

# Lecture #19

*Special Topics:*

# Self-Driving Database Management Systems

## Autonomous Systems I

Wan Shen Lim // 15-799 // Spring 2022

# LAST CLASS

---

- High-level NoisePage Pilot architecture.



# TODAY'S AGENDA

---

- UDO: Universal Database Optimization Using Reinforcement Learning. Junxiong Wang, Immanuel Trummer, Debabrota Basu. VLDB 2022.
- Commentary: I think of this as an intelligent configuration batching paper. Doesn't rely on models. Splits parameters into cheap / expensive to change. Introduces new RL algorithm with proofs.


# MOTIVATION

## Cost models

- Require training data.
- Error-prone estimates.

## Rely on sample runs instead

- Regime: high-quality, high-overhead optimization.



# MOTIVATION

---

## Primary use cases

- Tuning when a configuration can be reused over an extended period.
- Analysis tool for other tuning approaches.

## Generalizes to many tuning problems

- e.g., transaction query order, index selection, knob tuning.



# REINFORCEMENT LEARNING

---

## Learning from sample runs

- Classic use case for reinforcement learning.
- **Actions**      The tuning actions we've seen in class.
- **Reward**      The improvement in target metrics.



# REINFORCEMENT LEARNING

---

## Learning from sample runs

- **Algorithm**

- for config in search space

- try next action

- evaluate performance



Insert RL here

- What's the problem?



# REINFORCEMENT LEARNING

---

## Learning from sample runs

- **Algorithm**

- for config in search space

- try next action

- evaluate performance

- What's the problem?

- Trying actions can be expensive!

- High cost per iteration, slow convergence.
  - Key insight: yet not all actions are equally expensive.





# TERMINOLOGY

---

- **Parameter** Each tuning choice.  
Explicitly: a knob, a create index, etc.
- **Configuration** An assignment from parameters to values.
- **Heavy parameter** A parameter which is expensive to change.  
Currently, anything involving physical data structures or database restarts.
- **Light parameter** A parameter which is cheap to change.  
Anything which isn't heavy.

## UDO: KEY IDEAS

---

### **Give heavy parameters special treatment**

- Use RL algorithm that supports delayed rewards.
- e.g., create an expensive index once, batch evaluate similar configurations with that expensive index.
- For each heavy parameter, the optimization of light parameters is a separate Markov decision process.

### **Light parameters, business as usual**

- Use standard no-delay RL.

# UDO: SUPPORTING CAST

---

## Planning component

- Each configuration selected by RL is forwarded here.
- Decides when and in what order to evaluate configs.

## New RL algorithms that accept delayed feedback

- Variant of Monte Carlo Tree Search, delayed-HOO.

# MCTS, delayed-HOO

---

## Monte-Carlo Tree Search

- Recent publicity: AlphaGo, Total War, Tesla Autopilot.
- Game = tree, nodes are states, edges are actions.
  - Selection From root, select children until you reach a leaf.
  - Expansion Expand the leaf (if non-terminal) with children.
  - Simulation Rollout/playout until the game is won/lost.
  - Backprop Update weights.

## Delayed-Hierarchical Optimistic Optimization

- (same authors) [AAAI22] Procrastinated Tree Search.
- Proofs, details, regret bounds, etc. are there.

## FORMAL MODEL: DEFINITIONS

---

- **Parameter** As previously defined.  
Has a **value domain**, e.g., is index built 0/1, query position in txn.
- **Configuration**  $c = \text{vector [ parameter } \rightarrow \text{ value ]}$ .
- **Configuration space**  $C = \{\text{all possible } c\}$ .  
 $C\_H$  : heavy parameters,  $C\_L$ : light parameters,  
 $C = C\_H \times C\_L$ .
- **Benchmark metric**  $f : C \rightarrow \text{real number, stochastic}$ .
- **UDO instance**  $(f, C)$ , find  $c^* = \text{argmax } E[f(c)]$

## FORMAL MODEL: UDO $\rightarrow$ MDPs

---

Given a UDO instance  $(f, C)$ ,  
where the goal is to find  $c^* = \operatorname{argmax} E[f(c)]$ ,  
map  $(f, C)$  to **multiple** episodic Markov decision processes.

**Episodic MDP**  $(S, A, T, R, S_d, S_e)$ , in this case,  
( $S$  state space,  $A$  actions,  $T : S \times A \rightarrow S$  deterministic transition,  
 $R : S \rightarrow \text{real}$  stochastic reward,  
 $S_d$  episode start states,  $S_e$  episode end states)

# FORMAL MODEL: Heavy Parameter MDP

---

## Heavy Parameter MDP

- Each action changes one heavy parameter to a new value.
- Start state is default configuration.
- Reward is max over  $c_L$ ,  $f(c_H \circ c_L) - f(c_{\text{default}})$ .
- End state is all states that are  $N$  actions away (they use  $N=4$ ).

