Mean delay analysis of Multi Level Processor Sharing disciplines

S. Aalto[†] and U. Ayesta[‡]

[†] Helsinki University of Technology, Finland. e-mail: samuli.aalto@tkk.fi
[‡] CWI, The Netherlands. e-mail: urtzi.ayesta@cwi.nl

***Abstract*— Multilevel Processor-Sharing (MLPS) scheduling disciplines permit to model a wide variety of non-anticipating scheduling disciplines. Such disciplines 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. In this paper, we compare the mean delay in an M/G/1 queue among MLPS disciplines under the assumption that the service time distribution belongs to class Decreasing Hazard Rate (DHR). We are able to prove that, given an MLPS discipline, the mean delay is reduced whenever a level is added by splitting an existing one in several cases. The exceptions concern splitting the upper levels with PS internal discipline. Our numerical examples, however, indicate that the level splitting be advantageous even in these cases. Furthermore, we characterize the effect on the mean delay of changing internal disciplines within levels. By numerical means we demonstrate that the mean delay of an MLPS discipline can get close to the minimum optimal delay with just a few levels. As the number of levels increases in an MLPS discipline, the MLPS queue mimics closer and closer the behavior of a Foreground-Background queue, which is known to minimize the mean delay among all disciplines. Thus, our result provides a constructive way to demonstrate the optimality of FB.**

I. Introduction

Multilevel Processor-Sharing (MLPS) scheduling disciplines introduced by L. Kleinrock in the early 1970's [1] permit to model a wide variety of non-anticipating scheduling disciplines. A discipline is non-anticipating when the scheduler does not know the remaining service time of jobs. Such disciplines 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 [2], [3], [4], [5].

An MLPS scheduling 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^π is applied. The internal disciplines may vary in the set $\{\mathrm{FB}, \mathrm{PS}, \mathrm{FCFS}\}$, where FB refers to the Foreground-Background discipline, which gives priority to the job with the least attained service, PS to the Processor-Sharing discipline, which shares the service capacity evenly among all jobs, and FCFS to the ordinary First-Come-First-Served discipline. The FB discipline is also known as LAS (Least-Attained-Service).

Yashkov [6] has proven that, for M/G/1 queues, FB minimizes the mean delay, i.e., the mean sojourn time, among work-conserving and non-anticipating disciplines whenever the service time distribution is of type DHR (Decreasing Hazard Rate). Righter and Shanthikumar [7] proved that FB minimizes the queue length in the stochastic sense.

The PS discipline has been proposed as an appropriate model for the bandwidth sharing among (non-prioritized) TCP flows in a bottleneck router [8], [9], [10], [11]. According to a similar reasoning, PS is a relevant internal discipline for the MLPS disciplines that model the bandwidth sharing among prioritized TCP flows [2], [4]. On the other hand, flow sizes in the Internet have been modelled, e.g., by Pareto and hyperexponential distributions [12], [13] satisfying the DHR condition.

Even though the apparent desirable properties of FB, its deployment does not seem to be a simple task. For example, the deployment of FB would require to classify jobs into an infinite number of classes. On the contrary, an MLPS discipline with N thresholds would require just to maintain $N+1$ classes.

Regarding the mean delay in M/G/1 queues, the following results have been derived for the MLPS disciplines whose internal disciplines vary in the set $\{\mathrm{FB}, \mathrm{PS}\}$ (with FCFS excluded). Aalto *et al.* [15] proved that such MLPS disciplines with just *two* levels are better than PS with respect to the mean delay whenever the hazard rate of the service time distribution is decreasing, and vice versa if the hazard rate is increasing and bounded. In [16], similar results were found when comparing an MLPS discipline with *any* number of levels to the PS discipline. All in all, such MLPS disciplines look like a reasonable compromise between PS and FB, having a smaller overall mean delay than PS and a fairer conditional mean delay than FB.

From the previous analysis an important question remained widely open: Given any two MLPS disciplines, which one is better in the mean delay sense? Our main result states that given an MLPS discipline, the mean delay is reduced (under the DHR condition) if a level is added by splitting an existing one unless it is an upper level and the internal discipline is PS. Furthermore, we show that given an MLPS discipline, the mean delay is reduced (under the DHR condition) if an internal discipline is changed from FCFS to PS (or from PS to FB). These two results define a natural partial order among the MLPS disciplines. On the other hand, as the number of levels grows to infinity, the MLPS discipline mimics the behavior of the FB discipline. Thus, our results provide a constructive way to demonstrate the optimality of FB.

By numerical means we quantify the reduction on the mean delay after the addition of levels. We show that the mean delay of an MLPS discipline can get close to the minimum feasible delay (FB) with just a few levels. Our numerical examples indicate that the level splitting be advantageous even if the splitted level is an upper one provided with the PS internal discipline.

The rest of the paper is organized as follows. In Section II, after the introduction of the notation used, we present our new results related to the mean delay comparison among MLPS disciplines. The results are proved in Sections III and IV. In Section III we analyze the effect of changing an internal discipline, and in Section IV we consider the effect of splitting a level for an MLPS discipline. Illustrative numerical examples are given in Section V. Section VI concludes the paper.

II. New results

In this section we present the notation used and the new results of the paper. Some of the results are valid for each sample path, while the others are related to mean values only. The main new results are given in Theorems 1, 2 and 3 at the end of this section.

A. MLPS disciplines

We denote by MLPS the family of MLPS disciplines π for which $D_n^\pi \in \{\mathrm{FB}, \mathrm{PS}, \mathrm{FCFS}\}$ for all n, and by MLPS^* the family of MLPS disciplines π for which $D_n^\pi \in \{\mathrm{FB}, \mathrm{PS}\}$ for all n. Furthermore, we denote by $(N+1)$FCFS the family of MLPS disciplines π with $N+1$ levels for which $D_n^\pi = \mathrm{FCFS}$ for all n. Thus, for example, 1FCFS refers to the FCFS discipline alone, 2FCFS to the FCFS+FCFS disciplines, 3FCFS to the FCFS+FCFS+FCFS disciplines etc. All these disciplines belong to the class Π of work-conserving and non-anticipating disciplines, which do not idle when there are jobs waiting and neither use they any information about the remaining service times of jobs.

We note that any number of contiguous FB levels can always be considered as a single (but larger) FB level. So there is no need to split an FB level.

Among the disciplines $\{\mathrm{FB}, \mathrm{PS}, \mathrm{FCFS}\}$, we define the following order relation:

$$\mathrm{FB} \prec \mathrm{PS} \prec \mathrm{FCFS}.$$

In addition, we denote $D \preceq D'$ if and only if $D = D'$ or $D \prec D'$.

B. Sample path results

Consider a single server queueing system starting empty at time $t = 0$ and obeying a scheduling discipline π. We assume that the jobs arrive one at a time.

