FOREX Publication
Open Access | Rapid and quality publishing

International Journal of
Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

RIS-CFH: Physical Layer Authentication via Channel Fingerprint Hopping Using a Secret Pseudo-Random Phase Sequence Against an Adaptive Adversary

Mohammed Aboud Kadhim[1*], Murooj Nadhom[2], and Ameer Abdul Sattar Jaber[3]
[1,3]*Polytechnic College for Engineering Middle Technical University (MTU) Baghdad, Iraq*
[2]*College of Electromechanical Engineering, University of Technology, Baghdad, Iraq*

***Correspondence:** *mohammed_aboud@mtu.edu.iq*

ABSTRACT- Physical layer authentication (PLA) verifies a transmitter from the characteristics of its wireless channel without per-packet cryptographic signaling. Reconfigurable intelligent surfaces (RIS) enrich such channel fingerprints, but a fixed RIS configuration can be learned and replayed by an adversary that accumulates pilot observations. This paper proposes RIS channel fingerprint hopping (RIS-CFH), in which the RIS configuration applied to each of S pilot symbols in every slot is generated by a cryptographically secure pseudo-random generator (CSPRNG) whose seed is shared only by Alice, Bob and the RIS controller. Bob does not compare the current observation with that of the previous slot; instead, he predicts the legitimate fingerprint of the current secret configuration from an enrolled and continuously tracked estimate of the direct and cascaded channels. The null distribution of the resulting statistic is shown to be Gamma$(S, 1)$, which yields an exact Neyman–Pearson threshold, while estimation error and channel ageing are absorbed by online calibration. Against the strongest non-predicting adversary, which is granted genie knowledge of the channel at Bob and an unlimited observation record, the per-slot spoofing probability is bounded by a regularized incomplete gamma function whose argument decreases in proportion to the number of RIS elements N at fixed transmit power. By contrast, a finite codebook of M states have an irreducible $(1-\alpha)/M$ spoofing floor, and PRBS31-indexed hopping is broken by Berlekamp–Massey synthesis after about 16 slots. Monte Carlo simulations at 28 GHz with path loss, spatial correlation, 1-bit phase quantization, channel-estimation error and Doppler ($N = 256$, $S = 4$, direct-link SNR of 25 dB) give an empirical false-alarm rate of 4.7% (quasi-static channel) to 5.6% (with Doppler, $f_D T_s = 10^{-3}$) for a 5% target, an AUC of 0.9988 ± 0.0003, and a per-slot spoofing probability of 0.50% (95% confidence interval 0.41–0.61%) that does not increase with the adversary's observation length, provided the channel coherence time exceeds approximately $(N+1)/S$ slots.

Keywords: Physical Layer Authentication, Reconfigurable Intelligent Surface, Channel Fingerprint Hopping, Cryptographically Secure Pseudo-Random Generator, Spoofing Detection, Hypothesis Testing, 6G Security, Adaptive Adversary.

ARTICLE INFORMATION
Author(s): Mohammed Aboud Kadhim, Murooj Nadhom, and Ameer Abdul Sattar Jaber;
Received: 29/05/26; **Accepted**: 27/09/26; **Published**: 30/09/26;
E- ISSN: 2347-470X;
Paper Id: IJEER 2905B18;
Citation: 10.37391/ijeer.140322
Webpage-link:
https://ijeer.forexjournal.co.in/archive/volume-14/ijeer-140322.html

Publisher's Note: FOREX Publication stays neutral with regard to jurisdictional claims in Published maps and institutional affiliations.

1. INTRODUCTION

The proliferation of resource-constrained nodes in sixth-generation (6G) networks—encompassing industrial IoT sensors, wearable medical devices, and vehicular terminals—has rendered conventional public-key infrastructure (PKI) authentication impractical for many deployments [1]. Physical layer security (PLS) exploits the randomness and spatial uniqueness inherent in wireless channels to provide authentication and confidentiality at the waveform level, without resource-intensive upper-layer cryptographic procedures [2]. Reconfigurable intelligent surfaces (RIS) are composed of large two-dimensional arrays of electronically tunable reflecting meta-elements [2], which can create a reprogrammable propagation environment with greater diversity and fingerprint discriminability, making them a front-runner candidate for lightweight PLA in 6G-IoT networks [3]. Because RIS strengthens legitimate links while also creating a new attack surface, designing RIS with security awareness is an important area for research [4]. Physical layer authentication (PLA) treats the identity verification task as a binary hypothesis test: the receiver calculates a test statistic based on the instantaneous channel and compares it with a threshold chosen to guarantee a desired false alarm rate [5].

Gaussian process classification-based PLA leverages configurable RIS fingerprints to address inaccurate channel estimates in 6G IoT, achieving superior authentication accuracy over conventional schemes [5]. Deep-learning methods such as EsaNet extract frequency-independent channel fingerprints from massive-MIMO CSI, providing robust spoofing detection even under fast time-varying conditions [6]. These advances confirm that RIS-enriched channel fingerprints are a powerful authentication feature, but

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

all rely on static or slowly adapted configurations. In the PLA context, static-phase RIS-assisted authentication has been shown to widen the fingerprint gap between Alice and Eve [7], [8]. RIS phase hopping—randomizing the phase shifts of individual elements per symbol—was introduced in the communications-reliability context to convert slow fading into fast fading and improve outage performance [9]. These works, however, either assume a passive adversary with a fixed configuration or do not consider authentication at all, leaving the adversarial-patience vulnerability—whereby Eve accumulates pilot observations until her estimate converges to Alice's fingerprint—unaddressed.

The threat of adversarial RIS was quantified in [10], which showed that an Eve-controlled surface can partially reconstruct secret keys from physical layer key generation (PL-SKG) protocols by injecting a random reciprocal channel component. This result demonstrates that intelligent, adaptive adversaries can exploit RIS programmability in ways far beyond the classical passive-Eve model, and motivates a rigorous adaptive threat model in which Eve may accumulate an arbitrarily long record of pilot observations (up to 5000 slots in our evaluation), cluster and classify them, attempt to predict the configuration sequence, and is additionally granted genie knowledge of the legitimate channel at Bob.

This paper proposes RIS-CFH, in which the RIS configuration applied to every pilot symbol is generated by a CSPRNG whose seed is shared only by Alice, Bob and the RIS controller, and in which Bob authenticates each slot against the fingerprint predicted for the current secret configuration. The main contributions are as follows. (i) A consistent authentication model is formulated: the legitimate reference is not the observation of the previous slot but a state-specific prediction obtained from an enrolled and tracked estimate of the direct and cascaded channels, so that the legitimate test error contains only receiver noise, estimation error and ageing. (ii) The exact null distribution Gamma$(S, 1)$, the Neyman–Pearson threshold, the detection probability and a corrected closed-form AUC are derived, and an online calibration rule that keeps the false-alarm rate at its target under channel ageing is given. (iii) An explicit threat model is stated, a per-slot spoofing bound valid for any non-predicting adversary is proved (Proposition 1), and its scaling with N and its extension to multi-slot sessions are established; the analysis also shows why a finite codebook and a linear PRBS generator are insufficient. (iv) Algorithm 1 specifies enrollment, tracking, update, fallback and seed-refresh rules. (v) A reproducible Monte Carlo evaluation over SNR, N, M, K_R, Eve position, spatial correlation, mobility, phase resolution, channel-estimation error and observation length is reported with 95% confidence intervals, together with a precise overhead model.

2. LITERATURE REVIEW AND RELATED WORK

Gaussian-process-based PLA [5] and EsaNet [6] address the challenges of inaccurate fingerprints and fast channel variation, respectively, but require substantial training overhead and are sensitive to environment changes. Both [5] and [6] assume passive adversaries and static channel features; neither models the adversarial-patience vulnerability. Similarly, Amin et al. [7], [8] optimize a static RIS configuration for PLA discriminability; however, the convergence threat posed by an observing adversary is not considered. To the best of our knowledge, the combination studied here—cryptographically generated RIS configuration hopping, a state-specific predictive reference and an explicit adaptive-adversary analysis—has not been reported; this statement is limited to the literature surveyed in Table I and to the baselines implemented in *section 7*.

The adversarial-RIS threat model was elaborated in [10], which designed Eve-RIS schemes against CSI-based and cross-multiplication PL-SKG. An explainable adversarial learning framework [11] further quantified the mutual-information gap between legitimate pairs and a man-in-the-middle RIS, confirming that intelligent adversaries can crack almost all existing PL-SKG protocols. Both [10] and [11] focus on key generation rather than authentication; the proposed RIS-CFH addresses authentication specifically and provides an analytical spoofing bound that complements learning-based methods.

RIS-assisted specific emitter identification (SEI) was proposed in [12], where ON-OFF control of individual RIS elements generates spatially selective link signatures that resolve co-located transmitters. Their RSS-based spoofing detection framework and closed-form threshold derivation directly inform the PLA performance metrics used in this paper. Unlike [12], RIS-CFH uses full N-element phase control and a secret CSPRNG-driven configuration sequence, and it is evaluated against observing, clustering and sequence-predicting adversaries.

