Memory Defenses

The Elevation from Obscurity to Headlines

Rajeev Balasubramonian
School of Computing, University of Utah

CE Jr Seminar
Nov 19th 2019

The New York Times

Researchers Discover Two Major Flaws in the World's Computers

The two problems, called Meltdown and Spectre, could allow hackers to steal the entire memory contents of computers, including mobile devices, personal computers and servers running in so-called cloud computer networks.

Image sources: pinterest, gizmodo

Spectre Overview

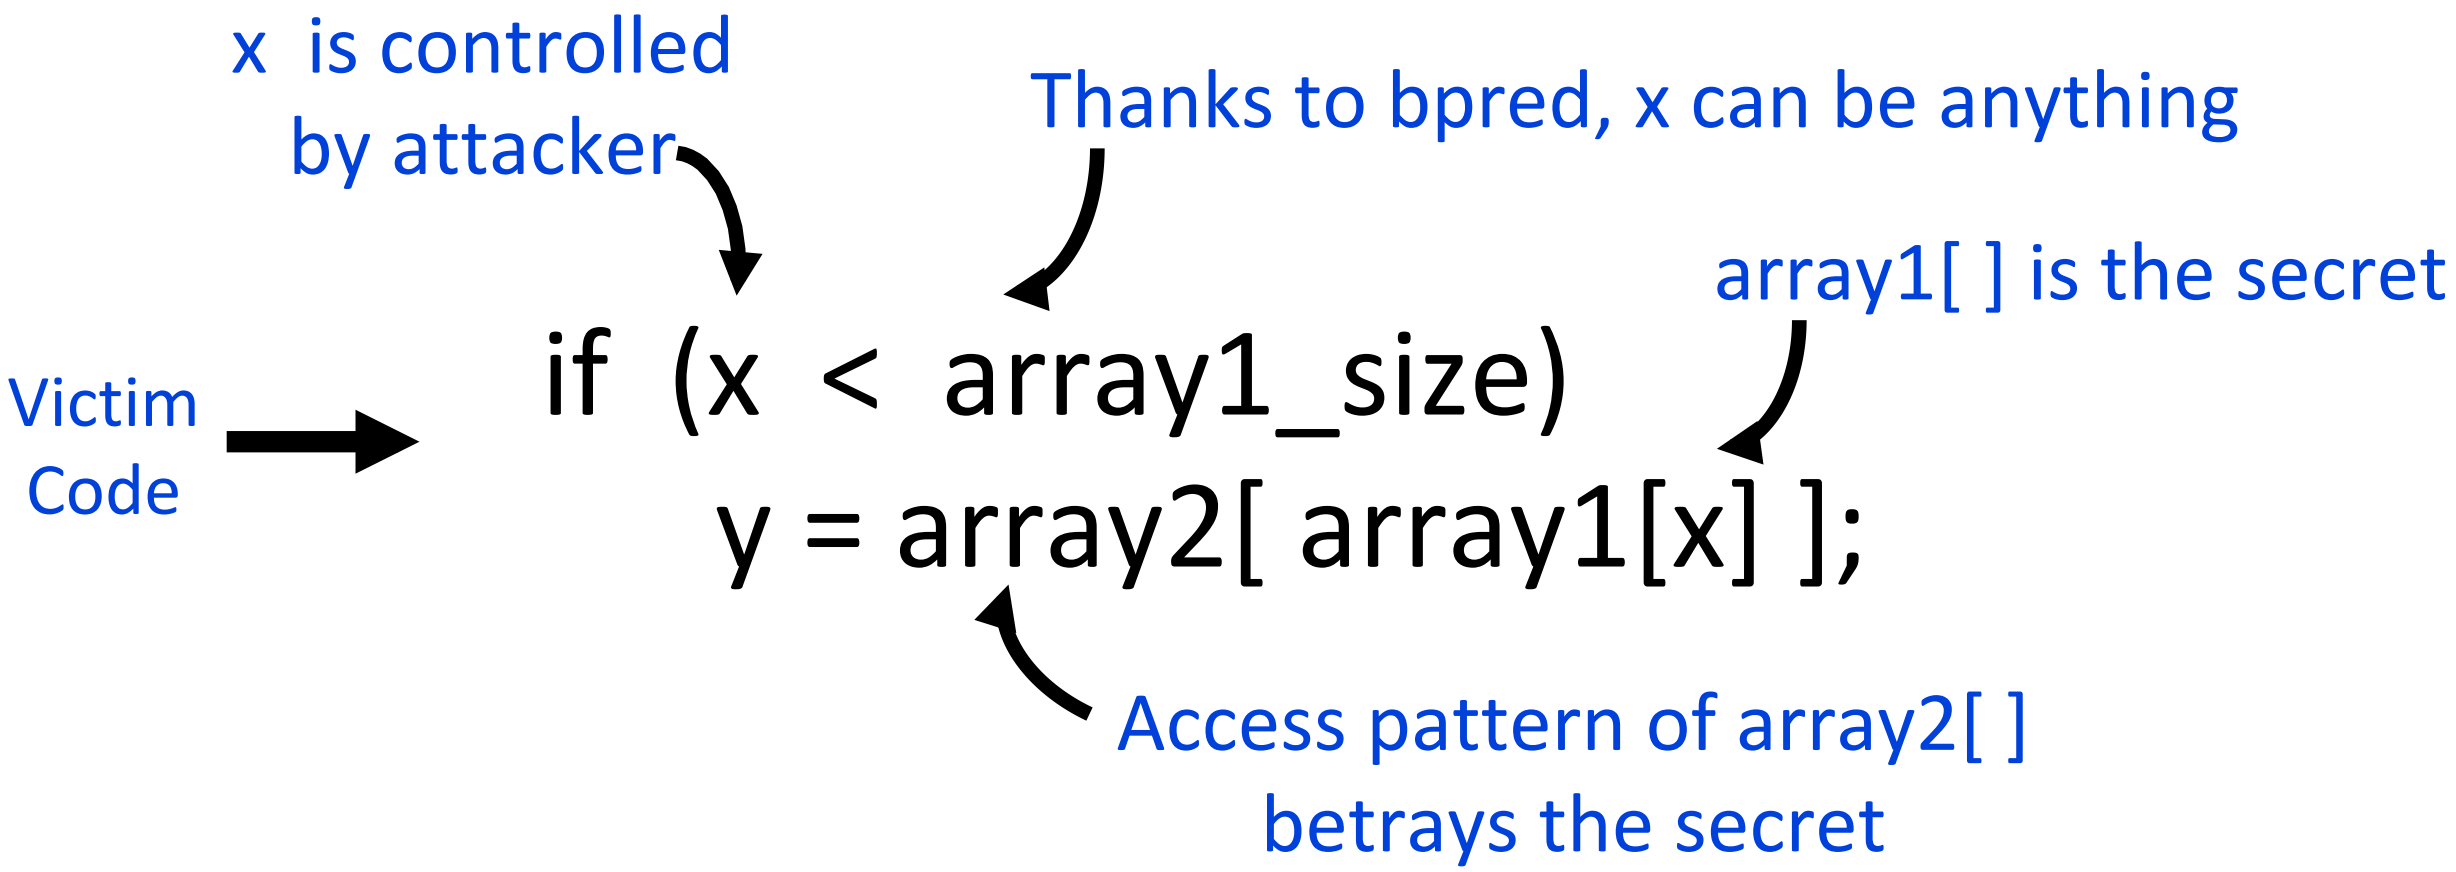

Spectre

What Did We Learn?

Speculation

+

Specific Code

+

No side channel defenses

The Wake Up Call

Say Yes to Side Channel Defenses

Overview

- Memory timing channels
  - The Fixed Service memory controller [MICRO 2015]
- Memory access patterns
  - Near-data ORAM [HPCA 2018]
  - Hierarchical ORAM [ASPLOS 2019]
- Memory integrity
  - Improving SGX with VAULT [ASPLOS 2018]

