This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2010 proceedings.

PHY-Authentication Protocol for Spoofing Detection in Wireless Networks

Liang Xiao *, Alex Reznik †, Wade Trappe ‡, Chunxuan Ye†, Yogendra Shah †, Larry Greenstein ‡ and Narayan Mandayam‡
*Department of Communication Engineering, Xiamen University, China
Email: lxiao@xmu.edu.cn
†InterDigital, King of Prussia, PA, USA
Email:{alex.reznik,chunxuan.ye,yogendra.shah}@interdigital.com
‡WINLAB, Rutgers University, North Brunswick, NJ, USA
Email:{trappe,ljg,narayan}@winlab.rutgers.edu

***Abstract*—We propose a PHY-authentication protocol to detect spoofing attacks in wireless networks, exploiting the rapid-decorrelation property of radio channels with distance. In this protocol, a PHY-authentication scheme that exploits channel estimations that already exist in most wireless systems, cooperates with any existing—either simple or advanced—higher-layer process, such as IEEE 802.11i. With little additional system overhead, our scheme reduces the workload of the higher-layer process, or provides some degree of spoofing protection for "naked" wireless systems, such as some sensor networks. We describe the performance of our approach as a function of the spoofing pattern and the snapshot performance that can be easily measured through field tests. We discuss the implementation issues of the authentication protocol on 802.11 testbeds and verify its performance via field tests in a typical office building.**

I. Introduction

Security has become a significant concern in wireless networks, as wireless platforms are being used to access an increasing amount of security-sensitive services, such as e-commerce and online banking. Although conventional cryptographic security mechanisms are essential to securing wireless networks, these techniques do not directly leverage the unique properties of the wireless domain to address security threats: First, the broadcast nature of radio implies that no physical connection is required for access to a wireless network. As a consequence, wireless networks are open to intrusion from the outside without the need for a physical connection. Second, mobile devices usually identify themselves by their MAC address, which may be easily spoofed.

Most existing wireless networks have some security mechanisms. For example, 802.11 systems are equipped with three link-layer security protocols, including the wired equivalent privacy (WEP), Wi-Fi Protected Access (WPA), and 802.11i. The most recent security standard, 802.11i, retains many WPA features, such as TKIP/Michael and 802.1x, and replaces the RC4 stream cipher with the advanced encryption standard (AES) in counter mode with CBC-MAC Protocol (CCMP) to provide better data encryption [1]. In spite of all these efforts, 802.11 systems still have security vulnerabilities, such as the vulnerability to spoofing attacks, where an attacker claims to be another node by using his identity, such as the MAC address. By spoofing some management frames/control messages, the attacker can further perform other attacks, such as session hijacking, man-in-the-middle and denial-of-service (DoS) attacks; and these can be performed even in presence of these advanced security mechanisms [2], [3].

To address this problem, a group of PHY-authentication techniques have been proposed to detect identity-based attacks in wireless networks, exploiting the received signal strength (RSS) [3]–[5] and channel impulse response (CIR) [6]. Note that channel responses contain more location-specific information than RSS, and hence channel response-based techniques can provide more fine grained information than those using RSS. Concurrent to these efforts, a scheme based on channel frequency response, called fingerprints-in-the-ether (FP), that detects spoofing and Sybil attacks using hypothesis testing was defined and further explored in [7]–[9]. Wireless communications are susceptible to channel fading, which varies over time in an "unpredictable" way and can be viewed as a fingerprint. More specifically, FP utilizes the rapid spatial decorrelation property of the channel response to discriminate among transmitters. Propagation theory has shown that the channel response decorrelates as the transmitter changes its location on the order of the signal wavelength [10].

By using both ray-tracing and stochastic channel modeling techniques, we have previously demonstrated that FP can accurately detect spoofing attacks for stationary or very slowly moving terminals [7], [8]. However, the previous work assumes that the correct reference channel response is always available. This assumption does not hold in general, since the channel response changes after a channel coherence time [11], and this negatively impacts the performance. In addition, in designing a full communication system we aim to supplement and improve the security mechanisms already existing in the wireless network.

Therefore, in this paper, we propose a hybrid authentication protocol to integrate the FP algorithm into existing wireless systems, cooperating with any existing higher-layer security mechanism, either "nominal" or advanced. We provide a performance bound for spoofing detection using this protocol in generalized scenarios and without assuming a reliable reference channel record. We implement our PHY-

978-1-4244-5638-3/10/$26.00 ©2010 IEEE

authentication scheme on InterDigital's 802.11 Physical Layer Security Platform (IPLSVP) and describe its performance in field tests.

