A Non-Volatile Embedded System

07. Januar 2020

Henriette Hofmeier

Ausgewählte Kapitel der Systemsoftware (AKSS '19/20)

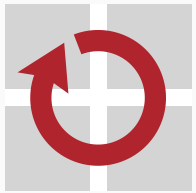

Chair in Distributed Systems
and Operating Systems

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG

FACULTY OF ENGINEERING

Motivation

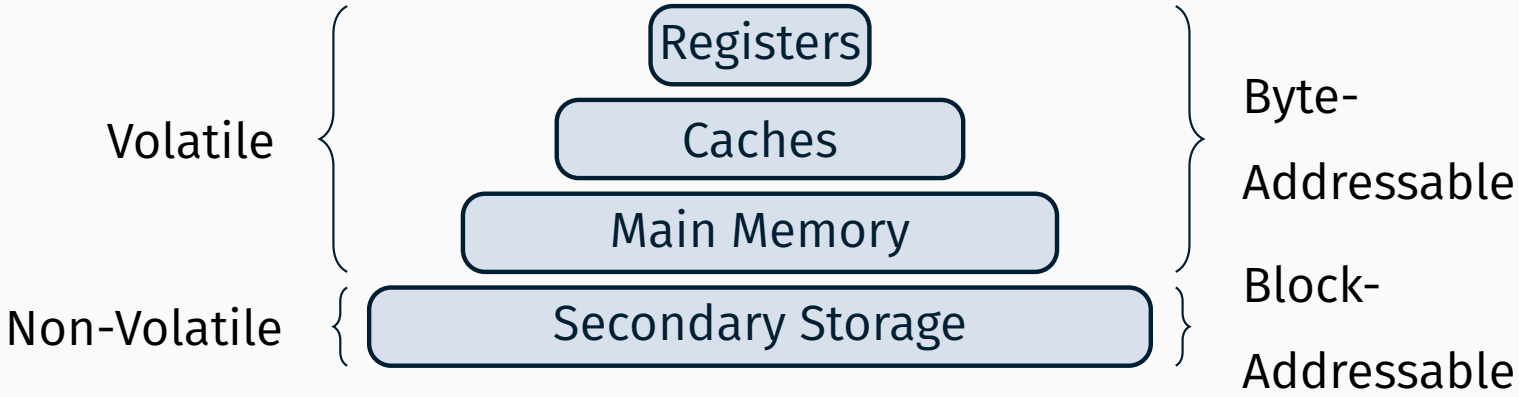
Registers
Caches
Main Memory
Secondary Storage
Volatile
Non-Volatile
Byte-
Addressable
Block-
Addressable

Motivation

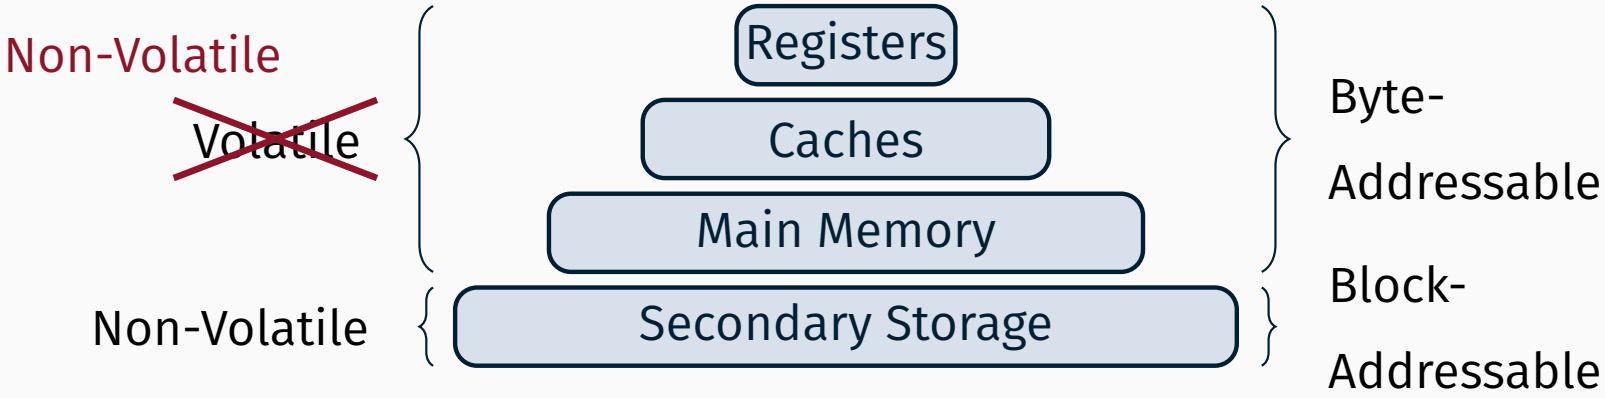
Non-Volatile
Volatile
Registers
Caches
Main Memory
Non-Volatile
Secondary Storage
Byte-
Addressable
Block-
Addressable

Motivation

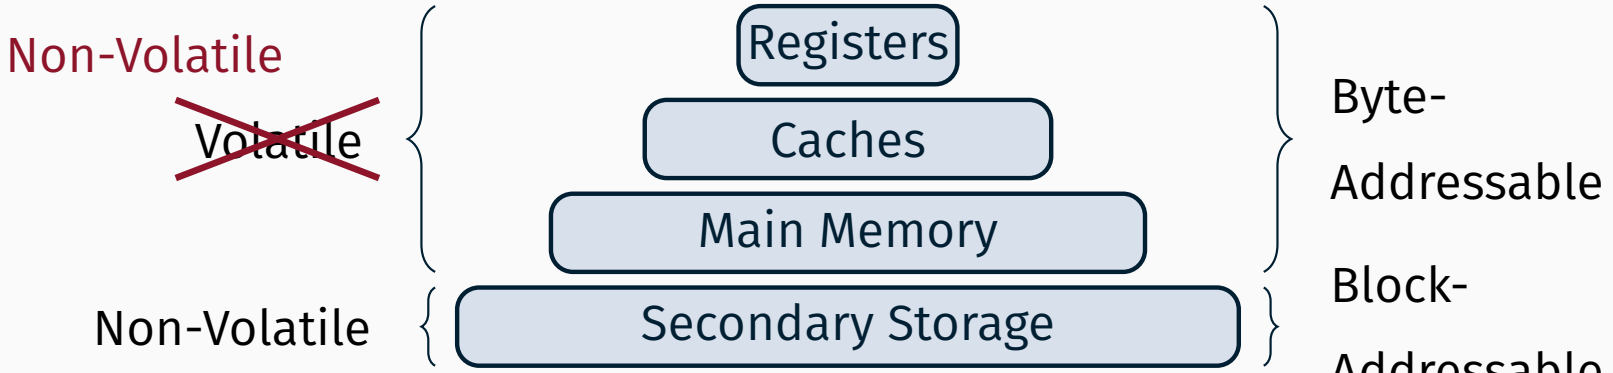
Non-Volatile
Volatile
Registers
Caches
Main Memory
Non-Volatile
Secondary Storage
Byte-
Addressable
Block-
Addressable

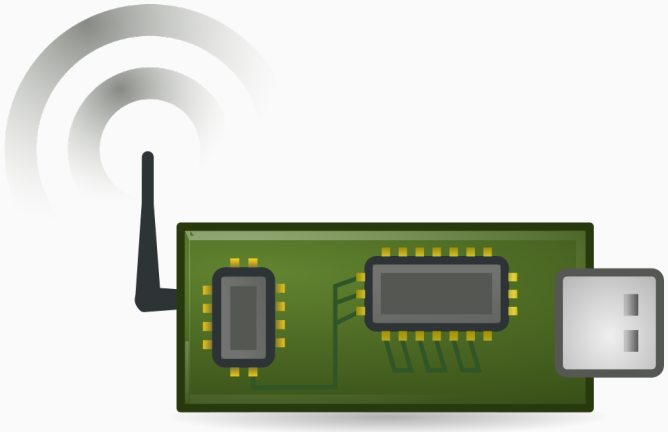

Motivation

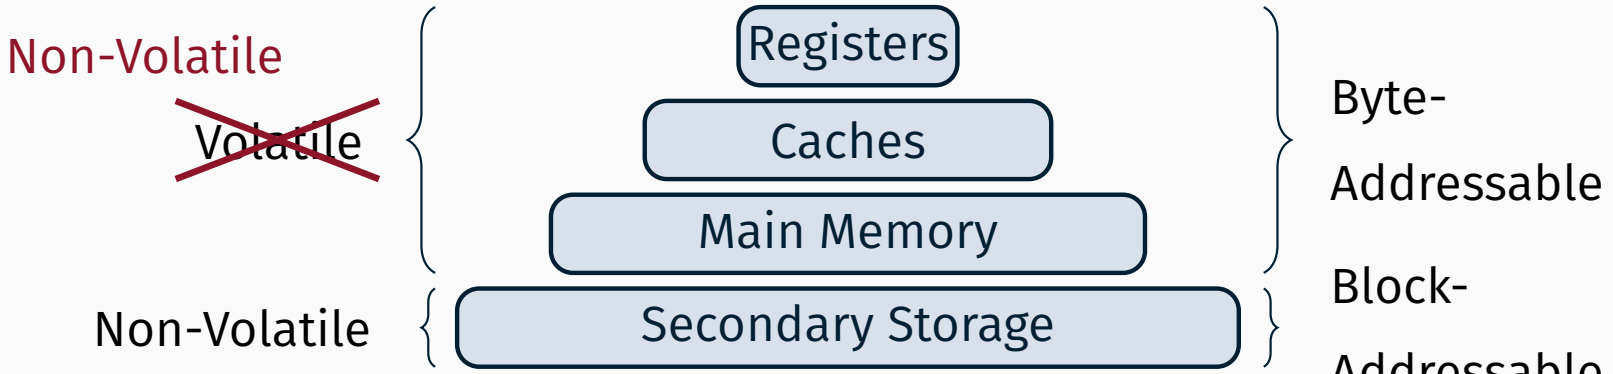
Non-Volatile
Volatile
Registers
Caches
Main Memory
Non-Volatile
Secondary Storage
Byte-
Addressable
Block-
Addressable

