# FadePrint - Satellite Spoofing Detection via Fading Fingerprinting

Gabriele Oligeri*, Savio Sciancalepore†, Alireza Sadighian*
*Division of Information and Computing Technology (ICT)
College of Science and Engineering (CSE), Hamad Bin Khalifa University (HBKU), Doha, Qatar
{sadighian, goligeri}@hbku.edu.qa
† Eindhoven University of Technology, and Eindhoven Artificial Intelligence Systems Institute (EAISI), Netherlands.
s.sciancalepore@tue.nl

***Abstract*—[1] While various methods exist to implement message authentication in different communication layers, the physical layer offers some unique and beneficial features for this purpose. Existing solutions authenticate transmitters at the physical layer by merging deep learning with physical-layer attributes, protecting against impersonation attacks. This approach requires a lengthy and resource-intensive training phase for every new transmitter that joins the network. However, for some scenarios (e.g. satellite communications), characterizing the channel experienced by the received signal might be effective in detecting impersonation. In this work, we propose FadePrint, a solution capable of detecting satellite spoofing attacks by fingerprinting the noise-fading process associated with the satellite communication channel. The fading characteristics of a satellite link differ significantly from terrestrial links (e.g., indoor), making it possible to distinguish between the two. Unlike other systems, FadePrint does not require retraining when new transducers are added to the network. We tested FadePrint with real satellite and indoor radio measurements and proved that FadePrint can effectively discriminate between a satellite transmitter and a fake indoor one, with an accuracy higher than 0.99 for all the considered configurations.**

***Index Terms*—Physical-Layer Security, Applications of AI for Security, Wireless Security.**

## I. Introduction

Spoofing is a malicious activity where an adversary transmits radio messages with a forged source identifier, thus posing significant challenges due to the broadcast nature of the radio spectrum [1]. It is especially effective in the satellite scenario, where transmitters broadcast messages that are neither encrypted nor authenticated, e.g., Global Positioning System (GPS). The insecurity of such communication technologies makes them prime targets for spoofers, who can use tools like Software Defined Radios (SDRs) and publicly accessible software to generate rogue signals [2], [3]. Factors such as predictable satellite trajectories, weak ground-level signal reception, potential absence of message authentication, and vulnerabilities in authentication secrets make satellites especially susceptible to this type of attack. Recent solutions for spoofing detection span from the application layer to the Physical Layer (PHY) and often involve shared secrets between communicating devices. A notable approach is using Artificial Intelligence (AI) to identify unique radio transducer fingerprints at the Physical Layer (PHY), which inherently counteracts spoofing [4]. However, while effective, PHY fingerprinting requires receivers to create a model for each transmitter, which is not feasible in scenarios with countless potential transmitters, e.g., satellite communication systems. However, a promising method of detecting spoofing in satellite contexts could be to identify the "type of link" experienced by the received signal. In particular, the noise and fading patterns of a satellite link differ drastically from those of a terrestrial link due to longer communication distances and fewer obstructions [5]. Therefore, determining if a signal originates from a terrestrial link, instead of a satellite link, could be an effective way to identify spoofers.

**Contribution.** In this paper, we introduce *FadePrint*, a technique to detect satellite spoofing, leveraging the PHY fingerprinting of the fading process that affects the communication link. Essentially, *FadePrint* generates a model that represents the expected pattern of a genuine satellite signal at the PHY level. During real-time operations, this model is used to discriminate the fading process of incoming signals, thereby pinpointing any terrestrial spoofing attempts. We tested *FadePrint* using real satellite data from the IRIDIUM satellite constellation and a set of real spoofing attacks, preliminarily carried out in indoor scenarios. We achieved very promising spoofing detection accuracy, larger than $0.99$. Compared to current solutions, *FadePrint* does not require the generation and management of multiple transmitter models, making it lightweight and scalable by design.

**Roadmap.** The paper is organized as follows. Section II summarizes related work, Section III presents the scenario and adversary model, Section IV introduces the preliminaries and the dataset, Section V provides the details of *FadePrint*, Section VI provides the performance assessment, and finally, Section VII tightens the conclusions and outlines future work.

