Title: Securing IoT devices through Power Auditing and Privacy-preserving CNN

Authors: Woosub Jung[1], Yizhou Feng[2], Sabbir Khan[2], Chunsheng Xin[2], Danella Zhao[2], Gang Zhou[1]

Affiliation: [1] William & Mary, [2] Old Dominion University

Commonwealth Cyber Initiative

cyberinitiative.org

Abstract

Due to highly increasing IoT botnets, IoT and legacy devices have become more vulnerable to various botnet attacks. The research community proposed several studies that focused on IoT security for a single IoT device via power auditing. In addition, they often aim at specific devices or protocols. Instead, this work solicits a universal security solution to the practical deployments on diverse IoT devices. Furthermore, recent classifier designs have suffered from overfitting problems against unseen patterns in power traces.

We propose a real-time classification system in a distributed setting for securing IoT devices via power consumption. We designed a power-auditing-based IoT security system that supports up to a hundred IoT devices in real-time. The system includes 1) hardware design for ubiquitous IoT sensing, 2) cloud-based CNN classifier implement tation, and 3) privacy-preserving CNN protocol design.

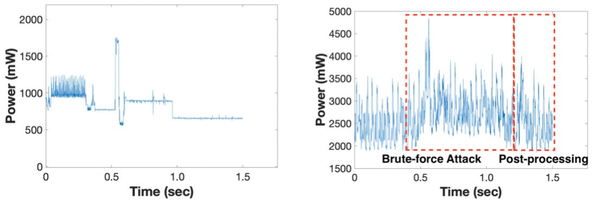

Figure 1. Power Auditing Traces for IoT Security

Introduction

On IoT devices, identifying and understanding various characteristics of abusive botnet attacks remains a challenge. Poor security on many IoT devices makes them soft targets, and often victims may not be even aware of whether they are infected. While network-based studies often targeted specific devices or protocols, IoT devices are not being controlled by just a few standardized operating systems or protocols. Therefore, new research is needed to find a universal security solution for diverse IoT devices in practical deployments. Furthermore, IoT devices are not capable of deploying sophisticated detection algorithms. It is also inefficient to deploy a dedicated system for a single IoT device. Thus, cloud-based mechanisms can be a solution to leverage their tremendous computing power to detect misused IoT devices.

In this work, we propose a distributed Power Auditing System that detects intrusion on IoT devices in real-time. We first utilize the power side-channel information for the detection of stealthy attacks. Figure 1 illustrates various displays of the power consumption data. Power consumption data is universal since it can capture accumulated tasks on heterogeneous devices, such as different hardware, vendors, operating systems, etc.

Next, we leverage the cloud environment for online classification to offload CNN inference computations. In our testbed, we use different types of IoT devices and measure power consumption data of them to identify IoT device behavior simultaneously. Overall, we demonstrate how power consumption data can tell malicious behavior on multiple IoT devices without data leakage.

Methodology

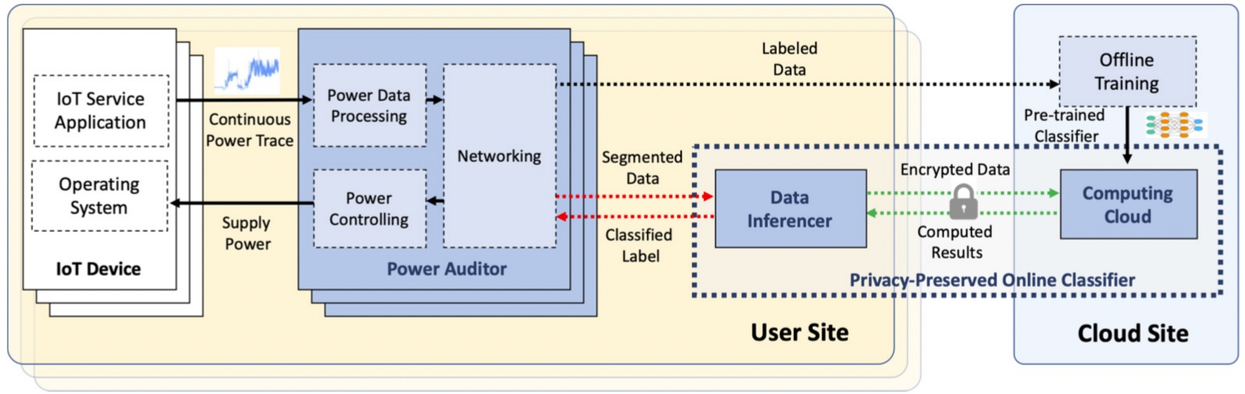

Figure 2. Online Distributed System Overview

Figure 2 illustrates the system overview. Our system includes a Power Auditor that can be interfaced with an IoT device, specifically between the IoT device itself and its power supply, as illustrated in Figure 3. The power consumption data of IoT devices are then transmitted to a cloud classifier for real-time intrusion detection.

Figure 4 provides an overview of the proposed 1-D CNN architecture. This CNN classifies time-series power-trace input as belonging to one of four behavior classes in IoT devices. Thus, we seek to distinguish malicious intrusions from other common behavior. This CNN model is partitioned and deployed in two non-colluding servers, the Data Inferencer and the Computing Cloud, to offload the computing requirements of the CNN inference computations to more capable means of performing them.

