Artificial Noise and Physical Layer Authentication: MISO Regime

Jake Bailey Perazzone†, Paul L. Yu*, Brian M. Sadler*, Rick S. Blum†

***Abstract*—We apply artificial noise to the fingerprint embedding authentication framework to improve information-theoretic authentication for the MISO channel. Instead of optimizing for secrecy capacity, we examine the trade-off between message rate, authentication, and key security. In this case, key security aims to limit an adversary's ability to obtain the key using a maximum likelihood decoder.**

I. Introduction

In the seminal work on the wire-tap channel [1], information-theoretic secrecy is made possible by assuming that the channel between two legitimate parties is less noisy than the channel between the transmitter and an adversary. The maximum achievable rate at which data can be sent while maintaining secrecy is known as secrecy capacity. Secrecy in this case is defined as limiting the mutual information between the message and the adversary's channel output to ϵ which goes to 0 as the block-length n goes to infinity. Since a less noisy channel is hard to guarantee in practical scenarios, much work has been done to determine and characterize other ways in which an advantage over an adversary can be utilized to facilitate information-theoretic secrecy.

In traditional (non-secret) communications, it is well documented that transitioning from single-input, single-output (SISO) to multiple-input, multiple-output (MIMO) systems leads to a dramatic increase in capacity. Likewise, secrecy capacity benefits from introducing multiple antennas at each terminal. The first work in this area came in the form of a practical achievable scheme that uses knowledge of the fading coefficients of a multiple-input and single-output (MISO) channel to create an advantage over the adversary. In that work, artificial noise (AN) is added to the transmitted signal that is designed such that it only affects the adversary and not the legitimate receiver [2]. The work was then extended to MIMO channels in [3] where it was shown that secrecy capacity can be guaranteed even if the adversary does not experience noise from the channel. A analytical approach to AN and MISO/MIMO secrecy capacity can be found in [4] and [5]. The use of artificial noise in the MISO case achieves secrecy capacity, but is suboptimal in the MIMO case.

† J.B. Perazzone and R.S. Blum are with Lehigh University, Bethlehem, PA. {jbp215,rblum}@lehigh.edu

* P.L. Yu and B.M. Sadler are with the US Army Research Lab, Adelphi, MD. {paul.l.yu.civ,brian.m.sadler6}@mail.mil

This material is based upon work partially supported by the U. S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-17-1-0331 and by the National Science Foundation under grants ECCS-1744129 and CNS-1702555.

Another approach to obtain non-zero secrecy capacity in MISO communications is artificial fast fading (AFF) [6]. It employs a precoding technique that is designed to prevent adversaries from being able to perform accurate channel estimation by inducing fast fading through randomized antenna gains for each symbol. The gains are designed such that the legitimate receiver experiences no fading and thus does not require pilot symbols. This forces the adversary to resort to blind channel estimation which does not perform well in fast fading-like environments. Artificial fast fading and artificial noise are compared in [7] where it is determined that AN outperforms AFF in low SNR regimes in terms of secrecy capacity. Since our work focuses on detecting a low power signal with low SNR, only AN is considered due to its superior performance and more manageable power allocation compared to AFF.

In this paper, we explore the utilization of artificial noise in the fingerprint embedding authentication framework [8] and analyze its ability to enhance security. Security is measured by our ability to limit the attack success probability of a computationally unlimited adversary [9]. In an analysis of authentication over noisy channels, it was indirectly shown that authentication security performance is closely related to the secrecy capacity of the channel [10]. Therefore, the increase in secrecy capacity afforded by artificial noise should increase authentication security performance. In [3], the goal is to maximize secrecy capacity by finding the optimal power allocation between the AN and the information bearing signal assuming a peak power constraint. In this case, they found the optimal power allocation using an exhaustive search over a discretized space. In the fingerprint embedding framework, however, maximizing the secrecy capacity is not necessarily the goal since a slightly different trade-off is present.

