

***Freewrite:***

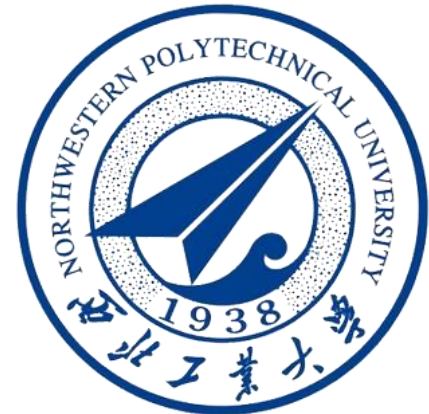
# **Creating (Almost) Zero-Cost Writes to SSD**

**Chunyi Liu, Fan Ni,  
Xingbo Wu, and *Song Jiang***



**University of Texas at  
Arlington, USA**

**Xiao Zhang**



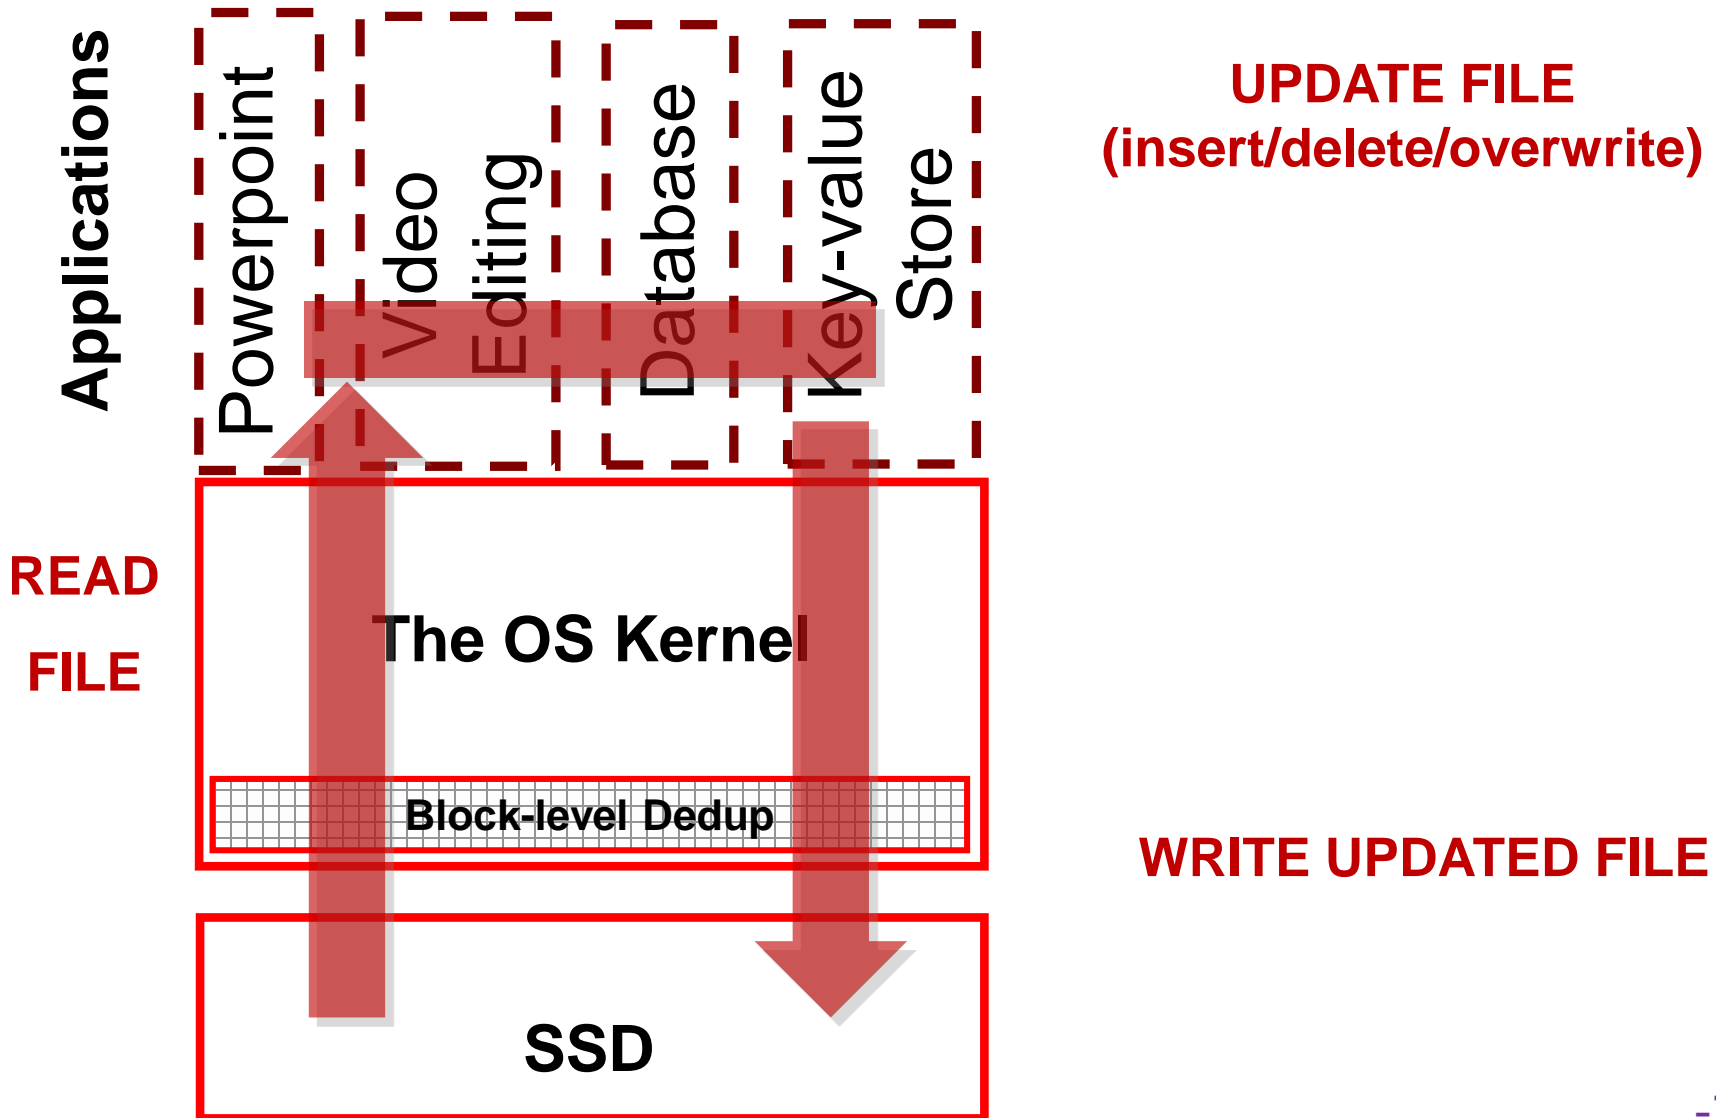
**Northwestern Polytechnical  
University, China**

