Optimal scheduling of jobs with a DHR tail in the M/G/1 queue

Samuli Aalto
TKK Helsinki University of Technology
Department of Communications and Networking
P.O.Box 3000
02015 TKK
Finland
samuli.aalto@tkk.fi

Urtzi Ayesta
LAAS-CNRS
Université de Toulouse
7 Avenue Colonel Roche
31077 Toulouse Cedex
France
urtzi@laas.fr

ABSTRACT

We consider the mean delay optimization in the M/G/1 queue for jobs with a service time distribution that has a tail with decreasing hazard rate (DHR). If the DHR property is valid for the whole distribution, then it is known that the Foreground-Background (FB) discipline, which gives priority to the job with least amount of attained service, is optimal among nonanticipating scheduling disciplines. However, FB may fail to be optimal if the DHR property is valid only for the tail of the distribution. An important example is the Pareto distribution bounded away from zero. In this paper we show that for a class of service time distributions with a DHR tail (including the Pareto distribution), the optimal nonanticipating discipline is a combination of FCFS and FB disciplines, which gives priority to the jobs with attained service less than some fixed threshold θ^*. These priority jobs are served in the FCFS manner. If there are no jobs with attained service less than θ^*, priority is given to the job with least amount of attained service.

Keywords

M/G/1, scheduling, mean delay, Pareto distribution, Gittins index

1. INTRODUCTION

We consider the optimal scheduling problem in the M/G/1 queue with the objective to minimize the mean delay (*i.e.*, sojourn time). We assume that jobs are served according to a preemptive, work conserving and nonanticipating scheduling discipline. A discipline is work conserving if it does not idle when there are jobs waiting, and nonanticipating if the remaining service times of jobs are unknown for the scheduler. Nonanticipating disciplines may utilize the attained service (age) information, but the remaining service times are unknown to such a scheduler. Thus, for example, the Shortest-Remaining-Processing-Time (SRPT) discipline does not belong to these disciplines, while the Foreground-Background (FB) discipline, which gives priority to the job with least amount of attained service [12], or any other age-based discipline is nonanticipating. Our motivation comes from the recent literature that deals with the performance of age-based scheduling on the Internet, see [3, 4, 6, 8, 10, 11, 13, 14, 15].

It is known that for the service times distributions belonging to the Decreasing-Hazard-Rate (DHR) class, FB is optimal [20, 16], whereas the ordinary First-Come-First-Served (FCFS) discipline minimizes the mean delay for the service time distributions belonging to the New-Better-than-Used-in-Expectation (NBUE) class [17]. Roughly said, DHR distributions have a high variation, whereas NBUE distributions are much less variable. A high variation is typical for flow sizes in the Internet, which have been modelled by distributions like hyperexponential or Pareto [5, 7]. While the hyperexponential distribution belongs to the DHR class, the Pareto service time distribution, defined by

$$P\{S > x\} = \left(\frac{k}{x}\right)^{\alpha}, \quad x \geq k > 0, \tag{1}$$

does not. The hazard rate for this distribution is first constant (zero), then jumps up to a positive value at $x = k$, and finally decreases from the argument value k on.

As noted in [19], FB may not be optimal for the Pareto service time distribution (1). The intuition behind this is as follows. Suppose that there is a job in the system that has attained service close but strictly less than k. When a new job arrives, the FB discipline would immediately start serving the new job. However, since the minimum service time is k, it might be better to keep on serving the job that is already in the system, and which might leave the system soon after getting k units of service.

In [1] we proved that FB is not optimal for a modification of the Pareto distribution. In [2] we made a similar observation (however, without any proofs due to the strict page limit) for a class of service time distributions for which the hazard rate is first constant and then decreasing.

In the present paper, we consider an even more general class of service time distributions, for which the head of the distribution has the NBUE property, while the tail satisfies the DHR requirement. This class includes both the Pareto distribution (1) and the modification considered in [1]. By applying the so called Gittins index approach [9], we prove that for this class of service time distributions the optimal discipline is a combination of the FCFS and FB disciplines.

ValueTools 2008, October 21-23, 2008, Athens, GREECE.

More precisely, it is an age-based discipline which gives priority to the jobs with attained service less than some fixed threshold θ^*. These priority jobs are served in the FCFS manner. If there are no jobs with attained service less than θ^*, the job with least amount of attained service will be served. We use notation FCFS + FB(θ^*) for this discipline.

The rest of the paper is organized as follows. In Section 2 we define the queueing model and introduce the relevant service time distribution classes. The Gittins index is introduced and some related results are presented in Section 3. The main results concerning the mean delay optimization are presented in Section 4. In Section 5 we give some numerical results to cast light on the optimal threshold θ^*, as well as the maximum gain achieved by the optimal discipline in the case of the Pareto distribution (1). In Section 6 we investigate, via a numerical example, what is the structure of the optimal scheduling discipline when the service time distribution is an upper bounded Pareto with a finite support. Section 7 concludes the paper.

2. PRELIMINARIES

Consider an M/G/1 queue with arrival rate λ, mean service time $E[S]$, and load $\rho = \lambda E[S] < 1$. Jobs are served according to a work conserving and nonanticipating scheduling discipline π. Let Π denote the family of such disciplines.

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

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

It is also right-continuous with left-limits. 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.

