

# A Spark Optimizer For Adaptive, Fine-Grained Parameter Tuning

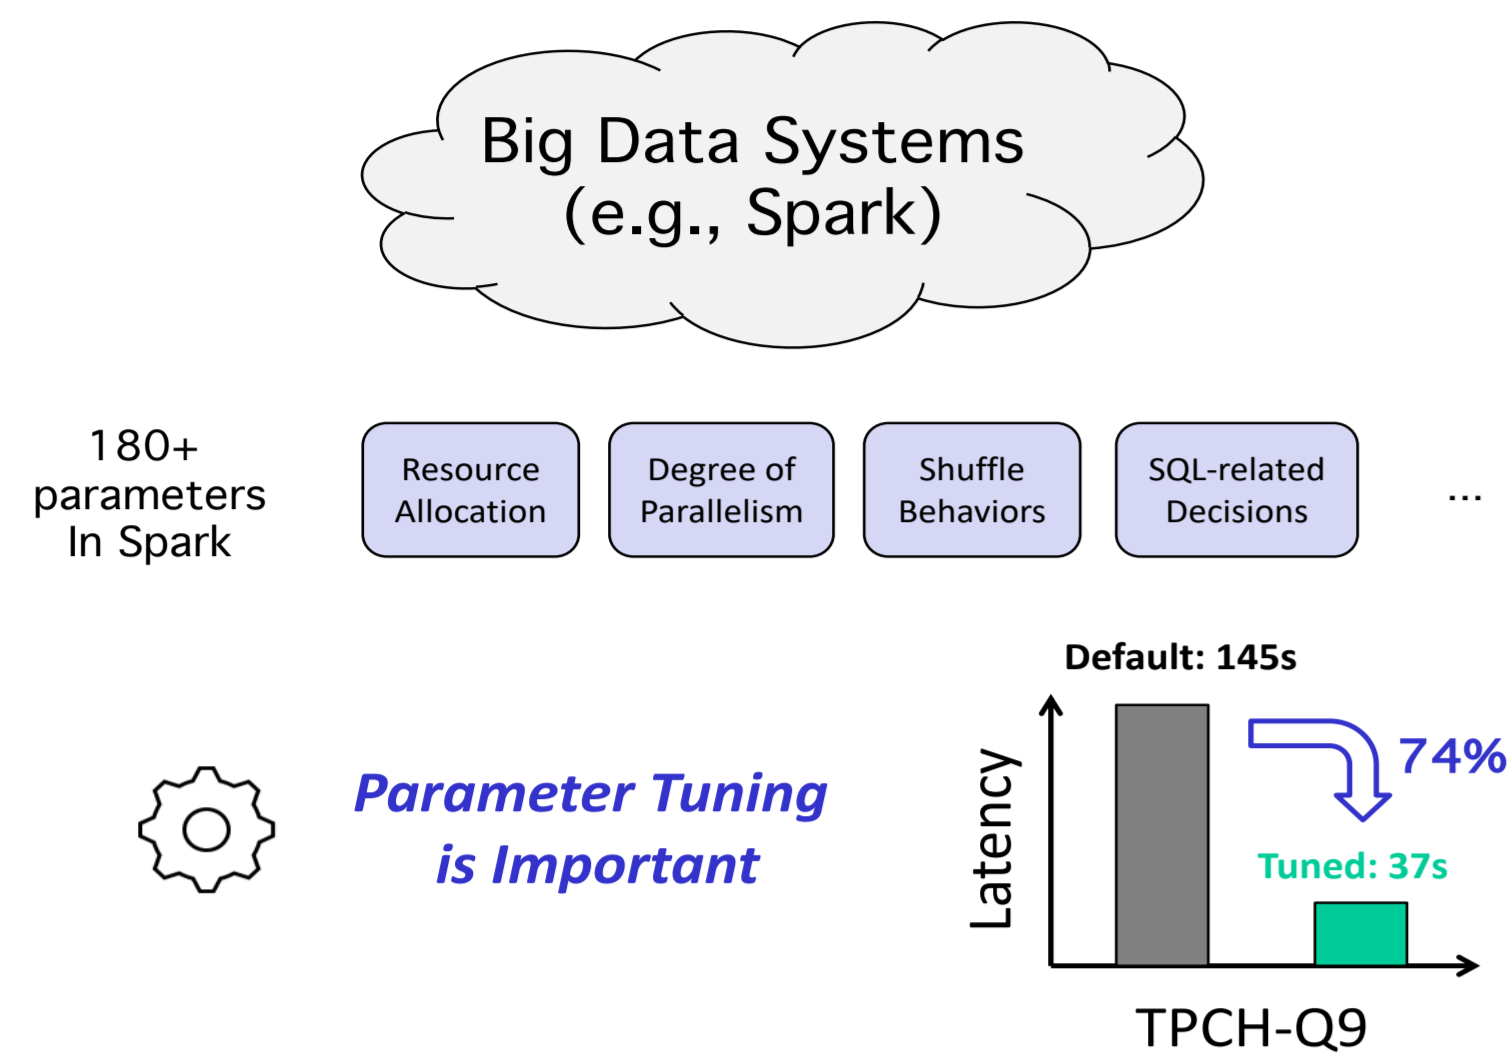


Chenghao Lyu<sup>1</sup>, Qi Fan<sup>2</sup>, Philippe Guyard<sup>2</sup>, Yanlei Diao<sup>1,2</sup>,  
<sup>1</sup>University of Massachusetts Amherst, <sup>2</sup>Ecole Polytechnique,

