Computing On Low-power Architectures: COLA workshop

Filippo Mantovani

Senior researcher at Barcelona Supercomputing Center
Principal investigator of the Mont-Blanc projects

Connecting dots

- **2nd of Sept 2015**
  Piero Altoè (E4 Engineering) email me the idea of organizing an event with “people who have dirty hands” working with low-power platforms involving ARM

- **18th of Sept 2015**
  Fabio Schifano (University of Ferrara) contact me asking if I am available in organizing an event focusing on low power architectures

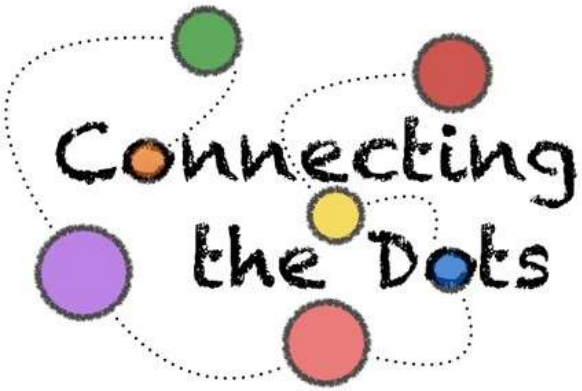

Scope of the workshop

Gather people with:

- Experience on low-power architectures-platforms
- Sensibility not only to compute power,
  but also to energy constraints
- Interest in minimizing power, cost and space
- Enthusiasm for disseminating and educating
  on energy efficiency issues
- Plans and requirements coming from real applications
- Open mind to new ideas and collaborations

All together in the same room with companies developing, integrating and deploying emerging technology.

Sponsors, contributors and organizers

- Sponsored by:

- With the contribution of:

Agenda

09:00 – 09:30	F. Mantovani (BSC)
09:30 – 09:55	G. North (ARM)
10:00 – 10:25	G. Lander (ARM)
10:30 – 11:00	Break
11:00 – 11:25	D. Cesini (INFN-CNAF)
11:30 – 11:55	E. Calore (INFN)
12:00 – 12:25	M. Michelotto (INFN)
12:30 – 12:55	D. Bastieri (UNIPD)
13:00 – 14:00	Lunch
14:00 – 14:25	L. Benini (ETHZ)
14:30 – 14:55	D. Tafani (LRZ)
15:00 – 15:25	L. Natvig (NTNU)
15:30 – 15:55	A. Lonardo (INFN)
16:00 – 16:30	Break
16:30 – 16:55	C. Cavazzoni (Cineca)
17:00 – 17:25	M. Manzali (UNIFE)
17:30 – 17:55	D. Abdurachmanov (CERN)
18:00 – 18:30	Panel day #1
20:30	Social dinner

09:00 – 09:25	P. Kumar (TUE)
09:30 – 09:55	S. Chessa (UNIPI)
10:00 – 10:25	R. De Petri (UNIPR)
10:30 – 11:00	Break
11:00 – 11:25	F. Pantaleo (CERN)
11:30 – 11:55	F. Spiga (U. Manchester)
12:00 – 12:25	C. Fantozzi (UNIPD)
12:30 – 12:55	Panel day #2
13:00 – 14:00	Lunch

COLA workshop impact

Library and system sw
- Open source support
- Ecosystem development
- ARM-64 bit support
- Scientific tool sets

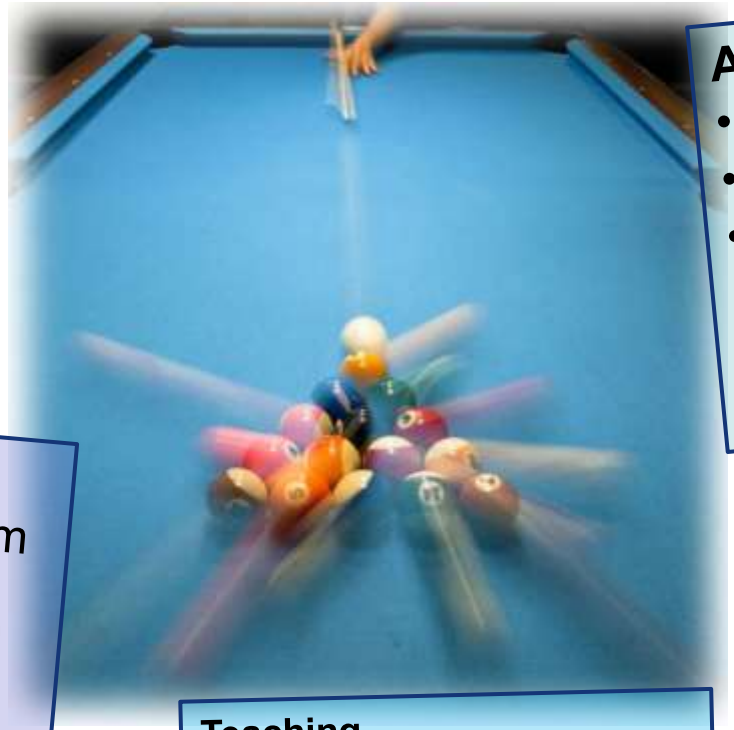

Applications
- HEP
- Neural networks
- Deep learning
- Material modeling
- Fluid dynamics
- Astrophysics

Prototyping
- Cost-efficient platform
- Heterogeneous platforms
- Benchmarking
- Energy to solution
- Data acquisition

Teaching
- Educational platforms
- Energy awareness

Data centers
- Power monitoring
- Energetic fair share
- Energy aware programming

Plus two panels chaired by companies for interacting, exchanging ideas, submit requests