In addition, define, for all x,

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

We note that

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

A service time distribution belongs to the class IMRL (Increasing Mean Residual Lifetime) if $1/H(x)$ is increasing for all x, *i.e.*, $1/H(x) \le 1/H(y)$ whenever $x \le y$. We note that IMRL is a weaker condition than DHR so that DHR $\subset$ IMRL. The class DMRL (Decreasing Mean Residual Lifetime) is defined correspondingly.

A service time distribution belongs to the class NWUE (New Worse than Used in Expectation) if $1/H(x) \ge 1/H(0)$ for all x. By definition, it is clear that NWUE is a weaker condition than IMRL so that IMRL $\subset$ NWUE. The class NBUE (New Better than Used in Expectation) 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.

Finally we introduce a new class of service time distributions, called NBUE + DHR(k), with a threshold parameter $k > 0$. A service time distribution belongs to the class NBUE + DHR(k) if

- *Condition A1:* $1/H(x) \le 1/H(0)$ for all $x < k$, and
- *Condition A2:* $h(x)$ is decreasing for all $x > k$.

Examples related to this distribution class are given below.

In particular, we note that a sufficient (but not necessary) condition for A1 is the following one:

- *Condition B1:* $h(x)$ is increasing for all $x < k$.

An even more stringent condition for A1 is as follows:

- *Condition C1:* $h(x)$ is constant for all $x < k$.

Thus the class NBUE + DHR(k) includes all the distributions considered in [2].

EXAMPLE 1. (Pareto) Let $\alpha > 1$ and define the tail distribution function $\overline{F}(x)$ as in (1) so that

$$\overline{F}(x) = \begin{cases} 1, & 0 \le x < k, \\ \left(\frac{k}{x}\right)^\alpha, & x \ge k. \end{cases}$$

Now the functions $h(x)$ and $H(x)$ are as follows:

$$h(x) = \begin{cases} 0, & 0 \le x < k, \\ \frac{\alpha}{x}, & x \ge k. \end{cases}$$

$$H(x) = \begin{cases} \frac{\alpha - 1}{k\alpha - x(\alpha - 1)}, & 0 \le x < k, \\ \frac{\alpha - 1}{x}, & x \ge k; \end{cases}$$

Here the hazard rate is constant zero for all $x < k$ satisfying Condition C1. The functions $h(x)$ and $H(x)$ are plotted in Figure 1 for parameter values $k = 1$, $\alpha = 2$.

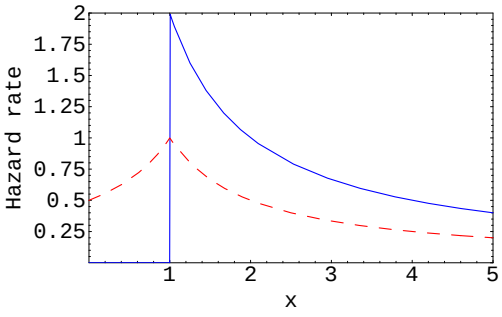

Figure 1: Functions $h(x)$ (solid curve) and $H(x)$ (dashed curve) for the Pareto distribution with parameters $k = 1$, $\alpha = 2$.

EXAMPLE 2. (Exponential+Pareto) Let $\alpha > 1$, $\mu > 0$, and define $\overline{F}(x)$ as follows:

$$\overline{F}(x) = \begin{cases} e^{-\mu x}, & 0 \le x < k, \\ e^{-\mu k}\left(\frac{k}{x}\right)^\alpha, & x \ge k. \end{cases}$$

Now $h(x)$ is as follows:

$$h(x) = \begin{cases} \mu, & 0 \le x < k, \\ \dfrac{\alpha}{x}, & x \ge k. \end{cases}$$

Here the hazard rate is a positive constant for all $x < k$ satisfying Condition C1. The Pareto distribution (1) presented in Example 1 can be considered as the limit of this distribution with $\mu \to 0$. In addition, the modified Pareto distribution constructed in [1] is a special case of this distribution with choices $\mu = \ln k$ and $\alpha = k$.

EXAMPLE 3. (Uniform+Pareto) Let $\alpha > 1$, $0 < p < 1$, and define $\overline{F}(x)$ as follows:

$$\overline{F}(x) = \begin{cases} 1 - \dfrac{p}{k}x, & 0 \le x < k, \\ (1-p)\left(\dfrac{k}{x}\right)^{\alpha}, & x \ge k. \end{cases}$$

Now $h(x)$ is as follows:

$$h(x) = \begin{cases} \dfrac{\frac{p}{k}}{1 - \frac{p}{k}x}, & 0 \le x < k, \\ \dfrac{\alpha}{x}, & x \ge k. \end{cases}$$

Here the hazard rate is (strictly) increasing for all $x < k$ satisfying Condition B1.

3. GITTINS INDEX

In this section we introduce the Gittins index for jobs and give some auxiliary results.

3.1 Function $J(a, \Delta)$

For any $a, \Delta \ge 0$, let

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

For a job that has attained service a and is assigned Δ units of service, Equation (2) can be interpreted as the ratio between (i) the probability that the job will complete with a quota of Δ (interpreted as payoff) and (ii) the expected processor time that a job with attained service a and service quota Δ will require from the server (interpreted as investment).

