


Towards Network Systems that Improve with Experience

Mohammad Alizadeh



Network Control Systems



Congestion Control



Traffic Engineering



Cluster Scheduling



Adaptive Video Streaming



Internet Telephony



Two Paradigms in Network Control

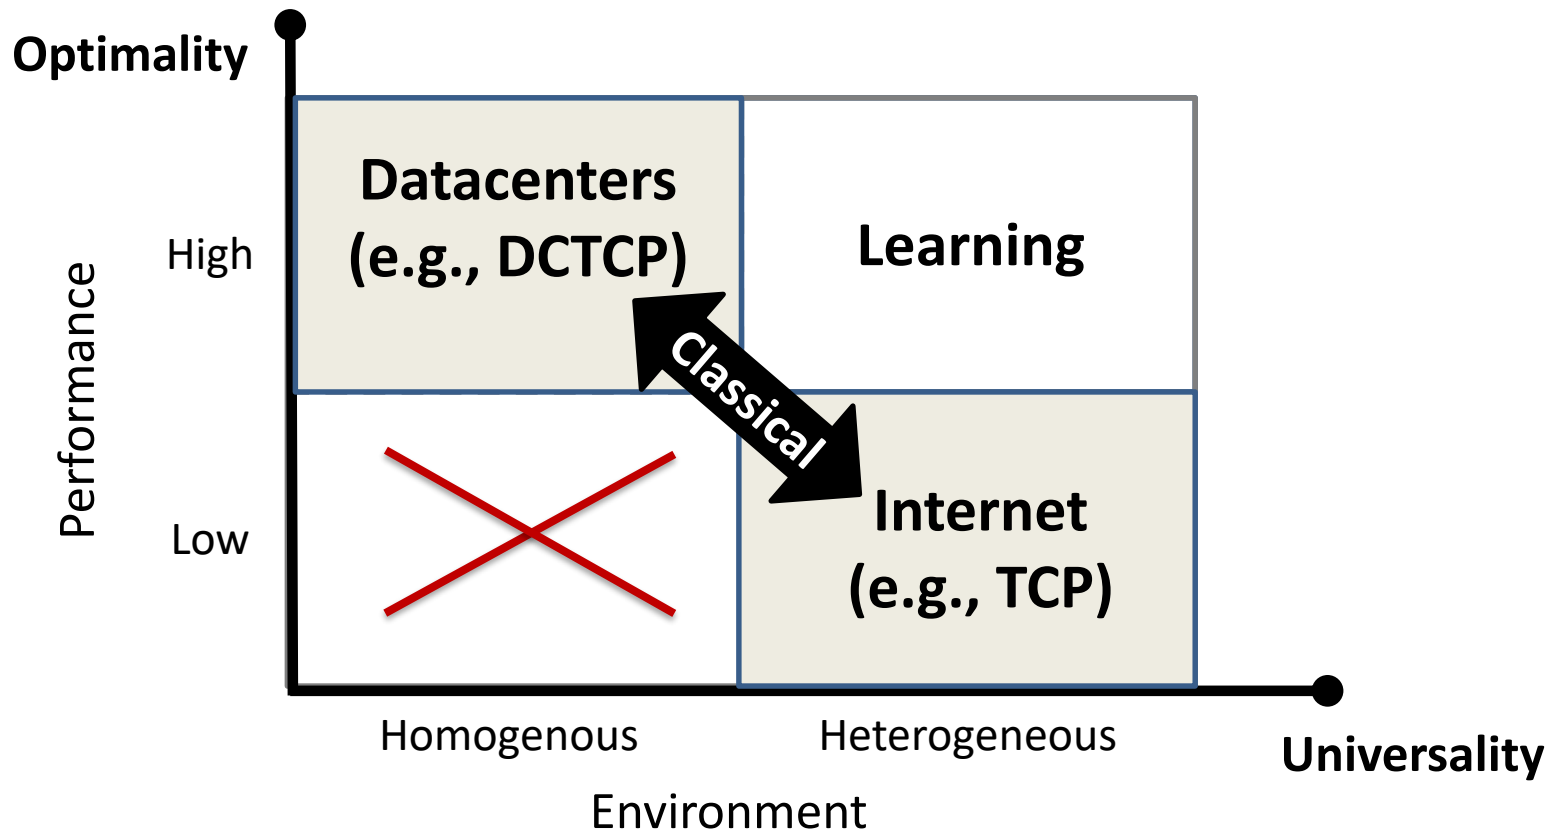
Classical Paradigm (1960 – now) [Specify & Build]

1. Specify operating environment & low-level design goals
2. Build one algorithm that achieves goals in (most) cases of interest

Learning Paradigm [Learn & Adapt]


