Performance Analysis of Scientific Computing Workloads on General Purpose TEEs

Ayaz Akram (yazakram@ucdavis.edu), Anna Giannakou, Venkatesh Akella, Jason Lowe-Power, Sean Peisert

Summary

Can TEEs enable secure scientific computing?

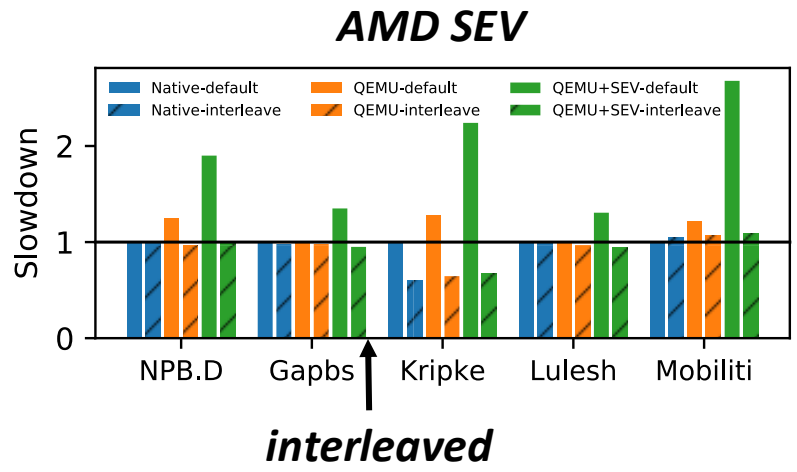

AMD SEV shows little performance degradation if used with **interleaved** NUMA allocation

- Irregular workloads can still show virtualization-based overheads

Intel SGX is inappropriate for scientific computing

- Incurs high performance overheads
- Programming model requires application modifications

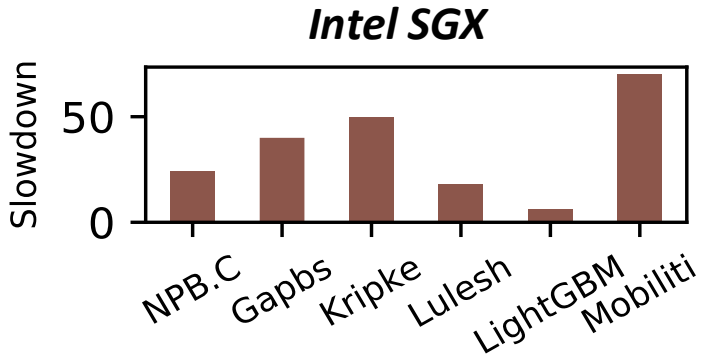

Outline

Why Secure High Performance Computing Matters?

Performance Analysis of General-Purpose TEEs for HPC

- What special configurations AMD SEV need?
- Why Intel SGX is not an appropriate fit for HPC?

Future Trends in TEEs

Secure High-Performance Computing

How to compute with large sensitive data?
- Biomedical data
- Proprietary data

Security threats in HPC centers
- External
- Internal

A usual tradeoff in HPC centers
- Risk acceptance vs data hosting

Usability challenge of secure environments

UCDAVIS

Motivation for this Work

Related Work

- Cloud[1] or General-Purpose[2] Computing Centric
- No focus on HPC

Distinction between HPC and Cloud Computing

- Resources allocated to a single user at a time
- Highly multithreaded apps, batched together
- Large working sets
- Scale across many nodes
- Perform limited types of I/O

[1] Gjerdrum et al., Performance of Trusted Computing in Cloud Infrastructures with Intel SGX, CLOSER 2017.

[2] Mofrad et al., A comparison study of intel SGX and AMD memory encryption technology, HASP 2018.

UCDAVIS

What are Trusted Execution Environments (TEEs)?

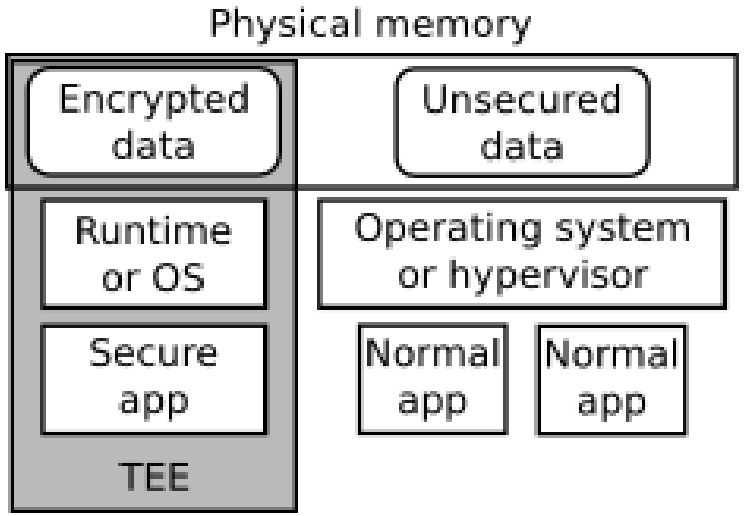

Trusted Execution Environments provide

- *hardware-enforced isolation*
- *cryptographic attestation to verify execution*
- *no significant usability challenges*

AMD SEV

arm TRUSTZONE

KeyStone

How TEEs Fit in Our Threat Model?

Threat Model

Not Trusted

- HPC System Admin.
- Operating System
- Hypervisor
- HPC users sharing the resources

Physical attacks and side-channels not within scope

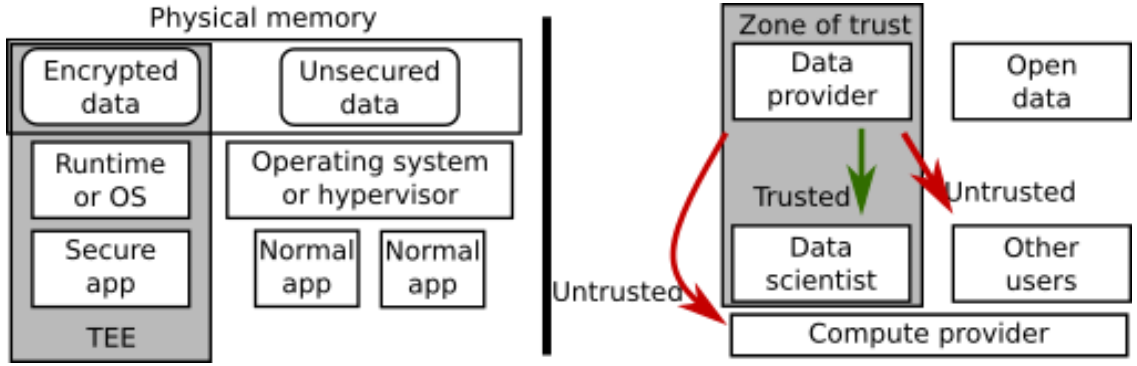

Outline

Why Secure High Performance Computing Matters?

Performance Analysis of General-Purpose TEEs for HPC

- What special configurations AMD SEV need?
- Why Intel SGX is not an appropriate fit for HPC?

Future Trends in TEEs

What are the Performance Implications of Current TEEs?

SEV does not incur significant performance degradation

- Default NUMA penalty can be high
- Interleaved NUMA policy improves performance

Irregular workloads performance suffers due to virtualization when running under SEV

SGX shows high performance overheads

- Does not support unmodified applications

We Analyze Two TEEs

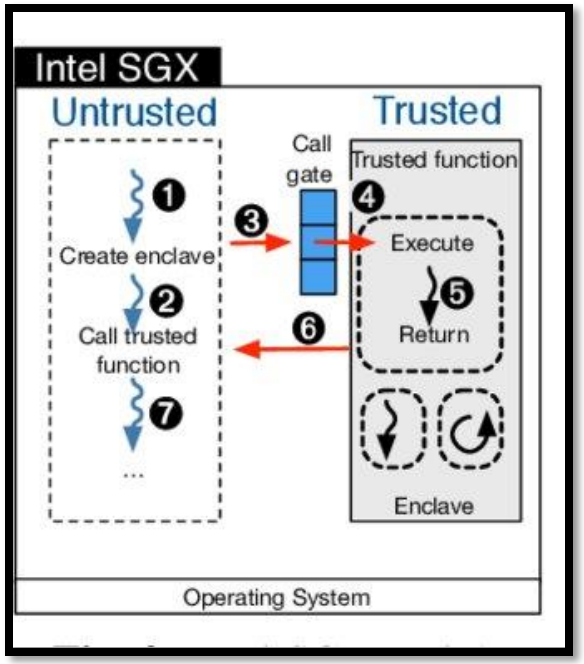

[1]

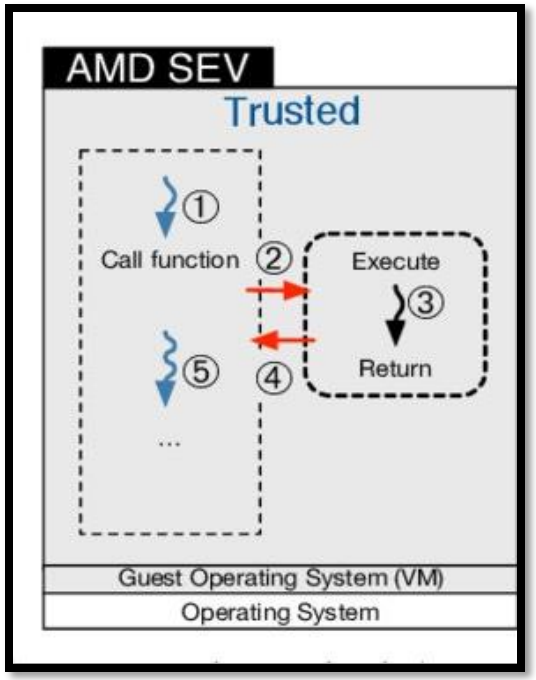

[1]

Technology	Integrity	TCB Size	Secure Memory Size	Application Changes
Intel SGX	Yes	Small	128 MB (useable: 94MB)	Required
AMD SEV	No	Large	Up to RAM size	Not Required

[1] Christian Göttel et al. "Security, performance and energy trade-offs of hardware-assisted memory protection mechanisms." IEEE Symposium on Reliable Distributed Systems (SRDS), 2018.

