Power Minimization in Parallel Vector Broadcast Channels with Zero-Forcing Beamforming

Christoph Hellings, Michael Joham, and Wolfgang Utschick

Global Communications Conference, 2010 (IEEE GLOBECOM 2010)

Miami, FL, USA, December 2010

Technische Universität München
Associate Institute for Signal Processing
http://www.msv.ei.tum.de

Power Minimization in Parallel Vector Broadcast Channels with Zero-Forcing Beamforming

Christoph Hellings, Michael Joham, and Wolfgang Utschick
Associate Institute for Signal Processing, Technische Universität München, 80290 Munich, Germany
Telephone: +49 89 289-28510, Fax: +49 89 289-28504, Email: {hellings,joham,utschick}@tum.de

***Abstract*—We consider the problem of power minimization under per-user quality of service (QoS) constraints (expressed in terms of rates) in parallel multiple-input single-output (MISO) broadcast channels employing linear zero-forcing precoding. Solving the arising scheduling problem by an exhaustive search is prohibitively complex due to its combinatorial nature so that the use of successive user allocation schemes has been proposed. We show that existing schemes lead to strongly suboptimal solutions in systems with a low number of degrees of freedom and that better performance can be achieved with a new scheduling criterion. By introducing additional correction steps, we end up with an efficient close-to-optimum algorithm.**

I. Introduction

Although the maximization of a weighted sum of rates in a multi-user communication system has attained considerable theoretical interest, it is hardly adapted to the needs of a system that has to guarantee the quality of service (QoS) for each user. Assuming transmission with error-free decoding, a reasonable QoS measure is the achievable sum rate of a user according to Shannon's capacity formula. At the same time, for reasons such as minimizing interference to neighboring communication systems or reducing energy consumption, the system operator might be interested in minimizing the total transmit power. These considerations lead to the problem of power minimization under per-user rate constraints.

As an example of a system consisting of parallel vector broadcast channels, we consider a multi-carrier downlink communication system with a multi-antenna base station.[1] Motivated by the high computational complexity of practical implementations of the non-linear, capacity-achieving dirty paper coding (DPC), we constrain the system to apply linear precoding. Thus, it is not possible to apply power minimization algorithms relying on the convex nature of the DPC capacity or power region, as those proposed in [1], [2] for multi-carrier and in [3], [4] for single-carrier MIMO broadcast channels.

Such single-carrier MIMO algorithms are in principle applicable to multi-carrier systems if the carrier-specific channels of the users are written into block-diagonal channel matrices (cf. e.g., [5]). However, unlike for the non-linear case, no algorithms exist that obtain the global optimimum of the non-convex power minimization problem for MIMO BCs with linear precoding. So, it would be necessary to use suboptimal schemes, such as the one proposed in [6]. Moreover, the dimensions of the resulting MIMO system might be prohibitively large, and the special structure might pose problems for some heuristic approaches, e.g., the coordinated transmit-receive processing from [7] can generate non-invertible effective channels due to the blockdiagonal structure of the channel matrices.

[1]The results also hold if groups of carriers within the frequency coherence interval are considered instead of individual carriers.

For the linear multi-carrier case, a convex power minimization problem has been formulated in [8] by allowing time-sharing between transmission strategies. However, systems without time sharing are in the focus of this paper. Additionally, we do not specialize to setups with the number of transmit antennas exceeding the total number of receive antennas as assumed in [9]. Furthermore, algorithms fulfilling per-stream QoS constraints, such as the power minimization algorithms in [10], [11], and optimizations with per-user MMSE constraints, as in [12], are not directly applicable to our setup as there is no explicit relation between these performance criteria and the per-user sum rate in the case of multi-stream transmission.

