Using Trusted Execution Environments On High-Performance Computing Platforms

Ayaz Akram, Anna Giannakou, Venkatesh Akella, Jason Lowe-Power, and Sean Peisert

yazakram@ucdavis.edu, agiannakou@lbl.gov, akella@ucdavis.edu,
jlowepower@ucdavis.edu, sppeisert@lbl.gov

University of California, Davis and Lawrence Berkeley National Lab

Open Source Enclave Workshop (OSEW), 2019

Using Trusted Execution Environments On High-Performance Computing Platforms

Ayaz Akram, Anna Giannakou,
Venkatesh Akella, Jason Lowe-Power, Sean Peisert

Secure High-Performance Computing

How to compute with large sensitive data?

- Biomedical data
- Proprietary data

Secure from both external and internal threats

- Integrity or confidentiality or both

High-Performance Computing Workloads

Common characteristics

- Large data sets (10s–100s GB per node)
- Limited user interaction (batch)
- Often highly multithreaded

Dedicated (super computers) or shared (cloud) nodes

Diverse compute, memory, and security requirements

UCDAVIS

We Analyze Two TEEs

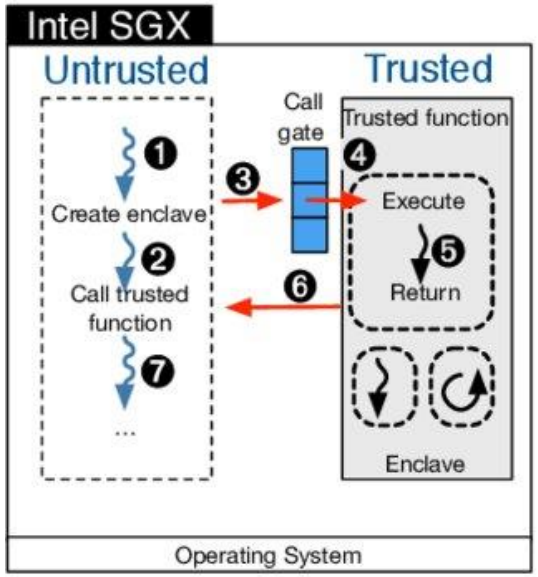

[1]

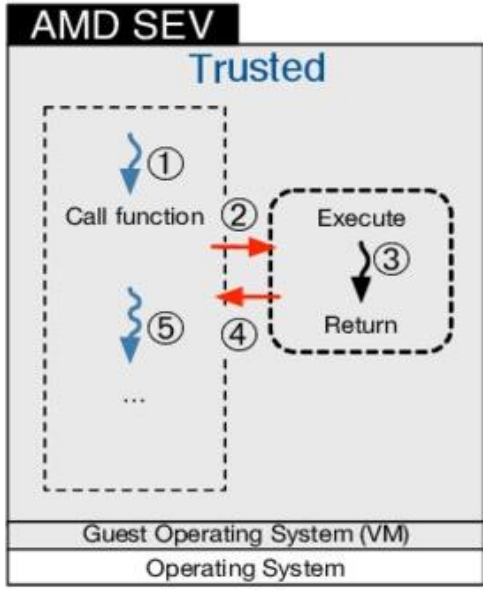

[1]

Technology	Ensures 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.

Methodology

- Benchmarks used: NAS parallel benchmarks, LightGBM and GAPBS
- Platforms used: Intel Core i7-8700 (12 threads/socket) for SGX and AMD EPYC 7451 (dual socket with 48 threads/socket) for SEV study
- Use of SCONE (SGX) and Kata (SEV) containers
- Measured slowdown of the used workloads under secure execution on both platforms
- Relate the slowdown to other collected metrics