Let A_i denote the arrival time of job i, S_i its service time, and $X_i^\pi(t)$ its attained service at time t. Let $\mathcal{A}(t)$ denote the set of jobs arrived until time t,

$$\mathcal{A}(t) = \{i : A_i \le t\},$$

and $A(t) = |\mathcal{A}(t)|$. Let $\mathcal{N}^\pi(t)$ denote the set of jobs in the system at time t,

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

and $N^\pi(t) = |\mathcal{N}^\pi(t)|$. Furthermore, for all $x \ge 0$, let $\mathcal{N}_x^\pi(t)$ denote the set of those jobs in the system whose attained service is less than x,

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

where $(S_i \wedge x) = \min\{S_i, x\}$, and $N_x^\pi(t) = |\mathcal{N}_x^\pi(t)|$.

Let $U_x^\pi(t)$ denote the unfinished truncated work with truncation threshold x at time t,

$$U_x^\pi(t) = \sum_{i \in \mathcal{N}_x^\pi(t)} ((S_i \wedge x) - X_i^\pi(t)). \tag{1}$$

An alternative expression is as follows:

$$U_x^\pi(t) = \sum_{i=1}^{A(t)} (S_i \wedge x) - \int_0^t \sigma_x^\pi(u)\, du, \qquad (2)$$

where $\sigma_x^\pi(t)$ refers to the total rate at which the jobs with attained service less than x are served at time t. The limiting value $U_\infty^\pi(t)$ is the ordinary unfinished work, which is the same for all work conserving disciplines.

Proposition 1: Let $\pi \in$ MLPS with thresholds $\{a_1, \ldots, a_N\}$ and $\pi' \in$ MLPS with thresholds $\{a'_1, \ldots, a'_{N'}\}$. Assume that there exist $n \in \{1, \ldots, N+1\}$ and $n' \in \{1, \ldots, N'+1\}$ such that $a_{n-1} = a'_{n'-1}$ and $a_n = a'_{n'}$.

(i) If $D_n^\pi = D_{n'}^{\pi'}$, then $U_x^\pi(t) = U_x^{\pi'}(t)$ for all $a_{n-1} \le x \le a_n$ and $t \ge 0$.

(ii) If $D_n^\pi \preceq D_{n'}^{\pi'}$, then $U_x^\pi(t) \le U_x^{\pi'}(t)$ for all $a_{n-1} \le x \le a_n$ and $t \ge 0$.

This is a fundamental locality result for MLPS disciplines saying that the unfinished truncated work within a level depends only on the internal discipline of that level (but not on the other internal disciplines). In [15, Prop. 8], the result was proved for class MLPS*. The generalization to all MLPS disciplines is proved below in Section III.

Proposition 2: Let $\pi, \pi' \in$ MLPS with the same thresholds $\{a_1, \ldots, a_N\}$ such that $D_n^\pi \preceq D_n^{\pi'}$ for all $n \in \{1, \ldots, N+1\}$. Then $U_x^\pi(t) \le U_x^{\pi'}(t)$ for all $x \ge 0$ and $t \ge 0$.

Proof: This is an immediate consequence of the locality result presented in Proposition 1. ■

This is the first step in ordering the MLPS disciplines. It allows us to compare different MLPS disciplines with the same thresholds.

C. Mean value results

Consider an M/G/1 queue obeying a scheduling discipline $\pi \in \Pi$. Let λ denote the arrival rate and S the service time of a job. We assume that $E[S] < \infty$ and $\rho = \lambda E[S] < 1$. Furthermore, we assume that the service time distribution is continuous with the corresponding density function denoted by $f(x)$. Let $F(x) = \int_0^x f(y)\, dy$ and $\overline{F}(x) = 1 - F(x)$. The corresponding hazard rate function is denoted by $h(x) = f(x)/\overline{F}(x)$.

Let U_x^π denote the unfinished truncated work with truncation threshold x and $T^\pi(y)$ the delay of a job with service time y. By [1, Eq. (4.60)],

$$\overline{U}_x^\pi = \lambda \int_0^x \overline{T}^\pi(y) \overline{F}(y)\, dy, \qquad (3)$$

where $\overline{U}_x^\pi = E[U_x^\pi]$ and $\overline{T}^\pi(y) = E[T^\pi(y)]$.

Taking the derivative, and integrating with respect to the density distribution we obtain that mean delay can be expressed as:

$$\overline{T}^\pi = \int_0^\infty \overline{T}^\pi(x) f(x)\, dx = \frac{1}{\lambda} \int_0^\infty (\overline{U}_x^\pi)' h(x)\, dx. \qquad (4)$$

where $\overline{T}^\pi = E[T^\pi]$ and $(\overline{U}_x^\pi)' = \frac{\partial}{\partial x} \overline{U}_x^\pi$.

The following result was proven in [15, Prop. 1].

Proposition 3: Let $\pi, \pi' \in \Pi$. If $\overline{U}_x^\pi \le \overline{U}_x^{\pi'}$ for all $x \ge 0$ and the hazard rate $h(x)$ is decreasing, then

$$\overline{T}^\pi \le \overline{T}^{\pi'}.$$

Hence, Proposition 3 allows to compare the mean delay performance of different scheduling disciplines by comparing their respective unfinished truncated work.

Proposition 4: Let $N \ge 1$, $n \in \{1, \ldots, N\}$, $\pi \in$ MLPS with thresholds $\{a_1, \ldots, a_N\}$, and $\pi' \in$ MLPS with thresholds $\{a_1, \ldots, a_{n-1}, a_{n+1}, \ldots, a_N\}$ such that $D_n^{\pi'} = D_n^\pi = D_{n+1}^\pi =$ FCFS and

$$D_m^{\pi'} = \begin{cases} D_m^\pi, & m = 1, \ldots, n-1, \\ D_{m+1}^\pi, & m = n+1, \ldots, N. \end{cases}$$

Then, $\overline{U}_x^\pi \le \overline{U}_x^{\pi'}$ for all $x \ge 0$.

This new result shows that, with respect to the mean unfinished truncated work, it is beneficial to split one level with FCFS as an internal discipline to two contiguous levels with the FCFS discipline applied at both levels. The proof is given in Subsection IV-C.

Proposition 5: Let $N \ge 1$, $\pi \in$ MLPS with thresholds $\{a_1, \ldots, a_N\}$, and $\pi' \in$ MLPS with thresholds $\{a_2, \ldots, a_N\}$ such that $D_1^{\pi'} = D_1^\pi = D_2^\pi =$ PS and

$$D_m^{\pi'} = D_{m+1}^\pi \quad \text{for } m = 2, \ldots, N.$$

Then, $\overline{U}_x^\pi \le \overline{U}_x^{\pi'}$ for all $x \ge 0$.

This is a counterpart of the previous result concerning a level with PS as an internal discipline instead of FCFS. The proof is given in Subsection IV-D. Unfortunately, we were not able to prove the result when the splitted level is an upper one ($n > 1$).

The following three theorems contain the main results of the present paper. They define a partial order among the MLPS disciplines (under the DHR condition) that allows to conclude whether an MLPS discipline is better or worse than another MLPS discipline with respect to the mean delay in certain cases.

