Discriminatory Processor Sharing Revisited

K. Avrachenkov[†], U. Ayesta[†,‡,⋆], P. Brown[‡], R. Núñez-Queija[⋆,⋆⋆]

[†] INRIA Sophia Antipolis, France. e-mail: {k.avrachenkov,uayesta}@sophia.inria.fr
[‡] France Telecom R&D. e-mail: patrick.brown@francetelecom.com
[⋆] CWI, Amsterdam, The Netherlands. e-mail: {urtzi.ayesta,sindo}@cwi.nl
[⋆⋆] Eindhoven University of Technology, The Netherlands. e-mail: r.nunez@tue.nl

***Abstract*— As a natural multi-class generalization of the well-known (egalitarian) Processor Sharing (PS) service discipline, Discriminatory Processor Sharing (DPS) is of great interest in many application areas, including telecommunications. Under DPS, the mean response time conditional on the service requirement is only known in closed form when all classes have exponential service requirement distributions. For generally distributed service requirements, Fayolle *et al.* [1] showed that the expected conditional response times satisfy a system of integro-differential equations. In this paper, we exploit that result to prove that, provided the system is stable, for each class the expected unconditional response time is finite and that the expected conditional response time has an asymptote. The asymptotic bias of each class is found in closed form, involving the mean service requirements of all classes and the second moments of all classes but the one under consideration. In the course of the development we prove two other results that are of independent interest: we establish a conservation law for the time average unfinished work of all classes and, using a stochastic coupling argument, we show that the response times of different classes are stochastically ordered according to the DPS weights. Finally, we study DPS as a tool to achieve size based scheduling and we provide guidelines as to how the weights of DPS must be chosen such that DPS outperforms PS.**

I. Introduction

Kleinrock [2] introduced and first studied the so-called Discriminatory Processor Sharing (DPS) discipline[1], where a single server is shared by M job classes. All jobs present in the system are served simultaneously with rates controlled by a vector of weights $\{g_k > 0; k = 1, \ldots, M\}$. If there are N_j jobs of class j present in the system, $j = 1, \ldots, M$, each class-k job is served at rate $g_k / \sum_{j=1}^{M} g_j N_j$. When all the weights are equal, this is equivalent to the standard Processor Sharing (PS) system.

By changing the DPS weights, one can effectively control the instantaneous service rates of different job classes. Thus, an appropriate choice of the DPS weights may enable differentiated quality of service among different job classes. The range of applications for DPS is broad. In the context of the Internet, one can think of a situation where $ADSL$ subscribers are offered different payment rates and in return obtain corresponding shares of the available bandwidth. The DPS system also appears as a natural candidate to model flow level sharing of TCP flows with different flow characteristics such as round trip times and packet loss probabilities [3], [4], [5], [6]. DPS is also a convenient theoretical abstraction of the weighted round-round robin discipline, which is used in operating systems for task scheduling.

Fayolle *et al.* [1] proved that previously found expressions for the expected conditional response time [2], [7] were not correct and obtained the expected conditional response times as the solution of a system of integro-differential equations (see Section II for more details). In the case of exponentially distributed service requirements, these lead to closed-form expressions for the expected conditional response times and the unconditional mean response times can be found as the solution to a system of linear equations. In particular, it is shown in [1] that, independent of the weights, for each class the slowdown under DPS tends to the (constant) slowdown of the PS model as the service requirement increases to infinity. Also for exponential service requirement distributions, Rege and Sengupta [8] obtained higher moments of the queue length distribution as the solutions

[1]In the original paper, Kleinrock [2] used the term Priority Processor-Sharing.

Corresponding author: Urtzi Ayesta.

This work is part of a French-Dutch *Van Gogh* research project funded by NWO (The Netherlands Organization for Scientific Research) and EGIDE under grant VGP 61-520.

R. Núñez-Queija is an EQUANET project member, funded by the Dutch Ministry of Economic Affairs via its agency SenterNovem.

to linear equations and, under the same assumptions, they also proved a heavy-traffic limit theorem. These results were recently extended to phase type distributions by Van Kessel *et al.* [6], [9]. For general service requirement distributions, Rege and Sengupta [10] proved a decomposition theorem for the response times conditional on service requirements. Grishechkin [11] further explored the relationship between Processor Sharing queues and Crump-Mode-Jagers branching processes. More recently, Bonald and Massoulie [4] discussed the applicability of DPS in modeling user-level performance in the Internet. Guo and Matta [12] have carried out an approximate analysis of DPS for general distributions. More recently, the asymptotic behavior of the delay distribution in a DPS system has been studied by Borst *et al.* [5], Altman *et al.* [3] study the behavior of DPS in overload and Bonald&Proutière [13] obtain insensitive bounds for the queue length distribution for certain parameter ranges.

This paper is organized as follows: in the next section, we review the DPS model and integro-differential equations that describe the expected conditional response times. In Section III we prove that the mean unconditional response time of each class is finite, provided the system is stable. In Section IV we obtain a conservation law for the mean unfinished work. Then, in Section V we show that the expected conditional response times for different job classes are stochastically ordered according to the DPS weights. Then, using the above mentioned conservation law, we study in Section VI the asymptotic behavior of the expected conditional response times. In Section VII we compare the unconditional mean sojourn time of a DPS discipline with the PS discipline in the context of size based scheduling. In Section VIII we investigate the behavior of DPS by means of numerical examples. The paper ends with some concluding remarks in the final section.

II. THE DPS MODEL

In the ensuing analysis, unless otherwise stated, we assume that $k \in \mathcal{N}$ is a job class index, where $\mathcal{N} = \{1, \dots, M\}$ is the set of indexes. Except for Proposition 1 below, we assume throughout the paper that the arrival process of class k is a Poisson process with rate λ_k and the service requirement distribution is $F_k(\cdot)$. We shall use $\overline{F}_k(x) = 1 - F_k(x)$ to denote the complementary class-k service requirement distribution and $E[X_k^n]$ to denote its n-th moment. For notational ease we shall write $E[X_k] = E[X_k^1]$. The aggregated job arrival process is Poisson with rate $\lambda = \sum_{j=1}^{M} \lambda_j$. Conditioning on the class a job belongs to, the n-th moment of the aggregated service requirement $E[X^n]$ is given by $E[X^n] = \sum_{j=1}^{M} \frac{\lambda_j}{\lambda} E[X_j^n]$. Define

$$\psi_k(x) = \lambda_k g_k \overline{F}_k(g_k x),$$

$$\Psi(x) = \sum_{k=1}^{M} \psi_k(x).$$

Steady state exists in the non-saturated regime, i.e., when $\rho = \sum_{k=1}^{M} \rho_k < 1$, where $\rho_k = \lambda_k E[X_k]$ for $k \in \mathcal{N}$. Let $T_k(t)$ be the expected conditional response time (time in the system) of a class-k job whose service requirement is t and let $T_k'(t)$ be its derivative. It was proven in [1] that $T_k(t)$ can be expressed as

$$T_k(t) = g_k \int_0^{t/g_k} a(x)dx + \int_0^{t/g_k} b(x)dx, \qquad (1)$$

where $a(t)$ is the unique solution of

$$a(t) = 1 + \int_0^t a(u)\Psi(t-u)du, \qquad (2)$$

and $b(t)$ is a function that satisfies

$$\begin{aligned} b(t) \;&=\; \sum_{j=1}^{M} \int_0^{\infty} g_j T_j'(g_j u)\psi_j(t+u)du \\ &\quad + \int_0^t b(u)\Psi(t-u)du. \end{aligned} \qquad (3)$$

