

Wireless Downlink Data Channels: User Performance and Cell Dimensioning

Thomas Bonald and Alexandre Proutière
France Telecom R&D
38-40 rue du Général Leclerc
92794 Issy-les-Moulineaux, France

{thomas.bonald,alexandre.proutiere}@francetelecom.com

ABSTRACT

We consider **wireless downlink data channels** where the transmission **power** of each base station is **time-shared between a dynamic number of active users** as in CDMA/HDR systems. We derive **analytical results relating user performance**, in terms of **blocking probability and data throughput**, to **cell size and traffic density**. These results are used to address a number of practically interesting issues, including the **trade-off** between **cell coverage and cell capacity** and the choice of efficient scheduling and admission control schemes.

Categories and Subject Descriptors

B.8 [Performance and Reliability]: Performance Analysis and Design Aids

General Terms

Performance

Keywords

CDMA/HDR systems, flow-level analysis, dimensioning

1. INTRODUCTION

Data services are expected to constitute a significant part of traffic in future CDMA networks. In this paper, we derive analytical performance results for downlink data channels, accounting for the random nature of traffic demand, and show the practical interest of these results in dimensioning cells and designing radio control algorithms.

To support high data rates, a number of new technologies have been standardized, such as HDR (High Data Rate) systems [6], corresponding to the CDMA2000 1xEV-DO standard, and their 3GPP equivalent, HSDPA (High Speed Downlink Packet Access) systems [1, 17]. Both systems are based on an **intra-cell interference cancellation** principle: **time is slotted and the base station (BS) transmits at full power** to

only one user in each slot. It has been shown that, assuming the feasible transmission rate is linear in the signal to interference-plus-noise (SINR) ratio, this maximizes overall throughput [18, 11, 4]. Note that this optimality principle is not strictly valid when only a discrete set of data rates is available [4].

A key component of such a TDMA-like scheme is the scheduling, i.e., deciding which user should be served in each time slot. In fact, it is not directly clear what a “good” scheduling strategy is, as the potential data rate of a user depends on her/his radio conditions, mainly determined by the distance to the BS and fading effects. To transmit always to the user with the highest potential rate maximizes overall throughput but typically results in the starvation of distant users. Another strategy, which realizes a reasonable trade-off between efficiency and fairness, consists in transmitting to the user with the highest potential rate *proportionally* to her/his current mean data rate [23]. This algorithm, termed Proportional Fair (PF) and implemented in HDR systems, has indeed been shown to fairly share the transmission resource [10]. Many other scheduling algorithms have been proposed and analyzed (see [8] and references therein).

As a general rule the evaluation of scheduling algorithms is performed with an assumed *static* population of users (see, e.g., [12]). We maintain that this may lead to misleading conclusions since the actual set of active users is *dynamic* and varies as a random process as new data flows are initiated and others complete. In particular, while users are generally assumed to be uniformly distributed in the cell, the location of active users in steady state *does* depend on the scheduling employed. This is due to the inherent “elasticity” of data transfers: the resource attributed to any user determines how long that user will stay active. Thus a scheduling scheme that favors near users results in a large proportion of active users being far from the BS.

A further issue when accounting for the statistical nature of traffic is that of admission control: situations may arise where it is preferable to block new demands rather than to further degrade the performance of ongoing data flows. Admission control has already been proposed for data services in a wireline context to preserve network efficiency in overload [5]. The design and evaluation of admission control schemes for the considered wireless network are still largely unexplored areas.

In evaluating scheduling and admission control it is important to understand how offered traffic impacts user perceived performance, in terms of transfer delays and blocking

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom'03, September 14–19, 2003, San Diego, California, USA.

Copyright 2003 ACM 1-58113-753-2/03/0009 ...\$5.00.

rates. Given the ever changing nature of data applications, control mechanisms should be designed so that performance depends mainly on an appropriately defined load factor, and not on precise traffic characteristics like the distribution of a typical transfer volume. This requirement underscores the interest of analytical modeling tools, in addition to simulation, in order to gain the necessary insight and identify efficient design choices.

Regarding dimensioning issues, it is not straightforward to define the traffic carrying capacity of a CDMA network handling data. Resource consumption depends in particular on the position of the users so that, in addition to the usual notion of intensity, traffic also has a spatial component. It is notably important to know how the capacity of a cell depends on its size, in analogy with the known trade-off in circuit switched CDMA networks [22, 25].

Related work. To our knowledge, very few papers address the issue of user performance for wireless data channels, accounting for the random nature of traffic and the inherent “elasticity” of data transfers. Most existing models indeed represent data transfers as circuit services (see, e.g., [2]). A notable exception is the recent work of Borst for CDMA/HDR systems [8]. User performance is explicitly evaluated and shown to be insensitive to the flow size distribution, with or without admission control, in a symmetric scenario where all users experience the same fast fading and the resource allocation is that realized by the PF scheduler. The radio channel is modeled at flow level by a processor-sharing queue which is indeed known to have the insensitivity property [20].

Contribution. In the present paper, we extend the analytical results of [8] to more general scheduling and admission control schemes and, accounting for the spatial component of offered traffic, apply them to a number of practical issues. We notably define the notion of “cell capacity”, critical for dimensioning purposes, and show that:

- sharing the transmission power in a fair way is efficient with respect to user performance;
- the waste of radio resources due to the granularity of feasible rates induced by coding constraints is typically not significant;
- the so-called “cell breathing” effect arises at very high loads only. In particular, a simple admission control independent of user locations is sufficient in practice;
- the number of active users is typically rather small in steady state. In particular, the impact of opportunistic schedulers like PF that take advantage of fast fading is typically much smaller than one would expect with an assumed static user population.

Outline. In the next section, we present the model used to derive the analytical results. In Sections 3 and 4, user performance is evaluated in terms of throughput and blocking rate. The following three sections are devoted to the impact of non-linear data rate *vs.* SINR dependency, fast fading and interference, respectively. Section 8 concludes the paper.

2. MODEL

In this section, we first present the model of the radio resource and the way this resource is shared. We then describe the characteristics of offered traffic.

2.1 Radio resource

We consider a cell with a single downlink channel whose resource is time-shared between active users. Denote by ϕ_u^b the fraction of time base station (BS) b transmits to user u , with $\sum_u \phi_u^b = 1$. The data rate of user u is then:

$$C_u = C \times \phi_u^b, \quad (1)$$

where C is the peak data rate, obtained in the absence of any other user in the cell, i.e., for $\phi_u^b = 1$.

In practice, the peak data rate depends in a complex way on the radio environment and varies over time due to user mobility, shadowing and multi-path reflections. Unless otherwise specified, we ignore these fading effects, i.e., we assume the peak data rate is approximately constant during data transfer. Section 6 is devoted to the impact of fast fading. To simplify the presentation, we also assume that the peak rate C depends on the distance r from BS b to user u only. This last assumption is not essential and the theoretical results derived in the following still hold in a more realistic radio environment, given any peak rate function of user’s location in the cell. We denote by C_0 the maximum peak rate (which depends on channel bandwidth and coding efficiency) and by r_0 the maximum distance at which this maximum peak rate is achieved:

$$C(r) = C_0 \text{ for all } r \leq r_0. \quad (2)$$

Rate vs. SINR dependency. For numerical applications, we use the peak rate function given by the following standard model. Let W be the cell chip rate. If P_u denotes the power received by user u from her/his BS, η the background noise and I_u the interference due to other BS, the user’s signal to interference-plus-noise ratio (SINR) and energy-per-bit to noise density ratio are respectively given by [24]:

$$\text{SINR}_u = \frac{P_u}{\eta + I_u}, \quad \frac{E_b}{N_0} = \frac{W}{C} \times \text{SINR}_u. \quad (3)$$

Given a target error probability, it is necessary that $E_b/N_0 \geq \delta$ for some threshold δ , which is assumed to be the same for all users. The peak data rate of user u is then the minimum of C_0 and $W \times \text{SINR}_u/\delta$. In particular, it is linear in the SINR up to the maximum peak rate C_0 . The assumption of a constant E_b/N_0 target is generally valid as long as the same type of modulation is used for all data rates [15]. In HDR systems for instance, δ is approximately equal to 2.5dB for all data rates except for the three highest, with a maximum value of 6.5dB [6]. The impact of such a rate-dependent target δ resulting in a non-linear rate *vs.* SINR dependency is evaluated in Section 5.