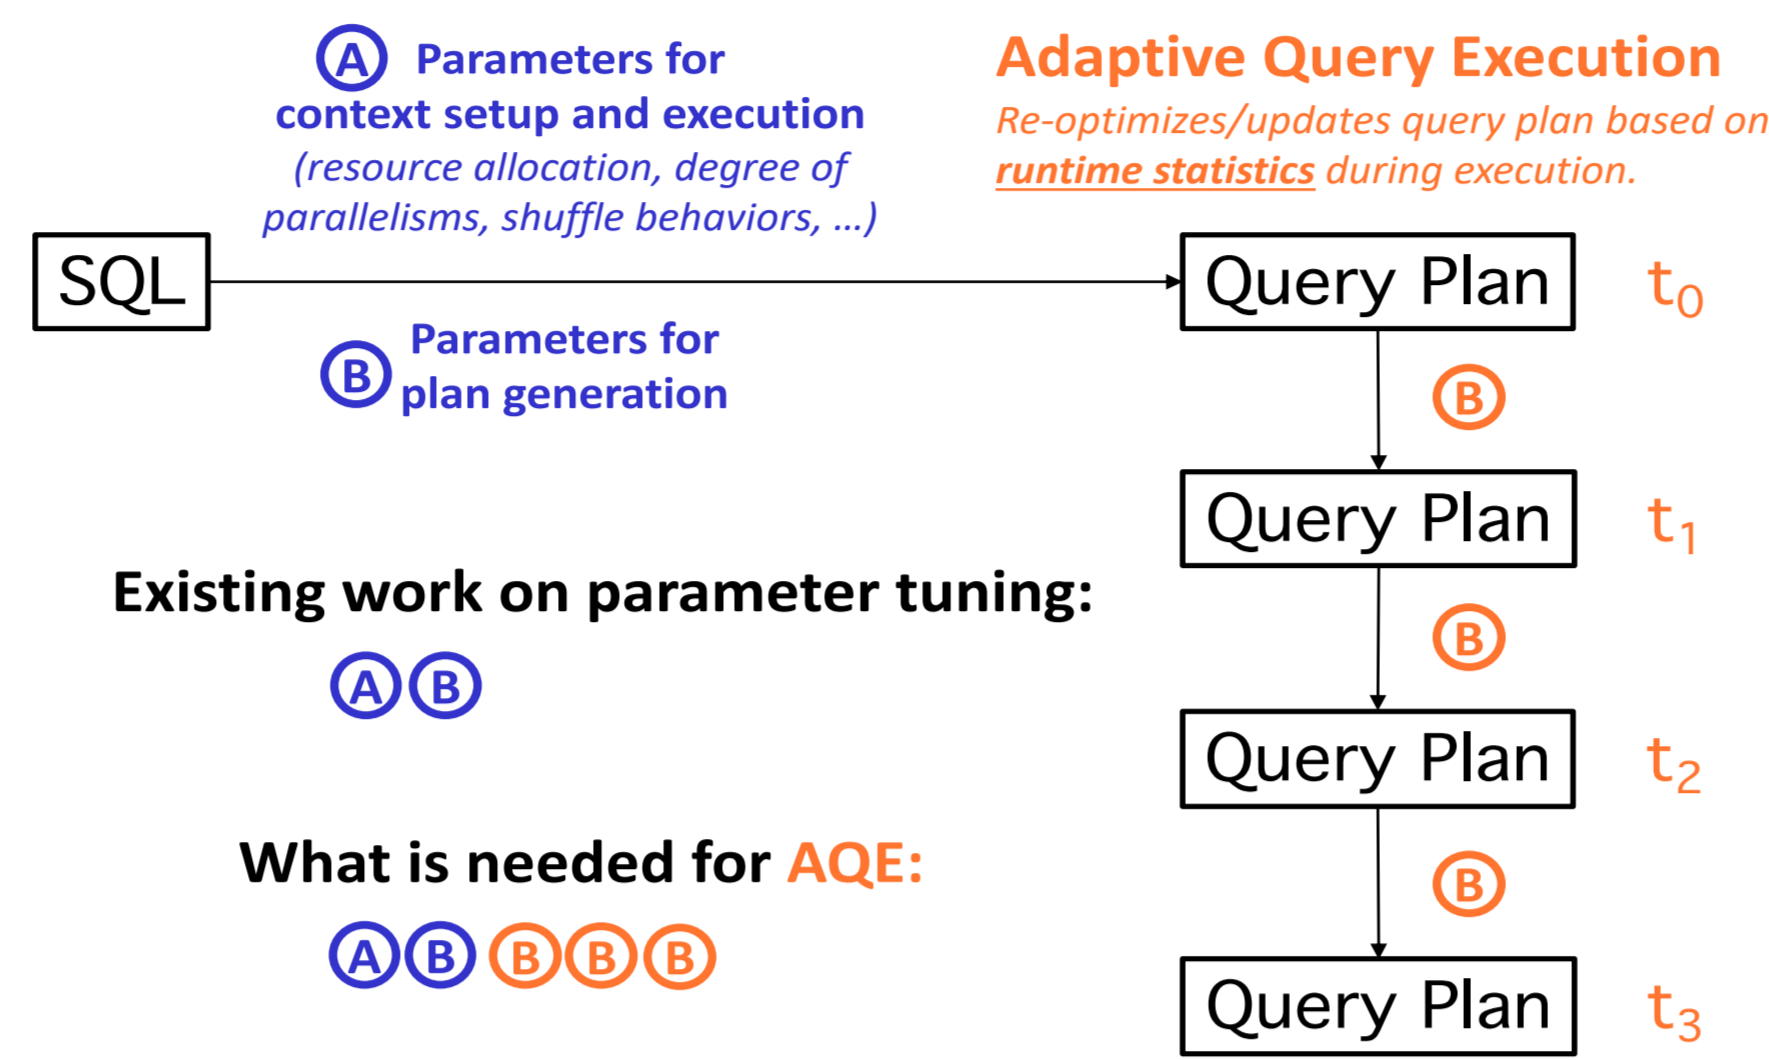


Paper Link

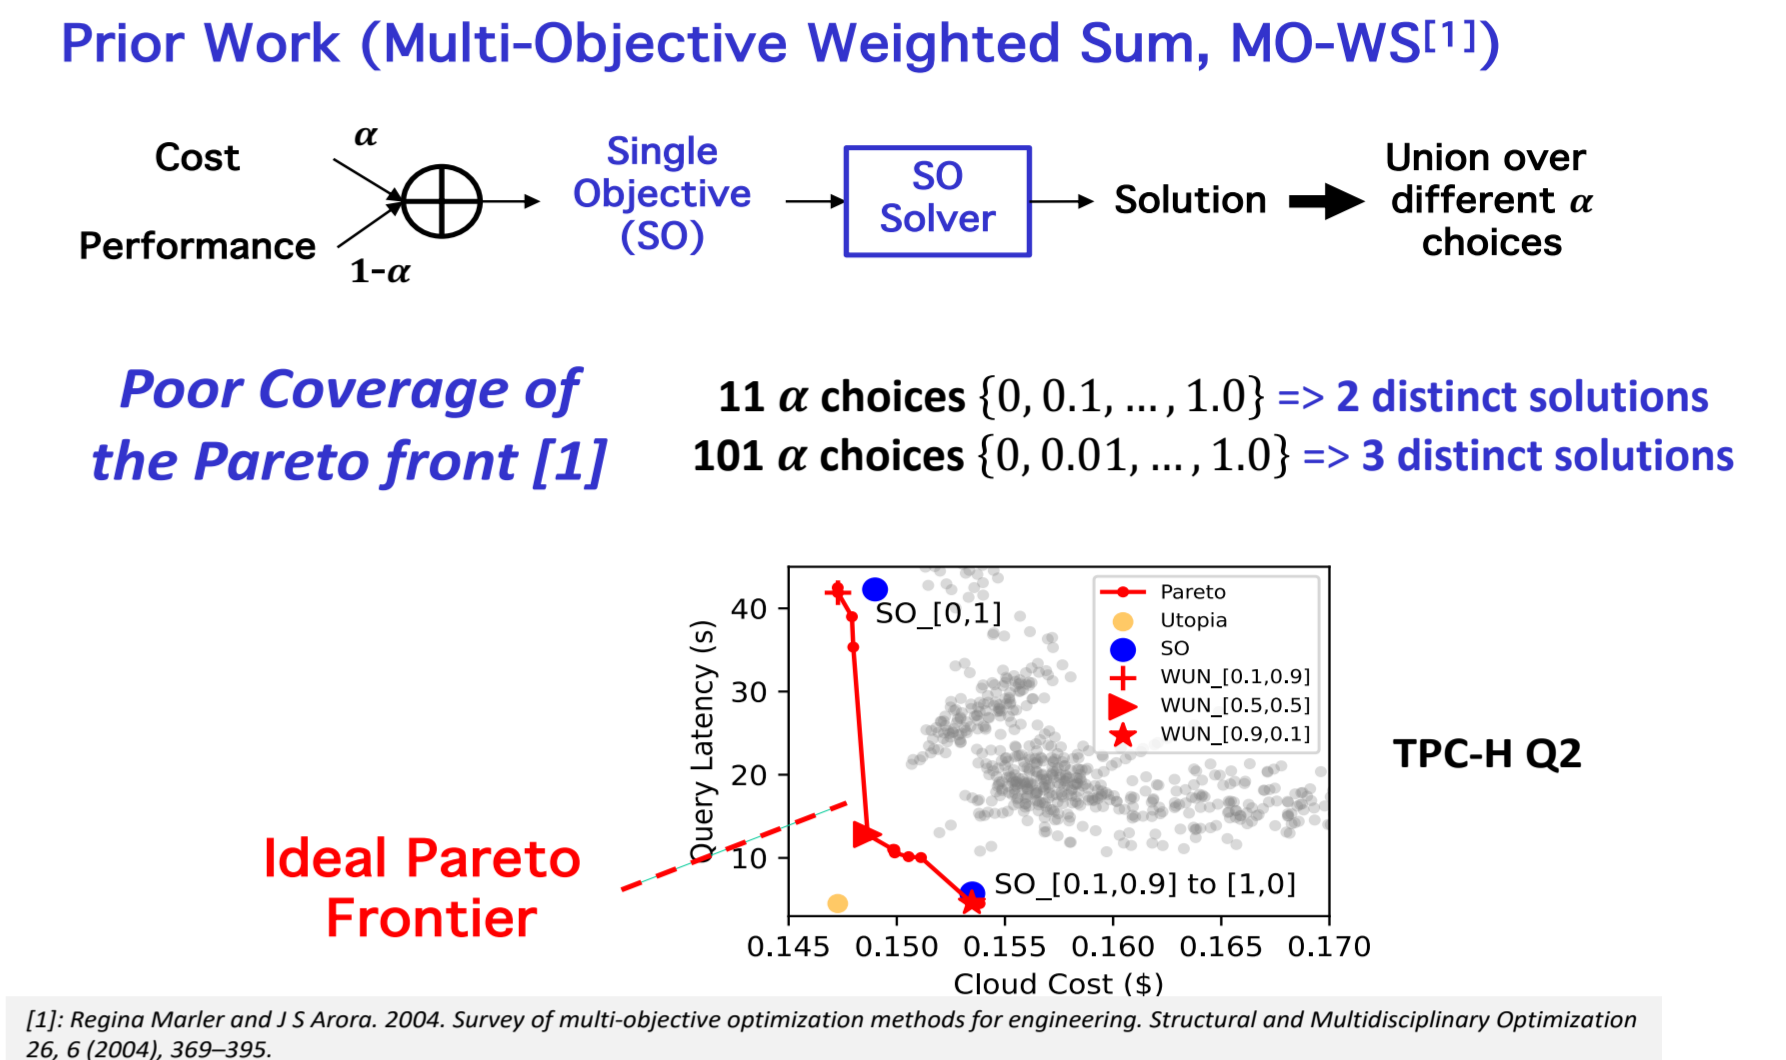