## II. Related Work

Spoofing and replay (meaconing) of satellites' messages is increasingly becoming an actual threat [6], [7]. Several recent contributions investigated PHY fingerprinting in many different forms and with many different objectives, also with

reference to GPS satellites [8]. Some work already noticed the correlation between the raw received messages and the experienced channel. To name a few, Shawabkha et al. [9] analyzed the impact of the wireless channel on the PHY authentication process, identifying the undesirable effects of fading on device identification. A few works use plaintext or side-information to infer traffic exchanged in a communication link. For example, Trinh et al. [10] classified mobile traffic using radio-link data, despite encryption. However, this fingerprinting applies to traffic features. Location-based fingerprinting has also been used in the avionics domain to authenticate aircraft, e.g., by using the Channel Impulse Response (CIR) [11], and carrier phase [12]. However, as reported by Strohmeier et al. [13], these approaches require the cooperation of the transmitter. On the contrary, our approach is completely opportunistic. Finally, other related works, such as Calvo et al. [14] and Tim et al. [15] leveraged PHY-layer characteristics for GNSS technologies. However, the proposed metrics are specific to the application scenarios and cannot be applied in other contexts, e.g., any (LEO) satellite constellation.

## III. Scenario and Adversary Model

Our reference scenario includes three entities: (i) a *satellite transmitter*, emitting Radio Frequency (RF) signals; (ii) a *User*, featuring a satellite receiver; and finally, (iii) an *Adversary*, able to perform *replay* and *spoofing* attacks. Users can receive satellite messages from both the legitimate satellite network and malicious terrestrial ground stations. To distinguish between these, we analyze the fading process that affects the received signal. Specifically, the satellite link features a Line-of-Sight (LoS) connection characterized by a low Signal-to-Noise Ratio (SNR) due to the large distance between the transmitter and receiver. In contrast, an indoor terrestrial link experiences a more complex fading process caused by obstructions (shadowing) and reflections (multipath). It is important to note that all bit sequences referenced in this study originate from genuine satellite communications [4]. This approach helps us emulate actual replay Attacks from adversarial transmitters.

**Adversary Model.** We consider an indoor spoofing transmitter that sends radio messages pretending to be a legitimate satellite transmitter. We assume that all authentication mechanisms (if present) are compromised, and the messages from terrestrial sources become indistinguishable from genuine satellite messages at any layer higher than the PHY.

## IV. Background and Measurements

The main use case considered in this paper is the IRIDIUM satellite constellation, initially operational in 1993 under Motorola and later revamped by Thales, rebranded as Iridium NEXT with updated satellites and technology. The IRIDIUM system comprises 66 operational satellites in Low Earth Orbit (LEO) at roughly 800 km above Earth, moving at an approximate speed of 7 km/s [2]. These satellites transmit within the 1,616 - 1,626.5 MHz range, and users can receive these signals using specialized equipment from companies such as Kyocera and Motorola. Our study employs a dataset from [4], containing 589 hours (24 days) of IRIDIUM Ring Alert (IRA) message recordings, totaling 102,318,546 physical-layer I-Q sample data (averaging 1,550,281 samples per satellite). Each IRA packet log includes the reception timestamp, satellite and beam ID, the satellite location details, and the raw I-Q samples. For a comprehensive understanding of the I-Q samples and the fingerprinting process at the physical layer, the reader can refer to [4].

## V. Methodology

Our methodology begins by converting I-Q sample batches into images and then applying anomaly detection through a Fully Convolutional Data Description (FCDD) network. As indicated by Liznerski et al. [16], FCDD is particularly suitable for detecting image anomalies, surpassing other top-tier methods. I-Q samples, sourced from an SDR, are represented by complex numbers: the real part represents the in-phase component (**I**), while the imaginary is the quadrature component (**Q**). We compute a bi-variate histogram on I-Q sample batches made up of $N = 10,000$ samples, organizing them into bi-dimensional bins. This pre-processing phase aims at generating matrices of size $225 \times 225$ (output of the bi-variate histogram) which will be the images to be considered as the input for the convolutional neural network. In fact, each matrix element is treated as an image pixel value, ranging from $[0, 255]$, entered into the CNN. If the histogram's output exceeds the maximum pixel value of $255$, adjustments to the I-Q batch size are necessary for calibration.

