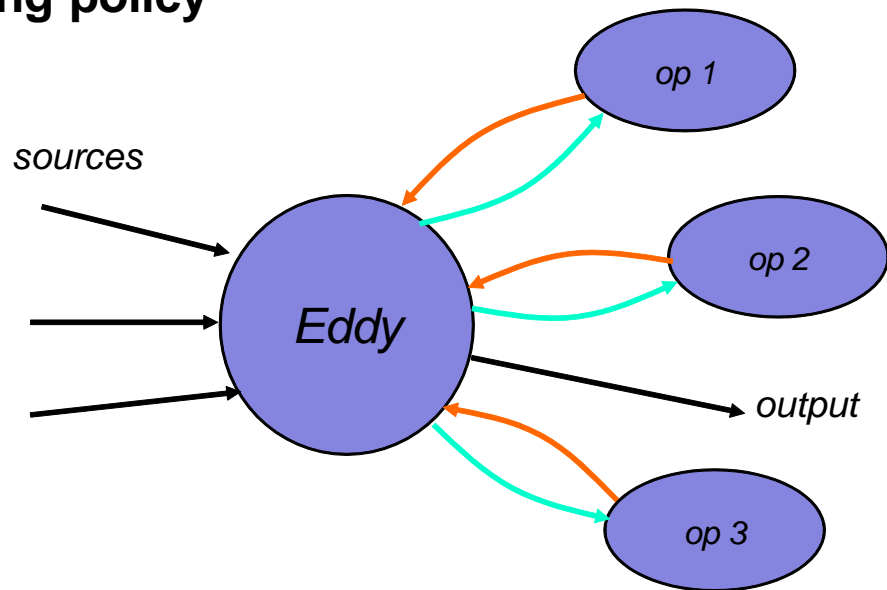


# A Reinforcement Learning Approach for Adaptive Query Processing

Kostas Tzoumas, Timos Sellis, Christian S. Jensen

## AQP with eddies

- Decides an operator to route an incoming tuple
- Uses additional meta-data for correctness (tuple descriptor)
- Uses a routing policy




## Problem statement

- The routing policy component
  - random, lottery scheduling, rank ordering, ...
- How does query optimization enter the picture?
  - Beginning with zero knowledge about the environment and the data
  - Optimization
    - find an optimal query plan
  - Adaptation
    - change it when the environment changes
  - It is a learning problem

## Reinforcement Learning

receive state  $s_t$   
 choose an action  $\alpha_t$   
 execute  $\alpha_t$   
 receive next state  $s_{t+1}$   
 receive reward  $r_{t+1}$   
 update Q values



the Q values  $Q(s, \alpha)$  state space  $S$ , action spaces  $A(s)$   
 $\alpha = \text{chooseAction}(Q, s)$   
 reward function  $r(s, \alpha)$   
 $Q(s, \alpha) = \text{update}(Q, r, s')$

## Choosing and updating

$\alpha = \text{chooseAction}(Q, s)$

- random
  - choose a random action
- greedy
  - choose the action with the highest  $Q(s, \alpha)$
- $\epsilon$ -greedy
  - greedy with prob.  $1-\epsilon$
  - random with prob.  $\epsilon$
- uniform
  - prob. to choose  $\alpha$  proportional to  $Q(s, \alpha)$
- simulated annealing
  - prob. proportional to  $\exp(Q(s, \alpha) / T)$

