Multi-User Scheduling in the 3GPP LTE Cellular Uplink

Narayan Prasad, Honghai Zhang, Hao Zhu, Sampath Rangarajan

Multi-User Scheduling in the 3GPP LTE Cellular Uplink

Narayan Prasad*, Honghai Zhang*, Hao Zhu+ and Sampath Rangarajan*
* NEC Labs America, Princeton; + University of Minnesota, Minneapolis;
{prasad, honghai, sampath}@nec-labs.com; zhuh@umn.edu

***Abstract*—We consider resource allocation in the 3GPP Long Term Evolution (LTE) cellular uplink, which will be the most widely deployed next generation cellular uplink. The key features of the 3GPP LTE uplink (UL) are that it is based on a modified form of the orthogonal frequency division multiplexing based multiple acess (OFDMA), which enables channel dependent frequency selective scheduling, and that it allows for multi-user (MU) scheduling wherein multiple users can be assigned the same time-frequency resource. In addition to the considerable spectral efficiency improvements that are possible by exploiting these two features, the LTE UL allows for transmit antenna selection together with the possibility of employing advanced receivers at the base-station, which promise further gains. However, several practical constraints that seek to maintain a low signaling overhead, are also imposed. In this paper, we show that the resulting resource allocation problem is APX-hard and then propose a *local ratio test (LRT)* based constant-factor deterministic polynomial-time approximation algorithm which can accommodate all the practical constraints. Detailed evaluations reveal that the proposed algorithm together with its proposed enhancements offers significant gains.**

I. Introduction

The next generation cellular systems will operate over wideband multi-path fading channels and have chosen OFDMA as their air-interface [1]. The motivating factors behind the choice of OFDMA are that it is an effective means to handle multi-path fading and that it allows for enhancing multi-user diversity gains via channel-dependent frequency-domain scheduling. The deployment of such cellular systems has begun and will accelerate in the coming years. Predominantly these cellular systems will be based on the 3GPP LTE standard [1] since an overwhelming majority of cellular operators have committed to LTE. Our focus in this paper is on the uplink (UL) in these LTE cellular systems and in particular on multi-user (MU) scheduling for the the LTE UL. The UL in LTE systems employs a modified form of OFDMA, referred to as the DFT-Spread-OFDMA [1]. The available system bandwidth is partitioned into multiple resource blocks (RBs), where each RB represents the minimum allocation unit and is a pre-defined set of consecutive subcarriers and OFDM symbols. The scheduler is a frequency domain packet scheduler, which in each scheduling interval assigns these RBs to the individual users. Anticipating a rapid growth in data traffic, the LTE UL has enabled MU scheduling along with transmit antenna selection. Unlike single-user (SU) scheduling, a key feature of MU scheduling is that an RB can be simultaneously assigned to more that one user in the same scheduling interval. MU scheduling is well supported by fundamental capacity and degrees of freedom based analysis [2], [3] and indeed, its promised gains need to be harvested in order to cater to the ever increasing traffic demands. However, several constraints have also been placed by the LTE standard on such MU scheduling (and the resulting MU transmissions). These constraints seek to balance the need to provide scheduling freedom with the need to ensure a low signaling overhead and respect device limitations. The design of an efficient and implementable MU scheduler for the LTE UL is thus an important problem.

In Fig. 1 we highlight the key constraints in LTE MU scheduling by depicting a feasible allocation. Notice first that all RBs assigned to a user must form a chunk of contiguous RBs and each user can be assigned at-most one such chunk. This restriction allows us to exploit frequency domain channel variations via localized assignments (there is complete freedom in choosing the location and size of each such chunk) while respecting strict limits on the per-user transmit peak-to-average-power-ratio (PAPR). Note also that there should be a complete overlap among any two users that share an RB. In other words, if any two users are co-scheduled on an RB then those two users must be co-scheduled on all their assigned RBs. This constraint is a consequence of Zadoff-Chu (ZC) sequences (and their cyclic shifts) being used as pilot sequences for channel estimation in the LTE UL [1]. In particular, two users who share an RB must have a complete overlap of their assigned pilot sequences in order to ensure inter-user pilot orthogonality at the UL receiver, which in turn is necessary for reliable channel estimation. In addition, to minimize the signaling overhead, each scheduled user can transmit with only one power level (or power spectral density (PSD)) on all its assigned RBs. This PSD is implicitly determined by the number of RBs assigned to that user, i.e., the user divides its total power equally among all its

978-3-901882-45-6/2012 - Copyright is with IFIP

assigned RBs subject possibly to a spectral mask constraint. While this constraint significantly decreases the signaling overhead involved in conveying the scheduling decisions to the users, it does not result in any significant performance degradation. This is due to the fact that the multi-user diversity effect ensures that each user is scheduled on the set of RBs on which it has relatively good channels. A constant power allocation over such *good* channels results in a negligible loss [4]. The LTE UL further assumes that each user can have multiple transmit antennas but is equipped with only one power amplifier due to cost constraints. Accordingly, it allows a basic precoding in the form of transmit antenna selection where each scheduled user can be informed about the transmit antenna it should employ in a scheduling interval.

Finally scheduling in LTE UL must respect control channel overhead constraints and interference limit constraints. The former constraints arise because the scheduling decisions are conveyed to the users on the downlink control channel, whose limited capacity in turn places a limit on the set of users that can be scheduled. The latter constraints are employed to mitigate intercell interference. In the sequel it is shown that both these types of constraints can be posed as column-sparse and generic knapsack (linear packing) constraints, respectively.

The goal of this work is to design practical MU resource allocation algorithms for the LTE cellular uplink, where the term resource refers to RBs, modulation and coding schemes (MCS), power levels as well as choice of transmit antennas. In particular, we consider the design of resource allocation algorithms via weighted sum rate utility maximization, which accounts for finite user queues (buffers) and practical MCS. In addition, the designed algorithms comply with all the aforementioned practical constraints. Our main contributions are as follows:

(1): We show that while the *complete* overlap constraint along with the at-most one chunk per scheduled user constraint make the resource allocation problem APX-hard, they greatly facilitate the use of local ratio test (LRT) based methods [5], [6]. We then design an LRT based polynomial time deterministic constant-factor approximation algorithm. A remarkable feature of this LRT based algorithm is that it is an end-to-end solution which can accommodate all constraints. Simulation results show that the proposed algorithm has good average performance, much superior to its worst-case guarantee and achieves more than 80% of a corresponding linear programming (LP) based upper bound.

