

Cooperative Interference Prediction for Enhanced Link Adaptation in the 3GPP LTE Uplink

Andreas Müller

Institute of Telecommunications
University of Stuttgart, Germany
E-mail: andreas.mueller@inue.uni-stuttgart.de

Philipp Frank

Intelligent Wireless Technologies & Networks
Deutsche Telekom Laboratories, Berlin, Germany
E-mail: philipp.frank@telekom.de

Abstract—The interference situation in the uplink of cellular networks, such as the 3GPP UTRAN Long Term Evolution (LTE), is usually highly volatile since from one transmission time interval to the other different users might be scheduled in nearby cells, thus often causing completely different levels of interference. This unstable behavior of the interference generally has a negative impact on the performance of fast link adaptation schemes since there is always an inherent delay between the time when the link adaptation is performed and the actual data transmission. In order to partially mitigate this problem, we propose in this paper a novel approach for predicting the interference situation during the actual data transmission already in advance, so that the link adaptation process can be considerably improved. This is accomplished by exchanging scheduling information between a set of cooperating base stations combined with multi-cell channel estimation. The performance of the proposed scheme is evaluated by means of extensive system-level simulations and we show that significant performance gains may be realized this way.

I. INTRODUCTION

The 3GPP UTRAN Long Term Evolution (LTE) represents the next major generation of cellular networks, offering peak data rates exceeding 300 Mbps in the downlink and 75 Mbps in the uplink combined with a very short radio access network round-trip time of less than 10 ms [1]–[3]. Key technologies that have been introduced into the standard for accomplishing this unprecedented performance include orthogonal frequency division multiple access (OFDMA) in the downlink and single-carrier frequency division multiple access (SC-FDMA) in the uplink, various open- and closed-loop multiple-input multiple-output transmission schemes, frequency-selective scheduling as well as adaptive modulation and coding based on a very short transmission time interval (TTI) of only 1 ms [1]–[5].

A general problem of any fast link adaptation scheme in wireless networks is that there is always an inherent delay between the time when the link adaptation is performed and the actual data transmission. As a consequence, the selected transmission parameters are often not really optimal anymore when the data transmission actually takes places. On the one hand, this is because the involved channels naturally change during that time, but at least for low to moderate user speeds, the impact of this effect should be only inferior. On the other hand and more importantly, however, the interference situation during the data transmission might be completely different from the interference situation during the time when the link adaptation has been performed. Therefore, the selected

modulation and coding schemes (MCSs) are frequently over- or underestimated, thus leading to a very high block error rate (BLER) or a rather low spectral efficiency, respectively. This particularly holds for the cellular uplink since from one TTI to the next completely different sets of users might be scheduled in nearby cells, thus causing completely different levels of interference. A standard approach to counteract this effect is to employ in addition to the conventional link adaptation scheme an outer-loop mechanism, which dynamically readjusts the target BLER based on the actually measured BLER such that the desired operating point can be achieved at least on average [6], [7]. However, even with such an outer-loop scheme, the performance is generally still significantly worse than with (theoretical) optimal link adaptation since *instantaneously* the selected MCSs often might still deviate considerably from the ones that actually would be optimal at that time.

In this paper, we therefore propose a novel approach aiming at partially mitigating this problem, which is capable of significantly improving the *instantaneous* MCS selection without having to adjust the actual target BLER. The basic idea is to predict the interference level that a base station (BS) will experience during a future data transmission already in advance when the link adaptation for this transmission is performed and to select appropriate MCSs based on this predicted rather than the currently measured interference level. This can be accomplished by exchanging scheduling decisions between a certain set of cooperating BSs via a fast backhaul network combined with multi-cell channel estimation.

The remainder of this paper is organized as follows: In Section II, we review the link adaptation mechanism of LTE while the actual interference prediction scheme is presented in Section III. Afterwards, we outline our simulation methodology in Section IV, followed by selected simulation results in Section V and finally some concluding remarks in Section VI.

