

Impact of inter-cell interference on flow level performance of scheduling schemes for the UMTS EUL

D. C. Dimitrova and G. Heijenk
University of Twente, The Netherlands
d.c.dimitrova@ewi.utwente.nl
geert.heijenk@utwente.nl

J.L. van den Berg
University of Twente,
TNO ICT, The Netherlands
J.L.vandenBerg@ewi.utwente.nl

R. Litjens
TNO ICT, The Netherlands
remco.litjens@tno.nl

Abstract—The UMTS Enhanced Uplink (EUL) is expected to provide higher capacity, increased data rates, and smaller latency on the communication link from users towards the network. A key mechanism in EUL traffic handling is the packet scheduler, for which a number of basic schemes can be identified (one-by-one, partial parallel, and full parallel). In this paper we analyze the interaction between the EUL scheduling scheme deployed in the network and the inter-cell interference. On the one hand, different scheduling schemes cause different inter-cell interference patterns on neighbouring cells. On the other hand, the different schemes are affected by inter-cell interference in different ways. The scheduling schemes are evaluated and compared under different approaches for reserving part of the allowed noise rise at the base station for inter-cell interference. For our analysis, we have developed a hybrid analytical/simulation approach allowing for fast evaluation of performance measures such as the mean flow transfer time and fairness expressing how the performance depends on the user's location. This approach takes into account both the packet-level characteristics and the flow-level dynamics due to the random user behaviour.

I. INTRODUCTION

UMTS (Universal Mobile Telecommunication System) networks are currently operational in many countries in the world. In order to fully exploit the available capacity and to provide increased data rates and reduced latencies to users, UMTS has been extended with the High-Speed Packet Access (HSPA) technology. For the downlink, from base station to mobile, the HSDPA (High-Speed Downlink Packet Access) technology has been standardized by 3GPP in their Release 5 [1] of the UMTS standard. HSDPA is currently used by many mobile operators. To provide an equivalent in the uplink, from mobile to base station, the Enhanced Uplink (EUL) has been specified in 3GPP Release 6 of the UMTS standard [2]. By using optimized scheduling schemes, the EUL should fully utilize the UMTS air interface capacity, exploiting the elasticity of the users' applications.

The enhanced uplink introduces a new transport channel called EDCH, see e.g. [3]. Channel access is coordinated by the base stations via packet scheduling based on time frames of fixed length (2 or 10 ms, termed TTI: Transmission Time Interval). Fast rate adaptation with an enhanced dynamic range and efficient time multiplexing through appropriate scheduling

schemes enable higher data transfer rates than usually provided on DCHs in 'plain' UMTS. Other key benefits offered by the EUL technology are an enhanced cell capacity and a reduced latency. In contrast to HSDPA for the downlink, due to limited transmit powers of the user terminals, a single uplink user cannot always use the total available channel resource on its own when it is scheduled (which would optimize throughput, cf. [4]) depending on its distance to the base station. Hence, it makes sense to consider scheduling schemes with simultaneous transmissions on the uplink, see e.g. [3].

In earlier work [5] we modelled and analyzed the flow-level performance of different scheduling schemes in a single-cell setting, not taking inter-cell interference into account. In the present paper, we will investigate the impact of inter-cell interference on the performance of different EUL scheduling schemes. On the one hand, different scheduling schemes cause different inter-cell interference patterns. On the other hand, the different schemes are affected by inter-cell interference in different ways. Next to comparing different scheduling schemes, we also identify and evaluate different approaches for reserving part of the allowed noise rise at the base station for inter-cell interference.

In our study, we take into account both the packet-level characteristics and the flow-level dynamics due to flow (file) transfer completions and initiations by the users at random time instants. In particular, we quantify the inter-cell interference from a neighbour cell on a reference cell, by modelling the flow-level behaviour of the neighbour cell as a continuous time Markov chain. For each state in this Markov chain we determine the detailed packet-level characteristics of the inter-cell interference process, which depend on the scheduling scheme used. Subsequently, we analyze the impact of the interference process from an independently behaving neighbour cell on the flow-level performance of the reference cell deploying the same scheduling scheme. This is done by extending the Markov chain describing the behaviour of the neighbour cell to a Markov chain modelling the behaviour of both cells simultaneously. The numerical evaluation of the resulting model is based on a hybrid analytical/simulation approach. In particular, the state transitions in the Markov

chain applying to the initiation or completion of flow transfers are simulated, while the packet-level behaviour in a particular state of the Markov chain visited during the simulation is analytically determined. This approach allows for fast evaluation of performance measures such as the mean flow transfer time and fairness, expressing how the performance depends on the user's location.

