EPFL

X-HEEP

EMBEDDED SYSTEMS LABORATORY

September 23, 2024

An Open-Source RISC-V-based GPGPU Accelerator for Machine Learning-based Edge Computing Applications

EPFL - Embedded Systems Laboratory (ESL)

simone.machetti@epfl.ch

Motivation

Analyzing the state-of-the art, we realized the need for a GPU…

Motivation

Analyzing the state-of-the art, we realized the need for a GPU…

Open-source

Motivation

Analyzing the state-of-the art, we realized the need for a GPU…

- **Open-source**
- **Natively Configurable**

Motivation

Analyzing the state-of-the art, we realized the need for a GPU…

- Open-source
- Natively Configurable
- RISC-V-based

Motivation

Analyzing the state-of-the art, we realized the need for a GPU…

- Open-source
- Natively Configurable
- RISC-V-based
- Fully Synthesizable

Motivation

Analyzing the state-of-the art, we realized the need for a GPU…

- Open-source
- Natively Configurable
- RISC-V-based
- Fully Synthesizable
- OpenCL Support

GPGPU Accelerator

GPGPU Accelerator

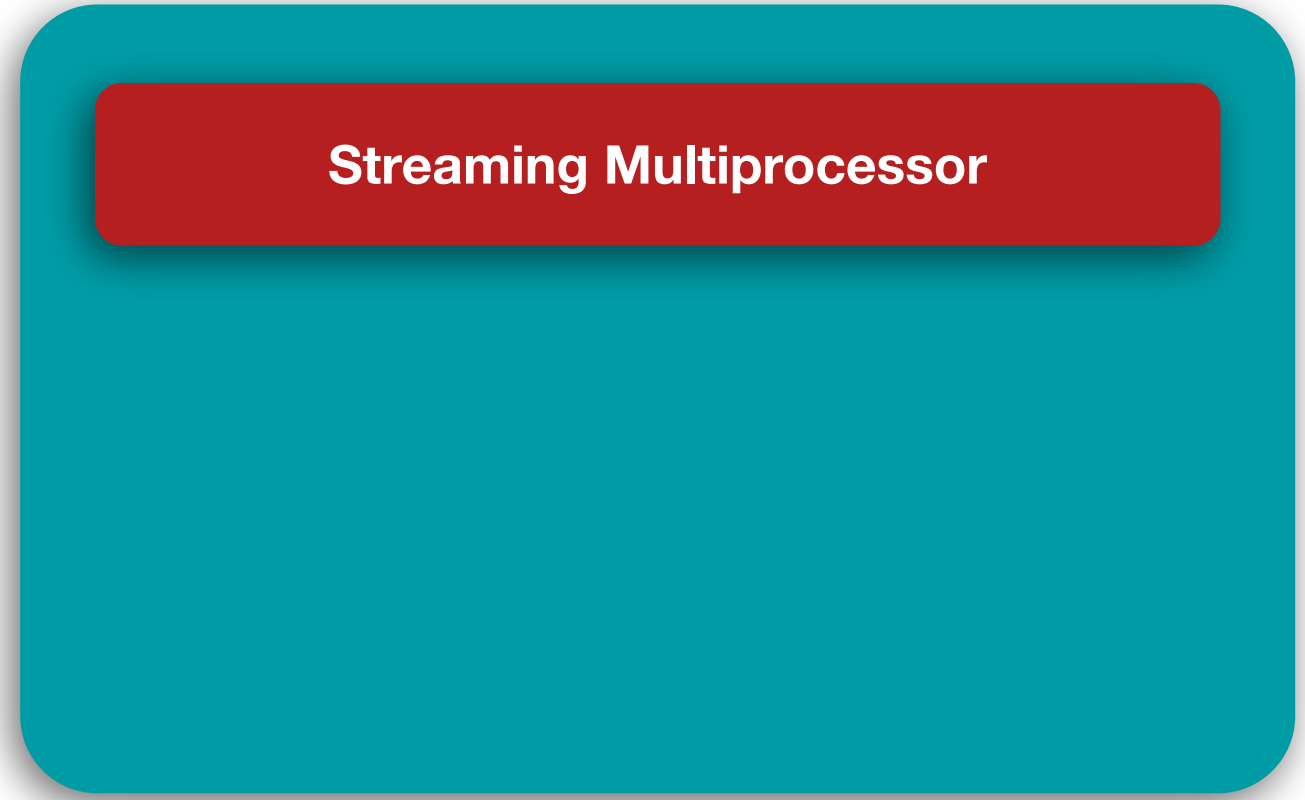

GPGPU Accelerator

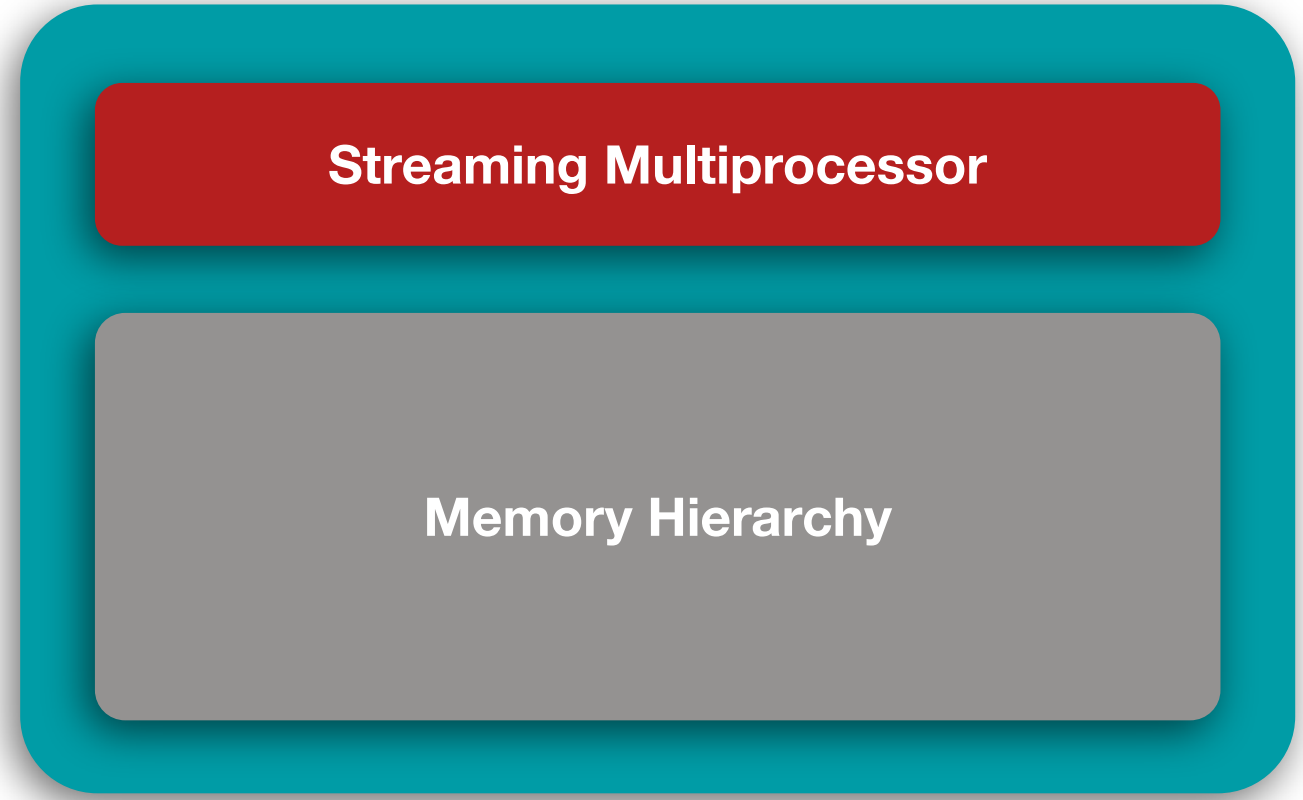

GPGPU Accelerator

Configurable

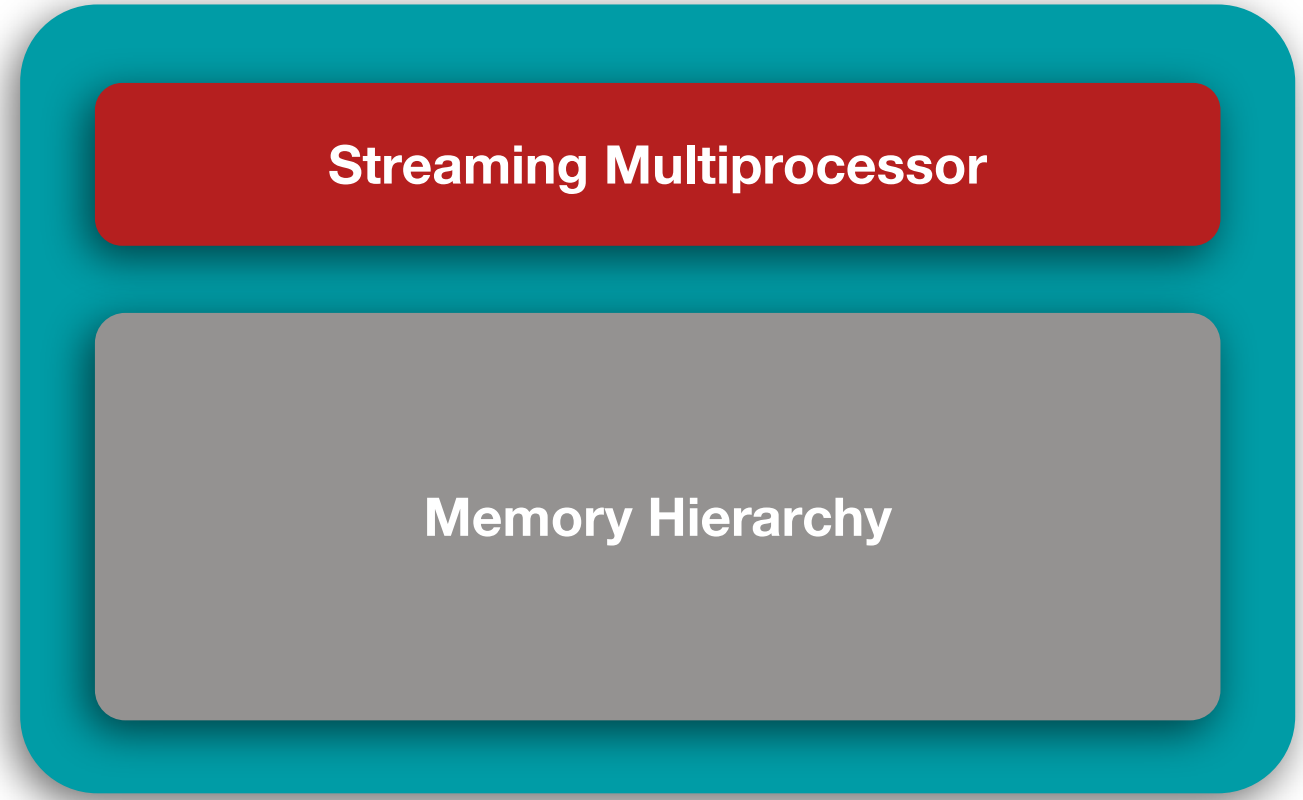

GPGPU Accelerator

Configurable

- Number of threads

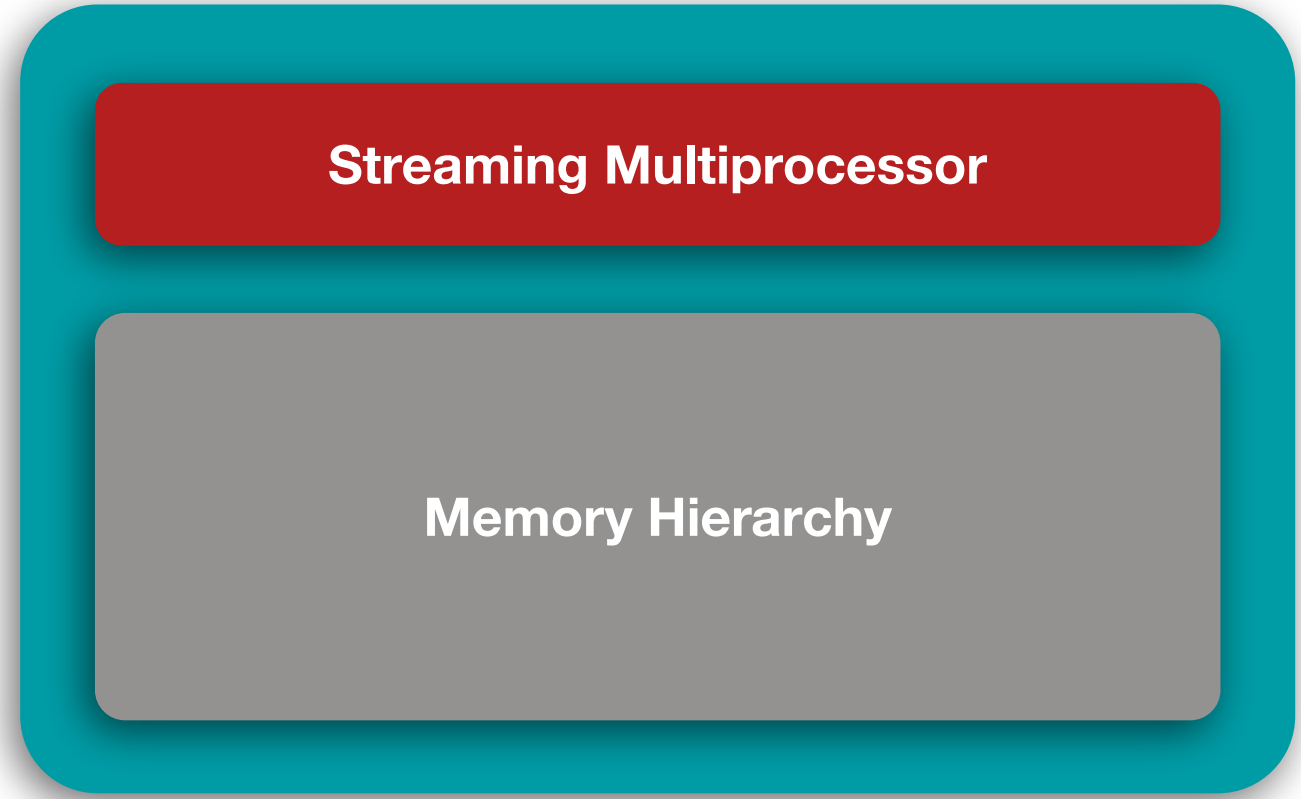

GPGPU Accelerator

Configurable

- Number of threads
- Number of warps

