



exploring Manycore Architectures for Next-GeneratiOn HPC systems

Deeply Heterogeneous HPC Systems with QoS Guarantees for Europe

Prague, EsD Workshop, May 2016

José Flich, Universitat Politècnica de València



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671668






Outline

- MANGO project and Consortium
- Capacity Computing
- Interconnect and heterogeneity as Enablers
- MANGO prototype as input
- Our plans


MANGO project and Consortium

- **MANGO: *exploring Manycore Architectures for Next-GeneratiOn HPC systems***

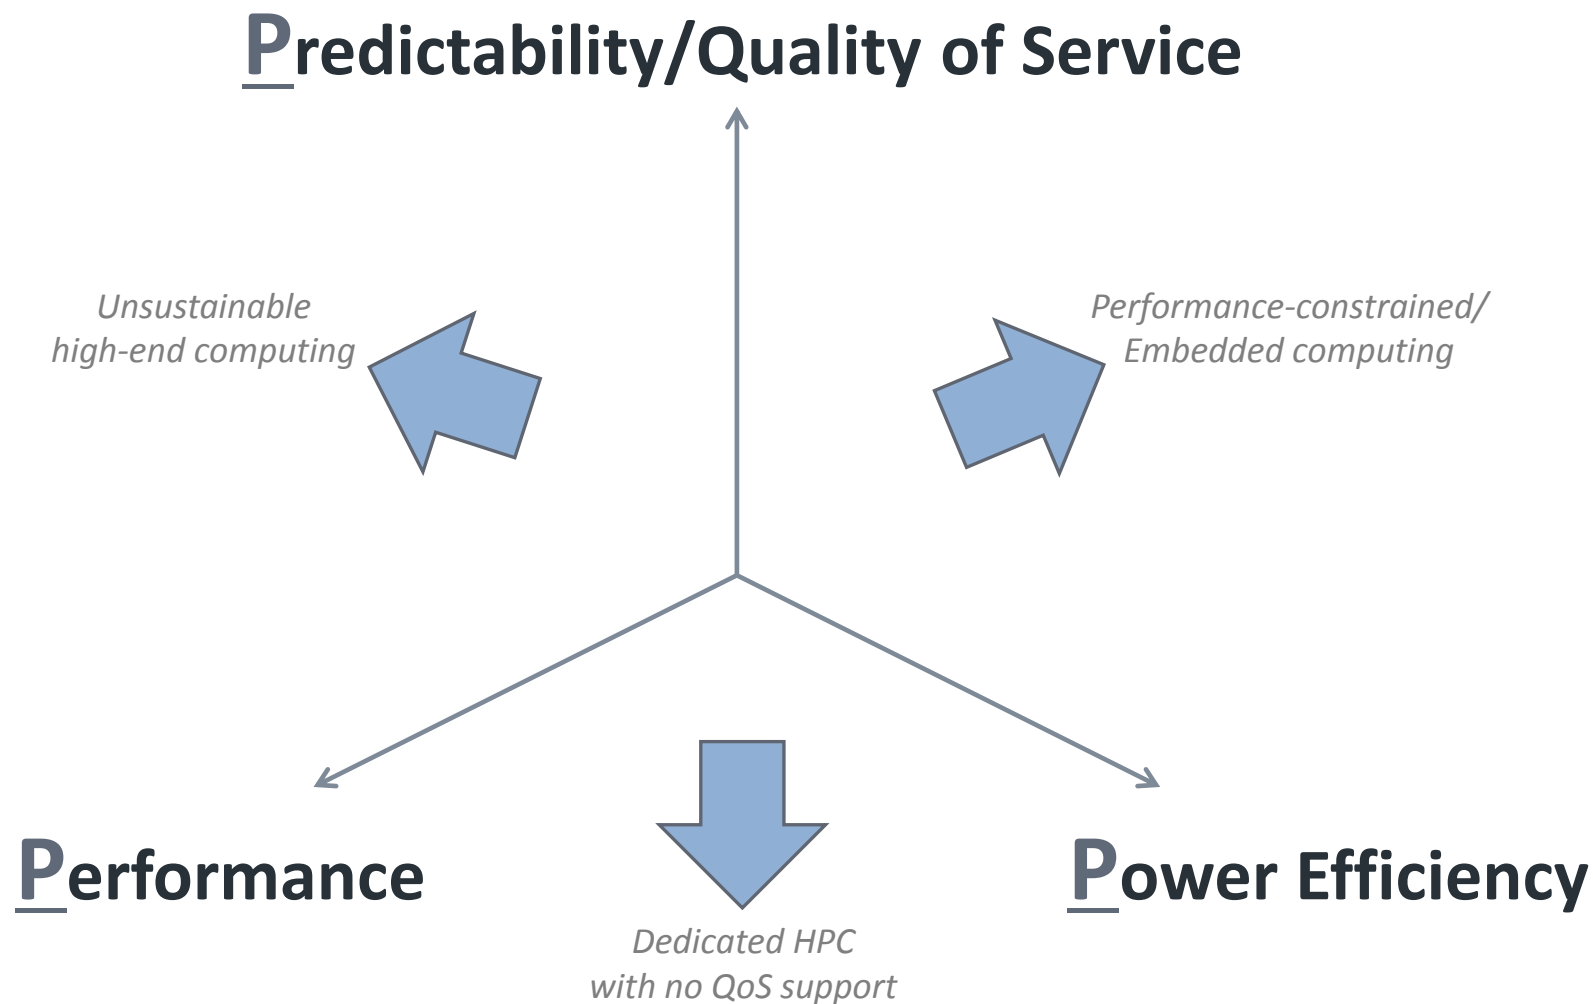
- started Oct. 2015, budget \approx 6M€
- Currently at M8



