Project: IEEE 802.16 Broadband Wireless Access Working Group http://ieee802.org/16

Title: Utility Fairness and the Proportional Fairness (PF) Criterion

Source:

Pantelis Monogioudis	monogiou@alcatel-lucent.com	+1 (973) 386-4804
Sudhir Ramakrishna	rsudhir@alcatel-lucent.com	+1 (973) 428-7726
Nandu Gopalakrishnan	nandu@alcatel-lucent.com	+1 (973) 884-6477

Alcatel-Lucent
67 Whippany Road,
Whippany, NJ, USA

Abstract:
If all users have identical utility function with respect to rate, fair allocations can be ensured by the current criterion. However, if multiple utility functions are present in the network, the current fairness criterion for best-effort traffic imposes significant limitations. The contribution suggests a new fairness criterion that is based on first principles optimal rate control theory.

Purpose: To modify the 802.16m simulation methodology document.

1 Introduction

The current criterion places a hard limit in the relative performance between users that belong to different utility function groups. This is a major problem for radio interfaces that use such groups to manage resources effectively. Such groups can be present in many ways: (a) by terminals having different peak transmission rates (or in general different capability sets), (b) by different multiple access options (especially in the uplink), (c) by geometry-based resource allocation, etc.

Let us first review the proportional fair criterion. It states that $a\%$ of the users should have time-average throughput T at least as high as $a\%$ of the user-averaged throughput $\bar{T}$. This is represented as shown in Table 1.

$\tilde{T}$	**CDF**
0.1	0.1
0.2	0.2
0.5	0.5

Table 1:

where $\tilde{T}$ is defined as the average normalized throughput of the k-th user.

$$\tilde{T}(k) = T(k)/\bar{T} \tag{1}$$

and $\bar{T} = \sum_{k=1}^{K} T(k)$ is the user-averaged throughput in a system with K users.

So, if an allocation to a user increases the average throughput then similar increases must be done to the other lower-ranked users to not violate the criterion. Looking it in another way it penalizes allocations that can be given to higher-ranked users for the benefit of lower-ranked users.

In microeconomics, a Pareto-optimal is any feasible allocation that results when it is not possible to increase the resource allocated to some users without decreasing the resources allocated to other users. Under this definition, the current fairness criterion may not lead to Pareto-optimal allocations in that it does not allow to increase resources to users that can translate the allocated resources to a larger utility than those who cannot.

2 Utility Fairness

In broad terms we associate the degree of satisfaction from a resource allocation with a utility function U that is a function of the achievable spectral efficiency or achievable rate x.

It is well known by the seminar work of Kelly, that proportional fairness is achieved by assuming sources that exhibit a uniform utility function of the form,

$$U(x_k) = \log(x_k) \tag{2}$$

where x_k is the transmission rate of the k-th user. It is then reasonable to suggest an objective based on utility that can take the form of,

$$\max \sum_k U_k(x_k) \tag{3}$$

In other words, maximize the sum of the utilities of the users. In the limit there may be K different utility forms, but in reality such utilities are far less than K. The objective is generic enough to accommodate different traffic types (elastic or inelastic) in addition to the other groupings mentioned above.

The application of utility fairness can be done as follows:

Let N_{sect} be the number of sectors in the system and N_s the cardinality of the set M that contains the k indices of the mobile stations that have non-zero throughput $x_k(s)$ during the communication with the serving sector s.

1. Gather into a vector **r** the average throughput of all MS in the system at the end of the simulation runs.

2. Compute the metric

$$U(\mathbf{r}) = N_{sect} U^{-1} \left(\frac{1}{N_s} \sum_{k \in M} U(x_k) \right) \tag{4}$$

where U^{-1} is the inverse utility function and the term in parenthesis effectively produces the per user throughput weighted implicitly by the utility function.

For $U(x_k) = \log(x_k)$ this results in a proportional weighting of the rate allocation without the need to check against an explicit criterion like before. Optimizing against this metric automatically results in maximizing the sector throughput subject to implicit fairness criterion determined by the utility function.