II. LINK ADAPTATION IN 3GPP LTE RELEASE 8

Fast link adaptation schemes in cellular networks generally require for proper operation accurate estimates of the current channel quality of the link between the user equipment (UE) for which the adaptation should be done and its associated serving BS. In the uplink of a 3GPP LTE Release 8 network as considered here, a BS may obtain accurate estimates of the signal-to-interference-plus-noise ratios (SINRs) on the various

subcarriers by evaluating the demodulation and sounding reference signals transmitted by the corresponding UEs. While demodulation reference signals are always transmitted along with actual data in order to enable coherent detection at the BS side, sounding reference signals are used to facilitate channel estimation also for subcarriers on which a certain UE currently does not transmit any data. In this regard, the sounding reference signals themselves are always scheduled by the BSs as well, which assign appropriate sounding sequences and hopping patterns to the various UEs [1].

Having estimates of the SINRs of all subcarriers allocated to a certain UE, the BS can determine the spectrally most efficient MCS for which a given target BLER is not exceeded. For that purpose, it may choose between several different modulation schemes as well as a variety of different channel coding rates [8]. Afterwards, the selected MCS is signaled as part of the scheduling grant to the corresponding UE and the actual data transmission starts then exactly three TTIs after the reception of this grant. Hence, the delay between the link adaptation and the data transmission is in any case higher than four TTIs since at least one additional TTI is needed for the processing at the BS side and the transmission of the scheduling grant, respectively. Due to the aforementioned volatile nature of the interference situation, the SINR values that actually will occur during the data transmission consequently may be quite different from the SINR values that have been used as input for the scheduling and link adaptation algorithms. Without any appropriate countermeasures as proposed in the following, the performance therefore would be often degraded compared to the idealized case with perfect link adaptation based on the channel conditions *during the actual data transmission*.

III. COOPERATIVE INTERFERENCE PREDICTION

A. Mode of Operation

The fundamental idea of our approach is to perform the link adaptation not based upon the currently estimated SINR values, but rather based upon predicted SINR values that are likely to occur during the associated future data transmissions. For that purpose, it is necessary that a BS can accurately predict the interference level that it will experience during such future data transmissions already a couple of TTIs in advance. This may be accomplished as follows. First of all, every BS performs conventional scheduling and power control and determines which UEs should transmit on which physical resource blocks (PRBs) and at which power levels. If the employed scheduling algorithm is channel-aware—what is the case for a proportional-fair scheduler, for example—the corresponding scheduling metrics are calculated as in conventional systems, taking into account only the currently observed channel and interference conditions, respectively. Afterwards, every BS exchanges the resource allocation tables that have been fixed during the scheduling process with a certain set of cooperating BSs via a fast backhaul network. For the case of a 3GPP LTE system as considered here, this could be realized by means of the X2-interface, for example [1], [9]. Hence, provided that the various BSs have reasonably accurate channel state

information (CSI) not only of the channels from the UEs located in their own cell, but also from those associated with any of their cooperating BSs, they eventually can accurately predict the interference level that will be generated by these UEs when the actual data transmission will take place. If, for example, the channel from the i -th interfering UE to the various antenna elements of a particular BS sector is denoted by \mathbf{h}_i , the expected contribution of this interferer to the overall interference covariance matrix simply would be given by $\mathbf{R}_i = E_{S,i} \mathbf{h}_i \mathbf{h}_i^H$, with $E_{S,i}$ as the mean energy per transmit symbol of that user. The predicted interference is then used as input to the link adaptation stage and afterwards the corresponding scheduling grants (including the assigned MCS) are signaled to all scheduled UEs, which finally transmit their data exactly three TTIs after the reception of these grants.

