

Proportional Fair Frequency-Domain Packet Scheduling for 3GPP LTE Uplink

Suk-Bok Lee* Ioannis Pefkianakis* Adam Meyerson* Shugong Xu† Songwu Lu*

*Computer Science Department
UCLA, CA 90095

†Sharp Laboratories of America
Camas, WA 98607

Abstract—With the power consumption issue of mobile handset taken into account, Single-carrier FDMA (SC-FDMA) has been selected for 3GPP Long-Term Evolution (LTE) uplink multiple access scheme. Like in OFDMA downlink, it enables multiple users to be served simultaneously in uplink as well. However, its single carrier property requires that all the subcarriers allocated to a single user must be *contiguous* in frequency within each time slot. This contiguous allocation constraint limits the scheduling flexibility, and frequency-domain packet scheduling algorithms in such system need to incorporate this constraint while trying to maximize their own scheduling objectives.

In this paper we explore this fundamental problem of LTE SC-FDMA uplink scheduling by adopting the conventional time-domain *Proportional Fair* algorithm to maximize its objective (i.e. proportional fair criteria) in the frequency-domain setting. We show the NP-hardness of the frequency-domain scheduling problem under this contiguous allocation constraint and present a set of practical algorithms fine tuned to this problem. We demonstrate that competitive performance can be achieved in terms of system throughput as well as fairness perspective, which is evaluated using 3GPP LTE system model simulations.

I. INTRODUCTION

In recent years Orthogonal Frequency Division Multiple Access (OFDMA) has been considered as a strong candidate for the broadband air interface for its robustness to multipath fading, higher spectral efficiency and bandwidth scalability, and it has been selected for 3GPP Long-Term Evolution (LTE) downlink (DL) radio access technology. However, one major disadvantage of OFDMA is that the instantaneous transmitted RF power can vary dramatically within a single OFDM symbol. Such an undesirable high peak-to-average power ratio (PAPR) is a serious concern for the uplink (UL), since power consumption is a key consideration for the mobile handsets. As a result of seeking an alternative to OFDMA, Single-carrier FDMA (SC-FDMA) has been selected for LTE uplink multiple access scheme. While keeping most of the advantages of OFDMA (e.g. the same degree of multipath protection), SC-FDMA has significantly lower PAPR, since the underlying waveform is essentially single-carrier. Thus, lower PAPR of SC-FDMA greatly benefits the mobile terminal in terms of transmit power efficiency.

As in DL OFDMA, multiple access in UL SC-FDMA is achieved by assigning different frequency portions of the system bandwidth to individual users based on their channel conditions. Such simultaneous frequency-domain multiplexing of users (inherently in concert with time-domain scheduling) is performed by *frequency domain packet scheduling* (FDPS). In LTE UL, the system bandwidth is divided into multiple subbands (i.e. groups of subcarriers) denoted as *physical*

resource blocks (RBs). In order to achieve large gain from multiuser frequency diversity, a scheduler needs to know the instantaneous radio channel conditions across all users and all RBs, which are fed as input for the frequency-domain adaptive user-to-RB allocation. For example, in LTE UL each user transmits a Sounding Reference Signal (SRS) to the scheduling node (i.e. base station) [1], which is used as *channel quality indicator* (CQI). With CQIs across all users and all RBs, a base station performs RB-to-user assignment at each time slot (e.g. in LTE every 1ms) according to the selected scheduling policy. Thus, in the time-frequency domain, an RB is considered as a minimum scheduling resolution, and also a minimum unit of the data-rate adaptation by *adaptive modulation and coding* (AMC) with a granularity of one sub-frame.

Most of the DL FDPS algorithms proposed so far adopt the well-known time-domain *Proportional Fair* (PF) algorithm as a basic scheduling principle and apply the PF algorithm directly over each RB one-by-one independently. However, such scheduling strategies cannot be employed in the UL SC-FDMA. Due to its single carrier property, SC-FDMA requires that all the RBs allocated to a single user must be *contiguous* in frequency within each time slot (i.e. sub-frame) [5], [6]. Thus, LTE UL FDPS algorithms should respect this constraint while trying to maximize their own scheduling objectives.

