Available online at www.sciencedirect.com
SCIENCE DIRECT®

Operations Research Letters 32 (2004) 73–76

Operations Research Letters

www.elsevier.com/locate/dsw

A note on comparing response times in the M/GI/1/FB and M/GI/1/PS queues☆

Adam Wierman*, Nikhil Bansal, Mor Harchol-Balter

Computer Science Department, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Received 20 September 2002; received in revised form 15 April 2003; accepted 16 April 2003

Abstract

We compare the overall mean response time (a.k.a. sojourn time) of the processor sharing (PS) and feedback (FB) queues under an M/GI/1 system. We show that FB outperforms PS under service distributions having decreasing failure rates; whereas PS outperforms FB under service distributions having increasing failure rates.

Keywords: Scheduling; M/G/1; Feedback (FB); Least attained service (LAS); Shortest elapsed time (SET); Processor sharing (PS); Sojourn time; Response time

Time-sharing scheduling policies, especially the processor-sharing (PS) discipline, are used quite frequently in modern systems. Under PS the processor is shared evenly among all jobs currently in the system. Although PS is commonly used, it provides a far from optimal mean response time (a.k.a sojourn time). The shortest-remaining-processing-time (SRPT) policy is known to be optimal with respect to overall mean response time [4]. This policy schedules the job having the smallest remaining size at all times; thus SRPT requires knowledge of the job's size (a.k.a service requirement). In the absence of this knowledge, the feedback (FB) policy [1] has long been proposed as an approximation to SRPT. Under FB the job with the least attained service (a.k.a age) gets the processor to itself. If several jobs all have the least attained service, they time-share the processor via PS. By biasing towards the jobs with small ages, FB is, in a sense, attempting to complete the short jobs as quickly as possible. The goal of this note is to understand under which service distributions FB improves upon the mean response time of PS.

The response times for both M/GI/1/PS and M/GI/1/FB are well known. We define $\rho \stackrel{\text{def}}{=} \lambda E[X]$ where λ is the arrival rate and X is a random variable sampled from the service distribution $F(x)$ with density function $f(x)$. Let $E[T]^P$ denote the mean response time under policy P and let $E[T(x)]^P$ denote

☆ This work was supported by NSF Career Grant CCR-0133077, by NSF ITR Grant 99-167 ANI-0081396, by Cisco Systems, Spinnaker Networks via Pittsburgh Digital Greenhouse Grant 01-1, and an IBM Research Fellowship.

* Corresponding author.

E-mail addresses: acw@cs.cmu.edu (A. Wierman), nikhil@cs.cmu.edu (N. Bansal), harchol@cs.cmu.edu (M. Harchol-Balter).

[1] Note that FB is sometimes referred to by three other names: Generalized foreground–background (FB), least-attained-service (LAS), and shortest-elapsed-time (SET).

0167-6377/03/$ - see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/S0167-6377(03)00061-0

the expected response time of a job of size x under policy P. Then the following classic results exist (see [6] or [2] for a proof of these):

$$E[T(x)]^{\mathrm{PS}} = \frac{x}{1-\rho},$$

$$E[T(x)]^{\mathrm{FB}} = \frac{\lambda \int_0^x t^2 f(t)\,\mathrm{d}t + \lambda x^2 \bar{F}(x)}{2(1-\rho_x)^2} + \frac{x}{1-\rho_x}$$
$$= \frac{\lambda \int_0^x t\bar{F}(t)\,\mathrm{d}t}{(1-\rho_x)^2} + \frac{x}{1-\rho_x},$$

where

$$\rho_x = \lambda \left(\int_0^x t f(t)\,\mathrm{d}t + x\bar{F}(x) \right) = \lambda \int_0^x \bar{F}(t)\,\mathrm{d}t.$$

Notice that ρ_x can be thought of as the load of jobs from service distribution $X_x \stackrel{\mathrm{def}}{=} \min(x, X)$.

