



Green Flash

High performance computing for real-time science

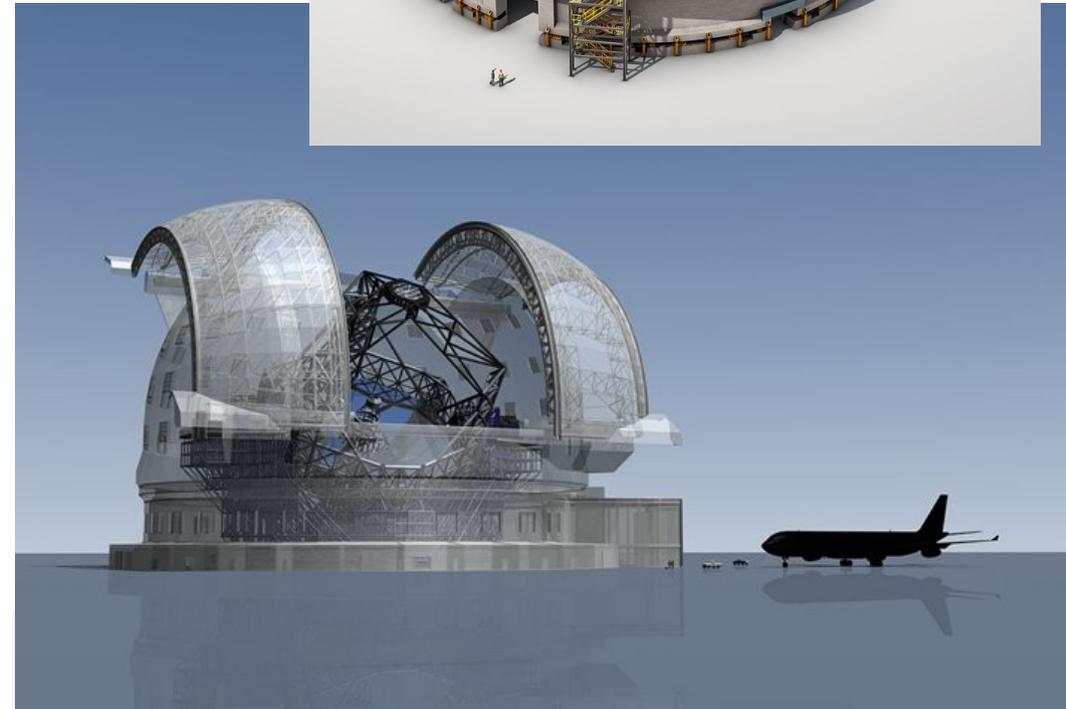
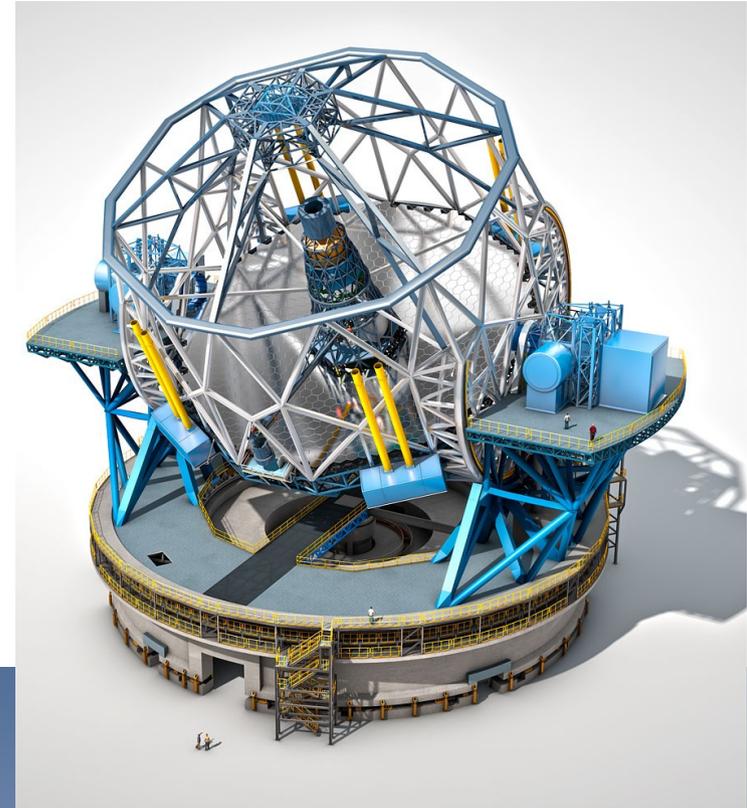
European HPC summit week 2016





European Extremely Large Telescope

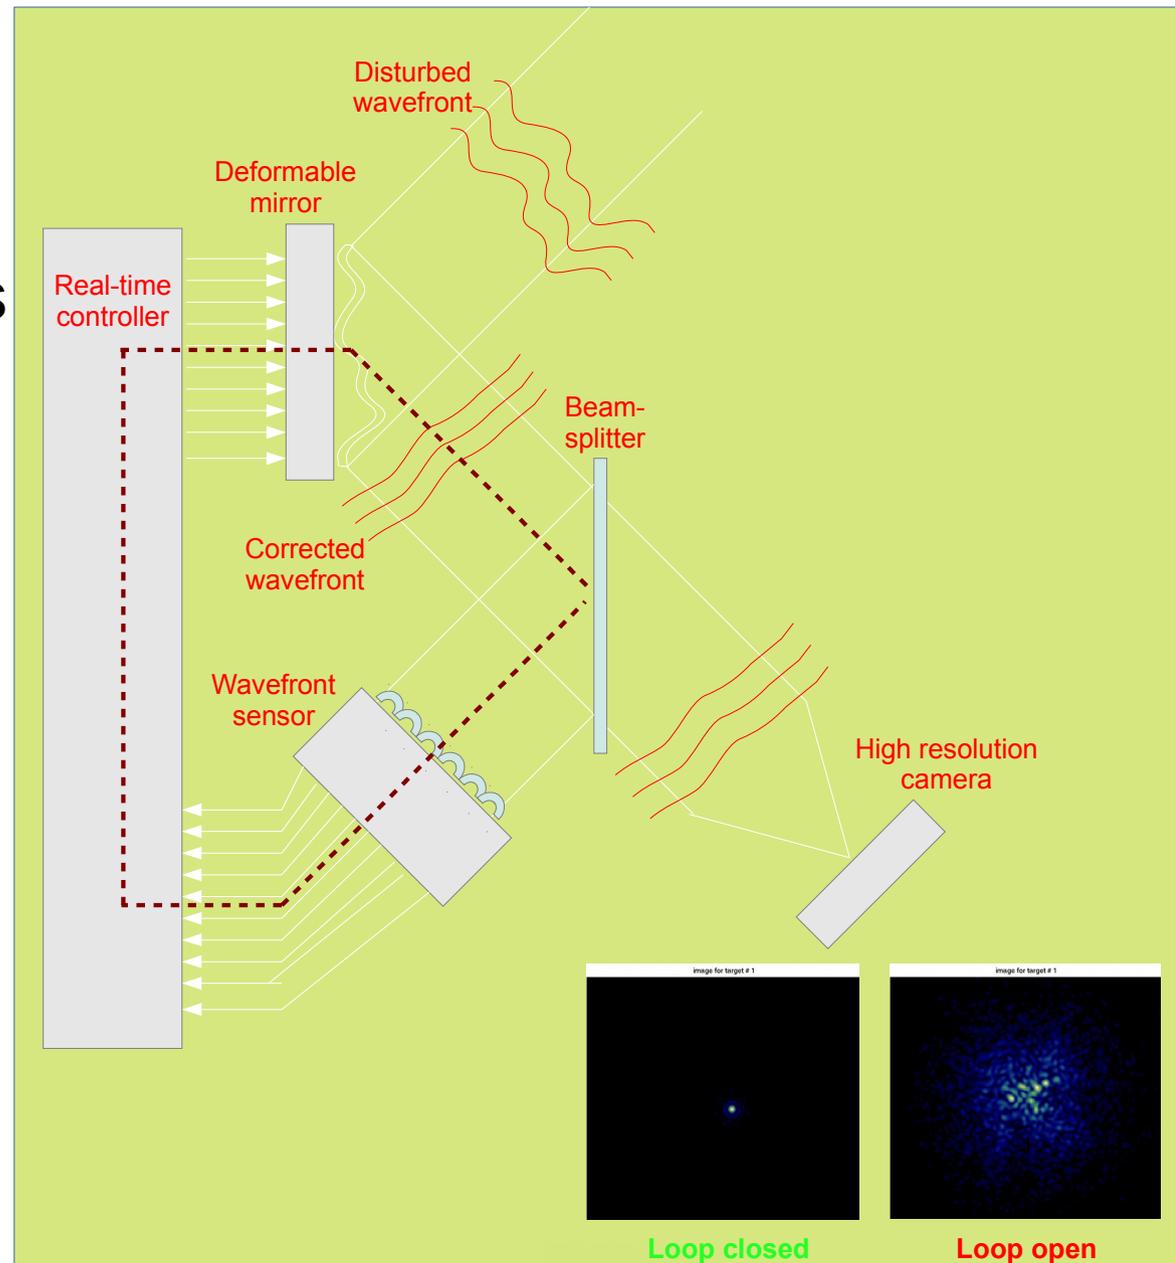
- 39m diameter telescope : x5 in diameter
=> x25 in system complexity
 - 100m dome, 2800 tones structure rotating @ 360°, seismic safe (Chile)
 - 1.2 G€ project, first light foreseen in 2024
 - Construction led by ESO (European Southern Observatory), international organisation funded by 15 European countries
 - Telescope components + science instruments built by european research labs + industrial partners





