On the Gittins index in the $M/G/1$ queue*

Samuli Aalto[1], Urtzi Ayesta[2,3,†], Rhonda Righter[4]

[1] TKK Helsinki University of Technology, Finland

[2] BCAM - Basque Center for Applied Mathematics, Spain

[3] LAAS-CNRS, France

[4] University of California at Berkeley, USA

Abstract

For an $M/G/1$ queue with the objective of minimizing the mean number of jobs in the system, the Gittins index rule is known to be optimal among the set of non-anticipating policies. We develop properties of the Gittins index. For a single-class queue it is known that when the service time distribution is of type Decreasing Hazard Rate (New Better than Used in Expectation), the Foreground-Background (First-Come-First-Served) discipline is optimal. By utilizing the Gittins index approach, we show that in fact, Foreground-Background and First-Come-First-Served are optimal *if and only if* the service time distribution is of type Decreasing Hazard Rate and New Better than Used in Expectation, respectively. For the multi-class case, where jobs of different classes have different service distributions, we obtain new results that characterize the optimal policy under various assumptions on the service time distributions. We also investigate distributions whose hazard rate and mean residual lifetime are not monotonic.

1 Introduction

We consider a single-server queue with Poisson arrivals and generally distributed service times. The objective is to minimize the mean delay (i.e., sojourn time), which by Little's law is equivalent to minimizing the mean number of jobs in the system, or the holding costs. The characteristics of the optimal scheduling discipline depend on the information available to the scheduler (server). Scheduling disciplines are often classified into

*The paper is intended for the Special Issue of *Queueing Systems* related to the Erlang Centennial Conference.

†Corresponding author: U. Ayesta, email: ayesta@bcamath.org

anticipating or non-anticipating. We say that a scheduling policy is anticipating if the remaining service times of jobs are known by the scheduler. A non-anticipating discipline does not have knowledge on the remaining service times of jobs, but instead it is assumed that it has knowledge of the service time already provided to each of the jobs present in the queue.

While the optimal anticipating discipline, Shortest-Remaining-Processing-Time (SRPT), minimizes the number of jobs sample-path wise [16, 19] for completely general service times and arrival process, the optimal *non-anticipating* discipline depends on the characteristics of the service time distribution. If the service time distribution belongs to the Decreasing Hazard Rate (DHR) class, the Foreground-Background (FB) policy (a.k.a. Least-Attained-Service (LAS)), in which the job with the least attained service is always served, is optimal [24, 25, 13], whereas the ordinary First-Come-First-Served (FCFS) discipline, or any other non-preemptive discipline, minimizes the mean delay for the service time distributions belonging to the New Better than Used in Expectation (NBUE) class [15]. These optimality results of FB and FCFS do not require the Poisson arrivals assumption. In multi-class systems with feedback the optimal scheduling policy among non-preemptive policies is given by Klimov's index policy [7, 8] (see also Meilijson and Weiss [10] and Meilijson and Yechiali [11]).

Gittins [5] considered the $M/G/1$ queue and proved that the so-called Gittins index policy minimizes the mean delay among all non-anticipating preemptive policies. The Gittins rule calculates, based on the attained service of jobs, the optimum quantum of service that a job should next obtain. Gittins derived this result as a byproduct of his groundbreaking results on the multi-armed bandit problem. The literature on multi-armed bandit related papers that build on Gittins' result is huge, see, e.g., [21, 23, 22, 20, 3]. However, Gittins' optimality result in the context of the $M/G/1$ queue has not been fully exploited, and it has not received the attention it deserves. We refer to [2] and [12] for two recent works on the application of Gittins' index to an $M/G/1$ queue. In [2] the optimal scheduling is studied when the service time distribution has a DHR tail, and in [12] the expected conditional sojourn time of the Gittins index policy is calculated for several examples of service time distributions.

In this paper we use the framework developed by Gittins to prove some new results related to optimal scheduling in the $M/G/1$ queue. In Section 2 we introduce the basic notation and the classes of service time distributions we use throughout the paper. In addition, we introduce the Gittins index as an optimal value of an efficiency function and the corresponding index policy. The Gittins index policy can be defined either as a continuous control serving always the job(s) with highest index, or as a decision process with discrete but random decision epochs related to the allocated service quantum and new arrivals. In Section 3 we give basic results for the Gittins index that are needed in the forthcoming sections. As a by-product we also get a verification for the equivalence of the two different definitions for the Gittins index policy. In Sections 4 and 5, we characterize the service time distribution classes mentioned in Section 2 by means of what we call the efficiency function and the Gittins index, respectively. The characterizations given in these sections seem to be novel. Section 6 includes the main results of the paper. Based on the characterizations given in earlier sections, we show that FB and FCFS are optimal if and only if the service time distribution is of type DHR and NBUE, respectively. Thus, we prove that the assumptions DHR and NBUE are both sufficient and necessary for the optimality of FB and FCFS, respectively. While their sufficiency is well-known, their necessity is proved here for the first time. In Section 7 we apply Gittins' framework to a multi-class single-server queue with holding costs and other classes of service time distributions besides NBUE and DHR. In Section 8 we consider service time distributions with non-monotonic hazard rates. Section 9 concludes the paper.

2 Preliminaries

Consider an $M/G/1$ queue with arrival rate λ, mean service time $\mathbb{E}[S]$, and load $\rho = \lambda\mathbb{E}[S] < 1$. Jobs are served according to a work conserving and non-anticipating scheduling discipline π. Let Π denote the family of such disciplines. We denote by $\mathbb{E}[T^\pi]$ and $\mathbb{E}[N^\pi]$ the mean delay and the mean number of jobs in steady-state under discipline $\pi \in \Pi$, respectively. Note that by Little's law the relation $\mathbb{E}[N^\pi] = \lambda\mathbb{E}[T^\pi]$ holds, thus minimizing the mean number of jobs results in minimizing the mean delay.

Let $F(x) = \mathbb{P}(S \le x)$, $x \ge 0$, denote the cumulative distribution of the service time of any job. Define $\overline{F}(x) = 1 - F(x)$, and assume that $\overline{F}(x) > 0$ for all $x \ge 0$. In addition, we only consider distributions with a density function $f(x)$ that is right-continuous with left-limits.

2.1 Service time distribution classes

The hazard rate $h(x)$, $x \ge 0$, of a service time distribution is defined by

$$h(x) = \frac{f(x)}{\overline{F}(x)} = \frac{f(x)}{\int_0^\infty f(x+y)\,\mathrm{d}y}.$$

Note that it is right-continuous ($h(x) = h(x^+)$) with left-limits $h(x^-)$. A service time distribution belongs to the class DHR (Decreasing Hazard Rate) if $h(x)$ is decreasing[1] for all x, i.e., $h(x) \ge h(y)$ whenever $x \le y$. The class IHR (Increasing Hazard Rate) is defined correspondingly.

[1]Throughout the paper we use the terms *decreasing* and *increasing* in their weak form so that the corresponding functions need not be *strictly* decreasing or increasing.

