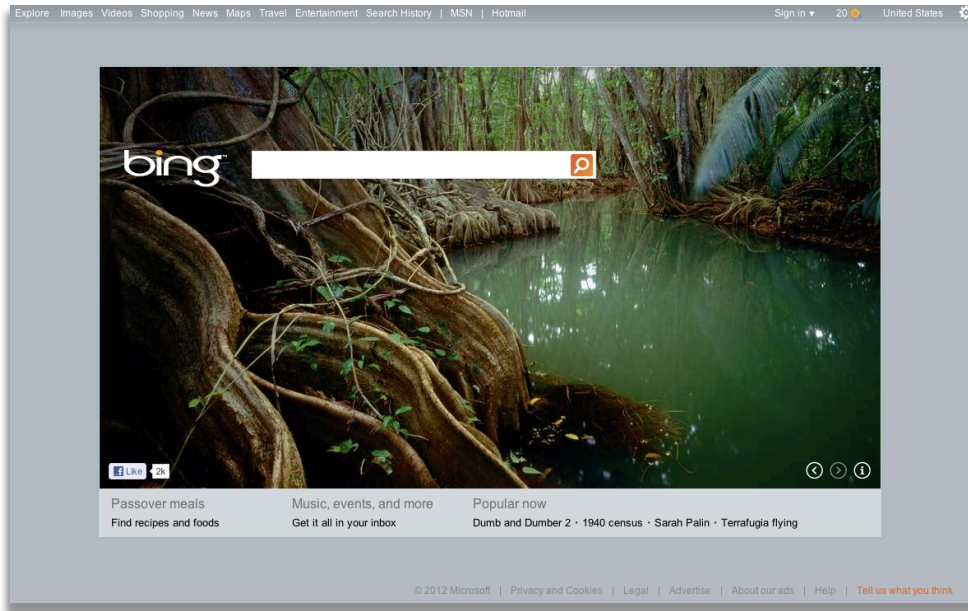


# Jockey

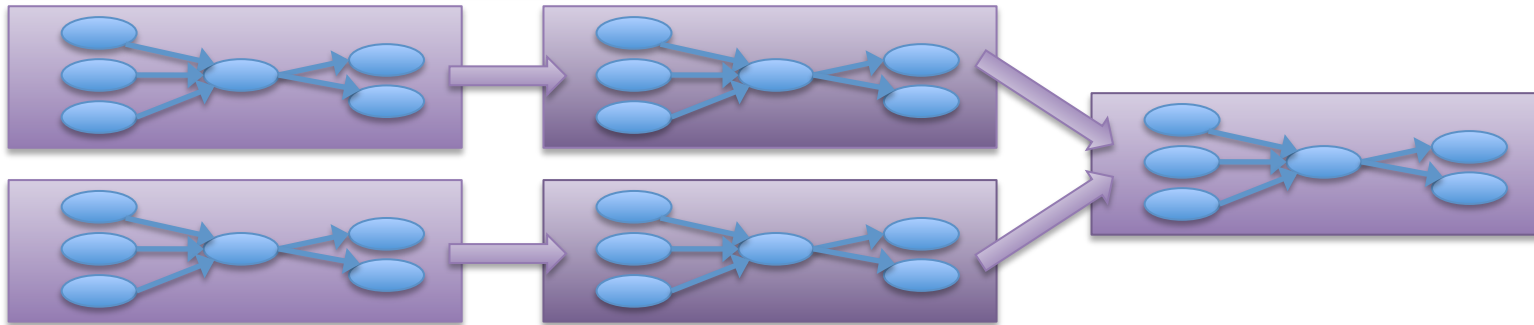
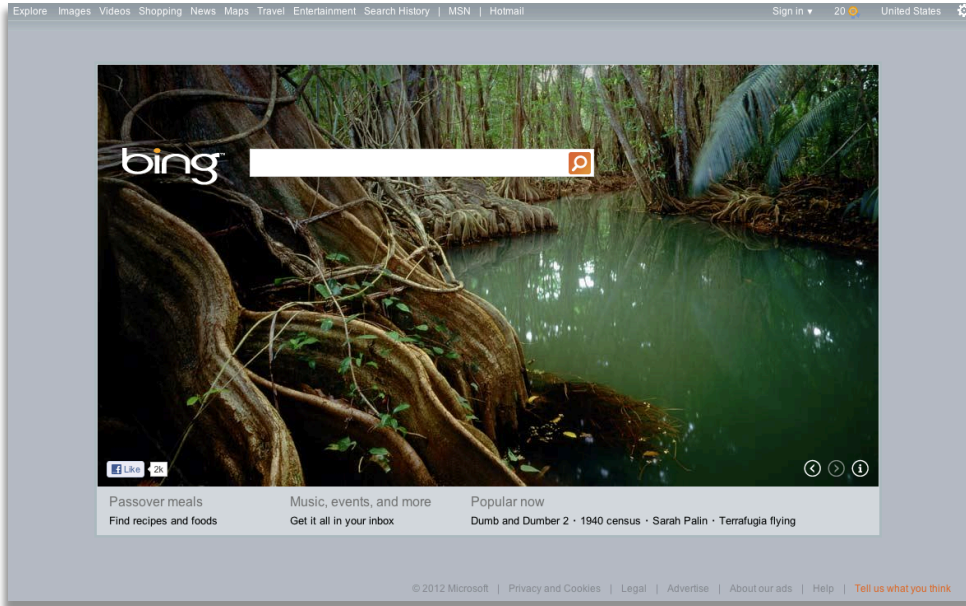
## Guaranteed Job Latency in Data Parallel Clusters

Andrew Ferguson, Peter Bodik, Srikanth  
Kandula, Eric Boutin, and Rodrigo Fonseca

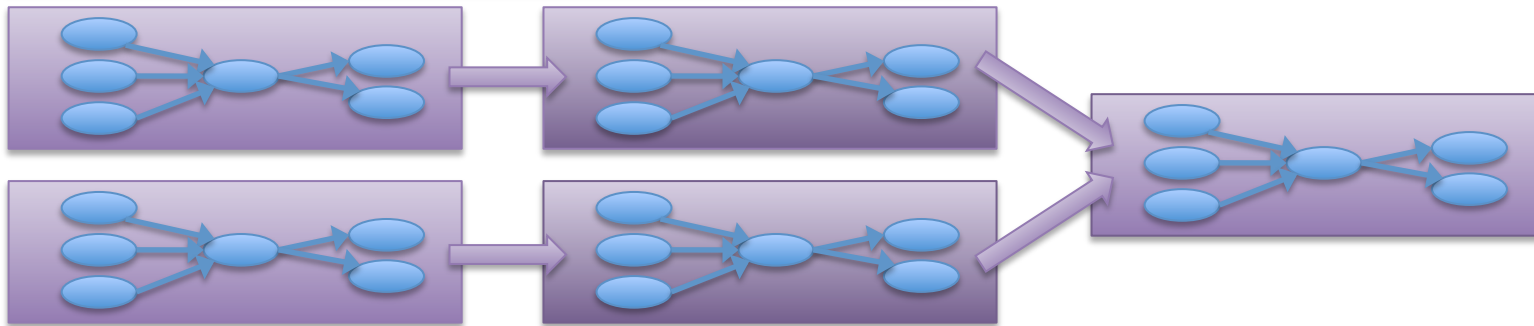
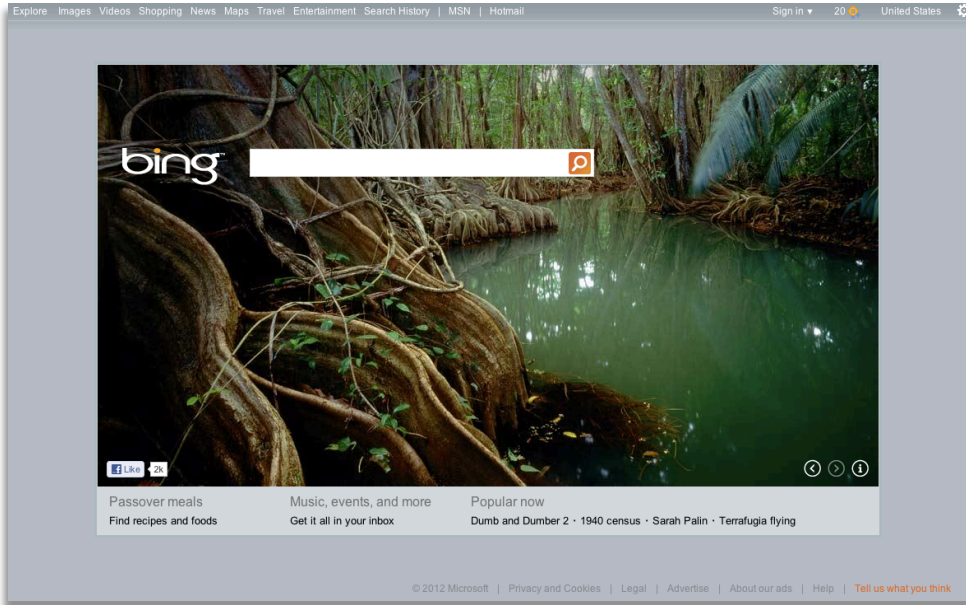




# DATA PARALLEL CLUSTERS

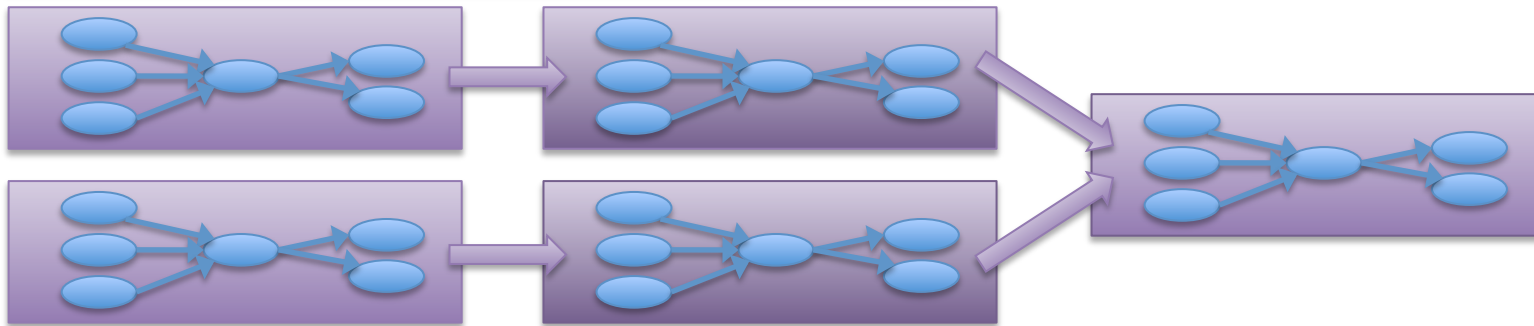
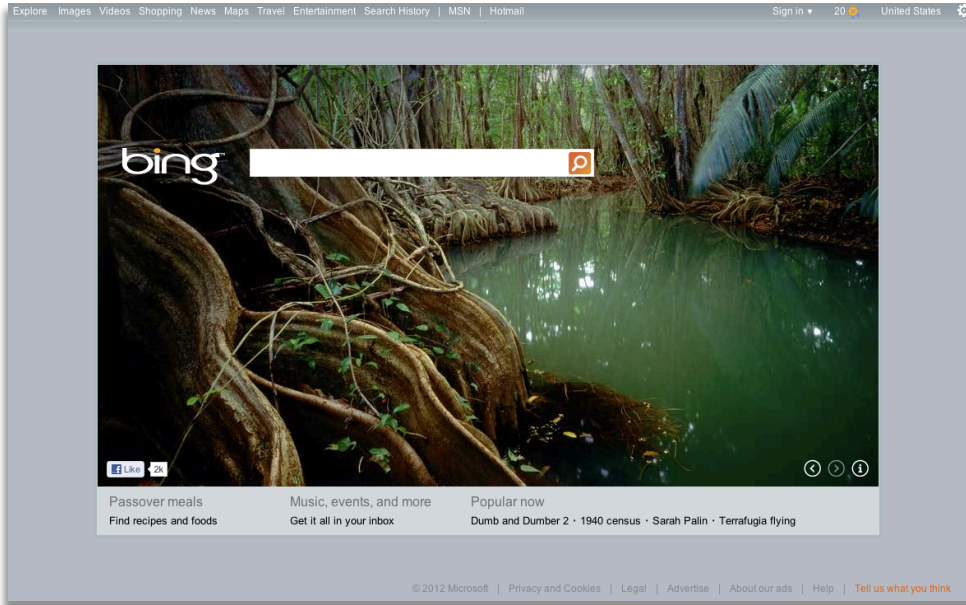