Adaptive optics

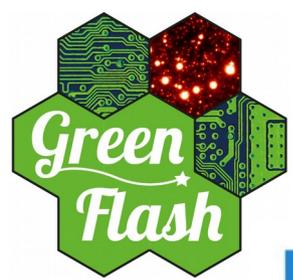
- Compensate in **real-time** the wavefront perturbations
- Using a wavefront sensor to measure them
- Using a deformable mirror to reshape the wavefront
- **Commands to the mirror must be computed in real-time (1ms rate)**



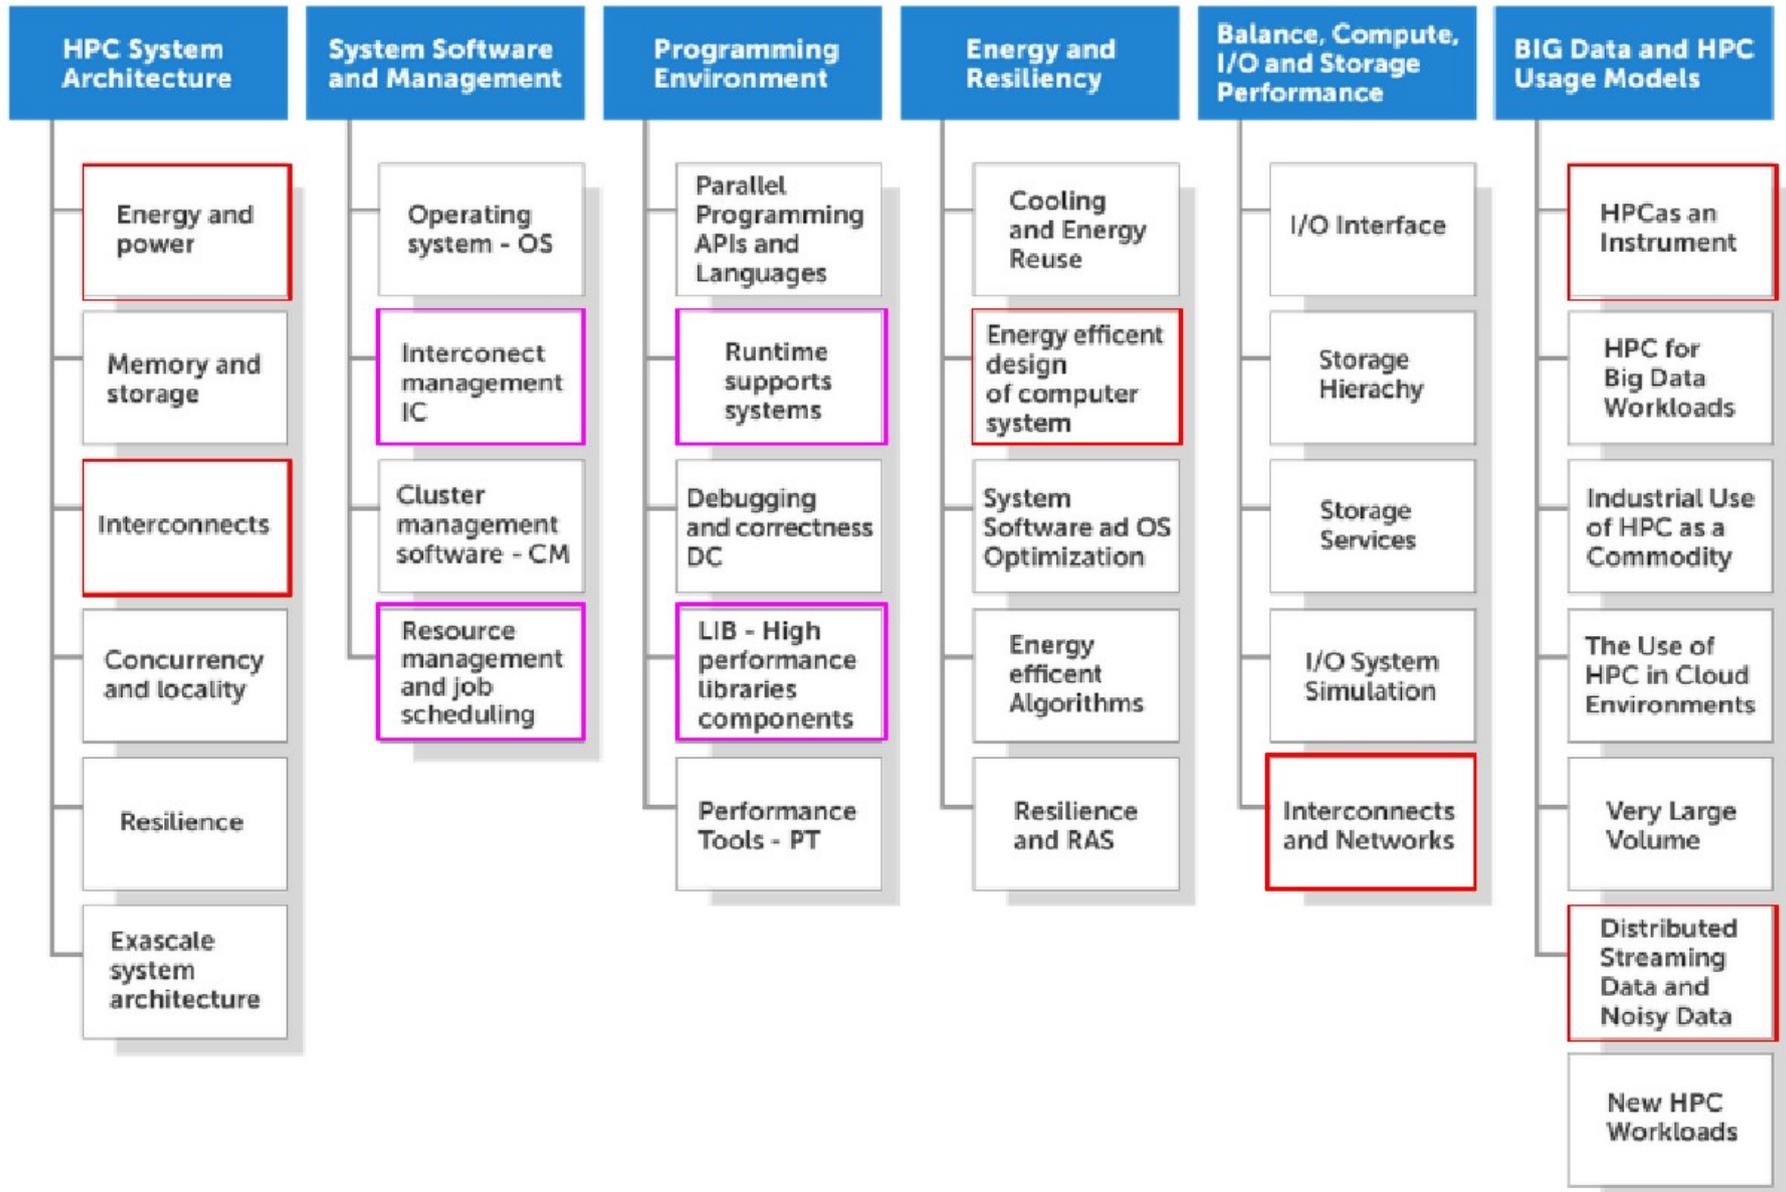


Introduction to Green Flash

- Program objectives: 3 research axes
 - 2 technological developments and 1 validation study
- Real-time HPC using accelerators and smart interconnects
 - Assess the determinism of accelerators performance
 - Develop a smart interconnect strategy to cope for strong data transfer bandwidth constraints
- Energy efficient platform based on FPGA for HPC
 - Prototype a main board, based on FPGA SoC and PCIe Gen3
 - Cluster such boards and assess performance in terms of energy efficiency and determinism
- AO RTC prototyping and performance assessment
 - Assemble a full functionality prototype for a scalable AO RTC targeting the MAORY system
 - Compare off-the-shelf solutions based on accelerators and new FPGA-based concept