In addition, we define, for all $x \geq 0$,

$$H(x) = \frac{\int_0^\infty f(x+y)\,\mathrm{d}y}{\int_0^\infty \overline{F}(x+y)\,\mathrm{d}y} = \frac{\overline{F}(x)}{\int_0^\infty \overline{F}(x+y)\,\mathrm{d}y}. \tag{1}$$

We note that the function $H(x)$ is related to the mean remaining service time as follows:

$$\mathbb{E}[S - x \mid S > x] = \frac{\int_0^\infty \overline{F}(x+y)\,\mathrm{d}y}{\overline{F}(x)} = \frac{1}{H(x)}.$$

A service time distribution belongs to the class IMRL (Increasing Mean Residual Lifetime) if $1/H(x)$ is increasing for all x, i.e., $1/H(x) \leq 1/H(y)$ whenever $x \leq y$. A service time distribution belongs to the class NWUE (New Worse than Used in Expectation) if $1/H(0) \leq 1/H(x)$ for all x. The classes DMRL (Decreasing Mean Residual Lifetime) and NBUE (New Better than Used in Expectation) are defined correspondingly.

The following two lemmas give a relation between the monotonicity properties of functions $h(x)$ and $H(x)$.

Lemma 1 *$H(x)$ is strictly increasing (decreasing) at x if and only if $h(x) < H(x)$ ($h(x) > H(x)$) at x. Thus $H(x)$ has a critical point at x if and only if $h(x) = H(x)$ at x.*

Proof: The result is straightforwardly obtained by differentiating (1). □

Lemma 2 *If the service time distribution is IHR or DHR, then $h(x) \geq h(y)$ implies that $H(x) \geq H(y)$.*

Proof: If the service time distribution is of type IHR [DHR], then its remaining service time $X_t = (S - t \mid S > t)$ is stochastically decreasing [increasing] in t, and therefore, its mean remaining service time, $\mathbb{E}[X_t] = 1/H(t)$ is decreasing [increasing] in t. □

From Lemma 2 it follows that DHR $\subset$ IMRL $\subset$ NWUE and IHR $\subset$ DMRL $\subset$ NBUE. In addition, all the distributions in the class NWUE [NBUE] have a greater [smaller] coefficient of variation than that of the exponential distribution (which is 1). We refer to [18, Chapter 1] for more details on this type of relations.

2.2 Gittins index and the efficiency function

The Gittins index $G(a)$ [5] is associated with the service time distribution of a job. It is defined for each value a of the attained service. For the definition, we first need to introduce the efficiency function $J(a, \Delta)$.

The *efficiency function* $J(a, \Delta)$, $a, \Delta \geq 0$, is defined as follows. For any $0 < \Delta < \infty$, let

$$J(a, \Delta) = \frac{\int_0^\Delta f(a+t)\,\mathrm{d}t}{\int_0^\Delta \overline{F}(a+t)\,\mathrm{d}t} = \frac{\overline{F}(a) - \overline{F}(a+\Delta)}{\int_0^\Delta \overline{F}(a+t)\,\mathrm{d}t}. \tag{2}$$

In addition, let

$$J(a,0)=\frac{f(a)}{\overline{F}(a)}=h(a),\quad J(a,\infty)=\frac{\overline{F}(a)}{\int_0^\infty \overline{F}(a+t)\,\mathrm{d}t}=H(a).$$

Note that, for any $\Delta>0$,

$$J(a,\Delta)=\frac{\mathbb{P}(S-a\le\Delta\mid S>a)}{\mathbb{E}[\min\{S-a,\Delta\}\mid S>a]}.$$

Thus, for a job that has attained service a and is assigned Δ units of service, the efficiency function $J(a,\Delta)$ is the ratio between (i) the probability that the job will complete with a service quantum of Δ (interpreted as reward) and (ii) the expected service time that a job with attained service a and service quantum Δ will require from the server (interpreted as investment).

Note further that $J(a,\Delta)$ is continuous with respect to Δ. In addition, the one-sided partial derivatives (with respect to Δ) read for any $\Delta>0$ as follows:

$$\frac{\partial}{\partial\Delta}J(a,\Delta)=\frac{\overline{F}(a+\Delta)(h(a+\Delta)-J(a,\Delta))}{\int_0^\Delta \overline{F}(a+t)\,\mathrm{d}t}. \tag{3}$$

Whether the partial derivative is continuous at point (a,Δ) or not, depends on the continuity of $f(x)$ at point $x=a+\Delta$.

In the next sections, we will utilize the following construction several times. Let $a\le x<a+\Delta$. Then

$$J(a,\Delta)=\frac{\int_a^x f(t)\,\mathrm{d}t}{\int_a^{a+\Delta}\overline{F}(t)\,\mathrm{d}t}+\frac{\int_x^{a+\Delta} f(t)\,\mathrm{d}t}{\int_a^{a+\Delta}\overline{F}(t)\,\mathrm{d}t}=pJ(a,x-a)+(1-p)J(x,a+\Delta-x), \tag{4}$$

where $p\in[0,1)$ refers to

$$p=\frac{\int_a^x \overline{F}(t)\,\mathrm{d}t}{\int_a^{a+\Delta}\overline{F}(t)\,\mathrm{d}t}$$

and where we suppress the dependency on a, Δ, and x.

The *Gittins index* $G(a)$, $a\ge 0$, is defined by

$$G(a)=\sup_{\Delta\ge 0}J(a,\Delta). \tag{5}$$

In addition, for any $a\ge 0$, we define the *optimum quantum of service* by

$$\Delta^*(a)=\sup\{\Delta\ge 0\mid G(a)=J(a,\Delta)\}.$$

Note that $J(a,\Delta^*(a))=G(a)$ by definition.

Note further that, because $J(a,0)=h(a)$ and $J(a,\infty)=H(a)$, we have $G(a)\ge\max\{h(a),H(a)\}$. If $G(a)=H(a)$, then necessarily $\Delta^*(a)=\infty$. However, if $G(a)=h(a)$, then $\Delta^*(a)=0$ only if $h(a)=J(a,0)>J(a,\Delta)$ for all $\Delta>0$.

2.3 Gittins index policy

The main interest of Gittins [5] was on the so-called multi-armed bandit problem. The classic multi-armed bandit problem refers to a system with M bandits (or projects). At every time slot the gambler chooses to activate one of the bandit processes, which will then yield a reward and undergo a Markovian state transition, while all the other bandit processes are passive, i.e., they yield no reward, and their states remain frozen. The objective is to choose a policy for activating the bandits so as to maximize the total discounted reward. In an important result, Gittins proved that the structure of the optimal policy is in fact rather simple. An index can be associated to each bandit based solely on its own state. Then, at every decision epoch the bandit with highest index is chosen.

Extending his framework by allowing the arrival and departure of new projects, Gittins proved the optimality of the Gittins index for minimizing sojourn times in an $M/G/1$ queue, where the index rule is defined as follows.

Definition 1 *The* Gittins index policy $\pi^* \in \Pi$ *is the discipline that at every instant of time gives service to the job with highest index $G(a)$ among all jobs present in the system.*

We now recall the important optimality result proved by Gittins for a multi-class $M/G/1$ queue [5, Theorem 3.28]. See also [26, Theorem 4.7].

Theorem 1 *Consider an $M/G/1$ multi-class queue. Then, $\mathbb{E}[N^{\pi^*}] \leq \mathbb{E}[N^{\pi}]$ for any $\pi \in \Pi$.*

In the multi-class case, the Gittins index of a job is calculated using the service time distribution of its own class. For the time being we will concentrate on the single-class case. In Section 7 we will consider the multi-class case.

Unfortunately the proof of Theorem 1 is rather obscure, which has certainly contributed to the fact that it has not received the attention it deserves. We refer to [21] for an interesting paper where the approach of Gittins is used in the context of optimal scheduling in a single server queue with random arrivals and departures.

