

# Speed Scaling with Sum-Power Constraint

Rahul Vaze



Tata Institute of Fundamental Research, Mumbai



Joint work Jayakrishnan Nair- IIT-Bombay


# Scheduling 101

**Job  
arrivals**




# Scheduling 101

**Job  
arrivals**




Single server

# Scheduling 101




# Scheduling 101



**Obj:** *min* total flow time

$$\sum_{j \in \text{Jobs}} \text{flow time}_j = \int n(t) dt$$

# Scheduling 101




**Obj: *min* total flow time**

$$\sum_{j \in \text{Jobs}} \text{flow time}_j = \int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Scheduling 101




**Obj: *min* total flow time**

$$\sum_{j \in \text{Jobs}} \text{flow time}_j = \int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$

**SRPT is optimal**

# Scheduling 101



**Obj: *min* total flow time**


$$\sum_{j \in \text{Jobs}} \text{flow time}_j = \int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$

**SRPT is optimal**



# Scheduling 101



**Obj: *min* total flow time**

$$\sum_{j \in \text{Jobs}} \text{flow time}_j = \int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$





**SRPT is optimal**

*min* # of outstanding jobs

# Parallel Scheduling

# Parallel Scheduling

packets  
arrivals



**Multiple Servers**





**Obj: *min* total flow time**

$$\int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling

packets  
arrivals



**Multiple Servers**

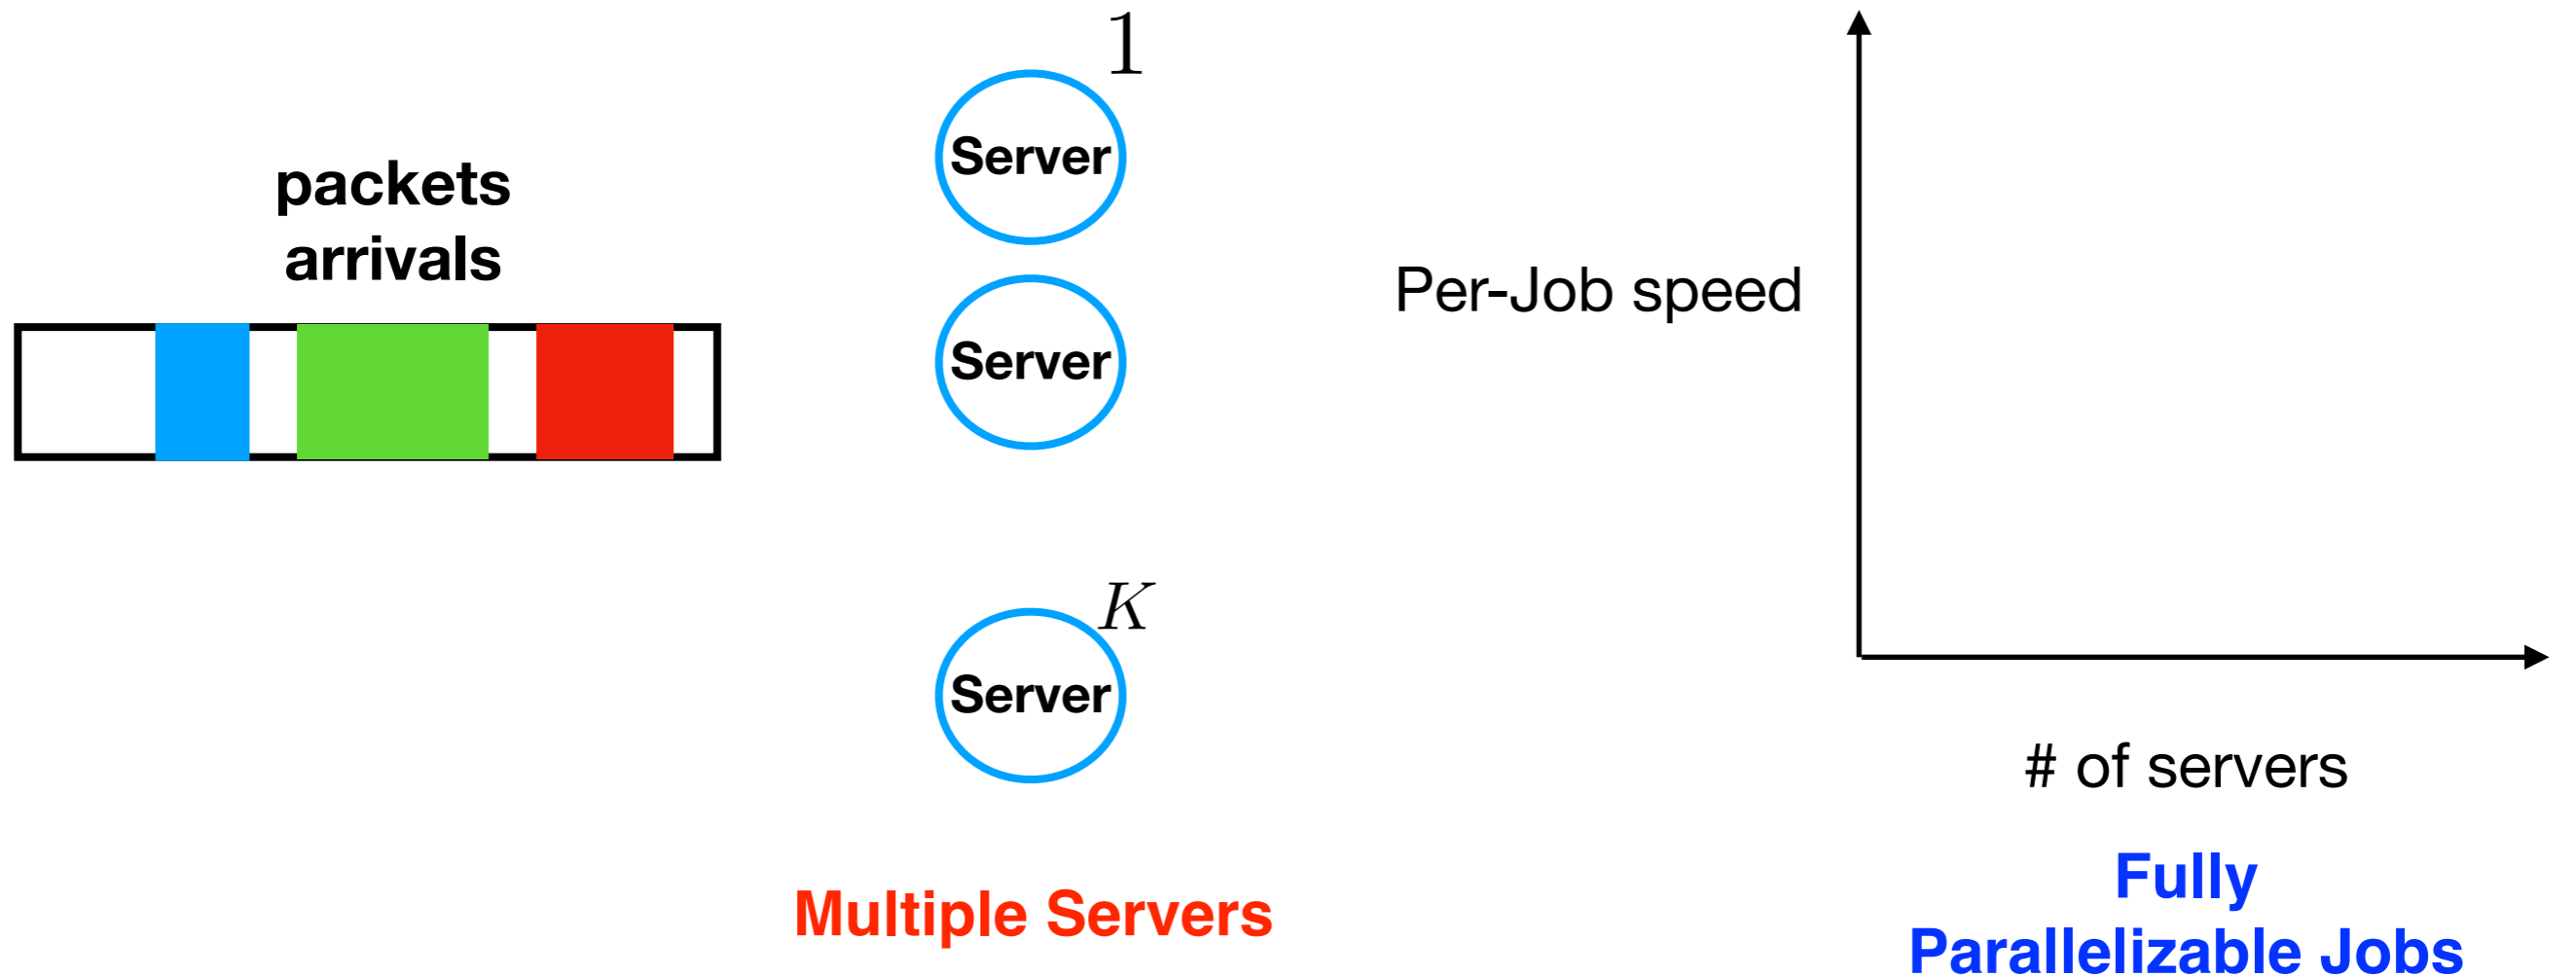
**Fully  
Parallelizable Jobs**

**Obj: *min* total flow time**

$$\int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling

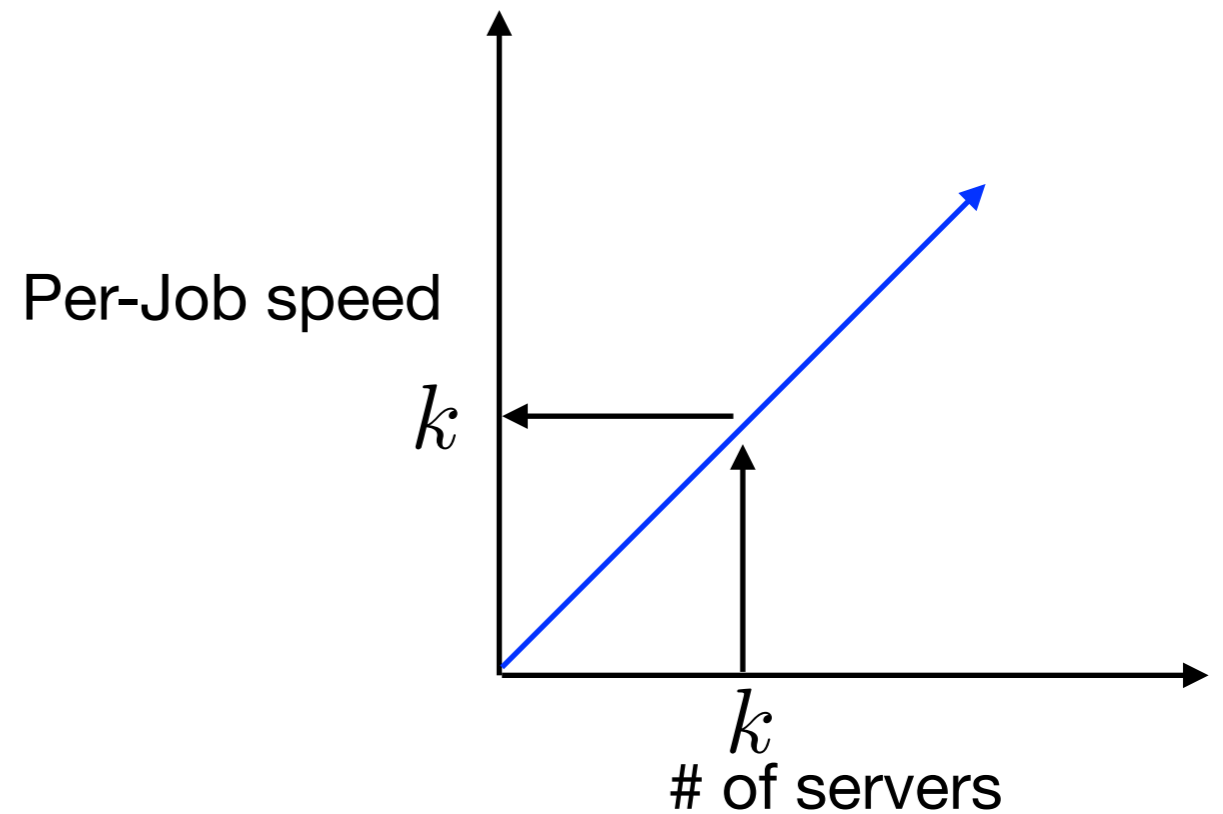



**Obj: *min* total flow time**

$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling





**Fully  
Parallelizable Jobs**

**Obj: *min* total flow time**

$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling



**Fully  
Parallelizable Jobs**

**Obj: *min* total flow time**

$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time t