Note that, for any a,

$$J(a,0) = \frac{f(a)}{\overline{F}(a)} = h(a),$$

$$J(a,\infty) = \frac{\overline{F}(a)}{\int_0^{\infty} \overline{F}(a+t)\,dt} = H(a).$$

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

$$\begin{aligned} \frac{\partial}{\partial\Delta} J(a,\Delta) &= \frac{f(a+\Delta)\int_0^{\Delta}\overline{F}(a+t)\,dt}{(\int_0^{\Delta}\overline{F}(a+t)\,dt)^2} - \\ &\qquad \frac{\overline{F}(a+\Delta)\int_0^{\Delta} f(a+t)\,dt}{(\int_0^{\Delta}\overline{F}(a+t)\,dt)^2} \\ &= \frac{\overline{F}(a+\Delta)(h(a+\Delta) - J(a,\Delta))}{\int_0^{\Delta}\overline{F}(a+t)\,dt}. \quad (3) \end{aligned}$$

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

3.2 Function $G(a)$

For any $a \ge 0$, let

$$G(a) = \sup_{\Delta \ge 0} J(a,\Delta), \quad (4)$$

which is called the *Gittins index* after the author of [9]. In addition, for any $a \ge 0$, let

$$\Delta^*(a) = \sup\{\Delta \ge 0 \mid J(a,\Delta) = G(a)\}. \quad (5)$$

By definition, $G(a) = J(a, \Delta^*(a))$ for all a.

Before studying the specific properties of the Gittins index for the service time distribution class NBUE + DHR(k), we present, in Lemmas 1 and 2, two general results that are needed later on.

LEMMA 1. *$G(x) \ge G(a)$ for all a, x such that $a \le x < a + \Delta^*(a)$.*

PROOF. Suppose that $a \le x < a + \Delta^*(a)$. Since $G(x) = \sup_{\Delta\ge 0} J(x,\Delta)$ and $G(a) = J(a,\Delta^*(a))$, it is sufficient to prove that

$$J(x, a + \Delta^*(a) - x) \ge J(a, \Delta^*(a)). \quad (6)$$

First note that

$$J(a,\Delta^*(a)) = \frac{\int_a^{a+\Delta^*(a)} f(t)\,dt}{\int_a^{a+\Delta^*(a)} \overline{F}(t)\,dt},$$

$$J(x, a+\Delta^*(a)-x) = \frac{\int_x^{a+\Delta^*(a)} f(t)\,dt}{\int_x^{a+\Delta^*(a)} \overline{F}(t)\,dt}.$$

Now,

$$\begin{aligned} &J(a,\Delta^*(a)) \\ &= \frac{\int_a^x f(t)\,dt}{\int_a^{a+\Delta^*(a)}\overline{F}(t)\,dt} + \frac{\int_x^{a+\Delta^*(a)} f(t)\,dt}{\int_a^{a+\Delta^*(a)}\overline{F}(t)\,dt} \\ &= pJ(a, x-a) + (1-p)J(x, a+\Delta^*(a)-x), \end{aligned}$$

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

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

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

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

from which (6) clearly follows. □

LEMMA 2. *$G(a+\Delta^*(a)) \le G(a)$ for all a such that $\Delta^*(a) < \infty$.*

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

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

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

$$G(a+\Delta^*(a)) = J(a+\Delta^*(a), 0) = h(a+\Delta^*(a)). \quad (7)$$

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

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

By (7) and (3), this implies that

$$G(a+\Delta^*(a)) = h(a+\Delta^*(a)) \le J(a,\Delta^*(a)) = G(a).$$

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

$$G(a+\Delta^*(a)) > G(a). \tag{8}$$

This is equivalent to

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

Now,

$$\begin{aligned} &J(a,\Delta^*(a)+\Delta^{**}(a)) \\ &= \frac{\int_a^{a+\Delta^*(a)} f(t)\,dt}{\int_a^{a+\Delta^*(a)+\Delta^{**}(a)} \overline{F}(t)\,dt} + \frac{\int_{a+\Delta^*(a)}^{a+\Delta^*(a)+\Delta^{**}(a)} f(t)\,dt}{\int_a^{a+\Delta^*(a)+\Delta^{**}(a)} \overline{F}(t)\,dt} \\ &= pJ(a,a+\Delta^*(a)) + (1-p)J(a+\Delta^*(a),\Delta^{**}(a)) \\ &= pG(a) + (1-p)G(a+\Delta^*(a)), \end{aligned}$$

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

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

Due to (8), we thus conclude that

$$J(a,\Delta^*(a)+\Delta^{**}(a)) > G(a) = \sup_{\Delta\ge 0} J(a,\Delta),$$

which is a contradiction. Thus the claim must be true also in this case. □

PROPOSITION 1. *Assume that the service time distribution belongs to the class* NBUE + DHR(k). *Then*

(i) $\Delta^*(0) \ge k$,

(ii) $G(a) \ge G(0)$ *for all* $a < \Delta^*(0)$,

(iii) $G(a)$ *is decreasing for all* $a > k$.

Furthermore, if $\Delta^*(0) < \infty$, *then*

(iv) $G(\Delta^*(0)) \le G(0)$.

