

Differentiated Bandwidth Sharing with Disparate Flow Sizes

Gijs van Kessel[†], Rudesindo Núñez-Queija^{†,‡}, Sem Borst^{†,‡,*}

[†]Department of Mathematics & Computer Science, Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands

[‡]CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

*Bell Laboratories, Lucent Technologies, P.O. Box 636, Murray Hill, NJ 07974, USA

Abstract—We consider a multi-class queueing system operating under the Discriminatory Processor-Sharing (DPS) discipline. The DPS discipline provides a natural approach for modeling the flow-level performance of differentiated bandwidth-sharing mechanisms. Motivated by the extreme diversity in flow sizes observed in the Internet, we examine the system performance in an asymptotic regime where the flow dynamics of the various classes occur on separate time scales. Specifically, from the perspective of a given class, the arrival and service completions of some of the competing classes (called *mice*) evolve on an extremely *fast* time scale. In contrast, the flow dynamics of the remaining classes (referred to as *elephants*) occur on a comparatively *slow* time scale. Assuming a strict separation of time scales, we obtain simple explicit expressions for various performance measures of interest, such as the distribution of the numbers of flows, mean delays, and flow throughputs. In particular, the latter performance measures are *insensitive*, in the sense that they only depend on the service requirement distributions through their first moments. Numerical experiments show that the limiting results provide remarkably accurate approximations in certain cases¹.

I. INTRODUCTION

Over the past few years, the Processor-Sharing (PS) discipline has been widely adopted as a convenient paradigm for analyzing the flow-level performance of dynamically interacting TCP transfers [5], [24]. While the PS model provides valuable insights, it critically relies on the assumption that the service capacity is shared in an egalitarian manner. The actual bandwidth shares may however show substantial variation among competing flows with heterogeneous characteristics. For instance, TCP flows that share a common bottleneck link but traverse heterogeneous routes, may experience diverse packet loss rates and round-trip delays. Because of TCP mechanics, these differences result in a significant discrepancy in the bandwidth shares, see for instance [2].

Besides TCP-related effects, the heterogeneity in bandwidth shares may also be due to deliberate service differentiation

among competing flows. As the Internet evolves to support an ever increasing range of services, there is a growing need for some form of service differentiation to satisfy the diverse requirements of heterogeneous applications. The ability to provide different bandwidth shares is arguably one of the most fundamental vehicles for service differentiation [9]. Discriminatory scheduling algorithms, such as Weighted Fair Queueing (WFQ), have been proposed as potential instruments to implement differentiated bandwidth sharing. *Equation-based* and general *Additive-Increase Multiplicative-Decrease* (AIMD) rate control algorithms also provide scope for differentiated bandwidth sharing [15], [16], [26], [31]. Special mechanisms of this sort are the algorithms proposed in [20], [23] for supporting low-priority data transfers by utilizing excess bandwidth only. This presents an extreme case with strict precedence between the low-priority transfers and regular TCP flows.

The *Discriminatory* Processor-Sharing (DPS) discipline provides a natural approach for modeling the flow-level performance of such differentiated bandwidth-sharing mechanisms. DPS is a multi-class extension of the ordinary egalitarian PS policy, where the various classes are assigned positive weight factors. The service capacity is shared among all users present in proportion to the respective class-dependent weight factors. In case all weight factors are equal, the DPS discipline reduces to the familiar egalitarian PS policy. Note that DPS shows some resemblance with the *Generalized* Processor-Sharing (GPS) discipline (or Generalized Head-Of-the-Line PS), where the service capacity is also shared in accordance with class-dependent weight factors. In GPS however, the capacity is not divided among all users present, but distributed across (non-empty) classes (e.g. the users at the head-of-the-line of the various classes), irrespective of the actual number of users present.

The results for DPS in the literature are surprisingly sparse. The DPS discipline was first studied by Kleinrock [21] and O'Donovan [25]. In a seminal paper, Fayolle, Mitrani & Iasnogorodski [14] obtain the conditional mean sojourn times as the solution of a system of integro-differential equations.

¹This work is part of the project EQUANET (End-to-end quality of service in next generation networks) supported by the Dutch Ministry of Economic Affairs via its agency SenterNovem.

For the case of exponentially distributed service requirements, they derive closed-form expressions and also determine the unconditional mean sojourn times from a system of linear equations. Rege & Sengupta [27] prove a decomposition theorem for the conditional sojourn time. They specifically show that the sojourn time of a customer which finds n customers upon arrival can be decomposed into $n + 1$ independent components, which can be characterized as the solution of a system of non-linear integral equations. In a further paper, Rege & Sengupta [28] obtain the moments of the queue length from a system of linear equations for the case of exponentially distributed service requirements. Moreover, they prove a heavy-traffic limit theorem for the joint queue length distribution. In recent work, Avrachenkov *et al.* [3] show that the mean queue lengths of all classes are finite under the usual stability condition, regardless of the higher-order moments of the service requirements. Assuming finite service requirement distributions with finite variance, they further determine for each class the asymptote of the conditional sojourn time distribution.

In the present paper, we investigate the performance of a DPS system in an asymptotic regime where the flow dynamics of the various classes occur on widely separate time scales. Specifically, from the point of view of every individual class, the interarrival times and service requirements of some of the competing classes are extremely *small*. In the limit, a complete separation of time scales arises, and the arrivals and service completions of these classes occur at an infinitely *fast* relative pace. For the remaining classes, the interarrival times and service requirements are comparatively *large*. In the limit, the flow dynamics of these classes then evolve on an infinitely *slow* relative time scale. The disparity in service characteristics reflects the extreme heterogeneity in flow sizes observed in the Internet, with a vast majority of short transfers (‘mice’) and a tiny fraction of huge flows (‘elephants’). The two limiting scenarios resemble the *fluid* and *quasi-stationary* regimes considered in [8], [13].

Assuming a strict separation of time scales, we derive simple closed-form results for various performance measures of interest, such as the distribution of the number of flows, mean delays, and flow throughputs. The performance experienced by every individual class may be characterized by that in a corresponding Processor-Sharing system with a *reduced service speed* and a *random number of permanent customers*. Specifically, the *reduction in service speed* simply corresponds to the total load of the *mice*, and is completely independent of the values of the DPS weights. The *permanent customers* represent the population of *elephants* encountered by a tagged flow of the given class, which is random but nearly static over the course of the sojourn time of an individual flow. Unlike the service speed, the distribution of the number of permanent customers *is* dependent on the specific values of the DPS weights. It is noteworthy that the above-mentioned performance measures are *insensitive*, in the sense that they only depend on the service requirement distributions through their first moments. (Insensitive *bounds* for the queue length

distributions were recently obtained in [8].) Numerical results indicate that the limiting regimes yield surprisingly accurate approximations in certain cases.

