

A Physical-Layer Technique to Enhance Authentication for Mobile Terminals

Liang Xiao, Larry Greenstein, Narayan Mandayam, Wade Trappe
Wireless Information Network Laboratory (WINLAB), Rutgers University
671 Rt. 1 South, North Brunswick, NJ 08902

Abstract—We propose an enhanced physical-layer authentication scheme for multi-carrier wireless systems, where transmission bursts consist of multiple frames. More specifically, it is based on the spatial variability characteristic of wireless channels, and able to work with moderate terminal mobility. For the authentication of the first frame in each data burst, the legal transmitter uses the saved channel response from the previous burst as the key for authentication of the first frame in the next burst. The key is obtained either via feedback from the receiver, or using the symmetric channel property of a TDD system. Then the authentication of the following frames in the burst is performed either by a Neyman-Pearson hypothesis test, or a least-squares adaptive channel estimator. Simulations in a typical indoor building show that the scheme based on the Neyman-Pearson test is more robust against terminal mobility, and is able to detect spoofing attacks efficiently with small system overhead when the terminal moves with a typical pedestrian speed.

I. INTRODUCTION

Wireless networks are susceptible to various attacks and threats. For example, in commodity networks, such as 802.11 networks, it is easy for a device to alter its MAC address and claim to be another device by simply issuing an `ifconfig` command. This weakness is a serious threat, and there are numerous attacks, ranging from session hijacking [1] to attacks on access control lists [2], that are facilitated by the fact that an adversarial device may masquerade as another device.

To address these challenges, many researchers have turned to using physical layer information to enhance wireless security, and the wireless channel has been explored as a form of fingerprint for wireless security. The reciprocity and rich multipath of the ultrawideband channel has been used as a means to establish encryption keys [3]. In [4], a practical scheme to discriminate between transmitters was proposed, which identifies mobile devices by tracking measurements of signal strength from multiple access points. A similar approach was considered for sensor networks in [5].

Concurrent with these efforts, the present authors have proposed a hypothesis test that exploits the spatial variability of propagation to enhance authentication [6]. This method combines channel measurement with hypothesis testing to determine whether the current and prior communication attempts are made by the same user (same channel response). The method was verified for a time-invariant channel in [6],

The authors may be reached at {lxiao, lfg, narayan, trappe}@winlab.rutgers.edu. This research is supported, in part, through a grant, CNS-0626439, from the National Science Foundation.

including the effects of receiver thermal noise; and for a time-varying channel in [7], where the terminals remain static and the variations are due to changes in the environment.

The physical-layer authentication, however, faces additional challenges as user mobility is introduced. Specifically, physical-layer authentication utilizes the differences between a measured (test) channel response and a prior channel response to discriminate between transmitters at different locations. Unfortunately, due to the rapid spatial decorrelation properties of the wireless multipath channel, even a minor movement of a mobile can lead to a quite different channel response, resulting in large false alarm rates. In this paper, we propose an enhanced scheme to solve this problem, which consists of two parts – inter-burst authentication and intra-burst authentication – and generates private keys from the channel response to relax the limit on user displacement between two bursts.

We begin in Section II by providing an overview of the authentication scheme. In Section III, we derive a Neyman-Pearson-test-based intra-burst authentication scheme. We present another practical intra-burst scheme in Section IV, based on the Least-Squares adaptive filter. In order to validate our ideas, we have performed simulations using the WiSE propagation tool [8] with a typical mobile velocity, and our results are presented in Section V. We conclude the paper in Section VI.

