DBA Bandits

Self-driving index tuning under ad-hoc, analytical workloads with safety guarantees

Malinga Perera

Joint work with Bastian Oetomo, Ben Rubinstein and Renata Borovica-Gajic

THE UNIVERSITY OF MELBOURNE

Index Tuning is Hard

Setting: TPC-H, SF10, DBMS-X, Tuning tool 5GB space for indexes*

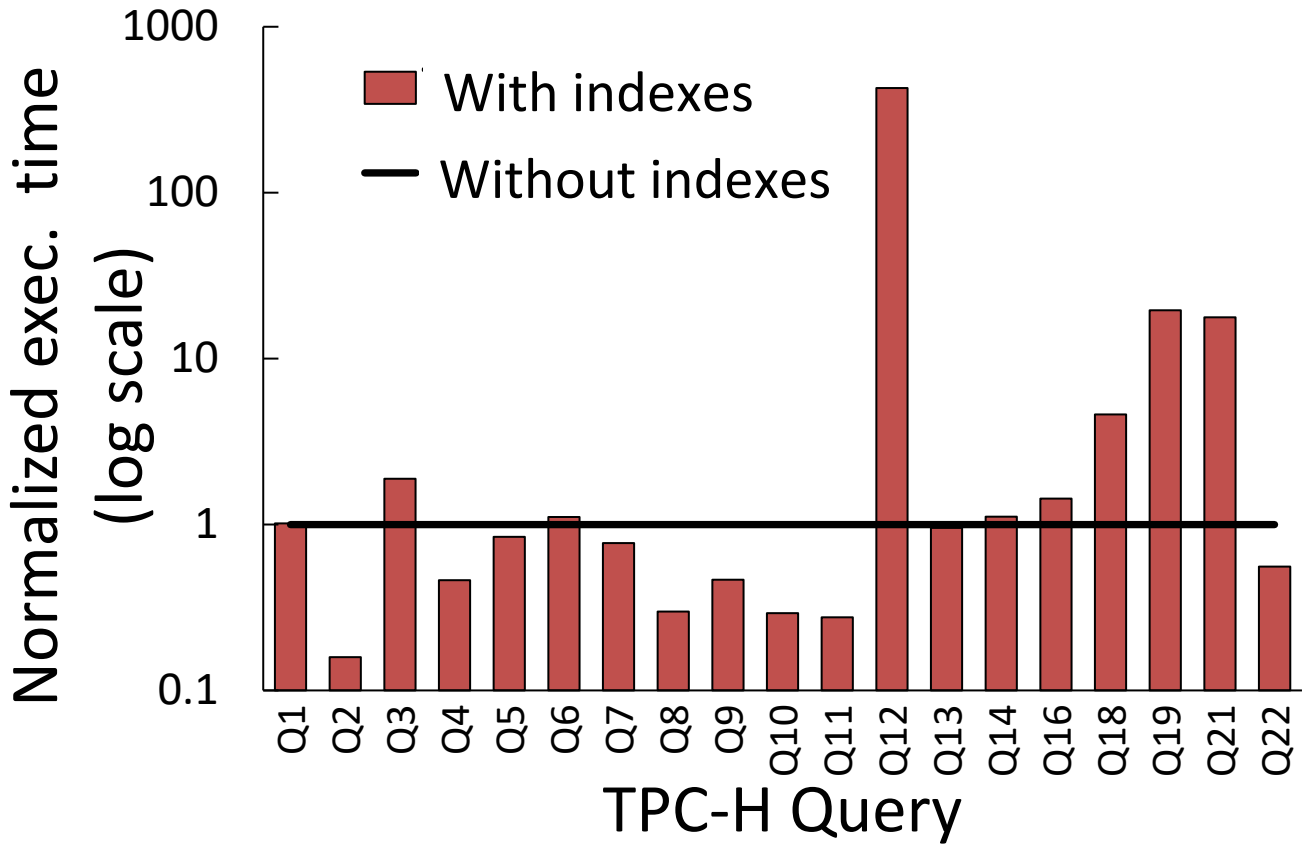

*R. Borovica-Gajic, S. Idreos, A. Ailamaki, M. Zukowski, and C. Fraser, "Smooth scan: Robust access path selection without cardinality estimation," The VLDB Journal, 2018.