In this paper we study this fundamental problem of UL frequency-domain packet scheduling under contiguous RB allocation constraint. We analyze this problem by adopting the widely employed PF algorithm to maximize its objective (i.e. proportional fair criteria) in the frequency-domain setting. The main goal of this paper is to investigate how to adapt the time-domain PF algorithm to this problem framework.

A. The Model

We consider a cellular network whose UL system bandwidth is divided into m RBs, and we have a single base station and n active wireless users. The base station can allocate m RBs to a set of n users. At each time slot multiple RBs (with the contiguity constraint) can be assigned to a single user, each RB however can be assigned to at most one user. In this paper we shall work in an *infinitely backlogged* model in which for each user there is always data available for service. Thus, the base station can schedule all the m RBs every time slot.

We define the indicator variable $x_i^c(t)$ to indicate whether or not RB c is assigned to user i at time slot t . We assume that channel conditions vary across RBs as well as users. The channel conditions typically depends on the channel frequency, so they may be different for different channels;

moreover, they also depends on the user location and the time slot. Therefore, each RB has *user-dependent* and *time-varying* channel condition. We use $r_i^c(t)$ to denote the instantaneous channel rate for user i on RB c at time t . This channel rates are estimated from the CQIs extracted from the UL channel sounding. Thus, if $x_i^c(t) = 1$, then user i can transmit data of size $r_i^c(t)$ on RB c at time slot t .

B. Problem Formulation

In the time-domain context, the well known Proportional Fair (PF) algorithm aims to maximize the logarithmic utility function $\sum_i \log R_i$, where R_i is the long-term service rate of user i . This objective is known as *proportional fair criteria*. In order to maximize $\sum_i \log R_i$, one should maximize $\sum_i d_i(t)/R_i(t)$ where $d_i(t)$ is total data transmitted to user i at time t (proven in [7], [10], [14]). Hence the time-domain PF algorithm always serves the user who maximizes $r_i(t)/R_i(t)$ at each time step t . Note that the PF algorithm achieves high throughput and maintains proportional fairness among all users by giving priority to users with a high-quality channel rate ($r_i(t)$) and a low current average service rate ($R_i(t)$).

We now adapt this time-domain PF metric to the frequency-domain setting with the utility function $\sum_i \log R_i$ as our objective. Let $\lambda_i^c(t) = r_i^c(t)/R_i(t)$ be the *PF metric value* that user i has on RB c at time slot t . We can establish a FDPS version of PF objective function when scheduling time slot t as follows:

$$\max \sum_i \sum_c x_i^c(t) \lambda_i^c(t) \quad (1)$$

It is fairly straightforward to see that objective (1) maximizes $\sum_i d_i(t)/R_i(t)$ at time step t , and therefore achieves proportional fairness, i.e. optimizing objective (1) maximizes the utility function $\sum_i \log R_i$ in the time and frequency domain context. For this reason, most of the proposed DL FDPS scheduling algorithms apply the PF algorithm directly over each RB one-by-one, i.e. for RB c the PF algorithm selects the best user who maximizes $r_i^c(t)/R_i(t)$ at time slot t . However, for LTE UL we need to incorporate the contiguous RB constraint into this objective (1) due to the physical layer requirement of SC-FDMA. The consequence is that we now cannot apply the PF algorithm on each RB one-by-one in isolation. In other words, the isolated local optimization of each RB hardly optimizes the objective (1). Figure 1 exemplifies the case. With the contiguity constraint we may need to serve users with suboptimal PF metric value λ_i^c for some RBs so as to optimize the PF objective (1).

		w/o contiguous requirement										w/ contiguous requirement											
		Max = 85										Max = 83											
user	carrier																						
A		8	7	6	5	4	3	4	5	6	7	8	8	7	6	5	4	3	4	5	6	7	8
B		1	8	1	8	2	8	3	8	2	7	1	1	8	1	8	2	8	3	8	2	7	1
C		6	6	6	5	5	6	4	4	6	6	5	6	6	6	5	5	6	4	4	6	6	5
D		3	4	5	6	7	8	9	8	7	6	5	3	4	5	6	7	8	9	8	7	6	5
E		7	8	6	3	6	4	5	8	2	8	6	7	8	6	3	6	4	5	8	2	8	6

Fig. 1 Maximizing the PF objective. The numbers denote the PF metric values λ_i^c . Dark-colored RBs represent assignment strategies maximizing the objective with/without the contiguity constraint.

