

# OrbID: Identifying Orbcomm Satellite RF Fingerprints

Cédric Solenthaler  
ETH Zurich  
csolenthaler@ethz.ch

Joshua Smailes  
University of Oxford  
joshua.smailes@cs.ox.ac.uk

Martin Strohmeier  
armasuisse Science + Technology  
martin.strohmeier@ar.admin.ch

**Abstract**—An increase in availability of Software Defined Radios (SDRs) has caused a dramatic shift in the threat landscape of legacy satellite systems, opening them up to easy spoofing attacks by low-budget adversaries. Physical-layer authentication methods can help improve the security of these systems by providing additional validation without modifying the space segment. This paper extends previous research on Radio Frequency Fingerprinting (RFF) of satellite communication to the Orbcomm satellite formation. The GPS and Iridium constellations are already well covered in prior research, but the feasibility of transferring techniques to other formations has not yet been examined, and raises previously undiscussed challenges.

In this paper, we collect a novel dataset containing 8992474 packets from the Orbcomm satellite constellation using different SDRs and locations. We use this dataset to train RFF systems based on convolutional neural networks. We achieve an ROC AUC score of 0.53 when distinguishing different satellites within the constellation, and 0.98 when distinguishing legitimate satellites from SDRs in a spoofing scenario. We also demonstrate the possibility of mixing datasets using different SDRs in different physical locations.

## I. INTRODUCTION

### A. Motivation

Over the past decades Software Defined Radios (SDRs) have proliferated and in the process have invalidated previously sensible wireless threat model assumptions. Nowadays, commonly available and cheap hardware enables attacks that were, just a few decades prior, firmly in the domain of nation state actors. In fact, many such wireless attacks are now occurring in the wild, such as GPS spoofing [1], and recent research even demonstrated arbitrary code execution through wireless signal injection [2]. While these developments are likely to precipitate improved wireless security in the future, some systems, such as satellite constellations, cannot simply be replaced due to accessibility, legacy, or cost considerations.

In the domain of satellite communication security, security measures that rely on the physical layer and are deployable purely in the ground segment provide valuable tools to improve the security of legacy systems. Physical-layer

identification techniques have shown promising results for satellite identification [3], [4], [5]. These works tend to focus on the GPS and Iridium constellations due to the ease and ubiquity of receiving signals from these satellites using off-the-shelf hardware and open-source software.

In this paper, we extend existing satellite fingerprinting techniques to work with the Orbcomm satellite constellation. These satellites present a typical example of the legacy systems that RFF is suited to securing. The Orbcomm constellation brings a number of unique challenges not present in the Iridium system:

- Communication channels are low-bandwidth and differ between satellites;
- Small frequency spacing between channels prevents high levels of oversampling;
- Messages contain identifying information which, if improperly masked, could allow the model to cut corners.

### B. Contributions

In this paper we provide the first application of RF fingerprinting to the Orbcomm constellation. In order to achieve this, we collected a dataset of 8992474 samples that were collected from multiple locations and SDRs. Using this dataset we demonstrate the effectiveness of RFF to authenticate communication, and separate legitimate communication from spoofing via attacker-controlled SDRs. We also introduce a novel preprocessing technique to anonymize messages without reducing the number of samples available to the fingerprinting model. Finally, we show the stability of the fingerprints over time, and evaluate transferability between SDRs.

## II. BACKGROUND

### A. Software Defined Radios

SDRs are radio devices that digitize the received signal, offloading the bulk of (traditionally analog) signal processing to a computer. SDRs utilize IQ demodulation in order to downconvert the signal and split it into I (In-phase) and Q (Quadrature) components. The I and Q signals, together with the original carrier frequency, can be used to fully reconstruct the received signal. After the IQ signals are sampled and digitized, they are commonly represented as a complex number where the real part corresponds to the in-phase component.

We focus on two commonly used off-the-shelf SDRs.

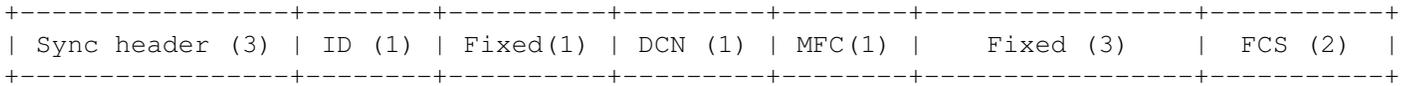


Fig. 1. Structure of an Orbcomm synchronization packet.

1) *HackRF One*: The HackRF One is a Semi-Duplex SDR that covers a frequency range from 1 MHz to 6000 MHz with a maximal bandwidth of 20 MHz and a resolution of 8 bits [6].

2) *RTL-SDR*: The RTL-SDR is a receive-only family of SDR devices that covers a frequency range from 500 kHz to 1.75 GHz with a maximal bandwidth of around 3 MHz and an effective resolution of 7 bits [7]. The RTL-SDR is notable due to its very low price point, in the range of 30 USD.

### B. Radio Frequency Fingerprinting

Radio Frequency Fingerprinting (RFF) aims to identify an emitter based on imperfections in the Electromagnetic (EM) signal [8]. These imperfections stem from variations in the fabrication of the RF chain and are thus hard to avoid. These variations can be used to discriminate signals from transmitters of the same model, and in particular can distinguish attacker-controlled transmitters from legitimate ones. This may be used to provide an extra security layer [8] against spoofing, which may be implemented independently of existing infrastructure.

