


Methods

When Does the Gittins Policy Have Asymptotically Optimal Response Time Tail in the M/G/1?

Ziv Scully,^{a,*} Lucas van Kreveld^b^aSchool of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853; ^bStochastic Operations Research, Eindhoven University of Technology, 5612 AZ Eindhoven, Netherlands

*Corresponding author

Contact: zivscully@cornell.edu,  <https://orcid.org/0000-0002-8547-1068> (ZS); l.r.v.kreveld@tue.nl,  <https://orcid.org/0000-0001-9588-6083> (LvK)

Received: January 23, 2022

Revised: May 5, 2023; December 5, 2023

Accepted: December 6, 2023

Published Online in Articles in Advance:
February 19, 2024

Area of Review: Stochastic Models

<https://doi.org/10.1287/opre.2022.0038>

Copyright: © 2024 INFORMS

Abstract. We consider scheduling in the M/G/1 queue with unknown job sizes. It is known that the Gittins policy minimizes mean response time in this setting. However, the behavior of the tail of response time under Gittins is poorly understood, even in the large-response-time limit. Characterizing Gittins's asymptotic tail behavior is important because if Gittins has optimal tail asymptotics, then it simultaneously provides optimal mean response time and good tail performance. In this work, we give the first comprehensive account of Gittins's asymptotic tail behavior. For heavy-tailed job sizes, we find that Gittins always has asymptotically optimal tail. The story for light-tailed job sizes is less clear-cut: Gittins's tail can be optimal, pessimal, or in between. To remedy this, we show that a modification of Gittins avoids pessimal tail behavior, while achieving near-optimal mean response time.

Funding: Z. Scully was supported by NSF [Grants CMMI-1938909, CMMI-2307008, CSR-1763701, DMS-2023528, and DMS-2022448]. L. van Kreveld was supported by the NWO through Gravitation grant NETWORKS-024.002.003.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2022.0038>.

Keywords: queueing theory • scheduling • M/G/1 queue • Gittins index • response time tail • heavy-tailed distributions • light-tailed distributions

1. Introduction

Scheduling to minimize response time (also known as sojourn time) of single-server queueing models is an important problem in queueing theory, with applications in computer systems, service systems, and beyond. In general, a queueing system will have a response time *distribution*, denoted T , and there are a variety of metrics that one might hope to minimize. There is significant work on minimizing *mean response time* $E[T]$, which is the average response time of all jobs in a long arrival sequence (Schrage 1968, Gittins 1989, Aalto et al. 2009, Gittins et al. 2011).

Much less is known about minimizing the *tail of response time* $P[T > t]$, which is the probability that a job has response time greater than a parameter $t \geq 0$. In light of the difficulty of studying the tail directly, theorists have studied the *asymptotic tail of response time*, which is the asymptotic decay of $P[T > t]$ in the $t \rightarrow \infty$ limit (Stolyar and Ramanan 2001, Núñez-Queija 2002, Borst et al. 2003, Boxma and Zwart 2007, Scully et al. 2020b). In this

work, we consider the preemptive M/G/1 queue and ask the following question.

Question 1.1. Does any scheduling policy simultaneously optimize the mean and asymptotic tail of response time in the M/G/1?

Our focus on the M/G/1, a classic, single-server queueing model (Kendall 1953, Cox and Smith 1961, Kleinrock 1976, Harchol-Balter 2013), is motivated by its balance between modeling flexibility and analytical tractability. The fact that the distribution of job sizes (also known as service times) may be general is particularly important for modeling computer systems, where the distribution can be far from exponential (Peterson 1996, Crovella and Bestavros 1997, Harchol-Balter and Downey 1997).

Prior work answers Question 1.1 when job sizes are known to the scheduler. In this setting, the *Shortest Remaining Processing Time* (SRPT) policy, which preemptively serves the job of least remaining size, always minimizes mean response time (Schrage 1968).

However, SRPT's tail performance depends on the job size distribution.¹

- If the job size distribution is *heavy-tailed* (roughly, power-law; see Definition 4.1), then SRPT is *tail-optimal*, meaning that it has the best possible asymptotic tail decay (Definition 4.2).

- If the job size distribution is *light-tailed* (roughly, subexponential; see Definition 5.2), then SRPT is *tail-pessimal*, meaning that it has the worst possible asymptotic tail decay (Definition 5.3).

This answers Question 1.1 for known job sizes: “yes, namely SRPT” in the heavy-tailed case, “no” in the light-tailed case.

Unfortunately, in practice, the scheduler often does not know job sizes, and, thus, one cannot implement SRPT. Instead, the scheduler often only knows the job size *distribution*. We study Question 1.1 in this unknown-size setting.

The question of minimizing mean response time with unknown job sizes was settled by Gittins (1989). He introduced a policy, now known as the *Gittins* policy, which leverages the job size distribution to minimize mean response time. Roughly speaking, Gittins uses each job's *age*—namely, the amount of time each job has been served so far—to figure out which job is most likely to complete after a small amount of service, then serves that job. For some job size distributions, Gittins reduces to a simpler policy, such as *First-Come, First-Served* (FCFS) or *Foreground-Background* (FB) (Aalto et al. 2009, 2011).

In the unknown-size setting, given that Gittins minimizes mean response time, Question 1.1 reduces to the following.

Question 1.2. For which job size distributions is Gittins tail-optimal for response time?

Unfortunately, the asymptotic tail behavior of Gittins is understood in only a few special cases.

- In the heavy-tailed case, Gittins has been shown to be tail-optimal, but only under an assumption on the job size distribution's hazard rate (Scully et al. 2020b, corollary 3.5).

- In the light-tailed case, Gittins sometimes reduces to FCFS or FB (Aalto et al. 2009, 2011). For light-tailed job sizes, FCFS is tail-optimal (Stolyar and Ramanan 2001, Boxma and Zwart 2007), but FB is tail-pessimal (Mandjes and Nuyens 2005).

This prior work leaves Question 1.2 largely open. We do not know whether Gittins is always tail-optimal in the heavy-tailed case, or whether it is sometimes sub-optimal, or even tail-pessimal. Moreover, we do not understand Gittins's asymptotic tail at all in the light-tailed case, aside from when Gittins happens to reduce to a simpler policy.

The prior work above does tell us an important fact: Gittins *can* be tail-pessimal. This prompts another question.

Question 1.3. For job size distributions for which Gittins is tail-pessimal, is there another policy that has near-optimal mean response time while not being tail-pessimal?

In this work, we answer Questions 1.1–1.3 for the M/G/1 with unknown job sizes, covering wide classes of heavy- and light-tailed job size distributions. The key tool we use to analyze Gittins's asymptotic response time tail is the *SOAP* (Schedule Ordered by Age-based Priority) framework (Scully and Harchol-Balter 2018, Scully et al. 2020b). SOAP gives a universal M/G/1 response time analysis of all *SOAP policies*, which are scheduling policies where a job's priority level is a function of its age (Definition 3.1). Underlying our Gittins results is a general tail analysis of SOAP policies.

Our main contributions, which we describe in more detail later (Sections 4 and 5), are as follows:

- *Heavy-tailed case:* We give a sufficient condition under which an arbitrary SOAP policy is tail-optimal (Section 7).

- *Heavy-tailed case:* We show that the above condition always applies to Gittins, implying that it is always tail-optimal (Section 8).

- *Light-tailed case:* We characterize when an arbitrary SOAP policy is tail-optimal, tail-pessimal, or in between (Section 9).

- *Light-tailed case:* We spell out how the above characterization applies to Gittins and show how to modify Gittins to avoid tail pessimality (Section 10).

- *General case:* At the core of our modification of Gittins which avoids tail pessimality is a general result, which states that slightly perturbing the Gittins rank function only slightly affects its mean response time (Theorem 5.10 and Online Appendix EC.2).²

The rest of the paper introduces definitions and notation (Sections 3 and 6) and concludes with some remarks about our motivating questions (Section 11).

2. Prior Work

2.1. Asymptotic Tail Analysis of Classic Scheduling Policies

Because of their frequent occurrence, light-tailed job size distributions have received a great amount of attention by queueing theorists. The performance of policies under light-tailed job sizes is generally measured in terms of the decay rate of the response time tail. In this sense, FCFS has proven to be optimal among all service policies (Stolyar and Ramanan 2001). Conversely, *Foreground-Background Processor Sharing* has the worst possible decay rate of the response time tail (Mandjes and Nuyens 2005).

On the other hand, it is shown that heavy-tailed job sizes can have a large impact on the performance characteristics of the queue. For this reason also, heavy-tailed

job sizes have been thoroughly investigated in the literature. For example, researchers have made the striking observation that, contrary to the light-tailed case, FB is optimal where FCFS has the worst possible response time tail (Borst et al. 2003). This dichotomy between light and heavy tails is not limited to FCFS and FB (Boxma and Zwart 2007).

Other noteworthy literature highlighting both light and heavy tails includes delicate asymptotic results for a two-class priority policy (Abate and Whitt 1997) and robust optimization using a limited PS policy (Nair et al. 2010).

With one exception, discussed in Section 2.3, the literature on this subject concerns only a few relatively simple policies. This paper considers policies in which the priority of a job can vary essentially arbitrarily with its age. This generality is needed to analyze the Gittins policy, where a job's priority can be nonmonotonic (Aalto et al. 2011).

2.2. Impossibility of Universal Tail-Optimal Scheduling

In light of the fact that both FCFS and FB can vary between tail-optimal and tail-pessimal for different job sizes, it is natural to ask whether there is a single policy that is always tail-optimal. Wierman and Zwart (2012) answer this question with an impossibility result, showing that no policy is tail-optimal for both heavy- and light-tailed job sizes, unless the policy has knowledge of the size distribution X or learns X over time.

One might worry that this impossibility result contradicts our results for Gittins, given that we show Gittins is always tail-optimal in the heavy-tailed case and sometimes tail-optimal in the light-tailed case. The reason there is no contradiction is that the Gittins policy changes based on the size distribution. For instance, there are some light-tailed distributions where Gittins reduces to FCFS and some heavy-tailed distributions where Gittins reduces to FB (Aalto et al. 2009).

2.3. Tail Optimality of Certain SOAP Policies in the Heavy-Tailed Case

We mention particularly the relation between this paper and the work of Scully et al. (2020b). Both this paper and the prior work study the response-time tail behavior of arbitrary SOAP policies, including the Gittins policy. There are two main factors that distinguish this paper from the prior work.

- Scully et al. (2020b) only study heavy-tailed job size distributions. In contrast, we study both the heavy- and light-tailed cases.
- Scully et al. (2020b) show that Gittins is tail-optimal, subject to a condition on the job size distribution's hazard rate (Scully et al. 2020b, corollary 3.5). However, their analysis is not sharp enough to completely characterize

under which (heavy-tailed) job size distributions Gittins is tail-optimal. In contrast, our analysis is sharper, allowing us to identify Gittins's tail performance under any job size distribution.

With this said, Scully et al. (2020b) lay an important technical foundation upon which we build to derive our heavy-tailed results. See Section 4.3 for a more technical discussion of what aspects of their work we use and what aspects we improve upon.

2.4. Beyond Asymptotic Tail Optimality

It is well known that FCFS has optimal tail decay rate under light-tailed job sizes. However, decay rate is a relatively crude tail performance measure, as it does not take into account the constant (or nonexponential term) in front of the exponent. Although this paper focuses just on decay rates, we mention that, very recently, a policy was introduced that has a better leading constant than FCFS (Grosz et al. 2021). An open question remains of what is the best possible leading constant in the response time tail. A by-product of our results—namely, that FCFS is the only SOAP policy with optimal decay rate—partially answers this question. Specifically, it follows that no SOAP policy is tail-optimal up to the leading constant.

2.5. Mean Response Time of Modified Gittins Policies

A recent study (Scully et al. 2022, theorem 7.2) shows that if one slightly modifies the prioritization rules of SRPT, then the mean response time of the resulting policy is only slightly worse than that of unmodified SRPT (which is optimal in case job sizes are known). It turns out, as shown in this paper, that essentially the same result holds for an approximate version of the Gittins policy, which can thus be seen as the unknown-job-sizes counterpart of Scully et al. (2022, theorem 7.2).³

3. Model, SOAP Policies, and the Gittins Policy

We consider an M/G/1 queue with arrival rate λ , job size distribution X , and load $\rho = \lambda E[X]$. For the tail of the job size distribution, we write $\bar{F}(t) = P[X > t]$. We denote the maximum job size by $x_{\max} = \inf\{t \geq 0 \mid \bar{F}(t) = 0\}$, allowing $x_{\max} = \infty$. We write T_π for the M/G/1's response time distribution under policy π . We allow policies to preempt jobs or share the processor without any overhead or loss of work (i.e., the model is preempt-resume).

Special attention is given in this paper to the Gittins policy. It assigns each job a *rank*—namely, a priority—based on the job's *age*—namely, the amount of time the job has been served so far. To analyze the Gittins policy, we make use of the *SOAP framework* (Scully and

Harchol-Balter 2018, Scully et al. 2018, 2020b), which gives a response time analysis of the following broad class of policies.

Definition 3.1. A SOAP policy is a policy π specified by a rank function $r_\pi : [0, x_{\max}) \rightarrow \mathbb{R}$. Policy π assigns rank $r_\pi(a)$ to a job at age a .⁴ When the policy being discussed is clear from context, we often omit the subscript and simply write $r(a)$. At every moment in time, a SOAP policy serves the job of minimum rank, breaking ties in FCFS order.⁵

Definition 3.2. The Gittins policy, denoted “Gtn” in subscripts for brevity, is the SOAP policy with rank function

$$r_{\text{Gtn}}(a) = \inf_{b>a} \frac{\mathbf{E}[\min\{S, b\} - a | S > a]}{\mathbf{P}[S \leq b | S > a]} = \inf_{b>a} \frac{\int_a^b \bar{F}(t) dt}{\bar{F}(a) - \bar{F}(b)}.$$

Note that the Gittins rank function depends on the job size distribution X by way of \bar{F} .

As is standard (Scully and Harchol-Balter 2018, Scully et al. 2018 appendix B), we assume that rank functions are *piecewise-continuous* and *piecewise-monotonic*, with finitely many pieces in any compact interval for both properties. This holds for Gittins under very mild conditions on the job size distribution X . For example, Aalto et al. (2011) show that the Gittins rank function is continuous and piecewise-monotonic, provided that X is a continuous distribution with continuous and piecewise-monotonic hazard rate. However, our results are not restricted to continuous job size distributions. Our generic SOAP results require no additional assumptions on X , and our Gittins results require only that X induces a piecewise-continuous and piecewise-monotonic Gittins rank function, which can occur even if X is not continuous.

4. Heavy-Tailed Job Sizes

In Section 4.1, we define which job size distributions are heavy-tailed, and we give our criterion for tail optimality in this scenario. The two main results in the heavy-tailed case are presented in Section 4.2:

- Theorem 4.6 gives a sufficient condition for a SOAP policy to be tail-optimal for heavy-tailed job sizes.
- Theorem 4.7 shows that for heavy-tailed job sizes, Gittins always satisfies this sufficient condition and is thus always tail-optimal.

4.1. Background on Heavy-Tailed Job Sizes

Roughly speaking, the heavy-tailed job size distributions we study are those that are asymptotically Pareto. The specific class that we study, described below, is slightly more general in that it also includes distributions whose tails oscillate between Pareto tails of different shape parameters.

Definition 4.1 (Heavy-Tailed Job Size Distribution). We say a job size distribution X is *nicely heavy-tailed* if $x_{\max} = \infty$ and both of the following hold:

i. The tail $\bar{F}(\cdot)$ is of intermediate regular variation (Cline 1994), meaning

$$\liminf_{\epsilon \downarrow 0} \liminf_{x \rightarrow \infty} \frac{\bar{F}((1 + \epsilon)x)}{\bar{F}(x)} = 1.$$

ii. There exist $\beta \geq \alpha > 1$ such that the upper and lower Matuszewska indices of $\bar{F}(\cdot)$ are in $(-\beta, -\alpha)$ (Bingham et al. 1987, section 2.1). This implies that for all sufficiently large $x_2 \geq x_1$,⁶

$$\Omega\left(\left(\frac{x_2}{x_1}\right)^{-\beta}\right) \leq \frac{\bar{F}(x_2)}{\bar{F}(x_1)} \leq O\left(\left(\frac{x_2}{x_1}\right)^{-\alpha}\right).$$

In informal discussion, we omit “nicely.” Although the above definition includes all distributions with power-law-like tail decay, it is noted that we do not consider tails of a less heavy order, such as those of the lognormal and Weibull distributions.

Definition 4.2 (Tail Optimality in Heavy-Tailed Case). Consider an M/G/1 with nicely heavy-tailed job size distribution X . We call a scheduling policy π *tail-optimal* among preemptive work-conserving policies if

$$\lim_{t \rightarrow \infty} \frac{\mathbf{P}[T_\pi > t]}{\bar{F}((1 - \rho)t)} = 1.$$

That is, tail optimality holds if large jobs have a response time of approximately $1/(1 - \rho)$ times their size, which is the best possible asymptotic tail decay in the heavy-tailed case (Boxma and Zwart 2007, Wierman and Zwart 2012).

4.2. Results for Heavy-Tailed Case

Let us focus on a tagged job of size x . For determining whether it will be delayed by other jobs, it is important to know the worst (highest) ever rank that it will ever have, as well as the ages at which other jobs will have rank lower than that worst ever rank.


