Detection and Performance Analysis of Greedy Individual and Colluding MAC Layer Attackers

Svetlana Radosavac and John S. Baras
Department of Electrical and Computer Engineering and the Institute for Systems research
University of Maryland, College Park 20742
Email:{svetlana,baras}@isr.umd.edu

***Abstract*— Selfish behavior at the Medium Access (MAC) Layer can have devastating side effects on the performance of wireless networks, with effects similar to those of Denial of Service (DoS) attacks. In this paper we consider the problem of misbehavior detection at the MAC layer, focusing on the back-off manipulation by colluding selfish nodes. We cast the problem within a minimax robust detection framework, providing a detection rule of optimum performance for the worst-case attack. We analyze the effects of a single optimal attacker with respect to the detection delay and average number of backoff slots and compare them with the effects of colluding attackers.**

I. Introduction

With the rise and flexibility of ubiquitous computing, new and unforeseeable ways of user interactions are expected, such as establishing collaborative networks with minimum or almost no central control. The communication protocols in different layers of an ad hoc network can also be subject to manipulation by selfish users due to the fact that they were designed under the assumption that all participating nodes obey the given specifications. However, when these protocols are implemented in an environment where each node has its own authority, nodes can deviate from the protocol specification in order to obtain a given goal, at the expense of honest participants. In this paper we focus on the effects of a single optimal attacker analyzed in [2] and derive the optimal strategy for colluding selfish users at the MAC layer in ad hoc networks.

Our approach is based on sequential detection procedures, placing the emphasis on the class of attacks that incur larger gain for the attackers and is able to cope with the uncertain environment of a wireless network. Hence, the minimax robust detection approach is adopted in order 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.

II. Background Work

Due to the popularity of the IEEE 802.11, most of the work in detecting MAC layer misbehavior has focused on this protocol. Most of the work in this area has been focused on detecting back-off manipulation [4], [1]. Due to the randomness introduced in the choice of the back-off, it is difficult to distinguish among legal and misbehaving nodes. The approach proposed in [4] focuses on adversaries that are unaware of the existing detection scheme. [2] addresses the issue of intelligent adversaries by providing a theoretical foundation for the design of optimal detection schemes. The authors in [**?**] presuppose a trustworthy receiver, who assigns the back-off value to be used to the sender. 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. However, the problems of applying this protocol in ad hoc networks are (i) the receiver might not be trusted and (ii) it cannot be applied in environments with no central authority. All of the above algorithms have only focused on individual misbehaving nodes and do not consider collusion.

III. IEEE 802.11 DCF

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 the 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). 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. An unsuccessful transmission causes the value of contention window to double. In case of a successful transmission the host resets its contention window (CW) 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, gaining significant advantage in channel sharing over legitimate 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, decreasing their chances of accessing the channel.

IV. Detection of Misbehaving Colluding Nodes

The problem of misbehavior in the existing literature has assumed the existence of a single attacker, facilitating efficient detection using various approaches. To illustrate the effects of a colluding attack we analyze the communication scenario in Fig. 1.

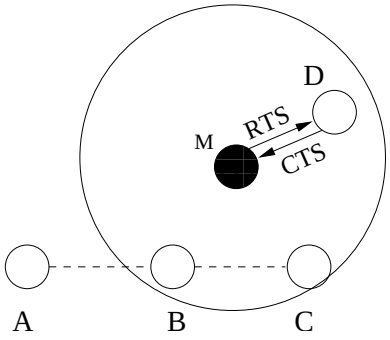

Fig. 1. Cooperation between malicious nodes M and D.

We assume that node C is in the wireless range of M and D and that it is capable of monitoring access times of its neighboring nodes. When M reserves the channel, any neighboring node can compute M's exact back-off values by listening to the sequence of its control messages. However, nodes D and M may collude and deny network access to nodes B and C. This effect can be easily achieved when back-off values of both sender and receiver are selected a priori (i.e. when both nodes select the back-off values using a pre-specified p.d.f.). Obviously, the already existing monitoring procedures do not work in this case due to the fact that both the sender and the receiver follow the specific sequence of back-off values that have been assigned a priori.

