

Application of Sequential Detection Schemes for Obtaining Performance Bounds of Greedy Users in the IEEE 802.11 MAC

Svetlana Radosavac and John S. Baras, University of Maryland, College Park

ABSTRACT

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 order to be able to evaluate the performance of the adversary, we define a set of security guarantees and requirements for an intrusion detection system and precisely define an adversary model. We then apply the proposed model for defining performance bounds of the worst case adversary and the quickest detection IDS in the IEEE 802.11 MAC.

INTRODUCTION

Deviation from legitimate protocol operation in wireless networks has received considerable attention in recent years. The increased level of sophistication of protocol components has led to the extreme where wireless network devices have become easily programmable. Consequently, it is feasible for a network peer to tamper with software and firmware, and abuse the protocol. The solution to the problem is the timely and reliable detection of such misbehavior instances. However, two difficulties arise: the random nature of some protocols and the nature of the wireless medium. Therefore, it is not easy to distinguish between misbehavior and an occasional protocol malfunction.

The goal of every system is to achieve robustness not only against a specific disruption, but also to maintain an acceptable performance level when such disruption occurs. This is not possible to achieve without careful design planning. In order to construct an intrusion detection system (IDS) that ensures robustness of a given system, the goal and capabilities of both the IDS and the adversary need to be defined. Only then is it possible to evaluate the performance of the sys-

tem under the worst case scenario, derive an optimal detection strategy, and determine whether the critical system parameters remain within acceptable boundaries.

In this article we first define a set of security guarantees and requirements that need to be satisfied by an IDS, and define an adversary model. We then provide motivation for employment of sequential detection methods for medium access control (MAC) layer misbehavior detection. Finally, we illustrate the efficiency of our scheme by evaluating the performance of our system against least favorable attacks. In this work the term *least favorable attack* refers to the attack that maximizes the gain of the adversary while minimizing the probability of detection, thereby achieving maximal detection delay. A *suboptimal* attack is any attack that does not satisfy the above conditions.

SECURITY GUARANTEES AND REQUIREMENTS

This section presents a set of security requirements an IDS must satisfy in order to derive a set of security guarantees (throughput, detection delay, etc.) that can then be delivered to a customer. It is important to notice that different IDSs will have different security constraints depending on the type of anomalies they aim to detect, their acceptable error tolerance (number of false alarms), and the cost involved in the construction of such systems. This consequently leads to a different set of security guarantees. Hence, each IDS can be defined by a set of security constraints it satisfies and security guarantees it delivers. In order to properly evaluate the performance of a given IDS in the remainder of the article, we first define the *goal* and *capabilities* of a detection system.

The goal of the detection system is to detect any deviation from normal behavior with the

Research supported by the U.S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011.

Svetlana Radosavac is now with DoCoMo Communications Laboratories USA, Inc.

minimum time delay and minimum probability of false alarm, P_{FA} . Decisions about the occurrence of misbehavior should be *robust* (they need to perform well for a wide range of attack strategies).

As for the capabilities of the detection system, we assume that the system is adaptive (can change its detection strategy depending on the wireless medium conditions) and intelligent (is capable of deriving a detection strategy that minimizes gain of an intelligent adversary under given security requirements).

In order to be able to obtain a set of security guarantees, an IDS needs to satisfy the following security requirements.

ROBUSTNESS AGAINST A CRITICAL CLASS OF ATTACKS

The basic requirement for every system is robustness against a critical class of attacks that is guaranteed by deploying a corresponding IDS. The critical class of attacks $\mathcal{F}$ represents the set of behaviors the system is not willing to tolerate that lead to suboptimal performance. The class of attacks outside of class $\mathcal{F}$ represent attacks that cause insignificant damage to the system and can be tolerated. Each system will have different tolerance levels for different behaviors; consequently, class $\mathcal{F}$ cannot be universally defined. We say that a system $\mathcal{S}$ is *robust* against a class of attacks $\mathcal{F}$ if its IDS can detect an adversary $A \in \mathcal{F}$ within a timeframe T , while maintaining the performance level of the system above the predefined threshold $\mathcal{P}_T$. The parameters T and $\mathcal{P}_T$ are not fixed and vary depending on how strict the security is required to be for a given system. A system $\mathcal{S}$ is *optimal* if its IDS is capable of constructing a universal detection strategy that minimizes the detection delay for the worst case attack scenario. Additionally, all nodes belonging to a robust IDS should be capable of identifying failures and anomalies in their neighborhood and notifying the rest of the network in a timely manner about such events.

