ETH zürich

Composite Enclaves: Towards Disaggregated Trusted Execution

Moritz Schneider, Aritra Dhar, Ivan Puddu, Kari Kostiainen, Srdjan Capkun

CHES 2022

Modern Computing

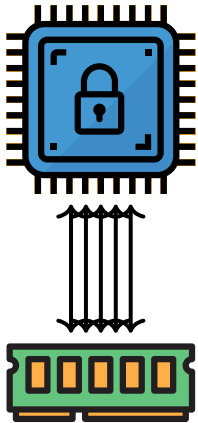

Modern Computing

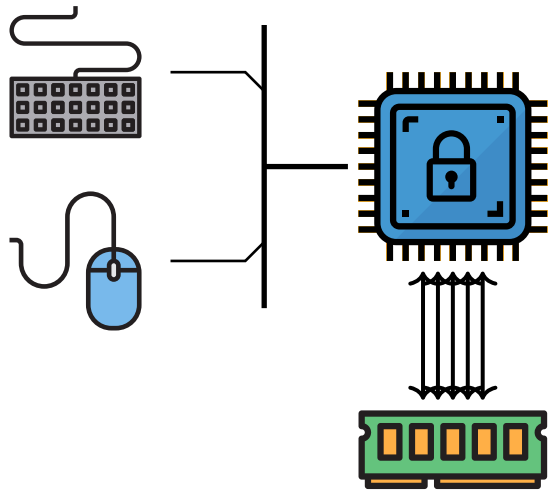

Modern Computing

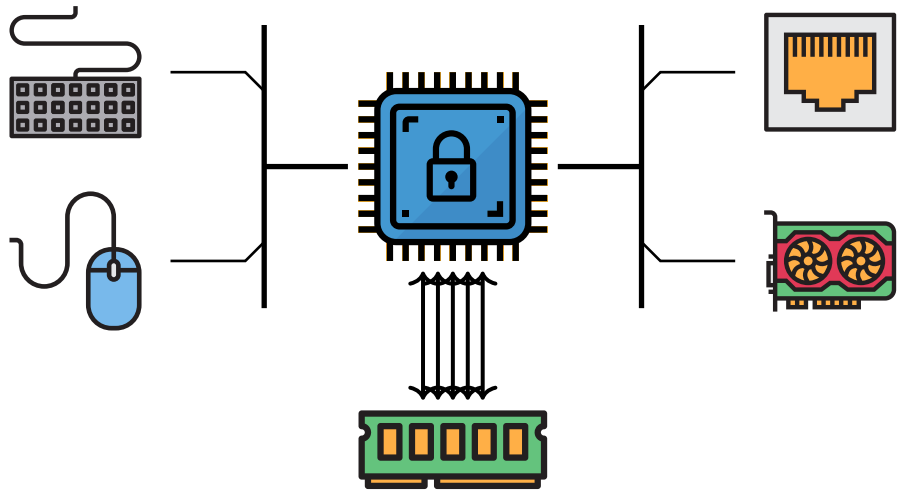

Modern Computing

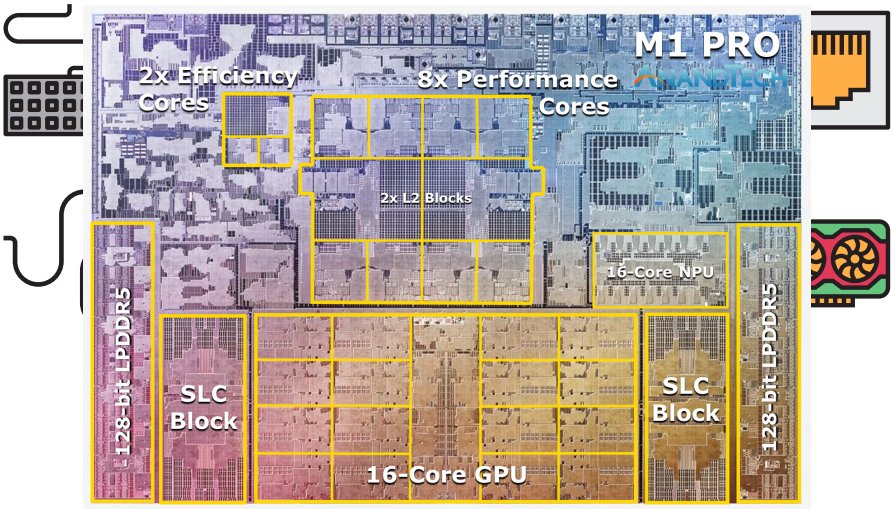

source: AnandTech

Trusted Computing?

