

User-Level Performance of Channel-Aware Scheduling Algorithms in Wireless Data Networks

Sem Borst^{*,†,‡}

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

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

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

Abstract—Channel-aware scheduling strategies, such as the Proportional Fair algorithm for the CDMA 1xEV-DO system, provide an effective mechanism for improving throughput performance in wireless data networks by exploiting channel fluctuations. The performance of channel-aware scheduling algorithms has mostly been explored at the packet level for a static user population, often assuming infinite backlogs. In the present paper, we focus on the performance at the flow level in a dynamic setting with random finite-size service demands. We show that in certain cases the user-level performance may be evaluated by means of a multi-class Processor-Sharing model where the total service rate varies with the total number of users. The latter model provides explicit formulas for the distribution of the number of active users of the various classes, the mean response times, the blocking probabilities, and the mean throughput. In addition we show that, in the presence of channel variations, greedy, myopic strategies which maximize throughput in a static scenario, may result in sub-optimal throughput performance for a dynamic user configuration and cause potential instability effects.

Index Terms—Channel-aware scheduling, elastic traffic, insensitivity, Processor Sharing, Proportional Fair scheduling, response time, stability, throughput optimization, wireless data networks.

I. INTRODUCTION

Next-generation wireless networks are expected to support a wide variety of data services. Data applications have fundamentally different traffic characteristics and different quality-of-service requirements than traditional voice services, calling for a significant departure from a conventional circuit-switched operation. In particular, the relative delay tolerance of data applications, combined with the bursty activity patterns, opens up the possibility of scheduling transmissions so as to obtain efficiency gains. An especially attractive approach, in fading environments, is to use channel-aware scheduling strategies, such as the Proportional Fair algorithm for the CDMA 1xEV-DO system [4], [11], [20], which harness channel variations so as to improve the throughput performance.

The performance of channel-aware scheduling algorithms has mostly been investigated at the packet level for a static user population, sometimes including packet-scale dynamics [3],

[17], but often assuming infinite backlogs [1], [8], [14], see also [15], [19] for related results. The assumption of a static user population is a reasonable modeling convention because of the separation of time scales: the scheduling algorithms operate at the packet level on which the user population evolves only relatively slowly. However, when examining throughput performance, and in particular comparing the throughput allocation among elastic traffic users under various strategies, it does not seem entirely satisfactory to assume that the user population is independent of the throughput characteristics and the parameter settings of the scheduling algorithm. For example, a scheduling algorithm that provides high throughput to users with favorable channel conditions, will tend to satisfy the service demands of these users sooner. As a result, the algorithm would tend to be left facing a user population with a higher fraction of users with poor channel conditions. Conversely, a scheduling algorithm that grants reasonable throughput to users with poor channel conditions, should to a certain degree benefit from that by seeing fewer of these users.

In order to capture the above interdependence between the scheduling algorithm and the user population, we move away from a static scenario with a fixed ensemble of users to a dynamic setting where elastic traffic users come and go as governed by the arrival and completion of service demands over time. The notion of finite-size service demands additionally allows us to consider user-perceived performance in terms of response times for file transfers for example, as opposed to delays experienced by individual packets. We will show that in certain cases the user-level performance may be evaluated by means of a multi-class Processor-Sharing model where the total service rate varies with the total number of users. The latter model provides explicit formulas for the distribution of the number of active users of the various classes, the mean response times, the blocking probabilities, and the mean throughput.

To put the above observations further into perspective, it is helpful to make a comparison with a situation where the transmission rates are possibly different across users but *constant*

over time. In that case, a standard work conservation argument implies that the ‘amount of work’ in the system (measured in transmission time rather than bits) is the same under any non-idling scheduling rule. In that sense, the throughput allocation among users corresponding to various scheduling strategies will have an impact on the delay characteristics, but no effect on the system throughput in case of finite-size service demands.

The above-mentioned work conservation property does not extend to a situation where the transmission rates vary over time, and it will no longer be the case that any non-idling scheduling strategy yields the same system throughput. As it turns out, it is not so much maximizing the instantaneous throughput in an absolute sense that determines stability then, but serving users at the right time so as to extract the maximum possible gains from the varying channel conditions. In particular, we will show that greedy, myopic strategies which maximize throughput in a static scenario, may result in sub-optimal throughput performance for a dynamic user configuration and cause potential instability phenomena. Of course, (in)stability is to a certain extent a theoretical concept that cannot occur in an actual system due to admission and flow control mechanisms and the inherent finiteness of buffers. However, it is plausible that instability effects will be reflected in poor performance in terms of long delays in practical circumstances as well.

The remainder of the paper is organized as follows. In Section II we recapitulate some relevant results from the literature for a static user population and state some preliminary facts. We extend the model to accommodate a dynamic user configuration in Section III. We describe how in certain symmetric cases the system behavior may be described by means of a multi-class Processor-Sharing model where the total service rate varies with the total number of users. We present exact results for the distribution of the number of active users of the various classes, the mean response times, the blocking probabilities, and the mean throughput. In Section IV we turn the attention to asymmetric scenarios and derive some stochastic majorization properties. We examine stability issues in Section V. In Section VI we discuss some numerical experiments that we conducted to illustrate the results.

II. STATIC USER POPULATION

We first review some relevant results from the literature for a static scenario with a population of M data users served by a single base station. The base station transmits in slots of some fixed duration. In each slot, the base station transmits to exactly one of the users.

We assume that the feasible rates for the various users vary over time according to some stationary discrete-time stochastic process $\{R_1(t), \dots, R_M(t)\}$, with $R_i(t)$ representing the feasible rate for user i in time slot t . In order to estimate the feasible rates, the base station relies on feedback information from the users on the instantaneous rates that can reliably be supported, as is for instance the case in the CDMA 1xEV-DO system (also known as HDR) [4]. The prediction of the feasible rates should be reasonably accurate when the feedback

delay is relatively short compared to the fading frequency. For convenience, we assume that the base station has perfect knowledge of the feasible rate $R_i(t)$ for every user i at the start of slot t , although the results may be extended to account for possible prediction errors.

Let (R_1, \dots, R_M) be a random vector with as distribution the joint stationary distribution of the feasible rates. We focus on the case where the feasible rates (R_1, \dots, R_M) have a discrete distribution on some finite set $\mathcal{R} \subseteq \mathbb{R}_+^M$. Let $p(r)$ be the stationary probability that the instantaneous feasible rate vector is $r \in \mathcal{R}$. With minor modifications, most of the results extend to scenarios with a continuous rate distribution.

Let T_i be the (long-term) throughput received by user i , and let $\mathcal{A} \subseteq \mathbb{R}_+^M$ be the set of achievable throughput vectors.

The next proposition provides a characterization of the set \mathcal{A} [3], [8].

Proposition 2.1: The set of achievable throughput vectors \mathcal{A} may be characterized as

$$\mathcal{A} = \{T \in \mathbb{R}_+^M : z(T) \geq 1\},$$

where $z(T)$ is the optimal value of the linear program

$$\begin{aligned} \max \quad & z \\ \text{sub} \quad & z \leq z_i = \sum_{r \in \mathcal{R}} p(r) x_i(r) r_i / T_i \quad i = 1, \dots, M \\ & \sum_{i=1}^M x_i(r) \leq 1 \quad r \in \mathcal{R} \\ & x_i(r) \geq 0 \quad i = 1, \dots, M, r \in \mathcal{R}. \end{aligned}$$

The variable $x_i(r)$ in the above linear program may be interpreted as the fraction of time slots allocated to user i in which the instantaneous rate vector is r . Thus, the term $\sum_{r \in \mathcal{R}} p(r) x_i(r) r_i$ represents the throughput received by user i , and the variable z_i measures the throughput as a fraction of the target throughput T_i .

