

Recent Progress in

Queueing and Scheduling Theory

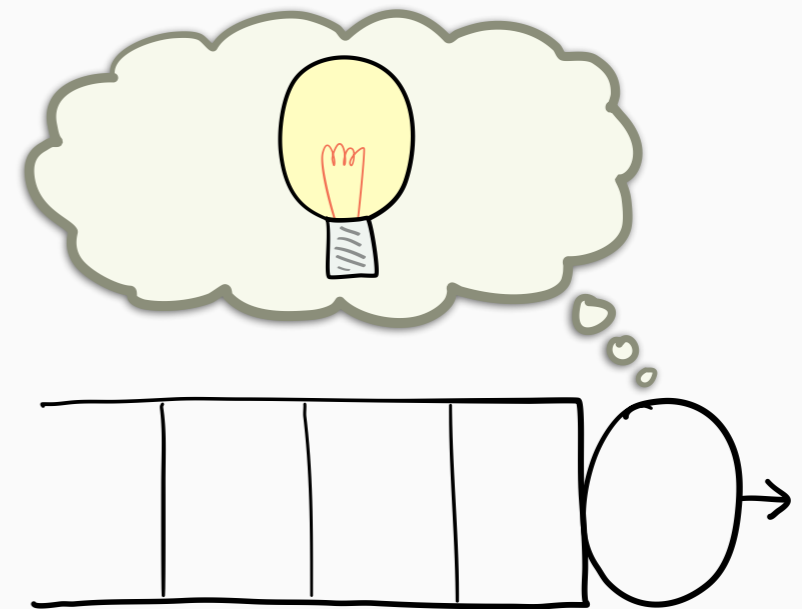
(for a TCS Audience)

Ziv Scully

Harvard & MIT → Cornell

`zivscully@cornell.edu`

`https://ziv.codes`



Collaborators




Isaac Grosf
CMU



Mor Harchol-Balter
CMU




Alan Scheller-Wolf
CMU




Michael Mitzenmacher
Harvard

Contention




Queueing




Delay

Contention



Queueing




Delay

healthcare


supply chains

Contention

healthcare



Queueing




Delay

your local supermarket


supply chains

Contention

healthcare



Queueing



Delay


your local supermarket

supply chains


Contention

call centers

healthcare



Queueing



Delay


your local supermarket

supply chains

Contention


call centers

healthcare



Queueing

transportation



Delay


your local supermarket

supply chains

Contention


call centers

healthcare



Queueing

transportation



Delay

databases


your local supermarket

supply chains

Contention


call centers

healthcare



Queueing

transportation



Delay

databases

networks


your local supermarket

supply chains

Contention


call centers

healthcare



Queueing

transportation



Delay

databases

networks

operating systems


your local supermarket

supply chains

Contention


call centers

healthcare



Queueing

transportation



databases

computer architecture

Delay

networks

operating systems


your local supermarket

supply chains

Contention

call centers


healthcare



transportation

Queueing

supercomputing



databases

computer architecture

Delay

networks

operating systems


your local supermarket

supply chains

Contention

call centers


healthcare



transportation

Queueing

supercomputing



databases

computer architecture

Delay

networks

operating systems



How to reduce delays?


your local supermarket

supply chains

Contention

call centers


healthcare



transportation

Queueing

supercomputing



databases

computer architecture

Delay


networks

operating systems



How to reduce delays?

Scheduling



Good news:

scheduling can reduce delay




Good news:

scheduling can reduce delay



Bad news:

limited understanding of scheduling



Good news:


scheduling can reduce delay



Bad news:

limited understanding of scheduling

evaluation



Good news:

scheduling can reduce delay




Bad news:

limited understanding of scheduling

design

evaluation




Good news:
scheduling can reduce delay




Bad news:
limited understanding of scheduling

design

evaluation




We need:
rigorous theory of scheduling




We need:

rigorous theory of scheduling



We need:

rigorous theory of scheduling




We need:

rigorous theory of scheduling



CS Theory




We need:

rigorous theory of scheduling



CS Theory

- Worst-case modeling
- Complex algorithms



We need:

rigorous theory of scheduling




CS Theory



Queueing Theory

- Worst-case modeling
- Complex algorithms



We need:

rigorous theory of scheduling




CS Theory

- Worst-case modeling
- Complex algorithms



Queueing Theory

- Stochastic modeling
- Simple algorithms



We need:

rigorous theory of scheduling



CS Theory




- Worst-case modeling
- Complex algorithms



Queueing Theory

- Stochastic modeling
- Simple algorithms



We need:

rigorous theory of scheduling



CS Theory




- Worst-case modeling
- Complex algorithms



Queueing Theory



- Stochastic modeling
- Simple algorithms



We need:

rigorous theory of scheduling



CS Theory




- Worst-case modeling
- Complex algorithms



Queueing Theory



- Stochastic modeling
- Simple algorithms



We need:

rigorous theory of scheduling



CS Theory



- Worst-case modeling
- Complex algorithms




Queueing Theory



- Stochastic modeling
- Simple algorithms





We need:

rigorous theory of scheduling



CS Theory



- Worst-case modeling
- Complex algorithms




Queueing Theory



- Stochastic modeling
- Simple algorithms



Best to learn from both



We need:

rigorous theory of scheduling



CS Theory



- Worst-case modeling
- Complex algorithms



Queueing Theory



- Stochastic modeling
- Simple algorithms



Best to learn from both

Today's talk

Today's talk

scheduling with

multiple servers

Today's talk

scheduling with

multiple servers

scheduling with

noisy predictions

Today's talk

scheduling with

multiple servers



TCS



Queueing

scheduling with

noisy predictions



TCS



Queueing

Today's talk

scheduling with

multiple servers



TCS



Queueing

scheduling with

noisy predictions



TCS



Queueing



Powered by **new tools**
in **queueing theory**

Today's talk

scheduling with

multiple servers



TCS



Queueing

scheduling with

noisy predictions



TCS




Queueing




Powered by **new tools**
in **queueing theory**

possible
crossover


What is scheduling?




What is scheduling?





What is scheduling?





What is scheduling?





What is scheduling?





What is scheduling?





What is scheduling?



What is scheduling?




What is scheduling?




What is scheduling?

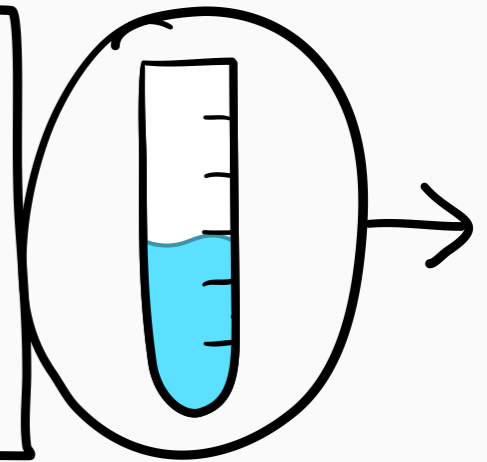
online arrivals




queue



server




job




What is scheduling?


online arrivals




queue



server



job




What is scheduling?

online arrivals


queue

server

Queueing model:



job



What is scheduling?


online arrivals

queue

server

Queueing model:

- λ = arrival rate (Poisson)



What is scheduling?


online arrivals

queue

server

Queueing model:

- λ = arrival rate (Poisson)
- S = job size distribution



What is scheduling?


online arrivals

queue

server


Queueing model:

- λ = arrival rate (Poisson)
- S = job size distribution
- $\rho = \lambda \mathbf{E}[S] = \text{load} < 1$




What is scheduling?


online arrivals




queue



server




job

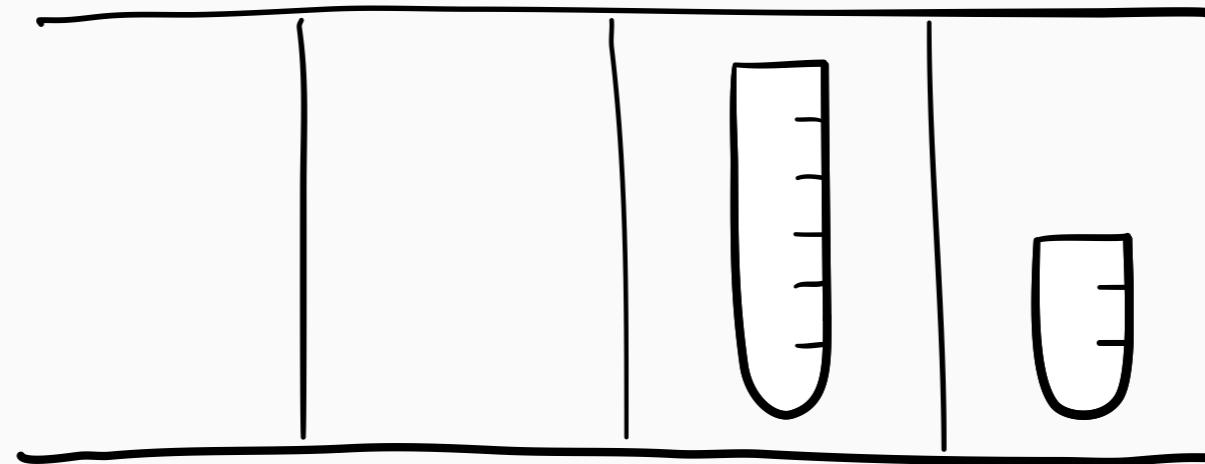


What is scheduling?


online arrivals




queue



server




job




What is scheduling?


online arrivals




queue



server




job




What is scheduling?


online arrivals




queue



server




job




What is scheduling?

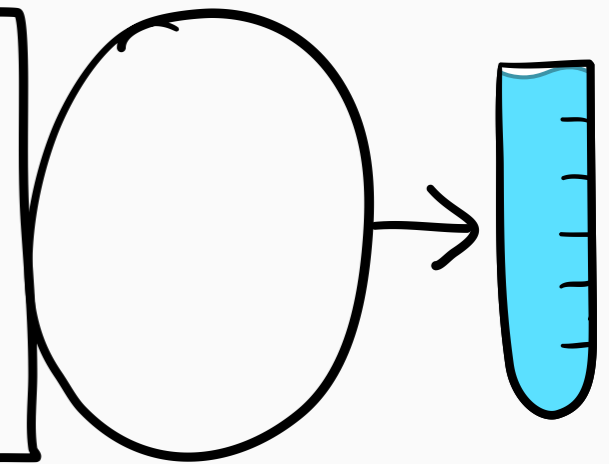
online arrivals




queue



server





job




What is scheduling?

online arrivals




job




What is scheduling?

online arrivals




queue

server




$T =$ response time

job




What is scheduling?

online arrivals




queue

server




$T =$ response time

job




What is scheduling?


online arrivals



queue




server



$T =$ response time

job




remaining size


age

What is scheduling?


online arrivals



queue




server




$T =$ response time

job




What is scheduling?

online arrivals




queue

server




$T =$ response time

job




What is scheduling?

online arrivals




queue

server




$T =$ response time

job




What is scheduling?

online arrivals




queue

server




$T =$ response time

job




What is scheduling?

online arrivals




queue

server



$T =$ response time


job



Question: schedule to minimize $E[T]$?


What is scheduling?

online arrivals




queue

server



$T =$ response time

job





SRPT

shortest remaining
processing time

Question: schedule
to minimize $E[T]$?