Definition 4.3. The *worst ever rank* of a job of size x is defined by $w_x = \sup_{0 \leq a < x} r(a)$.

Definition 4.4.

i. A *w-interval* is an interval (b, c) with $0 \leq b < c \leq x_{\max}$ such that $r(a) \leq w$ for all $a \in (b, c)$.

ii. A *w-interval* (b, c) is *right-maximal* if for all $c' \geq c$, the interval (b, c') is not a *w-interval*. This is equivalent to c satisfying either $r(c) \geq w$ or $c = x_{\max}$. We define *left-maximal* similarly, and we call a *w-interval maximal* if it is both left- and right-maximal.

Figure 1. (Color online) Illustration of Worst Ever Rank w_x (Purple Dotted Line) and Maximal w_x -Intervals (Shaded Orange Regions) for a SOAP Policy Given by Rank Function r (Solid Cyan Curve)



Note. A tagged job of size x always has rank w_x or better, so if another job has priority over the tagged job, that other job's age must be in a w_x -interval.

Note that the tagged job of size x , no matter its age, always has priority over jobs that have rank higher than w_x . Therefore, it can only be delayed by another job if that other job's age is in a w_x -interval. See Figure 1 for an illustration. To ensure that the tagged job does not wait too long behind other jobs, the w_x -intervals must be relatively short. We use the following condition to characterize the length of w_x -intervals.

Condition 4.5. *There exist $\zeta, \theta \in [0, \infty)$ and $\eta \in [\max\{1, \zeta + \theta\}, \infty]$ such that the following hold for any w_x -interval (b, c) with $b \geq x$:*

- i. *The w_x -interval's length is bounded by $c - b \leq O(b^{\zeta} x^{\theta})$.*
- ii. *The w_x -interval's endpoint is bounded by $c \leq O(x^{\eta})$.*

Note that to check that (i) and (ii) hold, it suffices to consider only right-maximal w_x -intervals.

At an intuitive level, we can think of Condition 4.5 as saying the following about how much other jobs delay the tagged job of size x :

- If ζ and θ are small enough, then w_x -intervals are relatively short, so jobs of age greater than x cannot delay the tagged job for too long. Figure 2 illustrates the difference between the roles of ζ and θ (see also Section 4.3).

- If η is small enough, then there are no w_x -intervals at sufficiently large ages, so jobs of sufficiently large age cannot delay the tagged job at all.

Rank functions that satisfy Condition 4.5 do so for many possible parameter values. For instance, if Condition 4.5 is satisfied with parameters (ζ, θ, η) , then it is also satisfied for $(\zeta + \delta, \theta - \delta + \epsilon, \eta + \epsilon)$ for all $\delta, \epsilon > 0$. But the idea is to find parameters that characterize a SOAP policy's rank function as tightly as possible because, as our main result below shows, if the parameters are small enough, then the policy is tail-optimal.

Theorem 4.6. *Consider an $M/G/1$ with any nicely heavy-tailed job size distribution under a SOAP policy. Condition 4.5 implies the policy is tail-optimal if its parameters satisfy*

$$\zeta + (\theta - 1)^+ - \frac{(1 - \theta)^+}{\eta} < \frac{\alpha - 1}{\beta}. \quad (4.1)$$

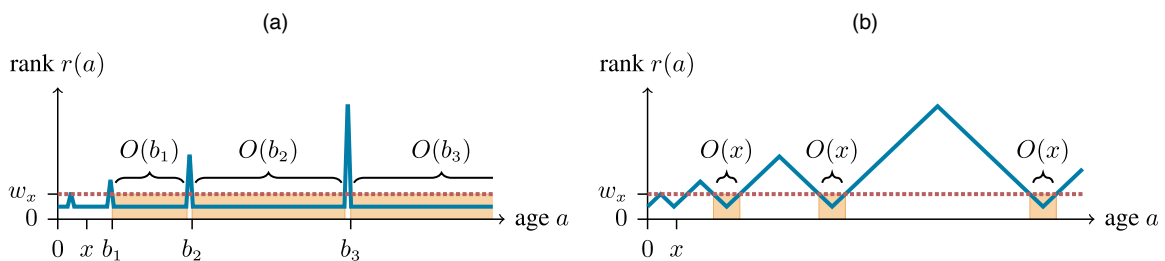
We can apply Theorem 4.6 to show that Gittins is tail-optimal. Specifically, we will show that Gittins satisfies Condition 4.5 with $\zeta = 0$, $\theta = 1$, and $\eta = \infty$. As such, it achieves a value of zero on the left-hand side of (4.1), so Gittins is tail-optimal, regardless of the values of $\beta \geq \alpha > 1$.

Theorem 4.7. *The Gittins policy is tail-optimal for any nicely heavy-tailed job size distribution.*

4.3. Comparing Our Tail Optimality Condition to That of Scully et al. (2020c)

Having formally stated our sufficient condition for tail optimality in the heavy-tailed case, we may now compare it in more detail to that of Scully et al. (2020b). Their condition (Scully et al. 2020b, assumption 3.2) and corresponding result (Scully et al. 2020b, theorem 3.3) are the same as our Condition 4.5 and Theorem 4.6, respectively, but restricted to the $\theta = 0$ case. This means, roughly speaking, that Scully et al. (2020b) only look at the lengths between “peaks” of the rank

Figure 2. (Color online) Illustration of Condition 4.5(i), with Examples Showing the Roles ζ and θ Play



Notes. (a) Rank function with $\zeta = 1$ and $\theta = 0$. (b) Rank function with $\zeta = 0$ and $\theta = 1$. Roughly speaking, one should think of $\zeta + \theta$ as characterizing how far apart different “peaks” of the rank function are, and one should think of $\zeta/(\zeta + \theta)$ as characterizing how “steep” the rank function is between peaks. Both (a) and (b) have the same peaks, as reflected by (a) and (b) having the same value of $\zeta + \theta$. But the slopes are much steeper in (a) than in (b), as reflected by the larger value of $\zeta/(\zeta + \theta)$ in (a) than in (b).

function, without looking directly at the length of w_x -intervals. For instance, the prior condition would treat the two rank functions shown in Figure 2 in the same way, even though the w_x -intervals are much longer in Figure 2(a) than in Figure 2(b).

Unfortunately, looking only at distances between peaks of the rank function is not enough to prove Gittins's tail optimality in the heavy-tailed case. For Gittins, it turns out that we cannot, in general, do better than setting $\eta = \infty$. If we had to set $\theta = 0$, then to make Gittins satisfy Condition 4.5, it turns out that we would need to set $\zeta = 1$, which is too large to satisfy (4.1). But the Gittins rank function looks less like Figure 2(a) and more like Figure 2(b): even though the peaks can be far apart, there are "gentle slopes" between them. Setting $\theta = 1$ and $\zeta = 0$ captures this behavior, and this satisfies (4.1). Thus, our refining of the sufficient condition of (Scully et al. 2020b, assumption 3.2) is necessary to prove Gittins's tail optimality in the heavy-tailed case, at least for this SOAP-based approach.

Underlying the tail optimality results of Scully et al. (2020b) is a busy period analysis combined with asymptotic response time bounds. We make use of their busy period analysis, which we distill into a simple statement (Lemma 7.4), but we replace their asymptotic response time bounds with a sharper analysis that accounts for the $\theta > 0$ possibility (Sections 7.1 and 7.2).

Finally, we reiterate that whereas Scully et al. (2020b) study only heavy-tailed size distributions, we also study light-tailed job size distributions. Our results for the light-tailed case are based on very different techniques, reflecting fundamental differences between the heavy- and light-tailed settings.

5. Light-Tailed Job Sizes

Similarly to the previous section, we first define the class of light-tailed distributions and state the corresponding tail-optimality criterion in Section 5.1. The main results in the light-tailed case, presented in Section 5.2, are summarized as follows:

- Theorem 5.5 classifies SOAP policies into tail-optimal, tail-intermediate, and tail-pessimal for light-tailed job sizes.
- Theorem 5.8 shows that for light-tailed job sizes, Gittins can be any of tail-optimal, tail-intermediate, or tail-pessimal.
- Theorem 5.10 shows that making a small change to the Gittins rank function results in only a small change to mean response time.
- Theorem 5.11 shows that for a wide class of light-tailed job size distributions for which Gittins is tail-pessimal, making a small change to Gittins's rank function results in a tail-optimal or -intermediate policy with mean response time arbitrarily close to Gittins's.

5.1. Background on Light-Tailed Job Sizes

Definition 5.1. The *decay rate* of random variable V , denoted $d(V)$, is

$$d(V) = \lim_{t \rightarrow \infty} \frac{-\log \mathbf{P}[V > t]}{t}.$$

That is, if the decay rate $d(V)$ is finite, then $\mathbf{P}[V > t] = \exp(-d(V)t \pm o(t))$. Higher decay rates thus correspond to asymptotically lighter tails.

Roughly speaking, the light-tailed job size distributions we study are those with positive decay rate. Our main tool for investigating the decay rate of a random variable V is via its Laplace-Stieltjes transform $\mathcal{L}[V]$, defined as

$$\mathcal{L}[V](s) = \mathbf{E}[\exp(-sV)] \in (0, \infty].$$

Under mild conditions on V (Nakagawa 2005, 2007; Mimica 2016), we can determine its decay rate in terms of the convergence of its Laplace-Stieltjes transform:

$$d(V) = -\inf\{s \leq 0 \mid \mathcal{L}[V](s) < \infty\}. \quad (5.1)$$

The specific class of light-tailed job size distributions we consider, described below, are those that allow us to use (5.1) throughout this work (Online Appendix EC.3). The class includes essentially all light-tailed distributions of practical interest, such as finite-support, phase-type, and Gaussian-tailed distributions. In the terminology of Abate and Whitt (1997), we consider all "Class I" distributions.⁸

Definition 5.2 (Light-Tailed Job Size Distribution). Given a job size distribution X , let

$$s^* = \inf\{s \leq 0 \mid \mathcal{L}[X](s) < \infty\}.$$

We say that X is *nicely light-tailed* if $s^* = -\infty$ or $s^* \in (-\infty, 0)$ and $\mathcal{L}[X](s^*) = \infty$. In informal discussion, we omit "nicely."

Definition 5.3 (Tail Optimality in Light-Tailed Case). Consider an M/G/1 with nicely light-tailed job size distribution X . We say a scheduling policy π is

- *Log-tail-optimal* if π maximizes $d(T_\pi)$,
- *Log-tail-pessimal* if π minimizes $d(T_\pi)$, and
- *Log-tail-intermediate* otherwise.

In each case, we mean minimizing or maximizing over preemptive work-conserving policies. In informal discussion, we omit "log-."

5.2. Results for Light-Tailed Case

We have seen that a job's worst ever rank plays an important role in the heavy-tailed setting. When the job sizes are light-tailed, we are interested in the age at which the rank function's *global* maximum occurs.

Definition 5.4. The *worst age*, denoted a^* , is the earliest age at which a job has the global maximum rank:

$$a^* = \inf\{a \in [0, x_{\max}) \mid \forall b \in [0, x_{\max}), r(a) \geq r(b)\}.$$

If the rank function has no maximum, we define $a^* = x_{\max}$.

As an example, FCFS has $a^* = 0$ because a job's priority is the worst before it starts service. In contrast, FB has $a^* = x_{\max}$ because a job's priority gets strictly worse with age.

We already know that FCFS and FB are tail-optimal and tail-pessimal, respectively. The theorem below fills in the gaps for all other SOAP policies, showing that the performance of the response time tail is completely determined by the worst age a^* .

Theorem 5.5. Consider an $M/G/1$ with any nicely light-tailed job size distribution under a SOAP policy. Let $x_{\max} = \inf\{x \geq 0 \mid \mathbf{P}[X > x] = 0\}$. The policy is

- Log-tail-optimal if $a^* = 0$,
- Log-tail-intermediate if $0 < a^* < x_{\max}$, and
- Log-tail-pessimal if $a^* = x_{\max}$.

To apply Theorem 5.5 to the Gittins policy, we need to characterize how the job size distribution X affects Gittins's worst age a^* .

Definition 5.6. We define two classes of distributions: NBUE and ENBUE.

• We say X is *New Better than Used in Expectation*, writing $X \in \text{NBUE}$, if for all ages $a \in [0, x_{\max})$,

$$\mathbf{E}[X] \geq \mathbf{E}[X - a \mid X > a].$$

• We say X is *Eventually New Better than Used in Expectation*, writing $X \in \text{ENBUE}$, if there exists $a_0 \in [0, x_{\max})$ such that $(X - a_0 \mid X > a_0) \in \text{NBUE}$. Put another way, $X \in \text{ENBUE}$ if there exists $a_0 \in [0, x_{\max})$ such that for all $a \in [a_0, x_{\max})$,

$$\mathbf{E}[X - a_0 \mid X > a_0] \geq \mathbf{E}[X - a \mid X > a].$$

Remark 5.7. It is well known that the NBUE class includes all distributions with (weakly) increasing hazard rate. Similarly, one can show that the ENBUE class includes all distributions with "eventually increasing" hazard rate, meaning that the hazard rate is increasing at all ages greater than some threshold $a_0 < x_{\max}$.

Results of Aalto et al. (2009, 2011) connect the classes NBUE and ENBUE to Gittins's worst age a^* , implying the following characterization.

Theorem 5.8. Consider an $M/G/1$ with any nicely light-tailed job size distribution X . Gittins is

- Log-tail-optimal if $X \in \text{NBUE}$,
- Log-tail-intermediate if $X \in \text{ENBUE} \setminus \text{NBUE}$, and
- Log-tail-pessimal if $X \notin \text{ENBUE}$.

More generally, Theorem 5.5 can imply an analogue of Theorem 5.8 for other SOAP policies, whose rank at age a is related to the expected remaining size $\mathbf{E}[X - a \mid X > a]$, such as the SERPT policy (Remark 10.1).

The fact that Gittins can be log-tail-pessimal is intriguing, considering that it is optimal for mean response time and tail-optimal under heavy-tailed job sizes. Fortunately, in most cases where Gittins is log-tail-pessimal, slightly tweaking Gittins yields a log-tail-intermediate policy without sacrificing much mean response time performance.

Definition 5.9. A SOAP policy π is a q -approximate Gittins policy if there exists a constant $m > 0$ such that for all ages $a \in [0, x_{\max})$,

$$\frac{r_\pi(a)}{r_{\text{Gtn}}(a)} \in [m, mq].$$

We may assume without loss of generality that $m = 1$ because the policy π' with rank function $r_{\pi'}(a) = r_\pi(a)/m$ has identical behavior to policy π .

Theorem 5.10. Consider an $M/G/1$ with any job size distribution. For any $q \geq 1$ and any q -approximate Gittins policy π ,

$$\mathbf{E}[T_\pi] \leq q\mathbf{E}[T_{\text{Gtn}}].$$

An important observation is that a q -approximate Gittins policy has near-optimal mean response time for q close to one. At the same time, changing the Gittins rank function even within a small factor q can decrease the worst age, and therefore improve the tail performance.

Theorem 5.11. Consider an $M/G/1$ with nicely light-tailed job size distribution $X \notin \text{ENBUE}$. Suppose that the expected remaining size of a job at all ages is uniformly bounded, meaning


$$\sup_{a \in [0, x_{\max})} \mathbf{E}[X - a \mid X > a] < \infty.$$

Then, for all $\epsilon > 0$, there exists a $(1 + \epsilon)$ -approximate Gittins policy that is log-tail-optimal or log-tail-intermediate.

As an example in which log-tail-pessimality may be avoided, consider a hyperexponential job size X with two different rates. That is, for $i = 1, 2$, with probability p_i the job size is sampled from an exponential with rate μ_i . Here, $p_1, p_2 > 0$ such that $p_1 + p_2 = 1$, and we assume without loss of generality that $\mu_1 > \mu_2$.

Because the hazard rate of a hyperexponential distribution is decreasing, r_{Gtn} is increasing. Therefore, Gittins reduces to FB, so it is log-tail-pessimal. Fortunately, though, $\mathbf{E}[X - a \mid X > a] < 1/\mu_2$ for all ages a ; hence, Theorem 5.11 can be applied: for any $q > 1$, the policy with

Figure 3. (Color online) The Rank Functions of the Gittins Policy (Translucent Cyan Curve) and a q -approximate Gittins Policy (Dotted Yellow-Green Curve) for a Hyperexponential Distribution with Rates μ_1 and μ_2



Notes. By Theorem 5.5, these have different tail asymptotics. The Gittins rank function attains its supremum of $1/\mu_2$ in the $a \rightarrow \infty$ limit, so it is tail-pessimist. But the q -approximate Gittins rank function, described in (5.2), attains its supremum at a finite age \tilde{a} , so it is tail-intermediate.

rank function

$$r(a) = \begin{cases} qr_{Gtn}(a) & \text{if } a \leq \tilde{a}, \\ r_{Gtn}(a) & \text{otherwise,} \end{cases} \quad (5.2)$$

where \tilde{a} is such that $qr_{Gtn}(\tilde{a}) = 1/\mu_2$ is a q -approximate Gittins policy with $a^* = \tilde{a} < \infty$. See Figure 3.

Although Theorem 5.11 implies that any approximation factor $q > 1$ suffices to prevent log-tail-pessimality, there is a tradeoff to be made when choosing q . Higher values of q result in worse guarantees for the mean (Theorem 5.10). But higher values of q lead to lower values of a^* , which, in turn, result in better tail decay (Lemma 9.8). We leave exploring this tradeoff quantitatively to future work.