PROOF. (i) Let $0 < x < k$. By A1 we have

$$J(x,\infty) = H(x) \ge H(0) = J(0,\infty),$$

which is equivalent to

$$\begin{aligned} &\frac{\overline{F}(x)}{\int_x^\infty \overline{F}(y)\,dy} \ge \frac{1}{\int_0^\infty \overline{F}(y)\,dy} \\ \Leftrightarrow\quad &\overline{F}(x)\int_0^\infty \overline{F}(y)\,dy \ge \int_0^\infty \overline{F}(y)\,dy - \int_0^x \overline{F}(y)\,dy \\ \Leftrightarrow\quad &\frac{1}{\int_0^\infty \overline{F}(y)\,dy} \ge \frac{1-\overline{F}(x)}{\int_0^x \overline{F}(y)\,dy}. \end{aligned}$$

Now the last inequality implies that

$$G(0) \ge J(0,\infty) \ge J(0,x).$$

Since this is true for any $0 < x < k$, we conclude that $\Delta^*(0) \ge k$.

(ii) This follows immediately from Lemma 1.

(iii) By A2, $h(x)$ is decreasing for all $x > k$. It follows that, for any $a > k$ and $\Delta > 0$,

$$\begin{aligned} J(a,\Delta) &= \frac{\int_0^\Delta h(a+t)\overline{F}(a+t)\,dt}{\int_0^\Delta \overline{F}(a+t)\,dt} \\ &\ge \frac{h(a+\Delta)\int_0^\Delta \overline{F}(a+t)\,dt}{\int_0^\Delta \overline{F}(a+t)\,dt} = h(a+\Delta). \end{aligned}$$

By (3), we deduce that $J(a,\Delta)$ is a decreasing function of Δ for any $a > k$. Thus, $G(a) = h(a)$, and thus decreasing, for all $a > k$.

(iv) This follows immediately from Lemma 2. □

As an illustration of Proposition 1, we have computed the functions $G(a)$ and $\Delta^*(a)$ for the Pareto distribution (1) with parameters $k = 1$, $\alpha = 2$. The results are plotted in Figure 2. In this case $\Delta^*(0) = 2+\sqrt{3} = 3.732$, and

$$G(0) = G(\Delta^*(0)) = h(\Delta^*(0)) = \frac{2}{2+\sqrt{3}} = 0.536$$

while $h(0) = 0$ and $H(0) = 1/2 = 0.500$. For clarity, we also mention that

$$\begin{cases} \Delta^*(a) > 0, & \text{for } a < 1, \\ \Delta^*(a) = 0, & \text{for } a \ge 1. \end{cases}$$

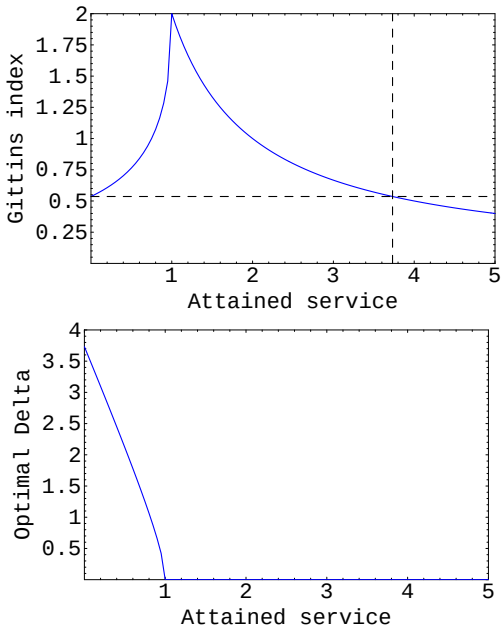

Figure 2: Pareto distribution with parameters $k = 1$, $\alpha = 2$. *Top*: Gittins index $G(a)$ as a function of attained service a with the horizontal dashed line equal to $G(0)$ and the vertical dashed line equal to $a = \Delta^*(0)$. *Bottom*: Optimal $\Delta^*(a)$ as a function of attained service a.

4. OPTIMAL DISCIPLINE

In this section we introduce the Gittins discipline and give the main contribution of the paper, which reveals the optimal scheduling discipline for the NBUE + DHR(k) service time distributions.

4.1 Gittins discipline

Let $\pi^* \in \Pi$ be such that service is always given to the job with the highest Gittins index. We call this discipline the *Gittins discipline*. In view of the interpretation of the function $J(a, \Delta)$, the Gittins discipline serves the job (with a certain quota) that returns the highest payoff/investment ratio among all the jobs present in the system.

To illustrate this rule, we consider the Pareto distribution related to Figure 2. Let us determine the structure of the Gittins discipline in this case. Assume that the queue is empty and that a new job arrives. Obviously this job will start being served immediately. From Figure 2, we see that

$$G(a) \geq G(0), \quad \text{for all } a < \Delta^*(0).$$

Thus, independently of the arrival process, the service of this job will not be interrupted until it gets $\Delta^*(0)$ units of service. Once this happens, if no new job has previously arrived, the original job will continue being served. However, since

$$G(a) \leq G(0), \quad \text{for all } a \geq \Delta^*(0),$$

it will be assigned an infinitesimally small quota each time. Eventually either the service of the original job is completed or a new job arrives. In the latter case, the server will immediately start serving the new job in a non-preemptive fashion until it gets $\Delta^*(0)$ units of service. Finally, consider the situation that all jobs in the system have obtained more units of service than $\Delta^*(0)$. From Figure 2 we observe that the Gittins index decreases as the attained service increases. Thus, among these jobs, priority is given to the job with least amount of service. In other words, these jobs will be served according to the FB discipline. Thus, we conclude that, in this example, the Gittins discipline turns out to be $\mathrm{FCFS} + \mathrm{FB}(\Delta^*(0))$.