This paper is organized as follows: We present our PHY-authentication protocol that integrates the PHY-authentication with the existing security mechanisms in wireless systems in Section II. We analyze the spoofing detection performance of the PHY-authentication in Section III. Section IV provides experimental results using an 802.11 platform. We conclude the paper in Section V.

II. Protocol Design of PHY-Authentication

Wireless communications are susceptible to channel fading, which vary over time in an unpredictable way. A PHY-authentication algorithm, called fingerprints-in-the-ether (FP), has been proposed to detect spoofing attacks in wireless networks. This algorithm is based on the spatial decorrelation property of the channel response, i.e., the channel response decorrelates as the location of the transmitter changes on the order of the signal wavelength [10]. Hence in FP, we reuse the pilot/premable-based channel estimation mechanism, which exists in most radio receivers, to discriminate among transmitters and thus detect spoofing attacks with little additional system overhead. The previous work in [7], [8] only covers a simplified "snapshot" scenario, while in this paper, we devise a PHY-authentication protocol to integrate FP into real systems under a generalized attack scenario.

We consider an Alice-Bob-Eve model where Alice and Bob are the legitimate transmitter and receiver, respectively. An adversary, Eve, injects spoofing messages into the networks, in hopes of spoofing Alice by using her identity, such as her MAC address in the messages. When receiving a new message that claims to be Alice, Bob has to decide whether it is a spoofing message sent by Eve. Suppose that a proportion, P_a, of all the messages that claim to be from Alice are actually sent by Alice, and the others are spoofing messages from Eve.

A. PHY-Authentication in "Snapshot" Scenario

In a simplified "snapshot" scenario previously assumed in [7], [8], Bob receives a new message claiming to be sent by Alice, and he has to check whether the claim is true. To this end, in FP, Bob utilizes the pilots/premables within the message to estimate the channel response and records the resulting channel vector, $\underline{\hat{H}}_1$, where the subscript "1" indicates that it is a new channel estimation for Alice. The vector contains M channel gains at different tones, or equivalently, their inverse Fourier transform. The size of the channel vector, M, increases with the system bandwidth, the number of pilot/preamble symbols, the number of antennas in the system, and how dispersive the channel is with respect to channel bandwidth.

Based on the uniqueness of the channel responses, Bob discriminates between Alice and Eve by comparing the new channel samples, $\underline{\hat{H}}_1$, with his reference channel record for Alice, $\underline{\hat{H}}_0$: If they are "close" to each other, Bob concludes that the new message is indeed from Alice; otherwise, Bob sends a spoofing alarm. To this end, we have derived a generalized likelihood ratio test under the generalized Rayleigh channel model. By assuming small channel time variation and estimation error, we have further simplified the test into

$$L(\underline{\hat{H}}_1, \underline{\hat{H}}_0) = \left\| \underline{\hat{H}}_1 - \underline{\hat{H}}_0 e^{jArg(\underline{\hat{H}}_1 \underline{\hat{H}}_0^H)} \right\|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \eta, \qquad (1)$$

where $||\cdot||$ denotes the Frobenius norm, the superscript ^{H} represents Hermitian transformation, and η is the test threshold [8].

The selection of the test threshold η significantly impacts the system performance: if η is too high, FP cannot detect any spoofing message. On the other hand, every message will be rejected if η is too low. By applying the Neyman-Pearson test, we have derived the test threshold for given false alarm rate α as $\eta = 0.5\sigma_A^2 \left(\frac{1+a^2}{\rho} + 1 - a^2 + b\right) F^{-1}_{\chi^2_{2M}}(1-\alpha)$, where $F^{-1}_{\chi^2_{2M}}(\cdot)$ is the inverse function of the cumulative distribution function of the chi-square distribution with order $2M$; σ_A^2 is the power of the receiver noise in the channel estimation; and a, b and ρ are all channel parameters corresponding to the specific propagation environment [8]. Since these parameters are usually unknown in real systems, we propose two practical strategies for threshold selection:

- *Pre-assigned threshold.* A fixed threshold η is used. We can choose a value corresponding to good performance in most situations. This strategy does not provide an optimal η for each specific scenario, and this pre-assigned threshold may require large scale field tests to determine.
- *Adaptive threshold.* Alice first sends $N_{tr} + 1$ training messages to inform Bob about the range of her channel. Based on these spoofing-free messages, Bob calculates the corresponding test statistics $L(k)$, $k = 2, \cdots, N_{tr}+1$. Then Bob searches the o-th percentile value of these N_{tr} test statistics, and sets it as the test threshold η in the subsequent FP test. In this way, the test threshold η is adjusted according to the channel property of the specific environment.

