Poster: Fingerprinting IoT Devices Using Latent Physical Side-Channels

Justin Feng, Tianyi Zhao, Shamik Sarkar, Dominic Konrad, Timothy Jacques, Danijela Cabric, and Nader Sehatbakhsh
UCLA
Los Angeles, California, USA

Title: Fingerprinting IoT Devices Using Latent Physical Side-Channels.

Authors: Justin Feng, Tianyi Zhao, Shamik Sarkar, Dominic Konrad, Timothy Jacques, Danijela Cabric, and Nader Sehatbakhsh.

Journal: Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, Volume 7, Issue 2.

Link: https://dl.acm.org/doi/abs/10.1145/3596247

DOI: 10.1145/3596247

Full Reference: Justin Feng, Tianyi Zhao, Shamik Sarkar, Dominic Konrad, Timothy Jacques, Danijela Cabric, and Nader Sehatbakhsh. Fingerprinting IoT devices using latent physical side-channels. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 7(2):1–26, 2023.

Abstract: The proliferation of low-end low-power internet-of-things (IoT) devices in "smart" environments necessitates secure identification and authentication of these devices via low-overhead fingerprinting methods. Previous work typically utilizes characteristics of the device's wireless modulation (WiFi, BLE, etc.) in the spectrum, or more recently, electromagnetic emanations from the device's DRAM to perform fingerprinting. The problem is that many devices, especially low-end IoT/embedded systems, may not have transmitter modules, DRAM, or other complex components, therefore making fingerprinting infeasible or challenging. To address this concern, we utilize electromagnetic emanations derived from the processor's clock to fingerprint. We present Digitus, an emanations-based fingerprinting system that can authenticate IoT devices at range. The advantage of Digitus is that we can authenticate low-power IoT devices using features intrinsic to their normal operation without the need for additional transmitters and/or other complex components such as DRAM. Our experiments demonstrate that we achieve $\geq 95\%$ accuracy on average, applicability in a wide range of IoT scenarios (range $\geq$ 5m, non-line-of-sight, etc.), as well as support for IoT applications such as finding hidden devices. Digitus represents a low-overhead solution for the authentication of low-end IoT devices.

Fingerprinting IoT Devices Using Latent Physical Side-Channels

Justin Feng, Tianyi Zhao, Shamik Sarkar, Dominic Konrad, Timothy Jacques, Danijela Cabric, Nader Sehatbakhsh
UCLA

Abstract

- Authenticating devices in an internet-of-things environment is essential for system security.
- Traditional fingerprinting approaches leverage artifacts found in wireless transmissions and/or DRAM.
- In this work, latent electromagnetic emanations from the processor's clock can be leveraged to authenticate.
- This approach applies to devices with or without radios and/or complex memory.
- Our experiments achieve ≥ 95% accuracy on a wide range of IoT scenarios.

Digitus

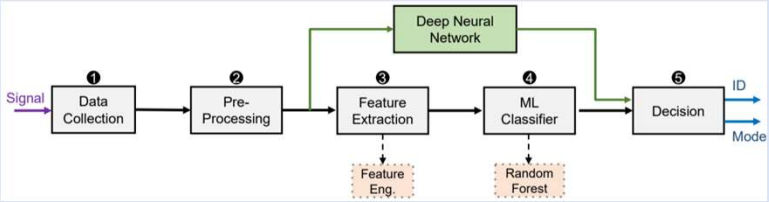

- Data collection and pre-processing are the first steps.
- We extract a varied set of features (time/frequency domain)
- Classification is completed via a feature-based classifier or a deep neural network.
- Finally, the decision is found (ID and mode).

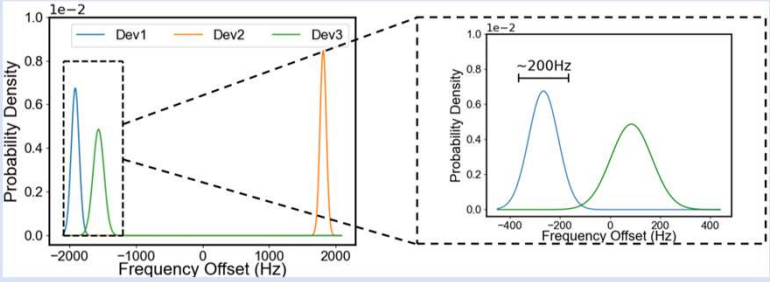

- This figure shows the distribution of center clock frequencies for three devices across our measurements.
- The measured clock drift across a month (~200Hz) is within the normal distribution, demonstrating feature stability.

