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

Adam Wierman [1] Nikhil Bansal [2] Mor Harchol-Balter [3]

September 2002

CMU-CS-02-177

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Two widely used scheduling policies used in the absence of knowledge of job sizes are Processor Sharing (PS) and Feedback (FB). While a lot of work has been done on comparing their performance on particular job size distributions, a general comparison has not been done. We compare the overall mean response time (a.k.a. sojourn time) of the PS and FB queues under an M/GI/1 system. We show that FB outperforms PS when the service distribution has a decreasing failure rate; but that when the failure rate of the service distribution is increasing, PS outperforms FB.

[1] Carnegie Mellon University, Computer Science Department. Email: acw@cs.cmu.edu.
[2] Carnegie Mellon University, Computer Science Department. Email: nikhil@cs.cmu.edu.
[3] Carnegie Mellon University, Computer Science Department. Email: harchol@cs.cmu.edu.
This work was supported by NSF Career Grant CCR-0133077 , Pittsburgh Digital Greenhouse Grant 01-1 and an IBM Research Fellowship.

Keywords: scheduling; M/G/1; FB; LAS; SET; feedback; least attained service; shortest elapsed time; PS; processor sharing; 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. This policy schedules the job having the smallest remaining size at all times; thus SRPT requires knowledge of the job sizes. 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 gets the processor to itself. If several jobs all have the least attained service (a.k.a age), 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 the expected response time of a job of size (a.k.a service requirement) x under policy P. Then the following classic results are well known (see [6] or [3] for a proof of these):

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

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

where

$$\rho_x = \lambda \left(\int_0^x t f(t)dt + x\bar{F}(x) \right) = \lambda \int_0^x \bar{F}(t)dt$$

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

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 [4] prove that for *any* service distribution $E[T]^{FB} \leq \frac{2-\rho}{2-2\rho} E[T]^{PS}$. In our paper we are concerned with understanding for *which* service distributions $E[T]^{FB} \leq E[T]^{PS}$. Coffman and Denning consider exactly this question and hypothesize the following relation [1, Pages 188-89]:

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

where $C^2 \stackrel{\text{def}}{=} \frac{Var(X)}{E[X]^2}$ 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.

[1] Note that FB is sometimes referred to by two other names: Least-Attained-Service (LAS) and Shortest-Elapsed-Time (SET).

Theorem 1 *Let $\mu(x) \stackrel{\text{def}}{=} f(x)/\bar{F}(x)$ 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 $E[T]^{FB} \leq E[T]^{PS}$.*
2. *If $\mu(x)$ is constant $E[T]^{FB} = E[T]^{PS}$.*
3. *If $\mu(x)$ is increasing $E[T]^{FB} \geq E[T]^{PS}$.*

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

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

Notice that Theorem 1 does not say anything about distributions which are both strictly increasing for some x and strictly decreasing for other values of x. Our example below shows that this cannot be avoided.

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

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

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

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

Lemma 2

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

[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 $x = 6$ with variance approaching 0. It is easy to see that Coffman and Denning's hypothesis does not hold for this continuous distribution either.

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

$$\begin{aligned} E[T(x)]^{FB} &= \frac{\lambda \int_0^x t\bar{F}(t)dt}{(1-\rho_x)^2} + \frac{x}{1-\rho_x} \\ &= \frac{x + \lambda \int_0^x t\bar{F}(t)dt - \lambda x \int_0^x \bar{F}(t)dt}{(1-\rho_x)^2} \\ &= \frac{x - \lambda \int_0^x (x-t)\bar{F}(t)dt}{(1-\rho_x)^2} \\ &= \frac{x - \lambda \int_0^x \int_t^x ds\bar{F}(t)dt}{(1-\rho_x)^2} \\ &= \frac{x - \lambda \int_0^x \int_0^s \bar{F}(t)dtds}{(1-\rho_x)^2} \\ &= \frac{\int_0^x (1-\rho_s)ds}{(1-\rho_x)^2} \end{aligned}$$

The second step follows from the first by observing that $\rho_x = \lambda \int_0^x \bar{F}(t)dt$. The fifth step follows from the fourth by an interchange of integrals. Finally, the last step follows from the fifth by writing x as $\int_0^x 1ds$ and noting that $\lambda \int_0^s \bar{F}(t)dt = \rho_s$. □

Notice that equation 1 gives us a particularly simple form for the response time under FB. We will use this to prove our main result (Theorem 1). Before our main result though we need the Chebyshev Integral Inequality [2], which states the following:

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, $(b-a)\int_a^b h(x)g(x)dx \geq \int_a^b h(x)dx \int_a^b g(x)dx$.
2. *Let $g(x)$ be a non-negative, integrable, decreasing function on $[a, b]$.*
   Then, $(b-a)\int_a^b h(x)g(x)dx \leq \int_a^b h(x)dx \int_a^b g(x)dx$.

