

# Load Balancing Guardrails

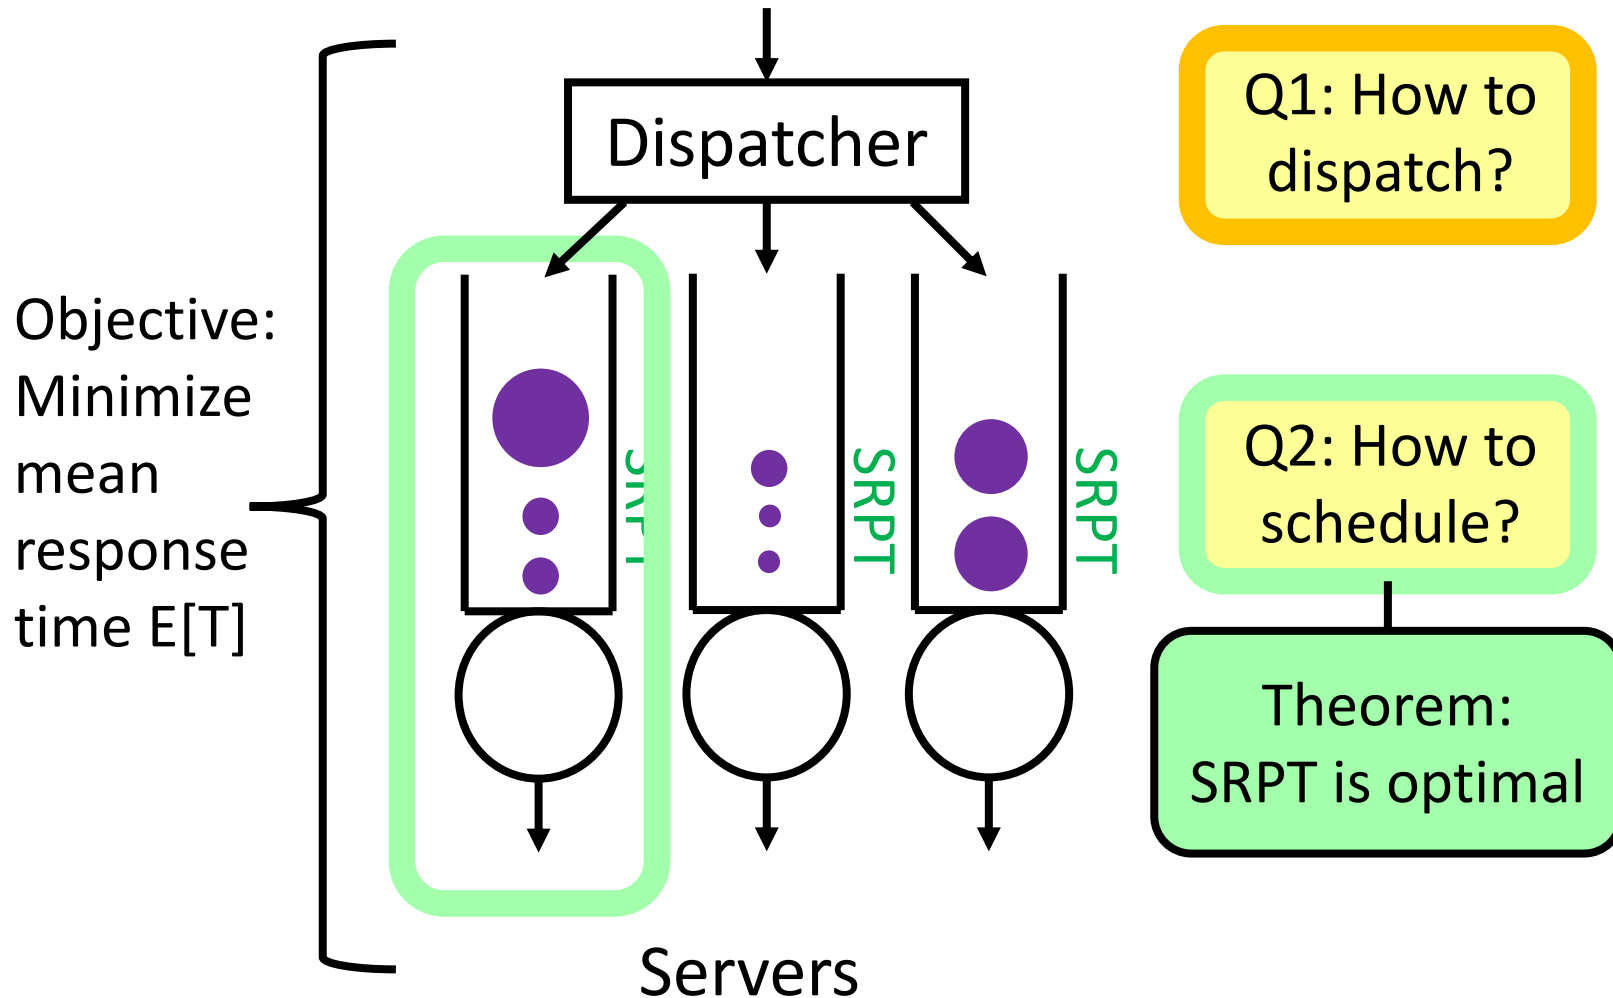
Keeping Your Heavy Traffic on the Road to Low Response Times

**Isaac Grosf (CMU)**

Ziv Scully (CMU)

Mor Harchol-Balter (CMU)

# Goal: Optimal Load Balancing



Q1: How to dispatch?


Q2: How to schedule?

Theorem: SRPT is optimal

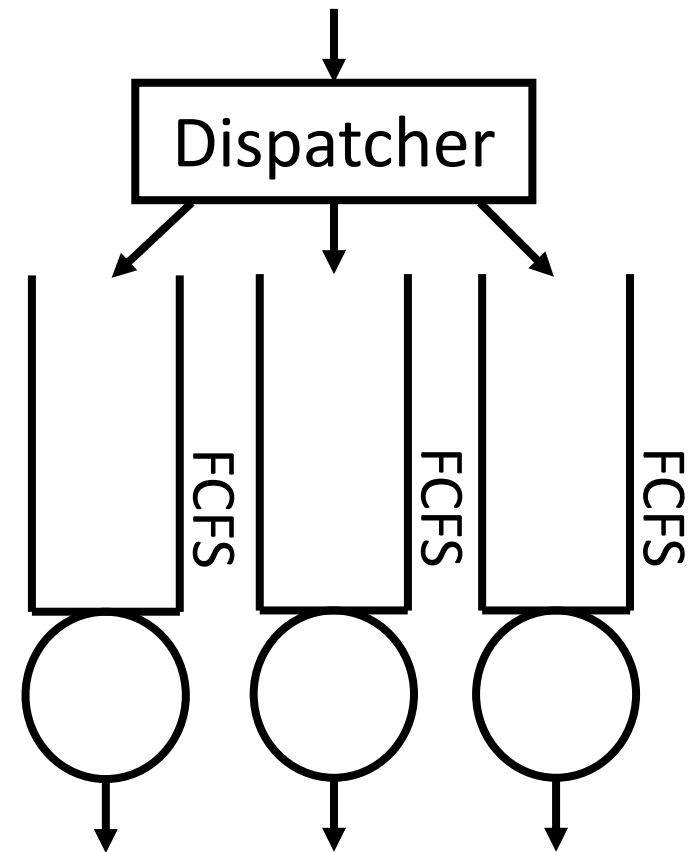
Assumptions:  
Stochastic Arrivals  
Known Sizes  
Preempt-Resume

# Prior Work on Dispatching

SRPT: Very little prior work



FCFS: Tons of prior work




# Prior Work on Dispatching - FCFS

**Join-Shortest-Queue (JSQ):** Winston, Weber, Whitt, Lin, Raghavendra, Foley, McDonald, Bramson, Lu, Prabhakar, Eschenfeldt, Gamarnik, ...