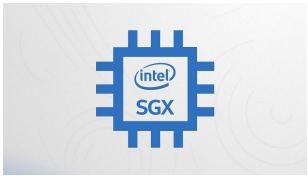

AMD SEV

ARM TrustZone

Keystone

Trusted Computing?

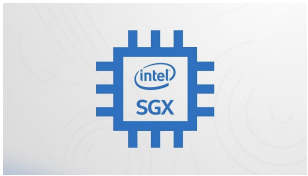

AMD SEV

ARM TrustZone

Keystone

NVIDIA CONFIDENTIAL COMPUTING

Secure data and AI models in use.

Trusted Computing?

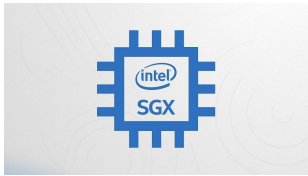

AMD SEV

ARM TrustZone

Keystone

NVIDIA CONFIDENTIAL COMPUTING

Secure data and AI models in use.

Modern vs Trusted Computing

- Specialized Computing devices
  - GPUs
  - Accelerators
  - Etc.
- CPU is mere coordinator

- Mostly limited to CPU
- No accelerator support*

Not Just Computing

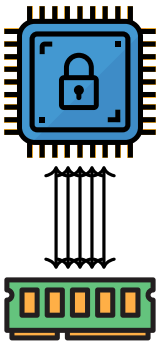

Not Just Computing

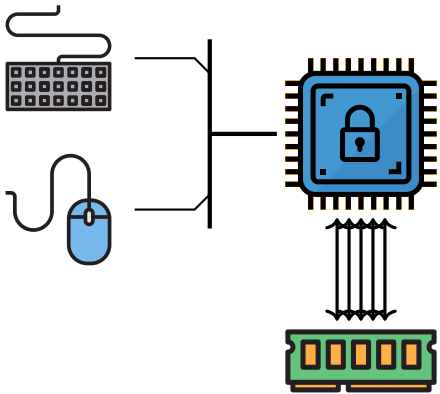

Not Just Computing

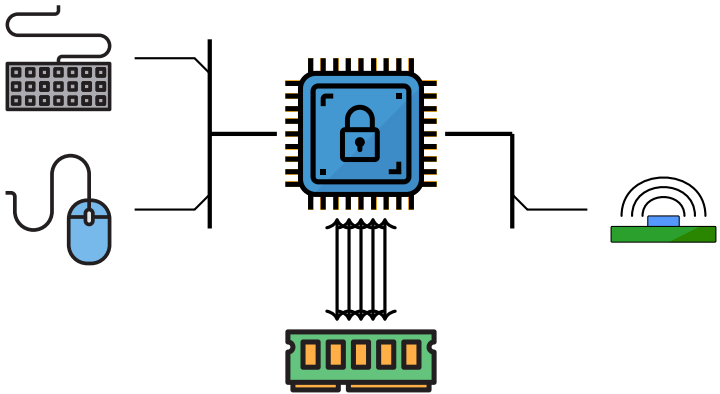

Not Just Computing

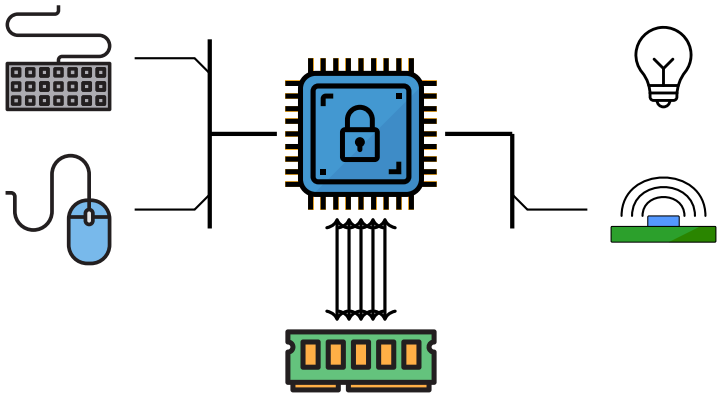

Not Just Computing

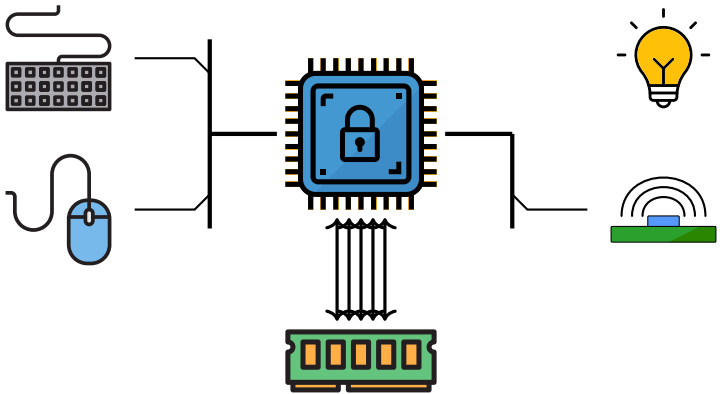

TEEs on Specialized Hardware

TEEs on Specialized Hardware

- Stavros Volos, Kapil Vaswani, and Rodrigo Bruno (2018). “Graviton: Trusted Execution Environments on GPUs”. In: OSDI 18. USENIX Association
- Insu Jang et al. (2019). “Heterogeneous isolated execution for commodity gpus”. In: ASPLOS