Cause for Sub-Optimal Plans

Cardinality errors

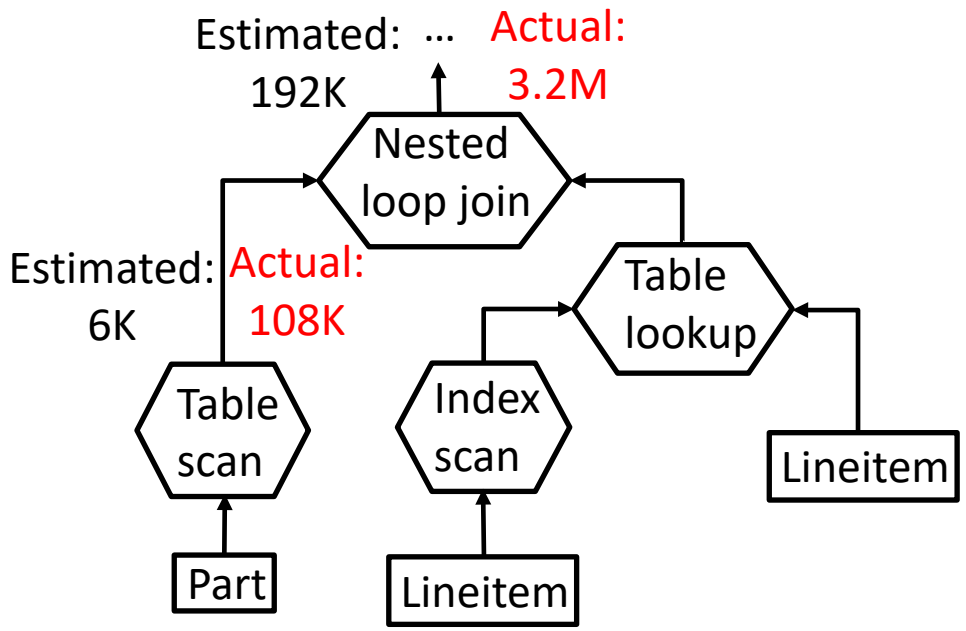

Optimizer's mistakes -> hurt predictability

Index Tuning Under Looking Glass

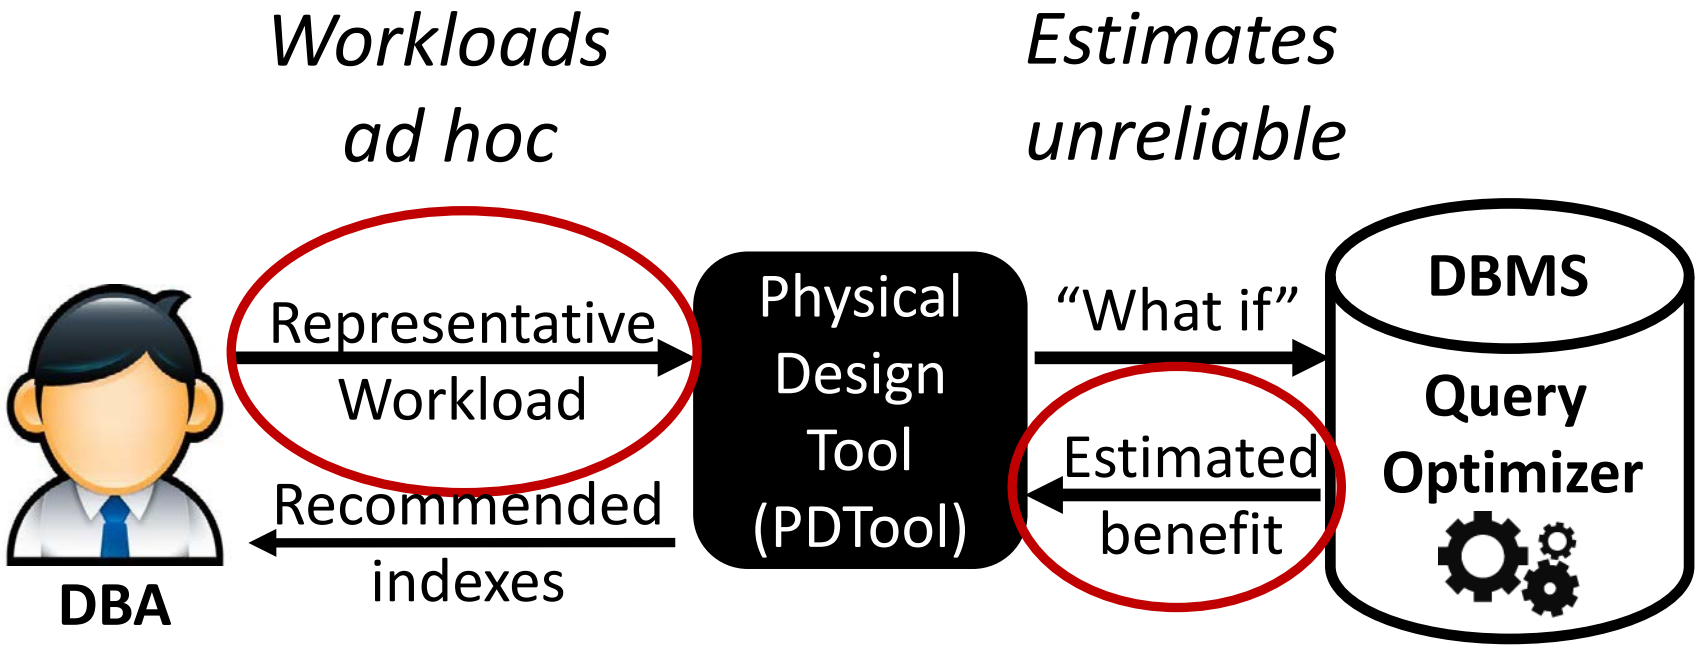

Broken pipeline….

Modern Applications: Untenable Status Quo

Bloomberg, Stock market°

Ever growing data

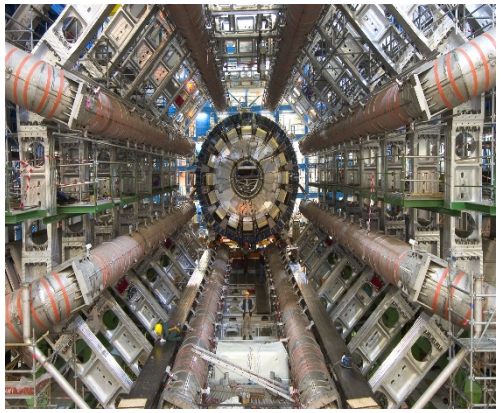

Atlas experiment, CERN*

(Ad hoc) data exploration

Strato Data Centre, cloudˆ

Multi-tenancy

Photo Credits: °Pexels, Muhammad Abdulaziz, *Wikimedia, Maximilien Brice, ˆStrato, Amber Case

Multi-Armed Bandits (MAB) to the Rescue

- Pull an arm (action) observe a reward
- Explore vs exploit
- Find a sequence of arms to maximize cumulative reward
- Many variants, but C^2UCB most interesting

Benefits of C^{2}UCB

- **UCB** *guarantees* to converge to optimal policy (optimism in the face of uncertainty)
- **Context** Allows the bandit to predict the reward of a new arm without trying even once
- **Combinatorial** pulls a set of arms per round given constraints

