

QTune: A Query-Aware Database Tuning System with Deep Reinforcement Learning

Paper Authors: Guoliang Li, Xuanhe Zhou, Shifu Li, Bo Gao

Presenter: Jiahao Gai

Background

- Knob tuning is an NP-hard problem and existing methods have several limitations.

- 1. Limited Scope and Time-Consuming for DBAs manual tuning**
- 2. Dependency on High-Quality Training Data**
- 3. Coarse-Grained Tuning**

Contributions of the paper

1. A query-aware database tuning system using DRL
2. A SQL query featurisation model
3. A DS-DDPG model
4. A DL based query clustering method
5. Experiments on various query workloads and databases outperforming SOTA.

Architecture

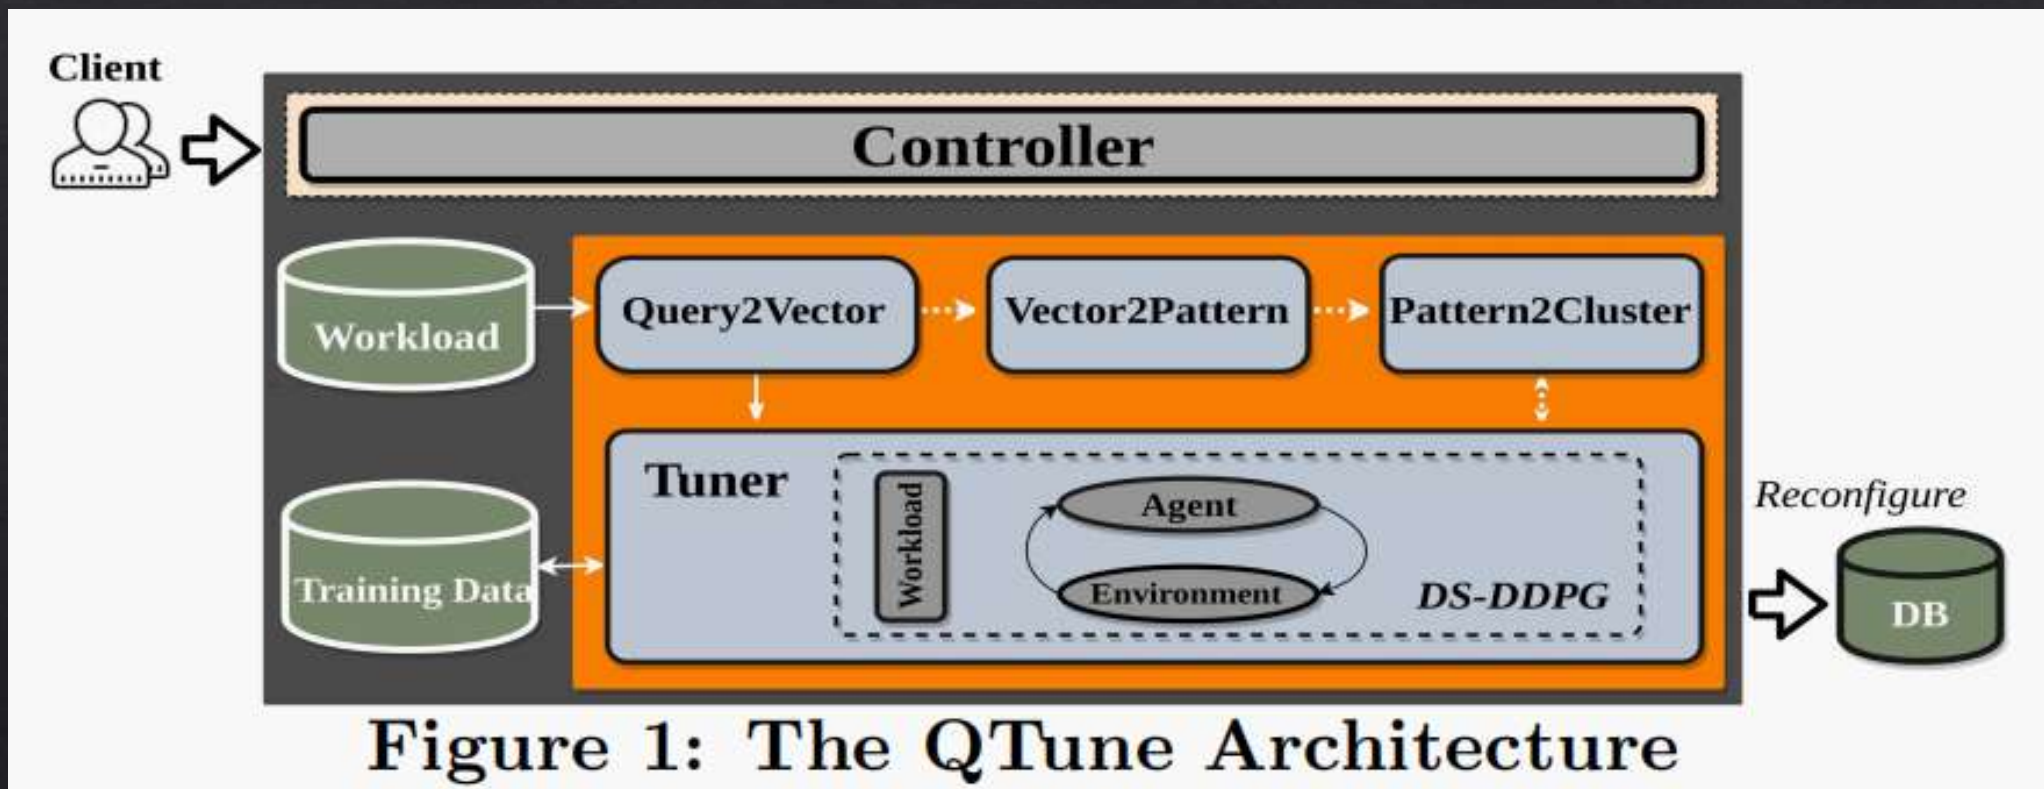


Figure 1: The QTune Architecture

Workflow

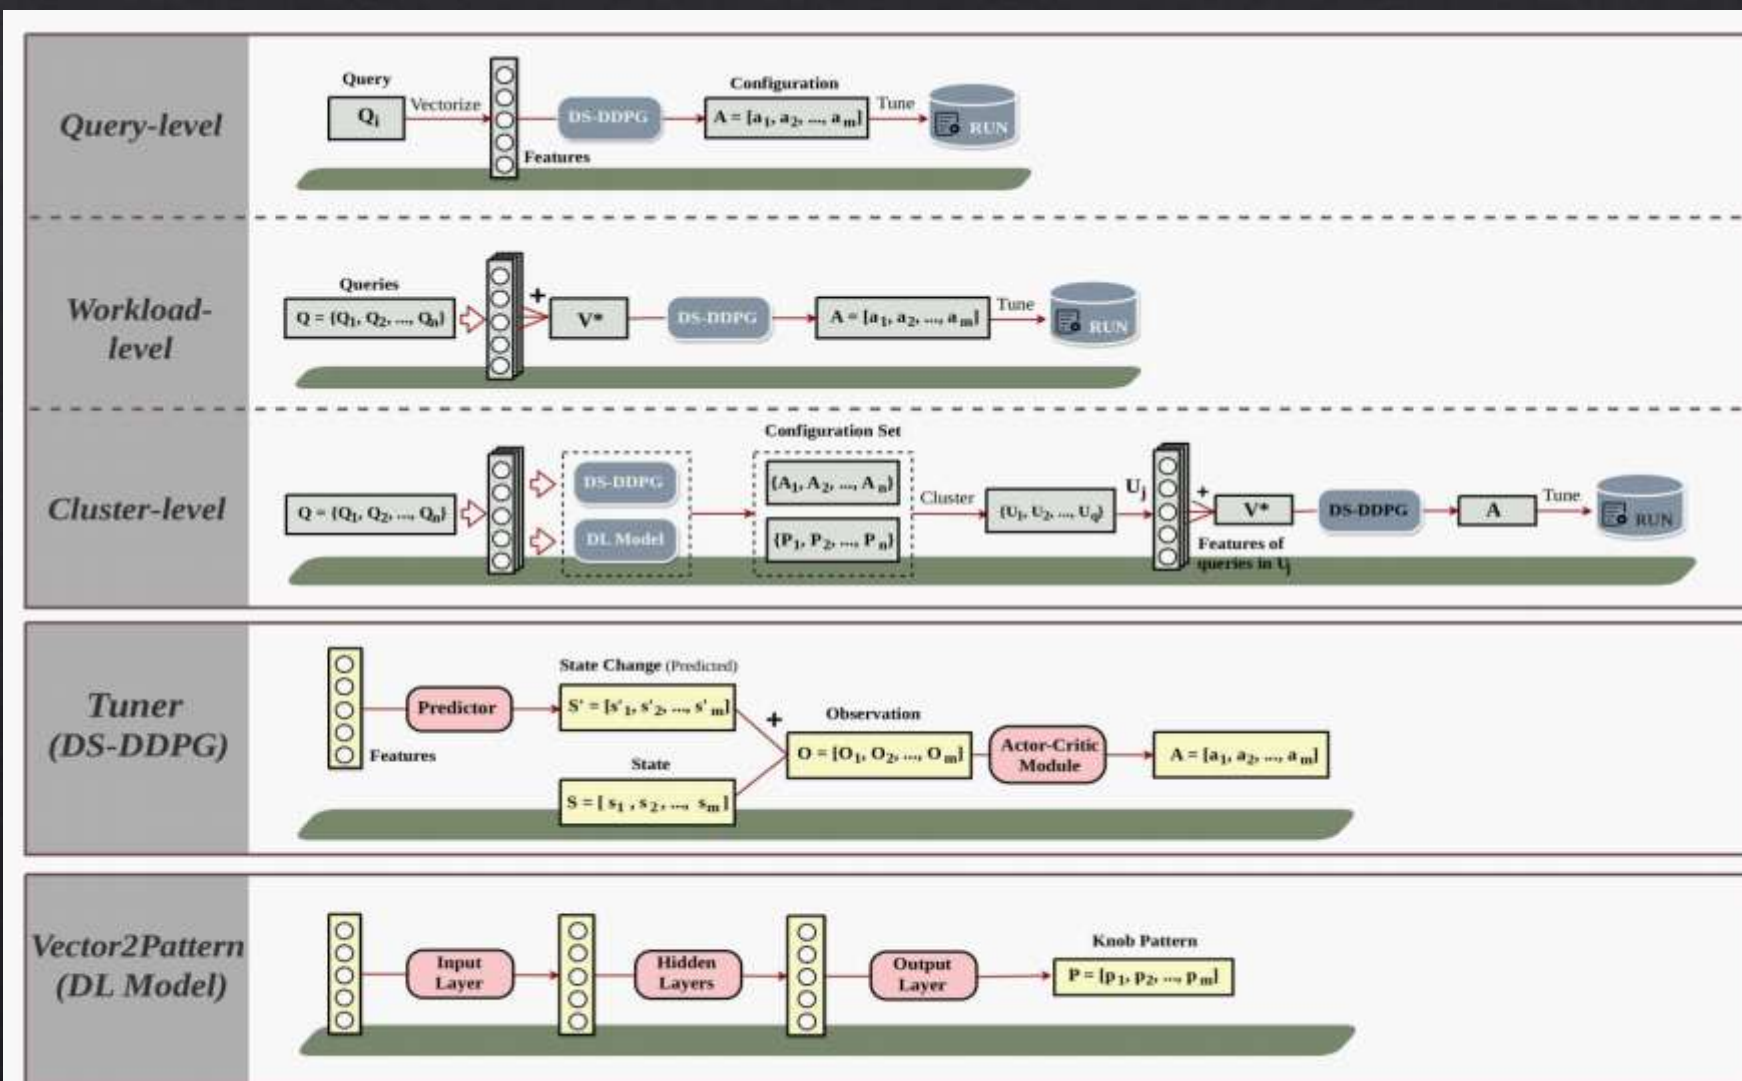


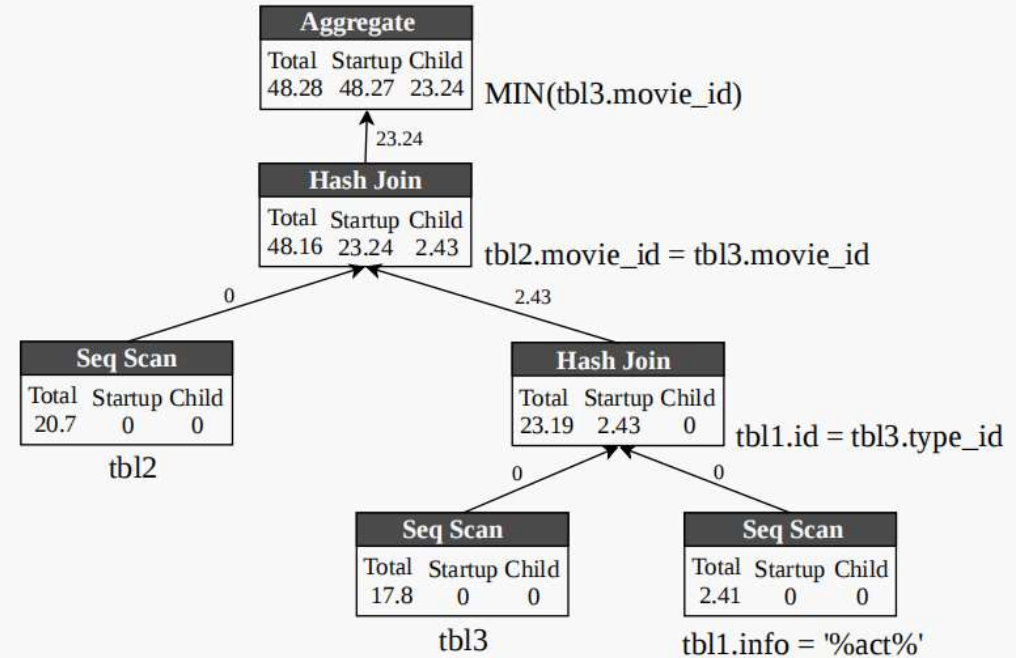
