A Scalable and Persistent Key-Value Store Using Non-Volatile Memory

연세대학교 컴퓨터과학과 김도영
2019년 06월

과제명: IoT 환경을 위한 고성능 플래시 메모리 스토리지 기반 인메모리 분산 DBMS 연구개발

과제번호: 2017-0-00477

과학기술정보통신부
Ministry of Science and ICT

연세대학교
YONSEI UNIVERSITY

IITP 정보통신기술진흥센터
Institute for Information & communications Technology Promotion

Contents

1. Introduction

- In-Memory Key-Value Store
- Non-Volatile Random-Access Memory (NVRAM)
- Redis – Logging Scheme
- Redis – Intel PMDK Proposed Method (NVRAM)

In-Memory Key-Value Store

ex) **Redis**, Memcached

Hold all key-value data in the DRAM

Provide logging schemes to prevent data loss

Byte Addressable

Non-Volatile Random-Access Memory (NVRAM)

ex) **SST-MRAM**, ReRAM, PCM(3DXPoint)

High performance comparable to DRAM

Data Persistence

High cost per capacity (Not Commercialized)

Byte Addressable

NVRAM

Technology	FeRAM	MRAM	ReRAM	PCM	DRAM	NAND Flash
Nonvolatile	Yes	Yes	Yes	Yes	No	Yes
Endurance	10^{12}	10^{12}	10^{6}	10^{8}	10^{15}	10^{3}
Write Time	100ns	~10ns	~50ns	~75ns	10ns	10µs
Read Time	70ns	10ns	10ns	20ns	10ns	25µs
Power Consumption	Low	Medium/Low	Low	Medium	Very High	Very High
Cell Size (f^2)	15-20	6-12	6-12	1-4	6-10	4
Cost ($/Gb)	$10/Gb	$30-70/Gb	Currently High	$0.16/Gb	$0.6/Gb	$0.03/Gb

SNIA persistent memory summit (1/24/2018)
Analysts Weigh In On Persistent Memory
https://www.snia.org/sites/default/files/PM-Summit/2018/presentations/14_PM_Summit_18_Analysts_Session_Oros_Final_Post_UPDATED_R2.pdf

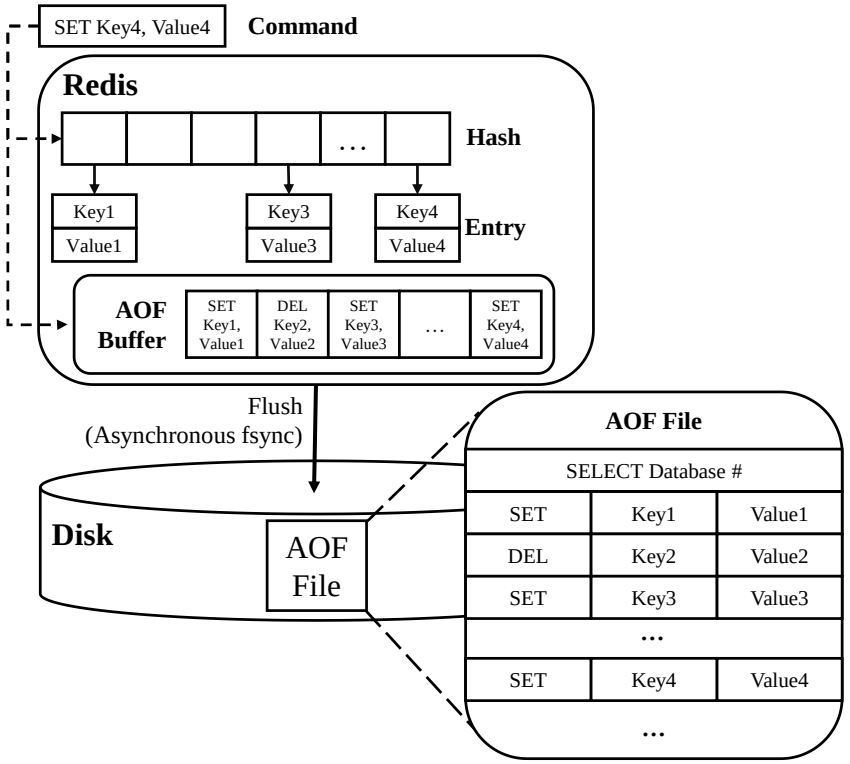

Redis – Logging Scheme
(Append-Only File)

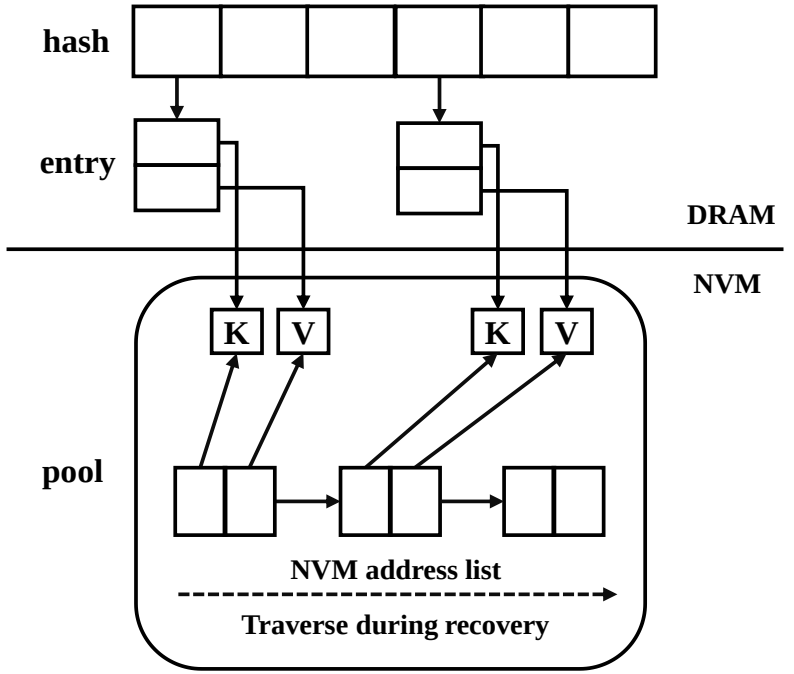

Redis – Intel PMDK Proposed
(Full NVRAM)

	AOF-Always	AOF-Everysec	PMDK (NVRAM)
Performance	Slow	Fast	Fast
Persistence	Strong	Weak	Strong
NVRAM Usage	N/A	N/A	High

2. NVRAM-DRAM Hybrid Redis

- Overview Structure
- Reverse-Tiered Structure
  - 1 Tier NVRAM, 2 Tier DRAM
  - Hot Data, Cold Data
  - Thread-safe LRU-clock Priority Queue
- Persistent Tiering