Memory Timing Channels

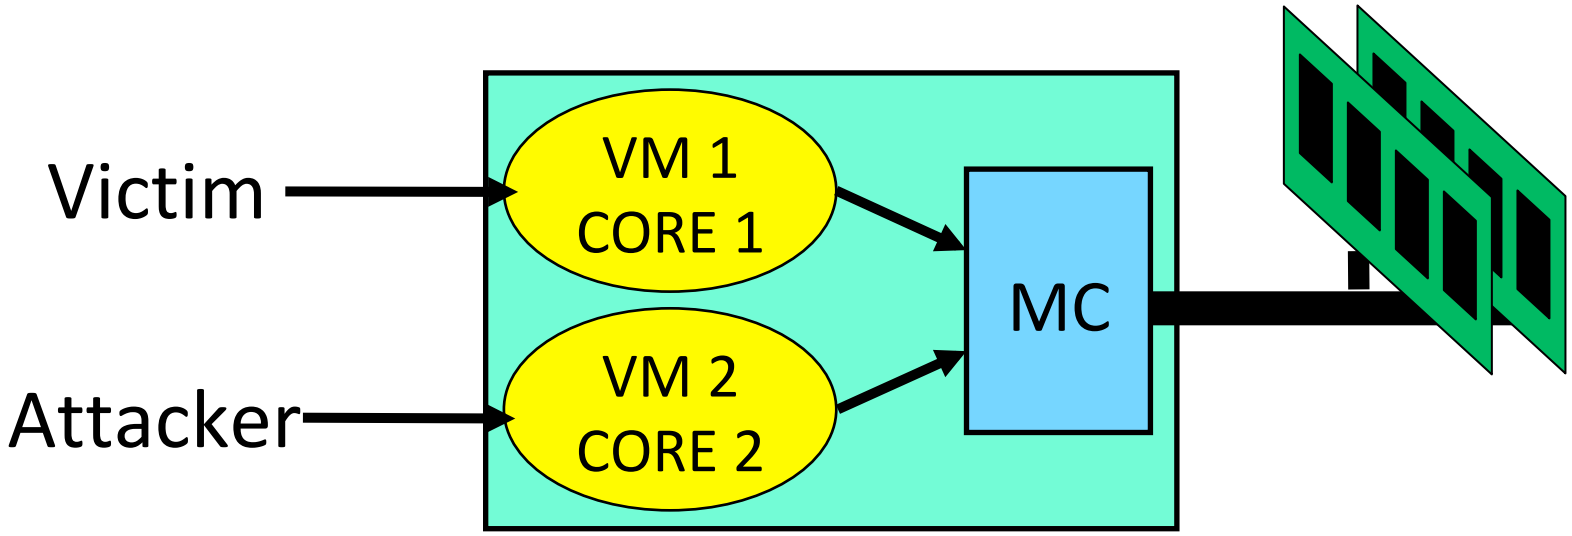

Two VMs sharing a processor and memory channel

Possible Attacks

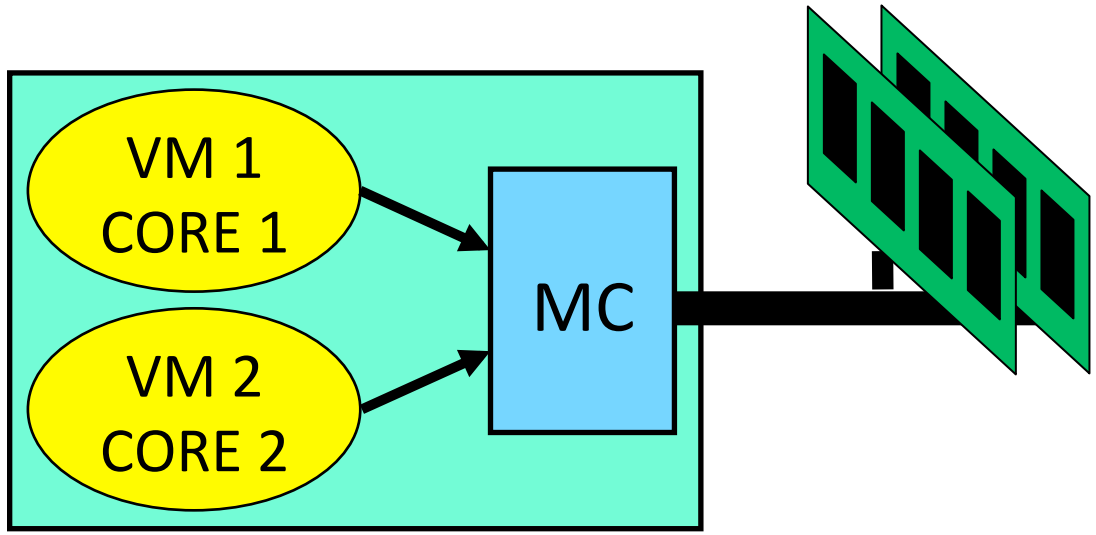

Attack 1: Bits in a key influence memory accesses

Attack 2: A victim can betray secrets through memory activity

Attack 3: A covert channel attack

Covert Channel Attack

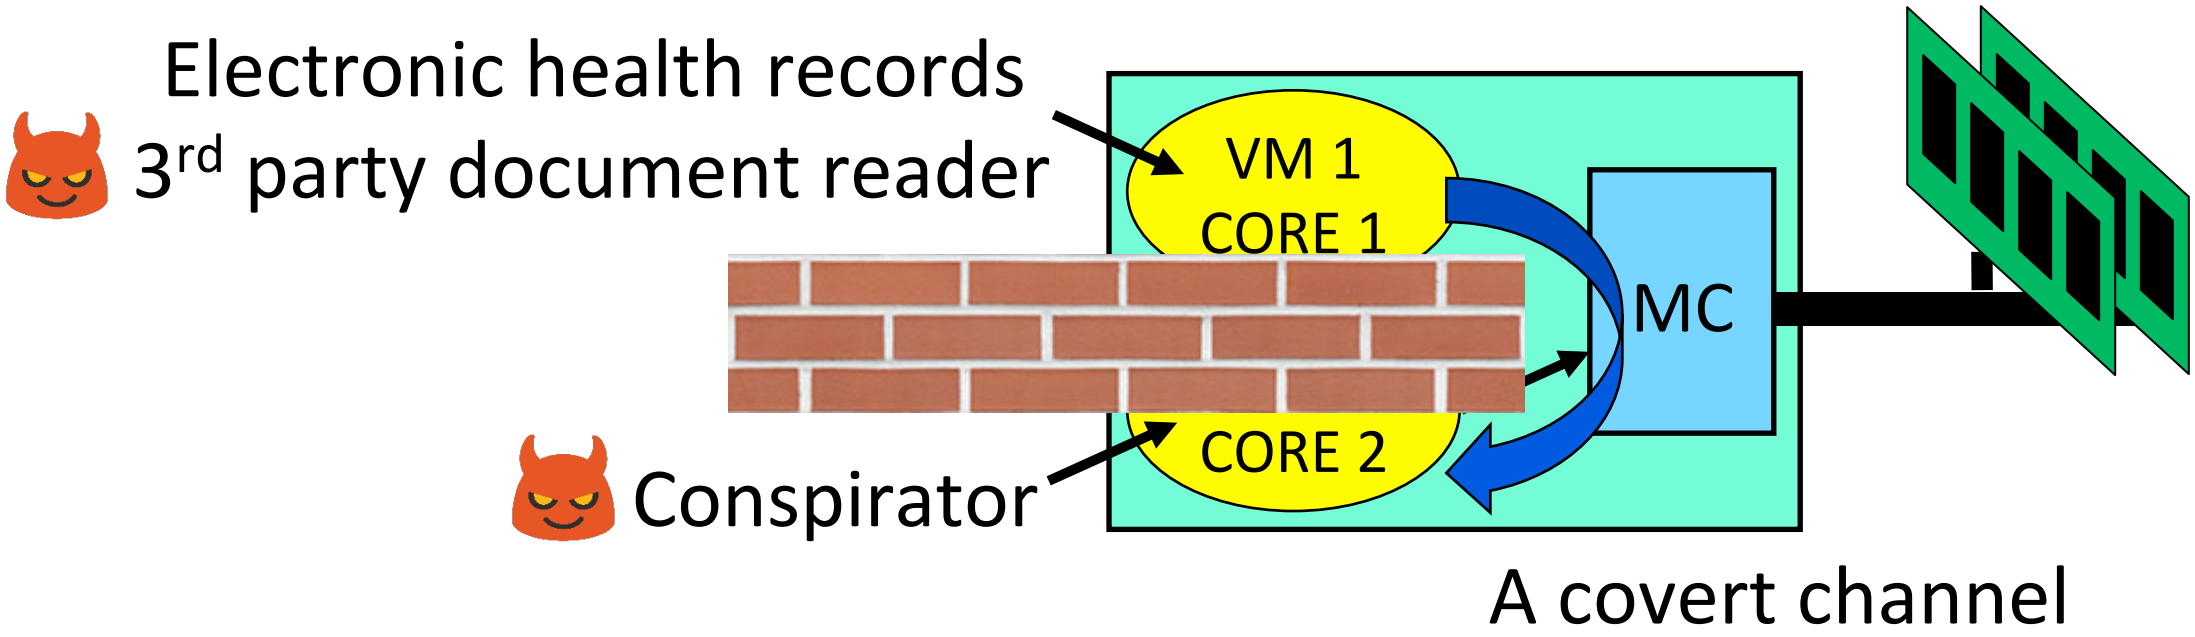

Fixed Service Memory Controller

VM-1 has its data in Rank-1
VM-2 has its data in Rank-2
…
VM-8 has its data in Rank-8