Multiserver scheduling



online arrivals



$T =$ response time

Multiserver scheduling

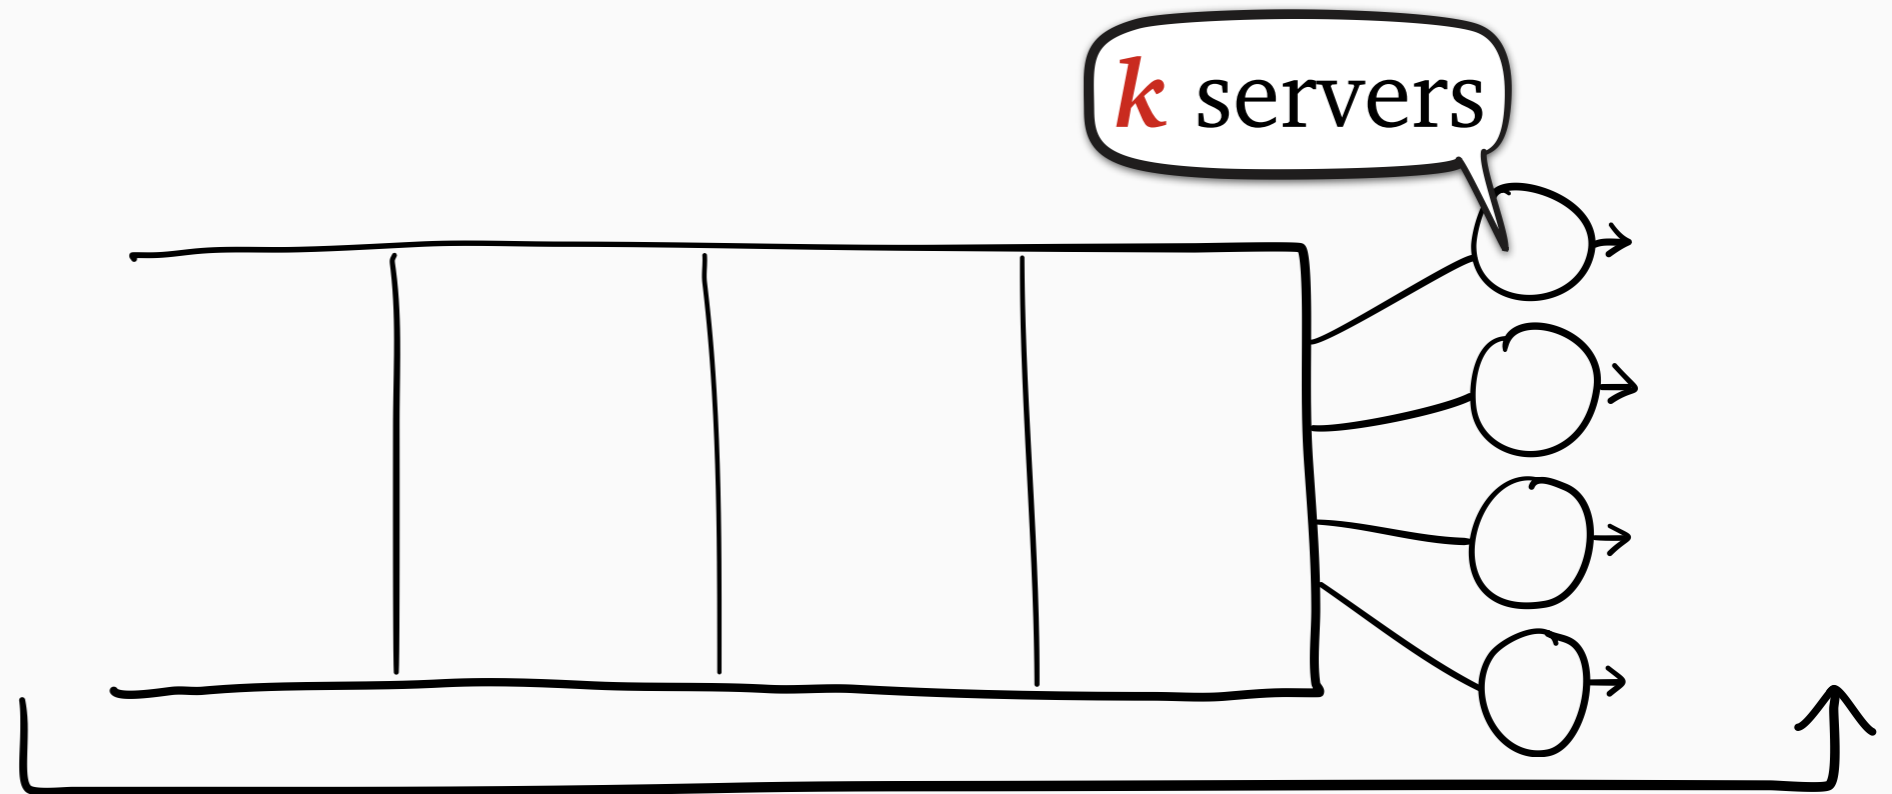

online arrivals



$T =$ response time

Multiserver scheduling



online arrivals



$T =$ response time

Multiserver scheduling

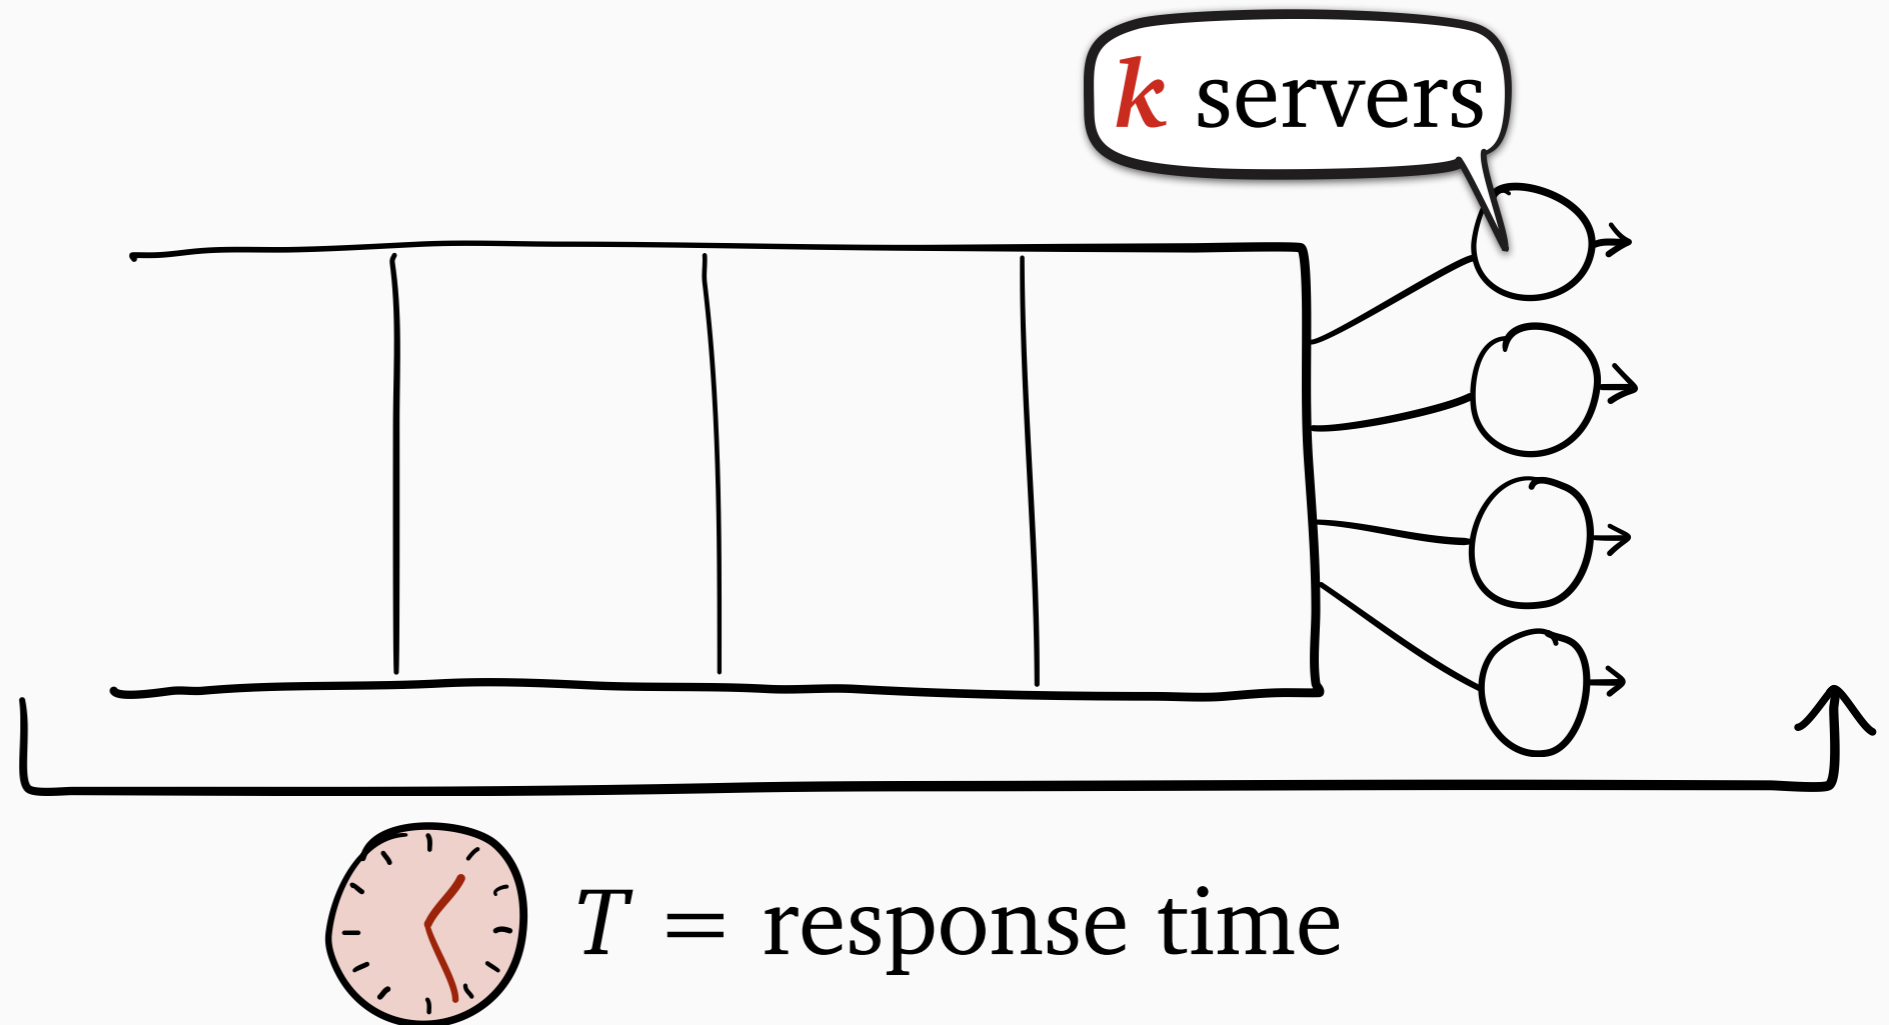

online arrivals



SRPT-1 (**single-server**): serves job of least remaining size

Multiserver scheduling

online arrivals



SRPT-1 (**single-server**): serves job of least remaining size

SRPT- k (**multiserver**): serves k jobs of least remaining size

How good is SRPT-*k*?

How good is SRPT-*k*?

TCS [Leonardi & Raz, 2007]: not great, but best possible

How good is SRPT- k ?

TCS [Leonardi & Raz, 2007]: not great, but best possible

Theorem: competitive ratio of SRPT- k vs. offline OPT- k is

$$\frac{\mathbb{E}[T_{\text{SRPT-}k}]}{\mathbb{E}[T_{\text{OPT-}k}]} \leq O\left(\min\left\{\log \frac{\# \text{ jobs}}{k}, \log \frac{\text{max size}}{\text{min size}}\right\}\right)$$

with matching lower bound

How good is SRPT-*k*?

TCS [Leonardi & Raz, 2007]: not great, but best possible

Theorem: competitive ratio of SRPT-*k* vs. offline OPT-*k* is

$$\frac{\mathbb{E}[T_{\text{SRPT-}k}]}{\mathbb{E}[T_{\text{OPT-}k}]} \leq O\left(\min\left\{\log \frac{\# \text{ jobs}}{k}, \log \frac{\text{max size}}{\text{min size}}\right\}\right)$$

with matching lower bound

How good is SRPT-*k*?

TCS [Leonardi & Raz, 2007]: not great, but best possible

Theorem: competitive ratio of SRPT-*k* vs. offline OPT-*k* is

$$\frac{E[T_{\text{SRPT-}k}]}{E[T_{\text{OPT-}k}]} \leq O\left(\min\left\{\log \frac{\# \text{ jobs}}{k}, \log \frac{\text{max size}}{\text{min size}}\right\}\right)$$

with matching lower bound

How good is SRPT-*k*?

TCS [Leonardi & Raz, 2007]: not great, but best possible

Theorem: competitive ratio of SRPT-*k* vs. offline OPT-*k* is

$$\frac{E[T_{\text{SRPT-}k}]}{E[T_{\text{OPT-}k}]} \leq O\left(\min\left\{\log \frac{\# \text{ jobs}}{k}, \log \frac{\text{max size}}{\text{min size}}\right\}\right)$$

with matching lower bound 

How good is SRPT-*k*?

TCS [Leonardi & Raz, 2007]: not great, but best possible

Theorem: competitive ratio of SRPT-*k* vs. offline OPT-*k* is

$$\frac{E[T_{\text{SRPT-}k}]}{E[T_{\text{OPT-}k}]} \leq O\left(\min\left\{\log \frac{\# \text{ jobs}}{k}, \log \frac{\text{max size}}{\text{min size}}\right\}\right)$$

with matching lower bound

How good is SRPT-*k*?

TCS [Leonardi & Raz, 2007]: not great, but best possible

Theorem: competitive ratio of SRPT-*k* vs. offline OPT-*k* is

$$\frac{E[T_{\text{SRPT-}k}]}{E[T_{\text{OPT-}k}]} \leq O\left(\min\left\{\log \frac{\# \text{ jobs}}{k}, \log \frac{\text{max size}}{\text{min size}}\right\}\right)$$


The equation includes three callout icons: a question mark above the first log term, a warning sign below the *k* in the denominator of the first log term, and a warning sign below the second log term.

with matching lower bound 

Queueing: decades-old open problem!


Queueing analysis of SRPT-1

[Schrage & Miller, 1966]



Queueing analysis of SRPT-1


[Schrage & Miller, 1966]




“tagged job”

Queueing analysis of SRPT-1

[Schrage & Miller, 1966]



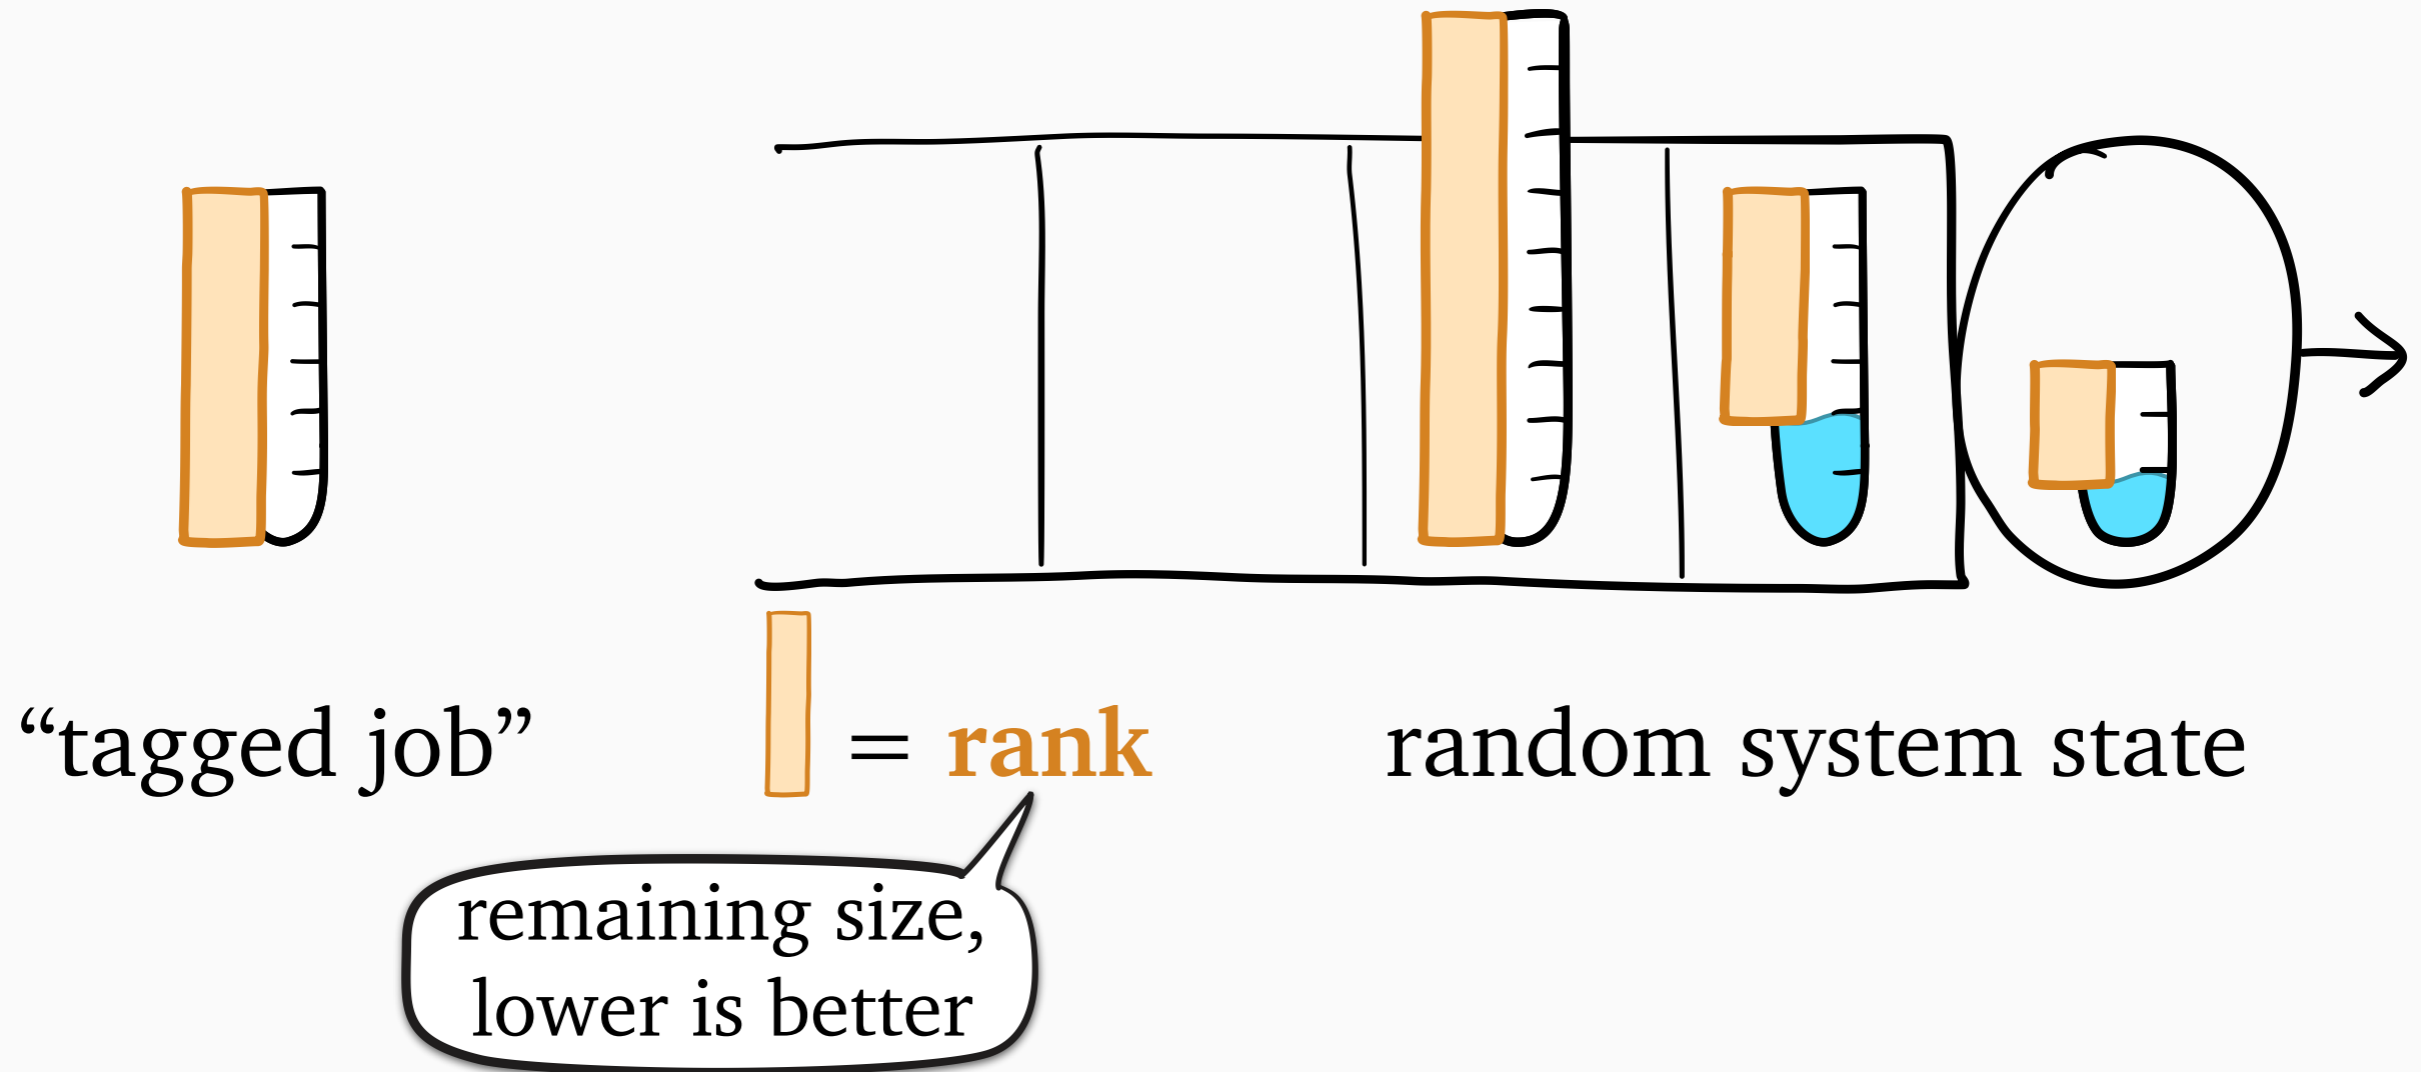
“tagged job”



