Bounding Mean Slowdown *in* Multiserver Systems

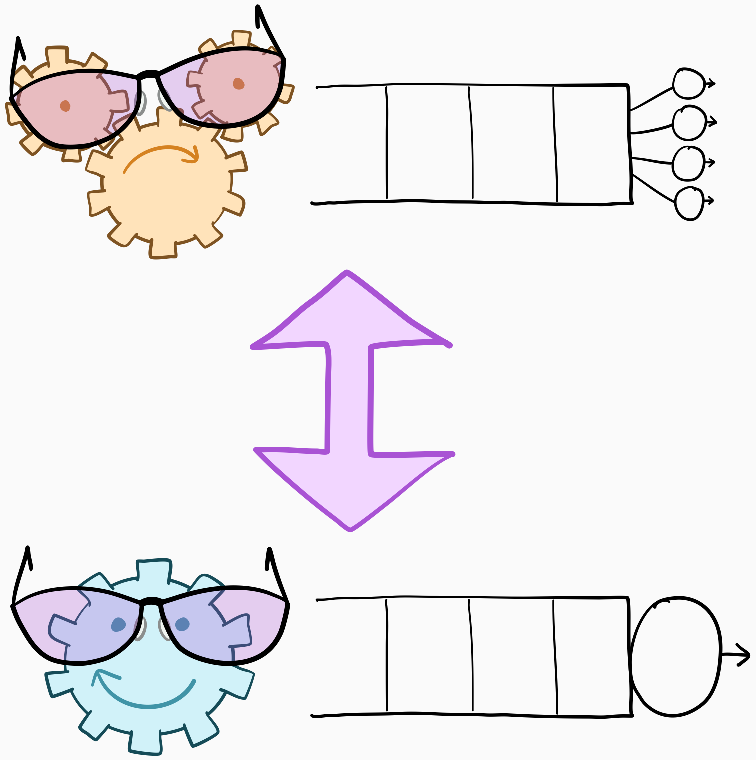

Ziv Scully
Carnegie Mellon University

This talk: near-optimal *multiserver* scheduling

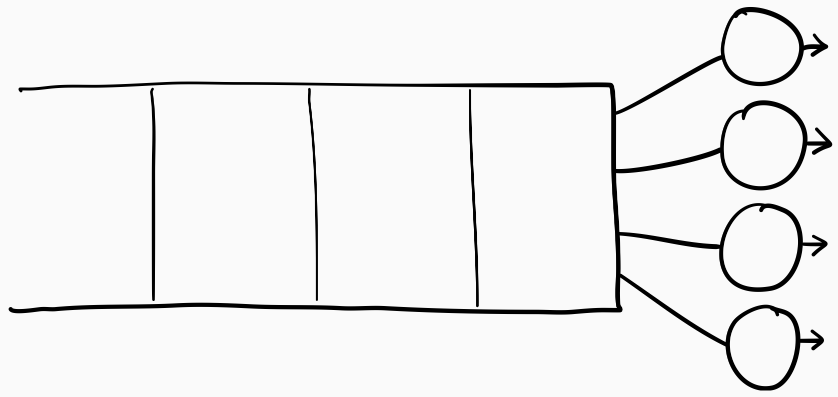

First: background on *single-server* scheduling

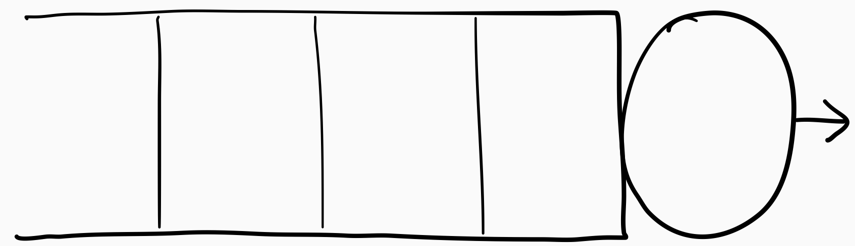

This talk: near-optimal *multiserver* scheduling

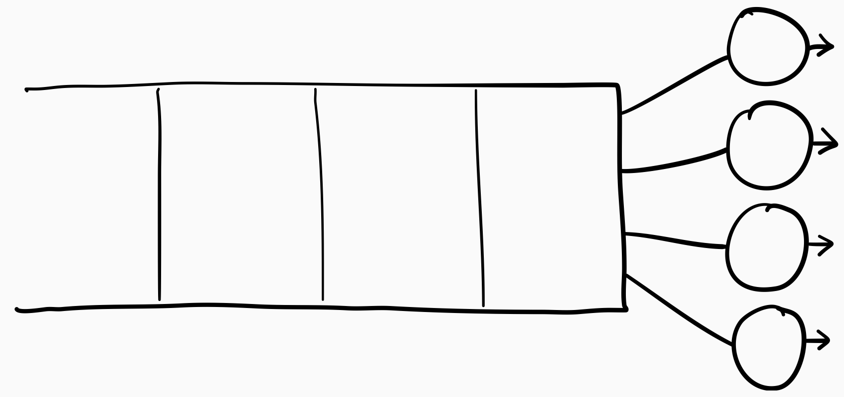

First: background on *single-server* scheduling

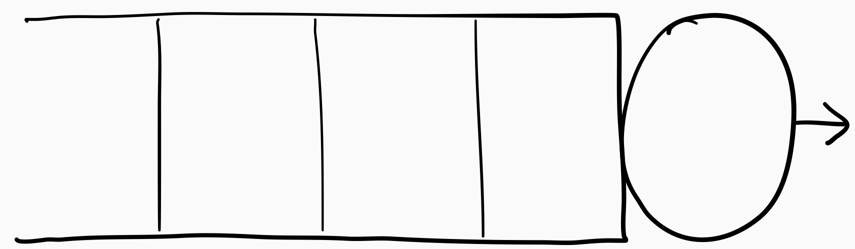

This talk: near-optimal *multiserver* scheduling

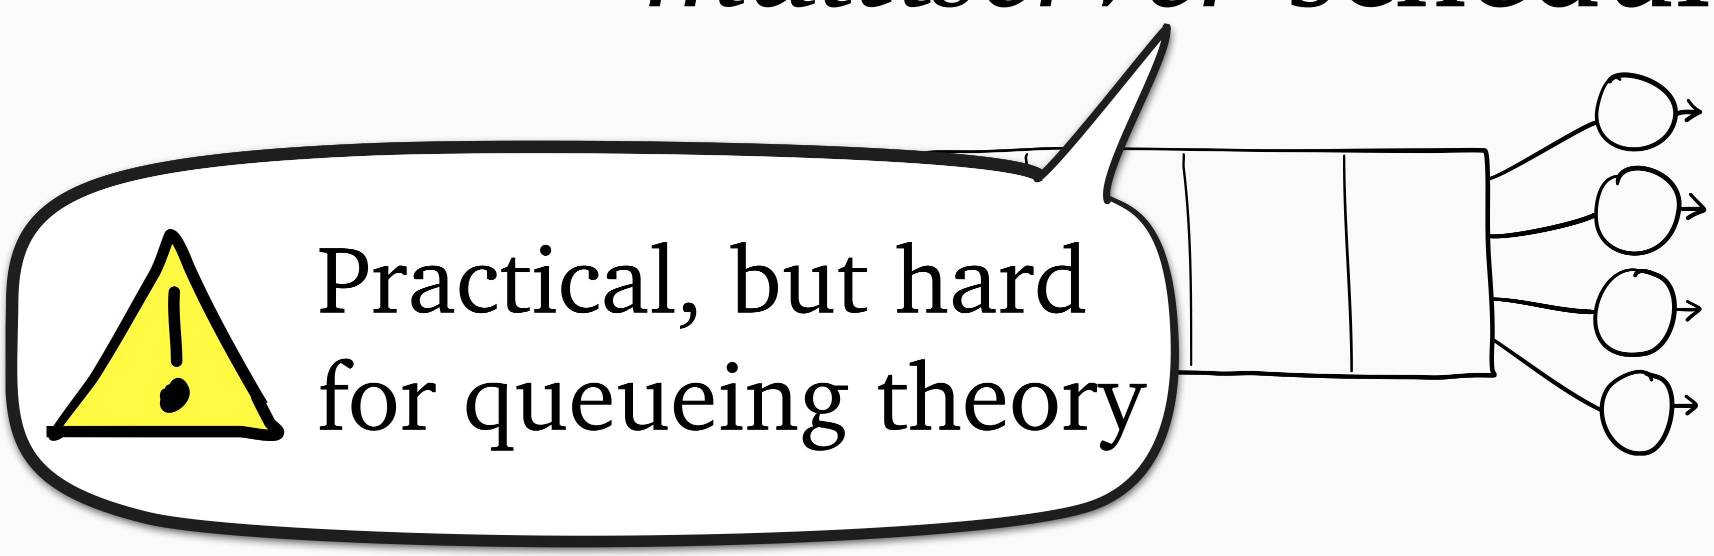

First: background on
single-server scheduling

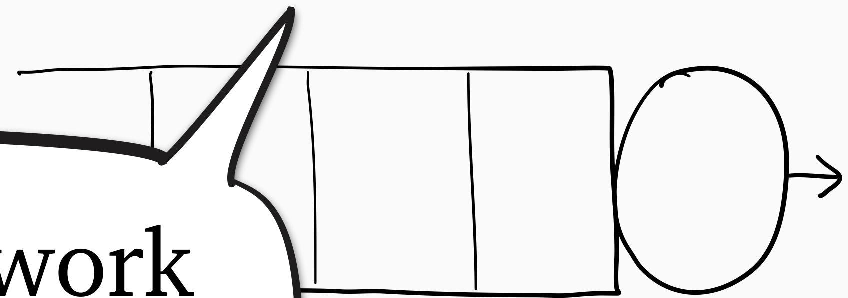

Lots of prior work
in queueing theory

This talk: near-optimal
multiserver scheduling

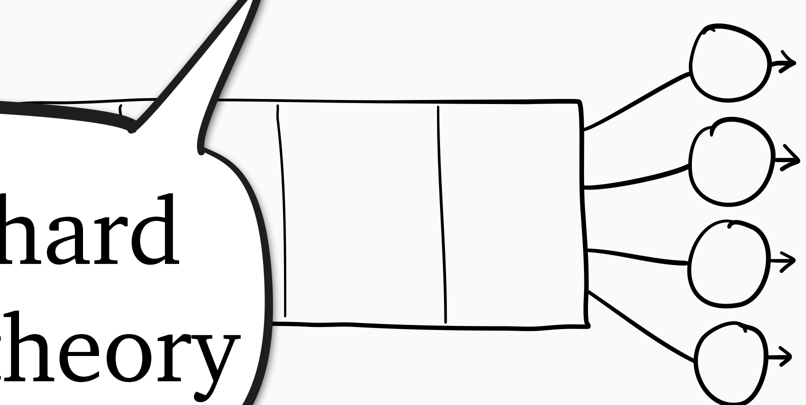

