# The Case for Phase-Aware Scheduling of Parallelizable Jobs

# **Benjamin Berg**, Justin Whitehouse, Benjamin Moseley, Weina Wang, Mor Harchol-Balter





1

#### Parallel Systems Have Resulted in Parallel Workloads



# Basic Model of a Database



- How to model the benefit of running on additional cores?
- How to split cores between queries?

Carnegie Mellon

University

# Job Size

- Job Size: Running time on single core
- Reflects amount of "inherent work"
- E.g. number of tuples to be processed

How does running time change on k cores?

s(k): How many times faster on k cores than on 1 core



# But, jobs change over time



# **Beyond Speedup Functions**





**Amdahl's Law:** Only p fraction of the job is parallelizable

Speedup functions measure *average* speedup







# Modeling Jobs with Phases



Idea: each job looks like a Continuous Time Markov Chain

# How to split cores between queries?

#### How should a scheduler leverage phase information?



## The Problem with Phase Unaware Scheduling



# Scheduling In Databases





#### Phase-Aware First-come-first-served (PA-FCFS)

NoisePage database [VLDB 2021] using morsel-driven parallelism [SIGMOD 2014]

# Improving phase Aware scheduling



### Two Key Ideas:

- Defer Parallelizable Work
- Work on shorter queries

# Proof: Deferring parallelizable work can <sup>12</sup> be sufficient for optimality

• Inelastic-First (IF) Policy: give strict priority to inelastic phases



- Phase sizes unknown
- Elastic phases ~  $\exp(\mu_E)$ , Inelastic phases ~  $\exp(\mu_I)$
- IF minimizes mean response time

What happens when we violate these assumptions?

# **Incorporating Query Size**



Process elastic phases in SRPT order
Why not Phase-Aware SRPT (PA-SRPT)?



# Does IF work in Practice?



# Conclusion



**Two Key Ideas:** 

- Defer Parallelizable Work
- Work on shorter jobs

- IF can be optimal
- IF-SRPT performs well in practice

# Questions?



**Two Key Ideas:** 

- Defer Parallelizable Work
- Work on shorter jobs

- IF can be optimal
- IF-SRPT performs well in practice