# DATA PARALLEL CLUSTERS



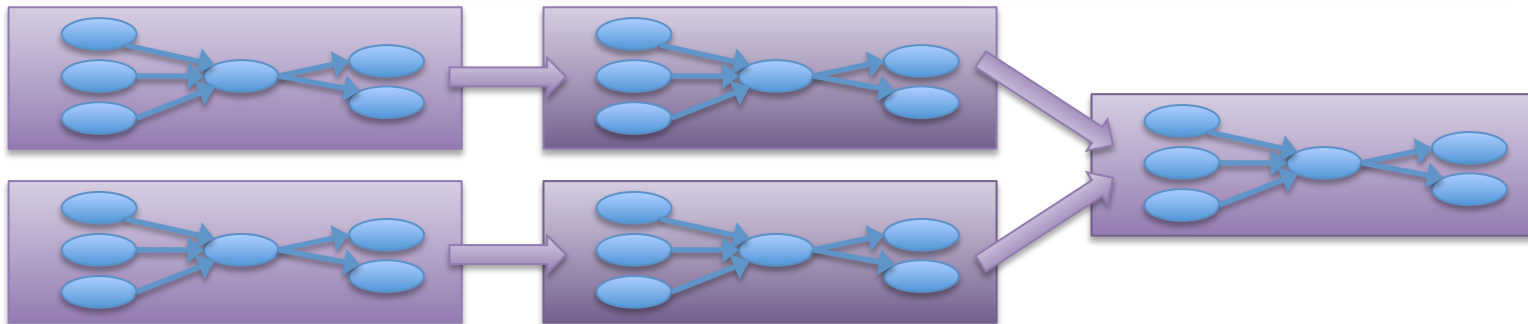
Predictability

# DATA PARALLEL CLUSTERS



Deadline

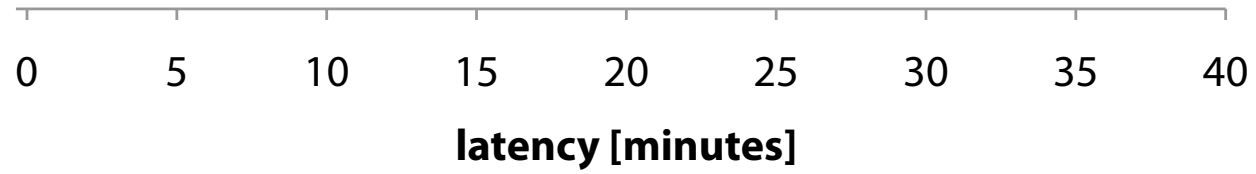
# DATA PARALLEL CLUSTERS



Deadline

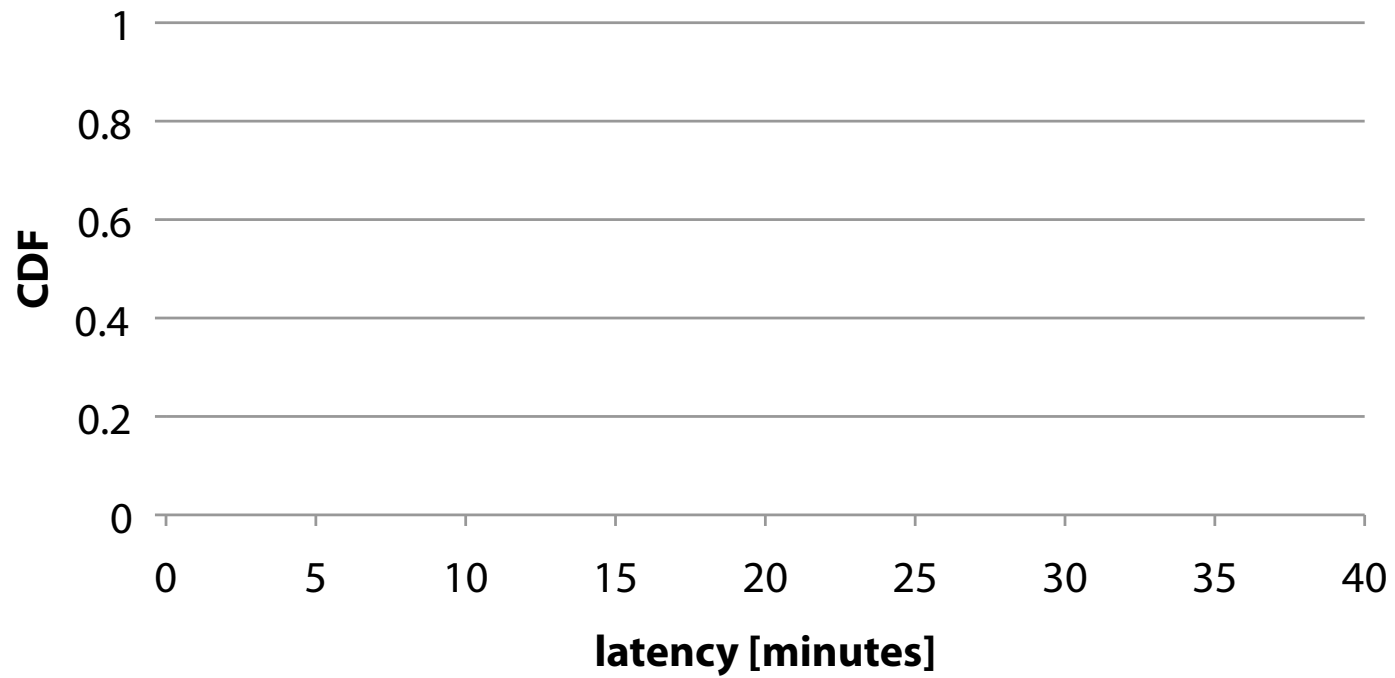
# DATA PARALLEL CLUSTERS

# VARIABLE LATENCY

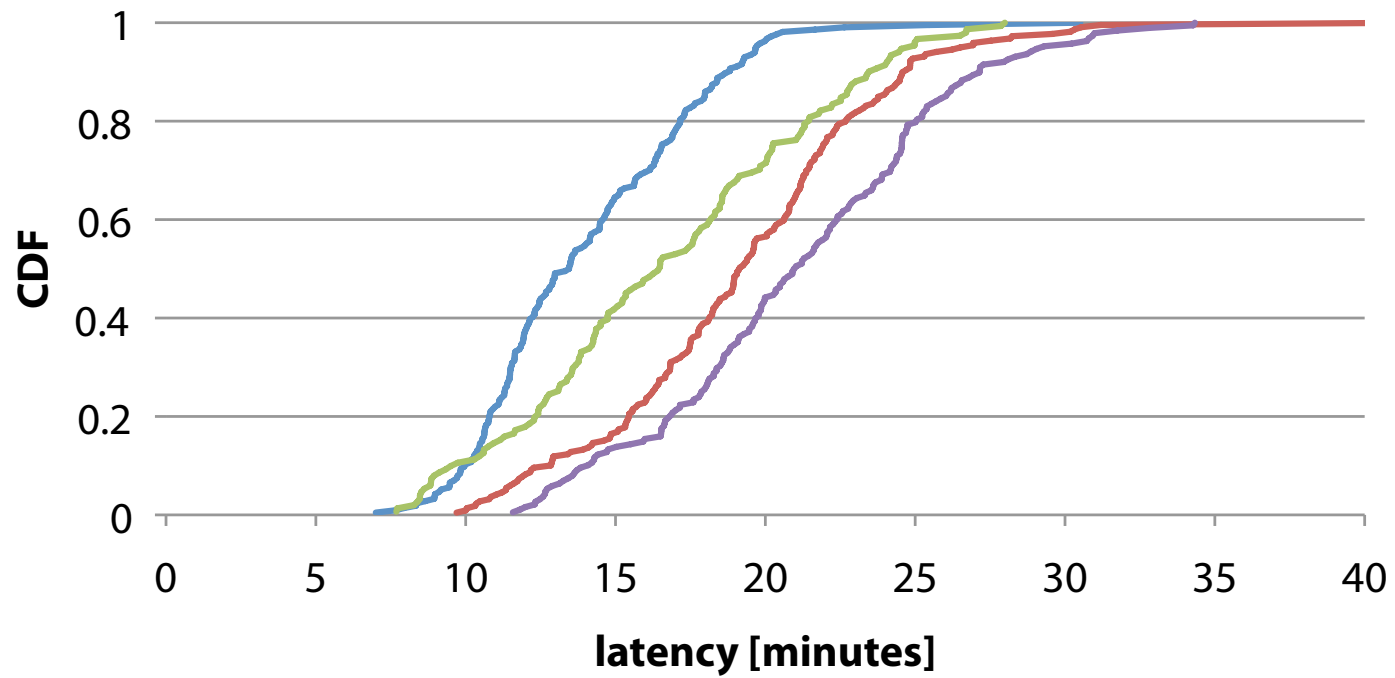


# VARIABLE LATENCY

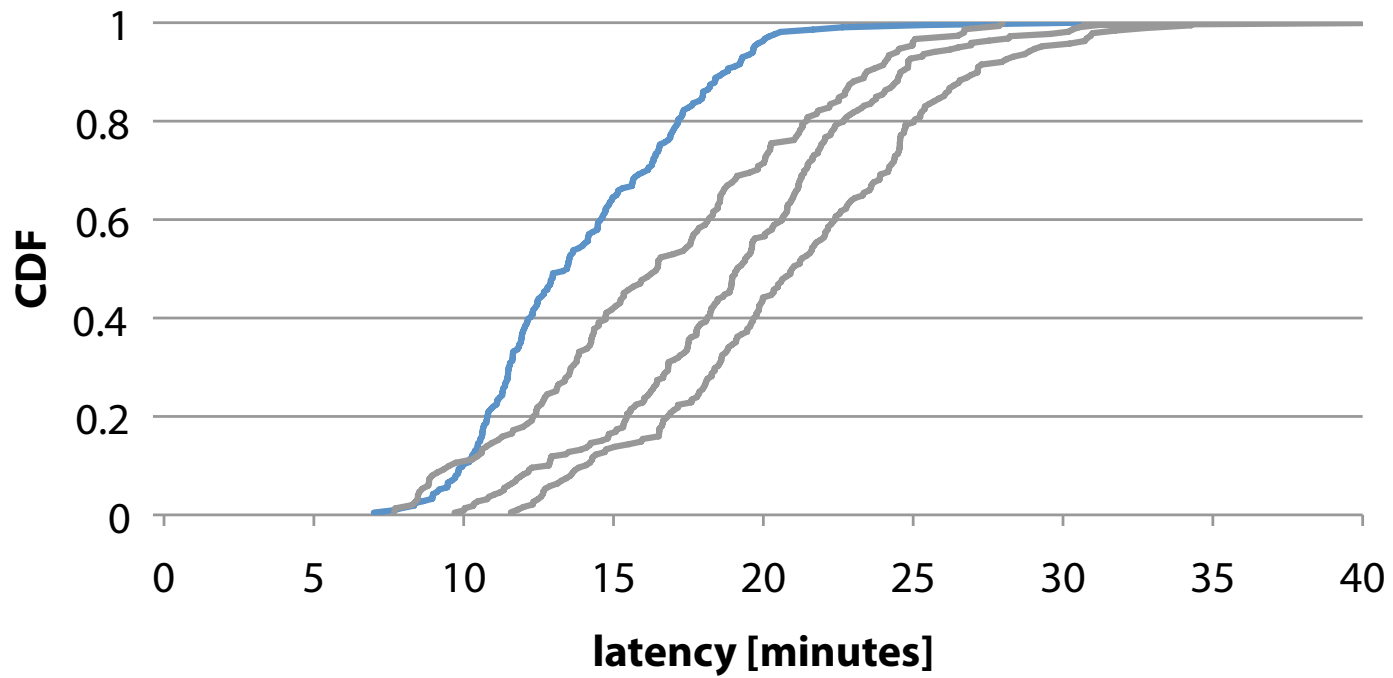




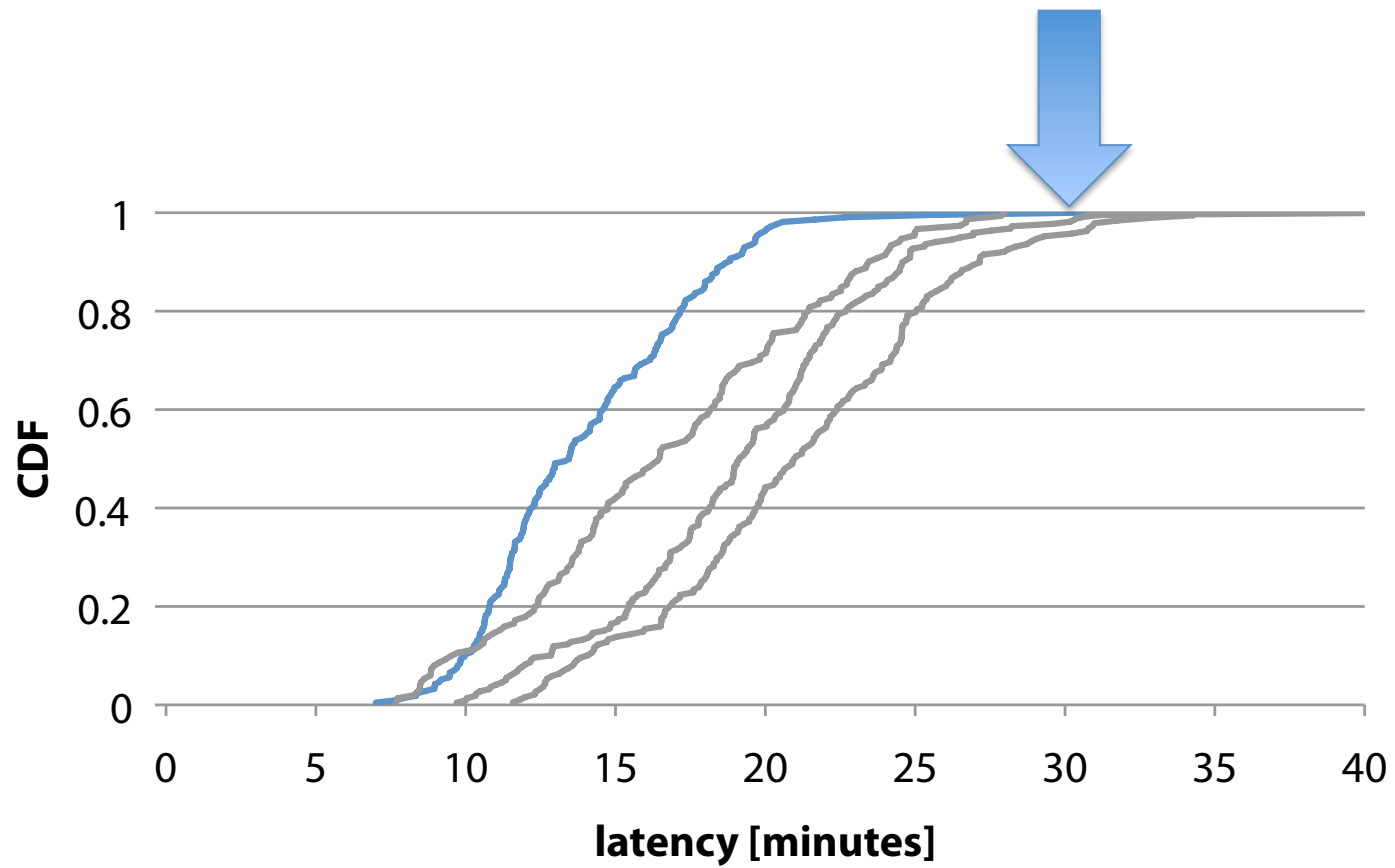
# VARIABLE LATENCY



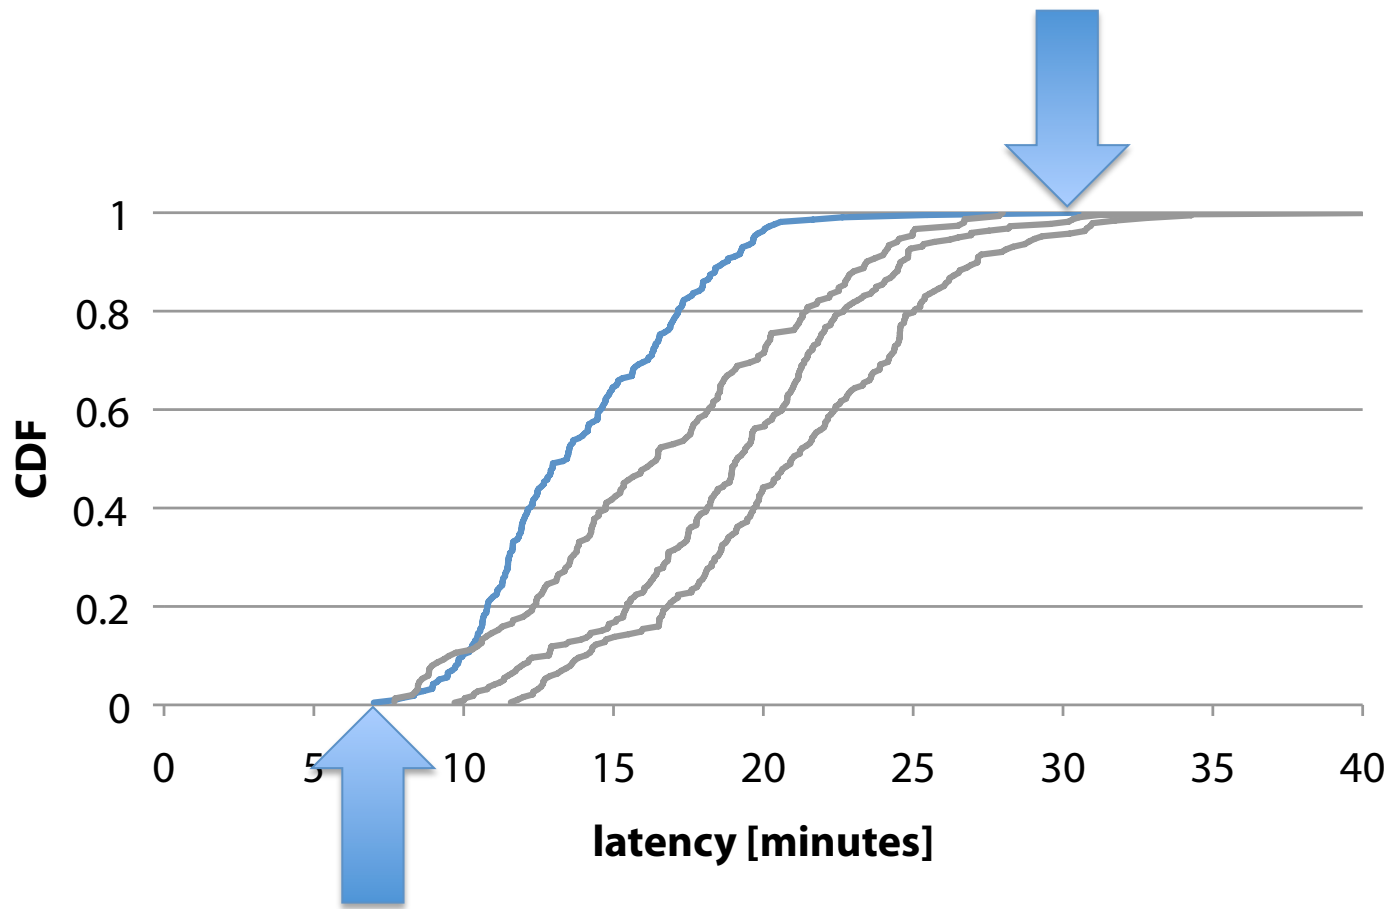
# VARIABLE LATENCY



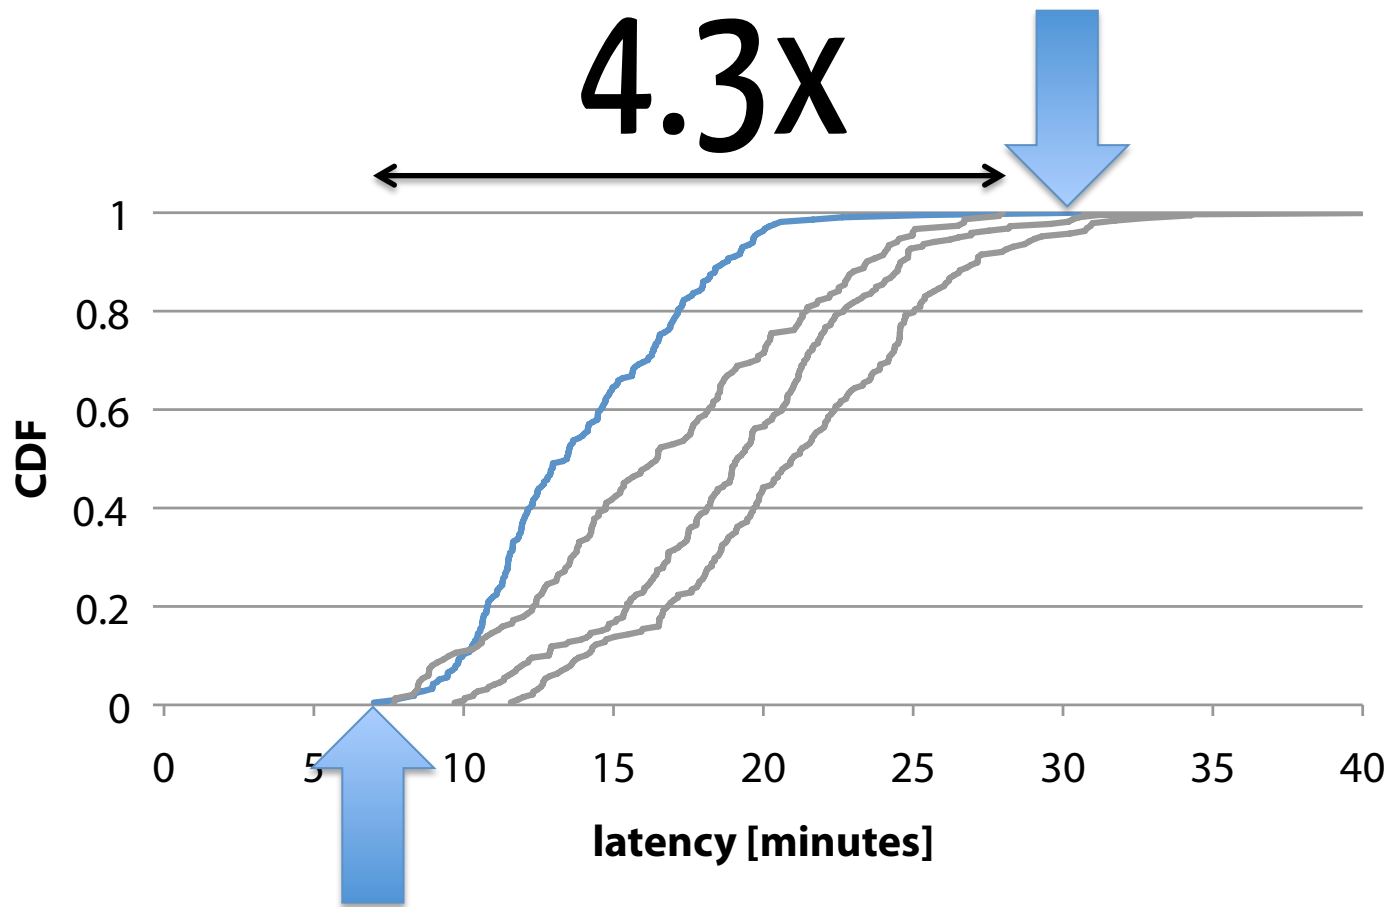
# VARIABLE LATENCY



# VARIABLE LATENCY



# VARIABLE LATENCY



**VARIABLE LATENCY**

# **Why does latency vary?**

- 1. Pipeline complexity**
- 2. Noisy execution environment**

# Cosmos



**MICROSOFT'S DATA PARALLEL CLUSTERS**



- **CosmosStore**



# Cosmos

## MICROSOFT'S DATA PARALLEL CLUSTERS

- **CosmosStore**
- **Dryad**



# **Cosmos**

## **MICROSOFT'S DATA PARALLEL CLUSTERS**

- **CosmosStore**
- **Dryad**
- **SCOPE**



# **Cosmos**

## **MICROSOFT'S DATA PARALLEL CLUSTERS**

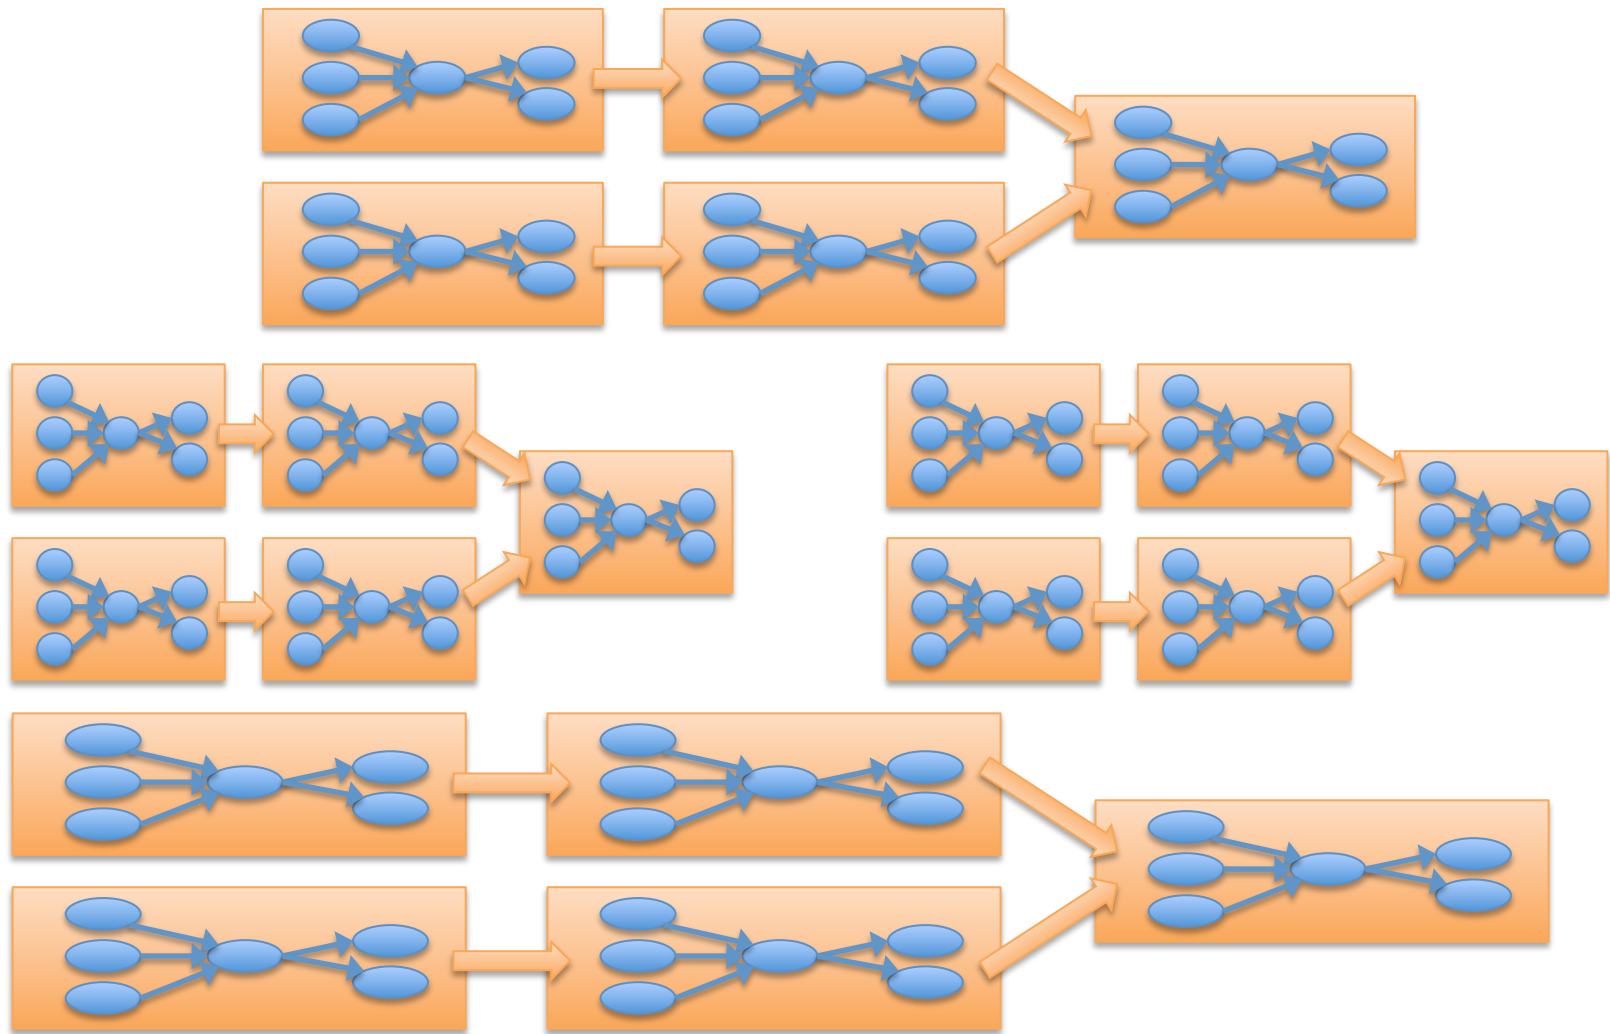
- CosmosStore
- **Dryad**
- SCOPE



# Cosmos

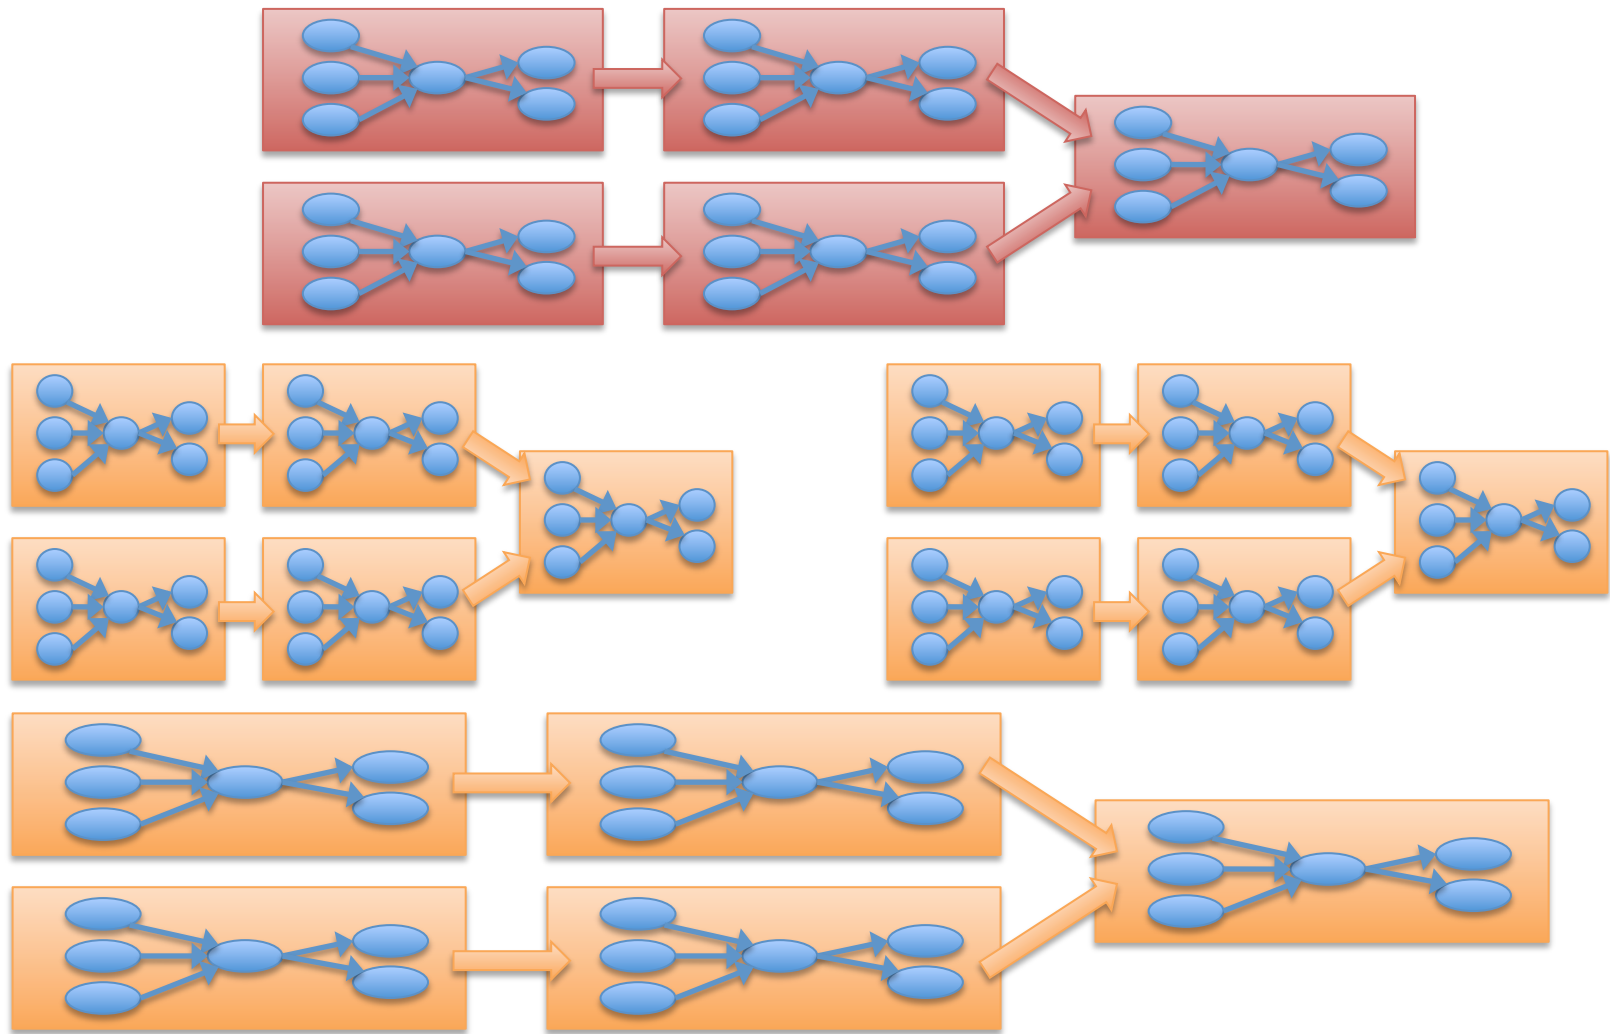
## MICROSOFT'S DATA PARALLEL CLUSTERS

Cosmos Cluster

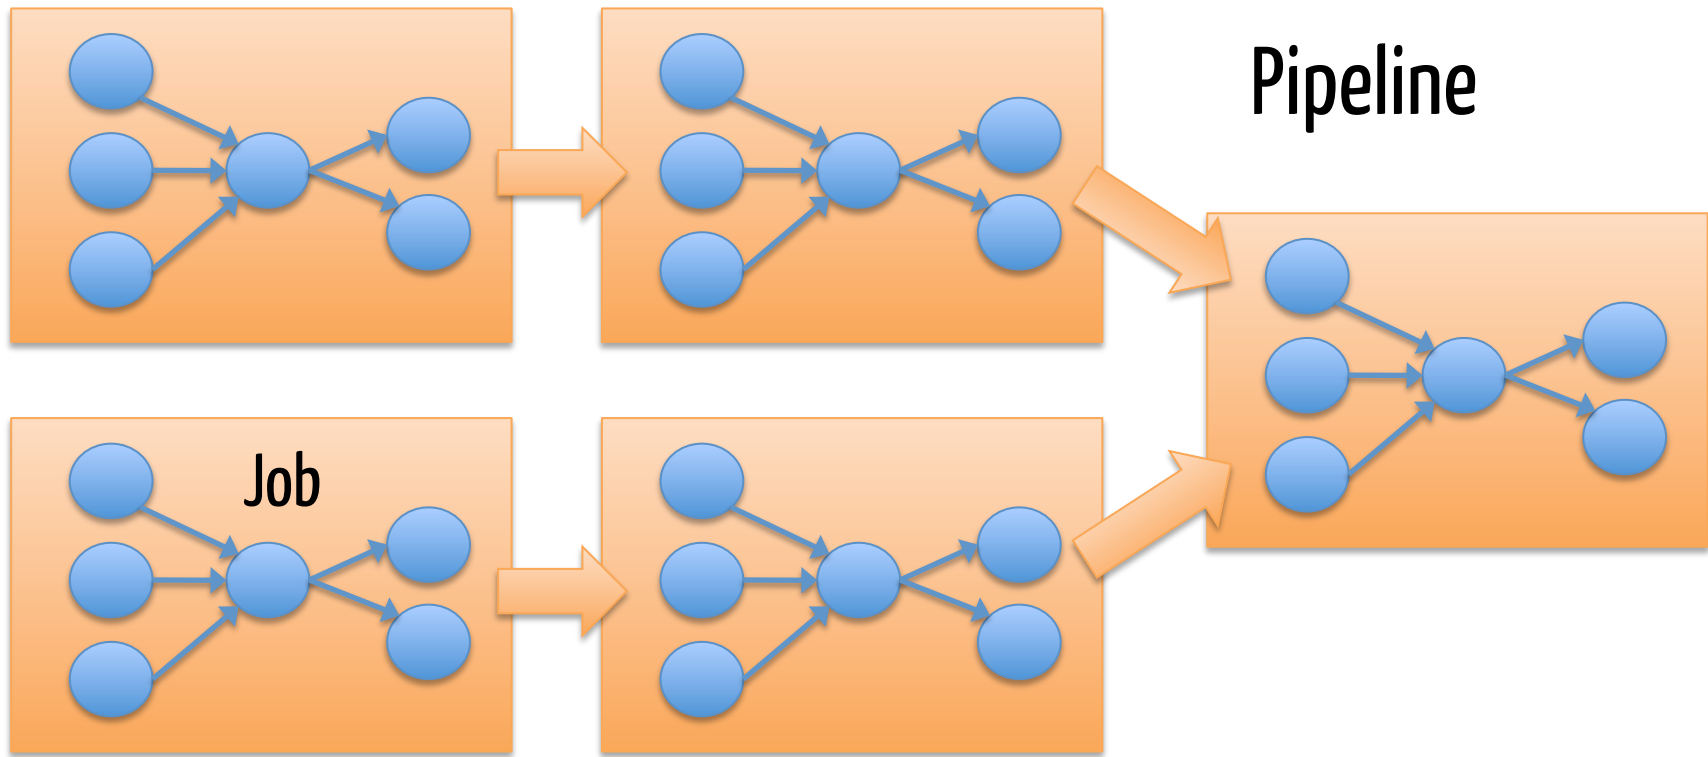


# DRYAD'S DAG WORKFLOW

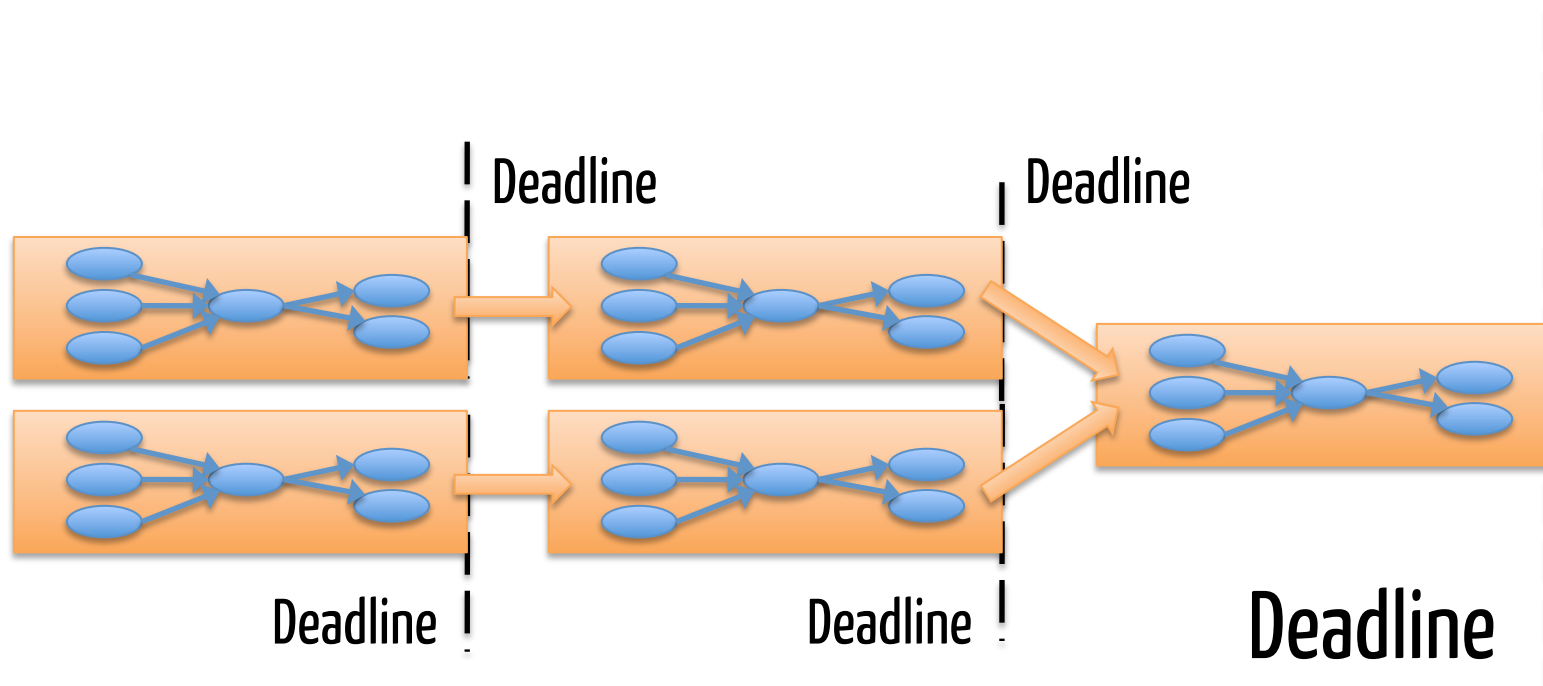
Cosmos Cluster



# DRYAD'S DAG WORKFLOW

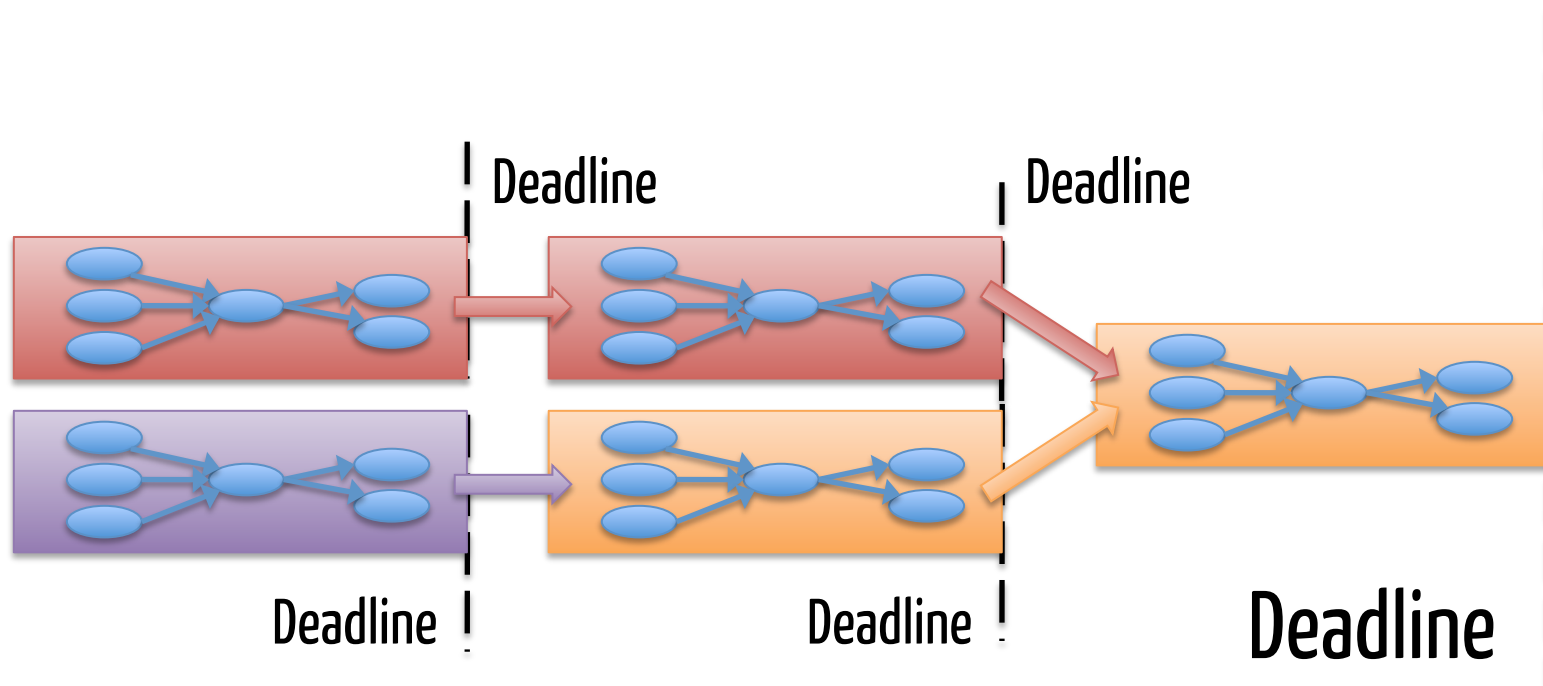


# DRYAD'S DAG WORKFLOW

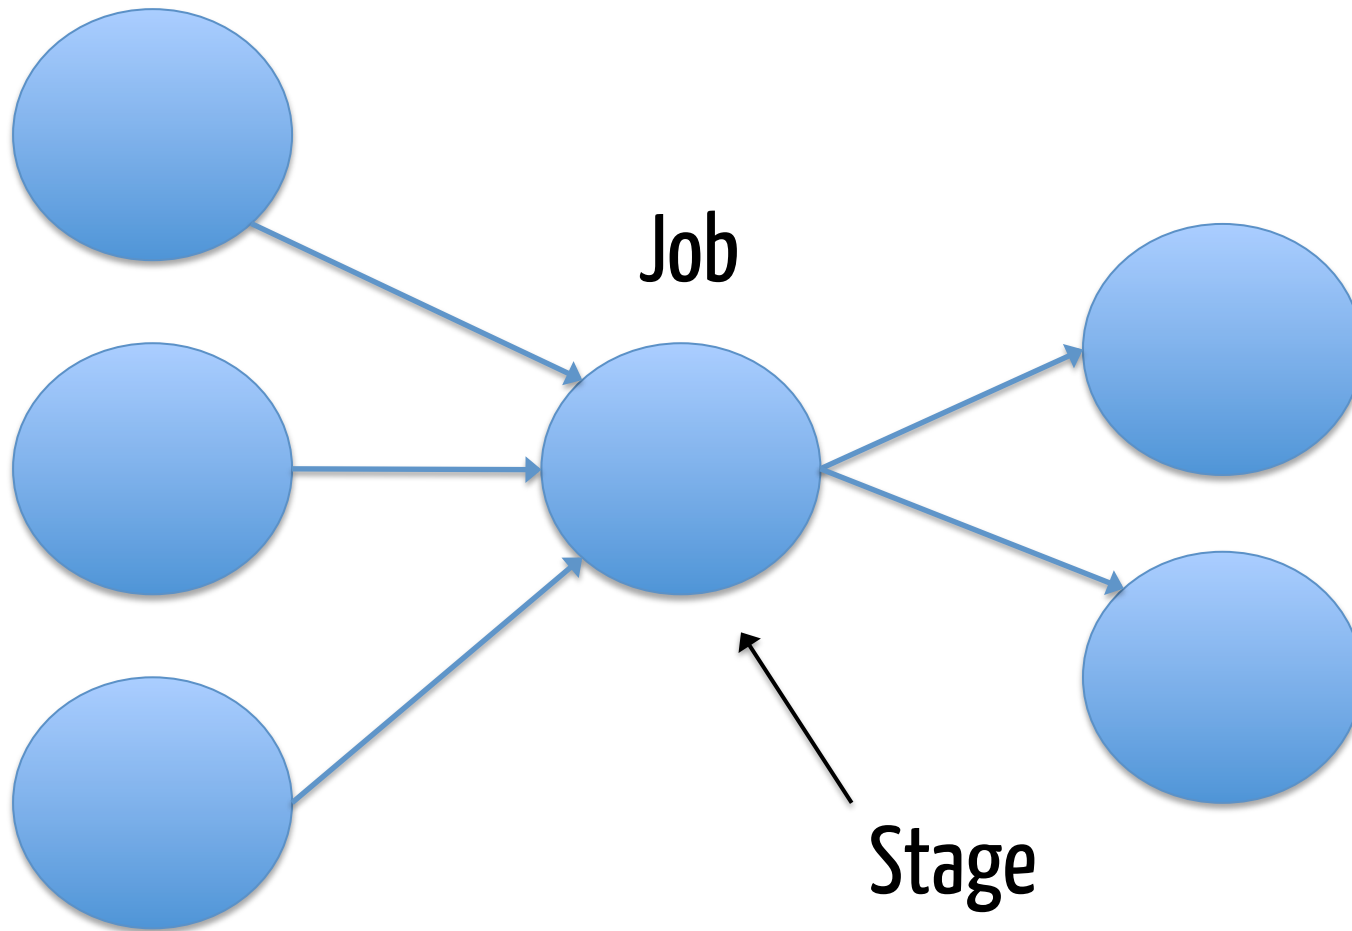


# DRYAD'S DAG WORKFLOW

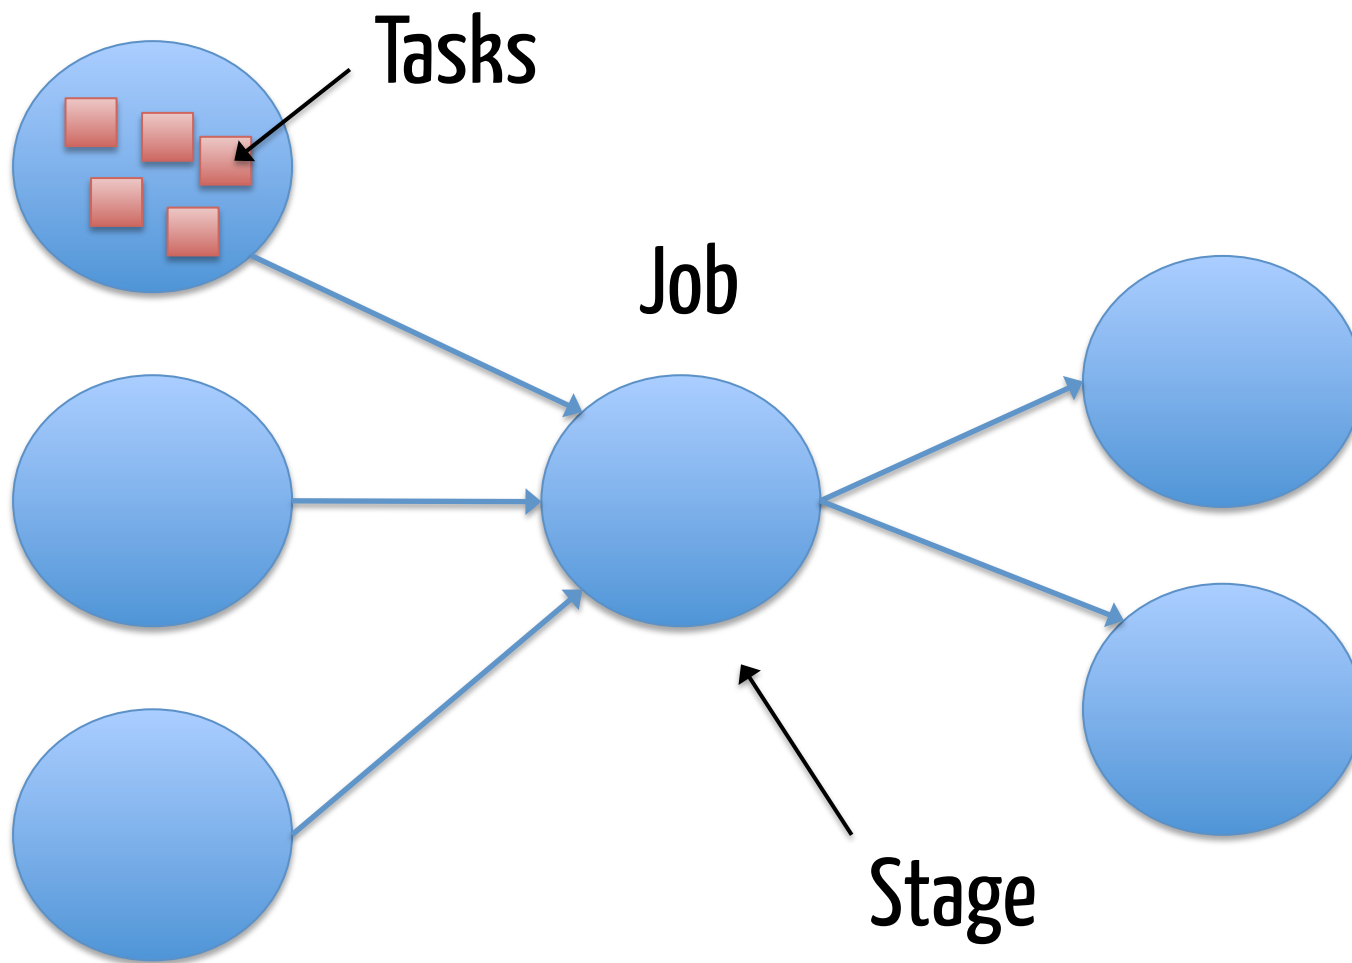




