MONT-BLANC

http://www.montblanc-project.eu

Exploiting mobile phone technology to build energy efficient supercomputers: the Mont Blanc project

Simon McIntosh-Smith, University of Bristol, UK

SEVENTH FRAMEWORK PROGRAMME

Disclaimer: Speaking for myself ... All references to unavailable products are speculative, taken from web sources. There is no commitment from ARM, Samsung, TI, Nvidia, Bull, or others, implied.

This project and the research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 288777.

Motivation: from MegaFLOPS to ExaFLOPS

- High Performance Computing (HPC) systems are now a common part of large-scale research infrastructure
- These systems have been growing larger and larger, and now themselves can consume considerable energy
  - Today we have systems capable of performing $O(10^{15})$ floating point operations per second, and we're working towards $O(10^{18})$
- Power consumption has grown by two orders of magnitude in the last two decades, from hundreds of kW in the '90s to tens of MWs today
- This is starting to cause **PROBLEMS**!

Power Consumption growing in the Top500

http://top500.org

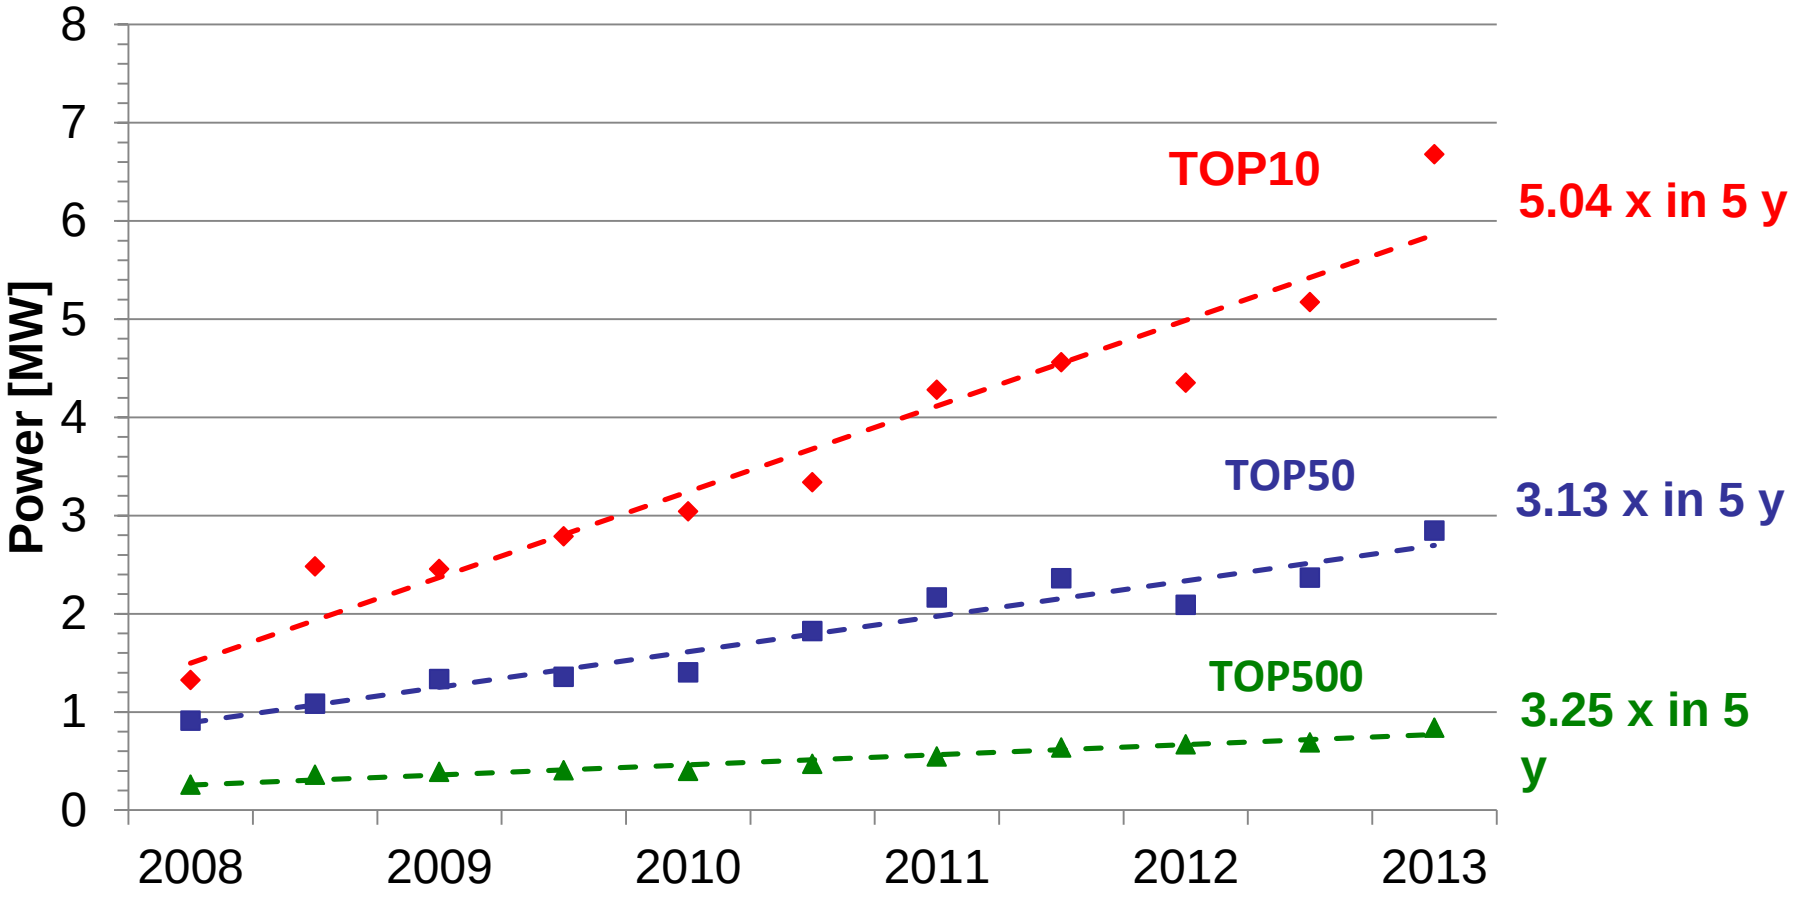

Graph courtesy of Erich Strohmaier, Lawrence Berkeley National Laboratory

Providing power is not the problem

- We can provide tens of MWs if they're really needed…

kWh cost for Medium size industries (€, 2011)

