathematical setup was first proposed by [Wald 1947] where he also introduced the Sequential Probability Ratio Test (SPRT) for its solution. The SPRT test is defined in terms of the log-likelihood ratio S_n

$$S_n = \ln \frac{f_1(x_1, \dots, x_n)}{f_0(x_1, \dots, x_n)}, \quad (4)$$

of the two joint probability density functions $f_i(x_1, \dots, x_n)$ of the data $\{x_1, \dots, x_n\}$ under hypothesis $\mathbf{H}_i, i = 0, 1$. The corresponding stopping time N and decision rule d_N are then given by

$$N = \inf_n \{n : S_n \notin [A, B]\} \quad (5)$$

$$d_N = \begin{cases} 1 & \text{if } S_N \geq B \\ 0 & \text{if } S_N \leq A, \end{cases} \quad (6)$$

where $A < 0 < B$ thresholds selected so as SPRT satisfies the two decision error probability constraints with equality. We can see that the SPRT test continues sampling as long as the log-likelihood ratio takes values within the interval (A, B) and stops taking more samples the first time it exceeds it. Once stopped, the decision function d_N decides in favor of hypothesis \mathbf{H}_1 when S_N exceeds the largest threshold and in favor of \mathbf{H}_0 when S_N is below the smallest threshold. If in particular the data are independent and identically distributed (i.i.d.) under both hypotheses then the log-likelihood ratio S_n takes the following simple form

$$S_n = \sum_{k=1}^n \ln \frac{f_1(x_k)}{f_0(x_k)} = S_{n-1} + \ln \frac{f_1(x_n)}{f_0(x_n)}, \quad S_0 = 0. \quad (7)$$

Here $f_i(x)$ is the common probability density function (pdf) of the samples under hypothesis $\mathbf{H}_i, i = 0, 1$. Notice that the recurrent relation in the right hand side of (7) allows for an efficient computation of the statistics S_n which requires only constant number of operations per time step and finite memory (we only need to store S_n as opposed to the whole sequence $\{x_n, \dots, x_1\}$).

Optimality of SPRT in the sense described in (3) is assured *only* when the data are i.i.d. under both hypotheses [Wald and Wolfowitz 1948]. For other data models there exists a very rich literature referring to asymptotic optimality results (see for example [Dragalin et al. 1999]). Concluding, we should also mention that the actual optimality of SPRT is significantly stronger than the one mentioned in (3). The SPRT not only minimizes the average delay under \mathbf{H}_1 but also *simultaneously* minimizes the alternative average delay $\mathbb{E}_0[N]$. This double optimality property is rather remarkable and not encountered in any other detection scheme.

It is clear from the previous discussion that our intention is to follow a sequential approach for the detection of attacks. Notice however that in order to be able to use the SPRT test it is necessary to specify both probability density functions $f_i(x)$, $i = 0, 1$ under the two hypotheses. Although the pdf $f_0(x)$ of a legitimate node is known, this is not the case for an attacker. Furthermore, specifying a candidate density $f_1(x)$ for an attacker without some proper analysis may result in serious performance degradation if the attacker's strategy diverges from our selection.

In order to be able to propose a specific detection rule we need to clarify and mathematically formulate the notion of an "attack". We should however place our main emphasis to attacks that incur large gains for the attacker (result in higher chances of channel access). An attack will then have 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 Min-max 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 min-max 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 needed to derive a decision.

A basic notion in min-max approaches is that of a *saddle point*. A strategy (detection rule) $\mathcal{D}^* = (N^*, d_N^*)$ and an operating point (attack) f_1^* in the uncertainty

class form a saddle point if:

- (1) For the attack f_1^* , any detection rule \mathcal{D} other than \mathcal{D}^* has worse performance. Namely \mathcal{D}^* is the optimal detection rule for attack f_1^* in terms of minimum (average) number of required observations.
- (2) For the detection rule \mathcal{D}^* , any attack f_1 from the uncertainty class, other than f_1^* gives better performance. Namely, detection rule \mathcal{D}^* has its worst performance for attack f_1^* .

Implicit in the min-max 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 (attainable) 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 i th 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 $i = 1, \dots, I_{\max}$ 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, for mathematical simplicity, we consider continuous distributions to represent attacks. The analysis for the discrete value case is very similar and is therefore omitted. 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.