# DRYAD'S DAG WORKFLOW



# DRYAD'S DAG WORKFLOW



# DRYAD'S DAG WORKFLOW



**Priorities?**

**EXPRESSING PERFORMANCE TARGETS**

**✘ Priorities? Not expressive enough**  
**Weights?**




**EXPRESSING PERFORMANCE TARGETS**

**✘ Priorities? Not expressive enough**

**✘ Weights? Difficult for users to set**

**Utility curves?**

## **EXPRESSING PERFORMANCE TARGETS**

-  **Priorities? Not expressive enough**
-  **Weights? Difficult for users to set**
-  **Utility curves? Capture deadline & penalty**

## **EXPRESSING PERFORMANCE TARGETS**



**OUR GOAL**

**Maximize utility**

**OUR GOAL**

**Maximize utility**  
**while minimizing resources**

**OUR GOAL**

**Maximize utility**  
**while minimizing resources**  
**by dynamically adjusting the allocation**

**OUR GOAL**

# Jockey



- **Large clusters**

**Jockey**



- **Large clusters**
- **Many users**

# **Jockey**



- **Large clusters**
- **Many users**
- **Prior execution**

# **Jockey**





f( , ) ->

## JOCKEY – MODEL

f(job state, ) ->

## JOCKEY – MODEL

**f(job state, allocation) ->**

**JOCKEY – MODEL**

**f(job state, allocation) -> remaining run time**

**JOCKEY – MODEL**



# JOCKEY – CONTROL LOOP



# JOCKEY – CONTROL LOOP



# JOCKEY – CONTROL LOOP

**f(job state, allocation) -> remaining run time**

**JOCKEY – MODEL**

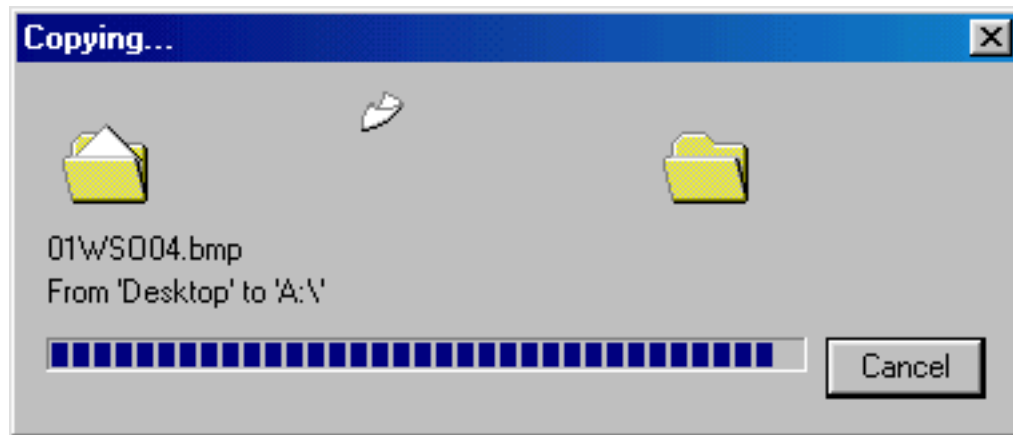


**f**(job state, allocation) -> remaining run time

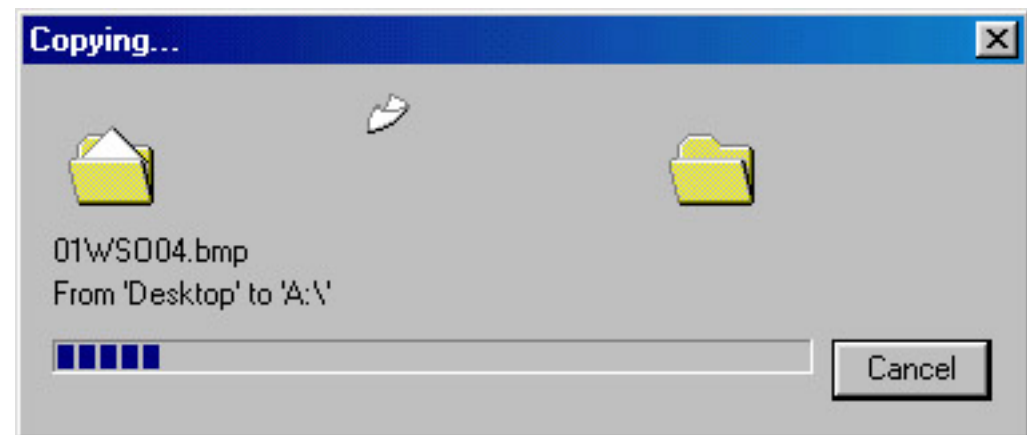
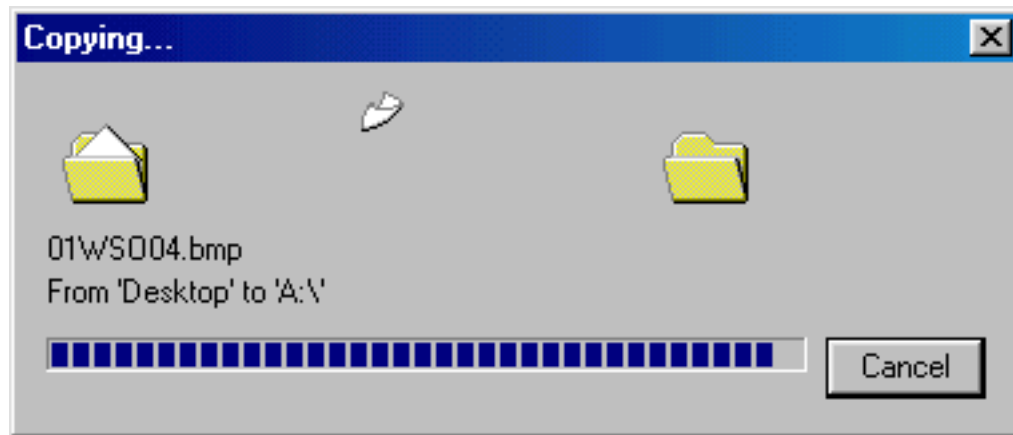


**f**(progress, allocation) -> remaining run time

**JOCKEY – MODEL**



# JOCKEY – PROGRESS INDICATOR



# JOCKEY – PROGRESS INDICATOR

# **JOCKEY – PROGRESS INDICATOR**

total running

**JOCKEY – PROGRESS INDICATOR**

total running  
+  
total queuing

**JOCKEY – PROGRESS INDICATOR**

**stage**

$$\begin{array}{c} \text{total running} \\ + \\ \text{total queuing} \end{array}$$