(2): We then propose an enhancement that can significantly reduce the complexity of the LRT based MU scheduling algorithm while offering identical performance, as well as an enhancement that can yield good performance improvements with a very small additional complexity. Indeed, in simulation examples that consider transmit antenna selection as well as different receiver choices, the proposed complexity reduction enhancement is shown to result in over 80% reduction in terms of metric computation complexity. On the other hand, the other enhancement improves the performance (in terms of the cell average spectral efficiency) by about 10% while incurring a very modest complexity increase of 2 to 3%. The performance of our algorithm with this latter enhancement is more than 90% of a corresponding LP upper bound.

(3): The performance of the proposed LRT based MU scheduling algorithm together with its enhancements are evaluated for different BS receiver options via elaborate system level simulations that fully conform to the 3GPP evaluation methodology. It is seen that the proposed LRT based MU scheduling algorithm along with an advanced BS receiver can yield over 27% improvement in cell average throughout along with over 10% cell edge throughput improvement compared to SU scheduling.

A. Related Work

Resource allocation for the OFDMA networks has been the subject of intense research [7]–[11] with most of the focus being on the downlink. A majority of OFDMA resource allocation problems hitherto considered are single-user (SU) scheduling problems, which attempt to maximize a system utility by assigning non-overlapping subcarriers to users along with transmit power levels for the assigned subcarriers. These problems have been formulated as *continuous optimization problems*, which are in general non-linear and non-convex. As a result several approaches based on the game theory [12], dual decomposition [7] or the analysis of optimality conditions [13] have been developed. Recent works have focused on emerging cellular standards and have modeled the resource allocation problems as constrained integer programs. Prominent examples are [10], [14] which consider the design of downlink SU-MIMO schedulers for LTE cellular systems and derive constant factor approximation algorithms.

Resource allocation for the DFT-Spread-OFDMA uplink has been relatively much less studied with [6], [15], [16] being the recent examples. In particular, [6], [15] show that the SU LTE UL scheduling problem is APX-hard and provide constant-factor approximation algorithms, whereas [16] extends the algorithms of [6], [15] to the SU-MIMO LTE-Advanced scheduling. The algorithm proposed in [6] is based on an innovative application of the LRT technique, which was developed earlier in [5]. However, we emphasize that the algorithms in [6], [15], [16] cannot incorporate MU scheduling and also cannot incorporate knapsack constraints. To the best of our knowledge the design of approximation algorithms for MU scheduling in the LTE uplink has not been considered before.

II. MU SCHEDULING IN THE LTE UL

Consider a single-cell with K users and one BS which is assumed to have $N_r \geq 1$ receive antennas. Suppose that user k has $N_t \geq 1$ transmit antennas and its power budget is P_k. We let N denote the total number of RBs.

We consider the problem of scheduling users in the frequency domain in a given scheduling interval. Let $\alpha_k,\ 1 \leq k \leq K$ denote the weight of the k^{th} user which is an input to the scheduling algorithm and is updated using the output of the scheduling algorithm in every scheduling interval, say according to the proportional fairness rule [17]. Letting r_k denote the rate assigned to the k^{th} user (in bits per N RBs), we consider the following weighted sum rate utility maximization problem,

$$\max \sum_{1 \leq k \leq K} \alpha_k r_k, \tag{1}$$

where the maximization is over the assignment of resources to the users **subject to:**

Decodability constraint: The rates assigned to the scheduled users should be decodable by the base-station receiver. Notice that unlike SU scheduling, MU scheduling allows for multiple users to be assigned the same RB. As a result the rate that can be achieved for user k need not be only a function of the resources assigned to the k^{th} user but can also depend on the those assigned to the other users as well.

One transmit antenna and one power level per user: Each user can transmit using only one power amplifier due to cost constraints so, only a basic precoding in the form of transmit antenna selection is possible. In addition, each scheduled user is allowed to transmit with only one power level (or power spectral density (PSD)) on all its assigned RBs.

At most one chunk per-user and at-most T users per RB: The set of RBs assigned to each scheduled user should form one chunk, where each chunk is a set of contiguous RBs. Further at-most T users can be co-scheduled on a given RB. T is expected to be small number typically two and no greater than four.

Complete overlap constraint: If any two users are assigned a common RB then those two users must be assigned the same set of RBs. Feasible RB allocation and co-scheduling of users in LTE MU UL is shown in Fig 1.

Finite buffers and finite MCS: Users in a practical UL will have bursty traffic which necessitates considering finite buffers. In addition, only a finite set of MCS (29 possibilities in the LTE network) can be employed.

Control channel overhead constraints: Every user that is given an UL grant (i.e., is scheduled on at least one RB) must be informed about its assigned MCS and the set of RBs on which it must transmit along with possibly the transmit antenna it should employ. This information is sent on the DL control channel of limited capacity which imposes a limit on the set of users that can be scheduled. In particular, the scheduling information of a user is encoded and formatted into one packet (henceforth referred to as a control packet), where the size of the control packet must be selected from a predetermined set of sizes. A longer (shorter) control packet is used for a cell edge (cell interior) user. In the LTE/LTE-A systems each user is assigned one search region when it enters the cell. In each scheduling interval it then searches for the control packet (containing the scheduling decisions made for it) only in that region of the downlink control channel, as well as a region common to all users. By placing restrictions on the location where a particular user's control packet can be sent and the size of that packet, the system can reduce the number of blind decoding attempts that have to be made by that user in order to receive its control packet. We note that a user is unaware of whether there is a control packet intended for it and consequently must check all possible locations where its control packet could be present assuming each possible packet size. Each control packet carries a CRC bit sequence scrambled using the unique user identifier which helps the user deduce whether the examined packet is meant for it [1].

Per sub-band interference limit constraints: Inter-cell interference mitigation can be performed by imposing interference limit constraints. In particular, on one or more subbands, the cell of interest can ensure that the total interference imposed by its scheduled users on a neighboring base-station is below a specified limit.