1. Learn operating environment & low-level goals
2. Learn algorithms that can adapt to network/workload conditions

Two Paradigms in Network Control




Our Work on Learning-Based Network Systems


Adaptive video streaming (Pensieve)



Cluster Scheduling (Decima)

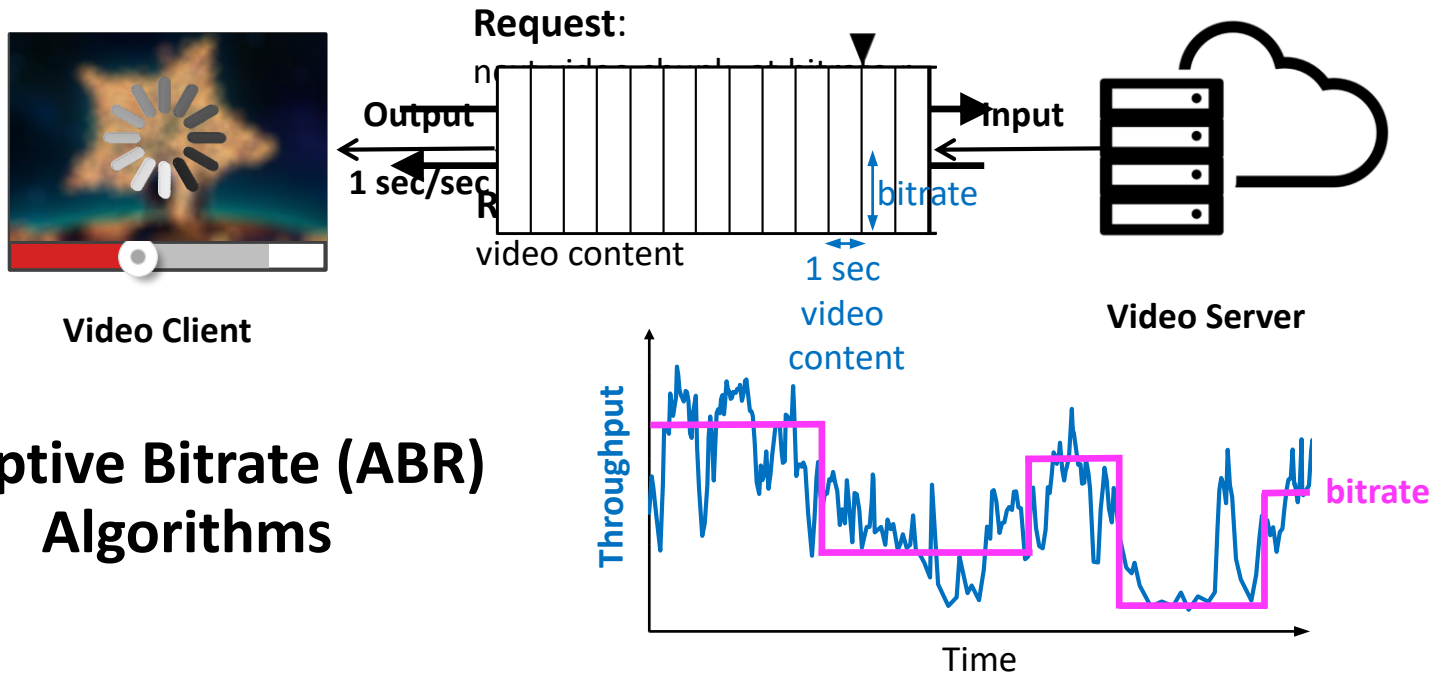


RL Techniques




Users start leaving if video doesn't play in 2 seconds

Dynamic Streaming over HTTP (DASH)



Adaptive Bitrate (ABR) Algorithms

Why is ABR Challenging?