Sevcik [17] studied the optimal scheduling discipline in a system without arrivals, that is, when all jobs to be processed are available at time zero, and via an interchange argument he proved the optimality of the Smallest-Rank (SR) policy (see also [4, Section 6.5] or [6, Section 3.9]). The reciprocal of his rank is the product of the Gittins index and the holding cost rate, so the SR policy is equivalent to serving the job with the largest Gittins index when all cost rates are identical.

Rather than monitoring the Gittins indices continuously over time and always serving the job with the highest value, an alternative is to only make decisions when there are new arrivals, the current job completes, or when the current job (with the highest index) completes its optimal quantum of service. We refer to this as the Gittins index quantum policy, as defined below.

Definition 2 *The* Gittins index quantum policy $\tilde{\pi}^* \in \Pi$ *is the discipline that at every decision epoch picks the job with highest index value* $G(a)$*, and assigns it a service quantum* $\Delta^*(a)$*. The next decision epoch will take place when (i) the previously selected job receives* $\Delta^*(a)$ *units of service, (ii) the previously selected job completes its service, or (iii) a new job arrives to the queue.*

Note that with Definition 2, if the job with the highest index has a zero optimal quantum of service, its index must be continuously updated while it is being served, as in Definition 1. If there are multiple jobs that are all tied with highest indices, and they all have zero optimal service quanta, then they will be processed together (processor sharing) as long as they have highest indices with zero optimal service quanta. Also note that, in practice, we can increase our service quantum for the highest indexed job(s) to the point at which its index is no higher than the next highest indexed job. That is, suppose that there is a single job with highest index, call it job 1 with index $G(a_1)$, and suppose the next highest index is $G(a_2)$ for job 2. Let

$$\hat{\Delta} = \sup\{\Delta \geq \Delta^*(a_1) \mid J(a_1, \Delta) \geq G(a_2)\}.$$

Then, under the Gittins policy job 1 will be served until either it completes, there is a new arrival, or it receives $\hat{\Delta}$ units of service. Also note that even if an arrival occurs before job 1 completes or before $\hat{\Delta}$, job 2 will not be served under the Gittins policy until either job 1 completes or it receives a total of $\hat{\Delta}$ units of service (possibly with several interruptions by arrivals).

The next theorem shows that the Gittins index policy and the Gittins index quantum policy are indeed equivalent. The proof is a direct consequence of Lemma 4 and Corollary 1, which are given in Section 3.

Theorem 2 *The Gittins index policy* π^* *(Definition 1) and the Gittins index quantum policy* $\tilde{\pi}^*$ *(Definition 2) are the same sample-path wise.*

Sevcik [17] showed this equivalence when there are no job arrivals and jobs can have different holding costs (in which case the Gittins index includes the holding cost as a multiplicative factor). Gittins showed the same equivalence in the classic multi-armed bandit problem, see [5, Lemma 2.5].

Main idea

Before getting into the technical details, we present now the main idea behind the results of the paper. From Definition 1 and Theorem 1, we see that the function $G(a)$ characterizes completely the optimal policy. For instance it's easy to see that

$$\text{FB is optimal} \iff G(a) \text{ is decreasing for all } a.$$

Then, our main goal will be to characterize the shape of $G(a)$ for particular classes of service time distributions. For example we will show that (see Proposition 5)

$$G(a) \text{ is decreasing for all } a \iff \text{the service time distribution is DHR.}$$

Combining these two facts we derive the desired characterization (see Theorem 3) that the policy FB is optimal if and only if the service time distribution is DHR. We proceed similarly in the case of distributions of type NBUE.

3 Basic properties of the Gittins index

In this section we give several results for the Gittins index $G(a)$ that are needed in the forthcoming sections. While the results are not all new, we present them with explicit proofs for completeness.

We start with some results indicating that, when given the optimum quantum of service $\Delta^*(a)$, the Gittins index of a job does not go below the starting value $G(a)$ until the end of the quantum.

Lemma 3 *If $\Delta^*(a) > 0$, then $J(x, a + \Delta^*(a) - x) \geq G(a)$ for all $a \leq x < a + \Delta^*(a)$.*

Proof: Suppose that $\Delta^*(a) > 0$ and $a \leq x < a + \Delta^*(a)$. Now

$$J(a, \Delta^*(a)) = pJ(a, x - a) + (1 - p)J(x, a + \Delta^*(a) - x),$$

where $p \in [0, 1)$ refers to

$$p = \frac{\int_a^x \overline{F}(t)\,\mathrm{d}t}{\int_a^{a+\Delta^*(a)} \overline{F}(t)\,\mathrm{d}t}.$$

Since $J(a, x - a) \leq G(a) = J(a, \Delta^*(a))$, we have

$$J(a, \Delta^*(a)) \leq pJ(a, \Delta^*(a)) + (1 - p)J(x, a + \Delta^*(a) - x),$$

from which the claim clearly follows. □

From this lemma we immediately obtain the following corollary.

Corollary 1 *If $\Delta^*(a) > 0$, then $G(x) \geq G(a)$ for all $a \leq x < a + \Delta^*(a)$.*

Theorem 2, as stated in the previous section, follows from Corollary 1 and Lemma 4 below.

Lemma 4 *If $\Delta^*(a) > 0$, then $x + \Delta^*(x) \leq a + \Delta^*(a)$ for all $a \leq x < a + \Delta^*(a)$.*

Proof: For $\Delta^*(a) = \infty$, the result is trivial. Thus, we may assume that $0 < \Delta^*(a) < \infty$.

Consider what happens if there is $a < x < a + \Delta^*(a)$ such that $x + \Delta^*(x) > a + \Delta^*(a)$. Now

$$J(a, x + \Delta^*(x) - a) = pJ(a, \Delta^*(a)) + (1-p)J(a + \Delta^*(a), x + \Delta^*(x) - a - \Delta^*(a)) \tag{6}$$

where $p \in (0, 1)$ refers to

$$p = \frac{\int_a^{a+\Delta^*(a)} \overline{F}(t)\,\mathrm{d}t}{\int_a^{x+\Delta^*(x)} \overline{F}(t)\,\mathrm{d}t}.$$

From the definition of $\Delta^*(a)$, we have that $J(a, x + \Delta^*(x) - a) < J(a, \Delta^*(a))$. Thus, (6) implies that

$$J(a + \Delta^*(a), x + \Delta^*(x) - a - \Delta^*(a)) < J(a, \Delta^*(a)).$$

On the other hand, Lemma 3 implies that $J(x, a + \Delta^*(a) - x) \geq J(a, \Delta^*(a))$. Thus, we conclude that

$$J(a + \Delta^*(a), x + \Delta^*(x) - a - \Delta^*(a)) < J(x, a + \Delta^*(a) - x). \tag{7}$$

Now let $p' \in (0, 1)$ refer to

$$p' = \frac{\int_x^{a+\Delta^*(a)} \overline{F}(t)\,\mathrm{d}t}{\int_x^{x+\Delta^*(x)} \overline{F}(t)\,\mathrm{d}t}.$$

Then, by (7), we get

$$\begin{aligned} J(x, \Delta^*(x)) &= p'J(x, a + \Delta^*(a) - x) + (1-p')J(a + \Delta^*(a), x + \Delta^*(x) - a - \Delta^*(a)) \\ &\stackrel{(7)}{<} J(x, a + \Delta^*(a) - x), \end{aligned}$$

which is a contradiction, since $J(x, \Delta^*(x)) = \sup_{\Delta \geq 0} J(x, \Delta)$.

Thus there does not exist $a \leq x < a + \Delta^*(a)$ such that $x + \Delta^*(x) > a + \Delta^*(a)$, and the result follows. □

The following two corollaries follow directly from Lemma 4.