RESILIENCE TO ATTACKS IN THE PRESENCE OF INTERFERENCE

Wireless networks with stringent security requirements require presence of mechanisms that can detect and isolate different classes of misbehavior, but that are also capable of functioning with acceptable performance in the presence of interference when a complete sequence of the adversary's actions cannot be obtained. We assume that a resilient system is capable of adjusting its optimal detection strategy as the operating conditions of the system change, while suffering minimal losses during the change in detection strategy.

ADVERSARY MODEL

The lack of a proper adversarial model can lead to significant decrease in system performance due to missed detection, detection delay, or a large number of false alarms. In order to properly evaluate the defense strategies and potential

damage caused by an adversary, a more formal definition of adversary capabilities and goals is provided.

INFORMATION AVAILABLE TO THE ADVERSARY

Throughout our work we adopt the strict assumption that an adversary is intelligent; that is, it knows everything the detection agent knows and can infer the same conclusions as the detection agent. This assumption enables the detector to obtain the upper bound on the detection delay.

CAPABILITIES OF THE ADVERSARY

We assume the adversary has full control over his/her actions. In order to describe the capabilities of the attacker, we use the class of attacks $\mathcal{F}$ defined earlier that describes his/her probable set of actions. It is important to note that $\mathcal{F}$ represents the *critical* class of attacks from the viewpoint of the IDS (i.e., the class of attacks that are a threat to the IDS). However, when observed from the attacker's side, $\mathcal{F}$ represents a set of actions available to the adversary and can be described as a *feasible* class of attacks.

GOAL OF THE ADVERSARY

We assume the existence of a *greedy* adversary whose objective is to design an access policy that maximizes his/her gain over the defined period of time while minimizing the probability of detection, P_D .

IEEE 802.11 MAC MISBEHAVIOR

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 backoff value b uniformly from the set $\{0, 1, \dots, W-1\}$, where W is the size of the contention window. The backoff counter decreases by one at each time slot that is sensed to be idle, and the node transmits after b idle slots. If the channel is perceived to be busy in one slot, the backoff counter freezes. After the backoff 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 CTS are required to adjust their network allocation vector (NAV) that indicates the duration for which they will defer transmission. An unsuccessful transmission instance due to collision or interference is denoted by a lack of CTS or acknowledgment (ACK) for the data sent and causes the value of the contention window to double. If the transmission is successful, the host resets its contention window to the minimum value W .

IEEE 802.11 DCF favors the node that selects the smallest backoff value among a set of contending nodes. Therefore, a malicious or selfish node may choose not to comply to protocol rules

Throughout our work we adopt the strict assumption that an adversary is intelligent, i.e. knows everything the detection agent knows and can infer the same conclusions as the detection agent. This assumption enables the detector to obtain the upper bound on the detection delay.

If the initial analysis shows that a given IDS cannot be employed in certain environments due to the strength of potential adversaries, the system administrator may choose to deploy multiple IDS units or to deploy a more expensive, but more sensitive IDS that will perform better than the original one.

by selecting small backoff intervals, thereby gaining significant advantage in channel sharing over regularly behaving honest nodes. Moreover, due to the exponential increase of the contention window after each unsuccessful transmission, non-malicious nodes are forced to select their future backoffs from larger intervals after every access failure. Therefore, the chance of their accessing the channel becomes even smaller. Although several other deviation strategies exist, this one is the most challenging to detect, and in this work we adhere to protocol deviations that occur due to manipulation of backoff values.

