Distributed Interference Pricing for the MIMO Interference Channel

Changxin Shi*, David A. Schmidt†, Randall A. Berry*, Michael L. Honig* and Wolfgang Utschick†

*Department of Electrical Engineering and Computer Science
Northwestern University, Evanston, Illinois 60208, USA
Email: cshi@u.northwestern.edu, {rberry, mh}@ece.northwestern.edu

†Associate Institute for Signal Processing
Technische Universität München, 80290 Munich, Germany
Email: {dschmidt, utschick}@tum.de

Abstract—We study distributed algorithms for updating transmit precoding matrices for a two-user Multi-Input/Multi-Output (MIMO) interference channel. Our objective is to maximize the sum rate with linear Minimum Mean Squared Error (MMSE) receivers, treating the interference as additive Gaussian noise. An iterative approach is considered in which given a set of precoding matrices and powers, each receiver announces an interference price (marginal decrease in rate due to an increase in interference) for each received beam, corresponding to a column of the precoding matrix. Given the interference prices from the neighboring receiver, and also knowledge of the appropriate cross-channel matrices, the transmitter can then update the beams and powers to maximize the rate minus the interference cost. Variations on this approach are presented in which beams are added sequentially (and then fixed), and in which all beams and associated powers are adjusted at each iteration. Numerical results are presented, which compare these algorithms with iterative water-filling (which requires no information exchange), and a centralized optimization algorithm, which finds locally optimal solutions. Our results show that the distributed algorithms perform close to the centralized algorithm, and by adapting the rank of the precoder matrices, achieve the optimal high-SNR slope.

I. INTRODUCTION

In both cellular and wireless ad-hoc networks situations commonly arise where two or more point-to-point communication links, consisting of one transmitter and one receiver, interfere with each other. Even for the elementary case of two interfering single-antenna links, however, the information theoretic capacity of the channel and the optimal transmission strategies are unknown, except for some special cases. In general, though, to achieve capacity the receivers must know the codebook used by the interfering transmitter.

In this work, we make the common assumption that the interfering codebook is not available, and that the interference behaves like additional (Gaussian) noise at the receivers, that cannot be distinguished from the thermal noise. Choosing suitable transmit strategies is then a power control problem when all terminals have one antenna. Distributed approaches

for maximizing a sum-utility objective have been presented in [1], [2].

When multiple antennas are available, the terminals have additional degrees of freedom and can avoid interference spatially. The problem of finding near-optimal transmit strategies, however, becomes far more difficult: for the special case of single-antenna receivers (Multi-Input/Single-Output (MISO) interference channel), it reduces to an optimization of the beam direction and can be viewed as a trade-off between an egoistical (matched filter) and an altruistic (forcing the interference to zero) solution [3]. A locally optimal solution with respect to sum utility can again be found in a distributed way by exchanging "interference prices" [4], [5].

For the general case of more than one receive antenna (Multi-Input/Multi-Output (MIMO) interference channel), prior work has so far focused on the non-cooperative strategy of alternatingly maximizing the throughput of the two links (iterative waterfilling) [6]–[8]. While this strategy does not require any communication between the links, the optimal high-SNR slope of the sum throughput, as derived in [9], is not achieved in general. Alternatively, a *centralized* gradient search was proposed in [7], where a locally sum-rate-optimal transmit strategy is found. (See also [10].)

In this paper, we propose algorithms that require a limited amount of information exchange between the links, but still achieve the optimal high-SNR sum rate slope. The key idea is to treat each transmitter's beams separately, and to associate a price with each beam as a measure for how much the utility will decrease per marginal increase in interference after the appropriate receive filter. This way, the individual beams can be updated either successively or iteratively taking into account the effects of the interference caused to the other beams.

In one proposed algorithm (Sequential Beamforming with Adaptive Rank), beams are added successively to both users and then kept fixed with a uniform power allocation, whereas in another algorithm (Full Beam-Power Adaptation) the beams are continuously updated both in terms of direction and beam power. Numerically, the proposed algorithms perform close to a centralized optimization algorithm.

The paper is organized as follows: in Section II, the system

model and the concept of *interference prices* are introduced. In Section III, two distributed algorithms based on exchanging interference prices are proposed. Simulation results are presented in Section IV, and conclusions are given in Section V.

