On the non-optimality of the FB discipline within the service time distribution class IMRL

Samuli Aalto
Helsinki University of Technology, Networking Laboratory
P.O. Box 3000, FIN-02015 TKK, Finland
email: samuli.aalto@tkk.fi

Urtzi Ayesta
CWI
Kruislaan 413, P.O. Box 94079, 1090GB Amsterdam, The Netherlands
email: u.ayesta@cwi.nl

Abstract

It is known that, within the service time distribution class DHR, the FB discipline minimizes the mean delay in the M/G/1 queue among all work-conserving and non-anticipating service disciplines. It is also believed that a similar result is valid within a more general distribution class IMRL. However, we point out a flaw in the existing proof of this latter result that cannot be overcome. Moreover, by constructing a counter-example, we demonstrate that FB is not optimal within class IMRL. On the other hand, we prove that the mean delay for FB is smaller than that of PS within class IMRL giving a weaker version of an earlier hypothesis.

Keywords Queueing theory, scheduling, M/G/1, IMRL, mean delay, FB, PS, MLPS

1 Introduction

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 non-anticipating service (scheduling) discipline π. Discipline π is work-conserving if it does not idle when there are jobs waiting and non-anticipating if the remaining service times of jobs are unknown for the server. Let Π denote the family of such service disciplines. For example, the well known disciplines FCFS and PS belong to this family, while SRPT (Shortest-Remaining-Processing-Time) does not.

Let $F(x) = P\{S \leq x\}$, $x \geq 0$, denote the cumulative service time distribution function of any job. We assume that $F(x) < 1$ for all x. If the service time distribution has density $f(x)$, the hazard rate $h(x)$ is defined by

$$h(x) = \frac{f(x)}{\int_x^\infty f(y)\,dy}. \quad (1)$$

A service time distribution belongs to class DHR (Decreasing Hazard Rate) if $h(x)$ is decreasing for all x, *i.e.*, $h(x) \geq h(y)$ whenever $x \leq y$.

Yashkov [10] has shown that, within class DHR, the mean delay is minimized by the FB (Foreground-Background) discipline, which gives priority to the job with the least attained service. In fact, Righter and Shanthikumar [7] have proved that FB minimizes, not only the mean delay and the mean queue length, but the queue length even in the stochastic sense. The FB discipline is also known as FBPS (Feedback Processor-Sharing), LAST (Least-Attained-Service-Time) and LAS (Least-Attained-Service).

Let then $\overline{F}(x)$ denote the corresponding tail distribution function, $\overline{F}(x) = 1 - F(x)$, and define

$$H(x) = \frac{\overline{F}(x)}{\int_x^\infty \overline{F}(y)\,dy}. \quad (2)$$

A service time distribution belongs to class IMRL (Increasing Mean Residual Lifetime) if $H(x)$ is decreasing for all x, *i.e.*, $H(x) \geq H(y)$ whenever $x \leq y$. This is due to the fact that

$$E[S - x \mid S > x] = \frac{\int_x^\infty \overline{F}(y)\,dy}{\overline{F}(x)} = \frac{1}{H(x)}. \quad (3)$$

Note that this conditional expectation is well defined since we assumed that $\overline{F}(x) > 0$ for all x. It is known that IMRL is a weaker condition than DHR. In other words, DHR $\subset$ IMRL. Righter *et al.* [8, Theorem 3.14] state that FB minimizes the mean delay even within class IMRL. (Unfortunately, there is a misprint in the abstract of [8] stating just the opposite. The right form is given in [8, Theorem 3.14].)

Still a more general class consists of those service time distributions for which $C^2[S] \geq 1$, where $C^2[S]$ refers to the squared cofficient of variation of the service time distribution, $C^2[S] = D^2[S]/E[S]^2$. Wierman *et al.* [9, Example 1] demonstrate, by constructing a counter-example, that FB is not optimal within this class. In particular, they evert the hypothesis by Coffman and Denning that the mean delay for FB would be smaller than that of PS (Processor-Sharing) whenever $C^2[S] > 1$. The distribution given in their counter-example, while having a greater squared cofficient of variation than 1, does not belong to class IMRL. (Unfortunately, there is also a misprint in [9, Example 1]. The corrected version reads as follows: $P\{S = 1\} = \frac{4}{5} + \epsilon$ and $P\{S = 6\} = \frac{1}{5} - \epsilon$. Then $C^2[S] > 1$ for any $0 < \epsilon < \frac{1}{10}$.)

In this paper, we prove that, contrary to [8, Theorem 3.14], FB does *not* minimize the mean delay within class IMRL. More specifically, we first identify a flaw in the proof of [8, Theorem 3.14] that cannot be overcome. Then we choose a service time distribution that belongs to IMRL but not to DHR, and construct a discipline for which the mean delay is smaller than that of FB. On the other hand, we prove that the mean delay for FB is smaller than that of PS within class IMRL giving a weaker version of the hypothesis by Coffman and Denning.

The rest of the paper is organized as follows. First, in Section 2, we recall an essential point in the proof of [8, Theorem 3.14], which a relationship between the mean delay and so called truncated workload. Then, in Section 3, we consider the sample paths of the truncated workload process, and prove that FB is not pathwisely optimal, which contradicts with [8, Lemma 3.5]. In Section 4, we prove that FB is neither optimal with respect to the mean truncated workload but still outperforms PS. In Section 5, we finally construct the counter-example demonstrating that FB does not minimize the mean delay within class IMRL, which contradicts with [8, Theorem 3.14].