MISBEHAVIOR DETECTION IN THE MAC LAYER

The nature of wireless network operation dictates that decisions about misbehavior should be made online as observations are revealed and not at a fixed observation interval. The sequence of backoff values belonging to a monitored node in the IEEE 802.11 MAC is revealed sequentially, and if a node starts to misbehave, at a certain (unknown) moment some probabilistic characteristics of this process change. A detector should decide as quickly as possible whether the change has happened or not; at the same time it should keep the number of decisions about change-points when they are not present to a minimum (minimize the number of false alarms). It is now obvious that two quantities are involved in decision making: the time at which the decision about the existence of a change-point is made and a decision rule. Intuitively, this setup gives rise to the sequential detection problem since a sequential decision rule consists of a stopping time N that indicates when to stop observing and a final decision rule d_N which, at the time of stopping, decides between hypotheses H_0 (legitimate behavior) and H_1 (misbehavior). We denote the above combination $D = (N, d_N)$.

In order to proceed with our analysis we first define the properties of an efficient detector following the framework from earlier. The starting point in defining a detector is minimization of the probability of false alarms $\mathbb{P}_0[d_N = 1]$. Additionally, each detector should be able to derive the decision as soon as possible (minimize the number of samples it collects from a misbehaving station) before calling the decision function $\mathbb{E}_1[N]$. Finally, it is also necessary to minimize the probability of deciding that a misbehaving node is acting normally $\mathbb{P}_0[d_N = 0]$. It is now easy to observe that $\mathbb{E}_1[N]$, $\mathbb{P}_0[d_N = 1]$, $\mathbb{P}_1[d_N = 0]$ form a multicriteria optimization problem. However, not all of the above quantities can be optimized at the same time. Therefore, a natural approach is to define the accuracy of each decision a priori and minimize the number of samples collected.

This scheme guarantees a minimum level of performance, which is the best minimum level possible over all classes of attacks.

OPTIMAL DETECTION STRATEGY

There has been extensive analysis of the impact of MAC layer misbehavior on 802.11 MAC and routing protocols. Denial-of-service (DoS)

attacks represent an extreme instance of misbehavior studied in this work. In [1] 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. Numerous misbehavior detection schemes have been proposed over the last couple of years, focusing on either changing the protocol in order to improve its resilience to a certain class of attacks [2] or measuring certain protocol parameters and detecting deviation from normal behavior [3]. The downside of the first approach is that the protocol needs to be changed, which is costly and time consuming. The second approach focuses on a very specific class of attacks, and exhibits superior performance for detection of such attacks and fails when the adversarial strategy changes (which is frequent when the adversary is adaptive). None of the proposed schemes attempts to construct an optimal detection scheme for a given class of attacks and derive performance bounds of the adversary. We believe this is a crucial step in IDS design since every IDS needs to satisfy a set of requirements, such as detection delay and cost, in order to be efficient. If the initial analysis shows that a given IDS cannot be employed in certain environments due to the strength of potential adversaries, the system administrator may choose to deploy multiple IDS units or deploy a more expensive but more sensitive IDS that will perform better than the original one.

As we have mentioned, the nature of the problem gives rise to employment of sequential detection techniques. This setup was first proposed by Wald [4], where he also introduced the Sequential Probability Ratio Test (SPRT) as a solution. The SPRT is defined in terms of the log-likelihood ratio S_n of two joint probability functions $f_i(x_1, \dots, x_n)$ under hypothesis H_i , $i = 0, 1$. The corresponding stopping time is $N = \inf_n \{S_n \notin [A, B]\}$ and decision rule $d_N = 1$ if $S_N \geq B$ and $d_N = 0$ if $S_N \leq A$, where $A < 0 < B$ are thresholds selected so that 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 when the threshold is exceeded. Once stopped, the decision function d_N decides in favor of hypothesis H_1 when S_N exceeds the largest threshold and in favor of H_0 when S_N is below the smallest threshold. The detection delay, $\mathbb{E}[N]$, is inversely proportional to S_N and is a function of P_D [4]. Wald proved that the SPRT achieves the shortest detection delay among all sequential and nonsequential tests. To illustrate this claim, we measure detection delay of both sequential and nonsequential detectors in our experiments. The effectiveness of the proposed statistics is easy to explain: the mathematical expectation of the log-likelihood ratio is negative before and positive after the change point. Hence, if an adversary that generates his/her backoff sequence according to the p.d.f. $f_1(x)$ deviates from the legitimate p.d.f. $f_0(x)$, the log-likelihood ratio will change and the sum S_n will shift toward the upper threshold. If the adversary's distribution $f_1(x)$ significantly deviates