Let me play my ball…

MONT BLANC

A pioneer European project

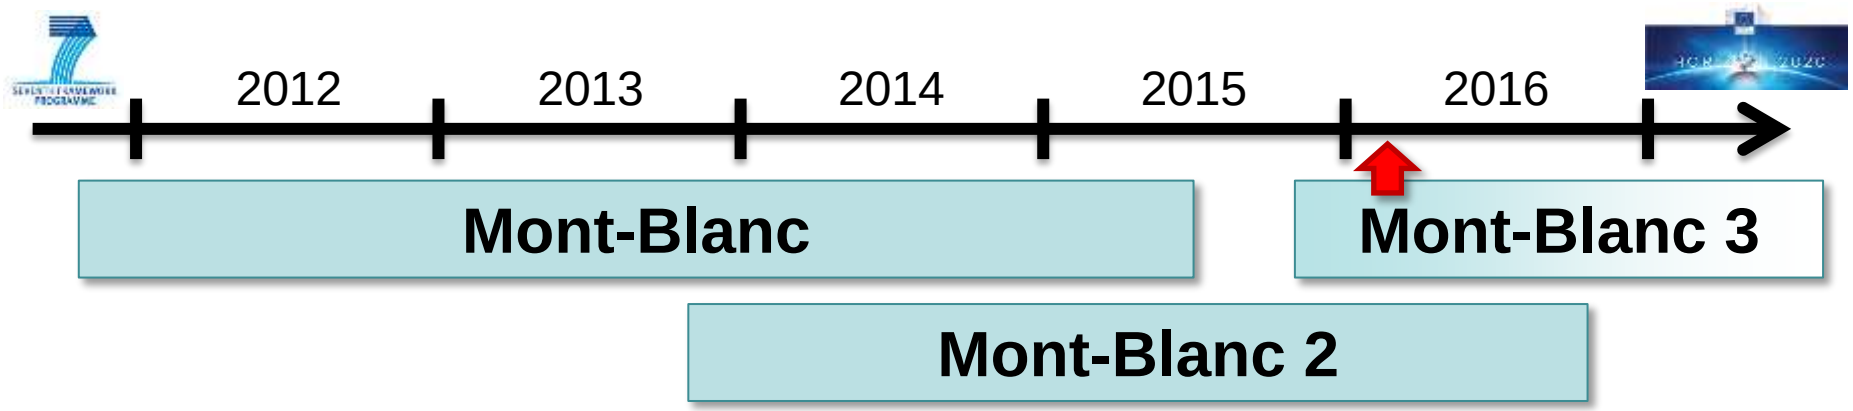

Vision: to leverage the fast growing market of mobile technology for scientific computation, HPC and non-HPC workload.

Mont-Blanc objectives

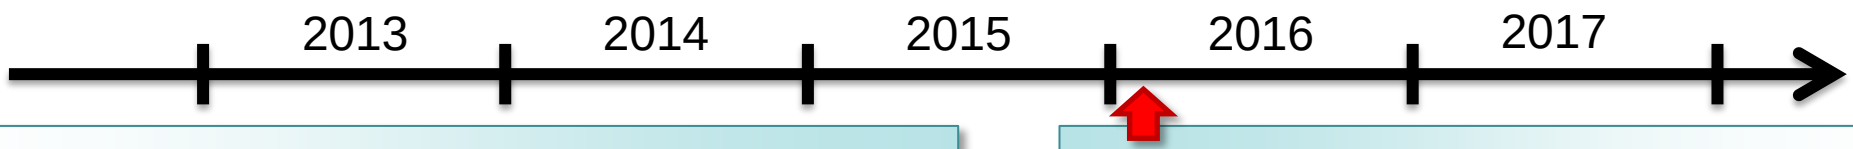

Mont-Blanc

Mont-Blanc 3

Mont-Blanc 2

- HPC prototype based on mobile embedded technology
- Port and test real scientific applications
- Learn from the experience, plan for future architecture

Extend

- Support hw and system sw:
  - OmpSs programming model
  - Productivity tools
- New scientific and industrial applications
- Next generation Mont-Blanc architecture

Explore

- ARM 64-bit
- Fault tolerance and resiliency
- Market of ARM-based platforms for mini-clusters

- Balanced ARM-based architecture targeting pre-exascale performance
- Focus on compute efficiency:
  - New high-performance ARM architecture
  - Throughput-oriented compute accelerators
- Co-design approach:
  - Architecture
  - System software
  - Applications