# The Achilles' Heel in SSD: Writes

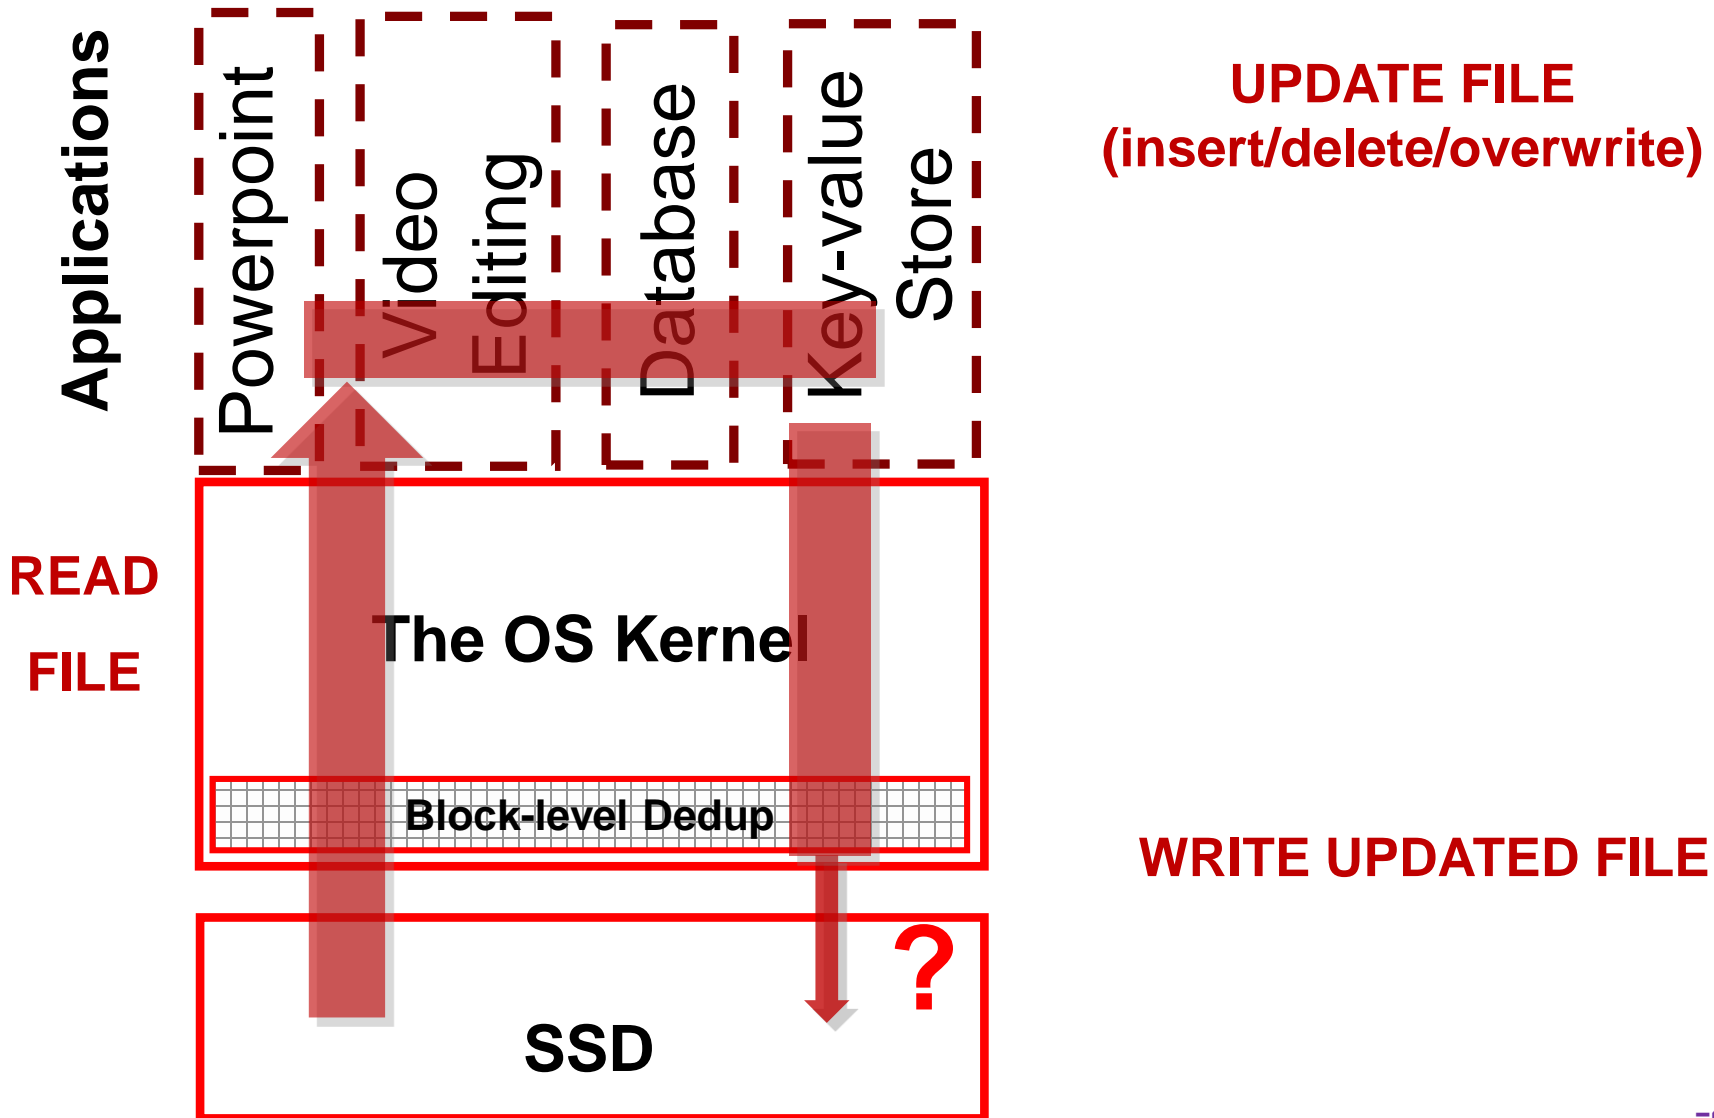
---

- SSD has high performance that is not sensitive to access sequentiality
- However, write operations are expensive and destructive.
  - Write latency can be higher than reads by many times.
  - Out-of-place writes may require mapping change, block erasure, and even page migrations.
  - Expensive garbage collections can be triggered.
  - Writes reduce its limited lifespan.
- The goal: **Remove Unnecessary Writes to the SSD**

