HyperFPGA: an Experimental Testbed for Future Heterogeneous Cluster Architectures

Werner Florian Samayoa

ICTP, UNITS, INFN werner_florian@live.com

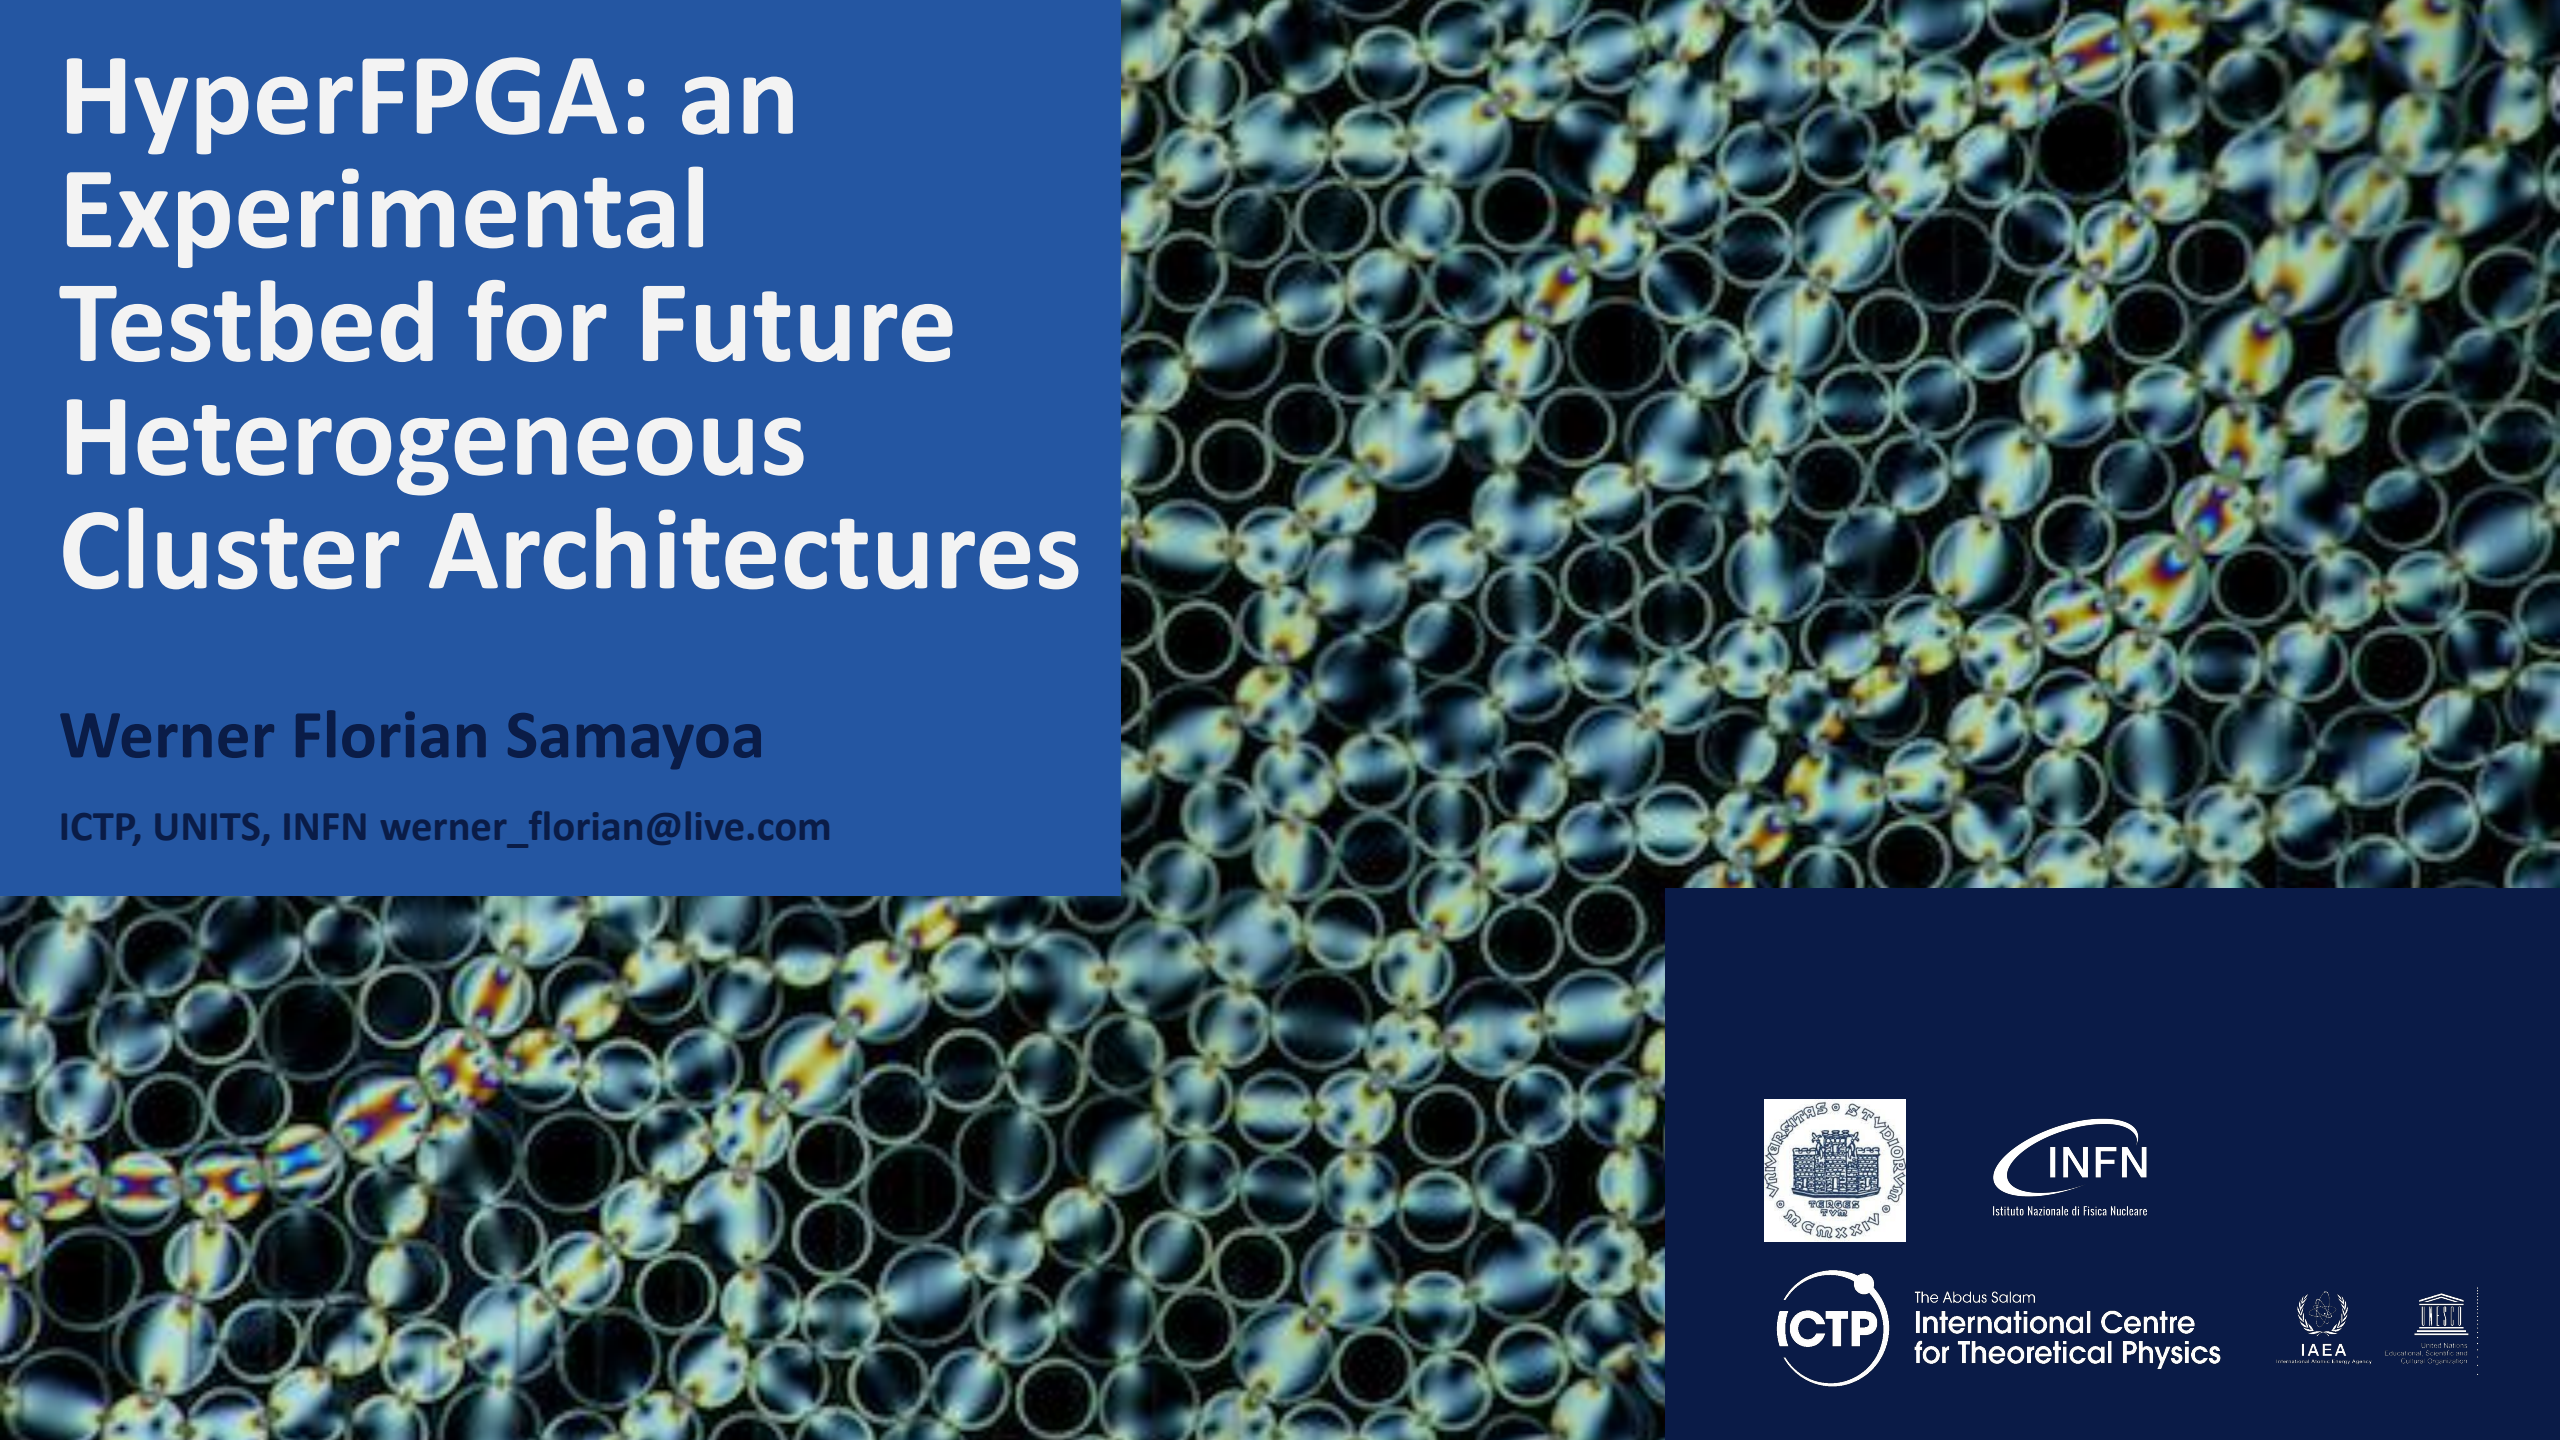

INFN
Istituto Nazionale di Fisica Nucleare

The Abdus Salam
International Centre
for Theoretical Physics

IAEA

UNESCO

Overview

- Trip to memory lane, *or what was before the HyperFPGA?*
- How the HyperFPGA came to be?
- The hardware of the HyperFPGA
- The software of the HyperFPGA
- Queens problem, *and how to use the HyperFPGA*
- Future works
- Conclusions

Are FPGAs good for general-purpose computing?