from $f_0(x)$, the shift will be larger, and the threshold will be crossed within a shorter time period. It then depends on the adversary's strategy: whether he/she chooses to occasionally follow the protocol and prolong detection or not. In order to explain the efficiency of the proposed detection system, we now revisit the optimal attack strategy presented in [5] and extend it with the framework presented earlier.

ADVERSARY MODEL

Following the framework presented earlier we now present an adversary model in the IEEE 802.11 MAC.

Capabilities of the Adversary — The adversary has full control over the probability mass function f_1 and the backoff values it generates. In addition, we assume that the adversary is intelligent (he/she knows everything the detection agent knows and can infer the same conclusions as the detection agent).

Goal of the Adversary — The objective of the adversary is to design an optimal access policy with the resulting probability of channel access P_A , while maximizing the detection delay $\mathbb{E}_1[N]$. The above policy results in generation of backoff sequences according to the pmf $f_1^*(x)$.

We assume an intelligent attacker is aware that an IDS is using the SPRT as a detection strategy and will stop misbehaving before it is detected. Although this may seem to be a disadvantage, it is actually not. The optimal IDS forces an adversary to either:

- Occasionally follow the protocol rules and shift below the threshold
- Apply a mild misbehavior strategy that is below the threshold at all times
- Relocate as soon as the threshold is approached

In the first two the attacker has to stop misbehaving or compromise by achieving a very mild advantage over other participants. In the third case the deployment of an optimal IDS forces an adversary to relocate frequently, therefore increasing the cost of launching an attack. It is important to note that the relocation space of an adversary is not infinite (a greedy user has to send packets to another node). Unless there is a set of collaborating adversaries, an adversary that chooses to employ aggressive misbehavior policy is quickly detected.

Although the access policy $f_1^*(x)$ of the adversary was derived under the assumption that the detection algorithm is known, $f_1^*(x)$ is a good adversarial policy against *any* detector. Namely, it is easy to see that the detection delay is inversely proportional to the Kullback-Leibler divergence between f_0 and f_1 . It is known that bounds on the probability of detection and false alarm for an optimal detector can be expressed in terms of the Kullback-Leibler divergence between the distribution of the two hypotheses [6]. Applying the results from information theory, the probability of detection of the optimal decision algorithm is lower bounded by $1 - 2^{-nD(f_1||f_0)}$.

It is now clear that an adversary that tries to minimize the probability of detection will

attempt to minimize the distance between distributions f_0 and f_1 , leading to the same $f_1^*(x)$ already obtained. It is essential to emphasize that by employing this adversarial strategy (i.e., choosing a *class* of attacks rather than a single access strategy), the adversary can easily adapt to the environment and maximize detection time against *any* IDS, not only the optimal one.

SPRT OPTIMALITY FOR ANY ADVERSARY IN F

Let $\Phi(D, f_1) = \mathbb{E}_1[N]$. We note that the solution for $f_1^*(x)$ was obtained in the form

$$\max_{f_1 \in \mathcal{F}} \min_{D \in \mathcal{D}_{a,b}} \Phi(D, f_1). \quad (1)$$

That is, we first minimized $\Phi(D, f_1)$ with the SPRT (minimization for any f_1) and then found the f_1^* that maximized $\Phi(\text{SPRT}, f_1^*)$.