II. SYSTEM MODEL

We consider a wireless network with two interfering, narrowband, time-invariant MIMO links, where each transmitter has N_T antennas and each receiver has N_R antennas. The received signal vector at receiver i, $i = 1, 2$, is given by

$$\mathbf{y}_i = \mathbf{H}_{ii}\mathbf{x}_i + \mathbf{H}_{ij}\mathbf{x}_j + \mathbf{n}_i, \tag{1}$$

where $\mathbf{H}_{ii}$ is the direct channel matrix, $\mathbf{H}_{ij}$ is the cross-channel matrix between transmitter j and receiver i, $\mathbf{x}_i$ is the transmit signal vector for user i, and $\mathbf{n}_i$ is complex Gaussian noise with covariance matrix $\mathbf{R}_{n_i}$. Let $\mathbf{U}_i$ be the precoding matrix for the transmitter i, and $\mathbf{u}_i^k$ its k-th column, which we refer to as the k-th beam of user i. Then, $\mathbf{x}_i = \sum_{k=1}^{N_i} \mathbf{u}_i^k x_i^k$, where x_i^k is the corresponding transmitted symbol,[1] and N_i is the number of beams for user i. The transmit covariance matrix for user i is $\mathbf{Q}_i = \mathrm{E}\{\mathbf{x}_i\mathbf{x}_i^\mathrm{H}\} = \mathbf{U}_i\mathbf{U}_i^\mathrm{H}$, where $(\cdot)^\mathrm{H}$ denotes Hermitian transpose. Assuming the interference is treated as additive Gaussian noise, the achievable rate is given by [11],

$$R_i = \log\det\left(\mathbf{I} + \mathbf{H}_{ii}^\mathrm{H}(\mathbf{R}_{n_i} + \mathbf{H}_{ij}\mathbf{Q}_j\mathbf{H}_{ij}^\mathrm{H})^{-1}\mathbf{H}_{ii}\mathbf{Q}_i\right). \tag{2}$$

Our objective is to maximize the sum rate $R_1 + R_2$ over the covariance matrices $\mathbf{Q}_1$ and $\mathbf{Q}_2$, or equivalently $\mathbf{U}_1$ and $\mathbf{U}_2$, given individual power constraints for both users. For transmitter i, the power is $\mathrm{E}\{\|\mathbf{x}_i\|_2^2\} = \mathrm{Tr}(\mathbf{Q}_i) = \sum_{k=1}^{N_i} \|\mathbf{u}_i^k\|_2^2$, where $\mathrm{Tr}(\cdot)$ is the trace operator. The optimization problem is therefore

$$\max_{\mathbf{Q}_1, \mathbf{Q}_2 \in \mathcal{S}} \quad R_1 + R_2 \tag{P0}$$
$$\text{s.t.} \quad \mathrm{Tr}(\mathbf{Q}_i) \le P_i^{\max} \quad (i = 1, 2)$$

where $\mathcal{S}$ is the cone of positive semi-definite matrices, R_1 and R_2 are given by (2), and $P_i^{\max}$ denotes the power constraint for user i.

To optimize the objective in a distributed manner, each user must know his own rate function, and also the level of interference level caused to the other user. As in [2], [4], [5], [12], the interference to another user can be represented by an *interference price*, which is the marginal decrease in rate for that user due to a marginal increase in interference. The cost of interfering is then the product of the price with the corresponding interference power. This method has been applied to single-input/single-output (SISO) and MISO channels; however, it is not straightforward to extend it to the MIMO case, because it is not clear how a marginal change in the precoding matrix $\mathbf{U}_i$ will influence the other user's rate R_j via (2). However, motivated by prior work for the MISO channel [4], [5], we can treat beams (columns) in $\mathbf{U}_i$

---

[1]We assume all symbols are independent and have unit variance.