For example, the sender can choose the back-off period equal to zero and transmit immediately upon the expiration of its DIFS period. Obviously, node C cannot detect misbehavior by observing whether nodes D and M deviate from agreed back-off values and other detection procedures need to be applied.

A. Detection and Attack Assumptions

We now consider detection strategies in the presence of an intelligent misbehaving node: a node that is aware of the existence of monitoring neighboring nodes and adapts its access policy in order to avoid detection. We assume that the goal of the misbehaving hosts is to choose an optimal attack strategy that minimizes the probability of detection P_D, while maximizing their gain (access to the channel).

However, it is difficult to come up with a universal access policy for misbehaving nodes due to the random nature of the wireless channel and the nature of the access protocol itself. Therefore, our desired detection procedure needs to be *robust* (needs to perform well for a wide range of attack strategies) and needs to be able to make online decisions as the observations are revealed (to facilitate the quickest attack detection). The first condition gives rise to the application of a minimax formulation that identifies the rule that optimizes worst-case performance over the class of allowed uncertainty conditions. 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 scheme guarantees a minimum level of performance which is the best minimum level possible over all classes of attacks.

The second condition implies that sequential detection procedures need to be used. 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 [5] 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). We refer the reader to [5] for a more detailed overview of the test. The quickest detection approach applied for the case of colluding attackers does not differ from the one in the case of a single attacker and is outlined in [3].

B. Minimax Robust Detection Approach

The need to consider most significant attacks that result in higher chances of channel access for the attacker has been addressed in the previous sections. For that reason the proposed optimal detection system focuses only on the most significant attacks, limiting the resource consumption of any kind for dealing with attacks whose effect on performance is rather marginal.

The approach should also cope with the encountered 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 the 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 that 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 required observation needed for deriving a decision.

A basic notion in minimax approaches is that of a 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.
2) For the detection rule d^*, any attack f other than f^* gives better performance.

We now describe formally our approach. Let hypothesis $\mathbf{H}_0$ denote legitimate operation and thus the corresponding pdf f_0 is the uniform one. Let also Hypothesis $\mathbf{H}_1$ correspond to misbehavior with unknown pdf $f(\cdot)$.

Given the maximum allowed false alarm rate (P_{FA}) and missed detection rate (P_M), the objective of a sequential 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 is therefore 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

$$\mathbb{E}[N] = \phi(d, f). \tag{1}$$

Let $\mathcal{D}$ denote the class of all (sequential and non-sequential) statistical hypothesis tests d 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 (taking values in Ω) 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 to find d and f such that

$$\mathbb{E}[N]^* = \min_{d \in \mathcal{D}} \max_{f \in \mathcal{F}_\eta} \phi(d, f), \tag{2}$$

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 a saddle point.

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

$$\phi(d^*, f) \leq \phi(d^*, f^*) \leq \phi(d, f^*) \; \forall d \in \mathcal{D}, \; \forall f \in \mathcal{F}_\eta. \tag{3}$$

A saddle point (d^*, f^*) of ϕ consists of a detection test d^* and an attack distribution f^*. In order to find the solution of problem (2), we find the saddle point of ϕ.

However, as we now show, finding the detection strategy satisfying the saddle point is easy (if we have f^*). First, recall that the optimal detection test in the sense of minimizing expected number of samples needed for detection is the SPRT. This means that the 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 (3) has been established. Therefore in the remainder of this paper we focus on how to obtain the worst attack distribution f^* satisfying the first equality of equation (3).

1) Definition of the Uncertainty Class: The underlying assumption in the minimax approach is that the attacker has full knowledge of the employed detection rule, enabling him to create a misbehavior strategy that maximizes the number of required samples for misbehavior detection delaying the detection as much as possible.

Following the same reasoning applied in [2], we 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.