However, an optimal detector needs to minimize *all* losses due to the worst case attacker. That is, the optimal test should be obtained by the following optimization problem:

$$\max_{D \in \mathcal{D}_{a,b}} \min_{f_1 \in \mathcal{F}} \Phi(D, f_1). \quad (2)$$


The proposed solution also satisfies this optimization problem since it forms a saddle point equilibrium, and we can claim that for every $D \in \mathcal{D}_{a,b}$ and every $f_1 \in \mathcal{F}$: $\Phi(D^*, f_1) \leq \Phi(D^*, f_1^*) \leq \Phi(D, f_1^*)$. More specifically, the existence of a saddle point ensures that for the attack f_1^* any detection rule $\mathcal{D}$ other than $\mathcal{D}^*$ has worse performance. $\mathcal{D}^*$ is the optimal detection rule for attack f_1^* in terms of minimum (average) number of required observations. In addition to that, it ensures that for the detection rule $\mathcal{D}^*$, any attack f_1 from the uncertainty class, other than f_1^* , gives better performance (the detection rule $\mathcal{D}^*$ has its worst performance for attack f_1^*).

As a consequence of the existence of a saddle point, no incentive for deviation from (D^*, f_1^*) for any of the players is offered.


EXPERIMENTAL RESULTS

We now proceed to experimental evaluation of the analyzed detection schemes. In this work we assume the existence of an intelligent adaptive attacker that is able to adjust its access strategy depending on the level of congestion. In order to minimize the probability of detection, the attacker chooses legitimate over selfish behavior when the congestion level is low and an adaptive selfish strategy otherwise. Therefore, we assume that all stations have packets to send at any given time. We assume that the attacker employs the least favorable misbehavior strategy for our detection algorithm, enabling us to estimate the maximal detection delay. This setting also represents the worst case scenario with regard to the number of false alarms per unit of time because the detection algorithm is forced to make a maximum number of decisions per unit of time (the number of alarms should be smaller in practice). The backoff distribution of an optimal attacker was implemented in the Opnet network simulator, and tests were performed for various levels of false alarms. The simulations were performed with nodes that followed the IEEE 802.11 access protocol. The corresponding scenario is present-


We assume that an intelligent attacker is aware that an IDS is using the SPRT as a detection strategy and will stop misbehaving before it is detected. Although this may seem to be a disadvantage, it actually is not.



■ **Figure 1.** Simulation scenario: two legitimate participants compete with the adversary.



■ **Figure 2.** Average number of data packets sent for legitimate and malicious nodes as a function of absolute gain.



■ **Figure 3.** Effects of the SPRT-based IDS on detection of the adversary from Fig. 1 that attempts to access the channel 60 percent of time.

ed in Fig. 1. We consider the scenario where one adaptive intelligent adversary competes with two legitimate stations for channel access. Consequently, in a fair setting each protocol participant is allowed to access the medium for 33 percent of the time under the assumption that each station is backlogged and has packets to send at any given time slot. The detection agent was implemented such that any observed backoff value $X_i > W$ was set up to be W .

In order to illustrate the power of the adversary that applies the backoff policy f_1^* , we introduce *absolute gain*, defined as $\eta/n+1$ for $1 < \eta < \eta+1$, where η represents the level of misbehavior ($\eta = 1$ represents legitimate behavior), and n represents the number of legitimate users competing for channel access. We observe that the maximum value of the absolute gain is equal to 1 and corresponds to the DoS attack, and the minimum value is equal to $1/n+1$ and corresponds to legitimate behavior of the observed node. To illustrate the effect of the worst case MAC layer attack, we observe the average number of data packets sent by each node as a function of absolute gain. The effects of the adversary on the performance of a legitimate node are presented in Fig. 2. With the increase of aggressiveness of the selfish node, the legitimate nodes are denied access to the channel by choosing larger backoff values and are not able to send data, since the selfish node gains a higher percentage of channel access. If the proposed SPRT-based detection system is employed, misbehavior is almost instantly detected. We represent the effects of the detection scheme on an adversary that attempts to access the channel 60 percent of time (instead of 33 percent, which corresponds to legitimate behavior) in Fig. 3. We now observe that misbehavior is detected almost instantly, and once the adversary is removed from the environment, the legitimate nodes resume their normal operation (we assume that the misbehaving node was replaced with a legitimate node).

