

# KVBench: A Key-Value Benchmarking Suite

*Zichen Zhu, Arpita Saha, Manos Athanassoulis, Subhadeep Sarkar*



Brandeis  
UNIVERSITY





# Key-Value Stores





FASTER



new key-value stores

# Key-Value Stores



new application requirements



run experiments



ORACLE  
NOSQL  
DATABASE




apples to apples comparison

# Key-Value Workload Generators & their Limitations

|           |          | operation |        |                       |                   |             |                        |                    |              | distribution |        |      |         |
|-----------|----------|-----------|--------|-----------------------|-------------------|-------------|------------------------|--------------------|--------------|--------------|--------|------|---------|
|           |          | insert    | update | non-empty point query | empty point query | range query | non-empty point delete | empty point delete | range delete | uniform      | normal | beta | Zipfian |
| benchmark | YCSB     | ✓         | ✓      | ✓                     | ✗                 | ✓           | ✗                      | ✗                  | ✗            | ✓            | ✗      | ✗    | ✓       |
|           | db_bench | ✓         | ✓      | ✓                     | ✓                 | ✓           | ✓                      | ✓                  | ✓            | ✗            | ✓      | ✗    | ✗       |
|           | KVBench  | ✓         | ✓      | ✓                     | ✓                 | ✓           | ✓                      | ✓                  | ✓            | ✓            | ✓      | ✓    | ✓       |

# Challenges



empty point queries  
not supported

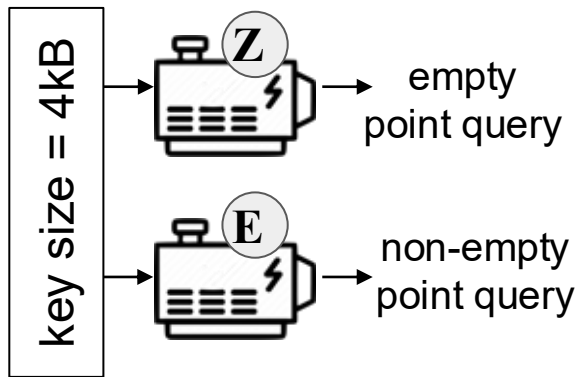


insert, update and point  
queries may have  
different distributions




deletes maybe  
interleaved with  
updates and queries

# Contributions



two generators for empty & non-empty point query

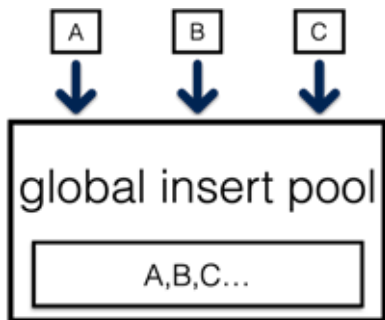
- mixed workload
- interleaved point queries



separate generator for update operation


- update and queries with different distributions

# Contributions



global insert pool to keep track of existing entries


- deletes interleaved
- deletion on existing keys




new set of benchmark workloads

- inspired by real-life workloads

# KVBench Suite



operations



distributions

# Operations



insert



update



query



delete

# Insert



- Create global insert pool
- Generate key until not matched in pool
- Insert new key into pool
- Output operation into workload file

```
HS = HashSet() //global insert pool
Key  $k$  = RandomKeyGenerator(dist)
while  $k \in$  HS : //check if key in pool
     $k$  =
    RandomKeyGenerator(dist)
HS.insert( $k$ ) //insert new key in pool
Value  $v$  = RandomValueGenerator()
workload_file.output("Insert",  $k$ ,  $v$ )
```

# Update



- Random key from global insert pool
- Random value generated
- Operation output to workload file

```
V = HS.toVector() //copy HS elements to vector  
Index idx = RandomIndexGenerator(|V|, dist)  
Key k = V.get(idx) //select key randomly from V  
Value v = RandomValueGenerator()  
workload_file.output("Update", k, v)
```

# Point Query



- *Empty point query:*  
Key generated randomly until not in pool
- *Non-empty point query:*  
Key selected randomly from insert pool
- Operation output to workload file

Empty point-query:

Key  $k = \text{RandomKeyGenerator}()$

*while*  $k \in HS$  : //generate key until not in pool

$k = \text{RandomKeyGenerator}()$

workload\_file.output("Query",  $k$ )


Non-empty point-query:

Index  $idx = \text{RandomIndexGenerator}(|V|, \text{dist})$

Key  $k = V.\text{get}(idx)$  //select key randomly from  $V$

workload\_file.output("Query",  $k$ )

# Range Query



- Random start key skipping  $Y * |V|$  entries
- Offset by selectivity for end key
- Operation output to workload file

```
V = V.sort() //sort vector
```

```
Index start_idx = RandomIndexGenerator(|V|*(1-Y), dist)
```

```
Key start_key = V.get(start_idx)
```

```
//random start key skipping  $Y * |V|$  entries
```

```
Index end_idx = start_idx + |V| * Y //offset by Y for end key
```

```
Key end_key = V.get(end_idx)
```

```
workload_file.output("RangeQuery", start_key, end_key)
```

# Point Delete



- Key selected randomly from insert pool
- Remove key from pool
- Delete operation output to workload file

Index  $idx = \text{RandomIndexGenerator}(|V|, \text{dist})$


Key  $k = V.\text{get}(idx)$  //random key from pool

$V.\text{remove}(k)$  //remove key from vector

$HS.\text{remove}(k)$  //remove key from hashset

$\text{workload\_file.output}(\text{"Delete"}, k)$

# Range Delete



- Random start key generated
- Offset by selectivity for end key
- Remove keys in range from pool
- Operation output to workload file

```
V = V.sort() //sort vector
```

```
Index start_idx = RandomIndexGenerator(|V|*(1-Y), dist)
```

```
Key start_key = V.get(start_idx)//random start key
```

```
Index end_idx = start_idx+|V|*Y//offset by Y for end key
```

```
Key end_key = V.get(end_idx)
```


```
for idx in range(start_idx, end_idx):
```

```
    HS.remove(V.get(idx))
```


```
V.remove(start_idx,end_idx) //remove all keys in range from pool
```

```
workload_file.output("RangeQuery", start_key, end_key)
```

