Self-organizing Dynamic Fractional Frequency Reuse on the Uplink of OFDMA Systems

Balaji Rengarajan
IMDEA Networks
Madrid, Spain
Email: balaji.rengarajan@imdea.org

Alexander L. Stolyar
Bell Labs, Alcatel-Lucent
Murray Hill, NJ 07974
Email: stolyar@research.bell-labs.com

Harish Viswanathan
Bell Labs, Alcatel-Lucent
Murray Hill, NJ 07974
Email: harish.viswanathan@alcatel-lucent.com

Abstract—Reverse link (or uplink) performance of cellular systems is becoming increasingly important with the emergence of new uplink-bandwidth intensive applications such as Video Share [14], where end users upload video clips captured through their mobile devices. In particular, it is important to design the system to provide good user throughput in most of the coverage area, including at the cell edge. Soft fractional frequency reuse (FFR) is one of the techniques for mitigating inter-cell interference in cellular systems, leading to overall spectral efficiency enhancements and/or cell edge throughput improvements. We propose a novel algorithm that dynamically creates efficient soft FFR patterns on the uplink of orthogonal frequency division multiple access (OFDMA) based cellular systems; this allows the system to "automatically" adapt to user traffic distribution and system layout.

Our algorithm is based on systematically ascending towards a local maximum of the system-wide sum of user *utilities*, which depend on user throughputs. We show that this can be done in a semi-autonomous fashion: each sector does its resource allocation independently, with only an infrequent periodic exchange of *interference costs* between neighboring sectors. The proposed algorithm, called Multi-sector Gradient for Uplink (MGR-UL), allocates in-sector resources (power, frequency, time-slots to each user) in a way that simultaneously takes into account both the benefit to its "own" users' utility and the cost of creating interference to neighboring sectors; along with that each sector estimates the cost of interference to itself. Extensive simulation results show that significant performance benefits (up to 69% in total throughput in some typical scenarios) can be achieved with respect to a baseline approach. Simulations also show the automatic formation of soft FFR patterns.

I. INTRODUCTION

Fourth generation cellular systems, through their high data rates and system capacity, are expected to enable the true mobile broadband experience. Many new applications will emerge as a result of affordable broadband to mobile devices. In particular, we are likely to see the emergence of applications that are uplink-bandwidth intensive. An example of such an application is Video Share [14], where mobile users share video clips captured through their mobile devices while on the go. Thus it is imperative to improve uplink data throughput and spectral efficiency in addition to that of the downlink.

Cellular spectral efficiency is limited by, among other things, out-of-cell interference, with cell-edge user throughputs suffering the most. Mitigating out-of-cell interference is thus a promising approach to improving throughput performance, especially at cell edge. Traditionally, frequency reuse

has been used to mitigate out-of-cell interference. However, traditional, "integer" frequency reuse (with each cell/sector confined to a fixed part of available frequency spectrum) suffers from inefficient utilization of bandwidth. More recently, soft fraction frequency reuse (FFR) schemes have been proposed for orthogonal frequency division multiple access (OFDMA) systems, where each cell uses all available frequencies, but an intelligent choice of user transmit powers and frequency assignment is employed so as to minimize loss from interference across neighboring cells. In addition, it is desirable for an FFR scheme to be *dynamic* (or, *adaptive*), because an efficient FFR "pattern" is highly dependent on user spatial distribution and traffic demand (as well as on the system layout), which change with time. Finally, it is important for a dynamic FFR technique to be distributed (or, autonomous) "as much as possible", namely that different cells perform their resource allocation independently, with no or infrequent exchange of signaling information between neighbor cells. Our prior work towards the design of distributed dynamic FFR schemes [10], [11] focused on the downlink, and in this paper we propose an uplink algorithm.

