MONT-BLANC

http://www.montblanc-project.eu

European scalable and power efficient HPC platform based on low-power embedded technology

Alex Ramirez

Barcelona Supercomputing Center

Technical Coordinator

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.

Mont-Blanc project goals

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

- Supported by EU FP7 with 16M€ under two projects:
  - Mont-Blanc: October 2011 – September 2014
    - 14.5 M€ budget (8.1 M€ EC contribution), 1095 Person-Month
  - Mont-Blanc 2: October 2013 – September 2016
    - 11.3 M€ budget (8.0 M€ EC contribution), 892 Person-Month

Mont-Blanc 1: Project objectives

- To deploy a prototype HPC system based on **currently available** energy-efficient embedded technology
  - Scalable to 50 PFLOPS on 7 MWatt
    - Competitive with Green500 leaders in 2014
  - Deploy a full HPC system software stack

- To design a next-generation HPC system and new embedded technologies targeting HPC systems that would **overcome most of the limitations** encountered in the prototype system
  - Scalable to 200 PFLOPS on 10 MWatt
    - Competitive with Top500 leaders in 2017

- To port and optimise a small number of **representative Exascale applications** capable of exploiting this new generation of HPC systems
  - Up to 11 full-scale applications

Mont-Blanc 2: Project objectives

- Complement the effort on the Mont-Blanc **system software stack**
  - Development tools: debugger, performance analysis
  - Resiliency
  - ARMv8 ISA
- Initial definition of the Mont-Blanc **Exascale architecture**
  - Performance & power models for DSE
- Continued tracking and **evaluation of ARM**-based products
  - Deployment and evaluation of small developer kit clusters
  - Evaluation of their suitability for HPC
- Continued **support** for the Mont-Blanc consortium
  - Mont-Blanc prototype(s) operation
  - OmpSs developer support
  - Increased dissemination effort