## FORMAL MODEL: Light Parameter MDP

---

**Light Parameter MDP**  $M_L[c_h]$  for each heavy parameter  $c_h$

- Actions are value changes for light parameters.
- End states are a fixed number of light parameter changes.
- Reward is  $f(c_H \circ c_L) - f(c_{\text{default}})$ .





# UDO OVERVIEW

## Iterate until the time limit is reached

- Note other stopping conditions could be used instead.

### Algorithm 1 UDO main function.

```

1: Input: Benchmark metric  $f$ , configuration space  $C$ , RL algorithms
    $\text{Alg}_H$  and  $\text{Alg}_L$  for heavy and light parameter optimization
2: Output: a suggested configuration for best performance
3: function UDO( $f, C, \text{Alg}_H, \text{Alg}_L$ )
4:   // Divide into heavy ( $C_H$ ) and light ( $C_L$ ) parameters
5:    $\langle C_H, C_L \rangle \leftarrow \text{SSA.SPLITPARAMETERS}(C)$ 
6:   // Until optimization time runs out
7:   for  $t \leftarrow 1, \dots, \text{Alg}_H.\text{Time}$  do
8:     // Select next heavy parameter configuration
9:      $c_{H,t} \leftarrow \text{RL.SELECT}(\text{Alg}_H, C_H, c_{H,t-1})$ 
10:    // Submit configuration for evaluation
11:     $\text{EVAL.SUBMIT}(c_{H,t}, t + \text{Alg}_H.\text{maxDelay})$ 
12:    // Receive newly evaluated light configurations
13:     $E \leftarrow \text{EVAL.RECEIVE}(\text{Alg}_L, f, C_L, t)$ 
14:    // Update statistics for heavy parameters
15:     $\text{RL.UPDATE}(\text{Alg}_H, E)$ 
16:   end for
17:   return best obtained configuration
18: end function

```

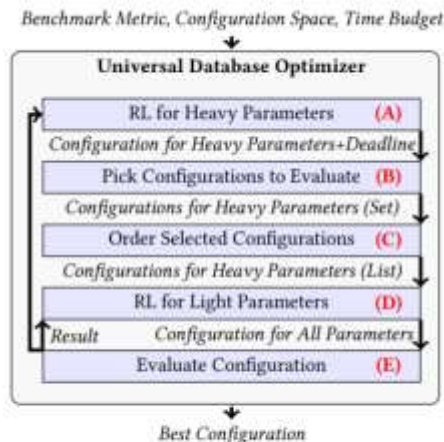


Figure 2: Overview of UDO system (rectangles represent processing steps, arrows represent data flow).

# EVALUATING CONFIGURATIONS: API

## EVAL.Submit(config, deadline)

- Deadline = max additional future configs that can be buffered before this specific config must be evaluated.

## EVAL.Receive(RL algorithm to use, $f$ , light config space, current time)

- Get the next set of evaluated configs.

### Algorithm 2 EVAL: Functions for evaluating configurations.

```

1: // Global variable representing evaluation requests
2:  $R \leftarrow \emptyset$ 

3: Input: heavy configuration  $c_H$  to evaluate and time  $t$ 
4: Effect: adds new evaluation request
5: procedure EVAL.SUBMIT( $c_H, t$ )
6:    $R \leftarrow R \cup \{(c_H, t)\}$ 
7: end procedure

8: Input: RL algorithm  $\text{Alg}_L$ , benchmark metric  $f$ , time  $t$ , and space  $C_L$ 
9: Output: evaluated configurations with reward values
10: function EVAL.RECEIVE( $\text{Alg}_L, f, C_L, t$ )
11:   // Choose configurations from  $R$  to evaluate now
12:    $N \leftarrow \text{PICKCONF}(R, t)$ 
13:   // Remove from pending requests
14:    $R \leftarrow R \setminus N$ 
15:   // Prepare evaluation plan
16:    $P \leftarrow \text{PLANCONF}(N)$ 
17:   // Collect evaluation results by executing plan
18:    $E \leftarrow \emptyset$ 
19:   for  $s \in P.steps$  do
20:     // Prepare evaluation of next configurations
21:      $\text{CHANGECONF}(s.hconf)$ 
22:     // Find (near-)optimal light parameter settings
23:      $c_L \leftarrow \text{RL.OPTIMIZE}(\text{Alg}_L, s.hconf, C_L, f)$ 
24:     // Take performance measurements on benchmark
25:      $b \leftarrow \text{EVALUATE}(f, s.hconf, c_L)$ 
26:     // Add performance result to set
27:      $E \leftarrow E \cup \{(c_L, s.hconf, b)\}$ 
28:   end for
29:   // Return evaluation results
30:   return  $E$ 
31: end function

```

# EVALUATING CONFIGURATIONS: PICKING

## PickConf-Threshold

- If you have “too much” work to do, you must do everything now.