We define the set $\boldsymbol{\mathcal{C}}$ as the set containing N length vectors such that any $\mathbf{c} \in \boldsymbol{\mathcal{C}}$ is binary-valued with $(\{0,1\})$ elements and contains a contiguous sequence of ones with the remaining elements being zero. Here we say an RB i belongs to $\mathbf{c}$ ($i \in \mathbf{c}$) if $\mathbf{c}$ contains a one in its i^{th} position, i.e., $c(i) = 1$. Note then that each $\mathbf{c} \in \boldsymbol{\mathcal{C}}$ denotes a valid assignment of RBs since it contains one contiguous chunk of RBs. Also $\mathbf{c}_1$ and $\mathbf{c}_2$ are said to intersect if there is some RB that belongs to both $\mathbf{c}_1$ and $\mathbf{c}_2$. For any $\mathbf{c} \in \boldsymbol{\mathcal{C}}$, we will use $\mathrm{Tail}(\mathbf{c})$ ($\mathrm{Head}(\mathbf{c})$) to return the largest (smallest) index that contains a one in $\mathbf{c}$. Thus, each $\mathbf{c} \in \boldsymbol{\mathcal{C}}$ has ones in all positions $\mathrm{Head}(\mathbf{c}), \cdots, \mathrm{Tail}(\mathbf{c})$ and zeros elsewhere. Further, we define $\{\mathcal{G}_1, \cdots, \mathcal{G}_L\}$ to be a partition of $\{1, \cdots, K\}$ with the understanding that all users that belong to a common set (or group) $\mathcal{G}_s$, for any $1 \leq s \leq L$, are mutually incompatible. In other words at-most one user from each group $\mathcal{G}_s$ can be scheduled in a scheduling interval. Notice that by choosing $L = K$ and $\mathcal{G}_s = \{s\}$, $1 \le s \le K$ we obtain the case where all users are mutually compatible. Let us define a family of subsets, $\boldsymbol{\mathcal{U}}$, as

$$\boldsymbol{\mathcal{U}} = \{\mathcal{U} \subseteq \{1, \cdots, K\} : |\mathcal{U}| \le T \ \& \ |\mathcal{U} \cap \mathcal{G}_s| \le 1 \ \forall \ 1 \le s \le L\}$$

and let $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{U}} \times \boldsymbol{\mathcal{C}}$.

We can now pose the resource allocation problem as

$$\max \sum_{(\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}}} p(\mathcal{U}, \mathbf{c}) \mathcal{X}(\mathcal{U}, \mathbf{c}), \ \ \text{s.t.}$$

$$\text{For each group } \mathcal{G}_s, \sum_{\substack{(\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}} \\ \mathcal{U}: \mathcal{U} \cap \mathcal{G}_s \neq \phi}} \mathcal{X}(\mathcal{U}, \mathbf{c}) \le 1;$$

$$\text{For each RB } i, \sum_{\substack{(\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}} \\ \mathbf{c}: i \in \mathbf{c}}} \mathcal{X}(\mathcal{U}, \mathbf{c}) \le 1;$$

$$\sum_{(\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}}} \beta^q(\mathcal{U}, \mathbf{c}) \mathcal{X}(\mathcal{U}, \mathbf{c}) \le 1, \ 1 \le q \le J;$$

$$\sum_{(\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}}} \alpha^q(\mathcal{U}, \mathbf{c}) \mathcal{X}(\mathcal{U}, \mathbf{c}) \le 1, \ q \in \mathcal{I}, \qquad (2)$$

where ϕ denotes the empty set and $\mathcal{X}(\mathcal{U}, \mathbf{c})$ is an indicator function that returns one if users in $\mathcal{U}$ are co-scheduled on the chunk indicated by $\mathbf{c}$. Note that the first constraint ensures that at-most one user is scheduled from each group and that each scheduled user is assigned at-most one chunk. In addition this constraint also enforces the complete overlap constraint. The second constraint enforces non-overlap among the assigned chunks. Note that $p(\mathcal{U}, \mathbf{c})$ denotes the weighted sum-rate obtained upon co-scheduling the users in $\mathcal{U}$ on the chunk indicated by $\mathbf{c}$. We emphasize that *there is complete freedom with respect to the computation of $p(\mathcal{U}, \mathbf{c})$. Indeed, it can accommodate finite buffer and practical MCS constraints, account for any particular receiver employed by the base station and can also incorporate any rule to assign a transmit antenna and a power level to each user in $\mathcal{U}$ over the chunk* $\mathbf{c}$.

The first set of J knapsack constraints in (2), where J is arbitrary but fixed, are generic knapsack constraints. Without loss of generality, we assume that the weight of the pair $(\mathcal{U}, \mathbf{c})$ in the q^{th} knapsack, $\beta^q(\mathcal{U}, \mathbf{c})$, lies in the interval $[0, 1]$. Notice that we can simply drop each vacuous constraint, i.e., each constraint q for which $\sum_{(\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}}} \beta^q(\mathcal{U}, \mathbf{c}) \le 1$. The second set of knapsack constraints are *column-sparse binary knapsack constraints*. In particular, for each $(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$ and $q \in \mathcal{I}$ we have that $\alpha^q(\mathcal{U}, \mathbf{c}) \in \{0, 1\}$. Further, we have that for each $(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$, $\sum_{q \in \mathcal{I}} \alpha^q(\mathcal{U}, \mathbf{c}) \le \Delta$, where Δ is arbitrary but fixed and denotes the column-sparsity level. Note that here the cardinality of $\mathcal{I}$ can scale polynomially in KN keeping Δ fixed. Together these two sets of knapsack constraints can enforce a variety of practical constraints, including the control channel and the interference limit constraints. For instance, defining a generic knapsack constraint as $\beta^1(\mathcal{U}, \mathbf{c}) = \frac{|\mathcal{U}|}{\tilde{K}}, \ \forall \ (\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$, for any given input $\tilde{K}$ can enforce that no more that $\tilde{K}$ can be scheduled in a given interval, which represents a coarse control channel constraint. In a similar vein, consider any given choice of a victim adjacent base-station and a sub-band with the constraint that the total interference caused to the victim BS by users scheduled in the cell of interest, over all the RBs in the subband, should be no greater than a specified upper bound. This constraint can readily modeled using a generic knapsack constraint where the weight of each $(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$ is simply the ratio of the total interference caused by users in $\mathcal{U}$ to the victim BS over RBs that are in $\mathbf{c}$ as well as the specified subband, and the specified upper bound. The interference is computed using the transmission parameters (such as the power levels, transmit antennas etc) that yield the metric $p(\mathcal{U}, \mathbf{c})$. A finer modeling of the LTE control channel constraints is more involved since it needs to employ the column-sparse knapsack constraints together with the user incompatibility constraints and is detailed in [18].