*FadePrint* involves a three-stage process:

- **Training.** We train the FCDD network using images derived from satellite I-Q samples. In every scenario, the network is trained on I-Q samples from all satellites, excluding two satellites, later utilized for calibration and testing.
- **Decision threshold estimation.** In our calibration phase, we test the trained model on two datasets: one composed of I-Q samples from terrestrial indoor measurements and the other of satellite samples not considered during the training process. Each evaluated image receives an *anomaly index*, enabling its categorization as either *normal* (derived from satellite link I-Q samples) or *anomaly* (originating from indoor terrestrial I-Q samples). After assessing all anomaly indices, we establish a threshold that jointly minimizes False Negativess (FNs) and False Positivess (FPs).
- **Testing.** In the testing phase, we evaluate two different image datasets from satellite and terrestrial links, calculating the anomaly indices for each image. It is important to note that the datasets considered during this phase were used in neither the training nor the calibration stages. Finally, we consider the threshold computed during the previous phase, and classify the images based on their origin, either from the satellite or the ground station, depending on the scenario.

Our data corpus includes the following datasets:

- *Satellite data.* The satellite dataset includes 66 streams of I-Q samples, i.e., one per satellite.
- *Indoor Terrestrial data.* We gathered five I-Q sample streams, each lasting 10 minutes, representing one stream for every TX-RX pair—we considered five distinct transmitting radios in each case. The indoor terrestrial data stream employs the same QPSK modulation and bits as the satellite stream, emulating genuine replay and spoofing attacks.

In accordance with the existing literature, we used three datasets: $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$ for training, calibration, and testing of the CNN, respectively. The training dataset, $\mathcal{X}$, comprises randomly shuffled images derived from I-Q samples of 64 of the 66 available satellite streams. The calibration dataset, $\mathcal{Y} = Y_S, Y_T$, is made up of two image sets: $Y_S$, created from a satellite stream not included in $\mathcal{X}$ and $Y_T$, derived from terrestrial measurements. The testing dataset is defined as $\mathcal{Z} = Z_S, Z_T$, where $Z_S$ contains images from a satellite stream, different from those in $\mathcal{X}$ and $\mathcal{Y}$, and $Z_T$ is a collection of images from the indoor scenario, different from those in $\mathcal{Y}$. Finally, note that we compute the threshold (output of the calibration process) by computing the anomaly indexes on the images coming from both the satellite and the indoor measurements, and then we apply the following Eq. 1.

$$thr = E[max(I_S), min(I_T)], \tag{1}$$

where $I_S$ and $I_T$ are the anomaly indexes associated with the images of the satellite and terrestrial datasets, respectively. This approach of setting the threshold—by taking the average of the highest anomaly score from the satellite images and the lowest anomaly score from the terrestrial images—strikes an optimal balance between False Positives (FP) and False Negatives (FN). It effectively minimizes the cumulative value of both, ensuring better accuracy.

## VI. Performance Evaluation

In this section, we provide the details of the performance evaluation of *FadePrint*. Our indoor test environment considers a typical office space. The setting was specifically chosen to represent a common indoor environment where radio transmissions might occur. Within this environment, the five transmitting devices were set up at a distance of 30 meters from the receiving device. In particular, there was no direct line of sight (NLoS) between the transmitters and the receiver—a condition that often exists in real-world indoor scenarios and can have significant effects on signal propagation. For our testing setup, the receiver was stationed inside an office room, whereas the transmitters were placed in an adjoining corridor that experienced frequent movement of people. This movement introduces additional variables to our testing, mimicking potential real-world interference such as the Doppler effect caused by moving objects. Our measurements were comprehensive. Each transmitter-receiver (TX-RX) pair was considered in individual sessions lasting 10 minutes, leading to a combined total of 50 minutes of measurements for all pairs. Each session yielded over 146 million I-Q samples, providing us with a vast dataset to analyze and draw insights. Figure 1 provides a graphical representation of our indoor testing setup. It outlines the floor plan of our office environment, highlighting the exact positions of both the transmitters and receiver, as well as the general movement patterns of people in the corridor. This illustration serves as a valuable reference point for understanding our experimental design. The results of our proposed solution are presented in Fig. 2, where we summarized a comprehensive overview of the system's performance. Fig. 2 is organized into three sections, which capture the three main aspects of our methodology.

