IOPS and QoS Analysis of DRAM-based and MRAM-based NVRAM Cards

Lorenzo Zuolo‡,**
Cristian Zambelli†, Terry Hulett*, Ben Cooke*, Rino Micheloni‡**, and Piero Olivo†

‡Microsemi Corporation – Via Torri Bianche, 1 – 20871 Vimercate (Italy)
†Dipartimento di Ingegneria – Università degli Studi di Ferrara – Via Saragat, 1 – 44122 Ferrara (Italy)
*Everspin Technologies, Chandler, Arizona 85224 (USA)

**Work done at Università degli Studi di Ferrara

The "Holy Grail of Storage": Persistent Memories

…"To be fast like memories and persistent/dense like storage"…

If we google "persistent memory," we find something about: 902,000 results.

Possibilities opened up by persistent memories:

- From "Three-tiers" to "Two-tiers" memory hierarchy
  - Tier 1: Processing/main memory
  - ~~Tier 2: Non-volatile memory~~
  - Tier 3: Long term archival memory
- "In-non-volatile memory processing"
  - eXec In Place (NOR Flash like)
- **… Infinite possibilities!**

Memories and Storage

Memories "Family"
- Fast like CPUs
- Byte addressable
- SRAMs
- DRAMs

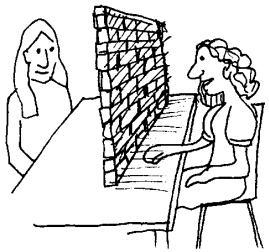

Storage "Family"
- Persistent/dense
- Block addressable
- Magnetic (HDDs)
- **NAND Flash (SSDs)**

…No communication up till today…

Memories + Storage…

Memories "Family"
- Fast like CPUs
- Byte addressable
- SRAMs
- DRAMs

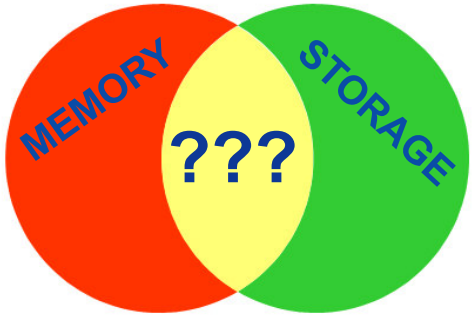

Storage "Family"
- Persistent/dense
- Block addressable
- Magnetic (HDDs)
- **NAND Flash (SSDs)**

Memory (DRAM) + Storage (NAND) == New storage tier

…Persistent DRAMs…

Persistent DRAMs combine the **performance of DRAMs** with the **data persistency of storage**

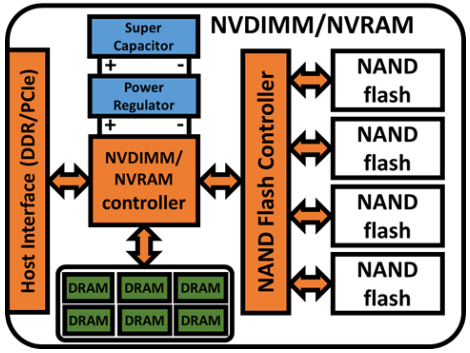

Persistent DRAMs Today

NVDIMMs

- Can be connected on the DDR bus

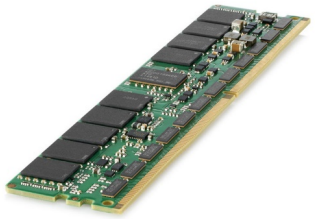

NVRAM cards

- Can be connected on the PCIe bus
- Can be seen either as storage or as memory

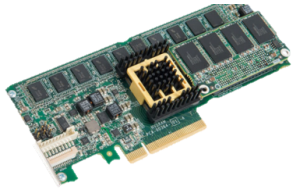

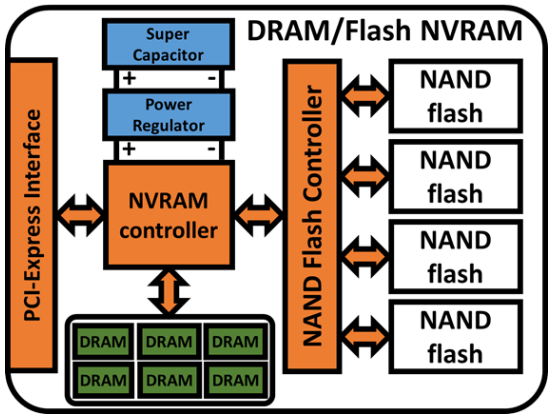

- The host sees the card either **as memory** (PCIe Base Address Registers – BARs) **or storage** (NVMe)
- If the power supply goes down, data are copied from the DRAM to the NAND Flash **(super capacitor)**