Performance Impact of SGX

High slowdown, especially for graph workloads

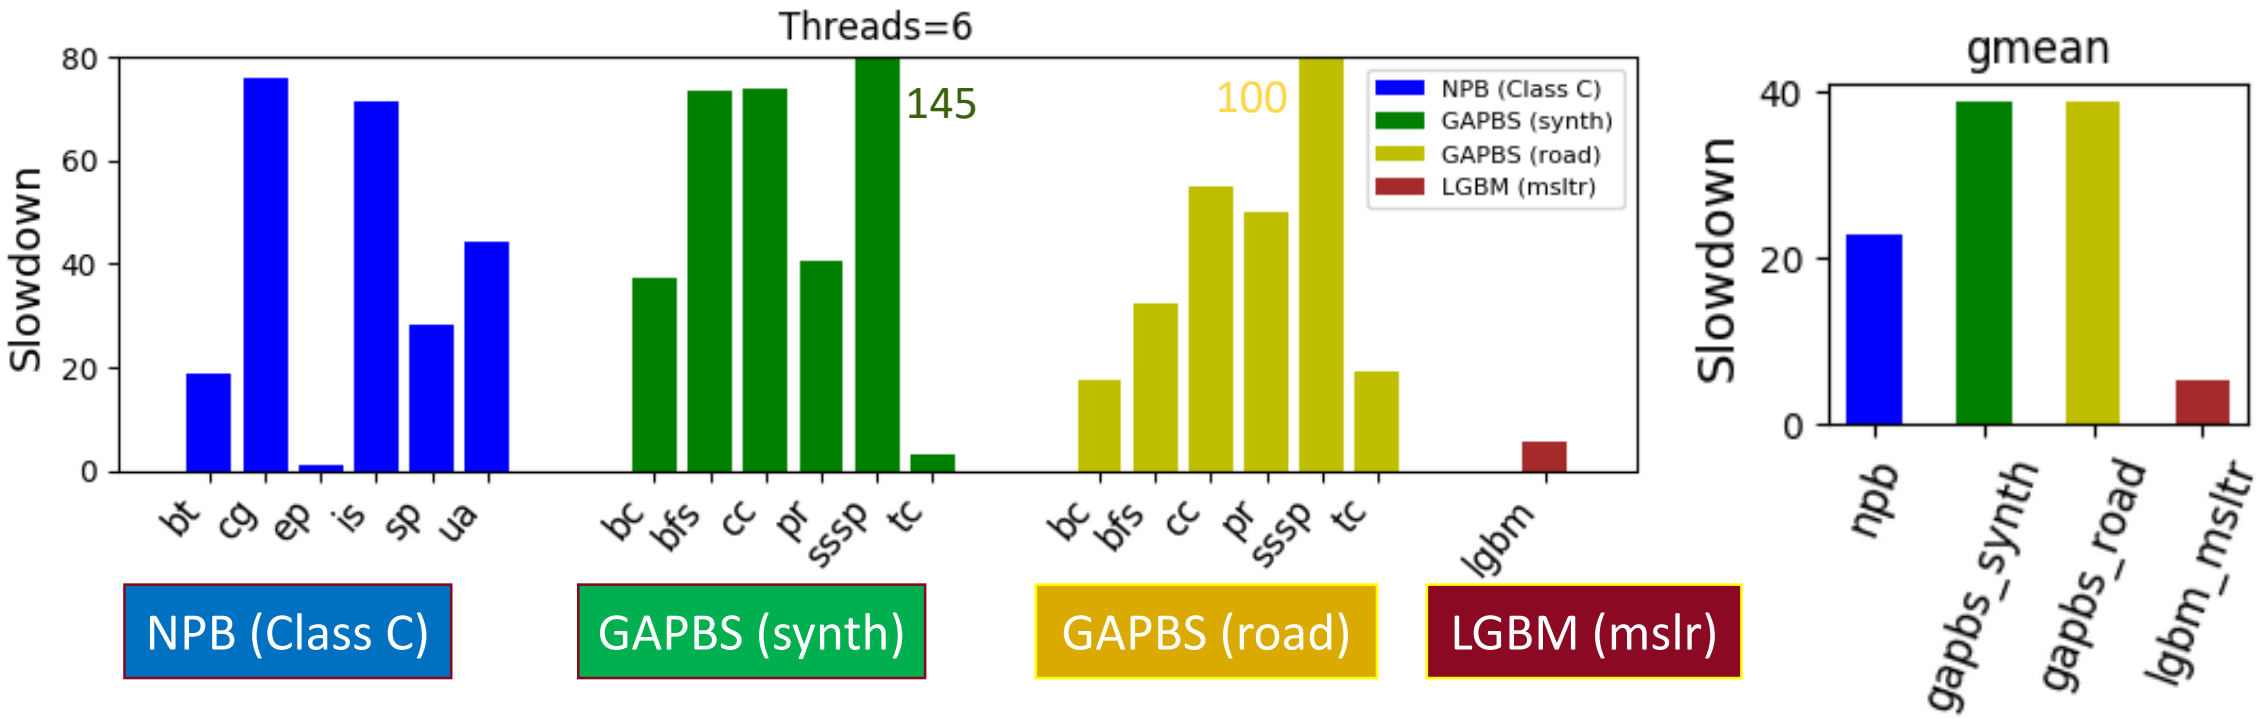

