

BSC researchers demonstrate that scaling-in on large-memory nodes saves up to 52% of energy consumed in HPC systems

Barcelona, 24th November 2016 – BSC researchers have received the award for Best Paper of the [MEMSYS 2016 conference](#) with the paper entitled “[Large-memory nodes for energy efficient High-Performance Computing](#)”. Energy consumption is by far the most important contributor to HPC cluster operational costs, and it accounts for a significant share of the total cost of ownership. In the awarded paper, BSC experts prove that a practical energy-saving approach is to scale-in the application on large-memory nodes. Achieving energy savings of up to 52% means that the investment in upgrading the hardware would be recovered in a system lifetime of less than five years. The study is an important finding of the BSC Memory Systems team using Samsung DRAM memory with technical support of Samsung Electronics Co., Ltd..

Traditionally, the HPC community focuses mainly on scale-out, referring to it simply as scaling. This refers to executing a problem on a fixed machine, but using a larger number of compute nodes and application processes. In this publication, BSC experts propose horizontal scaling in the opposite direction by scaling-in on large-memory nodes. Scaling-in refers to executing a fixed problem on a fixed machine but using a reduced number of applications processes and compute nodes. The proposal was evaluated using a set of HPC applications running on BSC’s MareNostrum supercomputer, and the authors obtained average energy savings of 36% from scaling-in on existing nodes, already a significant figure. “Scaling-in is limited by the per-node memory capacity, so we advocate increasing the per-node memory capacity to enable a greater degree of scaling-in. The presented results show that the additional energy savings, of up to 52%, mean that in many cases the investment in upgrading the hardware with more memory capacity would be recovered through electricity bill savings in a typical system lifetime of less than five years”, says Petar Radojkovic, Memory Systems team leader at BSC.




Figure 1: System scaling can be horizontal (scale-in or -out) and vertical (scale-up or -down). Traditionally, HPC community is focused mainly on scale-out, referring to it simply as scaling. BSC study analyzes scale-in on standard nodes, and a combined scale-up and scale-in approach on large-memory nodes.

“This study confirms that improvements in the memory system can have significant impact on the real world, improving power and energy, performance, and/or dollar cost. The MEMSYS program committee liked the study especially because, in contrast to most of the prior research, it can be applied immediately in the production systems, and it requires no changes to the system architecture, Operating System, system software or applications”, says Bruce Jacob, the MEMSYS 2016 Program Chair.



Figure 2: Darko holding his group's award for Best Paper during the MEMSYS 2016 conference, from left to right: Darko Zivanovic (PhD student at BSC), Bruce Jacob (MEMSYS 2016 Chair, U. Maryland) and Kazi Asifuzzaman (PhD student at BSC).

Scaling-in does not only reduce the energy consumption of HPC jobs, it also decreases the overall experimentation cost, improves system throughput and reduces the execution time for batches of jobs. “In modern HPC, the cost and energy consumption of the experiments has become highly visible and of prime importance. We therefore hope that the study will motivate the community to consider the trade-offs between horizontal and vertical scaling when provisioning and using HPC clusters. Maybe we could start this journey with some second thoughts about the way that we use the word *scalability*.” concludes Petar Radojkovic.

About the Barcelona Supercomputing Center (BSC)

Barcelona Supercomputing Center (BSC) is the national supercomputing centre in Spain. BSC specialises in high performance computing (HPC) and its mission is two-fold: to provide infrastructure and supercomputing services to European scientists, and to generate knowledge and technology to transfer to business and society. BSC is a Severo Ochoa Center of Excellence and a first level hosting member of the European research infrastructure PRACE (Partnership for Advanced Computing in Europe). BSC also manages the Spanish Supercomputing Network (RES). More information on page www.bsc.es

Press contact

Barcelona Supercomputing Center
Renata Giménez-Binder
Tel: (+34) 93 4015864
communication@bsc.es