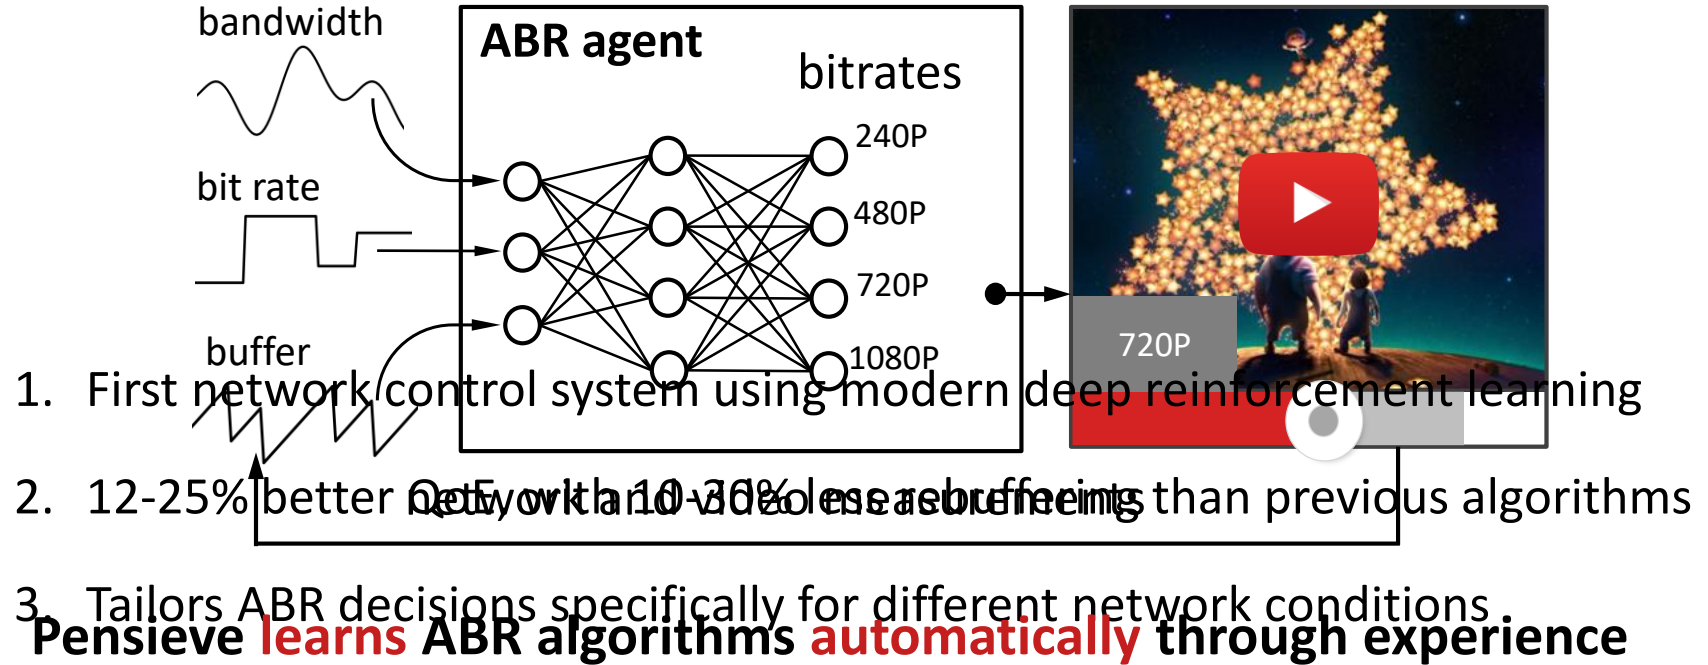
Network throughput
is variable & uncertain

Conflicting QoE goals


- Bitrate
- Rebuffering time
- Smoothness

Cascading effects
of decisions

Our System: Pensieve



Reinforcement Learning




Goal: maximize expected total future reward $\mathbb{E} \left[\sum_t r_t \right]$