Addressing HPC roadmap to exascale





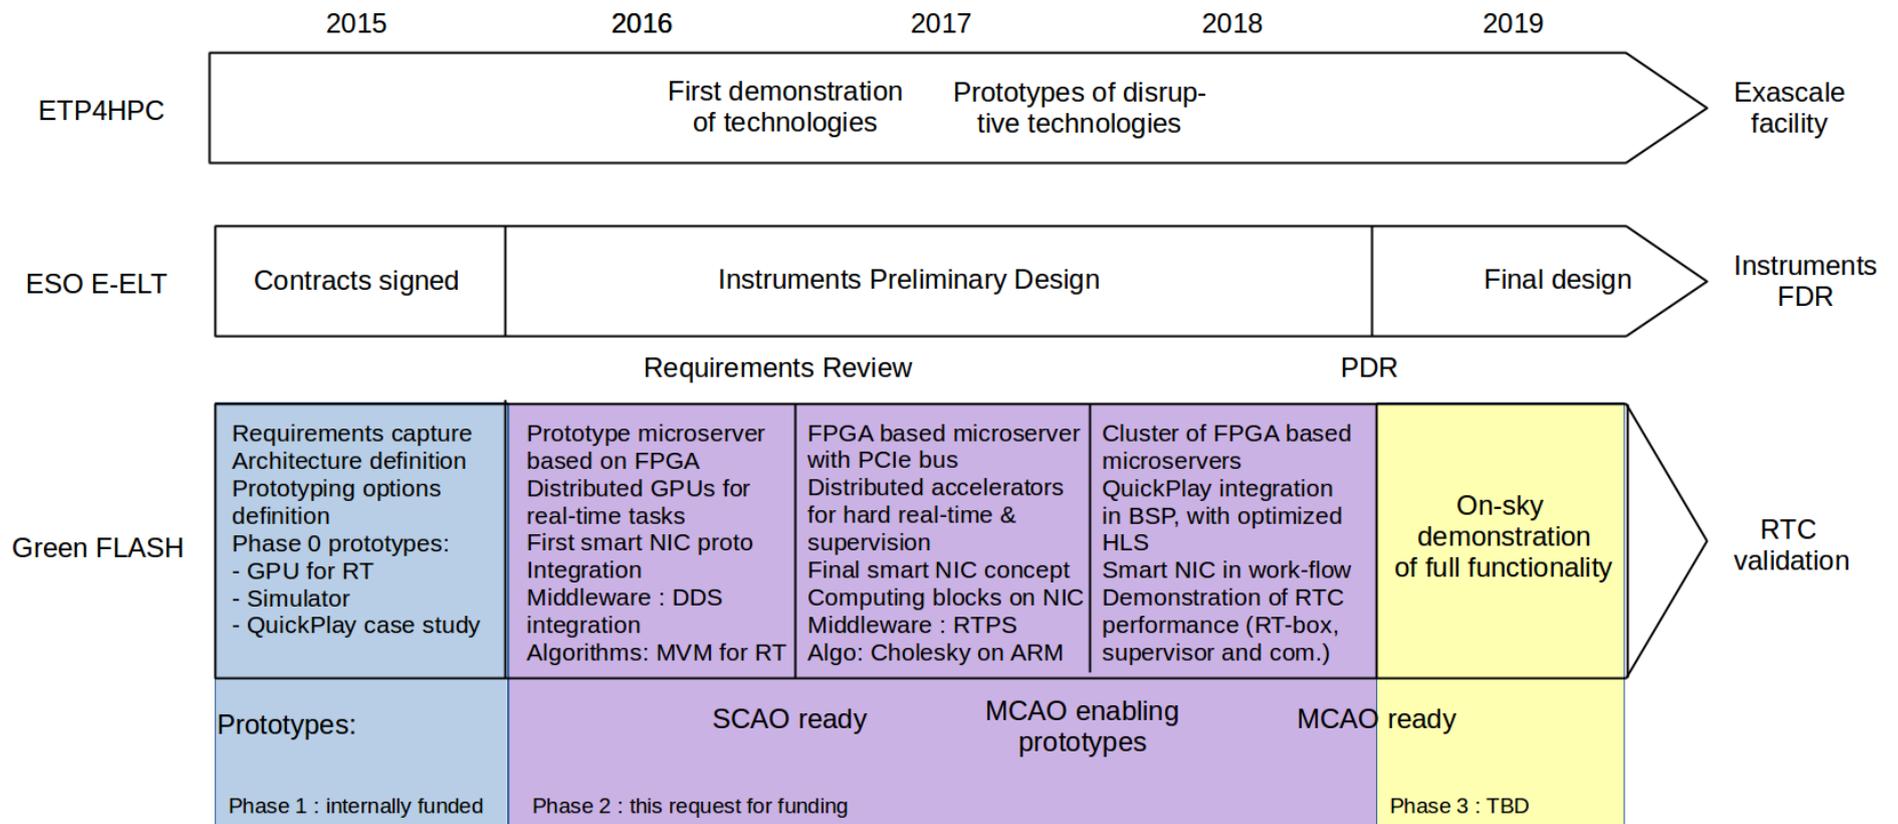
Green Flash project

- Partners
 - 2 academic partners
 - LESIA, Observatoire de Paris, P.I. Damien G.
 - CfAI, University of Durham
 - 2 industrial partners
 - Microgate : Italian SME designing FPGA solutions for various applications (including astronomical AO)
 - PLDA: french SME developing FPGA solutions (mostly IP cores, world leader in PCIe IP)
- 3.8M€ grant, 36 months work plan



