

Watch This Space: Securing Satellite Communication through Resilient Transmitter Fingerprinting

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

Sebastian Köhler
University of Oxford
Oxford, UK
sebastian.kohler@cs.ox.ac.uk

Simon Birnbach
University of Oxford
Oxford, UK
simon.birnbach@cs.ox.ac.uk

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

Ivan Martinovic
University of Oxford
Oxford, UK
ivan.martinovic@cs.ox.ac.uk

ABSTRACT

Due to an increase in the availability of cheap off-the-shelf radio hardware, signal spoofing and replay attacks on satellite ground systems have become more accessible than ever. This is particularly a problem for legacy systems, many of which do not offer cryptographic security and cannot be patched to support novel security measures.

Therefore, in this paper we explore *radio transmitter fingerprinting* in the context of satellite systems. We introduce the SATIQ system, proposing novel techniques for authenticating transmissions using characteristics of the transmitter hardware expressed as impairments on the downlinked radio signal. We look in particular at high sample rate fingerprinting, making device fingerprints difficult to forge without similarly high sample rate transmitting hardware, thus raising the required budget for spoofing and replay attacks. We also examine the difficulty of this approach with high levels of atmospheric noise and multipath scattering, and analyze potential solutions to this problem.

We focus on the *Iridium* satellite constellation, for which we collected 1 705 202 messages at a sample rate of 25 MS/s. We use this data to train a fingerprinting model consisting of an autoencoder combined with a Siamese neural network, enabling the model to learn an efficient encoding of the message headers that preserves identifying information.

We demonstrate the fingerprinting system’s robustness under attack by replaying messages using a Software-Defined Radio, achieving an Equal Error Rate of 0.120, and ROC AUC of 0.946. Finally, we analyze its stability over time by introducing a time gap between training and testing data, and its extensibility by introducing new transmitters which have not been seen before. We conclude that our techniques are useful for building fingerprinting systems that are stable over time, can be used immediately with new transmitters without retraining, and provide robustness against spoofing and replay attacks by raising the required budget for attacks.

CCS CONCEPTS

• **Security and privacy** → **Mobile and wireless security; Systems security**; • **Hardware** → *Wireless devices*; • **Computer systems organization** → *Sensor networks*; • **Computing methodologies** → *Machine learning*.

KEYWORDS

Radio security; Satellite security; Machine learning; Neural network; Fingerprinting; Systems security

ACM Reference Format:

Joshua Smailes, Sebastian Köhler, Simon Birnbach, Martin Strohmeier, and Ivan Martinovic. 2023. Watch This Space: Securing Satellite Communication through Resilient Transmitter Fingerprinting. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*, November 26–30, 2023, Copenhagen, Denmark. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3576915.3623135>

1 MOTIVATION

Recent years have seen a dramatic rise in the availability of cheap radio hardware, particularly Software-Defined Radios (SDRs). Not only have these devices become more widely available, but their capabilities have increased, with devices like the *HackRF One* (340 USD, Adafruit) able to transmit and receive in frequencies ranging from 1 MHz to 6 GHz [1]. As a result, the ability to carry out spoofing attacks, once exclusive to large-budget organizations and nation-state actors, is now within reach for even motivated hobbyists. This poses a particular threat to satellite systems, many of which were built under the assumption that tampering with signals would be prohibitively expensive for the vast majority of attackers.

Spoofing and replay attacks have been widely explored in wireless systems – attackers equipped with an SDR can overshadow legitimate communications or spoof messages outside normal communication. Many widely used systems are vulnerable to these attacks, including the ADS-B avionics protocol [2], the LTE telephony system [3, 4], and satellite systems including GPS [5]. Due to the critical nature of satellite systems, it is vital that operators can prevent or detect spoofing attacks to protect the systems and applications that rely on them.

There are a wide range of techniques for detection and prevention of spoofing attacks, the foremost of which is cryptography – a properly implemented cryptosystem with associated key management provides robust authentication, making spoofing attacks

CCS '23, November 26–30, 2023, Copenhagen, Denmark

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*, November 26–30, 2023, Copenhagen, Denmark, <https://doi.org/10.1145/3576915.3623135>.

near-impossible. However, there are a number of reasons why cryptography may not be desirable (or possible) in the context of satellite systems. Firstly, there are a huge number of legacy satellites currently in orbit. Many of these do not implement cryptography, and cannot be retrofitted to do so due to their limited onboard processing power. However, the data collected by these satellites is immensely useful in both scientific and private use cases – they are used for monitoring forest fires, land usage, population density, flooding, and more [6–9]. These satellites are often bespoke designs which would be prohibitively expensive to replace; it is important to ensure systems like these can be used for their entire projected lifespan (and beyond).

There are also a number of satellite systems which were initially built with cryptography, but which have become insecure post-launch due to leaked keys [10] or outdated cryptosystems [11]. Some of these satellites cannot be patched due to a lack of over-the-air update capabilities, so other methods must be used to authenticate their telemetry data.

Finally, some attacks can be carried out without violating any cryptographic properties of the system. The authors of [12] show that precisely timed message replays can cause Global Navigation Satellite Systems (GNSS) to misreport the location of the receiver to an attacker-specified location. Since these attacks are carried out by simply introducing delay to messages rather than altering message contents, conventional cryptography does not protect against them. We still want to be able to detect and prevent these attacks, so we must turn to non-cryptographic techniques for message authentication. In particular, we investigate radio transmitter fingerprinting, in which radio signal characteristics are used to identify the transmitters. This is achieved by identifying impairments on the signal which are created by small differences in the radio transmitter hardware. These impairments are unique to transmitters and consistent over time as we will show for the satellite case.

Contributions

In this paper we present SATIQ, a novel approach to fingerprinting satellite signals. We work with signals at a high sample rate to counteract problems with spoofing at lower sample rates. This makes the techniques more useful in a security context, requiring attackers to use more expensive radio hardware that works at high sample rates in order to successfully impersonate a device – this excludes a large number of low-budget adversaries. In doing this we can provide an additional level of confidence in the authenticity of the origin of satellite signals, particularly in systems where cryptography is either unavailable or ineffective. We verify that the system can detect replay attacks by replaying captured messages using an SDR.

We also use a *Siamese model* – unlike conventional classifiers, these compare two signals and produce a distance metric representing the likelihood of two messages having been sent from the same transmitter. This technique enables *one-shot learning*: new transmitters can be introduced without requiring the system to be retrained, and can be used immediately with only a small number of examples. This is particularly useful in the context of satellites in Low Earth Orbit (LEO), which must be replaced more frequently.

We captured a large high sample rate dataset of real-world data from the *Iridium* satellite constellation. We use this data to train and test our system, since decoders are readily available and previous works have demonstrated the constellation to be viable for fingerprinting [13]. Our baseline architecture can be used and enhanced for use in deployed satellite systems. To facilitate this, our full dataset and code have been made freely available.¹

2 BACKGROUND

In this section we discuss key concepts in radio and digital signal processing. We also explore existing countermeasures to spoofing attacks, and look at existing research in radio fingerprinting in general. Finally, we discuss related work that focuses on radio fingerprinting of satellites, as well as related work that considers fingerprinting in a security context.

2.1 Software-Defined Radios

Software-Defined Radios (SDRs) allow the use of software for signal processing tasks traditionally done using dedicated hardware. This is achieved by sampling raw signals into a digital form and sending the samples to a computer, where they can be processed further. This provides significant versatility to signal processing over traditional radio hardware, at the cost of additional processing power.

A concept at the core of digital signal processing is *IQ Sampling*. By taking a *carrier signal* at a given frequency and sampling the components of the incoming signal that are in phase (*I*) and out of phase (quadrature, *Q*) with the carrier before sampling, the incoming signal is downsampled, shifting the carrier frequency to 0 Hz [14]. This significantly reduces the sample rate required.

One benefit of sampling in this way is that samples can be represented as complex numbers, and plotted on the complex plane. In this representation, distance from the origin represents amplitude and angle from the horizontal axis represents phase relative to the carrier signal. This makes it a particularly useful representation of *Phase Shift Keying* (PSK), a form of signal modulation which encodes data in the phase of the signal. In this case, symbols appear as distinct points on the complex plane, producing a *constellation diagram*.