Most EUL performance studies in literature are based on dynamic system simulations, see e.g. [6], [7], [8]. The underlying simulation models incorporate many details of the channel operations and traffic behaviour, but running the simulations tends to require a lot of time. Most analytical studies focus on the performance of schedulers without taking into account the impact of the flow level dynamics, see e.g. [9]. Analytical studies on EUL performance capturing both the packet and flow level dynamics of the system are rare. Interesting references here are [10] and [11]. In particular, in [11] flow level performance metrics are analysed for two (rate-fair) scheduling disciplines assuming that the transmit powers of all mobiles are sufficient to reach the maximum bit rate, such that (in contrast to [5] and our present study) the users' performance does not depend on their location in the cell. In later research [12][13], the same authors consider a multi-cell scenario and choose for a log-normal approximation of the inter-cell interference process. However, this model does not capture the impact of the particular scheduling scheme used in the network.

The rest of the paper is organized as follows. Section II introduces the three different scheduling schemes and the different reservation schemes we will analyse in this paper. In Section III we describe the network scenario, state the modeling assumptions, and present the performance evaluation approach taken in this paper. Section IV presents and discusses numerical results illustrating the interaction between the scheduling scheme and the inter-cell interference, and its impact on the scheduler's performance. Finally, in Section V, conclusions and our plans for future work are given.

II. SCHEDULING SCHEMES FOR ENHANCED UPLINK

The presented analysis concentrates on a class of channel-oblivious schedulers which grant channel access to users irrespective of their channel conditions such that access fairness is supported. The considered scheduling schemes, denoted *one-by-one* (OBO), *partial parallel* (PP) and *full parallel* (FP), differ in their preference for granting consecutive or parallel channel access and in their capability to fully utilise available resources. The key radio resource in EUL cells is the noise rise target or, equivalently, the total received power budget (B) at the base station. The budget B is partially consumed by the (assumed fixed) thermal noise level N and some adaptively set reservation for the inter-cell interference B_{oc} . The remainder, denoted B' , is the intra-cell received power budget, which is distributed to the users in the cell by the packet scheduler.

The different packet schedulers are illustrated in Figure 1. In each TTI, the *OBO* scheduler assigns the full budget B' to a single user, serving all users having packets available for

transmission in round robin order [11]. The policy to limit the number of scheduled users to one, at any time, inherently achieves low intra-cell interference and hence higher instantaneous user data rates [4]. The drawback of the scheme is that a single user may not always be able to utilise the granted resource B' in full, depending on its maximum transmit power and the path loss to its serving base station. The *FP* scheduler [11] is in some sense the 'opposite' of the *OBO* scheduler, in that its primary objective is to fully utilise the available budget. It aims to do so by granting fractions of the budget to all users in proportion to their maximum attainable received power contribution (based on full power transmission). It is not hard to see that the *FP* scheduler improves resource utilisation, at the cost of lower instantaneous user data rates. The hybrid *PP* scheduler aims to achieve the best of both worlds. It effectively extends the *OBO* scheme by selecting (in each TTI) additional users for parallel transmission until the available resources are fully utilised. The users are hereby considered in round robin order, thus preserving the fair channel access for the different users. The scheduling cycle length c , i.e., the number of TTIs after which each user has been served once, under *OBO*, *FP*, and *PP* scheduling is equal to the number of users n , 1, and somewhere between 1 and n , respectively. A more detailed description of the three schedulers can be found in [5].