**SRPT is optimal**

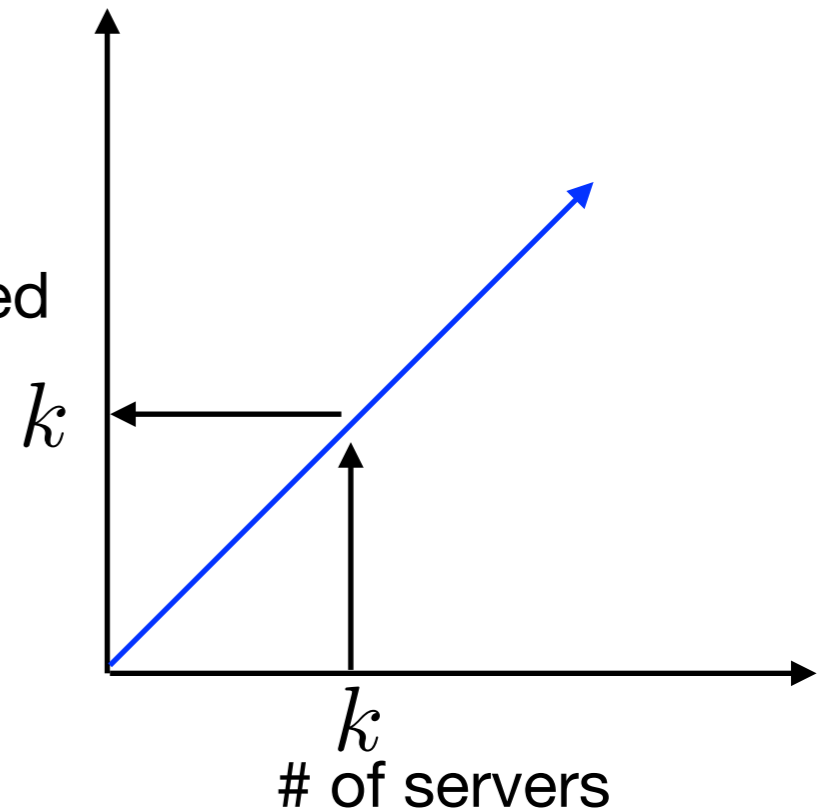
# Parallel Scheduling

packets  
arrivals



**Multiple Servers**

Per-Job speed



**Fully  
Parallelizable Jobs**

**Obj: *min* total flow time**


$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time  $t$

**SRPT is optimal**



# Parallel Scheduling- Reality





**Obj: *min* total flow time**

$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling- Reality

packets  
arrivals



**Multiple Servers**

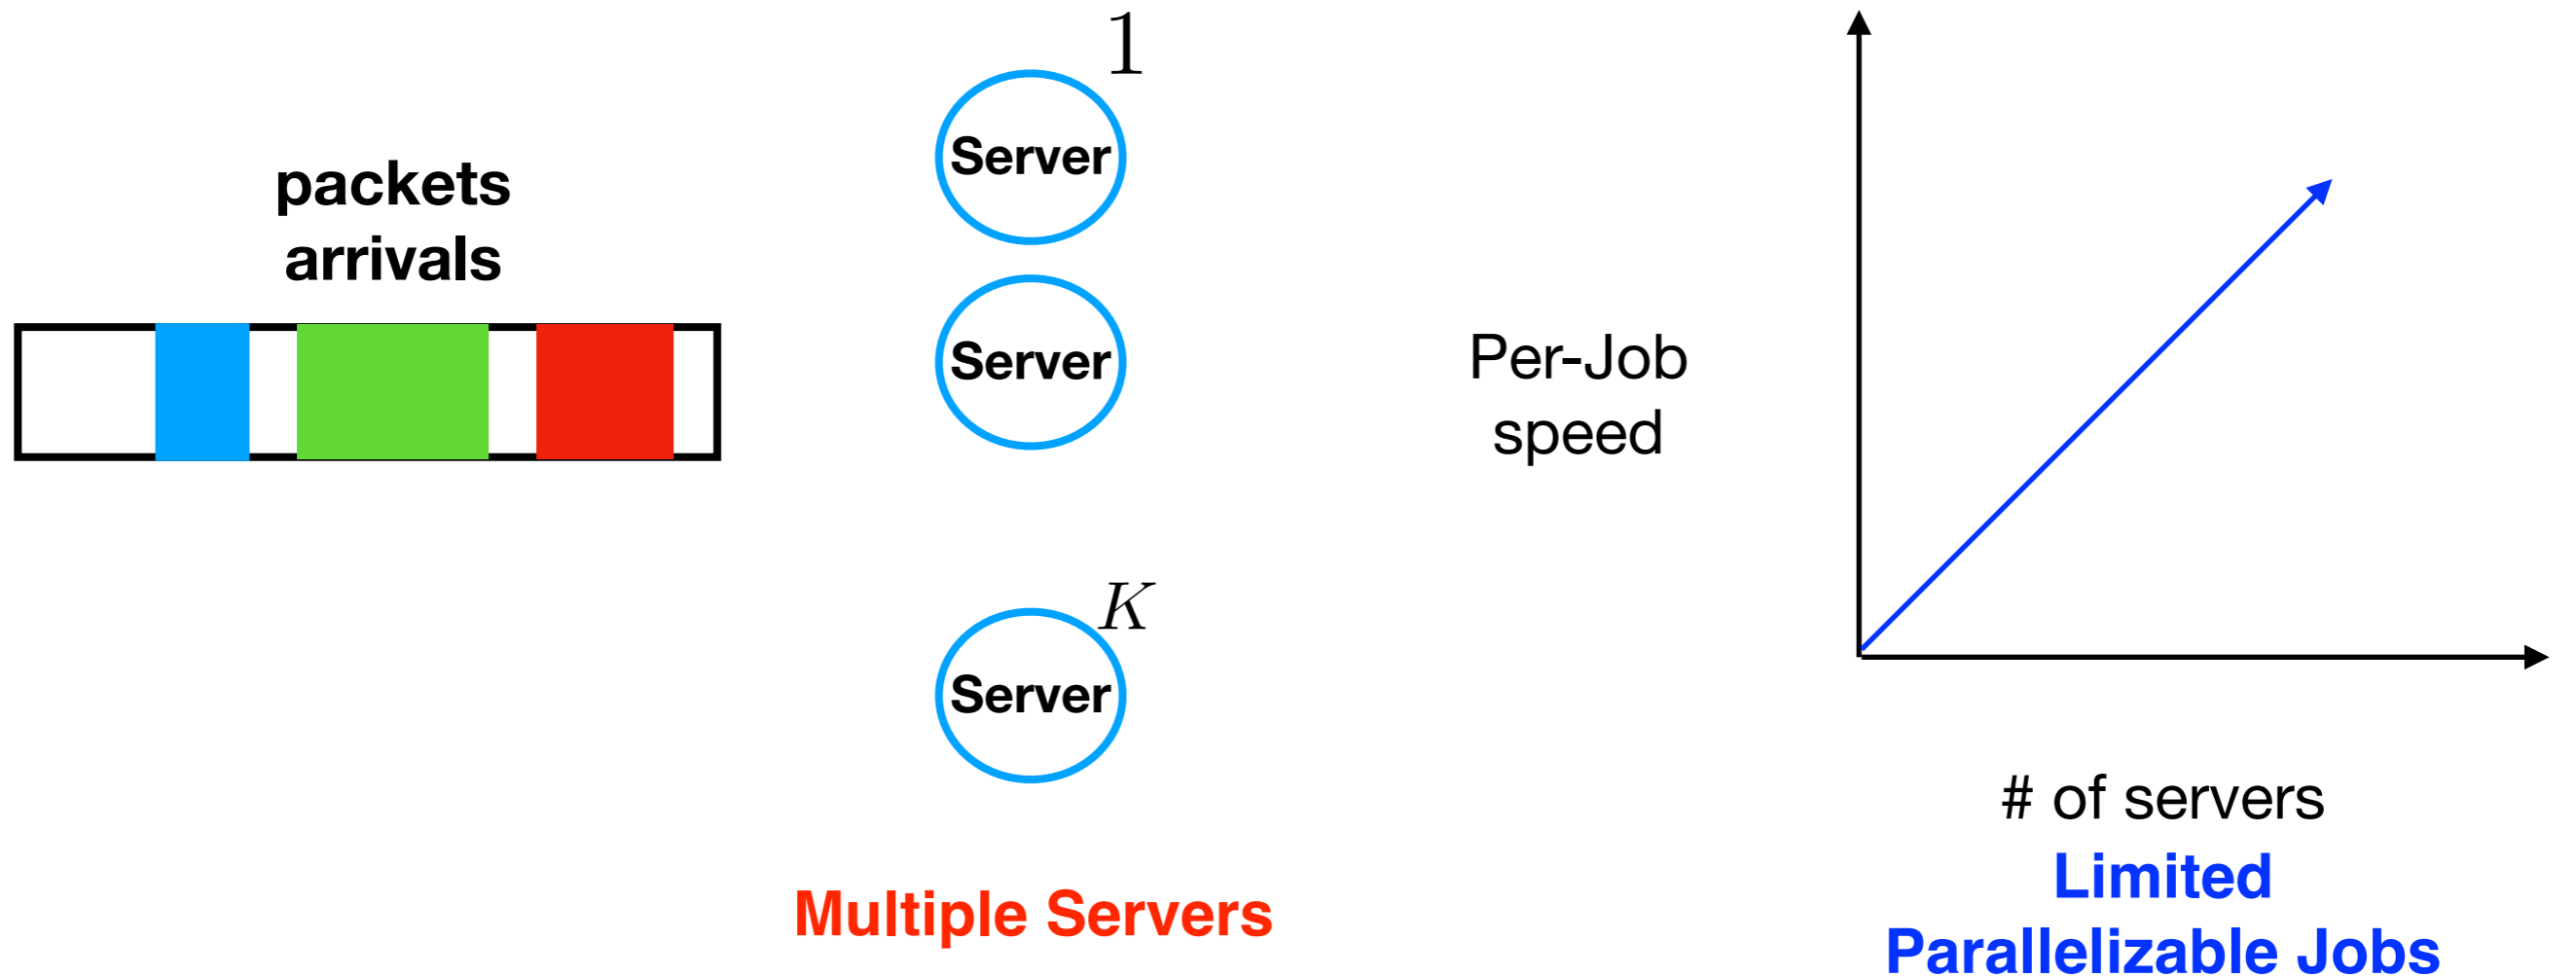
**Limited  
Parallelizable Jobs**

**Obj: *min* total flow time**

$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time t

# Parallel Scheduling- Reality

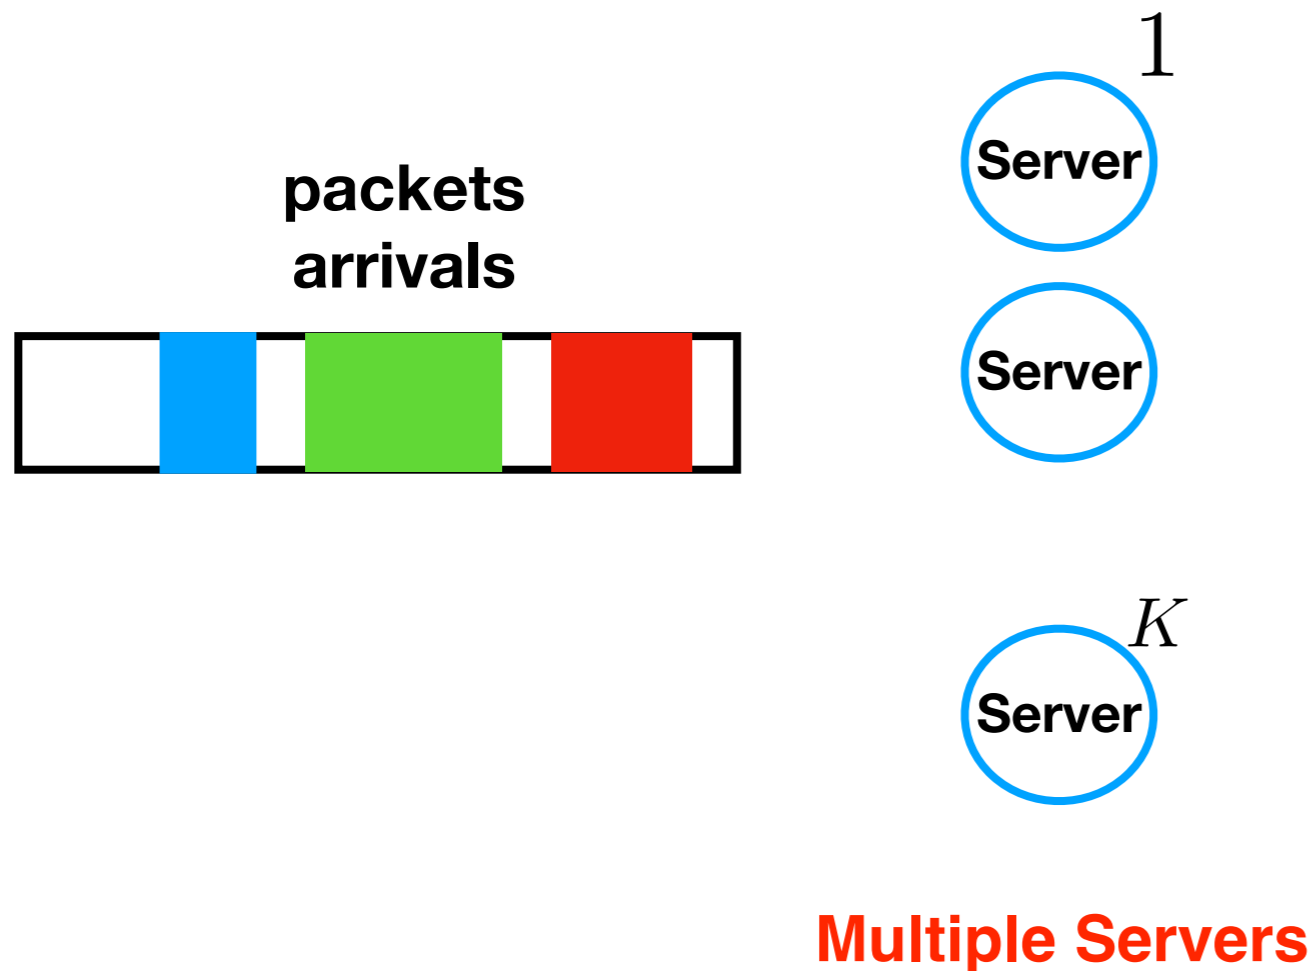


**Obj: *min* total flow time**

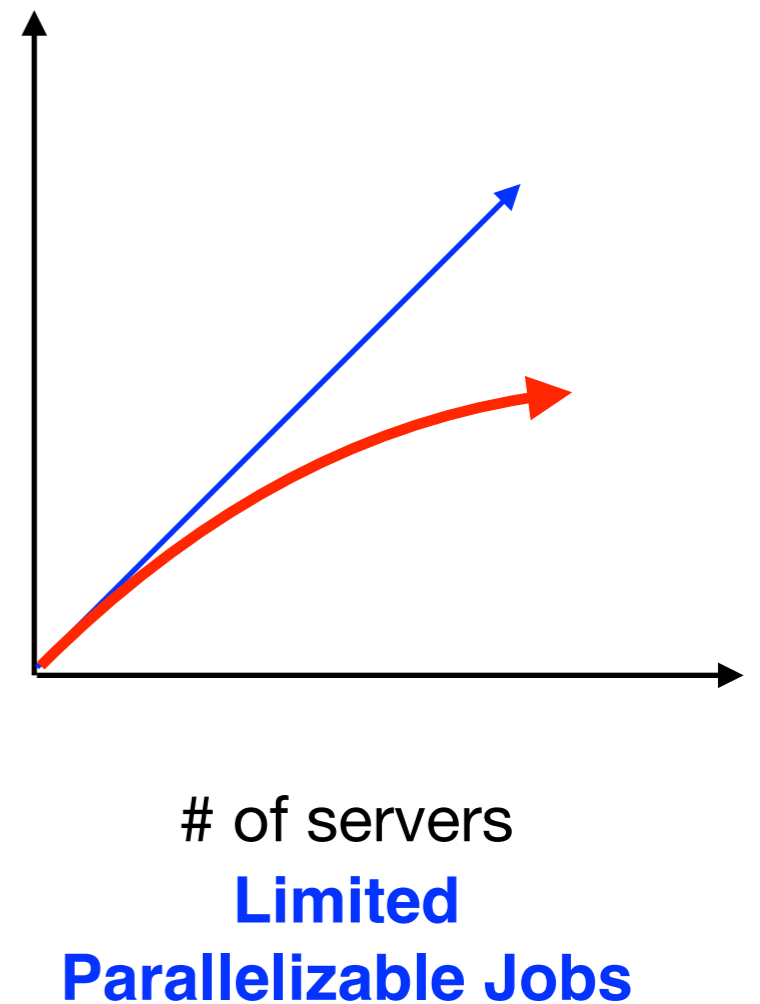
$$\int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling- Reality



Per-Job speed




**Obj: *min* total flow time**

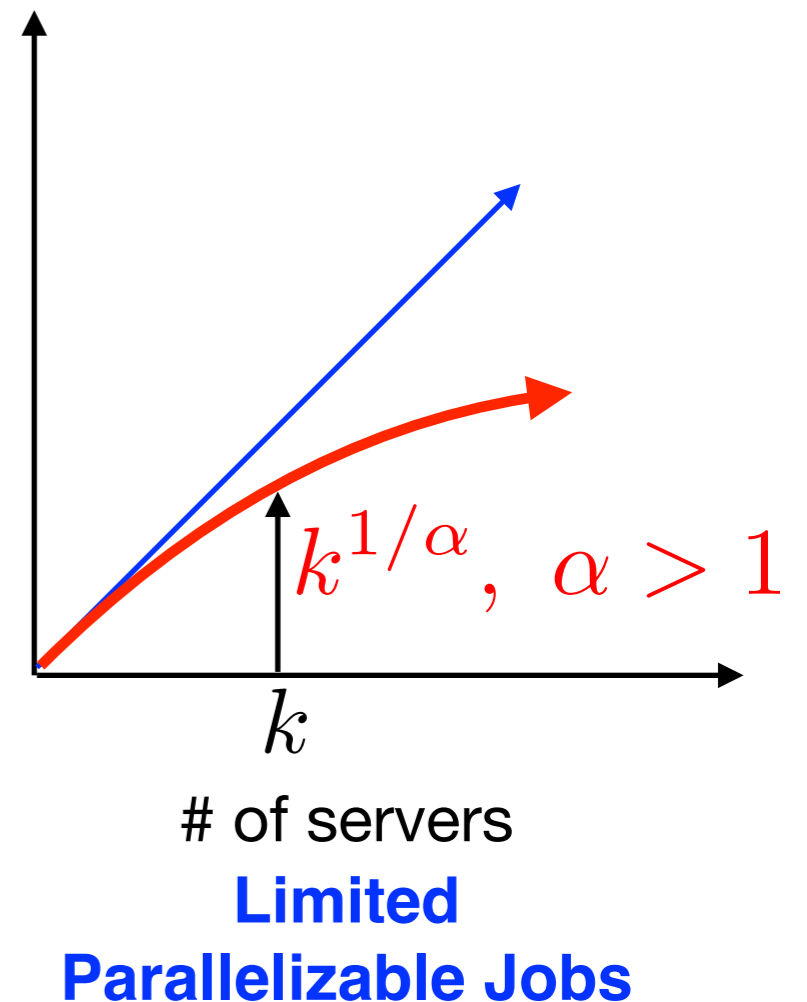
$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling- Reality



Per-Job speed




Obj: *min* total flow time


$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time t

# Parallel Scheduling- Reality



Per-Job speed



**Limited  
Parallelizable Jobs**


Diminishing returns by using more servers for the same job

**Obj: *min* total flow time**


$$\int n(t)dt$$

$n(t)$  number of outstanding jobs at time  $t$

# Parallel Scheduling- Reality



Per-Job speed



**Limited Parallelizable Jobs**

Diminishing returns by using more servers for the same job


**Obj: *min* total flow time**

$$\int n(t) dt$$


$n(t)$  number of outstanding jobs at time  $t$

**Open Question:  
Optimal Scheduling**

# Prior Work




Per-Job speed




**Limited  
Parallelizable Jobs**



# Prior Work



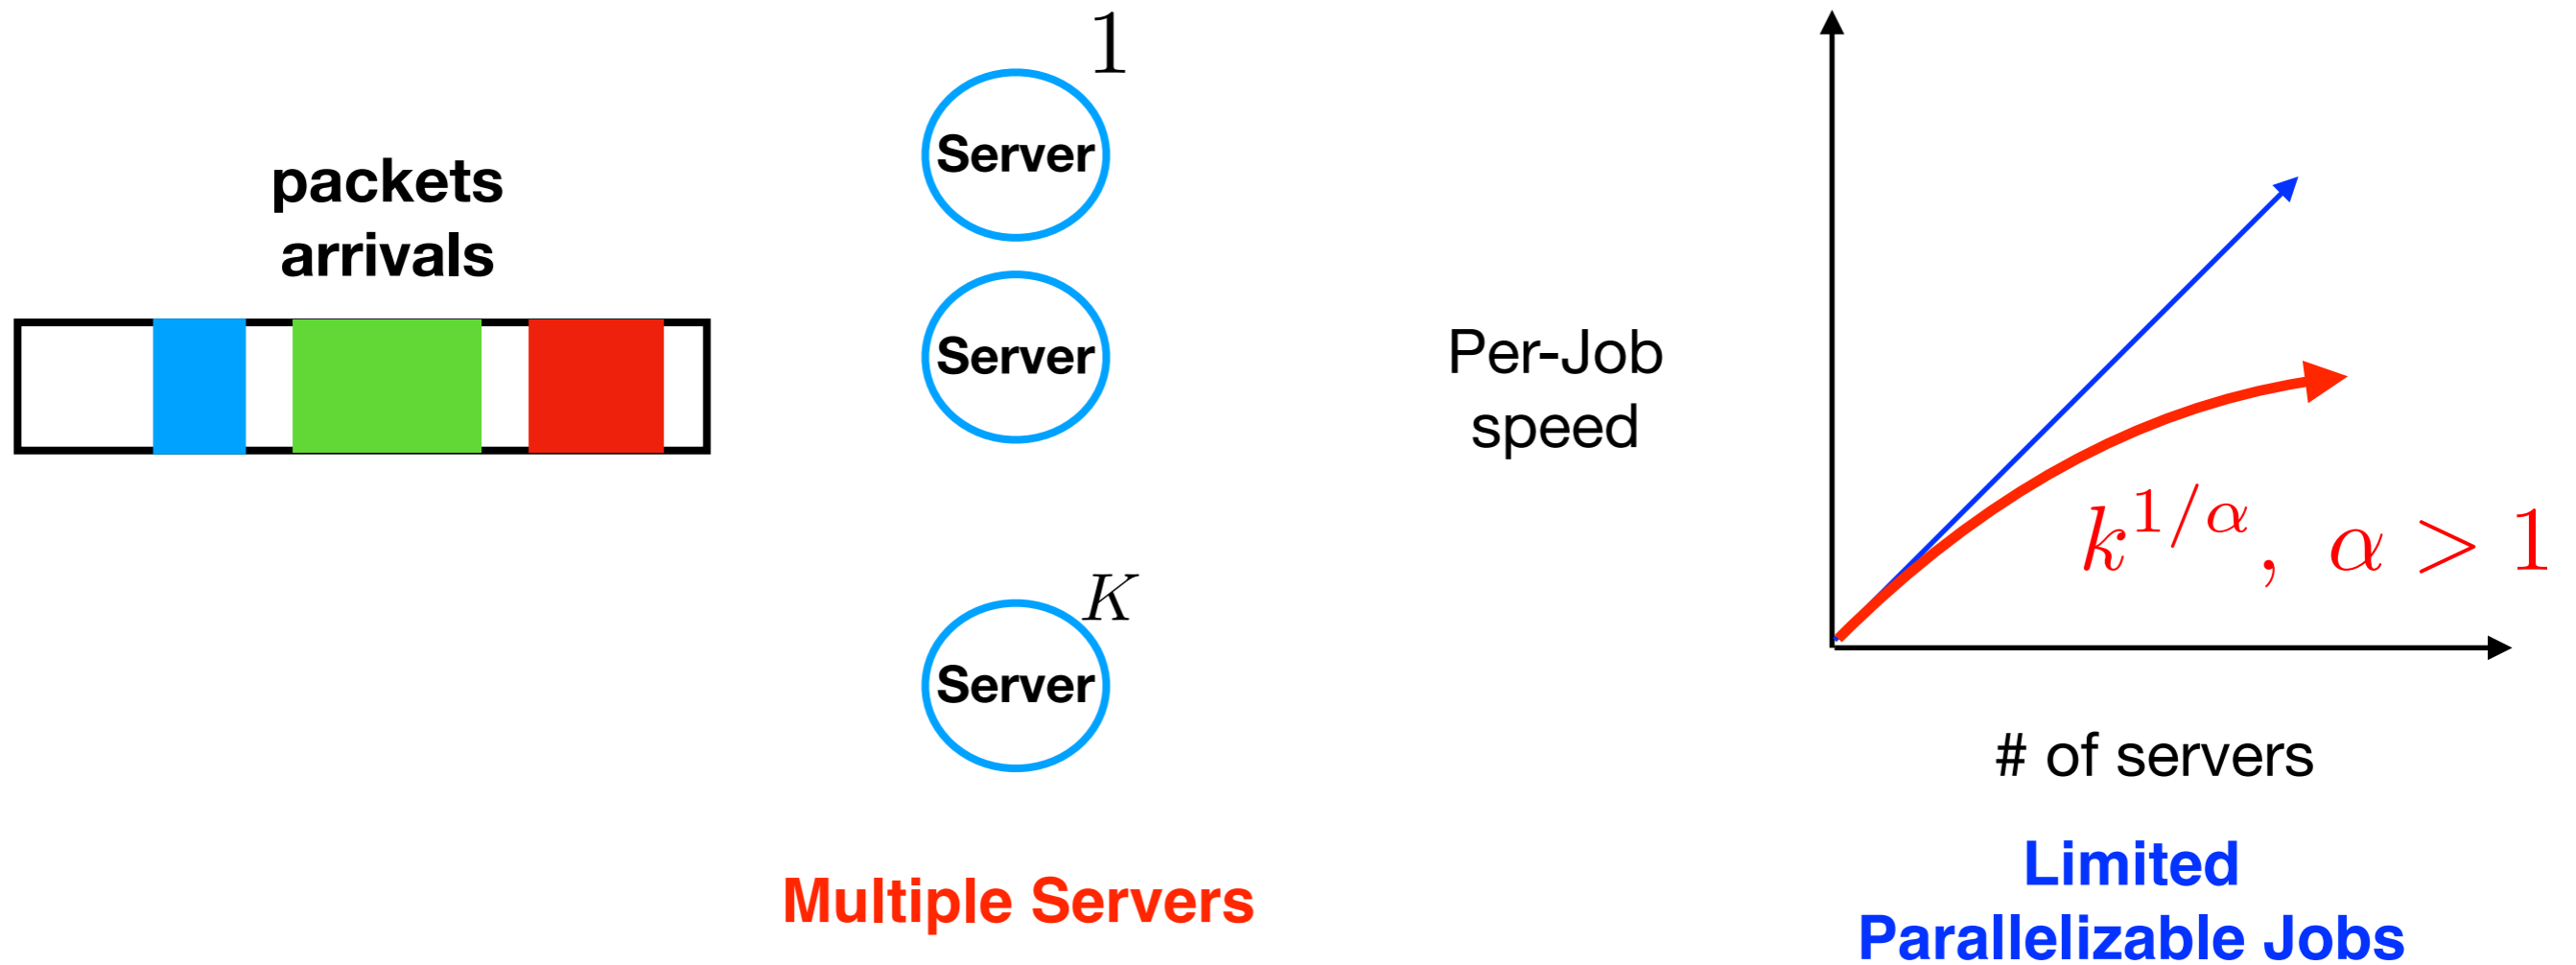
Per-Job speed



**Limited  
Parallelizable Jobs**


**All jobs available at time 0,**

# Prior Work



All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]


