Obscuring RF Noise to Combat RF Fingerprinting Surveillance by Adversaries

Zach Smith

This is a proposal to mature a wireless communications feature that would disarm the ability of a surveillance tool, called Radio Frequency (RF) fingerprinting, that identifies a wireless communications device based on its unique, hardware-specific wireless signature.

RF fingerprinting is a technology that has been in development for roughly ten years, and has been aided by recent advancements in machine learning techniques. It currently stands at Technology Readiness Level (TRL) 6 or above, and may already be deployed by US or foreign forces.

The broader impacts are that this technology would enable wireless military and civilian devices to evade the adversarial surveillance tracking technique, RF fingerprinting. The current state-of-the-art RF fingerprinting techniques are currently used by the US military and are researched and likely used by other countries, and could theoretically be deployed on foreign-owned satellites. The new technology we're proposing would disable adversarial forces' ability to use RF fingerprinting to uniquely identify and track US military and civilian wireless devices such as military radios, cell phones, and computers.

Contents

1 Introduction

A United States prime contractor has developed technology to uniquely identify the hardware that transmits RF (cellular, S-band, microwave, WiFi, Bluetooth, etc) using only the noise from the RF signal, and currently stands at TRL 6 or above. This is possible because there are necessarily unique imperfections in the hardware components of transmitters that are introduced during the manufacturing process. These imperfections are revealed in the raw signal the transmitter produces. Extracting the signal imperfections and then assigning a set of imperfections an identity is referred to as a RF fingerprinting. This technology can be used to identify and geolocate RF producing devices, and is currently known to be operationally tested in military and civilian settings to track the location of cell phones and radios, and therefore people over time.

A cell phone is the canonical example of a device that can be tracked using this technology, but other RF producing devices can also be tracked, including satellite dish antennas, software-defined radios such as an MPU5, radios on the Link16 network such as BATS-D, RFID tags and receivers, etc.

Because the underlying technology to achieve this surveillance capability is not classified, it's likely that other militaries either already have developed or will develop and use similar tech. Further, while we're currently enabled to use this technology both on the ground and airborne, space-based implementations could allow the U.S. and other governments to put this technology on satellites, which would enable global RF fingerprinting surveillance by both US and foreign forces.

The development of a small, uniquely parameterized transmitter firmware layer to remove and/or obscure the imperfections revealed in the RF signal would defend against adversarial ability to fingerprint and geolocate devices with this transmitter feature.

2 Current state of RF fingerprinting architecture and capability

2.1 Why do wireless transmitters produce unique noise that can be used for identification?

Modern wireless devices (cell phones, software-defined radios, etc) have the following transmitter architecture:

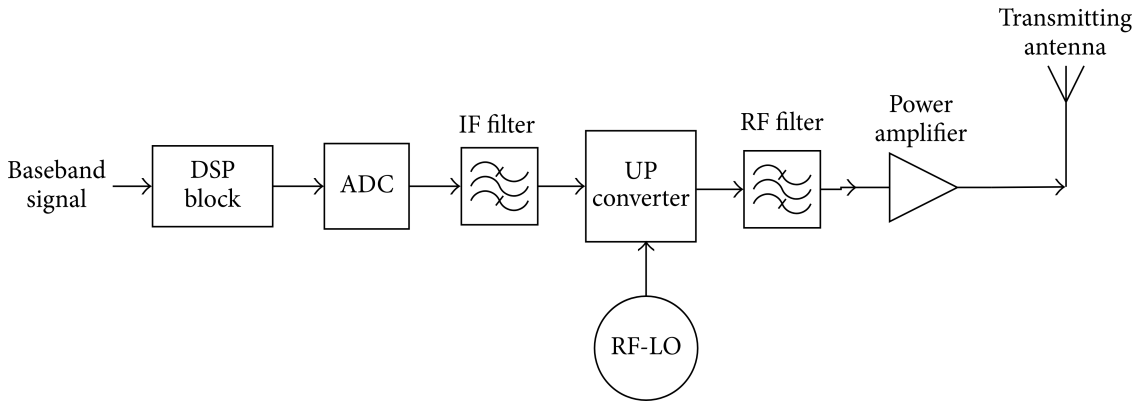

Figure 1: Architecture of a cell phone transmitter. The baseband generates the digital signal to be converted to analog and transmitted. [3]

Wireless devices use a transmitter to convert digital information into an analog signal, and then propagate that signal through space using an antenna. In the process of turning digital information into an analog signal, signal phase and amplitude imperfections are produced by imperfections in the quartz crystal oscillator, IF filter, UP converter, RF filter, amplifier, and antenna.