## Big Data Query Processing



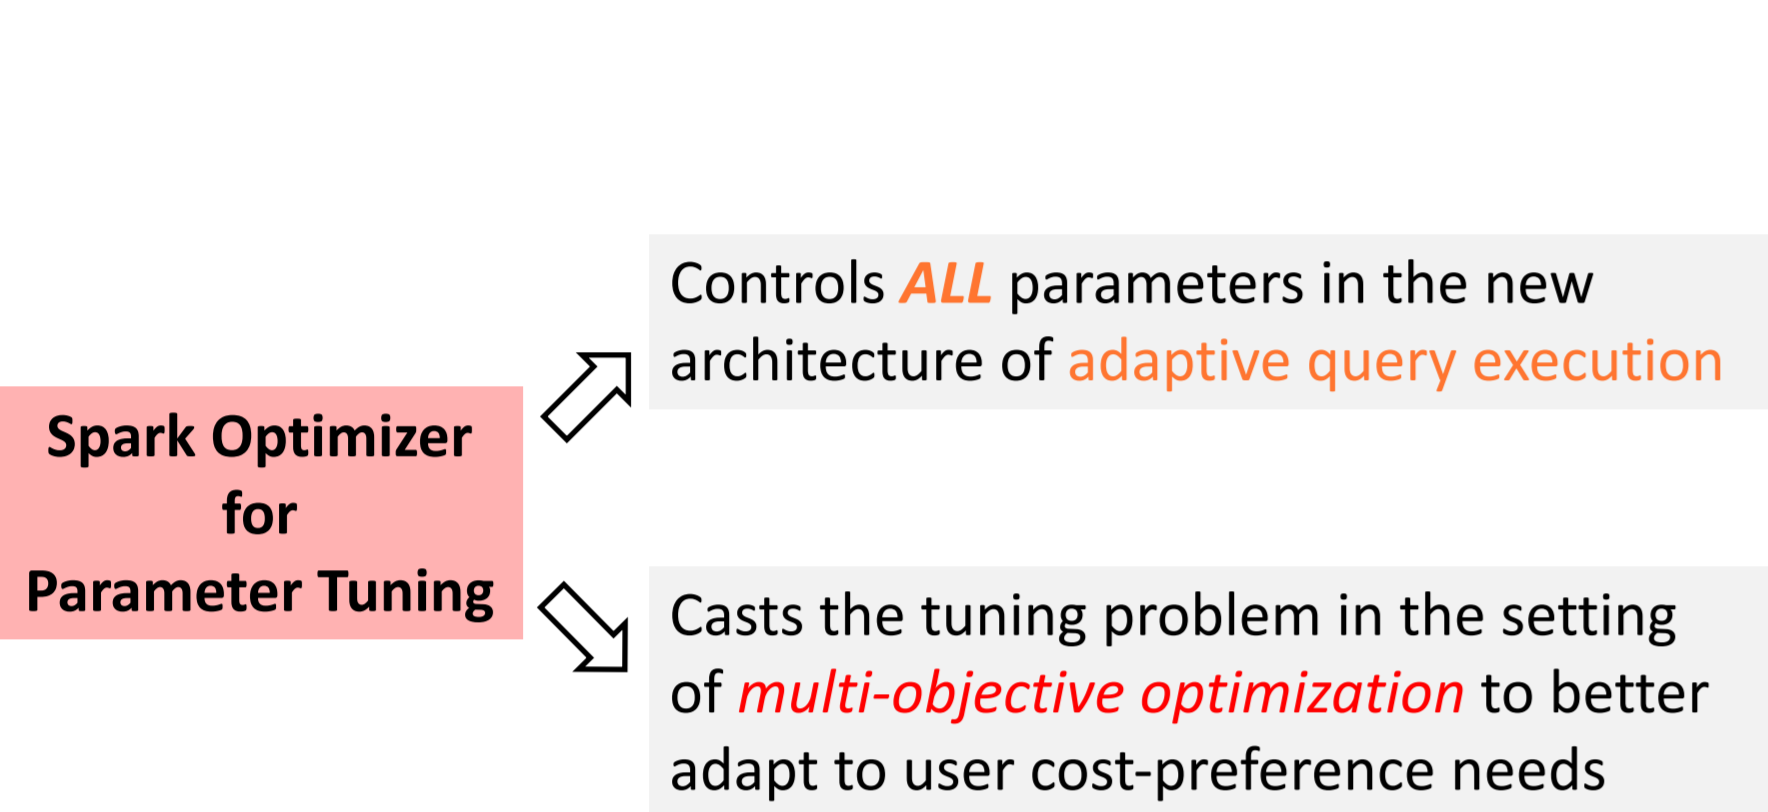
## Trend 1: Adaptive Query Execution (AQE)