  Note that we can use the higher-layer protocol throughout the training stage, in order to ensure that all the training messages for the threshold calculation are not spoofed. Moreover, this adaptive threshold policy ideally enables a better tradeoff between the false alarm and miss detection of spoofing attacks. However, the performance is sensitive to o.

The previous work on FP assumes that Bob has a reliable reference channel, $\underline{\hat{H}}_0$, in his spoofing detection process. However, this assumption does not always hold in practice:

First, the reference channel record might correspond to a spoofing message that has successfully fooled Bob. In particular, the FP test is prone to error propagation, unless an additional mechanism is used. For instance, once accepting one spoofing message from Eve, Bob saves the channel vector corresponding to the spoofing message as $\underline{\hat{H}}_0$. He is now likely to accept Eve and reject Alice in the following FP test. Second, as the channel response decorrelates after a channel coherence time [11], the use of stale reference channel data increases the false alarm rate of FP in spoofing detection.

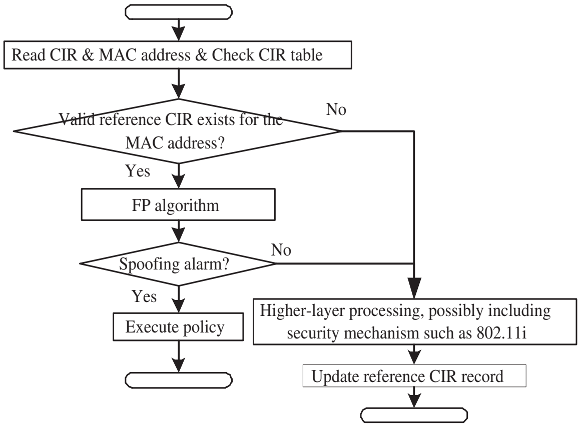
Read CIR & MAC address & Check CIR table
Valid reference CIR exists for the MAC address?
No
Yes
FP algorithm
Spoofing alarm?
No
Yes
Execute policy
Higher-layer processing, possibly including security mechanism such as 802.11i
Update reference CIR record

Fig. 1. Flow chart of the PHY-authentication protocol that integrates the PHY-authentication (FP) scheme with the higher-layer process, where the channel vector is obtained by the channel estimation using the pilot/preamble symbols in the message.

Hence, $\underline{\hat{H}}_0$ derived more than a channel coherence time earlier is no longer reliable. Finally, Bob sometimes does not have a valid reference channel at all, if Alice has kept silence during a long time or all her messages have been falsely rejected by Bob.

Therefore, we have to analyze the FP scheme without assuming a reliable $\underline{\hat{H}}_0(k)$. To do so, we build a hybrid PHY-authentication protocol to detect spoofing attacks in a general setting.

B. PHY-Authentication Protocol

In the PHY-authentication protocol, the FP test is integrated with a higher-layer security process that provides various degrees of security protection, possibly including key-based authentication or authorization. The higher-layer process can be either advanced, such as IEEE 802.11i, or simple. We intend to keep the existing process in the protocol design. In the PHY-authentication protocol, the receiver accepts a message only if both FP and the higher-layer process accept it. By filtering out most spoofing messages, FP reduces the workload of the higher-layer security mechanism or provides spoofing protection to some degree for "naked" wireless systems.

The flow chart of the protocol is presented in Fig. 1. When receiving a new message at time k, Bob collects the channel vector, $\underline{\hat{H}}_1(k)$, from channel estimation, identifies the claimed sender by checking the MAC address, and then checks whether he has a valid reference channel record for that user. If Bob has $\underline{\hat{H}}_0(k)$, the valid reference channel record, he performs the FP test to compare $\underline{\hat{H}}_0(k)$ and $\underline{\hat{H}}_1(k)$. If these two channel vectors are significantly different, FP issues a spoofing alarm and Bob processes the message according to an established security policy. Two exemplary policies are:

- Simply discard the message. In this strategy, the test threshold η has to be designed to ensure a low false alarm rate. We assume this policy in this section unless specified otherwise.
- Use a *simplified* higher-layer authentication algorithm to double-check the message, and discard the message if it fails this second check. This strategy has higher detection accuracy, though at the expense of higher system overhead, compared to the previous policy.

If the two channel vectors are similar, FP accepts the message and passes it to the higher-layer process. Same holds for messages where FP was suspended due to a lack of valid reference channel records.