2 Relationship between mean delay and truncated workload

In this section we recall a relationship between the mean delay and so called truncated workload, which the proof of [8, Theorem 3.14] is based on. We define the truncated workload as the sum of remaining service times of those jobs in the system whose attained service is less than a given truncation threshold x. The relationship is given below in Proposition 1.

Let $\overline{N}^{\pi}$ denote the mean number of jobs in the system and $\overline{V}_x^{\pi}$ the mean truncated workload with truncation threshold x. Righter *et al.* [8, Lemma 3.12] show that, for any $\pi \in \Pi$,

$$\overline{N}^{\pi} = \int_{0-}^{\infty} H(x)\, d\overline{V}_x^{\pi}, \tag{4}$$

where $H(x)$ is as given in (2). In fact, Righter *et al.* define the truncated workload as the sum of remaining service times of those jobs in the system whose attained service is less than *or equal to* a given truncation threshold x. However, since $H(x)$ is continuous from the right, formula (4) is valid also with our definition.

Assume then that the function $H(x)$ is monotonous so that the service time distribution belongs to either IMRL or DMRL. In this case the mean number of jobs in two systems with disciplines $\pi, \pi' \in \Pi$, respectively, may be compared as follows:

$$\overline{N}^{\pi} - \overline{N}^{\pi'} = -\int_{0}^{\infty} (\overline{V}_x^{\pi} - \overline{V}_x^{\pi'})\, dH(x). \tag{5}$$

This equation follows from (4) after a partial integration, and can be found from the proof of [8, Theorem 3.14]. Thus, if $\overline{V}_x^{\pi} \leq \overline{V}_x^{\pi'}$ for all x and the service time distribution belongs to class IMRL so that $-H(x)$ is increasing, then

$$\overline{N}^{\pi} \leq \overline{N}^{\pi'}.$$

Let then $\overline{T}^{\pi}$ denote the mean delay of a job. By applying Little's result, $\overline{N}^{\pi} = \lambda \overline{T}^{\pi}$, we finally get the following relationship between the mean delay and truncated workload.

Proposition 1 *Let $\pi, \pi' \in \Pi$. Assume that the service time distribution belongs to class IMRL. If $\overline{V}_x^{\pi} \leq \overline{V}_x^{\pi'}$ for all x, then $\overline{T}^{\pi} \leq \overline{T}^{\pi'}$.*

Righter *et al.* [8] use this result, which is valid *an sich*, to show that FB minimizes the mean delay within class IMRL. The problem lies in their Lemma 3.5, where they state that FB is optimal with respect to the truncated workload even in each *sample path*. Below in Proposition 3, we, however, show that this is not the case. Furthermore, in Proposition 4, we show that FB is neither optimal with respect to the *mean* truncated workload. Finally, in Theorem 1, we show that FB is not optimal with respect to the *mean delay*.

In all these cases, a better discipline is found from the family of MLPS disciplines, which were introduced by L. Kleinrock in the early 1970's [5]. They have recently attracted attention in the context of the Internet as an appropriate flow-level model for the bandwidth sharing obtained when priority is given to short TCP connections [4, 3, 2, 1, 6].

An MLPS discipline π is defined by a finite set of level thresholds $a_1 < \cdots < a_N$ defining $N+1$ levels, $N \geq 0$. A job belongs to level n if its attained service is at least a_{n-1} but less than a_n, where $a_0 = 0$ and $a_{N+1} = \infty$. Between these levels, a strict priority discipline is applied with the lowest level having the highest priority. Thus, those jobs with attained service less than a_1 are served first. Within each level n, an internal discipline D_n is applied. The internal disciplines may vary in the set $\{\mathrm{FB}, \mathrm{PS}, \mathrm{FCFS}\}$. We denote by MLPS the family of MLPS disciplines. All these disciplines belong to the family Π of work-conserving and non-anticipating disciplines.

In particular, we are interested in those MLPS disciplines π for which the internal discipline at the first level is FCFS. For any $x > 0$, let FCFS_x denote the subfamily of MLPS disciplines π such that $a_1 = x$ and $D_1 = \mathrm{FCFS}$. For example, the two-level discipline $\mathrm{FCFS} + \mathrm{FB}(x)$ with level threshold x, which applies FCFS at the first level and FB at the second level, belongs to this subfamily.

3 Sample path results for the truncated workload

In this section we show that FB is not pathwisely optimal with respect to the truncated workload, which contradicts with [8, Lemma 3.5].

Consider a single server queueing system starting empty at time $t = 0$ and obeying a service discipline $\pi \in \Pi$. We assume that jobs arrive one at a time. They are indexed with $i = 1, 2, \ldots$ according to their arrival order.