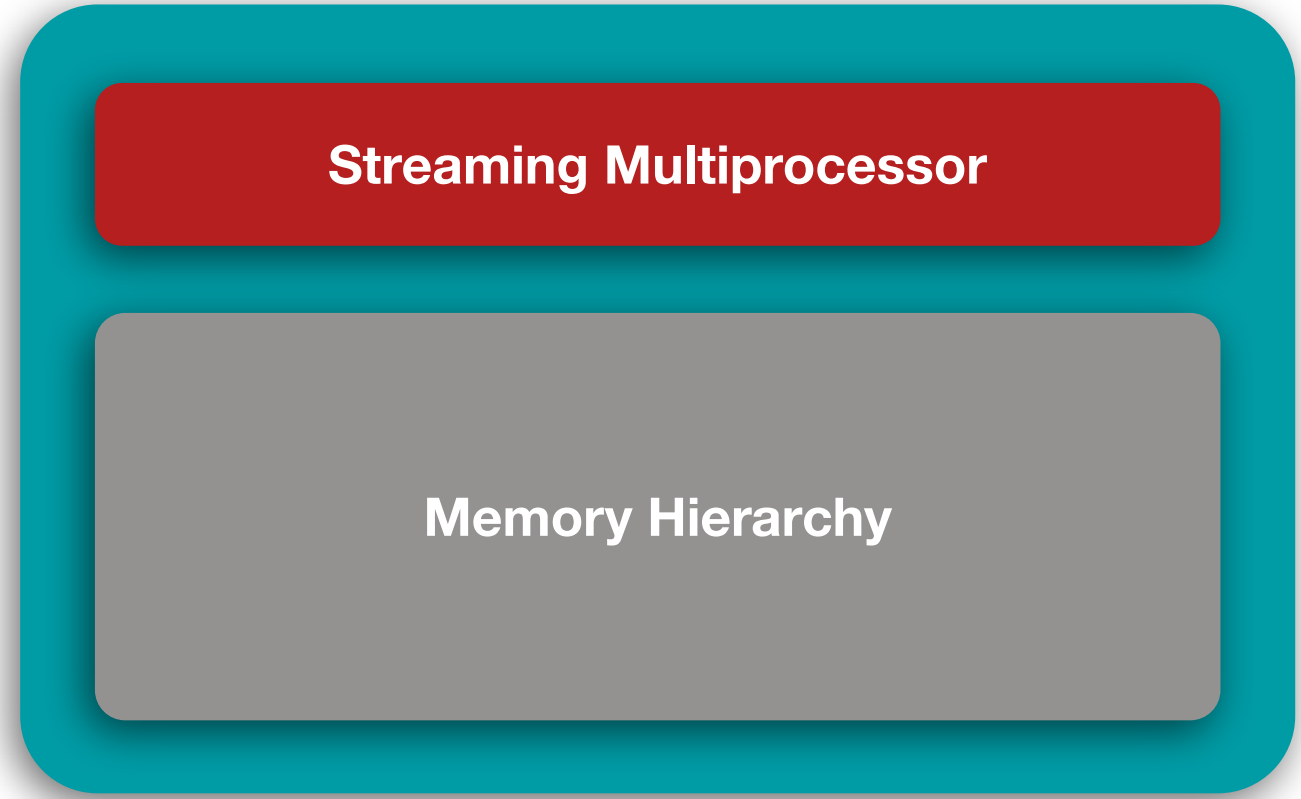

GPGPU Accelerator

Configurable

- Number of threads
- Number of warps
- Floating-point unit

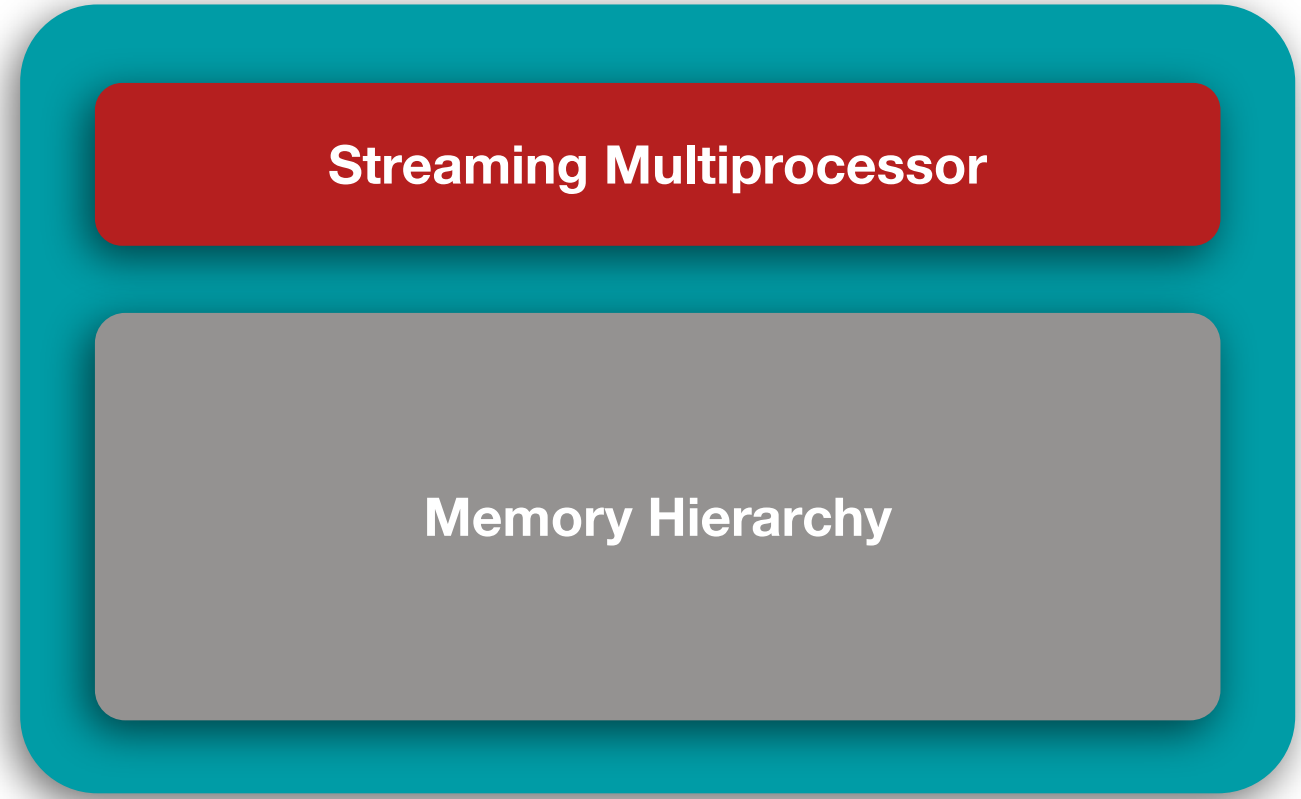

GPGPU Accelerator

Configurable

- Number of threads
- Number of warps
- Floating-point unit
- Memory hierarchy

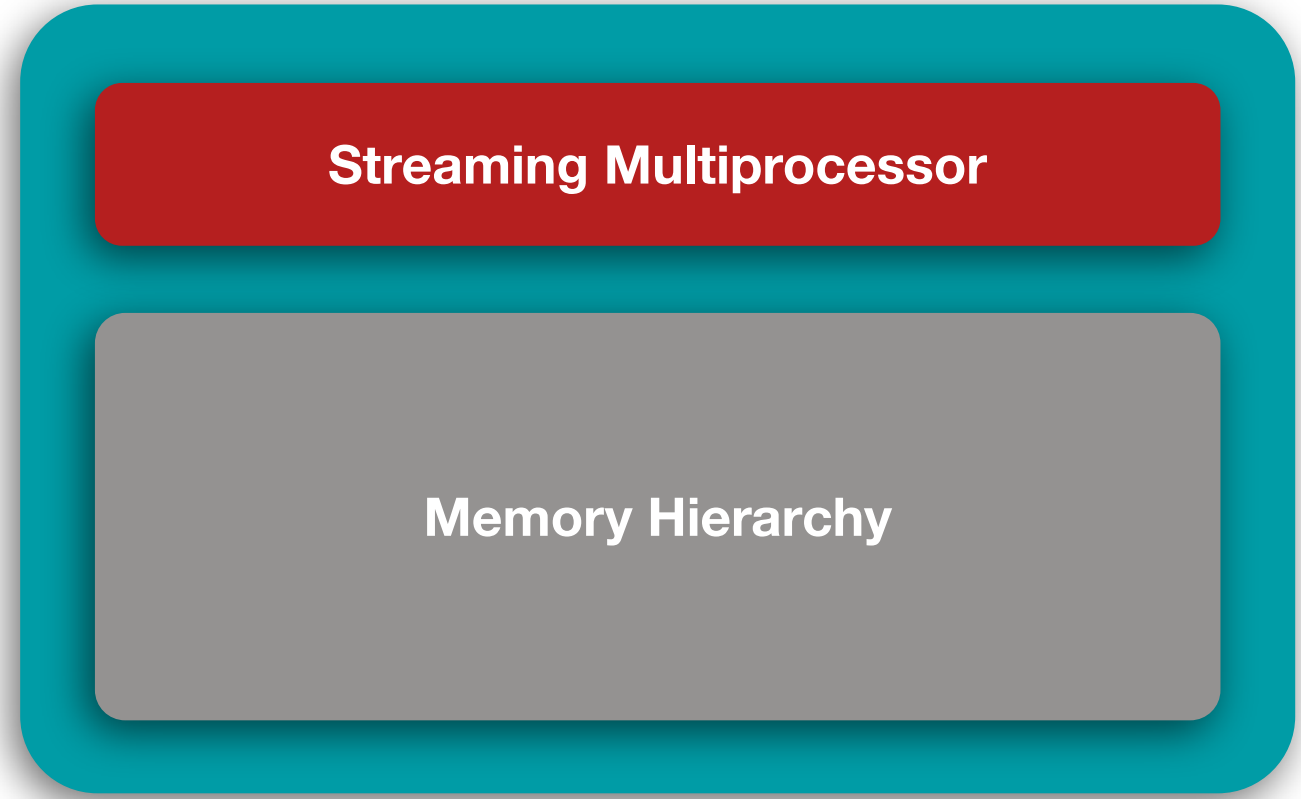

GPGPU Accelerator

Configurable

- Number of threads
- Number of warps
- Floating-point unit
- Memory hierarchy
  - Scratchpad-based