II. SYSTEM OVERVIEW

A. System Model

We borrow from the conventional terminology of the security community by introducing three different parties: Alice, Bob and Eve. For our purposes, these three entities may be thought of as wireless transmitters/receivers that are potentially located in spatially separated positions. Our two “legal” protagonists are the usual Alice and Bob, and for the sake of discussion throughout this paper, Alice will serve as the transmitter that initiates communication, while Bob will serve as the intended receiver. Their nefarious adversary, Eve, will serve as an active opponent who injects undesirable communications into the medium in the hopes of spoofing Alice. Our security objective, broadly speaking, is to provide authentication between Alice and Bob, despite the presence of Eve. More specifically, Bob has to differentiate between legitimate signals from Alice and illegitimate signals from Eve. For convenience, Bob is assumed to be stationary while Alice moves in any direction with a maximum velocity of v_a . However, our method is generic and our results can be easily extended to the case of mobility of all terminals.

Suppose Alice sends a signal to Bob with the frame structure shown in Fig. 1, where the whole session consists of several data bursts. Each burst has N_x frames (N_x may vary with the burst), while each frame, with M frequency subbands and duration T , consists of N_d data symbols and one pilot in each subband. (It is easy to see the compatibility of this format with the use of orthogonal frequency division multiplexing (OFDM), wherein the subbands are occupied by OFDM tones.) The number of pilots in the first symbol can, in fact, be less than the number of subbands, with the rest used for data. For concreteness, however, we assume initially that all subbands on the first symbol are used for pilots. In our numerical example later, we relax this assumption.

Bob uses the pilots for channel estimation, obtaining test vectors $\underline{H}_t(k) = [H_{t,1}(k), H_{t,2}(k), \dots, H_{t,M}(k)]^T$, where k is the frame index and the subscript t denotes “transmitter to be authenticated”. The frame duration T is assumed to be small enough to make the displacement of the transmitter (Alice) per frame much smaller than the channel decorrelation distance (i.e., $r = v_a T \ll \lambda/2$). Thus, two consecutive channel responses are highly correlated.

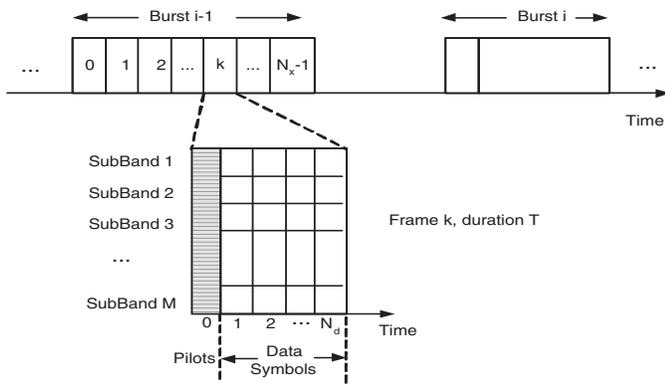


Fig. 1. Frame structure of the transmission from Alice to Bob. Each data burst consists of an arbitrary number of frames, while each frame has one pilot and N_d data symbols on each of M subbands. Frame 0 in each data burst contains the channel response value in the previous burst ($\underline{H}_A(-1)$) as a key for the inter-burst authentication. Bob uses the intra-burst authentication method in the following frames to authenticate Alice, and saves at least one frequency response as the key for the next burst.

B. Overview of the Authentication Process

As shown in [6] and [7], the rapid spatial decorrelation in a rich-scattering environment can be used to authenticate a transmitter. We enhance those earlier schemes to address the terminal mobility problem. Since terminal mobility may force self-decorrelation of Alice’s channel with respect to itself, we must employ a different strategy to bridge the gap between bursts of communications. To accomplish this, our improved process consists of two consecutive parts: an inter-burst authentication phase and an intra-burst authentication phase.

The inter-burst authentication is carried out using the first frame of each data burst to determine whether the current transmitter is still Alice. We note that at the outset of this protocol, in order for Bob to get an initial channel estimate for Alice, it may be necessary to employ a higher-layer

authentication protocol to bootstrap the association between Alice and a corresponding channel response. However, this is a one-time step, and generally the inter-burst process will focus on authenticating a subsequent data burst given that a prior data burst has been verified. We thus assume that Bob has an estimate of the Alice-Bob channel response of a particular frame in the previous data burst, which we shall denote as $\underline{H}_A(-1)$, where the subscript A corresponds to “Alice”. The time interval between two bursts may be so large that Alice has moved a significant distance. Thus the channel response of the first frame in the current burst, $\underline{H}_A(0)$, may be totally uncorrelated with $\underline{H}_A(-1)$.