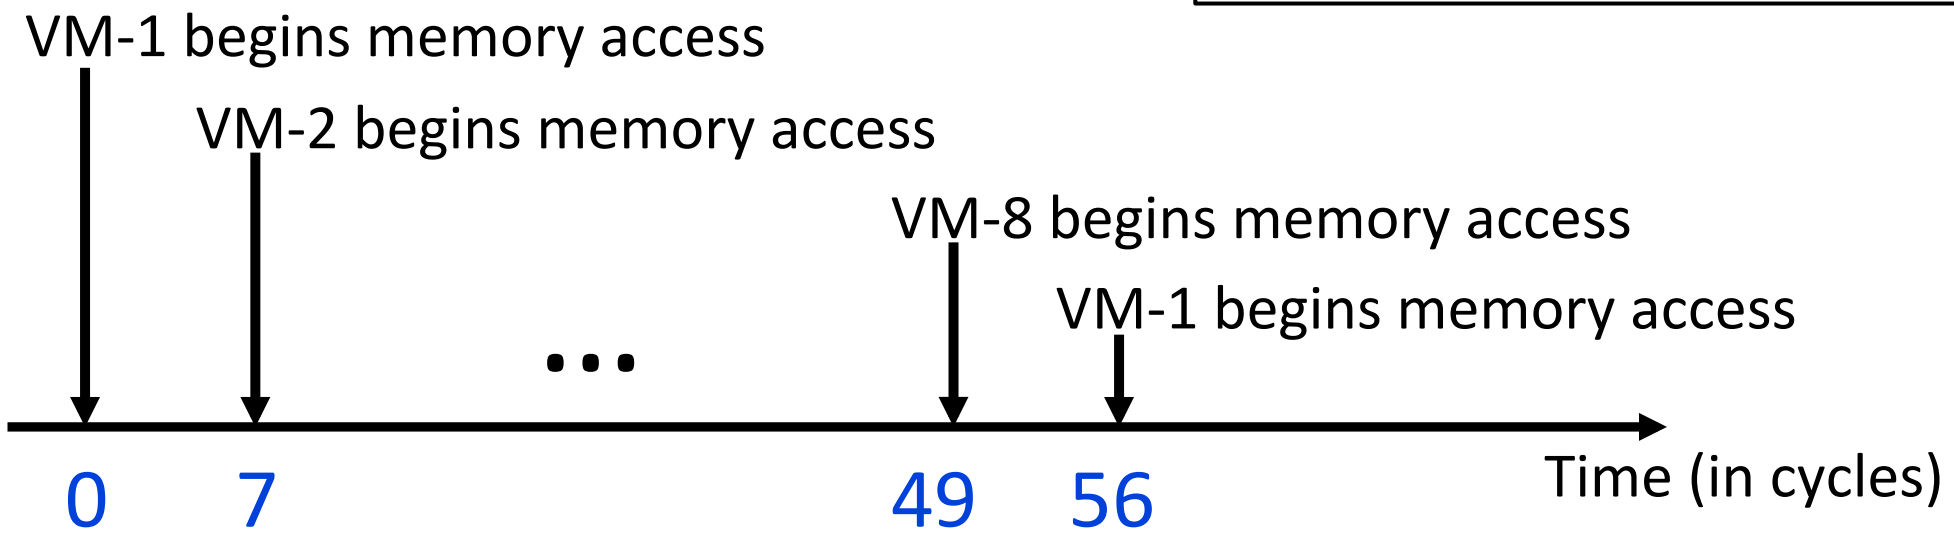

Fixed Service Details

- Deterministic schedule
- No resource contention
- Dummy accesses if nothing pending
- Lower bandwidth, higher latency
- Why 7? DRAM timing parameters, worst-case
- Rank partitioning: 7 cycle gap
- Bank partitioning: 15 cycle gap
- No partitioning: 43 cycle gap

Overcoming Worst-Case

- In one batch of requests, schedule all reads, followed by all writes (worst-case encountered once per batch)
- Impose constraints on banks that can be accessed – triple bank alternation

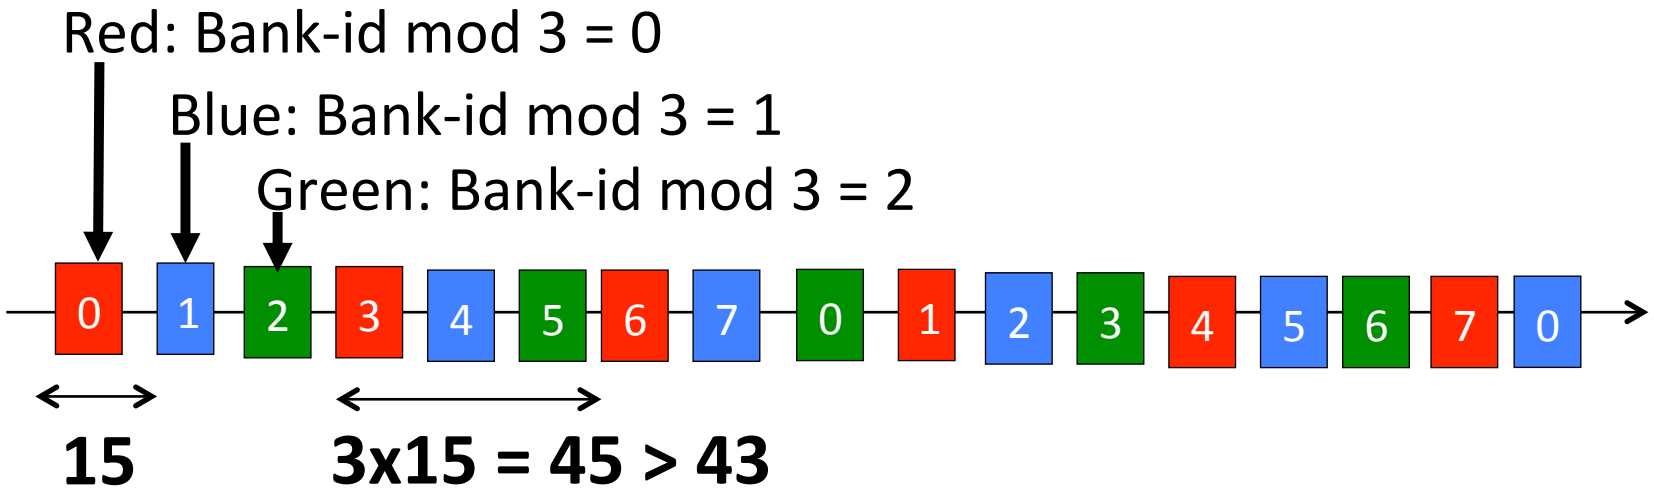

Results

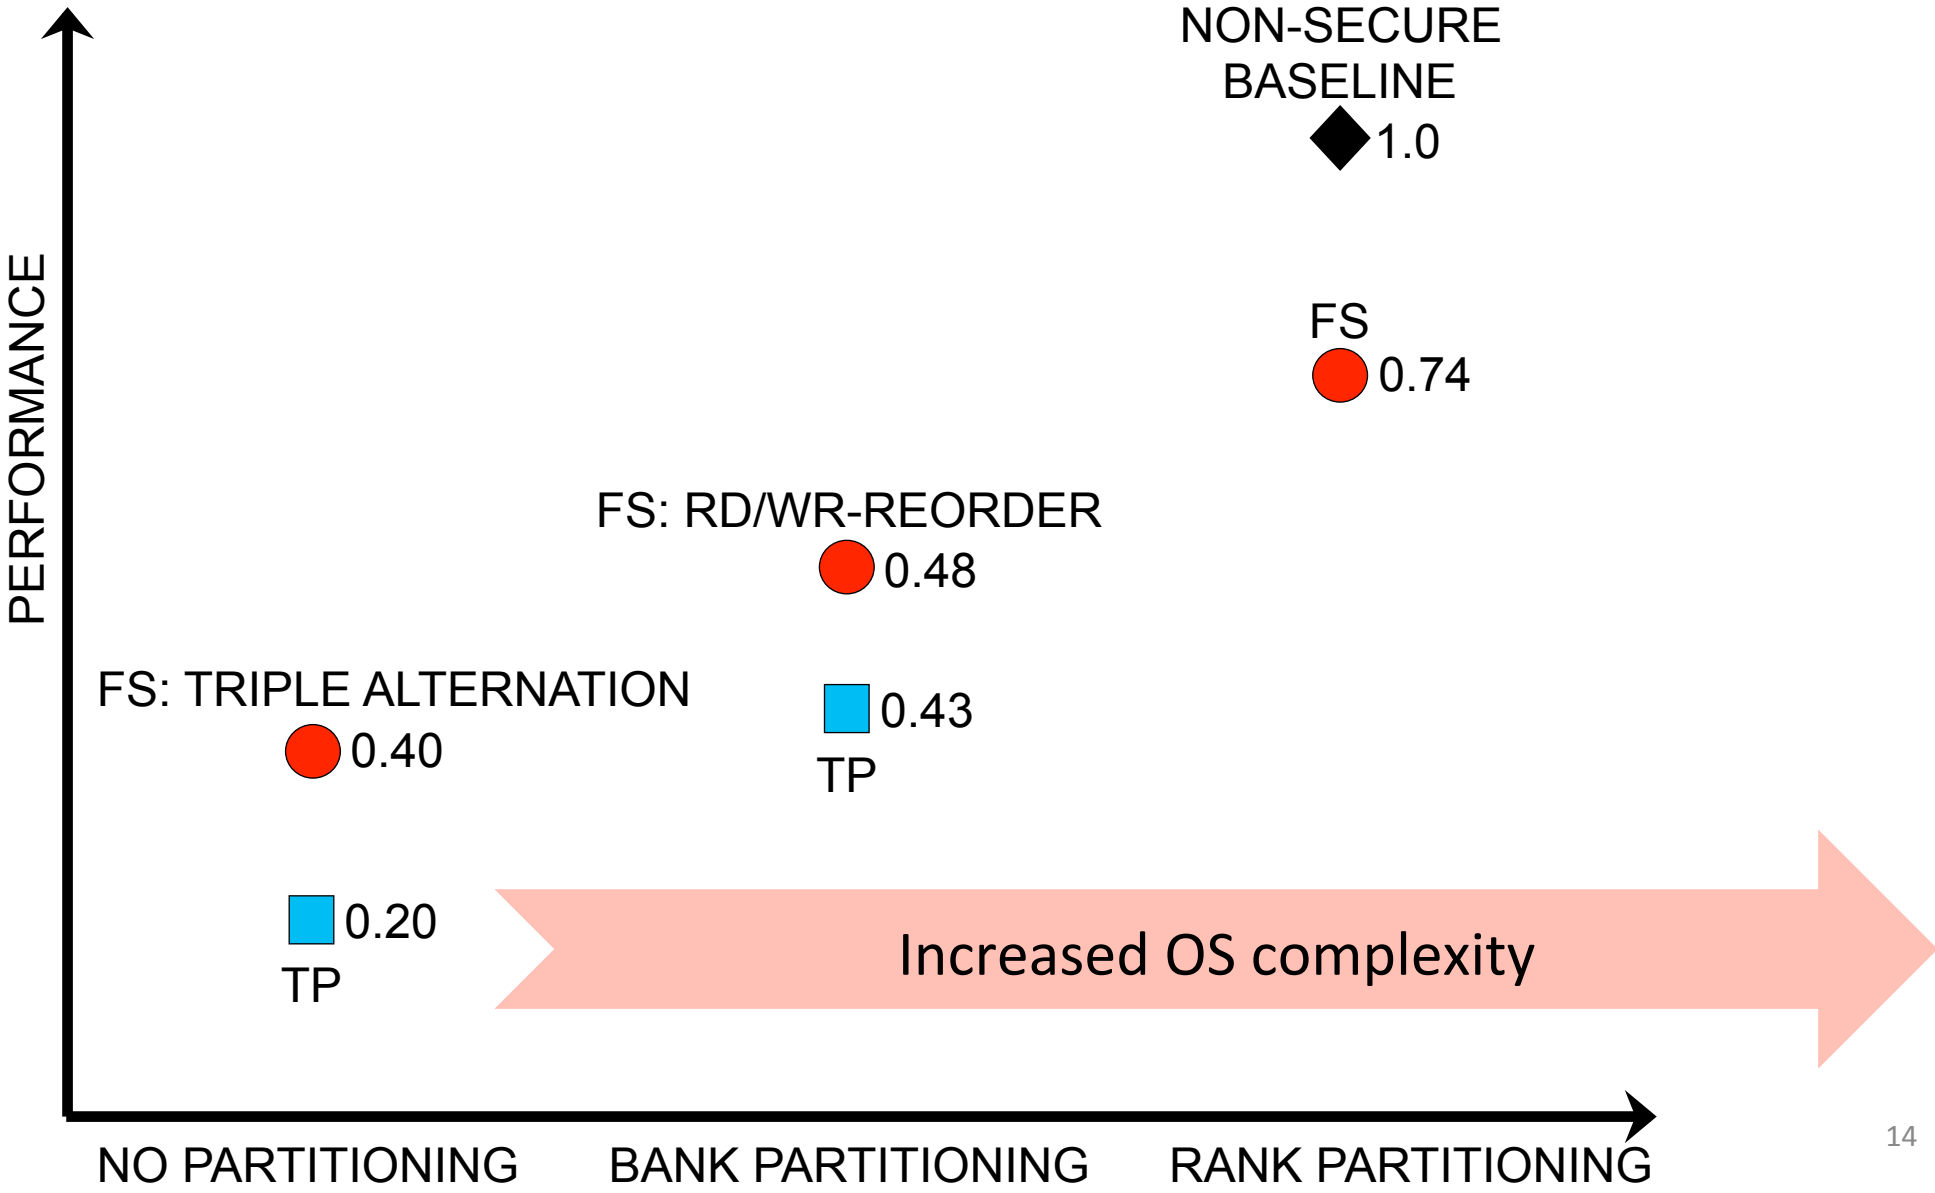
NON-SECURE
BASELINE
1.0
FS
0.74
FS: RD/WR-REORDER
0.48
FS: TRIPLE ALTERNATION
0.40
0.43
TP
0.20
TP
PERFORMANCE
Increased OS complexity
NO PARTITIONING
BANK PARTITIONING
RANK PARTITIONING

