Asymptotically Optimal Multiserver Scheduling

Isaac Grosof

Ziv Scully

Mor Harchol-Balter

Carnegie Mellon University

YEQT 2021

Two papers on Multiserver Scheduling

"SRPT for Multiserver Systems". Grosof, Scully, and Harchol-Balter. IFIP Performance 2018

"The Gittins Policy is Nearly Optimal in the M/G/k under Extremely General Conditions". Grosof, Scully, and Harchol-Balter. ACM SIGMETRICS 2021

View at my website: isaacg1.github.io

Multiserver scheduling: M/G/k

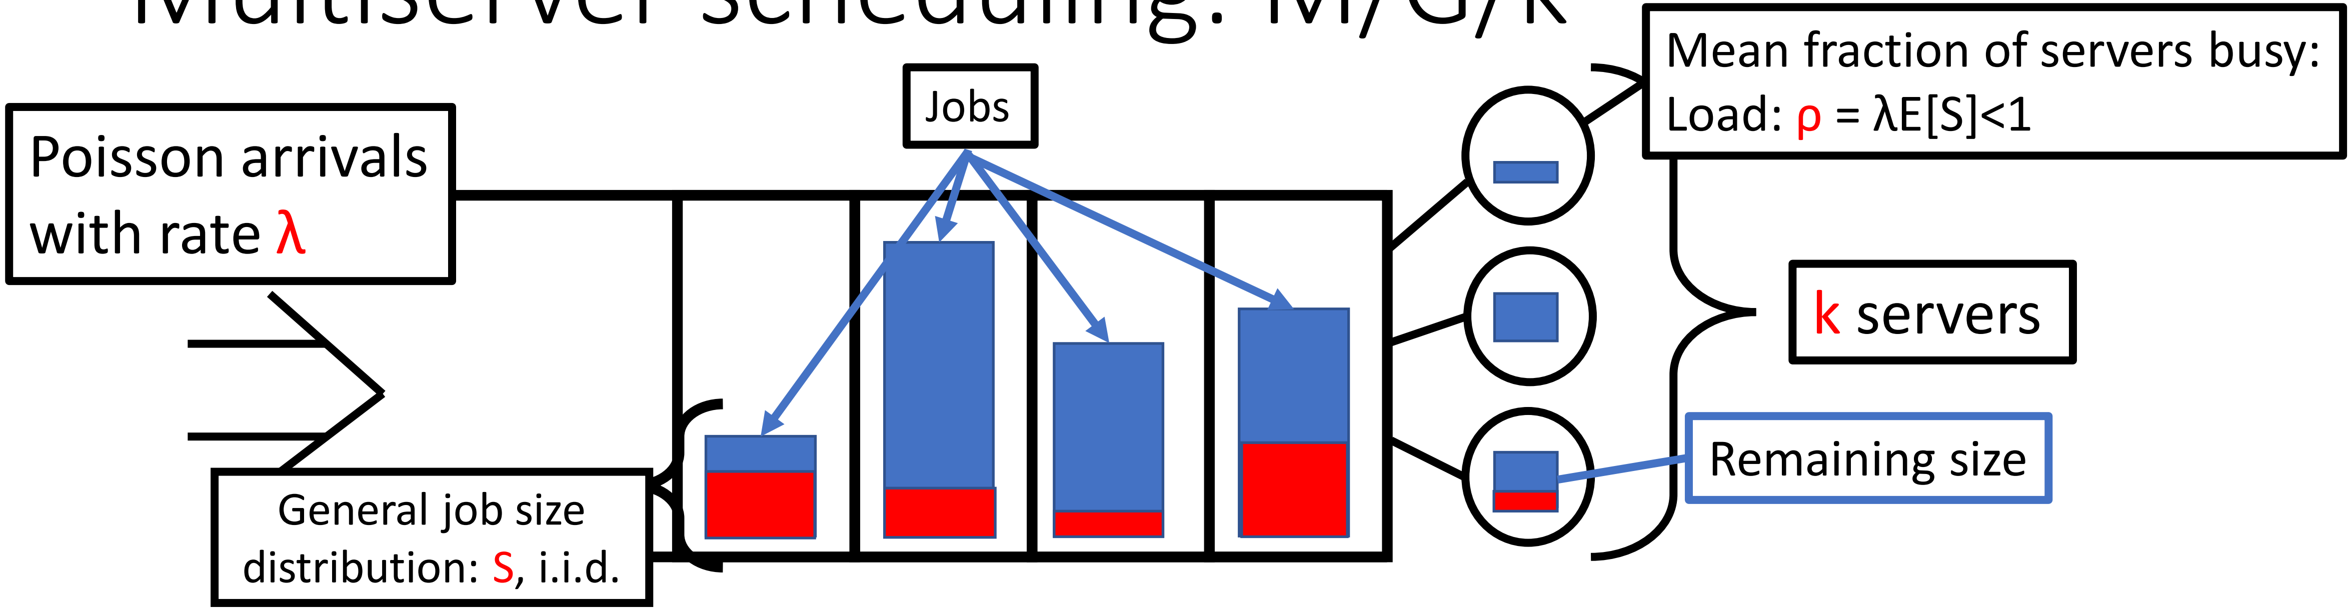

Q: How should we schedule?
Goal: Minimize mean response time E[T]
in $\rho \rightarrow 1$ limit

Response time: T
Time from arrival to completion

Known size M/G/k

"SRPT for Multiserver Systems". IFIP Performance 2018

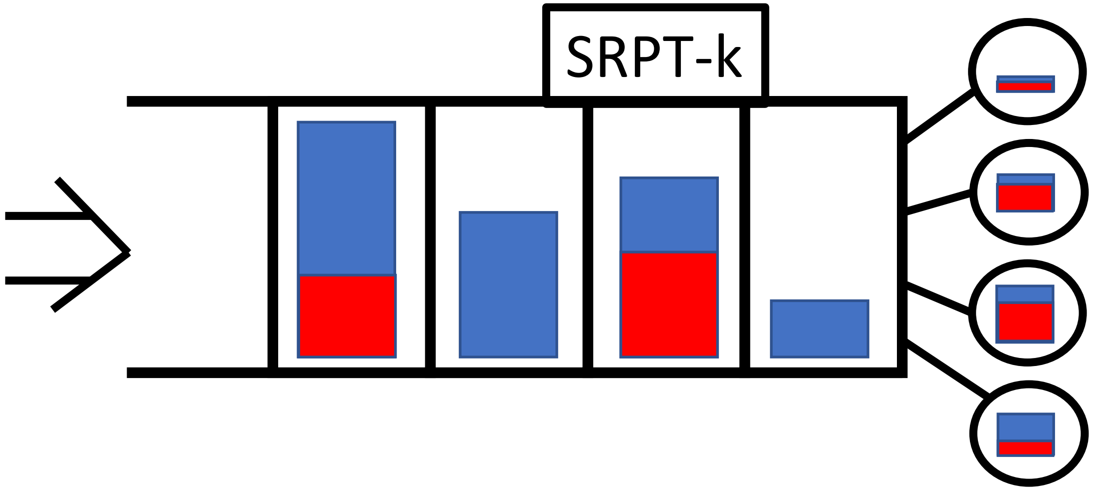

First paper to bound $E[T^{SRPT-k}]$

First paper to prove heavy-traffic optimality of SRPT-k.

Bonus: Also handles PSJF, SMART, FB, etc. See paper.

Proof structure

We want to prove heavy-traffic optimality:

$$\lim_{\rho \to 1} \frac{E[T^{SRPT-k}]}{E[T^{OPT-k}]} = 1$$

~~Solve the steady state?~~

~~State space collapse?~~

Directly bound $E[T^{SRPT-k}]$!