The next proposition provides a characterization of the optimal solution of the above linear program based on the complementary slackness conditions [3], [8].

Proposition 2.2: There exists a vector $w^* \in \mathbb{R}_+^M$ such that any optimal solution $x_i^*(r)$ to the above linear program satisfies

$$x_i^*(r) \left[w_i^* r_i - \max_{j=1, \dots, M} w_j^* r_j \right] = 0,$$

for all $i = 1, \dots, M$, $r \in \mathcal{R}$.

The above proposition shows that any feasible (non-dominated) throughput vector can be achieved by some weight-based strategy which allocates time slot t to a user i^* identified as

$$w_{i^*}^* R_i(t) = \max_{j=1, \dots, M} w_j^* R_j(t),$$

augmented with a suitable tie-breaking rule. In particular, any component-wise increasing function of the throughput vector

is maximized by some weight-based strategy.

We now state some simple auxiliary results that will play a crucial role in the further analysis.

Lemma 2.1: Any feasible throughput vector $T \in \mathcal{A}$ satisfies $\sum_{j=1}^M \alpha_j T_j \leq \mathbb{E}\{\max_{j=1,\dots,M} \alpha_j R_j\}$ for any vector $(\alpha_1, \dots, \alpha_M) \in \mathbb{R}_+^M$.

Proof

Note that the throughput function $\sum_{j=1}^M \alpha_j T_j$ is maximized by a weight-based strategy which assigns a weight $w_i = \alpha_i$ to user i (in fact, sample-path wise), and that the optimal value equals $\mathbb{E}\{\max_{j=1,\dots,M} \alpha_j R_j\}$. \square

We now consider a scenario where the distribution of the rate vector is symmetric in the sense that the relative fluctuations in the feasible rates for the various users around the respective time-average values are statistically identical. Specifically, we assume that $R_i \stackrel{d}{=} C_i Y_i Z$, where $C_i := \mathbb{E}\{R_i\}$ is the time-average rate of user i , Y_1, \dots, Y_M are independent and identically distributed copies, and Z represents a possible correlation component with unit mean. Define $G(M) := \mathbb{E}\{\max_{j=1,\dots,M} Y_j\}$.

Lemma 2.2: In the case of a symmetric rate distribution as described above, the weight-based strategy which assigns a weight $w_i = 1/C_i$ to user i , and breaks ties between users at random, provides each user a fraction $G(M)/M$ of its time-average rate.

Proof

Note that user i is selected when $\frac{R_i}{C_i} = \max_{j=1,\dots,M} \frac{R_j}{C_j}$, i.e., $Y_i = \max_{j=1,\dots,M} Y_j$, and possible ties are broken to its advantage. By symmetry considerations, user i thus receives a fraction $1/M$ of the time slots, and the expected rate when selected is

$$\mathbb{E}\{R_i | \frac{R_i}{C_i} = \max_{j=1,\dots,M} \frac{R_j}{C_j}\} = \mathbb{E}\{C_i Y_i Z | Y_i = \max_{j=1,\dots,M} Y_j\} = C_i \mathbb{E}\{Y_i | Y_i = \max_{j=1,\dots,M} Y_j\} = C_i \mathbb{E}\{\max_{j=1,\dots,M} Y_j\} = C_i G(M).$$

\square

Remark 2.1: The assumption that the relative rate fluctuations are statistically identical is roughly valid when the users for example have Rayleigh fading channels and the feasible rate is approximately linear in the SNR (signal-to-noise ratio). The latter approximation is reasonably accurate when the SNR is not too high. It is not necessary that the Doppler frequencies are identical, since only the instantaneous rate distribution affects the long-term average throughput achieved under a weight-based strategy. Of course, the Doppler frequencies do matter for the transient throughput behavior and also affect the ability to predict the feasible rate. Also, the assumption $R_i \stackrel{d}{=} C_i Y_i Z$ could be further relaxed. For instance, a somewhat milder condition would be that

$\mathbb{P}\{(R_1/C_1, \dots, R_M/C_M) \leq (t_{\pi(1)}, \dots, t_{\pi(M)})\}$ is invariant under permutations $\pi(1), \dots, \pi(M)$.

Remark 2.2: In certain cases, the Proportional Fair scheduling algorithm for the CDMA 1xEV-DO system mentioned earlier behaves approximately like a weight-based strategy. In Proportional Fair scheduling, the weights w_i are dynamically adapted and are inversely proportional to the exponentially smoothed throughputs W_i of the users. Thus, the expected rate of user i when selected is

$$\mathbb{E}\{R_i | w_i R_i = \max_{j=1,\dots,M} w_j R_j\} = \mathbb{E}\{R_i | \frac{R_i}{W_i} = \max_{j=1,\dots,M} \frac{R_j}{W_j}\}.$$

Now observe that both the instantaneous rate R_i and the exponentially smoothed throughput W_i scale linearly with the time-average rate C_i in case the relative rate fluctuations are statistically identical. As a result, the allocation of time slots only depends on the relative rate fluctuations and not on the time-average rates. Thus, each user receives a fraction $1/M$ of the time slots, and we may write $W_i \stackrel{d}{=} C_i V_i$, where the random variables V_1, \dots, V_M are identically distributed (but not independent). In addition, the exponentially smoothed throughputs will not show any significant variation when the time constant in the exponential smoothing is large, i.e., $V_1, \dots, V_M \approx V$ for some constant V . Substituting $R_i \stackrel{d}{=} C_i Y_i Z$ and $W_i \approx V C_i$ in the above formula, we find that the expected rate of user i when selected approximately equals $\mathbb{E}\{C_i Y_i Z | Y_i = \max_{j=1,\dots,M} Y_j\} = C_i G(M)$. In conclusion, in case the relative rate fluctuations are statistically identical and the time constant in the exponential smoothing is not too small, the Proportional Fair scheduling algorithm roughly behaves as the weight-based strategy which assigns a constant weight $w_i = 1/C_i$ to user i . We refer to [13], [18] for a rigorous justification of the above claims.

We would like to add that the above statements assume the users to have infinite backlogs. In situations with packet-scale dynamics, the Proportional Fair algorithm may be ill-behaved, and the throughput performance be degraded by convergence and fragmentation issues, giving rise to potential instability phenomena [2].

III. DYNAMIC USER CONFIGURATION

We now extend the model to accommodate a dynamic configuration of users. The user dynamics result from finite-size service demands that arrive randomly over time. We assume that the duration of the time slots is short relative to the size and arrival frequency of the service demands. Thus, the scheduling strategy operates on an extremely fast time scale compared to the user dynamics, making it natural to analyze the user-level performance in continuous rather than discrete time, and assume that the users are served simultaneously rather than in a time-slotted fashion. The continuous-time model naturally inherits its service characteristics from the discrete-time model. Specifically, we assume that the set of feasible service rate vectors in the continuous-time context for a given user population coincides with the set of achievable throughput vectors for that user population in a discrete-time setting.