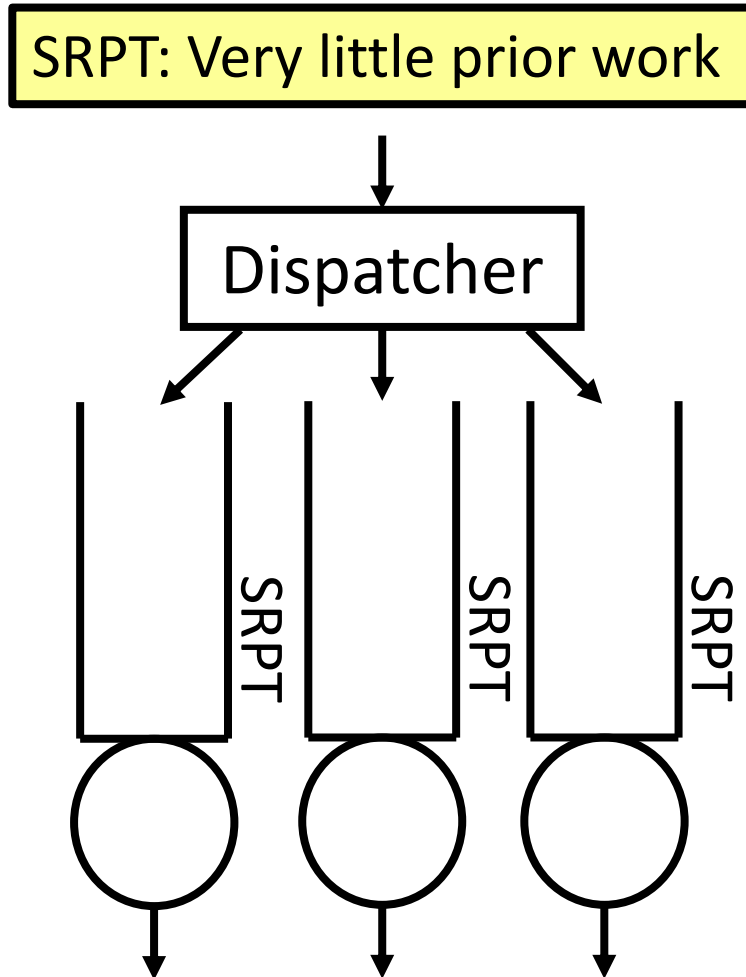
**Join-Shortest-of-d-Queues (JSQ-d):** Vvedenskaya, Dobrushin, Karpelevich, Mitzenmacher, Bramson, Ying, Srikant, Kang, Muckherjee, Borst, Leeuqaarden, ...

**Least-Work-Left (LWL):** Lee, Longton, Kingman, Takahashi, Daly, Tijms, Van Hoorn, Ma, Mark, Breur, Hokstad, Kimura, Gupta, Harchol-Balter, Dai, Zwart, Osogami, Whitt, ...

**Size-Interval-Task-Assignment (SITA):** Harchol-Balter, Crovella, Murta, Bachmat, Sarfati, Vesilo, Scheller-Wolf, ...



# Prior Work on Dispatching - SRPT

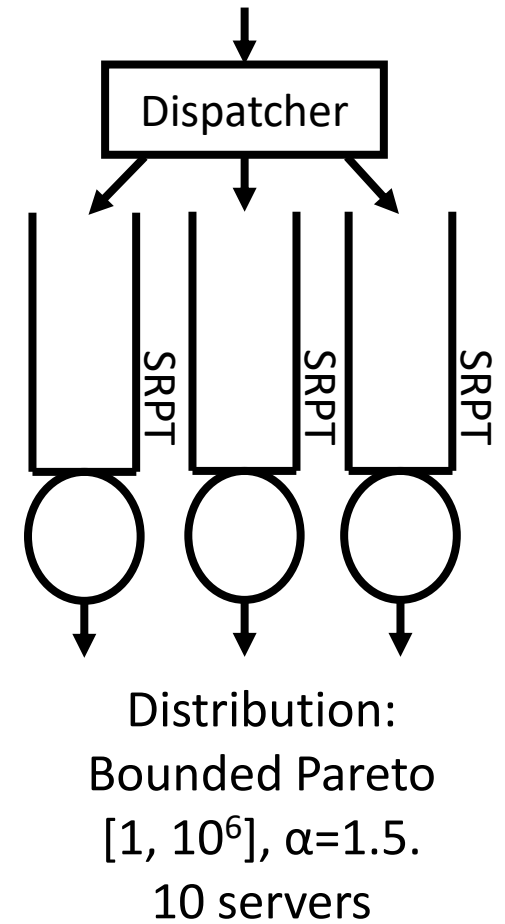
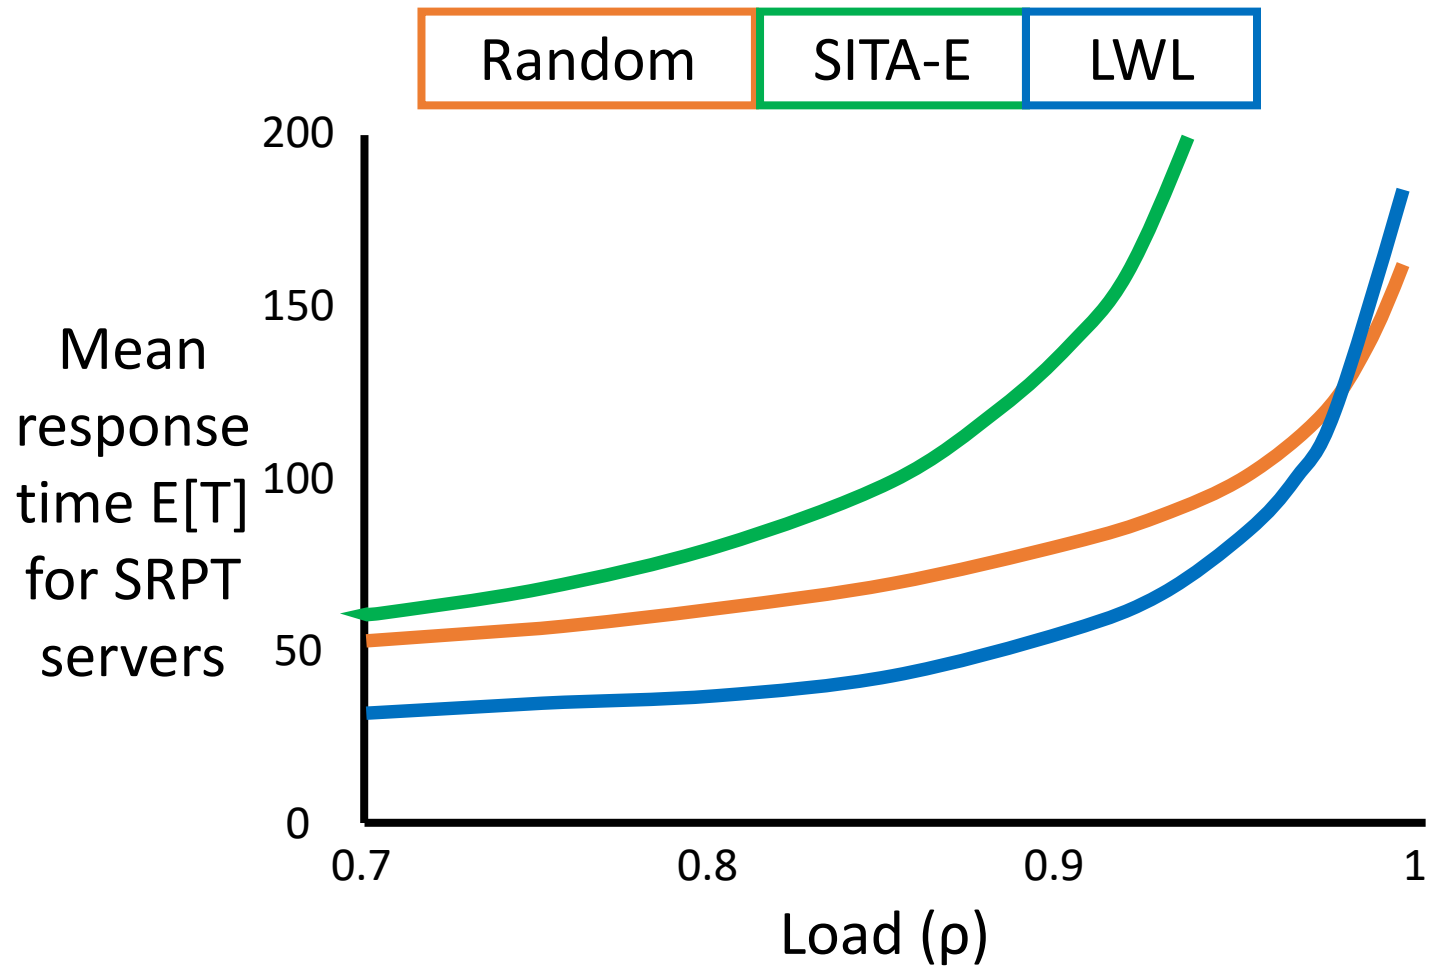