Leveraging a fast-growing market

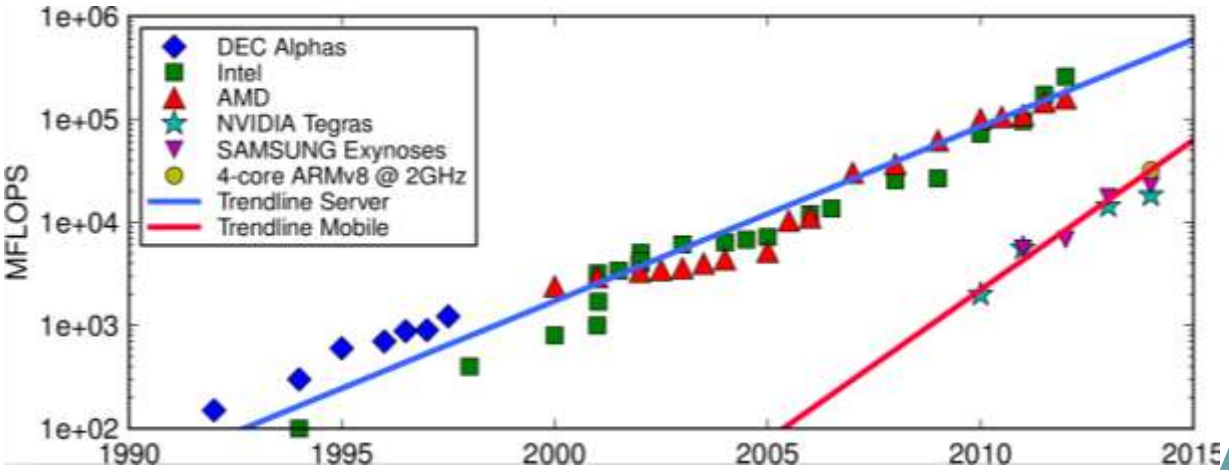

...and we are still ignoring tablets: >200M

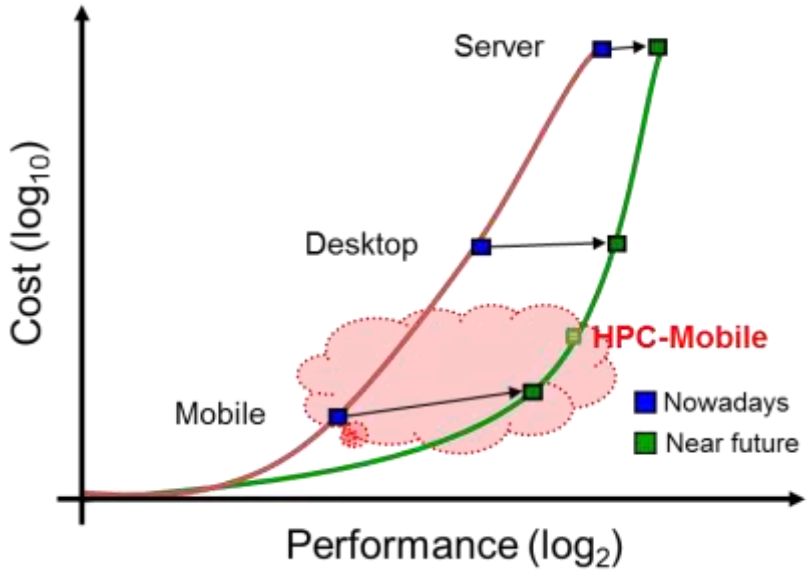

HPC
Jun 2015: 25 M cores

Server (+3%)
2013: 9.0 M
2014: 9.3 M

PC (-1 %)
2013: 316 M
2014: 314 M

Smartphone (+30%)
2013: 1000 M
2014: 1300 M

Highlight of Mont-Blanc contributions

- Prototyping
- Power monitoring
- System software
- Scheduling studies
- Memory fault statistics and analysis
- Dissemination activities

