

Implementing LSM Tree based Key Value Store

Sumatra Dhimoyee and Ayesha
Naeem



Why LSM?

They reduce the number of disk seeks by batching and sorting the writes in memory before flushing them to disk.

They reduce the amount of data written to disk by compacting and merging the files periodically, eliminating duplicates and stale data.

Memory Component

Memory Component



Vector

Disk Component


Entry

Key



Value

Tombstone flag

Timestamp





Block





4096 BYTES

SST






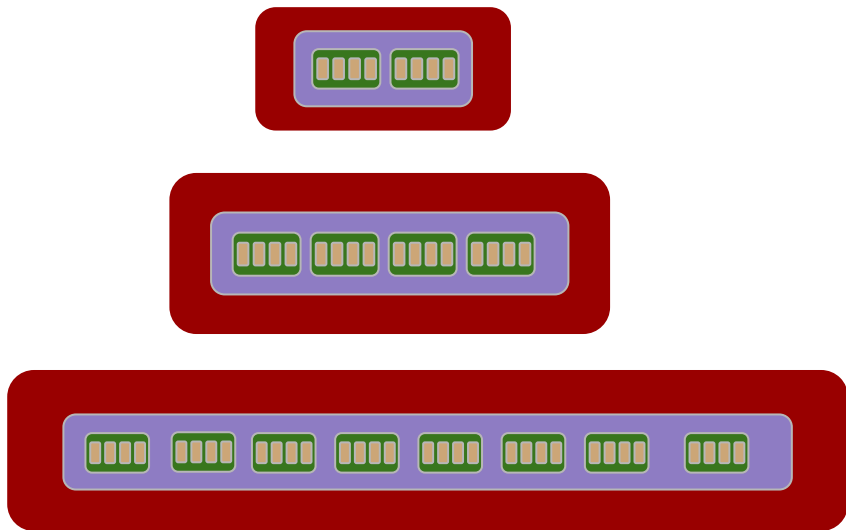
LEVEL IN A
LEVELED
LSM TREE



LEVEL IN A
TIERED LSM
TREE




Leveled




T = 2

Tiered

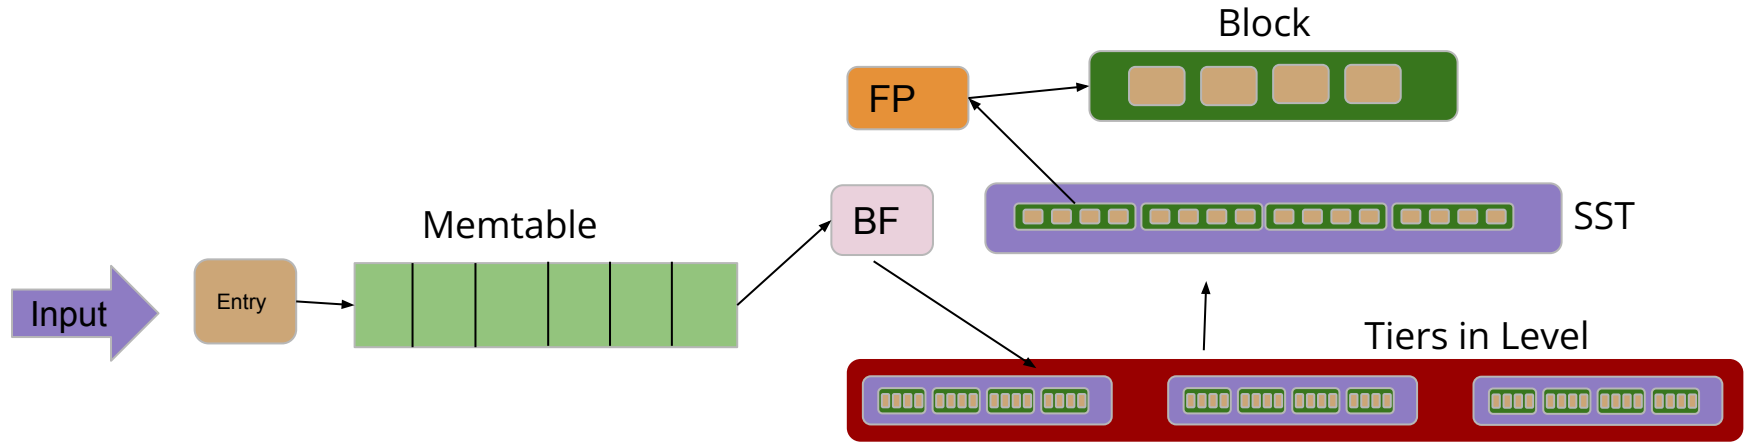


T = 2


Write Operation



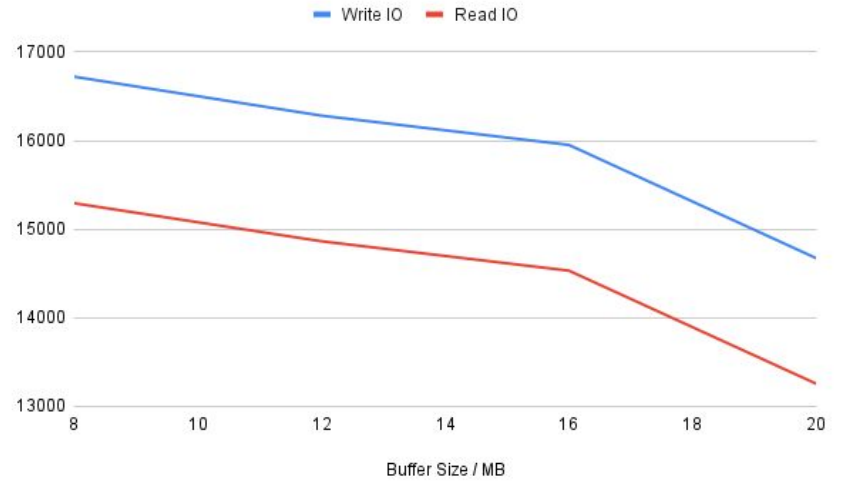
Read Operation



Results

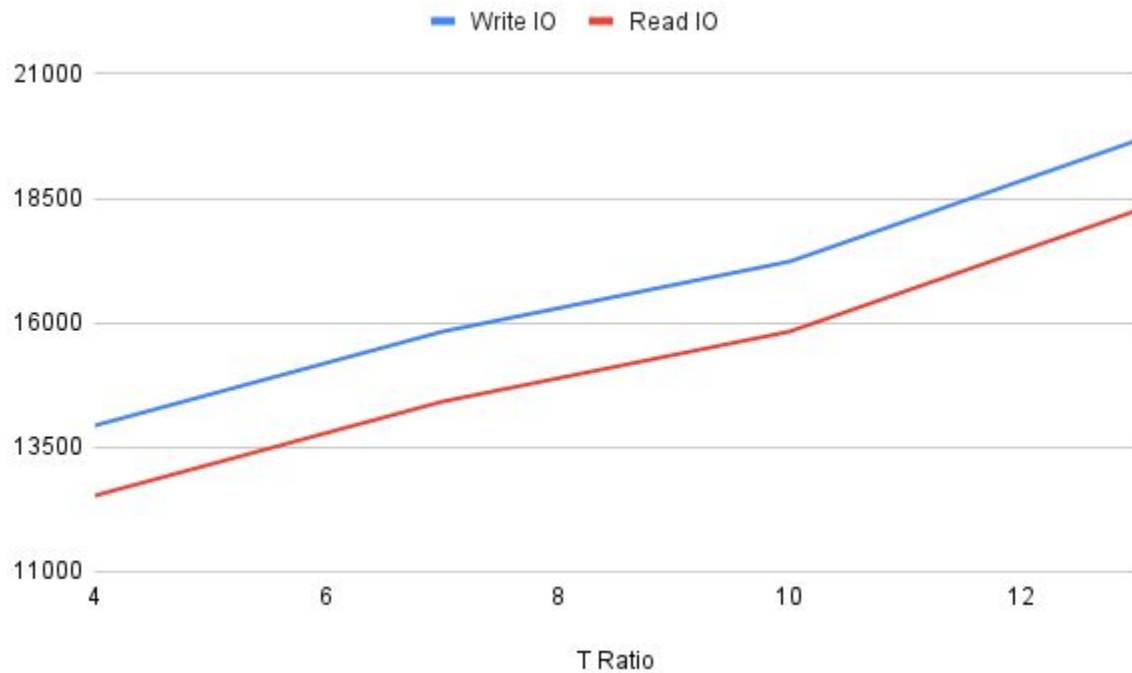


100K Inserts, Leveled Tree



100K Inserts, Tiered Tree

Results



100K Insert only - Leveled Tree

What's Remaining

Range Queries

Range Deletes

Thank you!