A trip down memory lane

In 1985

In 1993

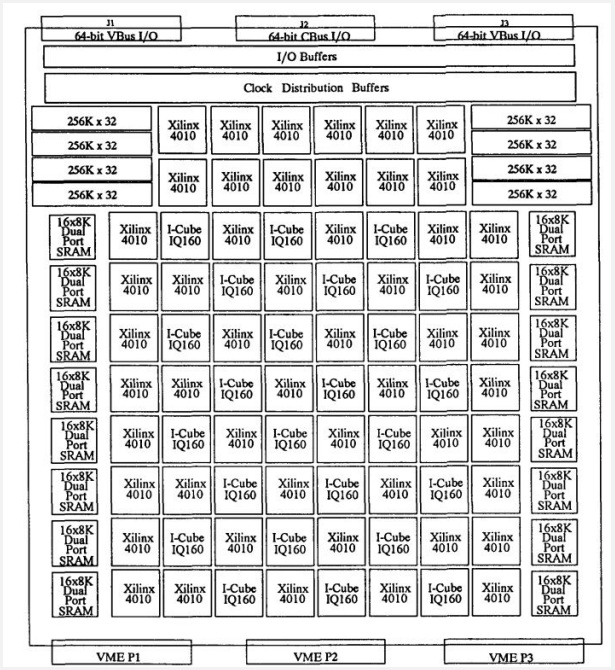

Today 2023

Cluster of FPGAs for manycore emulation

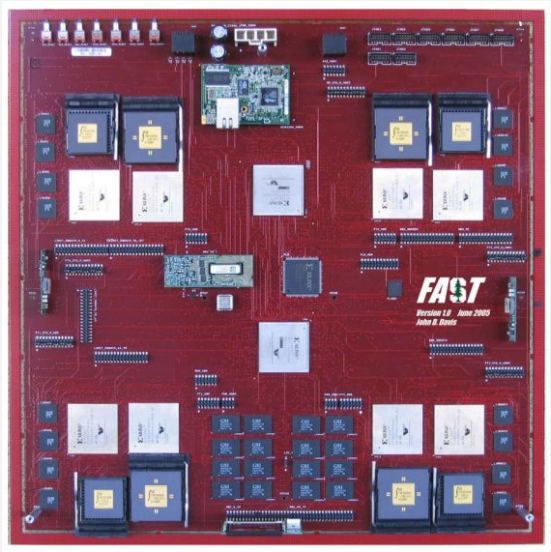

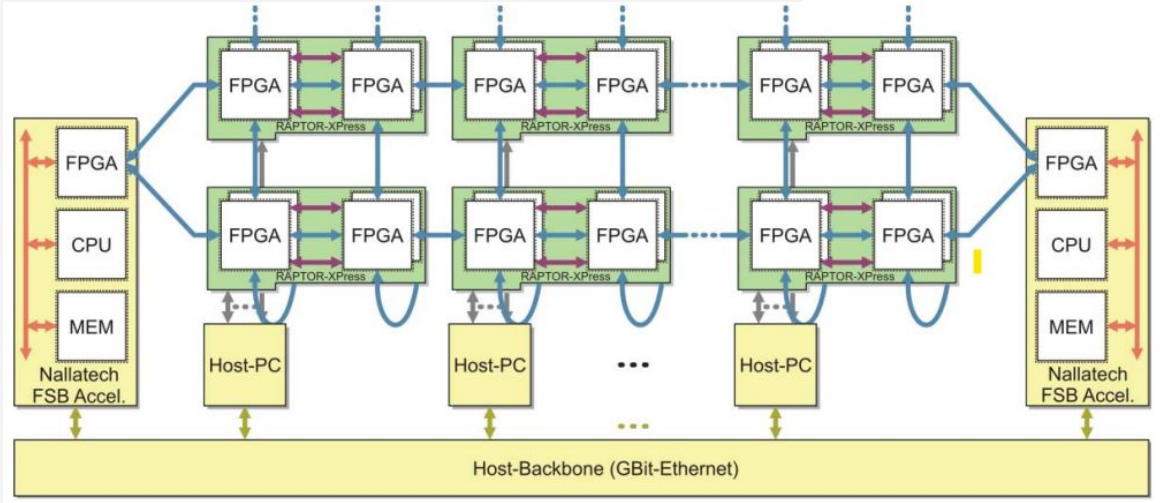

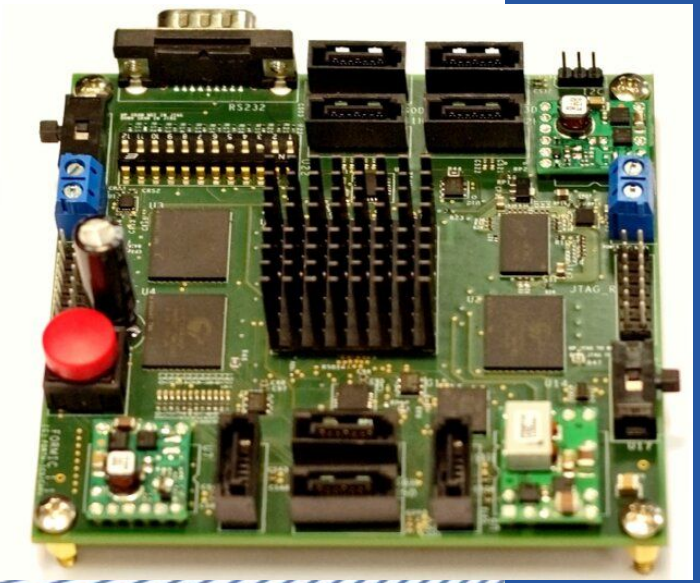

FAST 2002 — **RAPTOR 2010** — **Formic 2010**

Are FPGAs good for general-purpose computing?

Cluster of FPGAs for scientific computing

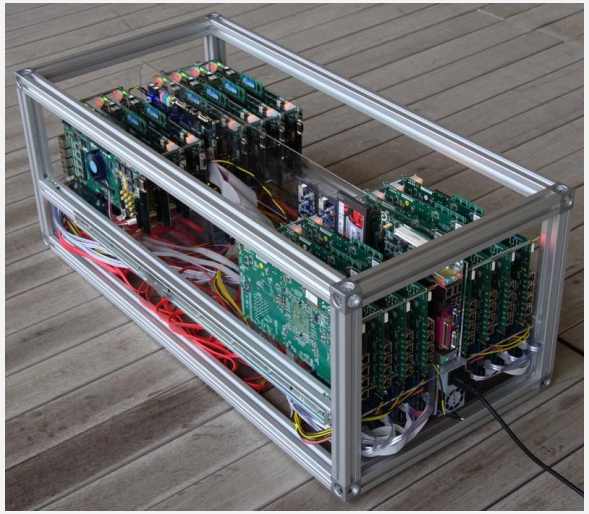

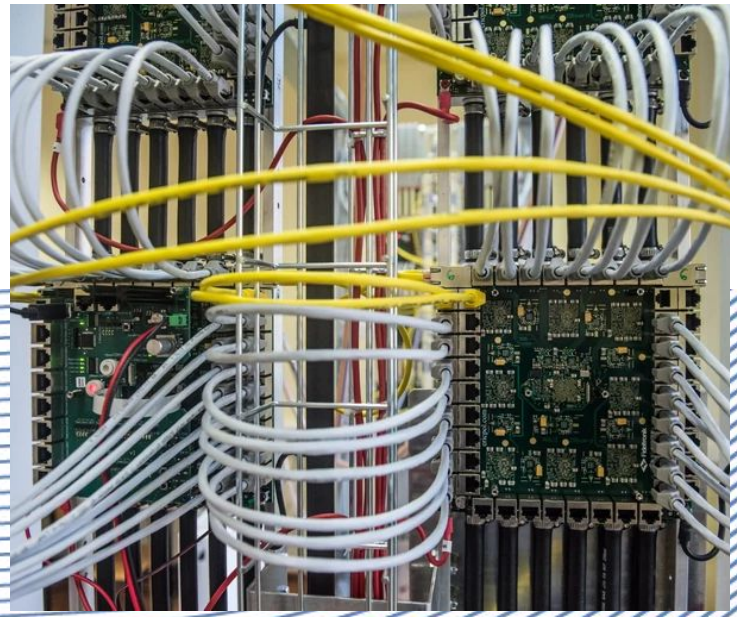

Janus
2006

Bluehive
2012

ARUZ
2018

Are FPGAs good for general-purpose computing?