The results for the ‘elephants’ convey that flows with relatively large service requirements will hardly benefit from having large DPS weights. At first sight, the ‘ineffectiveness’ of the weights may be perceived as a somewhat discouraging fact. However, some reflection shows that it is essentially inevitable that large flows suffer correspondingly long delays. This corroborates with the finding in [10] that the delay asymptotics in DPS queues with heavy-tailed service requirements are insensitive to the values of the weights. In fact, the insensitivity may also be interpreted as a positive characteristic, as it means that classes with large service requirements will not be significantly harmed by granting large weights to smaller flows, while the results for the ‘mice’ indicate that the latter flows will substantially benefit from that. In particular, it supports the view that low-priority service utilizing left-over bandwidth only is perfectly adequate for bulky transfers, since higher priority will not significantly improve the performance, and only cause harm to smaller flows. This is in agreement with classical results for single-server queues where overall performance is improved by giving preferential treatment to shorter jobs [29], [30].

The phenomenon that the efficacy of the DPS weights varies with the size of the service requirements has some intriguing consequences. Specifically, for fixed load and weight values, the performance of a given class does not necessarily improve when speeding up the flow dynamics of competing classes. This is rather surprising in view of the fact that in PS queues with time-varying capacity, accelerating the rate of the fluctuations usually *does* improve the performance. In particular, the fluid and quasi-stationary regimes mentioned above typically provide optimistic and pessimistic performance estimates in the latter scenario. For a certain parameter range, Bonald & Proutière [8] show that the two limiting regimes provide provable performance bounds. However, for DPS the roles of these regimes may be interchanged, making the fluid regime a pessimistic bound and the quasi-stationary regime an optimistic bound. This can be explained by the fact that the DPS weights are more effective for classes with smaller service requirements. Thus, in case the competing classes have larger DPS weights, ramping up their arrival and service processes may have the side-effect that DPS becomes more effective in favoring them, which may reduce or even nullify the positive effect of the finer time granularity, and negatively affect the class under consideration. Conversely, in case the competing classes have smaller DPS weights, slowing down their dynamics may help DPS to become more effective in favoring the class under consideration, which may limit or even mitigate the negative impact of the coarser time profile.

The remainder of the paper is organized as follows. In Section II we present a detailed model description and review some useful preliminary results. In order to illuminate the key ideas in the simplest possible setting, we start in Section III with the case of two classes. We then extend the results to

a scenario with an arbitrary number of classes in Section IV. In Section V we generalize the results of [28] to phase-type service requirements, and show that the moments of the queue length may be determined from sets of linear equations. In Section VI we use these results to compare the approximations motivated by the limiting regimes with the exact mean queue lengths and conclude the paper with a discussion of the main insights. Validation of DPS as a model for asymmetric bandwidth sharing by simulation is the subject of an on-going study within the project EQUANET.

II. MODEL DESCRIPTION AND PRELIMINARY RESULTS

We consider a single link of unit rate which is offered traffic from K distinct flow classes. Class- k flows arrive as a Poisson process of rate λ_k , and have generally distributed service requirements with mean β_k . Define $\lambda := \sum_{k=1}^K \lambda_k$ as the total arrival rate, and denote by $p_k := \lambda_k/\lambda$ the fraction of class- k flows. Define $\rho_k := \lambda_k \beta_k$ as the traffic intensity of class k , and denote by $\rho := \sum_{k=1}^K \rho_k$ the total traffic intensity. For stability, we assume $\rho < 1$.

The link rate is shared among the flows of the various classes in accordance with the Discriminatory Processor-Sharing (DPS) discipline which operates as follows. There are positive weights w_1, \dots, w_K associated with each of the classes. When there are n_k class- k flows active $k = 1, \dots, K$, each class- j flow receives service at rate $w_j / \sum_{k=1}^K w_k n_k$.

A. Egalitarian Processor Sharing

In the special case when all w_k are equal, and in particular in case $K = 1$, the DPS discipline reduces to the familiar egalitarian Processor-Sharing (PS) discipline. In that case, the joint equilibrium distribution of the numbers of active flows is [11], [18]:

$$\begin{aligned} & \mathbb{P}\{(N_1, \dots, N_K) = (n_1, \dots, n_K)\} \\ &= (1 - \rho) \rho^n \binom{n}{n_1 \dots n_K} \prod_{k=1}^K \left(\frac{\rho_k}{\rho} \right)^{n_k} \\ &= (1 - \rho) n! \prod_{k=1}^K \frac{\rho_k^{n_k}}{n_k!}, \end{aligned}$$

with $n := n_1 + \dots + n_K$. In particular, the marginal equilibrium distribution of the number of active class- k flows is geometric with parameter $\rho_k / (1 - \rho + \rho_k)$:

$$\mathbb{P}\{N_k = n_k\} = \frac{1 - \rho}{1 - \rho + \rho_k} \left(\frac{\rho_k}{1 - \rho + \rho_k} \right)^{n_k},$$

so the mean number of active class- k flows is:

$$\mathbb{E}\{N_k\} = \frac{\rho_k}{1 - \rho}. \quad (1)$$

Using Little's law, the mean delay of class- k flows is:

$$\mathbb{E}\{S_k\} = \frac{\beta_k}{1 - \rho}.$$

Note that the above performance measures are insensitive with respect to the service requirement distributions, in the sense that they only depend on the first moments of these distributions.

B. Discriminatory Processor Sharing

In the general case when not all w_k are equal, the system loses the symmetry properties of the egalitarian PS discipline, and does not admit similar closed-form results as presented above. For phase-type service requirements the moments of the number of active flows and thus the mean delays may however still be determined from a set of equations, as we shall show in Section V.

This set of equations may be used to evaluate the performance measures for essentially any specific system. Note that the class of phase-type distributions lie dense in the class of all positive random variables. However, the implicit nature of the equations renders little qualitative insight in the performance as a function of the system parameters in general and the weight values in particular.

C. Time-scale separation

In order to derive explicit results, we therefore evaluate the performance in an asymptotic regime where the flow dynamics of the various classes occur on widely separate time scales. Formally, we consider a sequence of systems indexed by a scaling factor r which governs the gap between the time scales, and determine the performance as the value of r approaches infinity. Specifically, class- k arrivals occur as a Poisson process of rate $\lambda_k^{(r)} := \lambda_k f_k(r)$, $k = 1, \dots, K$, and the class- k service requirements are distributed as $\mathbb{P}\{B_k^{(r)} < x\} := B_k(f_k(r)x)$ for some fixed distribution function $B_k(\cdot)$ with mean β_k , so that $\beta_k^{(r)} = \int_0^\infty x d\mathbb{P}\{B_k^{(r)} < x\} = \beta_k / f_k(r)$.

Thus $f_k(r)$ represents the time scale associated with class k as a function of r . Note that the traffic intensity of class k is $\rho_k^{(r)} = \lambda_k^{(r)} \beta_k^{(r)} \equiv \rho_k$, independent of r . We assume that $f_{k-1}(r)/f_k(r) \rightarrow \infty$ as $r \rightarrow \infty$, i.e., lower indexed classes operate on faster time scales. Let $N_k^{(r)}$ be the number of class- k flows in the r -th system, assuming it is in statistical equilibrium. Let $S_k^{(r)}$ be the class- k flow delay in the r -th system. Invoking Little's law, the class- k flow throughput may be expressed as $\tau_k^{(r)} := \beta_k^{(r)} / \mathbb{E}\{S_k^{(r)}\} = \rho_k / \mathbb{E}\{N_k^{(r)}\}$.