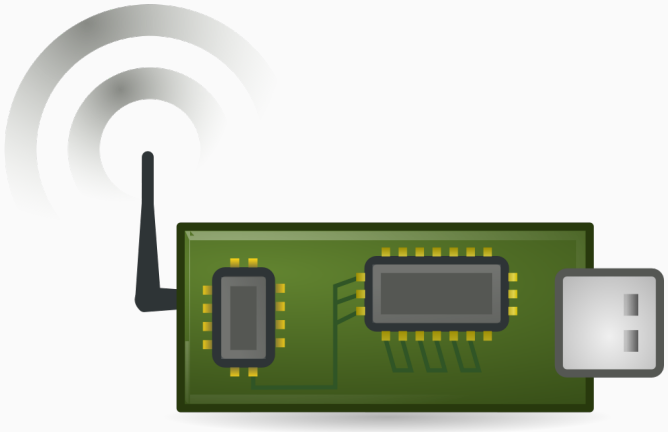

Overview

Motivation

Background

Non-Volatile Components

- Non-Volatile Processor
- Non-Volatile Caches
- Non-Volatile Main Memory

Non-Volatile Embedded System

Phase-Change RAM (PCRAM) [9, 11, 16]

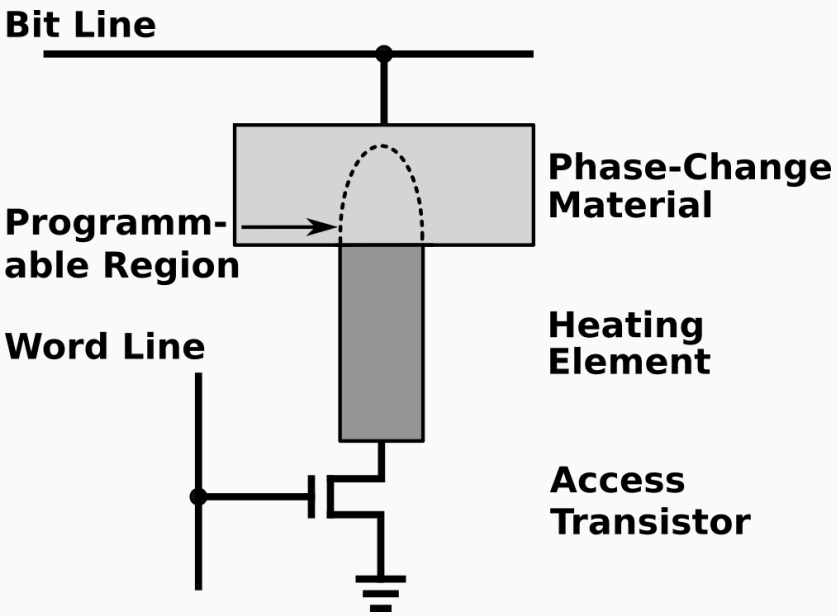

Advantages (PCRAM/DRAM) [8]

Disadvantages (PCRAM/DRAM) [8]

Phase-Change RAM (PCRAM) [9, 11, 16]

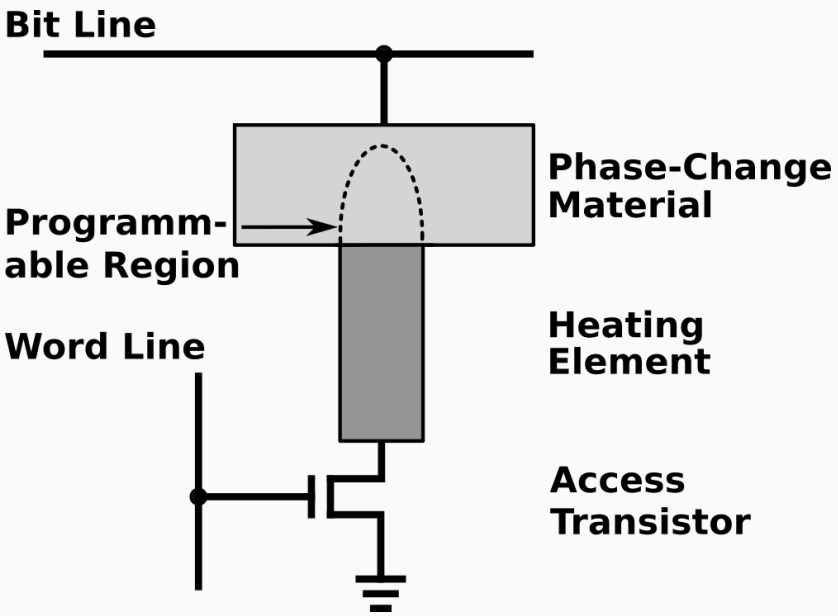

- ’o’: amorphous state (high resistance)

Advantages (PCRAM/DRAM) [8]

Disadvantages (PCRAM/DRAM) [8]

Phase-Change RAM (PCRAM) [9, 11, 16]

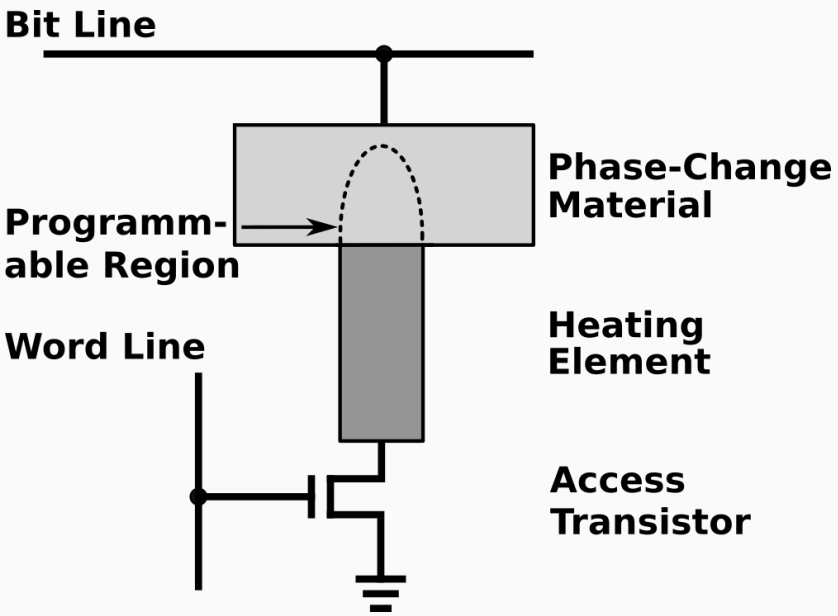

- ’0’: amorphous state (high resistance)
- ’1’: crystalline state (low resistance)

Advantages (PCRAM/DRAM) [8]

Disadvantages (PCRAM/DRAM) [8]

Phase-Change RAM (PCRAM) [9, 11, 16]

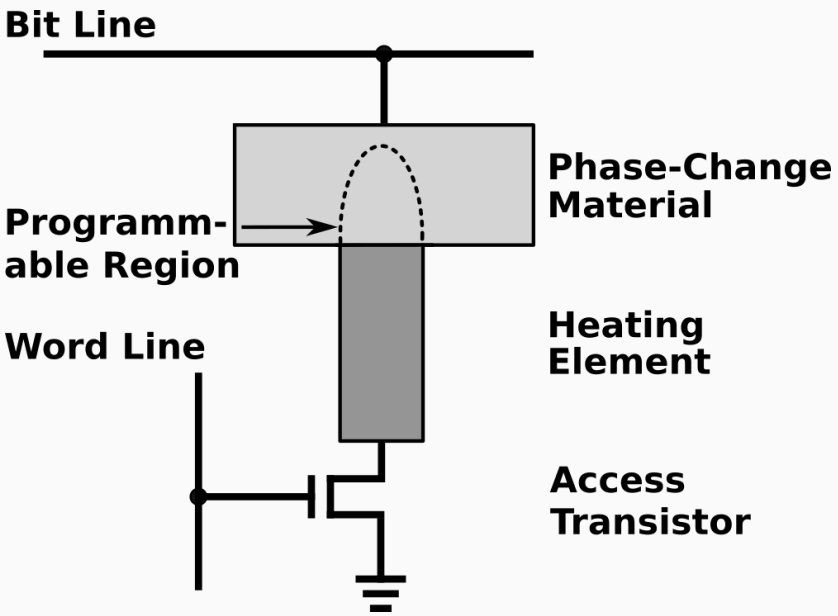

- ’0’: amorphous state (high resistance)
- ’1’: crystalline state (low resistance)
- Write: heat phase-change material

Advantages (PCRAM/DRAM) [8]

Disadvantages (PCRAM/DRAM) [8]

Phase-Change RAM (PCRAM) [9, 11, 16]

