

Blog

Microservers Brew in Europe's Labs

Christoforos Kachris

10/16/2014 06:00 AM EDT

0 comments [post a comment](#)

Tweet

4

Share

1

g+1

3

A handful of major
microserver projects,
mainly funded by the

European Union, presented data center servers based on energy-efficient embedded processors at the recent [HiPEAC EU Network of Excellence](#) event.

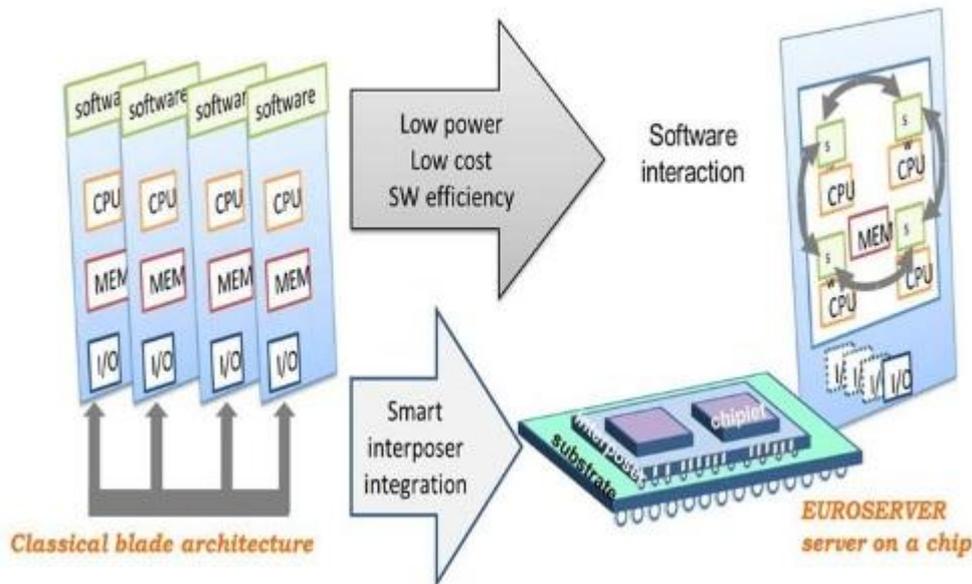
The [NanoStreams](#) project brings together European expertise in embedded systems design and high-performance computing (HPC) software to address the challenge of real-time analytics on fast data streams. NanoStreams uses an architecture and software stack that address the unique challenges of hybrid transactional-analytical workloads, encountered by emerging applications of real-time big-data analytics.

The NanoStreams processor is an amalgam of RISC cores and nano-cores, a new class of programmable, custom accelerators. Novel automatic compiler generation and parameterisation technology enables low-effort programming and integration of nano-cores into application-specific, many-core accelerators.

The project's proposed heterogeneous Analytics-on-Chip processor forms the backbone of the NanoStreams microserver. The system also leverages a hybrid DRAM-PCRAM memory system and a non-cache-coherent, scale-out architecture to achieve extreme energy-efficiency. It systems uses a mix of Calxeda SoC and Xilinx Zinq boards.

NanoStreams brings together a consortium of two academic institutions and three technology providers working in partnership with IBM and Credit Suisse. The main advantage of Nanostreams is that it is driven by the needs of real financial workloads and the proposed architecture is evaluated using real stock-exchange data.

In his talk on the project, Dimitrios Nikolopoulos, a professor and research director at Queen's University of Belfast, compared commodity x86 servers and microservers based on the Boston Viridis platform. He emphasized the need for fair comparisons of power consumption not only between the processors but between whole systems including the power supplies, storage and memory subsystems.



The Euroserver uses a processor on a 2.5-D chip stack with an interposer.

The [Euroserver](#) project applies state-of-the-art low-power ARM processors in a new server architecture that uses 3D "chipselets." The approach aims to reduce the acquisition cost of the system and scale the numbers of cores, memory capacity and I/O bandwidth it provides. New systems software supports both legacy and advanced features including system-wide virtualization to further reduce energy consumption.

The project takes a novel software approach of managing server resources as multiple coherent islands. It isolates and protects the multiple workloads from each other when they use shared server resources such as I/O, storage, memory, and interconnects. The main advantage of the architecture is it gives the user the option to move tasks and processes close to data instead of moving data around.

The Euroserver consortium includes ARM, STMicroelectronics, Eurotech, and OnApp, in addition to five academic institutions including CEA and FORTH.