2 NVRAM-DRAM Hybrid Redis

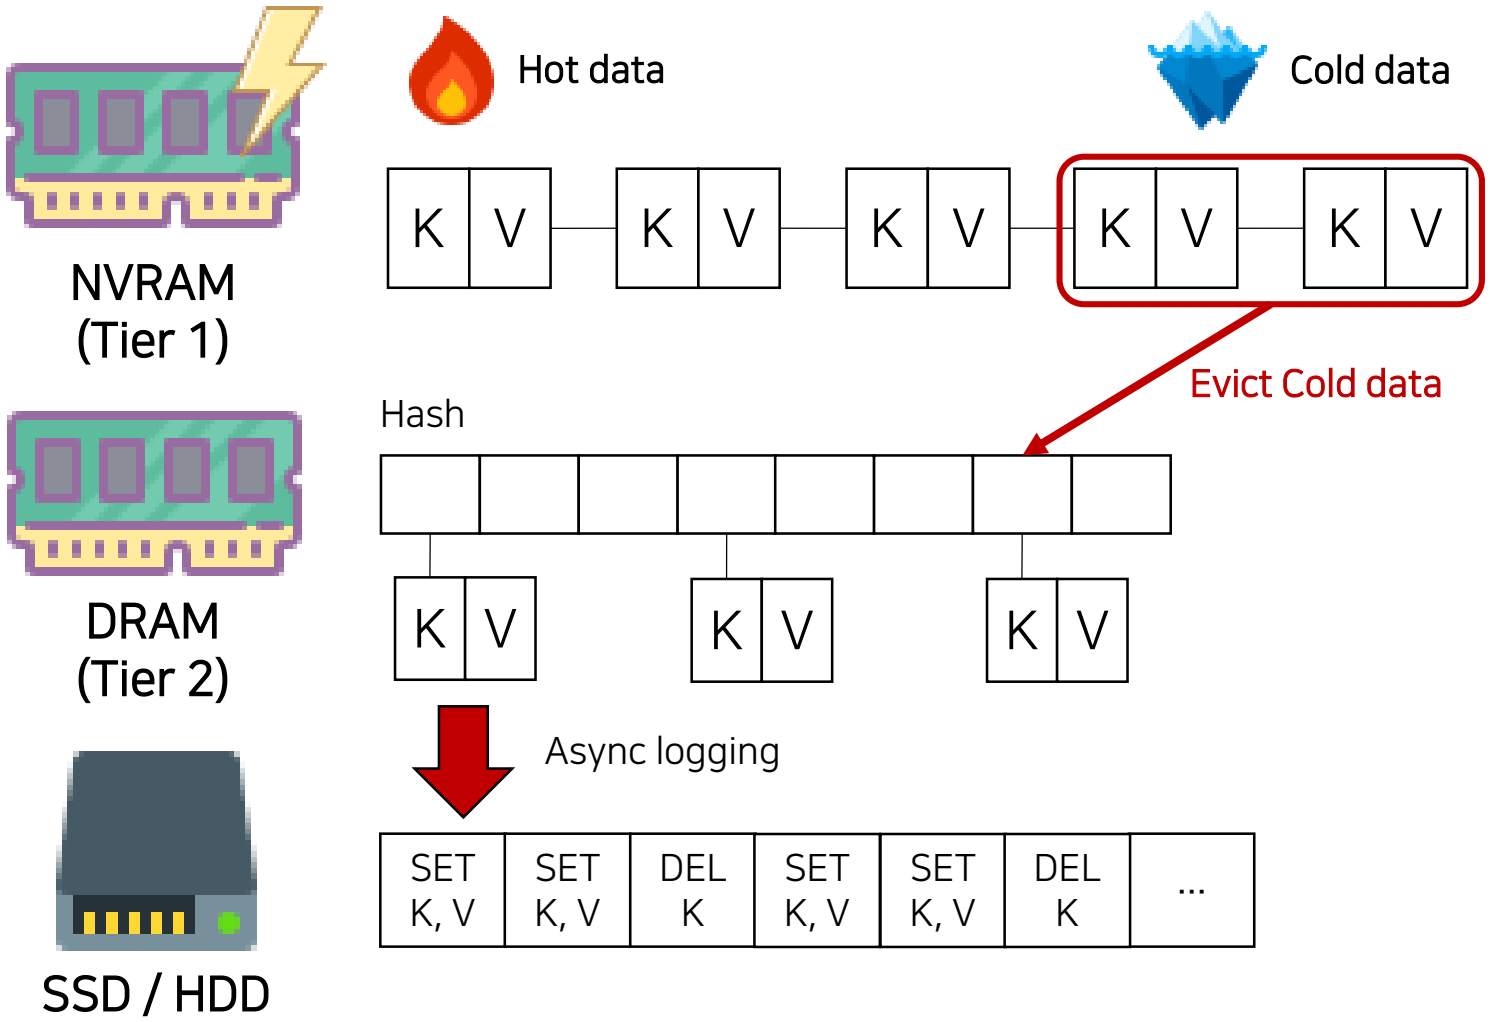

Abstract structure of our proposed method

	Hybrid Redis
Performance	Fast
Persistence	Strong
NVRAM Usage	Low

- In-Memory DB의 취약점 보안
  - Weak Persistency
  - Performance downgraded by logging
- NVRAM의 현실적 한계 보안
  - High cost per capacity

NVRAM-DRAM Hybrid Redis

- ✓ Hot data? Cold data?
  - 최근에 생성, 갱신된 데이터
    → Hot data
  - 생성, 갱신된 지 오래된 데이터
    → Cold data

- ✓ Hot data, Cold data로 나누는 이유
  - DRAM
    - 데이터 Update가 많은 Log를 발생
    - Cold data 저장 용이
  - NVRAM
    - 데이터 Update시 Log 필요 X
    - Hot data 저장 용이

- ✓ Hot, Cold 데이터로 나누는 방법
  - Priority Queue (LRU-Clock based)
  - Thread-Safe 해야함
    - Background thread에서 Eviction 진행
    - 안전한 데이터 갱신 필요