The difference in trade-off arises from the fact that we do not consider a secrecy constraint for the message. Instead, the authenticating tag and its corresponding key are to be kept secret, but since the legitimate receiver is tasked with detecting the tag, rather than decoding it, there is no associated rate. Therefore, secrecy capacity does not directly apply to our performance metrics. Alternatively, the interactions between message rate, authentication detection probability, and attack success probability are considered. This trio of metrics do not have an obvious objective function to maximize power allocations over, so we will instead demonstrate how each metric is affected by the system parameters with a focus on first achieving the desired authentication performance and then tuning the AN power allocation.

II. MISO Authentication with Artificial Noise

Authentication via physical layer fingerprinting is a framework that combines the process of hash-based message authentication codes (HMAC) [11] with physical layer security concepts and analysis. The framework, detailed in [8], is designed to utilize inherent noise in the wireless communication channel to protect an authentication tag created via HMAC methods to attain some degree of information-theoretic security. The effect of noise is strengthened by superimposing the tag on the message waveform at low power so that the tag SNR is kept low for an adversarial eavesdropper. This section summarizes the authentication process, including the adversary, in the MISO regime with the addition of artificial noise. The performance of the legitimate receiver and the adversary will also be detailed using an improved security metric that measures an adversary's ability to successfully impersonate a legitimate transmitter, as presented in [9].

In our model, we assume a MISO regime in which the legitimate transmitter, Alice, possesses N_T antennas while the legitimate receiver, Bob, and the adversary, Eve, both have only $N_R = N_E = 1$ antenna. We extend the analysis to $N_E, N_R > 1$ in future work. We also assume that Alice has full knowledge of the channel state information (CSI) h to Bob, but not the CSI g to Eve whereas Eve knows both h and g. For each transmission, the channel matrices h and g are both considered to be deterministic $(1 \times N_T)$ complex vectors and are constant throughout the block. To determine overall performance over many blocks, we will later assume that the channel matrices are composed of circularly symmetric i.i.d. complex Gaussian vectors.

A. Legitimate Transmitter Procedure

Alice's goal is to give Bob the ability to authenticate their communications while Eve's goal is to impersonate Alice by causing Bob to accept her messages as if they were from Alice. In order to facilitate authentication, before communication begins, Alice and Bob select and share a κ-bit key $\boldsymbol{k}$ drawn from a uniform distribution on $\mathcal{K}$ that is kept secret from Eve. Then, to send an authentic message $\boldsymbol{s}$ to Bob, Alice generates a tag $\boldsymbol{t} = f(\boldsymbol{s}, \boldsymbol{k})$ using the HMAC protocol [11]. The tag generating function $f(\cdot,\cdot)$ is assumed to be deterministic, but where outputs were selected uniformly over the tag space $\mathcal{T}$. A hash function is typically used as a good approximation of such a function. Both $\boldsymbol{s}$ and $\boldsymbol{t}$ are $(1 \times L)$ i.i.d. zero-mean complex vectors with unit variance where symbols are in the form of a desired modulation scheme, e.g. QPSK, QAM.

Next, Alice generates artificial noise that is designed to be orthogonal to h, but not g, so that the noise is canceled out for Bob but not for Eve. To do so, she determines the null space matrix Z of h and then generates an $((N_T-1)\times L)$ i.i.d. zero-mean complex Gaussian vector $\boldsymbol{w}$ with variance σ_w^2. Thus, $Z\boldsymbol{w}$ is also Gaussian distributed and orthogonal to h. Finally, Alice performs optimal precoding $h^\dagger/\|h\|$ and prepares to send

$$\boldsymbol{x} = \frac{h^\dagger}{\|h\|}(p_s\boldsymbol{s} + p_t\boldsymbol{t}) + Z\boldsymbol{w}\,, \tag{1}$$

where the power allocations are selected such that $E[\boldsymbol{x}^\dagger\boldsymbol{x}] \leq P_0 = p_s^2 + p_t^2 + \sigma_{\text{AN}}^2$, where $\sigma_{\text{AN}}^2 = (N_T-1)\sigma_w^2$.