To solve this problem, we assume that both Alice and Bob save at least one channel response in each data burst as the key in the authentication process for the next successive burst. Alice may obtain this $\underline{H}_A(-1)$ either by feedback from Bob, or by measurement of the reverse link pilots in a TDD system. In the first frame of each burst, Alice sends the saved $\underline{H}_A(-1)$ from the last burst to Bob. If it matches with Bob’s version, Bob will assume it is from Alice. The channel response $\underline{H}_A(-1)$ is not readily predicted by Eve. Thus she will fail the inter-burst authentication with high probability. Detailed performance analysis of how our scheme fits into a holistic cross-layer security framework is part of our ongoing research. Unless specified otherwise, we will focus on the intra-burst authentication in the remainder of the paper.

The intra-burst authentication happens within a data burst, after the first frame passes the inter-burst authentication process. For any frame index $k > 1$, Bob is assumed to obtain the Alice-Bob channel gain in the previous frame, $\underline{H}_A(k-1)$, and use the observation of the current channel gain, $\underline{H}_t(k)$, to determine whether the current transmitter is still Alice. In the null hypothesis, \mathcal{H}_0 , the claimant is Alice. Otherwise, in the alternative hypothesis, \mathcal{H}_1 , the claimant terminal is someone else. We use the notation $\tilde{\cdot}$ to denote accurate values without measurement error, and thus have

$$\mathcal{H}_0 : \tilde{\underline{H}}_t(k) = \tilde{\underline{H}}_A(k) \quad (1)$$

$$\mathcal{H}_1 : \tilde{\underline{H}}_t(k) \neq \tilde{\underline{H}}_A(k). \quad (2)$$

III. NEYMAN-PEARSON INTRA-BURST TEST

We now present an intra-burst authentication scheme based on the Neyman-Pearson (NP) test to choose between hypotheses (1) and (2). In particular, Bob uses channel estimates in two consecutive frames, $\underline{H}_A(k-1)$ and $\underline{H}_t(k)$, to determine whether they are from the same transmitter (Alice) or not. We first derive the NP test for an ideal case and then propose a practical method based on it.

A. NP Test for an Idealized Case

To gain insight, we begin with an idealized case wherein the set of channel response values form a Gaussian random vector. During frame k , Alice moves in an arbitrary direction from her location in the previous frame, with a maximum distance of $r = v_a T$. With $r \ll \lambda/2$, we can safely assume that $\underline{H}_A(k)$ is highly correlated with $\underline{H}_A(k-1)$.

For illustrative purposes, we use an autoregressive model of order 1 (AR-1) to characterize the temporal process of channel

response $\tilde{H}_A(k)$:

$$\tilde{H}_A(k) = \rho \tilde{H}_A(k-1) + \sqrt{(1-\rho^2)\sigma_A^2} \underline{\epsilon}(k), \quad (3)$$

where the AR coefficient ρ denotes the similarity of the channel responses in consecutive frames; the noise in the AR-1 model, $\underline{\epsilon}(k) \sim CN(0, \mathbf{I})$, is independent of $\tilde{H}_A(k-1)$; \mathbf{I} is an $M \times M$ identity matrix; and σ_A^2 is the variance over space of \tilde{H}_A .

Now considering random phase drift of the local oscillator (LO) and the additive thermal noise at the receiver, we model the measured channel gain to Alice as

$$\begin{aligned} H_A(k) &= \tilde{H}_A(k) e^{j\phi(k)} + \underline{N} \\ &\sim CN(\rho \tilde{H}_A(k-1) e^{j\phi(k)}, \sigma_0^2 \mathbf{I}), \end{aligned} \quad (4)$$

where we have added white thermal noise $\underline{N} \sim CN(0, \sigma_N^2 \mathbf{I})$; $\sigma_0^2 = \sigma_N^2 + \sigma_A^2$; and $\phi(k) \in [0, 2\pi)$ represents measurement errors in the phase of the channel response, considering the fact that the phase of Bob's receiver LO can change between one measurement and another. Since $r \ll \lambda/2$, we henceforth approximate ρ as 1.