Clearly, the performance of our approach strongly depends on the number of cooperating BSs. While with a large number of cooperation partners a BS generally should be able to predict the interference rather accurately, it would frequently underestimate the actual interference level if it cooperates with very few other BSs only. This is because in the basic scheme as described above no interference from non-cooperating cells is taken into account. In real-world scenarios, however, we expect that the set of cooperating BSs would be in most cases restricted to nearby neighbors only—on the one hand in order to keep the backhaul load limited and on the other hand because it is unrealistic that a BS may accurately estimate the channels from all UEs within a large number of cooperating cells. Therefore, it is essential that the impact of the interference caused by UEs in non-cooperating cells is taken into account as well, what might be efficiently done by making use of an additional outer loop link adaptation (OLLA) scheme. In this regard, we particularly consider an approach similar to the one presented in [6]. With this scheme, we always add a UE-specific offset Δ_{OLLA} to the predicted SINR values (in dB) before performing the actual link adaptation. This offset is permanently adjusted based on the outcome of previous transmission attempts of the corresponding UE. In particular, if an (initial) transmission attempt is successful, Δ_{OLLA} is increased by δ_{up} whereas in case that the transmission is not successful, it is decreased by δ_{down} . By relating δ_{up} to δ_{down} as

$$\delta_{\text{down}}|_{\text{dB}} = \left(\frac{1}{\text{BLER}_{\text{target}}} - 1 \right) \delta_{\text{up}}|_{\text{dB}}, \quad (1)$$

it is possible to achieve the average BLER target $\text{BLER}_{\text{target}}$. One of the main tasks of the OLLA consequently is to adjust the SINR offset Δ_{offset} such that it accounts for the average interference level generated by UEs in non-cooperating cells.

B. Practical Considerations

A crucial prerequisite for proper operation of the proposed approach in practice is that a BS is able to perform accurate multi-cell channel estimation. For that purpose, each BS has to be aware of the sounding reference signals and hopping sequences assigned to the various UEs in all cooperating cells. Hence, this information would have to be signaled in

addition to the actual resource allocation tables via the fast backhaul network. However, please note that this usually does not have to be done during every TTI since the sounding and hopping sequences normally are assigned in a semi-persistent manner, thus resulting in only minor additional backhaul load. It should also be noted that the requirements on the accuracy of the multi-cell channel estimation are generally much lower than those on the accuracy of the estimation of the desired link between a certain UE and its associated BS. On the one hand, this is because estimation errors made for different interfering channels may compensate each other—particularly if the number of cooperating BSs is relatively high—and on the other hand because it may be already sufficient for achieving a good performance to know whether on a certain PRB very high or very low interference has to be expected whereas the exact figures are only of secondary importance. In addition, if the channel from a certain UE in one of the cooperating cells cannot be estimated reliably since it is in a rather poor state, this is also not a big problem since in such a case this UE would cause only very little interference anyway.

Another prerequisite for the proposed approach is that cooperating BSs can exchange their resource allocation tables via a fast backhaul network. However, even if there are direct optical fiber links between all cooperating BSs, in general an additional delay is introduced because some time is always required for the processing of the exchanged information and the actual prediction of the interference. As a consequence, the overall latency may increase and the performance may degrade compared to the idealized case without any additional delay due to an increased mismatch between the channels used as the basis for the scheduling and link adaptation stages and those during the actual data transmission. Besides, the increased delay between scheduling and actual data transmission clearly affects also hybrid ARQ (HARQ) retransmissions. Since in the LTE uplink a synchronous HARQ protocol is used, this either would have to be replaced by an asynchronous one (but then HARQ blockings may occur) or the HARQ timing and hence the number of parallel HARQ processes would have to be increased appropriately, what we always assume in the following. Finally, it is clear that in principle even higher performance gains may be expected if not only the link adaptation is done in a cooperative fashion, but also the scheduling process itself. However, this would entail a significantly higher complexity and backhaul load as well as a higher cooperation delay. This is because changing the scheduling decision in one cell would result in different interference levels experienced by the other cooperating cells again, so that the whole process would have to be realized in an iterative fashion.