We only see distinct points on the constellation diagram if we sample at the exact symbol rate of the modulation scheme. If we instead *oversample* the signal by using a significantly higher sample rate, we start to see the points between the symbols, as the transmitter hardware modulates between them. This can be seen later on in Figure 8 (Section 4.3). The appearance of this interpolation between points is affected by a number of factors including atmospheric noise and multipath distortion, but also by small variations in the transmitter hardware – these can be used to fingerprint the transmitter.

2.2 Spoofing Countermeasures

In Section 3 we explore the threat of an attacker equipped with a software-defined radio, looking in particular at spoofing attacks.

¹The code is available at <https://github.com/ssloxford/SatIQ>, and the dataset and model weights at <https://zenodo.org/record/8220494> and <https://zenodo.org/record/8298532> respectively.

When implemented properly, modern cryptographic authentication can solve this problem, but it is also possible to make these systems more secure without the use of cryptography. As discussed in Section 1, this is desirable in legacy systems without cryptography, or systems where cryptographic authentication is undesirable or has been compromised. Additionally, some systems provide open data by design, leaving out cryptography on downlinked communications by choice.

There are a diverse range of approaches for the authentication of downlinked transmissions without the use of cryptography. These can be partitioned into *data inspection*, *timing analysis*, and *waveform analysis*.

Data Inspection. The integrity of data can be verified by receiving the same transmission at multiple locations, and comparing the data or its hash between receivers. This requires an attacker to be physically present at each location in order to successfully carry out spoofing attacks, where they would have previously only needed to be at a single location. This is particularly useful in cases where there are already large numbers of community-operated ground stations, such as with NASA’s Direct Readout Laboratory (168 operated worldwide at the time of writing) [15]. However, it is likely to be infeasible for smaller organizations to set up multiple ground stations.

Timing Analysis. These techniques involve looking at the timing of signals in order to verify their legitimacy. At the simplest level, this could be ensuring the signal is received at a time the satellite was known to be transmitting – this does not provide much real security. More advanced techniques in this area include Time Difference of Arrival (TDOA) analysis, looking at the time difference between multiple receivers to verify the transmitter’s location against where the satellite is known to be [16]. This is an effective technique, forcing any potential attackers to be physically present at every ground station, but is once again only feasible for larger organizations capable of operating many ground stations.

Waveform Analysis. Alongside data and timing analysis, the waveform itself can be inspected to detect attacks. A spoofed signal is likely to have different properties from the legitimate signal, particularly when spoofing requires overshadowing an ongoing transmission. These properties include amplitude, SNR, doppler shift, and signal distortion [17, 18]. With appropriate radio hardware it is possible to verify these parameters, making attacks more difficult to execute – the adversary must replicate the measured properties in order to successfully spoof messages.

Fingerprinting also falls into this category – by looking at unique impairments on the raw waveform we can identify the transmitter. When spoofing or replaying messages the attacker’s radio will impart a different fingerprint on the signal, allowing a fingerprinting system to detect when this has occurred. In order for the adversary to circumvent this system they will need to replicate the fingerprint of the legitimate transmitter. Depending on how the fingerprinter has been designed, this may raise the required budget to carry out attacks.

2.3 Fingerprinting

Radio fingerprinting is a mature field, with a large base of research looking at a wide range of techniques on many different systems – [19] provides a good overview of existing research. Fingerprinting techniques can be partitioned into two key areas: *transient fingerprinting* and *steady-state fingerprinting*.

2.3.1 Transient Fingerprinting. The *transient* of a radio signal occurs when the transmitter first powers on, or changes power levels following a signal lock. Various properties of the transient, such as its duration or the number of peaks which occur in the carrier signal, are characteristic to the transmitter and can be used to identify it, if properly extracted. The majority of historical fingerprinting research makes use of transient analysis, since almost all radio transmission involves a transient, and the transient typically exhibits the same characteristics every time the device powers up. Much of the work in transient fingerprinting revolves around novel techniques for precisely identifying the start and end of the transient [20, 21], or processing the transient to extract useful identifying features [22].

Transient fingerprinting has seen some use in security contexts – in [23], transient fingerprinting is used to identify devices even when all other identifying information has been removed, and to detect wormhole and device cloning attacks (these attack types are explored further in Section 3).

2.3.2 Steady-state Fingerprinting. In contrast to transient fingerprinting looking at a very brief portion of the signal, steady-state fingerprinting instead looks at the modulated portion of the signal. The features in the steady-state portion of the signal are different from the transient, often looking at how the IQ constellation is affected by hardware impairments. These impairments include quadrature errors, self-interference, amplitude clipping, and frequency offsets in the various stages of signal modulation (DAC, mixer, filter, upconverter, amplifier) [24]. Additionally, the signal is affected by properties of the wireless channel, including background noise, free space path loss, and multipath distortion, which are likely to change over time as the environment changes. These wireless channel properties are significantly more prominent in satellite communication, due to the long-distance radio links involved.

There is some variety in the techniques used, including observing features of the constellation at a low sample rate [13], analyzing features in the frequency domain [25], and looking at high sample rate signals to observe high frequency impairments [26]. Machine learning techniques are more commonly used to aid steady-state fingerprinting, particularly in approaches working at a high sample rate, in order to pull out features which may not be immediately obvious. Manual feature engineering is also used, but is less common than in transient analysis.

The majority of steady-state fingerprinting looks at a single protocol or class of devices at a time, but with the recent rapid increase in machine learning capabilities this is no longer a requirement, and there has been some work into generalizable fingerprinting techniques which do not require retraining to apply to new contexts [27].

There are also some interesting techniques extending the concept of fingerprinting – the authors of [28] train a convolutional neural

network to identify SDRs at 5 MS/s, then intentionally introduce signal impairments at the transmitter in order to further increase classification accuracy. This achieves incredible accuracy (greater than 0.995), but fingerprint forgery is not considered in this context.

2.4 Related Work

We have explored the existing body of radio fingerprinting research in the previous section. However, there is also some research that focuses specifically on satellite fingerprinting and fingerprinting in a security context.

2.4.1 Satellite Fingerprinting. Unlike signals from terrestrial devices, satellite signals have to travel hundreds of kilometers through the atmosphere, causing significant signal attenuation and channel noise. This adds additional challenge to fingerprinting in this context, particularly since many techniques rely on minimal presence of background noise. There are some works looking at fingerprinting in the presence of noise, either by adding noise to clean signals during model training (effectively training models to remove/ignore the noise) [29], or by smoothing out long signals at low sample rates to obtain average symbol positions [30]. In the context of satellites it is difficult to obtain signals without noise, and smoothing does not work with high sample rate signals (since important detail is lost) – we must find other methods of reducing or ignoring noise. We discuss this further in Section 4.

The authors of [13] design “PAST-AI”, analyzing heatmaps of low sample rate transmissions from the *Iridium* constellation to classify satellites. This technique achieves an accuracy of approximately 0.85, increasing to 1.00 for small subsets of the constellation. Although this technique achieves high accuracy, it is not as useful from a security context – the classifier works by processing large batches of consecutive messages, making it more difficult to detect individual message spoofing. Furthermore, fingerprinting at a very low sample rate (1 sample per symbol) makes fingerprint forgery significantly easier, since the attacker does not need to replicate as many features of the waveform. The work claims to be able to discriminate SDR-equipped attackers from legitimate satellites by solving the harder problem of discriminating between satellites with the same hardware. This is true, however, the assumption only holds at the sample rate used by the fingerprinting system. As PAST-AI operates at 1 sample per symbol, it cannot protect against SDR-based attacks at higher sample rates.

There has also been some fingerprinting work looking at other satellite systems – the authors of [24] make use of manual feature extraction to identify spoofing of GPS localization satellites. This technique is effective, but manual feature extraction is less likely to transfer easily to other satellites or constellations.

2.4.2 Security. Most fingerprinting is done for the purpose of security, preventing spoofing and replay attacks, but some techniques have been specifically proposed to provide better security properties. For instance, the authors of [26] address the problem of identifying devices which have not been seen in the training dataset by separating the feature extraction and classification components of the model, using clustering techniques on the extracted features to identify transmitters. This allows new transmitters to be introduced without retraining the feature extraction component. We take a

similar approach in our work, using an autoencoder to produce the fingerprints and a Siamese model in the place of clustering. The SATIQ system is fully described in Section 4.

