

Protecting Cache Memories from Eviction-Based Attacks

An algorithm to protect against eviction-based cache side channel attacks, while incurring negligible storage and performance loss, and without relying on any operating system support.

A Georgia Tech inventor successfully developed an algorithm that can protect caches against eviction-based attacks at an ultra-low cost, in terms of both storage and performance. The [MQ1] algorithm is based on applying an encrypted-address on cache memories for randomization and changing this randomized mapping periodically. The encryption would cause the lines that are mapped to the cache memories to get scattered to different sets of locations in an unpredictable fashion, which is only accessible via an encryption key. Furthermore, to increase security, the algorithm will dynamically change the encryption key, where the contents of the cache are gradually remapped from the old key to the new key. All-in-all, this algorithm provides cache randomizations without requiring any storage tables to memorize the mapping, nor will it slow down the system or require much storage.

Summary Bullets

- **Robust to attacks** – tolerates years of continuous eviction-based attacks
- **Negligible slowdown** – creates an miniscule slowdown of 1%
- **Storage efficient** – requires a storage overhead of less than 100 bytes for new structures

Solution Advantages

- **Robust to attacks** – tolerates years of continuous eviction-based attacks
- **Negligible slowdown** – creates an miniscule slowdown of 1%
- **Storage efficient** – requires a storage overhead of less than 100 bytes for new structures
- **Self-supporting** – does not need any OS support
- **Practical design** – easy to implement, simple design, that is amenable to commercial adoption

Potential Commercial Applications

- All processor manufactures with designs containing an on-chip cache

Background and More Information

Caches are structures within the memory chips; their purpose is to store recently-accessed data, so that when a user requests that data (for instance, a recent file being reopened), it is accessible much more quickly than if the data had to be retrieved from the main memory. Unfortunately, the timing difference between the cache-hit (when requested data is found successfully) and a cache-miss (when requested data cannot be found) can be used by an adversary to obtain unauthorized information from the system. One example of an attack is an “eviction-based attack”, which occurs when the adversary and the victim share some storage structures, such as the on-chip cache. While cache attacks have been demonstrated in the past at a smaller scale, the recent vulnerabilities show that cache attacks can affect hundreds of millions of processor systems, and highlight the need to develop efficient solutions to mitigate (or entirely prevent) such attacks

Inventors

- Dr. Moinuddin Qureshi
Professor – Georgia Tech College of Computing

IP Status

:

Publications

, -

Images

Visit the Technology here:

[Protecting Cache Memories from Eviction-Based Attacks](#)

<https://s3.sandbox.research.gatech.edu//index.php/print/pdf/node/3434>