Cluster of FPGAs in data centers

Microsoft Project Catapult
2014

IBM cloudFPGA
2015

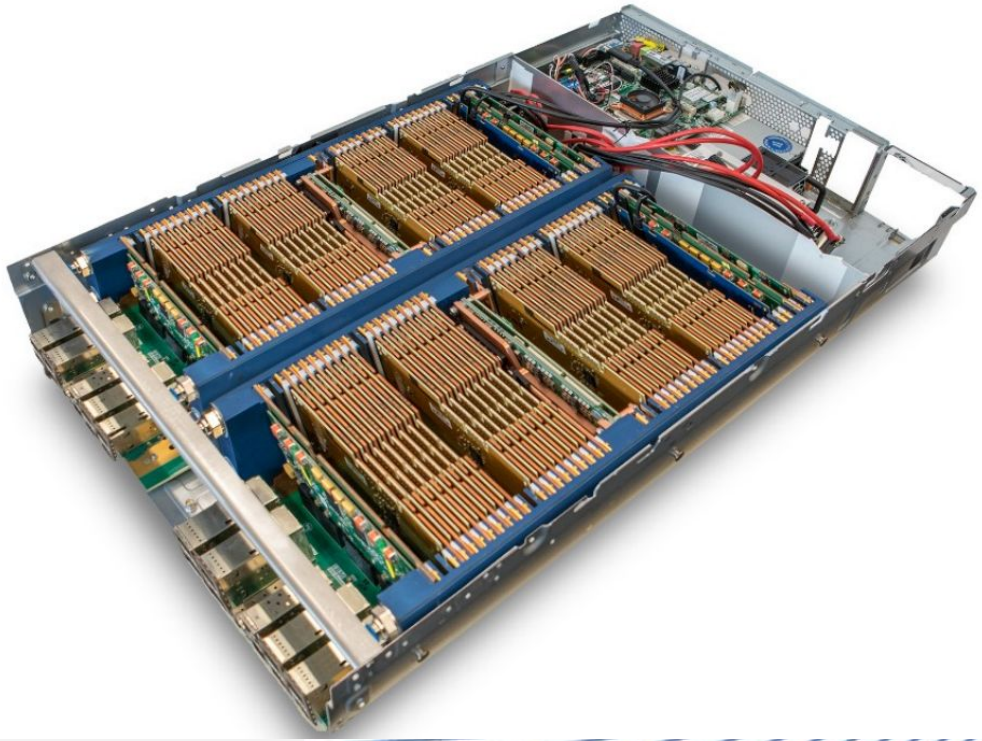

Cluster of FPGAs as communication infrastructure

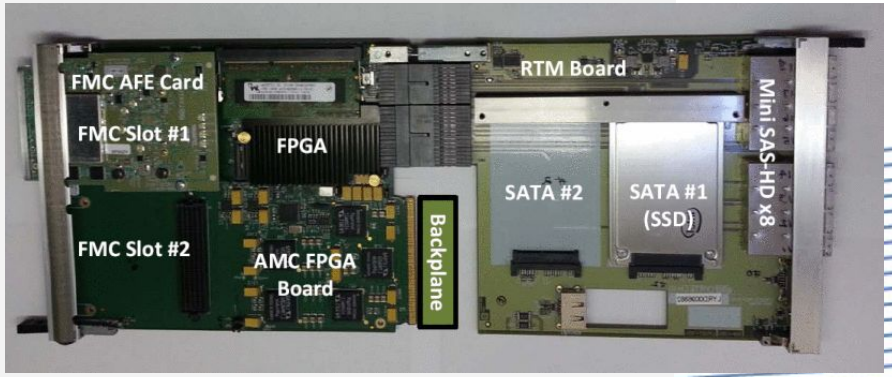

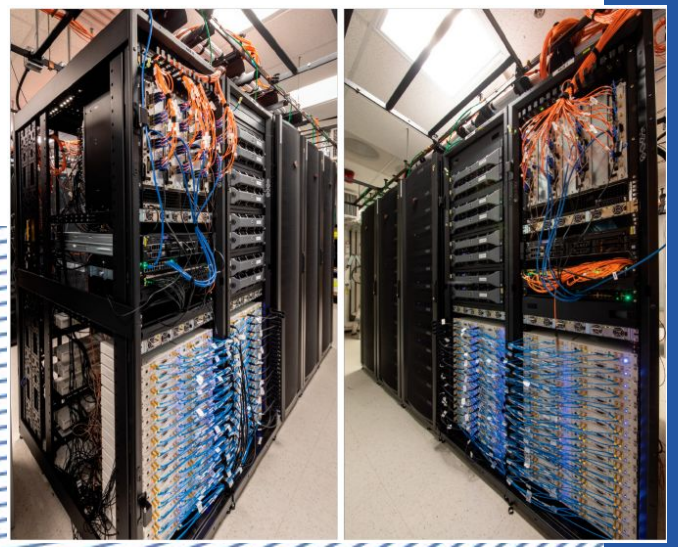

BEE4 2010 — **DWCE 2013** — **Colosseum 2021**

Cluster of FPGAs for general purpose applications

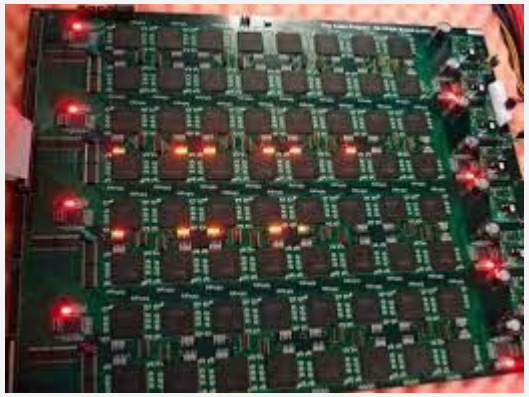

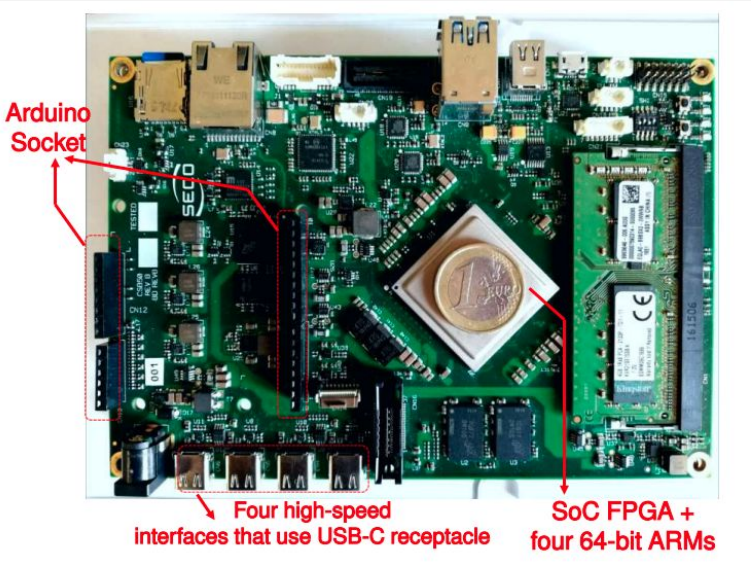

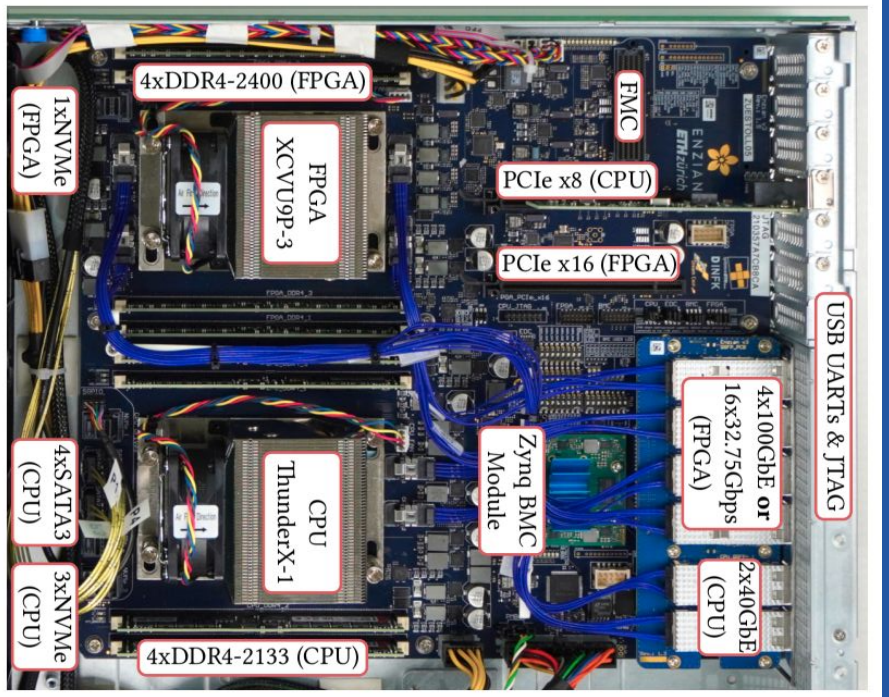

Cube 2009 — **Axiom 2017** — **Enzian 2020**

HyperFPGA design philosophy

HPC platform domains

Hardware

- Heterogeneous
  - FPGAs
  - CPUs
  - GPUs
- Experimental
  - Power monitoring
- Flexibility
  - Physical systems interface

Network

- Dual network
  - CPU network
  - FPGA dedicated network
- Experimental
  - Protocol independent
  - Physically independent
- Modularity
  - IF, Router, switch, etc.
- Flexibility
  - Topology

Software

- Open source
  - Administration
  - User interaction
- User diversity
  - Different levels of abs.
- Safe
  - OS abstraction
  - Memory protection
- Flexibility
  - Multiple parallel paradigms