## PickConf-Secretary

- Do everything that must be done.
- Then, secretary problem style, do whatever doesn't require “too much” reconfiguration work.

**Algorithm 3** PICKCONF: Methods for picking configurations to evaluate.

```

1: Input: Evaluation requests  $R$ , current timestamp  $t$ 
2: Output: Set of configurations to evaluate
3: function PICKCONF-THRESHOLD( $R, t$ )
4:   // Was size threshold reached?
5:   if  $|R| \geq \rho$  then
6:     // Return all requests
7:     return  $R$ 
8:   else
9:     return  $\emptyset$ 
10:  end if
11: end function

12: // Initialize maximal cost savings for each request
13:  $S \leftarrow \emptyset$ 

14: Input: Evaluation requests  $R$ , current timestamp  $t$ 
15: Output: Set of configurations to evaluate
16: function PICKCONF-SECRETARY( $R, t$ )
17:   // Add requests whose deadline is reached
18:    $E \leftarrow \{(c_{ij}, t_D) \in R \mid t_D \geq t\}$ 
19:   // Remove requests from pending set
20:    $R \leftarrow R \setminus E$ 
21:   // Iterate over requests
22:   for  $r = (c_H, t_D) \in R$  do
23:     // Calculate re-configuration cost savings
24:      $s \leftarrow \text{COSTSAVINGS}(r, E)$ 
25:     // Retrieve maximal savings so far
26:      $m \leftarrow S(r)$ 
27:     // Should we evaluate?
28:     if  $t - (t_D - \delta) \geq \delta/\epsilon \wedge s > m$  then
29:        $E \leftarrow E \cup \{r\}$ 
30:     end if
31:     // Update maximally possible savings
32:      $S(r) \leftarrow \max(m, s)$ 
33:   end for
34:   return  $E$ 
35: end function

```

# EVALUATING CONFIGURATIONS: ORDERING

## Ordering configurations is NP-HARD

- Hamiltonian graph.

## PlanConf

- Greedy algorithm.

## Integer linear programming

- Optimal solution.

---

**Algorithm 4** PLANCONF: Order configurations for evaluation.

---

```

1: Input: Evaluation requests  $R$ 
2: Output: Requests in suggested evaluation order
3: function PLANCONF-GREEDY( $R$ )
4:   // Initialize list of ordered requests
5:    $O \leftarrow []$ 
6:   // Iterate over all requests
7:   for  $r \in R$  do
8:     // Find optimal insertion point
9:      $i \leftarrow \arg \min_{i \in 0, \dots, |O|} C_R(O[i-1], O[i]) + C_R(O[i], O[i+1])$ 
10:    // Insert current request there
11:     $O.insert(i, r)$ 
12:   end for
13:   return  $O$ 
14: end function

```

---

# REINFORCEMENT LEARNING

Three main subroutines: RL.SELECT, RL.UPDATE, RL.OPTIMIZE.

RL.SELECT	RL.UPDATE	RL.OPTIMIZE
Pick the next action based on some statistics.	Update the statistics used by RL.SELECT.	Invoke the other two repeatedly for optimization.
See paper for details. $c_{t+1} \triangleq \underset{c}{\operatorname{argmax}} \hat{\mu}_c(t) + \sqrt{2.4 \hat{\sigma}_c^2(t) \frac{\log(v_{c_t})}{v_c}} + \frac{3b \log(v_{c_t})}{v_c}$	Update num visits to state-action pairs, present state, sample mean and variance of accumulated rewards.	<p><b>Algorithm 5</b> RL: Monte Carlo Tree Search optimization.</p> <pre> 1: <b>Input:</b> Algorithm Alg, configuration space <math>C</math>, state <math>c_0</math>, benchmark <math>B</math> 2: <b>Output:</b> Final parameter configuration 3: <b>function</b> RL.OPTIMIZE(Alg, <math>C</math>, <math>c_0</math>) 4:   Initialize <math>Stat \leftarrow \emptyset</math> 5:   <b>for</b> <math>t = 0, \dots, \text{Alg.Time}</math> <b>do</b> 6:     <math>\langle c_{t+1}, a_t \rangle \leftarrow \text{RL.SELECT}(\text{Alg}, C, c_t)</math> 7:     Evaluate the new configuration <math>r_t \leftarrow \text{B.EVALUATE}(c_{t+1})</math> 8:     Update <math>Stat \leftarrow Stat \cup \{ \langle c_t, a_t, c_{t+1}, r_t, t \rangle \}</math> 9:     RL.UPDATE(Alg, <math>Stat</math>) 10:   <b>end for</b> 11:   <b>return</b> Final parameter configuration <math>c_T</math> 12: <b>end function</b> </pre>