For phase-type service requirement distributions, it may be shown using the techniques developed in [1] that the limiting joint equilibrium distributions $\mathbb{P}\{(N_1^{(\infty)}, \dots, N_K^{(\infty)}) = (n_1, \dots, n_K)\} := \lim_{r \rightarrow \infty} \mathbb{P}\{(N_1^{(r)}, \dots, N_K^{(r)}) = (n_1, \dots, n_K)\}$ are well-defined. Hence, the limiting conditional equilibrium probabilities $\mathbb{P}\{N_k^{(\infty)}(n_{k+1}, \dots, n_K) = n_k\} := \lim_{r \rightarrow \infty} \mathbb{P}\{N_k^{(r)} = n_k | (N_{k+1}^{(r)}, \dots, N_K^{(r)}) = (n_{k+1}, \dots, n_K)\}$ are well-defined as well. We conjecture that these probabilities are in fact well-defined for all service requirement distributions with finite variance, although the speed of convergence may vary.

From a mathematical perspective, the time scale decomposition falls in the broader framework of perturbation analysis of Markov processes, and specifically relates to the concept of nearly complete decomposability. There exists an extensive body of literature on perturbation analysis of *finite*-state Markov processes. Some key references include the text books [12], [33] and the survey paper [4]. Perturbation analysis of *infinite*-state Markov processes is considered in [1], [7], [22].

III. TWO CLASSES

In order to introduce the central ideas in the simplest possible setting, we first focus on the case of two classes. Note that the DPS mechanics then only depend on the ratio of the weights $w := w_2/w_1$. In the next section we extend the arguments to a scenario with an arbitrary number of user classes.

As the value of r grows large, the gap between the time scales of the two classes widens, and the flow dynamics of class 1 evolve on an increasingly fast time scale compared to class 2. In the limit for $r \rightarrow \infty$, a complete time scale decomposition occurs, and the class-1 dynamics will average out on the relevant time scale for class 2. Thus class 1 takes away a constant service rate ρ_1 , and class 2 behaves as in a standard isolated PS system with reduced capacity $1 - \rho_1$. Specifically, the distribution of the number of active class-2 flows is geometric with parameter $\rho_2/(1 - \rho_1)$:

$$\mathbb{P}\{N_2^{(\infty)} = n_2\} = \frac{1 - \rho}{1 - \rho_1} \left(\frac{\rho_2}{1 - \rho_1} \right)^{n_2}, \quad (2)$$

so that:

$$\mathbb{E}\{N_2^{(\infty)}\} = \frac{\rho_2}{1 - \rho}, \quad (3)$$

and the flow throughput is $\tau_2^{(\infty)} = 1 - \rho$.

The above results reflect a remarkable degree of insensitivity. This is inherited from the fact that in the limiting regime class 2 behaves as in an ordinary isolated PS system which is known to exhibit insensitivity with respect to the service requirement distribution. Moreover, the reduction in capacity is simply the mean load of class 1, irrespective of the weights or the service requirement distributions. As a result, the class-2 performance measures only depend on the service requirement distributions through their first moments, and do not depend on the weight ratio w at all.

We now proceed to evaluate the performance of class 1. The class-2 dynamics will nearly vanish as $r \rightarrow \infty$ on the relevant time scale for class 1, which will reach some sort of statistical equilibrium for a given number of class-2 flows. Thus when there are n_2 active class-2 flows, class 1 behaves as in a standard PS system with wn_2 ‘permanent customers’. The conditional distribution of the number of active class-1 flows is

$$\mathbb{P}\{N_1^{(\infty)}(n_2) = n_1\} = P_{n_1}(\rho_1, wn_2), \quad n_1 = 0, 1, \dots,$$

where $P_n(u, x)$, $n = 0, 1, 2, \dots$ is the distribution of the number of regular customers in a PS system with traffic intensity u and x permanent customers [6], [11], [18]:

$$P_n(u, x) := \frac{\Gamma(n + x + 1)}{\Gamma(n + 1)\Gamma(x + 1)} u^n (1 - u)^{x+1}, \quad (4)$$

and $\Gamma(z+1) := \int_0^\infty y^z e^{-y} dy$ denotes the complete Γ -function. When x is a non-negative integer, $P_n(u, x)$ reduces to the familiar negative binomial distribution. The conditional mean number of class-1 flows is:

$$\mathbb{E}\{N_1^{(\infty)}(n_2)\} = (wn_2 + 1) \frac{\rho_1}{1 - \rho_1}.$$

Using (2), we obtain the unconditional distribution of the number of class-1 flows, and in particular the mean:

$$\mathbb{E}\{N_1^{(\infty)}\} = \left(w \frac{\rho_2}{1 - \rho} + 1 \right) \frac{\rho_1}{1 - \rho_1}, \quad (5)$$

yielding the flow throughput:

$$\tau_1^{(\infty)} = (1 - \rho_1) \left(w \frac{\rho_2}{1 - \rho} + 1 \right)^{-1}.$$

Just like for class 2, the class-1 performance measures only depend on the service requirement distributions through their first moments. This is induced by the fact that in the limiting regime class 1 behaves as in a PS system with permanent customers which retains the insensitivity with respect to the service requirement distribution. However, the class-1 performance measures *do* depend on the weight ratio w , which reflects the fact that the fraction of capacity claimed by the permanent customers depends on the weights.

The above observations suggest that it is advantageous to assign relatively large weights to small flows since their throughput performance will strongly benefit, without large flows being significantly affected. Conversely, it is unwise to allocate large weights to large flows because doing so will not substantially improve their throughput performance, but have major repercussions for smaller flows.

The phenomenon that smaller flows are more sensitive to the weight values than larger flows has some remarkable implications. Let us focus on the performance of a given class in a system with fixed load and weight values. The above results show that the performance of that class approaches that in an ordinary PS system when the competing class operates on an increasingly fast time scale. In contrast, when the competing class evolves on an increasingly slow time scale, the limiting performance of that class may either be better or worse than in a standard PS system, depending on the ratio of the weights being larger or smaller than one. Consequently, the performance of the class under consideration only improves with increasing granularity of the competing class when it has a larger weight. Otherwise, the performance actually degrades with increasing granularity of the competing class. The latter behavior is rather striking in view of the fact that in PS queues with time-varying capacity accelerating the rate of the fluctuations generally *does* improve the performance, see for instance [13]. In DPS queues however, cranking up the flow

dynamics of a competing class with a larger weight may have the side-effect that DPS becomes more effective in favoring it, which may offset or even nullify the positive effect of the finer granularity, and ultimately harm the class under consideration. Likewise, decreasing the pace of a competing class with a smaller weight may help DPS to become more effective in favoring the class under consideration, which may alleviate or even neutralize the negative impact of the coarser time profile.

IV. ARBITRARY NUMBER OF CLASSES

In the previous section we considered a two-class scenario in order to introduce the basic notions. We now extend the results to a scenario with an arbitrary number of user classes.