# KVBench Suite




operations



distributions


# Distributions

## Uniform



- `uniform_int_distribution()` in C++
- random keys for update and non-empty point queries


## Normal



- `normal_distribution` in C++
- mean percentile for  $\mu$  index and scaled  $\delta$
- keys chosen from vector accordingly


# Distributions

Zipfian



- Shuffle keys randomly
- Index selected according to Zipfian distribution

Beta



- Gamma\_distribution in C++
- 2 gamma to simulate a beta

# Mixed Workloads



## thresholding

- five thresholds
- # inserts before other operations



## interleaving

- deletes/inserts
- vector updated, sorted
- probability recalculated



## preloading

- existing insert-only workload
- read all inserts for HS and V

| Benchmark Workloads                |                           | KVB-I   | KVB-II  | KVB-III | KVB-IV  | KVB-V   |
|------------------------------------|---------------------------|---------|---------|---------|---------|---------|
| Preloading/Interleaving/Sequential |                           | P       | I       | P       | P       | S       |
| Operations                         | Insert                    | -       | 50%     | -       | -       | 95%     |
|                                    | Update                    | -       | 25%     | 50%     | 50%     | -       |
|                                    | Non-empty point query     | 20%     | -       | 25%     | -       | 5%      |
|                                    | Empty point query         | 80%     | 15%     | 25%     | -       | -       |
|                                    | Range query               | -       | -       | -       | -       | -       |
|                                    | Point delete              | -       | 10%     | -       | -       | -       |
|                                    | Range delete              | -       | -       | -       | 50%     | -       |
| Distributions                      | Insert distribution       | -       | uniform | -       | -       | Zipfian |
|                                    | Update distribution       | -       | uniform | Zipfian | Zipfian | -       |
|                                    | Non-empty pq distribution | beta    | uniform | uniform | -       | uniform |
|                                    | Empty pq distribution     | uniform | -       | uniform | -       | -       |

# Benchmark Setup

## Metrics

- Query latency
- Throughput

## System setup


- 2 Intel Xeon CPU@2.1GHz
- 375 GB RAM, 350 GB SSD

## Workload

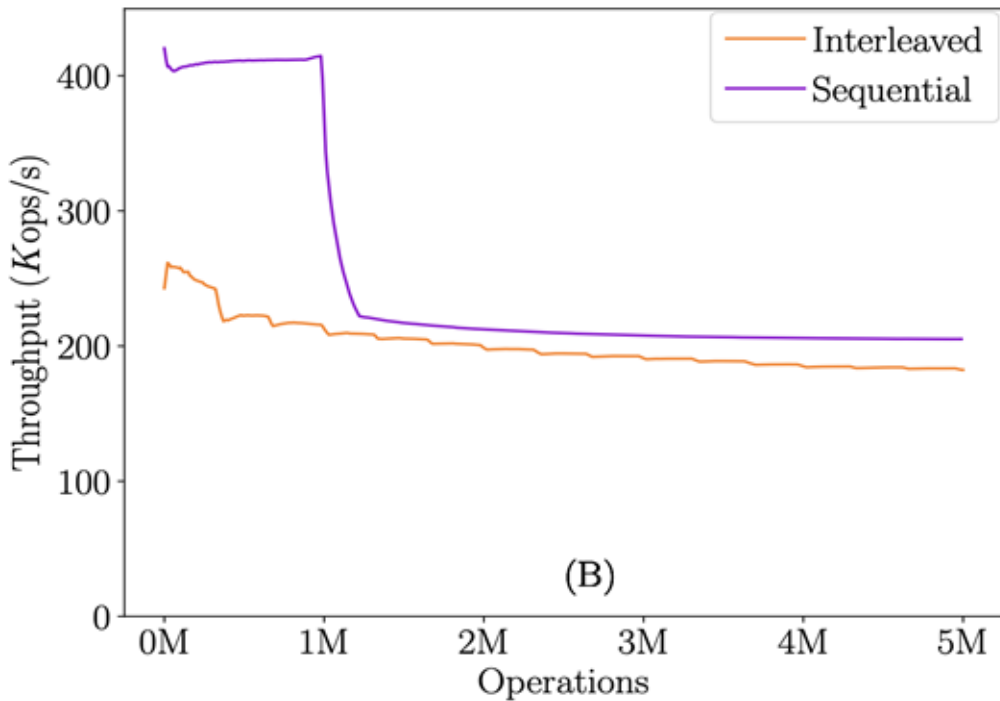
- entry size: 1024B
- # I, Q, U = 10M

## Index setup

- RocksDB
- LSM Tree (T =4 , P=4096, B=8)



The average latency per query decreases as there are more empty point queries for different distributions.



## Workload

# U = 0.9M  
# D = 0.1M  
# Q = 4M  
# I = 10M (P)

The throughput differs substantially between interleaved and sequential workloads even if they have the same workload composition.

# Future Work


increase efficiency of generating workload

more operations (queries on deleted keys, repeatedly deleted keys)

different selectivity and distributions for different range queries

integrate the workload generator into a general framework/ existing benchmark framework to test more databases

# KVBench: A Key-Value Benchmarking Suite



Distinguish between empty and non-empty point lookups



Different operations can have different distributions



Deletes can be mixed with other operation types



# Thanks!



**Brandeis**  
UNIVERSITY