Secrecy-rate performance of RIS-assisted MISO systems under Rician fading with spatially random eavesdroppers was analyzed in [13], confirming the role of the Rician K-factor and of the RIS size. Deep-reinforcement-learning-based PLS with RIS phase control [14] achieves improved secrecy rates at high computational complexity, whereas irregular RIS topologies [15] show that non-uniform element activation provides additional security degrees of freedom; that approach is orthogonal to RIS-CFH and could be combined with it. The smart-radio-environment paradigm of Di Renzo et al. [16], in which environmental objects are coated with reconfigurable meta-surfaces so that the propagation medium itself can be customized, underlies the passive-RIS model adopted here. Accurate cascaded-channel estimation is a prerequisite of the predictive reference used by RIS-CFH (*section 3.4*); the deep-learning estimator of Liu et al. [17] for double-RIS-aided massive MIMO illustrates how cascaded channels can be recovered from pilot observations with low normalized mean-square error. Dai et al. [18] proposed adaptive semantic communications based on online-learned nonlinear transform source-channel coding, which motivates the joint design discussed in *section 8*. Wu et al. [19] proposed token-based secret key generation that remains secure against an untrusted

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

internal eavesdropper; such a scheme is a natural way to establish and refresh the RIS-CFH seed without out-of-band exchange. Finally, Zhang et al. [20] compared active and passive RIS and showed that active RIS can offer substantially higher SNR at the cost of amplifier noise in the reflected signal; for RIS-CFH such noise would enter the residual variance and is absorbed by the online calibration of *section 4.4*.

Classical PLA in time-variant channels [21] compares the current channel response with a tracked estimate of the legitimate channel and explicitly accounts for channel variation between probes. RIS-CFH adopts the same principle, with the additional requirement that the legitimate reference must follow the secret RIS configuration. Table I summarizes the key distinctions between RIS-CFH and prior art under the attacker assumptions used in each work.

Table 1. Comparison of RIS-CFH with Prior Art

Work	Adversary	Observation model	Phase / state type	Metric	Key limitation
[5] Meng 2025	Passive	None	Configurable	Accuracy	Training overhead
[6] Gao 2024	Passive	None	Static (no RIS hopping)	Detection accuracy	Passive Eve only
[7], [8] Amin 2024/2026	Passive	None	Static, optimized	Accuracy	Fixed fingerprint is learnable
[9] Besser 2022	None	None	Random hopping	Outage	No security model
[10] Wei 2023	Active Eve-RIS	Limited	Adversarial	Key rate	Key generation, not authentication
[11] Wei 2025	Active MITM-RIS	Limited	Adversarial	MI gap	Key generation, not authentication
[12] Gao 2024	Co-located emitters	None	ON–OFF signatures	SEI accuracy	Static signatures
[14] Yan 2025	Passive	None	DRL-based	Secrecy rate	High complexity
This work (RIS-CFH)	**Adaptive; genie channel knowledge**	**Explicit (A1–A3), T_obs ≤ 5000**	**CSPRNG, fresh per pilot symbol**	**FAR, AUC, P_spoof with 95% CI**	**Needs coherence time ≫ (N+1)/S slots, S pilots/slot, authenticated enrollment**

3. SYSTEM MODEL

3.1. Network Topology and Propagation

The system comprises a single-antenna legitimate transmitter (Alice), a single-antenna receiver and authenticator (Bob), a passive RIS with $N = N_x N_y$ elements spaced $\lambda/2$, and a single-antenna adversary (Eve). The carrier frequency is $f_c = 28$ GHz ($\lambda \approx 10.7$ mm), representative of millimetre-wave bands in which LoS-dominant propagation justifies a Rician model [16]. The RIS lies in the y–z plane with its centre at the origin; Alice, Bob and Eve are located at distances d_{AR}, d_{RB} and d_{RE} from the RIS centre at the azimuths listed in Table II, and the direct Alice–Bob link suffers an additional blockage loss L_b. Following the measured path-loss model of Tang et al. [22], the cascaded power gain of element n is

$$\beta_n = \frac{G_e\, d_x d_y\, \lambda^2\, F(\theta_n^t) F(\theta_n^r)}{64\pi^3 \left(d_n^t\right)^2 (d_n^r)^2} \qquad (1)$$

where $G_e = 4$ is the element gain, $d_x = d_y = \lambda/2$, $F(\theta) = \cos\theta$ is the normalized element power pattern, and d_n^t, d_n^r are the distances from Alice to element n and from element n to Bob. The direct link has gain $\beta_d = (\lambda/4\pi d_{AB})^2 L_b$. Because per-element distances are used, the LoS phase profile is exact in the radiating near field (the Fraunhofer distance of the $N = 256$ array is 2.7 m, larger than d_{AR}).

3.2. Composite Channel Model

For a RIS configuration $\Phi = [\phi_1, \ldots, \phi_N]^T$, the end-to-end channel observed at Bob during slot t is

$$H(\Phi, t) = h_d(t) + \sum_{n=1}^{N} g_n(t)\, e^{j\phi_n} = \mathrm{x}(\Phi)^T \theta(t) \qquad (2)$$

where $g_n = \sqrt{\beta_n}\, [\mathrm{h}_r]_n^* [\mathrm{h}_t]_n$ is the cascaded coefficient of element n, $\theta = [h_d, g_1, \ldots, g_N]^T \in \mathbb{C}^{N+1}$ and $\mathrm{x}(\Phi) = \left[1, e^{j\phi_1}, \ldots, e^{j\phi_N}\right]^T$. The normalized Alice–RIS vector (and, analogously, the RIS–Bob vector) is

$$\mathrm{h}_t(t) = \sqrt{\frac{K_R}{K_R+1}}\, \mathrm{a}_t + \sqrt{\frac{1}{K_R+1}}\, \mathrm{R}^{1/2} \mathrm{w}_t(t) \qquad (3)$$

where $[\mathrm{a}_t]_n = e^{-j2\pi d_n^t/\lambda}$ is the spherical-wave LoS component, R with $[\mathrm{R}]_{mn} = \mathrm{sinc}(2 \parallel \mathrm{u}_m - \mathrm{u}_n \parallel/\lambda)$ is the spatial correlation of a planar RIS under isotropic scattering [23], and $\mathrm{w}_t \sim \mathcal{CN}(0, \mathrm{I})$.

$$\mathrm{w}(t) = \varepsilon\, \mathrm{w}(t-1) + \sqrt{1-\varepsilon^2}\, \mathrm{e}(t), \varepsilon = J_0(2\pi f_D T_s) \qquad (4)$$

The direct path is Rayleigh distributed. Temporal variation follows the first-order Gauss–Markov model applied to every NLoS component, with Jakes correlation coefficient ε [24], Doppler frequency f_D and slot duration T_s. Phase shifts are

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

quantized to b bits, $\phi_n \in \{2\pi k/2^b\}$. Eve's channels follow the same model at her own position, and her NLoS components are independent of Alice's because their separation greatly exceeds $\lambda/2$. With unit transmit power, Bob receives $y = H + n$, $n \sim \mathcal{CN}(0, \sigma^2)$, and the direct-link SNR is $\mathrm{SNR}_d = \beta_d/\sigma^2$. The quantity that governs security is the cascaded-to-residual power ratio

$$\bar{\rho} = \frac{\|g\|^2}{\sigma_T^2} = \frac{1}{\sigma_T^2}\sum_{n=1}^{N} \beta_n \, |[\mathrm{h}_r]_n|^2 |[\mathrm{h}_t]_n|^2 \qquad (5)$$

where σ_T^2 is the legitimate residual variance defined in *section 4.1*. At fixed transmit power $\bar{\rho}$ grows linearly with N, but its slope is set by the per-element gain (1); in the baseline geometry of Table II the average per-element cascaded gain is 38.7 dB below the direct-link gain. A normalization in which every element is as strong as the direct link therefore overestimates the fingerprint mismatch by about 39 dB, and it is not used here.

4.3. Secret Configuration Generation and Threat Model

Alice, Bob and the RIS controller share a 128-bit seed s_0 established at pairing (by a cryptographic key agreement or a PL-SKG protocol such as [19]) and refreshed at least once per session. The configuration applied during pilot symbol $s \in \{1, \ldots, S\}$ of slot t is

$$\Phi_{t,s} = \frac{2\pi}{2^b}\, \mathcal{G}(s_0, t, s) \qquad (6)$$

where $\mathcal{G}$ is a counter-mode deterministic random bit generator (CTR_DRBG with AES-128 [25]) that outputs Nb bits per pilot symbol from the seed and the shared slot counter. The configuration is never transmitted over the air: the controller runs $\mathcal{G}$ locally. A linear-feedback PRBS such as PRBS31 is deliberately not used, because its output is fully determined by 62 consecutive bits through Berlekamp–Massey synthesis [26], whereas the output of a CSPRNG is computationally indistinguishable from uniform even when all previous outputs are known. The prediction advantage of any efficient adversary over a uniform guess is denoted ϵ_G and is negligible. The adversary is characterized by the following explicit assumptions.

A1 (knowledge): Eve knows the protocol, N, b, S, α, η and the generator $\mathcal{G}$, and she is granted genie knowledge of the current channel vector $\theta(t)$ at Bob and of Bob's reference estimate; she does not know s_0.

A2 (observation): Eve may observe an unlimited number T_{obs} of past pilots, both at her own location and, as a further genie, at Bob's location, and she may cluster, classify and exploit temporal correlation in these observations.

A3 (transmission): Eve can synthesize any desired complex value at Bob's receiver, but she must commit to her pilot block before the configuration of the current pilot symbol is applied; the controller enforces this by switching $\Phi_{t,s}$ synchronously with the pilot symbol, so that same-symbol sensing is not possible.

A4 (enrollment): Enrollment takes place in an authenticated session (for example at pairing, using upper-layer credentials), so Eve cannot poison the initial reference.