Propagation model. In the first part of the paper, we neglect the interference term I_u , i.e., we consider a single cell in isolation. Equivalently, we assume that interference is constant over the considered cell. The impact of the interference generated by other BS is evaluated in Section 7. The power P_u received by user u is equal to $P \times \Gamma_u$ where P is the transmission power of the BS and Γ_u denotes the path

loss. In practice, path loss varies with respect to the user's location and radio conditions. Here we adopt a simple propagation model where the path loss Γ_u is a function Γ of the distance r from the BS to user u only:

$$\Gamma(r) = \begin{cases} 1 & \text{if } r \leq \varepsilon, \\ \left(\frac{\varepsilon}{r}\right)^\alpha & \text{otherwise,} \end{cases}$$

where ε denotes the maximum distance at which the full power P is received and α is the path loss exponent which characterizes the radio environment (typical values of α are between 2 and 5). Assuming that the maximum peak rate C_0 can be achieved (thus $r_0 > \varepsilon$), it follows from (2) that the peak rate function is:

$$C(r) = C_0 \times \begin{cases} 1 & \text{if } r \leq r_0, \\ \left(\frac{r_0}{r}\right)^\alpha & \text{otherwise.} \end{cases} \quad (4)$$

Coding constraints. Expression (4) corresponds to an ideal case where the set of achievable peak rates is continuous. In practice, coding constraints result in a discrete set of achievable peak rates $C_0 \equiv c_0 > c_1 > \dots > c_n$. In view of (4), these rates define a set of concentric "rings" of external radius $r_0 < r_1 < \dots < r_n$ corresponding to regions where these rates are achievable. Table 1 below gives the rates defined for HDR channels [6] with the corresponding radius (normalized so that $r_0 = 1$) evaluated from (4) with two values of the path loss exponent.

Ring k	Rate c_k (Kbit/s)	Radius r_k ($\alpha = 4$)	Radius r_k ($\alpha = 2$)
0	2457.6	1	1
1	1843.2	1.07	1.15
2	1228.8	1.19	1.41
3	921.6	1.28	1.63
4	614.4	1.41	2.00
5	307.2	1.68	2.83
6	204.8	1.86	3.46
7	153.6	2.00	4.00
8	102.6	2.21	4.90
9	76.8	2.37	5.61
10	38.4	2.82	7.94

Table 1: Rates and ring radius.

2.2 Traffic characteristics

We assume traffic demand is uniformly distributed in the cell. Data flows arrive as a Poisson process of intensity $\lambda \times ds$ in any area of surface ds . Flow sizes are independent and identically distributed (i.i.d.). We denote by σ the corresponding random variable and by $\rho = \lambda \times E[\sigma]$ the traffic density (in Kbit/s per surface unit). The traffic intensity generated by those users whose distance to the BS is between r and $r + dr$ is $d\rho(r) = \rho \times 2\pi r dr$. We consider two cases, depending on whether coding constraints are taken into account or not.

No coding constraint. In the ideal case where a continuous set of peak rates is achievable, the load generated by those users whose distance to the BS is between r and $r + dr$ is the ratio $d\rho(r)/C(r)$ of their traffic intensity to their peak

rate. This corresponds to their demand on the radio resource (here the time slot). Thus we define the *load* of a cell of radius R as:

$$\bar{\rho} = \int_0^R \frac{d\rho(r)}{C(r)}. \quad (5)$$

Users experience quality of service through the duration of data transfers. Note that, in view of (1), this does not only depend on user characteristics such as her/his peak rate, but also on the cell activity, i.e., on the dynamic number of active users who share the transmission resource. We are interested in the mean flow duration $T(r)$ for a user whose distance to the BS is r (recall that active users don't move during their data transfer in our model). Let $dx(r)$ be the mean number of active users whose distance to the BS is between r and $r + dr$. Applying Little's law [16], we get:

$$dx(r) = T(r)\lambda \times 2\pi r dr.$$

Thus the flow throughput $\gamma(r)$ of users whose distance to the BS is r , defined as the ratio of the mean flow size $E[\sigma]$ to the mean flow duration $T(r)$, is given by:

$$\gamma(r) = \frac{d\rho(r)}{dx(r)}. \quad (6)$$

This quantifies the average performance of data transfers at a distance r from the BS.

Coding constraints. When the set of achievable peak rates is discrete, the traffic intensity in ring 0 is given by:

$$\rho_0 = \rho\pi r_0^2,$$

while the traffic intensity in ring k , $k = 1, \dots, n$, is given by:

$$\rho_k = \rho\pi(r_k^2 - r_{k-1}^2).$$

As above, we define the cell load as:

$$\bar{\rho} = \sum_{k=0}^n \bar{\rho}_k,$$

where $\bar{\rho}_k = \rho_k/c_k$ denotes the load of ring k . The shape and the load contribution of each ring are illustrated in Figure 1 for the values given in Table 1 and $n = 4, 7, 10$, corresponding to cells of radius r_4, r_7, r_{10} , respectively¹. We observe that for large cells (corresponding to a large spectrum of possible peak rates), most load is concentrated in the outer ring.

We are interested in the mean flow duration T_k for a user in ring k . Denote by x_k the mean number of active users in ring k . By Little's law:

$$E[x_0] = T_0\lambda \times \pi r_0^2,$$

and for $k = 1, \dots, n$,

$$E[x_k] = T_k\lambda \times \pi(r_k^2 - r_{k-1}^2).$$

We deduce the flow throughput for a user in ring k , defined as the ratio of the mean flow size $E[\sigma]$ to the mean flow duration T_k :

$$\gamma_k = \frac{\rho_k}{E[x_k]}. \quad (7)$$

¹The maximum cell radius obtained for $\alpha = 2$ is in fact almost three times larger than that obtained for $\alpha = 4$; we changed the scale between the two sets of cells of Figure 1 for sake of readability.

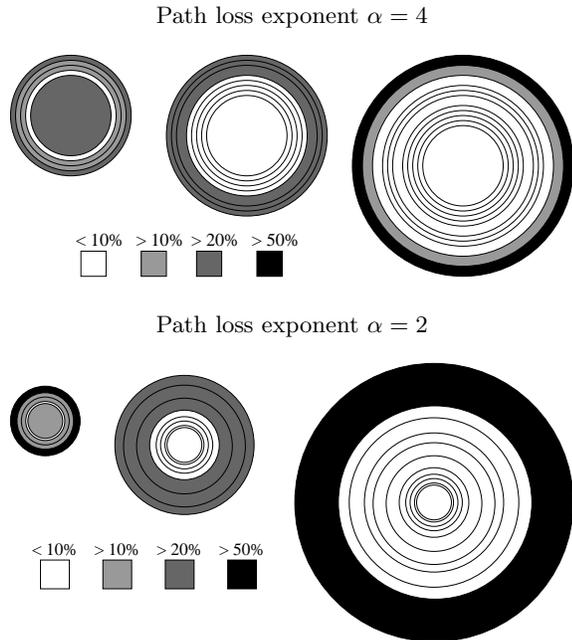


Figure 1: Ring shapes and load distribution.
($n = 4, 7, 10$)

3. THROUGHPUT PERFORMANCE (NO ADMISSION CONTROL)

In this section, we evaluate user performance in terms of flow throughput in the absence of admission control, i.e., when the number of active users is not limited. We first assume that the radio resource is fairly shared between active users, i.e., $\phi_u^b = 1/x$, when x users are active in the cell. This is the allocation realized by the PF scheduler or a simple round-robin scheduler, for instance (these are equivalent in the absence of fast fading). Other power allocations are considered in §3.3.

3.1 A continuous set of peak rates

We first consider the ideal case where a continuous set of peak rates is achievable. As the transmission resource is fairly shared between active users, the number of active users x evolves like the number of customers in a processor-sharing queue with Poisson arrivals of intensity $\lambda\pi R^2$ and i.i.d. service times [16]. Each service time is equal to the flow duration in the absence of any other user in the cell, i.e., $\sigma/C(r)$ for a user whose distance to the BS is r . Thus the distribution of the random variable $\bar{\sigma}$ representing a typical service time is given by:

$$d\bar{\sigma}(r) = \frac{\sigma}{C(r)} \frac{2rdr}{R^2}, \quad r \leq R.$$

In particular, the load of the processor-sharing queue corresponds to the cell load:

$$\lambda\pi R^2 \times \int_0^R \frac{E[\sigma]}{C(r)} \frac{2rdr}{R^2} = \int_0^R \frac{\rho}{C(r)} 2\pi r dr \equiv \bar{\rho}. \quad (8)$$

We conclude that the number of active users x tends to a finite stationary regime in underload ($\bar{\rho} < 1$), while it

grows indefinitely in overload ($\bar{\rho} > 1$). In the latter case, the data rate $C(r)/x$ tends to zero for all users, whatever their distance to the BS: the cell is saturated.

Cell capacity. We may define the cell capacity as the maximum traffic intensity for which the cell is not saturated. In view of (8), this is a function \bar{C} of the cell radius R :

$$\bar{C}(R) = \left(\int_0^R \frac{2rdr}{C(r)R^2} \right)^{-1}.$$

Note that \bar{C} is a decreasing function of R , equal to the maximum peak rate C_0 for $R \leq r_0$. Figure 2 gives the cell capacity with respect to its radius (normalized values so that $C_0 = 1$, $r_0 = 1$) for the peak rate function (4). We observe that the cell capacity decreases suddenly when $R \geq r_0$. The capacity of large cells is extremely small.