separately and apply a linear receive filter for each beam. The filtered signal is then

$$\hat{y}_i^k = \mathbf{g}_i^{k,\mathrm{H}}\left(\mathbf{H}_{ii}\mathbf{u}_i^k x_i^k + \sum_{l \ne k}\mathbf{H}_{ii}\mathbf{u}_i^l x_i^l + \sum_{l=1}^{N_j}\mathbf{H}_{ij}\mathbf{u}_j^l x_j^l + \mathbf{n}_i\right), \tag{3}$$

where $\mathbf{g}_i^k$ is the receive filter, and the second and third terms in the parentheses represent the interference from other beams of the same user and beams from other users, respectively. The linear MMSE receive filter for the k-th beam of user i is

$$\mathbf{g}_i^k = \left(\mathbf{R}_{n_i} + \sum_{l \ne k}\mathbf{H}_{ii}\mathbf{u}_i^l\mathbf{u}_i^{l,\mathrm{H}}\mathbf{H}_{ii}^\mathrm{H} + \sum_{l=1}^{N_j}\mathbf{H}_{ij}\mathbf{u}_j^l\mathbf{u}_j^{l,\mathrm{H}}\mathbf{H}_{ij}^\mathrm{H}\right)^{-1}\mathbf{H}_{ii}\mathbf{u}_i^k. \tag{4}$$

If we decode each beam separately, then the achievable rate for each beam is

$$R_i^k = \log\left(1 + s_i^k/\left(I_i^k + (\sigma_i^k)^2\right)\right) \tag{5}$$

where $s_i^k = |\mathbf{g}_i^{k,\mathrm{H}}\mathbf{H}_{ii}\mathbf{u}_i^k|^2$ is the power of the desired signal, $I_i^k = \sum_{l \ne k}|\mathbf{g}_i^{k,\mathrm{H}}\mathbf{H}_{ii}\mathbf{u}_i^l|^2 + \sum_{l=1}^{N_j}|\mathbf{g}_i^{k,\mathrm{H}}\mathbf{H}_{ij}\mathbf{u}_j^l|^2$ is the total interference power for the k-th beam, and $(\sigma_i^k)^2 = \mathrm{E}\{|\mathbf{g}_i^{k,\mathrm{H}}\mathbf{n}_i|^2\}$ is the noise power. We can then approximate the objective as the rate summed over beams,

$$\max_{\mathbf{u}_1^1\ldots,\mathbf{u}_1^{N_1},\mathbf{u}_2^1,\ldots,\mathbf{u}_2^{N_2}} \sum_{k=1}^{N_1}R_1^k + \sum_{k=1}^{N_2}R_2^k \tag{P1}$$
$$\text{s.t.} \quad \sum_{k=1}^{N_i}\|\mathbf{u}_i^k\|_2^2 \le P_i^{\max} \quad (i = 1, 2).$$

Following [5], for any given set of beams we can now define the interference price for the k-th beam of user i as

$$\pi_i^k = -\frac{\partial R_i^k}{\partial I_i^k} = \frac{1}{(\sigma_i^k)^2 + I_i^k} - \frac{1}{s_i^k + (\sigma_i^k)^2 + I_i^k}. \tag{6}$$

Given the interference prices for all other beams across both users, and assuming all other beams are fixed, user i can select beam $\mathbf{u}_i^k$ to optimize the *best response*, or *payoff*

$$\max_{\mathbf{u}_i^k} \quad R_i^k(\mathbf{u}_i^k) - \sum_{l \ne k}\pi_i^l|\mathbf{g}_i^{l,\mathrm{H}}\mathbf{H}_{ii}\mathbf{u}_i^k|^2 - \sum_{l=1}^{N_j}\pi_j^l|\mathbf{g}_j^{l,\mathrm{H}}\mathbf{H}_{ji}\mathbf{u}_i^k|^2$$
$$\text{s.t.} \quad \|\mathbf{u}_i^k\|_2^2 \le p_i^k \tag{P2}$$

where p_i^k is the predetermined power allocated to this beam. It can be seen that the necessary Karush-Kuhn-Tucker (KKT) conditions [13] for Problem P2 with respect to $\mathbf{u}_i^k$ match the KKT conditions for Problem P1. In other words, the global optimum for Problem P1 is also a locally optimal solution for $\mathbf{u}_i^k$ in Problem P2. Problem P2 for user i can also be viewed as a subproblem of P1, in which the loss in total utility due to user i's interference is replaced by a linear approximation.