TEEs on Specialized Hardware

- Stavros Volos, Kapil Vaswani, and Rodrigo Bruno (2018). “Graviton: Trusted Execution Environments on GPUs”. In: OSDI 18. USENIX Association
- Insu Jang et al. (2019). “Heterogeneous isolated execution for commodity gpus”. In: ASPLOS
- Sérgio Pereira et al. (2021). “Towards a Trusted Execution Environment via Reconfigurable FPGA”. In: arXiv

TEEs on Specialized Hardware

- Stavros Volos, Kapil Vaswani, and Rodrigo Bruno (2018). “Graviton: Trusted Execution Environments on GPUs”. In: OSDI 18. USENIX Association
- Insu Jang et al. (2019). “Heterogeneous isolated execution for commodity gpus”. In: ASPLOS
- Sérgio Pereira et al. (2021). “Towards a Trusted Execution Environment via Reconfigurable FPGA”. In: arXiv
- Many more...

TEEs on Specialized Hardware

- Stavros Volos, Kapil Vaswani, and Rodrigo Bruno (2018). “Graviton: Trusted Execution Environments on GPUs”. In: OSDI 18. USENIX Association
- Insu Jang et al. (2019). “Heterogeneous isolated execution for commodity gpus”. In: ASPLOS
- Sérgio Pereira et al. (2021). “Towards a Trusted Execution Environment via Reconfigurable FPGA”. In: arXiv
- Many more...
- **Nvidia Confidential Computing**