# THEORY

- Minimizes expected regret.
- See paper for details.
- Extended proofs in [AAAI22] Procrastinated Tree Search.

**THEOREM 6.1 (REGRET OF HOO (THEOREM 6, [9])).** *If the performance metric  $f$  is smooth around the optimal configuration (Assumption 2 in [9]) and the upper confidence bounds on performances of all the configurations at depth  $h$  create a partition shrinking at the rate  $c\rho^h$  with  $\rho \in (0, 1)$  (Assumption 1 in [9]), expected regret of HOO is*

$$\mathbb{E}[\text{Reg}_T] = O\left(T^{1-\frac{1}{d+2}} (\log T)^{\frac{1}{d+2}}\right) \quad (2)$$

for a horizon  $T > 1$ , and  $4/c$ -near-optimality dimension<sup>4</sup>  $d$  of  $f$ .

**THEOREM 6.2 (REGRET OF DELAYED-HOO).** *Under the same assumptions as Thm. 6.1, the expected regret of delayed-HOO is*

$$\mathbb{E}[\text{Reg}_T] = O\left((1+\tau)T^{1-\frac{1}{d+2}} (\log T)^{\frac{1}{d+2}}\right) \quad (3)$$

for delay  $\tau \geq 0$ , horizon  $T$ , and  $4/c$ -near-optimality dimension  $d$  of  $f$ .

**THEOREM 6.3 (REGRET OF UDO).** *If we use the delayed-HOO as the delayed-MCTS algorithm with delays  $\tau$  and 0, and time-horizons  $T_h$  and  $T_l$  for heavy and light parameters respectively, the expected regret of UDO is upper bounded by*

$$\mathbb{E}[\text{Reg}_T] = O\left((1+\tau)T_h^{1-\frac{1}{d+2}} (\text{HOO}^2(T_l) \log T_h)^{\frac{1}{d+2}}\right), \quad (4)$$

under the assumptions of Thm. 6.1. Here,  $\text{HOO}(T_l) \triangleq O\left([\log T_l/T_l]^{\frac{1}{d+2}}\right)$ .

Deviation in expected performance of the configuration returned by UDO from the optimum is  $O\left((1+\tau) [\text{HOO}^2(T_l)\text{HOO}(T_h)]^{\frac{1}{d+2}}\right)$ . Here,  $T_h$  and  $T_l$  are the number of steps allotted for the heavy and light parameters respectively. Deviation in expected performance of the configuration selected by UDO vanishes as  $T_h, T_l \rightarrow \infty$ .

# EXPERIMENTS: SETUP

---

## Hardware

- Server, 2x Intel Xeon Gold 5218, 2.3 GHz, 32 physical cores.
- 384 GB RAM.
- 1 TB HDD.

## DBMSs

- MySQL 5.7.29.
- PostgreSQL 10.15.



# EXPERIMENTS: SETUP

---

## UDO

- Delay = 10 for heavy MDP.
- $b = 3$  in UCB-V (RL.SELECT picking the next action).
- Per episode,
  - Up to 8 actions for TPC-H.
  - Up to 13 actions for TPC-C (four heavy parameter changes).



# EXPERIMENTS: WORKLOADS

---

## Workloads

- TPC-C (SF 10, 32 terminals), maximize throughput.
  - Reload snapshot every 10 iterations of main loop.
  - Standard mix for 5 seconds.
  - Parameters: 71 index, 19 reorder, 10 MySQL / 15 PostgreSQL knobs.
- TPC-H (SF 1), minimize latency.
  - Parameters: 99 index, 10 MySQL / 15 PostgreSQL knobs.

# EXPERIMENTS: IMPLEMENTATION

---



## UDO

- Python3 + OpenAI gym
- Gurobi for cost-based planning

## Baselines (targeted at no prior training data scenario)



- For RL comparisons, against Keras-RL's SARSA, DDPG.
- Some combination of MySQL-Tuner, PGTuner, Gaussian Process Regression, DDPG++, Quro, Dexter, EverSQL.
- When combining, optimize transaction code, then parameters, then index selection.