Our proposed algorithm, called Multi-sector Gradient for Uplink (MGR-UL), systematically pursues local maximization of the system-wide sum of user *utilities*, which are functions of user throughputs. We show that this can be done in a *semi-autonomous* fashion: each sector does its resource allocation independently, with only an infrequent periodic exchange of *interference costs* between neighboring sectors. In-sector resource allocation procedure includes power, time-slot, and frequency assignment to the users; this procedure utilizes interference costs (to neighbor sectors) as parameters; it also includes continuous estimation of the sector's "own" interference costs. Sectors send each other infrequent interference cost updates – this exchange is limited to neighbor sectors only. Thus, only a limited amount of information is exchanged between neighbor sectors, making practical implementation of the approach feasible.

Besides our earlier work on dynamic FFR for OFDMA downlink [10], [11], prior work on different aspects of resource allocation in OFDMA systems includes [1], [3], [4], [5], [7] – we refer reader to [10], [11] for a brief review. The concept of FFR for best effort traffic in the context of OFDMA systems has appeared in cellular network standardization technical

contributions [12], [13] and in [6].

The paper is organized as follows. In Section II we describe the basic system model under consideration and the problem that is being addressed. In Section III we consider an idealized model of system behavior; for such model we formally derive the (asymptotically) optimal scheme, pursuing the fastest system utility ascend. Section IV, first, outlines issues that need to be addressed by a resource allocation scheme in a real system, and then describes the MGR-UL algorithm designed for practical implementation. The simulation studies, comparing the performances of MGR-UL and a baseline algorithm in a realistic setting are given in Section V. We conclude with a summary and discussion of future work in Section VI.

II. BASIC MODEL AND PROBLEM STATEMENT

We consider the uplink (from mobile user to base station) of a multi-sector OFDMA system, where each user i is assigned to one of the sectors m; $\mathcal{I}$ and $\mathcal{M}$ are the finite sets of users and sectors, respectively. (Notation $i \in m$ will mean "user i is within sector m".) Frequency band is divided into J equal size subbands, indexed by j; $\mathcal{J} = \{1, \ldots, J\}$ is the set of subbands. Each subband consists of C resource blocks, each of bandwidth W. N_0 is spectral noise density.

The system operates in discrete time, over time-slots $t = 0, 1, \ldots$. In each time-slot each sector schedules a subset of "its" users to transmit, and for each scheduled user it assigns: the subbands, the number of resource blocks in each subband, and the power-per-resource-block (and per-subband) to use. The total power assigned to any user in any time-slot cannot exceed P^*.

A key feature of an OFDMA system is that transmissions of different users *within a sector* do not interfere with each other (because they use different frequencies), while transmissions of users within different sectors do interfere with each other *if they happen to use same frequencies*. This in particular means that two transmissions within same sector, or using different subbands in different sectors never interfere with each other. Two transmissions using the same subband in different sectors interfere with each other whenever the resource blocks used for the transmissions are overlapping in frequency. If different resource blocks within the same subband are used, then the transmissions in the different sectors do not interfere with each other.

In this paper, we consider best-effort user traffic. Each user i has an associated (concave, smooth, strictly increasing) utility function $U_i(X_i)$ of its average (over time) achieved rate X_i. ($U_i(X_i) = \log(X_i)$ for a proportional fair objective.) The goal is to find a scheduling strategy so as to maximize the total utility of the system, $\sum_i U_i(X_i)$. Note that here "scheduling" strategy is understood broadly, because it includes not only the choice of users to transmit, but also their subband, resource block and power assignments.

III. AN IDEALIZED MODEL OF SYSTEM-WIDE UTILITY MAXIMIZATION

To motivate our proposed scheme (given in Section IV), we first consider a further idealized model, for which we can formally derive a scheme that performs, essentially, a gradient ascend towards a local maximum of the system-wide utility. (More details of the derivation are given in [8]; we only sketch it here to save space.) The key additional idealization is, roughly speaking, to assume that a transmission in a time slot experiences the interference that is a time-averaged interference (over a certain window), as opposed to actual instantaneous one. This assumption is well suited for modeling a system where sectors do not (or cannot) have good estimates of actual *instantaneous* interference levels, and thus have to employ an autonomous or semi-autonomous control; this is the situation we are interested in in this paper. We will also assume that interference in a subband is spread uniformly over all resource blocks within the subband. Finally, we assume that the propagation gain G_i^m (path loss and shadowing) from user i to (base station of) sector m is constant.