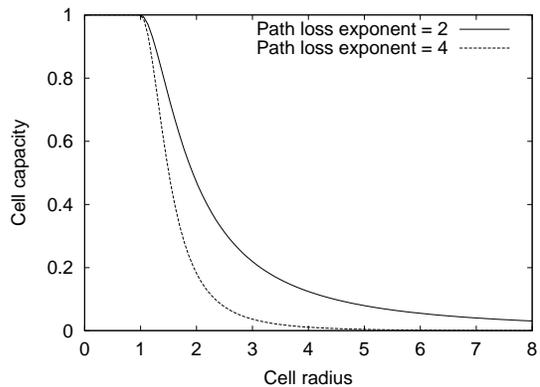


Figure 2: Cell capacity – defined as the maximum traffic intensity without saturation – with respect to cell radius, for path loss exponents $\alpha = 2$ and $\alpha = 4$.

Flow throughput. To evaluate user performance in underload, we use a key property of the processor-sharing queue: the stationary distribution of the number of customers is insensitive to the distribution of service times. This follows from the reversibility of the underlying Markov process [14, 20]. We deduce that the stationary distribution π of the number of active users is insensitive to the flow size distribution and given by:

$$\pi(x) = \bar{\rho}^x (1 - \bar{\rho}). \quad (9)$$

Note that the number of active users is typically rather small in steady state: the probability having more than x active users decreases geometrically with rate $\bar{\rho}$. In addition, the probability that the distance of an active user to the BS is between r and $r+dr$ is proportional to the load generated by these users, namely $d\rho(r)/C(r)$. In particular, the density of active users is inversely proportional to their peak rate $C(r)$. We deduce:

$$dx(r) = \frac{d\rho(r)}{\bar{\rho}C(r)} \times E[x],$$

where, in view of (9), the mean number of active users is:

$$E[x] = \frac{\bar{\rho}}{1 - \bar{\rho}}.$$

It then follows from (6) that:

$$\gamma(r) = C(r)(1 - \bar{\rho}).$$

Hence, the flow throughput is equal to the peak rate for $\bar{\rho} = 0$ and decreases *linearly* in the cell load.

3.2 A discrete set of peak rates

We now consider the practically interesting case where only a discrete set of peak rates is available. As above, the number of active users evolves like the number of customers in a processor-sharing queue with Poisson arrivals of rate $\lambda\pi R^2$ and i.i.d. service times. The random variable $\bar{\sigma}$ representing a typical service time is now equal to σ/c_k with probability p_k , $k = 0, 1, \dots, n$, with:

$$p_0 = \frac{r_0^2}{R^2}, \quad p_k = \frac{r_k^2 - r_{k-1}^2}{R^2}, \quad k = 1, \dots, n.$$

Again, the load of the processor-sharing queue corresponds to the cell load:

$$\lambda\pi R^2 \times \sum_{k=0}^n p_k \frac{E[\sigma]}{c_k} \equiv \bar{\rho}.$$

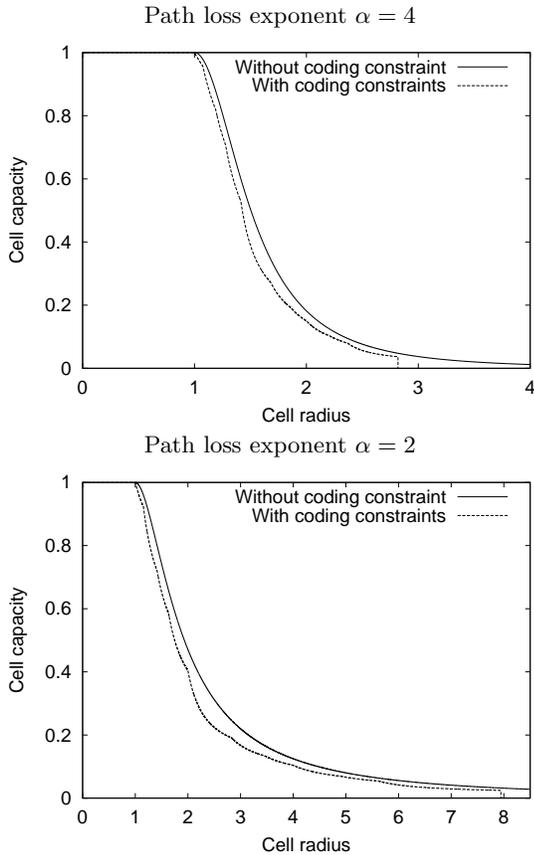


Figure 3: Cell capacity – defined as the maximum traffic intensity without saturation – with respect to cell radius, with and without coding constraints.

Cell capacity. As above, we define the cell capacity as maximum traffic intensity for which the cell is not saturated:

$$\bar{C}(R) = \left(\sum_{k=0}^n \frac{p_k}{c_k} \right)^{-1}.$$

Note that the cell capacity is always smaller than that obtained in the ideal case considered in §3.1. Figure 3 gives the cell capacity obtained for the values of Table 1. We observe that the difference with the ideal case is very small, indicating that resource wastage due to coding constraints is limited.

Flow throughput. The stationary distribution of the number of active users in each ring is that of the number of customers in a processor-sharing queue with classes of respective loads $\bar{\rho}_0, \dots, \bar{\rho}_n$:

$$\pi(x_0, \dots, x_n) = \frac{(x_0 + \dots + x_n)!}{x_0! \dots x_n!} \bar{\rho}_0^{x_0} \dots \bar{\rho}_n^{x_n} (1 - \bar{\rho}). \quad (10)$$

Again, the probability having more than x active users decreases geometrically with rate $\bar{\rho} = \bar{\rho}_0 + \dots + \bar{\rho}_n$. The probability an active user is in ring k is proportional to the load $\bar{\rho}_k = \rho_k/c_k$ generated by these users. In particular, the *density* of active users is inversely proportional to their peak rate, as in the absence of coding constraint. For instance, for the values given in Table 1, the mean number of active users per surface unit in ring 5 is always twice that in ring 4, whatever the values of the path loss exponent or the cell load.

In view of (10), the mean number of active users in ring k is:

$$E[x_k] = \frac{\bar{\rho}_k}{1 - \bar{\rho}},$$

and from (7):

$$\gamma_k = c_k(1 - \bar{\rho}).$$

Again, the flow throughput decreases *linearly* in the cell load. This is illustrated in Figure 4 for each ring of Table 1.

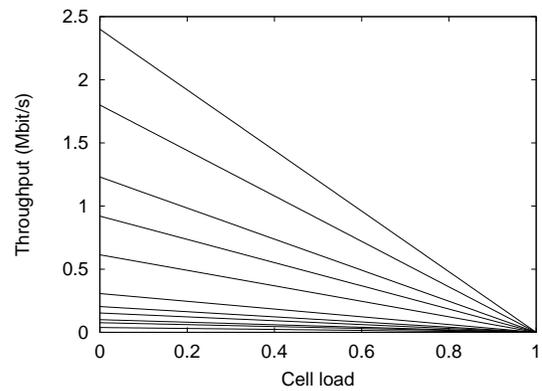


Figure 4: Flow throughput with respect to cell load for the 11 rings of Table 1.

3.3 Other power allocations

The allocation so far considered, referred to as “fair power” sharing, has the practically interesting property that user

performance can be explicitly evaluated and is insensitive to the flow size distribution. It may be considered as “unfair”, however, as the flow throughput experienced by a user is proportional to her/his peak data rate (refer to Figure 4).

Fair rate sharing. Consider another allocation, referred to as “fair rate” sharing, where the data rates of all active users are made equal. This is realized if BS b transmits to user u a fraction of time inversely proportional to her/his peak rate:

$$\phi_u^b = \frac{1/C(r(u))}{\sum_{u'} 1/C(r(u'))},$$

where $r(u')$ denotes the distance from BS b to user u' . The number of active users then evolves like the number of customers in a *discriminatory* processor-sharing queue [9] of load $\bar{\rho}$. Again, the cell is saturated in overload ($\bar{\rho} > 1$). In underload ($\bar{\rho} < 1$), the number of active users x tends to a finite stationary regime, but the stationary distribution is *sensitive* to the flow size distribution.

For an exponential flow size distribution, we obtain using [9] the results of Figure 5 for a 5-ring cell, with the values of Table 1. We observe that the gain in flow throughput for users in the outer ring is very limited. On the other hand, the impact on the performance of users in the inner disk is significant. This may be explained as follows. First, the fact that *active* users have the same data rate does not imply that users have the same flow throughput. When the cell load $\bar{\rho}$ is close to 0 for instance, an active user is typically alone in the cell so that the flow throughput is equal to the peak rate. Second, most load is concentrated in the outer rings (refer to Figure 1): the performance experienced by far users is essentially due to their own load and cannot be significantly improved; the performance experienced by near users, on the other hand, mainly depends on the load generated by far users thus may be significantly worsened.