5.3. Proof Organization

The remainder of this paper is organized as follows. Necessary background and notation on the SOAP framework is given in Section 6. Then, we prove our results for heavy tails, Theorems 4.6 and 4.7, respectively, in Sections 7 and 8. Similarly, proofs of our results for the light-tailed case are given in Sections 9 (Theorem 5.5) and 10 (Theorems 5.8 and 5.11). The remaining main result, Theorem 5.10, requires substantially more technical machinery for its proof, which is why we defer it to Online Appendix EC.2. Finally, Section 11 describes how our results answer the questions posed in Section 1.

6. SOAP Notation

We use the following notations related to SOAP policies, which are standard in the literature (Scully and Harchol-Balter 2018, Scully et al. 2018, 2020a,b). These definitions are necessary for writing down and working with the response time formulas of SOAP policies. All of these definitions are given in terms of a SOAP policy with rank function r . Throughout, it will always be clear from context which rank function is being referred to.

Definition 6.1. The k th maximal w -interval is $(b_k[w], c_k[w])$, where for all $k \geq 1$ ¹⁰

$$\begin{aligned} b_0[w] &= 0, & b_k[w] &= \inf\{a > c_{k-1}[w] \mid r(a) \leq w\}, \\ c_0[w] &= \inf\{a \geq 0 \mid r(a) > w\}, \\ c_k[w] &= \inf\{a > b_k[w] \mid r(a) > w\}. \end{aligned}$$

Additionally, let $K[w]$ be the maximum k such that $b_k[w] < x_{\max}$. It may be that $K[w] = \infty$.

One may easily check that $(b_k[w], c_k[w])$ is indeed a w -interval, with only one exception: it may be that $b_0[w] = c_0[w] = 0$, in which case the interval is empty and, thus, not a w -interval. See Figure 1, whose shaded orange regions are specifically the zeroth, first, and second maximal w_x -intervals for the pictured size x .

Definition 6.2. The k th w -relevant job segment is the random variable

$$X_k[w] = \max\{0, \min\{X, c_k[w]\} - b_k[w]\}.$$

For convenience, we define $X_k[w] = 0$ for $k > K[w]$.

The following lemma gives a convenient formula for moments of relevant job segments.

Lemma 6.3 (Scully et al. 2020b, lemma 6.16).

$$\mathbb{E}[X_k[w]^{p+1}] = \int_{b_k[w]}^{c_k[w]} (p+1)(t - b_k[w])^p \bar{F}(t) dt.$$

Definition 6.4. For a job of size x , we define¹¹

$$y_x = c_0[w_x -], \quad z_x = c_0[w_x].$$

See Figure 4. Intuitively, y_x is the earliest age at which a job of size x attains its worst ever rank w_x , and z_x is the earliest age at which the rank function exceeds w_x .


Note that it may be that $y_x = x = z_x$. This occurs if the rank function is strictly increasing at x , and x is a “running maximum,” meaning $r(a) < r(x)$ for all ages $a \in [0, x)$.

Our final piece of notation is a generalization of a job’s worst ever rank.

Definition 6.5. The *worst future rank* of a job of size x at age a , written $w_x(a)$, is

$$w_x(a) = \sup_{a \leq b < x} r(b).$$

Figure 4. (Color online) Illustration of y_x and z_x (Definition 6.4)



Note that the worst ever rank is simply the worst future rank at age zero—that is, $w_x = w_x(0)$.

One can write a formula characterizing $T(x)$, the response time distribution of jobs of size x , in terms of the notation from Definitions 6.2 and 6.5. We refer the reader to Scully and Harchol-Balter (2018) and Scully et al. (2018) for details, noting here one simple consequence of that formula.

Lemma 6.6. *Under any SOAP policy, the response time of jobs of size x is stochastically increasing in x . That is, for all $x_1 \geq x_0 > 0$, we have $T(x_0) \leq_{st} T(x_1)$.*

7. Heavy-Tailed Job Sizes: Tail Asymptotics of SOAP Policies

This section is devoted to proving Theorem 4.6, which gives a sufficient condition for a SOAP policy to be tail-optimal. Our first step is to invoke a result of Scully et al. (2020b).¹² Their result reduces the task of proving tail optimality to a much simpler task—namely, bounding various functions of moments of w_x -relevant job segments. To state their result, we use a “polynomially strict” version of little- o notation, which we write as $\check{o}(\cdot)$.

Definition 7.1. For $p > 0$, the notation $\check{o}(x^p)$ stands for $O(x^{p-\epsilon})$ for some unspecified $\epsilon > 0$.

Condition 7.2. *There exists $q > \beta$ such that for all $p \in (0, q)$,*

$$\sum_{k=0}^{K[w_x]} \mathbf{E}[X_k[w_x]^{p+1}] \leq \check{o}(x^p). \quad (7.1)$$

Condition 7.3.

$$\int_0^x \frac{1}{1 - \lambda \mathbf{E}[X_0[w_x(a)]]} da \geq \frac{x}{1 - \rho} - \check{o}(x).$$

Lemma 7.4 (Scully et al. 2020c). *Consider an M/G/1 with nicely heavy-tailed job size under a SOAP policy. Conditions 7.2 and 7.3 together imply tail optimality.*

Our proof of Theorem 4.6 is thus based on verifying Conditions 7.2 and 7.3.

Proof of Theorem 4.6. By Lemma 7.4, it suffices to show that Condition 4.5 and (4.1) together imply Conditions 7.2 and 7.3. We do so in Proposition 7.5 (see Section 7.1), which handles Condition 7.2, and Proposition 7.8 (see Section 7.2), which handles Condition 7.3. \square

7.1. Showing Condition 7.2

Our goal in this section is to prove the following.

Proposition 7.5. *If (4.1) holds, then Condition 4.5 implies Condition 7.2.*

That is, we want to show (7.1) under certain conditions. The first step is to compute the left-hand side of

(7.1), while assuming only Condition 4.5. The following lemma does so, separating out the $k = 0$ term because it has a slightly different form.

Lemma 7.6. *Suppose Condition 4.5 holds.*

i. *For all $p \geq 0$,*

$$\mathbf{E}[X_0[w_x]^{p+1}] \leq \begin{cases} O(1) & \text{if } p < \alpha - 1 \\ O(\log x) & \text{if } p = \alpha - 1 \\ O(x^{\max\{1, \zeta + \theta\}(p - \alpha + 1)}) & \text{if } p > \alpha - 1. \end{cases}$$

ii. *For all $p \geq 0$,*

$$\sum_{k=1}^{K[w_x]} \mathbf{E}[X_k[w_x]^{p+1}] \leq \begin{cases} O(x^{\theta p + \zeta p - \alpha + 1}) & \text{if } \zeta p < \alpha - 1 \\ O(x^{\theta p} \log x^\eta) & \text{if } \zeta p = \alpha - 1 \\ O(x^{\theta p + \eta(\zeta p - \alpha + 1)}) & \text{if } \zeta p > \alpha - 1. \end{cases}$$

The proof of Lemma 7.6 is largely computational, so we defer it to Online Appendix EC.1.1.

Proof of Lemma 7.5. Our goal is to choose $q > \beta$ such that for all $p \in (0, q)$, (7.1) holds. Specifically, we choose

$$q = \begin{cases} \frac{\alpha - 1}{d} & \text{if } d > 0 \\ \infty & \text{if } d \leq 0, \end{cases}$$

where

$$d = \text{left-hand side of (4.1)} = \zeta + (\theta - 1)^+ - \frac{(1 - \theta)^+}{\eta}.$$

We have $q > \beta$ by (4.1), and an analogue of (4.1) holds for all $p \in (0, q)$ —namely,

$$\zeta + (\theta - 1)^+ - \frac{(1 - \theta)^+}{\eta} < \frac{\alpha - 1}{p}. \quad (7.2)$$

By Lemma 7.6, it suffices to show that for all $p \in (0, q)$, both of the following hold:

$$p > \begin{cases} 0 & \text{if } p < \alpha - 1 \\ 0 & \text{if } p = \alpha - 1 \\ \max\{1, \zeta + \theta\}(p - \alpha + 1) & \text{if } p > \alpha - 1, \end{cases}$$

[see Lemma 7.6(i)]

$$p > \begin{cases} \theta p + \zeta p - \alpha + 1 & \text{if } \zeta p < \alpha - 1 \\ \theta p & \text{if } \zeta p = \alpha - 1 \\ \theta p + \eta(\zeta p - \alpha + 1) & \text{if } \zeta p > \alpha - 1. \end{cases}$$

[see Lemma 7.6(ii)]

Under the assumptions of Definition 4.1 and Condition 4.5, namely,

$$\alpha > 1, \quad \zeta > 0, \quad \theta > 0, \quad \text{and} \quad \eta \geq \max\{1, \zeta + \theta\}, \quad (7.3)$$

this is equivalent to showing both of

$$\max\{1, \zeta + \theta\} \left(1 - \frac{\alpha - 1}{p}\right)^+ < 1, \quad (7.4)$$

$$\theta - \left(\frac{\alpha - 1}{p} - \zeta\right)^+ + \eta \left(\zeta - \frac{\alpha - 1}{p}\right)^+ < 1. \quad (7.5)$$

We show below that both of these are implied by (7.2). To reduce clutter, let

$$v = \frac{\alpha - 1}{p}.$$

For (7.4), we compute

$$\begin{aligned} (7.4) &\Leftrightarrow (1 - v)^+ < \frac{1}{\max\{1, \zeta + \theta\}} \\ &\Leftrightarrow [\zeta + \theta \leq 1] \vee \left([\zeta + \theta > 1] \wedge \left[\frac{\zeta + \theta - 1}{\zeta + \theta} < v \right] \right) \\ &\hspace{15em} [\text{by (7.3)} \Rightarrow v > 0] \\ &\Leftrightarrow [\zeta + \theta \leq 1] \vee \left([\zeta + \theta > 1] \right. \\ &\quad \left. \wedge \left[\frac{\zeta + (\theta - 1)^+ - (1 - \theta)^+}{\zeta + \theta} < v \right] \right) \\ &\Leftrightarrow [\zeta + \theta \leq 1] \vee \left([\zeta + \theta > 1] \right. \\ &\quad \left. \wedge \left[\zeta + (\theta - 1)^+ - \frac{(1 - \theta)^+}{\eta} < v \right] \right) \\ &\hspace{15em} [\text{by (7.3)} \Rightarrow \eta \geq \zeta + \theta] \\ &\Leftrightarrow (7.2), \end{aligned}$$

and for (7.5), we compute

$$\begin{aligned} (7.5) &\Leftrightarrow ([\zeta < v] \wedge [\theta - v + \zeta < 1]) \\ &\quad \vee ([\zeta \geq v] \wedge [\theta - \eta v + \eta \zeta < 1]) \\ &\Leftrightarrow ([\zeta < v] \wedge [\zeta + \theta - 1 < v]) \\ &\quad \vee \left([\zeta \geq v] \wedge \left[\zeta - \frac{1 - \theta}{\eta} < v \right] \right) \\ &\hspace{15em} [\text{by (7.3)} \Rightarrow \eta > 0] \\ &\Leftrightarrow ([\theta \geq 1] \wedge [\zeta + \theta - 1 < v]) \\ &\quad \vee \left([\theta < 1] \wedge \left[\zeta - \frac{1 - \theta}{\eta} < v \right] \right) \\ &\Leftrightarrow (7.2). \quad \square \end{aligned}$$

Remark 7.7. Note that in addition to (7.2) implying (7.5), the reverse implication also holds. This suggests

that the precondition of Theorem 4.6, and in particular (4.1), cannot be easily relaxed.

7.2. Showing Condition 7.3

Our goal in this section is to prove Proposition 7.8 below. It is applicable to proving Theorem 4.6 because (4.1) implies its precondition.

Proposition 7.8. *If $\zeta < 1$ or $\eta < \infty$, then Condition 4.5 implies Condition 7.3.*

Proof. The $\eta < \infty$ case follows from a result of (Scully et al. 2020b, lemma 7.3), so we address only the $\zeta < 1$ case. We first observe that for all $\rho' \in [0, \rho]$, we have $\frac{1}{1-\rho'} \geq \frac{1}{1-\rho} - \frac{\rho-\rho'}{(1-\rho)^2}$. This means that Condition 7.3 holds if

$$\int_0^x (\mathbf{E}[X] - \mathbf{E}[X_0[w_x(a)-]]) da \leq \check{\delta}(x). \quad (7.6)$$

We can rewrite the integrand as

$$\begin{aligned} &\mathbf{E}[X] - \mathbf{E}[X_0[w_x(a)-]] \\ &= \int_0^\infty \bar{F}(t) dt - \int_0^{c_0[w_x(a)-]} \bar{F}(t) dt \quad [\text{by Lemma 6.3}] \\ &\leq \int_{c_0[w_x(a)-]}^\infty O(t^{-\alpha}) dt \quad [\text{by Definition 4.1}] \\ &\leq O(c_0[w_x(a)-]^{-(\alpha-1)}). \end{aligned}$$

Of course, the integrand is also bounded above by $\mathbf{E}[X]$, so

$$\begin{aligned} &\int_0^x (\mathbf{E}[X] - \mathbf{E}[X_0[w_x(a)-]]) da \\ &\leq \int_0^x O(\min\{1, c_0[w_x(a)-]^{-(\alpha-1)}\}) da. \quad (7.7) \end{aligned}$$

A job of size x attains its worst ever rank w_x at age y_x . This means that for all $a < y_x$, we have $w_x(a) = w_x$, which by Definition 6.4 implies $c_0[w_x(a)-] = y_x$. Splitting the integral in (7.7) at $a = y_x$ yields

$$\begin{aligned} &\int_0^x (\mathbf{E}[X] - \mathbf{E}[X_0[w_x(a)-]]) da \\ &\leq \int_0^{y_x} O(y_x^{-(\alpha-1)}) da + \int_{y_x}^x O(\min\{1, c_0[w_x(a)-]^{-(\alpha-1)}\}) da \\ &\leq y_x^{(2-\alpha)^+} + \int_{y_x}^x O(\min\{1, c_0[w_x(a)-]^{-(\alpha-1)}\}) da. \quad (7.8) \end{aligned}$$

Because $y_x \leq x$ and $\alpha > 1$, it suffices to show that the integral in (7.8) is $\check{\delta}(x)$.

The main remaining obstacle is bounding $c_0[w_x(a)-]$. We do so in Lemma 7.9, which states

$$c_0[w_x(a)-] \geq \Omega\left(\left(\frac{x-a}{x^\zeta}\right)^{1/\kappa}\right),$$

for some $\kappa \geq 2(\alpha - 1)$. Plugging this into (7.8) and substituting $u = x^{-\zeta}(x - a)$ gives

$$\begin{aligned} & \int_{y_x}^x O(\min\{1, c_0[w_x(a)-]^{-(\alpha-1)}\}) da \\ & \leq \int_0^x O\left(\min\left\{1, \left(\frac{x-a}{x^\zeta}\right)^{-(\alpha-1)/\kappa}\right\}\right) da \\ & \leq \int_{x-x^\zeta}^x O(1) da + x^\zeta \int_1^{x^{1-\zeta}} O(u^{-(\alpha-1)/\kappa}) du \\ & = O(x^\zeta) + O(x^{\zeta+(1-\zeta)(1-(\alpha-1)/\kappa)}). \end{aligned}$$

Because $\zeta < 1$ and $(\alpha - 1)/\kappa \in (0, 1/2]$, this is $\check{o}(x)$, as desired. \square

The following lemma bounds $c_0[w_x(a)-]$. We defer its proof to Online Appendix EC.1.1.

Lemma 7.9. *Suppose Condition 4.5 holds, and let $\kappa = 2 \max\{\alpha - 1, \theta\}$. For all $x \geq 0$ and $a \in (y_x, x)$,*

$$c_0[w_x(a)-] \geq \Omega\left(\left(\frac{x-a}{x^\zeta}\right)^{1/\kappa}\right).$$

8. Heavy-Tailed Job Sizes: Gittins Is Tail-Optimal

In this section, we prove Theorem 4.7—namely, that Gittins is tail-optimal for heavy-tailed job sizes. Specifically, we will show that Gittins satisfies Condition 4.5 with

$$\zeta = 0, \quad \theta = 1, \quad \text{and} \quad \eta = \infty.$$


This suffices because with the above values, (4.1) holds for all $\beta \geq \alpha > 1$, so Theorem 4.6 implies tail optimality. In fact, Condition 4.5 holds trivially when $\eta = \infty$, so only Condition 4.5 remains.

Our goal is thus show that for the Gittins rank function,¹³

$$\text{any right-maximal } w_x\text{-interval } (b, c) \text{ with } b \geq x \text{ has length } c - b \leq O(x). \quad (8.1)$$

In words, (8.1) says that whenever the Gittins rank function dips below the worst rank a job of size x ever

Figure 5. (Color online) Illustration of (8.1)



Note. Any w_x -interval (shaded orange regions)—namely, any interval where the Gittins rank function (solid cyan curve) is better than the worst ever rank w_x of a job of size x —has length $O(x)$.

has, it does so for an interval of length at most $O(x)$. See Figure 5.

8.1. The Time-Per-Completion Function

To make further progress, we need to consider the specific form of the Gittins rank function. One useful way of thinking about the Gittins rank function uses the following definition (Aalto et al. 2009, 2011).