A. Definition of system utility

Suppose that in each time slot, each sector m has finite set S_m of possible choices (decisions) of how to allocate resource blocks and transmit powers to its users $i \in m$. Associated with each choice $s \in S_m$, are the number of resource blocks $h_{ij}(s)$ and power-per-resource-block $p_{ij}(s)$, allocated to each user $i \in m$ in each subband j; for each s, the total power allocated to each user i cannot exceed maximum available mobile power, $\sum_j h_{ij}(s)p_{ij}(s) \leq P^*$, and the number of allocated resource blocks within each subband j cannot exceed C: $\sum_i h_{ij}(s) \leq C$.

Suppose each sector m does scheduling in a way such that decision $s \in S_m$ is picked in the fraction ϕ_s^m of all time-slots, $\sum_s \phi_s^m = 1$. Denote $\phi = \{\phi_s^m, \ s \in S_m, m \in \mathcal{M}\}$. Given ϕ, the system utility U is defined as

$$U = \sum_m \sum_{i \in m} U_i(x_i),$$

where each U_i is strictly increasing continuously differentiable function (e.g. $\log$) of user i average achieved rate

$$x_i = \sum_{s \in S_m, m \ni i} \phi_s^m \sum_j h_{ij}(s)r_{ij}(s), \qquad (1)$$

where user i rate per-resource-block in subband j is

$$r_{ij}(s) = W \log_2 \left(1 + f_{ij}(s)\right), \qquad (2)$$

$$f_{ij}(s) = \frac{G_i^m p_{ij}(s)}{N_0 W + I_j^m/C},$$

and, finally, the average interference in subband j of sector m is

$$I_j^m = \sum_{k \neq m} \sum_{s \in S_k} \phi_s^k \sum_{i \in k} G_i^m h_{ij}(s)p_{ij}(s). \qquad (3)$$

Clearly, the utility U is a continuously differentiable function of ϕ. It will be convenient, however, to view it as

$U(\phi) = V(\phi, I(\phi))$, where $V(\phi, I)$ is a function of ϕ and $I = \{I_j^m, m \in \mathcal{M}, j \in \mathcal{J}\}$, defined by $V(\phi, I) = \sum_m \sum_{i \in m} \bar{U}_i(x_i)$, (1) and (2); and $I(\phi)$ is a vector-function of ϕ, defined by (3).

We get the following expressions for partial derivatives:

$$\frac{\partial V}{\partial \phi_s^m} = \sum_{i \in m} U_i'(x_i) \sum_j h_{ij}(s) r_{ij}(s), \qquad (4)$$

$$a_j^m \doteq -\frac{\partial V}{\partial I_j^m} = \qquad (5)$$

$$\sum_{i \in m} U_i'(x_i) \sum_{s \in S_m} \phi_s^m h_{ij}(s) \frac{W[f_{ij}(s)]^2}{C(\ln 2)(1 + f_{ij}(s)) G_i^m p_{ij}(s)},$$

$$\frac{\partial I_j^m}{\partial \phi_s^k} = \sum_{i \in k} G_i^m h_{ij}(s) p_{ij}(s), \quad k \neq m, \qquad (6)$$