Pensieve Design

State s_t

Agent

Reward r_t

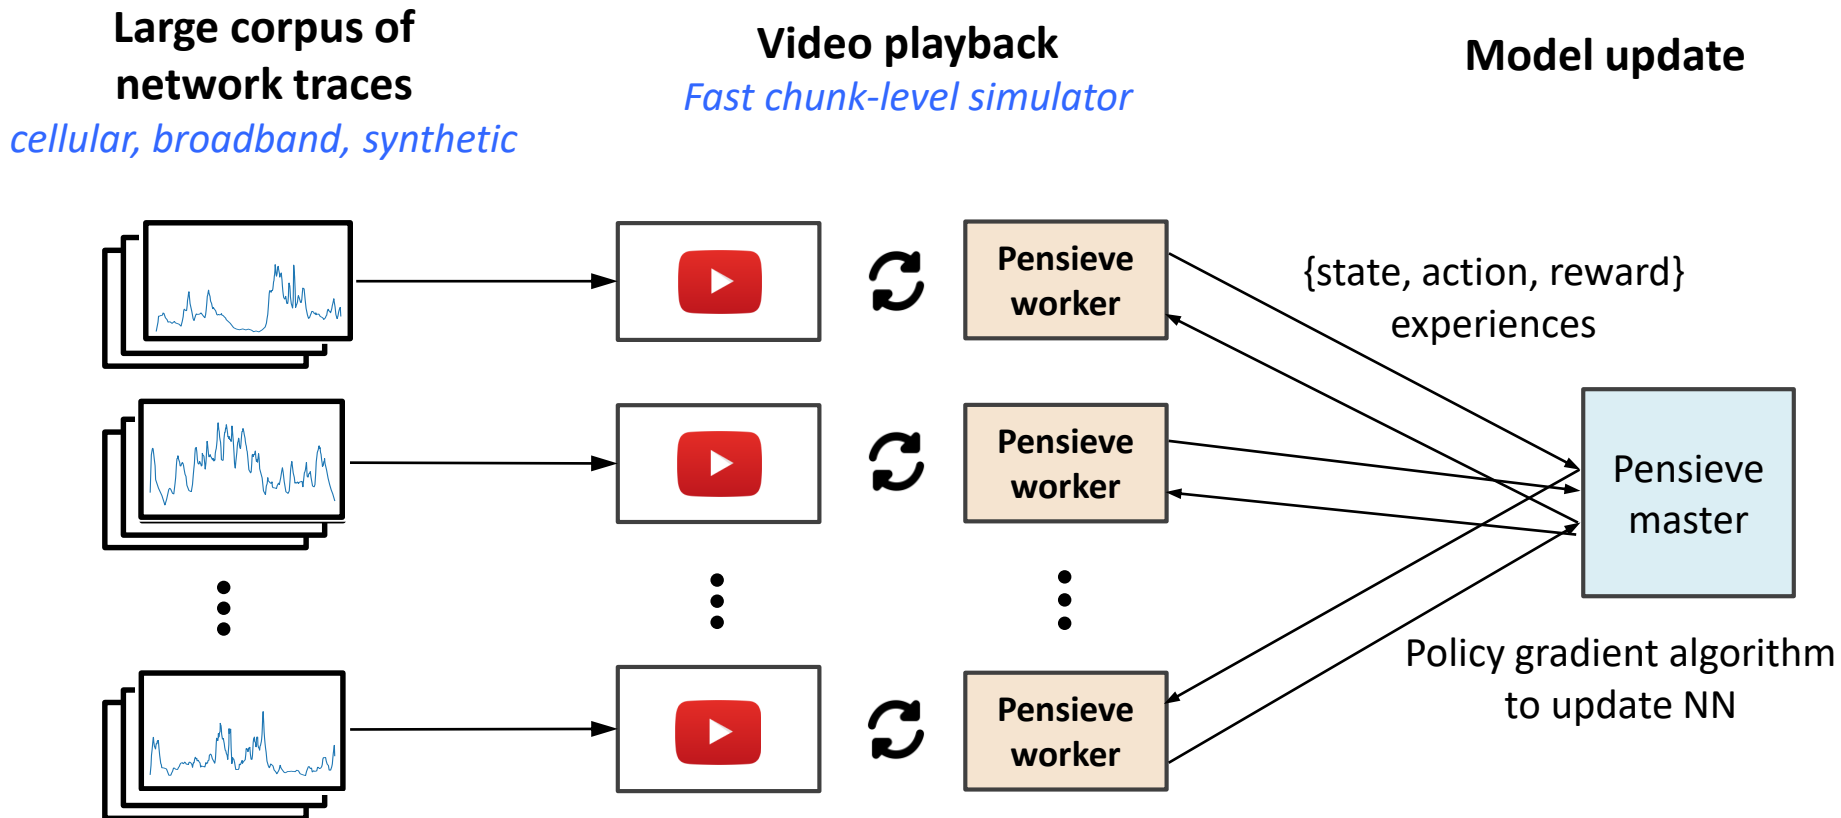


Action a_t

Environment

720P


Pensieve Training System



Pensieve

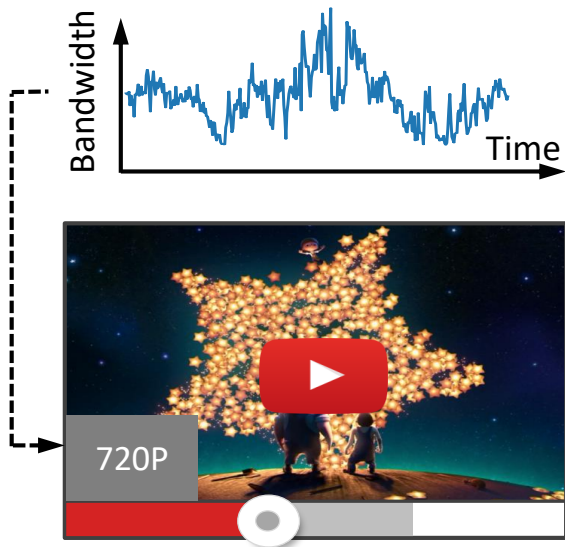


MPC




Our Work on Learning-Based Network Systems

Adaptive video streaming (Pensieve)



Cluster Scheduling (Decima)