A5 (generator): $\mathcal{G}$ is a secure CSPRNG (ϵ_G negligible); compromise of s_0 is outside the scope of the physical-layer analysis and is limited by seed refresh.

4.4. Enrollment and State-Specific Predictive Reference

During enrollment Alice transmits $P_g = L_g(N+1)$ pilots while the RIS applies a known set of orthogonal training configurations X_g (for example Hadamard columns, which a 1-bit RIS can realize), and Bob computes the least-squares estimate

$$\hat{\theta}(0) = \left(\mathrm{X}_g^H \mathrm{X}_g\right)^{-1} \mathrm{X}_g^H \mathrm{y}_g, \qquad \mathbb{E}\,\|\,\hat{\theta} - \theta\,\|^2 = \frac{\sigma^2}{L_g} \qquad (7)$$

since $\mathrm{X}_g^H \mathrm{X}_g = P_g \mathrm{I}$. During operation, Bob predicts the legitimate fingerprint of the current secret configuration as

$$\hat{H}_s(t) = \mathrm{x}\left(\Phi_{t,s}\right)^T \hat{\theta}(t) \qquad (8)$$

This resolves the question of how Bob knows the instantaneous legitimate channel for each new RIS state: the PRBS/CSPRNG identifies the configuration, and the tracked vector $\hat{\theta}$ maps that configuration to its complex fingerprint. After a slot is accepted, the estimate is updated by block normalized least-mean-squares (NLMS),

$$\hat{\theta}(t+1) = \hat{\theta}(t) + \frac{\mu}{N+1}\sum_{s=1}^{S} \mathrm{x}_s^* \, r_s(t),$$
$$r_s(t) = y_s(t) - \hat{H}_s(t) \qquad (9)$$

followed by a direct-path correction $\hat{h}_d \leftarrow \hat{h}_d + \mu_d \frac{1}{S}\sum_s r'_s$ computed with the updated residuals r'_s. The reference is updated only after an ACCEPT decision, so rejected (possibly adversarial) slots cannot contaminate it; after R_{max} consecutive rejections Bob falls back to upper-layer re-authentication and re-enrollment. Because the enrollment and tracking pilots reveal nothing about future configurations, the information that Eve could extract from them is already covered by the genie assumption A1.

4.5. Test Statistic and Decision Rule

Bob forms the normalized statistic

$$T(t) = \frac{1}{\hat{\sigma}_T^2}\sum_{s=1}^{S} \left| y_s(t) - \hat{H}_s(t) \right|^2 \qquad (10)$$

and decides ACCEPT (H_0, Alice) if $T(t) \le \eta$ and REJECT (H_1, spoofer) otherwise. The reference in *eq. (10)* is the fingerprint predicted for the current secret configuration; it is not the observation of the previous authenticated slot.

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

Consequently, the legitimate state change $H(\Phi_t, t) - H(\Phi_{t-1}, t-1)$ does not enter the legitimate test error.

4. MATHEMATICAL DERIVATIONS

4.1. Distribution of T under H₀ (Alice)

Under H_0, $y_s = \mathrm{x}_s^T\theta + n_s$ and $r_s = n_s - \mathrm{x}_s^T\epsilon$, where $\epsilon = \hat{\theta} - \theta \sim \mathcal{CN}\left(0, \left(\sigma^2/P_g\right)\mathrm{I}\right)$ is independent of the noise. For any configuration with unit-modulus entries, $\mathrm{x}_s^T\epsilon \sim \mathcal{CN}\left(0, \sigma^2/L_g\right)$, hence

$$r_s \sim \mathcal{CN}(0, \sigma_T^2), \qquad \sigma_T^2 = \sigma^2\left(1 + \frac{1}{L_g}\right) + \sigma_a^2 \tag{11}$$

where σ_a^2 is the ageing contribution of *section 4.4* (zero for a quasi-static channel). The cross-covariance of r_s and $r_{s'}$ equals $\left(\sigma^2/P_g\right)\mathrm{x}_s^T\mathrm{x}_{s'}^*$, whose magnitude relative to σ_T^2 is of order $1/\left(L_g\sqrt{N}\right)$ for independently generated configurations, so the S residuals are treated as independent. Each $2|r_s|^2/\sigma_T^2$ is then $\chi^2(2)$, and T is the sum of S independent Exp(1) variables:

$$f_T(\tau \mid H_0) = \frac{\tau^{S-1}e^{-\tau}}{(S-1)!}, \qquad \tau \geq 0 \tag{12}$$

that is, $T \sim \mathrm{Gamma}(S, 1)$, which reduces to Exp(1) for $S = 1$. The false-alarm rate is

$$P_{FA}(\eta) = \frac{\Gamma(S,\eta)}{\Gamma(S)} = e^{-\eta}\sum_{k=0}^{S-1}\frac{\eta^k}{k!} \tag{13}$$

and the Neyman–Pearson threshold for a target FAR α is

$$\eta^* = F^{-1}_{\Gamma(S,1)}(1-\alpha) = 1/2\, F^{-1}_{\chi^2_{2S}}(1-\alpha) \tag{14}$$

which gives $\eta^* = 2.996$, 7.754 and 13.148 for $S = 1$, 4 and 8 at $\alpha = 0.05$. For comparison, if the reference were the observation of the previous slot while the configuration changes, the residual would contain the legitimate difference $\mathrm{g}^T(e^{j\Phi_t} - e^{j\Phi_{t-1}})$ of variance $2\parallel \mathrm{g} \parallel^2$ in addition to $2\sigma^2$; the statistic would be scaled by $2(1+\parallel \mathrm{g} \parallel^2/\sigma^2)$ and the actual FAR would be $\exp[-\eta/(2 + 2\parallel \mathrm{g} \parallel^2/\sigma^2)]$, which tends to one as N grows. This is the inconsistency that the present formulation removes (*fig. 1*).

4.2. Distribution of T under H_1 and the Optimal Non-Predicting Attacker

Let Eve's committed pilot block produce the values $z_s + n_s$ at Bob, and let $u_s = z_s + n_s - \hat{H}_s$. Because $\Phi_{t,s}$ is (computationally) uniform and unknown to Eve, and N is large, the predicted reference $\hat{H}_s = \hat{h}_d + \sum_n \hat{g}_n\, e^{j\phi_n}$ is, conditioned on $\hat{\theta}$, approximately $\mathcal{CN}\left(\hat{h}_d, \parallel \hat{\mathrm{g}} \parallel^2\right)$ by the Lindeberg central limit theorem (for $b \geq 1$, $\mathbb{E}\left[e^{j\phi}\right] = 0$ and $\mathbb{E}\left|e^{j\phi}\right|^2 = 1$), independently across s. Therefore

$$u_s \sim \mathcal{CN}\left(z_s - \hat{h}_d,\ \sigma^2 + \parallel \hat{\mathrm{g}} \parallel^2\right) \tag{15}$$

and the statistic is a scaled non-central chi-square variable,

$$\frac{2\sigma_T^2}{\sigma^2+\|\hat{\mathrm{g}}\|^2}\, T_E \sim \chi^2_{2S}(\lambda),\ \lambda = \frac{2\sum_s |z_s - \hat{h}_d|^2}{\sigma^2+\|\hat{\mathrm{g}}\|^2} \tag{16}$$

The detection probability is

$$P_D(\mathrm{z}) = Q_S\left(\sqrt{\lambda}, \sqrt{\frac{2\eta\,\sigma_T^2}{\sigma^2+\|\hat{\mathrm{g}}\|^2}}\right) \tag{17}$$

where $Q_S(\cdot,\cdot)$ is the generalized Marcum Q-function of order S. Because $Q_S(a, b)$ increases with a, the detection probability is minimized, and the spoofing probability maximized, at $\lambda = 0$, that is, at $z_s = \hat{h}_d$; this is also a consequence of Anderson's theorem for symmetric unimodal densities [27]. The optimal non-predicting spoofer therefore replays the configuration-averaged fingerprint, and its success probability is

$$P^*_{spoof} = P\left(S, \frac{\eta}{\kappa+\bar{\rho}}\right), \qquad \kappa = \frac{\sigma^2}{\sigma_T^2}, \quad \bar{\rho} = \frac{\|\hat{\mathrm{g}}\|^2}{\sigma_T^2} \tag{18}$$

where $P(S, x) = \gamma(S, x)/\Gamma(S)$ is the regularized lower incomplete gamma function. The time-average (mean-replay) attacker converges to $z_s = h_d$ and therefore cannot exceed (18); a longer observation record cannot improve it. For 1-bit phases the reference is slightly improper; its circularity coefficient $|\sum_n \hat{g}_n^2|/\sum_n |\hat{g}_n|^2$ is below 0.13 in all simulated geometries, which changes (18) by less than 5%.

4.3. AUC Derivation

For the optimal attacker, $T_A \sim \Gamma(S, 1)$ and $T_E \sim c\,\Gamma(S, 1)$ with $c = \kappa + \bar{\rho}$, independent of T_A. Writing $T_E = cX$ and $T_A = Y$, with X, Y i.i.d. $\Gamma(S, 1)$, the event $T_E > T_A$ is equivalent to $Y/(X + Y) < c/(1 + c)$, and $Y/(X + Y) \sim \mathrm{Beta}(S, S)$. Hence

$$\mathrm{AUC} = P(T_E > T_A) = I_{c/(1+c)}(S, S), \qquad c = \kappa + \bar{\rho} \tag{19}$$

where $I_x(\cdot,\cdot)$ is the regularized incomplete beta function; for $S = 1$ this becomes $\mathrm{AUC} = c/(1 + c)$. The AUC tends to 1 as $\bar{\rho} \to \infty$ and to 0.5 when $\bar{\rho} \to 0$ and $\kappa \to 1$, as required.

