Multiscale Fairness and its Application to Resource Allocation in Wireless Networks[1]

Eitan Altman[a,1,*], Konstantin Avrachenkov[a], Sreenath Ramanath[a]

[a]INRIA - Sophia Antipolis, 2004 Route des Lucioles, 06902 Sophia Antipolis Cedex, France

Abstract

Fair resource allocation is usually studied in a static context, in which a fixed amount of resources is to be shared. In dynamic resource allocation one usually tries to assign resources instantaneously so that the average share of each user is split fairly. The exact definition of the average share may depend on the application, as different applications may require averaging over different time periods or time scales. Our main contribution is to introduce new refined definitions of fairness that take into account the time over which one averages the performance measures. We examine how the constraints on the averaging durations impact the amount of resources that each user gets. We also address how the spatial component, which arises due to mobility of users, influences resource sharing under different fairness criteria. We demonstrate these new concepts via example applications.

Keywords: Resource allocation; Multiscale fairness; α-fairness; T-scale fairness; Networking

1. Introduction

Let us consider some set S of resource that we wish to distribute among I users by assigning user i a subset S_i of it. We shall be interested in allocating subsets of the resource fairly among the users. The set S may actually correspond to one or to several resources. We shall consider standard fairness criteria for sharing the resources among users. We shall see, however, that the definition of a resource will have a major impact on the fair assignment.

We associate with each user i a measurable function x_i that maps each point in S to some real number. Then, we associate with each i a utility u_i which maps all measurable subsets S_i to the set of real numbers. We shall say that S is a resource if $u_i(S_i)$ can be written for each $S_i \subset S$ as

$$u_i(S_i) = f\left(\int_{S_i} x_i(s)ds\right)$$

As an example, consider I mobiles that wish to connect to a base station between 9h00 and 9h10 using a common channel. The time interval is divided into discrete time slots whose number is N. Assume that the utility for each mobile i of receiving subsets $\mathcal{N}_i$ of slots

[1]This article is a journal version of the paper [14] presented in Networking 2011, Valencia, Spain

*Corresponding author, Tel: +33 (0) 4 90 84 35 82, Fax:+33 (0) 4 90 84 35 01

Email addresses: eitan.altman@sophia.inria.fr (Eitan Altman), konstantin.avrachenkov@sophia.inria.fr (Konstantin Avrachenkov), sreenath.ramanath@sophia.inria.fr (Sreenath Ramanath)

depend only on the number of slots N_i it receives. Then the set of N slots is considered to be a resource.

Next assume that if mobile i receives the channel at time slot t then it can transmit at a throughput of X_t^i. Assume that the utility of user i is a function of the total throughput it has during this fraction of an hour. Then again the N slots are considered as a resource.

We adopt the idea that fair allocation should not be defined in terms of the object that is split but in terms of the utility that corresponds to the assignments. This is in line with the axiomatic approach for defining the Nash bargaining solution for example. With this in mind, we may discover that the set of N slots cannot always be considered as a resource to be assigned fairly. Indeed, a real-time application may consider the N slots as a set of n resources, each containing $B = N/n$ consecutive slots. A resource may correspond to the number of time slots during a period of 100 ms. The utility of the application is defined as a function of the instantaneous rate, i.e. the number of slots it receives during each period of 100 ms. (With a playout buffer that can store 100 ms of voice packets, the utility of the mobile depends only on how many slots are assigned to it during 100 ms and not which slots are actually assigned to it.)

Related work: Our work is based on the α-fairness notion introduced in [3]. This, as well as other fairness notions can be defined through a set of axioms, see [13]. This paper is inspired by several papers which already observed or derived fairness at different time-scales [9, 11, 12, 10, 5, 1]. However, we would like to mention that the T-scale fairness (a unifying generalization of long- and short- term fairness) and multiscale fairness are new concepts introduced in the present work.

Structure of the paper: In Section 2, we introduce a resource sharing model which is particularly suitable for wireless applications. We also define several fairness criteria. In Section 3, we apply these new concepts to study spectrum allocation in fading channels. Fair resource sharing taking into account the spatial component which arises due to mobility of users is addressed in Section 4. Section 5 concludes the paper and provides avenues for future research.

2. Resource Sharing model and fairness definitions

Consider n mobiles located at points $x_1, x_2, ..., x_n$, respectively. We assume that the utility U_i of mobile i depends on its location x_i and on the amount of resources s_i it gets.

Let $\mathbf{S}$ be the set of assignments; an assignment $s \in \mathbf{S}$ is a function from the vector x to a point in the n-dimensional simplex. Its ith component, $s_i(x)$ is the fraction of resource assigned to mobile i.

Definition 1. *An assignment s is α-fair if it is a solution of*

$$\begin{aligned} Z(x,s,\alpha) &:= \max_s \sum_i Z_i(x_i,s_i,\alpha) \text{ such that,} \\ \sum_i s_i &= 1,\ s_i \geq 0 \quad \forall i = 1,...,n \qquad (1) \\ \text{where,} \quad Z_i(x_i,s_i,\alpha) &:= \frac{(U_i(x_i,s_i))^{1-\alpha}}{1-\alpha} \text{ for } \alpha \neq 1 \quad \text{and} \\ Z_i(x_i,s_i,\alpha) &:= \log\left(U_i(x_i,s_i)\right) \text{ for } \alpha = 1 \end{aligned}$$

We shall assume throughout that U_i is non-negative, strictly increasing and is concave in s_i. Then for any $\alpha > 0$, $Z_i(x_i, s_i, \alpha)$ is strictly concave in s_i. We conclude that $Z(x_i, s_i, \alpha)$ is strictly concave in s for any $\alpha > 0$ and therefore there is a unique solution $s^*(\alpha)$ to (2).

Definition 2. *[3] We call $Z_i(s_i, \cdot, \alpha)$ the fairness utility of mobile i under s_i, and we call $Z(s, \cdot, \alpha)$ the instantaneous degree of α-fairness under s.*

In applications, the state X will be random, so that the instantaneous amount of resource assigned by an α-fair allocation will also be a random variable. Thus, in addition to instantaneous fairness we shall be interested in the expected amount assigned by being fair at each instant.

Definition 3. *We call $E[Z(s, X, \alpha)]$ the expected instantaneous degree of α-fairness under s.*

In Section 2.1 we introduce the expected long-term fairness in which the expected amount of resource is assigned fairly.

Definition 4. *We say that a utility is linear in the resource if it has the form:*

$$U_i(x_i, s_i) := s_i q_i(x_i).$$

For example, consider transmission between a mobile source and a base station, and assume
(i) that the base station is in the origin ($x = 0$) but at a height of one unit, whereas all mobiles are on the ground and have height 0. Thus, the distance between the base station and a mobile located on the ground at point x is $\sqrt{1 + ||x||^2}$.
(ii) that the Shannon capacity can be used to describe the utility. If the resource that is shared is the frequency (denoted by C below) then the utility has the linear form:

$$U(C, x) \quad := \quad Cq(x); \text{ with } q(x) = \log\left(1 + \frac{P(x^2+1)^{-\beta/2}}{\sigma^2}\right)$$

2.1. Fairness over time: Instantaneous Versus Long-term α-fairness

Next we consider the case where $x_i(t)$, $i = 1, ..., n$, may change in time.

