Subpage Migration in Heterogeneous Memory Systems

Shashank Adavally (University of North Texas)

Abstract

With the increasing demands for very large physical address spaces and the advent of memory technologies that can support large memories, there is a need to reduce the sizes of system tables such as TLBs and page tables. One can use very large (huge) pages instead of traditional 4K byte pages. However, huge pages are likely to lead to internal fragmentation and may make page migration strategies that aim to move heavily used pages to faster memories inefficient. If only a small portion of a huge page is heavily accessed, it may be worth migrating only that portion to a faster memory. This paper proposes two page migration techniques (i) subpage migration with Address Reconciliation (that is, updating physical addresses of migrated pages) and (ii) subpage migration with Reverse Migration (whereby no address reconciliation is needed).We observed speedup ranging up to 17% for some SPEC workloads.

Bio

Shashank Adavally received his master's degree and currently pursuing PhD at University of North Texas. His research experience is focused in the areas of Memory Accelerators, Conventional and Heterogeneous Memory Systems, Performance analysis but also interested in the field of Computer Architecture, High Performance Computing with a focus on Simulation and Modelling, Processor Micro-architecture, Memory Systems.