Power-Controlled Cross-Layer Scheduling

Johannes Gonter, Norbert Goertz
Institute of Telecommunications
Vienna University of Technology
Email: johannes.gonter@nt.tuwien.ac.at

***Abstract*—Power-Controlled Cross-Layer Scheduling is proposed for real-time multimedia-applications in cellular wireless networks. The new scheme borrows elements from Proportional Fair Scheduling, from which it differs in that it can satisfy hard delay constraints, and improves on exploiting multiuser diversity gains by selectively excluding users from scheduling decisions based on their rate demands.**
The performance of the new scheduler is compared to performance bounds from multiuser information theory. Then, the basic concept is extended to a real-world scenario, in which the long-term path losses of all users will change over time: this necessitates the introduction of a power control scheme, that, if carefully designed, "inverts" the long-term pathloss of the channel, but not the fast fading. An analysis of the outage probability characteristic shows near-optimum behavior for delay-constrained applications and reveals huge performance gains for demanding users, compared to classical proportional fair scheduling.

***Index Terms*—Cross-Layer Scheduling, Proportional Fair, Hard Delay Constraints, High Data Rate (HDR)**

I. Introduction

During the last decade, consumer communications has evolved away from speech and moved progressively towards the exchange of multimedia content. Together with more and more powerful user-equipment and emerging technologies like transmission of high-resolution 3D-video, this trend is and will be challenging even for the most elaborate infrastructure (including LTE advanced), particularly when high rates are requested with hard delay constraints, as e.g. in conversational video applications. There are hints that future contracts will be based on guaranteed data-rates rather than the absolute amount of data, which will lead to a much increased consumer-awareness with respect to Quality of Service (QoS) and Quality of Experience (QoE).
All of these developments call for new paradigms in wireless network management. Some of the most cost-efficient methods to make optimum use of the existing infrastructure and its allocated frequency spectrum are advanced scheduling concepts: we have developed a new approach we call "Power-Controlled Cross-Layer Scheduling" (PCCLS). The major difference to most other scheduling schemes presented in the literature is that, to the best of our knowledge, there is no other scheduler that guarantees (with very small outage) hard delay constraints for bit rates requested by the users and at the same time exploits multi-user diversity gains [1] with little loss compared to ergodic capacity limits [2]. We chose the term "Cross-Layer Scheduling" because the scheduler (which is classically [3], [4] based in the medium access layer) requires, in our new scheme, both knowledge of the channel coefficients (known in the physical layer) for each user but also information about all users' applications (known in the application layer), in particular their delay constraints and rate requests. PCCLS builds upon proportional fair scheduling (PFS) [5], which is also the baseline scheduler of the CDMA/HDR system [6]. The difference between the two scheduling approaches is, however, very fundamental: PFS aims to mitigate the "unfairness" of strictly opportunistic communication: users are scheduled whenever *their* instantaneous channel quality is good compared to *their* recent average. In principle, the number of time-slots across which the average is computed provides the possibility to integrate delay-constraints into the PFS-decision metric. However, the average rate a user will get still only depends on the channel quality and is, hence, a best-effort service - it was even shown in [7] that PFS is not stable at all, i.e. the scheme, though widely adopted, does not guarantee finite queue-lengths. Therefore, PFS would not be a suitable scheduler for applications with hard delay constraints (in the sense defined below), such as any real-time multimedia-application. The scheduler we propose in this paper adopts ideas from PFS and combines them with user-individual rate-monitoring and a "slow" power control. More specifically, the power control is designed to be slow compared to the fast fading, to ensure that the channel is not simply inverted, which might cause the use of a huge transmit power for a temporarily bad channel.
The paper is organized as follows. Section II introduces the system model and defines the delay constraints. Section III briefly mentions theoretical limits, and section IV introduces PCCLS, the new scheduling scheme. In section V, we discuss the algorithm, and subsequently provide numerical results in section VI. Section VII compares scheduling efficiency to Proportional Fair scheduling, and section VIII concludes the paper.