There are two main approaches to RFF, which focus on either the transient of the signal (the beginning of the signal as the transmitter powers up) or the steady state during normal transmission [8]. Within each of these, some approaches use a classifier (in which all transmitters are known and messages are labeled) and some use similarity-based techniques (in which fingerprints are compared to a database) [9].

RFF is used in the literature to identify transmitters operating various protocols, including WiFi [9], [10], [11], [12], or aviation communication [9], [12], as well as various satellite systems. Among these the Iridium constellation [3], [5], [13] and Global Navigation Satellite Systems (GNSSs) [14] stand out, in addition to some study of simulated CubeSats [15].

### C. Triplet Loss

The triplet loss function [16] is a measure of similarity between sparse representations and is calculated from a triplet of embeddings. It compares the distances from an anchor to a positive sample (same label as the anchor) and to a negative sample (different label). The triplet loss characterizes how much closer the positive example is to the anchor than the negative example. Triplet loss encourages a model to map semantically close inputs to metrically close embeddings [17].

Mathematically, the triplet loss  $L$  is calculated as described in Equation 1 [16] where  $p$  is the embedding of the anchor,  $p_+$  is the embedding of the positive example, and  $p_-$  is the embedding of the negative sample. The function  $d(x, y)$  describes a distance between two embeddings  $x$  and  $y$ . The margin  $m$  is a hyperparameter that describes the minimal desired difference between embeddings of different labels. Due

to the  $\max(\cdot)$  term, differences beyond the margin do not contribute to the loss anymore.

$$L = \max(m + d(p, p_+) - d(p, p_-), 0) \quad (1)$$

### D. Orbcomm

The Orbcomm satellite constellation was established from 1995 to 2000 [18] and provides communication services. The space segment consists of 30 satellites that communicate in the 137 MHz to 150 MHz (VHF) band [18] with a right-hand circular polarization. Not all of these satellites appear to be functional [19]. In practice, only 13 channels between 137 MHz to 138 MHz are used [20], one of which is not a user channel but a gateway channel [21].

The signal is encoded with Symmetrical Differential Phase Shift Keying (SDPSK), with 8 bits to a word and 600 words in a minor frame [18]. Minor frames consist of 12 or 24 word-long packets. Each packet contains a two word Fletcher checksum [22]. Orbcomm satellites transmit continuously [19] and with a transmission rate of 4800 bps.

There are several different types of used packets. Figure 1 shows the structure of a synchronization packet that contains a synchronization header (0x65A8F9), the satellite ID, a Downlink Channel Number (DCN), a Minor Frame Counter (MFC), as well as a Fletcher checksum (FCS) and four fixed words [22]. Both the ID and the DCN field (and the checksum too since only ID and DCN change between different synchronization packets) can be used to infer the identity of the satellite.

### E. Related Work

While there is an abundance of prior work [3], [4], [5], [9], [13], [14], [15], [23] using RFF on satellite communication (real or simulated), to the best of our knowledge there is no prior work evaluating the Orbcomm constellation. However, several papers do use Orbcomm to provide opportunistic positioning [19], [20], [21].

Some previous RFF works pre-process IQ samples into an image (or multiple images) to produce heatmaps, which are used in classifiers pre-trained on image recognition models [4], [5]. Other works such as SatIQ [3] perform 1D convolutions on the raw signal, similar to conventional signal processing. We use this approach in this paper.

Another important difference is into what space the model projects the input. Some previous works use a classification approach [4], [5], [11], whereas other works [3] opt for a more extendable similarity-based approach, enabling new transmitters to be added without retraining. The methodological difference to SatIQ [3] lies in not using a synchronized



Fig. 2. Turnstile antenna for the data acquisition system deployed on a rooftop in Zurich.

signal as input for the proposed model and making more use of shared weights in processing.

### III. THREAT MODEL

Satellite systems (especially GPS) are commonly spoofed in the wild [1]. Cheap, transmit-capable SDRs are readily available (e.g. HackRF for around 300 USD [6]) and are more than sufficient to perform many attacks. An attacker with reasonably low budget (2000 USD) is able to overshadow satellite systems [24] or even take over user ground stations [2].

Thus, we can assume the adversary has the capability to overshadow the legitimate signal or replay signals from the transmitter. Due to the low power of legitimate satellite signals on the ground, the attacker is assumed to be able to overshadow a satellite despite the user potentially utilizing a highly directional antenna. Alternatively the attacker might target a sidelobe, which has been shown to be practical for small aperture (user segment) satellite dishes [2].

Compared to prior work [5], we also consider non-terrestrial adversaries, because even a low-budget attacker can use an Uncrewed Aerial Vehicle (UAV) to carry an SDR and a Single Board Computer (SBC) above a receiver and thus fool systems based on detecting terrestrial links (e.g. via fading process [5]).

The user is unable to distinguish attacker messages from legitimate ones on the logical layer (e.g. due to absent/broken cryptography or leaked key material). Thus the user must attempt to authenticate the signal on the physical layer.

A sufficient test of this capability consists of distinguishing different transmitters of the same constellation. Previous work [3] indicates that the performance against a replay attacker (over a wired link) is better than the performance in distinguishing legitimate transmitters.

Out of scope for the purpose of this work is an adversary capable of hijacking the actual satellite hardware, thus taking over the transmitter. Neither is the adversary able to inject malicious samples into our training data. Such a capability would enable the adversary to create a backdoor in the machine learning model [25].

### IV. DATA COLLECTION

Data was collected from two physical locations. Figure 2 shows a turnstile antenna designed for frequencies from 135 MHz to 152 MHz and right-hand circular polarization [26] deployed on a rooftop in Zurich. It is connected via a coaxial cable to an RTL-SDR and a computer inside a weatherproof closet. Another turnstile antenna was deployed in St. Gallen and connected to a HackRF One. Preliminary experiments were also carried out with other antenna types (Dipole, V-Dipole, QHF) at St. Gallen. However, the turnstile antenna was found to deliver the best results.