Enclave Page Cache (EPC) Faults

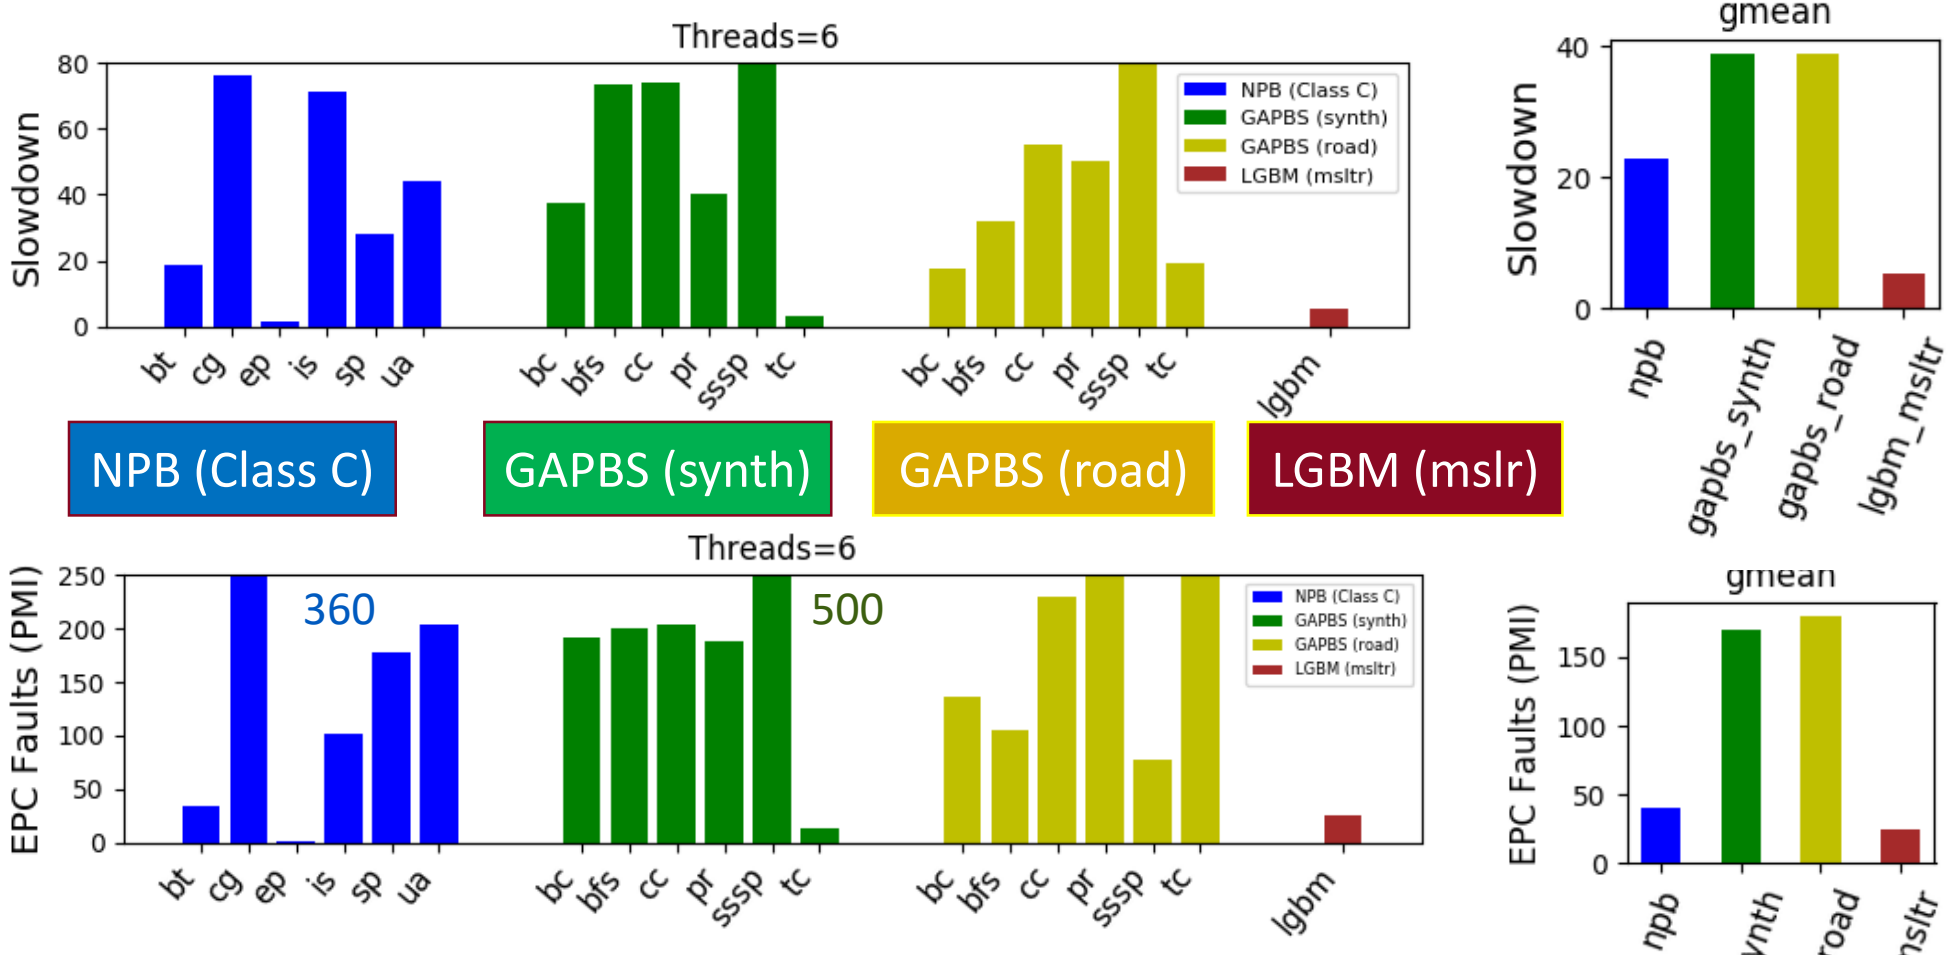