Let us first compute the probability P_1 of the attacker to access the channel as a function of the pdfs f_1 and f_0 . Following the IEEE 802.11 protocol, the back-off counter of any node freezes during the transmissions and reactivates during free periods. Therefore let us observe the back-off times during a fixed period T that *does not include* transmission intervals. Consider first the case of one misbehaving and one legitimate node and assume that within the time period T , we observe X_1, \dots, X_N , N samples of the attacker's back-off and Y_1, \dots, Y_M , M samples of the legitimate node's back-offs. It is then clear that the attacker's percentage of accessing the channel during the period T is $N/(N + M)$. In order to obtain the desired probability we simply need to compute the limit of this ratio as $T \rightarrow \infty$.

Notice that

$$\begin{aligned} X_1 + \cdots + X_N &\leq T < X_1 + \cdots + X_{N+1} \\ Y_1 + \cdots + Y_M &\leq T < Y_1 + \cdots + Y_{M+1}, \end{aligned}$$

which yields

$$\frac{\frac{N}{X_1 + \cdots + X_N}}{\frac{N}{N+1} \frac{N+1}{X_1 + \cdots + X_{N+1}} + \frac{M}{M+1} \frac{M+1}{Y_1 + \cdots + Y_{M+1}}} \geq \frac{\frac{N}{T}}{\frac{N}{T} + \frac{M}{T}} \geq \frac{\frac{N}{N+1} \frac{N+1}{X_1 + \cdots + X_{N+1}}}{\frac{N}{X_1 + \cdots + X_N} + \frac{M}{Y_1 + \cdots + Y_M}}. \quad (8)$$

Letting $T \rightarrow \infty$ results in $N, M \rightarrow \infty$ and from the previous double inequality, by applying the Law of Large Numbers, we conclude that

$$P_1 = \lim_{N, M \rightarrow \infty} \frac{N}{N+M} = \frac{\frac{1}{\mathbb{E}_1[X]}}{\frac{1}{\mathbb{E}_1[X]} + \frac{1}{\mathbb{E}_0[Y]}}. \quad (9)$$

Using exactly similar reasoning the probability P_1 , for the case of one misbehaving node against n legitimate ones, takes the form

$$P_1 = \frac{\frac{1}{\mathbb{E}_1[X]}}{\frac{1}{\mathbb{E}_1[X]} + \frac{n}{\mathbb{E}_0[Y]}} = \frac{1}{1 + n \frac{\mathbb{E}_1[X]}{\mathbb{E}_0[Y]}} = \frac{1}{1 + n \frac{2\mathbb{E}_1[X]}{W}}, \quad (10)$$

where in the last equality we have used the fact that the average back-off of a legitimate node is $W/2$ (because f_0 is uniform in $[0, W]$).

If the attacker were legitimate then $\mathbb{E}_1[X] = \mathbb{E}_0[Y]$ and his probability of accessing the channel, from Eq. (10), would have been $1/(n+1)$. It is therefore clear that whenever

$$\mathbb{E}_1[X] = \epsilon \mathbb{E}_0[Y], \quad \text{with } \epsilon \in (0, 1) \quad (11)$$

the attacker enjoys a gain as compared to any legitimate node since then

$$P_1 = \eta \frac{1}{n+1} > \frac{1}{n+1}, \quad \text{where } \eta = \frac{1+n}{1+\epsilon n} \in (1, n+1). \quad (12)$$

In other words his probability of accessing the channel is greater than the corresponding probability of any legitimate node by a factor $\eta > 1$.

Using the simple modeling introduced in the previous paragraph we are now able to quantify the notion of an ‘‘attack’’. Let η be a quantity that satisfies $1 < \eta < n+1$ and consider the class \mathcal{F}_η of all pdf’s that induce a probability P_1 of accessing the channel that is no less than $\eta/(n+1)$. Using (11) and (12) the class \mathcal{F}_η can be explicitly defined as

$$\mathcal{F}_\eta = \left\{ f_1(x) : \int_0^W x f_1(x) dx \leq \frac{1 - \frac{\eta}{n+1}}{n \frac{\eta}{n+1}} \frac{W}{2} \right\}, \quad 1 < \eta < n+1. \quad (13)$$

This class includes all possible attacks for which the incurred relative gain exceeds the legitimate one by $(\eta - 1) \times 100\%$. The class \mathcal{F}_η is the uncertainty class of the robust approach and η is a tunable parameter. Notice from (12) that since P_1 is a probability the *gain factor* η must not exceed $n+1$ in order for the previous inequality to produce a nonempty class \mathcal{F}_η .

By defining the class \mathcal{F}_η , 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. We define the severity of the attack by changing the gain factor η . Values of η larger but close to 1 are equivalent to low-impact attacks whereas values significantly larger than 1 are equivalent to DoS attacks.