For general service requirement distributions with finite second moments, it is known that $\frac{1}{t}T_k(t)$, i.e., the *slowdown* of class k, tends to $\frac{1}{1-\rho}$ as $t \to \infty$, for all classes [1] . In the case of exponential service requirement distributions, a closed form expression for the expected conditional response time was obtained

$$T_k(t) = \frac{t}{1-\rho} + \sum_{j=1}^{m} C_j \left(1 - e^{-\alpha_j t/g_k}\right), \qquad (4)$$

where m is equal to the number of different elements in the vector $v = (g_k E[X_k])_{k=1,\dots,M}$. The coefficient α_j is a positive root of a rational function and C_j can be expressed as a function of α_i, $i = 1, \dots, m$, and the elements of v [1]. Thus, the bias of the asymptote for $t \to \infty$ is given by $\sum_{j=1}^{m} C_j$ and its value may be positive, negative or zero. Unfortunately, the expression for the coefficients $C_j, i = 1, \dots, m$ given in [1] is not very insightful and it does not reveal how the bias depends on the arrival rates and the required service time distributions.

III. Finiteness of the expected unconditional response time

In this section, we show that the expected unconditional response times of all classes in a DPS system is finite whenever $\rho < 1$. Let $\mathcal{C}[0,\infty)$ be the space of continuous, bounded non-negative functions. In the next lemma we show that the functions $a(t), b(t) \in \mathcal{C}[0,\infty)$.

Lemma 1: Let the service time distributions of all classes have a finite mean and assume $\rho < 1$. Then, the solutions to the integral equations (2) and (3) exist and are unique. Furthermore, the solutions are continuous, bounded, and non-negative.

Proof: We first consider the solution for $a(t)$. Equation (2) is a defective renewal equation (see for example [14], [15]). Hence, $a(t)$ counts the number of arrivals in the interval $[0,t)$ of the defective renewal process with renewal density $\Psi(\cdot)$, starting with a renewal at time 0. Note that, by definition, the probability that the time between two consecutive renewals is smaller than t equals $\int_0^t \Psi(u)du$. The renewal equation is defective since

$$\begin{aligned}\int_0^\infty \Psi(u)du &= \sum_{j=1}^M \lambda_j g_j \int_0^\infty \overline{F}_j(g_j u)du \\ &= \rho < 1. \end{aligned} \tag{5}$$

As a consequence [14], $a(t)$ is a continuous, non-negative, non-decreasing, upper-bounded function and

$$\lim_{t\to\infty} a(t) = \frac{1}{1-\int_0^\infty \Psi(u)du} = \frac{1}{1-\rho}.$$

In particular, for all $t \geq 0$, $a(t) \leq 1/(1-\rho)$.

We now consider $b(t)$. As shown in [1], equation (3) can be expressed as

$$b(t) = c(t) + \int_0^\infty b(u)\Psi(t+u)du + \int_0^t b(u)\Psi(t-u)du, \tag{6}$$

where

$$c(t) = \sum_{j=1}^M g_j \int_0^\infty a(u)\psi_j(t+u)du. \tag{7}$$

Since $a(t)$ is a non-negative bounded function, so is $c(t)$. Furthermore, using that $a(t) \leq 1/(1-\rho)$,

$$c(t) \leq \sum_{j=1}^M \frac{g_j}{1-\rho} \int_0^\infty \psi_j(t+u)du \leq \frac{\sum_{j=1}^M g_j\rho_j}{1-\rho}.$$

Let us now consider the fixed point iterations

$$\begin{aligned} b_{n+1}(x) &= c(t) + \int_0^\infty b_n(u)\Psi(t+u)du + \\ &\quad + \int_0^t b_n(u)\Psi(t-u)du, \quad n = 0, 1, \ldots, \end{aligned}$$

on the complete functional space $\mathcal{C}[0,\infty)$ with the supremum metric $||b|| = \sup_t\{b(t)\} < \infty$. Define the integral operator $\mathcal{A}[\beta(x)]$ as follows:

$$\begin{aligned} \mathcal{A}[\beta(t)] &= c(t) + \int_0^\infty \beta(u)\Psi(t+u)du \\ &\quad + \int_0^t \beta(u)\Psi(t-u)du. \end{aligned}$$

The operator $\mathcal{A}[\beta(t)]$ maps the space $\mathcal{C}[0,\infty)$ onto itself.

If we show that the integral operator $\mathcal{A}[\beta(t)]$ is a contraction, then the integral equation (6) has a unique solution in $\mathcal{C}[0,\infty)$ [16, Section 2.8]. Let d be the distance in the metric space $\mathcal{C}[0,\infty)$, that is, $d(\beta_1,\beta_2) = \sup_t |\beta_1(t) - \beta_2(t)|$. Note that

$$\begin{aligned} & d(\mathcal{A}[\beta_1], \mathcal{A}[\beta_2]) = \sup_t\{|\mathcal{A}[\beta_1(t)] - \mathcal{A}[\beta_2(t)]|\} \\ \leq\ & \sup_t\{|\beta_1(t) - \beta_2(t)|\} \times \\ & \times \sup_t \left(\int_0^\infty \Psi(t+y)dy + \int_0^t \Psi(t-y)dy\right) \\ =\ & d(\beta_1,\beta_2)\int_0^\infty \Psi(u)du = \rho d(\beta_1,\beta_2). \end{aligned}$$

Thus, if $\rho < 1$, the operator $\mathcal{A}[\beta(t)]$ is indeed a contraction mapping on $\mathcal{C}[0,\infty)$. Thus its fixed point solution $b(t)$ is a continuous, non-negative bounded function. ■

Lemma 1 allows us to prove that in a stable DPS system, the mean queue lengths are finite irrespective of higher order characteristics of the service requirement distributions (Theorem 1 below). This extends the well-know property for ordinary multi-class PS with Poisson arrivals, in which case the mean queue lengths only depend on the arrival rates and the mean service requirements.

Remark. Note that, by Little's formula, this result is equivalent to the statement that the expected unconditional response time of class k is finite if its weight is strictly positive. It shows the benefits of time-sharing scheduling disciplines with respect to strict priority rules. Under strict priority disciplines, if the second moment of a class has an infinite second moment, the expected unconditional response times of all the classes with lower priority are infinite (see also Section VII).

The robustness of DPS comes from the fact that the common resource is shared among the different classes. For instance, the share of the server a given class gets depends on its weight as well as on the number of jobs present in the system. As a consequence, the share obtained by a class will increase proportionally as the number of jobs of this class in the system augments and as a consequence the DPS discipline prevents classes with small weights suffer from starvation.

Theorem 1: Let the service time distributions of all the classes have a finite mean and $\rho < 1$. Then, if the weight of class k is strictly positive, the expected number of class-k jobs present in the system is finite.