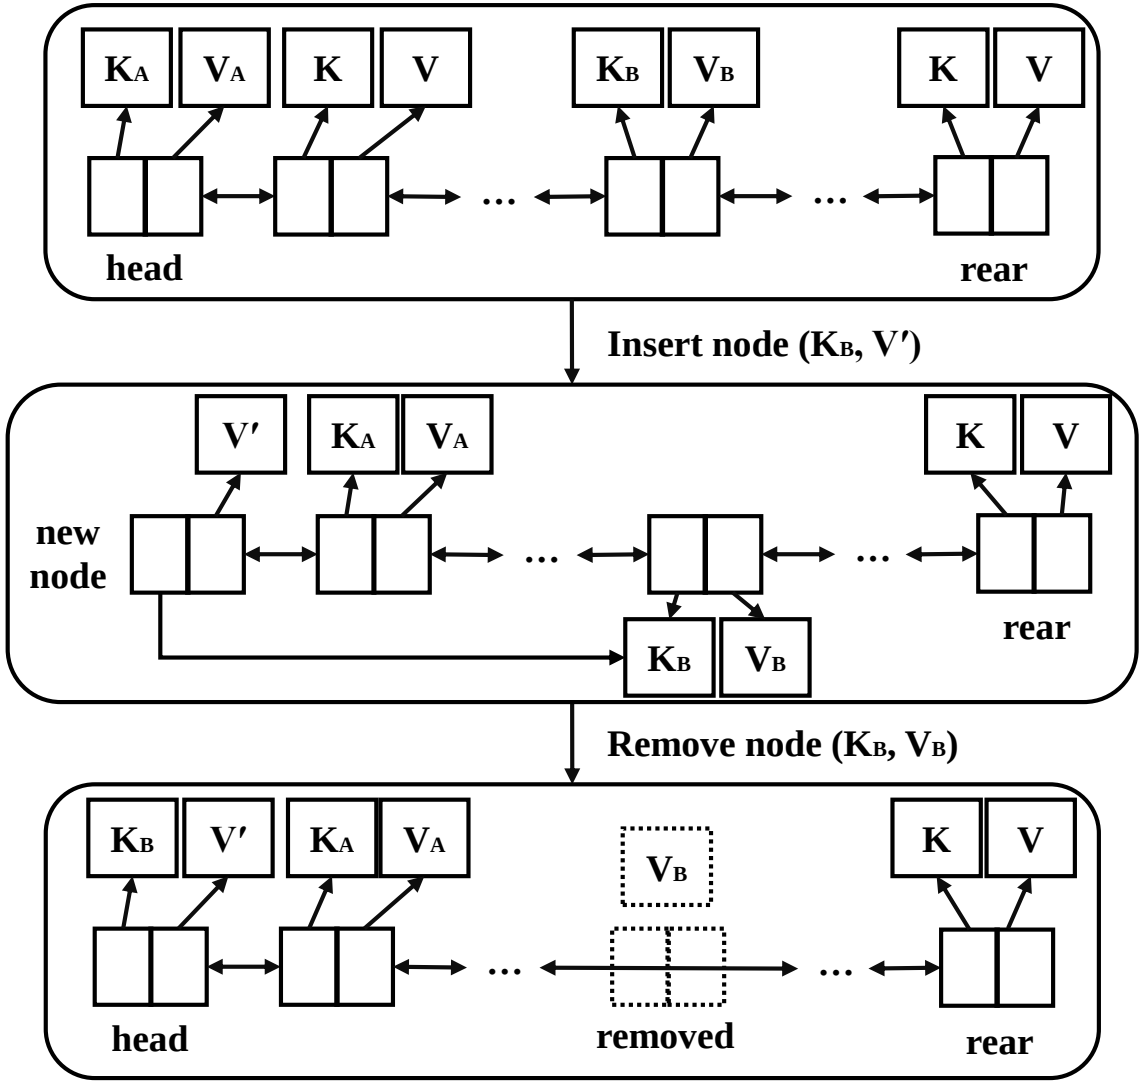

Thread-safe Priority Queue

NVRAM-DRAM Hybrid Redis

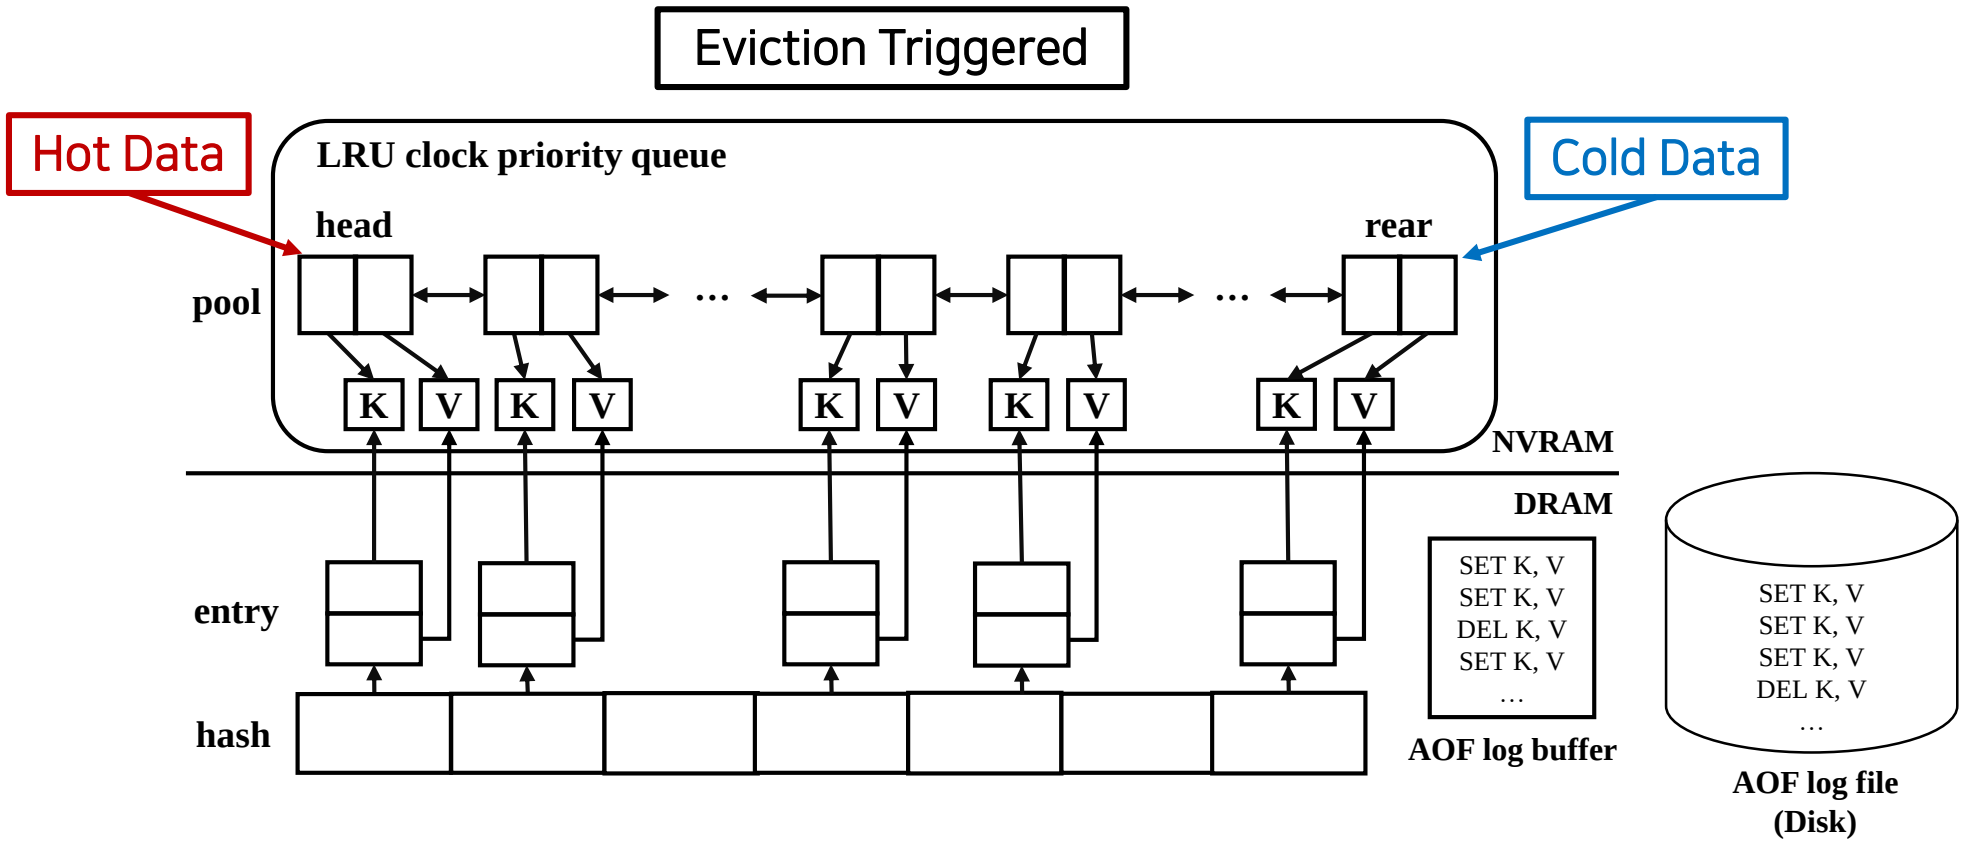