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

Hypothesis \mathbf{H}_0 concerns legitimate operation and thus the corresponding pdf $f_0(x)$, as was mentioned before, is the uniform one. Hypothesis \mathbf{H}_1 corresponds to misbehavior with unknown pdf $f_1(x) \in \mathcal{F}_\eta$.

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 $\mathbb{E}_1[N]$ needed until a decision is reached, where the average is taken with respect to the distribution $f_1(x)$ employed by the attacker. This expectation is clearly a function of the adopted detection rule \mathcal{D} and the pdf $f_1(x)$, that is,

$$\mathbb{E}_1[N] = \phi(\mathcal{D}, f_1). \quad (14)$$

Let $\mathcal{T}_{\alpha,\beta}$ denote the class of all sequential tests for which the false alarm and missed detection probabilities do not exceed some specified levels α and β respectively. Consider also the class \mathcal{F}_η of densities $f_1(x)$ as in (13) for some prescribed gain factor $\eta > 1$. In the context of the min-max robust detection framework, the goal is to optimize performance in the presence of worst-case attack, that is, solve the following min-max problem

$$\inf_{\mathcal{D} \in \mathcal{T}_{\alpha,\beta}} \sup_{f_1 \in \mathcal{F}_\eta} \phi(\mathcal{D}, f_1). \quad (15)$$

Solving a min-max problem is usually complicated, unless one can obtain a *saddle point* solution.

DEFINITION 3.1. A pair (\mathcal{D}^*, f_1^*) is called a *saddle point* of the function ϕ if

$$\phi(\mathcal{D}^*, f_1) \leq \phi(\mathcal{D}^*, f_1^*) \leq \phi(\mathcal{D}, f_1^*); \quad \forall \mathcal{D} \in \mathcal{T}_{\alpha,\beta}, \quad \forall f_1 \in \mathcal{F}_\eta. \quad (16)$$

As we can see a saddle point (\mathcal{D}^*, f_1^*) of ϕ consists of a detection scheme \mathcal{D}^* and an attack distribution f_1^* . Equation (16) is a formal statement of properties 1 and 2 that were mentioned in Section 3.2. The property that is important here in connection to the min-max problem (15) is the fact that the saddle point pair (\mathcal{D}^*, f_1^*) also solves the min-max problem, since one can prove that [Bertsekas 2003]

$$\inf_{\mathcal{D} \in \mathcal{T}_{\alpha,\beta}} \sup_{f_1 \in \mathcal{F}_\eta} \phi(\mathcal{D}, f_1) \geq \sup_{f_1 \in \mathcal{F}_\eta} \phi(\mathcal{D}^*, f_1) = \phi(\mathcal{D}^*, f_1^*). \quad (17)$$

Saddle point solutions are much easier to obtain than their min-max counterparts. Unfortunately saddle point solutions do not always exist. In view of (17), instead of solving (15) it is sufficient to find the saddle point that solves (16). The saddle point pair (\mathcal{D}^*, f_1^*) is specified in the next theorem.

THEOREM 3.2. *Let the gain factor $\eta \in (1, n + 1)$ and the maximal allowable decision error probabilities α, β be given. Then the pair (\mathcal{D}^*, f_1^*) which **asymptotically** (for small values of α, β) solves the saddle point problem defined in (16) is the following*

$$f_1^*(x) = \frac{\mu}{W} \frac{e^{\mu(1-\frac{x}{W})}}{e^\mu - 1}, \quad (18)$$

where $\mu > 0$ is the solution to the following equation in μ

$$2 \left(\frac{1}{\mu} - \frac{1}{e^\mu - 1} \right) = \frac{1 - \frac{\eta}{n+1}}{n \frac{\eta}{n+1}}. \quad (19)$$

The corresponding decision rule $\mathcal{D}^* = (N^*, d_{N^*})$ is the SPRT test that discriminates between $f_1^*(x)$ and $f_0(x)$ (the uniform density) and is given by

$$\begin{aligned} S_n^* &= S_{n-1}^* + \ln \frac{f_1^*(x_n)}{f_0(x_n)} \\ &= S_{n-1}^* + \mu \left(1 - \frac{x_n}{W} \right) + \ln \left(\frac{\mu}{e^\mu - 1} \right); \quad S_0^* = 0. \end{aligned} \quad (20)$$

$$N^* = \inf_n \{n : S_n^* \notin [A, B]\} \quad (21)$$

$$d_{N^*} = \begin{cases} 1 & \text{if } S_{N^*}^* \geq B \\ 0 & \text{if } S_{N^*}^* \leq A. \end{cases} \quad (22)$$