Commodity components drive HPC

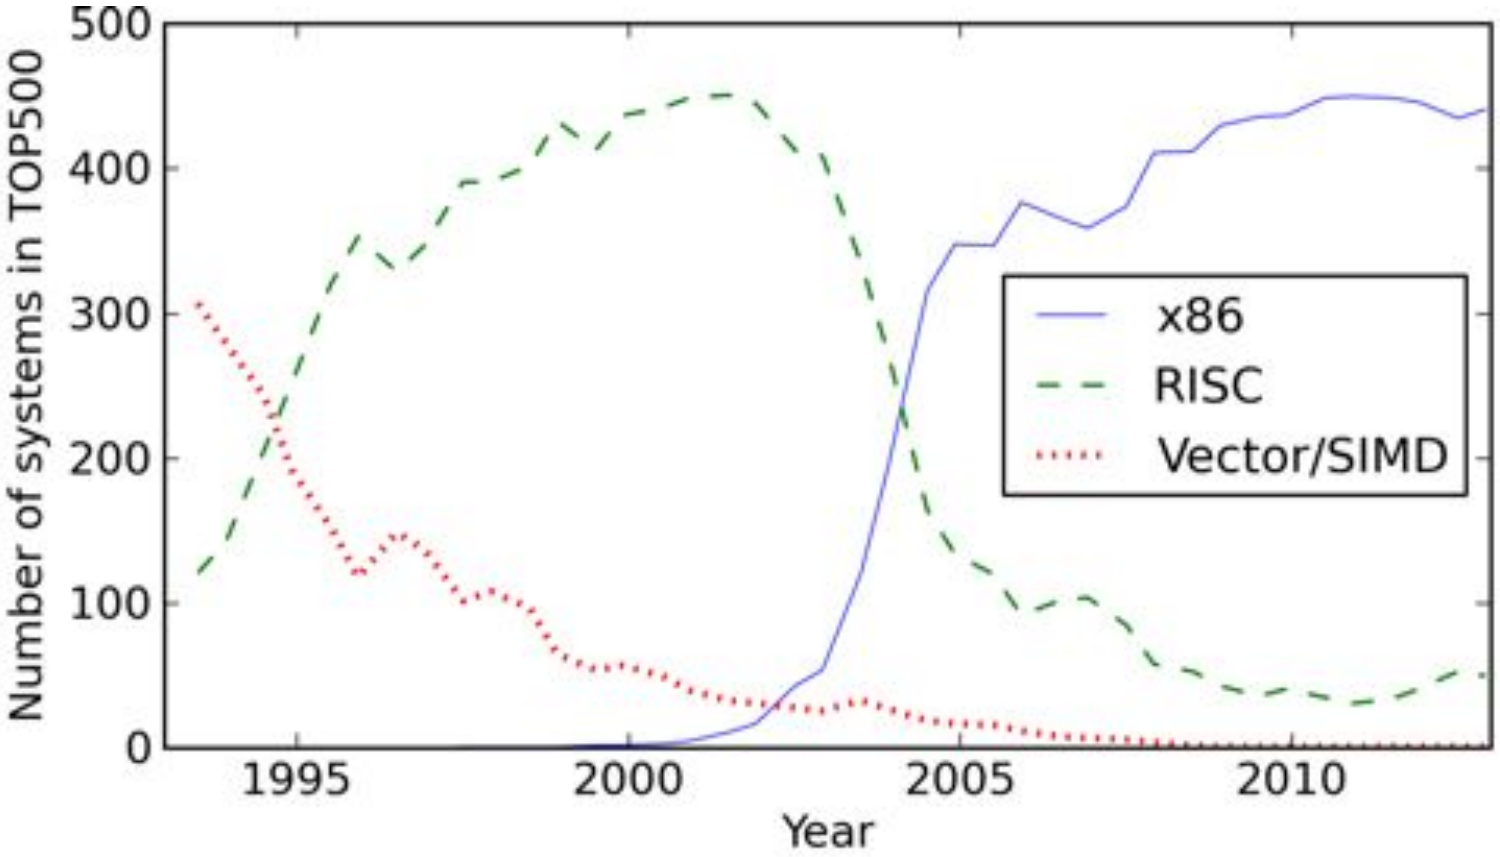

- RISC processors replaced vectors
- x86 processors replaced RISC
  - Vector processors survive as (widening) SIMD extensions

The Mont-Blanc prototype

- 6 BullX chassis
- 54 Compute blades
- 810 Compute cards
  - 1620 CPU
  - 810 GPU
  - 3.2 TB of DRAM
  - 52 TB of Flash
- 26 TFLOPS
- 18 KWatt

Limitations of current mobile processors for HPC

- 32-bit memory controller
  - Even if ARM Cortex-A15 offers 40-bit address space
- No ECC protection in memory
  - Limited scalability, errors will appear beyond a certain number of nodes
- No standard server I/O interfaces
  - Do NOT provide native Ethernet or PCI Express
  - Provide USB 3.0 and SATA (required for tablets)
- No network protocol off-load engine
  - TCP/IP, OpenMX, USB protocol stacks run on the CPU
- Thermal package not designed for sustained full-power operation

- **All these are implementation decisions, not unsolvable problems**
  - Only need a business case to jusitfy the cost of including the new features … such as the HPC and server markets