The hardware of the FFeCCa

- SoM and SoC-FPGA based

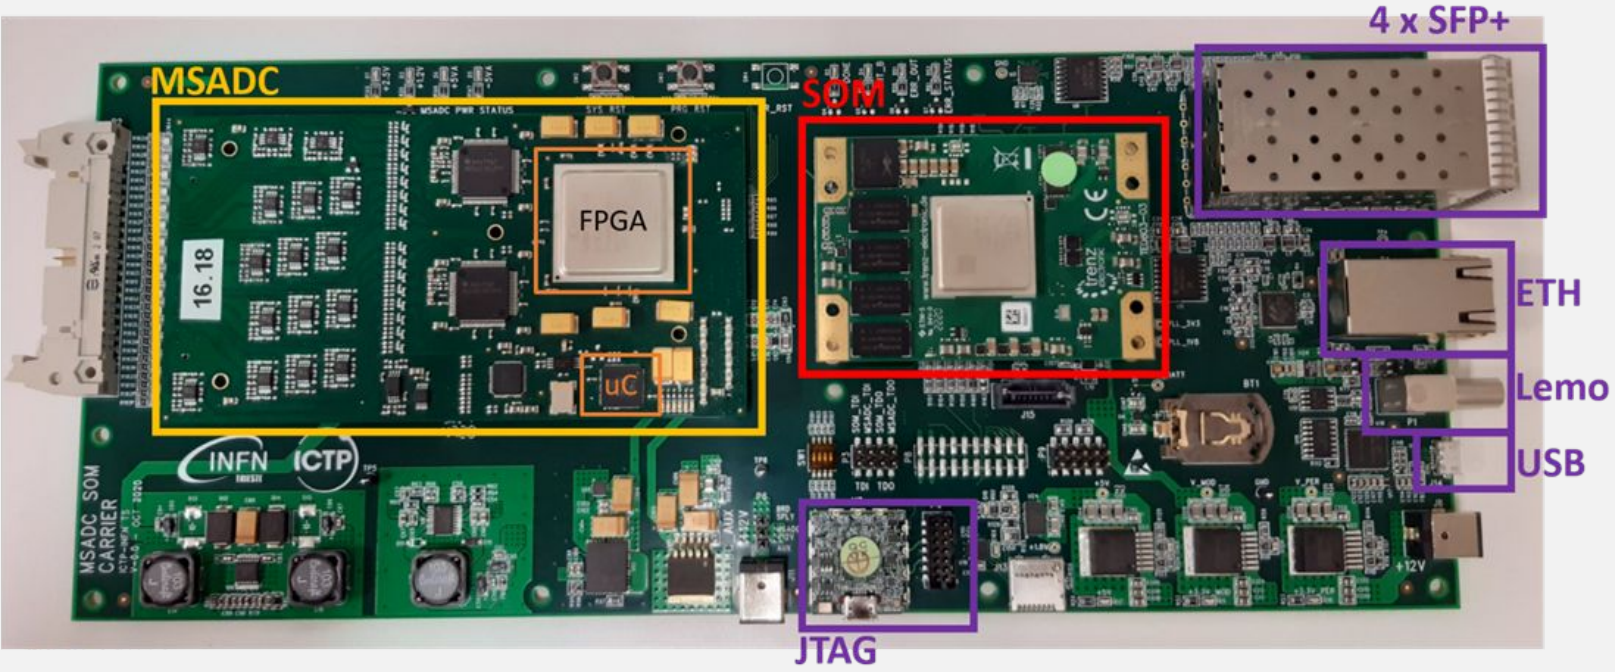

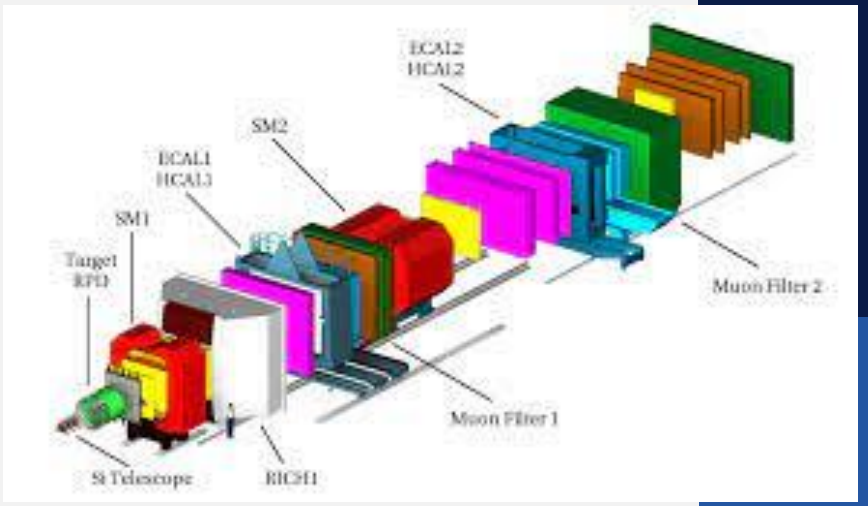

The hardware of the HyperFPGA

- SoM and SoC-FPGA based
  - Open pinout
  - Heterogeneous
  - Simpler layout
  - Vendor independence
  - Validated approach (AMBER ECAL 2 FEE)
- Flexible
  - HP, HD and GTH signals for interconnection

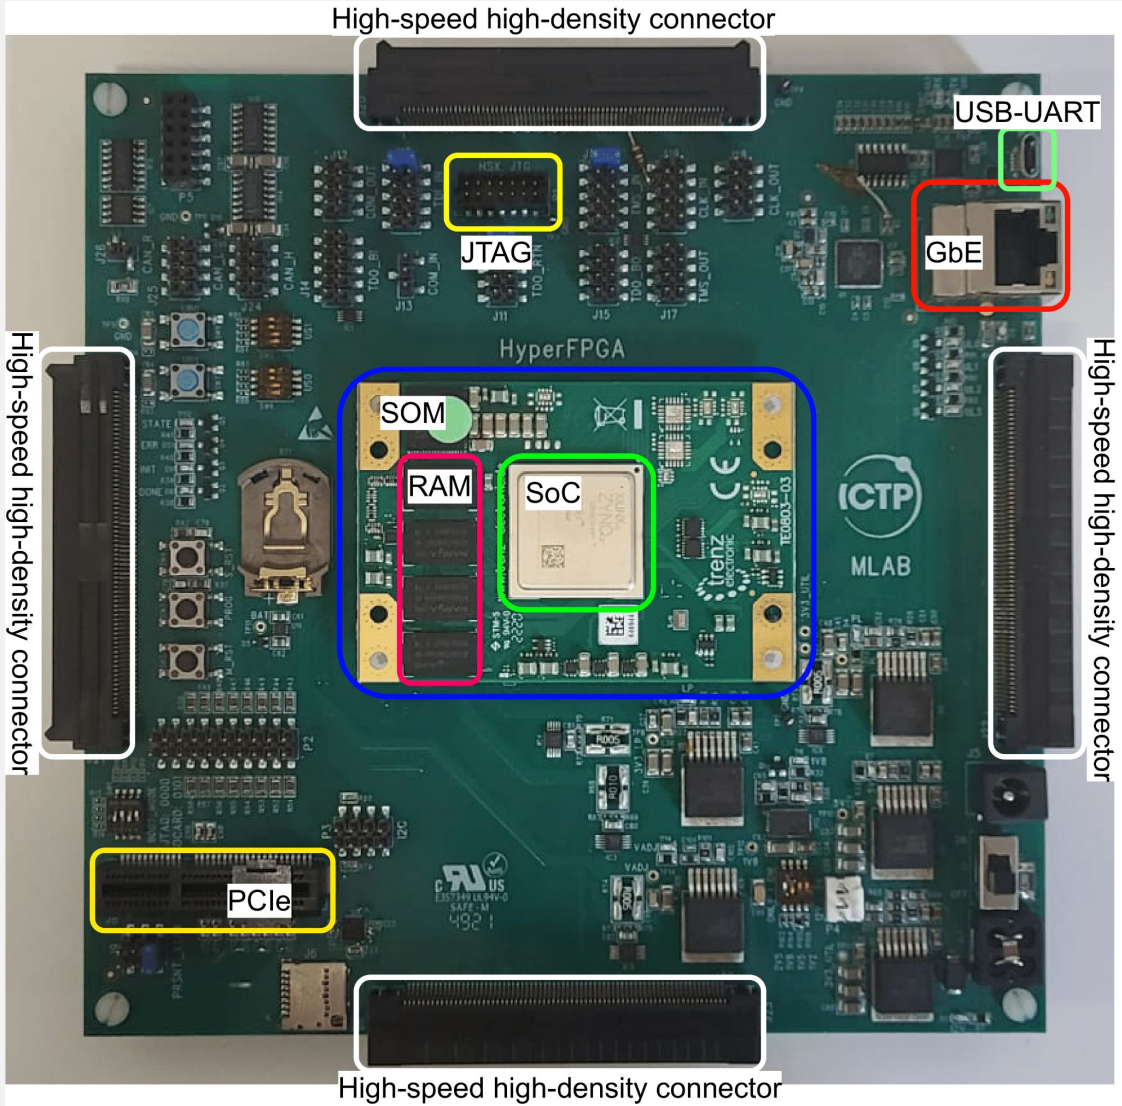

The software of the HyperFPGA

- XSA2Bit
  - Generate device tree overlay from an XSA and compile sources
- ComBlock
  - An abstraction layer between CPU and FPGA
- Linux Drivers
  - Secure access for read and write operations
- JupyterHub
  - Authentication, authorization, and accountability