# Prior Work



**Multiple Servers**

**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

# Prior Work




**Multiple Servers**

**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

All jobs get non-zero speed, while shorter jobs get more speed


# Prior Work



**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

All jobs get non-zero speed, while shorter jobs get more speed


# Prior Work



**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

All jobs get non-zero speed, while shorter jobs get more speed


# Prior Work



**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

All jobs get non-zero speed, while shorter jobs get more speed

# Prior Work




**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

All jobs get non-zero speed, while shorter jobs get more speed




# Prior Work



**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

All jobs get non-zero speed, while shorter jobs get more speed

# Prior Work



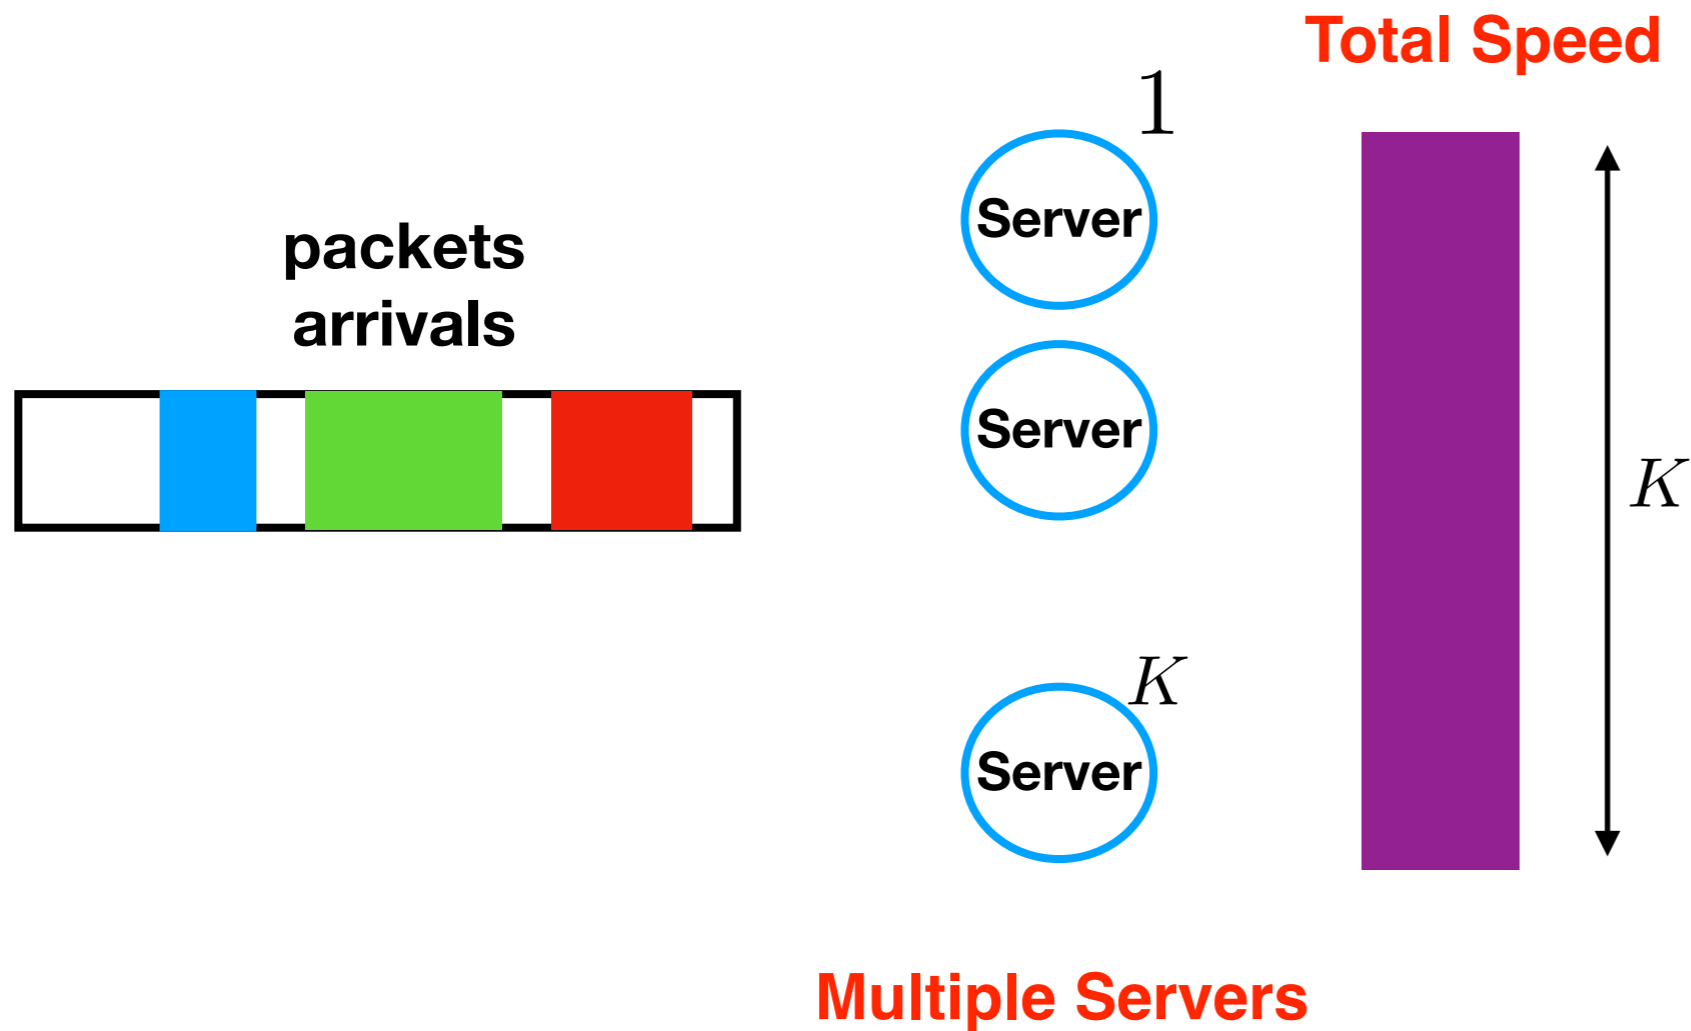
**All jobs available at time 0, Optimal Scheduling : heSRPT [Berg et al' 20]**

All jobs get non-zero speed, while shorter jobs get more speed

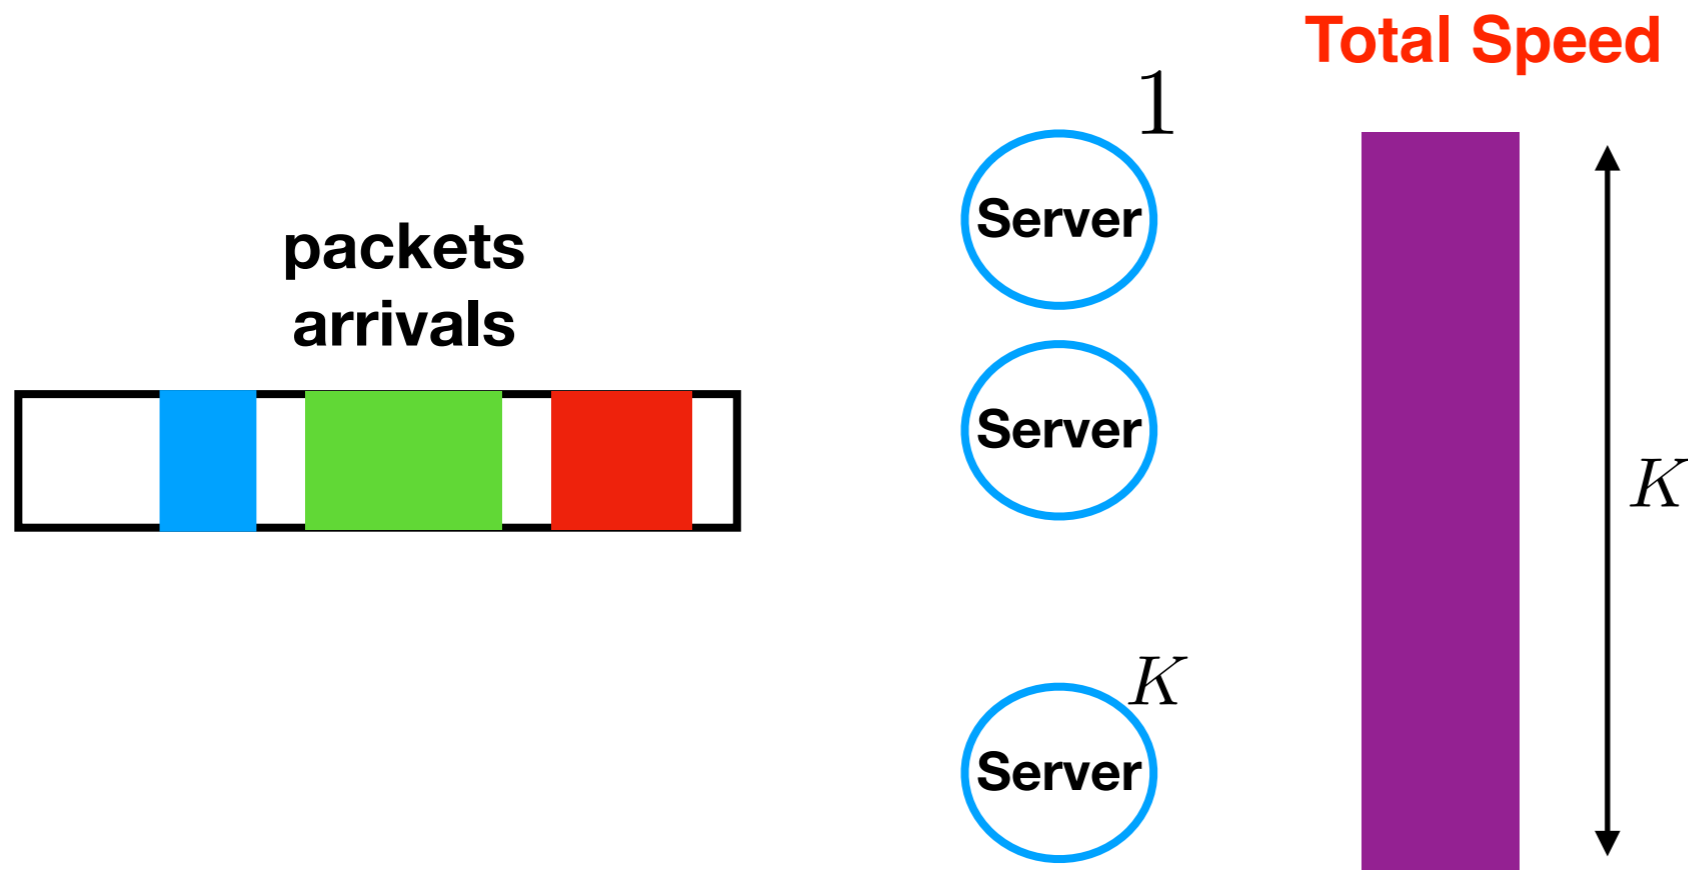
Among  $n(t)$  outstanding jobs  
 $i^{\text{th}}$  shortest job's speed

$$s_i(t) = K \left( \left( \frac{i}{n(t)} \right)^{\left( \frac{1}{1-1/\alpha} \right)} - \left( \frac{i-1}{n(t)} \right)^{\left( \frac{1}{1-1/\alpha} \right)} \right)$$