Fast convergence with guarantees

MAB for Index Tuning: An Example

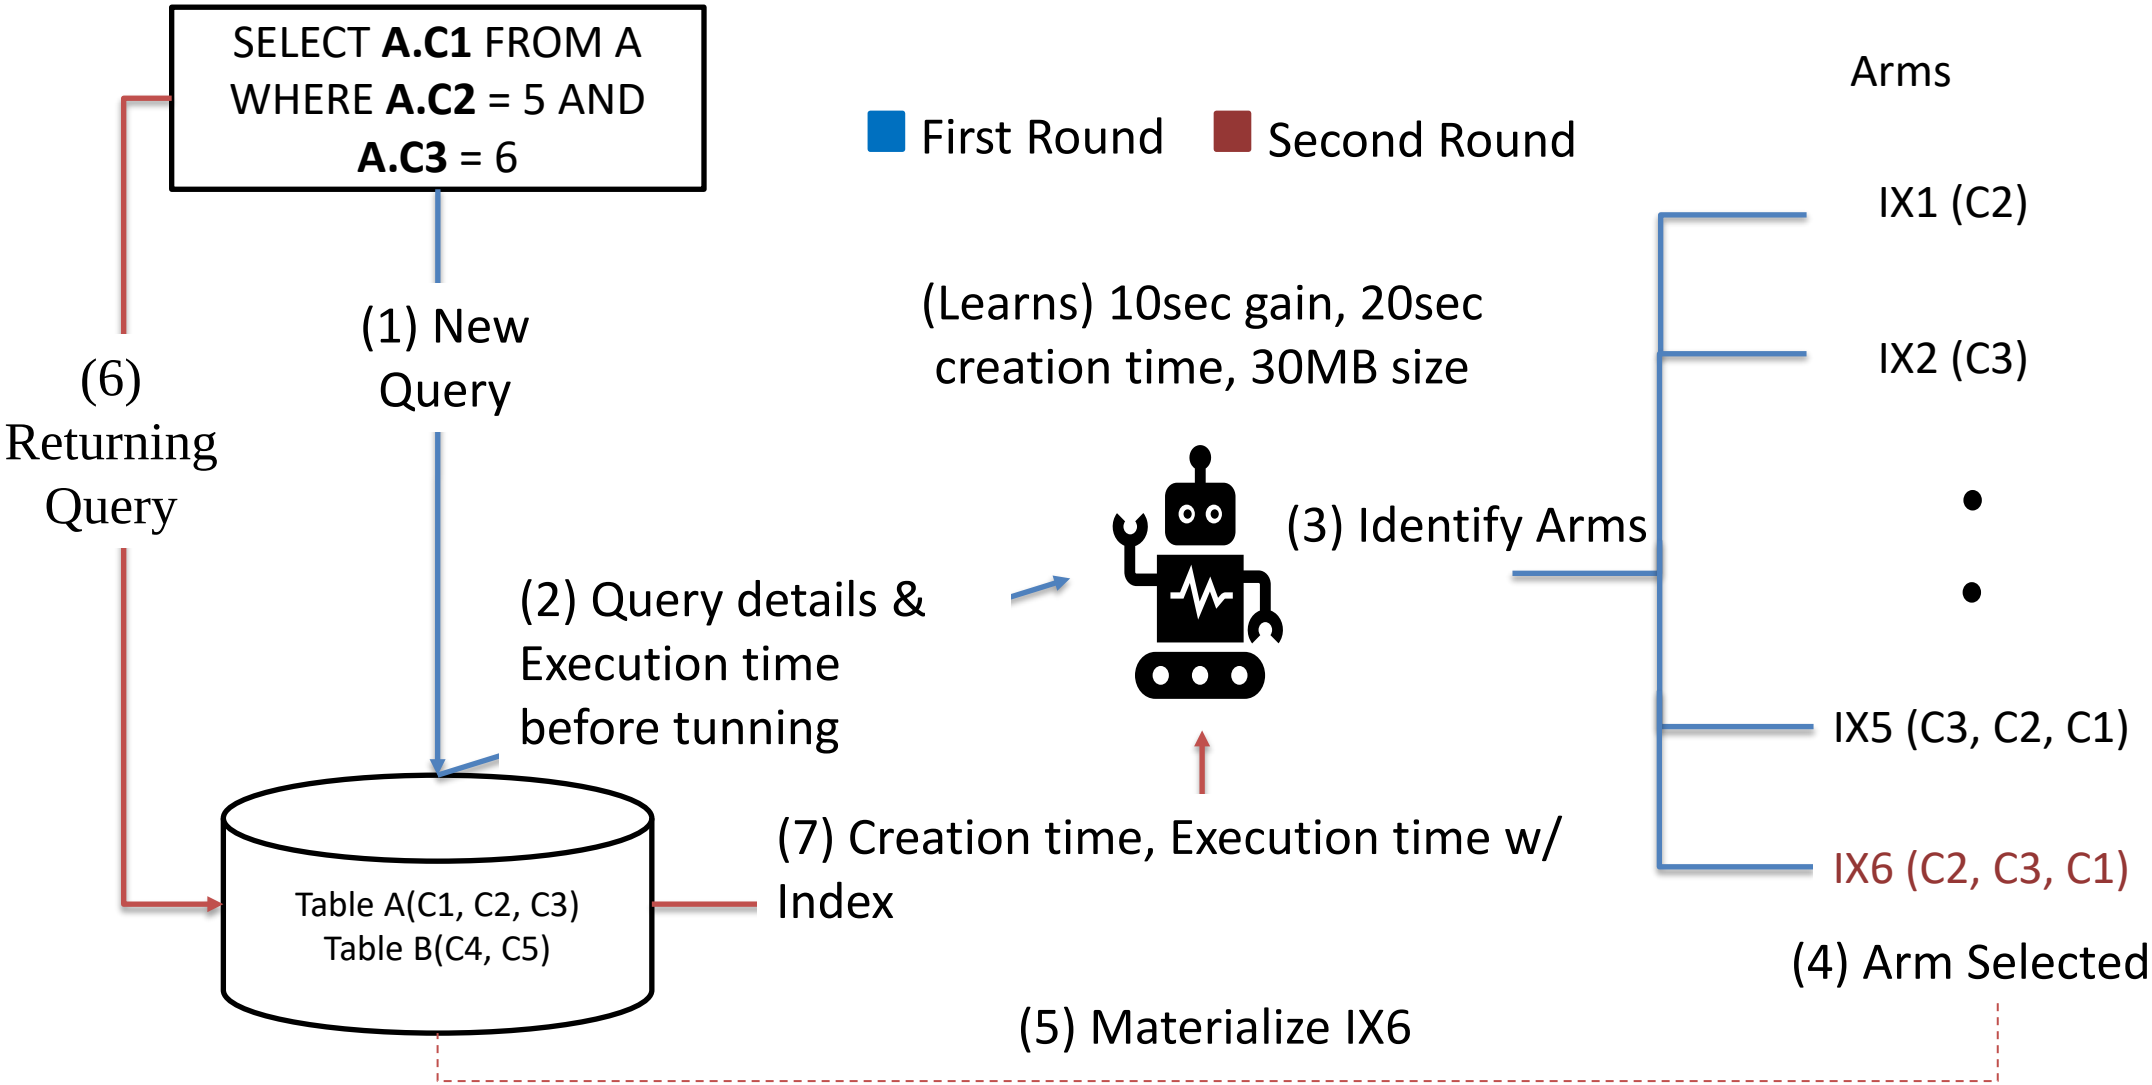

MAB Providing Up to 75% Speed-up

Setting: TPCH, TPCH skew, TPC DS, SSB (10GB); IMDb (6GB) datasets static (repetitive) vs random (ad hoc) queries, MAB vs PDTool, 25 rounds