In order to protect data leakage, we then propose a privacy-preserved inference protocol via Packed Homomorphic Encryption. This protocol enables CNN communications without leaking sensing data. In addition, we also developed a sliding window protocol between the Power Auditor and the Data Inferencer. This is because overlapping input instances of the CNN classifier can help to achieve better detection accuracy.

Overall, we design an online system in a distributed setting for securing IoT devices through power auditing and privacy-preserving CNNs.

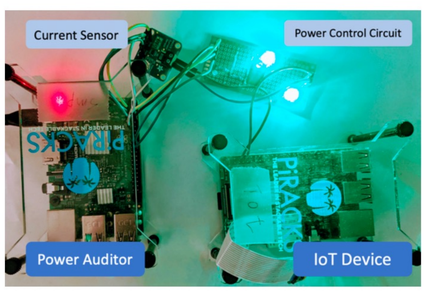

Figure 3. Power Auditing Device Prototype

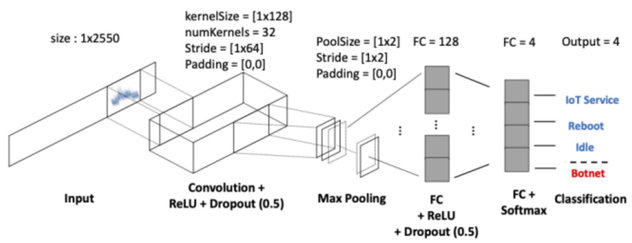

Figure 4. Online Distributed System Overview

Results

Our power auditing device design is lightweight enough for real-time inference. For example, the maximum CPU load in the Power Auditor is up to 20% of the Raspberry Pi's computing power. Memory (RAM) usage during online auditing is only 10MBytes out of Raspberry Pi's 1GB (1%). Based on our proposed sliding window design, the required network bandwidth is only 120Kbps. The power consumption of the Power Auditor is approximately 2W during the real-time inference.

Figure 5 shows the online classification results for each IoT device type. The results demonstrate the exceptional classification ability of the CNN classifier. For example, we achieved an overall accuracy of 98.95%. The Precision and Recall metric values for both tests are also above 98%.

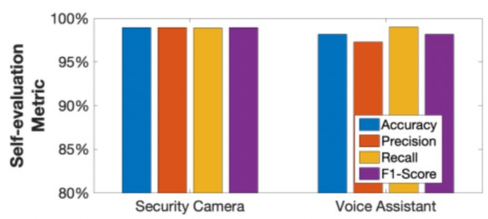

Figure 5. Online Classification Results

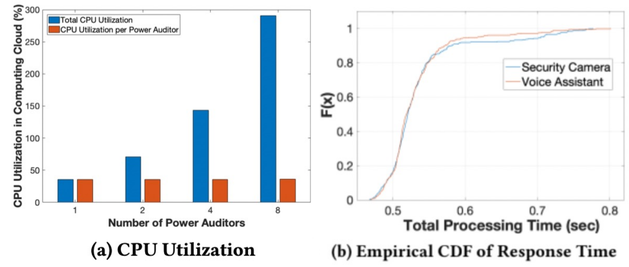

(a) CPU Utilization (b) Empirical CDF of Response Time

Figure 6. Scalable and Real-time System Results

Figure 6a deficits scalability performance results of four tests of 1, 2, 4, and 8 Power Auditors. As the number of Power Auditors increases, the CPU utilization of the Computing Cloud increases linearly. When we tested with eight Power Auditors, the total CPU utilization of the Computing Cloud server was less than 300% out of 3,200% (32 Cores).

Figure 6b demonstrates an empirical CDF function of the online classification response time. For both devices, we observe that over 80% of the inferences were done in 550ms or less. As we increased the number of Power Auditors, the average processing time does not change substantially.

Conclusion

In this research, we proposed a distributed online intrusion detection system for IoT devices via power auditing. We first developed a portable power-auditing device to measure power side-channel information of IoT devices in real-time. The one-dimensional CNN classifier was then designed and deployed in a distributed setting. The online CNN classifier predicted IoT devices' behavior with up to 98.9% accuracy, which outperforms the baseline classifier, especially in leave-one-out tests. In addition to the system components, we also designed distributed protocols to avoid data leakage and reduce networking redundancy. Finally, we evaluated the scalability of the system in a laboratory setting. Altogether, our system is the first online intrusion detection system that classifies multiple IoT devices' behavior via power traces.

In the future, we plan to enhance the performance of the inference protocol. Currently, the convolutional layer consumes the majority of the entire processing time. If we reduce that procedure, our system will be more reliable and scalable. In addition to the pre-trained classifier, we further plan to apply unsupervised learning so that users can use their dataset without labeling. This can expedite system deployment in a practical setting.

Acknowledgements

This research is partially supported by COVA CCI Cybersecurity Research and Innovation Funding, COVA CCI Cybersecurity Innovation Bridge Fund (Grant #HC-4Q21-005), COVA CCI Dissertation Fellowship, and NSF grant CNS-2120279.