Proof: Because of PASTA [17], the distribution at arrival epochs is equal to the stationary distribution. Hence, when a job arrives to the system, it finds on average $E[N_j]$ class-j jobs, $j = 1, \ldots, M$, where N_j has the time average distribution of the number of class-j jobs in the system. For the moment we do not exclude $E[N_j] = \infty$. Let $T_k'(t)$ be the derivative of the conditional response time. From [1, (2.1)] we have that

$$T_k'(0) \quad = \quad 1 + \sum_{j=1}^{M} \frac{g_j}{g_k} E[N_j]. \tag{8}$$

On the other hand, taking the derivative of equation (1), we obtain $T_k'(t) = a(t/g_k) + \frac{1}{g_k} b(t/g_k)$. Evaluating at $t = 0$ and noting that $a(0) = 1$ we have

$$T_k'(0) = 1 + \frac{1}{g_k} b(0), \tag{9}$$

so that

$$b(0) = \sum_{j=1}^{M} g_j E[N_j].$$

Since, $b(t)$ is a non-negative bounded function, if follows that if $g_j > 0$, then $E[N_j]$ is finite. ∎

Remark. Uniqueness, existence and boundedness of $b(t)$ was proven in [1] under the additional assumption that the integral $\int_0^\infty |b(u)| \Psi(u) du$ is bounded.

IV. Conservation law for DPS

The so-called *work conservation* property is fundamental to single-server (multi-class) systems. For any single server queue with M job classes, let $U_k(t)$ be the unfinished work at time t of class-k jobs, $k \in \mathcal{N}$, and let $U(t) = \sum_{j=1}^{M} U_j(t)$ denote the total unfinished work in the system. The unfinished work in the system $U(t)$ is a function that has vertical jumps at arrival epochs – the sizes of which are equal to the corresponding service requirements of customers – and remains constant when it hits the horizontal axis. When $U(t) > 0$, the unfinished work drains with a rate equal to the service rate. We say that the scheduling discipline is work-conserving if $U(t)$ decreases at rate 1 whenever $U(t) > 0$.

Then, a sample path argument shows that for any work-conserving discipline, the unfinished work in the system is the same, regardless of the scheduling discipline deployed. In the particular case of DPS, this implies that the total unfinished work in the system ($U(t) = \sum_{j=1}^{M} U_j(t)$) is independent of the particular values of the vectors $\{g_k; k = 1, \ldots, M\}$.

This property has led to the development of so-called work-conservation laws, see [18] for a survey and [19], [20] and references therein for the application of conservation laws to the design of optimal scheduling disciplines.

In Proposition 1 below, we state a conservation law for DPS. This result is key to the asymptotic analysis in Section VI.

Proposition 1: Consider a DPS system with M classes in which the superposed arrival process constitutes a renewal sequence. As before, let λ_j be the arrival rate of class-j jobs (not necessarily Poisson) and let the second moments of the service time requirement distributions be finite, i.e., $E[X_j^2] < \infty$, for all j. In addition, assume $\rho < 1$. Then,

$$\sum_{j=1}^{M} \lambda_j \int_0^\infty T_j(x) \overline{F}_j(x) dx = \overline{U}, \tag{10}$$

where $\overline{U}$ is the time average unfinished work in the system.

Proof: We consider an arbitrary class $j \in \mathcal{N}$. Let W_j^i, $i = 1, 2, \ldots$, be the *cumulative burden* of the i-th class-j job on the unfinished work in the system during its complete response time. Formally, $W_j^i := \int_{t=0}^{T_j^i} R_j^i(t_j^i + t) dt$ where t_j^i is the arrival time, T_j^i the response time and $R_j^i(t)$ the unfinished work of the i-th class-j customer at time t. In particular, $R_j^i(t_j^i)$ is the total service requirement of this job and $R_j^i(t_j^i + T_j^i) = 0$.

Since $\rho < 1$, the busy period has a finite length *with probability 1*. Furthermore, since the superposed arrival process is a renewal process, the begin points of the consecutive busy periods constitute regeneration points and, as a consequence, the sequence $\{W_j^n\}_{n=1}^{\infty}$ is a regenerative process with finite cycle lengths. Hence, the process $\{W_j^n\}_{n=1}^{\infty}$ is stationary and ergodic. Let $\tau_j^i(x)$ be the response time of the i-th class-j job. If $E[W_j]$ denotes the expected cumulative burden of an arbitrary class-j job then

$$\begin{aligned} E[W_j] &= E[\int_{x=0}^{\infty}\int_{t=0}^{\tau_j^i(x)} R_j^i(t_j^i+t)dtdF_j(x)] \\ &= E[\int_{x=0}^{\infty}\int_{u=0}^{x} \tau_j^i(u)dudF_j(x)] \\ &= \int_{u=0}^{\infty}\int_{x=u}^{\infty} T_j(u)dF_j(x)du \\ &= \int_{0}^{\infty} T_j(u)\overline{F}_j(u)du. \end{aligned}$$

In the second equality we use that, at any time, the share that each class-j job obtains is independent of the remaining service requirement, provided it is positive, cf. [21, page 164].

Applying the Palm inversion formula[2][22], [21] to the unfinished work of class j, we obtain $\overline{U_j} = \lambda_j E[W_j]$, where $\overline{U_j}$ is the time average unfinished class-j work in the system. Summing over all the classes and noting that $\sum_{j=1}^{M}\overline{U}_j = \overline{U}$ is equal to the time average of the total unfinished work in the system, we obtain (10). ■

Corollary 1: If, in addition to the assumptions in Proposition 1, we assume that class-j jobs arrive according to a Poisson process, then

$$\sum_{j=1}^{M}\lambda_j\int_{0}^{\infty} T_j(x)\overline{F}_j(x)dx = \frac{\sum_{j=1}^{M}\lambda_j E[X_j^2]}{2(1-\rho)}.$$

Proof: We are only left with the evaluation of $\overline{U}$. Note that $\overline{U}$, the time average unfinished work in the system, is independent of the scheduling discipline being deployed. Hence we may determine $\overline{U}$ assuming First-Come First-Served scheduling (or any other work-conserving discipline). The compound job arrival process is Poisson with rate $\lambda = \sum_{j=1}^{\infty}\lambda_j$ and the second moment of the service time requirement distribution is $E[X^2] = \sum_{j=1}^{\infty}\frac{\lambda_j}{\lambda}E[X_j^2]$. Then the time-average unfinished work is simply given by the Pollaczek-Khinchin formula,

$$\overline{U} = \frac{\lambda E[X^2]}{2(1-\rho)} = \frac{\sum_{j=1}^{M}\lambda_j E[X_j^2]}{2(1-\rho)},$$

and the result follows. ■

[2]The Palm inversion formula is commonly denoted as $H = \lambda G$ and, informally, states that the time average of a process (H), is equal to the sample rate (λ) times the expected contribution of each sample (G)

V. Ordering conditional response times

Using sample path arguments, we show that conditional response times are stochastically ordered according to the DPS weights.

Theorem 2: Let $\tau_k(x)$, $k = 1, 2, \ldots, M$, be the response time in steady state of a class-k job that requires x units of service. If $g_k \geq g_l$, then $\tau_k(x) \leq_{st} \tau_l(x)$, that is, $P(\tau_k(x) > y) \leq P(\tau_l(x) > y)$ for all $y \geq 0$. In particular, for all $x \geq 0$, we have that $T_k(x) = E[\tau_k(x)] \leq E[\tau_l(x)] = T_l(x)$.