Overview

- Memory timing channels
  - The Fixed Service memory controller [MICRO 2015]
- Memory access patterns
  - Near-data ORAM [HPCA 2018]
  - Hierarchical ORAM [ASPLOS 2019]
- Memory integrity
  - Improving SGX with VAULT [ASPLOS 2018]

Oblivious RAM

- Assumes that addresses are exposed

Image sources: vice.com

- PHANTOM [CCS'13]: Memory bandwidth overhead of …

Oblivious RAM

- Assumes that addresses are exposed

Image sources: vice.com

- PHANTOM [CCS'13]: Memory bandwidth overhead of …

2560x (about 280x today)

Path-ORAM

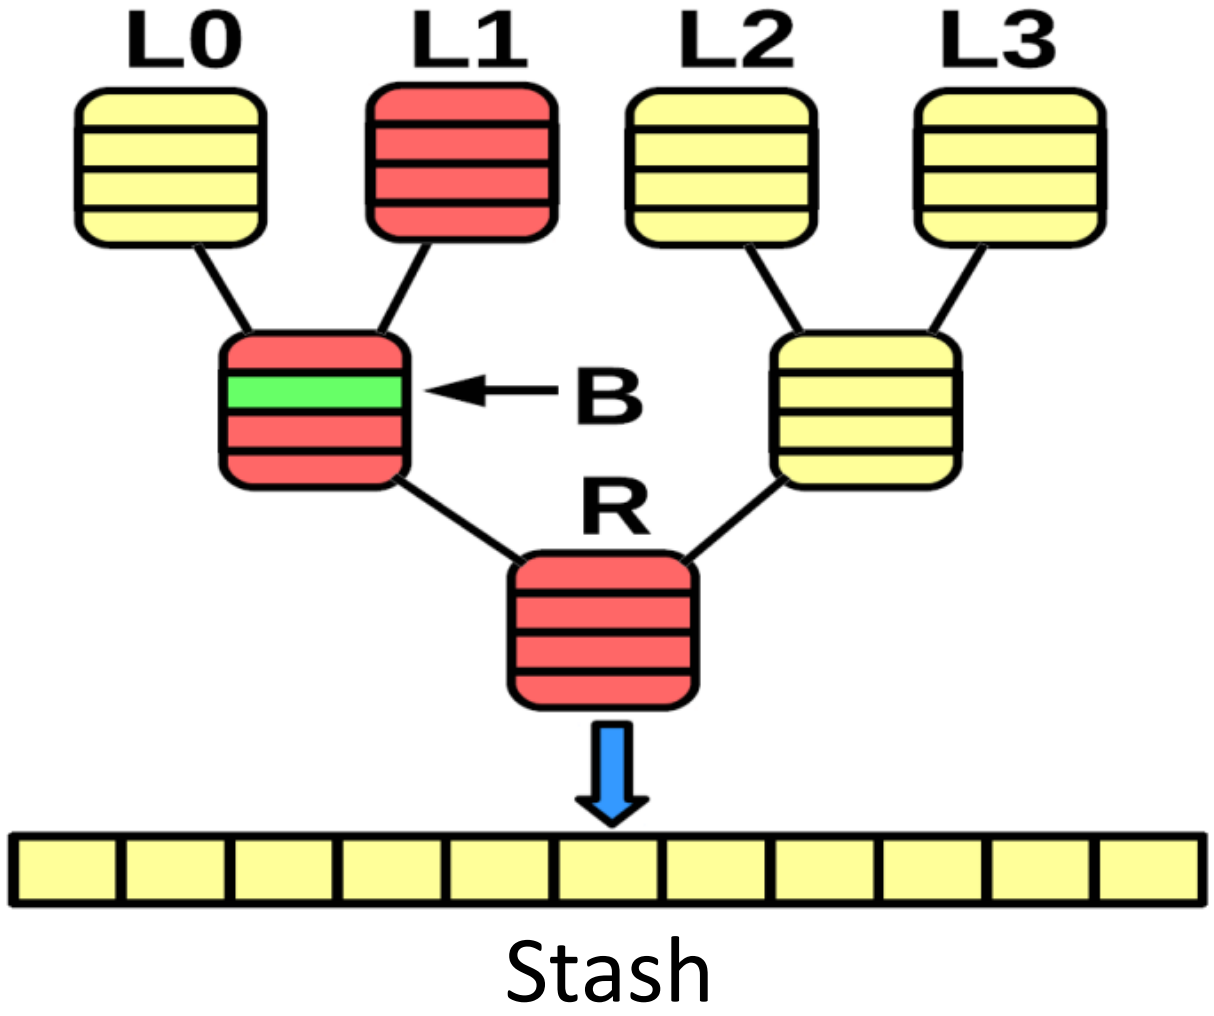

A Distributed ORAM

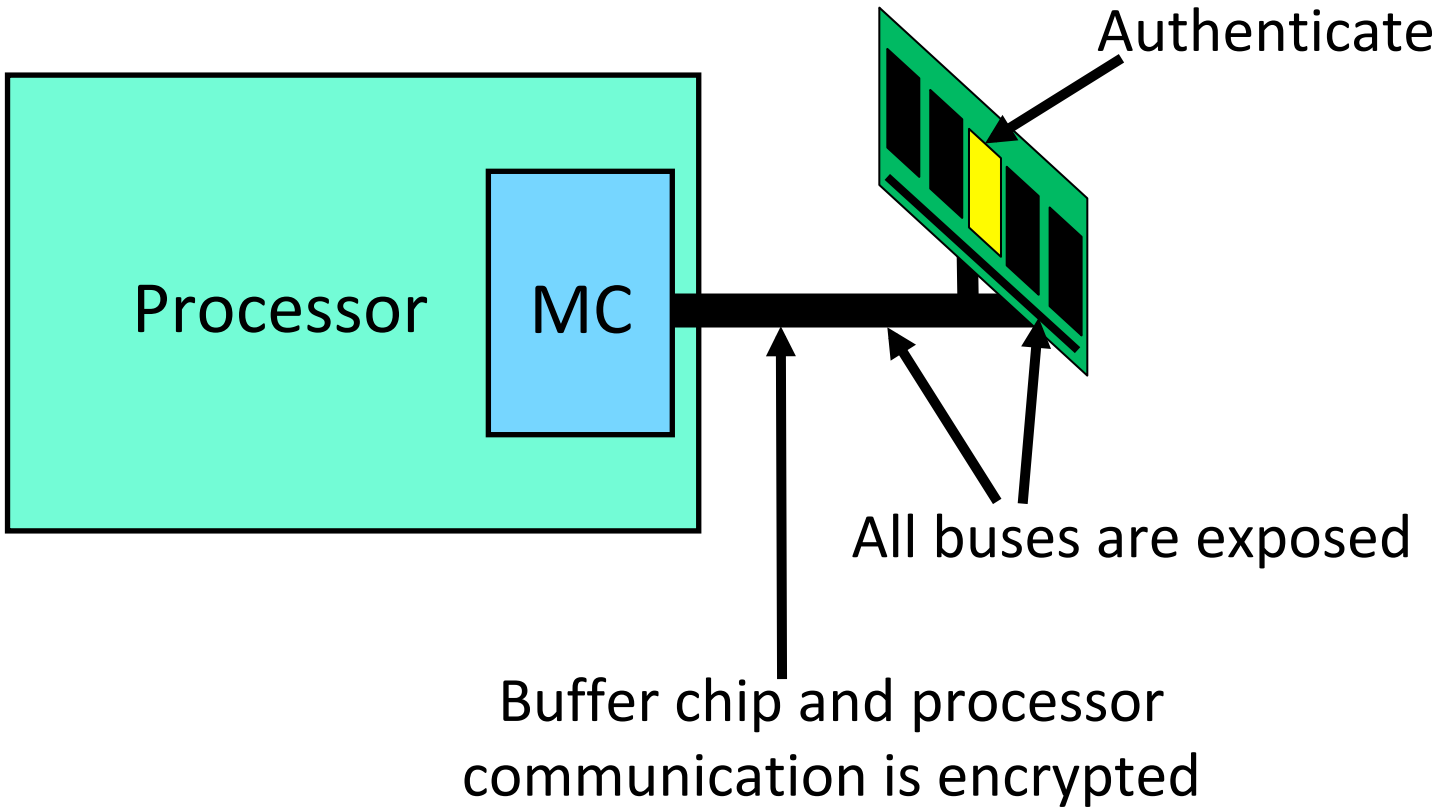

ORAM operations shift from Processor to SDIMM.
ORAM traffic pattern shifts from the memory bus to on-SDIMM “private” buses.