Workloads Evaluated

Traditional HPC

NAS Parallel Benchmarks

Graph Workloads

GAPBS (US road network)

Modern HPC

LULESH, Kripke, Mobiliti, LightGBM, BLASTN

UCDAVIS

Hardware Platforms Used

Feature	AMD SEV 1	AMD SEV 2	AMD SEV 3	Intel SGX
CPU	EPYC 7401P	EPYC 7702	EPYC 7402P	Core i7-8700
Sockets	1	2	1	1
Cores	24	128	24	6
NUMA	4 Nodes	2 Nodes	1 Node	1 Node
RAM	64GB	1TB	64GB	32GB

AMD Rome (7702)

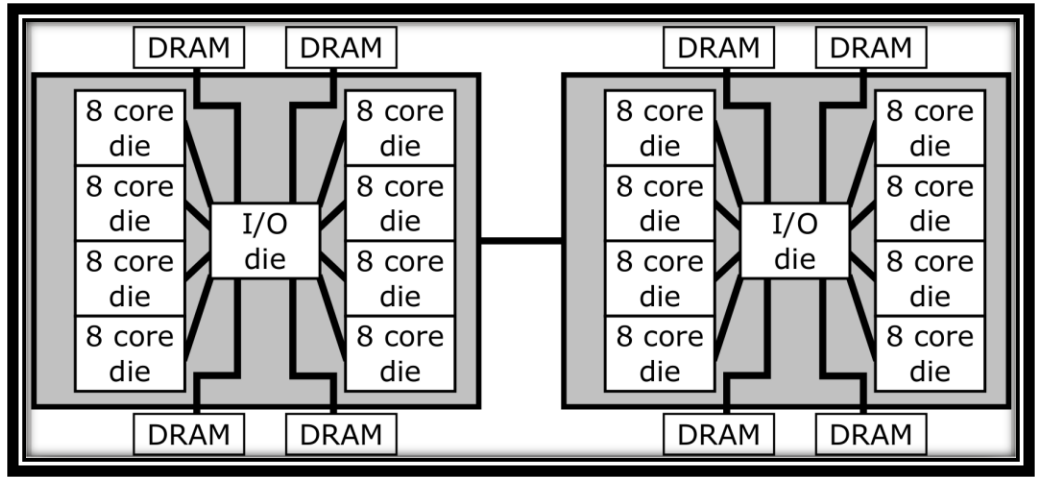

AMD Naples (7401)

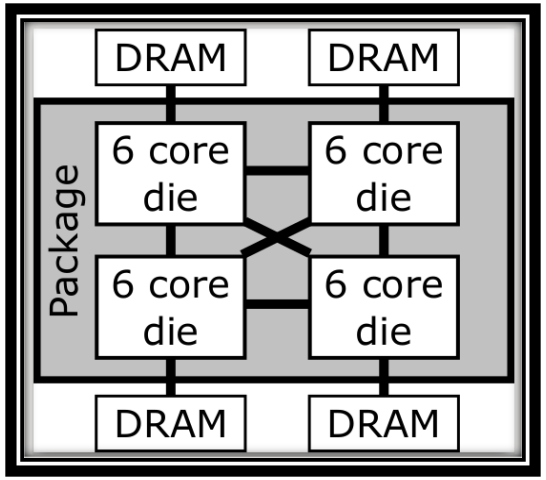

Outline

Why Secure High Performance Computing Matters?

Performance Analysis of General-Purpose TEEs for HPC

- **What special configurations AMD SEV need?**
- Why Intel SGX is not an appropriate fit for HPC?