PROOF. We first note that (19) is equivalent to

$$\int_0^W x f_1^*(x) dx = \frac{1 - \frac{\eta}{n+1}}{n \frac{\eta}{n+1}} \frac{W}{2} \quad (23)$$

which assures that $f_1^*(x)$ defined in (18) is a member of the uncertainty class \mathcal{F}_η . Let us now demonstrate that for any gain factor $\eta \in (1, n + 1)$ there always exists $\mu \in (0, \infty)$ so that (19) is true. Notice that for $\eta \in (1, n + 1)$ we have that $1/(n + 1) < \eta/(n + 1) < 1$. If we now call $\Phi(\mu) = 2 \left(\frac{1}{\mu} - \frac{1}{e^\mu - 1} \right)$ then $\Phi(\mu)$ is a continuous function of μ . Furthermore we observe that $\Phi(0) = 1 > \eta/(n + 1)$; while one can show that $\lim_{\mu \rightarrow \infty} \Phi(\mu) = 0 < \eta/(n + 1)$. Since we can find two values of μ one yielding a smaller and another a larger value than $\eta/(n + 1)$, due to continuity, we can argue that there exists $\mu > 0$ such that the equality in (19) is assured. In fact this μ is unique since it is also possible to show that $\Phi(\mu)$ is strictly decreasing.

Let us now proceed to the saddle point problem (16). We observe that the right hand side inequality suggests that \mathcal{D}^* must be the optimum detection structure for $f_1^*(x)$. Indeed this is how \mathcal{D}^* is defined, since it is selected as the SPRT test that optimally discriminates between $f_1^*(x)$ and the uniform $f_0(x)$.

In order to show that the left hand side is also true, we adopt an *asymptotic* approach. By considering that the two maximal error probabilities α, β are small, it is possible to use efficient approximations for the two thresholds A, B and the average detection delay function $\phi(\mathcal{D}^*, f_1)$. Specifically from [Wald 1947] we have that A, B can be approximated as

$$A = \ln \frac{\beta}{1 - \alpha}, \quad B = \ln \frac{1 - \beta}{\alpha}, \quad (24)$$

and the expected delay by the expression

$$\phi(\mathcal{D}^*, f_1) = \frac{A\beta + B(1 - \beta)}{\int_0^W \ln \frac{f_1^*(x)}{f_0(x)} f_1(x) dx}. \quad (25)$$

In fact these formulas become exact if the SPRT statistics S_n^* hits exactly (does not overshoot) the two thresholds A, B at the time of stopping. This for example happens in continuous-time and continuous-path processes.

Since the numerator in the previous formula is constant, the left hand side inequality in (16) is true if the denominator in (25) is minimized for $f_1(x) = f_1^*(x)$. Because we consider $f_1(x) \in \mathcal{F}_\eta$, inequality (13) applies, therefore we can write

$$\begin{aligned} \int_0^W \ln \frac{f_1^*(x)}{f_0(x)} f_1(x) dx &= \mu \int_0^W \left(1 - \frac{x}{W}\right) f_1(x) dx + \ln \left(\frac{\mu}{e^\mu - 1}\right) \\ &\geq \mu \left(1 - \frac{1 + n - \eta}{2n\eta}\right) + \ln \left(\frac{\mu}{e^\mu - 1}\right) \\ &= \mu \int_0^W \left(1 - \frac{x}{W}\right) f_1^*(x) dx + \ln \left(\frac{\mu}{e^\mu - 1}\right) \\ &= \int_0^W \ln \frac{f_1^*(x)}{f_0(x)} f_1^*(x) dx, \end{aligned} \quad (26)$$

where for the first inequality we used (13) and for the last two equalities we used (18),(23). This concludes the proof. \square

Regarding Theorem 3.2 we would like to point out that our selection of $f_1^*(x)$ was in fact the outcome of a rigorous analysis. We basically used the additional property enjoyed by the saddle point solution to solve not only the min-max problem in (15) but also its max-min version

$$\sup_{f_1 \in \mathcal{F}_\eta} \inf_{\mathcal{D} \in \mathcal{T}_{\alpha, \beta}} \phi(\mathcal{D}, f_1). \quad (27)$$

It turns out that this latter problem can be solved directly (using standard variational techniques), thus providing us with a suitable candidate pdf $f_1^*(x)$ for the saddle point problem (17). Of course we then need to go through the preceding proof in order to establish that $f_1^*(x)$ is indeed the correct pdf.

As was mentioned above, the min-max 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 Section 2, 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 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 Fig. 2, 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 back-off by choosing small back-off 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 observers 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.

