

# Scheduling for Reduced Tail Task Latencies in Highly Utilized Datacenters

**Smita Vijayakumar**, Evangelia Kalyvianaki

Department of Computer Science, University of Cambridge

July 2023

# Focus

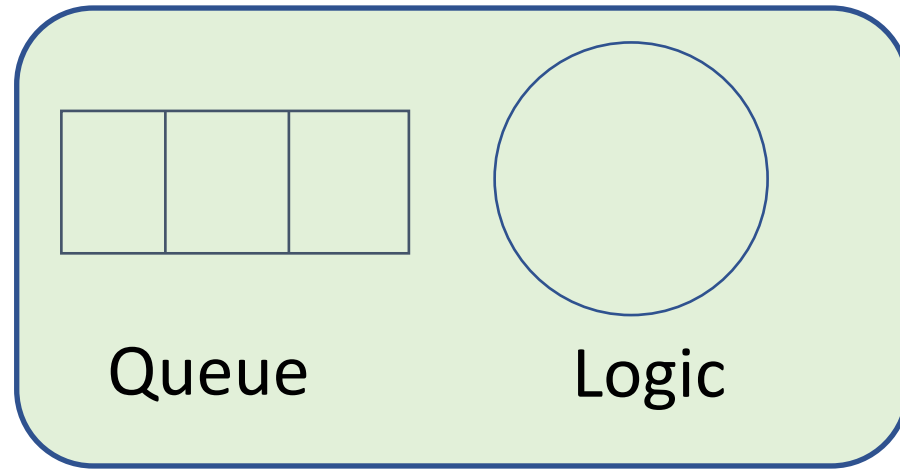
Impact of **scheduling** in **highly utilized datacenters**

