Using MRAM in High-Speed Enterprise Storage Caches

Krishna Thangaraj, Ranjana Godse, Adam McPadden & Gary A. Tressler

IBM

Abstract :

Using MRAM in High-Speed Enterprise Storage Caches

With the advent of persistent memory (PM), the performance gap between memory and storage is narrowing. Today's MRAM offers an obvious way to implement PM. MRAM has best-in-class write endurance and data retention, DRAM-like low latency access with low bit error rates, and extremely high performance. It is thus well-suited to holding critical cached data in enterprise applications. For example, a recent application uses MRAM as a write cache, read modify write buffer, and state dump cache in the event of a power loss. The results show greatly improved performance, higher availability, and lower complexity and simpler operation due to the elimination of power backup circuits.

Authors
Krishna Thangaraj, Ranjana Godse, Adam McPadden & Gary A. Tressler – IBM Systems. Poughkeepsie, NY.

Agenda :

- MRAM Overview
- IBM System implementations
- Qualification and Characterization of STT-MRAM for IBM Systems

MRAM overview :

- **Architecture:**
  - Perpendicular Magnetic Tunnel Junction (pMTJ)
  - Magnetic orientation of free layer decides the parallel or anti parallel state wrt fixed layer
- **256Mb DDR3 STT-MRAM** introduced through **IBM Flash Core Module (FCM)**
- Fast access time, persistence without batteries and super caps and low leakage make Spin Torque MRAM (STT-MRAM) attractive
- MRAM does not need refresh unlike DRAM
- Some memory vendors are also focused on eMRAM as a potential cache alternative to SRAM

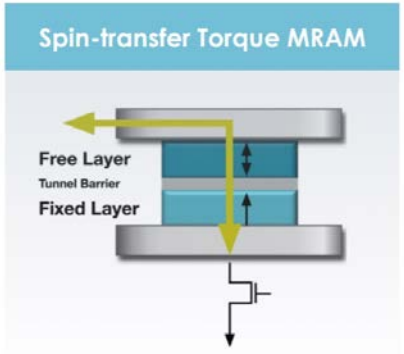

Cell structure for Perpendicular Magnetic Tunnel junction.
Ref: K. Conley, “The MRAM Revolution” MRAM Developer day, August 2018

IBM FlashCore Module
Ref: B. Yardley, “Spin Torque MRAM Is Ready for Applications Today”, MRAM Developer day, August 2018

Performance and $/GB Comparison

Features	DRAM	STT-MRAM	LL-NAND	PCM	ReRAM (Tgt specs)	TLC Flash	QLC Flash
Density	8/16 Gb	256 Mb	64/128 Gb	128 Gb	PCM Comparable	512 Gb	1 Tb
Latency	Low ns	Low ns	Low µs	Mid ns	Mid ns (tgt)	High µs	High µs
Cost							
Availability	✓	✓	✓	✓	R&D	✓	✓

Non-Volatile

Persistence comes with a price

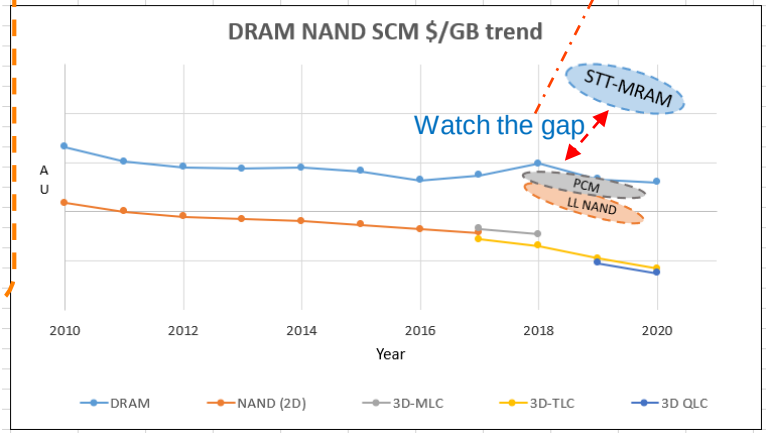

Storage System Implementations:

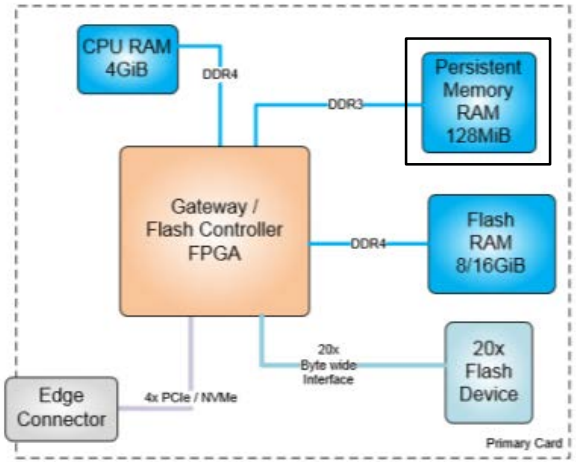

MRAM Use Case in Flash Core Module (1.0):

- Write Data Buffer/Cache
- Read Modify Write Buffer
- Flash Status Back Up on Power Loss
- Journal Buffer
- Flash Firmware Tables Storage
- NVMe Persistent Data (includes state dumps)

IBM's Flash Core Module (FCM) with STT-MRAM

- With the help of 256 Mb STT-MRAM low latency data persistence is achieved in IBM's Flash Core Module
- Significant real estate saving achieved through elimination of super-caps
- IBM to implement next generation of STT-MRAM in the next FCM generation (FCM 2)

Server System Implementation:

On Display during FMS @ BittWare and OpenCAPI booths

3TB Hybrid Memory Subsystem (HMS):

- Near term alternative to other Storage Class Memory based solutions
- High performance with today's OpenCAPI 25G interface
- Achieve competitive load/store performance at a significantly lower cost/GB, by combining DRAM (for caching, prefetch) and Low Latency LL NAND (for capacity)
- Target application requiring high memory capacity at lower cost than DRAM
- Optimized for sequential rather than random accesses (due to media latency)
- Innovative collaboration between BittWare – a Molex company, and IBM

MRAM Use Case in Hybrid Memory Subsystem (HMS) for Persistent Storage of:

- High speed logging of transactions
- Emergency power down data save
- Power up data restore
- Media management tables

Qualification of 256 Mb STT-MRAM for FCM :

- **IBM FlashSystem – Industry leader to use Persistent Memory in the form of STT-MRAM**
  - Qualification focused on Si Reliability, MTJ Reliability
  - 1st Level package Quality and Reliability
  - E2E Quality Review – Fab process controls, wafer parametric and functional test, Defect management
  - IBM Boxline Manufacturing test criteria, FA flow
  - Next steps: Qualification of higher density STT-MRAM for FCM 2

Characterization overview - 256 Mb STT-MRAM :

- Test platform and code environment packages shared to IBM characterization team by the MRAM vendor
- FPGA controller - DDR3 (667 MHz) SODIMM topology
- SODIMM populated with dies from **multiple wafer lots/corner samples**
- Testing focused on endurance - where the memory locations were regressively tested for RBER.
    - Data plots on RBER vs PE cycles (Write/Read) were constructed & analyzed.
- Proprietary FPGA codes/Python scripts and algorithms enabled access to multiple ECC levels