# UDO vs BASELINES



(a) TPC-C performance as a function of optimization time in MySQL. (b) TPC-C performance as a function of optimization time in Postgres.



**Figure 3: Comparing UDO to baselines on TPC-C.**




(a) TPC-H performance as a function of optimization time in MySQL. (b) TPC-H performance as a function of optimization time in Postgres.

**Figure 4: Comparing UDO to baselines on TPC-H.**



# UDO vs BASELINES




(a) TPC-C performance as a function of optimization time in MySQL.



(b) TPC-C performance as a function of optimization time in PostgreSQL.



(a) TPC-H performance as a function of optimization time in MySQL.



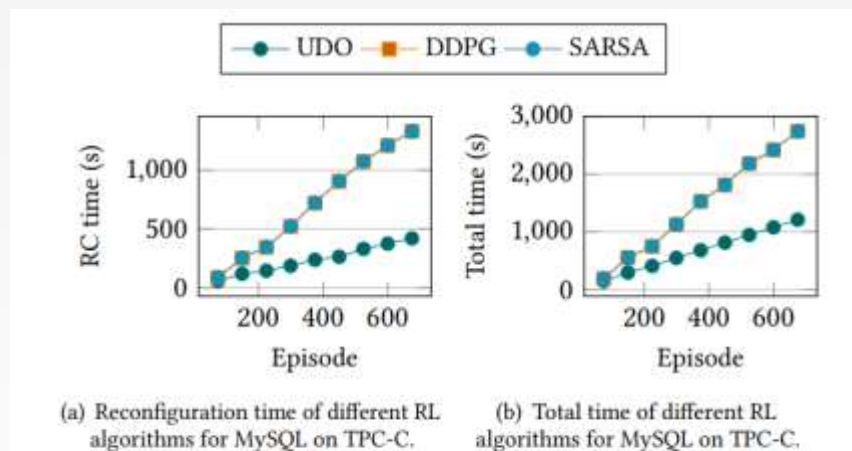
(b) TPC-H performance as a function of optimization time in PostgreSQL.

Figure 3: Comparing UDO to baselines on TPC-C.


Figure 3: Comparing UDO to baselines on TPC-H.

UDO is always the best  
Followed by DDPG++ with Dexter

# UDO vs BASELINES

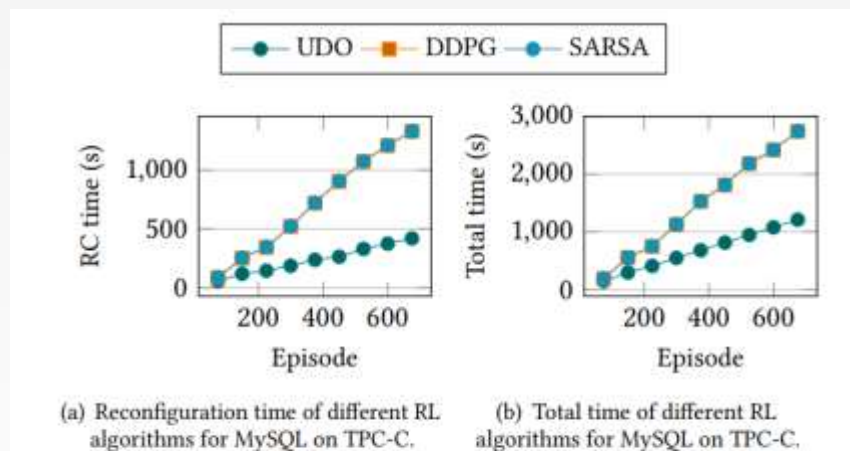


**Figure 5: Time spent per episode by different RL algorithms when optimizing MySQL for TPC-C.**




**Figure 6: Time spent per episode by different RL algorithms when optimizing Postgres for TPC-H.**

# UDO vs BASELINES




**Figure 5: Time spent per episode by different RL algorithms when optimizing MySQL for TPC-C.**




**Figure 6: Time spent per episode by different RL algorithms when optimizing Postgres for TPC-H.**

UDO can reduce reconfiguration time by a factor of 3


# UDO VARIANTS



**Figure 7: Impact of delayed feedback on UDO performance (MySQL on TPC-C).**




**Figure 8: Impact of evaluation time selection on UDO performance (MySQL on TPC-C).**



(a) Time spent in plan optimization.

(b) Time spent in reconfiguration.


**Figure 9: Impact of reconfiguration planning algorithm on UDO performance (MySQL on TPC-C).**




**Figure 10: Impact of search space design and search strategy on UDO performance (MySQL on TPC-C).**




# UDO VARIANTS



**Figure 7: Impact of delayed feedback on UDO performance (MySQL on TPC-C).**




**Figure 8: Impact of evaluation time selection on UDO performance (MySQL on TPC-C).**



(a) Time spent in plan optimization.