Proof: Consider two identical copies of the system, which we refer to as system I and system II. That is, each arriving customer is duplicated and one copy is placed in each of the two systems, the two copies having the same service requirement and belong to the same class. Denote the numbers of customers of class k, $k = 1, 2, \ldots, M$, in the two systems at time t by $N_k^I(t)$ and $N_k^{II}(t)$, respectively, and their residual service requirements by $R_{k,i}^I(t)$, for $i = 1, 2, \ldots, N_k^I(t)$ and $R_{k,i}^{II}(t)$, for $i = 1, 2, \ldots, N_k^{II}(t)$. Suppose the two systems are in steady state at time 0. At this time we place a customer with service requirement x in each system and we shall refer to the two copies as customer I and customer II for short. Customer I is assigned weight g_I and customer II gets weight g_{II}. Consistent with our previous notation, we denote the sojourn time of customer I by $\tau_I(x)$ and that of customer II by $\tau_{II}(x)$.

Note that for $t \leq 0$, $N_k^I(t) = N_k^{II}(t)$ and $R_{k,i}^I(t) = R_{k,i}^{II}(t)$, $i = 1, 2, \ldots, N_k^I(t)$, but that these quantities will differ after time 0 if $g_I \neq g_{II}$. Suppose that $g_I \leq g_{II}$. We now show that for all $k = 1, 2, \ldots, M$, and $t \in (0, \min\{\tau_I(x), \tau_{II}(x)\})$: $N_k^I(t) \leq N_k^{II}(t)$ and $R_{k,i}^I(t) \leq R_{k,i}^{II}(t)$, $i = 1, 2, \ldots, N_k^I(t)$. Here, the inequality $R_{k,i}^I(t) \leq R_{k,i}^{II}(t)$ refers to customers in the two systems that were each others copies upon arrival, which may require re-shuffling of the customers when a customer leaves. For instance, if the copy in I leaves before that in II, then the copy in II gets an index larger than $N_k^I(t)$ (we will show that the copy in II can not leave sooner).

Let the sequence $0 = t_0 < t_1 < t_2 \ldots < \min\{\tau_I(x), \tau_{II}(x)\}$ denote the time points at which the number of customers of some class changes in one of the two systems (or in both). Suppose that $N_k^I(t_n+) \le N_k^{II}(t_n+)$ and $R_{k,i}^I(t_n+) \le R_{k,i}^{II}(t_n+)$, $i = 1, 2, \ldots, N_k^I(t_n+)$. (This is trivially satisfied for $n = 0$.) We will show that this inequality is also valid if we replace n by $n+1$. Note that, for $t \in (t_n, t_{n+1})$ (in this period the numbers of customers do not change in either system), the service rates of customers of any class k in the two systems satisfy

$$\frac{g_k}{g_I + \sum_k g_k N_k^I(t)} \ge \frac{g_k}{g_{II} + \sum_k g_k N_k^{II}(t)}. \tag{11}$$

Thus, all customers in system I receive more service than their copies in system II, except (possibly) for customers I and II (below we will indeed show that the opposite is true for these customers). Thus, for $t \in (t_n, t_{n+1})$ we have $N_k^I(t) \le N_k^{II}(t)$ and $R_{k,i}^I(t) \le R_{k,i}^{II}(t)$, $i = 1, 2, \ldots, N_k^I(t)$.

Now, at time $t = t_{n+1}$ these relations are preserved, as we may argue by considering the different possibilities: (i) At time t_{n+1} a new customer arrives. Then a copy is placed in both systems, which does not violate the inequalities. (ii) At time t_{n+1} a customer in system I completes service and leaves while his copy still has residual service time left in system II. Then the numbers of customers in system I is further reduced, while the inequality regarding the residual services of the leaving customer and his copy is trivially satisfied. (iii) At time t_{n+1} a customer in system II completes service and leaves. If his copy in system I was also present, then necessarily it also completes service at time t_{n+1} (it must have zero service time left) and thus, the inequalities are preserved (the numbers of customers are reduced in both systems simultaneously). If there is no departure from system I, then that copy must have left before and thus, there is at least one more customer of that class in system II than in system I and, thus, the departing customer does not violate the inequalities.

Thus the ordering of the two systems is preserved while customers I and II are both present. It remains to show that customer II leaves before customer I, i.e., $\tau_I(x) \ge \tau_{II}(x)$. Let the total amounts of work at time t due to other customers than I and II in both systems be represented by

$$W^I(t) = \sum_{k=1}^{M} \sum_{i=1}^{N_k^I(t)} R_{k,i}^I(t)$$

and

$$W^{II}(t) = \sum_{k=1}^{M} \sum_{i=1}^{N_k^{II}(t)} R_{k,i}^{II}(t).$$

In the remainder let $0 < t < \min\{\tau_I(x), \tau_{II}(x)\}$. From the construction it follows that $W^I(0) = W^{II}(0)$ and $W^I(t) \le W^{II}(t)$. Let $A(t)$ be the total amount of work that arrived during $(0, t)$. Since the systems are work conserving, the amounts of service received by customers I and II over the period (0, t) are $B_I(t) = t - (W^I(0) + A(t) - W^I(t))$ and $B_{II}(t) = t - (W^{II}(0) + A(t) - W^{II}(t))$, so that $B_I(t) \le B_{II}(t)$. This is true for any $0 < t < \min\{\tau_I(x), \tau_{II}(x)\}$. Since customers I and II have the same service requirement x, customer I can not leave before customer II does. Finally set $g_I = g_l$ and $g_{II} = g_k$. ■

The next proposition provides several other relations among the expected conditional response times of different classes.

Proposition 2: If $g_k \ge g_l$, then for all $t \ge 0$:

(P1) $T_k(g_k t) \ge T_l(g_l t)$,

(P2) $\frac{T_k(g_k t)}{g_k t} \le \frac{T_k(g_l t)}{g_l t}$

Proof: We first show Property (P1). From equation (1) we have

$$T_k(g_k t) - T_l(g_l t) = (g_k - g_l) \int_0^t a(u) du,$$

and Property (P1) follows as a consequence of $a(t)$ being a non-negative function (see Lemma 1).

Property (P2) follows similarly since

$$\frac{1}{g_k} T_k(g_k t) - \frac{1}{g_l} T_l(g_l t) = (\frac{1}{g_k} - \frac{1}{g_l}) \int_0^t b(u) du,$$

and the fact that $b(u)$ is a non-negative function (see Lemma 1). ■

We provide an interpretation of Properties (P1) and (P2) of Proposition 2 with a particular example. Consider two classes , $k, l \in \mathcal{N}$, such that the respective weights satisfy $g_k = 2g_l$, that is, class-k jobs are served twice as fast as class-l jobs. Then, Property (P1) states that the expected conditional response time of a class-k job of size t, will be larger than that of a class-l job of size $t/2$. On the other hand, Property (P2) shows that the expected *slow down* of the class-k job with size t will be smaller than that of a class-l job of size $t/2$, and thus Property (P2) provides an upper bound on the degradation with respect to the response time predicted by Property (P1).

VI. Asymptotic analysis of the expected conditional response time

In this section we study the asymptotic behavior of the conditional expected response time, as the service requirement tends to infinity. It was proven in [1] that when $E[X_j^2] < \infty$, for all $j \in \mathcal{N}$, the slowdown in the DPS system approaches the slowdown of the PS system as the service requirement increases, that is $T_k(t)/t \to \frac{1}{1-\rho}$. Let $\mathcal{S}_k \subseteq \mathcal{N}$ denote the set of classes whose weights are equal to the weight of class k, that is, $\mathcal{S}_k = \{i \in \mathcal{N} | g_i = g_k\}$. Denote as $\mathcal{S}_k^c$ the complement set of $\mathcal{S}_k$, that is, $\mathcal{S}_k^c = \{i \in \mathcal{N} | g_i \neq g_k\}$.