Definition 8.1. The *time-per-completion* function for a given job size distribution S is defined for $0 \leq b < c \leq x_{\max}$ as

$$\phi(b, c) = \frac{\mathbb{E}[\min\{S, c\} - b | S > b]}{\mathbb{P}[S \leq c | S > b]} = \frac{\int_b^c \bar{F}(t) dt}{\bar{F}(b) - \bar{F}(c)}.$$

The intuition is that if we serve a job from age b until either age c or its completion, whichever comes first, then we can interpret $\phi(b, c)$ as

$$\phi(b, c) = \frac{\mathbb{E}[\text{time spent serving job from age } b \text{ to age } c]}{\mathbb{E}[\mathbb{1}(\text{job completes when served from age } b \text{ to age } c)]},$$

which is an expected amount of time over an expected number of completions.

One can rewrite the Gittins rank function (Definition 3.2) in terms of $\phi(b, c)$ as

$$r_{\text{Gtn}}(a) = \inf_{c > a} \phi(a, c). \quad (8.2)$$

That is, under Gittins, a job's rank at age a is the best possible time-per-completion ratio achievable on any interval starting at age a .

How does using the time-per-completion function help us show (8.1)? The key observation is that for the Gittins rank function to be low over an interval (b, c) , a job must be relatively likely to complete during (b, c) , which would imply that $\phi(b, c)$ is also low. The following lemma, which we prove in Online Appendix EC.2.2, formalizes this intuition.

Lemma 8.2. *Under Gittins, for any right-maximal w -interval (b, c) , we have $\phi(b, c) \leq w$.*

We note that although many results similar to Lemma 8.2 have been shown in prior work (Aalto et al. 2009, 2011; Scully et al. 2020b), to the best of our knowledge, Lemma 8.2 itself is new.

With Lemma 8.2 in hand, showing (8.1) amounts to

- Proving an upper bound on w_x , and
- Proving a lower bound on $\phi(b, c)$ for all right-maximal w_x -intervals (b, c) .

The first of these follows simply from prior work: Scully et al. (2020b, section 3.2) show that the Gittins rank function is bounded by $r_{\text{Gtn}}(a) \leq O(a)$, which implies

$$w_x = \sup_{0 \leq a < x} r_{\text{Gtn}}(a) \leq \sup_{0 \leq a < x} O(a) \leq O(x). \quad (8.3)$$

It thus remains only to bound $\phi(b, c)$ below. We begin with a bound on $\phi(b, c)$ from prior work.

Lemma 8.3 (Scully et al. 2021, lemma 6.8). *For any nicely heavy-tailed job size distribution and all $c > b \geq 0$, the time-per-completion function is bounded by*

$$\phi(b, c) \geq \Omega\left(\frac{b}{c}(c - b)\right).$$

Combining (8.3) and Lemmas 8.2 and 8.3, we find that for any right-maximal w_x -interval (b, c) ,

$$\begin{aligned} c - b &\leq \frac{c}{b} O(\phi(b, c)) && \text{[by Lemma 8.3]} \\ &\leq \frac{c}{b} O(w_x) && \text{[by Lemma 8.2]} \\ &\leq \left(1 + \frac{c - b}{b}\right) O(x). && \text{[by (8.3)]} \end{aligned}$$

Therefore, to show (8.1) and thereby Theorem 4.7, it suffices to show that

$$\begin{aligned} &\text{any right-maximal } w_x\text{-interval } (b, c) \\ &\text{with } b \geq x \text{ has length } c - b \leq O(b). \end{aligned} \quad (8.4)$$

We could equivalently write $c \leq O(b)$, but the $c - b \leq O(b)$ form emphasizes the progress we have made relative to (8.1): we have weakened our goal from proving an $O(x)$ bound to an $O(b)$ bound.

8.2. Bounding Lengths of w_x -Intervals

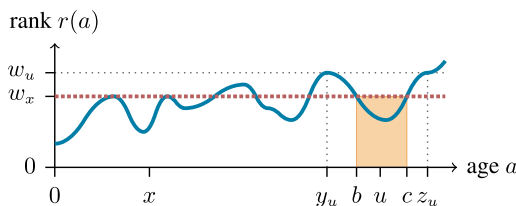
There is one more fact from prior work that we need to prove (8.4).

Lemma 8.4 (Scully et al. 2021, theorem 6.4). *Under Gittins, $y_x = \Theta(x)$ and $z_x = \Theta(x)$.*

In general, (y_x, z_x) is a right-maximal w_x -interval, but there can be plenty of other w_x -intervals starting at values greater than z_x . How can we use Lemma 8.4 to study such intervals? The key is to look not at y_x and z_x , but at y_u and z_u , where u is some point inside the w_x -interval whose length we wish to bound.

Specifically, consider a right-maximal w_x -interval (b, c) with $b \geq x$ and let $u \in (b, c)$ be an arbitrary size in the interval. Figure 6 illustrates the relationship

Figure 6. (Color online) A w_x -Interval (b, c) Contained Inside (y_u, z_u) , Where $u \in (b, c)$ Is an Arbitrary Point in the w_x -Interval



between the interval (b, c) , the worst ever rank w_u of size u , and the ages y_u and z_u (Definition 6.4). The figure suggests the following lemma, which is a slight generalization of a result of Scully et al. (2020b, lemma 6.18).

Lemma 8.5. *Under any SOAP policy, for any w_x -interval (b, c) with $b \geq x$ and any $u \in (b, c)$,*

$$y_u \leq b < c \leq z_u.$$

Proof. It is clear from Definition 6.4 that $y_u \leq u \leq z_u$. It thus suffices to show that $y_u \notin (b, c)$ and $z_u \notin (b, c)$. Both steps use the fact that $x \leq b < u$ implies $w_x \leq w_u$ (Definition 4.3).

We first show that $z_u \notin (b, c)$. By Definition 4.4, the rank function does not exceed w_x over the interval (b, c) . But Definition 6.4 implies that every neighborhood of z_u contains a point whose rank is greater than w_u . Because $w_u \geq w_x$, it must be that $z_u \notin (b, c)$.

If $w_u > w_x$, then the argument that $y_u \notin (b, c)$ is analogous to that for z_u . The difference is that this time, Definition 6.4 implies that for all $\epsilon > 0$, every neighborhood of y_u contains a point whose rank is greater than $w_u - \epsilon$. Because $w_u - \epsilon > w_x$ for small enough ϵ , it must be that $y_u \notin (b, c)$.

If, instead, $w_u = w_x$, then $y_u = y_x$, and we can reason more simply: Definition 6.4 implies $y_x \leq x$, and we have assumed $x \leq b$, so $y_u \leq b$. \square

Lemmas 8.4 and 8.5 combine to prove (8.4), from which Theorem 4.7 follows.

Proof of Theorem 4.7. Let (b, c) be a right-maximal w_x -interval with $b \geq x$, and let $u \in (b, c)$. We compute

$$\begin{aligned} c - b &\leq \frac{z_u}{y_u} \cdot \frac{b}{c} (c - b) && \text{[by Lemma 8.5]} \\ &\leq O(1) \cdot \frac{b}{c} (c - b) && \text{[by Lemma 8.4]} \\ &\leq O(\phi(b, c)) && \text{[by Lemma 8.3]} \\ &\leq O(w_x) && \text{[by Lemma 8.2]} \\ &\leq O(x). && \text{[by (8.3)]} \end{aligned}$$

The fact that $c - b \leq O(x)$ means that Gittins satisfies Condition 4.5 with $\zeta = 0$, $\theta = 1$, and $\eta = \infty$. These obey (4.1), so by Theorem 4.6, Gittins is tail-optimal. \square

9. Light-Tailed Job Sizes: Tail Asymptotics of SOAP Policies

In this section, we prove our main theorem for light-tailed job sizes.

Theorem 5.5. *Consider an $M/G/1$ with any nicely light-tailed job size distribution under a SOAP policy. Let $x_{\max} = \inf\{x \geq 0 \mid \mathbf{P}[X > x] = 0\}$. The policy is*

- Log-tail-optimal if $a^* = 0$,

- Log-tail-intermediate if $0 < a^* < x_{\max}$, and
- Log-tail-pessimal if $a^* = x_{\max}$.

Proof. The result follows from Propositions 9.1, 9.2, and 9.9, which we prove in the rest of this section. \square

9.1. Tail-Optimal Case

Proposition 9.1. *Consider an M/G/1 with any nicely light-tailed job size distribution under a SOAP policy. The policy is log-tail-optimal if $a^* = 0$.*

Proof. Suppose that $a^* = 0$. In this case, because of the FCFS tiebreaking (Definition 3.1), the oldest job in the system always has priority over all other jobs. Therefore, the SOAP policy is exactly FCFS, which is known to be log-tail-optimal (Stolyar and Ramanan 2001, Boxma and Zwart 2007). \square

9.2. Tail-Pessimal Case

Mandjes and Nuyens (2005) show that the FB policy, which has rank function $r(a) = a$, is log-tail-pessimal. Their argument focuses on analyzing the response time of very large jobs, showing that such jobs' response time tails have small decay rate. The fact that their argument focuses on just the very large jobs suggests that the essential property of FB is that it assigns the largest rank at the largest ages. Our tail pessimality result below shows this is indeed the case.

Proposition 9.2. *Consider an M/G/1 with any nicely light-tailed job size distribution under a SOAP policy. The policy is log-tail-pessimal if $a^* = x_{\max}$.*

We prove Proposition 9.2 in Online Appendix EC.1.2 by generalizing the proof that Mandjes and Nuyens (2005) give for FB, making use of several of their intermediate results along the way. We describe the approach below, after which we present the proof.

Consider first the case where $x_{\max} < \infty$, meaning we could have a job of size x_{\max} . Because $a^* = x_{\max}$, jobs of size x_{\max} complete at the end of the busy period during which they arrive. This is the latest time a job can complete under a work-conserving policy, so $d(T(x_{\max})) = d_{\min}$, where d_{\min} is the minimal, and thus pessimal, response time tail decay rate. If X has an atom at x_{\max} , meaning that $X = x_{\max}$ occurs with positive probability, then $d(T) = d(T(x_{\max})) = d_{\min}$.

Of course, for general job size distributions X , it may be that $X = x_{\max}$ occurs with probability zero. This is certainly the case if $x_{\max} = \infty$. Therefore, to generalize the argument above, instead of considering $T(x_{\max})$, we consider $T(x)$ in the $x \rightarrow x_{\max}$ limit. Although $d(T(x)) > d_{\min}$ for any fixed $x < x_{\max}$, we show that $\lim_{x \rightarrow x_{\max}} d(T(x)) = d_{\min}$. This means that for any $\epsilon > 0$, a positive fraction of jobs experience decay rate less than $d_{\min} + \epsilon$, implying $d(T) = d_{\min}$, as desired.

The last ingredient we need is notation for discussing the M/G/1 and, in particular, busy periods. This is because d_{\min} turns out to be a busy period's decay rate (Mandjes and Nuyens 2005, corollary 6).

Definition 9.3.

i. We denote by B the distribution of an M/G/1 busy period length with arrival rate λ and job size distribution X . More generally, we write $B(u)$ for a busy period with initial work u .

ii. We denote by B_a the distribution of an M/G/1 busy period length with arrival rate λ and truncated job size distribution $\min\{X, a\}$. More generally, we write $B_a(u)$ for a such busy period with initial work u .

iii. We denote by W the distribution of the total amount of work in an M/G/1.

It is known that $d_{\min} = d(B)$ (Mandjes and Nuyens 2005, corollary 6), so proving Proposition 9.2 amounts to showing $\lim_{x \rightarrow x_{\max}} d(T(x)) = d(B)$. To show this, we first prove $d(T(x)) \leq d(B_{y_x})$, so it suffices to show $\lim_{x \rightarrow x_{\max}} d(B_{y_x}) = d(B)$. This follows from the known fact that $\lim_{y \rightarrow x_{\max}} d(B_y) = d(B)$ (Mandjes and Nuyens 2005, proposition 8) and additional computation. See Online Appendix EC.1.2 for details.

9.3. Tail-Intermediate Case

We finally turn to the case where $0 < a^* < x_{\max}$, where we will show that the corresponding SOAP policy is log-tail-intermediate. We first simplify the problem of analyzing an arbitrary SOAP policy with $0 < a^* < x_{\max}$ to the problem of analyzing two policies, called *Step* and *Spike* (Definition 9.4 and Figure 7), with similar rank functions. We then bound the decay rates of Step and Spike.


For brevity, we give only the key definitions and lemma structure of the proof, building up to Proposition 9.9, which states the main result for the tail-intermediate case. The proofs of the individual steps are either largely computational or follow easily from prior work, so we defer the proofs to Online Appendix EC.1.2.

Definition 9.4. For a given value of $a^* \in (0, x_{\max})$ and $r^* > 0$, the *Step* and *Spike* policies are the SOAP policies given by the following rank functions, which are illustrated in Figure 7:

$$r_{\text{step}}(a) = r^* \mathbb{1}(a \geq a^*), \quad r_{\text{spike}}(a) = r^* \mathbb{1}(a = a^*).$$

We will compare the response time tail of a SOAP policy with $0 < a^* < x_{\max}$ and $r(a^*) = r^*$ to the response time tails of Step and Spike for those values of a^* and r^* . This comparison is possible because, as illustrated in Figure 7, all three policies have qualitatively similar rank functions.¹⁴

Figure 7. (Color online) Rank Functions of Multiple Policies for a Given Worst Age a^* and Maximum Rank r^* , Where $0 < a^* < x_{\max}$



Notes. (a) Rank of Step. (b) Rank of Spike. (c) Rank of generic policy. As shown in Lemma 9.5, for a fixed job size distribution and value of a^* , (a) the Step policy gives the worst possible tail decay, whereas (b) the Spike policy gives the best possible tail decay. This is because in (a), jobs of age a^* or greater have the worst possible rank r^* , whereas in (b), jobs of age a^* or greater have the best possible rank zero. Under (c), a generic SOAP policy π , jobs of age a^* or greater have rank somewhere in between.

For the remainder of this section, we consider SOAP policies π with $0 < a^* < x_{\max}$ and $r(a^*) = r^*$. We divide jobs in two classes:

- *Class 1*, jobs of size at most a^* ; and
- *Class 2*, jobs of size greater than a^* .

For each class $i \in \{1, 2\}$, let $\lambda^{(i)}$ be the arrival rate of class i jobs, and let $T_\pi^{(i)}$ be the response time of class i jobs under policy π .

It turns out that only class 2 jobs affect the asymptotic decay rate of response time. Moreover, it turns out that the Step and Spike policies represent the worst-case and best-case scenarios for class 2 jobs. These facts, expressed in the following lemma, imply that it suffices to analyze the response time decay rate for class 2 jobs under Step and Spike.

Lemma 9.5. *Let π be a SOAP policy with $0 < a^* < x_{\max}$. We have*

$$d(T_\pi) = d(T_\pi^{(2)}) \in [d(T_{\text{step}}^{(2)}), d(T_{\text{spike}}^{(2)})].$$

With Lemma 9.5 in hand, to show that π is tail-intermediate, it suffices to show that both Step and Spike are tail-intermediate by bounding $d(T_{\text{step}}^{(2)})$ and $d(T_{\text{spike}}^{(2)})$. We begin by characterizing $T_{\text{step}}^{(2)}$ and $T_{\text{spike}}^{(2)}$ in terms of the M/G/1 concepts defined in Definition 9.3.

Lemma 9.6. *The response time distributions of class 2 jobs under Step and Spike are*

$$T_{\text{step}}^{(2)} =_{\text{st}} B_{a^*}(W) + B_{a^*}(X^{(2)}),$$

$$T_{\text{spike}}^{(2)} =_{\text{st}} B_{a^*}(W) + B_{a^*}(a^*) + X^{(2)} - a^*,$$

where $X^{(2)} = (X|X > a^*)$ is the size distribution of class 2 jobs, and the random variables in each sum are mutually independent.

Combining Lemmas 9.5 and 9.6 reduces the question of analyzing the decay rate of π to analyzing the decay rate of the busy periods in Lemma 9.6. We will see soon that $B_{a^*}(W)$ is the dominant term, so both Step and Spike have decay rate $d(B_{a^*}(W))$. In order to prove that

$B_{a^*}(W)$ is indeed the dominant term, and in order to bound its decay rate, we make heavy use of Laplace-Stieltjes transforms (Section 5.1).

Recall from (5.1) that one can determine a random variable's decay rate by determining when its Laplace-Stieltjes transform converges.¹⁵ We therefore introduce notation to describe when a transform converges. For functions $f : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$, which diverge below a certain value and converge above it, define

$$\gamma(f) = \sup\{s \in \mathbb{R} : |f(s)| = \infty\} = \inf\{s \in \mathbb{R} : |f(s)| < \infty\},$$

to be the value at which f switches from diverging to converging. We can thus rewrite (5.1) as

$$d(V) = -\gamma(\mathcal{L}[V]). \tag{9.1}$$

Because we will be working with Laplace-Stieltjes transforms, we begin by recalling standard results for the transforms of the quantities defined in Definition 9.3. Let

$$\sigma^{-1}(s) = s - \lambda(1 - \mathcal{L}[X](s)),$$

$$\sigma_a^{-1}(s) = s - \lambda(1 - \mathcal{L}[\min\{X, a\}](s)).$$

We can define a partial inverse σ of σ^{-1} such that $\sigma(s)$ is the greatest real solution to

$$\sigma(s) = s + \lambda(1 - \mathcal{L}[X](\sigma(s))),$$

letting $\sigma(s) = -\infty$ if there is no real solution. We define σ_a similarly.