Heuristic QoS constrained optimization schemes applicable to the considered system model have been proposed in [13] and [14]. Like most approaches to find close-to-optimum linear transmit strategies for multi-carrier broadcast channels, they are based on a successive user allocation and on zero-forcing, i.e., it is assumed that the degrees of freedom at the base station have to be utilized to totally wipe out inter-stream interference (cf. Section III). Although suboptimal in general, zero-forcing is of practical interest, as it can be easily implemented and as it is known to be optimal in systems with low noise power. As shown in Section IV, the optimal zero-forcing power allocation can be explicitly calculated for a given allocation of users to subcarriers. Our solution to the arising scheduling problem follows the lines of the greedy user allocation scheme in [13], which will be presented in Section V together with some clarifications on implementational issues. Our main contributions are a performance analysis for cases where the number of users is close to the overall number of degrees of freedom (fully loaded system) and two extensions to the algorithm to drastically improve its behavior in those cases (cf. Section VI and the numerical results in Section VII). The algorithm from [14] is not further considered as it was developed for systems with much less users than degrees of freedom and does not guarantee feasible solutions in the fully loaded case we concentrate on.

II. System Model and Problem Formulation

We consider a multi-carrier downlink system where the base station is equipped with M antennas while the K receivers are equipped with a single antenna each. The N subcarriers are assumed to be orthogonal to each other, i.e., there is no inter-carrier interference. We further assume that $K \leq MN$, which is a necessary condition for the zero-forcing techniques proposed in this paper. The frequency flat vector channel between the base station and user k on subcarrier $n \in \{1, \ldots, N\}$ is denoted by $\boldsymbol{h}_k^{(n),\mathrm{H}} \in \mathbb{C}^{1\times M}$, and these channels are assumed to be perfectly known. The additive circularly symmetric complex Gaussian noise $\eta_k^{(n)} \sim \mathcal{CN}(0, \sigma_k^{(n),2})$ is assumed to be independent across users and across subcarriers and independent of the transmitted data symbols $s_k^{(n)} \sim \mathcal{CN}(0,1)$. The received signal of user k on carrier n can be written as

$$y_k^{(n)} = \boldsymbol{h}_k^{(n),\mathrm{H}} \sum_{k'=1}^{K} \boldsymbol{u}_{k'}^{(n)} \sqrt{p_{k'}^{(n)}} s_{k'}^{(n)} + \eta_k^{(n)} \tag{1}$$

with the transmit power $p_{k'}^{(n)} \geq 0$ and the unit norm beamformer $\boldsymbol{u}_{k'}^{(n)} \in \mathbb{C}^M$ of user k' on carrier n. The total rate of user k is $\sum_n r_k^{(n)}$, and $r_k^{(n)}$ is the rate achieved on carrier n.

We will focus on the question what total transmission power is necessary to fulfill a set of given non-zero per-user rate requirements ρ_k. The optimization problem reads as

$$\begin{aligned} &\min \quad \sum_{k=1}^{K}\sum_{n=1}^{N} p_k^{(n)} \\ &\text{s.t.: } \sum_{n=1}^{N} r_k^{(n)} \geq \rho_k \quad \forall k \quad \text{and} \quad p_k^{(n)} \geq 0 \quad \forall\, k, n \end{aligned} \tag{2}$$

where the minimization is over all powers and beamforming vectors, and $r_k^{(n)}$ is a function of those variables on carrier n. The difficulty of this problem is that there is an inherently strong coupling of different users' rates on a carrier while the constraints impose a coupling between the rates of streams on different carriers belonging to one user.

III. Zero-Forcing Beamforming

As the subcarriers are assumed to be orthogonal to each other, zero-forcing only needs to be performed in the spatial domain and can be achieved by choosing the beamformers to be the scaled columns of the Moore-Penrose pseudo-inverse of the joint channel matrix $\boldsymbol{H}^{(n)} = [\boldsymbol{h}_{k_1^{(n)}}^{(n)}, \ldots, \boldsymbol{h}_{k_{S^{(n)}}^{(n)}}^{(n)}]^{\mathrm{H}} \in \mathbb{C}^{S^{(n)}\times M}$ of the users scheduled on carrier n, where $k_s^{(n)}$ is the user corresponding to the s-th stream on subcarrier n. With the normalization of the beamforming vectors, the channel gains of the resulting independent scalar subchannels are

$$g_{k_s^{(n)}}^{(n)} = \sigma_{k_s^{(n)}}^{(n),-2} \left[\left(\boldsymbol{H}^{(n)} \boldsymbol{H}^{(n),\mathrm{H}} \right)^{-1} \right]_{s,s}^{-1} \tag{3}$$

where $[\boldsymbol{A}]_{s,s}$ is the s-th diagonal element of $\boldsymbol{A}$.