# Simple Algorithm EQUI




# Simple Algorithm EQUI



Multiple Servers

All jobs available at time 0,

# Simple Algorithm EQUI




**Multiple Servers**

**All jobs available at time 0,**

Assign equal speed to all jobs

# Simple Algorithm $EQU$




**Multiple Servers**

**All jobs available at time 0,**

Assign equal speed to all jobs

# Simple Algorithm EQUI




**Multiple Servers**

**All jobs available at time 0,**

Assign equal speed to all jobs

# Simple Algorithm $EQU$




**Multiple Servers**

**All jobs available at time 0,**

Assign equal speed to all jobs



# Simple Algorithm $EQU$

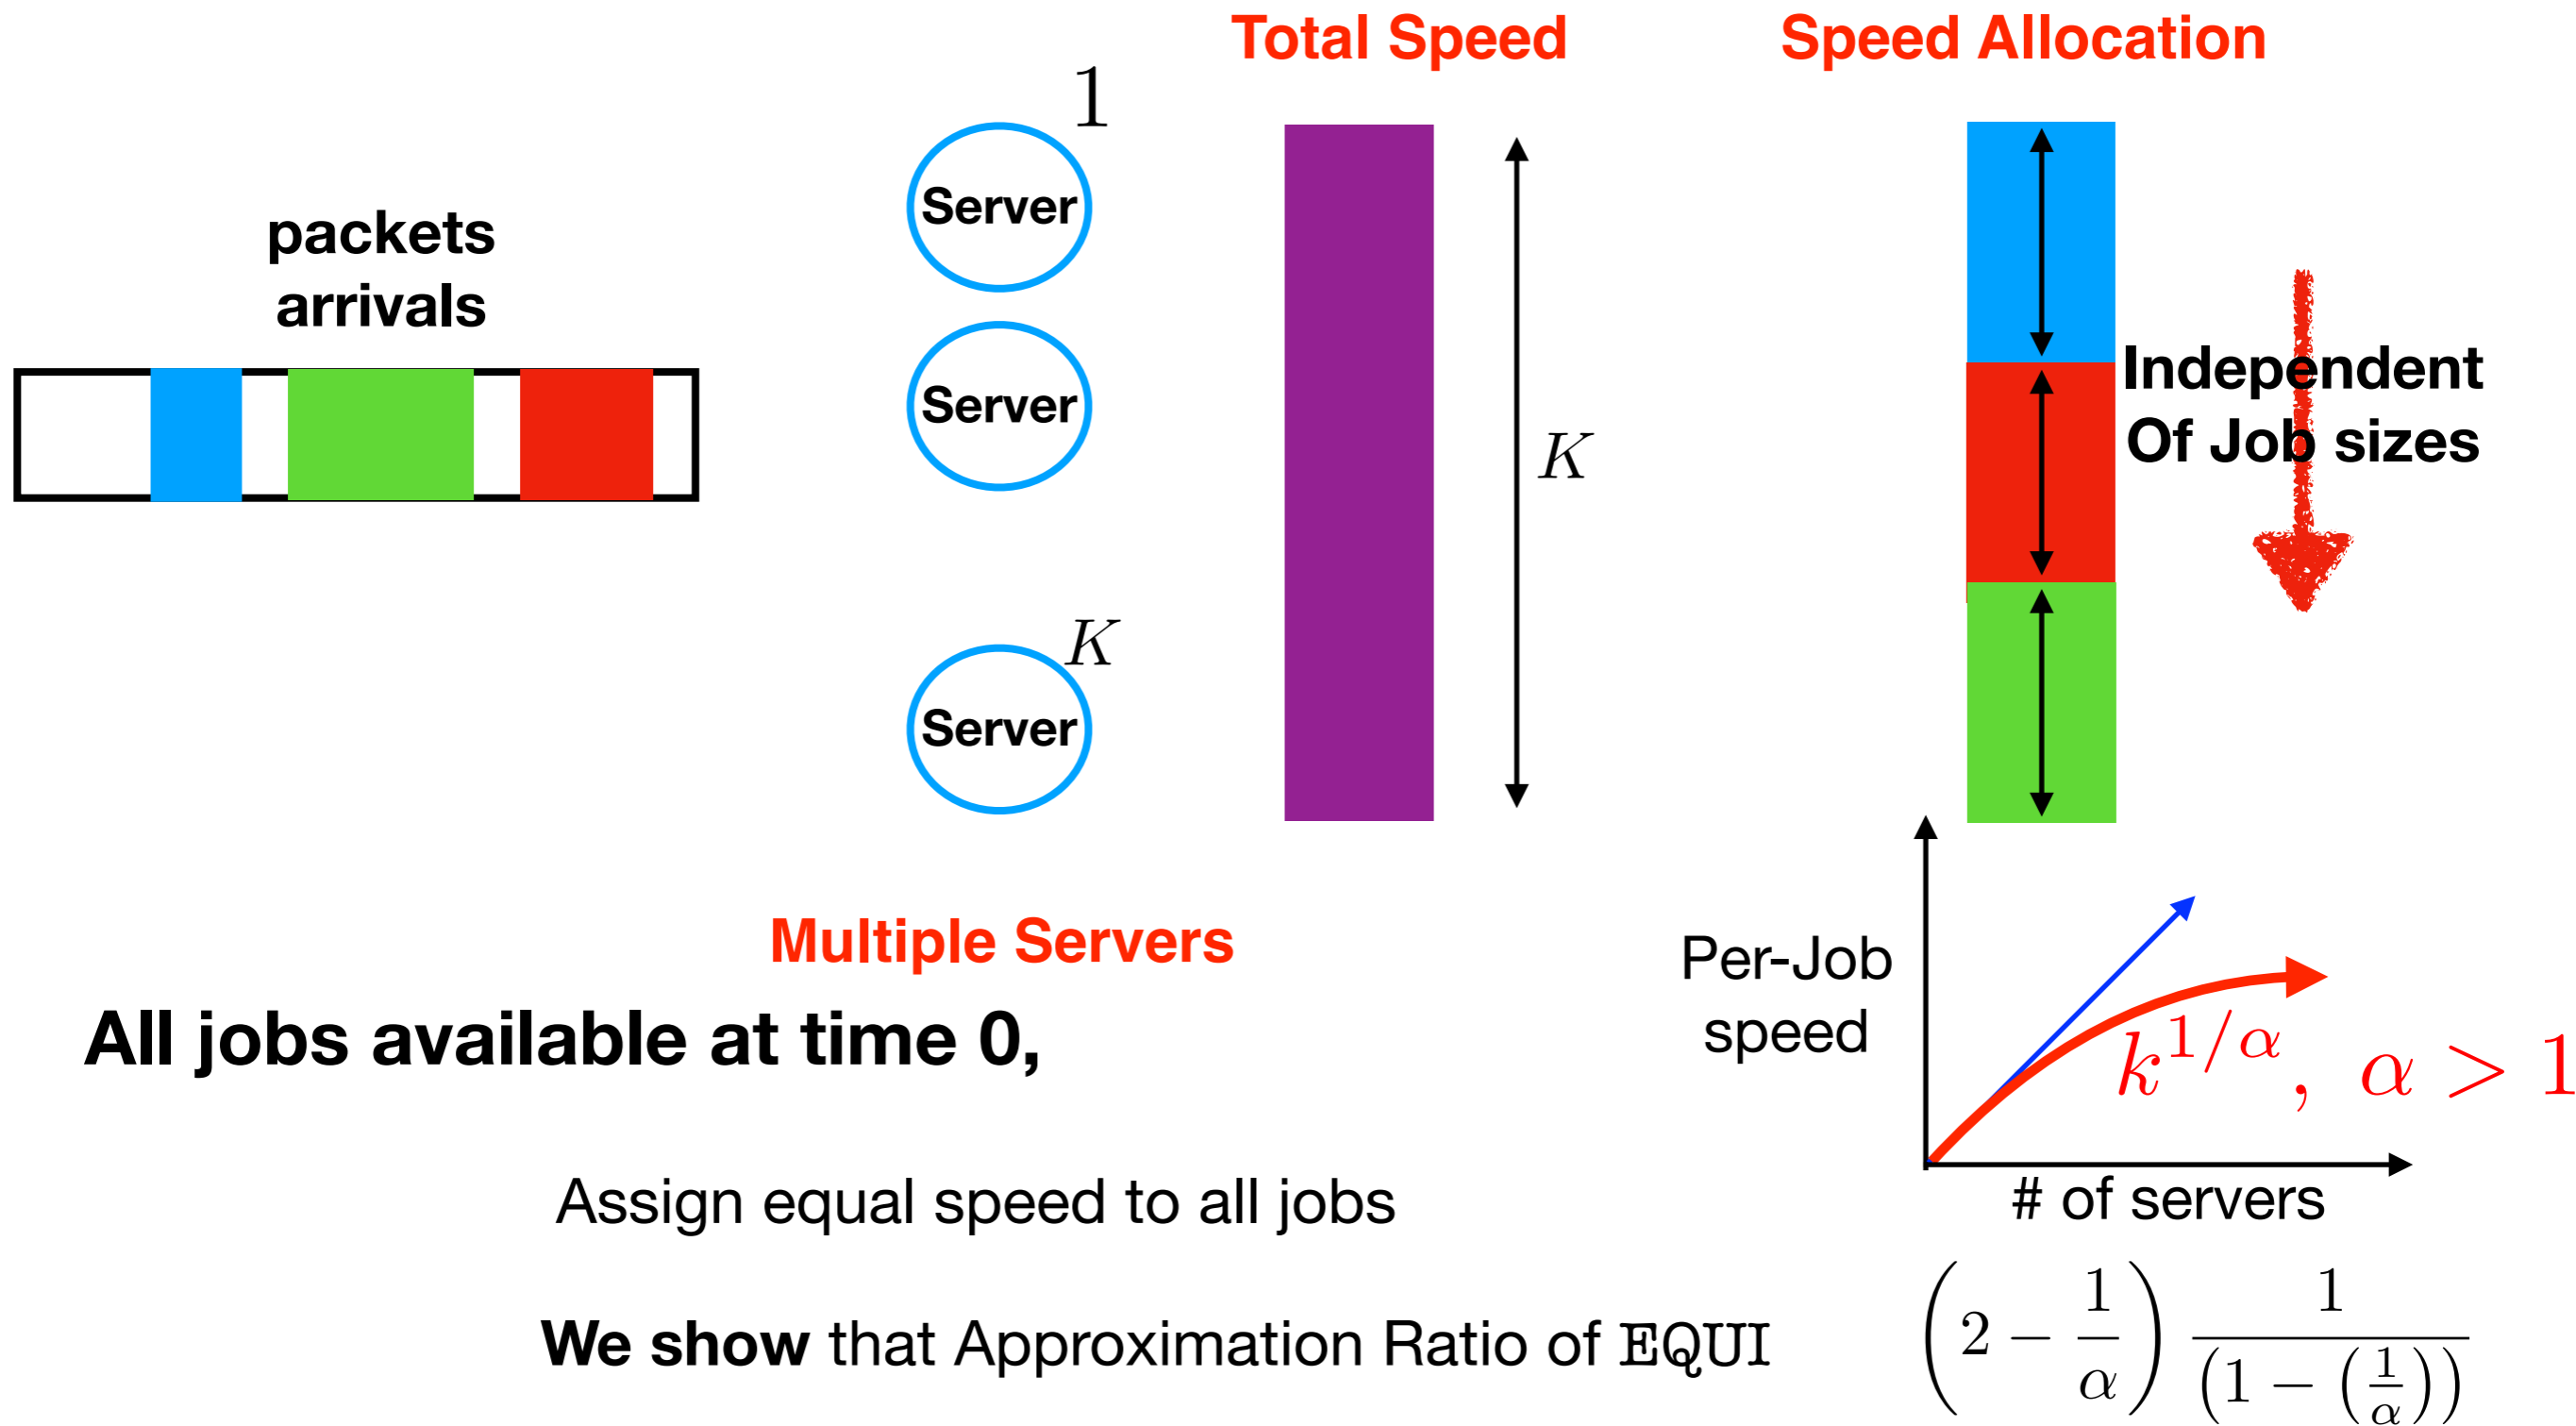


**Multiple Servers**


**All jobs available at time 0,**

Assign equal speed to all jobs

# Simple Algorithm $\mathbb{E}QUI$



# Simple Algorithm $\mathbb{E}QUI$



Multiple Servers

All jobs available at time 0,

Assign equal speed to all jobs

We show that Approximation Ratio of  $\mathbb{E}QUI$



$\mathbb{E}QUI$  is simpler than  $heSRPT$  to implement

$$\left(2 - \frac{1}{\alpha}\right) \frac{1}{\left(1 - \left(\frac{1}{\alpha}\right)\right)}$$

# Real Problem: Online

# Real Problem: Online



packets  
arrivals



**Multiple Servers**

# Real Problem: Online

packets  
arrivals



**Multiple Servers**

**Jobs arrive over time, arbitrary time and sizes**

# Metric for Online Algorithms

# Metric for Online Algorithms

**Competitive ratio**



# Metric for Online Algorithms

**Competitive ratio** ratio of the cost of an **online** and the **offline Opt** algorithm

$$r_{\text{ON}} = \max_{\sigma} \frac{v_{\text{ON}}(\sigma)}{v_{\text{OPT}}(\sigma)}$$

**Worst Case Input**

# Metric for Online Algorithms


**Competitive ratio** ratio of the cost of an **online** and the **offline Opt** algorithm

$$r_{\text{ON}} = \max_{\sigma} \frac{v_{\text{ON}}(\sigma)}{v_{\text{OPT}}(\sigma)}$$