NVRAM-DRAM Hybrid Redis

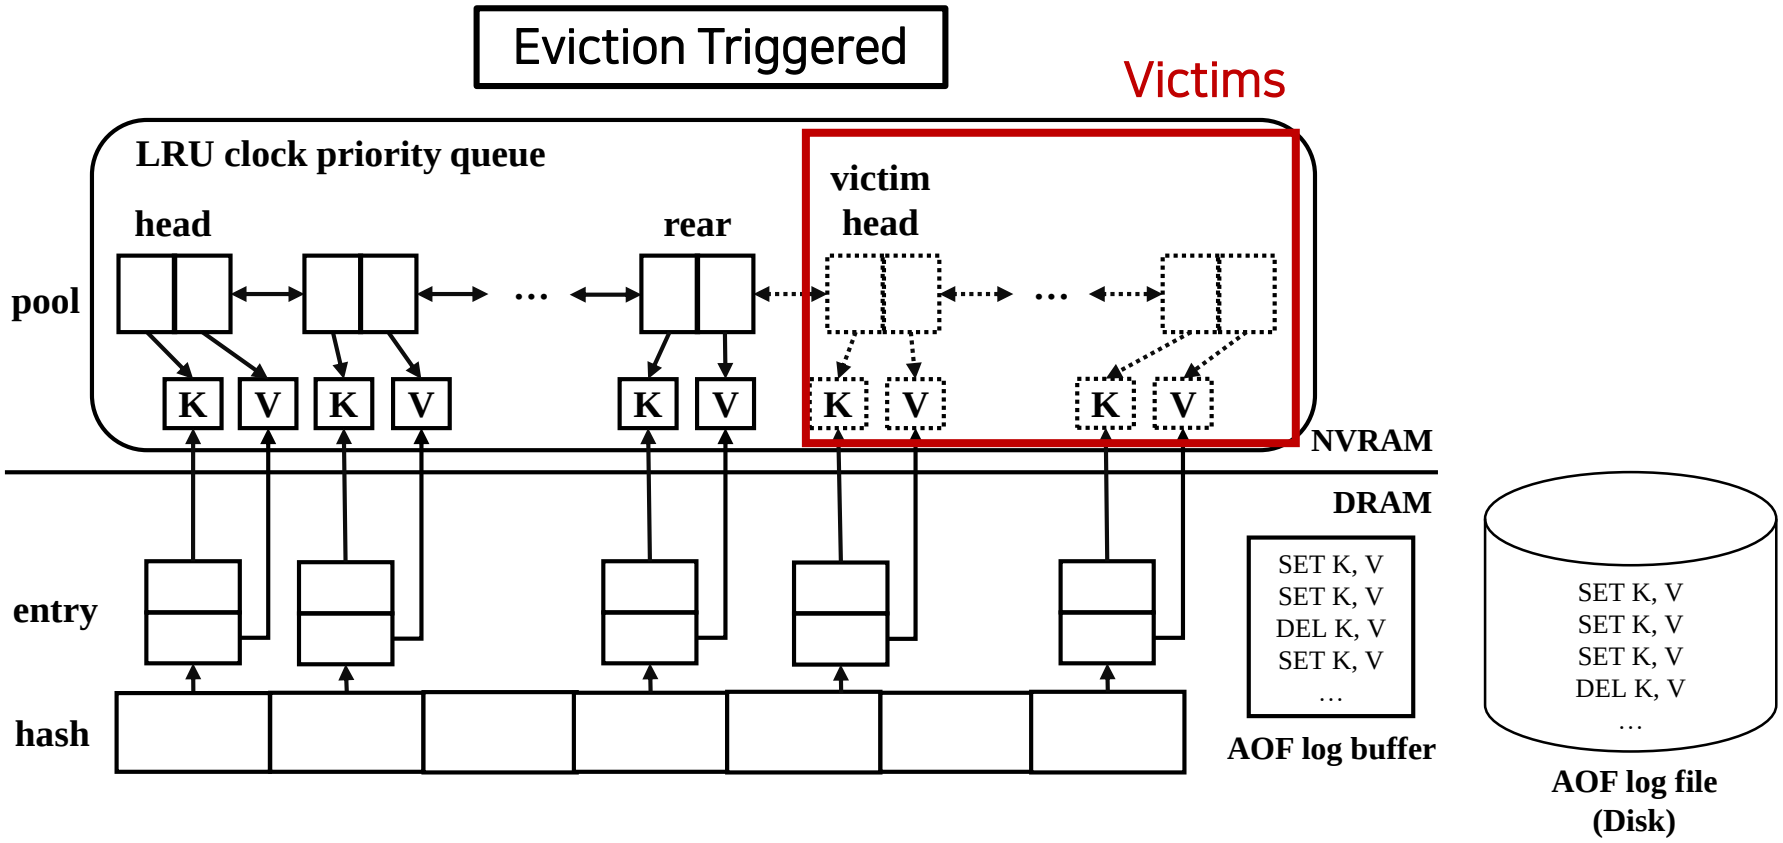

NVRAM-DRAM Hybrid Redis

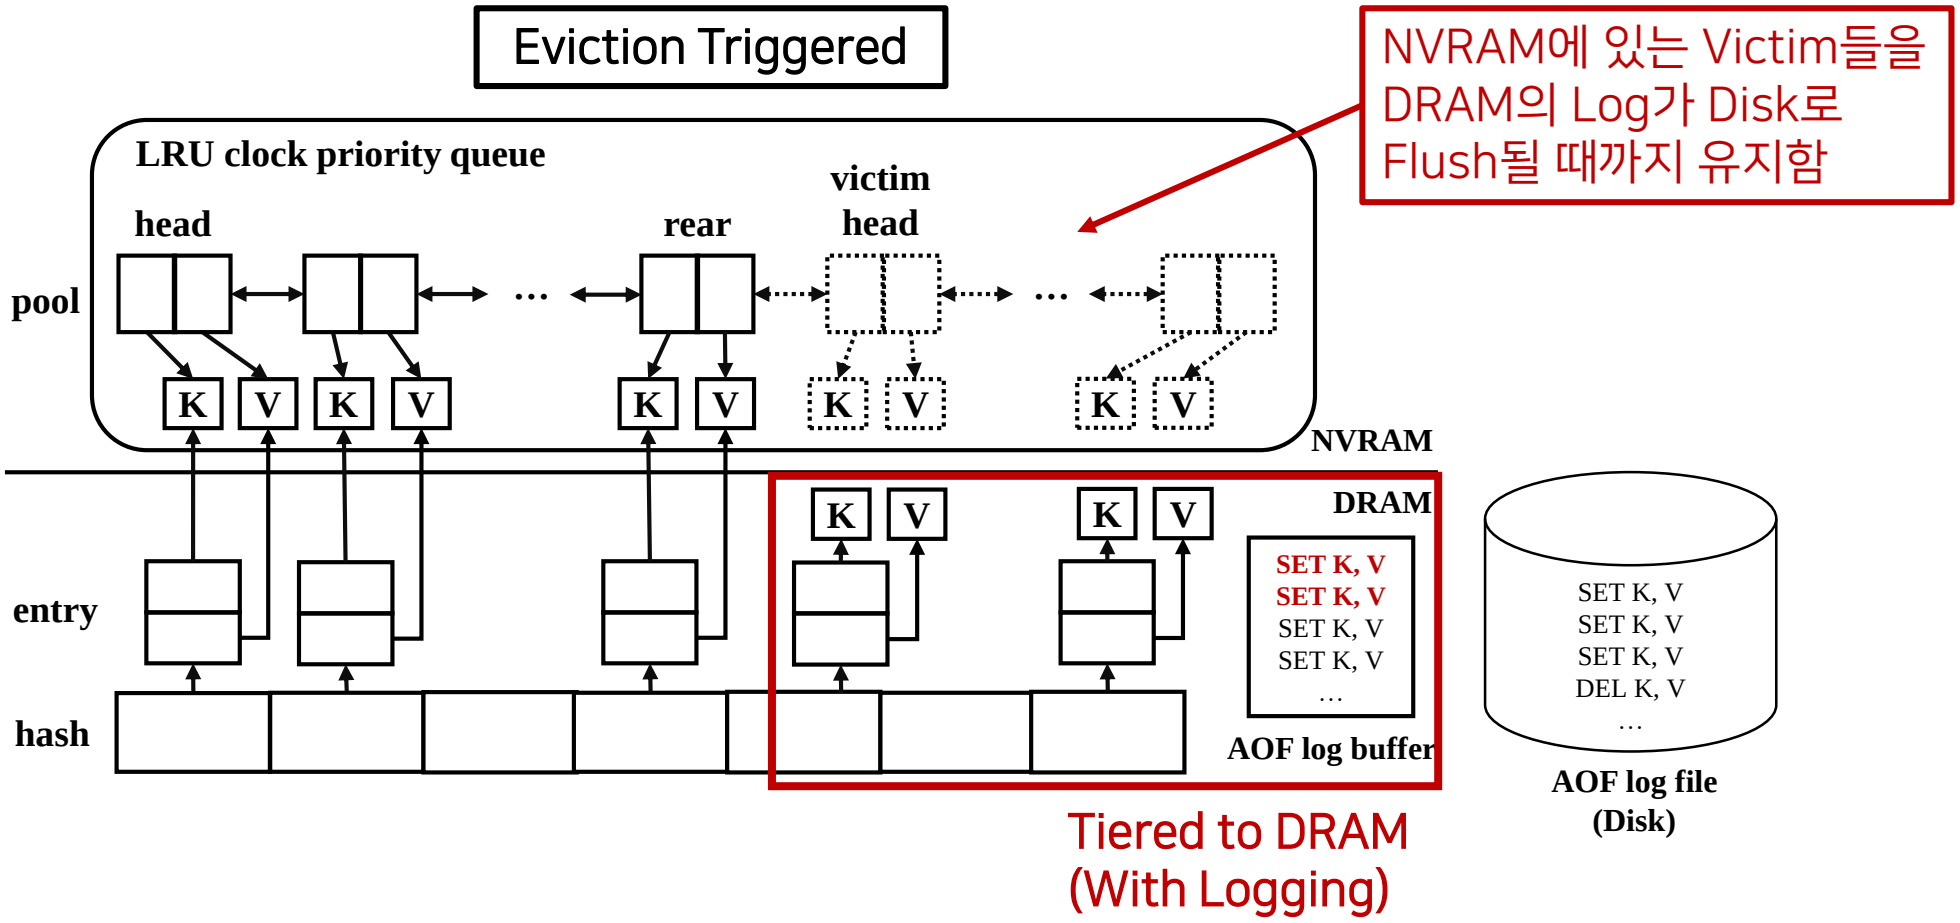