IV. SIMULATION METHODOLOGY

Extensive system-level simulations have been performed for analyzing the performance of the proposed approach, where the most important simulation parameters are given in Table I. Our deployment scenario corresponds to a hexagonal grid consisting of 19 sites with three sectors per site and we make use of the wrap-around technique in order to avoid any border

TABLE I
MOST IMPORTANT SIMULATION PARAMETERS.

Parameter	Value
Deployment scenario	19 sites with 3 sectors per site
Inter-site distance	500 m [Macro 1] or 1732 m [Macro 3]
Carrier freq. / bandwidth	2 GHz / 10 MHz
Path loss / shadow fading	According to 3GPP TR 25.814 [10]
Multipath fading	3GPP SCM
Penetration loss	20 dB
User speed	3 kmph (quasi-static)
Avg. number of UEs/sector	10
Target BLER	10%
Traffic model	Infinite full buffer
HARQ	Synchronous, non-adaptive
Parallel HARQ processes	8 + X (X = additional delay in TTI)
Scheduling algorithm	Proportional-fair according to [11]
Link-to-system interface	MIESM [12]
UE / BS antennas	1 (isotropic) / 2 (directional)
BS antenna spacing	10 times wavelength
BS noise figure	5 dB
UE category	5 (incl. 64-QAM support)
Power control	$P_0 = -58$ dBm, $\alpha = 0.6$ [Macro 1] $P_0 = -60$ dBm, $\alpha = 0.6$ [Macro 3] $P_{\max} = 24$ dBm
BS receiver type	Maximum ratio combiner (MRC)
Channel estimation	Ideal
Control channel overhead	Upper and lower 4 PRBs
Reference signals overhead	According to 3GPP TS 36.211 [4]
Default cooperation delay	2 TTIs
OLLA step size δ_{up}	0.01 dB

effects. Besides, we focus on the Urban Macro 1 and Macro 3 cases as specified in [10], with inter-site distances of 500 m and 1732 m, respectively, and if not stated otherwise always the assumptions made in [10] are used for our simulations as well. In particular, the carrier frequency is set to 2 GHz with a system bandwidth of 10 MHz and the multipath fading is modeled by means of the 3GPP spatial channel model (SCM). Furthermore, we make use of the mutual information equivalent SINR mapping (MIESM) for realizing the link-to-system interface [12] and on average there are always 10 uniformly distributed UEs in each sector. Transmit power control is performed using a simple open-loop scheme, where the transmit power of a certain UE is generally set to (in dBm)

$$P_{\text{TX}} = \min \{P_{\max}, P_0 + 10 \log_{10} M + \alpha \text{PL}\}, \quad (2)$$

with P_{\max} as the maximum transmit power, P_0 as a reference power level, M as the number of PRBs assigned to the UE, PL as the long-term attenuation of the channel between the UE and its serving BS and finally α as a constant path loss compensation factor. The values used for these parameters in our simulations are also given in Table I. Please note that (2) actually can be obtained from the power control formula explicitly specified in [8] by neglecting all (optional) short-term components. Furthermore, we make use of the variable bandwidth proportional-fair scheduler proposed in [11] as well as a synchronous, non-adaptive HARQ protocol with incremental redundancy. Finally, for the first study of the proposed approach in this paper, channel estimation is always assumed to be ideal, but the increasing inaccuracy of channel estimates with increasing time due to the time-varying nature of the channel is taken into account. While the absolute values

obtained this way may be generally overly optimistic, the assumption of perfect channel estimation allows us to quantify the maximum possible performance gains under various different conditions. Consideration of realistic channel estimation certainly would require a multitude of additional assumptions, for example regarding how BSs cooperate with each other in order to coordinate the scheduling of sounding reference signals and how the actual multi-cell channel estimation is performed, and depending on which assumptions are made, the performance may differ significantly. Due to the given space limitations, we therefore leave a more detailed treatment of the impact of realistic channel estimation for our future work.

