



exploring Manycore Architectures for Next-GeneratiOn HPC systems

MANGO project status update

Prague, EXDCI Workshop, May 2016

Mario Kovač, University of Zagreb



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

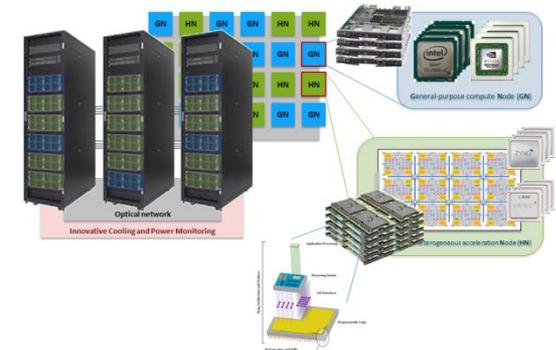




MANGO project and consortium

- **MANGO**: exploring **M**anycore **A**rchitectures for **N**ext-**G**enerati**O**n 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 scope

- Current HPC driven by Performance and Power optimization goals
- New applications and use cases emerging for HPC arena
 - With **real-time** constraints
 - Designs must be driven by **predictability** optimization goal
 - Including numerous applications in Data Center environments targeting **Capacity Computing** approach



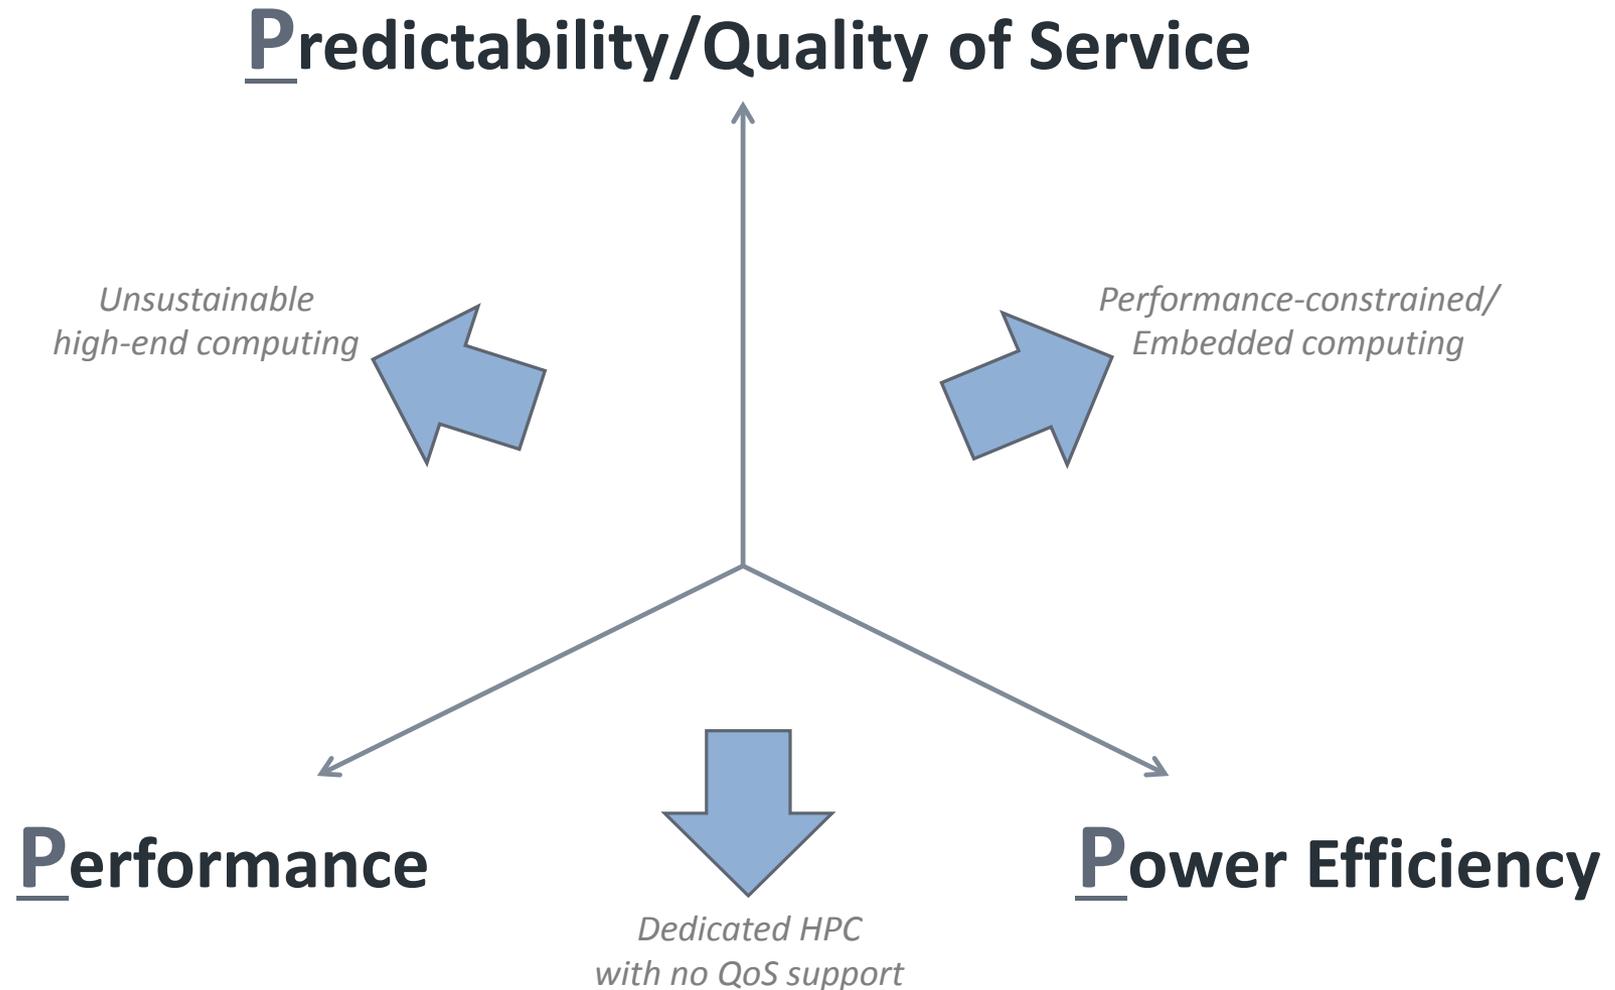
capacity computing



predictability

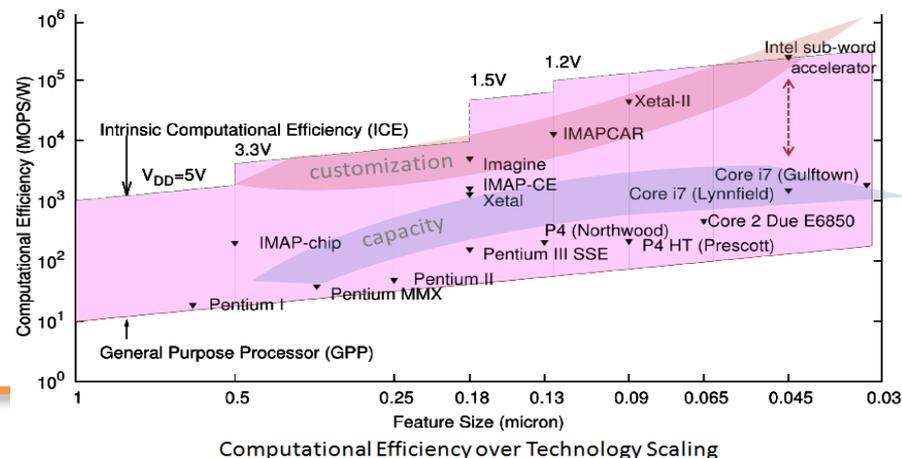


MANGO: Exploiting the *PPP* design space



MANGO scope

- Computational Efficiency can increase two orders of magnitude by deep customization of hardware
 - Need to fully match technology and architecture with applications
 - Efficient HPC systems will be achieved by **matching computing resources to applications** and underlying algorithms and addressing energy efficiency