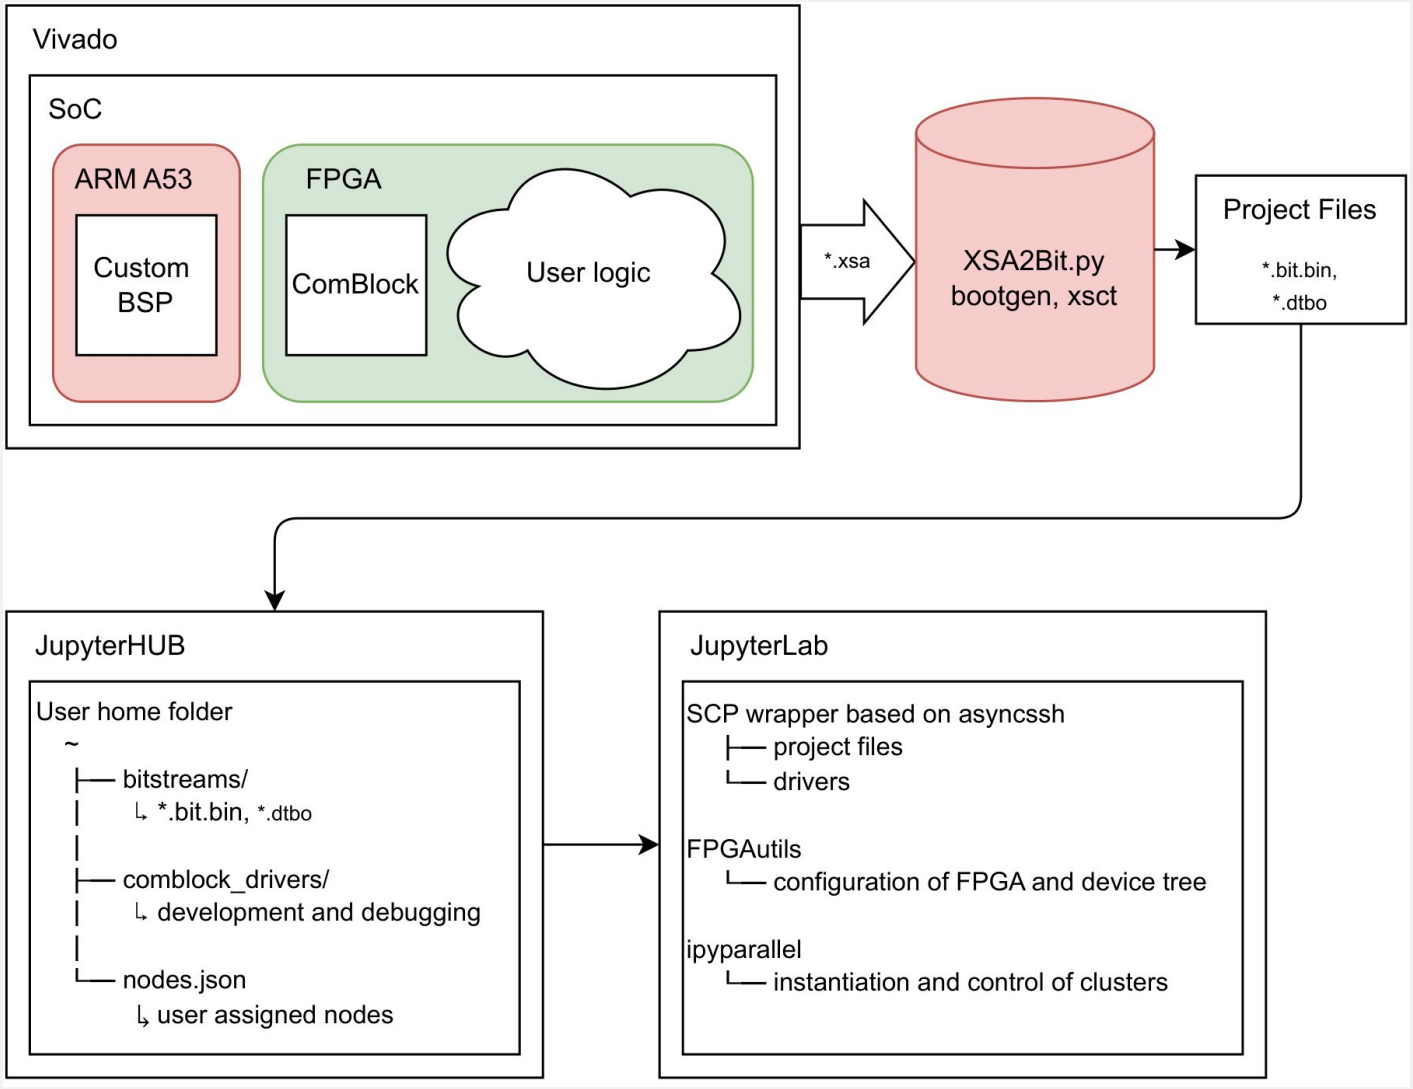

The boot sequence of the HyperFPGA

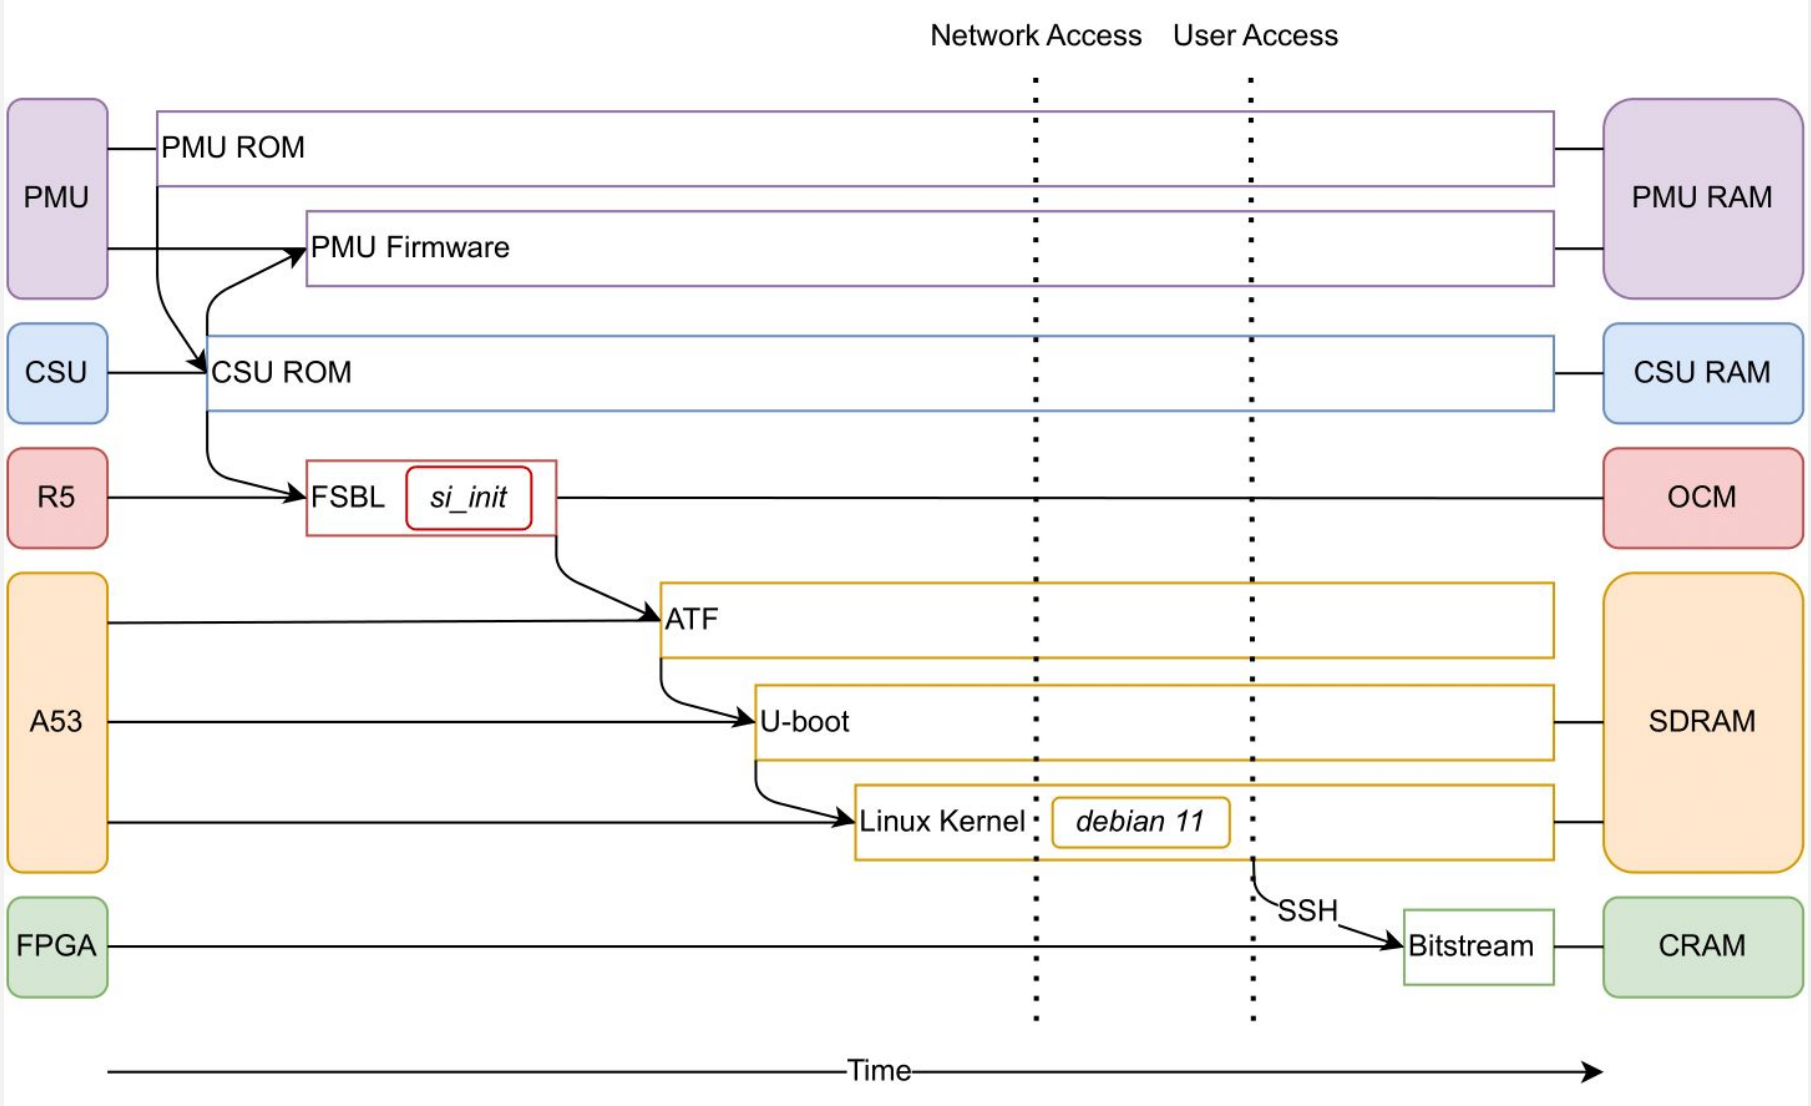

The software of the HyperFPGA

- One to many approach, ideal for parallel implementations
- One to one approach, ideal education