For the latter model, we consider a scenario where the relative fluctuations in the feasible rates for the various users around the respective time-average values are statistically identical as described in the previous section. Thus, we assume that the instantaneous rate of user i with time-average rate C_i is distributed as $R_i \stackrel{d}{=} C_i Y_i Z$, where Y, Y_1, Y_2, \dots are independent and identically distributed copies and Z represents a possible correlation component with unit mean. According to Lemma 2.2, we then have that under the strategy S^* which assigns a weight $w_i = 1/C_i$ to a user i with a time-average rate C_i , each user is continuously served at a fraction $G(n)/n$ of its time-average rate whenever there are n users in the system.

The above assumptions ignore the discrete nature of the time slots and neglect the transient fluctuations in the throughput. However, the law of large numbers suggests that these effects should be negligible in some suitable asymptotic sense in a limiting regime where the duration of the time slots shrinks relative to the time scale of the user dynamics.

To describe the service demands, we assume that users initiate file transfer requests randomly over time. We consider a scenario with K user classes. Class- k users submit file transfer requests as a Poisson process of rate λ_k . We assume that at most M users in total are admitted into the system simultaneously (possibly $M = \infty$). Users which submit requests when there are already M transfers in progress are denied access and abandon. Let (C_k, F_k) be a pair of random variables with as distribution the joint distribution of the time-average transmission rate and the file size of an arbitrary class- k user. We assume that the file size and time-average transmission rate are independent across users, but we allow for possible dependence between the file size and time-average transmission rate of a given user. Let $B_k := F_k/C_k$ be the normalized service requirement of a class- k user, with mean $\beta_k := \mathbb{E}\{B_k\} = \mathbb{E}\{F_k/C_k\}$. The normalized service requirement is the amount of time it would take to complete the file transfer if a user were the only user in the system. Note that the normalized service requirement encapsulates both the file size and the time-average transmission rate of a user, and is measured in transmission time rather than data volume. Define $\rho_k := \lambda_k \beta_k$ as the offered traffic associated with class- k users. Denote by $\rho := \sum_{k=1}^K \rho_k$ the total amount of offered traffic. Let B_k^r be a random variable representing the residual lifetime of B_k and $B_k^r(\cdot)$ the associated distribution function, i.e., $B_k^r(x) := \mathbb{P}\{B_k^r < x\} := \frac{1}{\beta_k} \int_{y=0}^x \mathbb{P}\{B_k > y\} dy$.

Let (N_1, \dots, N_K) be a random vector representing the number of users of the various classes in the system under strategy S^* at an arbitrary epoch in statistical equilibrium (assuming it exists). Denote by $N := N_1 + \dots + N_K$ the total number of users in the system. Given that there are n_k class- k users in the system, let $B_{k,i}^r$ be the remaining normalized service requirement of the i -th class- k user, $i = 1, \dots, n_k$, $k = 1, \dots, K$. Define $G^* := \sup_{M=1,2,\dots} G(M) = \lim_{M \rightarrow \infty} G(M)$. Note that $G^* = \infty$ when the distribution of Y has infinite support.

Proposition 3.1: Strategy S^* achieves stability for $\rho < G^*$ or $M < \infty$, in which case

$$\mathbb{P}\{N_k = n_k, B_{k,i}^r \leq t_{k,i}; i = 1, \dots, n_k, k = 1, \dots, K\} =$$

$$H^{-1} \frac{n! \rho^n}{\phi(n)} \prod_{k=1}^K \frac{1}{n_k!} \left(\frac{\rho_k}{\rho} \right)^{n_k} \prod_{i=1}^{n_k} B_{k,i}^r(t_{k,i}),$$

with $n = n_1 + \dots + n_K \leq M$, $\phi(n) := \prod_{i=1}^n G(i)$, and normalization constant

$$H := \sum_{n=0}^M \frac{\rho^n}{\phi(n)}.$$

In particular,

$$\mathbb{P}\{N = n\} = H^{-1} \frac{\rho^n}{\phi(n)},$$

$$\mathbb{E}\{N\} = H^{-1} \sum_{n=1}^M \frac{n \rho^n}{\phi(n)},$$

and

$$\mathbb{E}\{N_k\} = \frac{\rho_k}{\rho} \mathbb{E}\{N\}.$$

The blocking probability is given by

$$L = \mathbb{P}\{N = M\}.$$

Proof

According to Lemma 2.2, each user is served at a fraction $G(n)/n$ of its time-average rate whenever there are n users in the system. Thus, the normalized remaining service requirement of each user is reduced at rate $G(n)/n$, which means that the normalized remaining service requirements evolve in a similar probabilistic fashion as the remaining service requirements in a multi-class Processor-Sharing system with arrival rates λ_k , generic service requirements B_k , and service rate $G(n)$ when there are n users in total present. The statements then follow from results for the latter system [9], [12]. \square

Remark 3.1: Proposition 3.1 extends to the case where users generate sessions consisting of multiple file requests separated by arbitrarily distributed ‘think times’ [5], [7]. In that case, the offered traffic should be calculated so as to include the mean number of file requests per session.

Using Little’s law, we find that the mean transfer delay experienced by a class- k user is given by

$$\mathbb{E}\{S_k\} = \frac{\beta_k}{\rho(1-L)} \mathbb{E}\{N\}.$$

The above formula reflects the celebrated insensitivity property of the Processor-Sharing discipline, which shows that the mean delay of a class- k user only depends on the service requirement distribution of class k through its mean β_k . In fact, it may be shown that the conditional expected delay of any user with actual service requirement b is given by

$$\mathbb{E}\{S|B = b\} = \frac{b}{\rho(1-L)} \mathbb{E}\{N\}.$$

Thus, the expected transfer delay incurred by a user is proportional to its normalized service requirement, with factor of proportionality $\mathbb{E}\{N\}/(\rho(1-L))$. The latter property embodies a certain fairness principle, which means that users with larger service requirements tend to experience longer delays. Recall that the normalized service requirement encapsulates both the file size and the time-average transmission of a user, and is expressed in time units rather than data bits.

Proposition 3.2: No strategy achieves stability for $\rho > G^*$.

Proof

Define the normalized amount of work as the sum of the normalized remaining service requirements of all the users in the system. Let A_m and B_m be the arrival epoch and the normalized service requirement of the m -th arriving user, let X_m be the normalized amount of work in the system at time $t = A_m$, and let D_m be the reduction in the normalized amount of work between time epochs A_m and A_{m+1} . According to Lemma 2.1, taking $\alpha_i = 1/C_i$, no strategy is able to reduce the normalized amount of work at a rate higher than $G(M) \leq G^*$ when there are M users present. Hence,

$X_{m+1} = X_m + B_m - D_m \geq X_m + B_m - G^*(A_{m+1} - A_m)$, so that when $\rho > G^*$,

$$\begin{aligned} \mathbb{E}\{X_{m+1}\} &\geq \mathbb{E}\{X_m\} + \sum_{k=1}^K \frac{\lambda_k}{\lambda} \beta_k - \frac{1}{\lambda} G^* \\ &= \mathbb{E}\{X_m\} + \frac{1}{\lambda} [\rho - G^*] > \mathbb{E}\{X_m\}, \end{aligned}$$

with $\lambda := \sum_{k=1}^K \lambda_k$. Thus, the normalized workload process has positive drift when $\rho > G^*$ for any strategy. \square

Propositions 3.1 and 3.2 combined imply that strategy S^* achieves stability whenever feasible. The heuristic explanation is that the rate at which strategy S^* reduces the normalized amount of work will approach the maximum possible value G^* as the number of users tends to infinity. In fact, the proof of Proposition 3.2 shows that strategy S^* reduces the normalized amount of work at a higher rate than any other strategy, given the same number of users. (It is thus tempting to conjecture that strategy S^* actually minimizes the normalized amount of work among all strategies, but that does not appear to be true without further assumptions.) In particular, a weight-based strategy which assigns a weight $F(C)/C$ to a user with a time-average rate C reduces the normalized amount of work at a rate $\sum_{i=1}^M \mathbb{E}\{Y_i I_{\{F(C_i)Y_i = \max_{j=1,\dots,M} F(C_j)Y_j\}}\}$ when there are M users with time-average rates C_1, \dots, C_M . In general, there is no guarantee that the latter quantity under any circumstances approaches G^* when M tends to infinity. Intuitively, unless the weights are set inversely proportional to the time-average transmission rates, the relative rate fluctuations are not maximally exploited. We will examine these issues further in Section V.