We first evaluate the performance of class K which has the slowest flow dynamics. As the value of r grows large, the flow dynamics of all the other classes occur on increasingly fast time scales compared to class K . In the limit for $r \rightarrow \infty$, the flow dynamics of all the other classes will completely average out on the relevant time scale for class K . Thus the other classes take away a constant service rate $\sum_{k=1}^{K-1} \rho_k$, and class K evolves as in an ordinary isolated PS system with reduced capacity $1 - \sum_{k=1}^{K-1} \rho_k$. Specifically, the distribution of the number of active class- K flows is geometric with parameter $\rho_K / (1 - \sum_{k=1}^{K-1} \rho_k)$:

$$\mathbb{P}\{N_K^{(\infty)} = n_K\} = \frac{1 - \rho}{1 - \sum_{k=1}^{K-1} \rho_k} \left(\frac{\rho_K}{1 - \sum_{k=1}^{K-1} \rho_k} \right)^{n_K}, \quad (6)$$

so that

$$\mathbb{E}\{N_K^{(\infty)}\} = \frac{\rho_K}{1 - \rho},$$

and the flow throughput is $\tau_K^{(\infty)} = 1 - \rho$.

Next, we turn to the performance of class $K-1$, which has the slowest flow dynamics of all classes but K . As the value of r grows large, the flow dynamics of classes $1, \dots, K-2$ occur on increasingly fast time scales compared to class $K-1$, whereas those of class K evolve on a comparably slow time scale. In the limit for $r \rightarrow \infty$, the flow dynamics of classes $1, \dots, K-2$ will entirely average out, while those of class K will nearly freeze on the relevant time scale for class $K-1$. So classes $1, \dots, K-2$ will claim a constant service rate $\sum_{k=1}^{K-2} \rho_k$, and class $K-1$ will approach equilibrium for a given number of class- K flows. Thus when there are n_K active class- K flows, class $K-1$ behaves as in a standard isolated PS system of reduced capacity $1 - \sum_{k=1}^{K-2} \rho_k$ and $n_K w_K / w_{K-1}$ ‘permanent customers’. The conditional distribution of the number of active flows of class $K-1$ is:

$$\mathbb{P}\{N_{K-1}^{(\infty)}(n_K) = n_{K-1}\} = P_{n_{K-1}}(u, x),$$

where $P(\cdot, \cdot)$ is defined in (4), $x = n_K w_K / w_{K-1}$ and $u = \rho_{K-1} / (1 - \sum_{k=1}^{K-2} \rho_k)$. The conditional mean number of flows of class $K-1$ is:

$$\mathbb{E}\{N_{K-1}^{(\infty)}(n_K)\} = \left(\frac{w_K}{w_{K-1}} n_K + 1 \right) \frac{\rho_{K-1}}{1 - \sum_{k=1}^{K-1} \rho_k}.$$

Using (6), we obtain the unconditional distribution of the number of flows of class $K-1$, and in particular the mean:

$$\mathbb{E}\{N_{K-1}^{(\infty)}\} = \left(\frac{w_K}{w_{K-1}} \frac{\rho_K}{1 - \rho} + 1 \right) \frac{\rho_{K-1}}{1 - \sum_{k=1}^{K-1} \rho_k},$$

yielding the flow throughput:

$$\tau_{K-1}^{(\infty)} = \left(1 - \sum_{k=1}^{K-1} \rho_k \right) \left(\frac{w_K}{w_{K-1}} \frac{\rho_K}{1 - \rho} + 1 \right)^{-1}.$$

The above arguments readily extend to any given class l . As the value of r grows large, the flow dynamics of classes $1, \dots, l-1$ occur on increasingly fast time scales compared to class l , whereas those of classes $l+1, \dots, K$ evolve on comparably slow time scales. In the limit for $r \rightarrow \infty$, the flow dynamics of classes $1, \dots, l-1$ will completely average out, while those of classes $l+1, \dots, K$ will appear pseudo-static on the relevant time scale for class l . So classes $1, \dots, l-1$ will consume a constant service rate $\sum_{k=1}^{l-1} \rho_k$, and class l will converge to some sort of equilibrium for a given population of flows of classes $l+1, \dots, K$. Thus when there are n_k active class- k flows, $k = l+1, \dots, K$, class l evolves as in an ordinary isolated PS system of reduced capacity $1 - \sum_{k=1}^{l-1} \rho_k$ and $\sum_{k=l+1}^K n_k w_k / w_l$ ‘permanent users’. The conditional distribution of the number of active class- l flows is:

$$\mathbb{P}\{N_l^{(\infty)}(n_{l+1}, \dots, n_K) = n_l\} = P_{n_l}(u, x), \quad (7)$$

with $u = \rho_l / (1 - \sum_{k=1}^{l-1} \rho_k)$ and $x = \sum_{k=l+1}^K n_k w_k / w_l$. The conditional mean number of class- l flows is:

$$\mathbb{E}\{N_l^{(\infty)}(n_{l+1}, \dots, n_K)\} = \frac{\rho_l \left(\sum_{k=l+1}^K \frac{w_k}{w_l} n_k + 1 \right)}{1 - \sum_{k=1}^l \rho_k}.$$

Using (6), (7), the unconditional distribution and, in particular, the mean number of class- l flows may be recursively derived for any given class $l = K-1, K-2, \dots, 1$.

As in the two-class case, it is worth observing that the performance metrics of all classes are insensitive in the sense that they only depend on the service requirement distributions through their first moments. In addition, the performance characteristics of a given class l only depend on the weight factors of classes $l+1, \dots, K$ with slower flow dynamics, and

are independent of the weights of classes $1, \dots, l$ with faster flow dynamics.

The above observations offer further testimony that it makes sense to assign relatively large weights to smaller flows, since this will considerably boost their throughput performance, while only having a negligible impact on large flows. It is not worth allocating large weights to large flows, because this will only marginally improve their throughput performance, but inflict major penalties for smaller flows.

V. EXACT ANALYSIS FOR PHASE-TYPE SERVICE REQUIREMENTS

In this section we present an exact analysis of the mean queue lengths and mean sojourn times for the case where all classes have phase-type service requirement distributions. Similar results were recently obtained in [17].

Within each user *class*, we distinguish between users residing in different service phases, and refer to these as belonging to different customer *types*. Denoting the number of phases of the class- k phase-type distribution with m_k , the total number of types is $J := \sum_{k=1}^K m_k$. With slight abuse of terminology, we also refer to a class- i user in the j^{th} service phase as being of type $\sum_{k=1}^{i-1} m_k + j$. We use $k(j)$ to denote the user class to which type- j users belong.

Let p_{0j} be the probability that an arriving user starts as a type- j user, $j = 1, \dots, J$. Thus, $\sum_{j:k(j)=l} p_{0j} = p_l$ and, if $l = k(j)$, p_{0j}/p_l is the probability that a class- l user starts with service phase j . The service phase corresponding to type j has mean duration $1/\mu_j$, and its service weight is g_j . Although the DPS model implies that $g_i = g_j = w_{k(j)}$ if $k(i) = k(j)$, i.e., if types i and j belong to the same user class, this is in fact not necessary for the analysis in the present section. Furthermore, define p_{ij} ($i, j = 1, \dots, J$) as the probability that, after completing its current service phase a type- i user becomes a type- j user. In the DPS model, no transitions are possible between types belonging to different user classes, but this is again not essential for the analysis in the present section. Also, p_{i0} is the probability that a type- i user will leave the system after completing its current service phase. We denote the number of type- j users in the system by N'_j . By definition, $\sum_{i=1}^J p_{0i} = 1$, $\sum_{j=0}^J p_{ij} = 1$, and $\sum_{j:k(j)=l} N'_j = N_l$.

