

Main Memory in HPC: Do We Need More or Could We Live with Less?

Darko Zivanovic^{*‡}, Petar Radojković^{*}, Eduard Ayguadé^{*‡}

^{*}Barcelona Supercomputing Center (BSC), [‡]Universitat Politècnica de Catalunya

darko.zivanovic@bsc.es, petar.radojkovic@bsc.es, eduard.ayguade@bsc.es

Keywords— High-performance computing, Memory capacity requirements, Production HPC applications, HPL, HPCG, Large-memory nodes, Energy-efficiency, Scaling-in.

EXTENDED ABSTRACT

An important aspect of High-performance Computing (HPC) system design is the choice of main memory capacity. This choice becomes increasingly important now that 3D-stacked memories are entering the market. Compared with conventional DIMMs, 3D memory chiplets provide better performance and energy efficiency but lower memory capacities. Hybrid memory systems, that combine 3D-stacked DRAM with standard DIMMs, should bring the best of two worlds — the bandwidth, latency and energy-efficiency of 3D-stacked DRAM together with the capacity of DIMMs. However, they are still difficult to manage, so 3D-memories will only be employed in HPC if enough applications have sufficiently small memory footprints to fit inside 3D memories exclusively.

This study analyzes the memory capacity requirements of important HPC benchmarks and applications. We find that the High Performance Conjugate Gradients benchmark could be an important success story for 3D-stacked memories in HPC, but High-performance Linpack is likely to be constrained by 3D memory capacity.

The study also emphasizes that the analysis of memory footprints of production HPC applications is complex and that it requires an understanding of application scalability and target category, i.e., whether the users target capability or capacity computing. In HPC, capability computing refers to using large-scale HPC installations to solve a single, highly complex problem in the shortest possible time, while capacity computing refers to optimizing system efficiency to solve as many mid-size or smaller problems as possible at the same time at the lowest possible cost.

The results show that most of the HPC applications under study have per-core memory footprints in the range of hundreds of megabytes, and these applications represent use cases in HPC that require memory capacities that could be provided solely by 3D memories, which is a first step toward their adoption in HPC.

We also detect applications and use cases in capacity computing that still require gigabytes of memory per core, and for these use cases we propose scaling-in, i.e. reducing the number of nodes for the execution. We show that scaling-in leads to significant energy savings and we propose upgrading the memory capacity which enables greater degree of scaling-in. We show that additional energy savings, of up to 52%, mean that in many cases the investment in upgrading the memory system would be recovered in a typical system lifetime of less than five years.

A. ACKNOWLEDGEMENT

This work is published on the International Symposium in Memory Systems (MEMSYS) [1] and in ACM Transactions on Architecture and Code Optimization (TACO) [2].

References

- [1] D. Zivanovic, M. Radulovic, G. Llort, D. Zaragoza, J. Strassburg, P. M. Carpenter, P. Radojković, E. Ayguade, “Large-Memory Nodes for Energy Efficient High-Performance Computing,” in Proceedings of the International Symposium on Memory Systems (MEMSYS), Pages 3-9, October 2016.
- [2] D. Zivanovic, M. Pavlovic, M. Radulovic, H. Shin, J. Son, S. A. McKee, P. M. Carpenter, P. Radojković, E. Ayguade, “Main Memory in HPC: Do We Need More or Could We Live with Less?,” ACM Transactions on Architecture and Code Optimization (TACO), Volume 14 Issue 1, Article No. 3, March 2017.

Author biography



Darko Zivanovic is a PhD candidate at Barcelona Supercomputing Center and Polytechnic University of Catalonia. His research is focused on Memory Systems for High-performance Computing. He obtained his B.Sc. and M.Sc. degrees from the School of Electrical Engineering at the University of Belgrade in 2008 and 2010, respectively. During his Master he joined the Institute Mihajlo Pupin in Belgrade, where he worked from 2009 to 2011 as Embedded System Developer. From 2011 to 2013 he was part of Architecture and Compilers (ARCO) research group in Barcelona, and in 2013 he joined Barcelona Supercomputing Center in a pursue for his PhD degree.