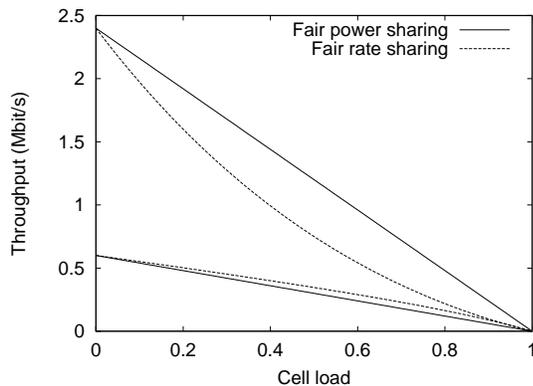


Figure 5: Throughput of users in the inner disk (upper curves) and the outer ring (lower curves) with respect to cell load in a 5-ring cell, for two different power allocations.

Service-dependent allocations. Other possible allocations are those which depend on the remaining or the processed service, such as the “shortest remaining processing time” or the “foreground-background” disciplines, respectively [16].

These allocations, which are known to outperform the processor sharing discipline for a heavy-tailed service distribution, have been proposed to be implemented in Web servers in a wireline context [3, 19] and more recently in scheduling algorithms in a wireless context [13]. They would here exacerbate the discrimination against far users, whose processing time is typically much higher than that of near users.

The above observations suggest that no significant gain is achieved by those allocations which do not share the transmission resource in a fair way and that the so-called “near-far” unfairness (in terms of flow throughput) is inherent to wireless data systems. The only way to achieve approximately fair throughput performance is to limit the cell size so that all users can achieve high data rates. Finally, it is worth noting that the key properties satisfied by the “fair power” allocation (explicit performance evaluation and insensitivity) still hold in the presence of admission control, as shown in the next section. The impact of admission control on user performance is typically extremely difficult to evaluate for other power allocations.

4. BLOCKING RATE (WITH ADMISSION CONTROL)

In the absence of admission control, we have seen that the cell is saturated in overload ($\bar{\rho} > 1$): the number of ongoing transfers grows indefinitely and the data rate of each transfer eventually tends to zero. Admission control is necessary to guarantee a minimum data rate c_{\min} for all users whatever the cell load. For a given target rate c_{\min} , $0 < c_{\min} < C_0$, the cell radius R cannot exceed R_{\max} , the maximum distance r such that $C(r) \geq c_{\min}$.

In the presence of admission control, users experience quality of service not only through flow throughput but also through blocking rate. In this section, cell capacity is evaluated in terms of the maximum traffic intensity for given minimum data rate c_{\min} and target blocking probability.

4.1 Admission control based on the number of active users

We first consider a simple admission criterion based on the number of active users. For a cell of radius R , the minimum rate c_{\min} is guaranteed if the number of active users does not exceed:

$$m = \frac{C(R)}{c_{\min}}.$$

No coding constraint. We first consider the case where the set of available peak rates is continuous. As the Markov process associated with the unconstrained system considered in §3.1 is reversible, the restriction of its stationary distribution to admissible states gives the stationary distribution of the Markov process associated with the constrained system [14]. In addition, the stationary distribution associated with the constrained system remains *insensitive* to the service time distribution. Assuming m is an integer for simplicity, it then follows from (9) that the stationary distribution of the number of active users is:

$$\pi(x) = \frac{\bar{\rho}^x}{1 + \bar{\rho} + \dots + \bar{\rho}^m}, \quad 0 \leq x \leq m. \quad (11)$$

Thus the blocking rate, which is independent of the user's distance to the BS, is given by:

$$B = \frac{\bar{\rho}^m}{1 + \bar{\rho} + \dots + \bar{\rho}^m}. \quad (12)$$

In addition, the probability that the distance of an active user to the BS is between r and $r + dr$ is proportional to the *actual* load generated by these users, namely:

$$(1 - B) \frac{d\rho(r)}{C(r)}.$$

In particular, the *density* of active users is inversely proportional to their peak rate $C(r)$ as in the absence of admission control (refer to §3.1). We also deduce:

$$dx(r) = \frac{1}{\bar{\rho}} \frac{d\rho(r)}{C(r)} \times E[x],$$

where, in view of (11), the mean number of active users is:

$$E[x] = \frac{\bar{\rho}}{1 - \bar{\rho}} \times \frac{1 - (m + 1)\bar{\rho}^m + m\bar{\rho}^{m+1}}{1 - \bar{\rho}^{m+1}}. \quad (13)$$

As the *actual* traffic intensity generated by those users whose distance to the BS is between r and $r + dr$ is $(1 - B)d\rho(r)$, expression (6) becomes:

$$\gamma(r) = (1 - B) \frac{d\rho(r)}{dx(r)}.$$

It then follows from (12) and (13) that:

$$\gamma(r) = C(r) \times \frac{(1 - \bar{\rho})(1 - \bar{\rho}^m)}{1 - (m + 1)\bar{\rho}^m + m\bar{\rho}^{m+1}}.$$

Hence, the flow throughput is equal to the peak rate for $\bar{\rho} = 0$ and decreases to $C(r)/m$ when the cell load $\bar{\rho}$ tends to infinity.

Coding constraints. The results are similar if the set of available peak rates is discrete. For a cell of radius $R = r_n$, the maximum number of users is $m = c_n/c_{\min}$. The blocking rate is the same in all rings and given by (12). The flow throughput in ring k is:

$$\gamma_k = c_k \times \frac{(1 - \bar{\rho})(1 - \bar{\rho}^m)}{1 - (m + 1)\bar{\rho}^m + m\bar{\rho}^{m+1}}.$$

Figure 6 gives the corresponding flow throughput for a 5-ring cell with the values of Table 1 and a maximum number of users $m = 12$, corresponding to a minimum data rate $c_{\min} = 51$ Kbit/s ($c_{\min}/C_0 \approx 0.02$). The flow throughput decreases from c_k to c_k/m for each ring k .

4.2 Admission control based on the minimum data rate

We now consider an admission control based on the minimum data rate: a new data transfer is accepted if and only if its rate and the rate of ongoing transfers would be larger than c_{\min} . In particular, the admission decision now depends on the location of users. The number of users cannot exceed:

$$M = \frac{C_0}{c_{\min}},$$

which corresponds to the best case where the distance from the BS to any active user does not exceed r_0 .

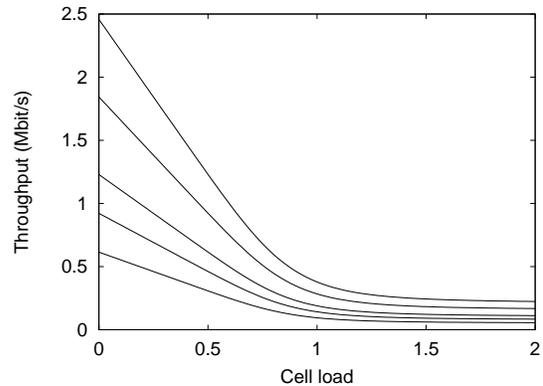


Figure 6: Flow throughput with respect to cell load in each ring of a 5-ring cell when the number of active users is limited to $m = 12$.

The objective here is to evaluate the gain in terms of blocking probability compared to the admission control based on a maximum number of active users m given by:

$$m = \frac{C(R)}{c_{\min}}.$$

No coding constraint. For a continuous set of available peak rates, admissible states are those for which

$$\frac{\min_{1 \leq u \leq x} C(r(u))}{x} \geq c_{\min},$$

where $r(u)$ denotes the distance from the BS to active user u . Assume M is an integer for simplicity. Let $L = M - m$, and for $j = 0, \dots, L$, define R_j as the maximum distance r such that:

$$C(r) \geq (M - j)c_{\min}.$$

Note that $r_0 \equiv R_0 < R_1 < \dots < R_L \equiv R$. Admissible states are those for which there are less than m active users or there are $x = M - j$ active users and the distance from the BS to any of these users does not exceed R_j , $j = 0, \dots, L$.

