When Does the Gittins Policy Have Asymptotically Optimal Response Time Tail?

Ziv Scully
Carnegie Mellon University
zscully@cs.cmu.edu

Lucas van Kreveld
University of Amsterdam
l.r.vankreveld@uva.nl

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.

1. INTRODUCTION

Scheduling to minimize response time (a.k.a. sojourn time) of single-server queueing systems 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 one might hope to minimize. There is significant work on minimizing *mean response time* $\mathbf{E}[T]$, which is the average response time of all jobs in a long arrival sequence [1, 7, 10].

Much less is known about minimizing the *tail of response time* $\mathbf{P}[T > t]$, which is the probability 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 $\mathbf{P}[T > t]$ in the $t \to \infty$ limit [4, 5, 9, 12, 13]. In this work, we study the M/G/1, a classic single-server queueing model, and ask the following question.

Question 1.1. Does any scheduling policy *simultaneously* optimize the mean and asymptotic tail of response time?

Prior work answers Question 1.1 when job sizes (a.k.a. service times) 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 [10]. However, SRPT's tail performance depends on the job size distribution.

- If the job size distribution is *heavy-tailed* (Definition 2.3), then SRPT is *tail-optimal*, meaning it has the best possible asymptotic tail decay (Definition 2.4).
- If the job size distribution is *light-tailed* (Definition 2.6), then SRPT is *tail-pessimal*, meaning it has the worst possible asymptotic tail decay (Definition 2.7).

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 [7]. 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) [1, 2].

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 does Gittins optimize the asymptotic tail of 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 [12, Corollary 3.5].
- In the light-tailed case, Gittins sometimes reduces to FCFS or FB [1, 2]. For light-tailed job sizes, FCFS is tail-optimal [5, 13], but FB is tail-pessimal [8].

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 suboptimal, or even tail-pessimal. And 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?

1.1 Contributions

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.

- (Section 3.1) If the job size distribution is heavy-tailed, then *Gittins is always tail-optimal.*
- (Section 3.2) If the job size distribution is light-tailed, then *Gittins can be tail-optimal, tail-pessimal, or in between.* In the cases where Gittins is tail-pessimal, we construct a modified policy that is not tail-pessimal but has mean response time arbitrarily close to Gittins's.

The key tool we use to analyze Gittins's asymptotic response time tail is the *SOAP* framework [11, 12]. 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. Underlying our Gittins results is a general tail analysis of SOAP policies.

2. MODEL AND PRELIMINARIES

We consider an M/G/1 queue with arrival rate λ, job size distribution X, and load $\rho = \lambda \mathbf{E}[X]$. For the tail of the job size distribution, we write $\overline{F}(t) = \mathbf{P}[X > t]$. We denote the maximum job size by $x_{\max} = \inf\{t \geq 0 \mid \overline{F}(t) = 0\}$, allowing $x_{\max} = \infty$. We write T_π for the M/G/1's response time distribution under policy π.

In the rest of this section, we introduce SOAP policies in general, the Gittins policy in particular, and background on heavy- and light-tailed job size distributions.

2.1 SOAP Policies and the Gittins Policy

The Gittins policy 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* [11, 12], which gives a response time analysis of the following broad class of policies.

Definition 2.1. A *SOAP policy* is a policy π specified by a *rank function* $r_\pi : [0, x_{\max}) \to \mathbb{R}$. Policy π assigns rank $r_\pi(a)$ to a job at age a.[1] 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.

Definition 2.2. The *Gittins* policy is the SOAP policy with rank function

$$r_{\text{Gittins}}(a) = \inf_{b > a} \frac{\int_a^b \overline{F}(t)\,\mathrm{d}t}{\overline{F}(a) - \overline{F}(b)}.$$

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

As is standard [11, Appendix B], we assume rank functions are piecewise-continuous and piecewise-monotonic with finitely many pieces in any compact interval. This holds for Gittins under very mild conditions [2].

2.2 Background on Heavy-Tailed Job Sizes

Definition 2.3 (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 $\overline{F}(\cdot)$ is of intermediate regular variation [6], meaning $\liminf_{\varepsilon \downarrow 0} \liminf_{x \to \infty} \overline{F}((1+\varepsilon)x)/\overline{F}(x) = 1$.

(ii) There exist $\beta \geq \alpha > 1$ such that the upper and lower Matuszewska indices of $\overline{F}(\cdot)$ are in $(-\beta, -\alpha)$ [3, Section 2.1]. This implies that there exists $C > 0$ such that for all sufficiently large $x_2 \geq x_1$,

$$\frac{1}{C}\left(\frac{x_2}{x_1}\right)^{-\beta} \leq \frac{\overline{F}(x_2)}{\overline{F}(x_1)} \leq C\left(\frac{x_2}{x_1}\right)^{-\alpha}.$$

In informal discussion, we omit "nicely".