4.2 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 Fig. 2, 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

$$\text{PER} = 1 - (1 - \text{BER})^{8m} \quad (28)$$

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 EXAMPLES

The goal of our examples 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 addition to that, we investigate the influence of the number of regular participants on the form of the least favorable distribution $f_1^*(x)$. 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 η .

Parameter η 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, η 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 Fig.3 we present the form of the least favorable attack pdf $f_1^*(x)$ as a function of the gain factor η and the number of legitimate nodes n .

Fig.3a depicts the form of the density for $n = 2$ legitimate nodes competing with one attacker for values of the gain factor $\eta = 1, 1.5, 2, 2.5$. We observe that as $\eta \rightarrow 3$ (the maximum possible gain for $n = 2$) the density tends to a Dirac delta function at $x = 0$ which corresponds to DoS attack, representing the extreme case of misbehavior where the attacker consumes all the available resources.

In Fig. 3b we fix the gain factor to $\eta = 1.5$ (the attacker enjoys 50% more access to the channel than a legitimate node) and plot $f_1^*(x)$ for number of legitimate nodes $n = 1, 2, 5, \infty$. We observe that as the number n of legitimate nodes increases, the attacker converges towards a less aggressive strategy. The interesting point is that the least favorable pdf converges very quickly to a limiting function as the number of legitimate nodes increases. This example confirms that it is possible to detect an attacker even if there is a large number of legitimate nodes present, since the attacker in order to maintain his relative gain must use a pdf which differs from the nominal uniform.

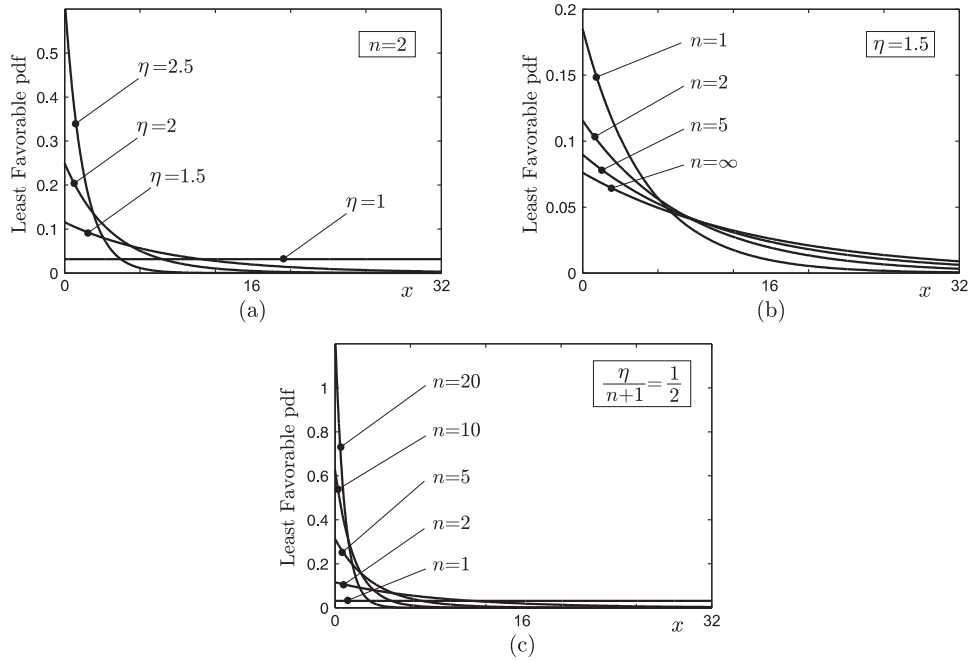


Fig. 3. Form of least favorable pdf $f_1^*(x)$: a) number of legitimate nodes $n = 2$, 1 malicious node and gain factor $\eta = 1, 1.5, 2, 2.5$; b) gain factor $\eta = 1.5$ and number of legitimate nodes $n = 1, 2, 5, \infty$; c) absolute gain $\frac{\eta}{n+1} = \frac{1}{2}$ and number of legitimate nodes $n = 1, 2, 5, 10, 20$.

Instead of fixing the attacker’s gain relatively to the gain of a legitimate node, let us examine what happens when the attacker follows a more aggressive policy and demands channel access for a *constant* percentage of time, regardless of the number of existing nodes. To achieve this goal, the gain factor η must be selected so that $\eta \frac{1}{n+1}$ is a constant. Fig. 3c depicts this specific scenario for $\frac{\eta}{n+1} = \frac{1}{2}$. In other words, the attacker has access to the channel 50% of the time, regardless of the number of competing nodes. We can see that when $n = 1$ the attacker behaves legitimately, but as the number n of legitimate nodes increases, the attacker quickly resorts to the strategies that are of DoS type in order to maintain this fixed access percentage. This is evident from the fact that the least favorable pdf tends to accumulate all its probability mass at small back-off values.