random system state


Queueing analysis of SRPT-1

[Schrage & Miller, 1966]




Queueing analysis of SRPT-1

[Schrage & Miller, 1966]



Queueing analysis of SRPT-1

[Schrage & Miller, 1966]




Key quantity:


$W(r)$ = “ r -work” = work relevant to job of rank r

Queueing analysis of SRPT-1

[Schrage & Miller, 1966]



remaining size,
lower is better




Key quantity:


$W(r)$ = " r -work" = work relevant to job of rank r

Queueing analysis of SRPT-1

[Schrage & Miller, 1966]



“tagged job”

 = rank


ranked queue state

remaining size,
lower is better


Key quantity:

$W(r)$ = “ r -work” = work relevant to job of rank r

Single-server system




Multiserver system




Single-server system


server is “choke point”




Multiserver system



Single-server system




server is “choke point”




rank ordering absolute


Multiserver system




Single-server system



server is “choke point”




rank ordering absolute




observed r -work determines T


Multiserver system




Single-server system



server is “choke point”

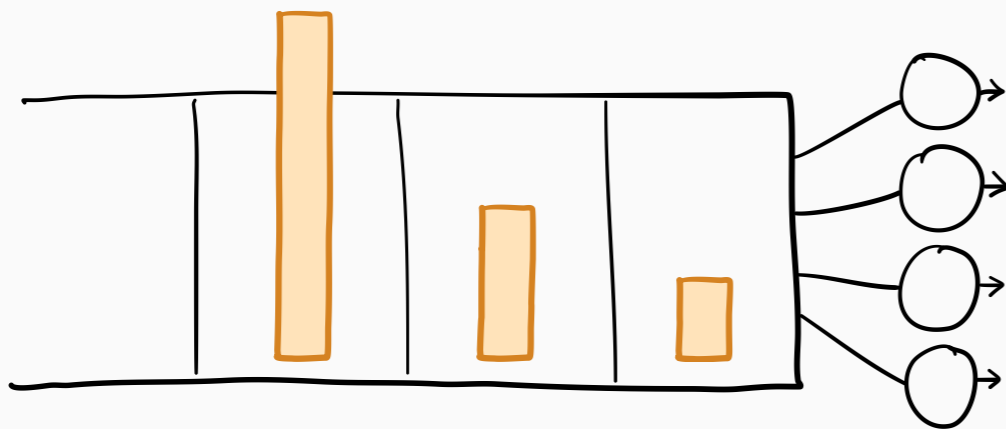


rank ordering absolute




observed *r*-work determines *T*

Multiserver system




no single “choke point”


Single-server system



server is “choke point”




rank ordering absolute



observed r -work determines T

Multiserver system

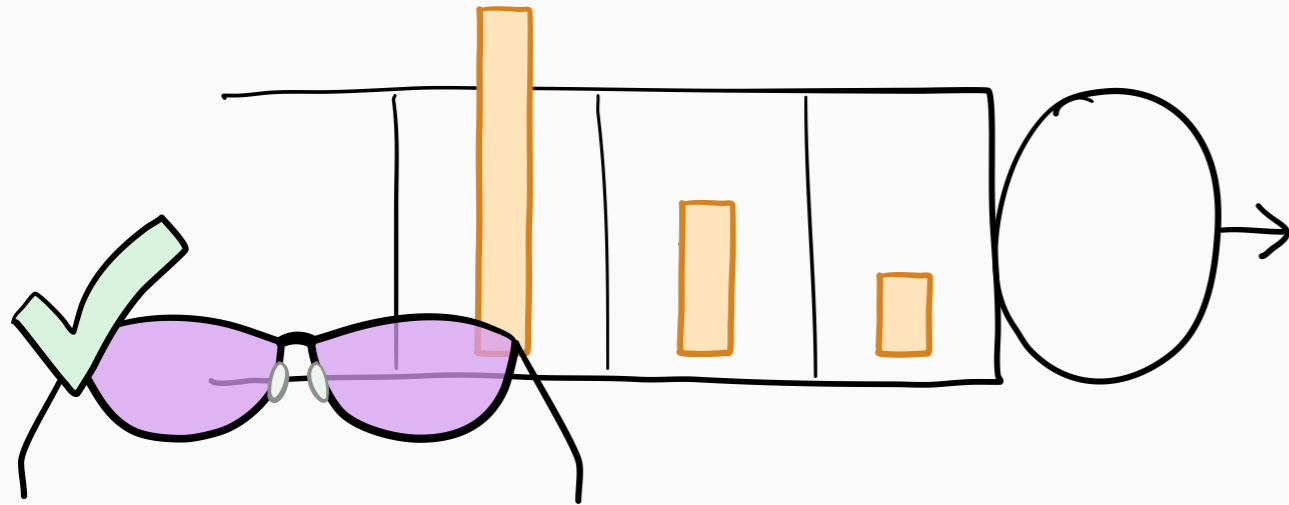


no single “choke point”

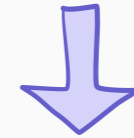


rank ordering *not* absolute

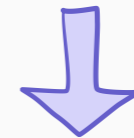
Single-server system



server is “choke point”




rank ordering absolute



observed r -work determines T

Multiserver system




no single “choke point”



rank ordering *not* absolute





observed r -work *not enough!*



r-work $W(r)$




tagged job
approach




response time T

single-server only (mostly)

tagged job
approach





r -work $W(r)$



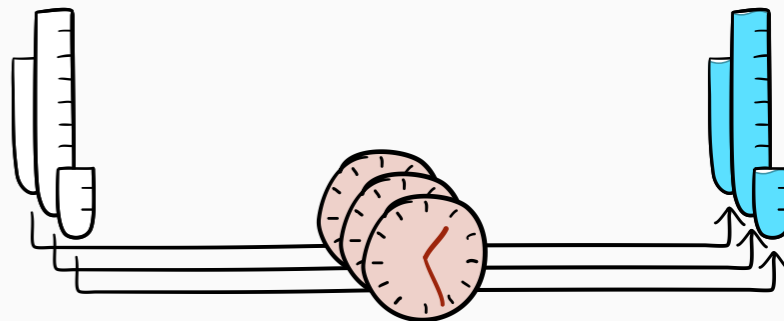
response time T

single-server only (mostly)

tagged job
approach




r -work $W(r)$



response time T





Little's law




number of jobs N

single-server only (mostly)

tagged job
approach





r -work $W(r)$



response time T

$$E[N] = \lambda E[T]$$

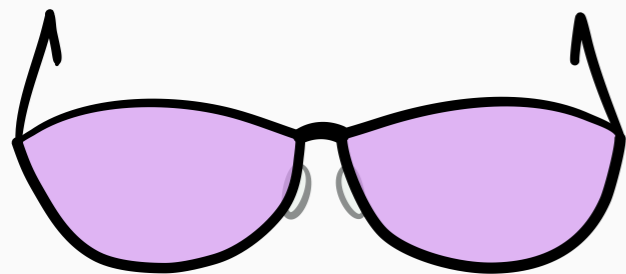
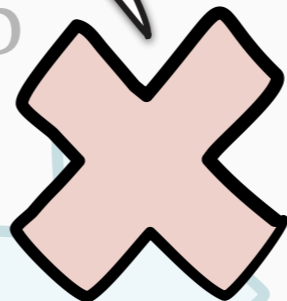
Little's law



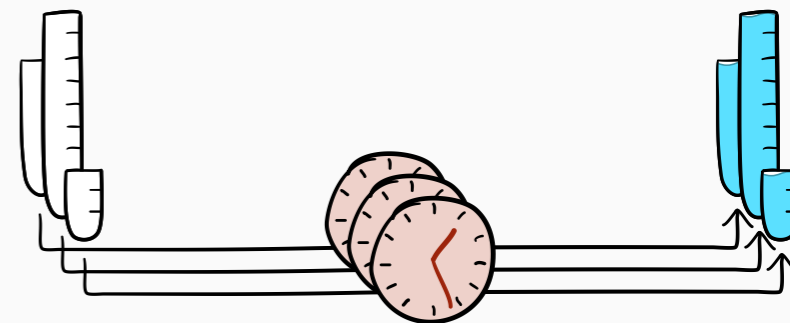
number of jobs N

single-server only (mostly)

tagged job approach



r -work $W(r)$

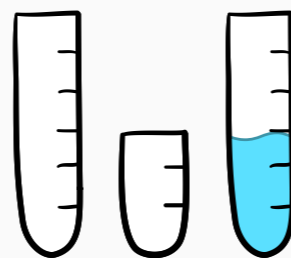


response time T

$$E[N] = \lambda E[T]$$

Little's law


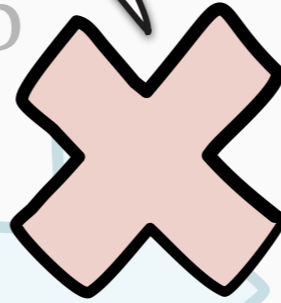
any number of servers




number of jobs N

single-server only (mostly)

tagged job approach



r -work $W(r)$



response time T




WINE



$$E[N] = \lambda E[T]$$

Little's law



any number of servers




number of jobs N

single-server only (mostly)

tagged job approach



r -work $W(r)$





response time T



WINE

any number of servers



number of jobs N

$$E[N] = \lambda E[T]$$


Little's law

any number of servers





WINE


Work Integral Number Equality



r -work $W(r)$




number of jobs N





WINE

Work Integral Number Equality



r-work $W(r)$



number of jobs N




What is *r*-work?





WINE

Work Integral Number Equality



r -work $W(r)$



number of jobs N



What is r -work?




How do we get number of jobs from r -work?





WINE

Work Integral Number Equality



r-work $W(r)$



number of jobs N

? What is *r*-work?

? How do we get number of jobs from *r*-work?

? How do we analyze *r*-work?