Note that the joint channel matrix $\boldsymbol{H}^{(n)}$ must have full row rank, which implies that the number of users $S^{(n)}$ scheduled on a carrier n has to be smaller than or equal to the number of transmit antennas ($S^{(n)} \leq M$). Additionally, as can be inferred from the matrix inversion in (3), the channel gains strongly depend on the condition of $\boldsymbol{H}^{(n)}$. Thus, finding an allocation of users to subcarriers that leads to well-conditioned joint channel matrices on all carriers is a crucial step.

We call a K-tuple $\mathcal{S} = (\mathcal{N}_1, \ldots, \mathcal{N}_K)$ of sets $\mathcal{N}_k \subseteq \{1, \ldots, N\}$ a valid user allocation if each carrier n is element of at most M sets $\mathcal{N}_k$. The tuple is called a feasible user allocation if $\mathcal{N}_k \neq \emptyset \;\forall k$, i.e., if all users are scheduled on at least one subcarrier. This condition is not only necessary but also sufficient for feasibility as the active streams are not interference-limited and can therefore carry any demanded rate when the transmit power is chosen sufficiently large.

Now, we rewrite problem (2) as a scheduling problem and a subsequent power allocation over the effective scalar channels:

$$\begin{aligned} &\min_{\mathcal{S}\in\mathbb{S}} \min_{\left\{p_k^{(n)}\right\}_{\substack{k\in\{1,\ldots,K\}\\ n\in\mathcal{N}_k}}} \sum_{k=1}^{K}\sum_{n\in\mathcal{N}_k} p_k^{(n)} \\ &\text{s.t.: } p_k^{(n)} \geq 0 \quad \forall k, \forall n \in \mathcal{N}_k \quad \text{and} \quad \sum_{n\in\mathcal{N}_k} r_k^{(n)} \geq \rho_k \quad \forall k \end{aligned} \tag{4}$$

where $r_k^{(n)} = \log_2(1 + p_k^{(n)} g_k^{(n)})$, $g_k^{(n)}$ is a function of $\mathcal{S}$, and $\mathbb{S}$ denotes the set of valid user allocations. The two parts of this optimization problem will be treated in the following sections.

IV. Power Allocation

The channel gains resulting from the zero-forcing beamforming [see (3)] are solely a function of $\mathcal{S}$ and can be considered as constants in the inner optimization. As a result, there is no coupling between users in the inner problem, and it can be written separately for each user k:

$$\begin{aligned} &\min_{\left\{p_k^{(n)}\right\}_{n\in\mathcal{N}_k}} \sum_{n\in\mathcal{N}_k} p_k^{(n)} \\ &\text{s.t.: } p_k^{(n)} \geq 0 \quad \forall n \in \mathcal{N}_k \quad \text{and} \quad \sum_{n\in\mathcal{N}_k} r_k^{(n)} \geq \rho_k. \end{aligned} \tag{5}$$

From the KKT conditions, we get the waterfilling equation

$$p_k^{(n)} = \max\left\{0,\ \mu_k - g_k^{(n),-1}\right\} \tag{6}$$

with the optimal water level

$$\mu_k = \left(2^{-\rho_k} \cdot \prod\nolimits_{n\in\mathcal{N}_{k,\mathrm{a}}} g_k^{(n)}\right)^{-1/|\mathcal{N}_{k,\mathrm{a}}|} \tag{7}$$

where $\mathcal{N}_{k,\mathrm{a}}$ is the set of user k's active subcarriers, i.e., those with a transmit power greater than zero. Sorting the entries of $\mathcal{N}_k$ with respect to the channel gains, the optimal water level can be found by a linear search for $|\mathcal{N}_{k,\mathrm{a}}| \in \{1, \ldots, |\mathcal{N}_k|\}$, which is equivalent to the solution proposed in [14].