$\partial I_j^m / \partial \phi_s^k = 0, \quad k = m.$

B. Dynamics of system utility

Consider a dynamic system, which evolves in discrete time $t = 0, 1, 2, \ldots$. At time $t + 1$ each sector m can choose one allocation $s_m(t + 1) \in S_m$; when such choice is made, the system "state" changes from $\phi(t)$ to $\phi(t + 1)$, where $\phi(t)$ is the set of (exponentially) averaged frequencies of choosing different decisions in the past up to time t. (See [8] for details.) Then, to maximize the increment $U(\phi(t + 1)) - U(\phi(t))$ of utility from slot t to slot $t+1$, each sector m can *independently* choose allocation

$$s_m(t + 1) \in \arg\max_{s \in S_m} \left(\sum_{i \in m} U_i'(x_i)(\sum_j h_{ij}(s) r_{ij}(s)) \right) - $$
$$\left(\sum_{k \neq m} \sum_j a_j^k \sum_{i \in m} G_i^k h_{ij}(s) p_{ij}(s) \right). \qquad (7)$$

The first term is interpreted as the utility gain in sector m due to allocation s, and the second term is the utility loss in other sectors due to interference caused by allocation s. We emphasize that, as long as the current interference costs a_j^k and gains G_i^k are "known" to sector m, the optimization is "local" to the sector. Thus, if the system dynamics is governed by scheduling rule (7), the state $\phi(t)$ is driven towards a local maximum of function $U(\phi)$.

IV. RESOURCE ALLOCATION SCHEME FOR A REAL SYSTEM

Motivated by the results in Section III, for the idealized model, our approach to system wide utility maximization in a real system will consist of two parts:

1. Each sector m continuously estimates and updates its current interference costs a_j^m, i.e. the sensitivity of its utility to the interference level in each subband j. These interference costs a_j^m are periodically communicated to other sectors (those neighbor sectors k that create non-negligible interference to m), which use them for making scheduling decisions (see item 2).

2. Each sector m makes scheduling decisions that maximizes its own sector utility subject to costs associated with interference caused to neighboring sectors. The main difference relative to traditional scheduling for utility maximization is the use of interference costs.

Thus, the approach attempts is to constantly "drive" the system towards an operating point where a local maximum of utility is attained.

However, any scheme for a real system has to take into account the following features/constraints.

(a) Interference experienced by a transmission in a given time slot and given resource block is not the time-average interference, but really is instantaneous, depending (among other things) on instantaneous user/resource-block/power scheduling assignment in neighboring sectors.

(b) Transmission of a data packet may not be completed within one time slot, due to insufficient instantaneous signal to interference ratio. Several HARQ retransmissions may be then required.

(c) To allow for HARQ feedback, time slots are divided into several "interlaces", say 5. This means that if a packet is transmitted in slot t and requires retransmission(s), they are scheduled in slots $t + 5, t + 10$, and so on. This in particular means that not all resource blocks may be available for scheduling new transmissions in a given time slot – some may be already taken by retransmissions.

(d) Retransmissions of a packet may not occur in the same physical resource block (same frequency) as original transmission, because a frequency hopping may be employed.

The details of the algorithm are given below. For a real system, notation G_i^ℓ will be used for the *average* propagation gain from user i to sector ℓ. When we talk about resource allocation in a given sector, it is usually denoted by m.

A. Interference costs

Non-negative cost a_j^m is the cost to the utility of sector m of unit interference increase in subband j. It is a dynamic quantity computed and maintained by each sector m as described below. It is changing slowly with time, and each sector sends periodic (infrequent) updates of its costs a_j^m to all other sectors ℓ. (In reality it sends it only to the neighbor sectors, those that cause sufficiently high interference to m.)

B. Scheduling objective

In each time slot, for each of its users $i \in m$ in each subband j, sector m needs to choose the number of resource blocks H_{ij} to be assigned and the transmit power P_{ij} per resource block, such that

$$\sum_j H_{ij} P_{ij} \leq P^* \quad \text{and} \quad \sum_j H_{ij} \leq C, \quad \forall i.$$