(b) Time spent in reconfiguration.


**Figure 9: Impact of reconfiguration planning algorithm on UDO performance (MySQL on TPC-C).**



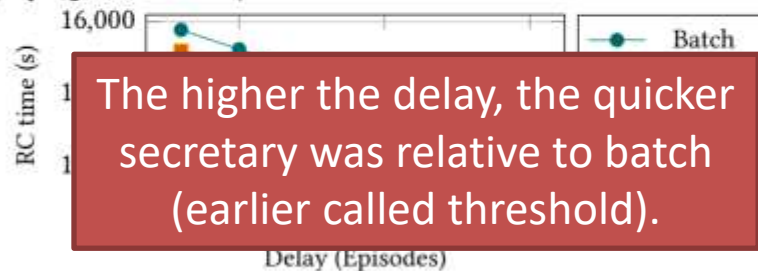
**Figure 10: Impact of search space design and search strategy on UDO performance (MySQL on TPC-C).**




# UDO VARIANTS



**Figure 7: Impact of delayed feedback on UDO performance (MySQL on TPC-C).**




**Figure 8: Impact of evaluation time selection on UDO performance (MySQL on TPC-C).**




(a) Time spent in plan optimization. (b) Time spent in reconfiguration.

**Figure 9: Impact of reconfiguration planning algorithm on UDO performance (MySQL on TPC-C).**

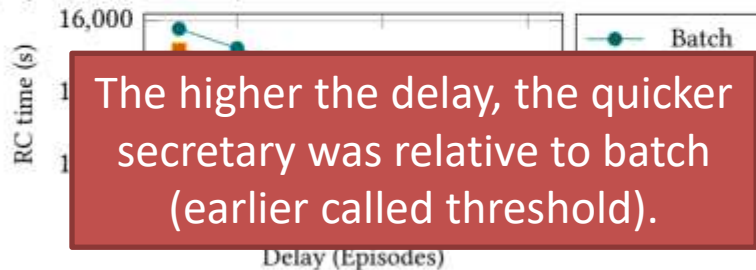


**Figure 10: Impact of search space design and search strategy on UDO performance (MySQL on TPC-C).**

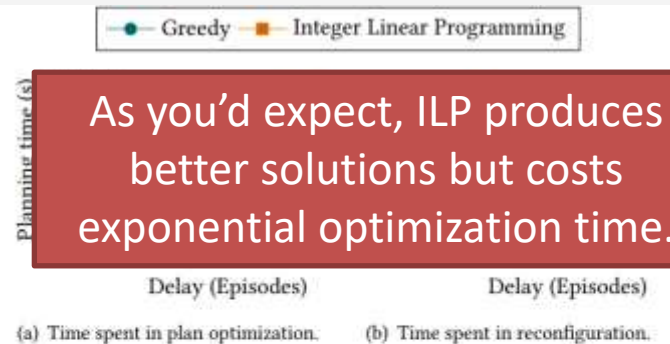
# UDO VARIANTS




**Figure 7: Impact of delayed feedback on UDO performance (MySQL on TPC-C).**



**Figure 8: Impact of evaluation time selection on UDO performance (MySQL on TPC-C).**




**Figure 9: Impact of reconfiguration planning algorithm on UDO performance (MySQL on TPC-C).**

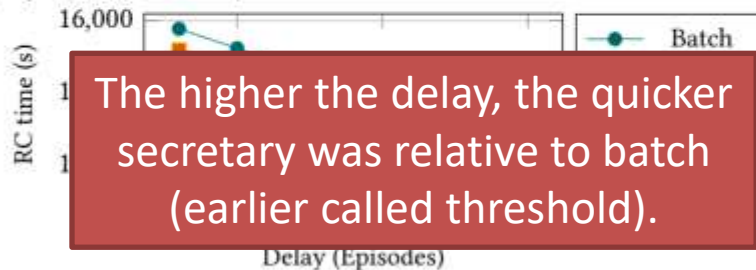


**Figure 10: Impact of search space design and search strategy on UDO performance (MySQL on TPC-C).**

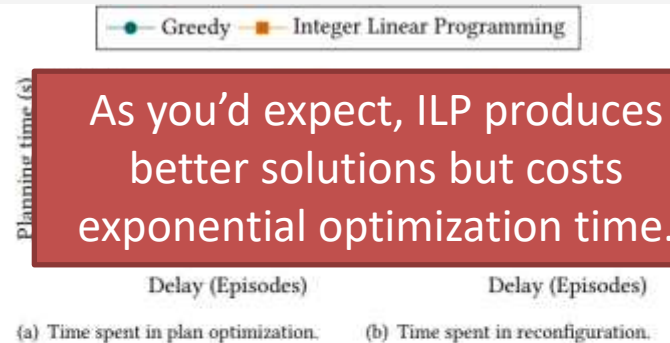