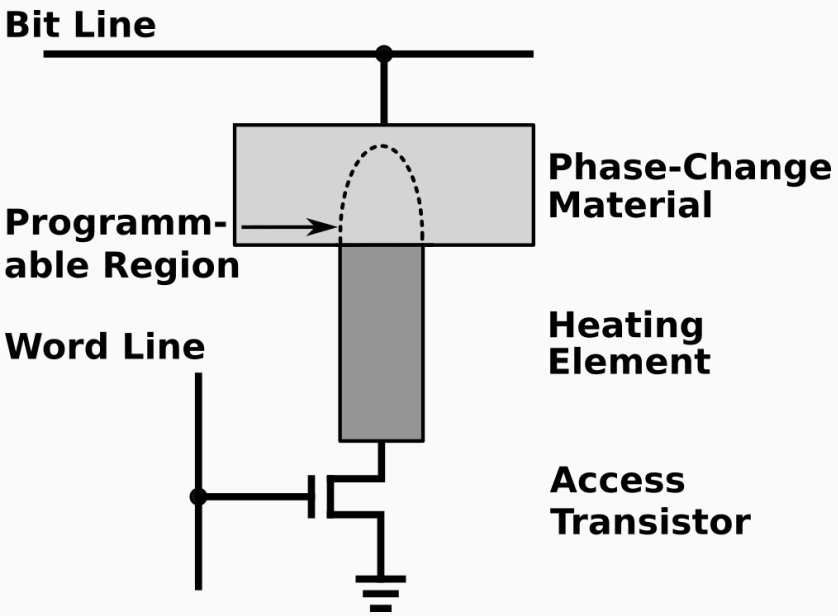

- ’0’: amorphous state (high resistance)
- ’1’: crystalline state (low resistance)
- Write: heat phase-change material
- Read: measure current through cell

Advantages (PCRAM/DRAM) [8]

Disadvantages (PCRAM/DRAM) [8]

Phase-Change RAM (PCRAM) [9, 11, 16]

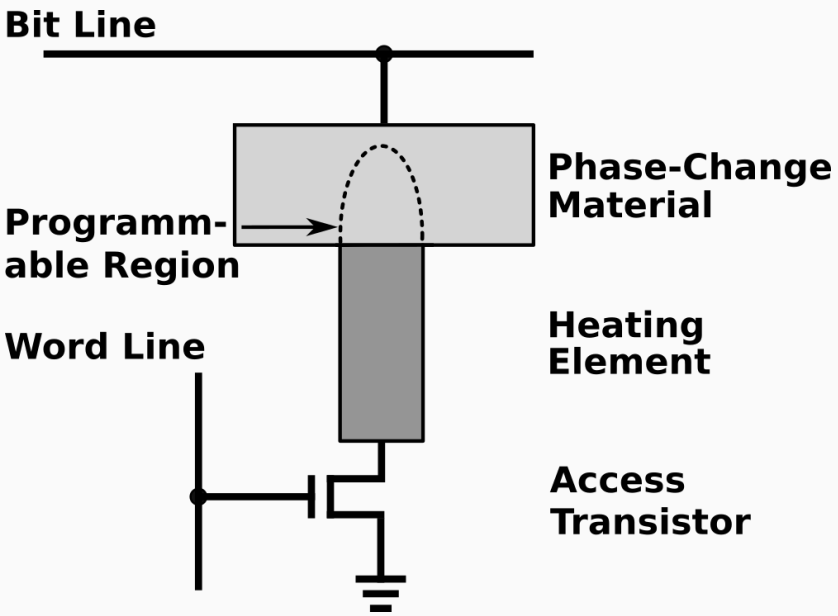

- ’0’: amorphous state (high resistance)
- ’1’: crystalline state (low resistance)
- Write: heat phase-change material
- Read: measure current through cell

Advantages (PCRAM/DRAM) [8]

- ✓ Read latency (50 ns/50 ns)
- ✓ No refresh operations
- ✓ Cell size ($4\,F^2$/$6\,F^2$)

Disadvantages (PCRAM/DRAM) [8]

Phase-Change RAM (PCRAM) [9, 11, 16]

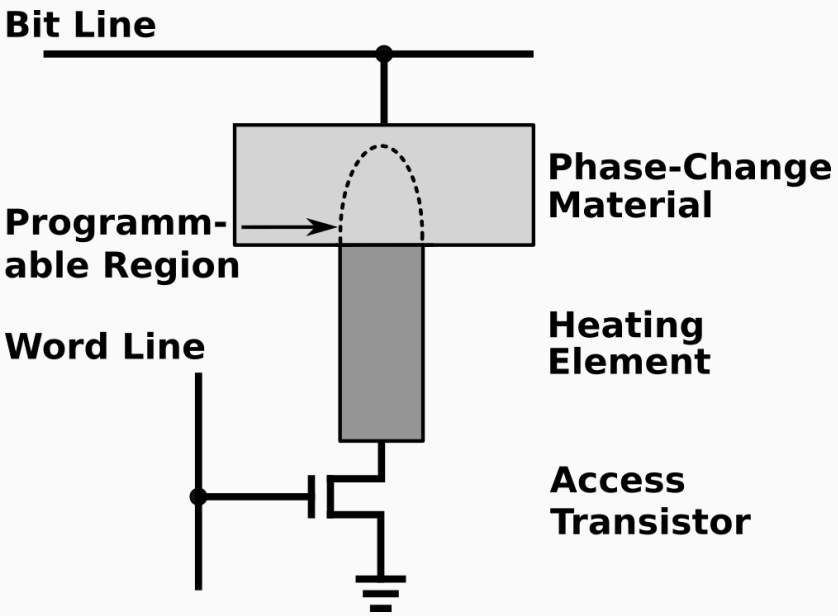

- ’0’: amorphous state (high resistance)
- ’1’: crystalline state (low resistance)
- Write: heat phase-change material
- Read: measure current through cell

Advantages (PCRAM/DRAM) [8]

- ✓ Read latency (50 ns/50 ns)
- ✓ No refresh operations
- ✓ Cell size ($4\,F^2$/$6\,F^2$)

Disadvantages (PCRAM/DRAM) [8]

- ↯ Write latency (500 ns/50 ns)
- ↯ Write energy
- ↯ Write endurance (10^8 cy/10^{15} cy)

Spin-Transfer Torque RAM (STT-RAM) [11, 9, 16]

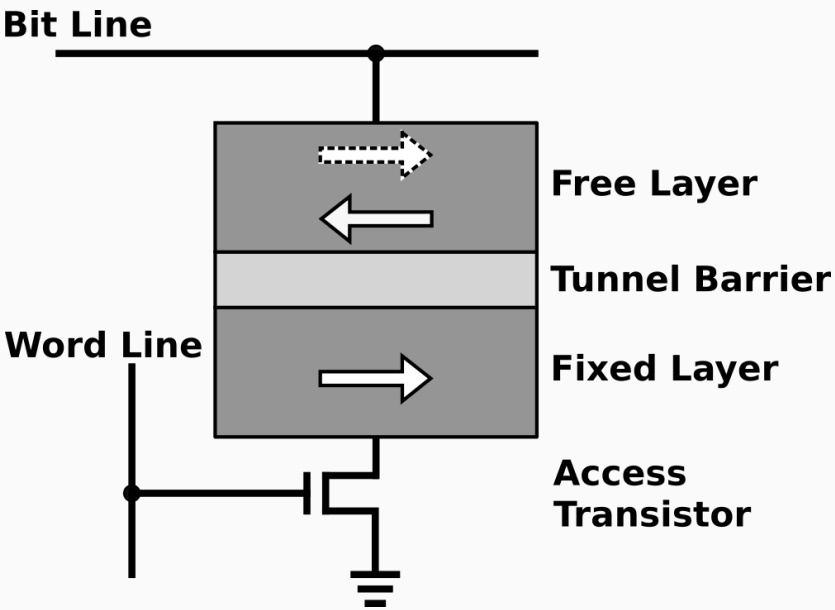

Advantages (STT-RAM/SRAM) [4, 13, 9]

Disadvantages (STT-RAM/SRAM) [4, 13, 9]

Spin-Transfer Torque RAM (STT-RAM) [11, 9, 16]

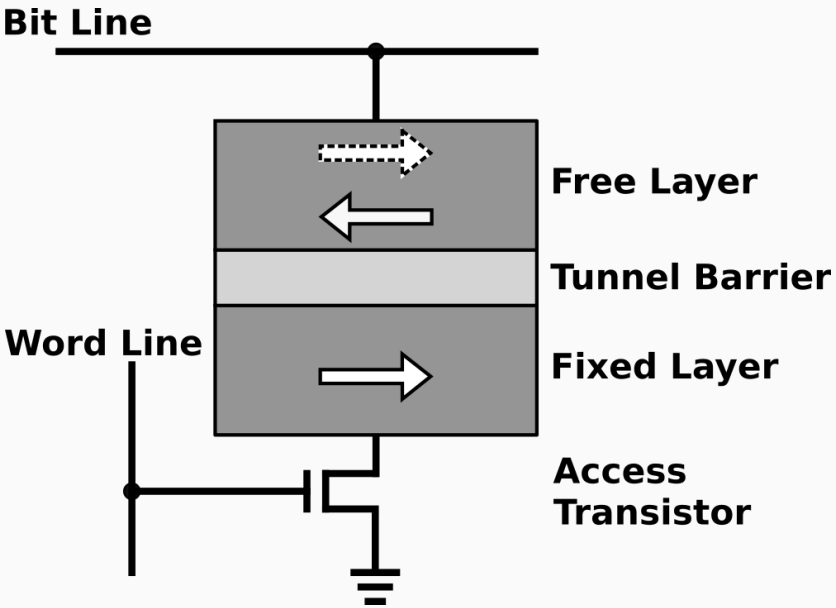

- ’o’: anti-parallel orientation (high resistance)

Advantages (STT-RAM/SRAM) [4, 13, 9]

Disadvantages (STT-RAM/SRAM) [4, 13, 9]

Spin-Transfer Torque RAM (STT-RAM) [11, 9, 16]