Seeking to maximize the PF objective (1) under this contiguity constraint, we present four variations of PF-FDPS algorithm (*Alg1* through *Alg4*). In this paper we explore the fundamental nature of this scheduling problem by investigating how well these algorithms fit into the problem framework.

C. Hardness Result

This contiguous RB allocation constraint is sufficient to make the problem hard.

Theorem 1: LTE UL PF-FDPS problem (i.e. maximizing objective (1) with the contiguous RB constraint) is NP-hard. The proof is omitted due to length constraints. It is presented in the full version of the paper [11].

II. HEURISTIC ALGORITHMS

In this section we present a set of greedy heuristic algorithms for objective (1) under contiguous RB constraint. Our heuristics do not give guaranteed error bound, and moreover we believe that no practical greedy algorithms can give an approximation to this particular problem.¹ Our heuristics fine-tuned to the typical instances of the problem might not perform well in their worst case scenarios, yet their overall performance is very good in practice, as shown in Section III.

A. Alg1: carrier-by-carrier in turn

As a starter, our first greedy heuristic *Alg1* schedules data from RB1 to RB m in sequence, and for each RB c it assigns the best user i who 1) has the maximum PF metric value λ_i^c on c and 2) satisfies the contiguity constraint.

Algorithm 1 : Carrier-by-carrier in turn

- 1: Let U be the set of schedulable users
 - 2: Let $A[m]$ be RB-to-user assignment status
 - 3: **for** RB $c = 1$ to m **do**
 - 4: pick the best user $i \in U$ with largest value λ_i^c
 - 5: assign RB c to user i (i.e. $A[c] \leftarrow i$)
 - 6: Let I be RBs already assigned to user i
 - 7: **if** $I = \emptyset$ **then**
 - 8: $U = U - \{A[c-1]\}$
 - 9: **end if**
 - 10: **end for**
-

Since *Alg1* schedules data from one end side RB, it is not likely to even have a chance to try users' high metric value frequency portions.

B. Alg2: largest-metric-value-RB-first

Viewing this scheduling problem as simply a packing problem, adhering to its rule of thumb "pack large items first" may help in our case. Adopting such a quite intuitive judgement, *Alg2* schedules RBs with largest metric value first. However, it is uncertain how our action should be in the case that, for a certain user i a candidate RB is not adjacent to RBs already assigned to i (e.g. RB3 is first assigned to i , then the next largest value one is RB5 of i . If RB4 is already assigned to other user, then the contiguity constraint prohibits i from being assigned to RB5. Should we however assign RB5 to i if RB4

¹We developed a randomized algorithm that gives $\frac{1}{2}$ -approximation for this problem, but is too complex to be employed for practical wireless scheduling. The algorithm is presented in the full version of the paper [11].

Algorithm 2 : largest-metric-value-RB-first

```
1: Let  $V$  be the sorted list of all the metric values  $\lambda_i^c$  in decreasing order
2: Let  $S$  be the set of not-yet-assigned RBs
3:  $k \leftarrow 1$ 
4: while  $S \neq \emptyset$  do
5:   pick RB  $c$  with  $k^{\text{th}}$  largest metric value  $\lambda_i^c \in V, c \in S$ 
6:   Let  $I$  be RBs already assigned to user  $i$ 
7:   if none is yet assigned to RBs between  $I$  and  $c$  then
8:     Let  $C'$  be all RBs located between  $I$  and  $c$ 
9:      $C' = C' \cup \{c\}$ 
10:    assign all RBs  $\in C'$  to user  $i$ 
11:     $S = S - C'; V = V - \{\lambda_i^c\}; k \leftarrow 1$ 
12:  else
13:     $k \leftarrow k + 1$ 
14:  end if
15: end while
```

is still unoccupied?). Strictly adhering to the argument “pack large items first”, *Alg2* assigns those candidate RBs anyway unless it clearly violates the contiguity constraint (i.e. it assigns RB5 to i).