Note that for a given K, N, an instance of the problem in (2) consists of a finite set $\mathcal{I}$ of indices, a partition $\{\mathcal{G}_1, \cdots, \mathcal{G}_L\}$, metrics $\{p(\mathcal{U}, \mathbf{c})\} \ \forall \ (\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$ and weights $\{\beta^q(\mathcal{U}, \mathbf{c})\}, \ \forall \ (\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}, 1 \le q \le J$ and $\{\alpha^q(\mathcal{U}, \mathbf{c})\}, \ \forall \ (\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}, q \in \mathcal{I}$. Then, in order to solve (2) for a given instance, we first partition the set $\boldsymbol{\mathcal{M}}$ into two parts as $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{M}}^{\text{narrow}} \cup \boldsymbol{\mathcal{M}}^{\text{wide}}$, where we define $\boldsymbol{\mathcal{M}}^{\text{narrow}} = \{(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}} : \beta^q(\mathcal{U}, \mathbf{c}) \le 1/2, \ \forall \ 1 \le q \le J\}$ so that $\boldsymbol{\mathcal{M}}^{\text{wide}} = \boldsymbol{\mathcal{M}} \setminus \boldsymbol{\mathcal{M}}^{\text{narrow}}$. We then define J sets, $\boldsymbol{\mathcal{V}}^{(1)}, \cdots, \boldsymbol{\mathcal{V}}^{(J)}$ that cover $\boldsymbol{\mathcal{M}}^{\text{wide}}$ (note that any two of these sets can mutually overlap) as $(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{V}}^{(q)}$ iff $\beta^q(\mathcal{U}, \mathbf{c}) > 1/2$ for $q = 1, \cdots, J$. Recall that T, J are fixed and note that the cardinality of $\boldsymbol{\mathcal{M}}$, $|\boldsymbol{\mathcal{M}}|$, is $O(K^T N^2)$ and that $\boldsymbol{\mathcal{M}}^{\text{narrow}}$ and $\{\boldsymbol{\mathcal{V}}^{(q)}\}$ can be determined in polynomial time. Next, we propose Algorithm I which possesses the optimality given below. The complexity of Algorithm I, which is essentially determined by that of its module Algorithm IIa, scales polynomially in KN (recall that T is a constant) A detailed discussion on the complexity along with steps to reduce it are deferred to the next section. We offer the following theorem which is proved in [18].

Theorem 1. *The problem in (2) is APX-hard, i.e., there is an $\epsilon > 0$ such that it is NP hard to obtain a $1 - \epsilon$ approximation algorithm for (2). Let $\hat{W}^{\text{opt}}$ denote the optimal weighted sum rate obtained upon solving (2) and let $\hat{W}$ denote the weighted sum rate obtained upon* *using Algorithm I. Then, we have that*

$$\hat{W} \geq \begin{cases} \frac{\hat{W}^{\text{opt}}}{1+T+\Delta+2J}, & \text{If } \boldsymbol{\mathcal{M}}^{\text{wide}} = \phi \\ \frac{\hat{W}^{\text{opt}}}{1+T+\Delta+3J}, & \text{Otherwise} \end{cases} \tag{3}$$

An interesting observation that follows from the proof of Theorem 1 is that any optimal allocation over $\boldsymbol{\mathcal{M}}^{\text{wide}}$ can include at-most one pair from each $\boldsymbol{\mathcal{V}}^{(q)},\ 1 \leq q \leq J$. Then since the number of pairs in each $\boldsymbol{\mathcal{V}}^{(q)},\ 1 \leq q \leq J$ is $O(K^T N^2)$, we can determine an optimal allocation yielding $\hat{W}^{\text{opt,wide}}$ via exhaustive enumeration with a high albeit polynomial complexity (recall that T and J are assumed to be fixed). Thus, by using exhaustive enumeration instead of Algorithm IIb, we can claim the following result.

Corollary 1. *Let $\hat{W}^{\text{opt}}$ denote the optimal weighted sum rate obtained upon solving (2) and let $\hat{W}$ denote the weighted sum rate obtained upon using Algorithm II albeit with exhaustive enumeration over $\boldsymbol{\mathcal{M}}^{\text{wide}}$. Then, we have that*

$$\hat{W} \geq \begin{cases} \frac{\hat{W}^{\text{opt}}}{1+T+\Delta+2J}, & \text{If } \boldsymbol{\mathcal{M}}^{\text{wide}} = \phi \\ \frac{\hat{W}^{\text{opt}}}{2+T+\Delta+2J}, & \text{Otherwise} \end{cases} \tag{4}$$

For notational simplicity, henceforth unless otherwise mentioned, we assume that all users are mutually compatible, i.e., $L = K$ with $\mathcal{G}_s = \{s\},\ 1 \leq s \leq K$.

III. Complexity Reduction

In this section we present key techniques to significantly reduce the complexity of our proposed local ratio test based multi-user scheduling algorithm. As noted before the complexity of Algorithm I is dominated by that of its component Algorithm IIa. Accordingly, we focus our attention on Algorithm IIa and without loss of generality we assume that $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{M}}^{\text{narrow}}$. Notice that hitherto we have assumed that all the metrics $\{p(\mathcal{U}, \mathbf{c}) : (\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}\}$ are available. In practise, computing these $O(K^T N^2)$ metrics, which are often complicated non-linear functions, is the main bottleneck and indeed must be accounted for in the complexity analysis. Before proceeding, we make the following assumption that is satisfied by all physically meaningful metrics.

Assumption 1. Sub-additivity*: We assume that for any $(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$*

$$p(\mathcal{U}, \mathbf{c}) \leq p(\mathcal{U}_1, \mathbf{c}) + p(\mathcal{U}_2, \mathbf{c}),\ \forall\ \mathcal{U}_1, \mathcal{U}_2 : \mathcal{U} = \mathcal{U}_1 \cup \mathcal{U}_2.$$

The following features can then be exploited for a significant reduction in complexity.