In order to obtain some intuition from our results, we consider the case of one attacker competing with $n \geq 1$ legitimate nodes. In Fig. 4a we depict the average required number of observation samples as a function of the parameter η . We fix the probability of detection and the probability of false alarm to 0.99 and 0.01 respectively and measure the Average Detection Delay $\mathbb{E}[N]$ for $1 < \eta < n+1$. The graph shows that low values of η prolong the detection procedure, since in that case the attacker does not deviate significantly from the protocol. On the other hand, a large η signifies a class of increasingly aggressive attacks for which the detection is achieved with very small delay. Due to the fact that the value of η is limited with the number of legitimate nodes, we cannot compare the performance of the

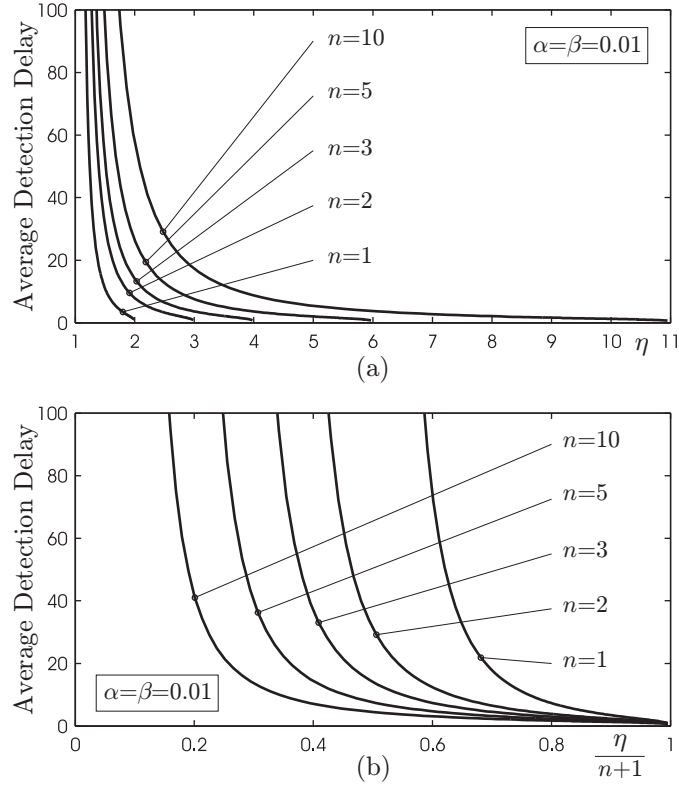


Fig. 4. Average Detection Delay $\mathbb{E}[N]$ as a function of (a) gain factor η ; (b) absolute gain $\frac{\eta}{n+1}$ for $\alpha = \beta = 0.01$

system for different values of n unless the absolute gain $\frac{\eta}{n+1}$ is used. In Fig. 4b we depict $\mathbb{E}[N]$ as a function of the absolute gain. It can be seen that detection becomes more efficient as the number of participating legitimate nodes increases. For example, for an absolute gain of 0.6, the IDS will require 10 times less samples to detect misbehavior for $n = 5$, than for the case of $n = 1$.

The results above provide useful insights about the response of the system with respect to the attack. A more aggressive attack policy 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. If the policy of a fixed gain is followed, the attackers behavior converges towards the DoS attack as n increases. The solution to this problem from the attackers point of view is offered in Section 6.

We now proceed to quantify 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




Fig. 5. Average detection delay for different values of SNR and $n=1, 3, 10$

the attack. Fig. 5 shows the average required number of samples needed for detection of an optimal attacker for different intensity of interference, with respect to the absolute gain $\eta \frac{1}{n+1}$. System performance is evaluated for $n = 1, 3$ and 10 . For large values of P_d it can be observed that intense interference conditions (reflected in the SINR values of 3-4 dB) can increase the number of required samples by 85% – 120% compared to the case of no interference. It is also worth mentioning that as the aggressiveness of an attacker increases, the number of samples needed for detection significantly decreases, regardless of the SINR values. More specifically, for $\text{SINR} > 8\text{dB}$, 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. Due to different lengths of RTS and CTS messages, the number of samples needed to detect misbehavior is lower when CTS messages are used in measurements. For example, for SINR values of 3-4 dB, $\alpha = \beta = 0.01$, we observe an increase of 85 – 100% in the number of required samples compared to that with no interference. Therefore, when assigning observer roles to nodes, emphasis should be given to those nodes that are located within range of the receiver. The amount of additional measurements needed for detection expressed in the form of PER for different values of SINR is presented in Fig. 6.