NVRAM cards are gaining a lot of traction because of:

- Peer-to-peer communication over fabric
- CAPI communication between an accelerator and the NVRAM card

2016:
The Year of the Storage Revolution

- As a reminder… the "holy grail of storage" is:

  "To be fast like memories and persistent/dense like storage"

- **At the 2016 Flash Memory Summit, Everspin presented a new class of STT-MRAMs (Spin-Torque Transfer Magnetic-RAMs)**

- 256 Mbit
- Fast like a DRAM (DDR3 compliant)
- Non-volatile
- Reliable (10^9 cycles)

PERSISTENT MEMORIES ARE IN MASS PRODUCTION TODAY!

"All-MRAM" NVRAM Card: Theoretical Architecture

- What can I do with MRAM chips? → Redesign the NVRAM card architecture

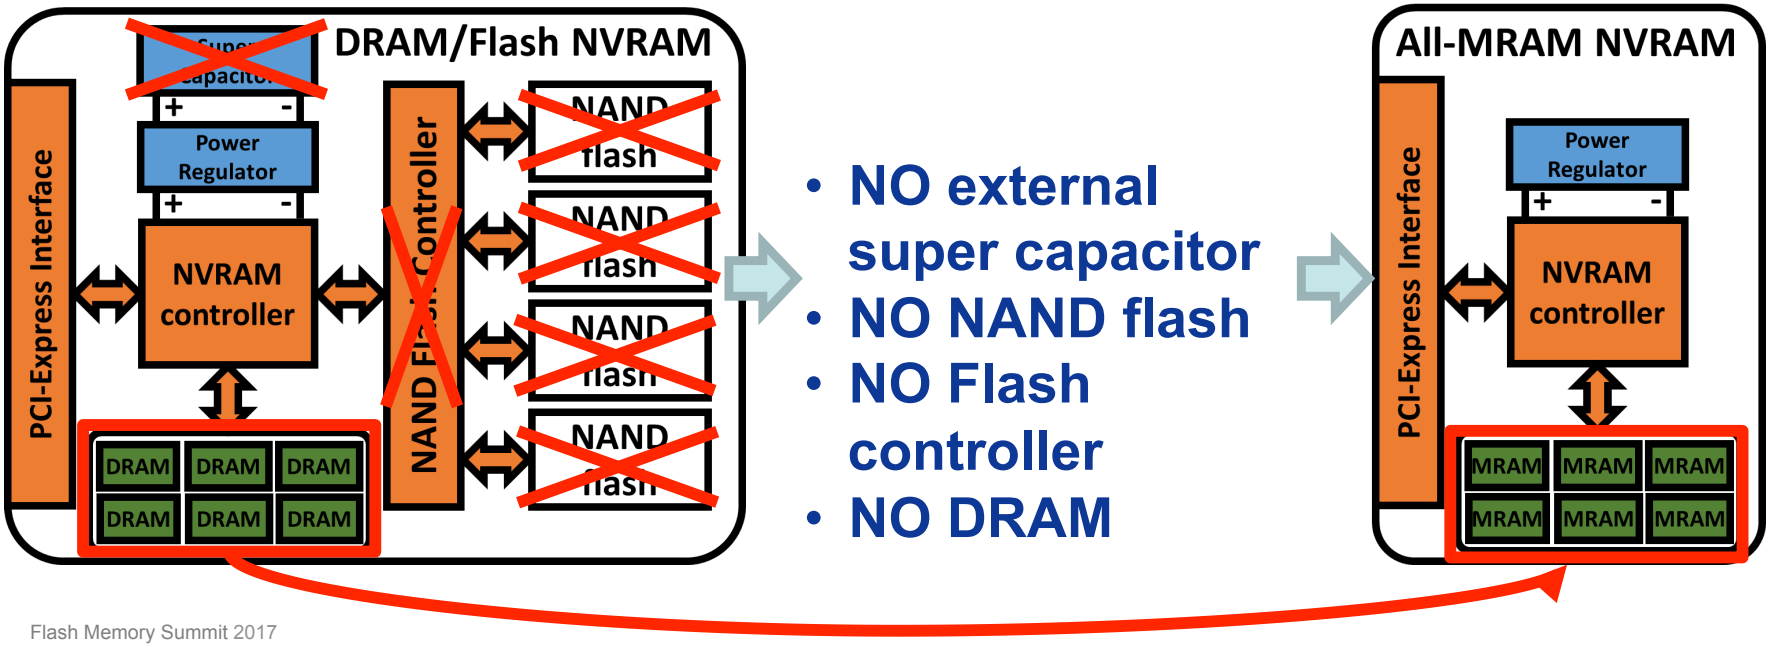

"All-MRAM" NVRAM Card: Realization

- **STT-MRAMs are physically different from DRAMs**
- First we need to **understand if we can replace DRAMs with MRAMs…**
- …without any HW modifications!
- **HOW? → Microsemi's Flashtec™ controller validation board**

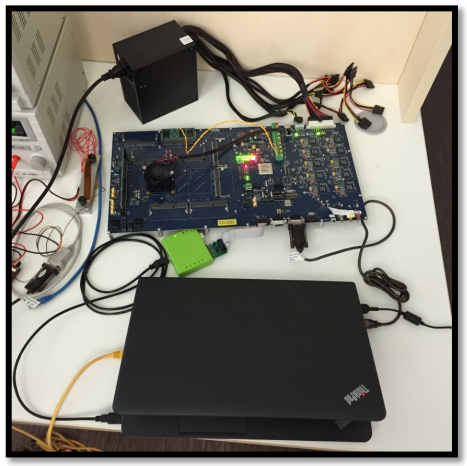

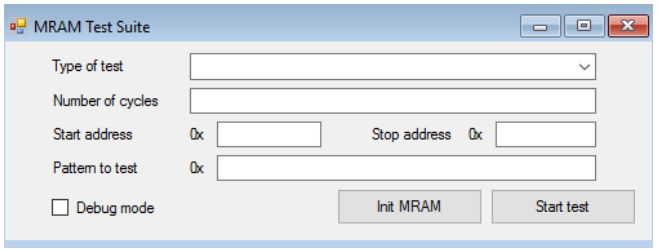

Firmware partially rewritten to take into account MRAM timings

…YES WE CAN DO IT!...

Tests performed on the MRAM:

1. Verify write and read operations:
   - Written BADC0FFE at address 0x80
   - Read address 0x80: result was BADC0FFE
2. Power cycle to check non-volatility:
   - Written BADC0FFE at address 0x80
   - Power cycled the board
   - Read address 0x80: result was BADC0FFE

"All-MRAM" NVRAM Card: Performance

- The Flashtec™ validation board was not designed for performance assessment
- **Hot questions are:**
  1. How an "All-MRAM" NVRAM card behaves when connected to a host system?
  2. Does the different storage paradigm exploited by MRAMs impact the performance/latency?
- **In these cases… Simulation is the way!**

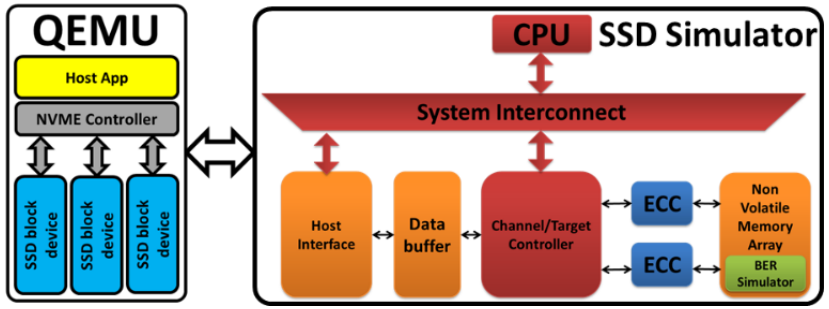

SSDExplorer

- Complete SSD/NVRAM card simulator
- **DRAM simulation with datasheet timings**
- Host emulation by means of Qemu
- Overall performance/latency figures

SSDExplorer Calibration

Before simulating the “All-MRAM” card, it is mandatory to:

- **Tune the simulator to track the behavior of a real DRAM/FLASH card**

NVRAM card parameter	Configuration
Codename	Flashtec™ (Microsemi)
Host Interface	PCI-Express Gen3 x8
Host protocol	NVMe 1.1
DRAM size	1 GB
DRAM controller	Single channel
Intel S2600GZ server	Dual Xeon E5-2680 v2 128 GBytes DRAM