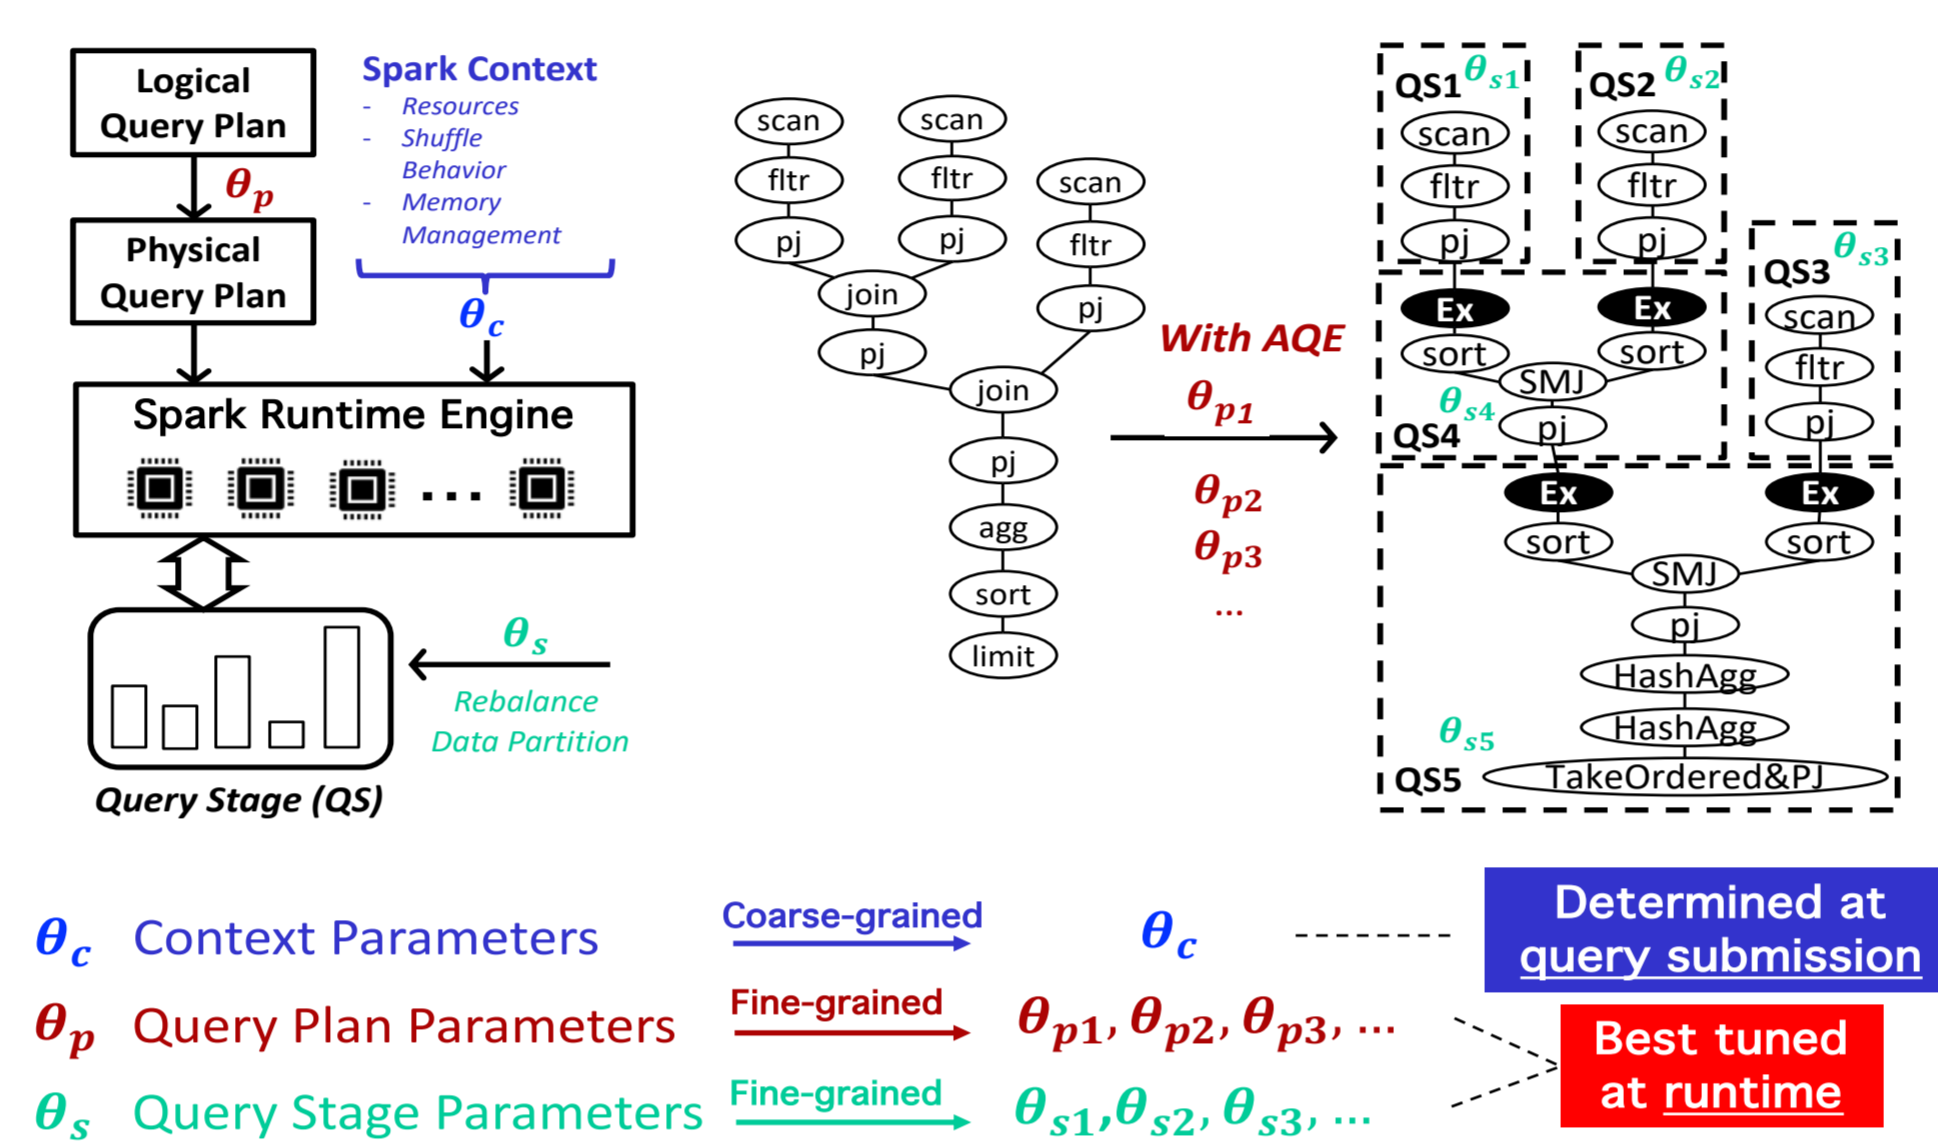
## Trend 2: Cost Performance Reasoning



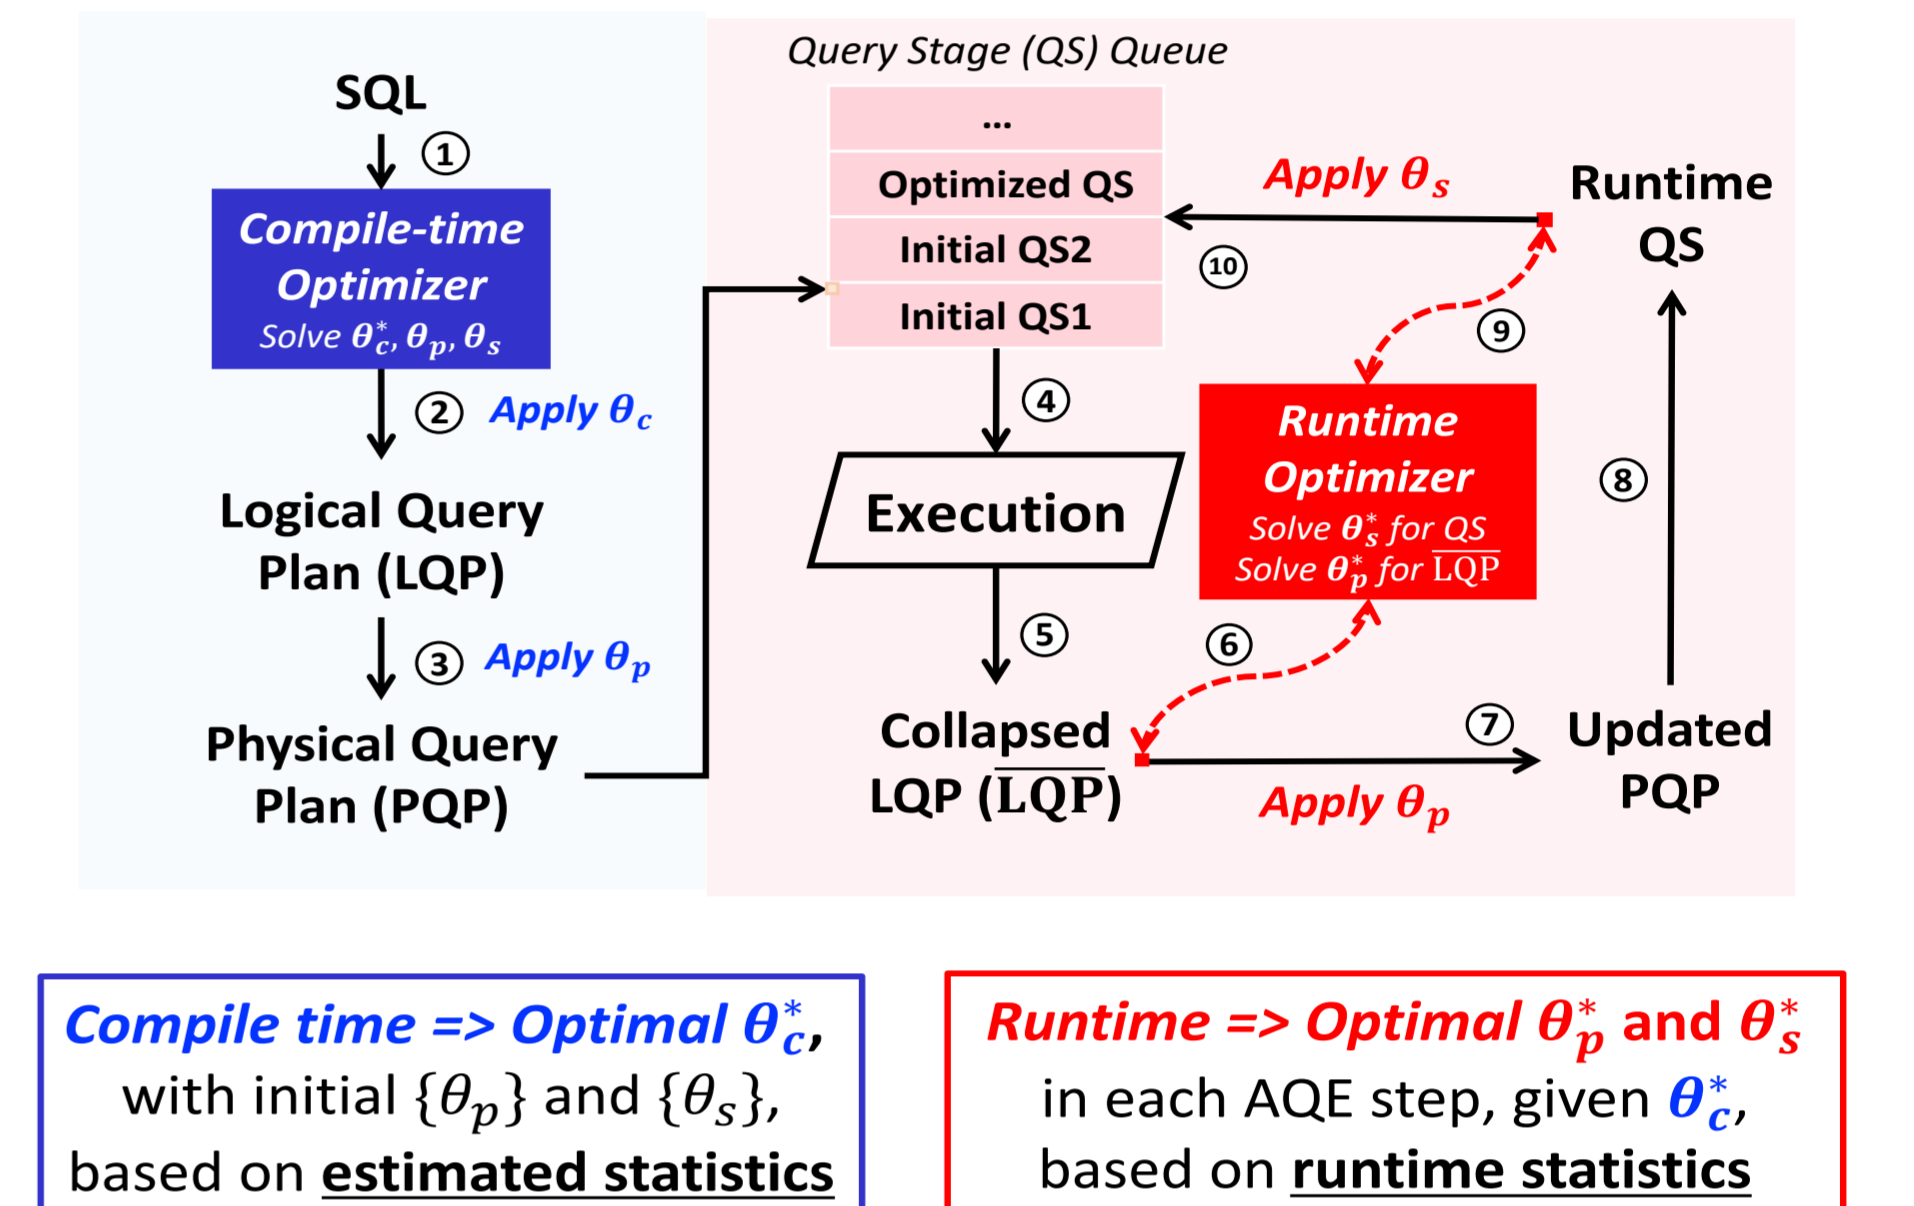
## GOAL



## Challenge 1: Complex Parameter Control



## Solution: Hybrid Multi-Granularity Tuning



## Challenge 2: High Overhead of Optimization

**Definition 3.3. Multi-Objective Optimization for Spark SQL**

$$\arg \min_{\theta_c, \{\theta_p\}, \{\theta_s\}} f(\theta_c, \{\theta_p\}, \{\theta_s\}) = \begin{bmatrix} f_1(LQP, \theta_c, \{\theta_p\}, \{\theta_s\}, \alpha, \beta, \gamma) \\ \dots \\ f_k(LQP, \theta_c, \{\theta_p\}, \{\theta_s\}, \alpha, \beta, \gamma) \end{bmatrix}$$

$k$  objectives

$\theta_c \in \Sigma_c$ ,  $\{\theta_p\} = \{\theta_{p1}, \theta_{p2}, \dots, \theta_{pt}, \dots\}, \forall \theta_{pt} \in \Sigma_p$ ,  $\{\theta_s\} = \{\theta_{s1}, \theta_{s2}, \dots, \theta_{st}, \dots\}, \forall \theta_{st} \in \Sigma_s$

Decision variables

$\alpha$  Input characteristics (# rows, size)

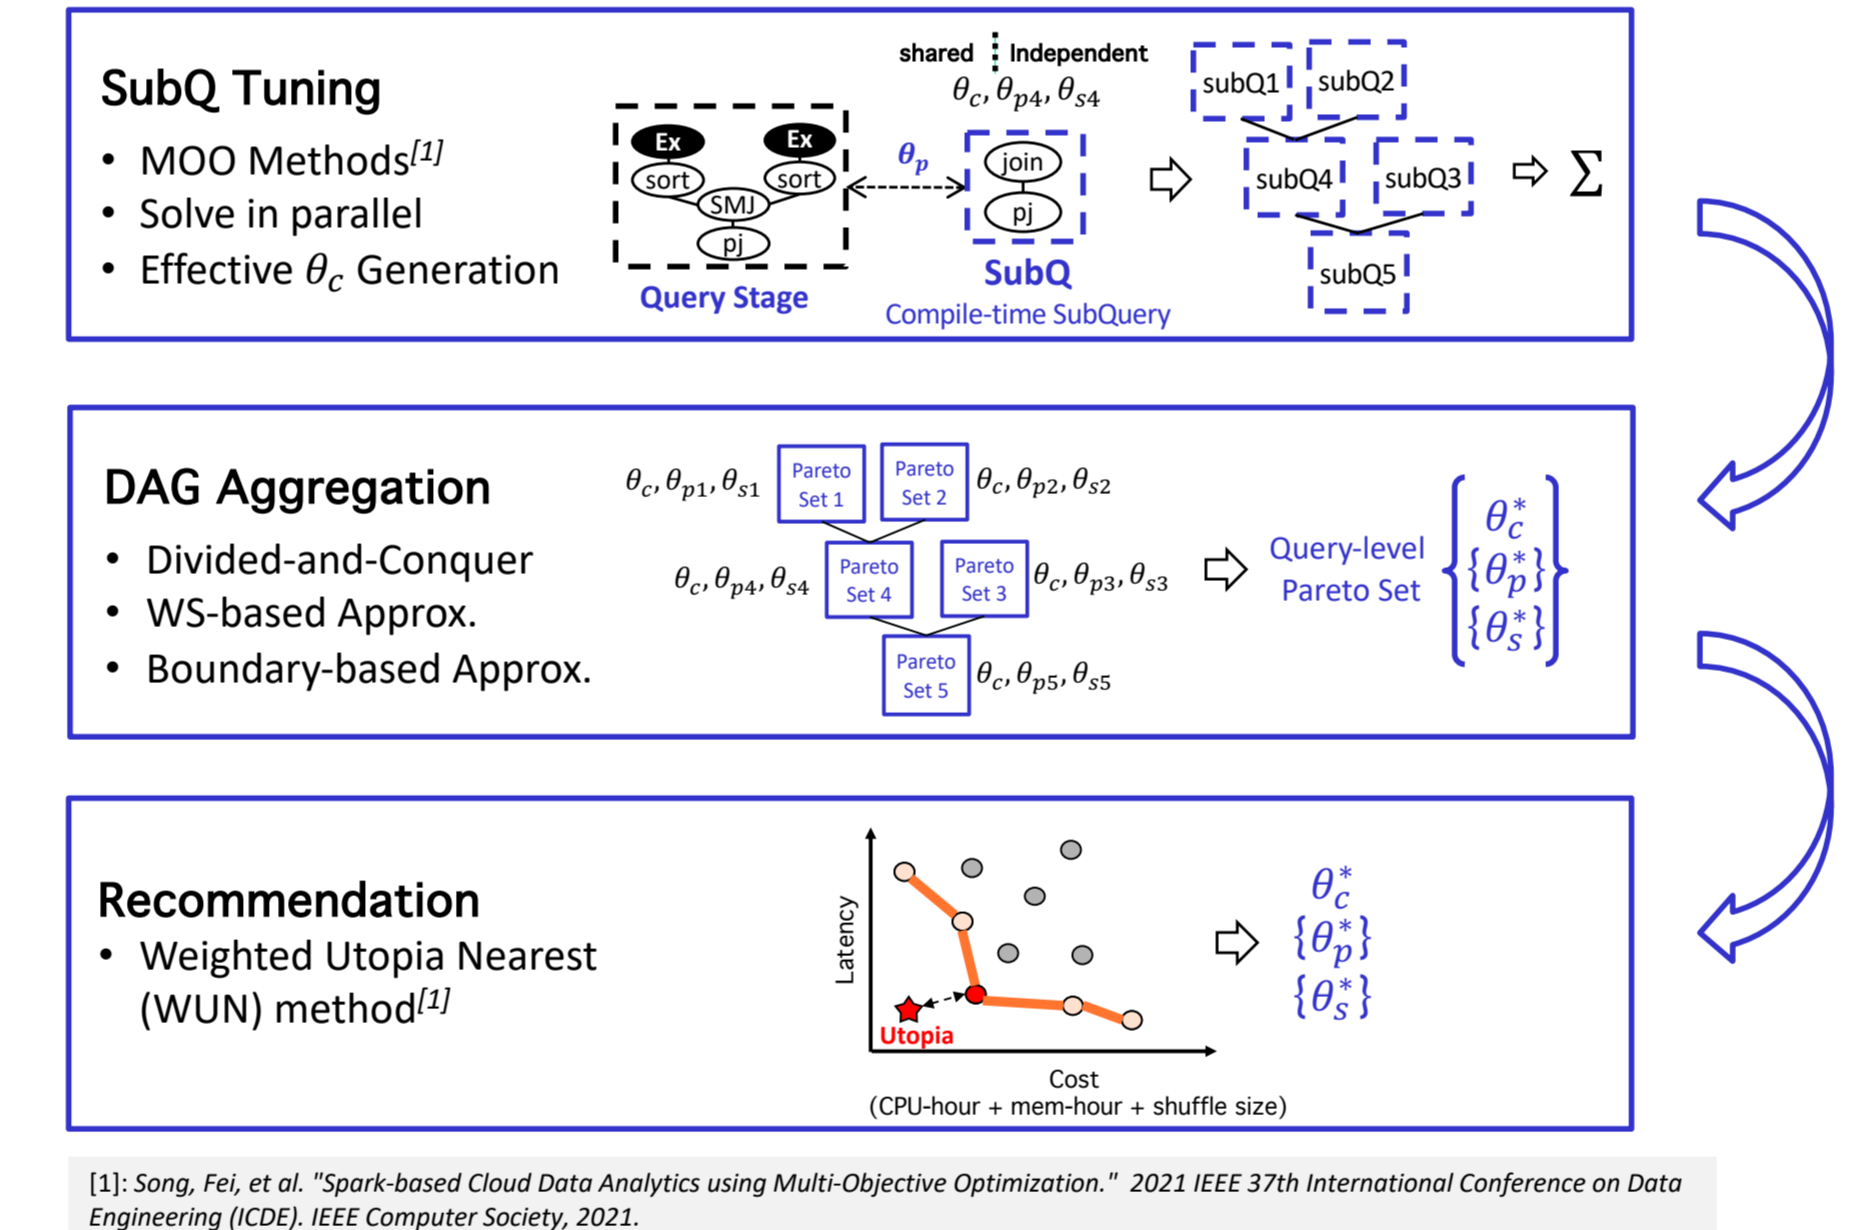
$\beta$  Distribution of input partitions