Assume that colluding and legitimate nodes 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 X_0 stand for the back-off value of legitimate user, hence it is uniformly distributed in $[0, W]$. Also, let the random variables X_1 and X_2 stand for the misbehaving nodes (attackers), with unknown pdf $f_{12}(x_1, x_2)$ with support $[0, W]$. The relative advantage of the attackers is quantified as the probability of accessing the channel, or equivalently the probability that their back-off is smaller than that of the legitimate node, $\Pr(X_0 < min(X_1, X_2))$.

If p denotes the access probability of each node, then the probability of successful channel access achieves fairness for $p^* = 1/n$, where n denotes the number of nodes competing for the channel access. If two nodes collude and compete with one legitimate node, they receive gain from their attack if $\Pr(X_0 < min(X_1, X_2)) \leq \frac{\eta}{3}$. In order to quantify this, let $\eta \in [0, 1]$ and define the class of attacks $\mathcal{F}_\eta$ as a set of functions $f_{12}(x_1, x_2)$ defined as

$$f_{12}(x_1, x_2) : \int_0^W \int_0^W \frac{\min(x_1, x_2)}{W} f_{12}(x_1, x_2) \, dx_1 \, dx_2 \leq \frac{\eta}{3}$$

This class includes attacks for which the incurred relative gain compared to legitimate operation exceeds a certain amount. The class $\mathcal{F}_\eta$ is the uncertainty class of the robust approach and the parameter η is a tunable parameter. By defining the class $\mathcal{F}_\eta$, 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.

2) Derivation of the worst-case attack: Assuming that the SPRT is used, we seek an attack distribution f^* such that $\phi(d^*, f^*) \geq \phi(d^*, f)$ for all other attacks $f \in \mathcal{F}_\eta$.

From [5] the average number of samples is

$$\mathbb{E}[N] = \frac{\mathbb{E}[S_N]}{\mathbb{E}[\Lambda]} = \frac{C}{\mathbb{E}_{12}\left[\ln \frac{f_{12}(X_1,X_2)}{f_0(X_1)f_0(X_2)}\right]} \tag{4}$$

where $f_0(x_i) = 1/W$ (denotes the uniform distribution of normal operation), $C = aP_D + b(1-P_D)$, and the expectation in the denominator is with respect to the unknown attack distribution f_{12}. Since C is a constant, the problem of finding the attack that maximizes the required number of observations reduces to the problem:

$$\min_{f_{12}} \int_0^W \int_0^W f_{12}(x_1x_2) \ln f_{12}(x_1x_2) dx_1 dx_2 \tag{5}$$

subject to the constraints,

$$\int_0^W \int_0^W f_{12}(x_1x_2) dx_1 dx_2 = 1 \tag{6}$$

$$\int_0^W \int_0^W \frac{\min(x_1x_2)}{W} f_{12}(x_1x_2) dx_1 dx_2 \leq \frac{\eta}{3} \tag{7}$$

The first constraint enforces the fact that f is a pdf and the second one holds due to the fact that $f \in \mathcal{F}_\eta$. By applying the Karush-Kuhn-Tucker (KKT) conditions, we find that the function $f_{12}^*(x_1,x_2)$ has the following form:

$$f_{12}^*(x_1,x_2) = e^{-1-\lambda} e^{-\mu \min(x_1,x_2)/W} \tag{8}$$

where λ and μ are the Lagrange multipliers that correspond to the constraints and are functions of W and η only. These can be obtained by plugging the Eq. 8 in the constraints (6) and (7).

Since $\phi(d^*, f^*) \geq \phi(d^*, f)$ for all $f \in \mathcal{F}_\eta$, we proved the left inequality in (3). We have now shown that the pair (d^*, f^*), where d^* is SPRT and $f^*(x)$ is the exponential density constitute a saddle point of ϕimplying that the problem

$$\max_{f \in \mathcal{F}_\eta} \min_{d \in \mathcal{D}} \phi(d, f) \tag{9}$$

has the same solution with (2).