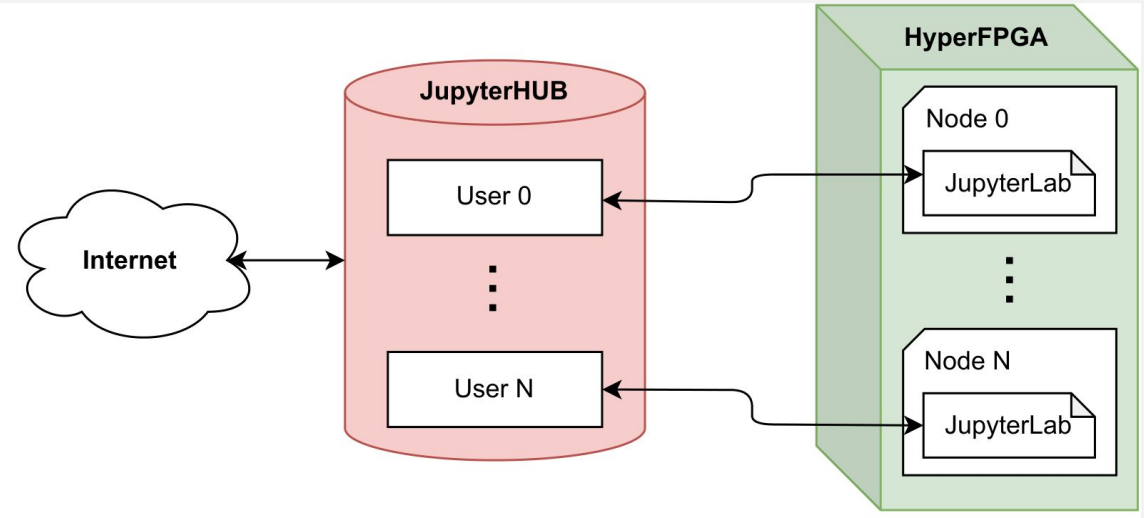

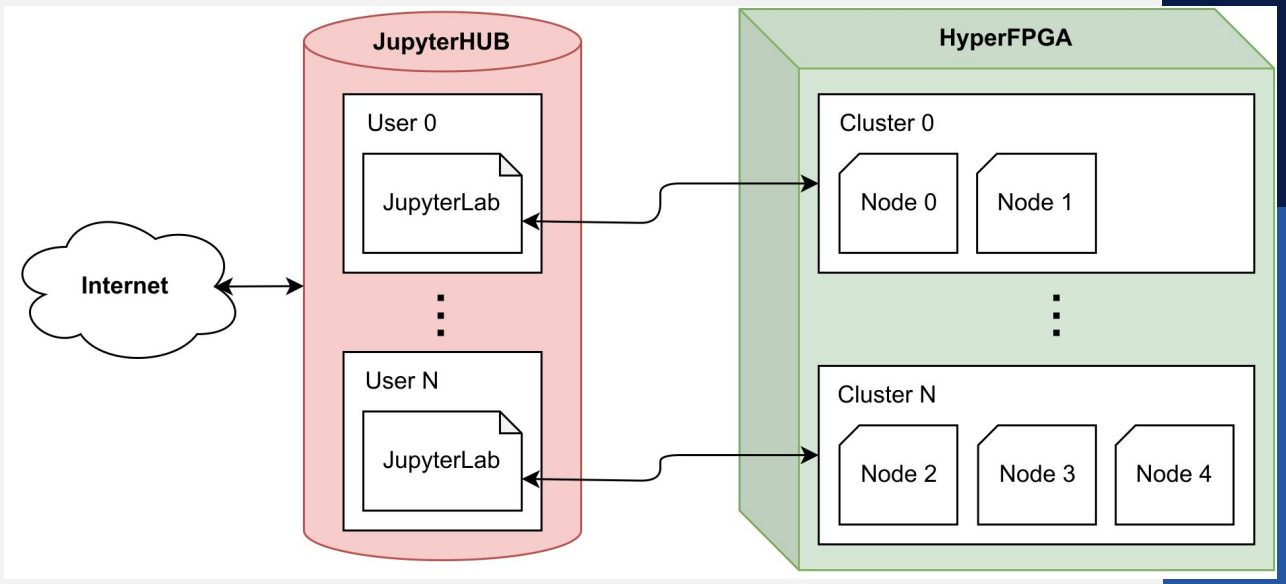

Example app. N-Queens problem

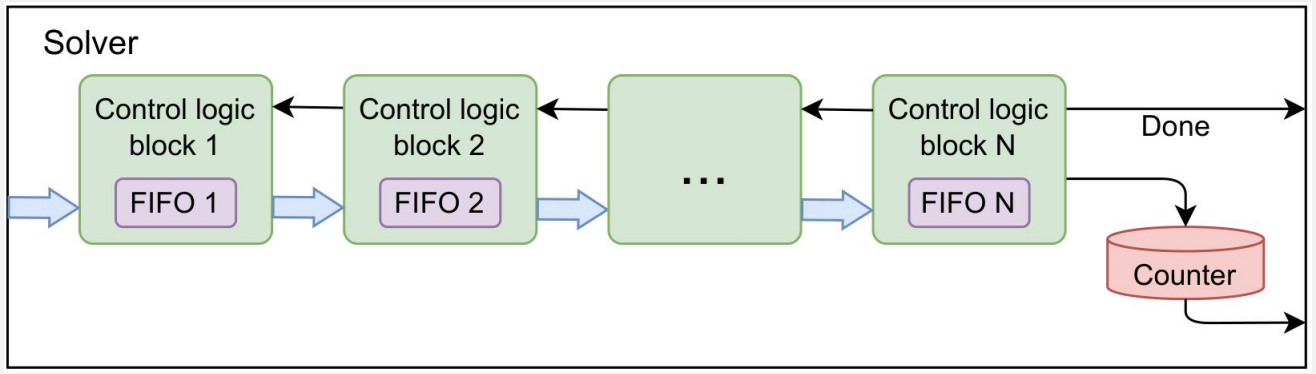

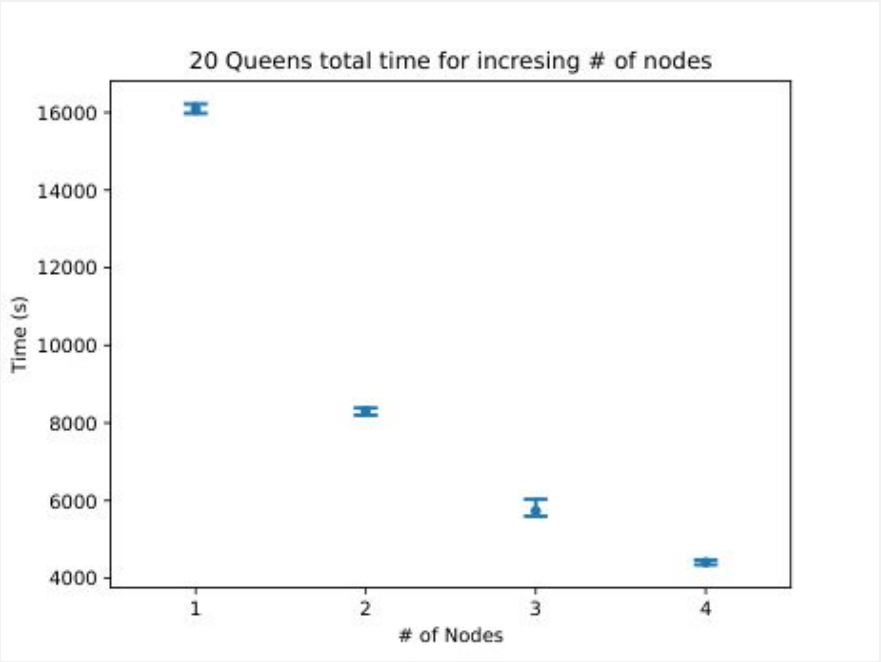

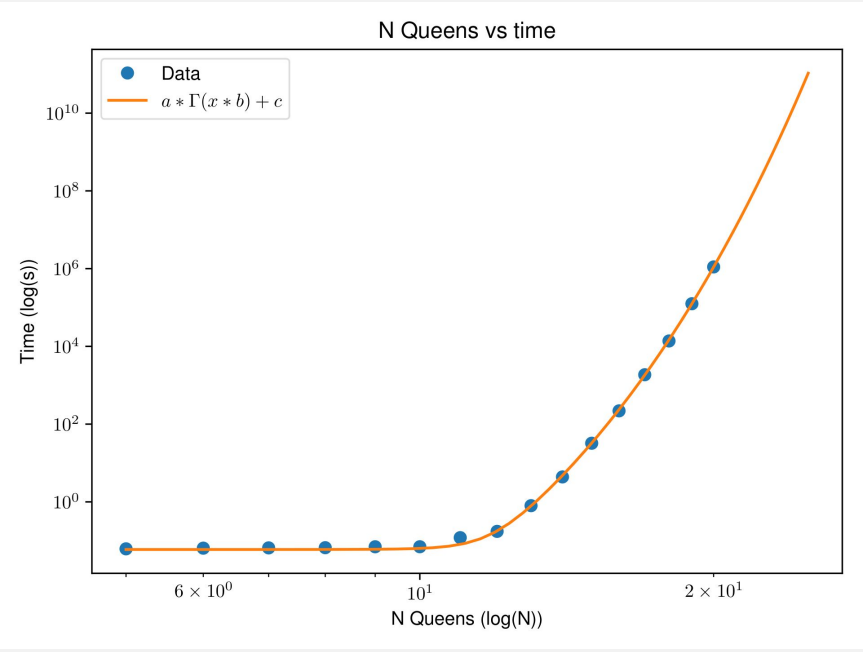

Future works

- More problems!!!
  - Quantum computer emulator
  - High-level trigger
- FPGA dedicated network
  - Modular interface for GTH, HP, and HD
- Novel distributed computing paradigms exploring direct communication between FPGAs
- High-level synthesis solutions for FPGA clusters
- Design space exploration tools
  - Automatic segmentation of the architecture
  - Better estimate of the resources, power, and performance

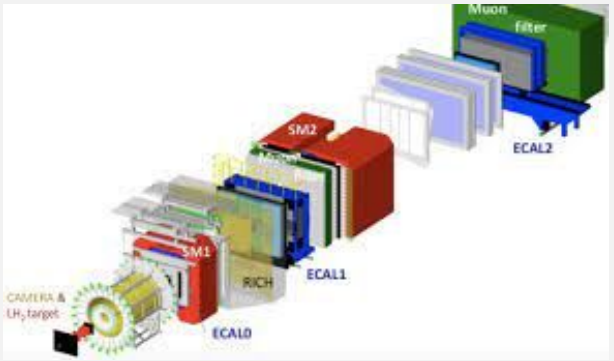

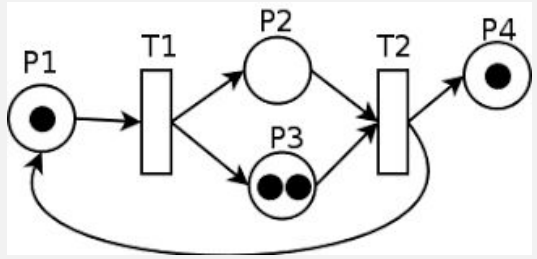

MLIR

Conclusions

- Experimental HyperFPGA cluster was developed on the principles of openness, modularity, and scalability.
- Vendor interoperability with an open SoM form factor and pinout.
- A custom Linux OS was built to accommodate remote management requirements using JupyterHub.
- Using only JupyterLab notebooks and the predefined interconnection layer consisting of the ComBlock and Linux kernel driver an application was shown.
- Backtracking algorithm for the N-Queens problem, a speed-up of about 50 times when compared to a Zedboard implementation, was obtained showing the potential of the hardware and the efficiency of the development environment.

Thank you!

ICTP Mlab Gitlab

https://gitlab.com/ictp-mlab/n-queens
https://gitlab.com/ictp-mlab/hyperfpga-hw
https://gitlab.com/ictp-mlab/hyperfpga-linux

Digital Object Identifier 10.1109/ACCESS.2023.3288431

SURVEY

A Survey on FPGA-Based Heterogeneous Clusters Architectures

WERNER FLORIAN SAMAYOA[1,2], MARIA LIZ CRESPO[1], ANDRES CICUTTIN[1], AND SERGIO CARRATO[2]

HyperFPGA: An Experimental Testbed for Heterogeneous Supercomputing

Werner Florian Samayoa, Maria Liz Crespo, Sergio Carrato, Agustin Silva, and 1 more

This is a preprint; it has not been peer reviewed by a journal.

https://doi.org/10.21203/rs.3.rs-3278560/v1
This work is licensed under a CC BY 4.0 License