There has also been work assessing attacks on fingerprinting systems – it has been shown that an arbitrary waveform generator with a sufficiently high sample rate can be used to impersonate devices, fooling fingerprinting systems [31]. Hardware that can achieve these sample rates is prohibitively expensive for the vast majority of adversaries (we discuss budget further in Section 3), so we do not consider this to be a major issue – although our techniques are likely vulnerable to impersonation at a very high sample rate, preventing spoofing below a certain transmitter sample rate is sufficient to exclude the vast majority of lower-budget attackers.

3 THREAT MODEL

Goal. In this paper we concentrate on attacks involving spoofing and message replay. In the case of spoofing, the adversary’s goal is to broadcast messages appearing to come from a satellite such that the ground system processes them alongside legitimate messages. Alternatively, the attacker may delay or advance messages (jamming the original and replaying a recording) to affect timing-based systems such as GPS [12]. Similarly, they could carry out “wormhole attacks”, in which messages are captured at one location and tunneled to another location, from where they are broadcast – this is also effective against GPS and other localization systems. Unlike spoofing, these attacks can be performed even on signed or encrypted messages, since they do not affect message contents.

Capabilities. The adversary’s capabilities and budget will vary depending on the type of adversary. We can reasonably expect any attacker to have access to an off-the-shelf SDR (as well as the appropriate amplifiers and antennas) enabling them to transmit messages within the vicinity of a single ground station. An equipment setup similar to the one given in Section 4.3 could be used to carry out attacks, costing approximately 6600 USD. By using a cheaper SDR with a lower sample rate and a different antenna (with a suitable amplifier), the budget can be significantly lowered to approximately 500 USD.² Alternatively, with a greater budget the attacker could afford a more powerful amplifier enabling messages to be broadcast from a greater distance, or multiple copies of the hardware to attack multiple ground stations simultaneously.

It has been demonstrated in [31] that device fingerprinting is vulnerable to signal replay attacks, provided the attacker has access to a high-end arbitrary waveform generator capable of transmitting signals with a sufficiently high sample rate.³ For this reason we do not consider attackers with nation-state level capabilities, since they are always capable of purchasing hardware that can fool a fingerprinting system.

In this work we are particularly interested in lower-budget attacks; a simple spoofing or replay attack with a cheap SDR can have a potentially devastating effect on improperly secured satellite systems. Through robust high-sample-rate device fingerprinting, we

²For example, a HackRF One can be used as the radio (340 USD, Adafruit), with a suitable power amplifier and passive antenna.

³The authors approached a reputable manufacturer, and received a quote for approximately 125 000 USD at academic institution rates. We therefore conservatively assume that in the best case this hardware would cost no less than 60 000 USD.

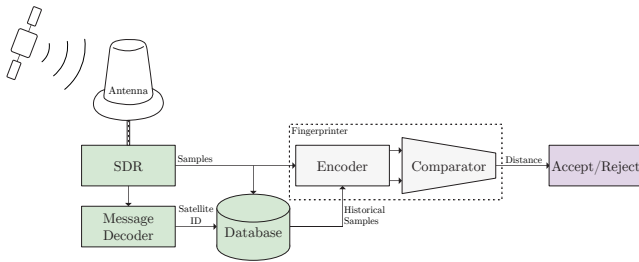


Figure 1: An overview of the end-to-end fingerprinting process used by SATIQ. Satellite signals are received, decoded, processed into fingerprints, and compared to historical fingerprints to determine a distance metric, which is used to accept or reject the message.




Figure 2: An overview of the Siamese autoencoder architecture used in SATIQ's fingerprinting model. Two inputs are passed into the encoder with identical weights, and the encodings are compared using the comparator (angular distance) to generate a distance metric.

aim to defeat these attacks by making it impossible to forge the fingerprint on spoofed signals using only cheap COTS radio hardware. In doing so, we increase the budget of attacks such that they can no longer be carried out by lower-budget attackers.

Of course, we do not expect fingerprinting techniques to be able to prevent all attacks; such a goal is unrealistic for this scope, requiring robust cryptography. As discussed in Section 1, there are a large number of satellites with outdated or no cryptography and no capability for in-flight patching, and we argue that it is still crucial to secure these satellites as well as possible despite this, so they can still be used to serve customers and for the advancement of science. We discuss in more detail in Section 5 how the threat model changes with the introduction of robust signal fingerprinting, addressing remaining concerns and potential avenues for future research.

We also do not consider attacks from within a compromised system, in which compromised satellites are communicating with other devices in the network. This further means that spoofing attacks launched from satellites within the constellation that our system protects are out of scope for this work. Existing systems such as PAST-AI [13] are already well-suited to tackle this type of attacker. However, as these existing systems are based on low-sample rate data, they cannot protect against SDR-based attacks – the focus of this work.

We are primarily concerned with attackers overshadowing messages on the downlink – whilst attacks on the uplink are possible,

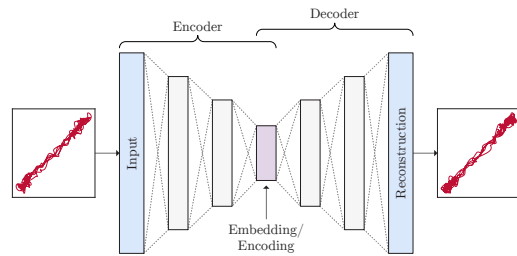


Figure 3: The layout of an autoencoder. Input is passed through an encoder to produce a low-dimensional encoding, and then through a decoder to produce output of the same dimension as the input. The model is trained to reconstruct the input as best as possible.

they have not been extensively explored due to the greater hardware cost of a suitable amplifier and directional dish. Furthermore, device fingerprinting on the space segment is currently infeasible, requiring large amounts of computational power, and cannot be carried out aboard the legacy satellites with which we are primarily concerned. Further work may consider fingerprinting the uplink of legacy satellite systems by capturing signals in-transit, but this is a very limited use case and is out of scope for this research.

4 SYSTEM DESIGN

Our system comprises two primary components: data collection and the SATIQ machine learning fingerprinting system. A representation of the end-to-end system can be seen in Figure 1 – messages are collected by an SDR and decoded into message contents, and the raw samples are ran through an encoder network (explained below), and compared against known example messages from the same transmitter. This produces a distance metric – low distances indicate the message is likely to be legitimate. This is used to accept or reject the message.

All code and datasets have been made openly available. The source code can be found at <https://github.com/ssloxford/SatIQ>, and the dataset and model weights at <https://zenodo.org/record/8220494> and <https://zenodo.org/record/8298532> respectively.

4.1 Design Decisions

Our fingerprinting model, illustrated in Figure 2, uses an autoencoder combined with a Siamese model to compare two input waveforms and produce a distance metric between the two inputs. In the following, we describe autoencoders and Siamese networks in general, and we explain why we chose them as the basis for our fingerprinting system.

Autoencoders. An autoencoder is a type of neural network which is used to learn an efficient encoding of data. This is achieved by simultaneously training an *encoder* and *decoder*, validating the accuracy of the encoding by comparing the output of the decoder to the input (reconstruction accuracy). The output of the encoder is restricted in size, thus forcing data to pass through a *bottleneck* at this portion of the network. This forces the model to discard less important information, producing an efficient encoding. This technique is particularly useful for dimensionality reduction, since