The computer runs a custom python script that acquires the samples. The script uses the Soapy abstraction library to ingest the IQ samples. Much of the decoding logic is based on an open-source project [27] that implements the decoding. It leverages ephemeris data for the satellite detection and prediction of the used channels, as well as the Doppler offset. If the receiver is able to synchronize to the incoming bitstream (i.e., estimate where word and packet boundaries are), all Fletcher checksums are checked. This is used to calculate the Packet Error Rate (PER). Packets are only accepted into the dataset if the receiver validates its checksum.

The possible oversampling is limited by the separation of the channels to a value such that no neighboring channel is included. That would be undesirable, because the presence or absence of neighboring bands could help an ML model to identify the transmitter, taking advantage of the variable spacing between downlink channels among the Orbcomm satellite fleet. While the satellites could be programmed to transmit in steps of 2.5 kHz [22], only a few are actually used. Various online forums report frequencies in use by Orbcomm and the smallest channel separation found was 20 kHz. With a bitrate of 4800 bps and a root raise cosine filter with  $\alpha = 0.4$  [22], the captured bandwidth should not be more than  $20\text{kHz} - (2.4\text{kHz} * (1 + \alpha)) = 16640\text{Hz}$ . To account for non-ideal filters, the signal is oversampled by a conservative factor of two.

The stored signal undergoes coarse frequency correction (Doppler shift prediction with ephemeris), fine frequency correction, but not timing nor phase recovery. The IQ samples and associated metadata (Signal to Noise Ratio (SNR), RX location, timestamp, used antenna, used SDR) are stored in a tfile record file. This data and the associated data collection code will be made available on publication.

Overall the data collection ran intermittently for 3 months (1.5 months thereof from two locations) and collected a total of 8992474 packets with 192 IQ samples each. This represents 50.0 h of continuous IQ recordings. The average SNR is 8.1 dB. Figure 3 summarizes the dataset and the distribution of the metadata of the collected samples.

### V. SYSTEM DESIGN

#### A. Data Augmentation

Because of concerns raised in previous research regarding the transferability of fingerprinting models across different

Collected Dataset Overview

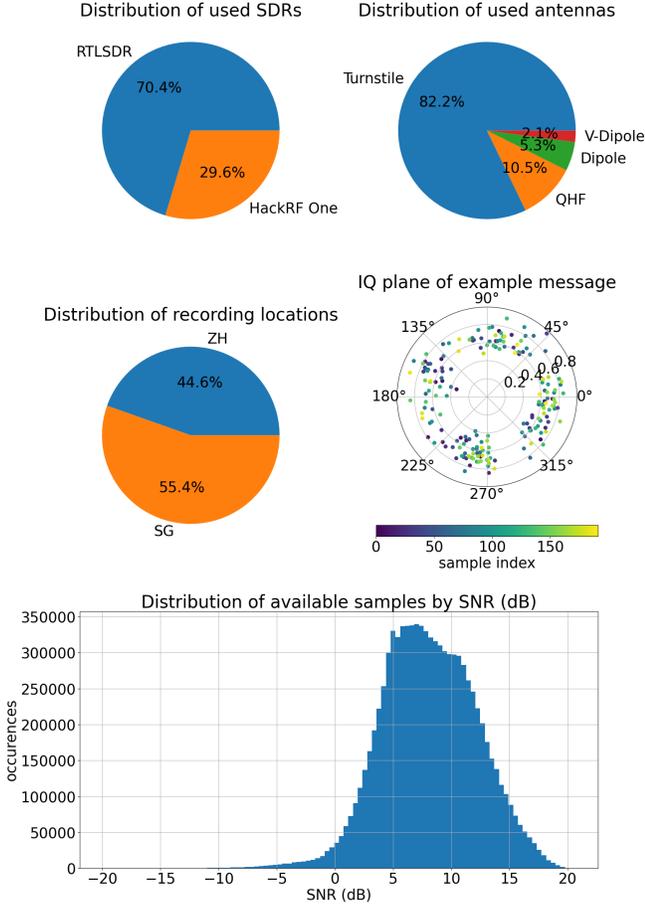


Fig. 3. Breakdown of collected dataset

SDR reboots [28], the data collected for this dataset stems from two different locations with two different physical SDRs.

The data is then augmented in the dataset pipeline. We apply the following operations: random phase shift ( $\Delta\phi \in [0, 2\pi)$ ), random scaling ( $\alpha \in [0.9, 1]$ ), and a random frequency shift ( $\Delta f/f_{\text{sampling}} \in [-0.01, 0.01]$ ). These transformations are illustrated in Figure 4. The intent is to improve training and prevent overfitting.

### B. Machine Learning Model

To provide an extendable system that can handle new categories without requiring extensive retraining, an encoder architecture is chosen that maps the received waveform (IQ samples) into an embedding space.

1) *Loss Function*: The triplet loss with a semi-hard online batching strategy is used. Semi-hard batching considers all positive samples in a batch but only the hardest negative example and was found to provide more stable training [29]. The chosen distance metric is the L2 metric with a margin of 0.7, because this performed best in preliminary experiments.