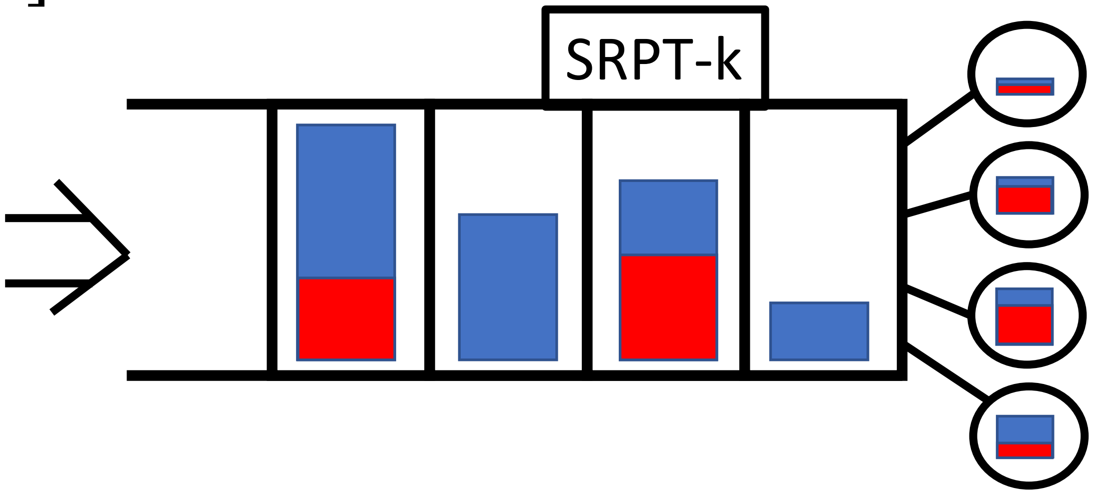

Idea: Compare to M/G/k to M/G/1

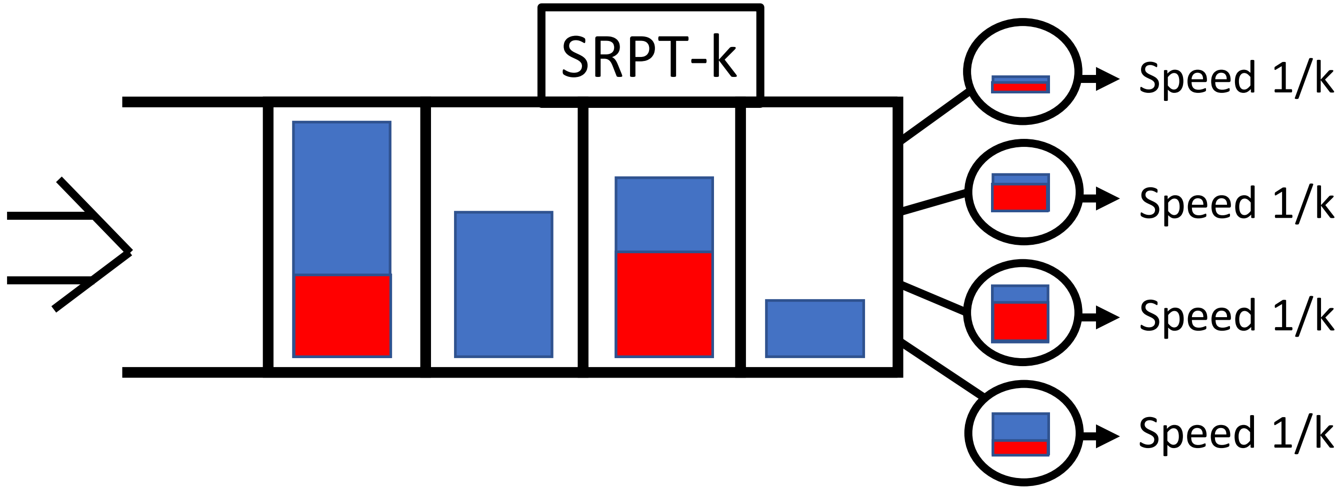

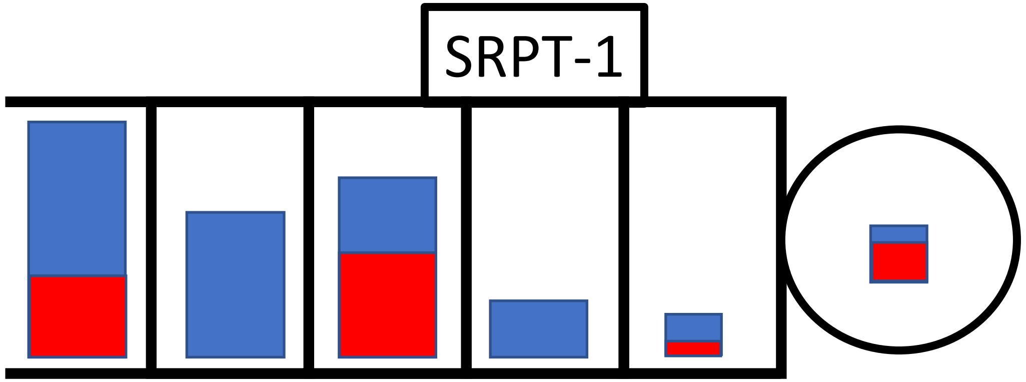

Idea: Use similarity of SRPT-k and SRPT-1 to bound $E[T^{SRPT-k}]$ relative to $E[T^{SRPT-1}]$