Definition 5. *We define an assignment to be instantaneous α-fair if at each time t each mobile is assigned a resource so as to be α-fair at that instant.*

Consider the instantaneous α-fair allocation and assume that time is discrete. We thus compute the instantaneous α-fair assignment over a period of T slots as the assignment that maximizes (for $\alpha \neq 1$)

$$\sum_{i=1}^{n} \frac{(U_i(x_i(t), s_i(t)))^{1-\alpha}}{1-\alpha} \text{ for every } t = 1, ..., T.$$

This is equivalent to maximizing

$$\sum_{t=1}^{T} \sum_{i=1}^{n} \frac{(U_i(x_i(t), s_i(t)))^{1-\alpha}}{1-\alpha}. \tag{2}$$

For $\alpha = 1$, we replace

$$\frac{(U_i(x_i(t), s_i(t)))^{1-\alpha}}{1-\alpha} \text{ by } \log[U_i(x_i(t), s_i(t))]$$

.

The optimization problem (2) corresponds to the α-fair assignment problem in which there are nT players instead of n players, where the utility of player $i = kn+j$ ($k = 0, ..., T-1, j = 1, ..., n$) is defined as

$$\mathcal{U}_i(x_i, s_i) = U_j(x_j(k+1), s_j(k+1)).$$

Definition 6. *Thus the expected instantaneous fairness criterion in the stationary and ergodic case regards assignments at different time slots of the same player as if it were a different player at each time slot!*

Note that when considering the proportional fair assignment, then the resulting assignment is the one that maximizes $\prod_{i=1}^n \prod_{t=1}^T U_i(x_i(t), s_i(t))$.

Definition 7. *Assume that the state process $X(t)$ is stationary ergodic. Let λ_i be the stationary probability measure of $X(0)$. The long-term α-fairness index of an assignment $s \in \mathbf{S}$ of a stationary process $X(t)$ is defined as*

$$\overline{Z}_\lambda(s) \;:=\; \sum_{i=1}^n \overline{Z}^i_\lambda(s); \;\; with \; Z^i_\lambda(s) = \frac{\Big(E_\lambda\left[U_i(X_i(0), s_i(X(0)))\right]\Big)^{1-\alpha}}{1-\alpha}.$$

An assignment s is long-term α-fair if it maximizes $Z_\lambda(s)$ over $s \in \mathbf{S}$.

As we see, instead of attempting to have a fair assignment of the resources at every t, it is the expected utility in the stationary regime that one assigns fairly according to the long-term fairness. Under stationarity and ergodicity conditions on the process $X(t)$ this amounts in an instantaneous assignment of the resources in a way that the time average amount allocated to the users are α-fair.

2.2. Fairness over time: T-scale α-fairness

Next we define fairness concepts that are in between the instantaneous and the expected fairness. They are related to fairness over a time interval T. Either continuous time is considered or discrete time where time is slotted and each slot is considered to be of one time unit. Below, we shall understand the integral to mean summation whenever time is discrete.

Definition 8. *The T-scale α-fairness index of $s \in \mathbf{S}$ is defined as*

$$Z_T(s) := \sum_{i=1}^n Z^i_T(s); \;\; with \; Z^i_T = \frac{\left[\frac{1}{T}\int_0^T U_i(X_i(t), s_i(X(t)))dt\right]^{1-\alpha}}{1-\alpha}.$$

The expected T-scale α-fairness index is its expectation. An assignment s is T-scale α-fair if it maximizes $Z_T(s)$ over $s \in \mathbf{S}$.

Definition 9. *The T-scale expected α-fairness index of $s \in \mathbf{S}$ is defined as*

$$Z_T(s) := \sum_{i=1}^n Z^i_T(s); \;\; with \; Z^i_T \;=\; \frac{\left[\frac{1}{T}\int_0^T E[U_i(X_i(t), s_i(X(t)))]dt\right]^{1-\alpha}}{1-\alpha}$$

2.3. Examples

We shall consider the following simple example of 2-scale fairness

Example 1. *Consider two time slots and two mobile stations. To whoever the first time slot will be allocated, that mobile would send or receive 25 units. At the second slot, a rate of 5 (resp. 10) units will be used if the slot is assigned to mobile 1 (resp. 2). We make the following observations. By [i,j] we shall denote the allocation that assigns slot 1 to mobile i and slot 2 to mobile j. The allocation [1,2] maximizes the global utility and moreover, the α-fair 2-scale utility for any α.*

Thus, we observe that the α-assignment is not monotone: The player with larger utilities received less at the α-fair utility, for all values of α!

Example 2. *(Example 1 continued) We now change a single utility in the last example: assume that if mobile 2 receives the first slot then it earns 10^2 units.*

(i) Now the global optimal solution is the assignment [2,2].

(ii) The proportional fair solution ($\alpha = 1$) is [2,1].

(iii) The maxmin fair assignment is [1,2].

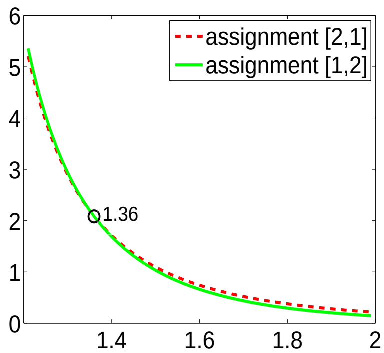

Figure 1: Performance index of [2,1] (dashed line) and [1,2] (solid line) assignments as a function of α (horizontal axis)

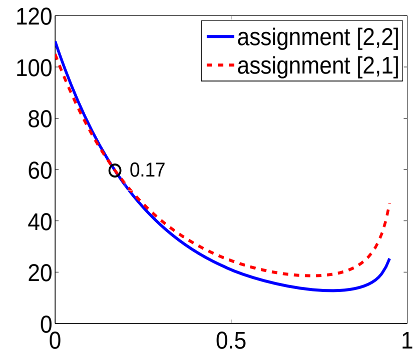

Figure 2: Performance index of [2,1] (dashed line) and [2,2] (solid line) assignments as a function of α (horizontal axis)

We depict in Figure 1 the performance index of the assignments [1,2] and [2,1]. We see that the max-min fair assignment [2,1] is 2-scale α-fair for all α larger than 1.36, whereas the assignment [1,2] is α-fair for $\alpha \in [1, 1.36]$.

For $\alpha < 1$ the two best assignments are [2,1] and [2,2]. The former is optimal over $\alpha \in [0.17, 1]$ and the latter over $\alpha \in [0, 0.17]$. This is seen from Figure 2.

Assume that the state processes is stationary ergodic. Then for any assignment $s \in \mathbf{S}$ we would have by the Strong Law of Large Numbers:

$$\lim_{T\to\infty} \frac{1}{T} \int_0^T U_i(X_i(t), s_i(X(t)))dt = E_\lambda \left[U_i(X_i(0), s_i(X(0)))\right].$$

Hence, for every i and s, we have,

$$\begin{aligned}\lim_{T\to\infty} Z^i_T(s) &= \lim_{T\to\infty} \frac{\left[\frac{1}{T}\int_0^T U_i(X_i(t), s_i(X(t)))dt\right]^{1-\alpha}}{1-\alpha} \\ &= \frac{\left(E_\lambda\left[U_i(X_i(0), s_i(X(0)))\right]\right)^{1-\alpha}}{1-\alpha} = \overline{Z}^i_\lambda(s).\end{aligned}$$

Assume that U_i is bounded. Then Z^i_T is bounded uniformly in T. The bounded convergence then implies that

$$\lim_{T\to\infty} E[Z^i_T(s)] = \overline{Z}^i_\lambda(s). \tag{3}$$

Theorem 1. *Assume that the convergence in (3) is uniform in s. Let $s^*(T)$ be the T-scale α-fair assignment and let s^* be the long-term α-fair assignment. Then the following holds:*

- $s^* = \lim_{T\to\infty} S^*(T)$
- *For any $\epsilon > 0$, s^* is an ϵ-optimal assignment for the T-scale criterion for all T large enough.*
- *For any $\epsilon > 0$, $s^*(T)$ is an ϵ-optimal assignment for the long-term fairness for all T large enough.*

Proof. According to [2], any accumulation point of $s^*(T)$ as $T \to \infty$ is an optimal solution to the problem of maximizing $\overline{Z}_T$ over S. Due to the strict concavity of $\overline{Z}_T$ in s it has a unique solution and it is coincides with any accumulation point of $s^*(T)$. This implies the first statement of the theorem. The other statements follow from Appendices A and B in [2].

