

Impact of size-based scheduling on flow level performance in wireless downlink data channels

Samuli Aalto and Pasi Lassila

TKK Helsinki University of Technology, Networking Laboratory, P.O. Box 3000,
FIN-02015 TKK, Finland
{samuli.aalto,pasi.lassila}@tkk.fi

Abstract. We analyze the impact of size-based scheduling on the flow level performance of elastic traffic in wireless downlink data channels. The impact is assessed by comparing the flow level delay of the simple RR scheduler to two optimized non-anticipating schedulers (FB and FB^o) and SRPT. The optimized distance-aware scheduler FB^o is derived by applying the Gittins index approach. Our results show that for Pareto-type file size distributions, the size-based information is more important than the location information. Additionally, FB not only decreases the overall mean delay, but it can also decrease considerably the mean delay of all users independently of their location.

Keywords: Scheduling, cellular system, HSPDA, HDR, mean delay, elastic traffic, Pareto distribution, Gittins index.

1 Introduction

We consider the optimal scheduling problem for downlink data traffic in a single cell of a cellular system. We assume that the traffic consists of elastic flows, such as file transfers using TCP, with each flow characterized by its size, *i.e.*, the total amount of bits to be transferred. An important performance measure for such elastic flows is the total time needed for transfer of all the bits, which we call flow level delay, or just briefly, *delay*.

Standardized systems like HSPDA and HDR minimize the intra-cell interference by having a time-slotted system where the base station transmits at full power to only one terminal in each slot. Typically the time slot is very short (milliseconds) compared to the flow level delay (ranging from seconds to minutes). The “air time” is shared in a fair way among the receiving terminals by the simple Round Robin (RR) scheduling discipline. For short time slots this is well approximated by the Processor Sharing (PS) discipline.

The cell capacity may be increased by applying *channel-aware* scheduling disciplines that utilize the fast fading effect. An example is the Proportionally Fair (PF) discipline implemented in HSPDA/HDR systems, where the time slot is scheduled to the terminal with the highest momentary receiving rate proportionally to its average receiving rate. This will lead to the PS discipline in the

limit as the time slot shrinks down to zero, however, with a higher cell capacity compared to the limit of the RR discipline [8, 2, 4]. The improvement in the cell capacity has been studied in [2] with the number of flows fixed and in [3] assuming a randomly varying number of flows. The latter concludes that the improvement in a dynamic setting, however, is not that large.

A fundamental result related to single server queueing systems says that the number of jobs is minimized pathwise by applying the SRPT (Shortest Remaining Processing Time) scheduling discipline [13]. This implies, by Little's result, that SRPT minimizes the mean delay as well. The benefit is achieved by utilizing the *size-dependent information*. However, applying SRPT as such in the downlink data traffic problem is intractable due to unpredictable factors such as fading effects. Hu *et al.* [9] developed heuristic algorithms that combine channel-aware and size-dependent scheduling.

In this paper we focus on non-anticipating disciplines for which the remaining service times are *not* known. If the service time distribution belongs to the class DHR (Decreasing Hazard Rate), the FB (Foreground Background) discipline is optimal among the non-anticipating disciplines in an M/G/1 queue [14, 11, 15]. FB is a size-based discipline giving full priority to the job with least amount of attained service, see [10]. If there are multiple jobs with the same least amount of attained service, then the service is shared evenly between these jobs.

What makes the difference here in the downlink data traffic problem, as compared to an ordinary M/G/1 queue, is the *location information*. Due to the slow fading effect, terminals far away from the base station have a lower receiving rate than the near-by terminals. Thus, an optimal scheduler that utilizes this information achieves even better delay performance than FB. As the main theoretical contribution of this paper, we determine the optimal non-anticipating and distance-aware discipline for DHR flow sizes by applying the so called Gittins index approach. In the case of Pareto distributions, which have been used to model flow sizes in the Internet [5, 6], the optimal discipline proves to be a simple modification of FB. We demonstrate that the improvement of the optimal policy is, however, marginal when compared to the ordinary FB discipline for Pareto flow size distributions. This implies that utilizing the size-dependent information is more important than utilizing the location information when scheduling the time slots for downlink data traffic. On the other hand, we also demonstrate, that there is a clear improvement in the delay performance when the simple RR discipline is replaced by FB.