Enclave Page Cache (EPC) Faults

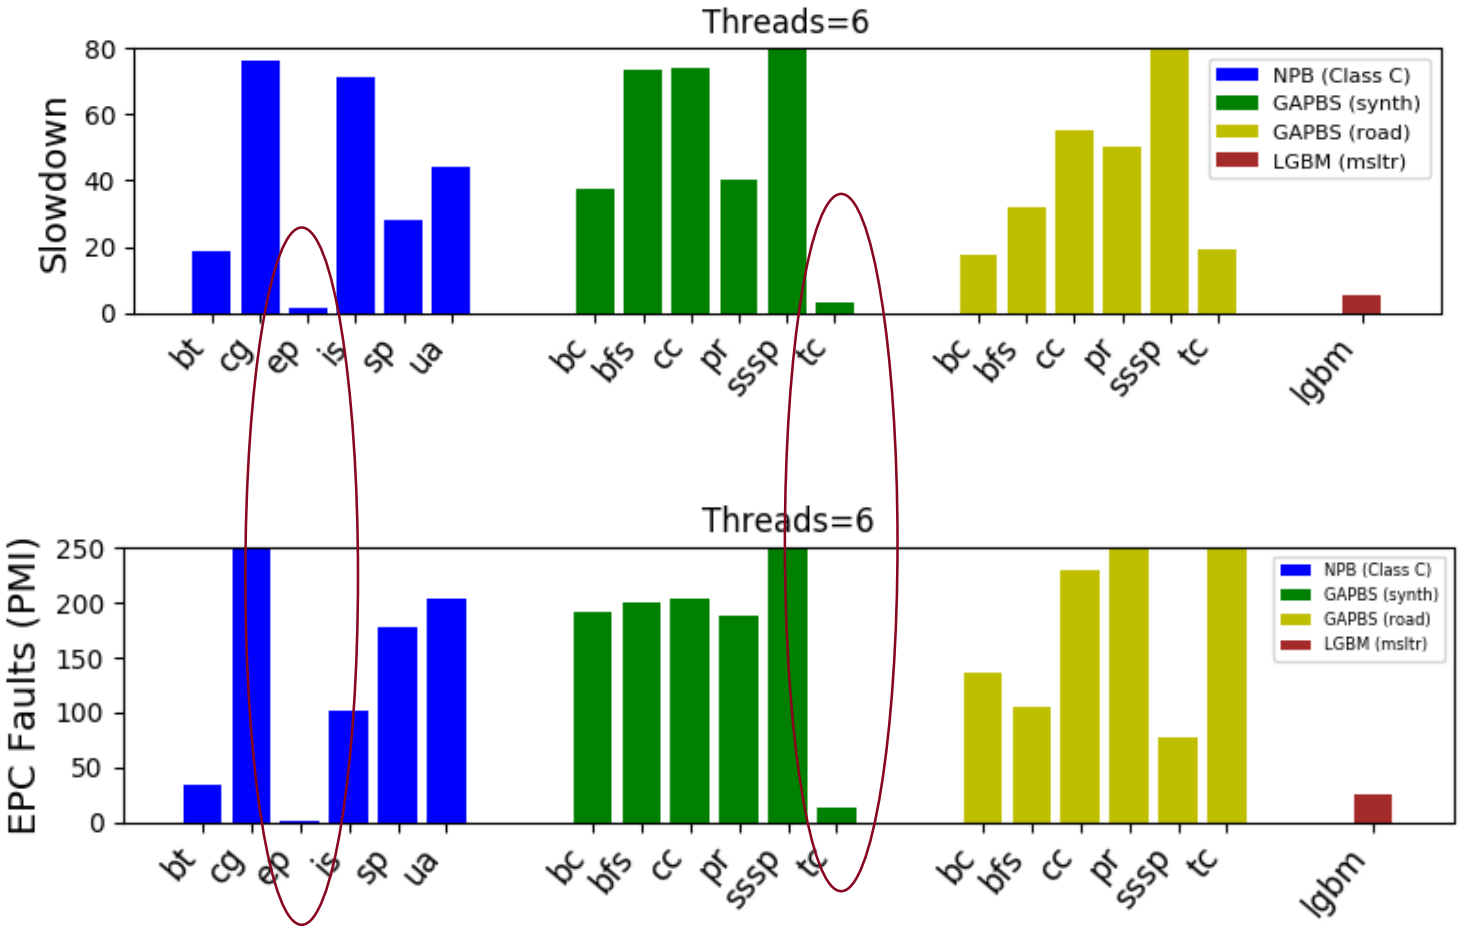

All the benchmarks have large resident memory except ep & tc_synth

Impact of Increasing Execution Threads (under SGX)