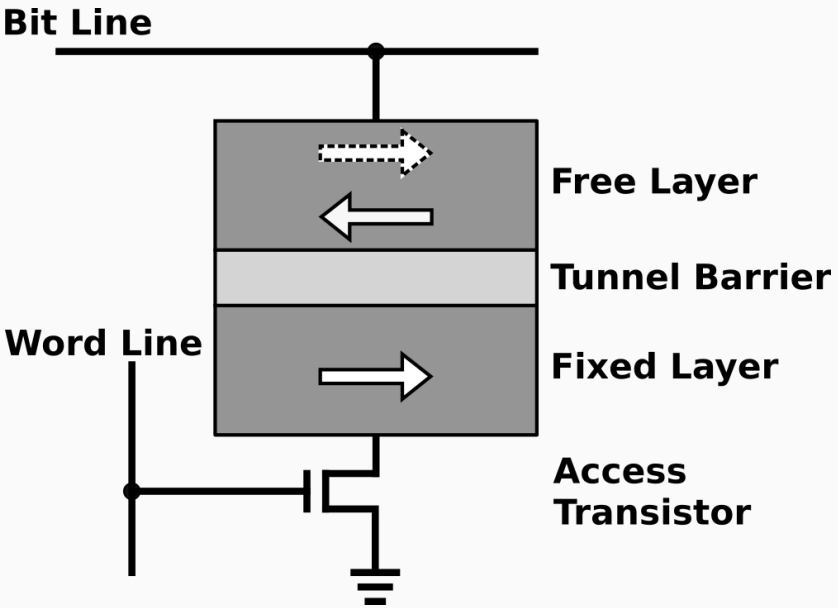

- ’0’: anti-parallel orientation (high resistance)
- ’1’: parallel orientation (low resistance)

Advantages (STT-RAM/SRAM) [4, 13, 9]

Disadvantages (STT-RAM/SRAM) [4, 13, 9]

Spin-Transfer Torque RAM (STT-RAM) [11, 9, 16]

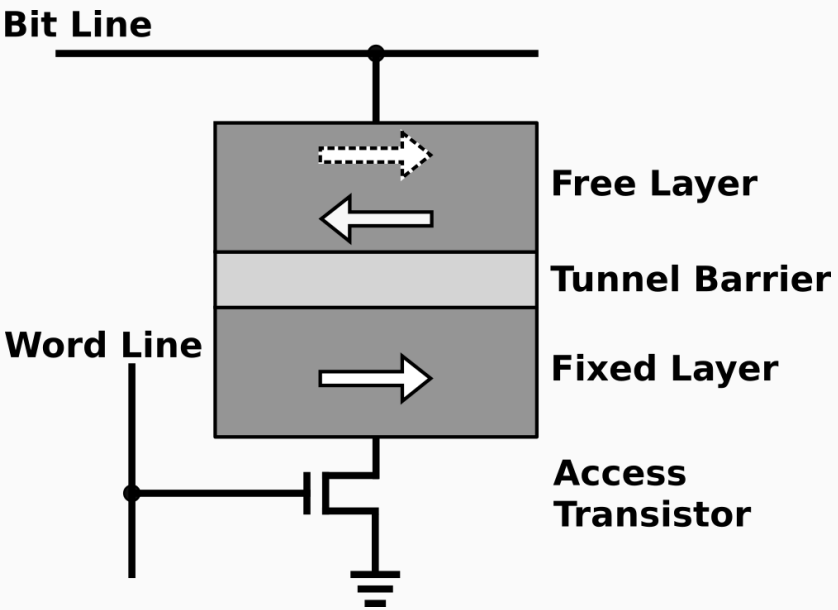

- ’0’: anti-parallel orientation (high resistance)
- ’1’: parallel orientation (low resistance)
- Write: apply current at fixed layer

Advantages (STT-RAM/SRAM) [4, 13, 9]

Disadvantages (STT-RAM/SRAM) [4, 13, 9]

Spin-Transfer Torque RAM (STT-RAM) [11, 9, 16]

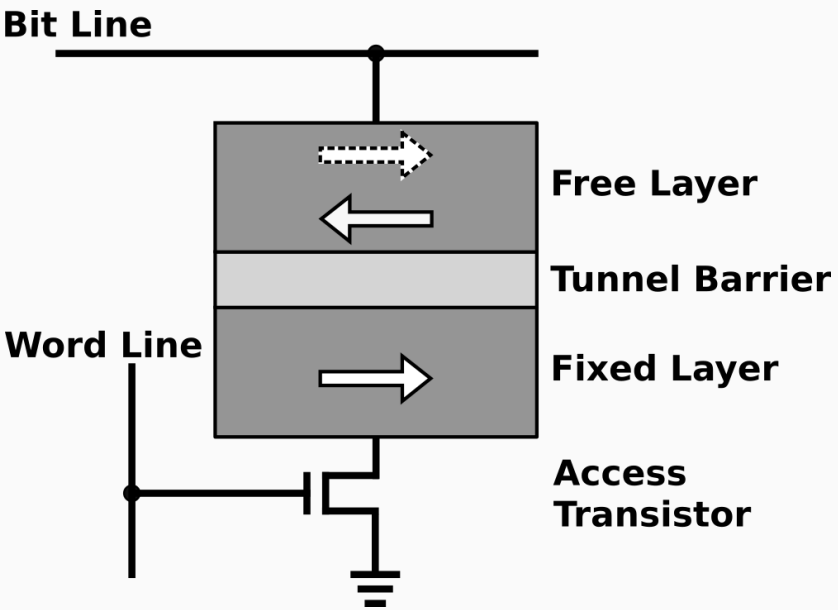

- ’0’: anti-parallel orientation (high resistance)
- ’1’: parallel orientation (low resistance)
- Write: apply current at fixed layer
- Read: measure current through cell

Advantages (STT-RAM/SRAM) [4, 13, 9]

Disadvantages (STT-RAM/SRAM) [4, 13, 9]

Spin-Transfer Torque RAM (STT-RAM) [11, 9, 16]

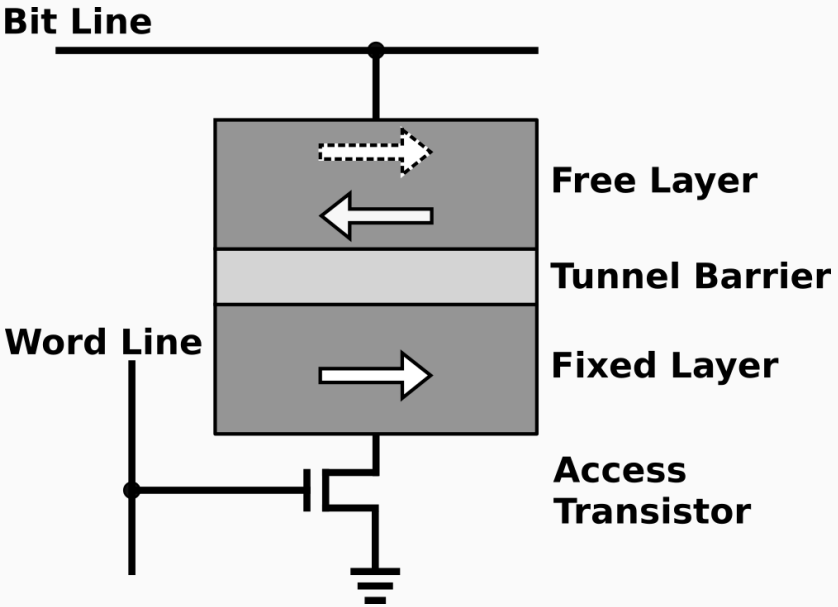

- ’0’: anti-parallel orientation (high resistance)
- ’1’: parallel orientation (low resistance)
- Write: apply current at fixed layer
- Read: measure current through cell

Advantages (STT-RAM/SRAM) [4, 13, 9]

- ✓ Read latency (1.34 ns/1.28 ns)
- ✓ Leakage power (1.82 mW/57.7 mW)
- ✓ Cell size (22 F^2/140 F^2)

Disadvantages (STT-RAM/SRAM) [4, 13, 9]

Spin-Transfer Torque RAM (STT-RAM) [11, 9, 16]

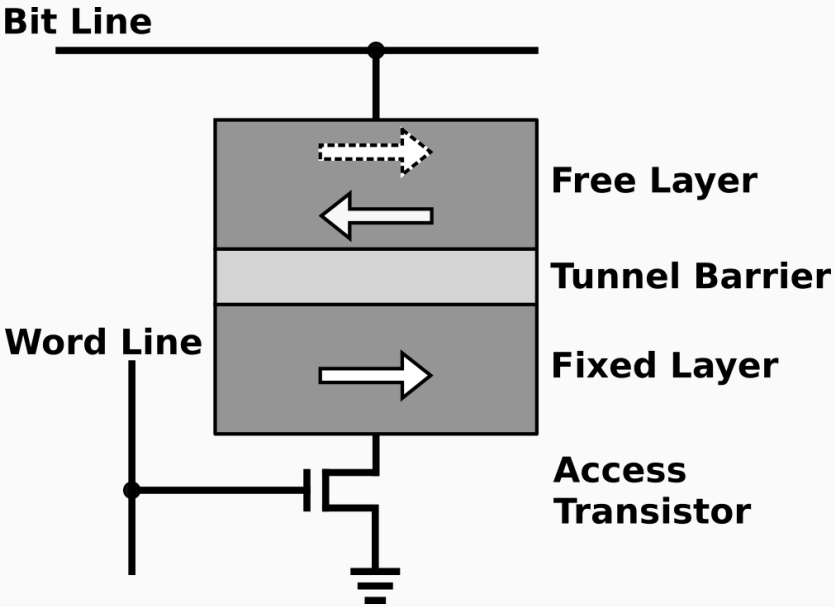

- ’0’: anti-parallel orientation (high resistance)
- ’1’: parallel orientation (low resistance)
- Write: apply current at fixed layer
- Read: measure current through cell

