

First Time Miss : Low Overhead Defense Against Shared Memory Cache Side Channels

Kartik Ramkrishnan

Stephen McCamant

Antonia Zhai

Pen Chung Yew

Department of Computer Science and Engineering
University of Minnesota, Twin Cities




Talk Outline

1. The Security Issue
2. Existing Mitigation Approaches
3. Our Approach
4. First Time Miss Implementation
5. Simulation Results
6. Conclusion

The Security Issue :

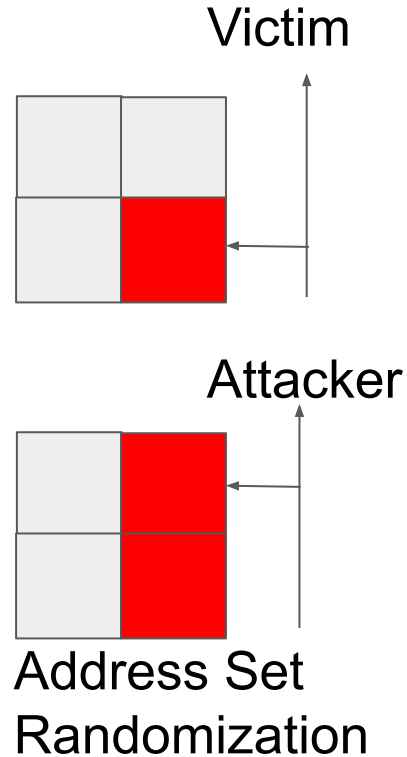
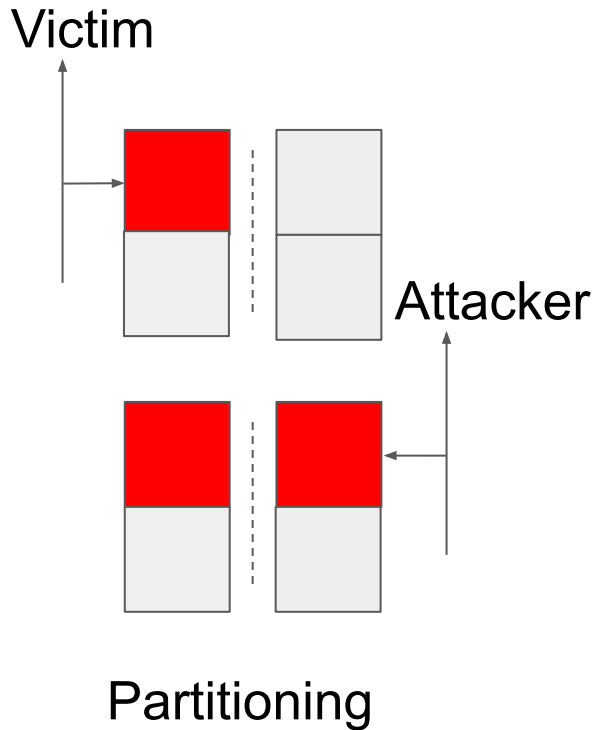
Information Leakage due to Extra Cache Hits

1. Attacker can detect extra cache hits due to victim's use of shared memory, for example, shared libraries.
2. Hit/Miss timing difference when accessing shared code leaks information.