**JOCKEY – PROGRESS INDICATOR**

Stage 1

total running  
+  
total queuing

+

Stage 2

total running  
+  
total queuing

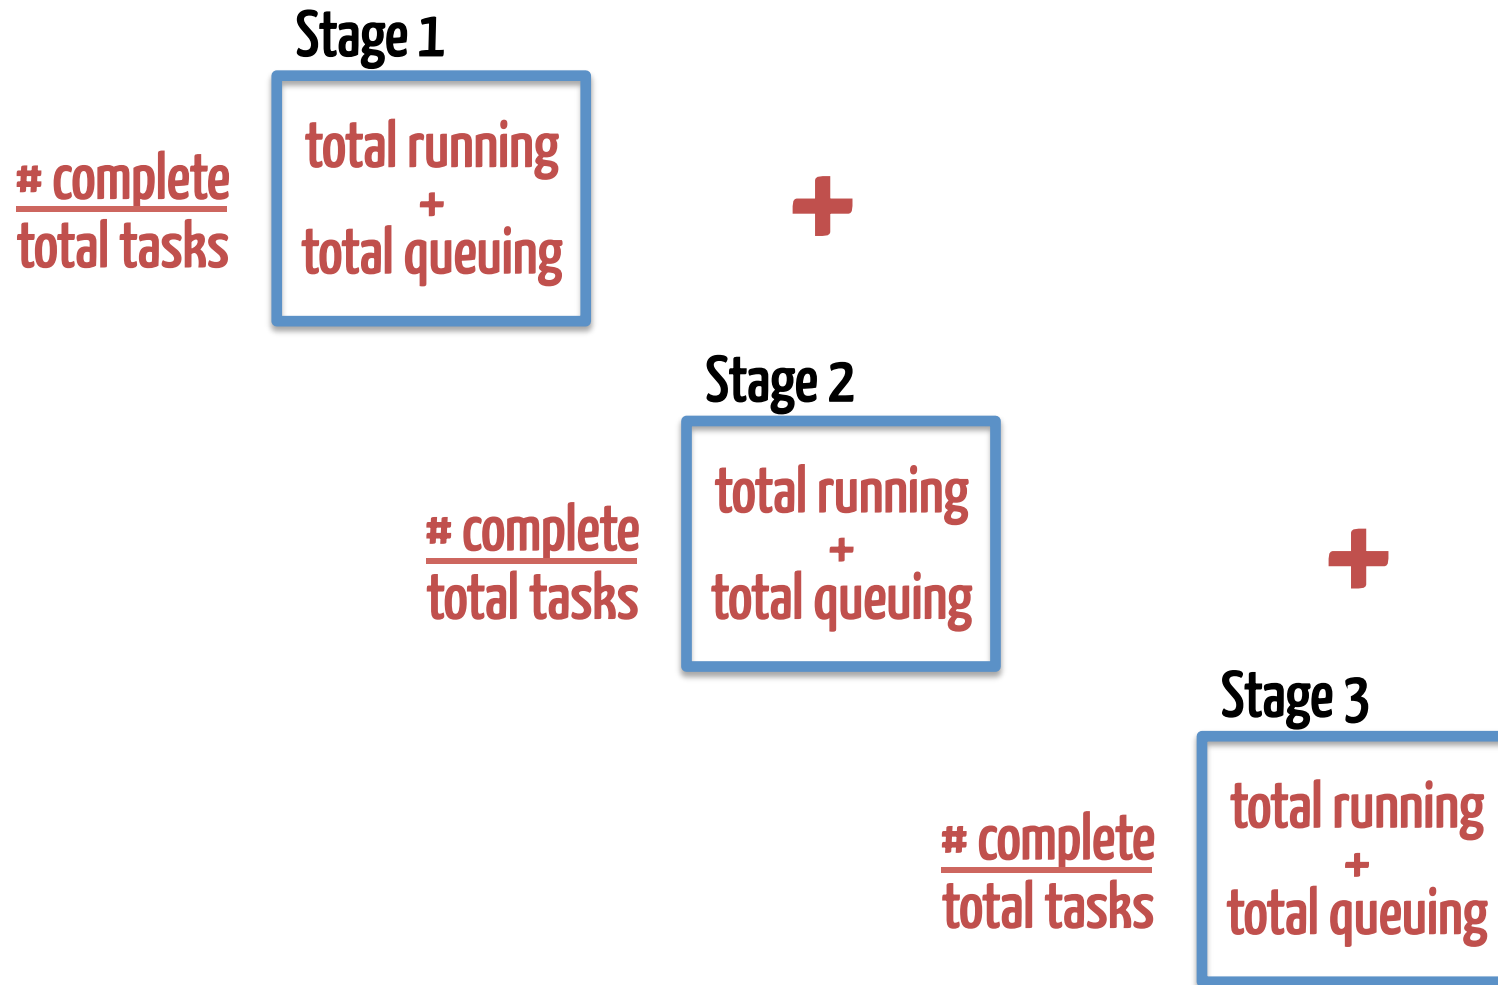
+

Stage 3

total running  
+  
total queuing

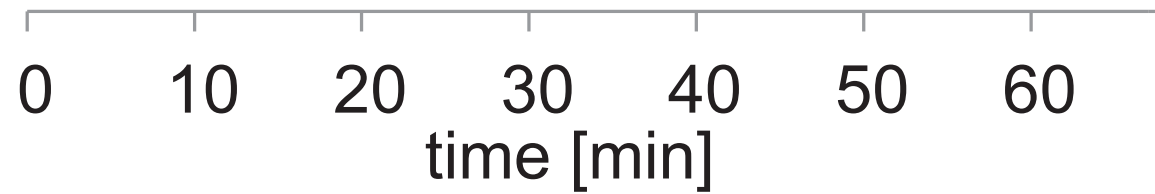
# JOCKEY – PROGRESS INDICATOR



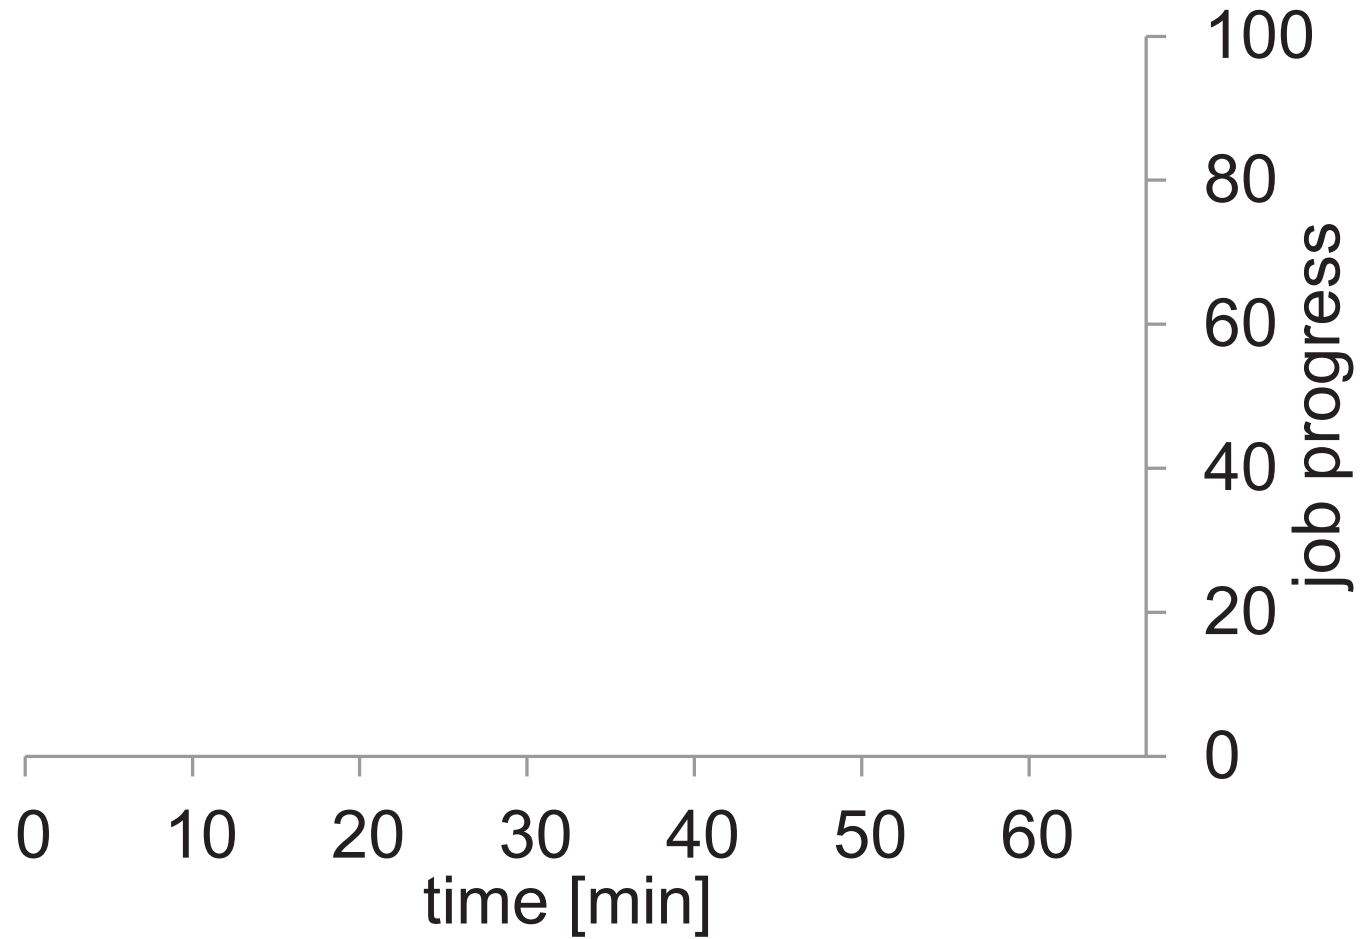


# JOCKEY – PROGRESS INDICATOR

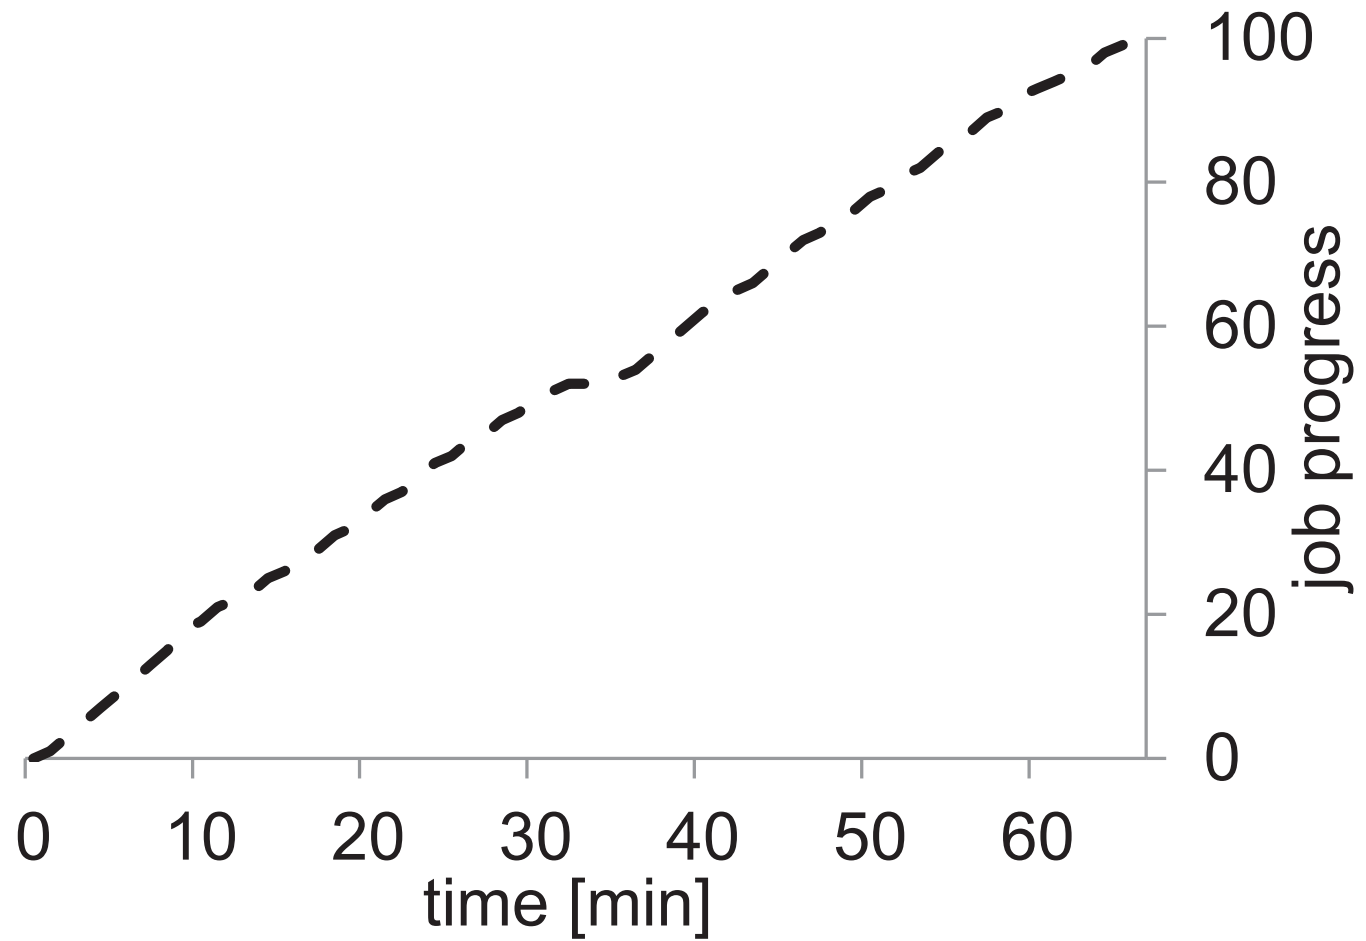
# **JOCKEY – PROGRESS INDICATOR**



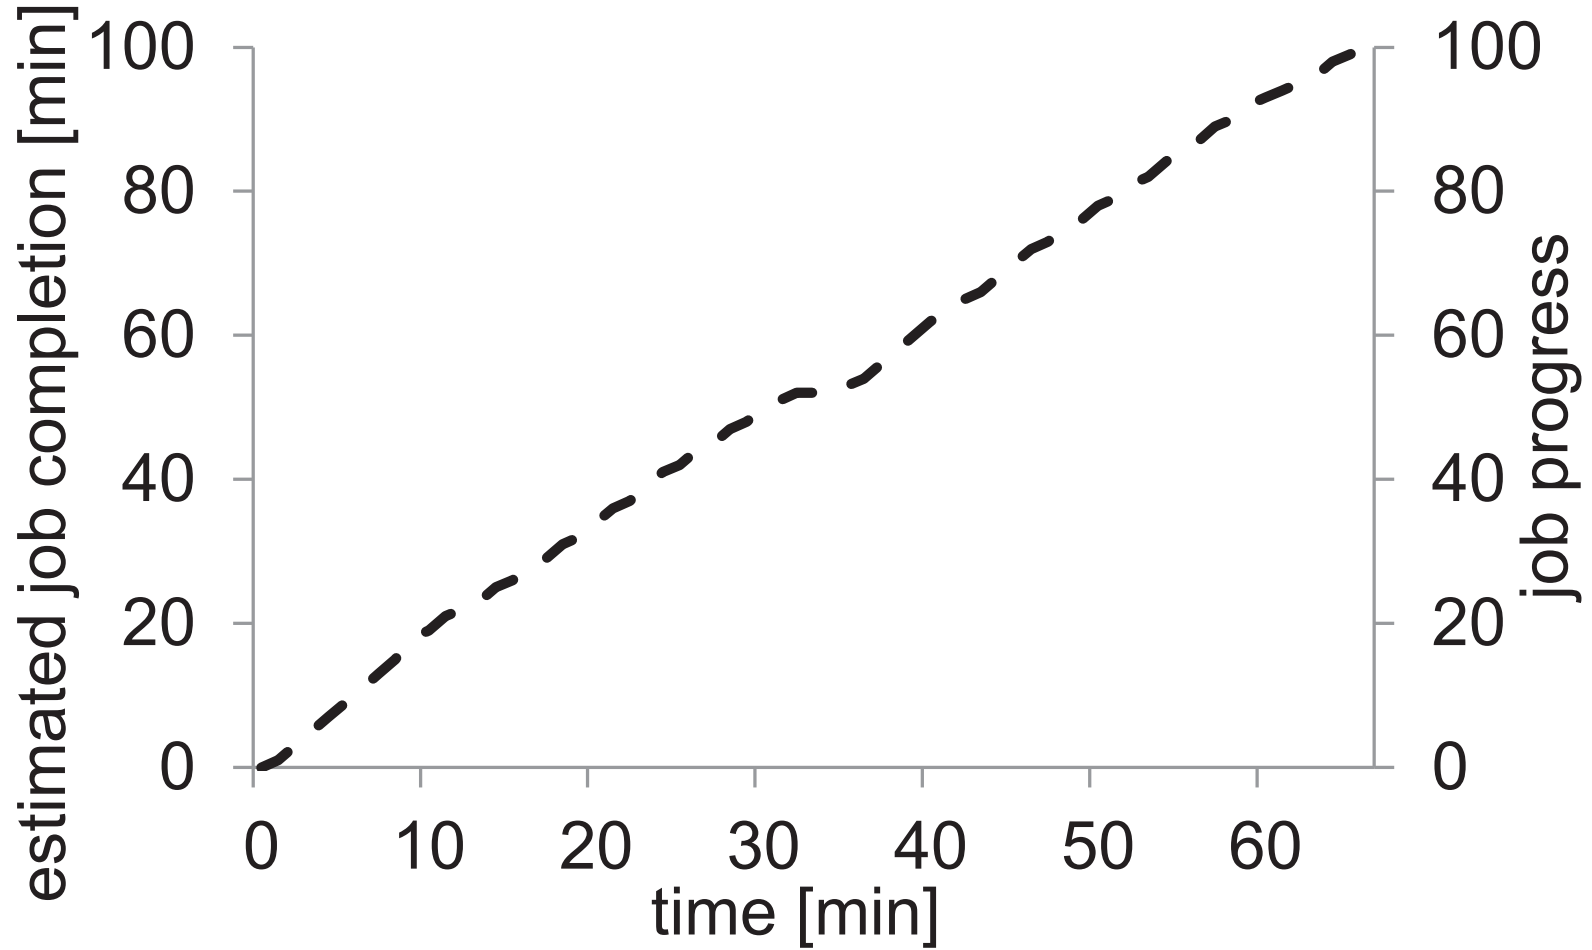
## **JOCKEY – PROGRESS INDICATOR**



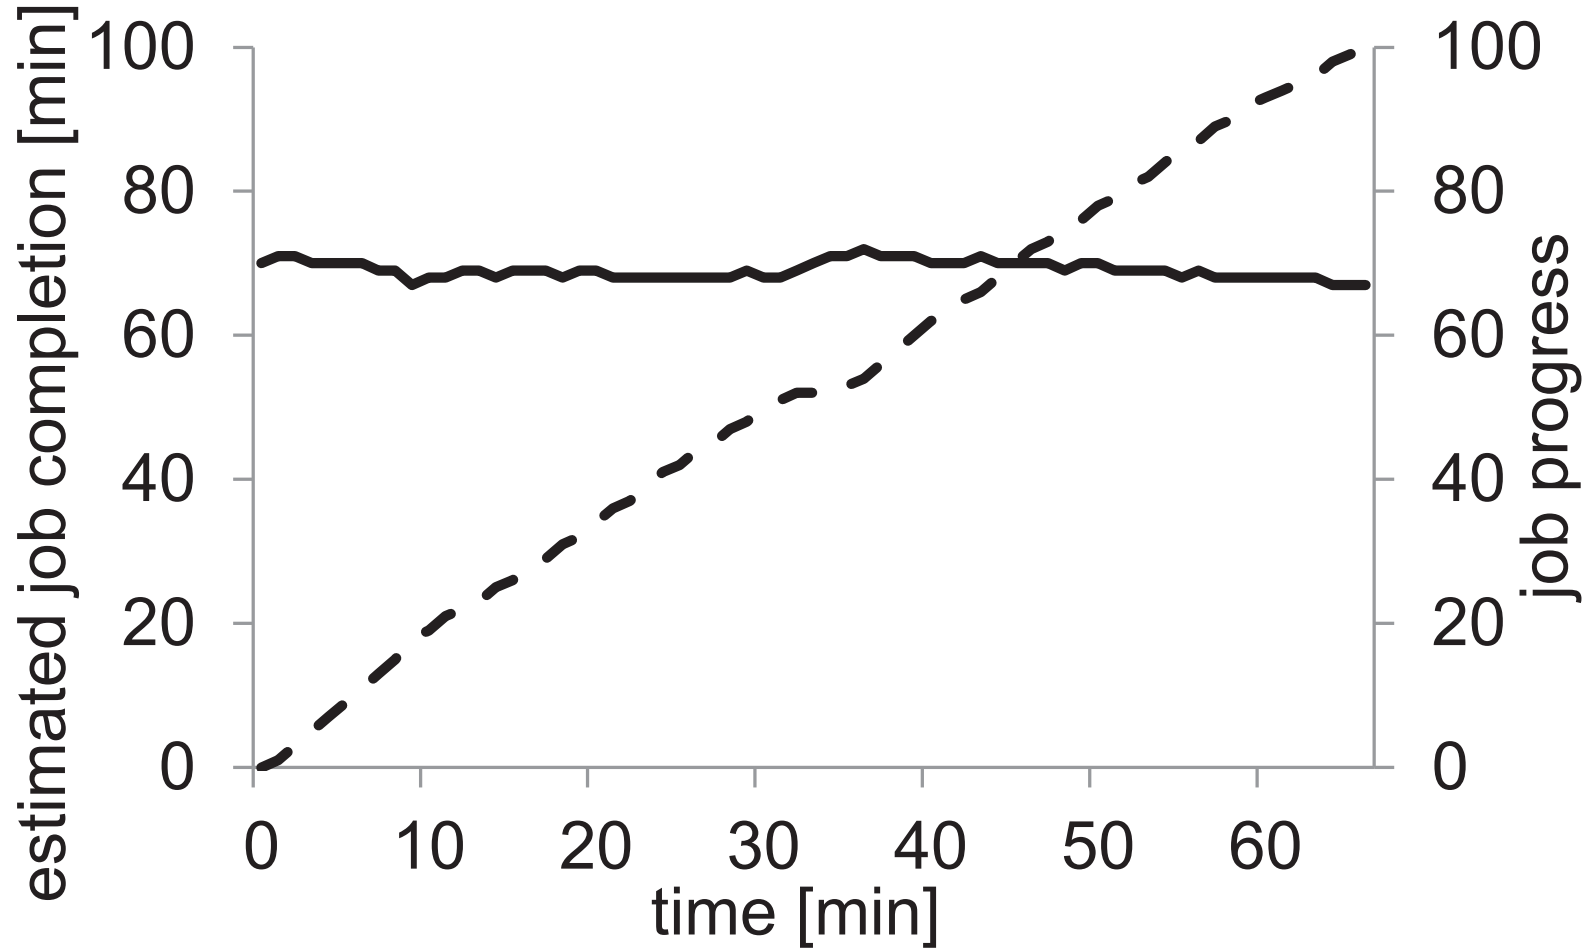
# JOCKEY – PROGRESS INDICATOR



## JOCKEY – PROGRESS INDICATOR



## JOCKEY – PROGRESS INDICATOR



## JOCKEY – PROGRESS INDICATOR

# **JOCKEY – CONTROL LOOP**



# Job model

1% complete			
2% complete			
3% complete			
4% complete			
5% complete			

## JOCKEY – CONTROL LOOP

# Job model

	10 nodes	20 nodes	30 nodes
1% complete			
2% complete			
3% complete			
4% complete			
5% complete			

## JOCKEY – CONTROL LOOP

# Job model

	10 nodes	20 nodes	30 nodes
1% complete	60 minutes	40 minutes	25 minutes
2% complete	59 minutes	39 minutes	24 minutes
3% complete	58 minutes	37 minutes	22 minutes
4% complete	56 minutes	36 minutes	21 minutes
5% complete	54 minutes	34 minutes	20 minutes

## JOCKEY – CONTROL LOOP

# Job model

	10 nodes	20 nodes	30 nodes
1% complete	60 minutes	40 minutes	25 minutes
2% complete	59 minutes	39 minutes	24 minutes
3% complete	58 minutes	37 minutes	22 minutes
4% complete	56 minutes	36 minutes	21 minutes
5% complete	54 minutes	34 minutes	20 minutes

**Completion:  
1%**

**Deadline:  
50 min.**

## JOCKEY – CONTROL LOOP

# Job model

	10 nodes	20 nodes	30 nodes
1% complete	60 minutes	40 minutes	25 minutes
2% complete	59 minutes	39 minutes	24 minutes
3% complete	58 minutes	37 minutes	22 minutes
4% complete	56 minutes	36 minutes	21 minutes
5% complete	54 minutes	34 minutes	20 minutes

**Completion:**

**1%**

**Deadline:**

**50 min.**

## JOCKEY – CONTROL LOOP

# Job model

	10 nodes	20 nodes	30 nodes
1% complete	60 minutes	40 minutes	25 minutes
2% complete	59 minutes	39 minutes	24 minutes
3% complete	58 minutes	37 minutes	22 minutes
4% complete	56 minutes	36 minutes	21 minutes
5% complete	54 minutes	34 minutes	20 minutes

**Completion:**

**3%**

**Deadline:**

**40 min.**

## JOCKEY – CONTROL LOOP

# Job model

	10 nodes	20 nodes	30 nodes
1% complete	60 minutes	40 minutes	25 minutes
2% complete	59 minutes	39 minutes	24 minutes
3% complete	58 minutes	37 minutes	22 minutes
4% complete	56 minutes	36 minutes	21 minutes
5% complete	54 minutes	34 minutes	20 minutes

**Completion:**

**5%**

**Deadline:**

**30 min.**

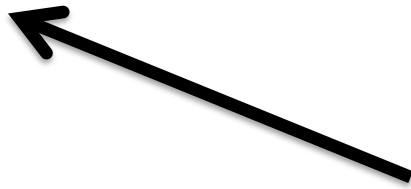
## JOCKEY – CONTROL LOOP

**f(progress, allocation) -> remaining run time**

**JOCKEY – MODEL**



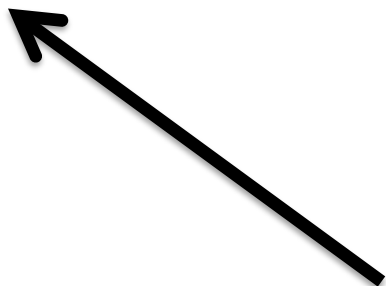
**f(progress, allocation) -> remaining run time**



**analytic model?**

**JOCKEY – MODEL**

**f(progress, allocation) -> remaining run time**

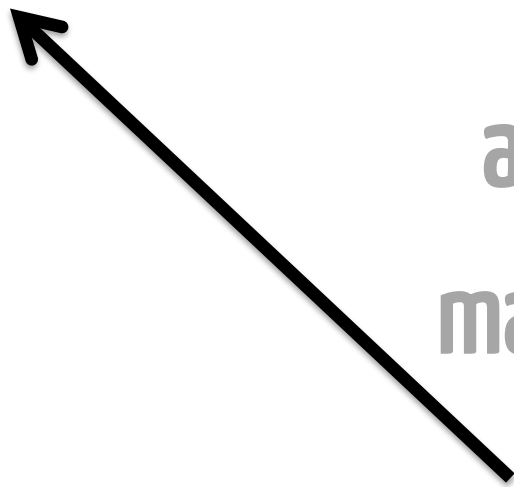


analytic model

**machine learning?**

**JOCKEY – MODEL**

**f(progress, allocation) -> remaining run time**



analytic model

machine learning

**simulator**

**JOCKEY – MODEL**

**Problem**

**Solution**

**JOCKEY**

Problem	Solution
Pipeline complexity	

**JOCKEY**

<b>Problem</b>	<b>Solution</b>
<b>Pipeline complexity</b>	<b>Use a simulator</b>

**JOCKEY**

<b>Problem</b>	<b>Solution</b>
<b>Pipeline complexity</b>	<b>Use a simulator</b>
<b>Noisy environment</b>	

**JOCKEY**

<b>Problem</b>	<b>Solution</b>
<b>Pipeline complexity</b>	<b>Use a simulator</b>
<b>Noisy environment</b>	<b>Dynamic control</b>