To perform the FP test, we have to maintain a table to record and update the reference channel data for active users. The channel sample $\underline{\hat{H}}_1(k)$ is saved as the new reference channel record at time $k+1$, i.e., $\underline{\hat{H}}_0(k+1) = \underline{\hat{H}}_1(k)$, if Bob accepts the message. In addition, channel response decorrelates after the channel coherence time, and hence the use of stale channel data increases the false alarm rate of FP. Accordingly, we set a timer for each record, and delete the outdated channel data in the channel table once their timers reach the maximum lifetime, N_T. The design goal is to set N_T less than the possible channel coherence time. In summary, the reference channel record is maintained according to

$$\underline{\hat{H}}_0(k) = \begin{cases} \underline{\hat{H}}_1(k-1), & \text{if accepting the } (k-1)\text{-th message} \\ \underline{\hat{H}}_0(k-1), & \text{if rejecting the } (k-1)\text{-th message,} \\ & \quad \text{and Elapsed Time } \le N_T \\ \text{No ref.,} & \quad \text{otherwise.} \end{cases} \tag{2}$$

In this protocol, a spoofing message can fool Bob only when both FP and the higher-layer processing fail to detect it. The FP decision, $I_1(k)$, depends on the difference between the new channel vector $\underline{\hat{H}}_1$ and the reference channel record $\underline{\hat{H}}_0$. Moreover, the successfully authenticated message is used to initialize $\underline{\hat{H}}_0$, the channel estimate for FP. More specifically, the reference channel is updated according to (2), based on the previous system decisions I_a and the timer lifetime limit N_T.

III. Performance Analysis

A. Performance Metrics

The PHY-authentication protocol consists of FP and a higher-layer process. For each of these authentication units, we define its performance metrics in spoofing detection, including the false alarm rate, as the probability of falsely rejecting Alice, and the miss rate, as the probability of falsely accepting Eve.

Let $I_1(k)$ denote the FP decision, where the value "0" represents no alarm, "1" denotes a spoofing alarm, and "2" indicates a FP suspension due to the lack of $\underline{\hat{H}}_0(k)$. If the channel samples, $\underline{\hat{H}}_1(k)$ and $\underline{\hat{H}}_0(k)$, are so different that $L \geq \eta$, FP sends a spoofing alarm with $I_1(k) = 1$. Thus, the FP decision function can be written as

$$I_1(k) = \begin{cases} 0 \text{ (No alarm)}, & L(\underline{\hat{H}}_1(k), \underline{\hat{H}}_0(k)) < \eta \\ 1 \text{ (Spoofing alarm)}, & L(\underline{\hat{H}}_1(k), \underline{\hat{H}}_0(k)) \geq \eta \\ 2 \text{ (FP suspension)}, & \text{No } \underline{\hat{H}}_0(k). \end{cases} . \tag{3}$$

We first consider the performance of FP in the snapshot scenario, and use $S(\underline{\hat{H}})$ =A(lice) or E(ve) to denote the actual sender of the message from which the channel vector $\underline{\hat{H}}$ is derived. The "snapshot" assumption indicates $S(\underline{\hat{H}}_0(k)) \equiv$ A. Thus the snapshot performance, including the false alarm rate α and the miss rate β, previously defined in [7], can be rewritten as

$$\alpha = \Pr(I_1(k) = 1|S(\underline{\hat{H}}_1(k)) = S(\underline{\hat{H}}_0(k))) \qquad (4)$$

$$\beta = \Pr(I_1(k) = 0|S(\underline{\hat{H}}_1(k)) \neq S(\underline{\hat{H}}_0(k))). \qquad (5)$$

Both of them can be conveniently measured through field tests under the simple snapshot scenario. We have even obtained their closed-form expression using stochastic channel modeling [8].

Note that the FP test decision $I_1(k)$ actually has *three* possible states in general, and both $I_1(k) = 0$ and 2 lead to the same process afterwards. Thus, we define the performance of FP, including the false alarm rate, P_{FA}, and the miss detection rate, P_M, as:

$$P_{FA} = \Pr(I_1(k) = 1|S(\underline{\hat{H}}_1(k)) = \mathrm{A}) \qquad (6)$$

$$P_M = \Pr(I_1(k) \neq 1|S(\underline{\hat{H}}_1(k)) = \mathrm{E}). \qquad (7)$$