V. SIMULATION RESULTS

Figures 1 and 2 show the performance that can be achieved with the proposed approach for the Macro 1 and Macro 3 cases, respectively. Both the average spectral efficiency as well as the cell-edge throughput—which we define as the 5th percentile of the UE throughput distribution—are depicted and we consider different numbers of cooperating BSs as well as the idealized case with perfect link adaptation, where the MCS is selected in retrospect based on the channel conditions during the actual data transmission. The case with 6 cooperating sectors actually corresponds to the situation where each BS receives resource allocation tables from all surrounding *sectors* whereas in case of 20 cooperating sectors each BS receives resource allocation tables from all sectors of all 6 surrounding *sites* as well as the other 2 sectors of the same site.

It can be seen from both figures that significant performance gains may be obtained, where the relative gains are always higher in terms of the cell-edge throughput than in terms of the average spectral efficiency. This is because UEs with a rather poor channel on average generally are scheduled on only very few PRBs since they often become power-limited. Hence, without interference prediction, variations of the interference level on these PRBs between the time when the link adaptation is performed and the actual data transmission often have a higher impact on the performance than for cell-center UEs, which are usually scheduled on more PRBs so that with higher probability changing interference conditions on some PRBs might be compensated by conversely changing interference conditions on other PRBs. In other words, if on some PRBs the interference level estimated during the link adaptation is higher than the one observed during the data transmission, this impact is partially compensated if on other PRBs it is exactly the other way around. Clearly, the more PRBs a certain UE gets, the more probable it becomes that such a compensation occurs. Hence, cell-edge UEs generally suffer more from the volatile nature of the interference and therefore they also benefit more from the proposed interference prediction scheme.

It can also be seen from Figs. 1 and 2 that with 6 cooperating sectors already most of the potential performance gains that are theoretically possible with interference prediction can be realized. This simply reflects the fact that the biggest share of the interference generally comes from the surrounding sectors, what is particularly true for the Macro 3 case with

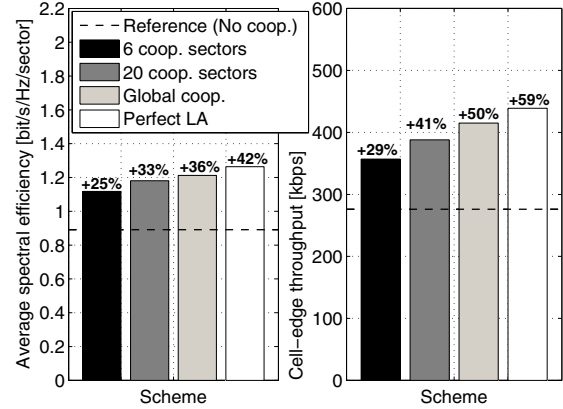


Fig. 1. System performance for the Macro 1 case with interference prediction and various numbers of cooperating sectors as well as (idealized) perfect link adaptation (LA).

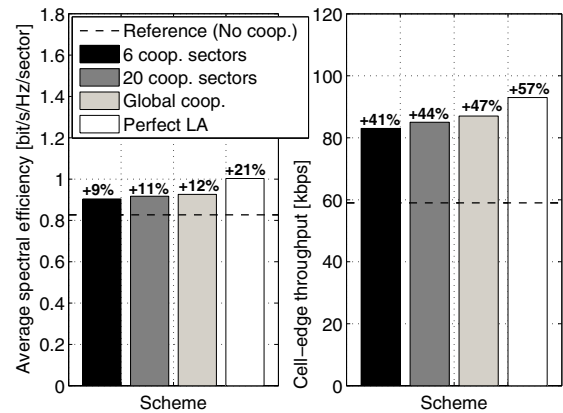


Fig. 2. System performance for the Macro 3 case with interference prediction and various numbers of cooperating sectors as well as (idealized) perfect LA.