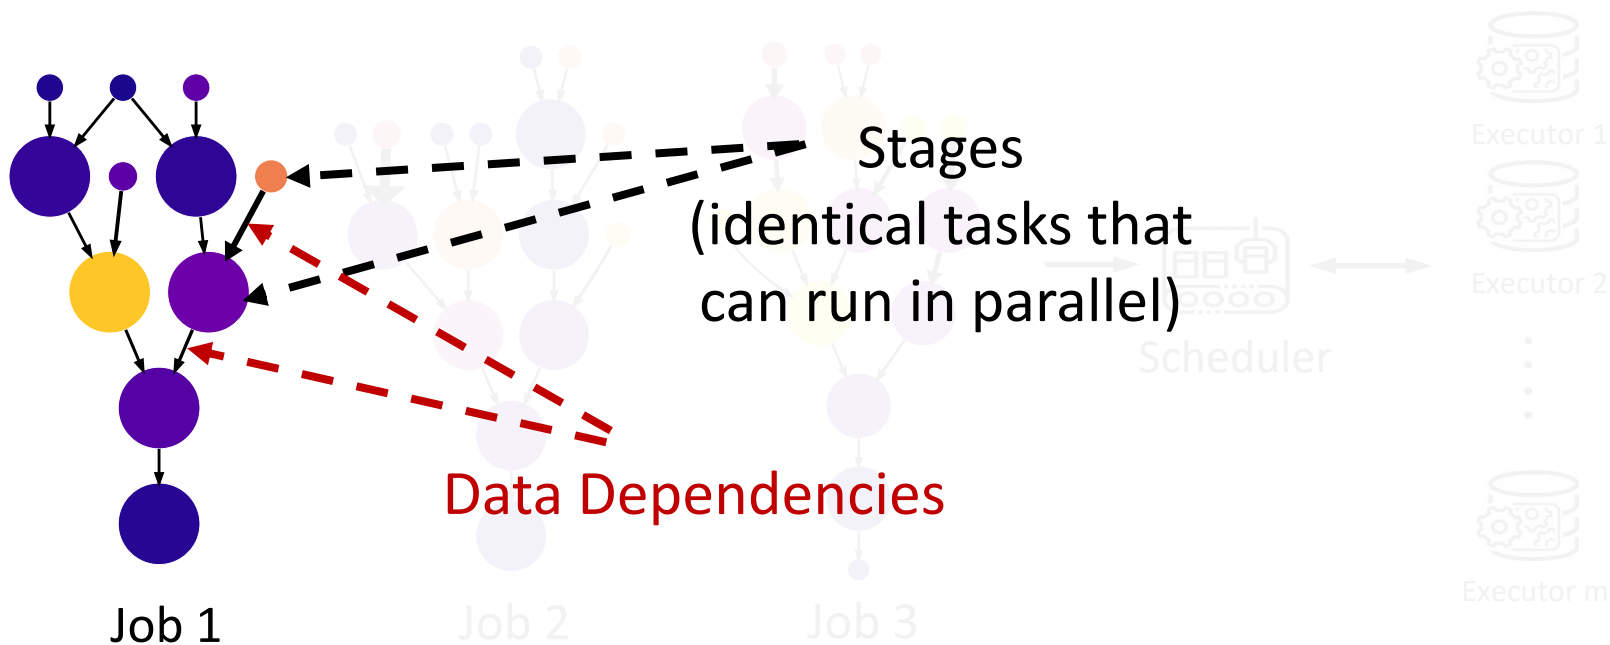


RL Techniques

Scheduling Graph-Structured Workloads

Many systems encode jobs as directed acyclic graphs (DAGs)

- Data processing (e.g., Spark, Hadoop), ML training (e.g., TensorFlow), ...



Designing Optimal Schedulers is Intractable

A lot of considerations for optimal performance:

- Job dependency structure
- Degree of parallelism
- Scheduling order
- Locality
- ...

No “one-size-fits-all” solution:

Best algorithm depends on specific **workload** and **system**

Decima: Technical Challenges

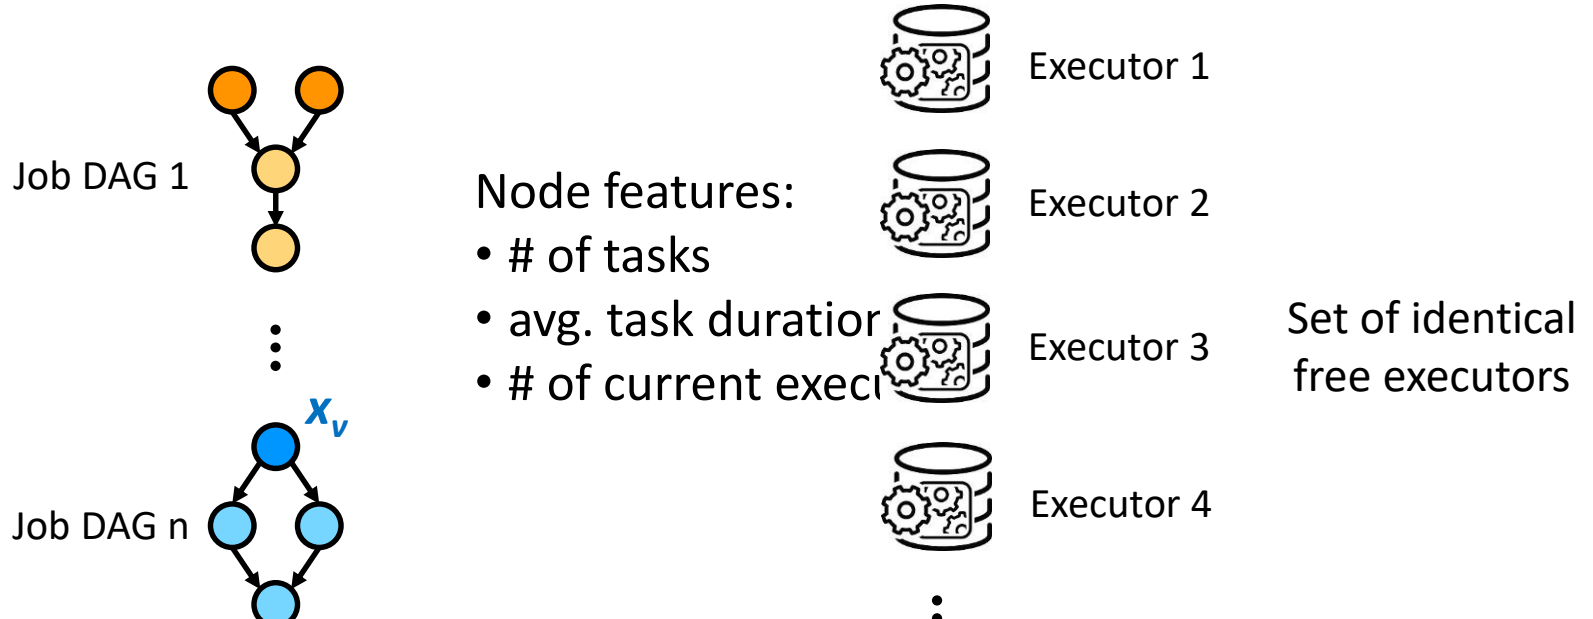
Challenge: Huge state and action space

→ Scalable Graph CNN to process any number of job DAGs

Challenge: Variance caused by stochastic job arrival process


→ Variance reduction technique for RL in input-driven systems

Decima Design




How to encode scheduling decisions as actions?