The Independent ORAM Protocol

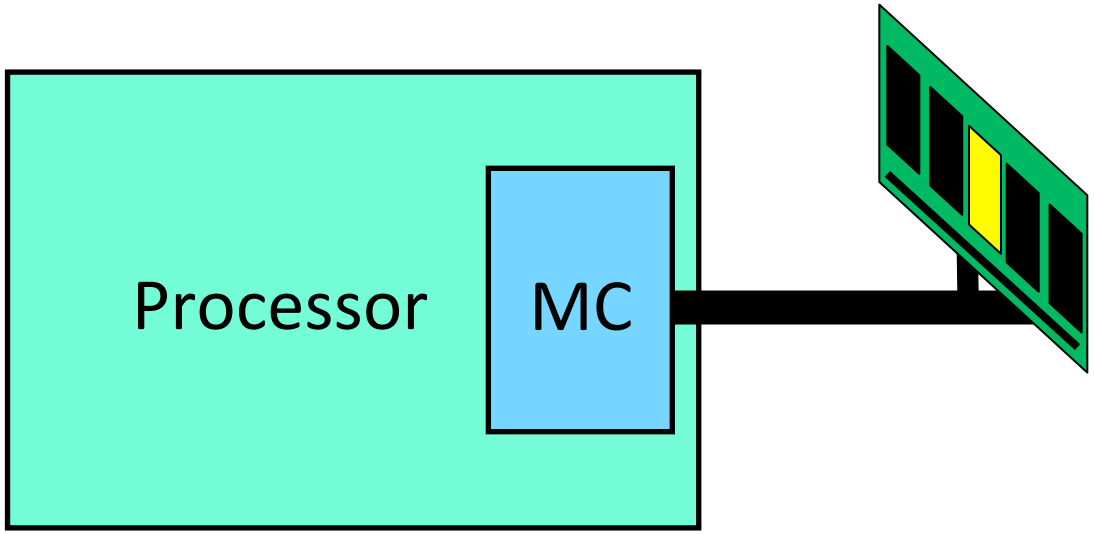

1. Each SDIMM handles a subtree of the ORAM tree.
2. Only traffic on shared memory channel: CPU requests and leaf-id re-assignments.
3. As much parallelism as the number of SDIMMs.

The Split ORAM Protocol

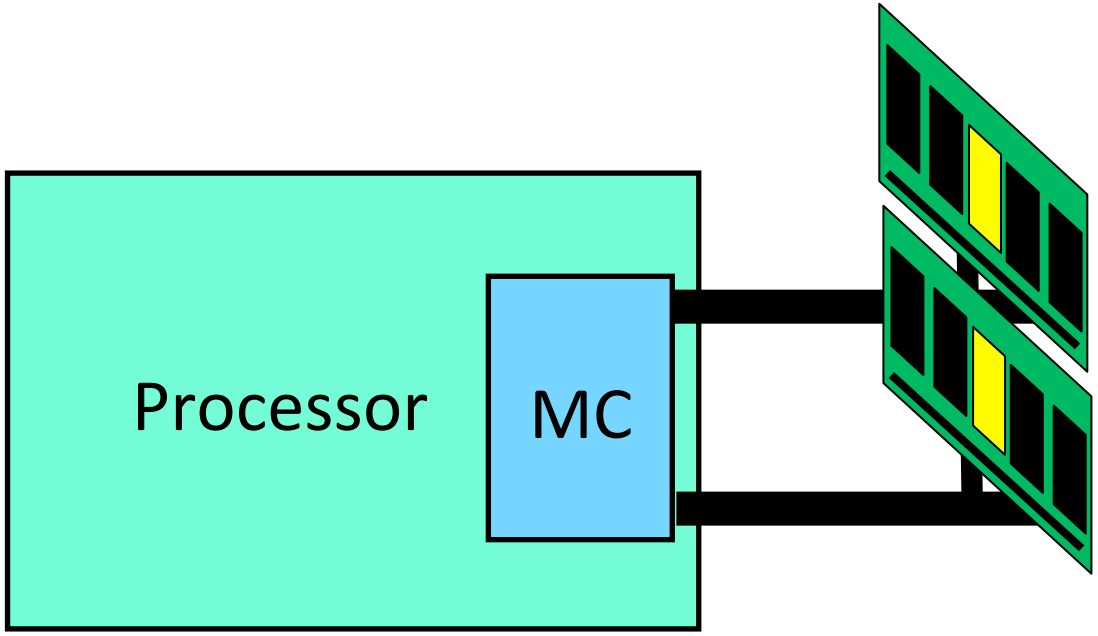

1. Each SDIMM handles a subset of every node.
2. Only metadata is sent to the processor.
3. The processor tells the SDIMMs how to shuffle data.
4. Lower latency per ORAM request, but lower parallelism as well.

ORAM Results Summary

- Can combine the Independent and Split protocols to find the best balance of latency and parallelism
- Bandwidth demands are reduced from 280x → 35x

  Execution time overheads from 5.2x → 2.7x
- Reduces memory energy by 2.5x

Overview

- Memory timing channels
  - The Fixed Service memory controller [MICRO 2015]
- Memory access patterns
  - Near-data ORAM [HPCA 2018]
  - Hierarchical ORAM [ASPLOS 2019]
- Memory integrity
  - Improving SGX with VAULT [ASPLOS 2018]

Intel SGX Basics

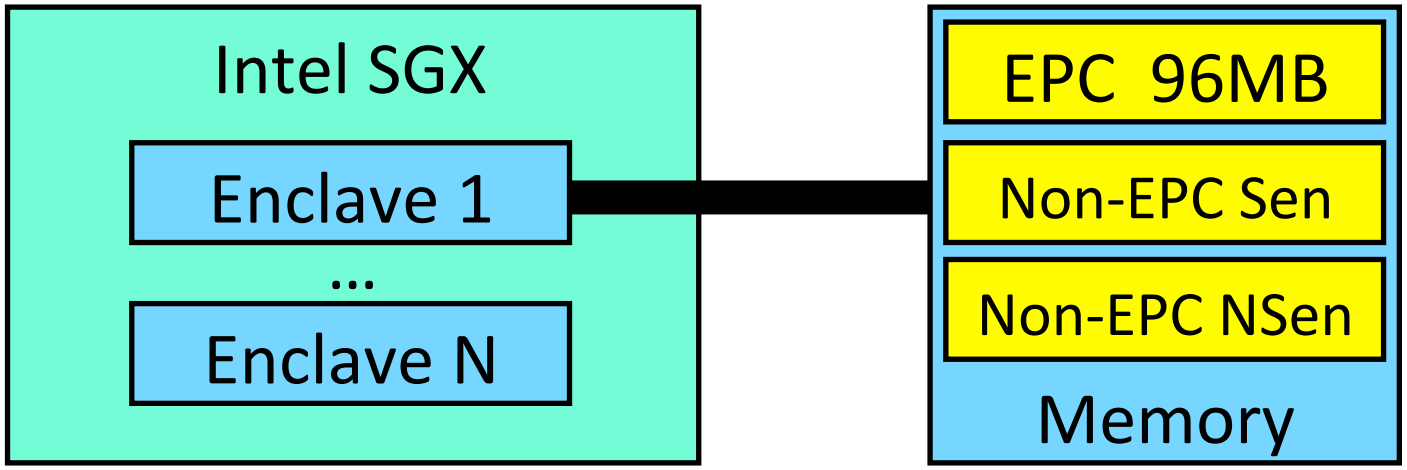

1. Enclave data is protected from malicious OS/ operator.
2. A per-block integrity tree protects EPC.
3. A per-page integrity tree protects non-EPC Sen.
4. This keeps overheads (bw and capacity) of integrity tree low.
5. Entails frequent paging between EPC and non-EPC.