Open Problems and Future Plans

Performance Impact of SEV for NPB (D) on AMD Naples

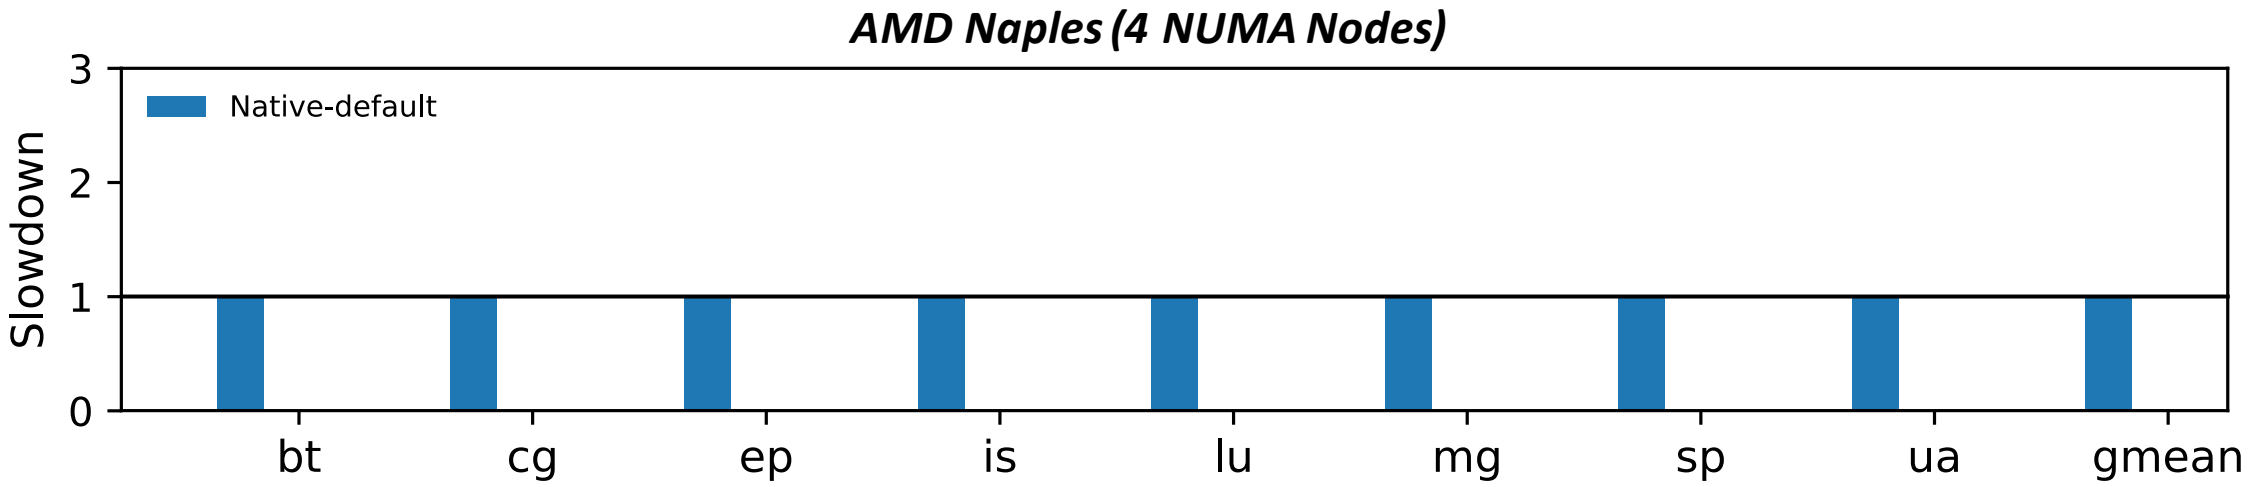

Performance Impact of SEV for NPB (D) on AMD Naples

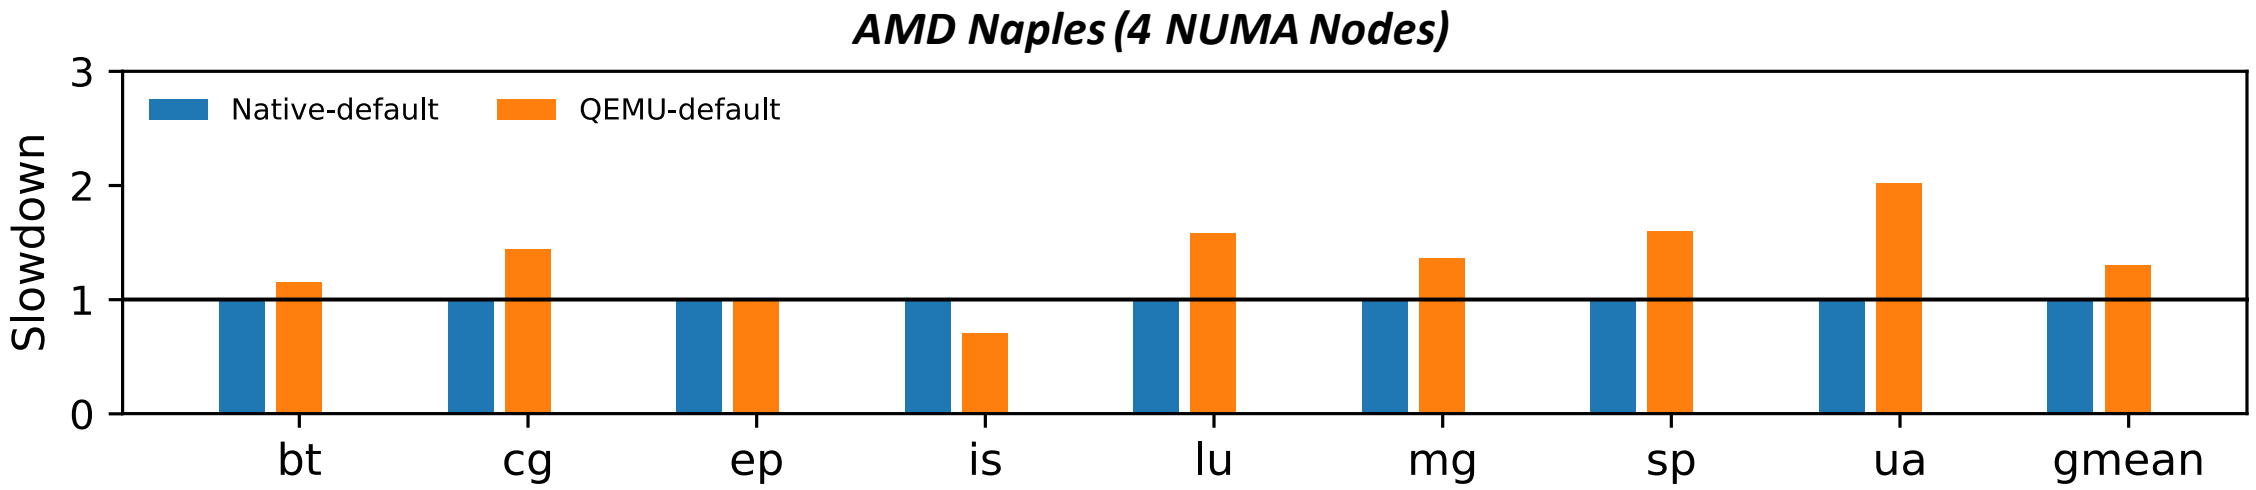

Performance Impact of SEV for NPB (D) on AMD Naples

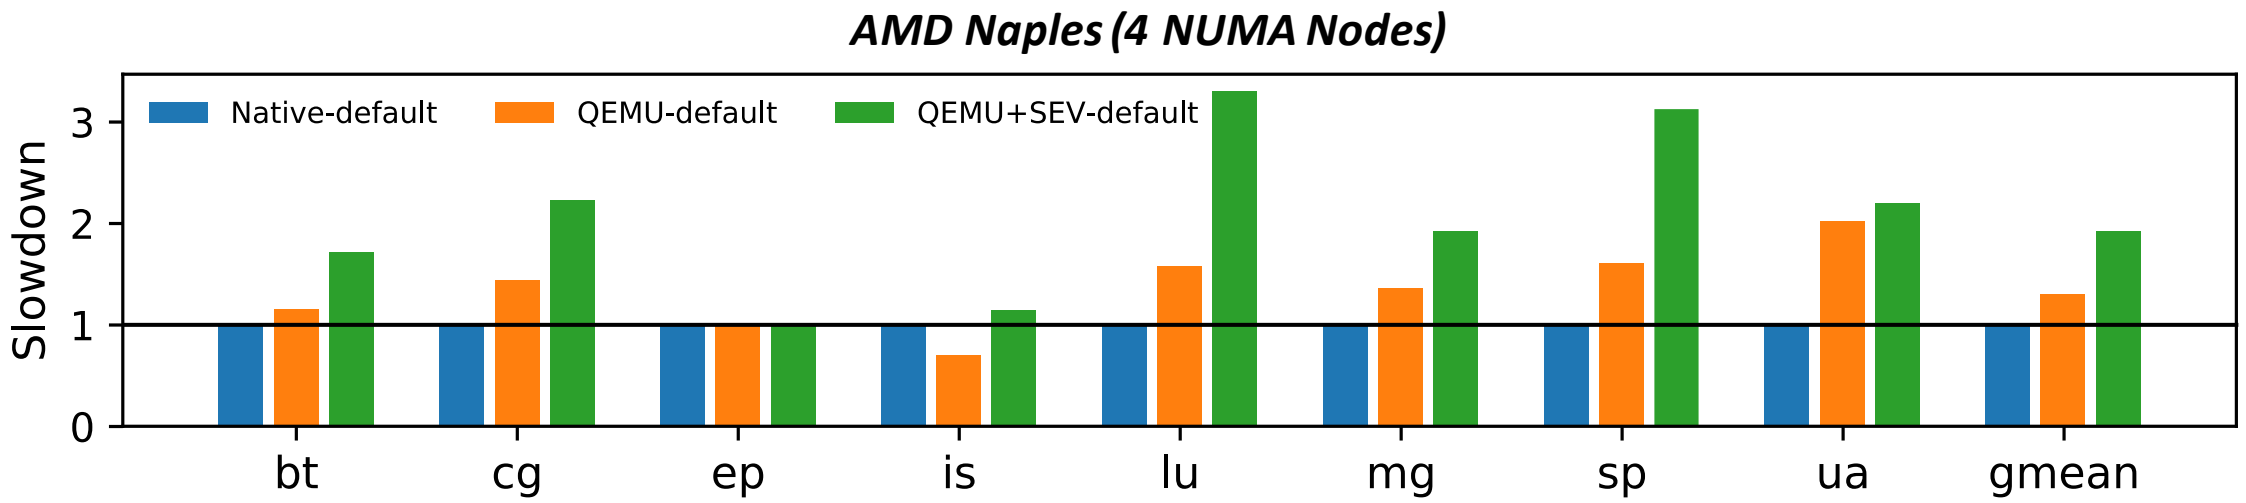

Comparing it to AMD Rome

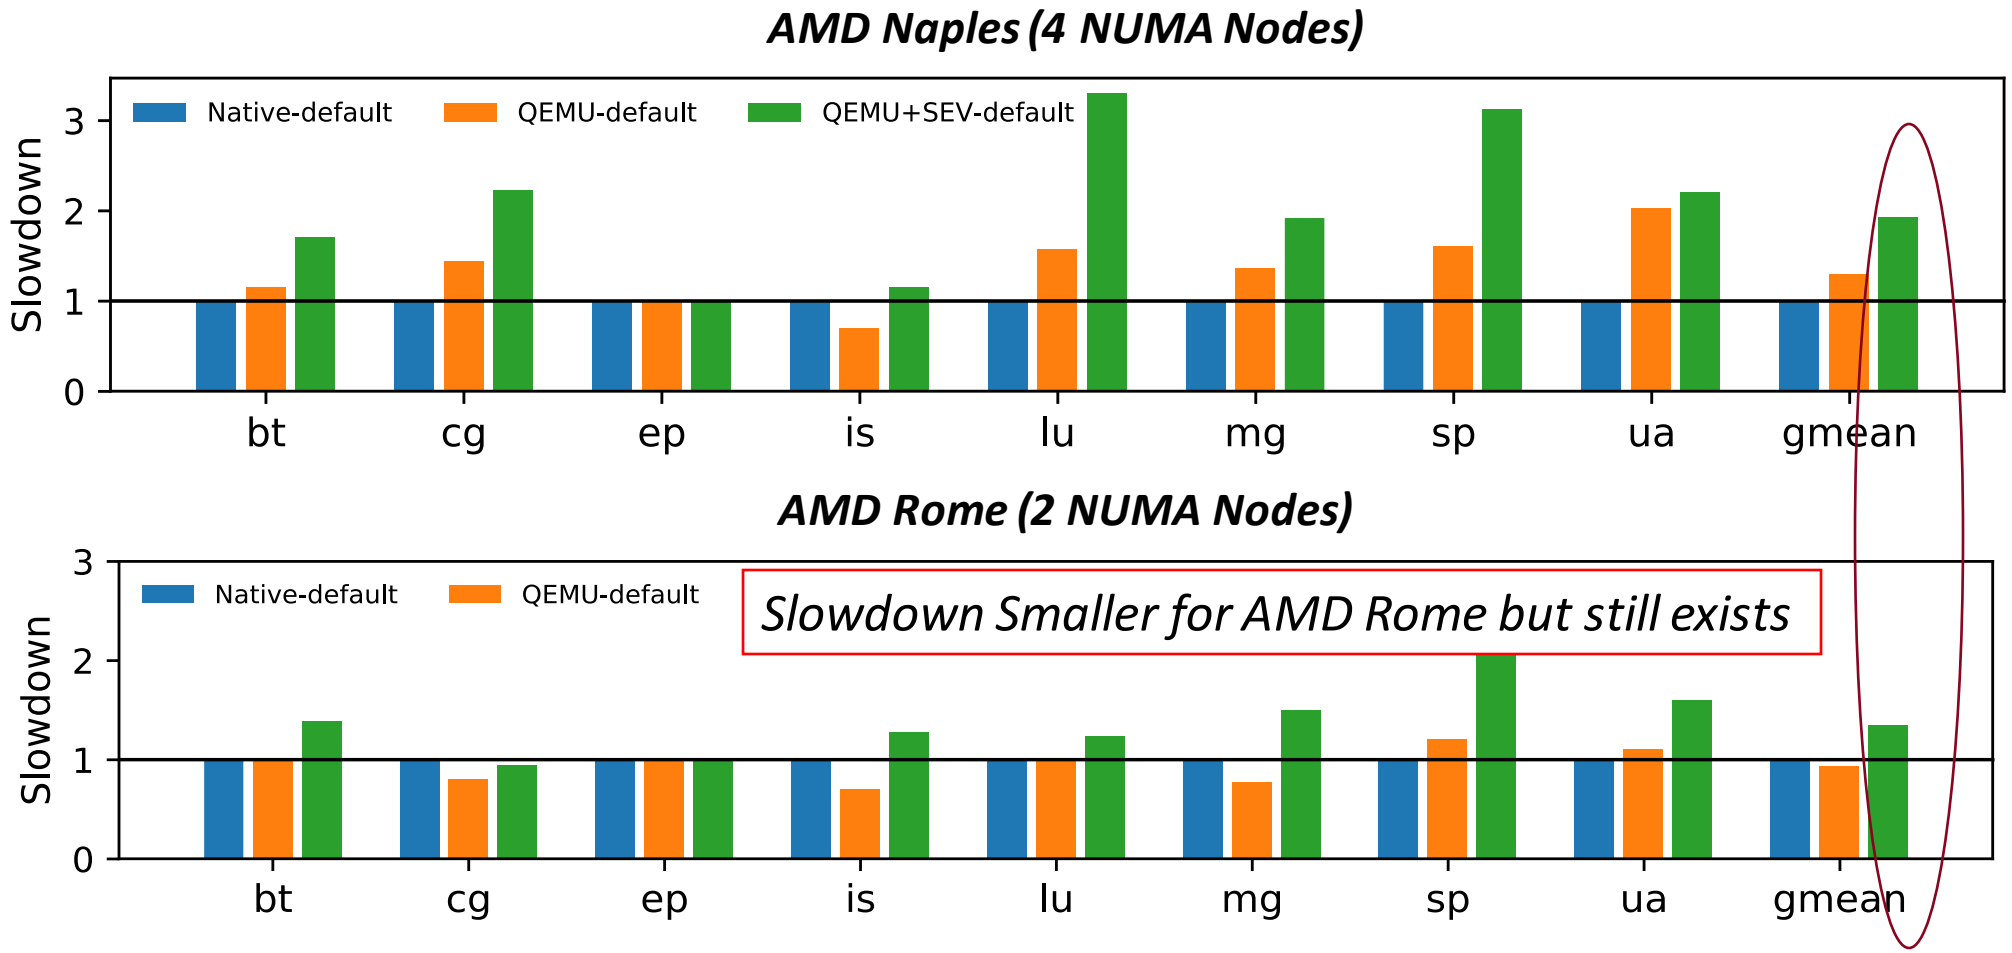

Memory Allocation on an AMD Naples System

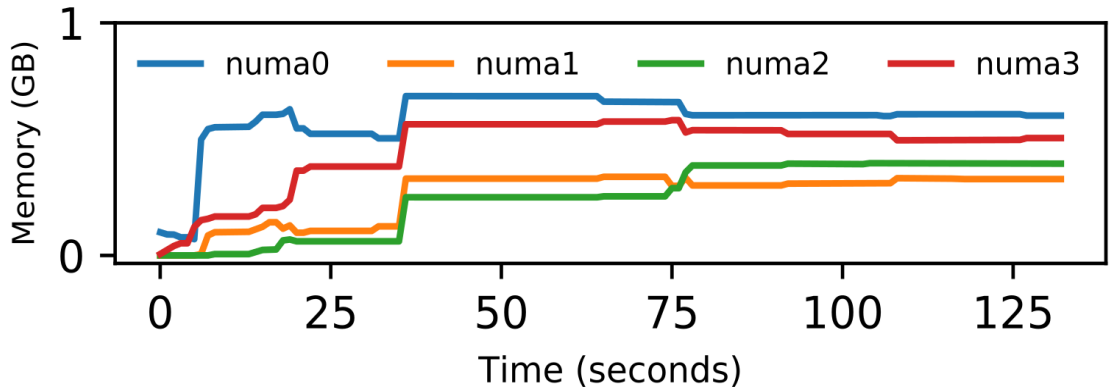

When only QEMU is used

- On demand memory allocation

Memory Allocation on an AMD Naples System

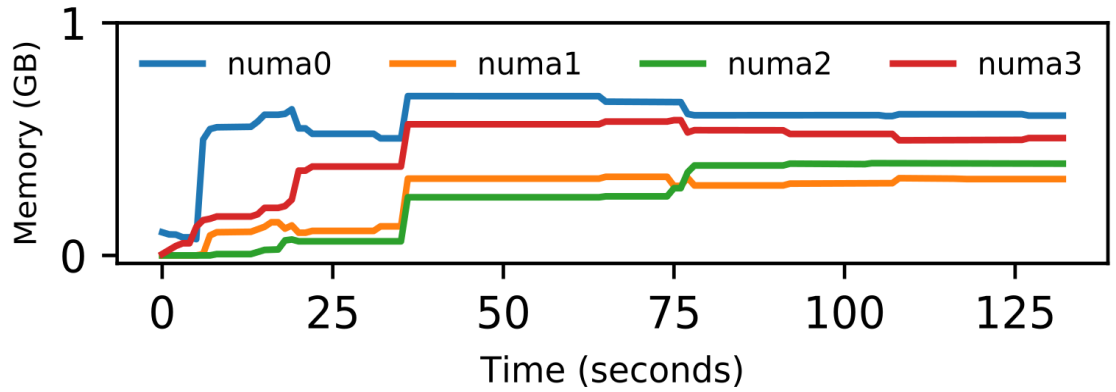

When only QEMU is used