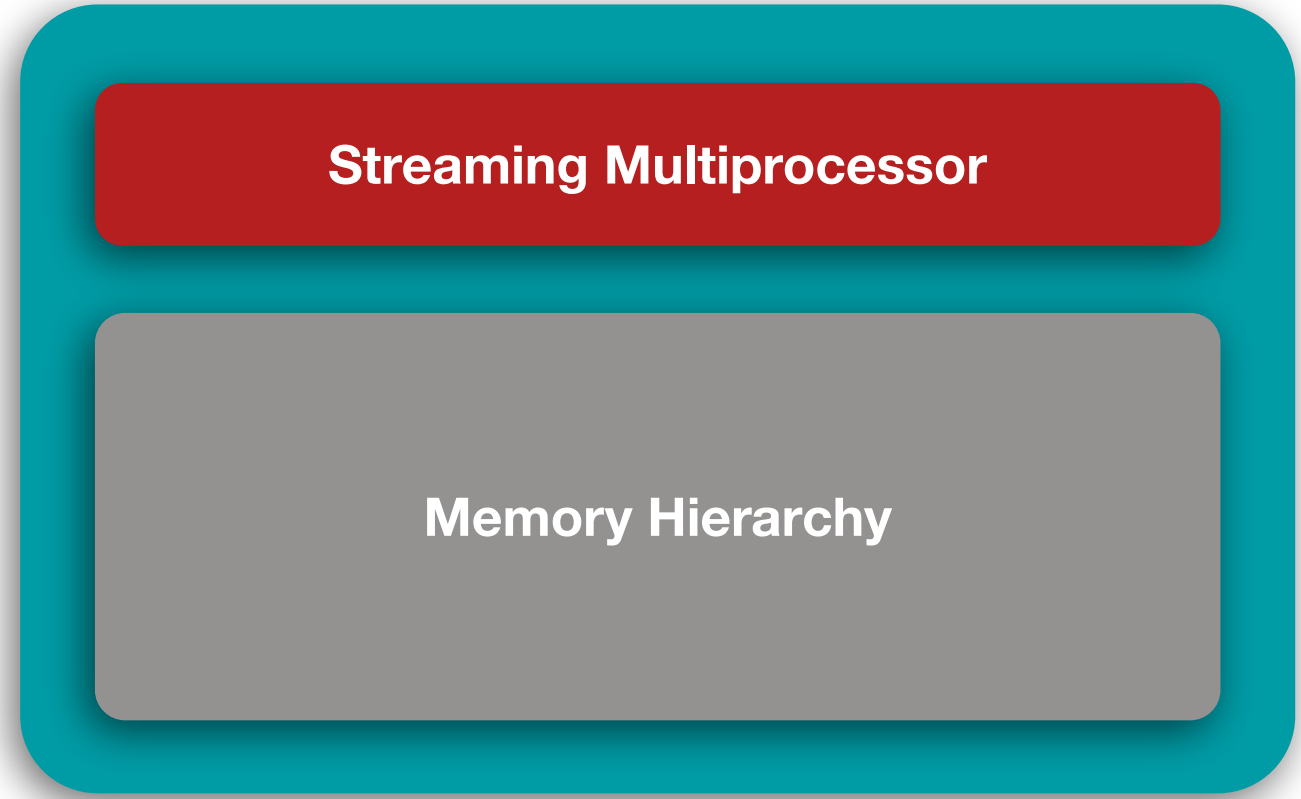

GPGPU Accelerator

Configurable

- Number of threads
- Number of warps
- Floating-point unit
- Memory hierarchy
  - Scratchpad-based
  - Cache-based

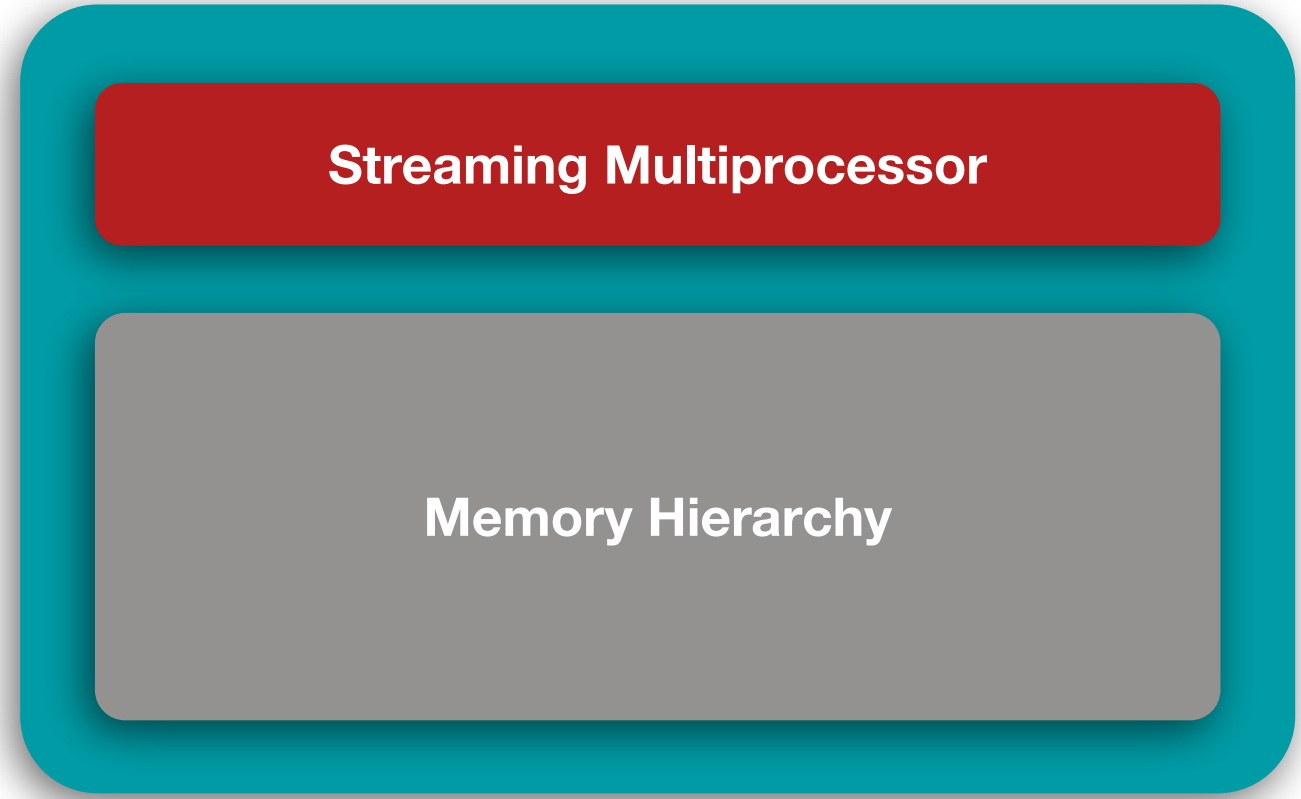

Accelerated Processing Unit (APU)