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. First recall the Markov chain for the `M/M/1/FCFS` discipline, where the state corresponds to the number of jobs in the system, and state $i > 0$ transitions to $i+1$ with rate λ and to state $i-1$ with rate μ. Notice that the `M/M/1/PS` discipline is represented by the exact same chain. When in state i the arrival rate is equivalent, and each of the i jobs is served at a rate of μ/i. Thus, by superposition of exponential distributions, the transition rate from i to $i-1$ is again μ. A similar argument can be made for `M/M/1/FB`. In state i, some number of jobs $j \leq i$ 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. It is also interesting to note that, since we did not make any assumptions about the arrival process, this result holds for 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

$$\begin{aligned} E[T]^{FB} &= \int_0^\infty E[T(x)]f(x)dx \\ &= \int_0^\infty \frac{\int_0^x (1-\rho_s)ds}{(1-\rho_x)^2} f(x)dx \\ &= \int_0^\infty (1-\rho_s)\int_s^\infty \frac{f(x)}{(1-\rho_x)^2}dxds \end{aligned}$$

The final step follows from the second by an interchange of integrals.

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

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

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

$$\int_{\rho_s}^{\rho} \frac{\mu(x)}{(1-\rho_x)^2}d\rho_x \geq \frac{1}{\rho-\rho_s}\int_{\rho_s}^{\rho}\mu(x)d\rho_x \int_{\rho_s}^{\rho}\frac{d\rho_x}{(1-\rho_x)^2}$$

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

$$\int_{\rho_s}^{\rho} \frac{\mu(x)}{(1-\rho_x)^2}d\rho_x \geq \frac{1}{\rho-\rho_s}\int_s^\infty \lambda f(x)dx \int_{\rho_s}^{\rho}\frac{d\rho_x}{(1-\rho_x)^2} \tag{3}$$

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

$$\int_{\rho_s}^{\rho} \frac{\mu(x)}{(1-\rho_x)^2}d\rho_x \quad\leq\quad \frac{1}{\rho-\rho_s}\int_s^\infty \lambda f(x)dx \int_{\rho_s}^{\rho}\frac{d\rho_x}{(1-\rho_x)^2}$$

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 Equations 2 and 3

$$\begin{aligned} E[T(x)]^{FB} &\geq \frac{1}{\lambda}\int_0^\infty \frac{1-\rho_s}{\rho-\rho_s}\int_s^\infty \lambda f(x)dx\int_{\rho_s}^{\rho}\frac{d\rho_x}{(1-\rho_x)^2}ds \\ &= \int_0^\infty \frac{1-\rho_s}{\rho-\rho_s}\bar{F}(s)\left(\frac{1}{1-\rho}-\frac{1}{1-\rho_s}\right)ds \\ &= \int_0^\infty \frac{1-\rho_s}{\rho-\rho_s}\bar{F}(s)\left(\frac{\rho-\rho_s}{(1-\rho)(1-\rho_s)}\right)ds \\ &= \int_0^\infty \frac{\bar{F}(s)}{1-\rho}ds \\ &= \frac{E[X]}{1-\rho} = E[T]^{PS} \end{aligned}$$

Which completes the final two cases of the proof. $\square$

References

[1] E.G. Coffman and P. Denning. *Operating System Theory*. Prentice Hall, 1973.

[2] I.S. Gradshteyn and I.M. Ryzhik. *Tables of Integrals, Series, and Products*. Academic Press, 2000.

[3] L. Kleinrock. *Queueing Systems*, volume II. Computer Applications. John Wiley & Sons, 1976.

[4] I. Rai, G. Urvoy-Keller, and E. Biersack. FB: An efficient scheduling policy for edge routers to speedup the internet access. Unpublished manuscript.

[5] D. Stoyan and D.J. Daley. *Comparison Methods for Queues and Other Stochastic Models*. John Wiley & Sons, 1983.

[6] Ronald W. Wolff. *Stochastic Modeling and the Theory of Queues*. Prentice Hall, 1989.