Defining r -work

$W(r)$ = work relevant to **rank** r

Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

$w_x(r)$ = r -work of *single job* of rem. size x




Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

$w_x(r)$ = r -work of *single job* of rem. size x




Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

$w_x(r)$ = r -work of *single job* of rem. size x




Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \end{cases}$




Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \end{cases}$




Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \end{cases}$




Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r

$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$



Defining r -work

for SRPT

$W(r)$ = work relevant to rank r

$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$



Defining r -work

for SRPT

$W(r)$ = work relevant to **rank** r
= total r -work of all jobs

$w_x(r)$ = r -work of *single job* of rem. size x = $\begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$



From r -work to number of jobs N

From r -work to number of jobs N

Goal: integral = N

$W(r)$



From r -work to number of jobs N

Goal: integral = N

$W(r)$



Suffices: integral = 1

$w_x(r)$




$$w_x(r) = r\text{-work of job of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

From r -work to number of jobs N


Goal: integral = N

$W(r)$



Suffices: integral = 1


$w_x(r)$




$$w_x(r) = r\text{-work of job of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

From r -work to number of jobs N

Goal: integral = N



Suffices: integral = 1




$$w_x(r) = r\text{-work of job of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

From r -work to number of jobs N


Goal: integral = N

$W(r)$



Suffices: integral = 1


$w_x(r)$



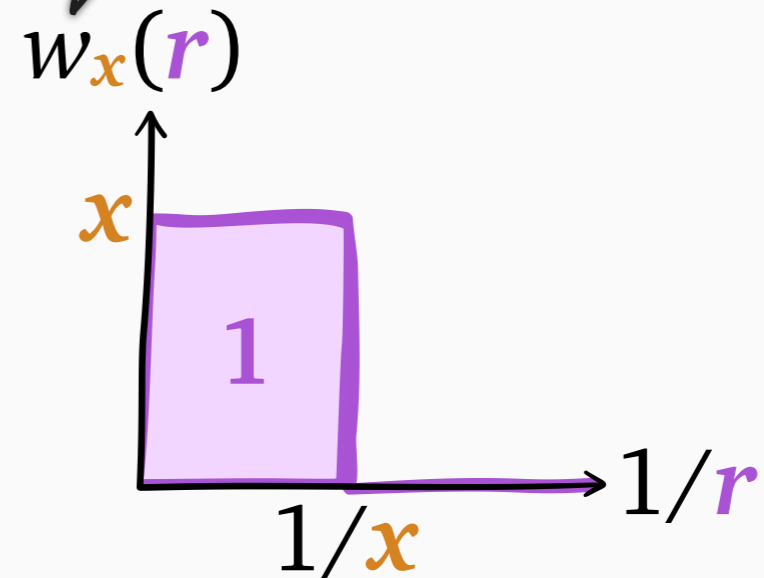
$$w_x(r) = r\text{-work of job of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

From r -work to number of jobs N

Goal: integral = N




Suffices: integral = 1




$$w_x(r) = r\text{-work of job of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

From r -work to number of jobs N

Goal: integral = N




Suffices: integral = 1




$$w_x(r) = r\text{-work of job of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

From r -work to number of jobs N

Goal: integral = N



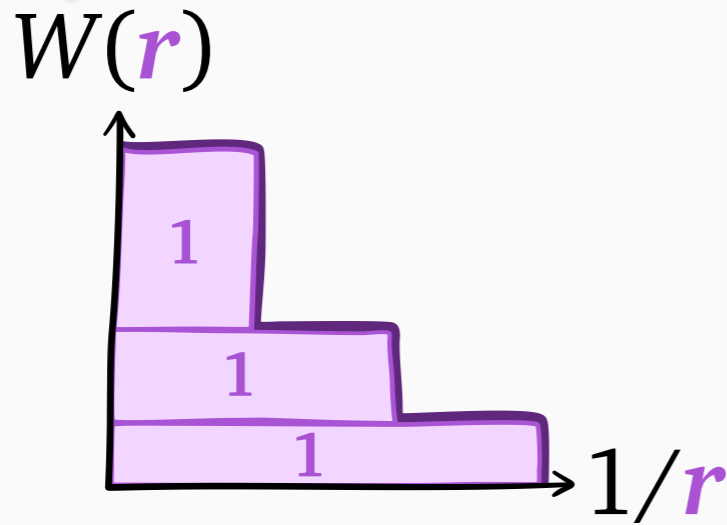
Suffices: integral = 1



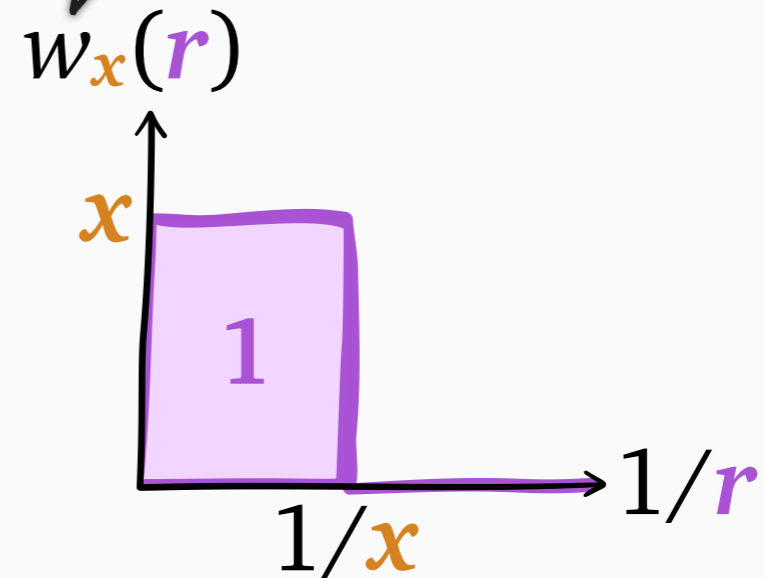
$$w_x(r) = r\text{-work of job of rem. size } x = \begin{cases} 0 & \text{if } r < x \\ x & \text{if } r \geq x \end{cases}$$

From r -work to number of jobs N

Goal: integral = N



Suffices: integral = 1




Theorem:


$$N = \int_0^{\infty} \frac{W(r)}{r^2} dr$$

From r -work to number of jobs N

Goal: integral = N




Suffices: integral = 1



NEW!

Theorem:


$$N = \int_0^{\infty} \frac{W(r)}{r^2} dr$$



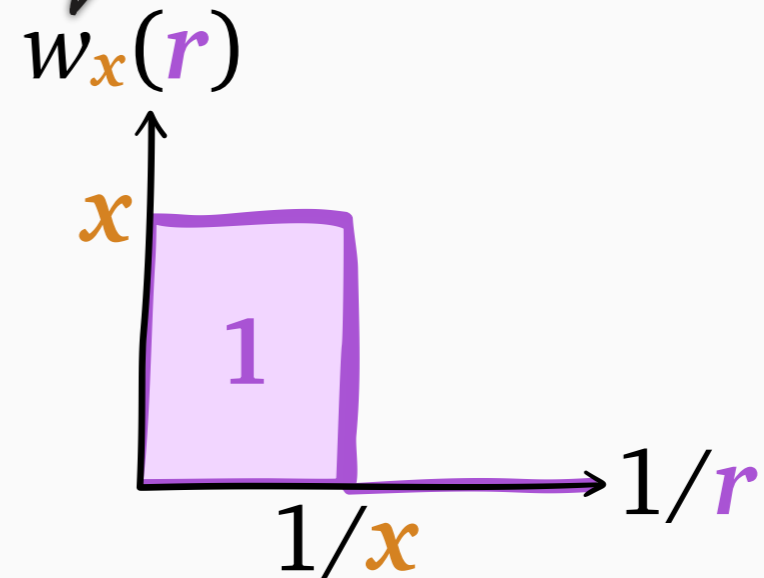
WINE

From r -work to number of jobs N

Goal: integral = N



Suffices: integral = 1



NEW!

Theorem:

$$N = \int_0^{\infty} \frac{W(r)}{r^2} dr$$


uses **rank** = rem. size




WINE

From r -work to number of jobs N

Goal: integral = N



Suffices: integral = 1




NEW!

Theorem: under *any* policy,


$$N = \int_0^{\infty} \frac{W(r)}{r^2} dr$$

uses **rank** = rem. size

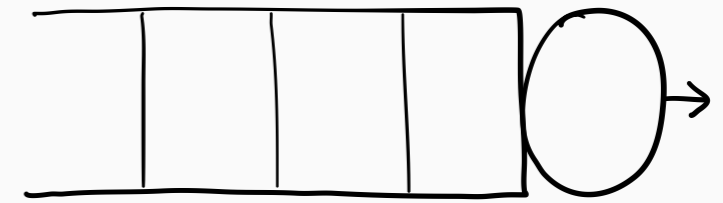


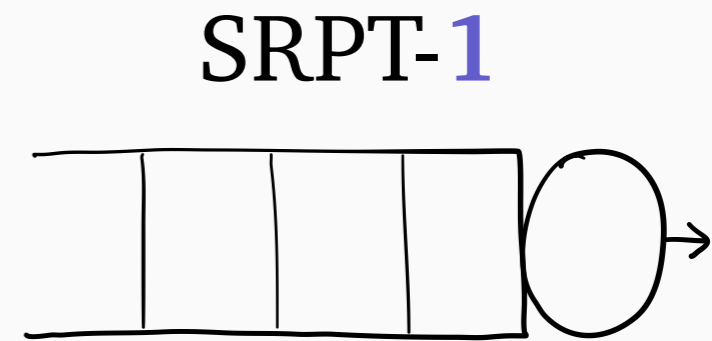
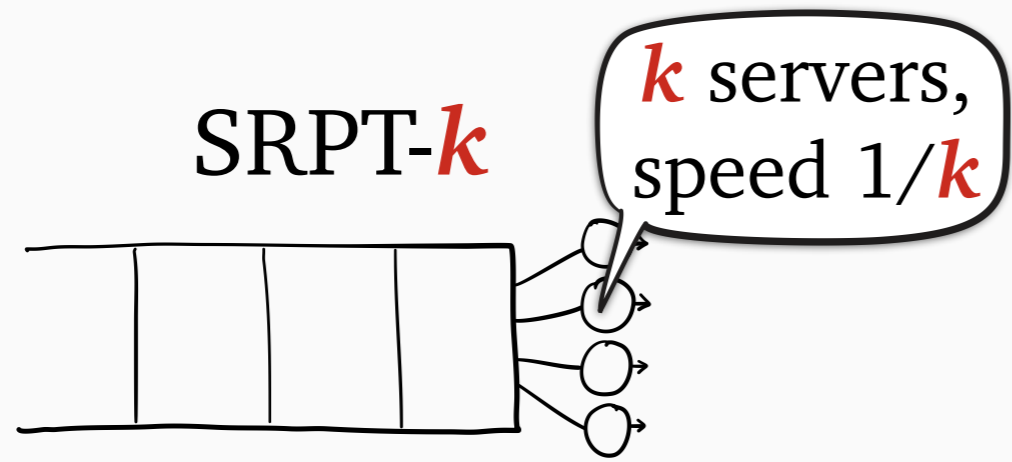
WINE

SRPT-*k*




SRPT-1



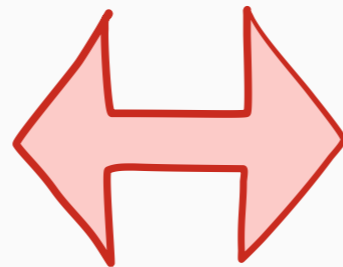



SRPT- k




k servers,
speed $1/k$

SRPT-1




SRPT- k




k servers,
speed $1/k$

SRPT-1




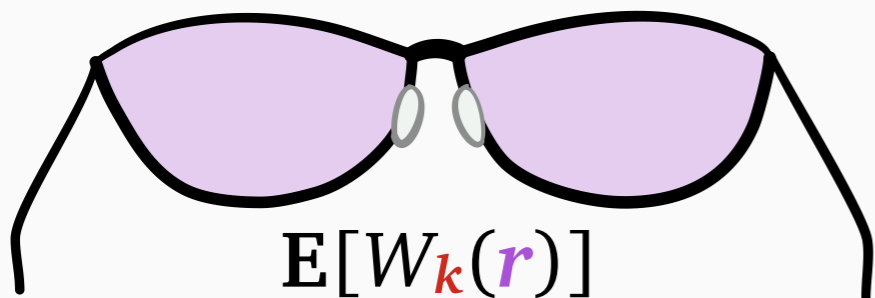
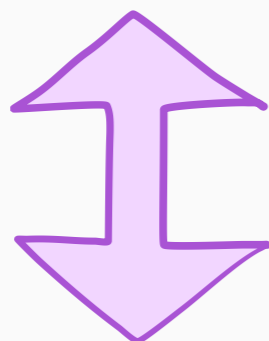




SRPT- k

k servers,
speed $1/k$

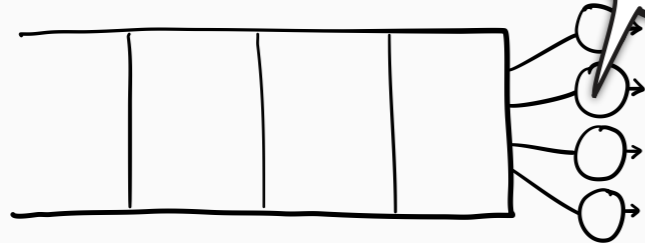


SRPT-1





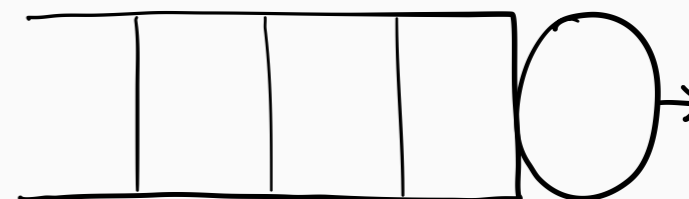


SRPT- k

k servers,
speed $1/k$




SRPT-1






SRPT- k

k servers,
speed $1/k$



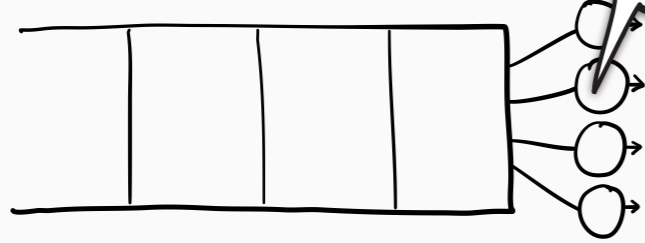
SRPT-1





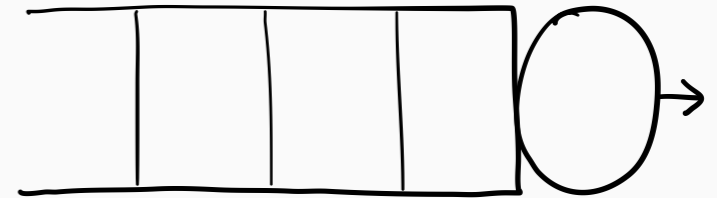
Lemma: r -work
decomposition

SRPT- k

k servers,
speed $1/k$




SRPT-1



Lemma: r -work
decomposition



SRPT- k





k servers,
speed $1/k$

beats OPT- k

SRPT-1



$E[W_k(r)]$



$E[W_1(r)]$

Lemma: r -work
decomposition

Additive bounds for SRPT- k

Lemma:

$$\mathbf{E}[N_k] = \mathbf{E}[N_1] + \int_0^\infty \frac{\mathbf{E}[W_k(r)] - \mathbf{E}[W_1(r)]}{r^2} dr$$

Additive bounds for SRPT- k

Lemma:

$$\mathbf{E}[N_k] = \mathbf{E}[N_1] + \int_0^\infty \frac{\mathbf{E}[W_k(r)] - \mathbf{E}[W_1(r)]}{r^2} dr$$

Lemma: in worst case,

$$W_k(r) \leq W_1(r) + kr$$

Additive bounds for SRPT- k

Lemma:

$$\mathbf{E}[N_k] = \mathbf{E}[N_1] + \int_0^\infty \frac{\mathbf{E}[W_k(r)] - \mathbf{E}[W_1(r)]}{r^2} dr$$

Lemma: in worst case,

$$W_k(r) \leq W_1(r) + kr$$



TCS

$$\mathbf{E}[N_k] \leq \mathbf{E}[N_1] + 2k + k \log \frac{\text{max size}}{\text{min size}}$$

Additive bounds for SRPT- k

Lemma:

$$\mathbf{E}[N_k] = \mathbf{E}[N_1] + \int_0^\infty \frac{\mathbf{E}[W_k(r)] - \mathbf{E}[W_1(r)]}{r^2} dr$$

Lemma: in worst case,

$$W_k(r) \leq W_1(r) + kr$$



TCS



Queueing

$$\mathbf{E}[N_k] \leq \mathbf{E}[N_1] + 2k + k \log \frac{\text{max size}}{\text{min size}}$$

$$\mathbf{E}[N_k] \leq \mathbf{E}[N_1] + 4(k-1) \log \frac{1}{1-\rho}$$

Additive bounds for SRPT- k

Lemma:

$$\mathbf{E}[N_k] = \mathbf{E}[N_1] + \int_0^\infty \frac{\mathbf{E}[W_k(r)] - \mathbf{E}[W_1(r)]}{r^2} dr$$