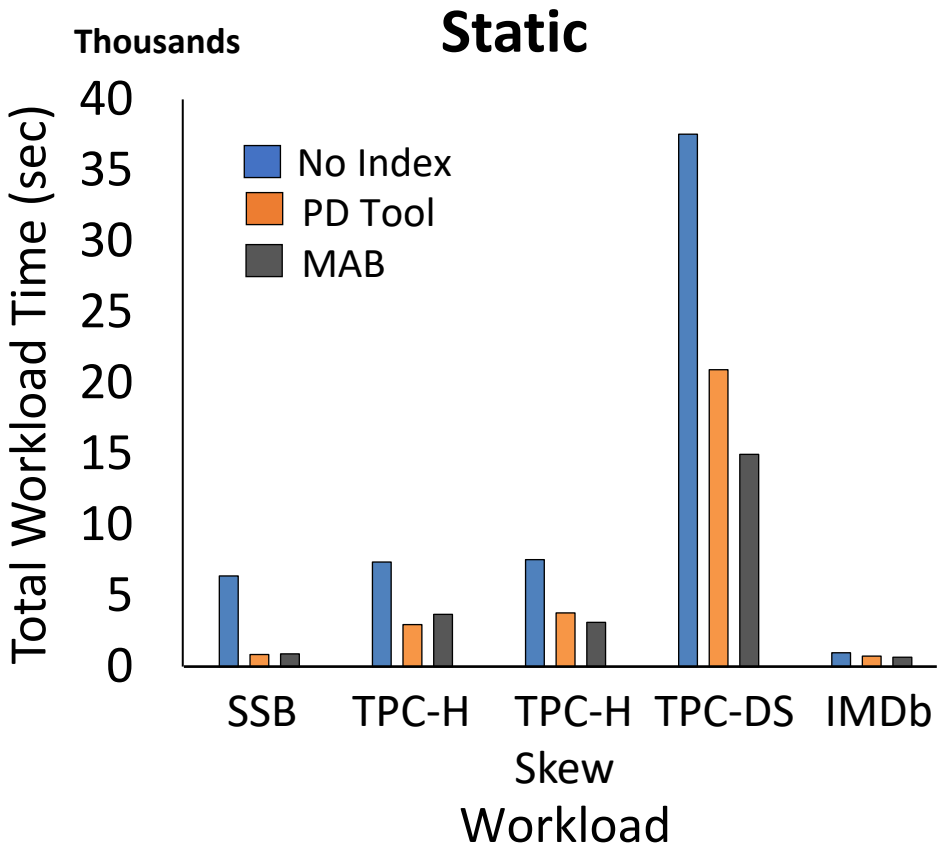

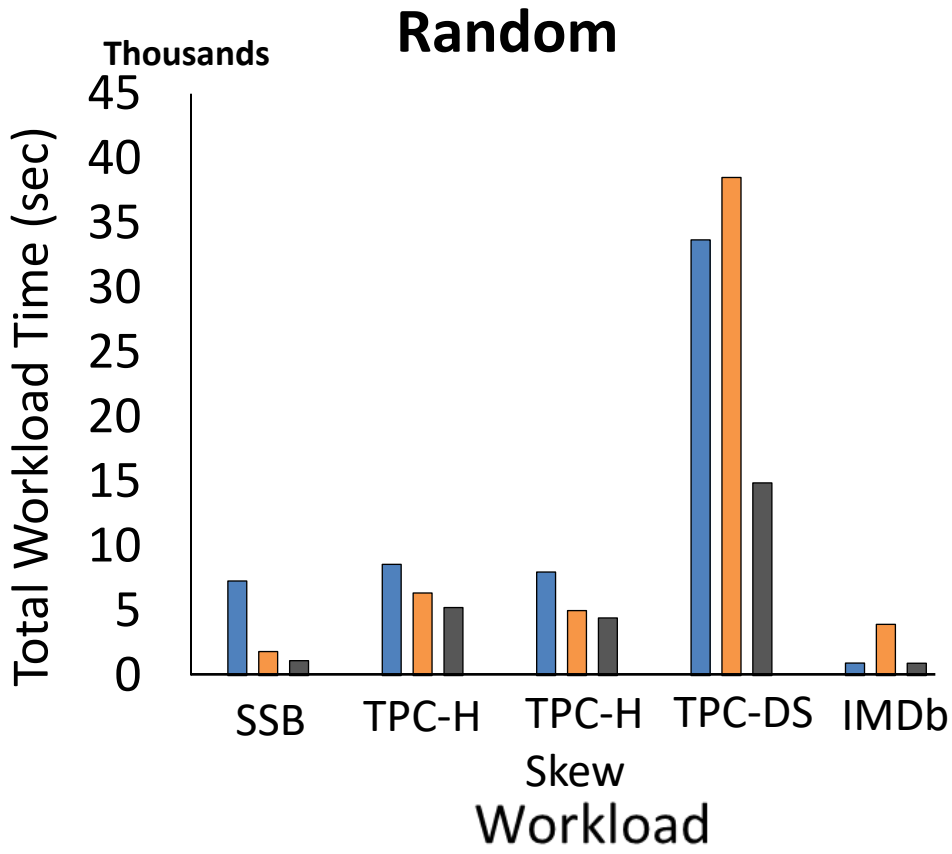

MAB with TPC-DS: Lightweight, Yet Efficient

Setting: TPC-DS, static vs ad hoc queries, MAB vs PDTool, 25 rounds

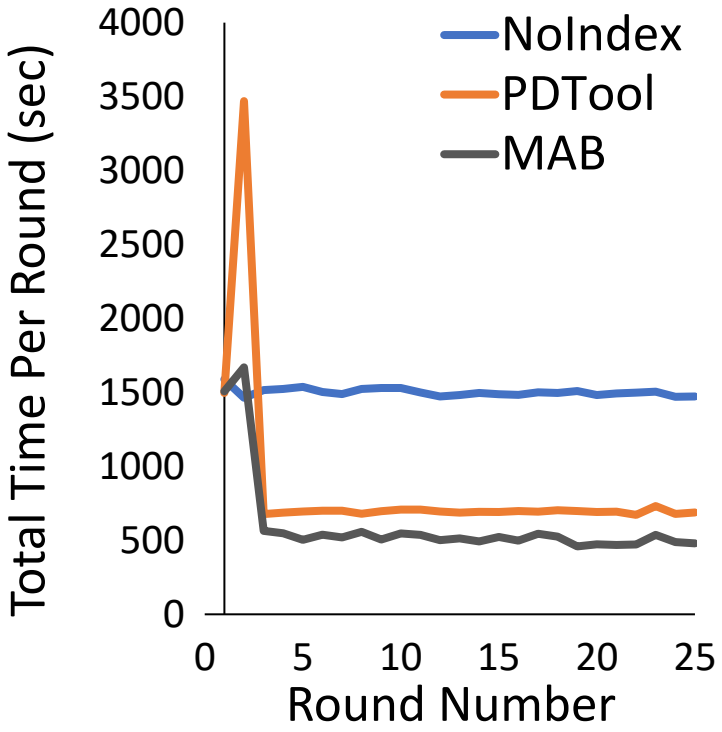

Static

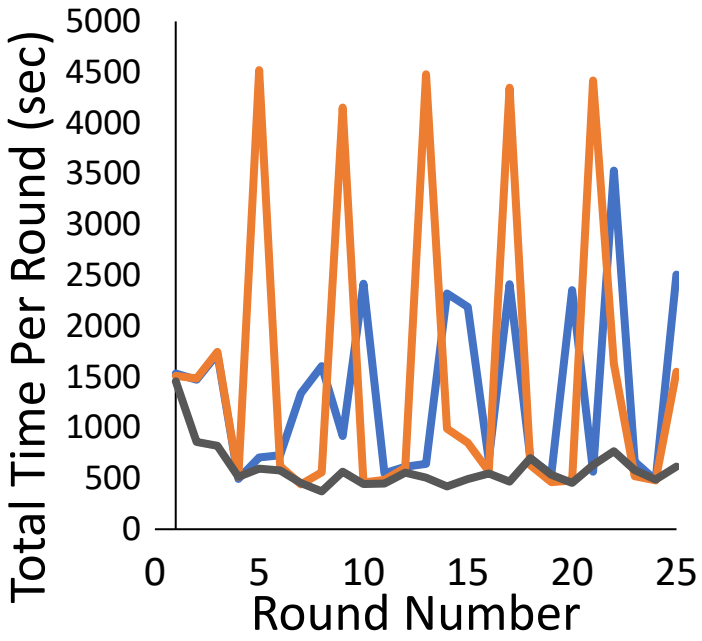

Random

Why not General RL?
Convergence Speed & Volatility

Setting: TPC-H and TPC-H Skew 10GB, 100 rounds *static*

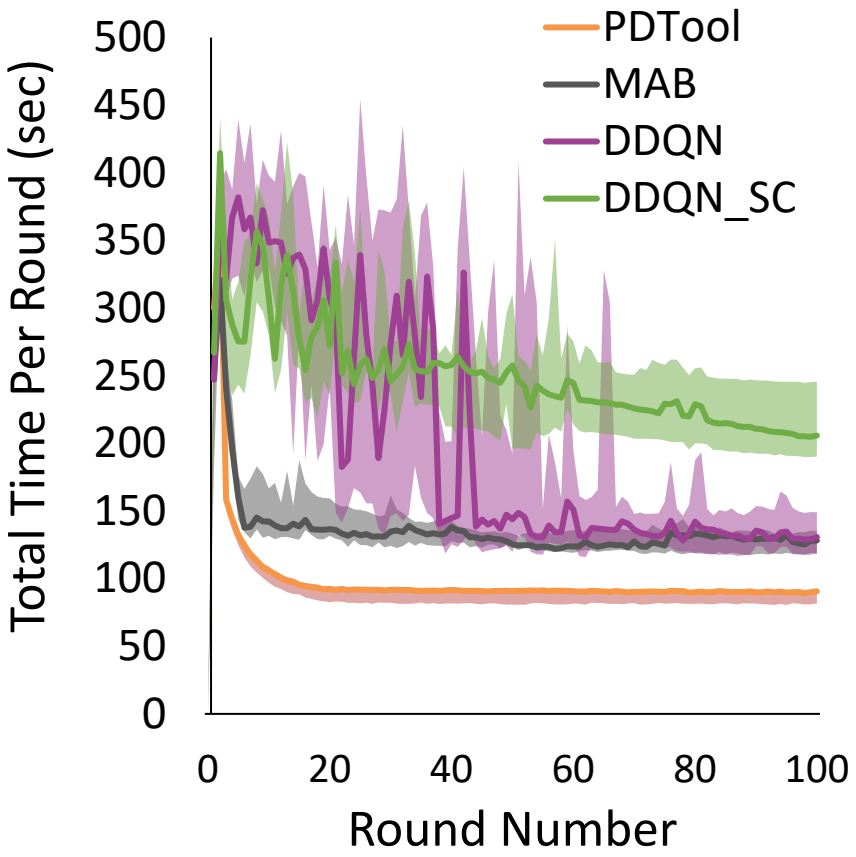

