LTE uplink scheduling - flow level analysis

D.C. Dimitrova[1⋆], J.L. van den Berg[1,2], G. Heijenk[1], and R. Litjens[2]

[1] University of Twente, Enschede, The Netherlands
[2] TNO ICT, Delft, The Netherlands

Abstract. Long Term Evolution (LTE) is a cellular technology developed to support data traffic at potentially high rates. It is foreseen to extend the capacity and improve the performance of current 3G cellular networks. A key mechanism in the LTE traffic handling is the packet scheduler, which is in charge of allocating resources to active flows in both the frequency and time dimension. In this paper we present a performance comparison of three distinct scheduling schemes for LTE uplink with main focus on the impact of flow-level dynamics resulting from the random user behaviour. We apply a combined analytical/simulation approach which enables fast evaluation of flow-level performance measures. The results show that by considering flow-level dynamics we are able to observe performance trends that would otherwise stay hidden if only packet-level analysis is performed.

1 Introduction

Currently mobile operators are making a shift towards the UTRA Long Term Evolution (LTE) with Orthogonal Frequency Division Multiple Access (OFDMA) as the core access technology. One of the key mechanisms for realising the potential efficiency of this technology is the packet scheduler, which coordinates the access to the shared channel resources. In OFDMA-based LTE systems this coordination refers to both the time dimension (allocation of time frames) and the frequency dimension (allocation of subcarriers). These two grades of freedom, together with specific system constraints, make scheduling in LTE a challenging optimization problem.

The LTE uplink scheduling problem can in general be formulated as a utility optimization problem, see e.g. [7]. The complexity of this problem depends on many aspects among which the considered utility function (mostly aggregated throughput maximisation), fairness requirements and specific system characteristics (e.g. regarding fast fading, multiple antennas), see [10, 12, 13]. As often the optimal solutions would be too complex for practical implementation, the proposed scheduling algorithms tend to be based on heuristics yielding reasonable system performance under practical circumstances, see e.g. [2, 15].

In the present paper we take different approach towards the scheduling problem in the uplink of LTE. Instead of searching for an optimal solution, we make

⋆ Corresponding author, d.c.dimitrova@ewi.utwente.nl, P.O. Box 217, 7500 AE Enschede, The Netherlands.

a step back towards the basis of scheduling. We take two basic types of schedulers - fair access and greedy access - and demonstrate the importance of taking into account the user behaviour for the performance of an LTE uplink. The initiations of finite sized flow transfers, occurring at random time instants and locations, is what we term *flow dynamics* and leads to a time - varying number of ongoing flow transfers. We consider this phenomenon very important for the final performance of both the overall system and the individual end users. Such combined analysis of LTE uplink with user behaviour is not commonly found in the literature.

On the one hand, most research on LTE scheduling has been treating the downlink scenario, some examples being [11, 14]. Considerable less work has been dedicated to the uplink, where the transmit power constraint of the mobile equipment plays an important role. On the other hand, studies that take into account flow dynamics are lacking. Most papers consider the performance of newly proposed scheduling schemes for scenarios with a fixed number of active users in the system. For example, the authors of [8] define a theoretically optimal scheduling scheme and its suboptimal, practically feasible counterpart. Unfortunately, [8] fails to present the flow dynamics in a real system and does not consider the impact of a user's location on performance. This factor is accounted for by [15], which proposes three channel-aware scheduling algorithms. The study shows the realised cell throughput for different number of users but does not specify whether the user population changes dynamically or is preset. Additionally, evaluating cell throughput does not give much insight on the performance of a single user.

In contrast to [8] and [15] we focus on the impact that flow behaviour has on the performance observed by the users while also accounting for the users' location in the cell. Our investigations are done for two classes of schedulers - a resource fair class, where the active users are scheduled in a Round Robin fashion and are all assigned an equal number of subcarriers, and a greedy class, which aims to maximise system capacity (best performance reference). Intermediate results for the resource fair class were published in [6]. The current study extends [6] by introducing a greedy class of scheduler and deployment limitations on the performance; a more detailed discussion of the study can be found in [4].