Remark 3.2: As mentioned in Section II, strategy S^* may be viewed as a proxy for the Proportional Fair scheduling algorithm in case the relative rate fluctuations are statistically identical and the time constant in the exponential smoothing is not too small. The latter statement assumed a static user population with infinite backlogs. With a dynamic user configuration, we need to assume that the duration of the time slots is relatively short compared to the backlog periods of the users, so that the throughput performance of the Proportional Fair algorithm is not substantially hampered by convergence or granularity issues. Otherwise, when the weights are initialized to zero, the algorithm may allocate time slots to arriving users almost regardless of the channel conditions, and thus fail to extract the maximum gains from the channel variations. The Proportional Fair algorithm may then result in sub-optimal throughput performance and potentially collapse into instability.

IV. ASYMMETRIC SCENARIOS

In the previous section we considered a scenario with K user classes where the relative rate fluctuations in the feasible rates are statistically identical for all users. We assumed that the system is operated according to the weight-based strategy S^* which assigns a weight $w_i = 1/C_i$ to a user i with a time-average transmission rate C_i .

We now consider a scenario where the relative fluctuations in the feasible rates around the respective time-average values for all users of a given class are statistically identical as before. However, we allow for the distributions of the fluctuations to vary across user classes. Thus, we assume that the instantaneous rate of a class- k user i is distributed as $R_{k,i} \stackrel{d}{=} C_{k,i} Y_{ki} Z$, where $Y_k, Y_{k1}, Y_{k2}, \dots$ are independent and identically distributed copies and Z represents a possible correlation component with unit mean.

The system is operated using a weight-based strategy S^α which assigns a weight $w_{k,i} = \alpha_k/C_{k,i}$ to a class- k user i with a time-average rate $C_{k,i}$. The parameters α_k allow for differentiation among the various user classes. The differentiation could be based on channel statistics, traffic characteristics, or Quality-of-Service requirements.

With the heterogeneous user classes, the system loses the symmetry properties of the ordinary Processor-Sharing discipline which facilitated the analysis in the previous section. In fact, asymmetric (discriminatory) versions of the Processor-Sharing discipline have remained largely intractable so far, even under exponentiality assumptions and when the service rates are constant [10], [16]. Therefore, we will not aim for full distributional results but focus on stochastic majorization properties and stability issues.

Note that strategy S^α allocates a time slot to a class- k user i when $w_{k,i} R_{k,i} = \max_{l=1,\dots,K} \max_{j=1,\dots,n_l} w_{l,j} R_{l,j}$, i.e., $\alpha_k Y_{ki} = \max_{l=1,\dots,K} \max_{j=1,\dots,n_l} \alpha_l Y_{lj}$. In order to avoid technicalities, we assume that $\mathbb{P}\{\alpha_k Y_k = \alpha_l Y_l\} = 0$ for $k \neq l$, so that there are no tie-breaking issues between user classes. Ties between users from the same class are broken at random.

Let $y_k := \inf\{y : \mathbb{P}\{Y_k > y\} = 0\}$ be the maximum value that Y_k can achieve. We assume that the user classes are

indexed such that $\alpha_1 y_1 > \alpha_2 y_2 > \dots > \alpha_K y_K$. Note that we do not allow for any ties, which implies that when the number of class- k users tends to infinity, classes $k+1, \dots, K$ will be completely starved, and hence be driven unstable as well.

Denote by $G_k^\alpha(n_1, \dots, n_K)$ the total normalized service rate for class k when there are n_l class- l users, $l = 1, \dots, K$. We may write

$$G_k^\alpha(n_1, \dots, n_K) = \mathbb{E}\{\max\{Y_{k1}, \dots, Y_{kn_k}\} \times$$

$$\mathbb{I}_{\{\alpha_k \max\{Y_{k1}, \dots, Y_{kn_k}\} \geq \alpha_l \max\{Y_{l1}, \dots, Y_{ln_l}\} \text{ for all } l=1, \dots, K\}}\}.$$

For any $1 \leq k \leq m \leq K$, define $G_k^\alpha(n_1, \dots, n_m) := \inf_{n_{m+1}, \dots, n_K=0,1,2,\dots} G_k^\alpha(n_1, \dots, n_K)$. The quantity $G_k^\alpha(n_1, \dots, n_m)$ may be interpreted as the minimum guaranteed total normalized service rate for class k when there are n_l class- l users, $l = 1, \dots, m$, competing against any number of class- p users, $p = m+1, \dots, K$. Also, denote $G_k^\alpha(n_1, \dots, n_{k-1}) := \sup_{n_k=0,1,2,\dots} G_k^\alpha(n_1, \dots, n_k)$. The latter quantity may be interpreted as the maximum of the minimum guaranteed total normalized service rate for class k when there are n_l class- l users, $l = 1, \dots, k-1$.

The next lemma lists some useful properties of the function $G_k^\alpha(n_1, \dots, n_K)$ and various derived quantities.

Lemma 4.1: The function $G_k^\alpha(n_1, \dots, n_K)$ satisfies the following properties:

- (i) For any $1 \leq k \leq m \leq K$, $G_k^\alpha(n_1, \dots, n_m) = \lim_{n_{m+1}, \dots, n_K \rightarrow \infty} G_k^\alpha(n_1, \dots, n_K)$;
- (ii) For any $1 \leq k \leq K$, $G_k^\alpha(n_1, \dots, n_{k-1}) = \lim_{n_k \rightarrow \infty} G_k^\alpha(n_1, \dots, n_k)$;
- (iii) For any $1 \leq k \leq m \leq K$, $G_k^\alpha(n_1, \dots, n_m)$ is increasing in n_k and decreasing in n_l , $l \neq k$, $l = 1, \dots, m$;
- (iv) For any $1 \leq k \leq m \leq K$, $G_k^\alpha(n_1, \dots, n_m)/n_k$ is decreasing in n_l , $l = 1, \dots, m$;
- (v) For any $1 \leq k \leq K$, $G_k^\alpha(n_1, \dots, n_{k-1}) = y_k \mathbb{P}\{\alpha_k y_k \geq \alpha_l \max_{l=1, \dots, k-1} \max\{Y_{l1}, \dots, Y_{ln_l}\}\} = \lim_{n_k \rightarrow \infty} G_k^\alpha(n_1, \dots, n_K) = \sup_{n_k=0,1,2,\dots} G_k^\alpha(n_1, \dots, n_K)$ for all (n_{k+1}, \dots, n_K) ;
- in particular, $G_1^\alpha = y_1 = \lim_{n_1 \rightarrow \infty} G_1^\alpha(n_1, \dots, n_K) = \sup_{n_1=0,1,2,\dots} G_1^\alpha(n_1, \dots, n_K)$ for all (n_2, \dots, n_K) ;
- (vi) The function $H^\alpha(n_1, \dots, n_m) := \sum_{k=1}^m \alpha_k G_k^\alpha(n_1, \dots, n_m)$ is increasing in n_l , $l = 1, \dots, m$.