TPC-H

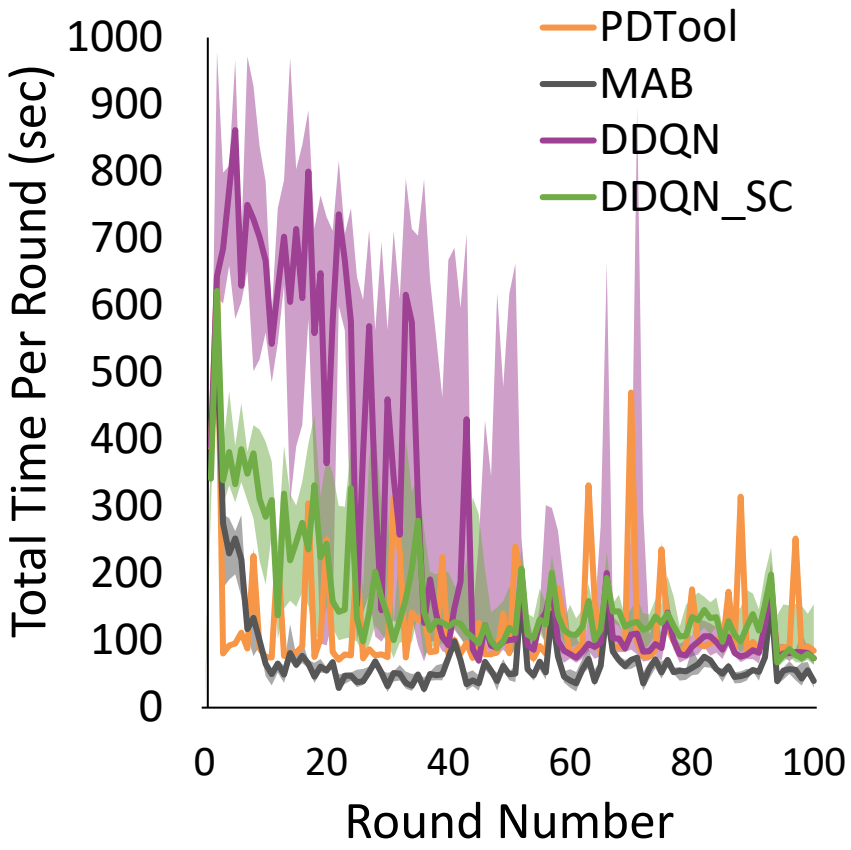

TPC-H Skew

Conclusions

- DBA Bandit is a lightweight MAB solution for index tuning
- C^2UCB enables strategic exploration *without* pulling all arms
- Safety bounds guarantee convergence to optimal choice (in hindsight)
- MAB mitigates tuning tool flaws: optimizer misestimates, unpredictable workloads
- Up to 75% improvement and 25% on average compared against a SOTA commercial tuning tool