Advantages (STT-RAM/SRAM) [4, 13, 9]

- ✓ Read latency (1.34 ns/1.28 ns)
- ✓ Leakage power (1.82 mW/57.7 mW)
- ✓ Cell size (22 F^2/140 F^2)

Disadvantages (STT-RAM/SRAM) [4, 13, 9]

- ↯ Write latency (10.22 ns/1.23 ns)
- ↯ Write energy (0.96 nJ/0.06 nJ)
- ↯ Write endurance (10^{12} cy/10^{16} cy)

Byte-Addressable Non-Volatile Memories

Characteristics [9]:

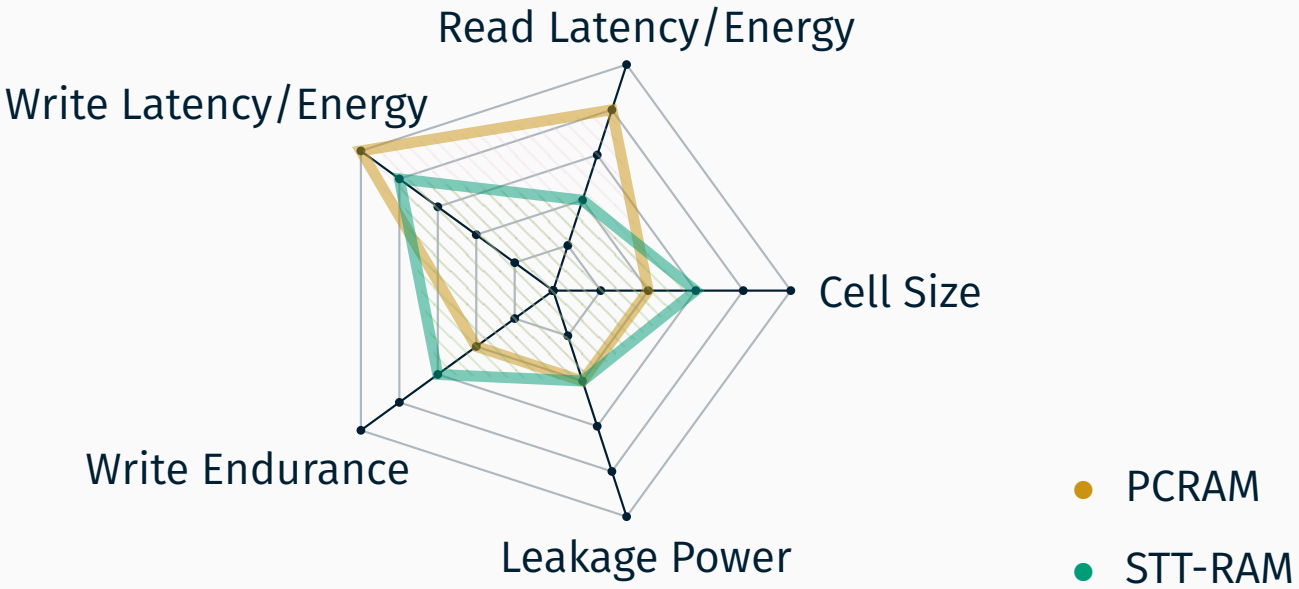

Byte-Addressable Non-Volatile Memories

Characteristics [9]:

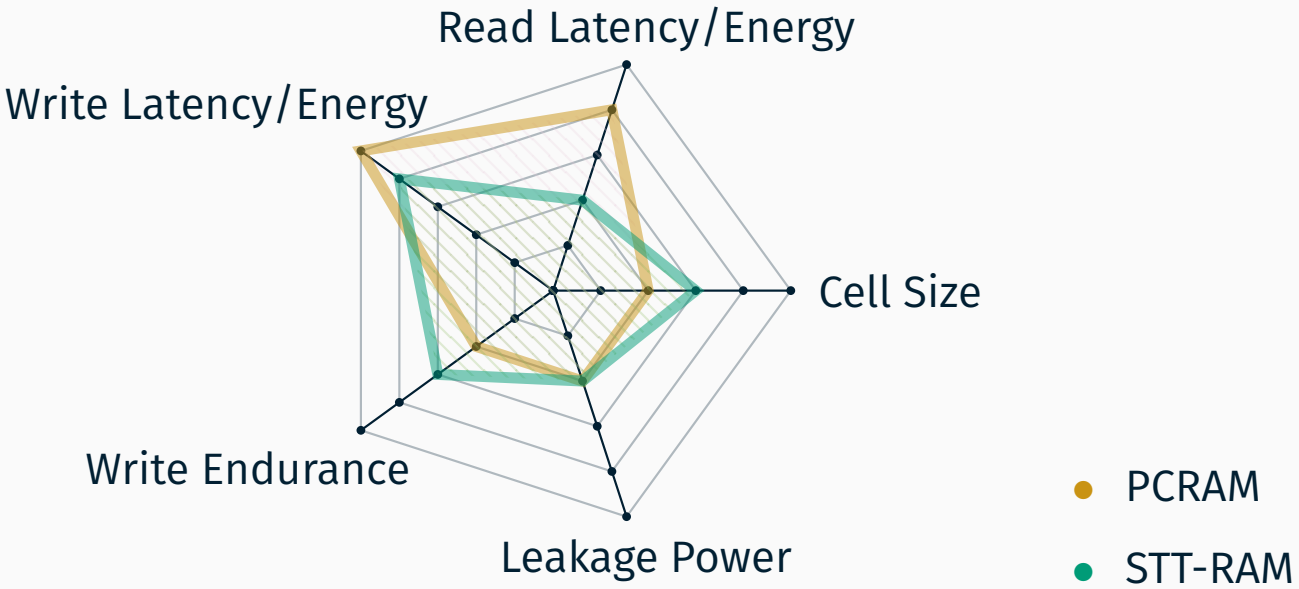

Read-Write Asymmetry [1, 19]

Expensive write operations require

$\Rightarrow$ reduced number of writes

$\Rightarrow$ changing physical properties to reduce latency

Non-Volatile Components

Non-Volatile Processor

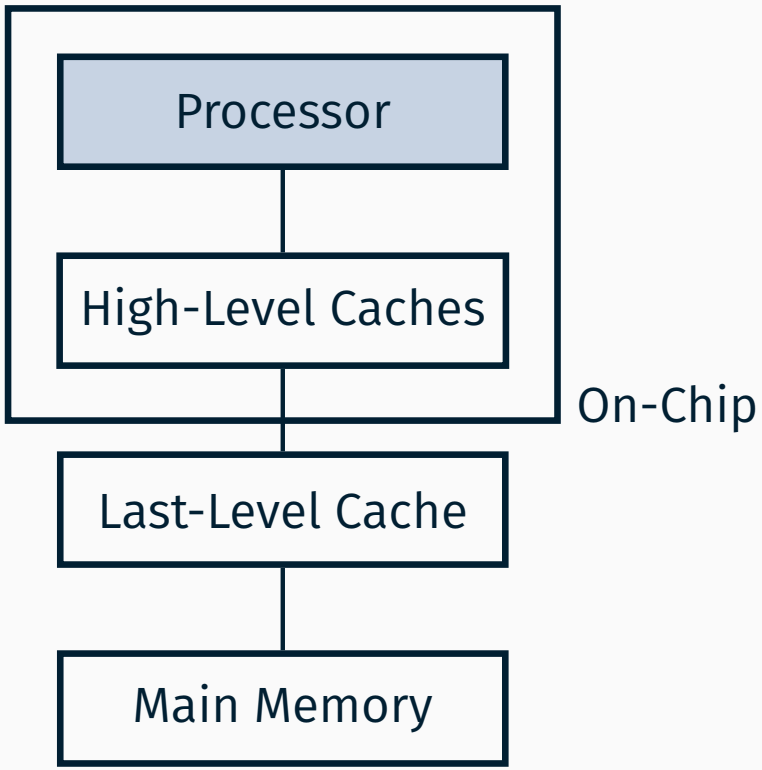

Non-Volatile Processor

Goal: Improved forward progress

Technology: STT-RAM / PCRAM

Back-Up:[7, 6]

Challenges:

Non-Volatile Processor

Goal: Improved forward progress

Technology: STT-RAM / PCRAM

Back-Up:[7, 6] Where?
→ duplicated memory components
vs. central NVM-block
What?
→ overhead per component
vs. forward progress improvement
When?
→ periodic vs. on-demand

Challenges:

Non-Volatile Processor

Goal: Improved forward progress

Technology: STT-RAM / PCRAM

Back-Up:[7, 6] Where?
→ duplicated memory components
vs. central NVM-block
What?
→ overhead per component
vs. forward progress improvement
When?
→ periodic vs. on-demand

Challenges:
- Limited write endurance
- Overhead due to back-up and restore operations