Bonald and Proutière [3] state that FB would exacerbate the discrimination against far terminals. This is indeed the case: the reduction in the mean delay is greater for the near terminals. Our observation, however, is that there is *some reduction* in the mean delay *for all terminals* even for the farthest ones. The losers are the huge flows independent of their location.

The rest of the paper is organized as follows. The model is explained in Section 2. In Section 3 we give the mean delay formulas for the reference disciplines, while in Section 4 the optimal non-anticipating and distance-aware discipline is

determined. Different disciplines are compared in Section 5 based on simulations and numerical evaluations. Section 6 concludes the paper.

2 Model

Consider downlink data traffic in a single cell of a cellular system. The traffic consists of elastic flows. We assume that the flows arrive at the base station according to a Poisson process with rate λ . The flow sizes X are assumed to be independent and identically distributed (positive real-valued) random variables with the cumulative distribution function denoted by $F_X(x) = P\{X \leq x\} = \int_0^x f_X(y) dy$, where $f_X(x)$ refers to the corresponding density function. The hazard rate $h_X(x)$ is defined by $h_X(x) = f_X(x)/\bar{F}_X(x)$, where $\bar{F}_X(x) = 1 - F_X(x)$. We assume that the flow size distribution belongs to the class DHR with a differentiable density function $f_X(x)$. Thus, $h'_X(x) \leq 0$ for all $x \geq 0$.

Each flow is associated with a receiving terminal. We assume that the terminals are independently and uniformly distributed in the cell area, which is a circular disk with radius r_1 . Thus, the distance R from the base station to the receiving terminal for a flow has the following distribution:

$$P\{R \leq r\} = \frac{1}{\pi r_1^2} \int_0^r 2\pi s ds = \left(\frac{r}{r_1}\right)^2, \quad r \leq r_1.$$

Furthermore, we assume that the transmission rate (bits per time unit) from the base station to the receiving terminal depends on the distance r between them as follows:

$$c(r) = \begin{cases} c_0, & r \leq r_0, \\ c_0 \left(\frac{r_0}{r}\right)^\alpha, & r > r_0. \end{cases} \quad (1)$$

Thus, we only consider the slow fading effect. Parameter α , called the attenuation factor, typically takes values in the range from 2 to 4. The fast fading effect is assumed to average out in the timescale of flow level delay. Note that we consider an ideal case where the set of achievable rates is continuous omitting possible coding constraints. No intercell interference is included in this single cell model.

Given the location of the receiving terminal, the service time for a flow is just the flow size divided by the constant rate. Let $F_r(t)$ denote the cumulative distribution function of the service time $S_r = X/c(r)$ for a flow with the terminal located at distance r from the base station. Now $E[S_r] = E[X]/c(r)$ and

$$F_r(t) = P\{S_r \leq t\} = P\{X \leq c(r)t\} = F_X(c(r)t).$$

The corresponding density function is clearly $f_r(t) = c(r)f_X(c(r)t)$ and the corresponding hazard rate $h_r(t) = c(r)h_X(c(r)t)$. Since the flow sizes X are DHR, also these conditional service times S_r belong to the class DHR.

On the other hand, if the location of the receiving terminal is *not* known, then the service time for a flow is the ratio between two independent random variables, $S = X/c(R)$, with mean

$$E[S] = E[X] \int_0^{r_1} \frac{1}{c(r)} P\{R \in dr\} = \frac{E[X]}{c_0} \left(\frac{\alpha}{\alpha+2} \left(\frac{r_0}{r_1}\right)^2 + \frac{2}{\alpha+2} \left(\frac{r_1}{r_0}\right)^\alpha \right).$$

The cumulative distribution function of the service time S is denoted by $F(t) = P\{S \leq t\} = \int_0^t f(s) ds$ with $f(t)$ referring to the corresponding density function. It is easy to see that

$$F(t) = F_X\left(c_0 t \left(\frac{r_0}{r_1}\right)^\alpha\right) + \left(\frac{r_0}{r_1}\right)^2 \int_{c_0 t (r_0/r_1)^\alpha}^{c_0 t} f_X(x) \left(\frac{c_0 t}{x}\right)^{\frac{2}{\alpha}} dx,$$

from which the density can be derived with the following result:

$$f(t) = f_X(c_0 t) c_0 \left(\frac{r_0}{r_1}\right)^2 + \frac{2}{\alpha t} \left(\frac{r_0}{r_1}\right)^2 \int_{c_0 t (r_0/r_1)^\alpha}^{c_0 t} f_X(x) \left(\frac{c_0 t}{x}\right)^{\frac{2}{\alpha}} dx.$$

The hazard rate of S is denoted by $h(t)$.