$\gamma$  Runtime status, encoded by statistics of the running tasks in parallel stages

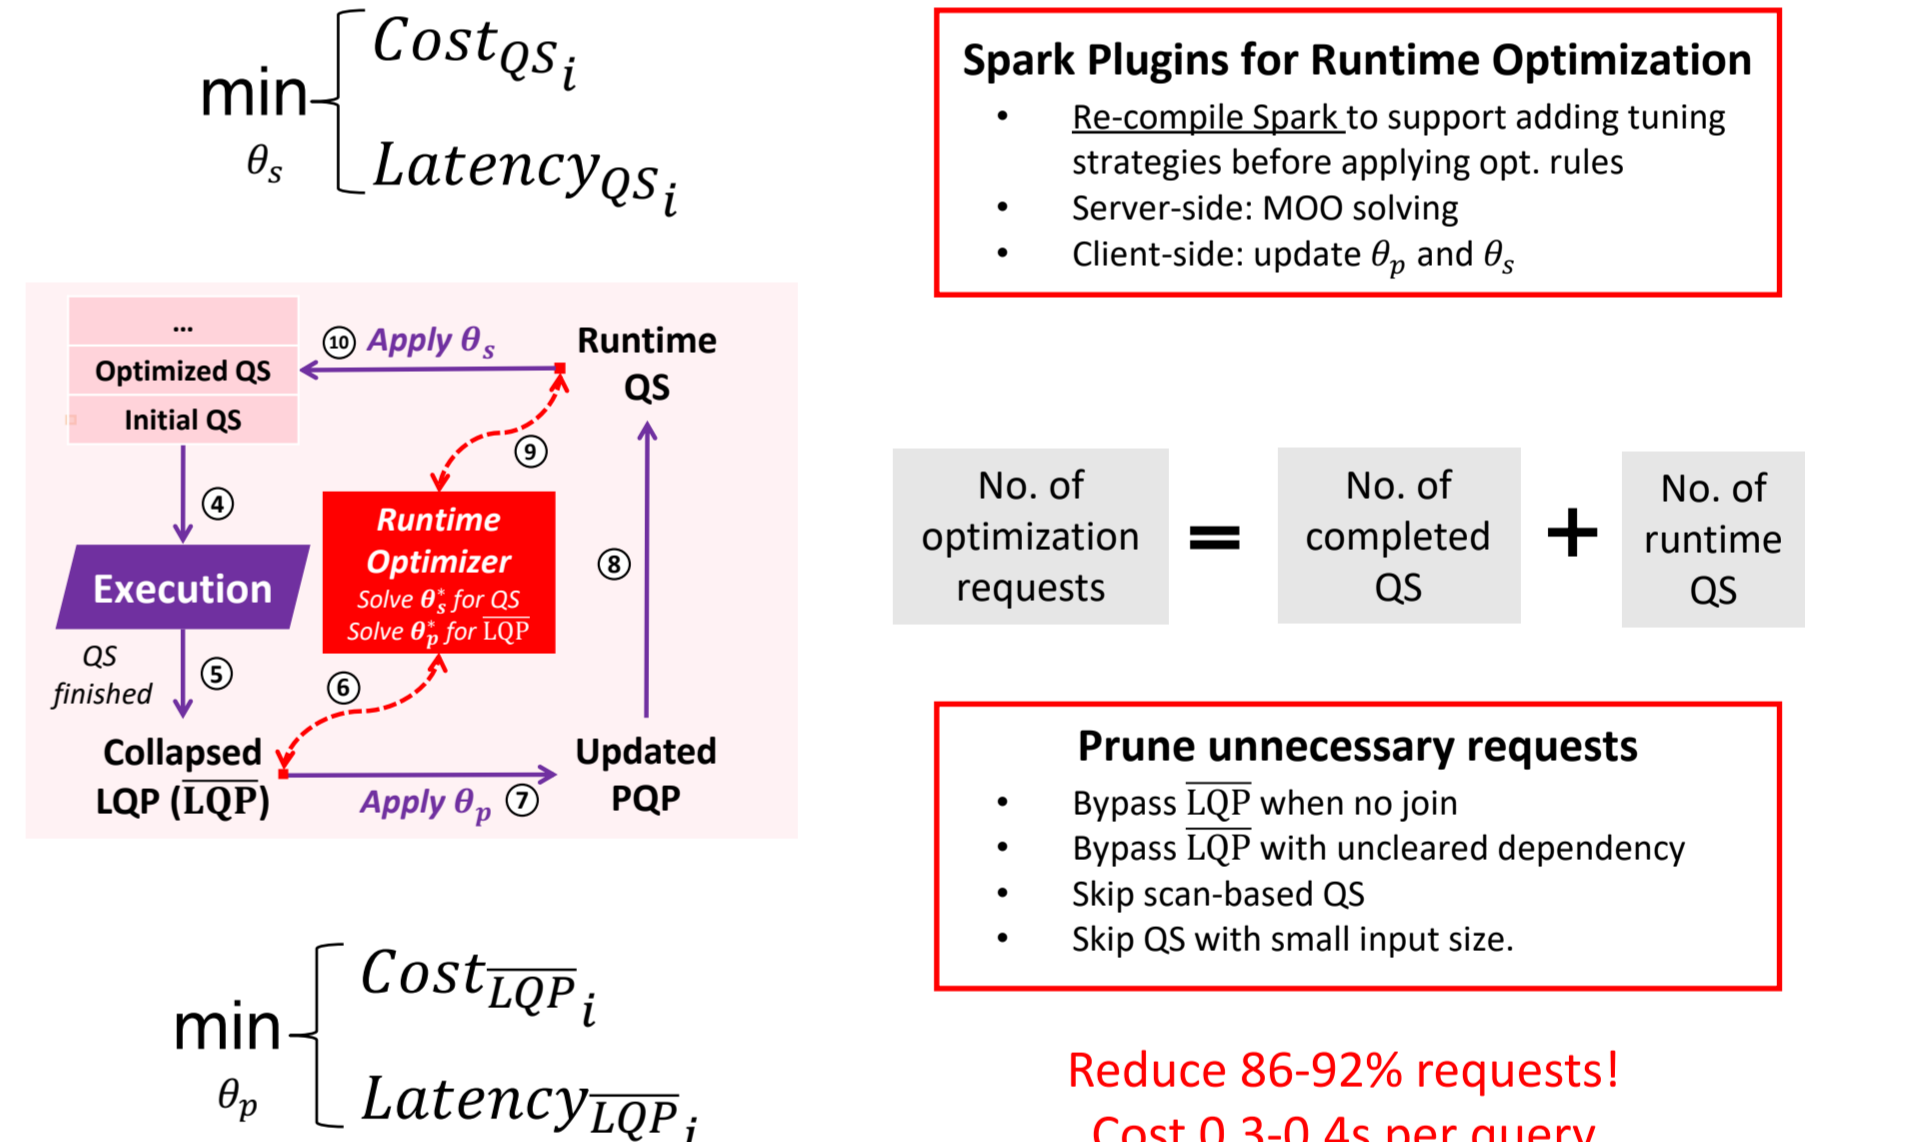
Non-decision variables

**Large Searching Space and High Overhead**

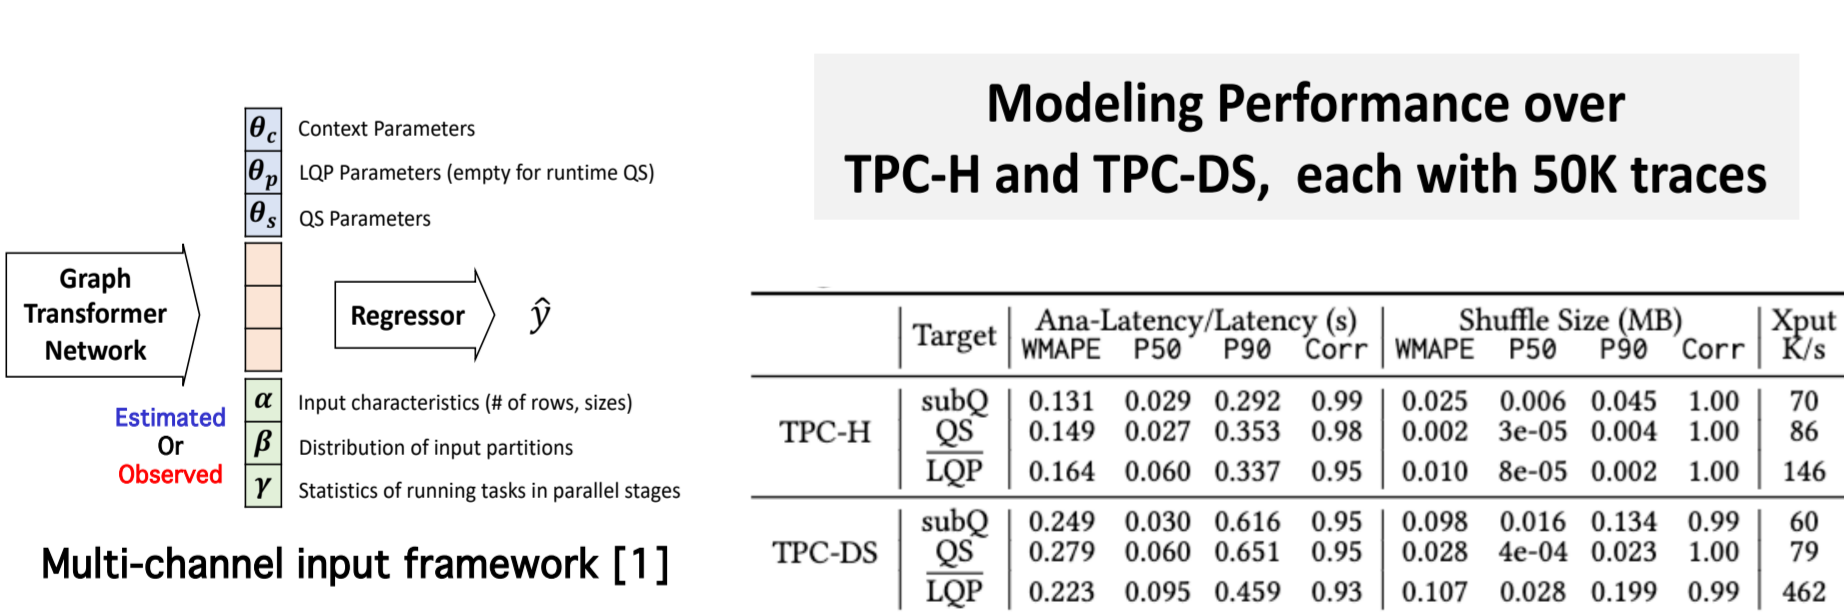
## Solution: Compile-Time Optimization



## Solution: Runtime Optimization



## Modeling Techniques and Results



**High Accuracy**

- Weighted Mean Absolute Percentage Error (WMAPE)
- P50Error, P90Error up to 0.65/0.19 for latency/shuffle size
- Pearson Correlation close to 1.

**High Inference Throughput**

- 60K-462K per second

## Benefits Over Query-Level MOO Against the Default

MO-WS<sup>[1]</sup> Multi-Objective Weighted Sum, Coarse-grained Tuning

HMOOC3 Ours, Compile-time, Fine-grained Tuning

HMOOC3+ Ours, Compile-time/Runtime, Fine-grained Tuning

Prioritize improving query speed

	TPC-H			TPC-DS		
	MO-WS	HMOOC3	HMOOC3+	MO-WS	HMOOC3	HMOOC3+
Total Lat Reduction	18%	61%	63%	25%	61%	65%