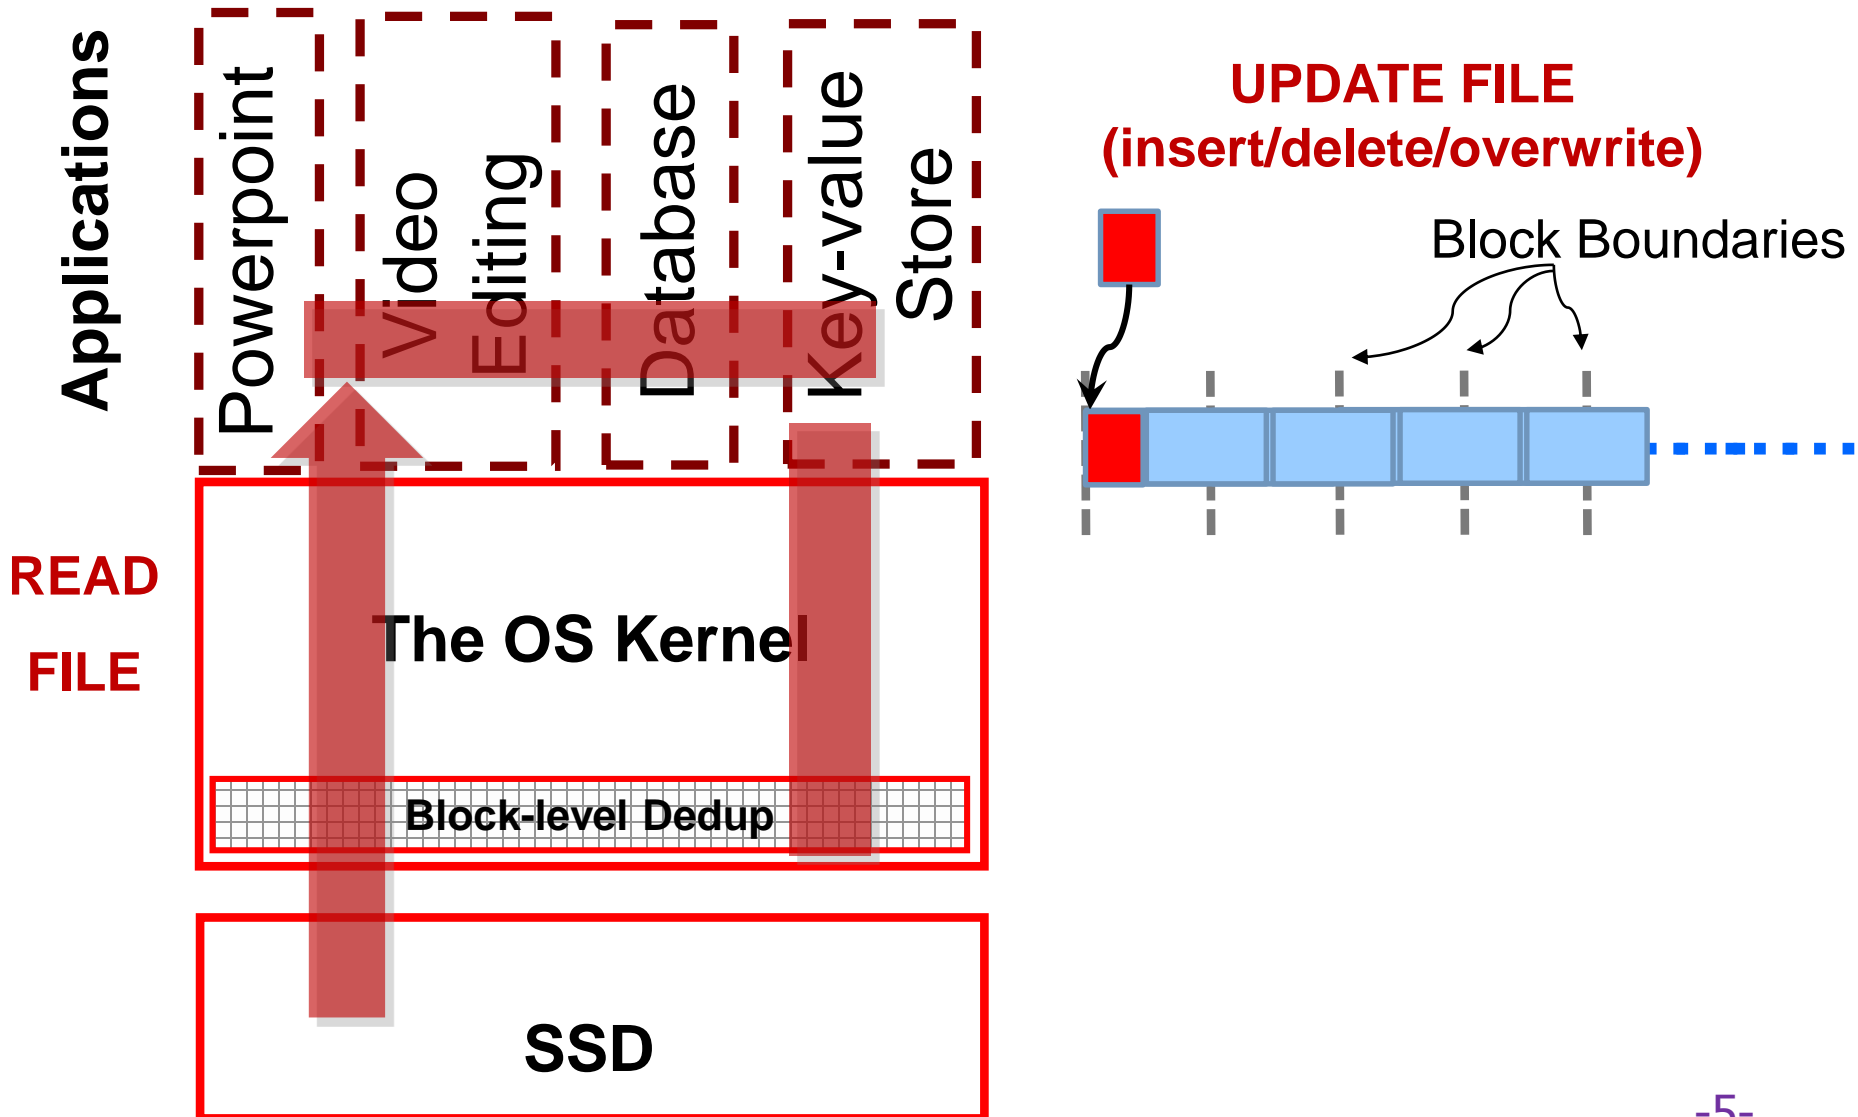
# The Opportunity and Challenge



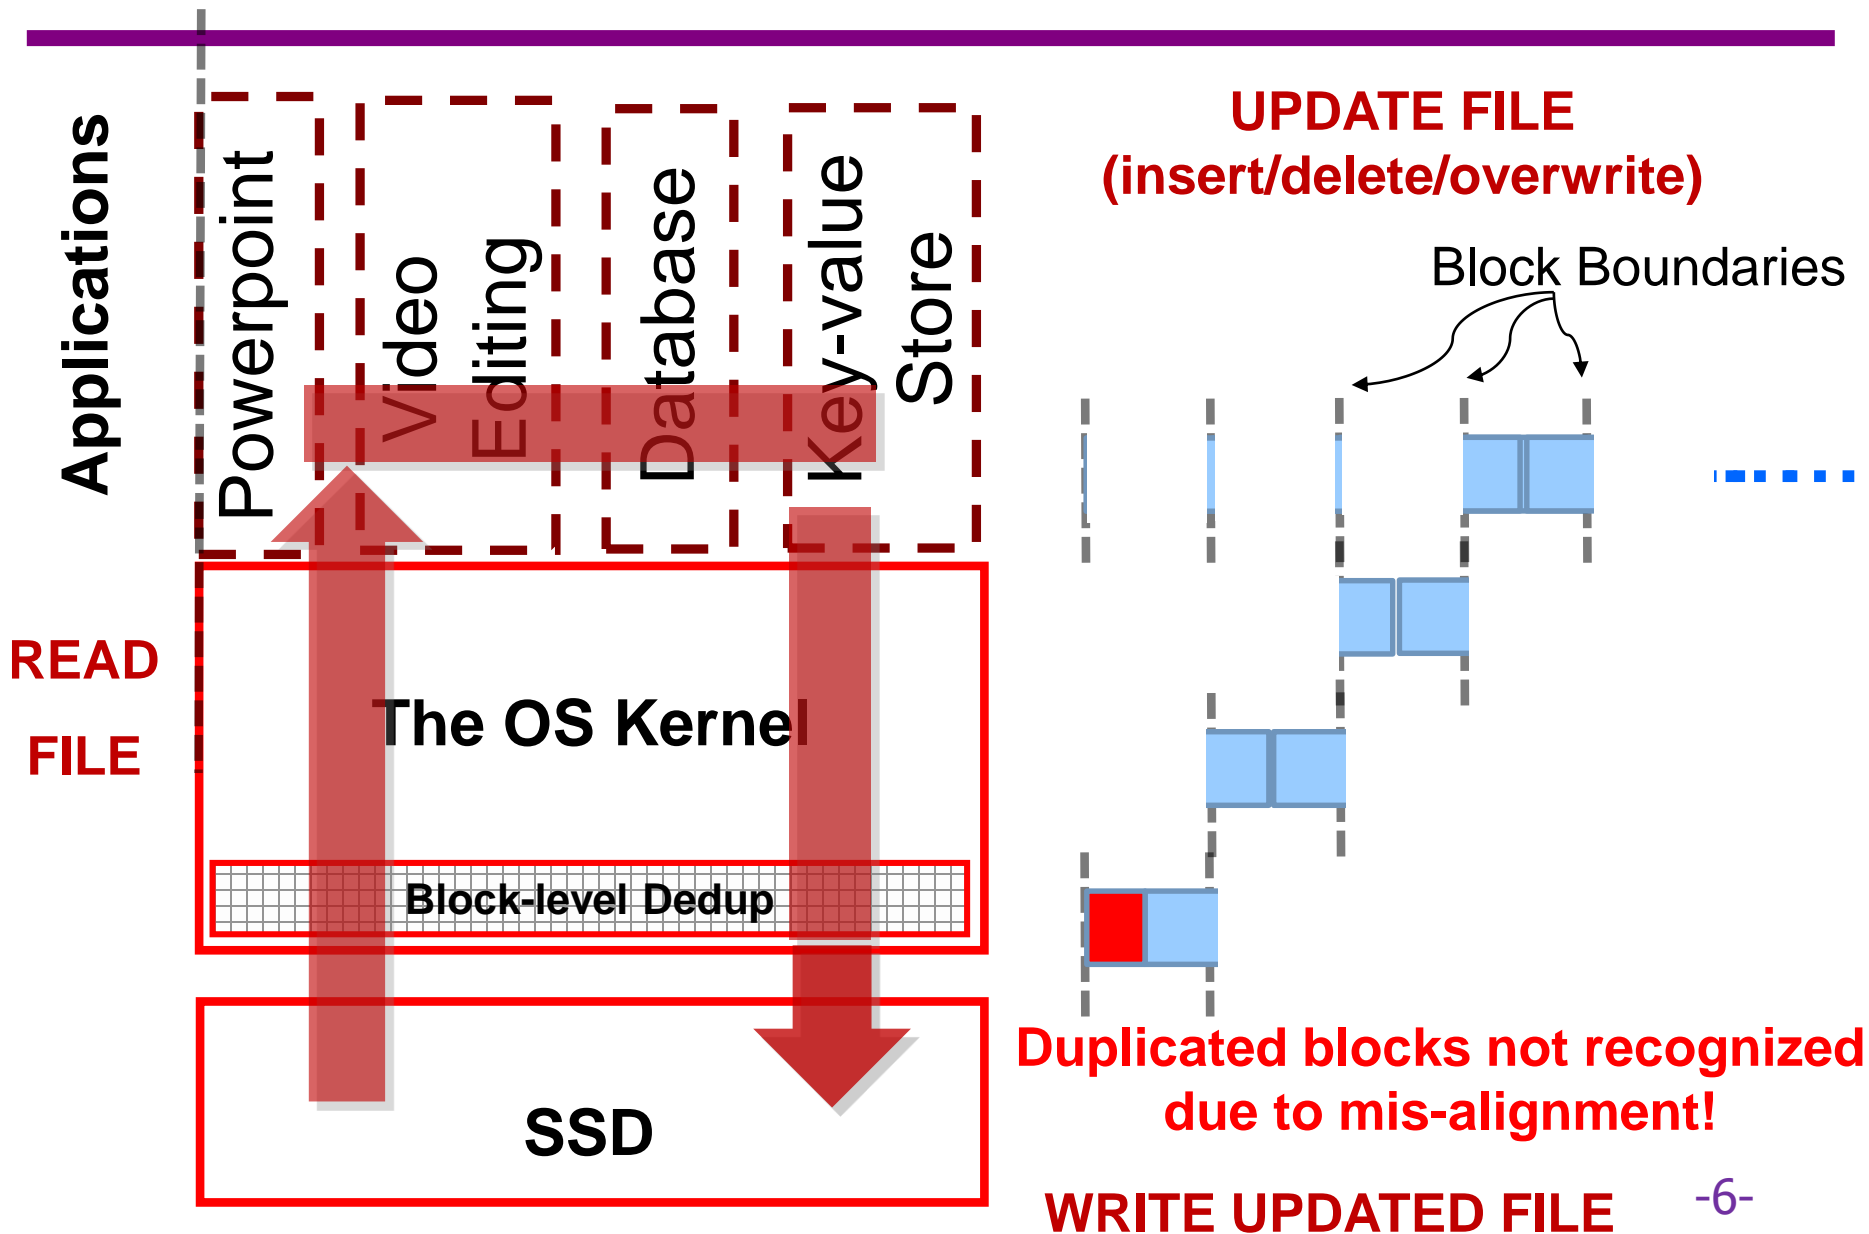
# The Opportunity and Challenge



# The