On demand memory allocation

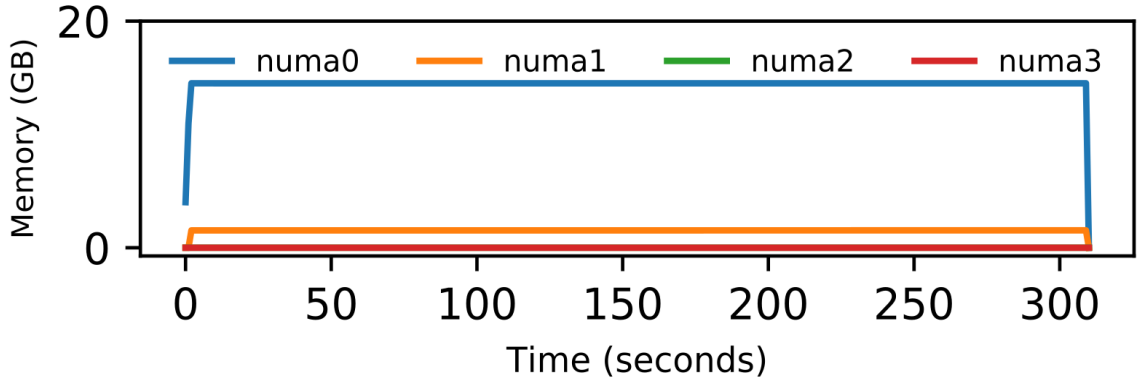

When QEMU+SEV is used

Memory allocation in the beginning

mlock is responsible

Why SEV requires locking pages to physical addresses?

SEV encryption implementation makes use of memory page's physical address

Due to default NUMA policy of "first touch" all memory gets allocated on a single node

Under-utilization of memory

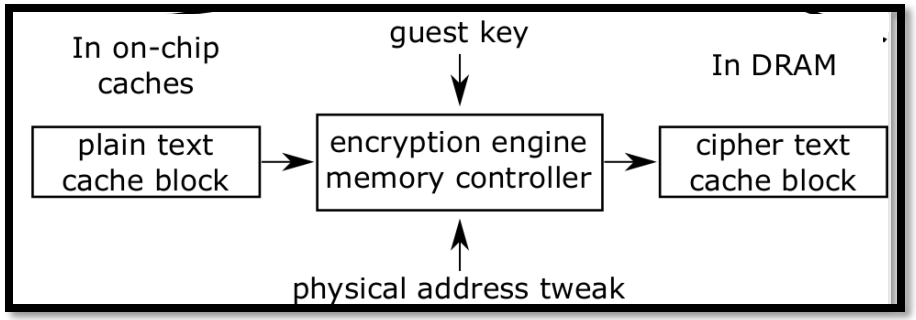

What about Interleaved NUMA Allocation?

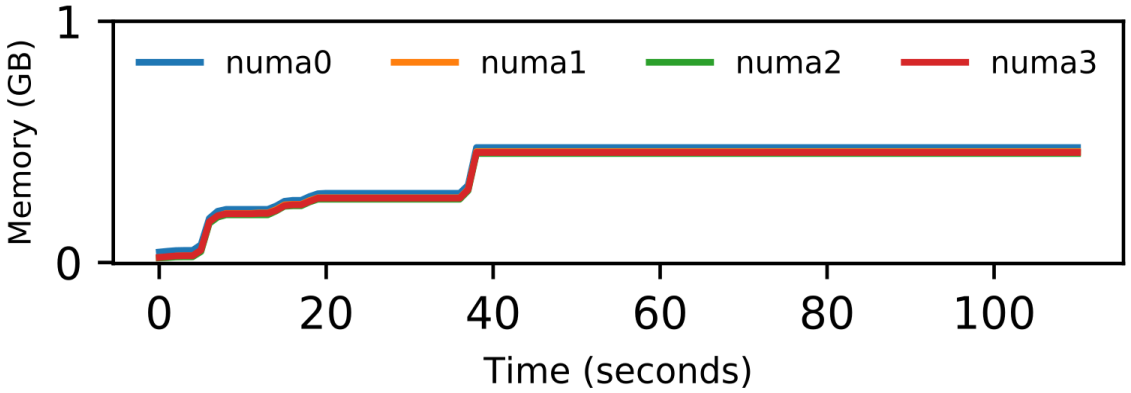

When only QEMU is used

Equal memory (on-demand) gets allocated on each node

What about Interleaved NUMA Allocation?