Non-Volatile Caches

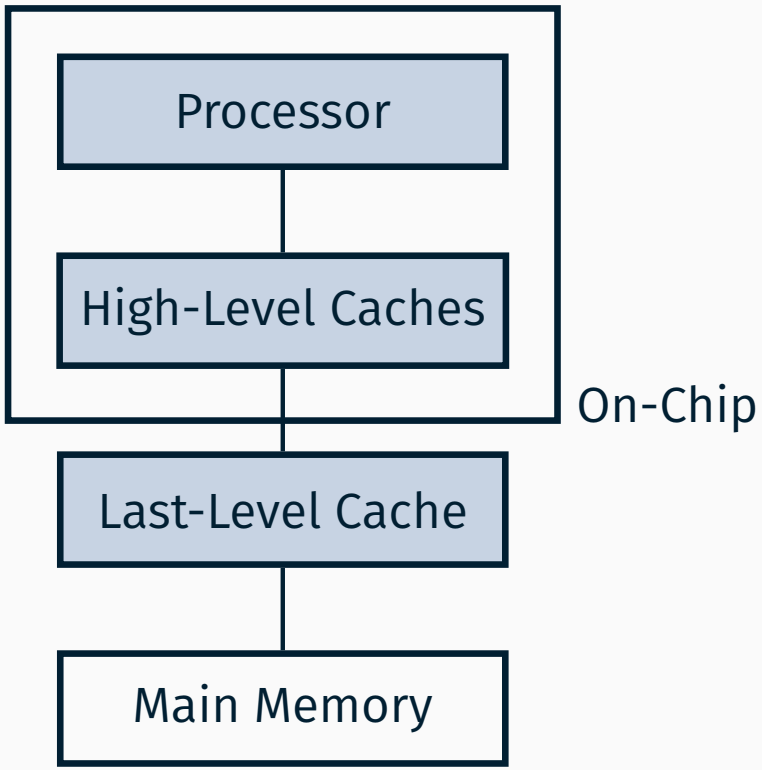

Non-Volatile Caches

Goal: Improved power consumption

Technology: STT-RAM

Challenges:

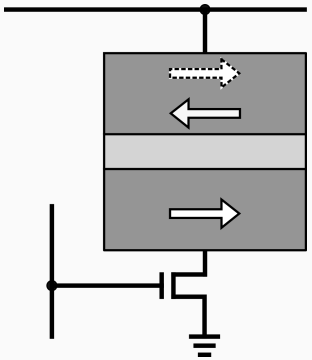

Applicability:

Non-Volatile Caches

Goal: Improved power consumption

Technology: STT-RAM

Challenges:

- Write latency and energy
  → retention relaxation [12, 13]
- Write endurance
  → wear leveling [3]
  → reducing the number of writes [20]

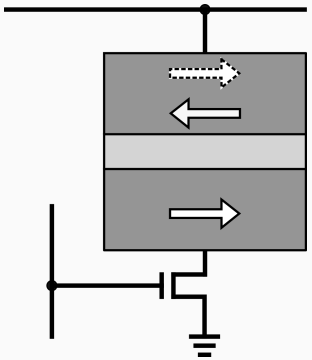

Applicability:

Non-Volatile Caches

Goal: Improved power consumption

Technology: STT-RAM

Challenges:

- Write latency and energy
  → retention relaxation [12, 13]
- Write endurance
  → wear leveling [3]
  → reducing the number of writes [20]

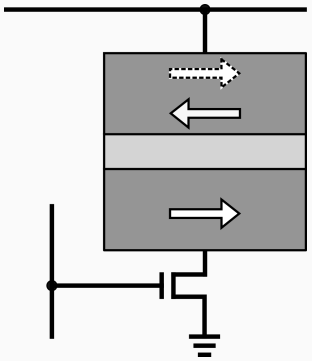

Applicability:

- Last-level cache [5]
- Higher-level cache with relaxed retention times [13]

Non-Volatile Main Memory

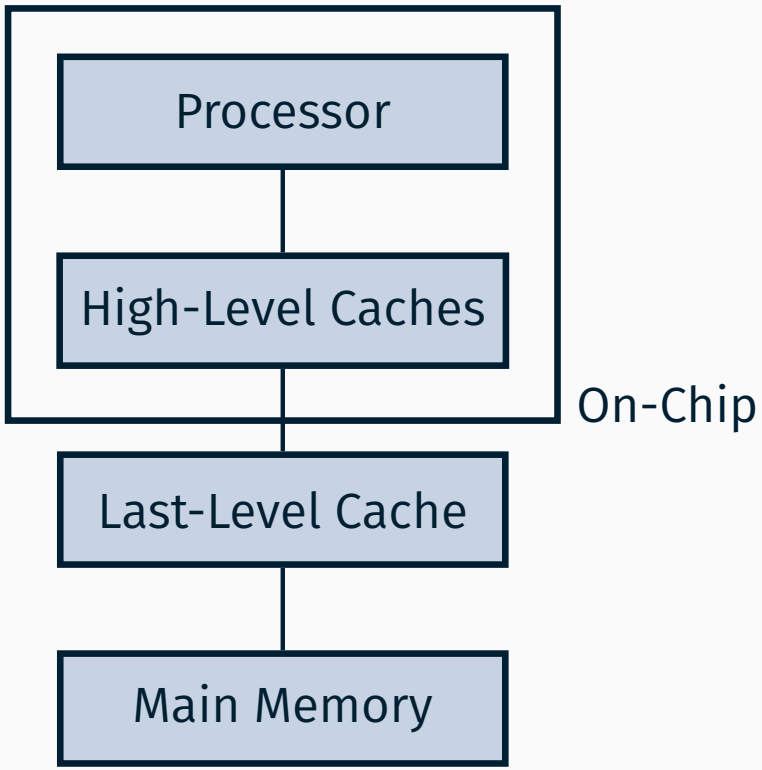

Non-Volatile Main Memory

Goal: Improved power consumption and non-volatility

Technology: PCRAM

Challenges:

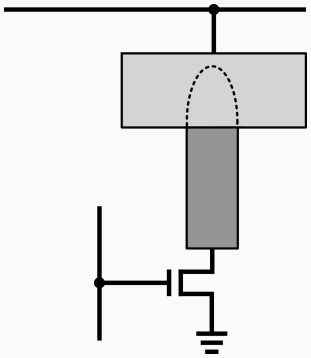

Applicability:

Non-Volatile Main Memory

Goal: Improved power consumption and non-volatility

Technology: PCRAM

Challenges:

- Consistency
  - → traditional transactions [14]
  - → memory-hierarchy based transactions [18, 10]
- Write endurance
  - → wear leveling [15]
  - → reducing the number of writes [10, 14]
  - → fault-tolerance [17, 2]

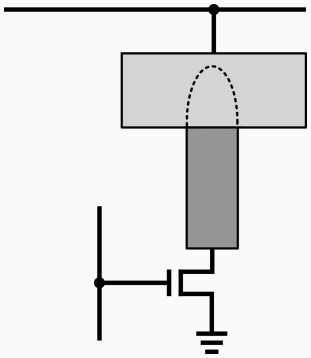

Applicability:

Non-Volatile Main Memory

Goal: Improved power consumption and non-volatility

Technology: PCRAM

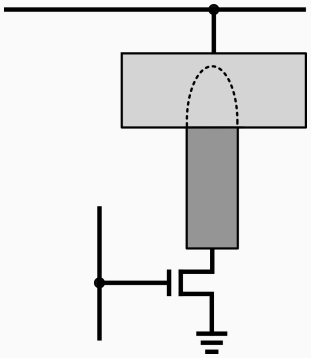

Challenges:

- Consistency
  - → traditional transactions [14]
  - → memory-hierarchy based transactions [18, 10]
- Write endurance
  - → wear leveling [15]
  - → reducing the number of writes [10, 14]
  - → fault-tolerance [17, 2]

Applicability:

- With volatile buffer [10, 14]
- With non-volatile last-level cache [18]