An important aspect of the presented analysis is to assess the schedulers' impact on the inter-cell interference characteristics. In a possible scheduling implementation, the fraction of the total received power budget B that is reserved for inter-cell interference, may be based on the observed inter-cell interference statistics from previous scheduling cycles. For example, some percentile value of the inter-cell interference distribution may be applied. In this light, it is important to note that *only* under *PP* scheduling, such a percentile is influenced by the precise scheduling order of the users. As illustrated by Figure 1, a user which is relatively close to the reference cell provides a much higher inter-cell interference level than a user at the other end of the neighbour cell. Hence it is readily seen that combining e.g. close-by users in a given TTI and far-away users in another TTI yields a highly variable inter-cell interference and hence requires substantial reservations. Note, however, that in practical implementations, the scheduler cannot base its scheduling order on the inter-cell interference impact, as it simply does not have this information at its disposal. For those reasons, we choose to consider two extreme versions of the *PP* scheduler which may provide us with bounds on the performance effects. For that purpose, all users are ranked according to their induced inter-cell interference levels (from high to low). In the '*worst*' version of *PP*, denoted *PP-w*, the applied scheduling order is simply from high to low (or vice versa, which makes no difference for our purposes). Alternatively, in the '*optimised*' version of *PP*, denoted *PP-o*, the induced inter-cell interference is smoothed by applying the order highest, lowest, second-highest, etc. to fill the available interference budget in consecutive TTIs of the scheduling cycle.

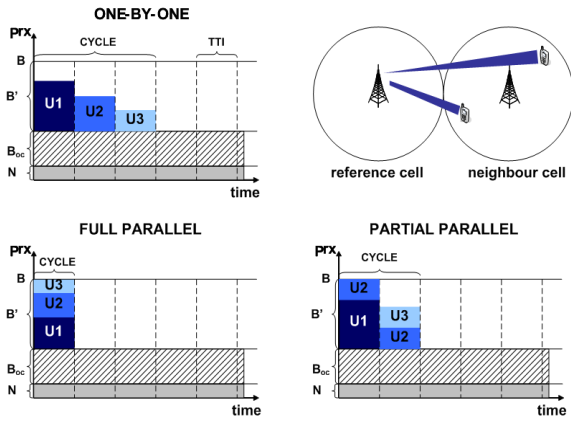


Fig. 1. Illustration of the packet handling of the considered EUL scheduling schemes.

III. MODELLING AND ANALYSIS

Our focus is on how the impact of inter-cell interference on EUL performance depends on the deployed scheduling scheme. In order to study this issue we will consider a simplified model of a small network scenario allowing for detailed analysis by simulation and analytical methods yet capturing the main UMTS/EUL system characteristics, in particular the specifics of the various scheduling schemes. The network scenario and main modelling assumptions are presented in Section III-A. The analysis is presented in Section III-B.

A. Scenario and main assumptions

We consider a scenario with one reference cell and one neighbour cell generating inter-cell interference, see Figure 2. The omnidirectional cells are split up into K concentric zones with equal areas, where zone i is characterized by a distance d_i to the base station and corresponding path loss denoted by $L(d_i)$, $i = 1, \dots, K$. In addition, in order to enable adequate modelling of inter-cell interference, the neighbour cell is split up in S sectors. The intersection of zones and sectors in the neighbour cell determines segments characterized by a distance d_{ij} and corresponding path loss $L(d_{ij})$ to the base station of the reference cell; $i = 1, \dots, K$, $j = 1, \dots, S$. It is assumed that the neighbour cell behaves independently of the reference cell. In particular, from the point of view of the reference cell, the inter-cell interference generated by the neighbour cell is an autonomous process determined by the evolution of the number of active users in each of the segments of the neighbour cell and the applied scheduling scheme. We will now first describe in more detail the assumptions concerning the reference cell; next we will consider the neighbour cell in more detail.

1) *Reference cell*: As explained in Section II, the total received power budget B at the reference cell, derived from an operator-specified noise rise target, is partially consumed by the constant thermal noise level N and a reservation level B_{oc} to cope with the anticipated inter-cell interference. The remainder of the budget - EDCH budget B' - is available for intra-cell interference originating from EDCH flows served by




Fig. 2. Modelling approach: division of the reference and neighbour cell in zones and sectors.

the reference cell. The reservation B_{oc} is dynamically updated based on the actual inter-cell interference generated by the neighbour cell. In particular, we assume that B_{oc} is adapted instantaneously at state changes of the neighbour cell. The specific choice of the value of B_{oc} as function of the state of the neighbour cell will be discussed later on.