B. Legitimate Receiver Procedure

Since $hZ = 0$, Bob receives

$$\begin{aligned}\boldsymbol{y} &= h\boldsymbol{x} + n_{\text{b}} &(2)\\ &= h(\frac{h^\dagger}{\|h\|}(p_s\boldsymbol{s} + p_t\boldsymbol{t}) + Z\boldsymbol{w}) + n_{\text{b}} &(3)\\ &= \|h\|(p_s\boldsymbol{s} + p_t\boldsymbol{t}) + n_{\text{b}}\,, &(4)\end{aligned}$$

where n_{b} is an L-length i.i.d. complex Gaussian noise vector with zero mean and variance σ_{b}^2.

Having received $\boldsymbol{y}$, Bob first decodes the message $\boldsymbol{s}$ as $\hat{\boldsymbol{s}}$ by treating the tag as noise due to its low power allocation. Assuming the message is decoded without error ($\hat{\boldsymbol{s}} = \boldsymbol{s}$), he then removes its contribution from $\boldsymbol{y}$ and normalizes the channel by dividing by $\|h\|p_t$ to obtain the tag estimate

$$\hat{\boldsymbol{t}}_{\text{b}} = \boldsymbol{t} + \frac{1}{\|h\|p_t}n_{\text{b}}\,. \tag{5}$$

Since Bob has access to the shared key $\boldsymbol{k}$, he can generate what he expects the correct tag to be by using the same process as Alice, i.e. he computes the expected tag $\tilde{\boldsymbol{t}} = g(\hat{\boldsymbol{s}}, \boldsymbol{k})$ using the decoded message. If $\boldsymbol{s}$ is decoded in error, then $\tilde{\boldsymbol{t}} \neq \hat{\boldsymbol{t}}_{\text{b}}$ and authentication will fail with high probability.

In order to determine the presence of $\tilde{\boldsymbol{t}}$ in $\hat{\boldsymbol{t}}_{\text{b}}$, Bob designs a hypothesis test in the classic Neyman-Pearson approach where the hypotheses are

$$\begin{cases} H_0 & \text{Invalid tag was sent} \\ H_1 & \text{Valid tag } \boldsymbol{t} \text{ was sent}\,. \end{cases} \tag{6}$$

H_0 corresponds to Eve superimposing a random tag and H_1 corresponds to Alice superimposing the expected valid tag. The optimal test for a given false alarm probability α is to compare the output of matched filter, tuned to the expected tag, to a designed threshold. Thus, Bob compares

$$\tau_{\text{b}} \triangleq \Re(\tilde{\boldsymbol{t}}\hat{\boldsymbol{t}}_{\text{b}}^\dagger) \tag{7}$$

to a threshold in order to determine the authenticity of the received message, where $\Re(\cdot)$ is the real part of its argument. Note that τ_{b} is the inner product since $\tilde{\boldsymbol{t}}$ and $\hat{\boldsymbol{t}}_{\text{b}}$ are row vectors. To properly design the threshold, Bob must establish the distribution of τ_{b} under both hypotheses.

With L sufficiently large, we use the central limit theorem to approximate the distribution as Gaussian. The mean and variance of the output of τ_{b} tuned to a random incorrect tag ($\tilde{\boldsymbol{t}} \neq \hat{\boldsymbol{t}}_{\text{b}}$) is

$$\begin{aligned}E[\tau_{\text{b}}|\mathcal{H}_0] &= \mu_0 = 0 &(8)\\ \text{var}(\tau_{\text{b}}|\mathcal{H}_0) &= \sigma_0^2 = \frac{L}{2}\left(1 + \frac{\sigma_{\text{b}}^2}{\|h\|^2p_t^2}\right), &(9)\end{aligned}$$

while for the correct tag it is

$$\begin{aligned}E[\tau_{\text{b}}|\mathcal{H}_1] &= \mu_1 = L &(10)\\ \text{var}(\tau_{\text{b}}|\mathcal{H}_1) &= \sigma_1^2 = \frac{L}{2}\left(\frac{\sigma_{\text{b}}^2}{\|h\|^2p_t^2}\right). &(11)\end{aligned}$$

Then, a threshold τ_0 designed to limit the probability of accepting an incorrect tag to α is computed as

$$\tau_0 = \Phi^{-1}(1-\alpha)\sigma_0\,, \tag{12}$$

where $\Phi^{-1}(\cdot)$ is the inverse standard normal CDF.