Random dispatch: Trivial



First Policy Iteration (FPI) Heuristic  
[Hyytiä & Aalto '12]

Multilayered Round Robin  
[Down & Wu '06]

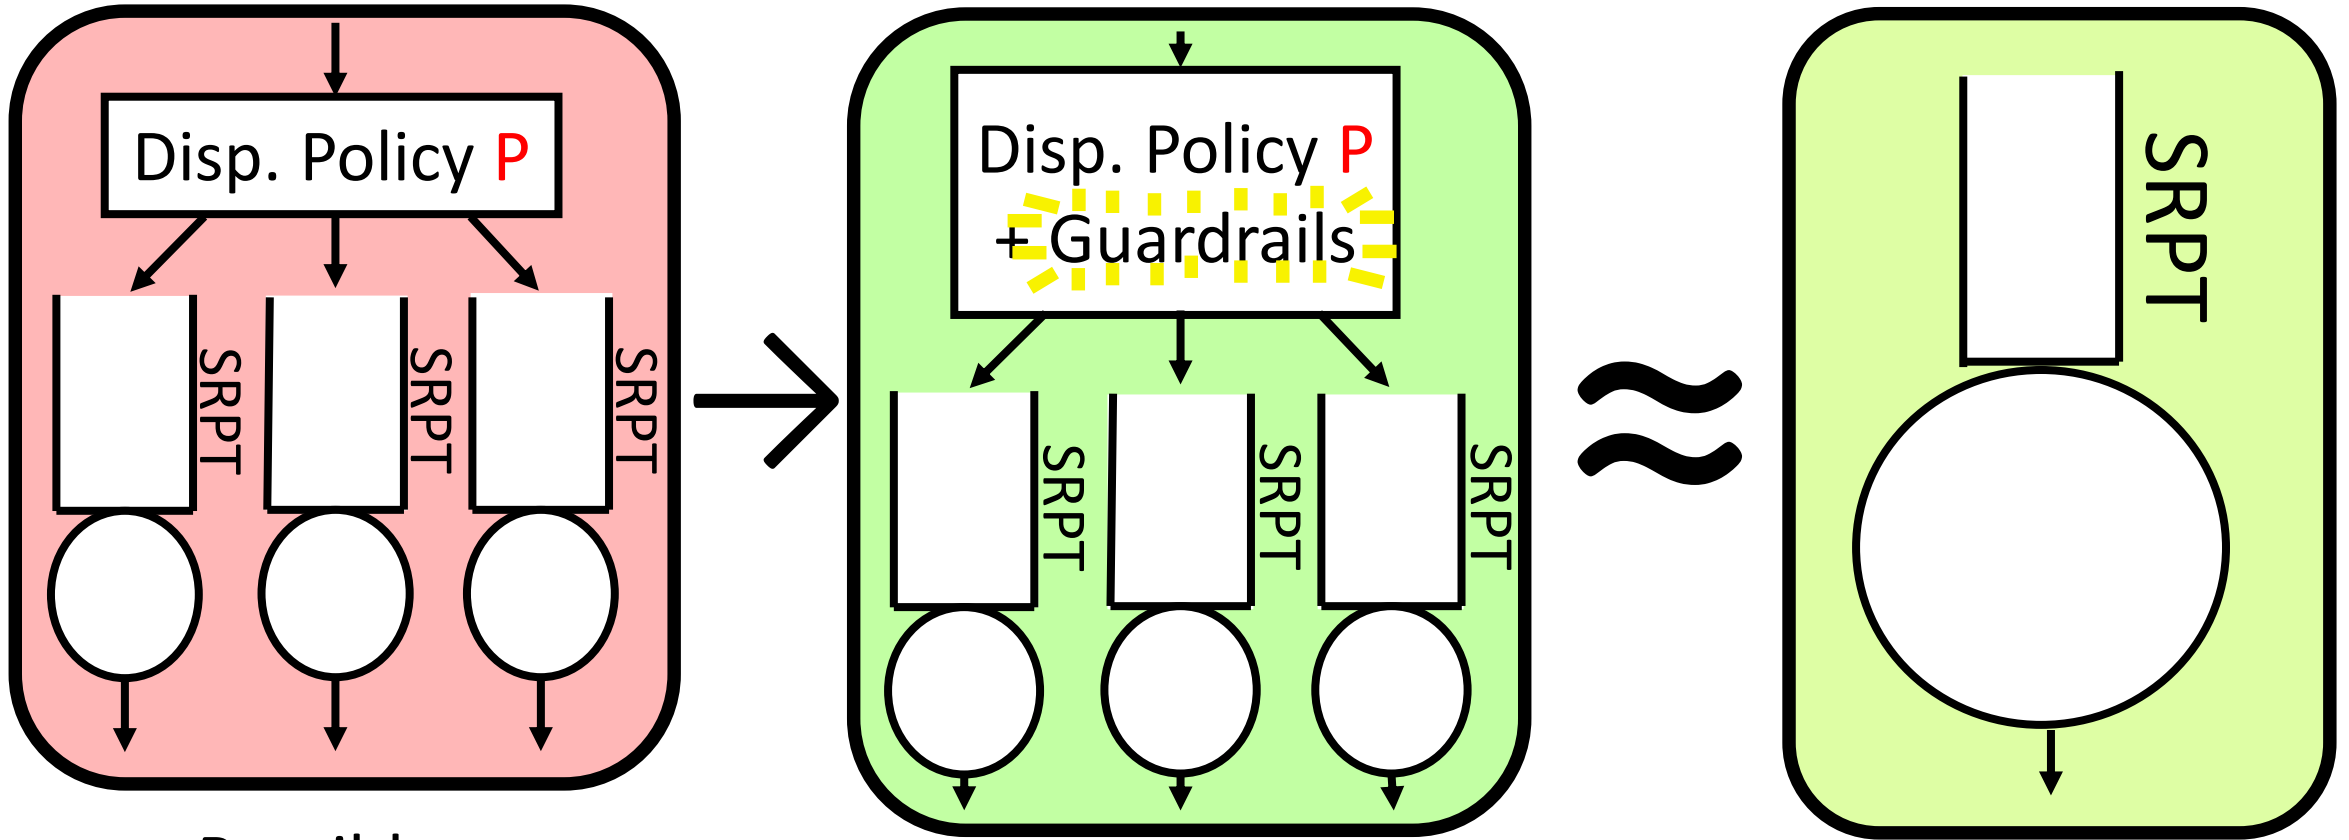
Good for FCFS  $\stackrel{?}{\Rightarrow}$  Good for SRPT