The price we pay for this a bit aggressive strategy is that we have to assign all the “in-between” RBs to a candidate user (i.e. it assigns RB5 to i , which as a result comes with assignment of RB4 to i , since i is already assigned RB3). The downside of this approach comes from this by-product assignment. Since the length of such “in-between” RBs is arbitrary, a potential improvement in those RBs is likely to be cancelled.

C. Alg3: riding peaks

Seeing the drawback of *Alg2*, we would like to utilize each user’s high valued RBs as much as possible. Let’s look at the PF metric values ($\lambda_i^c(t) = r_i^c(t)/R_i(t)$) at time slot t . One key observation is that, for each user i the denominator ($R_i(t)$) is constant for all RBs, so the resulting value for each RB c is dominated by channel rate ($r_i^c(t)$) only scaled down/up to the current service rate. Thus, at time slot t each user’s RB values fluctuate exactly as the channel rate changes between RB to RB. However, another fundamental physical layer characteristic is that in multi-carrier systems the channel SNR values (i.e. CQI) are correlated in both time and frequency (depending on the Doppler effect and the delay spread) [8], [12], [15]. In other words, if for each user i RB c has good channel rate, then the neighboring RBs ($c-1, c+1$) have high channel rate as well with high probability.

So the key idea of *Alg3* is to “ride users’ peaks” in




Fig. 2 *Alg3* rides peaks.

Algorithm 3 : riding peaks

```
1: Let  $V$  be the sorted list of all the metric values  $\lambda_i^c$  in decreasing order
2: Let  $S$  be the set of not-yet-assigned RBs
3:  $k \leftarrow 1$ 
4: while  $S \neq \emptyset$  do
5:   pick RB  $c$  with  $k^{\text{th}}$  largest metric value  $\lambda_i^c \in V, c \in S$ 
6:   Let  $I$  be RBs already assigned to user  $i$ 
7:   if ( $c$  is adjacent to  $I$ ) or ( $I = \emptyset$ ) then
8:     assign RB  $c$  to user  $i$ 
9:      $S = S - \{c\}; V = V - \{\lambda_i^c\}; k \leftarrow 1$ 
10:  else
11:     $k \leftarrow k + 1$ 
12:  end if
13: end while
```

frequency domain, by exploiting such correlations. Recall that the conventional PF algorithm rides peaks in time domain. *Alg3*, in fact, extends *Alg2*’s rule of thumb: 1) look at large value RBs first; 2) augment them by one neighbor RB. This second rule enforces a bit conservative contiguity condition (i.e. for a certain user i a candidate RB must be *adjacent* to RBs already assigned to i).

Figure 2 illustrates the “peak riding” of *Alg3*. In the beginning user A is first assigned to its high value RBs, while user B and C are assigned to their peak RBs a little bit later. In the end they are all assigned to the RBs around their peaks according to the rules. Note that *Alg2* fails to allocate user B to its high value RBs, since B ’s peak RB is surrounded by a bit higher A ’s peak RBs.

This “peak riding” approach so far seems quite good. There exist, of course the cases where it can lead to arbitrarily bad solutions. If for a certain user the channel rate across RBs changes arbitrarily, then sticking to peaks is not likely a good strategy. As mentioned earlier, we however can find typical instances displaying the frequency-domain correlation among RBs, and in fact, this approach can lead to a measurable improvement on both throughput and short-term fairness in the realistic UL SC-FDMA scenarios as shown in Section III.

D. Alg4: RB grouping

Given that the frequency domain exhibits a correlation (more precisely, correlation between two adjacent RBs), *Alg3* is expected to yield good performance. As mentioned in Section II-C, the channel quality values are indeed correlated in both time and frequency. However, in general the correlation in the frequency-domain is not as strong as the one in the time-domain (frequency-selective fading distortion) [12], [13]. That implies that we have the overall frequency correlation but its granularity may not be as small as one RB (i.e. the smooth lines in Figure 2 may need to be changed to the uneven ones). Figure 3 (overall fluctuation similar to Figure 2 but with some jitters) shows that such a condition incurs poor results by *Alg3*. Since *Alg3* relies on the strong frequency-domain correlation, it is easily cheated by the small-scale variation. In the figure, user B is falsely assigned to the abrupt peak, user A is trapped by the sudden drop, and in the end user C expands its region to that point.