Without *a priori* location information, at frame k , Eve is assumed to be randomly and uniformly distributed over the whole area of interest (e.g. a building). Since Eve is very likely to be far from Alice's previous location, her channel gain to Bob, $H_E(k)$, is independent of $H_A(k-1)$, where the subscript E denotes "Eve". Thus we model it as

$$H_E(k) \sim CN(0, \sigma_1^2 \mathbf{I}), \quad (5)$$

where the channel variance $\sigma_1^2 = \sigma_N^2 + \sigma_E^2$, and $\sigma_E^2 \gg \sigma_A^2$ is the channel variance due to the location uncertainty of Eve.

Considering (4) and (5), we build the corresponding log-likelihood ratio rule:

$$\begin{aligned} \ln \frac{P(\underline{H}_t(k) | \mathcal{H}_1)}{P(\underline{H}_t(k) | \mathcal{H}_0)} &= \ln \frac{P(\underline{H}_E(k))}{P(\underline{H}_A(k))} \\ &= \sum_{i=1}^M \frac{|H_{t,i}(k) - H_{A,i}(k-1) e^{j\phi(k)}|^2}{2\sigma_0^2} \\ &\quad - \sum_{i=1}^M \frac{|H_{t,i}(k)|^2}{2\sigma_1^2} + M \ln \frac{\sigma_0}{\sigma_1} \gtrsim_{\mathcal{H}_0}^{\mathcal{H}_1} J_1, \end{aligned} \quad (6)$$

where $P(\cdot)$ denotes a probability density function and J_1 is a suitably chosen decision threshold. Considering that $\sigma_1 \gg \sigma_0$, we can simplify as follows:

$$\|\underline{H}_t(k) - \underline{H}_A(k-1) e^{j\phi(k)}\|^2 \gtrsim_{\mathcal{H}_0}^{\mathcal{H}_1} J_2, \quad (7)$$

where J_2 is another threshold and $\|V\|$ is the norm of vector V . For the convenience of threshold determination, we normalize the test statistic via

$$\Lambda(\underline{H}_t(k)) = \frac{\|\underline{H}_t(k) - \underline{H}_A(k-1) e^{j\phi(k)}\|^2}{\sigma_0^2} \gtrsim_{\mathcal{H}_0}^{\mathcal{H}_1} \eta. \quad (8)$$

To do the Neyman-Pearson test, the threshold η is chosen to satisfy a constraint on the false alarm rate α , i.e.,

$$\alpha = \int_{\eta}^{\infty} p(\Lambda(\underline{H}_t(k)) | \mathcal{H}_0) d\Lambda. \quad (9)$$

Under \mathcal{H}_0 , we have $(\underline{H}_t(k) - \underline{H}_A(k-1) e^{j\phi(k)}) \sim CN(0, \sigma_0^2 \mathbf{I})$, and thus Λ is chi-square distributed with $2M$ degree of freedom, i.e., $\Lambda(\underline{H}_t(k) | \mathcal{H}_0) = \chi_{2M}^2$. Denoting the CDF of a random variable X as $F_X(\cdot)$, the test threshold given α is given by

$$\eta = F_{\chi_{2M}^2}^{-1}(1 - \alpha). \quad (10)$$

Similarly, the test statistic under \mathcal{H}_1 is $\Lambda(\underline{H}_t(k) | \mathcal{H}_1) = \sigma_1^2 \chi_{2M}^2 / \sigma_0^2$, and thus the miss rate can be written as

$$\beta = p(\Lambda < \eta | \mathcal{H}_1) = F_{\chi_{2M}^2}(\sigma_0^2 F_{\chi_{2M}^2}^{-1}(1 - \alpha) / \sigma_1^2), \quad (11)$$

which rises with $\sigma_0^2 / \sigma_1^2 = (\sigma_N^2 + \sigma_A^2) / (\sigma_E^2 + \sigma_N^2)$. Since $\sigma_E^2 > \sigma_A^2$, we can easily see that the miss rate β for given α rises with σ_N^2 ; and the smaller σ_A is, the greater is the rise of β . It means that the system performance degrades with thermal noise, and this degradation is more distinct as Alice moves slower.