- *On demand metric computation:* Notice in Algorithm IIa that the metric for any $(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$, where $\text{Tail}(\mathbf{c}) = j$ for some $j = 1, \cdots, N$, needs to be computed only at the j^{th} iteration at which point we need to determine

$$p'(\mathcal{U}, \mathbf{c}) = p(\mathcal{U}, \mathbf{c}) - \Gamma^{(j)}(\mathcal{U}, \mathbf{c}), \tag{5}$$

where the offset factor $\Gamma^{(j)}(\mathcal{U}, \mathbf{c})$ is given by

$$\Gamma^{(j)}(\mathcal{U}, \mathbf{c}) = \sum_{(\mathcal{U}_m^*, \mathbf{c}_m^*) \in \mathcal{S}} \Big(\tilde{p}(\mathcal{U}_m^*, \mathbf{c}_m^*)\mathcal{E}((\mathcal{U}, \mathbf{c}), (\mathcal{U}_m^*, \mathbf{c}_m^*)) + 2\tilde{p}(\mathcal{U}_m^*, \mathbf{c}_m^*) \max_{1 \leq q \leq J} \{\beta^q(\mathcal{U}_m^*, \mathbf{c}_m^*)\}\mathcal{E}^c((\mathcal{U}, \mathbf{c}), (\mathcal{U}_m^*, \mathbf{c}_m^*))\Big)$$

and where $\tilde{p}(\mathcal{U}_m^*, \mathbf{c}_m^*)$ is equal to the $p'(\mathcal{U}_m^*, \mathbf{c}_m^*)$ computed for the pair selected at the m^{th} iteration with $m \leq j-1$ and $\mathcal{E}((\mathcal{U}, \mathbf{c}), (\mathcal{U}_m^*, \mathbf{c}_m^*))$ denotes an indicator (with $\mathcal{E}^c((\mathcal{U}, \mathbf{c}), (\mathcal{U}_m^*, \mathbf{c}_m^*)) = 1 - \mathcal{E}((\mathcal{U}, \mathbf{c}), (\mathcal{U}_m^*, \mathbf{c}_m^*)))$ which is true when $\mathcal{U}_m^* \cap \mathcal{U} \neq \phi$ or $\mathbf{c} \cap \mathbf{c}_m^* \neq \phi$ or $\exists\ q \in \mathcal{I} : \alpha^q(\mathcal{U}_m^*, \mathbf{c}_m^*) = \alpha^q(\mathcal{U}, \mathbf{c}) = 1$. Further note that $p'(\mathcal{U}, \mathbf{c})$ in (5) is required only if it is strictly positive. Then, an important observation is that if at the j^{th} iteration, we have already computed $p(\mathcal{U}_1, \mathbf{c})$ and $p(\mathcal{U}_2, \mathbf{c})$ for some $\mathcal{U}_1, \mathcal{U}_2 : \mathcal{U} = \mathcal{U}_1 \cup \mathcal{U}_2$, then invoking the sub-additivity property we have that

$$p'(\mathcal{U}, \mathbf{c}) \leq p(\mathcal{U}_1, \mathbf{c}) + p(\mathcal{U}_2, \mathbf{c}) - \Gamma^{(j)}(\mathcal{U}, \mathbf{c}), \tag{6}$$

so that if the RHS in (6) is not strictly positive or if it is less than the greatest value of $p'(\mathcal{U}', \mathbf{c}')$ computed in the current iteration for some other pair $(\mathcal{U}', \mathbf{c}') : \text{Tail}(\mathbf{c}') = j$, then *we do not need to compute $p'(\mathcal{U}, \mathbf{c})$ and hence the metric $p(\mathcal{U}, \mathbf{c})$.*

- *Selective update* Note that in the j^{th} iteration, once the best pair $(\mathcal{U}_j^*, \mathbf{c}_j^*)$ is selected and it is determined that $p'(\mathcal{U}_j^*, \mathbf{c}_j^*) > 0$, we need to update the metrics for pairs $(\mathcal{U}', \mathbf{c}') : \text{Tail}(\mathbf{c}') \geq j+1$, since only such pairs will be considered in future iterations. Thus, the offset factors $\{\Gamma^{(j)}(\mathcal{U}', \mathbf{c}')\}$ need to be updated only for such pairs, via

$$\Gamma^{(j+1)}(\mathcal{U}', \mathbf{c}') = \Gamma^{(j)}(\mathcal{U}', \mathbf{c}') + p'(\mathcal{U}_j^*, \mathbf{c}_j^*)\mathcal{E}((\mathcal{U}', \mathbf{c}'), (\mathcal{U}_j^*, \mathbf{c}_j^*)) + 2p'(\mathcal{U}_j^*, \mathbf{c}_j^*) \max_{1 \leq q \leq J} \{\beta^q(\mathcal{U}', \mathbf{c}')\}\mathcal{E}^c((\mathcal{U}', \mathbf{c}'), p'(\mathcal{U}_j^*, \mathbf{c}_j^*)).$$

Further, if by exploiting sub-additivity we can deduce that $p'(\mathcal{U}', \mathbf{c}') \leq 0$ for any such pair, then we can drop such a pair along with its offset factor from future consideration.

IV. Improving Performance via a Second Phase

A potential drawback of the LRT based algorithm is that some RBs may remain un-utilized, i.e., they may not be assigned to any user. Notice that when the final stack $\mathcal{S}'$ is built in the while-loop of Algorithm IIa, an allocation or pair from the top of stack $\mathcal{S}$ is added to stack $\mathcal{S}'$ only if it does not conflict with those already in stack $\mathcal{S}'$. Often multiple pairs from $\mathcal{S}$ are dropped due to such conflicts resulting in spectral holes formed by unassigned RBs. To mitigate this problem, we perform a second phase. The second phase consists of running Algorithm IIa again albeit with modified metrics

$\{\check{p}(\mathcal{U},\mathbf{c}) : (\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}}^{\text{narrow}}\}$ which are obtained via the following steps.

1) Initialize $\check{p}(\mathcal{U},\mathbf{c}) = p(\mathcal{U},\mathbf{c}),\ \forall\ (\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}}^{\text{narrow}}$. Let $\mathcal{S}'$ be obtained as the output of Algorithm IIa when it is implemented first.
2) For each $(\mathcal{U},\mathbf{c}) \in \mathcal{S}'$, we ensure that any user in $\mathcal{U}$ is not scheduled by phase two in any other user set save $\mathcal{U}$, by setting

$$\check{p}(\mathcal{U}',\mathbf{c}') = 0 \text{ if } \mathcal{U}' \neq \mathcal{U} \;\&\; \mathcal{U}' \cap \mathcal{U} \neq \phi,\ \forall\ (\mathcal{U}',\mathbf{c}') \in \boldsymbol{\mathcal{M}}^{\text{narrow}}.$$

3) For each $(\mathcal{U},\mathbf{c}) \in \mathcal{S}'$, we ensure that no other user set save $\mathcal{U}$ is assigned any RB in $\mathbf{c}$, by setting

$$\check{p}(\mathcal{U}',\mathbf{c}') = 0 \text{ if } \mathcal{U}' \neq \mathcal{U} \;\&\; \mathbf{c}' \cap \mathbf{c} \neq \phi,\ \forall\ (\mathcal{U}',\mathbf{c}') \in \boldsymbol{\mathcal{M}}^{\text{narrow}}.$$

4) For each $(\mathcal{U},\mathbf{c}) \in \mathcal{S}'$, we ensure that the allocation $(\mathcal{U},\mathbf{c})$ is either unchanged by phase two or is expanded, by setting

$$\check{p}(\mathcal{U},\mathbf{c}') = \begin{cases} p(\mathcal{U},\mathbf{c}'), & \text{If Tail}(\mathbf{c}') \geq \text{Tail}(\mathbf{c}) \;\&\; \text{Head}(\mathbf{c}') \leq \text{Head}(\mathbf{c}) \\ 0, & \text{Otherwise} \end{cases}$$

A consequence of using the modified metrics is that the second phase has a significantly less complexity since a large fraction of the allocations are disallowed. While the second phase does not offer any improvement in the approximation factor, simulation results presented in the sequel reveal that it offers a good performance improvement with very low complexity addition.