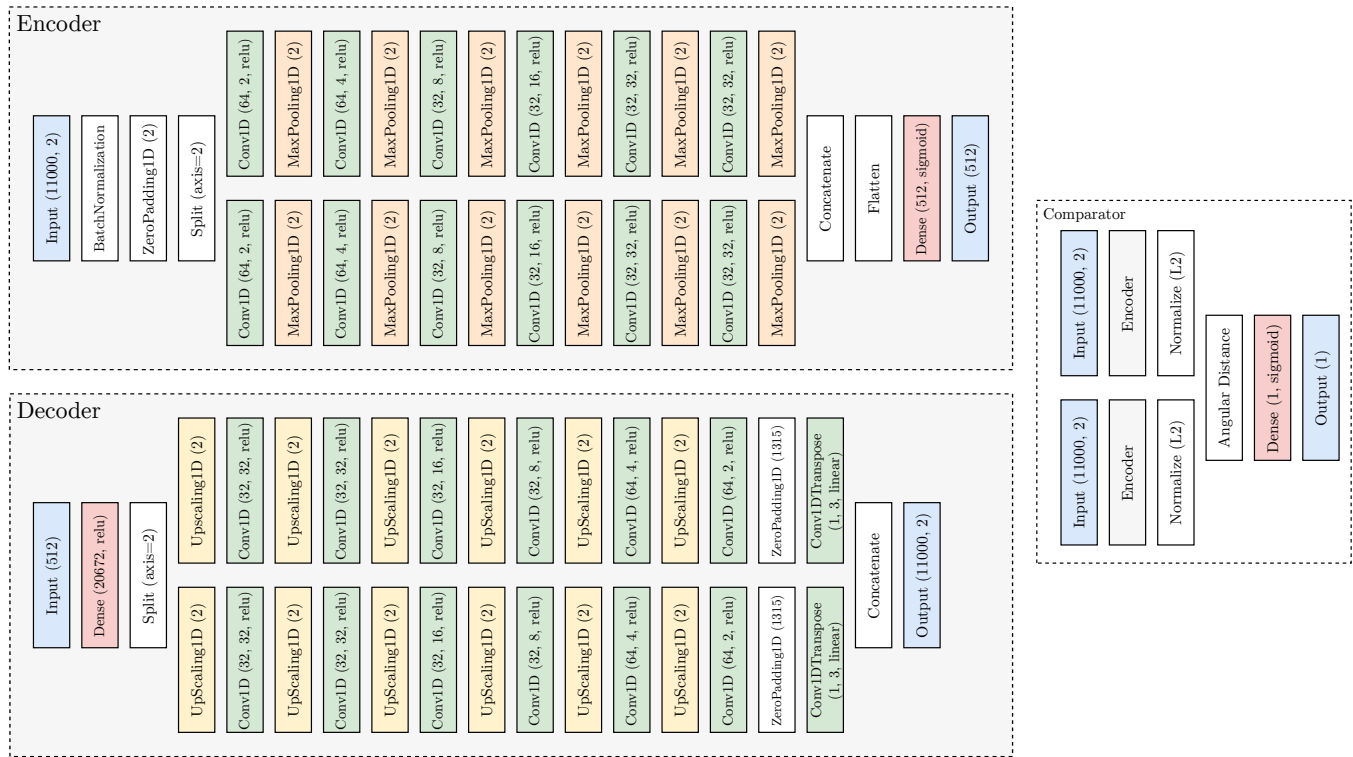


Figure 4: The layers of the Siamese neural network used in SATIQ. The encoder uses separate convolutional and max-pooling layers for the l and Q portions of the signal, before producing a final embedding using a dense layer. Similarly, the decoder uses separate convolutional and upsampling layers. The comparator uses two copies of the encoder with identical weights, computing a difference score between the outputs.

the encoder produces an embedding of the input in a significantly lower-dimensional space. The layout of an autoencoder can be seen in Figure 3.

Siamese Neural Networks. Siamese neural networks are designed to be effective for one-shot classification. To this end, they generate a similarity score between two inputs [32]. This is achieved by passing each input through the same “encoder” network to generate an embedding of the inputs in some feature space, followed by a comparison function to generate a distance metric in the feature space – the lower the distance, the more similar the samples are. The weights of the encoder network are shared between the two inputs.

We chose this approach due to some of its advantages over a simple classifier, particularly in the context of fingerprinting:

- The number of classes is not fixed – new classes can be introduced after training by comparing to examples from the new class.
- The one-shot (or in some cases few-shot) nature of the model means a new class can be identified using only a very small number of examples.
- The distance threshold can be raised to increase the acceptance rate of legitimate messages at the cost of increased false positives (or vice versa), granting fine-grained control of the level of protection granted by fingerprinting.

Past work has shown these models to be effective in a wide range of use cases, particularly in classification problems with huge numbers of classes such as malware detection or gait recognition [33, 34]. The architecture has also been demonstrated to be successful at detecting spoofing and replay attacks on other systems, such as face recognition and voice biometrics [35, 36]. Siamese networks have also seen use in radio systems, shown to be effective in areas such as automatic modulation classification [37]. Finally, some research has shown promise in using Siamese networks for fingerprinting radio transmitters [38, 39].

Our work builds upon these in a number of key aspects. Firstly, we deal with a more difficult scenario than the majority of terrestrial fingerprinting cases – the great distance of satellite transmitters introduces large amounts of atmospheric noise and multipath distortion, which dwarf the hardware impairments on the signal. Secondly, we consider in particular the security implications of radio fingerprinting, assessing the level of protection against spoofing and replay attacks granted by our approach, and the expected budget required to circumvent these techniques. Finally, we use an autoencoder alongside the Siamese network to provide encodings that better capture meaningful features in the input – prior work has shown this to be effective in areas such as signature verification, but to our knowledge we are the first to apply this architecture to radio fingerprinting [40].




Figure 5: An illustration of the triplet loss function. This function takes an anchor a , a positive sample p of the same class as the anchor, and a negative sample n of a different class. Optimizing this loss function minimizes the distance between the anchor and positive samples, and maximizes the distance between the anchor and negative samples.

In a non-adversarial classification setting, simple classifiers typically exhibit higher performance than Siamese networks, since they partition the entire input space into fixed categories. However, this is less effective in adversarial settings, since malicious transmitters will still be given a legitimate label. Our results in Section 5 show that Siamese architectures are particularly effective at detecting replay attacks, even with the levels of noise observed in satellite signals.

Siamese models can also require more data during training, since they must learn a distance metric that works for all transmitters. Furthermore, independently of the model architecture, greater levels of noise mean more training data is required. Our results show that our dataset was sufficient for this work, but more data may yield further performance improvements.

4.2 Fingerprinting Model

The Siamese network of SATIQ uses the encoder portion of the autoencoder to produce embeddings of two inputs, which are then compared to one another using an angular distance function. A triplet loss term encourages the model to produce embeddings that are close to one another for messages from the same transmitter, and different for messages from different transmitters:

$$d(u, v) = 1 - \frac{u \cdot v}{\|u\| \|v\|}$$

$$L(a, p, n) = \max(d(a, p) - d(a, n) + \alpha, 0)$$

This takes an anchor a , a positive example p belonging to the same class as the anchor, and a negative example n of a different class, and encourages the distance from a to p to be less than a to n . A margin α ensures that effort is not wasted on optimizing triplets for which this is already the case. This is illustrated in Figure 5.

Convolutional layers are used in the encoder to enable identification of position-invariant features and reduce the overall size of the model. For the comparator, we compute the angular distance between the embeddings – this tends to work better than L2 distance for high-dimensional data. A diagram of the fingerprinting model architecture used in this paper can be seen in Figure 4.

The inclusion of convolutional layers reduces the model’s size, and allows it to extract position-invariant features from the waveform. We also use separate layers for the in-phase (I) and quadrature (Q) portions of the signal – although the components are tightly coupled to one another, we find that they express different features and the model is able to perform better when the two are separated.

Following the convolutional layers, we concatenate the outputs and flatten, before using a fully connected layer to reduce the output to the correct size. The decoder uses a very similar architecture to the encoder, composed of alternating upsampling and convolutional layers.

4.3 Data Collection

For a fingerprinting model to be effective, a good dataset is essential. Community projects such as *SatNOGS* (an open-source network of ground stations [41]) provide data from a wide range of satellites. However, in this paper we collect our own data for a couple of important reasons: firstly, collecting our own data enables us to capture at a significantly higher sample rate than existing datasets, providing a good foundation for a fingerprinting system. We also capture signals from a specific constellation rather than a collection of individual satellites; by doing this, we ensure that the signal modulation and protocol does not vary between messages. We can therefore guarantee that the message header is always the same between messages, so differences in the captured waveform will be caused only by hardware differences and channel noise – the contents are always identical at the bit level. This gives us a consistent baseline upon which a fingerprinting model can be built.

We focus on the *Iridium* constellation, used in telecommunications. This constellation has a number of useful properties:

- The constellation contains a large number of satellites (66, each with 48 transmitters [42]), providing sufficient variety within the dataset;
- The transmitter hardware on each satellite is identical, so a fingerprinter will need to distinguish satellites only through differences introduced at time of manufacture, rather than distinguishing between entirely different components;
- The communication protocols are known and well documented [42], so no reverse engineering is required;
- Downlinked transmissions can be received using cheap and widely available COTS hardware.