**JOCKEY**



# **Jockey in Action**

- **Real job**

# **Jockey in Action**

- **Real job**
- **Production cluster**

# **Jockey in Action**

- **Real job**
- **Production cluster**
- **CPU load: ~80%**

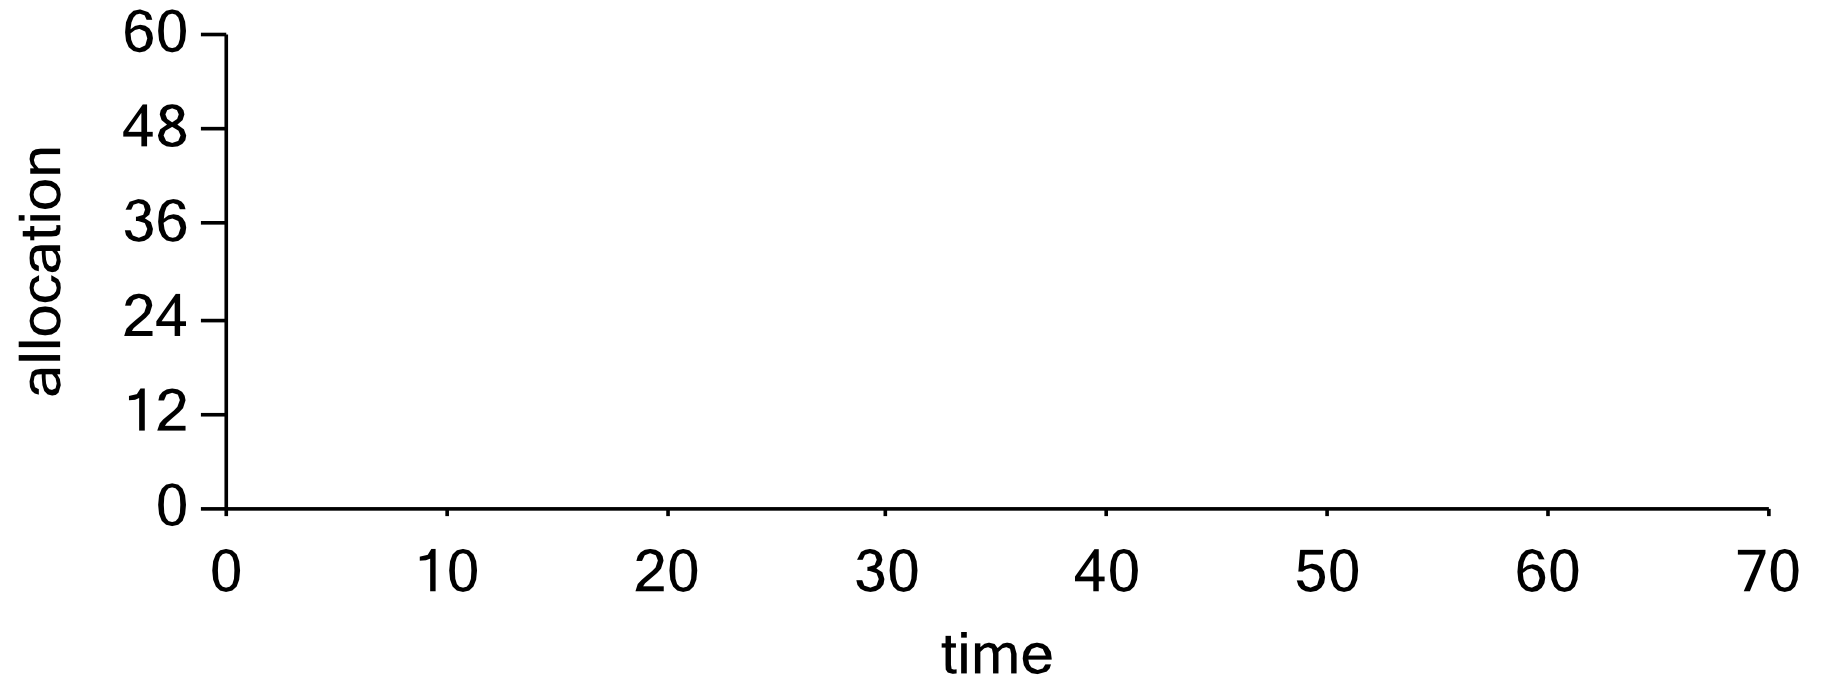
# **Jockey in Action**

# Jockey in Action

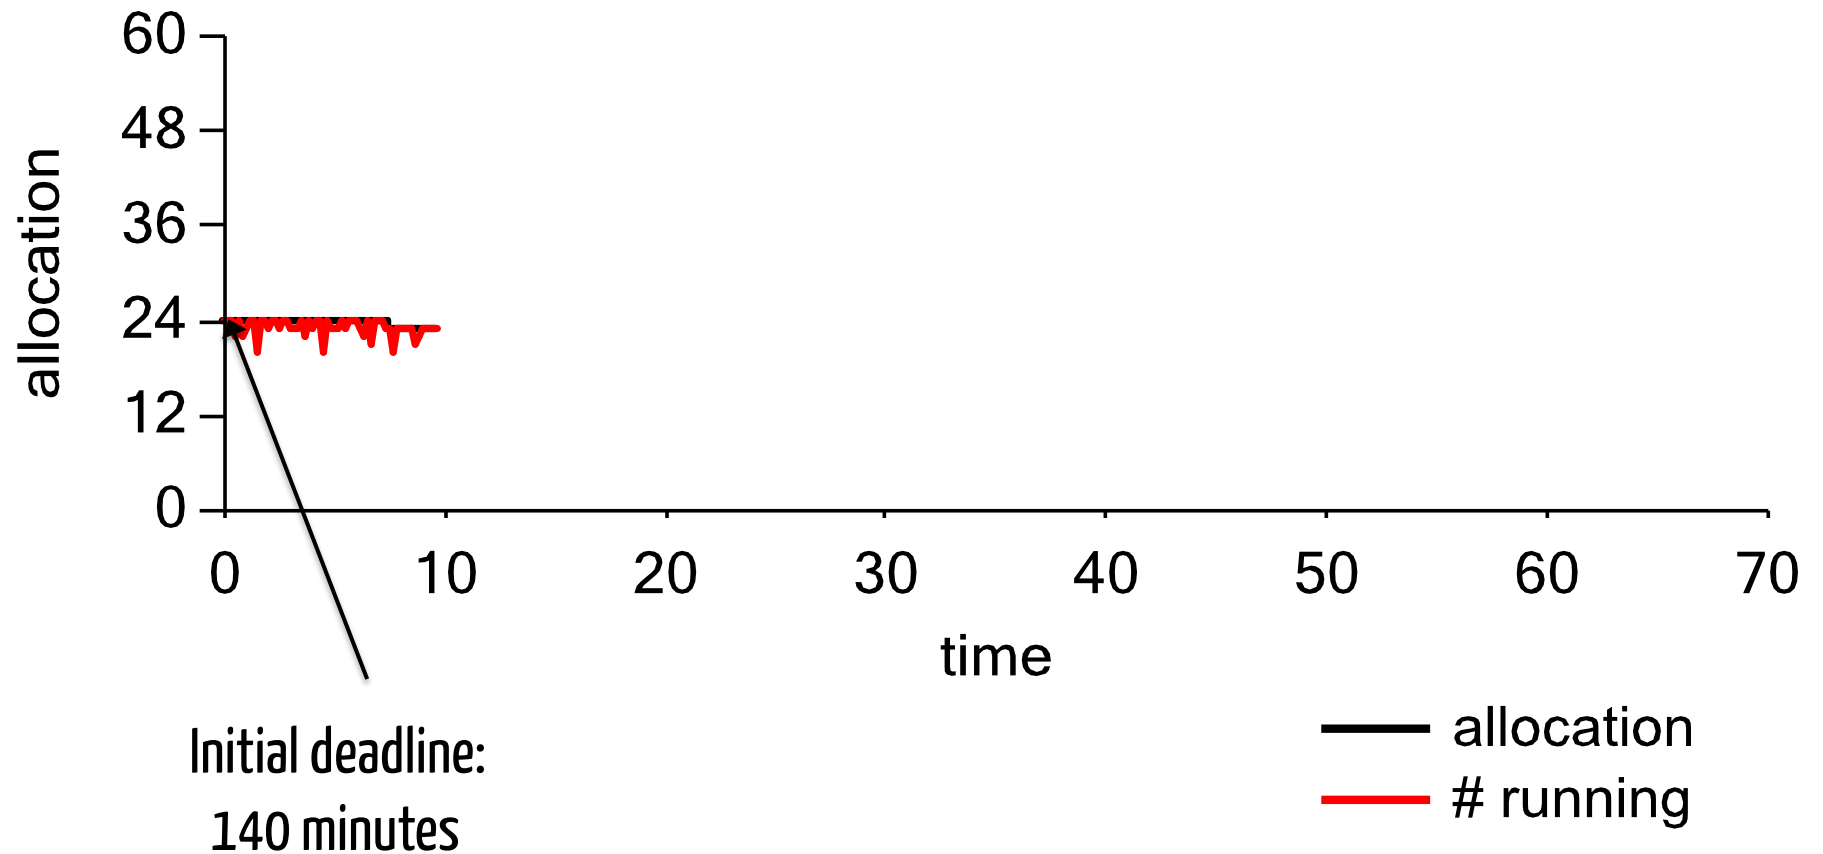
# Jockey in Action



# Jockey in Action

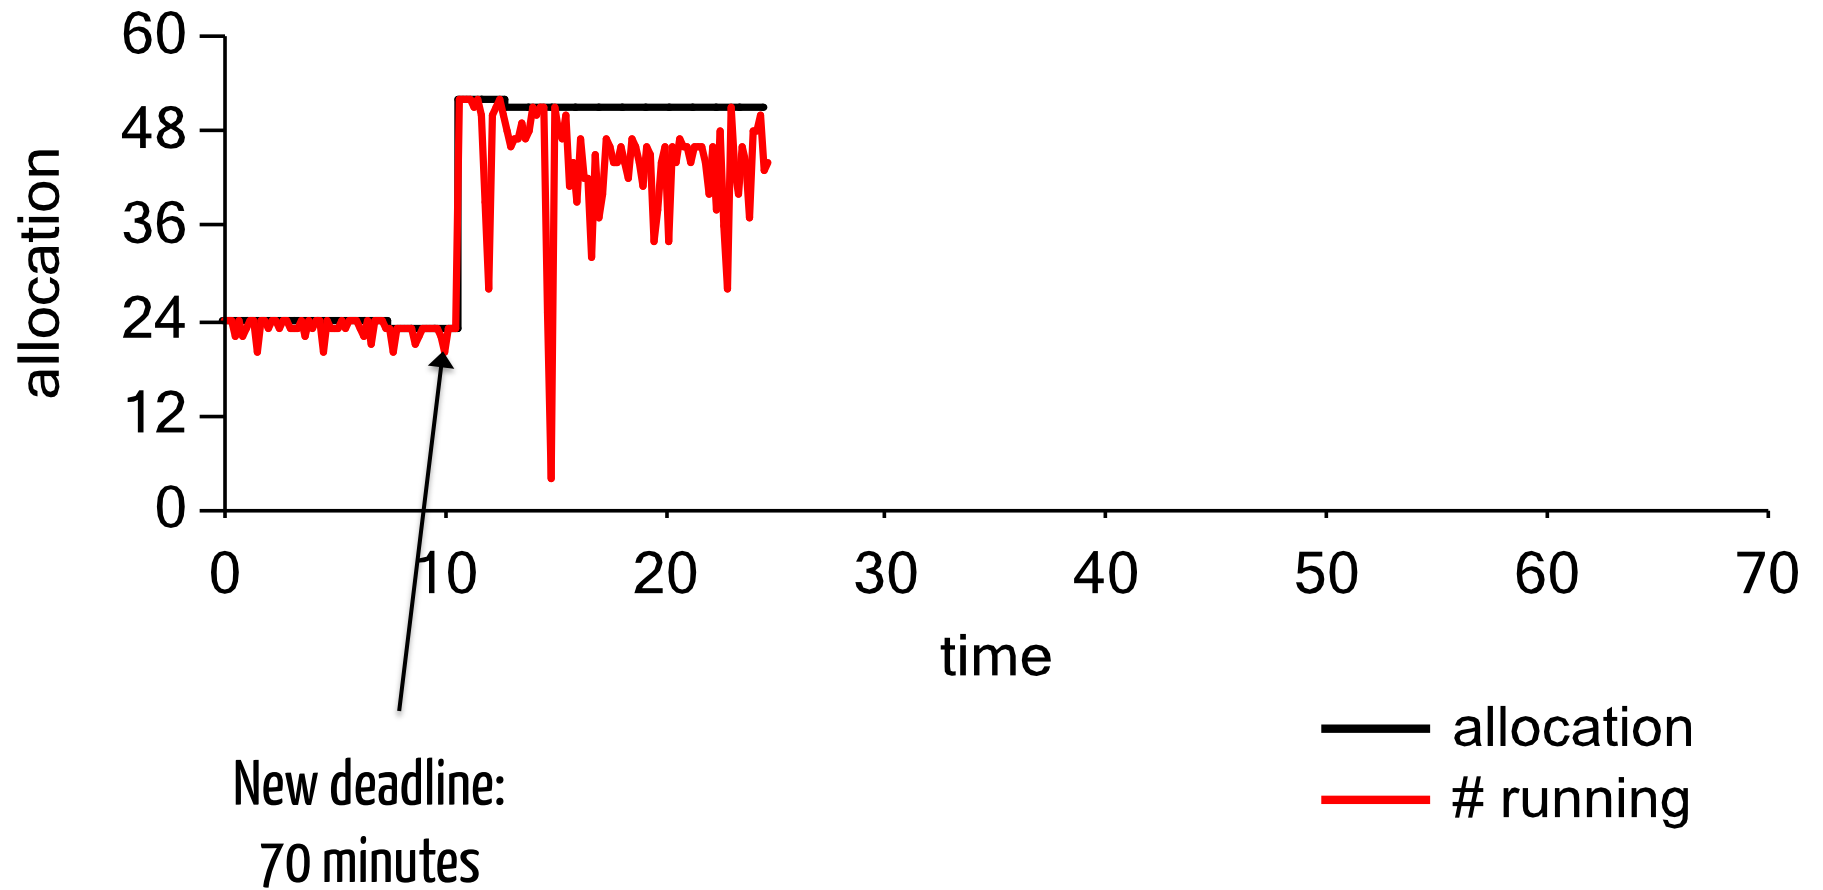


# Jockey in Action

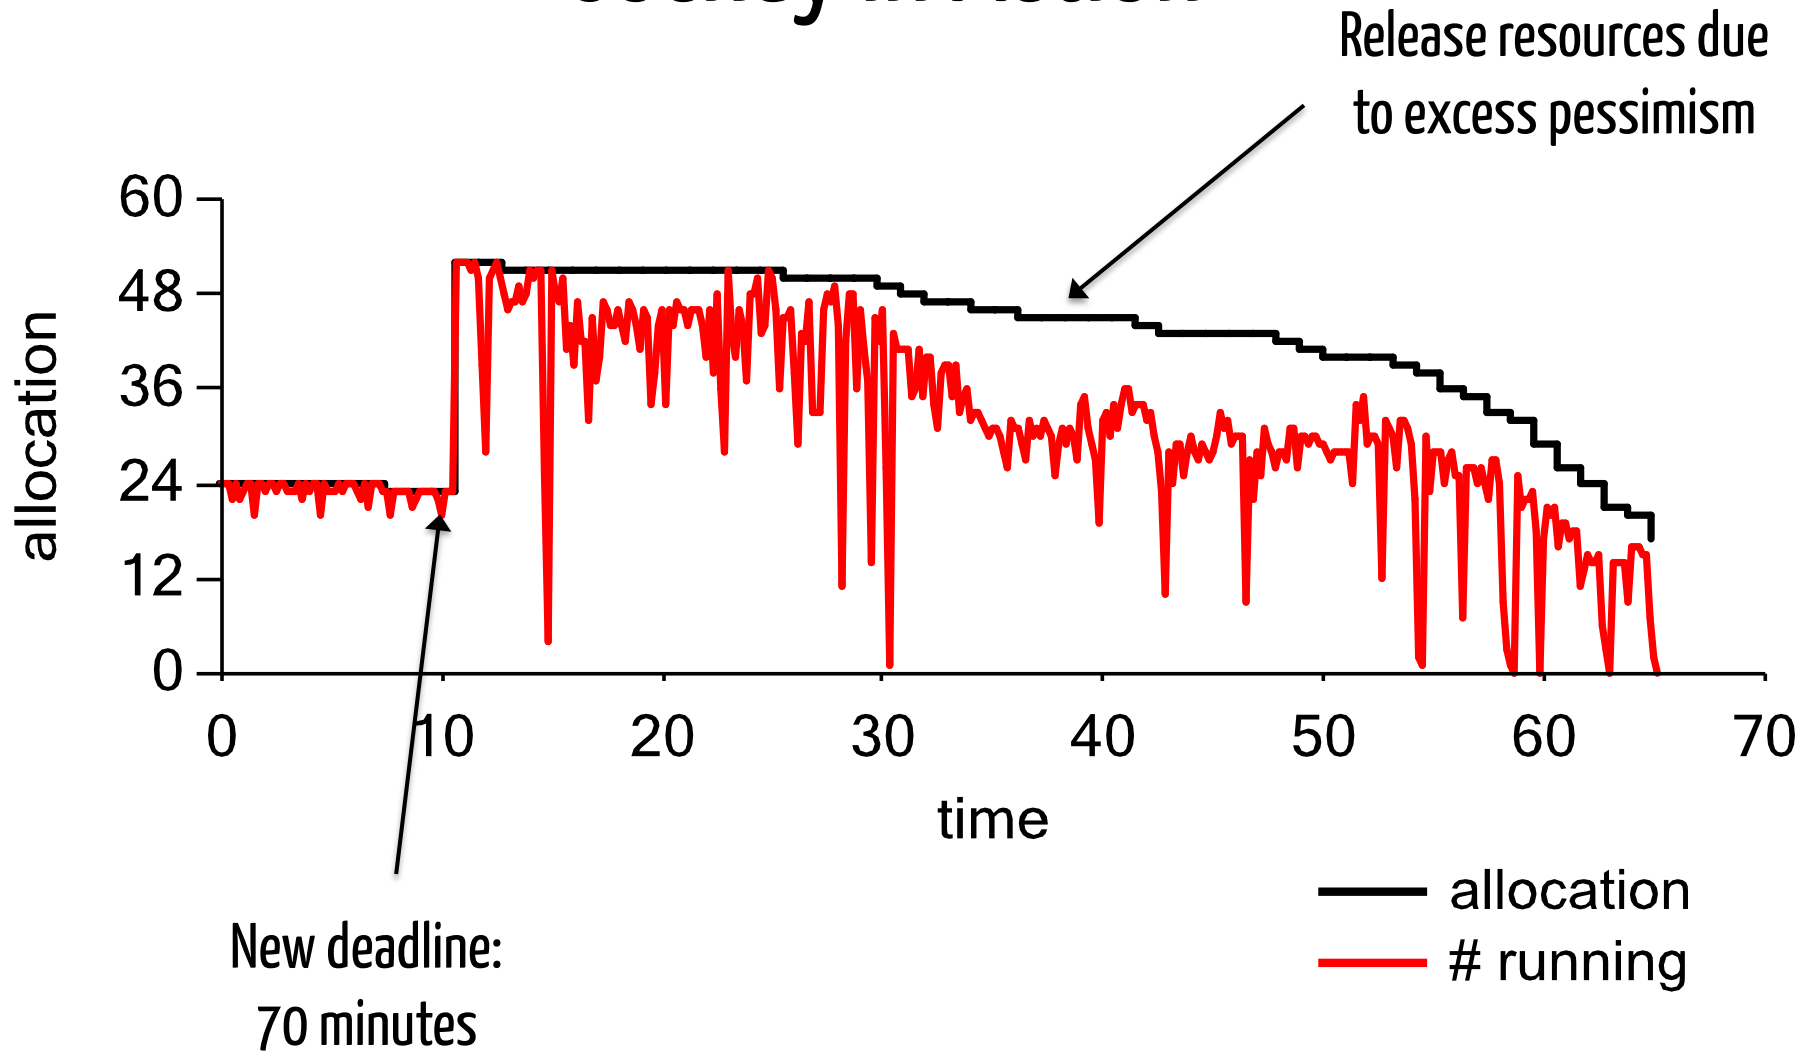




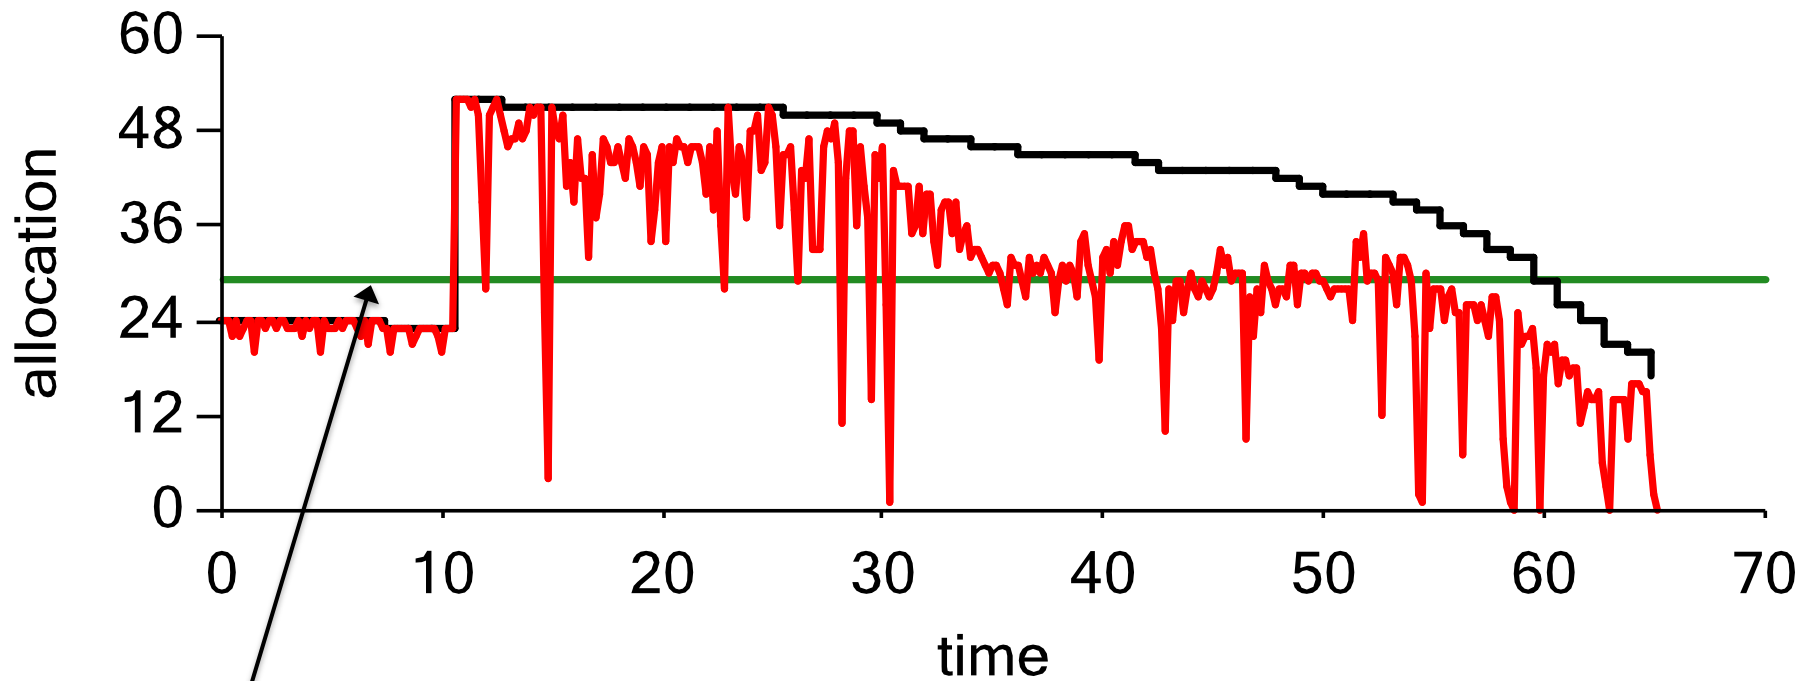
# Jockey in Action



# Jockey in Action



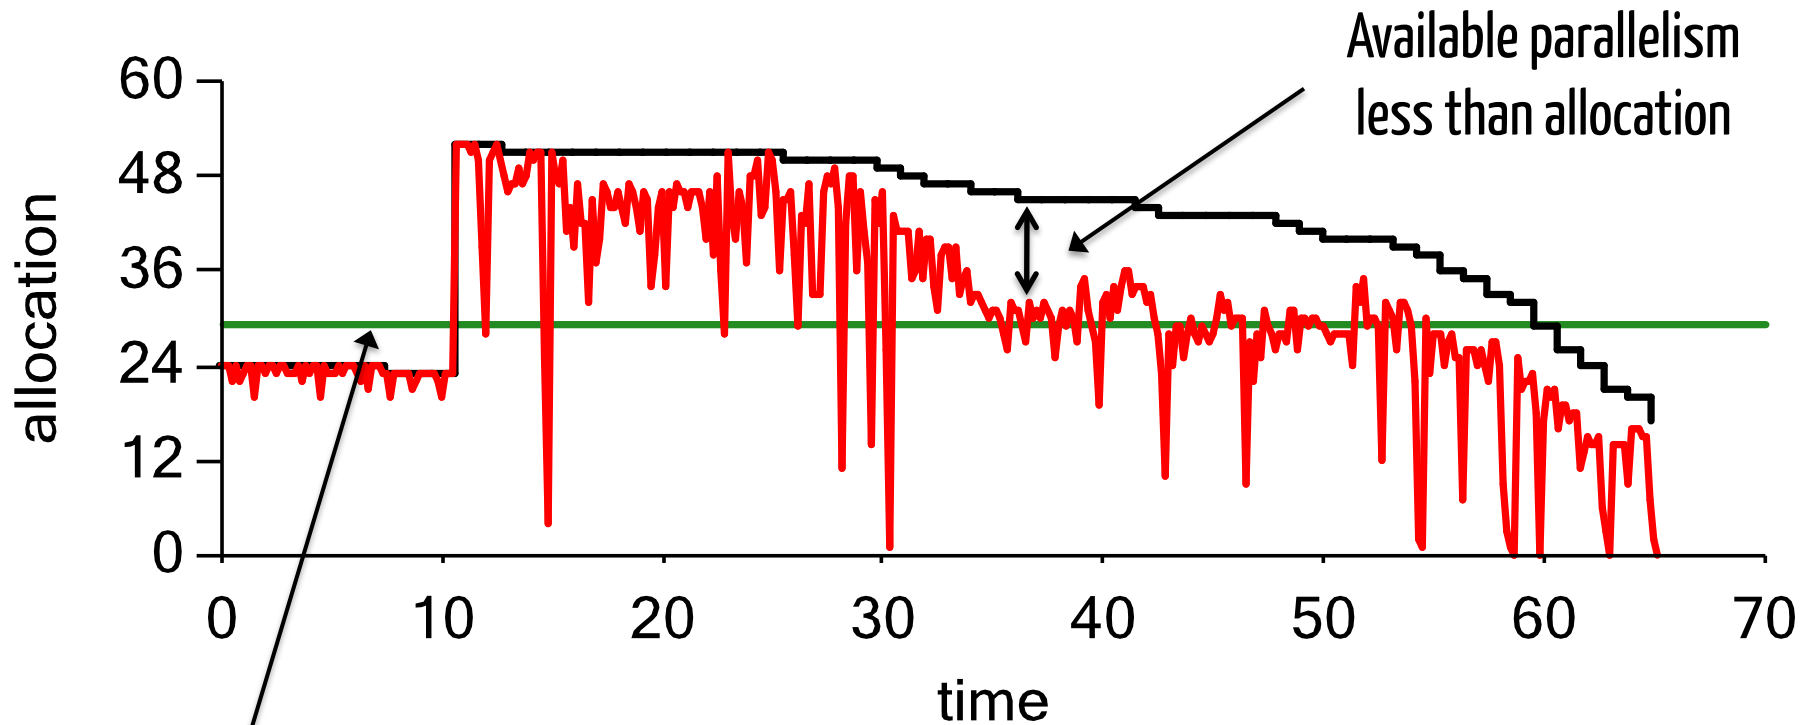
# Jockey in Action



“Oracle” allocation:  
Total allocation-hours  
Deadline

— allocation  
— # running  
— oracle tokens

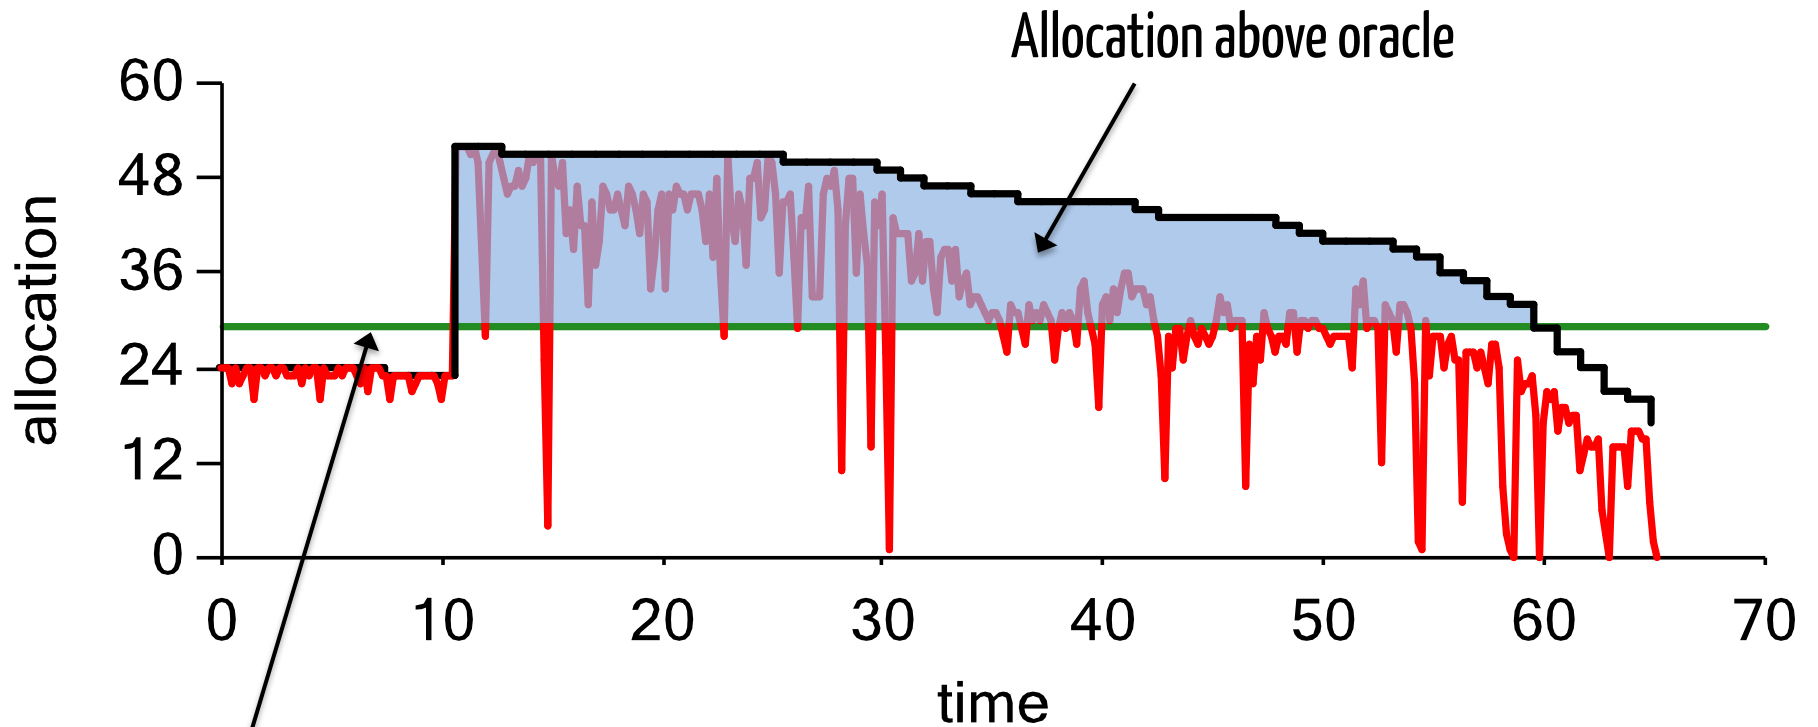
# Jockey in Action



“Oracle” allocation:  
 $\frac{\text{Total allocation-hours}}{\text{Deadline}}$

- allocation
- # running
- oracle tokens

# Jockey in Action



“Oracle” allocation:  
 $\frac{\text{Total allocation-hours}}{\text{Deadline}}$

- allocation
- # running
- oracle tokens

# Evaluation



- **Production cluster**

# Evaluation



- **Production cluster**
- **21 jobs**

# **Evaluation**





- **Production cluster**
- **21 jobs**
- **SLO met?**

# Evaluation



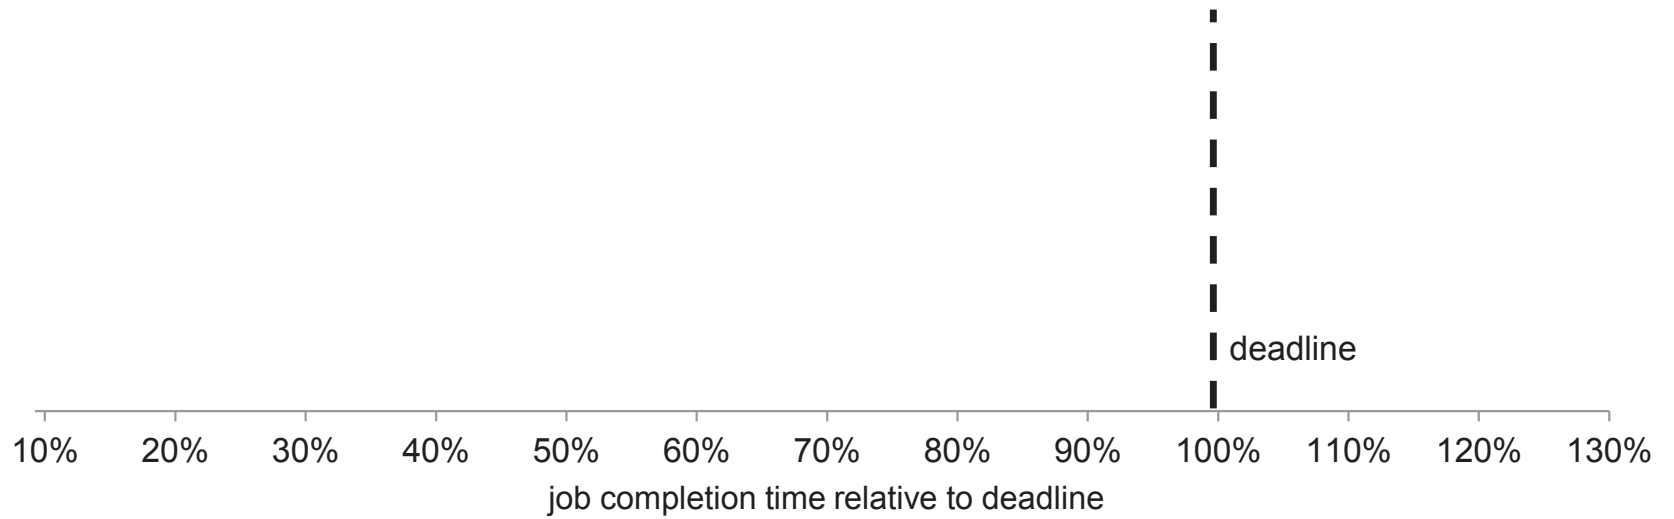
- **Production cluster**
- **21 jobs**
- **SLO met?**
- **Cluster impact?**