Theorem 1: Let $\pi, \pi' \in$ MLPS with the same thresholds $\{a_1, \ldots, a_N\}$ such that $D_n^\pi \preceq D_n^{\pi'}$ for all $n \in \{1, \ldots, N+1\}$. If the hazard rate $h(x)$ is decreasing, then $\overline{T}^\pi \le \overline{T}^{\pi'}$.

Proof: Due to Proposition 3, this new result is an immediate consequence of Proposition 2. ■

Theorem 2: Let $N \geq 1$, $n \in \{1,\dots,N\}$, $\pi \in$ MLPS with thresholds $\{a_1,\dots,a_N\}$, and $\pi' \in$ MLPS with thresholds $\{a_1,\dots,a_{n-1},a_{n+1},\dots,a_N\}$ such that $D_n^{\pi'} = D_n^{\pi} = D_{n+1}^{\pi} =$ FCFS and

$$D_m^{\pi'} = \begin{cases} D_m^{\pi}, & m = 1,\dots,n-1, \\ D_{m+1}^{\pi}, & m = n+1,\dots,N. \end{cases}$$

If the hazard rate $h(x)$ is decreasing, then $\overline{T}^{\pi} \leq \overline{T}^{\pi'}$.

Proof: Due to Proposition 3, this new result is an immediate consequence of Proposition 4. ■

Theorem 3: Let $N \geq 1$, $\pi \in$ MLPS with thresholds $\{a_1,\dots,a_N\}$, and $\pi' \in$ MLPS with thresholds $\{a_2,\dots,a_N\}$ such that $D_1^{\pi'} = D_1^{\pi} = D_2^{\pi} =$ PS and

$$D_m^{\pi'} = D_{m+1}^{\pi} \quad \text{for } m = 2,\dots,N.$$

If the hazard rate $h(x)$ is decreasing, then $\overline{T}^{\pi} \leq \overline{T}^{\pi'}$.

Proof: Due to Proposition 3, this new result is an immediate consequence of Proposition 5. ■

Consider, for example, disciplines π and π' defined as follows: π has three thresholds (10, 20, and 30) and applies PS as an internal discipline at all the four levels, while π' has only two thresholds (10 and 20) and applies FCFS as an internal discipline at all the three levels. By Theorem 2, π' is worse than π'' which has three thresholds (10, 20, and 30) and applies FCFS as an internal discipline at all the four levels. On the other hand, π'' is worse than π by Theorem 1. Thus, π' is worse than π with respect to the mean delay.

Finally we note that all the mean delay inequalities in Theorems 1, 2 and 3 would be reversed if we had assumed that the hazard rate $h(x)$ is increasing and bounded (instead of being decreasing). This can be readily seen since in this case the mean delay inequality in Proposition 3 is reversed as shown in [15, Prop. 2].

III. Changing internal disciplines

In this section we prove the sample path result presented in Proposition 1. We consider a single server queueing system starting empty at time $t = 0$ and assume that the jobs arrive one at a time.

Proposition 1 is a generalization of [15, Prop. 8(i)]. Regarding part (i) of Proposition 1, the proof presented in [15, Prop. 8(i)] is clearly valid as such for all MLPS disciplines. For part (ii) of Proposition 1, we need only consider the case of $D_n^{\pi} =$ PS and $D_{n'}^{\pi'} =$ FCFS.

Proposition 6: Let $\pi \in$ MLPS with thresholds $\{a_1,\dots,a_N\}$ and $\pi' \in$ MLPS with thresholds $\{a'_1,\dots,a'_{N'}\}$. Assume that there exist $n \in \{1,\dots,N+1\}$ and $n' \in \{1,\dots,N'+1\}$ such that $a_{n-1} = a'_{n'-1}$, $a_n = a'_{n'}$, $D_n^{\pi} =$ PS, and $D_{n'}^{\pi'} =$ FCFS. Then $U_x^{\pi}(t) \leq U_x^{\pi'}(t)$ for all $a_{n-1} \leq x \leq a_n$ and $t \geq 0$.

Proof: Let $a_{n-1} \leq x \leq a_n$. We prove the claim by induction with respect to the arrival epochs A_k.

1° During the interval $[0, A_1)$ both systems are empty. Thus the claim is trivially true for all $t < A_1$.

2° Let $k \in \{1,2,\dots\}$, and assume that the claim is true for all $t < A_k$. We will show that it is also true in the interval $[A_k, A_{k+1})$.

We divide the interval $[A_k, A_{k+1})$ into three consecutive periods I_1, I_2, and I_3, with the following starting (b) and ending (e) points:

$$\begin{aligned} I_1^{\mathrm{b}} &= A_k, & I_1^{\mathrm{e}} &= \sup\{I_1^{\mathrm{b}} < t \leq A_{k+1} \mid N_{a_{n-1}}^{\pi}(t) > 0\}, \\ I_2^{\mathrm{b}} &= I_1^{\mathrm{e}}, & I_2^{\mathrm{e}} &= \sup\{I_2^{\mathrm{b}} < t \leq A_{k+1} \mid N_{a_n}^{\pi}(t) > 0\}, \\ I_3^{\mathrm{b}} &= I_2^{\mathrm{e}}, & I_3^{\mathrm{e}} &= A_{k+1}. \end{aligned}$$

During the interval I_1 both systems give service only to those customers whose attained service is less than a_{n-1}. During interval the I_2 there are no longer such customers in either system, and both systems give service to those customers whose attained service is at least a_{n-1} but less than a_n. Finally, in the interval I_3, there are not any longer jobs with attained service less than a_n in either system. Note that I_1 is always of positive length, whereas I_2 and I_3 may vanish. The three intervals I_1 – I_3 are considered in 2.1° – 2.3°, respectively.

2.1° Consider the interval I_1. Since $x \geq a_{n-1}$ and, during this interval, strict priority is given (in both systems) to those customers with attained service time less than a_{n-1}, we have, for all $t \in I_1$,

$$\begin{aligned} U_x^{\pi}(t) &= U_x^{\pi}((A_k)^-) + (S_k \wedge x) - t + A_k \\ &\leq U_x^{\pi'}((A_k)^-) + (S_k \wedge x) - t + A_k \\ &= U_x^{\pi'}(t), \end{aligned}$$

where the inequality is due to the induction assumption. This is enough if the interval I_1 ends at time A_{k+1} when a new customer arrives. Otherwise we have to consider the interval I_2, too.