There is no free lunch

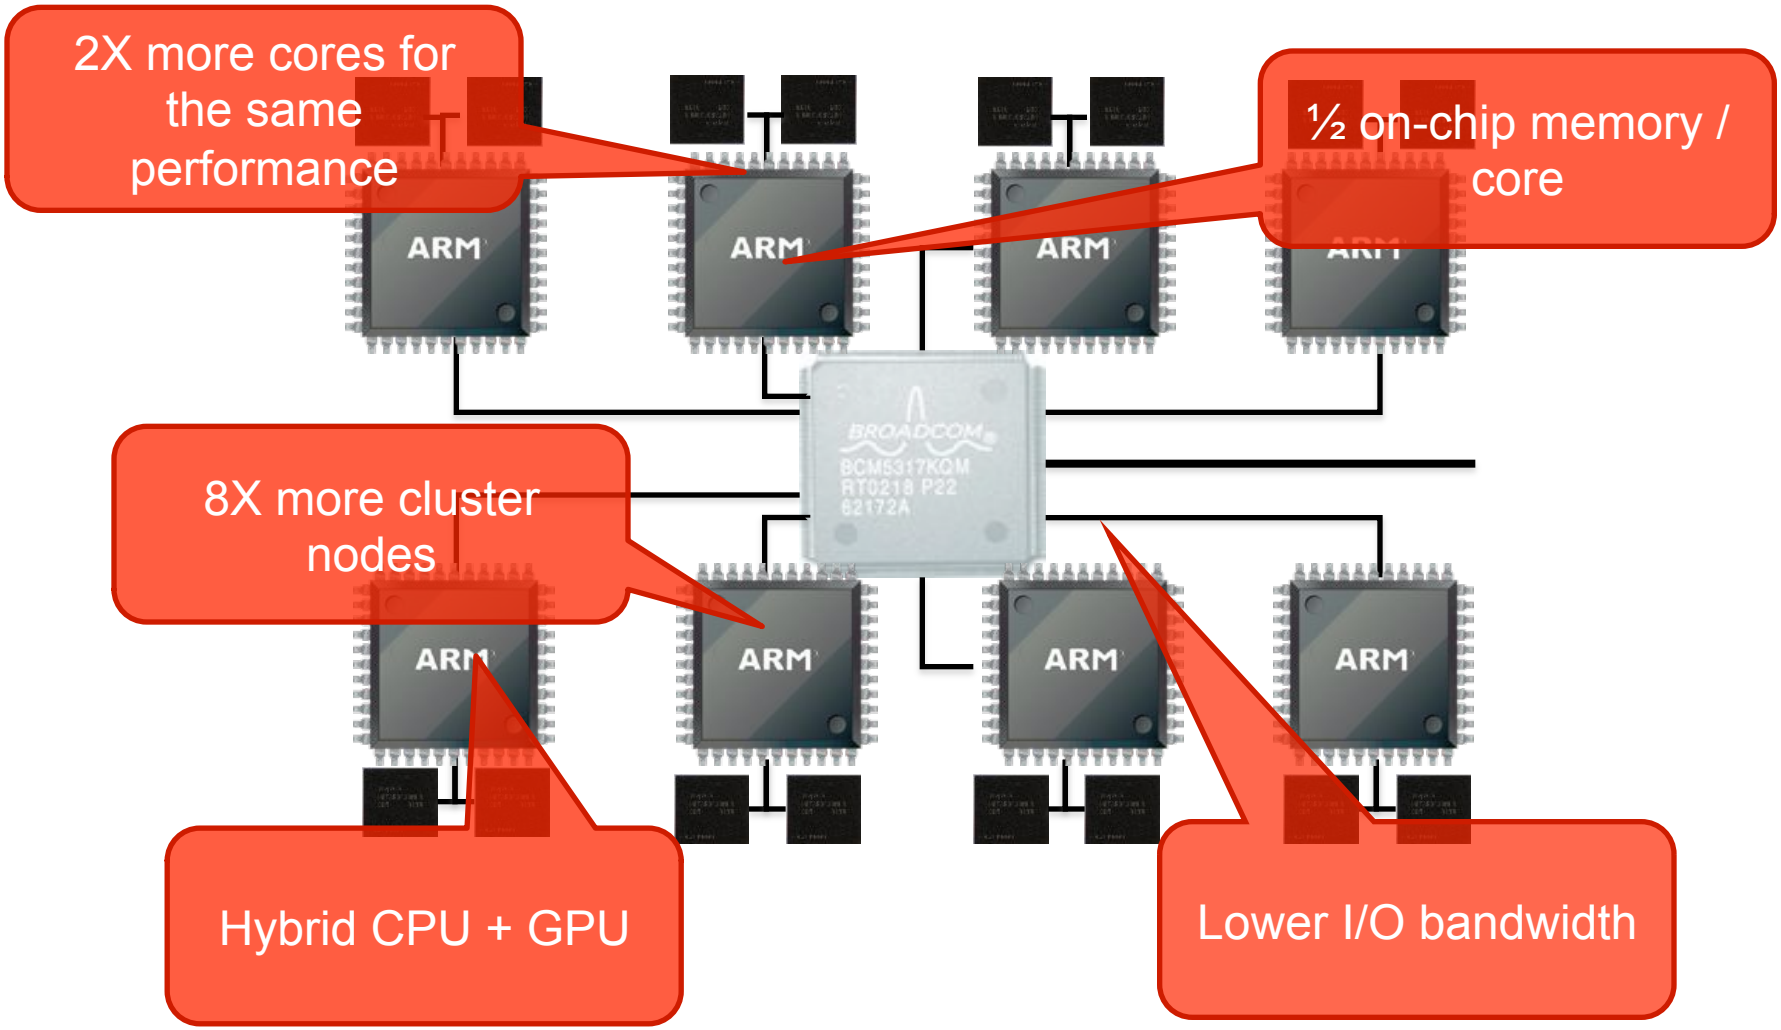

OmpSs runtime layer manages architecture complexity

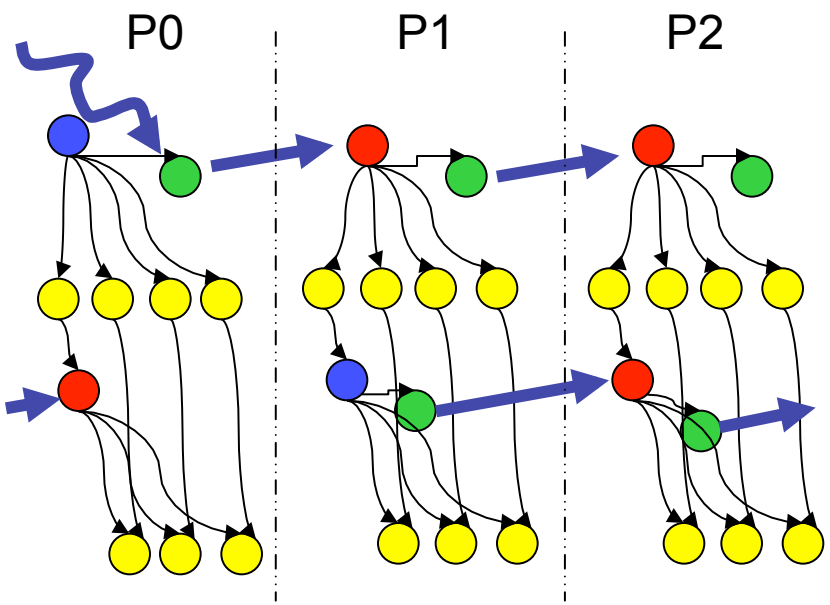

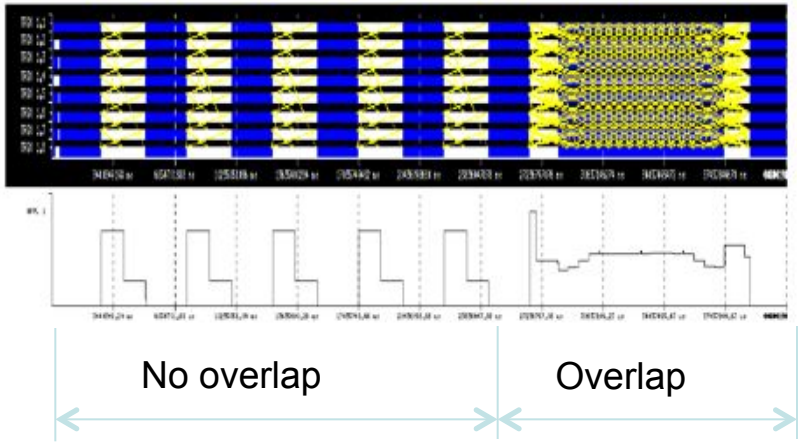

- Programmer exposed a simple architecture
- Task graph provides lookahead
  - Exploit knowledge about the future
- Automatically handle all of the architecture challenges
  - Strong scalability
  - Multiple address spaces
  - Low cache size
  - Low interconnect bandwidth
- Enjoy the positive aspects
  - Energy efficiency
  - Low cost