V. Results

In order to quantify the performance of the attackers we introduce the *misbehavior coefficient* $\varepsilon = 1 - \eta$, $\varepsilon \in [0,1]$ where the value of $\varepsilon = 0$ corresponds to normal behavior (no additional gain for the attackers). Alternatively, we say that ε defines the class of attacks of interest since it specifies the incurred relative gain of the attacker. 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.

A. Effects of a Single Optimal Attacker on the MAC Layer

Before moving towards more complex scenarios involving interaction of colluding nodes we evaluate the effects of the optimal single attacker scenario presented in [3] implemented in the network simulator Opnet. We evaluate the performance of the attacker by measuring the average backoff of all participants. As it has already been pointed out, the attacker attempts to access the channel more frequently than his competitors by choosing backoff values from the p.d.f. that differs from uniform. Consequently, due to the exponential nature of the backoff procedure, the legitimate nodes choose larger values of backoffs on average, increasing the probability of channel access of the malicious node, forcing the legitimate nodes to choose backoff from the interval $[0, CW_{max}]$. The results of the optimal attack strategy are presented in Fig. 2 and Fig. 3.

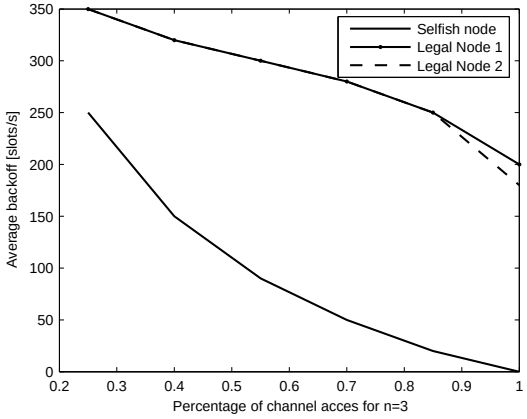

Fig. 2. Average number of backoff slots for n=3

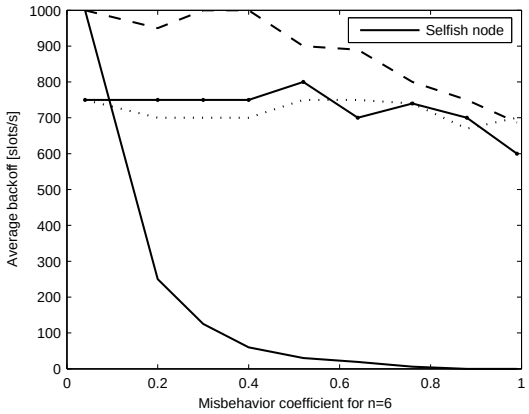

Fig. 3. Average number of backoff slots for n=6

The above results illustrate that the performance of the optimal attacker significantly depends on the number of legitimate nodes competing for channel access. The results

presented in Fig. 3 imply that the optimal attacker can choose less aggressive access strategies (and therefore significantly delay detection), achieving more devastating effects, when the number of legitimate participants increases. This behavior is due to the random nature of the backoff protocol where the probability of collision increases as the number of contenders increases. The above results imply that a single intelligent attacker is sufficient in scenarios that involve more than 4-5 legitimate nodes. However, depending on the goals and available resources, the attacker may choose to collaborate with additional malicious nodes in order to achieve more devastating effects in scenarios that involve low number of contending nodes. In order to illustrate the effects of the optimal attack in presence of $n > 5$ nodes, we analyze the average number of packets sent by legitimate and malicious nodes during the attack for $n = 6$. The results are presented in Fig. 4. As it can be seen, the data traffic sent by legitimate node is constantly low and independent from the misbehavior coefficient. This implies that in the presence of larger number of contending nodes even the least aggressive misbehavior strategy can cause more devastating effects than highly aggressive access strategy in the environments with low number of contending nodes.