Let us now return to the general setting. It is known that the Gittins discipline π^* is optimal with respect to the mean delay $\overline{T}$ for the M/G/1 queue, see [9, Theorem 3.28], [21, Theorem 4.7] and [18].

THEOREM 1. $\overline{T}^{\pi^*} \leq \overline{T}^{\pi}$ *for any* $\pi \in \Pi$.

Note that this result is, on one hand, very strong: given any fixed service time distribution, we are able to derive the optimal scheduling discipline by computing the Gittins index function $G(a)$. On the other hand, the result is very implicit: it does not tell straightforwardly under which conditions some certain fixed discipline is optimal. Thus, it remains an interesting problem to find out which forms the Gittins discipline takes under different assumptions concerning the service time distribution. Below we shed light on this problem.

4.2 New optimality result

Below we give the main result of this paper.

THEOREM 2. *Assume that the service time distribution belongs to the class* $\mathrm{NBUE} + \mathrm{DHR}(k)$. *Then there is* $\theta^* \geq k$ *such that* $\mathrm{FCFS} + \mathrm{FB}(\theta^*)$ *is optimal.*

PROOF. By Proposition 1, $\Delta^*(0) \geq k$ and $G(a) \geq G(0)$ for all $a < \Delta^*(0)$. Furthermore, if $\Delta^*(0) < \infty$, the Gittins index $G(a) \leq G(0)$ and $G(a)$ is a decreasing function of a for all $a \geq \Delta^*(0)$. Thus, the Gittins discipline is $\mathrm{FCFS} + \mathrm{FB}(\Delta^*(0))$ so that the claim follows from Theorem 1. □

Note, in particular, that the optimal threshold $\theta^* = \Delta^*(0)$ does not depend on the arrival rate λ at all (but only on the parameters of the service time distribution).

Another observation is that the two extreme cases, $k = 0$ and $k = \infty$, correspond to the two known optimality results. If $\theta^* = \infty$, the discipline $\mathrm{FCFS} + \mathrm{FB}(\theta^*)$ reduces to FCFS. For sure, this is the case when $k = \infty$ so that the service time distribution belongs to NBUE. On the other hand, if $k = 0$ so that the service time distribution belongs to DHR, then $G(a)$ is decreasing for all $a \geq 0$ (by Proposition 1), which implies that the Gittins discipline is FB.

Further we note that a sufficient condition for $\theta^* < \infty$ is as follows:

$$\lim_{x \to \infty} h(x) = 0,$$

which is due to the facts that $G(0) \geq H(0) > 0$ and $G(a) = h(a)$ for all $a \geq k$ (see the proof of Proposition 1(iii)). This is the case, for example, when the distribution has a power-law tail like in all our examples in Section 2.

Finally we note that, in fact, the priority jobs with attained service less than the optimal threshold θ^* may be served in any *non-preemptive* fashion. Thus, FCFS may be replaced, for example, by the Random-Order-Service (ROS) discipline, or the Last-Come-First-Served (LCFS) discipline that does not allow preemption.

5. NUMERICAL RESULTS

In this section we give some numerical results related to Theorem 2 to cast light on the optimal threshold θ^*, as well as the maximum gain achieved by the optimal discipline. Throughout this section we consider the Pareto service time distribution (1), which also appeared in Example 1 of Section 2.

We start by recalling the expressions for the conditional mean delay for disciplines FB and $\mathrm{FCFS} + \mathrm{FB}(\theta)$ found, *e.g.*, in [12]. The conditional mean delay for the FB discipline with service time requirement of x reads as follows:

$$\overline{T}^{\mathrm{FB}}(x) = \frac{\overline{W}_x + x}{1 - \rho_x}.$$

Here $\overline{W}_x$ refers to the mean workload (*i.e.*, unfinished work) in an M/G/1 queue with truncated service times $S \wedge x = \min\{S, x\}$ given by the Pollaczek-Khinchin formula,

$$\overline{W}_x = \frac{\lambda E[(S \wedge x)^2]}{2(1 - \rho_x)},$$

with $\rho_x = \lambda E[(S \wedge x)]$ referring to the truncated load. For the $\mathrm{FCFS} + \mathrm{FB}(\theta)$ discipline, the conditional mean delay reads as follows:

$$\overline{T}^{\mathrm{FCFS+FB}(\theta)}(x) = \begin{cases} \overline{W}_\theta + x, & 0 < x \leq \theta \\ \overline{T}^{\mathrm{FB}}(x), & x > \theta. \end{cases}$$

Given the conditional mean delay $\overline{T}^{\pi}(x)$ for all x, we get the (unconditional) mean delay by $\overline{T}^{\pi} = \int_0^\infty \overline{T}^{\pi}(x) f(x)\, dx$. Note that, since the density $f(x) = 0$ for all $x < k$, we have

$$\overline{T}^{\mathrm{FCFS+FB}(k)} = \overline{T}^{\mathrm{FB}}. \tag{9}$$

Furthermore, the mean delay for the FCFS discipline is known to be

$$\overline{T}^{\mathrm{FCFS}} = \frac{\lambda E[S^2]}{2(1 - \rho)} + E[S].$$

We note that for the Pareto service time distribution $\overline{T}^{\mathrm{FCFS}} < \infty$ only if $\alpha > 2$, while $\overline{T}^{\mathrm{FB}} < \infty$ for all $\alpha > 1$.