Next, the higher-layer authentication decision is denoted as $I_2(k)$, which equals 1 if a spoof is found, and 0 otherwise. Hence, we define the false alarm rate and miss rate of the higher-layer test, α_2 and β_2, respectively, by

$$\alpha_2 = \Pr(I_2(k) = 1|S(\underline{\hat{H}}_1(k)) = \mathrm{A}) \qquad (8)$$

$$\beta_2 = \Pr(I_2(k) = 0|S(\underline{\hat{H}}_1(k)) = \mathrm{E}). \qquad (9)$$

In our PHY-authentication protocol, Eve can successfully spoof Alice if and only if both FP and the higher-layer test miss it, with $I_1(k) \neq 1$ and $I_2(k) = 0$. The flow chart of the protocol shows that the final decision $I_a(k)$ is given by

$$I_a(k) = \begin{cases} 1 \text{ (Reject)}, & \text{if } I_1(k) = 1 \text{ or } I_2(k) = 1 \\ 0 \text{ (Accept)}, & \text{o.w.} \end{cases} . \qquad (10)$$

Let P_{FAA} and P_{MA} denote the overall performance of the authentication protocol, and by (6)-(10), we simplify them into

$$P_{FAA} = \Pr(I_a(k) = 1|S(\underline{\hat{H}}_1(k)) = \mathrm{A})$$
$$= (1 - \alpha_2)P_{FA} + \alpha_2 \qquad (11)$$

$$P_{MA} = \Pr(I_a(k) = 0|S(\underline{\hat{H}}_1(k)) = \mathrm{E}) = P_M\beta_2. \qquad (12)$$

It is clear that $P_{FAA} \geq \alpha_2$ and $P_{MA} \leq \beta_2$. For the FP with small P_{FA} and P_M, we have $P_{FAA} \approx \alpha_2$ while $P_{MA} \ll \beta_2$. Thus, the use of FP improves the system performance in terms of spoofing detection, compared to the original system only with a higher-layer process. As an extreme case, we consider a sensor network without any higher-layer security mechanism, i.e., $\alpha_2 = 0$ and $\beta_2 = 1$. If we implement the FP with $P_F = P_M = 10\%$, the overall false alarm rate P_{FAA} rises slightly from 0 to 10%, while the miss rate P_{MA} dramatically decreases from 1 to 10%.

Meanwhile, the use of FP significantly reduces the workload of the higher-layer process. Let C denote the overhead for the higher-layer functions to process a message, and P_a as the proportion of the legal messages. It is clear that the implementation of FP reduces the workload of the higher-layer functions from C to $C((1 - P_a)P_M + P_a(1 - P_{FA}))$, which is $0.74C$ with $P_a = 80\%$ and $P_F = P_M = 10\%$. This saving increases with the original computational overhead C and the portion of spoofing messages $1 - P_a$.

It is also shown in (11) and (12) that the overall system performance is directly determined by P_{FA} and P_M, instead of their "snapshot" counterparts, α and β. The generalized performance of FP depends on many factors, including α and β, the maximum timer lifetime (N_T), the higher-layer performance (α_2 and β_2), and the attack model, such as how often Eve injects spoofing messages. Consequently, it is very difficult to perform large-scale field tests to extensively measure P_{FA} and P_M, while their snapshot counterparts are much easier to obtain via field tests. Hence, we will consider two extreme cases of the higher-layer test in order to bound the performance of FP.

B. Upper-Bound with Ideal Higher-layer Process

For the authentication protocol shown in Fig. 1, the performance of FP can be upper bounded by an "ideal" higher-layer authentication test with $\alpha_2 = \beta_2 = 0$. It is clear by (11) and (12) that $P_{FAA} = P_{FA}$ and $P_{MA} = 0$, indicating that Eve has no chance to spoof Alice here. But note that Bob will still perform the higher-layer process to deal with some spoofing messages that are missed by FP. In this ideal case, the reference channel record, if it exists, always comes from Alice's message. For simplicity, we assume that the timer lifetime N_T is an integer, and that Bob receives exactly one message at each discrete time. Suppose the senders of these messages are i.i.d. Bernoulli distributed with

$$P_a = \Pr(S(\underline{\hat{H}}_1(k)) = A) = 1 - \Pr(S(\underline{\hat{H}}_1(k)) = E), \quad (13)$$

and we obtain the following statement. The discussion can be further extended to a general case.

Theorem 1: Without any assumption on the reference channel record, we can upper-bound the performance of the FP test based on the Alice-Bob-Eve attack model given by (13) as:

$$P_{FA} = \alpha - \alpha(1 - P_a(1 - P_{FA}))^{N_T} \qquad (14)$$

$$P_M = \beta + (1 - \beta)(1 - P_{FA}/\alpha), \qquad (15)$$

where $P_a \in [0, 1]$ denotes the fraction of all messages from Alice.