Our modelling and analysis approach is based on a time-scale decomposition and resembles, at high level, the approach we used previously in the context of UMTS/EUL, see [5]. The approach combines a packet-level analysis, which captures details of the scheduler and the propagation environment, and simulation, which models flow dynamics. In particular, we use continuous-time Markov chains to represent the change in number of active users. Depending on the scheduler's complexity the steady-state distribution of the Markov chain can be found either analytically (yielding insightful closed-form expressions) or by simulation. In this study, the analytical approach applies for some special cases of our resource fair schemes.

The rest of the paper is organised as follows. Section 2 provides a general introduction to LTE uplink scheduling and introduces the considered scheduling schemes. In Section 3 the network model is described. Subsequently, in Section 4 we present the performance evaluation approach and in Section 5 the numerical results at flow level. Finally, Section 6 concludes the presented work.

2 Scheduling

In this section we first give a general introduction to scheduling in the LTE uplink, necessary for understanding the proposed schemes. Subsequently, the proposed scheduling schemes are described.

2.1 LTE uplink scheduling

The radio access technology chosen for the LTE uplink is SC-FDMA (Single Carrier - Frequency Division Multiple Access), in which the radio spectrum is divided into nearly perfect mutually orthogonal sub-carriers. Hence, simultaneous transmissions from different mobile stations (MSs) do not cause intra-cell interference but they do compete for a share in the set of sub-carriers. The total bandwidth that can be allocated to a single MS depends on the resource availability, the radio link quality and the terminal's transmit power budget. The scheduling decision is taken by the packet scheduler, which is located at the base station (BS). Allocation of multiple sub-carries to the same user is possible as long as these sub-carriers are consecutive in the frequency domain. A key feature of packet scheduling in LTE networks is the possibility to schedule users in two dimensions, viz. in time and frequency. The aggregate bandwidth available for resource management is divided in sub-carriers of 15 kHz. Twelve consecutive sub-carriers are grouped to form what we term a 'sub-channel', with a bandwidth of 180 kHz; there are M sub-channels in the system bandwidth. In the time dimension, the access to the sub-channels is organised in time slots of 0.5 ms. Two slots of 0.5 ms form a TTI (Transmission Time Interval). The smallest scheduling unit in LTE is termed a *resource block* (RB) and has dimension of 180 kHz and 1 ms[3]).

The data rate that a user can realise is influenced by the number of RBs assigned to it by the scheduler, which determines the allocated bandwidth and the applied transmit power, and by its location, which determines the path loss and the *signal to interference plus noise ratio* (SINR). Some studies, e.g. [9], also argue that certain system characteristics such as the available bandwidth for signalling affect the performance. We investigate the issue in Section 5.2.

2.2 Scheduling schemes

In our analysis we focus on two (types of) *resource fair* scheduling schemes, which assign equal resource shares to all active users, independently of their respective channel conditions - *fair fixed assignment* (FFA) and *fair work-conserving*

[3] In fact, each scheduling entity of 180 kHz and 1 ms consists of two RBs, i.e. a RB has the duration of 0.5ms. In this study we use the term RB to refer to a 1 ms interval.

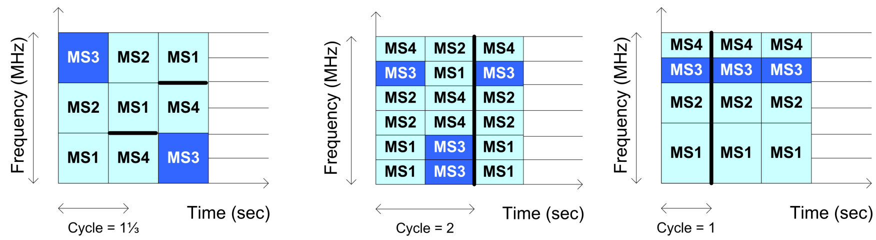

(a) FFA scheme (b) FWC scheme (c) MAV scheme

Fig. 1. Considered scheduling schemes for LTE uplink; examples with four users.