Finally, Bob's detection probability, given $\|h\|$, is

$$P_D(\|h\|) = 1 - \Phi\left(\frac{\tau_0 - \mu_1}{\sigma_1}\right)\,. \tag{13}$$

C. Eavesdropper Procedure

Following the approach in [9], we define the security of the authentication framework as the probability that an adversary successfully impersonates the legitimate transmitter[1]. The metric is dependent on the adversary's ability, or inability, to obtain the shared key used for authentication, so we measure security by the probability that the adversary can obtain the key using a maximum likelihood (ML) decoder. This section details Eve's procedure and performance using a bank of matched filters tuned to each possible tag as her ML decoder.

Due to the construction of the artificial noise, it is highly likely that $gZ \neq 0$, and therefore Eve receives

$$\boldsymbol{z} = g\frac{h^\dagger}{\|h\|}(p_s\boldsymbol{s} + p_t\boldsymbol{t}) + gZ\boldsymbol{w} + n_{\mathrm{e}}\,. \tag{14}$$

Since she knows h and g, she performs a similar procedure to Bob by first decoding $\boldsymbol{s}$ as $\hat{\boldsymbol{s}}_{\mathrm{e}}$ and then removing its contribution to $\boldsymbol{z}$. Eve does not perform channel normalization[2] and obtains a very noisy version of the tag

$$\tilde{\boldsymbol{t}}_{\mathrm{e}} = g\frac{h^\dagger}{\|h\|}p_t\boldsymbol{t} + gZ\boldsymbol{w} + n_{\mathrm{e}}\,. \tag{15}$$

In order to increase the probability of obtaining the correct key, Eve collects many observations and performs a joint test. We assume that she follows the procedure in [9] in which she treats each observation as a continuation of a long tag from one key. She does this by concatenating each observed noisy tag $\tilde{\boldsymbol{t}}_{\mathrm{e}}$ and tunes each matched filter to the concatenations of each possible expected tag $\hat{\boldsymbol{t}}_i$ produced by a given key and each message $\hat{\boldsymbol{s}}_{\mathrm{e}}$. Therefore, we have the possible expected tags

$$\hat{\boldsymbol{t}}_i = g\frac{h^\dagger}{\|h\|}p_t\left(f(\hat{\boldsymbol{s}}_{\mathrm{e}}^1, \boldsymbol{k}_i)||\cdots||f(\hat{\boldsymbol{s}}_{\mathrm{e}}^{N_o}, \boldsymbol{k}_i)\right)\ \forall \boldsymbol{k}_i \in \mathcal{K}, \tag{16}$$

and the observed tags

$$\tilde{t}_{\mathrm{e}} = \tilde{t}_{\mathrm{e}}^1||\cdots||\tilde{t}_{\mathrm{e}}^{N_o}\,, \tag{17}$$

where we assume Eve has access to the tag generating function $f(\cdot,\cdot)$ and N_o is the number of observations.

In trying to determine the correct key that was used, Eve faces a $|\mathcal{K}|$-ary hypothesis test. The optimal test is to choose the key with the associated tag that has the highest output in a bank of matched filters tuned to each hypothesis, that is each $\hat{\boldsymbol{t}}_i$ computed in (16). Thus, her test is

$$\arg\max_{k_i\in\mathcal{K}} \tau_{\mathrm{e}}(k_i) = \arg\max_{k_i\in\mathcal{K}} \Re(\hat{\boldsymbol{t}}_i\tilde{\boldsymbol{t}}_{\mathrm{e}}^\dagger). \tag{18}$$