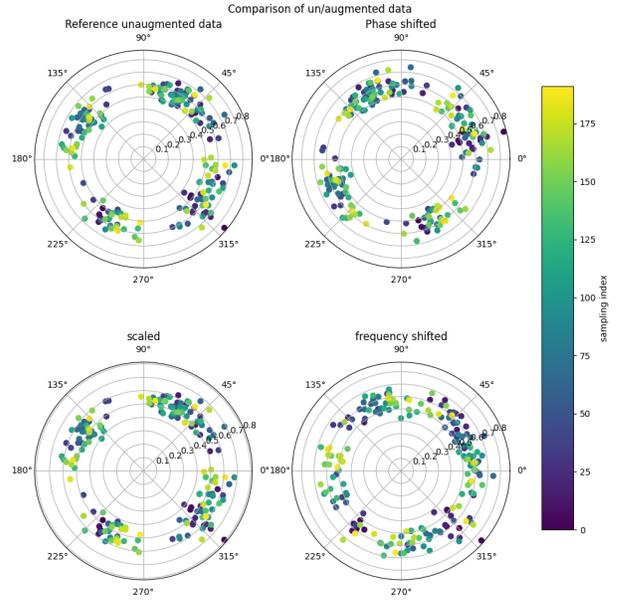


Fig. 4. Illustration of data augmentation

2) *Data Preprocessing*: To improve model performance and increase the model size, the signal is interpolated during preprocessing. This increases the input size of the model and enables bigger Conv + Maxpool layers with more parameters. The following list shows the preprocessing layers (and output tensor shapes) that are applied before the encoder.

- 1) Input (Batch size, 192, 2)
- 2) Interpolation (x64) to (Batch size, 12288, 2)
- 3) Random Bit Transition Shuffling (Batch size, 12288, 2)

Many Orbcomm packets contain fields with information that identifies the specific satellite (ID field or ephemeris data) or narrow down the possible satellites (e.g. information about used transmit frequencies) [22], further preprocessing is needed in order to prevent the model from learning to decode the Orbcomm signal. The data portion of the signal is easily spoofed by the attacker, so we must ensure the model cannot extract any identifying information from it.

To achieve this, the signal is split into different bit transitions and shuffled randomly. Afterwards the different bit transitions are stitched back together. This happens as a preprocessing layer inside the Keras model and is configured to only happen during training. This step decreases the quality of the signal since it introduces discontinuities but is vital to prevent the model from learning logical-layer features that do not offer any protection against replay or spoofing attacks.

3) *Embedding Network*: For the embedding network, a convolutional neural network is chosen. It consists of two separate branches of Conv1D, MaxPool1D and Dropout layers that apply to the I and Q part of the signal, respectively. Using weight sharing on the convolution kernels, the model size is reduced. This weight sharing is motivated by the fact that the received IQ signal can have any phase shift due to the wireless

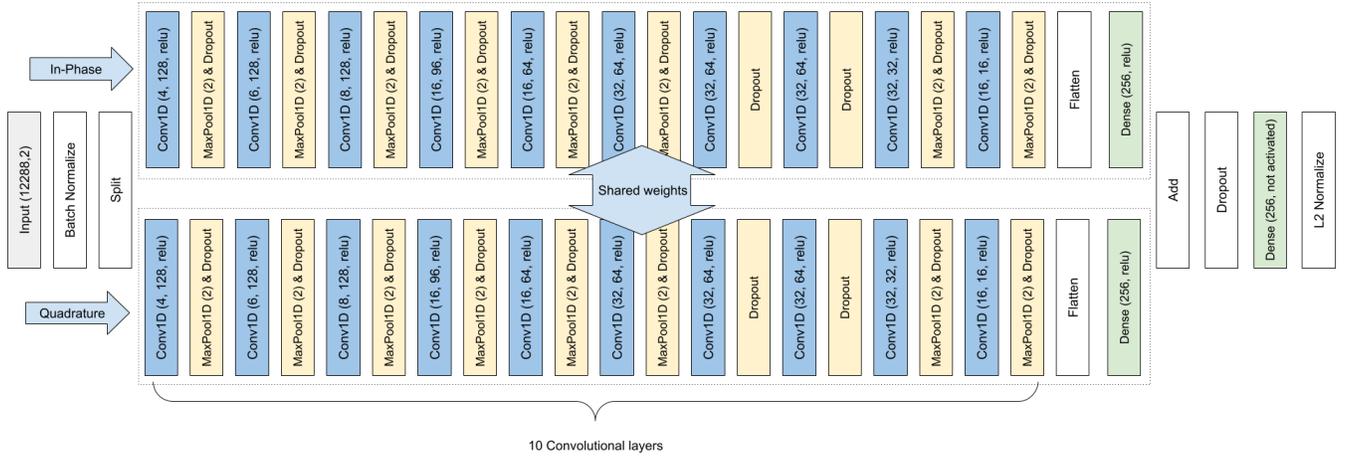


Fig. 5. Proposed Model Architecture

channel, and thus the IQ parts can be swapped (and one of them multiplied by  $-1$ ) if the phase shift is  $\pi/2$ .

After 10 convolutional layers, the processed I and Q signals pass through a dense layer that transform them to the desired embedding dimension. The final embedding dimension is 256. The signals from the I and Q branch are added and passed through a final dense layer. The full model architecture (excluding preprocessing layers) is shown in Figure 5. The final model contains 1091348 trainable parameters (4.2 MB).

4) *Training*: The model is implemented in Tensorflow with the help of Keras for many of the standard layers. Non-standard layers include the triplet loss function (where code from the now-deprecated tensorflow-addons project is used) as well as the preprocessing layers (both the interpolation and the bit transition shuffling layers are custom).

Of the available samples, 27% were used for testing and 9% for validation. The plots in Section VI were generated with validation data if not otherwise mentioned.