Other constellations are available which satisfy some of these properties, most notably *Starlink* and Planet Labs’ *Dove* constellation; transmissions from these are more difficult to capture, but could be useful for future evaluations of the versatility of our system. Furthermore, existing research in fingerprinting also uses Iridium satellites [13], providing evidence that some fingerprinting is plausible in this context and providing a baseline to which fingerprinting systems can be compared. We argue that testing on Iridium is sufficient to demonstrate and validate the effectiveness of SATIQ, as the properties of the signal used by SATIQ are not constellation-specific.

For the data collection itself, we use the following hardware:⁴

- Iridium Beam active antenna (1245 USD, Beam Communications)
- Mini-Circuits ZKL-33ULN-S+ low-noise amplifier (209 USD, Mini-Circuits)
- NooElec DC block (20 USD, Amazon)
- Mini-Circuits VBF-1560+, 1500–1620 MHz band pass filter (44 USD, Mini-Circuits)

⁴Prices are as recorded on 2023-04-21, and may not reflect current prices.




Figure 6: An overview of the data collection system. Signals are captured by the SDR and sent to GNURadio for processing. Raw IQ samples, demodulated bytes, and fully decoded message data are all sent to a database.




Figure 7: The distribution of the number of Iridium messages received per transmitter.

- USRP N210 SDR (3354 USD, Ettus Research)
- UBX 40 USRP daughterboard (1732 USD, Ettus Research)

The SDR is connected to a computer running *GNU Radio*, a software library to aid digital signal processing. We use components from the *gr-iridium* library, created and maintained by the Chaos Computer Club München e.V., to demodulate and decode messages [43]. Figure 6 illustrates our full data collection and processing pipeline.

Iridium downlink messages have a number of different message types, one of which is the *Iridium Ring Alert* (IRA) message – this is the only type of message we collect. These messages contain diagnostic information about the satellite, including the satellite ID, beam ID (identifier of the current transmitter), position, and altitude. They also contain the anonymous identifiers of subscribers currently receiving incoming paging calls. The messages are openly broadcast and do not contain any personally identifiable information, so we can decode and collect them without raising any ethical concerns.




Figure 8: A message header received from an Iridium satellite, shown in the time domain (left) and as a constellation plot (right).

We save the raw IQ samples from the message headers in a database, alongside demodulated bytes and decoded message contents. This gives us a consistent dataset of raw message headers which can be used for fingerprinting, using the decoded messages to label the data.

Our antenna was located on the roof of a building in Oxford, UK, giving it a good view of the sky in all directions. We collected messages at 25 MS/s, giving us an oversampling rate of 1000 (Iridium messages have a symbol rate of 25 000 symbols per second). This allows us to capture the high-frequency features characteristic to a transmitter for effective fingerprinting. In total we collected 1 010 464 messages over 23 days. For each transmitter we collected up to 872 messages, with a mean of 203 messages per transmitter – the distribution of message count per transmitter can be seen in Figure 7. We also collected a further 694 738 messages over 17 days immediately following the training data collection – this data will be used for evaluation in Section 5.

An example of the collected data can be seen in Figure 8. We can see the signal encodes 8 QPSK symbols, corresponding to the bit sequence for the Iridium message header: 11 00 00 11 11 00 11 00. However, unlike a constellation plot at 1 sample per symbol, we can see the movement between the two symbols and the impairment on the signal. It is clear that there is significant impairment; this is likely caused by a combination of channel noise, multipath distortion, and hardware characteristics of the transmitter. Our goal is to isolate the last of these, ignoring the channel noise and other characteristics – these will be the same between transmitters, and not useful in this context.

4.3.1 Data Preprocessing. We opt to use a minimal amount of preprocessing to avoid destroying data which might be useful for fingerprinting. On top of the band-pass filtering and phase synchronization performed by *gr-iridium*, we scale each waveform in the dataset so values range between $-1 - i$ and $1 + i$. This makes visualization easier, and removes magnitude as an additional factor the model needs to learn. We also remove all messages which do not decode as valid IRA messages – although this removes a large number of messages, it ensures that all data is labeled and eliminates the noisiest messages which do not properly decode.




Figure 9: The performance of the base SATIQ model, evaluated by comparing messages from different pairs of satellites without analyzing the attack case.

This leaves us with messages that are the most likely to contain meaningful identifying factors.

Finally, we shuffle the data and process it into “TFRecord” files – this format is optimized for use in TensorFlow, storing data as raw protocol buffers. This ensures we can load the data efficiently and minimize RAM usage by reading mostly from disk.

4.3.2 Dataset Construction. After data preprocessing, we split the data into training, validation, and testing sets, and use a generator to produce batches of examples. We set aside 50 000 messages for each of the validation and testing datasets, resulting in a training:validation:testing split of 90:5:5. This ensures we have sufficient data to evaluate the model without significantly reducing the size of the training dataset.

Since we are using a triplet loss function, we produce batches of inputs with 32 messages (4 from each of 8 different transmitters) – this ensures the loss function can select a large number of triplets in each batch.

4.4 Model Training

Model training was performed on a machine with the following hardware:

- 2× Intel Xeon E5-2660 v3 CPU (20 threads, 3.30 GHz)
- Nvidia TITAN RTX GPU (1770 MHz, 24 GB VRAM)
- 128 GB DDR4 RAM (2133 MHz)

On this hardware, up to 3 model configurations can be trained simultaneously. All models are trained for 100 epochs, taking 12–24 hours.

Running a trained model takes significantly fewer resources, taking a fraction of a second to verify messages – this makes it possible to run SATIQ in real-time, verifying messages as they arrive. Furthermore, much less RAM is required, as the training dataset does not need to be loaded. Hardware requirements are discussed further in Section 5.4.

4.5 Model Optimization

Figure 9 shows the base performance of SATIQ. In this section, models are trained and tested on the problem of differentiating satellites in the dataset, in order to achieve better performance later




Figure 10: ROC curves as we compare each incoming message to a larger number of anchors, taking the mean distance.




Figure 11: The AUC performance of the fingerprinting model as the size of the training dataset is changed. In each instance, the messages with the least noise are kept.

on when identifying adversaries in an attack scenario. Results in this section are primarily to assess relative performance and fine-tune the system. When assessing performance, we focus on two key metrics:

- *Equal Error Rate (EER)*: the error rate when the false positive rate (FPR) and false negative rate (FNR) are equal.
- *Area Under Curve (AUC)*: the area under the Receiver Operating Characteristic (ROC) curve, obtained by plotting the true positive rate (TPR) against the FPR. This can be intuitively thought of as the probability that the system can distinguish between two inputs of different classes [44].

In computing each of these, we vary the distance threshold below which we accept two messages as being from the same transmitter. By raising this threshold, we accept a greater number of legitimate messages, but open the system up to easier spoofing attacks. Conversely, by lowering the threshold the system is made more secure, but legitimate messages are less likely to be accepted as such.

Our base model has an AUC of 0.627 and EER of 0.405. This is sufficient performance to demonstrate the feasibility of our techniques, particularly in the more difficult case of distinguishing satellite transmitters with identical hardware. In Section 5 we go on




Figure 12: Hardware setup for the replay attacks. Raw samples are captured using an SDR, then replayed directly into the fingerprinting system’s SDR over a cable.

to demonstrate that performance is significantly better in a replay attack scenario, confirming its usefulness in a security context. For the rest of this section we continue to assess performance on the original (non-attacked) dataset.

4.5.1 Multiple anchors. To improve performance over the base system, we can compare each incoming message to a larger number of “anchors” (known messages from that transmitter), taking the mean distance between the message and each anchor in the embedding space. The results of this are shown in Figure 10. By taking 16 anchors for each incoming message, we can achieve an EER of 0.350 and AUC of 0.698 – a significant improvement! In practice, this can be implemented by saving a larger set of messages from each known transmitter, or by comparing multiple consecutive messages to the same set of anchors. Both of these techniques are practical – our observations suggest that during an Iridium phone call or web connection approximately 11 packets are exchanged per second, so an attacker will need to spoof many packets to have a meaningful impact on the victim. Such an attack would certainly be picked up by SATIQ, even if multiple consecutive messages are compared. We explore attack scenarios further in Section 5.