$$E[T^{OPT-1}] = E[T^{SRPT-1}]$$

To prove SRPT-k is optimal, it suffices to show that

$$\lim_{\rho \to 1} \frac{E[T^{SRPT-k}]}{E[T^{SRPT-1}]} = 1$$

Background for bound: Relevant work

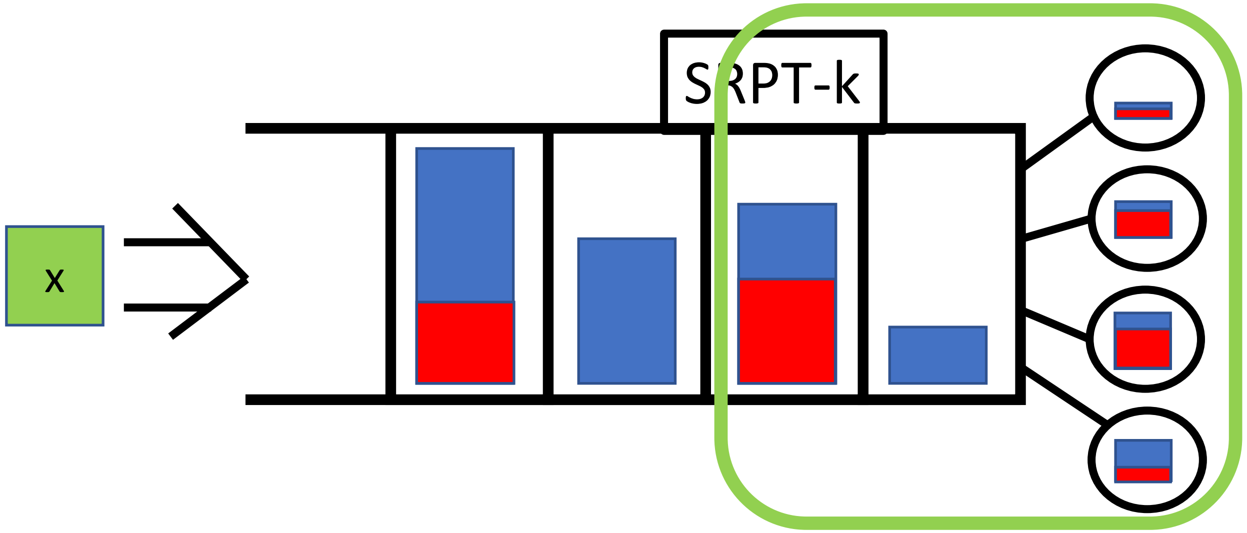

RelevantWork(x): Total remaining size of jobs with remaining size $\leq x$.

Sketch of SRPT-k bound

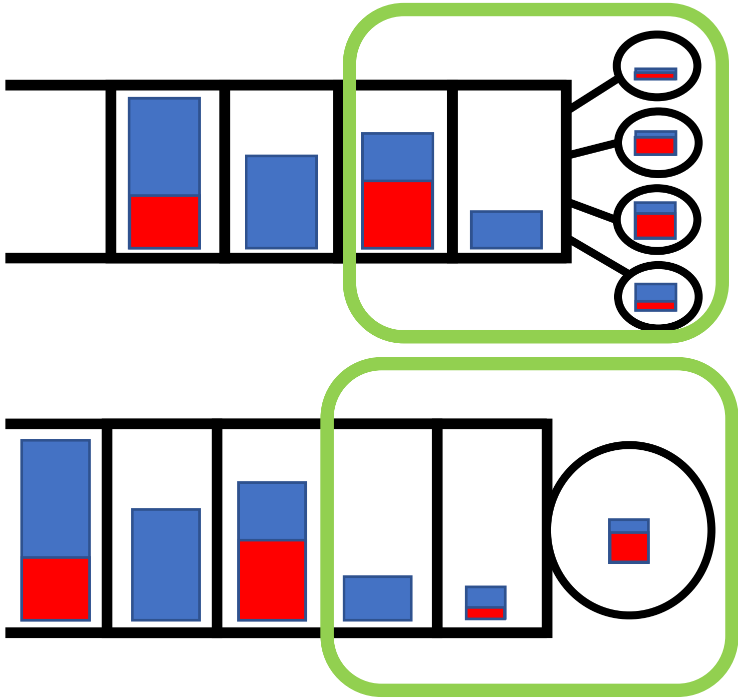

1. SRPT-k and SRPT-1 have similar RelevantWork(x).

   Sample-path argument, worst-case ideas.

2. Bound response time in terms of RelevantWork(x).

   Tagged job argument, M/G/1 ideas.

Bound:

$$E[T^{SRPT-k}] \leq E[T^{SRPT-1}] + \frac{2k}{\lambda} \ln \frac{1}{1-\rho}$$

SRPT-k Optimality (Performance '18)

$$E[T^{SRPT-k}] \leq E[T^{SRPT-1}] + \frac{2k}{\lambda} \ln \frac{1}{1-\rho}$$

Proven: $\lim_{\rho \to 1} \frac{E[T^{SRPT-k}]}{E[T^{SRPT-1}]} = 1$

Therefore: SRPT-k is heavy-traffic optimal.

