



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

MANGO

Green Days@Toulouse
2/3 Juillet 2018
Alexandre DRAY - EATON



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

<http://www.mango-project.eu>
<https://twitter.com/mangoeu>



Consortium overview



UNIVERSITAT
POLITÈCNICA
DE VALÈNCIA



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE



PHILIPS

THALES

CONCEPT

HETEROGENEOUS
NODES

APPROACH



NETWORK

PROGRAM
MODEL

RESOURCE
MANAGER

APPS

PLATFORM

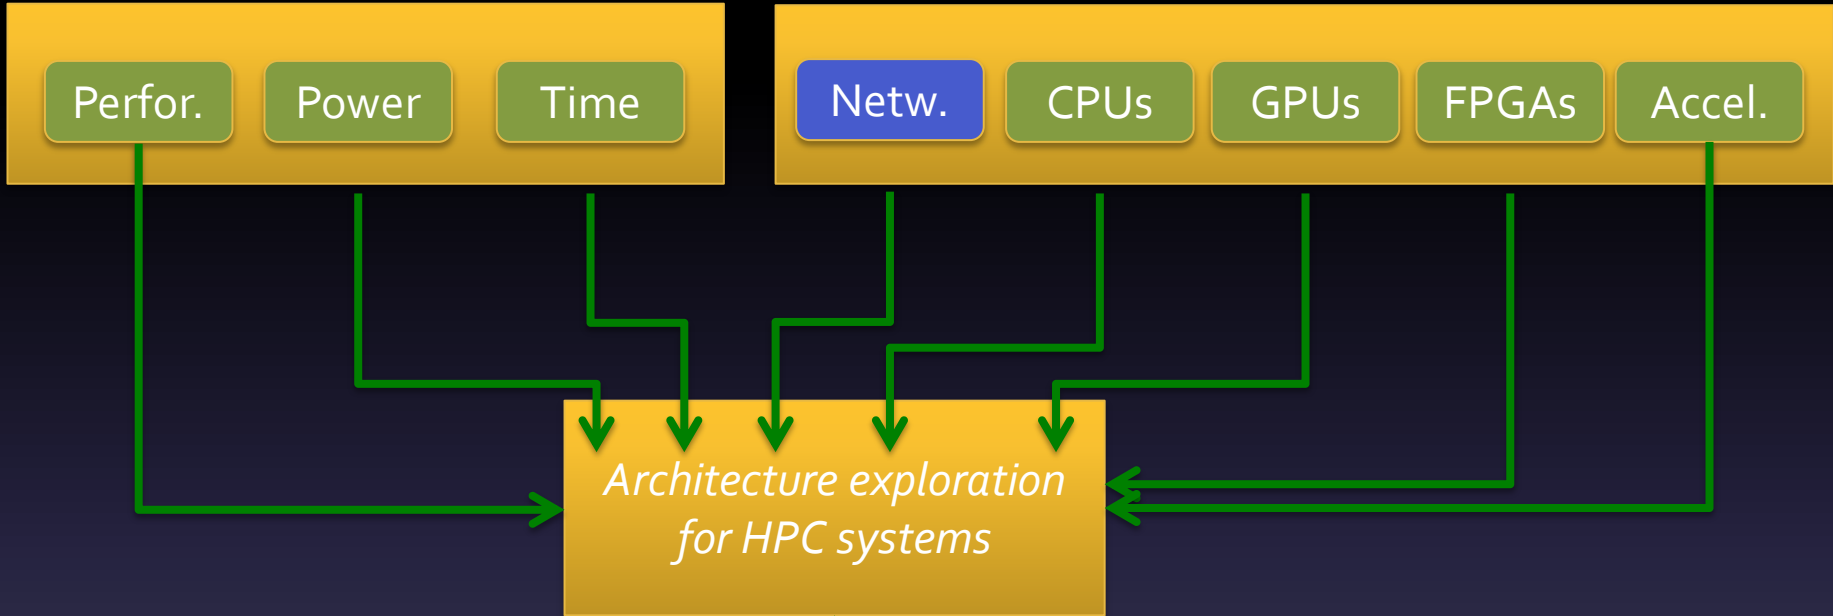
COOLING

CONCEPT




HPC requirements

Resources with different power/performance profiles



MANGO prototype



CONCEPT

HETEROGENEOUS
NODES

APPROACH



NETWORK

PROGRAM
MODEL

RESOURCE
MANAGER

APPS


PLATFORM

COOLING

Medical Volume Renderer


- Volume visualization of image slices coming from MRI-, CT- and Ultrasound scanners
- Per pixel-ray casting to visualize a 3D shaded volume
 - Time-critical and computational demanding
 - High data traffic, very strict latency requirements
- Goal: parallelize existing volume renderer and mapping to GPUs and FPGAs using MANGO platform.

PHILIPS



Volume renderer in action

PHILIPS



MRI slice of a pineapple

APPLICATIONS

MANGO H.265 Transcoder

- Transforming video bitstream in time critical environment



- Motion estimation
- In-loop filters
- Interpolation filter
- Exploitation of spatial and temporal dependencies

LDPC (Low Density Parity Check)

- Description

Low Density Parity Check code is a linear error correcting code, a method of transmitting a message over a noisy transmission channel

- Application domains

- Satellite communications for near and deep space
- 10 Giga-bit/s Ethernet (802.3an)
- Wi-Fi 802.11 standard as an optional part of 802.11n