# Good for FCFS $\not\Rightarrow$ Good for SRPT



# Our Contribution: Guardrails

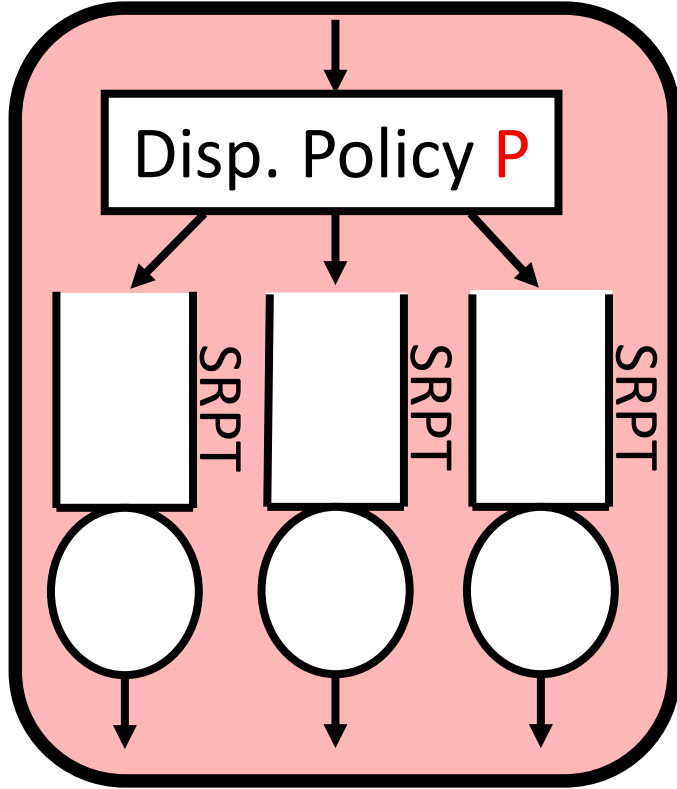


Possibly  
very  
bad

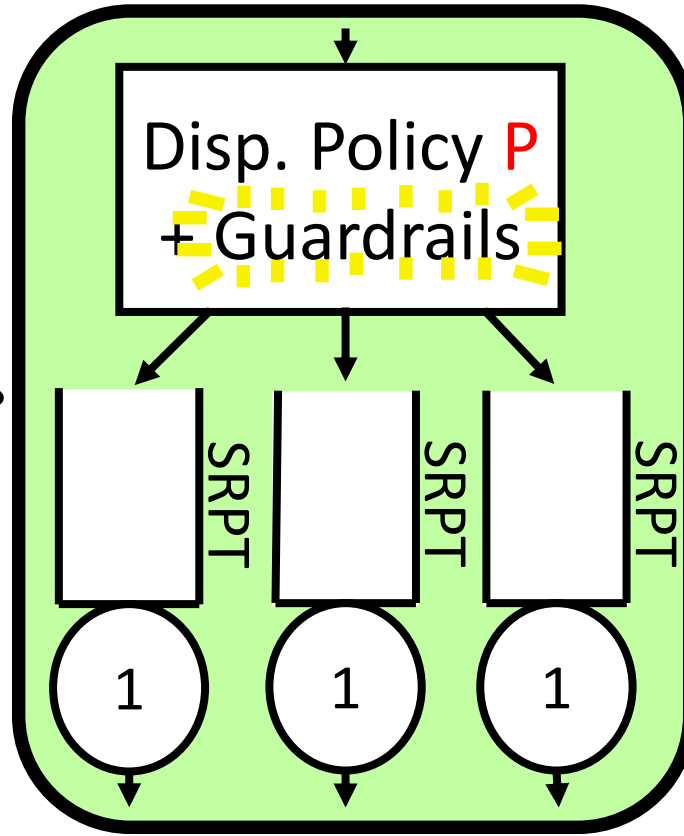

Guaranteed  
heavy traffic  
optimal




# Our Contribution: Guardrails





Possibly  
very  
bad




Guaranteed  
heavy traffic  
optimal



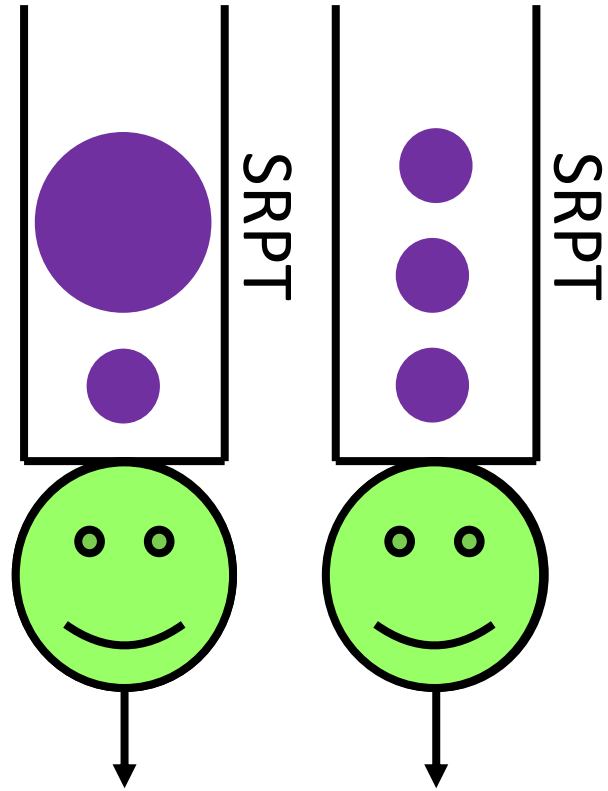
# Good for FCFS $\not\Rightarrow$ Good for SRPT