2.2° Consider the interval I_2. During this interval, strict priority is given (in both systems) to those customers with attained service time at least a_{n-1} but less than a_n. Let $t \in I_2$ and denote $\mathcal{M}(t) = \mathcal{N}_{a_n}^{\pi}(t) \cup \mathcal{N}_{a_n}^{\pi'}(t)$. Thus, $\mathcal{M}(t)$ comprises of the customers that are priority customers, at least, in one of the systems at time t. Since a_{n-1} and a_n are level thresholds for both π and π', we have, for all $i \in \mathcal{M}(t)$,

$$a_{n-1} \leq X_i^{\pi}(t) \leq a_n, \quad a_{n-1} \leq X_i^{\pi'}(t) \leq a_n.$$

Now,

$$U_x^{\pi}(t) = \sum_{i\in\mathcal{M}(t)} (S_i \wedge x) - \sum_{i\in\mathcal{M}(t)} (X_i^{\pi}(t) \wedge x), \quad (5)$$

$$U_x^{\pi'}(t) = \sum_{i\in\mathcal{M}(t)} (S_i \wedge x) - \sum_{i\in\mathcal{M}(t)} (X_i^{\pi'}(t) \wedge x). \quad (6)$$

Note that the first sum is the same for both disciplines. For $x = a_n$, we get

$$\begin{aligned} U_{a_n}^{\pi}(t) &= \sum_{i\in\mathcal{M}(t)} (S_i \wedge a_n) - \sum_{i\in\mathcal{M}(t)} (X_i^{\pi}(t) \wedge a_n) \\ &= \sum_{i\in\mathcal{M}(t)} (S_i \wedge a_n) - \sum_{i\in\mathcal{M}(t)} X_i^{\pi}(t), \\ U_{a_n}^{\pi'}(t) &= \sum_{i\in\mathcal{M}(t)} (S_i \wedge a_n) - \sum_{i\in\mathcal{M}(t)} (X_i^{\pi'}(t) \wedge a_n) \\ &= \sum_{i\in\mathcal{M}(t)} (S_i \wedge a_n) - \sum_{i\in\mathcal{M}(t)} X_i^{\pi'}(t). \end{aligned}$$

Since a_n is a level threshold for both disciplines, the unfinished truncated work with truncation threshold a_n is the same in both systems. Thus,

$$\sum_{i\in\mathcal{M}(t)} X_i^{\pi}(t) = \sum_{i\in\mathcal{M}(t)} X_i^{\pi'}(t). \quad (7)$$

Let then $I(t)$ denote the index of the customer who is served in the π' system at time t. Because of the internal FCFS discipline of the π' system, we observe that, for all $i \in \mathcal{M}(t)$ such that $i < I(t)$,

$$X_i^{\pi}(t) < (S_i \wedge a_n) = X_i^{\pi'}(t), \quad (8)$$

and, for all $i \in \mathcal{M}(t)$ such that $i > I(t)$,

$$X_i^{\pi}(t) \geq a_{n-1} = X_i^{\pi'}(t). \quad (9)$$

Furthermore, due to the internal PS discipline in the π system, we have, for all $i \in \mathcal{N}_{a_n}^{\pi}(t)$ and $j \in \mathcal{M}(t)$ such that $i < j$,

$$X_i^{\pi}(t) \geq X_j^{\pi}(t). \quad (10)$$

The next step is to utilize Lemma 1 presented in Appendix A. Consider customers $i \in \mathcal{M}(t)$. Let $N = |\mathcal{M}(t)|$, and re-index these customers in their arrival order from 1 to N. Let $r(i)$ denote the new index of customer i. Furthermore, let $a_{r(i)} = X_i^{\pi}(t)$ and $b_{r(i)} = X_i^{\pi'}(t)$ for all $i \in \mathcal{M}(t)$. In addition, let

$$m = \begin{cases} r(I(t)), & \text{if } X_{I(t)}^{\pi}(t) < X_{I(t)}^{\pi'}(t), \\ r(I(t)) - 1, & \text{if } X_{I(t)}^{\pi}(t) \geq X_{I(t)}^{\pi'}(t). \end{cases}$$

With the results given above in (7), (8), (9), and (10), it is easy to check that all the assumptions of Lemma 1 are valid with these choices. Now, by applying Lemma 1, we conclude that

$$\sum_{i\in\mathcal{M}(t)} (X_i^{\pi}(t) \wedge x) \geq \sum_{i\in\mathcal{M}(t)} (X_i^{\pi'}(t) \wedge x).$$

Thus, by (5) and (6), we have $U_x^{\pi}(t) \leq U_x^{\pi'}(t)$ for any $t \in I_2$. This is enough if the interval I_2 ends at time A_{k+1} when a new job arrives. Otherwise we have to consider the final interval I_3, too.

2.3° Consider finally the interval I_3, in which there are not any more customers with attained service less than a_n in either system. Since $x \leq a_n$, we deduce that, for all $t \in I_3$,

$$U_x^{\pi}(t) = U_x^{\pi'}(t) = 0.$$

This completes the proof. ■

IV. Splitting levels

In this section we consider the effect of splitting levels. We start with some preliminaries, and then present the proofs of Propositions 4 and 5 in Subsections IV-C and IV-D, respectively.

A. Preliminary results

The first preliminary result is valid for all sample paths. Thus, we consider a single server queueing system starting empty at time $t = 0$ and assume that the jobs arrive one at a time. Let $U_x^{\pi}(t; S \wedge b)$ denote the unfinished truncated work with truncation threshold x at time t in a modified system where the original service times $S_1, S_2, \ldots$ are replaced by their truncated versions $S_1 \wedge b, S_2 \wedge b, \ldots$.

Proposition 7: Let $N \geq 1$, $n \in \{1, \ldots, N\}$, $\pi \in$ MLPS with thresholds $\{a_1, \ldots, a_N\}$, and $\pi' \in$ MLPS with thresholds $\{a_1, \ldots, a_{n-1}\}$ such that

$$D_m^{\pi'} = D_m^{\pi} \quad \text{for } m = 1, \ldots, n.$$

Then, $U_x^{\pi}(t) = U_x^{\pi}(t; S \wedge a_n) = U_x^{\pi'}(t; S \wedge a_n)$ for all $x \leq a_n$ and $t \geq 0$.

Proof: Let $x \leq a_n$ and $t \geq 0$. Since the three systems follow the same rules as regards the jobs with attained service time less than a_n (which is a level threshold in the original system and the service time truncation threshold in the other systems), we surely have $\mathcal{N}_x^{\pi}(t) = \mathcal{N}_x^{\pi}(t; S \wedge a_n) = \mathcal{N}_x^{\pi'}(t; S \wedge a_n)$, and $X_i^{\pi}(t) = X_i^{\pi}(t; S \wedge a_n) = X_i^{\pi'}(t; S \wedge a_n)$ for all $i \in \mathcal{N}_x^{\pi}(t)$. The claim follows now from (1). ■

The second preliminary result concerns the mean unfinished truncated work in the system with non-truncated service times $S_1, S_2, \ldots$. Thus, we consider here an

M/G/1 queue with $\rho < 1$. Recall that $\overline{T}^{\pi}(x)$ refers to the conditional mean delay of a job with service time x.

Proposition 8: Let $\pi, \pi' \in \Pi$ and $a < b$. Assume that

(i) $\overline{U}_a^{\pi} \le \overline{U}_a^{\pi'}$, $\overline{U}_b^{\pi} \le \overline{U}_b^{\pi'}$, and

(ii) $\overline{T}^{\pi}(x) - \overline{T}^{\pi'}(x)$ is non-decreasing for all $a < x \le b$.

Then $\overline{U}_x^{\pi} \le \overline{U}_x^{\pi'}$ for all $a \le x \le b$.