- **Tail Task Completion Times**
- **Job Completion Time**

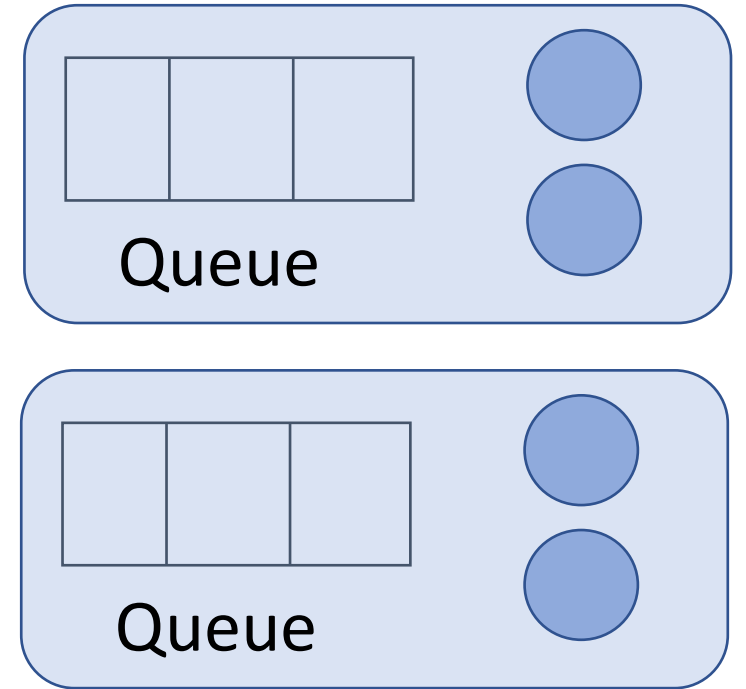