- **Universitat Politècnica de València (SPAIN)**
- **CeRICT / University of Naples (ITALY)**
- **Politecnico di Milano (ITALY)**
- **Zagreb University (CROATIA)**
- **Pro Design GmbH (GERMANY)**
- **Thales Communication & Security (FRANCE)**
- **EPFL (SWITZERLAND)**
- **Philips Medical Systems (NETHERLAND)**
- **Eaton Industries SAS (FRANCE)**







MANGO: Exploiting the *PPP* design space




Cases for time requirements

- Financial applications (high frequency trading)
- Biomedical applications
 - realtime biomedical diagnosis
- Multimedia applications
 - transcoding
- Video surveillance/security
- BigData is merging with HPC
 - CyberPhysical Systems
 - Smart Cities
 - Internet of Things



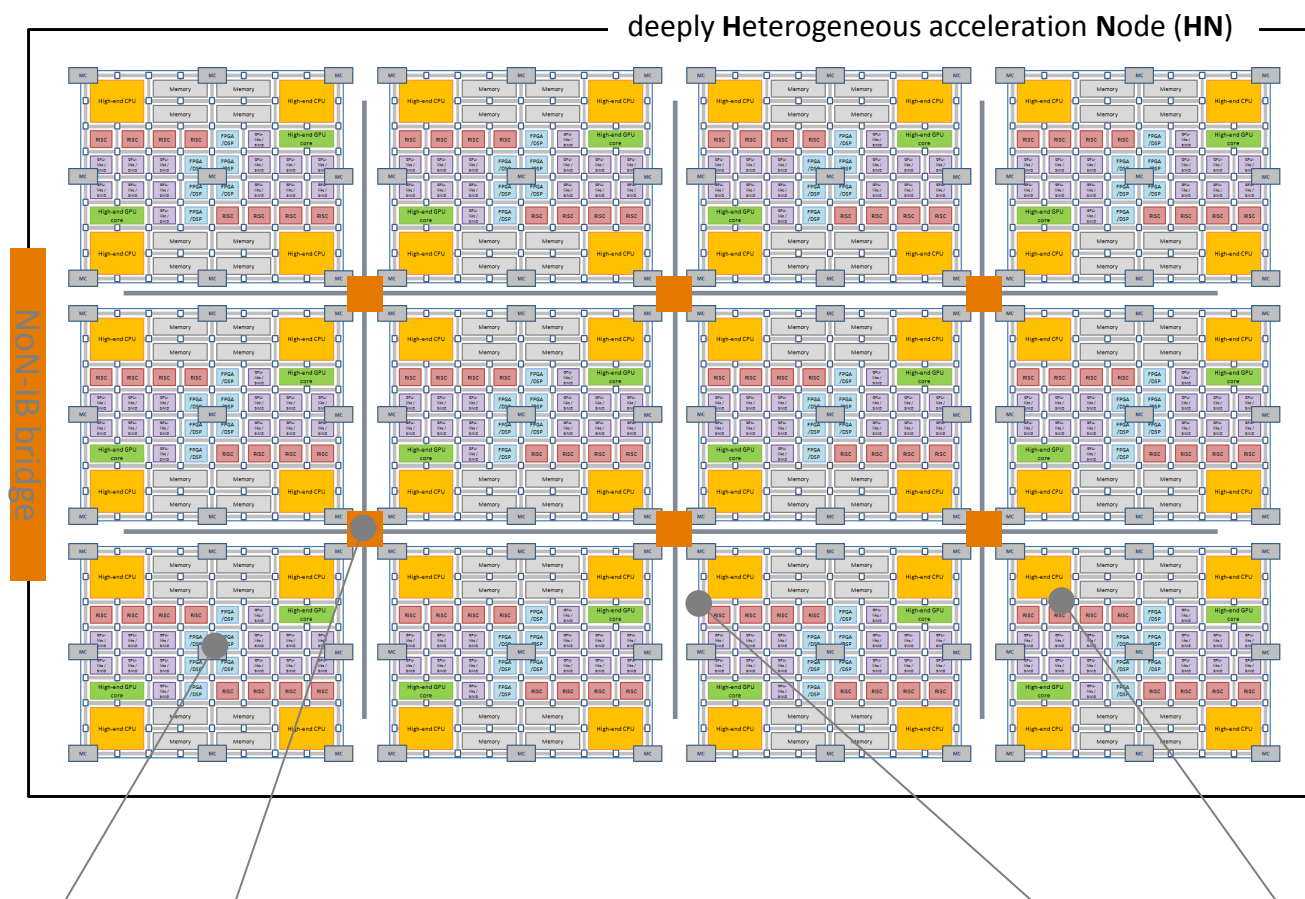
Capacity Computing

- Capacity Computing
 - Run as many applications as possible
 - Many application instances (users, data sets)
- Capability Computing
 - Run an application as fast as possible



New requirements

- Security
 - Guaranteeing applications do not interfere
- Virtualization of resources
 - Guaranteeing resources used exclusively by applications
- **Network** and **heterogeneity** to guarantee efficient capacity computing
- The network is at central to this approach
 - Needs to guarantee means of partitioning, reconfiguration, and isolation
 - Needs to guarantee proper bandwidth and latency allocation
 - Needs to guarantee jitter bounds

