Abstract of the Thesis

The unprecedented development in the processing speed of the Chip Multi-Processor (CMP) and the rise of modern data-intensive applications impose high pressure on the memory subsystem. It significantly increases the main memory footprint and necessitates designing of energy-efficient and high capacity main memory. Unfortunately, the traditional memory systems, built predominantly using DRAM are not scalable to the low nanometer regime. At this need of the hour, the Emerging Non-Volatile Memories (NVMs) like PCM, STT-RAM, ReRAM offer fascinating features like high density and low leakage power that are useful for building high capacity and energy-efficient memory systems. However, NVMs have asymmetric read/write operations, where writes are costly in terms of latency and energy. Also, frequent write operations to the NVM cells tend to wear out the memory cells, leading to a shortened memory lifetime. Furthermore, NVMs retain data even after the system is powered down. Hence, an attacker having physical access to the NVM DIMM can easily stream out the sensitive data stored in the NVM. Researchers have proposed encryption techniques to protect the sensitive NVM content. However, encryption algorithms lead to enormous bit-flips when the encrypted data is written in the NVM arrays. Hence, the lifetime issue of the NVM devices is further complicated by encryption-induced bit-flip spikes.

The contributions to this thesis revolve around designing policies to reduce write operations in the NVMs. In addition to extending the lifespan of standard NVMs, the contributions also cover reducing bit-flips caused by encryption and providing a strong security guarantee against data confidentiality-based attacks on NVMs. In particular, we have proposed policy to reduce write-back traffic of the evicted blocks from the Last Level Cache (LLC) to the NVM. While it reduces writes at a cache block level, the other contributions of the thesis focus on building efficient compression and encoding techniques to reduce bit-flips and proposing wear leveling algorithms to even out the bit-flips pressure across the memory space. In the last two contributions, we have designed a partial encryption-based encoding policy and propose techniques utilizing various compression and encoding techniques to reduce encryption-induced bit-flips while ensuring confidentiality in NVMs. The proposed techniques show significant improvement in lifetime, energy-efficiency and performance compared to the state-of-the-art techniques.