Proof: Assume that there is $a < x < b$ such that $\overline{U}_x^{\pi} > \overline{U}_x^{\pi'}$. Consider now what happens if $\overline{T}^{\pi}(x) \le \overline{T}^{\pi'}(x)$. Due to (ii), this would imply that $\overline{T}^{\pi}(y) \le \overline{T}^{\pi'}(y)$ for all $a < y < x$. Then, by (3),

$$\begin{aligned}\overline{U}_x^{\pi} &= \overline{U}_a^{\pi} + \lambda \int_a^x \overline{T}^{\pi}(y)\overline{F}(y)\,dy \\ &\le \overline{U}_a^{\pi'} + \lambda \int_a^x \overline{T}^{\pi'}(y)\overline{F}(y)\,dy \\ &= \overline{U}_x^{\pi'},\end{aligned}$$

which is impossible. Thus, we conclude that $\overline{T}^{\pi}(x) > \overline{T}^{\pi'}(x)$. Due to (ii), $\overline{T}^{\pi}(y) > \overline{T}^{\pi'}(y)$ for all $x < y < b$. Then, by (3),

$$\begin{aligned}\overline{U}_b^{\pi} &= \overline{U}_x^{\pi} + \lambda \int_x^b \overline{T}^{\pi}(y)\overline{F}(y)\,dy \\ &> \overline{U}_x^{\pi'} + \lambda \int_x^b \overline{T}^{\pi'}(y)\overline{F}(y)\,dy \\ &= \overline{U}_b^{\pi'},\end{aligned}$$

which is impossible. Thus, we conclude that $\overline{U}_x^{\pi} \le \overline{U}_x^{\pi'}$ for all $a \le x \le b$. ■

Hence, Proposition 8 allows to compare the mean unfinished truncated work of different scheduling disciplines by comparing their respective conditional mean delay within an interval.

B. Conditional mean delay for MLPS disciplines

In this subsection we recall the properties of the conditional mean delay for MLPS disciplines in an M/G/1 queue found, e.g., in [1].

Let $x \ge 0$, and replace, for a while, the service times S by their truncated versions $S \wedge x = \min\{S, x\}$. It is easy to see that

$$E[S\wedge x] = \int_0^x \overline{F}(z)\,dz, \quad E[(S\wedge x)^2] = 2\int_0^x z\overline{F}(z)\,dz.$$

Furthermore, let $\rho_x = \lambda E[(S \wedge x)]$ denote the truncated load. Clearly, $\rho_x \le \rho < 1$ for all x.

The mean workload (i.e., unfinished work) for a work conserving M/G/1 queue with truncated service times is, by the Pollaczek-Khinchin formula,

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

Of course, when $x \to \infty$, we get the ordinary Pollaczek-Khinchin formula,

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

It is worth to mention that the unfinished work W_x for an M/G/1 queue with truncated service times $(S \wedge x)$ differs, in general, from the unfinished truncated work U_x^{π} for an M/G/1 queue with original service times S. The former one is the same for all work-conserving disciplines, while the latter one depends on the discipline π used. Only for the FB discipline, $W_x = U_x^{\mathrm{FB}}$.

Let us then return to the original service times S. We consider the conditional mean delay for MLPS disciplines within an FCFS level. If there is just one level, then, for all $x \ge 0$,

$$\overline{T}^{\mathrm{FCFS}}(x) = \overline{W}_\infty + x. \tag{11}$$

If π is an MLPS discipline with FCFS applied at level n, then, for all $a_{n-1} < x \le a_n$,

$$\overline{T}^{\pi}(x) = \frac{\overline{W}_{a_n} + x}{1 - \rho_{a_{n-1}}}. \tag{12}$$

C. Splitting an FCFS level

In this subsection we present the proof of Proposition 4. Thus, we consider here an M/G/1 queue with $\rho < 1$. Let us recall the proposition, and then prove it.

Proposition 4: Let $N \ge 1$, $n \in \{1, \ldots, N\}$, $\pi \in$ MLPS with thresholds $\{a_1, \ldots, a_N\}$, and $\pi' \in$ MLPS with thresholds $\{a_1, \ldots, a_{n-1}, a_{n+1}, \ldots, a_N\}$ such that $D_n^{\pi'} = D_n^{\pi} = D_{n+1}^{\pi} = \mathrm{FCFS}$ and

$$D_m^{\pi'} = \begin{cases} D_m^{\pi}, & m = 1, \ldots, n-1, \\ D_{m+1}^{\pi}, & m = n+1, \ldots, N. \end{cases}$$

Then, $\overline{U}_x^{\pi} \le \overline{U}_x^{\pi'}$ for all $x \ge 0$.

Proof: By Proposition 1(i), we have $\overline{U}_x^{\pi} = \overline{U}_x^{\pi'}$ for all $x \le a_{n-1}$ and $x \ge a_{n+1}$. Thus, it is sufficient to consider the interval $a_{n-1} < x < a_{n+1}$. By (12), we have

$$\overline{T}^{\pi}(x) = \begin{cases} \dfrac{\overline{W}_{a_n} + x}{1 - \rho_{a_{n-1}}}, & a_{n-1} < x \le a_n, \\ \dfrac{\overline{W}_{a_{n+1}} + x}{1 - \rho_{a_n}}, & a_n < x \le a_{n+1}, \end{cases}$$

and

$$\overline{T}^{\pi'}(x) = \frac{\overline{W}_{a_{n+1}} + x}{1 - \rho_{a_{n-1}}}, \quad a_{n-1} < x \le a_{n+1}.$$

It is easy to verify that the difference $\overline{T}^{\pi}(x) - \overline{T}^{\pi'}(x)$ is non-decreasing for all $a_{n-1} < x \le a_{n+1}$. Assume first that $a_{n-1} < x \le a_n$. Now

$$\overline{T}^{\pi}(x) - \overline{T}^{\pi'}(x) = \frac{\overline{W}_{a_n} - \overline{W}_{a_{n+1}}}{1 - \rho_{a_{n-1}}},$$

which is a negative constant. Assume then that $a_n < x \le a_{n+1}$. Then

$$\overline{T}^{\pi}(x) - \overline{T}^{\pi'}(x) = \frac{(\overline{W}_{a_{n+1}} + x)(\rho_{a_n} - \rho_{a_{n-1}})}{(1 - \rho_{a_n})(1 - \rho_{a_{n-1}})},$$

which is positive and linearly increasing. Thus, by Proposition 8, we conclude that the claim is true. ∎

D. Splitting a PS level

In this subsection we present the proof of Proposition 4. Thus, we consider here an M/G/1 queue with $\rho < 1$. Let us recall the proposition, and then prove it.

Proposition 5: Let $N \ge 1$, $\pi \in \text{MLPS}$ *with thresholds* $\{a_1, \ldots, a_N\}$, *and* $\pi' \in \text{MLPS}$ *with thresholds* $\{a_2, \ldots, a_N\}$ *such that* $D_1^{\pi'} = D_1^{\pi} = D_2^{\pi} = \text{PS}$ *and*

$$D_m^{\pi'} = D_{m+1}^{\pi} \quad \text{for } m = 2, \ldots, N.$$

Then, $\overline{U}_x^{\pi} \le \overline{U}_x^{\pi'}$ *for all* $x \ge 0$.