Denoting by \bar{N}' and \bar{n} the vectors $(N'_1, N'_2, \dots, N'_J)$ and $(n_1, n_2, \dots, n_J) \geq \bar{0}$, respectively, the equilibrium distribution $\pi(\bar{n}) := \mathbb{P}\{\bar{N}' = \bar{n}\}$ satisfies, for $\bar{n} > \bar{0}$,

$$\begin{aligned} & \left(\Lambda + \frac{\sum_{i=1}^J g_i n_i \mu_i}{\sum_{i=1}^J g_i n_i} \right) \pi(\bar{n}) \\ &= \sum_{i=1}^J \Lambda p_{0i} \delta_{n_i} \pi(\bar{n} - \bar{e}_i) \end{aligned} \quad (8)$$

$$\begin{aligned} & + \sum_{i=1}^J \frac{g_i (n_i + 1) \mu_i}{\sum_{j=1}^J g_j n_j} \mu_i p_{i0} \pi(\bar{n} + \bar{e}_i) \\ & + \sum_{i=1}^J \sum_{j=1}^J \frac{g_i (n_i + 1) \mu_i p_{ij} \delta_{n_j} \pi(\bar{n} + \bar{e}_i - \bar{e}_j)}{\sum_{k=1}^J g_k n_k + g_i - g_j}, \end{aligned}$$

where $\delta_n = 1$ if $n > 0$, $\delta_n = 0$ otherwise and \bar{e}_i is a vector with i -th component equal to 1 and all other elements equal to 0. It will be notationally convenient to use the following transformation for $\bar{n} \neq \bar{0}$:

$$R(\bar{n}) = \frac{\pi(\bar{n})}{\sum_{j=1}^J n_j g_j},$$

with the convention that $R(\bar{0}) = 0$. Also, let $p(\bar{z})$ and $r(\bar{z})$ denote the generating functions of $\pi(\bar{n})$ and $R(\bar{n})$, respectively, where $\bar{z} = (z_1, \dots, z_J)$ and $|z_i| < 1$ for $i = 1, \dots, J$:

$$\begin{aligned} p(\bar{z}) &= \mathbb{E}\{z_1^{N'_1} \dots z_J^{N'_J}\} \\ &= \sum_{n_1=0}^{\infty} \dots \sum_{n_J=0}^{\infty} z_1^{n_1} \dots z_J^{n_J} \pi(\bar{n}), \\ r(\bar{z}) &= \mathbb{E}\left\{ \frac{z_1^{N'_1} \dots z_J^{N'_J}}{\sum_{i=1}^J N'_i g_i}; \sum_{j=1}^J N'_j > 0 \right\}; \\ &= \sum_{n_1=0}^{\infty} \dots \sum_{n_J=0}^{\infty} z_1^{n_1} \dots z_J^{n_J} R(\bar{n}). \end{aligned}$$

It follows that

$$p(\bar{z}) = \sum_{i=1}^J g_i z_i \frac{\partial r}{\partial z_i} + 1 - \rho. \quad (9)$$

From (8), we obtain the following partial differential equation for $r(\bar{z})$:

$$\begin{aligned} & \Lambda(1 - \rho) \left(1 - \sum_{j=1}^J p_{0j} z_j \right) \\ &= \sum_{i=1}^J \left\{ \mu_i g_i \left(p_{i0} + \sum_{j=1}^J p_{ij} z_j - z_i \right) - \Lambda g_i z_i \left(1 - \sum_{j=1}^J p_{0j} z_j \right) \right\} \frac{\partial r}{\partial z_i}. \end{aligned} \quad (10)$$

The above equation allows us to determine the moments of the queue length distribution by solving systems of linear equations. Define the following partial derivatives of $p(\bar{z})$ and $r(\bar{z})$:

$$\begin{aligned} L_{i_1 \dots i_j}^j &= \left. \frac{\partial^j p(\bar{z})}{\partial z_{i_1} \dots \partial z_{i_j}} \right|_{\bar{z} \rightarrow \bar{1}^-}, \\ R_{i_1 \dots i_j}^j &= \left. \frac{\partial^j r(\bar{z})}{\partial z_{i_1} \dots \partial z_{i_j}} \right|_{\bar{z} \rightarrow \bar{1}^-}. \end{aligned}$$

The following three theorems yield the mean queue lengths of each type.

The next three theorems determine the mean numbers of customers of each type ($L_i^1 = \mathbb{E}\{N_i^1\}$ for type i).

Theorem 1: The j^{th} moment of the queue length at phase i can be expressed in terms of $R_{i_1 \dots i_j}^j$ and $R_{i_1 \dots i_{j+1}}^{j+1}$ as follows:

$$L_{i_1 \dots i_j}^j = \sum_{i=1}^J g_i R_{i_1 \dots i_j i}^{j+1} + \sum_{l=1}^j g_{i_l} R_{i_1 \dots i_j}^j.$$

Proof: The theorem can be directly obtained by differentiating (9) with respect to z_{i_1}, \dots, z_{i_j} and letting $z_i \rightarrow 1$ for all i . ■

Let a_{ij} be the accumulated amount of work (from arrival until departure) received at phase j assuming a start in phase i . Obviously, the a_{ij} are determined by $a_{ij} = \sum_{k=1}^J p_{ik} a_{kj}$, if $i \neq j$, and $a_{ii} = \frac{1}{\mu_i} + \sum_{k=1}^J p_{ik} a_{ki}$.

Theorem 2: The R_j^1 , $j = 1, 2, \dots, J$, are given by

$$R_j^1 = \frac{1}{g_j} \Lambda \sum_{i=1}^J p_{0i} a_{ij}.$$

Proof: It follows from (10) that the R_j^1 satisfy the following set of J equations with J unknowns where $j = 1, \dots, J$,

$$\sum_{l=1}^J g_l \mu_l p_{lj} R_l^1 - \mu_j g_j R_j^1 = -\Lambda p_{0j},$$

which admits a unique solution, given in the theorem.

An alternative argument is as follows. The capacity dedicated to type j is $g_j R_j^1 = \mathbb{E}\{\delta_{N_j'} g_j N_j' / \sum_{i=1}^J g_i N_i^1\}$. This must be equal to the total amount of work requested in phase j per unit of time, which is $\Lambda \sum_{i=1}^J p_{0i} a_{ij}$. ■

Theorem 3: The coefficients R_{ij}^2 may be found from the following set of J^2 linear equations with J^2 unknowns where $i, j = 1, \dots, J$:

$$\begin{aligned} & (\mu_i g_i + \mu_j g_j) R_{ij}^2 \\ &= \sum_{l=1}^J g_l (\Lambda p_{0i} + \mu_l p_{li}) R_{jl}^2 + \Lambda p_{0i} g_j R_j^1 \\ &+ \sum_{l=1}^J g_l (\Lambda p_{0j} + \mu_l p_{lj}) R_{il}^2 + \Lambda p_{0j} g_i R_i^1. \end{aligned} \quad (11)$$

Proof: As in the proof for Theorem 2, the equations may be obtained from (10). ■

In the numerical experiments that we conducted, we found that the set of equations (11) are independent. This means that the coefficients R_{ij}^2 ($i, j = 1, \dots, J$) can be uniquely determined. For a system with two types, it can be analytically proven that the set of equations (11) are independent as long as $\Lambda < \mu$. As a side remark, note that the set of equations can be reduced

to one of $J(J+1)/2$ equations and equally many unknowns by using the fact that $R_{kl}^2 = R_{lk}^2$.