Opportunities and Challenge

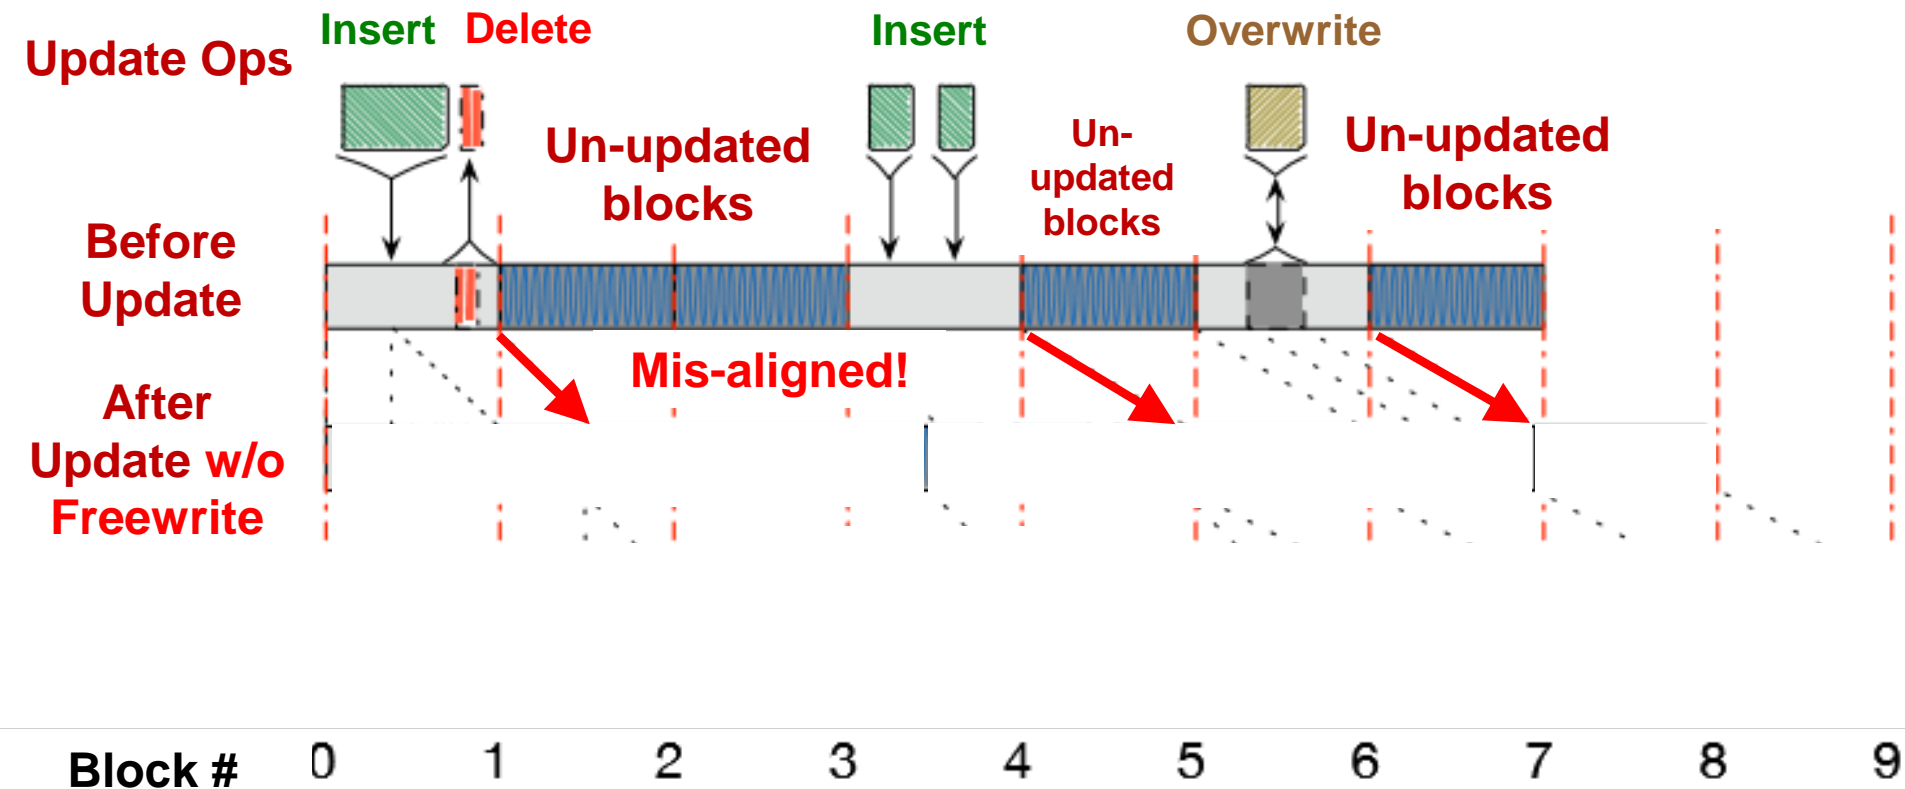


# The Opportunities and Challenge



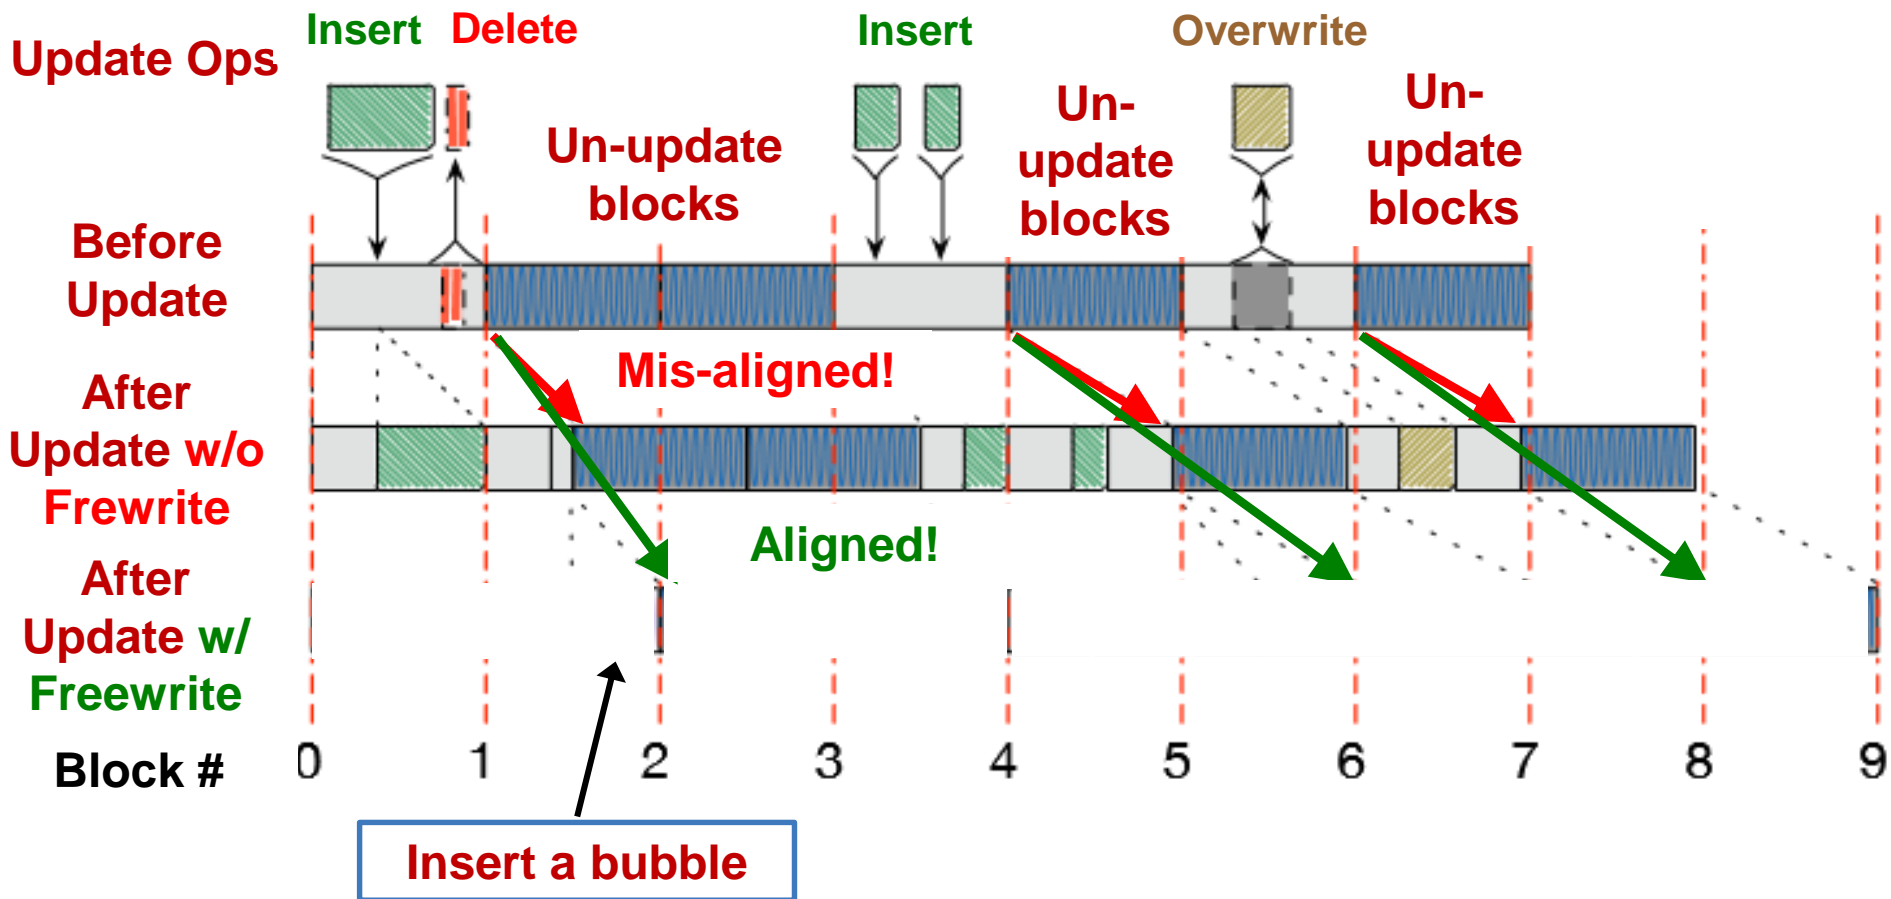