Let A_i denote the arrival epoch of job i and S_i its service time. In addition, let $X_i^\pi(t)$ denote the attained service of job i at time t. Let then $\mathcal{A}(t)$ and $\mathcal{N}^\pi(t)$ denote the set of jobs arrived and those in the system at time t, respectively,

$$\mathcal{A}(t) = \{i : A_i \leq t\}, \quad \mathcal{N}^\pi(t) = \{i \in \mathcal{A}(t) : X_i^\pi(t) < S_i\}.$$

For any $x > 0$, let $\mathcal{N}_x^\pi(t)$ denote the set of those jobs in the system whose attained service is less than x,

$$\mathcal{N}_x^\pi(t) = \{i \in \mathcal{A}(t) : X_i^\pi(t) < \min\{x, S_i\}\}.$$

Denote further $A(t) = |\mathcal{A}(t)|$, $N^\pi(t) = |\mathcal{N}^\pi(t)|$, and $N_x^\pi(t) = |\mathcal{N}_x^\pi(t)|$.

For any $x \geq 0$, let $V_x^\pi(t)$ denote the remaining workload in the system at time t contributed by those jobs with attained service less than x,

$$V_x^\pi(t) = \sum_{i \in \mathcal{N}_x^\pi(t)} (S_i - X_i^\pi(t)), \tag{6}$$

which is briefly called the *truncated workload* with truncation threshold x. Note that in the limit $x \to \infty$ the truncated workload equals the ordinary workload of this system, *i.e.*, the sum of remaining service times of all jobs, which is the same for all work-conserving disciplines.

To get another expression for the truncated workload $V_x^\pi(t)$, let $R_x^\pi(t)$ denote the total rate at which service is provided to the jobs with attained service less than x at time t. Whenever there are such jobs, the truncated workload is decreasing continuously with this rate. But it may also decrease discontinuously: whenever the attained service of a job reaches the truncation threshold x so that the job drops out from $\mathcal{N}_x^\pi(t)$, the truncated workload decreases by a step that equals the remaining service time of that job. Thus, we have

$$V_x^\pi(t) = \sum_{i \in \mathcal{A}(t)} S_i - \int_0^t R_x^\pi(u)\,du - \sum_{i \in \mathcal{A}(t) \setminus \mathcal{N}_x^\pi(t)} (S_i - \min\{x, S_i\}). \tag{7}$$

First we show that certain MLPS disciplines are not worse than FB with respect to the truncated workload.

Proposition 2 *Let $x > 0$ and $\pi^* \in \mathrm{FCFS}_x$. Then, for any $t \geq 0$,*

$$V_x^{\pi^*}(t) \leq V_x^{\mathrm{FB}}(t).$$

Proof. Both the disciplines, π^* and FB, give full priority for the jobs with attained service less than x. Thus, $R_x^{\pi^*}(t) = R_x^{\mathrm{FB}}(t)$ for all $t \geq 0$, and consequently

$$\int_0^t R_x^{\pi^*}(u)\,du = \int_0^t R_x^{\mathrm{FB}}(u)\,du.$$

It follows that the second term in (7) is the same for both disciplines.

Let us then consider the third term in (7). Note first that, for any $\pi \in \Pi$,

$$\sum_{i \in \mathcal{A}(t) \setminus \mathcal{N}_x^\pi(t)} (S_i - \min\{x, S_i\}) = \sum_{i \in \mathcal{A}_x(t) \setminus \mathcal{N}_x^\pi(t)} (S_i - \min\{x, S_i\}),$$

where

$$\mathcal{A}_x(t) = \{i \in \mathcal{A}(t) : S_i > x\}.$$

In the FB system the jobs with $S_i > x$ reach the attained service level x in batches. In the corresponding π^* system, where the discipline is FCFS below this level, the same jobs reach level x one-by-one in such a way that the job arrived last reaches this level not later than the whole batch in the FB system. This is due to the work-conserving principle. Thus,

$$\mathcal{A}_x(t) \setminus \mathcal{N}_x^{\mathrm{FB}}(t) \subset \mathcal{A}_x(t) \setminus \mathcal{N}_x^{\pi^*}(t),$$

which implies that

$$\sum_{i \in \mathcal{A}_x(t) \setminus \mathcal{N}_x^{\mathrm{FB}}(t)} (S_i - \min\{x, S_i\}) \leq \sum_{i \in \mathcal{A}_x(t) \setminus \mathcal{N}_x^{\pi^*}(t)} (S_i - \min\{x, S_i\}).$$

By combining these results with (7), we get

$$\begin{aligned} V_x^{\pi^*}(t) &= \sum_{i \in \mathcal{A}(t)} S_i - \int_0^t R_x^{\pi^*}(u)\,du - \sum_{i \in \mathcal{A}_x(t) \setminus \mathcal{N}_x^{\pi^*}(t)} (S_i - \min\{x, S_i\}) \\ &\leq \sum_{i \in \mathcal{A}(t)} S_i - \int_0^t R_x^{\mathrm{FB}}(u)\,du - \sum_{i \in \mathcal{A}_x(t) \setminus \mathcal{N}_x^{\mathrm{FB}}(t)} (S_i - \min\{x, S_i\}) \\ &= V_x^{\mathrm{FB}}(t), \end{aligned}$$

which completes the proof. □

Then we show that these disciplines are strictly better than FB.