Proof: We upper-bound the performance of FP by assuming perfect higher-layer process. By (2), if Bob has no valid reference channel record, he either has not received any messages from Alice or rejected all of Alice's messages during the past N_T time. In the i.i.d. attack model, (13), the FP decision $I_1(k)$ is independent to the sender of the next message, and thus we have

$$\Pr(\text{No } \underline{\hat{H}}_0(k))$$
$$= \prod_{n=1,\cdots,N_T} (1 - \Pr(S(\underline{\hat{H}}_1(k-n)) = A, I_1(k-n) \neq 1))$$
$$= (1 - P_a(1 - P_{FA}))^{N_T}. \qquad (16)$$

In addition, since $P_{MA} = \beta_2 = 0$, a reference channel record, $\underline{\hat{H}}_0(k)$, never results from a spoofing message:

$$\Pr(S(\underline{\hat{H}}_0(k)) = A) + \Pr(\text{No } \underline{\hat{H}}_0(k)) = 1. \qquad (17)$$

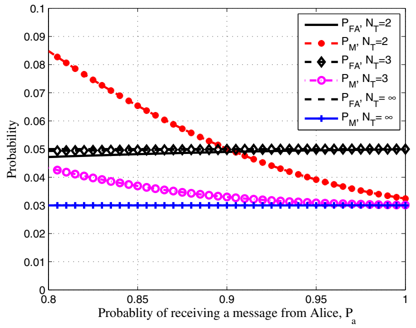

Fig. 2. Upper bound of the performance of FP in spoofing detection, including the false alarm rate P_{FA} and the miss rate P_M, given $\alpha = 0.05$, and $\beta = 0.03$, by (14) and (15).

From (16) and (17), we can rewrite the false alarm rate of FP, (6), as

$$\begin{aligned} P_{FA} =& \Pr(I_1(k) = 1 | S(\underline{\hat{H}}_1(k)) = S(\underline{\hat{H}}_0(k))) \Pr(S(\underline{\hat{H}}_0) = A) \\ &+ \Pr(I_1(k) = 1 | \text{No } \underline{\hat{H}}_0(k)) \Pr(\text{No } \underline{\hat{H}}_0(k)) \\ =& \alpha - \alpha(1 - P_a(1 - P_{FA}))^{N_T}. \end{aligned} \quad (18)$$

Similarly, by (16)-(18), the miss rate (7) can be simplified into

$$\begin{aligned} P_M =& \Pr(\text{No } \underline{\hat{H}}_0(k)) + \\ & \Pr(S(\underline{\hat{H}}_0(k)) = A) \Pr(I_1(k) \neq 1 | S(\underline{\hat{H}}_0(k)) \neq S(\underline{\hat{H}}_1(k))) \\ =& \beta + (1 - \beta)(1 - P_{FA}/\alpha). \end{aligned}$$

■

It is shown in (14) that P_{FA} is a function of α, β, P_a and N_T. Although (14) does not provide a closed-form expression of P_{FA}, it provides a convenient way to calculate P_{FA}, especially for small N_T. For example, given $N_T = 2$, P_{FA} is one of the roots of

$$\alpha P_a^2 P_{FA}^2 + (1 + 2\alpha(1 - P_a)P_a)P_{FA} + ((1 - P_a)^2 - 1)\alpha = 0,$$

which is within the range of 0 and 1. Second, this bound indicates $P_{FA} = \alpha$ and $P_M = \beta$, as N_T approaches infinity. Actually, the snapshot scenario always holds under this static channel environment. Finally, this bound implies that if Eve sends all the packets with $P_a = 0$, we have $P_{FA} = 0$ and $P_M = 1$, i.e., FP falsely accepts all spoofing messages. That is because FP only verifies whether the messages are sent by the same transmitter and does not provide user authorization. To address this problem, the higher-layer authentication can be performed to ensure the initial estimate of the channel is obtained from a validated message.

Figure 2 provides an upper bound of the performance of FP according to (14) and (15), as a function of P_a and N_T, given $\alpha = 0.05$ and $\beta = 0.03$. In Fig. 2, the performance improves with P_a, due to the error propagation property of FP, as well as the high cost of recovering from its previous wrong decisions. Moreover, the PHY-authentication scheme works efficiently when Bob receives many more messages from Alice than from Eve, e.g., $P_a > 0.8$. The performance improves with the channel record lifetime N_T, and in particular, $N_T = \infty$ (i.e., static channel) corresponds to the best performance. Actually, the performance gap between $N_T = 3$ and $N_T = \infty$ is very small, indicating that FP works well if Bob can receives about three messages on average during a channel coherence time.