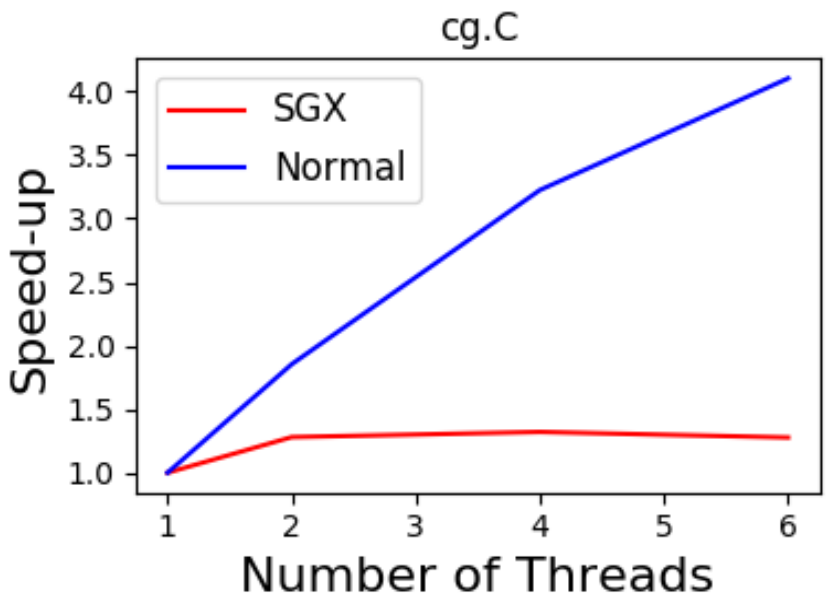

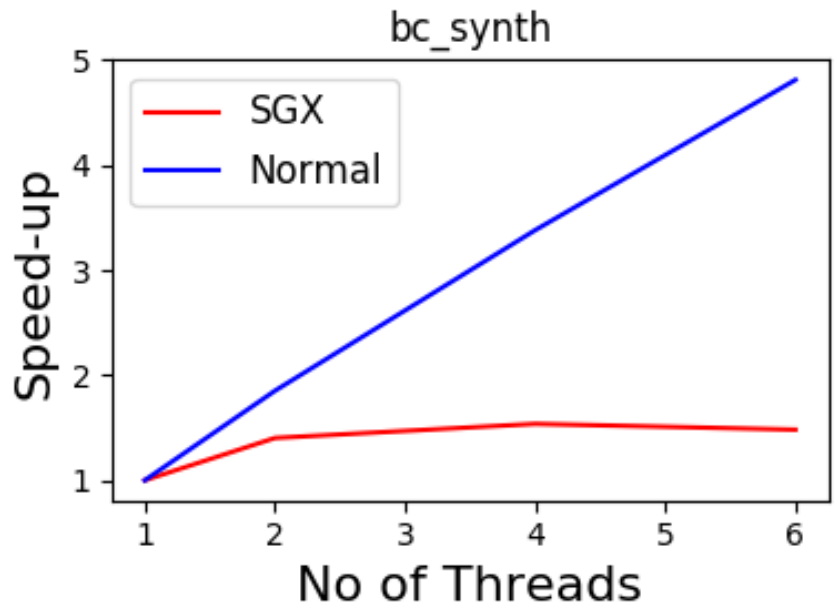

Don't scale well, as they have large resident memory

Impact of Increasing Execution Threads (under SGX)