Lemma: in worst case,

$$W_k(r) \leq W_1(r) + kr$$



TCS

$$\mathbf{E}[N_k] \leq \mathbf{E}[N_1] + 2k + k \log \frac{\text{max size}}{\text{min size}}$$



Queueing

$$\mathbf{E}[N_k] \leq \mathbf{E}[N_1] + 4(k-1) \log \frac{1}{1-\rho}$$

dominant term

Additive bounds for SRPT- k

Lemma:

$$\mathbf{E}[N_k] = \mathbf{E}[N_1] + \int_0^\infty \frac{\mathbf{E}[W_k(r)] - \mathbf{E}[W_1(r)]}{r^2} dr$$

Lemma: in worst case,

$$W_k(r) \leq W_1(r) + kr$$

uses stronger
stochastic lemma

TCS



Queueing

$$\mathbf{E}[N_k] \leq \mathbf{E}[N_1] + 2k + k \log \frac{\text{max size}}{\text{min size}}$$

$$\mathbf{E}[N_k] \leq \mathbf{E}[N_1] + 4(k-1) \log \frac{1}{1-\rho}$$

dominant term

Today's talk

scheduling with

multiple servers



TCS



Queueing



Today's talk

scheduling with

multiple servers

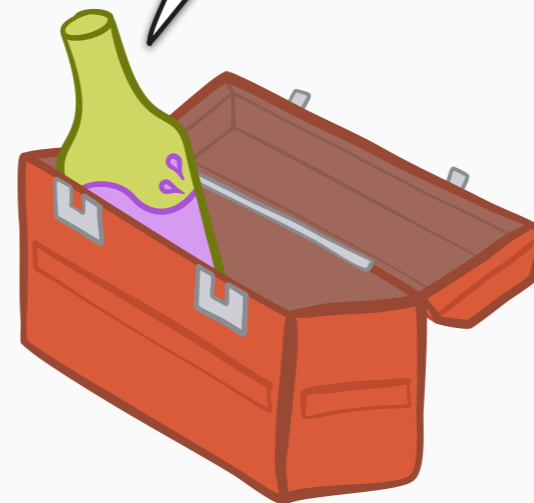


TCS



Queueing

$$\mathbf{E}[N] = \int_0^{\infty} \frac{\mathbf{E}[W(r)]}{r^2} dr$$



Today's talk

scheduling with

multiple servers




TCS



Queueing

$$\mathbf{E}[N] = \int_0^{\infty} \frac{\mathbf{E}[W(r)]}{r^2} dr$$



Today's talk

scheduling with

multiple servers



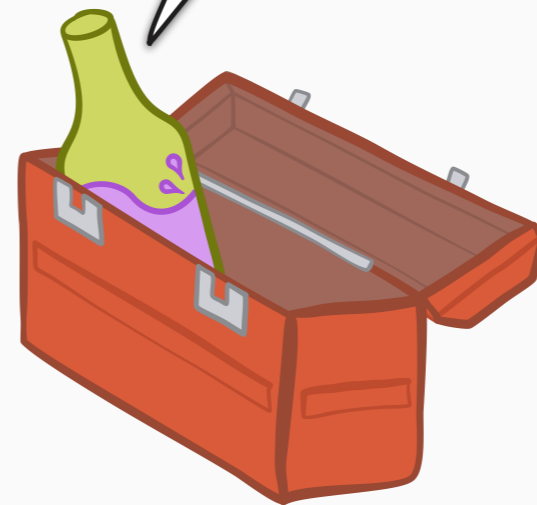
TCS



Queueing

- SRPT-*k* is good

$$E[N] = \int_0^{\infty} \frac{E[W(r)]}{r^2} dr$$



Today's talk

scheduling with

multiple servers




TCS



Queueing

- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

$$E[N] = \int_0^{\infty} \frac{E[W(r)]}{r^2} dr$$



Today's talk

scheduling with

multiple servers



TCS




Queueing

- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

TCS: RMLF

Queueing: Gittins

$$E[N] = \int_0^{\infty} \frac{E[W(r)]}{r^2} dr$$



Today's talk

scheduling with
multiple servers


 **TCS**   **Queueing**

- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

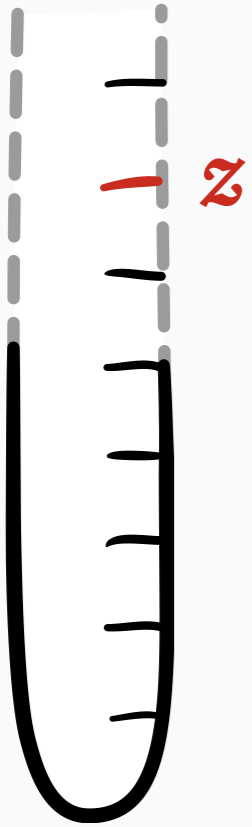
TCS: RMLF
Queueing: **Gittins**

scheduling with
noisy predictions


$$E[N] = \int_0^{\infty} \frac{E[W(r)]}{r^2} dr$$



Noisily predicted job sizes



Noisily predicted job sizes



Model: (β, α) -bounded noise

true size $s \Rightarrow$ **estimated** size $z \in [\beta s, \alpha s]$


Noisily predicted job sizes

below

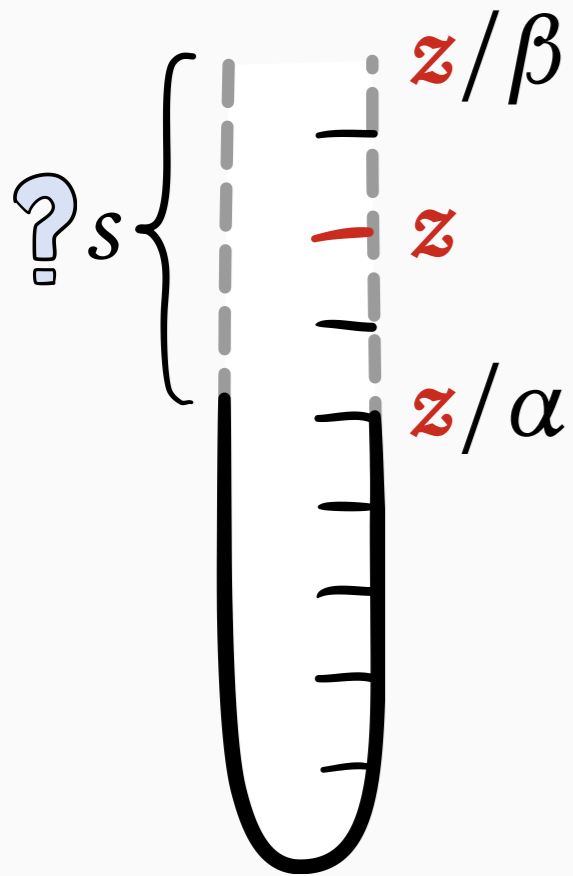
above

Model: (β, α) -bounded noise

true size $s \Rightarrow$ **estimated** size $z \in [\beta s, \alpha s]$



Noisily predicted job sizes




below above

Model: (β, α) -bounded noise

true size $s \Rightarrow$ **estimated** size $z \in [\beta s, \alpha s]$

Noisily predicted job sizes



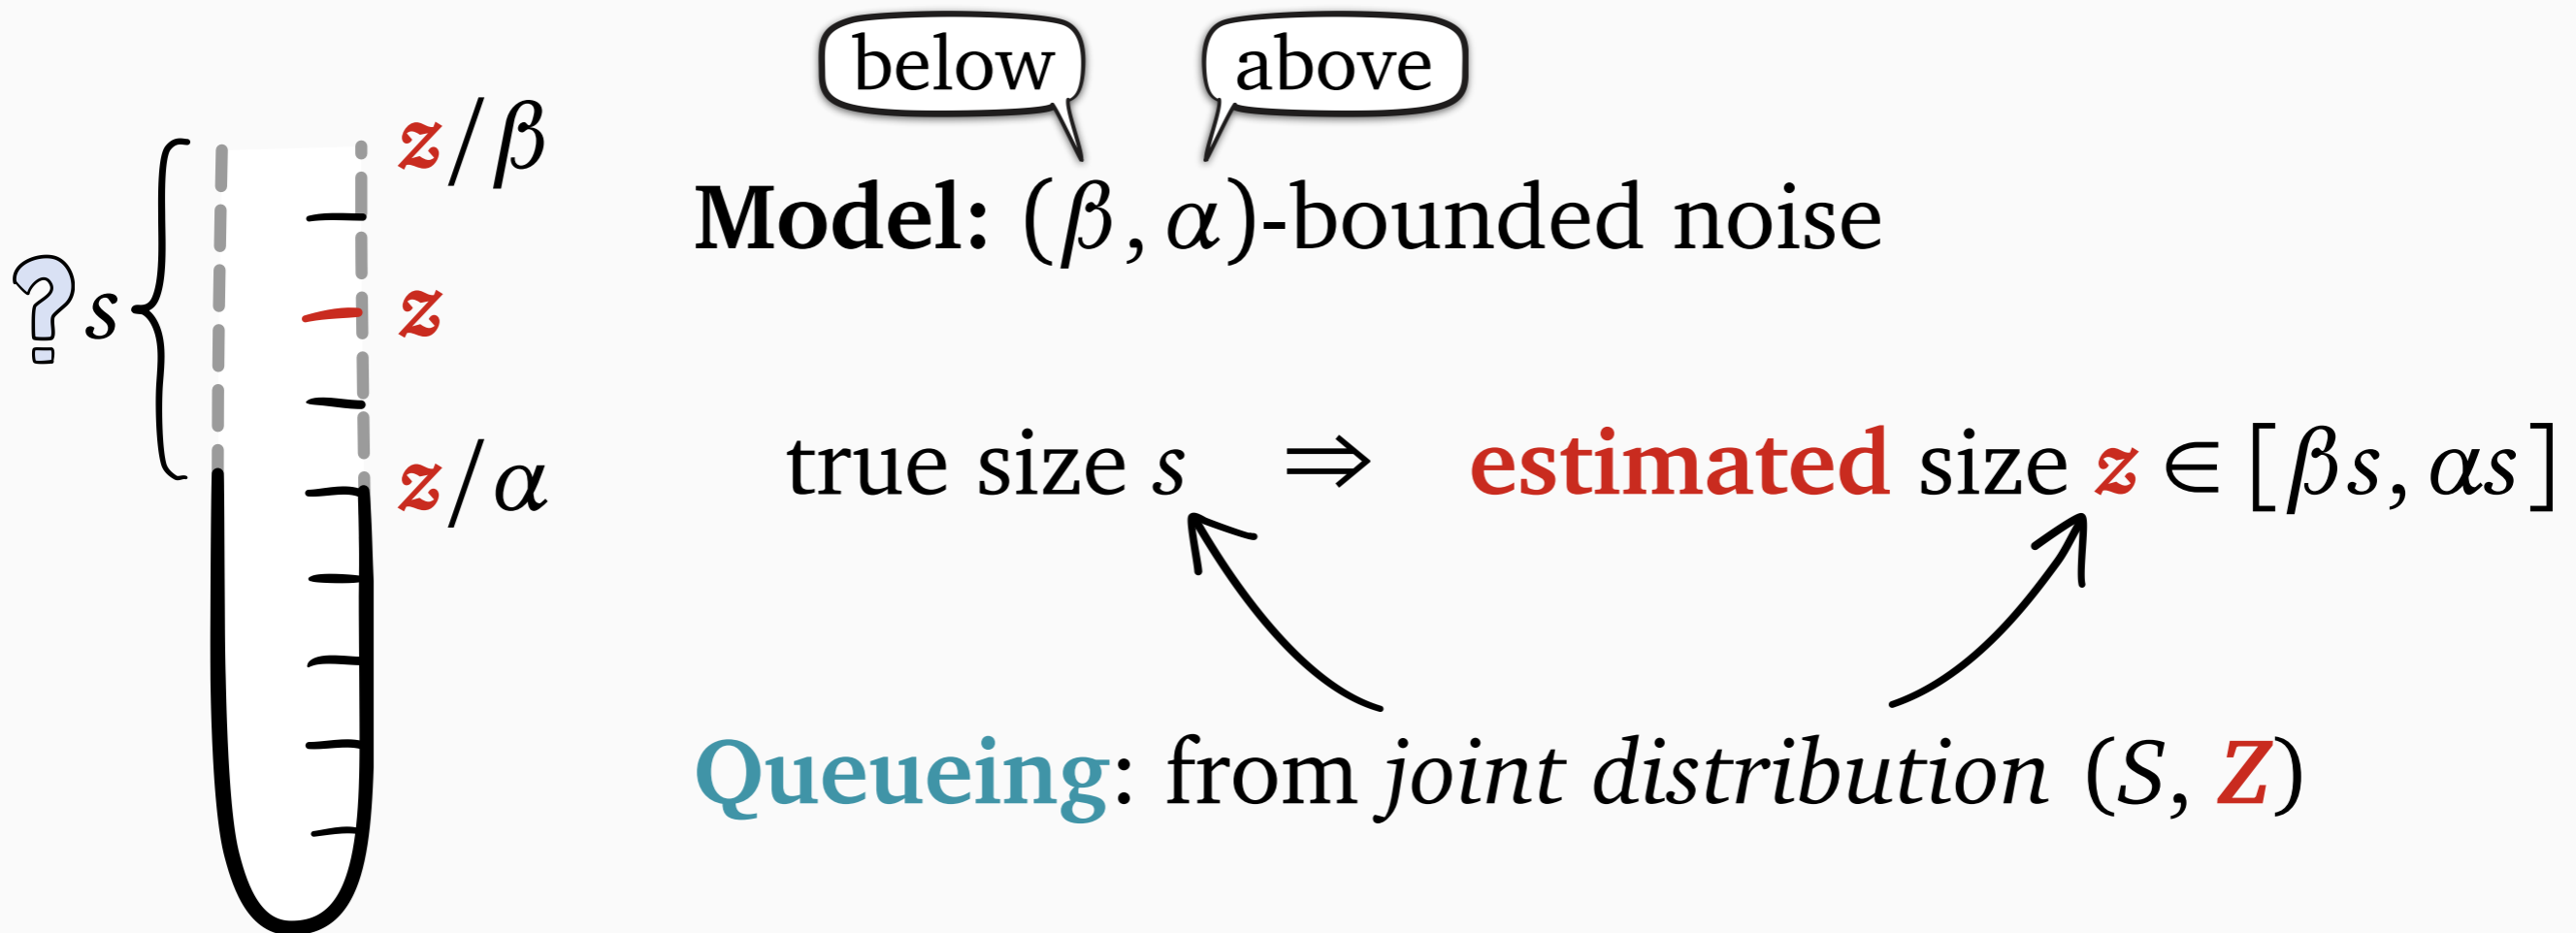
below above

Model: (β, α) -bounded noise

true size $s \Rightarrow$ **estimated** size $z \in [\beta s, \alpha s]$

Queueing: from *joint distribution* (S, Z)

Noisily predicted job sizes



Goal: design a policy with “good” $E[T]$ for

- *any* values of α, β
- *any* joint distribution (S, Z)

How well can we handle **noise**?