B. A Practical Method

In reality, the parameters σ_0 , σ_1 , and phase rotation $\phi(k)$, used in the test (10), are unknown. Therefore, instead of using σ_0^2 , we normalize the test statistic with a known parameter, $\|\underline{H}_A(k-1)\|^2$, i.e.,

$$\Lambda_0 = \frac{\|\underline{H}_t(k) - \underline{H}_A(k-1) e^{j\phi(k)}\|^2}{\|\underline{H}_A(k-1)\|^2} \gtrsim_{\mathcal{H}_0}^{\mathcal{H}_1} \eta. \quad (12)$$

Moreover, considering that $\phi(k)$ is also unknown, we modify Λ_0 into the following form:

$$\begin{aligned} \Lambda_1 &= \min_{\varphi} \frac{\|\underline{H}_t(k) - \underline{H}_A(k-1) e^{j\varphi}\|^2}{\|\underline{H}_A(k-1)\|^2} \\ &= \frac{\|\underline{H}_t(k) - \underline{H}_A(k-1) e^{j\varphi^*}\|^2}{\|\underline{H}_A(k-1)\|^2} \gtrsim_{\mathcal{H}_0}^{\mathcal{H}_1} \eta, \end{aligned} \quad (13)$$

where $\varphi^* = \text{Arg}(\underline{H}_A(k-1)^H \underline{H}_t(k))$. The new test statistic Λ_1 is a practical one, totally based on measured frequency samples at M subbands in consecutive time, $\underline{H}_A(k-1)$ and $\underline{H}_t(k)$. It represents their difference in both power (i.e., the distance effect) and shape (i.e., the multipath effect).

The test threshold η of Λ_1 has no closed-form expression and has to be determined by simulations, as we show later. For given threshold η , the false alarm rate α and the miss rate β are given by,

$$\alpha(\eta) = P[\Lambda_1 > \eta | \mathcal{H}_0] \quad (14)$$

$$\beta(\eta) = P[\Lambda_1 < \eta | \mathcal{H}_1]. \quad (15)$$

IV. LEAST-SQUARES ADAPTIVE FILTER

We now explore an alternative method for the intra-burst authentication, where M sets of linear least-squares adaptive filters are used independently to estimate the channel response for the M subbands. For the convenience of notion, we focus on the m -th subband, and ignore the frequency index m unless necessary.

The estimated channel response at time k , which is the output of the m -th adaptive linear filter with order L , can be written as

$$y(k) = \sum_{l=0}^{L-1} w_l^* u(k-l), \quad (16)$$

where $u(k)$ is the input of the adaptive filter at time k , and w_l is the l -th tap weight of the filter, which can be determined using various adaptive algorithms, like the recursive least-squares (RLS) algorithm [9].

If it is Alice transmitting during the time interval $[(k-L)T, kT]$, the filter inputs are $H_A(k-L), \dots, H_A(k-1)$, and the estimation error is $e(k) = H_A(k) - y(k)$. Because of the strong correlation of the inputs $H_A(k-L), \dots, H_A(k)$, the ensemble-averaged squared error of the channel estimation filter is usually quite small.

If on the other hand, Eve comes in at time k , due to the spatial variability of the channel response, the estimation error, $e(k) = H_E(k) - \sum_{l=0}^{L-1} w_l^* H_A(k-l-1)$, is very likely to jump to a much larger value.

