MEEP

MareNostrum Experimental Exascale Platform

www.meep-project.eu

MareNostrum Experimental Exascale Platform

Elias Perdomo (BSC-CNS) elias.perdomo@bsc.es
(on behalf of the MEEP Project Team meep-project@bsc.es)

What is MEEP[1]?

MEEP is a digital laboratory thar enables software/hardware co-design with pre-silicon validation of hardware architectures and a corresponding software development vehicle.

Main components of MEEP:

- HW:
  - 96 FPGA Infrastructure in 12 nodes (emulator for accelerators)
  - Open Source IPs for interfaces and communication (MEEP Shell)
  - Example Research Accelerator:
    - **ACME**: Accelerator design with disaggregated architecture
- SW: Suite of tools and software developed for HPC to be run on validated accelerators
- **Coyote**: Performance Modeling tool for many-core accelerators such as ACME.

MEEP — ACME — Coyote

- **HW&SW Co-design lab for HPC**
- **3 year project (Jan 20 - June 23)**
- **10M€ budget**
- **Based on RISC-V Open Source & European Processors**

MEEP FPGA INFRASTRUCTURE

MARENOSTRUM MN4

WHAT IS ACME[1]?

- Accelerated Compute and Memory Engine
- Self-Hosted Many-Core Accelerator for Exascale Systems
- Vector Length Agnostic
- Disaggregated Architecture:
  - Compute bound operations in VAS Tile
  - Memory bound operations in Memory Tile

ACME ARCHITECTURE

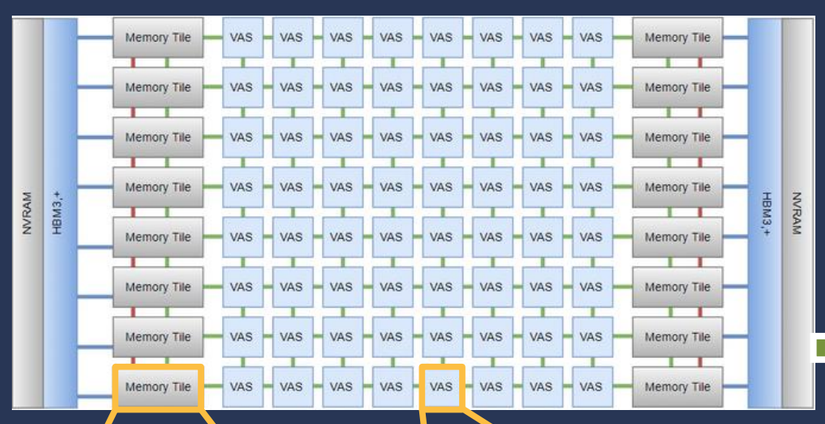

16 Memory Tiles

- Memory CPU (MCPU - RV64IMA)
- Vector Memory Operations (unit stride, non-unit, indexed)

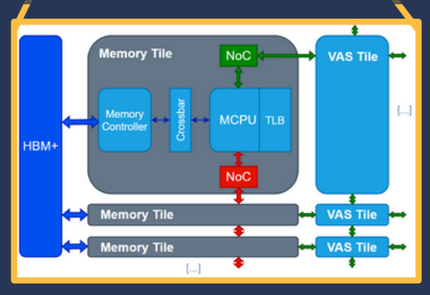

64 VAS Tiles

- VAS: Vector And Systolic Array
- 8 Cores sharing L2

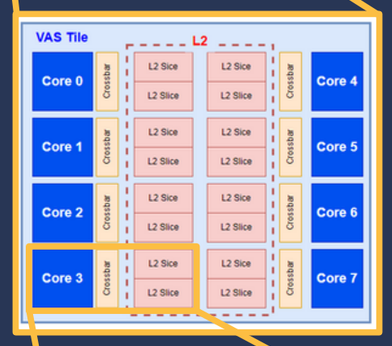

Core

- RV64GCV
- VPU 16-lanes
- Neural Network Systolic Array (SA-NN)
- High Efficiency Video Coding Systolic Array (SA-HEVC)