Definition: distortion is $\gamma = \frac{\alpha}{\beta}$

How well can we handle **noise**?

Definition: **distortion** is $\gamma = \frac{\alpha}{\beta}$

TCS [Azar, Leonardi, & Touitou, 2021 & 2022]:
pretty well, but need a sophisticated policy

$$\frac{\mathbf{E}[T_{\text{ZigZag}}]}{\mathbf{E}[T_{\text{SRPT}}]} \leq O(\gamma \log \gamma)$$

with nearly-matching $\Omega(\gamma)$ lower bound

How well can we handle **noise**?

Definition: distortion is $\gamma = \frac{\alpha}{\beta}$

TCS [Azar, Leonardi, & Touitou, 2021 & 2022]:
pretty well, but need a sophisticated policy

$$\frac{\mathbf{E}[T_{\text{ZigZag}}]}{\mathbf{E}[T_{\text{SRPT}}]} \leq O(\gamma \log \gamma)$$


with nearly-matching $\Omega(\gamma)$ lower bound



How well can we handle **noise**?

Definition: distortion is $\gamma = \frac{\alpha}{\beta}$

TCS [Azar, Leonardi, & Touitou, 2021 & 2022]:
pretty well, but need a sophisticated policy

$$\frac{\mathbf{E}[T_{\text{ZigZag}}]}{\mathbf{E}[T_{\text{SRPT}}]} \leq O(\gamma \log \gamma)$$



with nearly-matching $\Omega(\gamma)$ lower bound



How well can we handle **noise**?

Definition: distortion is $\gamma = \frac{\alpha}{\beta}$

TCS [Azar, Leonardi, & Touitou, 2021 & 2022]:
pretty well, but need a sophisticated policy


$$\frac{\mathbf{E}[T_{\text{ZigZag}}]}{\mathbf{E}[T_{\text{SRPT}}]} \leq O(\gamma \log \gamma)$$


with nearly-matching $\Omega(\gamma)$ lower bound



Queueing: can we do better with simpler policy?