Proposition 1. *If the flow size distribution $F_X(x)$ belongs to class DHR, so does the service time distribution $F(t)$.*

Proof. Since $F(t) = \int_0^{r_1} F_r(t) P\{R \in dr\}$, this is a special case of the result given in [1, Theorem 3.4]. \square

Throughout the paper we assume that the time slot used for scheduling is negligible compared to the flow sizes. So we have an M/G/1 queue with arrival rate λ , service times S , and load $\rho = \lambda E[S]$. For stability, we assume that $\rho < 1$.

3 Reference schedulers

In this section we give the mean delay formulas for the disciplines PS, FB, SRPT, and GR. The PS discipline (representing the limiting case of RR) is our main reference scheduler. FB is the optimal non-anticipating discipline that does not utilize location information. SRPT is the optimal (hypothetical) scheduler which gives the lower bound for the delay performance. The last one, GR, refers to the greedy distance-aware discipline that gives the full priority to the receiving terminal with the shortest distance from the base station.

PS The conditional mean delay of a flow with service time t is for the PS discipline as follows [10]:

$$E[T^{\text{PS}}(t)] = \frac{t}{1 - \rho}. \quad (2)$$

Thus, the mean delay of a flow is

$$E[T^{\text{PS}}] = \int_0^\infty E[T^{\text{PS}}(t)] f(t) dt = \frac{E[S]}{1 - \rho}. \quad (3)$$

In addition, the mean delay of a flow with the receiving terminal located at distance r from the base station becomes

$$E[T_r^{\text{PS}}] = \int_0^\infty E[T^{\text{PS}}(t)] f_r(t) dt = \frac{E[S_r]}{1 - \rho}. \quad (4)$$

FB The conditional mean delay for FB reads as follows [10]:

$$E[T^{\text{FB}}(t)] = \frac{\lambda E[(S \wedge t)^2]}{2(1 - \rho_t)^2} + \frac{t}{1 - \rho_t}, \quad (5)$$

where $S \wedge t = \min\{S, t\}$ and ρ_t refers to the truncated load, $\rho_t = \lambda E[(S \wedge t)]$. The mean delays $E[T^{\text{FB}}]$ and $E[T_r^{\text{FB}}]$ are calculated from this conditional mean delay $E[T^{\text{FB}}(t)]$ similarly as for the PS discipline, see (3) and (4).

SRPT The conditional mean delay formula for SRPT originates from [12]:

$$E[T^{\text{SRPT}}(t)] = \frac{\lambda E[(S \wedge t)^2]}{2(1 - \rho(t))^2} + \int_0^t \frac{1}{1 - \rho(s)} ds. \quad (6)$$

Here $\rho(t)$ refers to $\rho(t) = \lambda \int_0^t sf(s) ds$. Again, the mean delays $E[T^{\text{SRPT}}]$ and $E[T_r^{\text{SRPT}}]$ are calculated from this conditional mean delay $E[T^{\text{SRPT}}(t)]$ similarly as for the PS discipline, see (3) and (4).

GR Discipline GR results in a pre-emptive priority M/G/1 queue with a continuum of priority classes. By applying the well known results for priority queues [10], we conclude that the mean delay for a flow with the receiving terminal located at distance r from the base station is as follows:

$$E[T_r^{\text{GR}}] = \frac{E[S_r^*]}{(1 - \sigma_r)^2} + \frac{E[S_r]}{1 - \sigma_r}. \quad (7)$$