Prior work [LWZ'11]:
$\ln \frac{1}{1-\rho} = o(E[T^{SRPT-1}])$,
given finite variance.*

*Actually, $E[S^2 (\log S)^+] < \infty$

Unknown size M/G/k

“The Gittins Policy is Nearly Optimal in the M/G/k under Extremely General Conditions”. ACM SIGMETRICS 2021

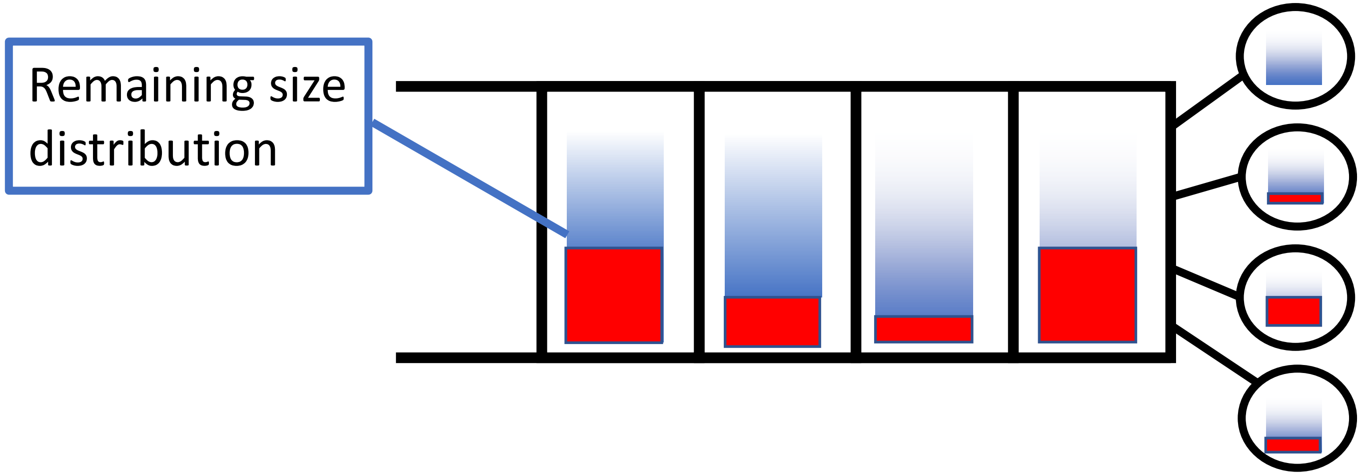

What’s a good policy for minimizing mean response time?

Unknown size M/G/1: Gittins

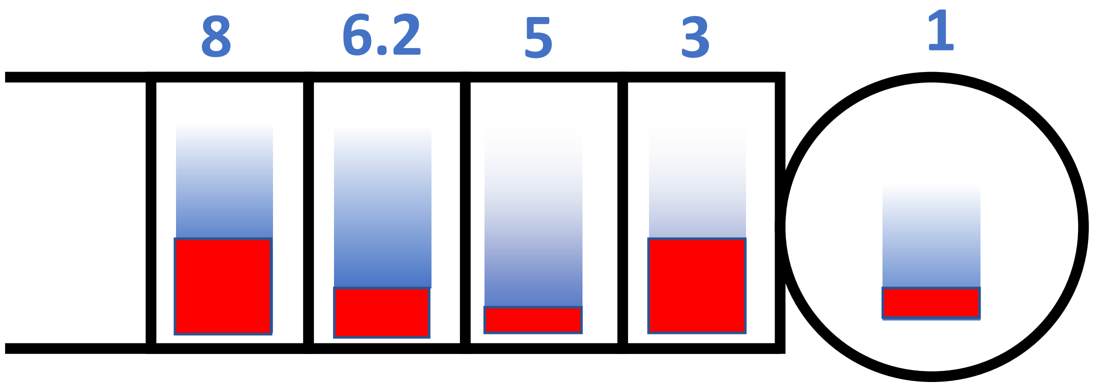

Optimal policy known: Gittins Policy [G ‘79]
Gittins also optimal for estimated sizes, staged jobs, more
Gittins assigns a rank to every remaining size distribution.
Over time, remaining size distributions change, ranks go up or down.