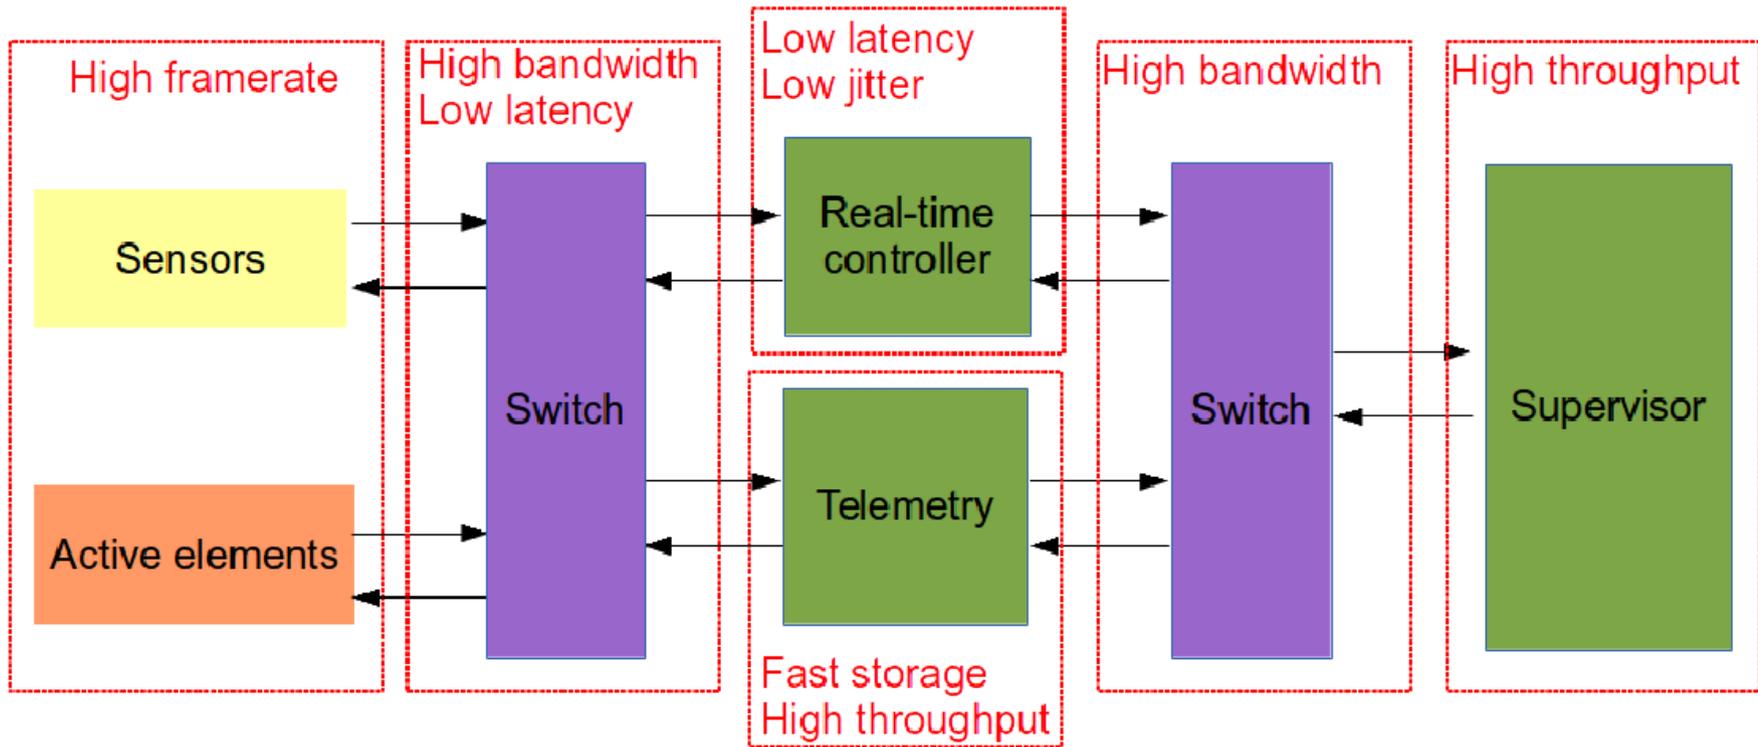
Project Management

- Good convergence with H2020 ETP4HPC / E-ELT project timeline



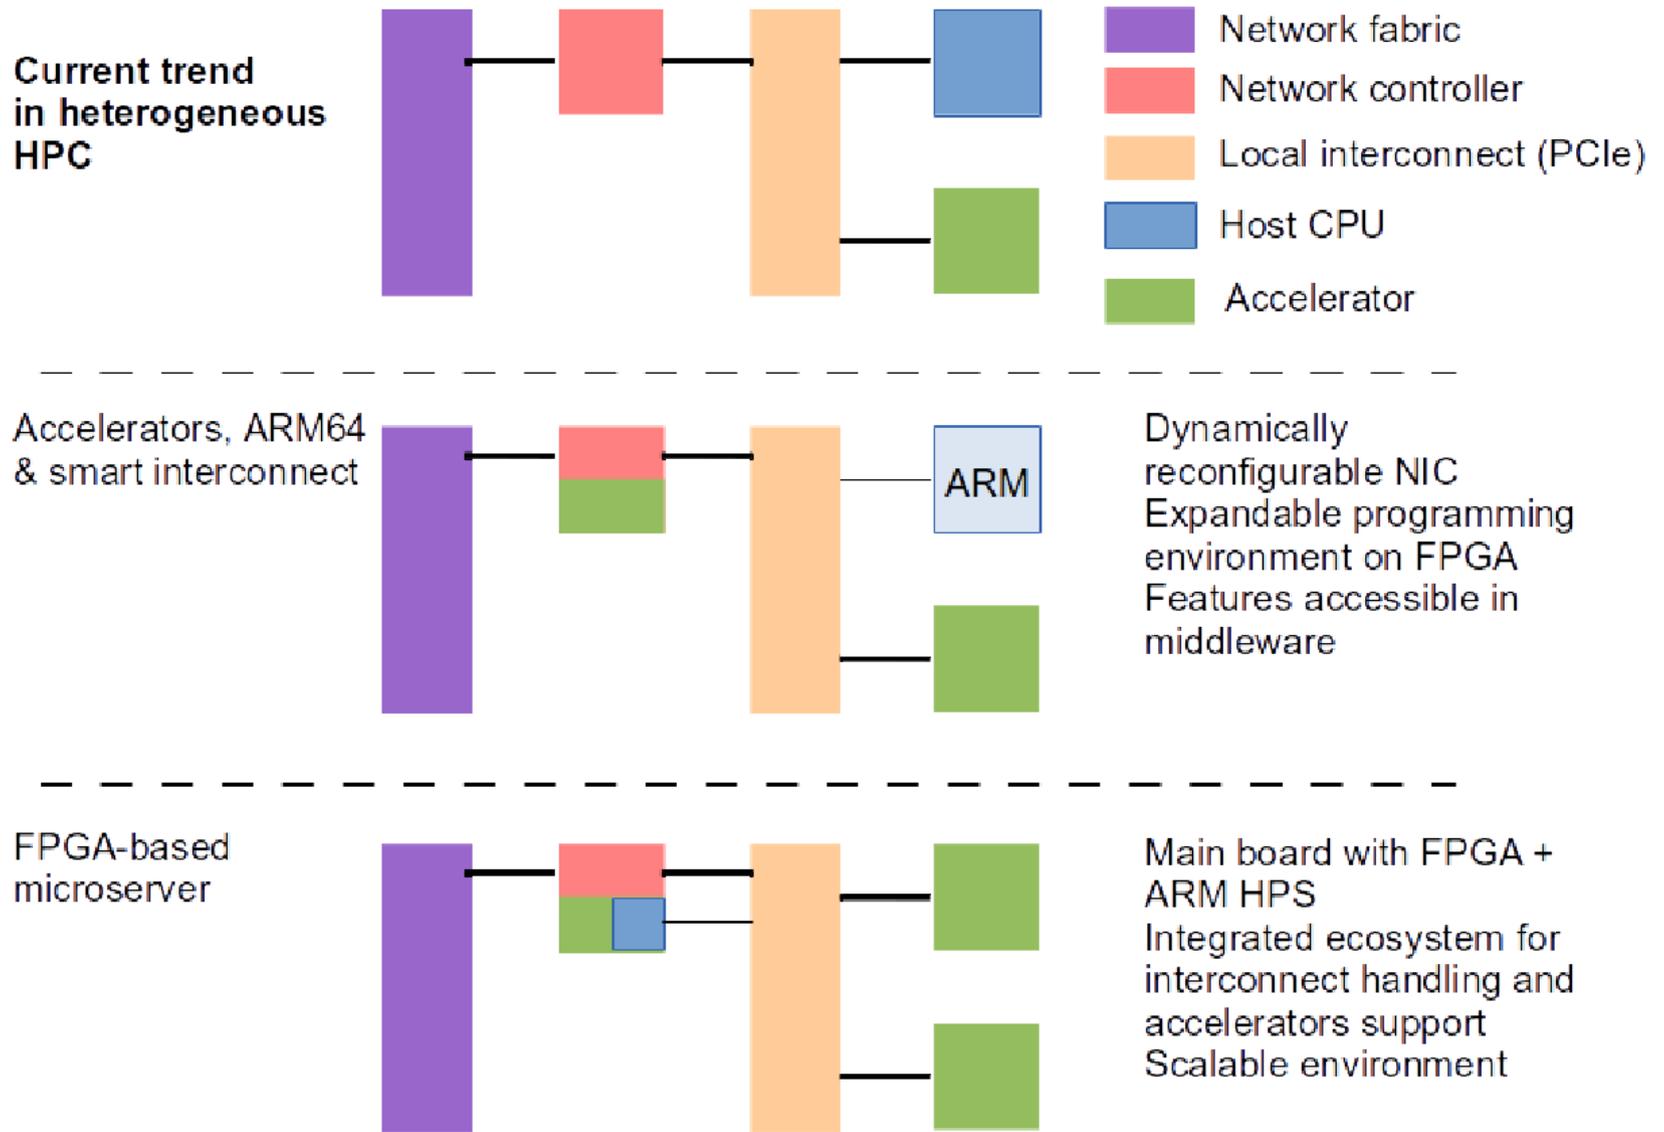


AO RTC concept



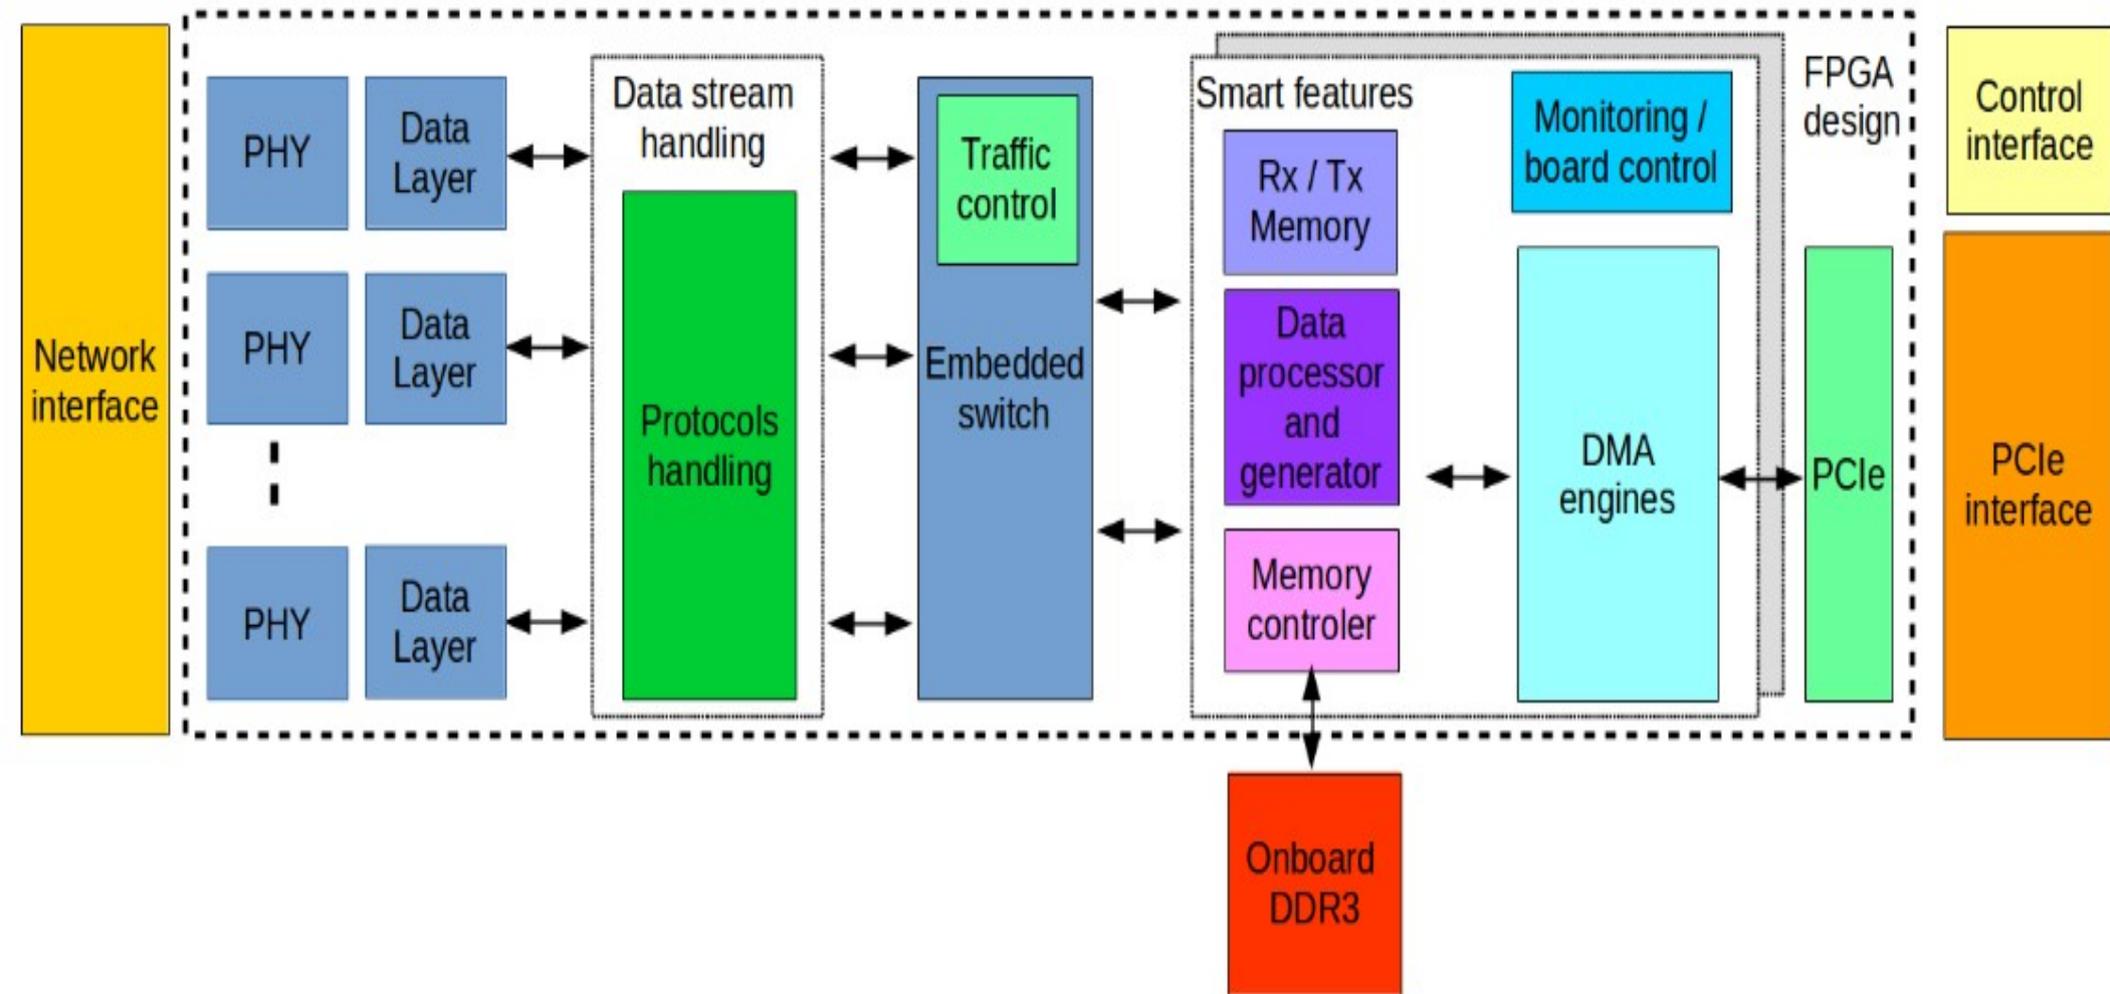


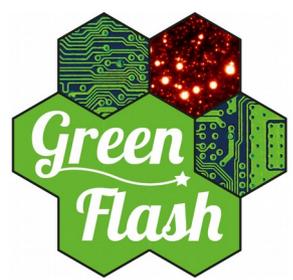
Assessing new HPC concepts





Smart Interconnect Subsystem Architecture





First prototype: hardware at work

- Reflex CES XpressKUS
 - Standard profile
 - XCKU040-2FFVA1156 (20 xceivers)
 - Up to PCIe 3.0 x8, with backward capability to PCIe 1.0 and 2.0
 - DDR3 SDRAM SODIMM support up to 8GB
 - HPC FMC interface
 - RoHS and REACH compliant
- Faster Technology FM-S14
 - Mezzanine Card (FMC) module
 - up to four SFP/SFP+ module
 - interfaces directly into Multi-Gigabit Transceivers (MGTs) of a Xilinx FPGA.



(* Reflex CES is a PLDA Group company)



Reducing Barriers to FPGA Adoption

QuickPlay Platform

- C/C++ Entry & Debug
(HDL Entry for Advanced Users)
- Hardware Abstraction
- 🏗️ Bridge the HW/SW Gaps



Ecosystem

- Boards, IP Cores, Libraries
- Value sharing
- 🏗️ QuickPlay as a channel.



Business Model

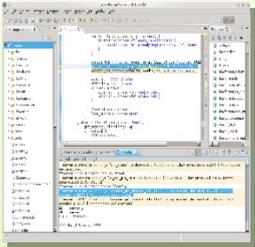
- Online Market Place
- Disruptive Licensing Mechanism
- 🏗️ Enable innovative business models by our partners