V. Simulation Results: Single cell Setup

We simulate an uplink wherein the BS is equipped with four receive antennas. The system has 20 RBs to serve 10 active users, all of whom have identical maximum transmit powers. The remaining relevant parameters are provided in [18]. For simplicity, we assume that there are no knapsack constraints and that $T = 2$. Notice then that since $\boldsymbol{\mathcal{M}} = \boldsymbol{\mathcal{M}}^{\text{narrow}}$ we can directly use Algorithm IIa.

In Fig. 2 we plot the normalized spectral efficiencies versus the average transmit SNR (dB), obtained by dividing each spectral efficiency by the one yielded by Algorithm IIa when only single user (SU) scheduling is allowed, which in turn can be emulated by setting all metrics $p(\mathcal{U},\mathbf{c}) : (\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}}$ in (2) to be zero whenever $|\mathcal{U}| \geq 2$. In all but one considered LRT based scheduling schemes, Algorithm IIa with the second phase described in Section III is employed. In the remaining one (denoted using "1Step" in the legend) the second phase is not employed. Further, in the legend, "AS" corresponds to the scenario where each user has two transmit antennas so that the BS can also exploit transmit antenna selection, whereas in all other cases each user has only one transmit antenna. Moreover, "MMSE" and "SIC" correspond to the scenarios where the BS employs the linear MMSE receiver and the advanced successive interference cancelation receiver, respectively.[1] For clarity, in the figure we plot the linear programming (LP) upper bounds only for the case when the BS employs the MMSE receiver. From Fig. 2 and the ones given in [18], we have the following observations:

- For both SIC and MMSE receivers, the performance of Algorithm IIa (without the second phase) is more than 80% of the respective LP upper bounds, which is much superior to the worst case guarantee $1/3$ (obtained by specializing the result in (3) by setting $\boldsymbol{\mathcal{M}}^{\text{wide}} = \phi, T = 2$ and $\Delta = J = 0$). Further, for both the receivers the performance of Algorithm IIa with the second phase is more than 90% of the respective LP upper bounds. The same conclusions can be drawn when antenna selection is also exploited by the BS.
- The SIC receiver results in a small gain (1.5% to 2.5%) over the MMSE receiver. Note that antenna selection seems to provide a much larger gain (6% to 8%) than the one offered by the advanced SIC receiver. This observation must be tempered by the facts that the simulated scenario is more favorable for antenna selection.
- MU scheduling offers substantial gains over SU scheduling (up-to 75% for the considered SNRs).

Next, in Fig. 3 we plot the normalized complexities for the aforementioned scheduling schemes. Here the complexity of a scheduling scheme is determined by the complexity of the metric computations made by it. In all cases the second phase is performed for Algorithm IIa and more importantly the sub-additivity property together with the on-demand metric computation feature are exploited, as described in Section III, to avoid redundant metric computations. All schemes compute the metrics $\{p(\mathcal{U},\mathbf{c}) : (\mathcal{U},\mathbf{c}) \in \boldsymbol{\mathcal{M}} \;\&\; |\mathcal{U}| = 1\}$ and each such metric is deemed to have unit complexity when each user has one transmit antenna and a complexity of two units when each user two transmit antennas. On the other hand, for each evaluated $p(\mathcal{U},\mathbf{c}) : |\mathcal{U}| = 2$, the complexity is taken to be two units when each user has one transmit antenna and the BS employs the MMSE receiver and one unit when each user has one transmit antenna and the BS employs the SIC receiver. The latter stems from the fact that with the SIC receiver, one of the users sees an interference free channel. Thus, its contribution to the metric is equal to the already computed single-user metric determined for the allocation when that user is scheduled alone on the corresponding chunk. Similarly, for each evaluated $p(\mathcal{U},\mathbf{c}) : |\mathcal{U}| = 2$, the complexity is taken to be eight units when each user

[1] Note that for SU scheduling MMSE and SIC receivers are equivalent.

has two transmit antennas and the BS employs transmit antenna selection together with the MMSE receiver and four units for the case when the BS employs transmit antenna selection together with the SIC receiver. Note that MMSE-Total, SIC-Total denote the complexities obtained by counting the corresponding complexities for all pairs $(\mathcal{U}, \mathbf{c}) \in \boldsymbol{\mathcal{M}}$, respectively, whereas MMSE-AS-Total, SIC-AS-Total denote the total complexities obtained when antenna selection is employed by the BS together with the MMSE receiver and the SIC receiver, respectively. Note that all complexities in Fig. 3 are normalized by MMSE-AS-Total. The key takeaway from Fig. 3 is that exploiting sub-additivity together with the on-demand metric computation can result in very significant complexity reduction. In particular, as per our definition of complexity, more than 80% reduction can be obtained for the MMSE receiver and more than 75% reduction can be obtained for the SIC receiver, with the respective gains being larger when antenna selection is also exploited. Further, we note that considering Algorithm IIa, the second phase itself adds a very small complexity overhead but results in a large performance improvement. To illustrate this, for the MMSE receiver the complexity overhead ranges from 2 to 4%, whereas the performance improvement ranges from 9 to 13%, respectively. We note that the complexity reductions are even larger when the complexity computed for each $p(\mathcal{U}, \mathbf{c})$ is also multiplied by the size of the chunk indicated by $\mathbf{c}$ [18].

VI. System Level Simulation Results

We now present the performance of our MU scheduling algorithms via detailed system level simulations. The simulation parameters conform to those used in 3GPP LTE evaluations and are given in [18]. In all cases inter-cell interference suppression (IRC) is employed by each base-station (BS).