Attacker Model

Goal: SGX Attacker Model

Attacker Model

Goal: SGX Attacker Model

- Local phyiscal adversary
- Controls entire SW stack

Attacker Model

Goal: SGX Attacker Model

- Local phyiscal adversary
- Controls entire SW stack

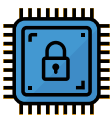

Attacker Model

Goal: SGX Attacker Model

- Local phyiscal adversary
- Controls entire SW stack

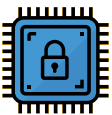

Attacker Model

Goal: SGX Attacker Model

- Local phyiscal adversary
- Controls entire SW stack

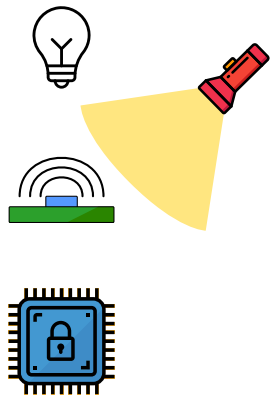

Attacker Model

Goal: SGX Attacker Model

- ~~Local phyiscal adversary~~
- Controls entire SW stack

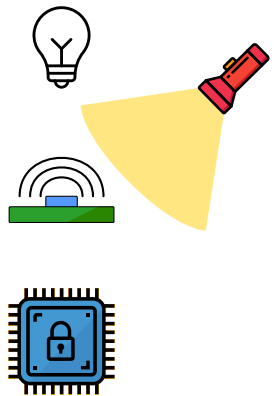

Interaction with Peripherals

Interaction with Peripherals