Scales normally under SGX and has a small memory footprint

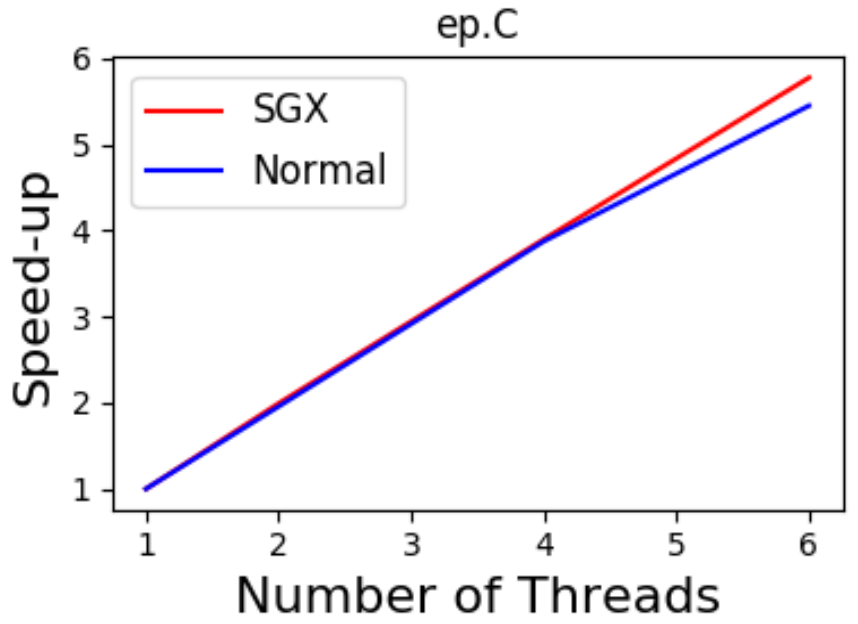

Performance Impact of SEV

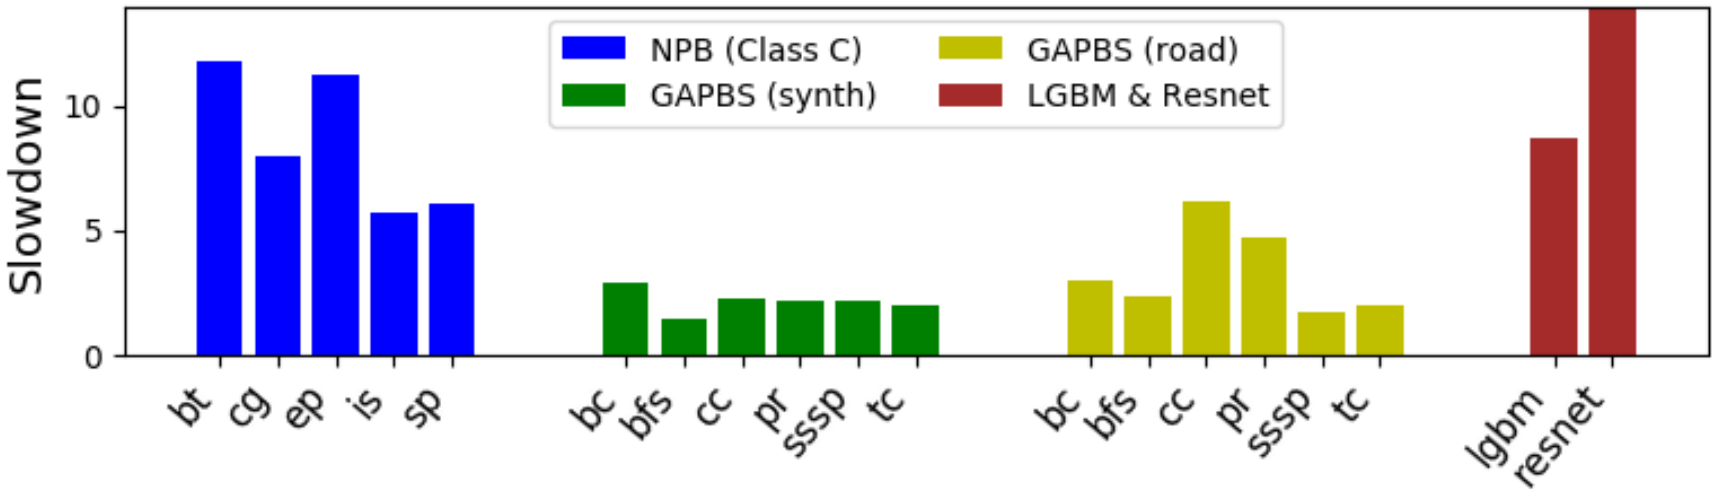

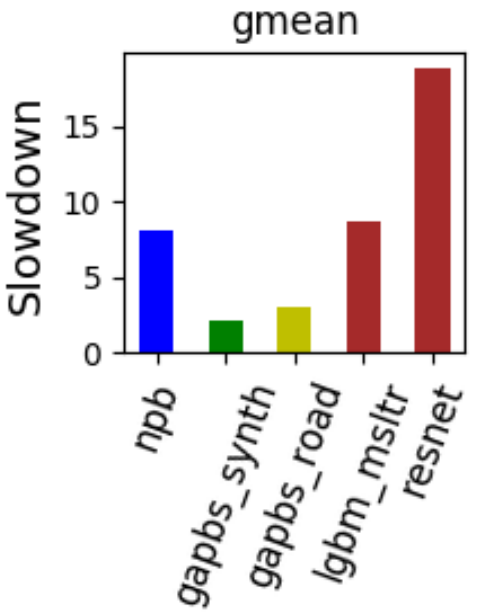

NPB (Class C)

GAPBS (synth)

GAPBS (road)

LGBM (mslr)