**Worst Case Input**

**Goal** online algorithm with least **CR**


# Prior Work : Online Case



**Multiple Servers**

**Jobs arrive over time, arbitrary time and sizes**

# Prior Work : Online Case




**Multiple Servers**

**Jobs arrive over time, arbitrary time and sizes**

Not much is known, unless resource augmentation is provided

# Prior Work : Online Case




**Multiple Servers**

**Jobs arrive over time, arbitrary time and sizes**

Not much is known, unless resource augmentation is provided

With resource augmentation an algorithm has more resources than OPT

# Prior Work : Online Case

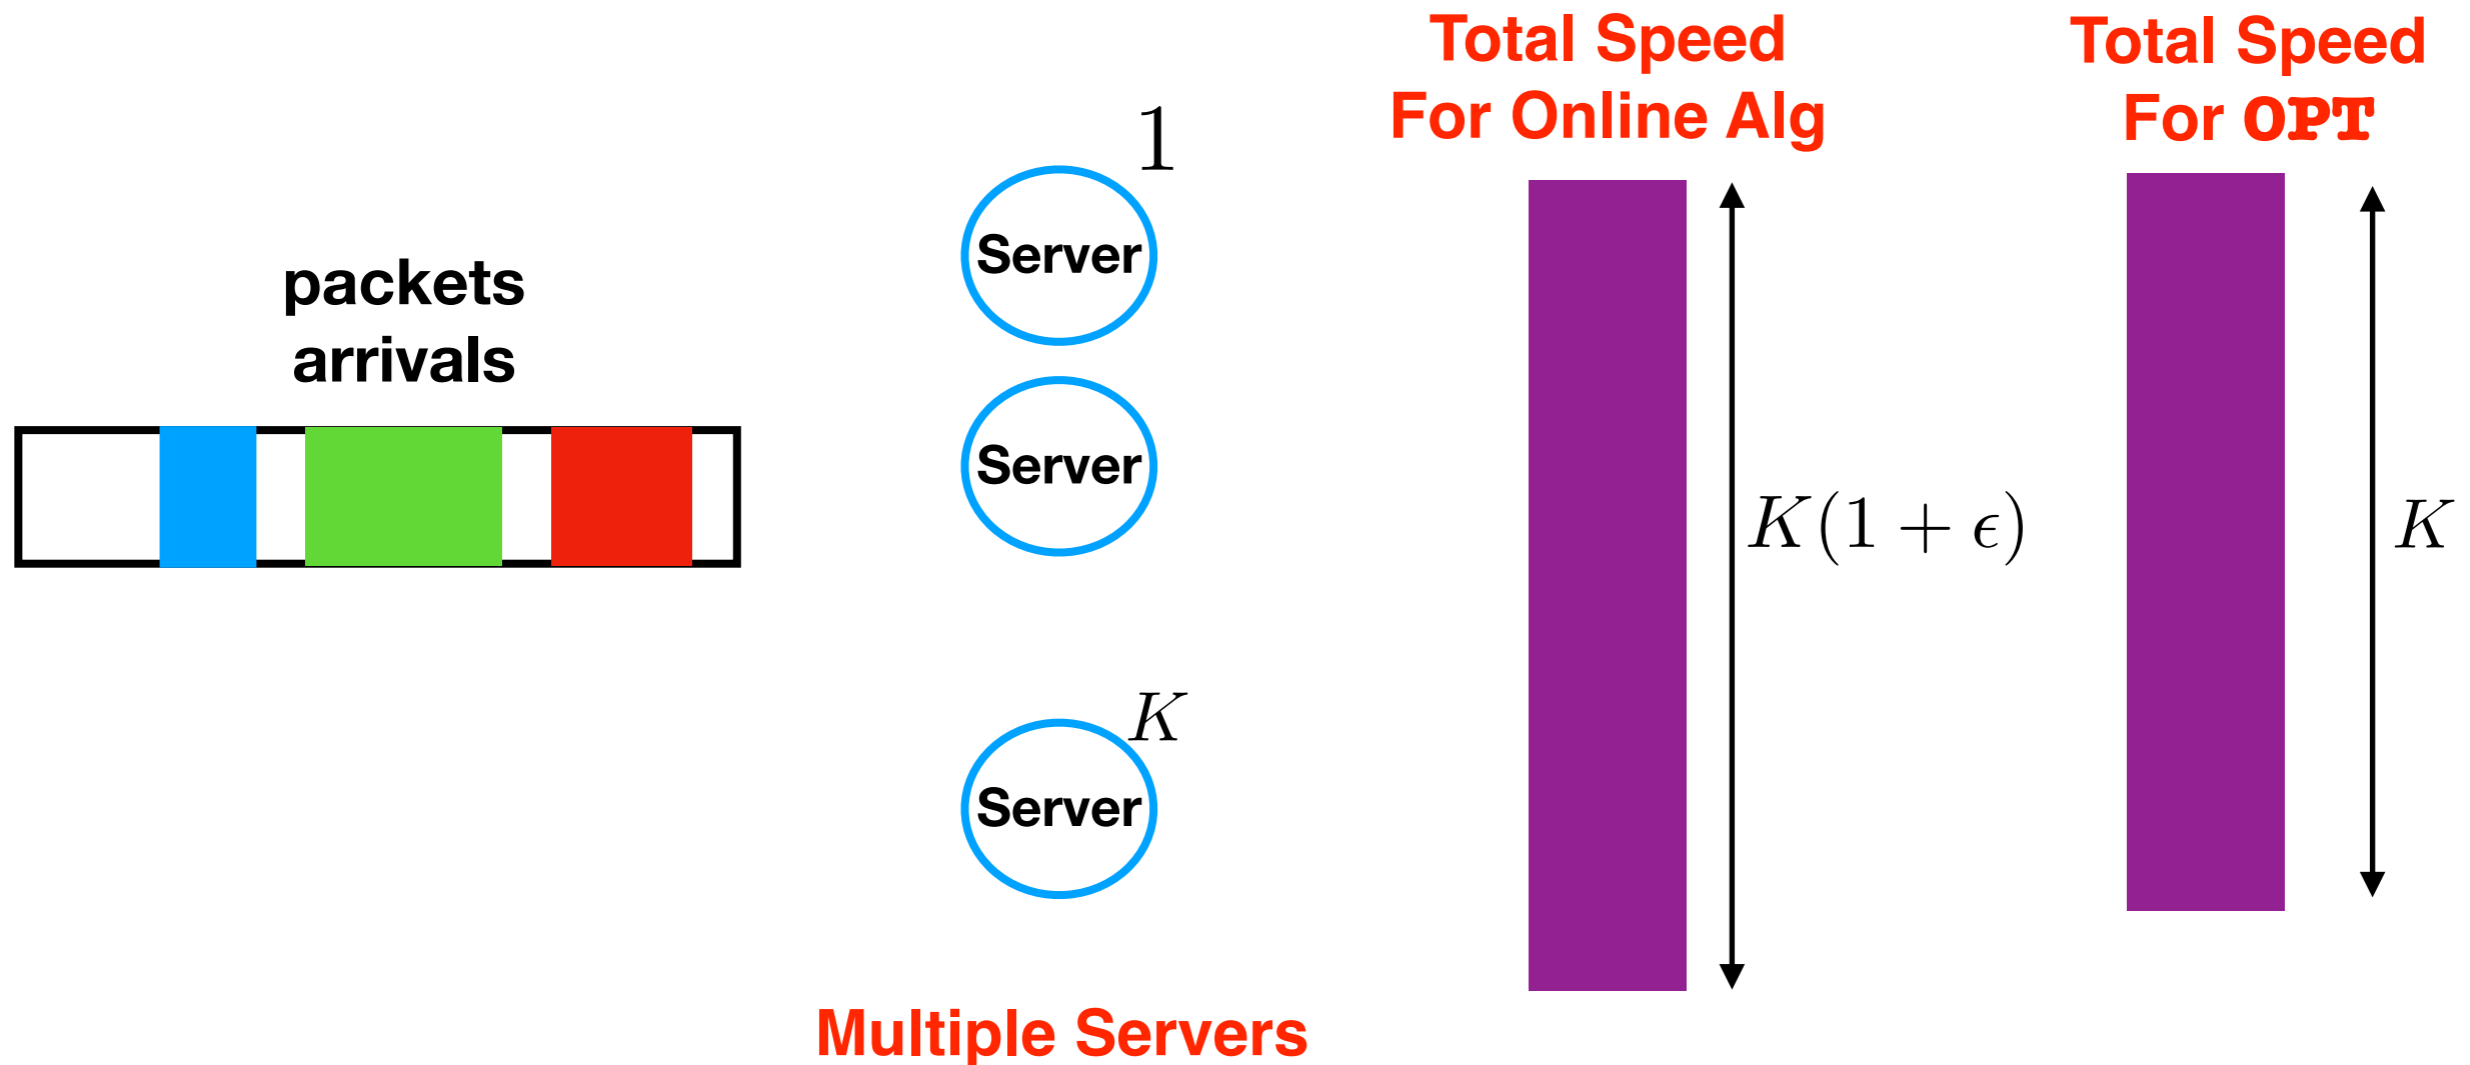


**Jobs arrive over time, arbitrary time and sizes**

Not much is known, unless resource augmentation is provided

With resource augmentation an algorithm has more resources than OPT

# Prior Work : Online Case

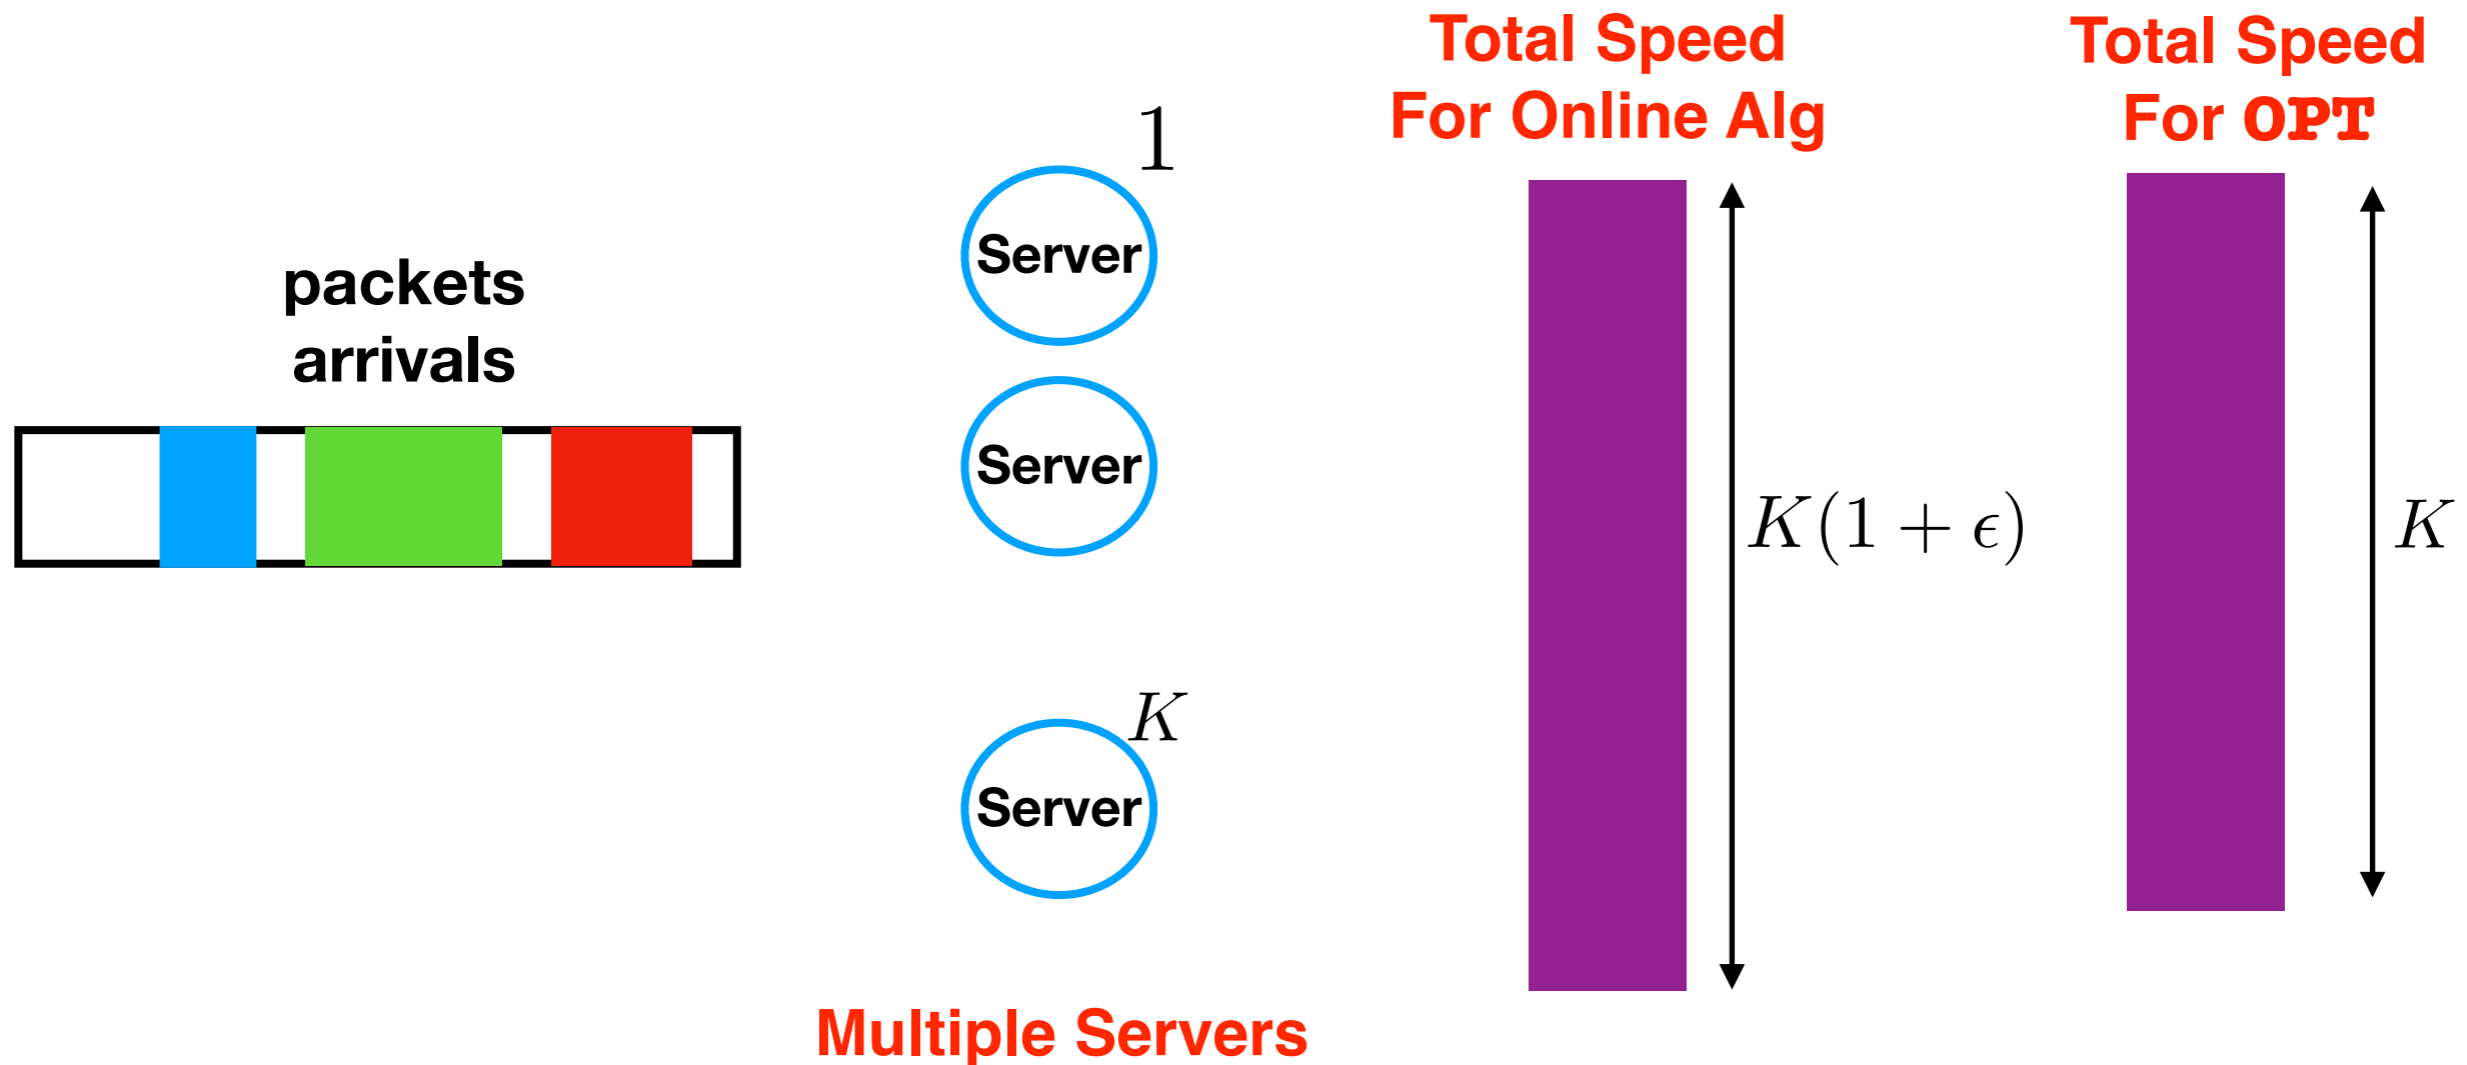


**Jobs arrive over time, arbitrary time and sizes**

Not much is known, unless resource augmentation is provided

With resource augmentation an algorithm has more resources than OPT

# Prior Work : Online Case



**Jobs arrive over time, arbitrary time and sizes**

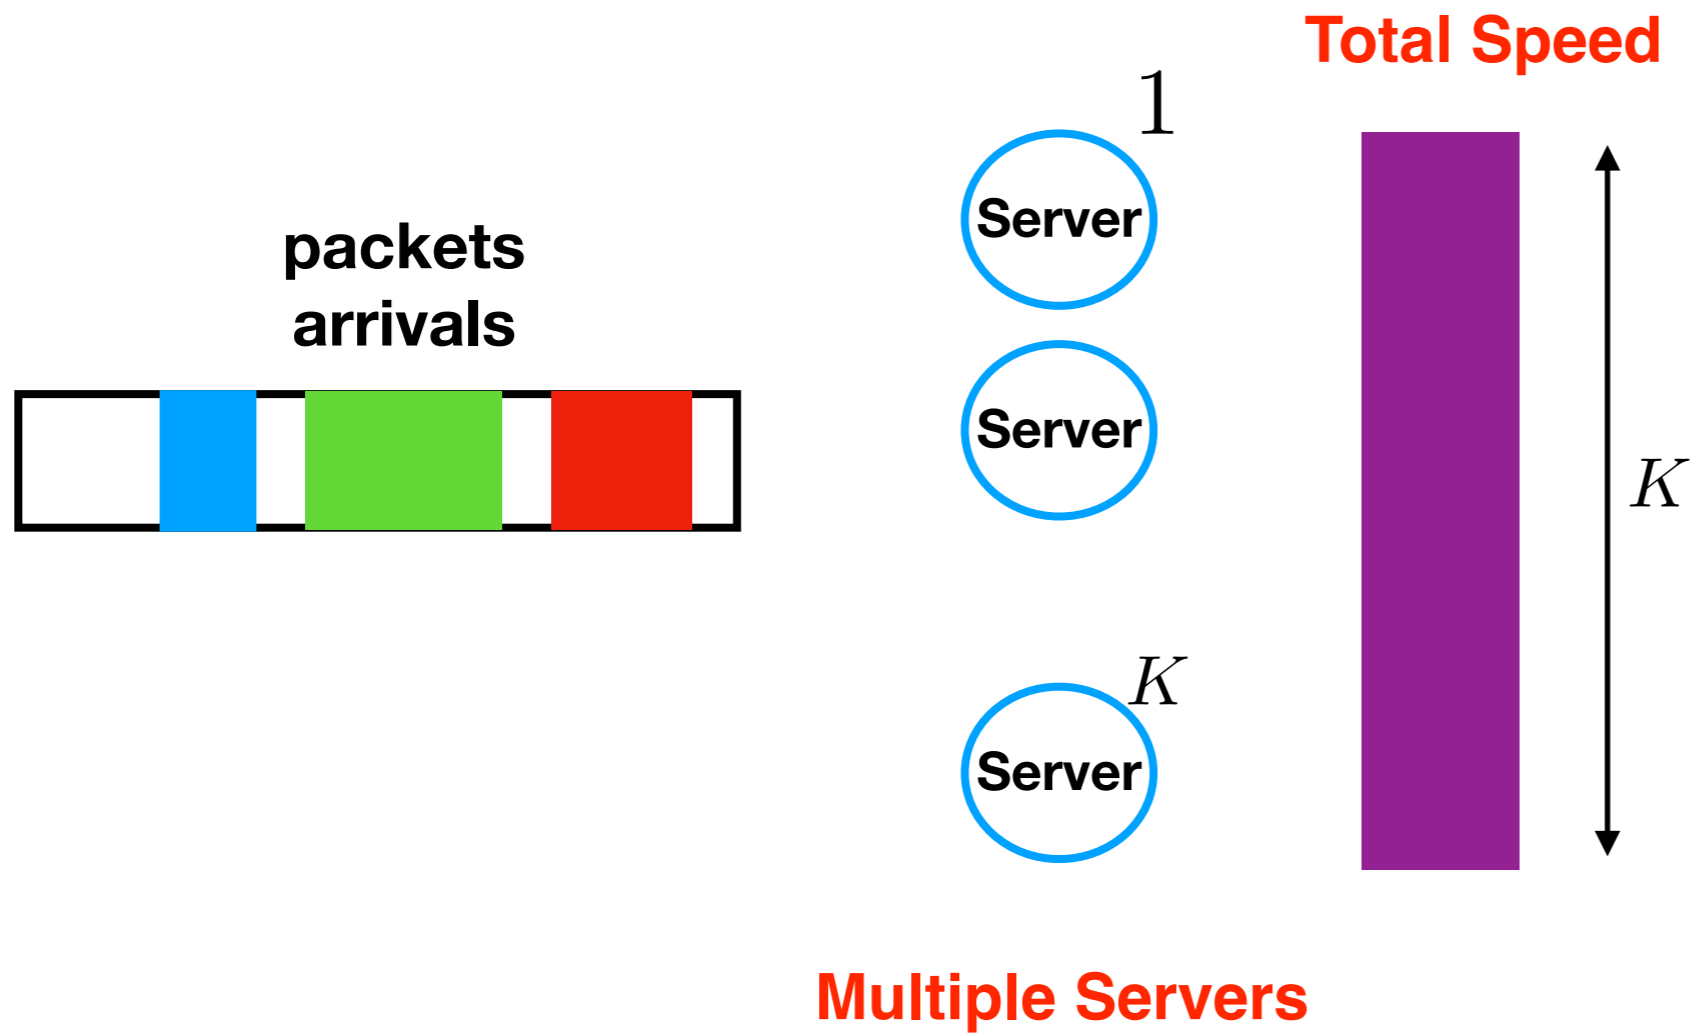
Not much is known, unless resource augmentation is provided

With resource augmentation an algorithm has more resources than OPT

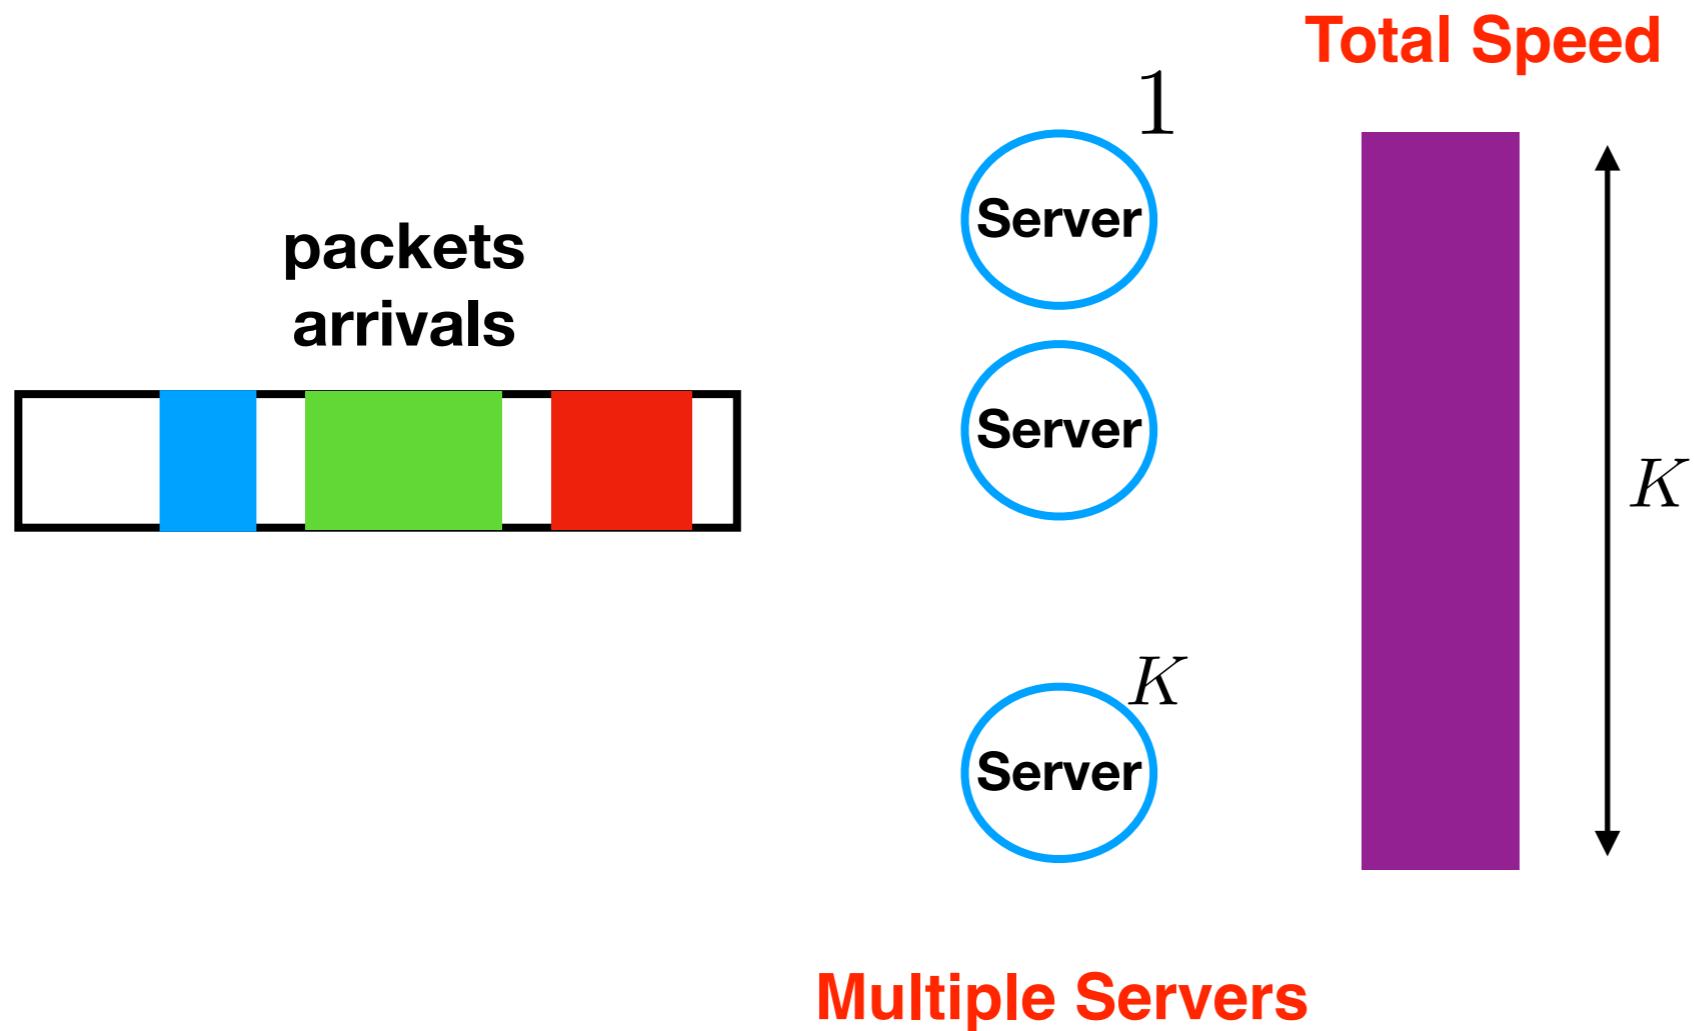
With resource augmentation constant competitive algorithms are known [Edmonds et al]




