Functional and Non-Functional Properties of NVRAM for Embedded Systems: An Analysis and Discussion

24. September 2020

Henriette Hofmeier, Stefan Reif, Tobias Langer, Laura Lawniczak, Timo Hönig and Wolfgang Schröder-Preikschat

Fachgruppe Betriebssysteme: Herbstreffen 2020

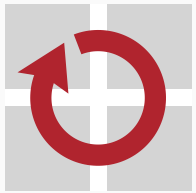

Chair in Distributed Systems
and Operating Systems

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG

FACULTY OF ENGINEERING

Motivation

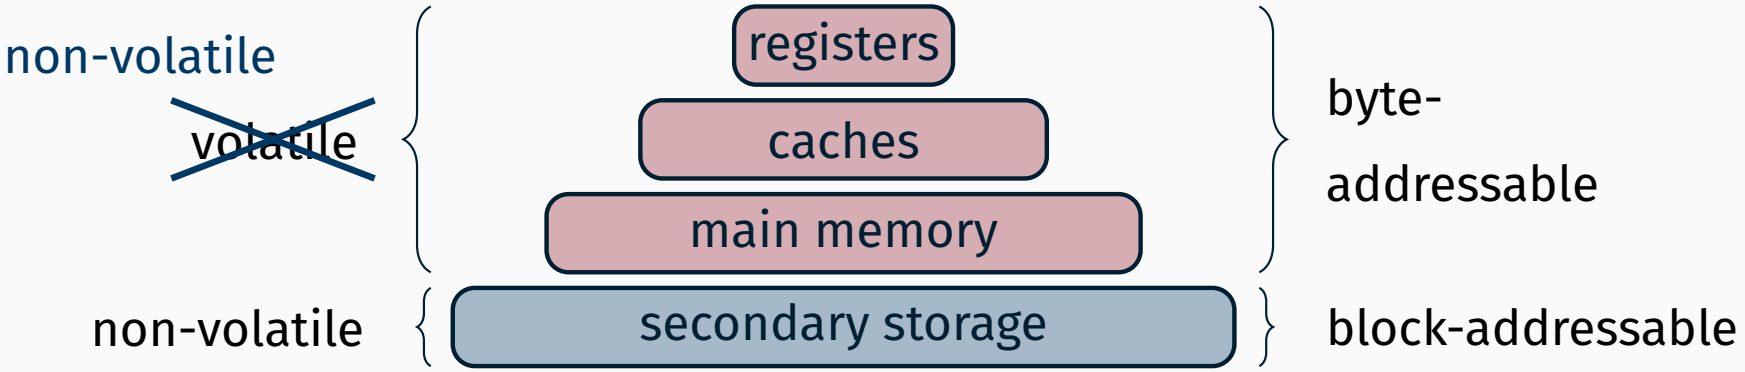
non-volatile
volatile
registers
caches
main memory
byte-addressable
non-volatile
secondary storage
block-addressable