Practical, but hard
for queueing theory

M/G/1 Queue

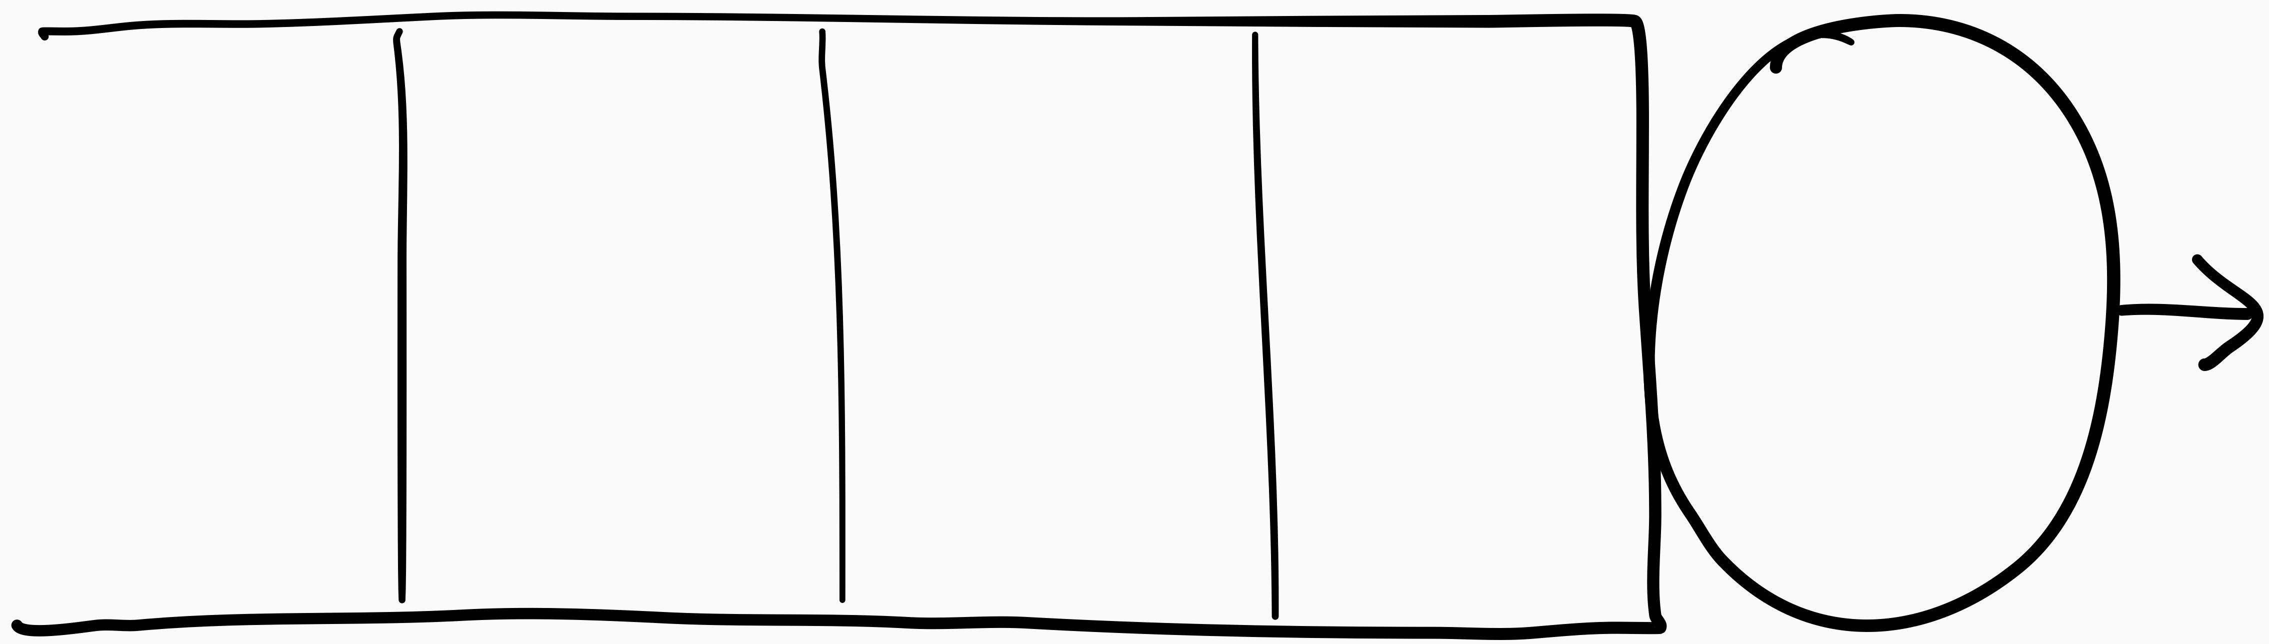

M/G/1 Queue

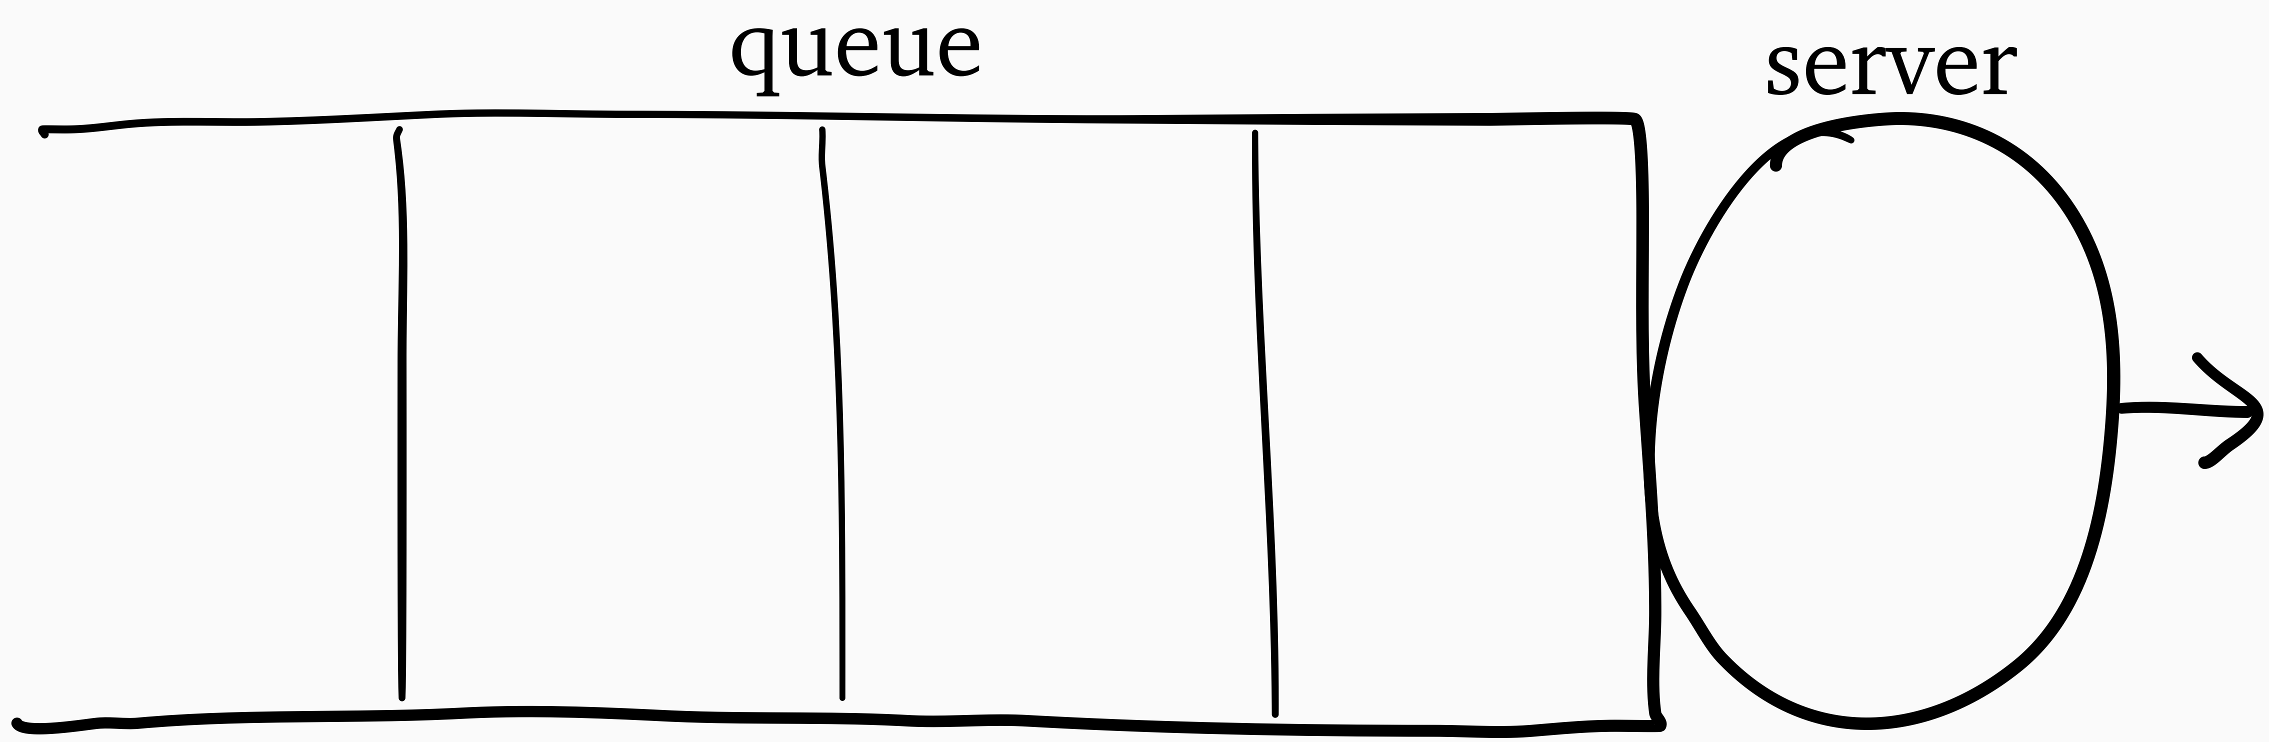

M/G/1 Queue

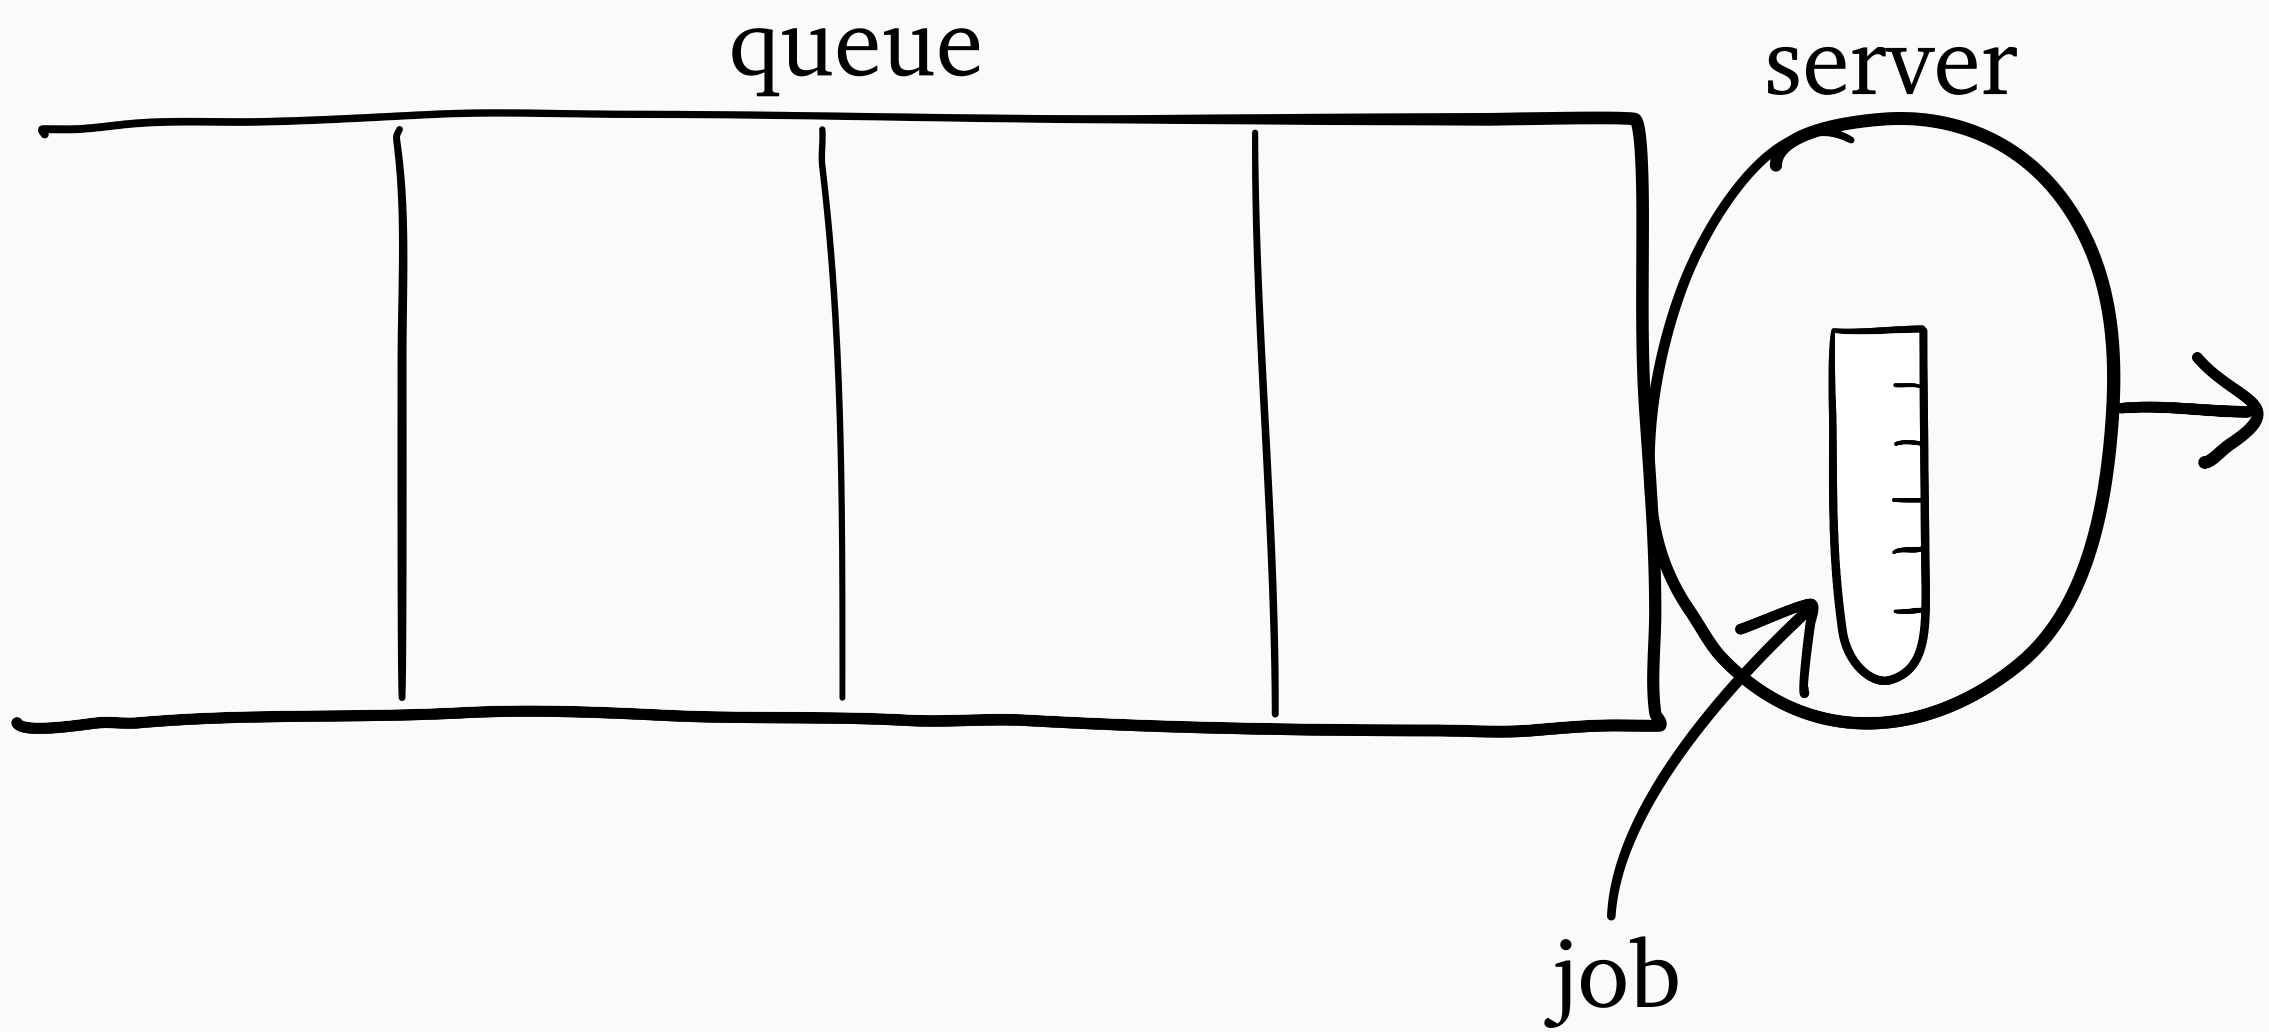

M/G/1 Queue

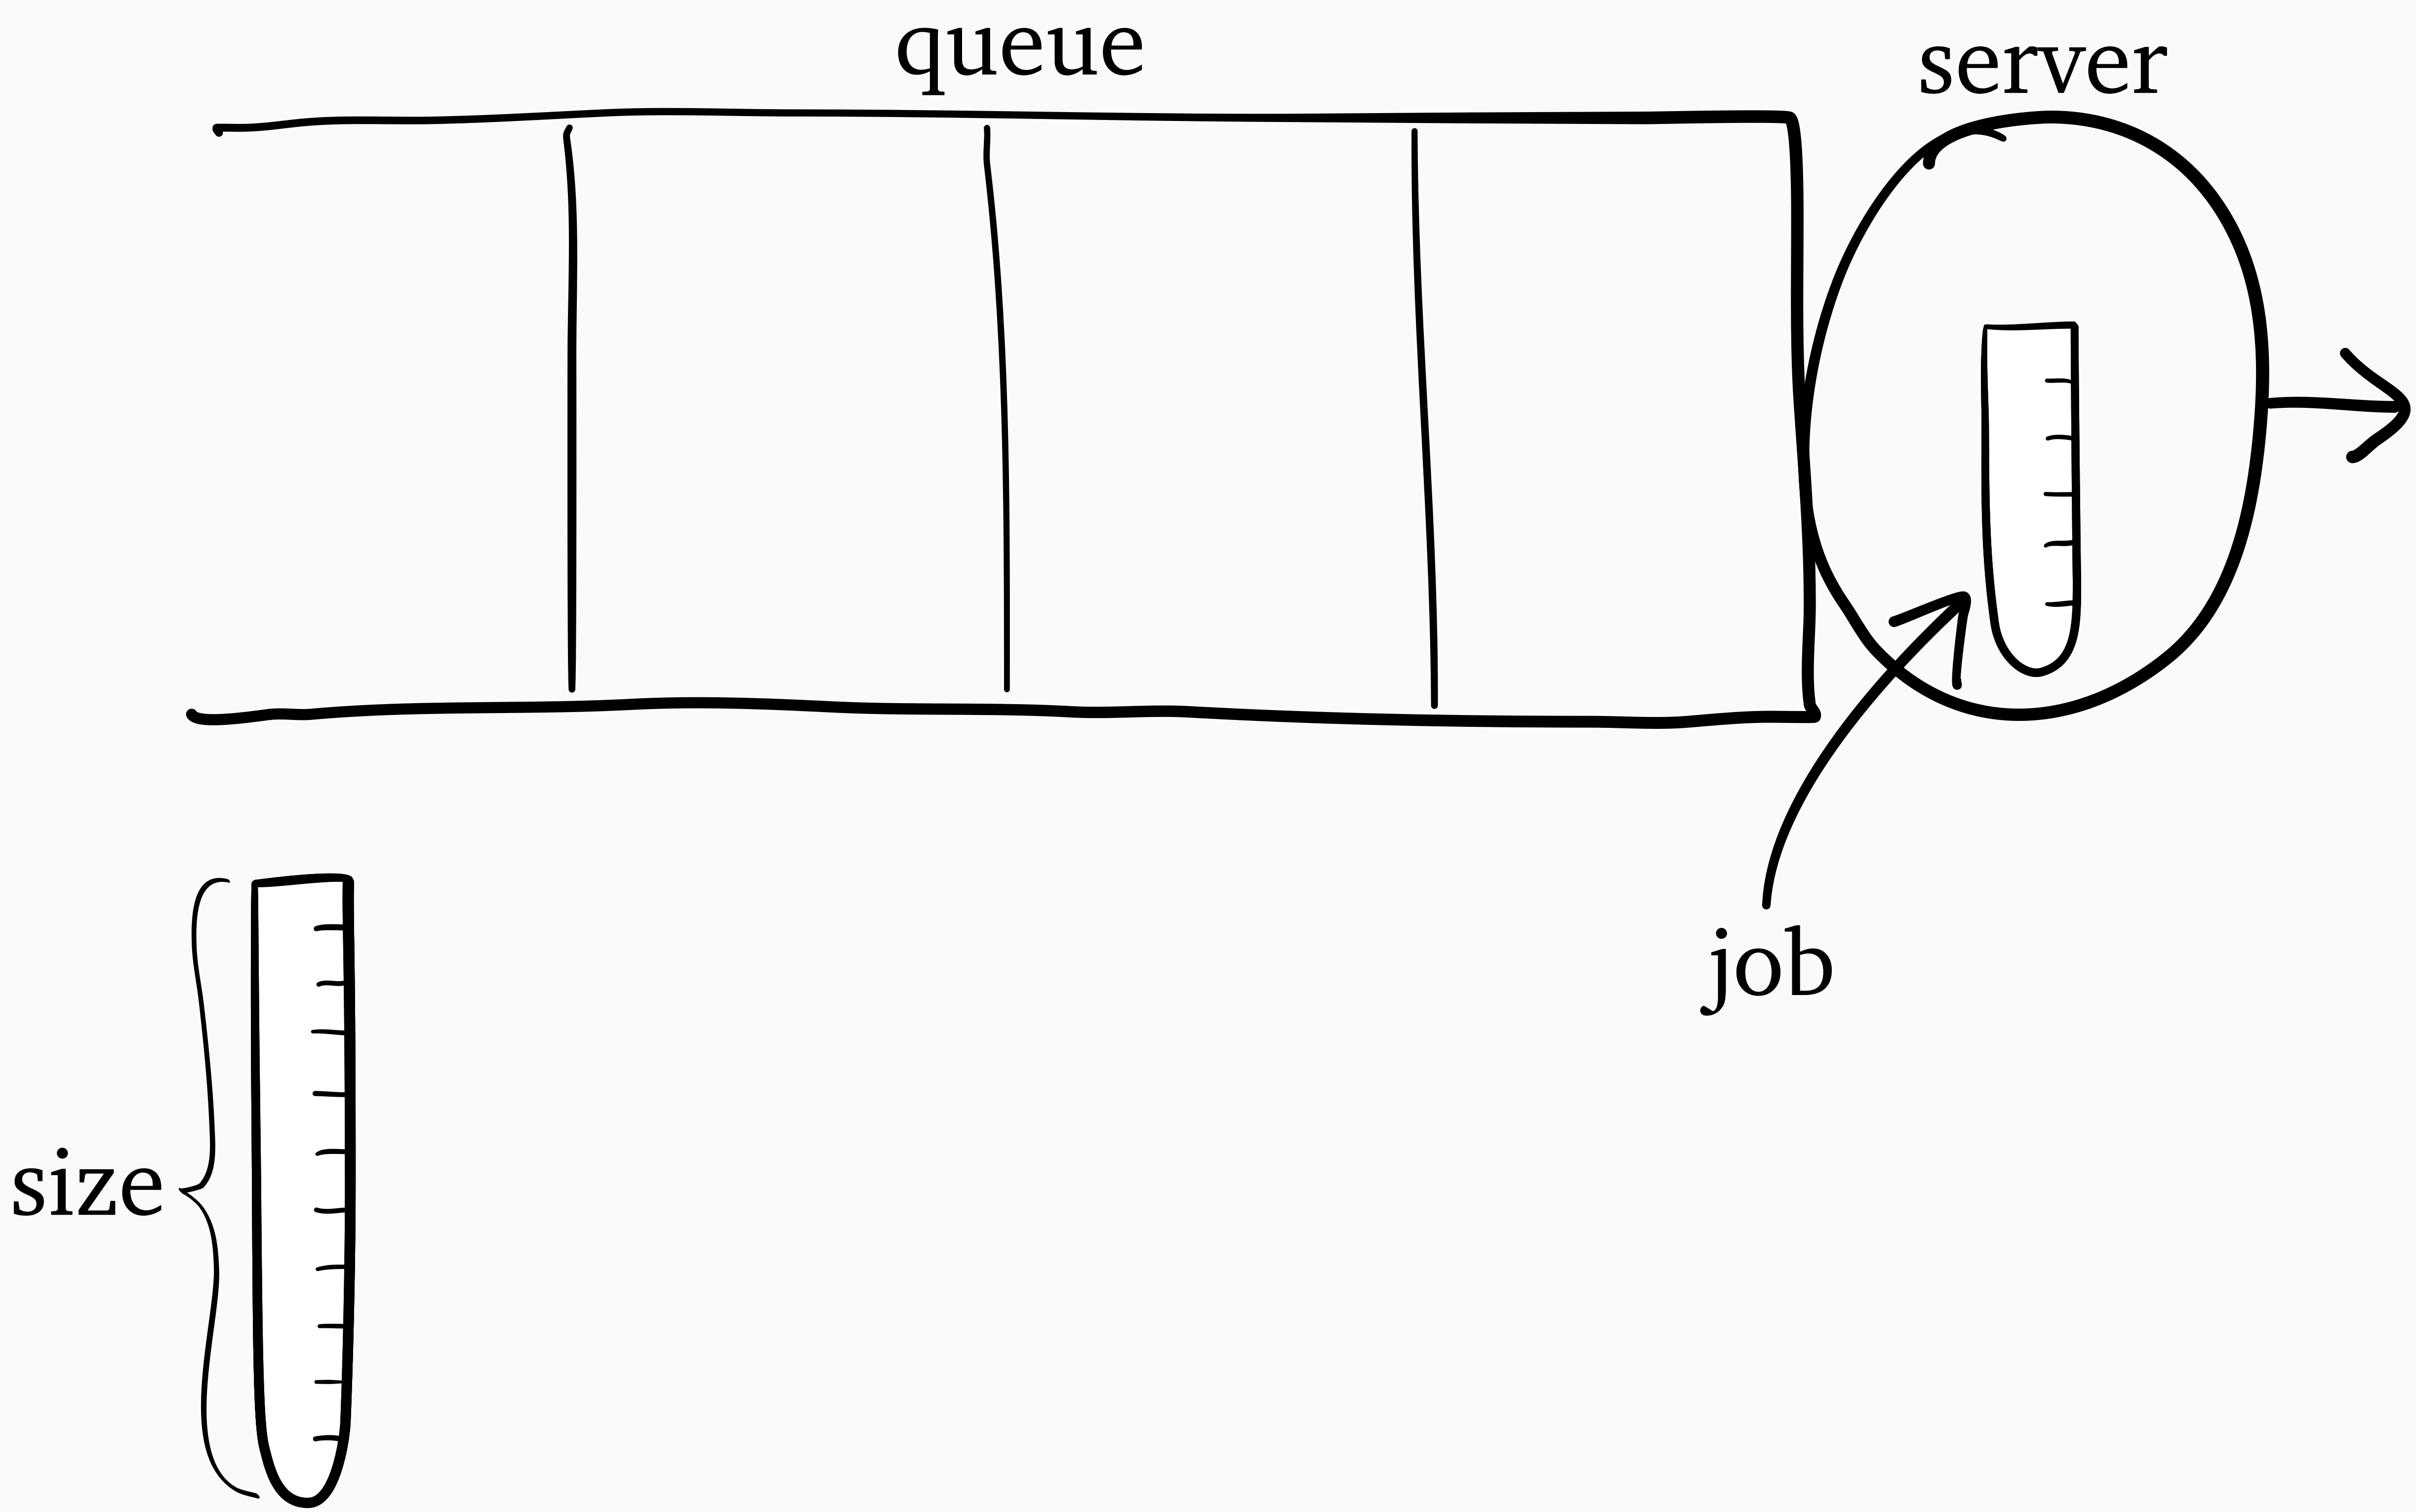

M/G/1 Queue

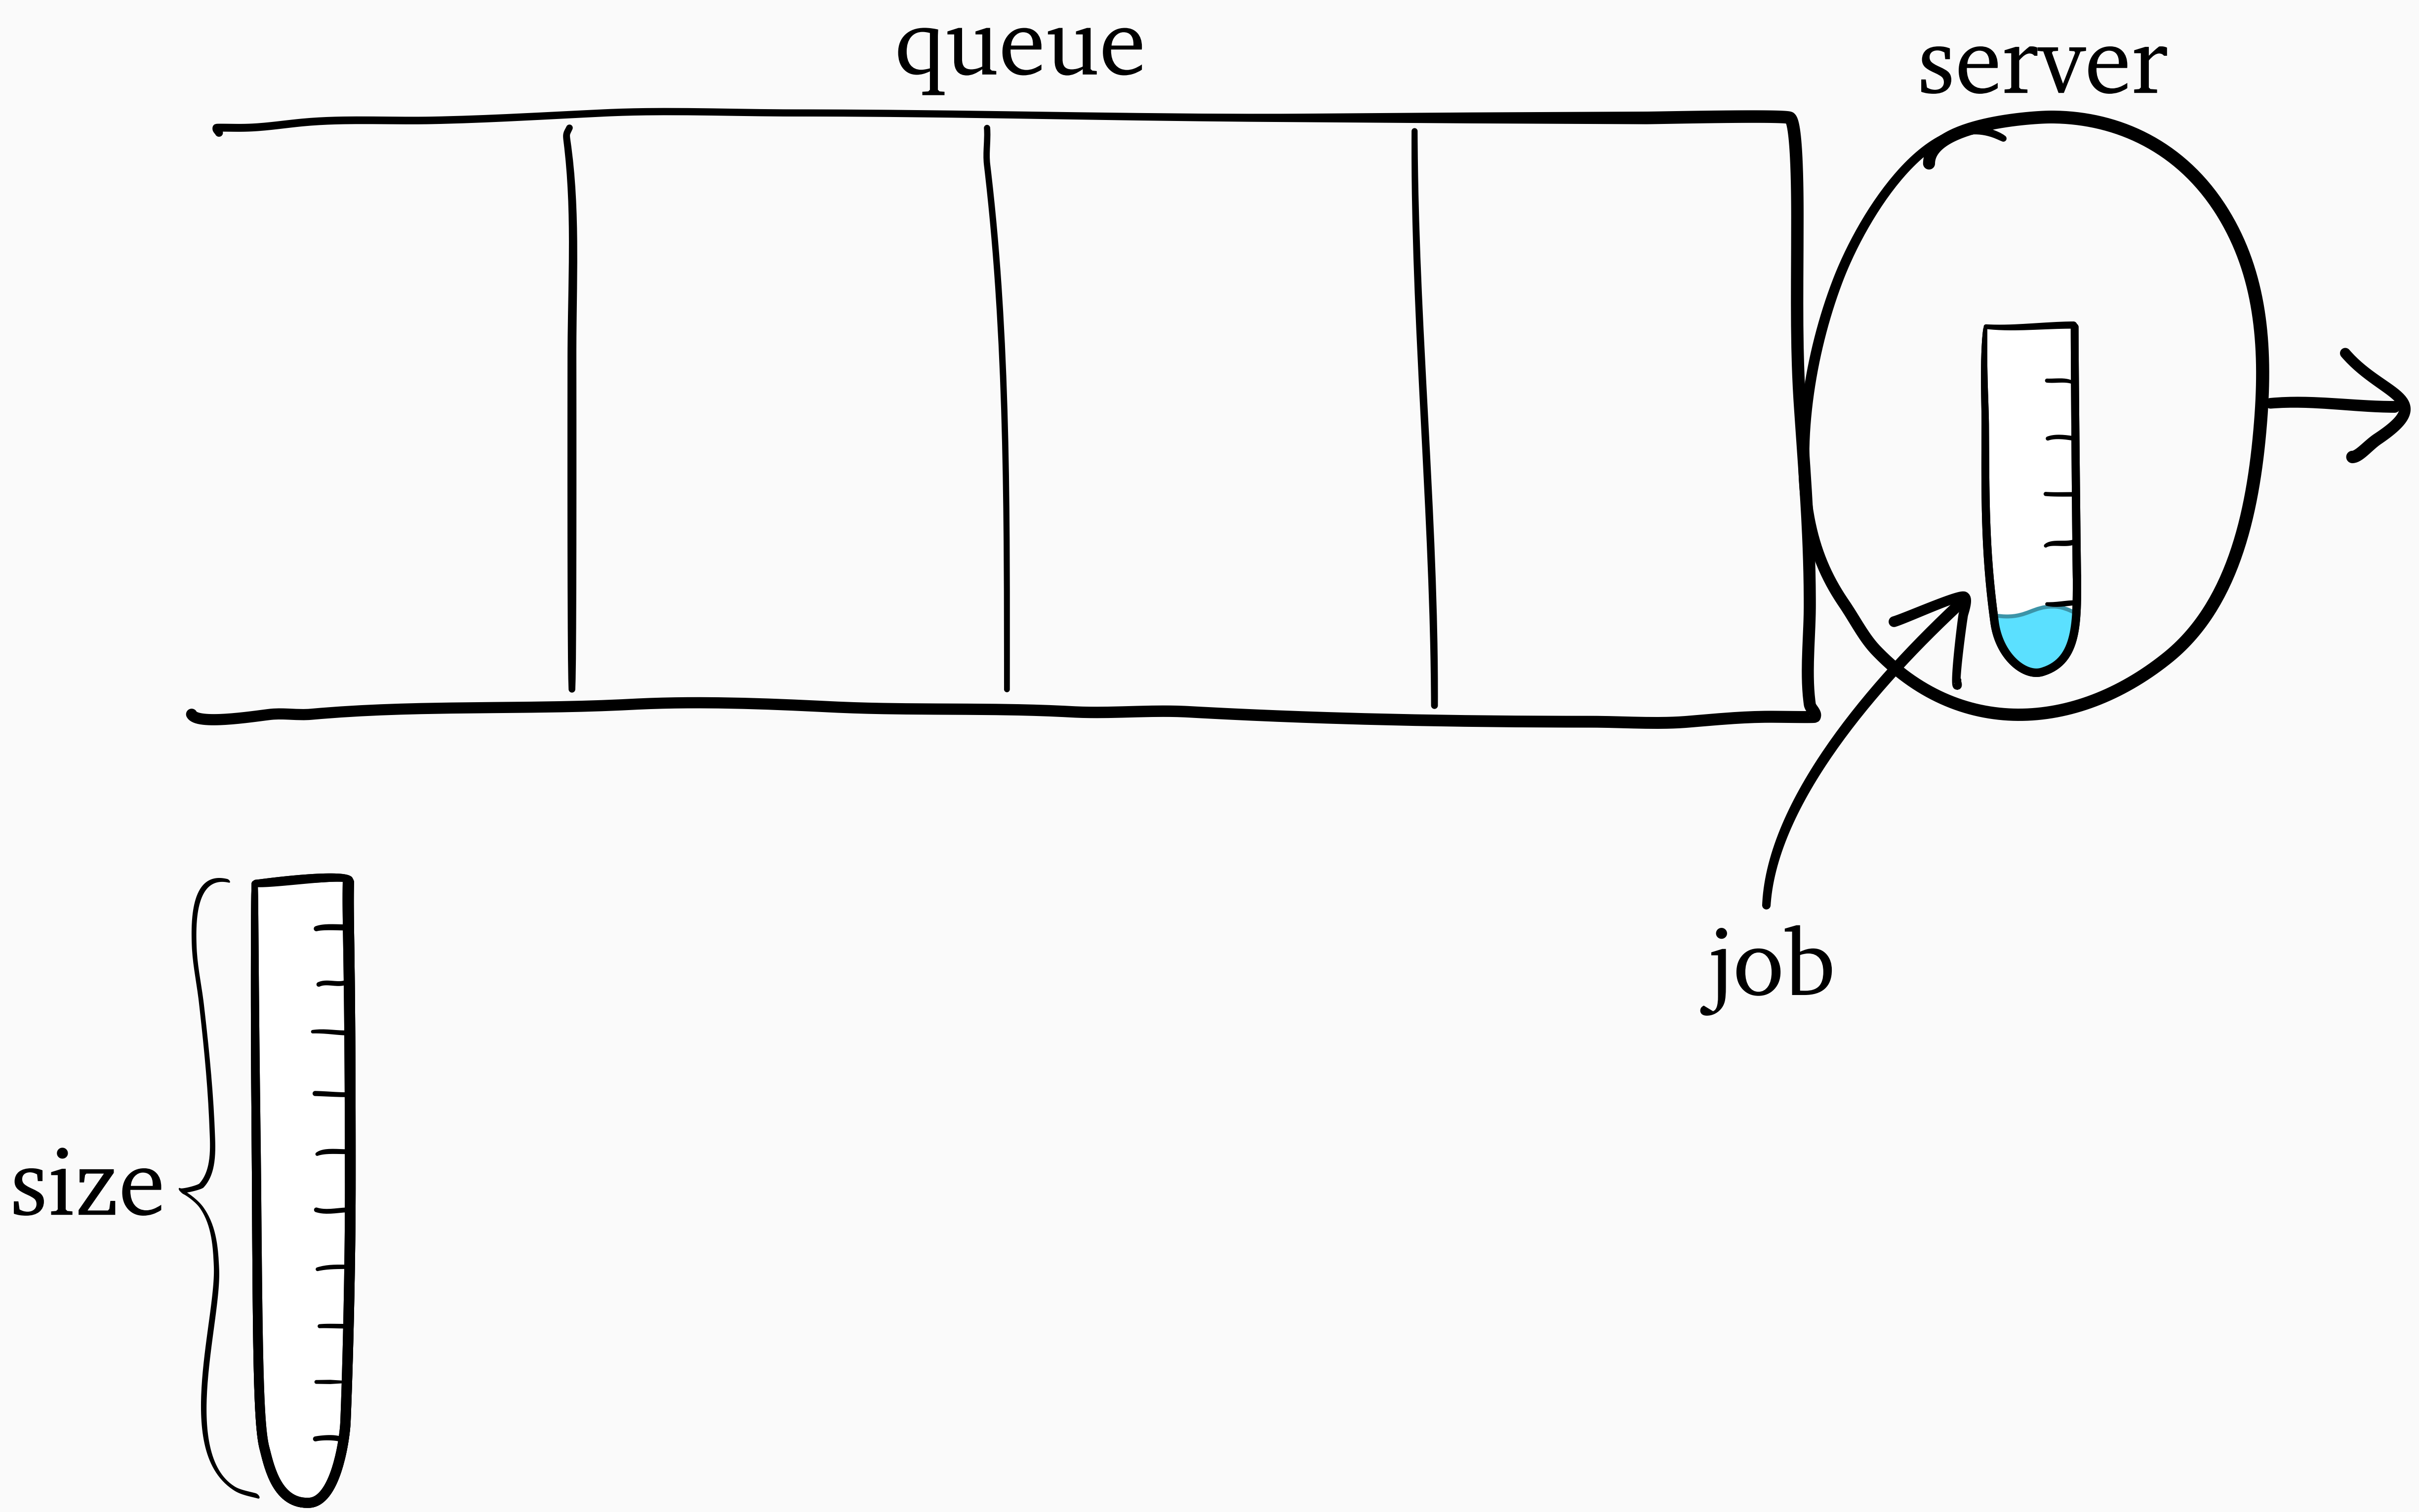

M/G/1 Queue

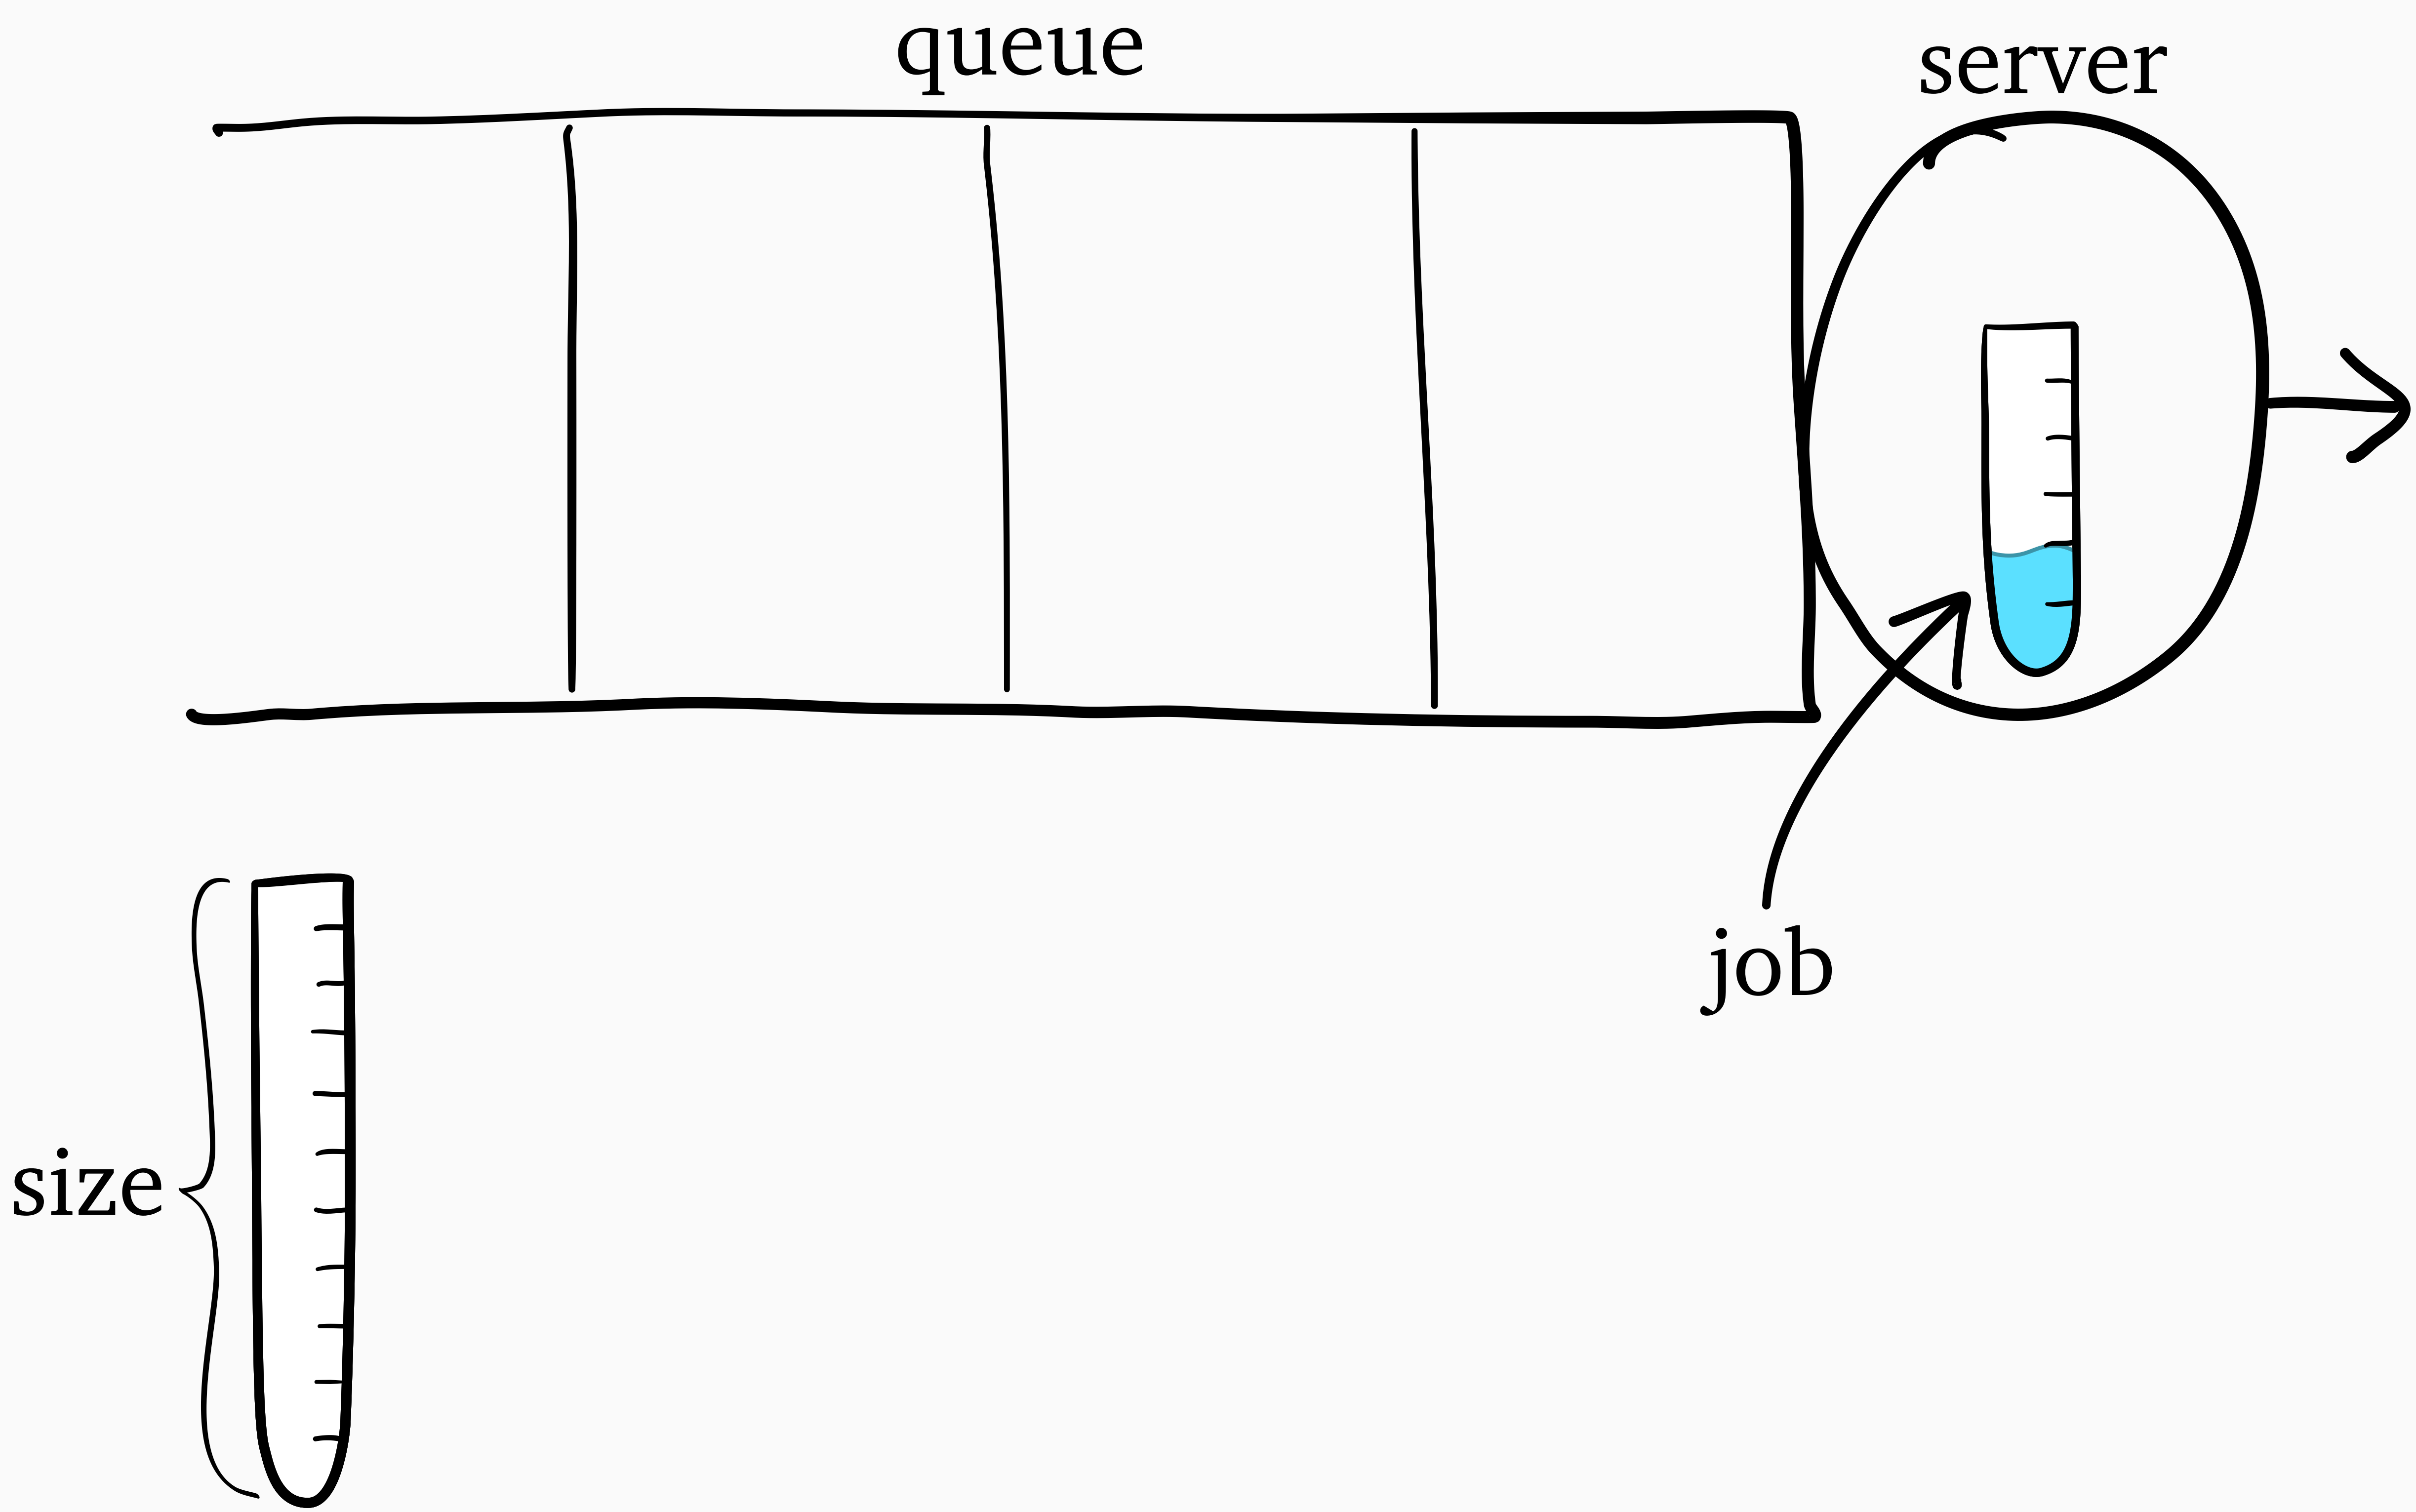

M/G/1 Queue

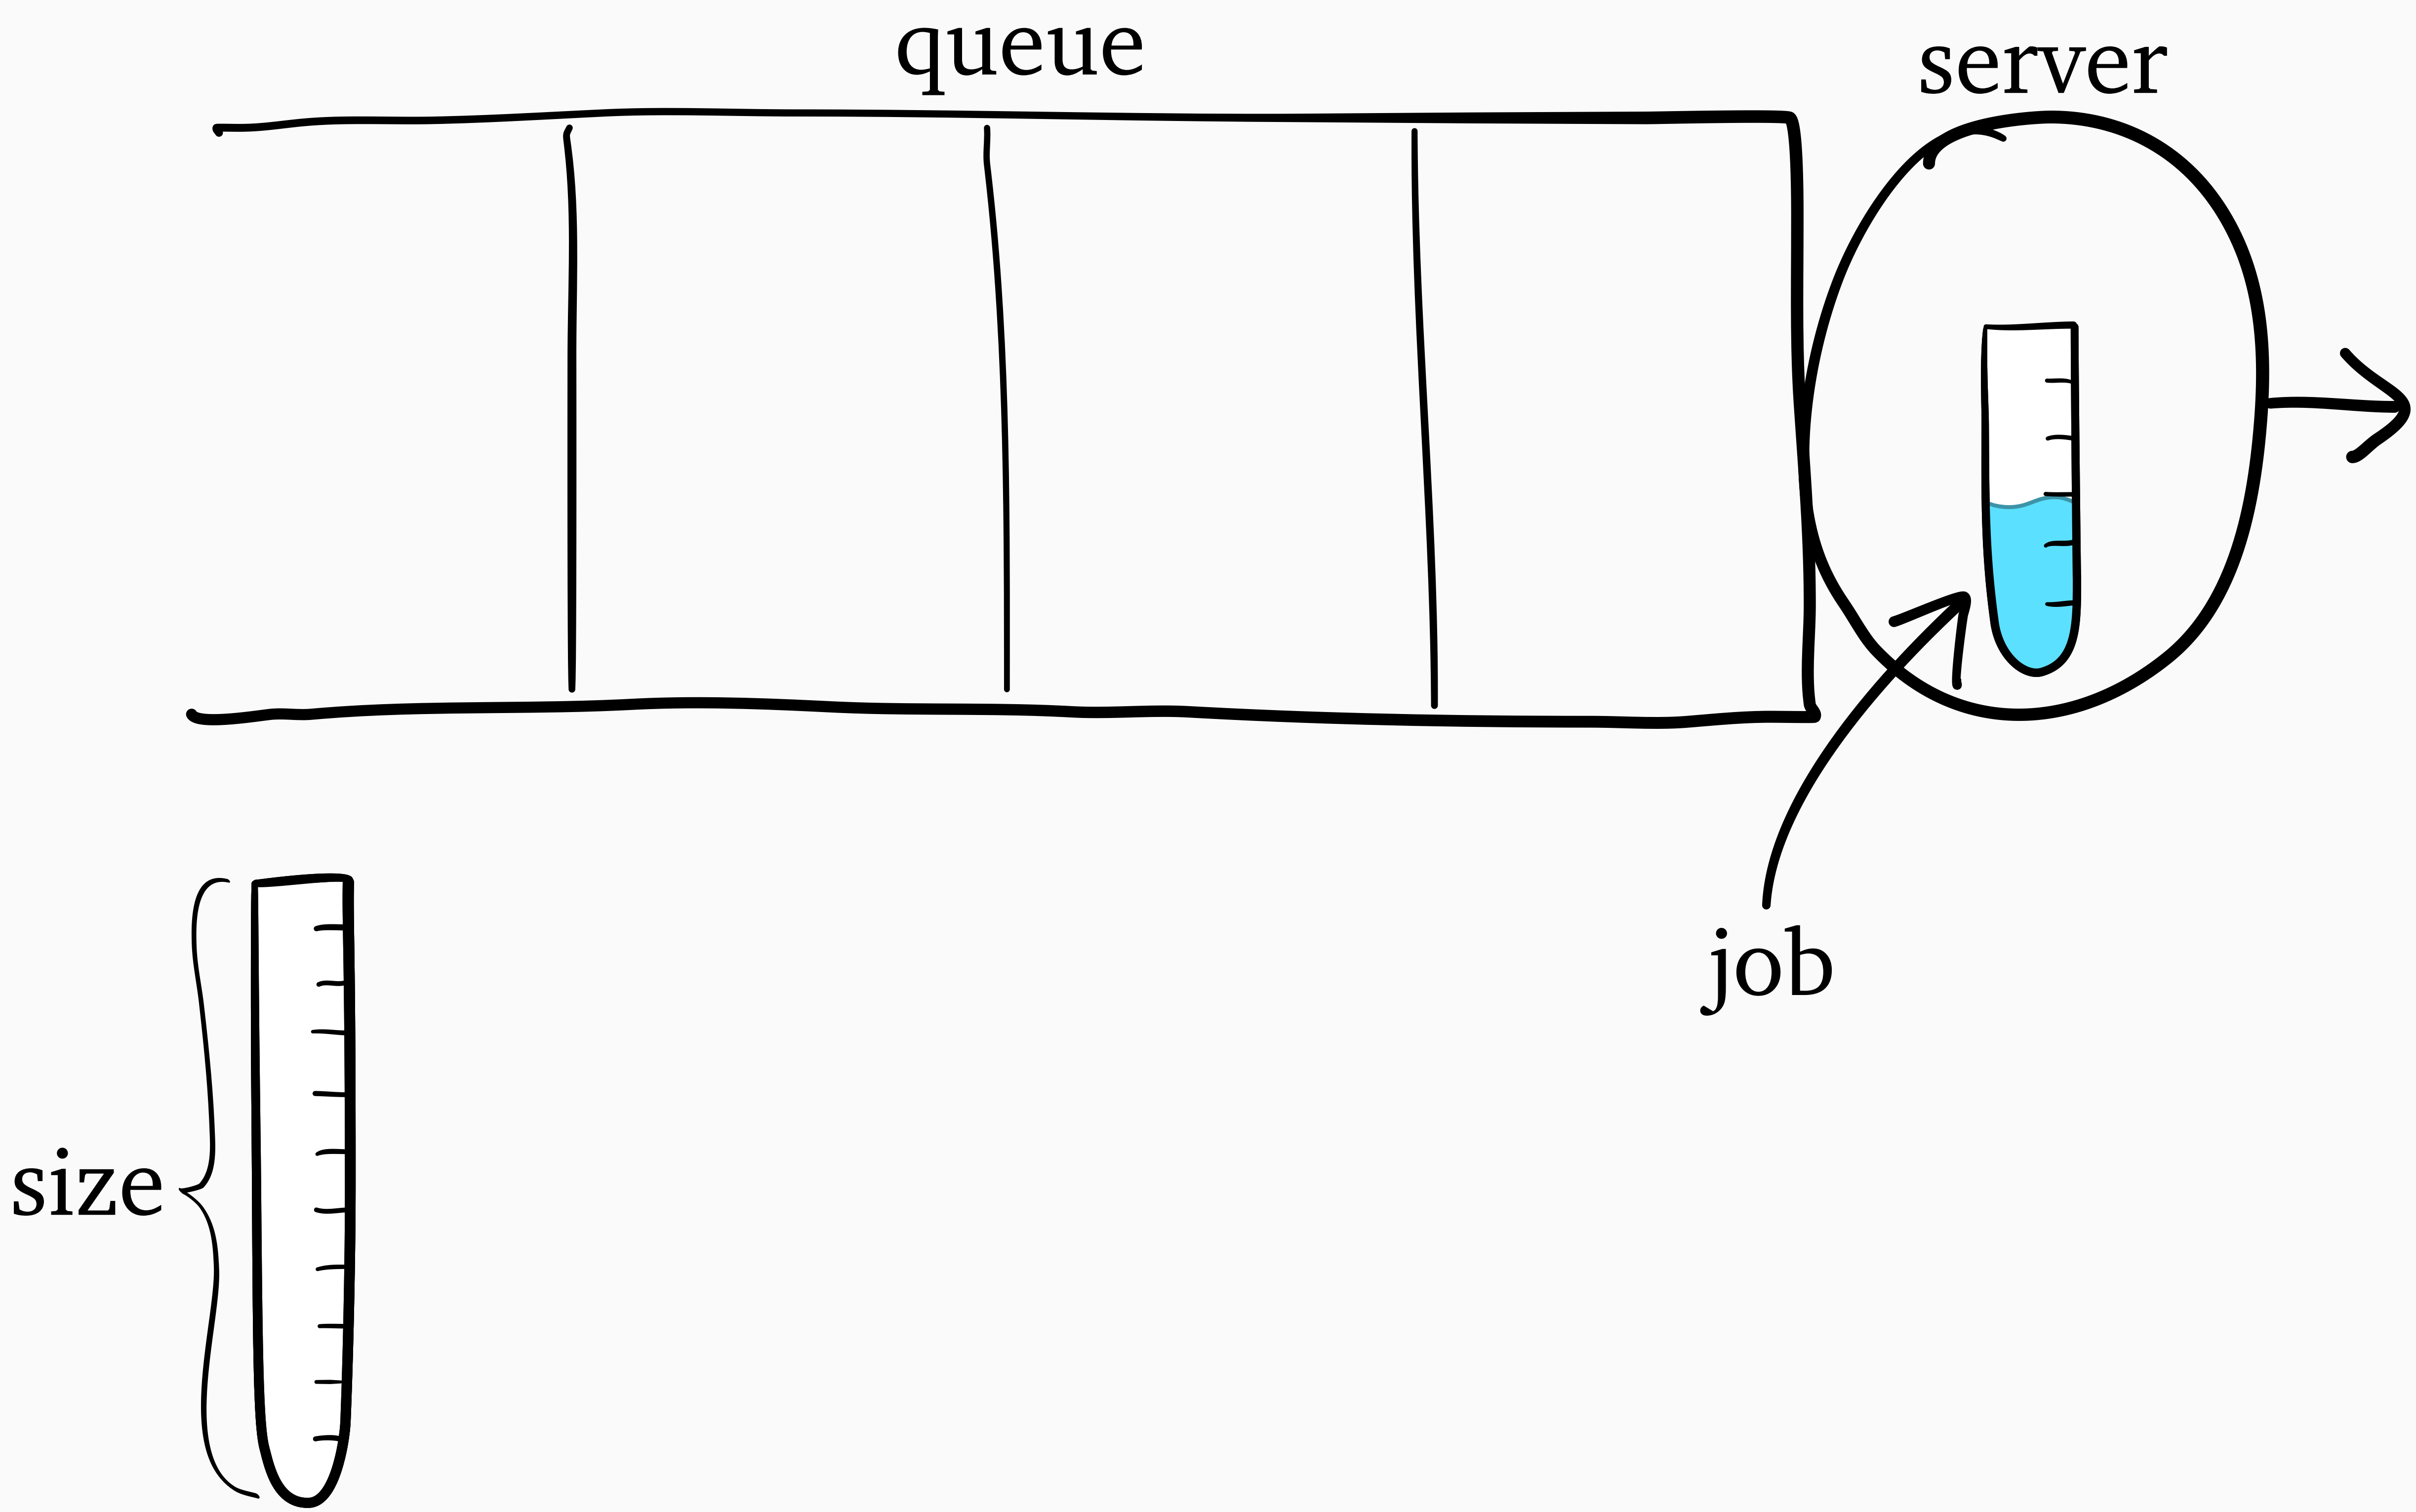

M/G/1 Queue

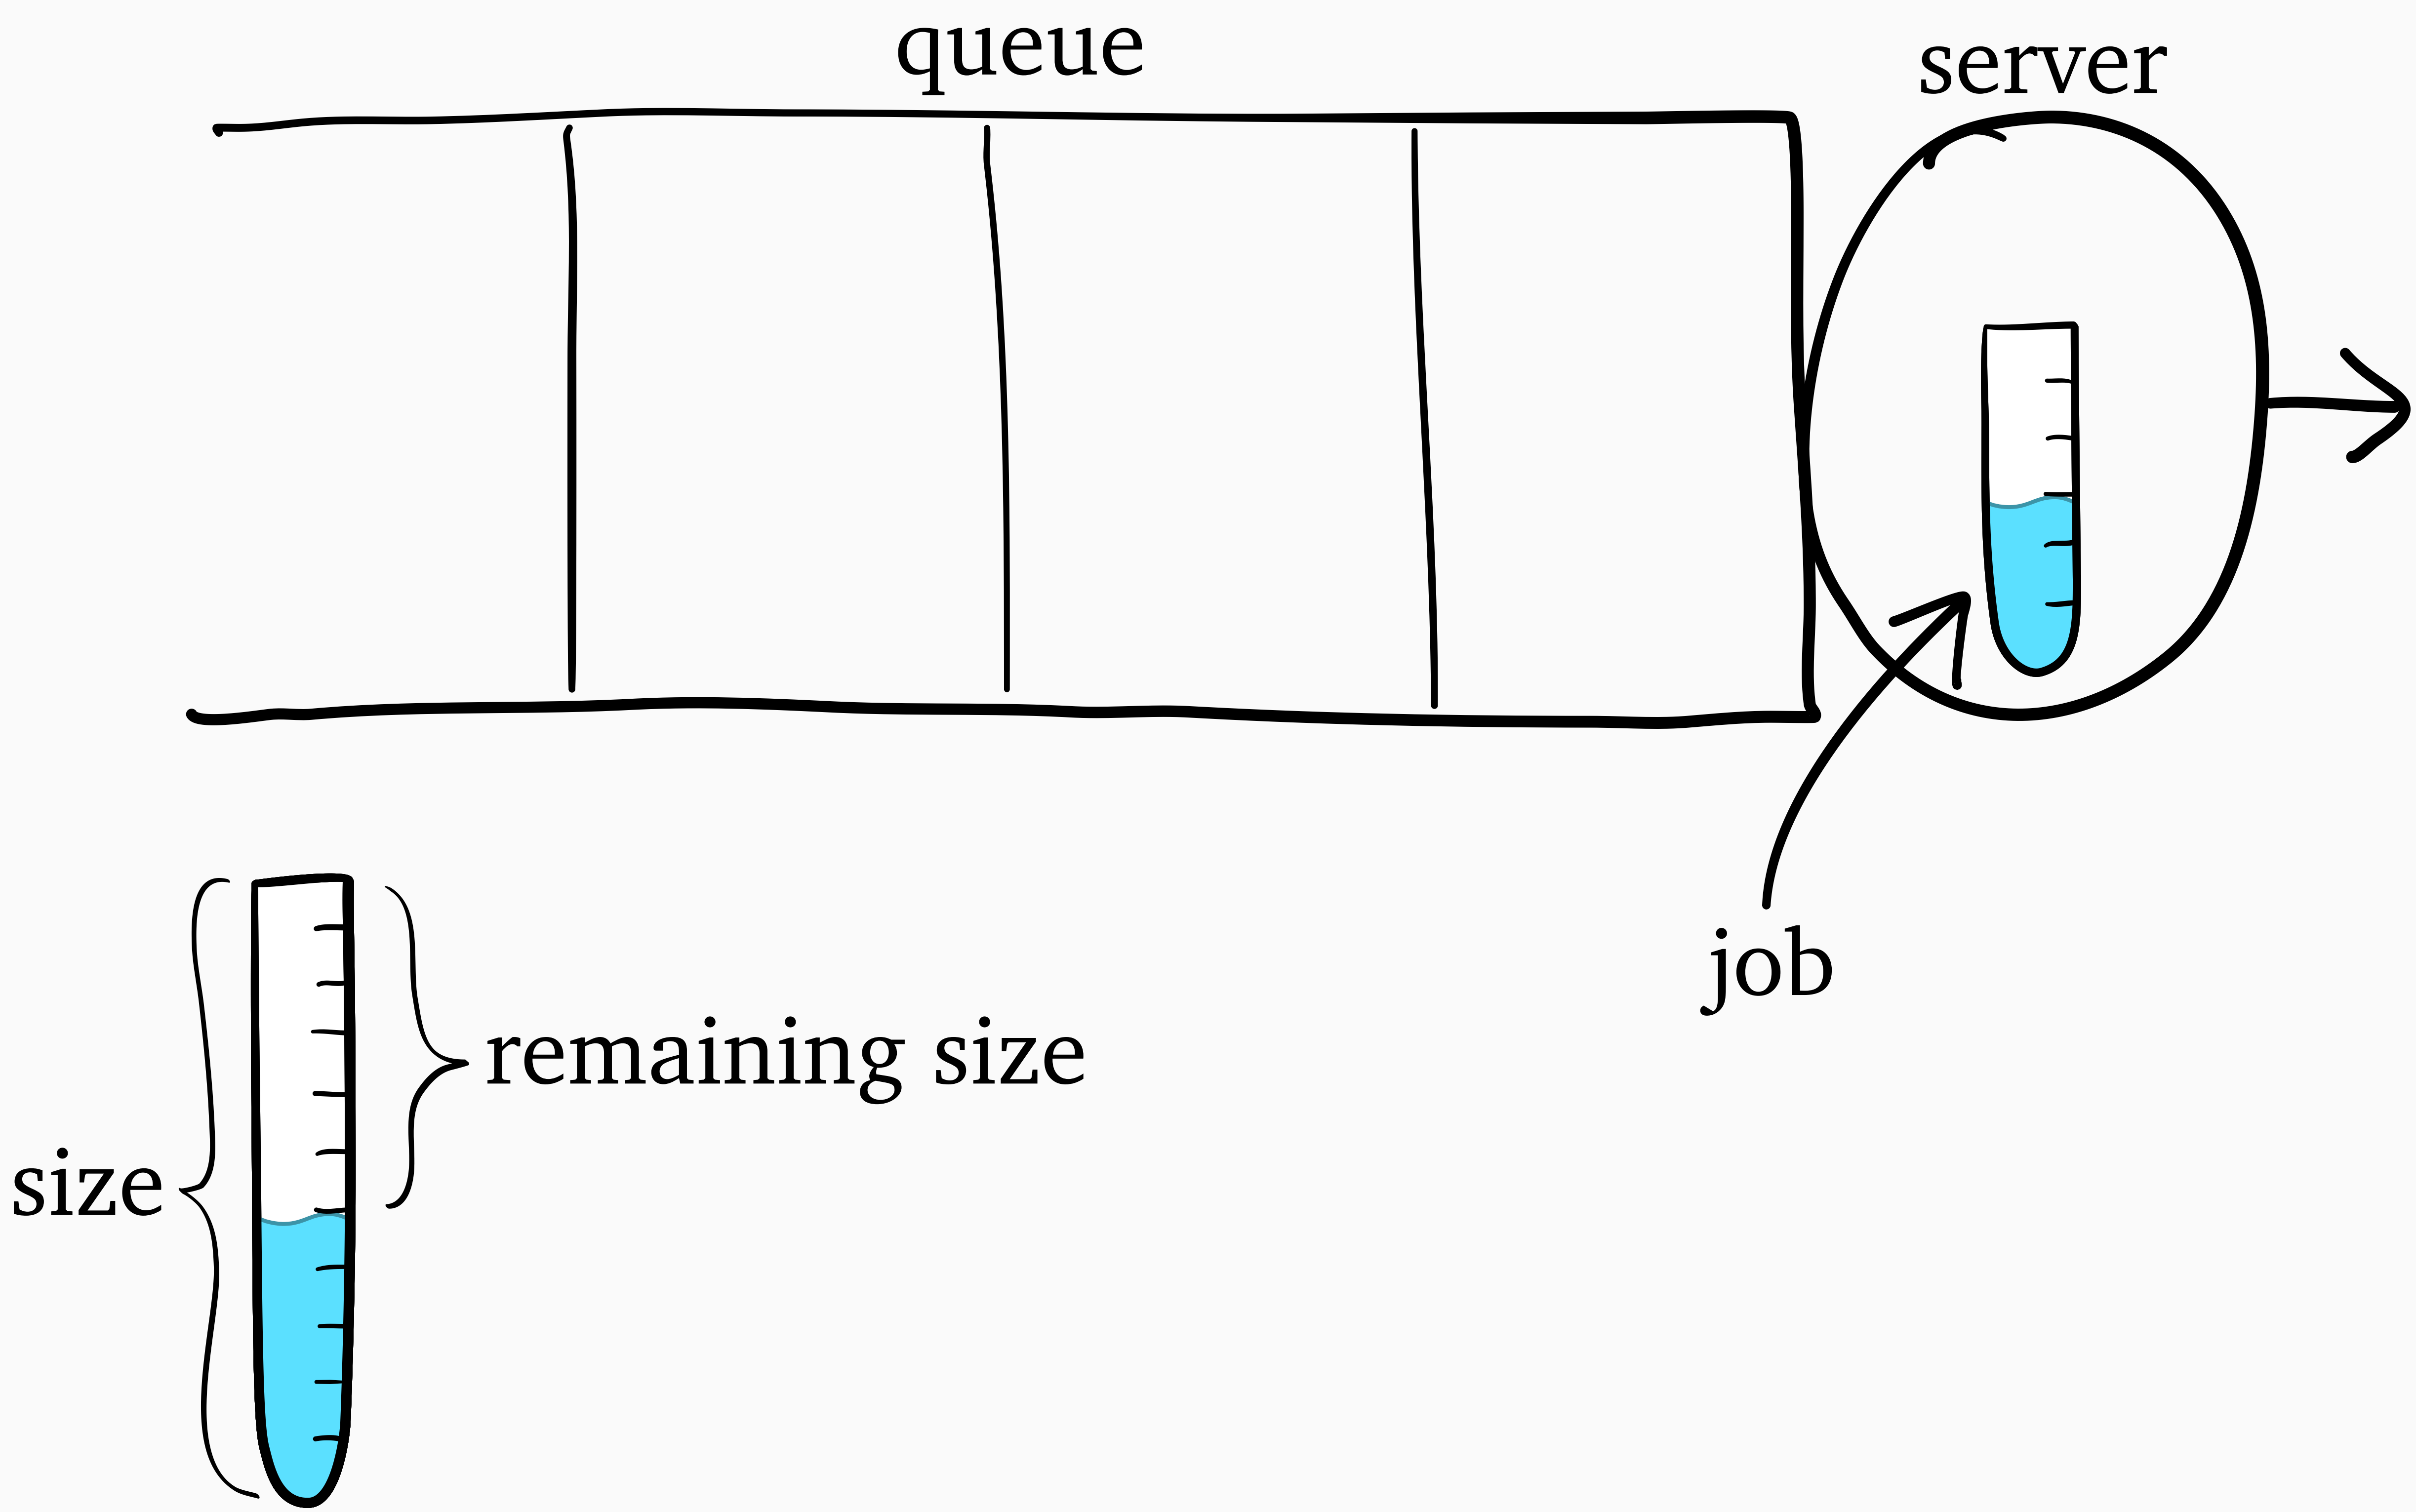

M/G/1 Queue

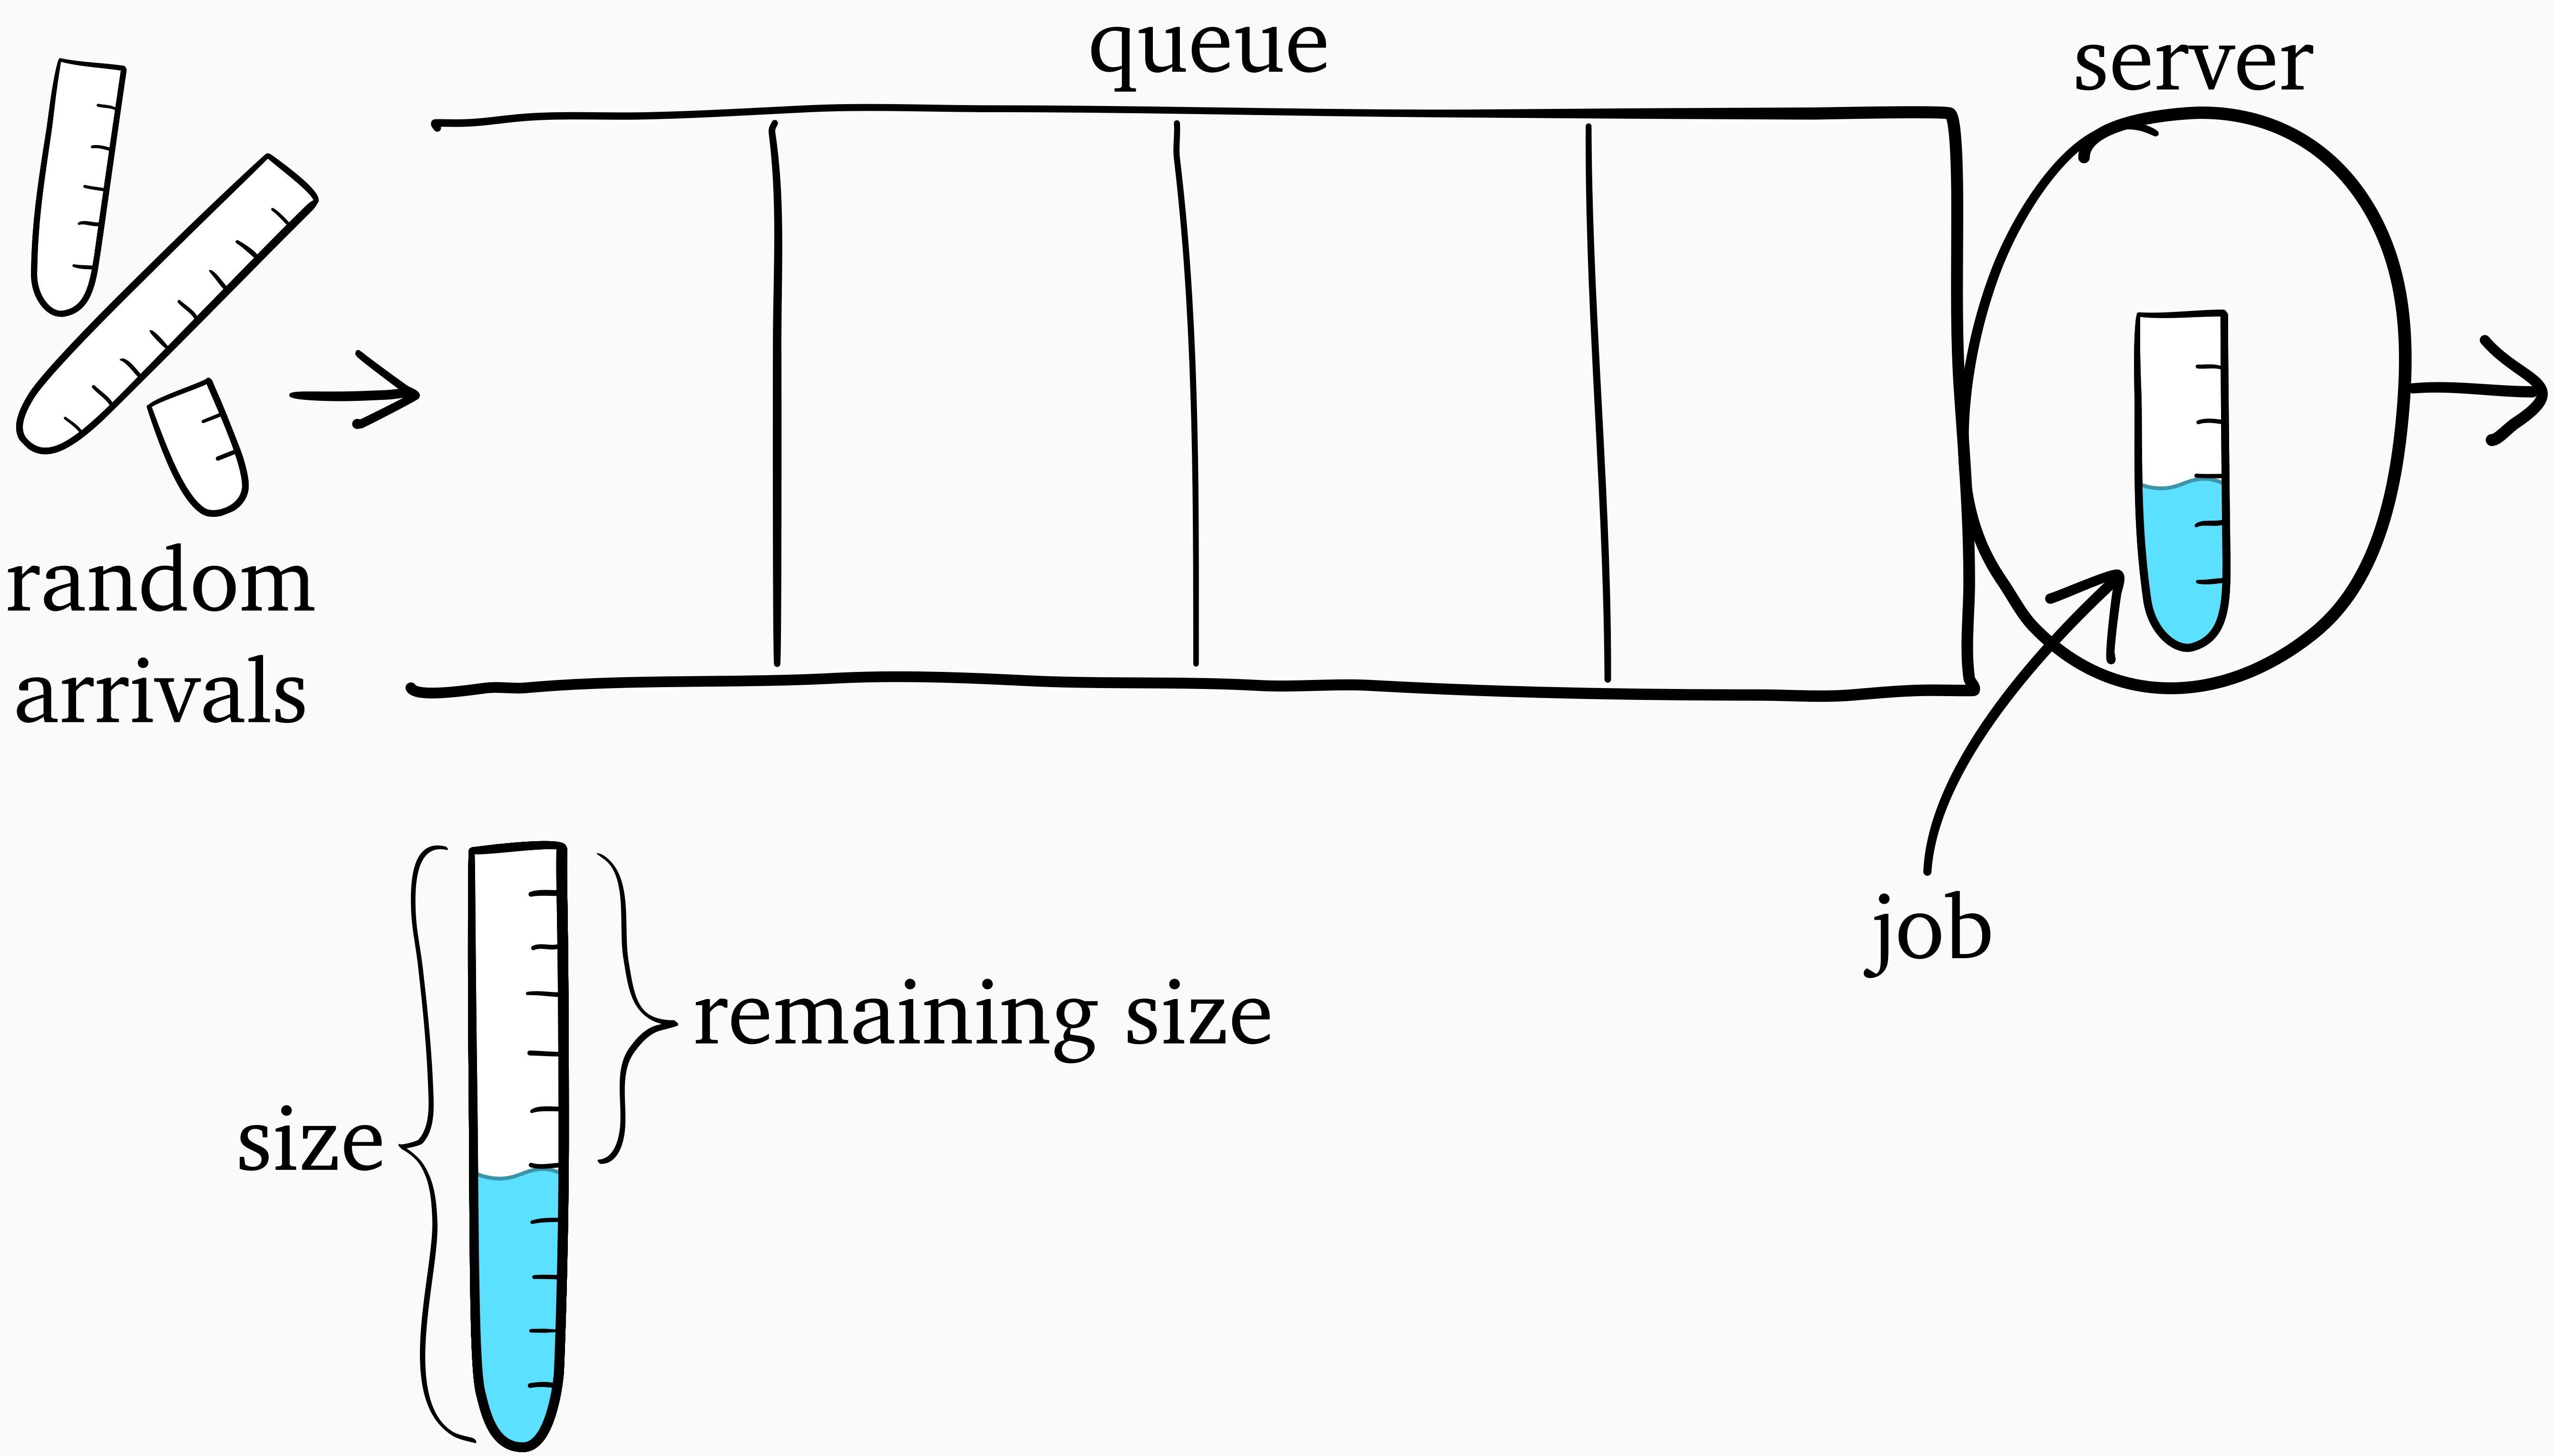

M/G/1 Queue

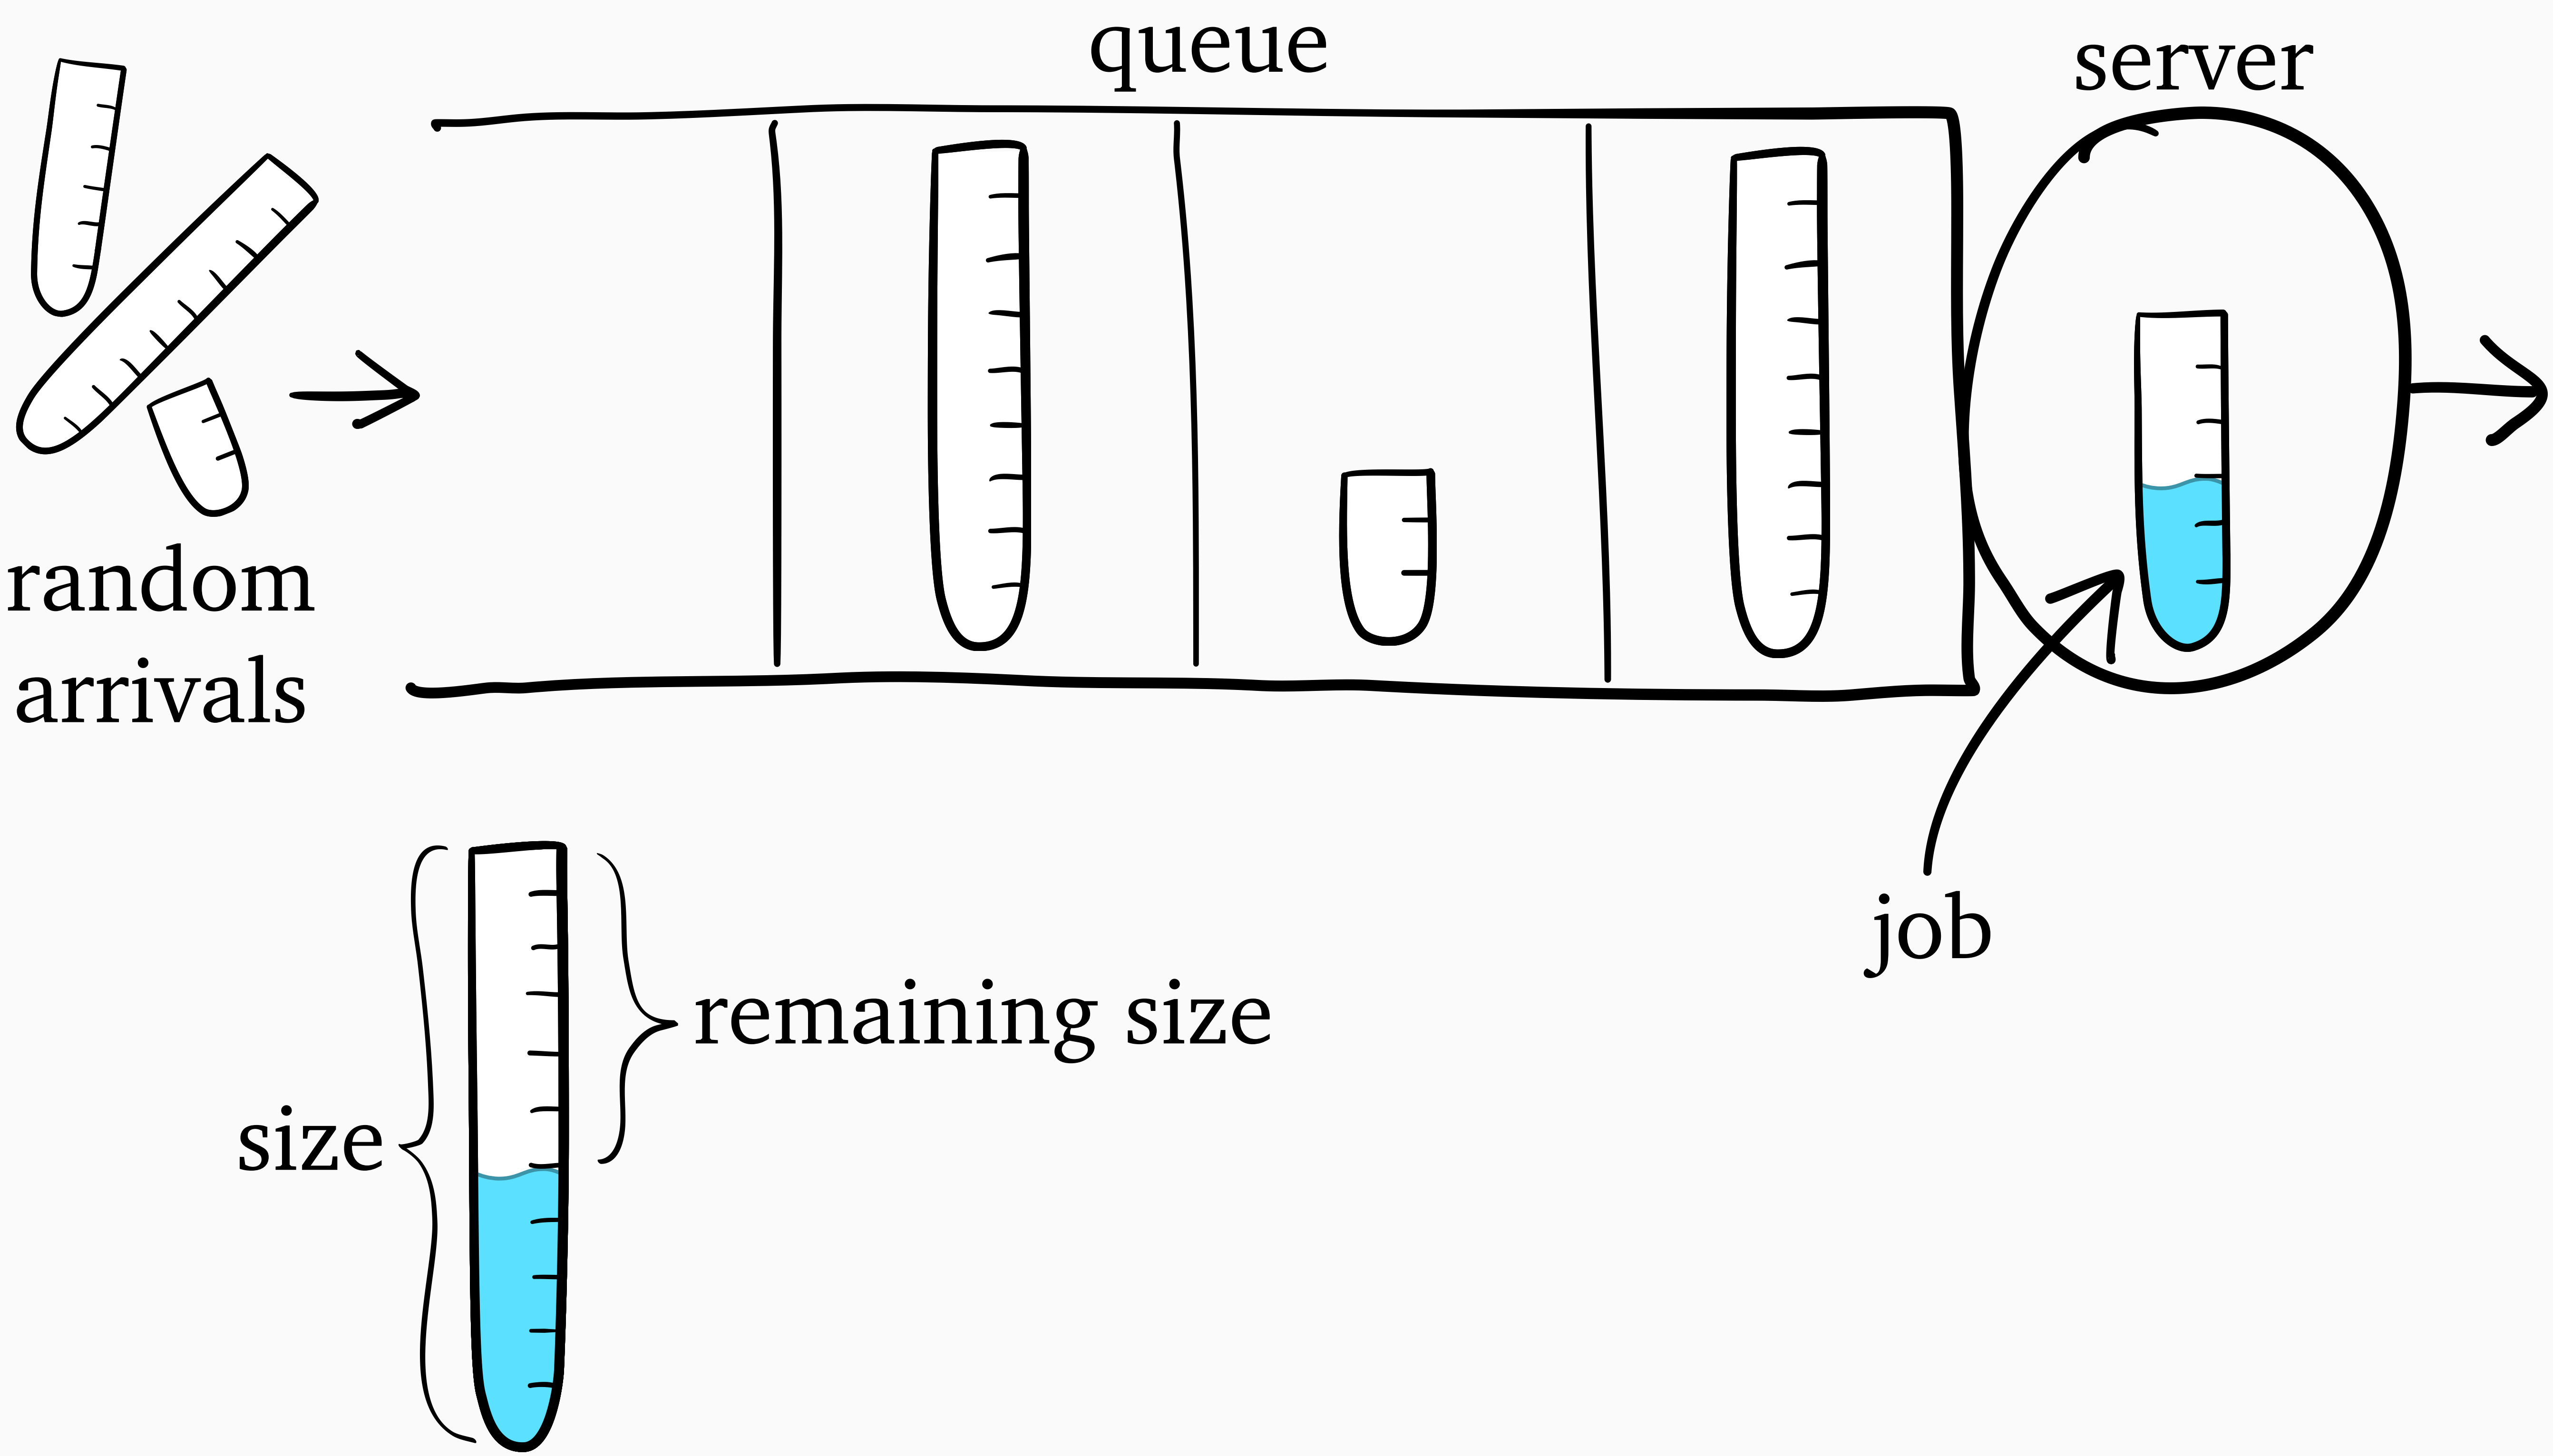

M/G/1 Queue

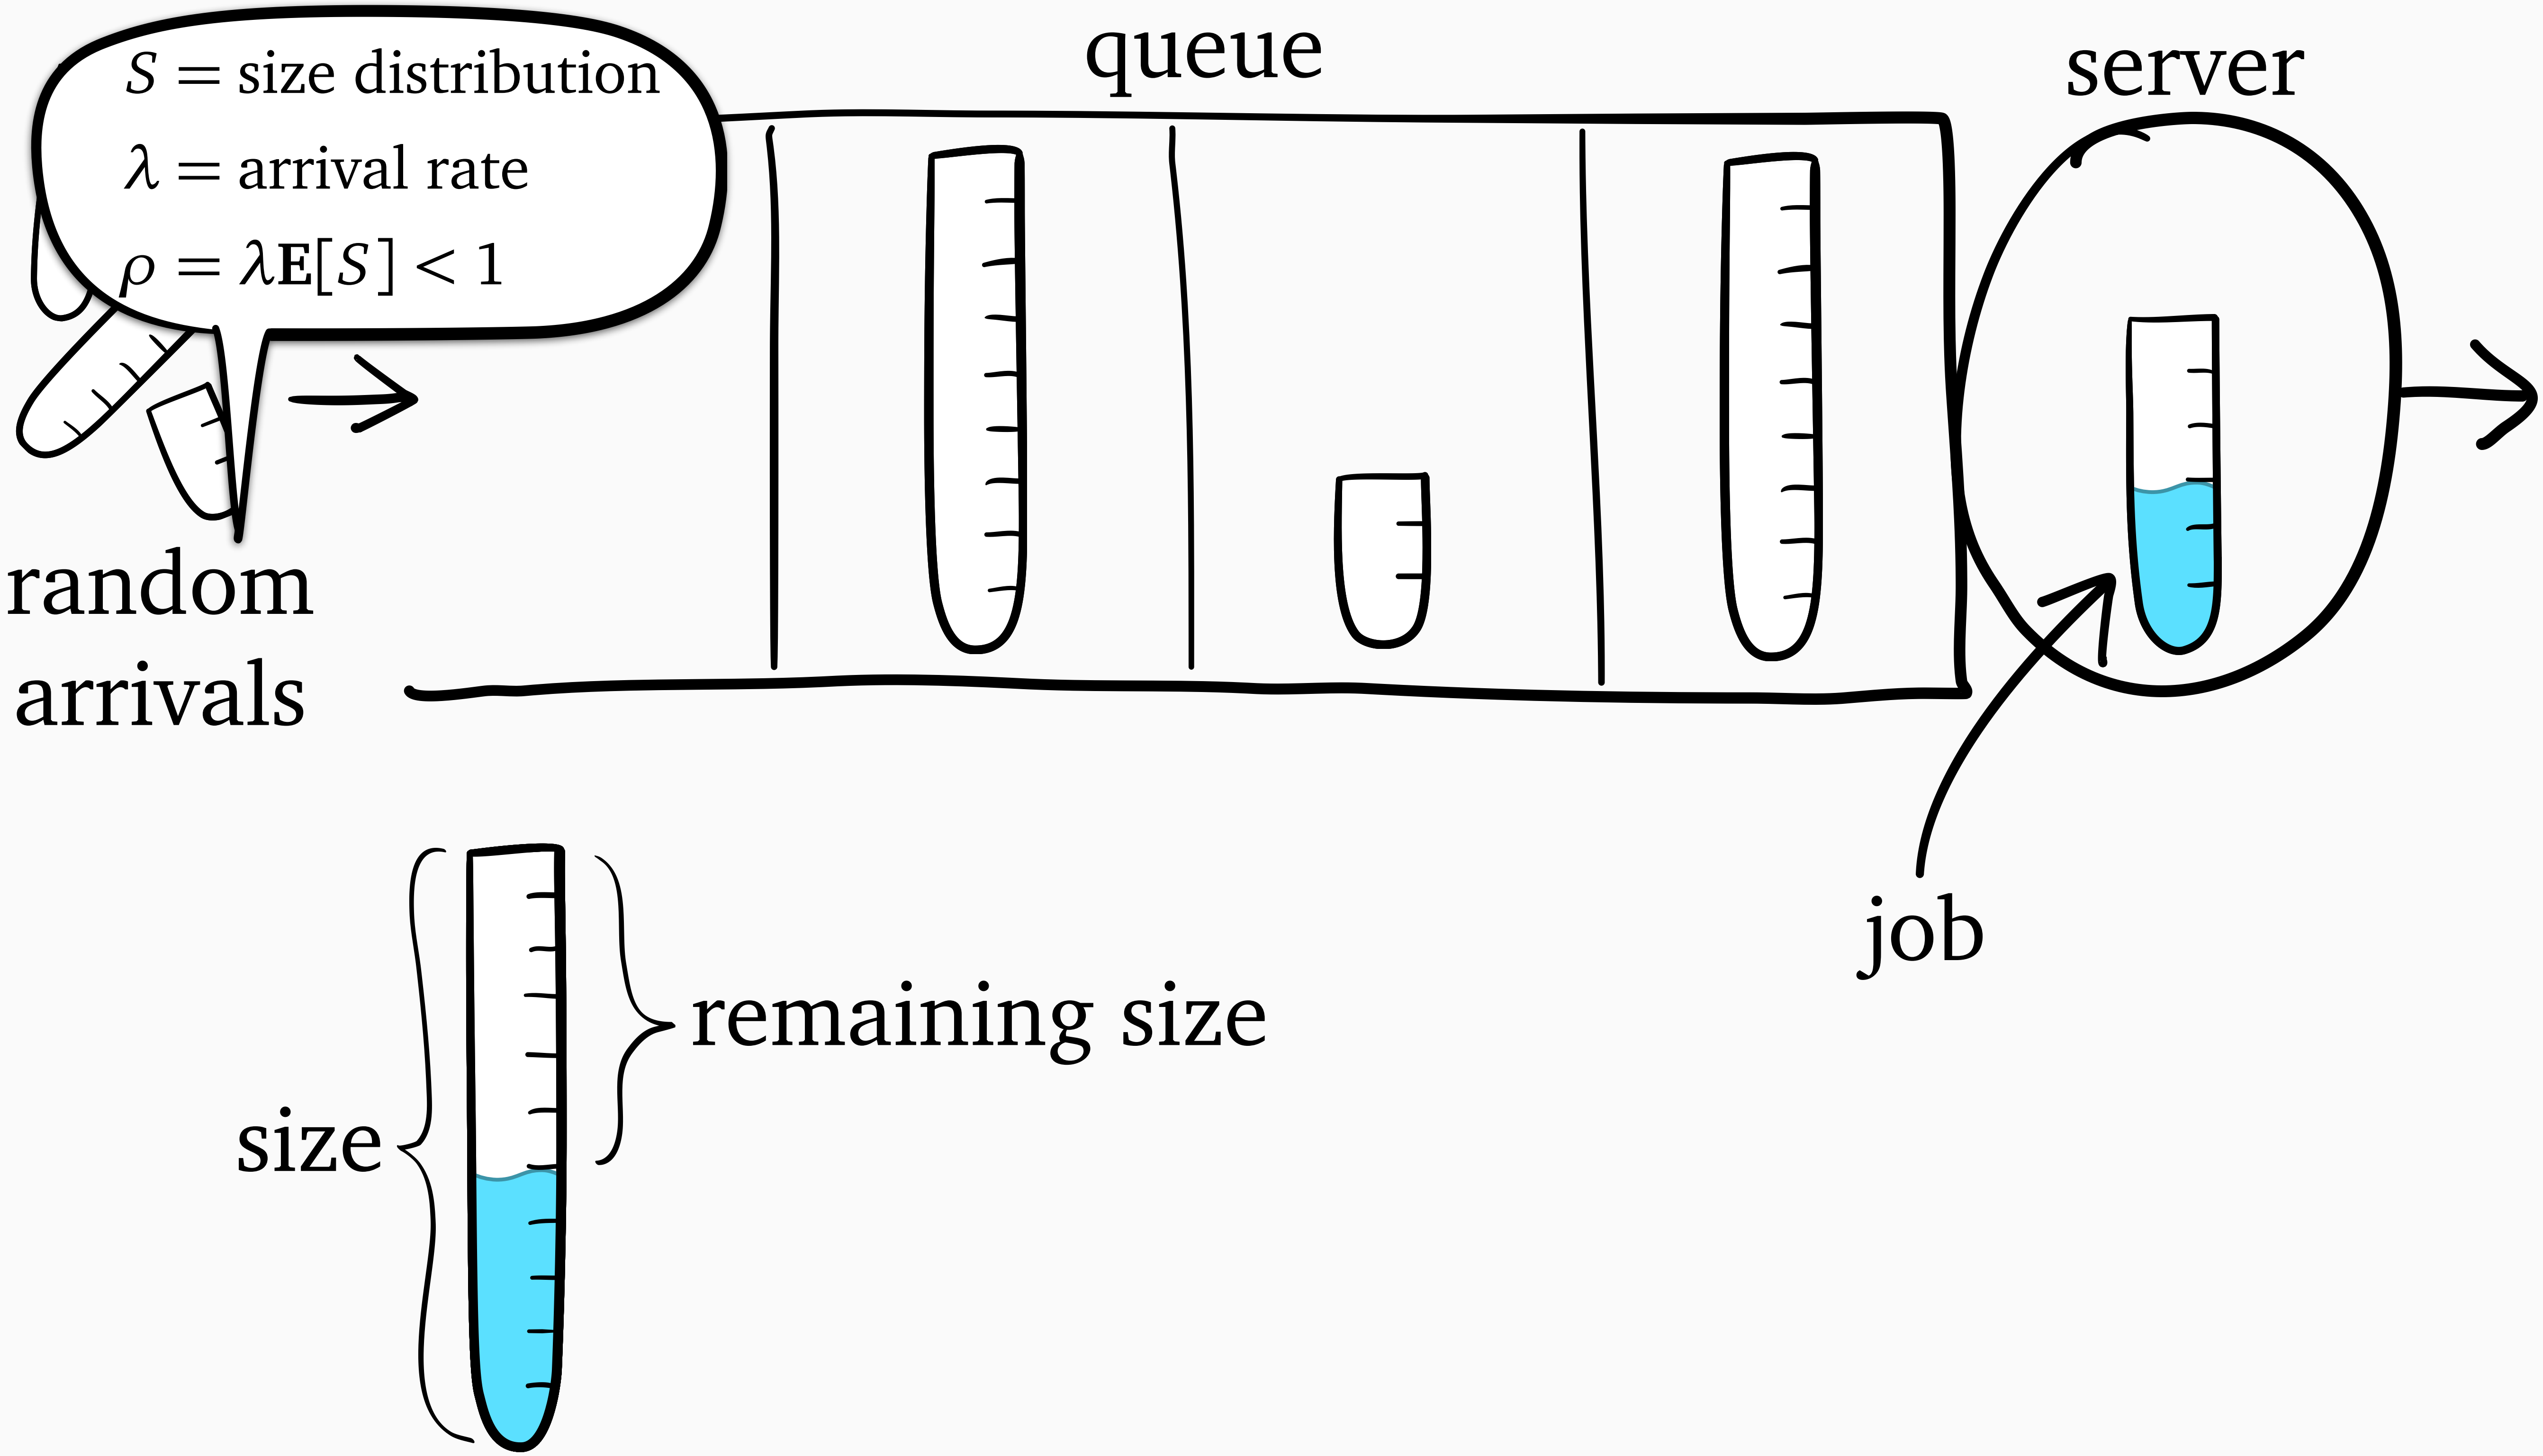

$S =$ size distribution
$\lambda =$ arrival rate
$\rho = \lambda \mathbf{E}[S] < 1$
queue
server
random arrivals
job
remaining size
size