2.4. Fairness over different time scales: Multiscale fairness

We consider real-time (RT) and non real-time (NRT) traffic. Resource allocation policy for RT traffic is *instantaneous-fair*, while for the NRT traffic, it is *expected-fair*. The available resources are divided amongst the RT and NRT traffic so as to guarantee a minimum quality of service (QoS) requirement for the RT traffic and to keep service time as short as possible for the NRT traffic.

The real-time traffic would like the allocation to be instantaneously α-fair. For $\alpha > 0$, this guarantees that at any time it receives a strictly positive allocation.

The non real-time traffic does not need to receive at each instant a positive amount of allocation. It may prefer the resources to be assigned according to the T-scale α-fair assignment where T may be of the order of the duration of the connection. Moreover, different non real-time applications may have different fairness requirements. For instance, bulk FTP transfer can prefer fairness over time scale longer than a time scale for some streaming application.

In order to be fair, we may assign part (say half) of the resource according to the instantaneous α-fairness and the rest of the resources according to the T-scale α-fairness. We thus combine fairness over different time scales.

We may now ask how to choose what part of the resource would be split according to the instantaneous assignment and what part according to the T-scale assignment. We propose to determine this part using the same α-fair criterion.

Specifically we define the multiscale fairness as follows:

Definition 10. *The multiscale α-fairness index of $s \in \mathbf{S}$ is defined as*

$$Z_{T_1,...,T_n}(s) := \sum_{i=1}^{n} Z^i_{T_i}(s); \ \textit{with } Z^i_{T_i} \quad = \quad \frac{\left[\frac{1}{T_i}\int_0^{T_i} U_i(X_i(t), s_i(X(t)))dt\right]^{1-\alpha}}{1-\alpha}$$

The expected multiscale α-fairness index is its expectation. An assignment s is multiscale α-fair if it maximizes $Z_{T_1,...,T_n}(s)$ over $s \in \mathbf{S}$. We also say that multiscale α-fair assignment is $(T_1, ..., T_n)$-scale fair assignment.

3. Application to spectrum allocation in fading channels

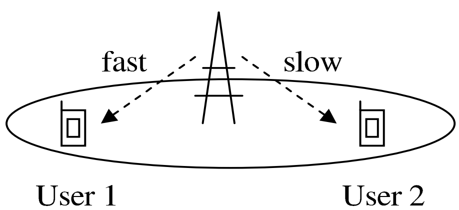

Figure 3: Spectrum allocation in random fading channels

We consider a fast-changing and a slowly-changing user (Figure 3), whose channels are modeled by the Gilbert model. The users can be either in a good or in a bad state. The dynamics of the users is described by a Markov chain $\{Y_i(t)\}_{t=0,1,...}$ i = 1,2, with the transition matrix and stationary distribution as:

$$P_i = \begin{bmatrix} 1-\epsilon_i\alpha_i & \alpha_i \\ \beta_i & 1-\epsilon_i\beta_i \end{bmatrix}; \quad \pi_i = \begin{bmatrix} \frac{\beta_i}{\alpha_i+\beta_i} & \frac{\alpha_i}{\alpha_i+\beta_i} \end{bmatrix}.$$

Let $\epsilon_1 = 1$ and $\epsilon_2 = \epsilon$. Note that the parameter ϵ does not have an effect on the stationary distribution, but, it influences for how long the slowly-changing user stays in some state. The smaller ϵ, the more seldom the user changes the states.

We assume that state 1 is a bad state and state 2 is a good state. Let h_{ij} represent the channel gain coefficient of user i in channel state j. The utility (achievable throughput via Shannon capacity) of user i in state j is given by

$$U_{ij} = s_{ij}\log_2(1 + \frac{|hij|^2 p_i}{\sigma^2})$$

where s_{ij} is the resource allocation and p_i is the power that corresponds to user i.

First, we would like to analyze T-scale fairness and to see the effect of the time scale on the resource allocation. Specifically, we consider the following optimization criterion

$$\sum_{i=1}^{2} \frac{1}{1-\alpha}\left[\frac{1}{T}\sum_{t=0}^{T} U_i(t)\right]^{1-\alpha} \quad \rightarrow \quad \max_{s_1,s_2} \tag{4}$$

with $U_i(t) = s_i(t)q_{i,Y_i(t)}$ and $s_1(t) + s_2(t) = 1$.

Let us consider several options for the time horizon T:

Instantaneous fairness. If we take $T = 1$ we obtain the instantaneous fairness. Namely, criterion (4) takes the form

$$\frac{1}{1-\alpha}\left[U_1^{1-\alpha}(0)+U_2^{1-\alpha}(0)\right] \quad \to \quad \max_{s_1,s_2}$$

The solution of the above optimization problem is given by

$$s_i(0)=\frac{q_{i,Y_i(0)}^{(1-\alpha)/\alpha}}{q_{1,Y_1(0)}^{(1-\alpha)/\alpha}+q_{2,Y_2(0)}^{(1-\alpha)/\alpha}}$$

This allocation results in the following expected throughputs

$$\theta_1=\sum_{i,j}\frac{q_{1,i}^{1/\alpha}}{q_{1,i}^{(1-\alpha)/\alpha}+q_{2,j}^{(1-\alpha)/\alpha}}\pi_{1,i}\pi_{2,j}, \theta_2=\sum_{i,j}\frac{q_{2,j}^{1/\alpha}}{q_{1,i}^{(1-\alpha)/\alpha}+q_{2,j}^{(1-\alpha)/\alpha}}\pi_{1,i}\pi_{2,j}. \tag{5}$$

Mid-term fairness. Let us take the time horizon as a function of the underlying dynamics time parameter ϵ, that is $T = T(\epsilon)$, satisfying the following conditions: (a) $T(\epsilon) \to \infty$ and (b) $T(\epsilon)\epsilon \to 0$. Condition (a) ensures that

$$\frac{1}{T(\epsilon)}\sum_{t=0}^{T(\epsilon)} 1\{Y_1(t)=i\} \to \pi_{1,i}, \quad \text{as} \quad \epsilon \to 0,$$

and condition (b) ensures that

$$\frac{1}{T(\epsilon)}\sum_{t=0}^{T(\epsilon)} 1\{Y_2(t)=i\} \to \delta_{Y_2(0),i}, \quad \text{as} \quad \epsilon \to 0.$$