Intel SGX Basics

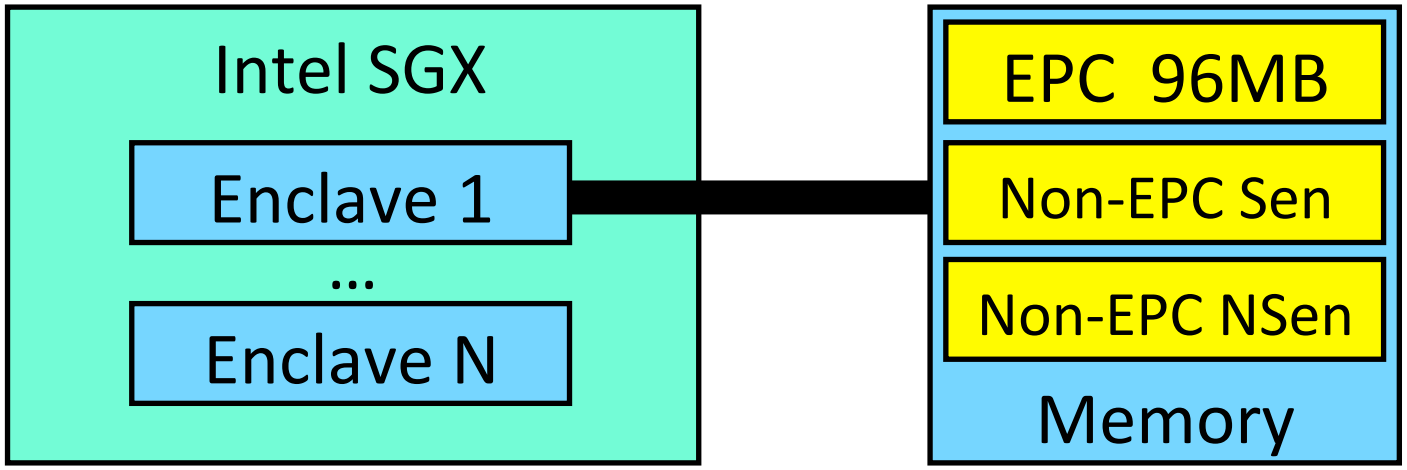

1. Enclave data is protected from malicious OS/ operator.
2. A per-block integrity tree protects EPC.
3. A per-page integrity tree protects non-EPC Sen.
4. This keeps overheads (bw and capacity) of integrity tree low.
5. Entails frequent paging between EPC and non-EPC.

VAULT: Unify EPC and non-EPC to reduce paging. New integrity tree for low bw. Better metadata for capacity.