To compute the best response update according to P2, transmitter i must know the effective cross-channels $\mathbf{H}_{ji}\mathbf{g}_j^l$, $l = 1, \cdots, N_j$. A distributed pricing algorithm for adjusting a

particular beam $\mathbf{u}_i^k$ consists of alternately solving P2 for the best response update and updating interference prices (across both users). If the receive filters $\mathbf{g}_i^k$ are fixed and viewed as part of the channel, then the MIMO model considered with interference pricing essentially reduces to multiple MISO channels. In [5] some conditions for the convergence of such a distributed algorithm are provided, but analyzing the convergence for the MIMO case considered is more difficult since the receive filters are also updated periodically.

In what follows we describe particular algorithms for updating all beams, which are based on distributed pricing for updating individual beams. Variations depend on the order in which beams are updated and prices are exchanged, and how power is allocated at each iteration. Assuming the algorithm converges, the fixed-point should approximate the solution to Problem P1 and Problem P0.

III. DISTRIBUTED PRICING ALGORITHMS

A. Sequential Beamforming with Adaptive Rank (SBAR)

In this algorithm, users add beams sequentially, i.e., one at a time, until the algorithm terminates. Here we assume that the power for a particular user is distributed uniformly among the beams. More specifically, starting from an initial condition of zero beams each, the users both attempt to add a single beam, corresponding to an iteration in an "outer loop". All beams added in previous iterations are fixed, and power is allocated uniformly over those beams and the new beam. Each user determines the current beam by solving Problem P2 (without considering future beams). The "inner loop" consists of exchanging interference prices and beam updates.

Once the beam updates in the inner loop have converged (or the maximum number of iterations has been reached), each user can independently decide whether to keep or discard the resulting beam according to a performance criterion. Hence the number of beams assigned to the two users in a particular iteration (and at termination) can be different. The algorithm terminates when both users discard the new beams. (Note that a user can choose to discard a beam in one iteration, but add a beam in a succeeding iteration.)

Each newly added beam causes interference to the fixed beams for both users (from previous iterations), so that the previous beams are no longer solutions to the corresponding best response problem P2. Therefore, to reduce the interference, we restrict each new beam to lie in the subspace orthogonal to the previous beams of that user at the channel output. [2] The k-th beam for user i (denoted with index k_i) is therefore determined by solving

$$\max_{\mathbf{u}_i^{k_i}} \quad R_i^{k_i}(\mathbf{u}_i^{k_i}) - \sum_{l=1}^{k_j} \pi_j^l |\mathbf{g}_j^{l,\mathrm{H}} \mathbf{H}_{ji} \mathbf{u}_i^{k_i}|^2 \quad \text{(P2')}$$

$$\text{s.t.} \quad \|\mathbf{u}_i^{k_i}\|_2^2 \le P_i^{\max}/k_i$$

$$\mathbf{g}_i^{l,\mathrm{H}} \mathbf{H}_{ii} \mathbf{u}_i^{k_i} = 0 \quad \forall\, l \in \{1, \ldots, k_i - 1\} \quad (7)$$

where $R_i^{k_i}$ is again given by (5).

An equivalent optimization problem can be obtained by dropping the orthogonality constraints (7) and replacing $\mathbf{u}_i^{k_i}$ by $\mathbf{\Pi}_i^{k_i} \mathbf{u}_i^{k_i}$ in the utility function, where $\mathbf{\Pi}_i^{k_i}$ is the projection matrix defined by

$$\mathbf{\Pi}_i^{k_i} = \mathbf{I} - \mathbf{A}_i^{k_i}(\mathbf{A}_i^{k_i,\mathrm{H}} \mathbf{A}_i^{k_i})^{-1} \mathbf{A}_i^{k_i,\mathrm{H}}$$

with

$$\mathbf{A}_i^{k_i} = \left[\mathbf{H}_{ii}^{\mathrm{H}} \mathbf{g}_i^1, \ldots, \mathbf{H}_{ii}^{\mathrm{H}} \mathbf{g}_i^{(k_i - 1)} \right].$$