4.4. Channel Ageing, Calibration and Hopping Rate

Between updates the channel evolves according to *eq. (4)*, and the tracking error of *eq. (9)* adds σ_a^2 to *eq. (11)*. Since σ_a^2 is not known a priori, Bob calibrates the residual variance online from accepted slots:

$$\hat{\sigma}_T^2 \leftarrow (1-\lambda_c)\hat{\sigma}_T^2 + \frac{\lambda_c}{c_S}\cdot\frac{1}{S}\sum_{s=1}^{S}|r_s|^2, \quad c_S = \frac{F_{\Gamma(S+1,1)}(\eta)}{1-\alpha} \tag{20}$$

where c_S corrects the truncation bias of averaging only accepted slots, because $\mathbb{E}[T \mid T \leq \eta] = S\,c_S$. With Gaussian residuals, *eq. (12)–(14)* remain valid with the calibrated $\hat{\sigma}_T^2$, and ageing reduces $\bar{\rho}$ in *eq. (18)*, that is, it lowers detection rather than inflating the FAR, as long as the tracker remains locked. Tracking requires the $N + 1$ unknowns of θ to be re-estimated within a coherence time, each slot providing S observations; with coherence time $T_c \approx 0.423/f_D$ this gives the rule of thumb.

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

$$T_c \gg \frac{N+1}{S} T_s \quad \Leftrightarrow \quad f_D T_s \ll \frac{0.423\,S}{N+1} \tag{21}$$

which equals 6.6×10^{-3} for $N = 256$ and $S = 4$. Condition *eq. (21)* is a necessary rather than a sufficient design guideline: a larger S also makes the statistic more sensitive to residual tracking lag, so the practical limit must be verified numerically (*fig. 2*). The hopping rate itself (one fresh configuration per pilot symbol) does not help Eve, because consecutive configurations are independent; slowing the hopping would allow Eve to reuse a configuration that she has sensed, which assumption A3 excludes only within a single pilot symbol. The RIS switching time must be shorter than the pilot symbol, which is satisfied by PIN-diode surfaces with microsecond-level switching.

5. SECURITY ANALYSIS

5.1. Static RIS Adversary

For a fixed configuration, Eve's sample mean of T_{obs} observations have error variance σ^2/T_{obs} per pilot symbol, so that $T_E = (1 + \kappa/T_{obs})\, \Gamma(S, 1)$ and

$$P_{spoof}^{ST}(T_{obs}) = P\left(S, \frac{\eta}{1+\kappa/T_{obs}}\right) \overset{T_{obs}\to\infty}{\to} 1 - \alpha \tag{22}$$

A static RIS therefore offers no protection against a patient adversary.

5.2. Finite Codebook of M Reused States

If the configurations are drawn from a finite codebook of M states that are reused, each with an enrolled per-state reference, a dictionary attacker can learn the M fingerprints by clustering Bob-side observations and then guess the state. When the guess is correct Eve passes with probability close to $1 - \alpha$; when it is wrong the two fingerprints differ by an approximately $\mathcal{CN}(0, 2 \parallel \mathrm{g} \parallel^2)$ term per pilot symbol. Hence

$$P_{spoof}^{CB} \approx \frac{1-\alpha}{M} + \frac{M-1}{M}\, \mathbb{E}\left[P\left(S, \frac{\eta}{\kappa+2\bar{\rho}}\right)\right] \tag{23}$$

whose floor $(1 - \alpha)/M$ does not vanish with N or with T_{obs}. This is why RIS-CFH draws a fresh configuration for every pilot symbol ($M \to \infty$ in effect) instead of indexing a small codebook.

5.3. Linear PRBS Generators

If the state index is produced by a linear PRBS such as PRBS31 ($x^{31} + x^{28} + 1$), the sequence has linear complexity 31, and Berlekamp–Massey synthesis recovers the generator from 62 correct output bits [26], that is, from 16 slots when $M = 16$. Eve can infer past indices by classifying her own observations against her state-dependent channels; in the baseline geometry her per-slot classification accuracy exceeds 99.9%, and a sliding-window synthesis with a consistency check predicted the future sequence in every simulated realization. The spoofing probability then returns to the static value $1 - \alpha$. An unknown seed is therefore not a security argument by itself; the generator must be cryptographically secure and the seed must be refreshed.

5.4. Proposition 1 — Per-Slot Spoofing Bound

Under assumptions A1–A5, for any Eve strategy that may be adaptive and may use all past observations and the genie channel knowledge,

$$P_{spoof} \le P\left(S, \frac{\eta}{\kappa + \bar{\rho}}\right) + \epsilon_G + \epsilon_N \tag{24}$$

where ϵ_N is the Berry–Esseen error of the Gaussian approximation in *eq. (15)*, of order $N^{-1/2}$.

Proof: Past observations reveal at most past configurations. By A5 the configuration of the current slot is independent of Eve's view up to a total-variation distance ϵ_G, and by A3 her pilot block is committed before it is applied. Conditioned on her view, the references $\widehat{H}_s$ are therefore distributed as in *eq. (15)* regardless of her history, T_E follows *eq. (16)* for her committed z, and *eq. (17)* is minimized at $z_s = \hat{h}_d$, which yields *eq. (18)*. Adding the two approximation terms gives *eq. (24)*.

5.5. Corollary 1 — Scaling with N

At fixed transmit power, $\mathbb{E} \parallel \mathrm{g} \parallel^2 = \sum_n \beta_n\, \mathbb{E}|[\mathrm{h}_r]_n|^2 |[\mathrm{h}_t]_n|^2$ grows linearly with N, so $\bar{\rho} \propto N$. Since $P(S, x) \approx x^S/S!$ for small x,

$$P_{spoof}^{*} \approx \frac{1}{S!}\left(\frac{\eta}{\bar{\rho}}\right)^S = \mathcal{O}(N^{-S}) \tag{25}$$

The decay is polynomial and its constant depends on the geometry through *eq. (1)*. Under a normalization that fixes the received SNR instead of the transmit power, $\bar{\rho}$ saturates once $\parallel \mathrm{g} \parallel^2$ becomes comparable to β_d, which in the baseline geometry happens only for N of the order of 7×10^3. The earlier element-normalized analysis that yielded a mismatch equal to N times the SNR is therefore not supported once path loss is included (*fig. 5*).

5.6. Corollary 2 — Session-Level Authentication

If a session is accepted when at least k of W consecutive slots pass, then, since the per-slot success probability conditioned on any history is at most the bound p of (24) (configurations are independent across slots),

$$P_{spoof}^{sess} \le \sum_{j=k}^{W} \binom{W}{j} p^j (1-p)^{W-j}, P_{FR}^{sess} = \sum_{j=0}^{k-1} \binom{W}{j} (1-\alpha)^j \alpha^{W-j} \tag{26}$$

6. RIS-CFH Protocol — Algorithm 1

Algorithm 1 summarizes the complete protocol. In the enrollment phase, which takes place in an authenticated session, Bob computes the threshold η^* from *eq. (14)* and the calibration constant c_S from (20), Alice transmits $L_g(N+1)$ training pilots with known orthogonal configurations, and Bob obtains the initial estimate $\hat{\theta}$ by least squares (7) and initializes $\hat{\sigma}_T^2 = \hat{\sigma}^2\left(1 + 1/L_g\right)$, where $\hat{\sigma}^2$ is obtained from a pilot-based noise calibration.

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

In the online phase, at the beginning of each slot the RIS controller and Bob independently advance their synchronized CSPRNG instances with the shared slot counter and obtain the S configurations of *eq. (6)*; nothing is transmitted over the air. The controller applies each configuration synchronously with the corresponding pilot symbol. Bob predicts the legitimate fingerprints with *eq. (8)*, computes $T(t)$ with *eq. (10)* and accepts when $T(t) \le \eta^*$. After an ACCEPT, Bob updates $\hat{\theta}$ with *eq. (9)* and $\hat{\sigma}_T^2$ with (20). After a REJECT, both are frozen, so that a spoofer cannot steer the reference; $R_{max} = 3$ consecutive rejections (probability $\alpha^3 \approx 1.25 \times 10^{-4}$ per slot for Alice) trigger upper-layer re-authentication and re-enrollment. The seed is refreshed every session.

The adversary is assumed to know everything in *algorithm 1* except s_0. As shown in *section 5*, her best non-predicting strategy is to replay the configuration-averaged fingerprint $\hat{h}_d$, whose success probability is bounded by *eq. (24)* independently of how many observations she has collected; the security of the protocol therefore rests on the unpredictability of the CSPRNG and on the cascaded power ratio $\bar{\rho}$, not on the secrecy of the codebook or of the channel.