# Scheduling



Job

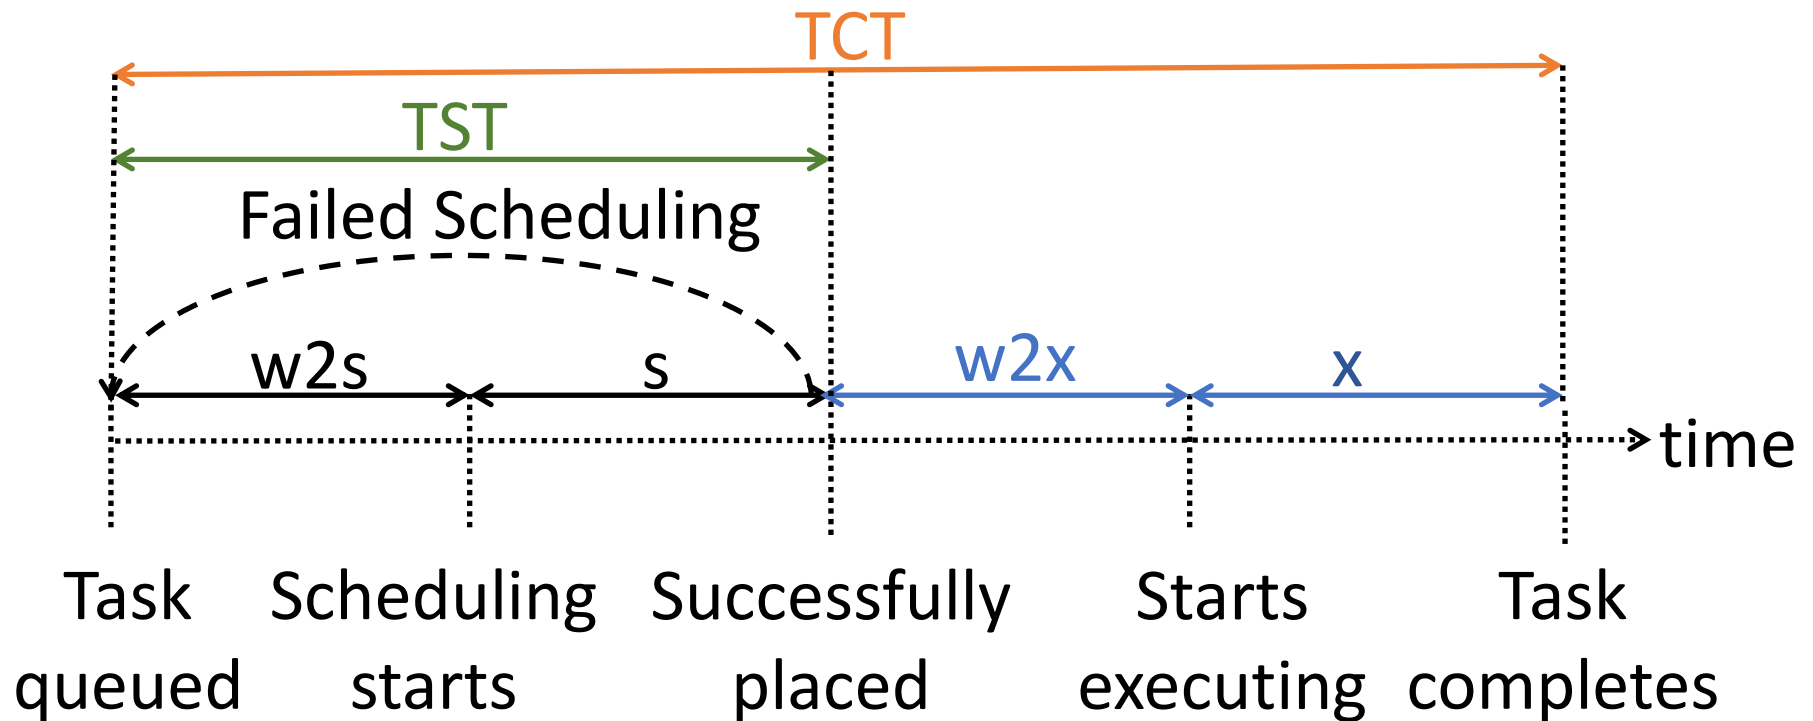
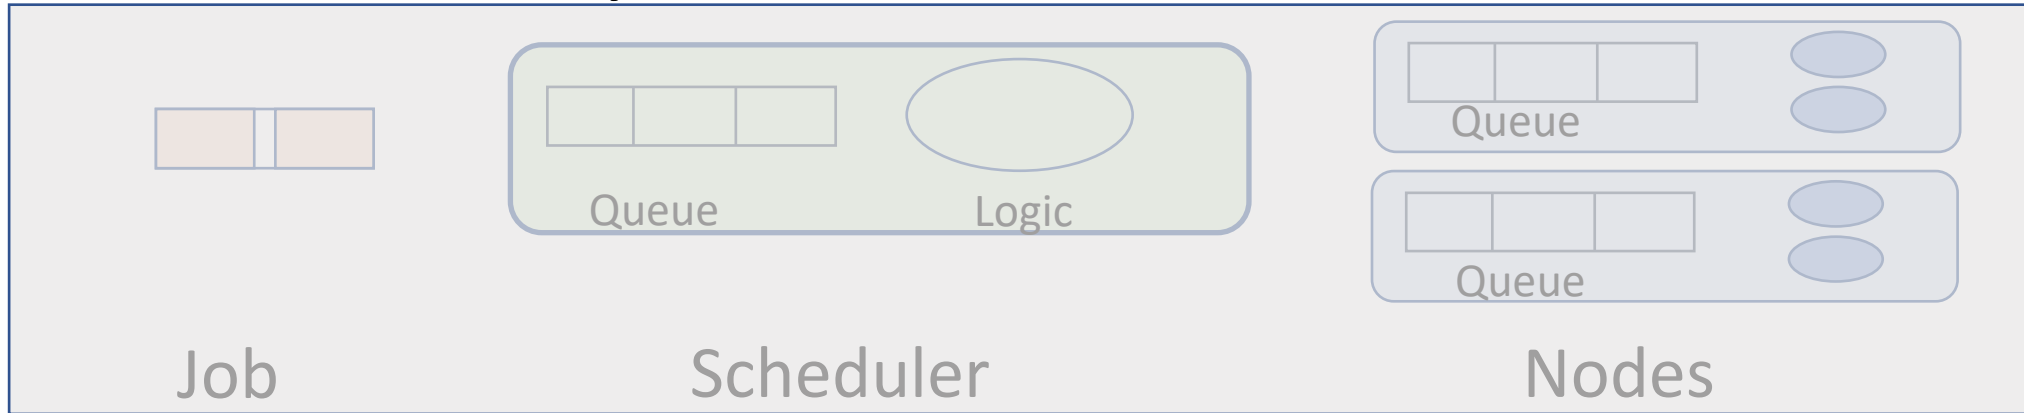