We briefly discuss some prior work comparing mean response times under M/GI/1/FB with those under M/GI/1/PS. Rai, Urvoy-Keller, and Biersack [3] prove that for *any* service distribution $E[T]^{\mathrm{FB}} \leqslant (2-\rho)/(2-2\rho)E[T]^{\mathrm{PS}}$. In our paper, we are concerned with understanding for *which* service distributions $E[T]^{\mathrm{FB}} \leqslant E[T]^{\mathrm{PS}}$. Coffman and Denning consider exactly this question and hypothesize the following relation [1, p. 188–189]:

$$E[T]^{\mathrm{FB}} < E[T]^{\mathrm{PS}} \text{ when } C > 1,$$

$$E[T]^{\mathrm{PS}} > E[T]^{\mathrm{FB}} \text{ when } C < 1,$$

where $C^2 \stackrel{\mathrm{def}}{=} \mathrm{Var}(X)/E[X]^2$ is the squared coefficient of variation of the service distribution. (Note that [1] makes the statement in terms of waiting times, but that this formulation is equivalent.)

It turns out that Coffman and Denning's hypothesis is not always true (see Example 1 below) and needs a slight refinement. Our following main theorem gives such a refinement.

Theorem 1. *Let $\mu(x) \stackrel{\mathrm{def}}{=} f(x)/\bar{F}(x)$ be the hazard rate of the service distribution. In an M/GI/1 system* FB *and* PS *relate as follows*:

1. *If $\mu(x)$ is decreasing*, $E[T]^{\mathrm{FB}} \leqslant E[T]^{\mathrm{PS}}$.
2. *If $\mu(x)$ is constant*, $E[T]^{\mathrm{FB}} = E[T]^{\mathrm{PS}}$.
3. *If $\mu(x)$ is increasing*, $E[T]^{\mathrm{FB}} \geqslant E[T]^{\mathrm{PS}}$.

Observe that this theorem is a refinement of Coffman and Denning's hypothesis because of the following well-known lemma [5, p. 16–19], which relates the hazard rate and the coefficient of variation.

Lemma 1. *When $\mu(x)$ is decreasing, $C \geqslant 1$ and when $\mu(x)$ is increasing $C \leqslant 1$.*

Notice that Theorem 1 does not say anything about distributions whose hazard rate is both strictly increasing for some x and strictly decreasing for other values of x. Our example below shows that it is exactly this situation (where hazard rate is both increasing and decreasing) which leads to a counterexample to the Coffman, Denning hypothesis.

Example 1. The following example gives a job size distribution where $C^2 > 1$ but $E[T]^{\mathrm{PS}} < E[T]^{\mathrm{FB}}$. Consider the discrete distribution

$$X = \begin{cases} 1 & \text{with probability } \frac{4}{5} - \varepsilon, \\ 6 & \text{with probability } \frac{1}{5} + \varepsilon. \end{cases}$$

It is easy to verify by simple calculation that $C^2 > 1$ for any $\varepsilon > 0$, but $E[T]^{\mathrm{PS}} < E[T]^{\mathrm{FB}}$ for small $\varepsilon > 0$.

Example 1 is counter to the hypothesis of Coffman and Denning, and moreover observe that this job size distribution belongs to a class where the hazard rate is neither always decreasing nor always increasing. [2]

Before proving Theorem 1, it is useful to describe the intuition behind the theorem. Intuitively, when the hazard rate of the service distribution is decreasing young jobs are likely to have small remaining times and old jobs are likely to have high remaining times. Thus, FB is mimicking SRPT by giving preference to jobs with small remaining times, and thus minimizing the number of jobs in the system, and equivalently the overall mean response time. Likewise, when the hazard rate of the service distribution is increasing, young jobs are likely to have larger remaining times, in which case FB maximizes the number of jobs in the system. In the case of constant hazard rate, a job's age

[2] Strictly speaking, the hazard rates are undefined as the distribution is discrete. However, we can approximate by a continuous distribution consisting of Gaussians at $x = 1$ and 6 with variance approaching 0. It is easy to see that Coffman and Denning's hypothesis does not hold for this continuous distribution either.

is independent of its remaining service time, in which case it would seem that FB scheduling should not improve upon PS.

The proof of Theorem 1 will rely on an alternative formulation of response times under FB as stated in the following Lemma.