Training was performed on an Nvidia A100 Graphical Processing Unit (GPU) with the adam optimizer [30]. A total of 25 epochs were trained over a time of 50 hours. The used batch size was 450. During training, every epoch that achieves a new best Area Under Curve (AUC) metric is saved, such that early stopping can be simulated. Actual early stopping was unhelpful as the AUC metric did not evolve monotonically.

## VI. EVALUATION

### A. Accuracy

We provide initial results of our approach. The Receiver Operating Characteristic (ROC) characteristics are calculated by taking the pairwise distance metric between a batch of  $N = 1024$  validation examples and then varying the detection threshold. This method yields  $(N^2 - N)/2$  binary comparisons per batch. The threshold decides if the pairwise distance is classified as belonging to the same transmitter or not. The threshold is evaluated over the entire validation dataset. The

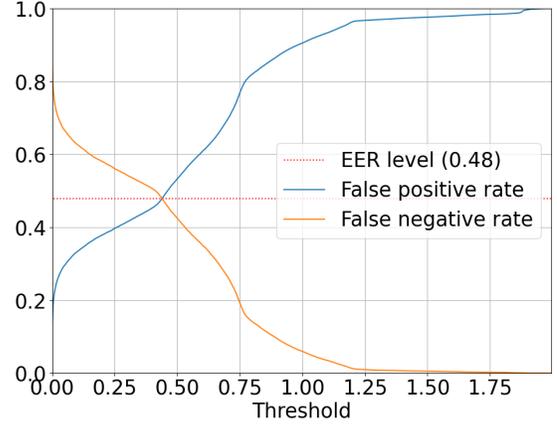


Fig. 6. EER of model

overall ROC AUC is 0.53 with an Equal Error Rate (EER) of 0.48. Figure 6 shows the EER of the ROC curve.

Figure 7 shows that the classification performance improves for higher SNR until it decreases again. We postulate that our SNR measurement is affected by interference which creates false entries with high SNR.

To improve the performance of the classifier, a signal can not only be compared to a single anchor (a ground truth embedding from a known satellite) but to several anchors instead. The anchors are randomly sampled and the mean distance between the tested embedding and the anchors. Figure 8 shows performance for different amounts of anchors. This scenario is different from the baseline performance, in that the model is used to decide if a message belongs to the same transmitter as the anchor or not. In the previous baseline scenario, a sampled subset of all pairwise distances is used to decide whether two signals belong to the same satellite.

The achievable baseline accuracy (at EER) is  $1 - \text{EER} = 0.52$ . When using three anchors the AUC is improved to 0.62

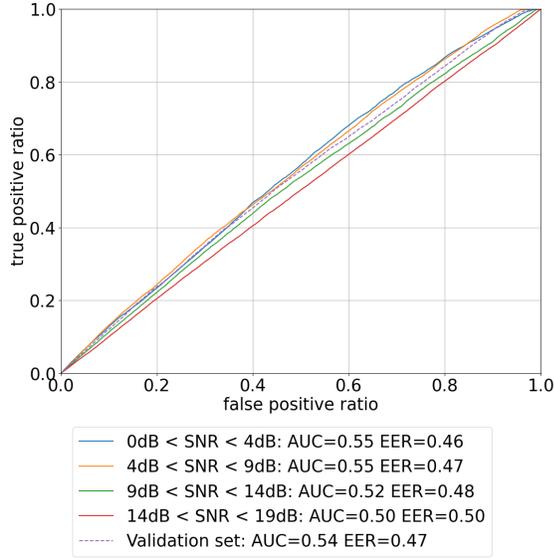


Fig. 7. ROC AUC for different minimal SNR

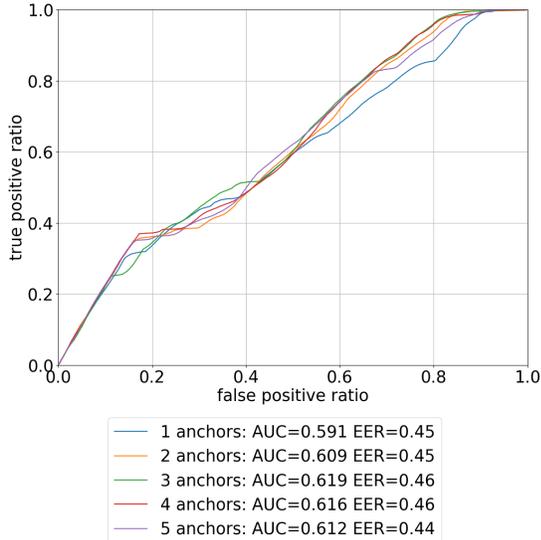


Fig. 8. Anchored ROC

and an accuracy of 0.54.

### B. Security

To assess the security of the proposed scheme, we collect a dataset of spoofed messages. We use a full-duplex USRP B200mini SDR to capture signals with a center frequency of 136.9 MHz and a sampling rate of 3 Msps. We use the same parameters to emit it on the transmit port, which is connected via a 10 dB attenuator to the data collection system. Figure 9 shows the spoofing setup connected to the collection system.