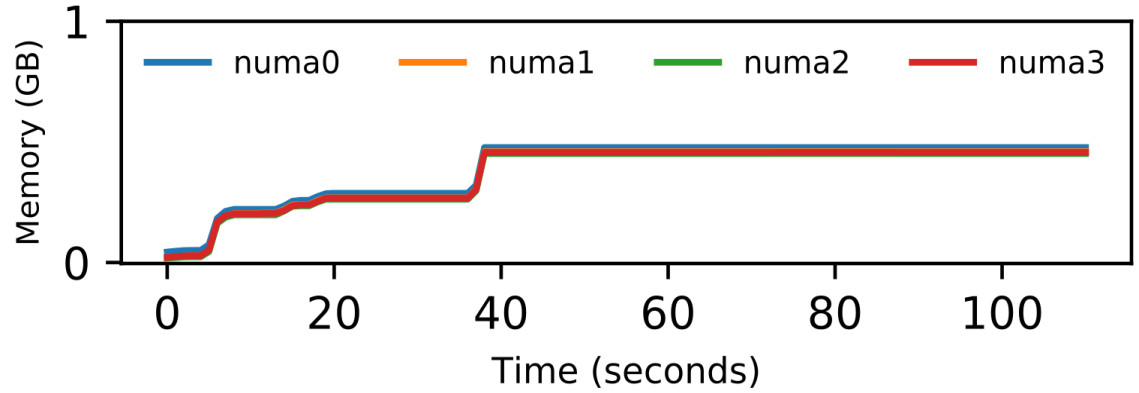

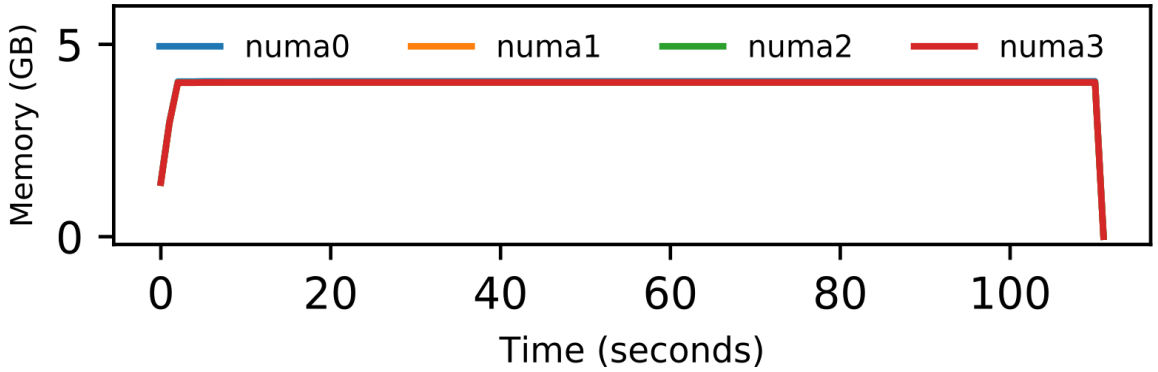

When only QEMU is used

Equal memory (on-demand) gets allocated on each node

When QEMU+SEV is used

Equal memory gets allocated on all nodes in the beginning

NUMA sensitive workloads prefer this

UCDAVIS

Performance Impact of SEV for NPB (D) on AMD Naples

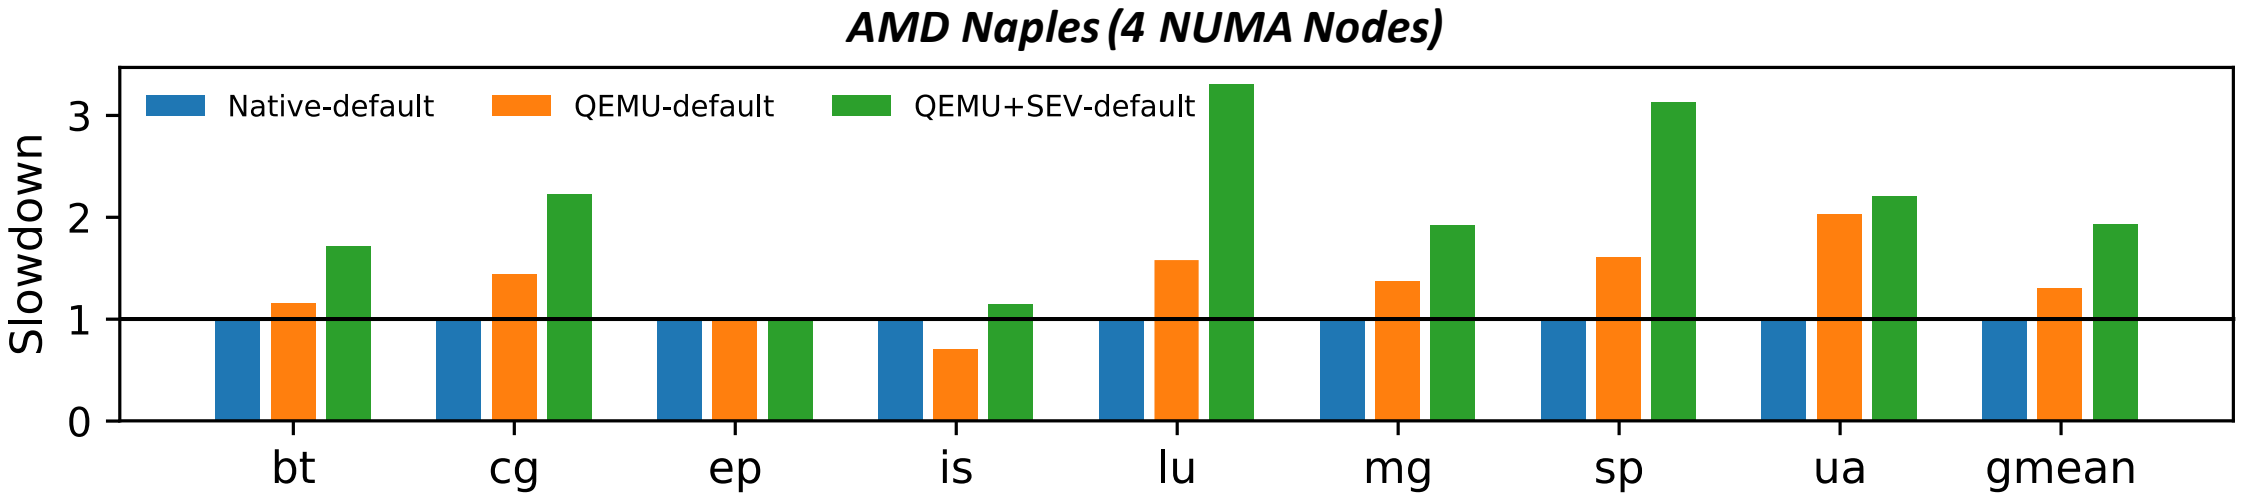

Performance Impact of SEV for NPB (D) on AMD Naples

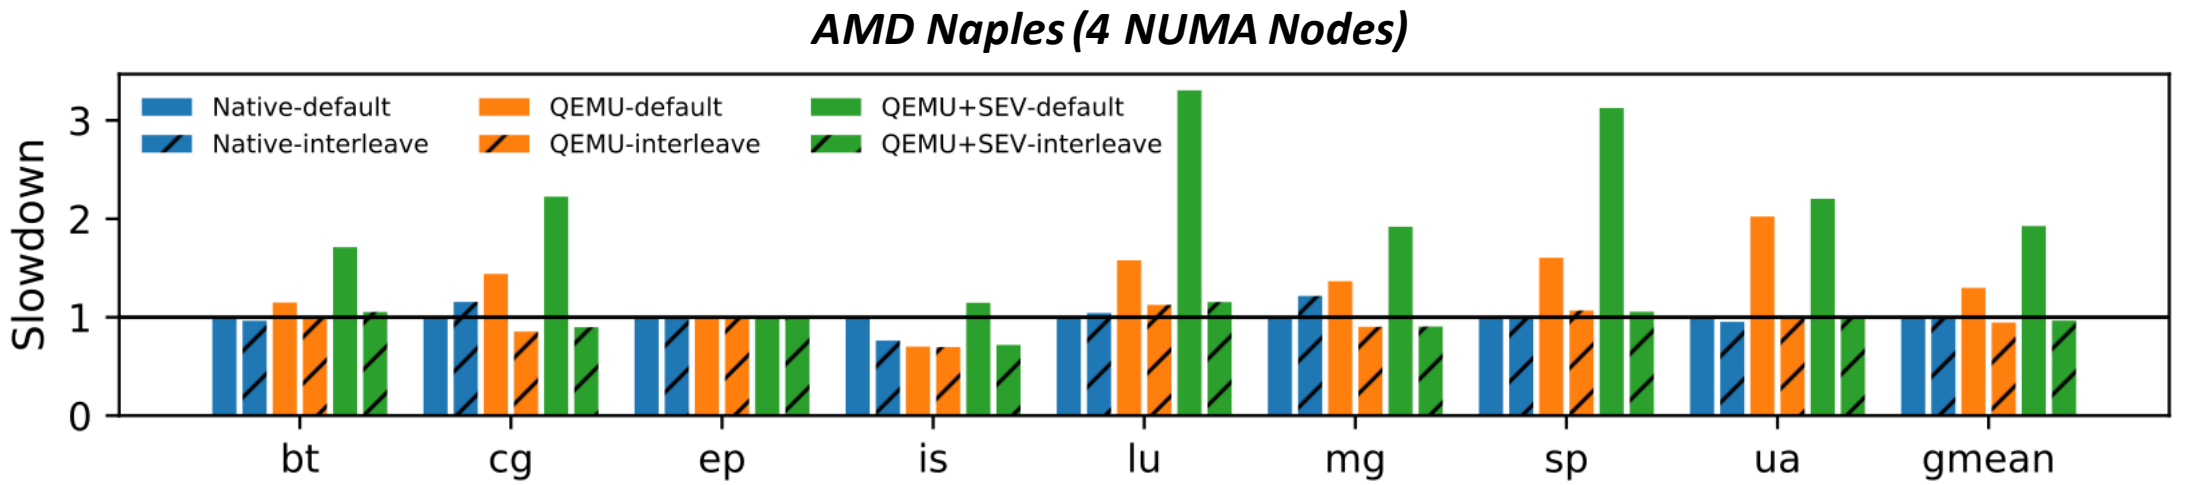

Hatched Bars show performance with interleaved allocation

Finding 1

SEV can be used for secure scientific computing without significant performance degradation for most workloads if it is configured correctly

