



November 22, 2013

Mont Blanc Forges Cluster from Smartphone Chips

Timothy Prickett Morgan



The Mont Blanc project, an effort by a number of European supercomputing centers and vendors that seeks to create an energy-efficient supercomputer based on ARM processors and GPU coprocessors, has put together its third prototype. That is one more step on the path to an exascale system.

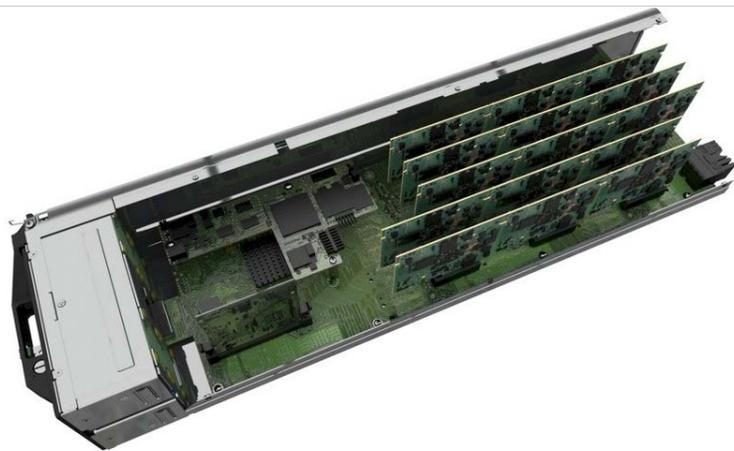
The third generation machine, which is being shown off at the SC13 conference in Denver this week, is by far the most elegant one that the Mont Blanc project has created thus far. This prototype supercomputer actually bears the name of the project this time around, and was preceded by the Tibidabo and Petraforca clusters, which were based on a different collection of ARM processors and GPU accelerators.

Just because this design is elegant, don't get the wrong idea, though. The [Mont Blanc](#) machine is still a prototype, cautions Alex Ramirez, leader of the Heterogeneous Architectures Research Group at BSC who heads up the Mont Blanc project.

"In order to make this a production product, we would have to go through at least one more generation," he says.

It stands to reason that the Mont Blanc project is waiting for the day when 64-bit ARM chips with integrated interconnects and faster GPUs are available before going into production. But for now, software can be ported to these prototypes and things can be learned about where the performance bottlenecks are and what reliability issues there might be.

The exact size of the Mont Blanc prototype cluster has not been determined yet, but Ramirez says it will have two or three racks of ARM-powered nodes. "It will be big enough to make scalability and reliability claims, but we are trying to keep the cost down on a machine that is not a production system," he says.

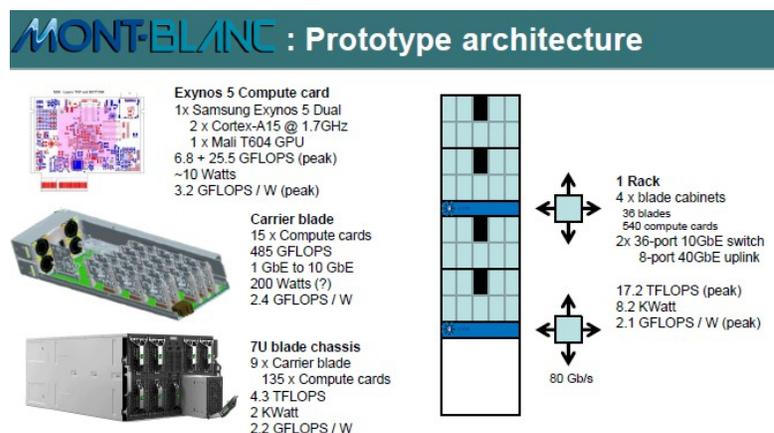


The server node in the Mont Blanc system is based on the Exynos 5 system-on-chip made by Samsung, which is a dual-core ARM Cortex-A15 with an ARM Mali-T604 GPU on the die. The ARM CPU portion of the system-on-chip has about twice the performance of the quad-core Cortex-A9 processor used on the Petraforca prototype that was put together earlier this year. (There were actually two versions, but the second one is more important.) That machine used Nvidia Tesla K20 GPU coprocessors to test out how a wimpy CPU and a brawny GPU might be married. Specifically, the ARM processors, which were Tegra 3 chips running at 1.3 GHz, were put into a Mini-ITX system board with one I/O slot that was linked to a PCI-Express switch that in turn had one GPU and one ConnectX-3 40 Gb/sec InfiniBand adapter card.