V. The Problem of User Allocation

Finding the optimal user allocation $\mathcal{S}$ is a combinatorial problem. A common approach to avoid an exhaustive search are greedy schemes that successively allocate resources to users by a series of locally optimal decisions, leading to a globally suboptimal solution in general. In [13], a greedy user allocation for problems with QoS constraints as well as a simplified low complexity allocation were proposed. We focus on the greedy scheme, which can be applied to the problem at hand as follows. In step i, for each subcarrier n that still has available spatial degrees of freedom ($S^{(n)} < M$) and for each user k that is not yet scheduled on n,[2] the user allocation $\mathcal{S}(k,n)$ resulting from assigning user k to carrier n is created:

$$\mathcal{S}(k,n) \leftarrow ({}^{(i-1)}\mathcal{N}_1, \ldots, {}^{(i-1)}\mathcal{N}_k \cup \{n\}, \ldots, {}^{(i-1)}\mathcal{N}_K). \quad (8)$$

The left superscript is the step index. For each $\mathcal{S}(k,n)$, the resulting total transmit power $P(k,n)$ is computed, and we set ${}^{(i)}\mathcal{S} \leftarrow \mathcal{S}(\mathrm{argmin}_{(k,n)} P(k,n))$. The algorithm terminates when the sum power increases from one step to the next.

Two issues to be considered besides the basic idea have been mentioned in [13] without a detailed study. First, the scheduler solving the outer optimization must ensure $\mathcal{N}_k = \mathcal{N}_{k,\mathrm{a}}\ \forall k$ to avoid the suboptimal situation that an unused stream imposes a zero-forcing constraint. We propose to repeatedly remove all zero power streams from $\mathcal{S}(k,n)$ and recalculate $P(k,n)$ until the allocated power is non-zero for all streams.[3] To avoid an endless loop of deallocation and reallocation, we let the algorithm also terminate when it arrives at a constellation which has already been considered, even if the total transmit power did not increase but stayed constant. Typically, this happens when no change in allocation has occured from one step to the next. Convergence is guaranteed as the total number of possible constellations is finite. Without the possibility of deallocation, the number of evaluations of the inner problem would be in the order of $\mathcal{O}(MKN^2)$ resulting from up to KN sum power computations for not more than MN streams. Numerical simulations suggest that the number of deallocated streams is small in general so that the average number of evaluations lies in the same order of magnitude.

The second issue mentioned in [13] is a more fundamental one, which is inherent to problems with QoS constraints: As long as not all users have at least one stream, none of the $\mathcal{S}(k,n)$ represents a feasible user allocation. Consequently, there is the need for an initialization phase with a heuristic criterion other than a pure greedy criterion. In [13], during this initialization phase, only users that have not yet been scheduled can be allocated, and decisions are made based on the greedy criterion while ignoring the constraints of inactive users. However, this technique leads to very suboptimal solutions if the system is close to being fully loaded ($K \approx MN$).

The allocation scheme discussed in this section is summarized in Algorithm 1. The minimization in Line (17) ignores pairs of (k,n) that have not been considered in the current iteration since all powers are initialized with ∞ in Line (3). For the same reason the termination criterion also applies in case that all possible MN streams have been allocated.