Once again, we use the central limit theorem, with sufficiently large L, to determine the distribution of the matched filter outputs. The mean and variance of a matched filter for a single observation tuned to a random incorrect tag, i.e. $\hat{\boldsymbol{t}}_i \neq \boldsymbol{t}$, is

$$E[\tau_{\mathrm{e}}|\mathcal{H}_0] = \mu_{0,\mathrm{e}} = 0 \tag{19}$$

$$\mathrm{var}(\tau_{\mathrm{e}}|\mathcal{H}_0) = \sigma_{0,\mathrm{e}}^2 = \frac{L}{2}\left|g\frac{h^\dagger}{\|h\|}\right|^2 p_t^2\left(p_t^2 + \|gZ\|^2\sigma_w^2 + \sigma_{\mathrm{e}}^2\right), \tag{20}$$

while for the correct tag it is

$$E[\tau_{\mathrm{e}}|\mathcal{H}_1] = \mu_{1,\mathrm{e}} = L\left|g\frac{h^\dagger}{\|h\|}\right|^2 p_t^2 \tag{21}$$

$$\mathrm{var}(\tau_{\mathrm{e}}|\mathcal{H}_1) = \sigma_{1,\mathrm{e}}^2 = \frac{L}{2}\left|g\frac{h^\dagger}{\|h\|}\right|^2 p_t^2\left(\|gZ\|^2\sigma_w^2 + \sigma_{\mathrm{e}}^2\right). \tag{22}$$

Then, her performance for a given collection of channel matrices is given by

$$P_K(h,g,N_o) = \int_{-\infty}^{\infty}\Phi\left(\frac{\tau - \sum_{i=1}^{N_o}\mu_{0,\mathrm{e}}(i)}{\sqrt{\sum_{i=1}^{N_o}\sigma_{0,\mathrm{e}}^2(i)}}\right)^{|\mathcal{K}|-1}\phi\left(\frac{\tau - \sum_{i=1}^{N_o}\mu_{1,\mathrm{e}}(i)}{\sqrt{\sum_{i=1}^{N_o}\sigma_{1,\mathrm{e}}^2(i)}}\right)d\tau\,, \tag{23}$$

where $\mu_{0,\mathrm{e}}(i)$, $\mu_{1,\mathrm{e}}(i)$, $\sigma_{0,\mathrm{e}}^2(i)$, and $\sigma_{1,\mathrm{e}}^2(i)$ are the means and variances for each observation $i = 1,\ldots,N_o$ due to varying h and g.

III. Authentication and Security Performance

The performance of the fingerprint embedding authentication framework with artificial noise is broken down into three parts. The first is the message rate which is dictated by the signal modulation constellation, error correction code, and error probability desired. Second is Bob's authentication detection probability which depends on the tag SNR and third is Eve's maximum likelihood key decoder success probability which depends on the tag SNR and AN power. In addition to improving communication and authentication performance, Alice's optimal beamforming provides an inherent advantage over Eve since the signal is directed towards Bob instead of her which increases security even without artificial noise. The results in this section, however, show that AN can significantly improve security performance, but at the expense of the message rate. To facilitate a fair trade-off between the three design goals, we assume that Alice has a peak power constraint P_0 that must be satisfied at all times. The constraint forces Alice to find a balance between allocating power to the message in order to prevent errors, allocating power to the tag to achieve Bob's desired performance, and allocating power to the artificial noise to antagonize Eve. The allocation must satisfy

$$P_0 \geq p_s^2 + p_t^2 + (N_T - 1)\sigma_w^2\,. \tag{24}$$

[1]We ignore the minor differences in the traditional substitution and impersonation attacks to assume a single attack strategy.

[2]We found that Eve's ML decoding performance for $N_o > 1$ is better when channel normalization is not performed at this step.

This section examines the trade-offs determined by the power allocations and system parameters. While both the AN and tag have deleterious effects on the message error performance, only the tag directly interferes with the message. Therefore, to minimize interference, we assume that Alice first chooses p_t^2 so that it attains a desired detection performance at Bob. Then, she chooses the maximum $\sigma_{\text{AN}}^2 = (N_T - 1)\sigma_w^2$ that leaves enough power to achieve her desired error performance of the message.

A. Communication and Authentication Performance

In Section II, both Bob and Eve's tests are designed with knowledge of deterministic h and g. In this section, we explore the average performance of their tests for stochastic h and g. We assume a Rayleigh slow-fading model in which the channel matrices are constant for the entirety of a given block. Both h and g are circularly symmetric i.i.d. complex Gaussian vectors. Note that Bob recomputes the threshold τ_0 for each channel realization.