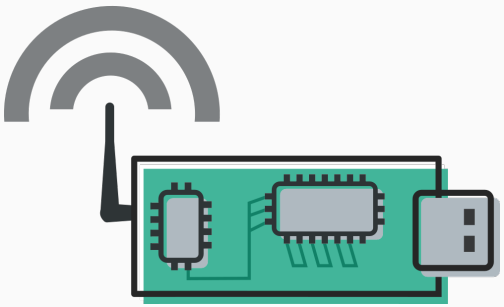

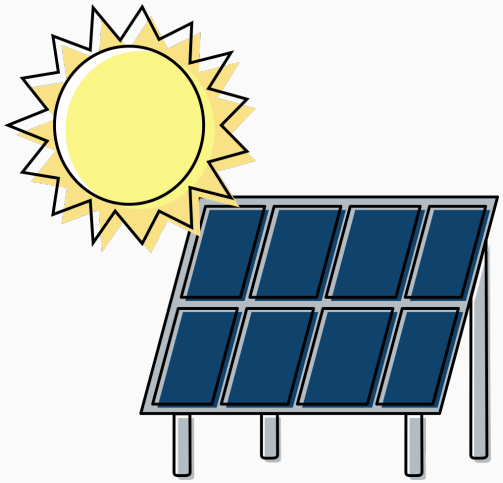

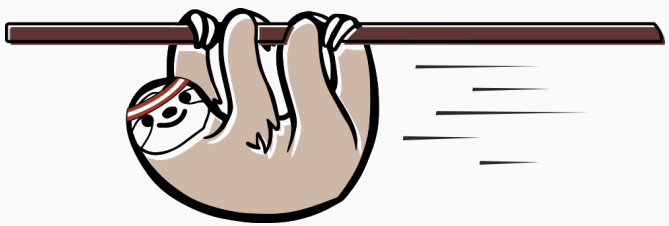

Overview

Motivation

Technologies

Components

Challenges

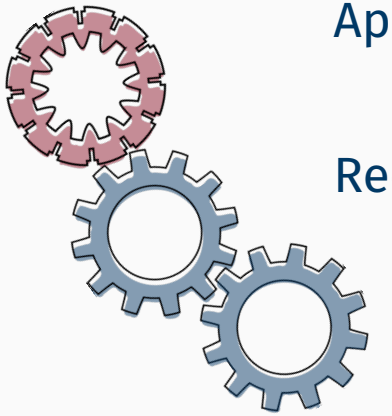

Approaches

Recap

Technologies

Technologies [11, 19, 22, 23]