Special Thanks

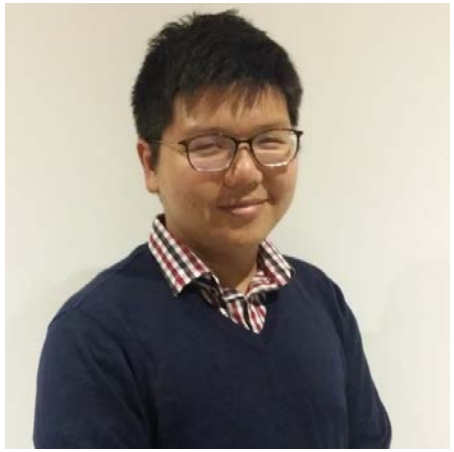

Bastian Oetomo

Ben Rubinstein

Renata Borovica-Gajic

Questions?

Thank you

Impact of Database Size: Promise for Big Data Era

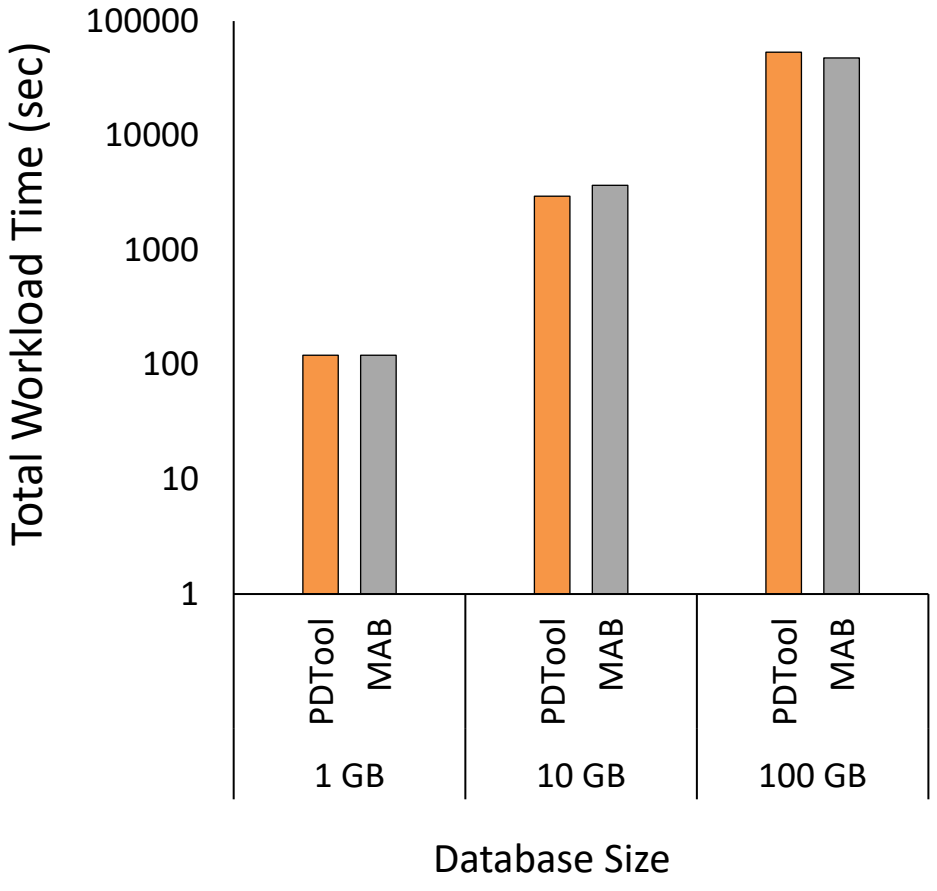

TPC-H

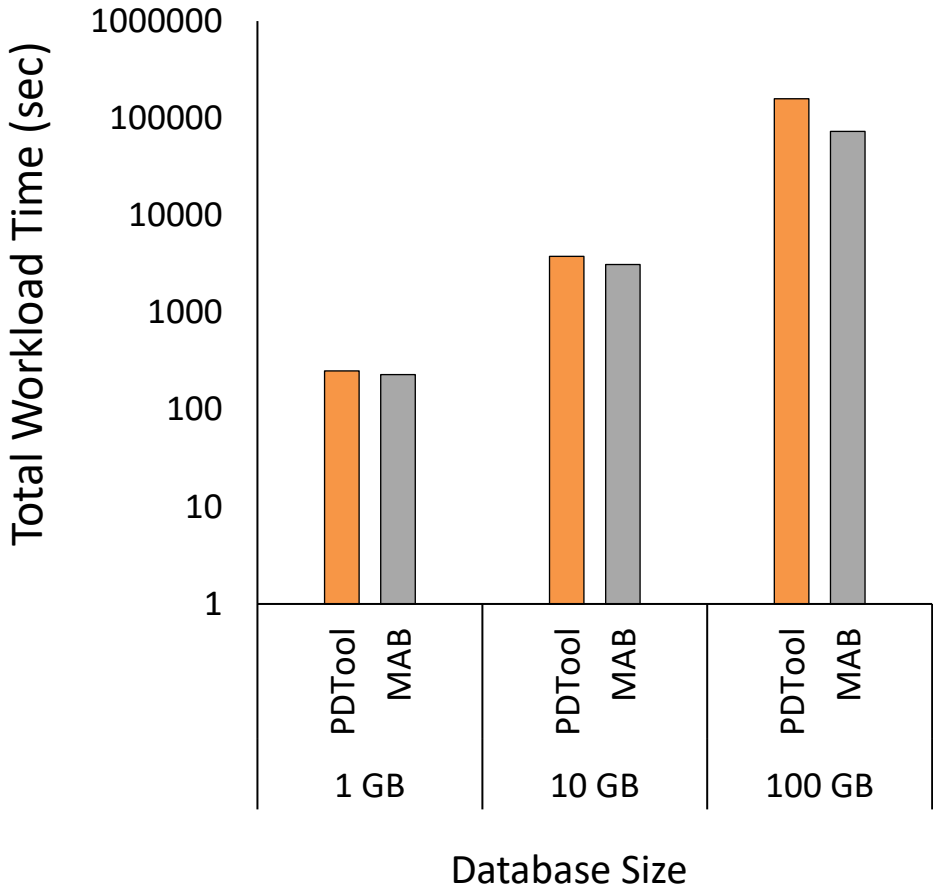

TPC-H Skew

Database Sizes: For MAB, Bigger Is Better

TABLE II
TOTAL END-TO-END WORKLOAD TIME FOR STATIC WORKLOADS UNDER DIFFERENT DATABASE SIZES (MIN)