As in §4.1, the reversibility property implies that the stationary distribution π of the number of active users is insensitive to the flow size distribution and given by the restriction to admissible states of the stationary distribution (9) associated with the unconstrained system. Noting that in this unconstrained system, the probability that the distance of an active user to the BS is less than R_j is equal to $\bar{\rho}_{\rightarrow j}/\bar{\rho}$, where

$$\bar{\rho}_{\rightarrow j} = \int_0^{R_j} \frac{d\rho(r)}{C(r)},$$

we deduce:

$$\pi(x) = \pi(0) \times \begin{cases} \bar{\rho}^x & \text{for } x < m, \\ \bar{\rho}_{\rightarrow i}^{M-i} & \text{for } x = M - i, i = 0, \dots, L, \end{cases}$$

where

$$\pi(0) = \left(1 + \bar{\rho} + \dots + \bar{\rho}^{m-1} + \sum_{i=0}^L \bar{\rho}_{\rightarrow i}^{M-i} \right)^{-1}.$$

We also deduce the blocking rate B_j for users whose distance to the BS is R_j , $j = 0, \dots, L$:

$$B_j = \frac{\sum_{i=0}^L \bar{\rho}_{\rightarrow i}^{M-i} - \sum_{i=j+1}^L \bar{\rho}_{\rightarrow i-1}^{M-i}}{1 + \bar{\rho} + \dots + \bar{\rho}^{m-1} + \sum_{i=0}^L \bar{\rho}_{\rightarrow i}^{M-i}}.$$

Noting that the blocking rate $B(r)$ for users whose distance to the BS is r is

$$B(r) = B_0 \text{ if } r \leq R_0,$$

and for $j = 1, \dots, L$,

$$B(r) = B_j \text{ if } R_{j-1} < r \leq R_j,$$

we get the mean blocking rate:

$$B = B_0 \frac{R_0^2}{R^2} + \sum_{j=1}^L B_j \frac{R_j^2 - R_{j-1}^2}{R^2}. \quad (14)$$

Figure 7 gives the corresponding cell capacity (normalized values so that $C_0 = 1$, $r_0 = 1$) compared to that obtained with the admission control based on a maximum number of users m . We observe that no significant gain is achieved. This is actually true for any reasonable values of the minimum data rate ($c_{\min} < 0.1$, say, corresponding to 246 Kbit/s for HDR channels) and target blocking rates ($B < 10\%$, say). We explain this result in §4.3 below.

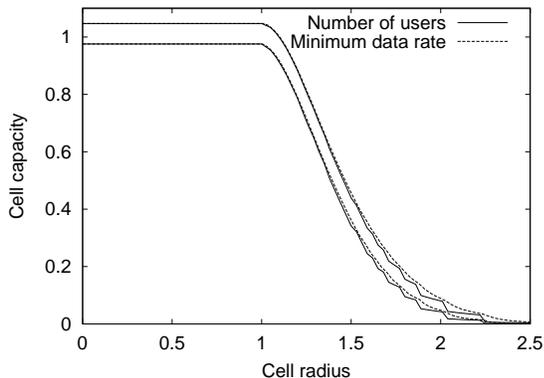


Figure 7: Comparison of the cell capacity – defined as the maximum traffic intensity for a minimum data rate $c_{\min} = 0.02$ and target blocking rates 1% (lower curves) and 5% (upper curves) – obtained for two different admission criteria, without coding constraints (path loss exponent $\alpha = 4$).

Coding constraints. For a discrete set of peak rates, admissible states are those for which

$$\frac{\min_{0 \leq k \leq n} c_k \mathbb{I}_{\{x_k > 0\}}}{x} \geq c_{\min},$$

where x_k denotes the number of active users in ring k . For $j = 0, \dots, L$, define k_j as the largest integer k such that:

$$c_k \geq (M - j)c_{\min}.$$

Note that $0 \equiv k_0 \leq k_1 \leq \dots \leq k_L \equiv n$. Admissible states are those for which there are less than m active users or there are $x = M - j$ active users and the distance from the BS to any of these users does not exceed r_{k_j} , $j = 0, \dots, L$.

Again, it follows from the reversibility of the Markov process associated with the unconstrained system considered in §3.2 that the stationary distribution of the number of active users is:

$$\pi(x) = \pi(0) \times \begin{cases} \bar{\rho}^x & \text{for } x < m, \\ \bar{\rho}_{\rightarrow k_i}^{M-i} & \text{for } x = M - i, i = 0, \dots, L, \end{cases}$$

where

$$\pi(0) = \left(1 + \bar{\rho} + \dots + \bar{\rho}^{m-1} + \sum_{i=0}^L \bar{\rho}_{\rightarrow k_i}^{M-i} \right)^{-1},$$

and for $k = 0, \dots, n$,

$$\bar{\rho}_{\rightarrow k} = \sum_{j=0}^k \bar{\rho}_j.$$

The blocking rate in ring k is given by:

$$B_k = \frac{\sum_{i=0}^L \bar{\rho}_{\rightarrow k_i}^{M-i} - \sum_{i=j_k+1}^L \bar{\rho}_{\rightarrow k_i-1}^{M-i}}{1 + \bar{\rho} + \dots + \bar{\rho}^{m-1} + \sum_{i=0}^L \bar{\rho}_{\rightarrow k_i}^{M-i}},$$

where j_k denotes the minimum integer j' such that $k_{j'} = k$, and the mean blocking rate is

$$B = B_0 \frac{r_0^2}{R^2} + \sum_{k=1}^n B_k \frac{r_k^2 - r_{k-1}^2}{R^2}. \quad (15)$$

Figure 8 is the analog of Figure 7 for the values of Table 1. Again, we observe that no significant gain is achieved with the admission control based on the minimum data rate. Furthermore, comparing Figures 7 and 8 indicates as in §3.2 that resource wastage due to coding constraints is limited.

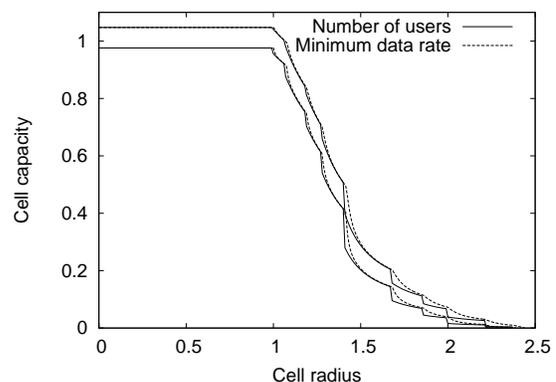


Figure 8: Comparison of the cell capacity – defined as the maximum traffic intensity for a minimum data rate $c_{\min} = 0.02$ and target blocking rates 1% (lower curves) and 5% (upper curves) – obtained for two different admission criteria, with coding constraints (path loss exponent $\alpha = 4$).

4.3 Cell breathing

Previous results suggest that a simple admission control based on the number of active users leads to a cell capacity similar to that obtained with more complex schemes such as that based on the minimum data rate where the admission decision depends on user locations. As the blocking rate $B(r)$ obtained with the latter is an increasing function of the distance r from the BS b to user u , one would expect the cell to “breathe”, i.e., the active users to be near the BS when the cell load increases (refer to Figure 9).

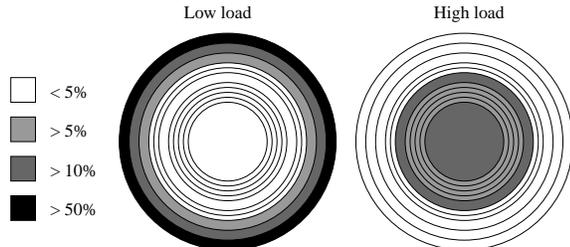


Figure 9: Cell breathing: repartition of active users with respect to cell load ($\alpha = 4$).

This “cell breathing” effect indeed arises as for circuit services [25], but at very high loads only (the results of Figure 9 were obtained for loads $\bar{\rho} = 0.5$ and $\bar{\rho} = 100$). This is illustrated by Figure 10 which gives the mean blocking rate and the blocking rate in each ring with respect to the cell load for the values of Table 1 ($\alpha = 4$): for a target mean blocking rate smaller than 10%, the blocking rate is approximately the same in all rings, and very well approximated by expression (12).

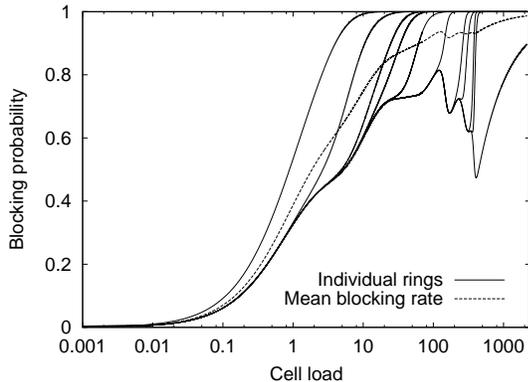


Figure 10: Blocking rate with respect to cell load for the 11 rings of Table 1 with a minimum data rate of 19.2 Kbit/s ($c_{\min} = c_{10}/2$).

The above observation can be explained by the fact that a significant part of the traffic intensity is generated in the outer rings: if the blocking rate in these rings is close to 1, the *mean* blocking rate is necessarily high (larger than the fraction of traffic intensity generated in these rings). At nominal loads corresponding to mean blocking rates smaller than 10%, the blocking rate is approximately the same for

all users whatever their location: the so-called “near-far” unfairness (in terms of blocking rates) does not hold.