The necessary condition for optimality, obtained by taking the derivative of the Lagrangian function, has the form

$$\mathbf{X}_i^{k_i} \mathbf{u}_i^{k_i} = \lambda \mathbf{u}_i^{k_i}, \quad \lambda \ge 0 \quad (8)$$

where

$$\mathbf{X}_i^{k_i} = \frac{\partial R_i^{k_i}}{\partial |\mathbf{g}_i^{k_i,\mathrm{H}} \mathbf{H}_{ii} \mathbf{u}_i^{k_i}|^2} \times \mathbf{\Pi}_i^{k_i} \mathbf{H}_{ii}^{\mathrm{H}} \mathbf{g}_i^{k_i} \mathbf{g}_i^{k_i,\mathrm{H}} \mathbf{H}_{ii} \mathbf{\Pi}_i^{k_i}$$

$$- \sum_{l=1}^{k_j} \pi_j^l \mathbf{\Pi}_i^{k_i} \mathbf{H}_{ji}^{\mathrm{H}} \mathbf{g}_j^l \mathbf{g}_j^{l,\mathrm{H}} \mathbf{H}_{ji} \mathbf{\Pi}_i^{k_i}, \quad (9)$$

which has the form of an eigenvector equation, except that $\mathbf{X}_i^{k_i}$ depends on $\mathbf{u}_i^{k_i}$. If we, however, assume $\mathbf{X}_i^{k_i}$ to be constant by evaluating (9) with the precoder $\mathbf{u}_i^{k_i}$ from the previous iteration, we obtain a suboptimal update for $\mathbf{u}_i^{k_i}$. If $\mathbf{X}_i^{k_i}$ does not have any positive eigenvalues, then the updated precoder is the zero vector, otherwise the updated precoder is the eigenvector corresponding to the dominant eigenvalue of $\mathbf{X}_i^{k_i}$, scaled so that $\|\mathbf{u}_i^{k_i}\|_2^2 = P_i^{\max}/k_i$.

Once the new beams converge to a fixed-point, or the maximum number of iterations in the inner loop is reached, each user decides whether or not to keep the beam. For the numerical results that follow, the user keeps the new beam only if it is nonzero, so that power is not wasted on an all-zero beam. Although quite simple, numerical results have shown that this criterion performs as well as evaluating more complicated metrics.

If the receive filter $\mathbf{g}_i^{k_i}$ is fixed in the inner loop, then the distributed pricing algorithm is the same as the the MISO-ADP algorithm presented in [5], which has been proven to converge from corner points of an appropriately defined constraint region. Changing the receive filters $\mathbf{g}_i^{k_i}$, however, greatly complicates the convergence analysis. Simulation results show reliable convergence at low SNRs[3] (cf. Section IV).

The SBAR algorithm is summarized as follows:

1) Start with zero beams for each user.
2) In the k-th iteration of the outer loop initialize the k_i-th beam ($k_i \le k$) for each user $i = 1, 2$ by selecting its direction and splitting the power P_i^{max} uniformly over the k_i beams. Then, based on the current profile, calculate the interference prices for all beams and the

---

[2]Another possibility is to enforce orthogonality to all previous beams of *both* users; however, that generally does not perform as well due to the loss in degrees of freedom.

[3]We define the SNR as P_i^{max}/n_0, where n_0 is the noise power at each receive antenna.

MMSE filter for the k_i-th beam from (6) and (4), respectively. [4]

3) Adapt the new beams in the inner loop according to the distributed pricing algorithm:

  3.a Compute the k_1-st beam for user 1 by computing the top eigenvector of the matrix $\mathbf{X}_1^{k_1}$ from (9) with $\mathbf{u}_1^{k_1}$ from the preceding (inner loop) iteration.

  3.b Update user 1's interference prices π_1^k for all beams $k = 1, \cdots, k_1$, and the receive filter for the k_1-th beam.

  3.c Adapt beam k_2 for user 2 as in step 3.a.

  3.d Update prices π_2^k, $k = 1, \cdots, k_2$, and the receive filter for the k_2-th beam.

  3.e Repeat from step 3.a until convergence.

4) If beam k_i is zero, discard the result obtained in the k-th iteration of the outer loop. Otherwise, keep the beam and proceed to the next iteration of the outer loop in which each user i optimizes a new beam. The outer loop terminates when both users discard the new beams.

5) To achieve the rate given by (2) with $\mathbf{Q}_i = \mathbf{U}_i \mathbf{U}_i^{\mathrm{H}}$, either the data streams associated with all beams for user i must be decoded jointly at the receiver, or the precoder and receiver must be redesigned to diagonalize the whitened channel.