Proposition 3 *Let $x > 0$ and $\pi^* \in \mathrm{FCFS}_x$. There is a sample path and $t \geq 0$ such that*

$$V_x^{\pi^*}(t) < V_x^{\mathrm{FB}}(t).$$

Proof. Assume that $A_1 = 0$, $S_1 = 3x/2$, $A_2 = x/2$, $S_2 = x/2$, and $A_3 = 3x$. Discipline FB serves job 1 in the intervals $[0, x/2)$ and $[x, 2x)$, and job 2 in the interval $[x/2, x)$. Whereas discipline π^* serves job 1 in the intervals $[0, x)$ and $[3x/2, 2x)$, and job 2 in the interval $[x, 3x/2)$. As a result,

$$V_x^{\pi^*}(5x/4) = x/4 < 3x/4 = V_x^{\mathrm{FB}}(5x/4).$$

In fact, $V_x^{\pi^*}(t) < V_x^{\mathrm{FB}}(t)$ for all $t \in (x, 3x/2)$, as can be seen from Fig. 1 where we have chosen $x = 2$. □

Remark 1 Note that Proposition 3 contradicts with [8, Lemma 3.5], which states that FB is optimal with respect to the truncated workload. This is essentially due to the fact that Righter *et al.* confuse the truncated workload $V_x^{\pi}(t)$ with the variable $U_x^{\pi}(t)$ defined by

$$U_x^{\pi}(t) = \sum_{i \in \mathcal{N}_x^{\pi}(t)} (\min\{S_i, x\} - X_i^{\pi}(t)), \tag{8}$$

which we call *unfinished truncated work* with truncation threshold x. So the remaining *truncated* service times are summed up instead of the ordinary remaining service times as in (6). It is true that FB is pathwisely optimal with respect to the unfinished truncated work [1, Proposition 5] but, as Proposition 3 reveals, not with respect to the truncated workload. The confusion in [8] can be explained as follows. As given in [1, Equation (16)], we have

$$U_x^{\pi}(t) = \sum_{i \in \mathcal{A}(t)} S_i - \int_0^t R_x^{\pi}(u)\,du. \tag{9}$$

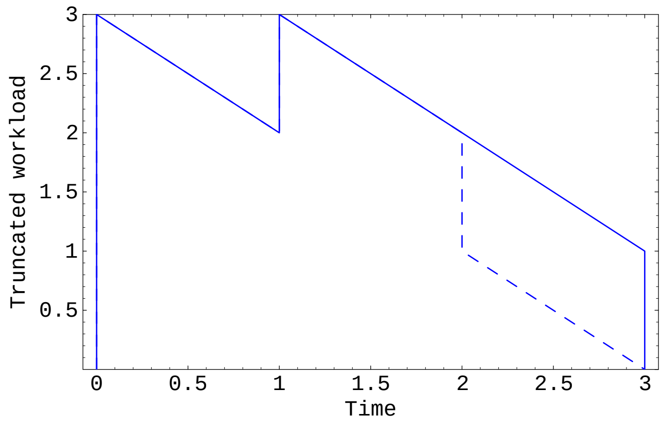

Figure 1: The truncated workload $V_x^\pi(t)$ as a function of time t for disciplines FB (solid line) and π^* (dashed line) with $x = 2$.

This corresponds to [8, Equation (3.1)]. Thus, Righter *et al.* miss the downward steps in the sample paths of the truncated workload process $V_x^\pi(t)$.

Remark 2 Strictly speaking, if the truncated workload were defined as in [8], *i.e.*, as the sum of remaining service times of those jobs in the system whose attained service is less than *or equal to* the given truncation threshold x, the disciplines π^* and FB considered in the proof of Proposition 3 would have the same truncated workload for all t. However, in that case, a slightly modified discipline $\pi^* \in \mathrm{FCFS}_{x+\epsilon}$ with sufficiently small $\epsilon > 0$ would serve as a counter-example for [8, Lemma 3.5], which is easily verified.

4 Mean value results for the truncated workload

Righter *et al.* [8] applied their Lemma 3.5 to prove that, for all $\pi \in \Pi$ and $x > 0$,

$$\overline{V}_x^{\mathrm{FB}} \leq \overline{V}_x^{\pi}, \tag{10}$$

which would be a sufficient (but not necessary) condition for the optimality of FB with respect to mean delay according to Proposition 1. However, as shown above, their Lemma 3.5 is not valid. So it remains an open question whether (10) is valid or not.

In this section we show that (10) is neither valid. We start with some preliminary results related to a modified queue where the service times are replaced by their truncated versions in Section 4.1. In Section 4.2, we derive some fundamental formulas for the mean truncated workload. Then, in Section 4.3, we give the mean truncated workload formulas for FB, PS and FCFS_x disciplines. Finally, in Section 4.4, we show that FB is not optimal with respect to the mean truncated workload but still outperforms PS.