Standard M/G/1 results (Harchol-Balter 2013) express the Laplace-Stieltjes transforms of random variables in Definition 9.3 using σ and σ_a . Specifically, for any non-negative random variable U ,

$$\mathcal{L}[B(U)](s) = \mathcal{L}[U](\sigma(s)), \quad \mathcal{L}[B_a(U)](s) = \mathcal{L}[U](\sigma_a(s)),$$

$$\mathcal{L}[W](s) = \frac{s(1 - \rho)}{\sigma^{-1}(s)}. \tag{9.2}$$

Therefore, to understand the decay rates of random variables in Definition 9.3, we need to understand σ^{-1} , σ , and σ_a . Figure 8 illustrates σ and σ^{-1} and some key

Figure 8. (Color online) Illustration of σ^{-1} (Solid Green Curve) and Key Values Associated with It



Notes. The partial inverse of σ^{-1} —namely, σ —corresponds to the branch going from the minimum of σ^{-1} to the right (orange highlight). As shown in Lemma 9.7, we have $s_0 < s_1 < s_2 < s_3$.

values associated with them. We make frequent use of relationships between these values and other properties of σ and σ^{-1} , which are summarized below.

Lemma 9.7. Consider an $M/G/1$ with nicely light-tailed job size distribution X , and define

$$s_0 = \gamma(\sigma^{-1}) = \gamma(\mathcal{L}[X]),$$

$$s_2 = \sigma(\gamma(\sigma)) = \arg \min_{s \geq s_0} \sigma^{-1}(s),$$

$$s_1 = \gamma(\mathcal{L}[W]) = \text{least root of } \sigma^{-1},$$

$$s_3 = \gamma(\sigma) = \min_{s \geq s_0} \sigma^{-1}(s).$$

Then, as illustrated in Figure 8, the following hold:

- i. σ^{-1} is convex on (s_0, ∞) , decreasing on (s_0, s_2) , and increasing on (s_2, ∞) ;
- ii. $s_0 < s_1 < s_2 < s_3 < 0$.
 Analogous statements hold for σ_a for all $a \in (0, x_{\max})$.

Together, (9.2) and Lemma 9.7 give us the last ingredients we need to compute the decay rate of T_π . We then show that this decay rate is neither optimal nor pessimal.

Lemma 9.8. Let π be a SOAP policy with $0 < a^* < x_{\max}$. Then, the decay rate of its response time is

$$d(T_\pi) = -\gamma(\mathcal{L}[W] \circ \sigma_{a^*}).$$

Proposition 9.9. Consider an $M/G/1$ with any nicely light-tailed job size distribution under a SOAP policy. The policy is log-tail-intermediate if $0 < a^* < x_{\max}$.

10. Light-Tailed Job Sizes: Gittins Can Be Tail-Optimal, Tail-Pessimal, or in Between

Theorem 5.8. Consider an $M/G/1$ with any nicely light-tailed job size distribution X . Gittins is

- Log-tail-optimal if $X \in \text{NBUE}$,
- Log-tail-intermediate if $X \in \text{ENBUE} \setminus \text{NBUE}$, and
- Log-tail-pessimal if $X \notin \text{ENBUE}$.

Proof. By Theorem 5.5, it suffices to determine the worst age a^* . Combining the following two prior results characterizes a^* in terms of whether X is in each of NBUE and ENBUE.

- A result of Aalto et al. (2009, proposition 7) implies $a^* = 0$ if and only if $X \in \text{NBUE}$.
- A result of Aalto et al. (2011, proposition 9) implies $a^* < x_{\max}$ if and only if $X \in \text{ENBUE}$ \square

Theorem 5.11. Consider an $M/G/1$ with nicely light-tailed job size distribution $X \notin \text{ENBUE}$. Suppose that the expected remaining size of a job at all ages is uniformly bounded, meaning

$$\sup_{a \in [0, x_{\max})} \mathbf{E}[X - a | X > a] < \infty.$$

Then, for all $\epsilon > 0$, there exists a $(1 + \epsilon)$ -approximate Gittins policy that is log-tail-optimal or log-tail-intermediate.

Proof. Suppose $\mathbf{E}[X - a | X > a]$ is uniformly bounded. Definition 3.2 implies

$$r_{\text{Gtn}}(a) \leq \frac{\int_a^\infty \bar{F}(t) dt}{\bar{F}(a)} = \mathbf{E}[X - a | X > a],$$

so $r_{\text{Gtn}}(a)$ is also uniformly bounded. This means that for any $\epsilon > 0$, there exists some sufficiently large age $a(\epsilon)$ such that increasing the rank at age $a(\epsilon) < x_{\max}$ from $r_{\text{Gtn}}(a(\epsilon))$ to $(1 + \epsilon)r_{\text{Gtn}}(a(\epsilon))$ and leaving all other ranks unchanged yields a new SOAP policy with worst age $a^* = a(\epsilon)$. By construction, the new policy is a $(1 + \epsilon)$ -approximate Gittins policy, and because its worst age is $a^* < x_{\max}$, Theorem 5.5 implies it is log-tail-optimal or log-tail-intermediate. \square

Recall from Theorem 5.10 that a $(1 + \epsilon)$ -approximate Gittins policy achieves mean response time within a factor of $1 + \epsilon$ of optimal. We defer its proof to Online Appendix EC.2. This means that Theorem 5.11, whose precondition applies to nonpathological light-tailed job size distributions, gives a non-tail-pessimal policy with near-optimal mean response time.

Remark 10.1. The *Shortest Expected Remaining Processing Time* (SERPT) policy, which has rank function $r_{\text{SERPT}}(a) = \mathbb{E}[X - a | X > a]$, is sometimes considered as a simpler alternative to Gittins (Scully and Harchol-Balter 2018, Scully et al. 2018, 2020a). Our results imply that SERPT has the same M/G/1 tail optimality properties as Gittins for the class of job size distributions we consider.

- Scully et al. (2020c) show that SERPT is always tail-optimal in the heavy-tailed case, which matches what we show for Gittins.

- Theorem 5.5 and Definition 5.6 imply that in the light-tailed case, SERPT is tail-optimal, tail-intermediate, and tail-pessimal under the same conditions as we show for Gittins in Theorem 5.8. In fact, one can show a stronger property: SERPT's and Gittins's response time distributions have the same decay rate. This follows from the fact that SERPT and Gittins have the same worst age a^* , as argued in the proof of Theorem 5.8.

11. Conclusion

In this paper, we have characterized the asymptotic tail performance of the response time in an M/G/1 queue under very broad conditions—namely, for every SOAP policy and for both heavy- and light-tailed job size distributions. In the heavy-tailed case, we characterize tail-optimal policies by a sufficient condition on the rank function (Theorem 4.6). This condition holds for a wide range of SOAP policies, and specifically for the Gittins policy (Theorem 4.7), providing the first proof of its tail optimality under general conditions. In the light-tailed case, we classify policies' performance as tail-optimal, tail-pessimal, or tail-intermediate. We show that the performance of a SOAP policy depends on the age at which the maximal rank is attained (Theorem 5.5). It turns out that the Gittins policy may belong to any of the three categories, depending on the job size distribution (Theorem 5.8). Finally, when Gittins has pessimal tail performance, boundedness of the expected remaining job size implies that there exists a slight modification of Gittins that has optimal or intermediate tail, while maintaining near-optimal mean response time (Theorem 5.11).

11.1. Returning to the Motivating Questions

We conclude by returning to Questions 1.1–1.3, restated below for convenience.

Question 1.1. Does any scheduling policy *simultaneously* optimize the mean and asymptotic tail of response time in the M/G/1?

Question 1.2. For which job size distributions is Gittins tail-optimal for response time?

Question 1.3. For job size distributions for which Gittins is tail-pessimal, is there another policy that has

near-optimal mean response time while not being tail-pessimal?

Our characterization of Gittins's tail asymptotics (Theorems 4.7 and 5.8) answers Question 1.2, and our modification in the case where Gittins is tail-pessimal (Theorem 5.11) answers Question 1.3 affirmatively. This leaves only Question 1.1. In cases where we have shown that Gittins is tail-optimal, the answer is clearly affirmative. We might hope to conclude that the answer is negative in cases where we have shown that Gittins is tail-pessimal or tail-intermediate, but the situation is still slightly unclear. The remaining ambiguity is due to the fact that we have only considered FCFS tiebreaking when two jobs have the same rank (Definition 3.1), as we explain in more detail below.

The Gittins policy minimizes mean response time with *arbitrary tiebreaking* between jobs of the same rank (Gittins 1989, Gittins et al. 2011, Scully and Harchol-Balter 2021). Moreover, these proofs can be extended to show that any “non-Gittins” policy is *strictly suboptimal* for mean response time, where a “non-Gittins” policy is one that for a nonvanishing fraction of time serves a job other than one of minimal Gittins rank. Therefore, to fully answer Question 1.1, one would have to consider Gittins under arbitrary tiebreaking rules. We conjecture that using a different tiebreaking rule cannot improve the asymptotic decay rate of Gittins's response time in the light-tailed case.

Finally, we note that we have, of course, only answered Questions 1.1–1.3 for the classes of heavy- and light-tailed job size distributions that we consider in this work (Definitions 4.1 and 5.2). Practically speaking, we believe the classes of distributions we consider are likely broad enough to draw useful conclusions. But it remains an open question whether we may extend our theory to broader classes of distributions. In particular, it seems likely that our proofs may hold mostly unchanged for additional light-tailed job size distributions, as discussed in Online Appendix EC.3.3.

Acknowledgments

The authors thank Adam Wierman, Onno Boxma, and Jan-Pieter Dorsman for helpful discussions. The authors also thank the anonymous referees for helpful comments that significantly improved the presentation. Z. Scully conducted this research in part while a graduate student at Carnegie Mellon University, in part while visiting the Simons Institute for the Theory of Computing, and in part while a Foundations of Data Science Institute postdoc at Harvard and MIT.

Endnotes

¹ Although we are referring to the job size distribution, we should clarify that here we are still discussing scheduling with known job sizes, specifically SRPT. But, starting shortly, we will shift attention to the case where job sizes are unknown, but we still know the job size distribution from which job sizes are drawn.

² During the review process, a generalization of this result was shown by Scully (2022, chapter 16).

³ In fact, both results are special cases of a more general result. See Scully (2022, chapter 16) for details.

⁴ The full SOAP definition is more general (Scully and Harchol-Balter 2018, Scully et al. 2018), but the given definition suffices for our unknown-size setting.

⁵ We give some brief remarks on other tiebreaking rules in Section 11.1.

⁶ We formally define the $O(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$ notations as follows. Suppose x_1, \dots, x_n are nonnegative variables. The notation $O(f(x_1, \dots, x_n))$ stands for an unspecified expression $g(x_1, \dots, x_n) \geq 0$ for which there exist constants $C, y_0, \dots, y_n > 0$ such that for all $x_1 \geq y_1, \dots, x_n \geq y_n$, we have $g(x_1, \dots, x_n) \leq Cf(x_1, \dots, x_n)$. The $\Omega(\cdot)$ notation is the same, but with the inequality reversed, and the $\Theta(\cdot)$ notation indicates that both inequalities hold, likely with different values for the constants. For all of these notations, the constants may depend on the job size distribution X .

⁷ This is trivially satisfied when $\eta = \infty$.

⁸ Our results can be generalized to some “Class II” distributions (Abate and Whitt 1997), which are also light-tailed. We comment on this in Online Appendix EC.3.3. However, working with Class II distributions generally requires additional regularity or smoothness assumptions (Abate and Whitt 1997, section 5), so for simplicity of presentation, we focus on Class I distributions.

⁹ This is a special case of a more general result (Scully 2022, chapter 16), which appeared while this work was in revision.

¹⁰ In the infimum expressions below, all sets of ages are implicitly assumed to be subsets of $[0, x_{\max}]$, and the infimum of an empty set is taken to be x_{\max} .

¹¹ Here and throughout, the postscript “–” denotes the limit from the left of a right-continuous function with left limits. In this case, $c_0[w_x -] = \lim_{\epsilon \downarrow 0} c_0[w_x - \epsilon]$.

¹² Although Scully et al. (2020b) do not explicitly state this result—namely, our Lemma 7.4—they prove it as an intermediate step toward another result (Scully et al. 2020b, section 4).

¹³ Recall from Definition 4.4(ii) that a w -interval (b, c) is right-maximal if, roughly speaking, c is as large as possible for a w -interval starting at b .

¹⁴ The choice of r^* does not actually affect the scheduling decisions of the Step and Spike policies, but it is convenient in discussion for all three rank functions to have the same maximum.

¹⁵ The validity of (5.1) for the light-tailed distributions we consider rests on the assumptions we make in Definition 5.2. See Online Appendix EC.3 for details.

References

Aalto S, Ayesta U, Righter R (2009) On the Gittins index in the $M/G/1$ queue. *Queueing Systems* 63(1–4):437–458.
Aalto S, Ayesta U, Righter R (2011) Properties of the Gittins index with application to optimal scheduling. *Probab. Engrg. Inform. Sci.* 25(3):269–288.
Abate J, Whitt W (1997) Asymptotics for $M/G/1$ low-priority waiting-time tail probabilities. *Queueing Systems* 25(1):173–233.
Bingham NH, Goldie CM, Teugels JL (1987) *Regular Variation, Encyclopedia of Mathematics and Its Applications*, vol. 27 (Cambridge University Press, Cambridge, UK).
Borst SC, Boxma OJ, Núñez-Queija R, Zwart B (2003) The impact of the service discipline on delay asymptotics. *Performance Eval.* 54(2):175–206.
Boxma OJ, Zwart B (2007) Tails in scheduling. *ACM SIGMETRICS Performance Eval. Rev.* 34(4):13–20.

Cline DBH (1994) Intermediate regular and Π variation. *Proc. London Math. Soc.* s3-68(3):594–616.
Cox DR, Smith WL (1961) *Queues* (Methuen, London).
Crovella ME, Bestavros A (1997) Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Trans. Networking* 5(6):835–846.
Gittins JC (1989) *Multi-Armed Bandit Allocation Indices*, 1st ed., *Wiley-Interscience Series in Systems and Optimization* (Wiley, Chichester, UK).
Gittins JC, Glazebrook KD, Weber RR (2011) *Multi-Armed Bandit Allocation Indices*, 2nd ed. (Wiley, Chichester, UK).
Grosf I, Yang K, Scully Z, Harchol-Balter M (2021) Nudge: Stochastically improving upon FCFS. *Proc. ACM Measurement Anal. Comput. Systems* 5(2):21.
Harchol-Balter M (2013) *Performance Modeling and Design of Computer Systems: Queueing Theory in Action* (Cambridge University Press, Cambridge, UK).
Harchol-Balter M, Downey AB (1997) Exploiting process lifetime distributions for dynamic load balancing. *ACM Trans. Comput. Systems* 15(3):253–285.
Kendall DG (1953) Stochastic processes occurring in the theory of queues and their analysis by the method of the imbedded Markov chain. *Ann. Math. Statist.* 24(3):338–354.
Kleinrock L (1976) *Computer Applications, Queueing Systems*, vol. 2 (Wiley, New York).
Mandjes M, Nuyens M (2005) Sojourn times in the $M/G/1$ FB queue with light-tailed service times. *Probab. Engrg. Inform. Sci.* 19(3):351–361.
Mimica A (2016) Exponential decay of measures and Tauberian theorems. *J. Math. Anal. Appl.* 440(1):266–285.
Nair J, Wierman A, Zwart B (2010) Tail-robust scheduling via limited processor sharing. *Performance Eval.* 67(11):978–995.
Nakagawa K (2005) Tail probability of random variable and Laplace transform. *Applicable Anal.* 84(5):499–522.
Nakagawa K (2007) Application of Tauberian theorem to the exponential decay of the tail probability of a random variable. *IEEE Trans. Inform. Theory* 53(9):3239–3249.
Núñez-Queija R (2002) Queues with equally heavy sojourn time and service requirement distributions. *Ann. Oper. Res.* 113(1/4):101–117.
Peterson DL (1996) Data center I/O patterns and power laws. *22nd International Computer Measurement Group Conference* (Computer Measurement Group, San Diego), 1034–1045.
Schrage LE (1968) A proof of the optimality of the shortest remaining processing time discipline. *Oper. Res.* 16(3):687–690.
Scully Z (2022) A new toolbox for scheduling theory. PhD thesis, Carnegie Mellon University, Pittsburgh.
Scully Z, Harchol-Balter M (2018) SOAP bubbles: Robust scheduling under adversarial noise. *56th Annu. Allerton Conf. Commun. Control Comput.* (IEEE, Monticello, IL), 144–154.
Scully Z, Harchol-Balter M (2021) The Gittins policy in the $M/G/1$ queue. *19th Internat. Sympos. Modeling Optim. Mobile Ad Hoc Wireless Networks WiOpt 2021* (IFIP, Philadelphia), 248–255.
Scully Z, Grosf I, Harchol-Balter M (2020a) The Gittins policy is nearly optimal in the $M/G/k$ under extremely general conditions. *Proc. ACM Measurement Anal. Comput. Systems* 4(3):43.
Scully Z, Grosf I, Harchol-Balter M (2021) Optimal multiserver scheduling with unknown job sizes in heavy traffic. *Performance Eval.* 145:102150.
Scully Z, Grosf I, Mitzenmacher M (2022) Uniform bounds for scheduling with job size estimates. *13th Innovations Theor. Comput. Sci. Conf. ITCS 2022, Leibniz International Proceedings in Informatics (LIPIcs)* (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Berkeley, CA), 114:1–114:30.
Scully Z, Harchol-Balter M, Scheller-Wolf A (2018) One clean analysis of all age-based scheduling policies. *Proc. ACM Measurement Anal. Comput. Systems* 2(1):1–30.