M/G/1 Queue

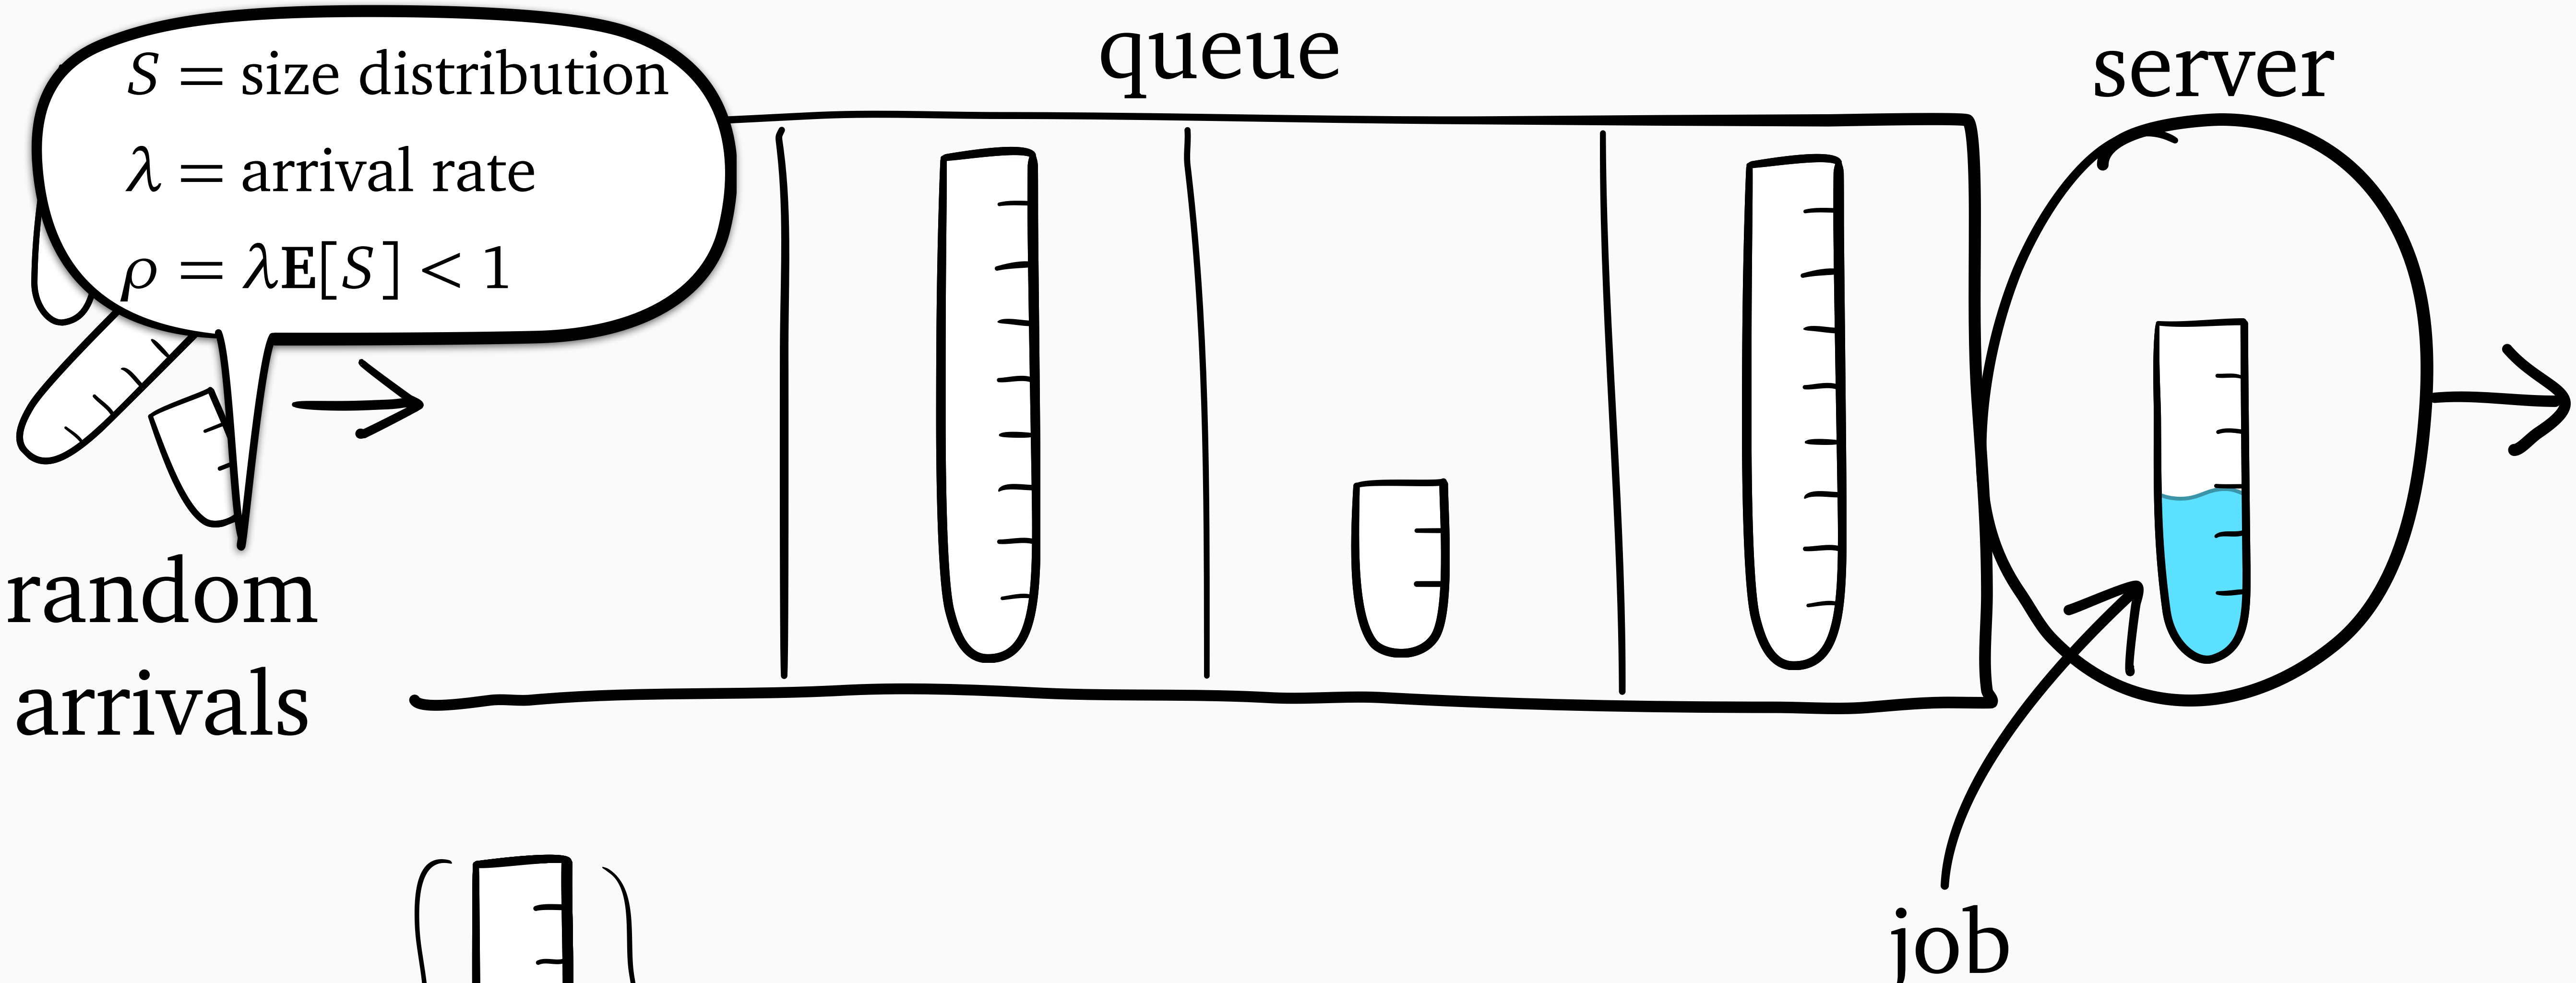

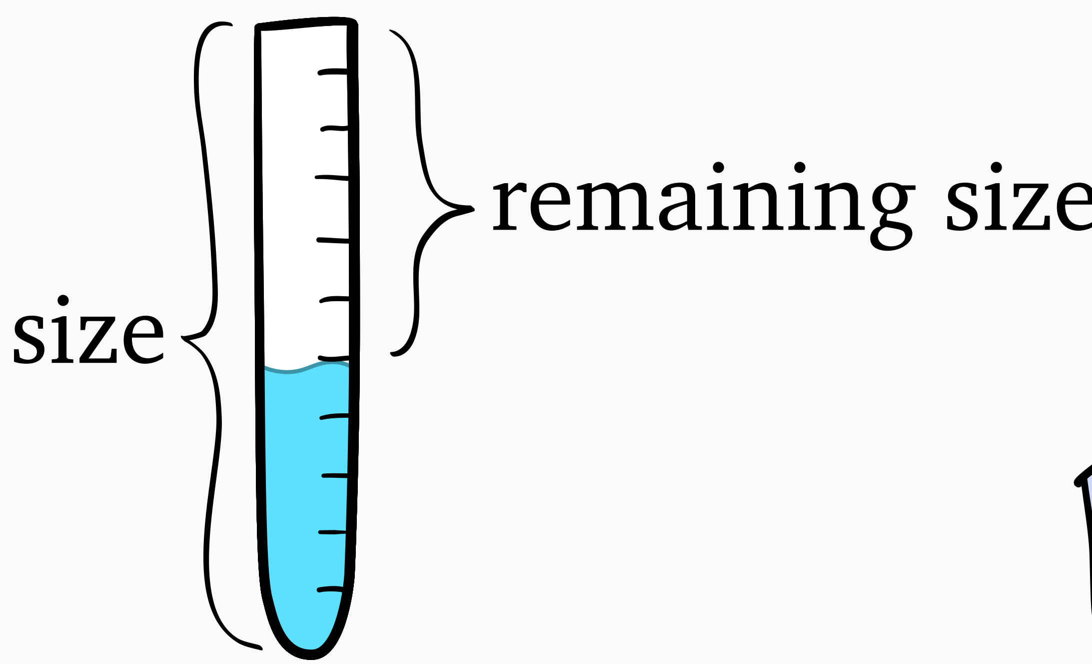

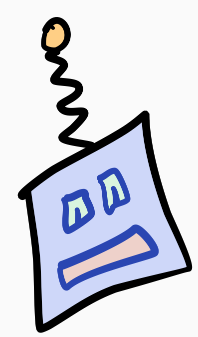

Scheduling policy:
picks which job to serve

M/G/1 Queue

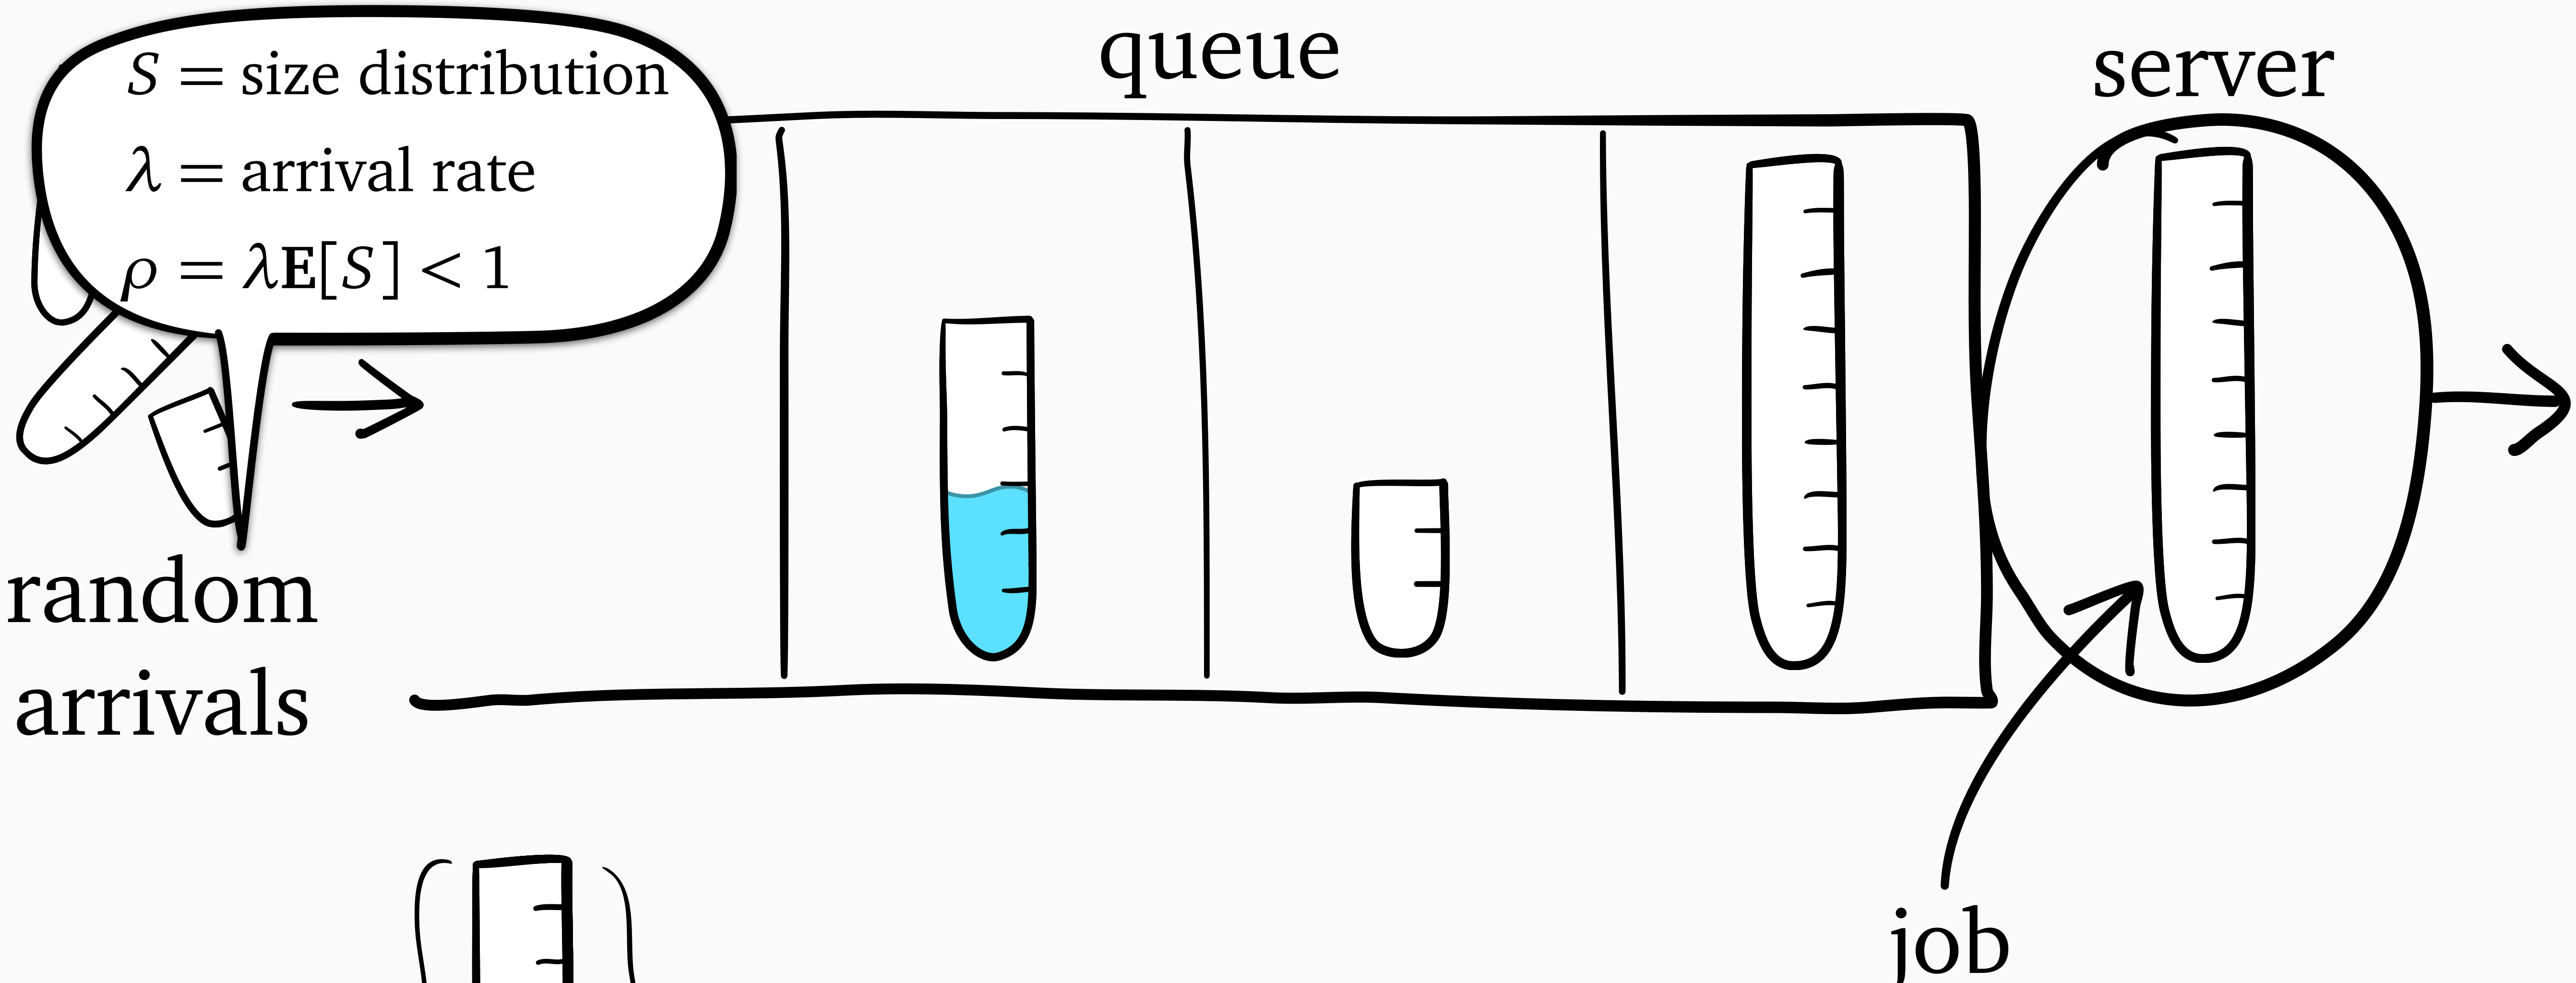

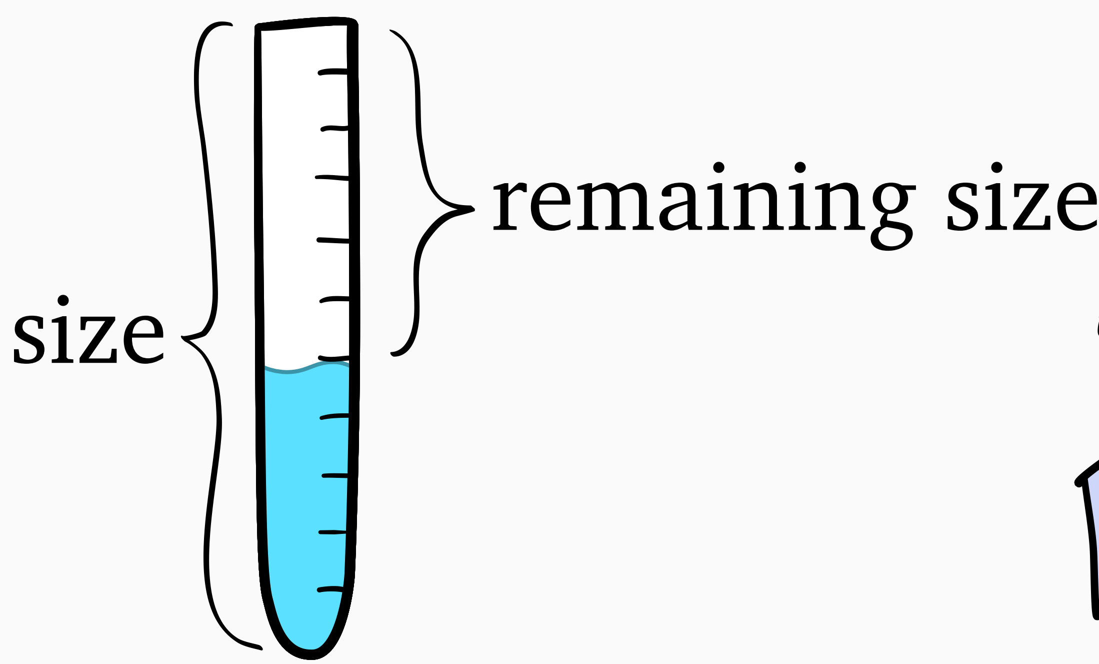

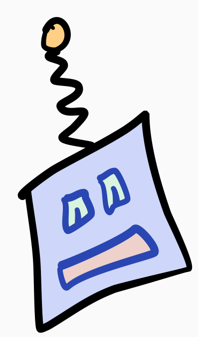

Scheduling policy:
picks which job to serve

M/G/1 Queue

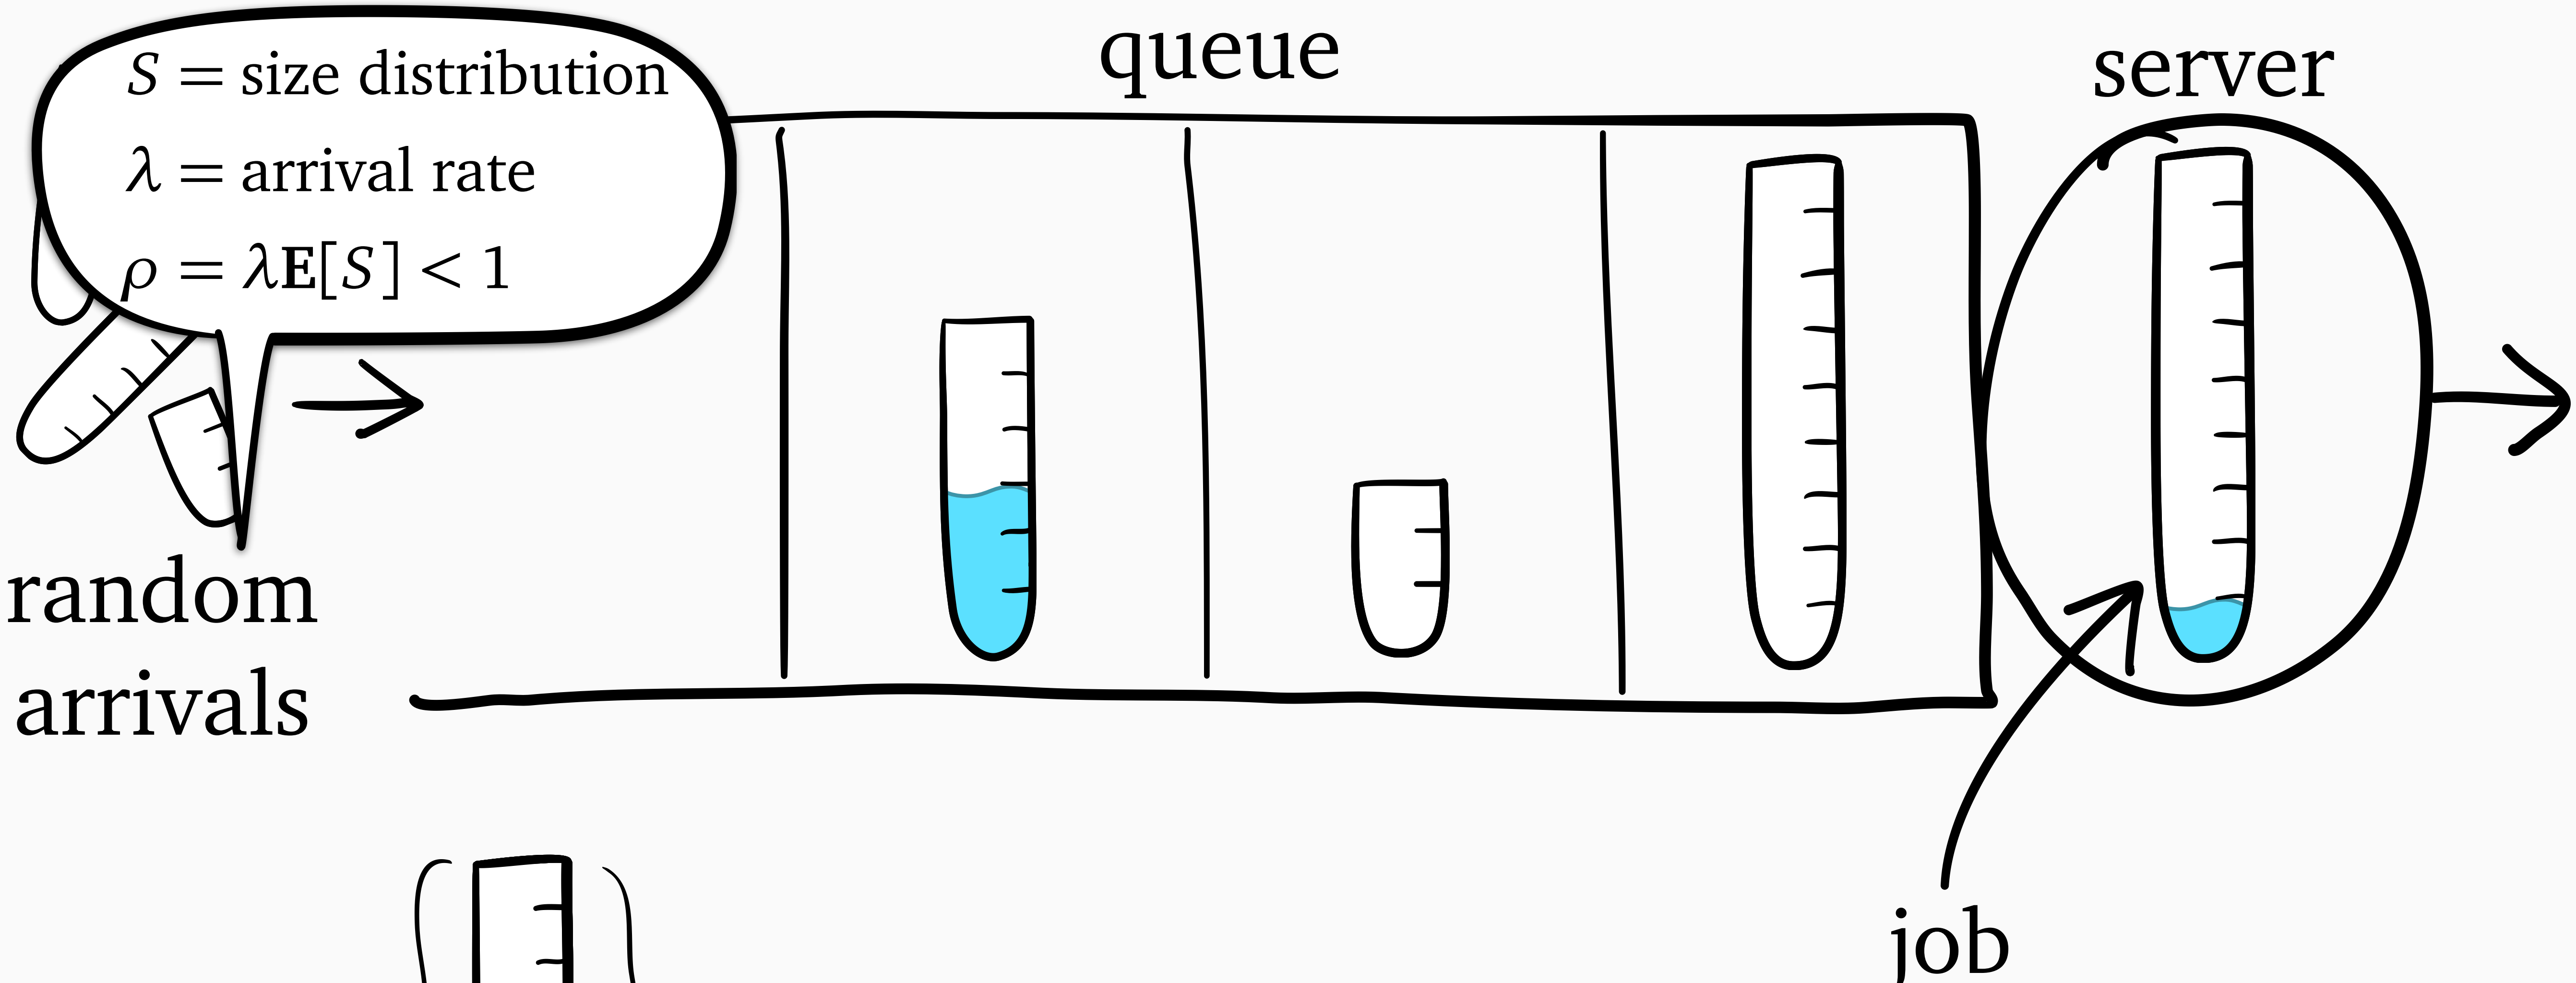

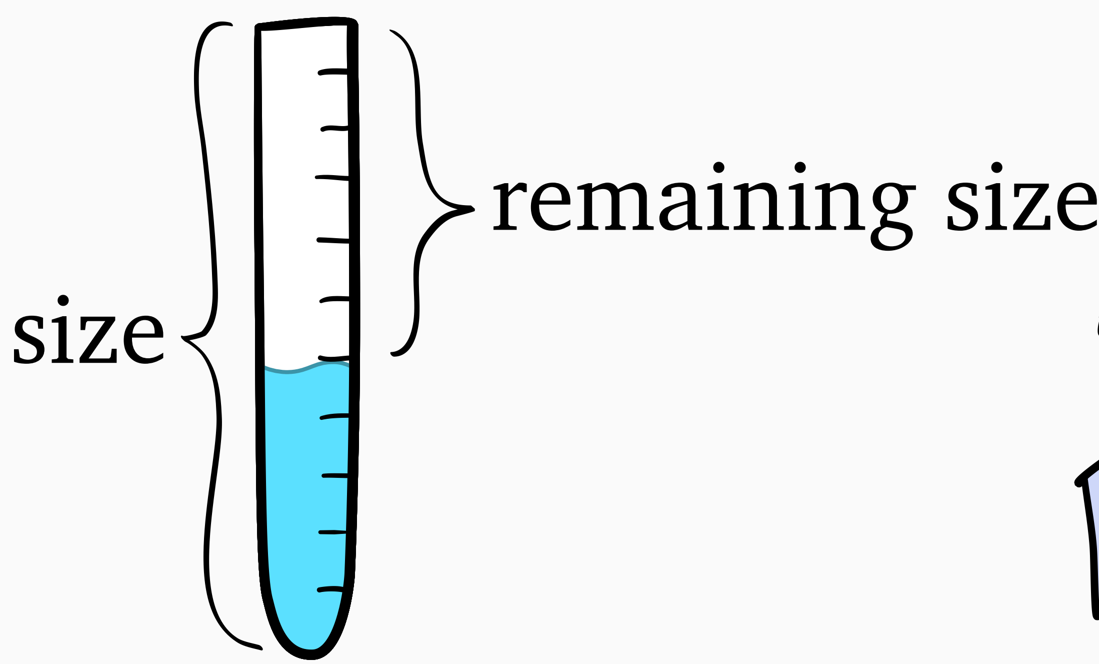

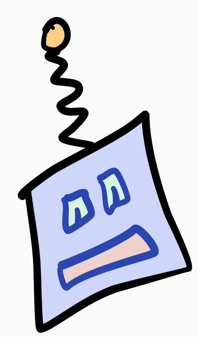

Scheduling policy:
picks which job to serve

M/G/1 Queue

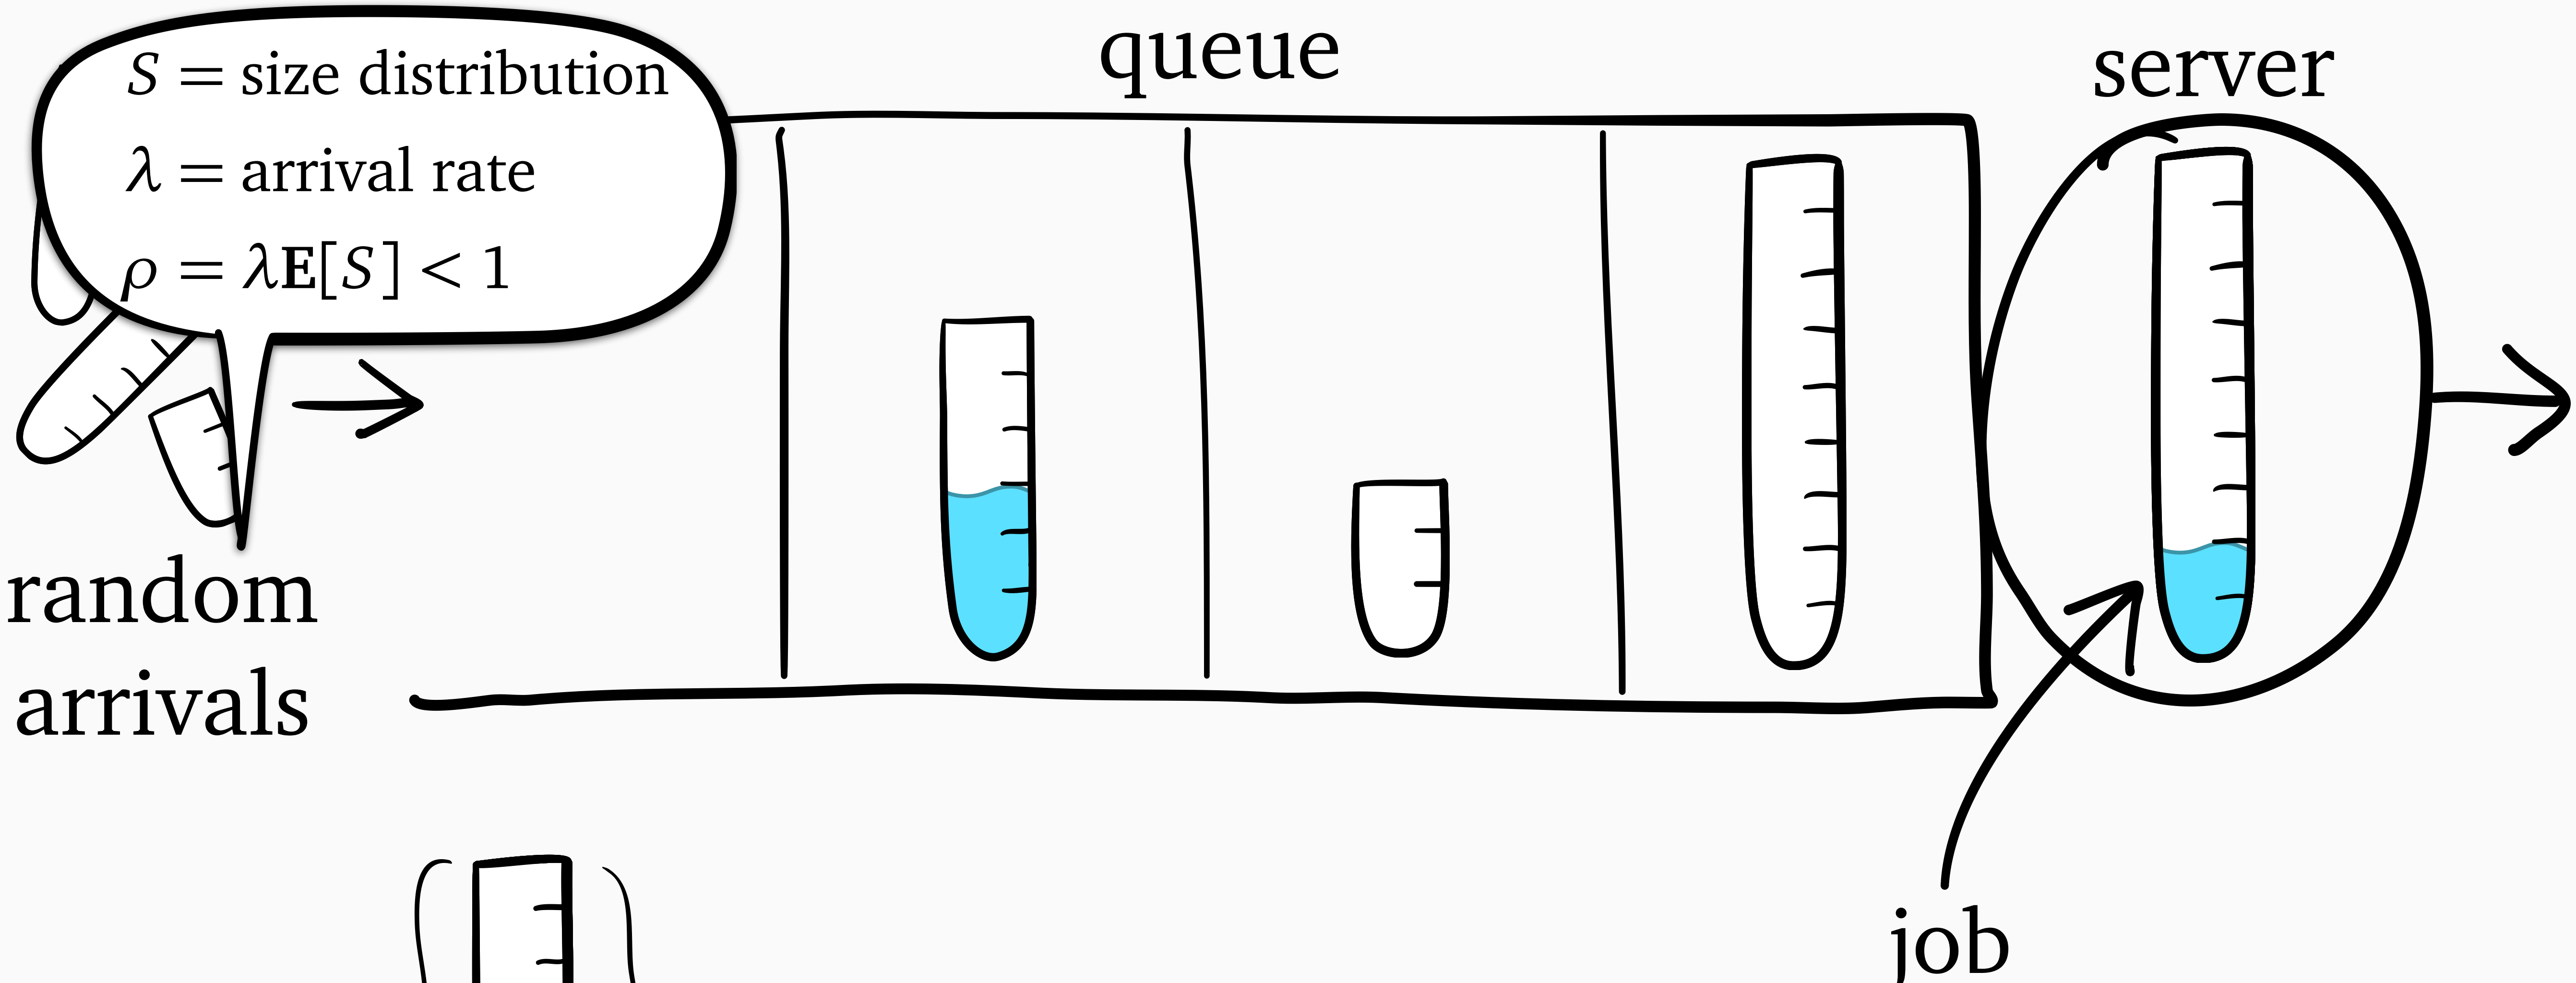

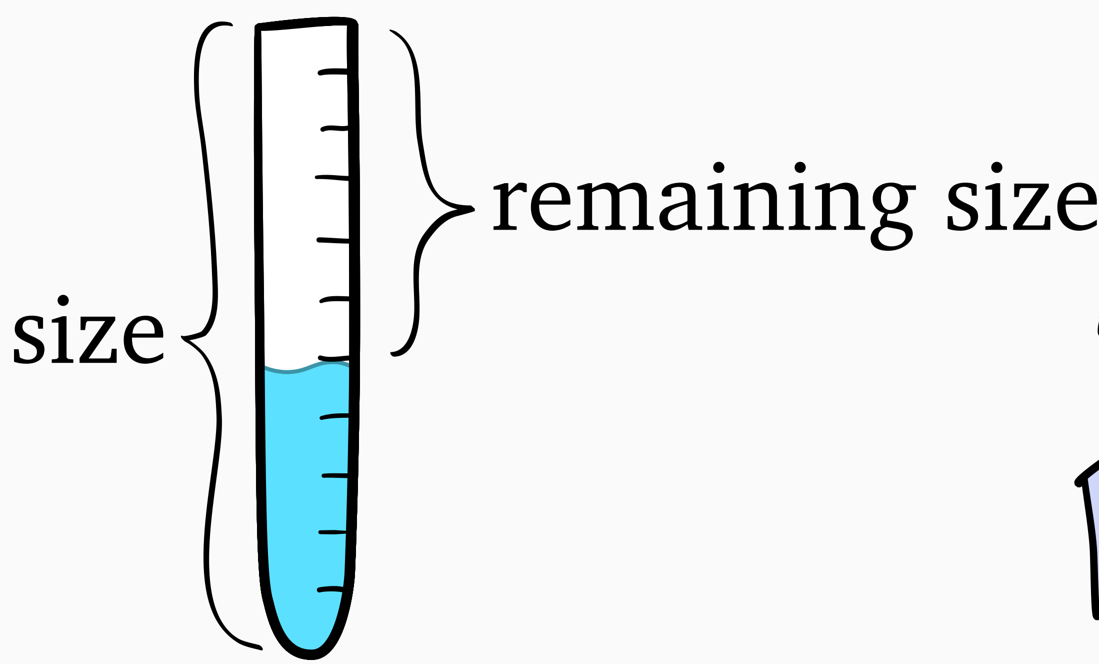

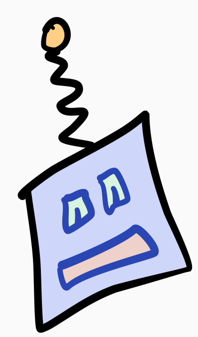

Scheduling policy:
picks which job to serve

M/G/1 Queue

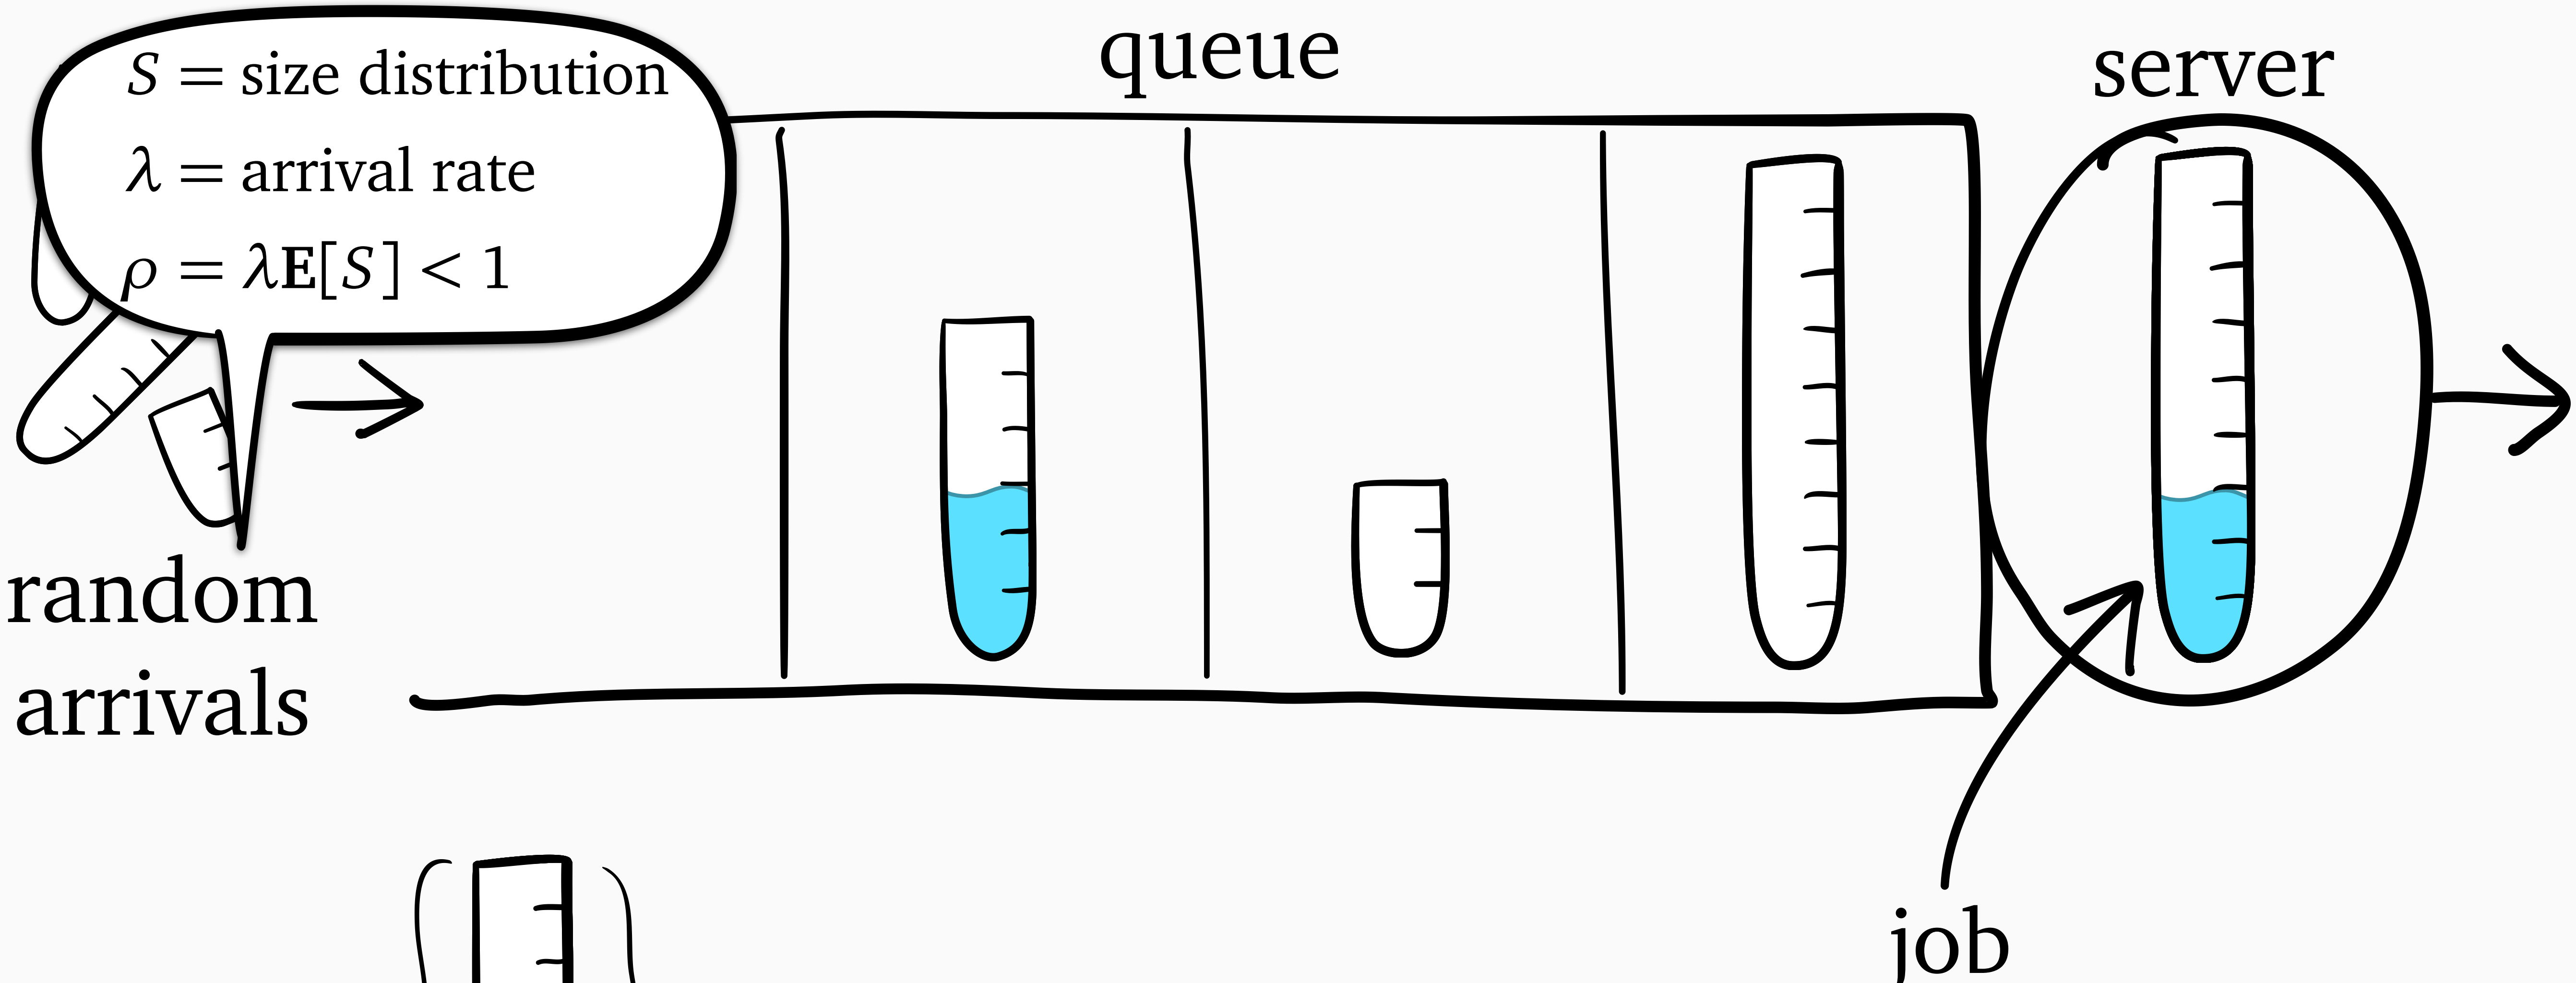

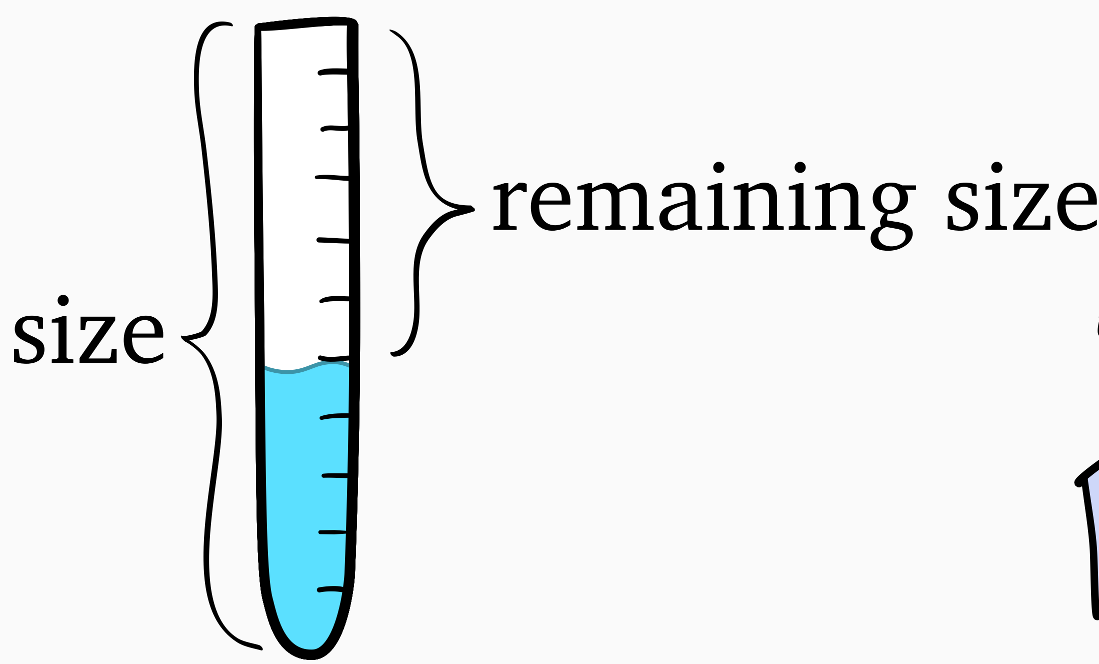

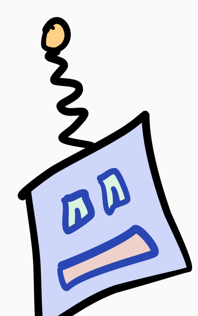

Scheduling policy:
picks which job to serve

Response Time and Slowdown

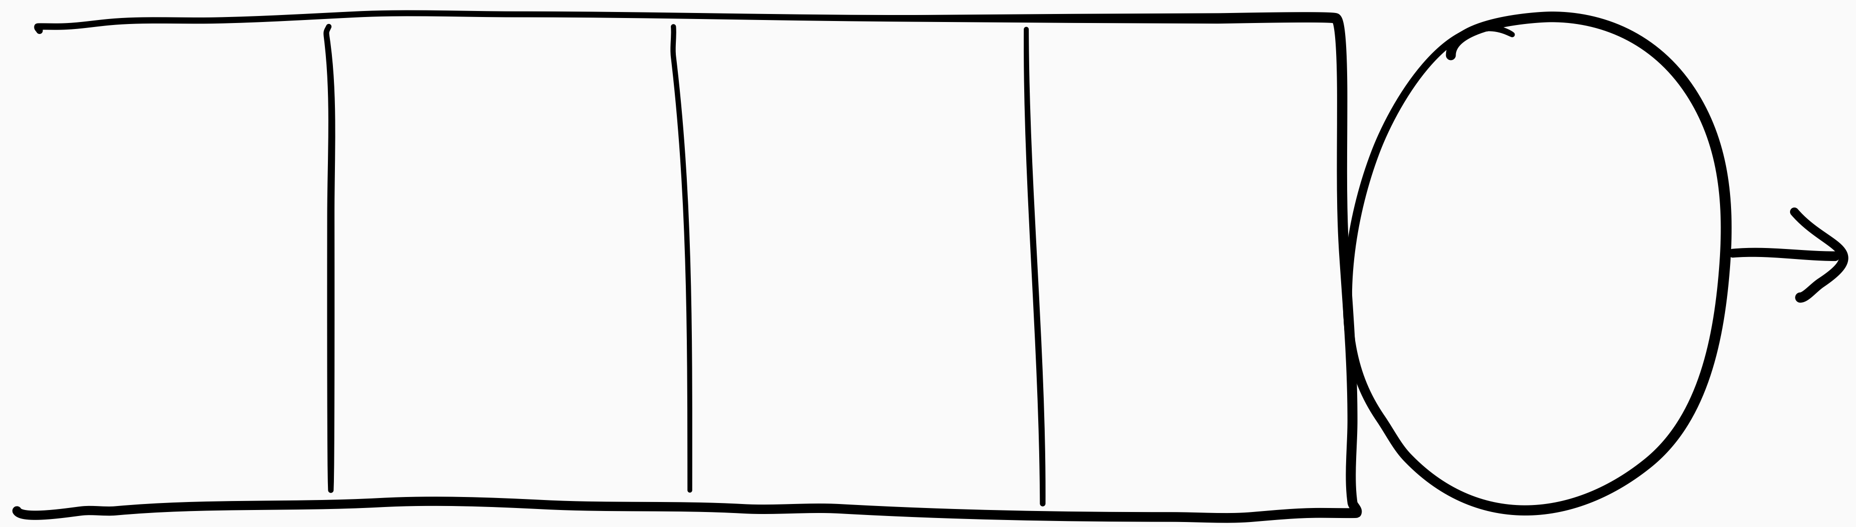

Response Time and Slowdown

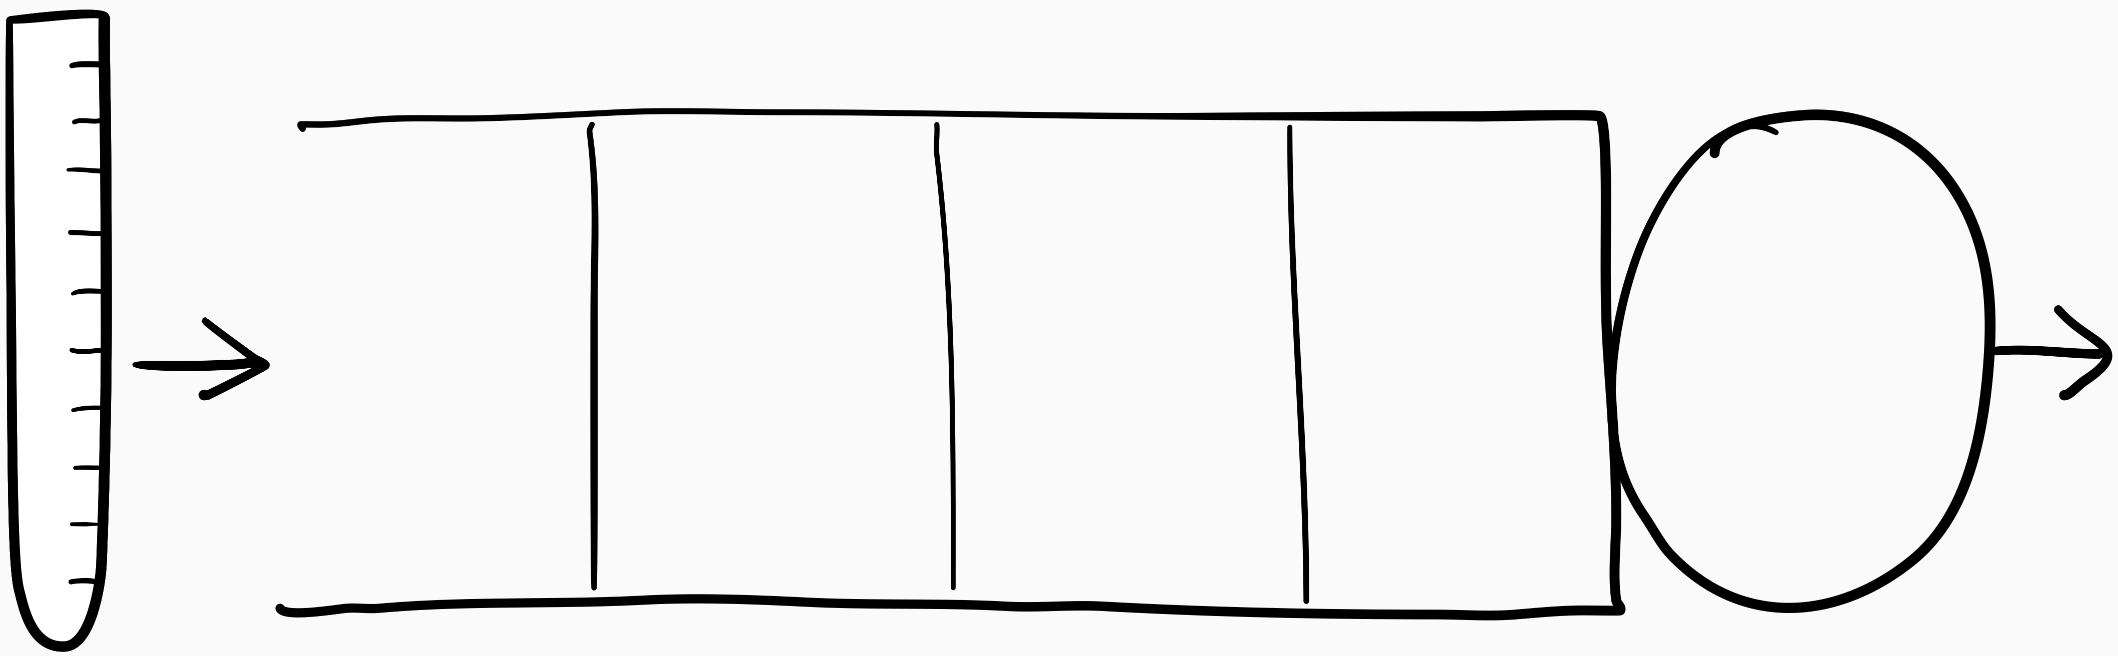

Response Time and Slowdown

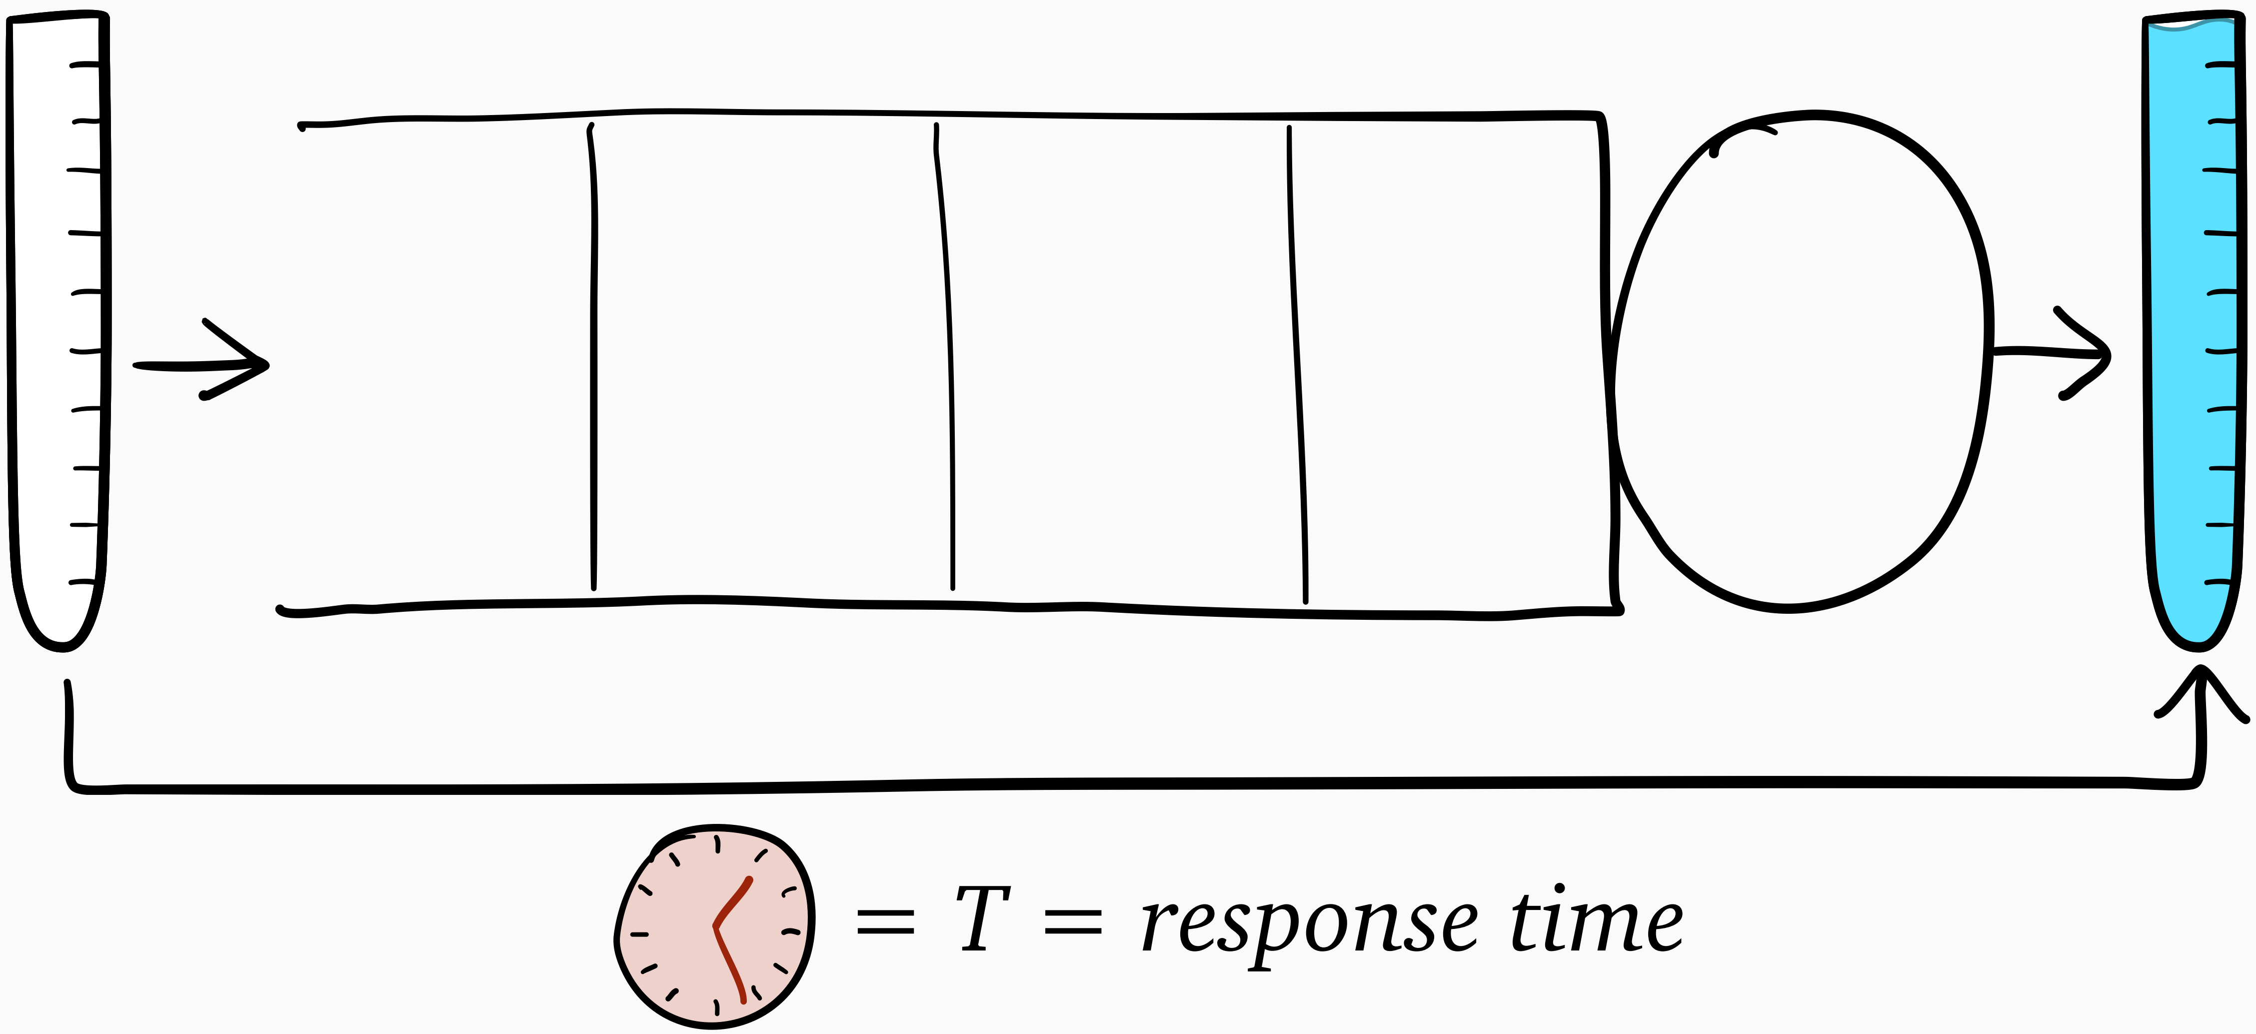

$= T =$ *response time*

Response Time and Slowdown

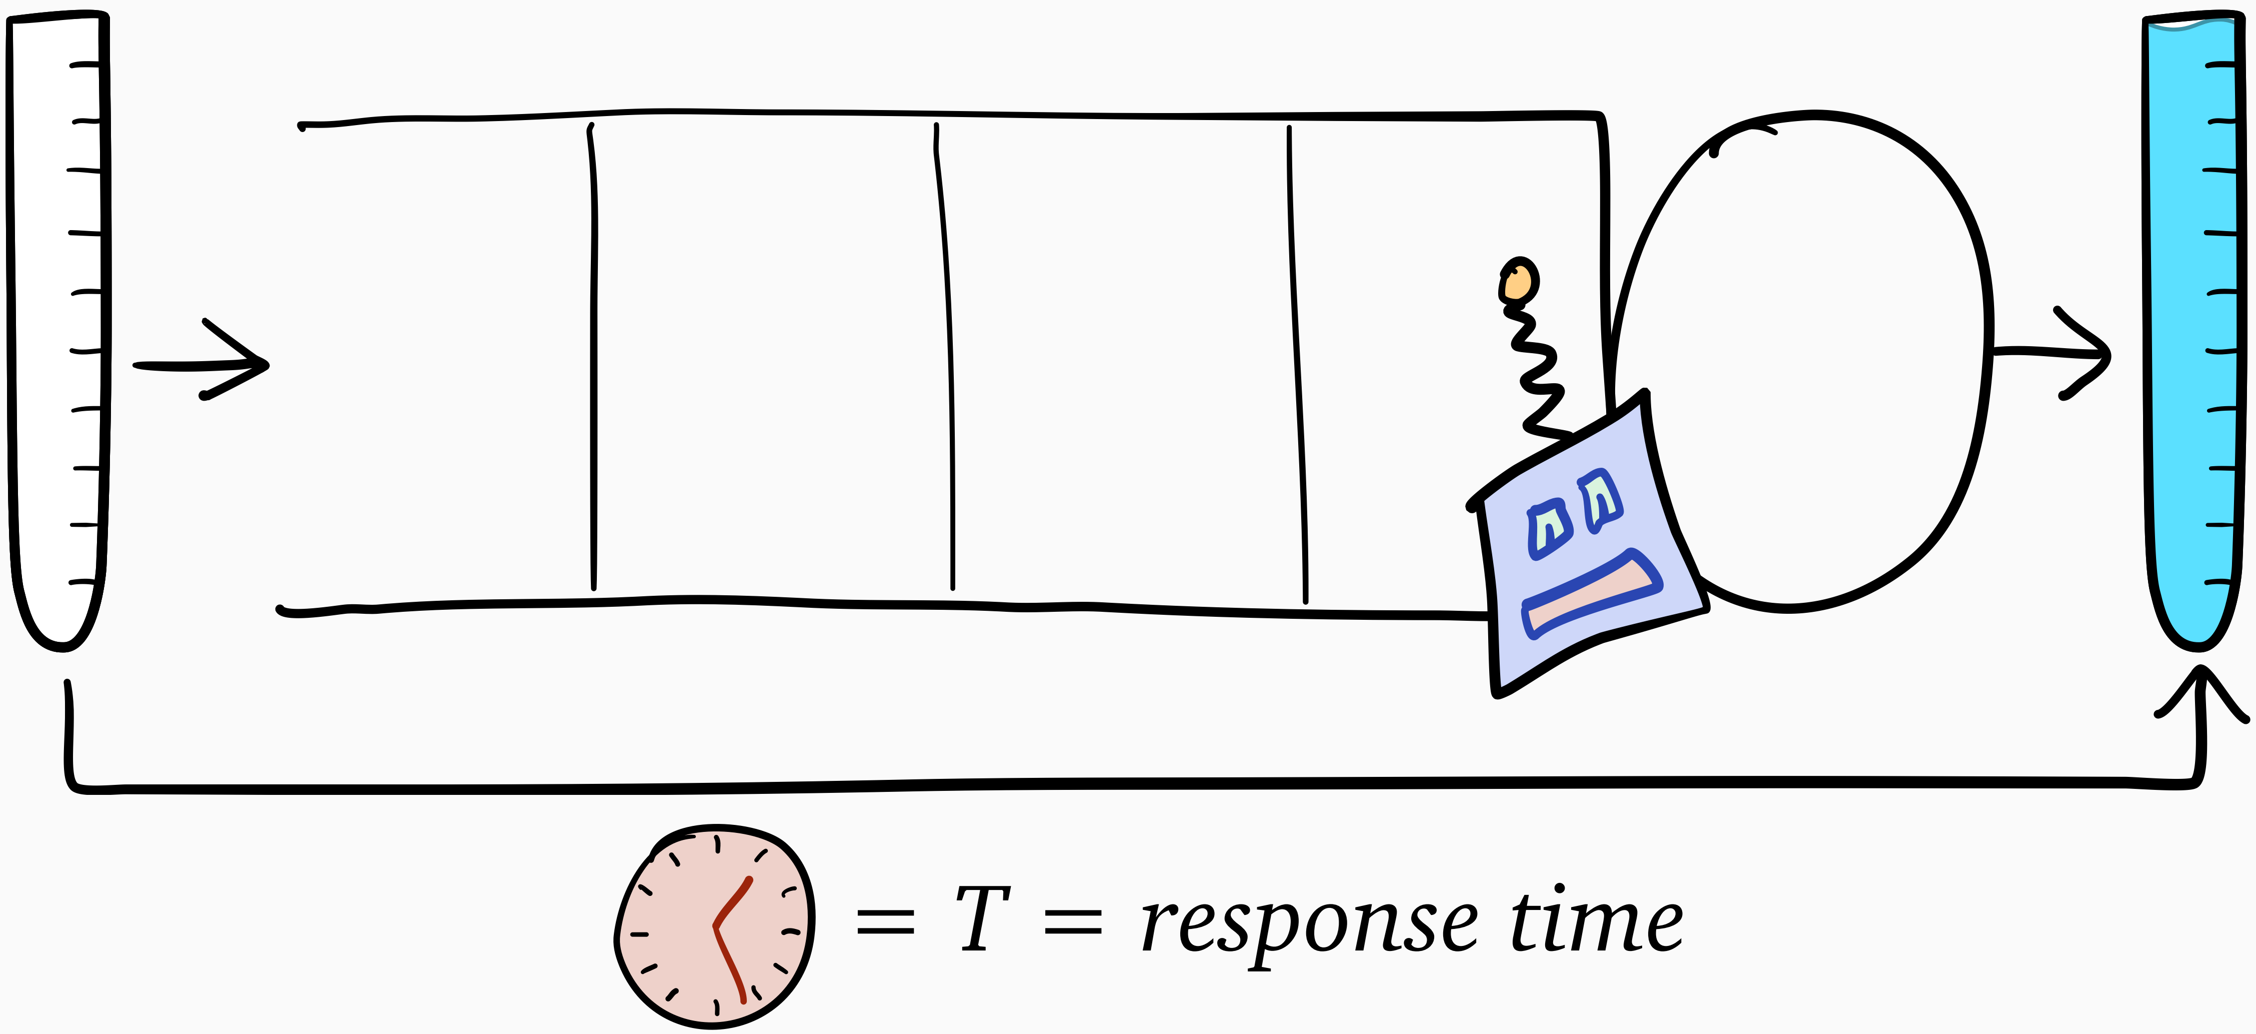

$= T =$ *response time*

Response Time and Slowdown

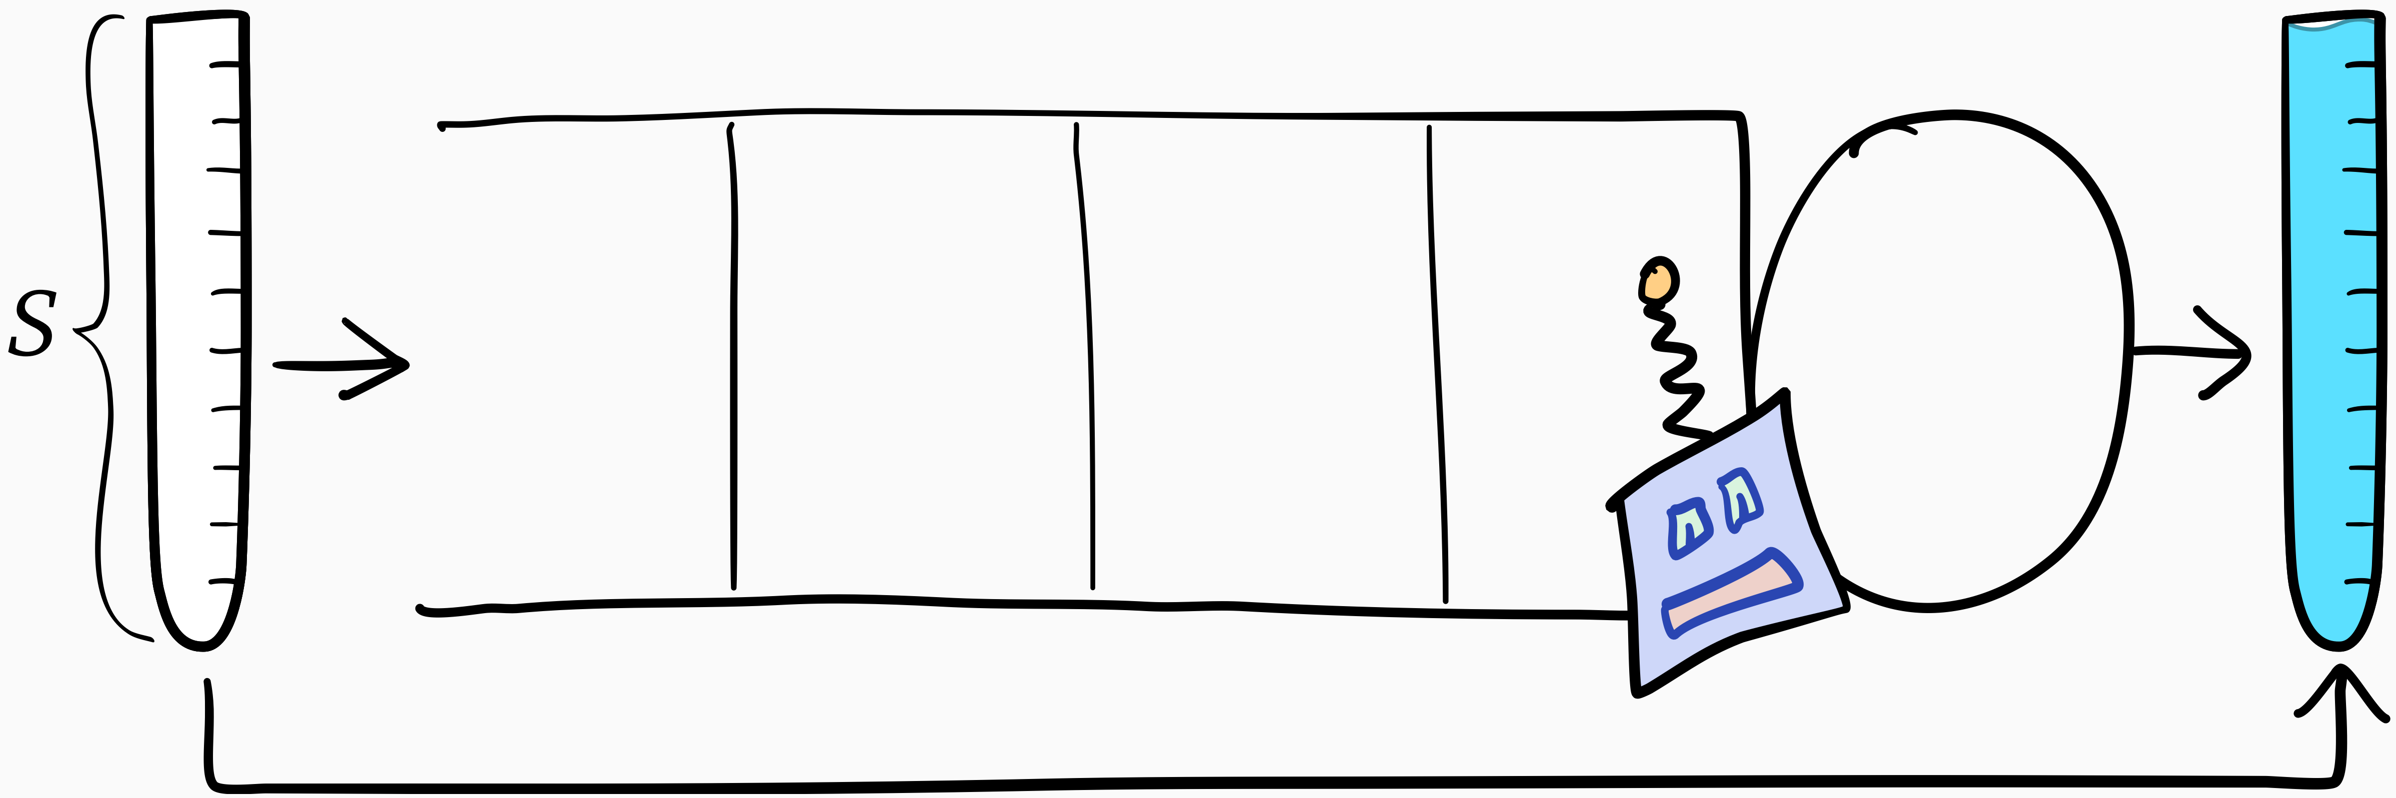

 $= T =$ *response time*

Response Time and Slowdown

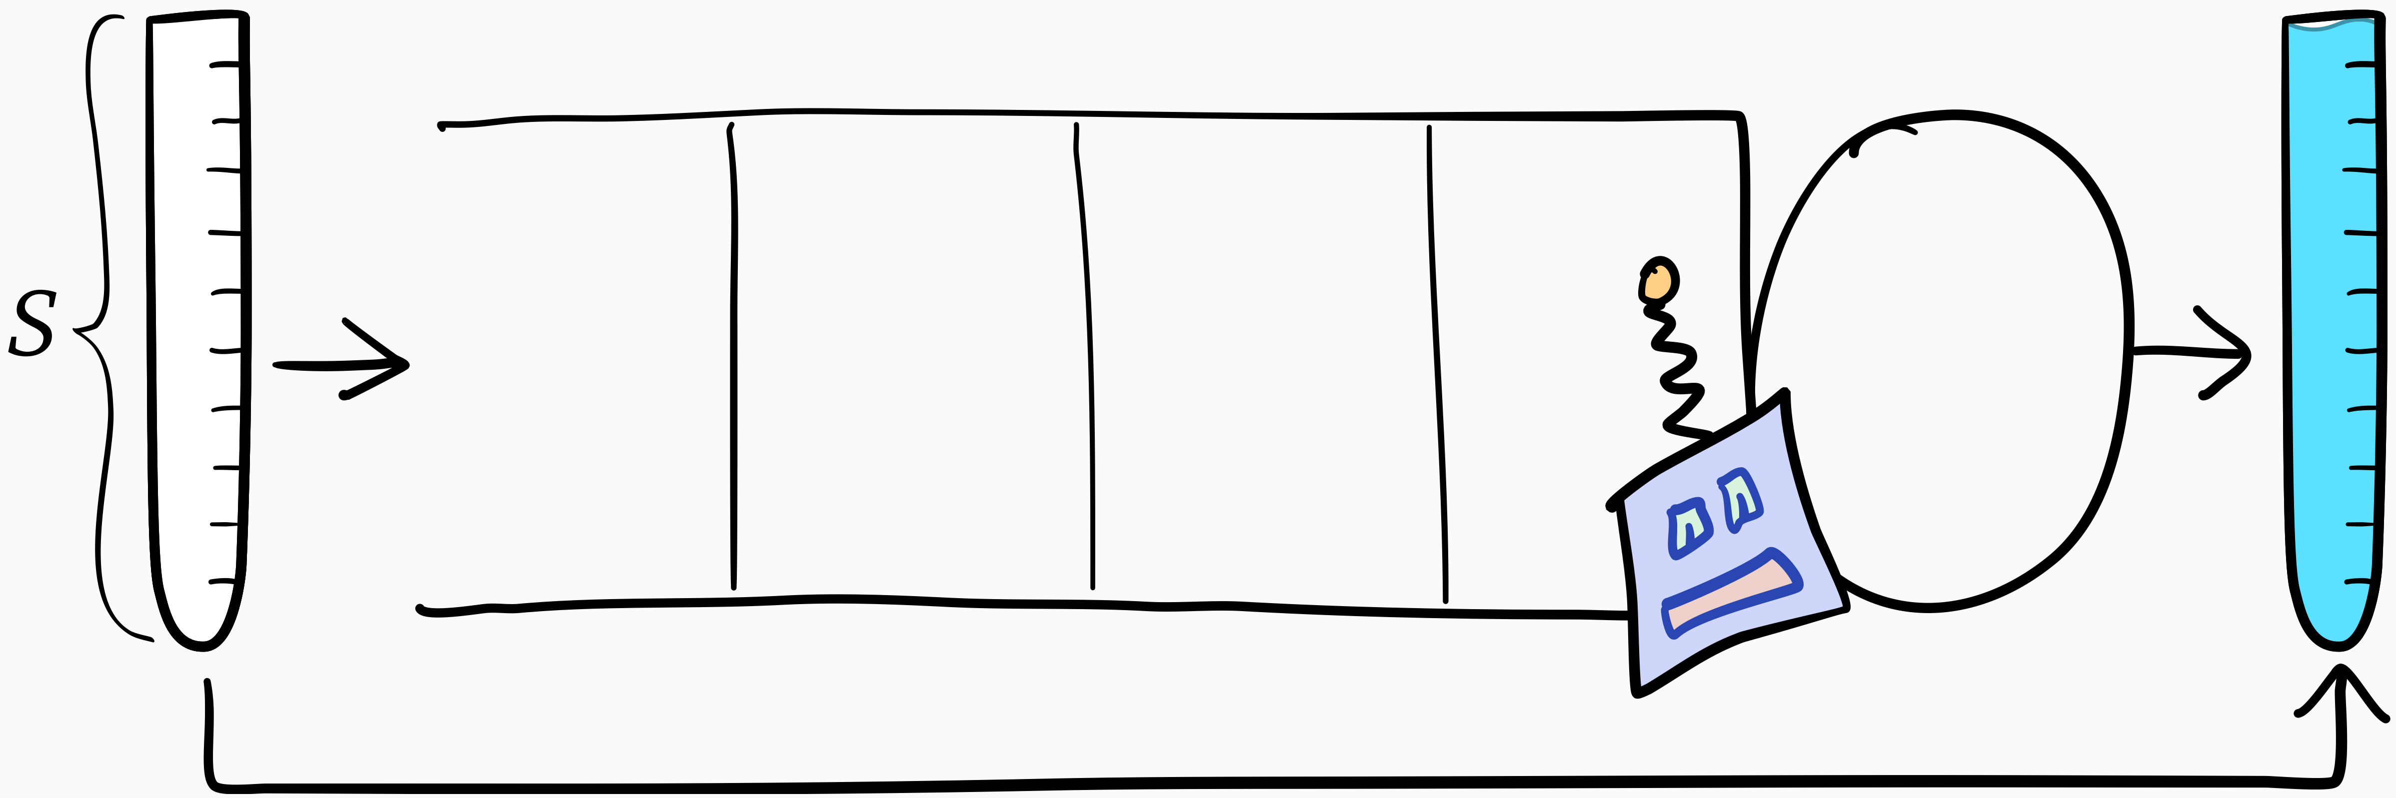