Avg Solving Time (s)	2.6	0.41	0.70	15	0.41	0.80

1. MO-WS: a total  $\downarrow$  of 18-25% latency with an average solving time of 2.6-15s

2. Ours: a total  $\downarrow$  of 63-65% latency with an average solving time of 0.7-0.8s

Multi-grained compile-time optimization ( $\downarrow$  61% with 0.41s)

## Adaptivity Comparison to SO with Fixed Weights

SO-FW Reduce to a Single Objectives with Fixed Weights

HMOOC3+\* Hybrid, Multi-granularity tuning (with multi-query plan search)

Performance-cost savings against default configuration over different preferences

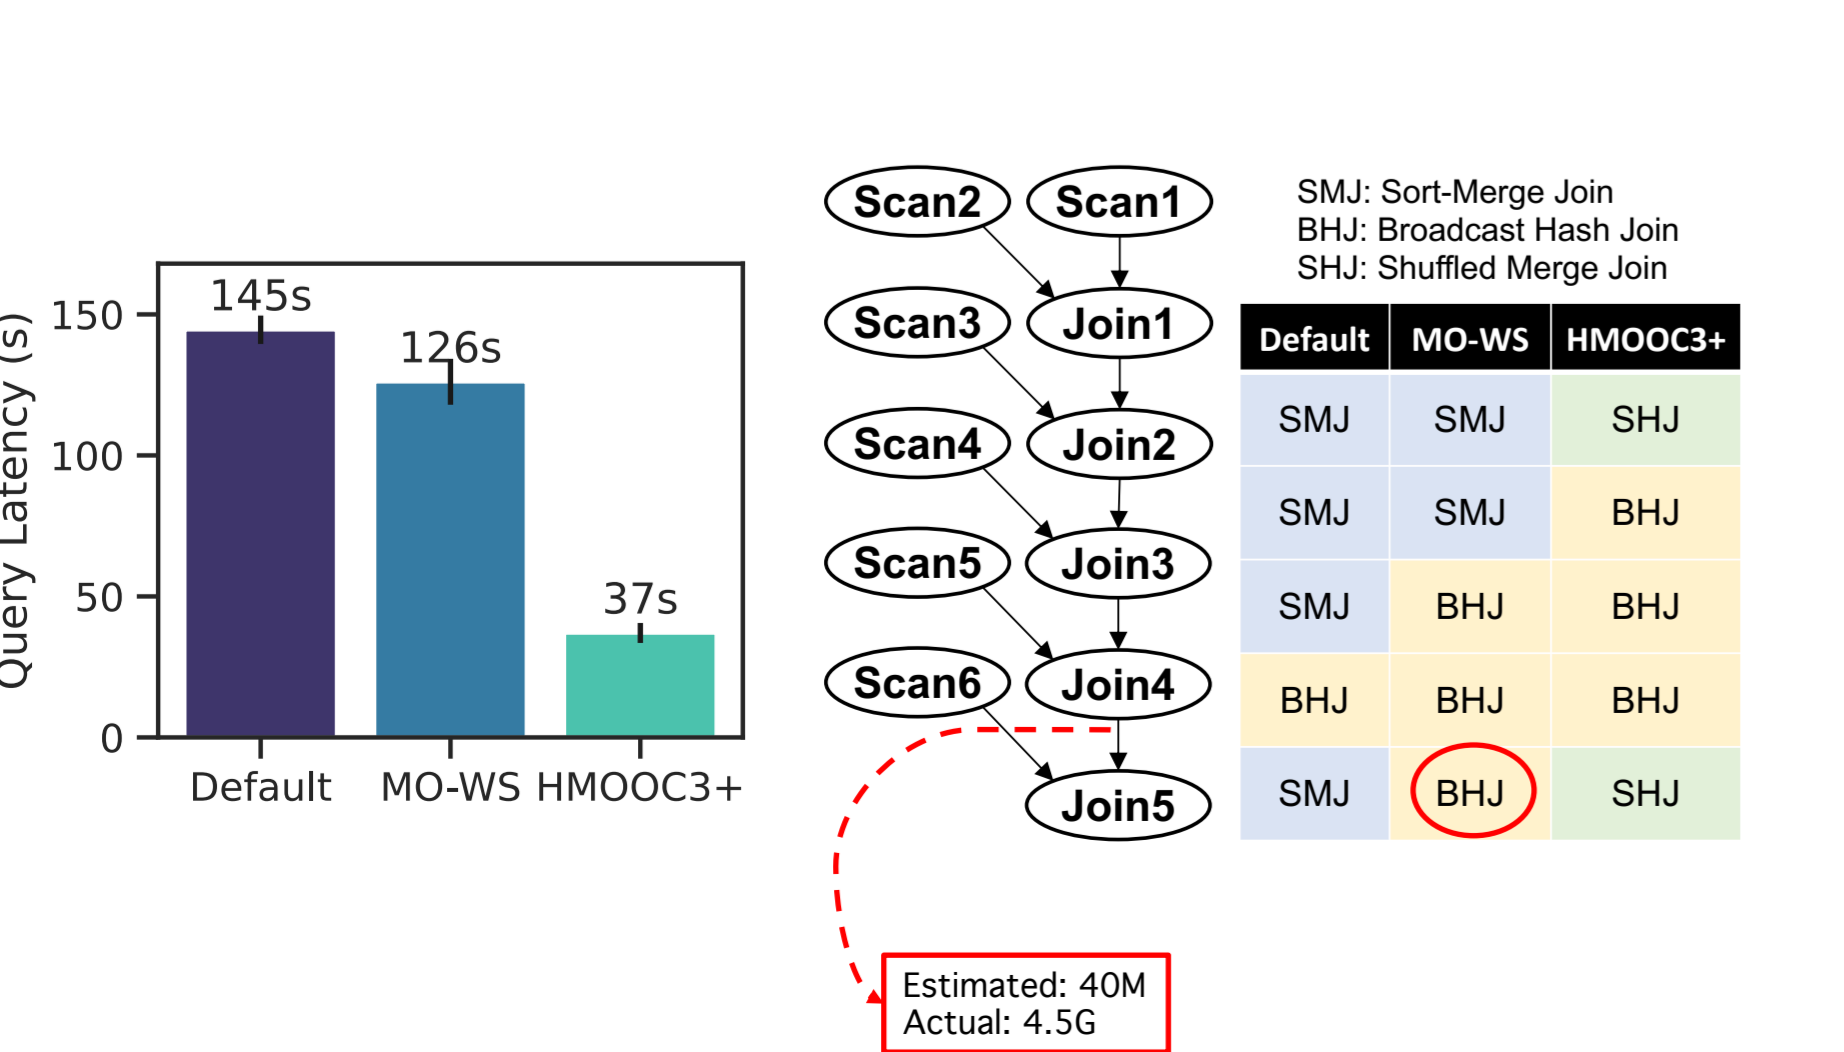
Prefs. Lat/Cost	SO-FW	TPC-H HMOOC3+*	SO-FW	TPC-DS HMOOC3+*
(0.0, 1.0)	20% / -11%	-15% / -10%	-6% / 64%	-45% / -22%
(0.1, 0.9)	1% / 1%	-31% / -5%	-28% / 105%	-57% / -7%
(0.5, 0.5)	-1% / 25%	-46% / -3%	-28% / 128%	-59% / 29%
(0.9, 0.1)	-13% / 27%	-51% / 0%	-34% / 139%	-59% / 55%
(1.0, 0.0)	-14% / 44%	-55% / 3%	-26% / 144%	-59% / 64%