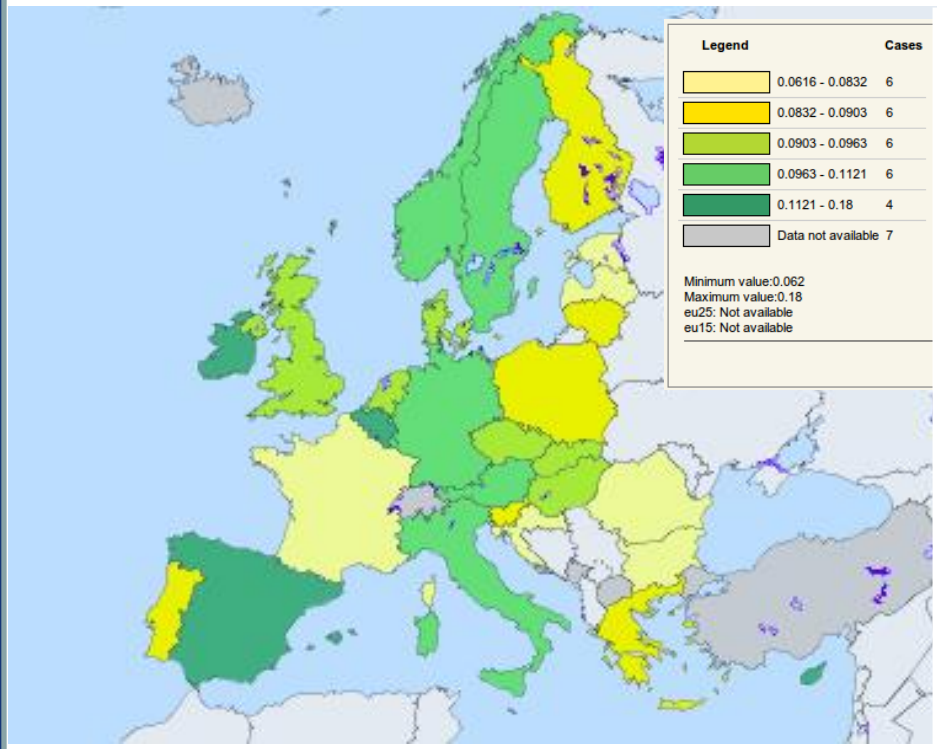

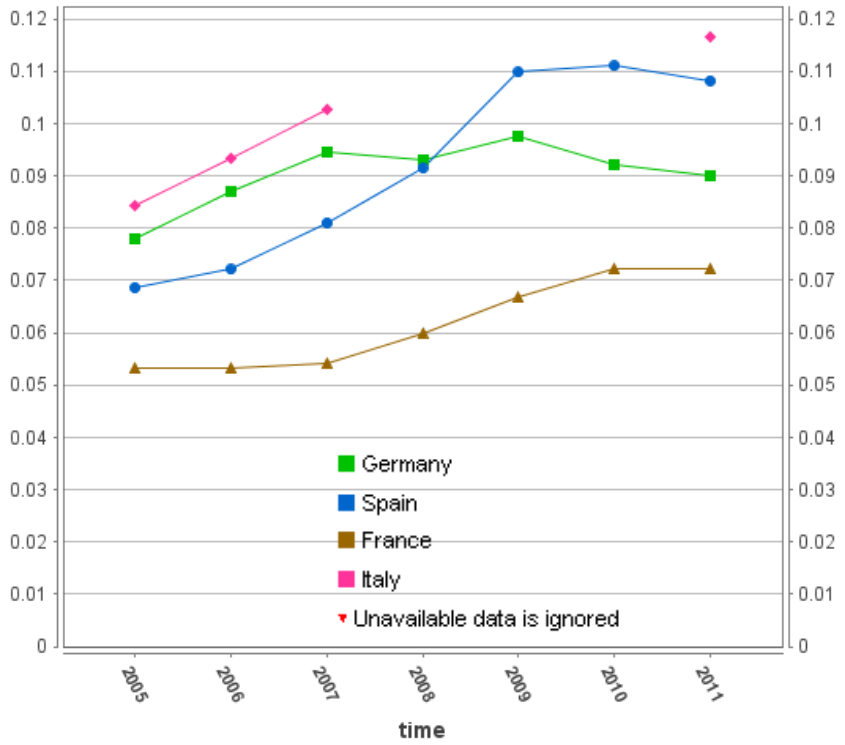

- The HPC power wall isn't due to power availability
  - We could provide 500 MW to an HPC facility
- It's the electricity bill at the end of the month! Also CO2…
  - 1 MWatt ~ 1 M EUR / year

MONT-BLANC

The problem … and the opportunity

- Europe represents ~40% of the HPC market
  - Yet, it does not have HPC technology of its own