Option 1: Assign all Executors in 1 Action




Problem: huge action space

Option 2: Assign 1 Executor per Action

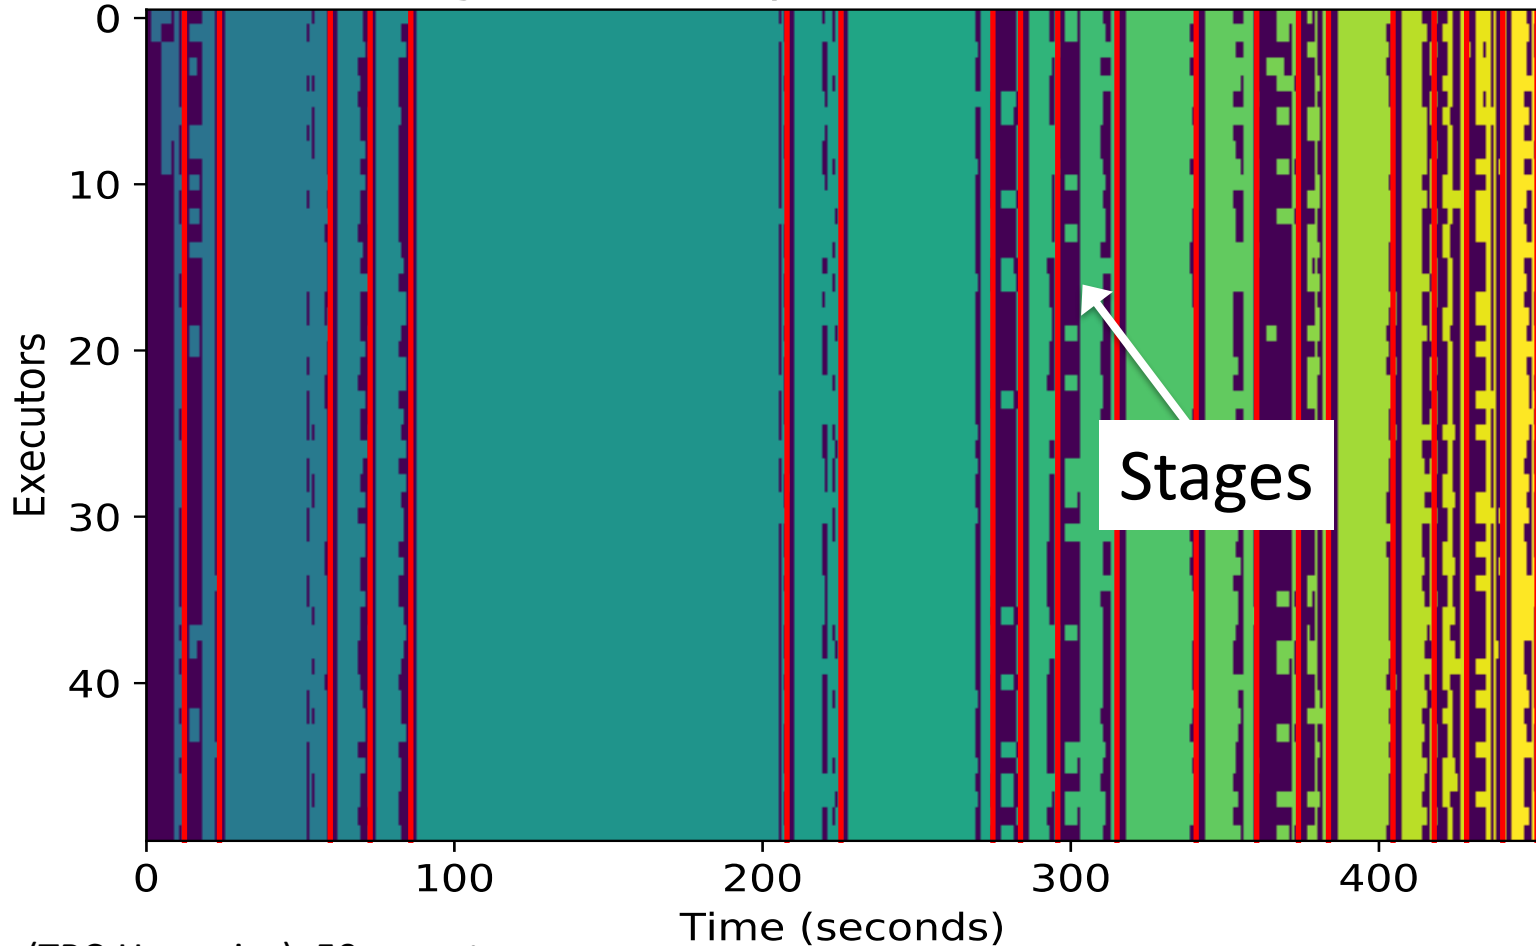


Decima: Assign Groups of Executors per Action



FIFO

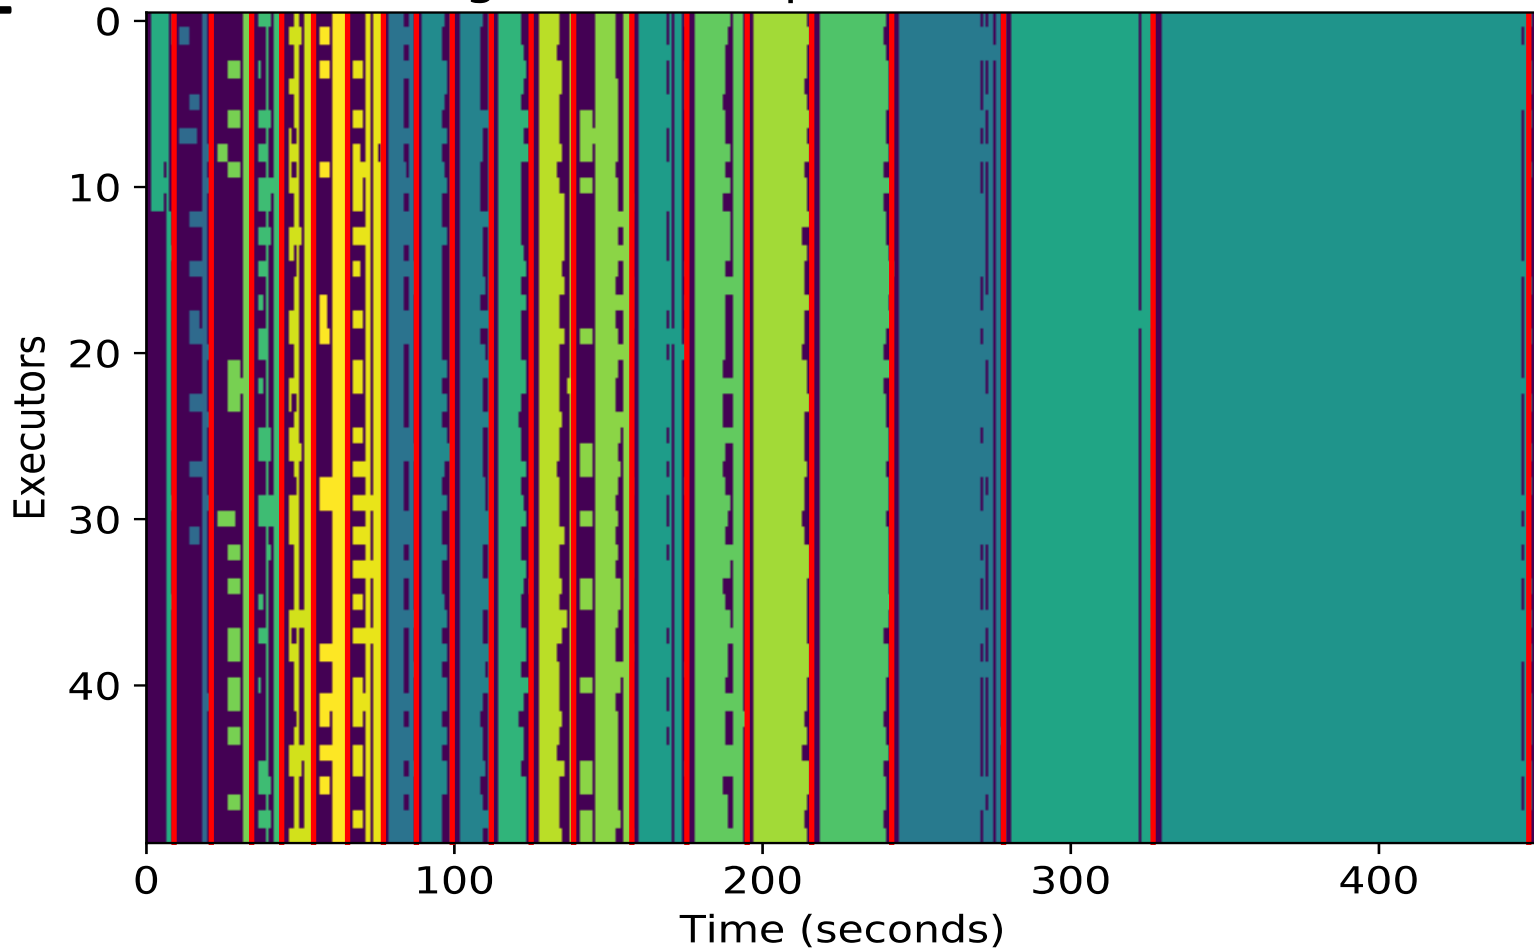
average DAG completion time: 272 sec



✧ 20 Spark jobs (TPC-H queries), 50 executors


Shortest- Job-First

average DAG completion time: 145 sec




Fair

average DAG completion time: 110 sec




average DAG completion time: 166 sec

Decima
it=0




average DAG completion time: 160 sec

Decima
it=3000




average DAG completion time: 148 sec

Decima
it=6000




Decima
it=9000

average DAG completion time: 145 sec




Decima
it=12000

average DAG completion time: 142 sec




Decima
it=15000

average DAG completion time: 126 sec




Decima
it=18000

average DAG completion time: 111 sec




average DAG completion time: 108 sec

Decima
it=21000




average DAG completion time: 107 sec

Decima
it=24000




Decima
it=27000

average DAG completion time: 93 sec




Decima
it=30000




Decima improves average job completion time by >2x over existing schedulers, and 19-31% over best hand-crafted heuristics

Our Work on Learning-Based Network Systems

Adaptive video streaming (Pensieve)





Cluster Scheduling (Decima)



RL Techniques


Input-Driven Environments

Dynamics driven by an exogenous stochastic **input process**




Input-Driven Environments

Dynamics driven by an exogenous stochastic **input process**



The input does not depend on the states and actions

Reinforcement Learning for Input-Driven MDPs



Must take the future input into account
when evaluating actions


Input-Dependent Baseline

Score for action $a_t = \sum_{t'=t}^{T-1} r_{t'} - b(s_t)z_t, z_{t+1}, \dots$


Expected return for trajectories from state s_t
with input sequence z_t, z_{t+1}, \dots

- Input-dependent baselines reduce variance without bias
- We use meta-learning to learn baseline efficiently

Walker2d with Wind




TRPO with standard baseline




TRPO with input-dependent baseline

Half-Cheetah on Floating Tiles




TRPO with standard baseline




TRPO with input-dependent baseline

7-DoF Arm Tracking Moving Object



TRPO with standard baseline



TRPO with input-dependent baseline

Summary

- Significant opportunity to build smarter networks that can adapt to workload and environment
- Network systems are an exciting domain for ML/AI; can lead to new techniques with broad applications
- Many challenges remain for learning-based systems
 - Safe training & exploration
 - Interpretability
 - ...