Finally, we implement the worst-case attack pdf characterized by Eq. 18 in the network simulator Opnet. We take advantage of the experimental setup and perform evaluation as a tradeoff between the average time to detection, T_d , and the average time to false alarm, T_{fa} , a quantity that is more meaningful and intuitive in practice. It is important to emphasize that the realistic false alarm rate used by actual intrusion detection systems is much lower than $\alpha = 0.01$ used in the theoretical analysis. We claim that this false alarm rate leads to an accurate estimate of the false alarm rates that need to be satisfied in actual anomaly detection systems [Cárdenas et al. 2006; Axelsson 1999]. Due to that fact we set $\beta = 0.01$ and vary α from 10^{-2} up to 10^{-10} (where $\alpha = 10^{-10}$ corresponds to one false alarm during the whole simulation period). The back-off distribution of an optimal attacker was implemented in the network simulator Opnet and tests were performed for various levels of false alarms. The backlogged environment in Opnet was created by employing a relatively high packet arrival rate per unit of time: the results were




Fig. 6. PER[%] as a function of SINR for RTS and CTS messages

collected for the exponential(0.01) packet arrival rate and the packet size was 2048 bytes. The results for both legitimate and malicious behavior were collected over a fixed period of 1.5min. We note that the simulations were performed with nodes that followed the standard IEEE 802.11 access protocol (with exponential back-off). The system's performance was evaluated for three values of absolute gain: 0.5, 0.6 and 0.8 and the results are presented in Fig. 7. By observing the tradeoff curves




Fig. 7. Tradeoff curves for $\frac{\eta}{n+1} = 0.5, 0.6, 0.8$ and $n = 2$.

in Fig. 7 we conclude that the system's detection delay decreases significantly as the attacker's absolute gain increases. To illustrate this claim, we observe the best case system performance, i. e. one false alarm over the whole simulation period of

1.5min, and note that the detection delay for the absolute gain of 80% is approximately 3.5 times shorter than in the case when the absolute gain is 50%. This again confirms the efficiency of our proposed detection system against most aggressive worst-case optimal attacks. In order to illustrate the influence of the number




Fig. 8. Tradeoff curve for $\frac{\eta}{n+1} = 0.5$ and $n = 2, 5$.

of legitimate competing nodes on the detection time, we compare the performance of the detection system for the case when $n = 2$ and $n = 5$. In order to obtain fair comparison, we use the same value of absolute gain, $\frac{\eta}{n+1} = 0.5$. The results are presented in Fig. 8. As expected, all nodes experience higher number of collisions in the congested environment, resulting in delayed detection. It is important to note that the traffic generation rate used in Fig. 8 is lower than the one used in Fig. 7. By observing the curves for $\frac{\eta}{n+1} = 0.5$ in both figures, we note that the detection system experiences larger delay when lower traffic rates are used, which is logical since all nodes access channel less frequently, generating smaller number of back-off samples within the same time interval.

Finally, it is important to address the issue of overhead of the proposed detection algorithm. The SPRT is highly efficient since no observation vectors need to be stored. The only storage complexity is the one needed for the pdfs f_1^* and f_0 , the thresholds “a” and “b” and the current statistic S_n . In addition to that, the SPRT algorithm is also time-efficient, since in order to compute the log-likelihood we only need to compute the ratio of two functions (f_1^* and f_0 , which are very simple to evaluate) and add this value to the current statistic S_n . Therefore, the overhead of the proposed algorithm is low and can be obtained by adding the two previously mentioned values.

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 min-max 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 $\hat{b} = b + d$.

The results provided in Section 5 confirm the necessity of collaboration among the attackers if a significant impact on the system is desired. Obviously, the strategy of an intelligent attacker depends on the number of legitimate nodes he competes with. As the number of legitimate nodes increases, the gain of the attacker who is trying to follow the min-max approach decreases. In the case of malicious nodes, whose main goal is to create disruptions in the network, the goal can be achieved by increasing the number of colluding attackers. However, this arises a serious efficiency issues since each level of disruption carries certain costs for malicious nodes. Therefore, the necessary parameters needed for estimating the efficiency of the attack can be described as follows:

—What is the minimum number of nodes that need to be involved in each attack

in order to create major disruption in the MANET functionality?