A number of additional assumptions are made at the *user* level. EDCH flows are generated according to a spatially uniform Poisson arrival process with rate λ . The flow size is exponentially distributed with mean F (in kbits). For the performance of EDCH flows it matters in which zone they appear. As a direct consequence of the uniformity assumption and equal area choice, the probability that a generated EDCH flow appears in zone i , is $1/K$ and equal for all zones, so that the EDCH flow arrival rate per zone λ_i is equal to λ/K , $i = 1, \dots, K$. All flows have the same maximum transmit power P_{max}^{tx} but different maximum received power at the base station $P_{i,max}^{rx}$ due to the zone-dependent path loss. As no user mobility is considered, shadowing is not relevant and is not considered. Further, we do not consider fading either. The bit rate at which an EDCH flow is served depends on the experienced signal-to-interference ratio C/I . Given a prefixed E_b/N_0 (energy-per-bit to interference-plus-noise-density ratio) requirement, the attainable bit rate is equal to $r = r_{chip} (C/I) / (E_b/N_0)$, where r_{chip} denotes the system chip rate, cf. [14] eq. (8.4). The signal level C is determined by the user's transmit power and the zone-dependent path loss. The interference level I comprises several distinct components: (i) the thermal noise level N ; (ii) the self-interference modelled by parameter ω , which is due to the effects of multipath fading; (iii) the interference I_{EDCH} originating from simultaneously active EDCH users in the reference cell; and (iv) the inter-cell interference I_{oc} originating from EDCH flow transmissions in the neighbour cell.

At a given time, the system state $\underline{n} \equiv (n_1, n_2, \dots, n_K)$ is described by the number of EDCH flows n_i in zone i , $i = 1, \dots, K$.

2) *Neighbour cell*: The assumptions with respect to the neighbour cell are similar to the assumptions for the reference cell described above. The main difference is that in the neighbour cell a constant inter-cell interference level is assumed. Hence, as stated before, the behaviour of the neighbour cell is considered to be an autonomous process that is not influenced by the behaviour of the reference cell. The additional split up of each of the zones of the neighbour cell into S equally sized segments allows for a more detailed description of its

behaviour, which is needed for characterising the inter-cell interference generated towards the reference cell, see Figure 2. The state \underline{n}_{oc} of the cell is described by the number of EDCH flows in each of the segments.

B. Analysis approach

The analysis, for each of the scheduling schemes described in Section 2, runs roughly as follows: First, the inter-cell interference process caused by the neighbour cell is determined. Next, the flow level behaviour of the reference cell is analysed, given the inter-cell interference process. Both steps in the analysis are based on a flow level modelling and analysis approach for a single cell scenario with fixed inter-cell interference presented in a previous paper, see [5]. We will first briefly summarise this approach; for details we refer to [5].

1) *Basic Analysis of a Single Cell:* Consider a single cell scenario as described for the reference cell in the previous subsection. It is assumed that the inter-cell interference reservation B_{oc} is fixed. Note that the system state \underline{n} fully determines the data rate $r_i(\underline{n})$ of a user in zone i , when it is scheduled in a particular TTI. $r_i(\underline{n})$ is called the ‘instantaneous throughput’ and is given by,

$$\begin{aligned} r_i(\underline{n}) &= \frac{r_{chip}}{E_b/N_0} \cdot \frac{C}{I} = \\ &= \frac{r_{chip}}{E_b/N_0} \cdot \frac{P_i^{rx}}{I_{EDCH}(\underline{n}) - \omega P_i^{rx} + B_{oc} + N}, \end{aligned} \quad (1)$$

where $I_{EDCH}(\underline{n})$, for the different scheduling schemes, is given by

$$I_{EDCH}(\underline{n}) = \begin{cases} P_i^{rx} & \text{for OBO} \\ \min\{\sum_{i=1}^K n_i P_i^{rx}, B'\} & \text{for PP and FP} \end{cases}$$

Since $I_{EDCH}(\underline{n})$ is defined to include the referenced user’s own signal, a fraction ω of the own signal must be subtracted from $I_{EDCH}(\underline{n})$ to model the effect of self-interference. The user’s average throughput in system state \underline{n} is determined by taking into account the frequency with which the user is scheduled for transmitting data when the system is in state \underline{n} . This is the so-called ‘state-dependent throughput’ $R_i(\underline{n})$ given by

$$R_i(\underline{n}) = \frac{r_i(\underline{n})}{c(\underline{n})}, \quad (2)$$

where $c(\underline{n})$ denotes the scheduling cycle length in state \underline{n} , which depends on the applied scheduling scheme as follows:

$$c(\underline{n}) \begin{cases} \sum_{i=1}^K n_i & \text{for OBO} \\ \frac{\sum_{i=1}^K n_i P_i^{rx}}{B'} & \text{for PP} \\ 1 & \text{for FP} \end{cases} \quad (3)$$