To deal with such small-scale variation, it would help to extend our unit of consideration (i.e. the number of contiguous

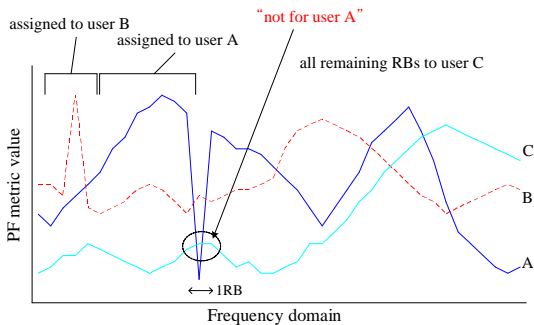


Fig. 3 *Alg3* suffers from small-scale variation.

RBs that we view at a time). This RB grouping might be helpful to catch a bit large-scale fluctuation. *Alg4* makes use of RB grouping to manage the weak frequency-domain correlation. The following questions may arise: “how big should a group be?”, “is it a variable size?”, and “freedom of positioning?”. The harder we try to set up good criteria regarding those questions, it becomes more a quagmire due to the NP-hard nature. Here we set up simple rules: 1) divide m RBs into n groups; 2) apply the “peak riding” over those RB groups. Thus, *Alg4* is an RB-grouping version of *Alg3*; *Alg4* “rides peaks” with the granularity of RB groups (one group = $\lceil \frac{m}{n} \rceil$ RBs). Notice that as n (i.e. the number of users) grows, the group size gets smaller (i.e. we see the smaller-scale fluctuation). As a ground for our choice of $\lceil \frac{m}{n} \rceil$, we argue that it would be beneficial to see the small-scale fluctuation with large number of users, since high multiuser frequency diversity can facilitate the potential improvement from the small-scale peaks.

One can easily find a bad example for *Alg4* and its inapproximability as well. However, such extremely bad instances are unlikely to happen in practice, and in fact, *Alg4* exhibits constantly better performance over *Alg3* on the real traces, particularly when the number of users is not large (as n grows, $\lceil \frac{m}{n} \rceil$ RBs becomes 1 RB).

III. SIMULATIONS

To evaluate the performance of our heuristics, SC-FDMA uplink system level simulations have been conducted based on 3GPP LTE system model. We use traces generated as specified in 3GPP deployment evaluation [2], based on Typical Urban channel model. Table 1 summarizes a list of the default simulation parameters and assumptions.

We analyze the performance of the algorithms in terms of throughput as well as short-term fairness², and assess how well they emulate the proportional fair criteria in this FDPS setting. However, since it is NP-hard to optimize objective (1) under the contiguity constraint, we do not have such an optimal algorithm in our hand. Thus, we use an algorithm that optimizes objective (1) without the constraint as our reference, and we refer to this algorithm as OPT^* . Note that OPT^* offers an upper bound of the optimum. We use Jain’s fairness index [9], measured by the data-rate fairness

²A well-known problem of the conventional time-domain PF scheduling is its poor short-term fairness.

TABLE I Simulation parameters


Parameter	Setting
System bandwidth	20 MHz
Subcarriers per RB	12
RB bandwidth	180 kHz
Number of RBs	96
Cell-level user distribution	Uniform
Number of active users in cell	10, 20, 30, 40, 50
Traffic model	Infinitely backlogged
Transmission time interval (TTI)	1 ms
Channel model	Typical Urban
User speed	3, 30, 120 km/h
User receiver	1x2/MMSE/ZF
Modulation/coding rate settings	QPSK: 1/3, 1/2, 2/3, 3/4 16QAM: 1/2, 2/3, 3/4
HARQ model	Ideal chase combining
HARQ Aak/Nack delay	8 ms
Max. number of HARQ retransmission	3

criterion³: $F_\phi(\Delta t) = \frac{[\sum_{i=1}^N \phi_i(\Delta t)]^2}{N \cdot \sum_{i=1}^N \phi_i(\Delta t)^2}$, where $\phi_i(\Delta t)$ denotes the actual data-rate user i achieved in time interval Δt , with N users in the system.