(FWC). Furthermore, we consider a greedy scheduling scheme - *maximum added value* (MAV) - as a reference for a strategy that aims at maximum system throughput, given the channel conditions of the active users.

The first scheduler is termed *fair fixed assignment* because it assigns the same, a priori specified, number of resource blocks to each mobile station (see Figure 1(a)). The number of assigned resource blocks per MS, denoted m is hence the same for each mobile station, independently from its location, and is an operator-specified parameter. If the number n of active users is such that the total number of requested resource blocks is less than the available number of resource blocks per TTI, i.e. if $n \cdot m < M$, then a number of resource blocks are left idle. Naturally this reflects a certain degree of resource inefficiency in the scheme, especially for situations with low traffic load and hence few active users. When the number of active users is such that $n \cdot m > M$, more than a single TTI is needed to serve all users at least once. We define a *cycle length* c as the number of TTIs necessary to serve all users at least once, see Figure 1(a). This cycle length can be expressed as $c = \max(1, n/M)$. According to this definition c is not necessarily integral (but at least one) and the start of a given cycle may fall within the same TTI as the end of the previous cycle.

The second scheme, the *fair work-conserving* scheme, aims to avoid the resource inefficiencies of the FFA scheme under low traffic loads, while still preserving the resource fairness property. The scheme's objective is to distribute the available resource blocks evenly over the active users within each individual TTI, see Figure 1(b). As a result the FWC scheduler is optimal in the class of resource-fair Round Robin schedulers. In principle each of the n MSs is assigned M/n resource blocks in each TTI. Since M/n needs not be integral, in an implementable version of the FWC scheduler, a scheduling cycle is defined of multiple TTIs during which user-specific resource block assignments appropriately vary between $\lfloor M/n \rfloor$ (low allocation) and $\lceil M/n \rceil$ (high allocation) in order to, on average, achieve the fair assignment of M/n resource blocks. More specifically, the cycle length is equal to the smallest integer c such that $c \cdot M/n$ is integral, which is at most equal to n.

Finally, the *maximum added value* (MAV) scheme has as main objective to maximise the total data rate realised given the active users present in the system. The scheme assigns RBs to those users that can make best use of it. In particular, scheduling decisions are based on a metric termed *added value*, which, for a particular user, is the gain in data rate that a new resource block can deliver to that user. Of all active users the one with the highest added value is assigned the resource block, see Figure 1(c). This procedure continues until all resource blocks have been assigned thus resulting in cycle length $c = 1$. In the MAV scheduling it is possible that cell edge users are deprived from service if the system is under high load - since other users can make better use of the available resource blocks cell edge users get none.

3 Model

We consider a single cell divided in K circular zones of equal area in order to differentiate between users' distances to the base station. Each zone is characterized by a distance d_i to the base station, measured from the outer edge of the zone. Mobile users are uniformly distributed over the zones and flow arrivals follow a Poisson process with rate λ. The arrival rate per zone $\lambda_i = \lambda/K$ is equal for all zones (due to equal area), where $i = 1, ..., K$. The distribution of the active users over the zones we term *state* $\underline{n} = (n_1, n_2, ..., n_K)$ where n_i defines the number of uses in zone i. Note that $n = \sum_{i=1}^{K} n_i$. All mobile stations are assumed to have the same maximum transmit power P_{max}^{tx}. This maximum power is equally distributed over the assigned RBs m_i, see [15], leading to transmit power per RB $P_i^{tx} = P_{max}^{tx}/m_i$. Note that for the FFA and FWC schemes m_i is the same for all zones but it differs in the case of the MAV scheduler. Each zone is characterized by a distinct path loss $L(d_i)$ defined by the Cost 231 Hata propagation model (given in dB), according to which