VI. NUMERICAL EXPERIMENTS

We now discuss the computational experiments that we performed to illustrate the results. We investigate the accuracy of the proposed limiting regimes as approximations to the system performance. For the sake of transparency, we focus on a system with two flow classes.

In the first two sets of experiments we examine the impact of the weight factors and the load on the system performance for exponentially distributed service requirements. In the third set of experiments we then investigate the sensitivity of the system performance with respect to the service requirement distributions.


In all of the experiments, we assume $\beta_1 = \beta_2 \equiv 1$, so that $\rho_i \equiv \lambda_i$, $i = 1, 2$, and take $f_1(r) = r$, $f_2(r) = 1$ to describe the time scales associated with the two classes as a function of the scaling factor r . Specifically, class-1 flows arrive as a Poisson process of rate $\lambda_1 r$ and have mean sizes $1/r$, while class-2 arrivals occur as a Poisson process of rate λ_2 and generate mean service requests of size 1. Thus the scaling factor r corresponds to the time scale of class 1, whereas the time scale of class 2 remains fixed.

For the sake of presentation, we focus on the performance of class 1, and consider both the limiting regimes $r \rightarrow \infty$ and $r \downarrow 0$, where the flow dynamics of class 1 occur on increasingly fast and slow time scales, respectively. The performance in these two regimes corresponds to that of classes 1 and 2, respectively, in the scenario $r \rightarrow \infty$ as analyzed in the previous sections. The various experiments show that the limiting results provide accurate approximations, even for parameter values far outside the asymptotic regime. The section will be concluded with a discussion of some useful rules of thumb.


A. Effect of the weight factors

We first examine the effect of the weight factors on the system performance. We specifically consider the mean number of active class-1 flows $\mathbb{E}\{N_1^{(r)}\}$ as a function of the ratio $w = w_2/w_1$ for several values of the scaling factor r . Recall that the mean delay is proportional to the mean number of active flows while the flow throughput is inversely proportional. For the limiting cases ‘ $r = 0$ ’ and ‘ $r = \infty$ ’, it follows from formulas (3) (with the roles of classes 1 and 2 interchanged as explained above) and (5) that $\mathbb{E}\{N_1^{(0)}\} = \rho_1/(1 - \rho)$ and $\mathbb{E}\{N_1^{(\infty)}\} = \rho_1[1 + w\rho_2/(1 - \rho)]/(1 - \rho_1)$, respectively. For $r \in (0, \infty)$, the value of $\mathbb{E}\{N_1^{(r)}\}$ is determined using Theorems 1, 2 and 3.

In Figures 1(a) and 1(b), the mean number of class-1 flows $\mathbb{E}\{N_1^{(r)}\}$ is plotted as a function of the ratio w for several values of the scaling factor r and $\rho_1 = \rho_2 = 0.3$ and $\rho_1 = 0.2, \rho_2 = 0.4$, respectively. The dashed horizontal lines correspond to the limiting case ‘ $r = 0$ ’, where the flow dynamics of class 1 occur on an infinitely slow time scale, and the performance becomes insensitive to the ratio w



(a) $\rho_1 = \rho_2 = 0.3$



(b) $\rho_1 = 0.2, \rho_2 = 0.4$


Fig. 1. Mean number of class-1 flows $\mathbb{E}\{N_1^{(r)}\}$ as a function of w

as noted earlier. Note that even for moderately small values of r the performance of class 1 already appears to be mostly insensitive. The solid slanted lines correspond to the limiting case ‘ $r = \infty$ ’, where the flow dynamics of class 1 occur on an infinitely fast time scale. The performance then becomes highly sensitive to the weights, and the mean number of class-1 flows varies with the ratio w in an affine manner as indicated earlier. The figures further show that for $w = 1$, the performance is completely insensitive to the scaling factor r , which is in agreement with the standard (multi-class) PS system, cf. (1). Note that in general the performance of class 1 may either improve or deteriorate with increasing value of r , depending on whether the ratio w is smaller or larger than 1. This corroborates with our earlier observations in that regard.


B. Impact of the load

We now investigate the impact of the load on the system performance, assuming equal class loads, i.e., $\rho_1 = \rho_2 = \rho/2$. First we study the mean number of active class-1 flows as a function of the load for $r = 0.1$ and several values of the ratio w .

Figure 2(a) shows the mean number of class-1 flows $\mathbb{E}\{N_1^{(r)}\}$ as a function of ρ for $r = 0.1$ and several values of the ratio w . In addition, the scaled version $\frac{1-\rho}{\rho}\mathbb{E}\{N_1^{(r)}\}$ is



(a) $\mathbb{E}\{N_1^{(r)}\}$ as a function of ρ



(b) $\frac{1-\rho}{\rho}\mathbb{E}\{N_1^{(r)}\}$ as a function of ρ

Fig. 2. Mean number of class-1 flows $\mathbb{E}\{N_1^{(r)}\}$ as a function of ρ ; $r = 0.1$

graphed in Figure 2(b). The scaling coefficient $\frac{1-\rho}{\rho}$ is motivated by the ordinary multi-class PS system where $\mathbb{E}\{N_k\} = \frac{\rho_k}{1-\rho}$, cf. (1).

The bold lines in Figure 2(b) correspond to the limiting scenarios ‘ $w = 0$ ’ and ‘ $w = \infty$ ’, where the behavior reduces to that in a strict-priority system. The dashed downward curve corresponds to the case ‘ $w = 0$ ’ where class 1 receives preemptive priority over class 2. Thus, class 1 is not affected by the presence of class 2, and behaves as in a standard isolated PS system with unit capacity, so that $\mathbb{E}\{N_1^{(r)}\} = \rho_1/(1-\rho_1) = \rho/(2-\rho)$. The solid slightly upward curve represents the opposite case ‘ $w = \infty$ ’, where class 2 receives preemptive priority over class 1, and $\mathbb{E}\{N_1^{(r)}\} = \lambda_1^{(r)}[(\rho_1/\mu_1^{(r)} + \rho_2/\mu_2^{(r)})/(1-\rho_1-\rho_2) + 1/\mu_1^{(r)}]/(1-\rho_2) = \rho(20-9\rho)/[(1-\rho)(40-20\rho)]$, see for example [32] Chapter 10. Note that for moderately large values of the ratio w the curves lie relatively close to the solid curve corresponding to the case ‘ $w = \infty$ ’. Since $r = 0.1$ is relatively small, this reflects again the property that the performance of a class with large service requirements is not too sensitive to the weight factors across a wide range of load values. Note that in the limit for $r \downarrow 0$ the performance of class 1 will become completely insensitive to the values of the weights.