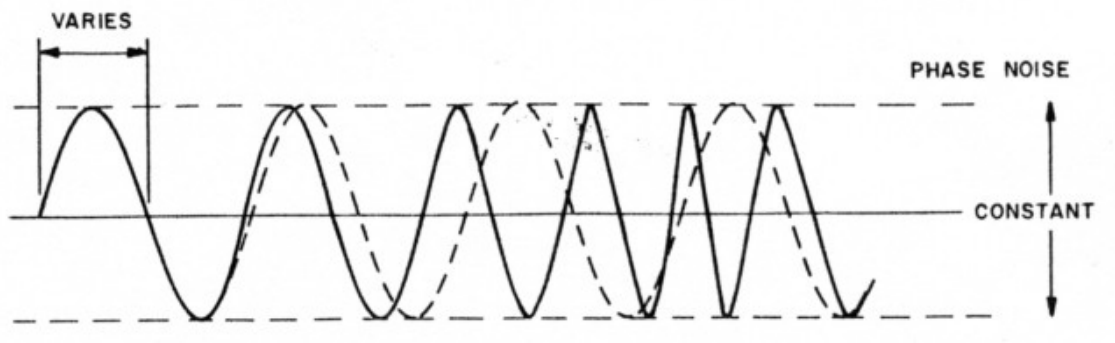

Figure 2: Phase noise [5]

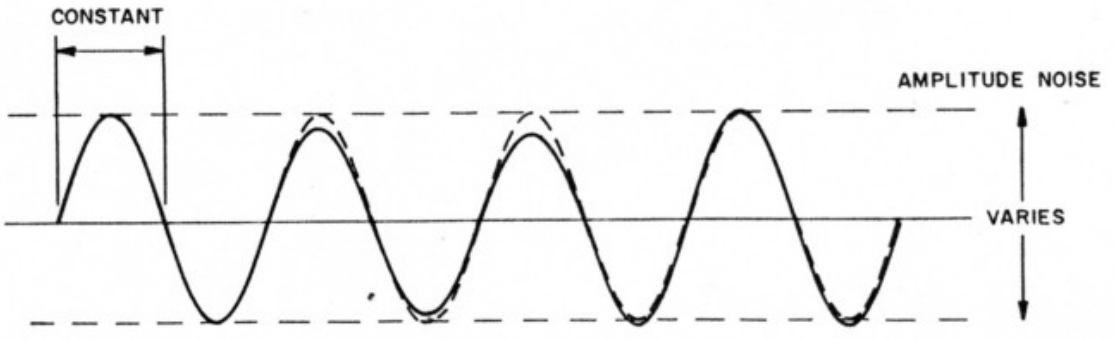

Figure 3: Amplitude noise [5]

These imperfections are part of what makes up the noise in the signal. In normal RF communications, when the RF signal (with noise from the hardware-imperfections) is received, the noise is ignored and the analog signal can be converted to digital. However, RF fingerprinting techniques use this noise to classify and recognize the originating transmitter using their unique noise signature.

2.2 How can the noise be used to uniquely identify a transmitter?

An antenna, digital receiver, and oscilloscope can be used to read the raw signal including the noise.

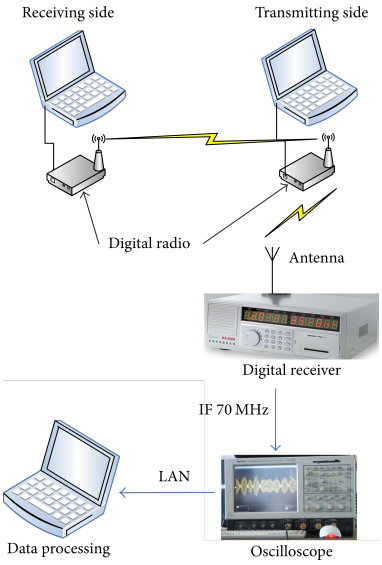

Figure 4: RF fingerprinter lab architecture [3]

Once the noise is extracted and isolated, it must be discretely described and classified. To extract a multidimensional feature vector from a time-series sample, we can calculate its multidimensional permutation entropy.