Proof: 1° For $N = 1$ so that $\pi \in 2\text{PS}$ with threshold a_1 and $\pi' = \text{PS}$, this has been proved in [15, Prop. 4].

2° Assume then that $N \ge 2$. By Proposition 1(i), we have $\overline{U}_x^{\pi} = \overline{U}_x^{\pi'}$ for all $x \ge a_2$. Thus, it is sufficient to consider the interval $0 \le x < a_2$.

Let $x \le a_2$. Recall from Subsection IV-A that $\overline{U}_x^{\pi}(S \wedge b)$ denotes the mean unfinished truncated work in a system where the original service times $S_1, S_2, \ldots$ are replaced by their truncated versions $S_1 \wedge b, S_2 \wedge b, \ldots$. In addition, let $\pi'' \in 2\text{PS}$ with threshold a_1 and $\pi''' = \text{PS}$. By Proposition 7 and 1°, we get

$$\overline{U}_x^{\pi} = \overline{U}_x^{\pi''}(S \wedge a_2) \le \overline{U}_x^{\pi'''}(S \wedge a_2) = \overline{U}_x^{\pi'},$$

which concludes the proof. ∎

In the preliminary version of this paper [17], we claim that the result presented above in Proposition 5 would be true for any level, not only for the lowest one ($n = 1$). The claim is essentially based on [17, Prop. 12], which we recall below (as a conjecture).

Conjecture 1: Let $\pi \in 2\text{PS}$ *with threshold* a *and* $\pi' \in 2\text{PS}$ *with threshold* a' *such that* $a \le a'$. *Then, for all* $x > a'$,

$$(\overline{T}^{\pi})'(x) \le (\overline{T}^{\pi'})'(x).$$

However, the proof of this result presented in [17] has a flaw that we have not been able to fix yet. The flaw, originally observed by one of the Infocom reviewers, is related to [17, Eq. (27)], whose left-hand-side is not simply $\beta(z)$. This is due to the fact that (with notation used in [17])

$$(g_x)'(a) = \beta(x - a) - \frac{\alpha''(x - a)}{1 - \rho_a}$$

and not $(g_x)'(a) = \beta(x - a)$ as erroneously claimed in [17]. As a consequence of the flaw, the logic behind the proof is broken. Thus, splitting a PS level that is not the lowest one is still an open problem.

V. Numerical examples

In this section we present illustrative numerical examples concerning various MLPS disciplines in the context of M/G/1 queues with hyperexponential and Pareto service time distributions, which belong to the class DHR. In addition, we consider the IHR case with a uniform distribution.

Hyperexponential service times: In the first example we have used a hyperexponential service time distribution with tail distribution function

$$\overline{F}(x) = pe^{-\mu_1 x} + (1 - p)e^{-\mu_2 x}, \quad x > 0,$$

which belongs to the class DHR. For the parameters we chose $\mu_1 = 1.0$, $\mu_2 = 0.1$, and $p = 0.8$, which results in the mean service time of $E[S] = 2.8$. The arrival rate is $\lambda = 0.3$, implying load $\rho = 0.84$ and mean unfinished work $\overline{U}_\infty = 39.0$.

In Figure 1, we have depicted the mean unfinished truncated work $\overline{U}_x^{\pi}$ as a function of the truncation threshold x for MLPS disciplines FCFS, 2FCFS(5), 3FCFS(5, 10), PS, and FB. As claimed in Proposition 4, the mean unfinished truncated work is decreasing uniformly when the number of levels is increased from one (FCFS) to two (2FCFS(5)), to three (3FCFS(5, 10)), and to infinity (FB).

In Figure 2, we have depicted the mean delay $\overline{T}^{\pi}$ as a function of the level threshold a for discipline 2FCFS(a). As claimed in Theorem 2, the mean delay for any 2FCFS(a) discipline is less than that of FCFS. It is also interesting to observe that, for a quite wide range of thresholds a, the two-level disciplines 2FCFS(a) give even lower mean delay than PS. In fact, the optimal two-level discipline is almost as good as FB. Thus, if the first threshold is chosen optimally, the gain that could be obtained by adding more levels is negligible.

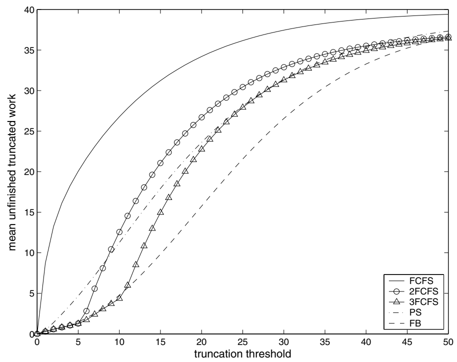

Fig. 1. Hyperexponential service times: Mean unfinished truncated work $\overline{U}_x^\pi$ as a function of the truncation threshold x for disciplines FCFS, 2FCFS(5), 3FCFS(5, 10), PS, and FB.

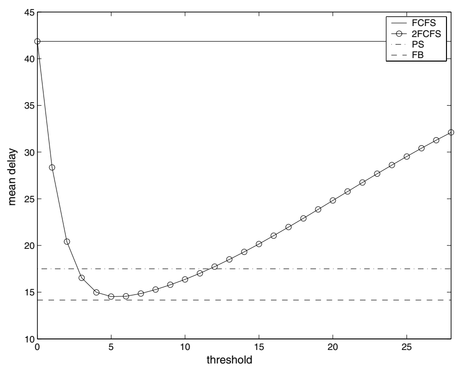

Fig. 2. Hyperexponential service times: Mean delay $\overline{T}^\pi$ as a function of the level threshold a for discipline 2FCFS(a). The three horizontal lines correspond to the mean delay of disciplines FCFS, PS and FB.

Pareto service times: The following examples are based on the Pareto service time distribution with tail distribution function

$$\overline{F}(x) = \left(\frac{1}{1+cx}\right)^\alpha, \quad x > 0,$$

for which the hazard rate is decreasing. We vary parameters c and α while keeping the mean service time fixed $E[S] = 30.0$.

First we consider the case $c = 1/36$ and $\alpha = 2.2$. The arrival rate is $\lambda = 0.03$, implying load $\rho = 0.9$. In Figure 3, we have depicted the mean delay $\overline{T}^\pi$ as a function of the level threshold a for disciplines $2\text{FCFS}(a)$ and $3\text{FCFS}(\min\{a, a^*\}, \max\{a, a^*\})$. The threshold value a^* used for the three-level disciplines is chosen to be the optimal one among the two-level disciplines $2\text{FCFS}(a)$. Note that again the optimal two-level discipline $2\text{FCFS}(a^*)$ gives even lower mean delay than PS. Having a third level gives some additional gain.