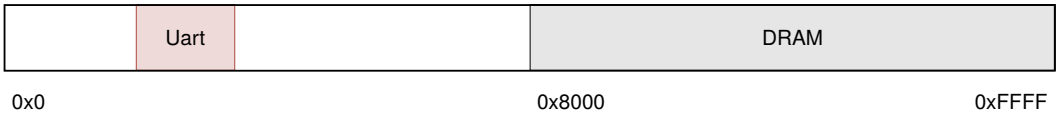

Interaction with Peripherals

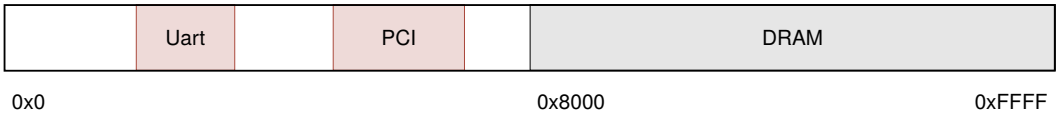

Interaction with Peripherals

Interaction with Peripherals

Interaction with Peripherals

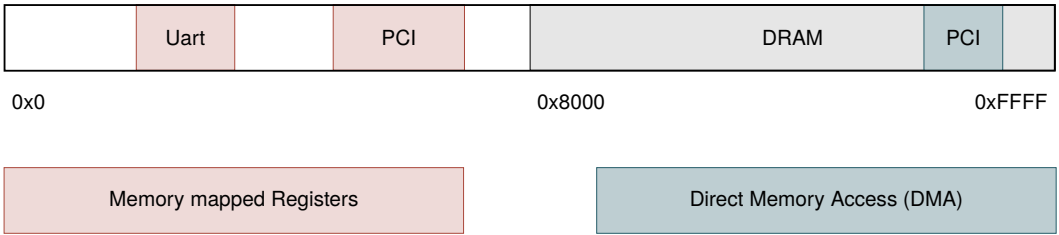

Interaction with Peripherals

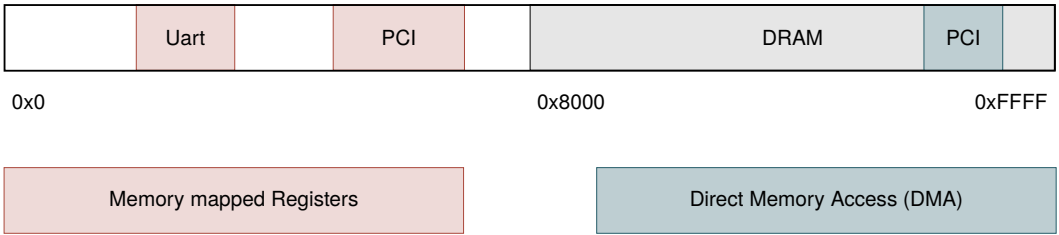

Memory Isolation Mechanisms can be used

Our Approach

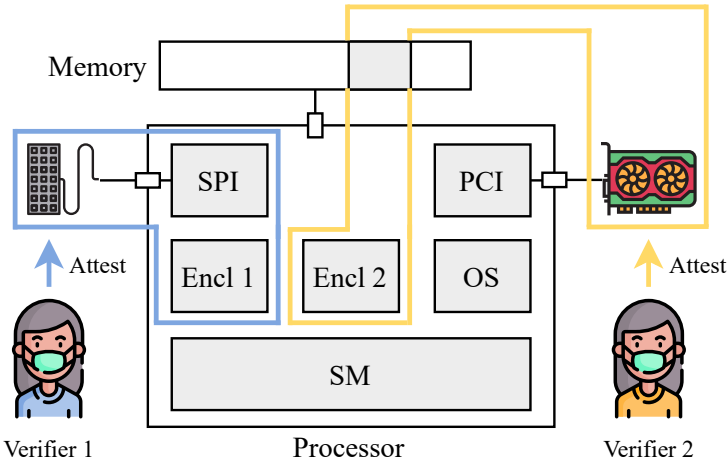

Main Problems

Main Problems

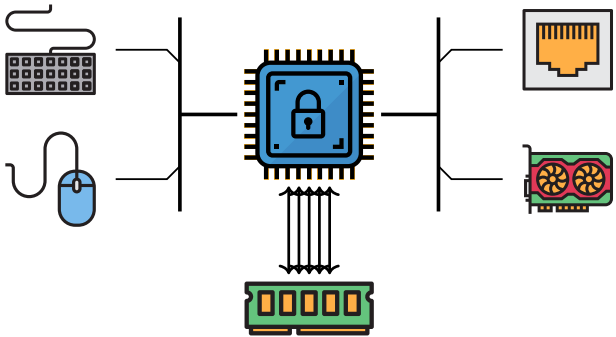

Main Problems

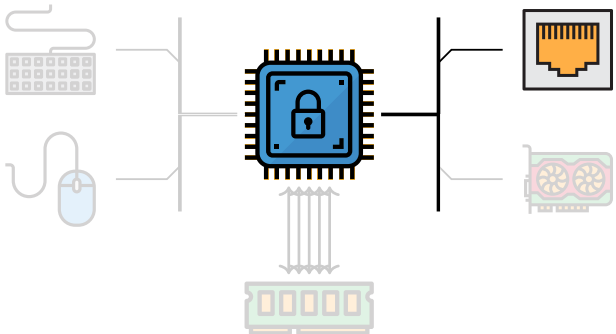

Main Problems

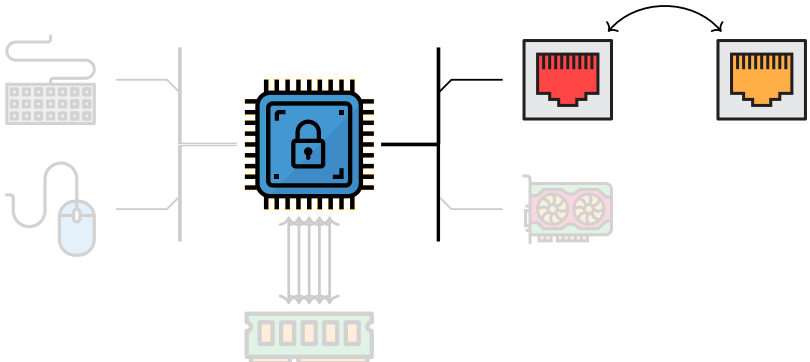

Our Approach

- MMIO statically assigned
- DMA dynamically created by TCB
- Explicit `connect` and `disconnect`
- Based on Keystone (Lee et al. 2020)