The consequence of Wald's theorem is that no other detection strategy can do better than the SPRT, and the consequence of the min-max formulation is that the worst SPRT performance is exhibited when misbehavior strategy $f_1^*(x)$ is employed (i.e., any other attack will be detected faster). To illustrate this, we implement the attack distributions $f_1^p(x)$ and $f_1^i(x)$ from [3, 5] in the OPNET network simulator and collect a corresponding sequence of backoff values for each scenario. We test the performance of the SPRT (sequential) and DOMINO (nonsequential) detection schemes against both attacks. We perform our evaluation in the form of trade-off curves between T_d (mean time between detections) and T_{fa} (mean time between false alarms) for both algorithms. Figure 4 confirms that the backoff distribution f_1^* is indeed least favorable since the detection delay of such attack strategies is larger than the corresponding detection delay for other strategies such as the DOMINO attack, $f_1^p(x)$, for any sensitivity of an IDS scheme (i.e., for any P_{fa}). Figure 5 compares the performance of the sequential (SPRT) and nonsequential (DOMINO) schemes for detection of least favorable attacks. Again, the sequential nature

of our scheme enables the quickest detection; consequently, the detection delay for the DOMINO scheme is up to 12 times higher than for the SPRT scheme. This confirms our claim that if a scheme is constructed for detection of a specific class of attacks, its performance against any other class of attacks becomes suboptimal. On the other hand, since the SPRT scheme was optimized for performance under the least favorable attack, it detects any other class of attacks within a shorter time period.

CONCLUSIONS


This article emphasizes the importance of deriving *optimal* detection schemes and establishing strict performance bounds for an IDS of interest for a given set of security requirements. In the case of IEEE 802.11 MAC, the detection scheme that provides the best performance for *any* class of attacks is of a sequential nature. However, using the proposed framework, it is possible to derive an optimal detection strategy for any other scenario and protocol.

It is also important to point out that the SPRT is a *parametric* statistic, while DOMINO belongs to the class of *nonparametric* statistics. Nonparametric statistics are easier to apply since they do not require exact models of distributions. They only require knowledge of some parameters (e.g., *nominal backoff* in DOMINO). Since such tests consider only certain parts of a distribution, they allow a very large class of probability distributions. The advantage of such tests is that they let us deal with unknown probability distributions, but at the same time they throw away a lot of information about the problem. A parametric statistic, on the other hand, needs a model for the distributions. If it has both models, it should perform better than the corresponding nonparametric statistic. Due to the fact that using the wrong distribution f_1 significantly deteriorates the performance of detection schemes, we used robust statistics, where the basic idea is to find the least favorable distribution f_1^* , which guarantees that any other distribution $f_1 \neq f_1^*$ results in a suboptimal adversarial strategy and is detected with smaller detection delay.


The SPRT not only exhibits the best performance over all sequential and nonsequential tests, but it is also highly efficient since no observation vectors need to be stored. The only storage complexity is the one needed for f_1 and f_0 , thresholds a and b , and the current statistics S_n . The SPRT algorithm is time-efficient since in order to compute the log-likelihood we only need to compute the ratio of two functions and add this value to S_n . The overhead is low and can be calculated by adding the two previously mentioned values. Consequently, the proposed approach can be used for online detection.