Proof

To prove Properties (i)-(iii), it suffices to show that $G_k^\alpha(n_1, \dots, n_K)$ is increasing in n_k and decreasing in n_l , $l \neq k$, $l = 1, \dots, m$, which follows immediately from the definition.

To check Property (iv), it is enough to verify that $G_k^\alpha(n_1, \dots, n_K)/n_k$ is decreasing in n_l , $l = 1, \dots, K$. We may write $G_k^\alpha(n_1, \dots, n_K)/n_k$ as

$$\mathbb{E}\{Y_k \mathbb{I}_{\{\alpha_k Y_k \geq \alpha_l \max\{Y_{l1}, \dots, Y_{ln_l}\} \text{ for all } l=1, \dots, K\}}\},$$

which yields the desired statement.

Property (v) follows immediately from the definition combined with Property (ii) and the fact that $\alpha_1 y_1 > \alpha_2 y_2 > \dots > \alpha_K y_K$.

To prove Property (vi), it suffices to show that $\sum_{k=1}^m \alpha_k G_k^\alpha(n_1, \dots, n_K)$ is increasing in n_l , $l = 1, \dots, m$. To do so, we may write

$$\begin{aligned} \sum_{k=1}^m \alpha_k G_k^\alpha(n_1, \dots, n_K) &= \sum_{k=1}^m \alpha_k \mathbb{E}\{\max\{Y_{k1}, \dots, Y_{kn_k}\} \times \\ &\mathbb{I}_{\{\alpha_k \max\{Y_{k1}, \dots, Y_{kn_k}\} \geq \alpha_l \max\{Y_{l1}, \dots, Y_{ln_l}\} \text{ for all } l=1, \dots, K\}}\} = \\ &\sum_{k=1}^m \mathbb{E}\{\alpha_k \max\{Y_{k1}, \dots, Y_{kn_k}\} \times \\ &\mathbb{I}_{\{\alpha_k \max\{Y_{k1}, \dots, Y_{kn_k}\} \geq \alpha_l \max\{Y_{l1}, \dots, Y_{ln_l}\} \text{ for all } l=1, \dots, K\}}\} = \\ &\mathbb{E}\{\max_{k=1, \dots, m} \alpha_k \max\{Y_{k1}, \dots, Y_{kn_k}\}\}, \end{aligned}$$

which yields the desired statement. \square

We now introduce two corresponding ‘restricted’ versions of the system. For any $m = 1, \dots, K$, the m -restricted version is a system with user classes $1, \dots, m$ only. In the m -restricted *integrated* system, each class- k user, $k = 1, \dots, m$, is served at a fraction $G_k^\alpha(n_1, \dots, n_m)/n_k$ of its time-average rate whenever there are n_l class- l users in the system, $l = 1, \dots, m$. In the m -restricted *segregated* system, each class- m user is also served at a fraction $G_m^\alpha(n_1, \dots, n_m)/n_m$ of its time-average service rate, while each class- k user, $k = 1, \dots, m-1$, is served at a fraction $G_k^\alpha(n_1, \dots, n_{m-1})/n_k$ of its time-average rate whenever there are n_l class- l users in the system, $l = 1, \dots, m$. Let $N_k(t)$, $\tilde{N}_k^{(m)}(t)$, and $\hat{N}_k^{(m)}(t)$ be the number of class- k users at time t in the original system, the m -restricted integrated system, and the m -restricted segregated system, $k = 1, \dots, m$, respectively. Note that each class- k user, $k = 1, \dots, m-1$, in the m -restricted segregated system receives the same service rate as in the $(m-1)$ -restricted integrated system. Hence, under equal initial conditions, $(\tilde{N}_1^{(m-1)}(t), \dots, \tilde{N}_{m-1}^{(m-1)}(t)) \stackrel{d}{=} (\hat{N}_1^{(m)}(t), \dots, \hat{N}_{m-1}^{(m)}(t))$ for all $t \geq 0$. Since all users in the m -restricted integrated system are worse off than in the original system, and are yet worse off in the m -restricted segregated system, it is further plausible that, again under equal initial conditions, $(N_1(t), \dots, N_m(t)) \leq_{st} (\tilde{N}_1^{(m)}(t), \dots, \tilde{N}_m^{(m)}(t)) \leq_{st} (\hat{N}_1^{(m)}(t), \dots, \hat{N}_m^{(m)}(t))$ for all $t \geq 0$.

In order to prove the latter stochastic ordering property, we now compare two systems restricted to user classes $1, \dots, m$, labeled I and II, where each class- k user is served at a fraction $G_k^I(n_1, \dots, n_m)/n_k$ and $G_k^{II}(n_1, \dots, n_m)/n_k$ of its time-average rate, $k = 1, \dots, m$, respectively. We assume that $G_k^I(n_1, \dots, n_m) \geq G_k^{II}(n_1, \dots, n_m)$ for all (n_1, \dots, n_m) and that $G_k^I(n_1, \dots, n_m)/n_k$ is decreasing in n_l for all $k, l = 1, \dots, m$. Let $B_{k,i}^I(t)$ and $B_{k,i}^{II}(t)$ be the remaining service

requirements at time t of the i -th arriving class- k user from time 0 onward in systems I and II, respectively. Users which have left by time t will simply be considered to have a zero remaining service requirement. To avoid excessive notation, we will view the users which are present at time 0 as having arrived at time 0.

By assumption, for an identical user population, the service rate of each user in system II is at most equal to that in system I. In addition, the service rate of each user in system I is decreasing in the number of users of the various classes. This suggests that system I should provide a lower bound for system II, as is confirmed by the next lemma.

Lemma 4.2: Assume that $M = \infty$. If $(N_1^I(0), \dots, N_m^I(0)) = (N_1^{II}(0), \dots, N_m^{II}(0))$ and $B_{k,i}^I(0) = B_{k,i}^{II}(0)$ for all $i = 1, \dots, N_k^I(0)$, $k = 1, \dots, m$, then $(N_1^I(t), \dots, N_m^I(t)) \leq_{\text{st}} (N_1^{II}(t), \dots, N_m^{II}(t))$ for all $t \geq 0$.

Proof

The proof uses stochastic coupling arguments. Specifically, we assume that the same users arrive to both systems at the same time epochs, with the same transmission rates and the same service requirements. Denote by $A_k(t)$ the number of class- k users arriving up to time t in both systems (including the users which are present at time 0). We will show that $B_{k,i}^I(t) \leq B_{k,i}^{II}(t)$ for all users $i = 1, \dots, A_k(t)$, $k \leq m$, and in particular $N_k^I(t) \leq N_k^{II}(t)$ for all t . The proof proceeds by forward induction on the time parameter t . Let $t_0 = 0, t_1, t_2, \dots$ be the event times, i.e., the time epochs at which users arrive or depart from either system. By assumption, the statement is true for all $t \leq t_0$. Now suppose that the statement is true for all $t \leq t_n$. Thus, $B_{k,i}^I(t_n) \leq B_{k,i}^{II}(t_n)$ for all users $i = 1, \dots, A_k(t_n)$, $k \leq m$, and in particular $N_k^I(t_n) \leq N_k^{II}(t_n)$. We will prove that the statement is then also true for all $t \in (t_n, t_{n+1}]$. Note that it suffices to show that the service rate of each user in system II is at most equal to that in system I. Using the dominance and monotonicity properties of the service rates, it thus suffices to show that $N_k^I(t_n^+) \leq N_k^{II}(t_n^+)$ for all $k = 1, \dots, m$. In order to do so, we distinguish between three different cases, depending on the type of event that occurs at time t_n .