5. IMPACT OF NON-LINEAR RATE VS. SINR DEPENDENCY

We have so far assumed that the target energy-per-bit to noise density ratio δ is constant, resulting in a linear rate vs. SINR dependency (up to the maximum peak rate C_0). As noted in §2.1, this assumption is not necessarily valid in real systems. In this section, we study the impact of the rate-dependent target δ of HDR channels.

5.1 Cell shape and load distribution

Table 2 below gives the target SINR for each data rate defined in HDR systems [6]. As the chip rate is $W = 1228.8$ Kchip/s, we get from (3) the corresponding Eb/N_0 target: $\delta_0 \approx 6.5$ dB, $\delta_1 \approx 5.4$ dB, $\delta_2 \approx 3$ dB, and $\delta_k \approx 2.5$ dB for all rings $k = 3, \dots, 10$. Thus the Eb/N_0 target is approximately constant except for the highest peak rates. This can notably be explained by the specific modulation used for these rates.

Ring	Rate (Kbit/s)	SINR (dB)	Radius ($\alpha = 4$)	Radius ($\alpha = 2$)
0	2457.6	9.5	0.79	0.63
1	1843.2	7.2	0.91	0.82
2	1228.8	3.0	1.16	1.34
3	921.6	1.3	1.27	1.62
4	614.4	-1.0	1.45	2.12
5	307.2	-4.0	1.73	2.99
6	204.8	-5.7	1.91	3.64
7	153.6	-6.5	2.00	3.99
8	102.6	-8.5	2.24	5.02
9	76.8	-9.5	2.37	5.63
10	38.4	-12.5	2.82	7.95

Table 2: Rates, target SINRs and ring radius for HDR channels.

In view of the propagation model of §2.1, the SINR is inversely proportional to r^α (for $r > \epsilon$). Table 2 gives the corresponding external ring radius, normalized so that r_0 would be equal to 1 if the Eb/N_0 target were constant and equal to 2.5dB. Comparing with the values of Table 1, we verify that the rate-dependent target δ of HDR channels essentially impacts the inner rings. In particular, the load distribution remains approximately the same, most load being concentrated in outer rings.

5.2 Cell capacity

As the load distribution is not significantly affected by the rate-dependent target δ , we expect the qualitative results of Section 4 to hold. In particular, the “cell breathing” effect arises at very high loads only and an admission decision based on the number of active users is sufficient. This is confirmed by the results of Figure 11 for instance, which show that no significant gain is achieved by an admission control based on the minimum data rate. We also observe that the impact of the rate-dependent target δ on cell capacity is mainly due to the reduction of inner rings (the cell capacity for a constant Eb/N_0 target is here defined for an admission control based on the number of users).

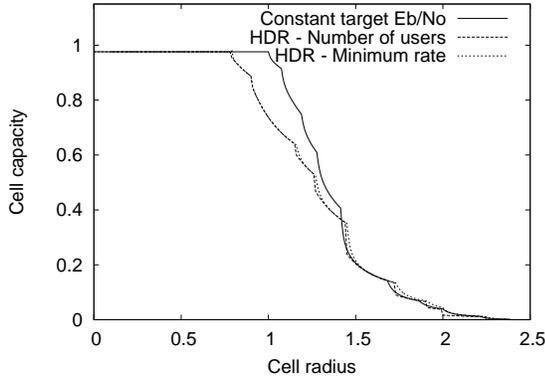


Figure 11: Comparison of the cell capacity – defined as the maximum traffic intensity for a minimum data rate $c_{\min} = 0.02$ and a target blocking rate 1% – obtained for a constant target energy-per-bit to noise ratio and for HDR channels with two admission criteria (path loss exponent $\alpha = 4$).

6. IMPACT OF FAST FADING

We now study the impact of fast fading on cell capacity. Fast fading is an extremely complex physical phenomenon involving multi-path reflections [21]. Here we assume that all users experience the same fast fading in the sense that the SINR of user u at time t is $\xi_u(t) \times \text{SINR}_u$ where $\xi_u(t)$ are i.i.d. copies of some stationary process $\xi(t)$ with unit mean and SINR_u denotes the SINR user u would get in the absence of fast fading. For numerical applications, we take the standard Rayleigh fading corresponding to an exponential marginal distribution of the process $\xi(t)$.

In the rest of the paper, we consider the ideal case where a continuous set of peak rates is available; we have verified as above that resource wastage due to coding constraints is limited. We also consider an admission control based on the number of users only; again, the impact on cell capacity of an admission decision based on the minimum data rate is limited.

6.1 Round-robin scheduling

We first consider a round-robin scheduling, that does *not* take advantage of the time-varying radio conditions of each user. Averaging over the fast fading variations, it follows from (4) that the data rate a user would get at distance r if she/he were alone in the cell is given by:

$$C(r) = E[C_0 \times \min(\xi \left(\frac{r_0}{r}\right)^\alpha, 1)],$$

where $\xi \equiv \xi(0)$ corresponds to the marginal distribution of the fading process. In particular, the blocking rate is given by (12), for the corresponding cell load:

$$\bar{\rho} = \int_0^R \frac{d\rho(r)}{C(r)}.$$

As illustrated in Figure 12 for Rayleigh fading, the impact of fast fading on cell capacity is significant for a cell radius $R \approx 1$ only (recall the convention $r_0 = 1$). This can be explained by the fact that, due to the maximum peak rate C_0 , the positive effects of fast fading do not compensate

its negative effects for those users whose distance to the BS is around r_0 . The nearest users almost always get the maximum peak rate C_0 while the farthest users almost never get the maximum peak rate C_0 : the former are insensitive to fast fading while for the latter, the negative and positive effects of fast fading cancel out.

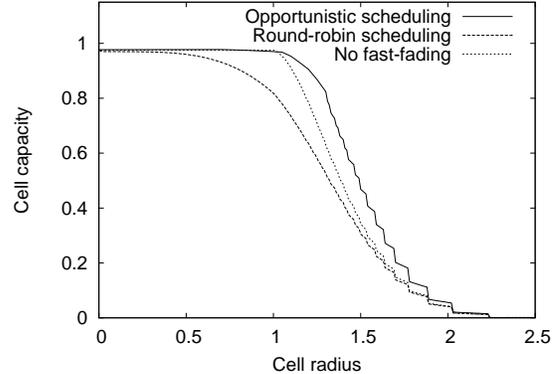


Figure 12: Impact of fast fading on the cell capacity, defined as the maximum traffic intensity for a minimum data rate $c_{\min} = 0.02$ and a target blocking rate 1% (Rayleigh fading, path loss exponent $\alpha = 4$).

6.2 Opportunistic scheduling

A number of so-called “opportunistic” schedulers have been proposed to take advantage of fast fading [23, 7]. The principle is to transmit to the various users when their radio conditions are relatively favorable, while ensuring fair access to the transmission resource. A typical example is the PF scheduler mentioned in Section 1. The impact of opportunistic scheduling on user performance has been evaluated in [8] in a symmetric scenario where fast fading equally impacts the data rates of active users. As observed above, fast fading does not affect all users in the same way, however, due notably to the maximum peak rate C_0 .

To simplify the analysis, we consider three user classes:

- near users, whose distance to the BS is $r < r'$ and for which the peak rate is almost always equal to C_0 (with probability $> 95\%$, say);
- far users, whose distance to the BS is $r > r''$ and for which the peak rate is almost never equal to C_0 (also with probability $> 95\%$);
- the other users, whose distance to the BS is $r \in [r', r'']$ and for which the peak rate is equal to C_0 with non-negligible probability.

For instance, the impact of Rayleigh fading on SINR is less than 4.8dB (resp. larger than -13 dB) with probability $> 95\%$: for a path loss exponent $\alpha = 4$, we deduce the corresponding distances: $r' \approx 0.47$, $r'' \approx 1.3$.