The above results follow from the theory of Markov chains with multiple time scales (see e.g., [4]). It turns out to be convenient to take the following notation for the resource allocation: We denote by $s(t)$ the allocation for the fast-changing user and by $1 - s(t)$ the resource allocation for the slowly-changing user. Thus, we have $s_1(t) = s(t)$ and $s_2(t) = 1 - s(t)$. We denote by $\bar{s}_{i,j} = E[s(t)|Y_1(t) = i, Y_2(t) = j]$. We note that since the fast-changing user achieves stationarity when $T(\epsilon) \to \infty$ we are able to solve (4) in stationary strategies. Then, criterion (4) takes the form

$$\begin{aligned}\frac{1}{1-\alpha} \quad & \left[(\pi_{1,1}q_{1,1}\bar{s}_{1,Y_2(0)}+\pi_{1,2}q_{1,2}\bar{s}_{2,Y_2(0)})^{1-\alpha}\right. \\ & + \quad \left.((1-\pi_{1,1}\bar{s}_{1,Y_2(0)}-\pi_{1,2}\bar{s}_{2,Y_2(0)})q_{2,Y_2(0)})^{1-\alpha}\right] \quad \to \max_{\bar{s}_{1,Y_2(0)},\ \bar{s}_{2,Y_2(0)}}\end{aligned}$$

The above nonlinear optimization problem can be solved numerically. The expected throughputs in the mid-term fairness case are given by

$$\begin{aligned}\theta_1 &= (\pi_{1,1}q_{1,1}\bar{s}_{1,1}+\pi_{1,2}q_{1,2}\bar{s}_{2,1})\pi_{2,1}+(\pi_{1,1}q_{1,1}\bar{s}_{1,2}+\pi_{1,2}q_{1,2}\bar{s}_{2,2})\pi_{2,2},\\ \theta_2 &= (1-\pi_{1,1}\bar{s}_{1,1}-\pi_{1,2}\bar{s}_{2,1})q_{2,1}\pi_{2,1}+(1-\pi_{1,1}\bar{s}_{1,2}-\pi_{1,2}\bar{s}_{2,2})q_{2,2}\pi_{2,2}.\end{aligned} \tag{6}$$

Long-term fairness. In the case of long-term fairness we set $T = \infty$ which results in the following criterion

$$\frac{1}{1-\alpha}\left[E[U_1]^{1-\alpha} + E[U_2]^{1-\alpha}\right] \quad \rightarrow \quad \max_{s_1,s_2}$$

Due to stationarity, we can solve the above optimization problem over sequences in stationary strategies. Namely, we have the following optimization problem

$$\begin{aligned}
\frac{1}{1-\alpha} \quad & [\quad ((\pi_{1,1}\pi_{2,1}\bar{s}_{1,1} + \pi_{1,1}\pi_{2,2}\bar{s}_{1,2})q_{1,1} + (\pi_{1,2}\pi_{2,1}\bar{s}_{2,1} + \pi_{1,2}\pi_{2,2}\bar{s}_{2,2})q_{1,2})^{1-\alpha} \\
& + \quad ((\pi_{2,1} - \pi_{1,1}\pi_{2,1}\bar{s}_{1,1} - \pi_{1,2}\pi_{2,1}\bar{s}_{2,1})q_{2,1} \\
& + \quad (\pi_{2,2} - \pi_{1,1}\pi_{2,1}\bar{s}_{1,1} - \pi_{1,2}\pi_{2,2}\bar{s}_{2,2})q_{2,2})^{1-\alpha}] \quad \rightarrow \max_{\bar{s}_{1,1},\ \bar{s}_{1,2},\ \bar{s}_{2,1},\ \bar{s}_{2,2}}
\end{aligned}$$

The expected throughputs in the long-term fairness case are given by

$$\begin{aligned}
\theta_1 = & \quad (\quad \pi_{1,1}\pi_{2,1}\bar{s}_{1,1} + \pi_{1,1}\pi_{2,2}\bar{s}_{1,2})q_{1,1} + (\pi_{1,2}\pi_{2,1}\bar{s}_{2,1} + \pi_{1,2}\pi_{2,2}\bar{s}_{2,2})q_{1,2} \\
\theta_2 = & \quad (\quad \pi_{2,1} - \pi_{1,1}\pi_{2,1}\bar{s}_{1,1} - \pi_{1,2}\pi_{2,1}\bar{s}_{2,1})q_{2,1} \\
& \quad + \quad (\pi_{2,2} - \pi_{1,1}\pi_{2,1}\bar{s}_{1,1} - \pi_{1,2}\pi_{2,2}\bar{s}_{2,2})q_{2,2}
\end{aligned} \tag{7}$$

Let us also consider the expected instantaneous fairness which is given by criterion

$$\frac{1}{1-\alpha}\left[E[U_1^{1-\alpha}(t)] + E[U_2^{1-\alpha}(t)]\right] \quad \rightarrow \quad \max_{s_1,s_2}$$

which is equivalent to

$$\begin{aligned}
\frac{1}{1-\alpha} \quad & \left[\sum_{ij} \pi_{1,i}\pi_{2,j} \int_0^1 (sq_{1,i})^{1-\alpha} dF_{ij}(s)\right. \\
& \left. + \sum_{ij} \pi_{1,i}\pi_{2,j} \int_0^1 ((1-s)q_{2,j})^{1-\alpha} dF_{ij}(s)\right] \quad \rightarrow \quad \max_{F_{ij}}
\end{aligned}$$

where $F_{ij}(s)$ is the distribution for $s(t)$ conditioned on the event $\{Y_1(t) = i, Y_2(t) = j\}$. The above criterion is maximized by

$$F_{ij}(s) = \begin{cases} 0, & \text{if } s < q_{1,i}^{(1-\alpha)/\alpha}/(q_{1,i}^{(1-\alpha)/\alpha} + q_{2,j}^{(1-\alpha)/\alpha}), \\ 1, & \text{if } s \geq q_{1,i}^{(1-\alpha)/\alpha}/(q_{1,i}^{(1-\alpha)/\alpha} + q_{2,j}^{(1-\alpha)/\alpha}). \end{cases}$$

Thus, we can see that the expected instantaneous fairness criterion is equivalent to instantaneous fairness.

Multiscale fairness: Next, let us consider multiscale fairness over time. Specifically, (T_1, T_2)-scale fairness is defined by the following criterion

$$\frac{1}{1-\alpha}\left[\left(\frac{1}{T_1}\sum_{t=0}^{T_1} U_1(t)\right)^{1-\alpha} + \left(\frac{1}{T_2}\sum_{t=0}^{T_2} U_2(t)\right)^{1-\alpha}\right] \quad \rightarrow \quad \max_{s_1,s_2}$$

In this particular example, there are 6 possible combinations of different time scales. It turns out that in this example only the $(1,\infty)$-scale fairness gives a new resource allocation. The other combinations of time scales reduce to some T-scale fairness. Thus, let us first consider the multiscale fairness when we apply instantaneous fairness to the fast-changing user and long-term fairness to the slowly-changing user. The $(1,\infty)$-scale fairness corresponds to the following optimization criterion

$$\frac{1}{1-\alpha}\left[U_1(0)^{1-\alpha}+E[U_2(t)]^{1-\alpha}\right] \quad \to \quad \max_{s_1,s_2}$$

which is equivalent to

$$\begin{aligned}\frac{1}{1-\alpha} \quad & \left[(q_{1,Y_1(0)}(\bar{s}_{Y_1(0),1}\pi_{2,1}+\bar{s}_{Y_1(0),2}\pi_{2,2}))^{1-\alpha}\right. \\ & \left. +(q_{2,1}(1-\bar{s}_{Y_1(0),1})\pi_{2,1}+q_{22}(1-\bar{s}_{Y_1(0),2})\pi_{2,2})^{1-\alpha}\right] \to \max_{\bar{s}_{Y_1(0),1},\ \bar{s}_{Y_1(0),2}}\end{aligned}$$

