



DATA CENTRE > HPC

Barcelona taps Tesla and Tegra for next-gen hybrid supercomputer

Singing the praises of ceepie-geepies for PRACE

By Timothy Prickett Morgan, 12th June 2013

4

[Free whitepaper – A new generation of competitive distinction](#)

RELATED STORIES

GTC 2013 Calm a CARMA drama chameleon: Barça super waves ARMs, GPUs


Can't wait for Nvidia? Try these Italian baby ARM clusters with GPU options

SC11 Spaniards prototype ARM-GPU hybrid supercomputer

SC11 ARM+GPU = dream super?

Barcelona boots Europe's most powerful

The Barcelona Supercomputing Center has been monkeying around with the combination of low-powered processors and relatively low-end graphics chips. Well, BSC is getting ready to take its ceepie-geepie prototyping up another notch by marrying a baby ARM processor aimed at smartphones and tablets with a full-on GPU coprocessor.



BSC is the Church of the Ceepie-Geepie, quite literally

The prototype cluster, which is to be called Pedraforca, will take the existing Tegra 3 processor from Nvidia - as implemented on the "Kayla" system [launched by the chipmaker](#) in conjunction with motherboard maker

SECO back in March at the GPU Technical Conference.

The predecessor Carma system from Nvidia (short for CUDA-ARM), also made with SECO, put a four-core Tegra 3 chip based on the Cortex-A9 processor running at

The Register uses cookies. Some may have been set already. [Read about managing our cookies.](#) Please click the button to accept our cookies. If you continue to use the site, we'll assume you're happy to accept the cookies anyway.

I'm fine with this [X]



processing at double precision.

Me gusta

4

Tweet

19

With the Kayla system board, which [comes in a MiniITX form factor](#), the CPU side is again a Tegra 3 chip from Nvidia, but the board has a PCI Express 2.0 x16 link to hook a full-on Tesla GPU to the ARM processor.

BSC collaborated with Nvidia and SECO to create the Kayla board, and built a prototype machine – its second ARM-GPU hybrid – on that card last fall. Tesla coprocessors based on Nvidia's GF108, GK104, and GK107 graphics processors are supported with the Kayla system.

It would be nice to have a faster PCI-Express link to hook the CPU and GPU together, and having only 2GB of memory for four cores might be a little skinny, too. A gigabit Ethernet link is not going to break any performance barriers, either. But at €349 per Kayla system, experimenting and seeing how software could run on such a ceepie-geepie is not exactly going to bust the budget, not even for a dense-packed rack of these little beastsies.

With the Pedraforca system, the third generation of ARM-based ceepie-geepies to be prototyped by BSC, the supercomputer center will again use a board that



has a Tegra 3 processor and Sumit Gupta, general manager of the Tesla Accelerated Computing business unit at Nvidia, says that the nodes will use a Tesla K20 to the ARM CPU's math homework.

The ARM processor nodes will be linked to each other using 40Gb/sec InfiniBand adapters and switches from Mellanox Technologies, giving it a substantial performance boost. This is not just because of the increase in bandwidth, but thanks to Remote Direct Access Memory (RDMA) in the InfiniBand protocol, which will allow the CPUs to talk to each other over the network without having to go through the network software stack in the Linux operating system on the cluster. And, thanks to the much-improved GPUDirect feature in the "Kepler" GPUs from Nvidia, the GPUs can talk over InfiniBand to each other without speaking to the CPU, too.

This, as it turns out, is important, as are [the Hyper-Q and Dynamic Parallelism features](#) of the [high-end Kepler GPUs from Nvidia](#).

"Fermi-class GPUs were too limited and we had to rely too much on the Tegra CPU with them," explains Alex Ramirez, leader of the Heterogeneous Architectures Research Group at BSC, to *EI Reg*. But with GPUDirect combined with InfiniBand and Hvner-Q (which allows the ARM CPU to queue up 32 MPI tasks at the same time on

The Register uses cookies. Some may have been set already. [Read about managing our cookies](#). Please click the button to accept our cookies. If you continue to use the site, we'll assume you're happy to accept the cookies anyway.

I'm fine with this [X]

CPU for permission), now BSC can start testing software in earnest on a ceepie-geepie.

For workloads that don't bother the CPU much, Ramirez says that the Pedraforca cluster should get just about the same performance as a Xeon cluster that is offloading most of its work to the GPU, but without all that Xeon heat and cost.

You might be wondering how a Kayla system from SECO is able to have both a Tesla K20 GPU coprocessor and an InfiniBand ConnectX-2 adapter card both plugged into them, since there is only one x16 slot. Ramirez says that BSC is getting a PCI-Express switch from PLX Technologies and putting it into that single slot and then plugging in the InfiniBand adapters and Tesla GPUs into the switch. It looks like there will be multiple PLX switches in the cluster, which Ramirez hopes to scale up to 128 nodes when it is installed in July.

The Pedraforca machine is partially funded by the Partnership for Advanced Computing in Europe (PRACE) initiative. The compute nodes will be manufactured by E4 Computer Engineering and Bull is being hired to do the system integration. ®

[Free whitepaper – A new generation of competitive distinction](#)

COMMENTS

[Post your comment](#) [House Rules](#) | [Send Corrections](#)

Latest Comments

Destroy All Monsters

Oh Yurop!

> Bull is being hired to do the system integration

I foresee failure.

▲ 0 ▼ 0

All Reader Comments (4)

Posted Wednesday 12th June 2013 21:02 GMT

[Report abuse](#)



tpm

Posted Wednesday 12th June 2013 20:20 GMT

[Report abuse](#)



↪ **Re: This article needs a good proof read**

I was having a high dyslexia morning. Apologies. Now you know why I call them ceepie-geepies. . .

▲ 0 ▼ 0

The Register uses cookies. Some may have been set already. [Read about managing our cookies.](#) Please click the button to accept our cookies. If you continue to use the site, we'll assume you're happy to accept the cookies anyway.

I'm fine with this [X]

Anonymous Coward

Posted Wednesday 12th June 2013 18:01 GMT

[Report abuse](#)

Tesla and tegra?

Xxx giga flops and 0-60 in 4 seconds??

▲ 0 ▼ 0



[◀ Previous Story](#)

[Next Story ▶](#)

WHITEPAPERS



Top 10 SIEM implementer's checklist

An automated discovery driven approach that shortens your path to comprehensive monitoring and prevents security blind spots.



The new Office Garage series:

News and tips for IT professionals who plan, implement, and maintain Office.



The 7 deadly sins of backup and recovery


This paper guides IT executives and system administrators in designing and implementing business backup and disaster recovery systems.



IT infrastructure monitoring strategies

Secure support mobile device access to sensitive corporate IT resources with confidence.

MORE FROM THE REGISTER



72

MacBook Air now uses PCIe flash... but who'd Apple buy it from?

Was Not Woz?



52

Users rage as Fasthosts virtual servers go titsup... again

'What on Earth did you do?'



143

Windows NT granddaddy OpenVMS taken out back, single gunshot heard

HP signals end to legendary big iron beast

MORE FROM THE REGISTER

[Send us News Tips](#)

[Week's Headlines](#)

[Reg Archive](#)

[Top 20 Stories](#)

[eBooks](#)

[Webcasts](#)

The Channel



© Copyright 19

The Register uses cookies. Some may have been set already. [Read about managing our cookies.](#) Please click the button to accept our cookies. If you continue to use the site, we'll assume you're happy to accept the cookies anyway.

I'm fine with this [X]