4.5.2 Reduced dataset. As satellites pass overhead, their down-linked signals are subject to different amounts of attenuation due to changing distances and different thicknesses of the atmosphere along the direct path between the satellite and the observer. As a result, our dataset contains messages of differing amplitudes and levels of noise. Although all messages are normalized before use, this is likely to have an effect on performance; it will be more difficult to extract identifying information from signals with high levels of background noise.

We can try to mitigate this by filtering the dataset, removing the messages with the greatest levels of noise. For this experiment we produced 9 datasets with different percentages of messages from the original dataset removed. Figure 11 shows the AUC of models trained on these smaller datasets. We can see that performance is greatest with the dataset composed of the least noisy 30%, with an AUC of 0.638 and EER of 0.407 – below this, the negative effect of the dataset’s small size exceeds the benefit provided by using cleaner inputs.

Care must be taken when using this approach in practice, as operators will need to decide whether messages above this noise

TPR	FPR	Threshold
0.999	0.978	1.343
0.990	0.818	1.198
0.950	0.307	1.055
0.900	0.150	0.977
0.861	0.100	0.932
0.805	0.050	0.885
0.672	0.010	0.795
0.424	0.000	0.665

Table 1: True positive (true accept) rates and false positive (false accept) rates for key threshold values, tested on replayed messages. Messages are tested against 32 anchors, and the mean distance is taken.

threshold are accepted or rejected by default, or wait for a suitably clean message before the transmitter can be authenticated. If multiple satellites are visible at once, the ground system can connect to the satellite with the best SNR – this is how real-world satellite systems work in practice, connecting to satellites that are more directly overhead. For the remainder of this paper, we use the full dataset.

5 EVALUATION

In the previous section we designed and trained the SATIQ system to distinguish between different legitimate satellite transmitters, without considering a real-world attack scenario. In this section we evaluate SATIQ for its robustness against replay attacks, extensibility to new satellite transmitters, and stability over time. We also discuss further techniques for assessing these types of model, including transferrability to new satellite constellations. Finally, we consider the practical questions surrounding a real-world deployment of SATIQ.

5.1 Security

We first evaluate the security properties of SATIQ, assessing its performance under an attack scenario. The simplest such scenario involves merely swapping the identifiers of transmitters in our existing dataset. This models an attack from within a compromised system, in which an attacker has gained control of a satellite currently in use – this matches our training scenario and results. As discussed in Section 3, this scenario is somewhat unrealistic, and not our primary concern.

The more interesting case is to evaluate SATIQ’s robustness under real-world replay attacks. We can realistically evaluate this by performing replay attacks over-the-wire. These do not perfectly model real-world spoofing and replay attacks, but any concessions made are in the attacker’s favor – for instance, they do not have to account for path loss or background noise as they would in a radio setting. A similar result could be achieved through the use of an RF-shielded box, but this has the potential to introduce further impairments through reflections and other effects. Our experimental setup is as shown in Figure 12. We first capture Iridium messages at 25 MS/s, saving the raw IQ samples to a file – this provides us with a dataset of samples identical to what the SDR




Figure 13: ROC curves showing the system’s performance when detecting replayed messages. Performance is significantly better than our training results, achieving a maximum AUC of 0.946.

would normally receive. We then replay these samples over a wire connected to the “victim” SDR, feeding the captured messages into the fingerprinting system.

For each replayed message, we take a number of “known good” messages from the same transmitter from our testing dataset. We randomly select a number of these messages to be our anchors using a “shuffle split” strategy. We compare the anchors to the replayed messages to obtain the false positive rate, and to the other known good messages to get the true positive rate. The results of this experiment are shown in Figure 13. We can see that SATIQ performs significantly better in this scenario, with a base AUC of 0.788. When we compare each message to 32 anchors this performance increases even further, with an AUC of 0.946 and an EER of 0.120! This indicates that the attacker’s SDR has introduced its own fingerprint, distorting the message and altering its features.

Furthermore, this performance is good enough to deploy in a real-world system – by adjusting the acceptance threshold we can achieve a high true acceptance rate while minimizing the number of spoofed messages that are accepted. These results are summarized in Table 1. By setting the threshold such that 90% of legitimate messages are accepted, we accept only 15% of the attacker’s messages. This performance is good enough to use in a real-world setting, particularly if we continuously fingerprint messages over the course of a communication session, taking the average acceptance rate over time as an indicator of attack – in order to have a meaningful impact the attacker will need to spoof multiple messages, which significantly raises the likelihood of detection.

This attack scenario assumes a well-equipped adversary with access to a high-end SDR, and eliminates all the difficulties of over-the-air replay attacks. Despite all these concessions, we are still able to detect the attack in the majority of cases. With an even higher budget (to transmit at an even higher sample rate) and careful effort to eliminate noise introduced by the radio, it will certainly be possible to circumvent this system [31], but our results show that it will take a concerted effort to do so – simple message replay is not




Figure 14: ROC curve showing the system’s performance on transmitters it has never seen before.

enough. We can therefore exclude a large proportion of attackers with all but the highest budgets, granting a real-world security benefit to ground systems.

5.2 Extensibility

Next we look at how well SATIQ performs on transmitters it has never seen before. This is of particular importance in satellite constellations, where satellites may need to be replaced at any time, and we want to minimize time and effort spent retraining the model. To test SATIQ’s extensibility, we trained the fingerprinting model on a dataset with 50 transmitters removed. The results of this analysis are shown in Figure 14. As expected, the base performance roughly matches our original model’s testing performance – an AUC of 0.637 when testing against 16 anchors. When tested on the 50 transmitters removed from its training data, the performance drops slightly to an AUC of 0.587. With this level of performance we can continue to use the system with slightly reduced accuracy – in the context of satellites, this is better than existing classifier-based systems, which require full retraining each time a satellite is launched or replaced. With further training and tuning, it may be possible to completely eliminate even this small performance drop, making a system that can be used indefinitely as transmitters are added and replaced.

It may also be possible to transfer a trained SATIQ model to a completely new satellite constellation, since many of the signal impairments will be common between hardware configurations. With a small amount of retraining, similar performance might be achieved across a wide range of satellite systems. Such analysis is beyond the scope of this paper, as it will require a new dataset for the second constellation.

5.3 Time Stability

Finally, we evaluate SATIQ’s stability over time. Existing works in radio fingerprinting observe a decrease in performance when there is a time gap between the training and testing data [45, 46]. This is thought to be caused by changes in the conditions of the wireless channel over time. To assess the effect of this on SATIQ, we test our




Figure 15: The relationship between the time difference between messages and their distance in the embedding space, separated by messages from the same transmitter and different transmitters.

model on the additional dataset collected after the training data, comprising 694 738 messages over 17 days.

Figure 15 shows how the distance between two messages in embedding space is affected by the time difference between the two messages. There is not a strong correlation between these two factors – further analysis is required to determine the exact relationship between time difference and performance.

We tested our trained model on the final 24 hours of the dataset, comprising 43 601 messages with a 16 day time gap between this data and our training dataset. As expected, we observe a moderate decrease in performance when comparing messages to anchors from our original test dataset, with an AUC of 0.580. However, if we use fresh anchors from the new dataset the overall AUC improves to 0.615 – this is only a small decrease from the training results, which compare older anchors to old messages. This performance can be improved further by testing against multiple anchors as in Section 4.5.1 – Figure 16 illustrates this, showing that a testing AUC of 0.659 can be achieved with an EER of 0.385 (close to our original AUC and EER of 0.698 and 0.350).

We therefore propose that a deployed system based on SATIQ replaces its anchor messages (used for testing incoming messages) from time to time, ensuring freshness and maximizing performance. However, caution must be exercised to ensure the anchors are not inadvertently replaced with spoofed messages, as this would entirely circumvent the security of the fingerprinting system. Secure refreshing of the anchors could be achieved by periodically verifying messages through other means – for instance, by precisely measuring angle of arrival.

5.4 Deployment Considerations

In this section we discuss what it would look like to deploy SATIQ in a real-world system. As we mention in Section 4.4, a trained SATIQ




Figure 16: ROC curves depicting performance on new data collected after the training data, comparing each incoming message against multiple anchors. At 32 anchors per message, an AUC of 0.659 is achieved.

model can be run in real-time, verifying messages as they arrive. We validate this by running the model on lightweight hardware. Our data collection observed roughly 0.5 ring alert messages per second, or 6 to 7 messages per second across all message types. This gives us approximately 150 ms to validate each incoming message. On the Intel Xeon CPU used for training (disabling the GPU), the model takes 4.85 s to validate a batch of 543 messages (8.9 ms per message), with a maximum RAM usage of 1.7 GB – this is more than fast enough to validate all incoming messages.