Figure 2: Workflow of QTune

Query2Vector

1. Query information (0/1)
2. Cost Information (real value)

```

SELECT      MIN(tbl3.movie_id)
FROM        tbl1, tbl2, tbl3
WHERE       tbl1.info = '%act%'
           AND  tbl1.id = tbl3.type_id
           AND  tbl2.movie_id = tbl3.movie_id
    
```



| Insert | Delete | Update | Select | tbl1 | tbl2 | tbl3 | ... | tbl8 | Hash_Join | Seq_Scan | Aggregate | ... |
|--------|--------|--------|--------|------|------|------|-----|------|-----------|----------|-----------|-----|
| 0 | 0 | 0 | 1 | 1 | 1 | 1 | ... | 0 | 68.92 | 40.91 | 25.04 | ... |

(1) DML (2) Tables (3) Operation Costs

| Insert | Delete | Update | Select | tbl1 | tbl2 | tbl3 | ... | tbl8 | Hash_Join | Seq_Scan | Aggregate | ... |
|--------|--------|--------|--------|------|------|------|-----|------|-----------|----------|-----------|-----|
| 0 | 0 | 0 | 1 | 1 | 1 | 1 | ... | 0 | 0.1401 | -0.166 | -0.2423 | ... |

Normalized Feature Vector

Figure 3: Character Encoding.

DS-DDPG model

1. Environment
2. Predictor
3. Actor
4. Critic

Training of DS-DDPG model

1. Training the Predictor
2. Training the Actor-Critic Module

Algorithm 1: Training DS-DDPG

Input: U : the query set $\{q_1, q_2, \dots, q_{|U|}\}$

Output: π_P, π_A, π_C

- 1 Generate training data T_P ;
- 2 TrainPredictor(π_P, T_P);
- 3 Generate training data T_A ;
- 4 TrainAgent(π_A, π_C, T_A);

Training the Predictor

Function TrainPredictor(π_P, T_P)

Input: π_P : The weights of a neural network; T_P :
The training set

- 1 Initiate the weights in π_P ;
 - 2 **while** *!converged* **do**
 - 3 **for** *each* $(v, S, I, \Delta S) \in T_P$ **do**
 - 4 Generate the output G of $\langle v, S, I \rangle$;
 - 5 Accumulate the backward propagation error:
 $E = E + \frac{1}{2} \|G - \Delta S\|^2$;
 - 6 Compute gradient $\nabla_{\theta_s}(E)$, update weights in π_P ;
-

Training the Agent

Function TrainAgent(π_A, π_C, T_A)

Input: π_A : The actor's policy; π_C : The critic's policy; T_A : training data

- 1 Initialize the actor π_A and the critic π_C ;
 - 2 **while** *!converged* **do**
 - 3 Get a training data
 $T_A^1 = (S'_1, A_1, R_1), (S'_2, A_2, R_2), \dots, (S'_t, A_t, R_t)$;
 - 4 **for** $i = t - 1$ *to* 1 **do**
 - 5 Update the weights in π_A with the
 action-value $Q(S'_i, A_i | \pi_C)$;
 - 6 Estimate an action-value
 $Y_i = R_i + \tau Q(S'_{i+1}, \pi_A(S'_{i+1} | \theta^{\pi_A}) | \pi_C)$;
 - 7 Update the weights in π_C by minimizing the
 loss value $L = (Q(S'_i, A_t | \pi_C) - Y_i)^2$;
-