Algorithm 1: RIS-CFH authentication with state-specific predictive reference

Input: N, b, S, α, L_g, μ, μ_d, λ_c, R_max; shared 128-bit seed s_0 (Alice, Bob, RIS controller)
Output: decision D(t) ∈ {ACCEPT, REJECT} for every slot t
Enrollment (authenticated session)
1: $\eta^* \leftarrow F^{-1}_\Gamma(S,1)(1-\alpha)$ [Eq. (14)]; c_S ← $F_\Gamma(S+1,1)(\eta^*)/(1-\alpha)$ [Eq. (20)]
2: Alice sends L_g(N+1) training pilots with known orthogonal configurations X_g; Bob computes $\hat{\theta}$ by LS [Eq. (7)]
3: $\hat{\sigma}_T^2 \leftarrow \hat{\sigma}^2(1 + 1/L_g)$; r_cnt ← 0
Online authentication (every slot t)
4: RIS controller and Bob: $\Phi_{t,s} \leftarrow (2\pi/2^b)\cdot\mathcal{G}(s_0, t, s)$, s = 1,…,S [Eq. (6)] (nothing sent over the air)
5: Controller applies $\Phi_{t,s}$ synchronously with pilot symbol s; the claimed transmitter sends S pilots
6: Bob receives y_s(t), predicts $\hat{H}_s(t) = x(\Phi_{t,s})^T \hat{\theta}$ [Eq. (8)] and computes T(t) [Eq. (10)]
7: if $T(t) \le \eta^*$ then D(t) ← ACCEPT; update $\hat{\theta}$ by Eq. (9); update $\hat{\sigma}_T^2$ by Eq. (20); r_cnt ← 0
8: else D(t) ← REJECT; $\hat{\theta}$ and $\hat{\sigma}_T^2$ are frozen; r_cnt ← r_cnt + 1
9: if r_cnt ≥ R_max then upper-layer re-authentication; on success re-enroll (Step 2)
10: refresh s_0 every session (e.g., PL-SKG [19] or cryptographic handshake)
11: return D(t)

Adversary (section 5): knows everything above except s_0; may use genie θ, unlimited observations, clustering and sequence prediction; must commit her pilot block before $\Phi_{t,s}$ is applied. Best non-predicting strategy: replay $\hat{h}_d$; success bounded by eq. (24).

7. SIMULATION RESULTS

All results were produced with a Python/NumPy implementation of *eq. (1)–(10)*. For each Monte Carlo realization an independent channel (LoS geometry, correlated Rician hops, Rayleigh direct path), an independent enrollment error, independent configuration bits and an independent 128-bit seed were drawn. The CSPRNG was implemented as a SHA-256 counter-mode generator, which stands in for the AES-128 CTR_DRBG of *eq. (6)*; the PRBS31 baseline uses the Fibonacci LFSR $x^{31} + x^{28} + 1$ with $\log_2 M$ bits per slot. Attack experiments use a quasi-static channel over the observation and attack windows, which is the most favourable case for Eve, whereas false-alarm and tracking experiments use the Gauss–Markov model *eq. (4)* over 3000 slots, of which the first 300 are discarded. The attacks evaluated are: A1, mean replay of T_{obs} Bob-side observations; A1*, the optimal non-predicting genie attacker of Section IV-B; A2, k-means clustering of T_{obs} Bob-side observations into M fingerprints followed by uniform state guessing; A2*, the same with a perfect (genie) dictionary; and A3, classification of Eve-side observations against her genie state-dependent channels followed by sliding-window Berlekamp–Massey prediction. All baselines use the same channel, SNR, enrollment error, S and threshold. Proportions are reported with Wilson 95% confidence intervals and AUC values with the Hanley–McNeil standard error [28]; analytical curves were evaluated on channel realizations generated with seeds independent of those used in the simulations, which provides an independent validation. The source code is available from the corresponding author as supplementary material.

Table 2. Simulation Parameters

Parameter	Symbol	Value
Carrier frequency / wavelength	f_c, λ	28 GHz, 10.7 mm
RIS elements (UPA, λ/2 spacing)	N	256 = 16×16 (Figs. 1–4, 6); 16–1024 (Fig. 5)
Element gain / pattern	G_e, F(θ)	4, cos θ [22]
Phase resolution	b	1 bit (2 bits and continuous in Table IV)
Distances from RIS centre	d_AR, d_RB, d_RE	2 m (1, 4 m in Fig. 5), 6 m, 6 m
Azimuths of Alice, Bob, Eve	—	30°, −40°, −10°
Direct-link blockage loss	L_b	20 dB, Rayleigh direct path

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

Parameter	Symbol	Value
Rician factor of RIS hops	K_R	5 (0, 10 in Table IV)
Spatial correlation	R	sinc model [23]
Direct-link SNR at Bob	SNR_d	25 dB (15, 35 dB in Table IV)
Pilot symbols per slot	S	4 (1 for comparison)
Target FAR / threshold	α, η*	5%; 7.754 (S = 4), 2.996 (S = 1), Eq. (14)
Enrollment pilots	L_g(N+1)	L_g = 20 (5, 100 in Table IV)
Tracking / calibration steps	μ, μ_d, λ_c	0.2, 0.3, 0.05; R_max = 3
Normalized Doppler	f_D T_s	10^{-3} (10^{-4}–3×10^{-2} in Fig. 2)
Codebook size (baselines)	M	16
Eve observations	T_obs	500 (Figs. 3, 6); 5–5000 (Fig. 4)
Monte Carlo size	—	40 × 2700 slots (Figs. 1, 2); 400 × 50 (Figs. 3, 6); 200 × 50 per point (Fig. 4); 150 × 300 per point (Fig. 5)
Seed / generator	s_0, $\mathcal{G}$	128 bits; SHA-256 counter mode (simulation), AES-128 CTR_DRBG [25] (protocol)

7.1. Validity of the Null Model — Fig. 1

Fig. 1(a) shows the empirical CDF of T under H_0 over 108,000 legitimate slots at $f_D T_s = 10^{-3}$. With the state-specific predictive reference, the empirical distribution coincides with Gamma(4,1) for $S = 4$ and with Exp(1) for $S = 1$ (Kolmogorov–Smirnov distances of 0.020 and 0.026, respectively). With the previous-slot reference, the statistic is shifted by more than an order of magnitude and the FAR at the nominal 5% threshold is 88.9%, which confirms that a previous-slot reference is inconsistent with slot-to-slot hopping. *Fig. 1(b)* compares the empirical FAR with the target over $\alpha \in [0.01, 0.2]$: the calibrated scheme follows the target with a small upward bias (5.6% for a 5% target with $S = 4$, and a relative excess of 7–24% over the whole range) caused by the finite averaging window of *eq. (20)* and the residual correlation of the tracking error; where a strict ceiling is required, the bias can be compensated by designing for a slightly smaller α.

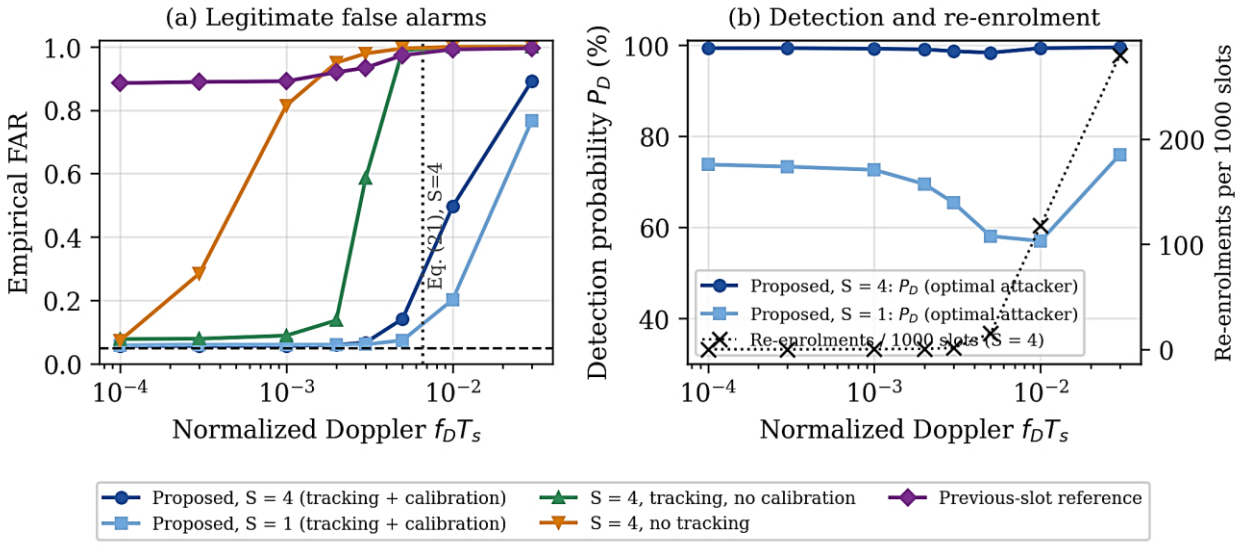


Fig. 2. Channel coherence versus hopping rate (N = 256, SNR_d = 25 dB, 108,000 slots per point). (a) Empirical legitimate FAR versus $f_D T_s$ for the proposed scheme (S = 4 and S = 1), for tracking without calibration, without tracking, and for a previous-slot reference; the dotted vertical line is the limit of Eq. (21) for S = 4. (b) Detection probability of the optimal non-predicting attacker and re-enrollment rate.