Significant ARM-based SoC

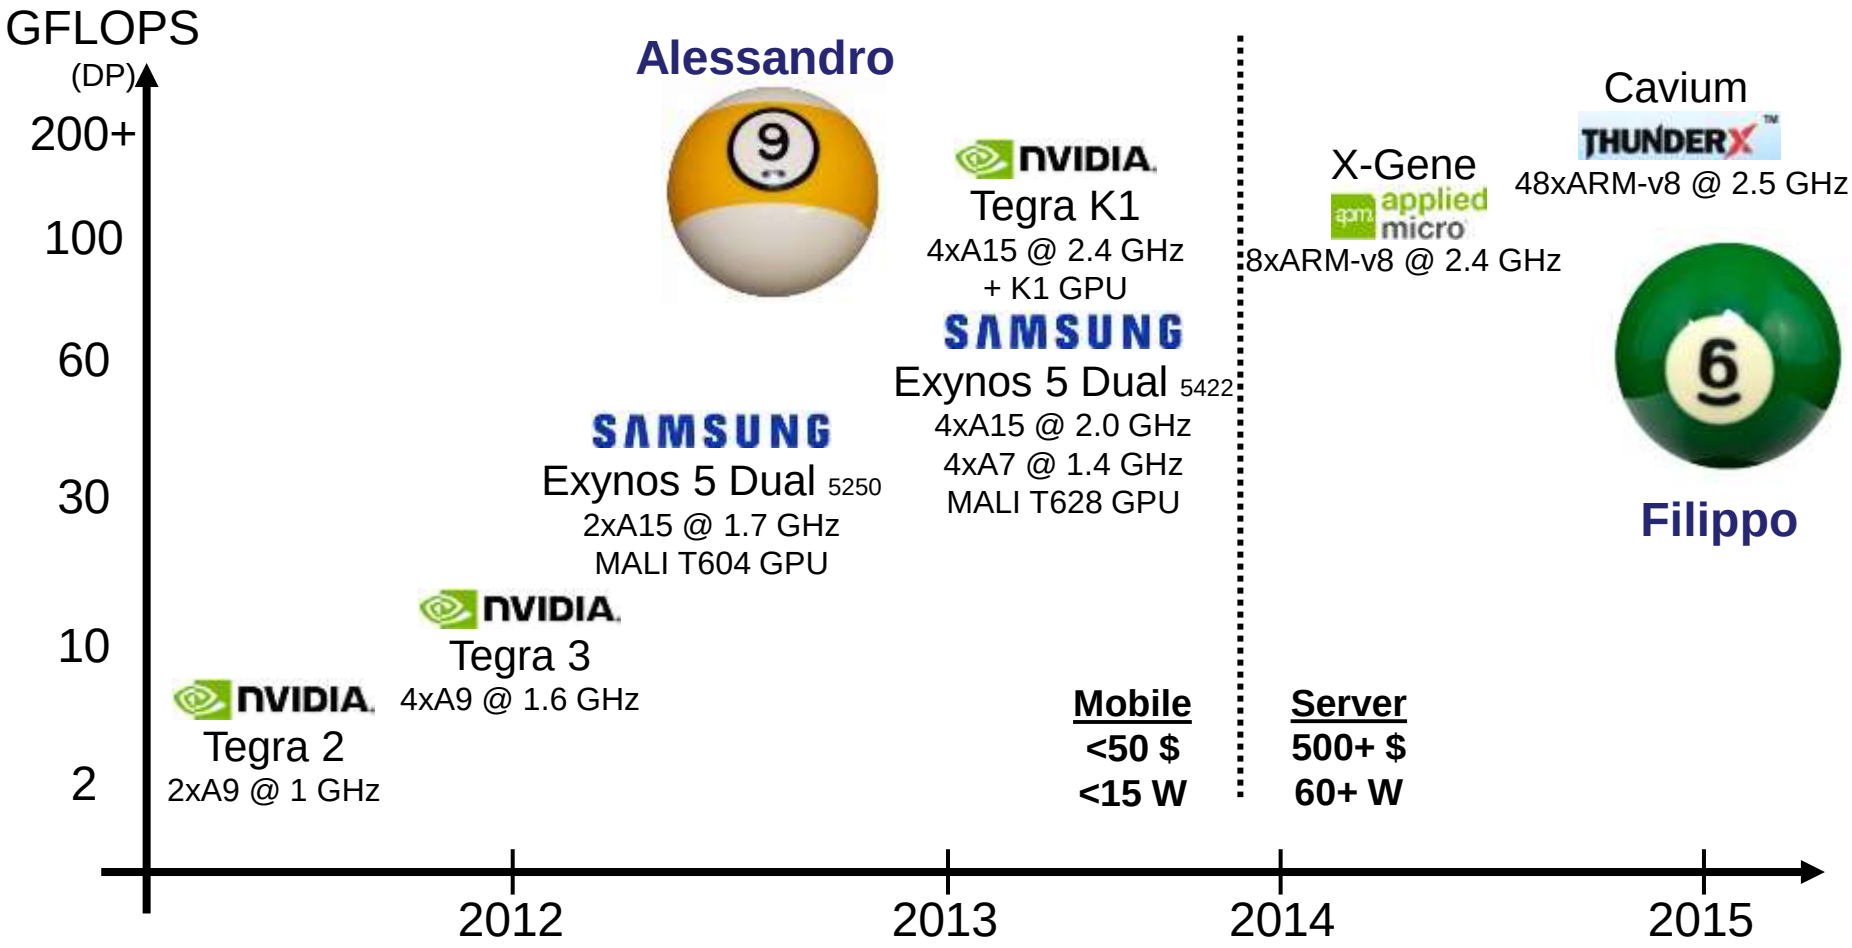

Disclaimer: Data from public sources

The Mont-Blanc prototype ecosystem

Prototypes are critical to accelerate software development
System software stack + applications

PRACE prototypes
- Tibidabo
- Carma
- Pedraforca

Mini-clusters
- Arndale
- Odroid XU
- Odroid XU-3
- NVIDIA Jetson

Mont-Blanc prototype

ARM 64-bit mini-clusters
- APM X-GENE2
- Cavium ThunderX

2012 2013 2014 2015 2016

Mont-Blanc Server-on-Module (SoM)