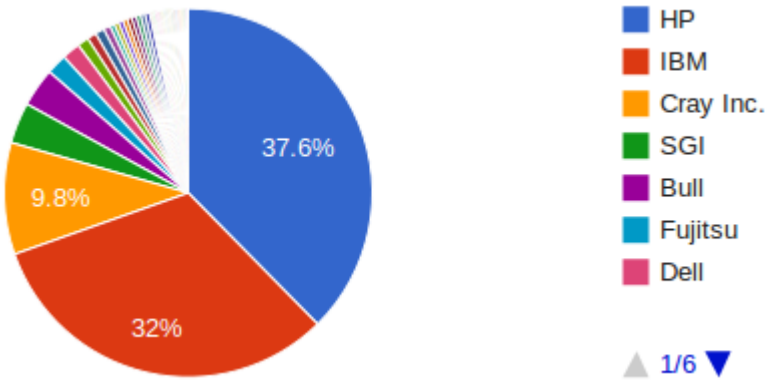

- Nobody yet knows how to build a sustainable EFLOPS supercomputer
  - Power defines performance
  - Consensus about it being revolutionary, not evolutionary
- Europe is very strong in embedded computing
  - The most energy-efficient computing technology today

Mont-Blanc project goals

- To develop a **European** Exascale approach
- Leverage **commodity** and embedded *power-efficient* technology

- Supported by two EU projects so far, 25.5M€ total:
  - FP7 Objective ICT-2011.9.13 Exascale computing, software and simulation:
    - 3-year IP Project (October 2011 - September 2014)
    - Total budget: 14.5 M€ (8.1 M€ EC contribution),
  - FP7 Objective ICT-2013.12.1 Exascale computing platforms, software and applications:
    - 3-year IP Project (October 2013 - September 2016)
    - Total budget: 11.5 M€ (8.0 M€ EC contribution),

Mont-Blanc 2 consortium

First, vector processors dominated HPC

- 1st Top500 list (June 1993) dominated by traditional HPC (vector) architectures
  - Cray vector,41%
  - MasPar SIMD, 11%
  - Convex/HP vector, 5%
- Fujitsu *Wind Tunnel* is #1 1993-1996, with 170 GFLOPS

Then, commodity displaced special purpose

- ASCI Red, Sandia
  - 1997, 1 TFLOPS
  - 9,298 cores @ 200 MHz
  - Intel Pentium Pro
    - Upgraded to Pentium II Xeon, 1999, 3.1 TFLOPS
- ASCI White, LLNL
  - 2001, 7.3 TFLOPS
  - 8,192 proc. @ 375 MHz,
  - IBM Power 3