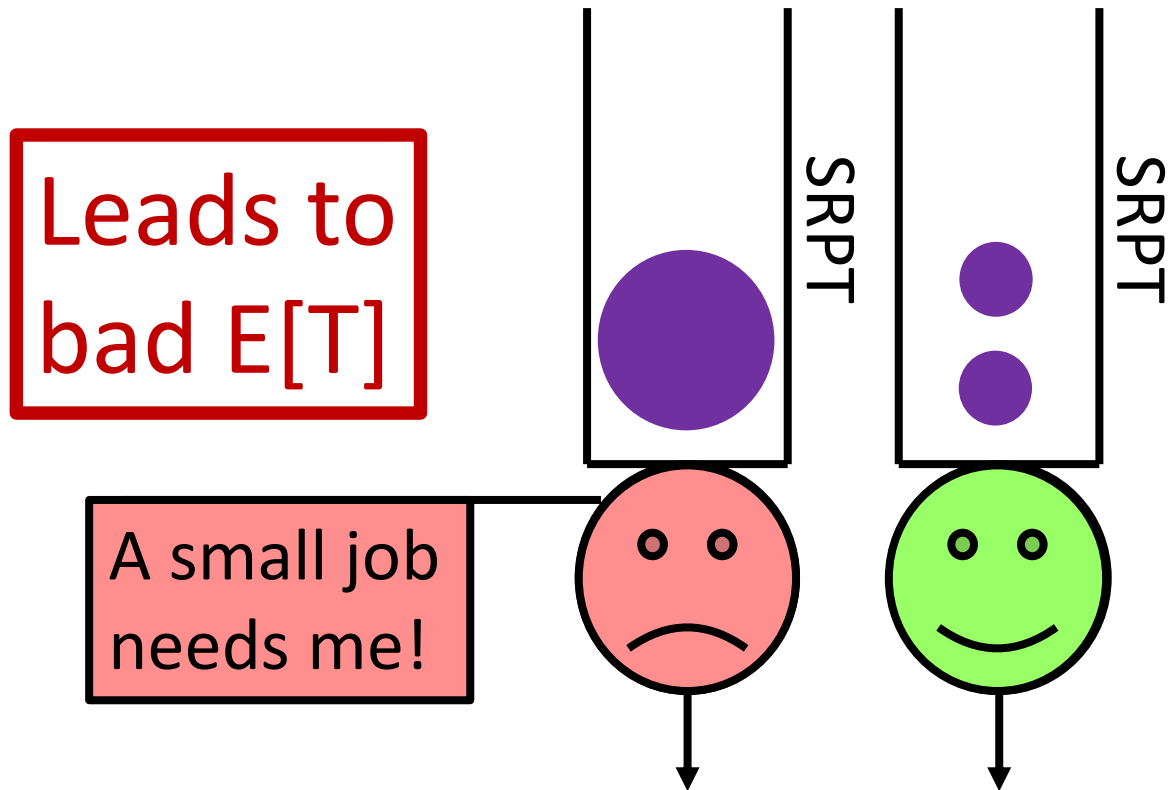
# Dispatching to SRPT Servers



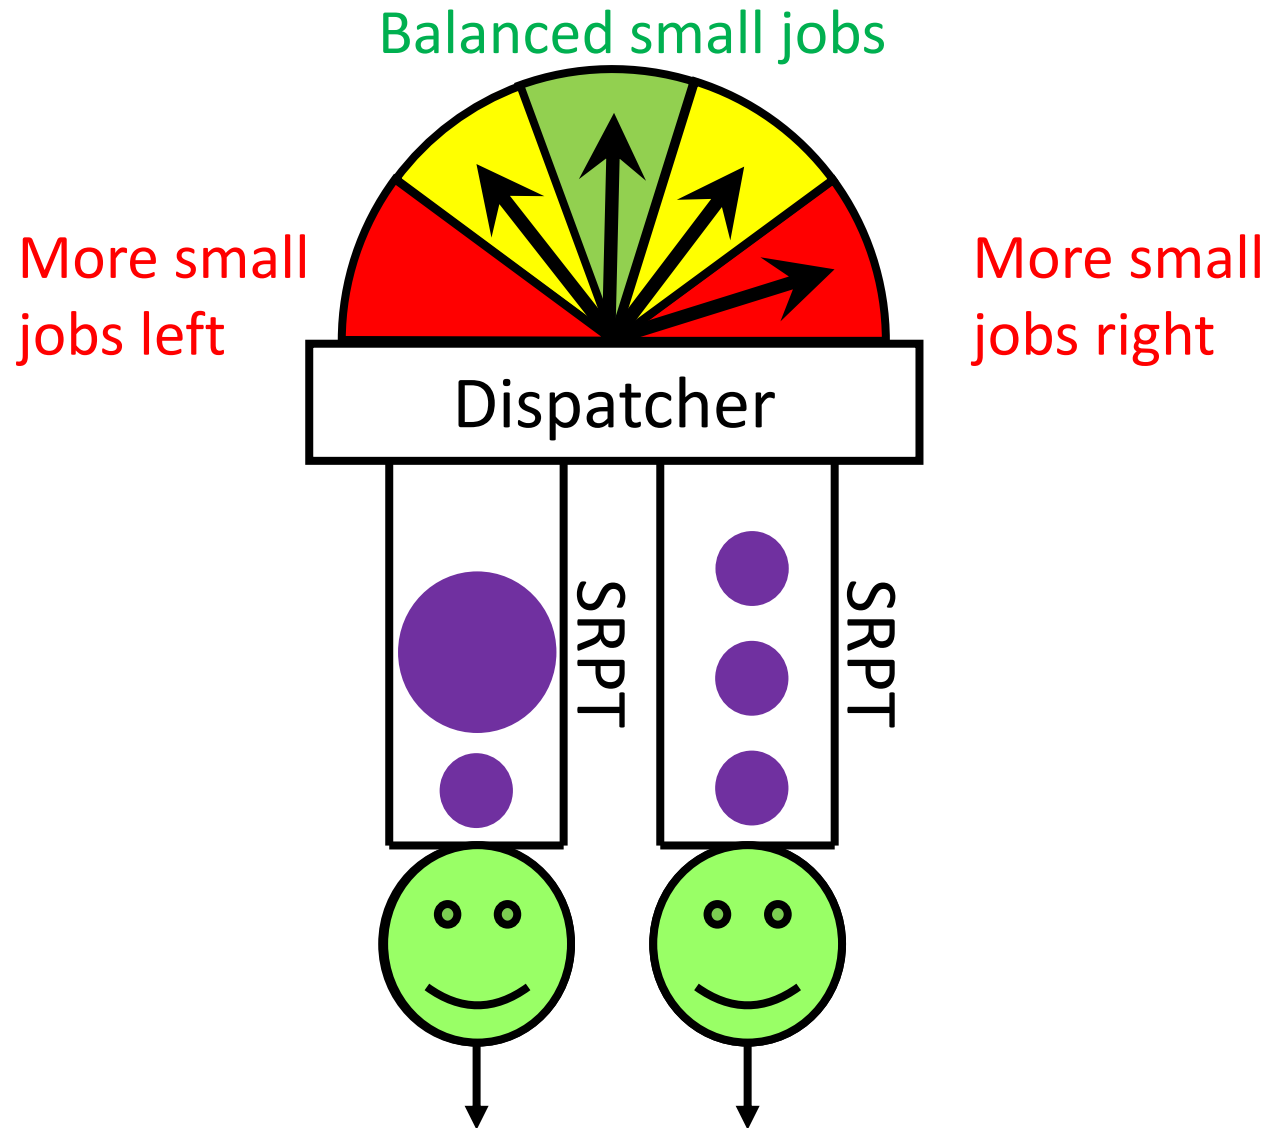
# Dispatching to SRPT Servers



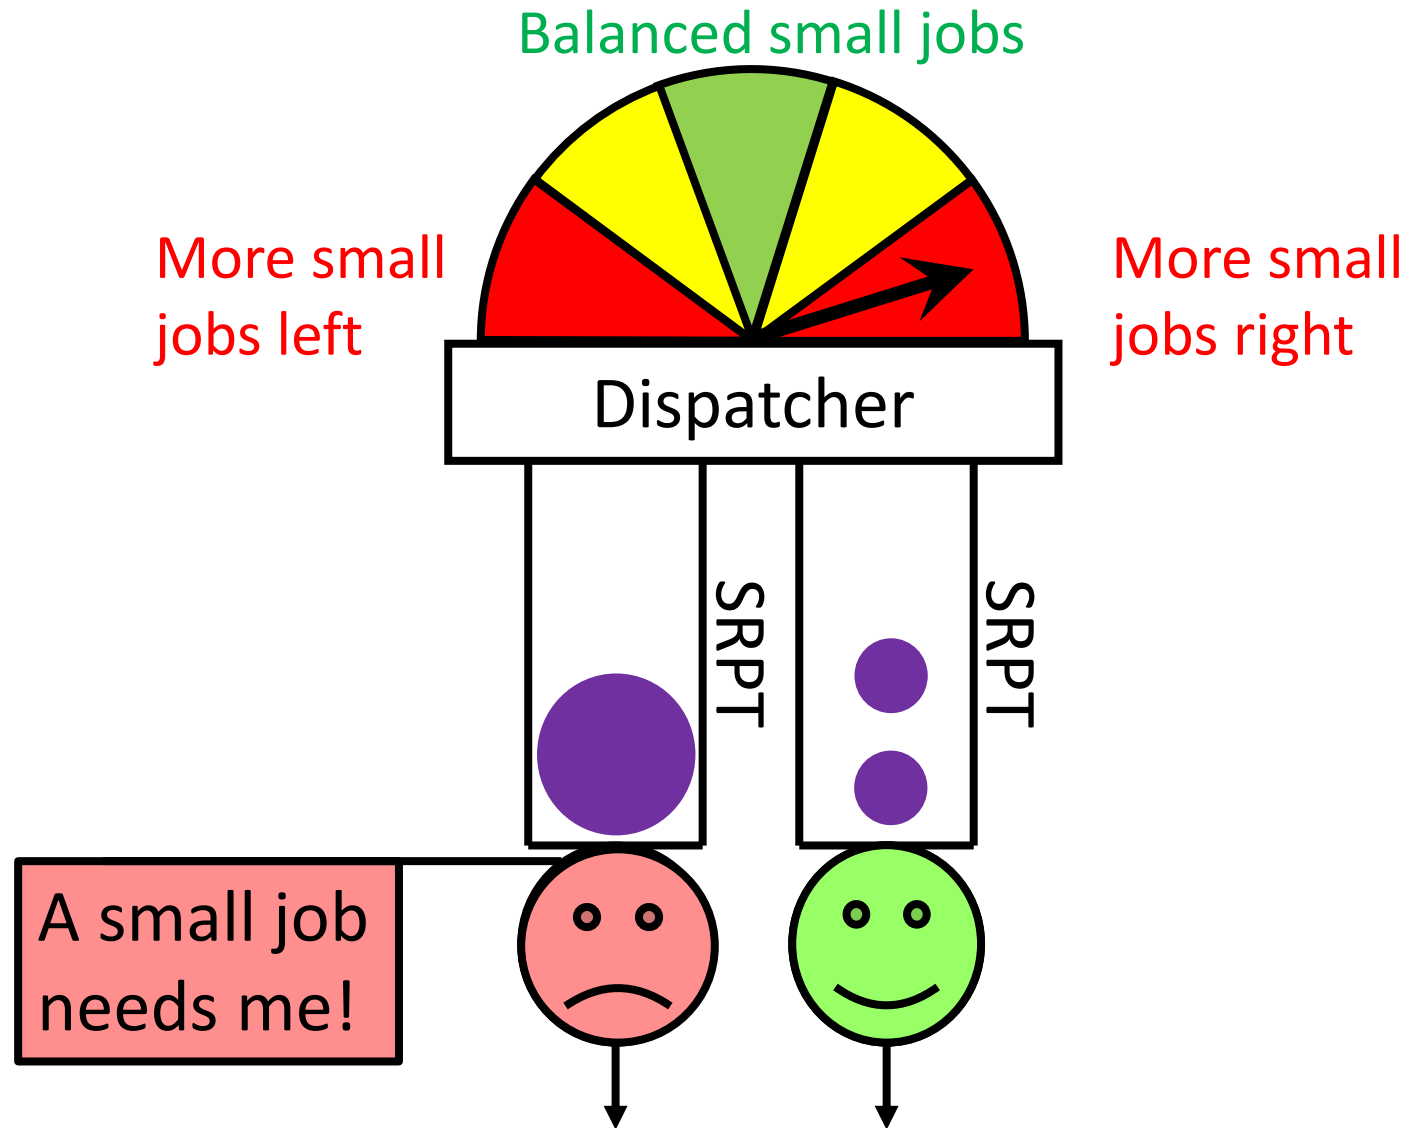
# Dispatching to SRPT Servers



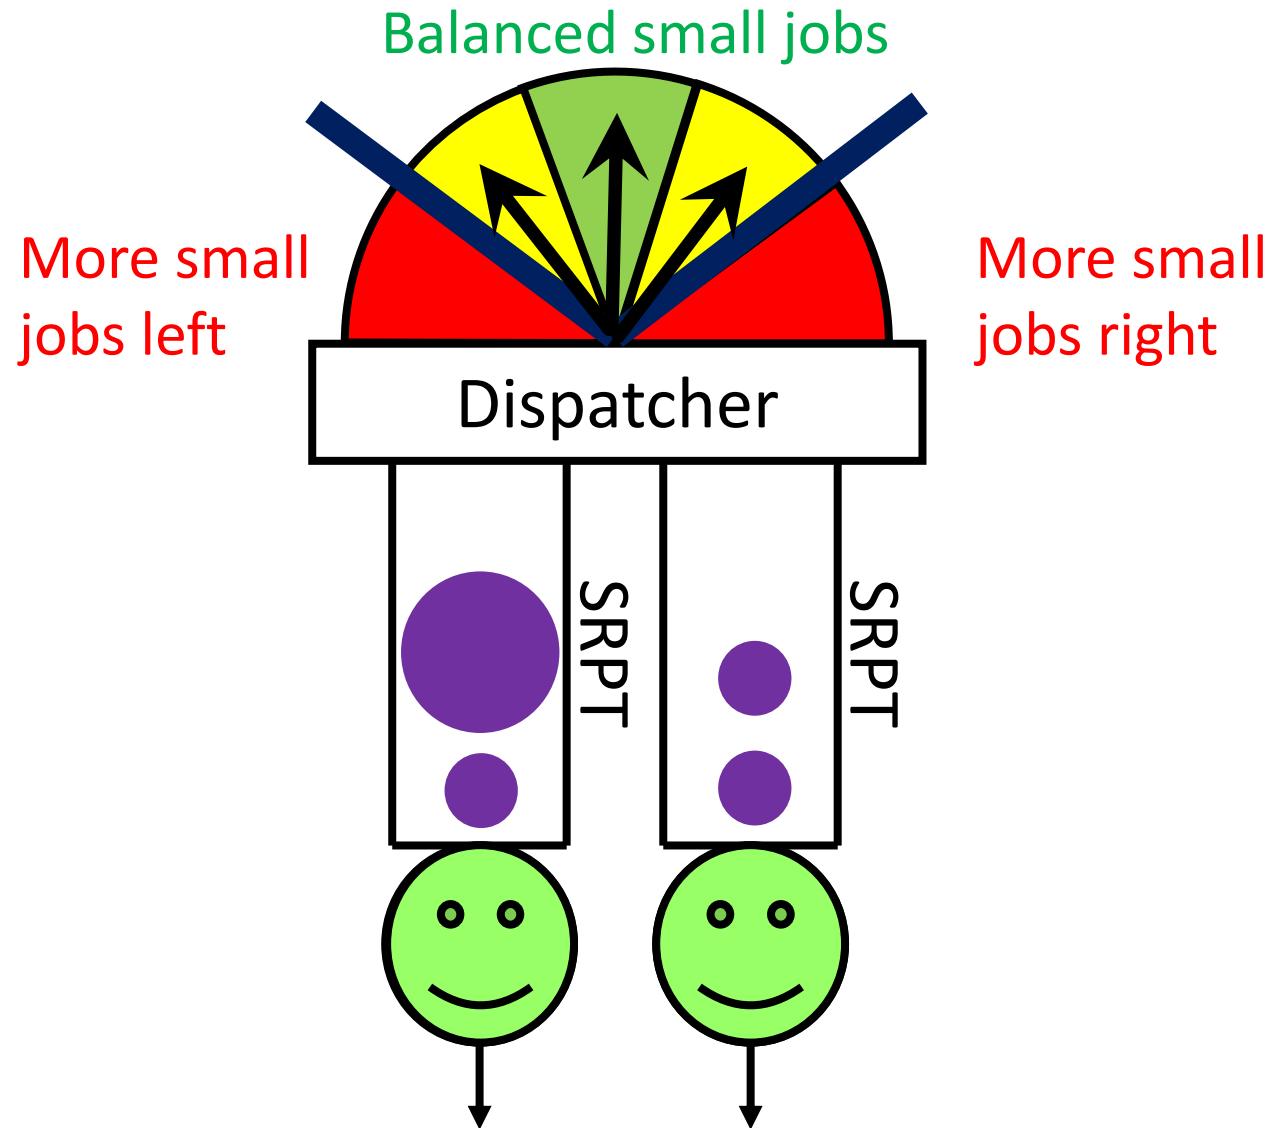
# Problem: Small Job Imbalance



# Problem: Small Job Imbalance




# Guardrails

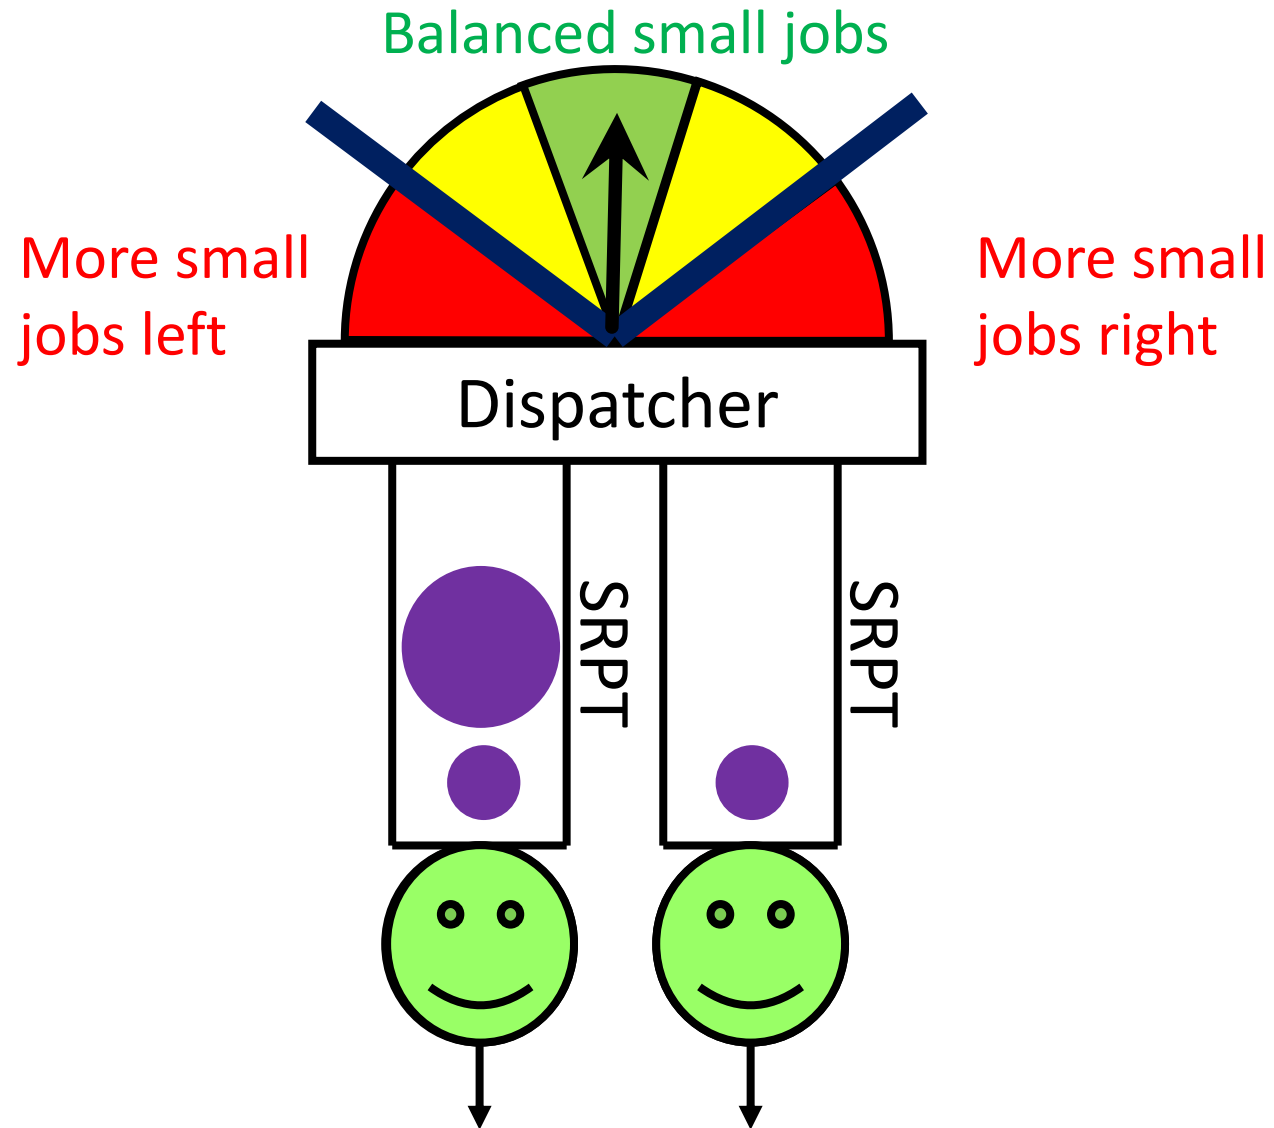