$$L(d_i) = L_{fix} + 10a \log_{10}(d_i), \tag{1}$$

where L_{fix} is a parameter that depends on system characteristics such as antenna height and a is the path loss exponent. In the rest of the analysis linear scale is used for $L(d_i)$. Users belonging to the same zone i have the same distance d_i and hence experience the same path loss. At this stage of the research we consider only thermal noise N from the components at the base station.Note that intra-cell interference in LTE is not an issue due to the orthogonality of the sub-carriers in LTE. We assume that the RBs used within the cell are not reused in neighbour cells, i.e. inter-cell interference is of no concern. Given a known path loss, the received power (per zone) at the base station P_i^{rx} can be expressed as

$$P_i^{rx} = \frac{P_i^{tx}}{L(d_i)}, \tag{2}$$

Eventually, for the signal-to-interference-plus-noise ratio measured at eNodeB from user of zone i we can write:

$$SINR_i = \frac{P_i^{rx}}{N} = \frac{P_{max}^{tx}/m_i}{L(d_i)N}. \tag{3}$$

Note that the SINR is lower bound to a minimum target level $SINR_{min}$, required for successful reception, and is upper bound by the highest supported modulation and coding scheme (MCS). In our model we work with 16QAM since it should be supported by all terminal classes.

4 Analysis

Our proposed evaluation approach consists basically of three steps. The first two steps take into account the details of the scheduler's behaviour, e.g. allocation of subcarriers, and the given state of the system, i.e. the number of active users and their distance to the base station. In the third step we create a continuous-time Markov chain, which describes the system behaviour at flow level. From the steady-state distribution of the Markov chain the performance measures, such as mean file transfer time T_i of a user in zone i, can be calculated.

4.1 Packet-level analysis

At the packet-level of the analysis approach we define the performance measure *instantaneous rate* r_i. It is the data rate realised by a user (from zone i) when it is scheduled and it is determined by the SINR as derived above, the possible modulation and coding schemes and the receiver characteristics related to that MCS. The instantaneous rate is calculated over all RBs that are allocated to a particular user. In our analysis we use the Shannon formula modified with a parameter σ to represent the limitations of implementation, see Annex A in [1]. Hence, for the instantaneous rate we can write:

$$r_i = (m_i \cdot 180\text{kHz})\sigma \log_2(1 + SINR_i), \tag{4}$$

where $m_i \cdot 180$ kHz is the bandwidth allocated to a user in a zone i. Note that both $SINR_i$ and r_i are calculated over the same RB allocation.

In the FFA scheme (with a fixed number of RB allocation per user in a cycle) the instantaneous rate of a particular MS is always the same when the MSs is served. In the case of the FWC and MAV schemes however the instantaneous rate depends on the total number of users in the system. Furthermore, for the FWC r_i depends on whether low or high allocation occurs in the specific TTI, see Section 2.2, and hence for the FWC scheme we calculate two instantaneous rates $r_{i,L}$ and $r_{i,H}$ respectively.

4.2 Flow-level analysis

The flow-level behaviour can be modelled by a K-dimensional Markov chain with state space $\underline{n} = (n_1, n_2, ..., n_K)$, $n_i \geq 0$ and $i = 1, ..., K$. The jumps in the Markov chain represent the initiation and completion of flow transfers. The corresponding transition rates in a particular state are determined from the (a-priori) given arrival rates λ_i and the long-term flow throughputs $R_i(\underline{\mathrm{n}})$ in that state. These throughputs can be derived from the instantaneous rates, see Equation (4), and from the cycle length. For the FFA scheduler the state-dependent throughput can be easily expressed as $R_i(\underline{\mathrm{n}}) = r_i/c$. The MAV scheduler has by definition a cycle length of a single TTI and thus $R_i(\underline{\mathrm{n}}) = r_i$. For the FWC scheme we need to consider the variation in low resource block allocation ($\lfloor M/n \rfloor$ blocks) and high resource block allocation ($\lceil M/n \rceil$ blocks). Each allocation applies for a fraction a_L and a_H, respectively, of the scheduling cycle as follow:

$$\text{Low allocation}: a_L = \left\lceil \frac{M}{n} \right\rceil - \frac{M}{n}, \tag{5}$$

$$\text{High allocation}: a_H = \frac{M}{n} - \left\lfloor \frac{M}{n} \right\rfloor. \tag{6}$$