# Freewrite: Retain Alignments to Maximize Duplicates

The approach: allow small unused spaces (bubbles) in a file to keep un-updated blocks aligned.



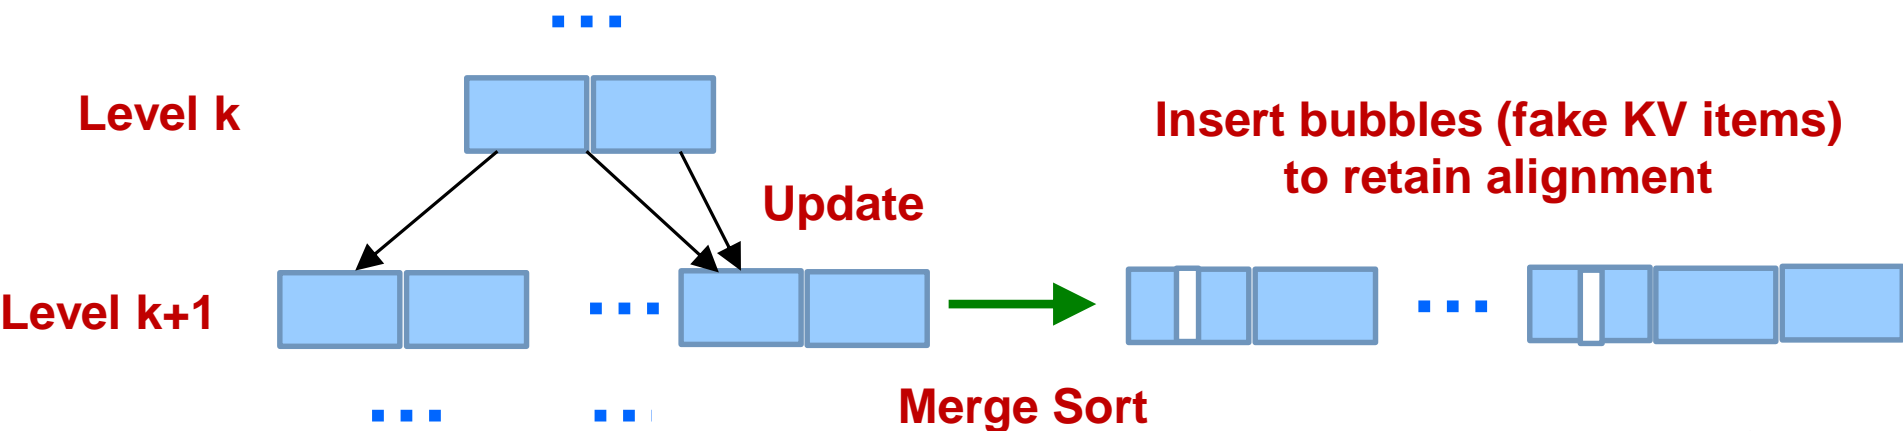
# Freewrite: Retain Alignments to Maximize Duplicates

The approach: allow small unused spaces (*bubbles*) in a file to keep un-updated blocks aligned.