For each of the schedulers, the system’s behaviour at flow level can be described by a K -dimensional continuous time Markov chain with transition rates λ_i (representing flow initiation in zone i) and $n_i R_i(\underline{n})/F$ (representing flow completion in zone i), $i = 1, \dots, K$. From the steady-state distribution of the Markov chain, which can be obtained by simulation in a time-efficient way, performance measures such as the zone-specific mean flow transfer time can be determined, see [5].




Fig. 3. Inter-cell interference process at the TTI level

2) Inter-cell Interference Generated by the Neighbour Cell:

For modelling the dynamics of the inter-cell interference generated by the neighbour cell towards the reference cell the evolution of the number of active users in the neighbour cell has to be tracked in more detail. In particular, one needs to know the distance of each active user to the base station of the reference cell. For that purpose the split up of the neighbour cell zones into segments (see Figure 2) is used. Extending the modelling and analysis approach presented in Section III-B1 the neighbour cell can be described by a $K * S$ dimensional continuous time Markov chain with transition rates that are straightforwardly linked to the transition rates λ_i and $n_i R_i(\underline{n})/F$ derived above.

For a given state \underline{n}_{oc} of the neighbour cell and given path loss $L(d_{ij})$ for the segments to the base station of the reference cell, the inter-cell interference generated in the TTIs of a scheduling cycle (and, hence, for the duration of that state!) is fully determined. As an example, Figure 3 shows, for one particular system state, the interference levels resulting from each of the four schedulers during a single scheduling cycle. For the FP scheme the cycle length equals 1, and hence the inter-cell interference in a given system state is just a single value. For the OBO, PP-o and PP-w schemes, the inter-cell interference may vary within a cycle, depending on the locations of the active users. Obviously, for a particular system state, the maximum level generated by the PP-w scheme within a cycle is never smaller than the maximum levels generated by the other scheduling schemes, cf. the definition of the schemes in Section II.

3) Inter-cell Interference Reservation at the Reference Cell:

In the description of the reference cell, in Section III.A, we assume that the inter-cell interference reservation B_{oc} is adapted instantaneously at state changes of the neighbour cell. The size of the reservation can be chosen in different ways depending on the desired performance. We have chosen to use a percentile of the interference level, given the state in the neighbour cell as the reservation B_{oc} . This percentile can be determined from the interference levels generated in a single scheduling cycle. For example, choosing a reservation equal to the maximum inter-cell interference generated in the scheduling cycle (i.e. the 100% percentile, which is dependent on the state of the neighbour cell) is the safest approach -

the total power budget will never be exceeded. However, a significant part of the total budget will then not be accessible for the EDCH users, which increases the mean flow transfer times. On the contrary, making no reservation at all (i.e. the 0% percentile) does not reduce the available power budget for the EDCH users, but it increases the risk that the total budget B is exceeded. This could cause degradation of the offered quality of service (QoS) via increased block error rates. A compromise option is to set the reservation at a particular percentile value between 0% and 100% of the inter-cell interference generated by the neighbour cell.

4) *Behaviour of Reference Cell*: Taking into account the inter-cell interference generated by the neighbour cell, the flow-level behaviour of the reference cell and neighbour cell can be modelled by a $(K + K * S)$ -dimensional continuous-time Markov chain. The first K dimensions describe the state of the reference cell, the other $K * S$ dimensions refer to the state of the neighbour cell. At each state transition in the neighbour cell sub-chain the inter-cell interference and, next, the corresponding required reservation B_{oc} are determined. Subsequently the new transition rates ('state-dependent throughputs') in the reference cell sub-chain can be calculated (as in the single cell case) from expressions (1), (2) and (3).