In the following example, we have a heavy tailed distribution with parameters $c = 1/24$ and $\alpha = 1.8$. In

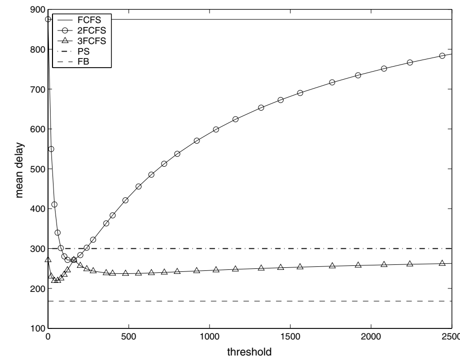

Fig. 3. Pareto service times: Mean delay $\overline{T}^\pi$ as a function of the level threshold a for disciplines $2\text{FCFS}(a)$ and $3\text{FCFS}(\min\{a, a^*\}, \max\{a, a^*\})$ with $\alpha = 2.2$. The three horizontal lines correspond to the mean delay of disciplines FCFS, PS, and FB.

this case $E[S^2] = \infty$ so that the mean delay of the FCFS discipline is unbounded. The arrival rate is $\lambda = 0.03$, implying load $\rho = 0.9$. In Figure 4, we have depicted the mean delay $\overline{T}^\pi$ as a function of the level threshold a for disciplines $2\text{PS}(a)$ and $3\text{PS}(\min\{a, a^*\}, \max\{a, a^*\})$. The threshold value a^* used for the three-level disciplines is chosen to be the optimal one among the two-level disciplines $2\text{PS}(a)$.

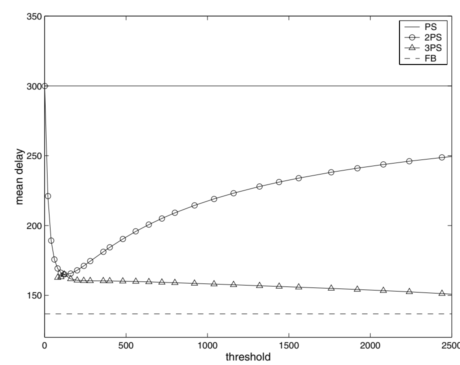

Fig. 4. Pareto service times: Mean delay $\overline{T}^\pi$ as a function of the level threshold a for disciplines $2\text{PS}(a)$ and $3\text{PS}(\min\{a, a^*\}, \max\{a, a^*\})$ with $\alpha = 1.8$. The two horizontal lines correspond to the mean delay of disciplines PS and FB.

Next we take an even heavier tailed distribution with parameters $c = 1/6$ and $\alpha = 1.2$. Thus, again $E[S^2] = \infty$. The arrival rate is now $\lambda = 0.02$, implying load $\rho = 0.6$. In Figure 5, we have again depicted the mean delay $\overline{T}^\pi$ as a function of the level threshold a for disciplines $2\text{PS}(a)$ and $3\text{PS}(\min\{a, a^*\}, \max\{a, a^*\})$. In this case, the performance of two-level or three-level policies is rather insensitive with respect to the chosen thresholds.

In the final Pareto example we have used parameters $c = 1/45$ and $\alpha = 2.5$ with $E[S^2] < \infty$. The arrival

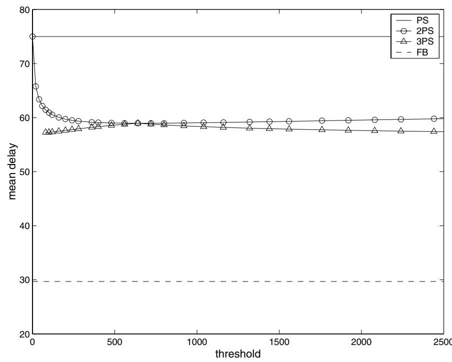

Fig. 5. Pareto service times: Mean delay $\overline{T}^{\pi}$ as a function of the level threshold a for disciplines $2\mathrm{PS}(a)$ and $3\mathrm{PS}(\min\{a, a^*\}, \max\{a, a^*\})$ with $\alpha = 1.2$. The two horizontal lines correspond to the mean delay of disciplines PS and FB.

rate is $\lambda = 0.03$, implying load $\rho = 0.9$. In Figure 6, we have depicted the mean delay $\overline{T}^{\pi}$ as a function of the level threshold a for disciplines $2\mathrm{FCFS}(a)$ and $2\mathrm{PS}(a)$. As claimed in Theorem 1, the mean delay for any $2\mathrm{PS}(a)$ discipline is less than that of the corresponding $2\mathrm{FCFS}(a)$ discipline. The shapes of the two curves are similar and the optimal level thresholds are not far away from each other. In addition, we observe that by choosing the level threshold optimally the 2FCFS discipline gives a remarkable gain compared to the FCFS discipline, performing even better than PS. The optimal 2PS discipline is almost as good as FB.

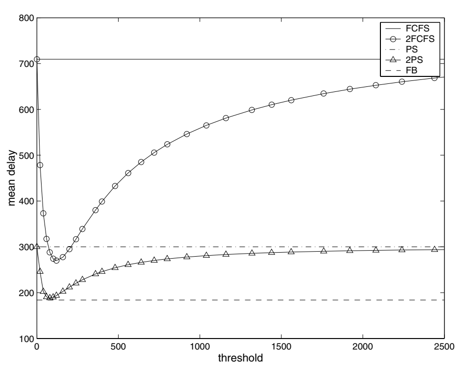

Fig. 6. Pareto service times: Mean delay $\overline{T}^{\pi}$ as a function of the level threshold a for disciplines $2\mathrm{FCFS}(a)$ and $2\mathrm{PS}(a)$ with $\alpha = 2.5$. The three horizontal lines correspond to the mean delay of disciplines FCFS, PS, and FB.

In conclusion, a small number of levels seems to give good results if at least one of the thresholds is reasonably chosen with respect to the prevailing traffic conditions. Thus, if these traffic conditions are known, two levels seem to be enough. On the other hand, if they are only partially known or they are rapidly changing, it would be reasonable to use an MLPS policy with a couple of levels for which the level thresholds are chosen from different magnitudes, for example, $a_n = a^n$ for a suitable constant a.

Uniform service times: In the final example we apply a uniform service time distribution with tail distribution function

$$\overline{F}(x) = \max\{1 - \frac{x}{P}, 0\}, \quad x > 0,$$

for which the hazard rate is *increasing*. With parameter $P = 500$, we have $E[S] = 250$. The arrival rate is $\lambda = 1/595$, implying load $\rho = 0.42$.

In Figure 7, we have depicted the mean delay $\overline{T}^{\pi}$ as a function of the level threshold a for disciplines $2\mathrm{FCFS}(a)$ and $2\mathrm{PS}(a)$. As noted at the end of Section II, the mean delay for any $2\mathrm{PS}(a)$ discipline is *greater* than that of the corresponding $2\mathrm{FCFS}(a)$ discipline. In this case, FCFS is optimal and FB the worst one [7].