4.1 Truncated service times

Let $x \geq 0$, and consider, for a while, a modified M/G/1 queue where the original service times S are replaced by their truncated versions $S \wedge x = \min\{S, x\}$. It is easy to see that

$$E[S \wedge x] = \int_0^x \overline{F}(y)\, dy, \quad E[(S \wedge x)^2] = 2\int_0^x y\overline{F}(y)\, dy. \tag{11}$$

Furthermore, let

$$\rho_x = \lambda E[(S \wedge x)] \tag{12}$$

denote the truncated load. The mean workload for a work conserving M/G/1 queue with truncated service times is, by the Pollaczek-Khinchin formula (cf. [5, Equation (4.26)]),

$$\overline{W}_x = \frac{\lambda E[(S \wedge x)^2]}{2(1-\rho_x)}. \tag{13}$$

Regarding the derivative, it is easy to verify that

$$\frac{\partial}{\partial x}\overline{W}_x = \lambda\, \frac{\overline{W}_x + x}{1-\rho_x}\, \overline{F}(x). \tag{14}$$

Letting $x \to \infty$, we get the ordinary Pollaczek-Khinchin formula:

$$\overline{W}_\infty = \frac{\lambda E[S^2]}{2(1-\rho)}.$$

4.2 Mean truncated workload

Consider then again the original M/G/1 queue with service times S and load ρ,

$$\rho = \lambda E[S] = \lambda \int_0^\infty \overline{F}(x)\, dx. \tag{15}$$

Note that $\rho > \rho_x$ for all $x > 0$, since we assumed that $\overline{F}(x) > 0$ for all x.

Let $\overline{V}_x^\pi$ denote the mean truncated workload with truncation threshold x. In addition, let $\overline{N}_x^\pi$ denote the mean number of those jobs in the system whose attained service is less than x. Since π is work-conserving and non-anticipating, we have (cf. [8, Equation (3.13)])

$$\overline{V}_x^\pi = \int_{0^-}^{x^-} E[S - y \mid S > y]\, d\overline{N}_y^\pi. \tag{16}$$

As noted earlier in Remark 1, the truncated workload V_x^π differs from the unfinished truncated work U_x^π. For a work-conserving and non-anticipating discipline π, the mean unfinished truncated work reads as follows:

$$\overline{U}_x^\pi = \int_{0^-}^{x^-} E[(S \wedge x) - y \mid S > y]\, d\overline{N}_y^\pi. \tag{17}$$

Note further that the limit $\overline{V}_\infty^\pi = \overline{U}_\infty^\pi = \overline{W}_\infty$ is the same for all $\pi \in \Pi$.

Let then $\overline{T}^{\pi}(x)$ denote the conditional mean delay of a job with service time x. We note that $\overline{T}^{\pi}(x)$ is increasing and continuous from the left, $\overline{T}^{\pi}(x^-) = \overline{T}^{\pi}(x) \leq \overline{T}^{\pi}(x^+)$. It is also known [5, Equation (4.11)] that

$$d\overline{N}^{\pi}_y = \lambda \overline{F}(y)\, d\overline{T}^{\pi}(y).$$

Thus, by (16) and (3),

$$\begin{aligned}
\overline{V}^{\pi}_x &\overset{(16)}{=} \lambda \int_{0^-}^{x^-} E[S-y \mid S>y]\, \overline{F}(y)\, d\overline{T}^{\pi}(y) \\
&\overset{(3)}{=} \lambda \int_{0^-}^{x^-} \int_y^{\infty} \overline{F}(t)\, dt\, d\overline{T}^{\pi}(y) \\
&= \lambda \int_0^x \overline{T}^{\pi}(y)\overline{F}(y)\, dy + \lambda \overline{T}^{\pi}(x) \int_x^{\infty} \overline{F}(y)\, dy \\
&= \lambda \int_0^x \overline{T}^{\pi}(y)\overline{F}(y)\, dy + \overline{T}^{\pi}(x)(\rho - \rho_x).
\end{aligned} \tag{18}$$

Note that $\overline{V}^{\pi}_x$ is increasing and continuous from the left, $\overline{V}^{\pi}_{x^-} = \overline{V}^{\pi}_x \leq \overline{V}^{\pi}_{x^+}$.

Since, by [5, Equation (4.60)],

$$\overline{U}^{\pi}_x = \lambda \int_0^x \overline{T}^{\pi}(y)\overline{F}(y)\, dy, \tag{19}$$

the mean truncated workload can also be given as follows:

$$\overline{V}^{\pi}_x = \overline{U}^{\pi}_x + \overline{T}^{\pi}(x)(\rho - \rho_x).$$

4.3 Mean truncated workload for some MLPS disciplines

Now we calculate the mean truncated workload for some MLPS disciplines in an M/G/1 queue. First we consider the FB discipline. By [5, Equation (4.27)],

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

Thus, by (14),

$$\lambda \overline{T}^{\mathrm{FB}}(x)\overline{F}(x) = \frac{\partial}{\partial x}\overline{W}_x$$

implying, by (18), that

$$\overline{V}^{\mathrm{FB}}_x = \overline{W}_x + \frac{\overline{W}_x + x}{1-\rho_x}(\rho - \rho_x) = \overline{W}_x\left(1 + \frac{\rho - \rho_x}{1-\rho_x}\right) + x\,\frac{\rho-\rho_x}{1-\rho_x}. \tag{20}$$

Then we consider the PS discipline. By [5, Equation (4.17)],

$$\overline{T}^{\mathrm{PS}}(x) = \frac{x}{1-\rho}.$$

Thus, by (18) and (13),

$$\overline{V}_x^{\mathrm{PS}} = \overline{W}_x \frac{1-\rho_x}{1-\rho} + x\frac{\rho-\rho_x}{1-\rho}. \tag{21}$$