For the analysis, 10000 real recordings and 10000 spoofed recordings are sampled from the datasets and transformed with the model into the embedding space. For every satellite with entries in both sets, the distance between all spoofed



Fig. 9. Signal Spoofing Setup

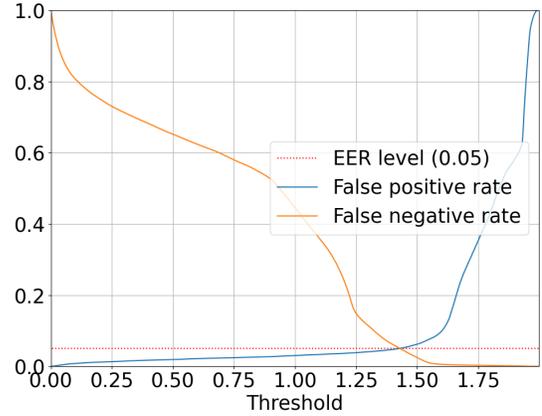


Fig. 10. EER plot of model used to discriminate real and spoofed signals

and real embeddings is calculated, as well as all pairwise distances between embeddings of real samples. The ROC for the binary discriminator (decision between spoofed and not spoofed) is 0.98. Figure 10 shows the false positive / negative rates depending on the used threshold and the EER level.

### C. Deployability

This subsection addresses some considerations that affect the deployability of the proposed system. The accuracies used in the following evaluations are calculated by using the EER threshold of the entire dataset to make a binary decision if two embeddings belong to the same transmitter.

1) *Transferability Between SDRs*: In a deployed system it would be preferable to use a single common set of anchors between all receivers, since deriving new data for each receiver will be cumbersome. In addition, it should be transferable to a new type of SDR device to simplify deployment or enable the system hardware to be upgraded.

Figure 11 depicts the AUC for signals collected from different locations and SDRs. It shows that while samples collected by the same SDR perform better, the proposed model learns to generalize beyond SDR boundaries.

2) *Time stability of the generated fingerprints*: An important requirement for a deployable system is how stable the generated fingerprint is as time goes by, since a low stability requires more frequent updating of fingerprints. [31] Figure 12 shows how the reception times of the two compared signals impact the model performance. Over the span of roughly 2

TABLE I  
COMPARISON OF DIFFERENT LOW EARTH ORBIT (LEO) RFF MODELS

Work	Target Constellation	Dataset Size	Best Achieved Result
This Work	Orbcomm	8992474 packets	AUC 0.62, EER 0.44
This Work	Iridium	SatIQ [3] Dataset	AUC 0.576, EER 0.448
SatIQ [3]	Iridium	$1.7 \cdot 10^6$ ring alerts	AUC 0.698, EER 0.350
PastAI [5]	Iridium	$1 \cdot 10^8$ IQ samples	Accuracy 0.82, up to $>0.9$ (for $>9$ excluded classes)
Zhu et al. [4]	Iridium	$1.98 \cdot 10^8$ IQ samples	Accuracy 0.924, F1-score 0.923

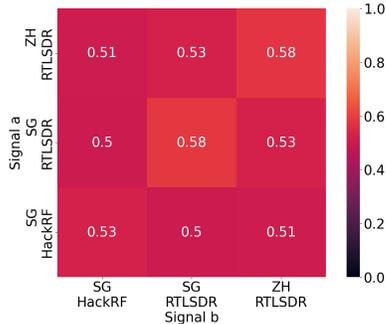


Fig. 11. AUC for different location and SDR combinations

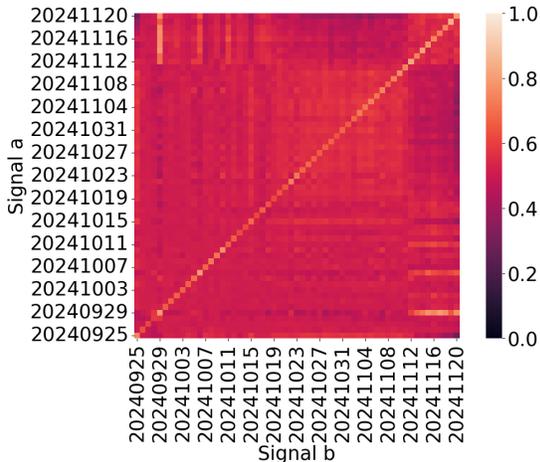


Fig. 12. Impact of reception date on model performance. The accuracy is calculated with the EER threshold of the model over the validation set.

months the performance does not show deterioration over time, except for very good performance on the same day. The good performance of recordings on the same day is presumably that this class contains samples from the same pass that are thus likely to have experienced the same channel and atmospheric distortion and are thus more similar.

3) *Usage Time*: The fingerprinting works on a packet level. Thus a packet consisting of 12 bytes (24 bytes for ephemeris packets [22]) has an air time of 20 ms (40 ms). The model takes an additional 30 ms on a single thread of an Intel NUC with an i5-1135G7 running at 2.40 GHz to make a prediction.

#### D. Comparison To Other Models

In order to test the merit of the proposed model itself, we use it to learn data published in the context of previous work on Iridium RFF. This allows the comparison of the results achieved by the proposed model and the SatIQ model [3].

The proposed model achieves a maximal validation AUC of 0.576 after training 100 epochs, compared with 0.627 after training for 100 epochs as achieved by the SatIQ model. In order to improve the comparison, no data augmentations were used during training.

The model sizes compare as follows: SatIQ consists of around 10 M trainable parameters (38.23 MB) (including an autodecoder stage that is not strictly relevant for inference). This work’s model contains 1091348 (4.2 MB) trainable parameters, however, due to the slightly smaller input dimension of the Iridium data, the compared model contains only 1,079,064 (4.12 MB) trainable parameters.

On a more conceptual level, Table I compares the results of the proposed model to previous research.

## VII. DISCUSSION

One interesting point is that the performance of the model against spoofed signals is significantly better than the performance in discriminating signals from the different satellites. We assume that this is because the task of discriminating between satellites of the same type is easier than to discriminate against a different type of transmitter (the attacker SDR in this case).