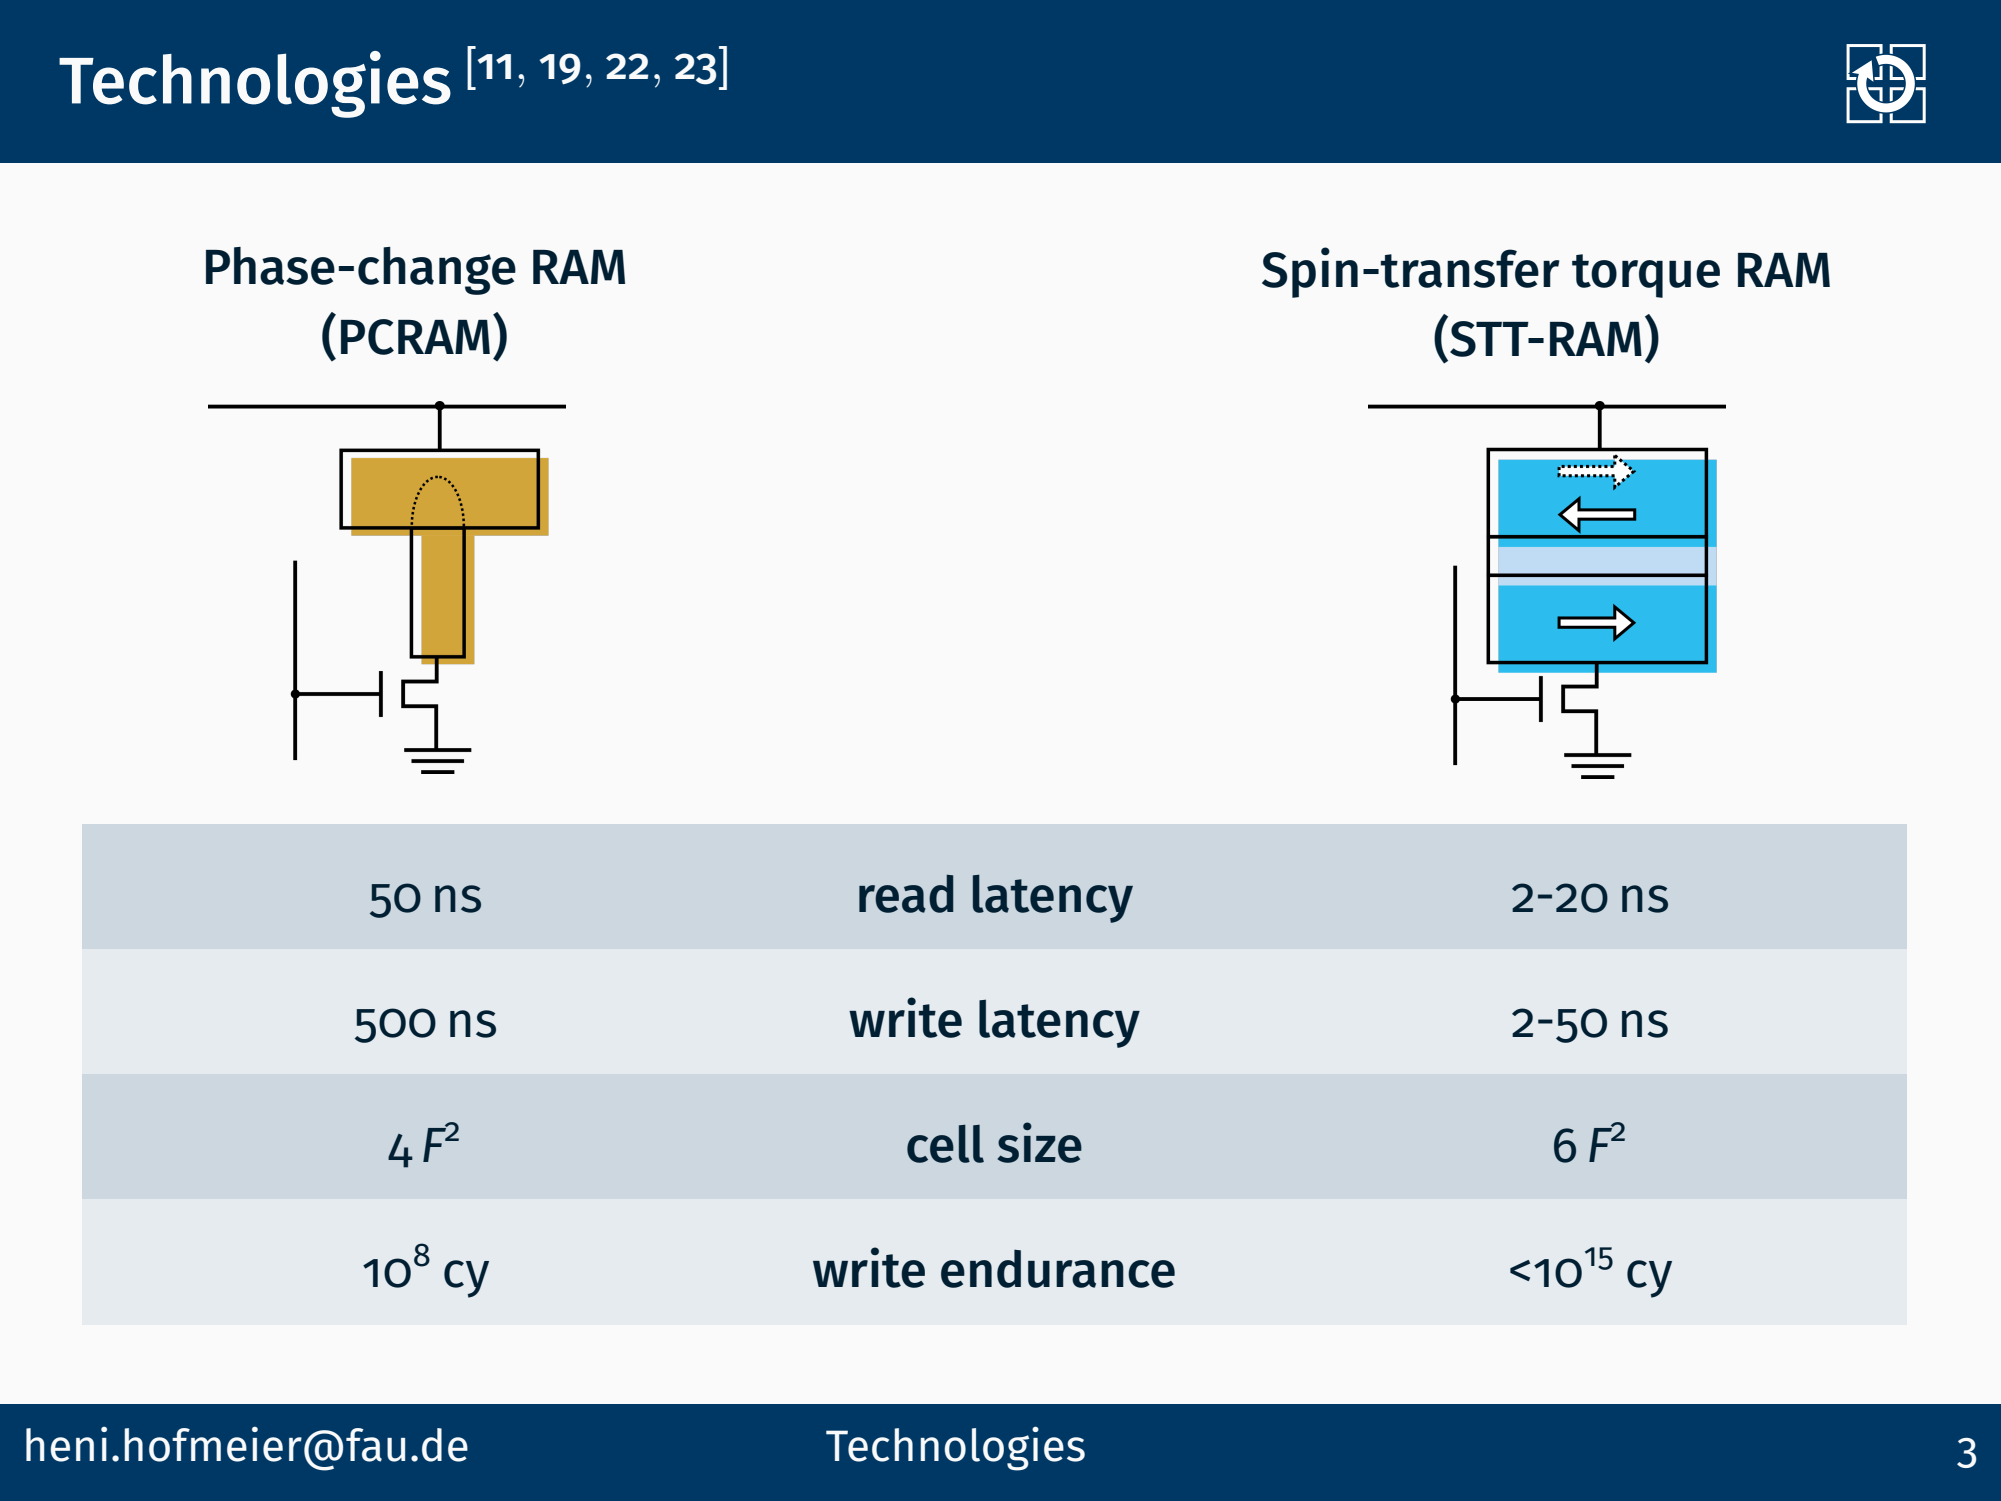

Phase-change RAM (PCRAM)

Spin-transfer torque RAM (STT-RAM)

50 ns	**read latency**	2-20 ns
500 ns	**write latency**	2-50 ns
$4\,F^2$	**cell size**	$6\,F^2$
10^8 cy	**write endurance**	$<10^{15}$ cy

Characteristics [12]

