DESIGN TECHNIQUES TO FACILITATE THE ADOPTION OF EMERGING MEMORY TECHNOLOGIES IN COMPUTING SYSTEMS

By

Qi Wu

A Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

DOCTOR OF PHILOSOPHY

Major Subject: Electrical Engineering

Approved by the
Examining Committee:

_____
Tong Zhang, Thesis Adviser

_____
James J.-Q. Lu, Member

_____
Gary Saulnier, Member

_____
Christopher D. Carothers, Member

Rensselaer Polytechnic Institute
Troy, New York

December 2011
(For Graduation December 2011)

ABSTRACT

The emerging three-dimensional (3D) integration and non-volatile memory based solid-state drive (SSD), which have recently received extensive attentions from both academia and industry, are promising candidates to tackle looming memory gap and storage gap in computing systems. Many critical problems need to be solved before these emerging technologies can be practically employed in computing systems. The purpose of this thesis aims to propose simple yet effective design techniques to address some critical problems and facilitate the adoptions of these technologies in real-life computing systems.

First, we proposed a 3D cached-DRAM design strategy which can reduce the 3D DRAM access time and energy consumption by up to 52.5% and 54.4% for a quad-core system, respectively. Second, we developed a through-DRAM TSV (through-silicon via) allocation strategy, which can well fit to the regular DRAM architecture, for 3D processor-DRAM integrated computing system. We also proposed a novel decoupling capacitor implementation method that can reduce the IR drop of power delivery network in 3D processor-DRAM integrated computing systems by up to 61.3%. Third, we proposed a cross-layer design strategy which can reduce the average SSD response time by 61.1%. The key is to cohesively exploit run-time data access workload variation and temporal locality at the system level and NAND flash memory write latency vs. data retention time trade-off at the device level. Fourth, we proposed a self-healing SSD design strategy that exploits heat-accelerated recovery of NAND flash memory cell wear-out to improve memory program/erase (P/E) cycling endurance and hence overall SSD lifetime. Simulation results show that SSD lifetime can be improved by over five times at reasonable performance and energy consumption overhead. Finally, we developed a time-aware design methodology to overcome the significant resistance drift of multi-level per cell (MLC) phase-change memory (PCM) for MLC PCM based data storage systems. Simulation results show that our proposed time-aware design techniques can improve the data retention limit by few orders of magnitude, and enable up to 97%

and 79% energy savings for PCM-based solid-state disk and PCM-based disk cache.