B. Full Beam-Power Adaptation (FBPA)

In this algorithm each user is assigned an initial set of beams, all of which are optimized in each iteration (i.e., once assigned, beams are not fixed as in the SBAR algorithm). In addition, the power allocation over the beams is also optimized. We expect this algorithm to perform better than the preceding SBAR algorithm, although the additional optimizations incur substantial additional complexity and iterations (information exchange) for convergence.

Given an initial choice of beams and powers across the beams, the beam directions are optimized sequentially. That is, fixing all other beams and powers across both users, the beam $\mathbf{u}_i^k$ is determined by solving Problem P2 as in the SBAR algorithm. After all beams for one user have been updated, the beam powers for this user are adjusted. This is equivalent to the power optimization problem with SISO channels. Following the algorithm proposed in [2], the powers for user i are determined solving the subproblem

$$\max_{\mathbf{p}_i} \quad \sum_{l=1}^{k_i} R_i^l \left(\frac{\mathbf{u}_i^l p_i^l}{\|\mathbf{u}_i^l\|_2^2} \right) - \sum_{m=1}^{k_j} \sum_{l=1}^{k_i} \pi_j^m p_i^l \frac{|\mathbf{g}_j^{m,\mathrm{H}} \mathbf{H}_{ji} \mathbf{u}_i^l|^2}{\|\mathbf{u}_i^l\|_2^2}$$

$$\text{s.t.} \quad \sum_{l=1}^{k_i} p_i^l \leq P_i^{\max} \qquad \text{(P3)}$$

where $\mathbf{p}_i = \{p_i^1, \ldots, p_i^{k_i}\}$ is the power profile for user i across the assigned k_i beams. Then, given the updated power profile, each user can re-optimize the beams. The algorithm is summarized as follows:

---

[4]The MMSE filters for all previous beams are not updated to reduce the amount of information exchanged.

1) The algorithm can be initialized with an arbitrary pair of precoding matrices $\mathbf{U}_1$ and $\mathbf{U}_2$. The interference prices and receive filters are determined from (6) and (4).

2) For user 1, adapt each beam sequentially by solving Problem P2 given the fixed power profile. Once a beam is updated, all interference prices and receive filters are re-computed.

3) User 1 updates his power profile by solving Problem P3. All interference prices and receive filters are then re-computed.

4) User 2 updates the beams and power allocation as in steps (2) and (3).

5) Repeat from step 2 until convergence.

The numerical results show that the algorithm can improve the sum rate compared to the initial value, and the average performance depends on the choice of the starting point. The FBPA algorithm converges more slowly than the SBAR algorithm. Therefore, in order to speed up convergence and yield a higher sum rate, we can combine these two algorithms by using the outcome of the SBAR algorithm as the initialization for the FBPA algorithm. This initialization is used to generate the following numerical results.

C. Extensions to Iterative Waterfilling

We also evaluate a simpler algorithm, in which the normalized *beams* are determined by IWF, but the *powers* across beams are then adapted by solving Problem P3. This is referred to as *iterative waterfilling with power pricing (IWF/PP)*.

The receive filters depend on both the beams and the powers, and should therefore be updated once a change in those variables occur. The information exchange for the IWF/PP algorithm can be further reduced by fixing all receive filters at their initial values (or at least not reporting updates). Relative to the SBAR and FBPA algorithms, this requires less information exchange, but is expected to incur a performance loss. Still, this can improve upon the performance of IWF.

IV. SIMULATION RESULTS

In this section, we discuss a typical performance plot for the previously proposed algorithms (SBAR, FBPA, and IWF/PP), and also compare them with the asynchronous IWF algorithm, as well as the result from a standard nonlinear optimization package (MATLAB) applied to Problem P0. The IWF algorithm is the best we can do without any information exchange, while the centralized solver requires knowledge of all channel matrices to yield a sum rate that is at least locally optimal.