# Evaluation



# Evaluation

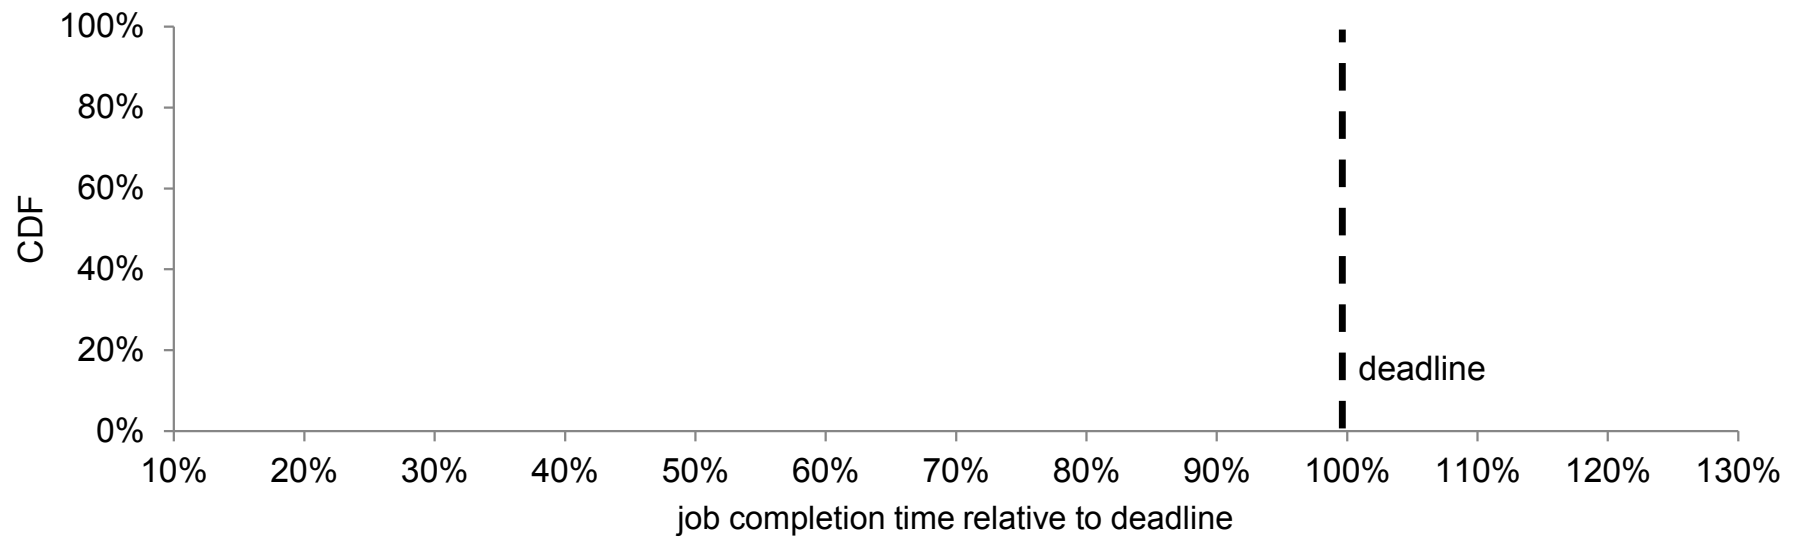
# Evaluation



Jobs which met the SLO



# Evaluation

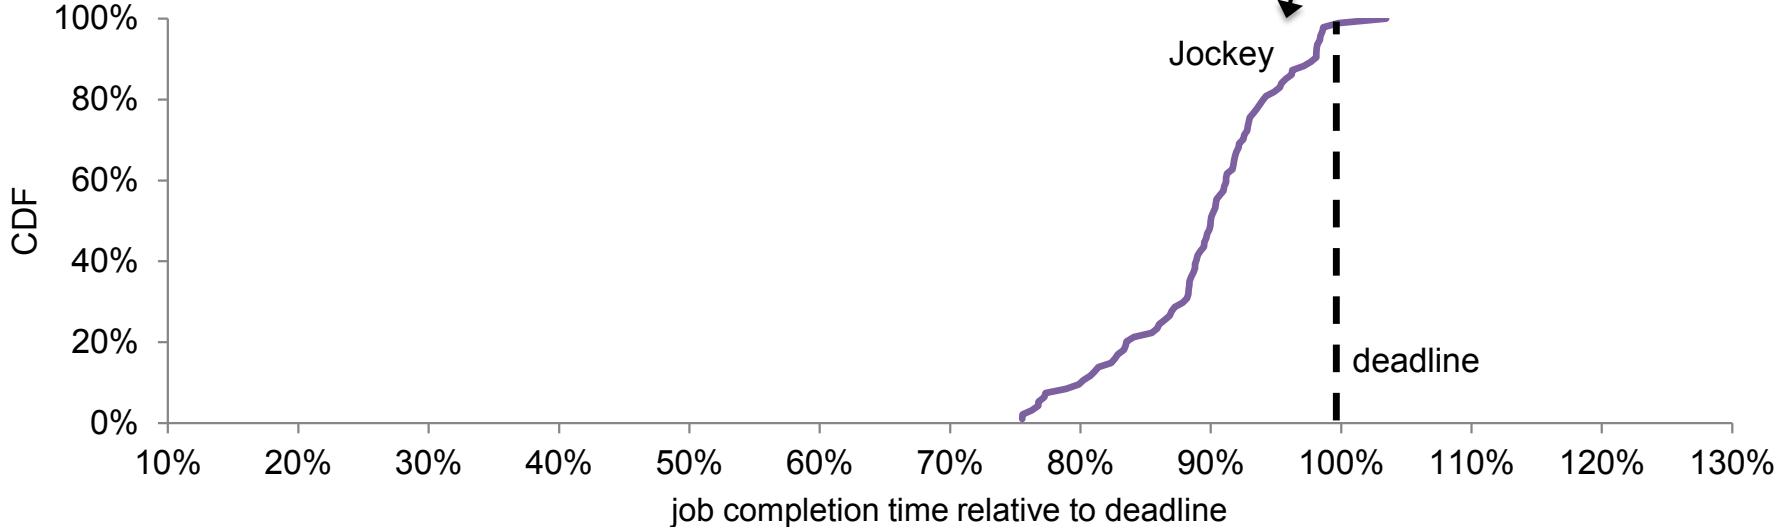


Jobs which met the SLO



# Evaluation

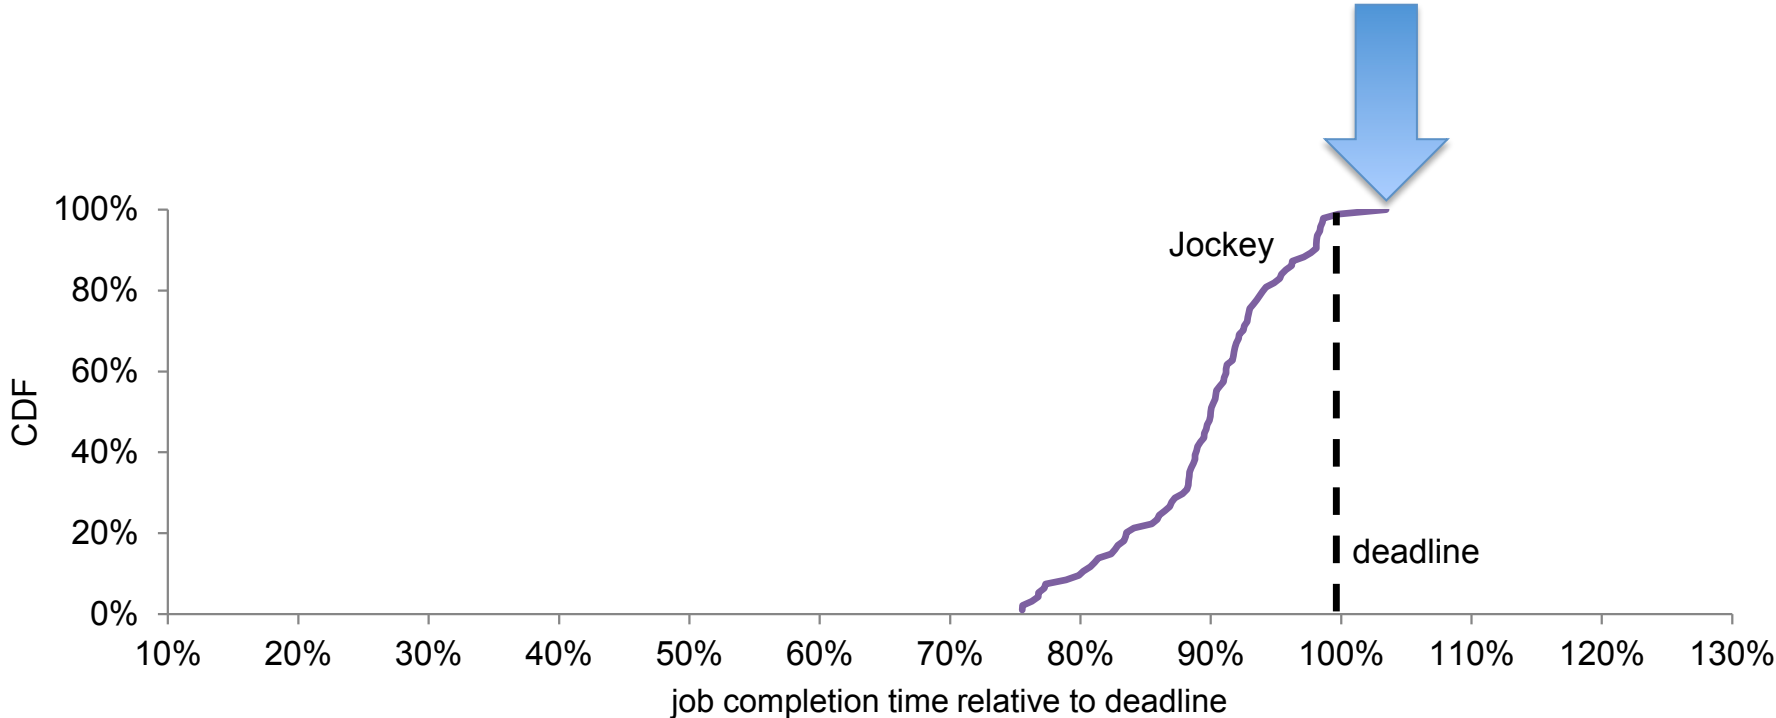
Missed 1 of 94 deadlines



Jobs which met the SLO



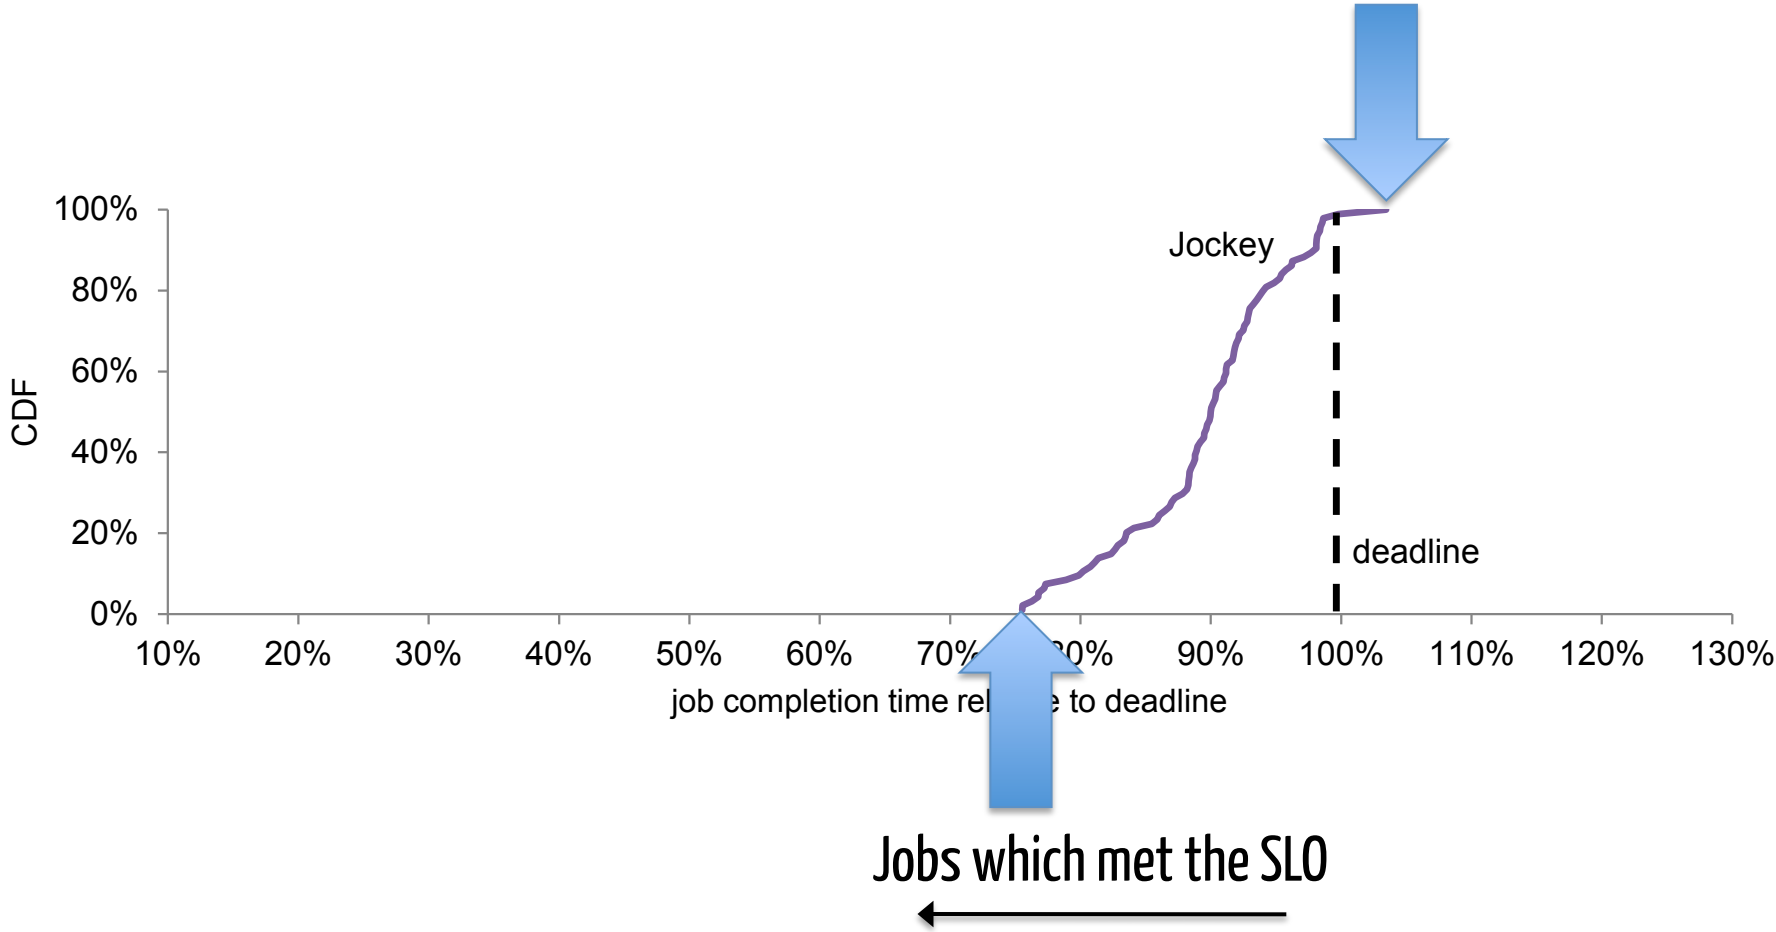
# Evaluation



Jobs which met the SLO

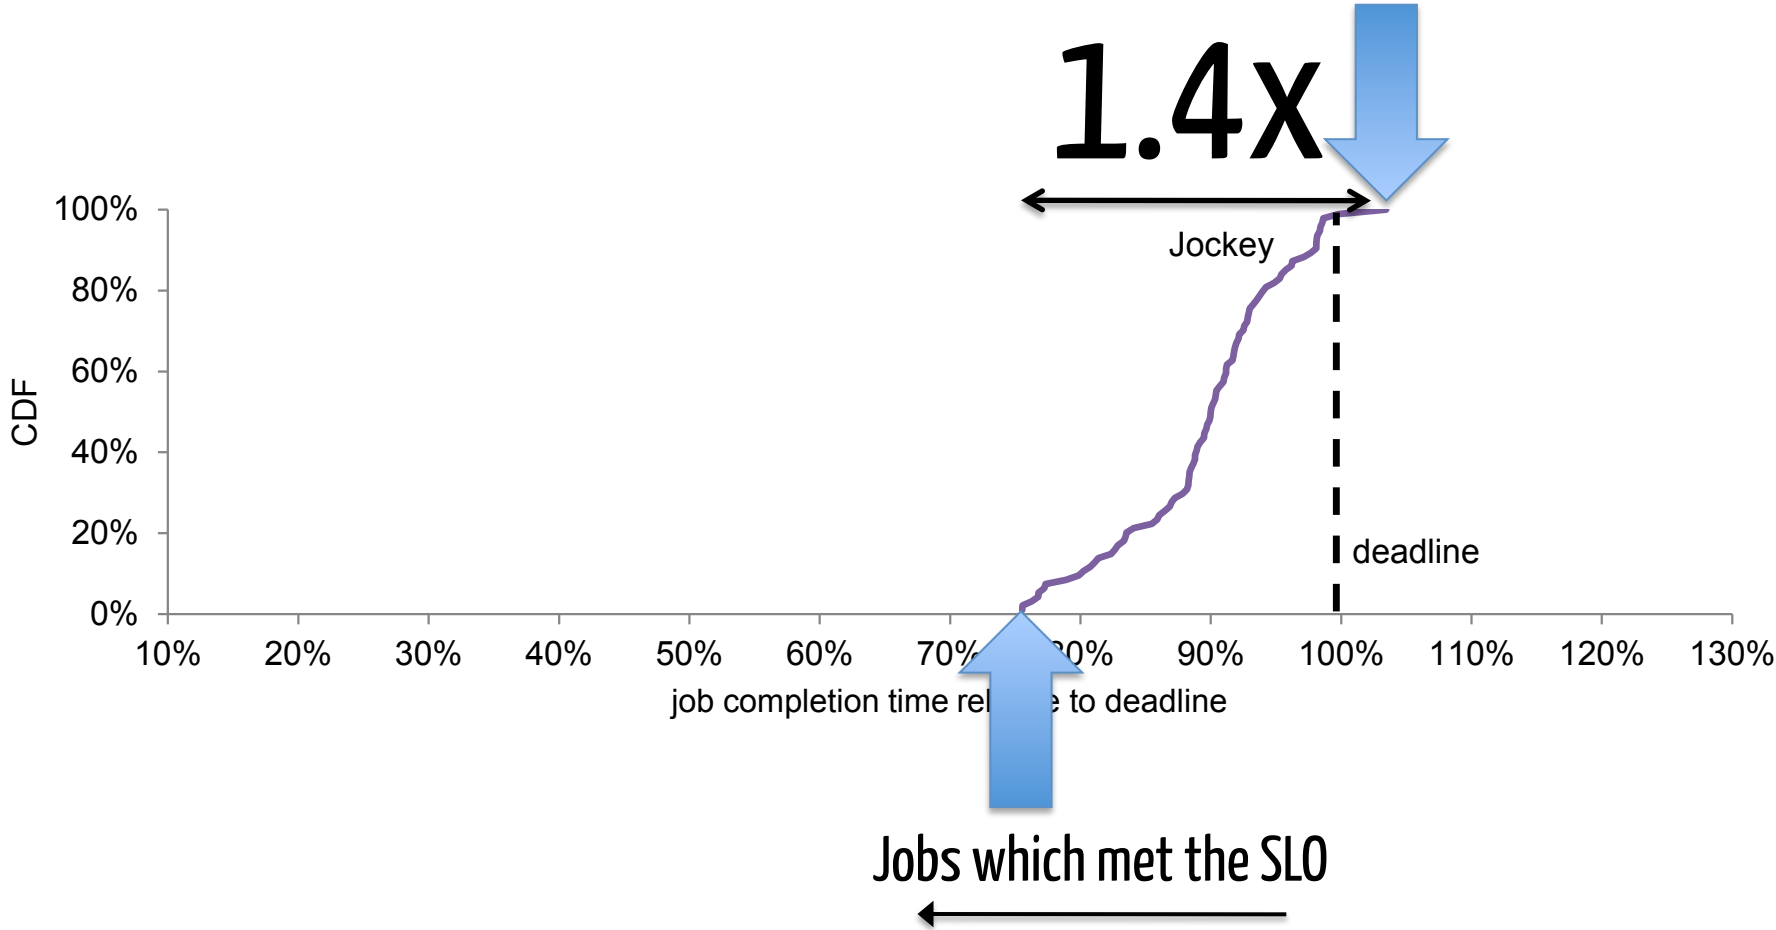


# Evaluation

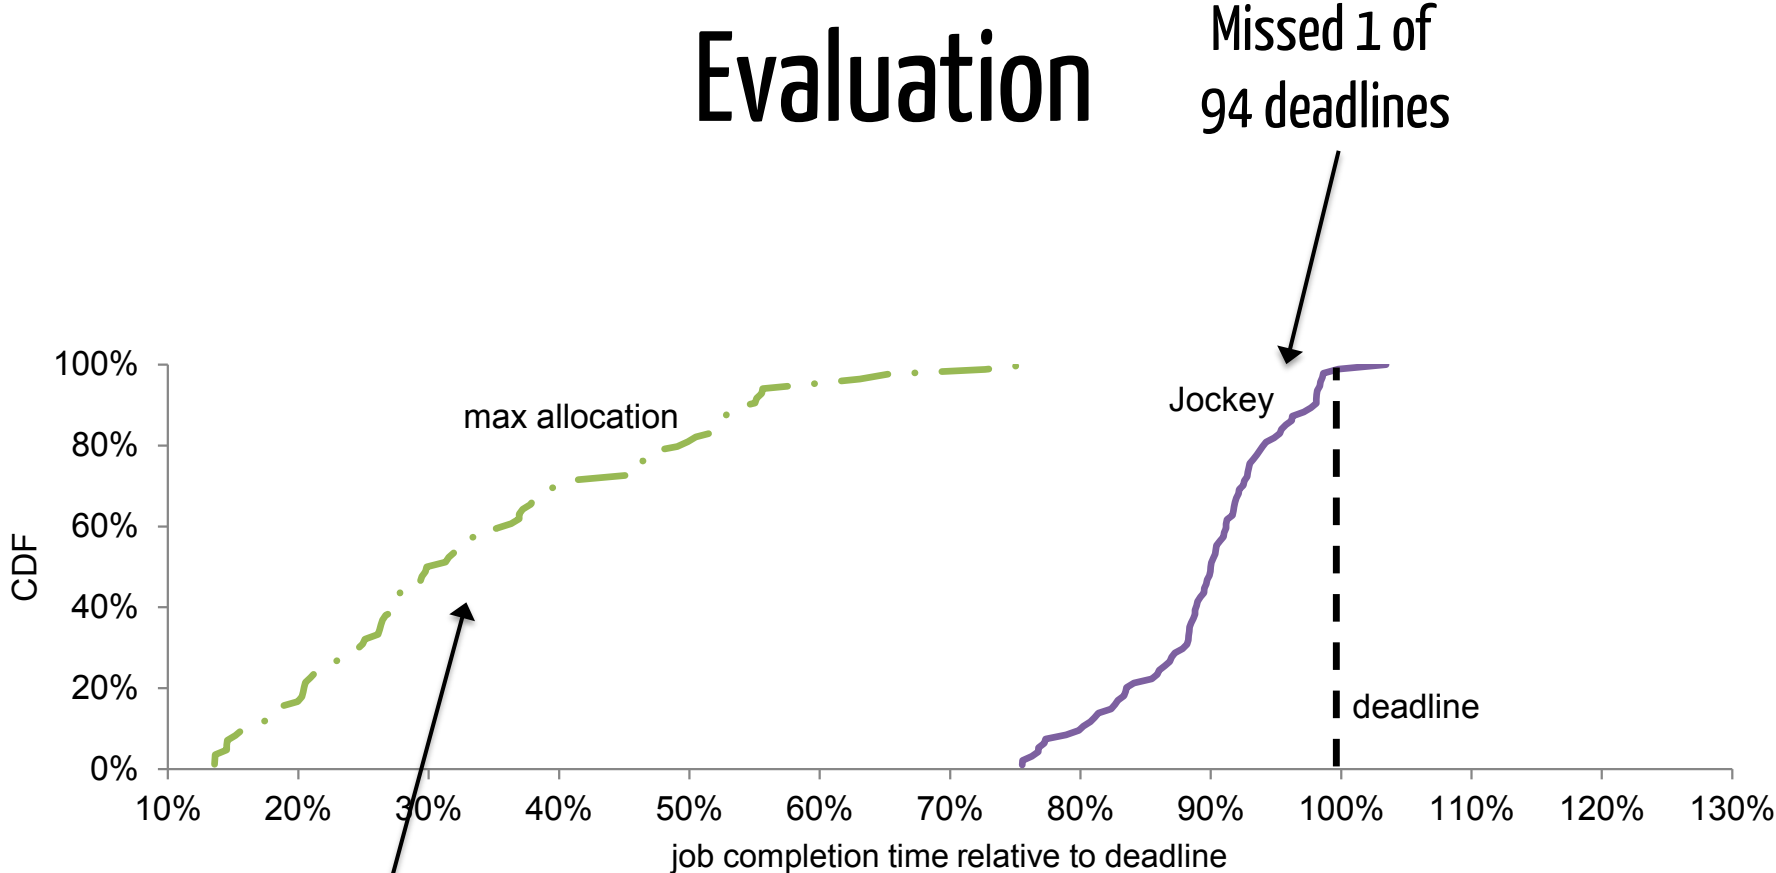




# Evaluation



# Evaluation



Missed 1 of 94 deadlines

max allocation

Jockey

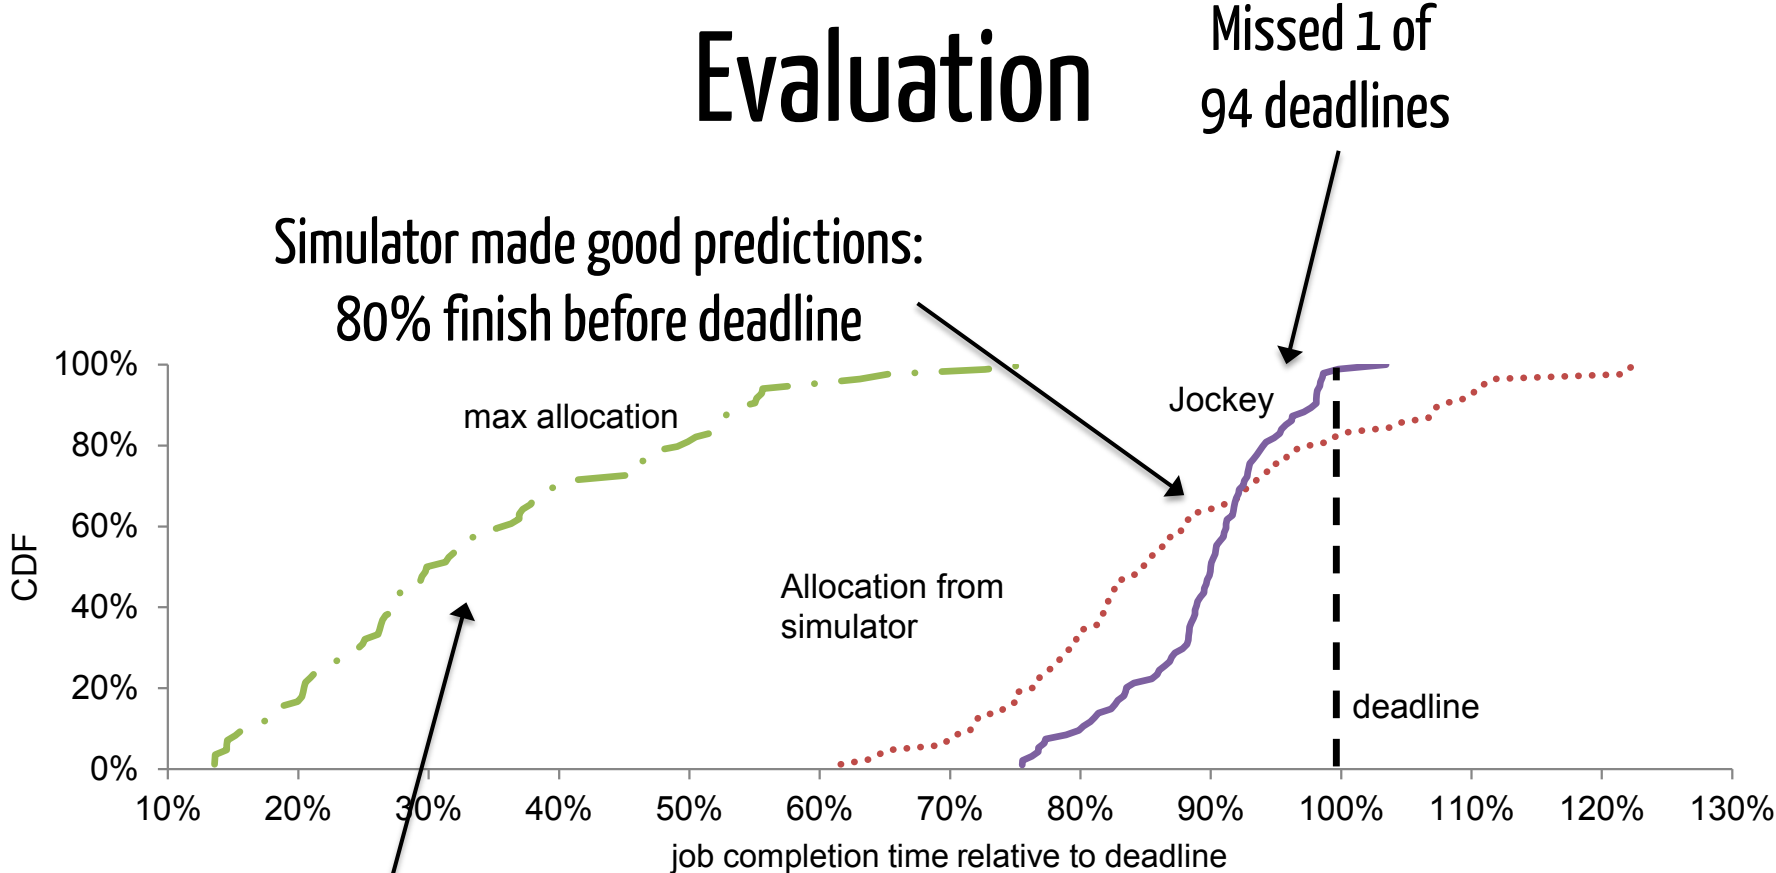
deadline

Allocated too many resources

Jobs which met the SLO



# Evaluation

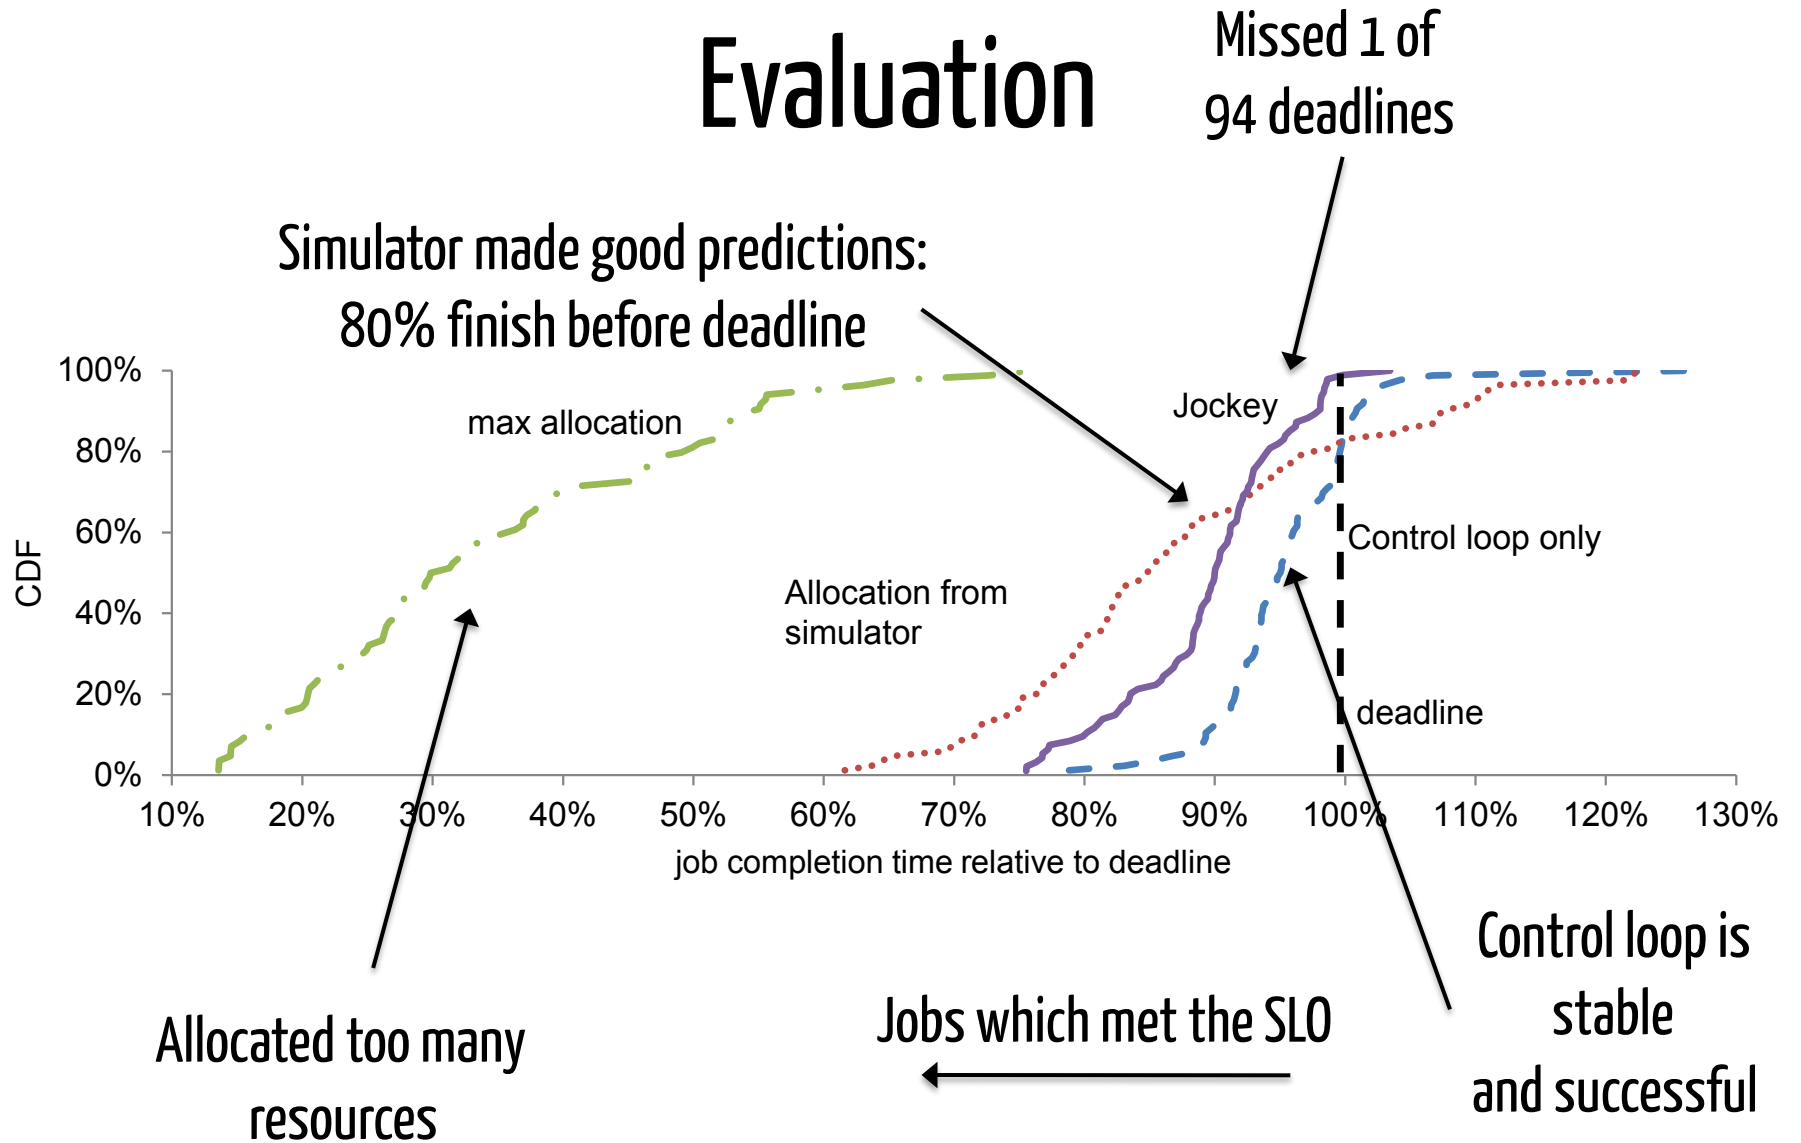