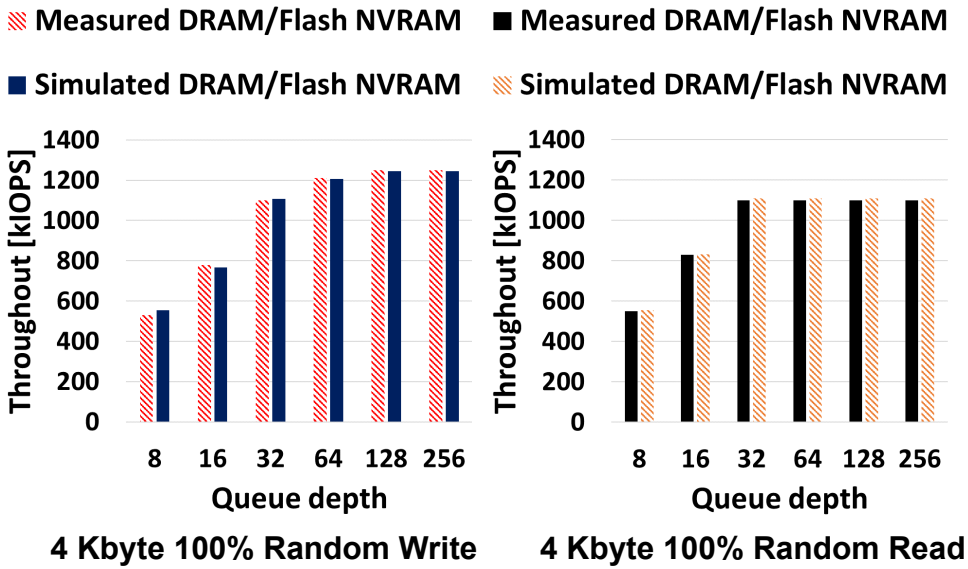

Simulation error is less than 1%

"All-MRAM" Card vs DRAM/Flash Card: Performance with 4KB 100% Random Write

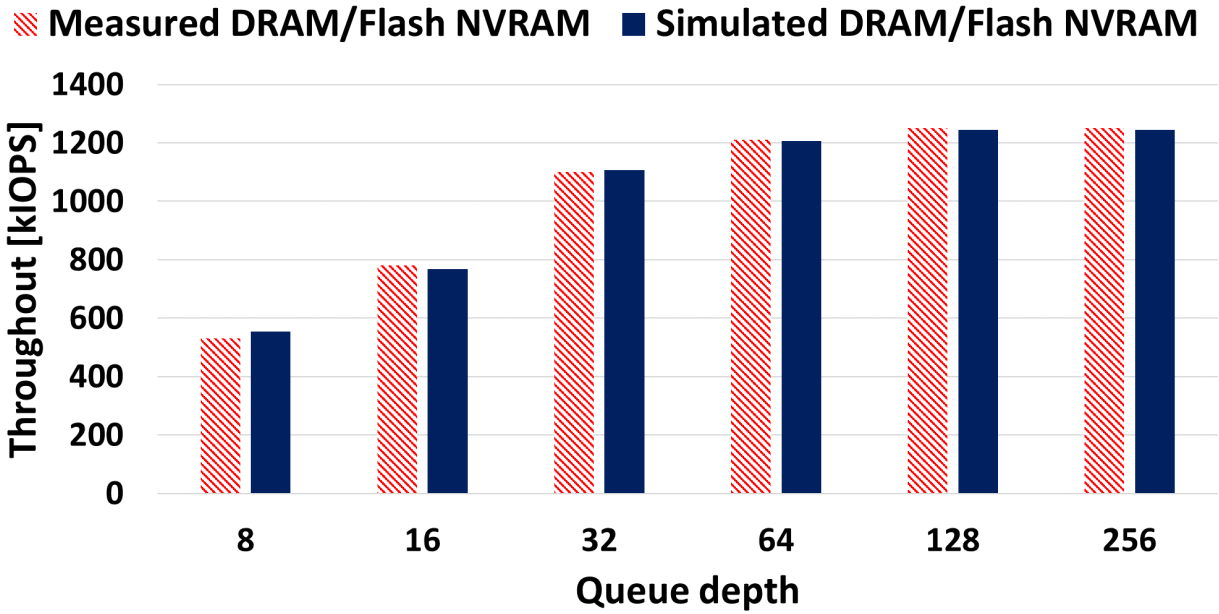

During write the "All-MRAM" card shows the same performance of a standard DRAM/Flash NVRAM card