Unknown size M/G/1: Gittins

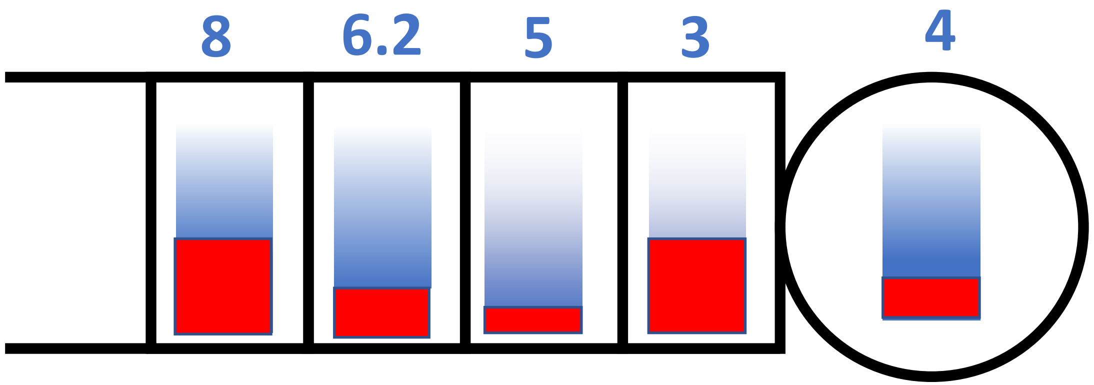

Optimal policy known: Gittins Policy [G '79]
Gittins also optimal for estimated sizes, staged jobs, more
Gittins assigns a rank to every remaining size distribution.
Over time, remaining size distributions change, ranks go up or down.

Unknown size M/G/1: Gittins

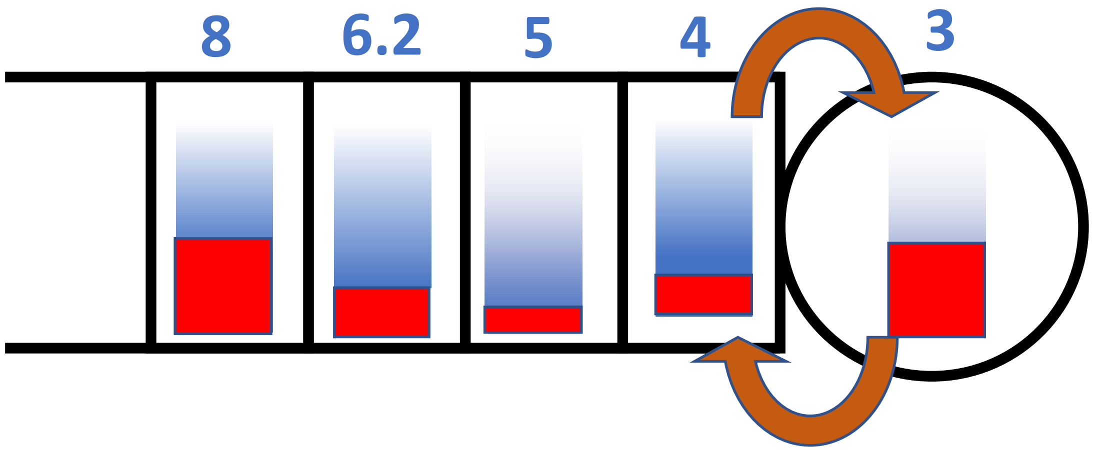

Optimal policy known: Gittins Policy [G '79]
Gittins also optimal for estimated sizes, staged jobs, more
Gittins assigns a rank to every remaining size distribution.
Over time, remaining size distributions change, ranks go up or down.

Unknown size M/G/k

"The Gittins Policy is Nearly Optimal in the M/G/k under Extremely General Conditions". ACM SIGMETRICS 2021