Workload	SF	PDTool	MAB
TPC-H	1	**2.02**	2.03
	10	**49.4**	61.38
	100	891.01	**793.40**
TPC-H Skew	1	4.17	**3.83**
	10	63.12	**51.99**
	100	2640.64	**1219.33**

Skewed Data: Where MAB Shines

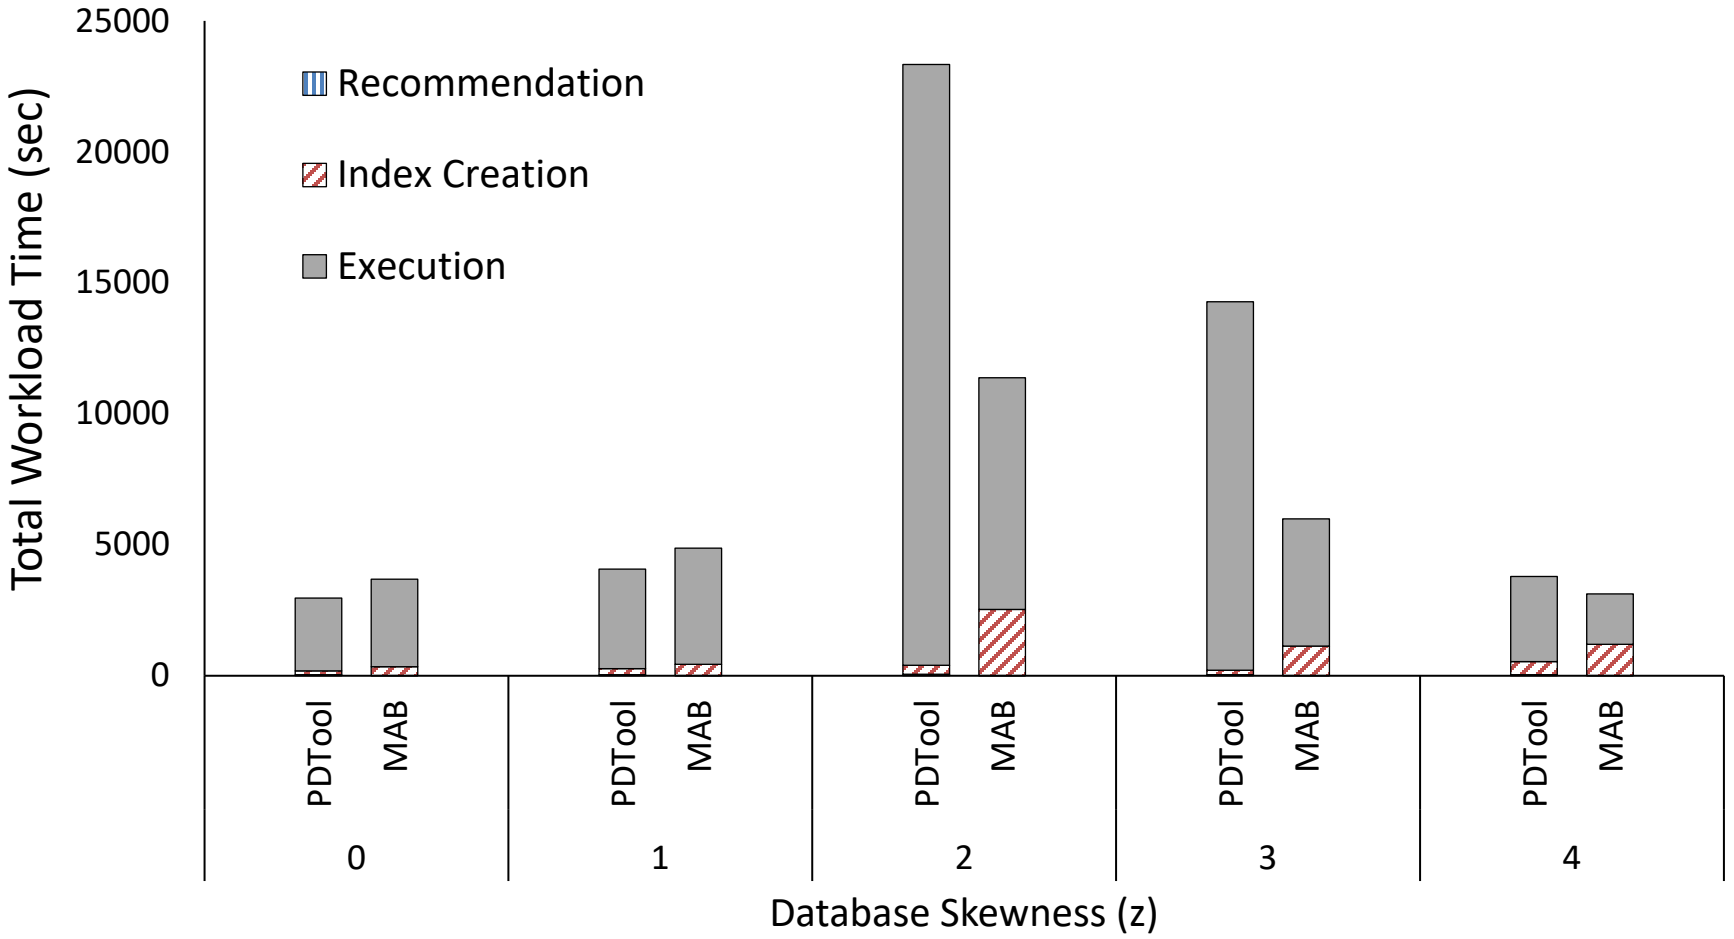

Main Results: MAB Wins 13/15

TABLE I
TOTAL TIME BREAKDOWN (MIN): THE BEST CHOICE IS IN BOLD TEXT.