SGX Overheads

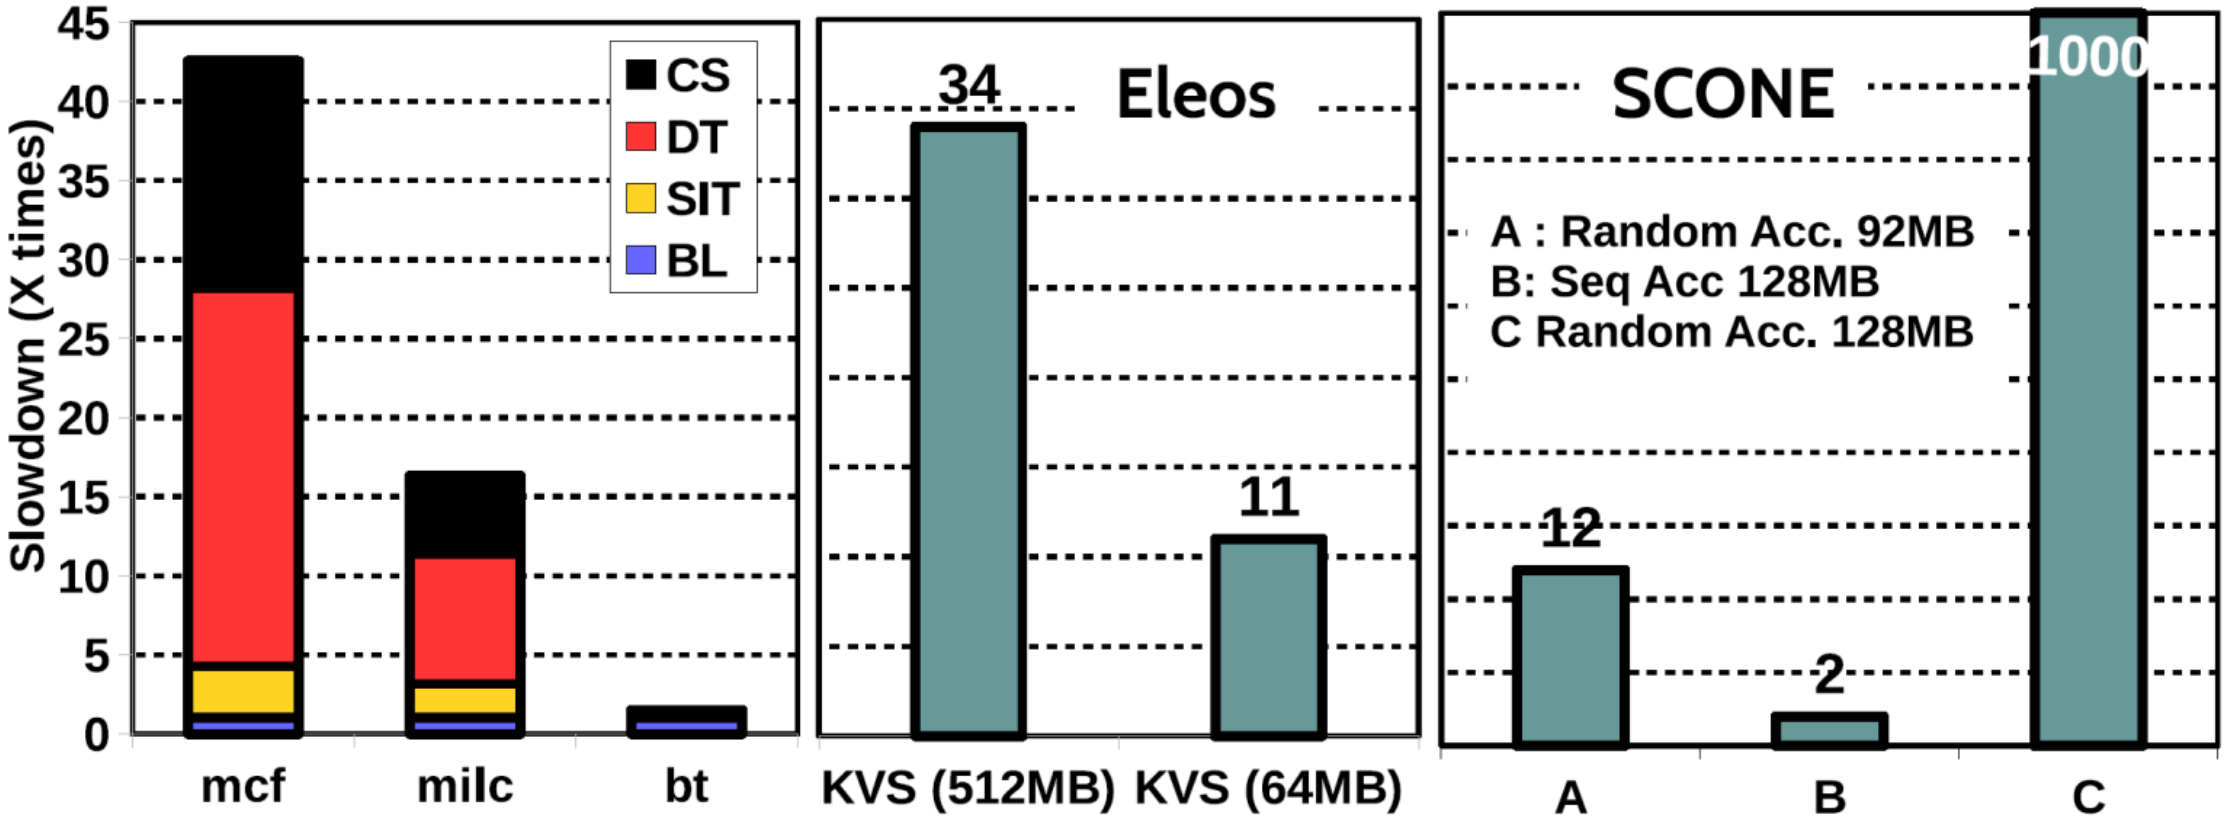

Bonsai Merkle Tree

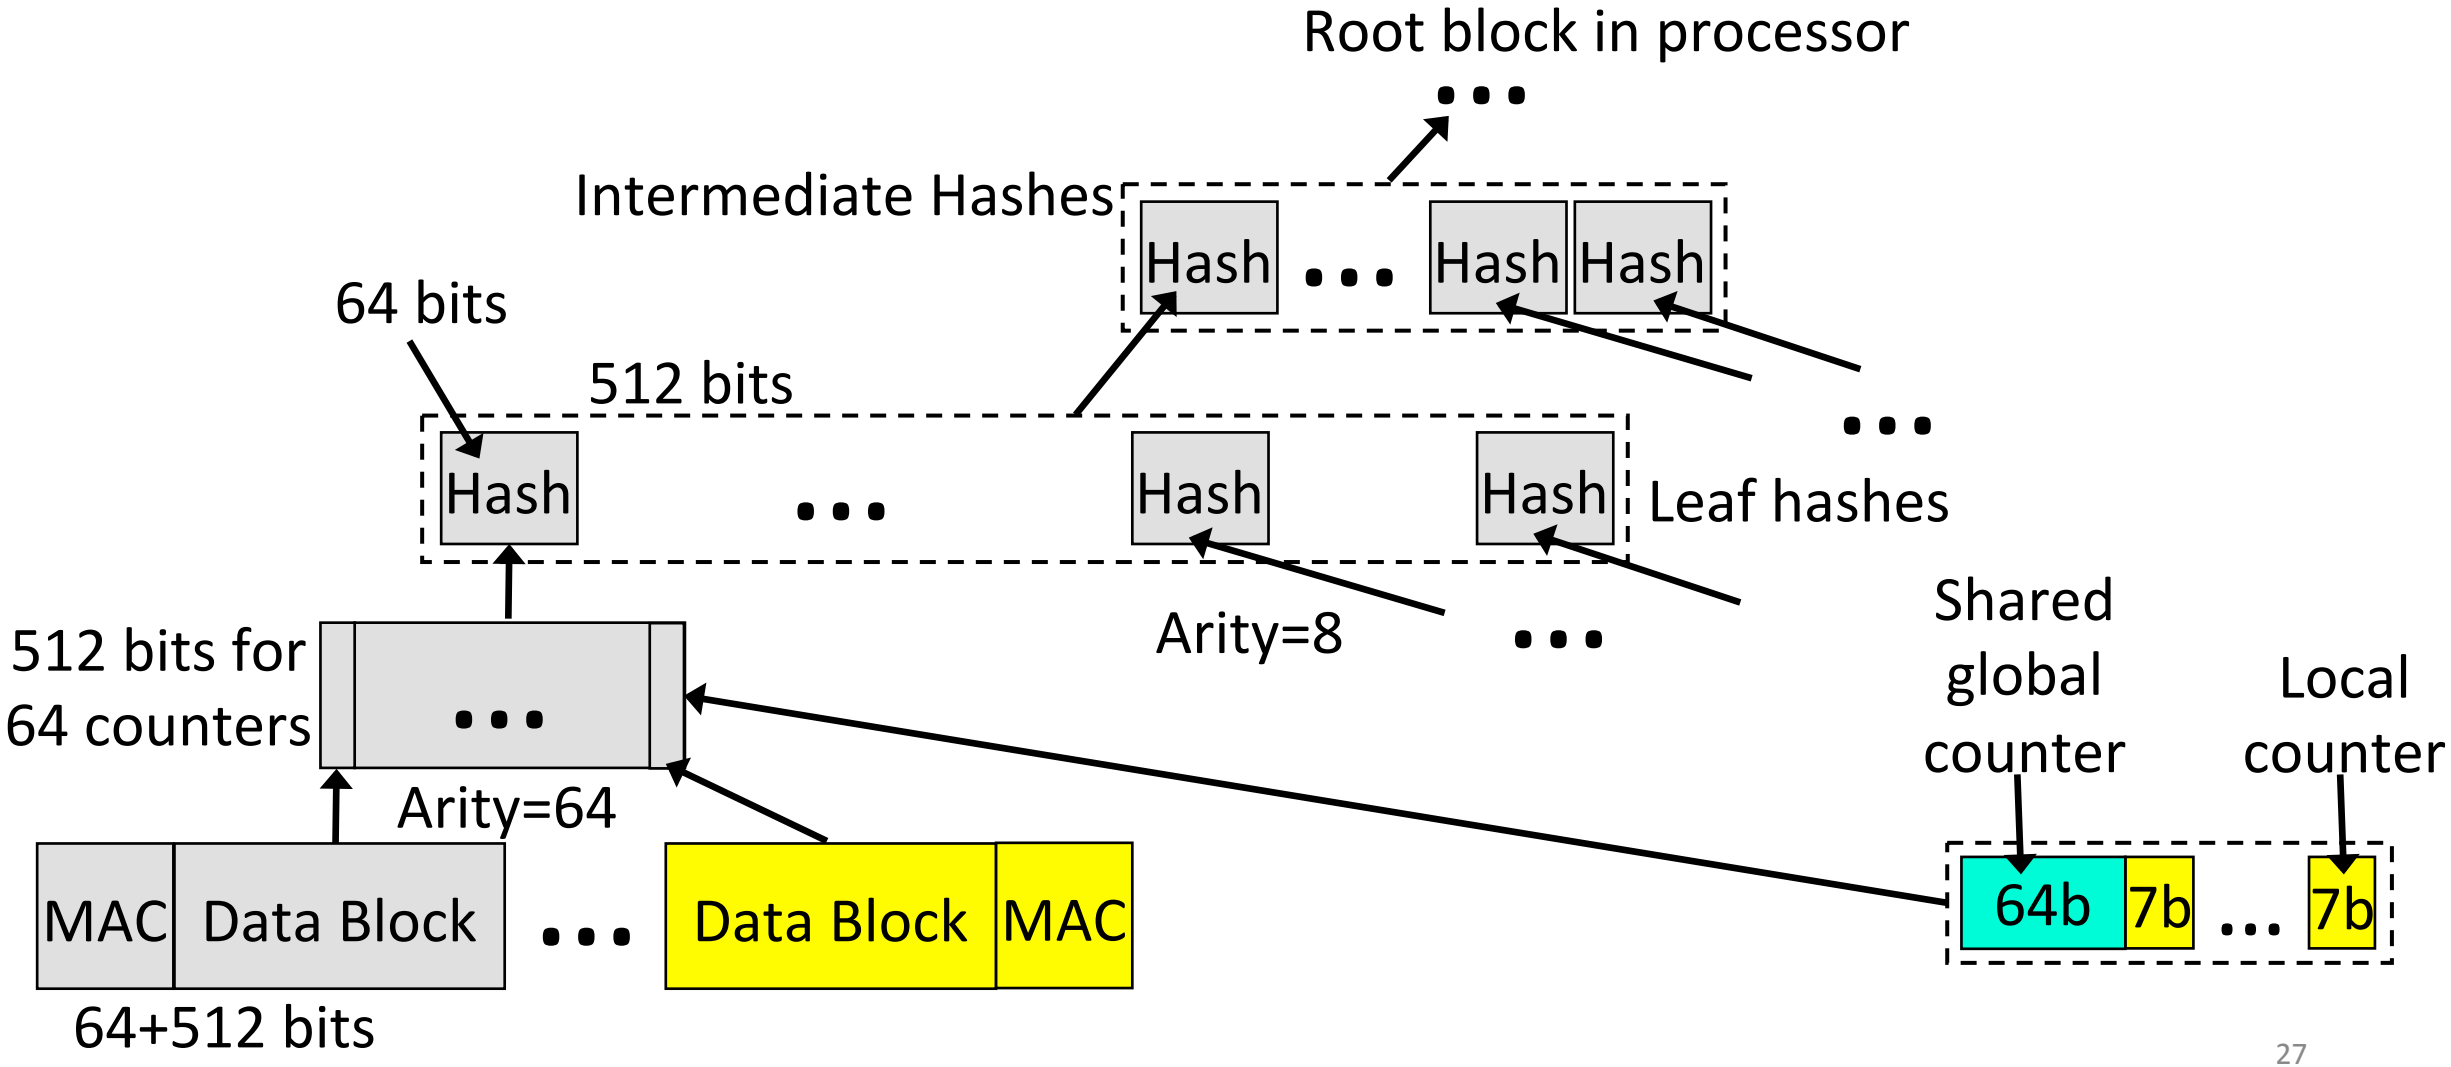

VAULT

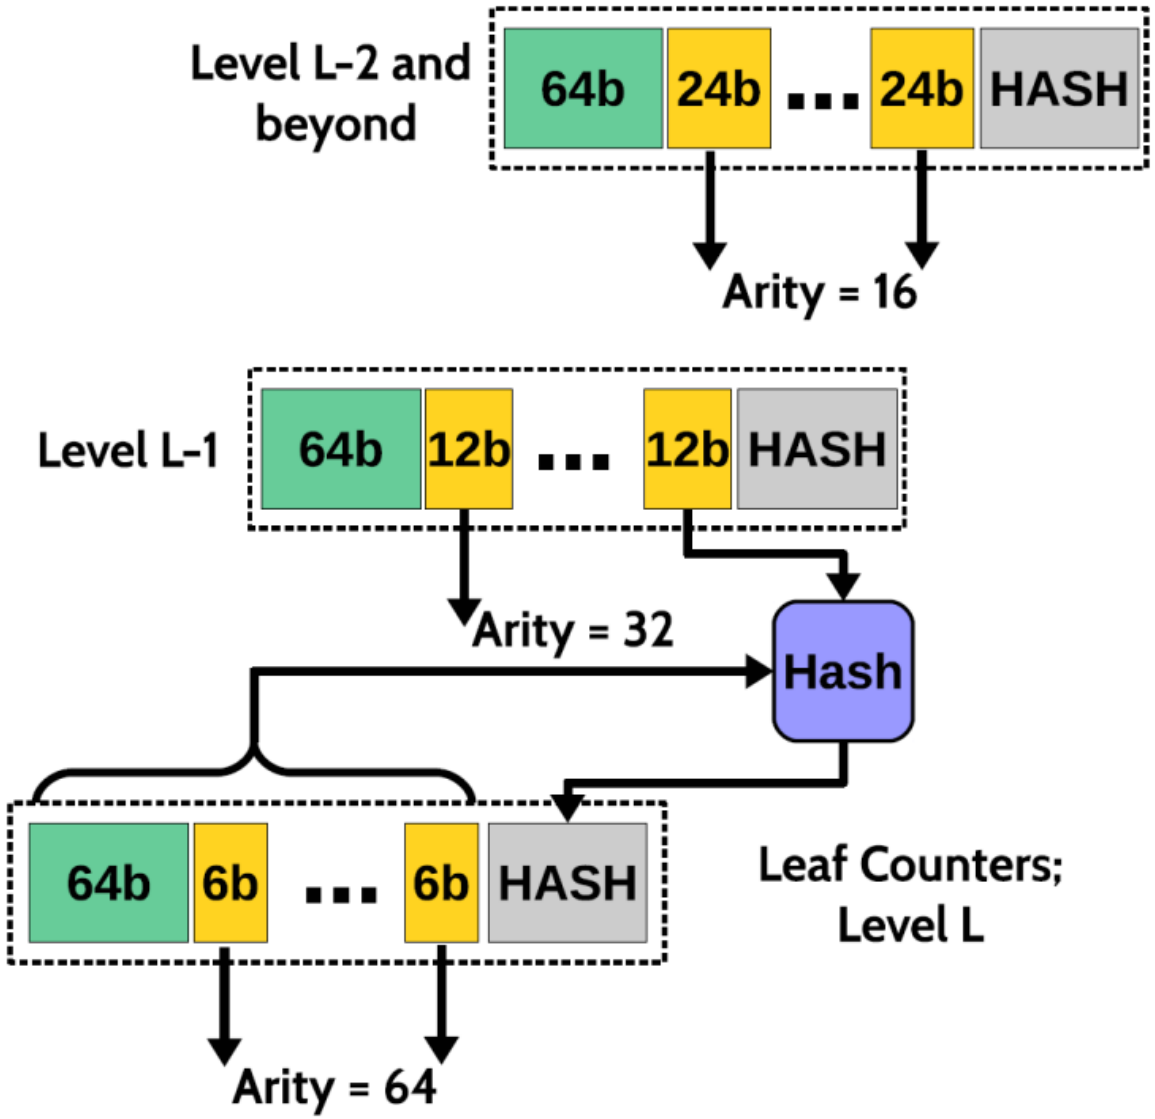

1. Small linkage counters ➔ high arity, compact/shallow tree, better cacheability.
2. Variable counter width to manage overflow.
3. Reduces bandwidth overhead for integrity verification.