MANGO scope

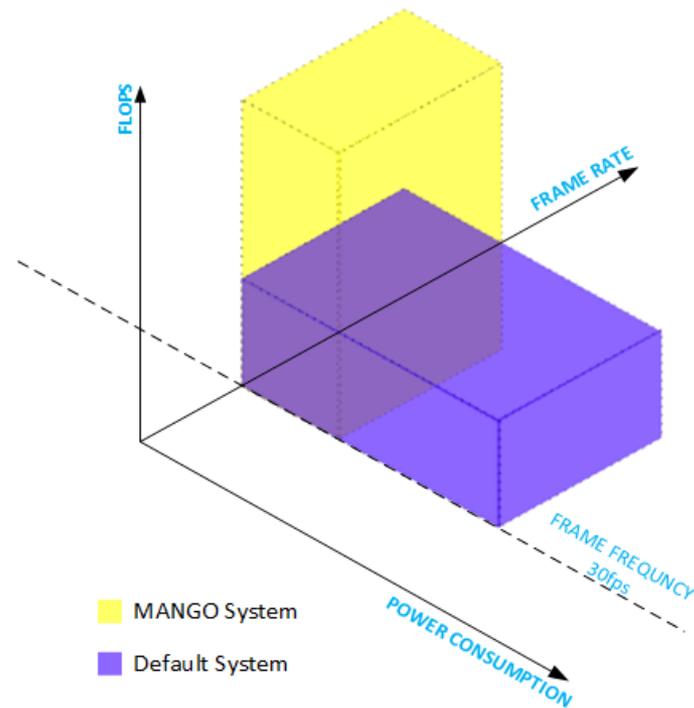
- Heterogeneity exploitation for HPC
 - Traditional devices available for HPC exploration
 - High-end processors, high-end GPUs
 - Non-conventional devices for HPC exploration
 - ARM cores, FPGAs, manycore chips (Kalray, Xeon Phi)
 - They exhibit different power/performance tradeoffs as well as different programmability complexities!
 - Key questions:
 - **How we combine them** for best achievement of computational efficiency
 - **How to program/manage them** for best achievement of computational efficiency



How do we combine them?



Predictability/QoS

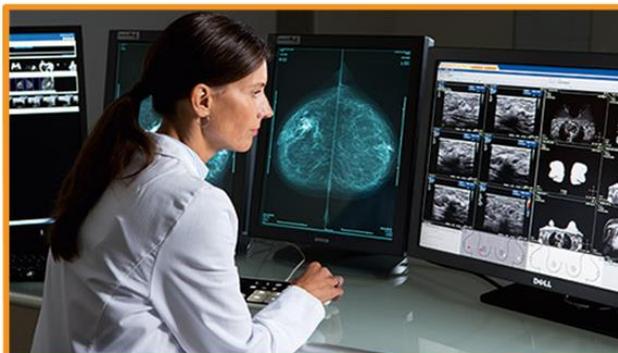
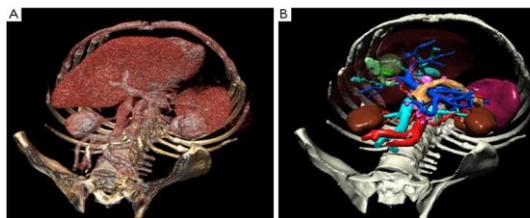




MANGO applications

- Chosen applications with stringent QoS and high-performance requirements:

- *Video transcoding*
- *Medical imaging*
- *Sensor data processing*
- *Security-related and cryptographic operations*





MANGO short-term goals

– Hardware

- Develop a **flexible prototype** for rapid exploration of archs
- **Explore** new deeply heterogeneous **manycore architectures**
- **Real-time support** exploring the PPP design space
- Provide a unified and simple (homogeneous) access to the system via a **smart interconnect**

– Software

- **Adapt** programming models and compiler support to the new architectures
- **Develop** the right resource manager to deal with the system

– Infrastructure

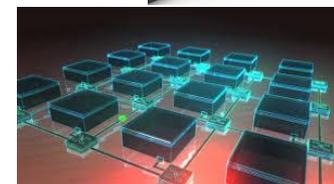
- Provide new **monitoring** tools to the system
- Provide new **cooling techniques** to the system

– Applications

- Analyze impact of on a set of **real applications**
 - video transcoding, medical imaging, security and surveillance applications