We also ran SATIQ on a Raspberry Pi 4B single-board computer, with 4 GB of RAM and a 1.8 GHz quad-core processor. On this hardware, it took 93.25 s to validate the same 543 messages (171 ms per message), using 2.0 GB of RAM. This demonstrates that SATIQ can run on lightweight hardware – although slightly too slow to validate all messages in real time, this can certainly validate incoming ring alert messages, and could run even faster with the addition of a dedicated AI accelerator.

We must also consider how deployed devices can capture high sample rate message fingerprints. Our current setup requires a moderately powerful computer to perform the necessary high sample rate message synchronization and decoding, but this can be reduced significantly with the use of a dedicated FPGA demodulator. Combined with the good performance on low-power CPUs shown above, it will be possible to integrate SATIQ into even lightweight devices – however, some instances may require the addition of hardware.

Additionally, it is important to decide what action is taken when a (potentially legitimate) message is rejected by SATIQ. The action taken depends on the level of security required. For some systems it may be sufficient to simply notify the user, but still accept the message, providing increased awareness of potential attacks without risking any impact to service through rejected legitimate messages. Alternatively, higher security applications may choose to reject all messages flagged by SATIQ, accepting the increased packet loss in exchange for greater security against spoofing and replay attacks.

It may be possible for some systems to request message retransmission in this case, to ensure messages are still delivered – this may occur at the transport or application layer, or systems may be modified to allow retransmission requests (if this has not already been implemented).

Finally, we must consider the logistics of such a deployment. As discussed in Section 5.3, we propose that SATIQ-based systems periodically refresh their anchor messages to maximize performance. An additional benefit of this approach is that no centralized database of fingerprints is required, significantly reducing the complexity of deploying SATIQ – there is no need to gather or maintain a database of fingerprints, and devices do not need to connect to a server to update their databases or verify incoming messages. SATIQ-based systems can therefore also be developed by either the satellite operator or the receiver manufacturer, or by a third party building devices to integrate with existing ground systems.

6 FUTURE WORK

Our research has revealed several promising avenues of future research. Firstly, it is likely that better performance can be achieved through fine-tuning and a larger dataset, using the same base model architecture as SATIQ. In particular, a dataset collected using multiple receiver configurations would likely produce a system that transfers more readily to different ground systems.

Alongside collecting additional data, it will also be useful to implement some of the other evaluation techniques, particularly in assessing the extent to which a trained system can be transferred to another constellation, and what degree of retraining is required. Furthermore, it would be particularly beneficial to standardize the other analyses such that they could be applied to other fingerprinting techniques. We currently have no method of empirically comparing fingerprinting techniques to each other in terms of their security properties – a standard suite of tests would remedy this. We consider a comparison between fingerprinting systems to be out of scope for this paper, as each fingerprinting technique relies on different kinds of data. For example, PAST-AI uses full messages rather than just message headers [13], and transient fingerprinting requires precise transient synchronization.

Another promising area of research would be to assess the effectiveness of fingerprinting in conjunction with other methods of spoofing detection, such as assessing SNR or distortion. Multiple fingerprinting methods could also be used in concert, providing even greater effectiveness than any model alone. However, some methods are likely to learn the same characteristics as each other (providing no mutual benefit), so a full analysis is needed in order to understand which methods are effective together.

Finally, it would be useful to assess the effectiveness of fingerprinting in systems which already have some amount of authentication. For instance, such an analysis could evaluate fingerprinting as a preventative measure against GNSS message delay/advancement attacks [12]. This would demonstrate that fingerprinting is not just useful in legacy systems, but has concrete usefulness even in new satellite systems.

7 CONCLUSION

In this paper we have contributed new methods towards radio signal fingerprinting in the context of satellite transmitters, providing novel techniques which can be used to build high sample rate fingerprinting systems. We have succeeded in demonstrating that satellite signals at high sample rates contain sufficient identifying information, and confirmed that our techniques combining autoencoders and Siamese models are feasible for fingerprinting.

We have also provided a large dataset of captured message headers from Iridium satellites, which can be used for further research and testing in satellite transmitter fingerprinting. This lays good groundwork for future research in this area.

We have laid out a clear path for future experimental work, both in increasing the accuracy of existing models, and in evaluating the performance of fingerprinting systems from a security context. Finally, we have discussed the potential for the extensibility of trained fingerprinting models, working across multiple constellations with little or no retraining.

Our work shows that high sample rate fingerprinting is possible and can improve the security against spoofing and replay attacks for the vast majority of low-budget attackers significantly. This will enable us to continue to use legacy satellite systems and their data with an increased degree of confidence in their integrity. We verified this by demonstrating that our system can successfully detect replay attacks even in the scenario of a powerful attacker sending messages via a wired channel, achieving an Equal Error Rate of 0.120 and ROC AUC of 0.946.

ACKNOWLEDGMENTS

We would like to thank armasuisse Science + Technology for their support during this work. Joshua and Sebastian were supported by the Engineering and Physical Sciences Research Council (EPSRC).

REFERENCES

- [1] Great Scott Gadgets. 2021. HackRF One. Retrieved Sept. 27, 2022 from <https://greatscottgadgets.com/hackrf/one/>.
- [2] Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. 2015. On the Security of the Automatic Dependent Surveillance–Broadcast Protocol. *IEEE Communications Surveys Tutorials*, 17, 2, 1066–1087. doi: 10.1109/COMST.2014.2365951.
- [3] Marc Lichtman, Roger Piqueras Jover, Mina Labib, Raghunandan Rao, Vuk Marojevic, and Jeffrey H Reed. 2016. LTE/LTE-A jamming, spoofing, and sniffing: threat assessment and mitigation. *IEEE Communications Magazine*, 54, 4, 54–61.
- [4] Gyuhoong Lee, Jihoon Lee, Jinsung Lee, Youngbin Im, Max Hollingsworth, Eric Wustrow, Dirk Grunwald, and Sangtae Ha. 2019. This is Your President Speaking: Spoofing Alerts in 4G LTE Networks. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, 404–416.
- [5] João Gaspar, Renato Ferreira, Pedro Sebastião, and Nuno Souto. 2020. Capture of UAVs Through GPS Spoofing Using Low-Cost SDR Platforms. *Wireless Personal Communications*, 115, 2729–2754.
- [6] NASA. 2022. Fire Information for Resource Management System (FIRMS). Retrieved May 3, 2022 from <https://earthdata.nasa.gov/earth-observation-data/near-real-time/firms>.
- [7] Esri. 2022. Esri releases updated land-cover map with new sets of global data. Retrieved May 3, 2022 from <https://www.esri.com/about/newsroom/announcements/esri-releases-updated-land-cover-map-with-new-sets-of-global-data/>.
- [8] Meta. 2022. High resolution population density maps. Retrieved May 3, 2022 from <https://dataforgood.facebook.com/dfg/tools/high-resolution-population-density-maps>.
- [9] Cloud to Street. 2022. Cloud to Street. Retrieved May 3, 2022 from <https://www.cloudtostreet.ai/>.