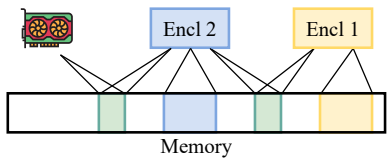

Implementation

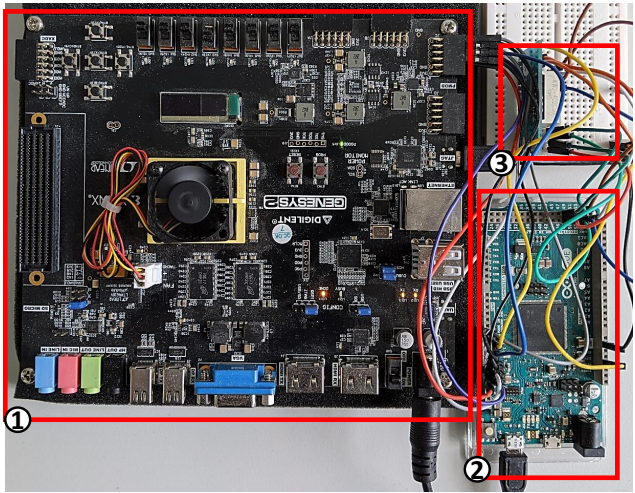

Context Switch Overhead

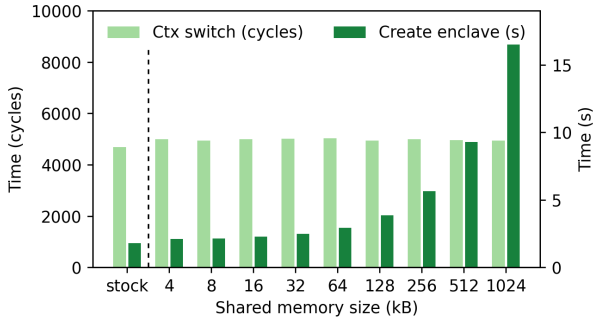

Accelerators and GPUs

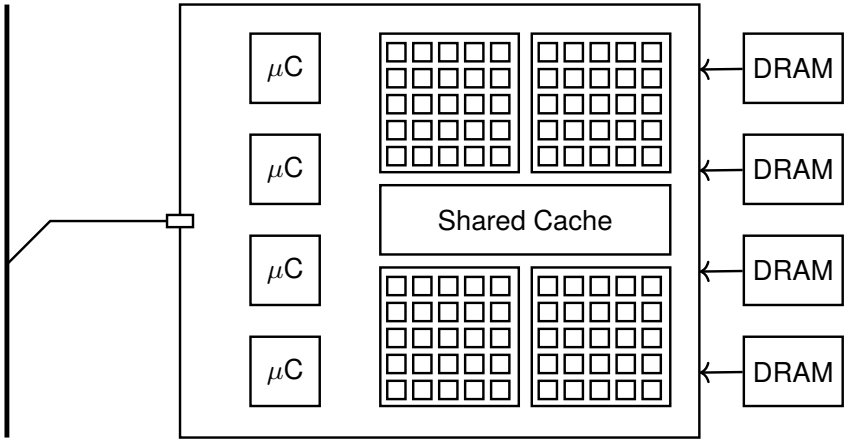

Accelerators and GPUs

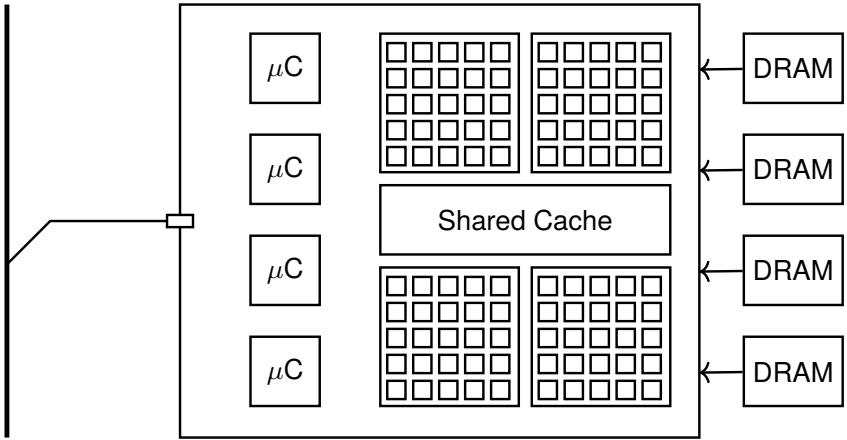

Accelerators and GPUs

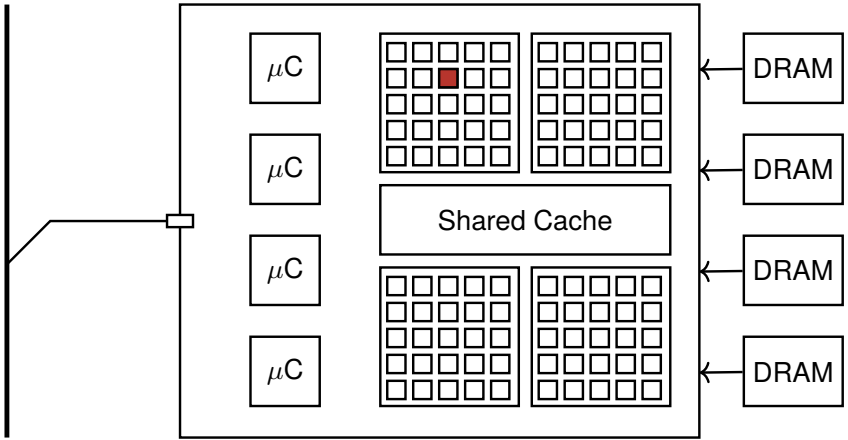

Compute Cluster

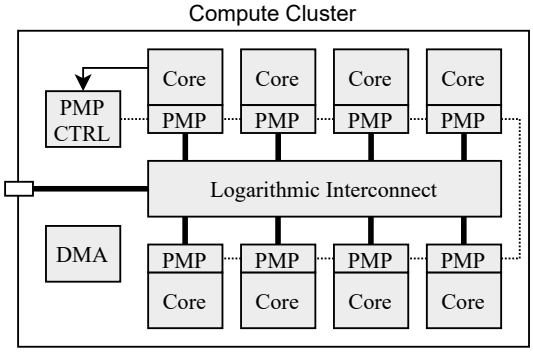

Hardware Overhead

Area [μm^2]	PMP Entries 0	PMP Entries 4*	Overhead