- Scully Z, Harchol-Balter M, Scheller-Wolf A (2020b) Simple near-optimal scheduling for the M/G/1. *Proc. ACM Measurement Anal. Comput. Systems* 4(1):11.
- Scully Z, van Kreveld L, Boxma O, Dorsman J-P, Wierman A (2020c) Characterizing policies with optimal response time tails under heavy-tailed job sizes. *Proc. ACM Measurement Anal. Comput. Systems* 4(2):30.
- Stolyar AL, Ramanan K (2001) Largest weighted delay first scheduling: Large deviations and optimality. *Ann. Appl. Probab.* 11(1):1–48.
- Wierman A, Zwart B (2012) Is tail-optimal scheduling possible? *Oper. Res.* 60(5):1249–1257.

Ziv Scully is an assistant professor in the School of Operations Research and Information Engineering at Cornell University. He researches decision making under uncertainty, with a focus on scheduling and resource allocation. His work has won the 2022 INFORMS George Nicholson Student Paper Competition, the 2022 SIGMETRICS Doctoral Dissertation Award, and multiple best paper awards at conferences.

Lucas van Kreveld is a university lecturer in Stochastic Operations Research at the Eindhoven University of Technology. His research is focused on queueing theory, with an emphasis on asymptotic methods, scheduling, and queueing networks.

EC.1. Deferred Proofs for Tail Asymptotics of SOAP Policies

EC.1.1. Proofs for Heavy-Tailed Job Sizes

LEMMA 7.6. *Suppose Condition 4.5 holds.*

(i) *For all $p \geq 0$,*

$$\mathbf{E}[X_0[w_x]^{p+1}] \leq \begin{cases} O(1) & \text{if } p < \alpha - 1 \\ O(\log x) & \text{if } p = \alpha - 1 \\ O(x^{\max\{1, \zeta + \theta\}(p - \alpha + 1)}) & \text{if } p > \alpha - 1. \end{cases}$$

(ii) *For all $p \geq 0$,*

$$\sum_{k=1}^{K[w_x]} \mathbf{E}[X_k[w_x]^{p+1}] \leq \begin{cases} O(x^{\theta p + \zeta p - \alpha + 1}) & \text{if } \zeta p < \alpha - 1 \\ O(x^{\theta p} \log x^\eta) & \text{if } \zeta p = \alpha - 1 \\ O(x^{\theta p + \eta(\zeta p - \alpha + 1)}) & \text{if } \zeta p > \alpha - 1. \end{cases}$$

Proof. We first show (i). Because $(x, c_0[w_x])$ is a w_x -interval, Condition 4.5 implies

$$c_0[w_x] - x = O(x^{\zeta + \theta}). \quad (\text{EC.1.1})$$

We compute

$$\begin{aligned} \mathbf{E}[X_0[w_x]^{p+1}] &= \int_0^{c_0[w_x]} (p+1)t^p \bar{F}(t) dt && \text{[by Lemma 6.3]} \\ &\leq \int_0^{O(x^{\max\{1, \zeta + \theta\})} O(t^{p-\alpha}) dt && \text{[by Definition 4.1 and (EC.1.1)]} \\ &= \begin{cases} O(1) & \text{if } p < \alpha - 1 \\ O(\log x) & \text{if } p = \alpha - 1 \\ O(x^{\max\{1, \zeta + \theta\}(p - \alpha + 1)}) & \text{if } p > \alpha - 1, \end{cases} \end{aligned}$$

thus proving (i).

We now show (ii), following a similar argument but with a more involved computation. Note that Definitions 6.1 and 6.5 together imply

$$b_k[w_x] \geq x \quad \text{for all } k \geq 1. \quad (\text{EC.1.2})$$

We compute

$$\begin{aligned} \sum_{k=1}^{K[w_x]} \mathbf{E}[X_k[w_x]^{p+1}] &= \sum_{k=1}^{K[w_x]} \int_{b_k[w_x]}^{c_k[w_x]} (p+1)(t - b_k[w_x])^p \bar{F}(t) dt && \text{[by Lemma 6.3]} \\ &\leq \sum_{k=1}^{K[w_x]} (p+1)(c_k[w_x] - b_k[w_x])^p \int_{b_k[w_x]}^{c_k[w_x]} \bar{F}(t) dt && \text{[by (EC.1.2)]} \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{k=1}^{K[w_x]} O(x^{\theta p} \cdot b_k[w_x]^{\zeta p}) \int_{b_k[w_x]}^{c_k[w_x]} O(t^{-\alpha}) dt && \text{[by Definition 4.1 and Condition 4.5]} \\
&\leq \sum_{k=1}^{K[w_x]} O(x^{\theta p}) \int_{b_k[w_x]}^{c_k[w_x]} O(t^{\zeta p - \alpha}) dt \\
&\leq O(x^{\theta p}) \int_x^{c_{K[w_x]}[w_x]} O(t^{\zeta p - \alpha}) dt && \text{[by (EC.1.2)]} \\
&\leq O(x^{\theta p}) \int_x^{O(x^\eta)} O(t^{\zeta p - \alpha}) dt && \text{[by Condition 4.5]} \\
&= \begin{cases} O(x^{\theta p + \zeta p - \alpha + 1}) & \text{if } \zeta p < \alpha - 1 \\ O(x^{\theta p} \log x^\eta) & \text{if } \zeta p = \alpha - 1 \\ O(x^{\theta p + \eta(\zeta p - \alpha + 1)}) & \text{if } \zeta p > \alpha - 1, \end{cases}
\end{aligned}$$

thus proving (ii). □

LEMMA 7.9. *Suppose Condition 4.5 holds, and let $\kappa = 2 \max\{\alpha - 1, \theta\}$. For all $x \geq 0$ and $a \in (y_x, x)$,*

$$c_0[w_x(a)-] \geq \Omega\left(\left(\frac{x-a}{x^\zeta}\right)^{1/\kappa}\right).$$

Proof. Because $\kappa > \theta \geq 0$, by Condition 4.5 and (EC.1.2), for all $u \geq 0$ and $k \geq 1$,

$$u \geq \Omega\left(\left(\frac{c_k[w_u] - b_k[w_u]}{b_k[w_u]^\zeta}\right)^{1/\kappa}\right). \quad \text{(EC.1.3)}$$

We now plug in $u = c_0[w_x(a)-]$ and make the following observations.

- By Definition 6.1, we know $u = c_0[w_x(a)-]$ is the earliest age at which a job has rank at least $w_x(a)$, so $w_u = w_x(a)$.
- By Definition 6.5, a job's rank is at most $w_x(a)$ between ages a and x , so there exists $k \geq 1$ such that

$$b_k[w_x(a)] \leq a < x \leq c_k[w_x(a)].$$

In particular, $x > b_k[w_x(a)]$ and $x - a \leq c_k[w_x(a)] - b_k[w_x(a)]$.

Applying these observations to (EC.1.3) with $u = c_0[w_x(a)-]$ yields the desired bound. □

EC.1.2. Proofs for Light-Tailed Job Sizes

PROPOSITION 9.2. *Consider an M/G/1 with any nicely light-tailed job size distribution under a SOAP policy. The policy is log-tail-pessimal if $a^* = x_{\max}$.*

Proof. Since no work-conserving policy has response time decay rate lower than a busy period's decay rate $d(B)$ (Mandjes and Nuyens, 2005, Corollary 6), it suffices to show $d(T) \leq d(B)$.

Recall that y_x denotes the (first) age of the maximum rank in the interval $[0, x]$. Since $T(x)$ is stochastically increasing in x (Lemma 6.6), it holds that $\mathbf{P}[T(x) > t] \geq \mathbf{P}[T(y_x) > t]$ for all $t \geq 0$. Additionally we have

that $\mathbf{P}[T(y_x) > t] \geq \mathbf{P}[T_{\text{FB}}(y_x) > t]$ for all $t, x \geq 0$, where $T_{\text{FB}}(x)$ is the response time for a job of size x under FB. The reason for this last inequality is that a job of size y_x must wait for all other jobs to receive up to y_x units of service before completing.¹ As a result, (Mandjes and Nuyens, 2005, Proposition 8) implies

$$d(T(x)) \leq d(T(y_x)) \leq d(T_{\text{FB}}(y_x)) = d(B_{y_x}).$$

Additionally, up to its last line the proof of (Mandjes and Nuyens, 2005, Lemma 9) is valid for arbitrary service policies. If $x_0 > 0$ is such that $\mathbf{P}[X \geq x_0] > 0$, we thus find

$$\begin{aligned} d(T) &\leq \mathbf{P}[X \geq x_0]^{-1} \int_{x_0}^{x_{\max}} d(T(x)) dF(x) \\ &\leq \mathbf{P}[X \geq x_0]^{-1} \int_{x_0}^{x_{\max}} d(B_{y_x}) dF(x). \end{aligned} \quad (\text{EC.1.4})$$

Our goal is to show $d(T) \leq d(B)$, or equivalently $d(T) < d(B) + \varepsilon$ for all $\varepsilon > 0$. By (EC.1.4), it suffices to show that $\lim_{x \rightarrow x_{\max}} d(B_{y_x}) = d(B)$. It is shown in (Mandjes and Nuyens, 2005, Lemma 10) that $\lim_{x \rightarrow x_{\max}} d(B_x) = d(B)$, so our task is to show that the limit still holds with y_x instead of x .

Consider arbitrary $\varepsilon > 0$. Because $\lim_{x \rightarrow \infty} d(B_x) = d(B)$, there exists $x_0 > 0$ such that $|d(B_x) - d(B)| < \varepsilon$ for all $x > x_0$. Because $a^* = x_{\max}$, there exists $x_1 > x_0$ such that $y_{x_1} = x_1$, and thus $|d(B_{y_{x_1}}) - d(B)| < \varepsilon$. But $d(B_{y_x})$ is decreasing in x , because y_x is increasing in x , and B_x is stochastically increasing in x . We conclude that for all $x > x_1$, we have $|d(B_{y_x}) - d(B)| < \varepsilon$. Our choice of $\varepsilon > 0$ was arbitrary, so $\lim_{x \rightarrow x_{\max}} d(B_{y_x}) = d(B)$, as desired. \square

LEMMA 9.5. *Let π be a SOAP policy with $0 < a^* < x_{\max}$. We have*

$$d(T_\pi) = d(T_\pi^{(2)}) \in [d(T_{\text{step}}^{(2)}), d(T_{\text{spike}}^{(2)})].$$

Proof. Clearly, T_π is a mixture of $T_\pi^{(1)}$ and $T_\pi^{(2)}$. Lemma 6.6 implies $T_\pi^{(2)} \geq_{\text{st}} T_\pi^{(1)}$, implying $d(T_\pi) = d(T_\pi^{(2)})$.

The same reasoning applies to Step and Spike. The lemma thus follows if we can show

$$T_{\text{spike}}^{(2)} \leq_{\text{st}} T_\pi^{(2)} \leq_{\text{st}} T_{\text{step}}^{(2)}. \quad (\text{EC.1.5})$$

The comparison in (EC.1.5) follows from a key fact from the SOAP analysis (Scully and Harchol-Balter, 2018) called the *Pessimism Principle*, which states that the response time of a particular job J is unaffected if, instead of following the usual rank function, job J follows its *worst future rank* function (Definition 6.5). The intuition is that any jobs that will get served ahead of job J in the future may as well be served ahead of it right now.

¹ One can give a more formal proof of the inequality using the SOAP analysis (Scully and Harchol-Balter, 2018).

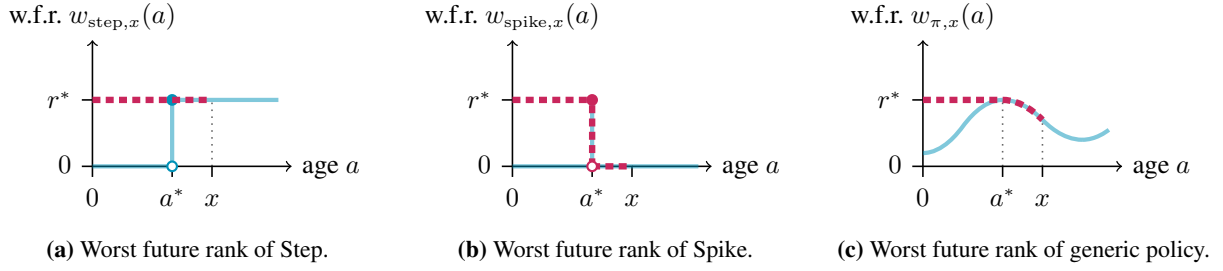


Figure EC.1 Worst future rank functions (Definition 6.5, abbreviated w.f.r., dotted magenta curves) of the policies shown in Figure 9.1, with the original rank functions (translucent cyan curves) for reference. We show the worst future rank functions for a class 2 job of size $x > a^*$.

We illustrate in Figure EC.1 the worst future rank under Step, Spike, and π . Notice that, for any size $x > a^*$, we have

$$\begin{aligned} a \in [0, a^*] &\Rightarrow r^* = w_{\text{spike},x}(a) = w_{\pi,x}(a) = w_{\text{step},x}(a) = r^*, \\ a \in (a^*, x) &\Rightarrow 0 = w_{\text{spike},x}(a) \leq w_{\pi,x}(a) \leq w_{\text{step},x}(a) = r^*. \end{aligned}$$

The Pessimism Principle says that we can compute a particular job J 's response time by imagining that it always has its worst future rank. Increasing a job's rank can only increase its response time, so the above worst future rank comparisons imply that for all $x > a^*$,

$$T_{\text{spike}}(x) \leq_{\text{st}} T_{\pi}(x) \leq_{\text{st}} T_{\text{step}}(x).$$

The desired (EC.1.5) follows because class 2 jobs are those of size greater than a^* . \square

LEMMA 9.6. *The response time distributions of class 2 jobs under Step and Spike are*

$$T_{\text{step}}^{(2)} =_{\text{st}} B_{a^*}(W) + B_{a^*}(X^{(2)}), \quad T_{\text{spike}}^{(2)} =_{\text{st}} B_{a^*}(W) + B_{a^*}(a^*) + X^{(2)} - a^*.$$

where $X^{(2)} = (X \mid X > a^*)$ is the size distribution of class 2 jobs, and the random variables in each sum are mutually independent.

Proof. This result follows easily from the SOAP analysis (Scully and Harchol-Balter, 2018). For completeness, we sketch the main ideas of how the SOAP analysis applies to Step and Spike. Consider a class 2 job J .

- Under Step, job J always has worst future rank r^* (Figure EC.1(a)). Job J is thus delayed by any jobs present when it arrives, plus the pre-age- a^* portion of any jobs that arrive while it is in the system.
- Under Spike, job J has worst future rank r^* only until age a^* (Figure EC.1(a)). Job J is thus delayed by any jobs present when it arrives, plus the pre-age- a^* portion of any jobs that arrive before it reaches age a^* . Once job J reaches age a^* , its worst future rank is 0, so no further arrivals delay it.

The reason in both cases for looking at the pre-age- a^* portion of new arrivals is because at age a^* , those new arrivals reach rank r^* , and thus job J has priority over them due to FCFS tiebreaking (Definition 3.1).

The delay due to jobs present when job J arrives corresponds to the W in each formula, and the delay due to new arrivals corresponds to the $B_{a^*}(\cdot)$ uses. The difference between the formulas is due to the fact that under Step, new arrivals delay job J until it completes, whereas under Spike, new arrivals delay job J only if they arrive before it reaches age a^* , with the last $X^{(2)} - a^*$ portion of job J's service occurring uninterrupted. \square

LEMMA 9.7. *Consider an M/G/1 with nicely light-tailed job size distribution X , and define*

$$\begin{aligned} s_0 &= \gamma(\sigma^{-1}) = \gamma(\mathcal{L}[X]), & s_2 &= \sigma(\gamma(\sigma)) = \arg \min_{s \geq s_0} \sigma^{-1}(s), \\ s_1 &= \gamma(\mathcal{L}[W]) = \text{least root of } \sigma^{-1}, & s_3 &= \gamma(\sigma) = \min_{s \geq s_0} \sigma^{-1}(s). \end{aligned}$$

Then, as illustrated in Figure 9.2, the following hold:

- (i) σ^{-1} is convex on (s_0, ∞) , decreasing on (s_0, s_2) , and increasing on (s_2, ∞) ;
- (ii) $s_0 < s_1 < s_2 < s_3 < 0$.

Analogous statements hold for σ_a for all $a \in (0, x_{\max}]$.

Proof. We prove the statement just for σ , as the argument for σ_a is analogous. The illustration in Figure 9.2 may provide helpful intuition for the arguments that follow.

We begin by observing some general properties of σ^{-1} . Because $\mathcal{L}[X]$ is convex on (s_0, ∞) , so is σ^{-1} . This, along with the definition of s_2 , implies (a). The slope of σ^{-1} at zero is

$$(\sigma^{-1})'(0) = 1 + \lambda \mathcal{L}[X]'(0) = 1 - \rho \in (0, 1),$$

and by Definition 5.2, we have $\sigma^{-1}(s_0) = \infty > 0$. Additionally, Definition 5.2 implies $s_0 < 0$. This means σ^{-1} is negative on a finite nonempty interval, namely $(s_1, 0)$, and nonnegative outside that interval.