Algorithm 1 Basic Algorithm

Require: $M,\ N,\ K,\ \rho_k,\ \boldsymbol{h}_k^{(n)},\ \sigma_k^{(n)}$
(1) $i \leftarrow 0,\ {}^{(0)}P = 0,\ {}^{(0)}\mathcal{S} \leftarrow (\emptyset, \ldots, \emptyset)$
(2) **repeat**
(3) $\quad i \leftarrow i+1,\ \mathcal{K} \leftarrow \{1, \ldots, K\},\ P(k,n) \leftarrow \infty\ \forall k, \forall n$
(4) $\quad$ **if** $i < K$ **then**
(5) $\quad\quad \mathcal{K} \leftarrow \{k \mid \mathcal{N}_k = \emptyset\}$
(6) $\quad$ **end if**
(7) $\quad$ **for all** $k \in \mathcal{K}$ and $n \notin {}^{(i-1)}\mathcal{N}_k$ with $S^{(n)} < M$ **do**
(8) $\quad\quad \mathcal{N}_k \leftarrow {}^{(i-1)}\mathcal{N}_k \cup \{n\},\ \mathcal{N}_{k'} \leftarrow {}^{(i-1)}\mathcal{N}_{k'}\ \forall k' \neq k$
(9) $\quad\quad$ **repeat**
(10) $\quad\quad\quad$ calculate all $g_{k'}^{(n')}$ from (3)
(11) $\quad\quad\quad$ compute $p_{k'}^{(n')}$ from (6) for all k' with $\mathcal{N}_{k'} \neq \emptyset$
(12) $\quad\quad\quad$ remove streams with zero power from all $\mathcal{N}_{k'}$
(13) $\quad\quad$ **until** no zero powers were allocated
(14) $\quad\quad \mathcal{S}(k,n) \leftarrow (\mathcal{N}_1, \ldots, \mathcal{N}_K)$
(15) $\quad\quad P(k,n) \leftarrow \sum_{k'=1}^{K} \sum_{n' \in \mathcal{N}_{k'}} p_{k'}^{(n')}$
(16) $\quad$ **end for**
(17) $\quad (k^*, n_{k^*}) \leftarrow \mathrm{argmin}_{(k,n)} P(k,n)$
(18) $\quad {}^{(i)}\mathcal{S} \leftarrow \mathcal{S}(k^*, n_{k^*}),\ {}^{(i)}P \leftarrow P(k^*, n_{k^*})$
(19) **until** $i > K$ and $({}^{(i)}P > {}^{(i-1)}P$ or $\exists i' < i : {}^{(i)}\mathcal{S} = {}^{(i')}\mathcal{S})$
(20) **return** $({}^{(i-1)}\mathcal{S}$, optimal $p_k^{(n)}$ for ${}^{(i-1)}\mathcal{S})$

VI. The Initialization Phase

If there are many degrees of freedom, there is still a wide range of reasonable allocations for the last users even if the first ones were allocated in a strongly suboptimal manner. Additionally, the deallocation of inactive streams can be seen as a correction of bad initial decisions. However, for cases with a low number of degrees of freedom, the initialization phase plays an important role, and in the extreme case of a fully loaded system, there is even no main phase. Thus, we will present two performance enhancing modifications to the initialization phase of the basic algorithm.

It can be observed that the basic algorithm tends to schedule users with low rate requirements first as this results in a low power increment in the current step. However, since in the long run, the constraints of all users have to be fulfilled, the difficult users with high requirements are considered last, when nearly no degrees of freedom are left. Clearly, this is not a preferable strategy as the numerical results show.

A. Look-Ahead (LA) Scheduling

To avoid such a shortsighted behavior, a criterion is needed that reflects the importance of scheduling a user on its optimal carrier. To this end, we introduce the degradation measure $\delta_k = P(k, \tilde{n}_k) - P(k, n_k)$ where n_k and $\tilde{n}_k$ are the best and second-best carrier for user k, respectively. By allocating user

$$k^* = \underset{k}{\mathrm{argmax}}\, \delta_k \quad \text{to carrier} \quad n_{k^*} = \underset{n}{\mathrm{argmin}}\, P(k^*, n) \quad (9)$$

during the initialization phase, we ensure that a user is immediately scheduled when the risk of a significant increase in sum power arises from not doing so. This decision rule, which can be interpreted as the avoidance of foreseeable critical situations, does not only have a clear rationale, but also proves to perform well in numerical simulations. Additionally,

[2] The last condition is specific to single-antenna receivers.

[3] In general, it is not necessary to remove all inactive streams, as deallocating one stream might suffice for other streams to get active. However, we rely on the greedy criterion to automatically reschedule unnecessarily deallocated streams in the subsequent steps.

while decisions of successive schemes usually are only optimal within one step, decisions based on (9) are optimal in the last *two* steps. This will be important later in this section.

Theorem 1. *Using criterion (9), the last two decisions of the initialization phase are optimal given a fixed allocation of previously scheduled users.*

Proof: We denote the users by k_1 and k_2, ${}^{(i-1)}P$ is the sum power before allocating the first user, and $\Delta P(k,n)$ is the power increment resulting from assigning user k to carrier n. We note that the last decision is optimal due to the rule $n_{k^*} = \operatorname{argmin}_n P(k^*, n)$.