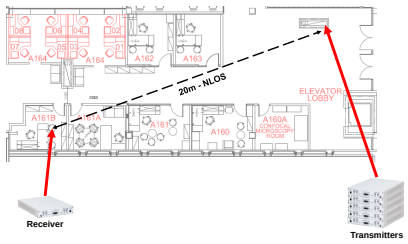


Fig. 1: Indoor scenario: the receiver is deployed 30 meters from the transmitters in an office scenario, with Non-Line-of-Sight (NLoS).

**Calibration Process.** In Fig. 2(a), we report the histograms associated with the anomaly scores computed on the calibration datasets. Blue bars depict scores derived from legitimate satellite transmissions, while the orange bars illustrate those originating from terrestrial sources, mimicking spoofing attacks. This distinct partitioning of scores between legitimate and spoofing sources highlights the system's capability to differentiate between the two. The calibration process was instrumental in determining an effective decision threshold, represented as the solid red line in Fig. 2(a). This threshold has been computed according to Eq. 1, and is equal to $thr = 2.9 \cdot 10^{-3}$.

**Testing Phase.** Once the model was calibrated, it was evaluated on a fresh dataset, distinct from the training and calibration data. Figure 2(b) illustrates the outcome of this testing phase. To ensure robustness and validate the effectiveness of the model under various conditions, the test was performed multiple times (specifically, 100 times). For each iteration, new permutations of satellite and terrestrial streams were used, effectively enhancing the validity of the results by eliminating biases related to particular satellites or terrestrial transmitters.

**Confusion Matrix.** The third section, Fig. 2(c), offers a detailed breakdown of the system's classification outcomes in the form of a confusion matrix. This matrix captures the true positives, false positives, true negatives, and false negatives to provide a granular perspective on the system's performance. In particular, we achieved an accuracy of more than 0.99 and only 19 false positives recorded out of a total of 150,370 samples; the model demonstrates outstanding performance in distinguishing between genuine satellite transmissions and

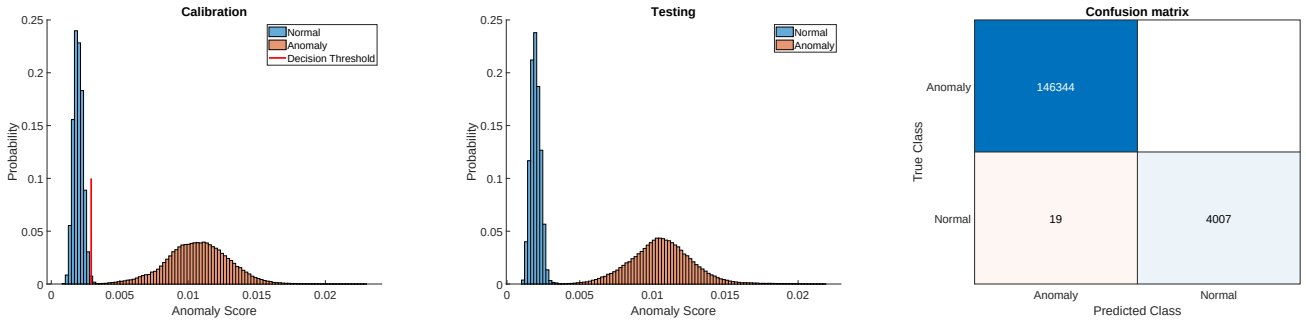


(a) (b) (c)

Fig. 2: Indoor scenario: (a) decision threshold calibration, (b) testing, and (c) resulting confusion matrix.

rogue terrestrial ones.

Our results prove the system's effectiveness in detecting spoofing attacks coming from indoor scenarios against satellite transmitters. The rigorous evaluation, both in terms of calibration and repeated testing, reaffirms its efficacy, making it a promising tool for protecting satellite communications.

## VII. Conclusion and Future Work

We have introduced *FadePrint*, a satellite spoofing detection method that leverages the unique fingerprinting of communication channel fading to distinguish between legitimate satellite transmitters and spoofers (terrestrial transmitters). Unlike existing techniques, *FadePrint* stands out as it does not necessitate retraining for each new transducer in the network. This is due to its reliance on the distinctive fading patterns inherent to satellite links, which are markedly different from terrestrial links. Our tests, which incorporated real satellite and indoor terrestrial measurements, demonstrated the effectiveness of *FadePrint*, achieving an accuracy over 0.99. Looking ahead, our future research will include outdoor and mobility scenarios.