a relatively large inter-site distance. Besides, since in the Macro 3 case the experienced interference level usually is less dominant compared to the noise level than for the Macro 1 case, the relative performance gains in terms of the average spectral efficiency are notably smaller as well. However, this obviously does not hold for the cell-edge throughput. This is because in the Macro 3 case the probability that cell-edge UEs become power-limited is higher than for the Macro 1 case, thus leading to a decreased probability of a compensation of varying interference levels on various PRBs as outlined before. Finally, it should be noted that the performance with global cooperation in Figs. 1 and 2 is still worse compared to the case with perfect link adaptation since even with global cooperation, we still have a certain delay between the link adaptation stage and the actual data transmission.

Figure 3 illustrates the fact that the link adaptation becomes more accurate with the proposed approach by showing the deviation between the MCS that would be optimal during the actual data transmission with the one that actually has

been selected. As can be seen, with interference prediction the probability that the perfect MCS is selected is almost doubled compared to the conventional case without this approach and aside from that the variations around the optimal MCS obviously can be significantly reduced as well.

Finally, Fig. 4 depicts the impact of the additional delay introduced by our approach, considering three different cases with 2, 6, and 20 cooperating sectors, respectively. In the first case, we actually have intra-site cooperation since scheduling decisions are only exchanged between the sectors belonging to the same site. As can be seen, with increasing delay the performance becomes steadily worse, what can be attributed to the fact that the channels of the various UEs change during that time. Hence, the channels used as the basis for the scheduling and link adaptation stages deviate more and more from the channels during the actual data transmission. For the case with 2 cooperating sectors, the performance actually might become even worse than without interference prediction if the delay exceeds 6 TTIs, but since no backhaul signaling is required in that case, the actual delay in practice usually should be much smaller. With 6 or 20 cooperating sectors, in contrast, even with 10 TTIs delay still moderate performance gains can be realized, but in presence of optical fiber links between cooperating BSs, the actual delay might be also in this case much smaller, thus still leading to considerable improvements.

VI. CONCLUSION

We have presented a novel approach for accurately predicting the interference level that will occur during a future data transmission in the uplink of cellular networks. Based on this predicted interference, it is possible to considerably improve the link adaptation process due to a generally more accurate selection of appropriate modulation and coding schemes. The proposed approach is based upon the exchange of scheduling decisions between different cooperating BSs via a fast backhaul network combined with multi-cell channel estimation. We have investigated the performance of our scheme by means of extensive system-level simulations for a 3GPP LTE system and it turned out that significant performance gains may be realized this way. Since the proposed approach is transparent to the UEs and causes only a moderate backhaul load, it represents a very attractive option for future LTE-Advanced systems.

ACKNOWLEDGEMENTS

The authors acknowledge the excellent cooperation with all partners within the EASY-C project and the support of the German Federal Ministry of Science and Education (BMBF).

REFERENCES

- [1] H. Holma and A. Toskala, *LTE for UMTS – OFDMA and SC-FDMA based radio access*. Chichester: John Wiley and Sons, 2009.
- [2] H. Ekstrom, A. Furuskär, J. Karlsson, M. Meyer, S. Parkvall, J. Torsner, and M. Wahlqvist, “Technical solutions for the 3G long-term evolution,” *IEEE Commun. Mag.*, vol. 44, no. 3, pp. 38 – 45, Mar. 2006.
- [3] D. Astely, E. Dahlman, A. Furuskär, Y. Jading, M. Lindström, and S. Parkvall, “LTE: The evolution of mobile broadband,” *IEEE Commun. Mag.*, vol. 47, no. 4, pp. 44 – 50, Apr. 2009.
- [4] 3GPP TS 36.211 V8.9.0, “Evolved universal terrestrial radio access (E-UTRA): Physical channels and modulation,” Dec. 2009.