In case of different optimal carriers n_{k_1} and n_{k_2}, both users will be assigned to their optimal carrier since allocating user k_1 to carrier n_{k_1} does not impair the optimality of n_{k_2} for user k_2 and vice versa. For $n_{k_1} = n_{k_2}$, it is either optimal that both users are scheduled on $n_{k_1} = n_{k_2}$ or one user k^* on its optimal carrier n_{k^*} and the other user on its second-best carrier. In the former case, the optimal solution is achieved independently of which user is scheduled first due to the optimality of the last decision. What remains to be shown is that the optimal k^* is chosen in the latter case. As assigning one user to a carrier does not change the situation on another carrier, we get

$$P_{\text{total}} = {}^{(i-1)}P + \Delta P(k_1, \tilde{n}_{k_1}) + \Delta P(k_2, \tilde{n}_{k_2}) + (\Delta P(k^*, n_{k^*}) - \Delta P(k^*, \tilde{n}_{k^*})) \quad (10)$$

at the end of the initialization phase. For this to be optimal, it is necessary that k^* is the user for which $(\Delta P(k^*, n_{k^*}) - \Delta P(k^*, \tilde{n}_{k^*}))$ is minimized. We further note that

$$\delta_k = P(k, \tilde{n}_k) - P(k, n_k) = \Delta P(k, \tilde{n}_k) - \Delta P(k, n_k), \quad (11)$$

which means that k^* has to be chosen such that δ_{k^*} is maximized, which is also the scheduling criterion. The scheduling of the other user is also optimal as it is the last decision. ■

B. Decide-and-Challenge (DC) Scheduling

Except for the case of zero power allocation (which will never happen in a fully loaded system), we so far have not deallocated streams. However, the concept of LA scheduling fails whenever a critical situation is unforeseeable. Consider the example when the last two available degrees of freedom are on the same subcarrier and the two last users to be scheduled have highly correlated channels on that particular carrier. In this case, the suboptimality of this allocation cannot be detected before one of the two users has been scheduled. From this example, it can be seen that to further improve performance it is necessary to re-evaluate decisions that have been made earlier and to allow a step back if necessary.

To this end, we propose a decision challenge phase after K streams have been allocated, i.e., at the end of the allocation in the fully loaded case or before starting the main phase in all other cases. It will be seen that this phase can be designed such that it neither risks to worsen the situation nor significantly increases the total computational complexity.

In general, it is not obvious which of the scheduling decisions shall be put into question. We therefore propose to deallocate the oldest and the newest stream. While the first allocation risks to have been a bad decision because it was made with the least amount of knowledge about the other allocations, the last one is potentially bad as the possible choices were already rather limited. Afterwards, the two users are reallocated according to the criterion chosen for the initialization phase, in our case LA scheduling.

Even if the user allocation does not change, i.e., if the old decision withstands the challenge, we have now a different oldest and potentially also a different newest stream. Thus, it makes sense to repeat the procedure $\lfloor \alpha K \rfloor$ times with $\alpha > 0$. In our simulations, $\alpha = 2$ turned out to be a good compromise between complexity and performance. The justification for choosing $\alpha > 1$ is that when a change in scheduling has occurred, it is worth retrying the deallocation and reallocation of a stream even if it has already been performed before.

From Theorem 1, it follows that the deallocation and reallocation of two users can never increase the sum power when the LA criterion is used. Thus, the sum power after the challenge phase is surely not higher than it was before. Note that if the system is not fully loaded, this does not imply that the final sum power with DC scheduling obtained after the main phase is guaranteed to be smaller than without DC scheduling, but at least on average this is what happens.

The increase in computational complexity is bounded from above by $2N\lfloor \alpha K \rfloor$ evaluations of the inner optimization and is negligible compared to the overall $\mathcal{O}(MKN^2)$ evaluations.

VII. Numerical Analysis

For our simulations, we have used 1000 realizations of random channel coefficients which are i.i.d. circularly symmetric complex Gaussian with zero mean and unit variance and per-user rate requirements which are the absolute values of i.i.d. real Gaussian random variables with zero mean and unit variance. Furthermore, we have fixed the noise power to $\sigma_k^{(n),2} = 1$ for all users and carriers.