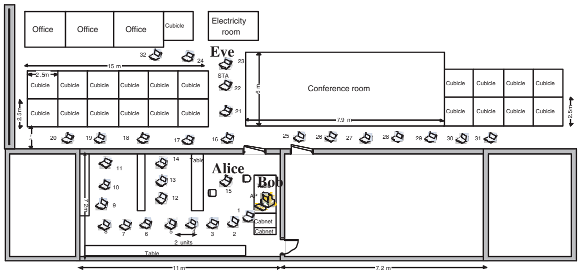

Fig. 3. The layout of the field test for the PHY-authentication, at the InterDigital office building. We placed the receiver, Bob, on a table of a room, and the transmitters at 32 different locations.

C. Lower-Bound with Nominal Higher-Layer Process

As another extreme case, the performance is lower bounded by a "nominal" higher-layer authentication with $I_2(k) \equiv 0$. In certain scenarios, e.g. where low-power lower-cost devices are highly desired, the conventional higher-layer authentication may be too expensive in terms of power, computation, or delay. As an extreme case, the FP scheme works independently, and Bob accepts any message that is not rejected by FP. We have $I_a(k) = 1$, if and only if $I_1(k) = 1$; and zero otherwise.

The error propagation property of FP is evident here, since any "non-empty" higher-layer authentication provides some degree of protection against spoofing messages missed by FP and thus reduces the error propagation. That is why this nominal higher-layer process lower-bounds the performance of FP.

Fortunately, FP can recover from a previous wrong decision by using the reference channel timers. As shown in (2), an "invalid" channel record expires, when Bob does not accept any new message during a period of N_T. Hence, if Eve keeps silent longer than N_T, Bob accepts the new message from Alice and resumes to the normal status. Hence, the PHY-authentication works independently and provides some degree of spoofing detection for "naked" wireless systems, against light spoofing attacks.

IV. Field Test and Implementation Issues in 802.11

We have verified the performance of the PHY-authentication scheme in 802.11 systems. The experiment was performed at InterDigital's King of Prussia office site, a typical indoor office environment, where Bob, located at a fixed location, continually sent probe messages to Alice and Eve, as shown in Fig. 3. Alice and Eve estimated the Bob-Alice and Bob-Eve channel vectors, respectively, based on these probe messages. By the channel reciprocity property, we had the Alice-Bob and Alice-Eve channel responses at the same time, and we then performed the "offline" FP test. Next, we computed the false alarm rate and miss rate, α and β. In the first test, Eve was located in the same room with Bob, while Alice was outside the room. Thus the Alice-Bob channel, $\underline{H}_A$, had larger

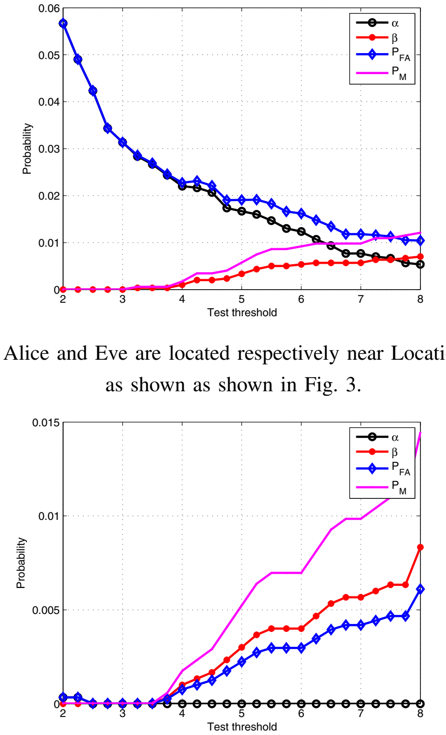

(a) Case 1: Alice and Eve are located respectively near Location 15 and 23, as shown as shown in Fig. 3.

(b) Case 2: Same as Case 1, except the locations of Alice and Eve are reversed.

Fig. 4. Performance of FP in two experiment scenarios each with three 802.11 radios, including both the snapshot performance, α and β, and the generalized performance, P_{FA} and P_M, calculated by (14) and (15), with $N_T = 2$ and $P_a = 70\%$ of the received messages coming from Alice. Note these two sub-figures use different scales in y-axis.

time variations than $\underline{H}_E$. In the second test, we exchanged the location of Alice and Eve, and thus $\underline{H}_E$ had larger channel variations.

