

Seamlessly embracing memory heterogeneity in HPC

João Pedro Barreto, Técnico Lisboa

Abstract

The increasing heterogeneity of the memory hierarchies of today's high-end systems is a tremendous opportunity for improving the efficiency, performance and fault tolerance of memory-intensive HPC applications. As we progress towards Exascale systems, this scenario is greatly amplified. However, in order to fully exploit the virtues of the upcoming heterogeneous memory architectures, we need to deeply rethink the way data is laid out, accessed and managed. One key question that needs to be answered is concerned data placement: at which locations of a heterogeneous memory system should an application place each data object?

This is one of the key problems that the EPEEC project addresses. In order to retain programming productivity, it is desirable that data placement is optimized without requiring disruptive changes to applications. This talk will present our main achievements and results on addressing the problem of data placement at the operating system's kernel level. The main goal is to seamlessly optimize the placement of the memory pages that HPC applications allocate and access. We will focus on two main scenarios: page placement for multi-socket NUMA systems; and page placement with hybrid memory systems combining DRAM with the Intel Optane persistent memory.

Bio

João Barreto is an Assistant Professor at the Computer and Information Systems Department at IST, where he received his Ph.D. degree in 2009. He is a senior researcher at INESC-ID since 2001, and his research interests are parallel and distributed systems and operating systems, in particular: concurrent programming, transactional memory, optimistic replication, and distributed data deduplication. He has participated in a number of international projects, including H2020 projects Cloud-TM and TRACE (as Technical Coordinator) and currently leads the INESC-ID team at the EPEEC project. He is author or co-author of over 25 peer-reviewed scientific publications, including a best paper award at the ACM/IFIP/USENIX 13th International Middleware Conference, and he has served as reviewer on a number of top international conferences and journals.