With the discrete time sequence {x(i), i = 1,2,...,N}, the phase space reconstruction can be calculated:

$$X_i = [x(i), x(i+l), ..., x(i+(m-1)l)]$$

where time delay is l and dimension embedding is m with $m >= 2$. Then there are $m!$ ways to arrange X_i so we define permutation $p_i = p_1 ... p_{m!}$. Permutation entropy, H_p, is

$$H_p = -\sum_{j=1}^{k} p_j \ln p_j$$

Then, we'll choose an unsupervised machine learning classifier such as k-means clustering or hierarchical clustering, we can classify the feature vector to give it a unique identity.

Whether we choose k-means or hierarchical clustering will depend on if we want to test a known or unknown number of RF devices. [3]

3 Prototype, Implementation, and Maturation (Phase II)

The goal of the prototype is to show that we can modify a wireless device so that it can't be uniquely identified using RF fingerprinting techniques when compared to other modified devices. A modified devices RF signature will still be differentiated from an unmodified device, but all modified devices of the same signal type will not be differentiated from each other. This Phase II will create a mature example of this technology, which could be readily applied to military radios and consumer cell phones.

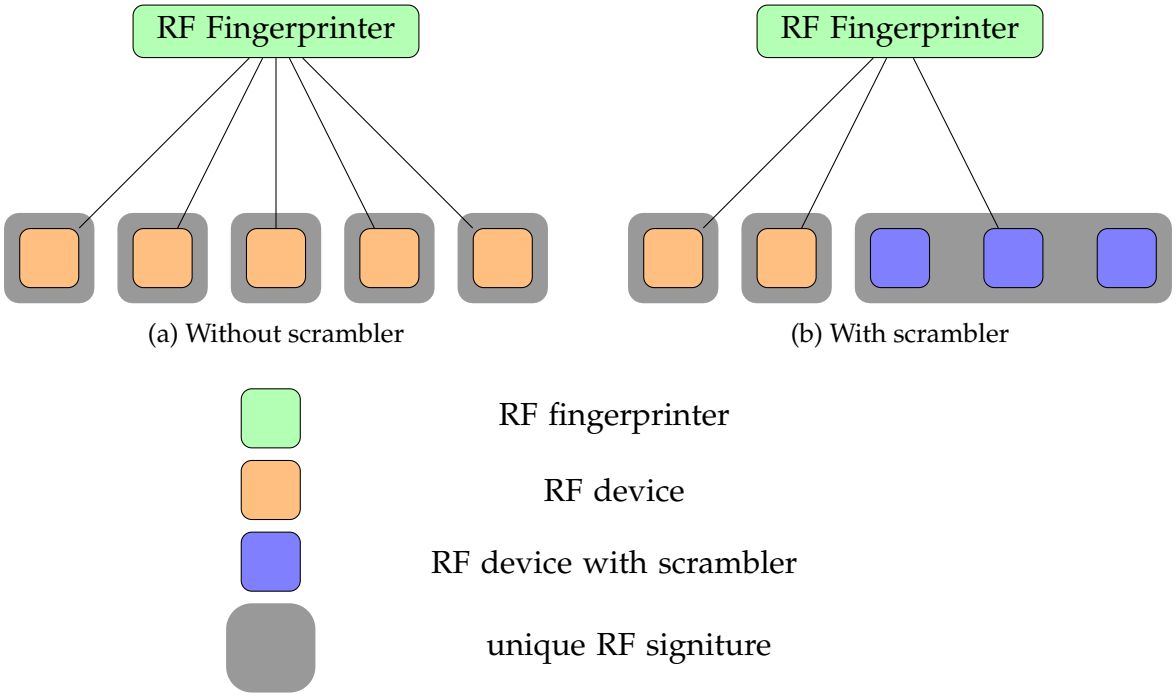

Initially, the scope of the prototype will be limited to a single model of transmitter, one instance and version of firmware, one fingerprint extraction technique, and one type of signal noise (one of phase or amplitude noise).

Core components:

- Transmitter
- Firmware (open source)
- Fingerprinter

3.1 How would we achieve this?

To develop a prototype of this feature, we'll need:

- an RF device with baseband chip with a Digital-to-Analog Converter (DAC) and firmware that can be modified
- an RF fingerprinter for that device

The baseband with DAC is the component that converts digital data to analog data. To modify (either correct or introduce randomness) the analog output, we will modify bits at the digital layer just before the digital information is converted to analog. The hypothesis is that either:

- we'll be able to use individual device noise data to precisely and individually modify that devices digital input stream (originating at the baseband before the DAC), resulting in the elimination of identifiable noise
- introducing some changing randomness to the signal via the baseband digital stream to target minor amplitude, frequency, and phase randomness will cause all devices (of same model or type) to produce the same noise, making each devices noise undifferentiated from other devices with this feature

We anticipate that the first hypothesis is more difficult to test, and with a lower probability of confirming the hypothesis. However, confirming the first hypothesis may result in a higher quality feature. The second hypothesis may be able to produce an equally capable feature with the benefit of simplifying the research and removing the need to have unique correction parameters for each devices baseband firmware. Because the second hypothesis will be much quicker to prototype and test, we will explore this during Phase II.