We first measure the system throughput of our algorithms with varying the number of active users in the cell. As shown in Figure 4(a), *Alg4* results in the highest throughput among our heuristics, followed by *Alg3*, *Alg2*, and *Alg1*. This trend seems to match with our expectation, since *Alg4* and *Alg3* contain more advanced heuristic idea than the other two. In general, *Alg3* performs better than *Alg1* and *Alg2* because *Alg3* seeks to take advantage of each users’ peak while both *Alg1* and *Alg2* are not so fine-tuned enough to effectively utilize multiuser frequency diversity. However, as seen from Figure 4(a), *Alg3* displays the poor performance with small number of active users (e.g. when $n = 10$, it yields even lower throughput than *Alg1* and *Alg2*). Such a result shows the implication of the weak frequency-domain correlation, by which *Alg3* is easily misled into bad solutions. On the other hand, *Alg4* constantly outperforms the other three algorithms in all scenarios. *Alg4* deals with this small-scale variations by widening its view to $\lceil \frac{m}{n} \rceil$ RBs. In the case of small number of active users, *Alg4* expands the RB-group size, and it rides each users’ aggregated peak by catching a bit large-scale fluctuation (it attains 84% of OPT^* while *Alg3* gets 77%). As n grows, *Alg4* adaptively lessens the view so as to exploit the small-scale fluctuation, and its performance gets similar to *Alg3* (when $n = 50$, *Alg4* and *Alg3* reach 95% of OPT^* while the other two get around 86%). It is worth stressing again that OPT^* does not represent the optimum of our objective but simply shows an upper bound of it, where the actual optimum lies between *Alg4* and OPT^* in general.


We now evaluate the short-term fairness of our algorithms with varying the number of active users. Figure 5(a) shows the short-term data-rate fairness $F_\phi(\Delta t)$, in the cell of 30 active users, with extending the time interval window Δt from 10 ms (i.e. 10 TTI) to 50 ms. In this setting, *Alg3* consistently outperforms other algorithms in all intervals, followed by *Alg4*, *Alg1*, and *Alg2*. To understand why *Alg3* provides better short-term fairness than others in this setting, we record

³ $F_\phi(\Delta t)=1$ implies that all users received equal data-rate within time Δt .



(a) cell throughput (b) fairness index (t=20ms)

Fig. 4 System throughput and fairness with varying num. of users



(a) 30 active user case (b) 50 active user case

Fig. 5 Short-term fairness with varying time interval


the number of users scheduled per one TTI for each algorithm. Figure 6(a) plots the average number of users scheduled per one TTI when 30 users are active in the cell. We can see that all of 30 users are likely assigned to all 96 RBs by *Alg3* and *Alg1*.⁴ However, the crucial difference is that *Alg1* is likely to allocate arbitrary rate on each user while *Alg3* seeks to assign users their peak RBs, which helps short-term “fair share” of the frequency resource. Figure 4(b) presents the short-term fairness of 20 ms interval window with increasing number of active users. Interestingly, *Alg1* offers the best fairness when the number of users is large (e.g. $n = 50$). See also Figure 5(b) and 6(b) for fairness and the average number of users scheduled per a TTI with 50 users. With the large number of users, *Alg1* is able to balance users’ rates, but those are not likely from peak RBs.

At this point we note that achieving both high throughput and fairness is a somewhat conflicting goal in general. For example, *Alg1* performs even better than *OPT** in terms of the short-term fairness while *OPT** yields 127% greater throughput over *Alg1*. Hence, to compare the algorithms by a comprehensive metric that takes both throughput and fairness into account. Such a balance is pursued by the proportional fair criteria (i.e. maximizing $\sum_i \log R_i$, where R_i is the long-term service rate for user i), which in fact is our ultimate objective function. Now we assess how well our heuristics emulate the proportional fair objective in our problem framework. In the following table we show the values of the PF criteria with 30 active users in the cell.