Definition 2.4 (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* if

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

Loosely speaking, tail-optimality holds if large jobs have a response time of approximately $1/(1-\rho)$ times their size.

2.3 Background on Light-Tailed Job Sizes

Definition 2.5. The *decay rate* of random variable V, denoted $d(V)$, is

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

Higher decay rates correspond to asymptotically lighter tails.

Definition 2.6 (Light-Tailed Job Size Distribution)**.** We say a job size distribution X is *nicely light-tailed* if $x_{\max} < \infty$ or $d(X) > 0$. In informal discussion, we omit "nicely".

Definition 2.7 (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 work-conserving policies. In informal discussion, we omit "log-".

3. MAIN RESULTS

Our main results characterize the asymptotic tail behavior of Gittins and other SOAP policies. The situation in the heavy-tailed case is simple: *Gittins is always tail-optimal.*

- Theorem 3.3 gives a sufficient condition under which a SOAP policy is tail-optimal for heavy-tailed job sizes.
- Theorem 3.4 shows that for heavy-tailed job sizes, Gittins always satisfies this sufficient condition, and is thus always tail-optimal.

The situation in the light-tailed case is more complicated: *Gittins can be optimal, pessimal, or in between.*

- Theorem 3.5 classifies SOAP policies into tail-optimal, -intermediate, and -pessimal for light-tailed job sizes.
- Theorem 3.7 shows that for light-tailed job sizes, Gittins can be any of tail-optimal, -intermediate, or -pessimal.

The fact that Gittins can be tail-pessimal raises a question: can we improve tail performance while only slightly degrading mean response time? We answer this affirmatively.

- Theorem 3.9 shows that making a small change to the Gittins rank function results in only a small change to mean response time.
- Theorem 3.10 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.

Our results use the following general definitions, which apply to any SOAP policy.

[1]The full SOAP definition is more general [11], but the given definition suffices for our unknown-size setting.

Definition 3.1.
(i) The *worst ever rank* of a job of size x is defined by $w_x = \sup_{0 \le a < x} r(a)$.
(ii) The *worst age* is $a^* = \inf\{a \ge 0 \mid \forall b \in (a, x_{\max}), r(a) \ge r(b)\}$, namely the first age at which a job has the maximum rank. If the rank function is unbounded, which only happens when $x_{\max} = \infty$ (Section 2.1), then $a^* = \infty$.
(iii) A *w-interval* is an interval (b, c) with $0 \le b < c \le x_{\max}$ such that $r(a) \le w$ for all $a \in (b, c)$.

3.1 Results for Heavy-Tailed Job Sizes

Condition 3.2. There exist constants $K > 0$; $\zeta \in [0, \infty)$; $\theta \in [0, \infty)$; and $\eta \in [\max\{1, \zeta + \theta\}, \infty]$ such that for sufficiently large x, the following hold for any w_x-interval (b, c):
(i) If $b \ge x$, then $c - b \le K b^\zeta x^\theta$.
(ii) $c \le K x^\eta$.

Condition 3.2 is a condition on SOAP policies. It is a sharper version of the condition used by Scully et al. [12, Assumption 3.2], mainly because it adds an extra parameter θ. As such, the following theorem generalizes their main result [12, Theorem 3.3].

Theorem 3.3. *Consider an M/G/1 with nicely heavy-tailed job size distribution under a SOAP policy. Condition 3.2 implies the policy is tail-optimal if*

$$\zeta + (\theta - 1)^+ - \frac{(1-\theta)^+}{\eta} < \frac{\alpha - 1}{\beta}. \tag{3.1}$$

Theorem 3.4. *Consider an M/G/1 with nicely heavy-tailed job size distribution. Gittins satisfies Condition 3.2 with $\zeta = 0$, $\theta = 1$, and $\eta = \infty$, so it is tail-optimal.*

3.2 Results for Light-Tailed Job Sizes

Theorem 3.5. *Consider an M/G/1 with nicely light-tailed job size distribution under a SOAP policy. 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 3.5 to the Gittins policy, we need to characterize how the job size distribution X affects Gittins's worst age a^*. Results of Aalto et al. [1, 2] connect the following classes of distributions to Gittins's worst age a^*.