"All-MRAM" Card vs DRAM/Flash Card: Performance with 4KB 100% Random Read

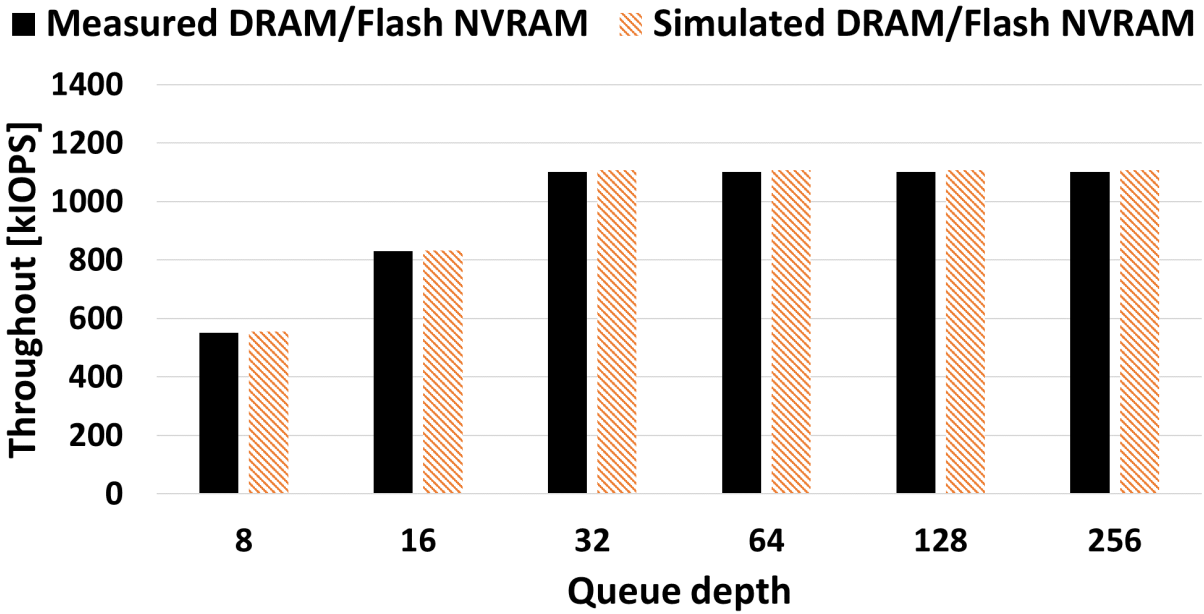

During read the "All-MRAM" card shows the same performance of a standard DRAM/Flash NVRAM card