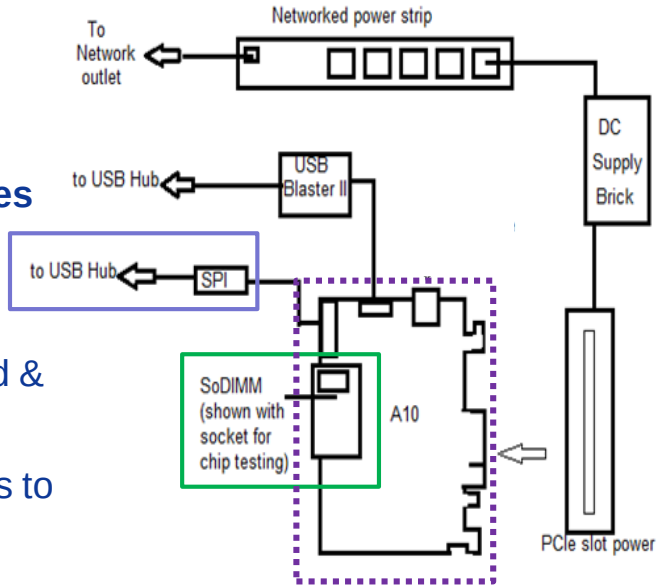

Block diagram of the characterization board

STT- MRAM Characterization– BER vs Cycling Endurance

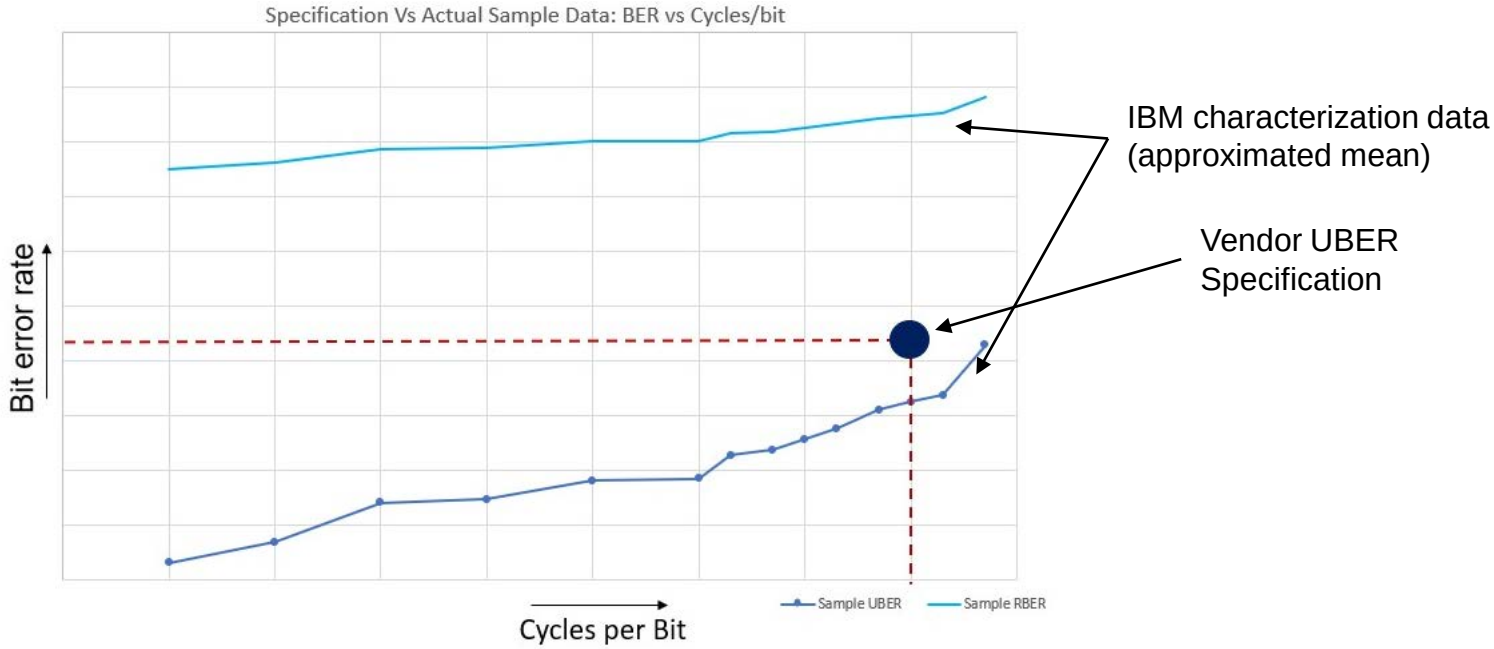

- IBM Characterization data is an average of samples tested
- The STT-MRAM BER information gathered supports vendor's specification

Summary

- IBM is the industry leader to implement persistent memory in the form of STT-MRAM as a cache application
- 256 Mb DDR3 STT-MRAM qualified and characterized for IBM Storage products
- IBM is on the path to enable next generation STT-MRAM
- eMRAM, stacked MRAM industry solutions are under investigation