Since Bob's performance is dependent on h through its norm $\|h\|$, which has a closed form expression for its distribution, we can compute his average detection performance as

$$E_h[P_D] = \int_{-\infty}^{\infty} P_D(\|h\|)P(\|h\|)d\|h\|, \tag{25}$$

where

$$P(\|h\|) = \frac{1}{\Gamma(N_T)}\|h\|^{N_T-1}e^{-\|h\|}, \tag{26}$$

and $\Gamma(\cdot)$ is the Gamma function. Equation (25) will be used to determine the tag power allocation required to achieve P_D.

The message error performance depends on the channel, coding structure, and the breakdown of the power allocation between the message, the tag, and the artificial noise. Bob observes the message with SNR of $\frac{p_s^2}{p_t^2+\sigma_{\text{b}}^2} = \frac{P_0-p_t^2-\sigma_{\text{AN}}^2}{p_t^2+\sigma_{\text{b}}^2}$. For the following plots, Equation (25) is used to determine the required tag power p_t^2 needed to obtain an average detection probability of $P_D = .9958$ for different noise powers and number of antennas. In Figure 1, the uncoded QPSK bit error rate (BER) is plotted versus varying artificial noise power allocation. The plot gives an indication of how much artificial noise can be allocated before communication breaks down. Further power can be allocated to AN if error correction codes or smaller alphabet modulation types are used leading to a trade-off between message rate and security. Figure 2 shows how the performance of different rate Reed-Solomon codes are affected by the additional allocation towards artificial noise. The $N_T = 10$ regime benefits more from the application of error coding. Using this analysis, Alice can determine the maximum amount of power she can allocate towards AN while maintaining the performance of both the message decoding and authentication detection.

Although we mainly focus on the impact of artificial noise on the authentication tag below in Section III-B, we note that it also affects Eve's ability to successfully decode the message. This adds a fair bit of security since correct decoding of the message is required for her to properly design her matched filter test in Equation (16). If she decodes the message in error, her codebook will be completely random and the decoder will be no better than a random key guess.

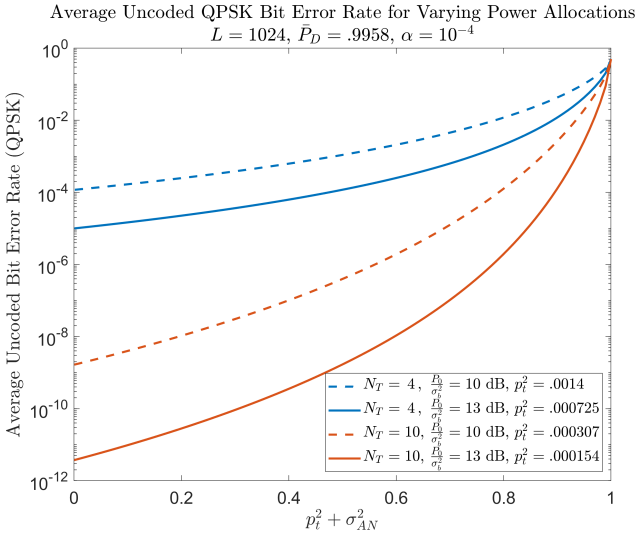

Fig. 1. Allocating more power towards artificial noise leads to larger bit errors. In this plot, the tag power is adjusted in each curve to achieve average detection performance of $P_D = .9958$. Greater antenna diversity leads to better error performance.