(i) Arrival of a class- l customer. Using the induction hypothesis, $N_k^I(t_n^+) = N_k^I(t_n^-) + I_{\{k=l\}} \leq N_k^{II}(t_n^-) + I_{\{k=l\}} = N_k^{II}(t_n^+)$. Also, by construction $B_{l,A_l}^I(t_n^+) = B_{l,A_l}^{II}(t_n^+)$.

(ii) Service completion of a class- l user in system I. Using the induction hypothesis, $N_k^I(t_n^+) = N_k^I(t_n^-) - I_{\{k=l\}} \leq N_k^{II}(t_n^-) - I_{\{k=l\}} = N_k^{II}(t_n^+)$.

(iii) Service completion of a class- l user in system II. Using the induction hypothesis, we conclude that the class- l user which completes service in system II must already have done so earlier in system I, i.e., $N_l^I(t_n^-) \leq N_l^{II}(t_n^-)$. Thus $N_k^I(t_n^+) = N_k^I(t_n^-) \leq N_k^{II}(t_n^-) - I_{\{k=l\}} = N_k^{II}(t_n^+)$.

Removing the conditioning completes the proof. \square

Applying Lemma 4.2, we obtain the following two corollaries.

Corollary 4.1: Under equal initial conditions, $(N_1(t), \dots, N_m(t)) \leq_{\text{st}} (\tilde{N}_1^{(m)}(t), \dots, \tilde{N}_m^{(m)}(t))$ for all $t \geq 0$.

Corollary 4.2: Under equal initial conditions, $(\tilde{N}_1^{(m)}(t), \dots, \tilde{N}_m^{(m)}(t)) \leq_{\text{st}} (\hat{N}_1^{(m)}(t), \dots, \hat{N}_m^{(m)}(t))$ for all $t \geq 0$.

V. STABILITY PROPERTIES

We now use the notion of the restricted system and the stochastic majorization properties derived in the previous section to establish necessary and sufficient conditions for stability of the various user classes. We assume that $M = \infty$ because otherwise stability is obviously not an issue. Let $(\tilde{N}_1^{(m)}, \dots, \tilde{N}_m^{(m)})$ be a random vector with as distribution the joint stationary distribution of $(\tilde{N}_1^{(m)}(t), \dots, \tilde{N}_m^{(m)}(t))$, assuming the m -restricted integrated system is stable. Denote $\pi^{(m)}(n_1, \dots, n_m) := \mathbb{P}\{(\tilde{N}_1^{(m)}, \dots, \tilde{N}_m^{(m)}) = (n_1, \dots, n_m)\}$, with the convention that $\pi^{(m)}(n_1, \dots, n_m) = 0$ for all (n_1, \dots, n_m) when the m -restricted integrated system is unstable. Define

$$G_k^\alpha := \sum_{n_1=0}^{\infty} \dots \sum_{n_{k-1}=0}^{\infty} \pi^{(k-1)}(n_1, \dots, n_{k-1}) G_k^\alpha(n_1, \dots, n_{k-1}),$$

with the convention that $G_1^\alpha = y_1$. The quantity G_k^α may be interpreted as the long-term average service rate for class k in the k -restricted segregated system when unstable. Also, define $K^* := \max\{k : \rho_l < G_l^\alpha \text{ for all } l = 1, \dots, k\}$. We will show that strategy S^α achieves stability for user classes $1, \dots, K^*$, and does not achieve stability for user classes $K^* + 1, \dots, K$, assuming that in fact $\rho_{K^*+1} > G_{K^*+1}^\alpha$. The result may be heuristically explained as follows. Suppose that some class $k \leq K^*$ were unstable. Let k^* be such class with the lowest index. According to Corollaries 4.1 and 4.2, class k^* must then be unstable in the k^* -restricted segregated system as well. In that system however, the long-term average service rate for class k^* when unstable will be equal to $G_{k^*}^\alpha > \rho_{k^*}$, which is not possible.

Conversely, suppose that some class $k > K^*$ were stable. First observe that classes $1, \dots, k-1$ would then all have to be stable as well, because otherwise class k would be starved due to the fact that $\alpha_1 y_1 > \alpha_2 y_2 > \dots > \alpha_K y_K$. In particular, class $K^* + 1$ would have to be stable. It is intuitively plausible, and can be rigorously shown, that the long-term average service rate for class $K^* + 1$ cannot be larger than the maximum possible long-term average service rate for class $K^* + 1$ in the $(K^* + 1)$ -restricted segregated system, which is equal to $G_{K^*+1}^\alpha < \rho_{K^*+1}$, precluding stability of class $K^* + 1$.

Proposition 5.1: Strategy S^α achieves stability for classes $1, \dots, K^*$.

Proposition 5.2: Strategy S^α does not achieve stability for classes $K^* + 1, \dots, K$.

The detailed proofs of the above two propositions may be found in [6].

Recall that Proposition 3.1 showed that strategy S^* achieves stability for $\rho < G^*$. The above two propositions suggest that strategy S^α may in general not achieve stability for $\rho < G^*$. To show that strategy S^α is not guaranteed to achieve stability whenever possible, it is useful to consider a two-class scenario where Y_1 and Y_2 are identically distributed as Y_0 with $\mathbb{P}\{\alpha_2 y < Y_0 < \alpha_1 y\}$ for some fixed y , so that $\mathbb{P}\{\alpha_1 Y_1 < \alpha_2 Y_2\} = 0$. Recall that the weights are set according to $w_{k,i} = \alpha_k / C_{k,i}$, so the above situation could correspond to two different scenarios: (i) both classes have the same time-average transmission rates $C_{k,i}$, but class 1 is assigned a larger weight and thus effectively receives priority over class 2; (ii) both classes are assigned the same weights $w_{k,i}$, but class 1 effectively receives priority over class 2 due to a larger time-average transmission rate (or a combination of these two scenarios). In either case, service of class 1 takes precedence over that of class 2, and we have $G_1^\alpha(n_1, n_2) = G(n_1)$ and $G_2^\alpha(n_1, n_2) = I_{\{n_1=0\}}G(n_2)$, with $G(n) := \mathbb{E}\{\max\{Y_{01}, \dots, Y_{0n}\}\}$. Thus, there are scheduling gains within both user classes, but not between classes.

Using Proposition 3.1, we deduce that class 1 is stable under strategy S^α as long as $\rho_1 < G^*$, in which case the probability that there are no class-1 users in the system is

$$\pi_0 = \left[\sum_{n=0}^{\infty} \frac{\rho^n}{\phi(n)} \right]^{-1},$$

with $\phi(n) = \prod_{i=1}^n G(i)$. Class 2 is stable if in addition $\rho_2 < \pi_0 G^*$. Now observe that $G(i) < G^*$ for all $i = 1, 2, \dots$ implies that $\pi_0 < 1 - \rho_1 / G^* - \gamma$ for some $\gamma > 0$. Hence, class 2 is stable under strategy S^α only if $\rho < (1 - \gamma)G^*$, which is a strictly stronger condition than for strategy S^* .