We prove, under the same conditions as in [1] that the expected conditional response time of class k has an asymptote. Furthermore, we provide a simple closed form expression for the asymptotic bias and we show that the value of the bias for class k depends only on the second moments of the service requirement distributions of classes $\mathcal{S}_j^c$.

Theorem 3: Let $E[X_j^2]$ be finite for all $j \in \mathcal{N}$. Then, the expected conditional response time of class k has an asymptote with slope $1/(1-\rho)$ and the following bias

$$\lim_{t\to\infty}\left(T_k(t) - \frac{t}{1-\rho}\right) = \frac{\sum_{j:j\in\mathcal{S}_k^c} \lambda_j (1-\frac{g_k}{g_j}) E[X_j^2]}{2(1-\rho)^2}. \quad (12)$$

Proof: In [1] it was shown that $T_k(t) \sim \frac{t}{1-\rho}$, as $t \to \infty$. To show existence of an asymptote for $T_k(t)$ with slope $1/(1-\rho)$, we need to show existence of $\lim_{t\to\infty}\left(T_k(t) - \frac{t}{1-\rho}\right)$. From equation (1) we observe that if the above limit exists, it can be calculated as follows:

$$\lim_{t\to\infty}\left(T_k(t) - \frac{t}{1-\rho}\right) = \lim_{t\to\infty}\left(g_k \int_0^{t/g_k}\left(a(x) - \frac{1}{1-\rho}\right)dx + \int_0^{t/g_k} b(x)dx\right).$$

We first show that $\int_0^\infty (a(x) - 1/(1-\rho))dx$ exists as well and we derive simple expressions for these integrals. From equation (2) it follows that

$$\begin{aligned} a(t) - \frac{1}{1-\rho} &= 1 + \int_0^t a(u)\Psi(t-u)du - \frac{1}{1-\rho} \\ &= \int_0^t \left(a(u) - \frac{1}{1-\rho}\right)\Psi(t-u)du \\ &\quad + \frac{1}{1-\rho}\int_0^t \Psi(u)du - \frac{\rho}{1-\rho}. \quad (13) \end{aligned}$$

Using equation (5) we can rewrite equation (13) as

$$\begin{aligned} a(t) - \frac{1}{1-\rho} &= \int_0^t \left(a(u) - \frac{1}{1-\rho}\right)\Psi(t-u)du \\ &\quad - \frac{1}{1-\rho}\int_t^\infty \Psi(u)du. \end{aligned}$$

Integrating with respect to t we have

$$\begin{aligned} &\int_0^\infty \left(a(t) - \frac{1}{1-\rho}\right)dt \\ &= \int_0^\infty \int_0^t \left(a(u) - \frac{1}{1-\rho}\right)\Psi(t-u)dudt \\ &- \frac{1}{1-\rho}\int_0^\infty \int_t^\infty \Psi(u)dudt \\ &= \int_0^\infty \left(a(u) - \frac{1}{1-\rho}\right)du \int_0^\infty \Psi(t)dt \\ &- \frac{1}{1-\rho}\int_0^\infty \Psi(u)udu, \end{aligned}$$

and thus

$$(1-\rho)\int_0^\infty \left(a(t) - \frac{1}{1-\rho}\right)dt = -\frac{1}{1-\rho}\int_0^\infty \Psi(u)udu.$$

Proceeding similarly as in equation (5), it follows that $\int_0^\infty \Psi(u)udu = \sum_{j=1}^M \frac{\lambda_j E[X_j^2]}{2g_j}$. In particular, the latter justifies the above calculations if the service requirement distributions have finite second moments. Thus, we have

$$\int_0^\infty \left(a(t) - \frac{1}{1-\rho}\right)dt = \frac{-1}{2(1-\rho)^2}\sum_{j=1}^M \frac{\lambda_j E[X_j^2]}{g_j}. \quad (14)$$

Let us now determine $\int_0^\infty b(t)dt$. Existence of $\int_0^\infty b(x)dx$ was proven in [1] provided that the service requirement distributions have a finite second moment. We first analyze the first term on the right hand side of equation (3). Integrating by parts we obtain

$$\begin{aligned} &\int_0^\infty g_j T_j'(g_j u)\psi_j(t+u)du \\ = &\int_0^\infty \lambda_j g_j^2 T_j'(g_j u)\overline{F}_j(g_j(t+u))du \\ = &\lambda_j g_j T_j(g_j u)\overline{F}_j(g_j(t+u)) \,|_{u=0}^{u=\infty} \\ &- \int_0^\infty \lambda_j g_j T_j(g_j u) d\overline{F}_j(g_j(t+u)). \quad (15) \end{aligned}$$

We evaluate the first term on the right hand side of the above equation. Integrating by parts we get

$$\int_0^\infty x dF_k(x) = \int_0^\infty \overline{F}_k(x)dx + \lim_{x\to\infty} x\overline{F}_k(x),$$

thus, if the service time requirement has a finite mean, $\lim_{x\to\infty} x\overline{F}_k(x) = 0$. Now, since $T_k(x) \sim x/(1-\rho)$ as $x \to \infty$, there exists some $L < \infty$ such that for all $x \geq 0$, $T_k(x) \leq Lx$. Consequently,

$$\lim_{x\to\infty} T_k(x)\overline{F}_k(x) \leq \lim_{x\to\infty} Lx\overline{F}_k(x) = 0.$$

On the other hand, it holds that $\lim_{t\to 0} T_j(g_j t) = 0$. Therefore equation (15) becomes

$$\int_0^\infty g_j T_j'(g_j u)\psi_j(t+u)du = -\int_0^\infty \lambda_j g_j T_j(g_j u) d\overline{F}_j(g_j(t+u)).$$

We now calculate the contribution of $b(t)$ to the bias of the asymptote. From equation (3) we have

$$\begin{aligned}\int_0^\infty b(t)dt &= \\ &-\sum_{j=1}^M \lambda_j g_j^2 \int_{u=0}^\infty T_j(g_j u)\int_{t=0}^\infty d\overline{F}_j(g_j(t+u))du \\ &+\int_0^\infty \int_0^t b(u)\Psi(t-u)dudt \\ =& \sum_{j=1}^M \lambda_j g_j \int_0^\infty T_j(g_j u)\overline{F}_j(g_j u)du \\ &+\int_0^\infty b(u)\int_t^\infty \Psi(t-u)dtdu\end{aligned}$$

and hence

$$(1-\rho)\int_0^\infty b(t)dt = \sum_{j=1}^M \lambda_j \int_0^\infty T_j(u)\overline{F}_j(u)du.$$

Recalling the work-conservation law of Proposition 1 we get

$$\int_0^\infty b(t)dt = \frac{\sum_{j=1}^M \lambda_j E[X_j^2]}{2(1-\rho)^2}. \tag{16}$$

Finally, combining equations (14) and (16) we obtain the desired expression for the asymptotic bias. ■

From Theorem 3 it is clear that depending on the weights g_k, $k \in \mathcal{N}$, the bias can be either positive or negative. It is easy to see that

$$\lim_{t\to\infty}(T_k(t) - T_i(t)) = \frac{g_i - g_k}{2(1-\rho)^2}\sum_{j=1}^M \frac{\lambda_j}{g_j}E[X_j^2]. \tag{17}$$

From equation (17) we observe that the bias for the class with the maximum weight is negative and that of the class with the minimum weight is positive. The vector of weights $\{g_j > 0; j = 1, \ldots, M\}$, may be chosen so that all but one class with the smallest (largest) weight have positive (negative) biases.