Scheduler

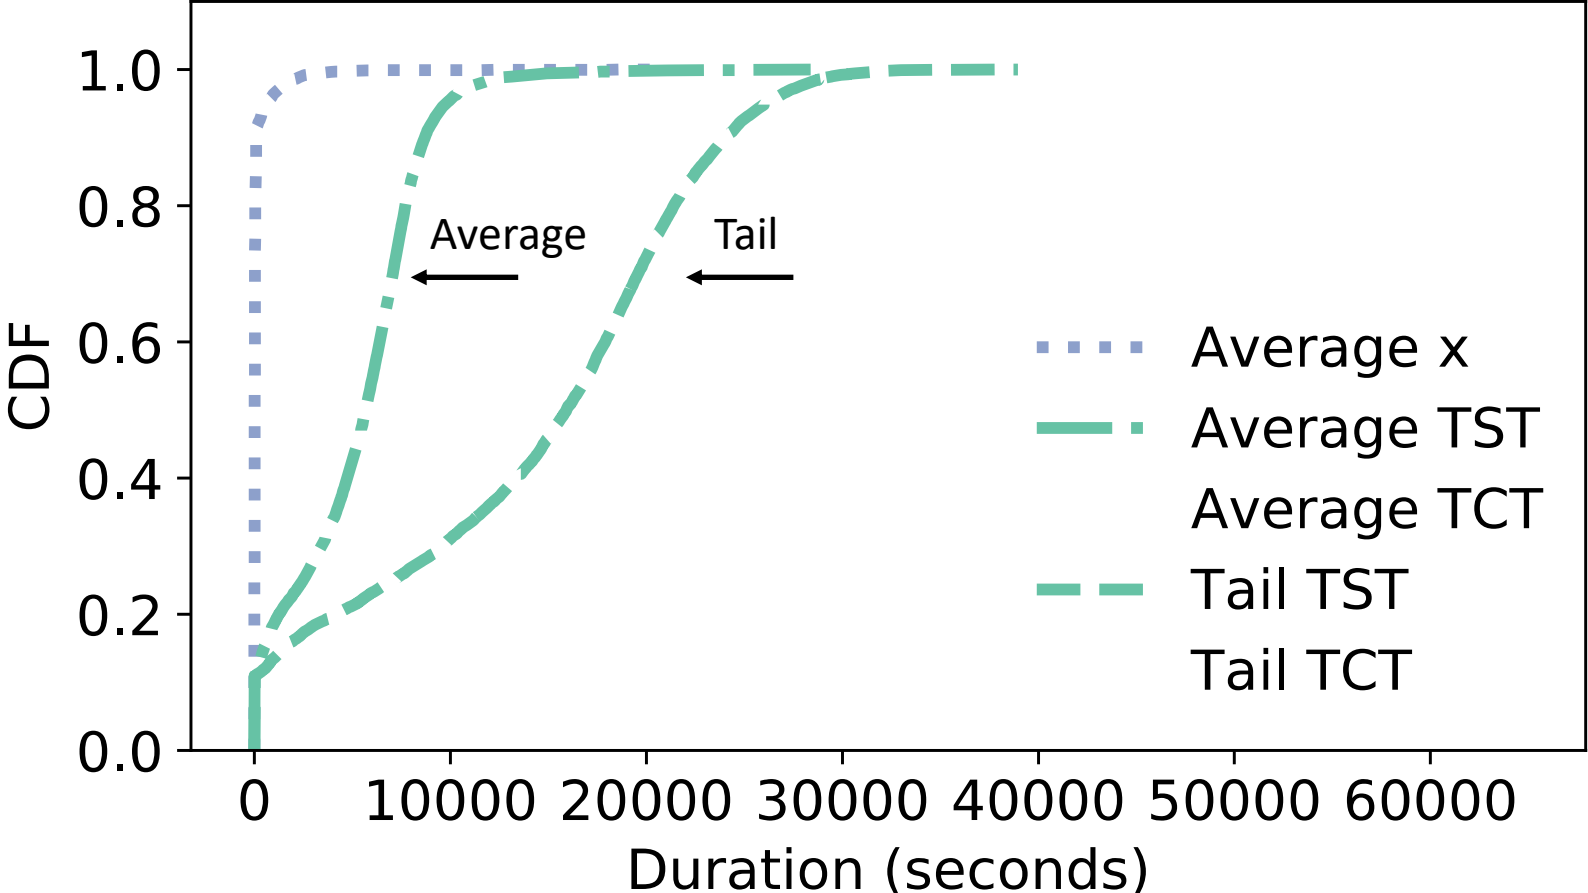


Nodes

# Task Lifecycle

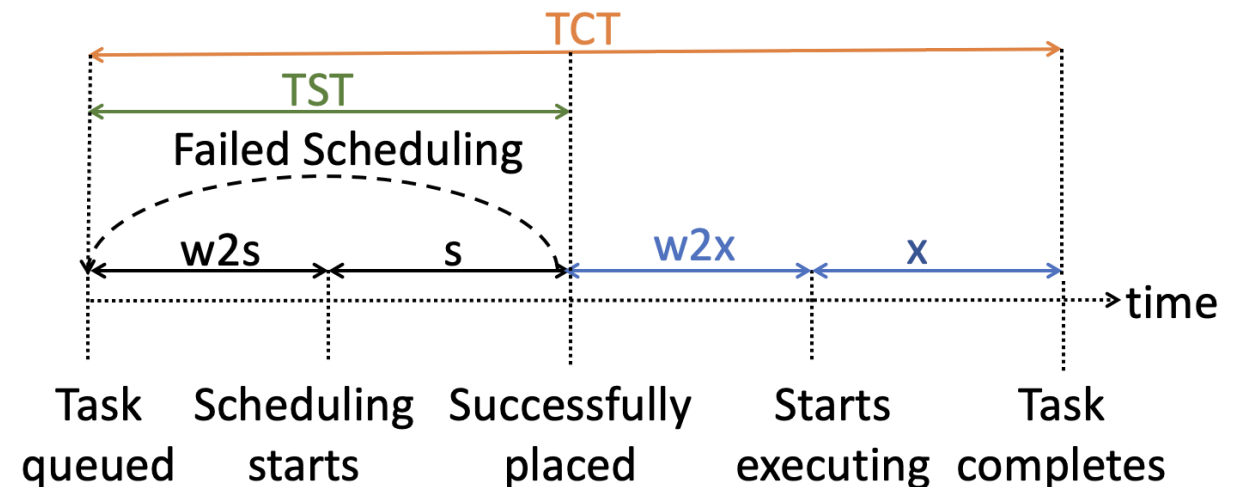


# Motivation



# Goals

- ↓ Job Completion Times
  - ↓ Tail Task Completion Times
    - Small Near-Constant Task Scheduling Time (TST)
    - Small **variance** in task start times ( $TST + w^2x$ )
- Scale to incoming rate
  - Small scheduler queue times ( $w^2s$ )