The objective is to maximize the quantity

$$\left(\sum_{i \in m} U_i'(X_i)(\sum_j H_{ij} R_{ij}) \right) - \left(\sum_{k \neq m} \sum_j a_j^k \sum_{i \in m} G_i^k H_{ij} P_{ij} \right),$$
$$(8)$$

where X_i is the current average rate of user i (updated as described below), R_{ij} is the rate per-resource-block

$$R_{ij} = W \log_2 \left(1 + \frac{G_i^m P_{ij}}{N_0 W + I_j^m / C} \right)$$
$$= \frac{W}{\ln 2} \ln \left(1 + F_{ij} \right), \qquad (9)$$

I_j^m is the average interference power to sector m in the entire subband j, and we denoted the user i SINR by

$$F_{ij} \doteq \frac{G_i^m P_{ij}}{N_0 W + I_j^m / C}.$$

C. Scheduling algorithm: Multi-cell Gradient for Uplink (MGR-UL)

The computational complexity of maximizing (8) is typically prohibitively high. Therefore, we propose a heuristic algorithm, consisting of the following steps, performed in each time-slot.

Step 1. We compute the transmission power P_{ij} per resource block, which is to be used *if user i is actually scheduled in subband j*. This power P_{ij} is determined so that it maximizes the utility value minus cost,

$$U_i'(X_i) R_{ij}(P_{ij}) - \sum_{\ell \neq m} a_j^\ell G_i^\ell P_{ij},$$

subject to being within $[0, P^*]$ and not causing more than the interference $\bar{P}$ per resource block to any other sector $\ell \neq m$. This gives

$$P_{ij} = \max \left(\min \left(Y, P^*, \min_{\ell \neq m} \frac{\bar{P}}{G_i^\ell} \right), 0 \right)$$

where

$$Y = \frac{U_i'(X_i) W}{(\ln 2) \sum_{\ell \neq m} b_j^\ell G_i^\ell} - \frac{1}{\hat{F}_{ij}}$$

and

$$\hat{F}_{ij} = \frac{G_i^m}{N_0 W + I_j^m / C}.$$

Estimates of the average gains G_i^ℓ to neighbors are obtained from downlink pilot strength measurements made at the mobile and fed back to the base station. (Such measurements are available, because they are required to enable handoffs.) The average propagation gain on the downlink will be the same as that on the uplink because of channel reciprocity for path loss and shadow fading. Fast fading, which is the only component of the gain that is different on the two links, is averaged out. The estimate of average "gain-to-interference-plus-noise" ratio $\hat{F}_{ij}$ is obtained by averaging (see Section IV-D) its instantaneous values $\hat{\Gamma}_{ij}$. (In turn, $\hat{\Gamma}_{ij}$ is obtained by measuring received pilot SINR for user i in subband j. The specific way we use in simulations is described in [8].)

Step 2. For the purposes of scheduling, the transmission rates are assumed to be

$$\hat{R}_{ij}(P_{ij}) = W \log_2 \left(1 + \hat{\Gamma}_{ij} P_{ij} \right).$$

Step 3. The packet size per-resource-block is chosen, based on $\hat{R}_{ij}(P_{ij})$; we chose it to be just equal to $\hat{R}_{ij}(P_{ij})$.

Step 4 (ACTUAL SCHEDULING). We pick a resource block at random among all still available resource blocks across all subbands. Suppose it happens to be in subband j. In this resource block we schedule user i for which the value of

$$Z_{ij} = U_i'(X_i) \hat{R}_{ij}(P_{ij}) - \sum_{\ell \neq m} a_j^\ell G_i^\ell P_{ij},$$

is maximal among those users i which still have at least κP_{ij} "leftover" power to transmit, where $\kappa \in [0, 1]$ is a parameter; we chose $\kappa = 0.9$. (The leftover power for the user is updated after it is scheduled in a resource block.) Then we pick another (unused yet) resource block at random, and so on, until we make allocation (perhaps, null) in each resource block.

Step 5 (OPTIONAL). For each allocated resource block and the corresponding user, we check if the scheduling objective can be improved by reallocating the user to a still available resource block in a different subband; if so, we make the "best" such reallocation. This step is performed for each of the resource blocks, allocated in Step 1. (In our simulations this step is moot, because the setting is such that all resource blocks are used in Step 4.)