VI. NUMERICAL EXPERIMENTS

We now present some numerical experiments that we performed to illustrate the results. We consider a system where users initiate file transfer requests as a Poisson process. At most $M = 20$ users are admitted into the system simultaneously. Users which generate download requests when there are already M transfers in progress are blocked and lost. The system operates in a time-slotted fashion, with a slot duration of 1.67 ms (600 slots per second) as in the CDMA 1xEV-DO system. Throughout, we assume that the users have independent Rayleigh fading channels.

We consider three different scenarios for the distribution of the mean SNR: (I) identical to 0 dB for all users; (II) a bi-modal distribution, either -2.0 dB or 4.0 dB with equal probability; (III) a linearized version of the distribution plotted in Figure 1 taken from [4].

The above assumptions determine how the instantaneous SNR values of the various users behave over time. It remains to specify how the instantaneous transmission rate of a user varies with the instantaneous SNR value. We distinguish between three different scenarios: (A) the instantaneous rate is linear in the instantaneous SNR (on an absolute scale): $R = C_1 \times SNR$, with $C_1 = 400$ Kbs; (B) the instantaneous rate is logarithmic in the instantaneous SNR (on an absolute




Fig. 1. CDF of SNR distribution from [4]

SNR \geq	Rate (Kbs)
- 12.5	38.4
- 9.5	76.8
- 8.5	102.6
- 6.5	153.6
- 5.7	204.8
- 4.0	307.2
- 1.0	614.4
1.3	921.6
3.0	1228.8
7.2	1843.2
9.5	2457.6

TABLE I
RATE (KBS) AS FUNCTION OF SNR IN 1xEV-DO SYSTEM

scale): $R = C_2 \times \log(1 + SNR)$, with $C_2 = 800$ Kbs; (C) the instantaneous rate is determined from the instantaneous SNR value (in dB) according to Table I as is used in the CDMA 1xEV-DO system [4].

In the first set of experiments, we examined the strategy S^* which assigns a weight $w_i = 1/C_i$ to a user i with a time-average transmission rate C_i . We determined the mean number of users, the mean response times, the blocking probabilities, and the mean throughput for varying arrival rates, comparing the analytical formulas given in Proposition 3.1 with simulation results. The simulations were run for 100,000,000 time slots, or equivalently, about 167,000 seconds of real time. Throughout, the (mean) file size is assumed to be 60 Kbytes (480 Kbits). We focus the discussion on results that we obtained for deterministic file sizes. Observe that Proposition 3.1 indicates that the above-mentioned performance metrics should be mostly insensitive to the file size distribution in case of a symmetric rate distribution. However, we will also present some results for exponentially distributed file sizes, which suggest that the performance metrics continue to be fairly insensitive even when the rate variations fail to be entirely symmetric.

We considered a total of nine cases obtained via pairwise

combination of the above scenarios for the mean SNR distribution and the rate variations. Note that the relative rate fluctuations are only statistically identical in case all users have identical mean SNR values or the instantaneous rate is linear in the instantaneous SNR value, i.e., in cases IA, IB, IC, IIA, and IIIA. In the remaining four cases, the relative rate fluctuations differ across users, and the notion of a gain factor $G(n)$ as defined in Section II is not strictly valid. In order to evaluate the analytical formulas, we used an approximate gain factor, which was computed as the gain factor that would have been obtained in a scenario with identical mean SNR values for all users, where the mean was calculated as the average SNR across users (on a logarithmic scale) The latter approximation is expected to be somewhat conservative, since the actual user populations will tend to be biased to low-SNR users, for whom the relative gain factor tends to be larger due to the concavity and the truncation of the transmission rate which limit the potential relative gain for the high-SNR users.

Figure 2 depicts the mean transfer delay as a function of the file arrival rate for cases IA-C, and indicates that the analytical formulas provide a highly accurate estimate in cases where the rate fluctuations are statistically identical: the analysis and simulation curves are nearly indistinguishable. This confirms that the separation of time scales underlying the analytical formulas is a reasonable assumption.

Figures 3-5 display the mean total number of users and the mean transfer delay for class-1 and class-2 users for cases IIA-C. In Case IIA, the rate fluctuations are still statistically identical, and Proposition 3.1 implies that the mean delays of class-1 and class-2 users should be inversely proportional to the time-average transmission rates, i.e., have a ratio of $10^{-0.2}$ to $10^{0.4}$, which is roughly 1 to 4, as is confirmed by Figures 4 and 5. Remarkably enough, Figure 3 shows that the analytical formulas continue to yield a fairly accurate prediction for the mean total number of transfers in progress in cases IIB and IIC, despite the fact that the rate fluctuations vary across users. However, Figures 4 and 5 reflect that the accuracy of the formulas for the mean per-class transfer delays is rather poor in these cases. The formulas consistently underestimate the delay for the high-SNR users and overestimate the delay for the low-SNR users. This is attributed to the fact that the relative gain from the rate variations is smaller for the high-SNR users due to the concavity and the truncation of the transmission rate as mentioned above.

Figure 6 plots the mean transfer delay as a function of the file arrival rate for cases IIIA-C, and indicates that the analytical formulas remain surprisingly accurate for a continuous distribution of the mean SNR as well. We repeated the latter experiment for exponentially distributed file sizes. The results as graphed in Figure 7 show that the mean transfer delay is fairly insensitive to the file size distribution, even when the symmetry conditions of Proposition 3.1 are not strictly satisfied.

In the second set of experiments, we evaluated similar performance metrics for varying weight factors used in the allocation of time slots. We considered a total of six cases obtained via pairwise combination of channel scenarios as




Fig. 2. Mean transfer delay as function of file arrival rate for Cases IA-C




Fig. 3. Mean total number of active users as function of file arrival rate for Cases IIA-C

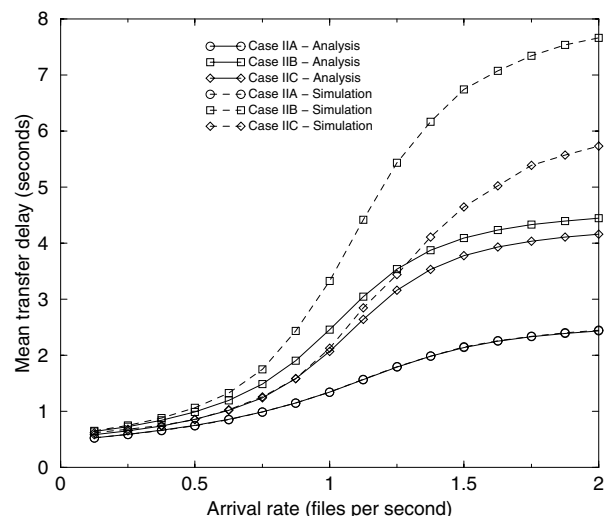


Fig. 4. Mean transfer delay for class-1 users as function of file arrival rate for Cases IIA-C




Fig. 5. Mean transfer delay for class-2 users as function of file arrival rate for cases IIA-C