Granularities of tuning

1. **Query-level:** can optimise the query latency; but low throughput
2. **Workload-level:** cannot optimise the query latency; high throughput
3. **Cluster-level:** can optimise both the latency and throughput

Granularities of tuning

1. **Query-level:** can optimise the query latency; but low throughput
2. **Workload-level:** cannot optimise the query latency; high throughput
3. **Cluster-level:** can optimise both the latency and throughput

Evaluation

1. Three query workloads JOB, TPC-H and Sysbench

Table 3: Workloads. RO, RW and WO denote read-only, read-write and write-only respectively.

| Name | Mode | Table | Cardinality | Size(G) | Query |
|----------|--------|-------|-------------|---------|---------|
| JOB | RO | 21 | 74,190,187 | 13.1 | 113 |
| TPC-H | RO | 8 | 158,157,939 | 50.0 | 22 |
| Sysbench | RO, RW | 3 | 4,000,000 | 11.5 | 474,000 |

2. Metrics: latency, throughput, as well as the training and tuning time

3. Three kinds of databases

Table 2: Database information

| Database | Knobs without restart | State Metrics |
|------------|-----------------------|---------------|
| PostgreSQL | 64 | 19 |
| MySQL | 260 | 63 |
| MongoDB | 70 | 515 |

Evaluating three types of tuning

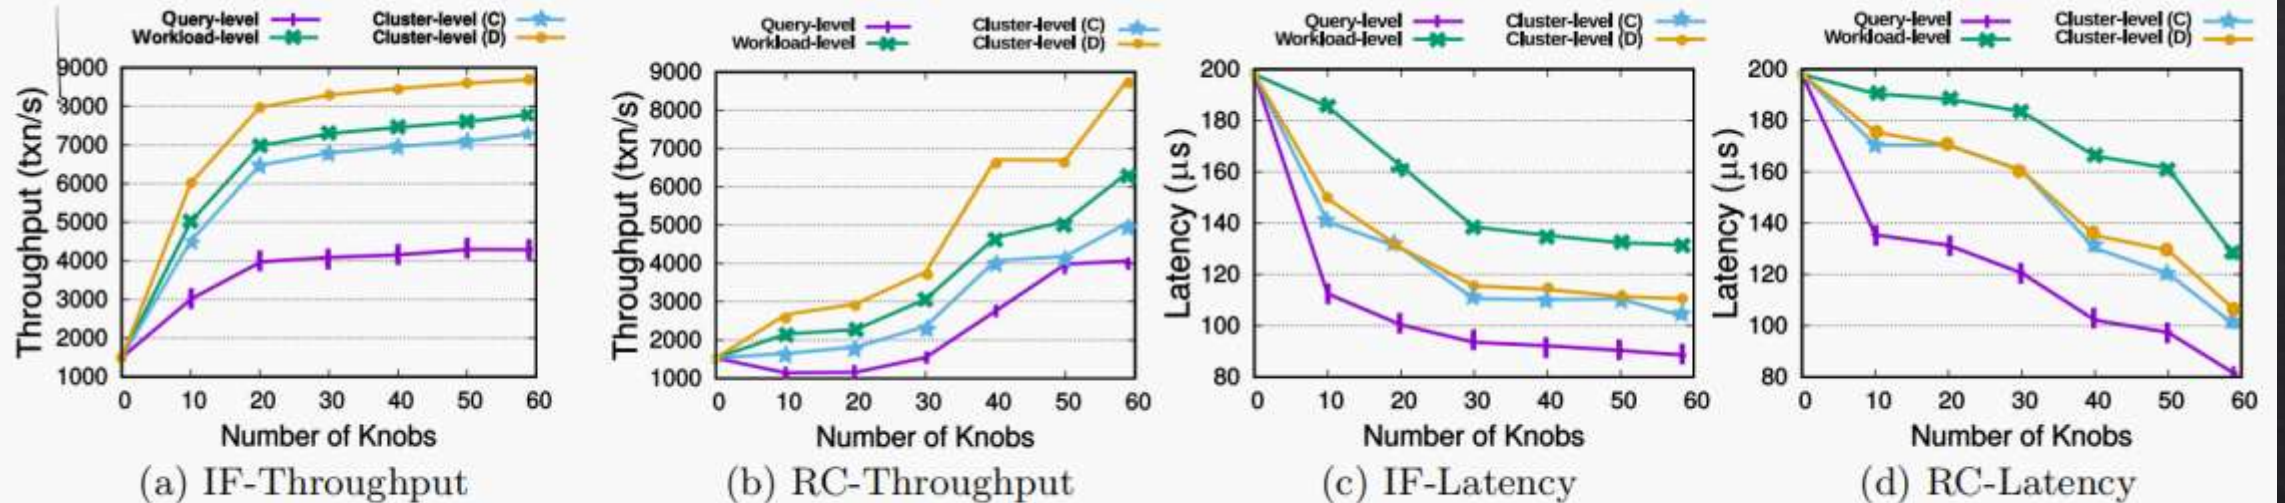
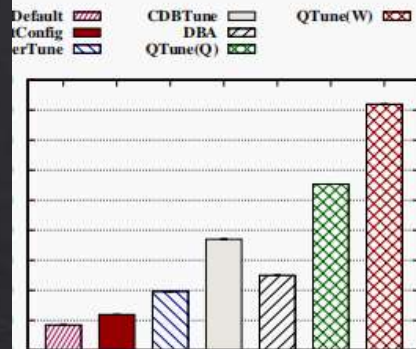