Therefore, we build another test statistic Λ_2 , using M parallel adaptive channel estimators. The null hypothesis \mathcal{H}_0 is accepted if the normalized squared sum of estimation error from these filters is less than a certain threshold η ; otherwise, the alternative hypothesis is chosen. Thus

$$\Lambda_2 = \frac{\sum_{m=1}^M |e_m(k)|^2}{\sum_{l=0}^{L-1} \sum_{m=1}^M |u_m(k-l)|^2 / L} \underset{\mathcal{H}_0}{\geq} \eta. \quad (17)$$

We normalize the estimation error to make η easier to determine. It does not have a closed-form expression but can be obtained through simulations.

Note that this test can be carried out only after the successful authentication of at least L frames, and even though the RLS algorithm converges fast, it still takes approximately $2L$ frames [9]. Since we have to take data after the algorithm converges, we usually choose $k > 3L$ in Eq. (17). Thus Λ_2 has larger system overhead ($3L$ frames) than Λ_1 (1 frame), as well as greater implementation complexity.

The use of RLS estimators in this context may not be practical or cost-effective, but the results we will present for this case are instructive. They will show that, even under the most favorable assumptions (RLS estimation), using least-squares adaptive filtering is not measurably superior to using the simpler NP test.

V. SIMULATIONS AND NUMERICAL RESULTS

A. Simulation Method

In order to test the proposed scheme, it is necessary to model “typical” channel responses and to capture the spatial variability of these responses. To that end, we make use of the Wireless System Engineering (WiSE) tool, a ray-tracing software package developed by Bell Laboratories [8]. One input to WiSE is the 3-dimensional plan of a specific building, including walls, floors, ceilings and their material properties. With this information, WiSE can predict the rays at any receiver from any transmitter, including their amplitudes, phases and delays. From this, it is straightforward to predict the transmit-receive path’s complex gain at any specified frequency, and we assume that M such gains are measured via pilots in the subbands.

We assume the total transmit power for the set of M pilots is P_T , i.e., P_T/M mW per pilot tone. For convenience only, we normalize all received tone power by P_T/M , so our

metrics deal essentially with path gains, as the above equations implicitly assume. The receiver noise power per pilot tone is $P_N = \kappa T N_F b$, where κT is the thermal noise density in mW/Hz; N_F is the receiver noise figure; and b is the noise bandwidth per tone in Hz. Given the normalization of received tone power by P_T/M , the noise power per tone is the dimensionless quantity

$$\sigma_N^2 = \frac{\kappa T N_F b}{P_T/M}. \quad (18)$$

We consider one particular office building, for which a top view of the first floor is shown in Fig. 2¹. This floor of this building is 120 meters long, 14 meters wide and 4 meters high. We consider the mobility of the legal transmitter Alice, and multiple possible positions of Eve. For our experiment, we randomly uniformly select N_A in-building locations for Alice, each corresponding to her position at the start (i.e., in Frame 0) of one of N_A data bursts. For each such location, we consider a set of N_E possible locations for Eve, which are also randomly uniformly selected. We assume that each burst has the same number of frames, N_x . Alice moves r mm per frame in arbitrary directions, and an arbitrary distance between neighboring data bursts. For each transmit-receive path, we use WiSE to generate the accurate channel gain \hat{H} ; and then generate N_n measured channel gain vectors, $\underline{H} = \hat{H} + \underline{H}_n$, based on independent vectors of additive white thermal noise, $\underline{N} \sim CN(\underline{0}, \sigma_N^2 \mathbf{I})$.

For each r , we collect $N_A(N_x - 1)N_n$ samples to calculate the false alarm rate α of Λ_1 , and $N_A N_E N_n$ samples for the miss rate β , for given threshold η . For the case of Λ_2 , we use $N_A(N_x - 3L)N_n$ and $N_A N_E N_n$ samples, respectively, to calculate α and β .