# UDO VARIANTS




**Figure 7: Impact of delayed feedback on UDO performance (MySQL on TPC-C).**



**Figure 8: Impact of evaluation time selection on UDO performance (MySQL on TPC-C).**




**Figure 9: Impact of reconfiguration planning algorithm on UDO performance (MySQL on TPC-C).**



**Figure 10: Impact of search space design and search strategy on UDO performance (MySQL on TPC-C).**

# SCENARIO VARIANTS



(a) TPC-H performance as a function of optimization time in MySQL. (b) TPC-H performance as a function of optimization time in Postgres.

Figure 11: Comparing UDO to baselines on TPC-H for SF 10.

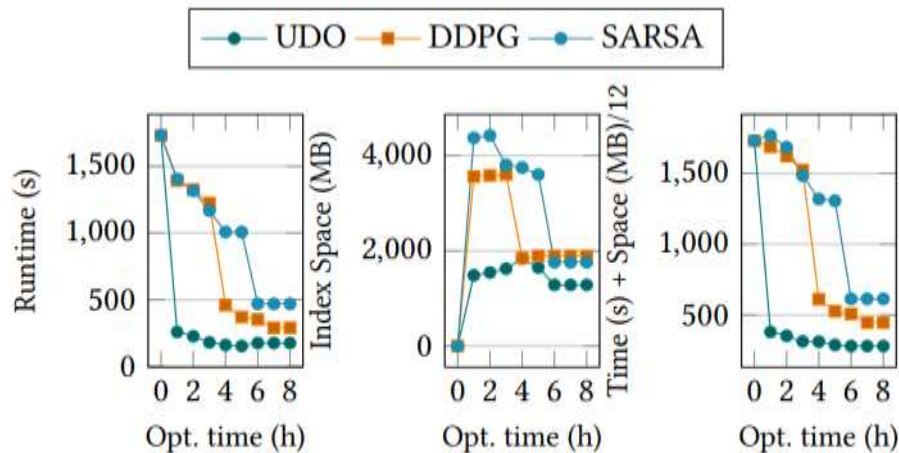
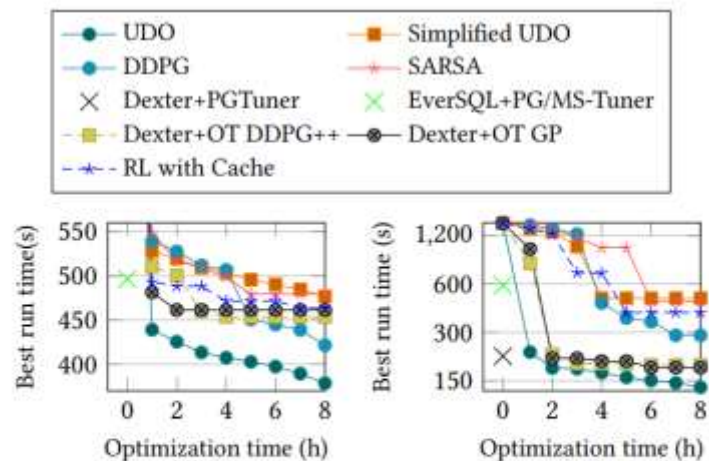


Figure 12: Optimizing weighted sum of run time and disk space for TPC-H SF 10 on Postgres.

# SCENARIO VARIANTS



(a) TPC-H performance as a function of optimization time in MySQL. (b) TPC-H performance as a function of optimization time in Postgres.

Figure 11: Comparing UDO to baselines on TPC-H for SF 10.

Higher scalefactor, similar trends.





Figure 12: Optimizing weighted sum of run time and disk space for TPC-H SF 10 on Postgres.



# SCENARIO VARIANTS



(a) TPC-H performance as a function of optimization time in MySQL. (b) TPC-H performance as a function of optimization time in Postgres.

Figure 11: Comparing UDO to baselines on TPC-H for SF 10.

Higher scalefactor, similar trends.