REFERENCES

- [1] V. Gupta, S. Krishnamurthy, and M. Faloutsos, "Denial of Service Attacks at the MAC Layer in Wireless Ad Hoc Networks," *Proc. IEEE MILCOM*, Oct. 7–11, 2002.
- [2] P. Kaysanur and N. Vaidya, "Detection and Handling of MAC Layer Misbehavior in Wireless Networks," *Proc. Int'l Conf. Dependable Sys. and Networks*, June 2003.



■ **Figure 4.** Trade-off curves of the SPRT scheme for detection of least favorable (f_1^*) and suboptimal (f_1^D) attacks.



■ **Figure 5.** Trade-off curves of SPRT and DOMINO schemes for detection of least favorable attacks.

- [3] M. Raya, J.-P. Hubaux, and I. Aad, "DOMINO: A System to Detect Greedy Behavior in IEEE 802.11 Hotspots," *Proc. MobiSys '04*, 2004, pp. 84–97.
- [4] A. Wald, *Sequential Analysis*, Wiley, 1947.
- [5] S. Radosavac, J. S. Baras, and I. Koutsopoulos, "A Framework for MAC Protocol Misbehavior Detection in Wireless Networks," *Proc. 4th ACM Wksp. Wireless Security*, 2005, pp. 33–42.
- [6] R. E. Blahut, *Principles and Practice of Information Theory*, Addison-Wesley, 1987.

BIOGRAPHIES

JOHN S. BARAS [F] (baras@isr.umd.edu) received a B.S. in electrical engineering with highest distinction from the National Technical University of Athens, Greece, in 1970. He received M.S. and Ph.D. degrees in applied mathematics from Harvard University, Cambridge, Massachusetts, in 1971 and 1973, respectively. Since 1973 he has been with

the Department of Electrical and Computer Engineering, University of Maryland at College Park, where he is currently a professor, a member of the Applied Mathematics and Scientific Computation Program Faculty, and an affiliate professor in the Department of Computer Science. From 1985 to 1991 he was the founding director of the Institute for Systems Research (ISR), one of the first six NSF Engineering Research Centers. In February 1990 he was appointed to the Lockheed Martin Chair in Systems Engineering. Since 1991 he has been the director of the Maryland Center for Hybrid Networks (HYNET), which he co-founded. He has held visiting research scholar positions at Stanford, MIT, Harvard, the Institut National de Recherche en Informatique et en Automatique, the University of California at Berkeley, Linköping University, and the Royal Institute of Technology in Sweden. He has received several awards throughout his career. His research interests include control, communication, and computing systems. He is a Foreign Member of the Royal Swedish Academy of Engineering Sciences. He is a member of ACM, SIAM, AMS, AIAA, ATA, and Sigma Xi. He has published more than 500 refereed publications, graduated 60 Ph.D. students, and sponsored 40 postdoctoral scholars. He was the editor of the book *Recent Advances in Stochastic Calculus* (Springer, 1990). He holds three patents and has three more patents pending. He has co-founded three small companies. He

was the initial principal architect of the ISR M.S. program in systems engineering. More recently he has been heavily involved in the development of new core courses for systems engineering, addressing the need for a new integrative approach to engineering. He has served on the IEEE Engineering R&D Committee, the Aerospace Industries Association advisory committee on advanced sensors, the IEEE Fellow evaluation committee, and the IEEE Control Systems Society Board of Governors (1991–1993). He is currently serving on the editorial boards of *Mathematics of Control, Signals, and Systems*; *Systems and Control: Foundations and Applications*; *IMA Journal of Mathematical Control and Information*; and *Systems Automation — Research and Applications*.

SVETLANA RADOSAVAC (svetlana@isr.umd.edu) received a B.S. degree in electrical engineering from the School of Electrical Engineering at the University of Belgrade, Serbia, in 1999, and M.S. and Ph.D. degrees in electrical and computer engineering from the University of Maryland, College Park in 2002 and 2007, respectively. She is currently working as a research engineer at DoCoMo Communications Laboratories USA, Inc., Palo Alto, California. Her research interests focus on wireless network security and information assurance, non-cooperative and cooperative dynamic games, detection of DDoS attacks, and Internet economy.