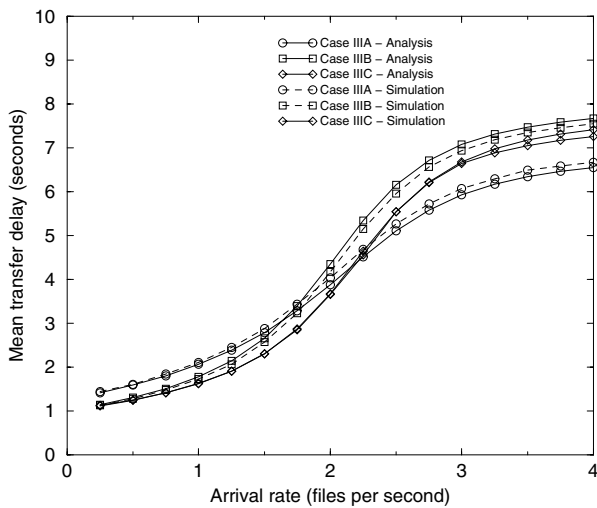


Fig. 6. Mean transfer delay as function of file arrival rate for Cases IIIA-C with deterministic file size




Fig. 7. Mean transfer delay as function of file arrival rate for Cases IIIA-C with exponentially distributed file sizes

before. In order to investigate the impact of the weight factors, we focused on a system with two user classes. The mean SNR values and the weights of the users are class-dependent. Within classes, the users have identical mean SNR values and equal weights. Throughout, the file size is deterministic and assumed to be 60 Kbytes (480 Kbits) as before.

Figures 8 and 9 depict the mean system throughput and the mean number of users as a function of $\log_2(w_2/w_1)$ for cases IA-C. Since all users have identical time-average transmission rates, Propositions 3.1 and 3.2 imply that strategy S^* which sets the weights equal for both classes is optimal from a stability perspective. Observe from the two figures that the minimum aggregate occupancy and maximum total throughput performance is achieved for equal weights as well. This demonstrates that the maximal stability guarantee of strategy S^* translates into superior performance in terms of overall throughput in systems with admission control where strictly speaking stability is not an issue. Further observe that Figures 8 and 9 indicate that differentiation between user classes can only be accomplished at the expense of the overall throughput performance.

Figures 10 and 11 plot similar performance characteristics for cases IIA-C. The time-average transmission rates of class-1 (4.0 dB) and class-2 (-2.0 dB) users in these three cases may be computed to be 1005 and 252, 840 and 344, and 915 and 348 Kbs, respectively. The strategy S^* in these three cases thus corresponds to a ratio between the weights of approximately 3.99, 2.44 and 2.63, respectively. As before, the two figures demonstrate that the minimum aggregate occupancy and the optimum total throughput performance is obtained for weight settings in that range. In addition, the prioritization of the high-SNR users incurs a relatively modest penalty in terms of the overall system performance, but preferential treatment for the low-SNR users can have major repercussions. Further observe that the strategy which sets the weights equal for both classes, and thus maximizes the throughput in a static scenario, fails to do so in a dynamic setting, although not by a wide margin. The intuitive explanation is that considering the absolute rates does not extract the maximum gains from the relative rate variations.

ACKNOWLEDGMENT

The author is grateful to Phil Whiting for helpful discussions.

REFERENCES

- [1] Agrawal, R., Bedekar, A., La, R.J., Subramanian, V. (2001). Class and channel condition based weighted proportional fair scheduler. In: *Teltraffice Engineering in the Internet Era, Proc. ITC-17*, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 553-565.
- [2] Andrews, D.M. (2002). Instability of the Proportional Fair scheduling algorithm for HDR. Technical memorandum, submitted for publication.
- [3] Andrews, D.M., Kumaran, K., Ramanan, K., Stolyar, A.L., Vijayakumar, R., Whiting, P.A. (2000). Scheduling in a queueing system with asynchronously varying service rates. Technical memorandum, submitted for publication.
- [4] Bender, P., Black, P., Grob, M., Padovani, R., Sindhushayana, N., Viterbi, A. (2000). CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users. *IEEE Commun. Mag.* **38** (7), 70-77.




Fig. 8. Mean system throughput as function of $\log_2(w_1/w_2)$ for Cases IA-C




Fig. 11. Mean total number of users (solid line), class-1 users (dashed line), and class-2 users (dotted line) as function of $\log_2(w_1/w_2)$ for cases IIA-C




Fig. 9. Mean total number of users (solid line), class-1 users (dashed line), and class-2 users (dotted line) as function of $\log_2(w_1/w_2)$ for cases IA-C




Fig. 10. Mean system throughput as function of $\log_2(w_1/w_2)$ for Cases IIA-C

- [5] Bonald, T., Proutière, A., Régnié, G., Roberts, J.W. (2001). Insensitivity results in statistical bandwidth sharing. In: *Teletraffic Engineering in the Internet Era, Proc. ITC-17*, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 125–136.
- [6] Borst, S.C. (2002). User-level performance of channel-aware scheduling algorithms in wireless data networks. CWI Report PNA-R0223.
- [7] Borst, S.C., Kumaran, K., Ramanan, K., Whiting, P.A. (2002). Queuing models for user-level performance of Proportional Fair scheduling. Technical memorandum, Bell Laboratories, Lucent Technologies.
- [8] Borst, S.C., Whiting, P.A. (2001). Dynamic rate control algorithms for HDR throughput optimization. In: *Proc. Infocom 2001*, 976–985.
- [9] Cohen, J.W. (1979). The multiple phase service network with generalized processor sharing. *Acta Informatica* **12**, 245–284.
- [10] Fayolle, G., Mitrani, I., Iasnogorodski, R. (1980). Sharing a processor among many job classes. *J. ACM* **27**, 519–532.
- [11] Jalali, A., Padovani, R., Pankaj, R. (2000). Data throughput of CDMA-HDR a high efficiency-high data rate personal communication wireless system. In: *Proc. 50th IEEE Veh. Techn. Conf.*, 1854–1858.
- [12] Kelly, F.P. (1979). *Reversibility and Stochastic Networks*. Wiley, Chichester.
- [13] Kushner, H.J., Whiting, P.A. (2002). Convergence of Proportional Fair sharing algorithms under general conditions. Technical memorandum, submitted for publication.
- [14] Liu, X., Chong, E.K.P., Shroff, N.B. (2001). A framework for opportunistic scheduling. Technical report, submitted for publication.
- [15] Neely, M.J., Modiano, E., Rohrs, C.E. (2002). Power and server allocation in a multi-beam satellite with time varying channels. In: *Proc. Infocom 2002*, 1451–1460.
- [16] Rege, K.M., Sengupta, B. (1996). Queue-length distribution for the discriminatory processor-sharing queue. *Oper. Res.* **44**, 653–657.
- [17] Shakkottai, S., Stolyar, A.L. (2001). Scheduling algorithms for a mixture of real-time and non-real time data in HDR. In: *Teletraffic Engineering in the Internet Era, Proc. ITC-17*, Salvador da Bahia, eds. J.M. de Souza, N.L.S. da Fonseca, E.A. de Souza e Silva (North-Holland, Amsterdam), 793–804.
- [18] Stolyar, A.L. (2002). Multi-user throughput dynamics under Proportional Fair and other gradient-based scheduling algorithms. Technical memorandum, submitted for publication.
- [19] Tassiulas, L., Ephremides, A. (1993). Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Trans. Inf. Theory* **30**, 466–478.
- [20] Tse, D. (2002). Multi-user diversity and Proportional Fair scheduling. In preparation.