NVRAM-DRAM Hybrid Redis

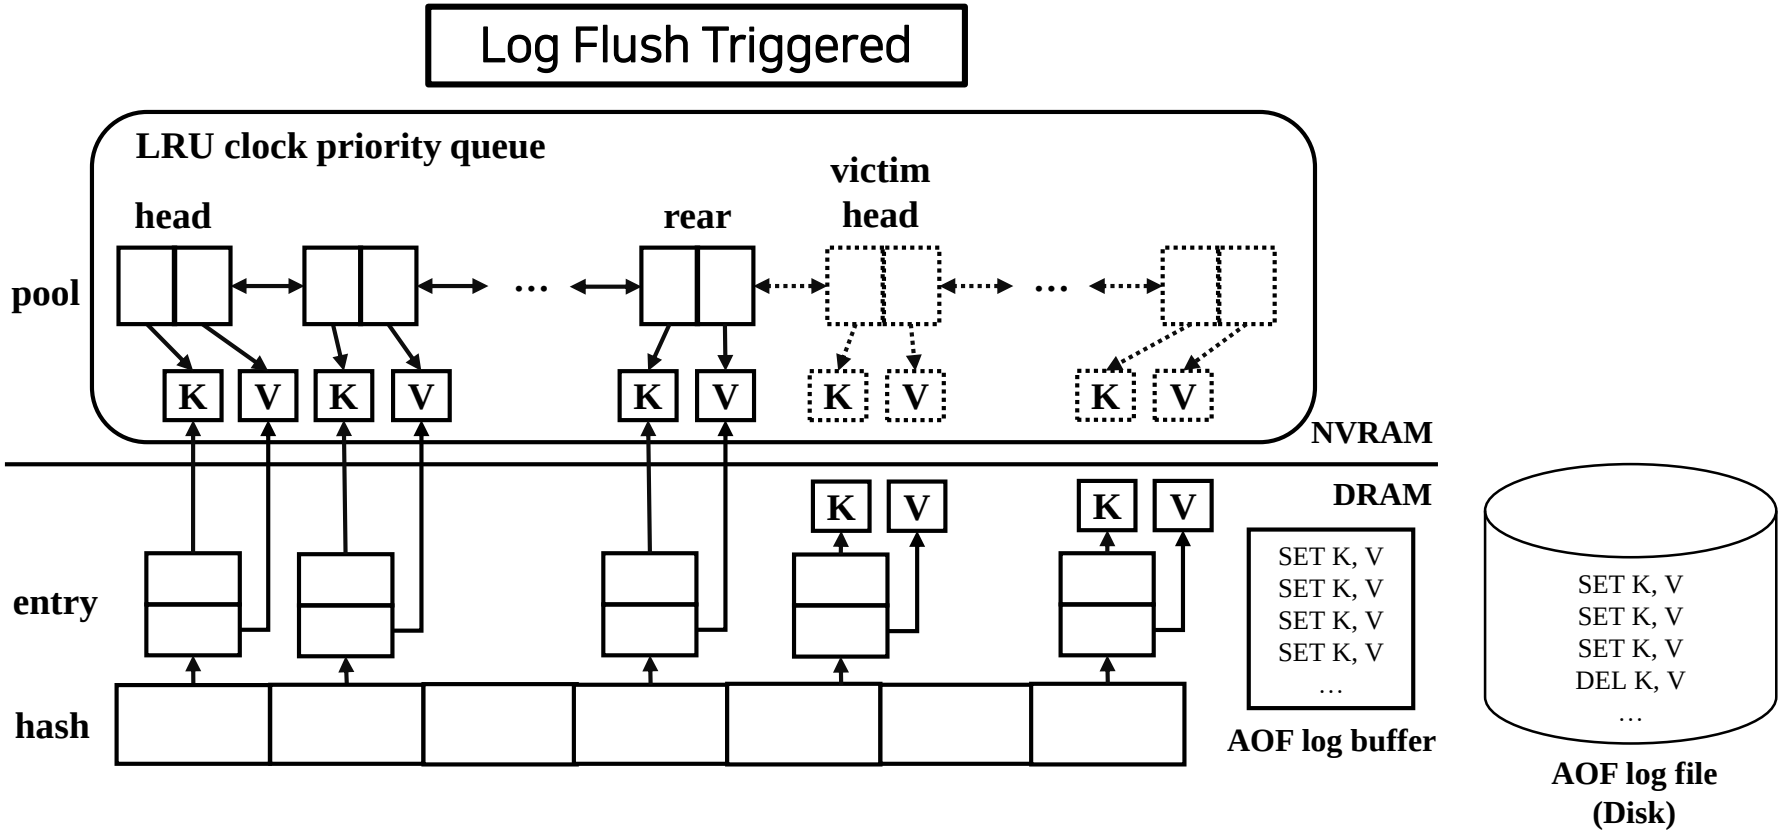

NVRAM-DRAM Hybrid Redis

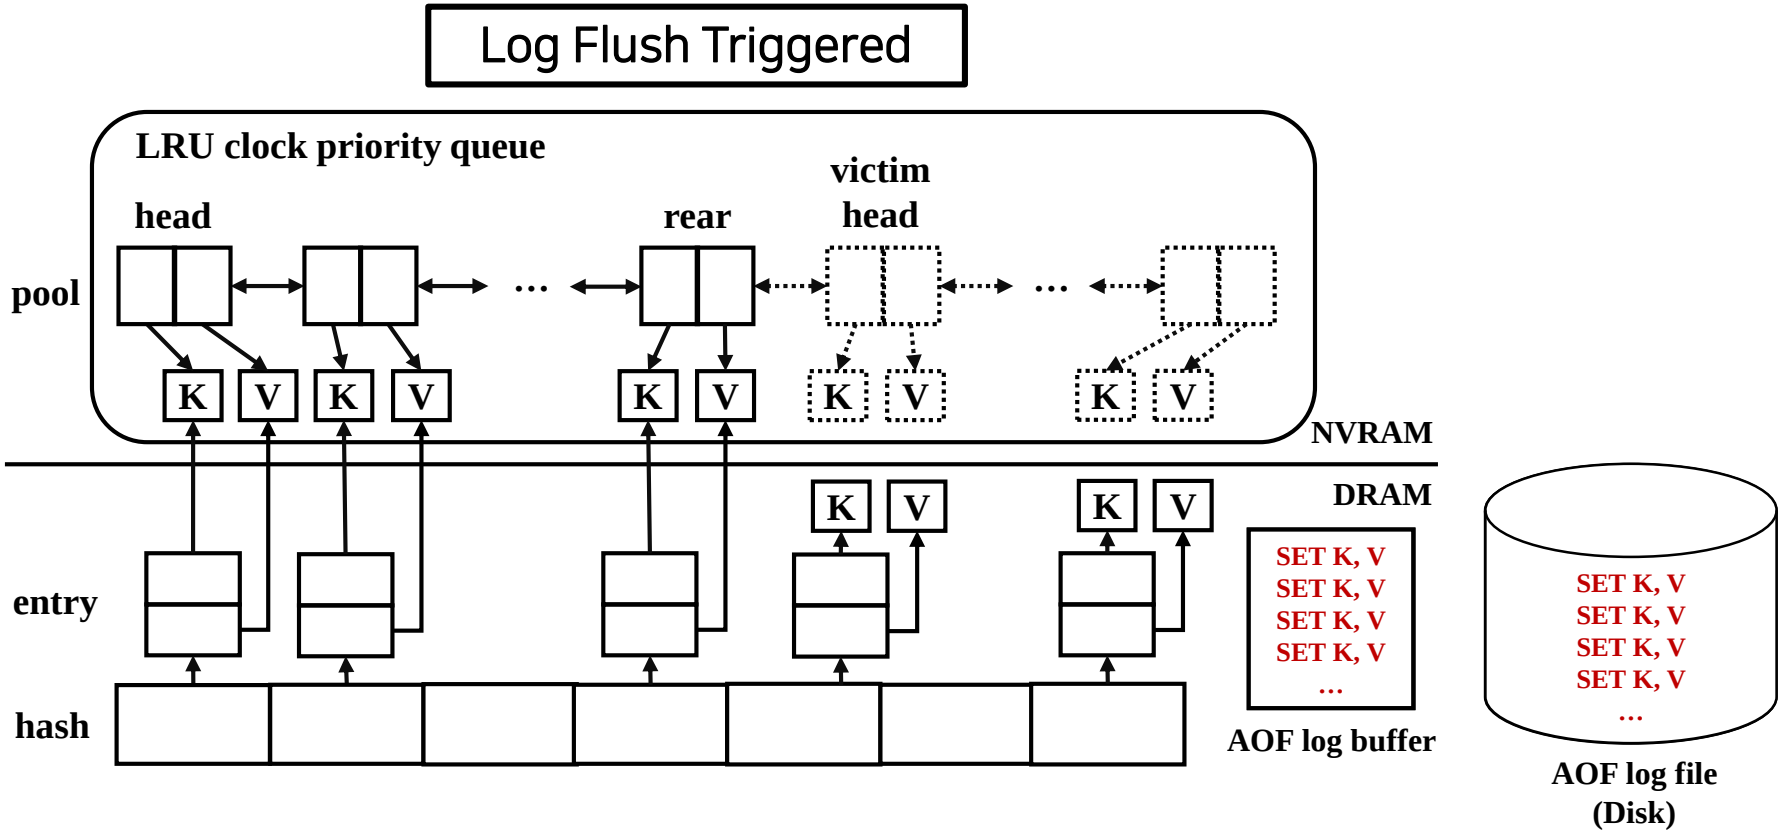

NVRAM-DRAM Hybrid Redis