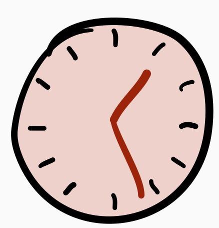 $= T =$ *response time*

$slowdown = Z = T/S =$ response time / size

Response Time and Slowdown

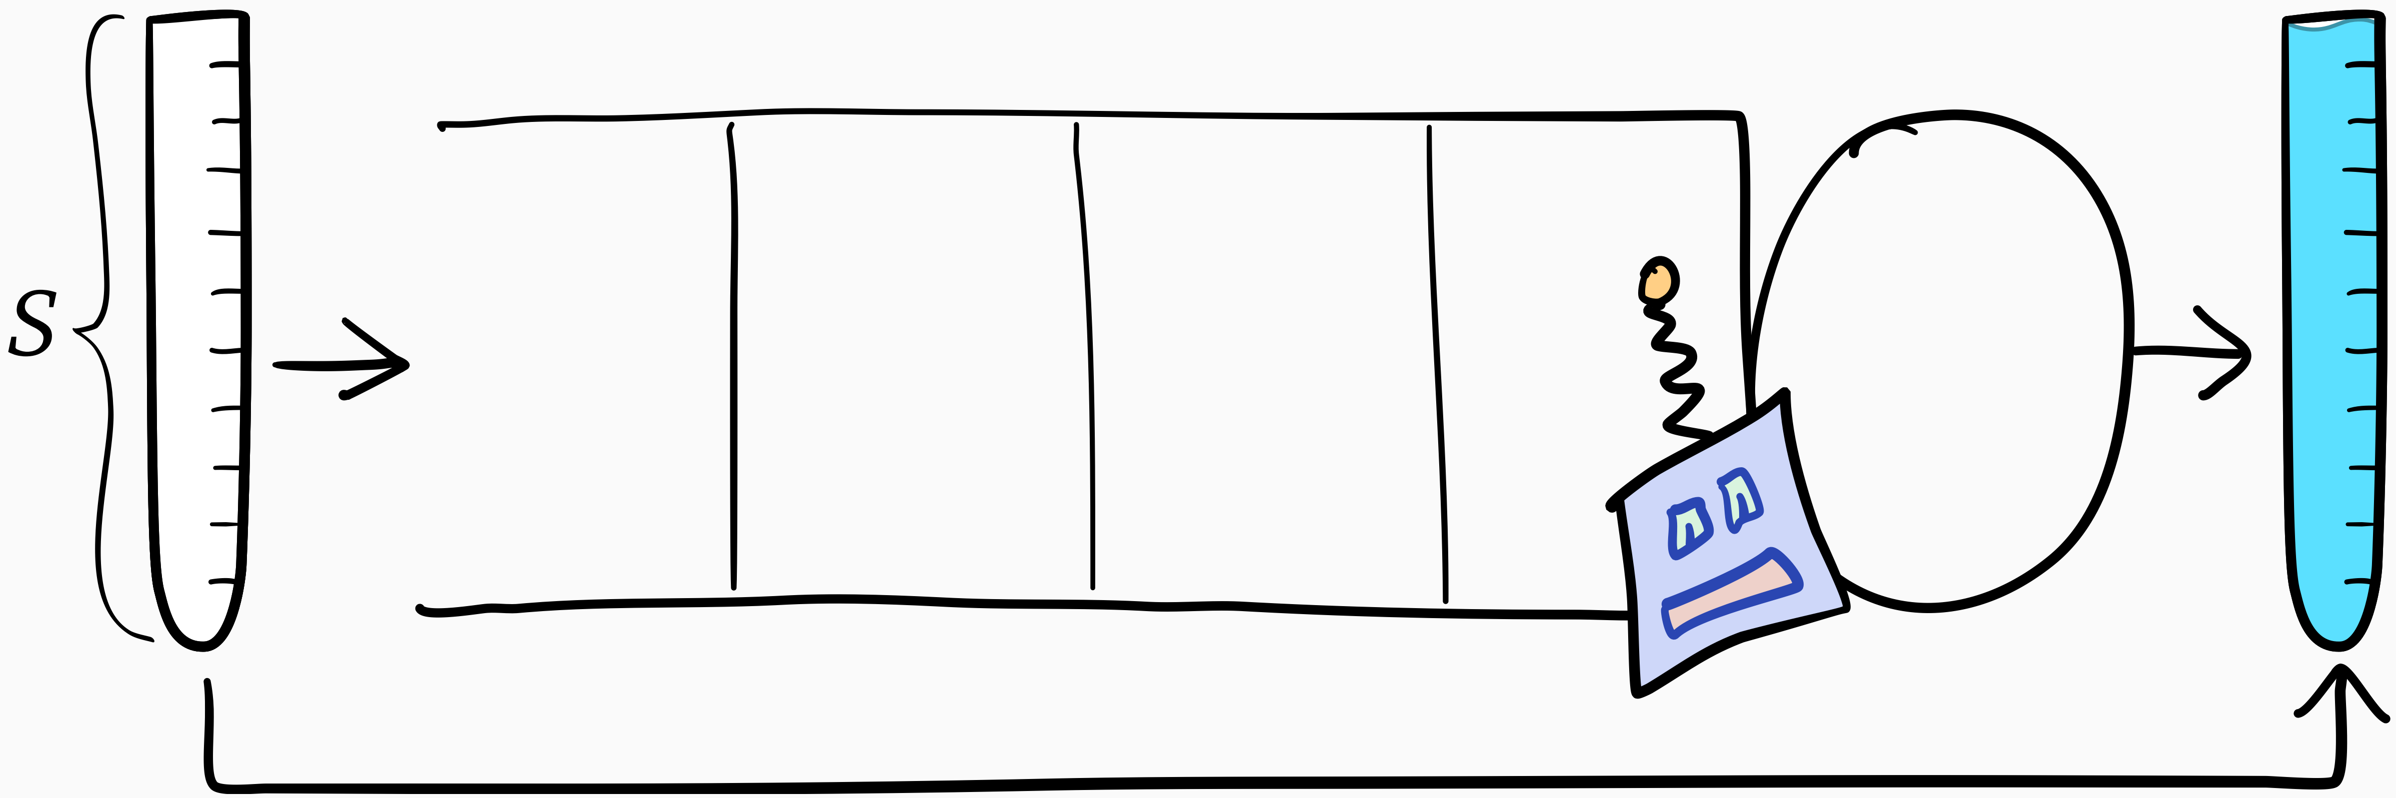

 $= T =$ *response time*

slowdown $= Z = T/S =$ response time / size

Goal: schedule to minimize *mean slowdown* $\mathbf{E}[Z]$

How to Schedule?

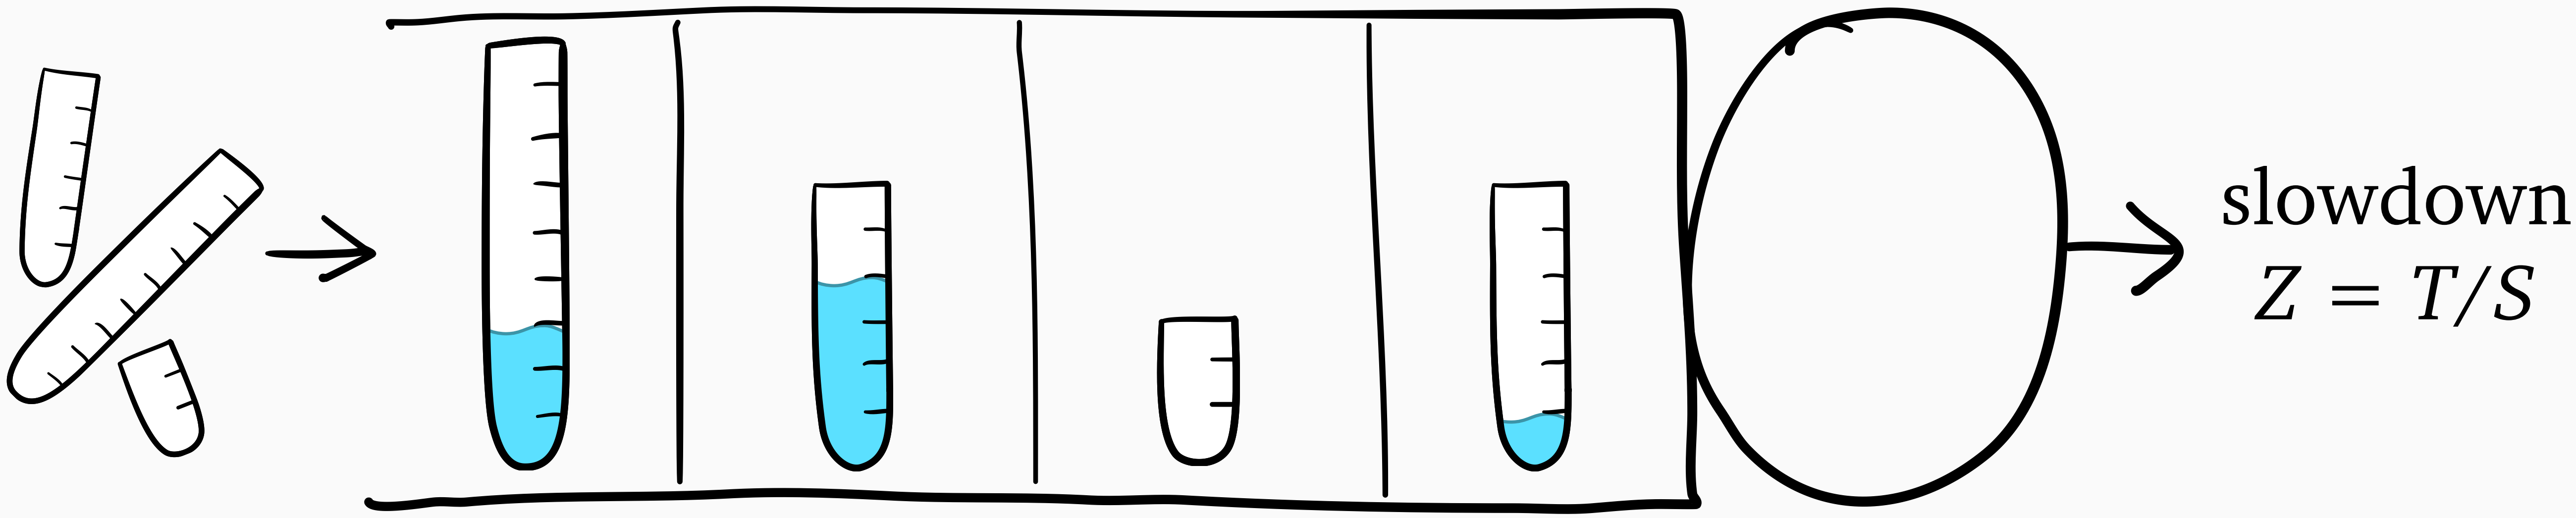

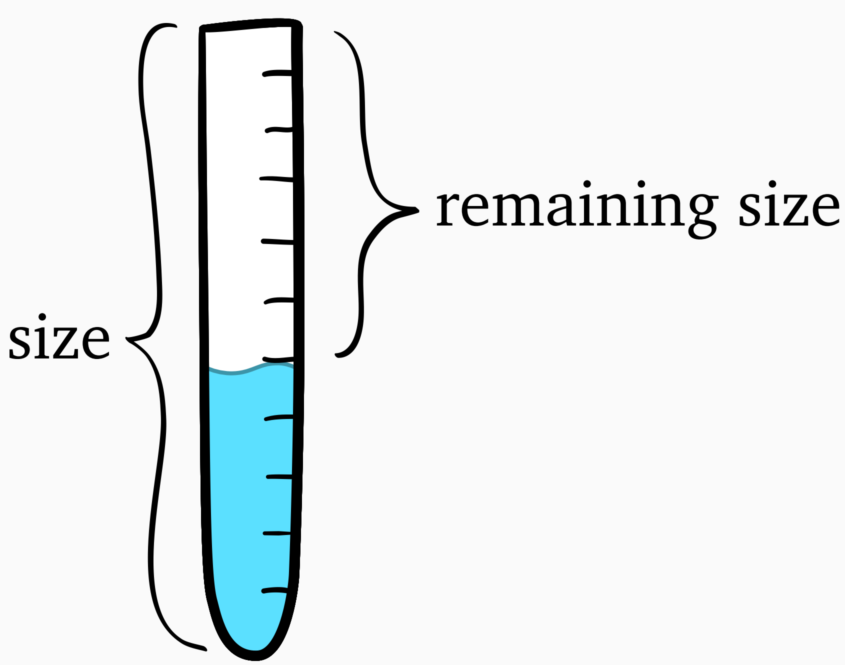

How to Schedule?

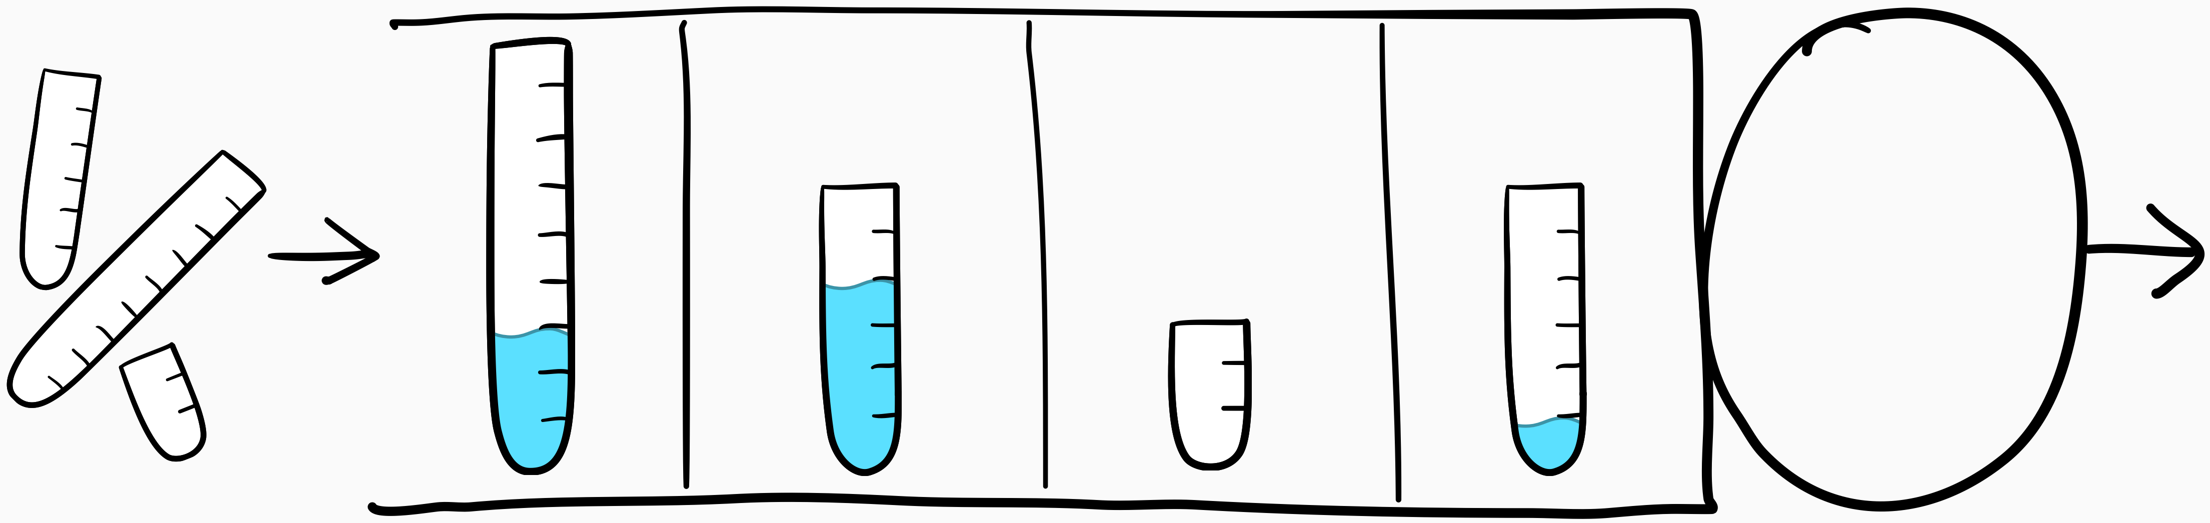

Want to prioritize small *size* and small *remaining size*

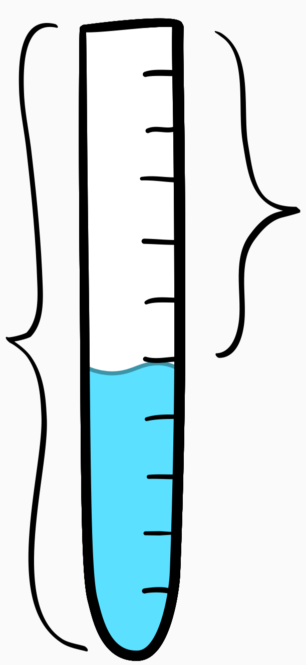

How to Schedule?

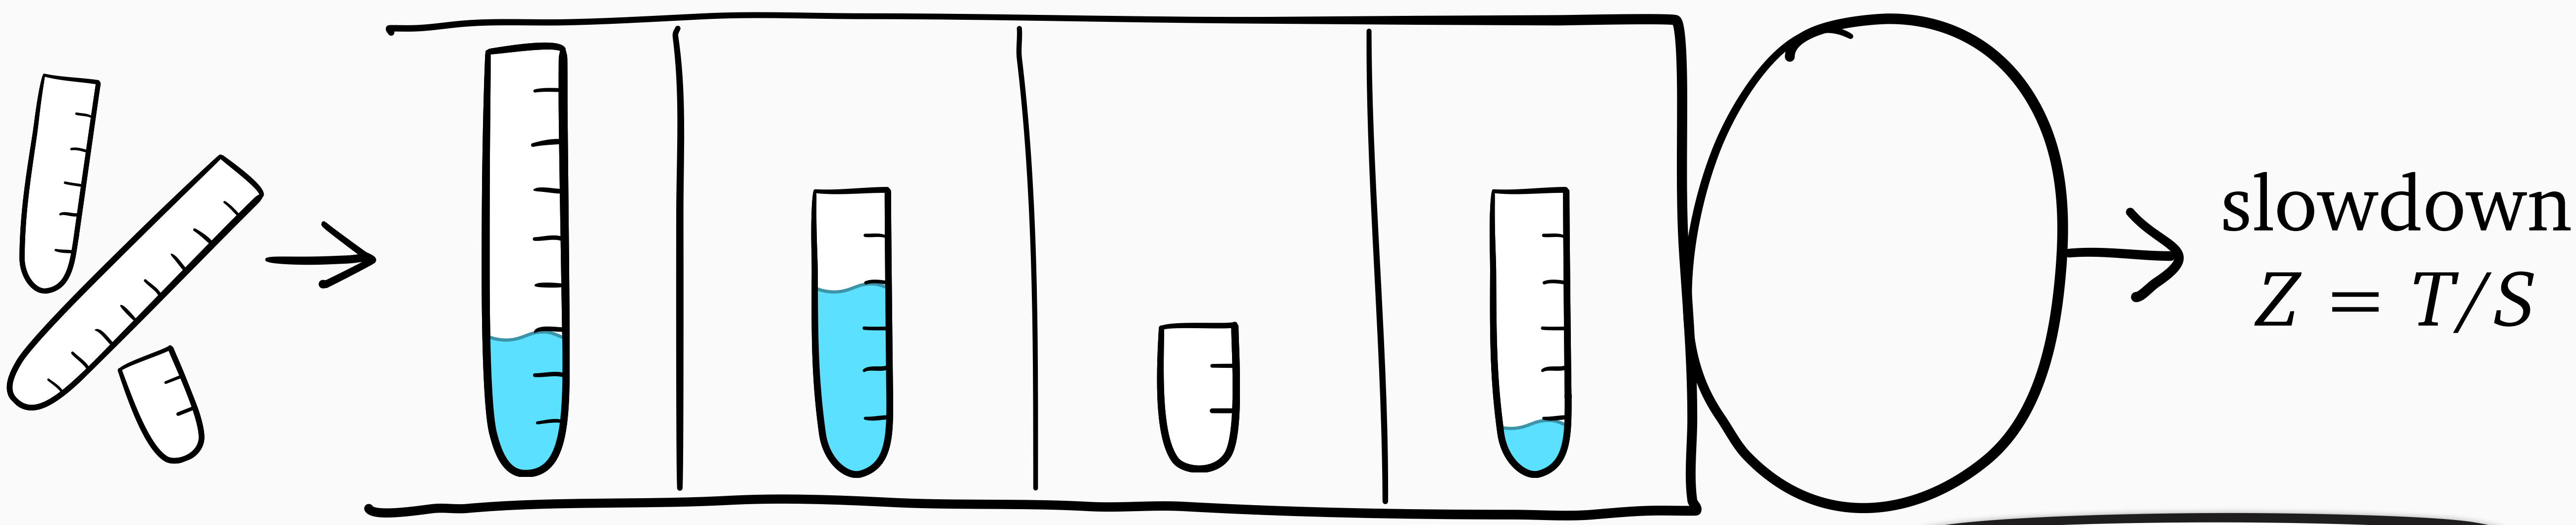

can finish job fast

Want to prioritize small *size* and small *remaining size*

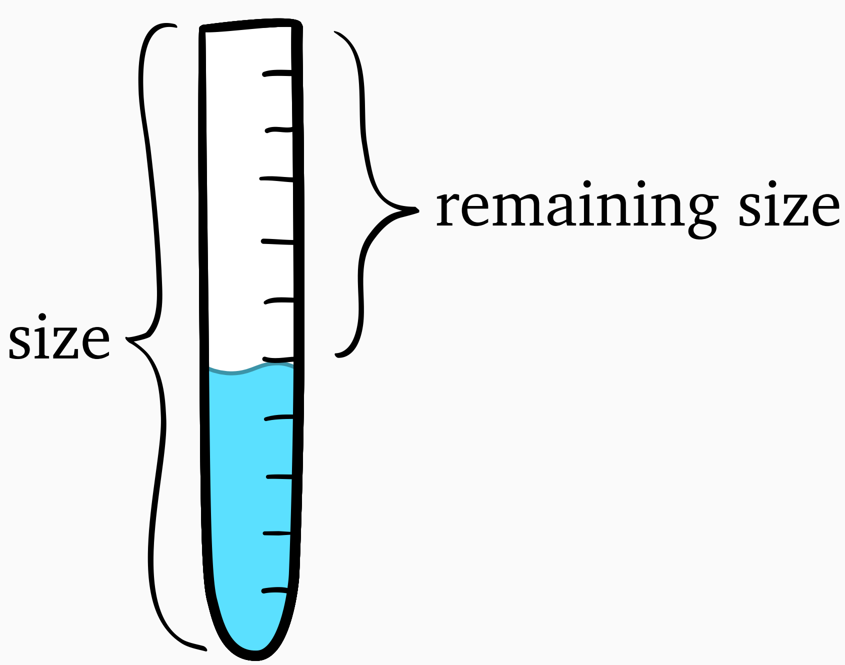

How to Schedule?

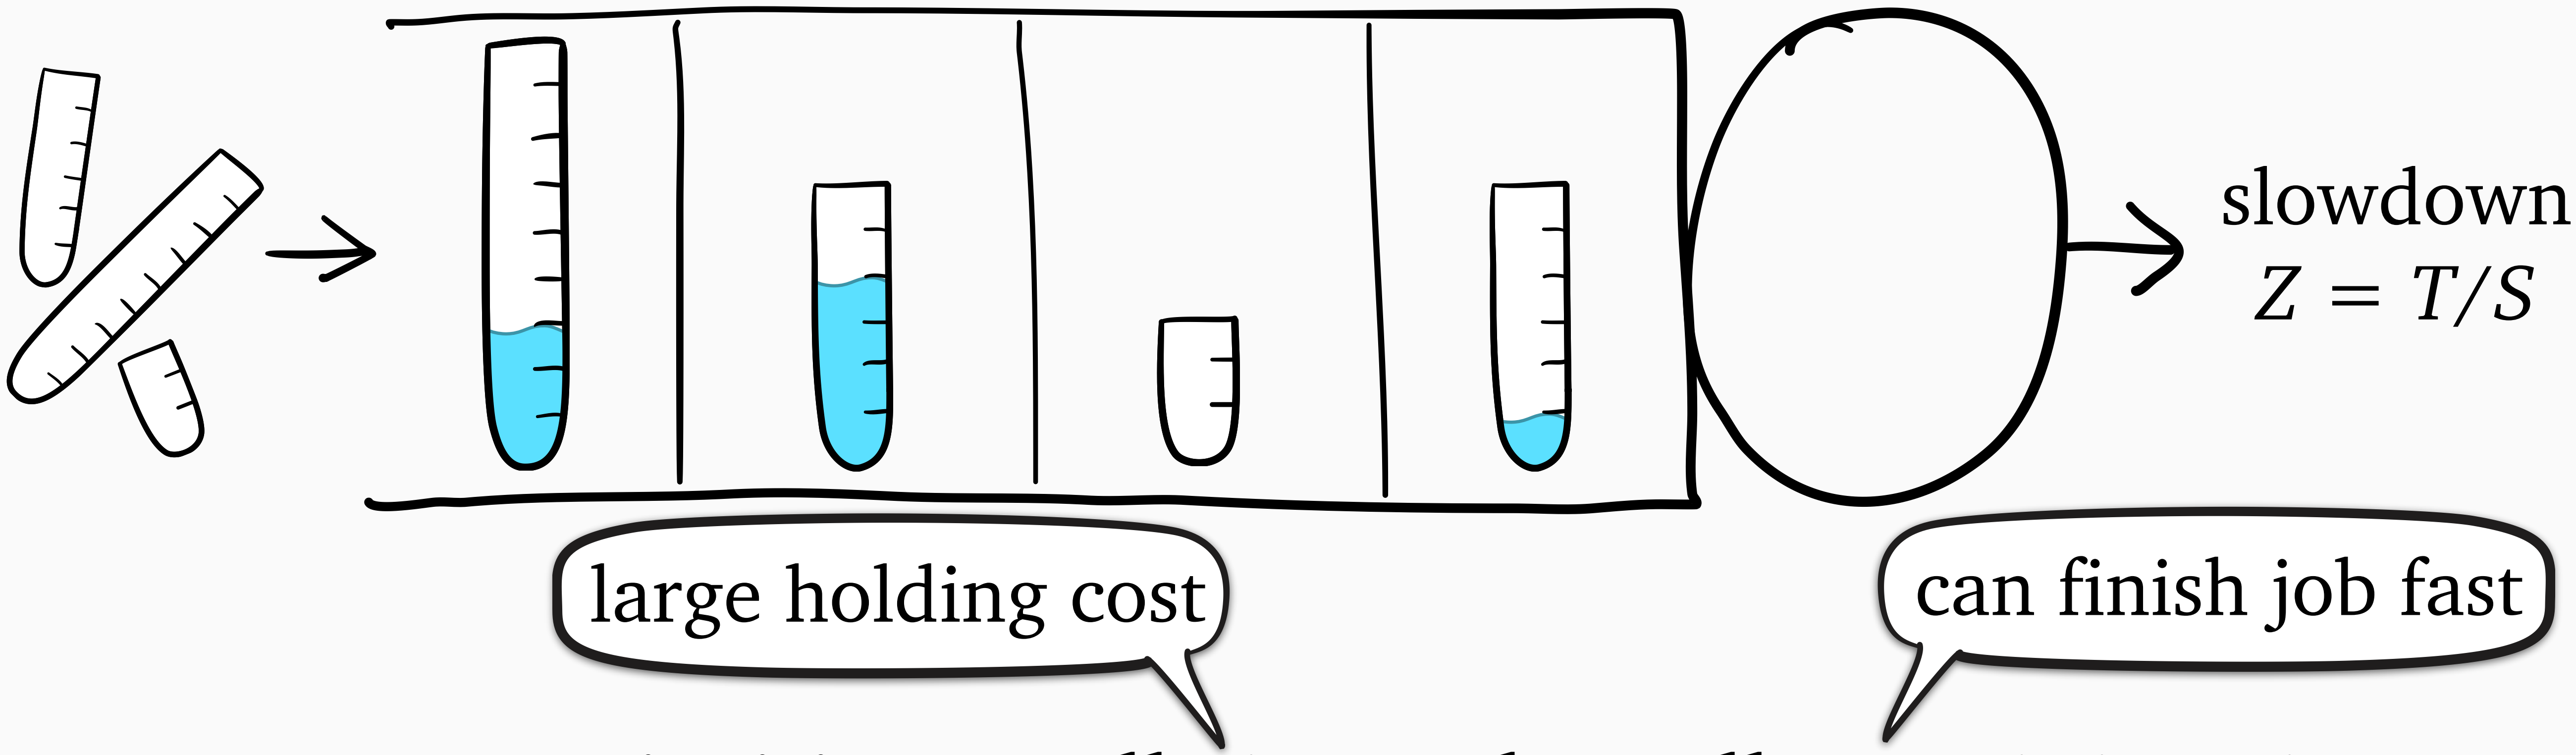

Want to prioritize small *size* and small *remaining size*

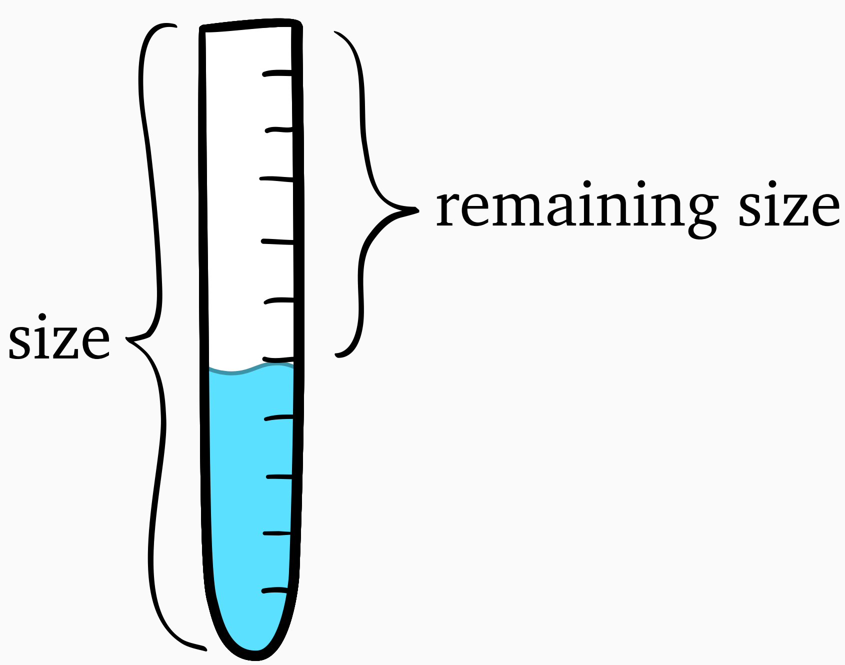

How to Schedule?

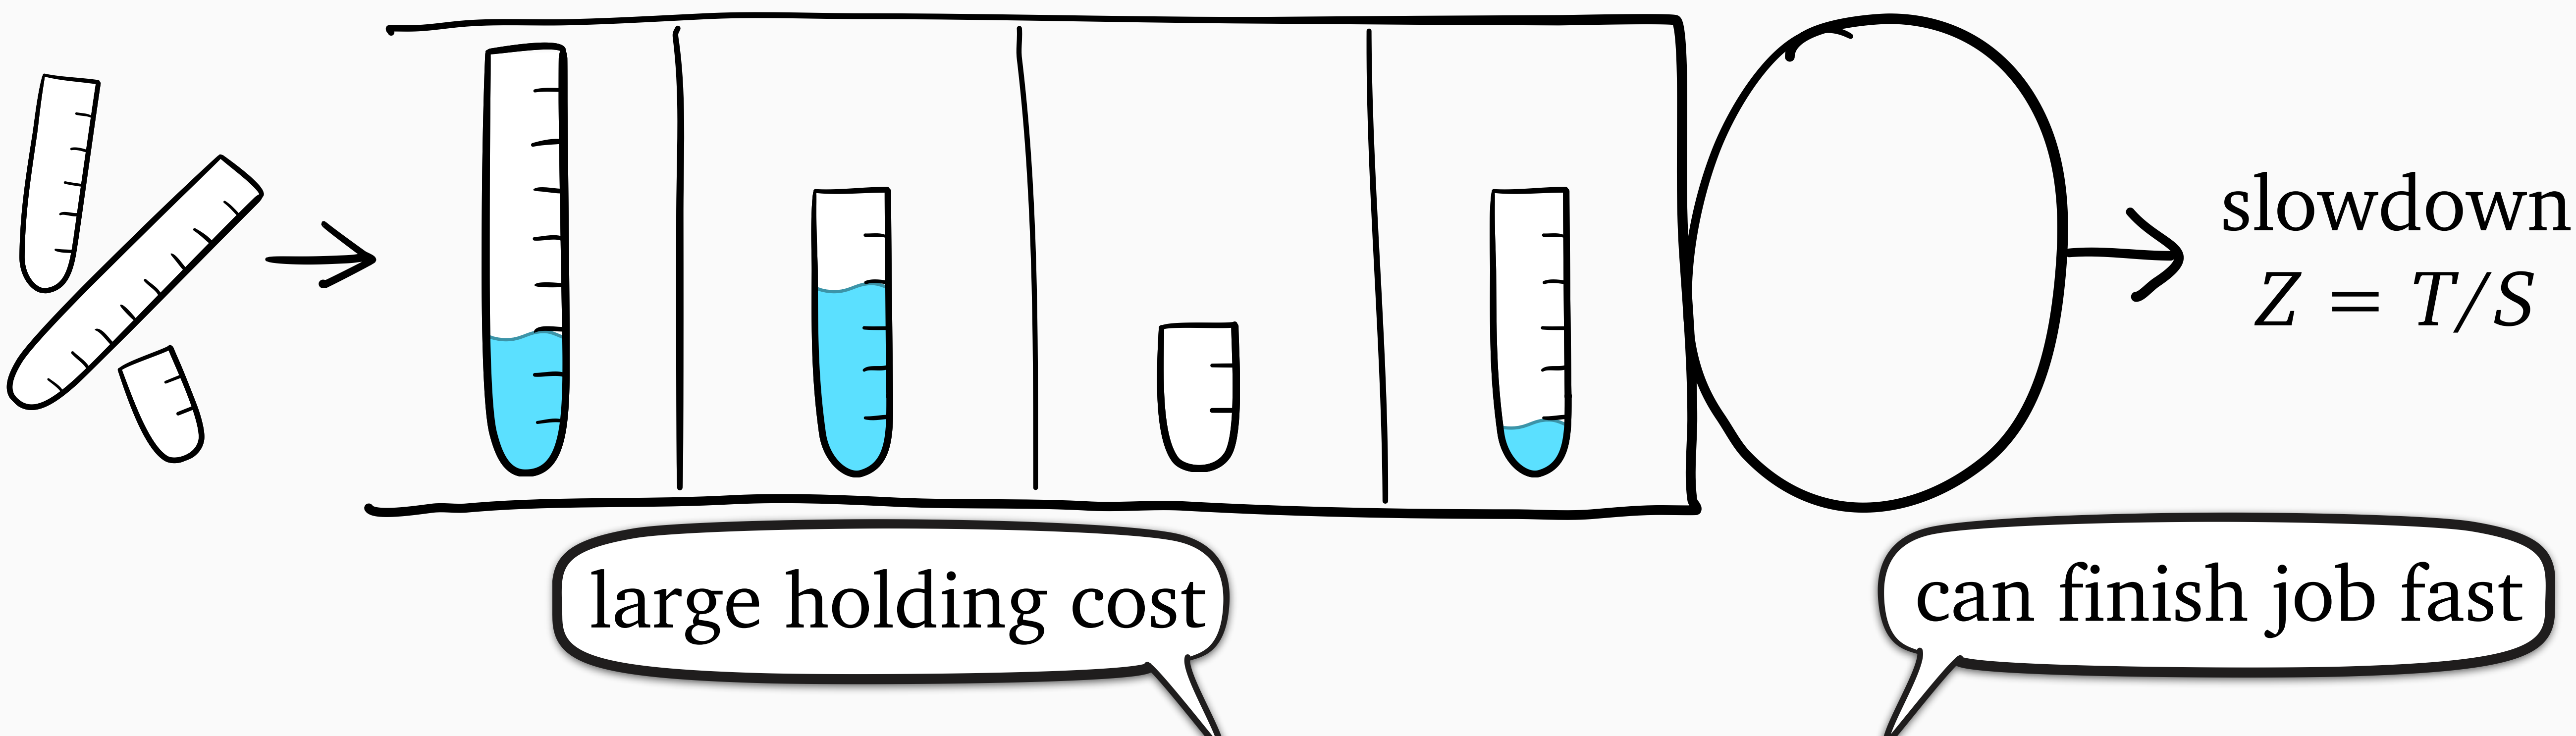

Want to prioritize small *size* and small *remaining size*

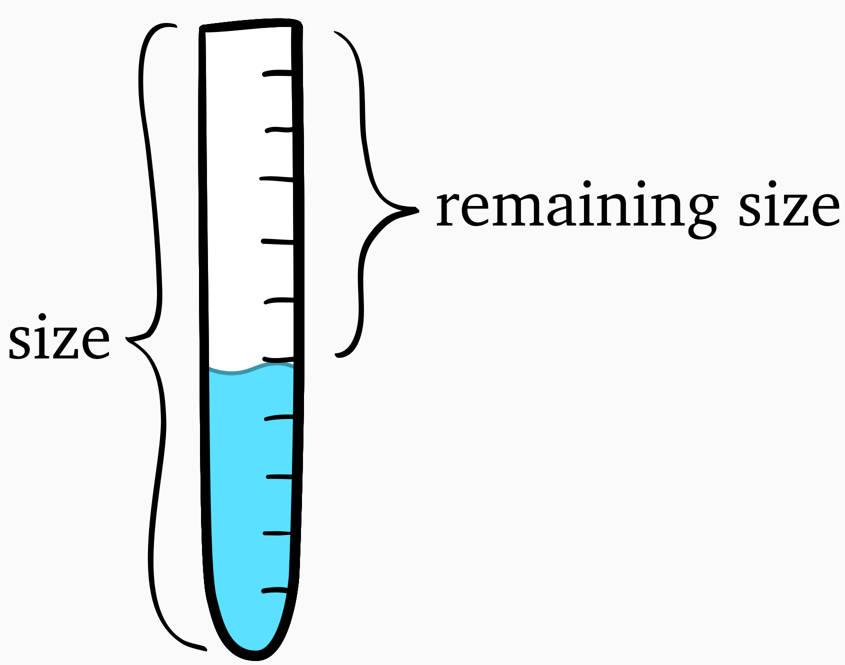

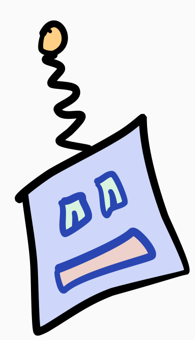

RS: always serve job of least **rank** = size · remaining size

How to Schedule?

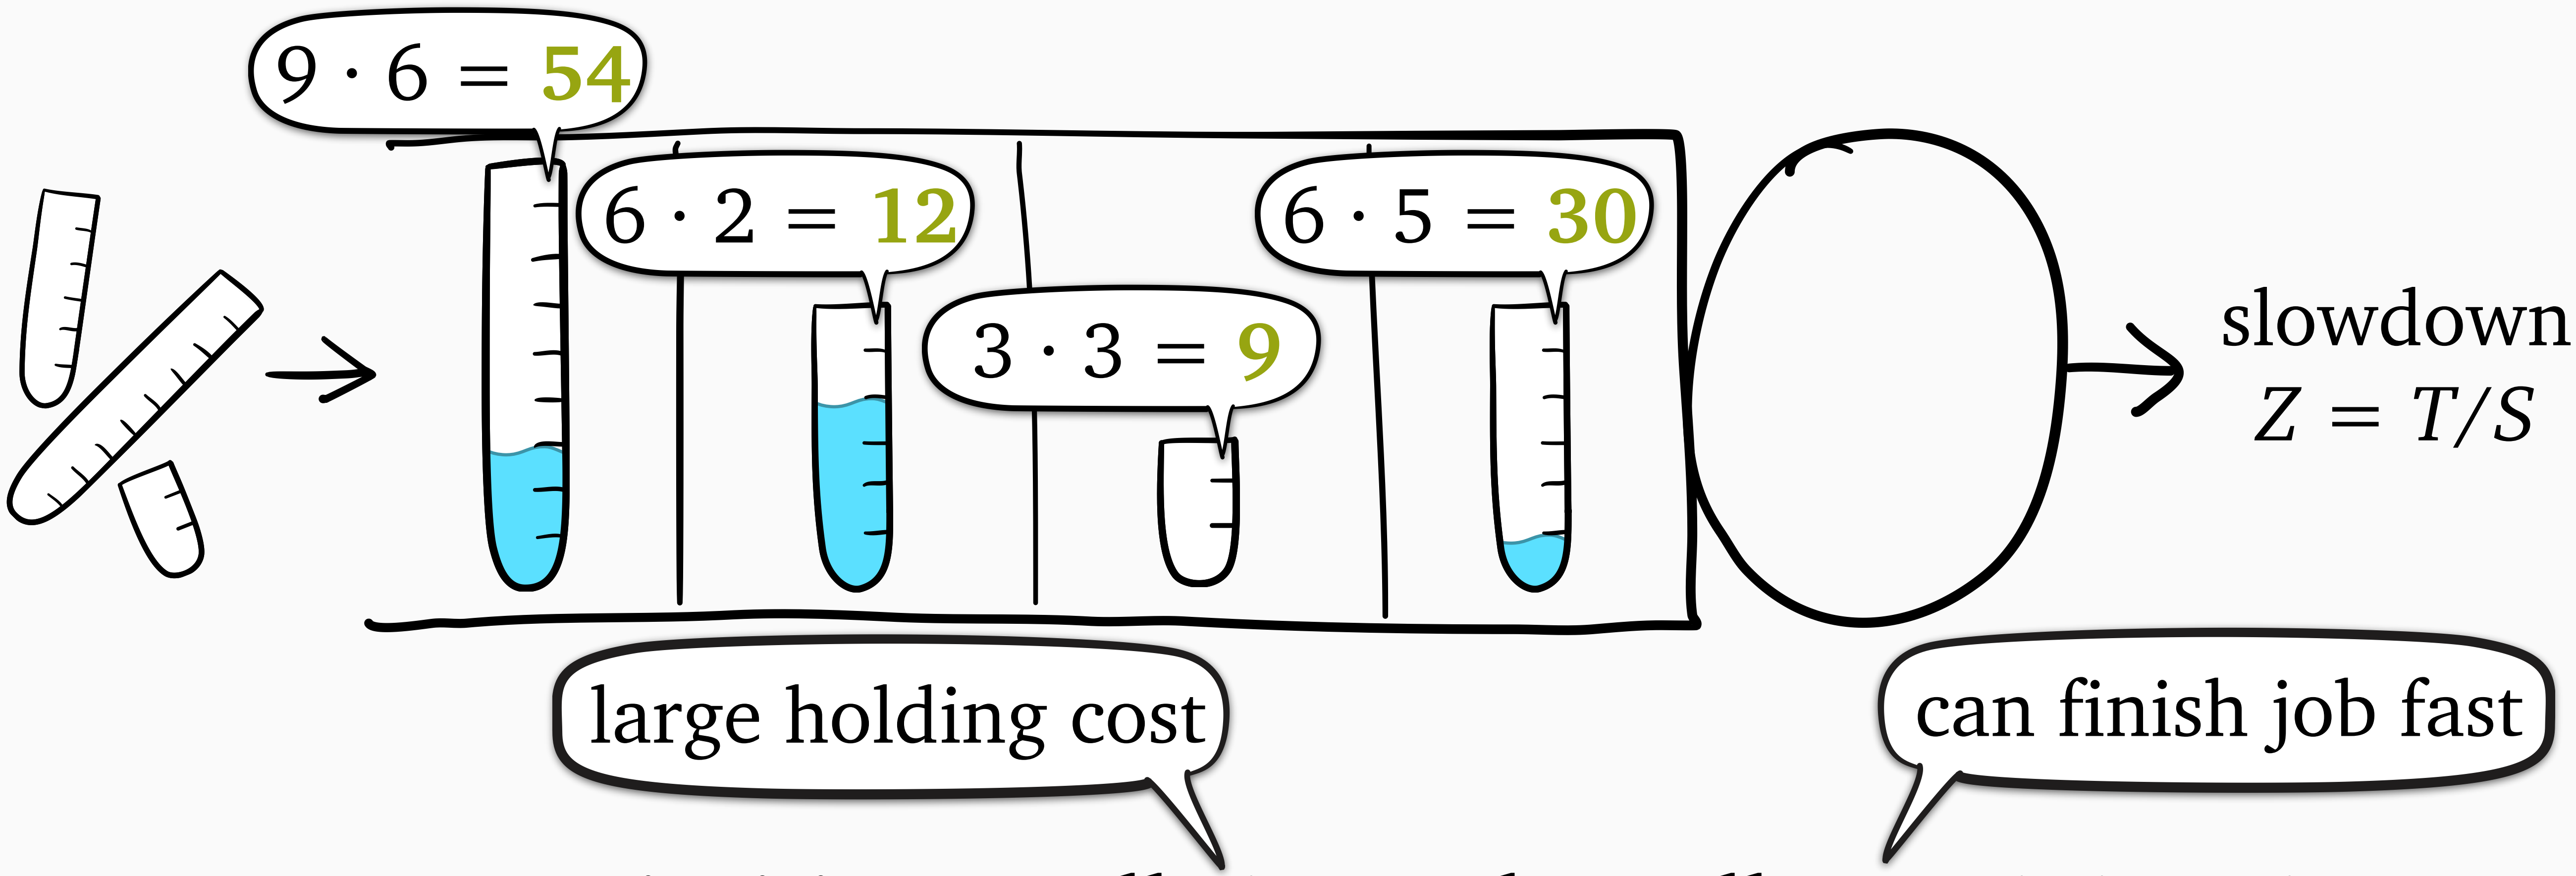

Want to prioritize small *size* and small *remaining size*

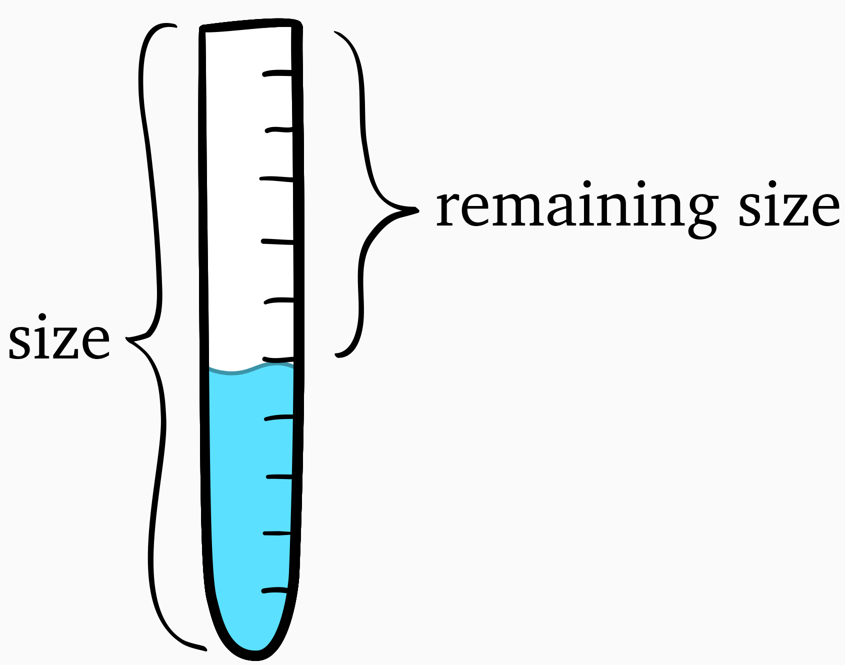

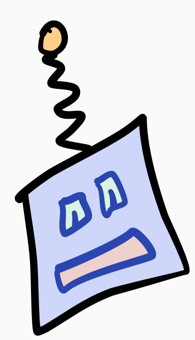

RS: always serve job of least
rank = size · remaining size

How to Schedule?