Here, σ_r is the expected load up to distance r defined by

$$\begin{aligned} \sigma_r &= \lambda \int_0^r E[S_z] P\{R \in dz\} = \lambda E[X] \int_0^r \frac{1}{c(z)} P\{R \in dz\} \\ &= \begin{cases} \lambda \frac{E[X]}{c_0} \left(\frac{r}{r_1}\right)^2, & r \leq r_0, \\ \lambda \frac{E[X]}{c_0} \left(\frac{\alpha}{\alpha+2} \left(\frac{r_0}{r_1}\right)^2 + \frac{2}{\alpha+2} \left(\frac{r}{r_0}\right)^\alpha \left(\frac{r}{r_1}\right)^2\right), & r > r_0, \end{cases} \end{aligned}$$

and S_r^* refers to the so called remaining service time for terminals up to distance r with mean

$$\begin{aligned} E[S_r^*] &= \frac{\lambda}{2} \int_0^r E[S_z^2] P\{R \in dz\} = \frac{\lambda}{2} E[X^2] \int_0^r \frac{1}{c^2(z)} P\{R \in dz\} \\ &= \begin{cases} \frac{\lambda}{2} \frac{E[X^2]}{c_0^2} \left(\frac{r}{r_1}\right)^2, & r \leq r_0, \\ \frac{\lambda}{2} \frac{E[X^2]}{c_0^2} \left(\frac{\alpha}{\alpha+1} \left(\frac{r_0}{r_1}\right)^2 + \frac{1}{\alpha+1} \left(\frac{r}{r_0}\right)^{2\alpha} \left(\frac{r}{r_1}\right)^2\right), & r > r_0. \end{cases} \end{aligned}$$

The mean delay is given by

$$E[T^{\text{GR}}] = \int_0^{r_1} E[T^{\text{GR}}(r)] P\{R \in dr\}. \quad (8)$$

4 Optimal non-anticipating distance-aware scheduler

In this section we determine the optimal non-anticipating and distance-aware discipline for DHR flow sizes. Let Π denote the family of non-anticipating scheduling disciplines that do not utilize location information. In addition, let Π° denote the whole family of non-anticipating scheduling disciplines including also those that utilize the location information. Thus, our purpose is to find the optimal discipline in Π° .

4.1 Gittins index

In this subsection we assume that the scheduler is not utilizing location information. We recall some results related to the so called Gittins index, see [7, 15], and apply them to DHR service times.

For any $a, \Delta \geq 0$, let

$$J(a, \Delta) = \frac{\int_0^\Delta f(a+t) dt}{\int_0^\Delta \bar{F}(a+t) dt}. \quad (9)$$

Note that $J(a, 0) = h(a)$ for any a . Function $J(a, \Delta)$ is clearly continuous with respect to both arguments. In addition, the one-sided partial derivatives with respect to Δ are defined for any pair (a, Δ) ,

$$\frac{\partial}{\partial \Delta} J(a, \Delta) = \frac{f(a+\Delta) \int_0^\Delta \bar{F}(a+t) dt - \bar{F}(a+\Delta) \int_0^\Delta f(a+t) dt}{(\int_0^\Delta \bar{F}(a+t) dt)^2} \quad (10)$$

For a flow with attained service a , the *Gittins index* is defined as follows [7, 15]:

$$G(a) = \sup_{\Delta \geq 0} J(a, \Delta), \quad (11)$$

Consider now a non-anticipating discipline π^* which always gives service to the job with the highest Gittins index. We call this discipline the *Gittins discipline*. It is known that the Gittins discipline is optimal with respect to the mean delay for an M/G/1 queue, see [7, Theorem 3.28], [15, Theorem 4.7].

Theorem 1. $E[T^{\pi^*}] \leq E[T^\pi]$ for any $\pi \in \Pi$.

This can be used to prove the optimality of the FB discipline for DHR flow sizes. Recall from Proposition 1 that in this case also the service times belong to the class DHR so that the hazard rate $h(t)$ is a decreasing function.

Proposition 2. *Function $J(a, \Delta)$ is decreasing with respect to Δ for any a .*

Proof. Let $a, \Delta \geq 0$. Since $h(t)$ is decreasing, we have $h(a+t) \geq h(a+\Delta)$ for all $0 \leq t \leq \Delta$, which is equivalent with

$$\frac{f(a+t)}{f(a+\Delta)} \geq \frac{\bar{F}(a+t)}{\bar{F}(a+\Delta)}. \quad (12)$$

By (10), we have

$$\frac{\partial}{\partial \Delta} J(a, \Delta) \leq 0 \iff \frac{1}{\int_0^\Delta \frac{f(a+t)}{f(a+\Delta)} dt} \leq \frac{1}{\int_0^\Delta \frac{\bar{F}(a+t)}{\bar{F}(a+\Delta)} dt}.$$

The claim follows from this by (12). \square

Proposition 3. $G(a) = h(a)$ for all a .