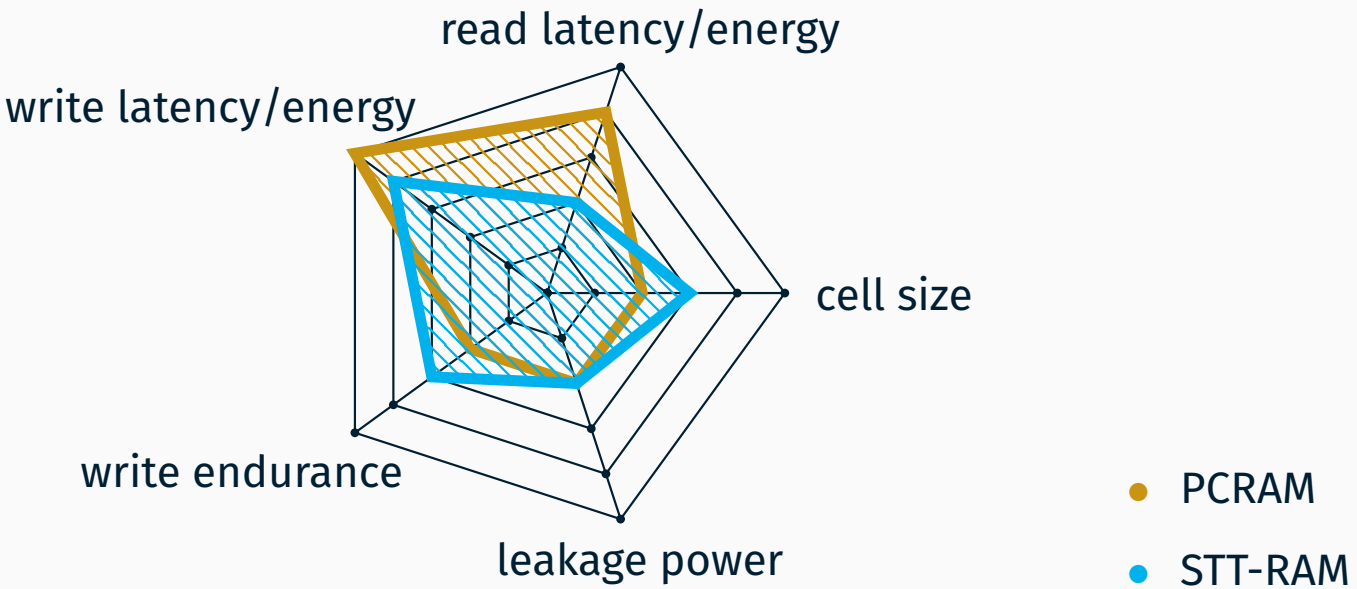

Read-write asymmetry [2, 25]

Expensive write operations require

⇒ reduced number of writes

⇒ changing physical properties to reduce latency

Components

Processor [9, 10]

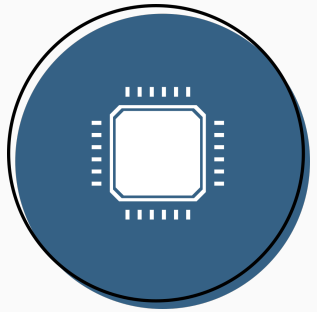

Goal: improved forward progress

Technology: STT-RAM / PCRAM

Implementation:
- replacing volatile components ***or***
- duplicating components ***or***
- adding NVM-block for processor state backup

Caches [8, 18, 24]

Goal:
- improved power consumption
- non-volatility

Technology: STT-RAM

Implementation:
- replacing last-level cache ***and/or***
- replacing higher-level cache
  (and changing physical attributes of STT-cells)

Main memory [13, 20, 24]

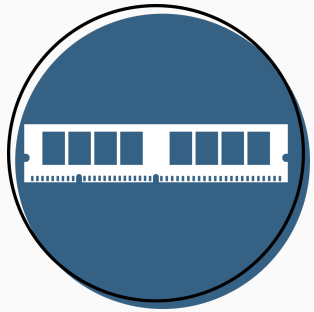

Goal:
- non-volatility
- improved power consumption

Technology: PCRAM

Implementation:
- replacing DRAM
  → possible combination with non-volatile last-level-cache
- hybrid architecture:
  → both volatile and non-volatile main memory ***or***
  → non-volatile main memory with volatile buffer