D. $\hat{F}_{ij}$ updates

$\hat{F}_{ij}$ is updated once per scheduling interval for all (ij):

$$\hat{F}_{ij} := (1 - \beta_0) \hat{F}_{ij} + \beta_0 \hat{\Gamma}_{ij},$$

where $\beta_0 > 0$ is a (small) averaging parameter.

E. Average rate X_i updates

X_i is updated once per scheduling interval for all users i:

$$X_i := \beta \overline{R}_i + (1 - \beta) X_i, \qquad (10)$$

where $\beta > 0$ is a (small) averaging parameter and $\overline{R}_i$ is the total size of all packets of user i whose transmission was *successfully completed* in this slot. If user i was not transmitting or did not complete any transmission, then $\overline{R}_i = 0$.

F. Calculation of costs a_j^m

As we allocate resource blocks during scheduling, and along with that, we update costs a_j^m as follows. If a resource block is assigned to user i (whether it is newly assigned, or user i HARQ transmission continues from the previous slot), we do

$$a_j^m := \beta_1 \frac{U_i'(X_i) W [F_{ij}]^2}{(\ln 2)(1 + F_{ij}) G_i^m P_{ij}} + (1 - \beta_1) a_j^m, \qquad (11)$$

where $F_{ij} = \hat{F}_{ij} P_{ij}$;
for an "unused" resource block (in subband j), we do

$$a_j^m := (1 - \beta_1) a_j^m. \qquad (12)$$

V. SIMULATIONS

A. System model for simulations

We consider a hexagonal grid of 19 base stations each with three sectors. The sector antennas are assumed to be oriented in a clover-leaf pattern so that the adjacent cell sectors are not facing each other directly. A wrap-around model for interference where the hexagonal arrangement is replicated by translation to create the same number of interfering cells around every one of the 19 cells is adopted.

Standard propagation parameters are used to determine the received signal power level for a given transmit power level. Parameter values can be found in [8]. For these parameters, with the site-to-site distance set at 856m, the cell edge SNR (signal to thermal noise ratio, when there is no interference from surrounding cells, assuming total available power is distributed uniformly over the entire bandwidth) turns out to be 22.4 dB. A small cell size, typical of urban morphology, has been chosen since gains are generally larger with smaller cells because of higher interference levels.

When fast fading is used in the simulations, the model is representative of frequency-selective Rayleigh fading with temporal characteristics captured through Jakes fading model with vehicle speed of 20 Km/hr and carrier frequency of 2 Ghz. The frequency-selectivity is modeled by simulating independent fading across sets of coherence bands. In our simulations we consider 6 sub-bands that are divided into three sets of coherence bands each with two sub-bands. For additional details on the simulation assumptions we refer the reader to [8].

B. Baseline algorithm

As a baseline against which we evaluate the performance of MGR-UL, we consider the following algorithm. (It will be referred to as 'Baseline'.) All other aspects of the system model are the same as in the previous sub-section.

Transmission power setting. Each sector m maintains a variable SINR target $\hat{F}_m$. The (potential, if scheduled) power P_{ij} per resource block for user i within *any* subband j (all subbands are "treated" the same way) is calculated to achieve SINR $\hat{F}_m$, or is set to the maximum power P^* if $\hat{F}_m$ is not achievable.

Scheduling. After the powers per resource block P_{ij} are calculated, scheduling within a time slot is done in exactly same way as for MGR-UL, except without taking into account the interference costs; namely, with $Z_{ij} = U_i'(X_i)\hat{R}_{ij}(P_{ij})$. (The average rate and pilot SINR updates are done the same way as well.)