Non-Volatile Embedded System

Non-Volatile Embedded System

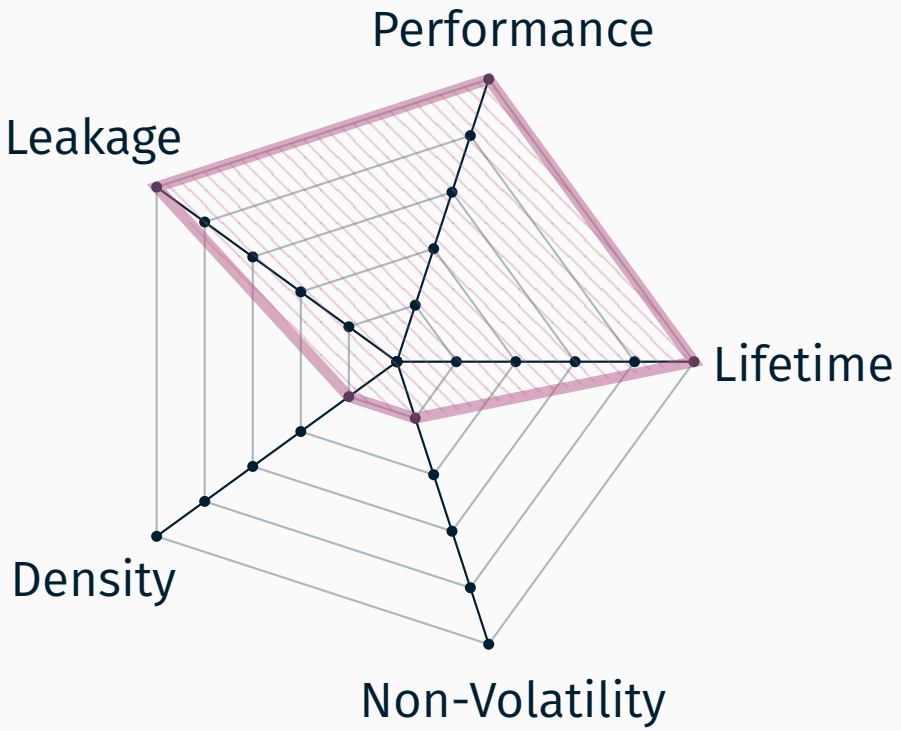

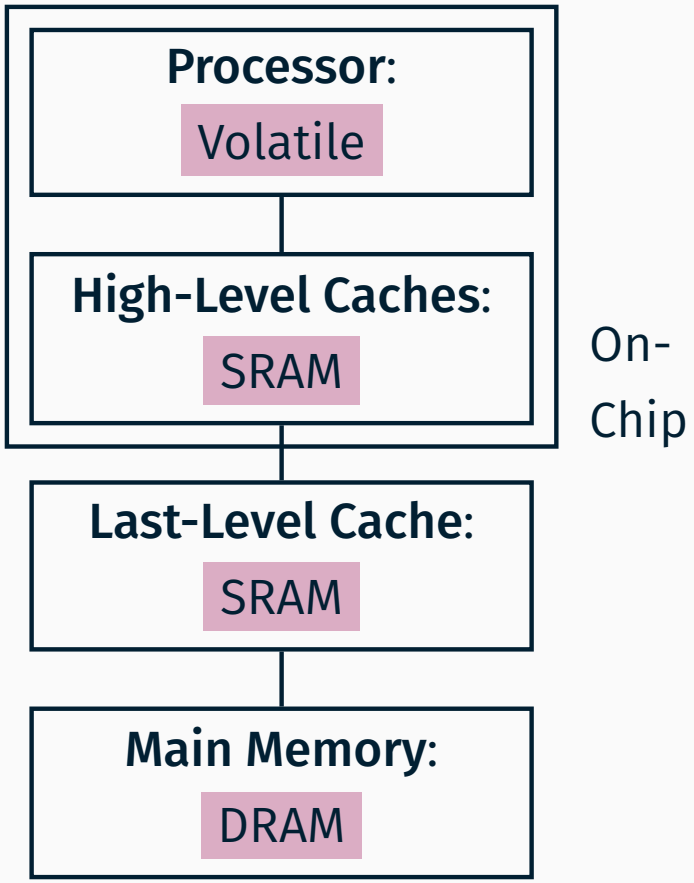

Non-Volatile Embedded System

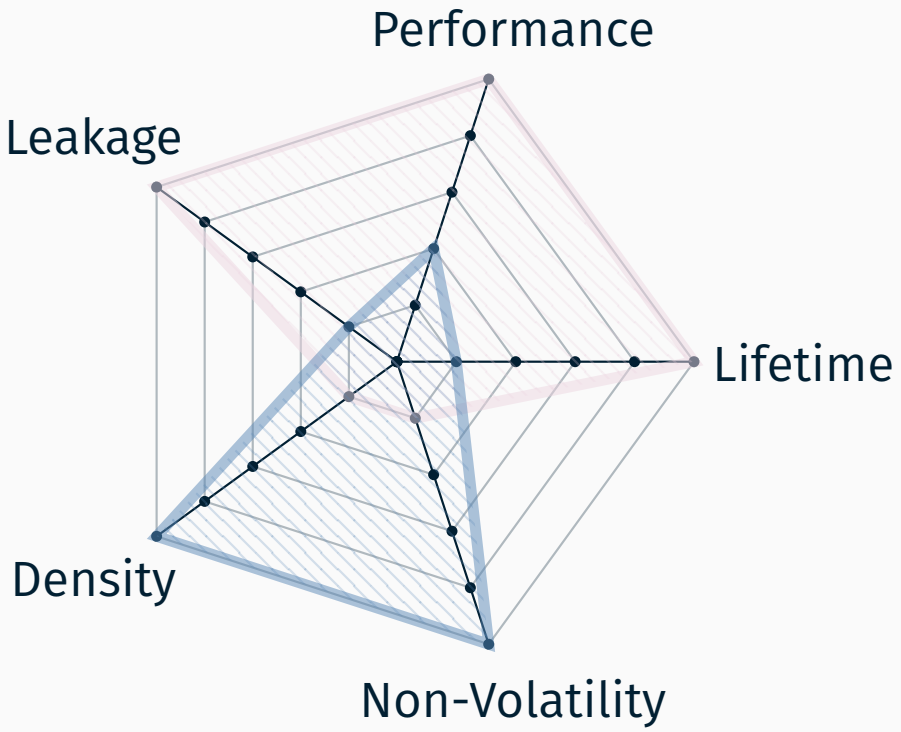

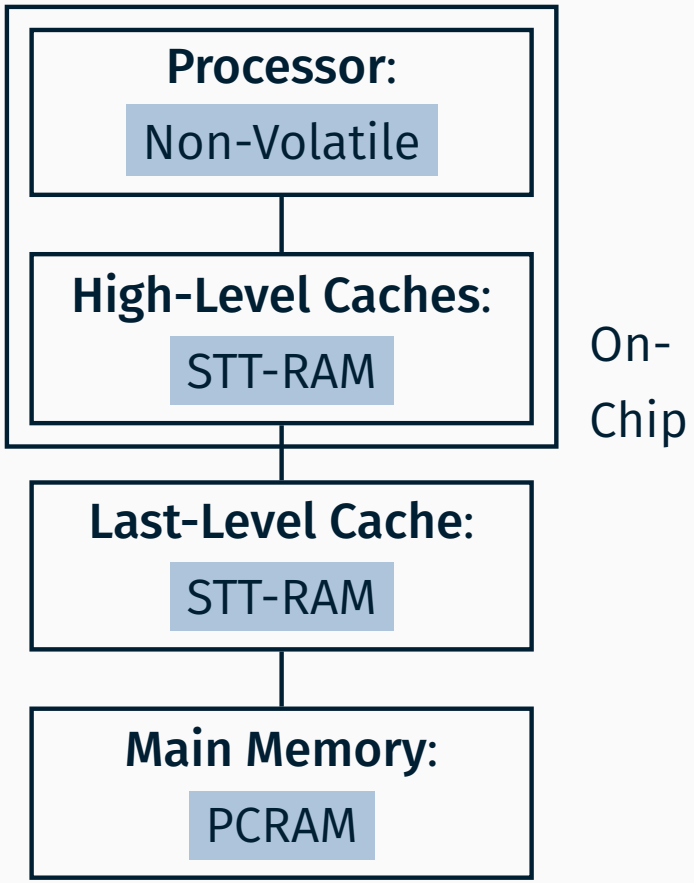

Non-Volatile Embedded System

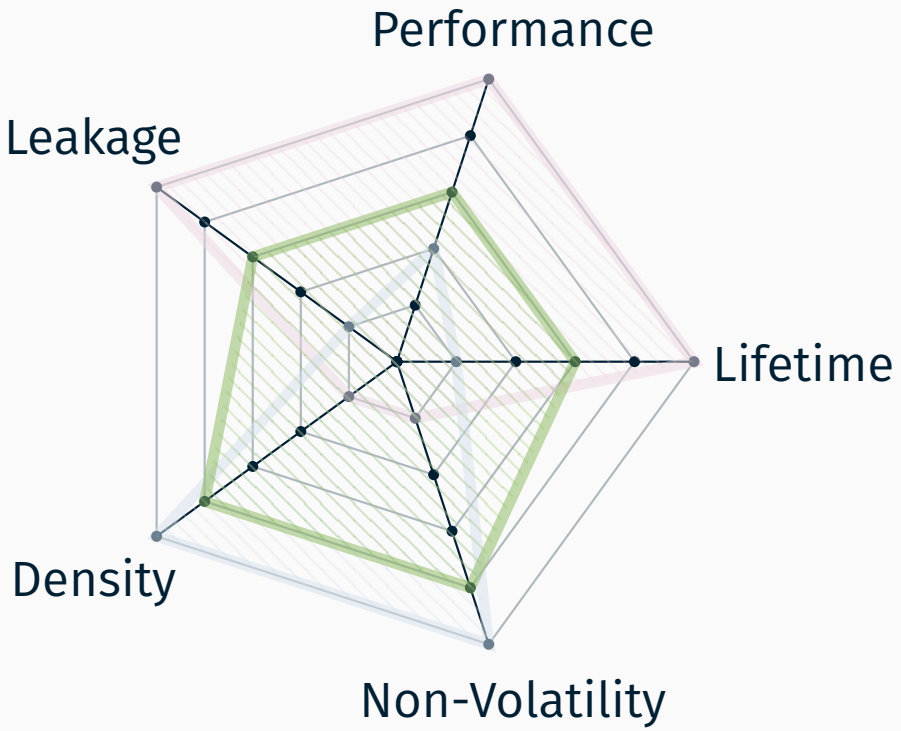

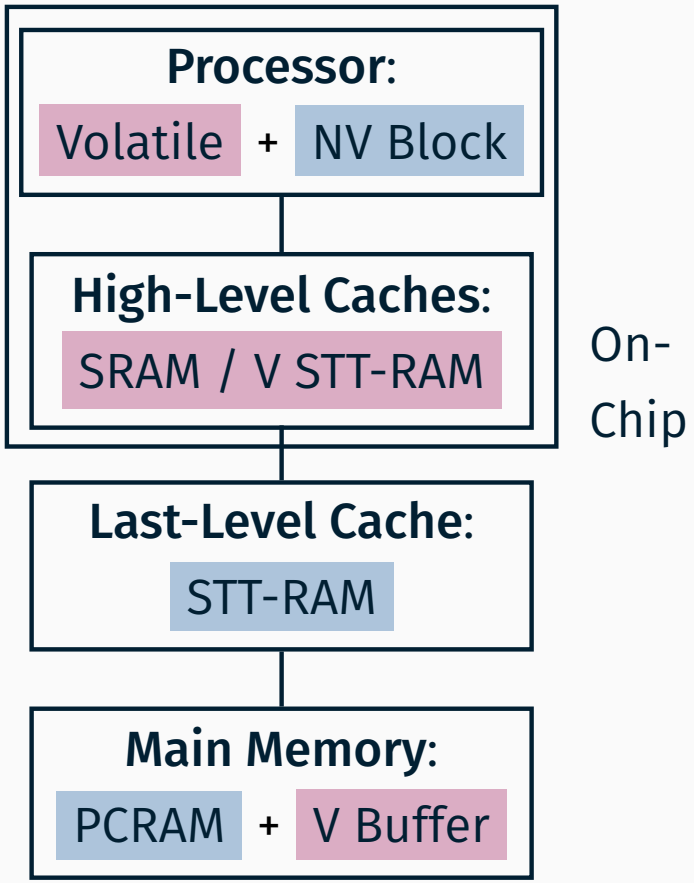

Non-Volatile Embedded System

Goal: Improved forward progress and power consumption

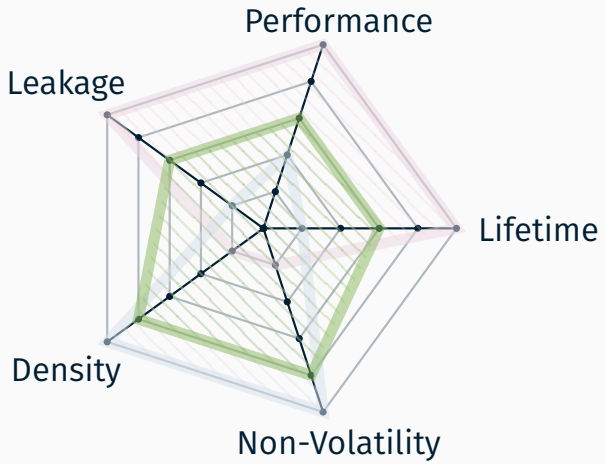

Composition:

Software:

Non-Volatile Embedded System

Goal: Improved forward progress and power consumption

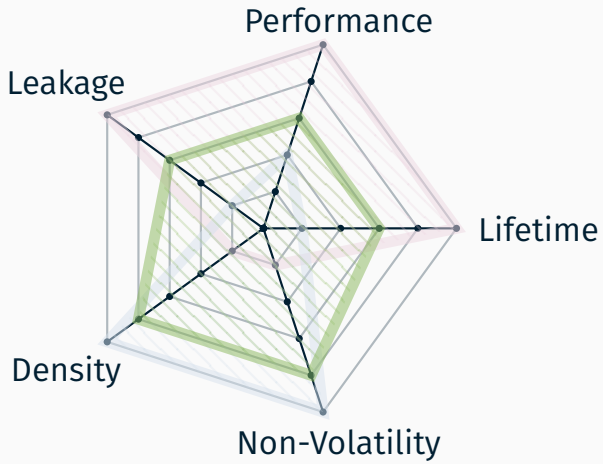

Composition:

- Processor
  → volatile flipflops
  → nv-block as back-up location
- Caches
  → higher level: relaxed STT-RAM or SRAM
  → last level: slightly relaxed STT-RAM
- Main Memory
  → PCRAM + volatile DRAM buffer

Software:

Non-Volatile Embedded System

Goal: Improved forward progress and power consumption

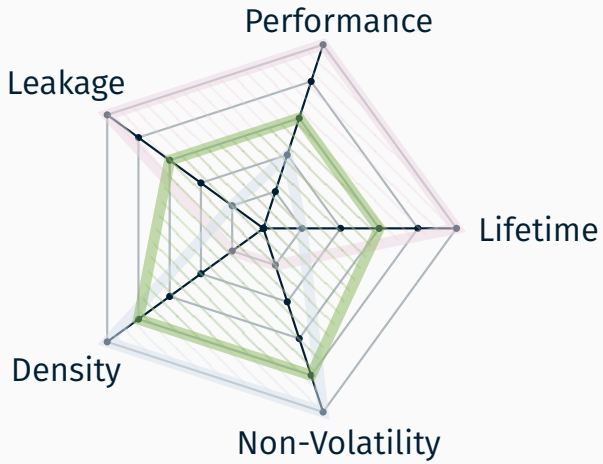

Composition:

- Processor
  → volatile flipflops
  → nv-block as back-up location
- Caches
  → higher level: relaxed STT-RAM or SRAM
  → last level: slightly relaxed STT-RAM
- Main Memory
  → PCRAM + volatile DRAM buffer

Software:

- Endurance-aware memory-allocation [10]
- OS-supported wear-leveling [15]

References (1)

[1] Rajendra Bishnoi, Fabian Oboril, Mojtaba Ebrahimi, and Mehdi B Tahoori.
Avoiding unnecessary write operations in stt-mram for low power implementation.
In *Proceedings of the 15th International Symposium on Quality Electronic Design (ISQED '14)*, pages 548–553, 2014.

[2] Yu Cai, Onur Mutlu, Erich F Haratsch, and Ken Mai.
Program interference in mlc nand flash memory: Characterization, modeling, and mitigation.
In *Proceedings of the 31st International Conference on Computer Design (ICCD '13)*, pages 123–130, 2013.

[3] Yiran Chen, Weng-Fai Wong, Hai Li, Cheng-Kok Koh, Yaojun Zhang, and Wujie Wen.
On-chip caches built on multilevel spin-transfer torque ram cells and its optimizations.
J. Emerg. Technol. Comput. Syst., 9:16:1–16:22, 2013.

[4] Navid Khoshavi and Ronald F Demara.
Read-tuned stt-ram and edram cache hierarchies for throughput and energy optimization.
IEEE Access, 6:14576–14590, 2018.

[5] Yong Li, Yiran Chen, and Alex K. Jones.
A software approach for combating asymmetries of non-volatile memories.
In *Proceedings of the 2012 ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED '12)*, pages 191–196, 2012.

[6] Kaisheng Ma, Xueqing Li, Shuangchen Li, Yongpan Liu, John Jack Sampson, Yuan Xie, and Vijaykrishnan Narayanan.
Nonvolatile processor architecture exploration for energy-harvesting applications.
IEEE Micro, 35(5):32–40, 2015.

References (2)

[7] Kaisheng Ma, Yang Zheng, Shuangchen Li, Karthik Swaminathan, Xueqing Li, Yongpan Liu, Jack Sampson, Yuan Xie, and Vijaykrishnan Narayanan.
Architecture exploration for ambient energy harvesting nonvolatile processors.
In *IEEE 21st International Symposium on High Performance Computer Architecture (HPCA '15)*, pages 526–537, 2015.

[8] Sparsh Mittal and Jeffrey S Vetter.
A survey of software techniques for using non-volatile memories for storage and main memory systems.