Configurations

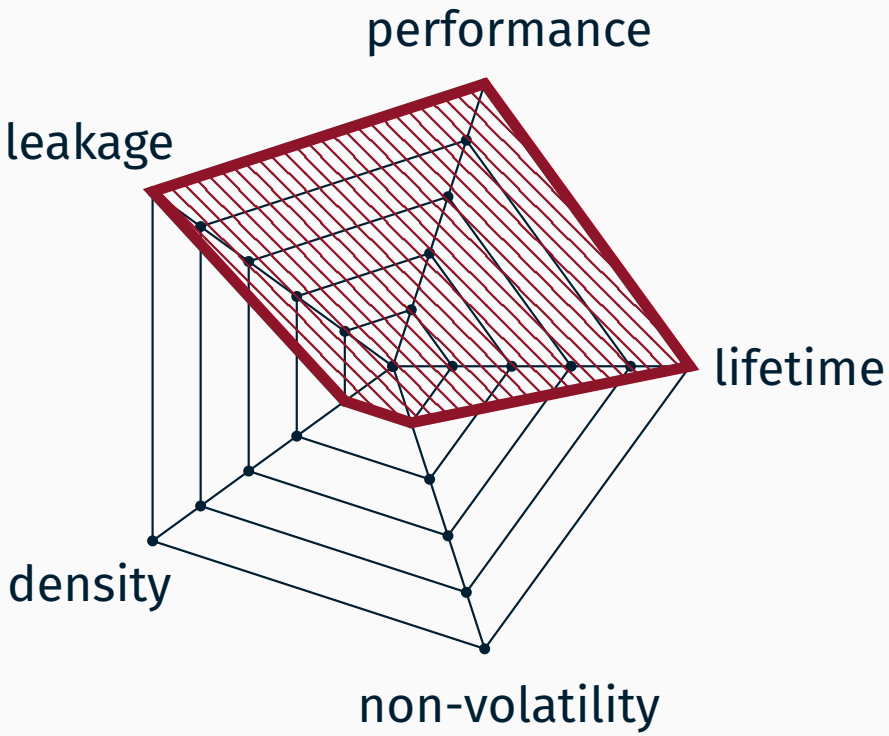

processor:
volatile

high-level caches:
SRAM

last-level cache:
SRAM

main memory:
DRAM

Configurations

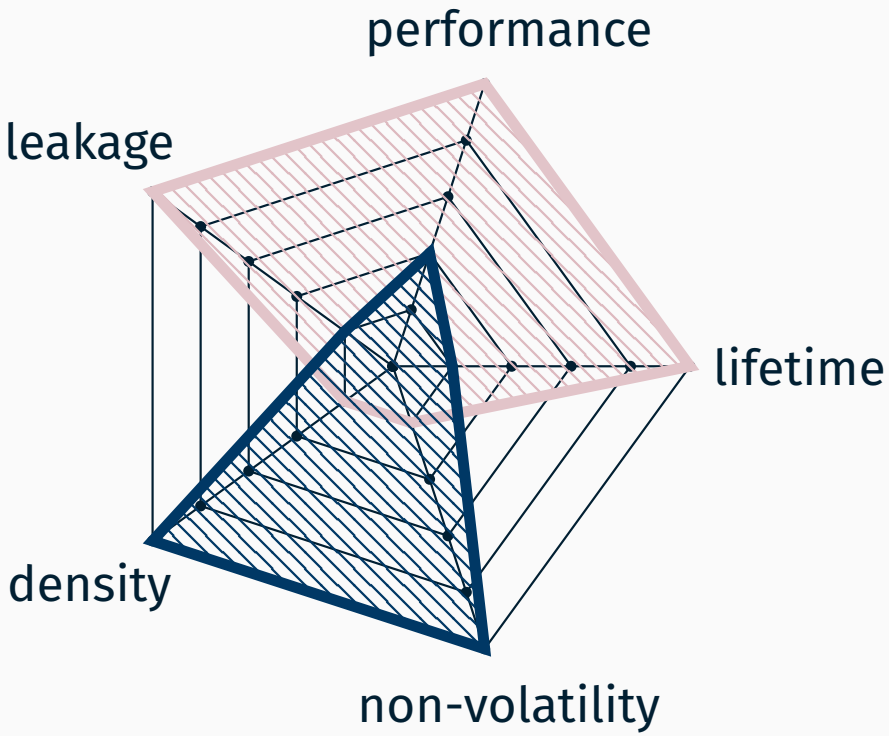

processor:
non-volatile

first-level cache:
STT-RAM

last-level cache:
STT-RAM

main memory:
PCRAM

Configurations

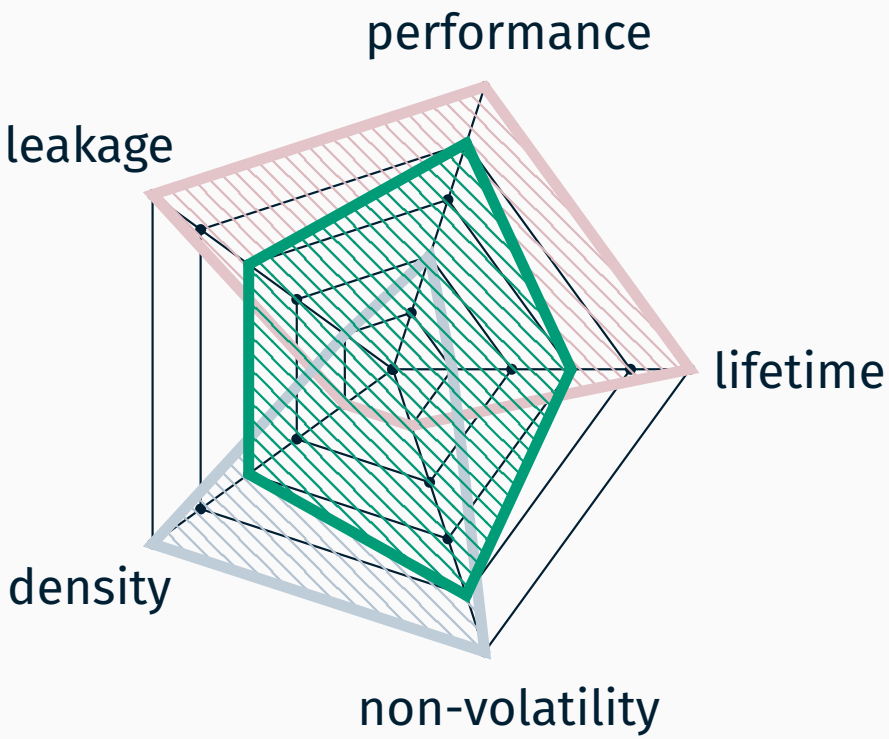

processor:
volatile + nv block

high-level caches:
SRAM / v STT-RAM

last-level cache:
STT-RAM

main memory:
PCRAM + v buffer

Challenges

Challenges

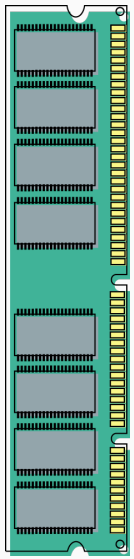

Hardware

→ low write endurance [1]

→ read-write asymmetry [2]

→ high write energy [3]

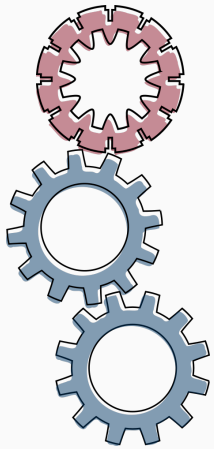

Software

→ inconsistency [4]

→ corruption [5]

Security [6]

1: [4, 13, 20, 21, 26] 2: [2, 25, 26] 3: [13, 20, 17, 18, 20, 26] 4: [1, 15, 14, 16, 5, 7, 20, 24] 5,6: [1]

Approaches

Approaches

Hybrid systems

→ NVRAM and DRAM DIMMs
→ persistence as a service
→ *Twizzler* [3]

Whole–system persistence[14]

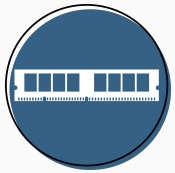

→ only NVRAM DIMMs
→ all application data persistent
→ all system data persistent