The PS model provides a reference model for comparisons with DPS, since the PS system reflects the system in the absence of priorities. Hence, the value of the bias can be considered as a measure of the improvement/degradation that large jobs experience with respect to egalitarian PS. For example, let us consider the case of only two job classes and. Let us assume that class 1 has an extremely large second moment (in comparison to class 2). Then, it is easy to see that if we give priority to class 1 ($g_1 > g_2$), the large jobs of class 2 will suffer a lot. On the other hand, if we give priority to class 2 ($g_1 < g_2$), large jobs of class 2 will gain substantially. This further reinforces the arguments in [9] that classes with typically short jobs should be assigned larger weights. See Sections VII and VIII for further discussion on the comparison of DPS and PS.

We note that when all the second moments of the service requirement distributions are finite, the value of the bias of class k does not depend on the value of the second moments of classes in $\mathcal{S}_k$. This suggests that Theorem 3 may still hold even when the second moment of some class belonging to $\mathcal{S}_k$ is infinite.

VII. Comparison of DPS and PS

As noted earlier, DPS is a generalization of egalitarian PS. Hence, a natural question is how the weights of the DPS system must be chosen so that DPS outperforms PS. Thus, in this section we compare the overall performance of DPS with respect to PS ($E[T^{PS}]$) and we provide guidelines as to how the weights should be chosen

In the context of single-class queues, it is well known that giving preferential treatment to short jobs reduces the overall expected unconditional response time of the system [23], [24], [25], [26], [27]. Hence, one would expect that in order for DPS to outperform PS, the DPS should give preferential treatment to the classes with smaller mean. We will show that this is indeed the case.

We take the expected unconditional response time as the metric of choice for the overall system performance. Let $E[T^{DPS}]$ and $E[T^{PS}]$ be the expected unconditional response time in the DPS and PS systems respectively. Then

$$E[T^{DPS}] = \sum_{k=1}^M \frac{\lambda_k}{\lambda}E[T_k^{DPS}],$$

and

$$E[T^{PS}] = \sum_{k=1}^{M} \frac{\lambda_k}{\lambda} E[T_k^{PS}] = \frac{E[X]}{1-\rho},$$

where $E[T_k^{DPS}]$ and $E[T_k^{PS}]$ denote the expected unconditional response times of class-k jobs in the DPS and PS system respectively, for all $k \in \mathcal{N}$. Note that by Little's law, reducing the expected unconditional response time is equivalent to reducing the mean total number of jobs in the system.

We assume that all the service requirement distributions are exponential with means μ_k^{-1}, for $k \in \mathcal{N}$. We re-index the classes such that $\mu_1^{-1} \leq \ldots \leq \mu_M^{-1}$. In this section we show that if the weights of DPS are chosen as $g_1 \geq g_2 \geq \ldots \geq g_M$, then $E[T^{DPS}] \leq E[T^{PS}]$. Note that this result can be seen as a multi-class counterpart of the classical single-class results on age-based scheduling.

The next lemma will be useful for the proof of the main result in this section. We first introduce some further notation. Let $\overline{U}_k^{DPS}$ and $\overline{U}_k^{PS}$ be the expected unfinished work of class $k \in \mathcal{N}$ in the DPS and PS systems respectively. Note, that by the memoryless property of exponential distributions and Little's law the following relations hold

$$\overline{U}_j^{\pi} = E[N_j^{\pi}]\mu_j^{-1} = \rho_j E[T_j^{\pi}],$$

where $E[N_j^{\pi}]$ denotes the average number of class-j jobs in the system and $\pi \in \{\text{PS, DPS}\}$.

In the ensuing analysis, we require the expected unconditional response times of the various classes in the DPS system satisfy the following condition.

Condition 1. There exists a $j^* \in \mathcal{N}$ such that for all $x \geq 0$,

$$\begin{cases} T_k^{DPS}(x) \leq T_k^{PS}(x) & \text{for all } k < j^*, \\ T_k^{DPS}(x) \geq T_k^{PS}(x) & \text{for all } k \geq j^*. \end{cases}$$

We note that the mean conditional response time under PS is equal for all classes, that is, $T_i^{PS}(x) = T_l^{PS}(x) = x/(1-\rho)$. Thus, Condition 1 states that the curve $x/(1-\rho)$ splits the set of classes into two groups, those that under DPS obtain a better service than under PS and viceversa. Through numerical analysis, we have verified that Condition 1 does not hold in general. Nevertheless, Condition 1 does not seem very restrictive under the regularity condition of Theorem 4, that is when $\mu_1^{-1} \leq \ldots \leq \mu_M^{-1}$ and $g_1 \geq g_2 \geq \ldots \geq g_M$. Under this additional assumption, we have not found any counterexample. See for example Figures 2 and 3.

Lemma 2: Assume Condition 1 is satisfied for an index $j^* \in \mathcal{N}$. Let $\Delta_k = \overline{U}_k^{DPS} - \overline{U}_k^{PS}$. Then,

$$\begin{cases} \Delta_k \leq 0 & \text{for all } k < j^*, \\ \Delta_k \geq 0 & \text{for all } k \geq j^*. \end{cases}$$

Proof: The proof is straightforward after noting that $\Delta_i = \lambda_i \int_0^\infty \left(T_i^{DPS}(x) - T_i^{PS}(x)\right) \overline{F}_i(x)dx$, for all $i = 1, \ldots, M$. ■

Now we show the main result of this section.

Theorem 4: Let the weights in DPS be chosen in decreasing order with respect to the mean job size, that is, $g_1 \geq g_2 \geq \ldots \geq g_M$ and $\mu_1^{-1} \leq \ldots \leq \mu_M^{-1}$. In addition let Condition 1 hold. Then

$$E[T^{DPS}] \leq E[T^{PS}].$$

Proof: We write the difference mean delay in terms of the difference in unfinished work, using the memoryless property of the exponential distribution and Little's law,

$$\begin{aligned} &E[T^{DPS}] - E[T^{PS}] \\ &= \sum_{k=1}^{M} \frac{\lambda_k}{\lambda} E[T_k^{DPS}] - \sum_{k=1}^{M} \frac{\lambda_k}{\lambda} E[T_k^{DPS}] \\ &= \frac{1}{\lambda} \sum_{k=1}^{M} \mu_k (\overline{U}_k^{DPS} - \overline{U}_k^{PS}) = \frac{1}{\lambda} \sum_{k=1}^{M} \mu_k \Delta_k. \end{aligned} \quad (18)$$

Hence, it suffices to show that $\sum_{k=1}^{M} \mu_k \Delta_k \leq 0$. From Lemma 2, let $j^* \in \mathcal{N}$ be such that $\Delta_i \leq 0$ for all $i < j^*$ and $\Delta_i \geq 0$ for all $i \geq j^*$.