The results presented in Figures 3 and 4 are related to the Pareto distribution with parameters $k = 1$, $\alpha = 2$. With these parameter values, we have

$$\overline{T}^{\mathrm{FB}} < \infty, \quad \overline{T}^{\mathrm{FCFS}} = \infty$$

for any load $0 < \rho < 1$.

In Figure 3 we depict the mean delay ratio

$$\frac{\overline{T}^{\mathrm{FCFS+FB}(\theta)}}{\overline{T}^{\mathrm{FB}}}$$

for $\theta \geq k$ and loads $\rho = 0.5$ and 0.8. The curves start from 1 as stated by (9). The maximum gain is achieved at point $\theta^* = \Delta^*(0) = 3.732$ (independently of the load), where the mean delay reduces by 18 % [13 %] for load 0.8 [0.5] as FB is replaced by the optimal $\mathrm{FCFS} + \mathrm{FB}(\theta^*)$ discipline.

In Figure 4 we plot the minimum mean delay ratio

$$\frac{\overline{T}^{\mathrm{FCFS+FB}(\theta^*)}}{\overline{T}^{\mathrm{FB}}}$$

for different loads ρ. Note that, even though the optimal threshold θ^* does not depend on the load, the gain obtained by the optimal discipline does. The maximum gain of 18 % is achieved with load $\rho = 0.8$.

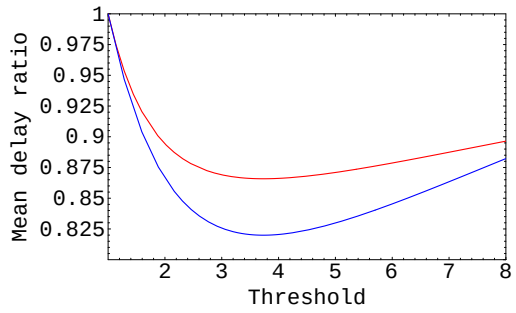

Figure 3: Pareto distribution with parameters $k = 1$, $\alpha = 2$. Mean delay ratio between $\mathrm{FCFS} + \mathrm{FB}(\theta)$ and FB as a function of threshold θ for loads $\rho = 0.5$ (upper curve) and $\rho = 0.8$ (lower curve).

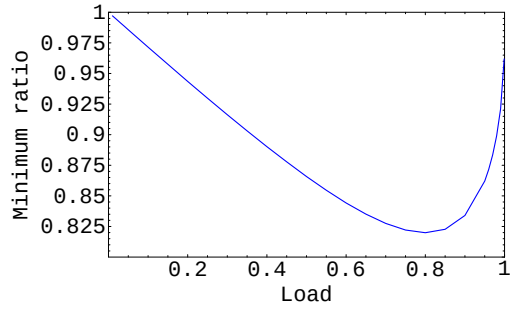

Figure 4: Pareto distribution with parameters $k = 1$, $\alpha = 2$. Minimum mean delay ratio as a function of load ρ.

The results presented in Figures 5 and 6 are related to the Pareto distributions with fixed $k = 1$ and varying α.

In Figure 5 we present the optimal threshold $\theta^* = \Delta^*(0)$ as a function of the shape parameter α. We recall that this optimal threshold does not depend on the load ρ as long as α is fixed. As seen from the figure, the optimal threshold becomes larger for larger values of α. The growth is rather linear and moderate.

In Figure 6 we compare the optimal $\mathrm{FCFS} + \mathrm{FB}(\theta^*)$ discipline to the FB and FCFS disciplines by plotting the minimum mean delay ratios

$$\frac{\overline{T}^{\mathrm{FCFS+FB}(\theta^*)}}{\overline{T}^{\mathrm{FB}}}, \quad \frac{\overline{T}^{\mathrm{FCFS+FB}(\theta^*)}}{\overline{T}^{\mathrm{FCFS}}}$$

as a function of the shape parameter α for loads $\rho = 0.5$ and 0.8. As α is varying, the arrival rate λ is modified accordingly to keep the load fixed. As we see, when compared to FB, the greater α, the greater the gain. But when compared to FCFS, the gain is maximal for $\alpha \leq 2$, and it starts to decrease as α increases from value $\alpha = 2$ on. In addition, one can observe that the two curves related to the same reference discipline with different loads do not cross each other. Thus, the gain achieved with the heavier load $\rho = 0.8$ is greater than that of the lighter load $\rho = 0.5$ when α ranges from 1.3 to 3.1.

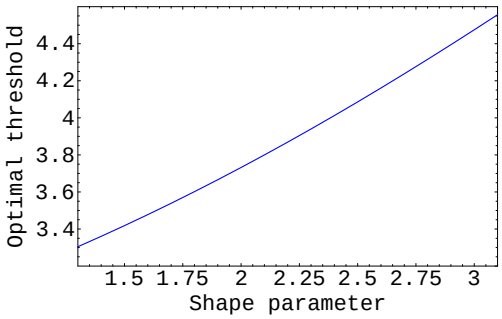

Figure 5: Pareto distributions with parameters fixed $k = 1$ and varying α. Optimal threshold $\theta^* = \Delta^*(0)$ as a function of α.