CPU + GPU + Memory + Local Storage + Network

Form factor: 8.5 x 5.6 cm

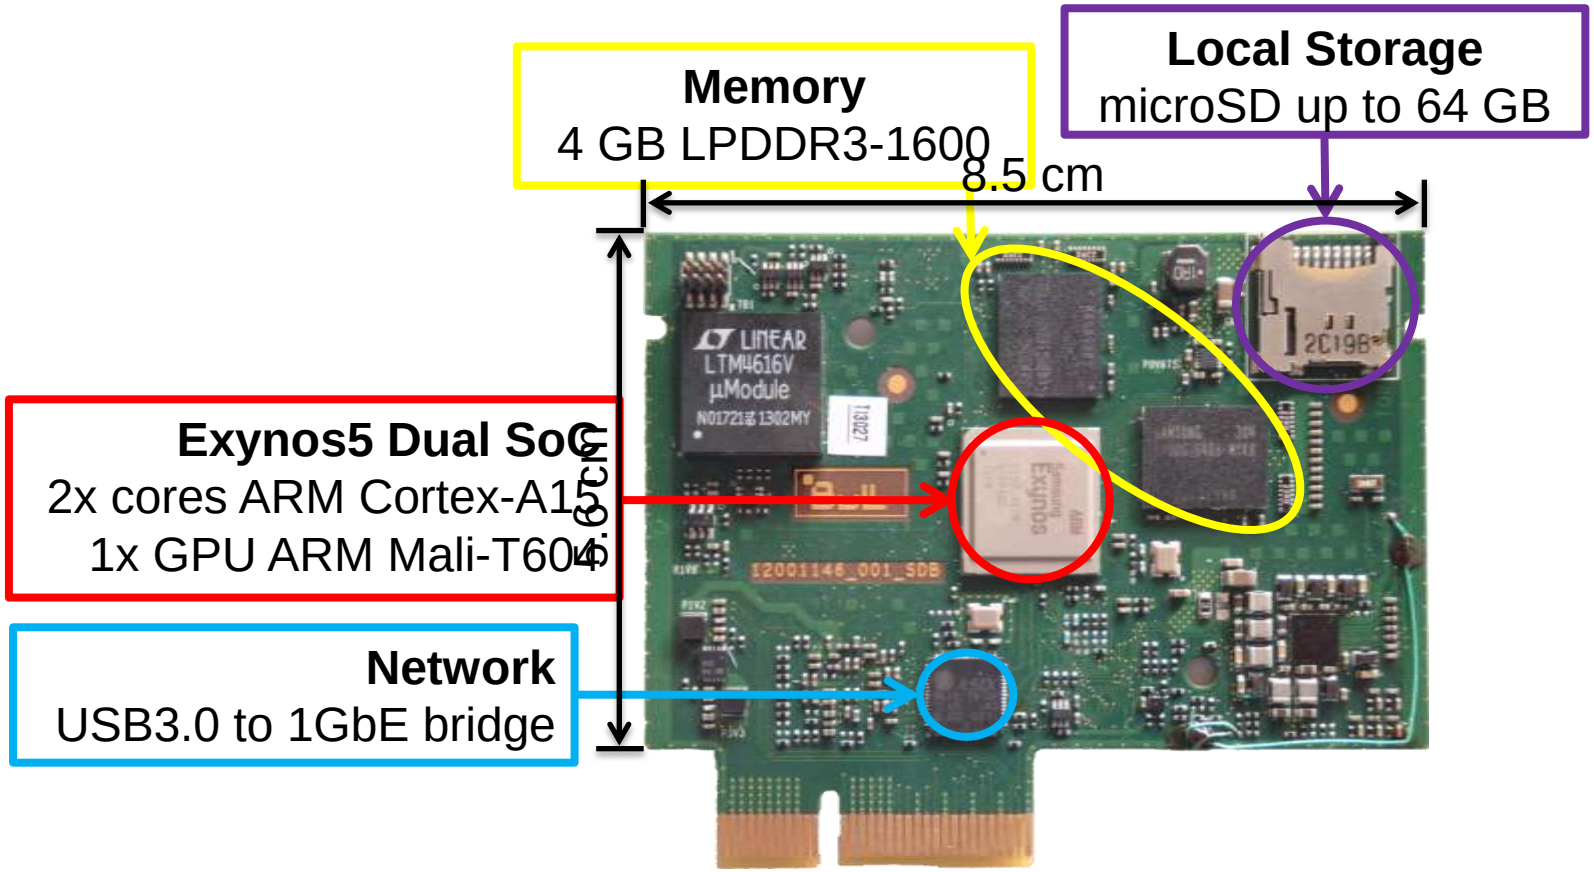

The Mont-Blanc prototype

Exynos 5 compute card
2 x Cortex-A15 @ 1.7GHz
1 x Mali T604 GPU
6.8 + 25.5 GFLOPS
15 Watts
2.1 GFLOPS/W

Carrier blade
15 x Compute cards
485 GFLOPS
1 GbE to 10 GbE
300 Watts
1.6 GFLOPS/W

Blade chassis 7U
9 x Carrier blade
135 x Compute cards
4.3 TFLOPS
2.7 kWatts
1.6 GFLOPS/W

Rack
8 BullX chassis
72 Compute blades
1080 Compute cards
2160 CPUs
1080 GPUs
4.3 TB of DRAM
17.2 TB of Flash

35 TFLOPS
24 kWatt

Daniele, David

Power monitor – HW infrastructure

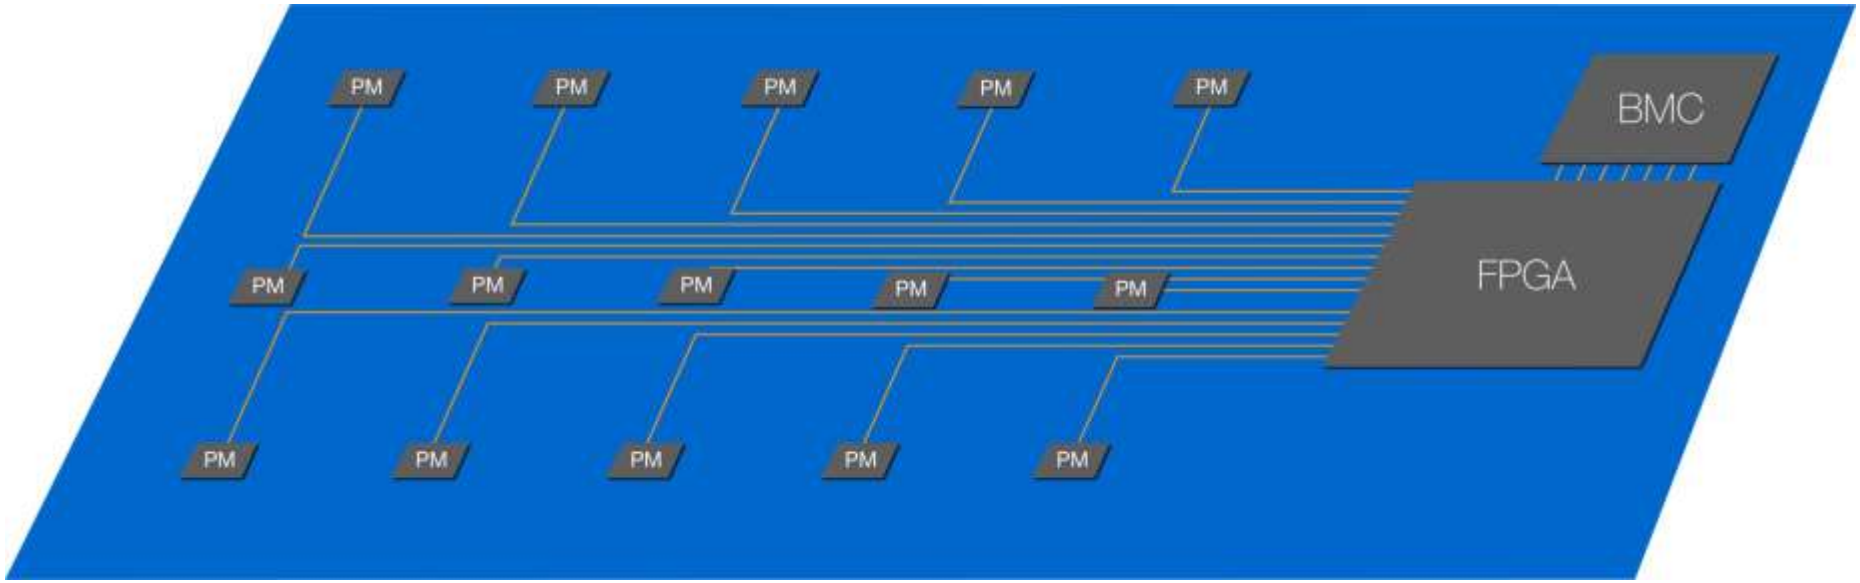

Power monitor – HW / SW interface

- Field Programmable Gate Array (FPGA)
  - Collects power consumption data from all 15 compute nodes
- Board Management Controller (BMC)
  - Collects 1s averaged data from FPGA
  - Stores measurement samples in FIFO
- Mont-Blanc Pusher
  - Collects measurement data from multiple BMCs using custom IPMI commands
  - Forwards data using MQTT protocol through Collect Agent into key-value store