Neverlast: Towards the Design and Implementation of the NVM-based Everlasting Operating System[6]

→ implemented on MSP-EXP430FR5739
→ all data resides in NVRAM
→ prototype including process management

To appear in *Proceedings of the 2021 Annual Hawaii International Conference on System Sciences (HICSS '21)*, 2021.

Recap

NVRAM ...

... is available in form of different **technologies**

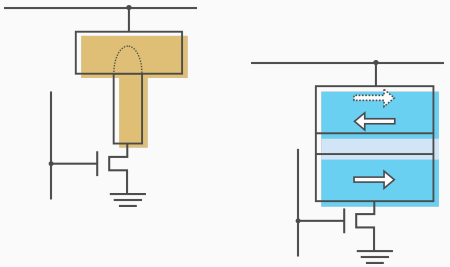

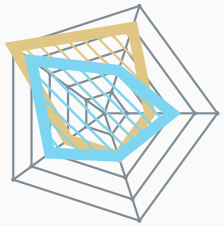

... suffers from **read-write-asymmetry**

... is used for **components** of non-volatile systems

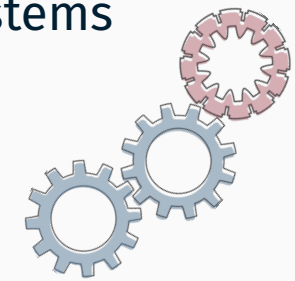

... poses **new challenges** to system developers

... allows for systems with

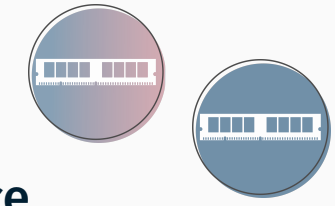

→ **hybrid** architectures

→ **whole-system persistence**

... leaves a wide field for **future research**

References (1)

[1] Katelin Bailey, Luis Ceze, Steven D Gribble, and Henry M Levy.
Operating system implications of fast, cheap, non-volatile memory.
In *Proceedings of the 2011 USENIX Workshop on Hot Topics in Operating Systems (HotOS '11)*, volume 13, pages 2–2, 2011.

