Zombie-Based Mitigation – Protecting Cache from Flush-Based Attack

An algorithm that tolerates flush-based cache attacks without requiring rewrite of applications or profile information, and retaining page sharing.

A Georgia Tech inventor developed an algorithm titled 'Zombie-Based Mitigation' (ZBM), which successfully guards against flush-based attacks. This is done by treating hits on invalidated-lines, also known as 'Zombie lines,' as misses to avoid any pattern leaks to the attacker. This eliminates any timing channel opportunity, which reveals the system's patterns, for the attacker as both the cache hit and misses would incur the same time period once the line is marked as a zombie line. A cache hit and a cache miss represent if the requested line can or cannot be found in the cache. Overall, the solution is based on marking the line as a "zombie" on flushes and protecting them until the lines are naturally evicted. Furthermore, ZBM requires only 1-bit per cache line, retains OS-based page sharing, does not require application rewrite, and does not incur slowdown.

Summary Bullets

- **Negligible performance** – performance is unchanged when the system is not under attack
- **Storage efficient** – low storage required for implementation and execution (1 bit per line)
- **Robust** – no restrictions on capacity benefits of page-sharing

Solution Advantages

- **Negligible performance** – performance is unchanged when the system is not under attack
- **Storage efficient** – low storage required for implementation and execution (1 bit per line)
- **Robust** – no restrictions on capacity benefits of page-sharing
- **Self-supporting** – no requirement of rewriting the software
- **Simple** – algorithm has a simple implementation and design

Potential Commercial Applications

Any device that has a processor and a cache. The steps for the solution are:

- Marking invalidated-lines
- Protecting invalidated-lines
- Determining victim selection on invalidated-lines misses

- Mitigating of invalidated-line hits

Background and More Information

OS-based page sharing is a commonly used optimization in modern systems to reduce memory redundancies. For instance, it is used to avoid redundant copies of pages across applications or for having multiple copies of the same data pages. Such sharing allows different programs accessing the same code to get routed to the same page. Unfortunately, recent vulnerabilities exploiting cache side-channels have universally impacted the entire computing industry, underscoring the importance of mitigations for next-generation hardware. In fact, such sharing can make the system vulnerable to Flush+Reload cache attacks, where an attacker inserts a cache line of shared data, and reloads it to reveal information stored in the system. Current proposals to mitigate Flush+Reload attacks are impractical as they rely on disabling page sharing, or using performance counters to detect deviation in the behavior of critical applications.

Inventors

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

IP Status

Patent has issued<o:p></o:p> : US11783032B2

Publications

-, -

Images

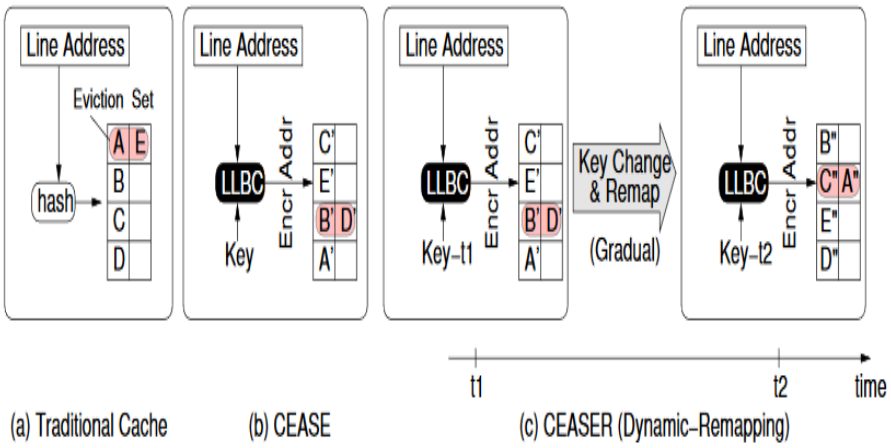

Visit the Technology here:
Zombie-Based Mitigation – Protecting Cache from Flush-Based Attack

---

https://s3.sandbox.research.gatech.edu//print/pdf/node/3435