II. System Model

We assume that J mobile users request data bit rates r_j (in kbits/s) for their applications with different hard delay constraints from the base station. Once the central network control has accepted to serve a user, we assume the required resources (power in particular) to be available. The challenge is to provide the services as requested by the users and at the same time to achieve multi-user diversity gains to increase the overall system power-efficiency.

The links from the base station to the mobile users are modeled as block-fading Gaussian broadcast channels. The current values of the channel power gains are assumed to be known at the transmitter. In the downlink, this can be practically realized by users estimating the channel power gains and returning them to the base station. The channel power gains include both the long-term path-loss (mainly depends on the distance to the base station) and the effects of "fast" fading, e.g., due to scattering and interference. We assume that the path-loss changes slowly, so that it can be viewed as constant across many time slots (block fading model).

A user-specific delay parameter τ_j defines the maximum number of successive time slots over which one block of source data of a user can be spread out. More specifically, the hard delay constraints used in this paper are met, when the inequality

$$\bar{R}_j(k) \geq r_j \cdot t_s \quad \forall j, k \tag{1}$$

is fulfilled. In (1) t_s is the channel-symbol time in seconds, $\bar{R}_j(k)$ is the average rate (in bits per channel-use) obtained by user j in time window $k = n \cdot \tau_j - 1$, $n = 1, 2, ...$; it is defined by

$$\bar{R}_j(k) \doteq \frac{1}{\tau_j} \sum_{i=0}^{\tau_j - 1} R_j(k-i) \ , \tag{2}$$

and $R_j(l)$ is the rate achieved (in bits per channel-use) by user j in time slot l.

III. Information Theoretical Limits

In the framework of Information Theory, the system model in Section II can be described by a block-fading Gaussian broadcast channel (e.g. [8]) where, in our scenario, a "block" is the same as a time-slot. When user j *is* scheduled we use the standard Gaussian capacity formula to compute the achieved rate (bits/channel-use), i.e.,

$$R_j(k) = \log_2 \left(1 + \frac{P_j(k) \cdot |H_j(k)|^2}{\sigma^2}\right) , \tag{3}$$

with $P_j(k)$ the transmit power allocated to user j in time slot k and $H_j(k)$ the corresponding channel coefficient; σ^2 is the variance of the Gaussian receiver noise we assume to be equal for all users (for simplicity). In a practical system we will not be able to achieve (3) but we can determine a power offset required by some practical coding and modulation scheme to achieve the rate $R_j(k)$ at some (low) bit-error rate that is acceptable for the application.

Although there has been work on block fading channels with rate-averaging over more than one block (time slot), most of the results (e.g., [9]) only apply to single user channels. Results for the fading broadcast channel are either restricted to single time slots (then, outage capacities are analyzed, [10]) or they apply only for averaging over an unlimited number of time slots (then, ergodic capacities can be stated, [2]). The dilemma with the practical setup we consider is that neither of the two cases really applies to our problem. In principle one would have to investigate zero-outage (delay limited) capacities of the fading broadcast channel with averaging of the achieved rates over τ time slots, with $1 < \tau \ll \infty$: unfortunately, there are no analytical results available from literature. What is known, however, is that ergodic capacity ($\tau \to \infty$) is a strict upper bound for zero-outage capacity with $\tau \ll \infty$, so we compare our simulation results below with ergodic capacity limits, assuming that the values of τ we use are "close enough" to infinity – the simulation results will confirm this point of view.

IV. Power-Controlled Cross-Layer Scheduling (PCCLS)

We initialize the transmit powers for all users by solving (3) for the transmit power, while considering the requested rates r_j and assuming all users are scheduled every J time slots:

$$P_j(1) = \left(2^{r_j \cdot (t_s \cdot J)} - 1\right) \cdot \frac{\sigma^2}{|\bar{H}_j|^2} \ . \tag{4}$$

By use of this initial power we equalize the *average* path-loss $|\bar{H}_j|^2$ of each user and we obtain "symmetric" fading channels for all users in the first iteration of the scheduler. Our experiments have shown that the long-term convergence of the power is not affected, as long as the initial assumptions are "reasonable" (as (4)).

