


Memory systems for HPC and AI



Next-generation memory systems for HPC and AI face significant challenges. These systems must deliver increased capacity and performance within a constrained power budget. We design and evaluate memory architectures that meet the requirements of these critical applications.

Summary

In state-of-the-art high-performance computing (HPC) and artificial intelligence (AI) architectures, the design of memory systems plays a crucial role in determining overall system performance. Memory systems have also become major contributors to system power demands, energy consumption, and the operational costs of large-scale computing infrastructures. Furthermore, as DRAM technology scales and main memory capacity increases, the likelihood of DRAM errors has grown, becoming a frequent cause of system failures in the field.

The Barcelona Supercomputing Center (BSC) recognizes the critical need for advanced research on memory systems for HPC and AI. For over a decade, our research has addressed challenges in designing and optimizing the utilization of next-generation memory systems for large-scale HPC and AI. Currently, our main research focuses on exploring the following areas:

- **Memory system benchmarking, simulation and application profiling**

The [Memory stress \(Mess\) framework](#) provides a unified view of the memory system benchmarking, simulation and application profiling. The Mess benchmark provides a holistic and detailed memory system characterization. It is based on hundreds of measurements that are represented as a family of bandwidth-latency curves. The benchmark increases the coverage of all the previous tools and leads to new findings in the behavior of the actual and simulated memory systems. The Mess memory simulator uses the bandwidth-latency concept for the memory performance simulation. The Mess simulator is fast, easy to integrate and it closely matches the actual system performance. Finally, the Mess application profiling positions the application in the bandwidth-latency space of the target memory system. This information can be correlated with other application runtime activities and the source code, leading to a better overall understanding of the application's behavior. The current Mess benchmark release covers all major CPU and GPU ISAs, x86, ARM, Power, RISC-V, and NVIDIA's PTX. We also release as open source the ZSim, gem5 and OpenPiton Metro-MPI integrated with the Mess simulator for DDR4, DDR5, Optane, HBM2, HBM2E and CXL memory expanders. The Mess application profiling is already integrated into a suite of production HPC performance analysis tools.

- **Performance prediction of future memory systems**

Novel memory systems are typically explored by hardware simulators that are slow and often have a simplified or obsolete abstraction of the CPU. BSC Memory team presents [PROFET](#), an analytical model that predicts how an application's performance and energy consumption change when executed on different (future) memory systems. The model is based on the instrumentation of an application execution on actual hardware, so it already considers CPU microarchitectural details. PROFET is evaluated on Intel and Huawei servers with DDR3, DDR4, HBM and Optane memory. We release the PROFET source code and all input data required for memory system and application profiling.

- **Processing in memory for HPC and AI**

By moving part of the computation to the memory devices, processing in memory (PIM) addresses a fundamental issue in the design of modern computing systems, the mismatch between the von Neumann architecture and the requirements of important data-centric applications. A number of industrial prototypes and products are under development or already available in the marketplace, and these devices show the potential for cost-effective and energy-efficient HPC and AI workload acceleration. Our research explores different PIM aspects, from technology, hardware, system software and programming environment, to updating of the algorithm and application. We aim to make PIM a reality in production HPC and AI.

- **FAiND the best hardware for your AI application**

The [FAiNDER.eu](#) is an open-source web platform designed to transform how researchers and developers navigate the rapidly evolving AI landscape. FAiNDER provides centralized, up-to-date information on the key system requirements of all major AI models to facilitate exploration and optimize hardware choices. With FAiNDER, users choose the depth at which they want to explore the AI world. For students and junior researchers, the tool offers a starting point for exploring any AI model. It provides a high-level view of the models' software architecture and their main features, and it includes a list of reliable sources for future exploration. For advanced users, the tool provides detailed information on all phases of the software pipeline, their main arithmetic operations, data types, memory footprint and the target hardware platforms, enabling quick comparisons and deep technical insights.

- **Detection, analysis and prediction of DRAM field errors**

We detect, analyze, and predict DRAM errors in the MareNostrum supercomputer, one of six Tier-0 systems in Europe. Our analysis provides a deeper understanding of DRAM errors in real-world conditions, leading to better utilization and enhanced stability of large-scale HPC systems.

Objectives

- Provide a unified view of the memory system benchmarking, simulation and application profiling.
- Understand HPC and AI applications' memory requirements and design future memory systems that meet them.
- Help researchers and developers to FAiND and build the best hardware for their AI applications.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 1 Jun 2025 - 14:28): <https://www.bsc.es/es/research-development/research-areas/computer-architecture-and-codesign/memory-systems-hpc-and-ai>