Core	5.7	6.7	15.5%
FPU	39.2	37.9	-3.3%
IPU	8.6	8.5	-1.4%
Total	53.5	53.2	-0.7%

Limitations

- What to do with drivers?
- Remote Attacker Model

Thanks

- Thanks to all the opensource projects that we used
  - Keystone
  - ETHZ RISC-V Cores

Backup/Nvidia Confidential Computing

Isolation on Nvidia GPUs

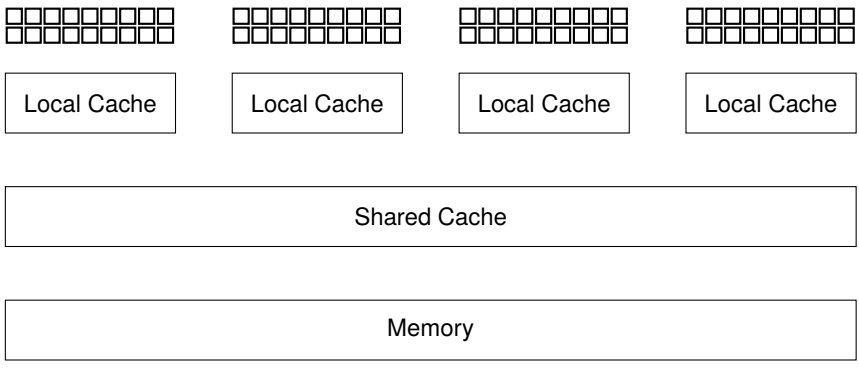

Isolation on Nvidia GPUs

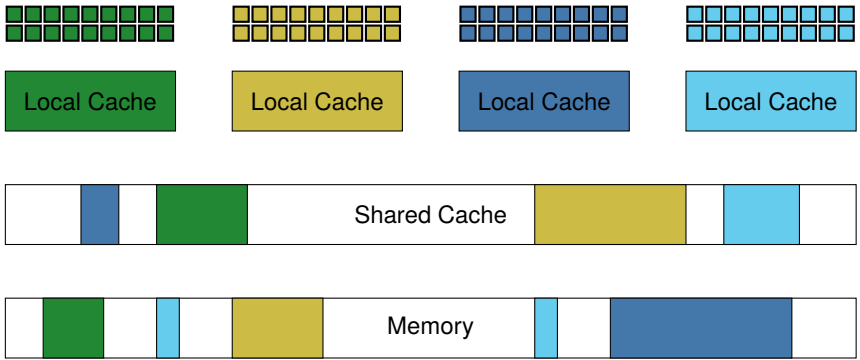

Isolation on Nvidia GPUs

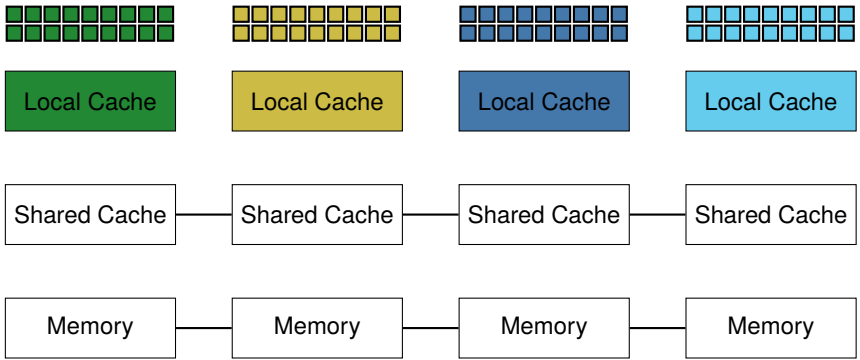

Isolation on Nvidia GPUs