The new scheduler works as follows for every time slot $k = 1, 2, ...$:

1) **Calculation of the rate** $R_j(k)$ user j *can* get in time slot k:

$$R_j(k) = \log_2 \left(1 + \frac{P_j(k) \cdot |H_j(k)|^2}{\sigma^2}\right) \cdot \mathcal{F}_j(k). \tag{5}$$

Here, $\mathcal{F}_j(k)$ is a "scheduling flag": only a user whose rate request hasn't been satisfied in the current delay-window will be considered for scheduling.

2) **Scheduling decision**: user $j_k^* \in \{1, 2, ..., J\}$ is scheduled in time slot k according to

$$j_k^* = \arg \max_{j=1,...,J} \frac{R_j(k)}{T_j(k)} \ . \tag{6}$$

$T_j(k)$ is the current rate-average for user j; as an initialisation, $T_j(1) = \log_2 \Big(1 + (P_j(1) \cdot |\bar{H}_j|^2)/\sigma^2\Big)$ can be used.

3) **Update of the rate-average** for time-slot $k+1$ for all users $j = 1, 2, ..., J$:

$$T_j(k+1) = \begin{cases} (1 - \frac{1}{a_j})T_j(k) + \frac{R_j(k)}{a_j}, & j = j_k^* \\ (1 - \frac{1}{a_j})T_j(k), & j \neq j_k^* \end{cases} \ . \tag{7}$$

Only if user j is served in time slot k (that is, $j = j_k^*$) the exponentially decaying average rate is increased by addition of $R_j(k)/a_j$, with $R_j(k)$ the current possible rate on user j's channel as determined in step 1.

The factor a_j defines the slope of the exponential decay in the computation of the rate average. In our experiments we have found that $a_j = \tau_j/10$ is a suitable choice. We are using an exponentially decaying window rather than a rectangular window for averaging, since this smoothes the average and eliminates panic-like scheduling decisions when a time-slot with a "high-rate peak" moves out of the averaging window.

4) **Power control:** The power for user j in the next time slot $k+1$ is calculated by

$$P_j(k+1) = \begin{cases} P_j(k) \cdot 2^{r_j \cdot t_s - T_j(k+1)}, & j = j_k^* \\ P_j(k), & j \neq j_k^* \end{cases} . \quad (8)$$

The *updated average* rate $T_j(k+1)$ is used in (8) to check whether the rate request, r_j, was matched and to correct the power if necessary. The rate-*average* is used with the aim to compensate for the *average long-term path-loss* but *not* for the fast fading, which is still to be exploited to obtain multi-user diversity gains [1]. For details about the power control mechanism, please see the Appendix.

5) **Increment slot counter**, i.e, $k := k+1$, and **go to step 1.**

V. Discussion

The major differences between PFS and PCCLS in Section IV are the introduction of the outer power-control loop and the user-individual scheduling exclusion. It is the task of the outer power-control loop to compensate for the slowly changing long-term path-loss so that every user can be served such that their rate- and delay-requests are met and not only on a best-effort basis as in PFS. The goal is *not* to combat fading but rather to exploit it, i.e., to schedule a user when the fading channel's fast fluctuations temporarily lead to a "good" channel. By introduction of the "schedule Flag" $\mathcal{F}_j(k)$, the system resources can be used much more efficiently compared to PFS, since PCCLS aims to use the minimum transmit power for every user to guarantee requested rate and delay requirement. The "schedule Flag" then guarantees an optimum channelization of resources to users who currently need them most. PCCLS does not have any queue-stability problems: it was reported [7] that PFS does not ensure queue stability, i.e., there is no guarantee that the length of a user's queue remains limited. This problem does not exist with PCCLS, as the rate requests are to be met within a specified time window: PCCLS will make sure that all arriving data will be delivered in time – otherwise, data packets would be useless due to violated delay constraints and could be dropped.