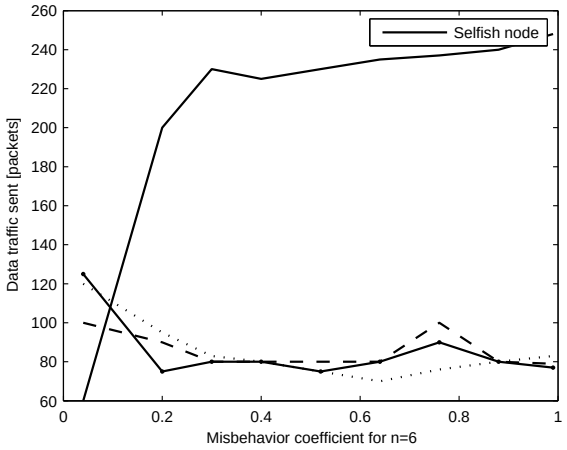

Fig. 4. Number of packets sent for different misbehavior coefficients for n=6

B. Effects of Colluding Optimal Attackers on the MAC Layer

As it has been mentioned in Sect. V-A, colluding attack strategies are useful only in environments with low number of contending nodes. It has already been pointed out that the attackers agree on the sequence of optimal backoffs by using the optimal pdf $f^*_{12}(x_1, x_2)$. IEEE 802.11 MAC is implemented and MATLAB is used to evaluate the performance of our scheme, taking into account the sequence of observed backoffs. In order to justify the introduction of colluding attacks, we compare the detection delay of colluding attackers with the detection delay of a single optimal attacker analyzed in [3]. The performance of the attackers for different probabilities of channel access is presented in Fig. 5. It can be seen that the colluding attackers incur larger detection delay for identical channel access probabilities. We also need to note that a large probability of access signifies a class of increasingly aggressive attacks for which the detection is achieved with very small delay. It is also interesting to mention that the detection delay between a single attacker and the colluding attackers decreases as the aggressiveness of the attack increases, bringing the difference close to zero for the DoS attacks.

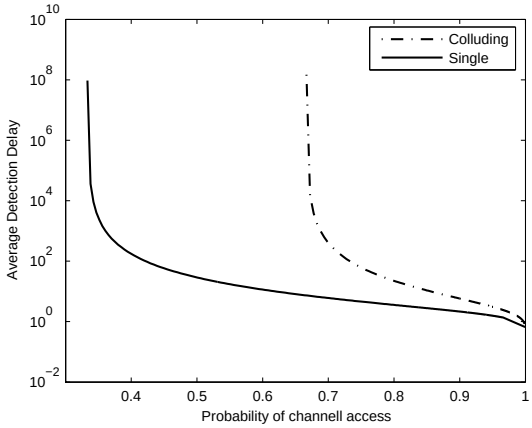

Fig. 5. Average Detection Delay for single attacker and colluding attackers for $P_{FA} = 0.01$ and $P_D = 0.99$.

VI. Conclusions and Future Work

Misbehavior at the MAC layer achieved by changing the back-off mechanism can lead to performance degradation and even DoS attacks in ad hoc networks. In this paper we have presented an algorithm that encompasses the case of intelligent colluding attackers that adapt their 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 the SPRT. We compare the performance of a single adaptive attacker in environments with low and high number of competing participants and conclude that a need for cooperation among malicious nodes exists only in scenarios that involve low number of contending nodes. It would be interesting to analyze the tradeoffs the attacker faces when switching from single to collaborative scenario.

References

[1] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: real vulnerabilities and practical solutions," in *Proc. of USENIX Security Symposium*, Washington, DC, August 2003, pp. 15–28.

[2] S. Radosavac, J. S. Baras, and I. Koutsopoulos, "A framework for mac protocol misbehavior detection in wireless a framework for MAC protocol misbehavior detection in wireless networks," in *Proceedings of the 4th ACM workshop on Proceedings of the 4th ACM workshop on Wireless security*, Cologne, Germany, September 2005, pp. 33–42.

[3] S. Radosavac, G. V. Moustakides, J. S. Baras, and I. Koutsopoulos, "An analytic framework for modeling and detecting access layer misbehavior in wireless networks," *submitted to ACM Transactions on Information and System Security (TISSEC)*, 2006.

[4] 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*, 2004, pp. 84–97.

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