# Guardrails




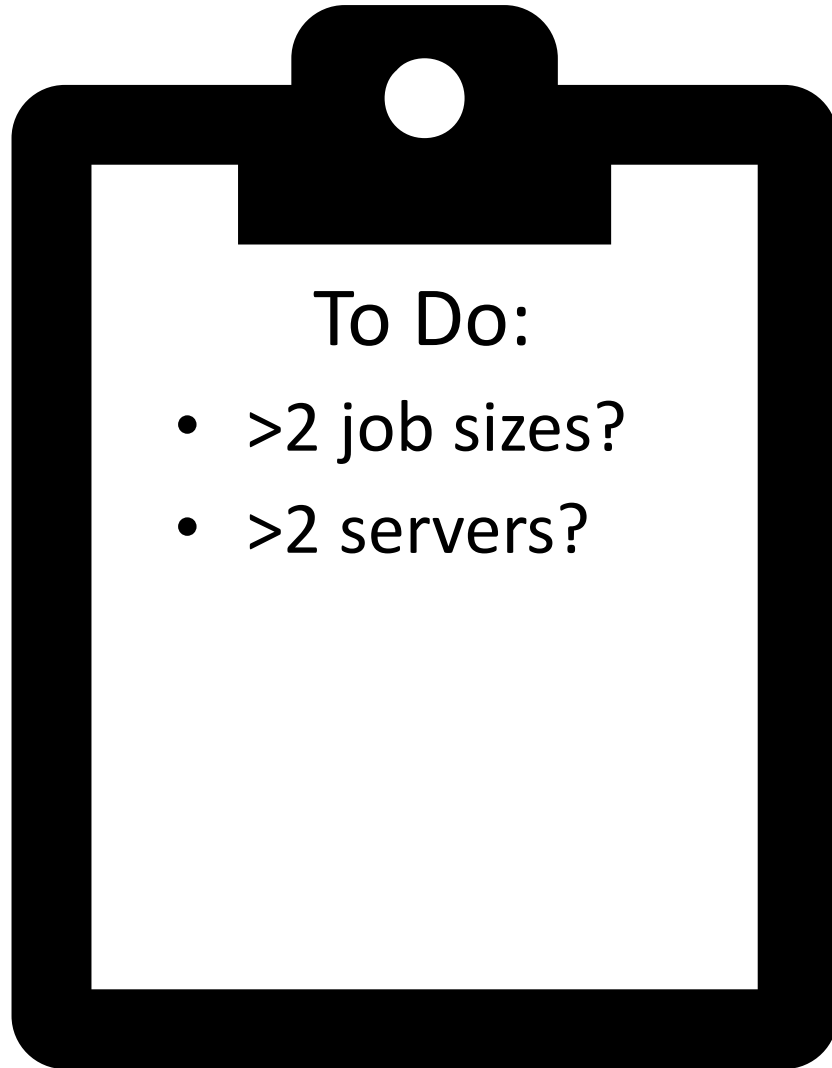
# Guardrails




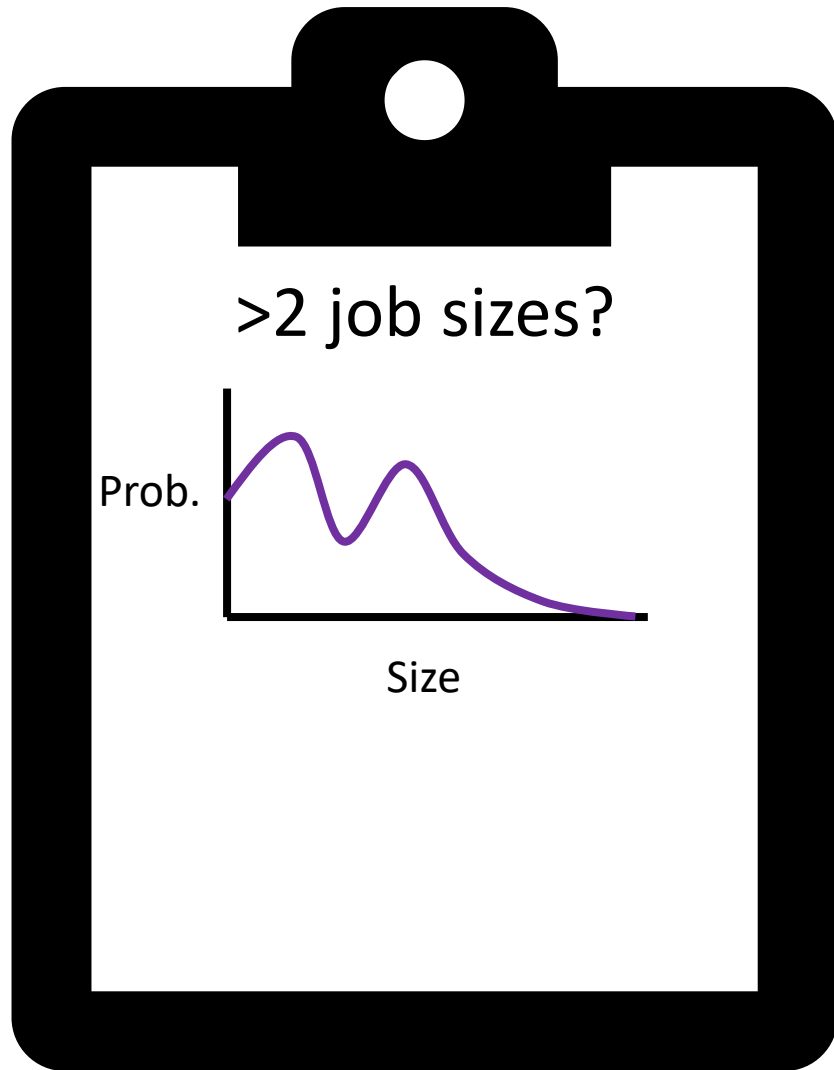
# Guardrails



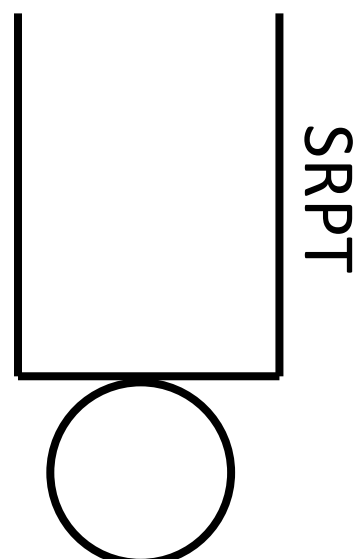

# Guardrails




# Guardrails



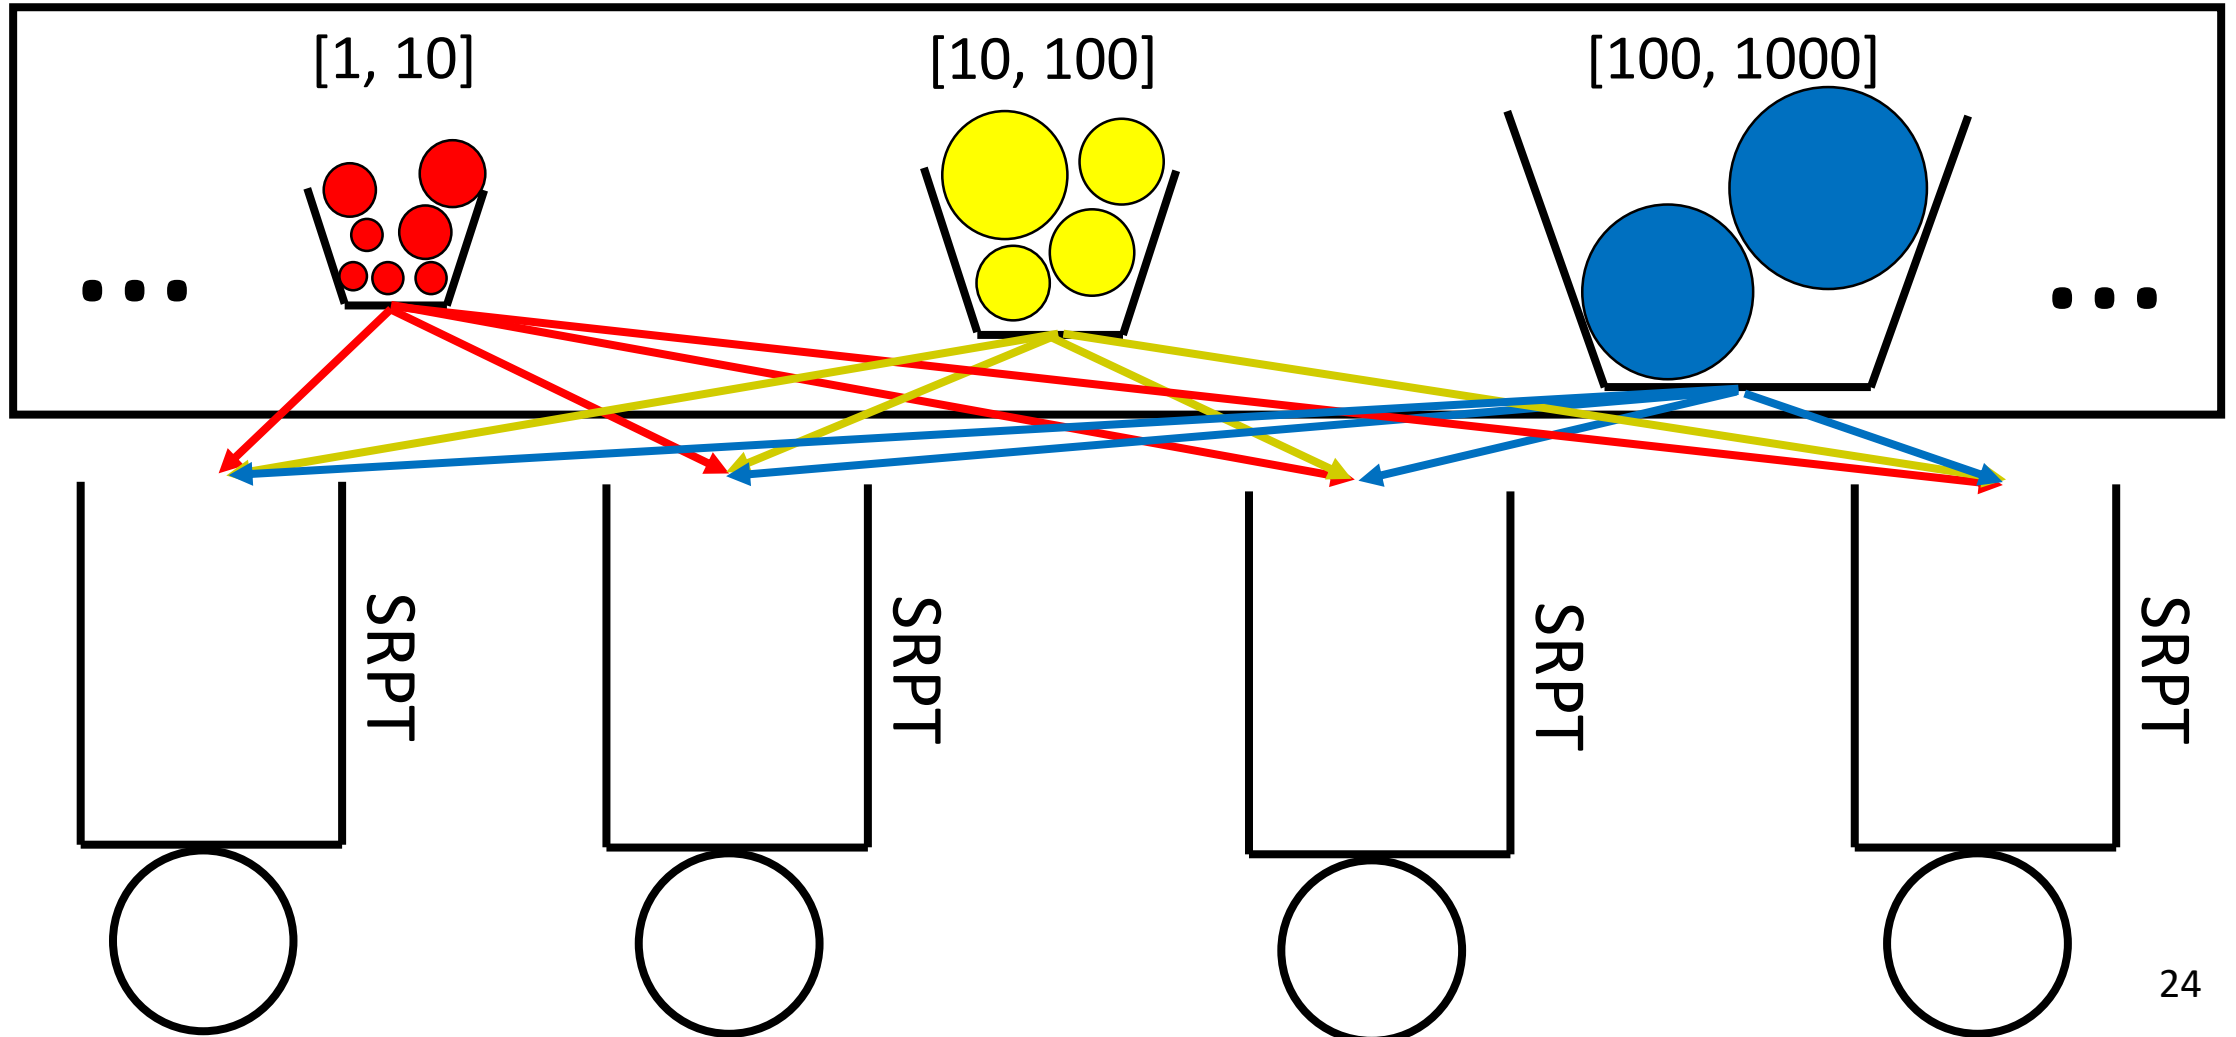