Finally we consider any $\pi \in \mathrm{FCFS}_x$. By [5, Equation (4.35)], we have, for all $y \le x$,

$$\overline{T}^{\pi}(y) = W_x + y.$$

Thus, by (18), for all $y \le x$,

$$\overline{V}_y^{\pi} = \overline{W}_x\rho + \overline{W}_y(1-\rho_y) + y(\rho-\rho_y).$$

In particular,

$$\overline{V}_x^{\pi} = \overline{W}_x(1+\rho-\rho_x) + x(\rho-\rho_x). \tag{22}$$

In the sequel we also need the following equality, which is valid for any $\pi \in \mathrm{FCFS}_x$,

$$\overline{U}_x^{\pi} = \overline{U}_x^{\mathrm{FB}} = \overline{W}_x. \tag{23}$$

This is easily verified by (19) and the conditional mean delay formulas given above. In fact, this property follows from the pathwise local optimality of these policies with respect to the unfinished truncated work, see [1, Section 3.1].

4.4 Comparison with the FB discipline

Now we are ready to show that FB is not optimal with respect to the mean truncated workload.

Proposition 4 *Consider any service time distribution. Let $x > 0$ and $\pi^* \in \mathrm{FCFS}_x$. Then*

$$\overline{V}_x^{\pi^*} < \overline{V}_x^{\mathrm{FB}}.$$

Proof. By (20) and (22),

$$\overline{V}_x^{\mathrm{FB}} - \overline{V}_x^{\pi^*} = \frac{(\overline{W}_x + x)(\rho-\rho_x)\rho_x}{1-\rho_x} > 0,$$

since $0 < \rho_x < \rho < 1$. □

As a numerical example, consider the exponential service time distribution with rate parameter equal to 1, $\overline{F}(x) = e^{-x}$ and $E[S] = 1$. Let $\lambda = 1/2$, so that we have a stable system with load $\rho = \lambda E[S] = 1/2 < 1$. Then

$$\overline{V}_1^{\mathrm{FB}} = 0.514, \qquad \overline{V}_1^{\mathrm{FCFS+FB(1)}} = 0.413,$$

implying that (10) is not valid in this case. This result is illustrated also in Fig. 2, where we have depicted the mean truncated workload $\overline{V}_x^{\pi}$ as a function of the truncation threshold x for disciplines FB and FCFS+FB(1).

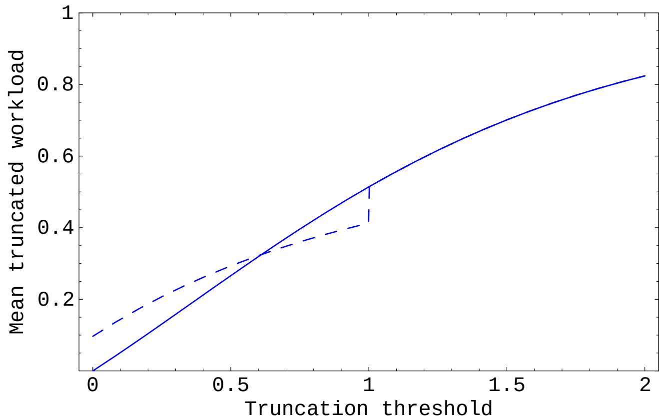

Figure 2: The mean truncated workload $\overline{V}_x^{\pi}$ as a function of the truncation threshold x for disciplines FB (solid line) and FCFS+FB(1) (dashed line).

Remark 3 Again strictly speaking, if the truncated workload were defined as in [8], *i.e.*, as the sum of remaining service times of those jobs in the system whose attained service is less than *or equal to* the given truncation threshold x, the disciplines π^* and FB considered in Proposition 4 could have the same mean truncated workload for the truncation threshold x. However, in that case, we can easily verify that, for any sufficiently small $\epsilon > 0$,

$$\overline{V}_{x-\epsilon}^{\mathrm{FB}} - \overline{V}_{x-\epsilon}^{\pi^*} > 0,$$

which is also demonstrated by the numerical example in Fig. 2.

Finally we show that, while not being optimal, FB still outperforms PS with respect to the mean truncated workload.

Proposition 5 *Consider any service time distribution. Let $x > 0$. Then*

$$\overline{V}_x^{\mathrm{FB}} < \overline{V}_x^{\mathrm{PS}}.$$

Proof. By (20) and (21),

$$\overline{V}_x^{\mathrm{PS}} - \overline{V}_x^{\mathrm{FB}} = \frac{(\overline{W}_x + x)(\rho - \rho_x)^2}{(1-\rho)(1-\rho_x)} > 0,$$

since $\rho_x < \rho < 1$. □

5 Mean delay results

In this section we justify our principal claim that FB does not minimize the mean delay within class IMRL, which contradicts with [8, Theorem 3.14]. On the other hand, we show that FB still outperforms PS with respect to the mean delay within class IMRL, which can be considered as a generalization of [9, Theorem 1].