The expected throughputs in the $(1,\infty)$-scale fairness case are given by

$$\begin{aligned}\theta_1 = & \ (\quad q_{1,1}(\bar{s}_{1,1}\pi_{1,1}\pi_{2,1}+\bar{s}_{1,2}\pi_{1,1}\pi_{2,2})+q_{1,2}(\bar{s}_{2,1}\pi_{1,2}\pi_{2,1}+\bar{s}_{2,2}\pi_{1,2}\pi_{2,2})), \\ \theta_2 = & \ (\quad q_{2,1}(1-\bar{s}_{1,1})\pi_{2,1}+q_{22}(1-\bar{s}_{1,2})\pi_{22})\pi_{1,1} \\ & + \quad (q_{2,1}(1-\bar{s}_{2,1})\pi_{2,1}+q_{22}(1-\bar{s}_{2,2})\pi_{22})\pi_{1,2}.\end{aligned}$$

As we have mentioned above, the other combinations of time scales reduce to some T-scale fairness. In particular, $(1,T(\epsilon))$-fairness reduces to the instantaneous fairness, $(T(\epsilon),\infty)$-fairness reduces to long-term fairness, and $(T(\epsilon),1)$-, $(\infty,1)$- and $(\infty,T(\epsilon))$-fairness all reduce to mid-term fairness.

Table 1: Case 1,2 & 3: Shannon capacity (q)/probability(π)

	Case-1		**Case-2**		**Case-3**	
	state-1 (bad)	state-2 (good)	state-1 (bad)	state-2 (good)	state-1 (bad)	state-2 (good)
User-1	2/0.2	8/0.8	3/0.1	9/0.9	3/0.9	9/0.1
User-2	2/0.2	8/0.8	1/0.3	7/0.7	1/0.3	7/0.7

Let us consider a numerical example. The parameters are given in Table 1. We consider three typical cases. The first case corresponds to the symmetric scenario. In the second case, the fast-changing user has in general better channel conditions. In the third scenario the slowly-changing user (user 2) is more often in the good channel state than the fast-changing user (user 1).

We plot the expected throughput of the mobiles for various fairness criteria for case-3 in Figure 4. Plots and explanation for case-1 and case-2 are provided in [18].

In the third scenario, the second user always gets better share in terms of throughput. This is expected as the second user spends on an average more time under better channel conditions (channel with a good state) and the long- or short- term throughput is the principal component of the optimization criteria. It is natural that long-term fairness gives the best efficiency for both types of users. However, we note that the $(1,\infty)$-scale fairness provides better control in terms of fairness. The $(1,\infty)$-scale fairness based allocation provides the second best efficiency after the long-term fairness based allocation. Thus, we conclude that multiscale fairness provides good sensitivity to the variation of the fairness parameter and

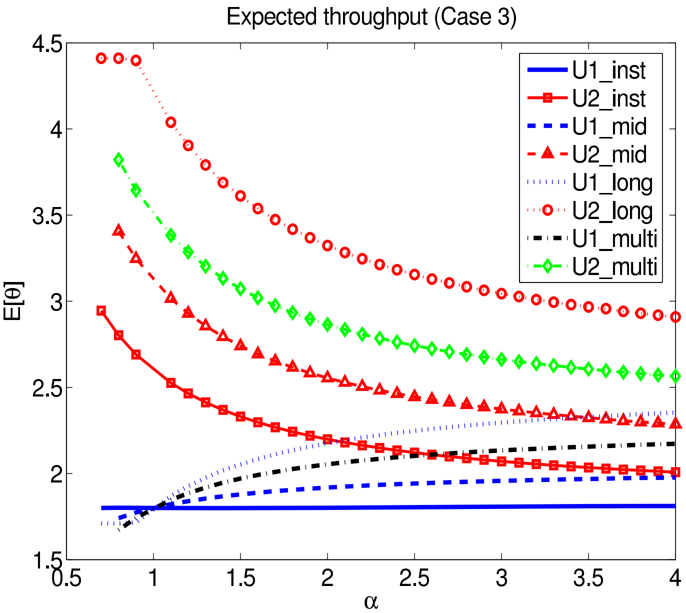

Figure 4: Throughput(θ) as a function of α for instantaneous, mid-term, long-term and (1,∞)-scale fairness criteria (Case 3).

at the same time good performance in expected throughput. Below we shall see that the multiscale fairness has another good property with respect to variance of the throughput. It is curious to observe that in this example the instantaneous fairness does not really help fast-changing user.

Coefficient of variation: We compute the coefficient of variation for short-term, mid-term, long-term and multiscale fairness. For this, we first compute the second moment of the throughput and then find the ratio of the standard deviation to its mean. For any user i, the coefficient of variation $\Gamma_i = \frac{\sqrt{\mathbf{E}[\theta_i^2]}}{\mathbf{E}[\theta_\mathbf{i}]}$.

In Figure 5, we plot the coefficient of variation in throughput for the various fairness criteria considered above. It is very interesting to observe that except the $(1,\infty)$-scale fairness criterion all the other fairness criteria behave similarly with respect to the coefficient of variation. Only in the case of $(1,\infty)$-scale fairness the coefficient of variation decreases for sort-term fairness oriented user. This is a very desirable property of the multiscale fairness as a short-term fairness oriented user is typically a user with a delay sensitive application. Similar to efficiency in throughput, multiscale fairness achieves efficiency in terms of overall variance as well.

4. A heterogeneous approach to resource allocation

In the previous sections, we discussed the notion of α-fairness to include time scale considerations for fair assignment of resources. Two extreme cases are elastic traffic (file transfer) and interactive voice. In this section, we consider time scale separation in the fairness that is related to the mobility of the users [2]. We solve the problem when the utilities are linear in the resources. We apply these results in the context of fair resource allocation in small cell networks in a dynamic setting and show how mobility and the constraints on the averaging durations impact the amount of resources each user gets.

[2]This work was presented in the proceedings of the workshop on Indoor Outdoor Femto Cells [15]

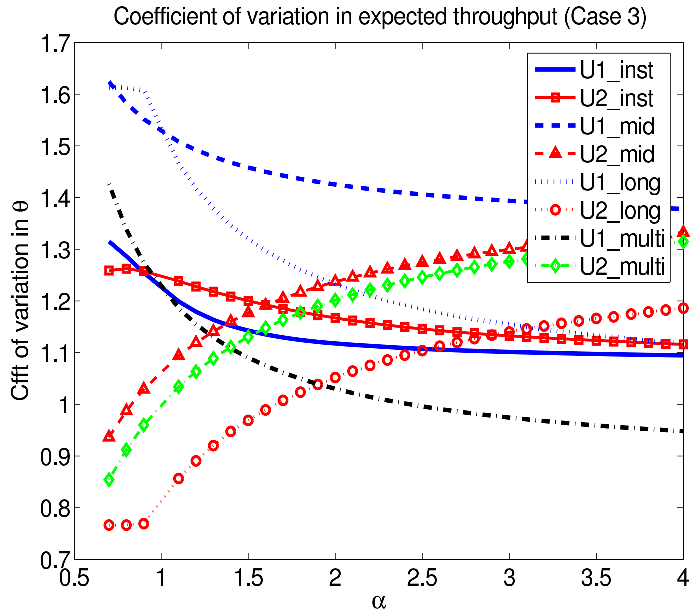

Figure 5: Coefficient of variation in expected throughput as a function of α for instantaneous, mid-term, long-term and (1,∞)-scale fairness criteria (Case 3).