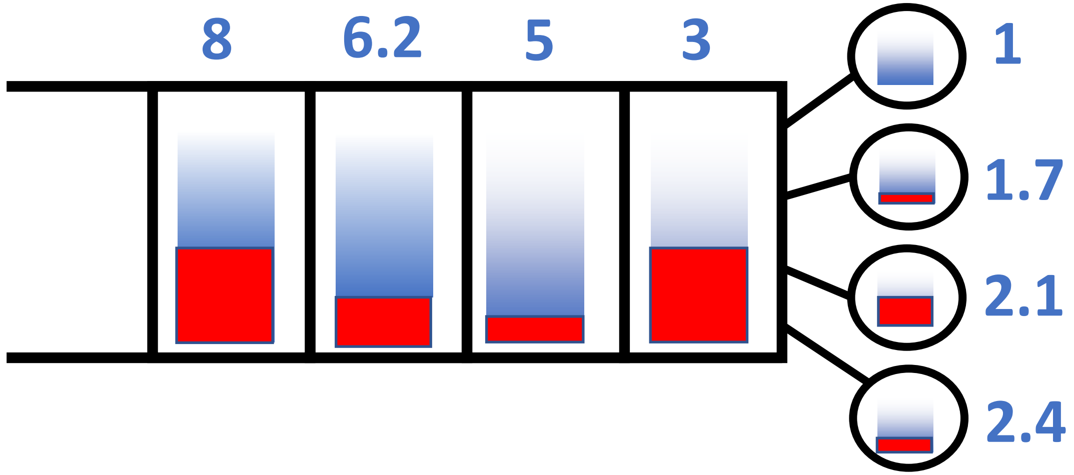

First paper to bound $E\left[T^{Gittins-k}\right]$

First paper to prove heavy-traffic optimality for Gittins-k.

Bonus: More settings. See paper.