Allocated too many resources

Jobs which met the SLO

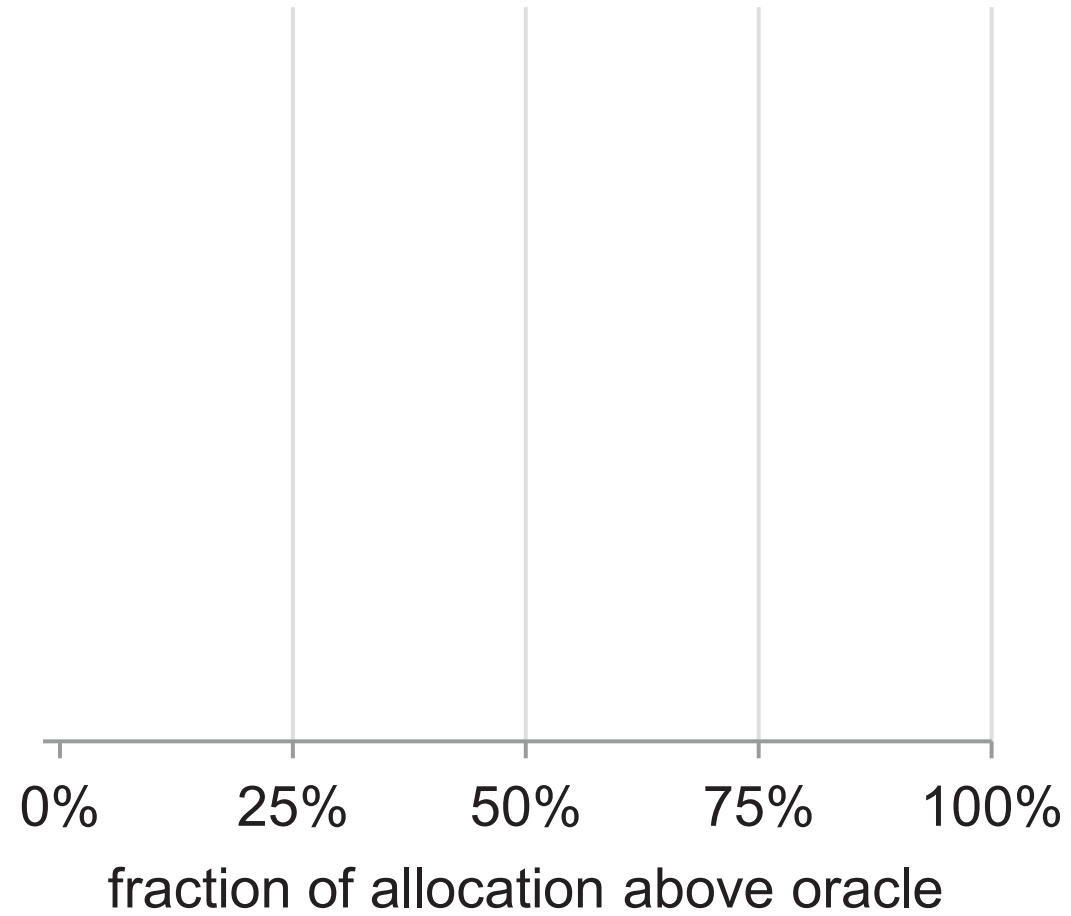


# Evaluation

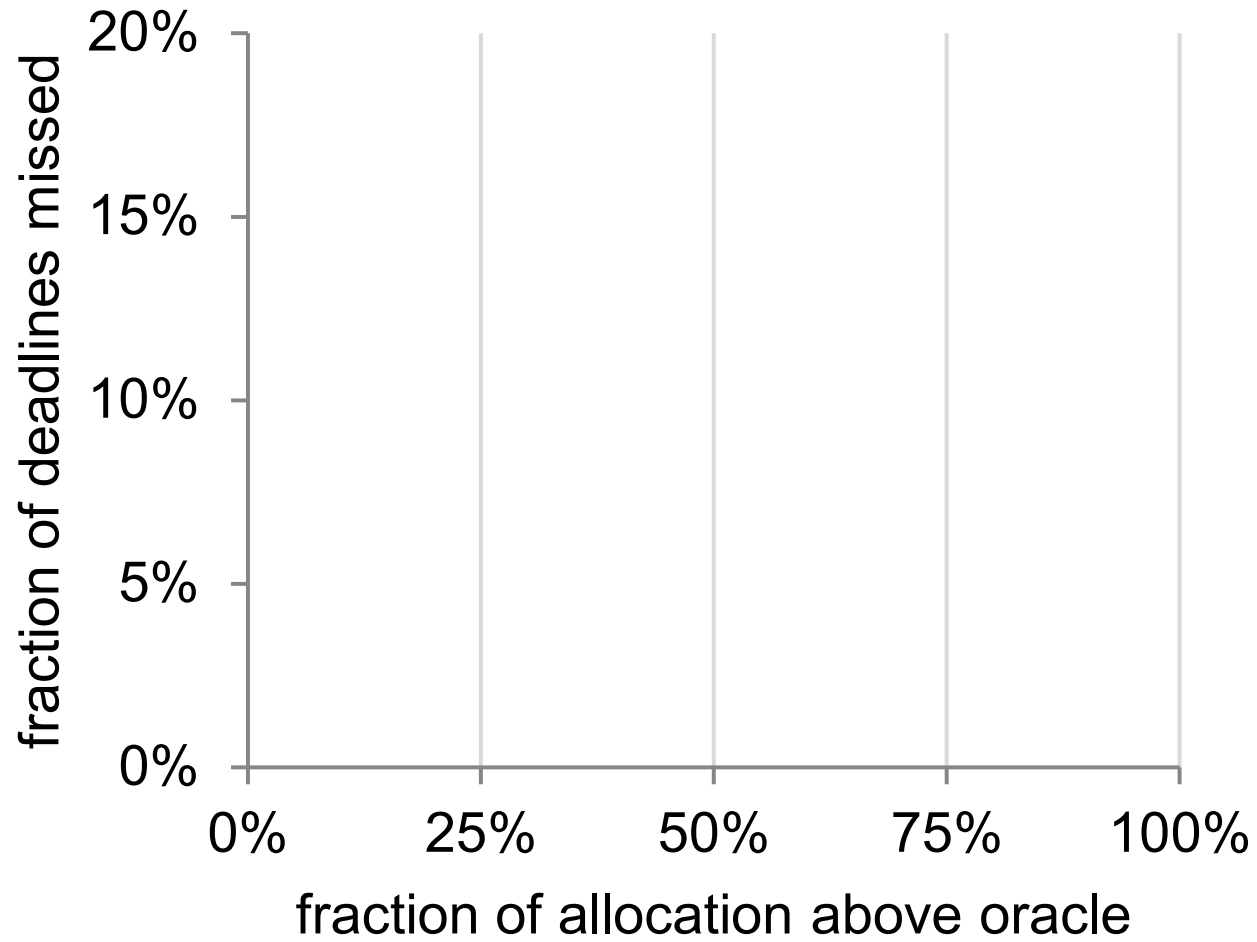


# Evaluation

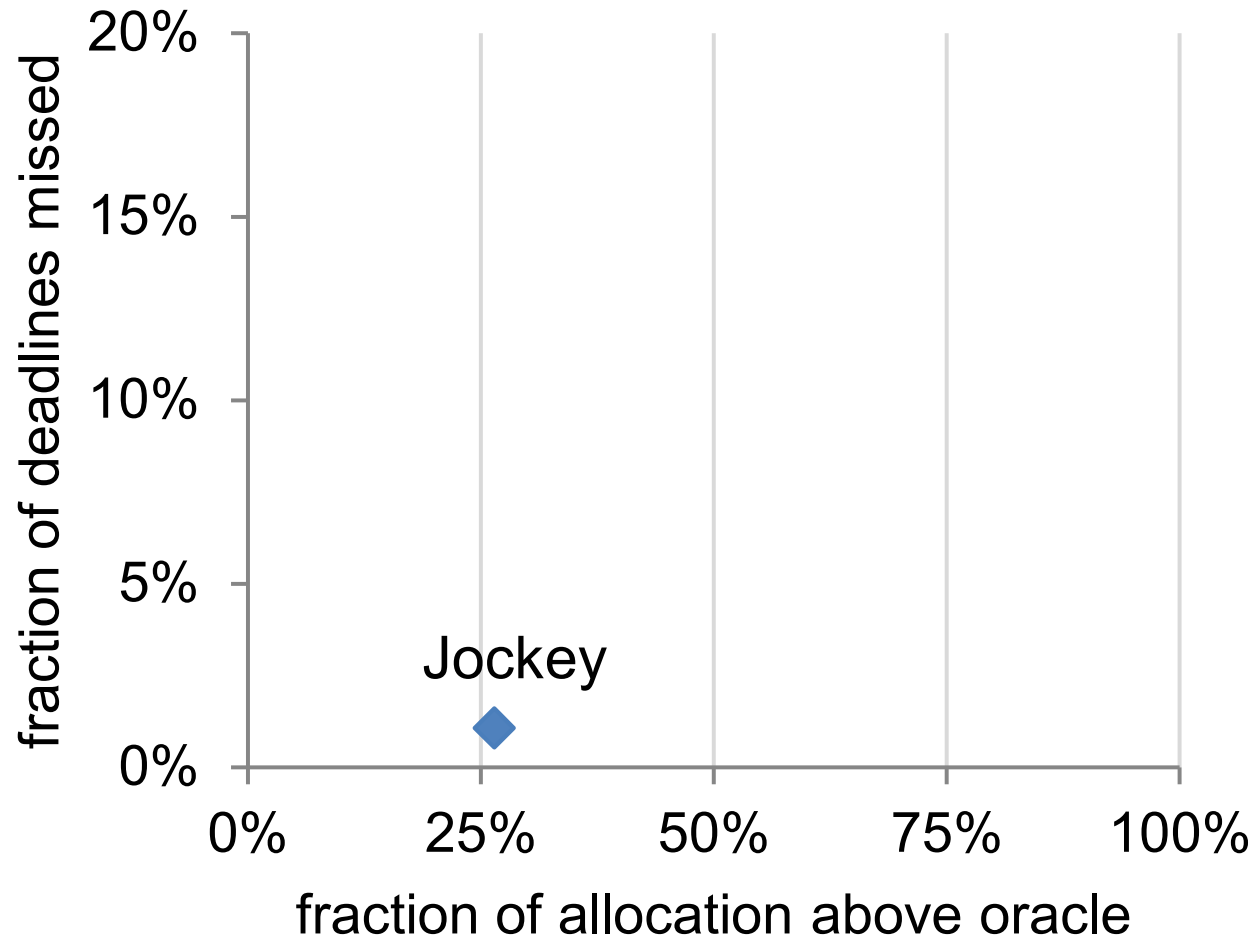
# Evaluation



# Evaluation

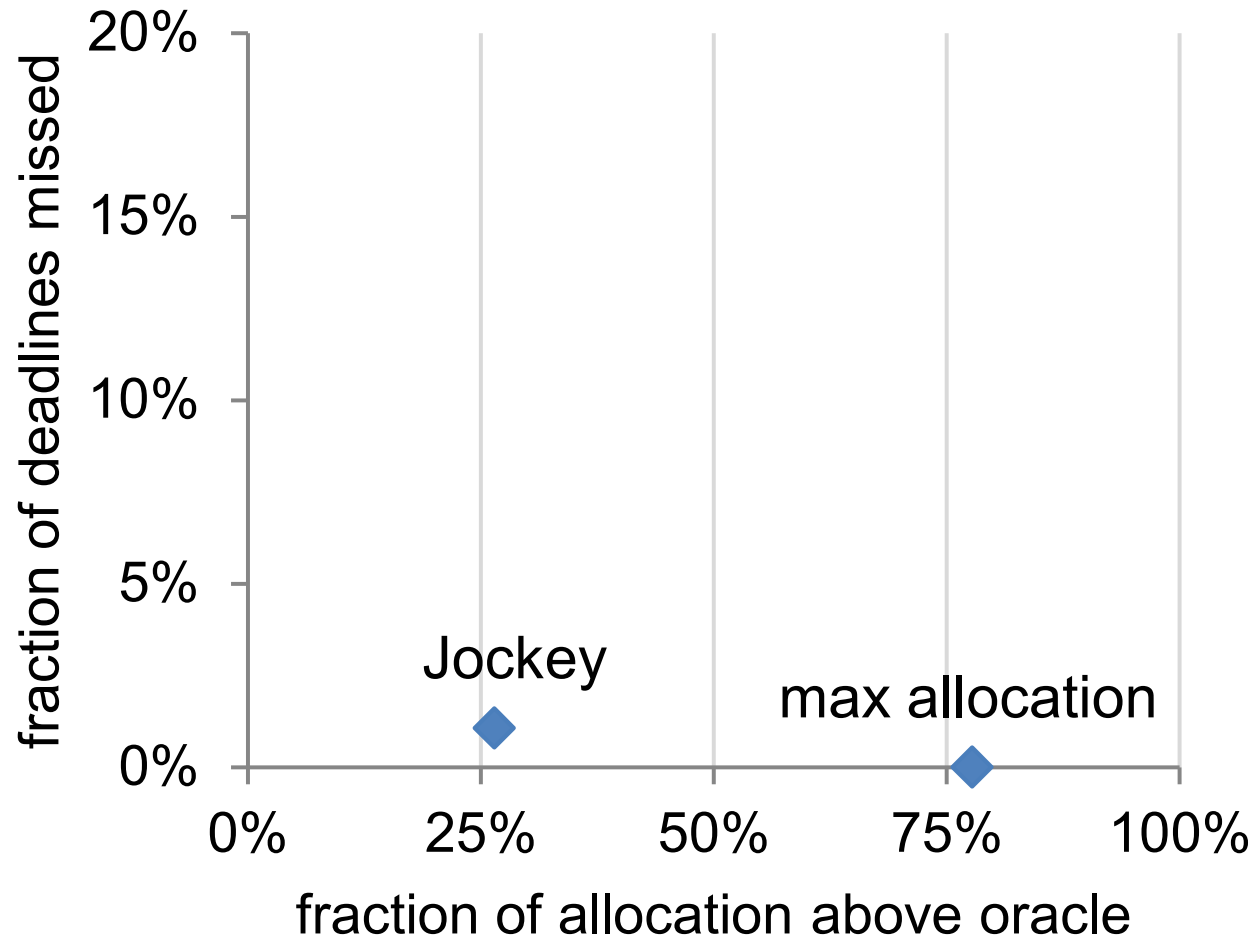


# Evaluation

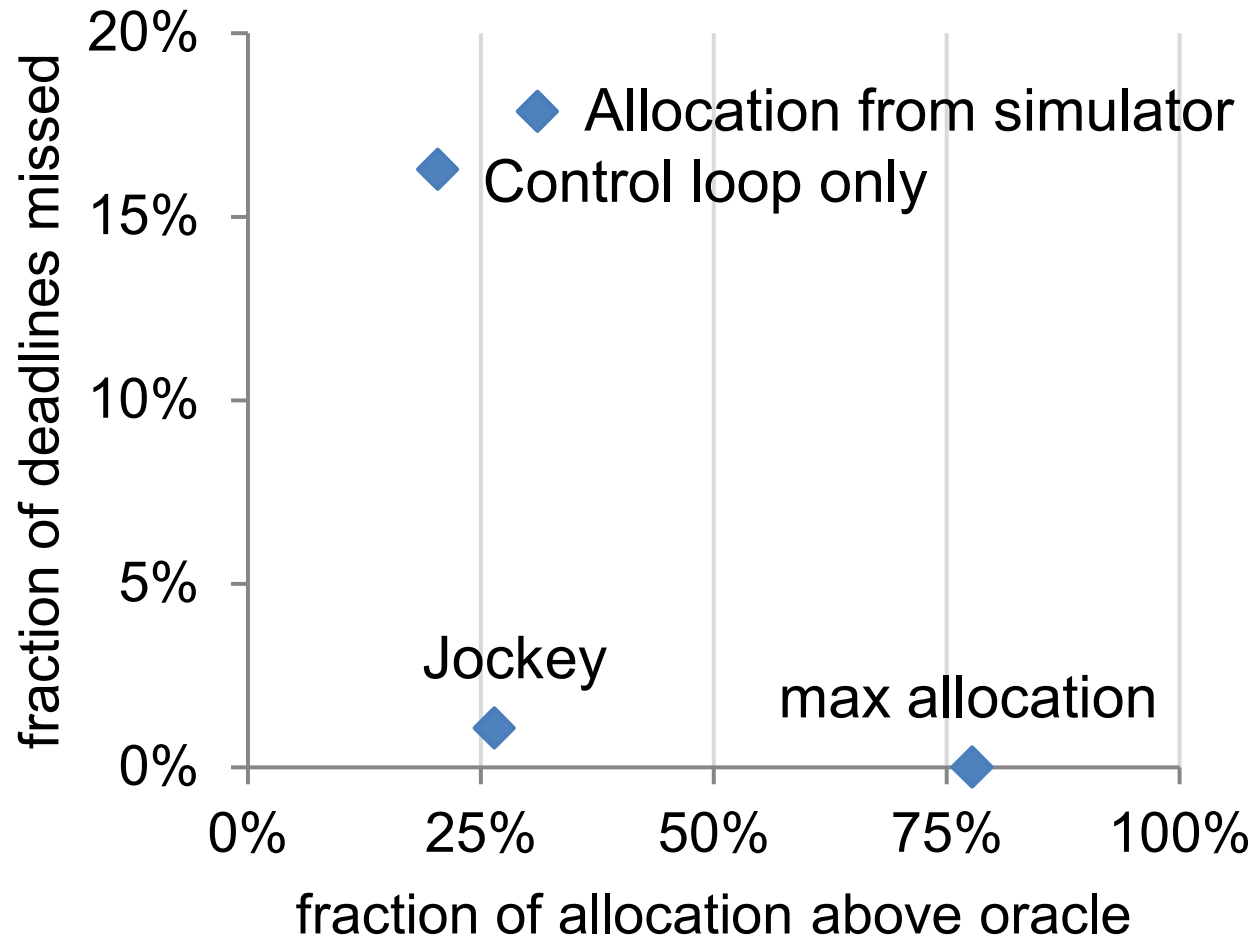




# Evaluation



# Evaluation



# Conclusion

**Data parallel jobs are complex,**

Data parallel jobs are complex,  
yet users demand deadlines.

Data parallel jobs are complex,  
yet users demand deadlines.  
Jobs run in shared, noisy clusters,

Data parallel jobs are complex,  
yet users demand deadlines.

Jobs run in shared, noisy clusters,  
making simple models inaccurate.