For the results in Fig. 1 and Fig. 2, we have performed simulations in relatively small systems in order to be able to compare the performance with the optimal solution obtained with an exhaustive search. Fig. 1 shows the average performance of the presented schemes in systems with $M = 2$ transmit antennas, $K = 4$ users, and $N \in \{2, 3, 4, 5\}$ subcarriers.[4] As can be seen, the LA and LA-DC schemes perform close to the optimal solution while the basic algorithm, which implements the scheme proposed in [13], shows a notable increase in transmit power. Note that the loss of the basic algorithm is more pronounced for a small number of subcarriers, i.e., if the number of degrees of freedom is scarce.

In Fig. 2, we have considered a system with $M = 4$ transmit antennas, $N = 3$ subcarriers, and $K = 12$ users, where we have replaced the random QoS requirements by $\rho = 2\rho_0$ for half of the users and by $\rho = \rho_0$ for the rest. Plotting the total transmit power P over ρ_0, it can be seen that there is no

[4]We have taken the geometric mean (equivalent to the arithmetic mean in the dB domain) as it is more robust against outliers.

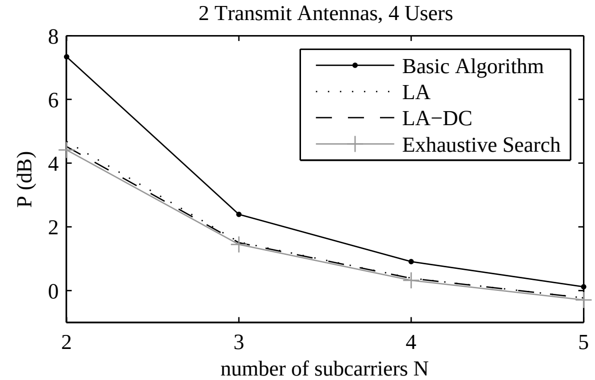

Fig. 1. Achieved Transmit Power for Different Numbers of Carriers.

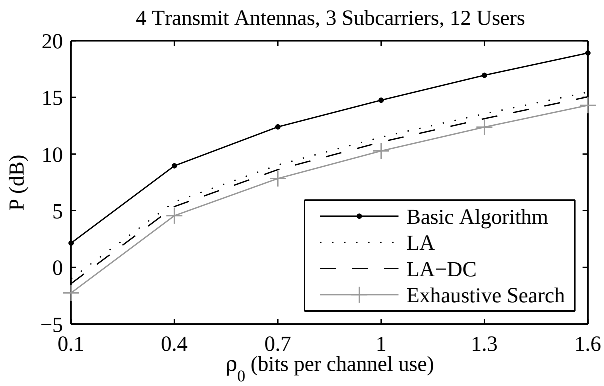

Fig. 2. Achieved Transmit Power for Different Rate Requirements.

qualitative difference between the behavior of the discussed schemes in the low and in the high rate regime.

To show that the results also hold for larger systems, a configuration with a significantly higher number of possible allocations is presented in Fig. 3. For a system with $M = 2$ transmit antennas, $N = 16$ subcarriers, $K = 32$ users, and random QoS requirements, a histogram of the total transmit powers of the various schemes is shown. Note that the optimal solution is not available for this setting due to the probitive complexity of the exhaustive search. While LA and LA-DC have a clear peak at low powers, the distribution of the powers resulting from the use of the basic algorithm is widely spread. Moreover, the additional gain in performance achieved by the DC scheme becomes visible in this plot.

VIII. Conclusion

In our discussion of successive allocation schemes for power minimization under per-user rate constraints, it became obvious that a scheduling criterion other than a pure greedy criterion is needed for the phase where not all users have been allocated their first stream. The difference between the total power resulting from a user's currently best and second-best scheduling has proven to be a criterion that, in most cases, is able to circumvent allocations which might lead to critically suboptimal situations in future steps. However, as not all those critical situations are foreseeable, a further increase

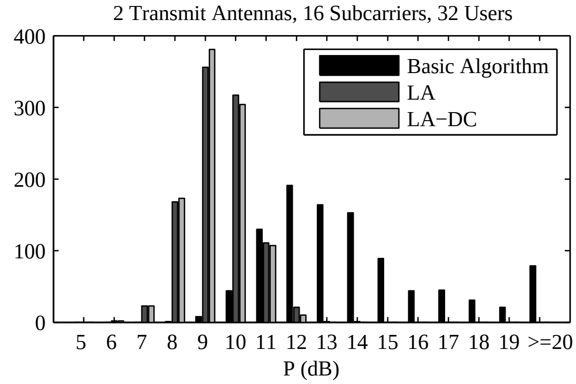

Fig. 3. Absolute Frequency of Occurence of Transmit Powers.

in performance can be achieved by systematically putting preceding scheduling decisions into question and allowing changes of those decisions if needed.