Porting applications to Mont-Blanc

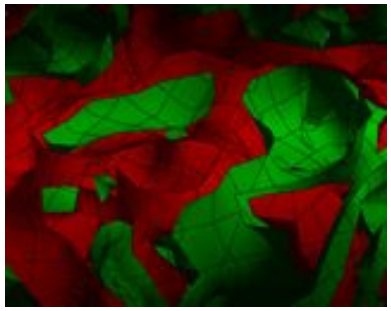

BQCD
Particle physics

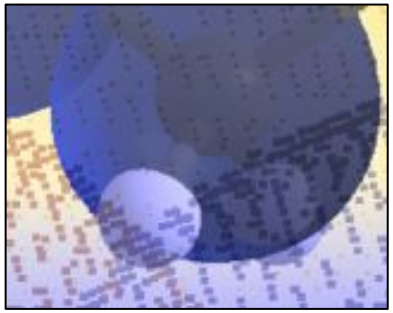

BigDFT *
Elect. Structure

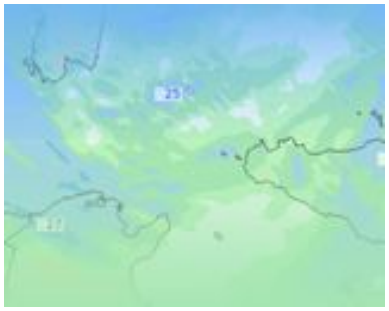

COSMO
Weather forecast

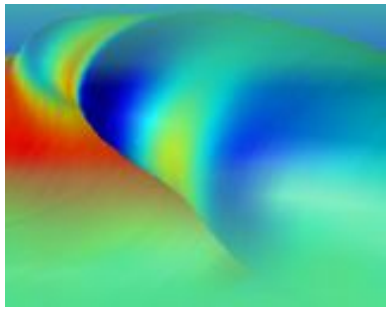

EUTERPE
Fusion

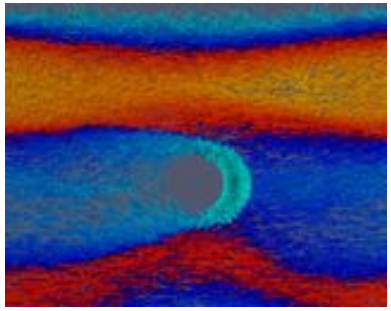

MP2C
Multi-particle collisions

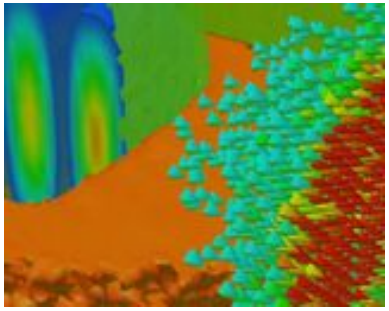

PEPC
Coulomb + Grav. Forces

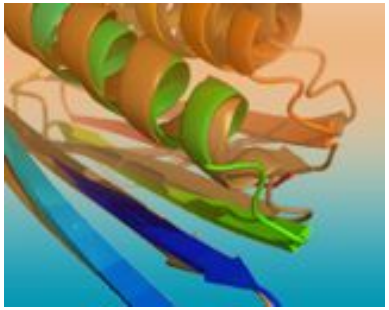

ProFASI
Protein folding

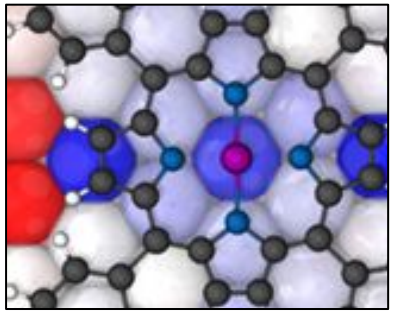

Quantum ESPRESSO *
Elect. Structure

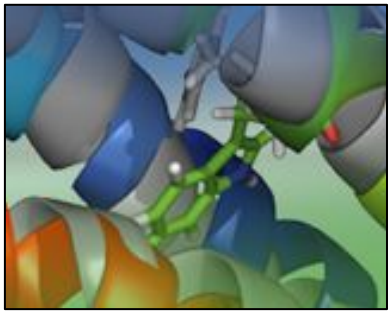

SMMP *
Protein folding

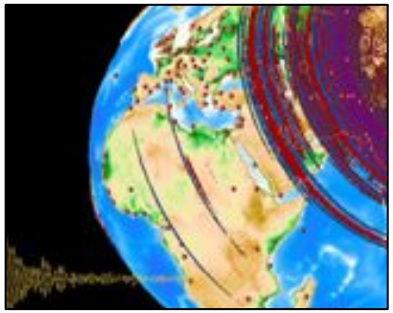

SPECFEM3D *
Wave propagation

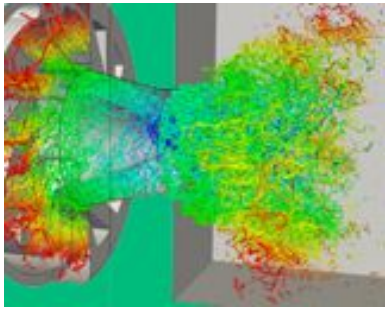

YALES2
Combustion

* Already GPU capable
(CUDA or OpenCL)

Conclusions

- Need sustainable EFLOPS technology
  - Limited power + space + cost
- Europe has a strong position in embedded computing
  - Energy efficiency
  - Commodity market
- BSC has a strong position in parallel programming models
  - OmpSs tasking model extends OpenMP 4.1
- Leverage on them to build a new class of sustainable computer
  - Faster, cheaper, more efficient


montblanc-project.eu


MontBlancEU


@MontBlanc_EU