The model performance shows an interesting dependence on the SNR, where the accuracy first increases but then decreases again with increasing SNR. As mentioned in Section VI, the problem might be that interference caused packets to have wrongly measured SNR. Another possibility is that due to the low occurrence of such high SNR examples (compare Figure 3), the model did not learn them well.

#### Improving Model Performance

In order to further improve the model performance, the following four approaches could be considered.

1) *Increase amount of IQ points per fingerprint*: A major difference to previous work is that only 192 In-phase Quadrature (IQ) samples are used to derive the fingerprint, whereas other work used 10’000 [5] or 11’000 [3].

A comparison from the literature is a study on IEEE802.11ac RFF [11], which uses 128 IQ sample points to feed a classifier network (that can distinguish 5 transmitters).

The used dataset contains 20 M samples collected from different distances. The achieved results show a good performance (>90% accuracy) for close (<12 m) distances. The dataset stops at 15m distance between transmitter and receiver with an accuracy just below 75% [11] and a decreasing trend. That study shows that fingerprinting is also possible with low-dimensional inputs but also that performance decreases with distance and, correlatedly, lower SNR. A major difference, though, is the channel, since the cited study collected the data in an indoor environment that is more affected by multi-path effects than the typical LEO satellite communication channel.

These comparisons indicate that a model with more input data could achieve better performance. Studying if fingerprinting on the level of minor frames (50 packets) leads to better results would provide more insight into the impact the amount of IQ input points have.

2) *Improve SNR*: Comparing to the literature, the captured SNRs (compare Figure 3) are rather low. An Iridium data collection campaign [5] yielded an SNR mode of 45 dB and 90% of captured samples fall within 40 dB to 60 dB SNR. In a simulation of LEO satellites, the identification accuracy drastically dropped between the 15 dB and 10 dB simulations (from > 80% down to < 40% across all simulated parameters) [23].

Another indication of RFF performance against noise comes from a study that compared the performance of the SatIQ RFF model against Jamming noise [32]. To compare those results the Gaussian jammer is interpreted as Gaussian channel noise instead. At Signal to Jamming Ratio (SJR) of  $-0.66$  dB (analog to SNR of 0.66 dB) the model on average starts to misidentify signals. These comparisons indicate that operating with higher SNR could improve the performance. This could be achieved by using antennae with more gain or using a preamplifier at the antenna feedpoint.

3) *Anchor Selection*: To make a classification if a signals belongs to a satellite, currently a random anchor is used. The performance could likely be improved by the selection of a suitable anchor.

4) *Model Training*: Training the model longer or with more data could also improve the performance of the model. This was not attempted due to timing constraints. More data would also enable usage of larger models before overfitting issues start manifesting.

## VIII. FUTURE WORK

This work raises a number of areas for future research.

One point of interest would be investigating how well the model generalizes to satellites from the observed constellation that it did not encounter during training.

Testing the model architecture on other VHF (space) communications is another near-term goal. Because the proposed model architecture does not require any particular input encoding of the IQ samples, it could conceivably be extended to cover multiple constellations at once.

Finally, although this paper and previous works show that machine learning based RFF systems perform well against

conventional spoofing, they may potentially introduce a vulnerability against adversarial samples designed to trick the model. Studying the vulnerability and the feasibility of injecting adversarial examples over a wireless channel may provide further insight into the security of ML-based RFF.

## IX. CONCLUSION

This paper extends traditional RFF and machine learning methods to the Orbcomm satellite constellation. We applied new methods to deal with the constraints of the Orbcomm system and demonstrated their effectiveness.

We have also collected and provided a large dataset of Orbcomm packet captures, which will be useful for future research. In addition, the initial assumption that training to distinguish different satellites provides good protection against spoofing attacks is validated by the much better performance of the model when used to detect spoofing compared to the baseline performance.

The final AUC of 0.98 against packet replays over a coax connection shows the promise of the proposed model.