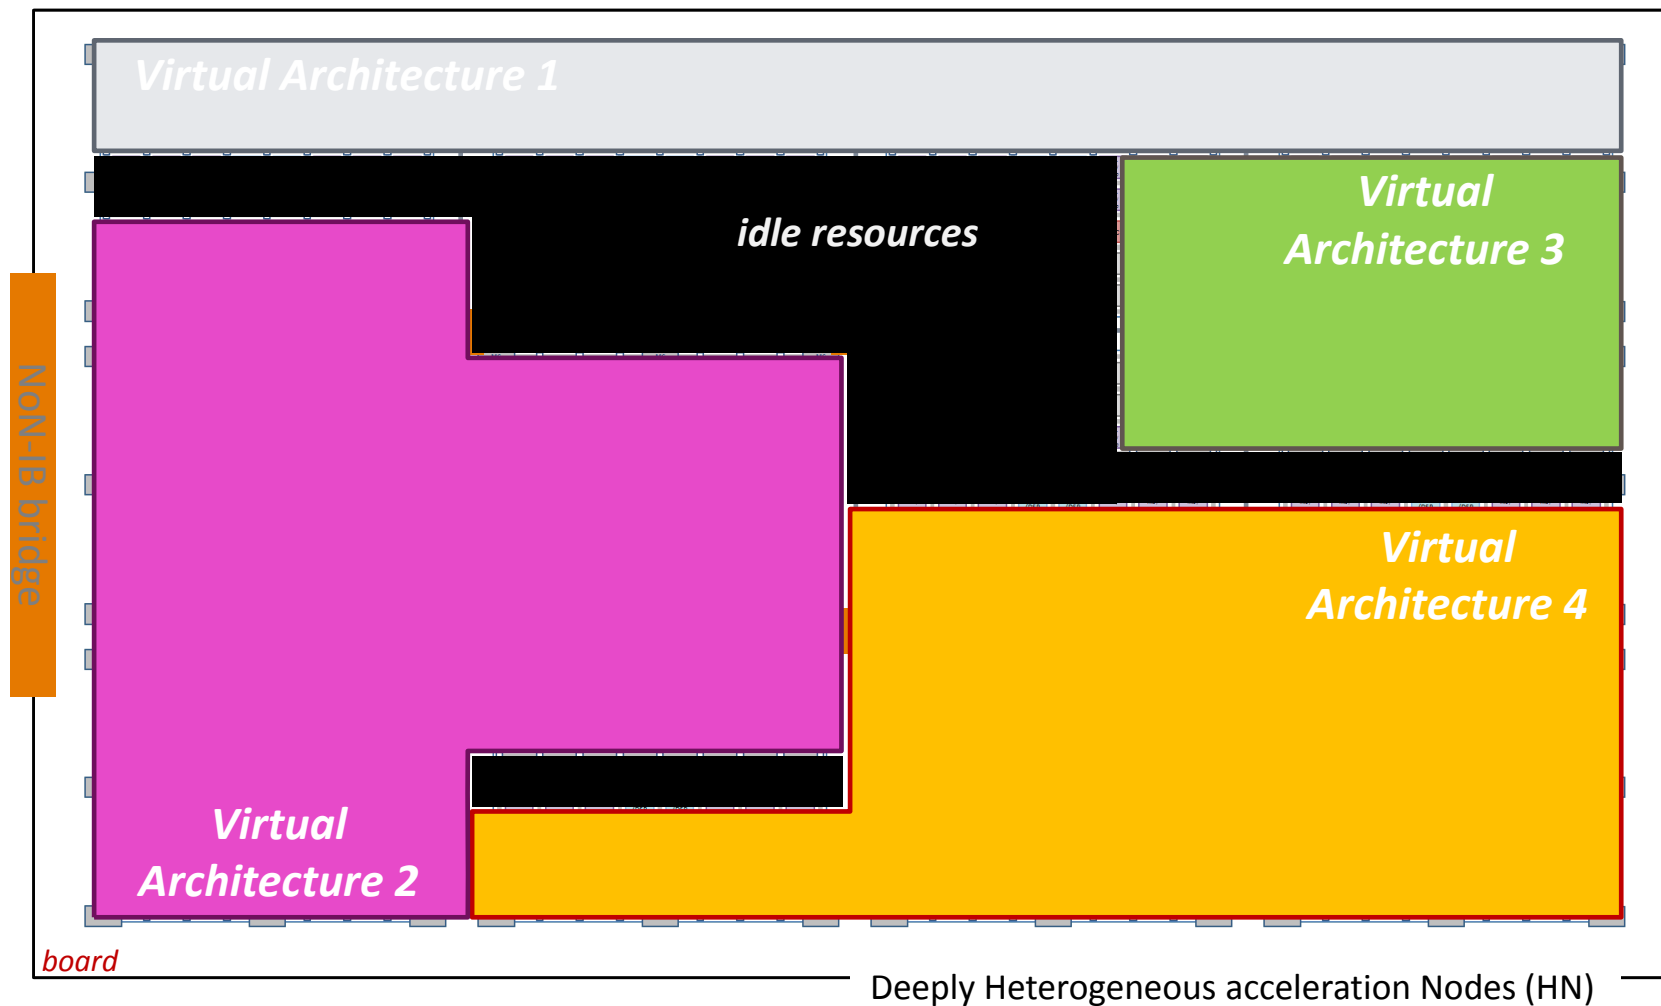


Advanced on-chip interconnect (**NoC**)