The immediate challenge is the near-necessarily closed-source requirement of baseband firmware (discussed in the Risks section). To circumvent this issue for the proof-of-concept, we'll use and modify open-source baseband firmware. [6]

To generate signal data and test modifications to the baseband firmware, we'll reimplement open-sourced fingerprinting tools. [8] [7]

Then, using the extracted signal fingerprint data, we'll modify the baseband firmware to correct or fuzz the digital input stream, and test using the fingerprinting techniques.

4 Path to TRL 9 (Phase III)

In a Phase III, we will identify specific RF devices with best product-market fit, reduce technical risk, and have external partnerships to diversify our funding profile.

Further, a Phase III will allow us to develop and test a fully integrated and marketable solution that meets TRL 9 and deliver value to our public and private partners.

5 Risks

5.1 Technical

The technical hypothesis is that RF transmitter firmware can be modified such that the device can still operate as intended, but that an RF fingerprinter cannot uniquely identify the transmitter based on its hardware imperfections.

We identify the following quantified technical risks:

- 10% chance of failing to show evidence that supports or refutes hypothesis
- given we do have evidence, 30% chance that the evidence refutes the hypothesis

This means we identify that there is a cumulative 40% risk that we won't have evidence to support our hypothesis, which is the only metric that matters to show technical promise. However, supporting, refuting, or absence of evidence does not indicate absolute viability or non-viability of the technical solution.

Ultimately, this is a highly tractable engineering pursuit and won't require any novel technical breakthroughs, just sound engineering.

5.2 Regulatory capture / bureaucratic

Baseband processors produce RF in this frequency and in a byte stream that is understood by cell phone transceiver towers. Because of this, the US regulates both the processor and hardware lock on the firmware, so that the process can only run specific firmware that can't be modified. Commercializing this product for cell phones would involve working in a highly regulated environment.

5.3 Closed-source

The proposed firmware layer we develop will likely be part of the component called the baseband processor. In the case of cell phones, these baseband processors are often hardware-locked, meaning that firmware can change, but that once it's changed, the hardware won't allow the new software to run. There are open-source projects where the firmware can be modified, so this won't be an issue during the Phase II prototype, but may become more of an issue as the product matures. This risk can likely be mitigated by working closely with the baseband processor and firmware creators when integrating beyond the initial prototype.

5.4 Commercial

Organizations already in the RF space might be better positioned and have more resources to develop this technology internally. They include the baseband processor / firmware creators, and organizations that integrate RF tech such as L3Harris and Apple, both of which highly value end-user privacy. However, large organizations have many competing priorities and often can't move as quickly as a smaller, focused startup. One of these organizations might be an ideal target acquirer of this technology in the future.

6 Commercialization Plan

6.1 Strategic vision

We envision this technology to be an important privacy feature that aids in protecting our war fighters, military assets, and civilians Position Location Information (PLI) information from foreign adversaries.
This technology cannot stand on it's own as a consumer product and must be incorporated as a feature into current and/or future RF products. The most viable path to successfully using this technology would be to partner with organizations that can use our technology to provide an end-user feature either through licensing and / or acquisition partners.

6.2 Financing / revenue model

6.2.1 Government customers

The US government has many potential beneficiaries of this technology. Because RF fingerprinting can theoretically be used by space-based assets, the US government has an interest in preventing tracking of any mobile RF producing device who's location is sensitive. The Department of Defense would likely be the primary government user of this technology, but would also include cell phones and radios of important individuals and their security teams, boats used by the coast guard, drones used by border and homeland security, etc.

Because of the national security interest, we believe the US government, particularly the Department of Defense, is aligned to support development of this technology.

The Department of Defense (DoD) has bought at least $500 million of radios recently, all with the security vulnerability that allows them to be uniquely identified by RF fingerprinting techniques. [2] [4] [1]

6.2.2 Non-government customers

Consumer product companies may incorporate this security feature to protect their customers' privacy. Apple in particular would be well suited to use this technology as they have developed a strong brand around protecting customer privacy.