We consider the case when each cell (or sector) has an average of 10 users and where there are no knapsack constraints. In Table IV we report the cell average and cell edge spectral efficiencies. The percentage gains shown for the MU scheduling schemes are over the baseline LRT based single-user scheduling scheme. Note that for all the three scheduling schemes we employed the second phase described in Section IV. As seen from Table IV, MU scheduling in conjunction with an advanced SIC receiver at the BS can result in very significant gains in terms of cell average throughout (about 27%) along with good cell edge gains. For the simpler MMSE receiver, we see significant cell average throughout gains (about 18%) but a degraded cell edge performance. We note that it is possible to tradeoff a small fraction of the cell average gains for a large cell edge performance improvement by altering the PF rule.

VII. Conclusions

We considered resource allocation in the 3GPP LTE cellular uplink which allows for transmit antenna selection for each scheduled user as well as multi-user scheduling, wherein multiple users can be assigned the same time-frequency resource. We showed that the resulting resource allocation problem, which must comply with several practical constraints, is NP-hard. We then proposed constant-factor polynomial-time approximation algorithms and demonstrated their performance via simulations.

References

[1] 3GPP, "TSG-RAN EUTRA, rel.8," *TR 36.101*, Jun 2011.

[2] W. Yu and W. Rhee, "Degrees of freedom in wireless multiuser spatial multiplex systems with multiple antennas," *IEEE Trans. Commun.*, vol. 54, pp. 1747–1753, Oct 2006.

[3] D. Tse and S. Hanly, "Multiaccess fading channels-part I: Polymatroid structure, optimal resource allocation, and throughput capacities," *IEEE Trans. Inform. Theory*, 1998.

[4] W. Yu and J. Cioffi, "Constant power water-filling: Performance bound and low-complexity implementation," *IEEE Trans. Commun.*, vol. 54, pp. 23–28, Jan. 2006.

[5] A. Barnoy, R. Bar-Yehuda, A. Freund, J. Naor, and B. Schieber, "A unified approach to approximating resource allocation and scheduling," *ACM Symposium on Theory of Computing*, 2000.

[6] H. Yang, F. Ren, C. Lin, and J. Zhang, "Frequency-domain packet scheduling for 3GPP LTE uplink," *IEEE Infocom*, 2010.

[7] W. Yu and R. Liu, "Dual methods for nonconvex spectrum optimization of multicarrier systems," *IEEE Trans. Commun.*, vol. 54, pp. 1310–1322, July 2006.

[8] W. Yu, R. Liu, and R. Cendrillon, "Dual optimization methods for multiuser orthogonal frequency division multiplex systems," *Proc. IEEE Global Telecommun. Conf. (Globecom)*, 2004.

[9] C. Wong, R. Cheng, K. Letaief, and R. Murch, "Multiuser OFDM with adaptive subcarrier, bit, and power allocation," *IEEE J. Select. Areas Commun.*, vol. 17, pp. 1747 – 1758, Oct. 1999.

[10] S. Lee, S. Choudhury, A. Khoshnevis, S. Xu, and S. Lu, "Downlink MIMO with frequency-domain packet scheduling for 3GPP LTE," *IEEE Infocom*, 2009.

[11] A. Abrardo, P. Detti, and M. Moretti, "Message passing resource allocation for the uplink of multicarrier systems," *Proc. IEEE Int. Conf. on Commun. (ICC)*, 2009.

[12] W. Noh, "A distributed resource control for fairness in ofdma systems: English-auction game with imperfect information," *Proc. IEEE Global Telecommun. Conf. (Globecom)*, 2008.

[13] K. Kim and Y. Han, "Joint subcarrier and power allocation in uplink OFDMA systems," *IEEE Commun. Let.*, vol. 9, pp. 526–528, Jan. 2006.

[14] H. Zhang, N. Prasad, and S. Rangarajan, "MIMO downlink scheduling in LTE and LTE-advanced sustems," in *Proc. 2012 IEEE INFOCOM Miniconference*, (Orlando, FL), Mar. 2012.

[15] M. Andrews and L. Zhang, "Multiserver scheduling with contiguity constraints," *Proc. IEEE Infocom*, 2009.

[16] N. Prasad, H. Zhang, M. Jiang, G. Yue, and S. Rangarajan, "Resource allocation in 4G MIMO cellular uplink," *IEEE Globecom*, 2011.

[17] Y. Liu and E. Knightly, "Opportunistic fair scheduling over multiple wireless channels," in *Proc. 2003 IEEE INFOCOM*, (San Francisco, CA), Mar. 2003.

[18] N. Prasad, H. Zhang, H. Zhu, and S. Rangarajan, "Multi-user scheduling in the 3GPP LTE cellular uplink," *Tech. Report: uploaded in arXiv cs.NI*, Jan. 2012.

Table I
Algorithm I: Algorithm for LTE UL MU-MIMO

Input $p(\mathcal{U},\mathbf{c}),\ \ \forall\,(\mathcal{U},\mathbf{c})\in\boldsymbol{\mathcal{M}}$ and $\boldsymbol{\mathcal{M}}^{\text{narrow}},\boldsymbol{\mathcal{M}}^{\text{wide}}$
Determine a feasible allocation over $\boldsymbol{\mathcal{M}}^{\text{narrow}}$ using Algorithm IIa and let $\hat{W}^{\text{narrow}}$ denote the corresponding weighted sum rate.
Determine a feasible allocation over $\boldsymbol{\mathcal{M}}^{\text{wide}}$ using Algorithm IIb and let $\hat{W}^{\text{wide}}$ denote the corresponding weighted sum rate.
Select and output the allocation resulting in $\hat{W}=\max\{\hat{W}^{\text{narrow}},\hat{W}^{\text{wide}}\}$.

Table II
Algorithm IIa: LRT based module $\boldsymbol{\mathcal{M}}^{\text{narrow}}$