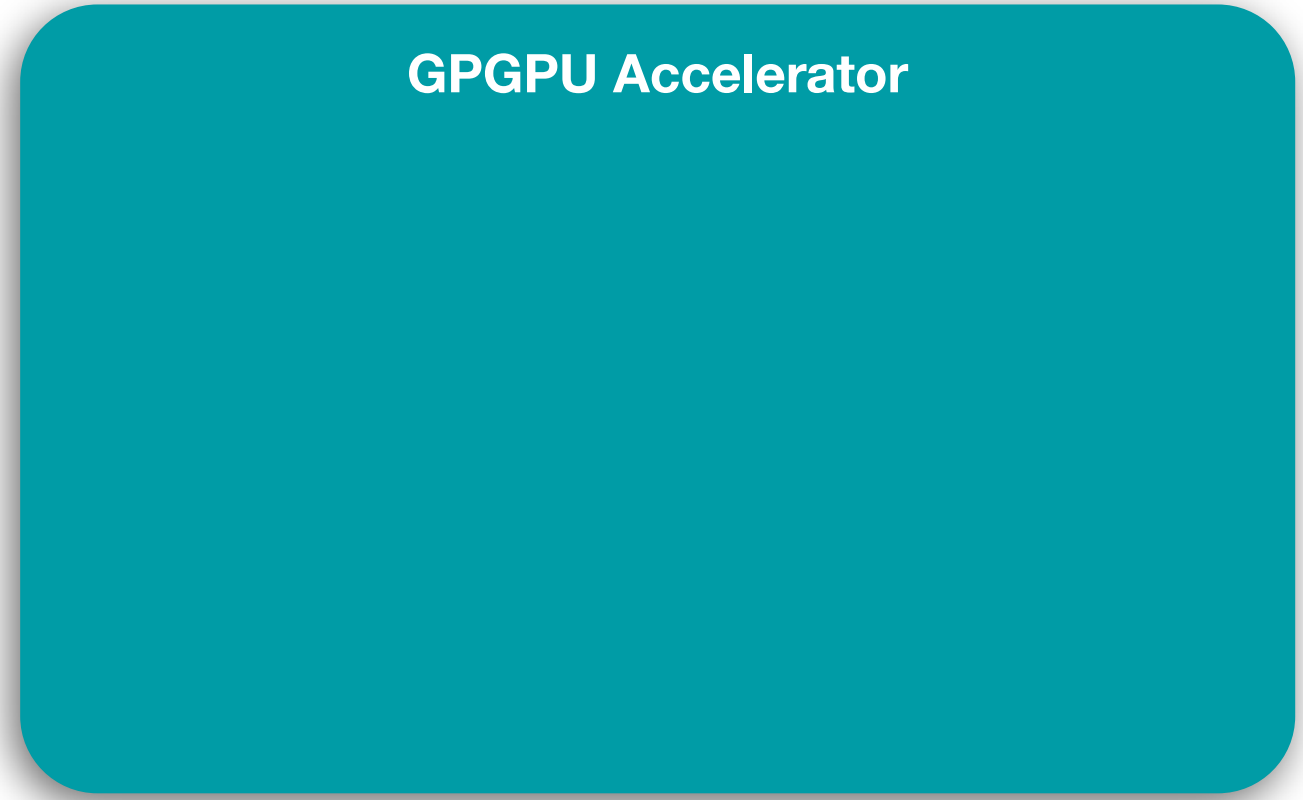

Accelerated Processing Unit (APU)