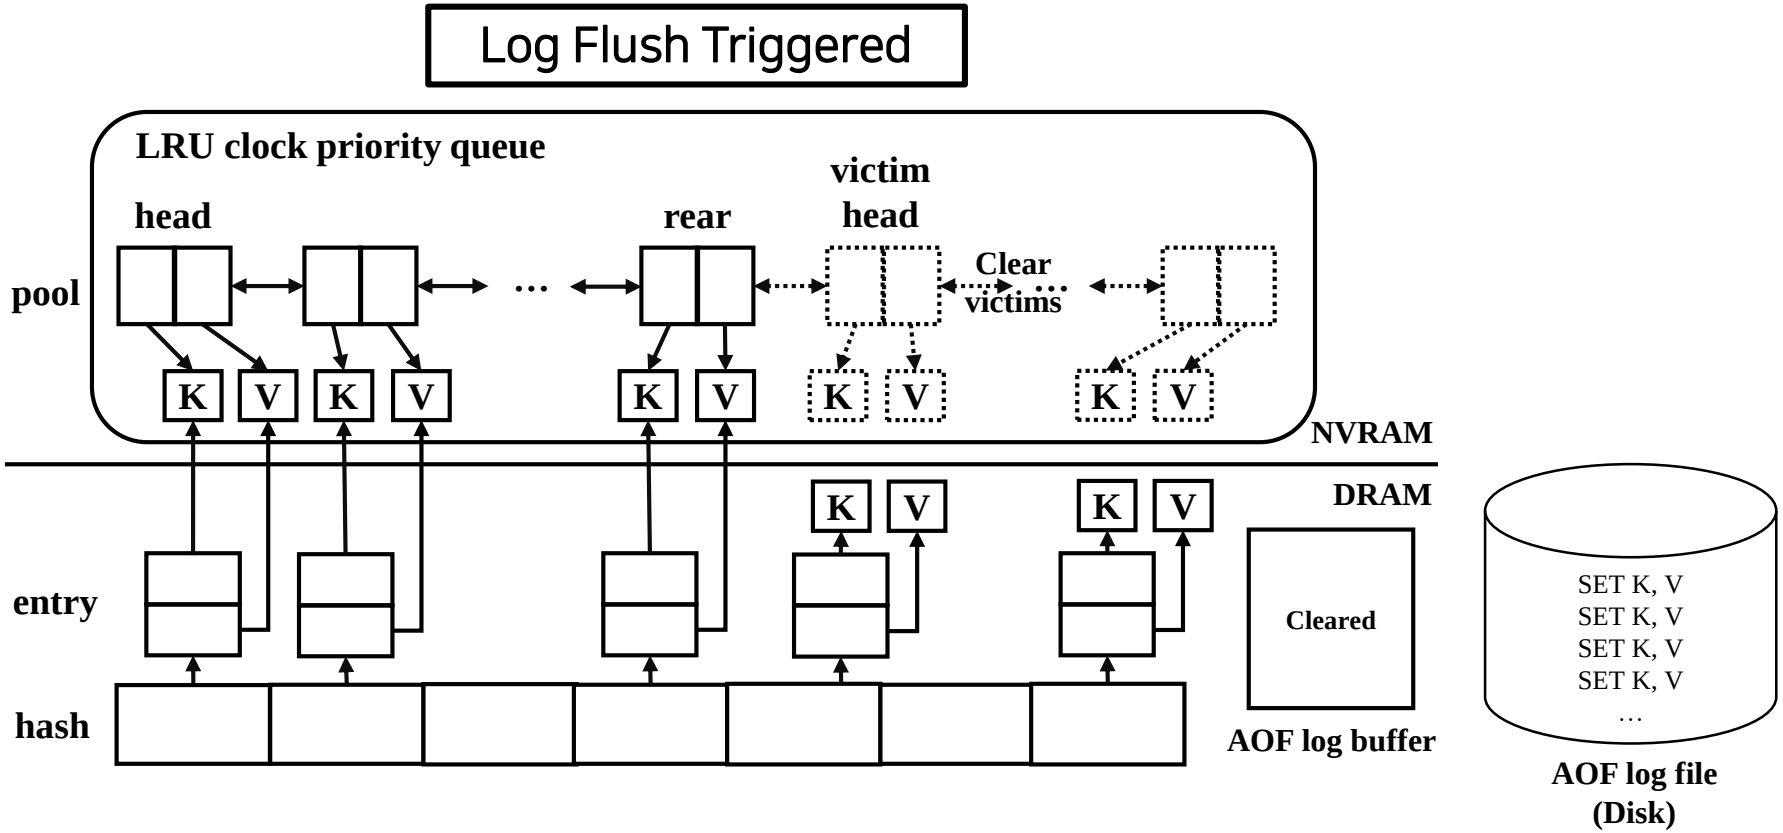

3. Evaluation

- NVRAM Latency
- NVRAM Usage test
- Memtier benchmark test

3 Evaluation

NVRAM Latency

	Read (ns)	Write (ns)
DRAM	50	50
MRAM	50	50
PCM	50	500
3D XPoint (Memory Mapped)	100	100
3D XPoint (Storage Mapped)	200	200