The power control in a real-world system would have to be adapted, as, in the basic concept proposed, the power required for some data rate is derived from the "Shannon-Capacity" that can not be achieved in practice. Hence, a power margin that depends on the specific coding and modulation schemes in use would have to be introduced. The scheduling decisions as such (i.e., which user is allowed to access the channel) would, however, still be justified.

VI. Numerical Results

We assume a generic system example summarised in Table I with $J = 8$ users with different rate requests and delay constraints. The system is assumed to operate with a channel symbol time of $0.4\mu s$ (parameters similar to CDMA2000) and a time-slot duration of $0.4ms$ (1000 channel symbols per slot). The profile of User 1 is that of a mobile phonecall (hard delay constraint of $\tau_j = 50$ time slots (i.e., $20ms$) but low rate request) while User 8 runs video streaming (relaxed delay constraints of $\tau_j = 1000$ time slots (i.e., $400ms$) but higher rate demand). For the average channel power gains (that are to account for path loss) convenient values were chosen, since they don't affect the relative performance of the scheduling scheme. Apart from the fixed path loss, the transmission is subject to Rayleigh fading; the fading coefficients are assumed to be constant within each time slot but they change randomly from one time slot to another.

Table II shows the theoretical limits for the scheduling problem described by Table I: it lists the channel SNRs (ratio of the transmit power and the receiver noise in dBs) required to achieve ergodic capacities that equal the rate requests of the users given in Table I; of course those SNRs are controlled by the allocated transmit power with the channel coefficients known at the scheduler. The general "code division" solution is optimal but it requires code superposition and successive interference cancellation at the decoder, which is impossible to implement in practice, as delay constraints would be violated. Therefore, the "TD" column contains the theoretical limits that apply with a time-division constraint (at most one user is scheduled per time slot). As obvious from the table, there is hardly any difference between the "CD" and the "TD" solution with respect to the average power used by the base station, so in practice one would certainly opt for the TD

User j	Rate Request r_j in kbits/s	Delay Limit τ_j slots	(Time/ms)	Average Channel Pow. Gain G_j/dB
1	50.0	50	(20)	-27.0
2	100.0	186		-26.0
3	150.0	321		-25.0
4	200.0	457	(183)	-24.0
5	250.0	593		-23.0
6	300.0	729		-22.0
7	350.0	864		-21.0
8	400.0	1000	(400)	-20.0

TABLE I
USER PROFILES FOR SIMULATIONS; THE CHANNEL SYSMBOL TIME IS SET TO 0.4μS.

User j	CD (optimal) SNR_j in dB	TD SNR_j in dB
1	17.21 (3.17)	17.71 (3.05)
2	17.57 (5.90)	18.00 (5.77)
3	17.53 (7.15)	17.90 (7.05)
4	17.31 (7.78)	17.66 (7.71)
5	16.96 (8.06)	17.32 (8.05)
6	16.50 (8.15)	16.93 (8.19)
7	15.93 (8.07)	16.48 (8.16)
8	15.22 (7.87)	16.02 (8.09)
Average	16.30	16.31

TABLE II
THEORETICAL LIMITS: USERS' CHANNEL SNRS REQUIRED TO ACHIEVE ERGODIC CAPACITIES ACCORDING TO THEIR RATE REQUESTS AND AVERAGE CHANNEL POWER GAINS GIVEN IN TABLE I: GENERAL CASE (CODE-DIVISION, CD) [2] AND WITH TIME-DIVISION CONSTRAINT (TD). SNRS WITHIN BRACKETS INCLUDE "ZERO POWER" WHEN USERS ARE NOT SCHEDULED.

solution as it is much more simple to implement (even if there was no problem with delay constraints). The theoretical limits for the SNRs in Table II are lower bounds for the scheduling and power-allocation problem described by Table I because the rate requests are seen as ergodic capacities that have to be achieved and ergodic theory ignores delay constraints as all rates are *long-term* averages. In contrast to that we wish to achieve the users' rate requests within each successive time window of length τ_j (number of time slots), i.e. we want the rate average (2) within each time-window to fulfill (1): that is what defines hard delay constraints in the given context. If we used a scheduling scheme derived from the information theoretical solution for the time-division constraint, this scheme will utterly fail in terms of the delay constraints but it will provide a lower limit of 16.31 dB for the average power used (see Table II).