7.2. Channel Coherence versus Hopping Rate — Fig. 2

Fig. 2(a) plots the legitimate FAR against the normalized Doppler $f_D T_s$. With tracking (9) and calibration (20), the FAR stays between 5.6% and 6.7% for $f_D T_s \le 3 \times 10^{-3}$ and rises sharply as (21) is approached (vertical dotted line), when the tracker can no longer follow the $N + 1$ channel coefficients. With $S = 1$ the FAR degrades somewhat later, because a single residual is less sensitive to tracking lag, but its detection capability is far weaker. Without calibration the FAR drifts upward earlier, and without tracking it exceeds 28% already at $f_D T_s = 3 \times 10^{-4}$; the previous-slot reference produces a FAR close to 90% at every Doppler value. Fig. 2(b) shows that detection of the optimal attacker degrades only mildly with Doppler for $S = 4$ (from 99.3% to 98.3%), whereas for $S = 1$ it is limited to about 73% even for a static channel. The right axis gives the re-enrollment rate triggered by

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

R_{max} consecutive rejections, which remains below 1.7 per 1000 slots within the operating region. For example, at 28 GHz with $T_s = 0.125$ ms, the condition $f_D T_s \leq 3 \times 10^{-3}$ corresponds to speeds up to about 0.26 m/s, which covers static and slowly moving IoT devices; faster mobility requires a larger S, a smaller N or a shorter slot.

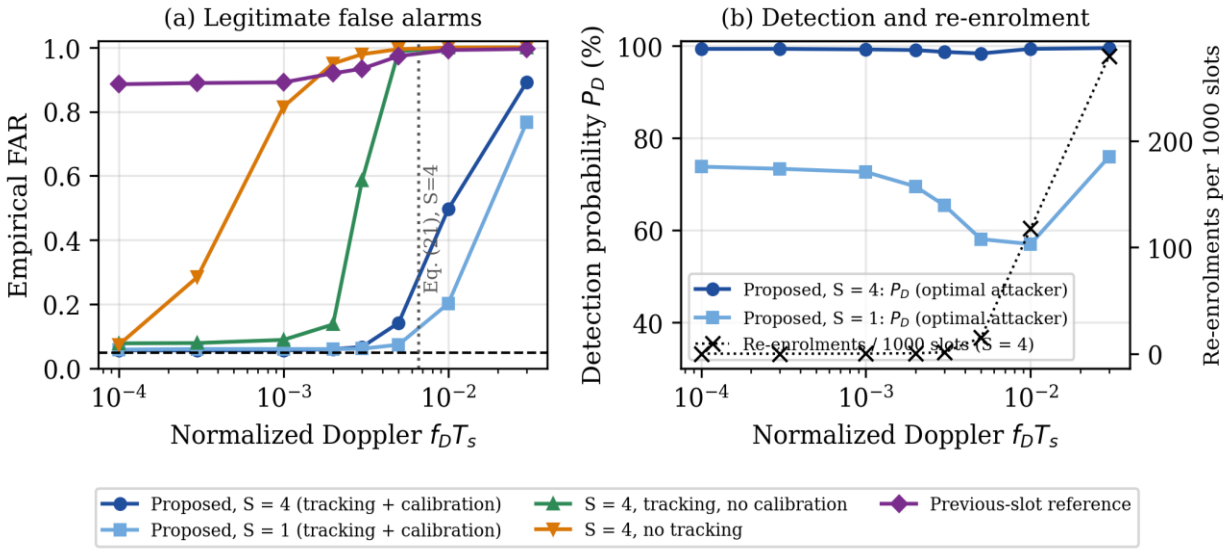


Figure 2. Channel coherence versus hopping rate (N = 256, SNR_d = 25 dB, 108,000 slots per point). (a) Empirical legitimate FAR versus $f_D T_s$ for the proposed scheme (S = 4 and S = 1), for tracking without calibration, without tracking, and for a previous-slot reference; the dotted vertical line is the limit of Eq. (21) for S = 4. (b) Detection probability of the optimal non-predicting attacker and re-enrollment rate

7.3. ROC Curves — Fig. 3

Fig. 3 shows ROC curves for $T_{obs} = 500$ over 20,000 attack slots per scheme. The proposed scheme with $S = 4$ achieves AUC = 0.9988 ± 0.0003 and $P_D = 99.50\%$ (95% interval 99.39–99.59%) at $P_{FA} = 0.05$ against the optimal attacker A1*, in agreement with the analytical ROC of (17)–(18). With a single pilot symbol ($S = 1$) the AUC falls to 0.920 ± 0.003, showing that the multi-symbol block is necessary at realistic link budgets. A CSPRNG-indexed codebook with $M = 16$ reaches AUC = 0.974 ± 0.002 under the k-means dictionary attack A2, but its ROC saturates at the $1 - (1 - \alpha)/M$ level predicted by (23). The static RIS and the PRBS31-indexed codebook both yield AUC ≈ 0.50, that is, complete failure.

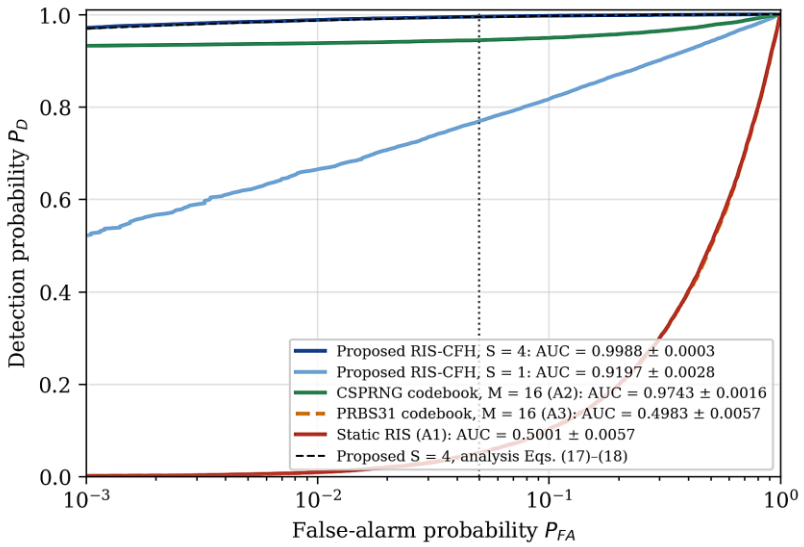


Figure 3. ROC curves under the adaptive adversary (T_{obs} = 500, N = 256, SNR_d = 25 dB, 20,000 attack slots per curve). AUC values are given with Hanley–McNeil 95% intervals. The dashed black curve is the analysis of Eqs. (17)–(18) for S = 4; the vertical dotted line marks P_{FA} = 0.05. The PRBS31 curve (dashed) coincides with the static-RIS curve

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

7.4. Spoofing Probability versus Observation Length — *Fig. 4*

Fig. 4 plots the per-slot spoofing probability at $\alpha = 0.05$ against $T_{obs} \in [5,5000]$ with 10,000 attack slots per point. The static RIS converges to $1-\alpha$ as predicted by (22). The PRBS31-indexed codebook remains at the $1/M$ guessing level for $T_{obs} < 16$ slots, when fewer than 62 bits are available, and jumps to 94.9% at $T_{obs} = 20$, after which Berlekamp–Massey prediction succeeds in 99.5% of the realizations at $T_{obs} = 20$ and in all realizations for $T_{obs} \geq 50$. The CSPRNG codebook with $M = 16$ approaches the $(1-\alpha)/M$ floor of (23) as the k-means dictionary improves. The proposed scheme stays between 0.29% and 0.65% for all T_{obs} under both A1 and A1*, fluctuating around the Proposition 1 value of 0.55% within the confidence intervals: additional observations do not increase Eve's success probability, and the mean-replay attacker A1 converges to the optimal attacker A1* as T_{obs} grows, as predicted in *section 4.2*.

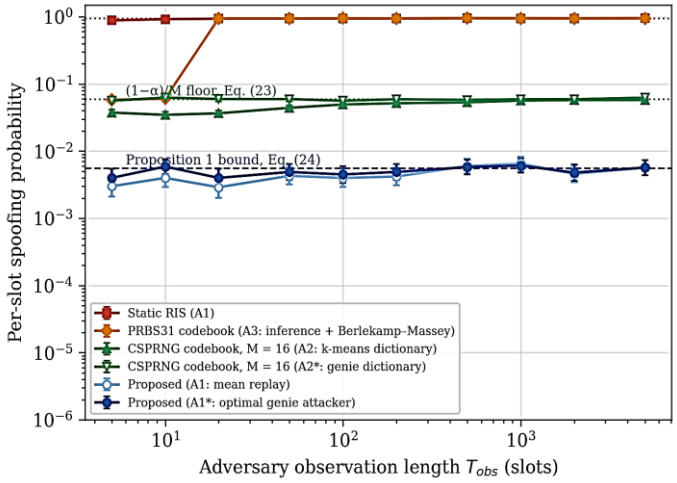


Figure 4. Per-slot spoofing probability versus adversary observation length (N = 256, M = 16, S = 4, SNR_d = 25 dB, 10,000 attack slots per point). Error bars are 95% Wilson intervals. Horizontal lines: the $(1-\alpha)/M$ floor of Eq. (23), the bound of Proposition 1, Eq. (24), and the static limit $1-\alpha$ of Eq. (22).

7.5. Scaling with the Number of RIS Elements — *Fig. 5*

Fig. 5 reassesses the N-scaling with the propagation model (1)–(3). With path loss included, the spoofing probability of the optimal attacker decreases polynomially with N, as stated in *eq. (25)*, and the simulated points agree with *eq. (18)* across three geometries. At $d_{AR} = 2$ m and $S = 4$ it falls from 0.69 at $N = 16$ to 5.8×10^{-3} at $N = 256$ and to about 10^{-4} at $N = 1024$, whereas at $d_{AR} = 4$ m an array of more than 400 elements is needed to reach P_{FA}-level spoofing. With $S = 1$ the decay is much slower (0.07 at $N = 1024$). The dashed curve reproduces the element-normalized model without path loss, which predicts spoofing probabilities below 10^{-3} even for $N = 16$; the comparison shows that that normalization overstates security by several orders of magnitude.