# Design Elements

Scheduling Logic Once Per Task



Small task scheduling times

Schedule Earliest Task First  
Using Estimated Node Wait Times




Smaller variance  
in task start times

Loosely coordinating decentralized  
schedulers



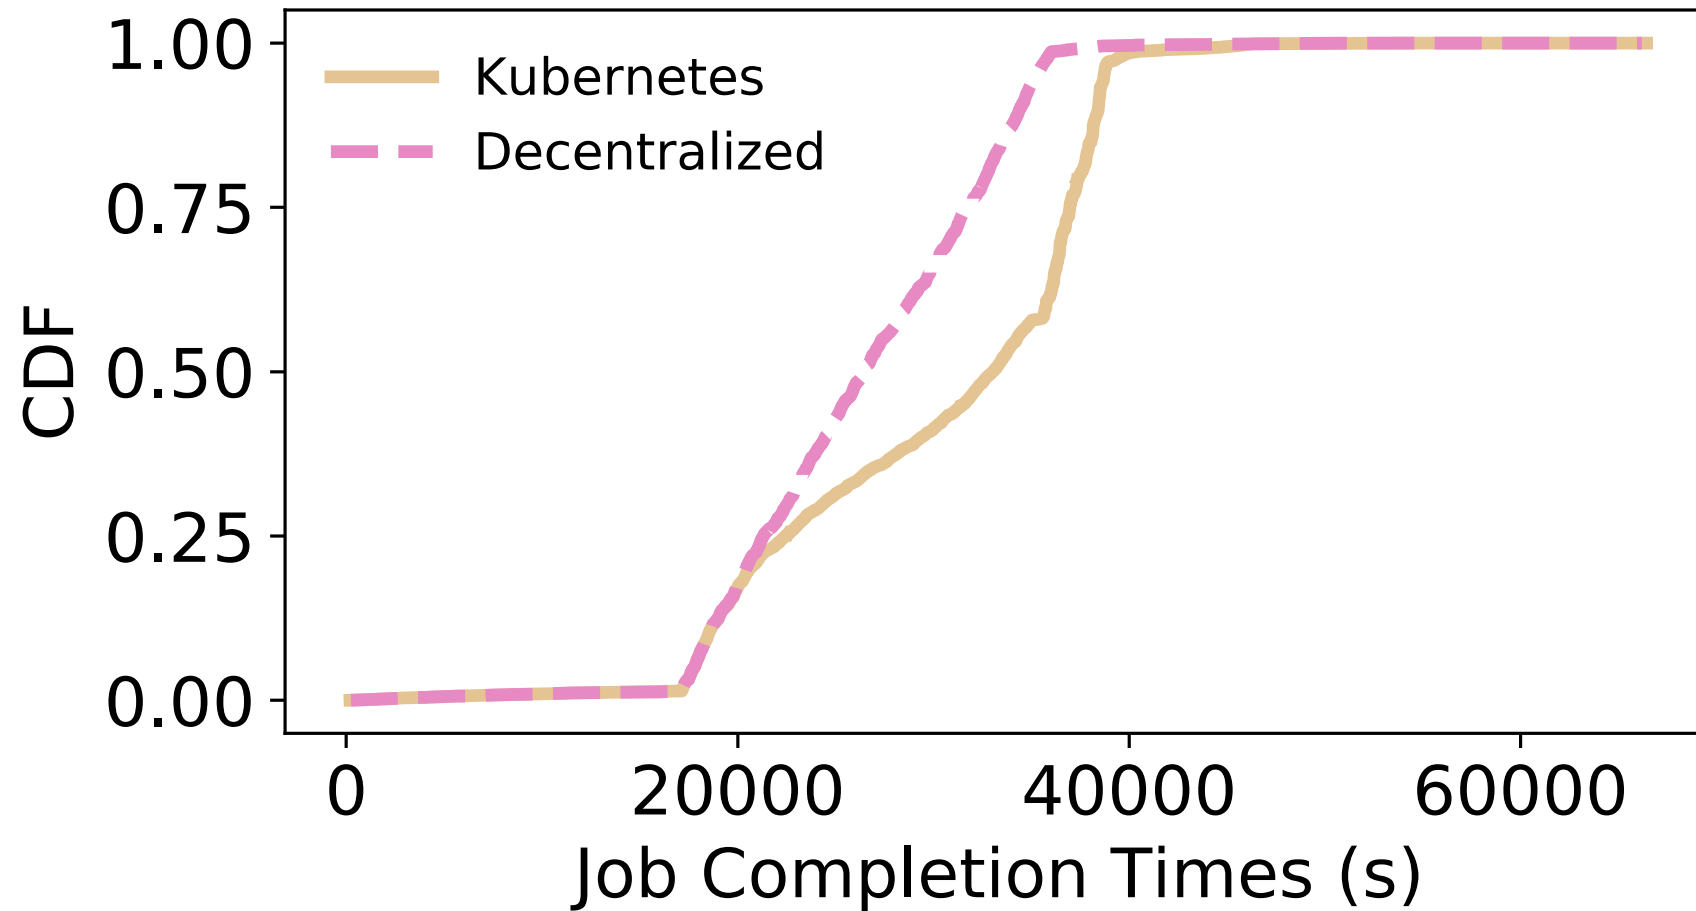
Scalability

# Results





# Results



# Conclusions

1. Schedulers affect completion times of tail tasks in highly utilized datacenters
2. Decentralized scheduling leads to better job completion times
  - 25% improvement in median JCT over kube-scheduler
  - 15-133% improvement in median JCT over current schedulers

Thank You!

Smita Vijayakumar

sv440@cam.ac.uk