Low-power RISC cores coexist with massively-parallel
heterogeneous accelerators (SIMD/GPU-like/FPGA)

Custom **Network-on-Node** extends the NoC across the chip boundary

Architecture-, NoC-, and NoN-level partitioning mechanisms for **QoS guarantees**




MANGO prototype as input


The MANGO platform




Compute platform



Emulation platform



XILINX
VIRTEX



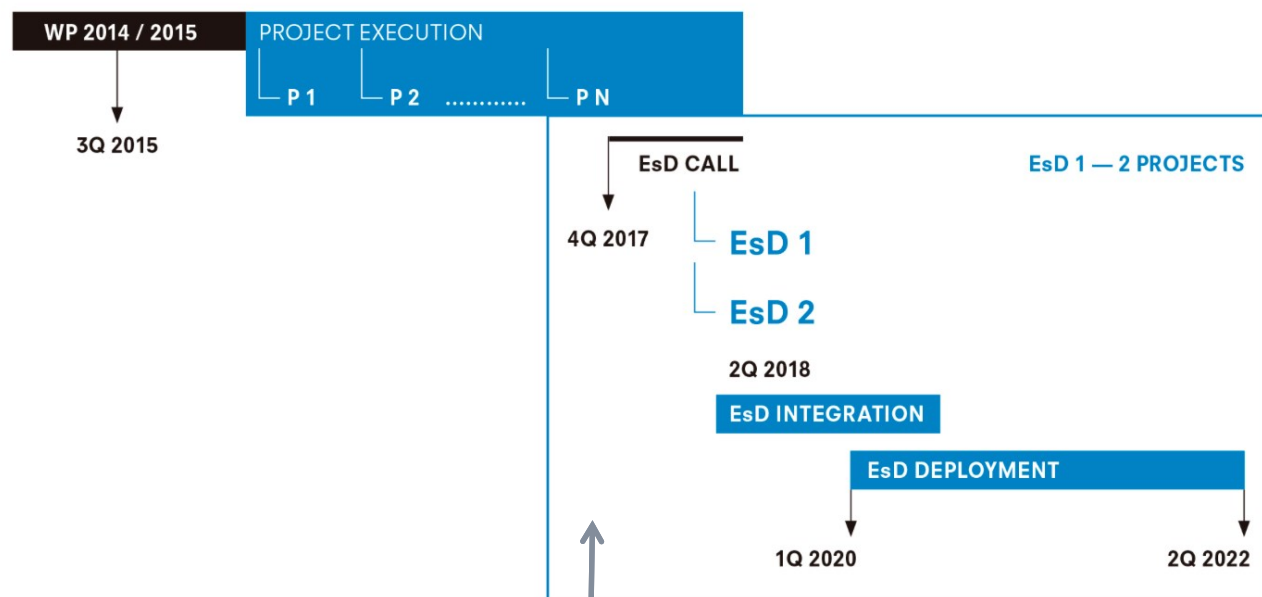
Innovative Cooling and Power Monitoring

MANGO prototype as input

- Pro-Design proFPGA quad V7 Prototyping system
 - Scalable up to 48 M ASIC gates capacity on one board
 - Modular with up to 4 x Xilinx Virtex XC7V2000T FPGAs, or Zynq-7000, or memory modules
 - Up to 4336 signals for I/O and inter FPGA connection
 - Up to 32 individually adjustable voltage regions
 - Up to 1.8 Gbps/12.5 Gbps point to point speed



Our plans...



- MANGO architectures deployed
- Ready for chip design/deployment
- MANGO hardware/software ecosystem ready for exploitation
- Advanced high-volume reconfigurable server

MANGO partners see collaboration opportunities with FPGA-related projects: GreenFLASH, EXTRA
Other collaboration opportunities are with global ecosystem: CoEs, FET HPC projects