Existing Approaches

1. Cache Partitioning may be expensive.
2. Address-Set Randomization also uses cache resources less efficiently.

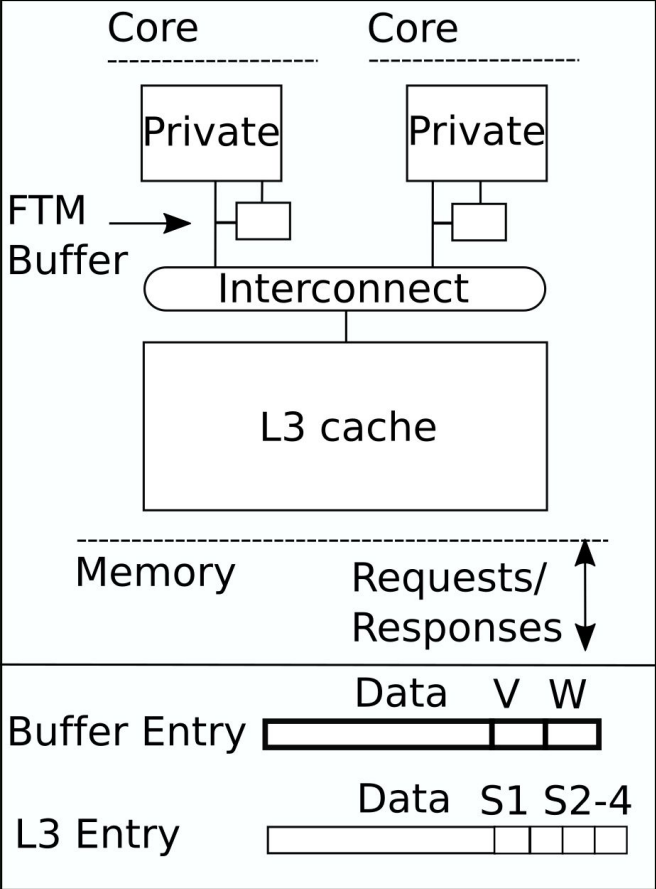


Our Approach - First Time Miss

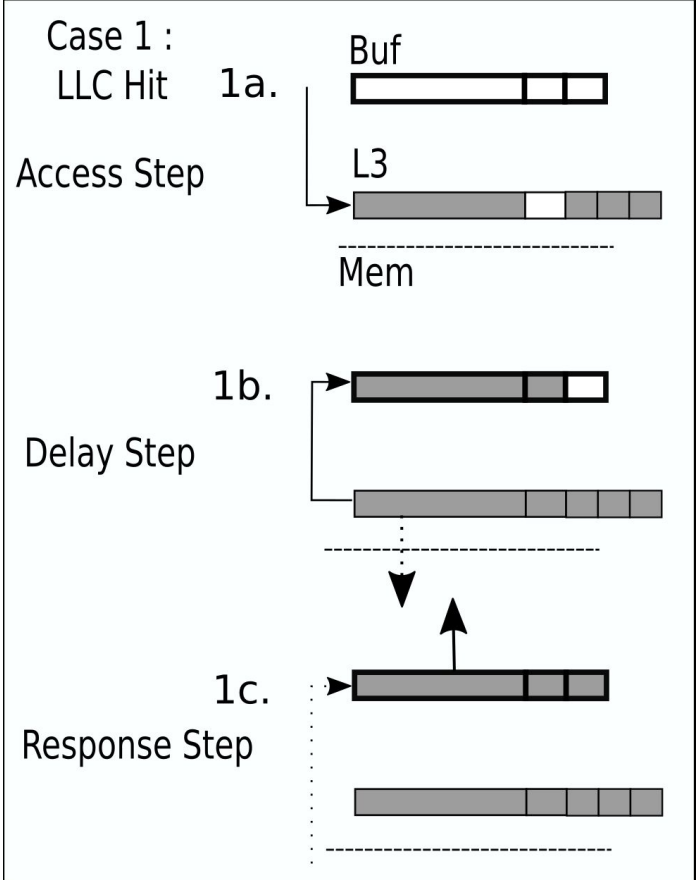
1. Delay malicious cache hits, making them appear to be cache misses.
2. This confuses the attacker's measurements and eliminates information leakage due to extra cache hits.



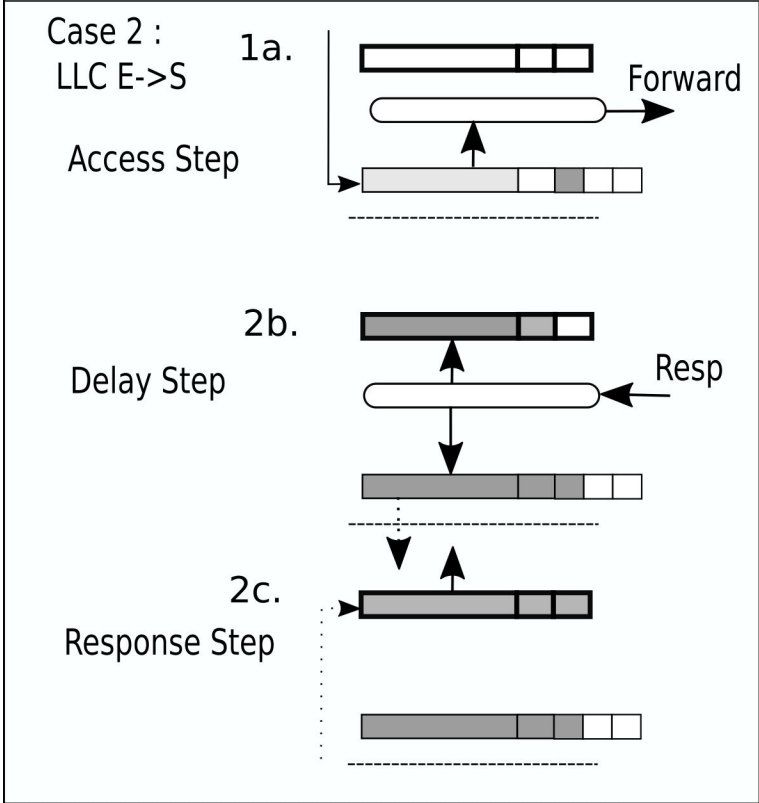
Implementation: Hardware Changes




Implementation : Operational Changes 1



Implementation : Operational Changes 2




Simulation Results - SPEC 2017 Stress Test (4 core)



Negligible performance overhead for FTM implementation due to low number of FTM

Simulation Result - PARSEC 3.0 (4 core)



↑
Negligible performance
overhead for FTM
implementation due to low
number of FTM

Conclusion

1. First Time Miss approach to mitigate shared memory cache side channels.
2. Low overhead implementation for the last level cache.
3. Evaluation using SPEC 2017 and PARSEC benchmarks indicates low performance overhead.

Thank You !