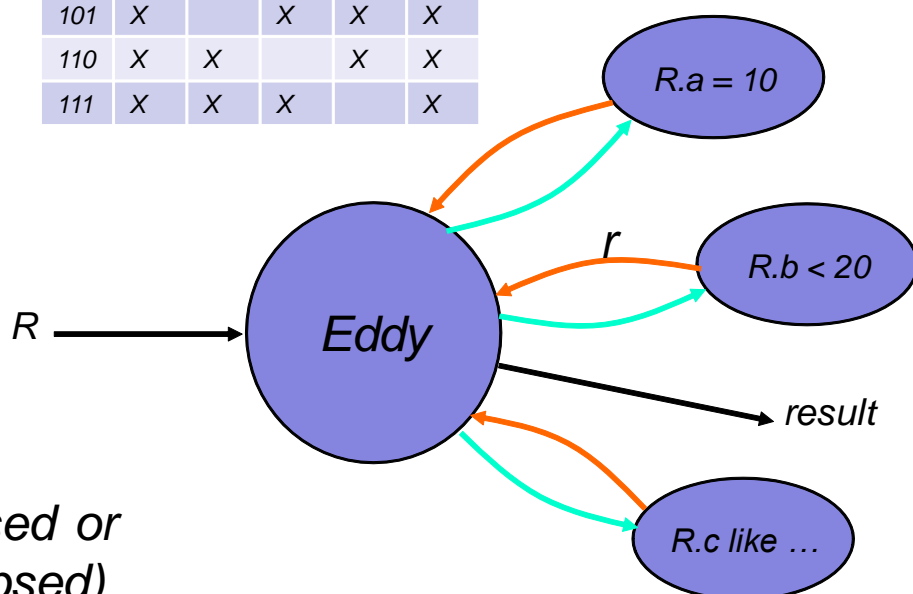
$Q(s, \alpha) = \text{update}(Q, r, s')$

- Monte Carlo
  - average of rewards
  - $Q = Q + 1/n(r - Q)$
- constant- $\alpha$  MC
  - weighted average
  - $Q = Q + \alpha(r - Q)$
- Q
  - bootstrapping
  - $Q = Q + \alpha[r + \gamma \max_{\alpha'} Q(s', \alpha') - Q]$

## Selections

receive state  $s=000$   
 $A(s) = \{\sigma_1, \sigma_2, \sigma_3\}$   
 choose action  $\alpha = \sigma_2$   
 send tuple to  $\sigma_2$   
 wait for  $\sigma_2$  to return  
 next state  $s'=010$   
 receive reward  $r$   
 update Q values

| Q    | $\sigma_1$ | $\sigma_2$ | $\sigma_3$ | out | get |
|------|------------|------------|------------|-----|-----|
| null | X          | X          | X          | X   | X   |
| 000  |            |            |            | X   | X   |
| 001  |            |            | X          | X   | X   |
| 010  | X          | X          | X          | X   | X   |
| 011  |            | X          | X          | X   | X   |
| 100  | X          |            |            | X   | X   |
| 101  | X          |            | X          | X   | X   |
| 110  | X          | X          |            | X   | X   |
| 111  | X          | X          | X          |     | X   |



$r = (20, 13, "abc")$   
 reward = -1 or 0  
 reward = - time elapsed or  $-\epsilon(\text{time elapsed})$

## Binary Symmetric Hash Joins

receive state  $s=S$   
 $A(s) = \{R \blacktriangleright \blacktriangleleft S, S \blacktriangleright \blacktriangleleft T\}$   
 choose action  $\alpha = R \blacktriangleright \blacktriangleleft S$   
 route tuple and push  $R \blacktriangleright \blacktriangleleft S$  in a stack  
 pop intermediate tuple  
 next state  $s'=RS$   
 receive reward  $r$   
 update Q values  
 $s=s'$   
 repeat

| Q    | $\blacktriangleright \blacktriangleleft 1$ | $\blacktriangleright \blacktriangleleft 2$ | out | get R | get S | get T |
|------|--|--|-----|-------|-------|-------|
| null | X  | X  | X   |       |       |       |
| R    |  | X  | X   | X     | X     | X     |
| S    | X  |  | X   | X     | X     | X     |
| T    | X  |  | X   | X     | X     | X     |
| RS   | X  |  | X   | X     | X     | X     |
| ST   |  | X  | X   | X     | X     | X     |
| RST  | X  | X  |     | X     | X     | X     |