Theorem 1 *There is a service time distribution belonging to class IMRL and* $\pi \in \Pi$ *such that*

$$\overline{T}^{\pi} < \overline{T}^{\mathrm{FB}}.$$

Proof. 1° First we need a distribution that belongs to IMRL but not to DHR. Here is a candidate (for any $c > 1$):

$$\overline{F}(x) = \begin{cases} c^{-x}, & 0 \leq x \leq c \\ x^{-c}, & x > c. \end{cases} \tag{24}$$

So we first have an exponential part and then a Pareto-type tail. The corresponding density function is

$$f(x) = \begin{cases} c^{-x}\ln c, & 0 \leq x \leq c \\ cx^{-c-1}, & x > c. \end{cases} \tag{25}$$

The hazard rate is as follows:

$$h(x) = \begin{cases} \ln c, & 0 \leq x \leq c \\ cx^{-1}, & x > c. \end{cases} \tag{26}$$

It is easy to see that $F(x)$ belongs to DHR if and only if $h(c-) \geq h(c+)$, which is equivalent to the requirement that $c \geq e$. In Fig. 3 we have depicted the function $h(x)$ for $c = 2$.

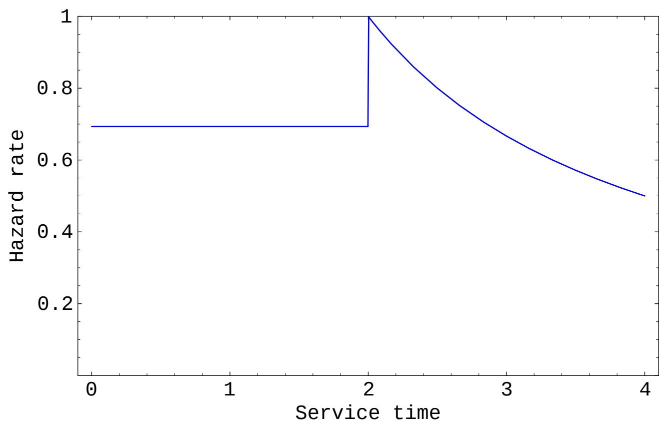

Figure 3: Hazard rate for the service time distribution defined in (24) with $c = 2$.

The mean residual lifetime function is as follows:

$$\frac{1}{H(x)} = \begin{cases} \dfrac{1}{\ln c} + \left(\dfrac{c}{c-1} - \dfrac{1}{\ln c}\right) c^{x-c}, & 0 \leq x \leq c \\ \dfrac{x}{c-1}, & x > c. \end{cases} \tag{27}$$

Since, for all $c > 1$,

$$\frac{c}{c-1} - \frac{1}{\ln c} > 0,$$

we deduce from (27) that $F(x)$ belongs to IMRL with any $c > 1$. In Fig. 4 we have depicted the function $1/H(x)$ for $c = 2$. All in all, $F(x)$ belongs to IMRL but not to DHR if and only if

$$1 < c < e \approx 2.71828.$$

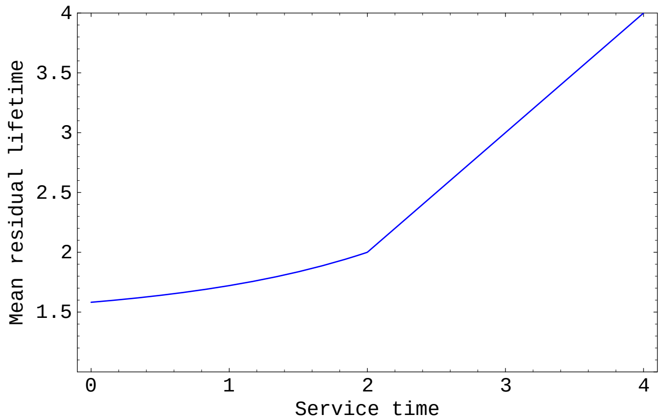

Figure 4: Mean residual lifetime for the service time distribution defined in (24) with $c = 2$.

2° Use then the following heuristic to construct a discipline that might be better than FB: whenever the hazard rate is increasing, FCFS is optimal, and whenever it is decreasing, FB is optimal. So let us compare disciplines FB and FCFS + FB($c + \epsilon$) with $\epsilon \geq 0$. The conditional mean delays read as follows:

$$\overline{T}^{\mathrm{FB}}(x) = \frac{\overline{W}_x + x}{1 - \rho_x}, \qquad x \geq 0, \tag{28}$$

and

$$\overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)}(x) = \begin{cases} \overline{W}_{c+\epsilon} + x, & 0 \leq x \leq c + \epsilon \\ \overline{T}^{\mathrm{FB}}(x), & x > c + \epsilon. \end{cases} \tag{29}$$

2.1° First we consider the FCFS + FB(c) discipline with $\epsilon = 0$ and show that

$$\overline{T}^{\mathrm{FCFS+FB}(c)} = \overline{T}^{\mathrm{FB}}. \tag{30}$$

By (28) and (29),

$$\overline{T}^{\mathrm{FCFS+FB}(c)} - \overline{T}^{\mathrm{FB}} = \int_0^c \left(\overline{T}^{\mathrm{FCFS+FB}(c)}(x) - \overline{T}^{\mathrm{FB}}(x)\right) f(x) dx.$$