Corollary 2 *If* $\Delta^*(a) > 0$, *then* $\lim_{x \to (a+\Delta^*(a))^-} \Delta^*(x) = 0$.

Corollary 3 *Let* $x < y$. *If* $\Delta^*(y) = \infty$, *then* $\Delta^*(x) \leq y - x$ *or* $\Delta^*(x) = \infty$.

In the remainder of this section we state and prove several technical lemmas on properties of the Gittins index.

Lemma 5 *If* $\Delta^*(a) < \infty$, *then* $G(a + \Delta^*(a)) \leq G(a)$.

Proof: Suppose that $\Delta^*(a) < \infty$. If $\Delta^*(a) = 0$, then the claim is trivially true. Thus we may assume that $\Delta^*(a) > 0$. Let $\Delta^{**}(a) = \Delta^*(a + \Delta^*(a))$ so that

$$\Delta^{**}(a) = \sup\{\Delta \geq 0 \mid J(a + \Delta^*(a), \Delta) = G(a + \Delta^*(a))\}.$$

1° Assume that $\Delta^{**}(a) = 0$ so that

$$G(a + \Delta^*(a)) = J(a + \Delta^*(a), 0) = h(a + \Delta^*(a)). \tag{8}$$

Due to optimality of $\Delta^*(a)$, we have

$$\frac{\partial}{\partial \Delta} J(a, \Delta)\bigg|_{\Delta \to \Delta^*(a)^+} \leq 0.$$

By (8) and (3), this implies that

$$G(a + \Delta^*(a)) \overset{(8)}{=} h(a + \Delta^*(a)) \overset{(3)}{\leq} J(a, \Delta^*(a)) = G(a).$$

2° Now assume that $\Delta^{**}(a) > 0$. Consider what happens if

$$G(a + \Delta^*(a)) > G(a). \tag{9}$$

Now

$$\begin{aligned} J(a, \Delta^*(a) + \Delta^{**}(a)) &= pJ(a, \Delta^*(a)) + (1-p)J(a + \Delta^*(a), \Delta^{**}(a)) \\ &= pG(a) + (1-p)G(a + \Delta^*(a)), \end{aligned} \tag{10}$$

where $p \in (0, 1)$ refers to

$$p = \frac{\int_a^{a+\Delta^*(a)} \overline{F}(t)\,\mathrm{d}t}{\int_a^{a+\Delta^*(a)+\Delta^{**}(a)} \overline{F}(t)\,\mathrm{d}t}.$$

Combining (9) and (10) we conclude that

$$J(a, \Delta^*(a) + \Delta^{**}(a)) > G(a),$$

which is a contradiction, since $G(a) = \sup_{\Delta \geq 0} J(a, \Delta)$. Thus, the claim must be true also in this case. □

Next we explore the close relation between the Gittins index and the hazard rate.

Lemma 6

(i) If $\Delta^*(a) > 0$, then $G(a) \leq h((a + \Delta^*(a))^-)$.

(ii) If $\Delta^*(a) < \infty$, then $G(a) \geq h(a + \Delta^*(a))$.

(iii) If $0 < \Delta^*(a) < \infty$ and $h(x)$ is continuous, then $G(a) = G(a + \Delta^*(a)) = h(a + \Delta^*(a))$.

Proof: (i) Let a be such that $\Delta^*(a) > 0$. Note first that

$$\frac{\partial}{\partial \Delta} J(a, \Delta)\bigg|_{\Delta \to \Delta^*(a)^-} \geq 0.$$

By (3), for any $\Delta > 0$,

$$\frac{\partial}{\partial x} J(a, x)\bigg|_{x \to \Delta^-} \geq 0 \quad \Longleftrightarrow \quad h((a+\Delta)^-) \geq J(a, \Delta).$$

Thus, $h((a + \Delta^*(a))^-) \geq J(a, \Delta^*(a)) = G(a)$.

(ii) This can be proved by a similar argument as (i).

(iii) It follows from Lemma 5, item (i), and the continuity assumption that

$$G(a + \Delta^*(a)) \leq G(a) \leq h((a + \Delta^*(a))^-) = h(a + \Delta^*(a)) \leq G(a + \Delta^*(a)),$$

which proves the claim. □

In the two following lemmas, we identify the Gittins index at point a assuming that $G(x) \leq G(a)$ or $G(x) \geq G(a)$ for all $x \geq a$.

Lemma 7 *If $G(x) \leq G(a)$ for all $x \geq a$, then $G(a) = h(a)$.*

Proof: Choose $a \geq 0$, and assume that $G(x) \leq G(a)$ for all $x \geq a$. If $\Delta^*(a) = 0$, then, by definition, $G(a) = J(a, 0) = h(a)$.

Now assume that $\Delta^*(a) > 0$. Then, the assumption above and Lemma 3 together imply that, for all $a \leq x < a + \Delta^*(a)$,

$$G(x) = G(a). \tag{11}$$

Let $0 < \epsilon < \Delta^*(a)$, and define

$$p = \frac{\int_a^{a+\epsilon} \overline{F}(t)\,\mathrm{d}t}{\int_a^{a+\Delta^*(a)} \overline{F}(t)\,\mathrm{d}t}.$$

Note that $p \in (0, 1)$. Now

$$\begin{aligned} G(a) &= J(a, \Delta^*(a)) \\ &= pJ(a, \epsilon) + (1-p)J(a+\epsilon, \Delta^*(a) - \epsilon) \\ &\leq pJ(a, \epsilon) + (1-p)G(a+\epsilon). \end{aligned}$$

On the other hand, $J(a, \epsilon) \leq G(a)$ and, by (11), $G(a + \epsilon) = G(a)$. Thus, we conclude that $J(a, \epsilon) = G(a)$. Since this is true for any $0 < \epsilon < \Delta^*(a)$, we have $h(a) = J(a, 0) = \lim_{\Delta \to 0^+} J(a, \Delta) = G(a)$. □

Lemma 8 *If $G(x) \geq G(a)$ for all $x \geq a$, then $G(a) = H(a)$.*

Proof: Let $a \geq 0$, and assume that $G(x) \geq G(a)$ for all $x \geq a$.

Consider what happens if $\Delta^*(a) < \infty$. Then $J(a,\Delta) < J(a,\Delta^*(a))$ for all $\Delta \in (\Delta^*(a), \infty]$. Thus, there are $c > 0$ and $M \in (\Delta^*(a), \infty)$ such that $J(a,\infty) < c < J(a,\Delta^*(a))$ and $J(a,\Delta) \leq c$ for all $\Delta \in [M,\infty]$. On the other hand, since $J(a,\Delta)$ is continuous (with respect to Δ) and $J(a,M) \leq c < J(a,\Delta^*(a))$, there is $m \in (\Delta^*(a), M]$ such that $J(a,m) = c$. There is also $m^* \in [m,M]$ such that

$$J(a,m^*) = \sup_{\Delta \in [m,M]} J(a,\Delta).$$

Clearly, for all $\Delta \in [m^*, \infty]$,

$$J(a,\Delta) \leq J(a,m^*). \tag{12}$$

Let $\Delta \in (0,\infty]$. Now, by (12),

$$\begin{aligned} J(a,m^*+\Delta) &= pJ(a,m^*) + (1-p)J(a+m^*,\Delta) \\ &\overset{(12)}{\geq} pJ(a,m^*+\Delta) + (1-p)J(a+m^*,\Delta), \end{aligned}$$

where $p \in (0,1)$ refers to

$$p = \frac{\int_a^{a+m^*} \overline{F}(t)\,\mathrm{d}t}{\int_a^{a+m^*+\Delta} \overline{F}(t)\,\mathrm{d}t}.$$