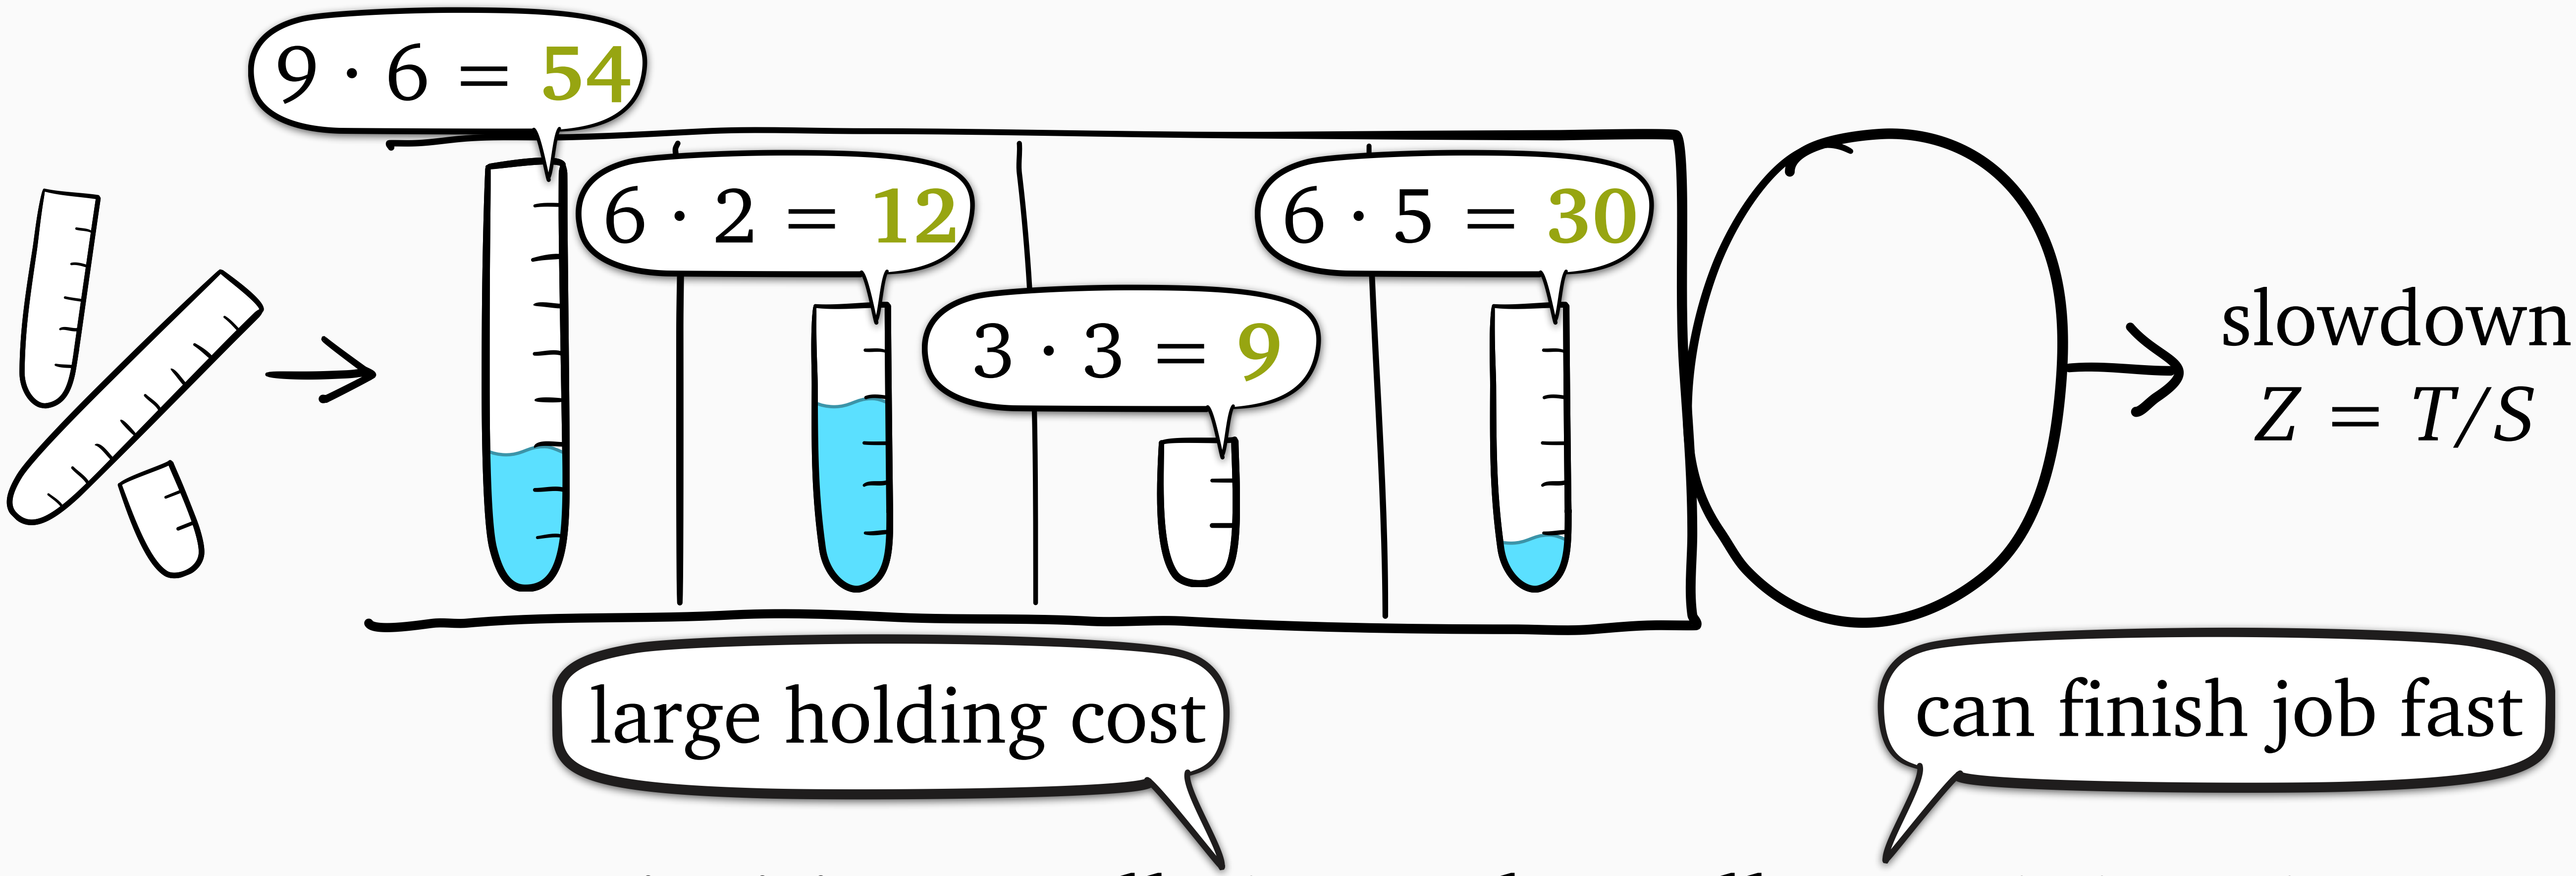

Want to prioritize small *size* and small *remaining size*

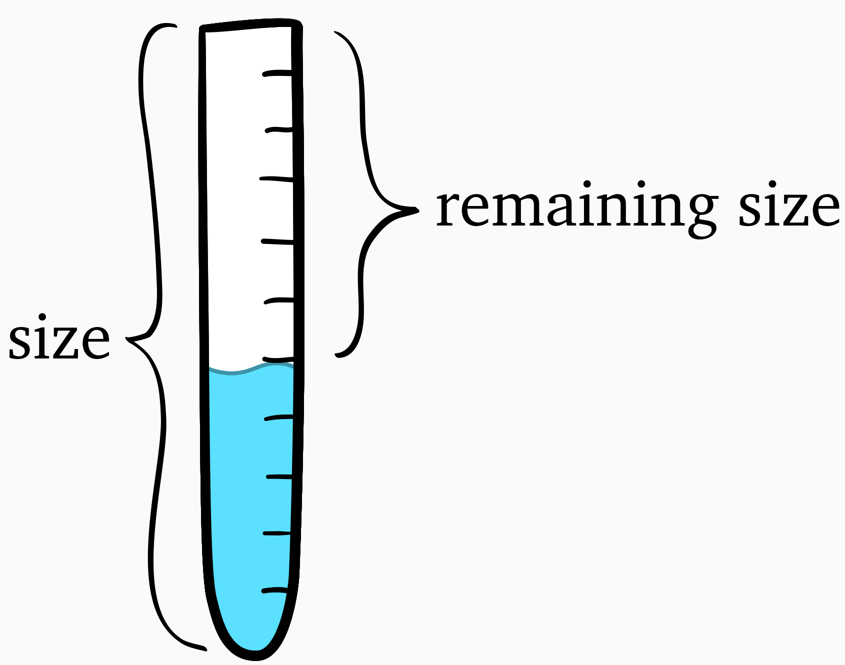

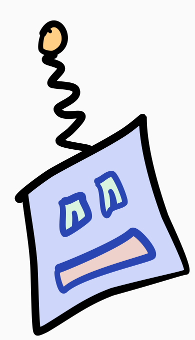

RS: always serve job of least
rank = size · remaining size

RS minimizes $\mathbf{E}[Z]$

First: background on *single-server* scheduling

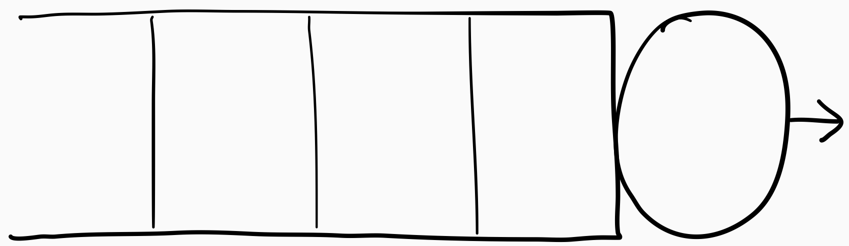

This talk: near-optimal *multiserver* scheduling

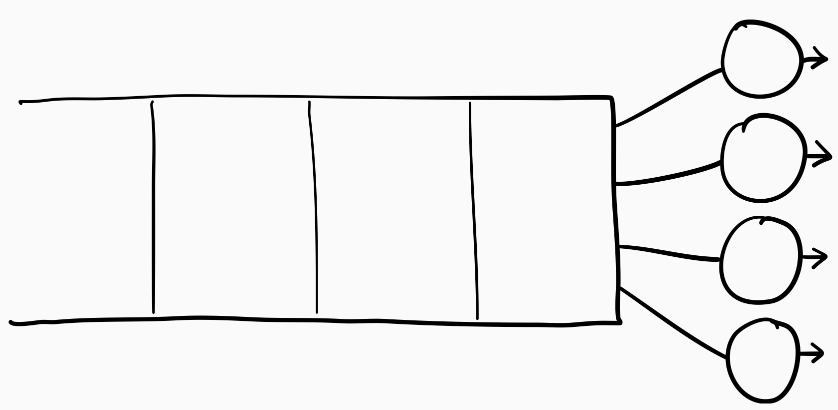

First: background on *single-server* scheduling

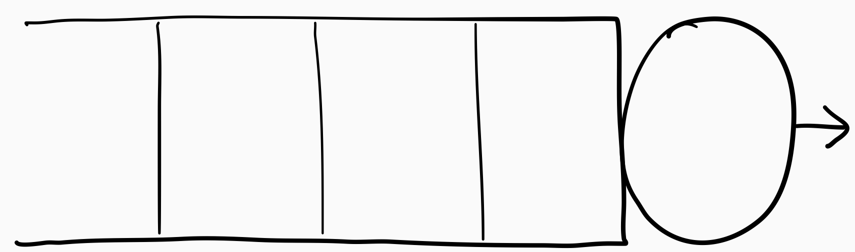

This talk: near-optimal *multiserver* scheduling

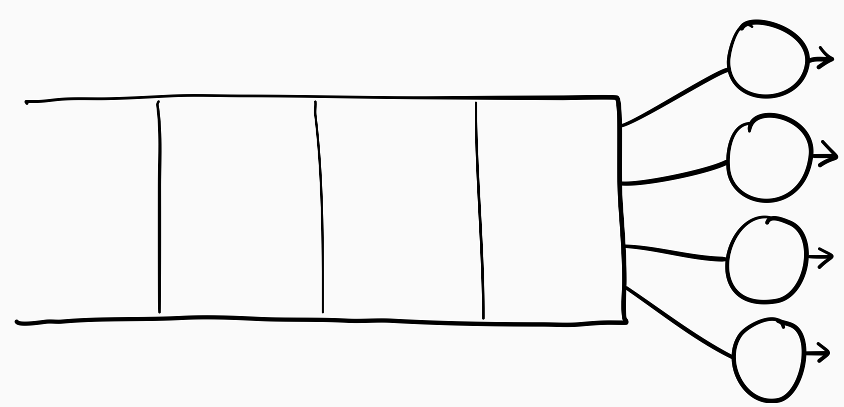

First: background on *single-server* scheduling

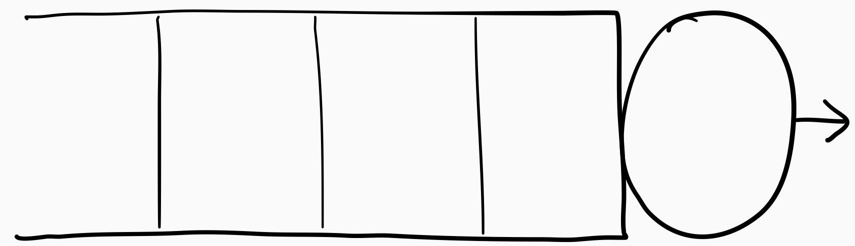

This talk: near-optimal *multiserver* scheduling

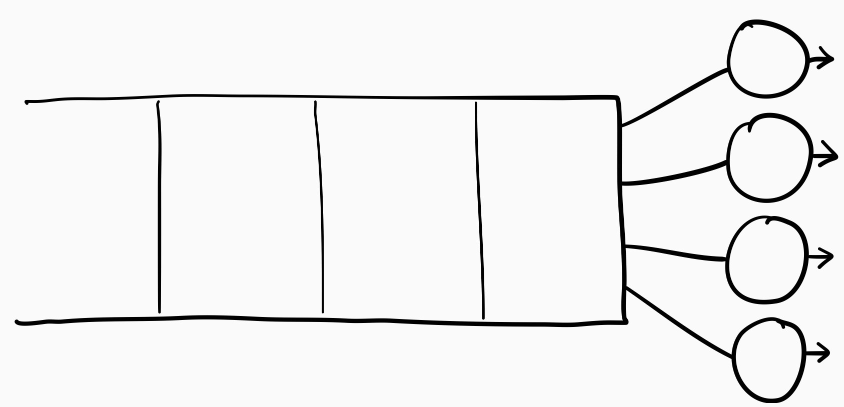

First: background on *single-server* scheduling

use **RS**!

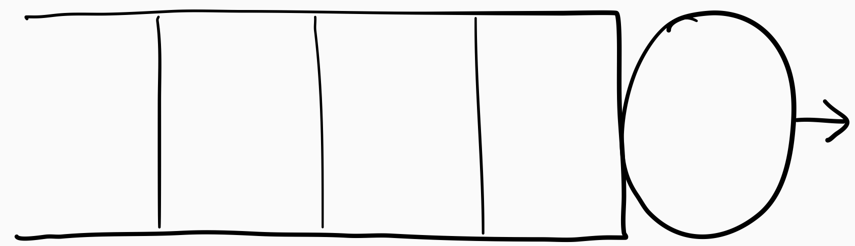

This talk: near-optimal *multiserver* scheduling

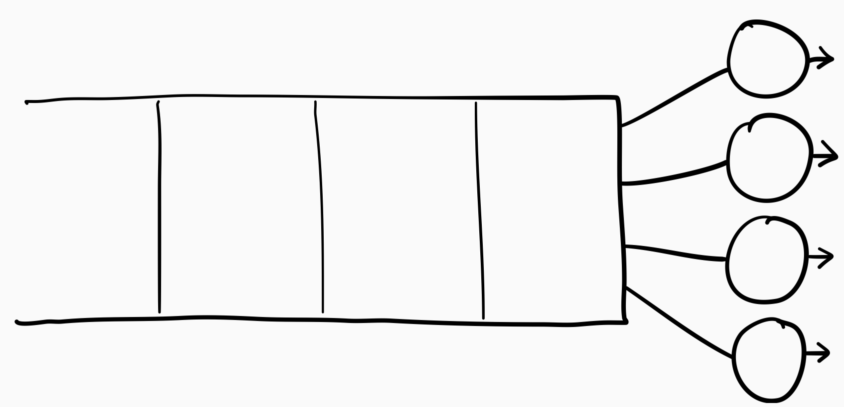

First: background on *single-server* scheduling

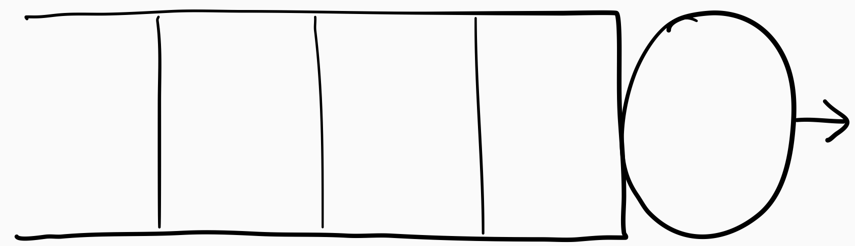

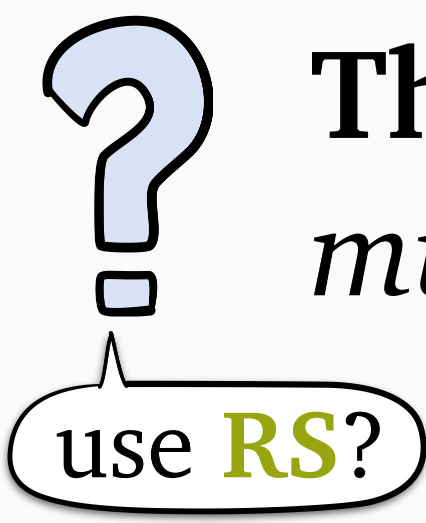

This talk: near-optimal *multiserver* scheduling

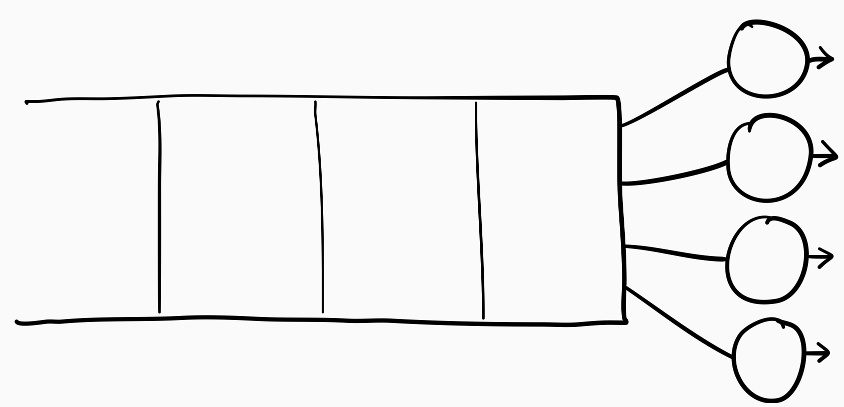

Multiserver Systems

Multiserver Systems

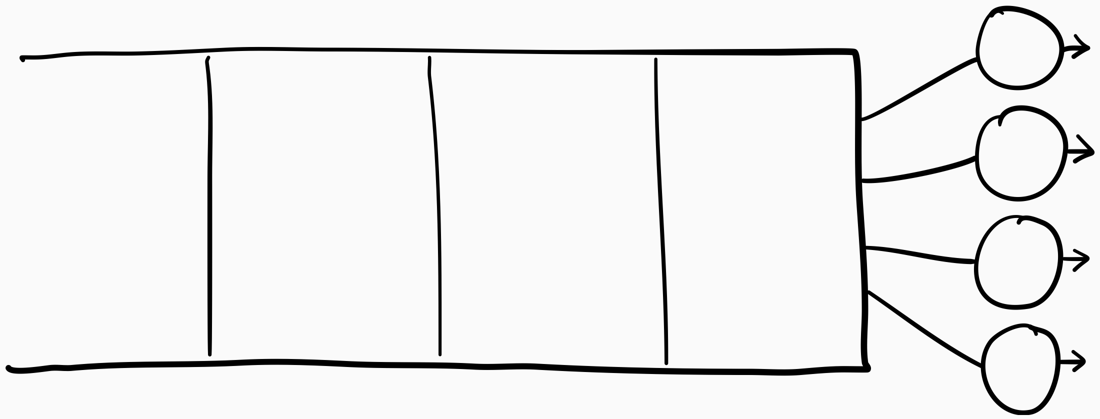

Central queue: M/G/k

Multiserver Systems

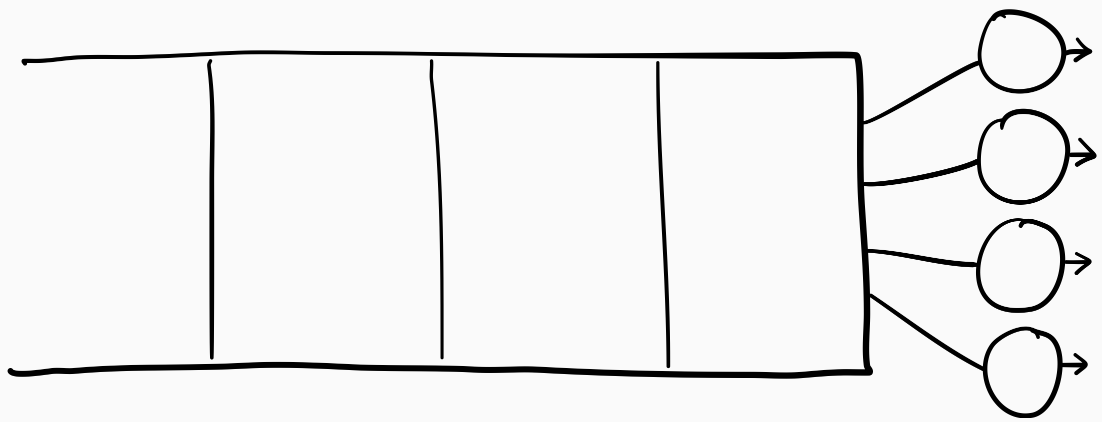

Central queue: M/G/k

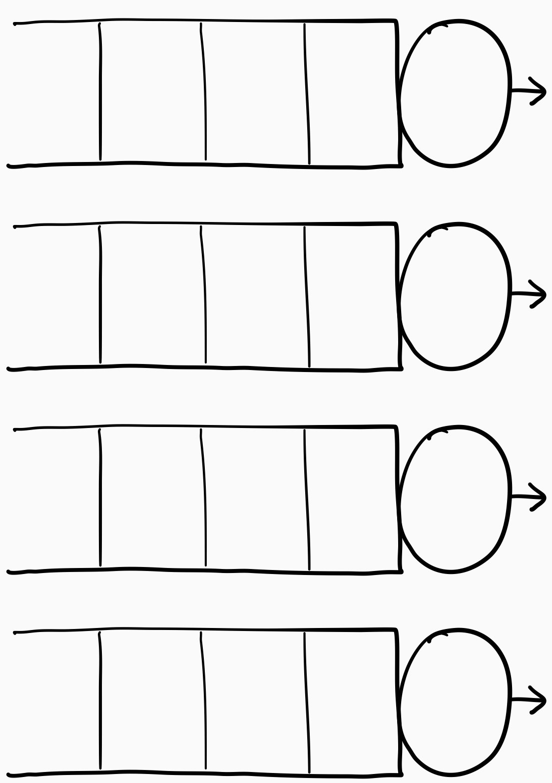

Load balancing: M/G/k/dispatch

Multiserver Systems

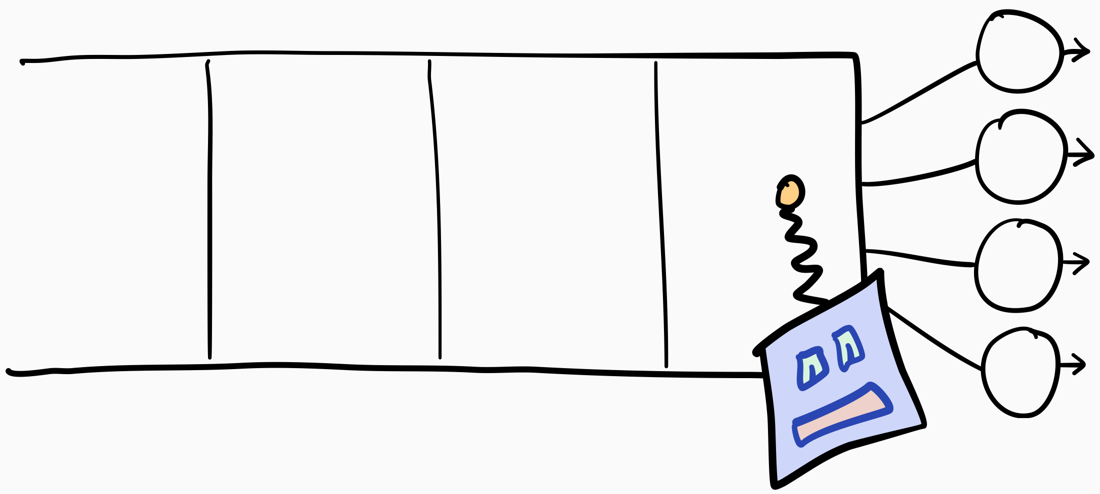

Central queue: M/G/k

- How to schedule?

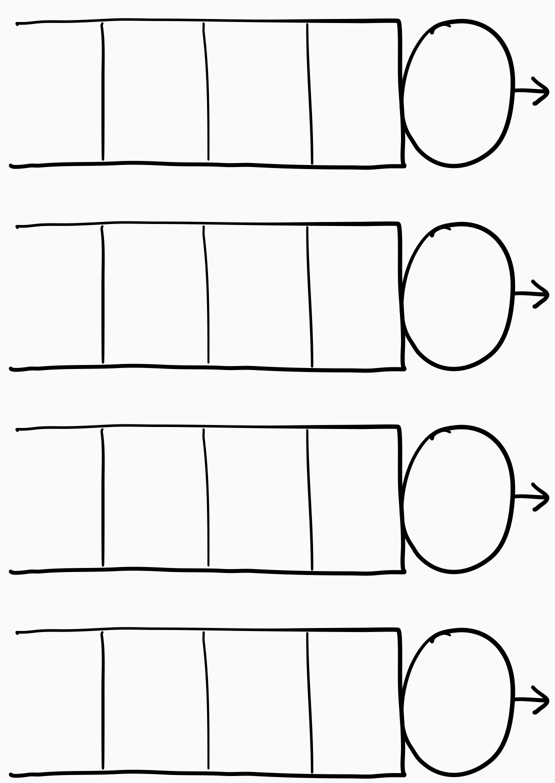

Load balancing: M/G/k/dispatch

Multiserver Systems

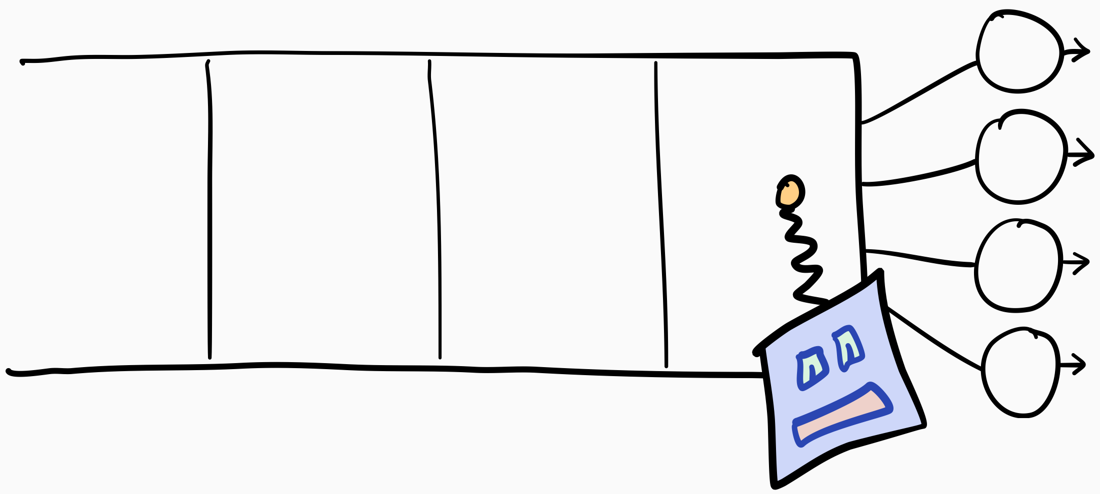

Central queue: M/G/k

- How to schedule?

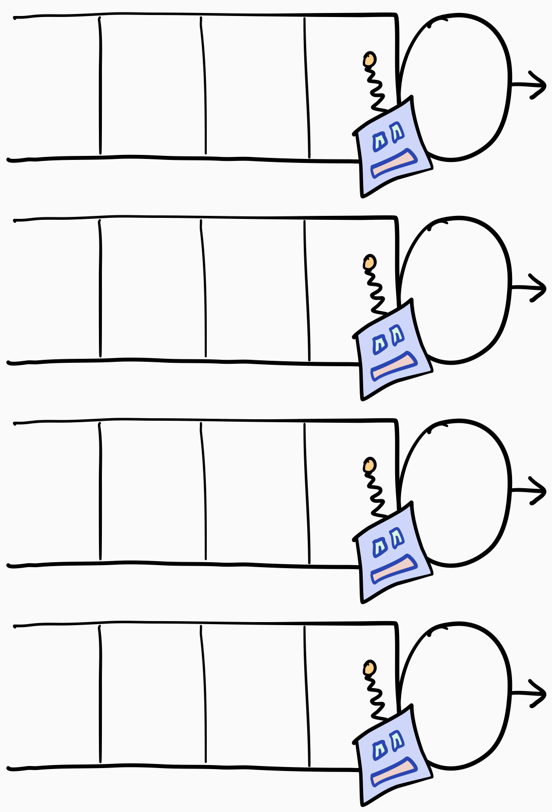

Load balancing: M/G/k/dispatch

- How to schedule?

Multiserver Systems

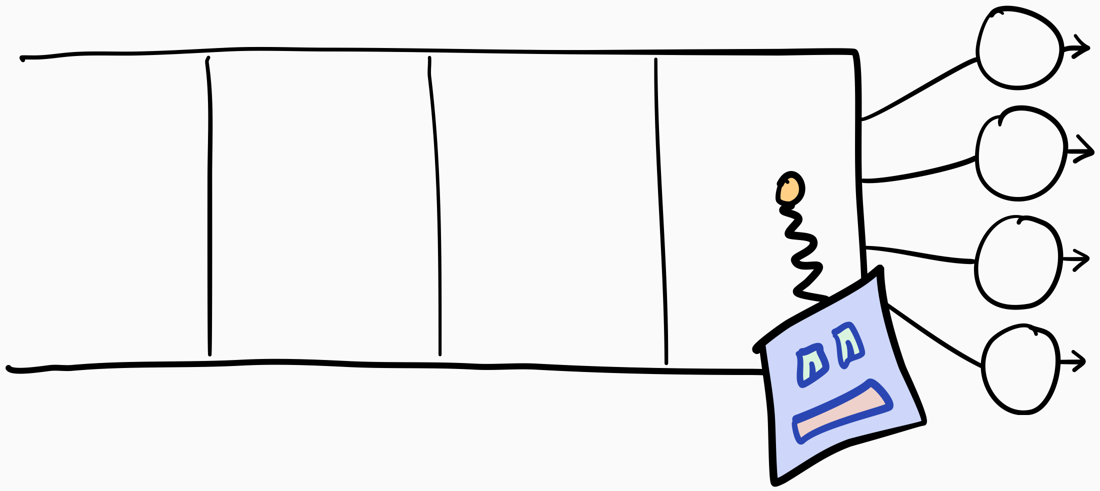

Central queue: M/G/k

- How to schedule?

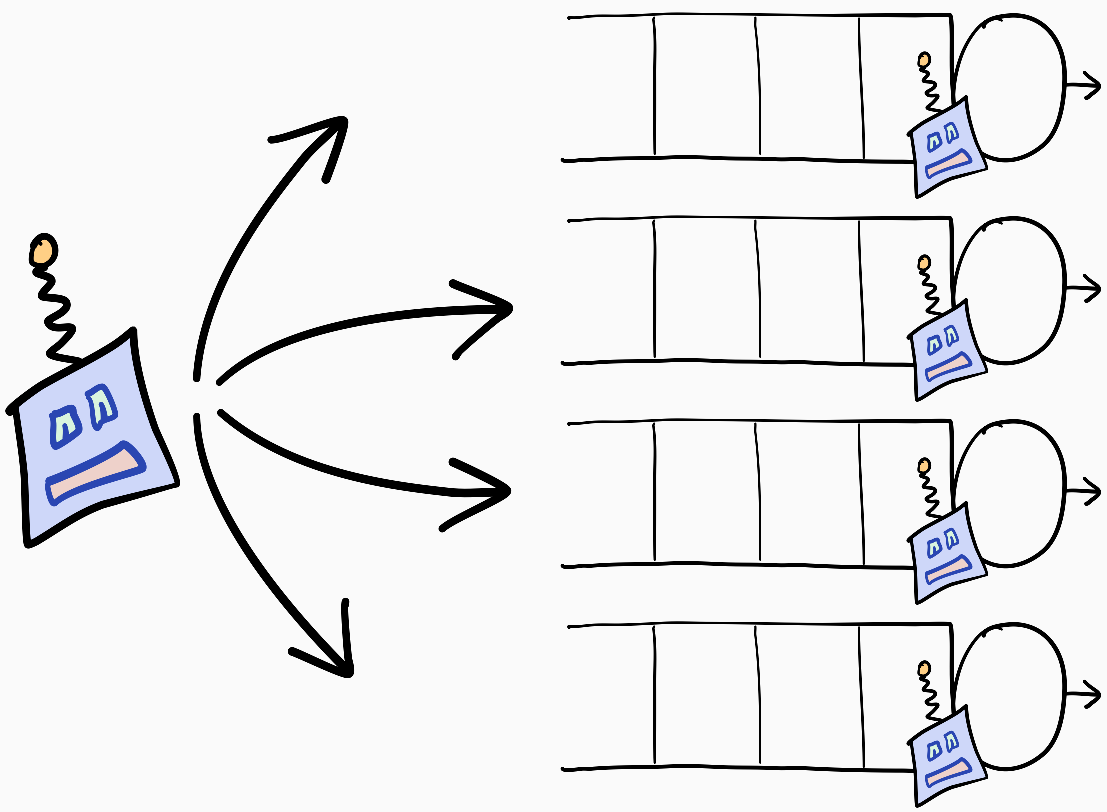

Load balancing: M/G/k/dispatch

- How to schedule?
- How to dispatch?

Comparing to M/G/1

k server of speed $1/k$

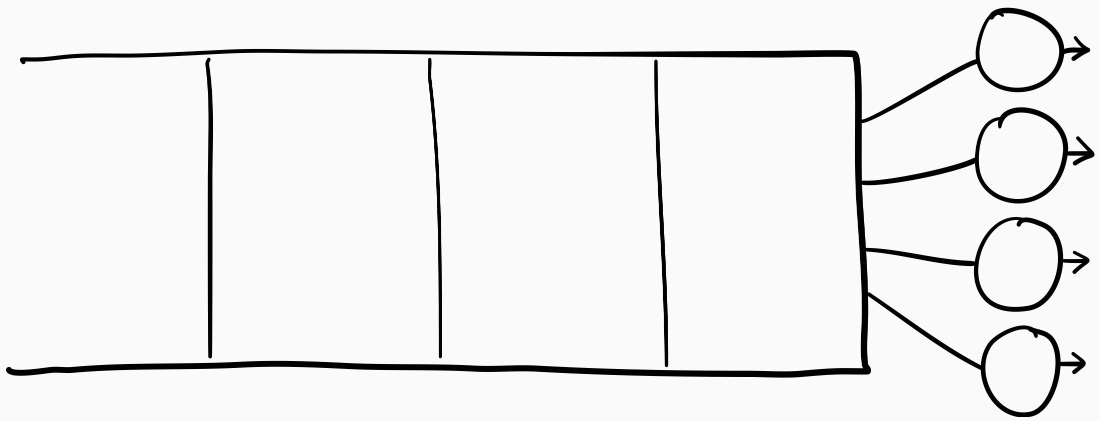

1 server of speed 1

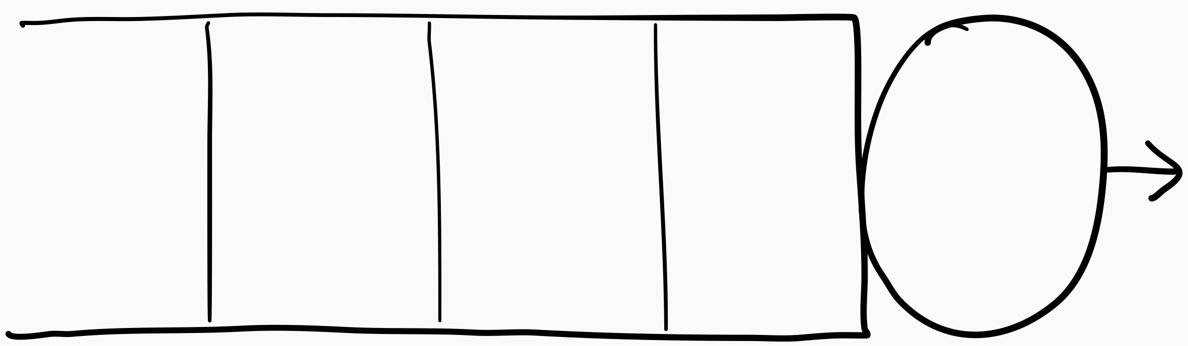

Comparing to M/G/1

k server of speed $1/k$

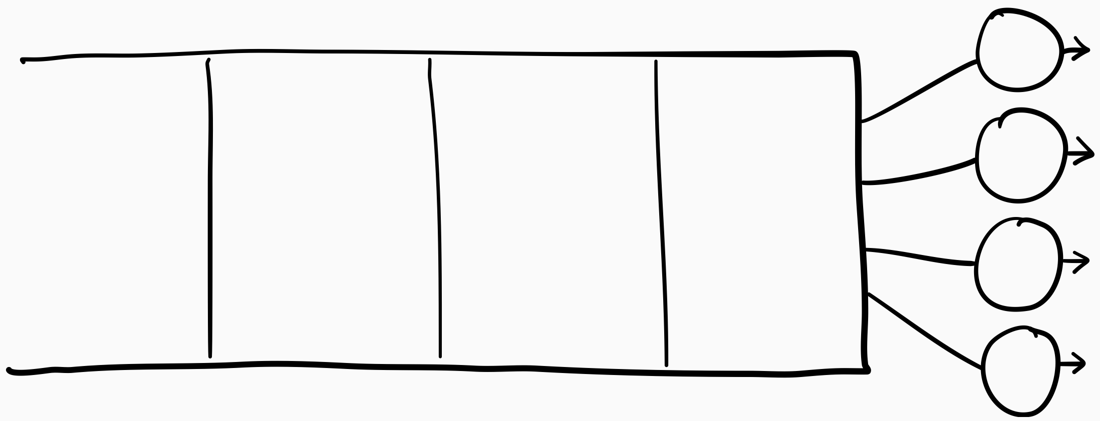

1 server of speed 1

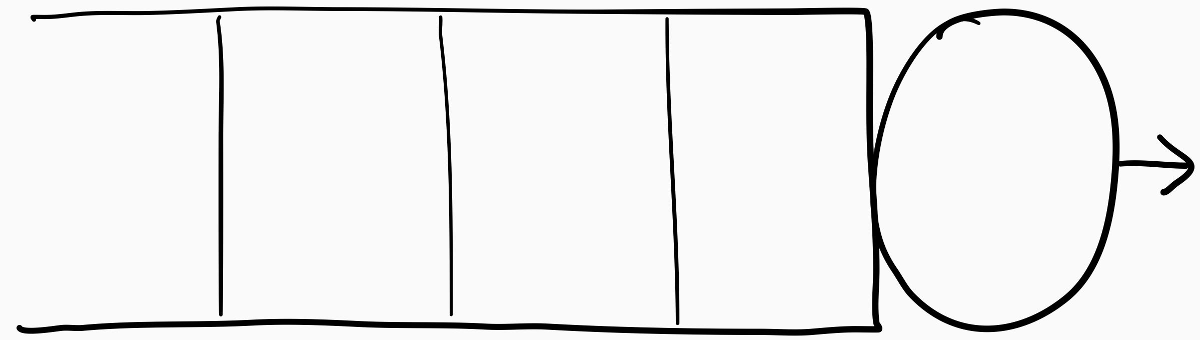

$$\mathbf{E}[Z_1^{\mathrm{Opt}}] \leq \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Comparing to M/G/1

k server of speed $1/k$

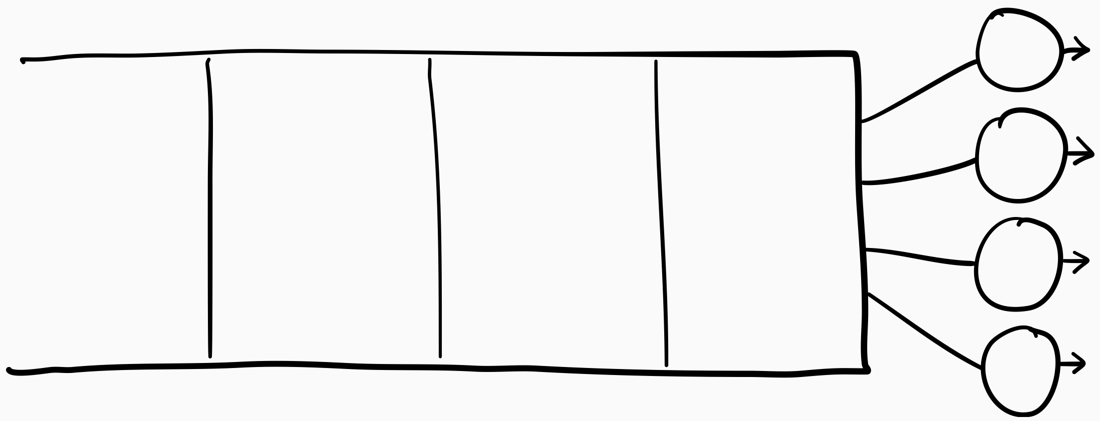

1 server of speed 1

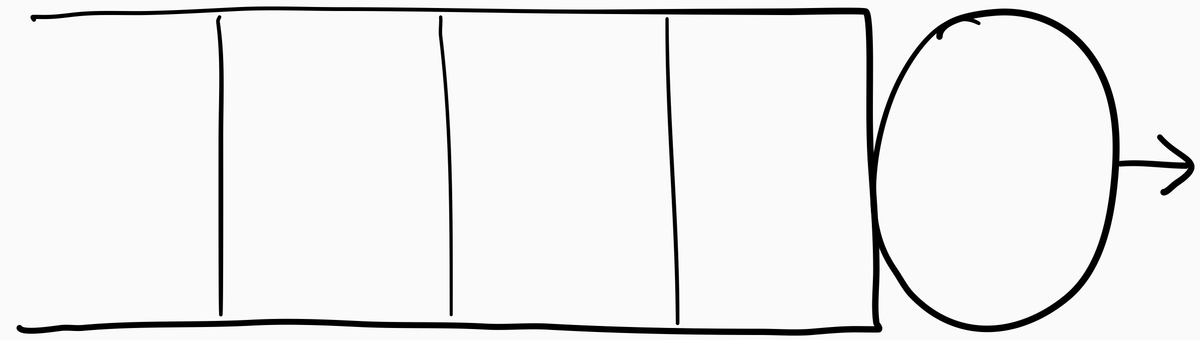

$$\mathbf{E}[Z_1^{\mathrm{RS}}] = \mathbf{E}[Z_1^{\mathrm{Opt}}] \le \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Comparing to M/G/1

k server of speed $1/k$

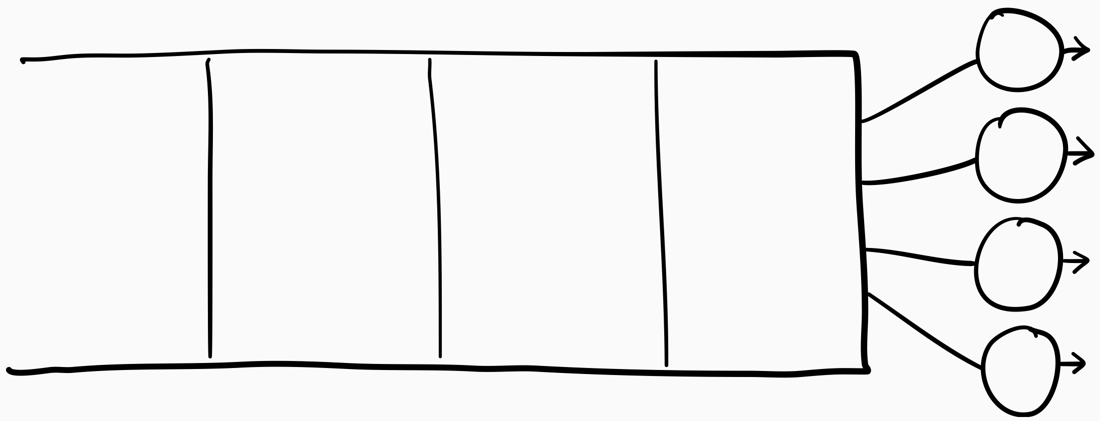

1 server of speed 1

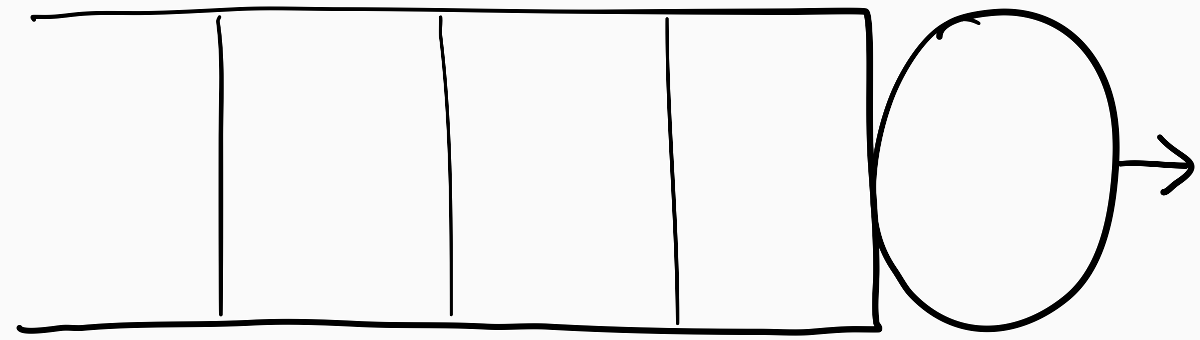

$$\mathbf{E}[Z_1^{\mathbf{RS}}] = \mathbf{E}[Z_1^{\mathrm{Opt}}] \leq \mathbf{E}[Z_k^{\mathrm{Opt}}] \leq \mathbf{E}[Z_k^{\mathbf{RS}}]$$

Comparing to M/G/1

k server of speed $1/k$

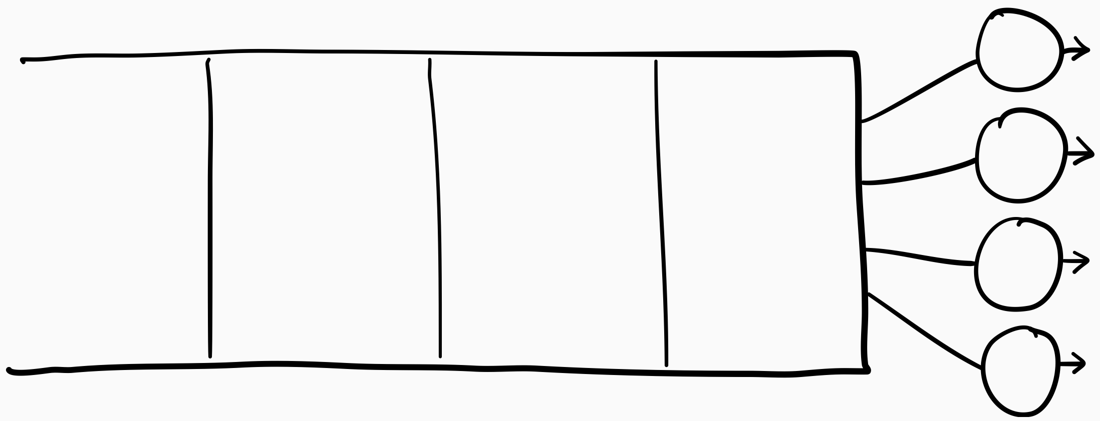

1 server of speed 1

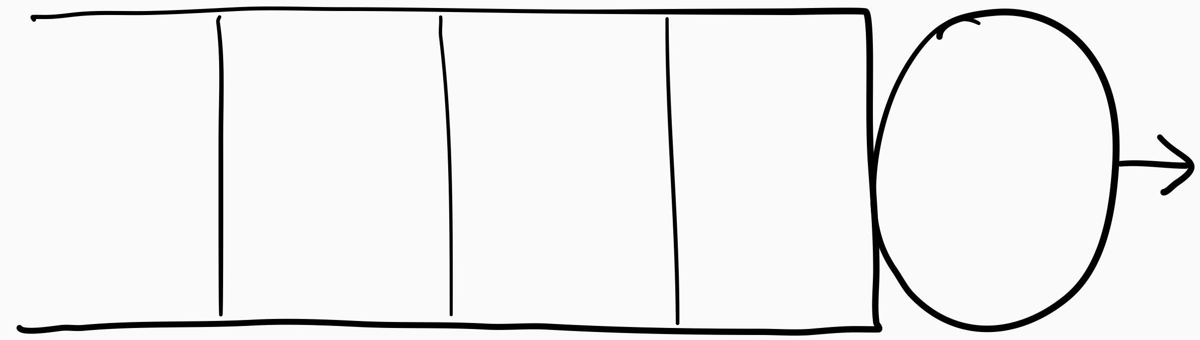

$$\mathbf{E}[Z_1^{\mathbf{RS}}] = \mathbf{E}[Z_1^{\mathrm{Opt}}] \le \mathbf{E}[Z_k^{\mathrm{Opt}}] \le \mathbf{E}[Z_k^{\mathbf{RS}}]$$

Goal: $\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + \text{“small”}$

Comparing to M/G/1

k server of speed $1/k$

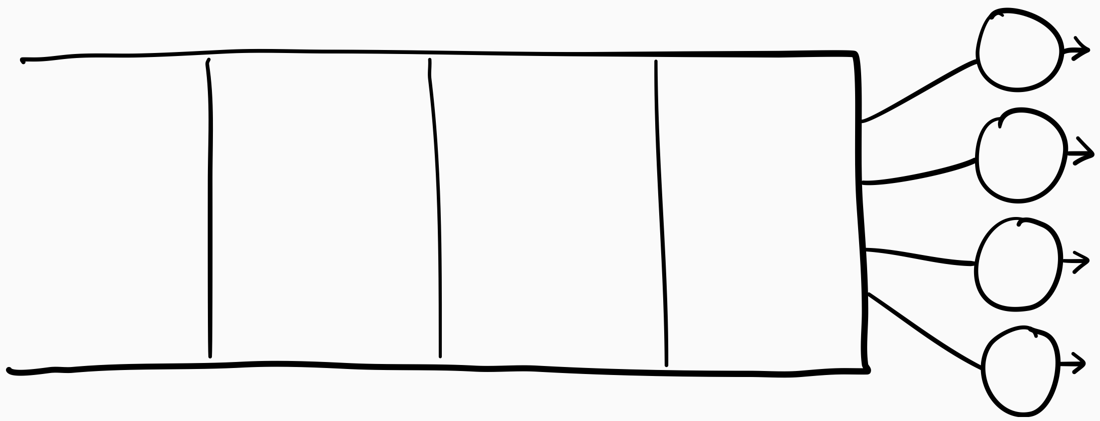

1 server of speed 1

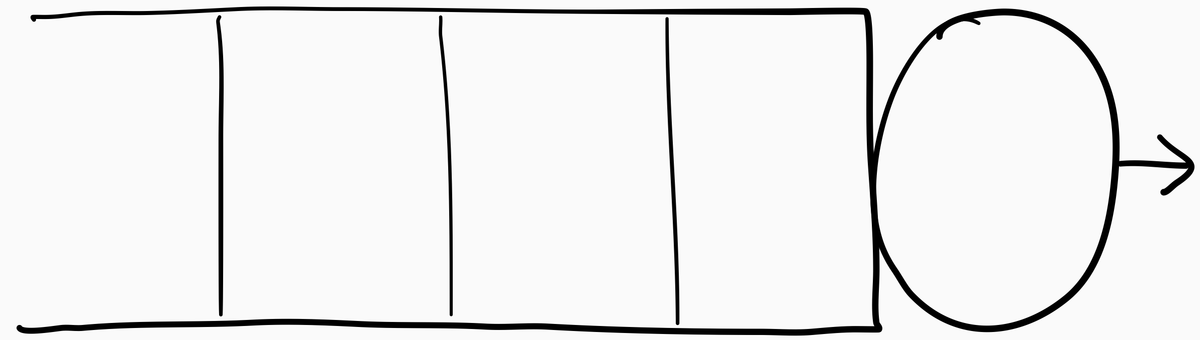

$$\mathbf{E}[Z_1^{\mathrm{RS}}] = \mathbf{E}[Z_1^{\mathrm{Opt}}] \le \mathbf{E}[Z_k^{\mathrm{Opt}}] \le \mathbf{E}[Z_k^{\mathrm{RS}}]$$

Goal: $\mathbf{E}[Z_k^{\mathrm{RS}}] \le \mathbf{E}[Z_1^{\mathrm{RS}}] + \text{“small”}$

Constant-factor approx.

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le c \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Comparing to M/G/1

k server of speed $1/k$

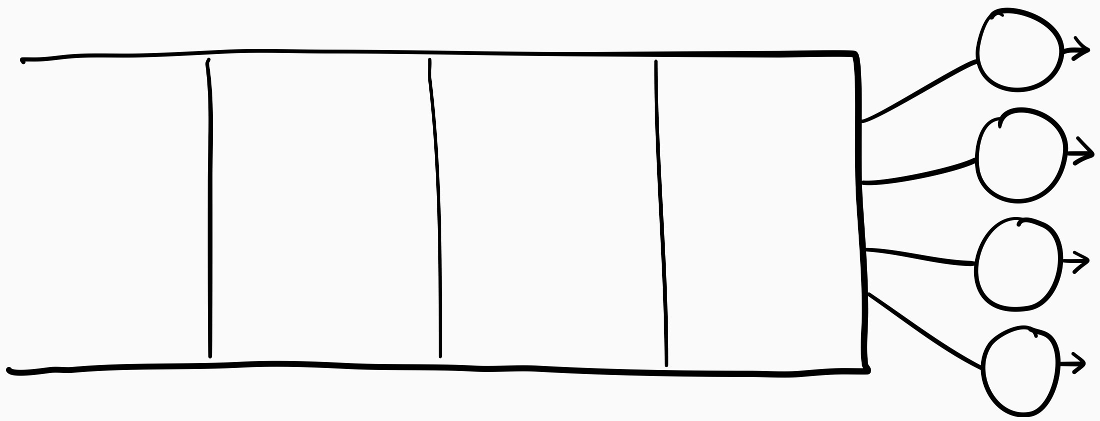

1 server of speed 1

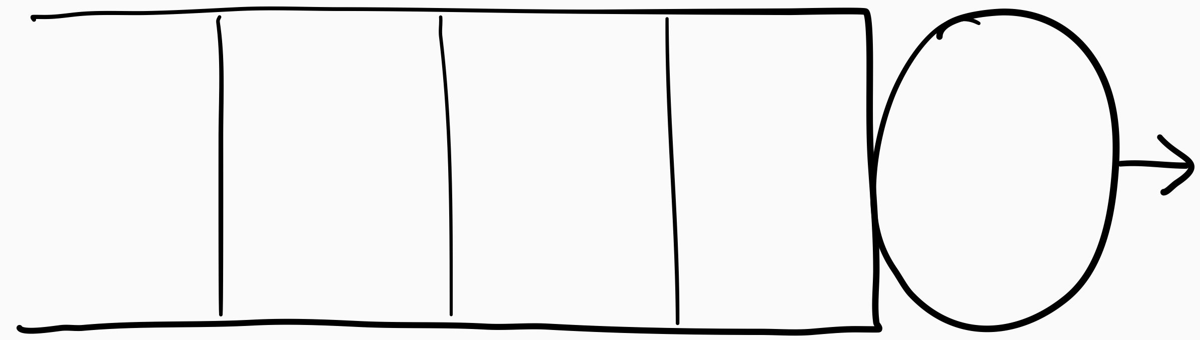

$$\mathbf{E}[Z_1^{\mathbf{RS}}] = \mathbf{E}[Z_1^{\mathrm{Opt}}] \leq \mathbf{E}[Z_k^{\mathrm{Opt}}] \leq \mathbf{E}[Z_k^{\mathbf{RS}}]$$

Goal: $\mathbf{E}[Z_k^{\mathbf{RS}}] \leq \mathbf{E}[Z_1^{\mathbf{RS}}] + \text{“small”}$

Constant-factor approx.

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \leq c \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Heavy-traffic optimality

$$\lim_{\rho \to 1} \frac{\mathbf{E}[Z_k^{\mathbf{RS}}]}{\mathbf{E}[Z_k^{\mathrm{Opt}}]} = 1$$

Main Results

Goal: $\mathbf{E}[Z_k^{\mathbf{RS}}] \leq \mathbf{E}[Z_1^{\mathbf{RS}}] + \text{“small”}$

Main Results

Goal: $\mathbf{E}[Z_k^{\mathbf{RS}}] \leq \mathbf{E}[Z_1^{\mathbf{RS}}] + \text{“small”}$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \leq \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

Main Results

Goal: $\mathbf{E}[Z_k^{\textbf{RS}}] \le \mathbf{E}[Z_1^{\textbf{RS}}] + \text{“small”}$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\textbf{RS}}] \le \mathbf{E}[Z_1^{\textbf{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\textbf{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\text{Opt}}]$$

Main Results

Goal: $\mathbf{E}[Z_k^{\mathrm{RS}}] \le \mathbf{E}[Z_1^{\mathrm{RS}}] + \text{“small”}$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le \mathbf{E}[Z_1^{\mathrm{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In M/G/k/dispatch with **guardrails** dispatching,

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le \frac{11\varepsilon}{8}\mathbf{E}[Z_1^{\mathrm{RS}}] + \frac{40}{\varepsilon}k \qquad \left(0 < \varepsilon \le \tfrac{3}{8}\right)$$

Main Results

Goal: $\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + \text{“small”}$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In M/G/k/dispatch with **guardrails** dispatching,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le \frac{11\varepsilon}{8}\mathbf{E}[Z_1^{\mathbf{RS}}] + \frac{40}{\varepsilon}k \qquad \left(0 < \varepsilon \le \tfrac{3}{8}\right)$$

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le 109 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Main Results

Goal: $\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + \text{“small”}$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In M/G/k/dispatch with **guardrails** dispatching,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le \frac{11\varepsilon}{8}\mathbf{E}[Z_1^{\mathbf{RS}}] + \frac{40}{\varepsilon}k \qquad \left(0 < \varepsilon \le \tfrac{3}{8}\right)$$

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le 109 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In both cases, $\displaystyle\lim_{\rho \to 1} \frac{\mathbf{E}[Z_k^{\mathbf{RS}}]}{\mathbf{E}[Z_k^{\mathrm{Opt}}]} = 1$

Main Results

Goal: $\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + \text{“small”}$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In M/G/k/dispatch with **guardrails** dispatching,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le \frac{11\varepsilon}{8}\mathbf{E}[Z_1^{\mathbf{RS}}] + \frac{40}{\varepsilon}k \qquad \left(0 < \varepsilon \le \tfrac{3}{8}\right)$$

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le 109 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In both cases, $\lim_{\rho \to 1} \frac{\mathbf{E}[Z_k^{\mathbf{RS}}]}{\mathbf{E}[Z_k^{\mathrm{Opt}}]} = 1$

if $\mathbf{E}[S^3] < \infty$

Main Results

Goal: $\mathbf{E}[Z_k^{\mathrm{RS}}] \le \mathbf{E}[Z_1^{\mathrm{RS}}] + \text{“small”}$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le \mathbf{E}[Z_1^{\mathrm{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In M/G/k/dispatch with **guardrails** dispatching,

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le \frac{11\varepsilon}{8}\mathbf{E}[Z_1^{\mathrm{RS}}] + \frac{40}{\varepsilon}k \qquad \left(0 < \varepsilon \le \tfrac{3}{8}\right)$$

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le 109 \cdot \mathbf{E}[Z_k^{\mathrm{Opt}}]$$

Theorem: In both cases, $\lim_{\rho \to 1} \frac{\mathbf{E}[Z_k^{\mathrm{RS}}]}{\mathbf{E}[Z_k^{\mathrm{Opt}}]} = 1$

if $\mathbf{E}[S^3] < \infty$

Main Results

Multiserver systems are complicated

Go

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\text{RS}}] \le \mathbf{E}[Z_1^{\text{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\text{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\text{Opt}}]$$

Theorem: In M/G/k/dispatch with **guardrails** dispatching,

$$\mathbf{E}[Z_k^{\text{RS}}] \le \frac{11\varepsilon}{8}\mathbf{E}[Z_1^{\text{RS}}] + \frac{40}{\varepsilon}k \qquad (0 < \varepsilon \le \tfrac{3}{8})$$

$$\mathbf{E}[Z_k^{\text{RS}}] \le 109 \cdot \mathbf{E}[Z_k^{\text{Opt}}]$$

Theorem: In both cases, $\lim_{\rho \to 1} \frac{\mathbf{E}[Z_k^{\text{RS}}]}{\mathbf{E}[Z_k^{\text{Opt}}]} = 1$

if $\mathbf{E}[S^3] < \infty$

Main Results

Go…

Multiserver systems are complicated

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\text{RS}}] \le \mathbf{E}[Z_1^{\text{RS}}] + 6k$$

$$\mathbf{E}[Z_k^{\text{RS}}] \le 7 \cdot \mathbf{E}[Z_k^{\text{Opt}}]$$

Theorem: In M/G/k/dispatch with guardrails dispatching,

Need to dispatch to queues using **RS**

Theorem: In both cases, $\displaystyle \lim_{\rho \to 1} \frac{\mathbf{E}[Z_k^{\text{RS}}]}{\mathbf{E}[Z_k^{\text{Opt}}]} = 1$

if $\mathbf{E}[S^3] < \infty$

Main Results

Go

Theorem: In M/G/k,

$E[Z_k^{RS}] \le E[Z_1^{RS}] + 6k$

Theorem: In M/G/k/dispatch with guardrails dispatching,

ases, $\lim_{\rho \to 1} \frac{E[Z_k^{RS}]}{E[Z_k^{Opt}]} = 1$

if $E[S^3] < \infty$

Multiserver systems are complicated

Need to dispatch to queues using **RS**

Heavy-traffic $\mathbf{E}[Z]$ poorly understood, even in M/G/1

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \le \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \leq \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

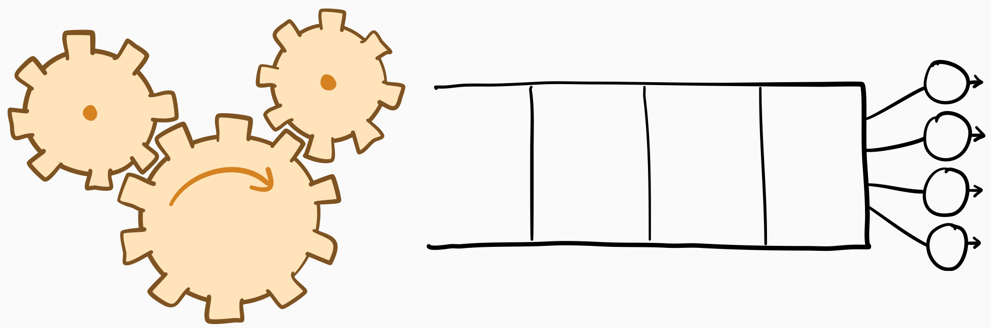

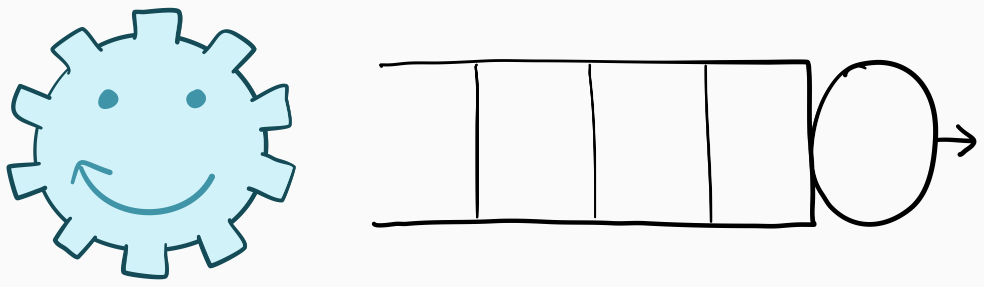

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \leq \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

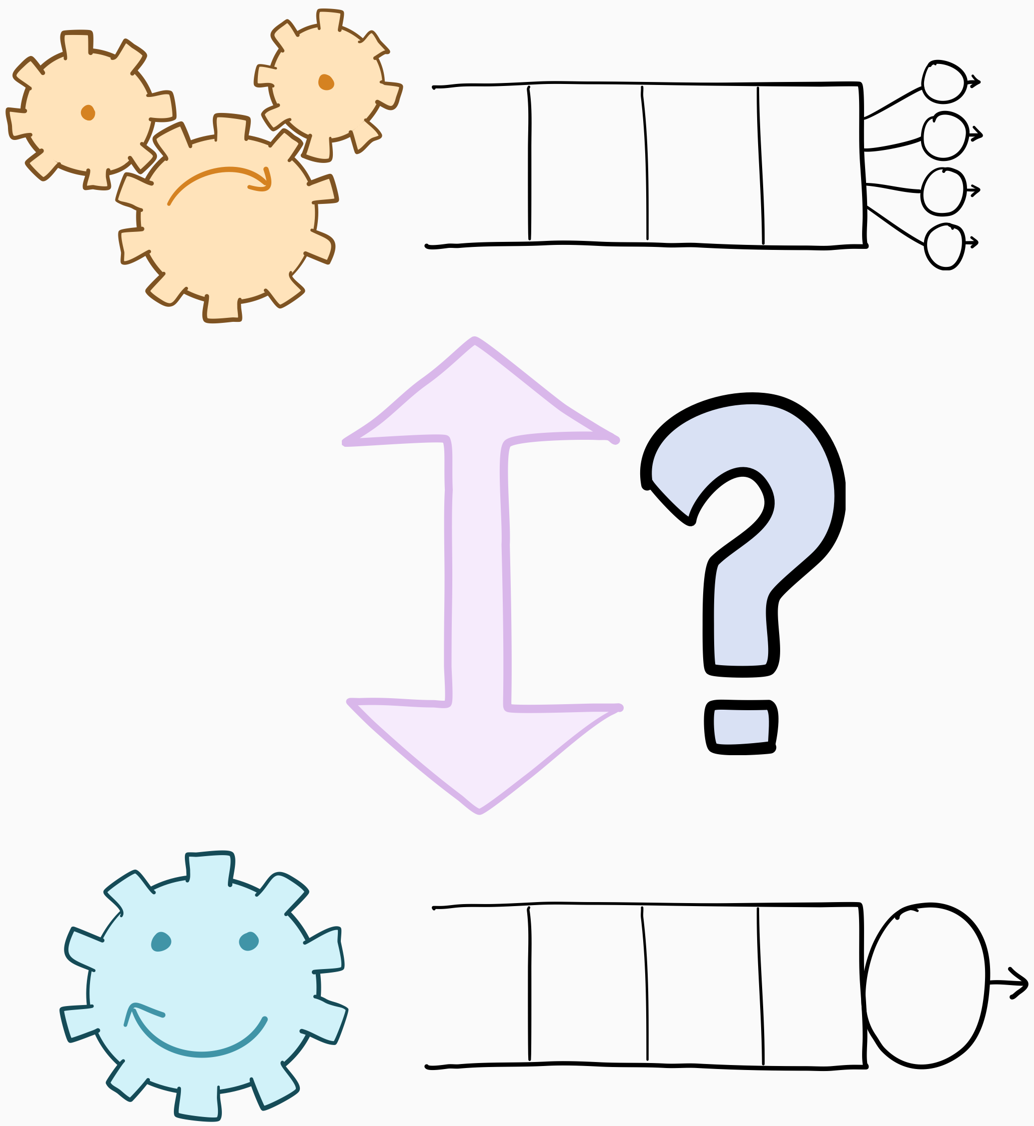

Theorem: In M/G/k,

$$\mathbf{E}[Z_k^{\mathbf{RS}}] \leq \mathbf{E}[Z_1^{\mathbf{RS}}] + 6k$$

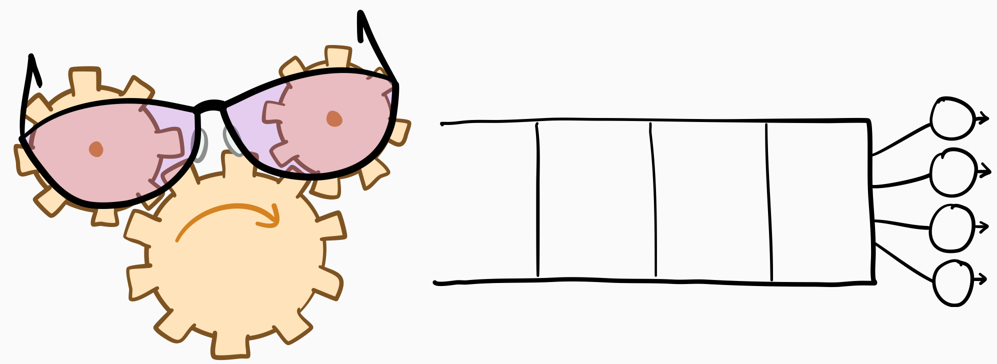

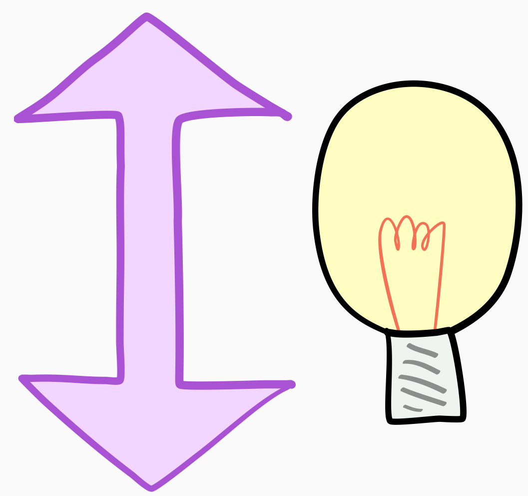

Key idea:
***r*-work**

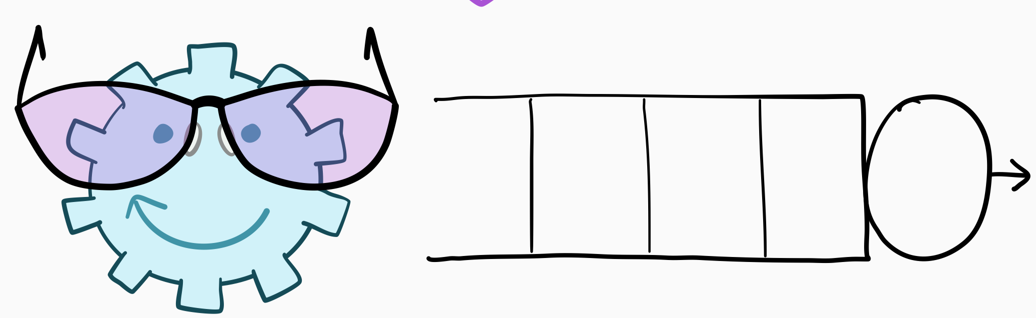

What is *r*-Work?