Proof. Let $a \geq 0$. By Proposition 2, $G(a) = J(a, 0) = h(a)$. \square

Theorem 2. $E[T^{\text{FB}}] \leq E[T^\pi]$ for any $\pi \in \Pi$.

Proof. By Proposition 3 and the fact that $h(a)$ is decreasing, the flow with least amount of service has the highest Gittins index. Thus, in this case the Gittins discipline corresponds to FB. The claim follows now from Theorem 1. \square

4.2 Utilizing location information

Consider now how a non-anticipating discipline can be improved if the scheduler is aware of the distances between the base station and the receiving terminals.

Recall that $h_r(t)$ refers to the hazard rate related to the cumulative distribution function of the service time S_r for a flow with the terminal located at distance r from the base station. As mentioned in Section 2, the service times S_r belong to the class DHR. Thus, the Gittins index for such a flow with attained service a is $G_r(a) = h_r(a) = c(r)h_X(c(r)a)$.

Let R_i and $A_i(t)$, respectively, denote the distance from the base station and the attained service time related to flow i at time t . Furthermore, let $\gamma_i(t)$ denote the proportion of time that is scheduled for flow i at time t . It follows that $A_i(t) = \int_0^t \gamma_i(s) ds$. According to the Gittins rule, the optimal scheduler transmits to terminal j such that

$$h_{R_j}(A_j(t)) = \max_i h_{R_i}(A_i(t)) = \max_i c(R_i) h_X(c(R_i)A_i(t)). \quad (13)$$

If this maximum is not unique, the service capacity shall be shared between the maximizing terminals $j \in \mathcal{J}(t)$ in such a way that

$$\frac{d}{dt} h_{R_j}(A_j(t)) = c(R_j)^2 h'_X(c(R_j)A_j(t)) \gamma_j(t) \quad (14)$$

is the same for all $j \in \mathcal{J}(t)$. The optimal shares $\gamma_j(t)$ can be determined from this condition together with the constraint

$$\sum_{j \in \mathcal{J}(t)} \gamma_j(t) = 1. \quad (15)$$

We denote this optimal non-anticipating distance-aware discipline by FB° .

Theorem 3. $E[T^{\text{FB}^\circ}] \leq E[T^\pi]$ for any $\pi \in \Pi^\circ$.

Proof. The claim follows from Theorem 1 together with the derivation made above. \square

4.3 Pareto flow sizes

Assume now that the flow size distribution is Pareto with shape parameter $\beta > 1$ and scale parameter $b > 0$ such that, for all $x \geq 0$,

$$\bar{F}_X(x) = \left(\frac{1}{1+bx} \right)^\beta, \quad h_X(x) = \frac{\beta b}{1+bx}, \quad h'_X(x) = \frac{-\beta b^2}{(1+bx)^2}.$$

By (13), the optimal distance-aware discipline FB° transmits to terminal j such that

$$\frac{1}{bc(R_j)} + A_j(t) = \min_i \left(\frac{1}{bc(R_i)} + A_i(t) \right). \quad (16)$$

Interestingly, the optimal rule is independent of the shape parameter β . In addition, the inverse of the Gittins index grows linearly with the amount of attained service. Within the constant transmission rate area ($r \leq r_0$), the flow with the least amount of attained service is the preferred one. In particular, if $r_0 = r_1$, then FB° reduces back to FB .

If the maximum is not unique, then

$$\frac{1}{c(R_j)} + b A_j(t) \quad (17)$$

is the same for all maximizing terminals $j \in \mathcal{J}(t)$. By (14), if the maximum is not unique, the service capacity shall be shared between the maximizing terminals $j \in \mathcal{J}(t)$ so that

$$c(R_j)^2 h'_X(c(R_j) A_j(t)) \gamma_j(t) = \frac{-\beta \gamma_j(t)}{\left(\frac{1}{c(R_j)} + b A_j(t) \right)^2} \quad (18)$$

is the same for all $j \in \mathcal{J}(t)$. But now by (17) we conclude that, in fact, the optimal shares $\gamma_j(t)$ are the same for all $j \in \mathcal{J}(t)$. Thus, by (15),

$$\gamma_j(t) = \frac{1}{|\mathcal{J}(t)|}, \quad (19)$$

where $|\mathcal{J}(t)|$ refers to the magnitude of the set $\mathcal{J}(t)$. Thus, for the Pareto distribution, FB° applies PS among the maximizing terminals.