Fig. 3. $\frac{1-\rho}{\rho} \mathbb{E}\{N_1^{(r)}\}$ as a function of ρ ; $w = 1/2$

Next, we fix the ratio $w = 1/2$ and study the mean number of class-1 flows as a function of the load for several values of the scaling factor r .

Figure 3 gives the scaled mean number of class-1 flows $\frac{\rho}{1-\rho} \mathbb{E}\{N_1^{(r)}\}$ as a function of ρ for $w = 1/2$ and several values of the scaling factor r .

The dashed horizontal line corresponds to the limiting case ‘ $r = 0$ ’, where the flow dynamics of class 1 occur on an infinitely slow time scale. The performance then becomes insensitive to the ratio w , and it follows from formula (3) (with the roles of classes 1 and 2 interchanged) that $\frac{1-\rho}{\rho} \mathbb{E}\{N_1^{(r)}\} = \frac{\rho_1}{\rho} = \frac{1}{2}$. The solid curve corresponds to the limiting case ‘ $r = \infty$ ’, where the flow dynamics of class 1 occur on an infinitely fast time scale. In that regime, the performance is highly sensitive, and formula (5) implies $\frac{1-\rho}{\rho} \mathbb{E}\{N_1^{(r)}\} = \frac{\rho_1}{\rho} \frac{1-\rho+\rho_2/2}{1-\rho_1} = \frac{1}{2} \frac{4-3\rho}{4-2\rho}$.

C. Sensitivity to the service requirement distribution

Again assuming equal class loads ($\rho_1 = \rho_2 = \rho/2$), we now examine the degree of sensitivity of the system performance with respect to the service requirement distributions. We assume the class-2 service requirements to be exponentially distributed as before and vary the characteristics of the class-1 service requirement distribution. For the latter distribution, we choose from the class of Erlang distributions (low variability, squared coefficient of variation $c^2 \leq 1$), the hyperexponential distribution (high variability, $c^2 \geq 1$) and the exponential distribution (intermediate case, $c^2 = 1$). (The squared coefficient of variation of a random variable X is defined as $c^2 = \text{Var}\{X\}/(\mathbb{E}\{X\})^2$.)

First we consider a relatively lightly-loaded system ($\rho = 0.25$; Tables I–V), and then repeat the experiment for a more heavily-loaded system ($\rho = 0.75$; Tables VI–X). For $\rho = 0.5$, $\rho = 0.9$ and $\rho = 0.95$, we refer to [19].

In almost all cases of Tables I–V, the mean number of class-1 flows is quite insensitive to the service requirement distribution, although less so for larger values of r (faster class-1 flow dynamics). This reinforces our earlier findings that the performance of a fast class is rather sensitive to the

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.15382	0.15465	0.15595	0.15826	0.16025
0.5	0.16426	0.16442	0.16468	0.16495	0.16512
1	0.16667	0.16667	0.16667	0.16667	0.16667
2	0.16793	0.16789	0.1678	0.16772	0.16766
10	0.16895	0.16889	0.16879	0.16868	0.16861

TABLE I

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 1/10$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.14762	0.14804	0.14881	0.15165	0.15761
0.5	0.15953	0.15999	0.16071	0.16200	0.16310
1	0.16667	0.16667	0.16667	0.16667	0.16667
2	0.17251	0.17209	0.17143	0.17073	0.17027
10	0.17807	0.17767	0.17687	0.17611	0.17563

TABLE II

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 1/2$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.14649	0.14673	0.14719	0.14928	0.15667
0.5	0.15766	0.15806	0.15873	0.16041	0.16237
1	0.16667	0.16667	0.16667	0.16667	0.16667
2	0.17650	0.17574	0.17460	0.17311	0.17203
10	0.18926	0.18817	0.18615	0.18424	0.18302

TABLE III

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 1$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.14588	0.14601	0.14626	0.14757	0.15550
0.5	0.15638	0.15665	0.15714	0.15881	0.16180
1	0.16667	0.16667	0.16667	0.16667	0.16667
2	0.18094	0.18001	0.17857	0.17601	0.17379
10	0.20997	0.20708	0.20238	0.19763	0.19452

TABLE IV

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 2$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.14537	0.14540	0.14545	0.14577	0.15117
0.5	0.15512	0.15519	0.15533	0.15606	0.16046
1	0.16667	0.16667	0.16667	0.16667	0.16667
2	0.18783	0.18736	0.18651	0.18335	0.17673
10	0.29512	0.28679	0.27381	0.25071	0.23080

TABLE V

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 10$)

weights, whereas that of a slow class is not too sensitive. The sensitivity increases when $\rho = 0.5$, but still remains limited. For $\rho = 0.75$, however, the sensitivity is considerably amplified, and becomes even stronger for $\rho = 0.9$ and $\rho = 0.95$ [19]. Evidently, for $w \approx 1$, the strict insensitivity to the service requirement distribution of the ordinary multi-class PS system manifests itself, regardless of the load.

A final remarkable observation is that the mean number of class-1 flows is amazingly close to that in a standard multi-class PS system for a large coefficient of variation as well.

D. Rules of thumb

The main insights from the numerical experiments may be summarized as follows.

- If a class has relatively slow dynamics, then its performance is mostly insensitive to the weights, whereas its performance is quite sensitive in case it has fast dynamics. As observed before, this suggests assigning relatively large weights to smaller flows, since this will substantially boost their performance, without seriously affecting the performance of large flows.
- The limiting scenarios ‘ $r = 0$ ’ (‘fluid’ regime) and ‘ $r = \infty$ ’ (‘quasi-stationary’ regime) provide lower and upper bounds for the performance, cf. [8]. *However, it depends on the ratio of the weights whether the fluid regime is a lower bound and the quasi-stationary regime an upper bound, or vice versa.*
- When the weights of the various classes are roughly equal, the performance impact of the service requirement distribution as well as that of the time scale parameter may be neglected, which is in agreement with the standard multi-class PS system.
- At low loads, the system performance is largely insensitive to the service requirement distribution, except for a class with a small weight and fast flow dynamics.

REFERENCES

- [1] Altman, E., Avrachenkov, K.E., Núñez-Queija, R. (2004). Perturbation analysis for denumerable Markov chains with application to queueing models. *Adv. Appl. Prob.* **36**, 839–853.
- [2] Altman, E., Jimenez, T., Kofman, D. (2004). DPS queues with stationary ergodic service times and the performance of TCP in overload. In: *Proc. IEEE Infocom 2004*, Hong Kong.
- [3] Avrachenkov, K.E., Ayesta, U., Brown, P., Núñez-Queija, R. (2005). Discriminatory Processor Sharing revisited. To appear in: *Proc. IEEE Infocom 2005*, Miami.
- [4] Avrachenkov, K.E., Filar, J.A., Haviv, M. (2002). Singular perturbations of Markov chains and decision processes. In: *Handbook of Markov Decision Processes: Methods and Applications*. Feinberg, E.A., Shwartz, A. (Eds.), Kluwer Academic Publishers, Boston, 113–150.
- [5] Ben Fredj, S., Bonald, T., Proutière, A., Régnié, G., Roberts, J.W. (2001). Statistical bandwidth sharing: a study of congestion at the flow level. In: *Proc. Sigcomm 2001*, 111–122.
- [6] Van den Berg, J.L., Boxma, O.J. (1991). The M/G/1 queue with processor sharing and its relation to a feedback queue. *Queueing Systems* **9**, 365–401.
- [7] Bielecki, T.R., Stettner, L. (1998). Ergodic control of singularly perturbed Markov process in discrete time with general state and compact action spaces. *Appl. Math. Opt.* **38**, 261–281.
- [8] Bonald, T., Proutière, A. (2004). On stochastic bounds for monotonic processor sharing networks. *Queueing Systems* **47**, 81–106.