We can now show the inequalities in (b).

- $s_0 < s_1$: Because $|\sigma^{-1}(s_1)| = 0 < \infty$, we have $s_0 = \gamma(\sigma^{-1}) \leq s_1$. But $\sigma^{-1}(s_0) > 0$, so $s_0 \neq s_1$.
- $s_3 < 0$: Because σ^{-1} is negative on some interval, its global minimum is negative.
- $s_1 < s_2$: Because $s_3 = \sigma^{-1}(s_2) < 0$, we must have $s_2 \in (s_1, 0)$.
- $s_2 < s_3$: Because σ^{-1} is convex with $\sigma^{-1}(0) = 0$ and $(\sigma^{-1})'(0) \in (0, 1)$, we have $s_2 < \sigma^{-1}(s_2)$. \square

In some of the proofs below, we use the fact that for sums of independent random variables $U, V \geq 0$, (9.1) implies

$$\begin{aligned} d(U + V) &= -\gamma(\mathcal{L}[U + V]) \\ &= -\max\{\gamma(\mathcal{L}[U]), \gamma(\mathcal{L}[V])\} \\ &= \min\{d(U), d(V)\}. \end{aligned} \tag{EC.1.6}$$

This is also shown by Mandjes and Nuyens (2005, Lemma 3) without relying on (9.1).

LEMMA 9.8. *Let π be a SOAP policy with $0 < a^* < x_{\max}$. Then the decay rate of its response time is*

$$d(T_\pi) = -\gamma(\mathcal{L}[W] \circ \sigma_{a^*}).$$

Proof. Combining Lemmas 9.5 and 9.6, we have

$$d(T_\pi) \in [d(B_{a^*}(W)), d(B_{a^*}(W) + B_{a^*}(X^{(2)}))].$$

By (9.1) and (9.2), the lower bound is

$$d(B_{a^*}(W)) = -\gamma(\mathcal{L}[W] \circ \sigma_{a^*}).$$

We aim to show that the upper bound matches this. Applying (9.1), (9.2), and (EC.1.6) to the upper bound, we see that it suffices to show

$$\gamma(\mathcal{L}[X^{(2)}] \circ \sigma_{a^*}) \leq \gamma(\mathcal{L}[W] \circ \sigma_{a^*}).$$

Lemma EC.1.1, which we state and prove below, implies the above if $\gamma(\mathcal{L}[X^{(2)}]) \leq \gamma(\mathcal{L}[W])$, which in turn is implied by Lemma 9.7(ii) and the fact that $\gamma(\mathcal{L}[X^{(2)}]) = \gamma(\mathcal{L}[X])$. \square

The following lemma, which is used in the proof above, relates $\gamma(f \circ \sigma)$ to $\gamma(f)$, thus relating the decay rate of a busy period to the decay rate of its initial work.

LEMMA EC.1.1. *Let $f : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ be a function for which $\gamma(f)$ is well defined and finite. Then $\gamma(f \circ \sigma)$ is finite, and*

$$\begin{aligned} \gamma(f \circ \sigma) &= \sigma^{-1}(\max\{\gamma(f), \sigma(\gamma(\sigma))\}) \\ &= \begin{cases} \sigma^{-1}(\gamma(f)) & \text{if } \gamma(f) > \sigma(\gamma(\sigma)) \\ \gamma(\sigma) & \text{otherwise.} \end{cases} \end{aligned}$$

In particular, $\gamma(f \circ \sigma)$ is a nondecreasing function of $\gamma(f)$. Analogous statements hold for σ_a for all $a \in [0, x_{\max}]$.

Proof. We prove the statement just for σ , as the proof for σ_a is analogous. There are two reasons $f(\sigma(s))$ can be infinite:

- We can have $\sigma(s)$ infinite, which happens if and only if $s < \gamma(\sigma)$.
- We can have $\sigma(s)$ finite but $f(\sigma(s))$ infinite, which happens if $-\infty < \sigma(s) < \gamma(f)$ and only if $-\infty < \sigma(s) \leq \gamma(f)$.

Recalling that $\sigma(\gamma(\sigma))$ is the minimum finite value $\sigma(s)$ can take on (see Figure 9.2), we see that the latter reason can occur for some $s > \gamma(\sigma)$ if and only if $\sigma(\gamma(\sigma)) < \gamma(f)$, implying the desired formula.

The finiteness of $\gamma(f \circ \sigma)$ follows from finiteness of $\sigma^{-1}(\gamma(f))$ when $\gamma(f) > \gamma(\sigma)$, which by Lemma 9.7(ii) includes all cases when $\gamma(f) > \sigma(\gamma(\sigma))$. The monotonicity of $\gamma(f \circ \sigma)$ in $\gamma(f)$ follows Lemma 9.7(i). \square

PROPOSITION 9.9. *Consider an M/G/1 with any nicely light-tailed job size distribution under a SOAP policy. The policy is log-tail-intermediate if $0 < a^* < x_{\max}$.*

Proof. The optimal decay rate is that of FCFS. A special case of a result of Stolyar and Ramanan (2001, Theorem 2.2), together with (9.1) and (EC.1.6) implies this is

$$d(T_{\text{FCFS}}) = d(W + X) = d(W) = -\gamma(\mathcal{L}[W]).$$

The pessimal decay rate is that of FB. A result of Mandjes and Nuyens (2005, Theorem 1) states $d(T_{\text{FB}}) = d(B)$. Together with (9.1) and (9.2) and Lemma EC.1.1, this implies

$$\begin{aligned} d(T_{\text{FB}}) &= d(B) = -\gamma(\mathcal{L}[X] \circ \sigma) \\ &= -\gamma(\sigma) = -\gamma(\mathcal{L}[W] \circ \sigma). \end{aligned}$$

Above, we use the fact that $\gamma(\mathcal{L}[X]) < \gamma(\mathcal{L}[W]) < \sigma(\gamma(\sigma))$, as shown in Lemma 9.7(ii), when applying Lemma EC.1.1.

Having computed the optimal and pessimal decay rates in Lemma 9.8, it suffices to show that in the $0 < a^* < x_{\max}$ case, we have

$$\gamma(\mathcal{L}[W]) < \gamma(\mathcal{L}[W] \circ \sigma_{a^*}) < \gamma(\mathcal{L}[W] \circ \sigma),$$

which we may rewrite as

$$\gamma(\mathcal{L}[W] \circ \sigma_0) < \gamma(\mathcal{L}[W] \circ \sigma_{a^*}) < \gamma(\mathcal{L}[W] \circ \sigma_{x_{\max}}).$$

Lemma EC.1.2, which we state and prove below, implies $\gamma(\mathcal{L}[W] \circ \sigma_a)$ is strictly increasing in a . Therefore, the above holds if $0 < a^* < x_{\max}$, as desired. \square

It remains only to prove the strict monotonicity of $\gamma(\mathcal{L}[W] \circ \sigma_a)$ in a . We prove a more general statement below.

LEMMA EC.1.2. *Let $f : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ be a function for which $\gamma(f) < 0$, and let $0 \leq a < b \leq x_{\max}$. Then*

$$\gamma(f \circ \sigma_a) < \gamma(f \circ \sigma_b).$$

Proof. We begin by comparing $\sigma_a^{-1}(s)$ with $\sigma_b^{-1}(s)$ for all $s < 0$, computing²

$$\begin{aligned} &a < b \\ \Rightarrow &\min\{X, a\} <_{\text{st}} \min\{X, b\} \\ \Rightarrow &\mathcal{L}[\min\{X, a\}](s) < \mathcal{L}[\min\{X, b\}](s) \\ \Rightarrow &\sigma_a^{-1}(s) < \sigma_b^{-1}(s). \end{aligned} \tag{EC.1.7}$$

²Two clarifications about the computation below. First, the notation $U <_{\text{st}} V$ means that $\mathbf{P}[U > t] \leq \mathbf{P}[V > t]$ for all $t \in \mathbb{R}$, and the set of $t \in \mathbb{R}$ such that $\mathbf{P}[U > t] < \mathbf{P}[V > t]$ has positive Lebesgue measure. Second, because $a < \infty$, the left-hand sides of the last two steps are always finite for all $s < 0$.

There are two important implications of (EC.1.7). The first implication is that the global minimum of σ_a^{-1} is less than that of σ_b^{-1} . But these global minima are $\gamma(\sigma_a)$ and $\gamma(\sigma_b)$, respectively (see Figure 9.2), so

$$\gamma(\sigma_a) < \gamma(\sigma_b). \quad (\text{EC.1.8})$$

This means $\sigma_a(s)$ is finite whenever $\sigma_b(s)$ is. This contributes to the second implication of (EC.1.7): by Lemma 9.7(i),

$$\sigma_a(s) > \sigma_b(s) \quad \text{for all } s \in [\gamma(\sigma_b), 0). \quad (\text{EC.1.9})$$

Note that $f(\sigma_a(s))$ diverges only if $s \leq \gamma(\sigma_a)$ or $\sigma_a(s) \leq \gamma(f)$, while $f(\sigma_b(s))$ diverges if $s < \gamma(\sigma_b)$ or $\sigma_b(s) < \gamma(f)$. Therefore, (EC.1.8) and (EC.1.9) together imply that there exists a value of s such that $f(\sigma_b(s))$ diverges while $f(\sigma_a(s))$ does not. \square

EC.2. Properties of the Gittins Policy via the ‘‘Gittins Game’’

The goal of this section is to prove two key remaining properties of the Gittins policy, Theorem 5.10 and Lemma 8.2. To prove both of these properties, we will use a different perspective on the Gittins policy called the ‘‘Gittins game’’ (Scully et al., 2020a). The Gittins game gives an alternative way to define the Gittins rank function. While it is less direct than the definitions we have used so far (Definitions 3.2 and 8.1), the intermediate steps it introduces turn out to be crucial for proving Theorem 5.10 and Lemma 8.2.

Aside from Theorem 5.10 and Lemma 8.2, most of the definitions and results in this section are due to Scully et al. (2020a), who actually study a much more general job model than ours. For simplicity, we restate the key definitions and results in our setting. However, the statements and proofs of Theorem 5.10 and Lemma 8.2 are straightforward to translate to the more general job model of Scully et al. (2020a).

EC.2.1. The Gittins Game

The Gittins game is an optimization problem. Its inputs are a job at some age b and a *penalty* w . During the game, we serve the job for as long as we like. If the job completes, the game ends. At any moment before the job completes, we may choose to give up, in which case we pay the penalty w and the game immediately ends. The goal of the game is to minimize the expected sum of the time spent serving the job plus the penalty paid.

We can think of the Gittins game with penalty w as an optimal stopping problem whose state is the age b of the job. Standard optimal stopping theory (Peskir and Shiryaev, 2006; Shiryaev, 2008) implies that the optimal strategy thus has the following form: serve the job until it reaches some age $c \geq b$, then give up. A possible policy here is never giving up, which is represented by $c = \infty$.

Suppose we start serving a job at age b and stop if it reaches age c . The expected amount of time we spend serving the job is

$$\text{service}(b, c) = \mathbf{E}[\min\{S, c\} \mid S > b] = \int_b^c \frac{\bar{F}(t)}{\bar{F}(b)} dt,$$

and the probability the job finishes before reaching age b is

$$\text{done}(b, c) = \mathbf{P}[S \leq c \mid S > b] = 1 - \frac{\bar{F}(c)}{\bar{F}(b)}.$$

We can write the time-per-completion function as $\varphi(b, c) = \text{service}(b, c) / \text{done}(b, c)$ (see Definition 8.1).

Suppose we employ the stop-at-age- c policy in the Gittins game starting from age b with penalty w . The expected cost of the Gittins game with this policy is

$$\text{game}(w; b, c) = \text{service}(b, c) + w(1 - \text{done}(b, c))$$

The *optimal* cost of the Gittins game is therefore

$$\text{game}^*(w; b) = \inf_{c \geq b} \text{game}(w; b, c).$$

The lemma below follows immediately from the definition of $\text{game}^*(w; b)$ as an infimum of $\text{game}(w; b, c)$, each of which is a linear function of w (Scully et al., 2020a, Lemmas 5.2 and 5.3).

LEMMA EC.2.1. *For all ages b , the optimal cost $\text{game}^*(w; b)$ is increasing and concave as a function of w . Because giving up immediately is always a possible policy, it is also bounded above by $\text{game}^*(w; b) \leq w$.*

EC.2.2. Relating the Gittins Game to the Gittins Rank Function

The Gittins game is intimately connected to the Gittins rank function, and it is this connection that is important for proving Lemma 8.2. The following lemmas state two such connections. They are the same or very similar to many previous results in the literature on Gittins in the M/G/1 (Aalto et al., 2009; Gittins, 1989; Gittins et al., 2011; Scully et al., 2020a, 2018a), but we sketch their proofs for completeness.

LEMMA EC.2.2. *The Gittins rank function can be expressed in terms of the Gittins game as*

$$\begin{aligned} r(a) &= \inf\{w \geq 0 \mid \text{game}^*(w; a) < w\} \\ &= \max\{w \geq 0 \mid \text{game}^*(w; a) = w\}. \end{aligned}$$

Proof. The infimum and maximum are equivalent by Lemma EC.2.1. The infimum is equal to the rank $r(a) = \inf_{c > a} \varphi(a, c)$ because, by the fact that we can write $\text{game}(w; b, c)$ as

$$\text{game}(w; b, c) = w - (w - \varphi(b, c)) \text{done}(b, c), \tag{EC.2.1}$$

we have $\text{game}(w; b, c) < w$ if and only if $\varphi(b, c) < w$.³ □

³ Recall that $\text{done}(b, c) \in [0, 1]$ and that if $\text{done}(b, c) = 0$, then $\varphi(b, c) = \infty$ (Definition 8.1).

LEMMA EC.2.3. *In the Gittins game with penalty w with the job currently at age a , it is optimal to continue serving the job if and only if $r(a) \leq w$,⁴ and it is optimal to give up if and only if $r(a) \geq w$.*

Proof. Giving up incurs cost w , so by the maximum in Lemma EC.2.2, it is optimal to give up if and only if $r(a) \geq w$. This means it is optimal to continue serving the job if $r(a) < w$. The fact that serving is optimal in the $r(a) = w$ edge case follows from the fact that if $\varphi(a, c) = w$ for some $c > a$,⁵ then by (EC.2.1), we have $\text{game}(w; a, c) = w$. \square

We are now ready to prove Lemma EC.2.3, which we restate below. Recall that a w -interval is one in which the Gittins rank is bounded above by w . The key to the proof is that Lemma EC.2.3 relates w -intervals to optimal play the Gittins game.

LEMMA 8.2. *Under Gittins, for any right-maximal w -interval (b, c) , we have $\varphi(b, c) \leq w$.*

Proof. Consider playing the Gittins game starting from age b . By Lemma EC.2.3, giving up if the job reaches age c is an optimal policy. Specifically, because (b, c) is a w -interval, it is optimal to continue serving the job until at least age c , and because the w -interval is right-maximal, it is optimal to give up if the job reaches age c (which never happens if $c = x_{\max}$). This means $\text{game}^*(w; b) = \text{game}(w; b, c)$. Combining Lemma EC.2.1 and (EC.2.1) implies $\varphi(b, c) \leq w$. \square

We note that Lemma 8.2 is similar, but not identical, to properties of Gittins in the M/G/1 studied by Aalto et al. (2009, 2011). Related properties have also been shown for versions of Gittins in discrete-time settings (Dumitriu et al., 2003; Gittins, 1989; Gittins et al., 2011).

EC.2.3. Relating the Gittins Game to Mean Response Time

It remains only to prove Theorem 5.10, which bounds the mean response time of q -approximate Gittins policies. To do so, we use a result of Scully et al. (2020a) that relates the Gittins game to a system's mean response time.

DEFINITION EC.2.4. Let $r : [0, x_{\max}) \rightarrow \mathbb{R}$ be the rank function of some SOAP policy, and let $w \in \mathbb{R}$.

- (i) The (r, w) -relevant work of a job is the amount of service the job requires to either complete or reach rank at least w according to r , meaning reaching an age a satisfying $r(a) \geq w$.
- (ii) The (r, w) -relevant work of the system is the total (r, w) -relevant work of all jobs present. We denote the steady-state distribution of (r, w) -relevant work under policy π by $W_\pi(r, w)$. Note that r need not be the rank function of policy π .

⁴ Strictly speaking, it is optimal to continue serving the job if and only if the rank is upper bounded in a “forward neighborhood” of a , meaning there exists $\varepsilon > 0$ such that for all $\delta \in [0, \varepsilon)$, we have $r(a + \delta) \leq w$. For non-pathological job size distributions, this holds in the $r(a) < w$ case (Aalto et al., 2011), so it only needs to be checked when $r(a) = w$.

⁵ The $c > a$ restriction is why we need the rank to be bounded not just at a but in a forward neighborhood of a .

The (r_{Gtn}, w) -relevant work of a job is related to the Gittins game via Lemma EC.2.3: it is the amount of time we would serve the job when optimally playing the Gittins game with penalty w . It turns out that mean (r_{Gtn}, w) -relevant work directly translates into mean response time.