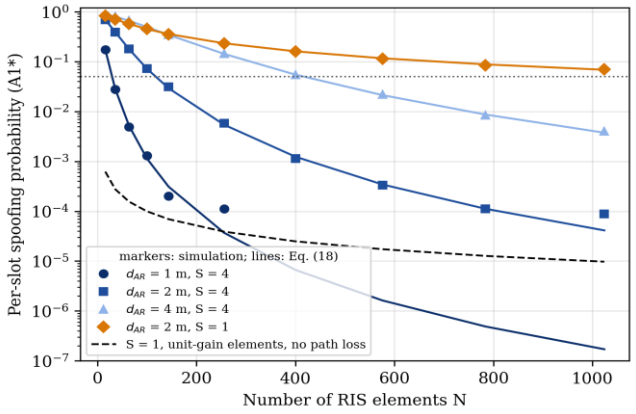


Figure 5. Per-slot spoofing probability of the optimal non-predicting attacker versus N (SNR_d = 25 dB, 45,000 attack slots per point). Markers: simulation (points with zero events omitted); solid lines: Eq. (18) averaged over independent channel realizations; dashed: S = 1 with unit-gain elements and no path loss; dotted horizontal line: α = 0.05

7.6. Session-Level Performance — Fig. 6

Fig. 6 applies the k-of-W rule of Corollary 2 with $k = \lceil 0.75W \rceil$. For the proposed scheme, a window of $W = 4$ slots reduce the session spoofing probability from 5.0×10^{-3} to 5.0×10^{-7}, while the legitimate session rejection is 1.4%; with $W = 8$ these become 4×10^{-13} and 0.58%, respectively. The codebook scheme with $M = 16$ requires about twice the window length to reach comparable spoofing levels because of its $(1-\alpha)/M$ per-slot floor. The simulated legitimate rejection agrees with (26), which supports the independence assumption.

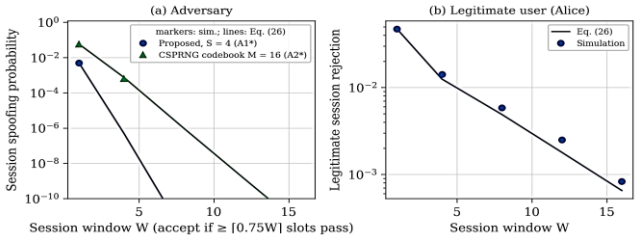


Fig. 6. Session-level authentication with the k-of-W rule, k = ⌈0.75W⌉ (N = 256, S = 4, 400 realizations of 50 consecutive slots). (a) Session spoofing probability for the proposed scheme (A1*) and the CSPRNG codebook with M = 16 (A2*). (b) Legitimate session rejection. Markers: simulation; lines: Eq. (26)

7.7. Overhead

The index-information ratio $\log_2(M)/\log_2(64)$ used previously is not a communication overhead and has been removed. Table III instead separates over-the-air signalling, pilot, enrollment, storage, generator, computation and synchronization costs for $N = 256$, $b = 1$, $S = 4$, $M = 16$ and $L_g = L_e = 20$. RIS-CFH requires no over-the-air signalling per slot; its costs are S pilot symbols per slot, a one-time enrollment of $L_g(N+1) = 5140$ pilots, $NbS = 1024$ generator bits per slot (eight AES blocks) and $\mathcal{O}(SN)$ complex multiply-accumulate operations per slot for prediction and tracking, about 2000 for the baseline, which is negligible for an access point.

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

Table 3. Overhead Definition and Comparison

Cost item	Static RIS	Codebook hopping (M = 16)	Proposed RIS-CFH
Over-the-air signalling per slot	0	0	0
Pilot symbols per slot	S = 4	S = 4	S = 4
One-time enrollment pilots	S·L_e = 80	M·S·L_e = 1280	L_g(N+1) = 5140
Storage at Bob	S complex refs (256 bit)	M·S refs (4096 bit) + seed	N+1 coefficients (16.4 kbit) + seed
Storage at RIS controller	N·b = 256 bit	M·S·N·b = 16.4 kbit + seed	128-bit seed + DRBG state
Generator output per slot	0	$\log_2 M$ = 4 bit	S·N·b = 1024 bit (8 AES blocks)
Computation at Bob per slot	O(S)	O(S)	O(S(N+1)) ≈ 2×10³ complex MACs
RIS switching events per slot	0	S	S
Synchronization	None	Shared slot counter	Shared slot counter
Per-slot spoofing floor (T_obs → ∞)	1 − α	(1 − α)/M	Eq. (24), no floor

7.8. Sensitivity Analysis

Table IV varies one parameter at a time around the baseline, with 10,000 legitimate and 10,000 attack slots per row. The FAR remains at the 5% target in every row. The spoofing probability of the proposed scheme is essentially insensitive to phase resolution, K_R, spatial correlation and Eve's position, because it depends only on $\bar{\rho}$ and on the unpredictability of the configurations; it is governed by the link budget, rising to 52% at $\mathrm{SNR}_d = 15$ dB (where $\bar{\rho} \approx 0$ dB) and falling below 10^{-4} at 35 dB. Coarse enrollment ($L_g = 5$) slightly increases spoofing through a larger σ_T^2. The PRBS31 codebook is broken in every configuration in which Eve's state classification is accurate, and remains highly vulnerable even at $\mathrm{SNR}_d = 15$ dB. The codebook scheme stays at its $(1 - \alpha)/M$ floor whenever $\bar{\rho}$ is large.

Table 4. Sensitivity Analysis (Per-Slot, α = 0.05, S = 4, M = 16, T_obs = 500)

Configuration	$\bar{\rho}$ (dB)	FAR (%) [95% CI]	Proposed P_spoof, A1* (%) [95% CI]	Eq. (18) (%)	Codebook M=16, A2* (%)	PRBS31, A3 (%)	Eve state acc. (%)
Baseline	10.1	5.06% [4.65, 5.51]	0.37% [0.27, 0.51]	0.54	6.07	94.8	99.9
b = 2 bits	10.1	5.15% [4.73, 5.60]	0.57% [0.44, 0.74]	0.54	6.08	94.8	99.9
Continuous phase	10.1	4.71% [4.31, 5.14]	0.58% [0.45, 0.75]	0.54	5.99	95.4	100.0
K_R = 0 (Rayleigh)	10.1	5.06% [4.65, 5.51]	0.56% [0.43, 0.73]	0.59	6.08	94.8	99.9
K_R = 10	10.1	5.06% [4.65, 5.51]	0.38% [0.28, 0.52]	0.54	6.07	94.8	100.0
Uncorrelated elements	10.1	5.06% [4.65, 5.51]	0.51% [0.39, 0.67]	0.54	6.05	94.8	100.0
L_g = 5 (coarse CSI)	9.6	4.94% [4.53, 5.38]	0.60% [0.47, 0.77]	0.82	6.17	94.8	99.9
L_g = 100 (fine CSI)	10.3	5.25% [4.83, 5.70]	0.29% [0.20, 0.42]	0.48	6.07	94.9	99.9
SNR_d = 15 dB	0.1	4.97% [4.56, 5.41]	52.13% [51.15, 53.11]	53.04	29.79	88.8	80.4
SNR_d = 35 dB	20.1	5.06% [4.65, 5.51]	0 (< 3.8e-04)	1.2e-04	6.03	94.8	100.0
Eve 2 m from RIS	10.1	5.06% [4.65, 5.51]	0.37% [0.27, 0.51]	0.54	6.07	94.8	100.0
Eve 15 m from RIS	10.1	5.06% [4.65, 5.51]	0.37% [0.27, 0.51]	0.54	6.07	94.8	89.6

FOREX Publication
Open Access | Rapid and quality publishing
International Journal of Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

8. DISCUSSION

The results corroborate the analysis of *sections 4* and *5*. The decisive observation is the behaviour with the observation length in Fig. 4: the static RIS converges to $1-\alpha$ spoofing, a linear PRBS generator is broken after a few tens of slots, a finite codebook converges to its $(1-\alpha)/M$ floor, and only the proposed scheme, which never reuses a configuration and derives it from a CSPRNG, keeps a spoofing probability that does not grow with T_{obs}. This property follows from Proposition 1, whose assumptions A1–A5 are stated explicitly; in particular the result does not require Eve to be ignorant of the channel, the geometry or the generator, but it does require that she cannot sense the configuration of a pilot symbol before committing her own transmission.

Compared with Table I, the distinctive aspects of RIS-CFH are (a) a state-specific predictive reference that makes the legitimate null distribution exact under hopping, (b) an explicit adaptive threat model that includes genie channel knowledge, clustering and sequence prediction, and (c) an analytical spoofing bound. The learning-based schemes [5], [6], [14] were not re-implemented because they target different threat models, and a quantitative comparison with them remains future work. The improvement of the proposed scheme over the strongest implemented baseline (CSPRNG codebook with $M = 16$) is an AUC of 0.9988 against 0.9743 and a per-slot spoofing probability of 0.50% against 5.6%, both with non-overlapping 95% confidence intervals.