- [10] sam210723. 2020. Receiving images from geostationary weather satellite GEO-KOMPSAT-2A. Retrieved May 9, 2022 from <https://vksdr.com/xrit-rx>.
- [11] sam210723. 2018. COMS-1 LRIT key decryption. Retrieved May 9, 2022 from <https://vksdr.com/lrit-key-dec>.
- [12] Maryam Motallebighomi, Harshad Sathaye, Mridula Singh, and Aanjhan Ranganathan. 2022. Cryptography Is Not Enough: Relay Attacks on Authenticated GNSS Signals. *arXiv preprint arXiv:2204.11641*. arXiv: 2204.11641.
- [13] Gabriele Oligeri, Simone Raponi, Savio Sciancalepore, and Roberto Di Pietro. 2020. PAST-AI: Physical-layer Authentication of Satellite Transmitters via Deep Learning. *arXiv:2010.05470 [cs]*, (Oct. 2020). eprint: 2010.05470.
- [14] Marc Lichtman. 2021. IQ Sampling. In *PySDR: A Guide to SDR and DSP Using Python*.
- [15] NASA. 2022. X-band Direct Readout Sites Worldwide. Retrieved Sept. 27, 2022 from <https://directreadout.sci.gsfc.nasa.gov/?id=dspContent%5C&cid=78>.
- [16] Eric Jedermann, Martin Strohmeier, Matthias Schäfer, Jens Schmitt, and Vincent Lenders. 2021. Orbit-based authentication using tdoa signatures in satellite networks. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 175–180.
- [17] Mohsen Riahi Manesh, Jonathan Kenney, Wen Chen Hu, Vijaya Kumar Devabhaktuni, and Naima Kaabouch. 2019. Detection of GPS Spoofing Attacks on Unmanned Aerial Systems. In *2019 16th IEEE Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 1–6. ISBN: 1-5386-5553-5.
- [18] Damian Miralles, Aurelie Bornot, Paul Rouquette, Nathan Levisne, Dennis M Akos, Yu-Hsuan Chen, Sherman Lo, and Todd Walter. 2020. An Assessment of GPS Spoofing Detection Via Radio Power and Signal Quality Monitoring for Aviation Safety Operations. *IEEE Intelligent Transportation Systems Magazine*, 12, 3, 136–146.
- [19] Naeimeh Soltanieh, Yaser Norouzi, Yang Yang, and Nemai Chandra Karmakar. 2020. A Review of Radio Frequency Fingerprinting Techniques. *IEEE Journal of Radio Frequency Identification*, 4, 3, 222–233.
- [20] Jeyanthi Hall, Michel Barbeau, and Evangelos Kranakis. 2003. Detection of transient in radio frequency fingerprinting using signal phase. *Wireless and Optical Communications*, 13–18.
- [21] Lianfen Huang, Minghui Gao, Caidan Zhao, and Xiongpeng Wu. 2013. Detection of Wi-Fi transmitter transients using statistical method. In *2013 IEEE International Conference on Signal Processing, Communication and Computing (ICSPCC 2013)*. IEEE, 1–5. ISBN: 1-4799-1027-9.
- [22] K. J. Ellis and N. Serinken. 2001. Characteristics of radio transmitter fingerprints. *Radio Science*, 36, 4, (July 2001), 585–597. doi: 10.1029/2000RS002345.
- [23] Kasper Bonne Rasmussen and Srdjan Capkun. 2007. Implications of radio fingerprinting on the security of sensor networks. In *2007 Third International Conference on Security and Privacy in Communications Networks and the Workshops-SecureComm 2007*. IEEE, 331–340. ISBN: 1-4244-0974-8.
- [24] Mahsa Foruhandeh, Abdullah Z. Mohammed, Gregor Kildow, Paul Berges, and Ryan Gerdes. 2020. Spotr: GPS spoofing detection via device fingerprinting. In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '20)*. Association for Computing Machinery, (July 2020), 242–253. doi: 10.1145/3395351.3399353.
- [25] Irwin O Kennedy, Patricia Scanlon, Francis J Mullany, Milind M Buddhikot, Keith E Nolan, and Thomas W Rondeau. 2008. Radio transmitter fingerprinting: A steady state frequency domain approach. In *2008 IEEE 68th Vehicular Technology Conference*. IEEE, 1–5. ISBN: 1-4244-1722-8.
- [26] Joshua Bassej, Damilola Adesina, Xiangfang Li, Lijun Qian, Alexander Aved, and Timothy Kroecker. 2019. Intrusion detection for IoT devices based on RF fingerprinting using deep learning. In *2019 Fourth International Conference on Fog and Mobile Edge Computing (FMEC)*. IEEE, 98–104. ISBN: 1-72811-796-8.
- [27] Francesco Restuccia, Salvatore D’Oro, Amani Al-Shawabka, Mauro Belgiovine, Luca Angioloni, Stratis Ioannidis, Kaushik Chowdhury, and Tommaso Melodia. 2019. DeepRadiolD: Real-time channel-resilient optimization of deep learning-based radio fingerprinting algorithms. In *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 51–60.
- [28] Kunal Sankhe, Mauro Belgiovine, Fan Zhou, Luca Angioloni, Frank Restuccia, Salvatore D’Oro, Tommaso Melodia, Stratis Ioannidis, and Kaushik Chowdhury. 2020. No Radio Left Behind: Radio Fingerprinting Through Deep Learning of Physical-Layer Hardware Impairments. *IEEE Transactions on Cognitive Communications and Networking*, 6, 1, (Mar. 2020), 165–178. doi: 10.1109/TCCN.2019.2949308.
- [29] ÖH Tekbaş, Oktay Üreten, and Nur Serinken. 2004. Improvement of transmitter identification system for low SNR transients. *Electronics Letters*, 40, 3, 182–183.
- [30] Weidong Wang and Lu Gan. 2022. Radio Frequency Fingerprinting Improved by Statistical Noise Reduction. *IEEE Transactions on Cognitive Communications and Networking*.
- [31] Boris Danev, Heinrich Luecken, Srdjan Capkun, and Karim El Defrawy. 2010. Attacks on physical-layer identification. In *Proceedings of the Third ACM Conference on Wireless Network Security*, 89–98.
- [32] Davide Chicco. 2021. Siamese Neural Networks: An Overview. *Artificial Neural Networks*, 73–94.
- [33] Jinting Zhu, Julian Jang-Jaccard, and Paul A Watters. 2020. Multi-Loss Siamese Neural Network With Batch Normalization Layer for Malware Detection. *IEEE Access*, 8, 171542–171550.
- [34] Cheng Zhang, Wu Liu, Huadong Ma, and Huiyuan Fu. 2016. Siamese Neural Network Based Gait Recognition for Human Identification. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2832–2836. ISBN: 1-4799-9988-1.
- [35] Mingtao Pei, Bin Yan, Huiling Hao, and Meng Zhao. 2023. Person-Specific Face Spoofing Detection Based on a Siamese Network. *Pattern Recognition*, 135, 109148.
- [36] Kaavya Sriskandaraja, Vidhyasaharan Sethu, and Eliathamby Ambikairajah. 2018. Deep Siamese Architecture Based Replay Detection for Secure Voice Biometric. In *Interspeech*, 671–675.
- [37] Yu Mao, Yang-Yang Dong, Ting Sun, Xian Rao, and Chun-Xi Dong. 2021. Attentive Siamese Networks for Automatic Modulation Classification Based on Multitiming Constellation Diagrams. *IEEE Transactions on Neural Networks and Learning Systems*.
- [38] Louis Morge-Rollet, Frédéric Le Roy, Denis Le Jeune, and Roland Gautier. 2020. Siamese Network on I/Q Signals for RF fingerprinting. In *Actes de La Conférence CAID 2020*, 152.
- [39] Zachary Langford, Logan Eisenbeiser, and Matthew Vondal. 2019. Robust Signal Classification Using Siamese Networks. In *Proceedings of the ACM Workshop on Wireless Security and Machine Learning*, 1–5.
- [40] Kian Ahrabian and Bagher BabaAli. 2019. Usage of Autoencoders and Siamese Networks for Online Handwritten Signature Verification. *Neural Computing and Applications*, 31, 9321–9334.
- [41] SatNOGS. 2022. SatNOGS: Open Source global network of satellite ground-stations. Retrieved Oct. 5, 2022 from <https://satnogs.org/>.
- [42] Dan Veeneman. 2021. Iridium: Technical Details. Retrieved Sept. 29, 2022 from <http://www.decodesystems.com/iridium.html>.
- [43] Tobias Schneider and Stefan Zehl. 2022. gr-iridium: GNU Radio Iridium Out Of Tree Module. Chaos Computer Club München. (Sept. 2022).
- [44] Sarang Narkhede. 2018. Understanding AUC - ROC Curve. *Towards Data Science*, 26, 1, 220–227.
- [45] Bechir Hamdaoui and Abdurrahman Elmaghub. 2022. Deep-Learning-Based Device Fingerprinting for Increased LoRa-IoT Security: Sensitivity to Network Deployment Changes. *IEEE Network*, 36, 3, (May 2022), 204–210. doi: 10.1109/MNET.001.2100553.
- [46] Amani Al-Shawabka et al. 2020. Exposing the Fingerprint: Dissecting the Impact of the Wireless Channel on Radio Fingerprinting. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. (July 2020), 646–655. doi: 10.1109/INFOCOM41043.2020.9152529.