QuickPlay™
SDK



SW Application Development



**QuickPlay
Communication API**

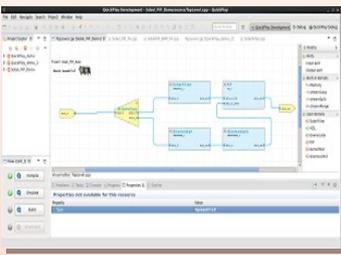
QuickStore

IP Cores

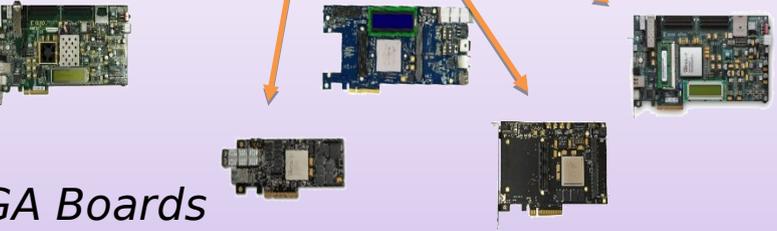


QuickPlay™

Unified FPGA IDE



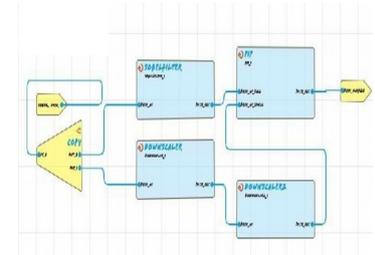
FPGA Boards



1

MODEL

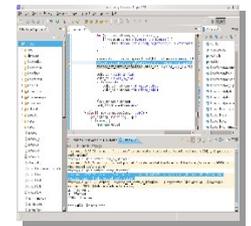
- C/C++ functional modeling



2

VERIFY & VALIDATE

- Desktop execution of system functional model



3

BUILD

- Hardware implementation: HLS, Logic Synthesis, P&R

4