Lemma 2.

$$E[T(x)]^{\mathrm{FB}} = \frac{\int_0^x (1-\rho_s)\,\mathrm{d}s}{(1-\rho_x)^2} \tag{1}$$

Proof. To derive this new expression we can combine terms and interchange integrals as follows:

$$\begin{aligned} E[T(x)]^{\mathrm{FB}} &= \frac{x}{1-\rho_x} + \frac{\lambda\int_0^x t\bar{F}(t)\,\mathrm{d}t}{(1-\rho_x)^2} \\ &= \frac{x - x\rho_x + \lambda\int_0^x t\bar{F}(t)\,\mathrm{d}t}{(1-\rho_x)^2} \\ &= \frac{x - \int_0^x \rho_s\,\mathrm{d}s}{(1-\rho_x)^2} \\ &= \frac{\int_0^x (1-\rho_s)\,\mathrm{d}s}{(1-\rho_x)^2}. \end{aligned}$$

The third equation follows from the second by observing that $\int_0^x \rho_s\,\mathrm{d}s = x\rho_x - \lambda\int_0^x t\bar{F}(t)\,\mathrm{d}t$. The final line then follows by writing x as $\int_0^x 1\,\mathrm{d}s$. □

Notice that Eq. (1) gives us a particularly simple form for the response time under FB. This simple form, combined with the Chebyshev Integral Inequality (stated below), will allow us to prove Theorem 1.

Theorem 2 (Chebyshev Integral Inequality). *Let $h(x)$ be a non-negative, integrable, increasing function on $[a,b]$.*

1. *Let $g(x)$ be a non-negative, integrable, increasing function on $[a,b]$.*
   Then, $(b-a)\int_a^b h(x)g(x)\mathrm{d}x \geqslant \int_a^b h(x)\mathrm{d}x \int_a^b g(x)\mathrm{d}x$.
2. *Let $g(x)$ be a non-negative, integrable, decreasing function on $[a,b]$.*
   Then, $(b-a)\int_a^b h(x)g(x)\mathrm{d}x \leqslant \int_a^b h(x)\mathrm{d}x \int_a^b g(x)\mathrm{d}x$.

Using Lemma 2 in combination with the Chebyshev Integral Inequality, we will now prove Theorem 1.

Proof of Theorem 1. We will start with the case where $\mu(x)$ is constant. Notice that this implies that the service distribution is exponential with some rate μ. Recall the Markov chain for the `M/M/1/FCFS` discipline, where the state corresponds to the number of jobs in the system, and for all $i>0$, the Markov chain moves from state i to state $i+1$ with rate λ, and the Markov chain moves from state i to state $i-1$ with rate μ. Notice that the `M/M/1/PS` discipline is represented by the exact same chain. The key point is that when the Markov chain is in state i, each of the i jobs in the system are served at rate μ/i, which, by superposition of exponential distributions, again results in a total transition rate of μ from state i to $i-1$. A similar argument can be made for `M/M/1/FB`. In state i, some number of jobs $j \leqslant i$ will share the processor evenly, and thus the total completion rate of j jobs receiving μ/j service is μ. In fact, any work conserving policy that does not depend on the job sizes can be represented by this same chain. Thus, the mean queue length and the mean sojourn time are also the same for all work conserving policies that do not make use of job size. It is also interesting to note that, since we did not make any assumptions about the arrival process in the above argument, the mean queue length and the mean sojourn time are also the same for all work conserving policies that do not make use of job size under any arbitrary sequence of arrivals.

We now prove the remaining two cases. Using Lemma 1, we can write the mean response time under FB as

$$\begin{aligned} E[T]^{\mathrm{FB}} &= \int_0^\infty E[T(x)]f(x)\,\mathrm{d}x \\ &= \int_0^\infty \frac{\int_0^x (1-\rho_s)\,\mathrm{d}s}{(1-\rho_x)^2} f(x)\,\mathrm{d}x \\ &= \int_0^\infty (1-\rho_s) \int_s^\infty \frac{f(x)}{(1-\rho_x)^2}\,\mathrm{d}x\,\mathrm{d}s. \end{aligned}$$

Finally, observing that $\mathrm{d}\rho_x/\mathrm{d}x = \lambda\bar{F}(x)$ and that $f(x) = \mu(x)\bar{F}(x)$, we get

$$E[T]^{\mathrm{FB}} = \frac{1}{\lambda}\int_0^\infty (1-\rho_s)\int_{\rho_s}^{\rho} \frac{\mu(x)}{(1-\rho_x)^2}\,\mathrm{d}\rho_x\,\mathrm{d}s. \tag{2}$$