Eventually for the state dependent throughput we can write for the FWC scheme:

$$R_i(\underline{\mathrm{n}}) = a_L r_{i,L} + a_H r_{i,H}. \tag{7}$$

The eventual transition rates in the Markov chain is given by $n_i R_i(\underline{\mathrm{n}})/F$, where F is the mean flow size, and are scheduler specific.

The steady-state distribution of the Markov chain can be found either by simulating the (state transitions of) the Markov chain or, in special cases, by analytical approaches leading to explicit closed-form expressions. In particular, in our study the model of the FFA scheduler appeared to be similar to a M/M/1 processor sharing (PS) queuing model with multiple classes of customers and state dependent service rates. We will further discuss this below. The Markov chains for the FWC and the MAV scheduler are of more complex form and not trivial to solve, which is why we selected a simulation approach for these cases.

Explicit solution for the FFA scheme We argue that the Markov chain of the FFA scheduler is similar to the Markov chain describing the behaviour of a M/M/1 PS queuing model with multiple classes of customers and state-dependent service rates. This queuing model is described and analysed in [3], Section 7. In [3] each 'task', given there are k active tasks, receives a service portion $f(k)$. The Markov chain of the FFA scheduler turns out to be the same as the Markov chain of the M/M/1 PS model. In particular, it is recognised that the cycle length of the FFA scheme, which depends on the number of active users n, actually determines the service portions $f(\cdot)$. Using the expression for c, given in Section 2.2, we have for the FFA scheme:

$$f(n) = \begin{cases} 1 & \text{for } n = 1, ..., L, \\ \frac{L}{n} & \text{for } n > L, \end{cases} \tag{8}$$

where $L = \lfloor M/m \rfloor$, i.e. the maximum number of MSs that can be served in a TTI. The two situations in Equation (8) occur due to the limited number of RBs per TTI.

Given the above relationship between the Markov chains, we can write for the mean flow transfer time T_i, see [3], of zone i:

$$T_i = \tau_i \frac{\sum_{j=0}^{\infty} \frac{\rho^j}{j!} \Phi(j+1)}{\sum_{j=0}^{\infty} \frac{\rho^j}{j!} \Phi(j)}, \tag{9}$$

where $\Phi(n) = \left(\prod_{j=1}^{n} f(j)\right)^{-1}$ and $\tau_i = F/r_i$ represents the average service requirement of mobile station in zone i; ρ is the system load defined as $\rho = \sum_{i=1}^{K} \lambda_i F/r_i$. Substituting $f(n)$ in (9) we get:

$$T_i = \frac{\tau_i * A}{\sum_{j=0}^{L} \frac{\rho^j}{j!} + \frac{L^L}{L!}(\rho/L)^{L+1} \frac{1}{1-\rho/L}}, i = 1, ..., K, \tag{10}$$

with

$$A = \sum_{j=0}^{L-1} \frac{\rho^j}{j!} + \frac{L^L}{L!\rho}\left((\frac{\rho}{L})^{L+1} \frac{L}{1-\rho/L} + (\frac{\rho}{L})^{L+1} \frac{1}{(1-\rho/L)^2} \right).$$

Note that the impact of the distance of each zone is taken in the specific flow size τ_i, expressed in time.

5 Numerical results

In Sections 5.2 and 5.3 we present a quantitative evaluation of the three LTE uplink schedulers introduced in Section 2.2. Beforehand, in Section 5.1 we present the parameter settings.