Acknowledgment

This work was supported by Deutsche Forschungsgemeinschaft (DFG) under fund Jo 724/1-1.

References

[1] G. Wunder and T. Michel, "Minimum Rates Scheduling for MIMO OFDM Broadcast Channels," in *Proc. ISSSTA 2006*, Aug. 2006, pp. 510–514.

[2] M. Mohseni, R. Zhang, and J. Cioffi, "Optimized transmission for fading multiple-access and broadcast channels with multiple antennas," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1627–1639, Aug. 2006.

[3] W. W. L. Ho and Y.-C. Liang, "Efficient Power Minimization for MIMO Broadcast Channels with BD-GMD," in *Proc. ICC 2007*, Jun. 2007, pp. 2791–2796.

[4] C.-H. F. Fung, W. Yu, and T. J. Lim, "Multi-antenna downlink precoding with individual rate constraints: power minimization and user ordering," in *Proc. ICCS 2004*, Sep. 2004, pp. 45–49.

[5] P. Tejera, W. Utschick, G. Bauch, and J. Nossek, "Efficient Implementation of Successive Encoding Schemes for the MIMO OFDM Broadcast Channel," in *Proc. ICC 2006*, vol. 12, Jun. 2006, pp. 5354–5359.

[6] A. Mezghani, M. Joham, R. Hunger, and W. Utschick, "Transceiver design for multi-user MIMO systems," in *Proc. ITG WSA 2006*, Ulm, Germany, Mar. 2006.

[7] Q. H. Spencer, A. L. Swindlehurst, and M. Haardt, "Zero-forcing methods for downlink spatial multiplexing in multiuser MIMO channels," *IEEE Trans. Signal Process.*, vol. 52, no. 2, pp. 461–471, Feb. 2004.

[8] N. Hassan and M. Assaad, "Low complexity margin adaptive resource allocation in downlink MIMO-OFDMA system," *IEEE Trans. Wireless Commun.*, vol. 8, no. 7, pp. 3365–3371, Jul. 2009.

[9] Q. H. Spencer, A. L. Swindlehurst, and M. Haardt, "Fast power minimization with QoS constraints in multi-user MIMO downlinks," in *Proc. ICASSP 2003*, vol. 4, Apr. 2003, pp. IV–816–IV–819.

[10] A. Khachan, A. Tenenbaum, and R. Adve, "Linear Processing for the Downlink in Multiuser MIMO Systems with Multiple Data Streams," in *Proc. ICC 2006*, vol. 9, Jun. 2006, pp. 4113–4118.

[11] M. Codreanu, A. Tölli, M. Juntti, and M. Latva-aho, "Joint Design of Tx-Rx Beamformers in MIMO Downlink Channel," *IEEE Trans. Signal Process.*, vol. 55, no. 9, pp. 4639–4655, Sep. 2007.

[12] M. Schubert, S. Shi, and H. Boche, "Iterative transceiver optimization for linear multiuser MIMO channels with per-user MMSE requirements," in *Proc. EUSIPCO 2006*, Sep. 2006.

[13] C. Guthy, W. Utschick, and G. Dietl, "Spatial Resource Allocation for the Multiuser Multicarrier MIMO Broadcast Channel – A QoS Optimization Perspective," in *Proc. ICASSP 2010*, Mar. 2010.

[14] F. She, W. Chen, H. Luo, T. Huang, and X. Wang, "Joint Power Allocation and Scheduling of Multi-Antenna OFDM System in Broadcast Channel," in *Proc. ICC 2009*, Jun. 2009, pp. 1–5.