A simple, practical approach would be Round-Robin scheduling, where every user is scheduled in a fixed repetitive pattern, and the delay constraints are fulfilled by equalizing the channel in each slot k with suitable transmit power. Such as scheme will consume 31.8 dB average transmit power (for the problem described by Table I), so almost 16 dB more power is needed compared to the ergodic limits in Table II; perhaps even worse, the peak power would be unlimited. Proportional fair scheduling (PFS) can not be used in the given setup as PFS is a best-effort scheme that will serve the users with rates according to their channel qualities. Hence, it is impossible to give any rate guarantees or fulfill delay constraints.

Fig. 1 shows the results for the PCCLS-scheme presented in Section IV. The achieved average rates (2) within time

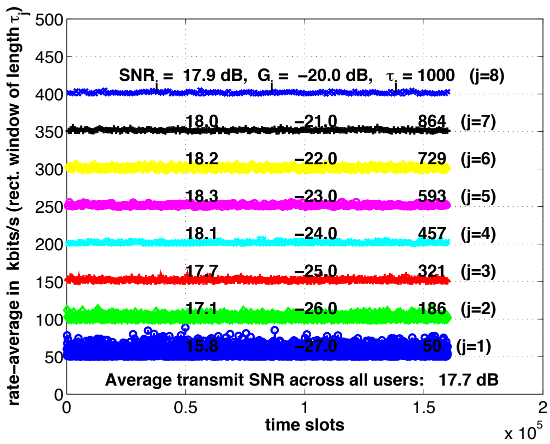

Fig. 1. Simulation results for Power-Controlled Cross-Layer Scheduling. The average rates (2) are plotted on the y-axis; on the x-axis the time slot indices $k = n\tau_j - 1$ are given with $n = 0, 1, 2, \ldots$.

windows of size τ_j are depicted. While the delay constraints according to (1) are nicely met for the high-rate user ($j = 8$, 400 kbits/s), the achieved average rates show some fluctuations for user $j = 1$ with hard delay constraints (rate request of 50 kbits/s). But even in the latter case, the rate requests are fulfilled, as the value of (2) does not drop below the requested rate. The most significant result is that PCCLS fulfills the delay constraints and at the same time the average power used leads to a channel SNR of only 17.7 dB compared to 16.31 dB for the theoretical limits. Hence, we achieve a *practical* power gain of more than 14 dB compared to Round-Robin scheduling with channel inversion.

VII. Outage Analysis

We define an outage as the event that the actually achieved average rate for user j is smaller than the requested rate r_j, within the delay window of length τ_j. The probability of such an event is the outage probability. Based on our recently published analytical approach to calculate the outage-probability for PFS in delay-constrained applications [11], we are going to compare the outage-probability for PCCLS and PFS for the above system parameters. The optimum outage probability characteristic for a scheduler guaranteeing delay constraints as well as a certain rate is a step-function. The desired outage probability is zero for the requested rate and one for every bigger rate. PCCLS is designed to achieve this goal by optimizing the transmit power for every user and excluding users from the scheduling decision if their rate requests have been satisfied. Therefore, PCCLS can be seen to spend as much power as needed, but not more energy than required on every user. The following figure 2 illustrates the benefits of this behavior.