- What are quantitative metrics and relationships between the number of attack nodes and the magnitude of the disruption occurred?

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.

The addition of mobility is a very challenging perspective. Our work assumes a stationary network where the node relative positions do not change with time. In a network of mobile nodes, one would expect the detection performance to deteriorate since potential attackers move in and out of range of an observer node with an IDS system, hence the sequence of observations is intermittent. In that case, interesting topics to consider would be the impact of specific mobility patterns on the detection performance, how to engineer mobility patterns of defender nodes in order to alleviate the impact of attacks. Another interesting problem is the extent to which information can be passed in the system from nodes that have received an interrupted sequence of backoffs, such that the prior history of the attacker will be at the disposal of other observers that will apply the detection algorithm. On the other hand, a spatial dimension of the payoff of the attacker (besides the temporal one of channel access) might need to be incorporated in the model in order to account for the spatial pattern of channel denial. Intuitively, the payoff of the attacker is smaller if it misbehaves at different spatial locations in the network. Nevertheless, these movements by the attacker might be necessary in order to avoid detection.

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.

REFERENCES

- AXELSSON, S. 1999. The base-rate fallacy and its implications for the difficulty of intrusion detection. In *Proc. of the 6th ACM Conference on Computer and Communications Security (CCS '99)*. 1–7.
- BELLARDO, J. AND SAVAGE, S. 2003. 802.11 denial-of-service attacks: real vulnerabilities and practical solutions. In *Proc. of USENIX Security Symposium*. San Antonio, TX.
- BERTSEKAS, D. 2003. *Convex analysis and optimization*. Athena Scientific.

- BUCHEGGER, S. AND BOUDEC, J.-Y. L. 2002. Performance Analysis of the CONFIDANT Protocol. In *Proceedings of MobiHoc*. Lausanne.
- CÁRDENAS, A. A., BARAS, J. S., AND SEAMON, K. 2006. A framework for the evaluation of intrusion detection systems. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*. Oakland, CA.
- CÁRDENAS, A. A., RADOSAVAC, S. R., AND BARAS, J. S. 2004. Detection and prevention of MAC layer misbehavior in ad hoc networks. In *Proceedings of SASN '04*. 17–22.
- DRAGALIN, V., TARTAKOVSKY, A., AND VEERAVALLI, V. 1999. Multihypothesis Sequential Probability Ratio Tests - Part I: Asymptotic optimality. *IEEE Trans. on Information Theory* 45, 7 (Nov.), 2448 – 2461.
- GUPTA, V., KRISHNAMURTHY, S., AND FALOUTSOS, M. 2002. Denial of service attacks at the MAC layer in wireless ad hoc networks. In *Proc. of MILCOM*.
- IEEE. 1999. IEEE wireless LAN medium access control (MAC) and physical layer (PHY) specifications.
- KASSAM, S. AND POOR, H. 1985. Robust techniques for signal processing : a survey. *Proc. of the IEEE* 73, 3 (March), 433–481.
- KYASANUR, P. AND VAIDYA, N. 2003. Detection and handling of MAC layer misbehavior in wireless networks. In *Proc. of International Conference on Dependable Systems and Networks*.
- MARTI, S., GIULI, T. J., LAI, K., AND BAKER, M. 2000. Mitigating routing misbehavior in mobile ad hoc networks. In *Mobile Computing and Networking*. 255–265.
- RADOSAVAC, S., BARAS, J. S., AND KOUTSOPOULOS, I. 2005. A framework for MAC protocol misbehavior detection in wireless networks. In *Proceedings of the 4th ACM workshop on Wireless Security (WiSe 05)*. 33–42.
- RAYA, M., HUBAUX, J.-P., AND AAD, I. 2004. DOMINO: A system to detect greedy behavior in IEEE 802.11 Hotspots. In *Proceedings of MobiSys '04*. 84–97.
- ČAGALJ, M., GANERIWAL, S., AAD, I., AND HUBAUX, J.-P. 2005. On selfish behavior in CSMA/CA networks. In *Proceedings of the IEEE Infocom*.
- VERDU, S. AND H.V.POOR. 1984. On minimax robustness: a general approach and applications. *IEEE Trans. on Information Theory* 30, 2 (March), 328–340.
- WALD, A. 1947. *Sequential Analysis*. John Wiley and Sons, New York.
- WALD, A. AND WOLFOWITZ, J. 1948. Optimum character of the sequential probability ratio test. *Ann. Math. Statist.* 19, 326 – 339.

August 2006