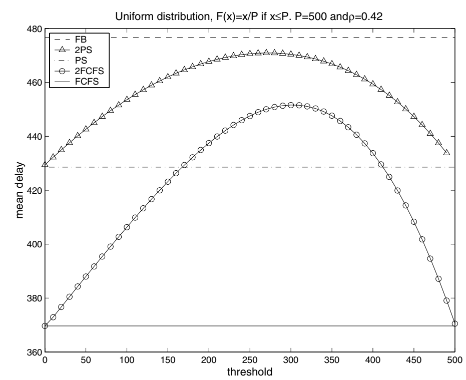

Fig. 7. Uniform service times: Mean delay $\overline{T}^{\pi}$ as a function of the level threshold a for disciplines $2\mathrm{FCFS}(a)$ and $2\mathrm{PS}(a)$. The three horizontal lines correspond to the mean delay of disciplines FB, PS, and FCFS.

VI. Conclusions

This paper compares the mean delay among MLPS disciplines. The MLPS disciplines form an important subset within the set of non-anticipating scheduling disciplines. Previous analysis considered FB and PS as internal disciplines and showed that such MLPS disciplines are better than PS with respect to the mean delay whenever the hazard rate of the service time distribution is decreasing. In this paper we allow FCFS as an internal discipline as well. Under the same DHR condition as above, we show that given an MLPS discipline, the mean delay is reduced whenever a level is added by splitting an existing one unless it is an upper level and the internal discipline is PS. As the number of levels increases, the

MLPS discipline emulates closer and closer the behavior of an FB discipline, which is known to be optimal. We also show that given an MLPS discipline, the mean delay is reduced (under the DHR condition) if an internal discipline is changed from FCFS to PS (or from PS to FB).

By numerical analysis we have evaluated the mean delay of MLPS disciplines in the presence of distributions of practical interest as Pareto and hyperexponential. Our results demonstrate that if the level threshold is chosen appropriately, the mean delay of a two-level MLPS can be close to that of FB. This result is important in view of recent work that proposes to provide differential treatment to flows on the Internet based in just two classes: mice and elephants. Our numerical examples indicate that the level splitting be advantageous even if the splitted level is an upper one provided with the PS internal discipline. However, proving this is still a open problem left for further research.

Appendix A: Truncation Lemma

In this appendix we present a lemma that is needed for the proof of Proposition 6 presented in Section III.

Lemma 1: Assume that there are two sets of real numbers, $(a_1, \dots, a_N)$ and $(b_1, \dots, b_N)$, and an index $1 \le m \le N$ that satisfy the following conditions:

(a) $\sum_{i=1}^N a_i = \sum_{i=1}^N b_i$,
(b) $a_i \le b_i$ for all $i \in \{1, \dots, m\}$,
(c) $a_i \ge b_i$ for all $i \in \{m+1, \dots, N\}$,
(d) $a_i \ge a_j$ for all $i \in \{1, \dots, m\}$ and $j \in \{i+1, \dots, N\}$.

Then, for any $x \ge 0$,

$$\sum_{i=1}^N (a_i \wedge x) \ge \sum_{i=1}^N (b_i \wedge x).$$

Proof: Note first that, for any $n \le m$,

$$\begin{aligned}\sum_{i=n+1}^N a_i &= \sum_{i=1}^N a_i - \sum_{i=1}^n a_i \\ &\overset{(a)}{=} \sum_{i=1}^N b_i - \sum_{i=1}^n a_i \\ &\overset{(b)}{\ge} \sum_{i=1}^N b_i - \sum_{i=1}^n b_i \\ &= \sum_{i=n+1}^N b_i. \end{aligned} \tag{13}$$

Denote $A = \max\{a_{m+1}, \dots, a_N\}$, and consider separately cases $x > A$ and $x \le A$.

1° Let $x > A$ and define

$$p = \begin{cases} 0 & \text{if } a_1 < x, \\ \max\{i \in \{1, \dots, m\} : a_i \ge x\}, & \text{otherwise.} \end{cases}$$

By the definition above and assumptions (b) and (d), we deduce that $b_i \ge a_i \ge x$ for all $i \in \{1, \dots, p\}$, while $a_i < x$ for all $i \in \{p+1, \dots, m\}$. Furthermore $a_i \le A < x$ for all $i \in \{m+1, \dots, N\}$. Thus,

$$\begin{aligned}\sum_{i=1}^N (a_i \wedge x) &= \sum_{i=1}^p x + \sum_{i=p+1}^N a_i \\ &\overset{(13)}{\ge} \sum_{i=1}^p x + \sum_{i=p+1}^N b_i \ge \sum_{i=1}^N (b_i \wedge x).\end{aligned}$$

2° Let then $x \le A$. By assumption (d), $a_i \ge A \ge x$ for all $i \in \{1, \dots, m\}$. Thus,

$$\begin{aligned}\sum_{i=1}^N (a_i \wedge x) &= \sum_{i=1}^m x + \sum_{i=m+1}^N (a_i \wedge x) \\ &\overset{(c)}{\ge} \sum_{i=1}^m x + \sum_{i=m+1}^N (b_i \wedge x) \ge \sum_{i=1}^N (b_i \wedge x).\end{aligned}$$

This completes the proof. ■

References

[1] L. Kleinrock. *Queueing Systems, Volume II: Computer Applications*. John Wiley & Sons, New York, 1976.
[2] L. Guo and I. Matta. Differentiated control of web traffic: A numerical analysis. In *SPIE ITCOM'2002*, Boston, MA, 2002.
[3] H. Feng and V. Misra. Mixed scheduling disciplines for network flows. *ACM SIGMETRICS Performance Evaluation Review*, 31:36–39, 2003.
[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*, pp. 762–773, Hong Kong, 2004.
[5] 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*, pp. 106–117, New York, NY, 2004.
[6] S. Yashkov. Processor-sharing queues: Some progress in analysis. *Queueing Systems*, 2:1–17, 1987.
[7] 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.
[8] D. Heyman, T. Lakshman, and A. Neidhardt. A new method for analysing feedback-based protocols with applications to engineering web traffic over the Internet. In *ACM SIGMETRICS 1997*, pp. 24–38, Seattle, WA, 1997.
[9] M. Nabe, M. Murata, and H. Miyahara. Analysis and modelling of world wide web traffic for capacity dimensioning of internet access lines. *Performance Evaluation*, 34(4):249–271, 1998.
[10] L. Massoulié and J. Roberts. Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems*, 15(1-2):185–201, 2000.

[11] S. Ben Fredj, T. Bonald, A. Proutiere, G. Regnie, and J. Roberts. Statistical bandwidth sharing: A study of congestion at flow level. In *ACM SIGCOMM 2001*, pp. 111–122, San Diego, CA, 2001.

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

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

[14] K. Avrachenkov, U. Ayesta, and P. Brown. Batch arrival processor sharing with application to multilevel processor sharing scheduling. In *Queueing Systems*, 50:459-480, 2005.

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

[16] S. Aalto, U. Ayesta, and E. Nyberg-Oksanen. M/G/1/MLPS compared to M/G/1/PS. In *Operations Research Letters*, 33:519–524, 2005.

[17] S. Aalto, U. Ayesta. Mean delay comparison among multilevel proccessor-sharing scheduling disciplines. Report 11-04/05-fall, Institut Mittag-Leffler, 2005.