“All-MRAM” Card vs DRAM/Flash Card: Latency with 4KB 100% Random Write

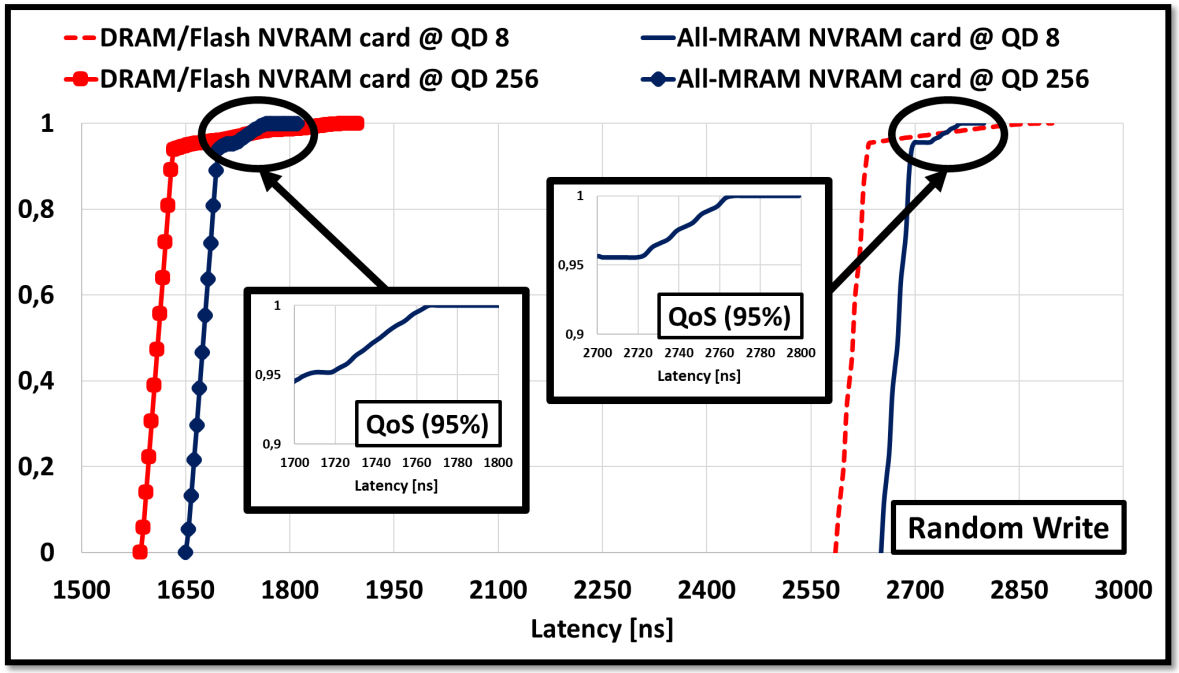

During write the “All-MRAM” card shows almost the same latency as a standard DRAM/Flash NVRAM card

“All-MRAM” Card vs DRAM/Flash Card: Latency with 4KB 100% Random Read

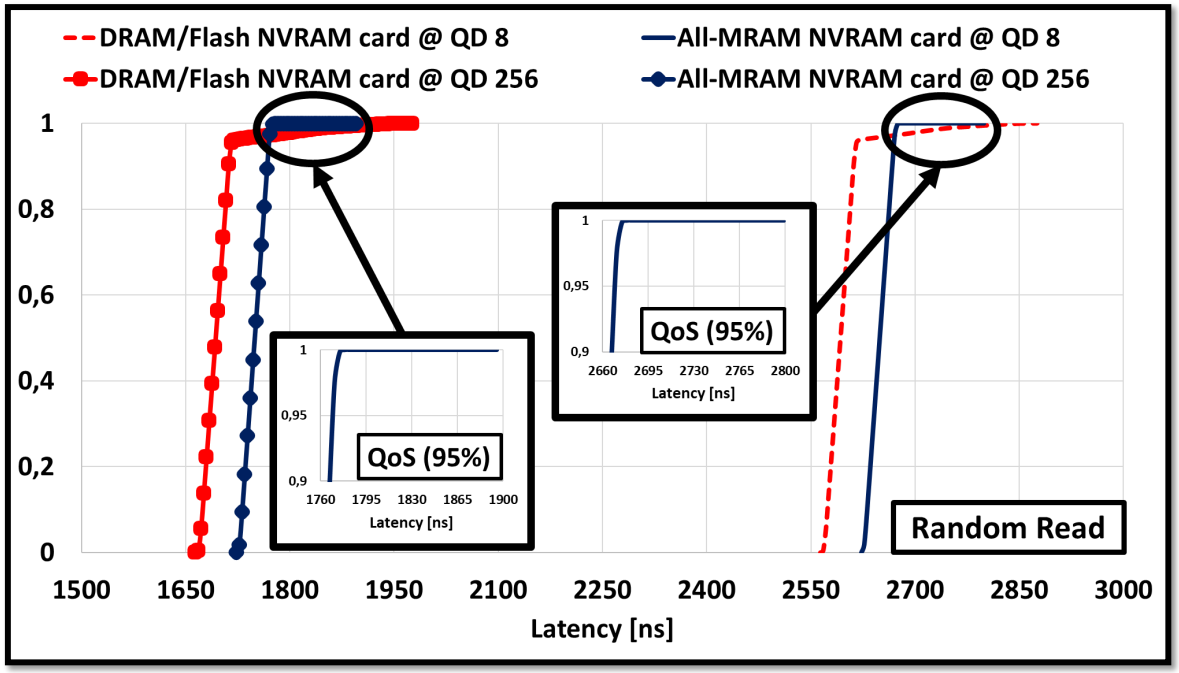

During read the “All-MRAM” card shows almost the same latency as a standard DRAM/Flash NVRAM card

Conclusions

- **<u>Final take away:</u>** with respect to the DRAM/Flash legacy NVRAM card **the "all-MRAM" NVRAM solution does not introduce any significant performance degradation**
- We have shown that **STT-MRAM is a viable candidate for replacing DRAM** inside NVRAM cards.
  - And with the upcoming **1 Gbits** version this architecture will become even more appealing

…But power consumption and costs have to be evaluated…


Flash Memory Summit


Microsemi


EVERSPIN
TECHNOLOGIES
The MRAM Company™


UNIVERSITÀ
DEGLI STUDI
DI FERRARA
- EX LABORE FRUCTUS -

Thank you

Q&A