The dual-core Exynos 5 chip from Samsung is used in smartphones, runs at 1.7 GHz, and has a quad-core Mali-T604 GPU that supports OpenCL 1.1. It has a dual-channel DDR3 memory controller and a USB 3.0 to 1 Gb/sec Ethernet bridge. Each Mont Blanc node is a daughter card made by Samsung that has the CPU and GPU, 4 GB of memory (1.6 GHz DDR3), a microSD slot for flash storage, and a 1 Gb/sec Ethernet network interface. All of this is crammed onto a daughter card that is 3.3 by 3.2 inches that has 6.8 gigaflops of compute on the CPU and 25.5 gigaflops of compute on the GPU for something around 10 watts of power. That works out to around 3.2 gigaflops per watt at peak theoretical performance.

The Mont Blanc system is using the Bull B505 blade server carrier and the related blade server chassis and racks to house multiple ARM server nodes. In this case, the blade carrier is fitted with a custom backplane that has a Broadcom Ethernet crossbar switch on it that links fifteen of these ARM compute nodes together. Every blade in the carrier has an Ethernet bridge chip, made by ASIX Electronics, that converts the USB port into Ethernet and then lets it hook into that Broadcom switch in the carrier.

Here is how you stack up the Mont Blanc rack:



In this particular setup, says Ramirez, the location had some power density and heat density restrictions, so it was limited to four Bull blade server chassis. But the system is designed to support up to six chassis if the datacenter has enough power and cooling.

Each blade has fifteen nodes, and is a cluster in its own right. The blade delivers on the order of 485 gigaflops of compute and will burn about 200 watts. (Ramirez is estimating

because he has not actually been able to do the wall power test yet because the machines just came out of the factory a few days prior to SC13.) That works out to 2.4 gigaflops per watt or so after the overhead of the network is added in.

The 7U blade chassis can hold nine carrier blades, for a total of 135 compute nodes. That works out to 4.3 teraflops in the aggregate per chassis at around 2 kilowatts of power, or 2.2 gigaflops per watt. With two 36 port 10 Gb/sec Ethernet switches to link the chassis together and 40 Gb/sec uplinks to hook into other racks, a four-chassis rack would deliver 17.2 teraflops of computing in an 8.2 kilowatt power envelope, or about 2.1 gigaflops per watt. With six blade chassis, you can get 25.8 teraflops into a rack. That is 810 chips in total per rack, by the way, with a total of 1,620 ARM cores and 3,240 Mali GPU cores.

This Mont Blanc effort will get very interesting next year, when many different ARMv8 processors, sporting 64-bit memory addressing and integrated interconnects, become available from a variety of vendors, including AppliedMicro, Calxeda, AMD, and maybe others like Samsung. Many of the components that had to be woven together in this third prototype will be unnecessary, and the thermal efficiency of the cluster will presumably rise dramatically once these features are integrated on the chips. These future ARM chips will also come with server features, such as ECC memory protection and standard I/O interfaces like PCI-Express.

“There will be enough providers that at least one of them will have exactly the kind of part you want at any given time,” says Ramirez, a bit like a kid in a candy store.

The Mont Blanc project was established in October 2011 and is a five-year effort that is coordinated by the Barcelona Supercomputer Center in Spain. British chip maker ARM Holdings, French server maker Bull, French chip maker STMicroelectronics, and British compiler tool maker Alinea are vendor participants in the Mont Blanc consortium. The University of Bristol in England, the University of Stuttgart in Germany, and the CINECA consortium of universities in Italy are academic members of the group, and the CEA, BADW-LRZ, Juelich, and BSC supercomputer centers are also members. So are a number of other institutions that promote HPC in Europe, including Inria, GENCI, and CNRS.

Mont Blanc was originally a three year project with a relatively modest budget of €14.5 million, and it has secured an additional €8.1 million in funding from the European Commission to extend it two more years. The funds are not just being used to create an exascale design, but also to create a parallel programming environment that will run on hybrid ARM-GPU machines as well as creating check pointing software to run on the clusters.

Share this: 

Tags: [Mont Blanc](#), [SC13](#)

Only registered users may comment. Register using the form below.

Check off newsletters you would like to receive *

HPCwire

EnterpriseTech

Datanami

Digital Manufacturing Report

Technology Conferences & Events

Advanced Computing Job Bank

Technology Product Showcase

Email *

Name *

First

Last