Performance Impact of SEV for GAPBS and Other HPC Workloads

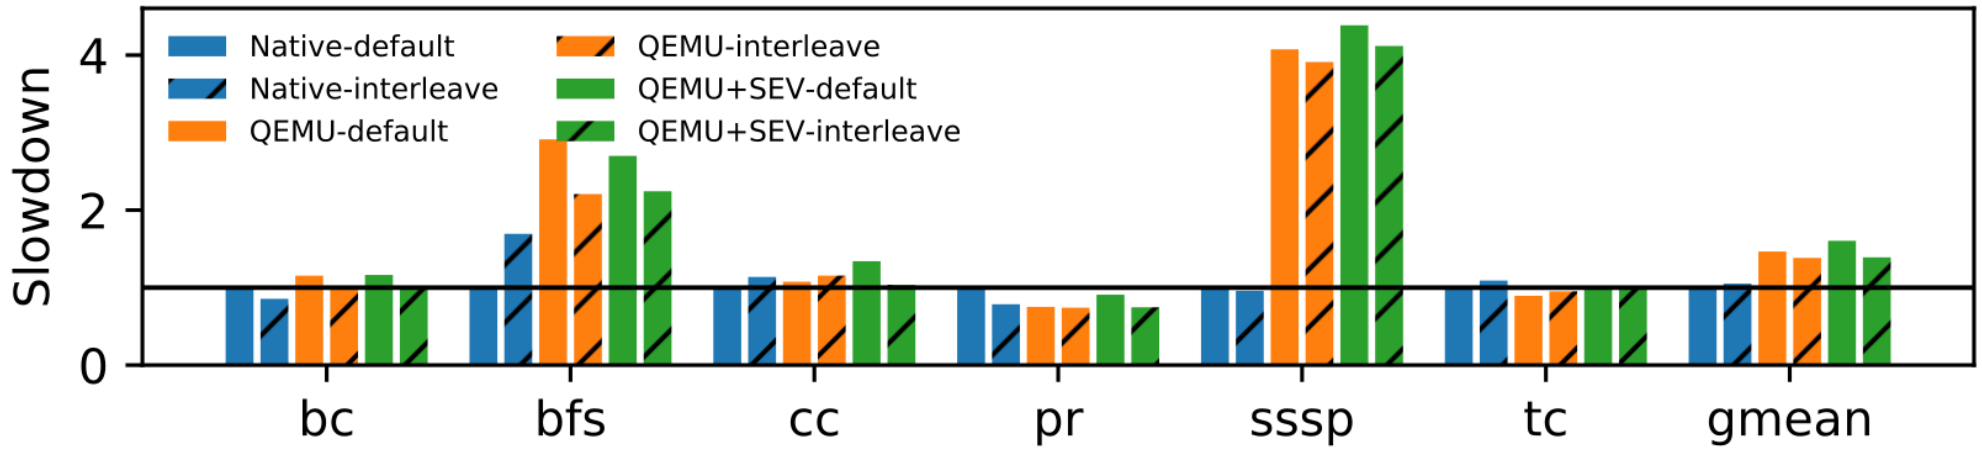

GAPBS on AMD Rome (128 Core Machine)

High thread contention leads to high KVM exits

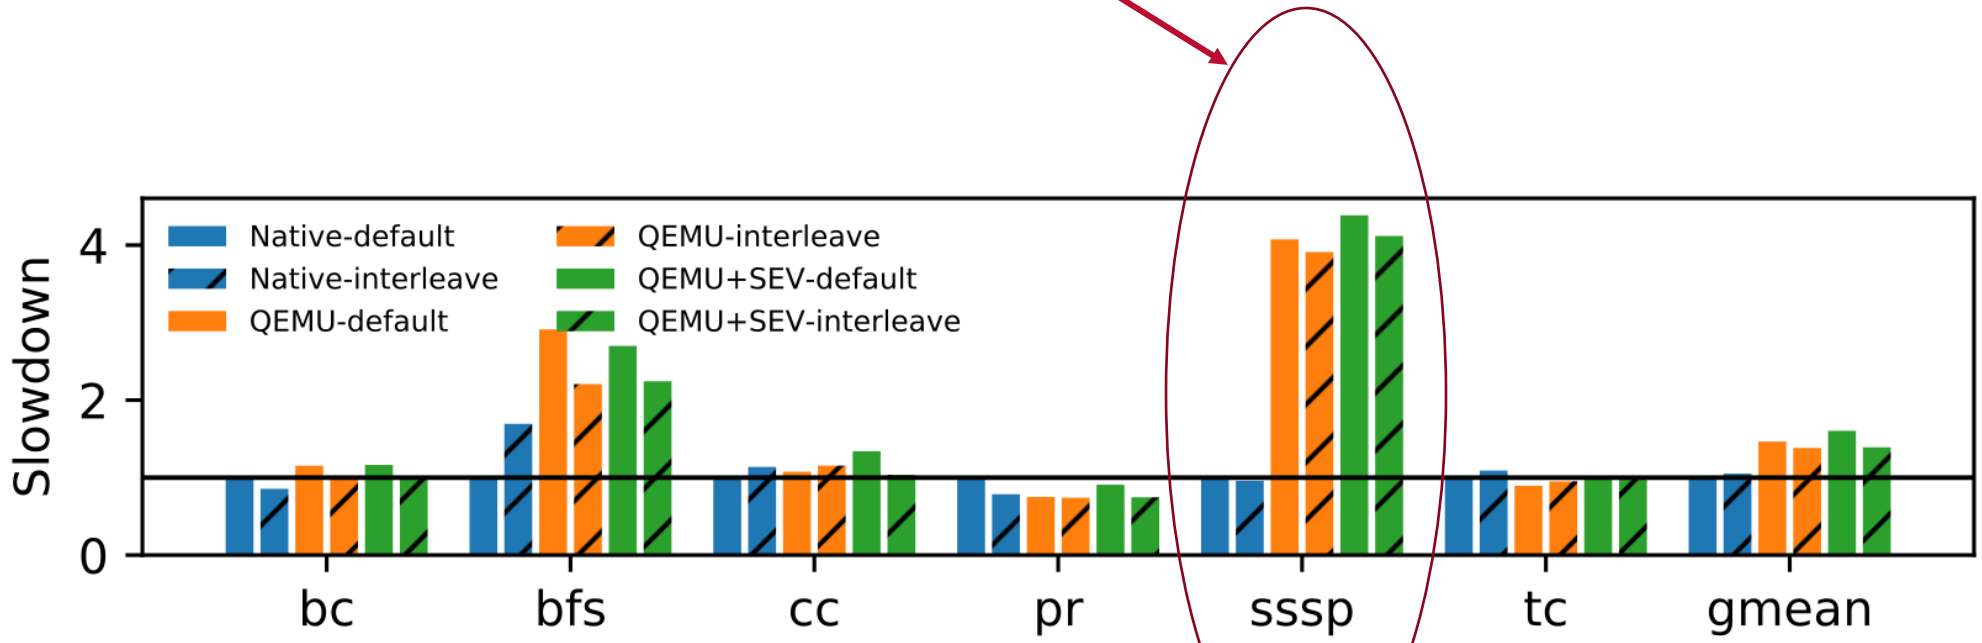

GAPBS on AMD Rome (128 Core Machine)

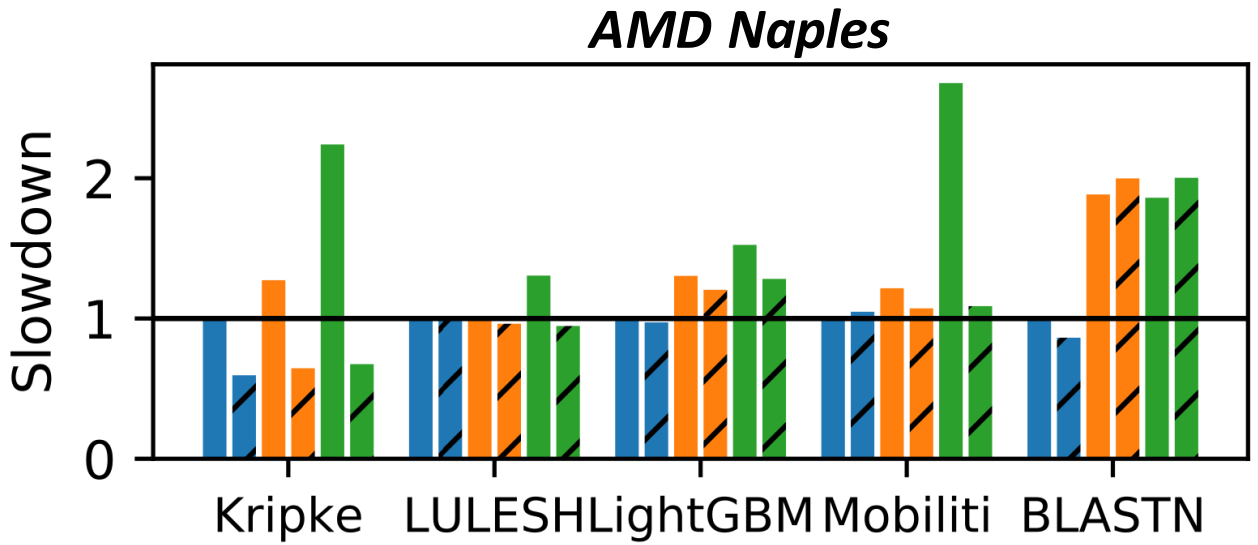
AMD Naples
Slowdown
0
1
2
Kripke
LULESH
LightGBM
Mobiliti
BLASTN