Setup

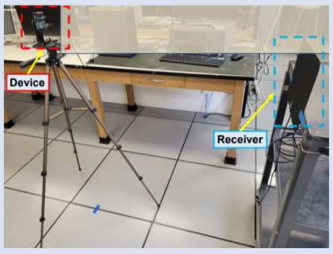

(a)

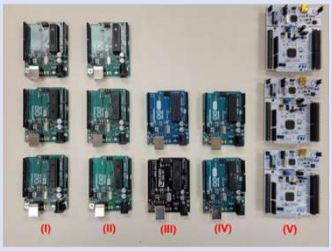

(b)

- (a) Here is a 1m line-of-sight setup.
- (b) The devices tested in the various experiments. Arduino Uno and STM32F411RE were tested.
- More complex environments and scenarios were tested.

Side-Channel Emanations

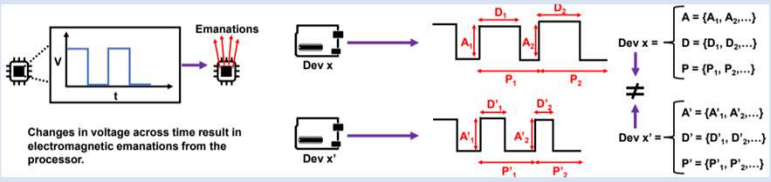

- The key insight is each device's processor has a unique voltage vs. time characteristic due to clock imperfections.
- These differences (in amplitude, duty cycle, period, etc.) result in unique fingerprints observable as electromagnetic energy.

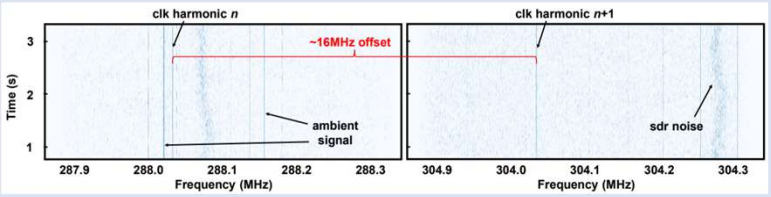

- Above is a spectrum collected using an SDR and antenna.
- A device's fingerprint is measured by analyzing the energy in the spectrum near the device's clock frequency harmonics.
- Ambient signal and noises are also measured that are filtered out during preprocessing.

Results

Experiment	Feature-Based Test Accuracy	Deep Learning Test Accuracy
Baseline Test	95.1%	97.5%
Temporal Test	91.5%	94.6%
Long-Term Test	96.3%	95.6%
Temperature	90.2%	91.2%
Attenuation (3m)	99.4%	93.9%
Attenuation (7.5m)	89.6%	80.1%
Non-Line-of-Sight	92.7%	94.8%
Intra-Device	96.1%	97.0%
Spoofing HackRF	96.3%	88.1%
Spoofing USRP	80.5%	87.2%
Spoofing (w/ training)	96.0%	97.2%

- A wide variety of environments were tested to demonstrate Digitus' robustness in real-world scenarios.
- We compare a feature-based vs. deep learning classifier.

Experiment	Description	FB Test Accuracy	DL Test Accuracy
Hidden Devices	Closed Set	99.0%	98.9%
	Open Set	66.0%	84.0%
Device State	Singular Device	72.1%	92.5%
	Multiple Devices, Target Device	80.0%	≥99.9%

- We examine Digitus' performance on detecting hidden devices and device state.
- Deep learning provides benefits in more complex scenarios.

Future Work

- Developing finer-grained characterization is a future work.
- Integration of side-channel-based analysis within existing systems.
- Combining with our prior work, SideComm.