At this point we will apply the Chebyshev Integral Inequality. First, we will deal with the case when $\mu(x)$ is increasing. Note that ρ_x is increasing and hence $1/(1-\rho_x)^2$ is increasing. Thus setting $h(x) = \mu(x)$,

$g(x) = 1/(1-\rho_x)^2$, $a = \rho_s$ and $b = \rho$ in Theorem 2, we have that

$$\int_{\rho_s}^{\rho} \frac{\mu(x)}{(1-\rho_x)^2}\,\mathrm{d}\rho_x \geqslant \frac{1}{\rho-\rho_s}\int_{\rho_s}^{\rho} \mu(x)\,\mathrm{d}\rho_x \int_{\rho_s}^{\rho} \frac{\mathrm{d}\rho_x}{(1-\rho_x)^2}.$$

Rewriting $\int_{\rho_s}^{\rho} \mu(x)\,\mathrm{d}\rho_x$ as $\int_s^{\infty} \lambda f(x)\,\mathrm{d}x$ we get,

$$\int_{\rho_s}^{\rho} \frac{\mu(x)}{(1-\rho_x)^2}\,\mathrm{d}\rho_x \geqslant \frac{1}{\rho-\rho_s}\int_{s}^{\infty} \lambda f(x)\,\mathrm{d}x \int_{\rho_s}^{\rho} \frac{\mathrm{d}\rho_x}{(1-\rho_x)^2}. \qquad (3)$$

Conversely, when $\mu(x)$ is decreasing, using an identical argument we have

$$\int_{\rho_s}^{\rho} \frac{\mu(x)}{(1-\rho_x)^2}\,\mathrm{d}\rho_x \leqslant \frac{1}{\rho-\rho_s}\int_{s}^{\infty} \lambda f(x)\,\mathrm{d}x \int_{\rho_s}^{\rho} \frac{\mathrm{d}\rho_x}{(1-\rho_x)^2}.$$

Now, we can simply evaluate the integral to obtain our bounds. We will consider only the case of increasing $\mu(x)$ (the decreasing case follows identically). Using Eqs. (2) and(3)

$$E[T]^{\mathrm{FB}} \geqslant \frac{1}{\lambda}\int_0^{\infty} \frac{1-\rho_s}{\rho-\rho_s}\int_s^{\infty} \lambda f(x)\,\mathrm{d}x \times \int_{\rho_s}^{\rho} \frac{\mathrm{d}\rho_x}{(1-\rho_x)^2}\,\mathrm{d}s$$

$$= \int_0^{\infty} \frac{1-\rho_s}{\rho-\rho_s}\bar{F}(s)\left(\frac{1}{1-\rho} - \frac{1}{1-\rho_s}\right)\mathrm{d}s$$

$$= \int_0^{\infty} \frac{1-\rho_s}{\rho-\rho_s}\bar{F}(s)\left(\frac{\rho-\rho_s}{(1-\rho)(1-\rho_s)}\right)\mathrm{d}s$$

$$= \int_0^{\infty} \frac{\bar{F}(s)}{1-\rho}\,\mathrm{d}s$$

$$= \frac{E[X]}{1-\rho} = E[T]^{\mathrm{PS}}$$

This completes the final two cases of the proof. □

References

[1] E.G. Coffman, P. Denning, Operating System Theory, Prentice-Hall, Englewood Cliffs, NJ, 1973.

[2] L. Kleinrock, Queueing Systems, Vol. II, Computer Applications, Wiley, New York, 1976.

[3] I. Rai, G. Urvoy-Keller, E. Biersack, Analysis of LAS scheduling for job size distributions with high variance, in: Proceedings of ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems, San Diego, CA, 2003.

[4] L.E. Schrage, L.W. Miller, The queue M/G/1 with the shortest remaining processing time discipline, Oper. Res. 14 (1966) 670–684.

[5] D. Stoyan, D.J. Daley, Comparison Methods for Queues and Other Stochastic Models, Wiley, New York, 1983.

[6] R.W. Wolff, Stochastic Modeling and the Theory of Queues, Prentice-Hall, Englewood Cliffs, NJ, 1989.