# Our Algorithm



# Our Algorithm LCFS-EQUI

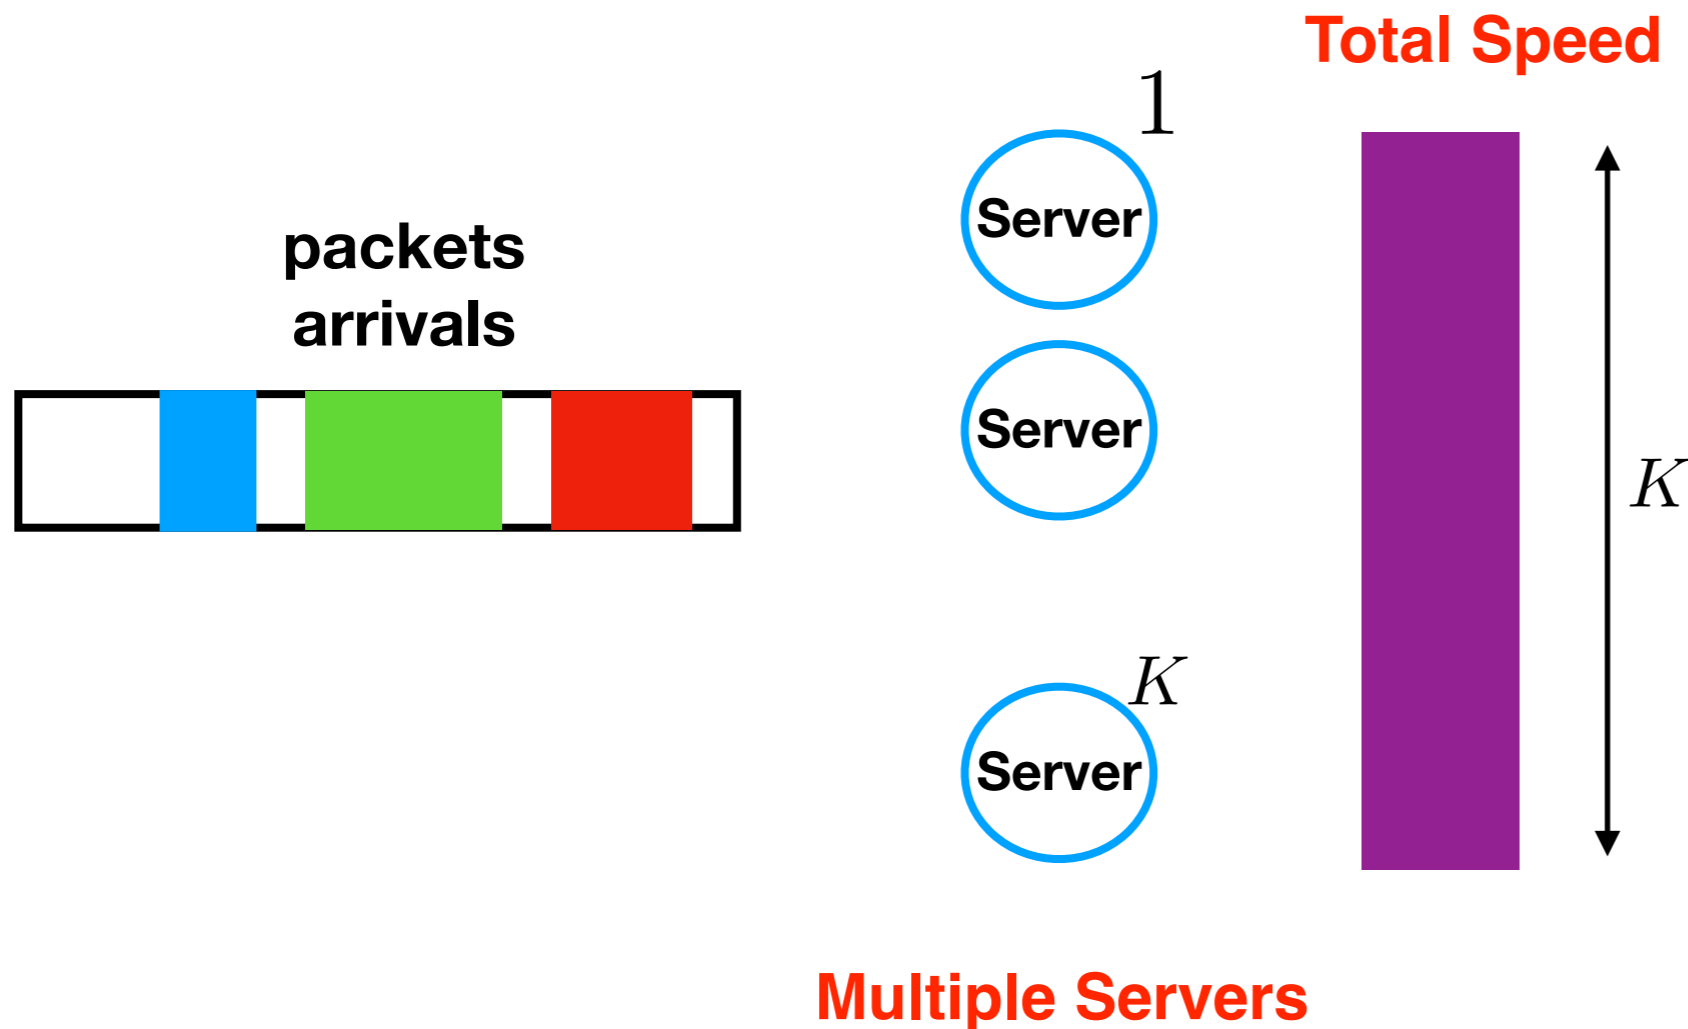


# Our Algorithm LCFS-EQUI



Process the  $\beta$ -fraction of the **most recently arrived** jobs

# Our Algorithm LCFS-EQUI

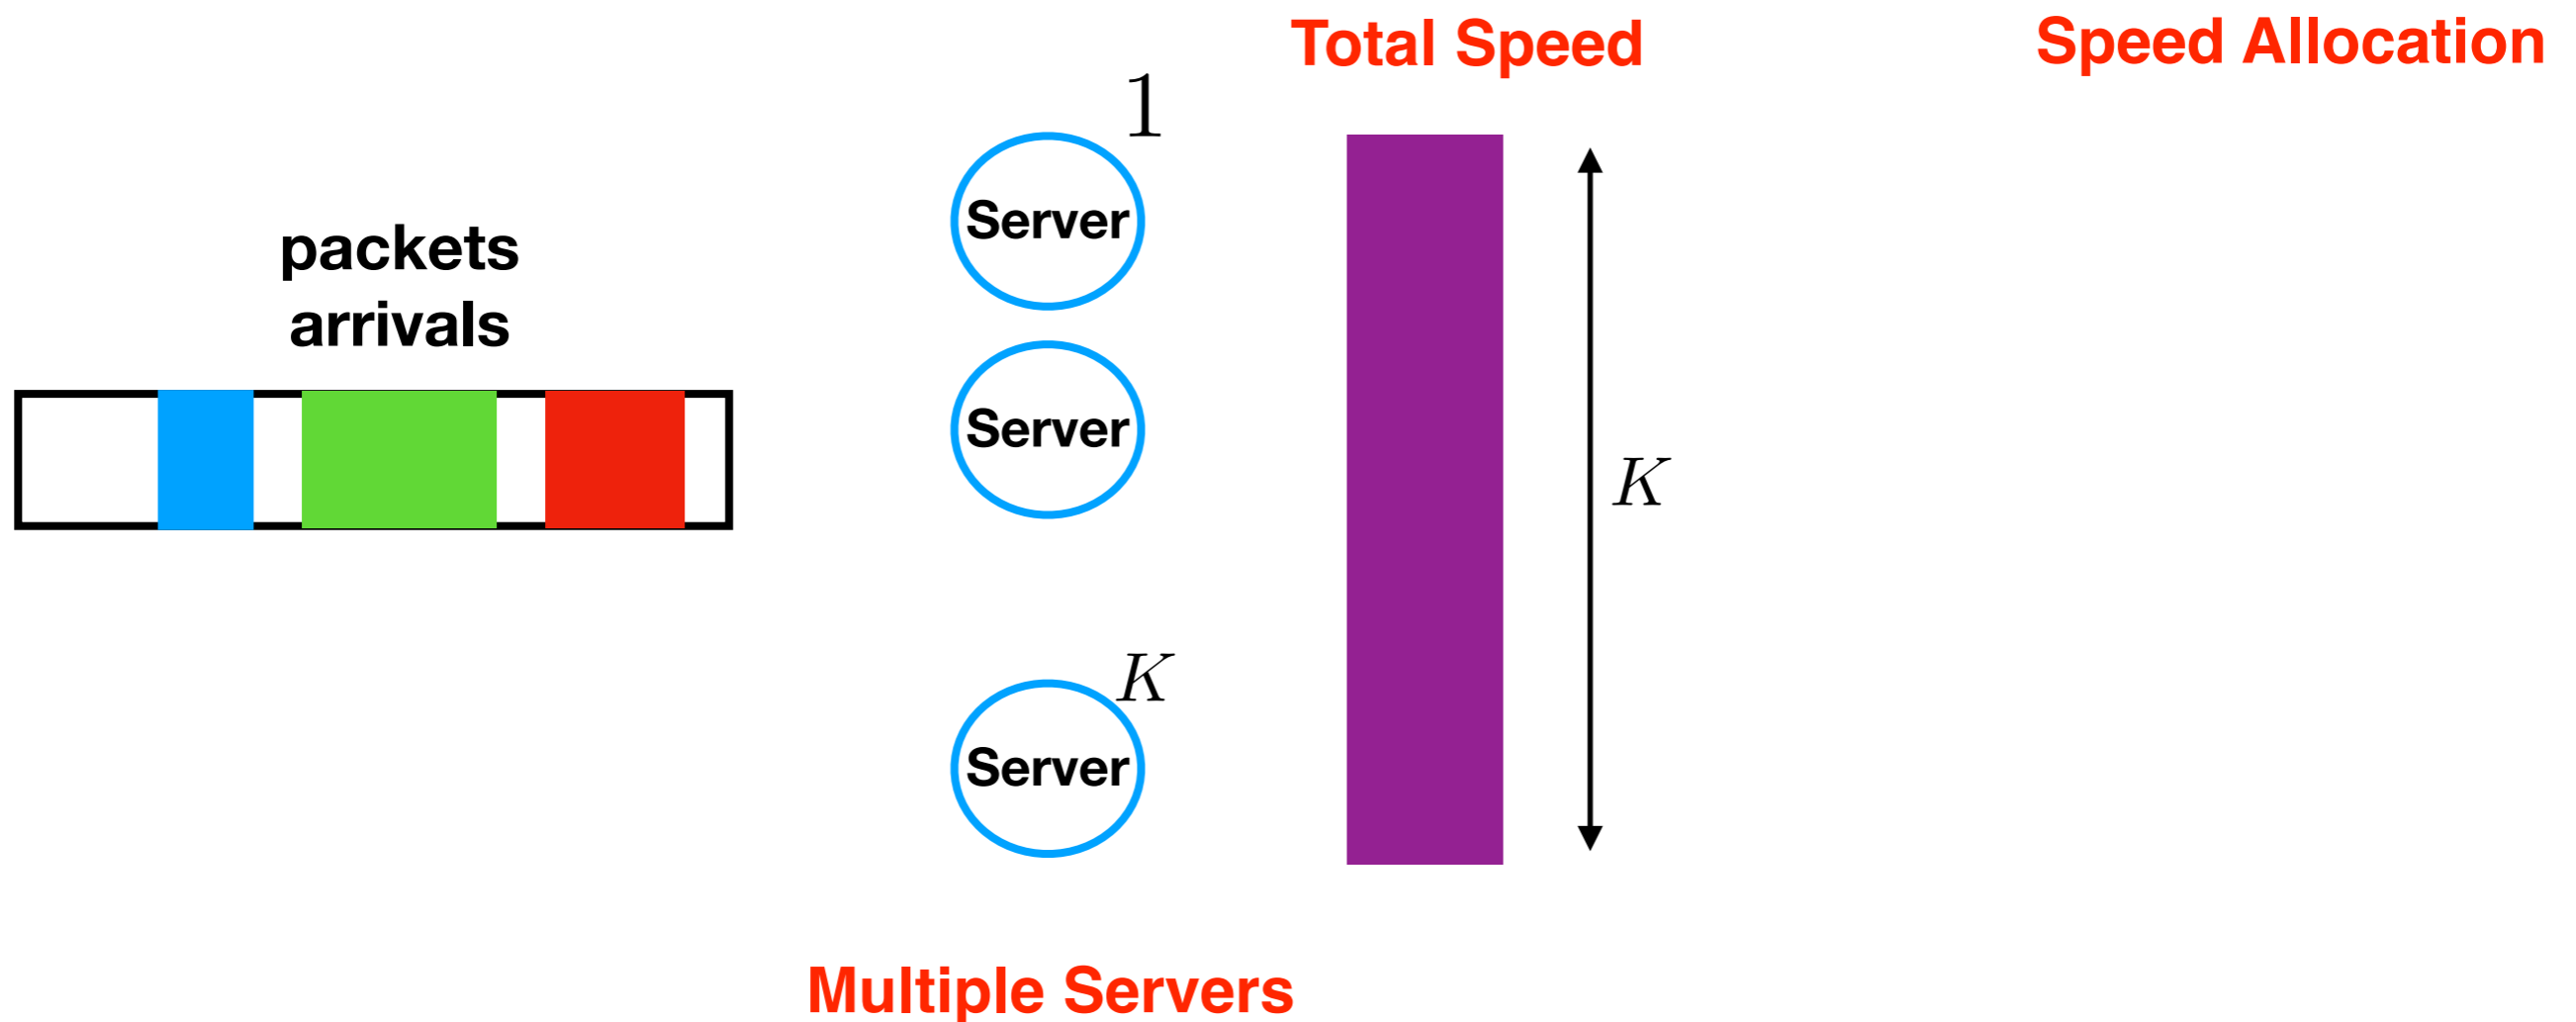


**Multiple Servers**

Process the  $\beta$ -fraction of the **most recently arrived** jobs


Assign **equal speed** to all jobs being processed

# Our Algorithm LCFS-EQUI



Process the  $\beta$ -fraction of the **most recently arrived** jobs  
Assign **equal speed** to all jobs being processed