IEEE Transactions on Parallel and Distributed Systems, 27(5):1537–1550, 2015.

[9] Sparsh Mittal, Jeffrey S Vetter, and Dong Li.
A survey of architectural approaches for managing embedded dram and non-volatile on-chip caches.
IEEE Transactions on Parallel and Distributed Systems, 26(6):1524–1537, 2014.

[10] Iulian Moraru, David G. Andersen, Michael Kaminsky, Niraj Tolia, Parthasarathy Ranganathan, and Nathan Binkert.
Consistent, durable, and safe memory management for byte-addressable non volatile main memory.
In *Proceedings of the 1st ACM SIGOPS Conference on Timely Results in Operating Systems (TRIOS '13)*, pages 1:1–1:17, 2013.

[11] Rodrigue Rizk, Dominick Rizk, Ashok Kumar, and Magdy Bayoumi.
Demystifying emerging nonvolatile memory technologies: Understanding advantages, challenges, trends, and novel applications.
In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS '19)*, pages 1–5, 2019.

References (3)

[12] Clinton W Smullen, Vidyabhushan Mohan, Anurag Nigam, Sudhanva Gurumurthi, and Mircea R Stan.
Relaxing non-volatility for fast and energy-efficient stt-ram caches.
In *Proceedings of the 17th IEEE International Symposium on High Performance Computer Architecture (HPCA '11)*, pages 50–61, 2011.

[13] Zhenyu Sun, Xiuyuan Bi, Hai (Helen) Li, Weng-Fai Wong, Zhong-Liang Ong, Xiaochun Zhu, and Wenqing Wu.
Multi retention level stt-ram cache designs with a dynamic refresh scheme.
In *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '11)*, pages 329–338, 2011.

[14] Haris Volos, Andres Jaan Tack, and Michael M. Swift.
Mnemosyne: Lightweight persistent memory.
In *Proceedings of the 16th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '14)*, pages 91–104, 2011.

[15] Chundong Wang and Weng-Fai Wong.
Saw: System-assisted wear leveling on the write endurance of nand flash devices.
In *Proceedings of the 50th Annual Design Automation Conference (DAC '13)*, page 164, 2013.

[16] L Wang, C-H Yang, and J Wen.
Physical principles and current status of emerging non-volatile solid state memories.
Electronic Materials Letters, 11(4):505–543, 2015.

[17] Doe Hyun Yoon, Naveen Muralimanohar, Jichuan Chang, Parthasarathy Ranganathan, Norman P Jouppi, and Mattan Erez.
Free-p: Protecting non-volatile memory against both hard and soft errors.
In *Proceedings of the 17th International Symposium on High Performance Computer Architecture (HPCA '11)*, pages 466–477, 2011.

References (4)

[18] Jishen Zhao, Sheng Li, Doe Hyun Yoon, Yuan Xie, and Norman P Jouppi.
Kiln: Closing the performance gap between systems with and without persistence support.
In *Proceedings of the 46th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '13)*, pages 421–432, 2013.

[19] Ping Zhou, Bo Zhao, Jun Yang, and Youtao Zhang.
A durable and energy efficient main memory using phase change memory technology.
In *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA '09)*, pages 14–23, 2009.

[20] Ping Zhou, Bo Zhao, Jun Yang, and Youtao Zhang.
Energy reduction for stt-ram using early write termination.
In *Proceedings of the International Conference on Computer-Aided Design (ICCAD '09)*, pages 264–268, 2009.