	<i>OPT*</i>	<i>Alg1</i>	<i>Alg2</i>	<i>Alg3</i>	<i>Alg4</i>
$\sum_i \log R_i$	223.1	216.5	218.9	220.6	221.6

We can see that *Alg4* has the highest value of $\sum_i \log R_i$, followed by *Alg3*, *Alg2* and *Alg1*. We obtain the same trend

⁴This result seems quite intuitive in the sense that *Alg3* and *Alg1* make assignment decision on one single RB at a time while *Alg2* and *Alg4* assign potentially multiple RBs to a certain user at a time.



(a) 30 active user case (b) 50 active user case

Fig. 6 Average num. of users scheduled per 1 TTI

(with similar gaps between values) in all other scenarios. As we underlined earlier, *OPT** simply represents an upper bound of the optimum of our objective, so the actual optimum has a value of $\sum_i \log R_i$ between *Alg4* and *OPT**. Therefore, among our heuristics *Alg4* has the value of the PF criteria closest to the actual optimum, and it emulates best the PF criteria in UL FDPS setting.

IV. CONCLUSIONS

Due to its single carrier property of SC-FDMA, LTE UL requires the RBs allocated to a single user to be contiguous in frequency. In this paper we explored this fundamental problem of frequency-domain scheduling under contiguous RB allocation constraint. We investigated how to adapt the time-domain PF algorithm to this problem framework. We first showed the NP-hard nature of this problem, then presented a set of practical algorithms fine tuned to this problem. Among them, an algorithm that exploits the frequency-domain correlations in concert with an adaptive RB grouping technique emulates best the PF criteria in the LTE UL FDPS context.

Finally we believe that no practical wireless scheduling algorithms can give an approximation to this particular problem, but whether there actually exists such an algorithm or not still remains as an open problem.

REFERENCES

- [1] Physical layer aspects for evolved Universal Terrestrial Radio Access (UTRA). 3GPP TR 25.814
- [2] Technical specification group radio access networks - Deployment aspects. 3GPP TR 25.943
- [3] Requirements for Evolved UTRA and UTRAN (Release 7). 3 GPP TR 25.913 V7.0.0, June 2005.
- [4] TSG-RAN WG1 #42, R1-050737, “Bandwidth of resource blocks for DL OFDMA”, London, UK, Sep, 2005.
- [5] Moray Rumney. 3GPP LTE: Introducing SIngle-Carrier FDMA Agilent Measurement Journal, 2008.
- [6] 3GPP TSG-RAN WG2 Meeting #57, R2-070585, “Resource fragmentation in LTE uplink”, St. Louis, USA, Feb, 2007.
- [7] M. Andrews. A survey of scheduling theory in wireless data networks. *IMA*, 2005.
- [8] E. Biglieri, J. Proakis, and S. Shamai. Fading channels: Information-theoretic and communications aspects. *IEEE Trans. on Information Theory*, 2000.
- [9] R. Jain, D. M. Chiu, and W. Hawe. A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems. DEC Research Report TR-301.
- [10] H. Kushner and P. Whiting. Asymptotic properties of proportional-fair sharing algorithms. *Allerton*, 2002.
- [11] S.-B. Lee, I. Pefkianakis, A. Meyerson, S. Xu, and S. Lu. Proportional Fair Frequency-Domain Packet Scheduling for 3GPP LTE Uplink. *UCLA TR-090001*, 2009.
- [12] B. Sklar. Rayleigh fading channels in mobile digital communication systems, Part I: Characterization. *IEEE Communications Magazine*, 1997.
- [13] B. Sklar. Rayleigh fading channels in mobile digital communication systems, Part II: Mitigation. *IEEE Communications Magazine*, 1997.
- [14] D. Tse. Multiuser diversity in wireless networks. <http://www.eecs.berkeley.edu/dtse/stanford416.ps>, 2002.
- [15] W. Wang, T. Ottosson, M. Sternad, A. Ahlen, and A. Svensson. Impact of multiuser diversity and channel variability on adaptive OFDM. *IEEE VTC*, 2003.