[2] 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 2014 International Symposium on Quality Electronic Design (ISQED '14)*, pages 548–553, 2014.

[3] Daniel Bittman, Peter Alvaro, Pankaj Mehra, Darrell DE Long, and Ethan L Miller.
Twizzler: a data-centric os for non-volatile memory.
In *Proceedings of the 2020 USENIX Annual Technical Conference (ATC '20)*, pages 65–80, 2020.

[4] 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.
ACM Journal on Emerging Technologies in Computing Systems, 9:1–22, 2013.

[5] Joel Coburn, Adrian M. Caulfield, Ameen Akel, Laura M. Grupp, Rajesh K. Gupta, Ranjit Jhala, and Steven Swanson.
Nv-heaps: Making persistent objects fast and safe with next-generation, non-volatile memories.
In *Proceedings of the 2011 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '11)*, pages 105–118, 2011.

[6] Christian Eichler, Henriette Hofmeier, Stefan Reif, Timo Hönig, and Wolfgang Schröder-Preikschat.
Neverlast: Towards the design and implementation of the nvm-based everlasting operating system.
In *Proceedings of the 2021 Annual Hawaii International Conference on System Science (HICSS '21)*, 2021.
To appear.

References (2)

[7] Marcel Köppen, Jana Traue, Christoph Borchert, Jörg Nolte, and Olaf Spinczyk.
Cache-Line Transactions: Building Blocks for Persistent Kernel Data Structures Enabled by AspectC++.
In *Proceedings of the 2019 Workshop on Programming Languages and Operating Systems (PLOS '19)*, pages 38–44, 2019.

[8] 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.

[9] 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.

[10] 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 *Proceedings of the 2015 IEEE International Symposium on High Performance Computer Architecture (HPCA '15)*, pages 526–537, 2015.

[11] 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.

[12] 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.

References (3)

[13] 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 2013 ACM SIGOPS Conference on Timely Results in Operating Systems (TRIOS '13)*, pages 1–17, 2013.

[14] Dushyanth Narayanan and Orion Hodson.
Whole-system persistence.
In *Proceedings of the 2012 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '12)*, pages 401–410, 2012.

[15] Benjamin Ransford and Brandon Lucia.
Nonvolatile memory is a broken time machine.
In *Proceedings of the 2014 Workshop on Memory Systems Performance and Correctness (MSPC '14)*, page 5, 2014.

[16] Benjamin Ransford, Jacob Sorber, and Kevin Fu.
Mementos: system support for long-running computation on RFID-scale devices.
In *Proceedings of the 2011 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '11)*, page 159, 2011.

[17] 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 2011 IEEE International Symposium on High Performance Computer Architecture (HPCA '11)*, pages 50–61, 2011.

References (4)

[18] 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 2011 Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '11)*, pages 329–338, 2011.

[19] Amoghavarsha Suresh, Pietro Cicotti, and Laura Carrington.
Evaluation of emerging memory technologies for hpc, data intensive applications.
In *Proceedings of the 2014 IEEE International Conference on Cluster Computing (CLUSTER '14)*, pages 239–247, 2014.

[20] Haris Volos, Andres Jaan Tack, and Michael M. Swift.
Mnemosyne: Lightweight persistent memory.
In *Proceedings of the 2014 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '14)*, pages 91–104, 2011.

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

[22] 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.

[23] Fei Xia, De-Jun Jiang, Jin Xiong, and Ning-Hui Sun.
A survey of phase change memory systems.
Journal of Computer Science and Technology, 30(1):121–144, 2015.

References (5)

[24] 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 2013 Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '13)*, pages 421–432, 2013.

[25] 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 2009 Annual International Symposium on Computer Architecture (ISCA '09)*, pages 14–23, 2009.

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