Uses a 245GB database, much larger than the memory size of the test platform

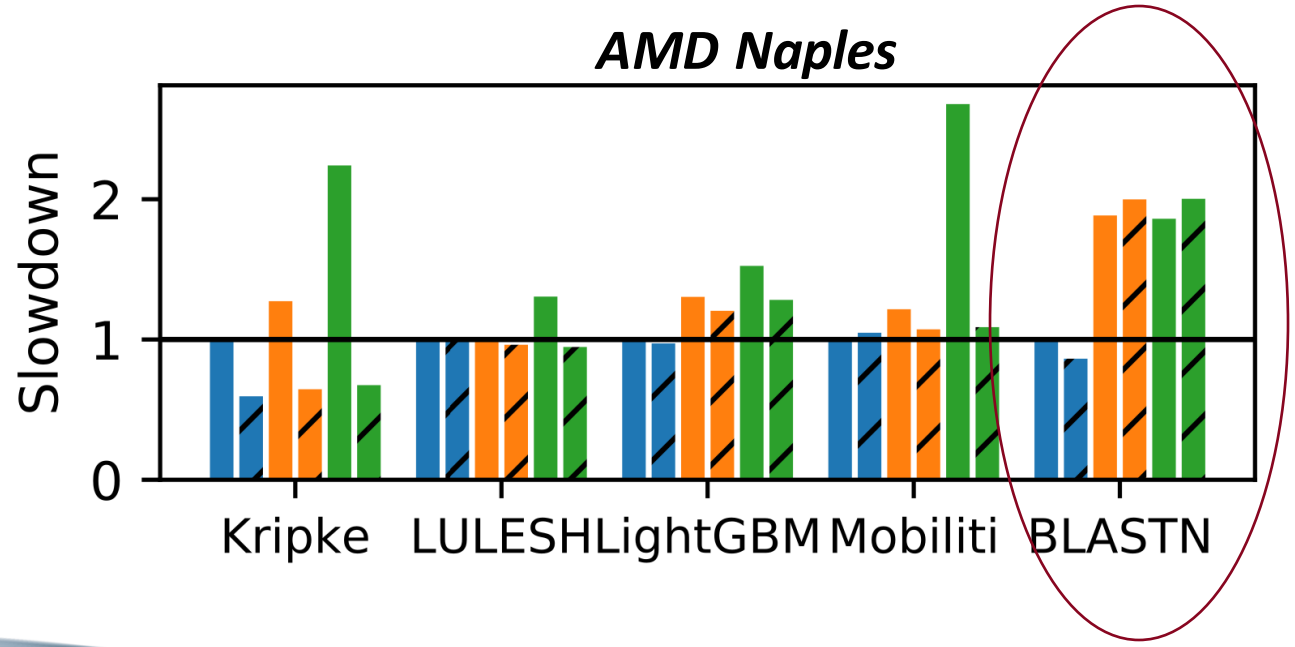

Finding 2

In some cases SEV overhead is mainly because of virtualization, which is a requirement of the SEV programming model

Finding 3

SEV initialization is slow and depends on the memory footprint of the VM

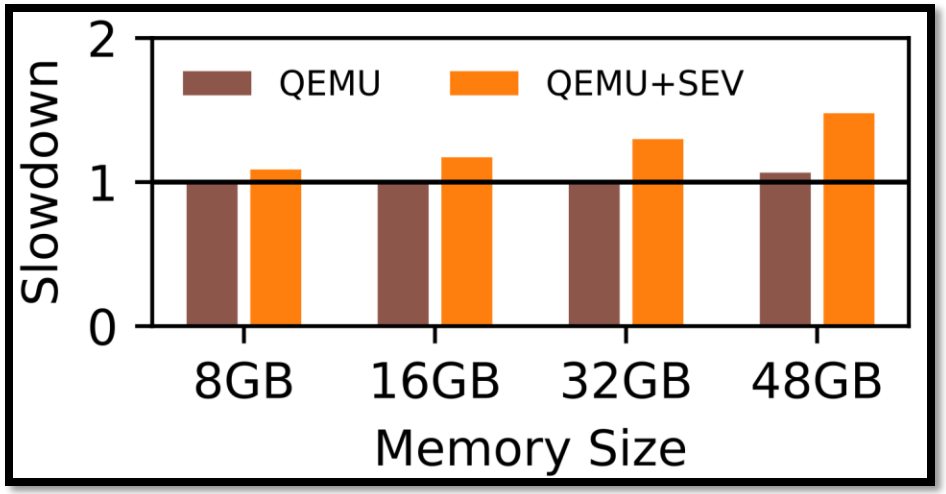

Outline

Why Secure High Performance Computing Matters?

Performance Analysis of General-Purpose TEEs for HPC

- What special configurations AMD SEV need?
- **Why Intel SGX is not an appropriate fit for HPC?**