NVRAM 사용량 평가

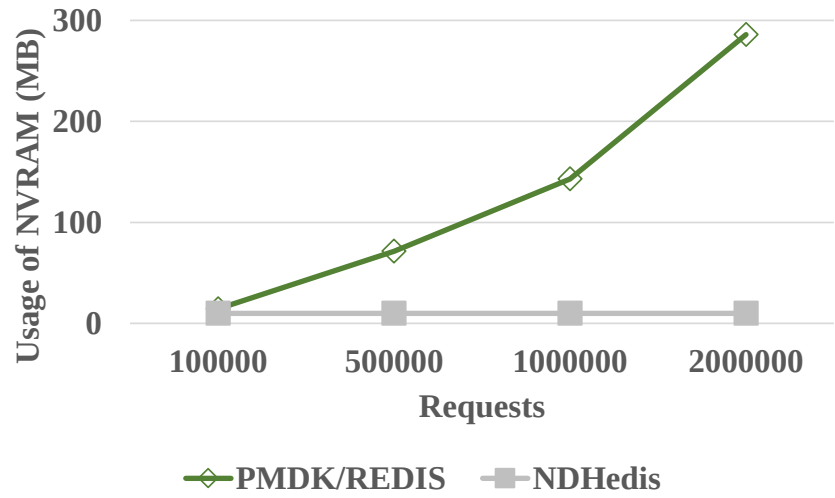

Memtier-benchmark 성능 평가

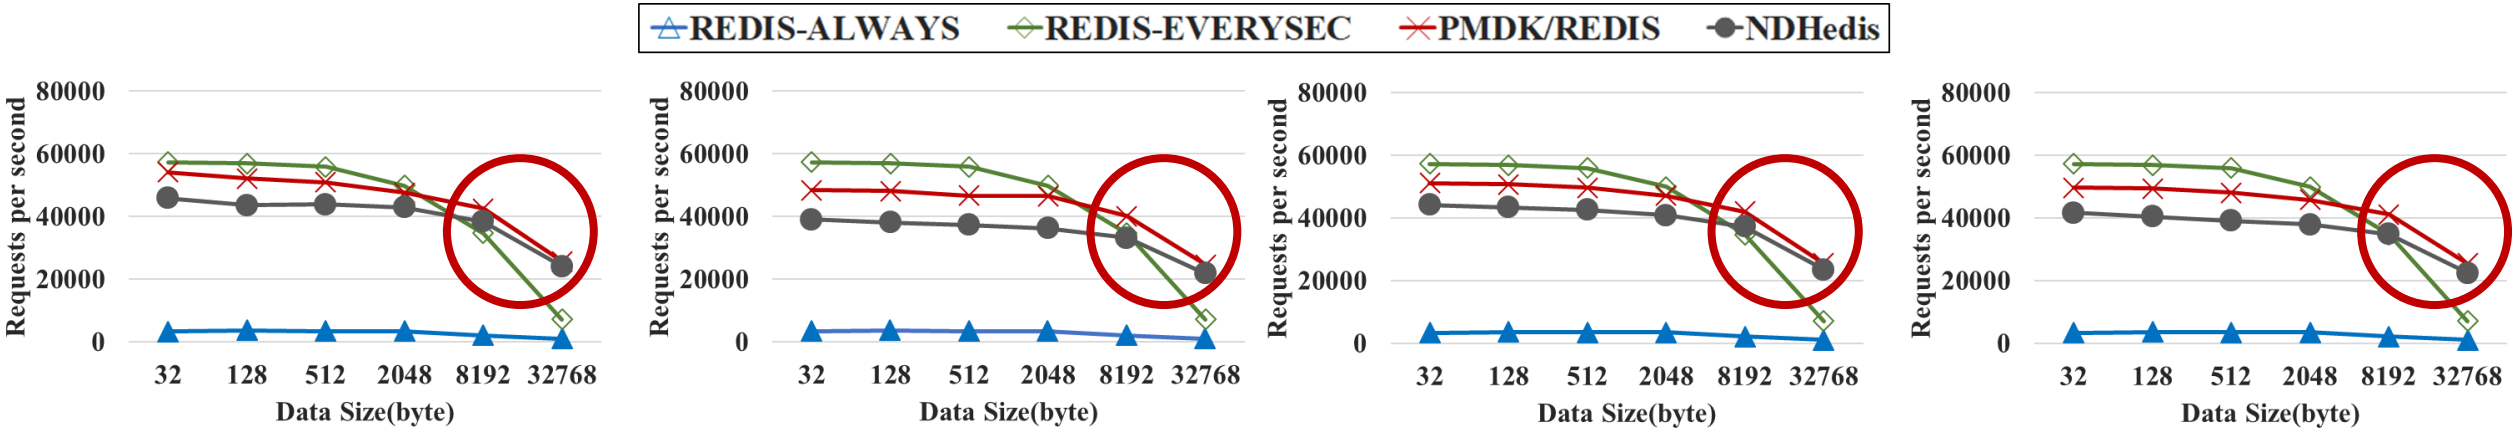

(a) DRAM, MRAM (b) PCM (c) 3D XPoint (Memory Mapped) (d) 3D Xpoint (Storage Mapped)

4 Summary

- NVRAM-DRAM Hybrid System
  - NVRAM의 Cost per Capacity 한계 극복
  - DRAM의 Volatile 한계 극복, Log 의존도 개선
- Fast Performance, Data persistence, Low NVRAM Usage
- Cost per Capacity가 높지만, 성능이 뛰어난 NVRAM에 적합

5 SAC 2019 Best Poster

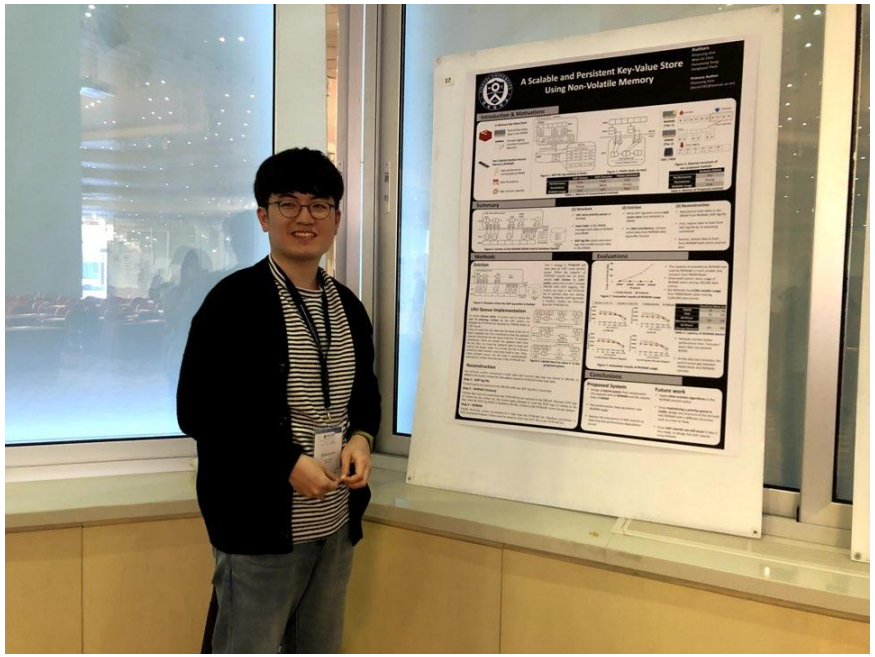

ACM DL DIGITAL LIBRARY Yonsei University

A scalable and persistent key-value store using non-volatile memory

Full Text: PDF

Authors:
Doyoung Kim Yonsei University, Yonsei-ro, Seodaemun-gu, Seoul, Republic of Korea
Won Gi Choi Yonsei University, Yonsei-ro, Seodaemun-gu, Seoul, Republic of Korea
Hanseung Sung Yonsei University, Yonsei-ro, Seodaemun-gu, Seoul, Republic of Korea
Sanghyun Park Yonsei University, Yonsei-ro, Seodaemun-gu, Seoul, Republic of Korea

2019 Article

Best Poster

Bibliometrics
· Citation Count: 0
· Downloads (cumulative): 52
· Downloads (12 Months): 52
· Downloads (6 Weeks): 52

Published in:

· Proceeding
SAC '19 Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing
Pages 464-467

Limassol, Cyprus — April 08 - 12, 2019
ACM New York, NY, USA ©2019
table of contents ISBN: 978-1-4503-5933-7 doi>10.1145/3297280.3298991

https://dl.acm.org/citation.cfm?id=3298991

Q & A