4.4 Exponential flow sizes

Assume now that the flow size distribution is exponential with rate μ so that, for all $x \geq 0$,

$$\bar{F}_X(x) = e^{-\mu x}, \quad h_X(x) = \mu, \quad h'_X(x) = 0.$$

By (13), FB° transmits to terminal j such that

$$c(R_j) = \max_i c(R_i). \quad (20)$$

It is easy to see that, in this case, it does not matter how the service is shared among the flows with the same transmission rate. Thus, for exponential flow sizes, FB° is equivalent with the greedy discipline GR .

5 Numerical results

In the following we give numerical evidence of the amount of performance gains achievable with the various scheduling disciplines. As the baseline scheduler we use the PS discipline (corresponding to the limit of the RR scheduler). The idea is to compare the performance of the other disciplines with respect to PS, *i.e.*, we are only interested in the relative performance of the policies. Thus, in our numerical examples we have scaled the parameters such that job sizes have unit length, $E[X] = 1$, the maximum transmission rate $c_0 = 1$, and the cell radius $r_1 = 1$ (unit circle). We evaluate the relative performance for exponentially distributed and Pareto distributed flow sizes. The parameters affecting the performance are the attenuation factor α and the radius of the constant rate area r_0 , which jointly define the variability of the transmission rates. Additionally, for the Pareto distribution we have the shape parameter β .

In practice, the variability in the rates (or the range of the possible rates) is determined by the capabilities of the technology and the amount of transmission power the base station has available. To have a realistic scenario for the variability in the rates, we use the parameters from Table 1 in [3], where it is given that in current HDR/HSPDA systems for $\alpha = 2$ the ratio of the cell radius to the constant rate area's radius $r_1/r_0 = 7.94$.

5.1 Overall mean delay

We first compare the optimal non-anticipating distance-aware scheduler FB° with the plain FB. We have simulated the system at a fixed load, $\rho = 0.9$, with Pareto distributed flow sizes under both policies using the same stochastic input to minimize the variance. In the simulations we vary the loss exponent $\alpha = 2$ and 4, and the Pareto shape parameter $\beta = 2$ and 3. With exponentially distributed flow sizes, the FB° discipline corresponds to the GR discipline, and the results can be obtained analytically.

The results in Table 1 depict the relative mean delay difference $\Delta = (E[T^{\text{FB}^\circ}] - E[T^{\text{FB}}])/E[T^{\text{FB}}]$, together with the confidence intervals for simulation results. As seen from the results, FB° performs better than FB, especially for exponential flow sizes. However, for Pareto flow sizes the difference is rather small so that one does not benefit much from the location information.

	Exponential		Pareto, $\beta = 3$		Pareto, $\beta = 2$	
	$\alpha = 2$	$\alpha = 4$	$\alpha = 2$	$\alpha = 4$	$\alpha = 2$	$\alpha = 4$
Δ	15.7%	17.3%	$2.0\% \pm 0.1\%$	$2.6\% \pm 0.1\%$	$0.7\% \pm 0.1\%$	$1.0\% \pm 0.1\%$

Table 1. Mean delay comparison between FB° and FB.

Next we study the mean overall delays as a function of the load ρ . Since the delay difference between FB° and FB is so small, results for FB° are not

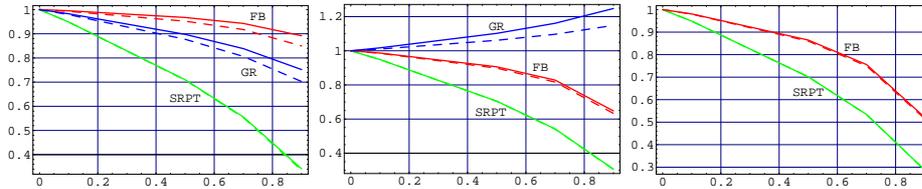


Fig. 1. Mean delay ratio as a function of ρ for exponential distribution (left), Pareto with $\beta = 3$ (middle) and Pareto with $\beta = 2$ (right).

shown in the Pareto cases. The idea is to compare how much better/worse the disciplines GR, FB and SRPT perform than PS by considering the normalized results $E[T^\pi]/E[T^{\text{PS}}]$ for $\pi = \text{GR, FB, SRPT}$.