VAULT+SMC

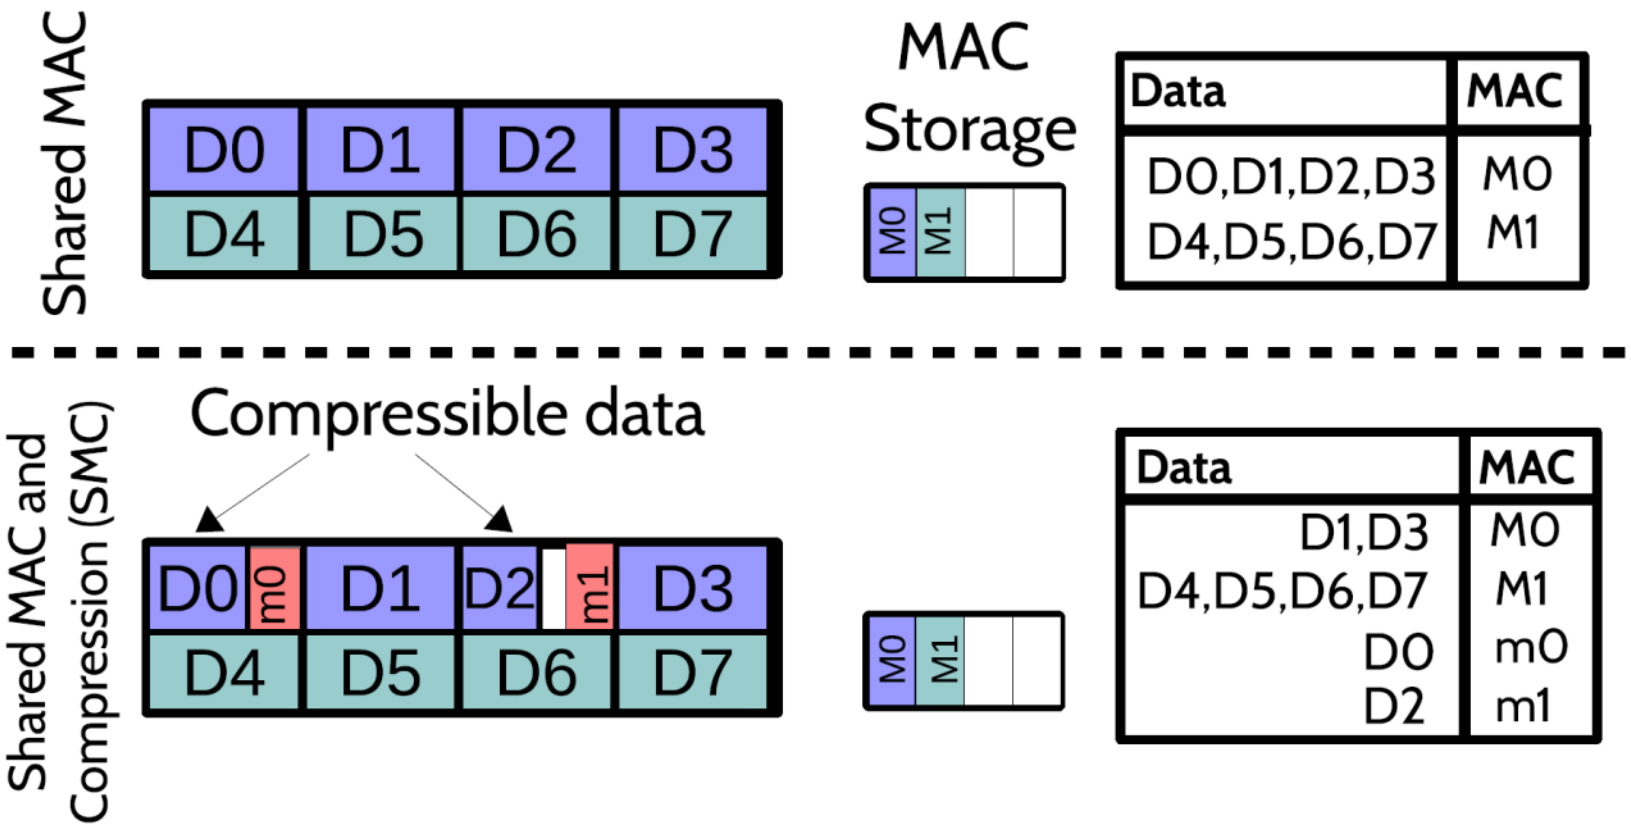

1. MAC storage and bw overheads are high.
2. Sharing a MAC among 4 blocks reduces storage, but incr bw.
3. A block is compressed and the MAC is embedded in the block ➔reduces bw and storage.

Integrity Results Summary

- 3.7x performance improvement over SGX – primarily because of lower paging overheads
- A large effective EPC is palatable – 4.7% storage overhead and a more scalable tree (34% better than the SGX tree)

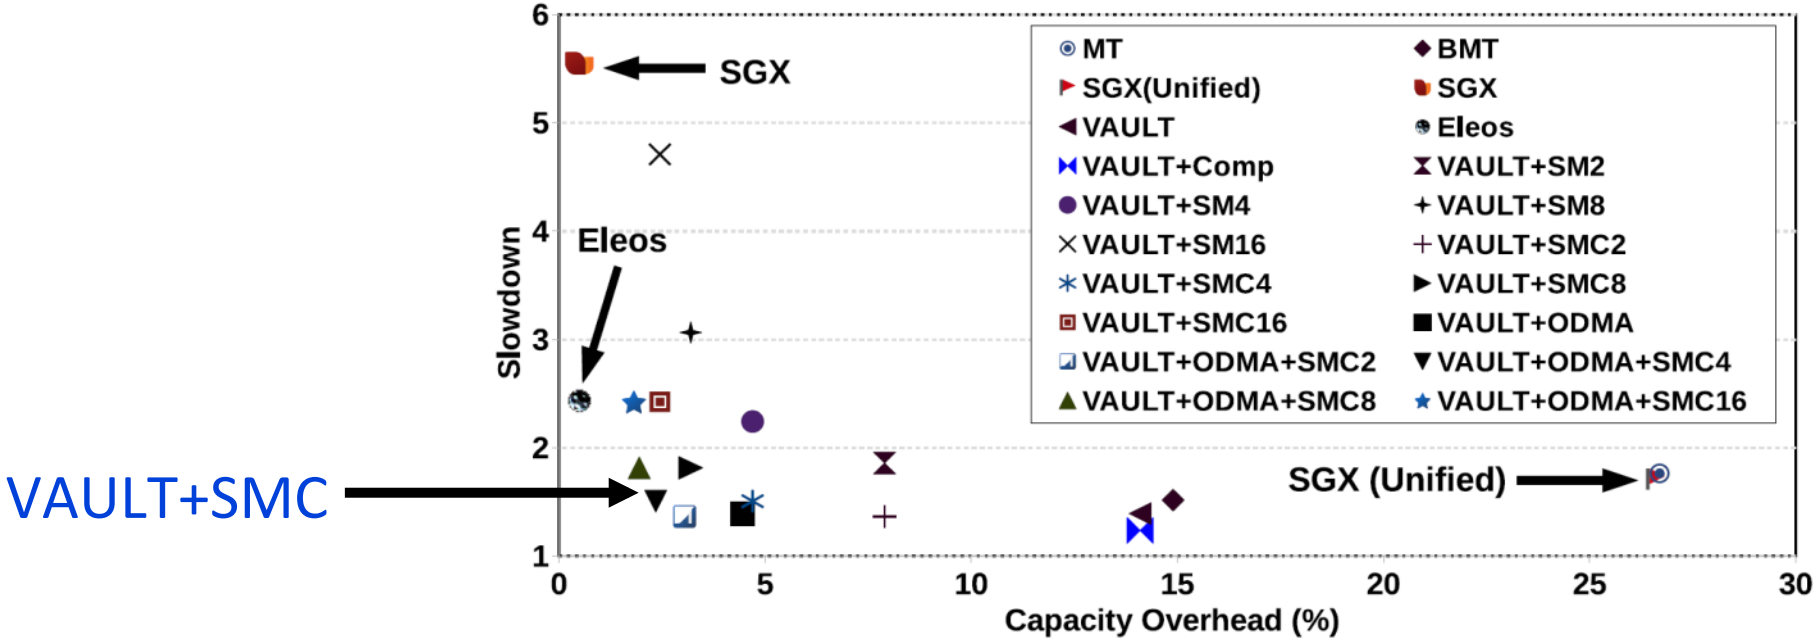

Big Finish

- Memory defenses were purely academic pursuits
- Integrity trees now a part of Intel SGX: overheads of 2x – 40x
- VAULT improves integrity overhead to 1.5x – 2.5x
- FS eliminates timing channels with overhead of 2x
- SDIMM improves ORAM overhead to 2.7x
- An array of memory defenses is now commercially viable

  … and strategic given latent vulnerabilities

Acks: Ali Shafiee, Meysam Taassori, Chandru Nagarajan, Akhila Gundu, Manju Shevgoor, Andrew Vuong, Mohit Tiwari, Feifei Li, NSF, Intel.