We denote by x the number of near users, x'' the number of far users, and x' the number of other users. For near users, we have $C(r) \approx C_0$ and there is no scheduling gain. For far users, we have $C(r) \approx C_0 \times \left(\frac{r_0}{r}\right)^\alpha$ and the distribution of the feasible rate is approximately that of $\xi \times C(r)$. The

transmission rate of any user at distance $r > r''$ from the BS is then:

$$G'' \times \frac{C(r)}{x + x' + x''},$$

where G'' denotes the scheduling gain. In the presence of i far users, a conservative approximation of the scheduling gain $G''(i)$ is what one would obtain in the absence of any other users, corresponding to a symmetric scenario as considered in [8]. Denoting by $\xi_1, \xi_2, \dots, \xi_i$ i.i.d. copies of ξ , we deduce:

$$G''(i) = E[\max(\xi_1, \xi_2, \dots, \xi_i)].$$

For Rayleigh fading, we obtain [7]:

$$G''(i) = 1 + \frac{1}{2} + \dots + \frac{1}{i}.$$

For the other users, we assume that the distribution of the feasible rate is approximately the same, given by:

$$\xi' \times C' = \int_{r'}^{r''} C_0 \times \min\left(\xi \left(\frac{r_0}{r}\right)^\alpha, 1\right) \frac{2rdr}{r'^{1/2} - r'^2},$$

where ξ' is a unit mean random variable representing the variations around the mean rate C' . The transmission rate of any user at distance $r \in [r', r'']$ from the BS is then:

$$G' \times \frac{C'}{x + x' + x''},$$

where G' denotes the scheduling gain. In the presence of i such users, a conservative approximation of the scheduling gain $G'(i)$ is what one would obtain in the absence of any other users, corresponding again to a symmetric scenario as considered in [8]. Denoting by $\xi'_1, \xi'_2, \dots, \xi'_i$ i.i.d. copies of ξ' , we deduce:

$$G'(i) = E[\max(\xi'_1, \xi'_2, \dots, \xi'_i)].$$

The stationary distribution of the number of active users of each class is insensitive to the flow size distribution and given by:

$$\pi(x, x', x'') = \pi(0) \frac{(x + x' + x'')!}{x!x'!x''!} \bar{\rho}^x \prod_{i=1}^{x'} \frac{\bar{\rho}'}{G'(i)} \prod_{i=1}^{x''} \frac{\bar{\rho}''}{G''(i)}.$$

where $\pi(0)$ follows from the usual normalizing condition and $\bar{\rho}, \bar{\rho}'$ and $\bar{\rho}''$ denote the corresponding class loads:

$$\bar{\rho} = \int_0^{r'} \frac{d\rho(r)}{C_0}, \bar{\rho}' = \int_{r'}^{r''} \frac{d\rho(r)}{C'}, \bar{\rho}'' = \int_{r''}^R \frac{d\rho(r)}{C_0 \times \left(\frac{r_0}{r}\right)^\alpha}.$$

Denoting by m the maximum number of users, the blocking rate is independent of user location and given by:

$$B = \frac{\sum_{x+x'+x''=m} \pi(x, x', x'')}{\sum_{x+x'+x'' \leq m} \pi(x, x', x'')}. \quad (16)$$

As illustrated in Figure 12 for Rayleigh fading, the impact of opportunistic scheduling on cell capacity is relatively limited, especially for large cells. This is a rather counter-intuitive result in view of the high scheduling gains (e.g., $G'(i) \approx 1.5$ and $G''(i) \approx 2.9$ for $i = 10$). This may notably be explained by the fact that, even in the absence of scheduling gain and admission control, the number of active users is typically rather small in steady state (cf. Section 3). The number of active users is here further limited by admission control, especially for large cells.

7. IMPACT OF INTERFERENCE

In order to assess the impact of interference on cell capacity, we consider two types of homogeneous networks: linear networks, where BS are equidistant and placed on a common infinite line; hexagonal networks, where cells are hexagons of the same size and cover the entire plane. In both cases, denote by $2 \times R$ the distance between two BS. We assume that the BS are always active and transmit at the same power P . The interference suffered by user u served by BS b is then:

$$I_u = P \times \sum_{b' \neq b} \Gamma(r_u^{b'}),$$

where $r_u^{b'}$ denotes the distance from BS b' to user u . Let $\bar{\eta} = \eta/P$ and $\bar{I} = I_u/P$, and denote by r the distance from BS b to user u . The peak data rate of user u is then the minimum of C_0 and:

$$\frac{W}{\delta} \times \frac{\Gamma(r)}{\bar{\eta} + \bar{I}}.$$

To compare the results with those obtained for isolated cells, we still denote by r_0 the maximum distance at which this maximum peak rate is achieved in the absence of interference, i.e.,

$$C_0 = \frac{W}{\delta} \times \frac{\Gamma(r_0)}{\bar{\eta}}. \quad (17)$$

7.1 Linear networks

We first consider a linear network. A reasonable approximation consists in considering that interference is generated by the 2 closest BS only. The interference term \bar{I} is then:

$$\bar{I}(r) = \Gamma(2R - r) + \Gamma(2R + r).$$

We deduce the peak rate function:

$$C(r) = \min\left(C_0, \frac{W}{\delta} \times \frac{\Gamma(r)}{\bar{\eta} + \bar{I}(r)}\right).$$

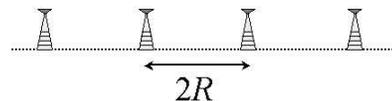


Figure 13: A linear network.

Note that, in view of (17), this function is entirely determined by the maximum peak rate C_0 , the distance r_0 and the ratio W/δ . Previous results still hold, with the cell load given by:

$$\bar{\rho} = \int_0^R \frac{\rho 2dr}{C(r)}. \quad (18)$$

The cell capacity, defined as the maximum traffic intensity without saturation, is given by:

$$\bar{C}(R) = \left(\int_0^R \frac{dr}{C(r)R} \right)^{-1}.$$

Figure 14 gives the cell capacity with respect to the cell radius (normalized values $C_0 = 1$, $r_0 = 1$) for HDR parameters ($W = C_0/2$, $\delta = 2.5$ dB) and negligible ϵ . This is a

decreasing function of the cell radius, with a maximum cell capacity equal to:

$$\bar{C}(0) \approx \begin{cases} 0.86 & \text{for } \alpha = 4, \\ 0.70 & \text{for } \alpha = 2. \end{cases}$$

Thus the impact of interference on cell capacity is significant, especially for small cells. This is notably due to the fact that, for HDR parameters, the maximum peak rate C_0 is achievable at a distance r strictly smaller than R , whatever the cell radius R .

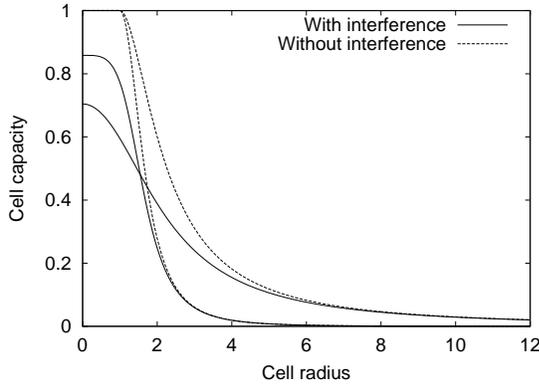


Figure 14: Cell capacity in a linear network – defined as the maximum traffic intensity without saturation for path loss exponent $\alpha = 4$ (lower curves) and $\alpha = 2$ (upper curves) – with and without interference.

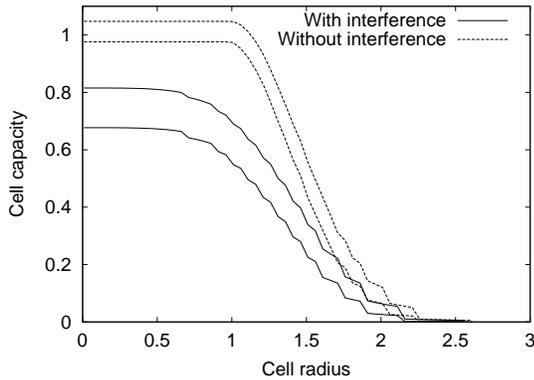


Figure 15: Cell capacity in a linear network – defined as the maximum traffic intensity for a minimum data rate $c_{\min} = 0.02$ and target blocking rates 1% (lower curves) and 5% (upper curves) – with and without interference.

In the presence of an admission control based on a maximum number of active users $m = C(R)/c_{\min}$, the blocking rate is given by expression (12) for the cell load (18). Figure 15 gives the corresponding cell capacity for a path loss exponent $\alpha = 4$. Again, we observe that the impact of interference on cell capacity is significant.

7.2 Hexagonal networks

Now consider a hexagonal network. The interference suffered by a user u served by BS b is almost entirely generated by the 6 surrounding BS. A conservative approximation of the interference term \bar{I} is given by the following function of the distance r from BS b to user u :

$$\bar{I}(r) = \Gamma(2R - r) + 2\Gamma(\sqrt{(R - r)^2 + 3R^2}) + 2\Gamma(\sqrt{(R + r)^2 + 3R^2}) + \Gamma(2R + r).$$

This approximation is obtained assuming that user u is on a segment from BS b to a neighbor BS b' . We deduce the peak rate function:

$$C(r) = \min \left(C_0, \frac{W}{\delta} \times \frac{\Gamma(r)}{\bar{\eta} + \bar{I}(r)} \right).$$

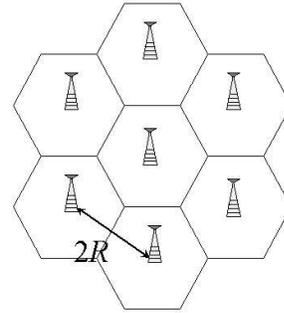


Figure 16: A hexagonal network.