We note that by the work-conservation law (see Section IV), $\sum_{k=1}^{M} \overline{U}_k^{PS} = \sum_{k=1}^{M} \overline{U}_k^{DPS}$, and therefore

$$\sum_{k=1}^{M} \Delta_k = 0. \quad (19)$$

Note further that by the conservation law necessarily $j^* > 1$. Then we have

$$\begin{aligned} \sum_{k=1}^{M} \mu_k \Delta_k &= \sum_{k=1}^{j^*-1} \mu_k \Delta_k + \sum_{k=j^*}^{M} \mu_k \Delta_k \\ &\leq \mu_{j^*-1} \sum_{k=1}^{j^*-1} \Delta_k + \mu_{j^*} \sum_{k=j^*}^{M} \Delta_k. \end{aligned}$$

By equation (19)

$$\sum_{k=1}^{j^*-1} \Delta_k = -\sum_{k=j^*}^{M} \Delta_k,$$

and hence

$$\sum_{k=1}^{M} \mu_k \Delta_k \quad \leq \quad (\mu_{j^*} - \mu_{j^*-1}) \sum_{k=j^*}^{M} \Delta_k \leq 0,$$

which completes the proof of the theorem. ■

From equations (18) and (19), we note that, as expected, if $\mu_i^{-1} = \mu_j^{-1}$ for all $i, j \in \mathcal{N}$, then $E[T^{DPS}] = E[T^{PS}]$ independent of the weights. The counterpart of Theorem 4 can be proved along the same lines, here we only state the result for completeness.

Theorem 5: Let the weights in DPS be chosen in increasing order with respect to the mean job size, that is, $g_1 \geq g_2 \geq \ldots \geq g_M$ and $\mu_1^{-1} \geq \ldots \geq \mu_M^{-1}$. In addition let Condition 1 hold. Then

$$E[T^{DPS}] \geq E[T^{PS}].$$

In the particular case of two classes Condition 1 is not necessary. It was shown in [1] that

$$E[T_1^{DPS}] = \frac{\mu_1^{-1}}{1-\rho}\left(1 + \frac{\rho_2(g_2 - g_1)}{\mu_1^{-1} D}\right)$$
$$E[T_2^{DPS}] = \frac{\mu_2^{-1}}{1-\rho}\left(1 + \frac{\rho_1(g_1 - g_2)}{\mu_2^{-1} D}\right),$$

where $D = \frac{g_1(1-\rho_1)}{\mu_1^{-1}} + \frac{g_2(1-\rho_2)}{\mu_2^{-1}}$, and where $\rho_i = \lambda_i \mu_i^{-1}$, $i = 1, 2$. Next, we determine the difference

$$\begin{aligned}
&E[T^{DPS}] - E[T^{PS}] \\
&= -\frac{\lambda_1\mu_1^{-1} + \lambda_2\mu_2^{-1}}{\lambda(1-\rho)} + \frac{\rho_1}{\lambda(1-\rho)}\left(1 + \frac{\rho_2(g_2 - g_1)}{\mu_1^{-1} D}\right) \\
&+ \frac{\rho_2}{\lambda(1-\rho)}\left(1 + \frac{\rho_1(g_1 - g_2)}{\mu_2^{-1} D}\right) \\
&= \frac{-\rho_1\rho_2}{\lambda(1-\rho)D}(g_1 - g_2)(\mu_1 - \mu_2).
\end{aligned}$$

Hence, we note that if $\mu_1 \geq \mu_2$ and $g_1 \geq g_2$, then $E[T^{DPS}] \leq E[T^{PS}]$.

VIII. NUMERICAL EXAMPLES

In this section, we further support our discussion of appropriate weight setting by numerical experiments. Our main tool is equation (4). Let us consider a DPS system with three classes. The mean service times are given by $\mu_1^{-1} = 5000$, $\mu_2^{-1} = 20$ and $\mu_3^{-1} = 2$; the arrival rates are $\lambda_1 = 0.0001$, $\lambda_2 = 0.008$ and $\lambda_3 = 0.1$. Hence the partial loads are $\rho_1 = 0.5$, $\rho_2 = 0.16$ $\rho_1 = 0.1$. Recall that the second moment of an exponentially distributed random variable equals twice the square of the mean. Thus, the second moment of class 1 is much larger than that of classes 2 and 3. In Figure 1 we plot the respective complementary distributions.

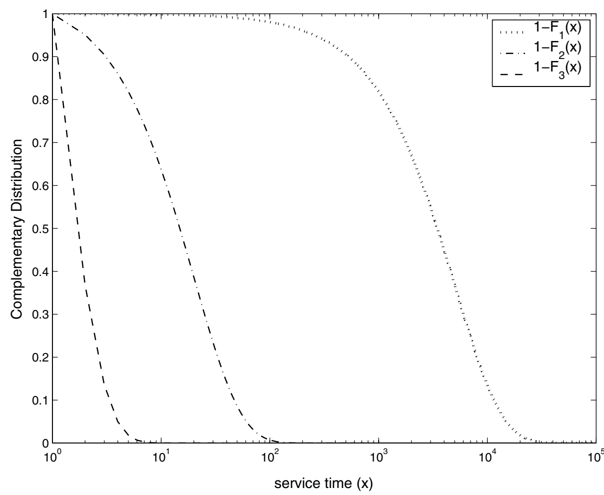

Fig. 1. Complementary distributions of the three classes.

In Figures 2 and 3 we plot the mean conditional response time for all three classes for different choices of the weights. In Figure 2 we choose the weights such that ($g_1 << g_2 << g_3$) and in Figure 3 we choose the weights such that ($g_1 >> g_2 >> g_3$).

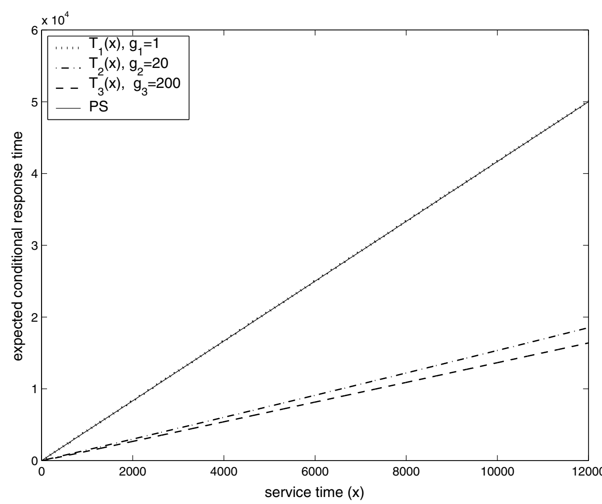

Fig. 2. Comparison of the expected conditional response times under DPS and PS.

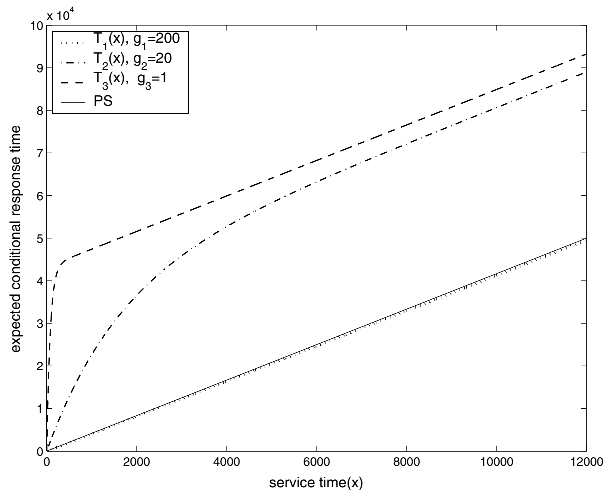

Fig. 3. Comparison of the expected conditional response times under DPS and PS.