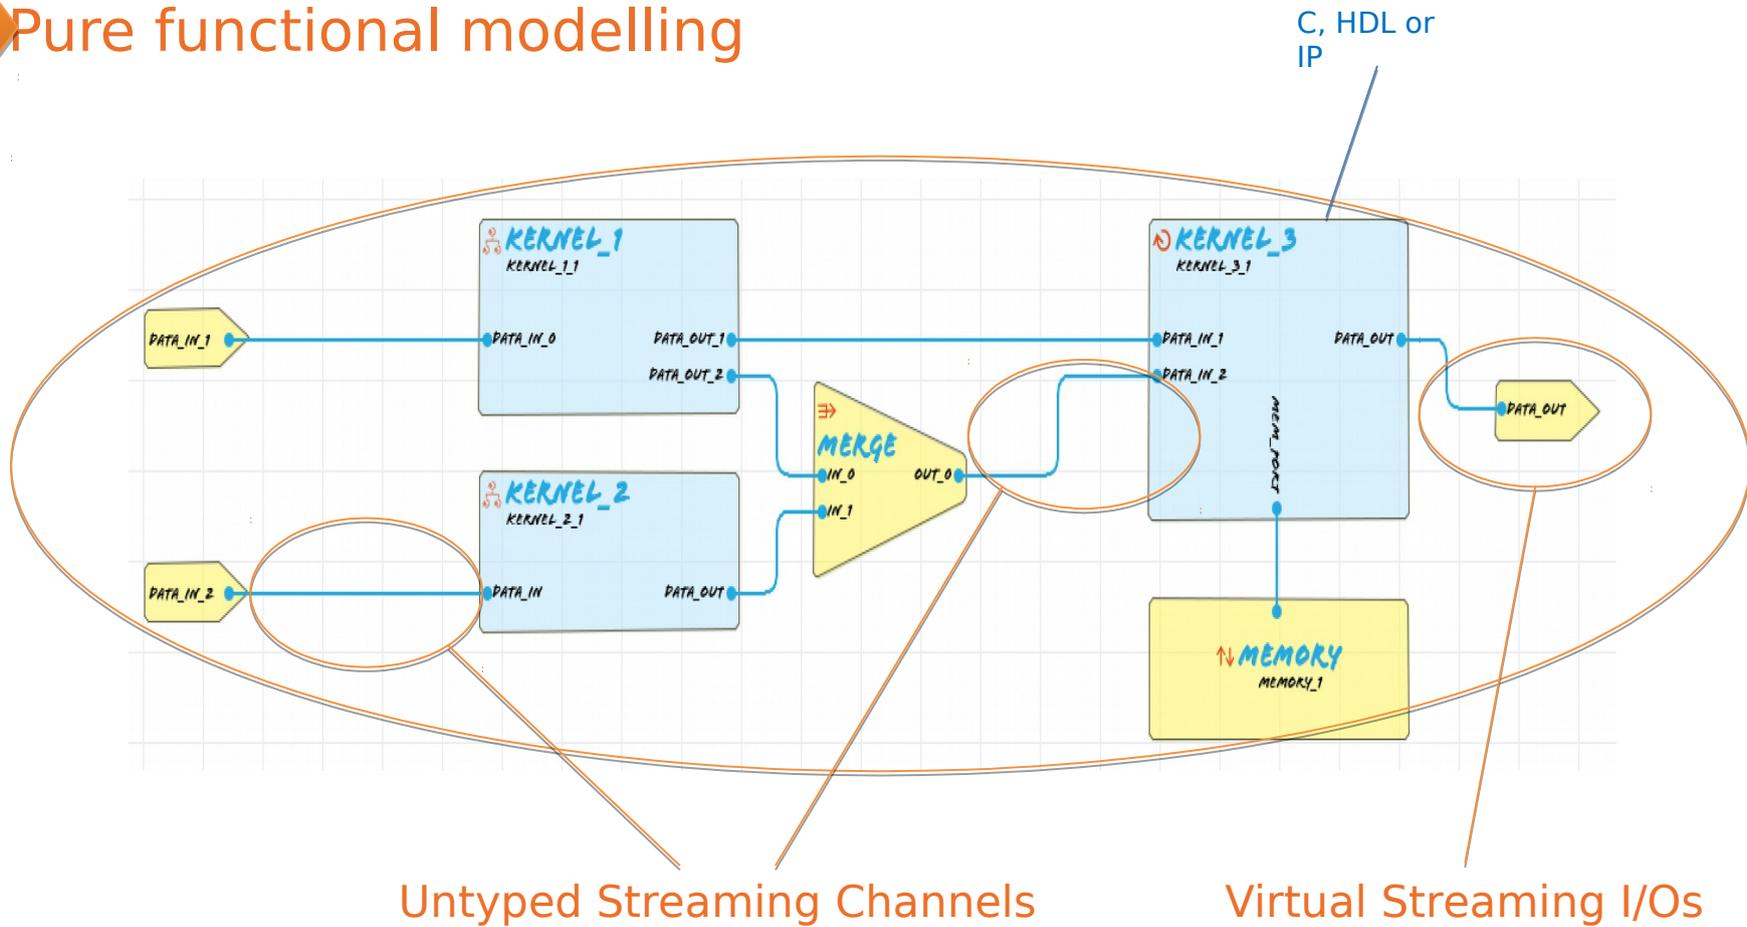
EXECUTE

- FPGA based system hardware execution



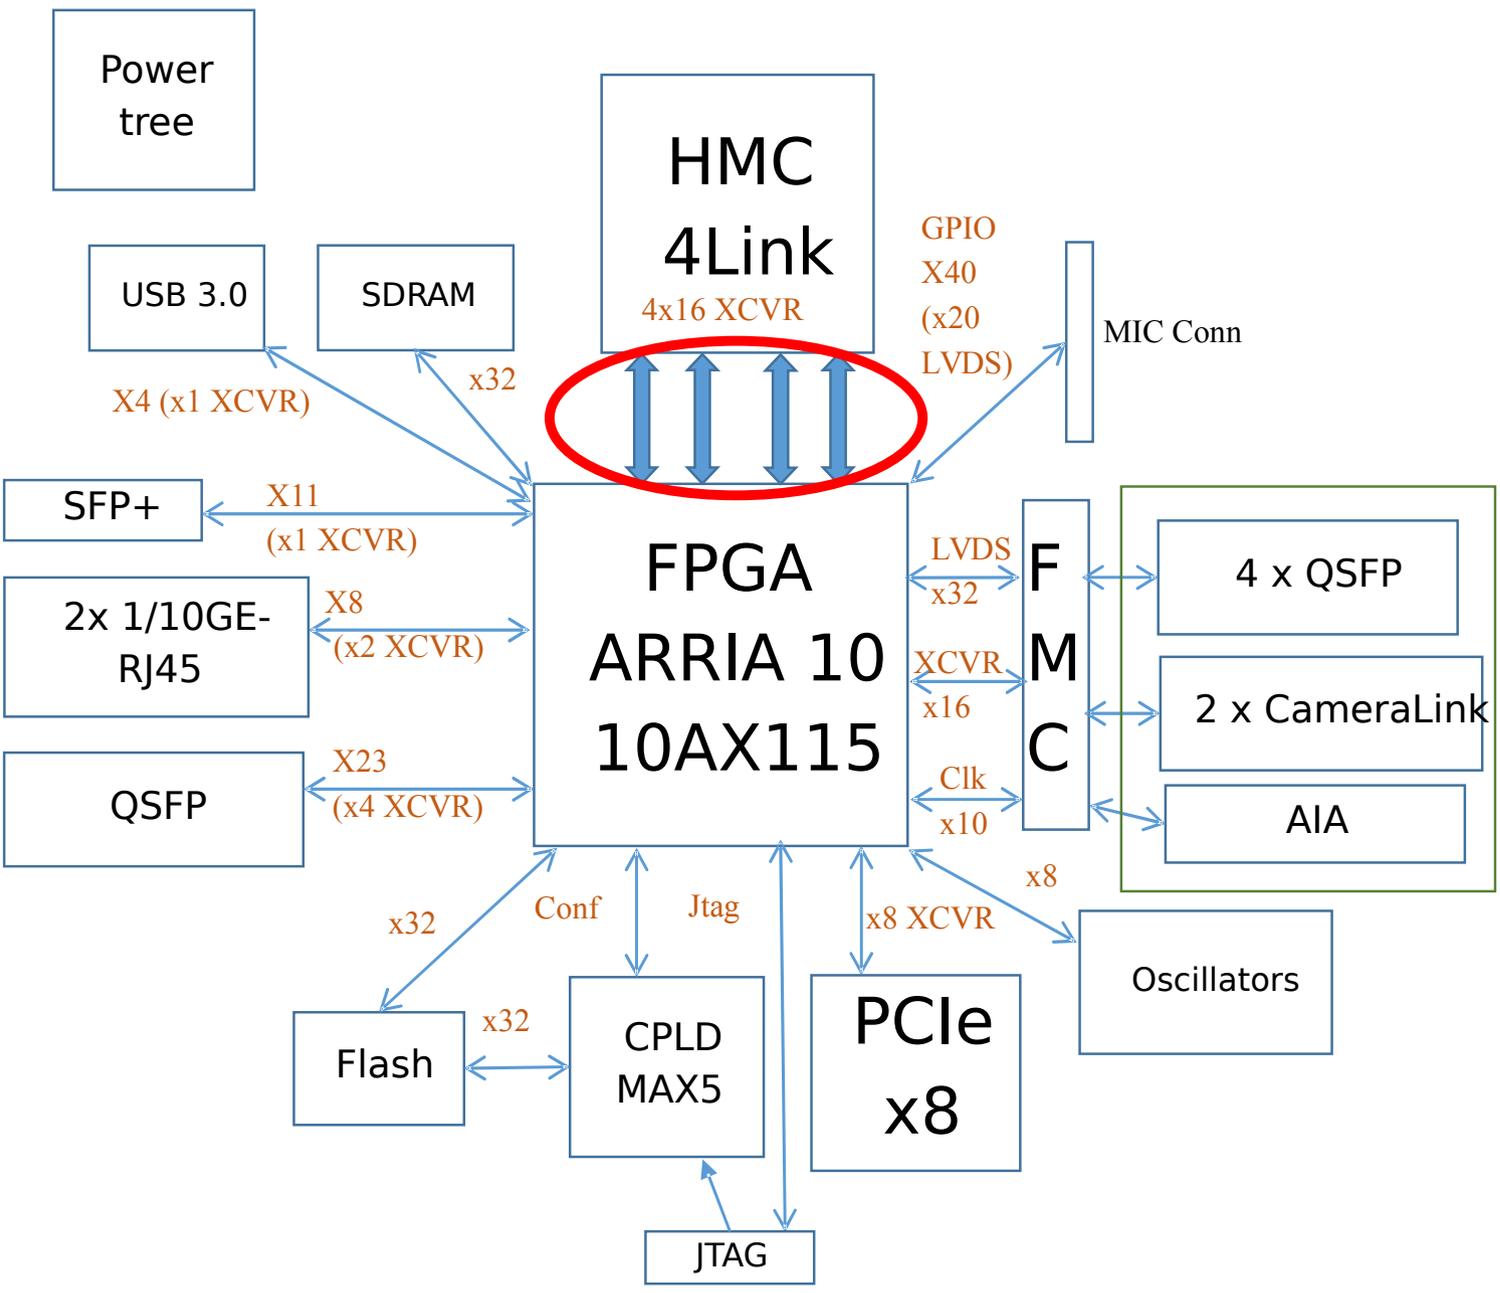
➔ “Design the way you think”

➔ Pure functional modelling





Scheme 1 - Single FPGA system



ARRIA 10AX115

- 1518 DSP blocks
- 6.6MB int. RAM
- max. 96 XCVR

PROS:

- Max. Bandwidth between HMC and FPGA - 4 Links
- Easier power tree because only one FPGA and HMC
- Shorter length of board possible than 2 FPGA version

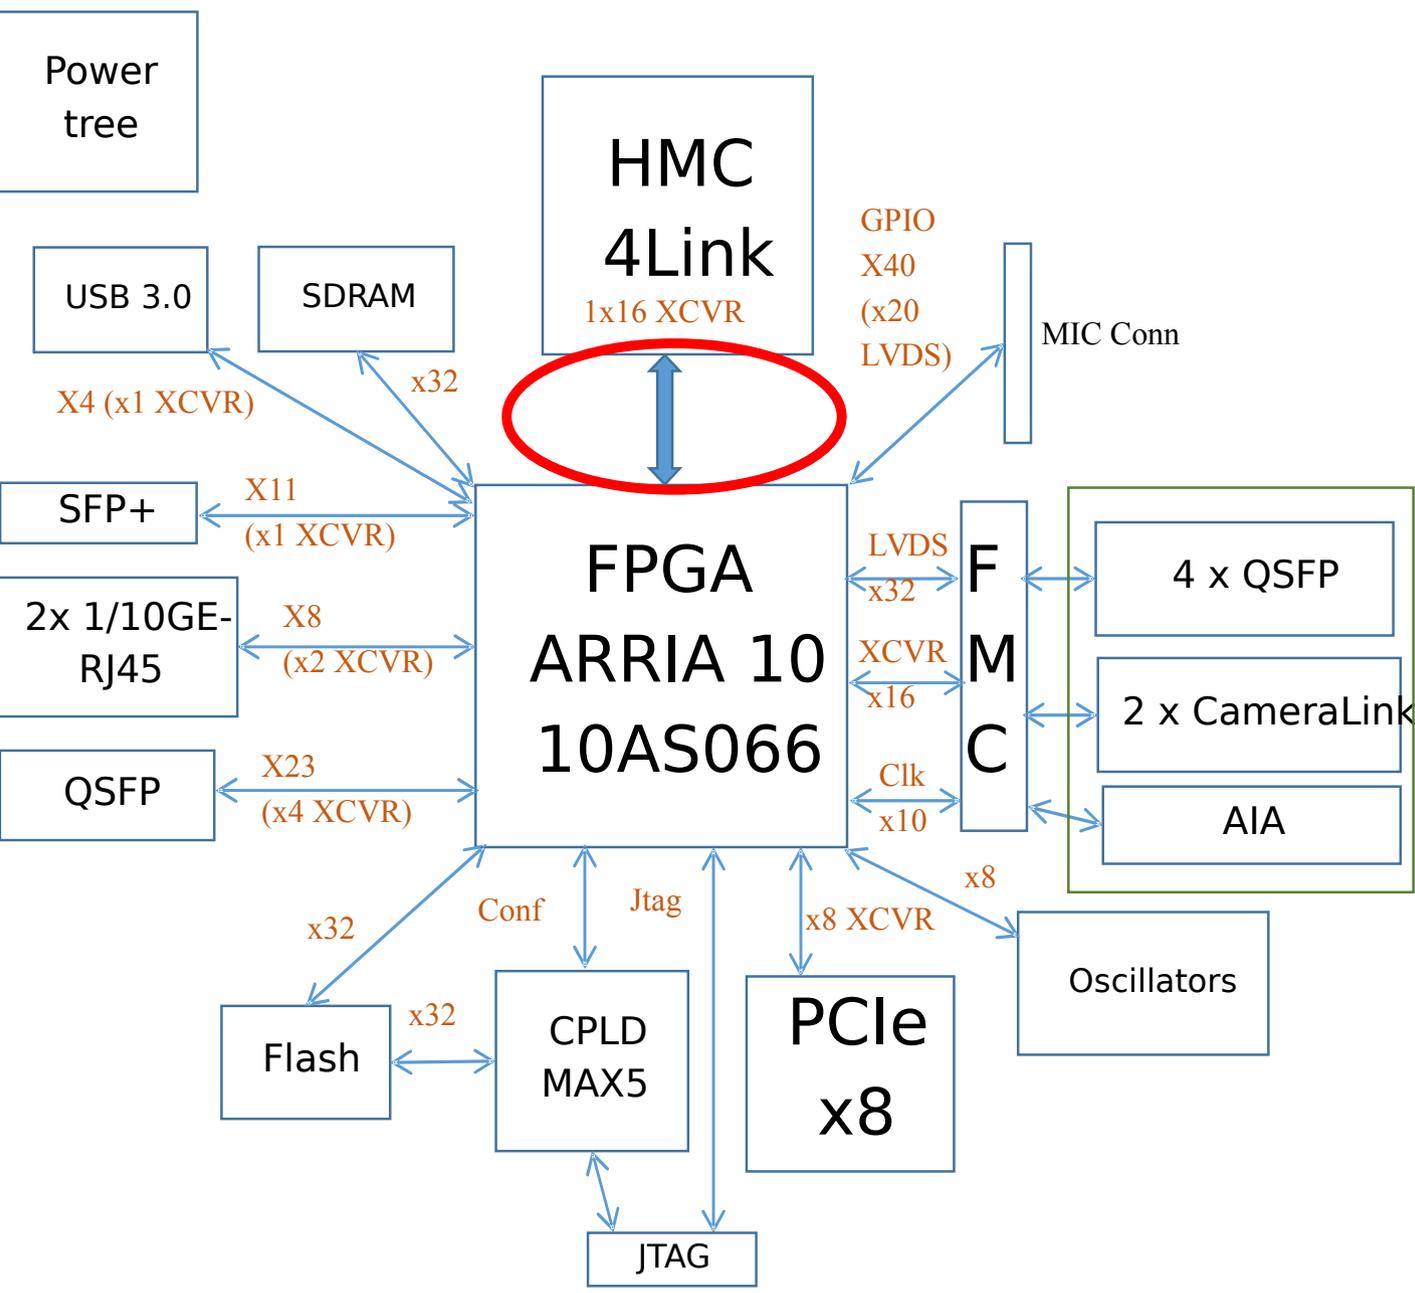
CONS:

- No ARM co-processor only NIOS II in logic
- No powerful PCIe root port because no ARM
- More expensive than scheme 2





Scheme 2 - FPGA with ARM coprocessor system



ARRIA 10AS066 SoC

- 1.5GHz ARM dual-core Cortex-A9 co-processor
- 1855 DSP blocks
- 5.2MB int. RAM
- max. 48 XCVR

PROS:

- ARM Co-Processor for stand-alone real-time box
- Powerful PCIe root port because of ARM and OS
- Shorter length of board possible than 2 FPGA version
- Less expensive than Scheme 1

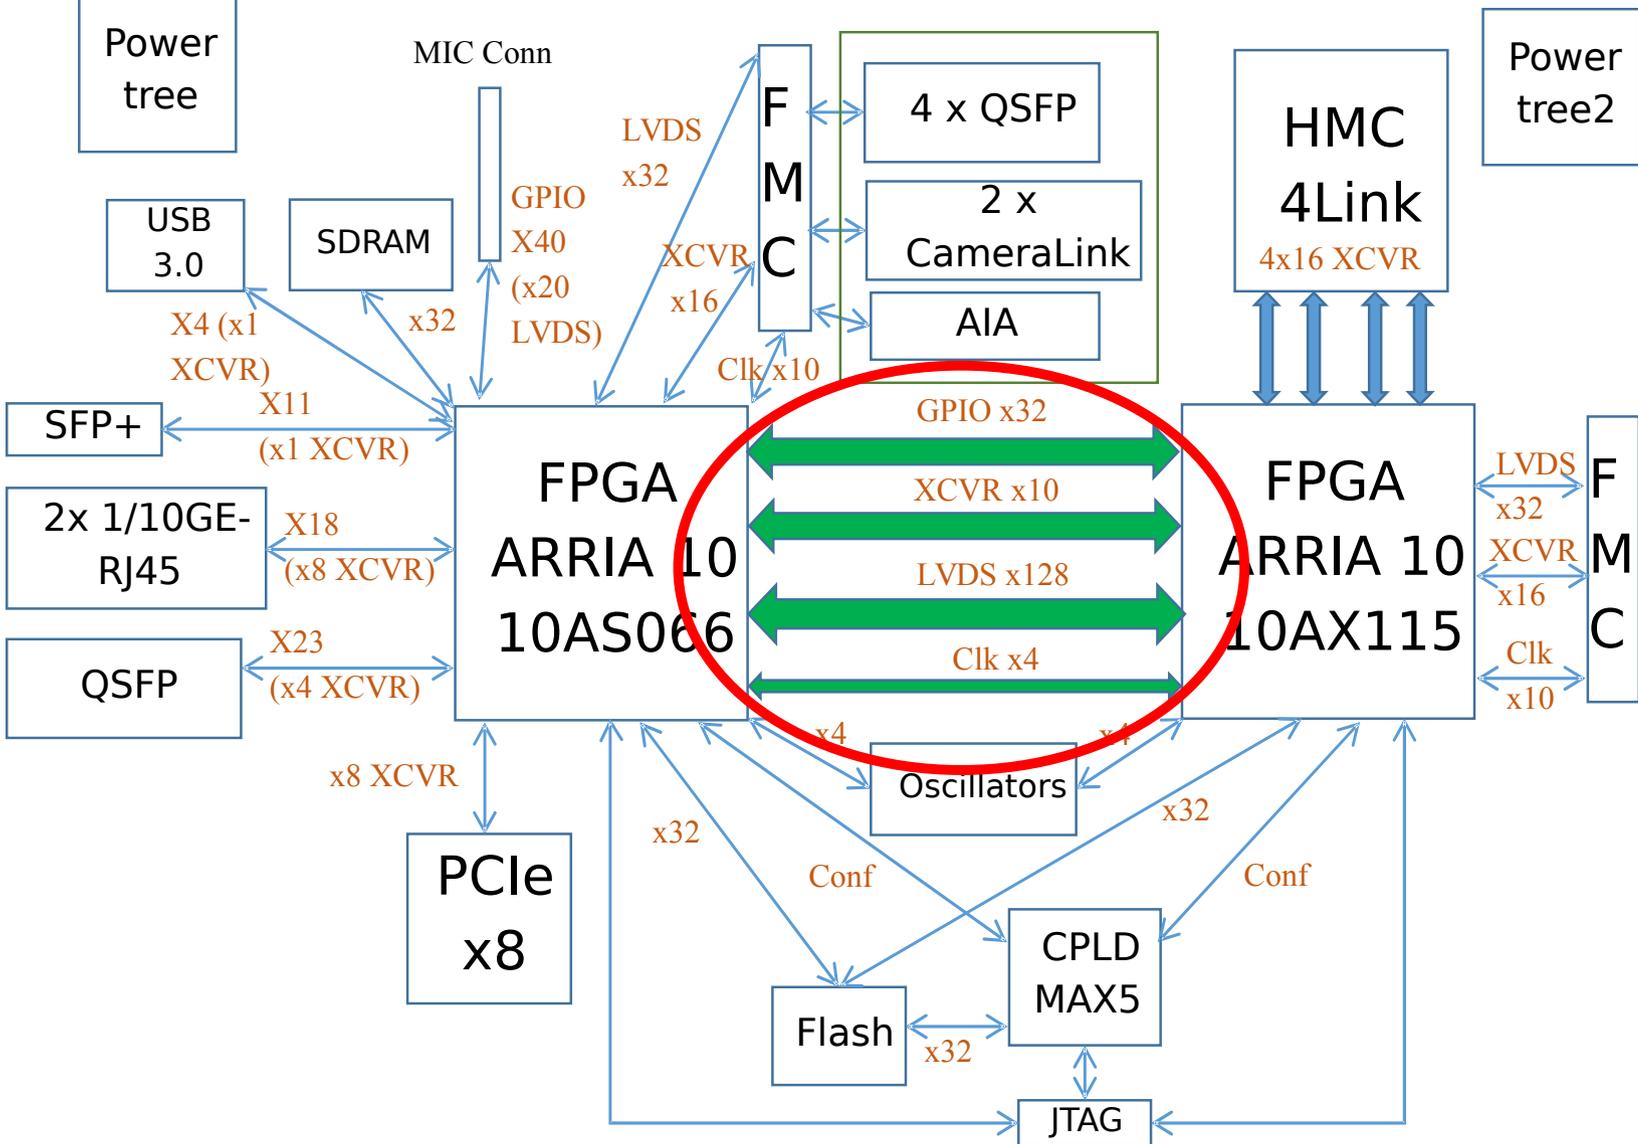
CONS:

- Less Transceiver
- 1/4 Bandwidth between HMC and FPGA - 1 Link
- Reduced computation because ext. mem bottleneck, less int. mem