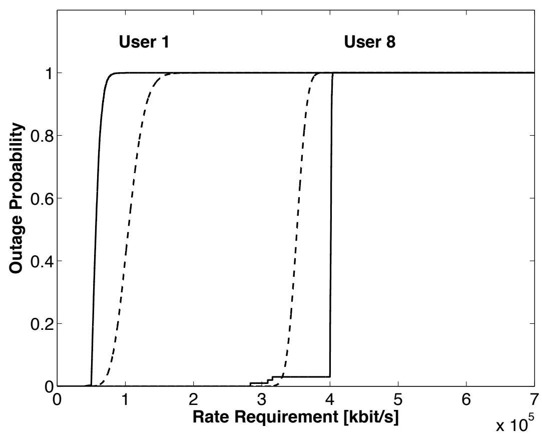

Fig. 2. Outage Probability comparison between Power-Controlled Cross-Layer Scheduling and Proportional Fair Scheduling, based on the same transmit powers for both schemes.The dashed lines show the outage probabilities for PFS, the solid lines for PCCLS.

The curves were obtained by simulation, where first PCCLS was run on 100,000 blocks, and subsequently PFS was performed on the *same* channel coefficients, with the transmit power previously optimized in the PCCLS process. In addition, for PFS the EWMA (exponentially-weighted, moving-average) filter was initialized with an optimized seed, so that the whole sequence of 100,000 blocks could be used for the outage-assessment. For "User 1", the characteristic behavior of PCCLS, as compared to PFS, is already visible: PCCLS optimizes the transmit-power, initialized as in (4), such that the probability of an outage is zero for the requested rate (when the starting-phase it considered as well, the probability of an outage is close to zero). Due to users no longer being scheduled as soon as their respective rate-requests are met within the current delay-window, the probability of an outage increases rapidly above the requested rate, which is $50kbits/s$ in case of "User 1". PFS, on the other hand, is a best-effort service with a good performance for "User 1", however, energy is wasted for rates above the required rate. This is hurting users with high rate demands, such as "User 8": with a rate requirement of $400kbits/s$, this user suffers from energy being spent on users whose rate-requests have already been satisfied. In figure 2, the initializing phase has been taken into account for PCCLS in order to show the behavior for adverse changes in pathloss, but even then the huge gain of cross-layer scheduling as compared to PFS is obvious. PCCLS can, while reducing the energy being spent to a minimum (even no users may be served temporarily), provide service with close-to-zero outage-probability, where PFS causes a waste of energy and will almost certainly fail to provide the requested rates to demanding users.

VIII. Conclusions

We have proposed a new scheduling scheme which we call "Power-Controlled Cross-Layer Scheduling" (PCCLS). It is based on Proportional Fair scheduling (PFS) but differs fundamentally in that PCCLS is designed to guarantee hard delay constraints. This is achieved by adopting a cross-layer approach: PCCLS takes into account both knowledge about the physical channel (the channel coefficients) and application requirements (requested data rates and delay constraints). In contrast to PFS, which is a best-effort scheme that does not guarantee any delay constraints and has been shown to be potentially unstable [7], PCCLS applies power control to ensure a low outage probability and selectively switches off users if their rate-requests have been satisfied. Thus, by comparing the outage probabilities between PCCLS and PFS, the new scheme can be seen to reduce the energy spent in the whole system while at the sam time significantly reducing the probability of an outage for users with high requested data rates. We therefore believe that PCCLS is suited to increase the performance of cellular wireless networks beyond what is possible with PFS, taking resource-efficiency to the next level.

Appendix

We start with the rate R_1 we get from some transmission power P_1 with some unknown channel coefficient H on a complex Gaussian channel with a noise power of $N_0/(2t_s)$ in each real sub-channel:

$$R_1 = \log_2\big(1 + \frac{G \cdot P_1}{\sigma^2}\big)\ . \tag{9}$$

This formula is an approximation as the channel coefficient H in (9) is actually an *average* coefficient because we use the rate *average* for power correction in (8).