Thus, we have, for any $\Delta \in (0,\infty]$,

$$J(a+m^*,\Delta) \leq J(a,m^*+\Delta) \leq J(a,m^*). \tag{13}$$

By continuity, we also have

$$J(a+m^*,0) = \lim_{\Delta \to 0^+} J(a+m^*,\Delta) \leq J(a,m^*). \tag{14}$$

From (13) and (14) it follows that

$$G(a+m^*) = \sup_{\Delta \geq 0} J(a+m^*,\Delta) \leq J(a,m^*) < J(a,\Delta^*(a)) = G(a),$$

which contradicts our assumption that $G(x) \geq G(a)$ for all $x \geq a$. Thus, necessarily $\Delta^*(a) = \infty$, which is equivalent to $G(a) = H(a)$. □

4 Connections between the efficiency function and the distribution classes

In this section we characterize the service time distribution classes mentioned in Section 2 by means of the efficiency function $J(a,\Delta)$. We start with the classes DHR, IMRL, and NWUE. To the best of our knowledge characterizations given in this section are novel.

Proposition 1 *Let $a < b$. If $h(x)$ is decreasing for all $x \in [a,b]$, then $J(a,\Delta)$ is decreasing (with respect to Δ) for all $\Delta \in [0, b-a]$.*

Proof: Let $\Delta \in (0, b-a]$. Now $h(a+t) \geq h(a+\Delta)$ for all $0 \leq t \leq \Delta$. Thus,

$$J(a,\Delta) = \frac{\int_0^\Delta h(a+t)\overline{F}(a+t)\,\mathrm{d}t}{\int_0^\Delta \overline{F}(a+t)\,\mathrm{d}t} \geq h(a+\Delta),$$

which implies, by (3), that $J(a,\Delta)$ is decreasing with respect to Δ for any $\Delta > 0$. For $\Delta = 0$, the monotonicity property follows from the continuity of $J(a,\Delta)$ with respect to Δ. □

With $a = 0$ and $b \to \infty$, we get the following corollary.

Corollary 4 *If the service time distribution belongs to DHR, then $J(a,\Delta)$ is decreasing (with respect to Δ) for all a, Δ.*

Proposition 2 *Let $a < b$. Now $1/H(b) \geq 1/H(a)$ if and only if $J(a, b-a) \geq J(a,\infty)$.*

Proof: Let $\Delta = b - a$. Now

$$\begin{aligned}
J(a,\Delta) \geq J(a,\infty) \quad &\iff \quad \frac{\overline{F}(a) - \overline{F}(a+\Delta)}{\int_0^\Delta \overline{F}(a+t)\,\mathrm{d}t} \geq \frac{\overline{F}(a)}{\int_0^\infty \overline{F}(a+t)\,\mathrm{d}t} \\
&\iff \quad \left(\overline{F}(a) - \overline{F}(a+\Delta)\right) \int_0^\infty \overline{F}(a+t)\,\mathrm{d}t \geq \overline{F}(a) \int_0^\Delta \overline{F}(a+t)\,\mathrm{d}t \\
&\iff \quad \overline{F}(a) \int_0^\infty \overline{F}(a+\Delta+t)\,\mathrm{d}t \geq \overline{F}(a+\Delta) \int_0^\infty \overline{F}(a+t)\,\mathrm{d}t \\
&\iff \quad \frac{\int_0^\infty \overline{F}(a+\Delta+t)\,\mathrm{d}t}{\overline{F}(a+\Delta)} \geq \frac{\int_0^\infty \overline{F}(a+t)\,\mathrm{d}t}{\overline{F}(a)} \\
&\iff \quad \frac{1}{H(a+\Delta)} \geq \frac{1}{H(a)},
\end{aligned}$$

which completes the proof. □

We have immediately the following two corollaries.

Corollary 5 *The service time distribution belongs to IMRL if and only if $J(a,\Delta) \geq J(a,\infty)$ for all a, Δ.*

Corollary 6 *The service time distribution belongs to NWUE if and only if $J(0,\Delta) \geq J(0,\infty)$ for all Δ.*

Using similar proofs as above (with reversed inequalities), we get the following additional two results.

Proposition 3 *Let $a < b$. If $h(x)$ is increasing for $x \in [a,b]$, then $J(a,\Delta)$ is increasing (with respect to Δ) for all $\Delta \in [0, b-a]$.*

Proposition 4 *Let $a < b$. Now $1/H(b) \leq 1/H(a)$ if and only if $J(a, b-a) \leq J(a, \infty)$.*

Thus the following characterizations for IHR, DMRL and NBUE can be derived:

Corollary 7 *If the service time distribution belongs to IHR, then $J(a, \Delta)$ is increasing (with respect to Δ) for all a, Δ.*

Corollary 8 *The service time distribution belongs to DMRL if and only if $J(a, \Delta) \leq J(a, \infty)$ for all a, Δ.*

Corollary 9 *The service time distribution belongs to NBUE if and only if $J(0, \Delta) \leq J(0, \infty)$ for all Δ.*

5 Connections between the Gittins index and the distribution classes

In this section we characterize the service time distribution classes by means of the Gittins index $G(a)$. These results prove to be very useful in the following sections when considering the optimality of various scheduling disciplines.

Proposition 5 *The following two statements are equivalent:*

(i) The service time distribution belongs to DHR (i.e., $h(a)$ is decreasing for all a).

(ii) $G(a)$ is decreasing for all a.

In this case, $G(a) = h(a)$ for all a.

Proof: 1° Assume that $h(a)$ is decreasing for all a. Then, by Corollary 4, $G(a) = J(a, 0) = h(a)$ for all a. Thus, $G(a)$ is decreasing for all a.

2° Now assume that $G(a)$ is decreasing for all a. Then, by Lemma 7, $G(a) = h(a)$ for all a. Thus, $h(a)$ is decreasing for all a. □

Proposition 6 *The following three statements are equivalent:*

(i) The service time distribution belongs to DMRL (i.e., $H(a)$ is increasing for all a).

(ii) $G(a)$ is increasing for all a.

(iii) $G(a) = H(a)$ for all a.

Proof: Note first that (i) and (iii) are equivalent by Corollary 8. Thus it remains to prove the equivalence of (ii) and (iii).

1° Assume that $G(a) = H(a)$ for all a. Then, by Corollary 8, $H(a)$, and thus also $G(a)$, is increasing for all a.

2° Now assume that $G(a)$ is increasing for all a. Then, by Lemma 8, $G(a) = H(a)$ for all a. □

Proposition 7 *The following three statements are equivalent:*

(i) The service time distribution belongs to NBUE (i.e., $H(a) \geq H(0)$ for all a).

(ii) $G(a) \geq G(0)$ for all a.

(iii) $G(0) = H(0)$.

Proof: Note first that (i) and (iii) are equivalent by Corollary 9. Thus it remains to prove the equivalence of (ii) and (iii).

1° Assume that $G(0) = H(0)$. Then, by Corollary 9, $H(a) \geq H(0)$ for all a. Thus, $G(a) \geq H(a) \geq H(0) = G(0)$ for all a.

2° Now assume that $G(a) \geq G(0)$ for all a. Then, by Lemma 8, $G(0) = H(0)$. □

6 The optimal policy for DHR and NBUE service times

We present now the main results of the present paper. As we have seen in Theorem 1 the Gittins policy minimizes the mean delay among all non-anticipating policies. We use the optimality of the Gittins policy to characterize *completely* under which assumptions the policies FB and FCFS are optimal.