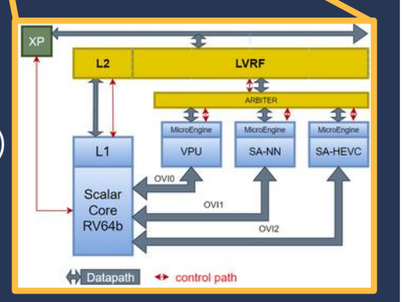

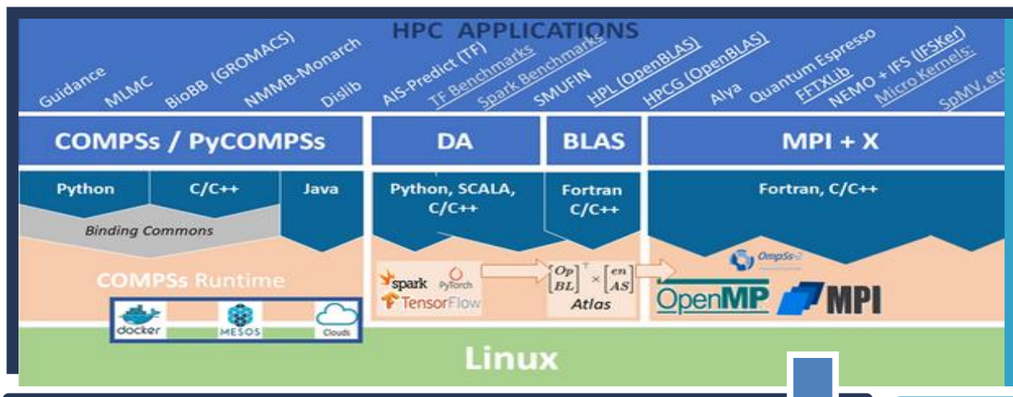

SW STACK

- Workloads to accelerate in MEEP platform
- SW Support at different levels
  - Compiler
  - Linux OS on ACME
  - Services & extensions for Host/FPGA communication

MEEP FPGA SHELL

- Flexible and custom wrapper for any design
- Guarantees standard communication interfaces between emulated accelerator & outer system (PCIe, Aurora, Ethernet, HBM)

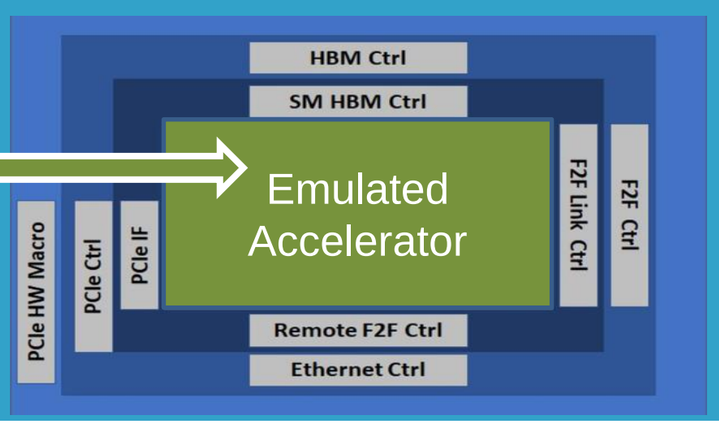

COYOTE[2]

- Execution-driven Simulator
- Spike[3] for RISC-V cores
  - Multi-core System
  - Vector Extension
- Sparta[4] Framework for Timing
- BookSim[5] for NoC Simulations
- Focused on Data Movement & Modeling of the Memory Hierarchy