Revision to relevant work

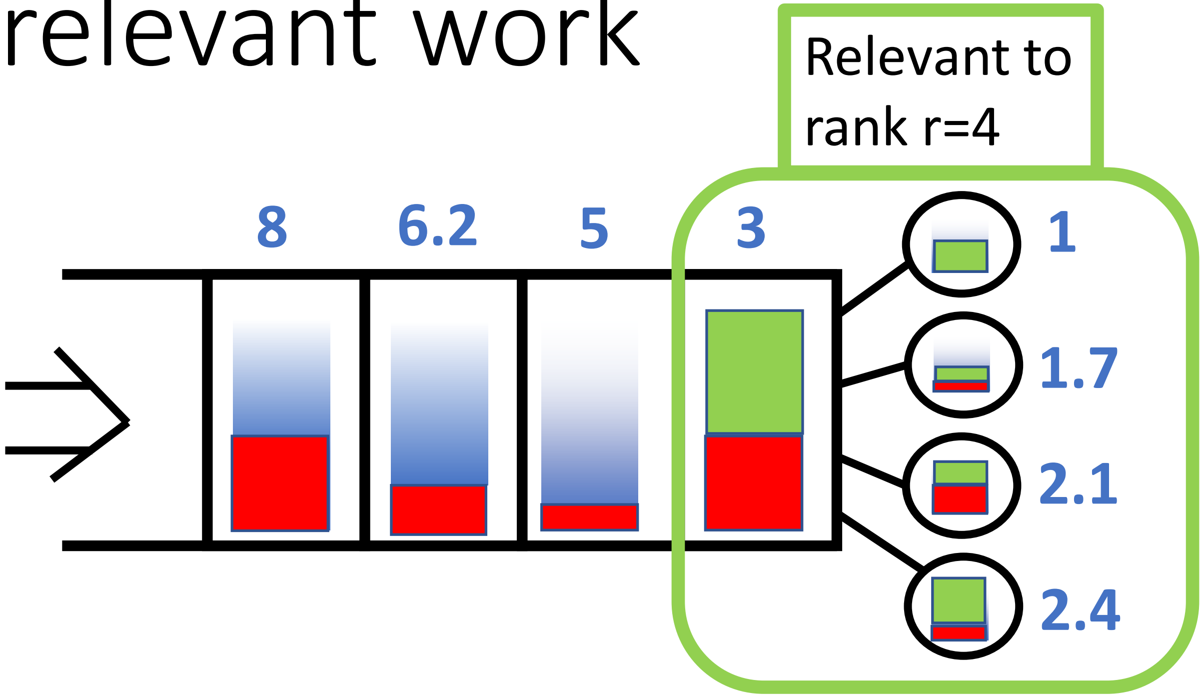

RelevantWork(r): Total service until all jobs complete or reach rank >r

Sketch of Gittins-k bound

1. Gittins-k and Gittins-1 have similar RelevantWork(r)

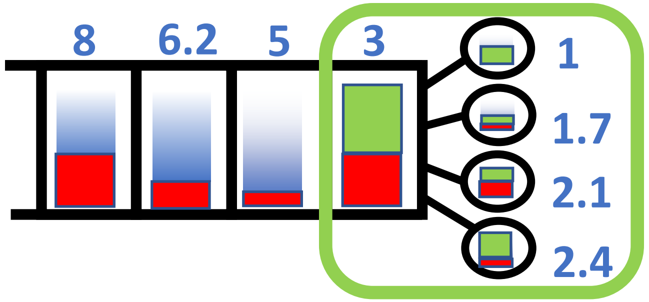

2. Bound response time in terms of RelevantWork(r)

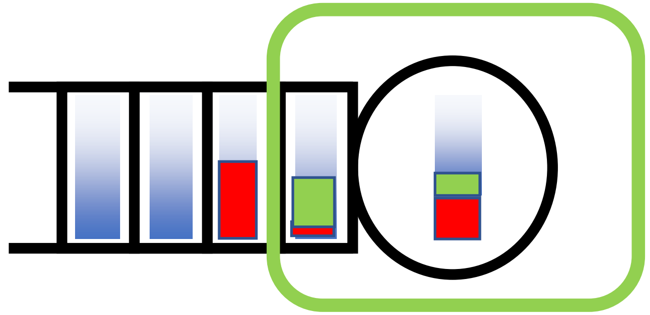

3. Bound

Sketch of Gittins-k bound

1. Gittins-k and Gittins-1 have similar RelevantWork(r)

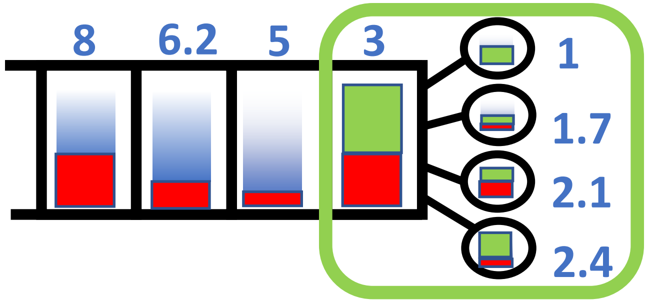

SRPT method relied on deterministic remaining size.

New method: Conservation of $W(r)^2$. Palm Calculus.

2. Bound response time in terms of RelevantWork(r)

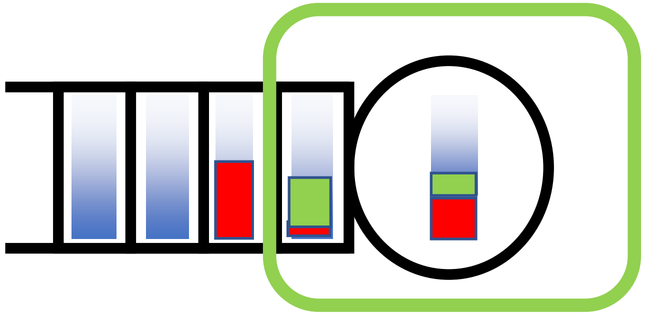

SRPT method relied on remaining size only decreasing.

New method: Gittins-specific response time formula.

3. Bound

$$E\left[T^{Gittins-k}\right] \le E\left[T^{Gittins-1}\right] + (k-1)E[S]\left(\ln\frac{1}{1-\rho} + \ln\frac{E[S^2]}{E[S]^2} + 4.811\right)$$

Gittins-k Optimality (SIGMETRICS '21)

$$E\left[T^{Gittins-k}\right] \leq E\left[T^{Gittins-1}\right] + (k-1)E[S]\left(\ln\frac{1}{1-\rho} + \ln\frac{E[S^2]}{E[S]^2} + 4.811\right)$$