W = work = total remaining size of all jobs

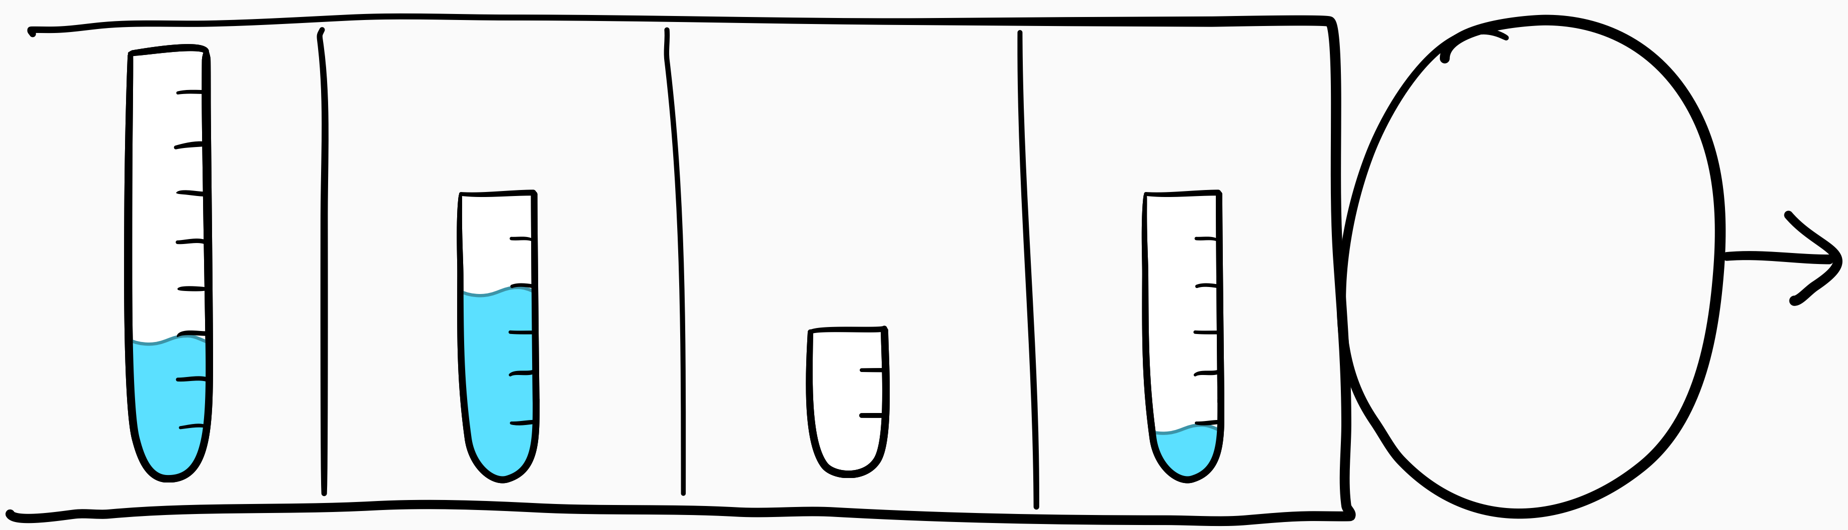

What is *r*-Work?

W = work = total remaining size of all jobs

$W(r)$ = *r*-work = total remaining size of all jobs that have **rank** $\leq r$

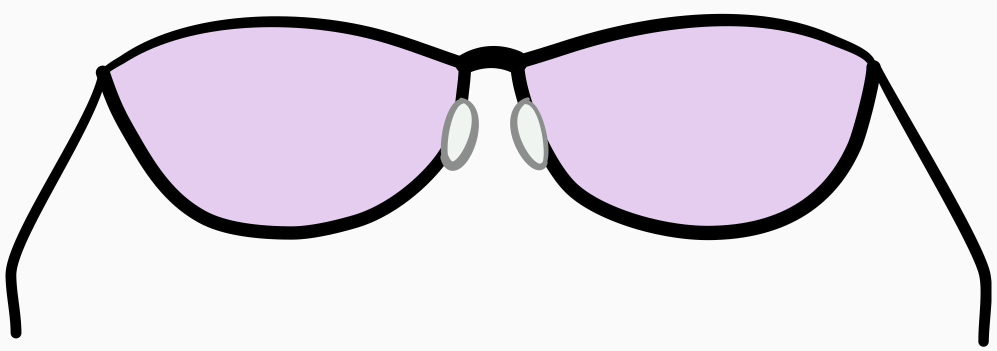

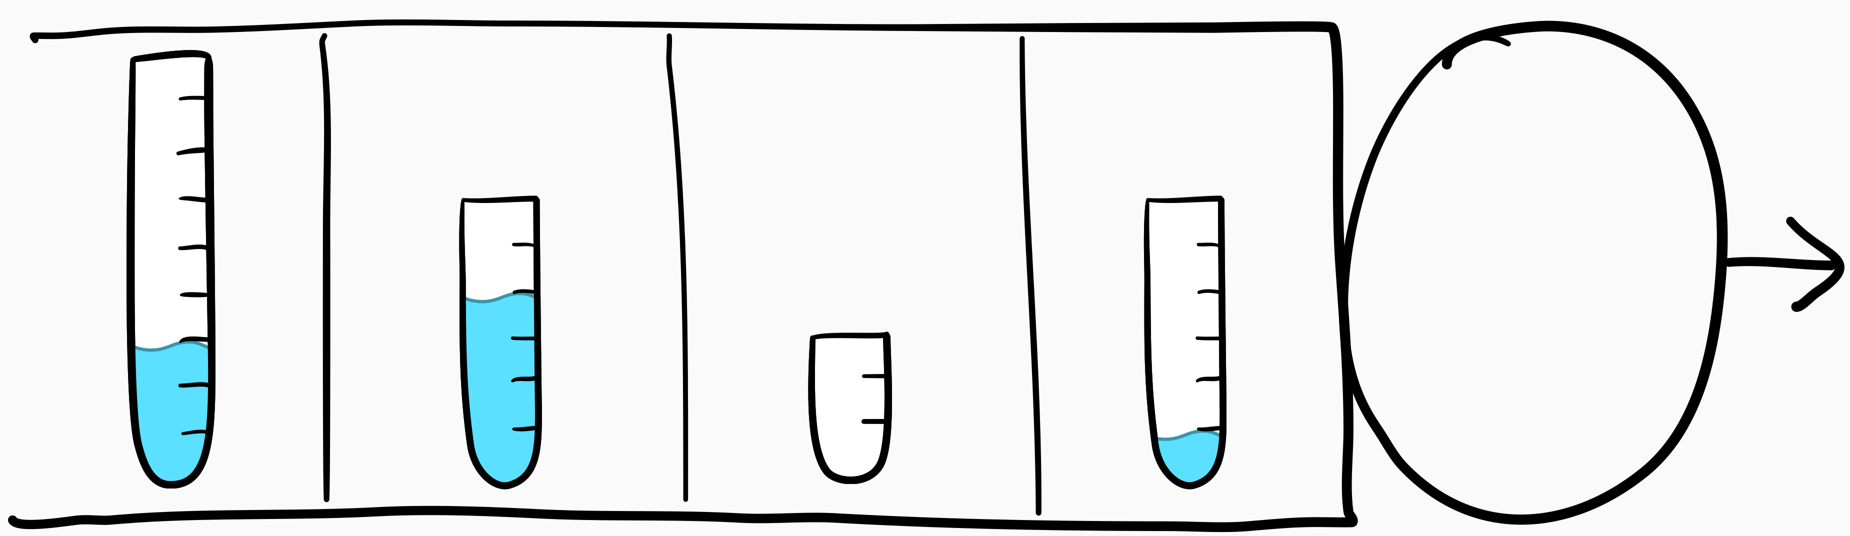

What is *r*-Work?

W = work = total remaining size of all jobs

$W(r)$ = *r*-work = total remaining size of all jobs
that have **rank** $\leq r$

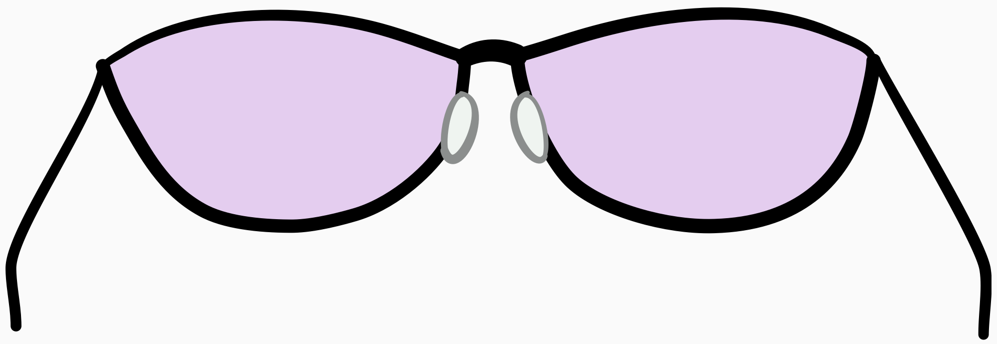

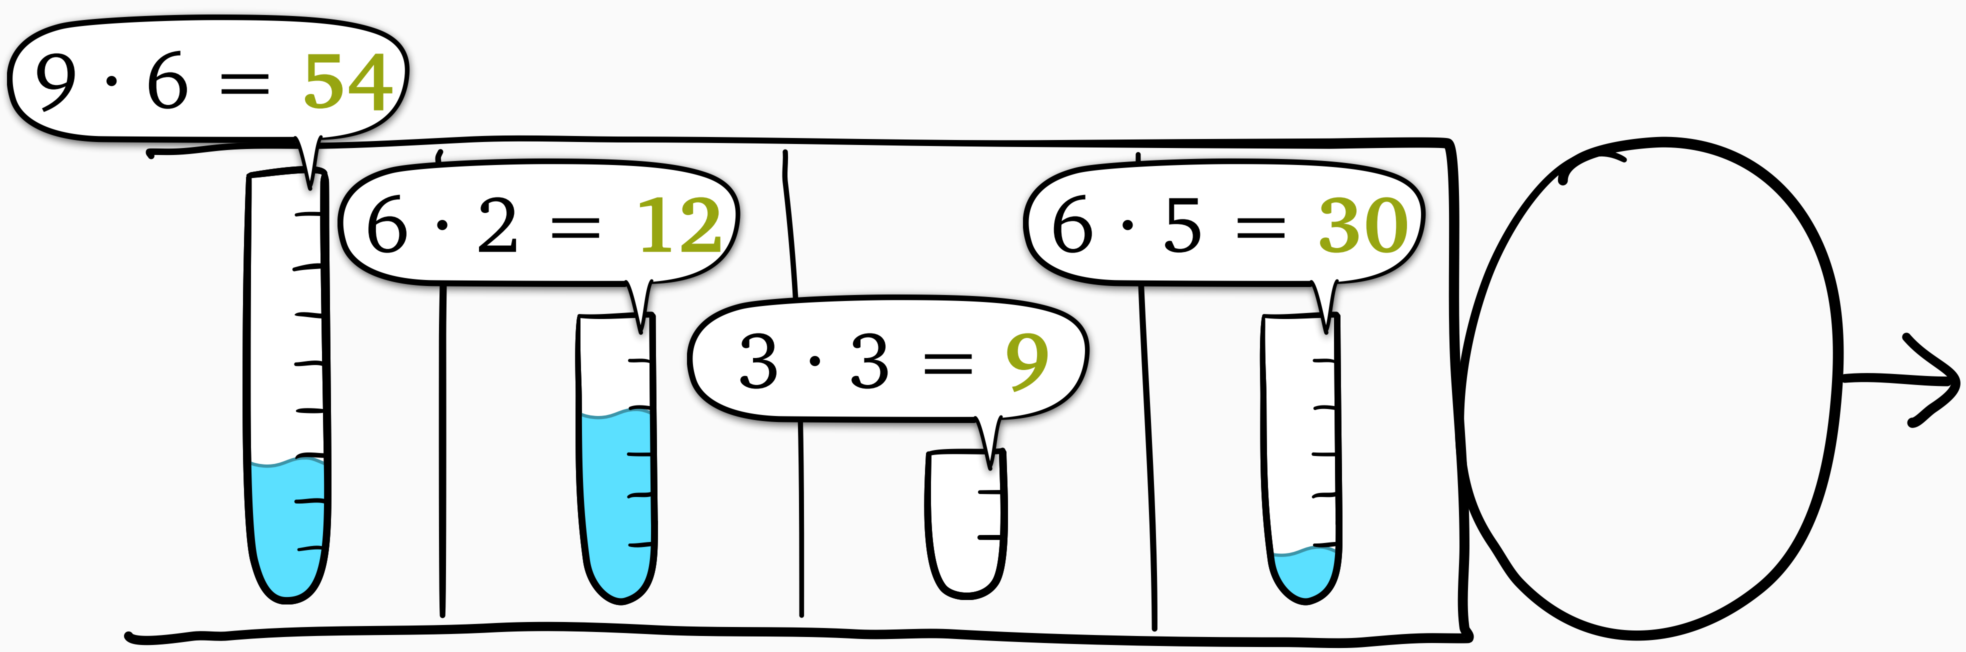

What is *r*-Work?

W = work = total remaining size of all jobs

$W(r)$ = r-**work** = total remaining size of all jobs that have **rank** $\leq r$

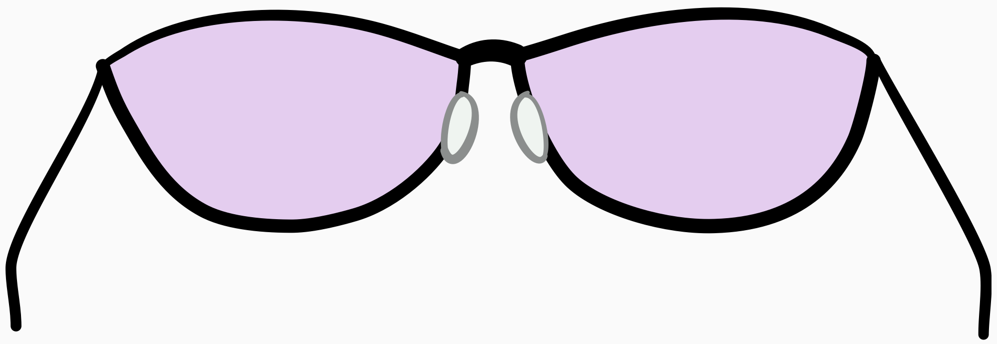

20-work = ?

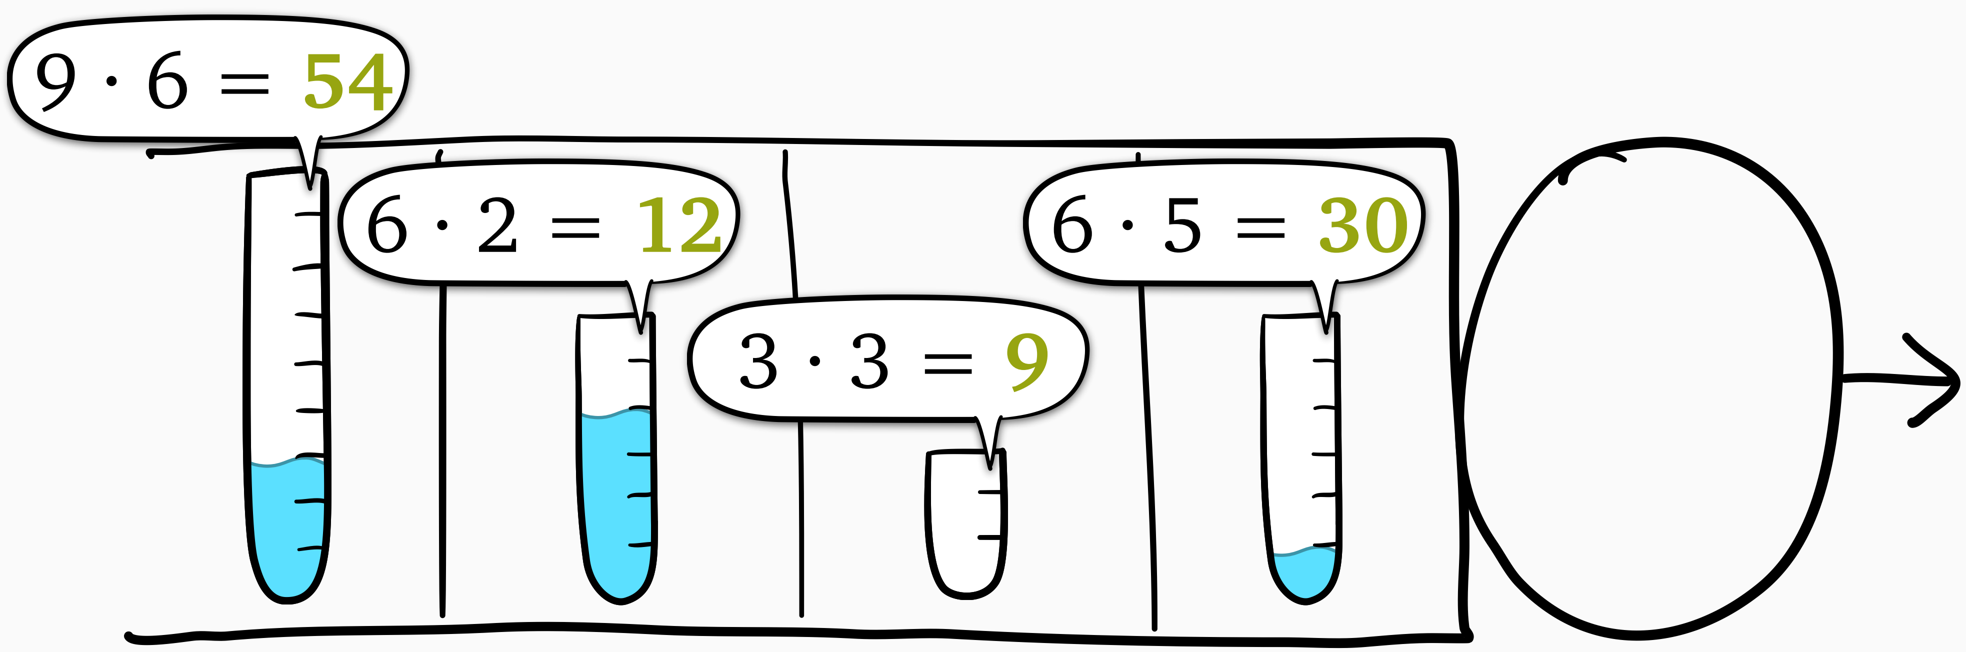

What is *r*-Work?

W = work = total remaining size of all jobs

$W(r)$ = r-work = total remaining size of all jobs that have **rank** $\leq r$

20-work = ?

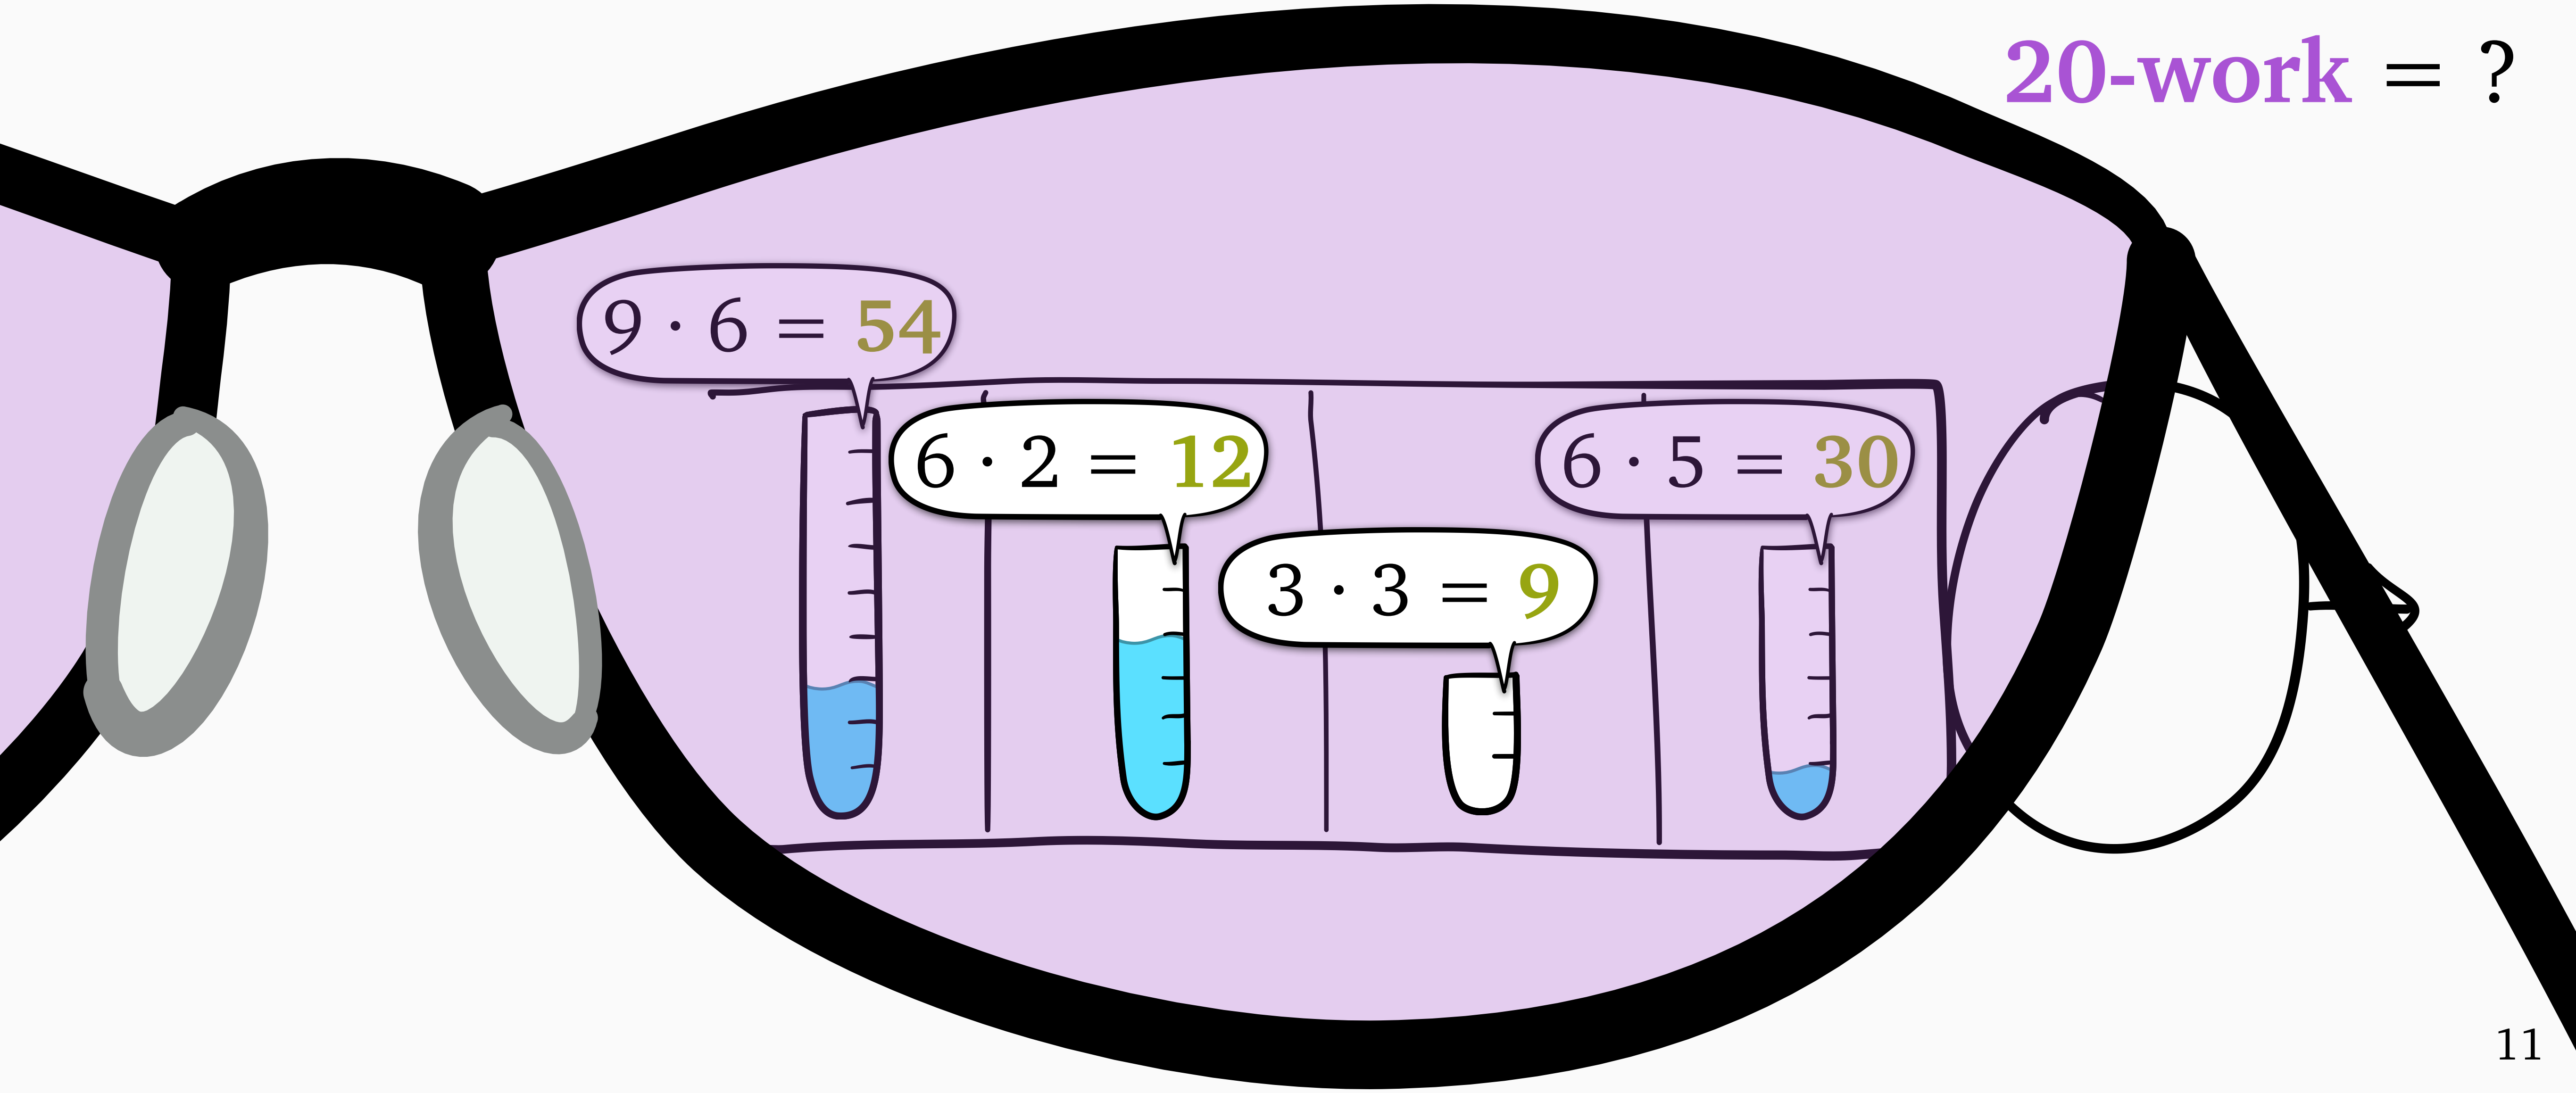

What is *r*-Work?