6. IMPACT OF AN UPPER BOUND

In this section we consider the case where the service time distribution is upper bounded (contrary to the previous sections), and investigate the impact on the structure of the optimal scheduling discipline by a numerical example. In particular we consider an upper bounded Pareto distribution (*cf.*, for example, [6]) defined by

$$\overline{F}(x) = \begin{cases} 1, & 0 \leq x < k, \\ 1 - \dfrac{1 - (k/x)^\alpha}{1 - (k/p)^\alpha}, & k \leq x < p, \\ 0, & x \geq p. \end{cases}$$

The hazard rate is given by

$$h(x) = \begin{cases} 0, & 0 \leq x < k, \\ \dfrac{\alpha}{x\,(1 - (x/p)^\alpha)}, & k \leq x < p. \end{cases}$$

It is easy to see that $\lim_{x \to p} h(x) = \infty$. Intuitively, this means that as the attained service gets close to p, the prob-

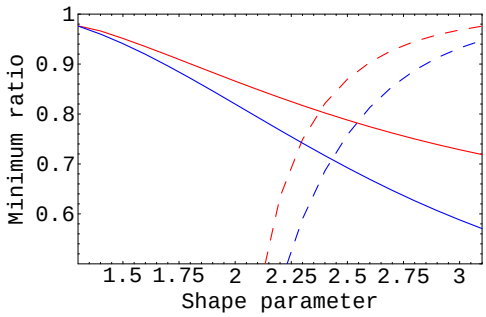

Figure 6: Pareto distributions with parameters fixed $k = 1$ and varying α. Minimum mean delay ratios between FCFS + FB(θ^*) and FB (solid lines) and between FCFS + FB(θ^*) and FCFS (dashed lines) as a function of α for loads $\rho = 0.5$ (upper curves) and $\rho = 0.8$ (lower curves).

ability that the job will depart approaches 1. We also note that the hazard rate is strictly decreasing for all $k < x < p/(1+\alpha)^{1/\alpha}$, and strictly increasing for all $x > p/(1+\alpha)^{1/\alpha}$.

Let us consider a particular example with parameters $k = 1$, $\alpha = 2$, and $p = 20$. For these values $p/(1+\alpha)^{1/\alpha} = 11.547$. The hazard rate $h(x)$ is plotted in Figure 7.

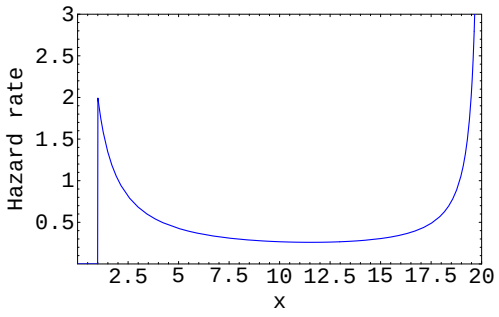

Figure 7: Hazard rate $h(x)$ for the upper bounded Pareto distribution with parameters $k = 1$, $p = 20$, $\alpha = 2$.

The Gittins index $G(a)$ as well as the optimal quota $\Delta^*(a)$ for this example are plotted in Figure 8, allowing us to determine what is the optimal scheduling discipline. We expect that our qualitative conclusions will also be valid for a wide variety of the choice of the parameters.

The Gittins index has a clear structure and can be divided into three parts: (i) $0 \le a < \Delta^*(0) = 3.855$, (ii) $3.855 \le a < 8.284$, and (iii) $8.284 \le a \le p = 20$. In the interval (i) we have $G(a) \ge G(0)$ for all a, and in the interval (ii) the Gittins index $G(a)$ is decreasing for all a. Hence, in the intervals (i) and (ii), the optimal policy is equal to what we have already seen previously, that is, non-preemptive service until $\Delta^*(0) = 3.855$, and FB from that on. However, in the interval (iii), $G(a)$ is increasing for all a, and as a consequence the structure of the optimal discipline changes when a job attains service equal to 8.284.

To explain this let us assume that there are two jobs in the system with the same amount of attained service a with

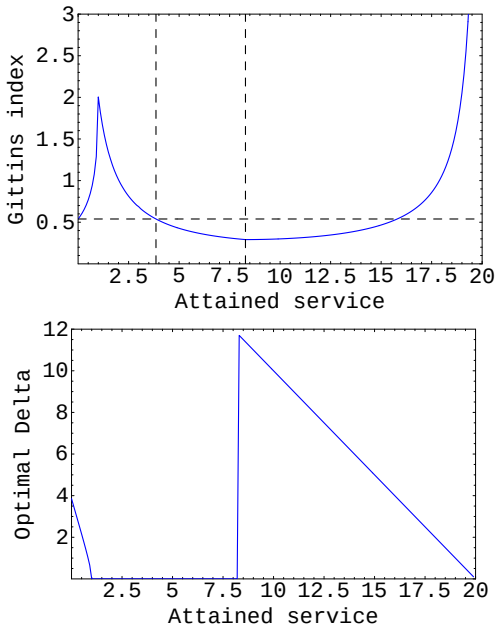

Figure 8: Upper bounded Pareto distribution with parameters $k = 1$, $p = 20$, $\alpha = 2$. *Top*: Gittins index $G(a)$ as a function of attained service a with the horizontal dashed line equal to $G(0)$ and the two dashed vertical lines equal to $a = \Delta^*(0) = 3.855$ and $a = 8.284$, respectively. *Bottom*: Optimal $\Delta^*(a)$ as a function of attained service a.

