

[Back to Table of contents](#)

Primeur live 20120621

### TOP500

*Europe shoots up in TOP500: 4 systems in the top 10*

*ISC'12 caps 27th conference with record attendance, awards and more*

### Hardware

*CRESTA is focusing on co-design to try and materialize exascale*

*PRACE Prototypes evaluation in ISC'12 BOF session*

*DEEP - advanced high performance computing with Knights Corner booster*

*Mont Blanc, first European supercomputer built with embedded technology*

### Applications

*The PRACE Award winners 2012 announced*

### Company news

*AMD announces Cloudera-certified Collax server cluster to handle Big Data at significantly lower cost*

*AMD supercomputing leadership continues with broader developer ecosystem and latest TOP500*

*Supermicro FatTwin takes centre stage at International Supercomputing Conference 2012*

*Acer to resell Altair's PBS Professional and expand options for HPC users*

## Mont Blanc, first European supercomputer built with embedded technology

20 Jun 2012 Hamburg - *Alex Ramirez from the Barcelona Supercomputer Center presented the EC-funded Mont Blanc project. Mont Blanc is a European HPC Platform based on embedded energy-efficient technology. Many important decisions are still pending on how to build this powerful European supercomputer. The road to the top of the HPC mountain is a long one but the Mont Blanc team is very persistent and strong.*

Mont Blanc will be the prototype of a new class of energy-efficient supercomputers, built from mobile and low-cost technology and built on European industry strengths, Alex Ramirez stated. The leading energy-efficiency prototype will be ready by 2014 and the leading performance next-generation design is planned for 2017. The machine will be capable of running a portfolio of Petascale applications.

At the start there will be 14 candidate applications. The partners will make a selection out of this.

Are the mobile-killers coming? Alex Ramirez asked the audience. Where is the sweet spot? Maybe in the low-end. Today we have a 1.8 ration in performance and 1.100 in cost. Tomorrow this will be a 1.2 ration in performance but still 1.100 in cost. This is the same reason why microprocessors killed the vector computers, explained Alex Ramirez. It is not so much the performance but much more the lower cost and power.

The prototypes are critical to push the software development. The team will start from COTS and move on to new developments.

The ARM processors do not have enough power to generate Gigaflops. Another solution will be needed. The compute element is homogeneous multicore. The architecture features an accelerator. It is discrete and integrated, including the interconnect and memory system.

Alex Ramirez said that they already know how to build a homogeneous multicore system. ARM cores are more energy-efficient but low power. High multicore density is required before CPU becomes the main power sink. The rest of the power goes to memory, interconnect, storage, and power supply.

The ARM multicore efficiency projections involve increasing multicore density which improves energy efficiency. This improves the ration of CPU power versus "glue" power. The 16-core ARM Cortex-A15 at 1 GHz would be competitive with BlueGene/Q.

The other options are heterogeneous multicore with an accelerator. The hybrid CPU and GPU systems are becoming common. The performance is dominated by the accelerator. The team wants to replace high-end CPU for low-power CPU to improve the energy efficiency.

Pedraforca has ARM and GPU. Alex Ramirez said that you can increase the performance density of an ARM cluster by adding a GPU accelerator.

The integration of CPU and GPU is a great solution. Prototypes for all the other options are needed. If we want to be better we must be different, Alex Ramirez stated. Integrated GPU has many advantages. Shared memory with CPU even allows separate address spaces today. The prototype soon will be cache coherent. No power is wasted on the PCIe bus and no power is wasted on the GDDR5 memory. Higher energy efficiency and lower cost are necessary.

The interconnect is Ethernet. There is a higher probability of finding an ARM SoC with integrated Ethernet NIC. The prototype will be open standard. It will be actively developed with multiple vendors at a lower cost and with a wider range of options. This provides an opportunity for customization.

The available interfaces in ARM SoC involve server class SoC with integrated Ethernet. It is best in terms of latency and power with a higher integration density but no integrated GPU accelerator.

Mobile class SoC has an external NIC with PCIe low latency and high bandwidth. The low latency interface is technically possible but there is no LLI Ethernet NIC available. The USB3.0 has a potential latency impact and is restricted to 1 GbE.

The 16-core Cortex-A15 at 2Ghz is already limited by 1 GbE. The team will be increasing the performance with GPU, explained Alex Ramirez.

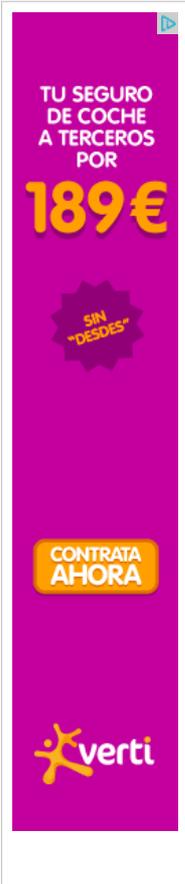
There is overlap communication with computation. The OmpSs runtime system automatically overlaps MPI messages with computation. The team should move MPI out of the critical path. There is asynchronous execution and the team has to avoid that resources are being idle. They can sustain the performance with a lower bandwidth network.

The memory system involves a performance versus power versus density trade-off. This means DDR3 versus PDDR3. There is an error correction code. A one tablet or phone crash every 6 months is not an issue but a 5000 node system may become unusable. This shows how badly the team needs ECC in order to measure memory errors.

The Memtester benchmark is similar to memtest86. It works in the user space with OS running. There are tests for various patterns and operations.

Alex Ramirez concluded that the Mont Blanc will have ARM multicore and integrated GPU Ethernet NIC with a high density packaging. He stressed that many important decisions are still pending. The team will be busy contacting providers and comparing alternatives.

Leslie Versweyeld



TU SEGURO DE COCHE A TERCEROS POR **189€**

SIN "DESDES"

CONTRATA AHORA

verti