Figure 6: Performance by increasing knobs in Important First (IF) and Randomly Choosing (RC) respectively when running Sysbench (RO) on PostgreSQL.


Comparison with Existing Techniques




(a) Sysbench (RW)




(b) JOB (RO)



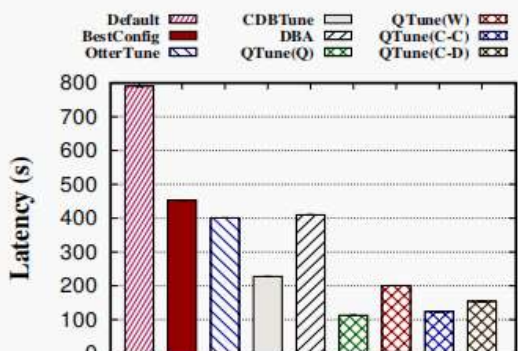
(c) TPC-H (RO)



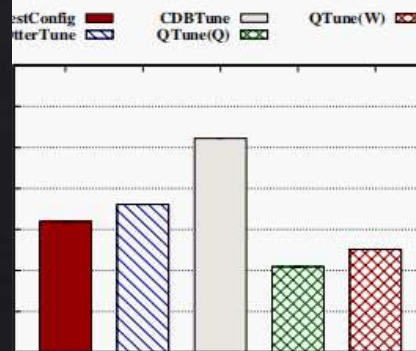
(d) Sysbench (RW)



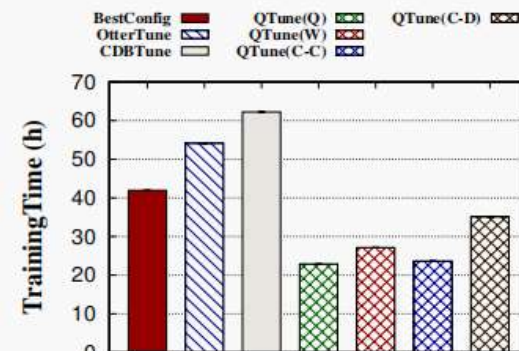
(e) JOB (RO)



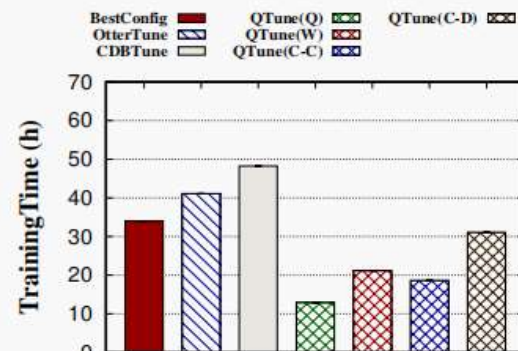
(f) TPC-H (RO)



(g) Sysbench (RW)



(h) JOB (RO)



(i) TPC-H (RO)

Reviews

Pros

- Comprehensive evaluations

Cons

- Feature vectorisation makes the database hard to add and delete future tables
- It is unclear why they did not provide cluster-level evaluations on the Sysbench dataset

Questions?