Next page: Leveraging Exynos, attacking MapReduce

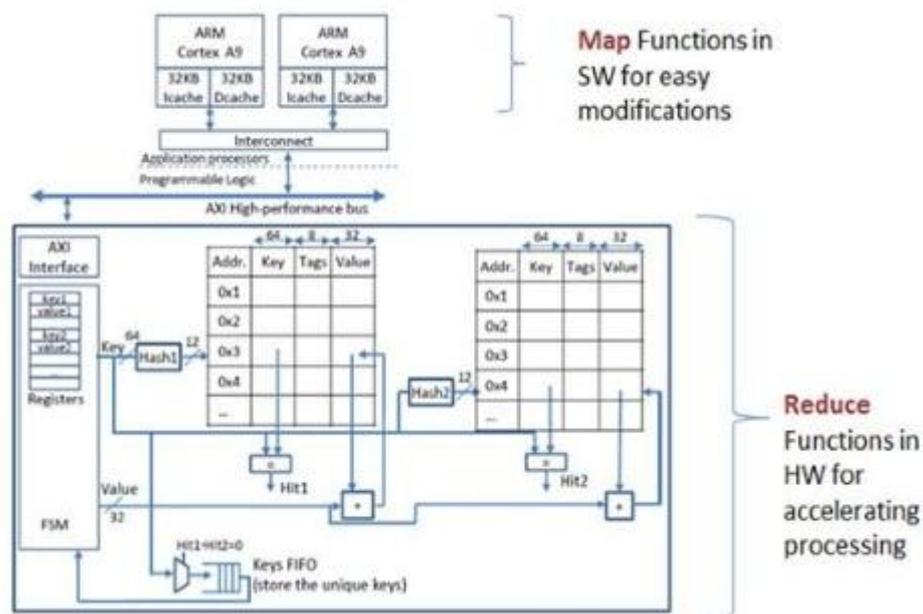
Leveraging Exynos, attacking MapReduce

The [Mont-Blanc](#) project aims to design a new type of computer architecture capable of setting future HPC standards, built from energy efficient solutions used in embedded and mobile devices. Mont-Blanc's primary objective is to deploy a prototype HPC system based on currently available energy-efficient embedded technology that can scaled to 50 PFlops while consuming 7 MW power. The Mont-Blanc consortium joins ARM, Bull and STMicroelectronics with research supercomputing centers such as the Barcelona Supercomputing Center and Inria.

The Mont-Blanc rack is based on a server-on-module (SoM), a small card that includes a Samsung Exynos 5 SoC using two Cortex A15 cores clocked at 1.7GHz and one Mali T604 GPU. The module also sports 4 GBytes DRAM and a micro-SD slot that can host up to 64 GBytes flash. Fifteen of these cards can be used to form a carrier blade that can sustain up to 485 GFlops consuming about 300 W. A blade chassis can host up to nine carrier blades (135 compute cards) sustaining 4.3 TFlops while consuming less than 3 KW.

The prototype under development showed that the processing gap between high performance processors and SoCs targeting smartphones and tablets is getting smaller and smaller. It supports typical programming frameworks used on supercomputing applications

and therefore it will be evaluated using real science applications, such as a particle physics simulation, protein folding, and weather forecasting.



The MapReduce coprocessor uses a mix of software and hardware-accelerated blocks.

Finally, [GreenCenter](#) is a national project that aims to develop a novel MapReduce coprocessor that can be incorporated to future server and microserver processors. The MapReduce coprocessor is used as a so-called "computational storage" module in which the key/value pairs of the MapReduce application can be stored and processed.

The module can be hosted close to the central processor and can be used to alleviate the processor for typical reduce tasks such as accumulators and average calculations. It also offloads the cache from storing the key/value pairs.

In the same way that TCP/IP offload engines are used to alleviate processors from the typical packet processing, the MapReduce coprocessor can alleviate cores from the typical MapReduce tasks. The flexibility of MapReduce remains unchanged because the Map tasks are still performed in software, while the Reduce tasks are handled by the MapReduce coprocessor.

— Christoforos Kachris is a Senior Researcher at Athens Information Technology in Peania, Greece.

[EMAIL THIS](#) [PRINT](#) [COMMENT](#)

Copyright © 2014 UBM Electronics, A UBM company, All rights reserved. [Privacy Policy](#) | [Terms of Service](#)