SINR target $\hat{F}_m$ updates. Each sector k measures the average (over time) Interference-over-Thermal-noise level I_{oT}^k. If I_{oT}^k exceeds some maximum level (parameter) I_{oT}^* in a time slot, sector k sends an "overload message" (one bit) to is neighboring sectors m indicating such event. In every time slot, sector m increases its $\hat{F}_m$ by $\delta_1 > 0$, unless it receives an overload message from at least one of its neighbors – then it decreases $\hat{F}_m$ by $\delta_2 > 0$.
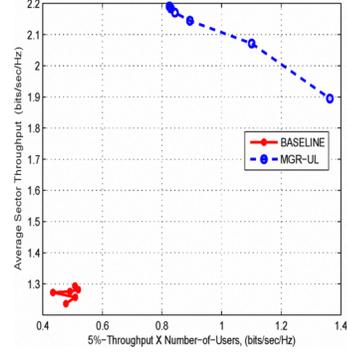

Fig. 1. Average sector throughput Vs. Normalized 5-% edge throughput. No fast fading, 6 subbands
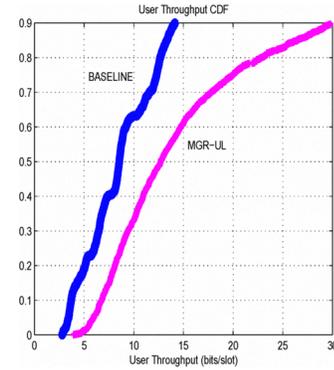

Fig. 2. Cumulative distribution function of user throughputs. No fast fading, 6 subbands

C. Results and discussion

Figure 1 shows the performance of the MGR-UL algorithm relative to the Baseline algorithm for the case of no fast fading. Different points on the curves correspond to different values of minimum rate parameter R_{min}, which range from 0 to 10 bits/slot. (The algorithm with $R_{min} > 0$ attempts to provide the average rate of at least R_{min} to each user; it is done by a dynamic adjustment of the marginal utilities $U_i'(X_i)$, as described in [8].) From the figure, we observe that the MGR-UL algorithm provides substantial improvement in both the average sector throughput and the 5-percentile throughput. The sector throughput improvement, corresponding to $R_{\min} = 0$, is 69% while the highest edge throughput improvement is 176% with $R_{\min} = 10$.

Notice that for the Baseline scheme changing R_{min} parameter has little effect on the system performance. In particular, increasing R_{min} practically does not increase cell edge user throughput. This is because Baseline, which by definition tries to equalize *all* user SINRs within a sector, has a tendency towards equalizing user throughputs as well; as a result, attempts to increase edge user throughputs within the framework of Baseline are ineffective. On the other hand, the interference-cost based power setting in MGR-UL allows "good" users (those close to serving base station) to achieve much higher
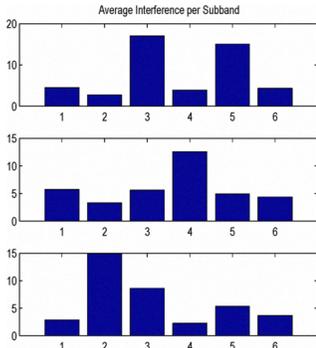
Fig. 3. Average interference per subband in sectors 1-3. No fast fading, 6 subbands

SINRs. This not only improves overall spectral efficiency, but also allows to "free up" more or less time-frequency resources for the cell edge users, when necessary, thus allowing a trade-off between total and cell edge throughputs.

Figure 2 shows the cumulative distribution function of the user throughputs for the MGR-UL and the Baseline algorithms for the six sub-band case. From the curves it is clear that the MGR-UL algorithm improves throughput performance for all users and is not simply trading off throughput of "good" users with that of cell edge users.

Figure 3 shows the average interference across the six sub-bands for the three co-located sectors. It is clear that the MGR-UL algorithm indeed results in a soft reuse pattern where some sub-bands are preferred in certain sectors.

While we have chosen to illustrate the basic results for the case of no fast fading, simulations with fast fading show that MGR-UL can achieve either the average throughput improvement of 36% (with a small cell edge throughput improvement as well), or the edge throughput improvement of 38% (with still a 20% average throughput increase).