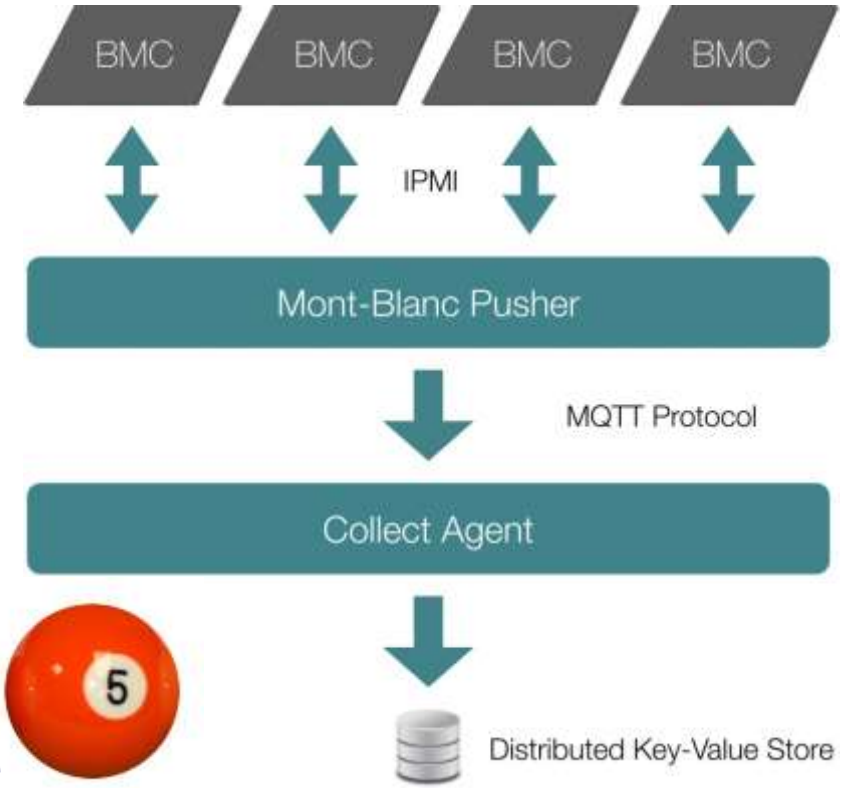

Daniele

Power profile study

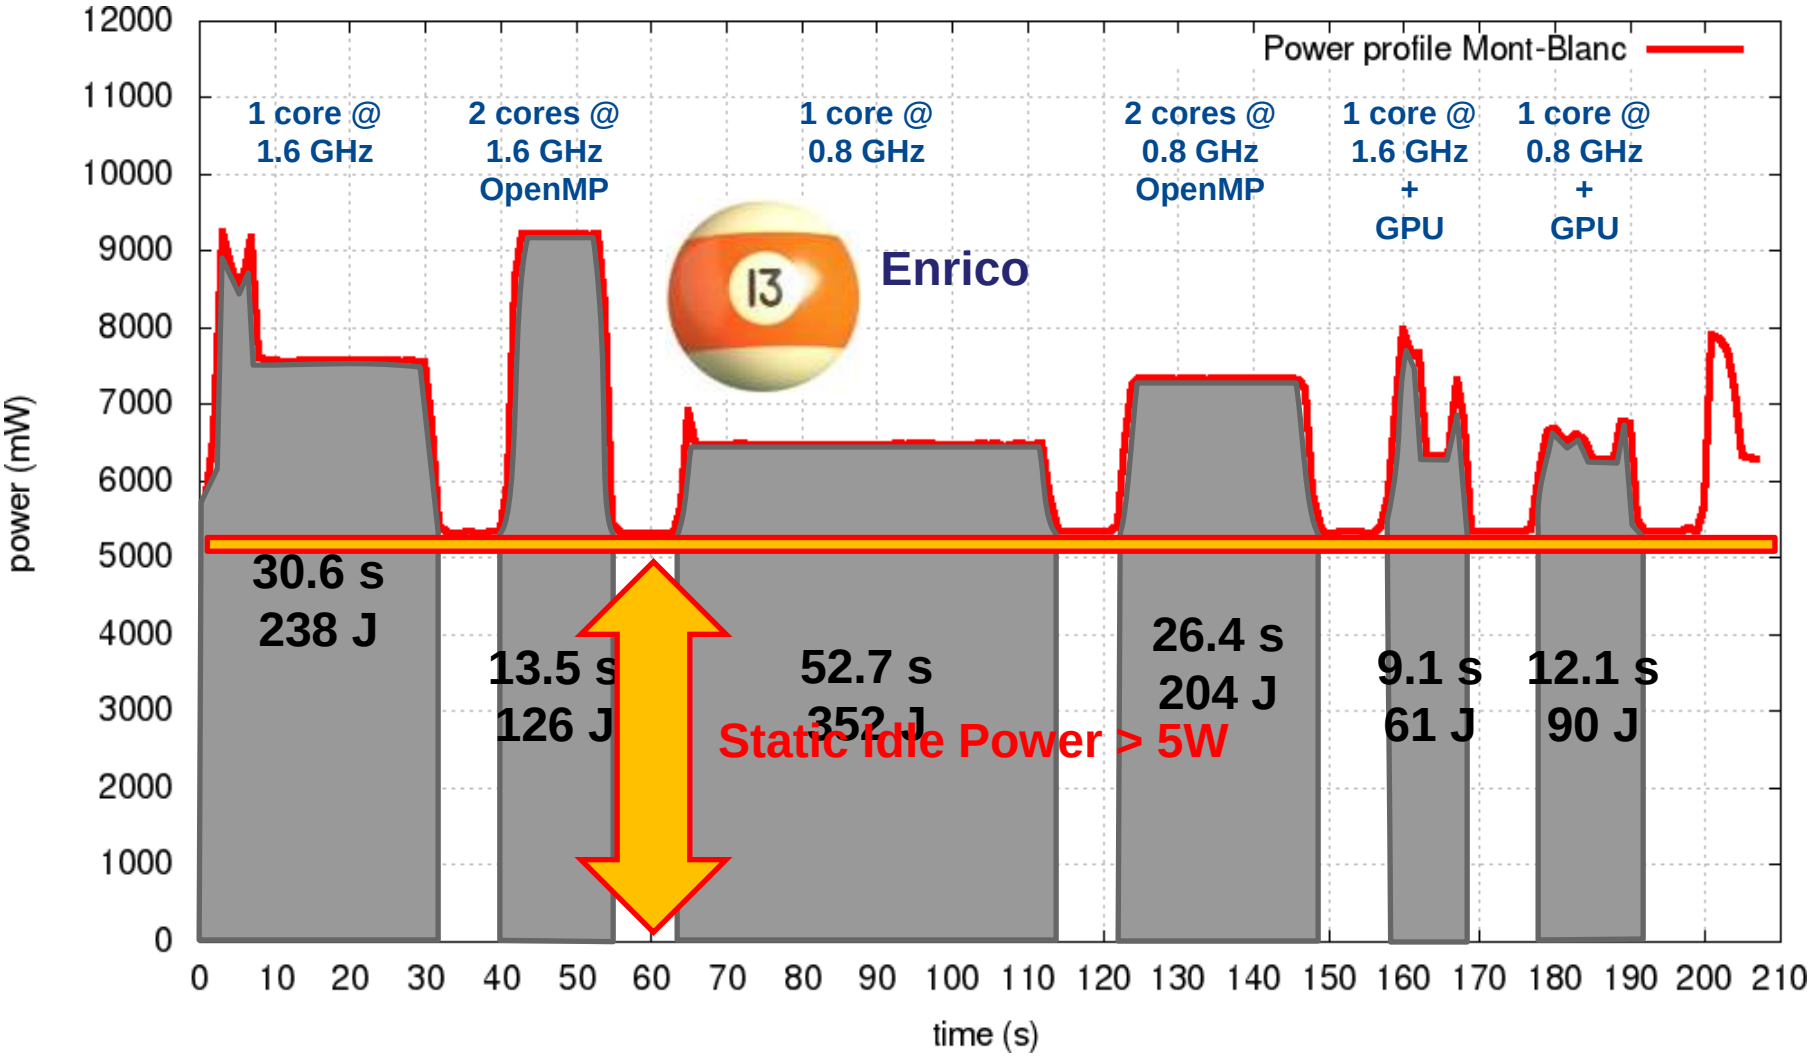

Full system software stack for ARM

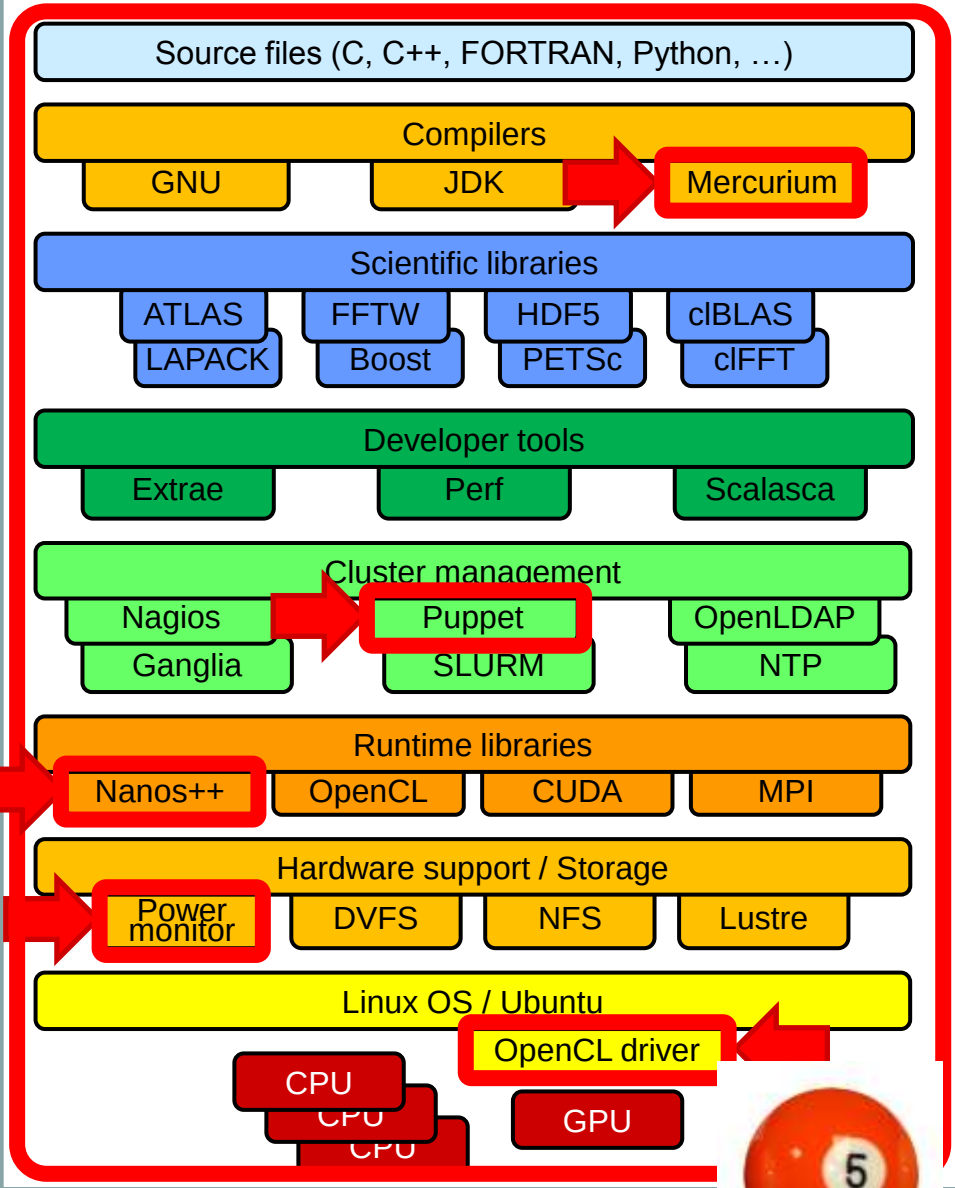

1. Fine grained power monitoring tool
2. OmpSs programming model ported to ARM + OpenCL support + FORTRAN support
3. Automatically deployed through Puppet and distributed through github
4. Tested on commercial ARM-based platforms

Criticality-Aware Task Scheduler

CATS dynamically assigns critical tasks to fast cores to improve performance in a heterogeneous system, e.g. big.LITTLE