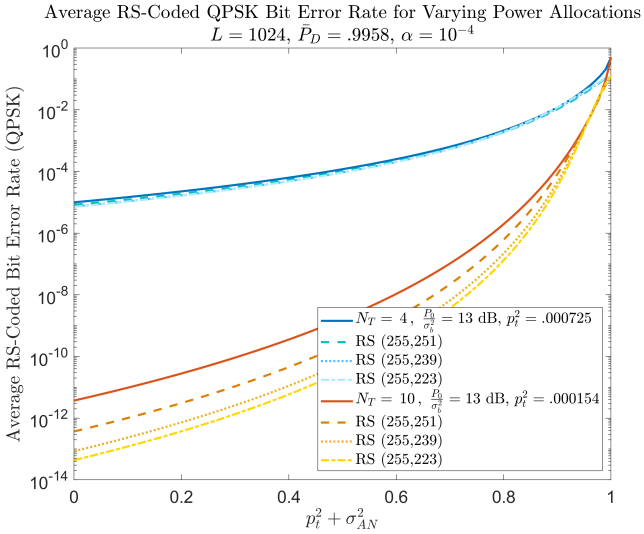

Fig. 2. Reducing the coding rate leads to better error performance and allows for larger artificial noise allocation. In this plot, the tag power is adjusted in each curve to achieve average detection performance of $P_D = .9958$.

B. Security Performance

After p_t^2 and σ_{AN}^2 are chosen, we can determine the system's security performance $P_K(h, g, N_o)$. While the distribution of Bob's SNR has a closed form expression, Eve's does not. Additionally, her performance depends directly on many realizations of h and g making it unwieldy to compute

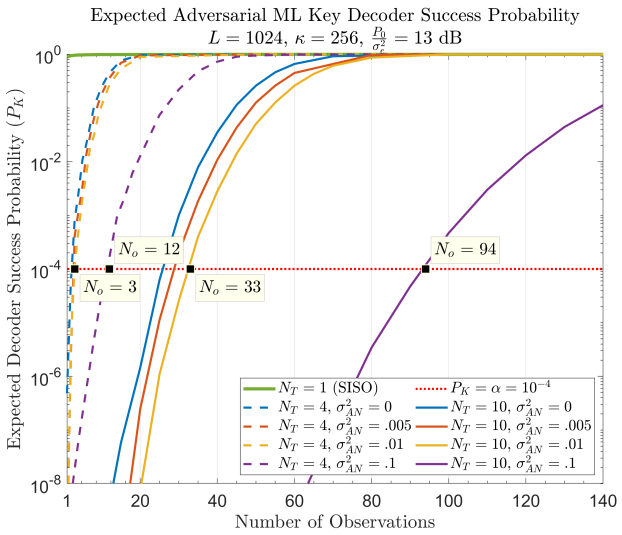

Fig. 3. Expected probability of correctly obtaining the key using an ML matched filter decoder for $N_T = 4, 10$. The lifespan of the key increases with σ^2_{AN} and N_T.

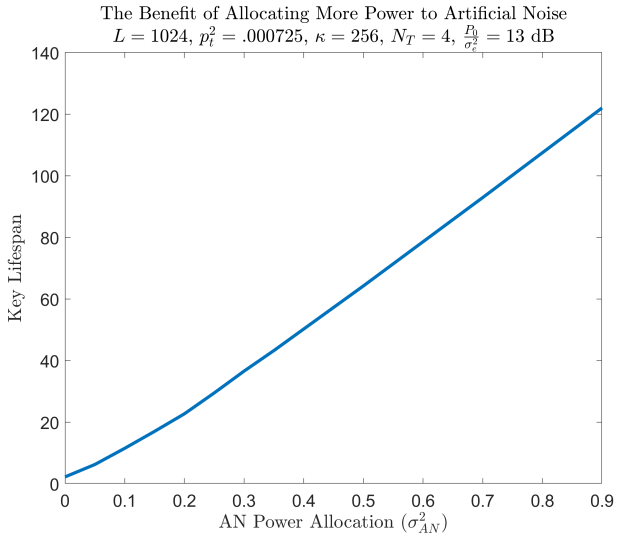

Fig. 4. The increase in key lifespan by allocating more power to artificial noise is slightly faster than linear for low σ^2_{AN}.

numerically. We, therefore, resort to Monte Carlo simulations to produce the expected performance over h and g.