Scheduling with **rank** functions




Scheduling with **rank** functions


rank




age a



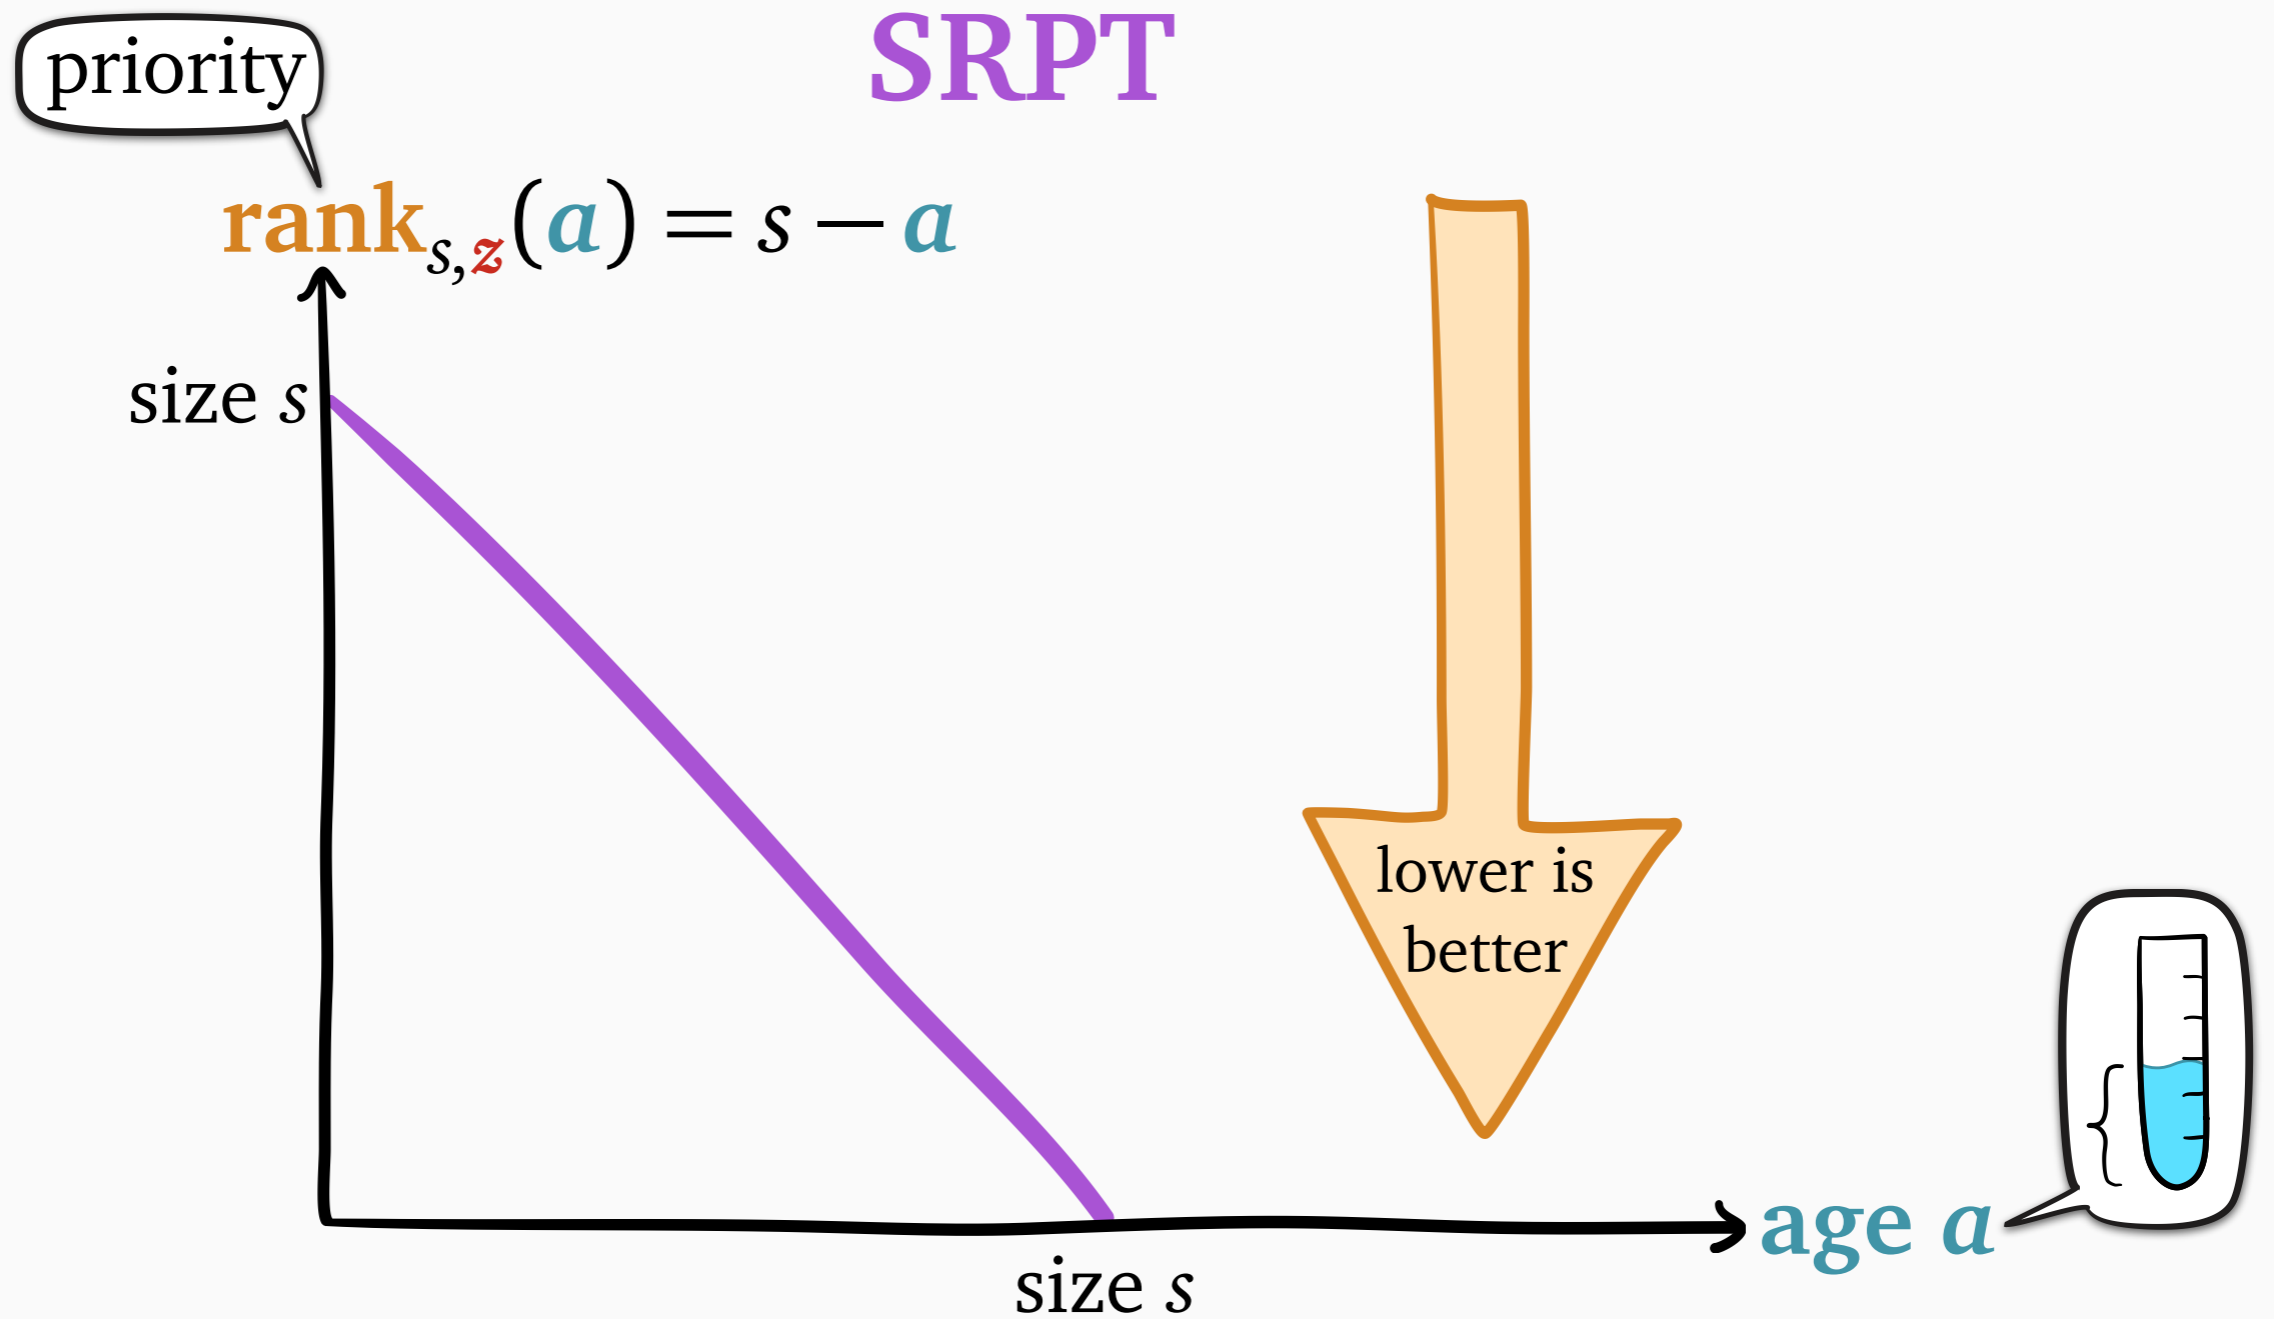
Scheduling with **rank** functions



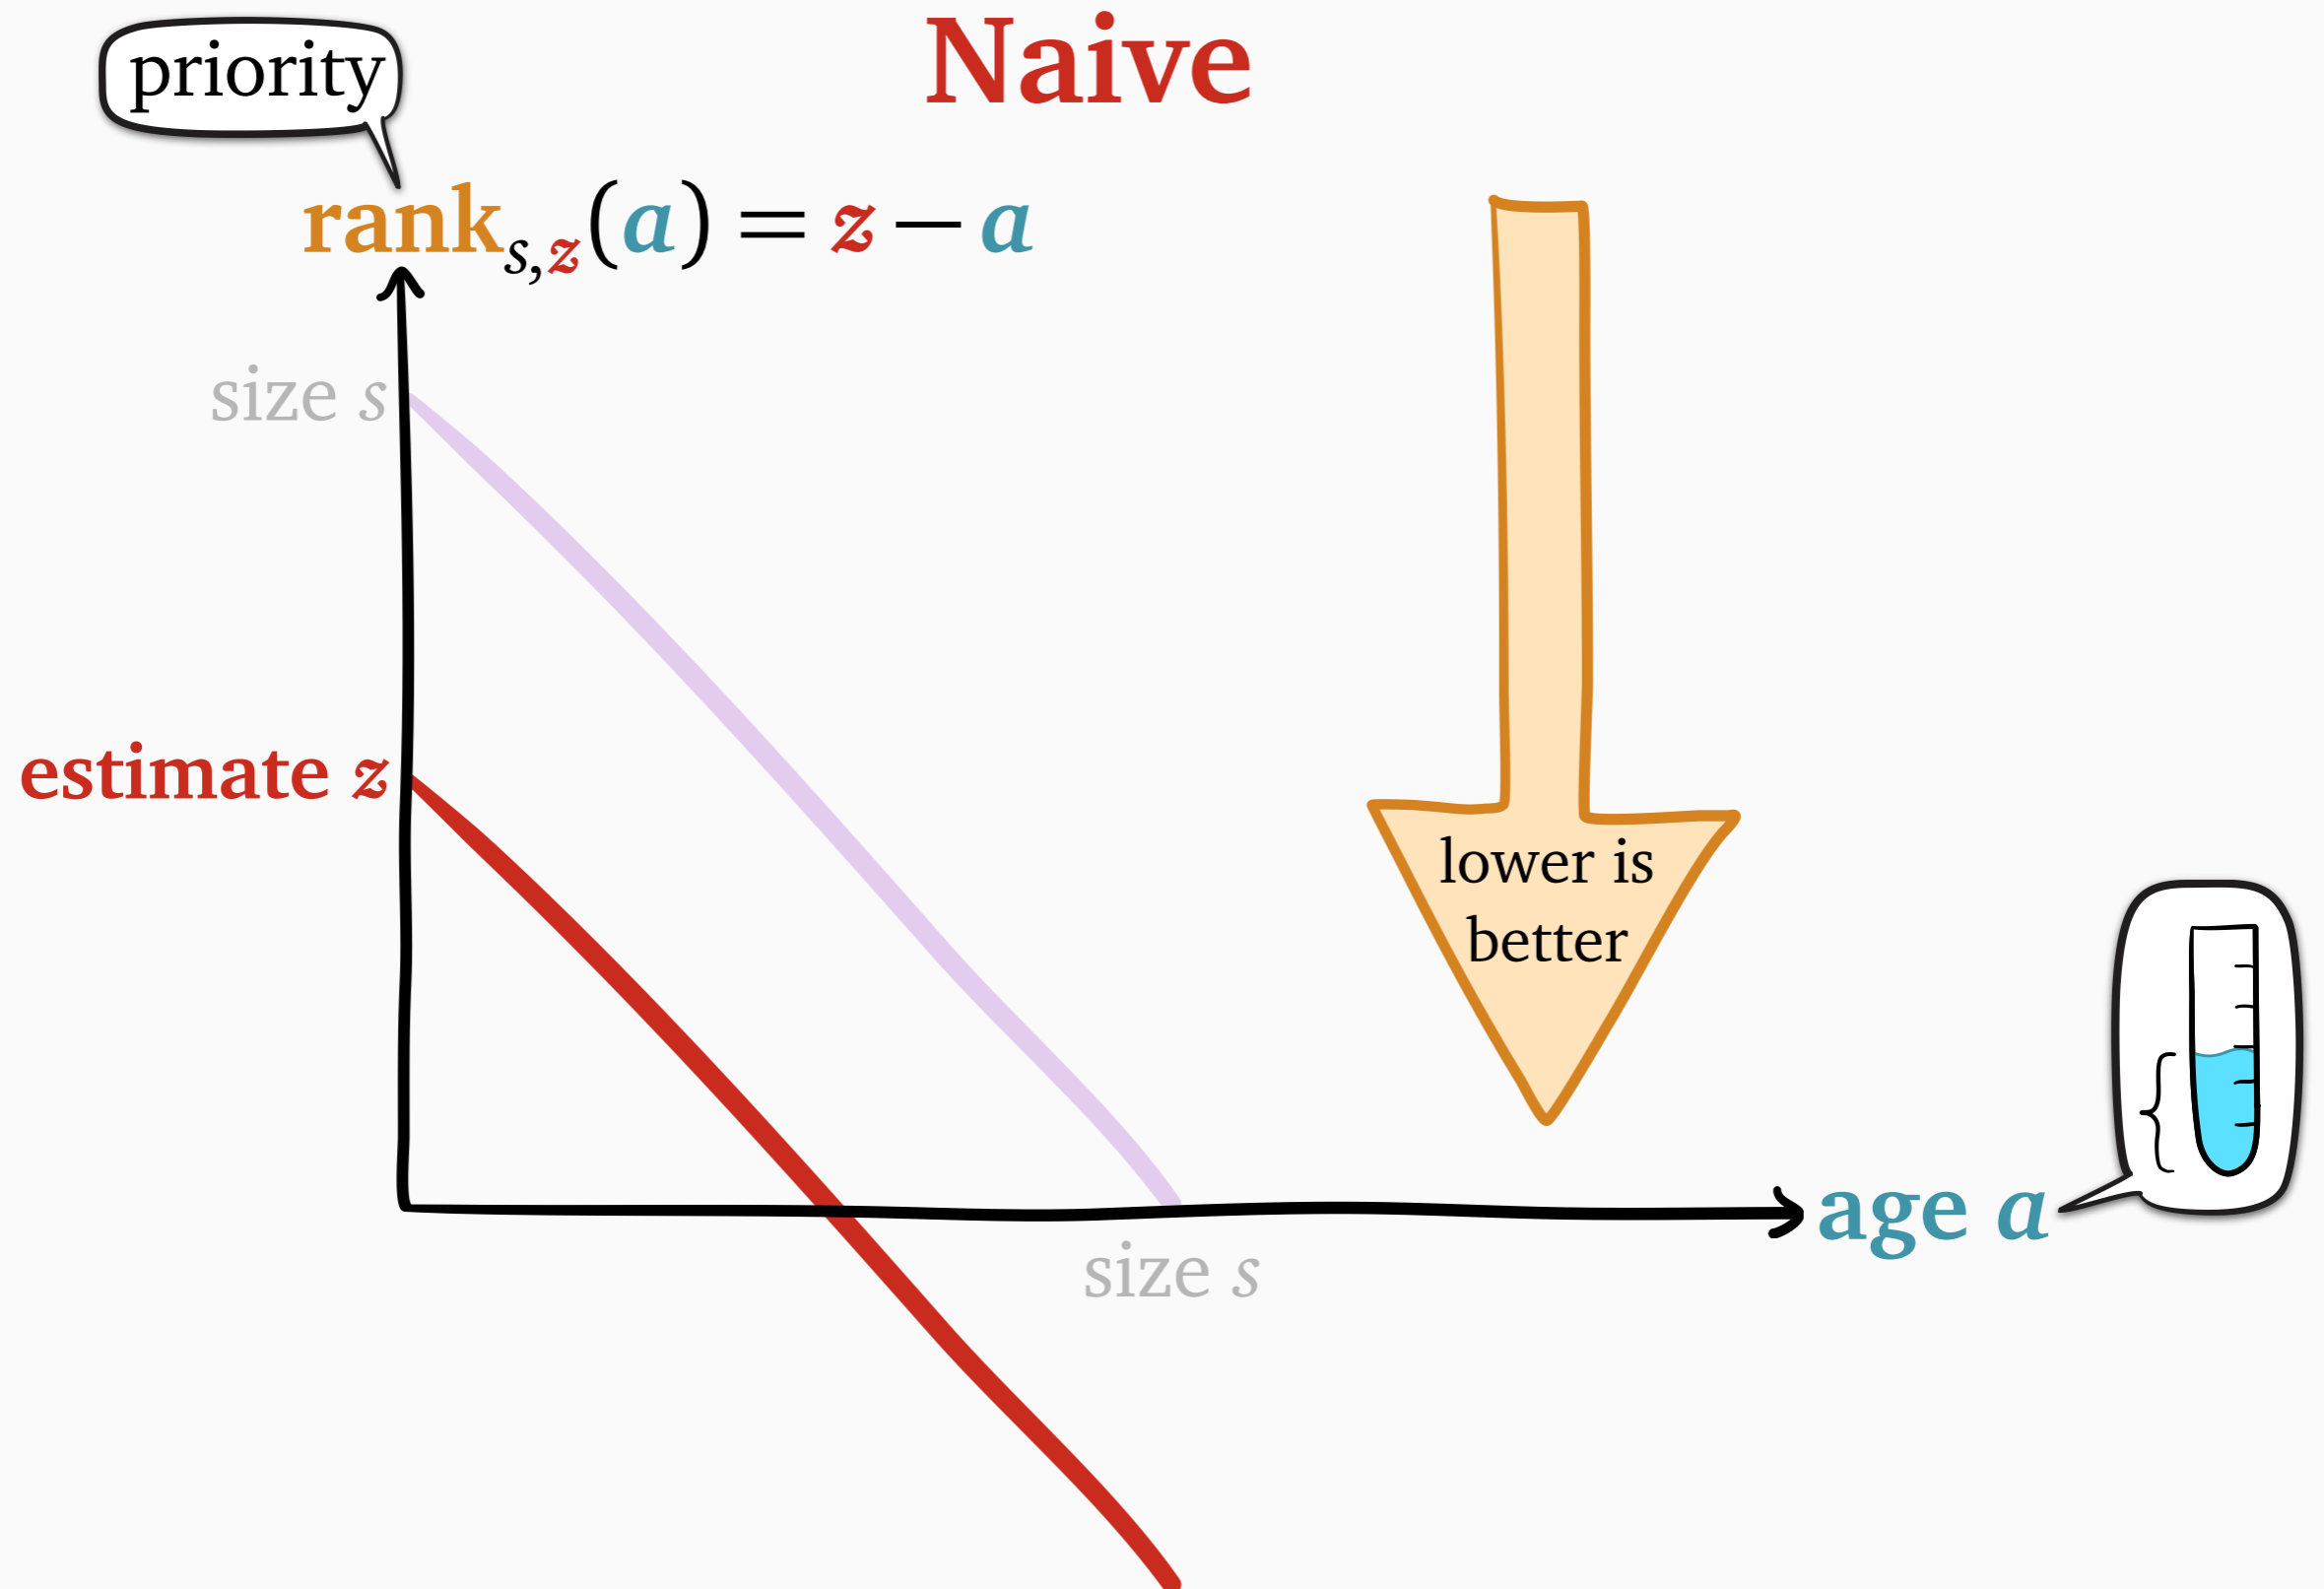
Scheduling with **rank** functions



Scheduling with **rank** functions




Scheduling with **rank** functions





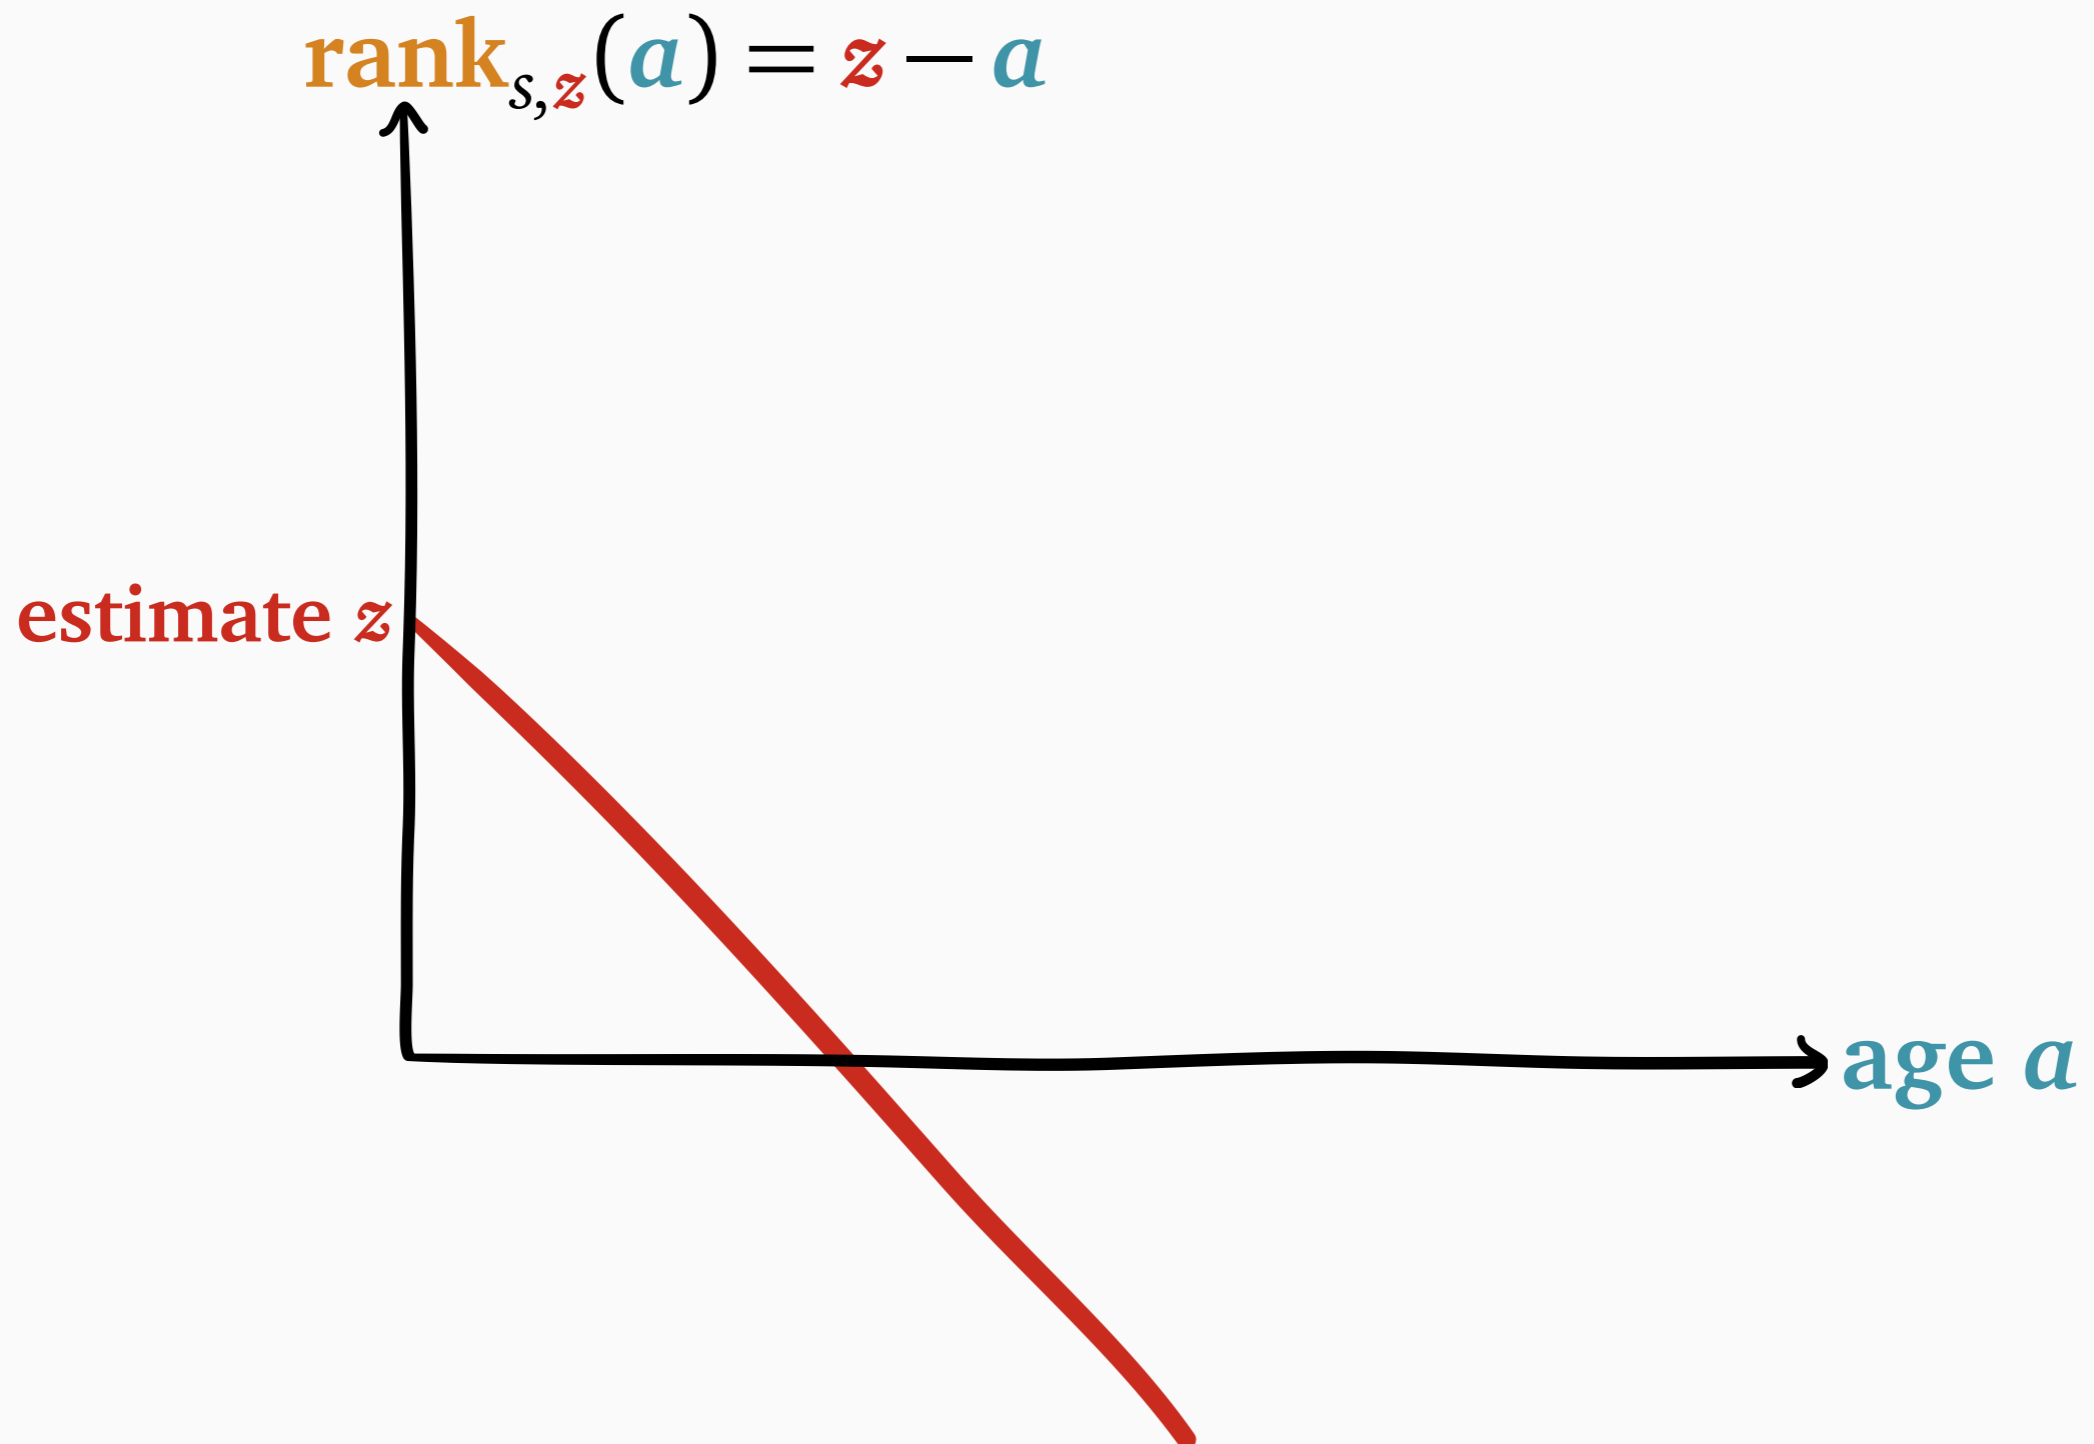
Policy design space:
rank functions

What's the right **rank** function?




What's the right **rank** function?

Naive





What's the right **rank** function?

Naive



What's the right **rank** function?




What's the right **rank** function?

Checkmark


$$\text{rank}_{s,z}(a) = |z - a|$$

estimate z



Not these...


Naive



What's the right **rank** function?


Checkmark

$$\text{rank}_{s,z}(a) = |z - a|$$



Not these...

Naive




What if $\beta < \frac{1}{2}$?

What's the right **rank** function?


Checkmark

$$\text{rank}_{s,z}(a) = |z - a|$$



Not these...


Naive



What if $\beta < \frac{1}{2}$?


What's the right **rank** function?

Checkmark



Not these...

Naive



What if $\beta < \frac{1}{2}$?

What's the right **rank** function?


Checkmark

$$\text{rank}_{s,z}(a) = |z - a|$$

new worst **rank** \Rightarrow
preemption likely




estimate z





What if $\beta < \frac{1}{2}$?

Not these...

Naive




What's the right **rank** function?



What's the right **rank** function?


Radical

$$\text{rank}_{s,z}(a) = \min\{|z - a|, z\}$$




Not these...

Naive




Checkmark



What's the right **rank** function?


Radical

$$\text{rank}_{s,z}(a) = \min\{|z - a|, z\}$$




Not these...

Naive




Checkmark



What's the right **rank** function?


Radical

$$\text{rank}_{s,z}(a) = \min\{|z - a|, z\}$$







Not these...

Naive



Checkmark








Theorem: in **queueing** model,

$$\frac{\mathbf{E}[T_{\text{Radical}}]}{\mathbf{E}[T_{\text{SRPT}}]} \leq C_{\alpha, \beta} \cdot \gamma$$

where

$$C_{\alpha, \beta} \leq 3.5$$

$$C_{\alpha, \beta} \rightarrow 1 \quad \text{as } \alpha, \beta \rightarrow 1$$







Theorem: in **queueing** model,






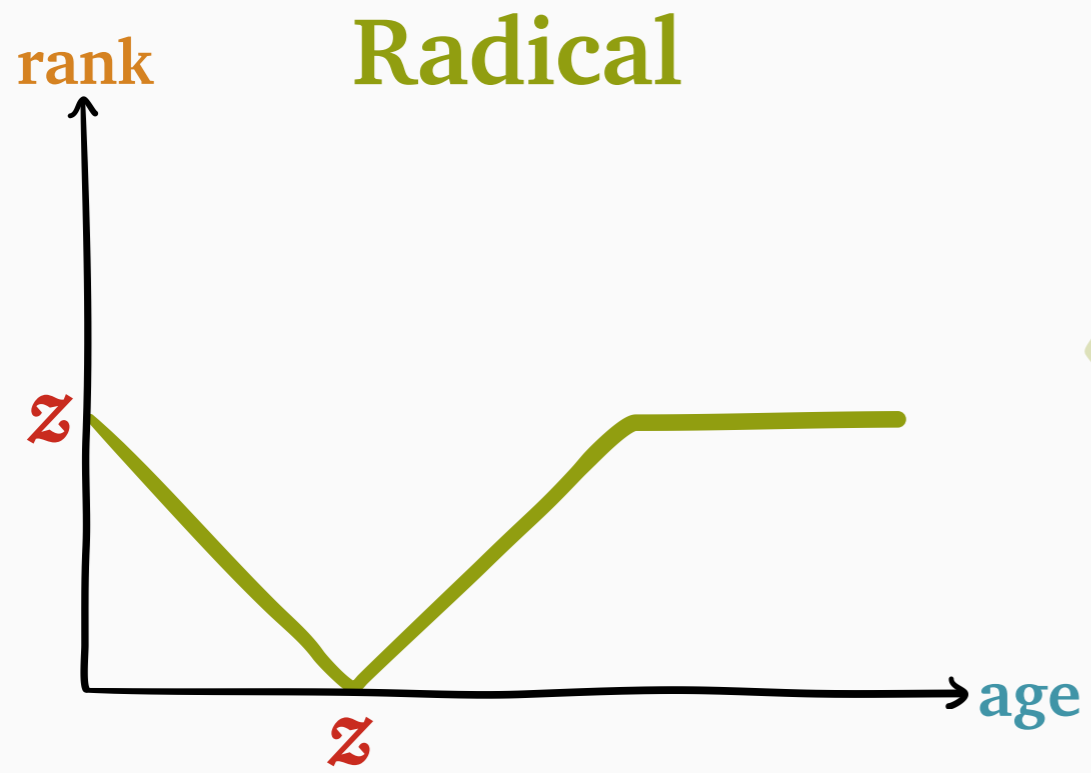
$$\frac{\mathbf{E}[T_{\text{Radical}}]}{\mathbf{E}[T_{\text{SRPT}}]} \leq C_{\alpha, \beta} \cdot \gamma$$





where

$$C_{\alpha, \beta} \leq 3.5$$


$$C_{\alpha, \beta} \rightarrow 1 \quad \text{as } \alpha, \beta \rightarrow 1$$












 **SOAP**





WINE




Schedule **O**rdered by **A**ge-based **P**riority




SOAP

Schedule **O**rdered by **A**ge-based **P**riority

stochastic arrival
process λ , (S , Z)




any **rank** function




SOAP

Schedule **O**rdered by **A**ge-based **P**riority


stochastic arrival
process λ , (S , Z)







any **rank** function



response time
distribution T









Scale




WINE





rank functions
close enough







SOAP

Scale




WINE






rank functions
close enough



SOAP



WINE

Comparing r -work

Lemma:

$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Comparing r -work

Lemma:

$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

Comparing r -work

Lemma:

$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

- **SRPT** minimizes mean r -work

Comparing r -work

Lemma:



$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

- **SRPT** minimizes mean r -work

Comparing r -work

Lemma:



$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

- **SRPT** minimizes mean r -work
- **Scale** minimizes mean *Scale*-flavored r -work

Comparing r -work


Lemma:



$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

- **SRPT** minimizes mean r -work
- **Scale** minimizes mean *Scale*-flavored r -work



filters using **Scale's rank**
instead of **SRPT's rank**

Comparing r -work


Lemma:



$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

- **SRPT** minimizes mean r -work
- **Scale** minimizes mean *Scale*-flavored r -work
- Under any policy,
 r -work \leq **Scale**-flavored αr -work $\leq \frac{\alpha}{\beta} r$ -work



filters using **Scale's rank**
instead of **SRPT's rank**

Comparing r -work

Lemma:

$$\mathbb{E}[W_{\text{SRPT}}(r)] \leq \mathbb{E}[W_{\text{Scale}}(r)] \leq \mathbb{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

- **SRPT** minimizes mean r -work
- **Scale** minimizes mean *Scale*-flavored r -work