LEMMA EC.2.5 (Scully et al. (2020a, Theorem 6.3)). *Under any nonclairvoyant scheduling policy π , the mean response time can be written in terms of (r_{Gtn}, w) -relevant work as*

$$\mathbf{E}[T_\pi] = \frac{1}{\lambda} \int_0^\infty \frac{\mathbf{E}[W_\pi(r_{\text{Gtn}}, w)]}{w^2} dw.$$

With Lemma EC.2.5 in hand, the proof of Theorem 5.10, restated below, reduces to bounding the mean amount of (r_{Gtn}, w) -relevant work under q -approximate Gittins policies.

THEOREM 5.10. *Consider an M/G/1 with any job size distribution. For any $q \geq 1$ and any q -approximate Gittins policy π ,*⁶

$$\mathbf{E}[T_\pi] \leq q\mathbf{E}[T_{\text{Gtn}}].$$

Proof. Recall from Definition 5.9 that we may assume $r_\pi(a)/r_{\text{Gtn}}(a) \in [1, q]$ for all ages a without loss of generality. We will prove

$$\mathbf{E}[W_\pi(r_{\text{Gtn}}, w)] \leq \mathbf{E}[W_\pi(r_\pi, qw)] \leq \mathbf{E}[W_{\text{Gtn}}(r_{\text{Gtn}}, w)], \quad (\text{EC.2.2})$$

from which the theorem follows by the computation below:

$$\begin{aligned} \mathbf{E}[T_\pi] &= \frac{1}{\lambda} \int_0^\infty \frac{\mathbf{E}[W_\pi(r_{\text{Gtn}}, w)]}{w^2} dw && \text{[by Lemma EC.2.5]} \\ &\leq \frac{1}{\lambda} \int_0^\infty \frac{\mathbf{E}[W_{\text{Gtn}}(r_{\text{Gtn}}, qw)]}{w^2} dw && \text{[by (EC.2.2)]} \\ &= \frac{1}{\lambda} \int_0^\infty \frac{\mathbf{E}[W_{\text{Gtn}}(r_{\text{Gtn}}, w')]}{(w'/q)^2} d(w'/q) && \text{[by substituting } w' = qw\text{]} \\ &= q\mathbf{E}[T_{\text{Gtn}}]. && \text{[by Lemma EC.2.5]} \end{aligned}$$

To show the left-hand inequality of (EC.2.2), it suffices to show that an arbitrary job's (r_{Gtn}, w) -relevant work is upper bounded by its (r_π, qw) -relevant work (Definition EC.2.4). This is indeed the case: $r_{\text{Gtn}}(a) \leq w$ implies $r_\pi(a) \leq qr_{\text{Gtn}}(a) \leq qw$, so the job will reach rank w under Gittins after at most as much service as it needs to reach rank qw under π .

To show the right-hand inequality of (EC.2.2) we use a property of SOAP policies due to Scully and Harchol-Balter (2018, proof of Lemma 5.2). The property implies that for any rank w and SOAP policy π , we can express $\mathbf{E}[W_\pi(r_\pi, w)]$ in terms of just the job size distribution X , arrival rate λ , and the set of ages $A_\pi[w] = \{a \in [0, x_{\text{max}}) \mid r_\pi(a) < w\}$.⁷ In particular, for any fixed job size distribution, arrival rate, and

⁶This is a special case of a more general result (Scully, 2022, Chapter 16), which appeared while this work was in revision.

⁷Scully and Harchol-Balter (2018) actually focus on $\mathbf{E}[W_\pi(r_\pi, w+)] = \lim_{w' \downarrow w} \mathbf{E}[W_\pi(r_\pi, w')]$ as opposed to $\mathbf{E}[W_\pi(r_\pi, w)]$, but the same reasoning applies to $\mathbf{E}[W_\pi(r_\pi, w+)]$.

rank w , $\mathbf{E}[W_\pi(r_\pi, w)]$ is a nondecreasing function of $A_\pi[w]$, where we order sets by the usual subset partial ordering. We have $r_\pi(a) \geq r_{\text{Gtn}}(a)$, which means $A_\pi[w] \subseteq A_{\text{Gtn}}[w]$, which implies the right-hand inequality of (EC.2.2), as desired. \square

We note that one can use the techniques of Scully et al. (2018b) to generalize the statement and proof of Theorem 5.10 beyond SOAP policies. It turns out that Theorem 5.10 still holds even if we allow q -approximate Gittins policies to *adversarially* assign ranks to jobs, provided that the assigned ranks are still within a factor- q window around the rank Gittins would assign.

EC.3. Relationship Between Decay Rate and Laplace-Stieltjes Transform

The goal of this appendix is to justify our computation of decay rates (Definition 5.1) by means of Laplace-Stieltjes transform convergence (Section 9.3). Our specific goal is to justify our use of (9.1), which states $d(V) = -\gamma(\mathcal{L}[V])$. As a reminder,

$$d(V) = \lim_{t \rightarrow \infty} \frac{-\log \mathbf{P}[V > t]}{t},$$

$$\gamma(f) = \inf\{s \in \mathbb{R} \mid |f(s)| < \infty\}.$$

EC.3.1. Sufficient Condition for Computing Decay Rates

Our main tool for translating between $d(V)$ and $\gamma(\mathcal{L}[V])$ is a result of Mimica (2016), restated as Lemma EC.3.2 below, which gives a sufficient condition for $d(V) = -\gamma(\mathcal{L}[V])$. The result rests on the following definition.

DEFINITION EC.3.1. We say a function $f : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ is *regularly varying from the right at s^* with negative index*, or simply “regularly varying at s^* ”, if there exists $\alpha > 0$ such that for all $c > 0$,

$$\lim_{s \downarrow 0} \frac{f(s^* + cs)}{f(s^* + s)} = c^{-\alpha}.$$

In particular, f having a pole of finite order at s^* suffices.

It turns out being regularly varying at the singularity is the condition we need to express decay rate in terms of Laplace-Stieltjes transform convergence.

LEMMA EC.3.2 (special case of Mimica (2016, Corollary 1.3)). *Let V be a non-negative random variable with $\gamma(\mathcal{L}[V]) > -\infty$. If either $\mathcal{L}[V]$ or $\mathcal{L}[V]'$ is regularly varying at $\gamma(\mathcal{L}[V])$, then*

$$d(V) = -\gamma(\mathcal{L}[V]).$$

EC.3.2. Showing the Sufficient Condition for Computing Decay Rates Holds

It remains to show that the precondition of Lemma EC.3.2 holds whenever we apply (9.1) in Section 9.3. It turns out that all of the Laplace-Stieltjes transforms to which we apply (9.1) have a common form, so we will show that Lemma EC.3.2 applies to all functions of that form. To describe the form, we need the following definition.

DEFINITION EC.3.3. Consider an M/G/1 with arrival rate λ , job size distribution X , and load $\rho = \lambda\mathbf{E}[X]$.

(i) We define the function

$$\sigma_X^{-1}(s) = s - \lambda(1 - \mathcal{L}[X](s)).$$

Note that $\sigma_X^{-1}(s) = \infty$ if and only if $\mathcal{L}[X](s) = \infty$.

(ii) We define σ_X to be the the inverse of σ_X^{-1} , choosing the branch that passes through the origin. That is, for $s \geq \inf_r \sigma_X^{-1}(r)$, we define $\sigma_X(s)$ to be the greatest real solution to

$$\sigma_X(s) = s + \lambda(1 - \mathcal{L}[X](\sigma_X(s))).$$

If $s < \inf_r \sigma_X^{-1}(r)$, then no such solution exists, so we define $\sigma_X(s) = -\infty$.

(iii) We define the *work-in-system transform*

$$\mathcal{L}[W_X](s) = \frac{s(1 - \rho)}{\sigma_X^{-1}(s)}.$$

Note that all of the above definitions depend on both λ and X , However, because the following discussion considers a fixed arrival rate λ and varies only the job size distribution X , we keep λ implicit to reduce clutter. Additionally, we assume in all uses of the above definitions that $\rho < 1$.

One may recognize the functions defined in Definition EC.3.3 as core to the theory of the M/G/1 with job size distribution X (Harchol-Balter, 2013).

- The work-in-system transform is, as suggested by its name, the Laplace-Stieltjes transform of the equilibrium distribution W_X of the total workload in the M/G/1.
- The function σ_X is related to busy periods in the M/G/1. Specifically, the length of a busy period started by initial workload V has Laplace-Stieltjes transform $\mathcal{L}[V](\sigma_X(s))$.

It turns out that throughout Section 9.3, all of the Laplace-Stieltjes transforms to which we apply (9.1) are of the form $\mathcal{L}[W_X]$ or $\mathcal{L}[W_X] \circ \sigma_Y$, the latter meaning $s \mapsto \mathcal{L}[W_X](\sigma_Y(s))$, for nicely light-tailed job size distributions X and Y (Definition 5.2). Specifically, X is the system's job size distribution, and Y is either X or a truncation $\min\{X, a^*\}$. Therefore, to justify the uses of (9.1) using Lemma EC.3.2, it suffices to prove Propositions EC.3.4 and EC.3.5 below.⁸

⁸ While Definition EC.3.3 assumes a single arrival rate λ , Proposition EC.3.5 easily generalizes to the case where $\mathcal{L}[W_X]$ and σ_Y are defined using different arrival rates.

PROPOSITION EC.3.4. *For any nicely light-tailed job size distribution X ,*

- (i) $\gamma(\mathcal{L}[W_X]) \in (-\infty, 0)$; and
- (ii) $\mathcal{L}[W_X]$ has a first-order pole at $\gamma(\mathcal{L}[W_X])$, so it is regularly varying at $\gamma(\mathcal{L}[W_X])$.

PROPOSITION EC.3.5. *For any nicely light-tailed job size distributions X and Y ,*

- (i) $\gamma(\mathcal{L}[W_X] \circ \sigma_Y) \in (-\infty, 0)$, and
- (ii) either $\mathcal{L}[W_X] \circ \sigma_Y$ or $(\mathcal{L}[W_X] \circ \sigma_Y)'$ is regularly varying at $\gamma(\mathcal{L}[W_X] \circ \sigma_Y)$.

Our approach is as follows. We first prove Proposition EC.3.4. We then prove a lemma characterizing σ_X , which we use in conjunction with Proposition EC.3.4 to prove Proposition EC.3.5

Proof of Proposition EC.3.4. Recall from Definition EC.3.3 that $\mathcal{L}[W_X](s) = s(1 - \rho)/\sigma_X^{-1}(s)$, so we focus on σ_X^{-1} . Because $\mathcal{L}[X]$ is a mixture of exponentials, σ_X^{-1} is convex, so it has at most two real roots. It is well-known that under the assumption on X made in Definition 5.2, σ_X^{-1} has a first-order root at 0 and a negative first-order root (Abate and Whitt, 1997; Mandjes and Nuyens, 2005), the latter of which is $\gamma(\mathcal{L}[W_X])$, but we give a brief proof for completeness. One can compute $\sigma_X^{-1}(0) = 0$ and $(\sigma_X^{-1})'(0) = 1 - \rho$, so σ_X^{-1} has a first-order root at 0. Definition 5.2 implies $\mathcal{L}[X](\gamma(\mathcal{L}[X])) = \infty$, so $\sigma_X^{-1}(\gamma(\mathcal{L}[X])) = \infty$. This means σ_X^{-1} has another first-order root in $(\gamma(\mathcal{L}[X]), 0)$. \square

LEMMA EC.3.6. *For any nicely light-tailed job size distribution X ,*

- (i) $\gamma(\sigma_X) \in (-\infty, 0)$;
- (ii) $\sigma_X(\gamma(\sigma_X)) \in (-\infty, 0)$; and
- (iii) there exist $C_0, C_1 > 0$ such that in the $s \downarrow 0$ limit,

$$\begin{aligned}\sigma_X(\gamma(\sigma_X) + s) &= \sigma_X(\gamma(\sigma_X)) + C_0\sqrt{s} \pm \Theta(s), \\ \sigma_X'(\gamma(\sigma_X) + s) &= \frac{C_1}{\sqrt{s}} \pm \Theta(1),\end{aligned}$$

so σ_X' is regularly varying at $\gamma(\sigma_X)$.

Proof. As in the proof of Proposition EC.3.4, we again use the fact that σ_X^{-1} is convex, has roots at a negative number and at zero, and is negative between its roots. Specifically, this fact implies that σ_X^{-1} has a finite negative global minimum. By Definition EC.3.3, this minimum is $\gamma(\sigma_X)$, and the value at which the minimum is attained is $\sigma_X(\gamma(\sigma_X))$ proving (i) and (ii).

It remains only to prove (iii). The fact that Laplace-Stieltjes transforms are analytic in the interior of their domains of convergence implies that σ_X^{-1} can be written as a Taylor series about $\gamma(\sigma_X)$ whose first nonzero coefficient is quadratic, i.e. for some constant $K > 0$,

$$\sigma_X^{-1}(s) = Ks^2 \pm \Theta(s^3).$$

An extension of the Lagrange inversion theorem (DLMF, §1.10(vii)) implies that the inverse of σ_X^{-1} , namely σ_X , may thus be written in the desired form. The desired form for σ'_X , which completes (iii), then follows from

$$\begin{aligned} (\sigma_X^{-1})'(s) &= 2Ks \pm \Theta(s^2), \\ \sigma'_X(s) &= \frac{1}{(\sigma_X^{-1})'(\sigma_X(s))}. \end{aligned} \quad \square$$

Proof of Proposition EC.3.5. There are three cases to consider:

- $\gamma(\mathcal{L}[W_X]) > \sigma_Y(\gamma(\sigma_Y))$,
- $\gamma(\mathcal{L}[W_X]) < \sigma_Y(\gamma(\sigma_Y))$, and
- $\gamma(\mathcal{L}[W_X]) = \sigma_Y(\gamma(\sigma_Y))$.

For an intuitive grasp of these cases, it is helpful to imagine decreasing s starting at $s = 0$, tracking the behavior of $\mathcal{L}[W_X](\sigma_Y(s))$ as s decreases.

If $\gamma(\mathcal{L}[W_X]) > \sigma_Y(\gamma(\sigma_Y))$, then at some point before $s = s^*$ reaches $\gamma(\sigma_Y)$, meaning for some $s^* \in (-\gamma(\sigma_Y), 0)$, we have $\gamma(\mathcal{L}[W_X]) = \sigma_Y(s^*)$. This means $\gamma(\mathcal{L}[W_X] \circ \sigma_Y) = s^*$. The Lagrange inversion theorem (DLMF, §1.10(vii)) and the fact that $s > \gamma(\sigma_Y)$ imply that σ_Y can be linearly approximated near s^* , so the result follows from Proposition EC.3.4.

If $\gamma(\mathcal{L}[W_X]) < \sigma_Y(\gamma(\sigma_Y))$, then in contrast to the previous case, s reaches $\gamma(\sigma_Y)$, the last point at which $\sigma_Y(s)$ is finite, before $\sigma(s)$ reaches the pole of $\mathcal{L}[W_x]$. This means $\gamma(\mathcal{L}[W_X] \circ \sigma_Y) = \gamma(\sigma_Y)$. Similarly to the previous case, we can linearly approximate $\mathcal{L}[W_x]$ near $\gamma(\sigma_Y)$, so the result follows from Lemma EC.3.6.

If $\gamma(\mathcal{L}[W_X]) = \sigma_Y(\gamma(\sigma_Y))$, then roughly speaking, both of the previous cases' events happen simultaneously: just as s reaches $\gamma(\sigma_Y)$, the last point at which $\sigma_Y(s)$ is finite, $\sigma_Y(s)$ reaches the pole of $\mathcal{L}[W_X]$. Combining Proposition EC.3.4 and Lemma EC.3.6 implies that in the $s \downarrow \gamma(\sigma_Y)$ limit, we can approximate $\mathcal{L}[W_X](\sigma_Y(s))$ as

$$\mathcal{L}[W_X](\sigma_Y(\gamma(s))) = \frac{K_0}{\sigma_Y(\gamma(s))} \pm \Theta(1) = \frac{K_1}{\sqrt{s - \gamma(\sigma_Y)}} \pm \Theta(1)$$

for some constants $K_0, K_1 > 0$, from which the result follows. □

EC.3.3. Expanding the Definition of Nicely Light-Tailed Job Size Distributions

The class of light-tailed distributions we consider in Definition 5.2, namely what Abate and Whitt (1997) call ‘‘Class I’’ distributions, is well behaved enough for Propositions EC.3.4 and EC.3.5 to hold. More generally, our results apply to any job size distribution with positive decay rate for which one can show Propositions EC.3.4 and EC.3.5. In particular, this includes many distributions that Abate and Whitt (1997) call ‘‘Class II’’. These are job size distributions X such that $\mathcal{L}[X](\gamma(\mathcal{L}[X])) < \infty$.

In order to prove Propositions EC.3.4 and EC.3.5 for Class II job size distributions, one would need to assume a regularity condition. We believe it would suffice to assume that $\mathcal{L}[X]'$ is regularly varying at $\gamma(\mathcal{L}[X])$. The main change to the proofs would be additional casework. For example, it may be that $\mathcal{L}[W_X]$ still has a first-order pole, or it may be that it diverges without a pole because $\mathcal{L}[X]$ does. See Abate and Whitt (1997) and references therein for additional discussion.

More generally, it likely suffices to assume that some higher-order derivative $\mathcal{L}[X]^{(n)}$ is regularly varying at $\gamma(\mathcal{L}[X])$, as the result of Mimica (2016, Corollary 1.3) we use applies to higher derivatives as well. Other results of Mimica (2016) may allow one to relax the assumption even further.