Resource allocation algorithms often try to achieve fairness and efficiency. In particular, opportunistic scheduling algorithms in both uplink and downlink achieve higher throughputs by giving preference to mobiles with better relative channel conditions. Preference according to relative channel conditions rather than absolute channel conditions mean that the radio conditions of each mobile are normalized by the averaged conditions of that channel until then, and mobiles with best normalized radio conditions are selected for transmission; this guarantee fairness.

Averaging the radio conditions is done in practice using some low pass filter. Thus averaging is done over some effective period T. A longer T achieves a larger opportunistic gain at the cost of longer starvation periods. In other words, efficiency is obtained at the cost of being more unfair over a short time scale. Where as elastic traffic may prefer to be insensitive to short time scale unfairness, interactive real-time applications may need averaging over shorter time.

In the previous sections, we introduced the concept of T-scale and multiscale fairness [14]. We considered abstract time varying channels, time scales were related to the channel coherence times. In this section, we assume that user mobility is the source of channel variations and the timescale of fairness is determined by spatio-temporal behavior of the user. For example, such situations arise when resource allocation has to take into account that a user can be located at any random location over his region of mobility. The fairness concepts developed with this notion can be applied for example in small cell networks (e.g., pico or femto cells) [7, 8]. The scheduler in this case takes into account the expected utility of the user over his region of mobility to compute the resource allocation.

In subsequent paragraphs, we shall discuss *instantaneous fair share of resources* and how this can be applied for short-term fair share of resources in an example application (see sub section 4.1.1) the case of *long-term fair share of resources* follows this sub section (see sub section 4.1.2).

Consider the linear resource allocation model. Recall from definition 4 that a utility is linear in the resource if it has the form:

$$U_i(x_i, s_i) := s_i q_i(x_i).$$

Then, the α-fair utility can be written as:

$$Z_i(s, x, \alpha) := s_i^{1-\alpha} v_i(x), \qquad v_i := \frac{(q_i(x_i))^{1-\alpha}}{1-\alpha}$$

In the case of linear resources the instantaneous α-fairness has a nice explicit expression.

Theorem 2. *(i) The α-fair share is given by*

$$s_i^* = \frac{v_i(x_i)^{1/\alpha}}{\sum_j v_j(x_j)^{1/\alpha}} = \frac{q_i(x_i)^{1/\alpha-1}}{\sum_j q_j(x_j)^{1/\alpha-1}}$$

(ii) The utility for mobile i under the fair assignment is then

$$U_i(s^*, x) = \frac{v_i(x_i)^{1/\alpha}}{\sum_j v_j(x_j)^{1/\alpha}} q_i(x_i) = \frac{q_i(x_i)^{1/\alpha}}{\sum_j q_j(x_j)^{1/\alpha-1}}$$

(iii) The optimal value Z is given by

$$Z = \frac{1}{1-\alpha} \sum_i \left(\frac{q_i(x_i)^{1+1/\alpha}}{\sum_j q_j(x_j)^{1/\alpha}} \right)^{1-\alpha}$$

Proof. We relax the constraint and use KKT condition. s is optimal if and only if there is some $\lambda > 0$ such that s maximize L^λ s.t. $s_i \geq 0$ for all i, where

$$L^\lambda = \sum_i s_i^{1-\alpha} v_i + \lambda(1 - \sum_i s_i)$$

Equating the derivative w.r.t. s_i to zero gives

$$\begin{aligned} s_i^{-\alpha} v_i(x_i) &= \frac{\lambda}{1-\alpha} \\ \text{so that } s_i &= \left(\frac{1-\alpha}{\lambda} v_i(x_i) \right)^{1/\alpha} \end{aligned}$$

Since the sum of s_i is 1, we conclude that

$$\frac{\lambda}{1-\alpha} = \left(\sum_j v_j(x_j)^{1/\alpha} \right)^{\alpha}$$

Substituting in the previous equation yields (i), from which (ii), (iii) follows. $\diamond$

Example 3. *Consider as an example a path loss $\beta = 2$ and let the base station be located one unit above the mobiles. We assume that $q_i(x)$ is proportional to the attenuation between the mobile and the base station: $q_i(x) = c_i q(x)$ where $q(x) = (1+x^2)^{-\beta/2}$. For $\beta = \alpha = 2$ we have*

$$s_i^*(x) = \frac{c_i^{-1/2}(1+x_i^2)^{1/2}}{\sum_j c_j^{-1/2}(1+x_j^2)^{1/2}}.$$

Furthermore,

$$U_i(s^*, x) = s_i^*(x) q_i(x_i) = \frac{c_i^{1/2}(1+x_i^2)^{-1/2}}{\sum_j c_j^{-1/2}(1+x_j^2)^{1/2}}.$$

Let us consider a simple example of allocation amongst indoor and outdoor users.

Example 4. *Say, we have in total N time slots. The allocation happens in a bundle of 6 slots, such that, either we allocate all of it to an outdoor user located at x_1 or fair share them amongst 3 indoor users located at (y_1, y_2, y_3) with $y_i \in (0, L)$, with any user getting two consecutive slots. Now the question is "Given that we fair-share among the indoor users, how do we fair share between the outdoor and the indoor users?".*

In this example, we assume any user gets a throughput $q \in [0,1]$. Let $\{U_1, U_2\}$, represent the utility of user 1 and sum utility of users $2-4$ and let $\{T_1, 1-T_1\}$ represent their respective assignment of resources. Now, utility of user 1,

$$U_1(T_1) = 6T_1 q_1(x_1).$$

Let $\bar{s} = \{s_1, s_2, 1-s_1-s_2\}$ represent the assignment of resources for the indoor users. Then, utility of users $2-4$ is

$$\begin{aligned} u_2(T_1, \bar{s}) &= 6s_1(1-T_1)q_2(y_2), \\ u_3(T_1, \bar{s}) &= 6s_2(1-T_1)q_3(y_3) \text{ and} \\ u_4(T_1, \bar{s}) &= 6(1-s_1-s_2)(1-T_1)q_4(y_4). \end{aligned}$$

Now the α-fair share $\bar{s}^* = \{s_1^*, s_2^*, 1-(s_1^* - s_2^*)\}$ is given by,

$$\bar{s}^* = \arg\max_{\bar{s}} \sum_{i=2}^{4} \frac{E[u_i(T_1, \bar{s})]^{1-\alpha_1}}{1-\alpha_1}$$

The sum utility of users $2-4$ is,

$$U_2(T_1) = \sum_{i=2}^{4} 6\bar{s}_i{}^*(1-T_1)q_i(y_i)$$

The α-fair share between the outdoor and indoor users is,

$$T_1^* = \arg\max_{T_1} \frac{E\left[U_1(T_1)\right]^{1-\alpha} + E\left[U_2(1-T_1)\right]^{1-\alpha}}{1-\alpha}$$

4.1. Application to resource allocation in small cell networks

Small cell networks, typically constitute pico and femto cells, which are small portable base stations, typically designed for use in home, small business, commercial centers, hot spots, etc. These base stations connect to the service providers network via DSL or cable broadband and allow service providers to improve capacity and extend service coverage to address the growing mobile broadband needs. These deployments address wide cellular markets that include GSM, WiMAX and LTE technologies. Typically the service range extends to few tens to couple of hundreds of meters covering indoor and outdoor partitions typical of residential homes and offices. For the state of art and current research trends in small cell networks, see [16, 17, 7, 8, 19, 6] and the references therein.

In our example application, we consider an indoor-outdoor partition as shown in Figure 6. Let Ω be the line segment $[-L, L]$, with the indoor partition spanning $[0, L]$. Assume that the base station is located just to the left of the wall. Mobile 1 is at some point $x \leq 0$ outdoor and user 2 remains always indoor and is located at some Y_t which is uniformly distributed over $[0, L]$.