$\ln\frac{1}{1-\rho} = o\left(E\left[T^{Gittins-1}\right]\right)$, given finite variance.*

Proven: $\lim_{\rho\to 1}\frac{E\left[T^{Gittins-k}\right]}{E\left[T^{Gittins-1}\right]} = 1$

Gittins-k is heavy-traffic optimal.

*Actually, $E[S^2(\log S)^+] < \infty$

Future directions

- We proved heavy traffic optimality for SRPT-k and Gittins-k.
  Heavy traffic optimality often indicates good performance at all loads.
  Outside of heavy traffic: Better upper bounds? Better lower bounds?
- We bounded response time in terms of work for monotonic policies like SRPT, and for the Gittins family of policies.
  Can we derive a tight bound for non-monotonic, non-Gittins policies?
- We proved that Gittins is heavy-traffic optimal under unknown sizes.
  However, Gittins is complicated and hard to implement.
  Can optimality be achieved by a simple policy?

Conclusion

"SRPT for Multiserver Systems". IFIP Performance 2018

Known size: First bound on $E[T^{SRPT-k}]$, first heavy-traffic optimality.

"The Gittins Policy is Nearly Optimal in the M/G/k under Extremely General Conditions. ACM SIGMETRICS 2021

Unknown size: First bound on $E\left[T^{Gittins-k}\right]$, first heavy-traffic optimality.