W = work = total remaining size of all jobs

$W(r)$ = r-work = total remaining size of all jobs that have **rank** $\leq r$

20-work = 5

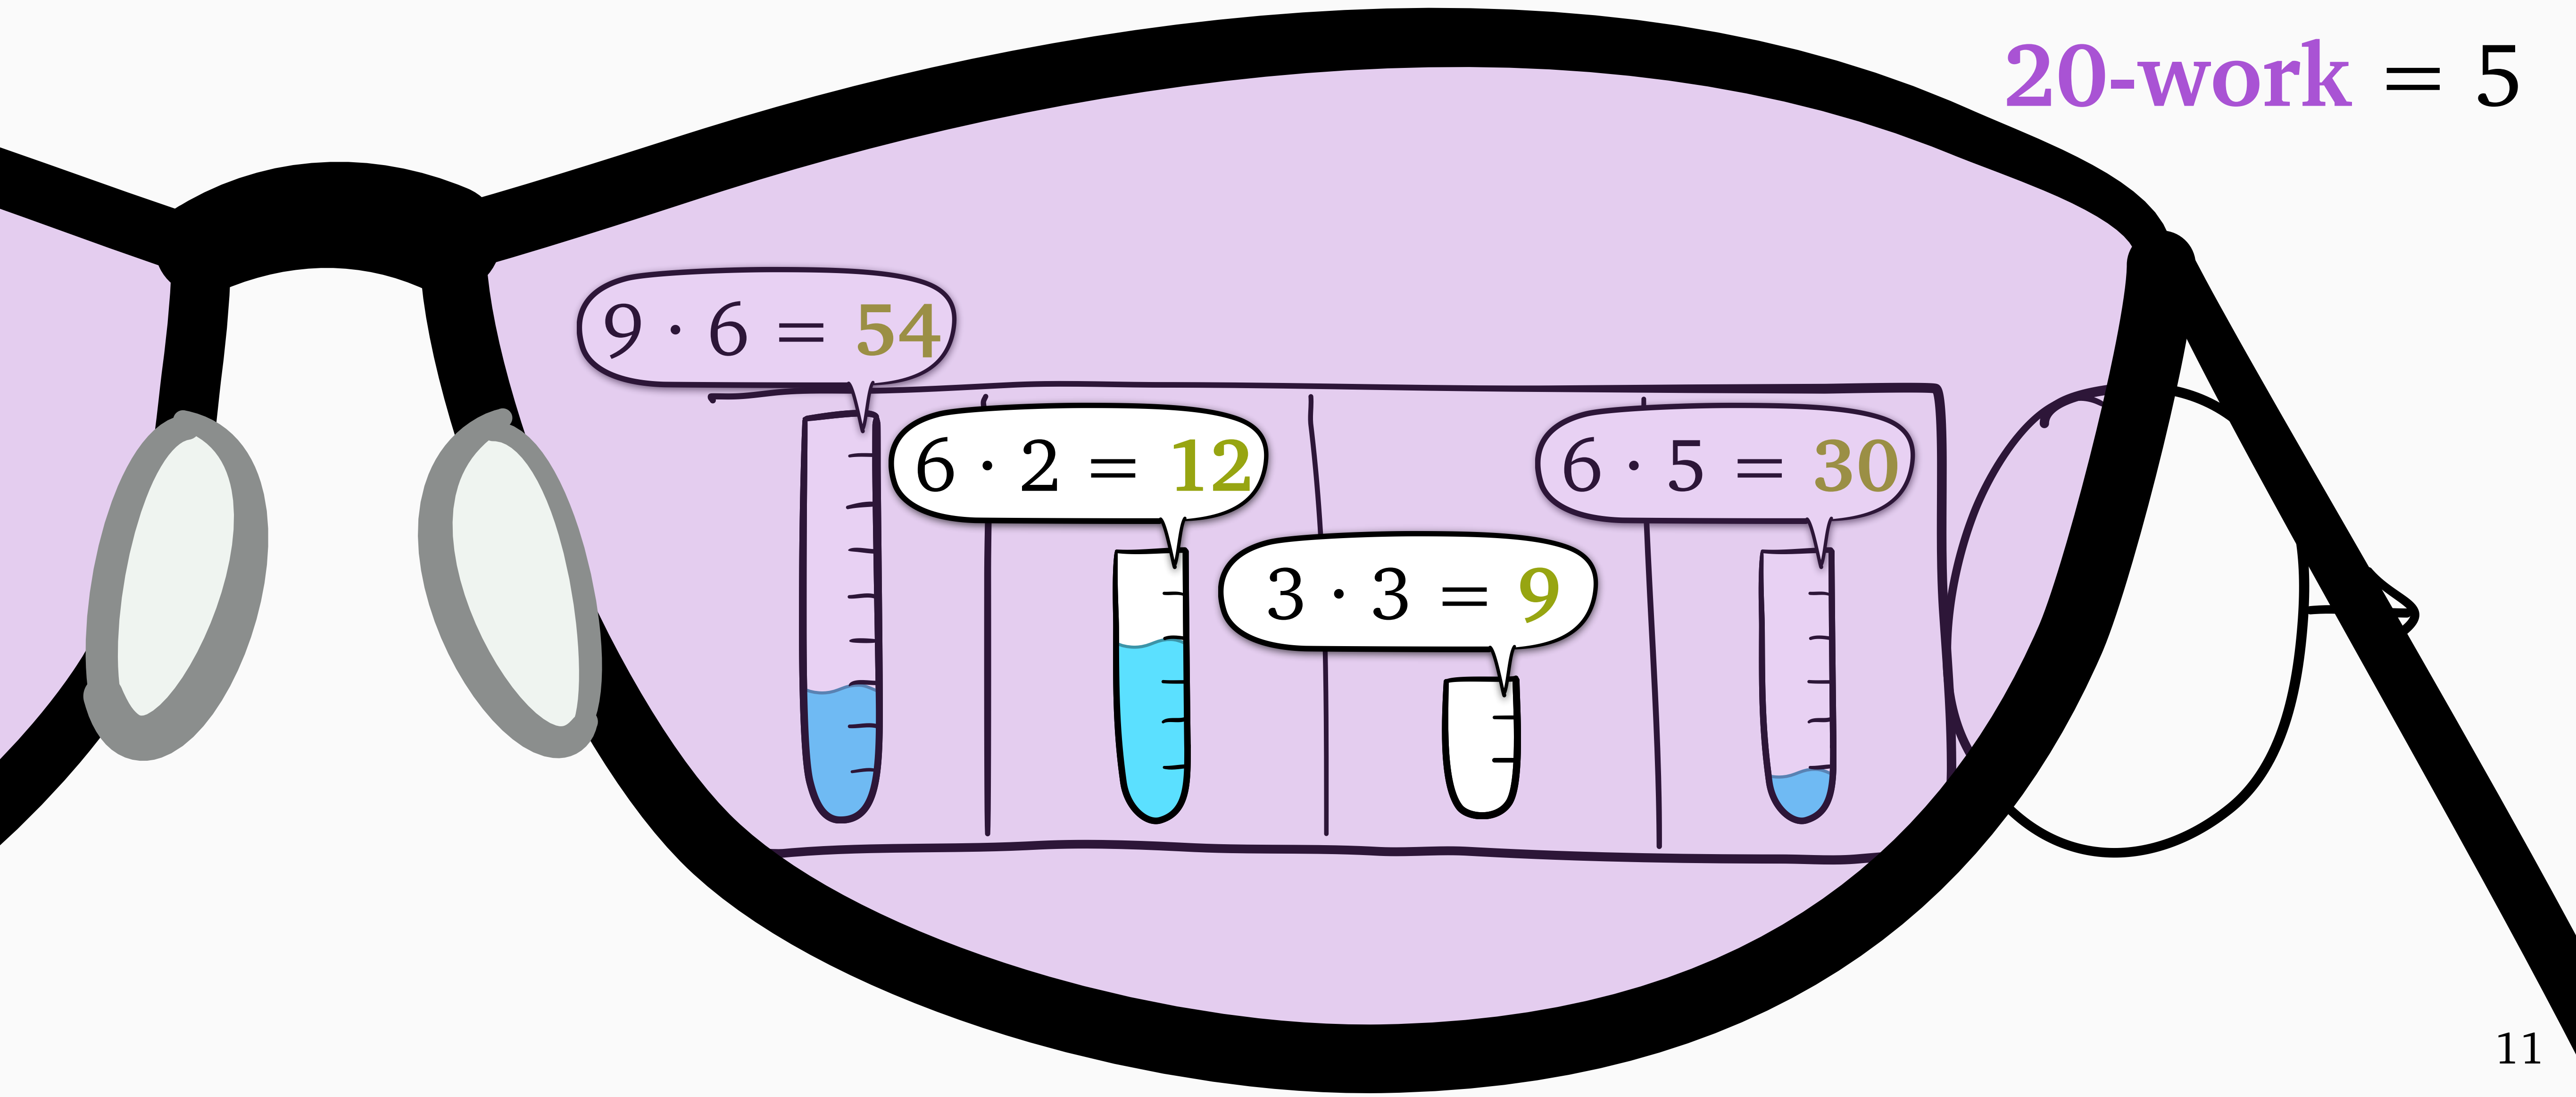

Slowdown via *r*-Work

$\mathbf{E}[Z_k]$

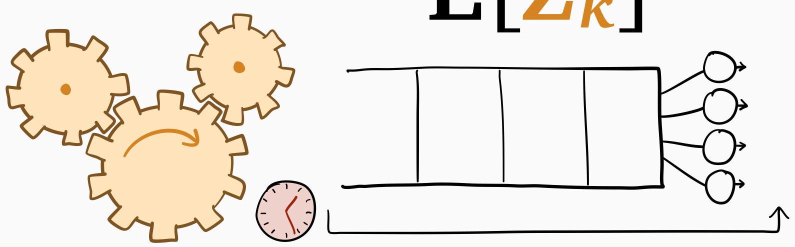

$\mathbf{E}[Z_1]$

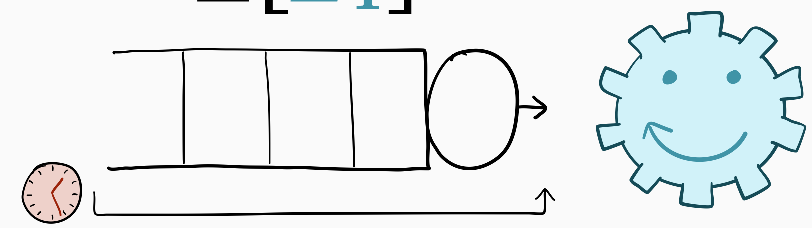

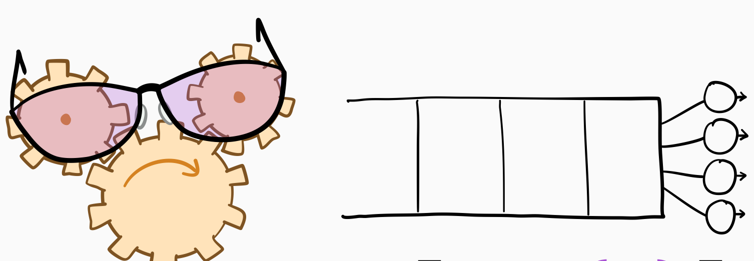

$\mathbf{E}[W_k(r)]$

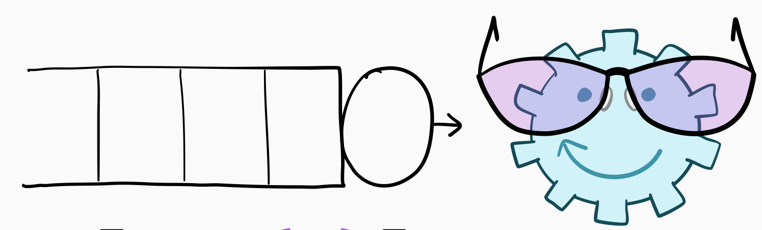

$\mathbf{E}[W_1(r)]$

Slowdown via *r*-Work

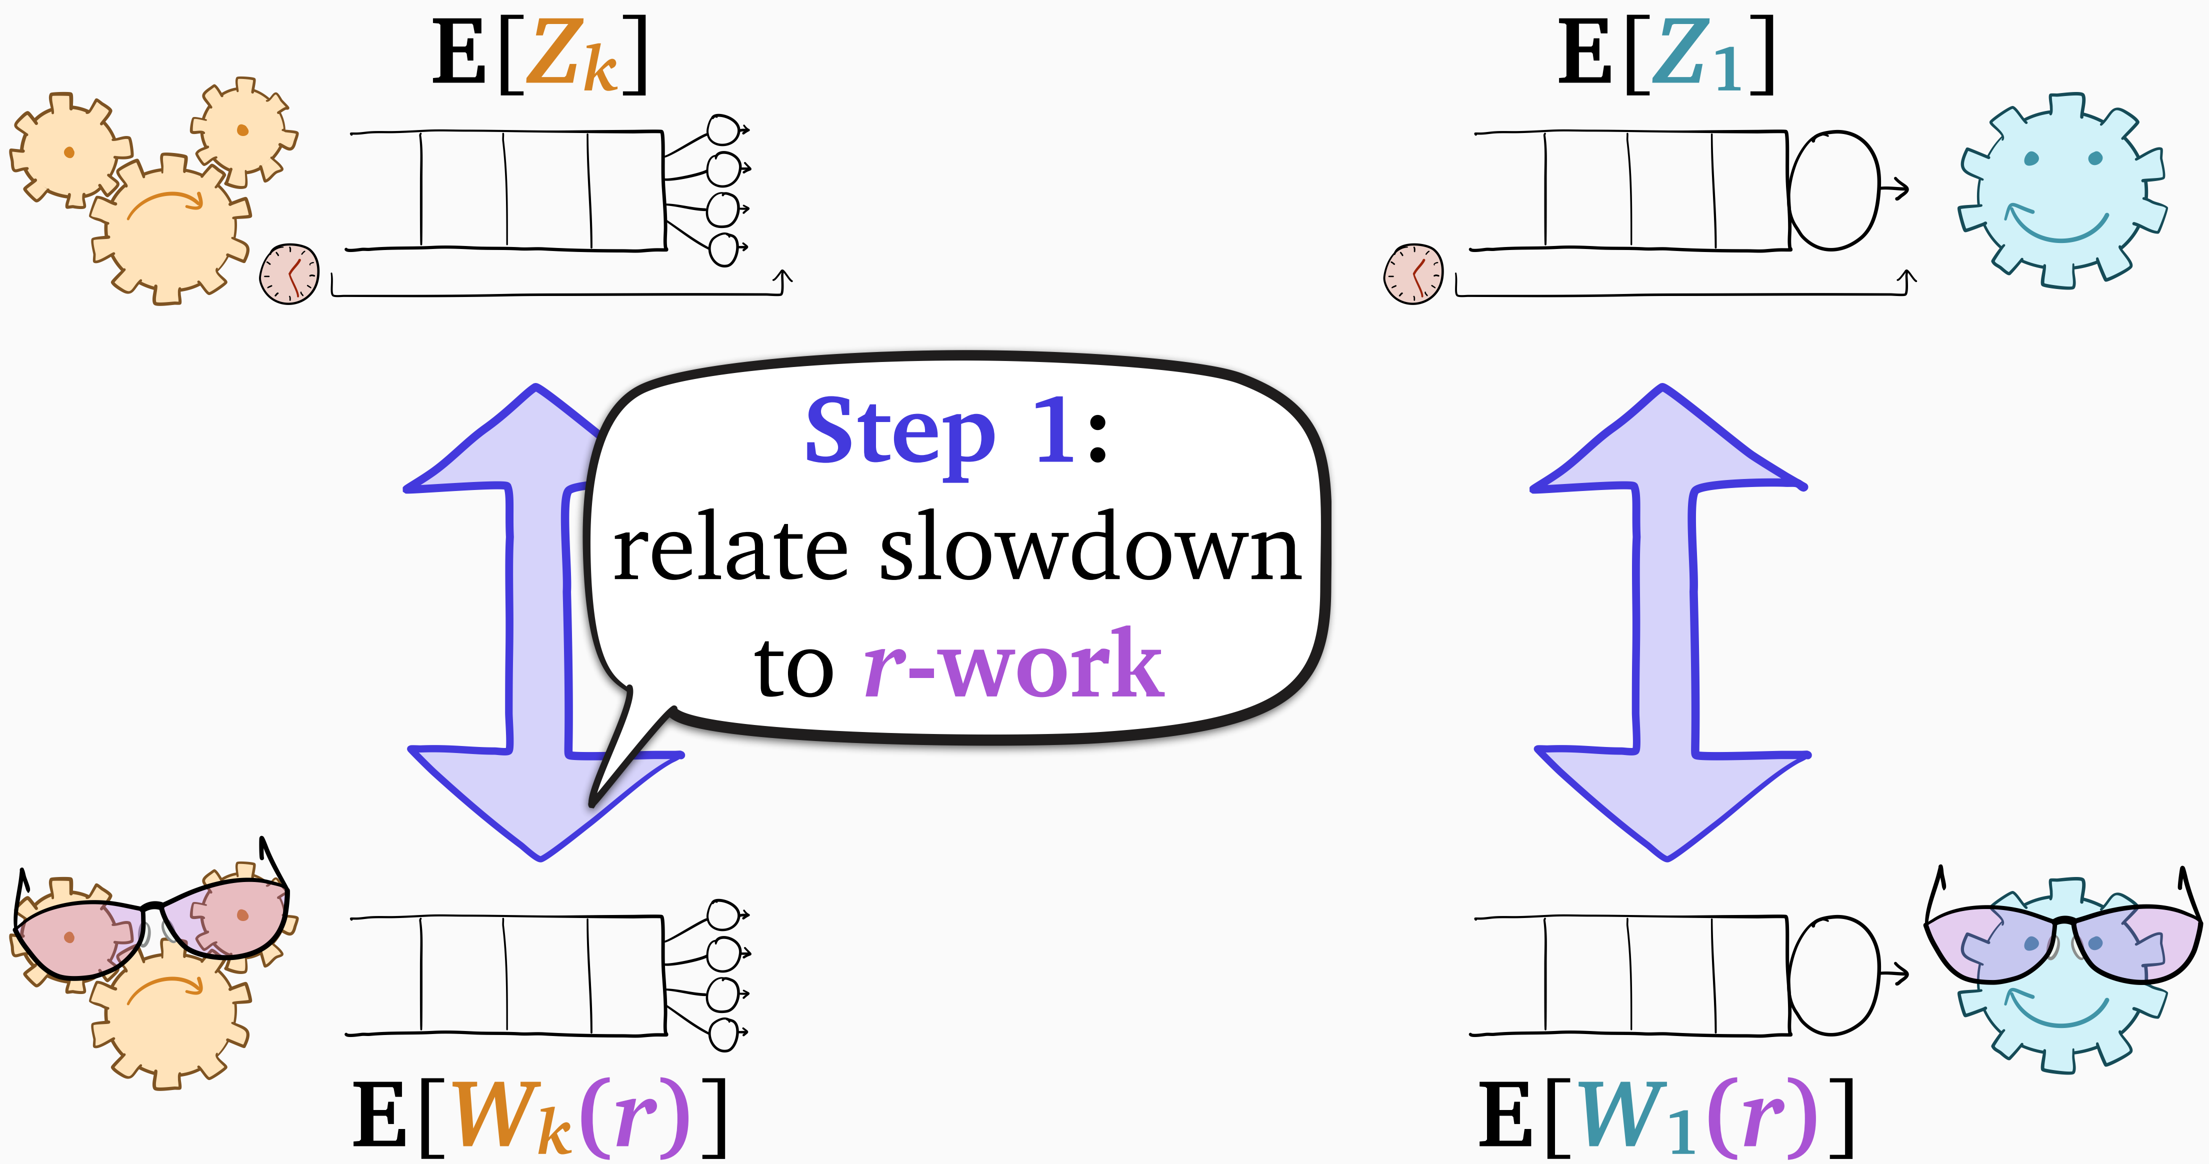

Slowdown via *r*-Work

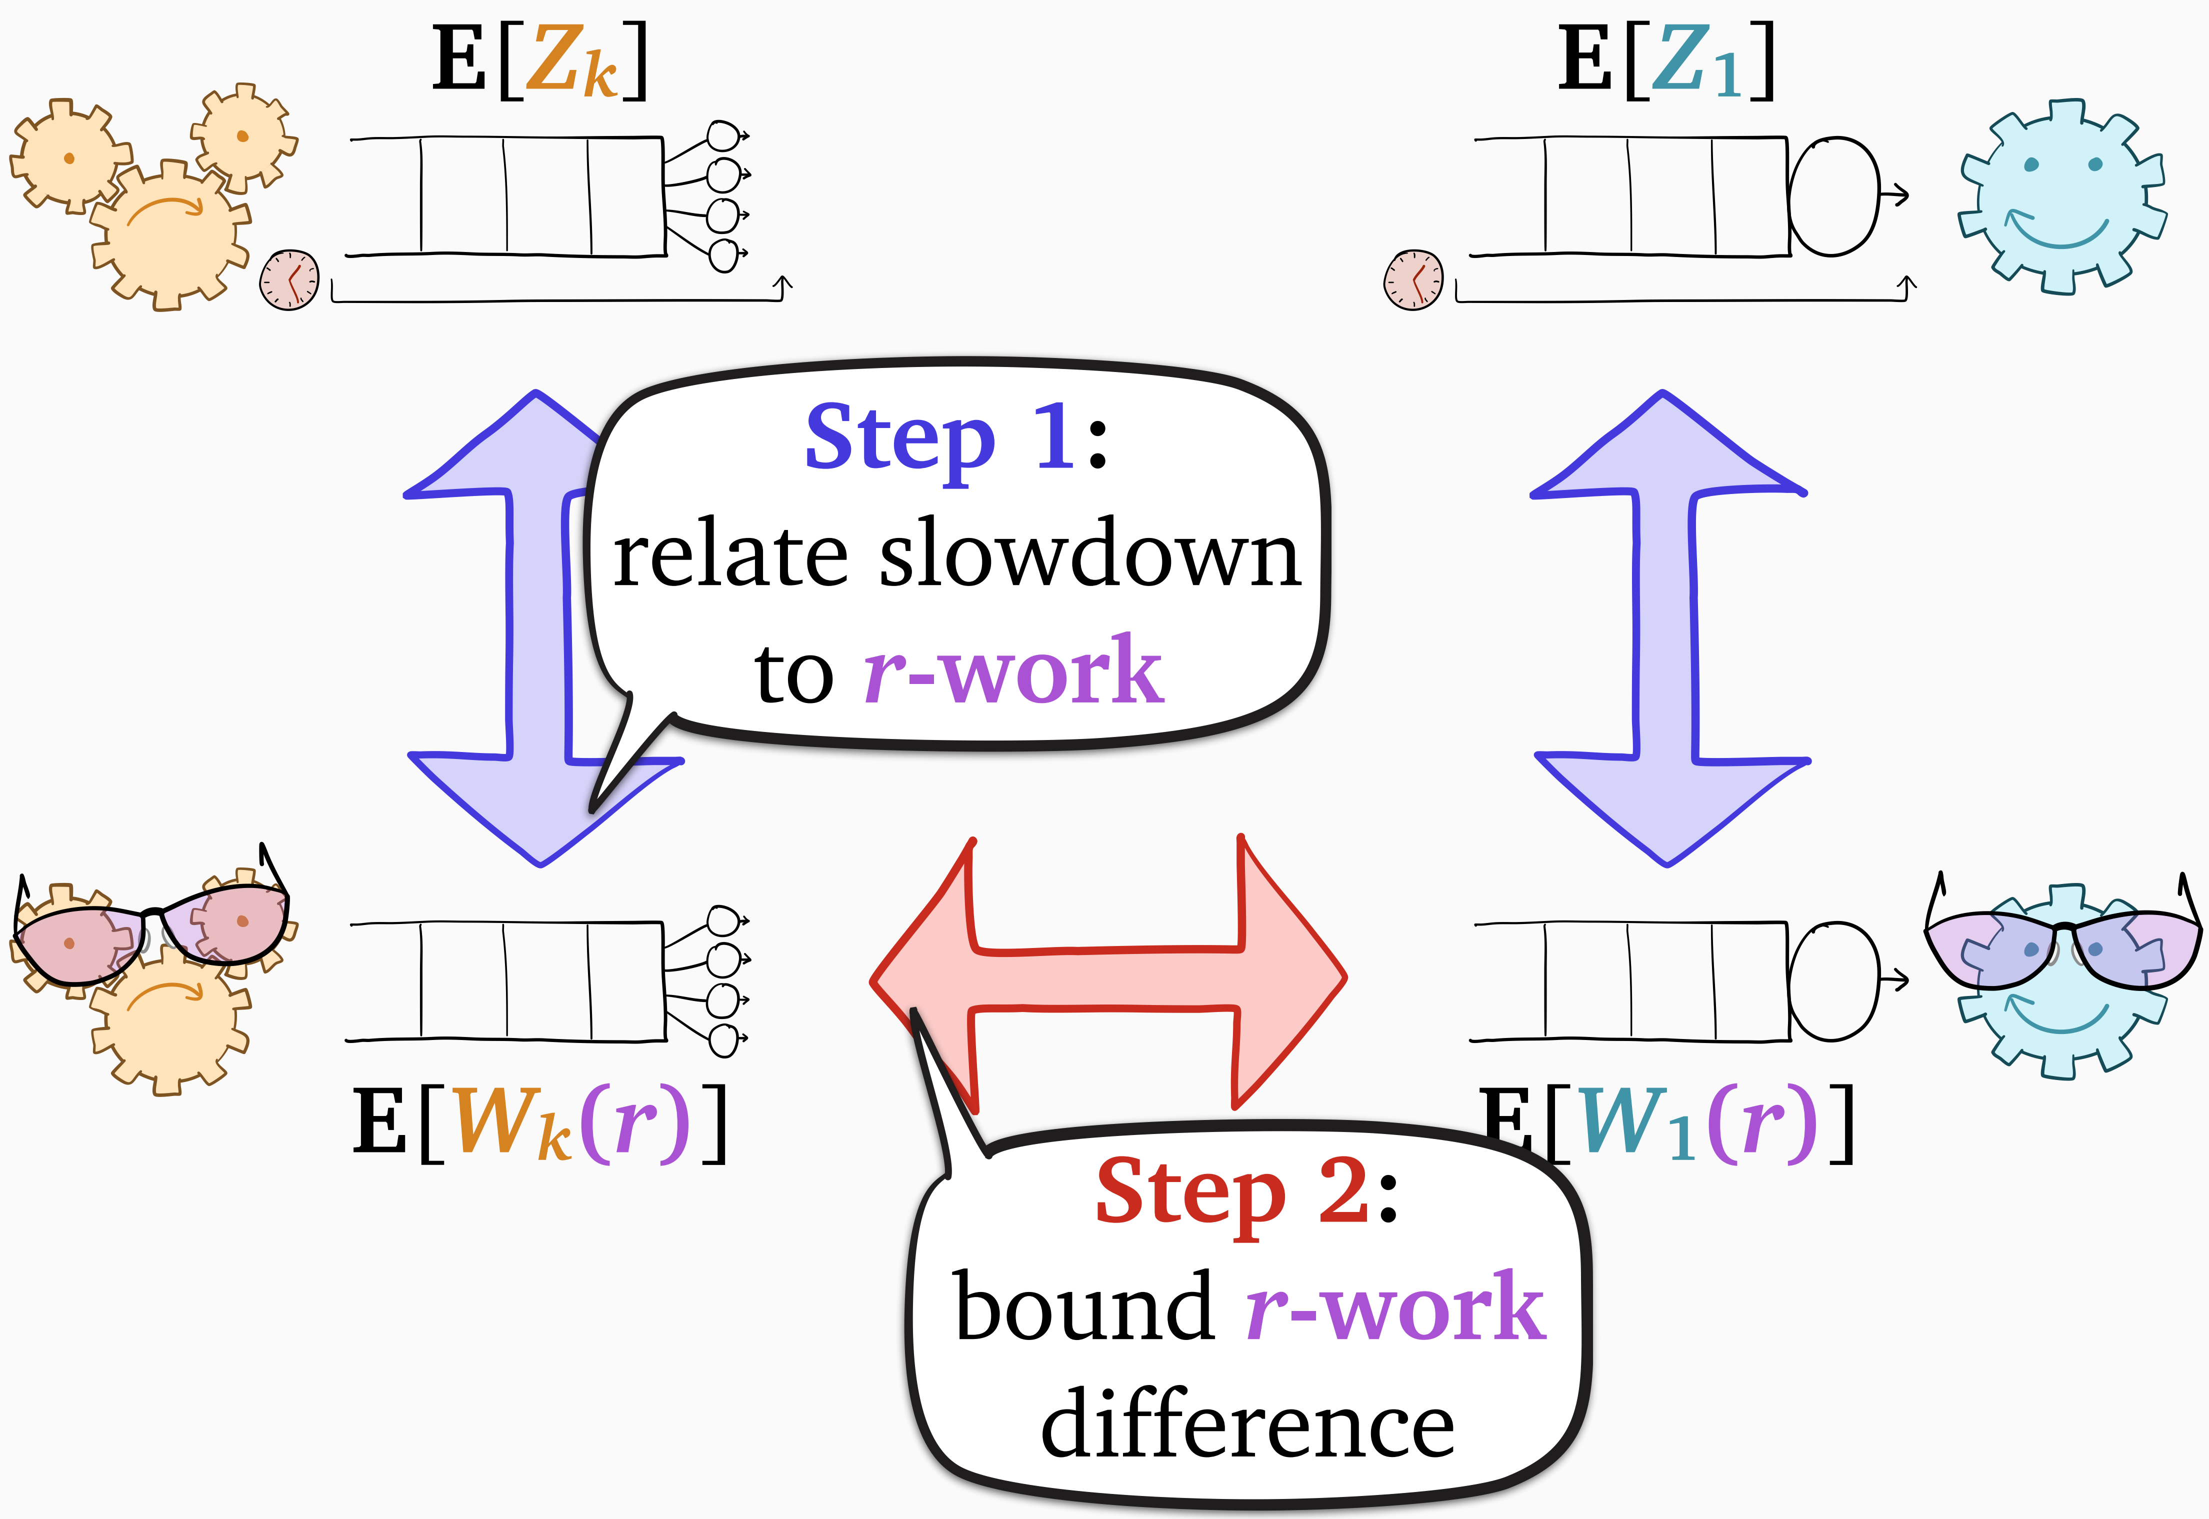

Slowdown via *r*-Work

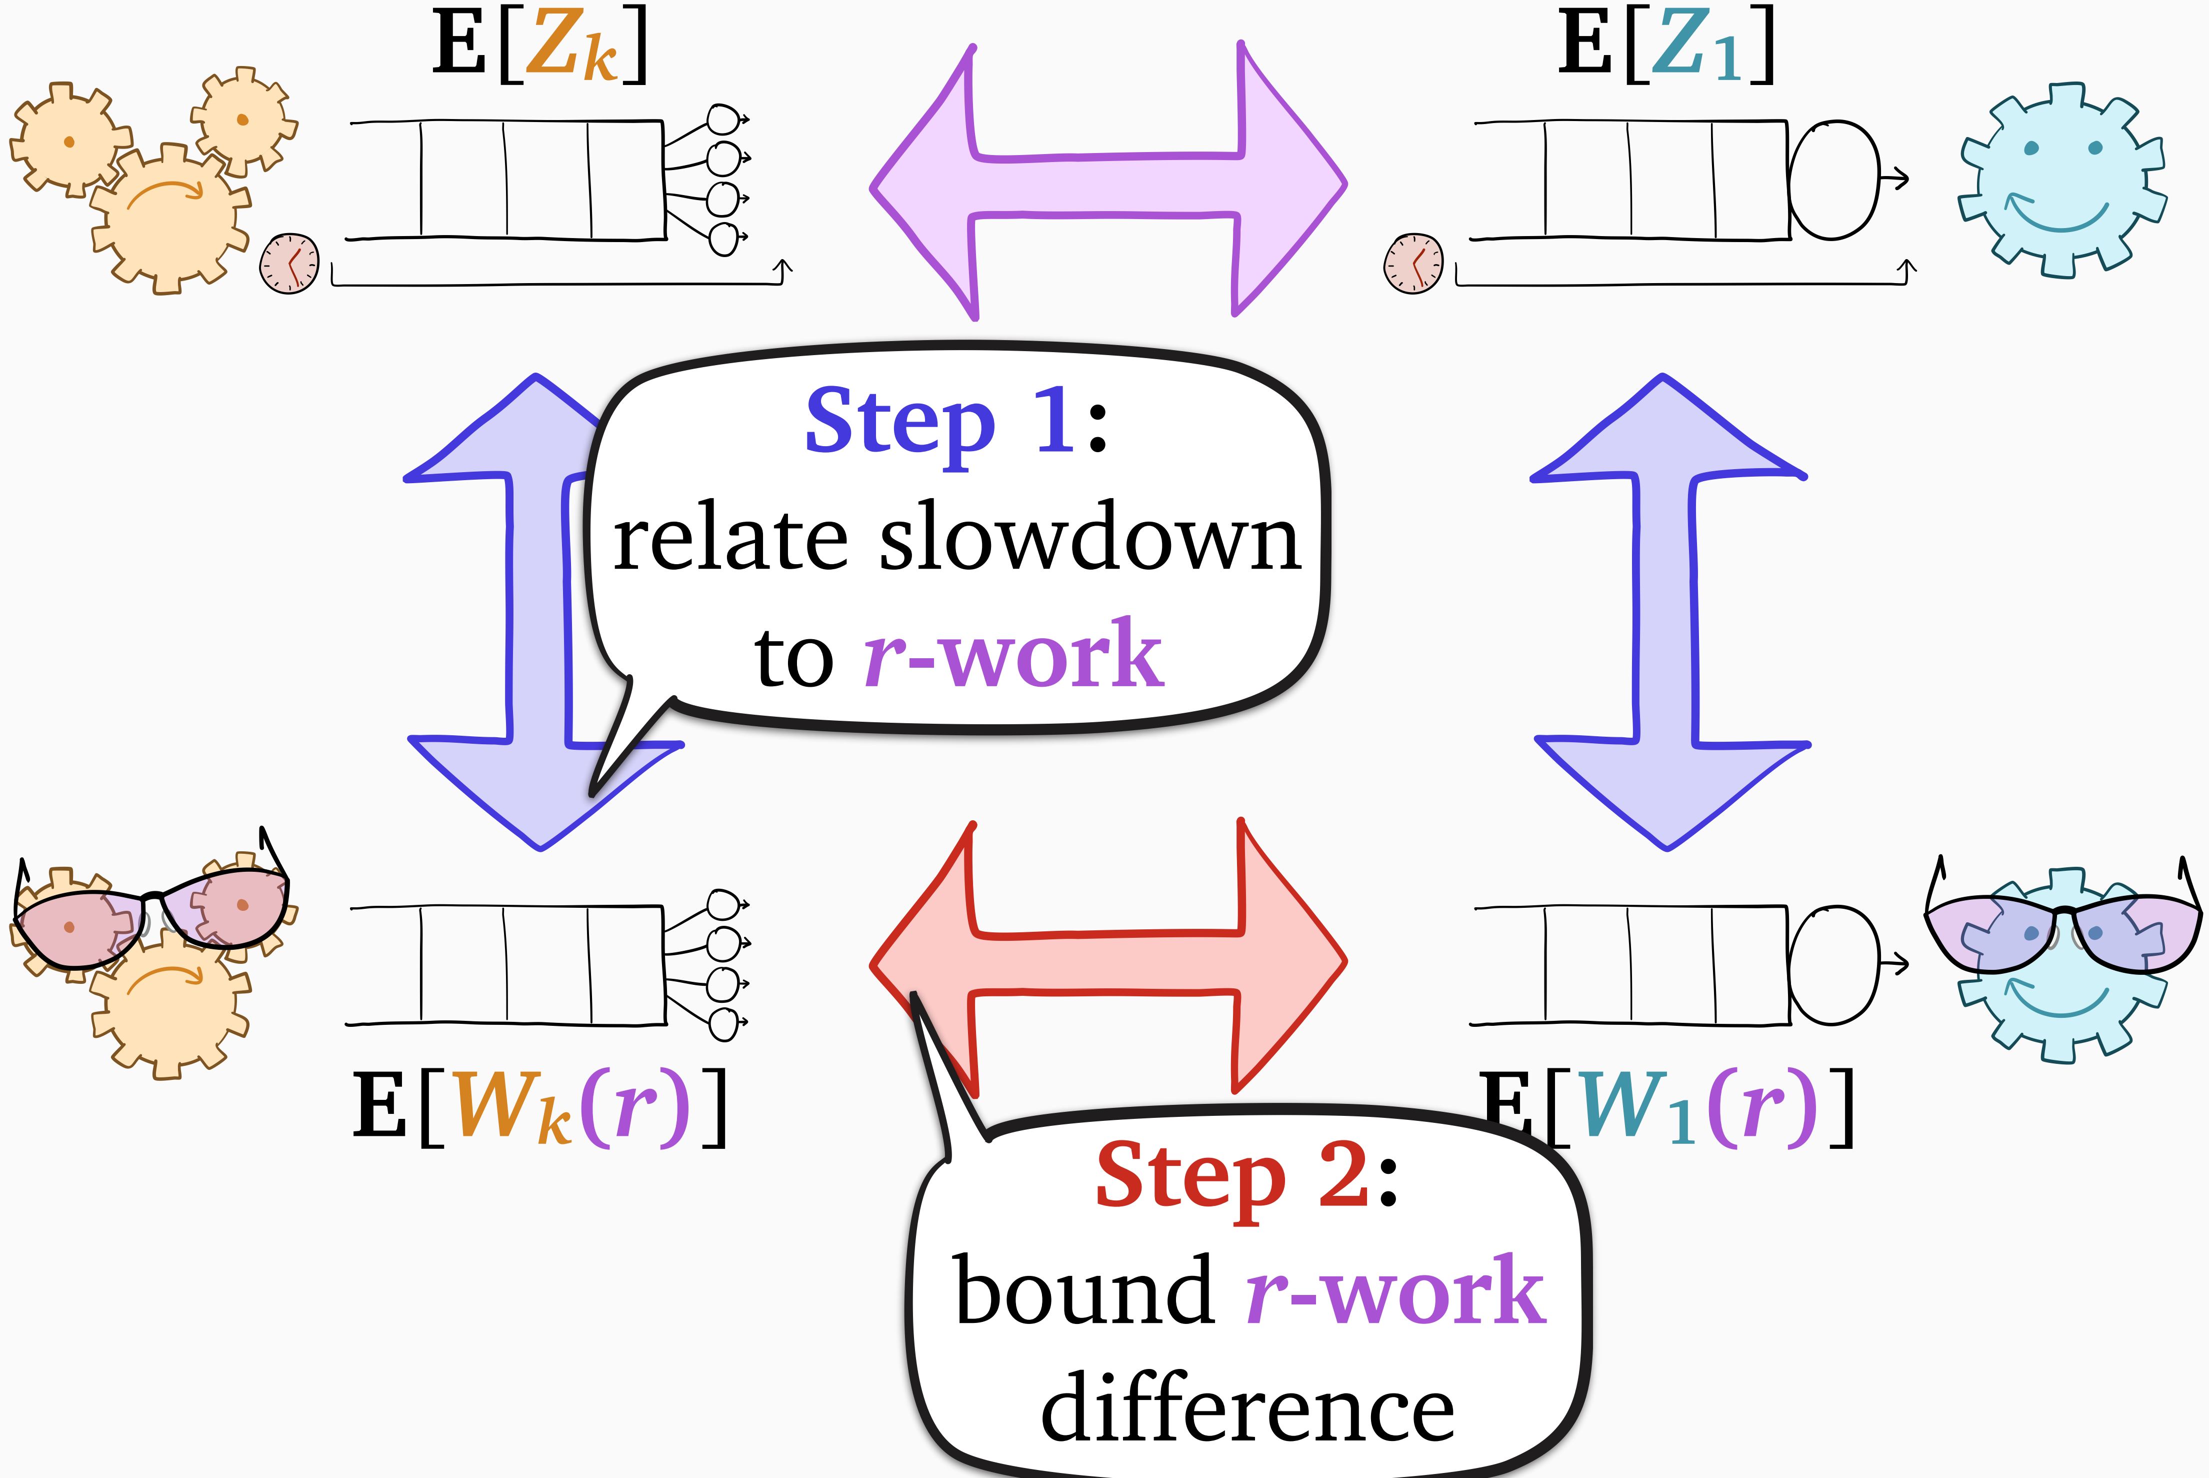

Step 1: $\mathbf{E}[Z]$ to $\mathbf{E}[H]$

Holding cost of job of size $s = 1/s$

$H =$ *total holding cost* of all jobs in system

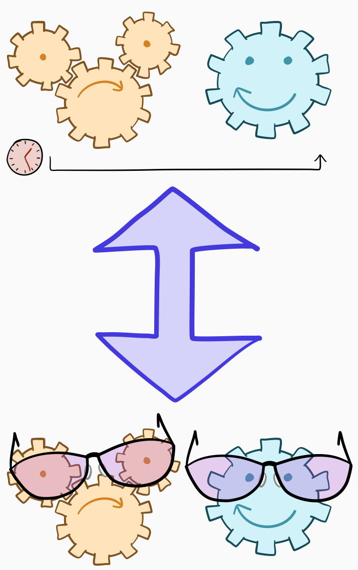

Step 1: $\mathbf{E}[Z]$ to $\mathbf{E}[H]$

Holding cost of job of size $s = 1/s$

$H =$ *total holding cost* of all jobs in system

Generalized Little's law:

$$\mathbf{E}[H] = \lambda \mathbf{E}[Z]$$

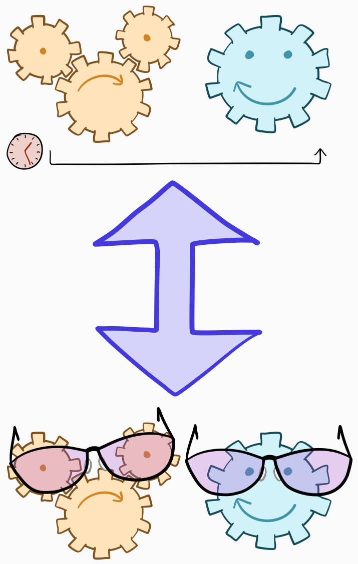

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

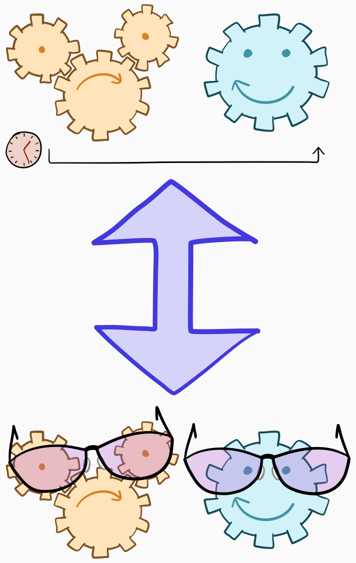

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\, \mathrm{d}r = \int_0^\infty W(r)\, \mathrm{d}(1/r)$$

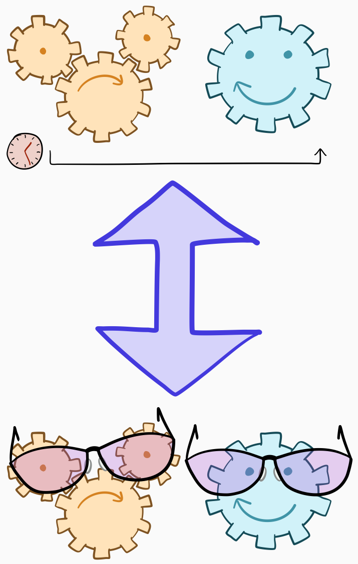

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

One job's ***r*-work**:

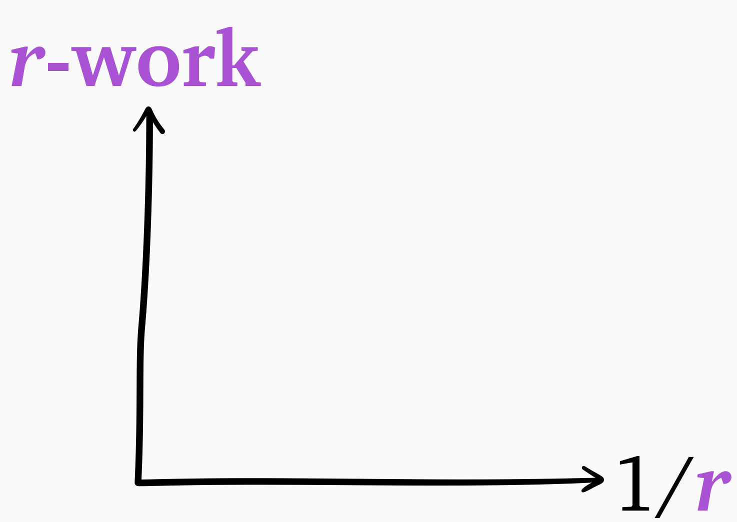

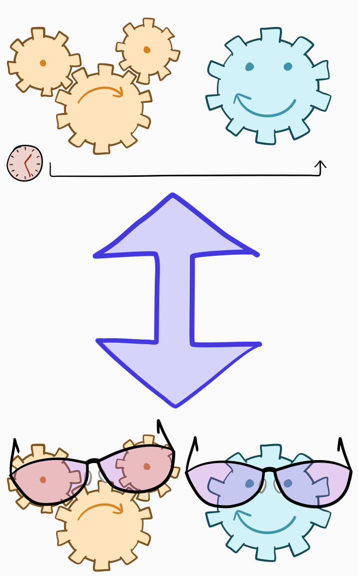

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's **r-work**:

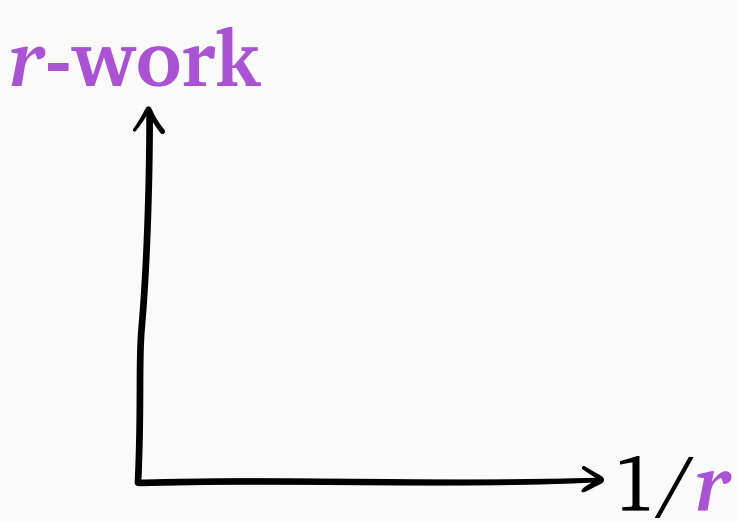

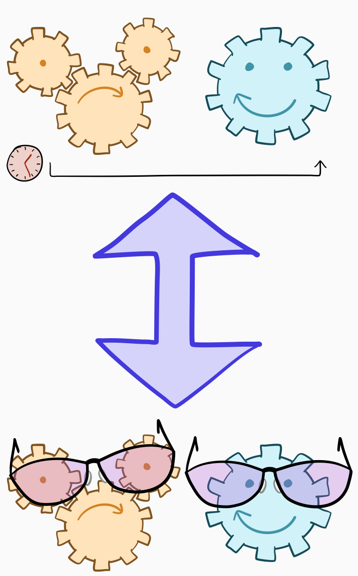

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's **r-work**:

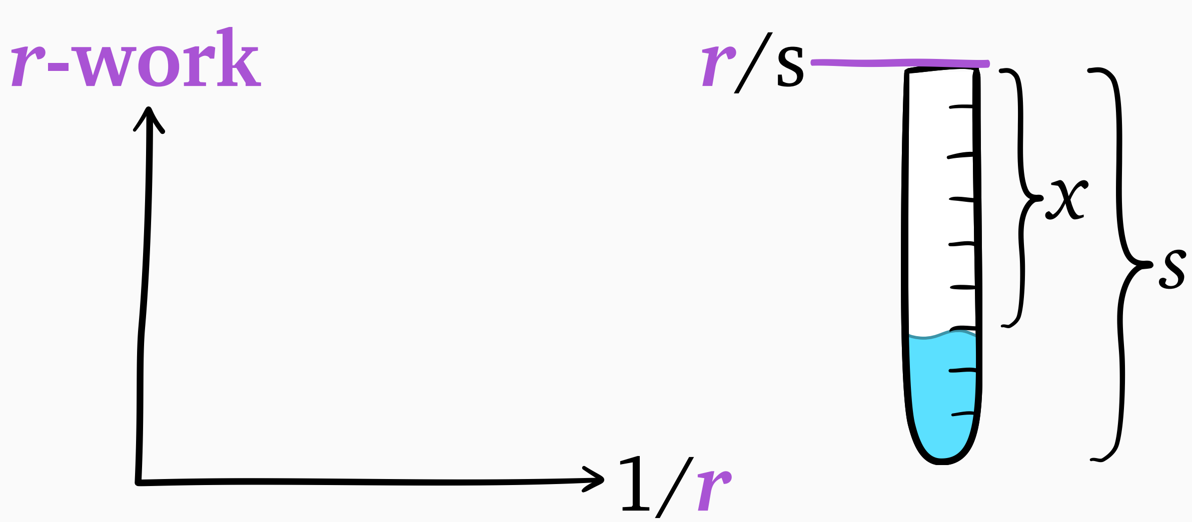

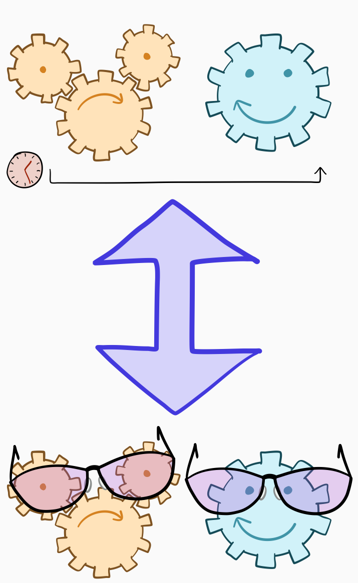

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's ***r*-work**:

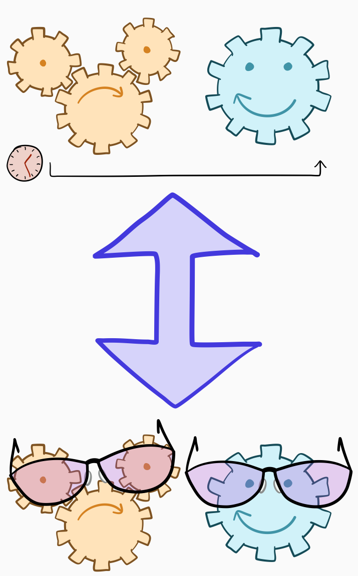

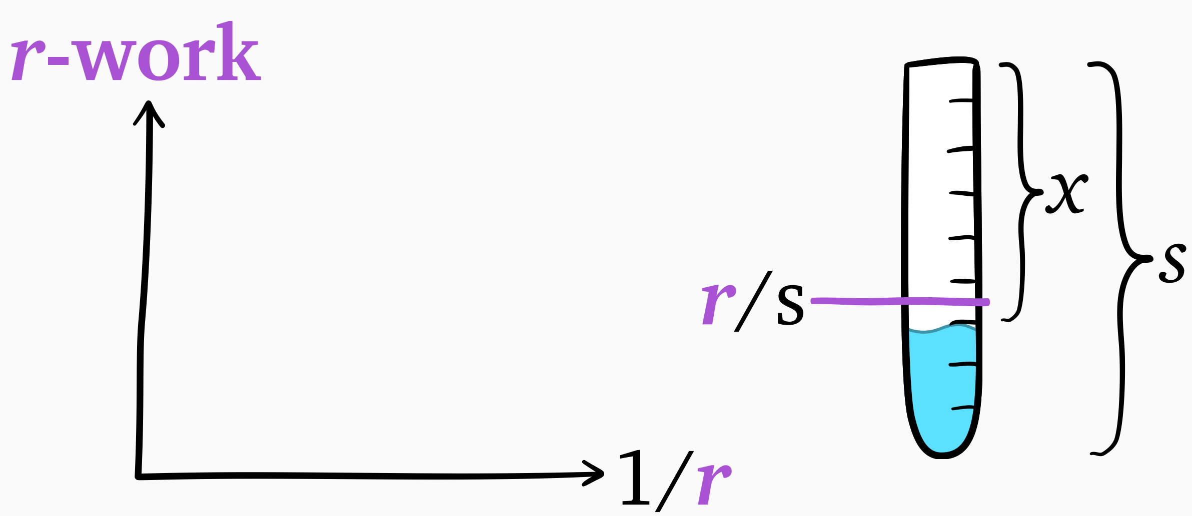

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's **r-work**:

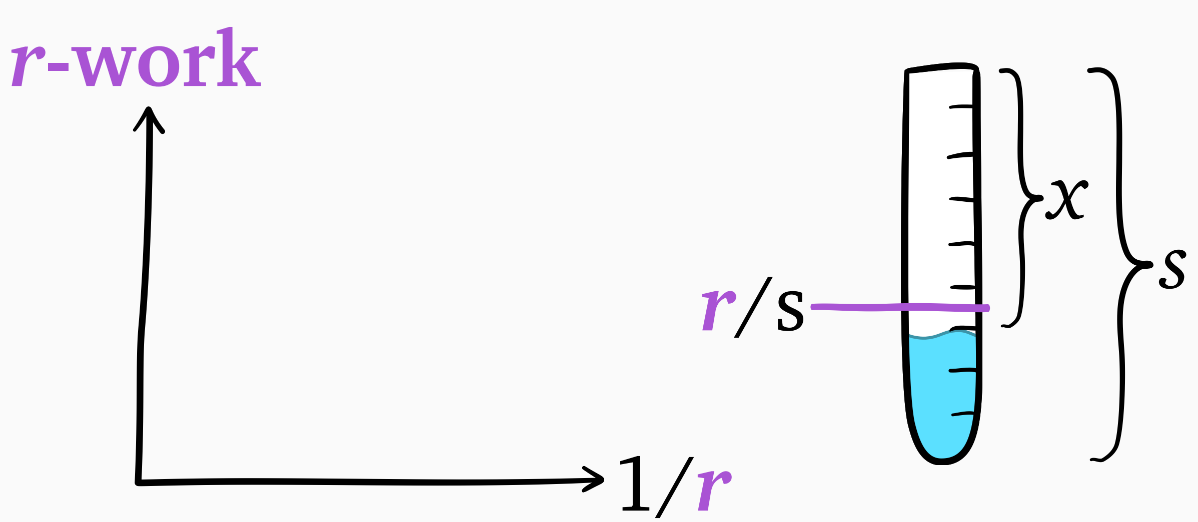

$r < sx$: **r-work** $= 0$

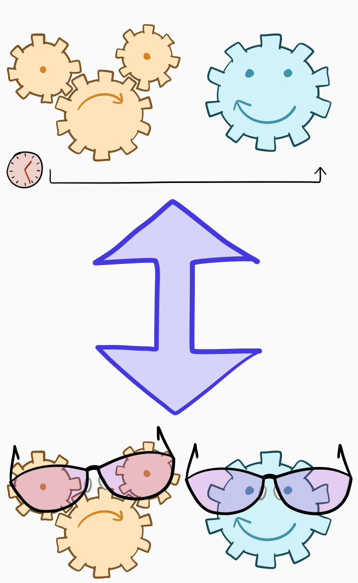

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's ***r*-work**:

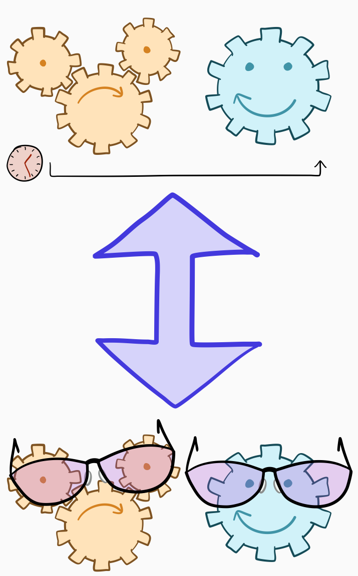

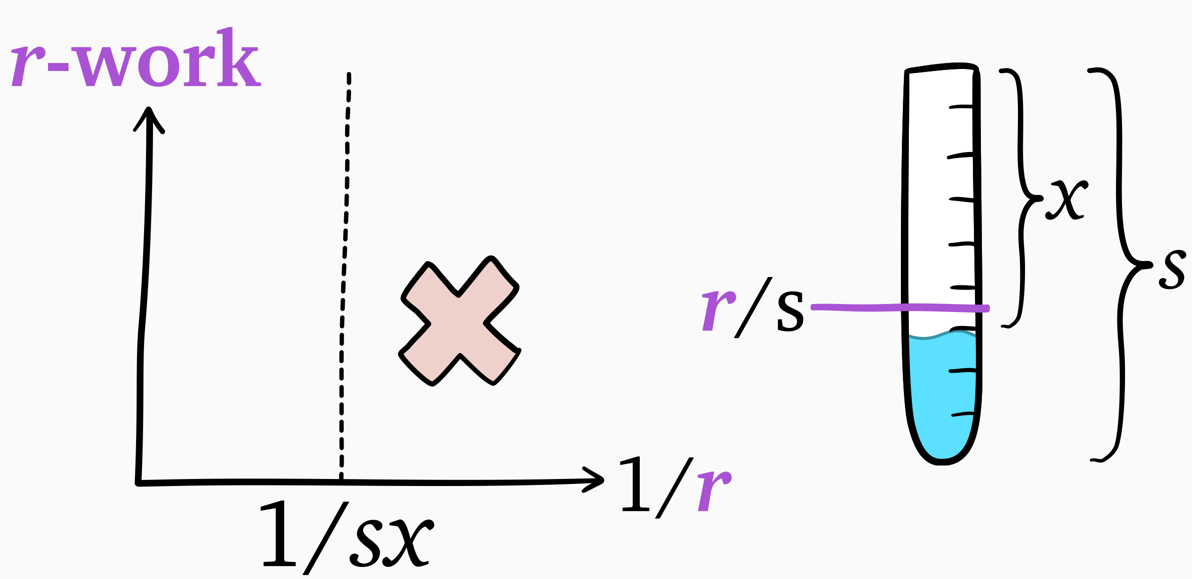

$r < sx$: r-work $= 0$

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's ***r*-work**:

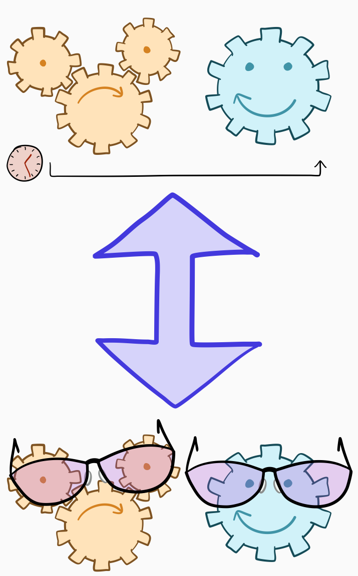

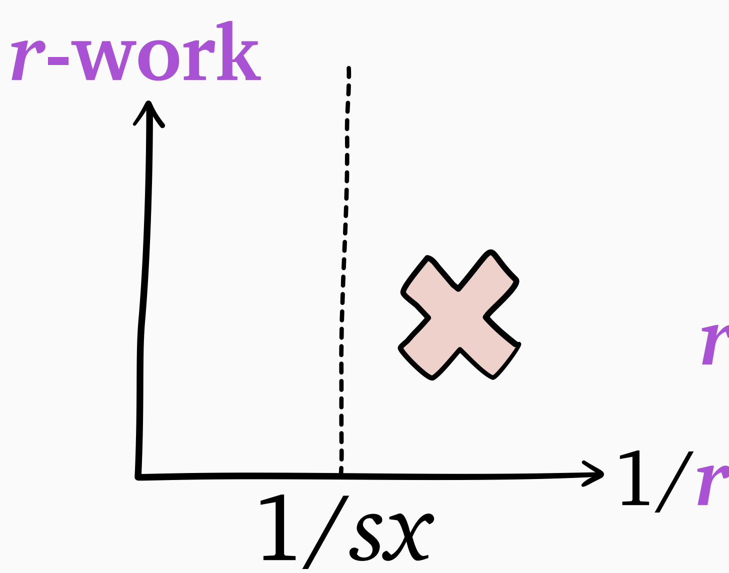

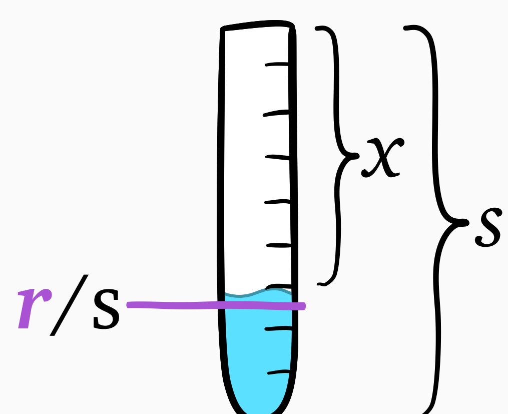

$r < sx$: r-work $= 0$

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's r-work:

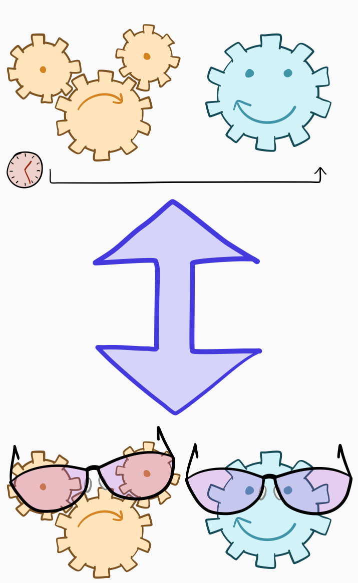

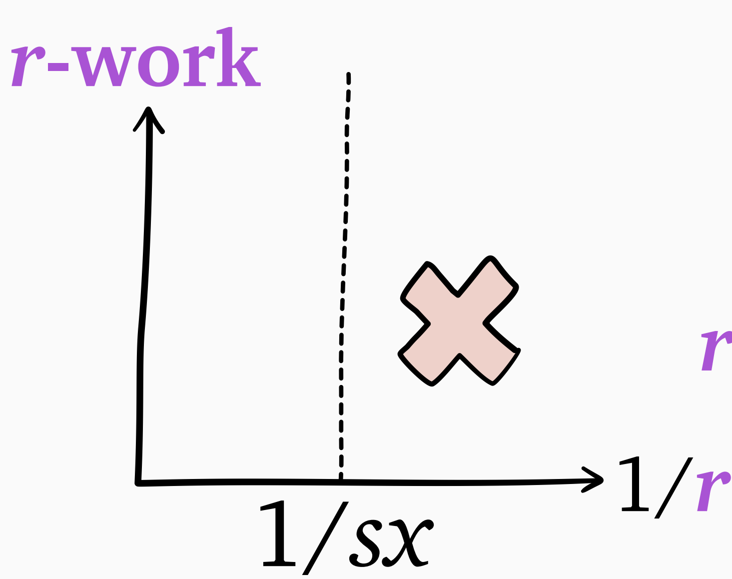

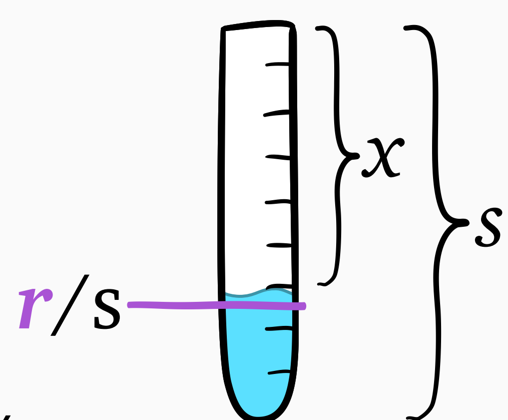

$r < sx$: r-work $= 0$

$r \geq sx$: r-work $= x$

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's **r-work**:

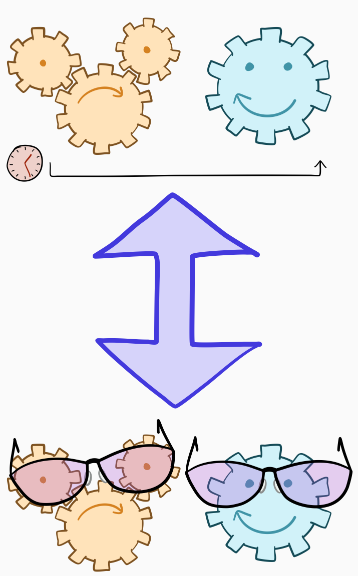

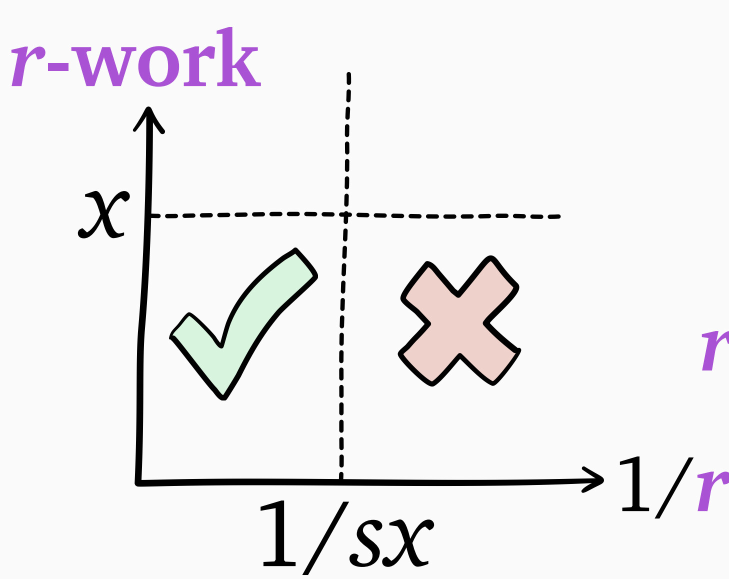

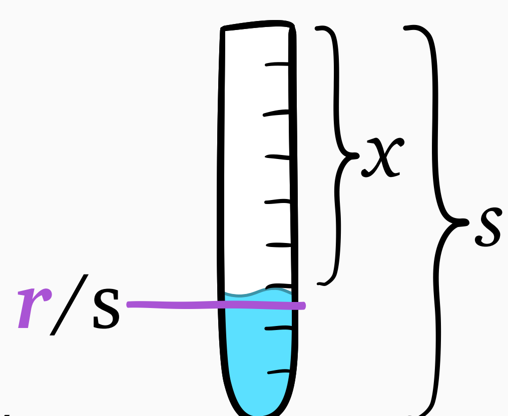

$r < sx$: **r-work** $= 0$

$r \geq sx$: **r-work** $= x$

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's r-work:

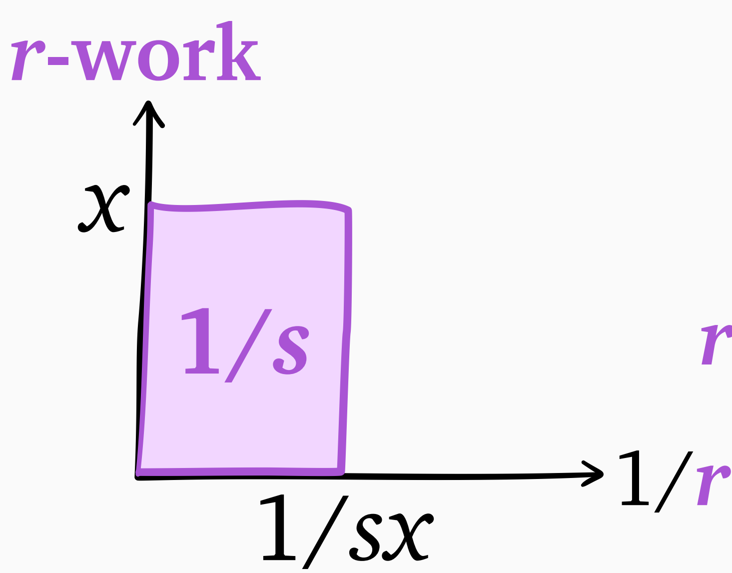

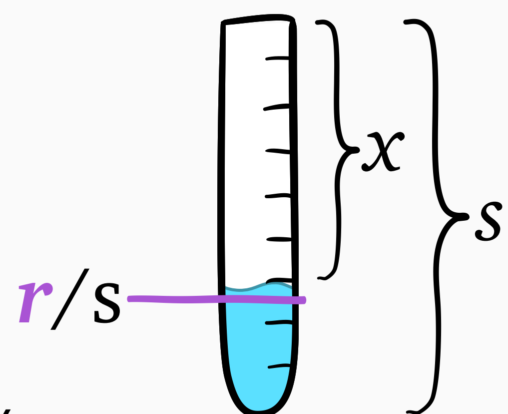

$r < sx$: r-work $= 0$

$r \geq sx$: r-work $= x$

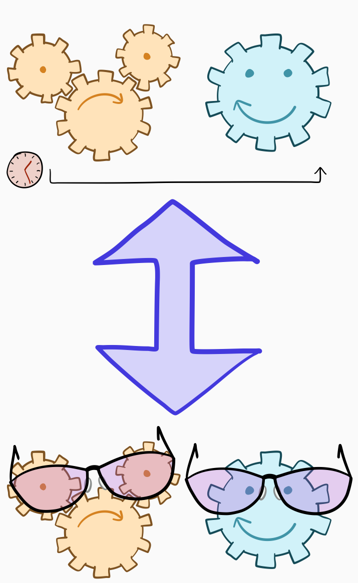

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's ***r*-work**:

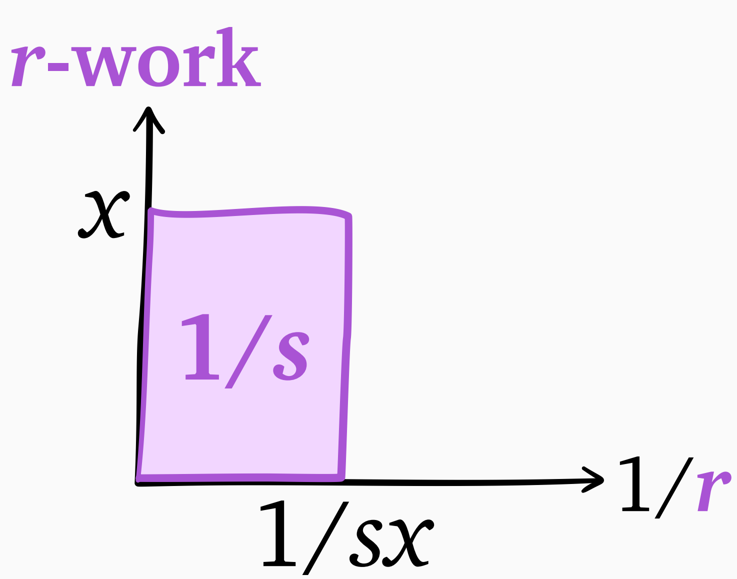

All jobs' ***r*-work**:

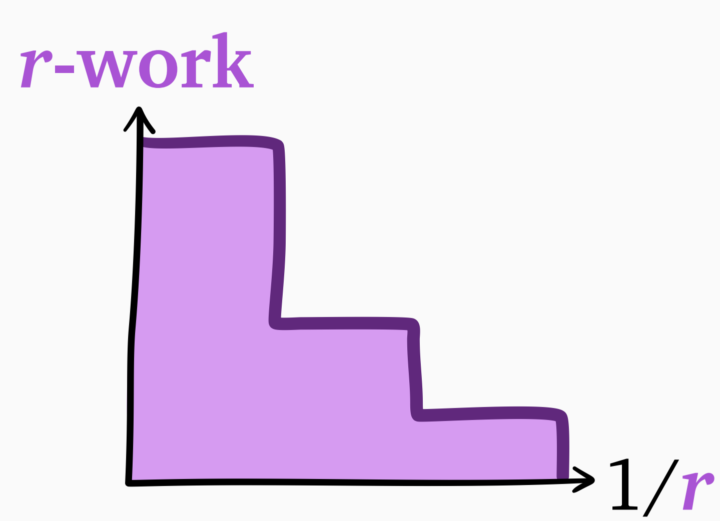

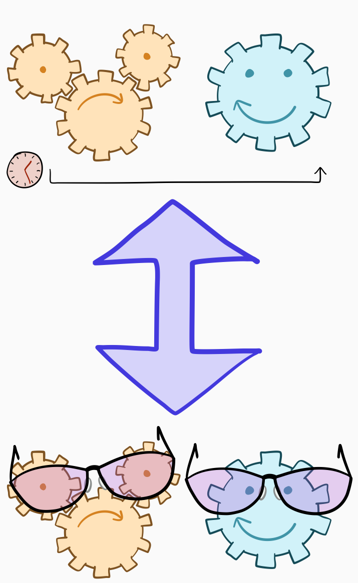

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's r-**work**:

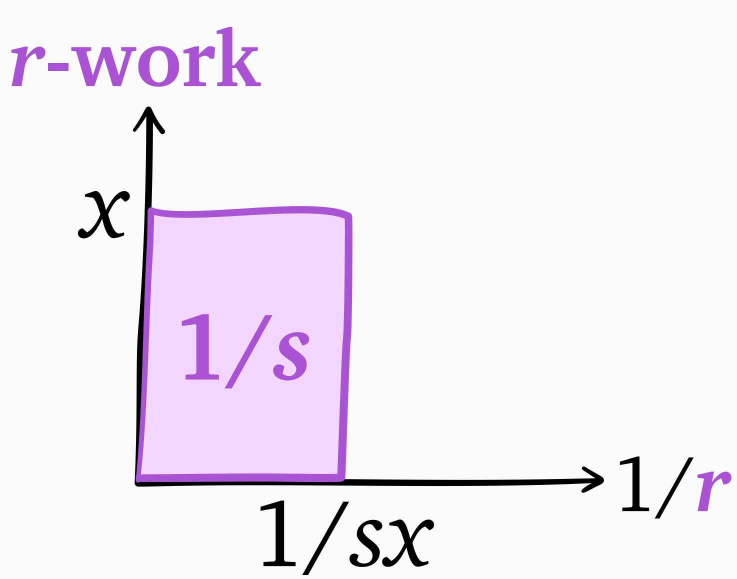

All jobs' r-**work**:

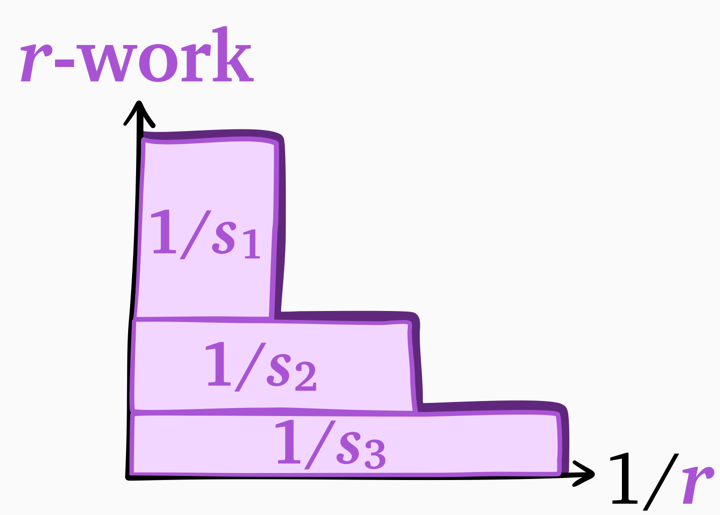

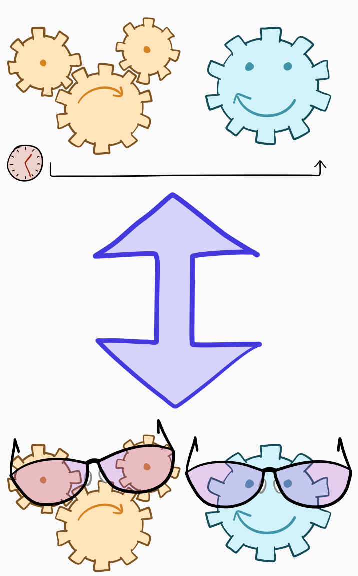

Step 1: $\mathbf{E}[H]$ to $\mathbf{E}[W(r)]$

Theorem: In basically any queueing system,

$$H = \int_0^\infty \frac{W(r)}{r^2}\,\mathrm{d}r = \int_0^\infty W(r)\,\mathrm{d}(1/r)$$

Proof:

size s,
remaining size x

One job's r-**work**:

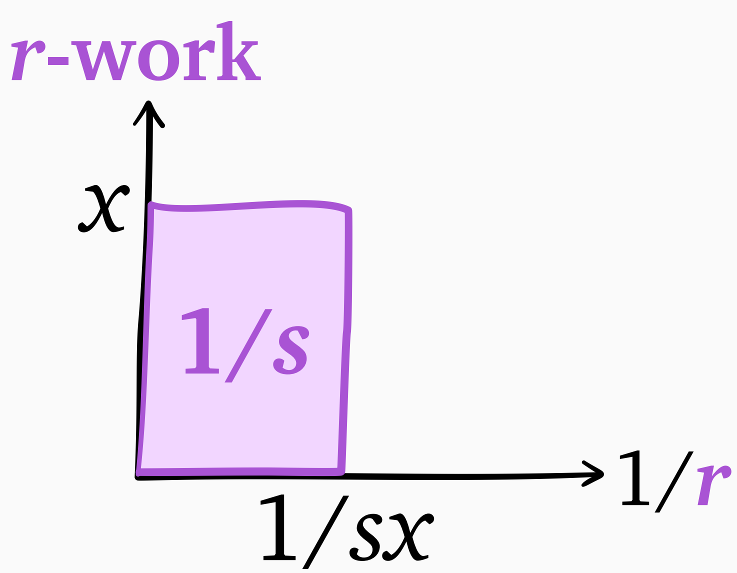

All jobs' r-**work**:

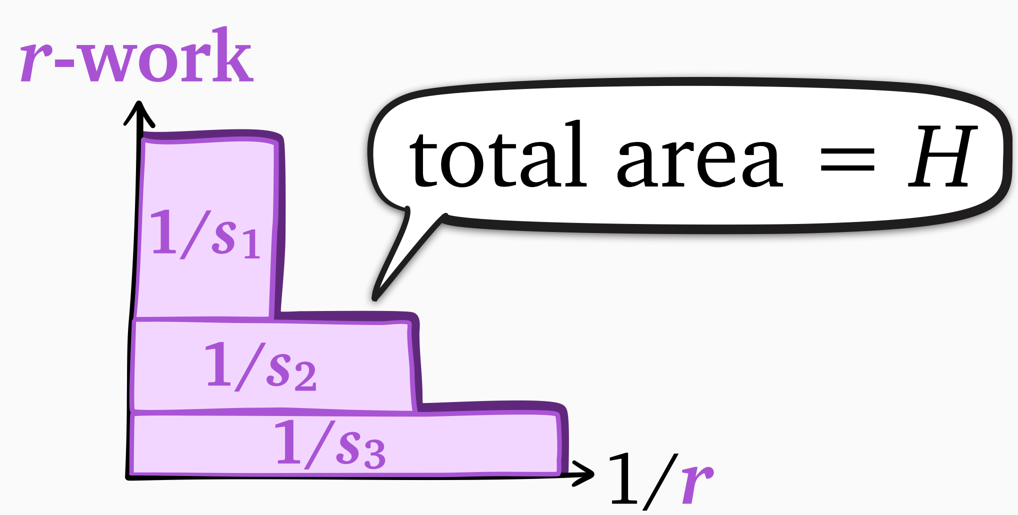

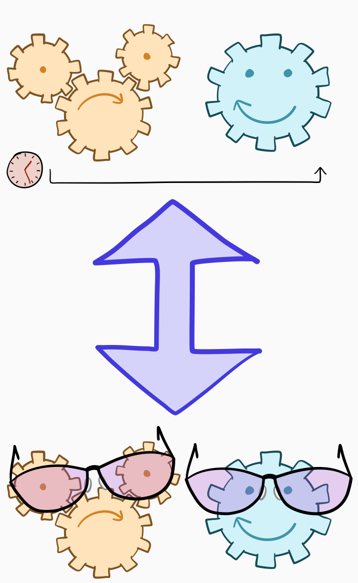

Slowdown via *r*-Work

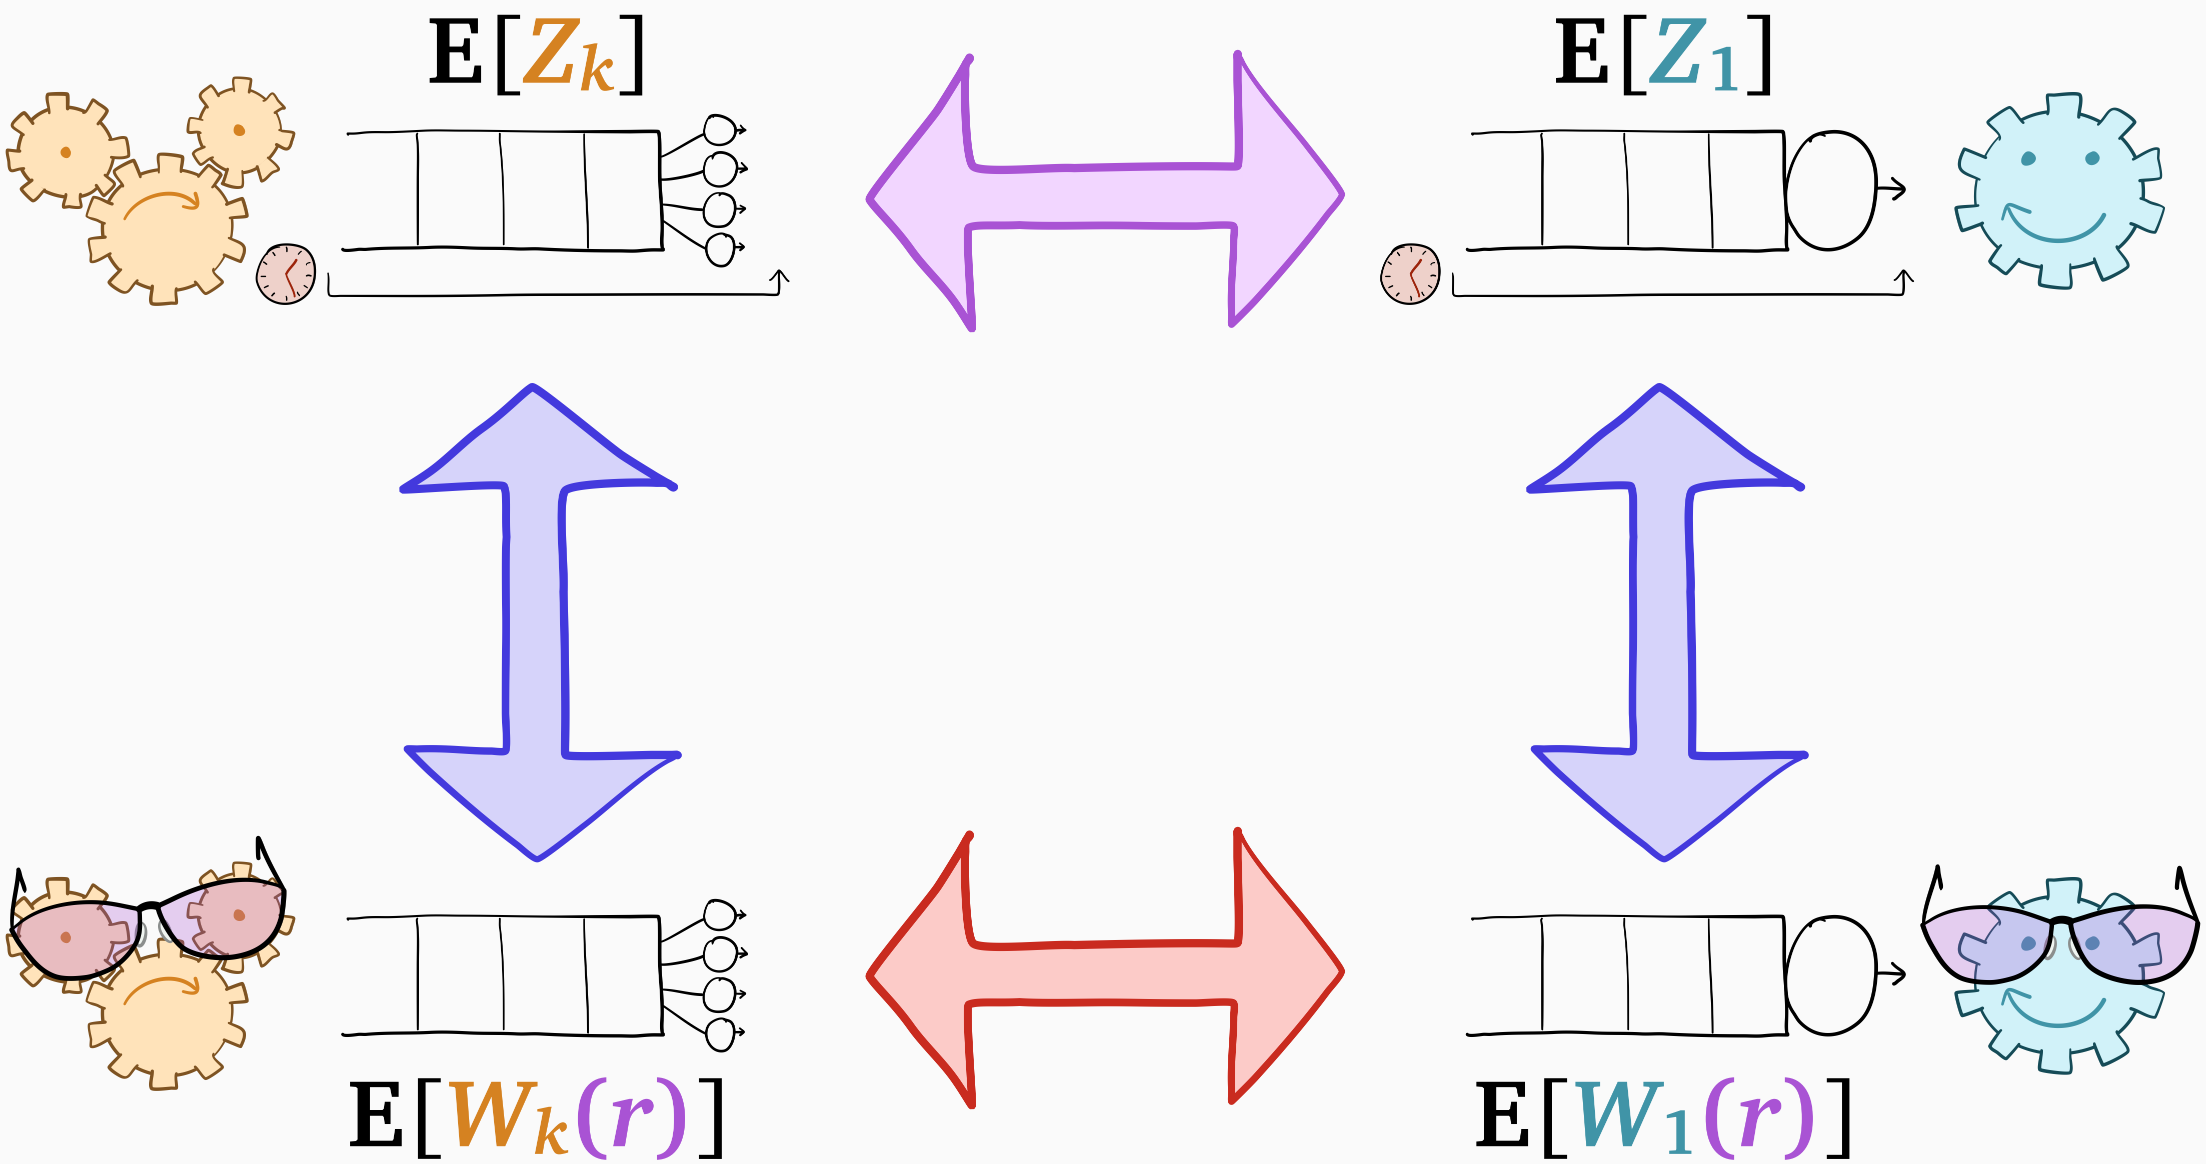

Slowdown via *r*-Work

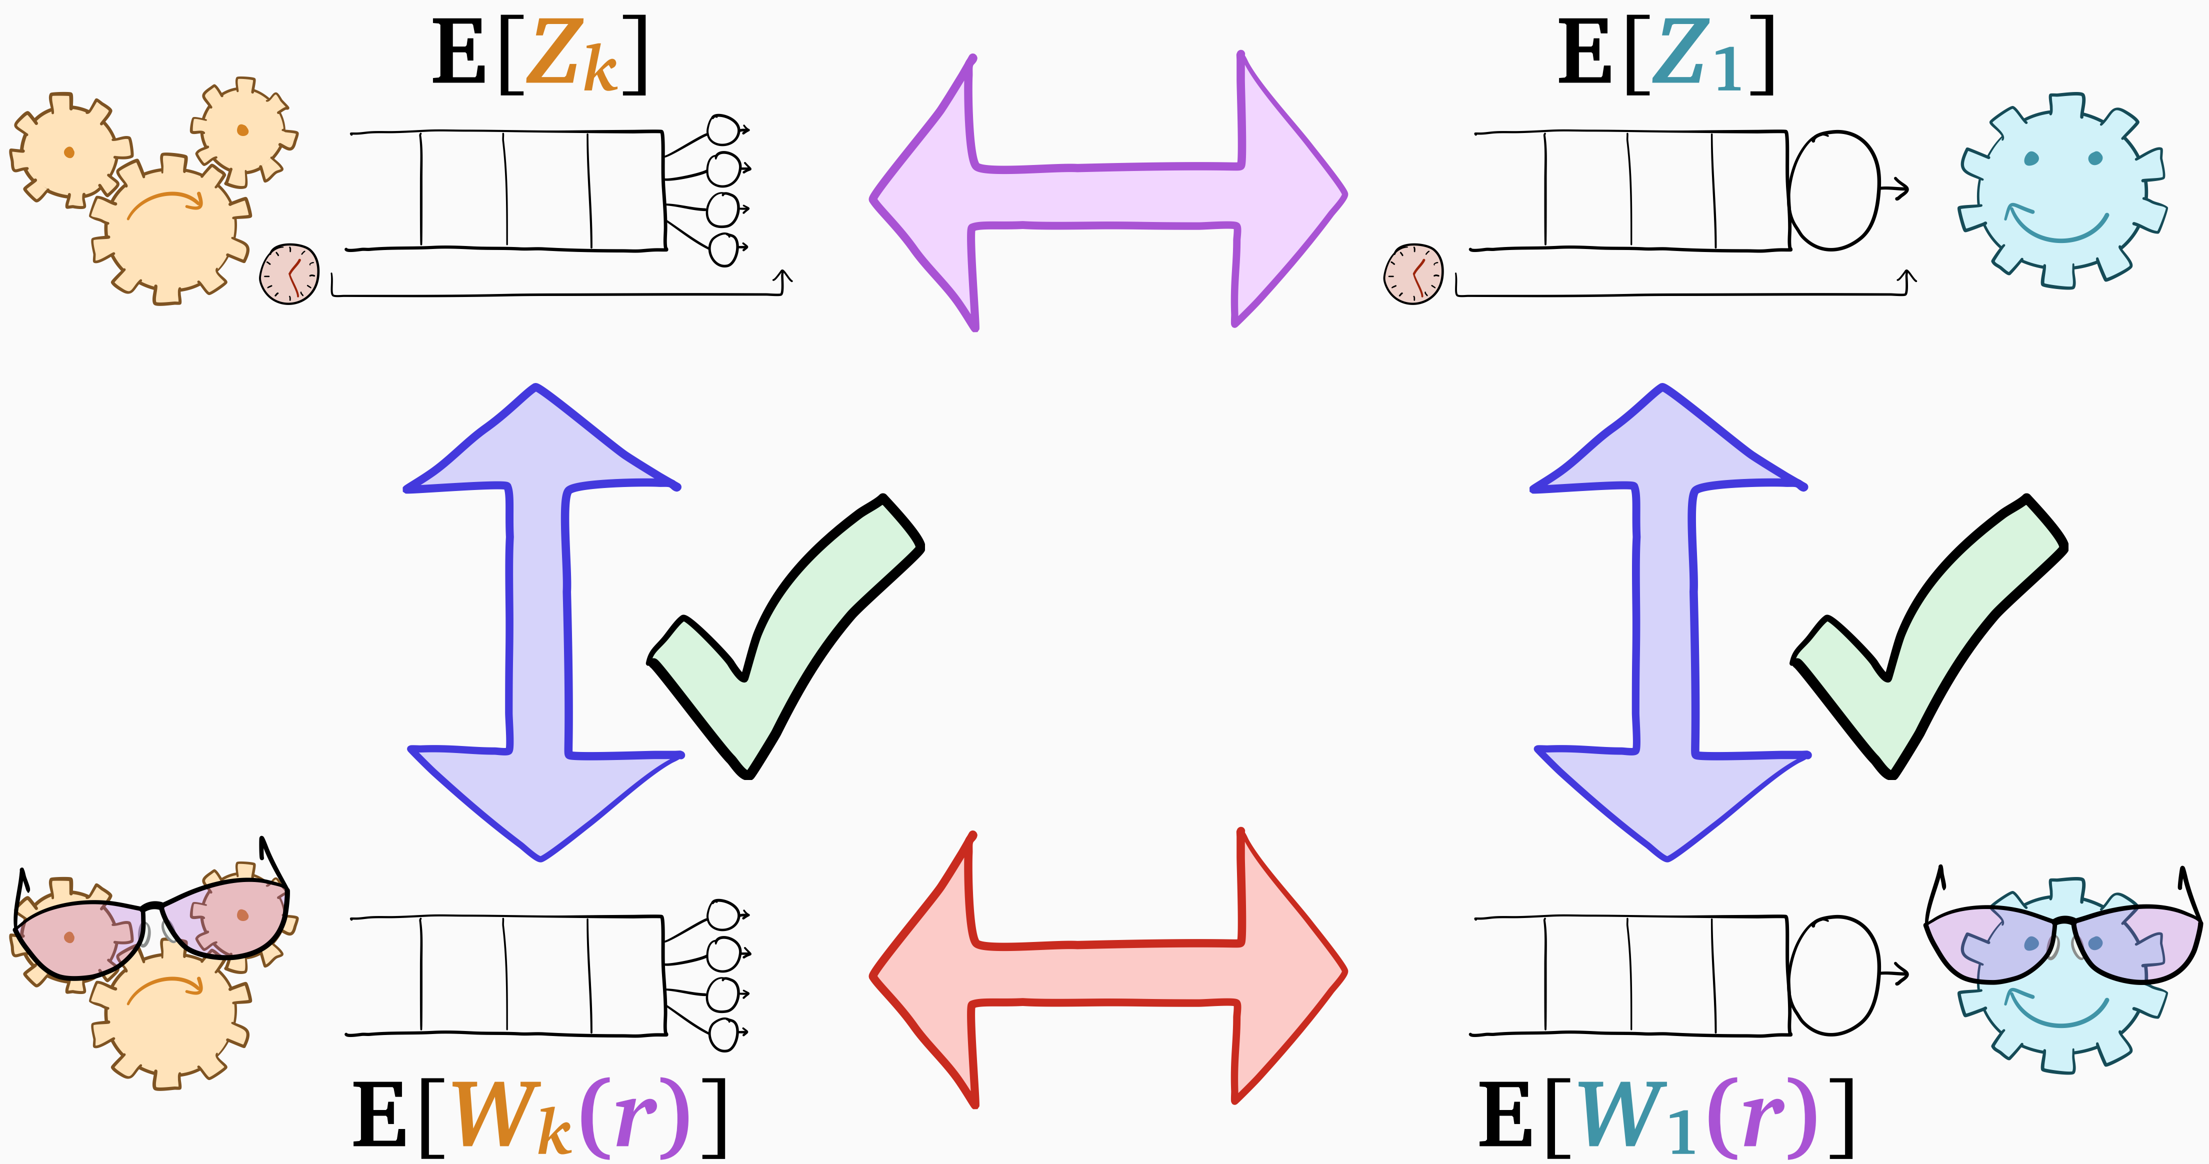

Step 2: $\mathbf{E}[W]$ Difference

Step 2: $\mathbf{E}[W]$ Difference

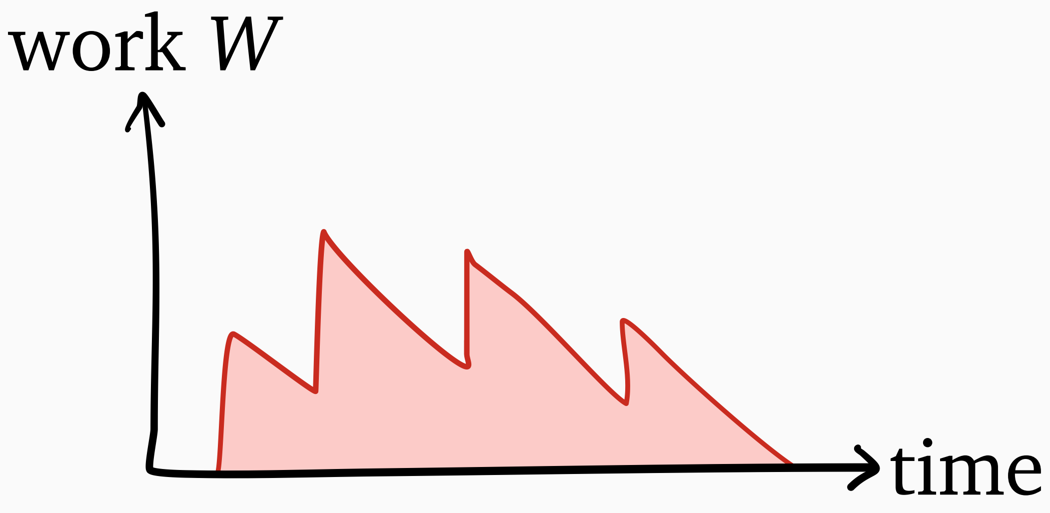

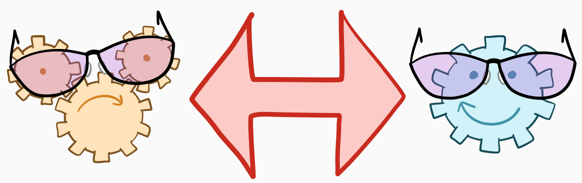

Step 2: $\mathbf{E}[W]$ Difference

In steady-state system, for any f,
$\mathbf{E}[f(W)]$ constant w.r.t. time

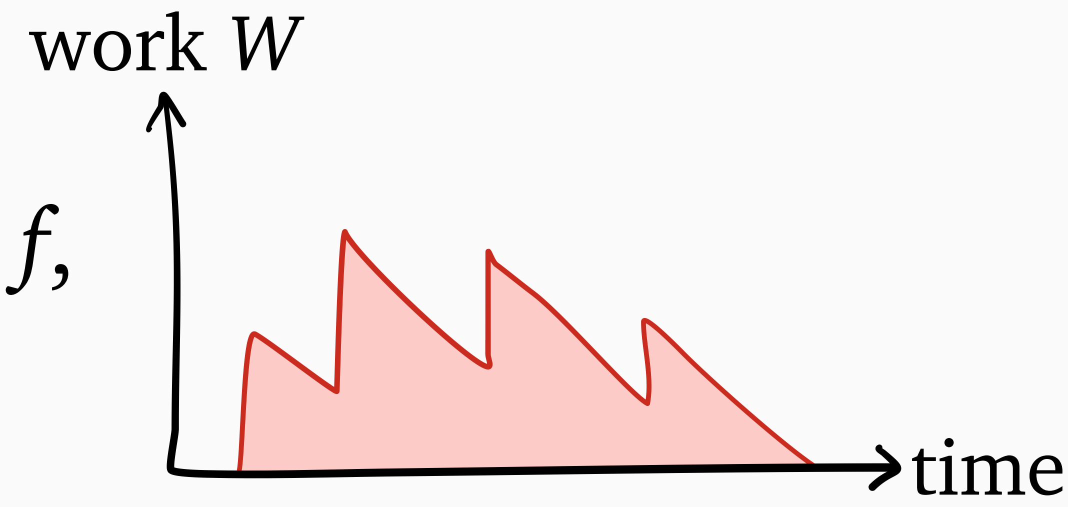

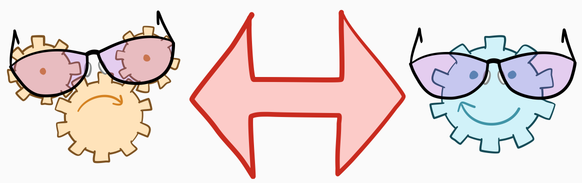

Step 2: $\mathbf{E}[W]$ Difference

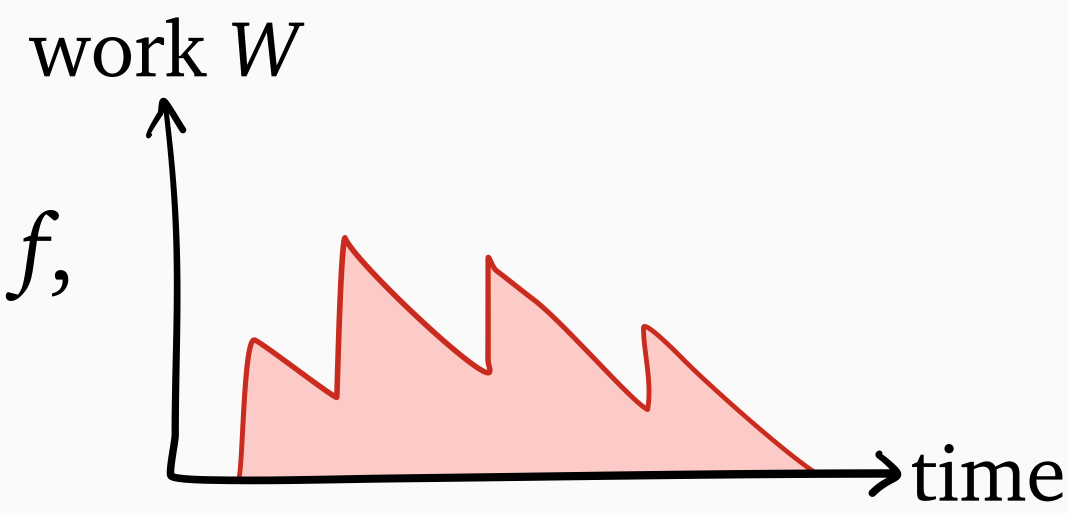

In steady-state system, for any f,
$\mathbf{E}[\,f(W)]$ constant w.r.t. time

we use $f(w) = w^2$

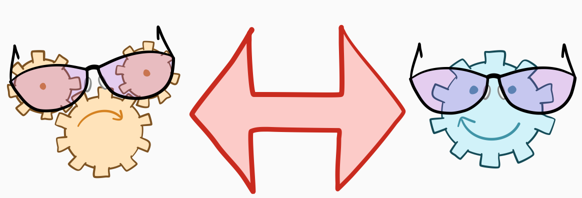

Step 2: $\mathbf{E}[W]$ Difference

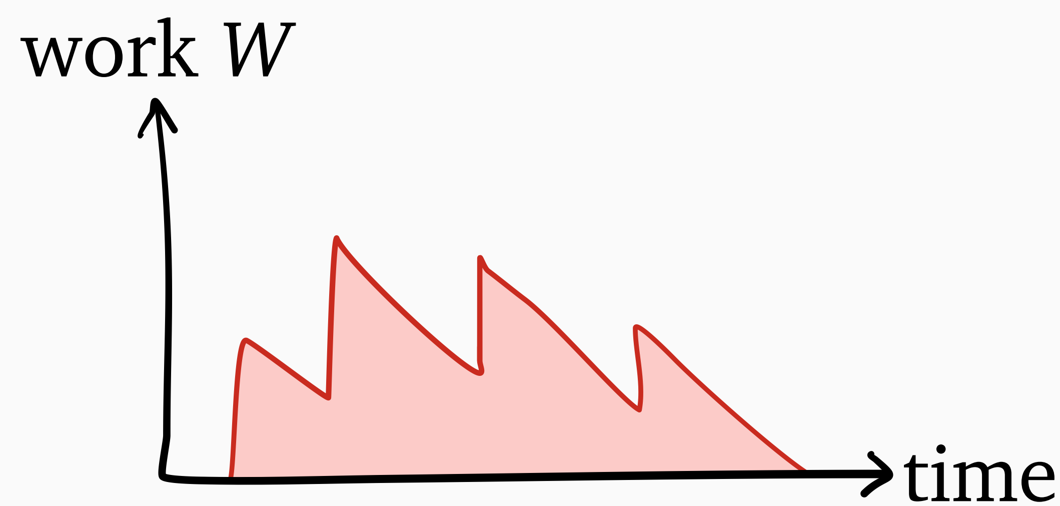

$$\mathbf{E}[W^2 \text{ decrease rate}] = 2\mathbf{E}[BW]$$

$$\mathbf{E}[W^2 \text{ increase rate}] = \lambda\mathbf{E}[(W+S)^2 - W^2]$$

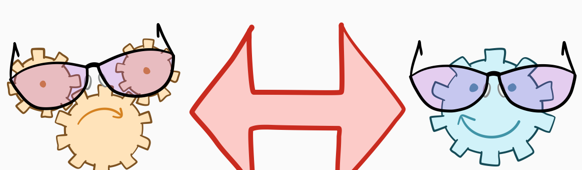

Step 2: $\mathbf{E}[W]$ Difference

B = service rate, a.k.a. fraction of servers busy

$$\mathbf{E}[W^2 \text{ decrease rate}] = 2\mathbf{E}[BW]$$

$$\mathbf{E}[W^2 \text{ increase rate}] = \lambda\mathbf{E}[(W+S)^2 - W^2]$$

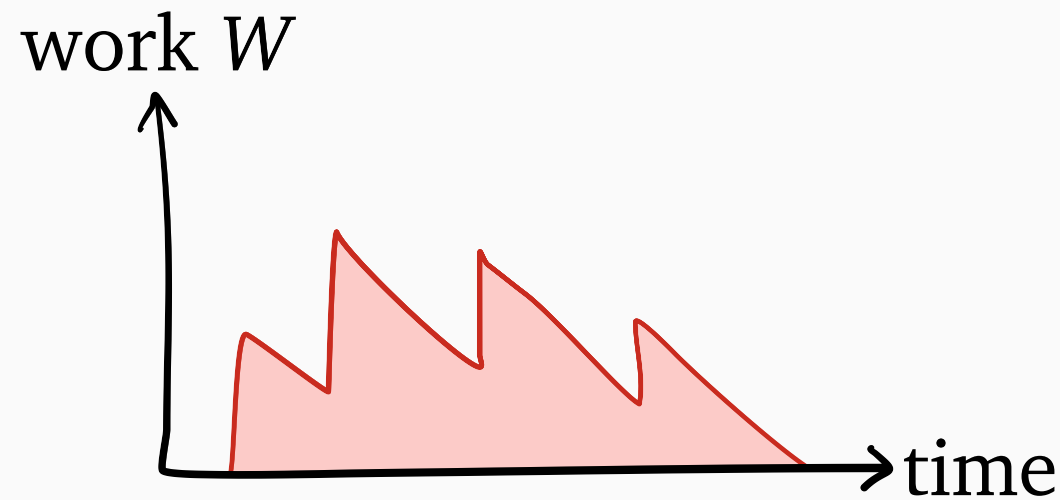

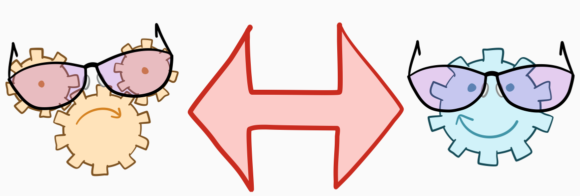

Step 2: $\mathbf{E}[W]$ Difference

B = service rate, a.k.a. fraction of servers busy

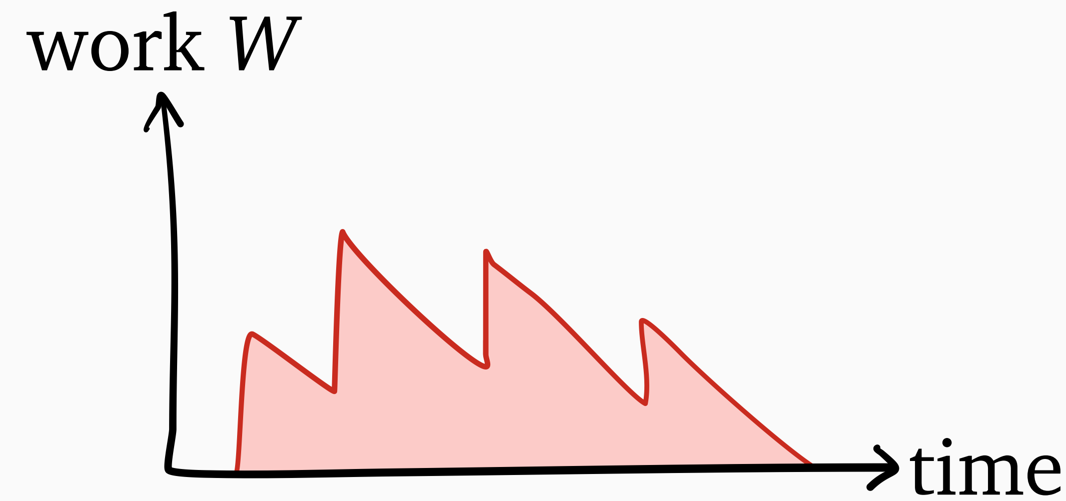

$$\mathbf{E}[W^2 \text{ decrease rate}] = 2\mathbf{E}[BW]$$

$$\mathbf{E}[W^2 \text{ increase rate}] = \lambda\mathbf{E}[(W+S)^2 - W^2]$$

$$\mathbf{E}[W] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2]}{1-\rho} + \frac{\mathbf{E}[(1-B)W]}{1-\rho}$$

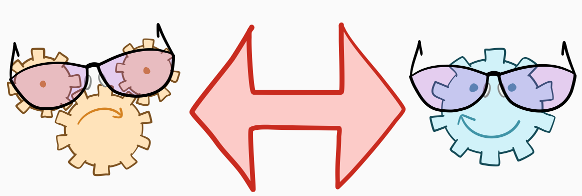

Step 2: $\mathbf{E}[W]$ Difference

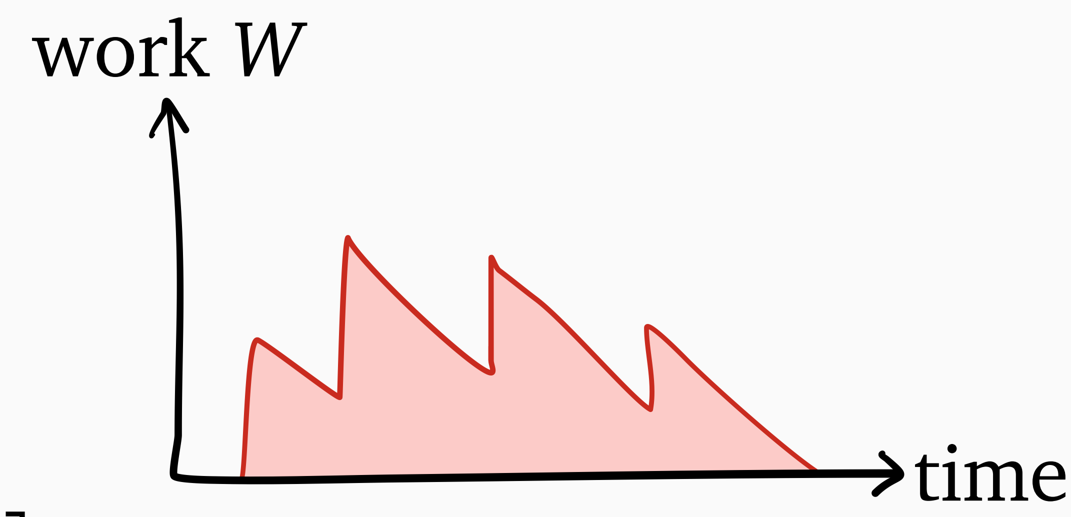

B = service rate, a.k.a. fraction of servers busy

$$\mathbf{E}[W^2 \text{ decrease rate}] = 2\mathbf{E}[BW]$$

$$\mathbf{E}[W^2 \text{ increase rate}] = \lambda\mathbf{E}[(W+S)^2 - W^2]$$

M/G/1: $(1 - B)W = 0$

$$\mathbf{E}[W] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2]}{1-\rho} + \frac{\mathbf{E}[(1-B)W]}{1-\rho}$$

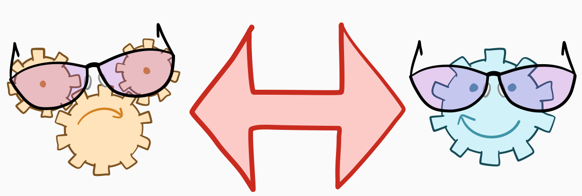

Step 2: $\mathbf{E}[W]$ Difference

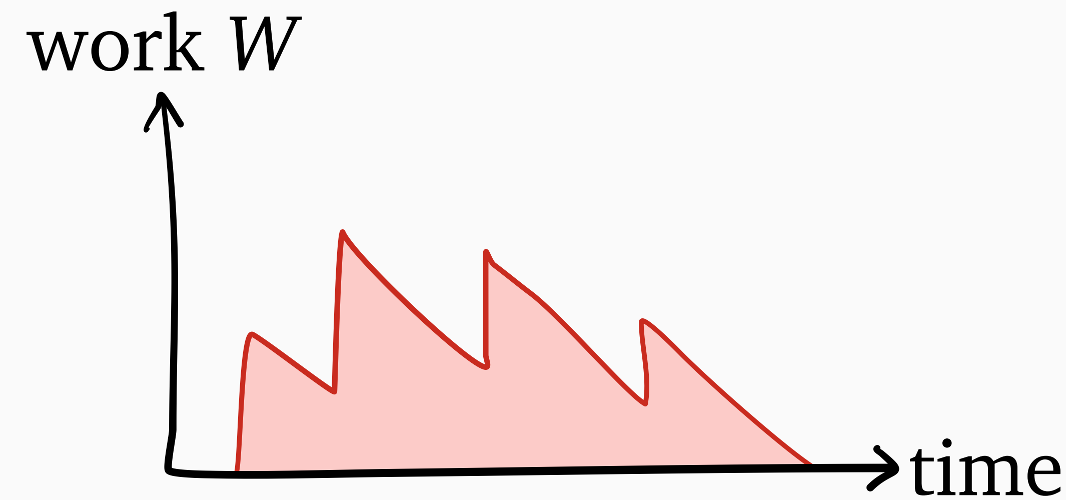

B = service rate, a.k.a. fraction of servers busy

$$\mathbf{E}[W^2 \text{ decrease rate}] = 2\mathbf{E}[BW]$$

$$\mathbf{E}[W^2 \text{ increase rate}] = \lambda\mathbf{E}[(W+S)^2 - W^2]$$

M/G/1: $(1 - B)W = 0$

$$\mathbf{E}[W] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2]}{1-\rho} + \frac{\mathbf{E}[(1-B)W]}{1-\rho}$$

Theorem:

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

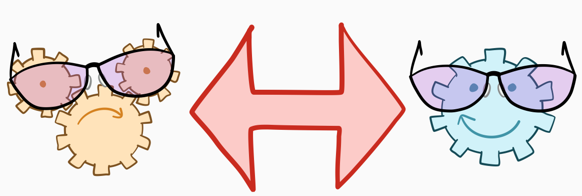

Step 2: $\mathbf{E}[W]$ Difference

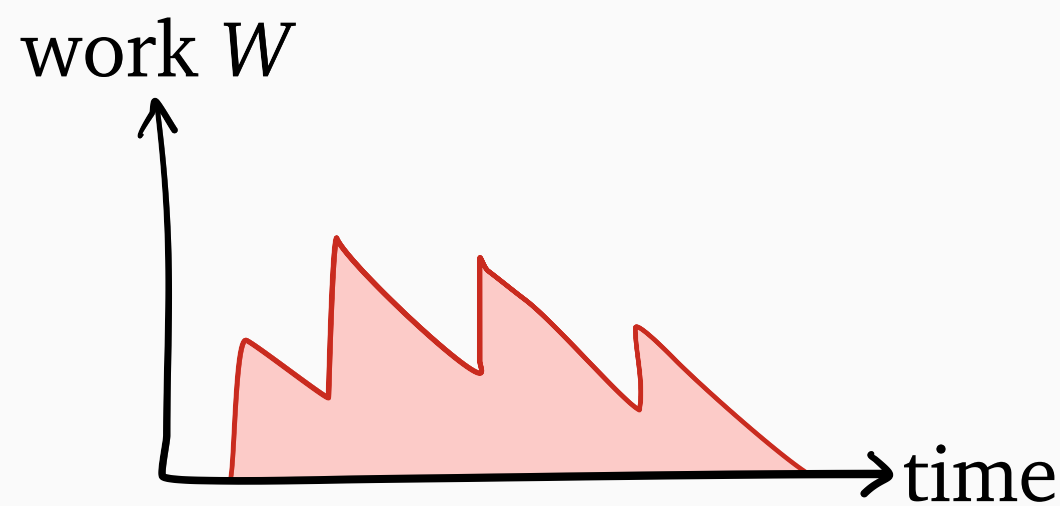

B = service rate, a.k.a. fraction of servers busy

$$\mathbf{E}[W^2 \text{ decrease rate}] = 2\mathbf{E}[BW]$$

$$\mathbf{E}[W^2 \text{ increase rate}] = \lambda \mathbf{E}[(W+S)^2 - W^2]$$

M/G/1: $(1 - B)W = 0$

$$\mathbf{E}[W] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2]}{1-\rho} + \frac{\mathbf{E}[(1-B)W]}{1-\rho}$$

Theorem:

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

(similar holds for ***r*-work**)

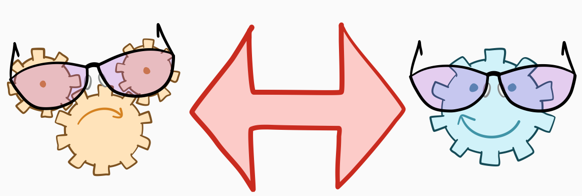

Step 2: Bound in M/G/k

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1 - B_k)W_k]}{1 - \rho}$$

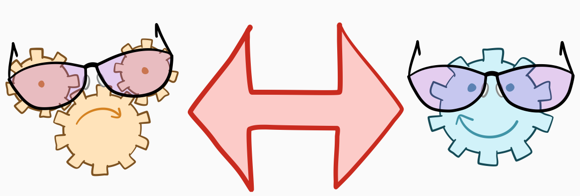

Step 2: Bound in M/G/k

Suppose $S \leq s_{\max}$ with probability 1

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1 - B_k)W_k]}{1 - \rho}$$

Step 2: Bound in M/G/k

Suppose $S \leq s_{\max}$ with probability 1

$\leq (k-1)s_{\max}$

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

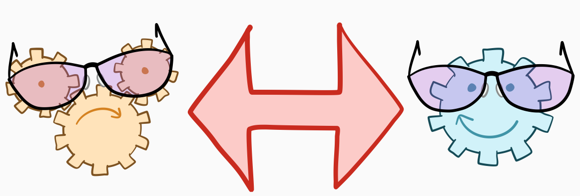

Step 2: Bound in M/G/k

Suppose $S \le s_{\max}$ with probability 1

$\mathbf{E}[B] = \rho$

$\le (k-1)s_{\max}$

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1 - B_k)W_k]}{1 - \rho}$$

Step 2: Bound in M/G/k

Suppose $S \le s_{\max}$ with probability 1

$\mathbf{E}[B] = \rho$

$\le (k-1)s_{\max}$

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

$$\le \mathbf{E}[W_1] + (k-1)s_{\max}$$

Step 2: Bound in M/G/k

Suppose $S \le s_{\max}$ with probability 1

$\mathbf{E}[B] = \rho$

$\le (k-1)s_{\max}$

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

$$\le \mathbf{E}[W_1] + (k-1)s_{\max}$$

“work of $\le k-1$ jobs”

Step 2: Bound in M/G/k

Suppose $S \leq s_{\max}$ with probability 1

$\mathbf{E}[B] = \rho$

$\leq (k-1)s_{\max}$

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

$$\leq \mathbf{E}[W_1] + (k-1)s_{\max}$$

"work of $\leq k-1$ jobs"

$\mathbf{E}[W_k(r)] = \mathbf{E}[W_1(r)] +$ "r-**work** of $\leq k-1$ jobs"

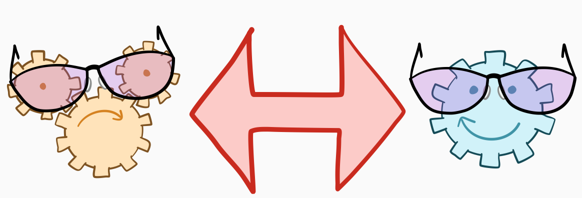

Step 2: Bound in M/G/k

Suppose $S \le s_{\max}$ with probability 1

$\mathbf{E}[B] = \rho$

$\le (k-1)s_{\max}$

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

$$\le \mathbf{E}[W_1] + (k-1)s_{\max}$$

"work of $\le k-1$ jobs"

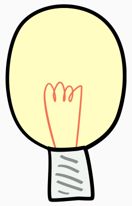

Single job's r-**work** is at most $\sqrt{r}$

$$\mathbf{E}[W_k(r)] = \mathbf{E}[W_1(r)] + \text{“}r\text{-work of } \le k-1 \text{ jobs”}$$

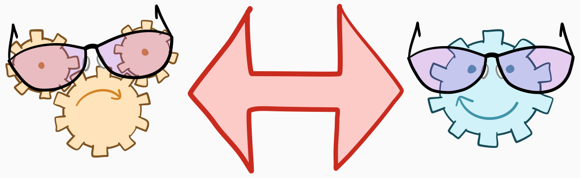

Step 2: Bound in M/G/k

Suppose $S \leq s_{\max}$ with probability 1

$\mathbf{E}[B] = \rho$

$\leq (k-1)s_{\max}$

$$\mathbf{E}[W_k] = \mathbf{E}[W_1] + \frac{\mathbf{E}[(1-B_k)W_k]}{1-\rho}$$

$$\leq \mathbf{E}[W_1] + (k-1)s_{\max}$$

"work of $\leq k-1$ jobs"

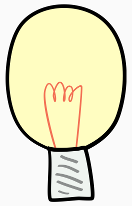

Single job's r-work is at most $\sqrt{r}$

$$\mathbf{E}[W_k(r)] = \mathbf{E}[W_1(r)] + \text{"}r\text{-work of} \leq k-1 \text{ jobs"}$$

$$\leq \mathbf{E}[W_1(r)] + (k-1)\sqrt{r}$$

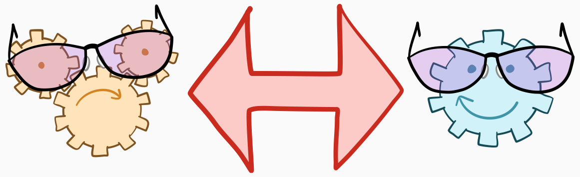

Slowdown via *r*-Work

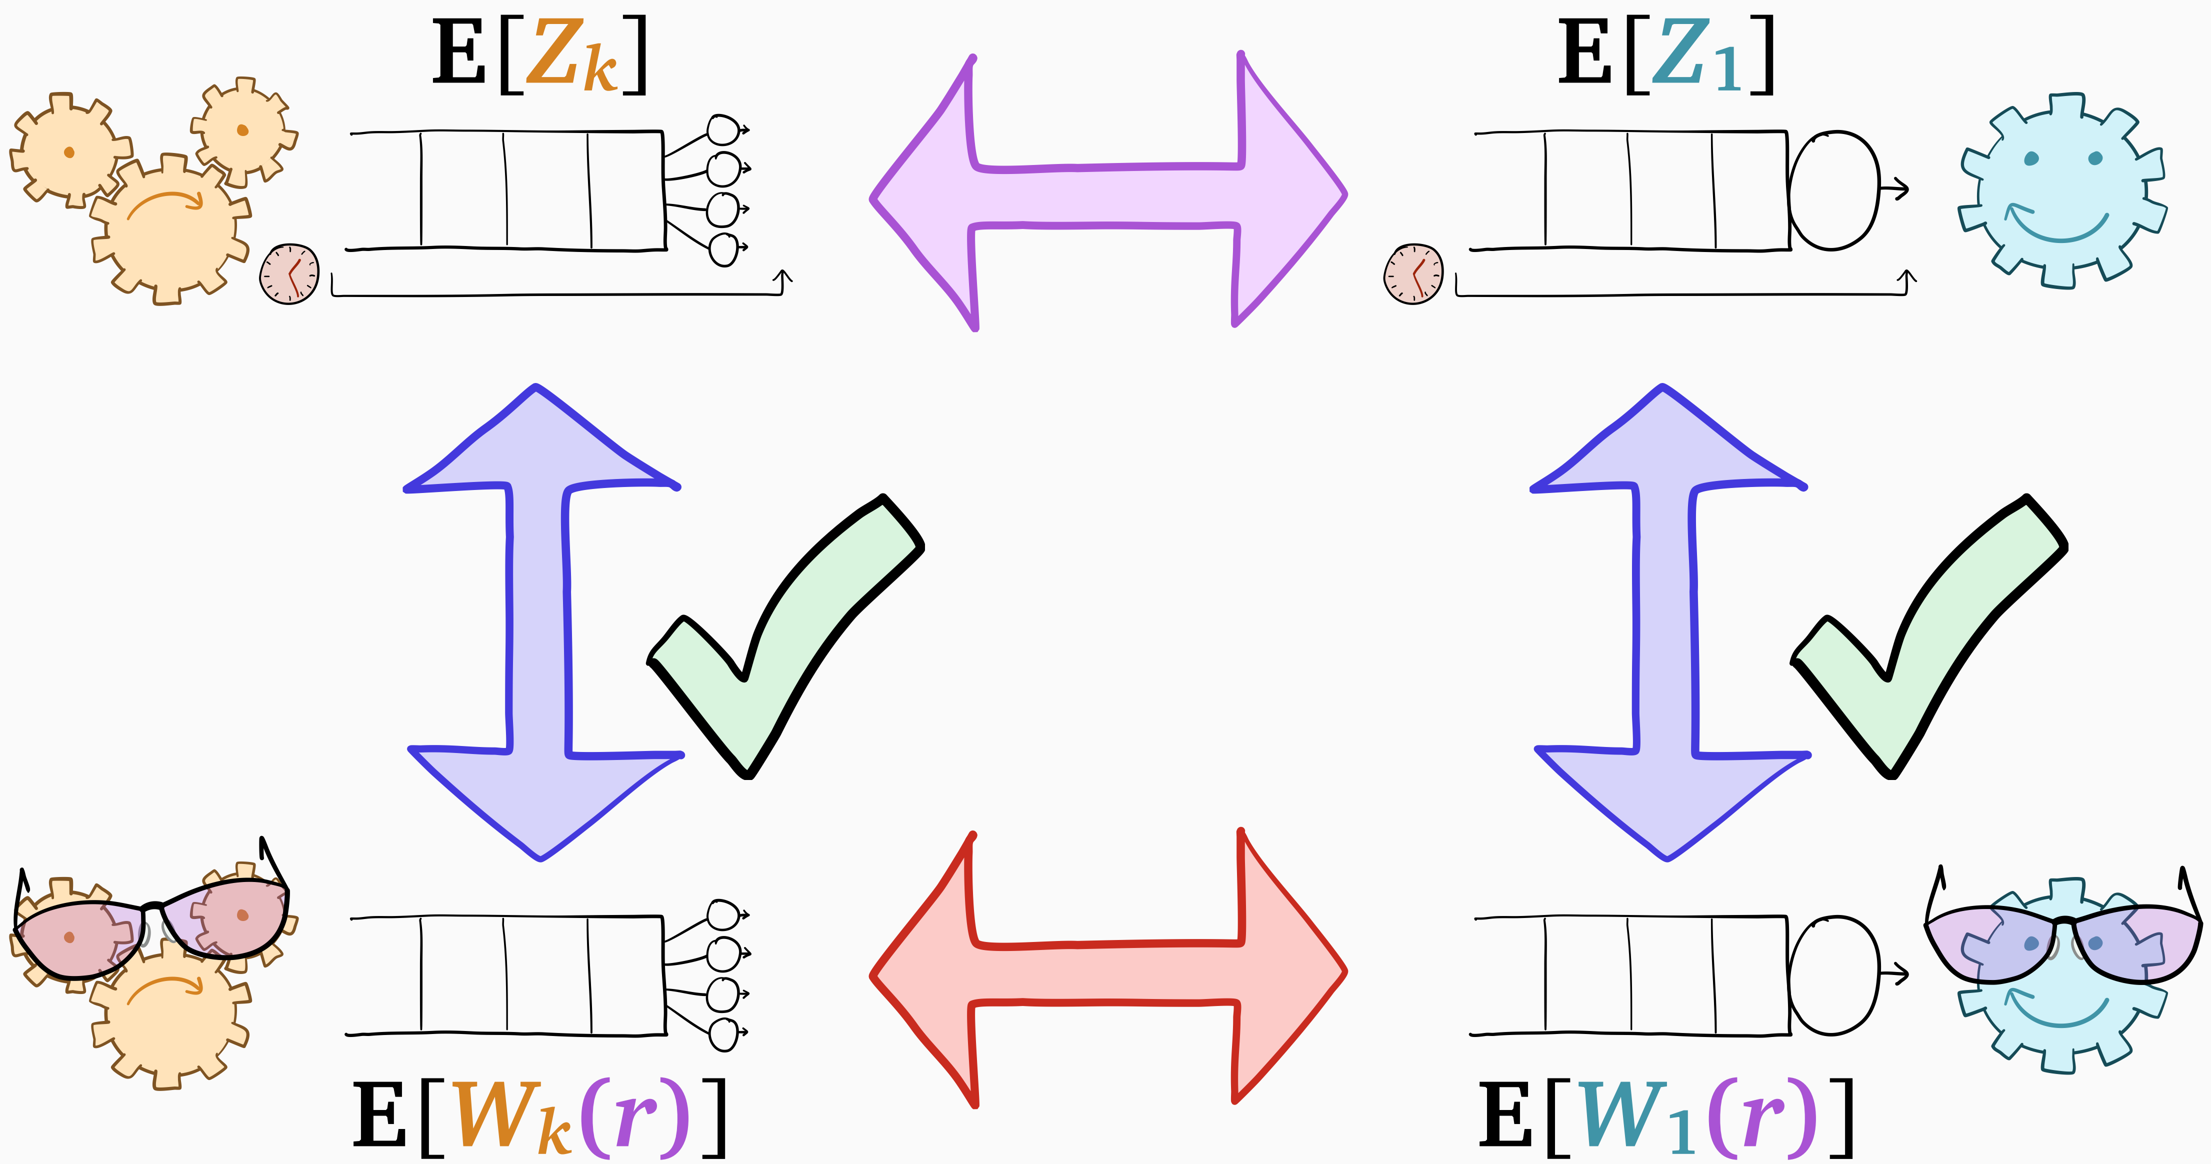

Slowdown via *r*-Work

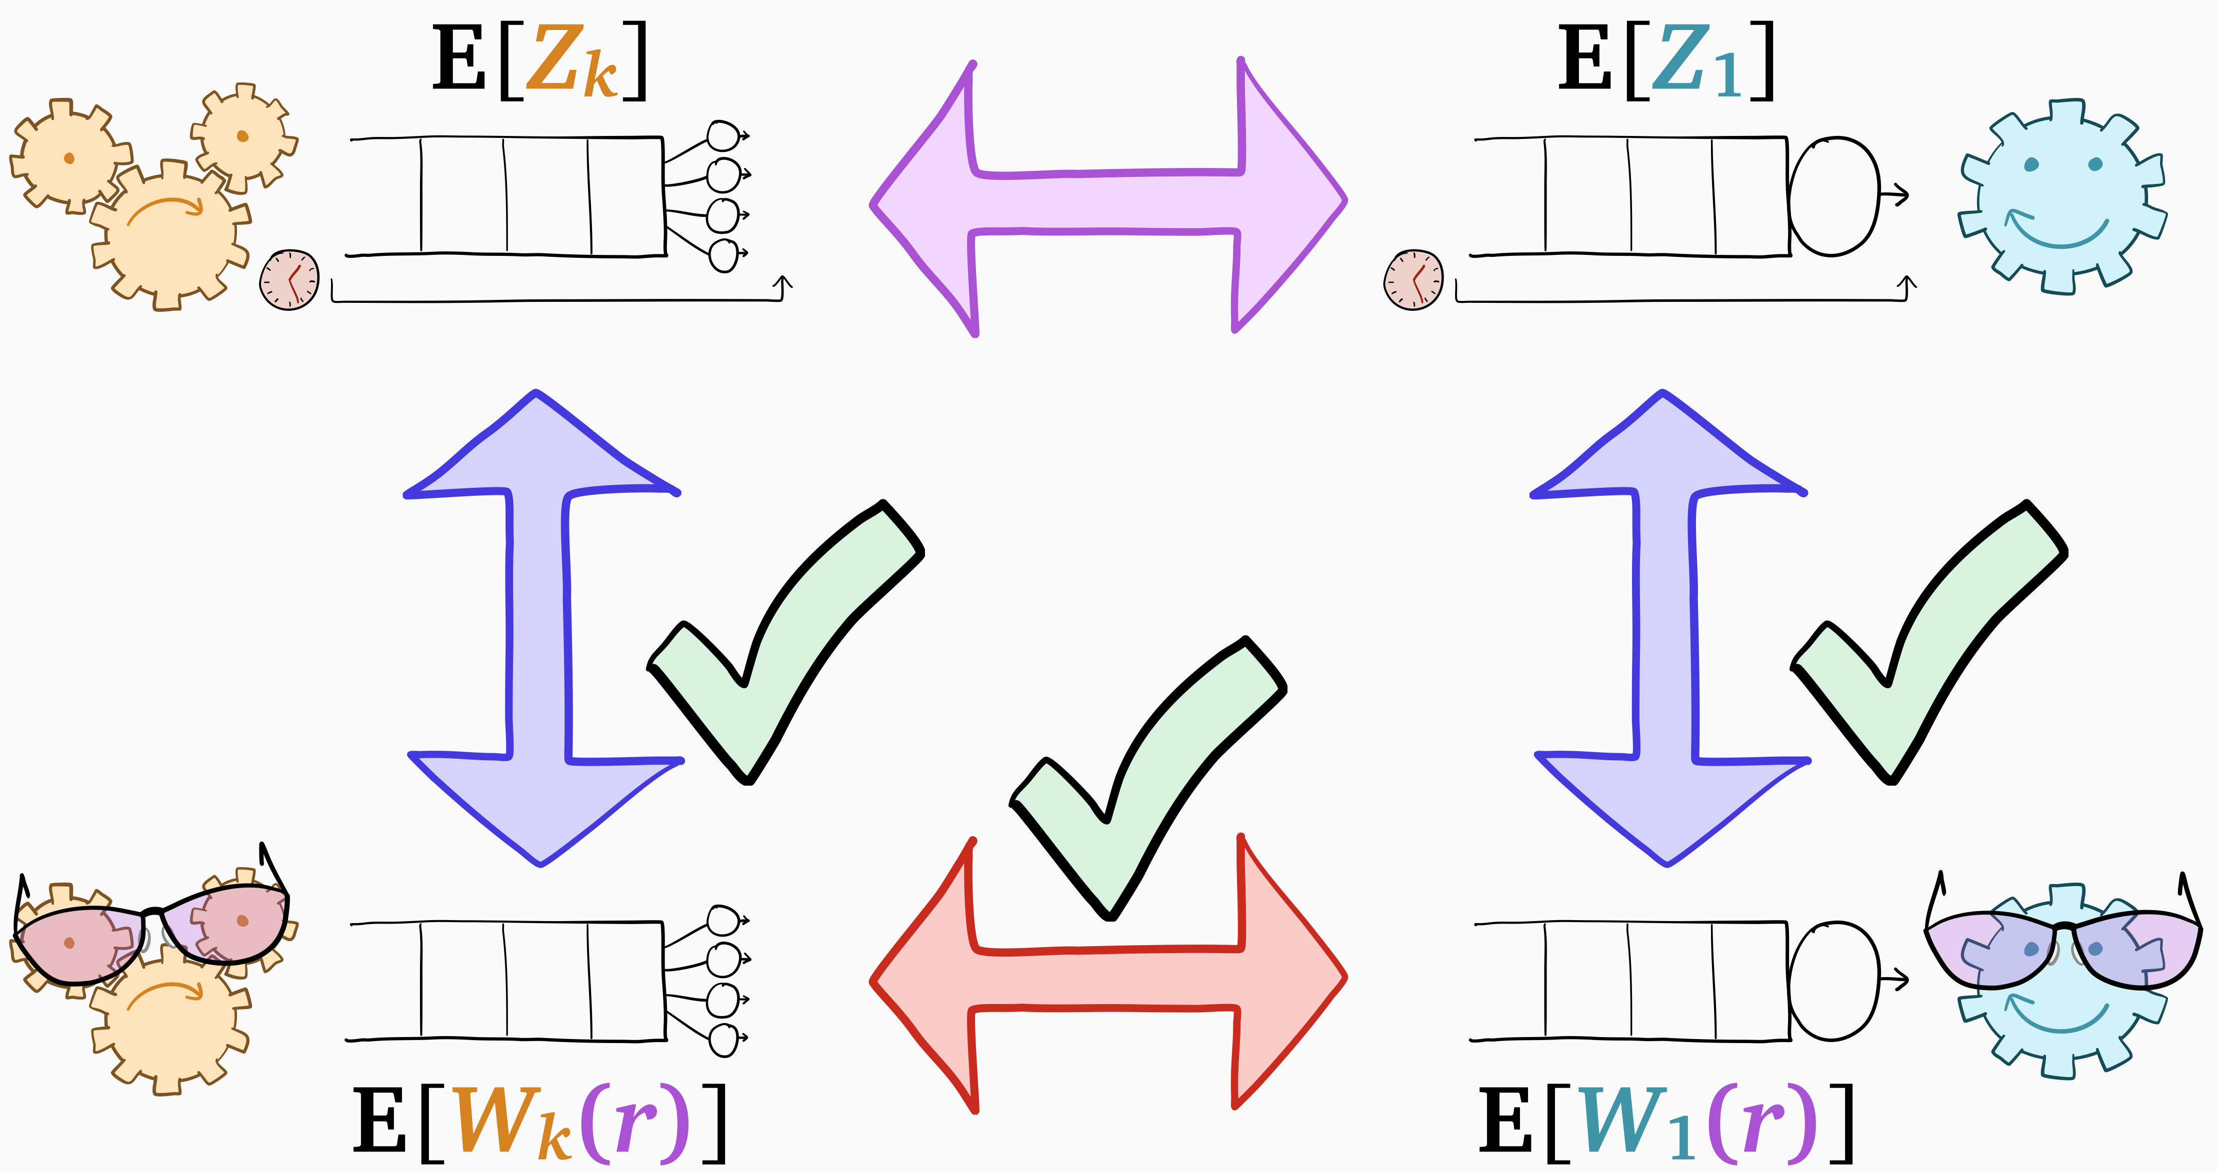

Slowdown via r-Work

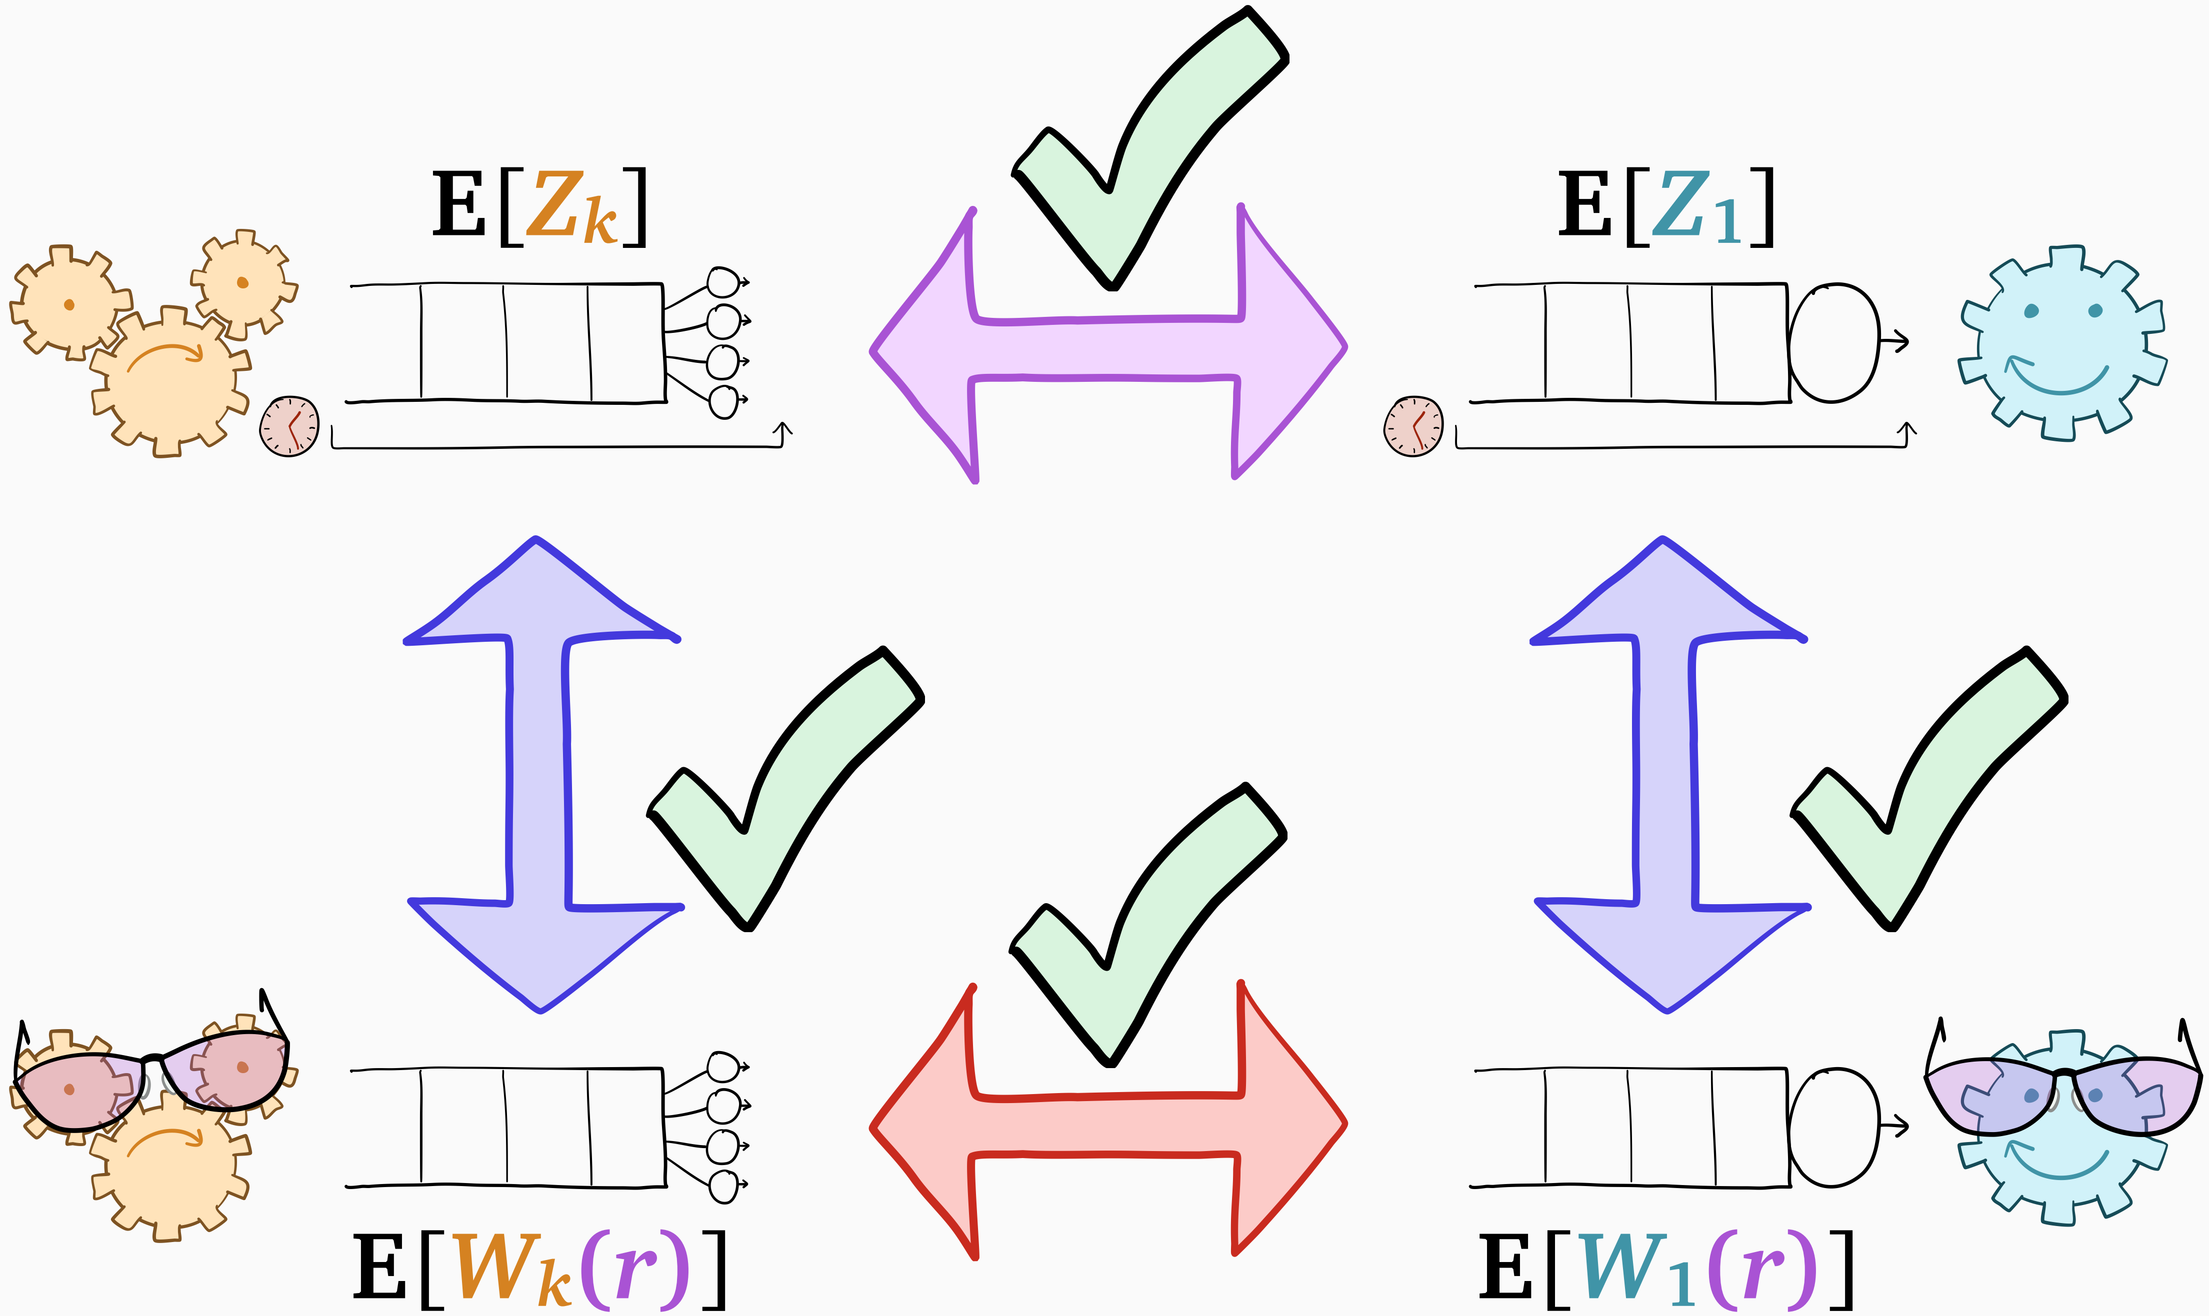

Summary

Minimize mean slowdown
in multiserver systems

Summary

Minimize mean slowdown in multiserver systems

Multiserver systems very hard to analyze directly

Summary

Minimize mean slowdown in multiserver systems

Multiserver systems very hard to analyze directly

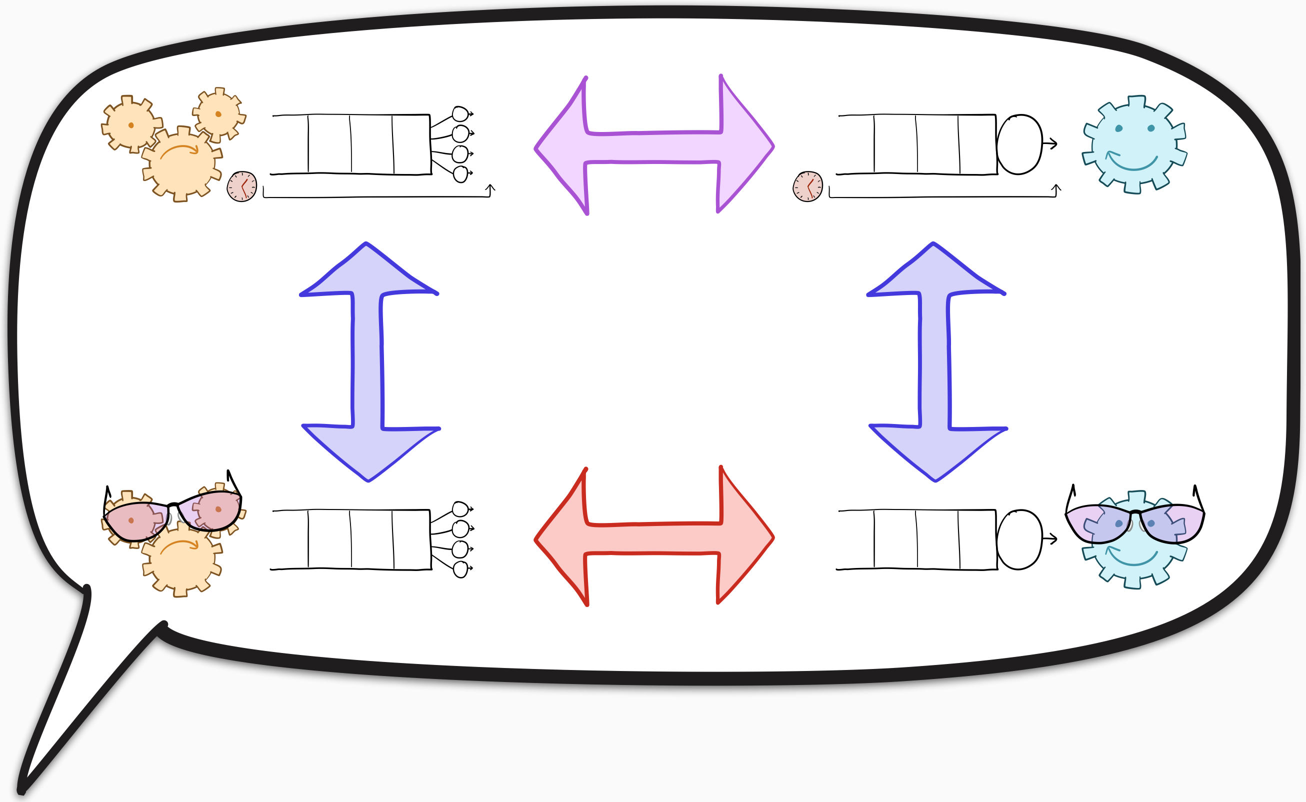

New technique based on relating $\mathbf{E}[Z]$ to ***r*-work**

Summary

Minimize mean slowdown in multiserver systems

Multiserver systems very hard to analyze directly

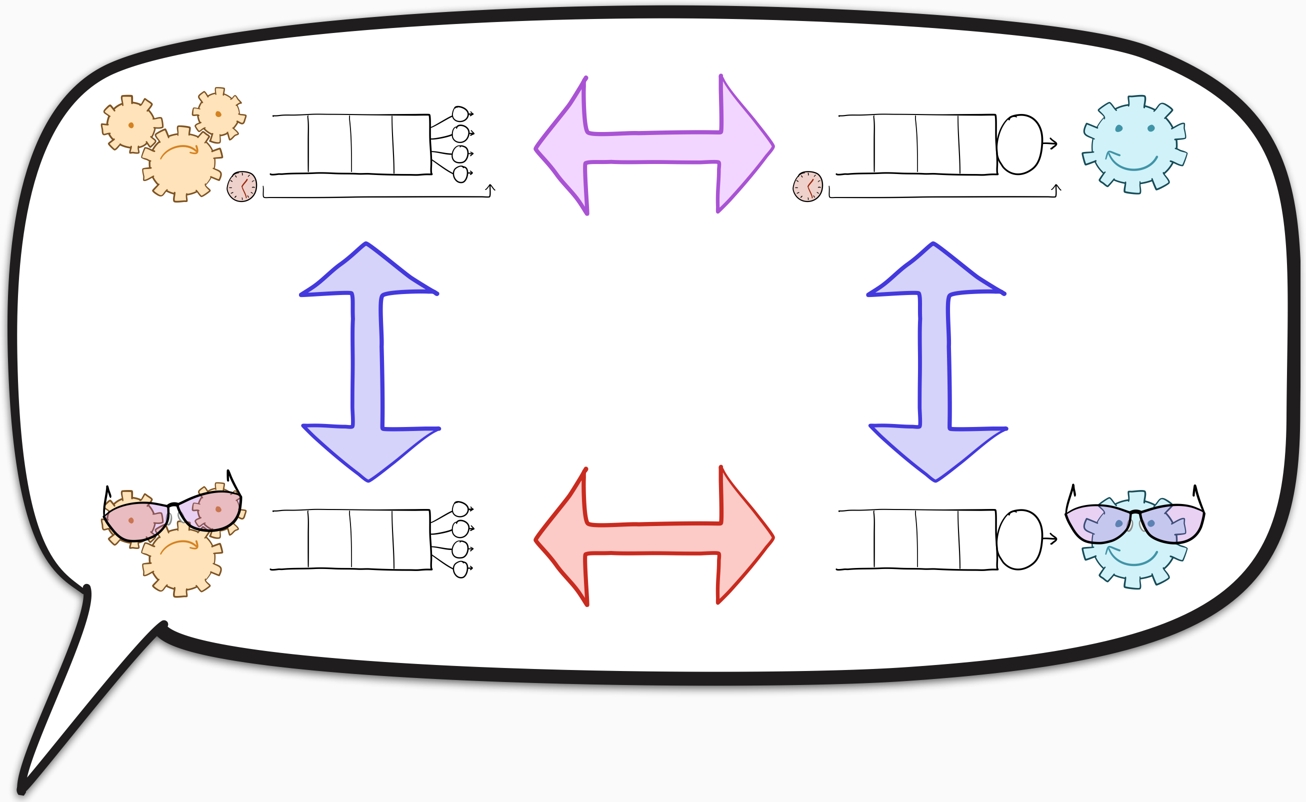

New technique based on relating $\mathbf{E}[Z]$ to ***r*-work**

$\mathbf{E}[Z_k^{\mathrm{RS}}] \leq \mathbf{E}[Z_1^{\mathrm{RS}}] + (6 \text{ or } 54)k$

RS has "near-optimal" $\mathbf{E}[Z]$ in the M/G/k and M/G/k/dispatch

Summary

Minimize mean slowdown in multiserver systems

Multiserver systems very hard to analyze directly

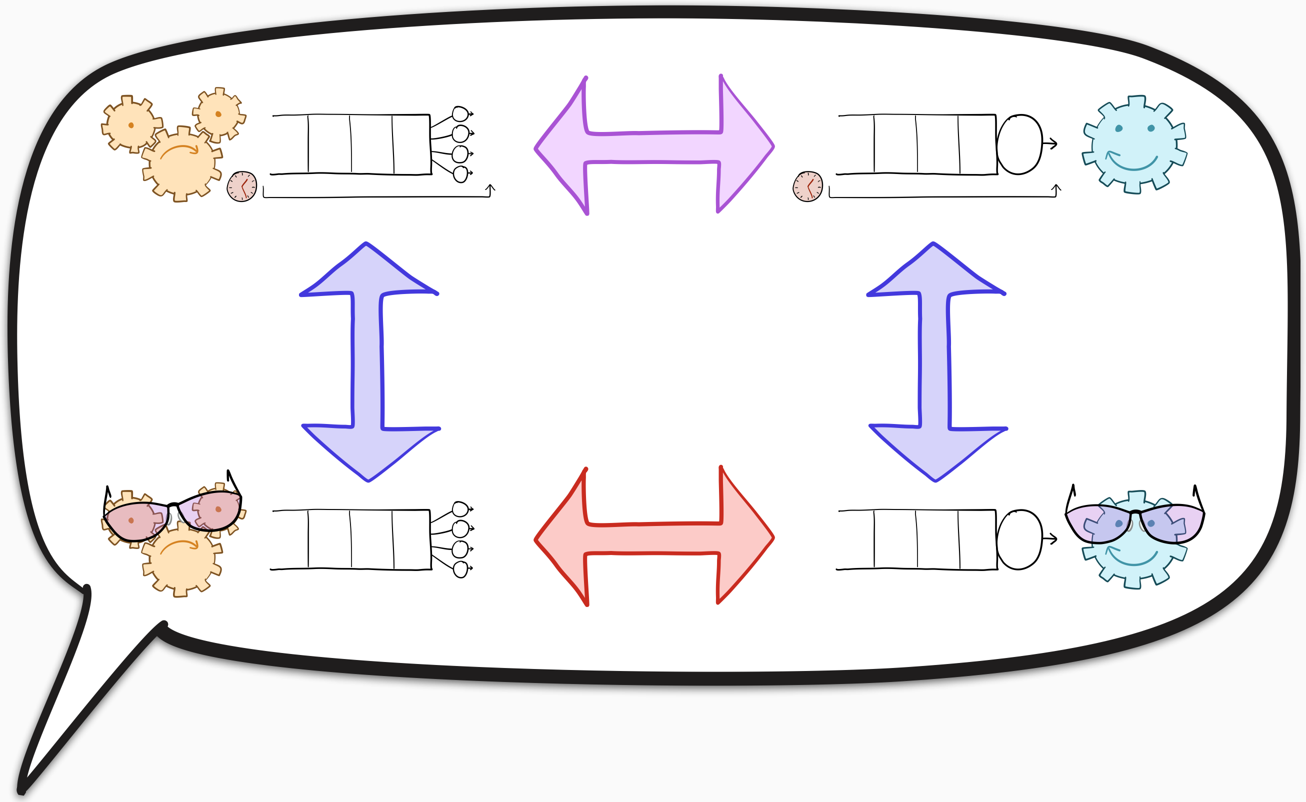

New technique based on relating $\mathbf{E}[Z]$ to ***r*-work**

$$\mathbf{E}[Z_k^{\mathrm{RS}}] \le \mathbf{E}[Z_1^{\mathrm{RS}}] + (6 \text{ or } 54)k$$

RS has "near-optimal" $\mathbf{E}[Z]$ in the M/G/k and M/G/k/dispatch

Get in touch: `zscully@cs.cmu.edu`