Initialize $p'(\mathcal{U},\mathbf{c})\leftarrow p(\mathcal{U},\mathbf{c}),\ \forall\ (\mathcal{U},\mathbf{c})\in\boldsymbol{\mathcal{M}}^{\text{narrow}}$, stack $\mathcal{S}=\phi$
For $j=1,\cdots,N$
Determine $(\mathcal{U}^*,\mathbf{c}^*)=\arg\max_{\substack{(\mathcal{U},\mathbf{c})\in\boldsymbol{\mathcal{M}}^{\text{narrow}}\\ \text{Tail}(\mathbf{c})=j}}\ p'(\mathcal{U},\mathbf{c})$
If $p'(\mathcal{U}^*,\mathbf{c}^*)>0$ **Then**
Set $\hat{p}=p'(\mathcal{U}^*,\mathbf{c}^*)$ and Push $(\mathcal{U}^*,\mathbf{c}^*)$ into $\mathcal{S}$.
For each $(\mathcal{U},\mathbf{c})\in\boldsymbol{\mathcal{M}}^{\text{narrow}}$ such that $p'(\mathcal{U},\mathbf{c})>0$
If $\exists\,\mathcal{G}_s:\mathcal{U}\cap\mathcal{G}_s\neq\phi\ \&\ \mathcal{U}_j^*\cap\mathcal{G}_s\neq\phi$ or $\mathbf{c}^*\cap\mathbf{c}\neq\phi$ **Then**
Update $p'(\mathcal{U},\mathbf{c})\leftarrow p'(\mathcal{U},\mathbf{c})-\hat{p}$
Else If $\exists\,q\in\mathcal{I}:\alpha^q(\mathcal{U},\mathbf{c})=\alpha^q(\mathcal{U}_j^*,\mathbf{c}_j^*)=1$ **Then**
Update $p'(\mathcal{U},\mathbf{c})\leftarrow p'(\mathcal{U},\mathbf{c})-\hat{p}$
Else
Update $p'(\mathcal{U},\mathbf{c})\leftarrow p'(\mathcal{U},\mathbf{c})-2\hat{p}\max_{1\leq q\leq J}\beta^q(\mathcal{U},\mathbf{c})$.
End If
End For
End If
End For
Set stack $\mathcal{S}'=\phi$
While $\mathcal{S}\neq\phi$
Obtain $(\mathcal{U},\mathbf{c})=\text{Pop}\ \mathcal{S}$
If $(\mathcal{U},\mathbf{c})\cup\mathcal{S}'$ is valid **Then** %% *$(\mathcal{U},\mathbf{c})\cup\mathcal{S}'$ is deemed valid if no user in $\mathcal{U}$ is incompatible with any user present in $\mathcal{S}'$ and no chunk in $\mathcal{S}'$ has an overlap with $\mathbf{c}$ and all knapsack constraints are satisfied by $(\mathcal{U},\mathbf{c})\cup\mathcal{S}'$.*
Update $\mathcal{S}'\leftarrow(\mathcal{U},\mathbf{c})\cup\mathcal{S}'$
End While
Output $\mathcal{S}'$ and $\hat{W}^{\text{narrow}}=\sum_{(\mathcal{U},\mathbf{c})\in\mathcal{S}'}p(\mathcal{U},\mathbf{c})$.

Table III
Algorithm IIb: Greedy module over $\boldsymbol{\mathcal{M}}^{\text{wide}}$

Input $p(\mathcal{U},\mathbf{c}),\ \ \forall\,(\mathcal{U},\mathbf{c})\in\boldsymbol{\mathcal{M}}^{\text{wide}}$ and $\{\boldsymbol{\mathcal{V}}^{(q)}\}_{q=1}^{J}$.
Set $\mathcal{S}=\phi$ and $\boldsymbol{\mathcal{M}}'=\boldsymbol{\mathcal{M}}^{\text{wide}}$.
Repeat
Determine $(\mathcal{U}^*,\mathbf{c}^*)=\arg\max_{\substack{(\mathcal{U},\mathbf{c})\in\boldsymbol{\mathcal{M}}'\\ \mathcal{S}\cup(\mathcal{U},\mathbf{c})\text{ is valid}}}\ p(\mathcal{U},\mathbf{c})$.
Update $\mathcal{S}\leftarrow\mathcal{S}\cup(\mathcal{U}^*,\mathbf{c}^*)$ and $\boldsymbol{\mathcal{M}}'=\boldsymbol{\mathcal{M}}'\setminus\{\boldsymbol{\mathcal{V}}^{(q)}:(\mathcal{U}^*,\mathbf{c}^*)\in\boldsymbol{\mathcal{V}}^{(q)}\}$
Until $(\mathcal{U}^*,\mathbf{c}^*)=\phi$ or $\boldsymbol{\mathcal{M}}'=\phi$.
Output $\mathcal{S}$ and $\hat{W}^{\text{wide}}=\sum_{(\mathcal{U},\mathbf{c})\in\mathcal{S}}p(\mathcal{U},\mathbf{c})$.

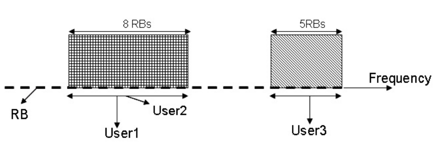

Figure 1. A Feasible RB Allocation in the LTE UL

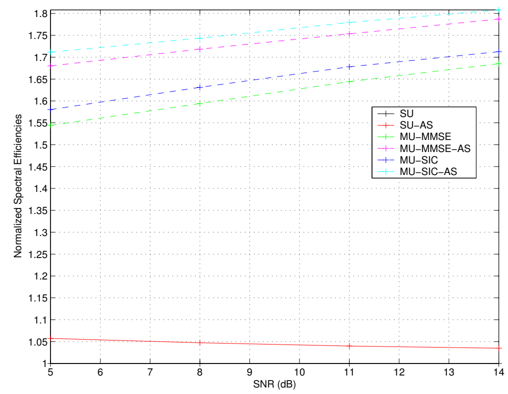

Figure 2. Normalized spectral efficiency versus SNR (dB)

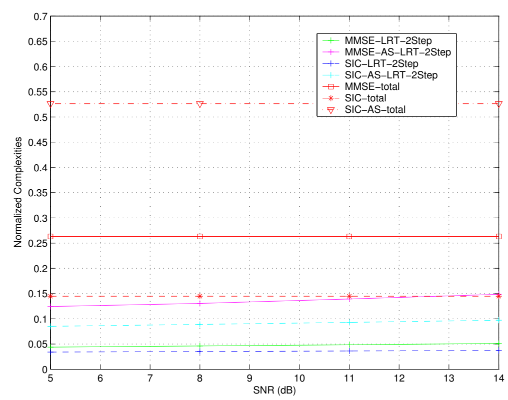

Figure 3. Normalized complexity versus SNR (dB)

Scheduling method	cell average	5% cell-edge
LRT SU	1.6214	0.0655
LRT MU with MMSE	1.9246 (18.70%)	0.0524
LRT MU with SIC	2.0651 (27.37%)	0.0745

Table IV
SPECTRAL EFFICIENCY OF LRT BASED SU AND MU UL SCHEDULING SCHEMES. AN AVERAGE OF 10 USERS ARE PRESENT IN EACH CELL AND ALL ASSOCIATED ACTIVE USERS CAN BE SCHEDULED IN EACH INTERVAL.