# Our Algorithm LCFS-EQUI



Process the  $\beta$ -fraction of the **most recently arrived** jobs

Assign **equal speed** to all jobs being processed

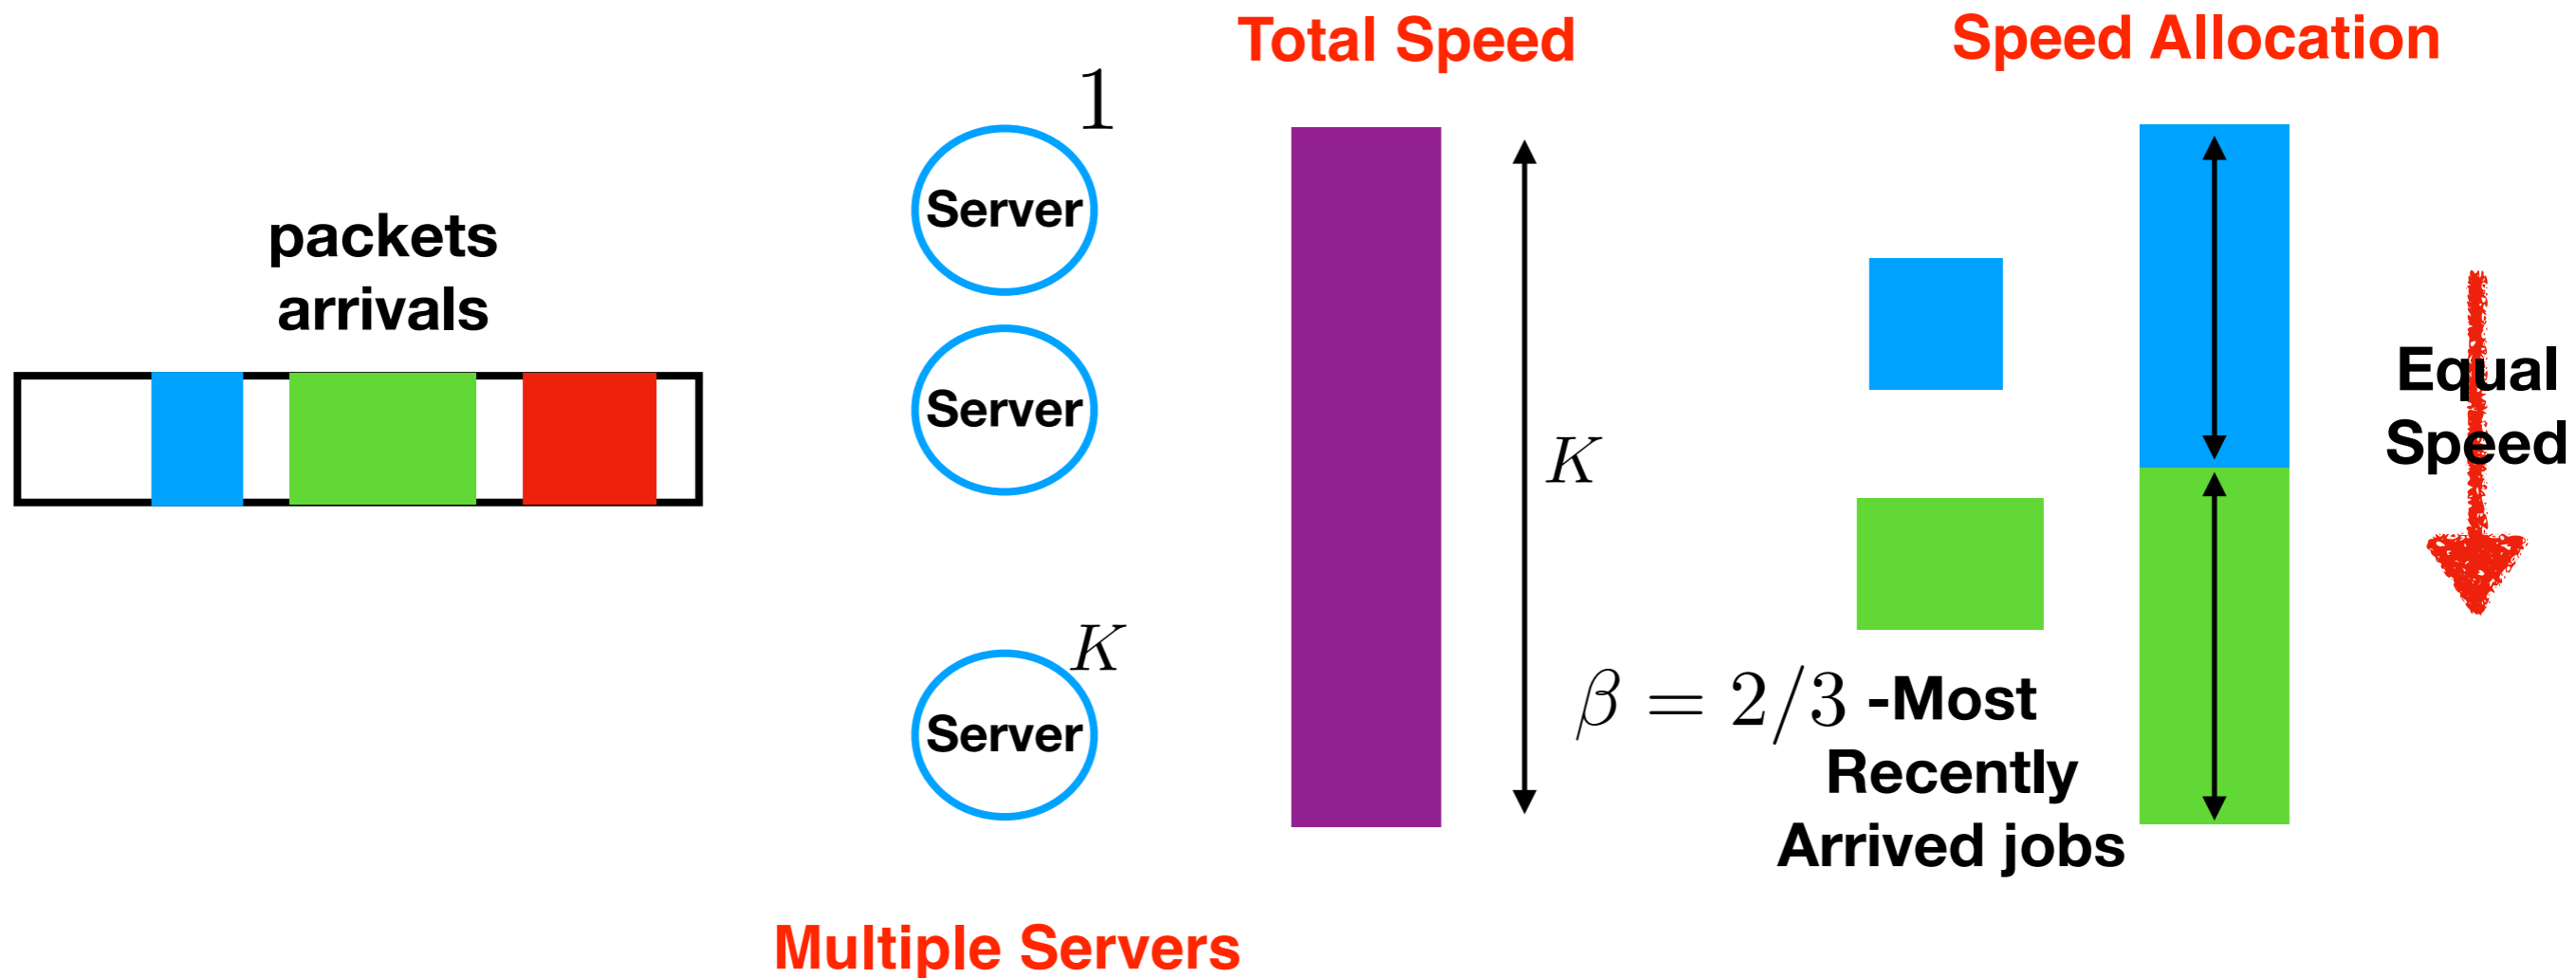
# Our Algorithm LCFS-EQUI



Process the  $\beta$ -fraction of the **most recently arrived** jobs

Assign **equal speed** to all jobs being processed

# Our Algorithm LCFS-EQUI




Process the  $\beta$ -fraction of the **most recently arrived** jobs

Assign **equal speed** to all jobs being processed


**Algorithm is non-clairvoyant : no information needed about remaining job size**



# Intuition LCFS-EQUI

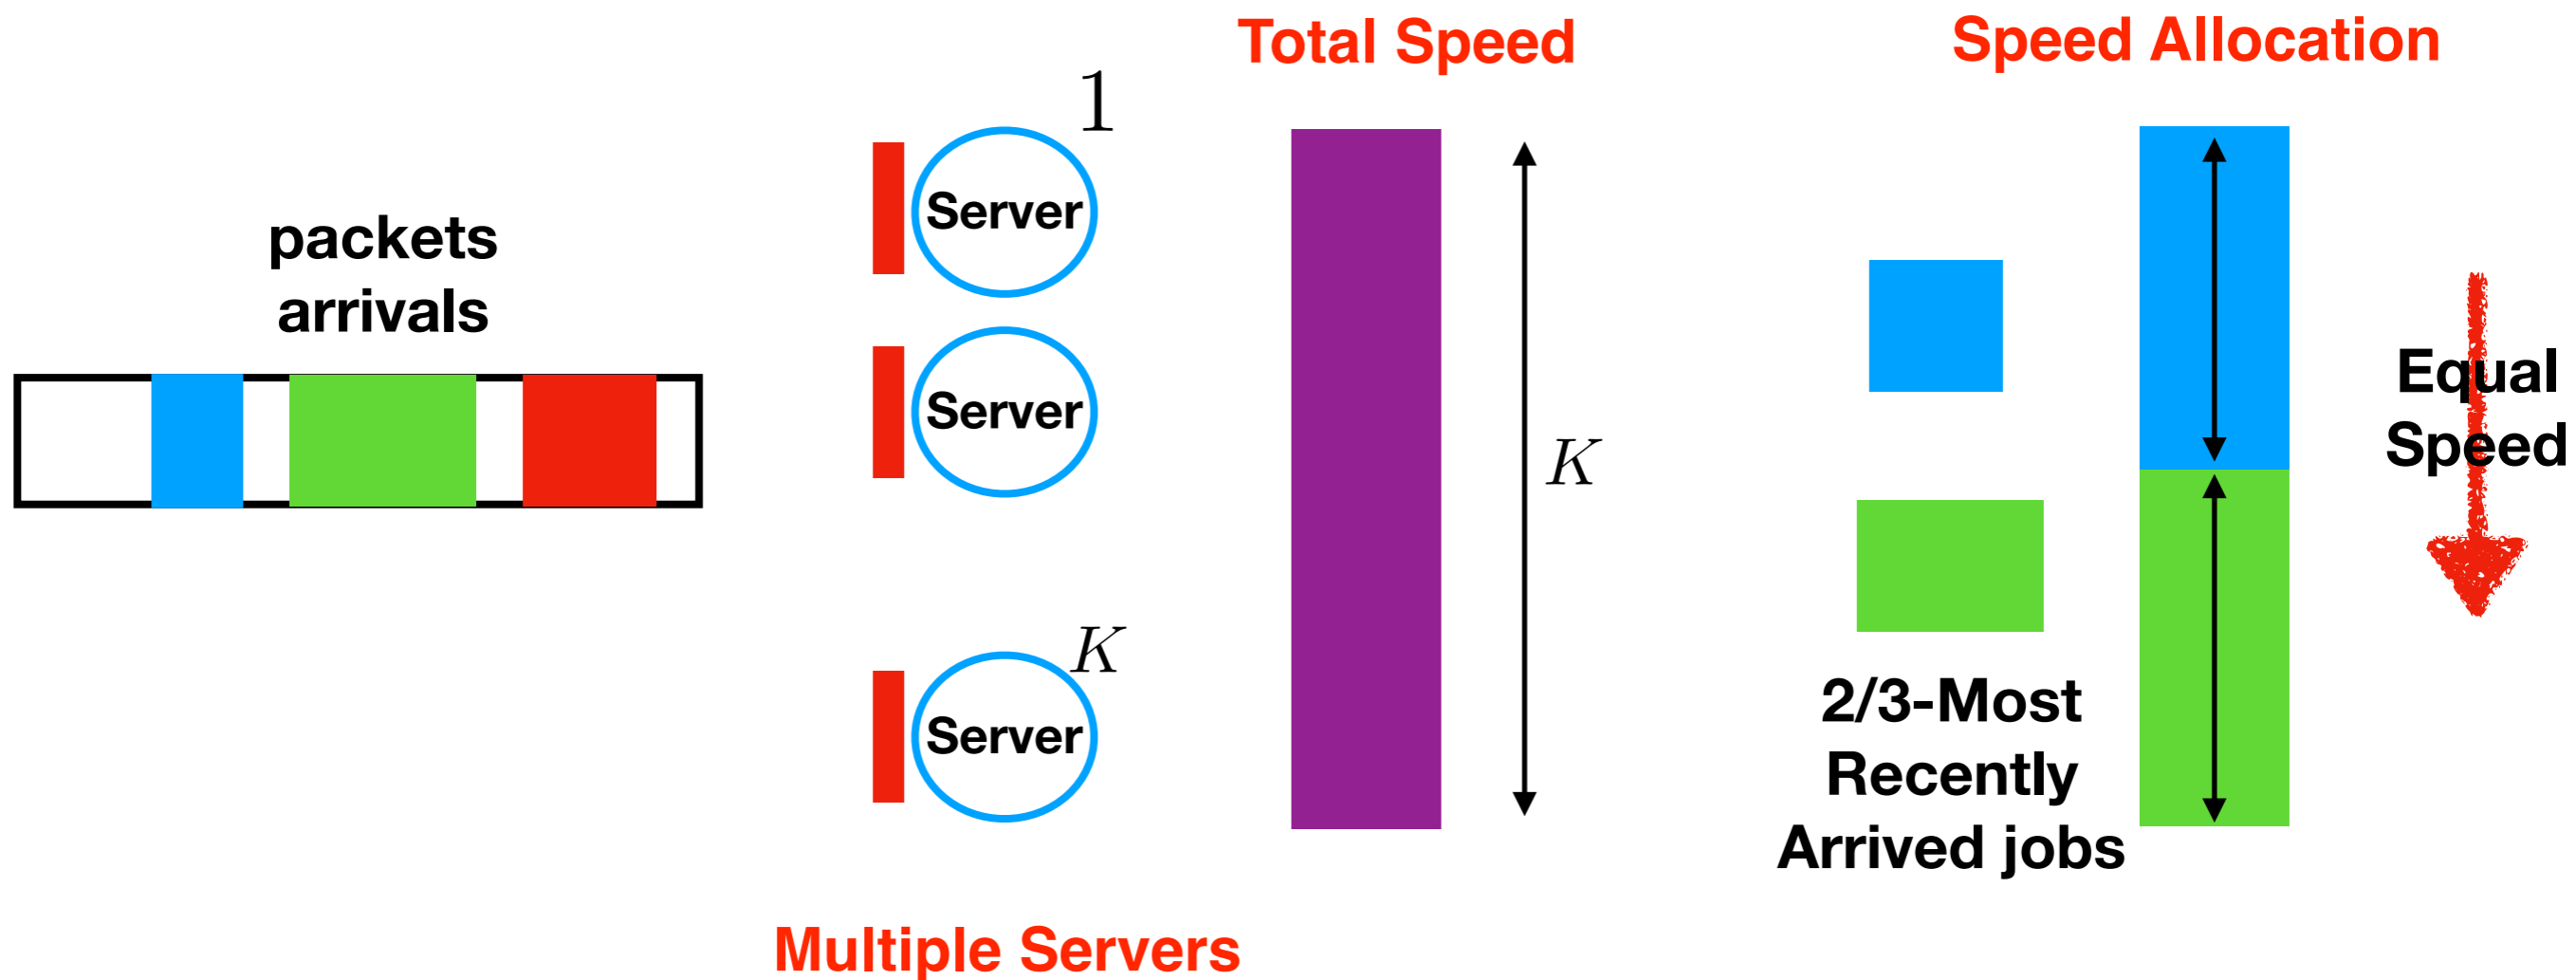


# Intuition LCFS-EQUI



**SRPT- short jobs remain in system for short time**

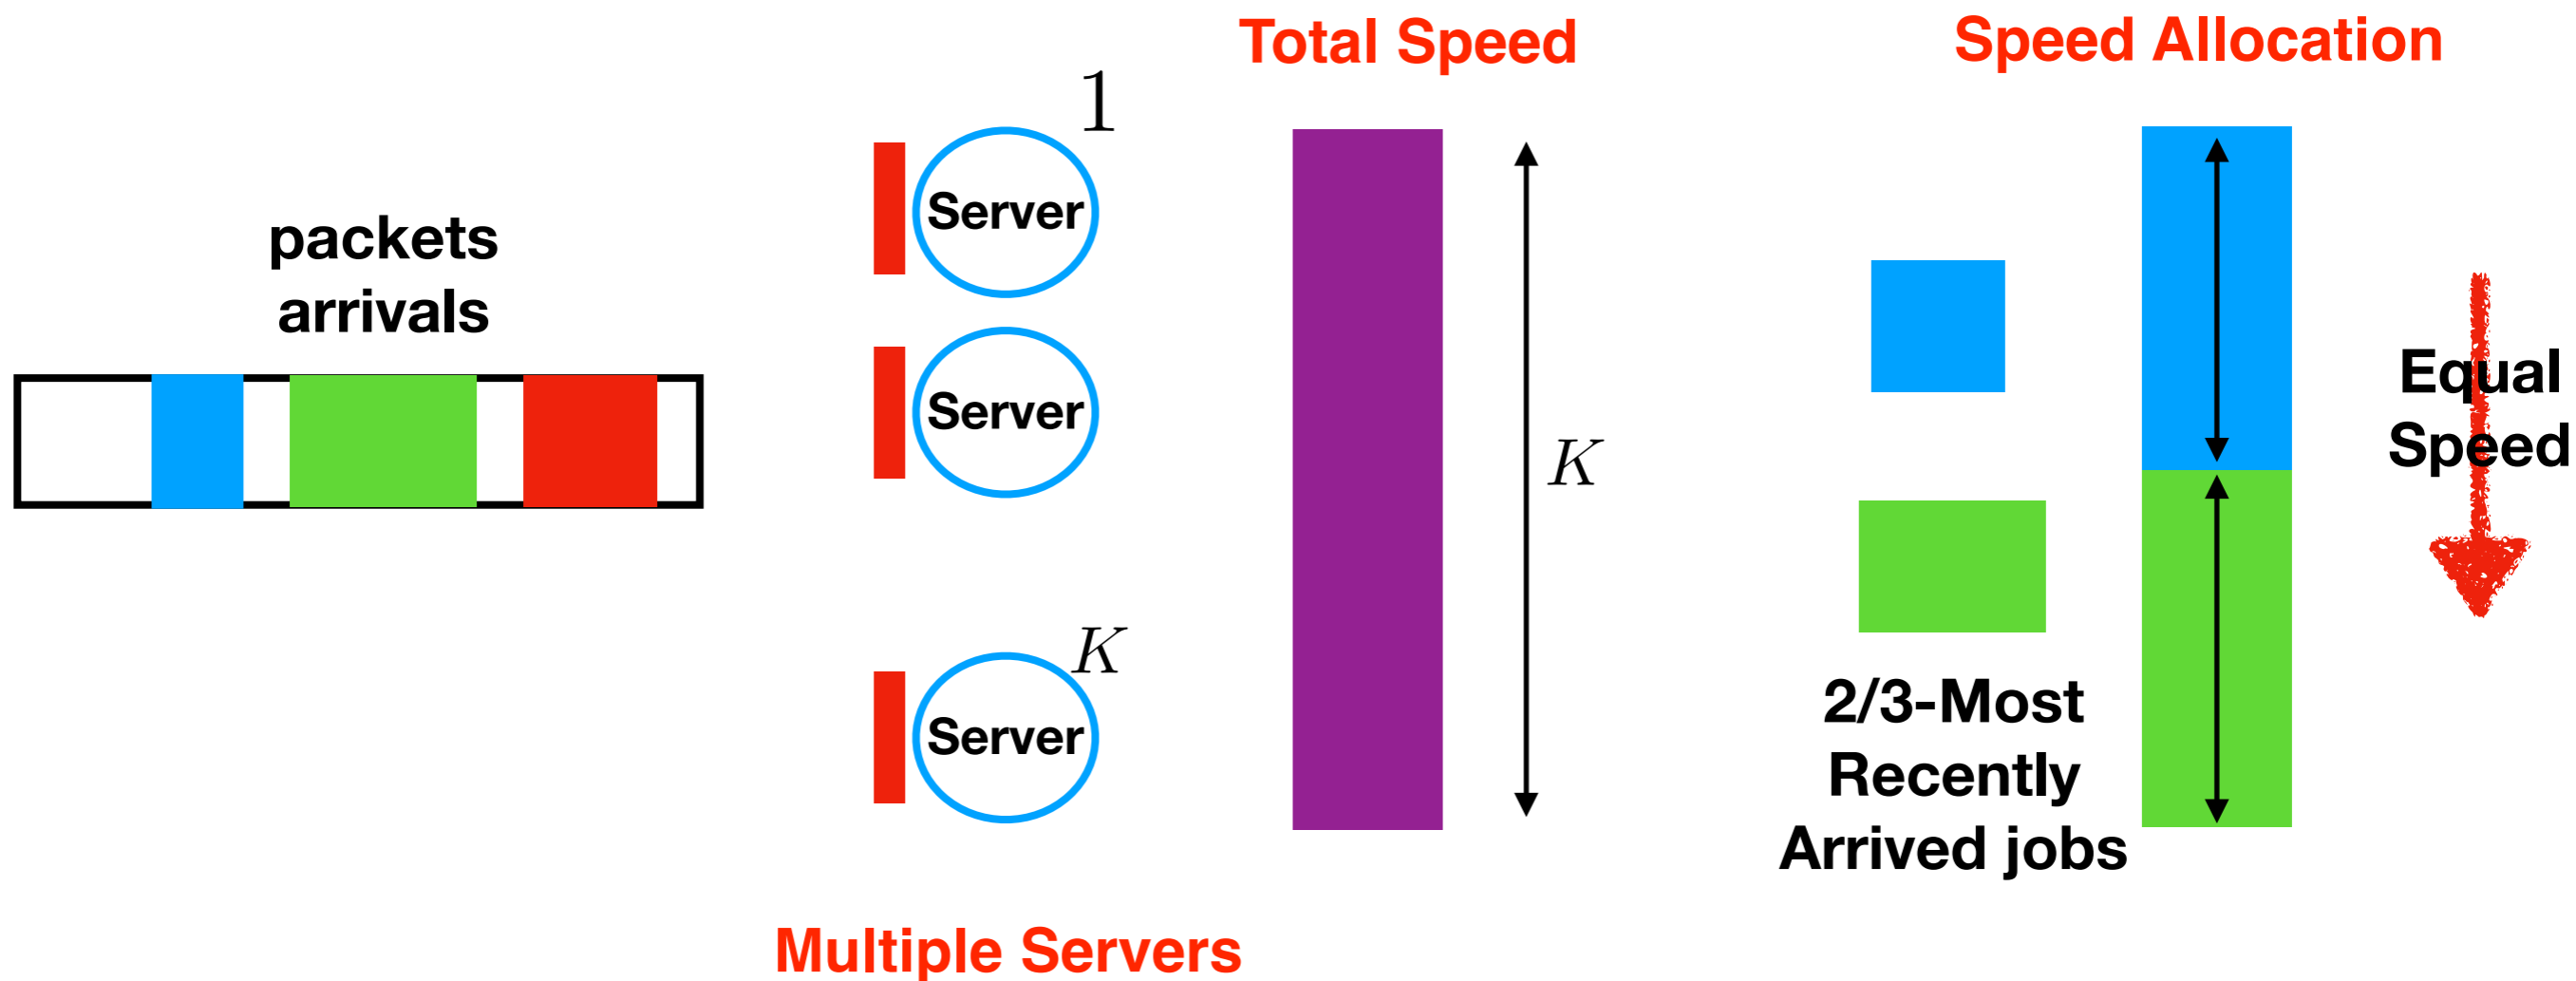
# Intuition LCFS-EQUI



**SRPT-** short jobs remain in system for short time

**LCFS-** w/o job size information, long jobs remain in system for long

# Intuition LCFS-EQUI




**SRPT-** short jobs remain in system for short time

**LCFS-** w/o job size information, long jobs remain in system for long

Technical - construction of potential (Lyapunov) function is easy

# Intuition LCFS-EQUI



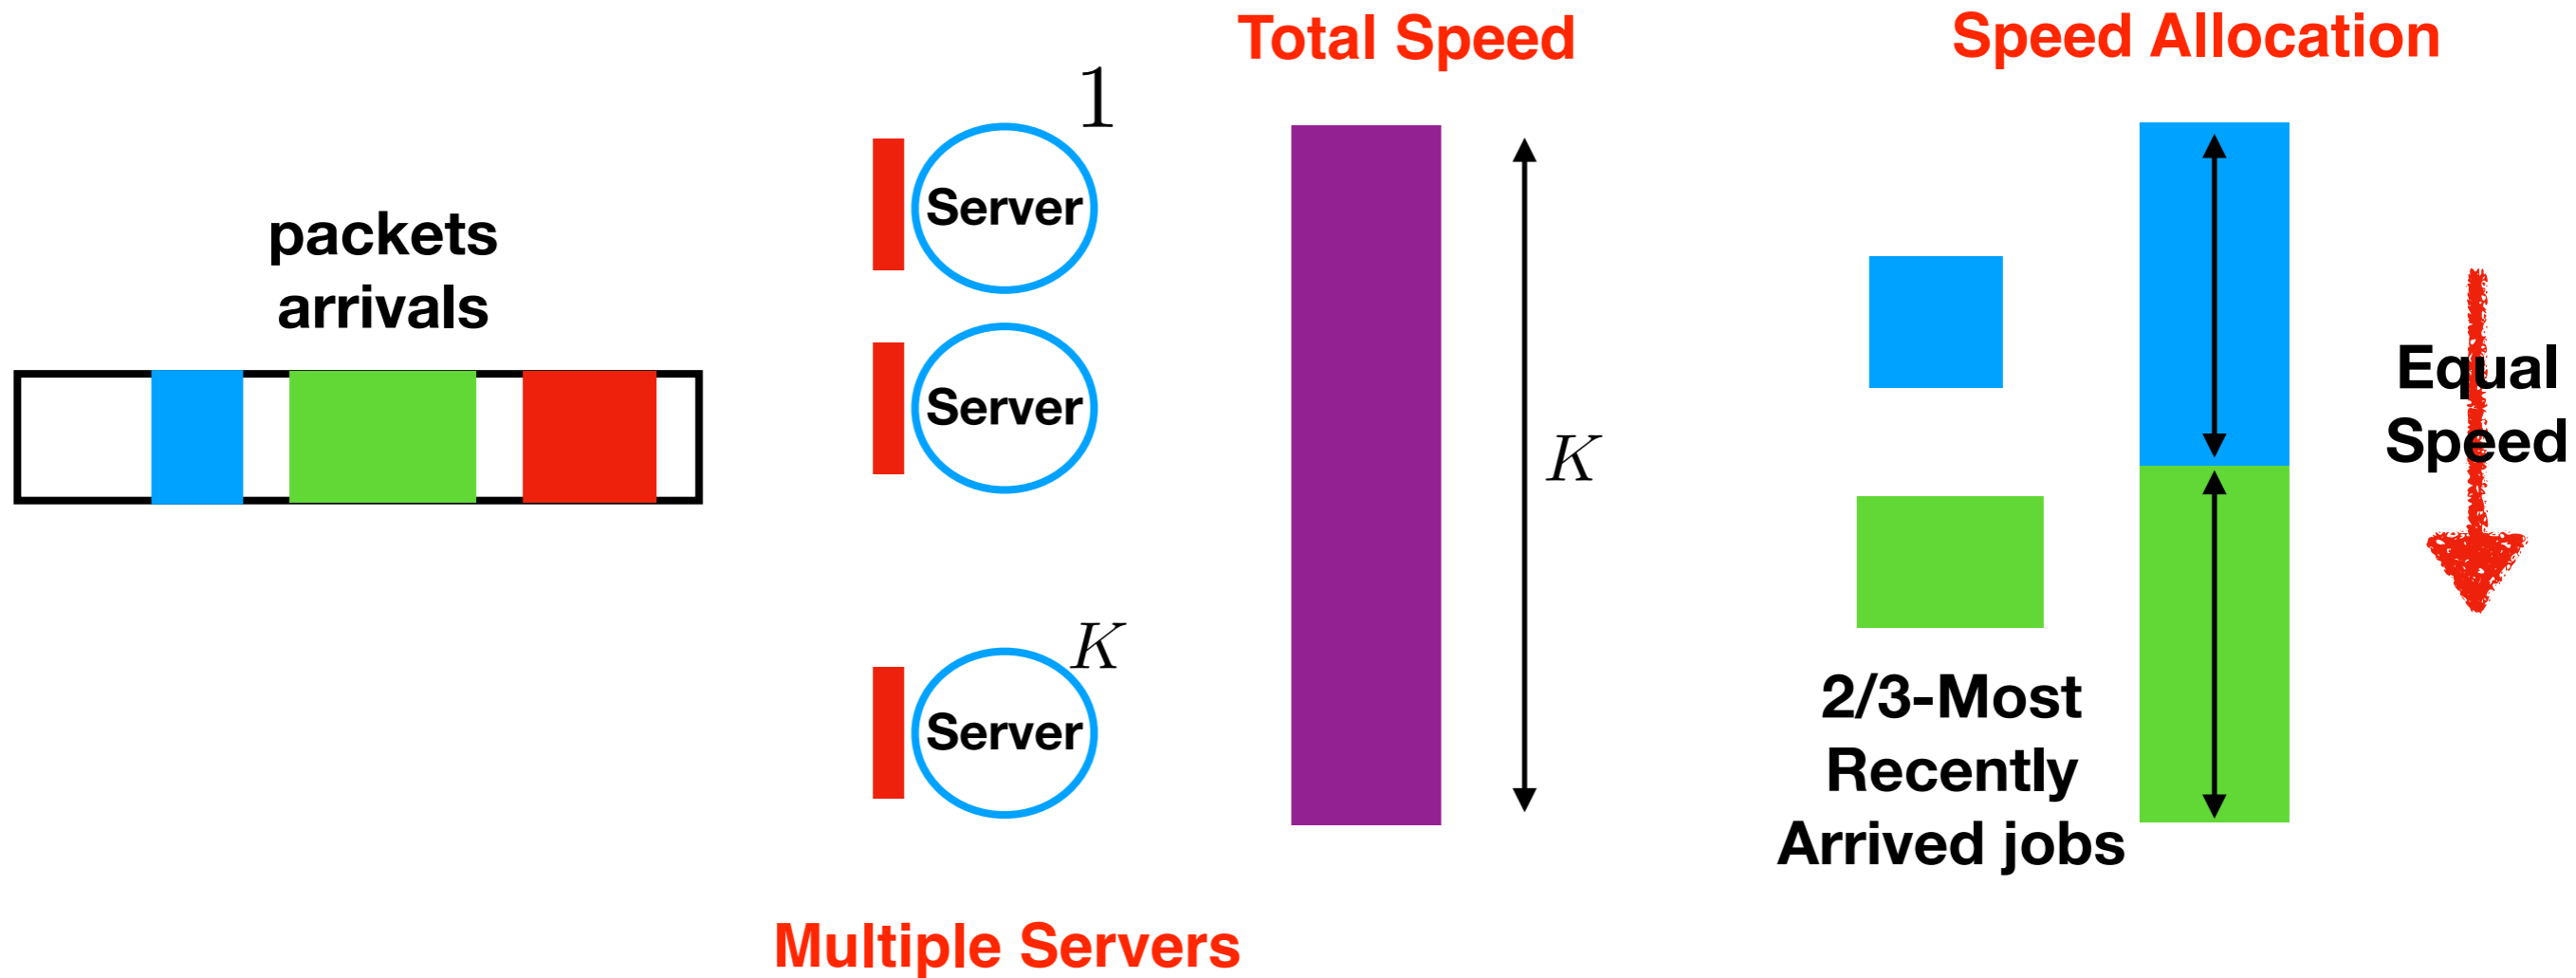
**SRPT-** short jobs remain in system for short time