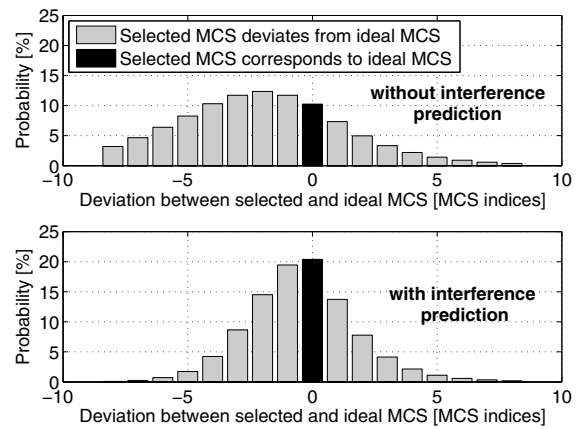


Fig. 3. Accuracy of the link adaptation process without and with interference prediction for the Macro 1 scenario. In the latter case, each BS always receives scheduling information with 2 TTIs delay from all 6 surrounding sectors.

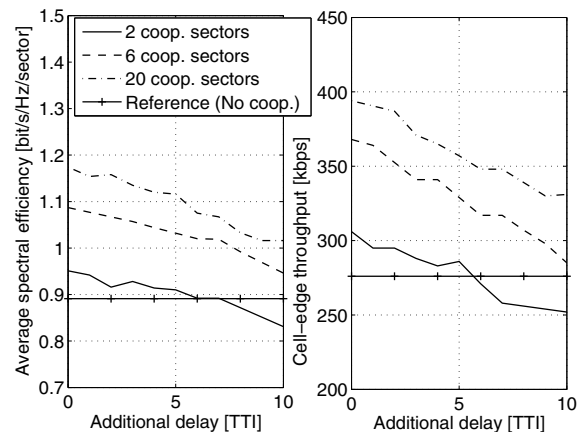


Fig. 4. Impact of the additional delay due to interference prediction on the system performance for the Macro 1 case and three different cluster sizes.

- [5] K. I. Pedersen, T. E. Kolding, F. Frederiksen, I. Z. Kovacs, D. Laselva, and P. E. Mogensen, “An overview of downlink radio resource management for UTRAN long-term evolution,” *IEEE Commun. Mag.*, vol. 47, no. 7, pp. 86 – 93, July 2009.
- [6] M. Nakamura, Y. Awad, and S. Vadgama, “Adaptive control of link adaptation for high speed downlink packet access (HSDPA) in W-CDMA,” in *Proc. Int. Symp. Wireless Pers. Multim. Commun.*, Oct. 2002.
- [7] A. Pokhariyal, K. I. Pedersen, G. Monghal, I. Z. Kovacs, C. Rosa, T. E. Kolding, and P. E. Mogensen, “HARQ aware frequency domain packet scheduler with different degrees of fairness for the UTRAN long term evolution,” in *Proc. IEEE Veh. Technol. Conf.*, Apr. 2007.
- [8] 3GPP TS 36.213 V8.8.0, “Evolved universal terrestrial radio access (E-UTRA): Physical layer procedures,” Sept. 2009.
- [9] 3GPP TS 36.420 V8.1.0, “Evolved universal terrestrial radio access network (E-UTRAN): X2 general aspects and principles,” Dec. 2008.
- [10] 3GPP TR 25.814 V7.1.0, “Physical layer aspects for evolved universal terrestrial radio access (UTRA),” Sept. 2006.
- [11] F. D. Calabrese, C. Rosa, M. Anas, P. H. Michaelsen, K. I. Pedersen, and P. E. Mogensen, “Adaptive transmission bandwidth based packet scheduling for LTE uplink,” in *Proc. IEEE Veh. Technol. Conf.*, Sept. 2008.
- [12] K. Brueninghaus, D. Astely, T. Sälzer, S. Visuri, A. Alexiou, S. Karger, and G.-A. Seraji, “Link performance models for system level simulations of broadband radio access systems,” in *Proc. IEEE Int. Symp. Pers., Indoor and Mob. Radio Commun.*, Sept. 2005.