$3.855 < a < 8.284$. These two jobs will be served with a FB discipline, that is, they will get served in parallel and each one will get a half of the service capacity. When they reach an attained service equal to 8.284, the optimal discipline will pick up one of the two jobs, and since $G(a)$ is increasing for all $a \ge 8.284$ it will start serving it in a non-preemptive fashion. Suppose that when the job has attained service of 12.5 units, a new job arrives. Since the Gittins index at this point satisfies $G(12.5) < G(0)$, the new arrival will start being served. Interestingly, the new job will be served until it obtains 6.436 units of service, where $G(6.436) = G(12.5) = 0.347$. At this moment, the scheduler will switch jobs, and it will start serving the job whose service had been interrupted. Once the Gittins index of this job exceeds $G(0)$, its service will not be interrupted anymore by new arrivals, and hence it will continue being served until it leaves the system.

Therefore, we conclude that the optimal policy in each of the intervals will be (i) FCFS, (ii) FB, and (iii) FCFS. But interestingly, the priority given is not in that order. We can only say that the jobs with attained service belonging to the interval (i) have a strict priority over the jobs whose attained service belongs to the interval (ii), but their relation to the jobs with attained service belonging to the interval (iii) changes dynamically.

7. CONCLUSIONS

In this paper we have studied the problem of optimal scheduling in the M/G/1 queue when the service time distribution has a DHR tail. In particular, we have considered the

Pareto distribution, which is used in certain queueing theory related problems, for example, when modeling flow sizes in the Internet. Contrary to what is optimal for pure DHR service time distributions, we show that the optimal scheduling discipline for the Pareto distribution is not FB, but rather a two-level scheduling discipline, where jobs that have attained service less than a certain threshold are served in a non-preemptive way. The gain achieved by the optimal discipline (as compared to FB) is most remarkable for medium, or heavier than medium, loads. Only for the lightest or the heaviest loads possible, the gain remains negligible.

In fact we have found the optimal discipline not only for the Pareto service time distribution but for a whole class of distributions that satisfy the NBUE condition for small service times and the DHR condition for the tail. In addition, we have briefly illustrated the case where the service time distribution has a finite support (like the upper bounded Pareto distribution). Continuing into this direction is interesting but also challenging.

8. 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. Mean delay optimization for the M/G/1 queue with Pareto type service times. In *ACM SIGMETRICS 2007*, pages 383–384, San Diego, CA, 2007.

[3] S. Aalto, U. Ayesta, and E. Nyberg-Oksanen. Two-level processor-sharing scheduling disciplines: mean delay analysis. In *ACM SIGMETRICS/PERFORMANCE 2004*, pages 97–105, New York, NY, 2004.

[4] K. Avrachenkov, U. Ayesta, P. Brown, and E. Nyberg. Differentiation between short and long TCP flows: predictability of the response time. In *IEEE Infocom 2004*, pages 762–773, Hong Kong, 2004.

[5] M. Crovella and A. Bestavros. Self-similarity in world wide web traffic: evidence and possible causes. In *ACM SIGMETRICS 1996*, pages 160–169, Philadelphia, PA, 1996.

[6] M. E. Crovella, R. Frangioso, and M. Harchol-Balter. Connection scheduling in web servers. In *USITS'99*, Boulder, CO, 1999.

[7] A. Feldmann and W. Whitt. Fitting mixtures of exponentials to long-tail distributions to analyze network performance models. In *IEEE Infocom 1997*, pages 1096–1104, Kobe, Japan, 1997.

[8] H. Feng and V. Misra. Mixed scheduling disciplines for network flows. *ACM SIGMETRICS Performance Evaluation Review*, 31(2):36–39, 2003.

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

[10] L. Guo and I. Matta. Scheduling flows with unknown sizes: approximate analysis. In *ACM SIGMETRICS 2002*, pages 276–277, Marina del Rey, CA, 2002.

[11] M. Harchol-Balter, K. Sigman, and A. Wierman. Asymptotic convergence of scheduling policies with respect to slowdown. *Performance Evaluation*, 49:241–256, 2002.

[12] L. Kleinrock. *Queueing Systems, Volume II: Computer Applications*. Wiley, New York, 1976.

[13] I. Rai, E. Biersack, and G. Urvoy-Keller. Size-based scheduling to improve the performance of short TCP flows. *IEEE Network*, 19:12–17, 2005.

[14] I. Rai, G. Urvoy-Keller, and E. Biersack. Analysis of LAS scheduling for job size distributions with high variance. In *ACM SIGMETRICS 2004*, pages 218–228, San Diego, CA, 2003.

[15] I. Rai, G. Urvoy-Keller, M. Vernon, and E. Biersack. Performance analysis of LAS-based scheduling disciplines in a packet switched network. In *ACM SIGMETRICS/PERFORMANCE 2004*, pages 106–117, New York, NY, 2004.

[16] R. Righter and J. Shanthikumar. Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures. *Probability in the Engineering and Informational Sciences*, 3:323–333, 1989.

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

[18] 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.

[19] R. Wu and D. G. Down. Scheduling multi-server systems using foreground-background processing. In *the Forty-second Allerton Conference*, 2004.

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

[21] S. Yashkov. Mathematical problems in the theory of shared-processor systems. *Journal of Mathematical Sciences*, 58:101–147, 1992.