Transition from Vector parallelism to Message-Passing Programming Models

The killer microprocessors

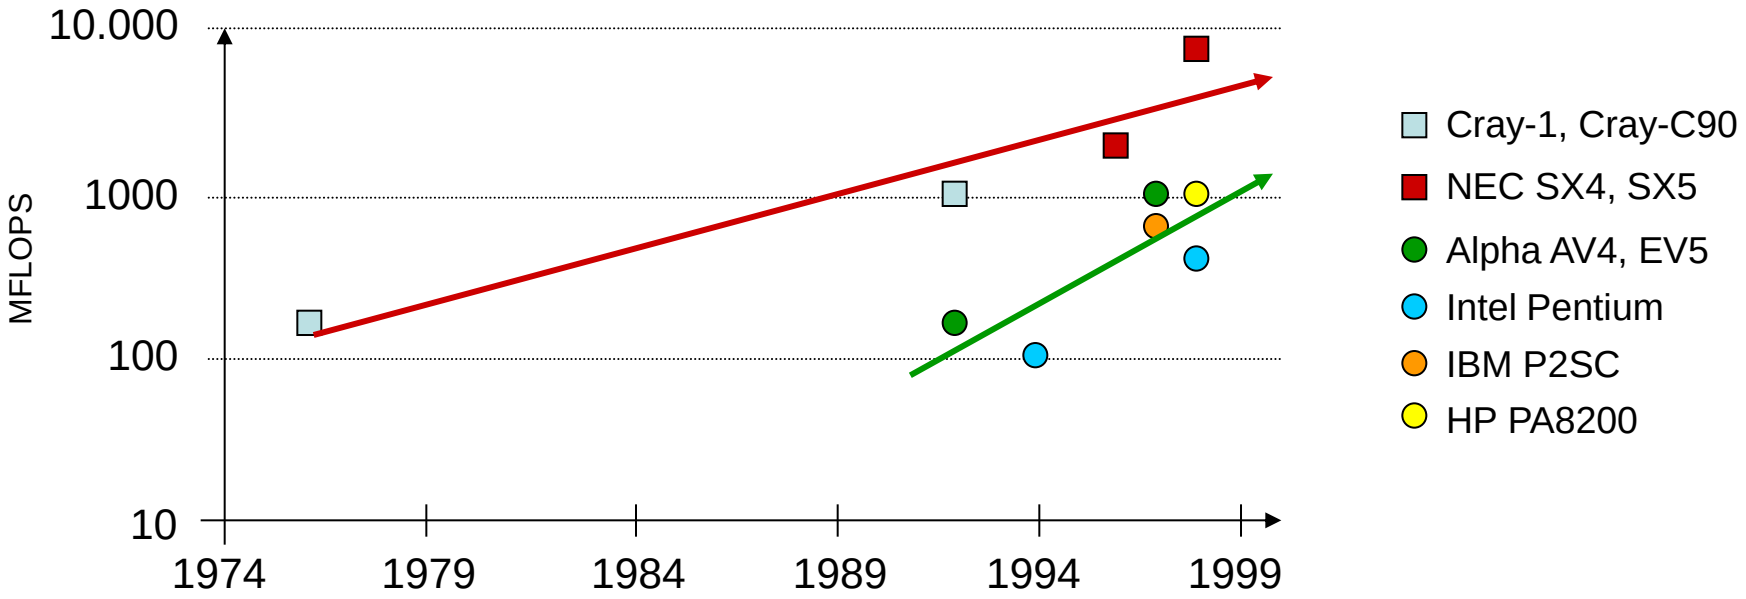

- Microprocessors killed the Vector supercomputers
  - They were not faster ...
  - ... but they were significantly cheaper (and more energy efficient)
- Needed 10 microprocessors to achieve the performance of 1 Vector CPU – but this was **still** a win!