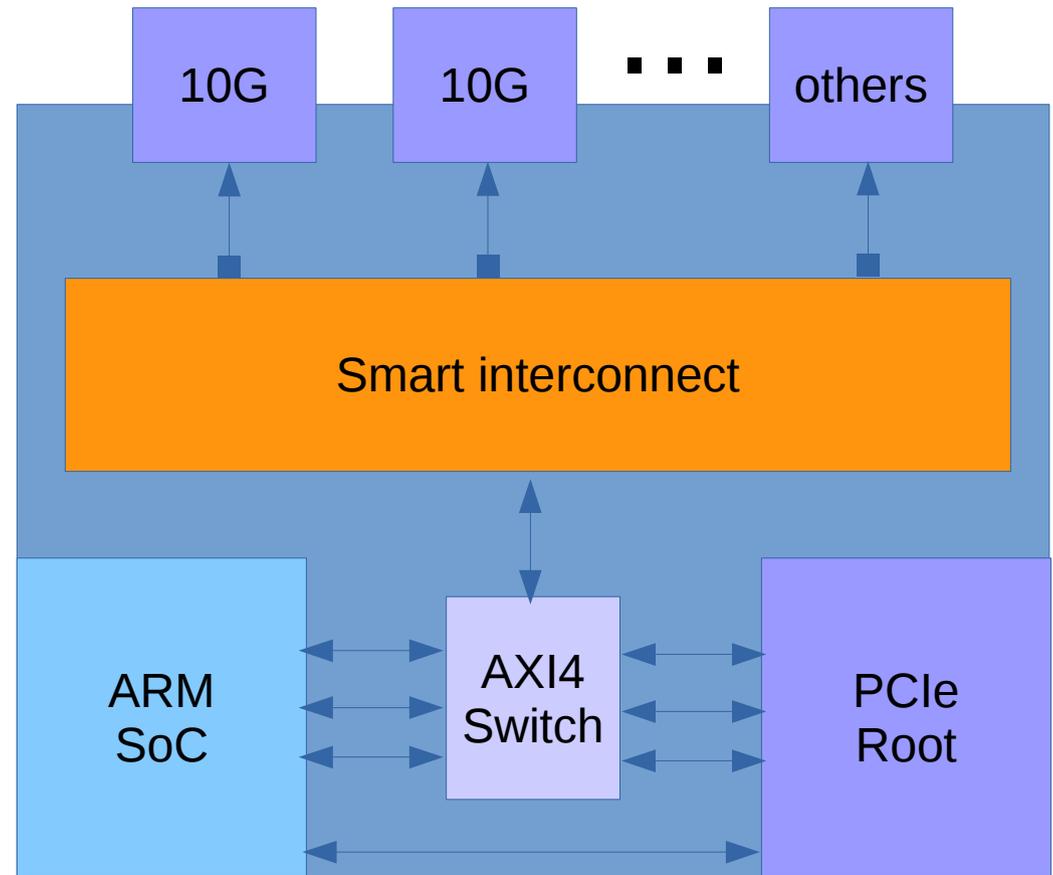
Scheme 3 - Two FPGA System





FPGA design of microserver board

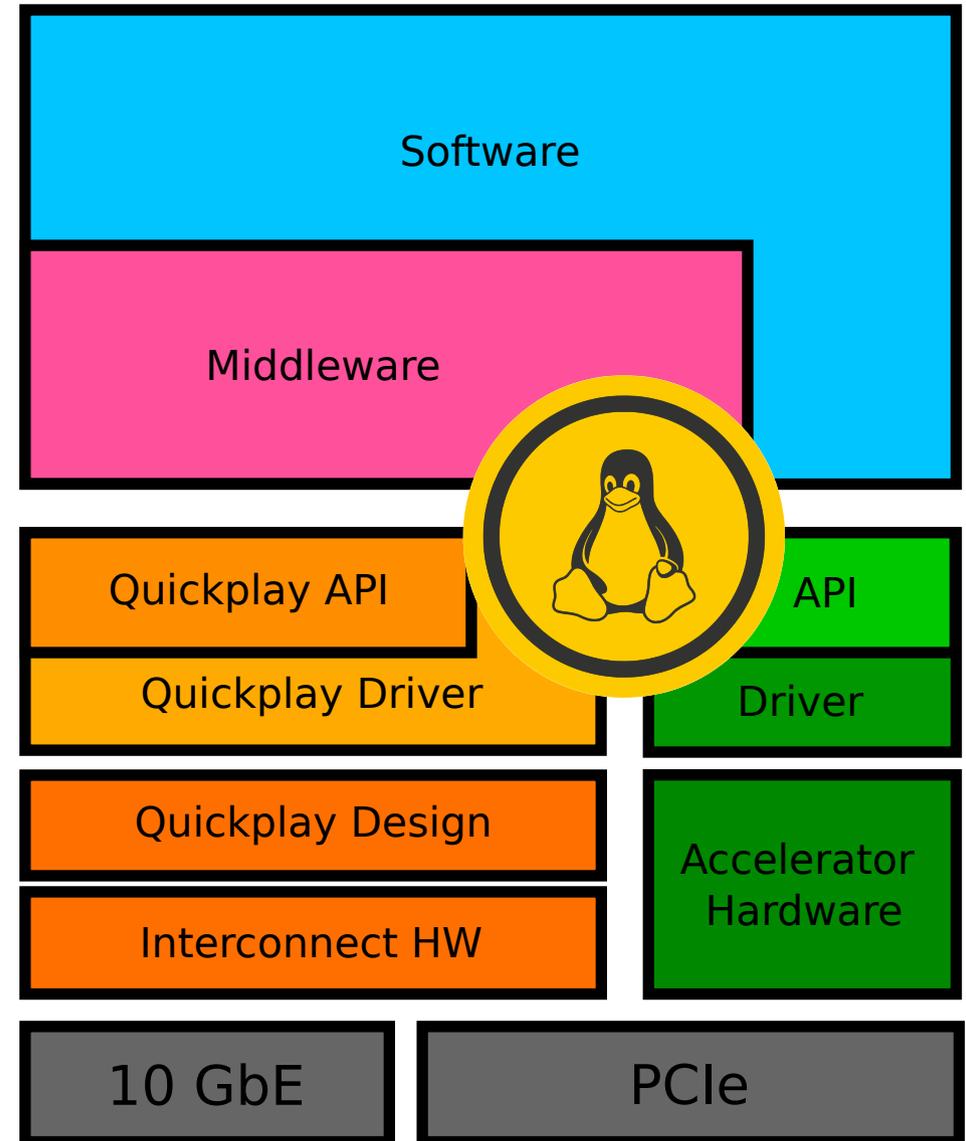
- Integration of Quick Play
- Coordination of efforts between 2 european SMEs
- Bring together smart interconnect concept with energy efficient compute platform





SW / MW stack

- Under development at academic partners
- Run a standard HPC ecosystem on the prototype board
- Performance assessed w/r AO application (mainly linear algebra)





Working on the application side

- 2016 Gordon Bell award submission : 3.5 PFLOPs (60% of peak perf) on custom application

Extreme Scale Simulations of Multi-object Adaptive Optics for the European Extremely Large Telescope

Hatem Ltaief¹, Damien Gratadour^{2,3}, Ali Charara¹, Saber Feki⁴, Eric Gendron², Bilel Hadri⁴, Ahmad Abdelfattah⁵, and David Keyes¹

¹Extreme Computing Research Center, KAUST, Saudi Arabia

²LESIA, Observatoire de Paris, France

³Université Paris Diderot, France

⁴KAUST Supercomputing Lab, KAUST, Saudi Arabia

⁵Innovative Computing Laboratory, Knoxville TN, USA

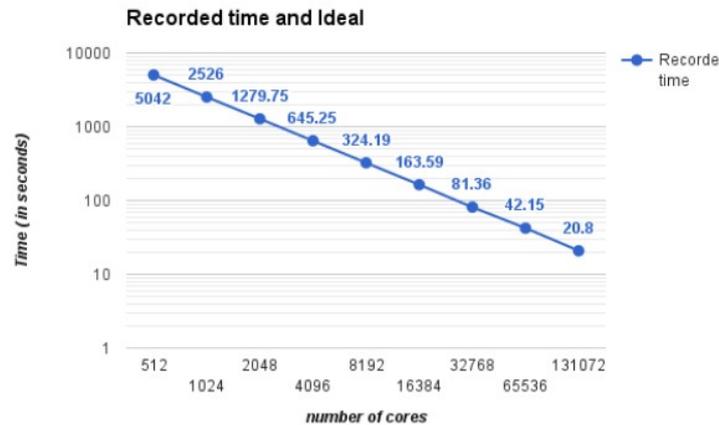
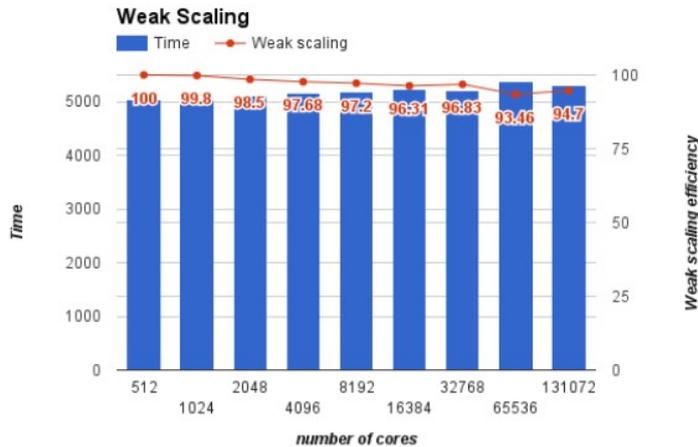
ABSTRACT

The European Extremely Large Telescope (E-ELT) is one of today's most challenging projects in ground based astronomy. Addressing one of the key science cases for the E-ELT,

1. PERFORMANCE ATTRIBUTES

Category of achievement	Scalability Time-to-solution Peak performance
Type of method	Direct solve
Results	Whole application
Precision	Double precision
System scale	two third of full system
Metrics	Time Flops/s

Table 1: List of Performance Attributes.



(a) Nearly perfect weak scalability of the code from 512 up to

(b) Time to solution per snapshot: weak scalability from 512



That's it for today ! Thank you