Theorem 3 *FB minimizes stochastically the number of jobs in the system if and only if the service time distribution belongs to DHR.*

Proof: 1° If the service time distribution is of type DHR, then from [13] we know that FB minimizes stochastically the number of jobs in the system.

2° Assume now that FB minimizes stochastically the number of jobs in the system. Then, in particular, FB minimizes the mean number of jobs in the system. Thus, by Theorem 1, we conclude that the Gittins index $G(a)$ is decreasing for all a, which implies, by Proposition 5, that the service time distribution is of type DHR. □

The optimality of FB when the service time distribution is of type DHR is not a new result, see [25, Section 6], [13, Corollary 2.1.2], [26, Theorem 4.8a]. It is also known that FB is not necessarily optimal for IMRL service times, see [1]. But to the best of our knowledge, that DHR is a sufficient and necessary condition for the stochastic optimality of FB is a new result.

Theorem 4 *FCFS minimizes the mean number of jobs in the system if and only if the service time distribution belongs to NBUE.*

Proof: 1° Assume that the service time distribution belongs to NBUE. Then, by Proposition 7, $G(a) \geq G(0)$ for all $a \geq 0$. Thus, any non-preemptive discipline (like FCFS) is a Gittins discipline in this case. The claim follows now from Theorem 1.

2° Assume now that the optimal discipline is FCFS. Thus, by Theorem 1, we conclude that $G(a) \geq G(0)$ for all a, which implies, by Proposition 7, that the service time distribution is of type NBUE. □

The optimality of FCFS in this case was already known, see [15]. However, to the best of our knowledge, the "if and only if" characterization stated in Theorem 4 had not been obtained before. Yashkov [26, Theorem 4.8b] uses the Gittins index to state the optimality of FCFS only within the class IHR, which is a strict subclass of NBUE.

Note that when FB is optimal, we have shown that the distribution must belong to DHR, and, using the argument of [13], it can be shown that FB minimizes the number in system in a stochastic sample path sense, i.e., all departures are jointly stochastically earlier. This is stronger than stochastically minimizing the number in system at any time. Righter and Shanthikumar also showed that for IHR distributions, FCFS stochastically minimizes the number in system at any time. However, for IHR distributions, FCFS does not minimize the number in the stochastic sample path sense. For this, ILR (Increasing in Likelihood Ratio) distributions are required [14]. Also, for the the DHR, IHR, ILR and NBUE results, the arrival process can be arbitrary.

7 Multi-class queues

We now permit multi-class queues. Let K denote the number of classes. Jobs of class i, $i = 1, \ldots, K$, have distribution F_i ($\overline{F}_i$, f_i, h_i, H_i), satisfying the conditions given in Section 2, and arriving jobs may have already received some service elsewhere. Class i jobs arrive to the system according to a Poisson process of rate λ_i. Let A_i be the distribution of ages upon arrival for class i. We assume, for stability, that

$$\sum_i \lambda_i \int_0^\infty 1/H_i(a) \, \mathrm{d}A_i(a) < 1.$$

Let $\mathbb{E}[N_i]$ be the long-run mean number of jobs of class i in the system, and let $\mathbb{E}[T_i]$ be the long-run mean sojourn time for class i jobs. From Theorem 1 we know that the Gittins index policy minimizes the number of jobs and the mean sojourn time in this queue.

We first consider DHR service time distributions. By HHR (Highest-Hazard-Rate) we mean the service discipline that always serves the job with the current largest hazard rate, $h_i(a)$. When there is a single class and the service times are of type DHR, the HHR discipline clearly coincides with FB.

Theorem 5 *If all service time distributions belong to DHR, then HHR stochastically minimizes the number of jobs in the system at any time.*

The proof, which is similar to that one of Theorem 3 (utilizing Proposition 5), is omitted.

Corollary 10 *If all service time distributions belong to DHR, and if the hazard rates for different classes are non-overlapping, i.e., the classes can be ordered such that for all $i < j$, $h_i(x) \geq h_j(y)$ for all x, y, then class i should have strict priority over class j whenever $i < j$, and within a class FB should be applied, i.e., the job with the least attained service (from either that server or elsewhere) should have priority.*

Theorem 6 slightly generalizes Theorem 3 to allow arrivals that have already received some service.

Theorem 6 *If there is a single class, then FB stochastically minimizes the number of jobs in the system at any time if and only if the service time distribution belongs to DHR.*

Proof: The result follows from Theorems 5 and 1 together with Proposition 5. □

Corollary 11 *If there is a single class, and if FB minimizes the mean number of jobs in the system at any time then it also stochastically minimizes the number of jobs in the system at any time.*

Now we consider DMRL service time distributions. By SERPT (Shortest-Expected-Remaining-Processing-Time-first) we mean the service discipline that always serves the job with the current shortest expected remaining processing time, $1/H_i(a)$. When there is a single class and the service times are of type DMRL, the SERPT discipline clearly coincides with MAS (Most-Attained-Service).

Theorem 7 *If all service time distributions belong to DMRL, then SERPT minimizes the mean number of jobs in the system.*

Proof: The result follows from Theorem 1 and Proposition 6. □

Note that when service time distributions are of type DMRL, under SERPT jobs will be preempted only by higher priority arrivals. That is, the relative priorities of jobs never change so that, within a class, MAS is optimal. Also note that under the conditions of Corollary 10, HHR coincides with SERPT from Lemma 2, so SERPT is optimal in this case too, but now the relative priorities of different jobs in the same class change as jobs are served.

Theorem 8 *If SERPT minimizes the mean number of jobs in the system, and relative priorities of jobs within a class do not change, then all service time distributions must be of type DMRL.*

Proof: The result follows from Theorem 1 and Proposition 6. □

Corollary 12 *If there is a single class, then MAS minimizes the mean number of jobs in the system if and only if the service time distribution belongs to DMRL.*

When there is a single class and jobs receive no service elsewhere, then we do not need to order two jobs with two different non-zero attained services, so we can weaken the assumption on the distribution to NBUE, as already seen in our previous Theorem 4.

Under the conditions of Corollary 10 or Theorem 7, the job classes can be ordered so that class i should have strict priority over class j whenever $i < j$, so we can strengthen our results to include agreeable weights. That is, suppose class i has weight or holding cost c_i such that $c_i \geq c_j$ whenever $i < j$. Note that this means, under the conditions of Theorem 9, that $c_i \geq c_j$ whenever $H_i(x) \geq H_j(y)$ for any x, y.

Theorem 9 *If service times are of type DHR with non-overlapping hazard rates, or if they are of type DMRL, and if the weights c_i are agreeable with SERPT, then SERPT minimizes $\sum_i c_i \mathbb{E}[N_i]$ and $\sum_i c_i \mathbb{E}[T_i]$.*

Proof: Under our service time distribution conditions, we can order the job classes so that SERPT gives strict preemptive priority to class i over class j when $i < j$. That is, job classes with higher priority (lower index) are unaffected by lower priority classes. Therefore, from Corollary 10 and Theorem 7 (and Little's law), we know that SERPT minimizes $\sum_{k=1}^{i} \mathbb{E}[N_k]$ and $\sum_{k=1}^{i} \mathbb{E}[T_k]$ for all $1 \leq i \leq K$.