The killer mobile processors™

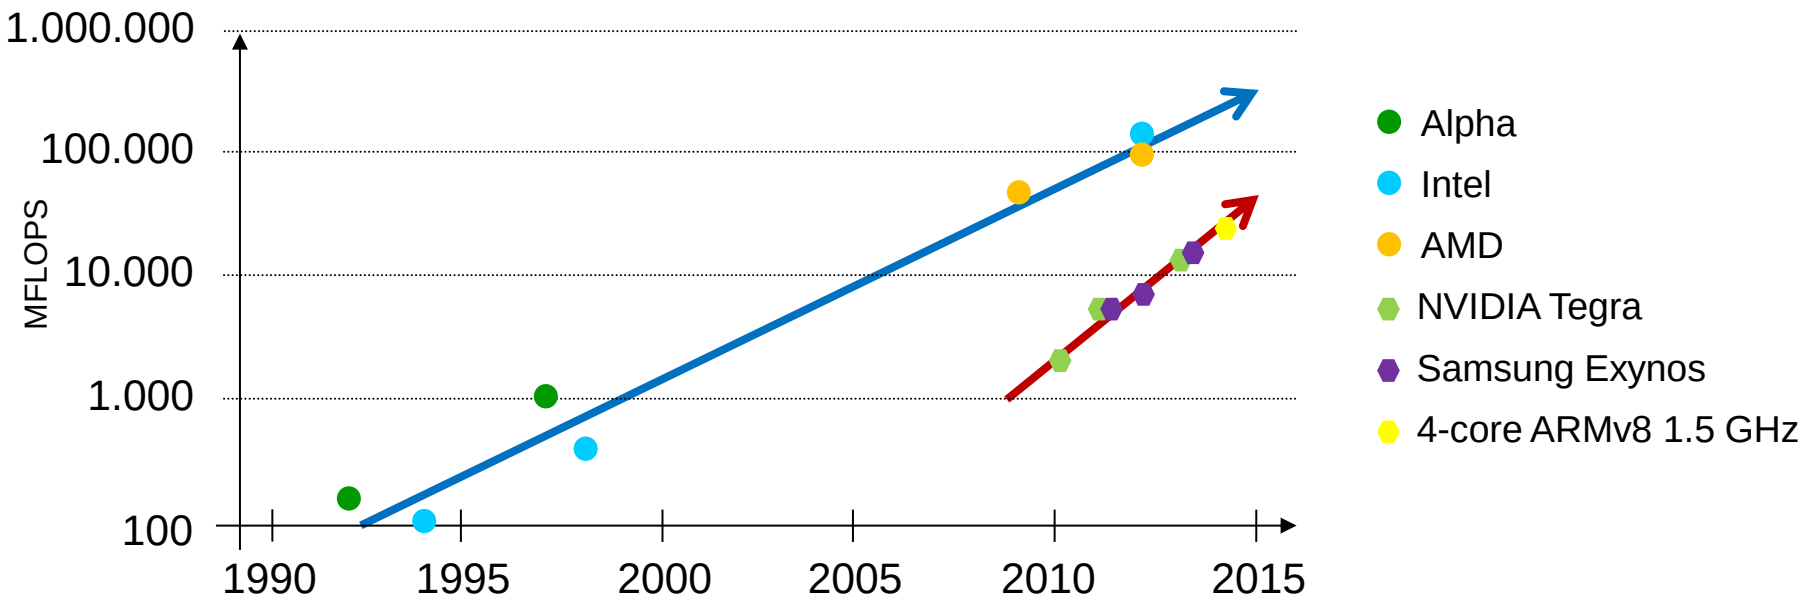

- Microprocessors killed the Vector supercomputers
  - They were not faster ...
  - ... but they were significantly cheaper and greener
- History may be about to repeat itself …
  - Mobile processors are not faster …
  - … but they are significantly cheaper (and greener?)