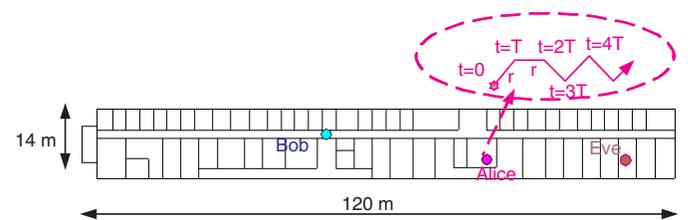


Fig. 2. System topology assumed in the simulations. The receiver, Bob, is fixed at a location within the hall way of a 120 m \times 14 m \times 4 m office building. We randomly uniformly select N_A locations for Alice inside the building, representing her positions at the start of each of N_A data bursts. For each of these, we consider a set of N_E positions for Eve, which are also randomly uniformly selected. Each burst has the same number of frames, N_x , and Alice moves a distance of r from frame to frame, in an arbitrary direction. The independence among her N_A selected starting locations means that her position is independent from one burst to another.

B. Simulation Results

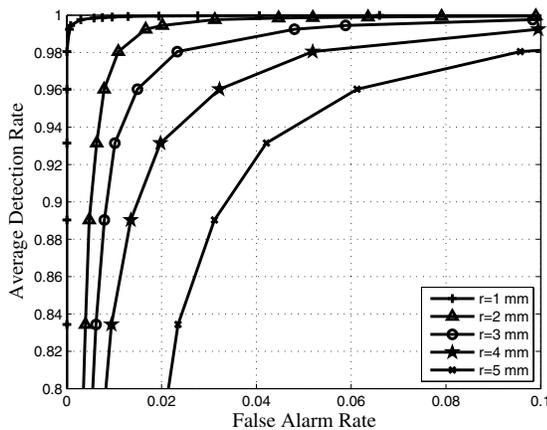
We assume $P_T = 10$ mW, $N_F = 10$ (10 dB noise figure), $\kappa T = 10^{-17.4}$ mW/Hz, $b = 0.25$ MHz, $M = 3^2$, $N_A = 50$, $N_E = 1000$, $N_n = 5$, and $N_x = 100$. The center frequencies

¹As in [6], this is the Bell Labs building at Crawford Hill in Holmdel, NJ.

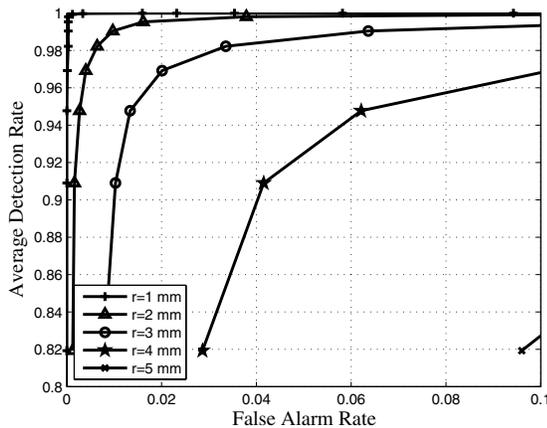
²Here we depart from our initial assumption that the number of pilots used to measure the channel is equal to the number of subbands in the signal format. Previous studies [6], [7] have shown that only a few measurements say, 3-10, are needed; in an OFDM format, however, the number of subbands (tones) is generally much larger

of the subbands are at 4.75, 5.0, and 5.25 GHz. The per tone signal-to-noise ratio (SNR) in the channel estimation ranges from 1.7 dB to 69 dB, with a median value of 30 dB. To implement the Λ_2 test, we use the RLS algorithm [9], with the filter order $L = 2$, forgetting factor $\lambda = 0.9995$, and regularization parameter $\delta = 10^{-10}$.

Figure 3 presents the receiver operating characteristic (ROC) curves of the intra-burst authentication method, i.e., the detection rate, $1 - \beta$, as a function of the miss detection rate, α , for the NP-based statistic Λ_1 and the adaptive filter based statistic Λ_2 , with Alice displacement per frame $r \in \{1, 2, 3, 4, 5\}$ mm. This corresponds to the frame duration $T \in \{0.70, 1.4, 2.1, 2.8, 3.5\}$ ms given a typical pedestrian velocity $v_a = 1.43$ m/s.