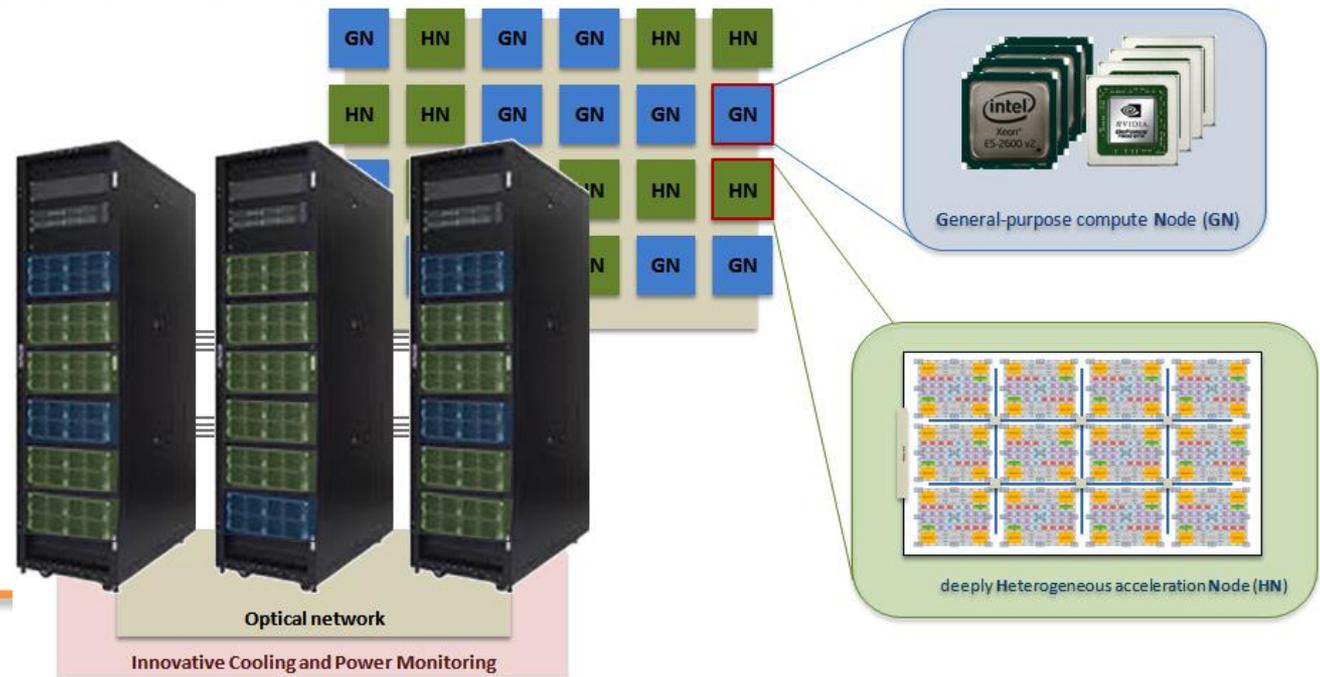


Real-Time



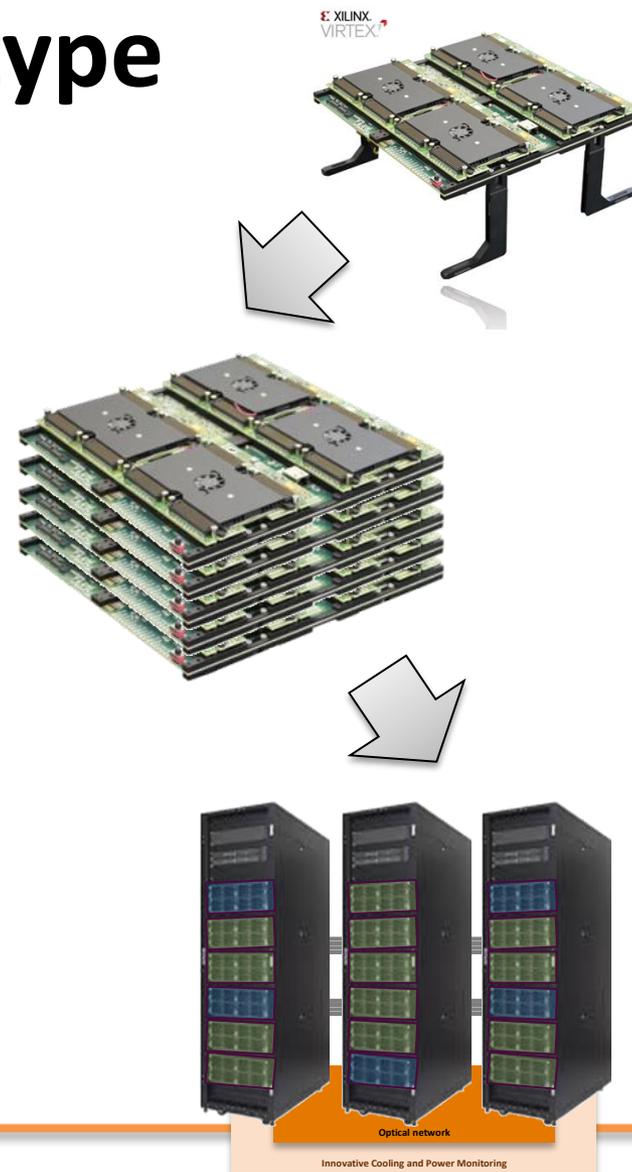
MANGO prototype

- Deeply heterogeneous prototype
 - High-end CPUs, GPGPUs, FPGAs (at this time Xilinx, including ZYNQ ARM-based)
 - On-chip DRAM modules as scratchpad memories or L1 caches
 - Massive parallel medium/low-end RISC cores
 - Fine-grained and FPGA-based custom accelerators



MANGO prototype

- Phase 1 – *Network of HN emulators*
 - Pro-Design proFPGA quad V7 Prototyping system
- Phase 2 – *Dedicated chassis*
 - standard connectivity and form factor
- Phase 3 – *Rack assembly*
 - rack collecting up to 16 blades
 - high-end CPUs, e.g. Intel Xeon chips, and GPUs +
 - 64 HN nodes