CONCEPT

HETEROGENEOUS
NODES

APPROACH



NETWORK

RESOURCE
MANAGER

PROGRAM
MODEL

APPS

PLATFORM

COOLING

RESOURCE MANAGER

Thermal-Aware Task Mapping and Allocation

Applications



*Global
Resource
Manager*





MANGO Platform

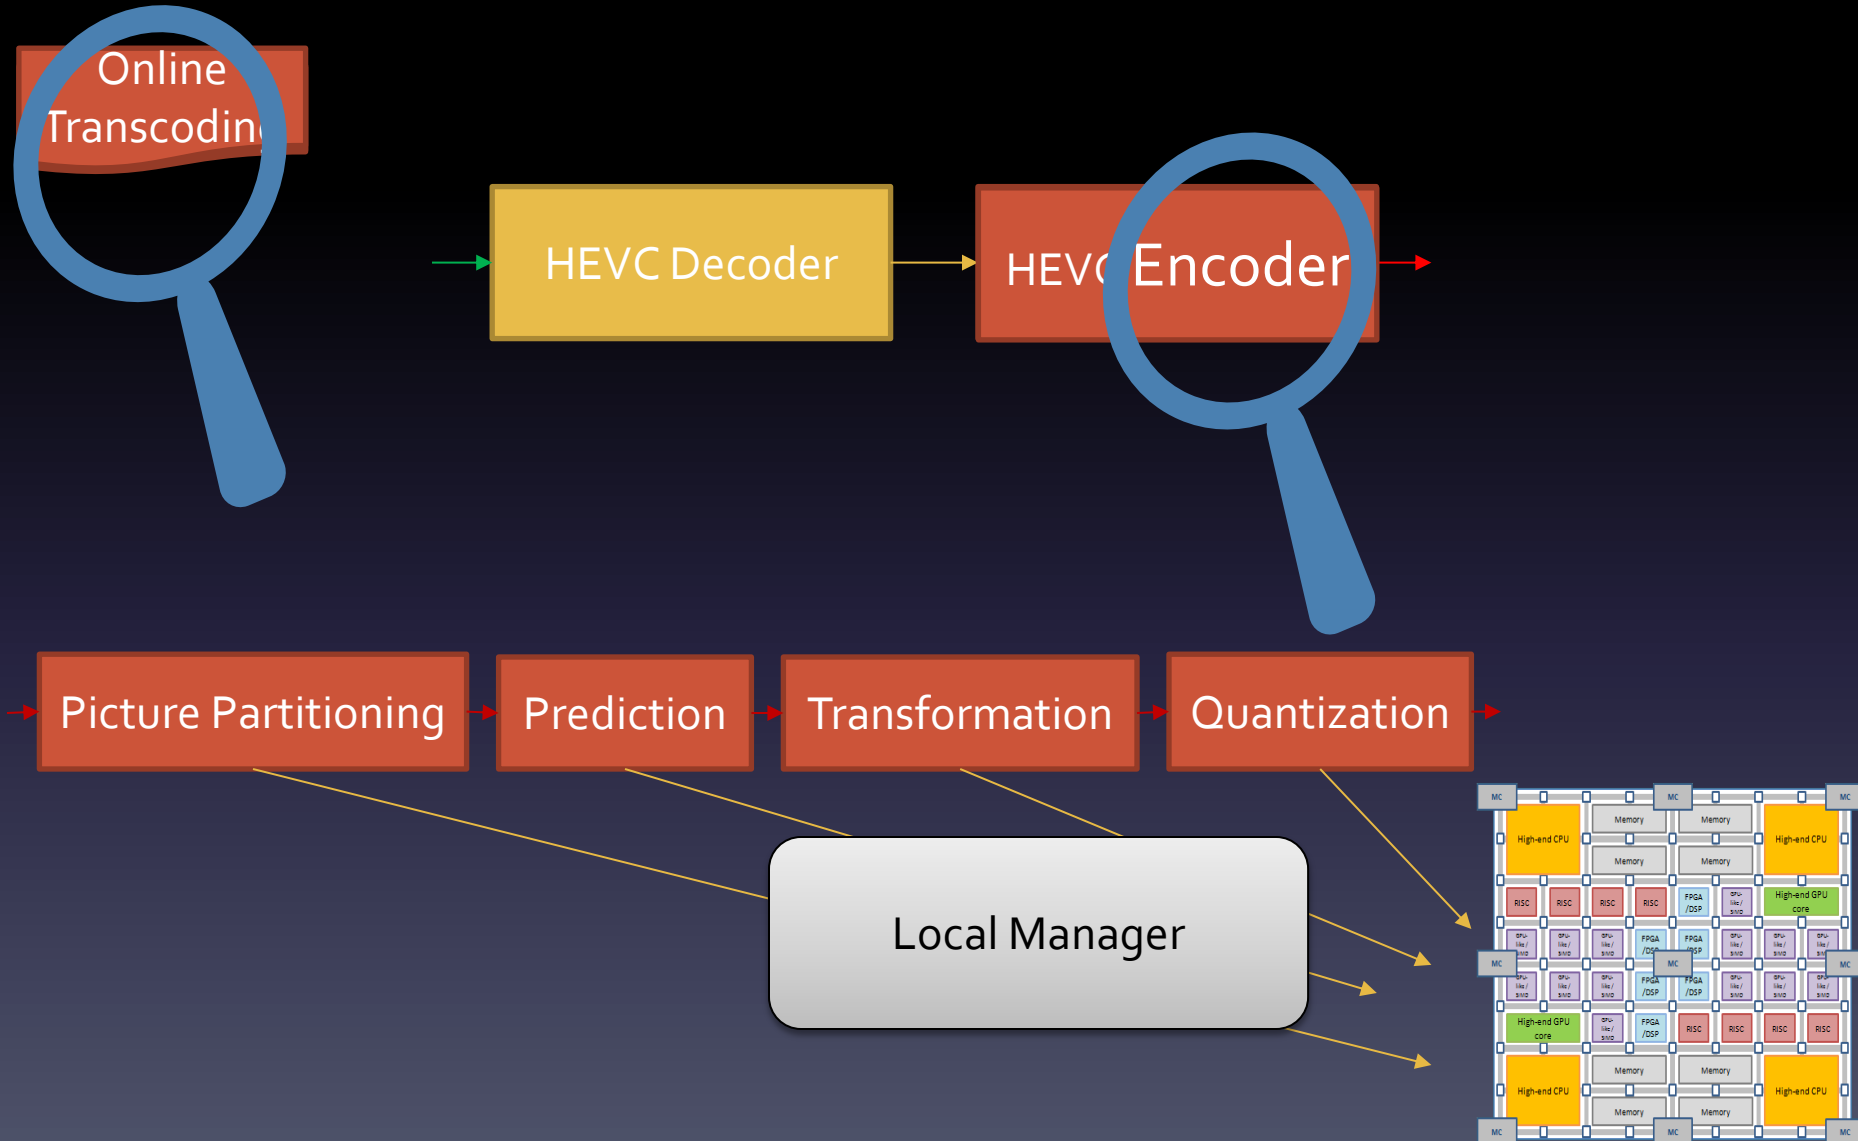
RESOURCE MANAGER



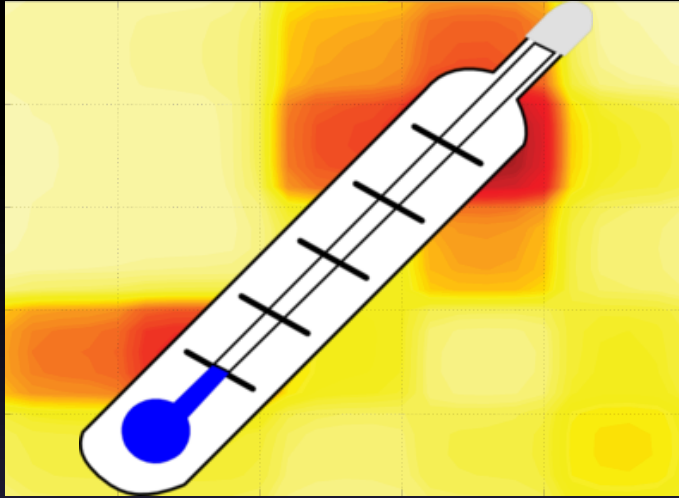
*Global Resource
Manager*



Intra Application Mapping




Thermal-Aware Task Allocation



Thermal-Aware Mapping and Allocation

Random Mapping and Allocation



CONCEPT

HETEROGENEOUS
NODES

APPROACH



NETWORK

PROGRAM
MODEL

RESOURCE
MANAGER


APPS

PLATFORM


COOLING

PLATFORM

- FPGA Hardware used by MANGO:
 - Up to 256 FPGAs (Stratix 10, Virtex-7, Zynq7000, Ultrascale, Ultrascale+, Ultrascale+ Zynq)
 - Long distance optical Gigabit communication with 40GBit per link, up to 12 Links per FPGA, <1us latency
 - Short distance general purpose communication with 120 GBit/s per link, up to 8 Links per FPGA, < 100ns latency
 - DDR4 memory up to 2400MWords/s