- Under any policy,
 r -work \leq **Scale**-flavored αr -work $\leq \frac{\alpha}{\beta} r$ -work



filters using **Scale's rank**
instead of **SRPT's rank**

Comparing r -work

Lemma:

$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$

Key steps:

- **SRPT** minimizes mean r -work
- **Scale** minimizes mean *Scale-flavored* r -work
- Under any policy,
 r -work \leq **Scale**-flavored αr -work $\leq \frac{\alpha}{\beta} r$ -work


$$\mathbf{E}[N] = \int_0^{\infty} \frac{\mathbf{E}[W(r)]}{r^2} dr$$



Comparing r -work

Lemma:

$$\mathbf{E}[W_{\text{SRPT}}(r)] \leq \mathbf{E}[W_{\text{Scale}}(r)] \leq \mathbf{E}[W_{\text{SRPT}}(\gamma r)]$$






$$\mathbf{E}[N_{\text{SRPT}}] \leq \mathbf{E}[N_{\text{Scale}}] \leq \gamma \mathbf{E}[N_{\text{SRPT}}]$$

Key steps:

$$\mathbf{E}[N] = \int_0^{\infty} \frac{\mathbf{E}[W(r)]}{r^2} dr$$

- **SRPT** minimizes mean r -work
- **Scale** minimizes mean *Scale*-flavored r -work
- Under any policy,
 r -work \leq **Scale**-flavored αr -work $\leq \frac{\alpha}{\beta} r$ -work





rank functions close enough







SOAP

Scale



WINE






rank functions
close enough

r -work amounts
close enough

SOAP

Scale

WINE



Today's talk

scheduling with
noisy predictions



TCS



Queueing

Today's talk

scheduling with
noisy predictions



- **TCS**: need to be careful

Today's talk

scheduling with
noisy predictions



- **TCS**: need to be careful
- **Queueing**: simple **rank**-based policy suffices

Today's talk

scheduling with

multiple servers



TCS



Queueing

- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

scheduling with

noisy predictions



TCS



Queueing

- **TCS**: need to be careful
- **Queueing**: simple **rank**-based policy suffices

Today's talk

scheduling with

multiple servers



TCS



Queueing

- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

scheduling with

noisy predictions



TCS



Queueing

- **TCS**: need to be careful
- **Queueing**: simple **rank**-based policy suffices



New tools:

Today's talk

scheduling with

multiple servers



TCS



Queueing



scheduling with

noisy predictions



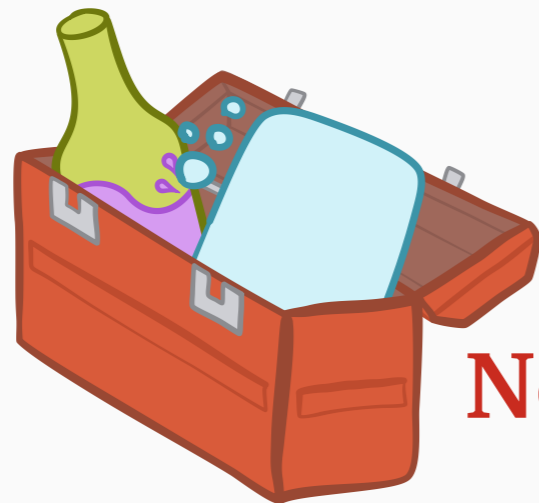
TCS



Queueing

- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

- **TCS**: need to be careful
- **Queueing**: simple **rank**-based policy suffices



New tools: WINE and SOAP

Today's talk

scheduling with

multiple servers



TCS



Queueing

- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

scheduling with

noisy predictions




TCS



Queueing

- **TCS**: need to be careful
- **Queueing**: simple **rank**-based policy suffices



potential
TCS use?

New tools: WINE and SOAP

Today's talk

scheduling with
multiple servers

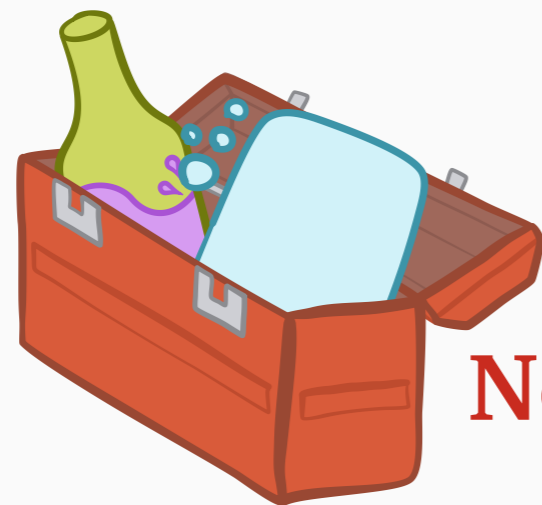


- SRPT-*k* is good
- In general: good to adapt optimal single-server policy

scheduling with
noisy predictions



- **TCS**: need to be careful
- **Queueing**: simple **rank**-based policy suffices



potential
TCS use?

detailed system
modeling?

New tools: WINE and SOAP

References

Scully, Harchol-Balter, and Scheller-Wolf (2018). “SOAP: One Clean Analysis of All Age-Based Scheduling Policies.” *Proc. ACM Meas. Anal. Comput. Syst.* (SIGMETRICS 2018).

- Introduces **rank** functions and the general **SOAP** analysis
- **Finalist: 2019 INFORMS APS Best Student Paper Prize**

Grosof, Scully, and Harchol-Balter (2018). “SRPT for Multiserver Systems.” *Perform. Eval.* (PERFORMANCE 2018).

- First **queueing** analysis of SRPT-**k**
- Uses tagged job method plus **worst-case** r-work decomposition
- **Winner: PERFORMANCE 2018 Best Student Paper Award**

Scully, Grosof, and Harchol-Balter (2020). “The Gittins Policy is Nearly Optimal in the M/G/k under Extremely General Conditions.” *Proc. ACM Meas. Anal. Comput. Syst.* (SIGMETRICS 2021).

- First **queueing** analysis of **Gittins-k**
- Introduces **WINE**
- **Winner: 2022 INFORMS George Nicholson Student Paper Competition**



Scully, Grosof, and Mitzenmacher (2022). “Uniform Bounds for Scheduling with Job Size Estimates.” *13th Innovations in Theoretical Computer Science Conference* (ITCS 2022).

- First **queueing** competitive ratios for **noisy predictions**
- Uses both **SOAP** and **WINE**