Without loss of generality we let $c_{K+1} = 0$. By assumption $c_i - c_{i+1} \geq 0$ for all $1 \leq i \leq K$, then we have that SERPT minimizes

$$\sum_{k=1}^{K} c_k \mathbb{E}[N_k] = \sum_{k=1}^{K} \mathbb{E}[N_k] \sum_{i=1}^{K} (c_i - c_{i+1}) = \sum_{i=1}^{K} (c_i - c_{i+1}) \sum_{k=1}^{i} \mathbb{E}[N_k]$$

and similarly

$$\sum_{k=1}^{K} c_k \mathbb{E}[T_k] = \sum_{i=1}^{K} (c_i - c_{i+1}) \sum_{k=1}^{i} \mathbb{E}[T_k].$$

□

8 Non-monotonic hazard rate and mean residual lifetime

In this section we consider service time distributions where $h(x)$ and $H(x)$ are non-monotonic. Suppose (for a while) there is a single class of jobs, and all jobs arrive with 0 attained service. Note that as a general principle, from Proposition 3 we know that if it is optimal to work on a job whose hazard rate is increasing, then we should continue to work on it as long as its hazard rate continues to increase, or until a new job arrives. We now consider some special cases.

We denote as FCFS+FB(θ^*) the policy that serves non-preemptively the jobs that have received no service until their attained service reaches θ^*. Jobs with attained service greater than θ^* are served according to FB, and jobs with no attained service have priority over those with attained service of at least θ^*. We note that this discipline belongs to the class of MLPS (Multilevel-Processor-Sharing) disciplines [6, Sect. 4.7].

We have the following result:

Theorem 10 *Suppose there is a single class of jobs with 0 attained service upon arrival. If $h(x)$ increases for $x \in [0, k)$ and decreases for $x \in [k, \infty)$, then FCFS+FB(θ^*) is optimal, where $\theta^* = \Delta^*(0) \geq k$.*

Proof: We know from Proposition 3, that $J(x, \Delta)$ is increasing with respect to Δ, for $\Delta \in [0, k-x]$ and $x \in [0, k]$. Hence, $\theta^* := \Delta^*(0) \geq k$. Also, from Lemma 3, for $x \leq \theta^*$, $G(x) \geq G(0)$. For $x \geq k$, $G(x) = h(x)$ from Proposition 1, and among the jobs with attained service at least k, HHR = FB is optimal. Also, if $\theta^* < \infty$, $G(0) \geq h(\theta^*)$ from Lemma 6, so $G(0) \geq G(x)$ for $x \geq \theta^*$. Therefore, jobs that have received no service should be served non-preemptively (e.g., according to FCFS) until their attained service reaches θ^*, and then they should be processed according to FB, and jobs with no attained service should have priority over those with attained service of at least θ^*. This is precisely the policy that we have denoted as FCFS+FB(θ^*). □

In fact, FCFS+FB(θ^*) can be shown to be optimal in a more general class of service time distributions, NBUE+DHR, where $1/H(x) \leq 1/H(0)$ for $x \in [0, k)$ and $h(x)$ is decreasing for $x \in [k, \infty)$, see [2].

Theorem 10 easily extends to multiple classes with non-overlapping hazard rates, where there is strict priority among the classes, and within a class the policy is as above.

In the following result we suppose that we have multiple classes and the service distributions are such that $h_i(x)$ decreases for $x \in [0, k_i)$, and $H_i(x)$ increases for $x \in [k_i, \infty)$, and we permit arrivals of jobs with partially completed service.

Theorem 11 *If $h_i(x)$ decreases for $x \in [0, k_i)$ and if $H_i(x)$ increases for $x \in [k_i, \infty)$, then $G_i(x) = h_i(x)$ for $x < \theta_i^*$, and $G_i(x) = H_i(x)$ for $x \geq \theta_i^*$, where $\theta_i^* = \inf\{x : h_i(x) \leq H_i(x)\} \leq k_i$. The job with the highest value of $G_i(x)$ should be served. Within class i, jobs with attained service at least θ_i^* should be served according to MAS (so they will only be preempted by new arrivals), and jobs with attained service less than θ_i^* should be served according to FB.*

Proof: From Proposition 4, for $x \in [k_i, \infty)$, $G_i(x) = H_i(x)$, $\Delta^*(x) = \infty$, and jobs of class i with attained service at least k_i should be served according to MAS. From Proposition 1, for $x \in [0, k_i)$, $J_i(x, \Delta)$ is decreasing in Δ for $\Delta \leq k_i - x$, so $\Delta_i^*(x) \notin (0, k_i - x)$. Also, from Corollary 3, for $x \in [0, k_i)$, if $\Delta_i^*(x) \geq k_i - x$, then $\Delta_i^*(x) = \infty$. We therefore have, for $x \in [0, k_i)$, that either $\Delta_i^*(x) = 0$ or ∞, and $G_i(x) = \max\{h_i(x), H_i(x)\}$. Let $\theta_i^* = \inf\{x : h_i(x) \leq H_i(x)\}$, so $\theta_i^* \leq k_i$, $\Delta_i^*(\theta_i^*) = \infty$, $G_i(x) = h_i(x)$ for $x < \theta_i^*$, and $G_i(x) = H_i(x)$ for $x \geq \theta_i^*$. □

Suppose now that all jobs are from a single class and arrive with 0 attained service. We define the policy FB+FCFS(θ^*)[2] as the one that serves the jobs with attained service less than θ^* according to FB, and the jobs with attained service greater than θ^* according to FCFS. Note that there will be at most one job with attained service (strictly) greater than θ^*. If there are jobs with attained service less than θ^* and a job with attained service greater than θ^*, then which group is given priority depends on which is greater of the highest hazard rate of the former jobs and the reciprocal of the mean remaining service time of the latter job.

A minor modification of the proof of Theorem 11 gives us the following.

Corollary 13 *Suppose there is a single class of jobs with 0 attained service upon arrival. Also suppose $h(x)$ decreases for $x \in [0, k)$, and let $\theta^* = \min\{\inf\{x : h(x) \leq H(x)\}, k\}$. If $H(\theta^*) \leq H(x)$ for $x \in [\theta^*, \infty)$ then FB+FCFS(θ^*) is optimal.*

Example: *Bounded Pareto*

Let us consider a specific distribution with a bathtub hazard rate and we illustrate the result of Corollary 13. We refer to [9, Chapter 3] for various examples of distributions with a bathtub hazard rate. Here we will

[2] This is not an MLPS discipline since it does not give an absolute priority to jobs with attained service less than θ^*.

consider a bounded Pareto distribution. The bounded Pareto distribution is defined by

$$\overline{F}(x) = 1 - \frac{1 - \frac{1}{(1+x)^\alpha}}{1 - \frac{1}{(1+p)^\alpha}}, \ 0 \le x \le p,$$

and the hazard rate is

$$h(x) = \frac{\alpha}{(1+x)\left(1 - \left(\frac{1+x}{1+p}\right)^\alpha\right)}, \ 0 \le x \le p,$$

We set the parameters $\alpha = 2$ and $p = 20$. In Figure 1 we depict the functions $h(x)$ and $H(x)$. The hazard rate function is decreasing in $[0, 11.1)$ and increasing in $(11.1, 20)$.

