DRAM

Refresh

- J. Liu, B. Jaiyen, Y. Kim, C. Wilkerson, and O. Mutlu. An experimental study of data retention behavior in modern dram devices: Implications for retention time profiling mechanisms. In *International Symposium on Computer Architecture*, 2013.

- J. Stuecheli, D. Kaseridis, H. Hunter, and L. John. Elastic refresh: Techniques to mitigate refresh penalties in high density memory. In *International Symposium on Microarchitecture*, 2010.

Performance Optimization

- Y. Kim, V. Seshadri, D. Lee, J. Liu, and O. Mutlu. A case for exploiting subarray-level parallelism (SALP) in DRAM. In *International Symposium on Computer Architecture*, 2012.

- D. Lee, Y. Kim, V. Seshadri, J. Liu, L. Subramanian, and O. Mutlu. Tiered-latency dram: A low latency and low cost dram architecture. In *International Symposium on High Performance Computer Architecture*, 2013.

Turn Around Penalty From Write To Read

Generally, read operations are given higher priority than writes. When the memory system is servicing reads, the DIMMs drive the off-chip data bus and data is propagated from the DIMMs to the processor. Since writes are not on the critical path for program execution, they are buffered at the processor's memory controller. When the write buffer is nearly full (reaches a high water mark), writes have to be drained. The data bus is turned around so that the processor is now the data bus driver and data is propagated from the processor to the DIMMs. This bus turnaround delay (*tWTR*) has been of the order of 7.5 ns for multiple DDR generations. Frequent bus turnarounds add turnaround latency and cause bus underutilization which eventually impacts queuing delay. Therefore, to amortize the cost of bus turnaround, a number of writes are drained in a single batch until a low water mark is reached. During this time, reads have no option but to wait at the memory controller; the uni-directional nature of the bus prevents reads from opportunistically reading data out of idle banks. Thus, modern main memory systems offer nearly zero read-write parallelism within a single channel.

- J. Stuecheli, D. Kaseridis, D. Daly, H. C. Hunter, and L. K. John. The virtual write queue: Coordinating dram and last-level cache policies. In *International Symposium on Computer Architecture*, 2010

- N. Chatterjee, N. Muralimanohar, R. Balasubramonian, A. Davis, and N. Jouppi. Staged reads: Mitigating the impact of dram writes on dram reads. In *International Symposium on High Performance Computer Architecture*, 2012.

Power Optimization

- H. Zheng, J. Lin, Z. Zhang, E. Gorbatov, H. David, and Z. Zhu. Mini-rank: Adaptive DRAM architecture for improving memory power efficiency. In *International Symposium on Microarchitecture*, 2008.
- M. Shevgoor, J.-S. Kim, N. Chatterjee, R. Balasubramonian, A. Davis, and A. N. Udipi. Quantifying the relationship between the power delivery network and architectural policies in a 3d-stacked memory device. In *International Symposium on Microarchitecture*, 2013.

Phase Change Memories

Performance

1. **Improving Read Performance of Phase Change Memories via Write Cancellation and Write Pausing.**
   Moinuddin K. Qureshi, Michele Franceschini and Luis Lastras

2. **PreSET: Improving Read-Write Performance of Phase Change Memories by Exploiting Asymmetry in Write Times** Moinuddin Qureshi (GT) and Michele Franceschini, Ashish Jagmohan, and Luis Lastras (IBM)

Endurance

1. **Enhancing Lifetime and Security of Phase Change Memories via Start-Gap Wear Leveling.**
   Moinuddin K. Qureshi, John Karidis, Michele Franceschini, Viji Srinivasan, Luis Lastras and Bulent Abali
   Appears in the International Symposium on Microarchitecture (MICRO) 2009.

Security

1. **i-NVMM: A secure non-volatile main memory system with incremental encryption**
   Chhabra, S.; Solihin, Y.

Hard Errors

1. Stuart Schechter, Gabriel Loh, Karin Strauss, and Doug Burger, **Use ECP, not ECC, for Hard Failures in Resistive Memories**, in *Proceedings of the 37th International Symposium on Computer Architecture*, June 2010

2. **Pay-As-You-Go: Low-Overhead Hard-Error Correction for Phase Change Memories**
   Moinuddin Qureshi (*IBM*)

3. **Efficient Scrub Mechanisms for Error-Prone Emerging Memories** , M. Awasthi, M. Shevgoor, K. Sudan, B. Rajendran, R. Balasubramonian, V. Srinivasan, *18th International*

Symposium on High-Performance Computer Architecture (HPCA-18) , New Orleans, February 2012

MLC

1. **Morphable Memory System: A Robust Architecture for Exploiting Multi-Level Phase Change Memories.**
   Moinuddin K. Qureshi, Michele Franceschini, Luis Last