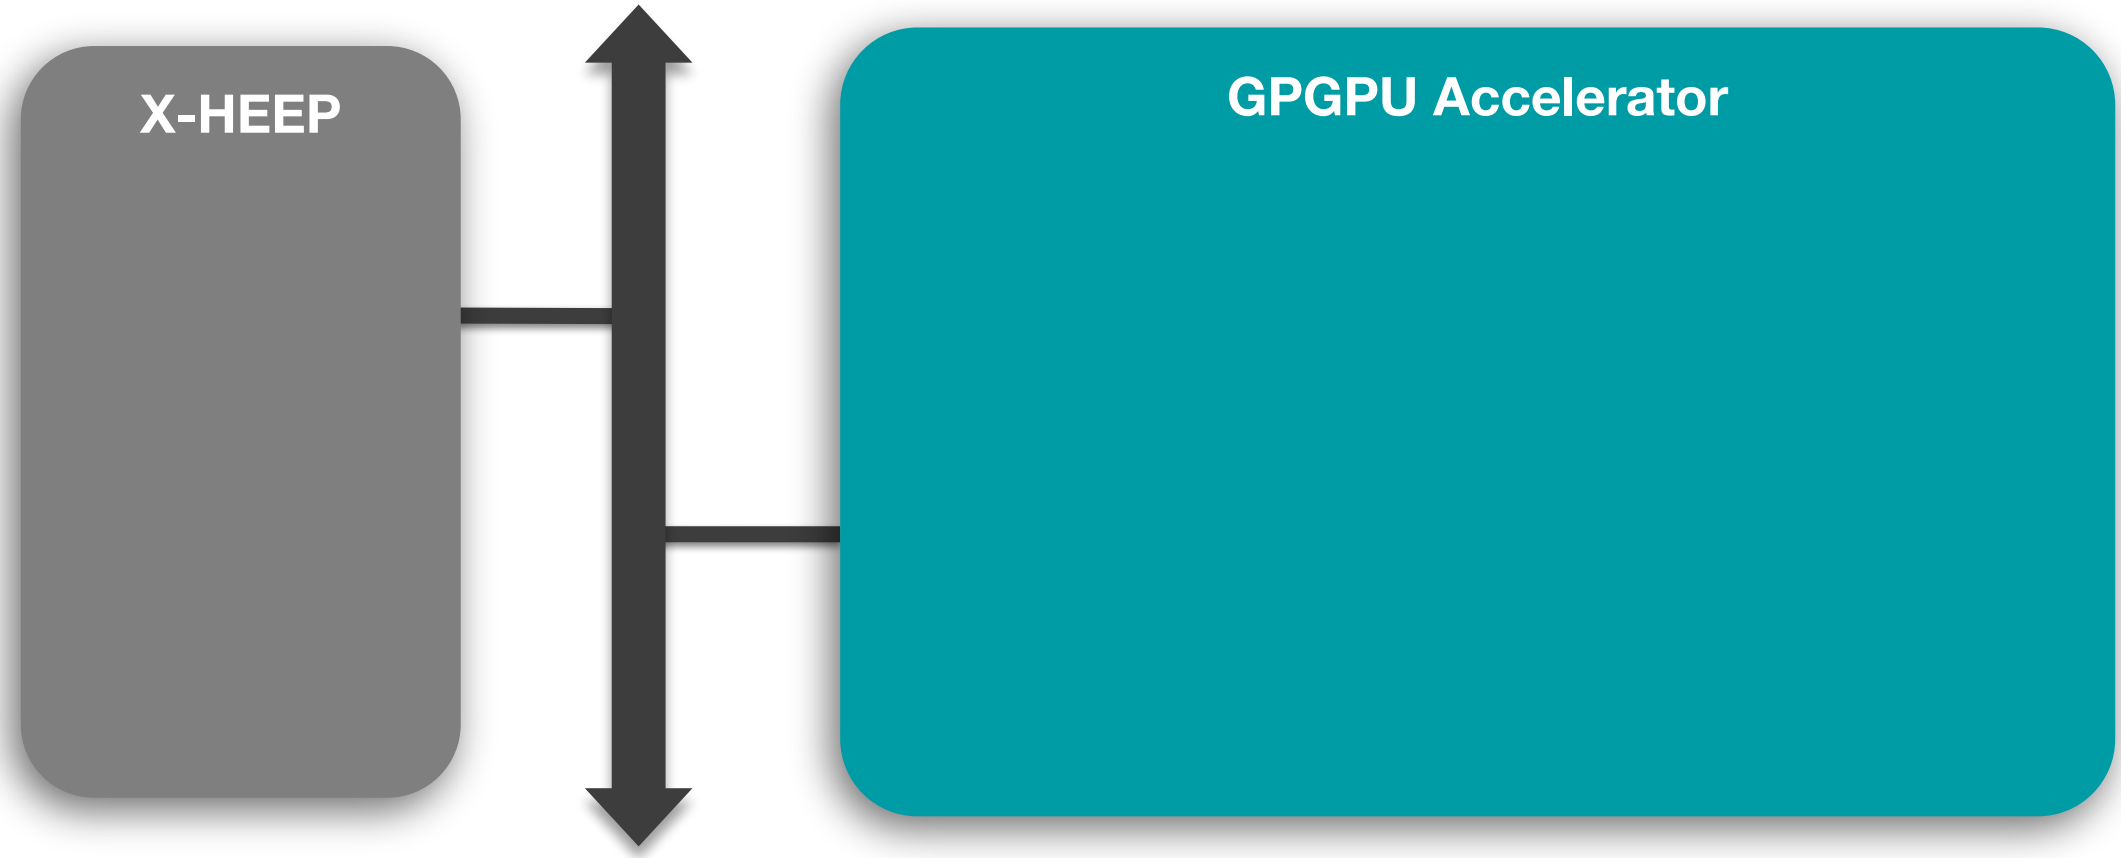

Accelerated Processing Unit (APU)

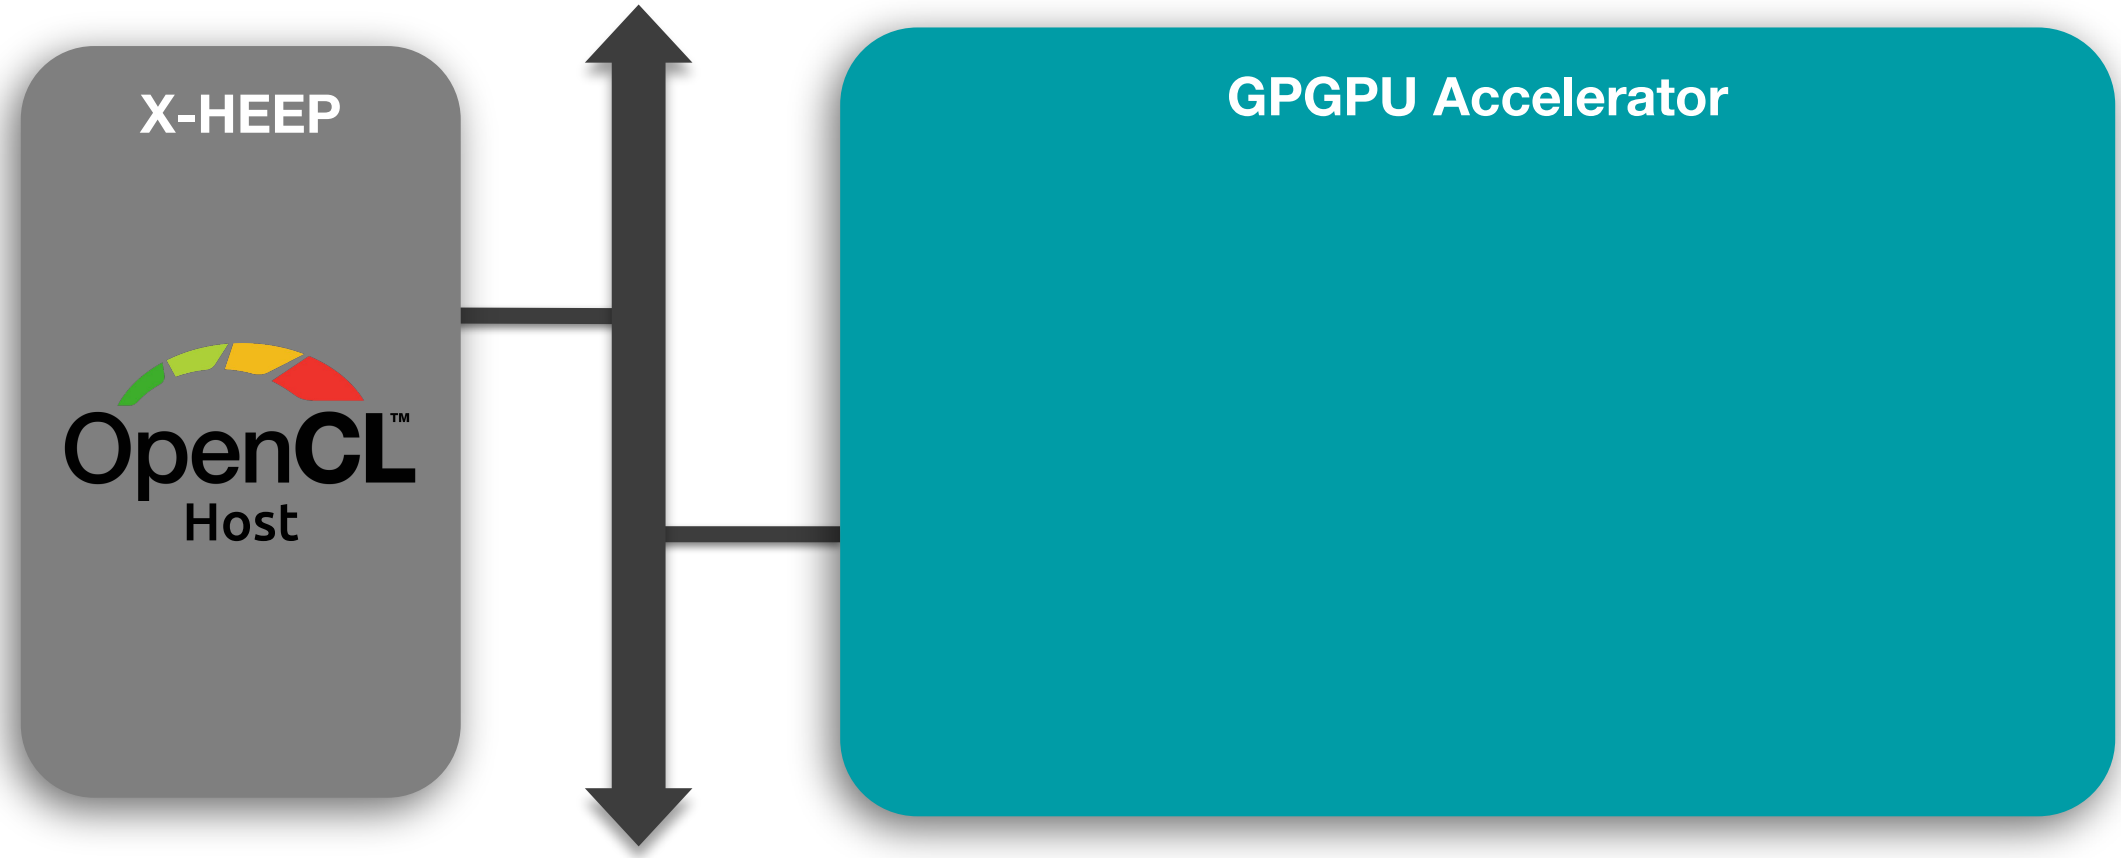

Accelerated Processing Unit (APU)