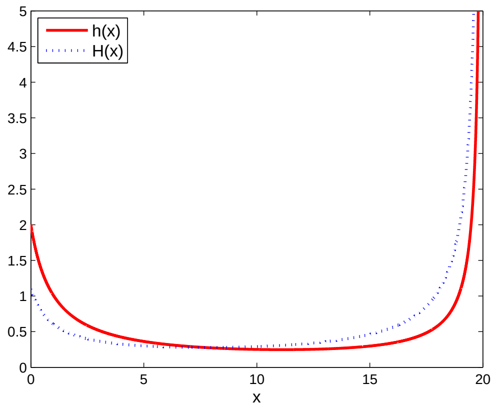

Figure 1: $h(x)$ and $H(x)$ for a bounded Pareto distribution

The value of $\theta^* \approx 7.7$. From Lemma 1 we know that this point satisfies $h(\theta^*) = H(\theta^*)$, and in addition $H(x)$ is decreasing in $[0, \theta^*)$ and increasing in $[\theta^*, p)$. Thus the assumption of Corollary 13 is satisfied since $H(\theta^*) \le H(x)$ for $x \in [\theta^*, p)$. In Figures 2 and 3 we depict the values of $\Delta^*(a)$ and $G(a)$, respectively. We can observe that the optimum quantum becomes strictly positive for the first time exactly at θ^*. The curve $G(a)$ is decreasing in $[0, \theta^*)$ and increasing in $[\theta^*, p]$, and as a consequence, the optimal policy is precisely FB+FCFS(θ^*) as stated in Corollary 13.

9 Conclusions

Gittins [5] considered an $M/G/1$ queue and proved that the so-called Gittins index policy minimizes the mean delay. The Gittins index policy determines, depending on the service times of jobs, which job should be served next. Gittins derived this result as a by-product of his groundbreaking results on the multi-armed bandit problem. Gittins' results on the multi-armed bandit problem have had a profound impact and it is

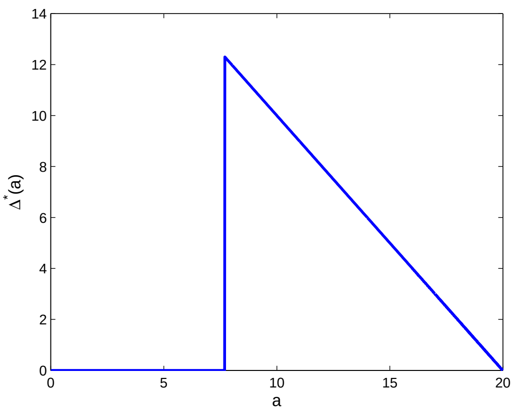

Figure 2: Optimal quantum $\Delta^*(a)$

extremely highly cited. However, and despite the large body of literature on scheduling disciplines in single server queues, Gittins' work in the $M/G/1$ context has not received much attention.

In the current paper we have used the Gittins index policy to obtain several new characterizations of the optimal scheduling discipline in an $M/G/1$ queue. For example we have shown that the policy FB is optimal if and only if the service time distribution is of type DHR (previously only the reverse implication was known). We have also extended Gittins' result to the multi-class case, which has allowed us to characterize the optimal policy for several particular cases of service time distributions.

In future research we plan to further investigate the structure of the optimal policy when the Gittins index of the service time distribution is not monotonic. The conditions under which the Gittins index policy gives significantly better performance than the FB and FCFS policies should also be addressed. Finally the applicability of our results in real systems like the Internet should also be more carefully evaluated.

References

[1] S. Aalto and U. Ayesta. On the nonoptimality of the foreground-background discipline for IMRL service times. *Journal of Applied Probability*, 43:523–534, 2006.

[2] S. Aalto and U. Ayesta. Optimal scheduling of jobs with a DHR tail in the M/G/1 queue. In *Proceedings of ValueTools 2008*, Athens, Greece, 2008.

[3] D. Bertsimas and J. Nino-Mora. Restless bandits, linear programming relaxations, and a primal-dual index heuristic. *Operations Research*, 48:80–90, 2000.

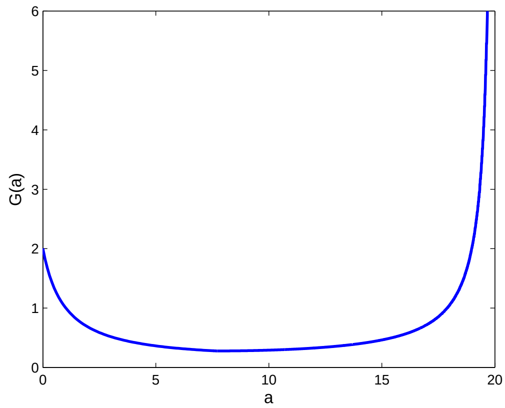

Figure 3: Gittins index $G(a)$

[4] E. Gelenbe and I. Mitrani. *Analysis and Synthesis of Computer Systems.* London: Academic Press, 1980.

[5] J.C. Gittins. *Multi-armed Bandit Allocation Indices.* Wiley, Chichester, 1989.

[6] L. Kleinrock. *Queueing Systems, vol. 2.* John Wiley and Sons, 1976.

[7] G.P. Klimov. Time-sharing service systems. I. *Theory of Probability and Its Applications*, 19:532–551, 1974.

[8] G.P. Klimov. Time-sharing service systems. II. *Theory of Probability and Its Applications*, 23:314–321, 1978.

[9] C.-D. Lai and M. Xie. *Stochastic Ageing and Dependence for Reliability.* Springer New York, 2006.

[10] I. Meilijson and G. Weiss. Multiple feedback at a single server station. *Stochastic Processes and Applications*, 5:195–205, 1977.

[11] I. Meilijson and U. Yechiali. On optimal right-of-way policies at a single server station when insertion of idle times is permitted. *Stochastic Processes and Applications*, 6:25–32, 1977.

[12] N. Osipova, U. Ayesta, and K.E. Avrachenkov. Optimal policy for multi-class scheduling in a single server queue. In *Proceedings of ITC-21*, 2009.

[13] R. Righter and J.G. Shanthikumar. Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures. *Probability in the Engineering and Informational Sciences*, 3:323–334, 1989.

[14] R. Righter and J.G. Shanthikumar. Extremal properties of the FIFO discipline in queueing networks. *Journal of Applied Probability*, 29:967–978, 1992.

[15] R. Righter, J.G. Shanthikumar, and G. Yamazaki. On extremal service disciplines in single-stage queueing systems. *Journal of Applied Probability*, 27:409–416, 1990.

[16] L.E. Schrage. A proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 16(3):687–690, 1968.

[17] K. Sevcik. Scheduling for minimum total loss using service time distributions. *Journal of the ACM*, 21:66–75, 1974.

[18] M. Shaked and J.G. Shanthikumar. *Stochastic orders*. Springer, 2007.

[19] D.R. Smith. A new proof of the optimality of the shortest remaining processing time discipline. *Operations Research*, 26:197–199, 1978.

[20] J.N. Tsitsiklis. A short proof of the gittins index theorem. In *IEEE CDC*, pages 389–390, 1993.

[21] P. Varaiya, J. Walrand, and C. Buyukkoc. Extensions of the multiarmed bandit problem: the discounted case. *IEEE Transactions on Automatic Control*, 30:426–439, 1985.

[22] R. Weber. On the gittins index for multiarmed bandits. *Annals of Applied Probability*, 2(4):1024–1033, 1992.

[23] P. Whittle. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability*, 25:287–298, 1988.

[24] S.F. Yashkov. On feedback sharing a processor among jobs with minimal serviced length (in russian). *Technika Sredstv Svyazi. Ser. ASU*, 2:51–62, 1978.

[25] S.F. Yashkov. Processor sharing queues: Some progress in analysis. *Queueing Systems*, 2:1–17, 1987.

[26] S.F. Yashkov. Mathematical problems in the theory of processor sharing queueing systems. *Journal of Soviet Mathematics*, 58:101–147, 1992.