Mont-Blanc roadmap to Exascale

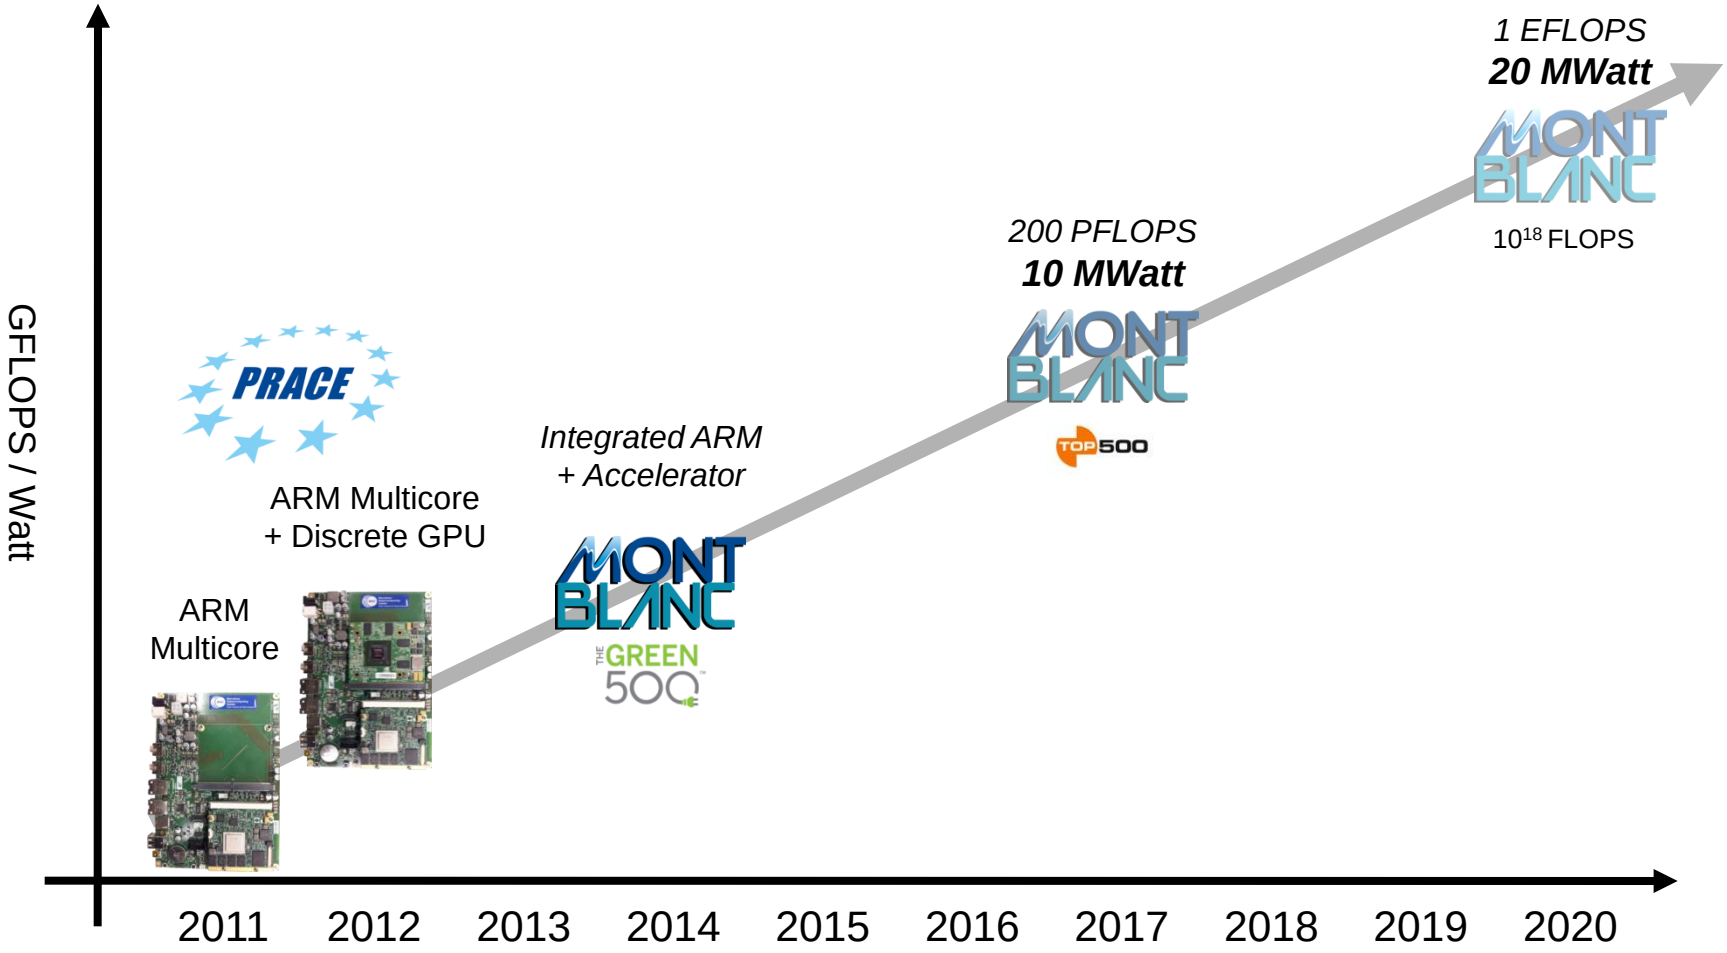

- The goal is for Europe to become a serious force in the HPC Exascale race
  - Mont-Blanc NG system by 2017, 200 PFLOPS on **10 Mwatt**
  - Mont-Blanc EX system by ~2020, 1 EFLOPS on **20 MWatt**