The results are shown in Figure 1. In the graphs, solid lines correspond to $\alpha = 2$ and the dashed lines to $\alpha = 4$. The scheduling disciplines are indicated next to each curve. The left figure corresponds to exponentially distributed flows sizes. The middle and right figures correspond to the Pareto case with shape parameters $\beta = 3$ and 2 , respectively. Under the exponential distribution (left figure), all disciplines achieve a better performance than PS. In the Pareto cases, for $\beta = 3$ (middle figure), the mean delay for GR is finite, but the delay is greater than under PS for both $\alpha = 2$ and 4 . For $\beta = 2$ (right figure), $E[X^2] \rightarrow \infty$, and correspondingly the mean delay for GR becomes infinite. For FB and SRPT, the delays are always smaller than under PS and the benefit increases with load. Notably, the benefit does not seem to be affected by the value of α .

In conclusion, it is clear that the size-based scheduling mechanisms FB, FB^o and SRPT yield substantial performance gains in terms of the overall mean delay. In addition, due to the marginal difference between FB and FB^o for Pareto flow sizes, we conclude that utilizing the size-dependent information is more important than utilizing the location information.

5.2 Near-far unfairness

Near-far unfairness refers to the inherent property of the HSPDA/DHR systems that users far away from the base station experience much worse performance than users near the base station. This is also addressed in [3], where it is remarked that size-based schedulers only exacerbate the near-far unfairness property. Thus, next we examine in more detail the near-far unfairness issue in terms of the conditional mean delays of the different disciplines when the receiving terminal is at a given distance from the base station.

The results as a function of the distance r are given in Figure 2 for the exponential distribution (left) and Pareto distributions with $\beta = 3$ (middle) and $\beta = 2$ (right). The results show the normalized conditional delay of the other disciplines with respect to PS, *i.e.*, we plot $E[T_r^\pi]/E[T_r^{\text{PS}}]$ for $\pi = \text{GR, FB}$. Results for SRPT are not shown to keep the figures clear. However, under SRPT the conditional delays are always smaller than under FB by a similar margin as

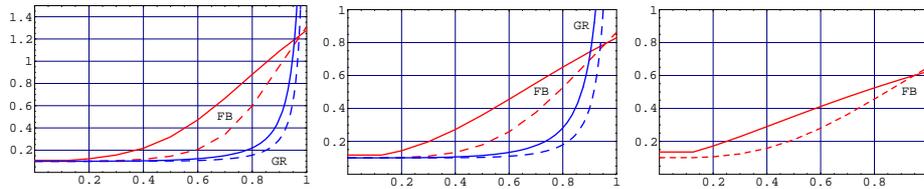


Fig. 2. Conditional mean delay ratio as a function of distance r for exponential distribution (left) and Pareto distributions with $\beta = 3$ (middle) and $\beta = 2$ (right).

in the overall delays in Figure 1. From the figures we can see that GR results in extreme near-far unfairness, *i.e.*, the performance benefits are achieved at the expense of the far users (again, for $\beta = 2$ the GR discipline does not even have a finite mean). On the other hand, the FB discipline makes the performance slightly worse for the users on the border in the exponential case. However, in the Pareto cases the conditional delays are uniformly in the cell better than under PS. Thus, roughly speaking, everybody “wins”.

6 Conclusions

We have analyzed the impact of size-based scheduling on the flow level performance of elastic traffic in the downlink data channels in a single cell. In such systems the service time of a flow is determined both by its size and transmission rate (as determined by the random location of a user in the cell). Both information can be used to minimize the flow level delay. The baseline scheduling discipline was provided by the simple RR policy (corresponding to a PS system in the limit), which does not employ any information about the rates nor the sizes. The other studied schedulers were the GR scheduler (only rate information), plain FB scheduler (distribution of rates and sizes), FB° (distribution of sizes and exact knowledge of rate) and SRPT (exact knowledge of size and rate). Notably, the FB° policy was derived by applying the Gittins index approach, and it yields the optimal non-anticipating scheduling discipline that utilizes exact knowledge of a user’s rate. In the case of Pareto distributions, which have been used to model flow sizes in the Internet [5, 6], the optimal discipline FB° proves to be a simple modification of FB.