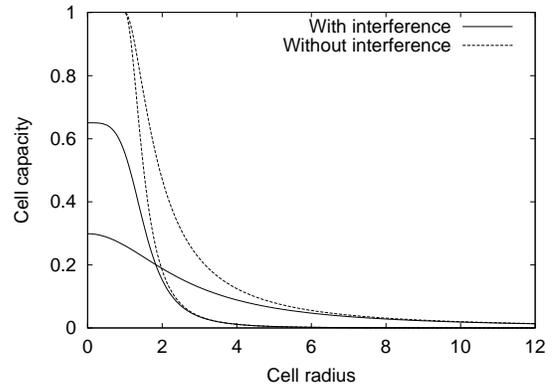


Figure 17: Cell capacity in a hexagonal network – defined as the maximum traffic intensity without saturation – for path loss exponent $\alpha = 4$ (lower curves) and $\alpha = 2$ (upper curves) – with and without interference.

For simplicity, we approximate the hexagonal cells by circular cells of radius R . Figure 17 compares the cell capacity – defined as the maximum traffic intensity without saturation – with and without interference. This is a decreasing function of the cell radius, with a maximum cell capacity equal to:

$$\bar{C}(0) \approx \begin{cases} 0.65 & \text{for } \alpha = 4, \\ 0.29 & \text{for } \alpha = 2. \end{cases}$$

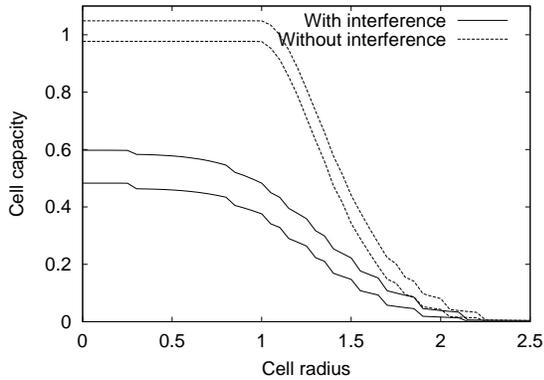


Figure 18: Cell capacity in a hexagonal network – defined as the maximum traffic intensity for a minimum data rate $c_{\min} = 0.02$ and target blocking rates 1% (lower curves) and 5% (upper curves) – with and without interference.

In the presence of an admission control based on the number of active users, the blocking rate is given by expression (12). Figure 18 gives the corresponding cell capacity for a path loss exponent $\alpha = 4$. Again, we observe that the impact of interference on cell capacity is significant.

8. CONCLUSION

We have derived analytical results relating user performance, in terms of blocking probability and data throughput, to cell size and traffic density, accounting for the random nature of traffic and the way the radio resource is shared. We have notably shown that the performance of the “fair power” sharing allocation (realized by a round-robin or a PF scheduler) cannot be significantly improved. Furthermore, this allocation has the practically interesting property that user performance can be explicitly evaluated, with or without admission control, and is insensitive to detailed traffic characteristics such as the flow size distribution.

We have observed that the impact of HDR feasible rate constraints on cell capacity is generally negligible. We conclude that the optimality principle of a TDMA-like strategy proved for a continuous set of available data rates remains valid in a more realistic situation where only a discrete set of data rates is available.

Another key observation is that, due to the relatively small number of active users in steady state, the impact of opportunistic schedulers that take advantage of fast fading is much more limited than one would expect from the analysis of a static scenario with a fixed user population. In particular, the cell capacity does not much differ from that obtained with a simple round-robin scheduler, especially for large cells.

Concerning admission control, we have studied the difference in terms of cell capacity between a scheme based on the number of active users and another scheme based on the minimum data rate. The former leads to a blocking rate which does not depend on user locations, unlike the latter. We have shown that provided the target blocking rate is not too high (less than 10 % typically), these two schemes are in fact equivalent: it is not necessary to base the admission decision on user locations.

Some issues need to be further explored. For instance, we assumed in Section 7 that BS are always active and observed an important decrease of the cell capacity due to interference. It would be useful to study how *idle periods* of BS could improve capacity. Other interesting issues include the impact of mobility and the integration of data services with voice and video services.

Acknowledgment. We are grateful to Sem Borst and Jim Roberts for fruitful discussions on the work presented in this paper.

9. REFERENCES

- [1] 3GPP TS 25.848, Physical layer aspects of UTRA High-Speed Downlink Packet Access, Release 4, 2001.
- [2] E. Altman, Capacity of Multi-service Cellular Networks with Transmission-Rate Control: A Queueing Analysis, in: *Proc. of ACM Mobicom*, 2002.
- [3] N. Bansal and M. Harchol-Balder, Analysis of SRPT Scheduling: Investigating Unfairness, in: *Proc. of ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*, 2001.
- [4] A. Bedekar, S. Borst, K. Ramanan, P. Whiting, E.M. Yeh, Downlink scheduling in CDMA data networks, in: *Proc. of IEEE Globecom*, 1999.
- [5] S. Ben Fredj, T. Bonald, A. Proutière, G. Régnié and J. Roberts, Statistical Bandwidth Sharing: A Study of congestion at flow level, in: *Proc. of ACM Sigcomm*, 2001.
- [6] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana and A. Viterbi, CDMA/HDR: A Bandwidth-Efficient High-Speed Wireless Data Service for Nomadic Users, *IEEE Communications Magazine*, 70–77, July 2000.
- [7] F. Berggren and R. Jntti, Asymptotically fair scheduling in fading channels, in: *Proc. of IEEE VTC Fall*, 2002.
- [8] S. Borst, User-level performance of channel-aware scheduling algorithms in wireless data networks, in: *Proc. of IEEE Infocom*, 2003.
- [9] G. Fayolle, I. Mitran and R. Iasnogorodski, Sharing a processor among many classes, *Journal of the ACM* (27) 519–532, 1980.
- [10] J.M. Holtzman, Asymptotic Analysis of Proportional Fair Algorithm, in: *Proc. of 12th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, Sept-Oct 2001.
- [11] J.M. Holtzman and S. Ramakrishna, A Scheme for throughput maximization in a dual class CDMA system, *IEEE Journal on Selected Areas of Communications* (40,2) 830–841, 1998.
- [12] A. Jalali, R. Padovani and B. Pankaj, Data throughput of CDMA-HDR a High Efficiency - High Data Rate Personal Communication Wireless System, in: *Proc. IEEE VTC Spring*, 2000.
- [13] N. Joshi, S.R. Kadaba, S. Patel and G.S. Sundaram, Downlink scheduling in CDMA Data Networks, in: *Proc. of ACM Mobicom*, 2000.
- [14] F. Kelly, *Reversibility and Stochastic Networks*, Wiley and Sons, 1979.

- [15] S-L. Kim, Z. Rosberg, J. Zander, Combined Power Control and Transmission Rate Selection In Cellular Networks, in: *Proc. of IEEE VTC Fall*, 1999.
- [16] L. Kleinrock, *Queueing systems*, Wiley and Sons, 1976.
- [17] R. Love, A. Gosh, R. Nikides, L. Jalloul, M. Cudak and B. Classon, High-Speed Downlink Packet Access Performance, in: *Proc. of IEEE VTC Spring*, 2001.
- [18] J.G. Pottie, System design choices in personal communications, *IEEE Personal Communications*, 50–67, 1995.
- [19] I.A. Rai, G. Urvoy-Keller and E. W. Biersack, Size-based scheduling with differentiated services to improve response time of highly varying flows, in: *Proc. of 15th ITC Specialist Seminar, Internet Traffic Engineering and Traffic Management*, Würzburg, 2002.
- [20] R.F. Serfozo, *Introduction to Stochastic Networks*, Springer Verlag, 1999.
- [21] B. Sklar, Rayleigh Fading Channels in Mobile Digital Communication Systems, Part I: Characterization, *IEEE Communications Magazine*, 90–100, July 1997.
- [22] V.V. Veeravalli, A. Sendonaris, The Coverage-Capacity Tradeoff in Cellular CDMA Systems, *IEEE Trans. on Vehicular Technology*, 1443–1451, 1999.
- [23] P. Viswanath, D. Tse and R. Laroia, Opportunistic Beamforming Using Dumb Antennas, *IEEE Trans. on Information Theory* (48) 1277–1294, June 2002.
- [24] A.J. Viterbi, *CDMA: Principles of Spread Spectrum Communication*, Addison-Wesley, 1995.
- [25] S-T Yang and A. Ephremides, Resolving the CDMA Cell breathing Effect and Near-Far Unfair Access Problem by Bandwidth-Space Partitioning, in: *Proc. of IEEE VTC Spring*, 2001.