The cell phone market is a $80 billion / year industry, and 1.5 billion cell phones sold every year. In addition there are many other products that would benefit from this technology including laptops, wireless headphones, tablets, cars, key fobs, wireless speakers, and many more.

The demand of the non-government sector isn't as pressing as the government need. We predict non-government interests will show support after successful proof-of-concept.

7 Resource allocation

7.1 Work Breakdown Structure

1. **RF fingerprinting lab part selection** : 100 hours
   To build a testing environment, we'll use COTS and open-source tools to create an RF fingerprinter. Assess COTS hardware and open-source software solutions.

2. **Implement a minimal RF fingerprinting lab**: 300 hours
   Design and build the RF signature testing pipeline using the parts / software from part selection. This will allow quicker testing and a tighter feedback loop as we modify the firmware.

3. **Building and installing open-source baseband firmware**: 300 hours
   Instead of writing baseband firmware ourselves, we'll start with an open-source baseband implementation. The quality, documentation, and ease-of-use of open-source software varies. Regardless, starting with an open-source implementation will likely be much faster that writing our own baseband firmware from scratch.

4. **Collecting data / RF fingerprinting unmodified transmitter**: 100 hours
   We'll establish an RF signature baseline, so that when we collect an RF signature after firmware modification, we can compare to the baseline.

5. **Experiment modifications to obscure uniquely identifiable signal noise**: 800 hours
   This is the primary section and most valuable part of the project. Here we'll iterate through the development cycle by modifying the firmware, testing in the RF fingerprinting lab, assessing the results, and repeating this processes as necessary.

6. **Identify Phase III transition partners**: 100 hours
   As we mature the technology and prior to closing out Phase II, we'll establish relationships with DoD and commercial transition partners, and likely use financial instruments such as STRATFI / TACFI to raise funds and align our interests with public and private interests.

7. **Data analysis and report**: 100 hours
   We'll quantify and visualize the test results as a communication tool and for posterity.

8 Broad impact / benefit to society

Modern warfare is fought through the accumulation and manipulation of data. The surveillance of US war fighters by foreign adversaries poses a security risk to national security.

When this technology is developed and deployed, society will benefit because our war fighters and civilians privacy will be more protected. Adversaries will be disarmed of one weapon to access sensitive personnel location data of US soldiers in combat and

US civilians. While there are many other privacy vulnerabilities, we must take a proactive approach to protect against the entire surface area of attack and adversarial surveillance.

9 About the team

Zach Smith is a technologist and engineer. He has led development efforts across wide range of projects, including experimental military networks, agricultural robots, RF antenna design, and computer vision tools.

Notably, he's developed and tested a novel military network architecture as part of an SBIR Phase II grant with Hanscom AFB and USSOCOM alongside Alex Fleming (Principle Investigator) at iMetalx. Zach led the team to successful demonstrations at the Eglin AFB 46th Test Squadron and Myakka Test Range and brought the technology from TRL 1 to TRL 6.

Further, he has developed and implemented farming technology alongside Daniel Theobald of Vecna Robotics, including an RF antenna design concept for farming applications, a solar-powered farming vehicle, an autonomous irrigation system, and a farming robot for organic farms.

During his time at the University of Texas, Austin, he developed a synthetic data pipeline to train and improve a machine learning model used for an aerial computer vision project.

References

[1] Us army to field trellisware technology for all tactical radios that comprise integrated tactical network capability set 21. businesswire.com/news/home/20201030005153/en/US-Army-to-Field-TrellisWare

[2] Usmc orders additional an/prc-160 hf radios from l3harris. naval-technology.com/news/usmc-orders-additional-an-prc-160-hf-radios-from-l3harris/.

[3] Shouyun Deng. Radio frequency fingerprint extraction based on multidimension permutation entropy. *International Journal of Antennas and Propagation*, 2017(1538728).

[4] Mariana Iriarte. U.s. army $406 million software-defined radio contract won by raytheon. militaryembedded.com/comms/sdr/u-s-army-406-million-software-defined-radio-contract-won-by-raytheon.

[5] Rob Kaloeijer. Phase noise and its effects on amateur communications 1. robkalmeijer.nl/techniek/electronica/radiotechniek/hambladen/qst/1988/03/page14/index.html.

[6] open sdr. openwifi. github.com/open-sdr/openwifi.

[7] Smart Home Privacy Project. Robust deep-learning-based radio fingerprinting with fine-tuning. github.com/SmartHomePrivacyProject/RadioFingerprinting.

[8] Luc Wachter. Rf fingerprinting for nfc device identification. github.com/Laykel/nfc-rf-fingerprinting.