1. Record Nr.   UNINA9910139122103321

   Titolo   2010 IEEE International Symposium on Performance Analysis of Systems and Software

   Pubbl/distr/stampa   [Place of publication not identified], : IEEE, 2010

   ISBN   9781424460243
                          1424460247

   Descrizione fisica   1 online resource (ix, 248 pages) : illustrations

   Disciplina   004.24

   Soggetti   Computer systems - Evaluation

   Lingua di pubblicazione Inglese

   Formato   Materiale a stampa

   Livello bibliografico Monografia

   Note generali   Bibliographic Level Mode of Issuance: Monograph

   Sommario/riassunto   As processor performance continues to outgrow memory capacity and bandwidth, system and application performance has become constrained by the memory subsystem. Promising new technologies like Phase Change Memory (PCM) and Flash have emerged which may add capacity at a cost cheaper than conventional DRAM, but at the cost of added latency and poor endurance. It is likely that systems leveraging these new memory technologies in the memory subsystem would require an innovative memory system architecture to gain the benefit of added capacity while mitigating the costs of latency and potential device wear-out. One such proposed architecture is a hierarchical memory sub-system with a faster but costly memory (e.g., DRAM) acting as a cache for a slower but cheaper memory e.g., solid state memory like NAND flash, NOR flash or PCM. The memory subsystem is now a hybrid of two different memory technologies, exploiting the cost effectiveness and non-volatility of solid state memory devices with the speed of traditional DRAM. In order to study the performance tradeoffs with such hierarchical architectures one needs to first study the effect of having a last level cache, which is much larger than the caches in existing systems. Existing tools and methodologies for cache evaluation fall short. We develop a multi-processor system prototype that runs applications with a coherently-

attached FPGA which can emulate different memory architectures for long periods of time. The output of the system is not a memory trace, but the performance results of the emulated memory system design which may be used at any time to evaluate the design tradeoffs. The large cache will filter out references going to the solid state memory. Thus the miss ratio of the large cache is an important metric. The sensitivity of the miss ratio to configuration parameters like cache size and line size needs to be evaluated to identify the right set of parameters.