**LCFS-** w/o job size information, long jobs remain in system for long

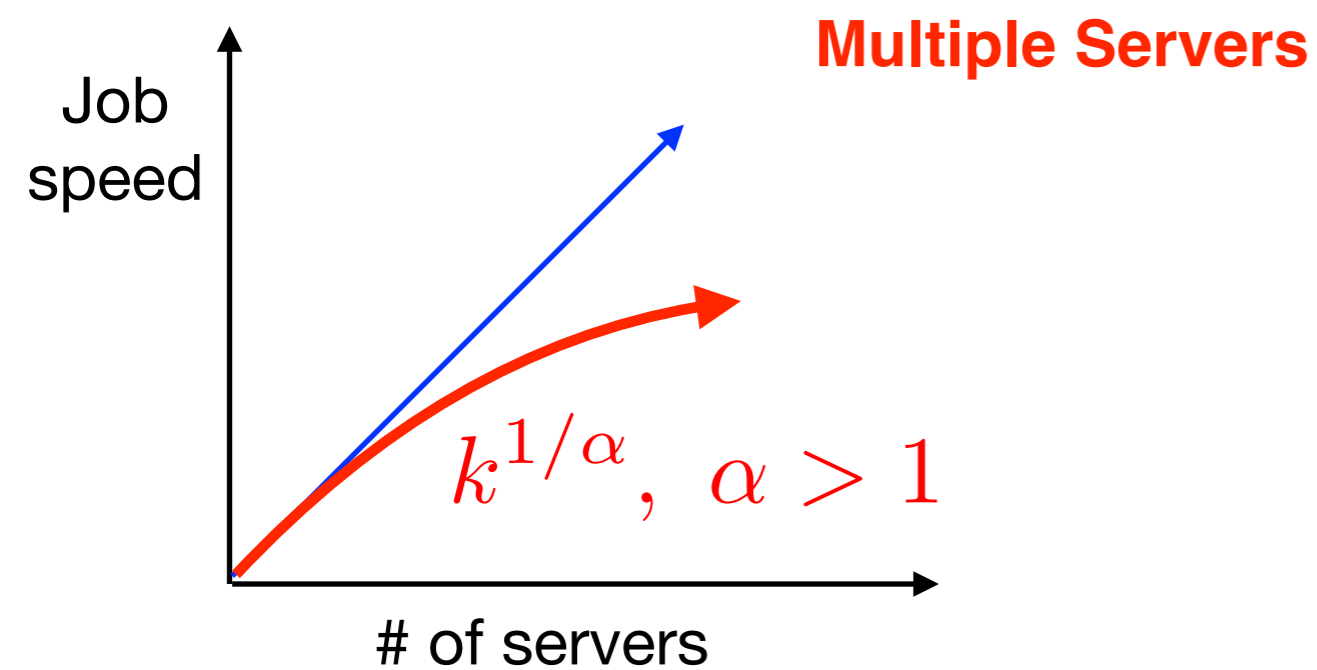

Technical - construction of potential (Lyapunov) function is easy

**EQUI :** with equal speed, analysis is easy


# Guarantee LCFS-EQUI



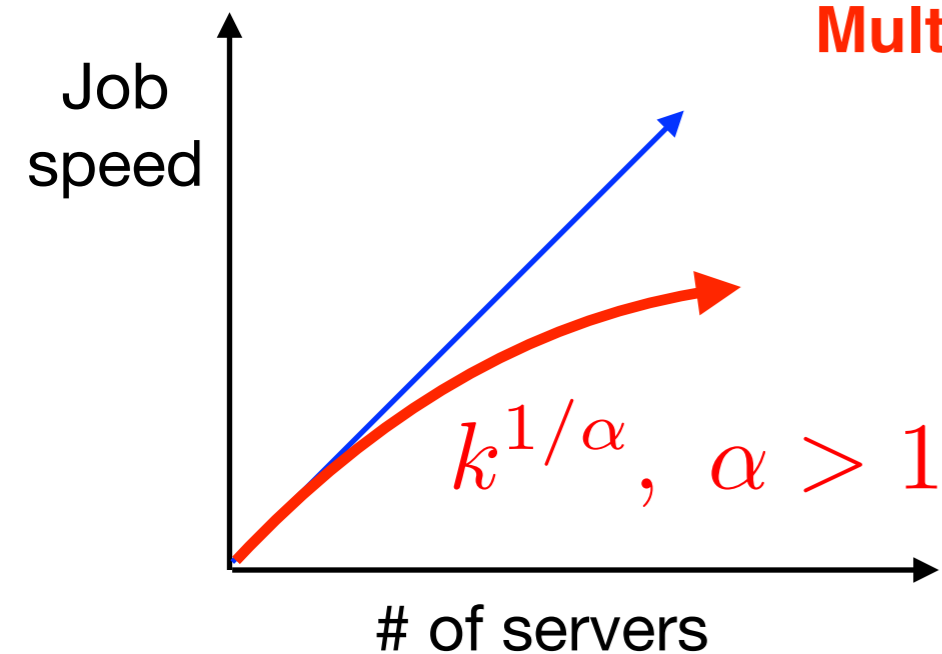
# Guarantee LCFS-EQUI



# Guarantee LCFS-EQUI



## Multiple Servers





For a suitable choice of fraction  $\beta$   
**Competitive ratio is a constant that only depends on exponent  $\alpha$**



# General Problem


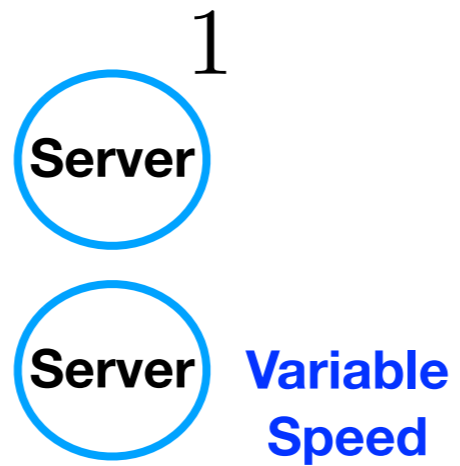
packets  
arrivals



**Unlimited # of Servers**

# General Problem



packets  
arrivals



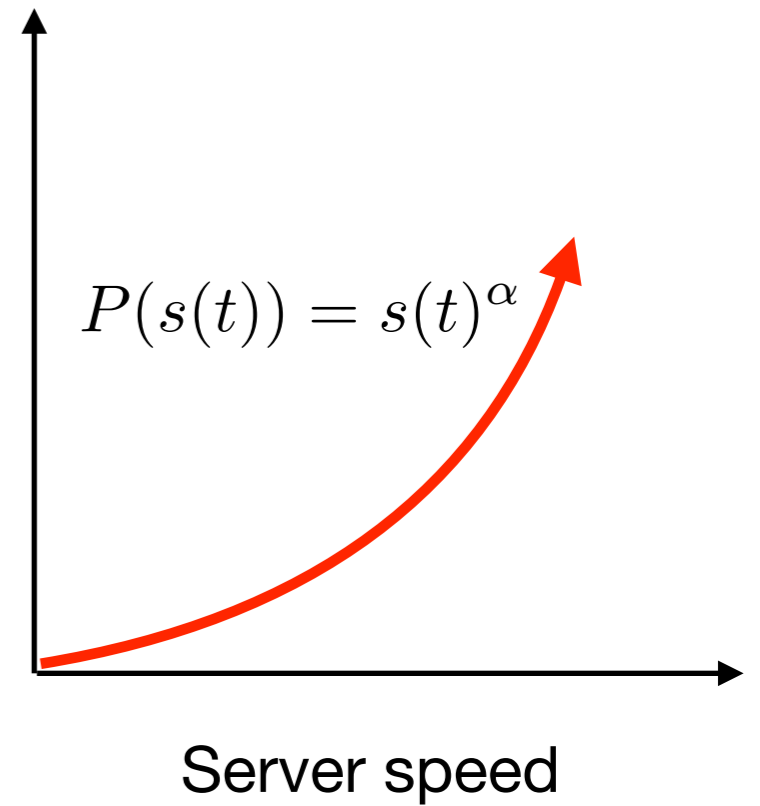
**Unlimited # of Servers**

# General Problem

packets  
arrivals




Per-Server  
Power  
Consumption




**Unlimited # of Servers**

# General Problem

packets  
arrivals



Per-Server  
Power  
Consumption



**Unlimited # of Servers**

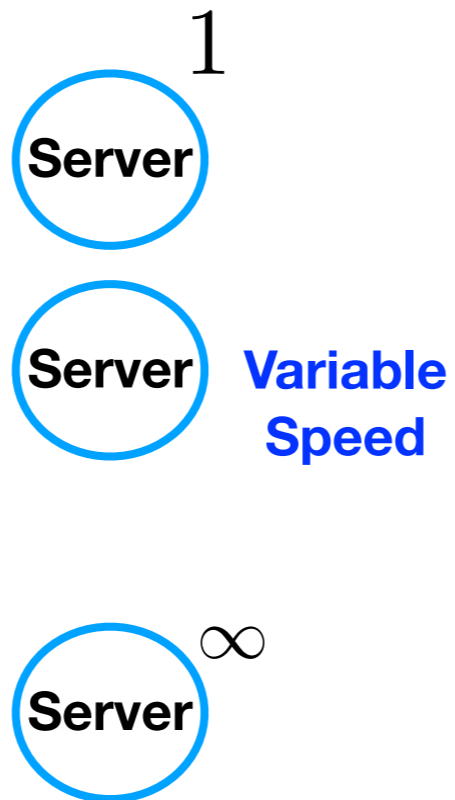
**Obj: *min* total flow time**

$$\int n(t) dt$$


$n(t)$  number of outstanding jobs at time  $t$

# General Problem

packets  
arrivals



Per-Server  
Power  
Consumption



**Unlimited # of Servers**

**Obj: *min* total flow time**


$$\int n(t) dt$$

$n(t)$  number of outstanding jobs at time  $t$

**subj: sum-power constraint  
across all servers**  
 $\mathcal{P}$

# General Problem

packets  
arrivals



Per-Server  
Power  
Consumption

$$P(s(t)) = s(t)^\alpha$$

Server speed

**Unlimited # of Servers**

**Obj: *min* total flow time**

$$\int n(t) dt$$



$n(t)$  number of outstanding jobs at time  $t$

**subj: sum-power constraint  
across all servers**  
 $\mathcal{P}$


**Decision : speed of each job/server**

# Result - General Problem

packets  
arrivals

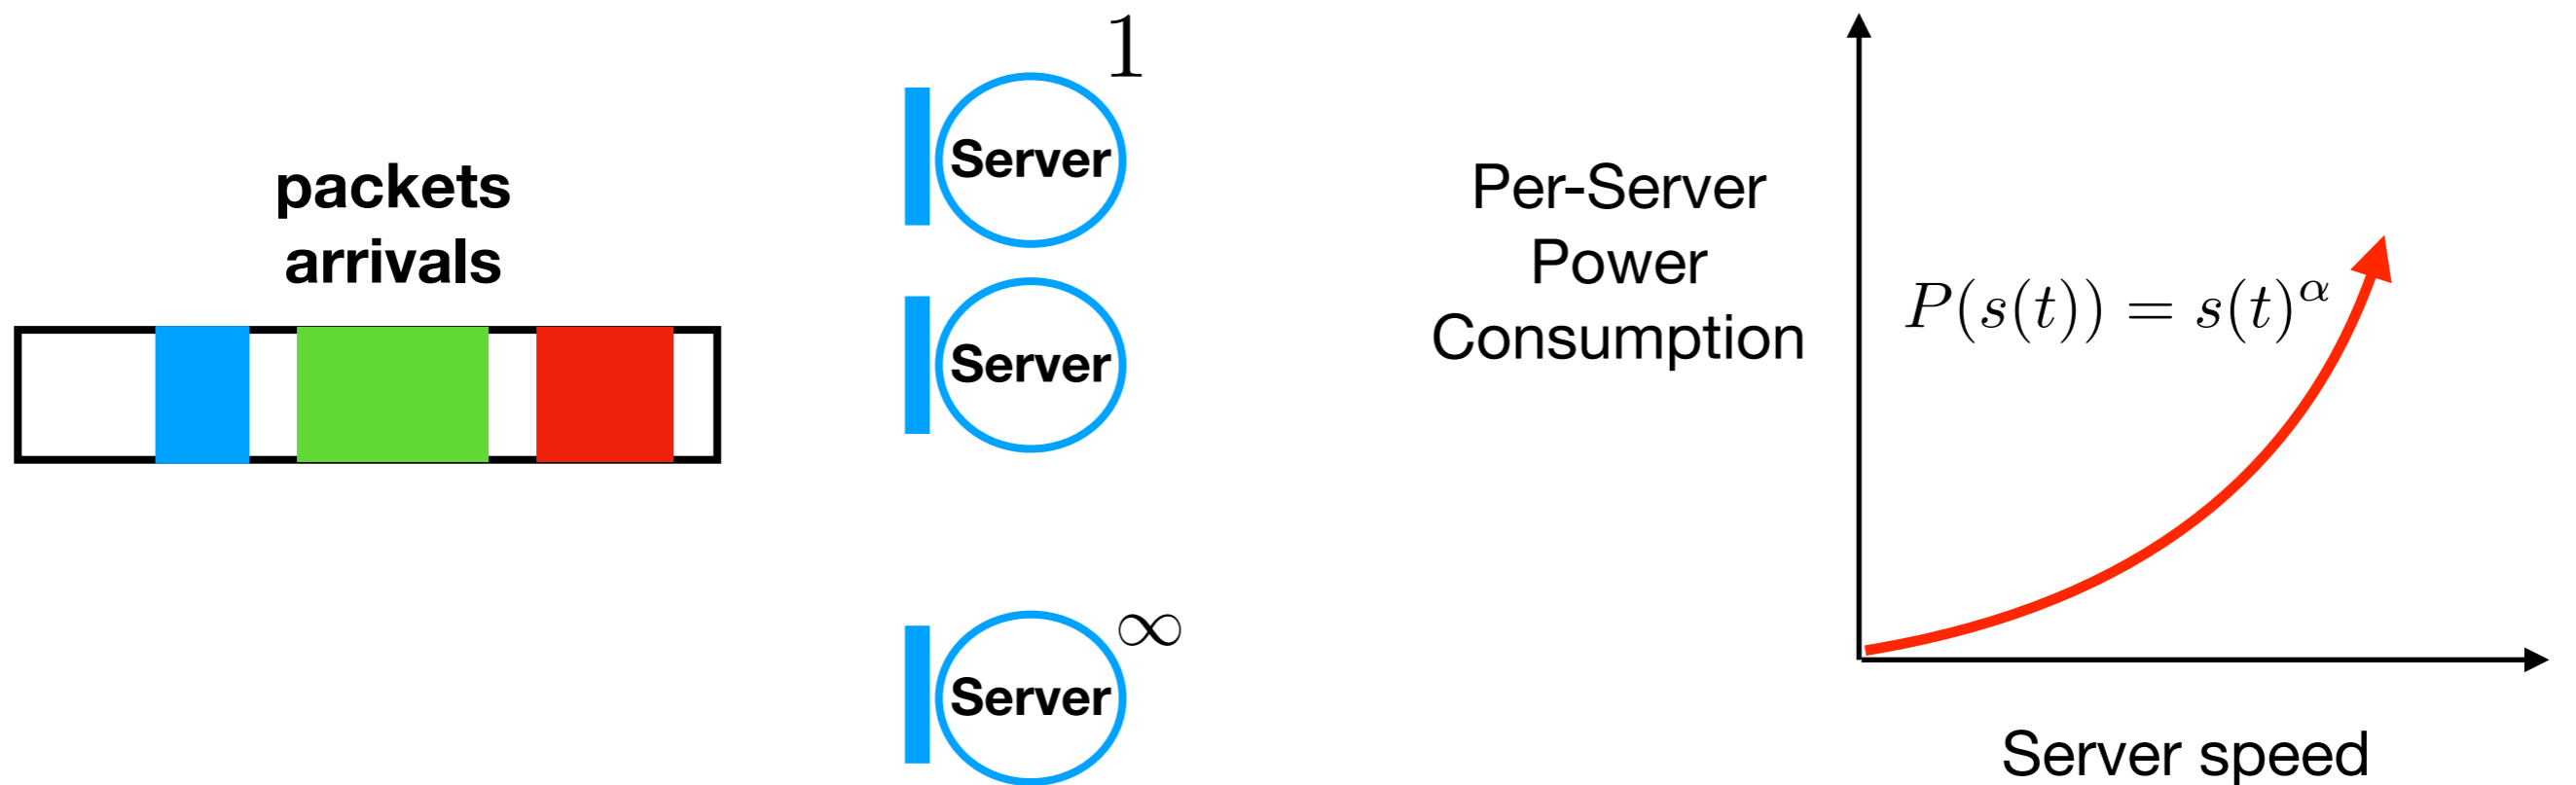


Per-Server  
Power  
Consumption



**Unlimited # of Servers**

# Result - General Problem



**Unlimited # of Servers**

**Competitive ratio is a constant that only depends on exponent  $\alpha$**



# Conclusions

# Conclusions

**An online algorithm for an important job-scheduling problem**

# Conclusions

**An online algorithm for an important job-scheduling problem**

**Fractional LCFS+Equal Speed** algorithm is constant competitive

# Conclusions

**An online algorithm for an important job-scheduling problem**

**Fractional LCFS+Equal Speed** algorithm is constant competitive

**Past approaches needed resource augmentation**

# Conclusions

**An online algorithm for an important job-scheduling problem**

**Fractional LCFS+Equal Speed** algorithm is constant competitive

**Past approaches needed resource augmentation**

## Open Questions

# Conclusions

An online algorithm for an important job-scheduling problem

**Fractional LCFS+Equal Speed** algorithm is constant competitive

Past approaches needed resource augmentation

## Open Questions

Competitive ratio of heSRPT algorithm (**that is locally optimal**) ?

# Conclusions

**An online algorithm for an important job-scheduling problem**

**Fractional LCFS+Equal Speed** algorithm is constant competitive

**Past approaches needed resource augmentation**

## Open Questions

**Competitive ratio of heSRPT algorithm (**that is locally optimal**) ?**

**Lower Bound on the competitive ratio of any online algorithm?**