The results showed that the GR discipline can sometimes (*e.g.*, for exponentially distributed flow sizes) decrease the overall mean delay compared with RR. However, it is always at the expense of the users that are on the border of the cell. Comparing plain FB and FB° , it appears that for Pareto distributed flow sizes, knowledge of the location is not that important. Additionally, for Pareto distributed files we demonstrated that considerable gains can be achieved by applying size-based scheduling (FB and SRPT), both in terms of the overall mean delay, as well as the conditional mean delay at a given distance from the base station. The latter property implies that age-based scheduling can increase the performance of all users, *i.e.*, everybody wins compared with RR.

In this paper the baseline policy was provided by RR. In practical systems, channel-aware schedulers, such as PF, are used, and they are known to achieve a better performance under fast fading conditions. However, as shown in [3], the performance gain in a dynamic traffic situation may not be that great. Thus, it can be argued that the significant gains achievable with FB may be higher than those achievable with channel-aware schedulers. Nevertheless, a detailed performance comparison with channel-aware scheduling remains as future work. Additionally, multi-cell scenarios offer another area of extensions, as well as the inclusion of coding constraints in the rate function (1).

References

1. Barlow, R.E., Marshall, A.W., Proschan, F.: Properties of probability distributions with monotone hazard rate. *The Annals of Mathematical Statistics* **34** (1963) 375–389.
2. Berggren, F., Jäntti, R.: Asymptotically fair scheduling on fading channels. In: *Proceedings of IEEE VTC 2002-Fall, Vancouver, Canada* (2002) 1934–1938.
3. Bonald, T., Proutière, A.: Wireless downlink data channels: user performance and cell dimensioning. In: *Proceedings of ACM MobiCom 2003, San Diego, CA* (2003) 339–352.
4. Borst, S.: User-level performance of channel-aware scheduling algorithms in wireless data networks. In: *Proceedings of IEEE Infocom 2003, San Francisco, CA* (2003) 321–331.
5. Crovella, M., Bestavros, A.: Self-similarity in world wide web traffic: evidence and possible causes. In: *Proceedings of ACM SIGMETRICS 1996, Philadelphia, PA* (1996) 160–169.
6. Feldmann, A., Whitt, W.: Fitting mixtures of exponentials to long-tail distributions to analyze network performance models. In: *Proceedings of IEEE Infocom 1997, Kobe, Japan* (1997) 1096–1104.
7. Gittins, J.C.: *Multi-armed Bandit Allocation Indices*, Wiley, Chichester (1989).
8. Holzman, J.M.: CDMA forward link waterfilling power control. In: *Proceedings of IEEE VTC 2000-Spring, Tokyo, Japan* (2000) 1663–1667.
9. Hu, M., Zhang, J., Sadowsky, J.: Traffic aided opportunistic scheduling for downlink transmissions: algorithms and performance bounds. In: *Proceedings of IEEE Infocom 2004, Hong Kong* (2004) 1652–1661.
10. Kleinrock, L.: *Queueing Systems, Volume II: Computer Applications*. Wiley, New York (1976).
11. Righter, R., Shanthikumar, J.G.: Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures. *Probability in the Engineering and Informational Sciences* **3** (1989) 323–333.
12. Schrage, L.E., Miller, L.W.: The queue M/G/1 with the shortest remaining processing time discipline. *Operations Research* **14** (1966) 670–683.
13. Schrage, L.E.: A proof of the optimality of the shortest remaining processing time discipline. *Operations Research* **16** (1968) 687–690.
14. Yashkov, S.F.: Processor-sharing queues: Some progress in analysis. *Queueing Systems* **2** (1987) 1–17.
15. Yashkov, S.F.: Mathematical problems in the theory of shared-processor systems. *Journal of Mathematical Sciences* **58** (1992) 101–147.