VI. CONCLUSIONS AND FUTURE WORK

We proposed a resource allocation scheme for the uplink of OFDMA systems, whose underlying objective is to continuously maximize the overall system utility, which in turn depends on the achieved user throughputs. "In the process of doing that", the scheme dynamically creates an efficient FFR pattern, appropriate for each given spatial user traffic distribution. The scheme is semi-autonomous, with the inter-cell communication limited to infrequent exchange of interference costs. Our simulations show that the scheme achieves substantial performance improvements over a baseline scheme.

The idea of using interference costs (to neighbors) can potentially be used for the purposes other than dynamic FFR. For example, using constant, a priori fixed, interference costs amounts to a form of *static* FFR; in this case *no* inter-cell communication is required. (This may be an attractive option, if inter-cell exchange is undesirable or infeasible.) It might be of interest to compare the performance of such static scheme to that of other static schemes proposed in the literature.

REFERENCES

[1] E. Altman, K. Avrachenkov, and A. Garnaev, "Closed form solutions for water-filling problem in optimization and game frameworks," in *Proceeding of INFOCOM'2008*, Phoenix, April 14-18, 2008.

[2] M. Andrews, L. Qian, A. L. Stolyar, "Optimal Utility Based Multi-User Throughput Allocation subject to Throughput Constraints," in *Proceeding of INFOCOM'2005*, Miami, March 13-17, 2005.

[3] T. Bonald, S. C. Borst, and A. Proutiere, "Inter-cell scheduling in wireless data networks," in *Proceedings of European Wireless Conference*, 2005.

[4] S. T. Chung, S. J. Kim, J. Lee, and J.M. Cioffi, "A game theoretic approach to power allocation in frequency-selective Gaussian interference channels," in *Proceedings of the IEEE International Symposium on Information Theory*, pp 316-316, July 2003.

[5] S. Das, H. Viswanathan, and G. Rittenhouse, "Dynamic load balancing through coordinated scheduling in packet data systems," in *Proceedings of INFOCOM*, 2003.

[6] S. Das and H. Viswanathan, "Interference mitigation through intelligent scheduling," in *Proceedings of the Asilomar Conference on Signals and Systems*, Asilomar, CA, November 2006.

[7] A. Gjendemsjo, D. Gesbert, G. E. Oien, and S. G. Kiani, "Optimal power allocation and scheduling for two-cell capacity maximization," in *Proceedings of the IEEE RAWNET (WiOpt)*, April 2006.

[8] B. Rengarajan, A. L. Stolyar, H. Viswanathan, "A Semi-autonomous Algorithm for Self-organizing Dynamic Fractional Frequency Reuse on the Uplink of OFDMA Systems," *Bell Labs, Alcatel-Lucent, Technical Memo*, December, 2009. http://ect.bell-labs.com/who/stolyar/publications/new_rl_dffr.pdf

[9] A.L. Stolyar, "On the Asymptotic Optimality of the Gradient Scheduling Algorithm for Multi-User Throughput Allocation," *Operations Research*, 2005, Vol. 53, No.1, pp. 12-25.

[10] A. L. Stolyar, H. Viswanathan, "Self-organizing Dynamic Fractional Frequency Reuse in OFDMA Systems," in *Proceedings of INFOCOM'2008*, Phoenix, April 14-18, 2008.

[11] A. L. Stolyar, H. Viswanathan, "Self-organizing Dynamic Fractional Frequency Reuse for Best-Effort Traffic Through Distributed Inter-cell Coordination," in *Proceedings of INFOCOM'2009*, Rio-de-Janeiro, April 19-25, 2009.

[12] Third Generation Partnership Project 2, "Ultra Mobile Broadband Technical Specifications," http://www.3gpp2.org, March 2007

[13] Third Generation Partnership Project, Radio Access Network Work Group 1 Contributions, http://www.3gpp.org, September 2005

[14] Video Share from AT&T, http://www.wireless.att.com/learn/messaging-internet/media-entertainment/attvideoshare.jsp