


LSM Background

LSMs are Everywhere

fast ingestion + competitive lookup


RocksDB

cassandra

BigTable

levelDB

influxdb

mongoDB


SCYLLA

How do LSMs realize Range Query


Why they need a lot of computational resources

?

Range Queries in LSM-trees




Compaction in LSM

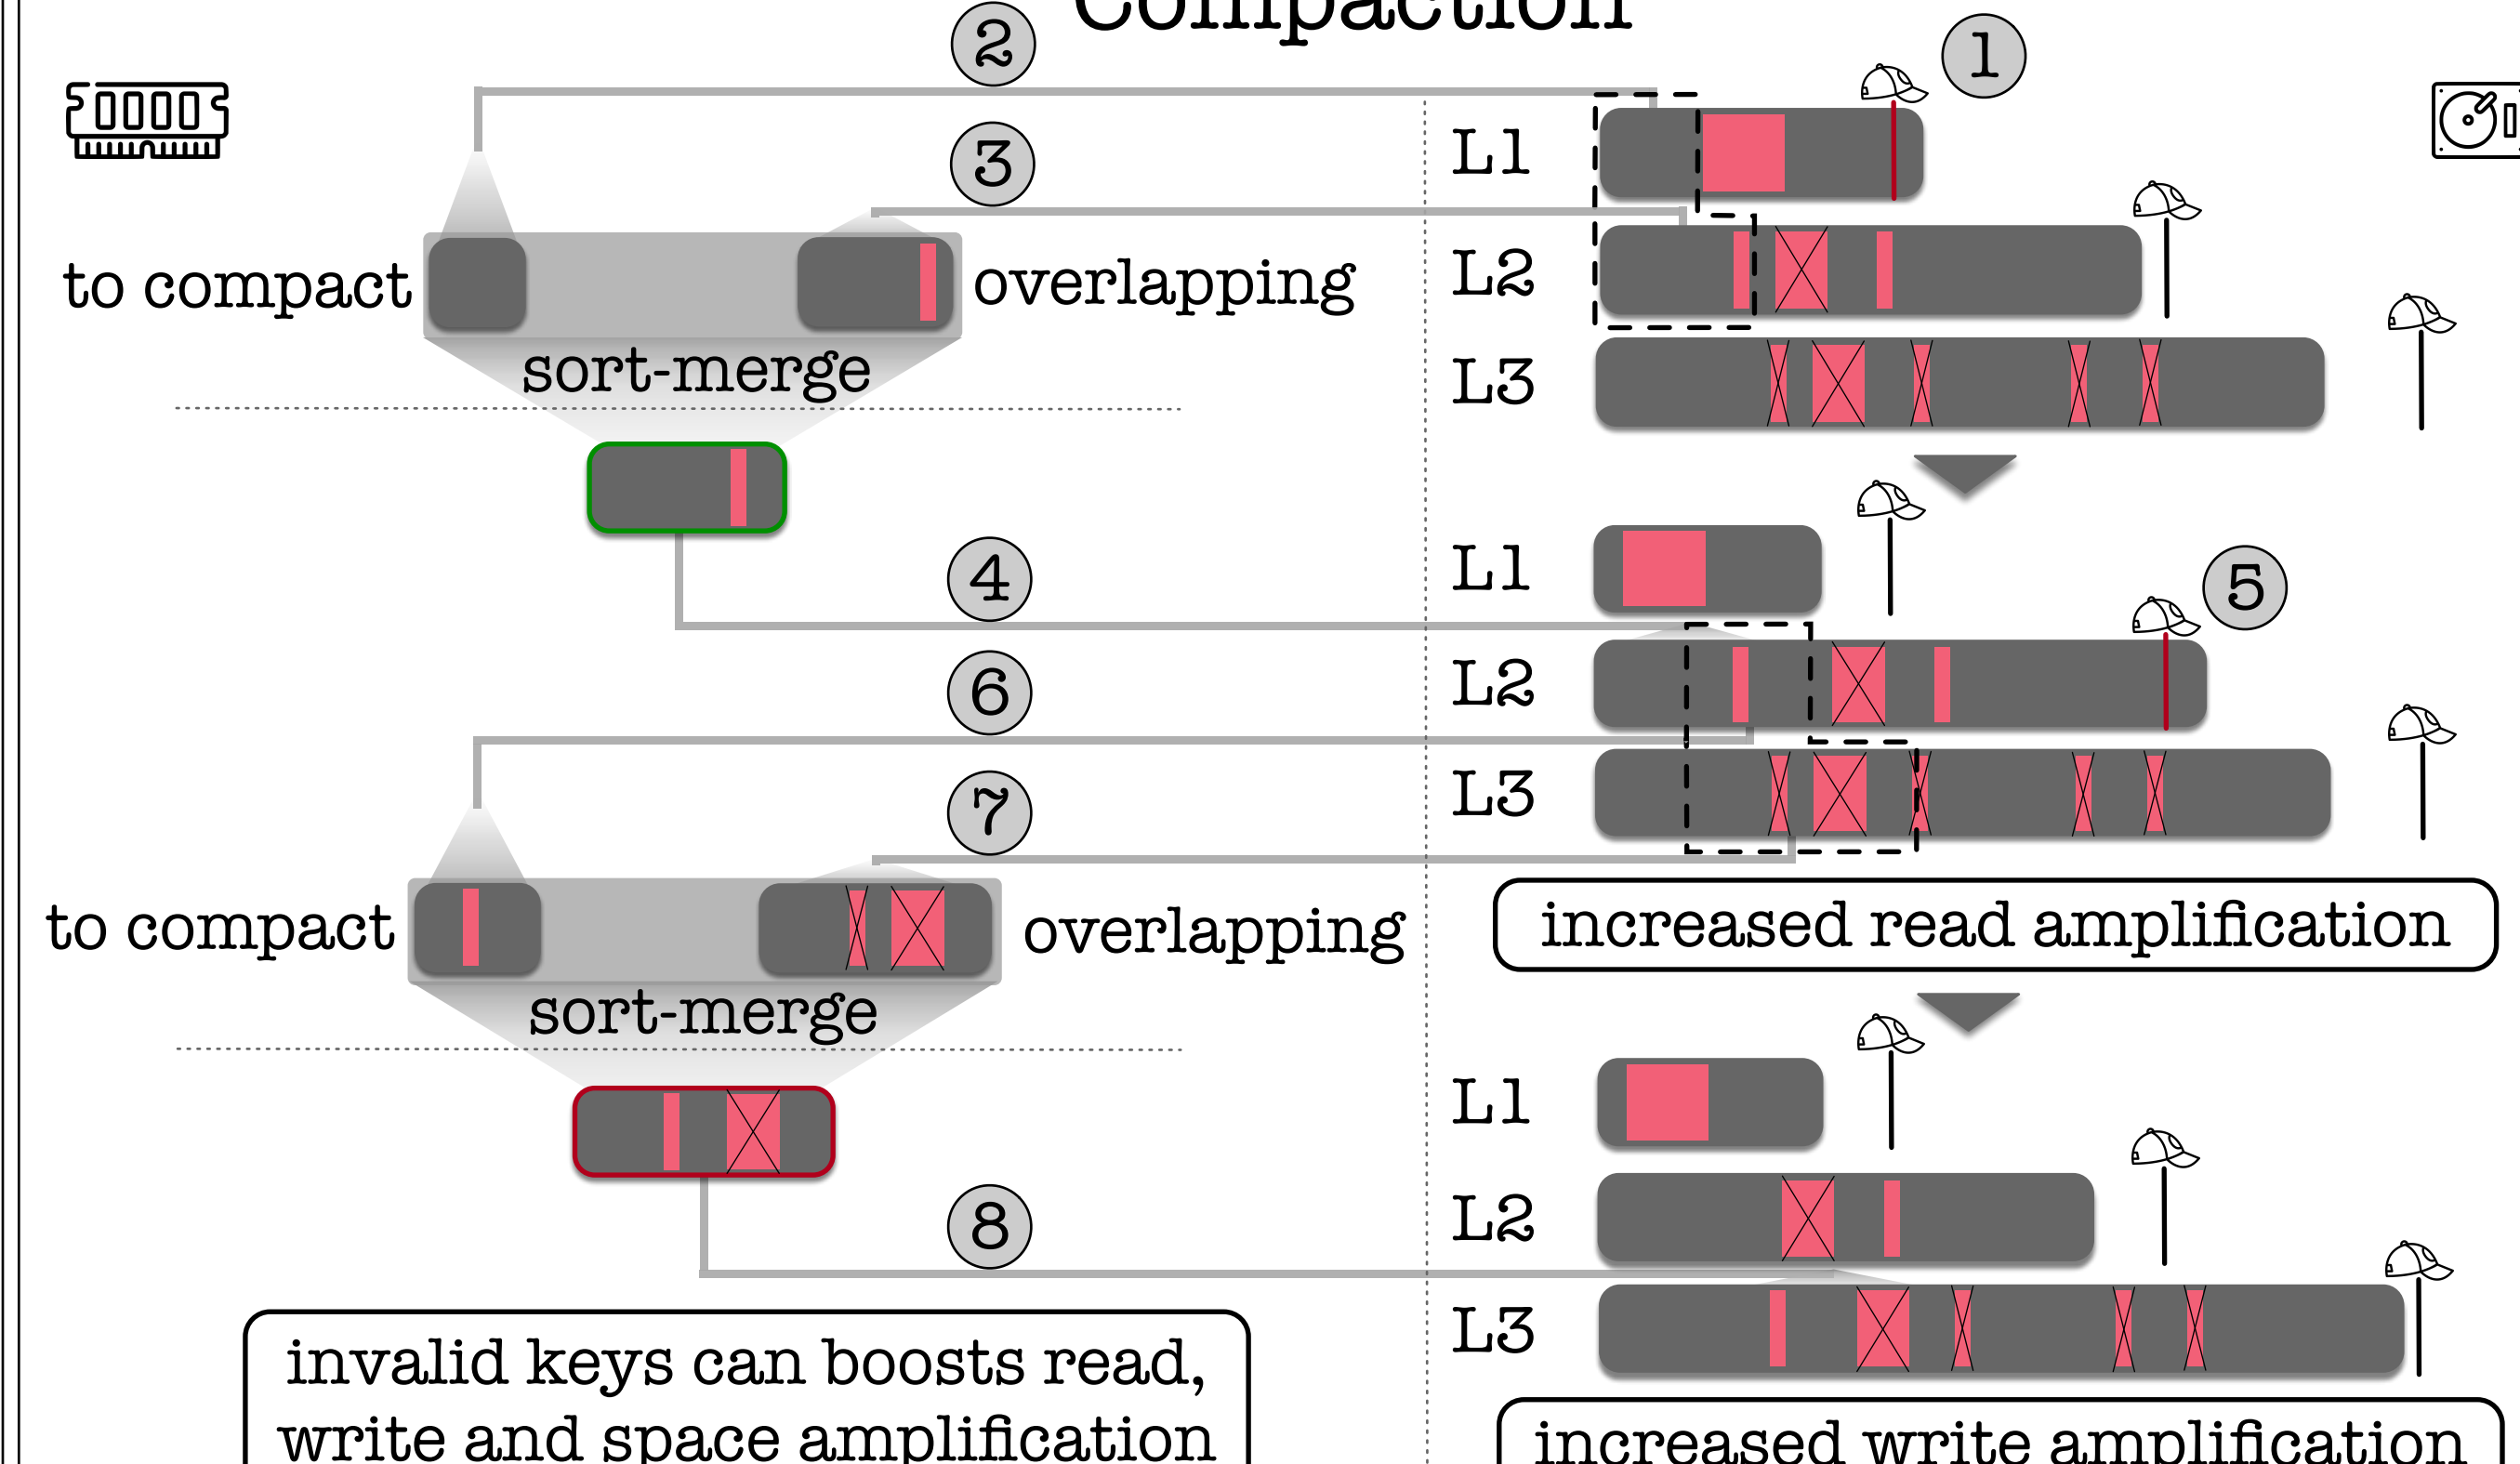


Problem

Range Query




Compaction




RangeReduce

Overview







Executing Range Queries



Evaluation

Performance Evaluation

Workload: 5M inserts initially, then 10 epochs of 125K updates with 100 interleaved range queries (10% selectivity).

RangeReduce is 3.6% faster than Rocksdb for range queries

RangeReduce reads 14.3% less data for range queries

RangeReduce achieves slightly lower space amplification

write amplification debt is 0.8x of Rocksdb