5.1 Parameter settings

The cell with cell radius of 1km is divided in ten zones, i.e. $K = 10$. A system of 10 MHz bandwidth is studied, which, given that a RB has 180 kHz bandwidth (and including control overhead), results in 50 RBs available per TTI.

Mobile stations have maximum transmit power $P_{\max}^{tx} = 0.2$ Watt. The lower bound on the SINR is -10dB while the upper bound on performance is determined by a 16QAM modulation that corresponds to SINR of 15dB. For the path loss we have used $PL_{fix} = 141.6$ and path loss exponent of $a = 3.53$, height of the mobile station 1.5m, height of the eNodeB antenna 30m and system frequency 2.6GHz. The thermal noise per subcarrier (180kHz) is -121.45dBm and with noise figure of 5dB the effective noise level per resource block is $N = -116.45dBm$. The attenuation of implementation σ is taken at 0.4, see [1] and Equation (4). The average file size F is 1Mbit and the rate λ at which users become active changes depending on the discussed scenario.

5.2 Fair allocation schedulers

In this section we compare the performance of the two fair allocation strategies FFA and FWC. First, the impact of the scheduling policy in combination with the RB allocation strategy is investigated. Second, we evaluate the impact of certain practical limitations on performance.

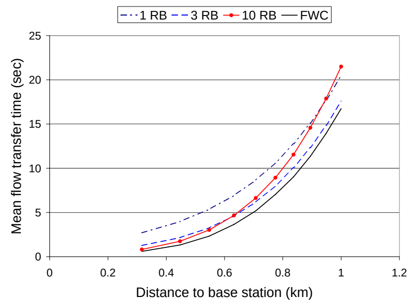

Fig. 2. Performance evaluation of the scheduling strategy, given a specific arrival rate.

Impact of scheduling policy In this subsection we investigate how different features of the scheduler, namely number of assigned RBs and user selection for service, affect the flow-level performance. Evaluation is performed for arrival rate of $\lambda = 0.5$ flows/sec and number of assigned RBs in the FFA scheme changing from one to three to ten[4].

The impact of the number of assigned RBs on performance is observed for the FFA scheme in Figure 2. The general trend is that increase in allocation translates to lower mean flow transfer times, e.g. $m = 1$ vs. $m = 3$. However, for remote MSs high allocation worsens performance, e.g. $m = 3$ vs. $m = 10$. While close-by MSs have sufficient power capacity to reach $SINR_{min}$ for all of its allocated RBs remote users lack this ability (due to high path loss). They actually use fewer RBs in order to guarantee $SINR_{min}$.

The impact of the scheduling policy is investigated by comparing the results for FFA with $m = 1$ and the FWC scheme, see Figure 2. Our choice is motivated by the similar system capacity of the two schemes. The FWC scheme visibly outperforms the FFA with $m = 1$ due to its more efficient distribution of RBs over the active users. Recall that the FWC schemes keeps on scheduling users in the same TTI even if all users have been served once. In the FFA scheme however once each user is assigned a resource block in the TTI the scheduling stops thus leaving RBs unused.

Impact of PDCCH limitations The potential performance gains that the freedom to schedule in the frequency dimension brings could be diminished by practical limitations. One such limitation, according to [9], is the radio resource set up for the Physical Downlink Control CHannel (PDCCH). PDCCH carries information from the base station towards a mobile (control packet) about its scheduled uplink transmission. The resource per TTI set aside for these control messages is limited and therefore, given a fixed size of the control package, only a limited number of uplink transmissions can be served within a TTI. The

[4] These showed to be the most interesting assignments within the range one to ten RBs with a step of one.