(a) NP-based statistic, Λ_1 .



(b) Adaptive filter based statistic, Λ_2 .

Fig. 3. Receiver operating characteristic (ROC) curves of the intra-burst authentication method, i.e., the average detection rate, $P_D = 1 - \beta$, as a function of false alarm rate, α , with Alice's displacement per frame $r \in \{1, 2, 3, 4, 5\}$ mm in arbitrary directions, and Eve randomly placed in the building with topology shown in Fig. 2.

It is shown in Fig. 3 that both Λ_1 and Λ_2 have good authentication performance, given that $r \leq 2$ mm. For example, Λ_1 and Λ_2 result in detection rate greater than 0.98 and 0.99, respectively, with $\alpha = 0.01$, $r \leq 2$ mm and $\eta = 0.1$. The performance degrades as Alice moves faster, since it leads to smaller correlation between successive channel realizations

of Alice's channel to Bob. In addition, although Λ_2 has better performance under smaller terminal velocity (e.g., $r \leq 2$ mm), Λ_1 is more robust against terminal mobility. For instance, the detection rates of Λ_1 and Λ_2 are around 0.96 and below 0.8, respectively, given false alarm rate of 0.06, transmitter speed of 1.43 m/s, and frame duration of 3.5 ms. Considering that Λ_2 has larger system overhead than Λ_1 , we believe Λ_1 is a better statistic to use than Λ_2 .

VI. CONCLUSION

We have proposed an enhanced physical layer technique to authenticate mobile transmitters in a wireless in-building environment. It utilizes the channel responses as keys to discriminate between a legitimate user and a would-be intruder. To address the terminal mobility problem, the authentication process is divided into two parts: the inter-burst authentication uses the channel response in the previous burst as a key for the first frame, solving the problem of possibly long intervals between bursts. The intra-burst authentication, on the other hand, compares the channel response in two consecutive frames via either of two practical methods: one is based on the Neyman-Pearson test; and the other uses adaptive filters. The NP-based method is more robust against terminal mobility, and more efficient in terms of system overhead and implementation complexity. Simulation results using the ray-tracing software WiSE show that the proposed scheme can detect spoofing attacks efficiently under slow terminal velocity. For instance, the detection rate is around 0.96, given a false alarm rate of 0.06, when the transmitter moves at a speed of 1.43 m/s and the frame duration equals to 3.5 ms.

In our ongoing research, we are working to integrate physical layer authentication into a holistic cross-layer framework for wireless security. We aim to quantify the net benefit in thus augmenting traditional "higher-layer" network security mechanisms with physical layer methods.

REFERENCES

- [1] A. Mishra and W. A. Arbaugh, "An initial security analysis of the IEEE 802.1x standard," Tech. Rep. CS-TR-4328, University of Maryland, College Park, 2002.
- [2] A. Mishra, M. Shin, and W. A. Arbaugh, "Your 802.11 network has no clothes," *IEEE Communications Magazine*, pp. 44 – 51, 2002.
- [3] R. Wilson, D. Tse, and R. Scholtz, "Channel identification: Secret sharing using reciprocity in UWB channels," *IEEE Transactions on Information Theory*, pp. 364–375, September 2007.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] L. Xiao, L. Greenstein, N. Mandayam, and W. Trappe, "Using the physical layer for wireless authentication in time-variant channels," *IEEE Trans. on Communications*, to appear.
- [8] S. J. Fortune, D. H. Gay, B. W. Kernighan, O. Landron, M. H. Wright, and R. A. Valenzuela, "WiSE design of indoor wireless systems: Practical computation and optimization," *IEEE Computational Science and Engineering*, March 1995.
- [9] S. Haykin, *Adaptive Filter Theory*, Prentice Hall, Englewood Cliffs, 1986.