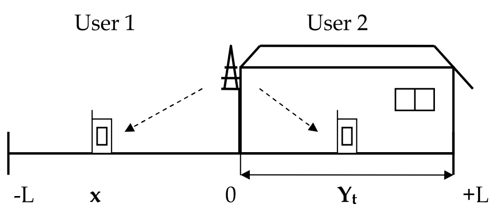

Figure 6: Indoor-outdoor scenario: User 1 located at x, user 2 located indoor at Y_t

We let $q_i(x) = c_i q(x)$ with $c_1 = 1$ and c_2 is equal to some large fixed number. Thus the presence of the wall between the base station and mobile 2 is modeled by a multiplicative attenuation by some constant c_2. Assume that the mobility pattern of mobile 2 is uniform over the indoor part $[0, L]$.

We consider allocation of the fraction of time between the two mobiles. We apply the results obtained in Section-III to obtain the short-term fair resource allocation.

4.1.1. Short-term fair share of resources

We compute the expected utility for each user when assigning the channel so as to achieve instantaneous fairness. The expected utility for mobile 1 under the instantaneous optimal fairness $\mathbf{s}$ is given by

$$\begin{aligned}
U_1(s^*, x) &= s^*(x) q_1(x), \text{ where } s^*(x) := \frac{a}{a+b} \\
\text{for } \beta = 2,\ a &:= c_1^{-1/\alpha} \log\left(1 + \frac{1}{x^2}\right)^{-1/\alpha} \\
b &:= c_2^{-1/\alpha} \left[\log(1 + \frac{1}{L^2}) + \frac{2}{L}\tan^{-1}(L)\right]^{-1/\alpha}
\end{aligned}$$

Note that mobile 2 has a mobility pattern which is uniform over the indoor part [0,L] and hence its utility is given by $\frac{1}{L}\int_0^L \log(1 + \frac{1}{x^2})dx = \log(1 + \frac{1}{L^2}) + \frac{2}{L}\tan^{-1}(L)$, which is the second term in the denominator.

Similarly, the utility of mobile 2 is given by

$$U_2(s^*, x) = [1 - s^*(x)] \quad c_2 q_2(x).$$

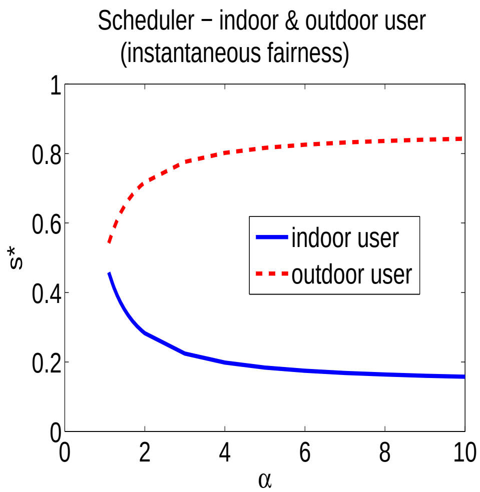

Figure 7: Scheduler s^* for the indoor and outdoor user with instantaneous fairness as a function of α for $\alpha > 1$. Wall attenuation 6 dB, path-loss $\beta = 3$, position of outdoor user $x = -3$.

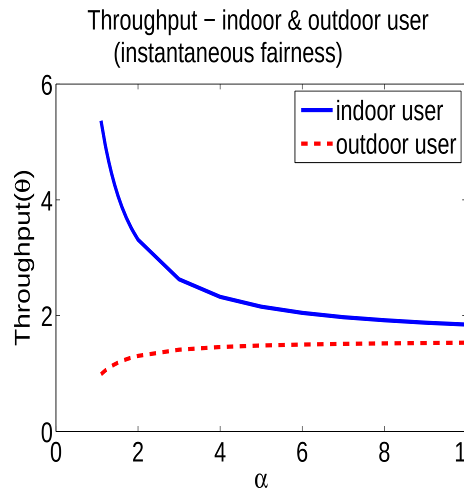

Figure 8: Throughput θ for the indoor and outdoor user with instantaneous fairness as a function of α for $\alpha > 1$. Wall attenuation 6 dB, path-loss $\beta = 3$, position of outdoor user $x = -3$.

Numerical example: In Figure 7 and Figure 8, we plot the scheduler and the instantaneous throughput for the indoor and outdoor user, as a function of α. We fix the location of the outdoor user at $x = -3$, path-loss $\beta = 3$. We set $L = 3$ for this example. The indoor user is located at some point which is uniformly distributed over $[0, L]$.

Observation: When the fairness index α is small, we observe that the instantaneous throughput achieved is higher as the outdoor user is located at the boundary $(-L)$. But, as the fairness index increases, the throughput of the indoor and the outdoor user starts to converge. Notice that the outer user is scheduled more often as α increases, which results in an increase in the outdoor users throughput. Thus, the instantaneous fair scheduler that takes into account the region of mobility for the indoor user favors the outdoor user as the degree of fairness, α, increases.

Discussion: In the above analysis if the resource allocation had not taken into account user 2's (indoor user) mobility, the throughput of the indoor user would have been less for smaller α and would tend to increase as the degree of fairness α increases. Also, one would have noticed that the indoor user is scheduled more often with an increase in α contrary to the observation that we saw in the numerical example. Thus, for small values of α, considering the region of mobility for the indoor user favors the outdoor user.

4.1.2. Long-term fair share of resources

Next we consider the long-term fairness. The long-term allocation $s \in \mathbf{S}$ (which is a function of x and Y_t) is given by maximizing

$$Z(s) := \frac{1}{1-\alpha}\left[\left(\frac{1}{L}\int_0^L dy s_1(x,y)q(x)\right)^{1-\alpha} + \left(\frac{1}{L}\int_0^L dy c_2 s_2(x,y)q(y)\right)^{1-\alpha}\right]$$

Theorem 3. *The long-term α-fair policy is given by $s_2(x_2) = 1$ for $x_2 \leq l(\alpha)$ and is otherwise zero, where $l(\alpha)$ is the solution of the fixed-point equation*

$$l(\alpha) = \left(1 - \frac{1}{\beta - 1}\left(c_2 \frac{q(l(\alpha))}{q(x)}\right)^{1-\frac{1}{\alpha}}\right)^{-1} L$$

where $q(x)$ is a monotone decreasing function of the form $x^{-\beta}$

Proof. It is easy to see that α-fair policy has to have the form mentioned in the theorem statement. If not, for example say there exists an optimal policy which allocates mobile 2 in two disjoint intervals. Then, one can construct a better policy by shifting the right most interval to the end of the left interval and this contradicts the optimality. Thus the optimization simplifies to one-dimensional optimization

$$\max_s Z(s) = \max_{l \in [0,L]} Z(s^l).$$

$$\text{where, } s_2^l(x) = 1 \text{ for } \{x \leq l\}.$$

It is easy to see that

$$Z(s^l) = \frac{1}{1-\alpha}\left[\left(\frac{L-l}{L}\, q(x)\right)^{1-\alpha} + \left(\frac{1}{L}\int_0^l c_2\, q(y)dy\right)^{1-\alpha}\right].$$

The optimal $l(\alpha)$ is obtained by differentiating the above equation w.r.t l and equating to zero, which results in the fixed-point equation

$$\left((L - l(\alpha))q(x)\right)^{-\alpha} q(x) - \left(c_2 \int_0^{l(\alpha)} q(y)dy\right)^{-\alpha} q(l(\alpha)) = 0.$$

Specifically when $q(x) = x^{-\beta}$ then the fixed-point equation simplifies to

$$l(\alpha) = \left(1 - \frac{1}{\beta - 1}\left(c_2 \frac{q(l(\alpha))}{q(x)}\right)^{1-\frac{1}{\alpha}}\right)^{-1} L$$

◇

Numerical example: We plot in Figure 9 a numerical example to observe how $l(\alpha)$ varies with α for $\alpha > 1$. In this example, we consider path loss $\beta = 2$, location of the outdoor user $x = -2$ and wall attenuation of 6 dB.