Since, by (25), $f(x) = \overline{F}(x) \ln c$ for any $x \le c$, we get

$$\overline{T}^{\mathrm{FCFS+FB}(c)} - \overline{T}^{\mathrm{FB}} = \ln c \int_0^c \left(\overline{T}^{\mathrm{FCFS+FB}(c)}(x) - \overline{T}^{\mathrm{FB}}(x) \right) \overline{F}(x) dx,$$

implying, by (19) and (23), that

$$\overline{T}^{\mathrm{FCFS+FB}(c)} - \overline{T}^{\mathrm{FB}} = \frac{\ln c}{\lambda} \left(\overline{U}_c^{\mathrm{FCFS+FB}(c)} - \overline{U}_c^{\mathrm{FB}} \right) = 0.$$

2.2° If we now can prove that

$$\left. \frac{\partial}{\partial \epsilon} \overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} \right|_{\epsilon = 0^+} < 0, \tag{31}$$

it follows from (30) that there is $\delta > 0$ such that, for any $0 < \epsilon < \delta$,

$$\overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} < \overline{T}^{\mathrm{FB}} \tag{32}$$

revealing the non-optimality of FB for this distribution.

By (28) and (29),

$$\overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} = \int_0^{c+\epsilon} (\overline{W}_{c+\epsilon} + x) f(x)\, dx + \int_{c+\epsilon}^{\infty} \frac{\overline{W}_x + x}{1 - \rho_x} f(x)\, dx.$$

Thus, by (14), we get the desired result as follows:

$$\begin{aligned}
\frac{\partial}{\partial \epsilon} \overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} &= \left(\frac{\partial}{\partial \epsilon} \overline{W}_{c+\epsilon} \right) F(c+\epsilon) + (\overline{W}_{c+\epsilon} + c + \epsilon) f(c+\epsilon) - \frac{\overline{W}_{c+\epsilon} + c + \epsilon}{1 - \rho_{c+\epsilon}} f(c+\epsilon) \\
&\overset{(14)}{=} \lambda \frac{\overline{W}_{c+\epsilon} + c + \epsilon}{1 - \rho_{c+\epsilon}} \overline{F}(c+\epsilon)\, F(c+\epsilon) - \frac{\overline{W}_{c+\epsilon} + c + \epsilon}{1 - \rho_{c+\epsilon}} f(c+\epsilon)\, \rho_{c+\epsilon} \\
&= \lambda \frac{\overline{W}_{c+\epsilon} + c + \epsilon}{1 - \rho_{c+\epsilon}} \left((c+\epsilon)^{-c} \left(1 - (c+\epsilon)^{-c}\right) \,-\, c (c+\epsilon)^{-c-1} \frac{\rho_{c+\epsilon}}{\lambda} \right) \\
&\overset{\epsilon \to 0^+}{\to} \lambda \frac{\overline{W}_c + c}{1 - \rho_c} \left(c^{-c} \left(1 - c^{-c}\right) \,-\, c^{-c} \frac{\rho_c}{\lambda} \right).
\end{aligned}$$

Since

$$\frac{\rho_c}{\lambda} = E[S \wedge c] \overset{(11)}{=} \int_0^c \overline{F}(x)\, dx \overset{(24)}{=} \int_0^c c^{-x}\, dx = \frac{1}{\ln c}(1 - c^{-c}),$$

we finally get

$$\left. \frac{\partial}{\partial \epsilon} \overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} \right|_{\epsilon = 0^+} = \lambda \frac{\overline{W}_c + c}{1 - \rho_c} c^{-c} (1 - c^{-c}) \left(1 \,-\, \frac{1}{\ln c} \right) < 0,$$

where the inequality follows from the fact that $1/\ln c > 1$ for all $1 < c < e$. □

As numerical examples, we have computed the following results:

$$
\begin{aligned}
&c = 2.0, \quad \lambda = 0.5, \quad \rho = 0.791, \quad \epsilon = 1.0: \quad \overline{T}^{\mathrm{FB}} = 5.901, \quad \overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} = 5.811,\\
&c = 2.1, \quad \lambda = 0.5, \quad \rho = 0.733, \quad \epsilon = 0.7: \quad \overline{T}^{\mathrm{FB}} = 4.640, \quad \overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} = 4.584,\\
&c = 2.5, \quad \lambda = 0.5, \quad \rho = 0.575, \quad \epsilon = 0.2: \quad \overline{T}^{\mathrm{FB}} = 2.561, \quad \overline{T}^{\mathrm{FCFS+FB}(c+\epsilon)} = 2.558.
\end{aligned}
$$

So in all these cases $\mathrm{FCFS} + \mathrm{FB}(c + \epsilon)$ is found to be better than FB.

Theorem 2 *Assume that the service time distribution belongs to class IMRL. Then*

$$\overline{T}^{\mathrm{FB}} \leq \overline{T}^{\mathrm{PS}}.$$

Proof. Due to Proposition 1, this follows immediately from Proposition 5. □

References

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

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

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

[4] Guo, L. and Matta, I. (2002). Differentiated control of web traffic: A numerical analysis. In *SPIE ITCOM'2002*, Boston.

[5] Kleinrock, L. (1976). *Queueing Systems, Volume II: Computer Applications.* John Wiley & Sons, New York.

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

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

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

[9] Wierman, A., Bansal, N. and Harchol-Balter, M. (2004). A note on comparing response times in the M/GI/1/FB and M/GI/1/PS queues. *Operations Research Letters* **32,** 73–76.

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