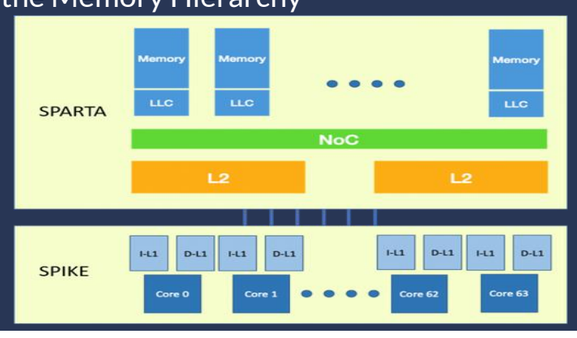

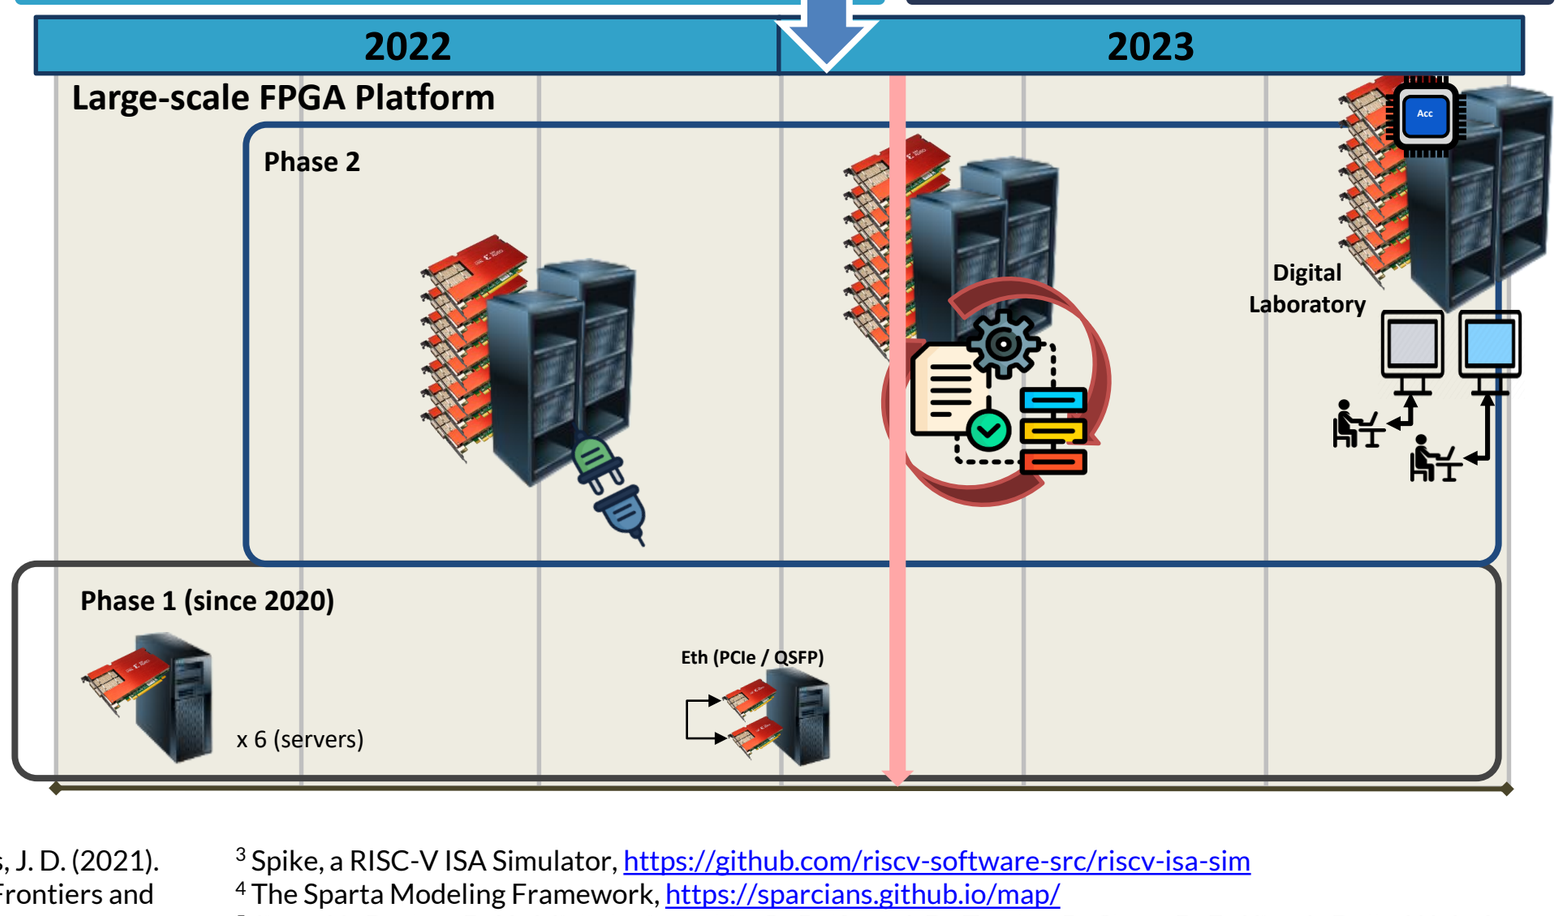

References:

[1] Fell, A., Mazure, D. J., Garcia, T. C., Perez, B., Teruel, X., Wilson, P., & Davis, J. D. (2021). *The MareNostrum Experimental Exascale Platform (MEEP)*. Supercomputing Frontiers and Innovations, 8(1), 62–81. https://doi.org/10.14529/jsfi210105

[2] Perez, B., Fell, A., Davis, J. D.: *Coyote: An Open Source Simulation Tool to Enable RISC-V in HPC*. Design, Automation, and Test in Europe, DATE (2021)

[3] Spike, a RISC-V ISA Simulator, https://github.com/riscv-software-src/riscv-isa-sim

[4] The Sparta Modeling Framework, https://sparcians.github.io/map/

[5] Jiang, N., Becker, D. U., Michelogiannakis, G., Balfour, J. D., Towles, B., Shaw, D. E., Kim, J., Dally, W. J.: *A detailed and flexible cycle-accurate Network-on-Chip simulator*, *ISPASS*, *pp. 86-96*, *IEEE Computer Society*, *2013*.

EuroHPC Joint Undertaking

The MEEP Project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 946002. The JU receives support from the European Union's Horizon 2020 research and innovation program and Spain, Croatia and Turkey

MEEP PARTNERS

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

University of Zagreb Faculty of Electrical Engineering and Computing

TÜBİTAK İLGEM