



LSM Design Space

### LSMs Everywhere

Fast Ingestion  
+  
Good Reads  
+  
Great Space Utilization

LSM Configuration

### Memory Buffer Implementation

### Data Layout on Disk

### Other Design Choices

- memory
  - Memory buffer
  - Bloom filter
  - Fence pointer
- tuning
  - Size ratio
  - File size
  - Page size
- data structure
  - Static/Dynamic
  - Prefix length
  - Hash buckets

Problem

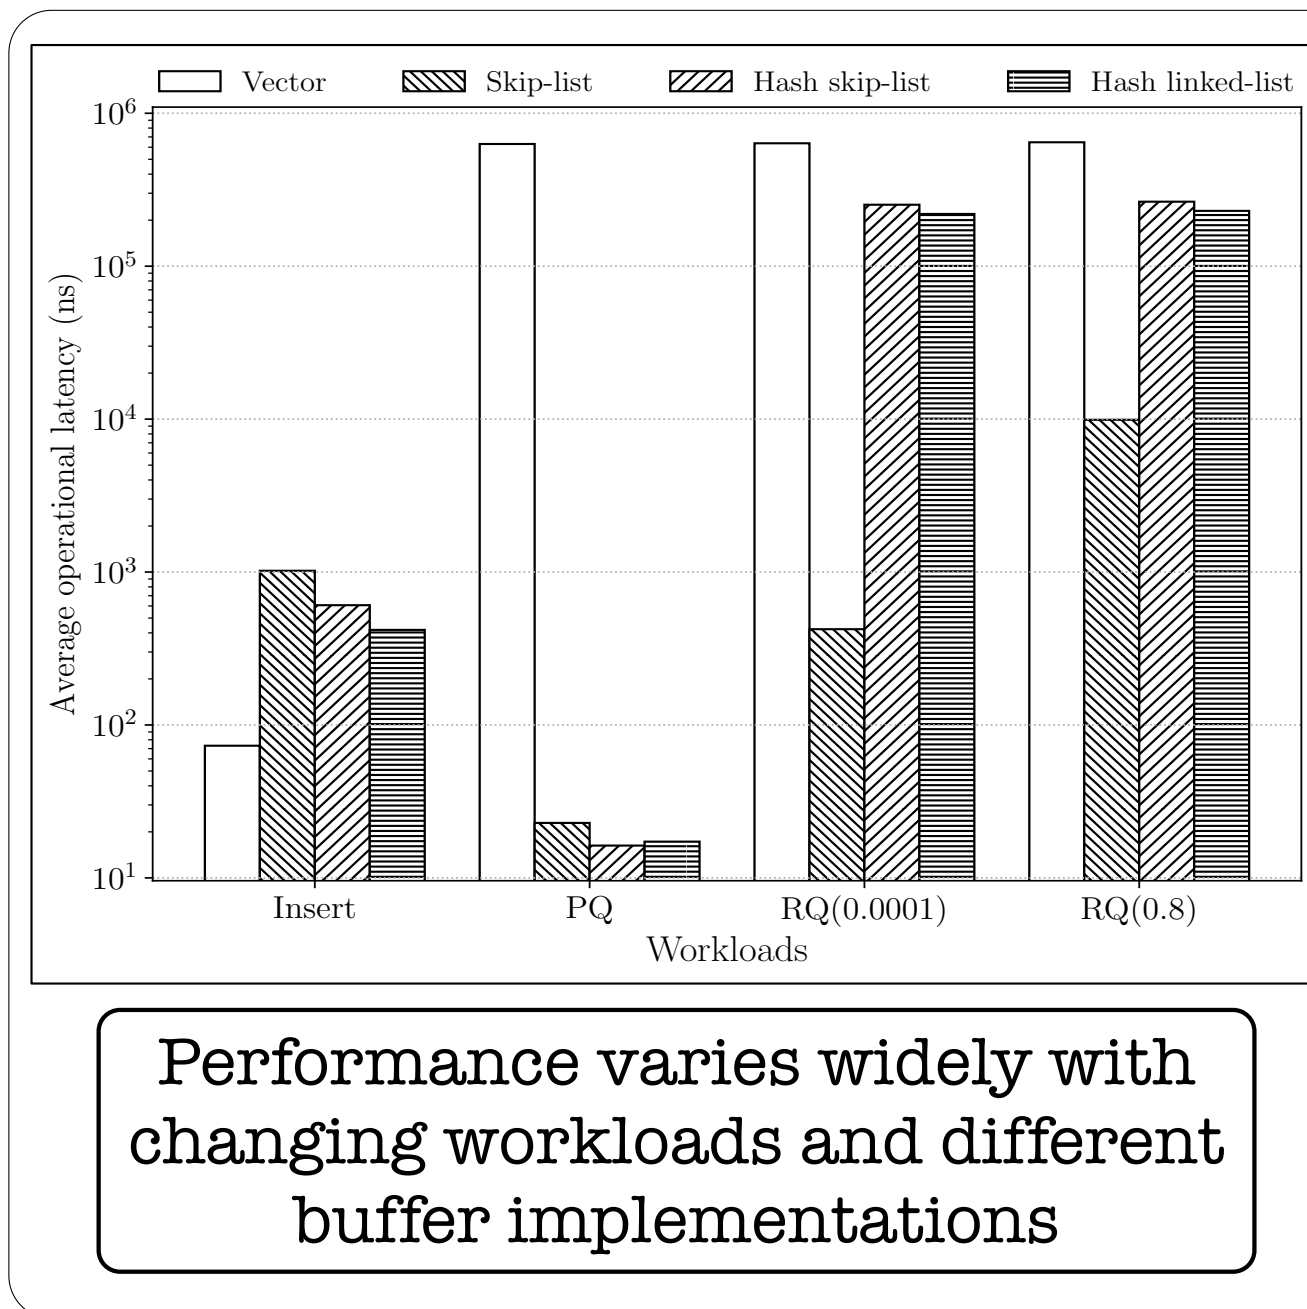
### Tuning LSMs

workload

target performance

LSM tuning

What is the optimal LSM configuration?



### Design Challenges

- Hand tuning not feasible
- Vast design space
- Hard to model interactions
- Dynamic workloads
- Lack of benchmark dataset

State of the Art

### Reinforcement Learning: RUSKEY

FLSM Tree

Level statistics

Current state  
Action (tuning)  
Next state  
Reward

Actor-Critic Network

Compaction Strategy

Single design choice

Single-level learning

### Bayesian Optimization: LlamaTune

Projection matrix for low-dimensional tuning

LHS to generate initial low-dimensional points

Points projected to high-dimension (REMBO/HeSBO)

Initiate BO surrogate model

Exploration/exploitation

Optimized function

Inefficient collection of samples by running workload on DBMS

### Machine Learning: OtterTune

Factor analysis for dimensionality reduction

k-means clustering for metric identification

Lasso regression to identify important knobs

Euclidean distance for mapping workload

Gaussian process for optimal configuration

Data Repository

new observations

Model requires feature engineering; not trained end-to-end

Solution Approach

### Workload Change

Reads: 5%  
Writes: 90%  
Updates: 5%

Point Reads: 65%  
Range Reads: 25%  
Writes: 5%  
Updates: 5%

### LSM Tuning Predictions from Learned Model

Workload composition

Distribution of operations

Target performance

Buffer Implementation

Data Layout

Memory Allocation

Size Ratio

Neural network model

Online learning

Real-time predictions

### Lazy Transition to Optimal Design

Buffer Implementation

Data Layout on Disk

Memory Allocation

Size Ratio

Vector

Skiplist

Tiering

Leveling

T=3

T=5

buffer filter index