# Use Case I: Compaction in LevelDB

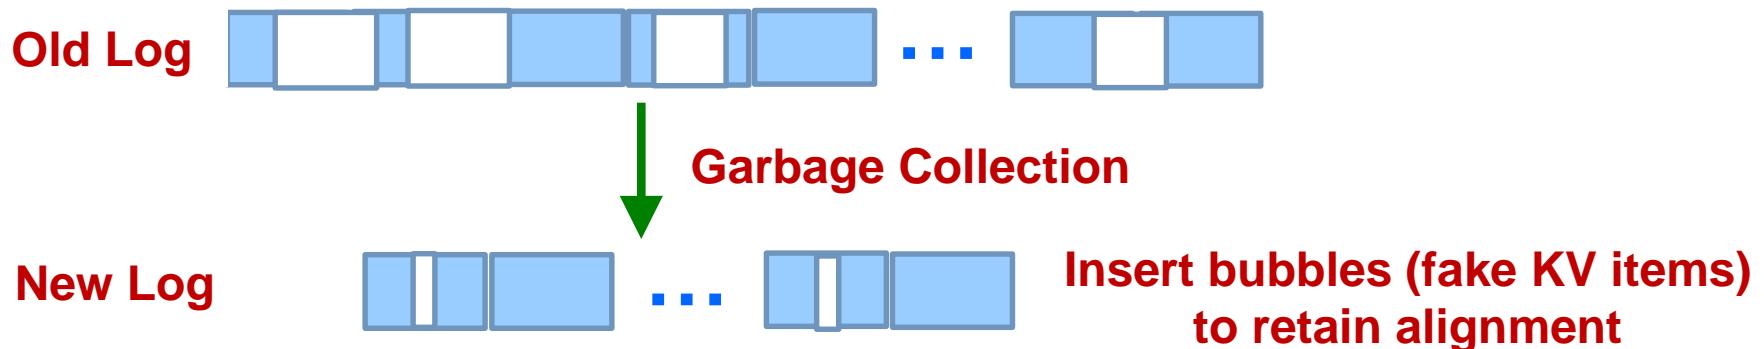
- LevelDB is a key-value store in which KV items are incrementally sorted in a multi-level structure.
- Level  $k+1$  is much larger than Level  $k$ .
- Adjacent levels are constantly **merge-sorted**.
- Freewrite inserts bubbles into the new Level  $k+1$  to reduce its writes.



# Use Case II: Garbage Collection in FAWN

---

- FAWN is a key-value store always appending KV items in a log file.
- GC is periodically performed in the log file to remove invalidated items.
- Freewrite inserts bubbles into the new log file to reduce its writes.

