
The Role of New Memory Structures in Data Intensive Computing

Mazin Yousif (U. Arizona)
Alexandre Ferreira (U. Pittsburgh)

Abstract

The technical challenges/opportunities and the respective architectural ramifications for using Phase Change Memory (PCM) in Data Intensive Computing (DIC) in general and server in-specific are the focus of this talk. PCM is a solid-state memory (similar to Flash) made of germanium-antimony alloy that can be changed rapidly between crystalline to amorphous, representing binary “1” and “0,” respectively. Switching from one form to another is a thermal-based process (not electrical process as in DRAM). PCM is non-volatile, slower than DRAM, has asymmetric latencies (reads are faster than writes), is low power, is highly dense and unlike DRAMs, PCM has limited number of writes. Such features have profound impact on the architectures of servers and consequently on DIC.

This talk will first provide a background on PCM to be followed by possible server architectures with PCM and means to mitigate PCM’s inefficiencies such as endurance and long asymmetrical Read/Write latencies, including adding DRAM cache and processing capabilities to PCM dies. We will also look at the implication of PCM’s non-volatility to performance and reliability in DIC, as well as possible research to bring these capabilities to reality.