Simulating the state changes in the Markov chain described above we obtain the steady state distribution, from which performance measures like the mean flow transfer time in a particular zone can be derived. In addition, we can also determine the overall budget excess probability on TTI level. To achieve this, in each state of the Markov chain visited during the simulation, the budget excess probability is measured over a (single) scheduling cycle as the ratio of the number of TTIs in which B is exceeded and the cycle length. Appropriately weighing these state-dependent values we obtain the overall budget excess probability.

IV. NUMERICAL RESULTS

In this section we present a quantitative evaluation of the inter-cell interference and its impact on the performance of the scheduling schemes. In the comparison of the scheduling schemes we are particularly interested in the mean flow transfer time experienced by the EDCH flows. The simulations are carried out with a generic simulator developed in Matlab for deriving the steady-state distribution of the multi-dimensional Markov chain describing the system behaviour at flow level. These Markov chain simulations require relatively short running times. For example, our simulations, with 95% confidence intervals of about 1%, took typically about 10 minutes.

A. Parameter Settings

In the numerical experiments we apply a system chip rate r_{chip} of 3840 kchips/s, a thermal noise level N of -105.66 dBm and a noise rise target η of 6 dB. Hence the total received power budget is equal to $B = \eta \cdot N$. A self-interference of 10% is considered, i.e. $\omega = 0.9$. The assumed path loss is given by $L(d) = 123.2 + 35.2 \log_{10}(d)$ (in dB) with d the distance in kilometer.

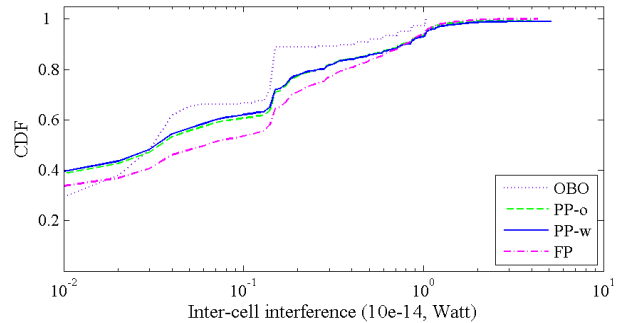


Fig. 4. CDF of the inter-cell interference levels for each of the different schedulers

Both cells are split in $K = 10$ zones¹ and the neighbour cell in $S = 6$ sectors (see Figure 2). Given a cell radius of 1.65 km we applied straightforward link budget calculations to determine the zone radii, the path losses and the maximum received powers at the reference cell, originating from users in both the reference and neighbour cell. In the calculation of the instantaneous rate we have used an E_b/N_0 target of 5 dB for EDCH transmissions and $P_{max}^{tx} = 0.125$ Watt. The applied mean file size is $F = 1000$ kbit the call arrival rate is set to $\lambda = 0.7$ for both cells.

B. Inter-cell Interference Levels

For each of the scheduling schemes, the induced inter-cell interference in a particular system state can be characterised by a unique CDF, see Section III. Deconditioning the state-dependent CDF with respect to the different system states we obtain a scheduler-specific overall CDF, as depicted in Figure 4. Observe that due to the limited set of user locations and therefore possible inter-cell interference values, these (overall) CDFs have a discrete step-like form. This is best observed for the OBO scheduler, as its single-user policy leads to a CDF with a very limited set of possible values.

An important difference between all schemes is the maximum inter-cell interference level that is observed. As expected the OBO scheme has the lowest maximum, due to the single-user policy, while PP-w has the highest maximum, because it jointly schedules users that are close to the reference cell in the same TTI.

C. Performance of the Reference Cell

We will now consider numerical results for the performance of the reference cell, in particular the budget excess probabilities and the expected flow transfer times.

1) *Budget Excess Probability*: As described in Section III-B3, several approaches exist to reserve part of the total interference budget for inter-cell interference. Each reservation option influences the performance of the schedulers differently. Table I presents the so-called 'budget excess probability', i.e. the fraction of time that the actually realised total interference

¹Extensive numerical experiments showed that this granularity is sufficient for our purposes.