The practical limits are equally clear. Security depends on $\bar{\rho}$, that is, on the cascaded power relative to the residual variance, which in turn depends on the RIS–device distances through (1); a nearby RIS, a sufficiently large N and a multi-symbol pilot block are needed, and at $\mathrm{SNR}_d = 15$ dB the baseline geometry offers little protection (Table IV). Tracking requires the coherence condition (21), which restricts the scheme to static or slowly moving devices unless S is increased. Active RIS [20] would add amplifier noise to σ_T^2; the calibration rule (20) absorbs it at the cost of a smaller $\bar{\rho}$. Finally, the CTR_DRBG seed can be established and refreshed with PL-SKG schemes such as [19], and a joint design with semantic communications [18], in which the configuration sequence also carries task-related information, is a direction for future work.

9. CONCLUSION

This paper proposed RIS-CFH, a PLA framework in which the RIS configuration of every pilot symbol is generated by a CSPRNG shared by Alice, Bob and the RIS controller, and in which Bob authenticates each slot against the fingerprint predicted for the current secret configuration from an enrolled and tracked estimate of the direct and cascaded channels. With this reference, the legitimate test statistic follows $\mathrm{Gamma}(S, 1)$ exactly (Eq. (12)), the Neyman–Pearson threshold is given by Eq. (14), the AUC by Eq. (19), and online calibration (Eq. (20)) keeps the false-alarm rate at its target under channel ageing within the coherence condition of Eq. (21). Against any non-predicting adversary with genie channel knowledge and unlimited observations, the per-slot spoofing probability is bounded by Eq. (24) and decreases as $\mathcal{O}(N^{-S})$ at fixed transmit power (Corollary 1), while a finite codebook has an irreducible $(1-\alpha)/M$ floor and a linear PRBS generator is broken by Berlekamp–Massey synthesis.

Under realistic 28 GHz propagation with path loss, spatial correlation, 1-bit phases, estimation error and Doppler ($N = 256$, $S = 4$, $\mathrm{SNR}_d = 25$ dB), Monte Carlo simulation gave an empirical FAR of 4.7–5.6% for a 5% target, an AUC of 0.9988 $\pm$ 0.0003 and a per-slot spoofing probability of 0.50% that does not increase with the adversary's observation length; a 3-of-4 session rule lowers spoofing to about 5×10^{-7} with 1.4% legitimate session rejection. These results were obtained by simulation under the stated assumptions and should be confirmed experimentally. Future work includes hardware validation with a 1-bit RIS prototype, multi-antenna and multi-user extensions, hardware impairments such as phase errors and timing jitter, and co-design with semantic communications [18] and secret key generation [19].

REFERENCES

[1] Y. Liu *et al.*, "Reconfigurable intelligent surfaces: Principles and opportunities," *IEEE Commun. Surveys Tuts.*, vol. 23, no. 3, pp. 1546–1577, 3rd Quart. 2021.

[2] S. Zhang, W. Huang, and Y. Liu, "A systematic survey on physical layer security oriented to reconfigurable intelligent surface empowered 6G," *Comput. Secur.*, vol. 148, Art. no. 104100, 2025.

[3] Khalid, W., Rehman, M. A. U., Van Chien, T., Kaleem, Z., Lee, H., & Yu, H. (2023). Reconfigurable intelligent surface for physical layer security in 6G-IoT: Designs, issues, and advances. *IEEE Internet of Things Journal*, *11*(2), 3599-3613.

[4] S. Zhang, D. Zhu, and Y. Liu, "Artificial intelligence empowered physical layer security for 6G: State-of-the-art, challenges, and opportunities," *Comput. Netw.*, vol. 242, Art. no. 110255, 2024.

[5] Meng, R., Zhu, F., Cheng, X., Xu, X., Wang, B., Dong, C.,& Zhang, P. (2025). Efficient Gaussian process classification-based physical-layer authentication with configurable fingerprints for 6G-enabled IoT. *IEEE Transactions on Consumer Electronics*, *71*(2), 2603-2617.

[6] N. Gao, Q. Huang, C. Li, S. Jin, and M. Matthaiou, "EsaNet: Environment semantics enabled physical layer authentication," *IEEE Wireless Commun. Lett.*, vol. 13, no. 1, pp. 178–182, Jan. 2024.

[7] H. Amin, W. Aman, and S. Al-Kuwari, "On the potential of re-configurable intelligent surface (RIS)-assisted physical layer authentication (PLA)," 2024, *arXiv:2405.00426*.

[8] H. Amin, W. Aman, S. Al-Kuwari, and M. Qaraqe, "Enhancing the performance of physical layer authentication through RIS in wireless networks," *Comput. Netw.*, Art. no. 112174, 2026.

[9] K.-L. Besser and E. A. Jorswieck, "Reconfigurable intelligent surface phase hopping for ultra-reliable communications," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 9082–9095, Nov. 2022.

[10] Z. Wei, B. Li, and W. Guo, "Adversarial reconfigurable intelligent surface against physical layer key generation," *IEEE Trans. Inf. Forensics Security*, vol. 18, pp. 2368–2381, 2023.

[11] Z. Wei, W. Hu, J. Zhang, W. Guo, and J. A. McCann, "Explainable adversarial learning framework on physical layer key generation combating malicious reconfigurable intelligent surface," *IEEE Trans. Wireless Commun.*, vol. 24, no. 4, pp. 3529–3545, Apr. 2025.

[12] N. Gao *et al.*, "RIS-assisted wireless link signatures for specific emitter identification," *IEEE Trans. Wireless Commun.*, vol. 23, no. 12, pp. 17872–17883, Dec. 2024.

International Journal of
Electrical and Electronics Research (IJEER)
Research Article | Volume 14, Issue 3 | Pages 882-894 | e-ISSN: 2347-470X

[13] W. Shi, J. Xu, W. Xu, C. Yuen, A. L. Swindlehurst, and C. Zhao, "On secrecy performance of RIS-assisted MISO systems over Rician channels with spatially random eavesdroppers," *IEEE Trans. Wireless Commun.*, vol. 23, no. 8, pp. 8357–8371, Aug. 2024.

[14] Yan, Y., Jing, T., Gao, Q., Huo, Y., Fan, X., Wang, Y., & Wang, Y. (2025). An initial phase-added and RIS-assisted physical layer security scheme based on deep reinforcement learning. *IEEE Transactions on Vehicular Technology*, *74*(10), 15991-16006.

[15] E. O. Frimpong, B.-H. Oh, T. Kim, and I. Bang, "Physical-layer security with irregular reconfigurable intelligent surfaces for 6G networks," *Sensors*, vol. 23, no. 4, Art. no. 1881, 2023.

[16] M. Di Renzo *et al.*, "Smart radio environments empowered by reconfigurable AI meta-surfaces: An idea whose time has come," *EURASIP J. Wireless Commun. Netw.*, vol. 2019, no. 1, Art. no. 129, 2019.

[17] Liu, M., Li, X., Ning, B., Huang, C., Sun, S., & Yuen, C. (2022). Deep learning-based channel estimation for double-RIS aided massive MIMO system. *IEEE Wireless Communications Letters*, *12*(1), 70-74.

[18] J. Dai, J., Wang, S., Yang, K., Tan, K., Qin, X., Si, Z., ... & Zhang, P. (2023). Toward adaptive semantic communications: Efficient data transmission via online learned nonlinear transform source-channel coding. *IEEE Journal on Selected Areas in Communications*, *41*(8), 2609-2627.

[19] Wu, H., Fang, Y., Li, N., Yuan, X., Wei, Z., Nan, G., & Tao, X. (2025). Secret key generation with untrusted internal eavesdropper: Token-based anti-eavesdropping. *IEEE Transactions on Information Forensics and Security*, *20*, 2523-2537.

[20] Zhang, Z., Dai, L., Chen, X., Liu, C., Yang, F., Schober, R., & Poor, H. V. (2022). Active RIS vs. passive RIS: Which will prevail in 6G?. *IEEE Transactions on Communications*, *71*(3), 1707-1725.

[21] L. Xiao, L. J. Greenstein, N. B. Mandayam, and W. Trappe, "Using the physical layer for wireless authentication in time-variant channels," *IEEE Trans. Wireless Commun.*, vol. 7, no. 7, pp. 2571–2579, Jul. 2008.

[22] Tang, W., Chen, M. Z., Chen, X., Dai, J. Y., Han, Y., Di Renzo, M., ... & Cui, T. J. (2020). Wireless communications with reconfigurable intelligent surface: Path loss modeling and experimental measurement. *IEEE transactions on wireless communications*, *20*(1), 421-439.

[23] Björnson, E., & Sanguinetti, L. (2020). Rayleigh fading modeling and channel hardening for reconfigurable intelligent surfaces. *IEEE Wireless Communications Letters*, *10*(4), 830-834.

[24] Jakes, W. C., & Cox, D. C. (Eds.). (1994). *Microwave mobile communications*. Wiley-IEEE press.

[25] Barker, E. B., & Kelsey, J. M. (2007). *Recommendation for random number generation using deterministic random bit generators (revised)* (pp. 800-900). Washington, DC, USA: US Department of Commerce, Technology Administration, National Institute of Standards and Technology, Computer Security Division, Information Technology Laboratory.

[26] J. L. Massey, "Shift-register synthesis and BCH decoding," *IEEE Trans. Inf. Theory*, vol. IT-15, no. 1, pp. 122–127, Jan. 1969.

[27] T. W. Anderson, "The integral of a symmetric unimodal function over a symmetric convex set and some probability inequalities," *Proc. Amer. Math. Soc.*, vol. 6, no. 2, pp. 170–176, Apr. 1955.

[28] J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology*, vol. 143, no. 1, pp. 29–36, Apr. 1982.

© 2026 by the Mohammed Aboud Kadhim, Murooj Nadhom, and Ameer Abdul Sattar Jaber. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0/).