Future Trends in TEEs

Performance Impact of SGX

Large slowdown (upto--126x) specially for graph workloads

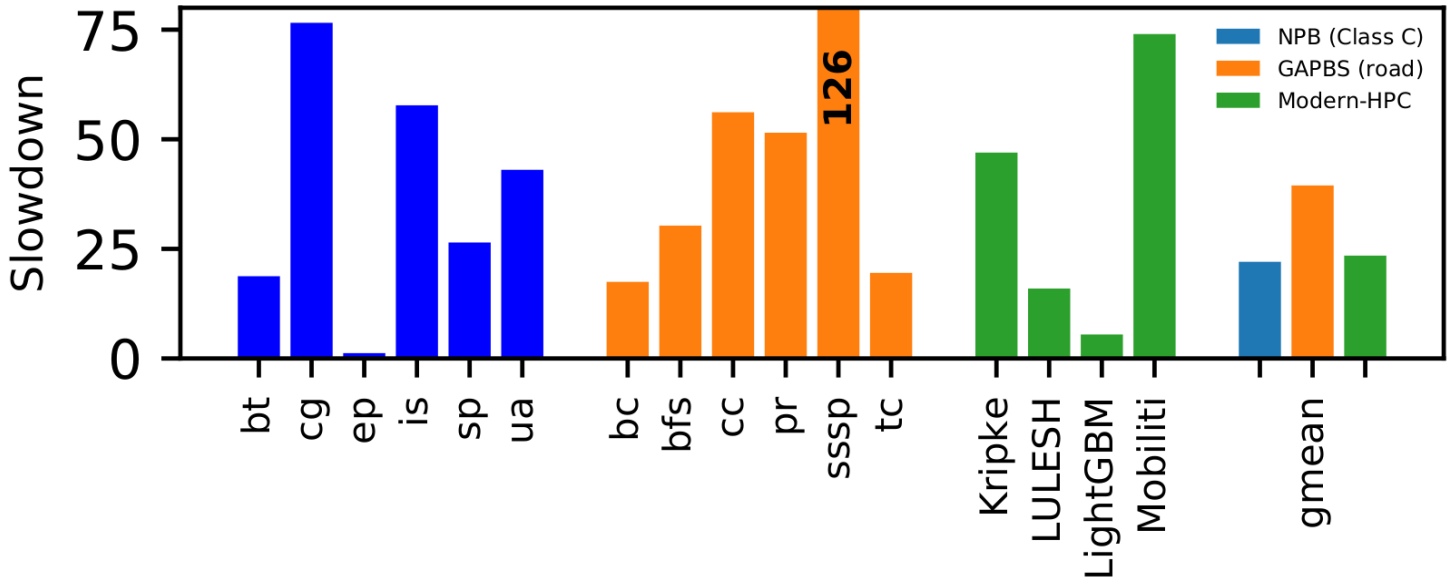

Enclave Page Cache (EPC) Faults

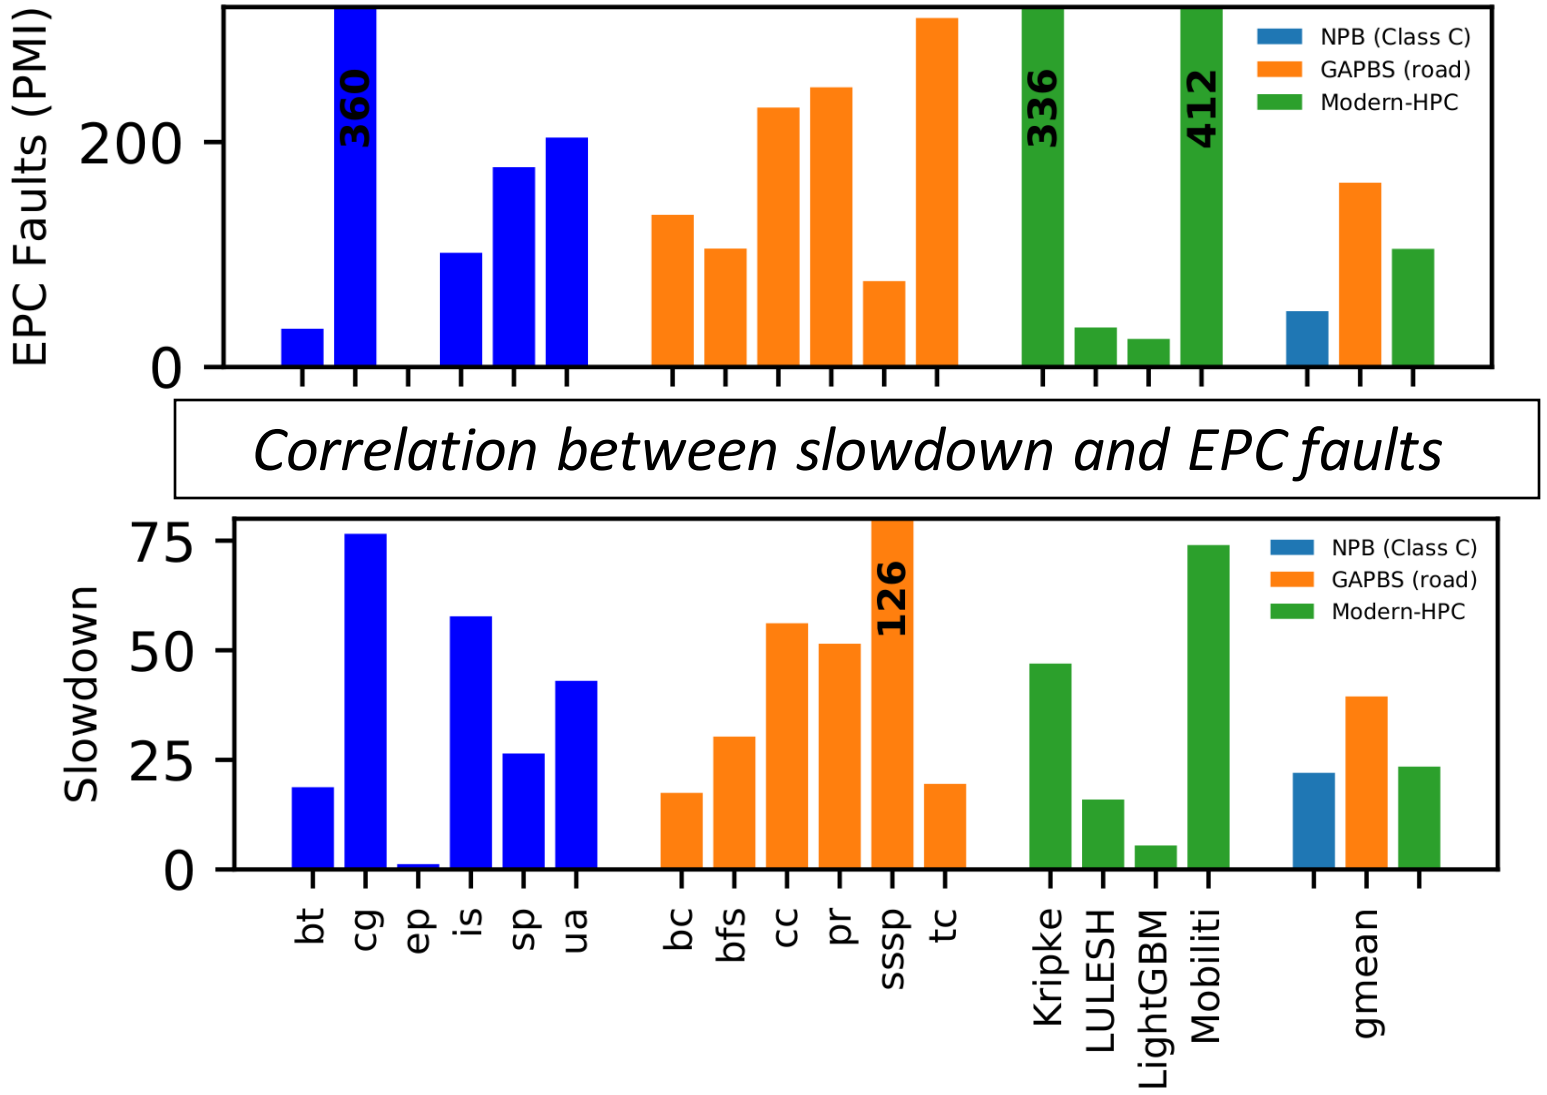

Impact of Increasing Execution Threads (under SGX)

Doesn't scale well because of high resident memory

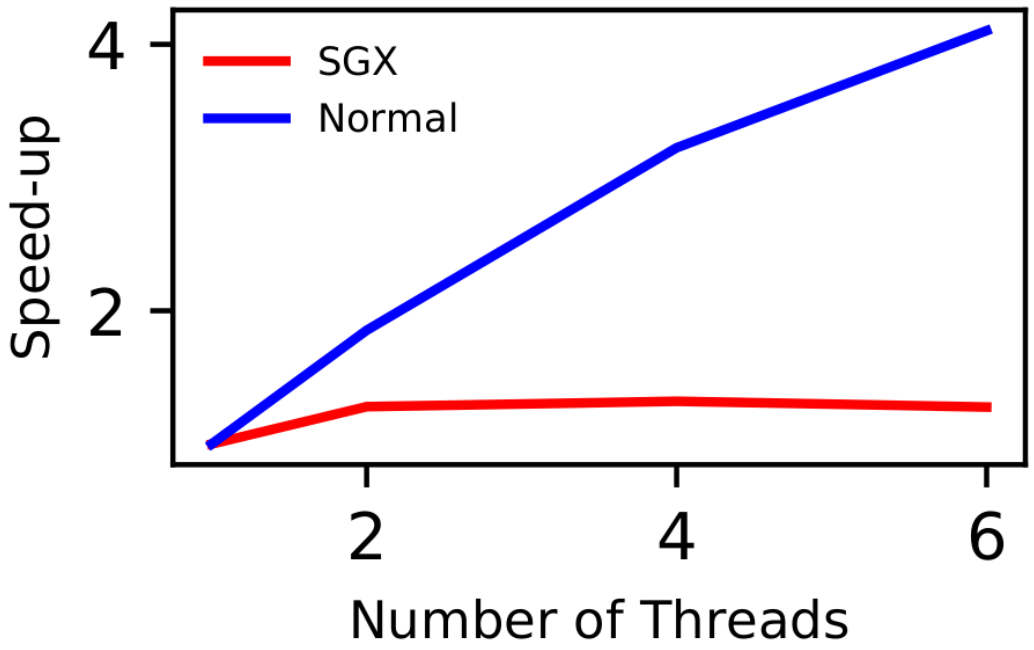

Scales normally under SGX and has small resident mem

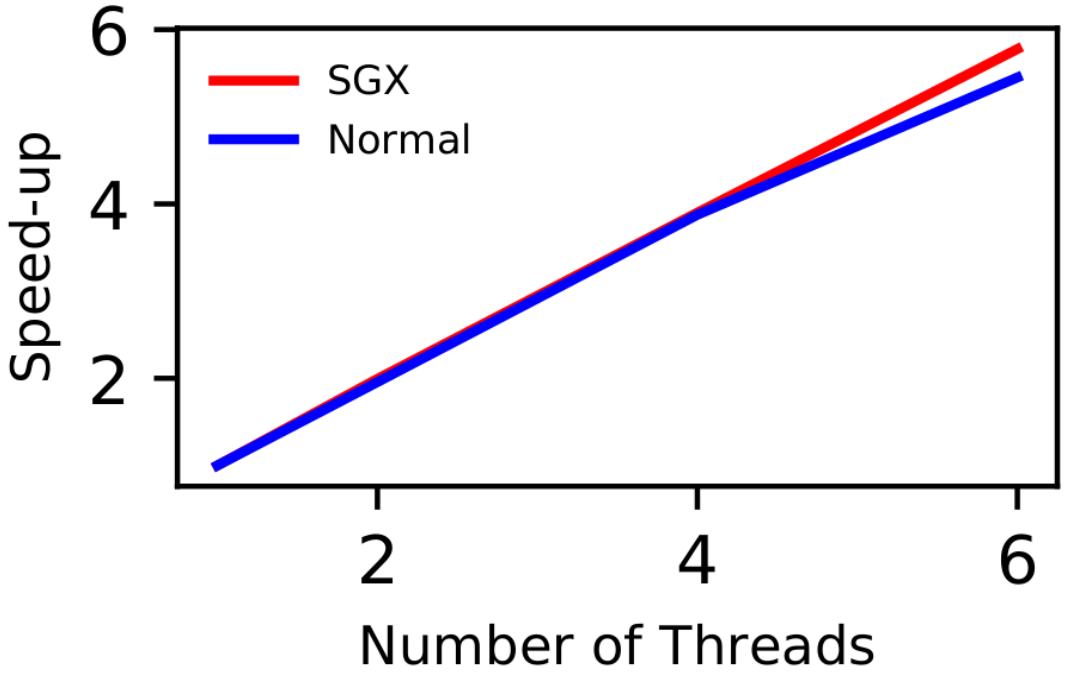

Finding 4

SGX is inappropriate for unmodified HPC workloads because of its limited secure memory, poor thread scalability and its unsuitable programming model for HPC

UCDAVIS

Outline

Why Secure High Performance Computing Matters?

Performance Analysis of General-Purpose TEEs for HPC

- What special configurations AMD SEV need?
- Why Intel SGX is not an appropriate fit for HPC?

Future Trends in TEEs

Future Trends in TEEs

Enhancements to SEV in the form of SEV-ES and SEV-SNP

Most TEEs seem to be following SEV like design

Intel has introduced

- MKTME (multi key total memory encryption)
- TDX (Trust Domain Extension)

ARM v9's Confidential Compute Architecture has introduced ARM Realms

Summary

Can TEEs enable secure scientific computing?

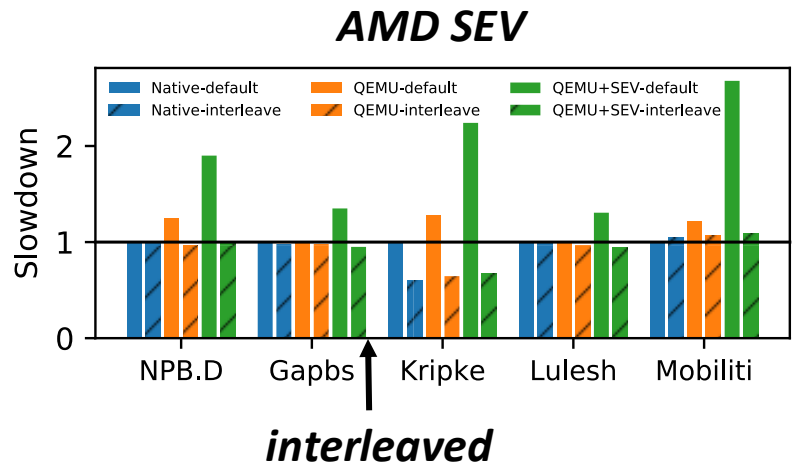

AMD SEV shows little performance degradation if used with **interleaved** NUMA allocation

- Irregular workloads can still show virtualization-based overheads

Intel SGX is inappropriate for scientific computing

- Incurs high performance overheads
- Programming model requires application modifications

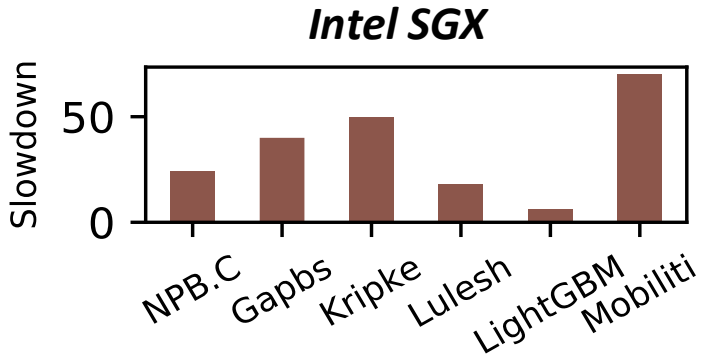

Performance Analysis of Scientific Computing Workloads on General Purpose TEEs

Ayaz Akram (yazakram@ucdavis.edu), Anna Giannakou, Venkatesh Akella, Jason Lowe-Power, Sean Peisert