TABLE I
BUDGET EXCESS PROBABILITY

reservation level	OBO	FP	PP-w	PP-o
no reservation	5.00%	33.00%	12.40%	11.94%
20% percentile	2.29%	0.00%	1.28%	1.29%
40% percentile	1.69%	0.00%	1.16%	1.15%
60% percentile	0.81%	0.00%	0.14%	0.13%
80% percentile	0.19%	0.00%	0.00%	0.00%
100% percentile	0.00%	0.00%	0.00%	0.00%

level exceeds the target level B . Such ‘budget excess’ occurs if the actual inter-cell interference exceeds the reservation level. On the one hand, a high budget excess probability is undesired, since it may correspond to high block error rates. On the other hand, if a low excess probability is achieved by an excessive reservation level, the available budget for intra-cell scheduling is unnecessarily small, causing low transfer rates and long flow transfer times. Observe that, as expected, the budget excess probability is decreasing in the reservation level. In the extreme case that the maximum reservation level is applied, the budget excess probability is 0. If no reservation is applied, the OBO scheme induces the lowest budget excess probability. This is due to its single-user scheduling policy and the fact that a user can mostly not fill up the budget on its own, which causes fewer excess events compared to PP and FP, whose goal is maximum utilization. When using percentile-based reservations, these percentiles are generally extracted from a discrete value-based CDF, which is non-trivial. In our implementation the next-higher percentile is selected as a reservation level, cf. Section II-B. In particular, in the case of FP scheduling, the inter-cell interference level in a given system state, and hence the reservation in that state, is inherently deterministic. Although Table I indicates that for a given reservation percentile the PP schemes lead to somewhat lower budget excess probabilities, the small differences must be regarded with care in light of the above-mentioned discretisation effect. For instance, considering a target reservation level of 20% and a system state with 6 users in the reference cell, the OBO scheduler establishes a cycle length of 6 TTIs and thus effectively bases the reservation on the 33% percentile, while the PP scheduler may have a cycle length of e.g. 2 TTIs and thus effectively apply a reservation based on the 50% percentile. This higher relative reservation, combined with the fact that PP generally establishes a higher inter-cell interference due to parallel transmissions, results in even higher absolute reservation and consequently in low budget excess probability.

2) *Expected flow transfer times*: We now assess the impact of the different scheduling schemes on the mean flow transfer times in the reference cell, conditional on the distance from the user to the base station. It is stressed that aside from the intra-cell impact of the scheduler, viz. the different policies according to which the budget is shared by the active users, the schedulers also affect the level of this shared budget via the reservation that is applied to cope with the scheduler-specific inter-cell interference effects. Although it is readily

understood that the mean flow transfer times will increase in the reservation level, we are particularly interested in the performance impact of the scheduling scheme.

As Figure 5 shows, indeed, using no reservation results in the lowest mean flow transfer times, but this comes at a cost of a high budget excess probability as shown in Table I. Applying a percentile-based reservation for inter-cell interference reduces the budget excess probability but leads to an increase in the mean flow transfer times, since a smaller part of the total budget B is available to the EDCH users in the reference cell. The impact of the percentile value on the flow transfer times is negligible for the PP scheme, but is significant for the OBO scheme. The longer cycle length of OBO, i.e. the larger set of inter-cell interference values, allows choosing different percentiles to result in different reservation values, while in PP the different percentiles might result in the same value.

Compare now the difference in the achieved flow transfer times under no and maximum reservation levels. It is observed that the relative impact is largest under OBO scheduling. This seems to be somewhat counterintuitive in light of the relatively low inter-cell interference levels established under OBO scheduling (see Figure 4). Note however that under OBO scheduling, the highest inter-cell interference levels are caused by user(s) close to the reference cell and hence relatively far from their serving base station. Such users have low transfer rates, long holding times and hence have a long-lasting negative effect on the budget availability in the reference cell. In contrast, the maximum inter-cell interference levels of the PP and FP schemes, even if higher in absolute value, have a shorter effect on the budget availability in the reference cell, because the schemes serve the flows faster.

If we concentrate on the distinct flavours of the PP scheduler, worst and optimized, we observe that the performance difference in terms of the location-dependent mean transfer times is negligible. Apparently, the fact that PP-w requires higher absolute reservation levels due to its larger inter-cell interference variability (see Figure 4) hardly affects the performance at flow level. This is most likely because the neighbour cell states inducing higher reservation levels do not occur very often and therefore their contribution to the mean flow transfer times remains small. Hence the exact ordering policy of a PP scheduler is of importance for the inter-cell interference pattern in individual system states, but hardly affects the performance at flow level.