As Theorem 3 predicts, the bias of class 1 in Figure 2 is very small and positive, whereas the bias for classes 2 and 3 is negative and large in absolute value. Thus, the performance class-1 jobs perceive is similar to that in absence of discrimination, whereas for classes 2 and 3 the performance is significantly improved. From Theorem 4 we know that the set of parameters of Figure 2 will lead to a decrease of the expected unconditional response time of the DPS system (with respect to PS). In this particular example the expected unconditional response times are $E[T^{DPS}] = 38.3$ and $E[T^{PS}] = 56.8$, that is, with DPS the expected unconditional response time is reduced by 33%.

On the other hand, as we see from Figure 3, the bias for class 1 is negative and small (in absolute value). Hence, class-1 jobs will hardly notice any reduction in the expected response time (note further that mean of class 1 is 5000). On the other hand, by giving preference to class 1 we inflict an important degradation to classes 2 and 3. It is interesting to see that with this choice of weights not only classes 2 and 3 suffer, but the overall system performance is severely degraded. For this particular set of parameters, the expected unconditional response time in DPS is $E[T^{DPS}] = 2085$.

In view of Theorem 3, we conjecture that the phenomena pointed out here will be even more dramatic in the case of classes with infinite second moment.

IX. Conclusion

In this paper we have analyzed a DPS system. The DPS model provides a natural framework for the characterization of multi-class time-sharing systems. We have demonstrated several important properties of DPS.

We have shown that in a stable DPS system it is sufficient that a class' weight is positive to ensure a finite mean number of jobs of that class. This property holds regardless of the characteristics of the required service time distributions. It implies that a DPS system will outperform any strict priority policy if the second moment of one class' service requirement is infinite.

We have proved a new work-conservation law for single server multi-class systems. Existing work-conservation laws have proven to be highly valuable for performance analysis and system optimization. The new conservation law opens up new research lines to further understand these issues in DPS systems, as well as in similar multi-class single server queues, such as Generalized Processor Sharing.

Our closed-form expression for the asymptote points out an insightful dependence of the DPS model with respect to the second moment of the distribution of the required service times. This plays an important role in optimization of weight setting. dependence might be useful for example in determining optimal choice of weights.

Finally, we have studied the impact of the weights on overall system performance, in the context of size based scheduling disciplines. In the case of exponential service requirement distributions, we provide guidelines as to how the weights should be chosen such that DPS outperforms PS.

References

[1] G. Fayolle, I. Mitrani, and R. Iasnogorodski, "Sharing a processor among many job classes," *J. ACM*, vol. 27, no. 3, pp. 519–532, 1980.

[2] L. Kleinrock, "Time-shared systems: A theoretical treatment," *J.ACM*, vol. 14, no. 2, pp. 242–261, 1967.

[3] E. Altman, T. Jimenez, and D. Kofman, "DPS queues with stationary ergodic service times and the performance of TCP in overload," in *Proc. IEEE INFOCOM*, 2004.

[4] T. Bonald and L. Massoulie, "Impact of fairness on internet performance," in *ACM Sigmetrics*, Cambridge, Massachusetts, USA, 2001.

[5] S.C. Borst, D.T.M.B. Van Ooteghem, and A.P. Zwart, "Tail asymptotics for discriminatory processor-sharing queues with heavy-tailed service requirements," Tech. Rep. 2003-25, Eindhoven University of Technology, SPOR, 2003, to appear in *Perf. Eval.*

[6] G. van Kessel, R. Núñez-Queija, and S.C. Borst, "Asymptotic regimes for discriminatory processor sharing," *Perf. Eval. Rev.*, vol. 32, no. 2, pp. 44–46, 2004.

[7] T.M. O'Donovan, "Direct solutions of M/G/1 processor-sharing models," *Oper. Res.*, vol. 22, pp. 1232–1235, 1974.

[8] K.M. Rege and B. Sengupta, “Queue length distribution for the discriminatory processor-sharing queue,” *Oper.Res.*, vol. 44, pp. 653–657, 1996.

[9] G. van Kessel, R. Núñez-Queija, and S.C. Borst, “Differentiated bandwidth sharing with disparate flow sizes,” in *Proc. IEEE INFOCOM*, 2005.

[10] K.M. Rege and B. Sengupta, “A decomposition theorem and related results for the discriminatory processor sharing queue,” *Queueing Systems*, vol. 18, pp. 333–351, 1994.

[11] S. Grishechkin, “On a relationship between processor-sharing queues and Crump-Mode-Jagers branching processes,” *Adv. Appl. Prob.*, vol. 24, pp. 653–698, 1992.

[12] L. Guo and I. Matta, “Scheduling flows with unknown sizes: Approximate analysis,” *Perf. Eval. Rev.*, vol. 30, no. 1, pp. 276–277, 2002, Extended version available as a Boston University Technical Report BU-CS-2002-009.

[13] T. Bonald and A. Proutière, “On stochastic bounds for monotonic processor sharing networks,” *Queueing Systems*, vol. 47, pp. 81–106, 2004.

[14] S. Asmussen, *Applied Probability and Queues*, Springer, 2002.

[15] R.W. Wolff, *Stochastic Modeling and the theory of Queues*, Prentice-Hall, 1989.

[16] A.N. Kolmogorov and S.V. Foming, *Introductory Real Analysis*, Dover Publications, 1970.

[17] R.W. Wolff, “Poisson arrivals see time averages,” *Oper. Res.*, vol. 30, pp. 223–231, 1982.

[18] O.J. Boxma, “Workloads and waiting times in single-server systems with multiple customer classes,” *Queueing Systems*, vol. 5, pp. 185–214, 1989.

[19] J. Shanthikumar and D. Yao, “Multiclass queueing systems: Polymatroidal structure and optimal scheduling control,” *Oper. Res.*, vol. 40, no. 2, pp. 293–299, 1992.

[20] T.C. Green and S. Stidham, “Sample-path conservation laws, with application to scheduling queues and fluid systems,” *Queueing Systems*, vol. 36, pp. 175–199, 2000.

[21] F. Baccelli and P. Bremaud, *Elements of queuing theory: Palm martingale calculus and stochastic recurrences*, Springer, 2003.

[22] D. Brumelle, “On the relation between customer and time average in queues,” *J. Appl. Prob.*, vol. 2, pp. 508–520, 1971.

[23] L.E. Schrage, “The queue M/G/1 with feedback to lower priority queues,” *Management Science*, vol. 13, pp. 466–471, 1967.

[24] R. Righter and G. Shanthikumar, “Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures,” *Prob. Eng. Inf. Sc.*, vol. 3, pp. 323–334, 1989.

[25] R. Righter, G. Shanthikumar, and G. Yamazaki, “On extremal service disciplines in single-stage queueing systems,” *J. Appl. Prob.*, vol. 27, pp. 409–416, 1990.

[26] H. Feng and V. Misra, “Mixed scheduling disciplines for network flows,” *Perf. Eval. Rev.*, vol. 31, no. 2, pp. 36–39, 2003.

[27] S. Aalto, U. Ayesta, and E. Nyberg, “Two-level processor sharing scheduling disciplines: Mean delay analysis,” in *Proc. ACM SIGMETRICS*, New York, 2004, pp. 97–105.