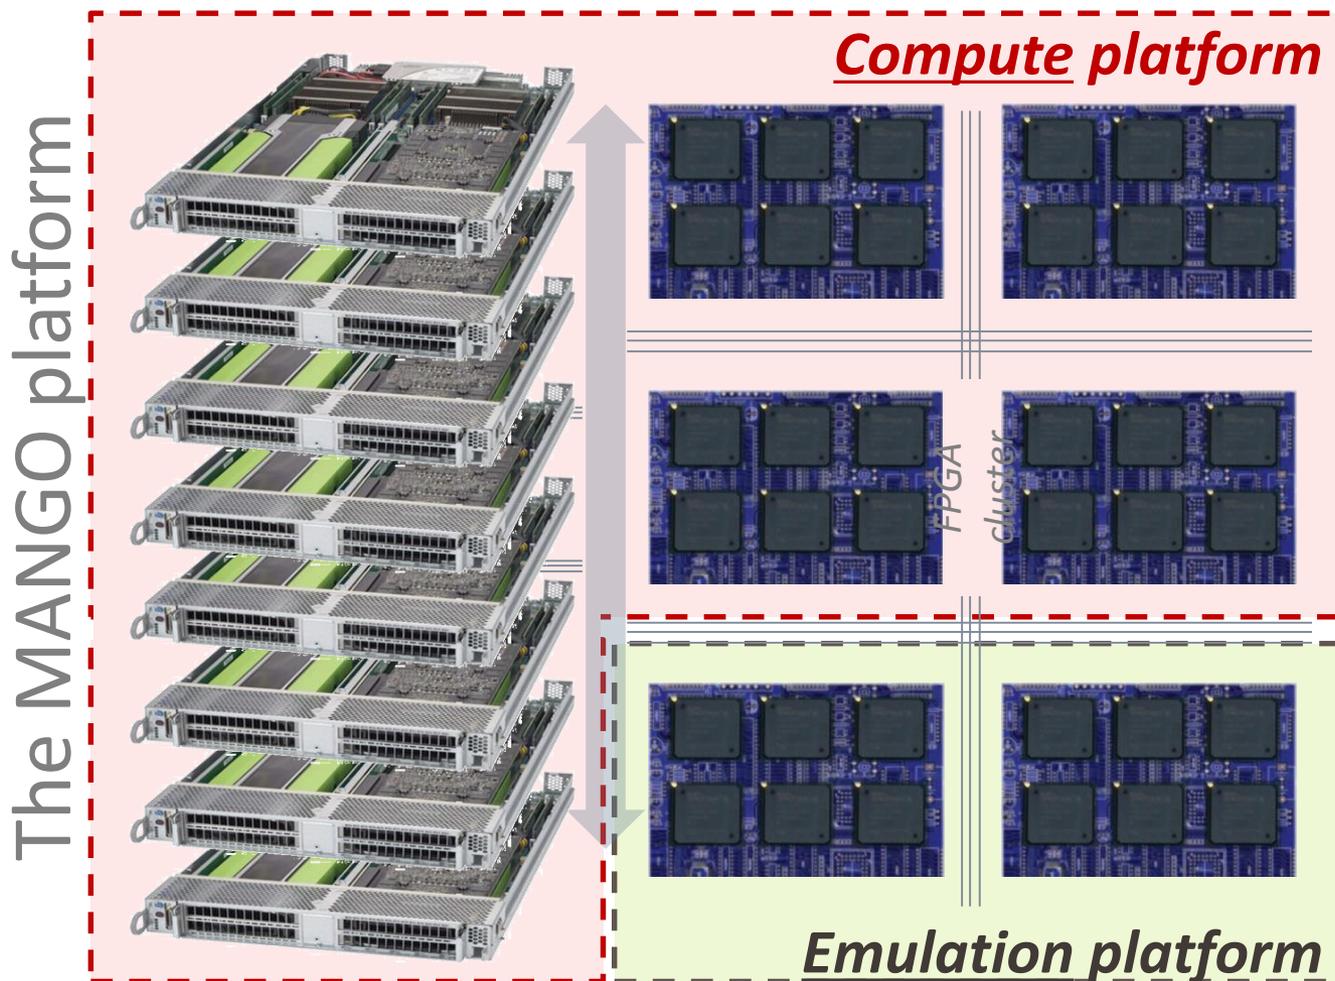


Phase 1 – Network of HN emulators

- **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 (or UltraScale)
 - 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



Compute vs Emulation platform





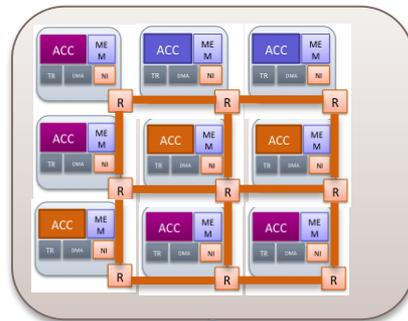
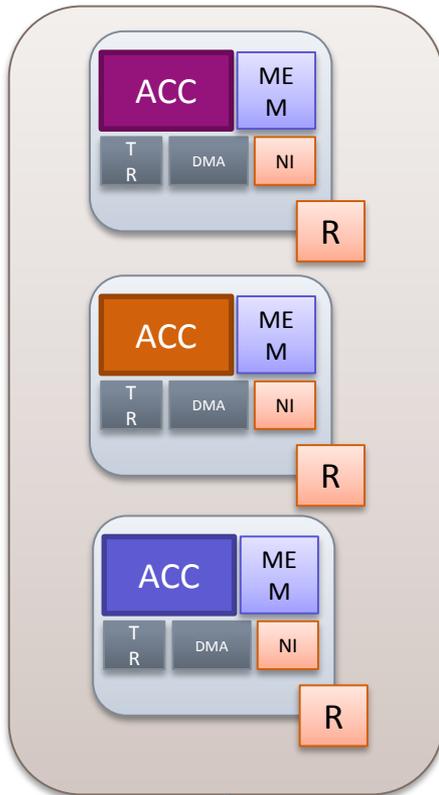
Current status

tile

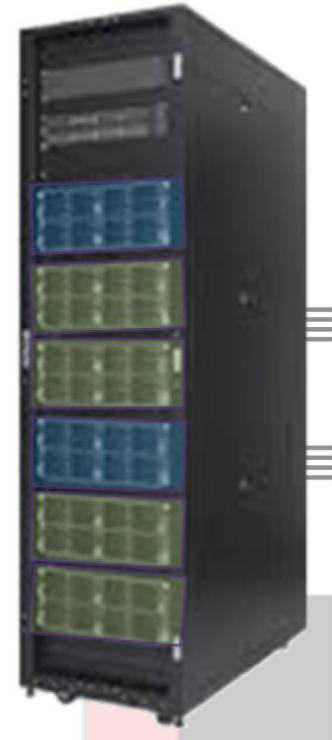
FPGA

BOARDS

SYSTEM



XILINX
UltraSCALE



Input specifications for applications, hardware, software

Basic infrastructure for configuration and monitoring

Initial accelerators developed (manycore PEAK, GPU-like,...)



MANGO offer to the ecosystem

3P (Performance/Power/**Predictability**) approach

- HN Architecture and novel Interconnect
- QoS sensitive SW stack
- Novel power and cooling approach
- Deep application level optimization libraries



exploring Manycore Architectures for Next-GeneratiOn HPC systems

Thank you...

Prague, EXDCI Workshop, May 2016

Mario Kovač, University of Zagreb



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