Observation: We observe that as α increases, the value of $l(\alpha)$ monotonically decreases and starts to saturate. Also, it is interesting to note that the indoor user is scheduled when its *mobility* and the *fairness* of resource allocation, $(l(\alpha), \alpha)$, lie within the dashed region below the curve. Thus, with a long-term fair scheduler, as the mobility region increases, the range of fairness applicable decreases.

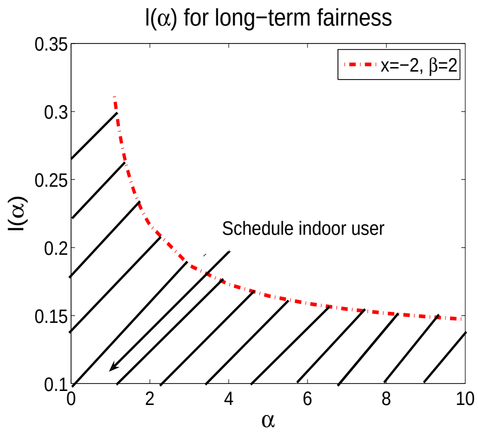

Figure 9: $l(\alpha)$ for long-term fairness as a function of α ($\alpha > 1$) and wall attenuation of 6 dB ,path-loss $\beta = 2$, position of outdoor user $x = -2$.

5. Conclusions and future research

We have introduced T-scale fairness and multiscale fairness. The notion of T-scale fairness allows one to address in a flexible manner requirements of emerging applications (like You Tube) which demand quality of service requirement between strict real-time traffic and best-effort traffic. The notion of multiscale fairness allows one to use a single optimization criterion for resource allocation when different applications are present in the network. We compared the new fairness notions with the previously known criteria of instantaneous and long-term fairness criteria and illustrated them with an example application of spectrum allocation when users with different dynamics are present in the system. We demonstrated that the multiscale fairness provides a versatile framework for resource allocation.

Next, we investigated how the spatial component, which arises due to mobility of users, influences resource sharing under different fairness criteria. By considering that utilities are linear in resources, we derived explicit expressions for the short-term and long-term fair resource allocation. We applied this in the context of fair resource allocation in indoor-outdoor small cells in a dynamic setting. In the case of short-term fairness, we observed that the scheduler starts to schedule the outdoor user more often as the degree of fairness increases, which is contrary to the case if the region of mobility had not been considered. When considering long-term, we observed how the resource allocation to the mobile user depends on a combination of its region of mobility and fairness of resource allocation. As the region of mobility increases, the degree of fairness applicable decreases. We envisage that these concepts can be used to derive some very interesting insights into resource allocation. Especially, one can for example consider the base station to be located inside and study how the resource allocation changes. Also, the indoor user can be static and the outdoor user can be mobile. Further one can include certain mobility models for users that are mobile and study fair resource allocation.

Also, it would be interesting to investigate how multiscale fairness criterion allows to allocate resources when a number of applications with different QoS requirements are present in the network. It is also interesting to investigate T-scale fairness in the non-stationary regime.

Acknowledgement

This work was done in the framework of the INRIA and Alcatel-Lucent Bell Labs Joint Research Lab on Self Organized Networks and the ANR Ecoscells project.

[1] F.P. Kelly, A.K. Maulloo and D.K.H. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability", Journal of the Operational Research Society, Vol. 49, No. 3, pp. 237-252, 1998.

[2] M. Tidball, A. Lombardi, O. Pourtallier and E. Altman, "Continuity of Optimal Values and Solutions for Control of Markov Chains with Constraints", SIAM Journal of Control and Optimization, Vol. 38, No.4, pp. 1204-1222, 2000.

[3] J. Mo and J.Walrand, "Fair End-to-end Window-based Congestion Control", IEEE/ACM Transactions on Networking, Vol. 8, No. 5, pp. 556-567, 2000.

[4] J. Filar, H.A. Krieger and Z. Syed, "Cesaro limits of analytically perturbed stochastic matrices", Linear Algebra Applications, Vol. 353, pp. 227-243, 2002.

[5] V. Ramaiyan, A. Kumar and E. Altman, "Fixed Point Analysis of Single Cell IEEE 802.11e WLANs: Uniqueness and Multistability", IEEE/ACM Transactions on Networking, Vol. 16, No. 5, pp. 1080-1093, 2008.

[6] "Femtocell Networks", Special issue in EURASIP Journal on Wireless Communication and Networking. 2010.

[7] H. Claussen, "The future of small cell networks", IEEE COMMSOC MMTC E-Letter, Sep 2010.

[8] J. Hoydis, M. Kobayashi and M. Debbah, "Green small-cell networks", IEEE Vehicular Technology Magazine, Mar 2011.

[9] E. Altman, K. Avrachenkov and B. J. Prabhu, "Fairness in MIMD Congestion Control Algorithms", In Proceedings of IEEE INFOCOM, 2005, Miami, USA.

[10] M. Bredel and M. Fidler, "Understanding Fairness and its Impact on Quality of Service in IEEE 802.11", Arxive, 2008.

[11] E. Altman, K. Avrachenkov and A. Garnaev, "Generalized Alpha-Fair Resource Allocation in Wireless Networks", In Proceedings of 47th IEEE Conference on Decision and Control, 2008, Cancun, Mexico.

[12] E. Altman, K. Avrachenkov and A. Garnaev, "Alpha-Fair Resource Allocation under Incomplete Information and Presence of a Jammer", In Proceedings of NETCOOP 2009 (LNCS 5894, pp. 219-233, 2009).

[13] T. Lan, D. Kao, M. Chiang and A. Sabharwal, "An Axiomatic Theory of Fairness for Resource Allocation", In Proceedings of IEEE INFOCOM, 2010, San Diego, USA.

[14] E. Altman, K. Avrachenkov, S. Ramanath, "Multiscale Fairness and its Application to Resource Allocation in Wireless Networks", In Proceedings of Networking 2011, May 9-13, Valencia, Spain (LNCS 6641, pp. 225-237, 2011).

[15] S. Ramanath, E. Altman, K. Avrachenkov, "A heterogeneous approach to fair resource allocation and its application in femtocell networks", In Proceedings of the workshop on Indoor and Outdoor Femto Cells (IOFC), WiOpt 2011, May 13, Princeton USA.

[16] Gregory Davi, "Using picocells to build high-throughput 802.11 networks", RF Design, Jul 2004.

[17] "Picocell Mesh: Bringing Low-Cost Coverage, Capacity and Symmetry to Mobile WiMAX", A Tropos Networks white paper.

[18] E. Altman, K. Avrachenkov, S. Ramanath, "Multiscale Fairness and its Application to Dynamic Resource Allocation in Wireless Networks", INRIA research report number RR-7783. Available at http://hal.inria.fr/inria-00515430/en/

[19] "Femtocell", Wikipedia link: http://en.wikipedia.org/wiki/Femtocell.