JackHammer: Efficient Rowhammer on Heterogeneous FPGA-CPU Platforms

Zane Weissman, Thore Tiemann, Daniel Moghimi
Advisors: Berk Sunar, Thomas Eisenbarth

WPI

UNIVERSITÄT ZU LÜBECK

tom's HARDWARE

US Edition RSS Search

Reviews Best Picks Raspberry Pi CPUs Graphics Cards Laptops More Forums

Home > News

Newly Demoed JackHammer Cyberattack Uses FPGA-CPU Combo to Attack Memory

By Nathaniel Mott January 02, 2020

Jack be nimble, Jack be swift, Jack will smash your RAM to bits.

ZDNet IOT CLOUD AI SECURITY MORE NEWSLETTERS ALL WRITERS

FPGA cards can be abused for faster and more reliable Rowhammer attacks

Researchers expand Rowhammer attacks to FPGA-CPU hybrid platforms.

Abstract

We studied two new **heterogeneous FPGA-CPU platforms** from Intel: the **integrated Arria 10 GX** which shares a chip with its host CPU, and the **Arria 10 GX PAC expansion card** which connects the FPGA to the CPU via a PCIe interface.

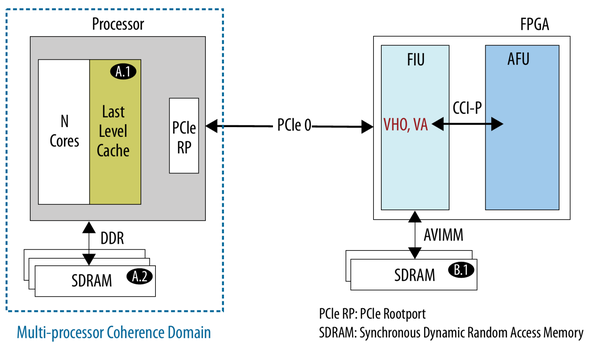

We demonstrate **JackHammer**, a Rowhammer attack from the FPGA to the host's main memory, performing **twice as fast** as a conventional CPU Rowhammer and causing **four times as many faults**, as well as a **cache covert channel across FPGA and CPU.**

JackHammer

In the Rowhammer exploit, the electromagnetic effect of **repeated accesses to certain memory addresses causes stored bits in physically adjacent locations to flip their values**. JackHammer is our hardware Rowhammer implementation for Arria 10 GX FPGAs. It uses the PCIe interface between the FPGA and its host CPU to the Arria 10 GX's **simpler memory access architecture** compared to complicated modern CPUs, and because **memory reads from the FPGA bypass the CPU cache**, eliminating the need for time consuming cache flushes between RAM accesses.

JackHammer vs. CPU Rowhammer

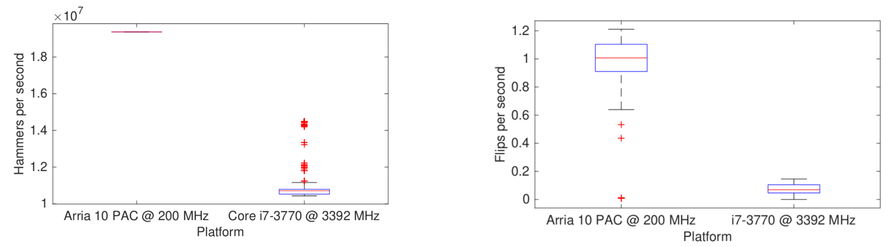

Fault Attack on WolfSSL RSA

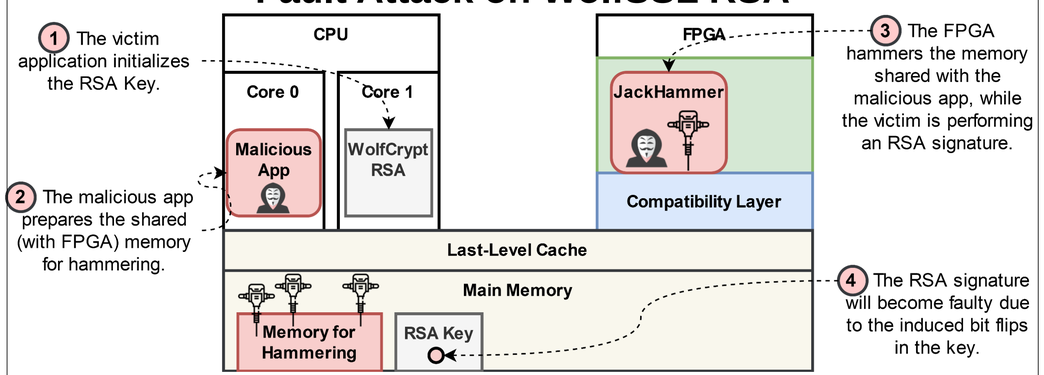

We constructed a **fault injection attack against the RSA signing function in WolfSSL**, outlined in the figure above. When using JackHammer instead of a conventional CPU Rowhammer, **a key can be recovered an average of 17% faster**. With some typical defenses against Rowhammer exploits in place, JackHammer is **over three times more likely to cause a fault** than the same attack with CPU Rowhammer.

Caching Behavior and Covert Channel

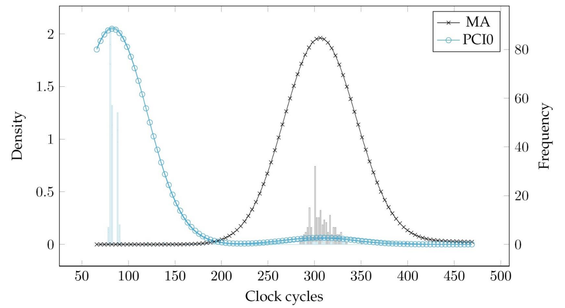

The blue line in this figure is the distribution of CPU memory access latency when accessing cache lines recently written by the FPGA; a typical distribution is shown in black. Our measurements show that **cache lines written by the FPGA are placed in the last-level cache of its host CPU.**

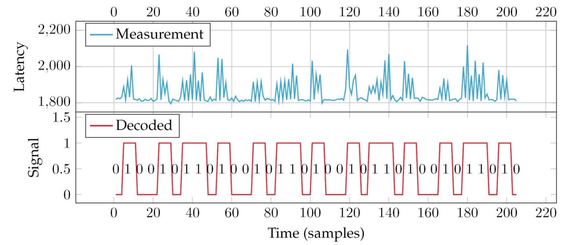

We constructed a covert channel with the FPGA as the sender and a cooperative CPU program as receiver. The **FPGA sends binary messages by writing to a cache line** when transmitting a 1 and staying quiet otherwise. The **receiver continuously probes a the cache set** to detect access latency fluctuations to receive the messages. An example decoding of the latency measurements is shown in the figure above. While using heavily redundant encoding, we still achieve a **throughput of 94.98 kBit/s.**

References

- Weissman et al., JackHammer. Available on arxiv.org
- Acceleration Stack for Intel® Xeon® CPU with FPGAs Core Cache Interface (CCI-P) Reference Manual
- Special thanks to Intel's Evan Custodio, Alpa Trivedi, and Sayak Ray for their guidance and support