## REFERENCES

- [1] [Online; accessed 09-December-2024]. [Online]. Available: <https://spoofing.skai-data-services.com/>
- [2] R. Bisping, J. Willbold, M. Strohmeier, and V. Lenders, "Wireless signal injection attacks on VSAT satellite modems," in *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Aug. 2024.
- [3] J. Smiles, S. Köhler, S. Birnbach, M. Strohmeier, and I. Martinovic, "Watch This Space: Securing Satellite Communication through Resilient Transmitter Fingerprinting," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023.
- [4] S. Zhu, Y. Zhang, J. Zhu, Y. Chen, Y. Shen, and X. Jiang, "3d convolution-based radio frequency fingerprinting for satellite authentication," in *GLOBECOM 2023 - 2023 IEEE Global Communications Conference*, 12 2023, pp. 7586–7591.
- [5] G. Oligeri, S. Sciancalepore, S. Raponi, and R. D. Pietro, "Past-ai: Physical-layer authentication of satellite transmitters via deep learning," *IEEE Transactions on Information Forensics and Security*, vol. 18, 2023.
- [6] "HackRF One - Great Scott Gadgets — greatscottgadgets.com," <https://greatscottgadgets.com/hackrf/one/>, [Accessed 02-12-2024].
- [7] "About RTL-SDR — rtl-sdr.com," <https://www.rtl-sdr.com/about-rtl-sdr/>, [Accessed 02-12-2024].
- [8] N. Soltanieh, Y. Norouzi, Y. Yang, and N. C. Karmakar, "A review of radio frequency fingerprinting techniques," *IEEE Journal of Radio Frequency Identification*, vol. 4, no. 3, pp. 222–233, 2020.
- [9] I. Agadacos, N. Agadacos, J. Polakis, and M. Amer, "Chameleons' oblivion: Complex-valued deep neural networks for protocol-agnostic rf device fingerprinting," in *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, 09 2020, pp. 322–338.
- [10] V. Brik, S. Banerjee, M. Gruteser, and S. Oh, "Wireless device identification with radiometric signatures," in *Proceedings of the 14th ACM International Conference on Mobile Computing and Networking*, ser. MobiCom '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 116–127.
- [11] S. Riyaz, K. Sankhe, S. Ioannidis, and K. Chowdhury, "Deep learning convolutional neural networks for radio identification," *IEEE Communications Magazine*, vol. 56, no. 9, pp. 146–152, 2018.
- [12] T. Jian, B. C. Rendon, E. Ojuba, N. Soltani, Z. Wang, K. Sankhe, A. Gritsenko, J. Dy, K. Chowdhury, and S. Ioannidis, "Deep learning for rf fingerprinting: A massive experimental study," *IEEE Internet of Things Magazine*, vol. 3, no. 1, pp. 50–57, 2020.
- [13] G. Oligeri, S. Sciancalepore, and A. Sadighian, "Satprint: Satellite link fingerprinting," in *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*, 04 2024, p. 177–185.
- [14] H. Wang, W. Cheng, C. Xu, M. Zhang, and L. Hu, "Method for identifying pseudo gps signal based on radio frequency fingerprint," in *2018 10th International Conference on Communications, Circuits and Systems (ICCCAS)*, 2018, pp. 354–358.

- [15] S. Qiu, K. Sava, and W. Guo, "Robust satellite antenna fingerprinting under degradation using recurrent neural network," *Modern Physics Letters B*, vol. 36, no. 12, p. 2250043, 2022.
- [16] G. Chechik, V. Sharma, U. Shalit, and S. Bengio, "Large scale online learning of image similarity through ranking," *Journal of Machine Learning Research*, vol. 11, pp. 1109–1135, 03 2010.
- [17] A. Hermans, L. Beyler, and B. Leibe, "In defense of the triplet loss for person re-identification," 2017.
- [18] J. Harms, "The orbcomm experience," [https://connectivity.esa.int/sites/default/files/1\\_The\\_Orbcomm\\_Experience.pdf](https://connectivity.esa.int/sites/default/files/1_The_Orbcomm_Experience.pdf), OHB Technology, Tech. Rep., 2004.
- [19] F. Farhangian and R. Landry, "Multi-constellation software-defined receiver for doppler positioning with leo satellites," *Sensors*, vol. 20, p. 5866, 10 2020.
- [20] Y. Xie, G. Li, H. Qin, C. Zhao, M. Chen, and W. Zhou, "Carrier phase tracking and positioning algorithm with additional system parameters based on orbcomm signals," *GPS Solutions*, vol. 28, no. 4, p. 184, Aug 2024. [Online]. Available: <https://doi.org/10.1007/s10291-024-01721-8>
- [21] M. Orabi, J. Khalife, and Z. Kassas, "Opportunistic navigation with doppler measurements from iridium next and orbcomm leo satellites," in *2021 IEEE Aerospace Conference (50100)*, 04 2021, pp. 1–9.
- [22] M. Kenny, "Ever wondered what is on the orbcomm satellite downlink?" Tech. Rep., 2002. [Online]. Available: <https://github.com/fbieberly/ORBCOMM-receiver/blob/master/literature/Orbcomm.pdf>
- [23] Q. Jiang and J. Sha, "Rf fingerprinting identification based on spiking neural network for leo-mimo systems," *IEEE Wireless Communications Letters*, vol. 12, no. 2, pp. 287–291, 2023.
- [24] E. Salkield, M. Szakály, J. Smailes, S. Köhler, S. Birnbach, M. Strohmeier, and I. Martinovic, "Satellite spoofing from a to z: On the requirements of satellite downlink overshadowing attacks," in *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, ser. WiSec '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 341–352.
- [25] T. Zhao, N. Wang, Y. Wu, W. Zhang, and X. Wang, "Backdoor attacks against low-earth orbit satellite fingerprinting," in *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2024, pp. 01–06.
- [26] *TA-1 Kreuzdipol 135-152 MHz*, [https://www.wimo.com/media/akeneo\\_connector/media\\_files/1/8/18350\\_TA1\\_8b7b.pdf](https://www.wimo.com/media/akeneo_connector/media_files/1/8/18350_TA1_8b7b.pdf), WiMo Antennen und Elektronik GmbH.
- [27] F. Bieberly, "orbcomm-receiver," <https://github.com/fbieberly/ORBCOMM-receiver>, [Accessed 16-10-2024].
- [28] M. Irfan, S. Sciancalepore, and G. Oligeri, "On the reliability of radio frequency fingerprinting," 2024. [Online]. Available: <https://arxiv.org/abs/2408.09179>
- [29] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Jun. 2015.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [31] S. Čapkun, *The Cyber Security Body of Knowledge v1.0, 2019*. University of Bristol, 2019, ch. Physical Layer & Telecommunications Security, kA Version 1.0. [Online]. Available: <https://www.cybok.org/>
- [32] J. Smailes, E. Salkield, S. Köhler, S. Birnbach, M. Strohmeier, and I. Martinovic, "Sticky fingers: Resilience of satellite fingerprinting against jamming attacks," in *2nd Workshop on Security of Space and Satellite Systems (SpaceSec)*, Mar 2024.