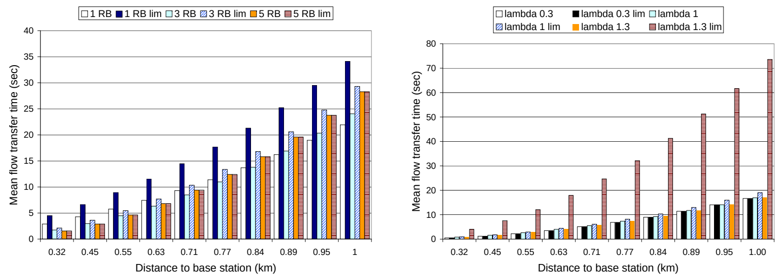

(a) FFA scheme (b) FWC scheme

Fig. 3. Impact of PDCCH limitation on performance of (a) a FFA scheme and (b) a FWC scheme.

limit proposed by [9] is eight to ten users; in our evaluation we have chosen for ten users. In the rest of the discussion we will use the abbreviations FFA-lim and FWC-lim to refer to the PDCCH limited versions of the two resource fair schemes.

Each of the access fair schemes we discuss, is affected differently by the PDCCH limitation. The FFA scheme will simply serve at the most ten users in a TTI according to the chosen RB allocation policy. Any unused resource blocks will be therefore waisted, which suggests negative impact on performance. Figure 3(a) shows that, independently of the RB allocation policy, i.e. $m = 1, 3$ or 5, the FFA-lim scheme performs worse (or equal) than the original FFA scheme. The difference is biggest for $m = 1$ and non-existent for $m = 5$. The later observation can be explained by the fact that with $m = 5$ or more the maximum number of users that can be served in one TTI is ten or less. Hence the RB allocation policy self limits the number of simultaneous transmissions in one TTI to ten (or less).

For the FWC scheme the limitation of up to ten users implies that each user will get assigned at least five RBs. Although the mobile stations can fully utilise all allocated RBs remote users cannot reach the maximum possible data rate (highest modulation scheme). Hence, at high loads with FWC-lim MSs use the available RBs much less efficiently than in the original FWC scheme. At low load the two implementations behave practically the same. Based on the results for both FFA-lim and FWC-lim we can conclude that the effects of the limitation are scheduler specific but generally leads to increased mean flow transfer times.

5.3 MAV greedy scheduling

The performance of the two access fair (FFA with $m = 1$ and FWC) schemes and the greedy scheduler (MAV) for an arrival rate $\lambda =$ is presented in Figure 4(a). The MAV scheme outperforms the other two for close-by users but its performance is worse for remote users. The reason is the preference of MAV to

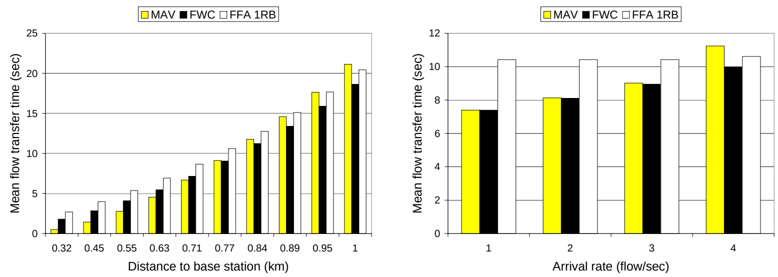

(a) Single arrival rate (b) Range of arrival rates

Fig. 4. Mean flow transfer times for the three schedulers. (a) Impact of user's location for high load, i.e. $\lambda = 3$; and (b) impact of flow arrival rate on the overall flow transfer time.

schedule users that can make best use of the available RBs, i.e. close-by users with low path loss.

Figure 4(b) presents the average mean flow transfer time over all zones for various arrival rates. Surprisingly, the MAV scheme does not always have the best performance! This is explained as follows. Although MAV tries to maximise the total data rate in each state, such strategy is vulnerable to reaching states where the available resources cannot be efficiently used, i.e. states with mainly remote users. Situations as described above are more probable to occur for high load, which corresponds to the results presented in Figure 4(b). This hidden inefficiency of MAV also implies that for remote users it is more beneficial to schedule them frequently even if with few RBs per user.