An increasing demand of saving latency

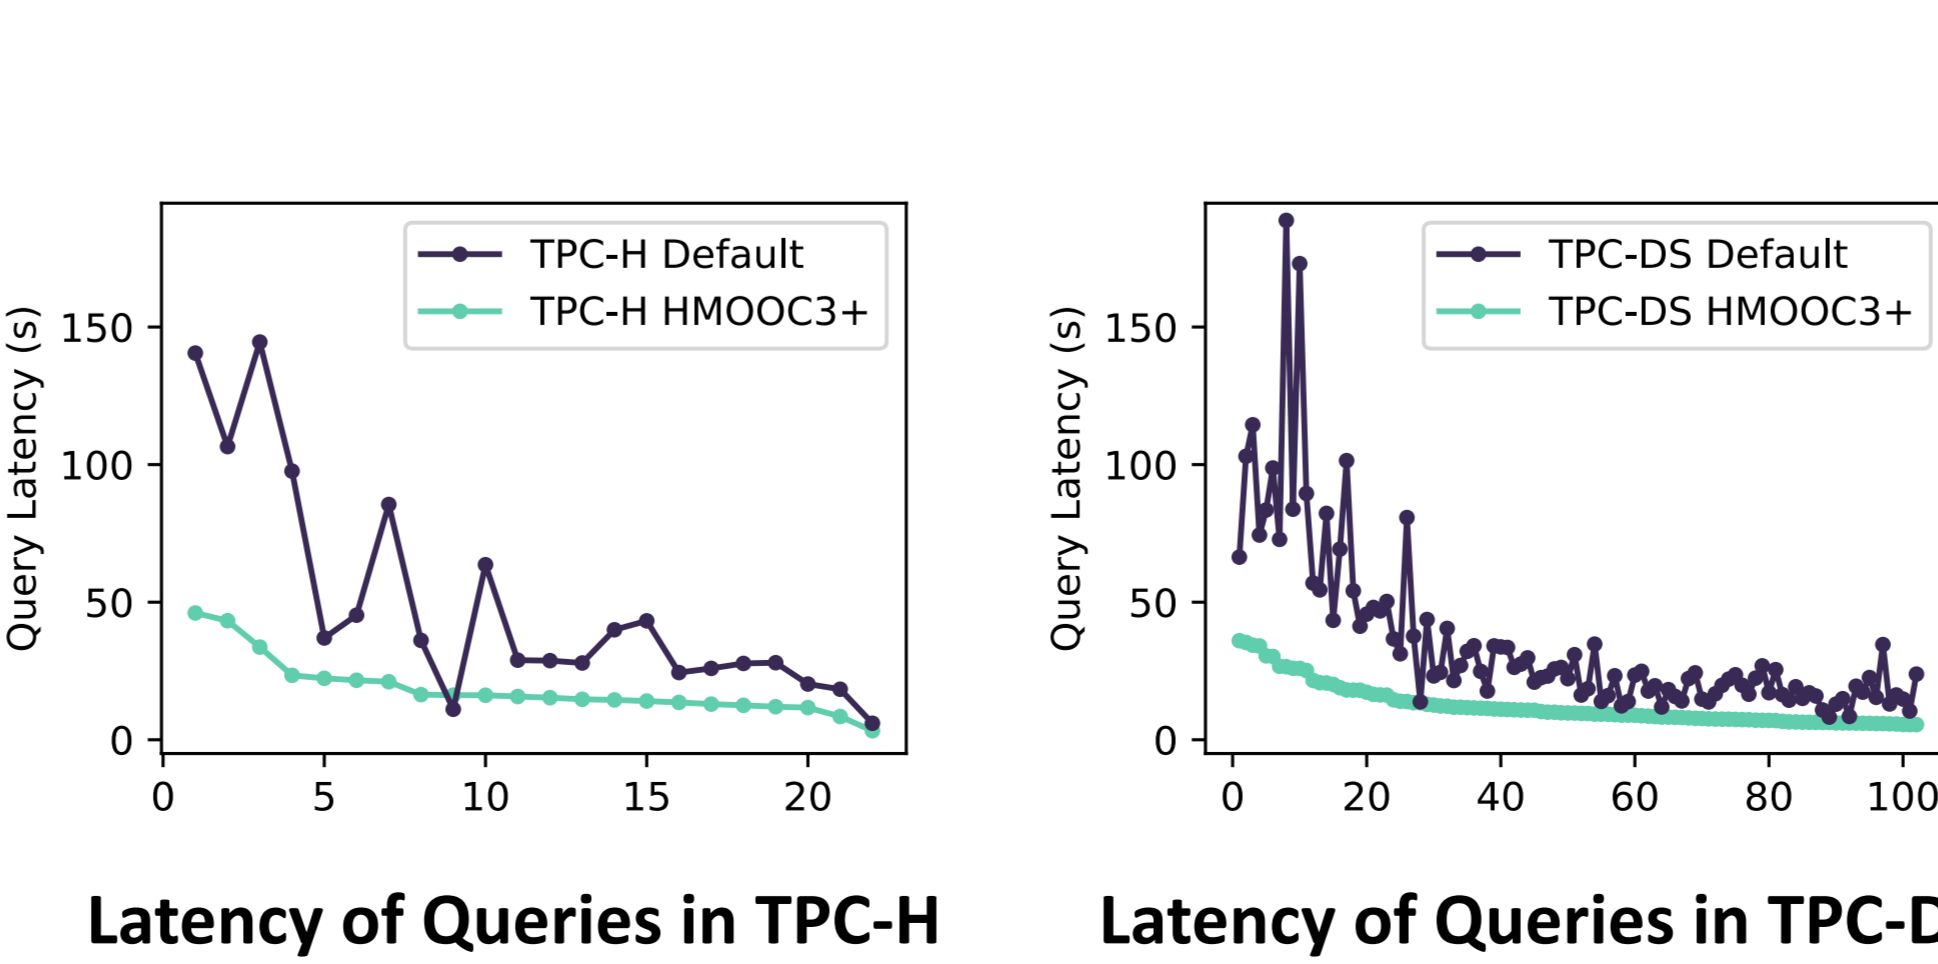
1. Ours: Up to 55-59%  $\downarrow$  latency and up to 10-22%  $\downarrow$  cost Superior adaptivity

2. SO-FW: at most 14-34%  $\downarrow$  latency and rare  $\downarrow$  cost Not adapting well

## Example of Tuning TPC-H Q9



## E2E Performance Comparison for Individual Queries



## Thanks For Your Attention

❖ **Code Implementation**

- [Client-side] Spark plugins to support runtime tuning: <https://t.ly/xMTxu>
- [Server-side] Modeling and MOO algorithms: <https://t.ly/XqfL1>

❖ **Potential Future Work**

- Extended to other systems who support runtime adaptivity with observed statistics (e.g., Presto, Greenplum, etc.)

❖ **Contact Me**

- Email: [chenghao@cs.umass.edu](mailto:chenghao@cs.umass.edu)
- I am expecting a full-time job in 2024/2025



Résumé

WeChat