Workload		Recommendation		Creation		Execution		Total	
		PDTool (#)	MAB	PDTool	MAB	PDTool	MAB	PDTool	MAB
Static	SSB	0.34 (0.34)	**0.02**	**0.95**	1.86	**12.9**	13.15	**14.19**	15.03
	TPC-H	0.6 (0.6)	**0.08**	**2.45**	5.66	**46.35**	55.64	**49.4**	61.38
	TPC-H Sk.	0.58 (0.58)	**0.11**	**8.37**	19.82	54.17	**32.06**	63.12	**51.99**
	TPC-DS	44.86 (44.86)	**1.53**	**1.45**	5.94	302.63	**242.15**	348.94	**249.62**
	IMDB	0.34 (0.34)	**0.31**	**1.1**	1.3	11.01	**9.42**	12.41	**11.03**
Dynamic	SSB	1.28 (0.32)	**0.05**	**1.5**	2.21	**5.42**	5.69	8.2	**7.95**
	TPC-H	1.55 (0.32)	**0.12**	**9.36**	9.74	26.35	**25.14**	37.25	**35**
	TPC-H Sk.	1.65 (0.41)	**0.16**	**14.98**	20.96	85.49	**21.44**	102.11	**42.56**
	TPC-DS	11.13 (2.78)	**1.66**	**6.08**	16.48	187.08	**155.65**	204.29	**173.79**
	IMDB	3.09 (0.77)	**0.29**	**1.59**	2.24	11.21	**7.93**	15.89	**10.46**
Random	SSB	2.83 (0.57)	**0.02**	**1.77**	2.37	26.59	**16.83**	30.85	**19.22**
	TPC-H	7.55 (1.51)	**0.08**	14.68	**7.06**	84.14	**80.43**	106.37	**87.57**
	TPC-H Sk.	3.3 (0.66)	**0.08**	**31.74**	34.68	48.71	**39.44**	83.75	**74.2**
	TPC-DS	310.22 (62.04)	**1.4**	**8.23**	19.81	323.57	**227.02**	642.01	**248.24**
	IMDB	14.74 (2.94)	**0.28**	2.72	**1.14**	48.55	**14.47**	66.01	**15.89**

The average time of a single PDTool invocation

Fast Recommendation Times, Workload Complexity Helps

TABLE III
RECOMMENDATION TIMES (MIN) VS. WORKLOAD COMPLEXITY

Workload	SSB	TPC-H	TPC-DS
PDTool	0.84	1.36	44.86
MAB	0.05	0.14	1.53

Strong Performance Under a Range of Round Sizes

TABLE IV
TPC-H SKEW BENCHMARK UNDER DIFFERENT ROUND SIZES (MIN)

Round size	Recommendation	Creation	Execution	Total
Single Query	1.11	27.77	30.16	59.04
0.5x	0.13	22.39	30.39	52.92
1x	0.11	19.82	32.06	51.99
2x	0.08	12.66	43.53	56.27

Why not General RL? Efficient Exploration

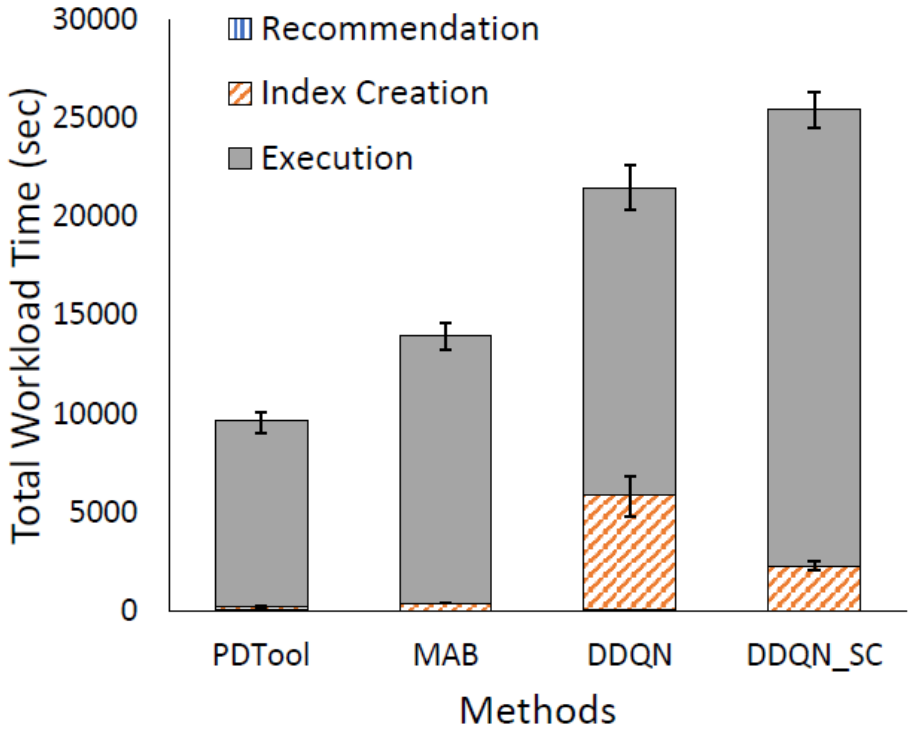

TPC-H

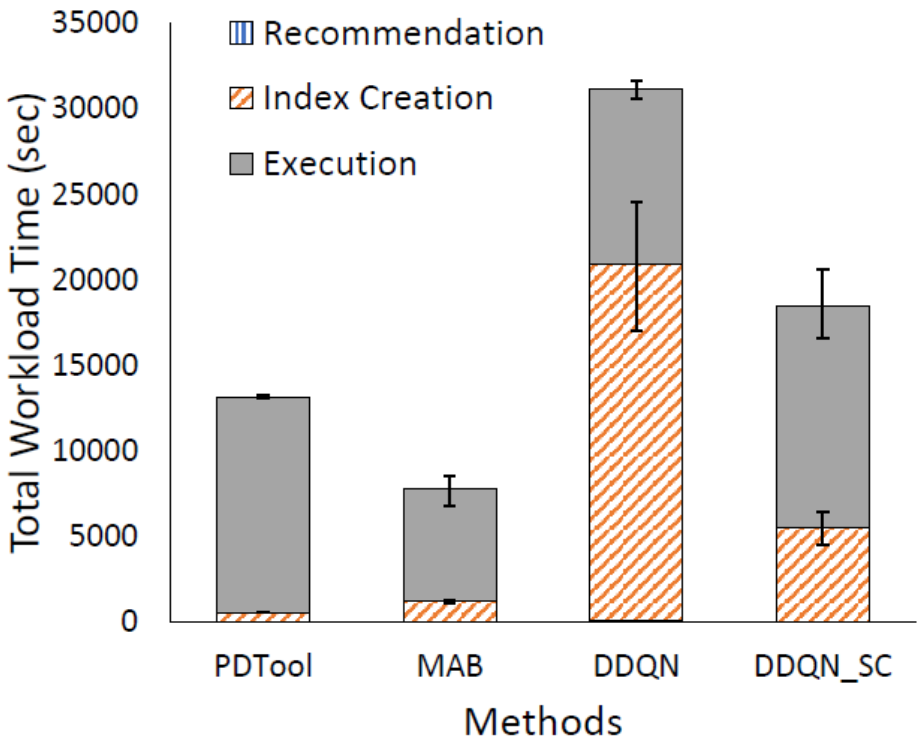

TPC-H Skew