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	1.00276	1.03869	1.095	1.22931	1.36284
0.5	1.39755	1.40916	1.425	1.44891	1.46528
1	1.5	1.5	1.5	1.5	1.5
2	1.55782	1.55142	1.54286	1.53087	1.52316
10	1.60646	1.59485	1.5802	1.55917	1.54581

TABLE VI

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 1/10$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.77845	0.79484	0.825	0.95440	1.29302
0.5	1.22375	1.24372	1.275	1.34961	1.4238
1	1.5	1.5	1.5	1.5	1.5
2	1.74468	1.71746	1.68	1.62132	1.58016
10	2.02041	1.96275	1.88571	1.77784	1.70844

TABLE VII

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 1/2$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.73662	0.74592	0.76364	0.85309	1.2502
0.5	1.1564	1.17254	1.2	1.28611	1.40749
1	1.5	1.5	1.5	1.5	1.5
2	1.89287	1.8549	1.8	1.69764	1.61646
10	2.50462	2.39061	2.23636	2.01802	1.87491

TABLE VIII

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 1$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.71396	0.71892	0.72857	0.78245	1.18575
0.5	1.10988	1.12059	1.14	1.21787	1.39103
1	1.5	1.5	1.5	1.5	1.5
2	2.05249	2.01256	1.95	1.80077	1.6524
10	3.34403	3.13514	2.85	2.41965	2.12488

TABLE IX

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 2$)

w	Service requirement distribution				
	Erl(4) $c^2 = 1/4$	Erl(2) $c^2 = 1/2$	Exp $c^2 = 1$	Hyp(2) $c^2 = 2$	Hyp(2) $c^2 = 10$
0.1	0.6949	0.69595	0.69802	0.71072	0.95197
0.5	1.06331	1.06607	1.07143	1.10136	1.32542
1	1.5	1.5	1.5	1.5	1.5
2	2.30172	2.28352	2.25	2.10622	1.72998
10	6.47244	6.11307	5.55	4.20694	2.8716

TABLE X

MEAN NUMBER OF CLASS-1 FLOWS $\mathbb{E}\{N_1^{(r)}\}$ ($r = 10$)

- [9] Bonald, T., Roberts, J.W. (2000). Performance of bandwidth sharing mechanisms for service differentiation in the Internet. In: *Proc. ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, Monterey CA, USA, 22.1–22.10.
- [10] Borst, S.C., Van Ooteghem, D.T.M.B., Zwart, A.P. (2004). Tail asymptotics for Discriminatory Processor-Sharing queues with heavy-tailed service requirements. (SPOR-Report 2003-25, Eindhoven University of Technology.) To appear in *Perf. Eval.*
- [11] Cohen, J.W. (1979). The multiple phase service network with generalized processor sharing. *Acta Informatica* **12**, 245–284.
- [12] Courtois, P.J. (1977). *Decomposability: Queueing and Computer System Applications*. Academic Press, New York.
- [13] Delcoigne, F., Proutière, A., Régnié, G. (2004). Modeling integration of streaming and data traffic. *Perf. Eval.* **55**, 185–209.
- [14] Fayolle, G., Mitrani, I., Iasnogorodski, R. (1980). Sharing a processor among many job classes. *J. ACM* **27**, 519–532.
- [15] Floyd, S., Handley, M., Padhye, J. (2000). A comparison of equation-based and AIMD congestion control. Preprint.
- [16] Floyd, S., Handley, M., Padhye, J., Widmer, J. (2000). Equation-based congestion control for unicast applications. In: *ACM Comp. Commun. Rev.* **30** (*Proc. ACM SIGCOMM 2000*), 43–56.
- [17] Haviv, M., Van der Wal, J. (2004). Waiting times in queues and in processor sharing systems with relative priorities. Submitted for publication.
- [18] Kelly, F.P. (1979). *Reversibility and Stochastic Networks*. Wiley, Chichester.
- [19] Van Kessel, G. (2004). *Limiting regimes and approximations for Discriminatory Processor Sharing*. MSc Thesis, Eindhoven University of Technology.
- [20] Key, P.B., Massoulié, L., Wang, B. (2004). Emulating low-priority transport at the application layer: a background transfer service. In: *Proc. Sigmetrics / Performance 2004*, 118–129.
- [21] Kleinrock, L. (1967). Time-shared systems: a theoretical treatment. *J. ACM* **14**, 242–261.
- [22] Korolyuk, V.S., Turbin, A.F. (1993). *Mathematical Foundations of the State Lumping of Large Systems*. Kluwer Academic Publishers, Dordrecht.
- [23] Kuzmanovic, A., Knightly, E.W. (2003). TCP-LP: a distributed algorithm for low-priority data transfer. In: *Proc. IEEE Infocom 2003*, San Francisco CA, USA.
- [24] Núñez-Queija, R. (2000). *Processor-Sharing Models for Integrated-Services Networks*. PhD Thesis, Eindhoven University of Technology.
- [25] O'Donovan, T.M. (1974). Direct solutions of M/G/1 processor-sharing models. *Oper. Res.* **22**, 1232–1235.
- [26] Padhye, J., Kurose, J., Towsley, D., Koodli, R. (2000). A model-based TCP-friendly rate control protocol. Preprint.
- [27] Rege, K.M., Sengupta, B. (1994). A decomposition theorem and related results for the discriminatory processor-sharing queue. *Queueing Systems* **18**, 333–351.
- [28] Rege, K.M., Sengupta, B. (1996). Queue length distribution for the discriminatory processor-sharing queue. *Oper. Res.* **44**, 653–657.
- [29] Righter, R., Shanthikumar, G. (1989). Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures. *Prob. Eng. Inf. Sc.* **3**, 323–334.
- [30] Schrage, L.E. (1967). The queue M/G/1 with feedback to lower priority queues. *Mgmt. Sc.* **13**, 466–471.
- [31] Vojnovic, M., Le Boudec, J.Y., Boutremans, C. (2000). Global fairness of additive-increase and multiplicative-decrease with heterogeneous round-trip times. In: *Proc. IEEE Infocom 2000*, Tel Aviv, Israel, 1303–1312.
- [32] Wolff, R.W. (1991). *Stochastic Modeling and the Theory of Queues*. Prentice Hall, Englewood Cliffs.
- [33] Yin, G., Zhang, Q. (1998). *Continuous-Time Markov Chains and Applications: A Singular Perturbation Approach*. Series: Applications of Mathematics, 37, Springer-Verlag, New York.