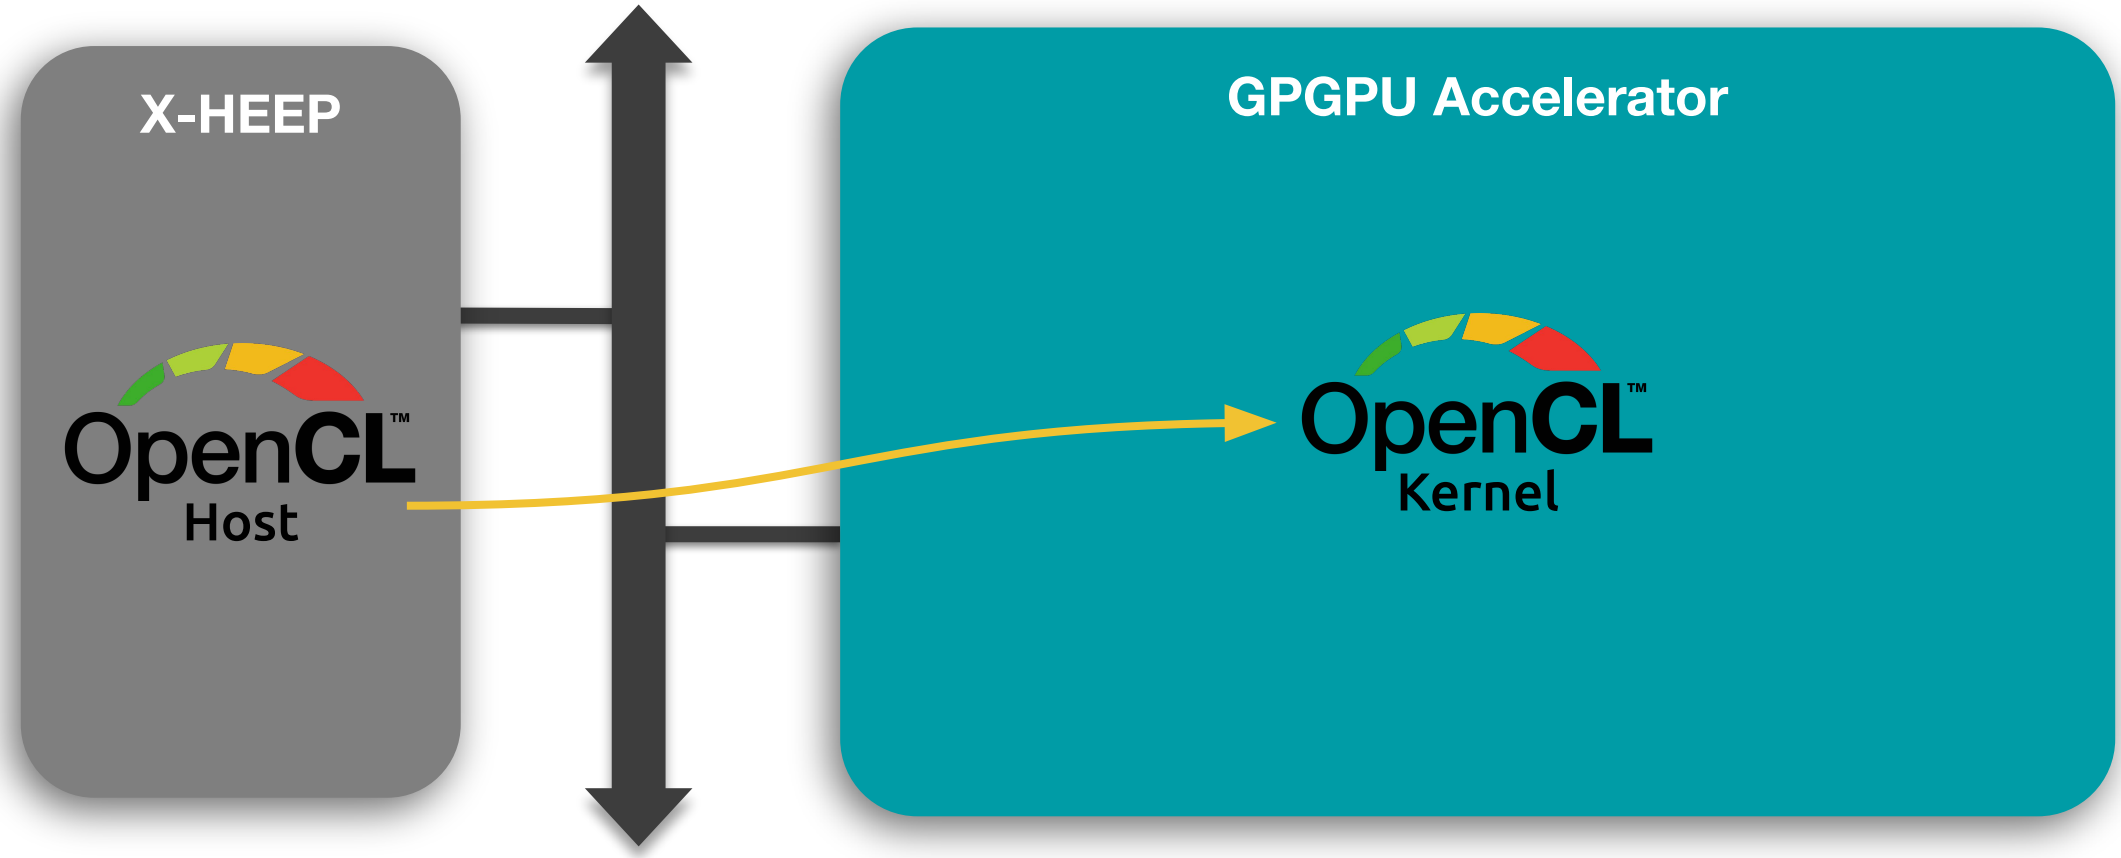

Conclusion

The APU code and documentation will be 100% open-source and the first version will be released very soon…

Conclusion

The APU code and documentation will be 100% open-source and the first version will be released very soon…

Conclusion

The APU code and documentation will be 100% open-source and the first version will be released very soon…

Stay Tuned!

EPFL

X-HEEP

Thank you for your attention!

EPFL - Embedded Systems Laboratory (ESL)

simone.machetti@epfl.ch