As shown in Fig. 4, both α and β are mostly below 5% in both cases. The performance criteria, P_{FA} and P_M, are in the same range, when $N_T = 2$ and 30% of the messages received by Bob are sent by Eve. As mentioned in Fig. 2, both P_{FA} and P_M decrease with P_a. Thus Fig. 4 verifies the performance of FP for spoofing messages under 30%. Moreover, FP exhibits significantly different performance, after the locations of Alice and Eve are reversed. The false alarm rate (either α or P_{FA}) in Case 1 is larger than the counterpart in Case 2. This indicates that the performance of the FP test is more sensitive to the channel time variation of Alice than that of Eve. Figure 4 also shows that unlike α, P_{FA} does not always decrease with the test threshold η, since P_{FA} depends on both β and α. Hence a bad selection of η may lead to a large P_{FA}, as well as large P_M.

V. Conclusion & Future Work

We have devised a hybrid PHY-authentication protocol to detect spoofing attacks for wireless networks. In this scheme, the PHY-authentication algorithm, FP, is integrated with the already existing higher-layer security mechanisms. Analysis results show that our scheme requires little additional system overhead, and it can significantly reduce the workload of the higher-layer security mechanism. For example, the workload of the higher-layer process is reduced by about 26%, if 20% of all the received messages are spoofed and both the false alarm rate and the miss rate of FP are 10%. Moreover, this protocol provides spoofing protection for “naked” wireless systems, with slightly increased false alarm rate but dramatically decreased miss rate.

We have upper-bounded the performance of FP in a generalized spoofing attack pattern as a function of the “snapshot” performance that is much easier to measure, and implemented the PHY-authentication protocol on an 802.11 testbed. Results via field tests have confirmed the efficacy of our protocol with stationary terminals. For example, both the averaged miss rate and false alarm rate in the spoofing detection are mostly smaller than 5%, if 30% of all messages are spoofed.

Future work may include: (1) further quantify all the benefits that the FP scheme brings to wireless security, e.g., the computational time that FP saves for 802.11i when subjected to a series of spoofing attacks; (2) perform some “online” authentication tests using the 802.11 testbed for realistic communication scenarios; and (3) apply or modify FP in a sensor-based cognitive radio network, where the “police” sensors might exploit the received channel information to detect transmitters who pretend to be primary users.

Acknowledgment

The work of the first author was supported in part by the Natural Science Foundation of Fujian Province of China (No.2010J01347) and Tsinghua-Qualcomm joint research center. The work of Wade Trappe was supported, in part, through a grant ARO W911NF-09-1-0089.

References

[1] H. Yang, F. Ricciato, S. Lu, and L. Zhang, “Securing a wireless world,” *Proceedings of the IEEE*, vol. 94, no. 2, pp. 442–454, February 2006.

[2] J. Bellardo and S. Savage, “802.11 denial-of-service attacks: real vulnerabilities and practical solutions,” in *Proc. USENIX Security Symposium*, 2003, pp. 15–28.

[3] D. Faria and D. Cheriton, “Detecting identity-based attacks in wireless networks using signalprints,” in *Proc. ACM Workshop on Wireless Security*, 2006, pp. 43 – 52.

[4] M. Demirbas and Y. Song, “An RSSI-based scheme for sybil attack detection in wireless sensor networks,” in *Proc. IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, June 2006.

[5] Y. Chen, W. Trappe, and R. Martin, “Detecting and localizing wireless spoofing attacks,” in *Proc. Sensor, Mesh and Ad Hoc Communications and Networks*, 2007, pp. 193–202.

[6] N. Patwari and S. Kasera, “Robust location distinction using temporal link signatures,” in *Proc. ACM International Conference on Mobile Computing and Networking*, 2007, pp. 111 – 122.

[7] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe, “Fingerprints in the ether: Using the physical layer for wireless authentication,” in *Proc. IEEE International Conference on Communications (ICC)*, June 2007, pp. 4646–4651.

[8] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe, “Channel-based spoofing detection in frequency-selective rayleigh channels,” *IEEE Trans. on Wireless Communications*, vol. 8, pp. 5948–5956, Dec. 2009.

[9] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe, “Channel-based detection of sybil attacks in wireless networks,” *IEEE Transactions on Information Forensics & Security*, vol. 4, pp. 492–503, Sept. 2009.

[10] W. C. Jakes, *Microwave Mobile Communications*, chapter 7.2.1, IEEE Press, 1994.

[11] A. Goldsmith, *Wireless Communications*, chapter 3, Cambridge University Press, 2005.