Transition Recovery Attack on Embedded State Machines

Using power analysis strategies and data clustering to detect state transitions

Summer 2022

Team Members

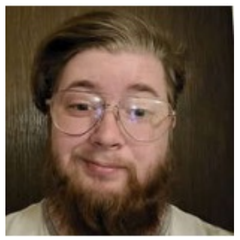

Clay Carper
3rd year PHD student in Computer Science. Studies side-channel attacks.

Andey Robins
2nd year Ph.D. student in Computer Science. Studies program synthesis and AI security

Problem Statement

Given a deterministic finite automata, written in C, can we differentiate between two states using collected power-trace data?

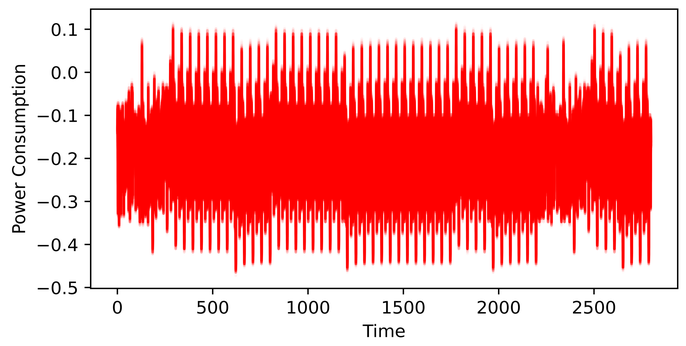

Overlapping power traces for 1001101_2

Results

Through application of the K-nearest neighbors algorithm, we were able to correctly classify the 256 possible classes with an accuracy of 81.02%. Additionally, we found that for this data the most suitable value for k was 171. As a general trend, accuracy increases as the number of neighbors (k) increases. It is worth noting that we were able to distinguish between the individual states S1 and S2 with 100% accuracy when k = 2.

Background

Recent developments in side-channel analysis have begun integrating various learning-based algorithms to exfiltrate useful information. An interesting, ongoing problem in the differential power analysis consists of detecting state transitions within embedded finite state machines. This body of work applies the K-nearest neighbors algorithm to a large power trace data set to identify state transitions and individual states.

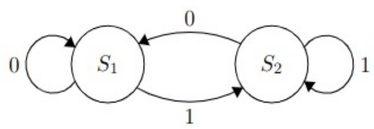

Finite state machine used for these experiments

Methods

Each of the states in our finite state machine were represented by separate, common numerical operations. The first is the infamous Quake fast inverse square root approximation, with the second being a combination of bit shifting and XORing. The transition between the two states in the machine was controlled by an 8-bit transition vector. For example, 10101010_2 would see the finite state machine alternate between state S1 and S2, creating a maximal number of state transitions. The 256 unique 8-bit transition vectors were each applied 200 times and power traces were also collected.

Challenges & Future Work

This method shows promising results for small finite state machines. However, future work looks to extend this method to larger, more complex structures. With ~40 GB of data from these experiments, future work will explore questions of learned transferability while reducing data collection requirements.

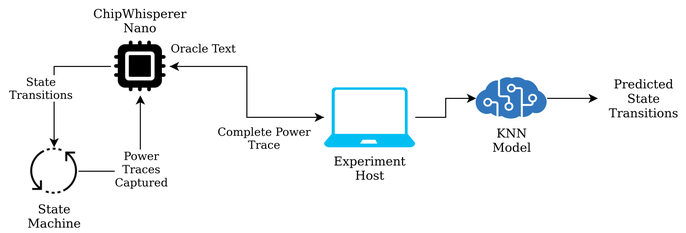

Project organization and workflow summary

Advisor: Dr. Mike Borowczak
Group Members:
- Clay Carper(ccarper2@uwyo.edu)
- Andey Robins (jtuttle5@uwyo.edu)

University of Wyoming

UW College of Engineering and Physical Sciences

UW Cybersecurity Education and Research Center

UW School of Computing

UW College of Engineering and Physical Sciences Electrical Engineering and Computer Science

CEDAR Cybersecurity Education And Research Center