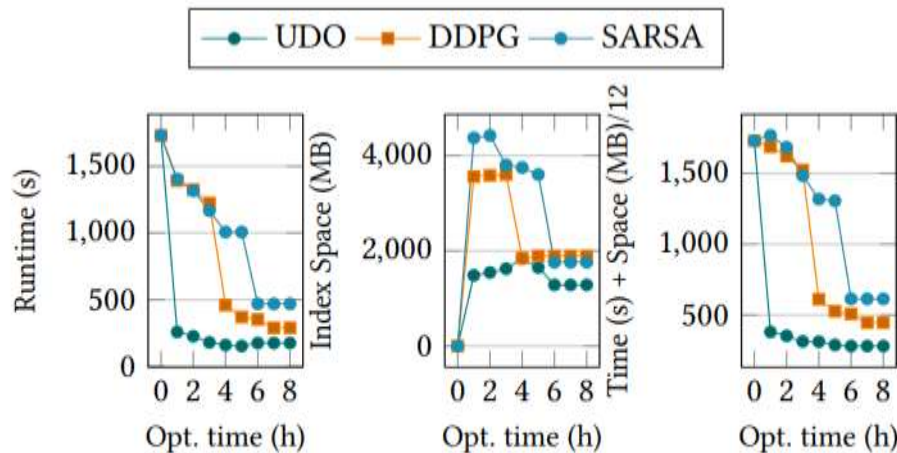



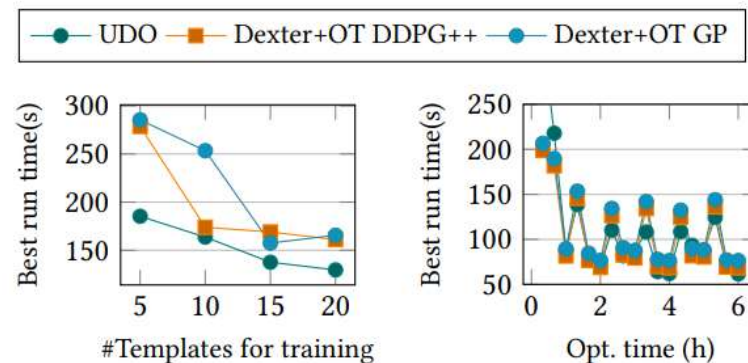
Figure 12: Optimizing weighted sum of run time and disk space for TPC-H SF 10 on Postgres.

Multi-objective optimization.

# SCENARIO VARIANTS




**Figure 13: Comparing UDO to baselines for index recommendation (TPC-H SF 10).**




**Figure 14: Performance for non-representative training sets and changing workloads (TPC-H SF 10, Postgres).**

# SCENARIO VARIANTS



**Figure 13: Comparing UDO to baselines for index recommendation (TPC-H SF 10).**




**Figure 14: Performance for non-representative training sets and changing workloads (TPC-H SF 10, Postgres).**

Restricted to just indexes, still good.




# SCENARIO VARIANTS



**Figure 13: Comparing UDO to baselines for index recommendation (TPC-H SF 10).**


Restricted to just indexes, still good.



**Figure 14: Performance for non-representative training sets and changing workloads (TPC-H SF 10, Postgres).**

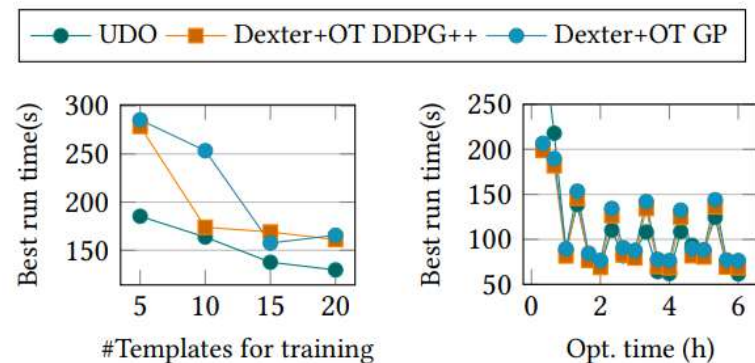
Workload  
shift is  
difficult.

# SCENARIO VARIANTS



**Figure 13: Comparing UDO to baselines for index recommendation (TPC-H SF 10).**

Restricted to just indexes, still good.



**Figure 14: Performance for non-representative training sets and changing workloads (TPC-H SF 10, Postgres).**

Workload  
shift is  
difficult.

Pingpong  
between even  
/ odd TPC-H.

## PARTING THOUGHTS

---

Parameters are not equal cost.

- Batch light parameters, multiple MDPs.
- delayed-HOO to account for delayed rewards.

Thoughts and commentary.

- Good use of both DBMS and RL domain knowledge; fig 10 cautionary of xkcd1838.
- *Universal*, counterexamples?
- Where do cost models still play a role? Other parts of the quality/overhead regime?