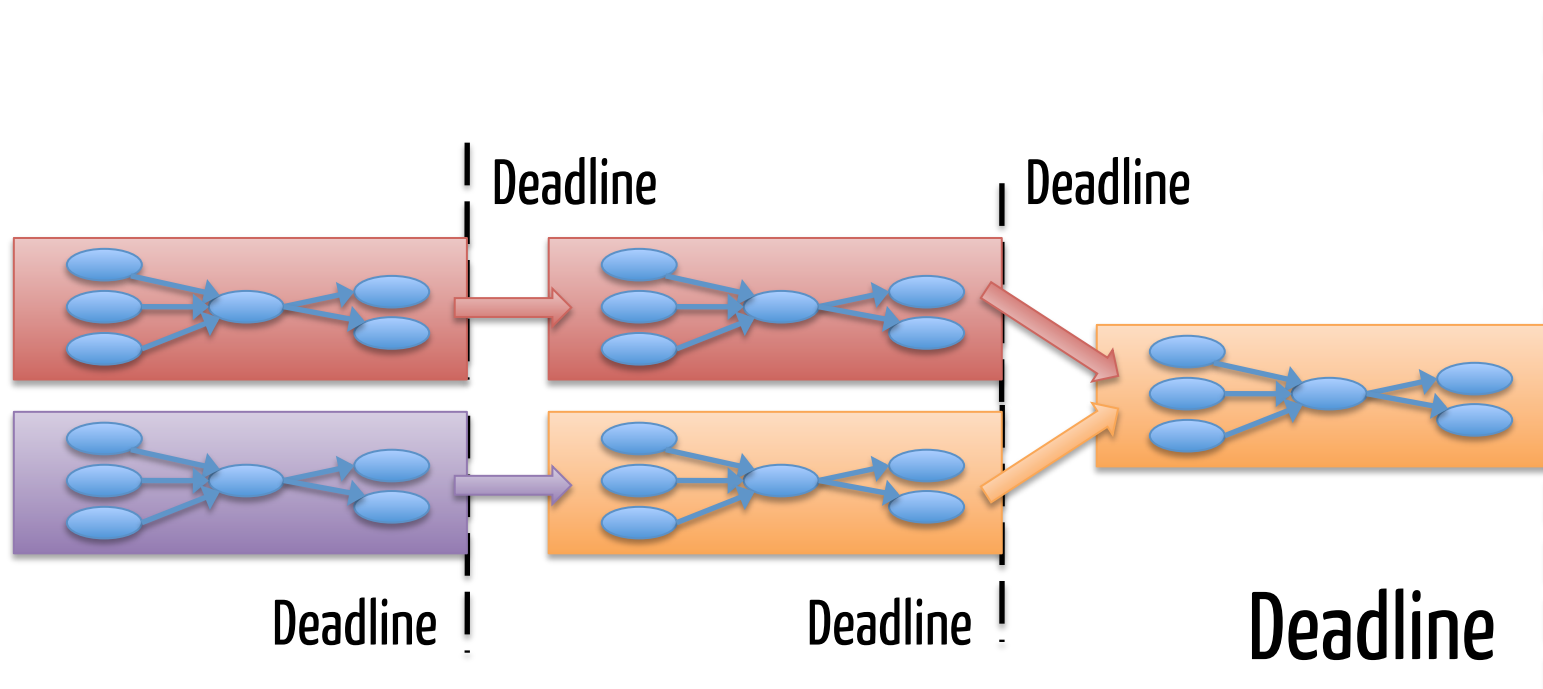
# Jockey

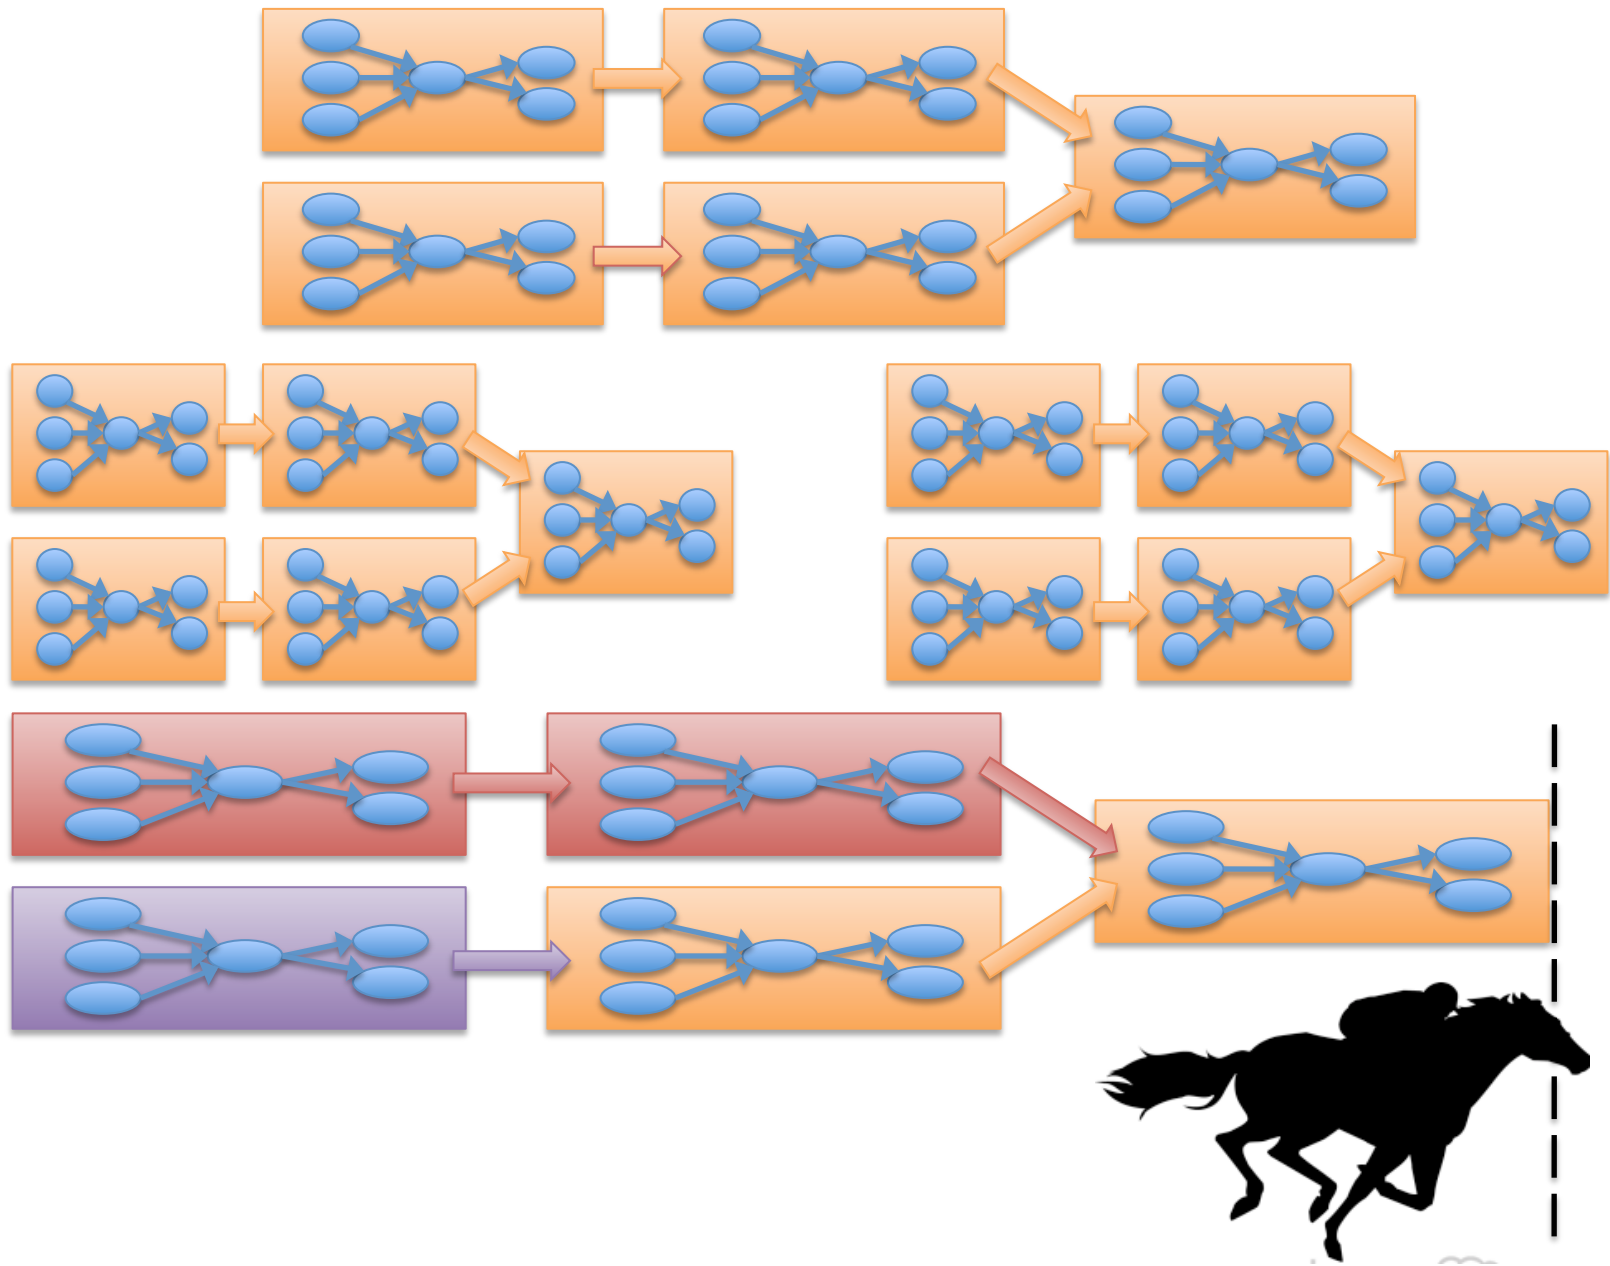




**simulator**

**control-loop**





# Questions?

Andrew Ferguson  
adf@cs.brown.edu

## Co-authors

- Peter Bodík  
(Microsoft Research)
- Srikanth Kandula  
(Microsoft Research)
- Eric Boutín  
(Microsoft)
- Rodrigo Fonseca  
(Brown)



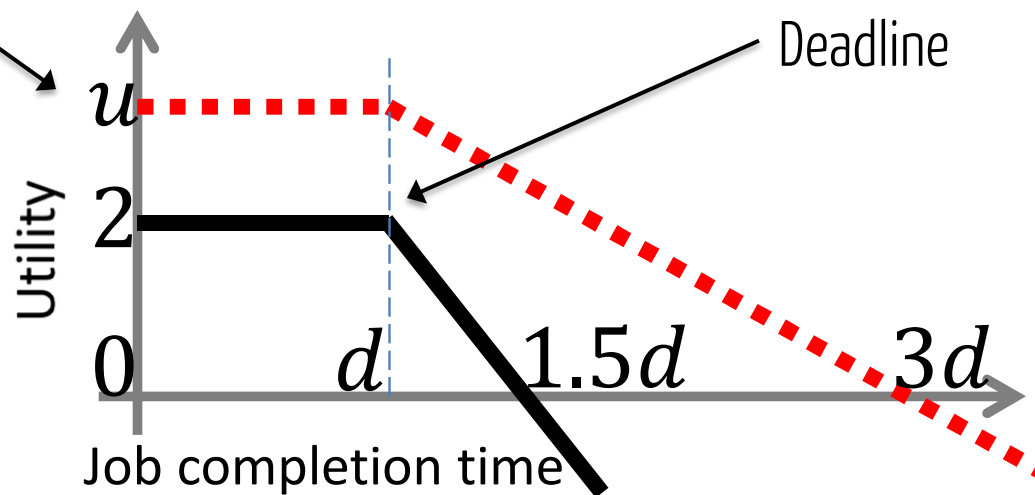
# Questions?

Andrew Ferguson  
adf@cs.brown.edu

# Backup Slides

# Utility Curves

For single jobs,  
scale doesn't matter



For multiple jobs,  
use financial penalties



# Jockey

## Resource allocation control loop

Utility      Run Time      Prediction

$$U_a = U(t_r + C(p, a))$$
$$A^r = \arg \min_a \{a : U_a = \max_b U_b\}$$

1. Slack

2. Hysteresis

$$A_t^s = A_{t-1}^s + \alpha(A^r - A_{t-1}^s)$$

3. Dead Zone

# Cosmos

## Resource sharing in Cosmos

- Resources are allocated with a form of fair sharing across business groups and their jobs. (Like Hadoop FairScheduler or CapacityScheduler)
- Each job is guaranteed a number of *tokens* as dictated by cluster policy; each running or initializing task uses one token. Token released on task completion.
- A token is a guaranteed share of CPU and memory
- To increase efficiency, unused tokens are re-allocated to jobs with available work

# Jockey

## Progress indicator

- Can use many features of the job to build a progress indicator
- Earlier work (ParaTimer) concentrated on fraction of tasks completed
- Our indicator is very simple, but we found it performs best for Jockey's needs

Fraction of completed vertices

Total vertex initialization time

$$\sum_{\text{stage } s} f_s (Q_s + T_s)$$

Total vertex run time

# Comparison with ARIA

- ARIA uses analytic models
- Designed for 3 stages: Map, Shuffle, Reduce
- Jockey's control loop is robust due to control-theory improvements
- ARIA tested on small (66-node) cluster without a network bottleneck
- We believe Jockey is a better match for production DAG frameworks such as Hive, Pig, etc.

# Jockey

## Latency prediction: $C(p, a)$

- Event-based simulator
  - Same scheduling logic as actual Job Manager
  - Captures important features of job progress
  - Does not model input size variation or speculative re-execution of stragglers
  - Inputs: job algebra, distributions of task timings, probabilities of failures, allocation
- Analytic model
  - Inspired by Amdahl's Law:  $T = S + P/N$
  - $S$  is remaining work on critical path,  $P$  is all remaining work,  $N$  is number of machines

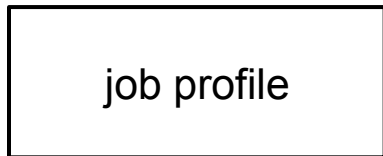
# Jockey

## Resource allocation control loop

- Executes in Dryad's Job Manager
- Inputs: fraction of completed tasks in each stage, time job has spent running, utility function, precomputed values (for speedup)
- Output: Number of tokens to allocate
- Improved with techniques from control-theory

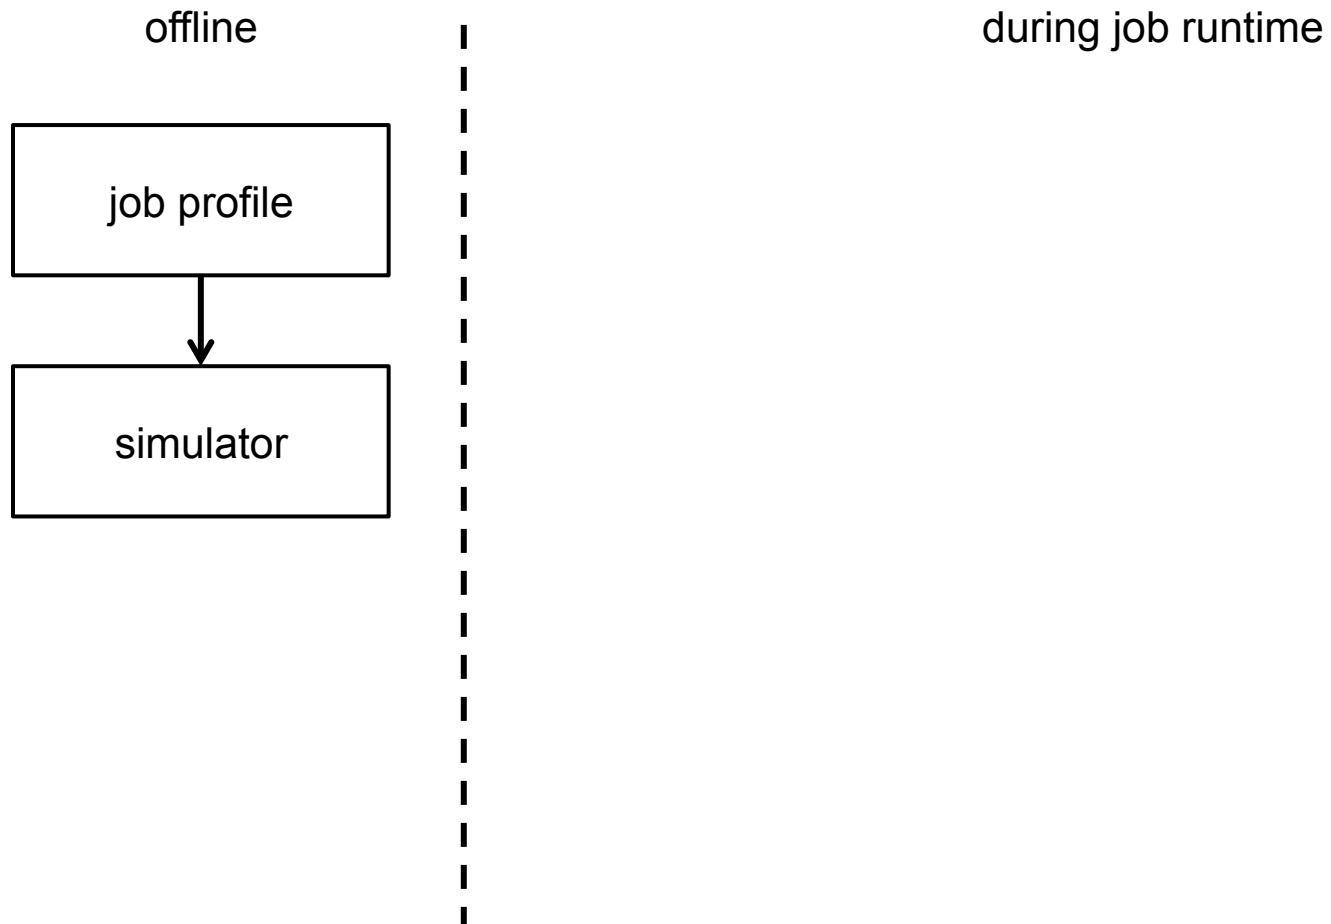
# Jockey

offline



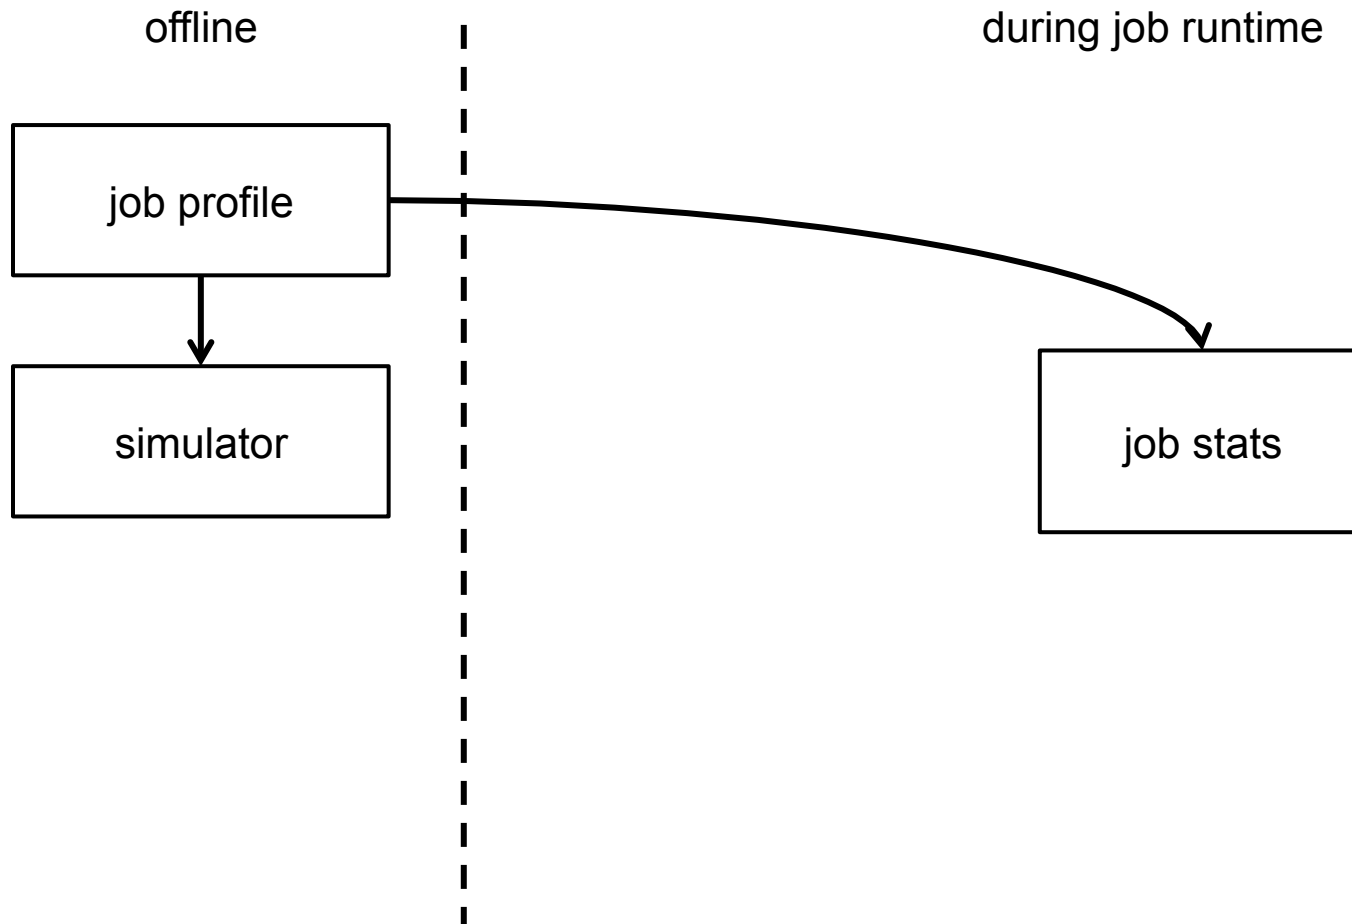
during job runtime

# Jockey

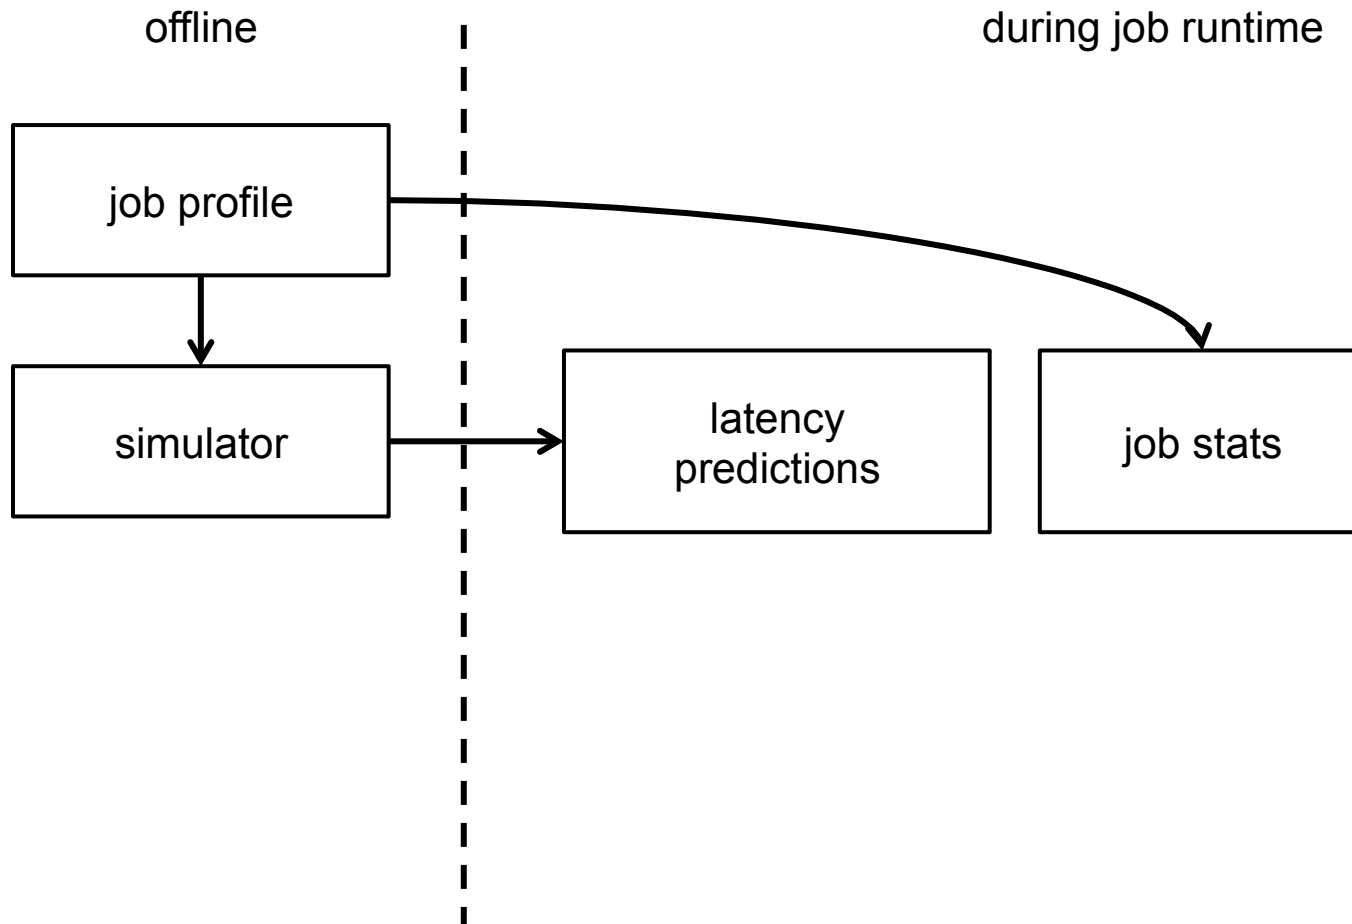




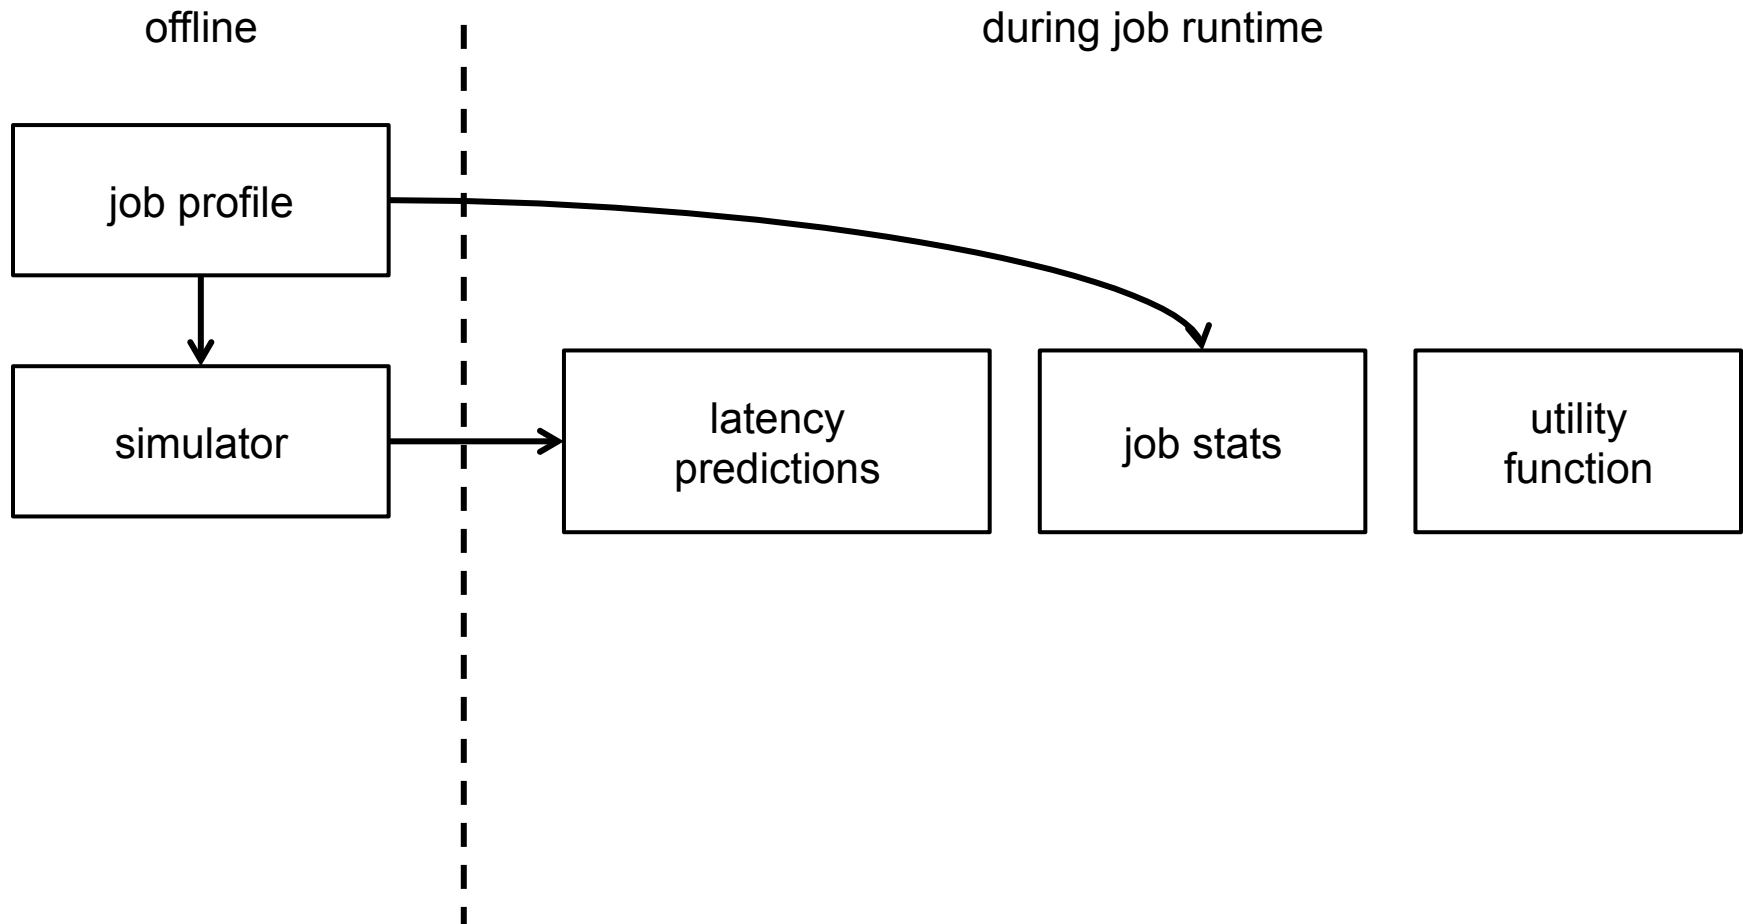
# Jockey



# Jockey



# Jockey



# Jockey