# Guardrails: Bucketing



# Guardrails: Bucketing



# Guardrails: Bucketing





# Precise Dispatching Requirement

Job of size  $x$  has rank  $r \leftrightarrow x \in [c^r, c^{r+1})^*$

$V_i^r(t)$  = Volume of rank  $r$  work dispatched to server  $i$  by time  $t$ .

Guardrail requirement:

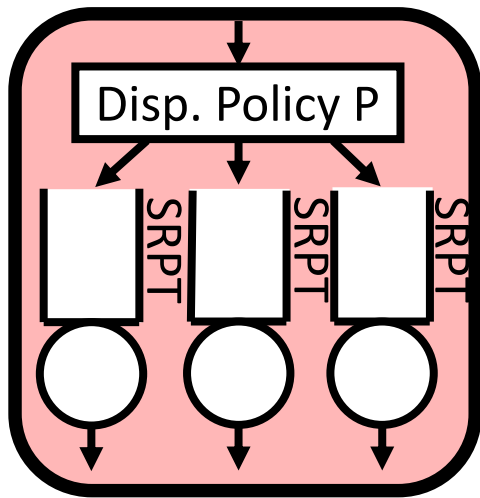
$\forall$  ranks  $r$ ,  $\forall$  servers  $i, j$ ,  $\forall$  times  $t$ ,

$$|V_i^r(t) - V_j^r(t)| \leq c^{r+1}$$



\*  $c$  is chosen as a function of load  $\rho$ .

# The Guardrail Theorem

$$\frac{E[\text{Resp. Time of Disp. Policy P with Guardrails}]}{E[\text{Resp. Time of Single SRPT Superserver}]}$$



Possibly  
Very bad



Guaranteed heavy  
traffic optimal

≈  
w.r.t.  
 $E[T]$