Definition 3.6. We define two classes of distributions.
- We say X is *New Better than Used in Expectation*, writing $X \in \mathsf{NBUE}$, if for all $a \ge 0$,
$$\mathbf{E}[X] \ge \mathbf{E}[X - a \mid X > a].$$
- We say X is *Eventually New Better than Used in Expectation*, writing $X \in \mathsf{ENBUE}$, if there exists $a_0 \ge 0$ such that $(X - a_0 \mid X > a_0) \in \mathsf{NBUE}$. That is, $X \in \mathsf{ENBUE}$ if there exists $a_0 \ge 0$ such that for all $a \ge a_0$,
$$\mathbf{E}[X - a_0 \mid X > a_0] \ge \mathbf{E}[X - a \mid X > a].$$

Theorem 3.7. *Consider an M/G/1 with nicely light-tailed job size distribution X. Gittins is*
- *log-tail-optimal if $X \in \mathsf{NBUE}$,*
- *log-tail-intermediate if $X \in \mathsf{ENBUE} \setminus \mathsf{NBUE}$, and*
- *log-tail-pessimal if $X \notin \mathsf{ENBUE}$.*

Fortunately, in most cases where Gittins is tail-pessimal, modifying Gittins yields a tail-intermediate policy with near-optimal mean response time.

Definition 3.8. A SOAP policy π is a *q-approximate Gittins policy* if there exists a constant $m > 0$ such that for all $a \ge 0$,

$$m \le \frac{r_\pi(a)}{r_{\text{Gittins}}(a)} \le mq.$$

Theorem 3.9. *Consider an M/G/1. For any $q \ge 1$ and any q-approximate Gittins policy π,*

$$\mathbf{E}[T_\pi] \le q\mathbf{E}[T_{\text{Gittins}}].$$

Theorem 3.10. *Consider an M/G/1 with nicely light-tailed job size distribution X. Suppose a job's expected remaining size is uniformly bounded at all ages, meaning $x_{\max} < \infty$ or $\sup_{a \ge 0} \mathbf{E}[X - a \mid X > a] < \infty$. Then for all $\varepsilon > 0$, there exists a $(1+\varepsilon)$-approximate Gittins policy that is log-tail-optimal or log-tail-intermediate.*

Acknowledgements. We thank Adam Wierman, Onno Boxma, and Jan-Pieter Dorsman for helpful discussions. Ziv Scully was supported by NSF grant nos. CMMI-1938909 and CSR-1763701. Lucas van Kreveld was supported by the NWO through Gravitation grant NETWORKS-024.002.003.

References

[1] Samuli Aalto, Urtzi Ayesta, and Rhonda Righter. 2009. On the Gittins Index in the M/G/1 Queue. *Queueing Systems* 63, 1-4 (Dec. 2009), 437–458.
[2] Samuli Aalto, Urtzi Ayesta, and Rhonda Righter. 2011. Properties of the Gittins Index with Application to Optimal Scheduling. *Probability in the Engineering and Informational Sciences* 25, 3 (July 2011), 269–288.
[3] Nicholas H. Bingham, Charles M. Goldie, and Jef L. Teugels. 1987. *Regular Variation.* Number 27 in Encyclopedia of Mathematics and Its Applications. Cambridge University Press, Cambridge, UK.
[4] Sem C. Borst, Onno J. Boxma, Rudesindo Núñez-Queija, and Bert Zwart. 2003. The Impact of the Service Discipline on Delay Asymptotics. *Performance Evaluation* 54, 2 (Oct. 2003), 175–206.
[5] Onno J. Boxma and Bert Zwart. 2007. Tails in Scheduling. *ACM SIGMETRICS Performance Evaluation Review* 34, 4 (March 2007), 13–20.
[6] Daren B. H. Cline. 1994. Intermediate Regular and Π Variation. *Proceedings of the London Mathematical Society* s3-68, 3 (May 1994), 594–616.
[7] John C. Gittins. 1989. *Multi-Armed Bandit Allocation Indices* (first ed.). Wiley, Chichester, UK.
[8] Michel Mandjes and Misja Nuyens. 2005. Sojourn Times in the M/G/1 FB Queue with Light-Tailed Service Times. *Probability in the Engineering and Informational Sciences* 19, 3 (2005), 351–361.
[9] Rudesindo Núñez-Queija. 2002. Queues with Equally Heavy Sojourn Time and Service Requirement Distributions. *Annals of Operations Research* 113, 1/4 (July 2002), 101–117.
[10] Linus E. Schrage. 1968. A Proof of the Optimality of the Shortest Remaining Processing Time Discipline. *Operations Research* 16, 3 (June 1968), 687–690.
[11] Ziv Scully, Mor Harchol-Balter, and Alan Scheller-Wolf. 2018. SOAP: One Clean Analysis of All Age-Based Scheduling Policies. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 1, Article 16 (April 2018), 30 pages.
[12] Ziv Scully, Lucas van Kreveld, Onno J. Boxma, Jan-Pieter Dorsman, and Adam Wierman. 2020. Characterizing Policies with Optimal Response Time Tails under Heavy-Tailed Job Sizes. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 2, Article 30 (June 2020), 33 pages.
[13] Alexander L. Stolyar and Kavita Ramanan. 2001. Largest Weighted Delay First Scheduling: Large Deviations and Optimality. *Annals of Applied Probability* 11, 1 (2001), 1–48.