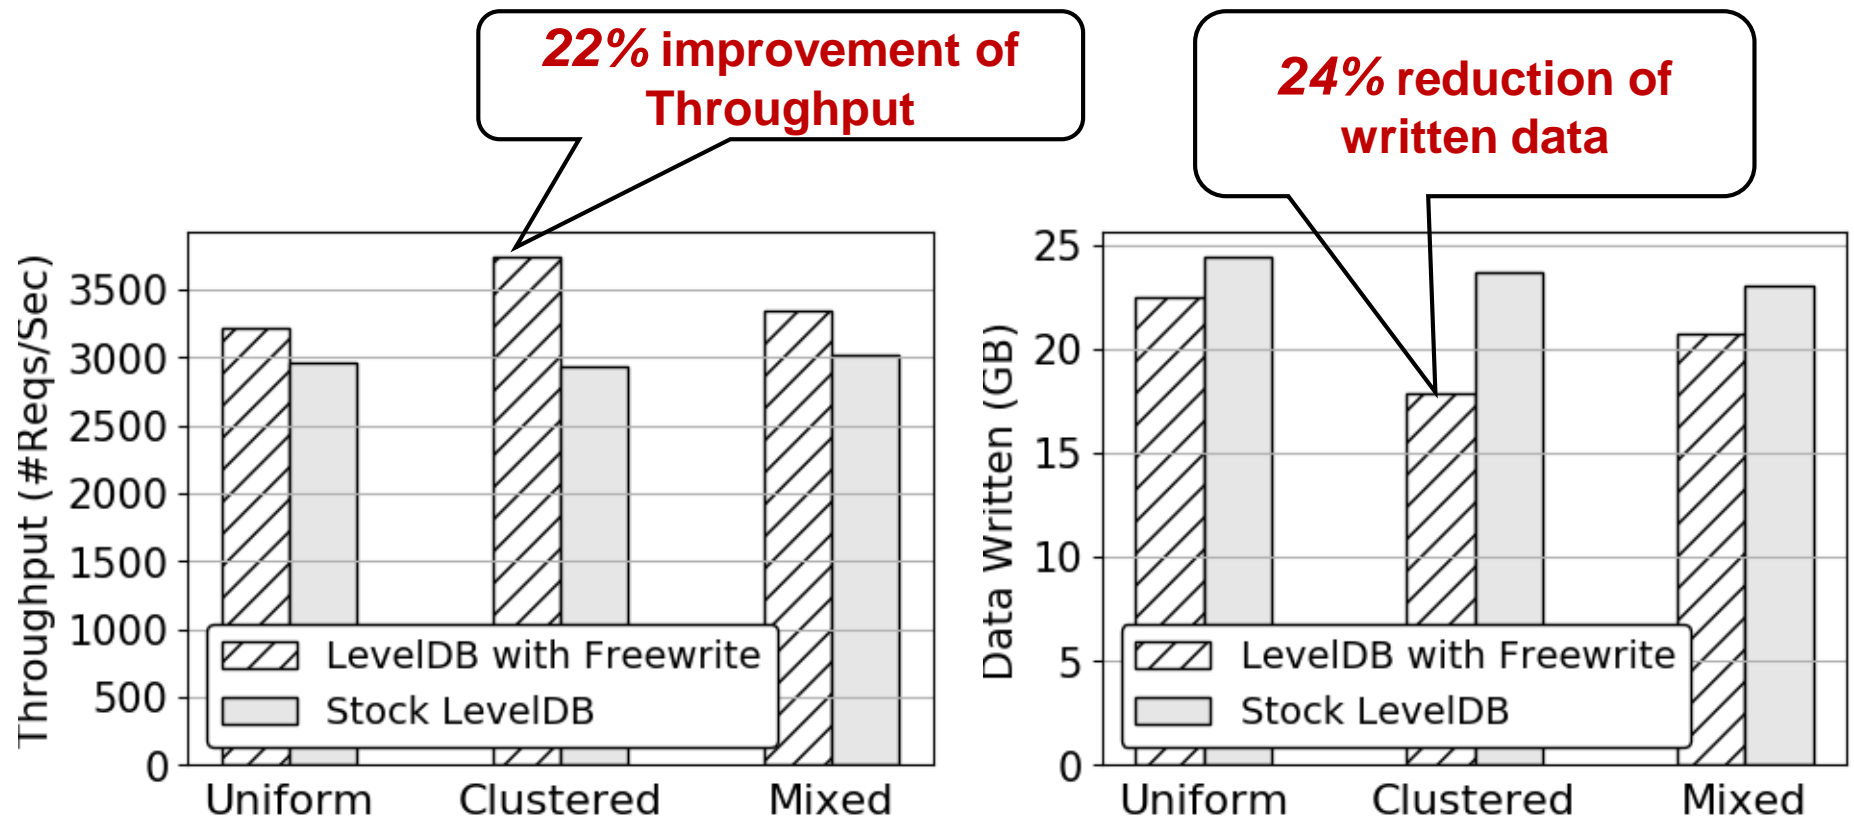


# Experimental Setup

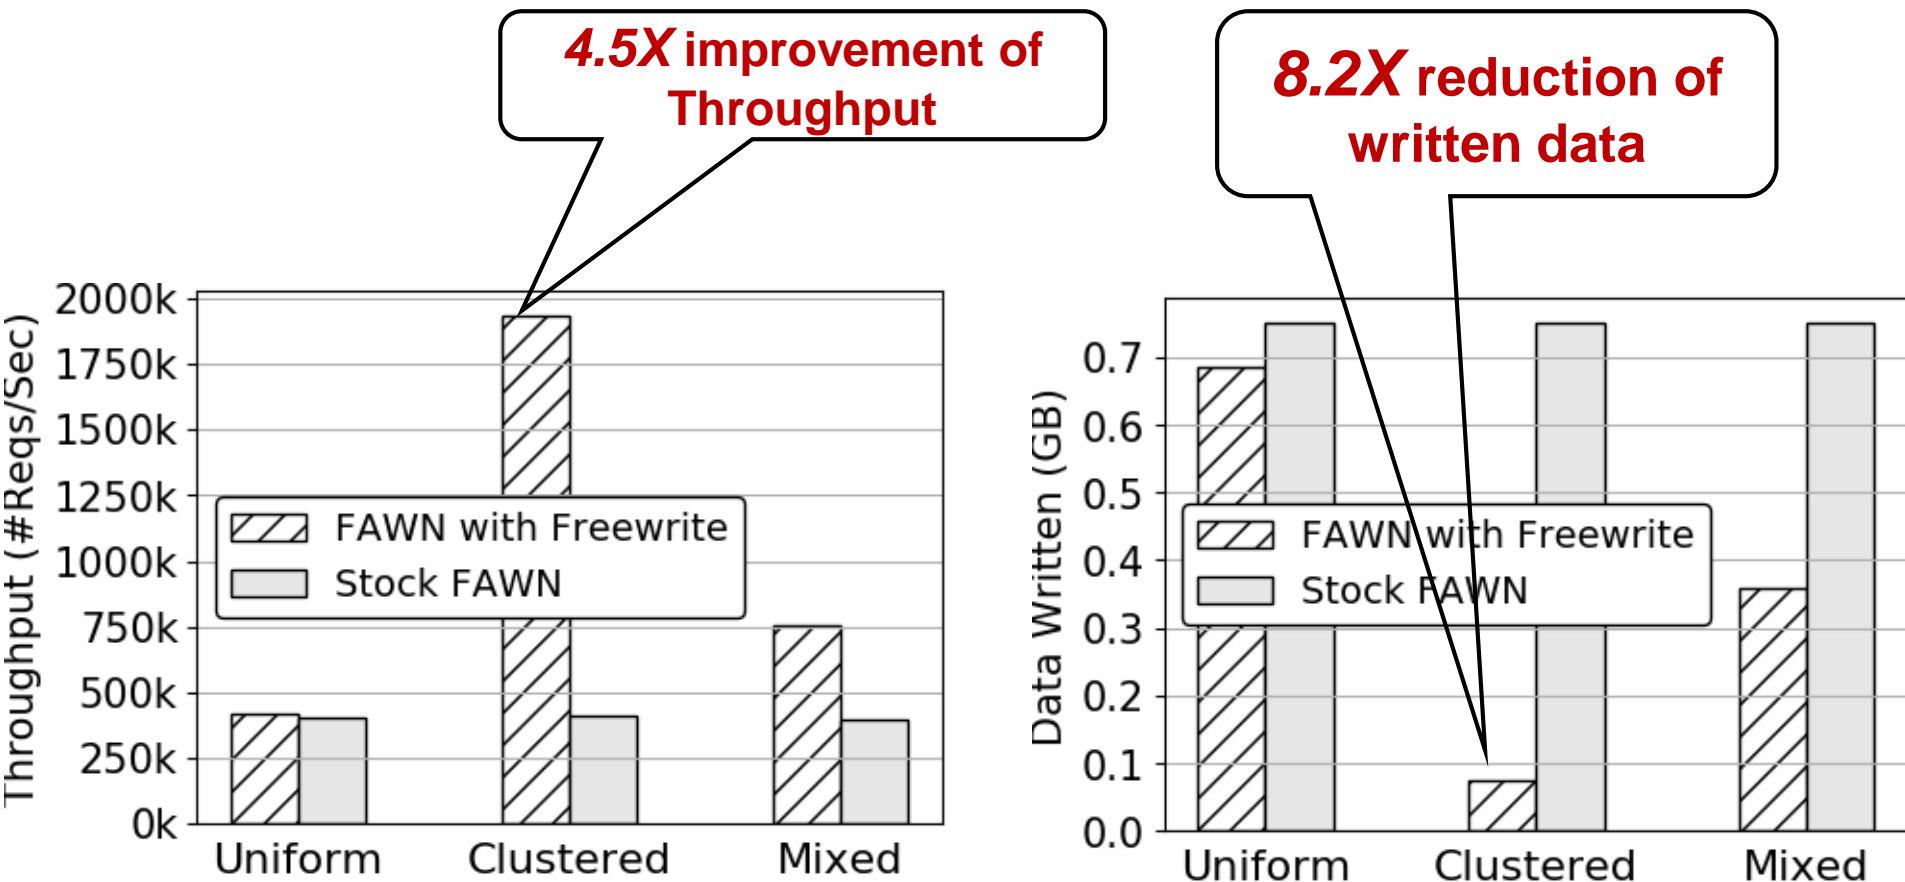
---

- Implemented at the application level (LevelDB and FAWN)
- Use Dmddedup as block-level dedup.
- Use Samsung 840 EVO 1GB SSD.
- Use *put* requests with 4% repeated keys.
- Threshold of bubble size is 12.5% of block size (512B), and total bubble space overhead is capped at 5%.
- Three key distributions: *uniform*, *clustered*, and *mixed*.

# Experiment Results (LevelDB)



# Experiment Results (FAWN)



# Summary

---

- Freewrite creates **deduplication opportunities** to leverage the **readily available** block-level deduplication system.
- Freewrite improves **write performance** and helps to extend **SSD lifetime**.
- The overhead of using bubbles is **very limited**.