If R_1 is the rate which was delivered and R_2 is the rate we actually want to achieve, we seek the factor with which to scale the power P_1 to get the requested rate R_2. Therefore, we solve (9) for the power which gives

$$P_1 = \big(2^{R_1} - 1\big) \cdot \frac{\sigma^2}{\bar{g}}\ , \quad \bar{g} \neq 0\ . \tag{10}$$

We take the ratio of the required power P_2 for the rate R_2 and the power P_1 with which we achieved rate R_1:

$$\frac{P_2}{P_1} = \frac{\big(2^{R_2} - 1\big)}{\big(2^{R_1} - 1\big)} \doteq f(R_1, R_2)\ . \tag{11}$$

Although the function f in (11) would work well as a power scaling rule in (8), it turns out to be fairly "agressive" leading to an "unsmooth" behaviour of the power adaptation. A less aggressive scaling functions, can be found by (optimistically) assuming that the rate difference $\Delta R \doteq R_2 - R_1$ is small (it is the job of the scheduler to ensure exactly this). We introduce the rate difference in (11) an obtain

$$\frac{P_2}{P_1} = \frac{\big(2^{R_1+\Delta R} - 1\big)}{\big(2^{R_1} - 1\big)} = 2^{\Delta R}\frac{\big(2^{R_1} - 2^{-\Delta R}\big)}{\big(2^{R_1} - 1\big)} \tag{12}$$

As we assume that $\Delta R \approx 0$, we obtain from a Taylor series expansion $2^{-\Delta R} \approx 1 - \Delta R \cdot \log(2)$, so we get

$$\frac{P_2}{P_1} = 2^{\Delta R}\frac{\big(2^{R_1} - 1 + \Delta R \cdot \log(2)\big)}{\big(2^{R_1} - 1\big)}\ . \tag{13}$$

By neglecting additive term $\Delta R \cdot \log(2)$ in the numerator we find the simple rule

$$P_2 = P_1 \cdot 2^{R_2 - R_1}\ . \tag{14}$$

References

[1] R. Knopp and P.A. Humblet, "Information capacity and power control in single-cell multiuser communications," in *Proceedings IEEE International Conference on Communications (ICC)*, Seattle, WA, USA, June 1995, pp. 331–335.

[2] L. Li and A. Goldsmith, "Capacity and optimal resource allocation for fading broadcast channels – Part I: Ergodic capacity," *IEEE Transactions on Information Theory*, vol. 47, no. 3, pp. 1083–1102, Mar. 2001.

[3] H. Zimmermann, "OSI reference model – the ISO model of architecture for open systems interconnection," *IEEE Transactions on Communications*, vol. COM–28, no. 4, pp. 425–432, Apr. 1980.

[4] ISO/IEC, "Information technology – open systems interconnection – basic reference model: The basic model," *ISO/IEC 7498–1:1994(E)*, 1994.

[5] P. Viswanath, D.N.C. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1277–1294, June 2002.

[6] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi, "A bandwidth efficient high speed data service for nomadic users," *IEEE Communications Magazine*, July 2000.

[7] M. Andrews, "Instability of the proportional fair scheduling algorithm for HDR," *IEEE Transactions on Wireless Communications*, vol. 3, no. 5, pp. 1422–1426, Sept. 2004.

[8] L. H. Ozarow, S. Shamai, and A. D. Wyner, "Information theoretic considerations for cellular mobile radio," *IEEE Transactions on Vehicular Technology*, vol. 43, no. 2, pp. 359–378, May 1994.

[9] G. Caire, G. Taricco, and E. Biglieri, "Optimum power control over fading channels," *IEEE Transactions on Information Theory*, vol. 45, no. 5, pp. 1468–1489, July 1999.

[10] L. Li and A. Goldsmith, "Capacity and optimal resource allocation for fading broadcast channels – Part II: Outage capacity," *IEEE Transactions on Information Theory*, vol. 47, no. 3, pp. 1103–1127, Mar. 2001.

[11] J. Gonter, N. Goertz, and A. Winkelbauer, "Analytical outage probability for max-based schedulers in delay-constrained applications," in *Proceedings 2012 IFIP Wireless Days conference*, Nov. 2012, Accepted.