- Scheduling information discoverable at runtime
  - No need of profiling
- Applies to task-based programming models supporting task dependencies
- Evaluation based on Odroid-XU3 + kernels + OmpSs

- Samsung Exynos 5422
- 2GB LPDDR3@933MHz
- 4x Cortex-A15@2.0GHz
- 4x Cortex-A7@1.4GHz

Michele, Denis, Carlo, Matteo, Pinaki, Roberto, Carlo, Felice, Stefano

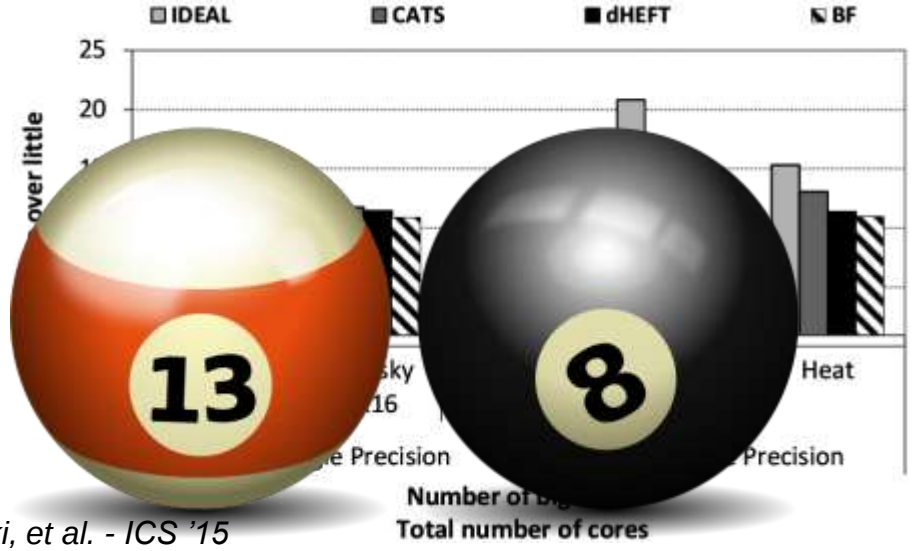

K. Chronaki, et al. - ICS '15

Memory fault statistics and analysis

Fact: Memory of Mont-Blanc prototype is not ECC protected
Can we survive with this? What do we learn from this?

- Developed a simple in-house memory scanner
- Scanning user-space (~3GB/node) when nodes are idle
- Collected statistics for ~2 Mhours over ~900 nodes
- Early statistics shows MTBF ~150 hours
- On going study…

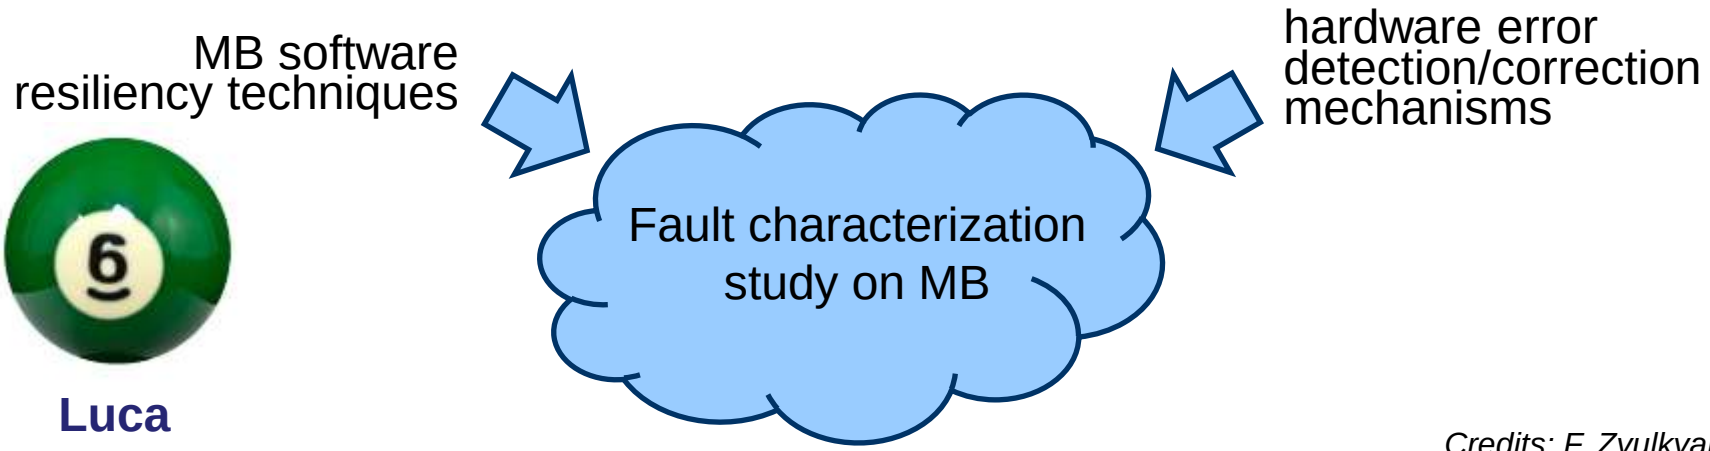

Credits: F. Zyulkyarov

End-User Group

- Develops a synergy among industry, research centers and partners of the project
- Validates the HPC technologies of the project
- Provides feedback to the project

ROLLS ROYCE

CybeleTech

DASSAULT AVIATION

Cenaero
Excellence in Simulation Technologies for Aeronautics

Pharmacelera™

MUREX™

HIGH END COMPUTE

ArcelorMittal

Termo Fluids

Mont-Blanc provides EUG members with:

- Remote access to Mont-Blanc prototype platforms
- Support in platform evaluation and performance analysis
- Invitation to the Mont-Blanc training program

Computing On Low-power Architectures

#COLAworkshop

cola-info@fe.infn.it

UNIVERSITÀ
DEGLI STUDI
DI FERRARA
- EX LABORE FRUCTUS -

montblanc-project.eu

MontBlancEU

@MontBlanc_EU

"The secret is to win going as slowly as possible."
Niki Lauda