We can conclude that applying the maximum reservation level seems to be an acceptable option for the FP and PP schedulers, while for the OBO scheme lower reservation percentiles are deemed more suitable. Regardless of what reservation level is applied, the PP scheme is noted to outperform the other schemes, while for the current scenarios situation the FP scheme performs second best and the OBO scheme yields the worst performance.

V. CONCLUSIONS


We have presented a modeling and analysis approach for evaluating the impact of inter-cell interference on the flow level performance of different scheduling schemes for the UMTS EUL. Further, we have compared different strategies for adaptively reserving part of the noise rise budget of a cell for interference from the neighbouring cell. Due to the hybrid analytical/simulation approach used, we were able to obtain results, such as mean flow transfer time and budget excess probability, fast.

The numerical results show that the different scheduling schemes have different inter-cell interference patterns, and also different maximum interference levels. The extent to which the schedulers' performance is affected by the chosen reservation scheme varies greatly. The OBO scheduler is very sensitive to the chosen reservation percentile, whereas the PP schedulers hardly show any sensitivity. It is especially remarkable that the ordering policy of the PP scheduler hardly influences its performance. Overall, the two PP schemes, and to a lesser extent, the FP scheme, clearly outperform the OBO scheme in terms of mean flow transfer time.


Currently we extend our analysis towards multi-cell scenario and power adaptation due to mutual inter-cell influence. Finally, we try to account for user mobility in our modeling approach.

REFERENCES


- [1] G. T. 25.308, "High Speed Downlink Packet Access (HSDPA); Overall Description."
- [2] G. T. 25.309, "FDD Enhanced Uplink; Overall Description."
- [3] H. Holma and A. Toskala, *HSDPA/HSUPA for UMTS*. John Wiley & Sons Ltd, 2006.
- [4] S. Ramakrishna and J. M. Holtzman, "A scheme for throughput maximization in a dual-class CDMA system." ICUPC '97, San Diego, USA, 1997.
- [5] D. C. Dimitrova, H. van den Berg, G. Heijenk, and R. Litjens, "Flow-level performance comparison of packet scheduling schemes for umts eul," vol. 5031. WWIC '08, Tampere, Finland, 2008.
- [6] C. Rosa, J. Outes, T. Sorensen, J. Wigard, and P. Mogensen, "Combined time and code division scheduling for enhanced uplink packet access in WCDMA." IEEE VTC '04 (Fall), Los Angeles, USA, 2004.
- [7] K. Helmersson, E. Englund, M. Edvardsson, C. Edholm, S. Parkvall, M. Samuelsson, Y.-P. Wang, and J.-F. Cheng, "System performance of WCDMA enhanced uplink." IEEE VTC '05 (Spring), Stockholm, Sweden, 2005.
- [8] C. Li and S. Papavassiliou, "On the fairness and throughput trade-off of multi-user uplink scheduling in WCDMA systems." IEEE VTC '05 (Fall), Dallas, USA, 2005.
- [9] K. Kumaran and L. Qian, "Uplink scheduling in CDMA packet-data systems," in *Wireless Networks*, vol. 12, 2006, pp. 33–43.
- [10] G. Fodor and M. Telek, "Performance analysis of the uplink of a CDMA cell supporting elastic services." Networking 2005, Waterloo, Canada, 2005.
- [11] A. Mäder and D. Staehle, "An analytical model for best-effort traffic over the UMTS enhanced uplink." IEEE VTC '06 (Fall), Montreal, Canada, 2006.
- [12] T. Liu, A. Mäder, and D. Staehle, "Analytical other-cell interference characterization over HSUPA-enabled multi-cell UMTS networks," in *IEEE VTC Fall '07*, Baltimore, USA, oct 2007.
- [13] T. Liu, A. Mäder, D. Staehle, and D. Everitt, "Analytic modeling of the UMTS enhanced uplink in multi-cell environments with volume-based best-effort traffic," in *IEEE ISCT '07*, Sydney, Australia, oct 2007.
- [14] H. Holma and A. Toskala, *WCDMA for UMTS*. John Wiley & Sons Ltd, 2001.




(a) OBO scheme



(b) FP scheme



(c) Worst PP scheme



(d) Optimized PP scheme

Fig. 5. Performance comparison of the three scheduling schemes