Performance Impact of SEV

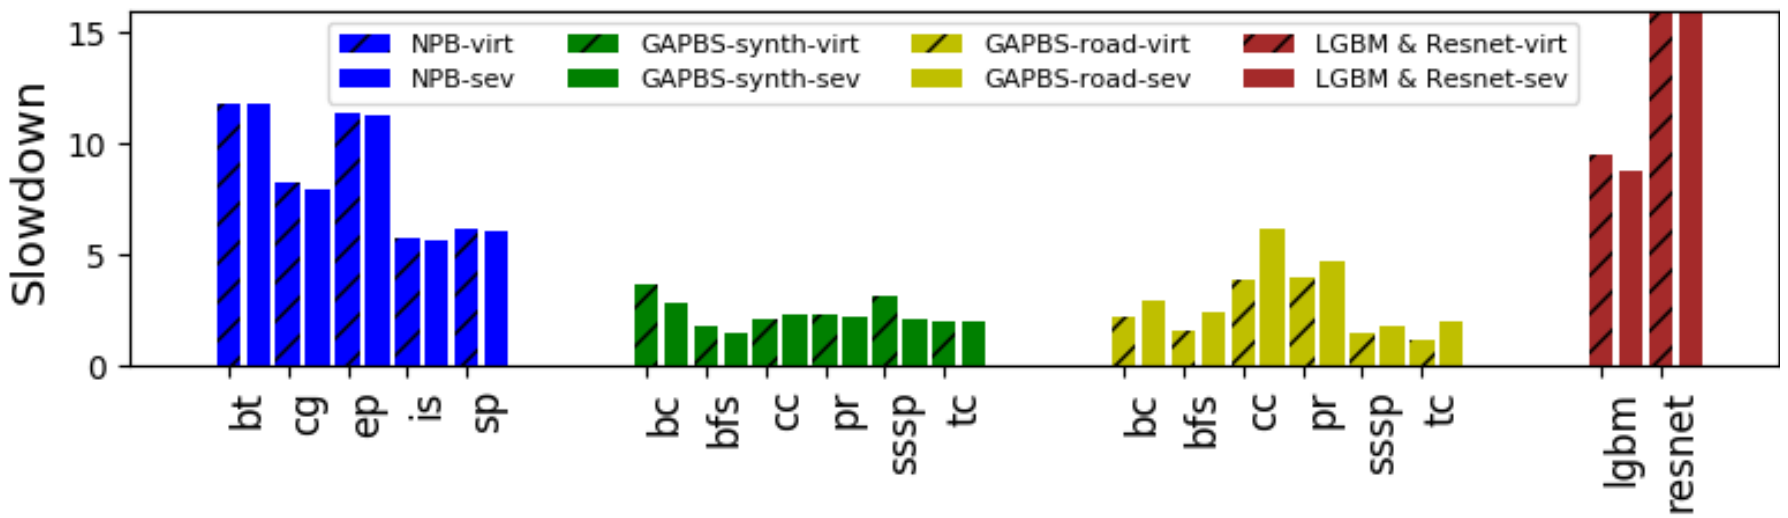

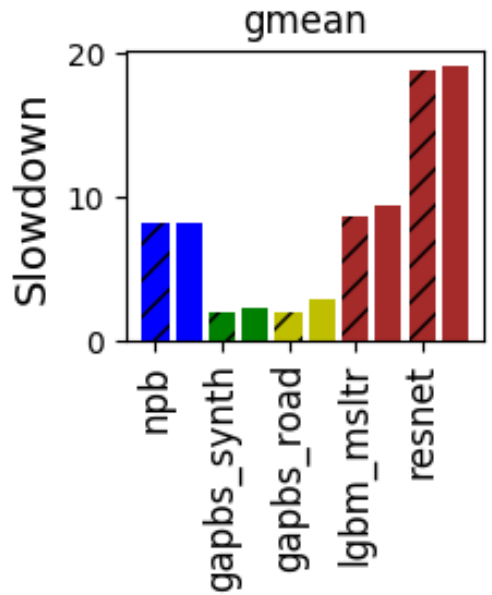

Virtualization appears to be the biggest reason of slowdown

Preliminary Takeaways

Future TEEs should support HPC apps

Smaller slowdowns for SEV

Performance issues for SGX

- EPC faults
- Multiple execution threads

Dynamic choice of threat model

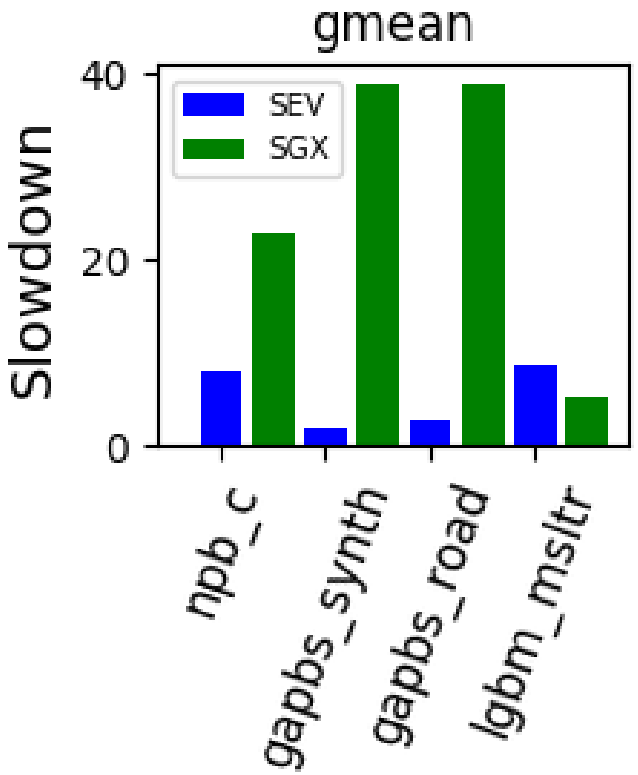

SEV and SGX slowdowns