FPGA-FPGA cables DDR-4 memory



OSFP PCIe gen3



proFPGA quad
Stratix® 10



proFPGA quad
XILINX Virtex-7



proFPGA quad
XILINX Ultrascale



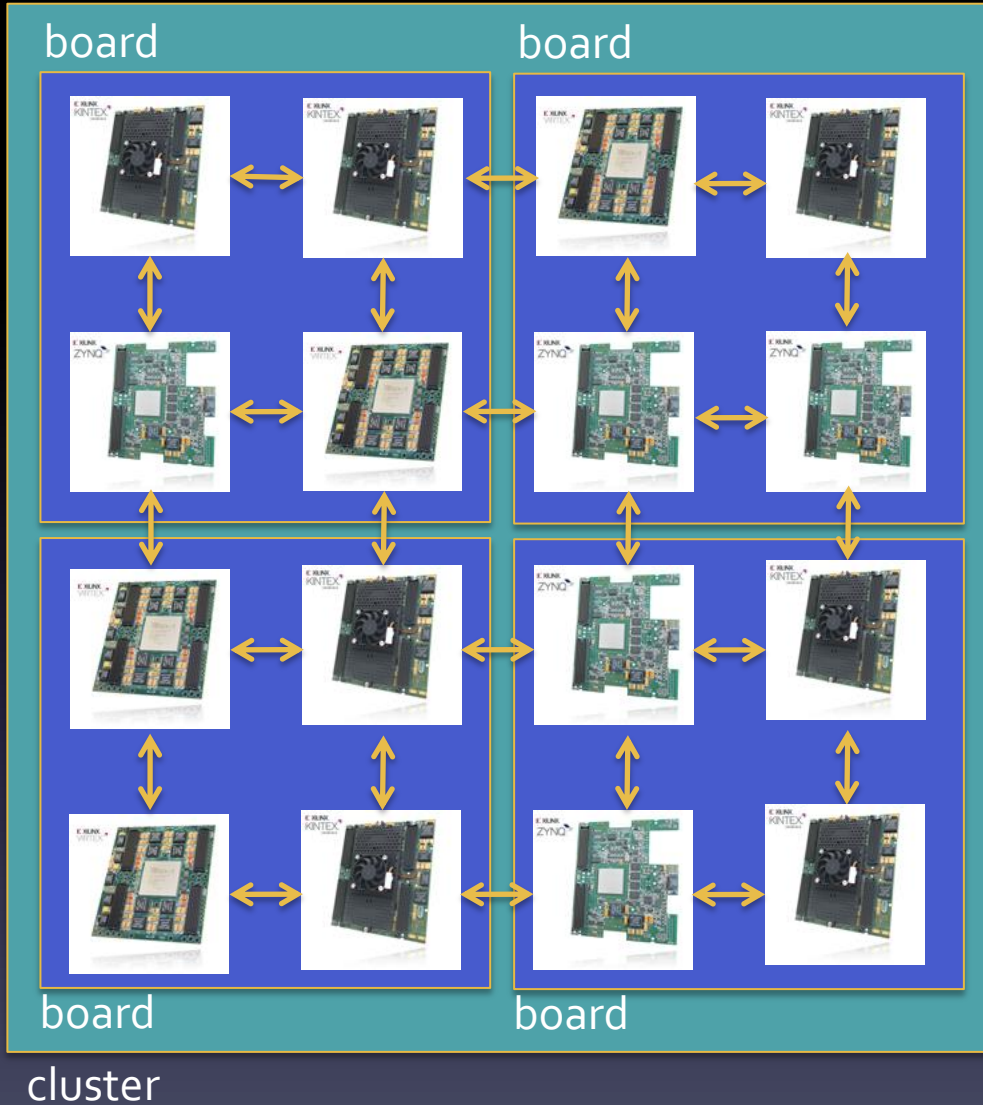


proFPGA Zynq
7000

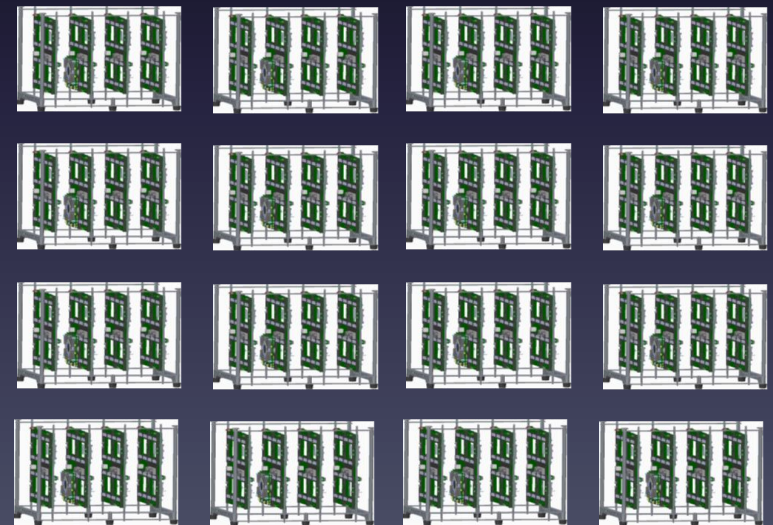


proFPGA Kintex
Ultrascale

PLATFORM




16x clusters




MANGO system (optically interconnected)

PLATFORM


Final prototype layout



FPGA boards (Air cooling)



*FPGA boards
(Thermosyphon cooling)*




CPU/GPU Servers





12V DC custom UPS

From initial concept...



... to the real demo !



CONCEPT

HETEROGENEOUS
NODES

APPROACH

MANGU


NETWORK

RESOURCE
MANAGER

PROGRAM
MODEL



APPS

PLATFORM

COOLING


COOLING

Principle







COOLING

Real prototype on MANGO FPGA



Thermosyphon vs. Fan





Thank you