The paradox is that the MAV scheme, which aims at optimising throughput on a per TTI basis, achieves a lower system capacity (overall throughout), when analysed at the flow level.

6 Conclusion

In this paper we presented an investigation on the impact that flow dynamics (changing number of users) has on performance given the complex scheduling environment of LTE uplink. Two low complexity access fair scheduling schemes are examined - both designed to provide equal channel access. Additionally, as a reference base for optimal system performance, a greedy resource allocation scheme is considered. All schemes are evaluated by a hybrid analysis approach, which accounts for packet-level details such as scheduler's specifics as well as for the dynamic behaviour of users at flow level, i.e. flow initiations and completions. Due to its hybrid nature the approach allows fast evaluation while considering sufficient details of the investigated model. The most valuable contribution of our research is that considering flow dynamics reveals trends that would be otherwise left unobserved. An excellent example is our finding that the greedy scheduler although designed to maximise system throughput for a given number of users seems to be, contrary to expectations, less efficient than the access fair schemes in the long term.

References

1. 3GPP TS 36.942. LTE; evolved universal terrestrial radio access (E-UTRA); radio frequency (RF) system scenarios.
2. M. Al-Rawi, R. Jntti, J. Torsner, and M. Sagfors. On the performance of heuristic opportunistic scheduling in the uplink of 3G LTE networks. PIMRC 2008.
3. J. W. Cohen. The multiple phase service network with generalized processor sharing. volume 12, pages 254–284. Springer Berlin/Heidelberg, 1979.
4. D. C. Dimitrova. *Analysing uplink scheduling in mobile networks. A flow-level perspective*. 2010.
5. D.C. Dimitrova, J.L. van den Berg, G. Heijenk, and R. Litjens. Flow level performance comparison of packet scheduling schemes for UMTS EUL. WWIC '08, 2008.
6. D.C. Dimitrova, J.L. van den Berg, G. Heijenk, and R. Litjens. Scheduling strategies for LTE uplink with flow behaviour analysis. Proceedings 4th ERCIM, 2010.
7. L. Gao and S. Cui. Efficient subcarrier, power and rate allocation with fairness considerations for OFDMA uplink. volume Vol. 7, pages 1507–1511. IEEE Transactions on Wireless Communications, 2008.
8. A. M. El Hajj, E. Yaacoub, and Z. Dawy. On uplink OFDMA resource allocation with ergodic sum-rate maimization. PIMRC 2009, 2009.
9. H. Holma and A. Toskala. *LTE for UMTS, OFDMA and SC-FDMA Based Radio Access*. John Wiley & Sons, 2009.
10. J. Huang, V.G. Subramanian, R. Agrawal, and R. Berry. Joint scheduling and resource allocation in uplink OFDM systems for broadband wireless access networks. volume Vol. 27, pages 226–234. IEEE Journal on Selected Areas in Communications, 2009.
11. R. Kwan, C. Leung, and J. Zhang. Multiuser scheduling on the downlink of an LTE cellular system. *Rec. Lett. Commun.*, pages 1–4, 2008.
12. S. B. Lee, I. Pefkianakis, A. Meyerson, S. Xu, and S. Lu. Proportional fair frequency-domain packet scheduling for 3GPP LTE uplink. IEEE INFOCOM 2009 mini-symposium.
13. L.A. Maestro Ruiz de Temino, G. Berardinelli, S. Frattasi, and P. Mogensen. Channel-aware scheduling algorithms for SC-FDMA in LTE uplink. Proceedings PIMRC 2008.
14. C. Wengerter, J. Ohlhorst, and A. G. E. von Elbwart. Fairness and throughput analysis for generalized proportional fair frequency scheduling in OFDMA. Vehicular Technology Conference, 2005. VTC 2005-Spring.
15. E. Yaacoub, H. Al-Asadi, and Z. Dawy. Low complexity scheduling algorithms for the lte uplink. In *Computers and Communications, 2009. ISCC 2009. IEEE Symposium on*, pages 266 –270.