In Figure 3, the average success probability of Eve's ML key decoder for a 256 bit key is plotted for various artificial noise power allocations, a normalized transmit SNR of $\frac{P_0}{\sigma_e^2} = 13$ dB, and $N_T = 4, 10$ transmit antennas.The curves are compared to the false alarm rate α to show at which point intelligent use of observations to design an attack outperforms a naive random tag attack. We refer to the average number of observations required for $P_K(h, g, N_o)$ to exceed α as the lifespan of the key. At the end of a key's lifespan, Alice and Bob would replenish their secret key to reset Eve's attack success probability. In this case, tag powers of $p_t^2 = .000725$ and $p_t^2 = .000154$ are required to achieve an average detection probability of $P_D = .9958$ for $N_T = 4$ and $N_T = 10$, respectively. The increased antenna diversity allows Alice to lower her tag power and allocate more power towards AN, both of which are detrimental to Eve's ML decoding performance. In addition to lowering the necessary tag power, more transmit antennas allows Alice to add AN to more dimensions, increasing the chance of having at least one large component at Eve. This results in much better performance for the $N_T = 10$ case where the lifespan of the key can increase to 94 observations/uses by dedicating 10% of the power to AN.

For comparison, Eve's performance in the SISO regime is also presented in Figure 3 in which $p_t^2 = .0869$ is required for the same average P_D. Since Alice cannot add AN or perform beamforming in SISO, Eve is able to observe the tag with high SNR, especially when the channel gain is favorable.

To examine the rate at which security performance increases with additional AN power allocation, we plot the average lifespan of the key for increasing σ^2_{AN} in Figure 4. The increase is slightly better than linear for low AN power allocations, but begins to increase linearly beyond $\sigma^2_{\text{AN}} \approx .3$. Comparing with Figure 1, if Bob only requires an average uncoded BER of 10^{-4}, he can afford to dedicate 47% of the total transmit power to the tag and AN to achieve a key lifespan of 60 transmissions/observations.

REFERENCES

[1] A. D. Wyner, "The wire-tap channel," *Bell Labs Technical Journal*, vol. 54, no. 8, pp. 1355–1387, 1975.

[2] R. Negi and S. Goel, "Secret communication using artificial noise," in *IEEE Vehicular Technology Conference*, vol. 62, no. 3. Citeseer, 2005, p. 1906.

[3] S. Goel and R. Negi, "Guaranteeing secrecy using artificial noise," *IEEE transactions on wireless communications*, vol. 7, no. 6, 2008.

[4] A. Khisti and G. W. Wornell, "Secure transmission with multiple antennas i: The misome wiretap channel," *IEEE Transactions on Information Theory*, vol. 56, no. 7, pp. 3088–3104, 2010.

[5] ——, "Secure transmission with multiple antennaspart ii: The mimome wiretap channel," *IEEE Transactions on Information Theory*, vol. 56, no. 11, pp. 5515–5532, 2010.

[6] X. Li, J. Hwu, and E. P. Ratazzi, "Using antenna array redundancy and channel diversity for secure wireless transmissions." *JCM*, vol. 2, no. 3, pp. 24–32, 2007.

[7] H.-M. Wang, T. Zheng, and X.-G. Xia, "Secure miso wiretap channels with multiantenna passive eavesdropper: Artificial noise vs. artificial fast fading," *IEEE Transactions on Wireless Communications*, vol. 14, no. 1, pp. 94–106, 2015.

[8] P. L. Yu, B. M. Sadler, G. Verma, and J. S. Baras, "Fingerprinting by design: Embedding and authentication," in *Digital Fingerprinting*, C. Wang, R. M. Gerdes, Y. Guan, and S. K. Kasera, Eds. Springer, 2016, pp. 69–88.

[9] J. B. Perazzone, L. Y. Paul, B. M. Sadler, and R. S. Blum, "Physical layer authentication via fingerprint embedding: Min-entropy analysis," in *Information Sciences and Systems (CISS), 2019 53rd Annual Conference on*. IEEE, 2019, pp. 1–6.

[10] L. Lai, H. El Gamal, and H. V. Poor, "Authentication over noisy channels," *IEEE Transactions on Information Theory*, vol. 55, no. 2, pp. 906–916, 2009.

[11] A. J. Menezes, P. C. Van Oorschot, and S. A. Vanstone, *Handbook of applied cryptography*. CRC press, 1996.