## Acknowledgements

The authors would like to thank the anonymous Reviewers for their constructive comments on the article. This publication was made possible by the NPRP12C-0814-190012-SP165 awards from the Qatar National Research Fund (a member of Qatar Foundation). This work has also been partially supported by the INTERSECT project, Grant No. NWA.1162.18.301, funded by the Netherlands Organization for Scientific Research (NWO). The content herein is solely the responsibility of the authors.

## References

[1] Schmidt, D. et al., "A Survey and Analysis of the GNSS Spoofing Threat and Countermeasures," *ACM Comput. Surveys*, vol. 48, no. 4, pp. 1–31, 2016.

[2] G. Oligeri, S. Sciancalepore, and R. Di Pietro, "GNSS Spoofing Detection via Opportunistic IRIDIUM Signals," in *ACM Conf. on Security and Privacy in Wirel. and Mob. Netws.*, 2020, p. 42–52.

[3] S. Raponi, S. Sciancalepore, G. Oligeri, and R. Di Pietro, "Road traffic poisoning of navigation apps: Threats and countermeasures," *IEEE Security & Privacy*, vol. 20, no. 3, pp. 71–79, 2022.

[4] G. Oligeri, S. Sciancalepore, S. Raponi, and R. Di Pietro, "PAST-AI: Physical-layer Authentication of Satellite Transmitters via Deep Learning," *IEEE Trans. on Informat. Forensics and Security*, vol. 18, pp. 274–289, 2022.

[5] Vinogradov, E. et al., "Tutorial on UAV: A blue sky view on wireless communication," *arXiv preprint arXiv:1901.02306*, 2019.

[6] G. Oligeri, S. Sciancalepore, O. A. Ibrahim, and R. Di Pietro, "GPS Spoofing Detection via Crowd-Sourced Information for Connected Vehicles," *Computer Networks*, vol. 216, p. 109230, 2022.

[7] G. Oligeri, S. Sciancalepore, O. A. Ibrahim, and R. Di Pietro, "Drive me not: GPS spoofing detection via cellular network: (architectures, models, and experiments)," in *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*, 2019, pp. 12–22.

[8] Foruhandeh, M. et al., "Spotr: GPS Spoofing Detection via Device Fingerprinting," in *ACM Conf. on Security and Privacy in Wirel. and Mob. Netwks.*, 2020, pp. 242–253.

[9] Al-Shawabka, A. et al., "Exposing the fingerprint: Dissecting the impact of the wireless channel on radio fingerprinting," in *INFOCOM 2020*. IEEE, 2020, pp. 646–655.

[10] Trinh, D. et al., "Mobile Traffic Classification Through Physical Control Channel Fingerprinting: A Deep Learning Approach," *IEEE Trans. on Netw. and Serv. Managem.*, vol. 18, no. 2, pp. 1946–1961, 2021.

[11] Zhang, J. et al., "Mobility assisted secret key generation using wireless link signatures," in *IEEE INFOCOM*, 2010, pp. 1–5.

[12] Wang, Q. et al., "Fast and scalable secret key generation exploiting channel phase randomness in wireless networks," in *IEEE INFOCOM*, 2011, pp. 1422–1430.

[13] Strohmeier, M. et al., "Security of ADS-B: State of the Art and Beyond," *arXiv preprint arXiv:1307.3664*, 2013.

[14] Calvo-Palomino, R. et al., "Short: LSTM-based GNSS Spoofing Detection Using Low-cost Spectrum Sensors," in *Int. Symp. on "A World of Wirel., Mob. and Multimedia Netwks."*, 2020, pp. 273–276.

[15] Sun, C. et al., "Robust Spoofing Detection for GNSS Instrumentation Using Q-Channel Signal Quality Monitoring Metric," *IEEE Trans. on Instrumentat. and Measurem.*, vol. 70, pp. 1–15, 2021.

[16] Liznerski, P. et al., "Explainable Deep One-Class Classification," in *Int. Conf. on Learn. Representat.*, 2021.