Tibidabo: The first ARM HPC multicore cluster

Q7 Tegra 2
2 x Cortex-A9 @ 1GHz
2 GFLOPS
5 Watts (?)
0.4 GFLOPS / W

Q7 carrier board
2 x Cortex-A9
2 GFLOPS
1 GbE + 100 MbE
7 Watts
0.3 GFLOPS / W

1U Rackable blade
8 nodes
16 GFLOPS
65 Watts
0.25 GFLOPS / W

2 Racks
32 blade containers
256 nodes
512 cores
9x 48-port 1GbE switch

512 GFLOPS
3.4 Kwatt
0.15 GFLOPS / W

- Proof of concept
  - It is possible to deploy a cluster of smartphone processors
- Enable software stack development

Samsung-based Mont Blanc Daughter Card

- Compute card based on the Samsung Exynos 5 Dual mobile phone chip
- Each daughter card is a full HPC compute node
  - (CPU + GPU) + DRAM + NAND storage + Ethernet NIC

MONT-BLANC : Prototype architecture

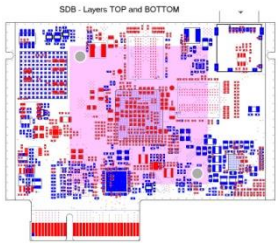

Exynos 5 Compute card
1x Samsung Exynos 5 Dual
2 x Cortex-A15 @ 1.7GHz
1 x Mali T604 GPU
6.8 + 25.5 GFLOPS (peak)
~10 Watts
3.2 GFLOPS / W (peak)

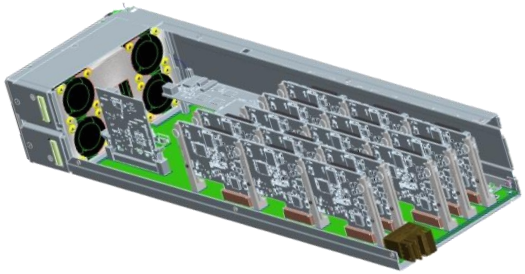

Carrier blade
15 x Compute cards
485 GFLOPS
1 GbE to 10 GbE
200 Watts (?)
2.4 GFLOPS / W

7U blade chassis
9 x Carrier blade
135 x Compute cards
4.3 TFLOPS
2 KWatt
2.2 GFLOPS / W

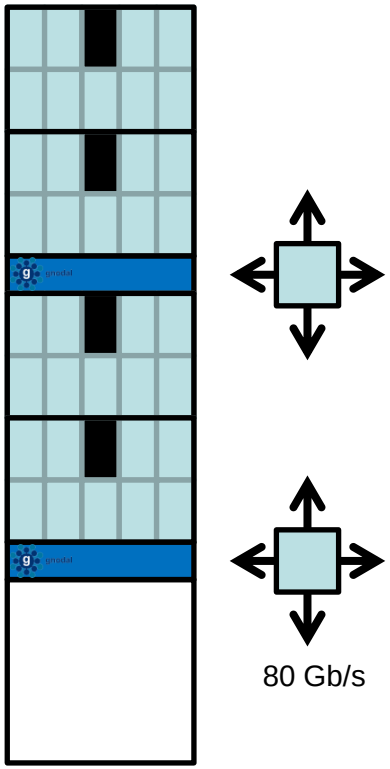

1 Rack
4 x blade cabinets
36 blades
540 compute cards
2x 36-port 10GbE switch
8-port 40GbE uplink

17.2 TFLOPS (peak)
8.2 KWatt
2.1 GFLOPS / W (peak)

80 Gb/s

Conclusions

- The bigger the science, the larger the supporting infrastructure!
- Our HPC systems have grown to tens of MWs to process all the data we are generating
- The processors developed for the mobile device space might yield a significant jump in energy efficiency for computation
- The Mont Blanc project is exploring Europe's strengths in mobile technology to exploit the opportunity arising in energy-dominated HPC