We consider a two-user system with 4 transmit antennas and 4 receive antennas for each user. The direct channels and the cross-channels are complex Gaussian with the same variance. In Fig. 1, we show the sum rate performance of the algorithms versus SNR in dB (defined as in Footnote 3) averaged over 400 channel realizations. We observe that all distributed algorithms perform significantly better than the non-cooperative IWF algorithm, especially in the high-SNR setting, and that these algorithms appear to achieve the optimal
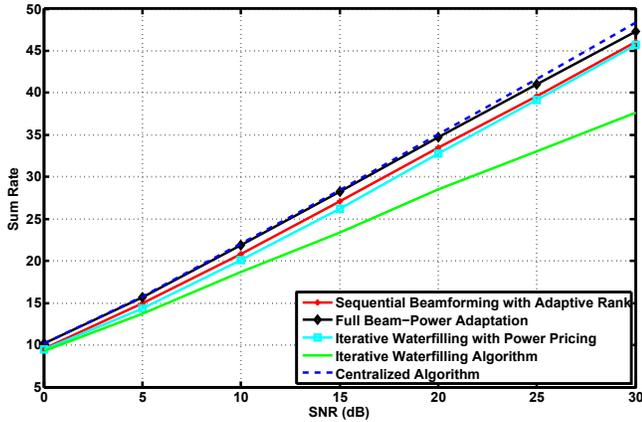
Fig. 1. Illustration of the average performance of all proposed distributed algorithms.

Alg.	0 dB	5 dB	10 dB	15 dB	20 dB	25 dB	30 dB
SBAR	400	400	400	399	387	357	256
FBPA	400	400	400	398	378	315	211
IWF/PP	400	400	400	400	400	400	400
IWF	400	397	393	388	393	389	396

high-SNR slope. Furthermore, it is also verified (but not shown in the figure) that if only exchanging the interference prices is allowed with fixed receive filters, we can still improve the sum rate by adapting powers. The performance is slightly worse than the IWF/PP algorithm, and the gain compared with the IWF algorithm is more than 10% when SNR is high ($\geq 20 dB$).

The gain in sum rate performance is accompanied by an increase in information exchange between the users. While IWF requires no communication between the users, for the three distributed algorithms each time a beam and a receive filter is updated, the scalar price and the filter vector should be communicated to the other user. The total amount of information exchange therefore depends on the iterations needed for convergence. The FBPA algorithm generally requires more iterations and information exchange than the SBAR algorithm, and offers a (small) performance improvement.

The simulations have shown that the algorithms do not always converge. Table I shows the number of trials (channel realizations) for which each algorithm terminates according to the convergence criterion, as opposed to reaching the maximum number of iterations (in our case, 500 updates for each beam or power). Specifically, convergence is assumed to occur when the difference in sum rate between two consecutive iterations is less than a threshold (e. g., 0.1% in our simulation). The table shows that convergence becomes more problematic at high SNRs for all algorithms, and that the FBPA is the least likely to converge, followed by the SBAR algorithm and IWF. (The entries for the IWF/PP algorithm are for the power updates only. The beams are fixed according to the outcome of IWF.) Although convergence is not guaranteed, in all cases observed the fluctuations in sum rate, which occur due to lack of convergence, are quite small (but noticeable).

V. CONCLUSIONS

We have presented distributed algorithms for selecting the users' precoding matrices to maximize the sum rate for the two-user MIMO interference channel. Each beam (column of the precoding matrix) is adapted separately to maximize a

payoff function, given a set of interference prices for received beams at the neighboring receiver. In the SBAR algorithm each user attempts to add beams sequentially, whereas the FBPA algorithm adapts an initial set of beams along with the corresponding power allocation. Numerically, our proposed distributed algorithms perform significantly better than the asynchronous IWF algorithm and in some cases (e. g., a "Z"-channel in which only one cross-channel matrix has large gains, and the other has small gains) perform as well as the centralized numerical solver.

Determining analytical conditions for convergence of the distributed algorithms remains an open problem. Although those appear to be difficult to obtain, due to the complexity of the updates, it may be possible to gain insight by examining some special cases. Another interesting topic for further work is how to reduce the amount of information exchange between the users without significantly sacrificing sum rate performance.

REFERENCES

[1] M. Chiang, "Balancing Transport and Physical Layers in Wireless Multihop Networks: Jointly Optimal Congestion Control and Power Control," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, January 2005.
[2] J. Huang, R. A. Berry, and M. L. Honig, "Distributed Interference Compensation for Wireless Networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 5, pp. 1074-1084, May 2006.
[3] E. Larsson and E. Jorswieck, "Competition Versus Cooperation on the MISO Interference Channel," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 7, pp. 1059-1069, Sept. 2008.
[4] D. A. Schmidt, A. Gründinger, W. Utschick, and M. L. Honig, "Distributed Precoder Optimization for Interfering MISO Channels," in *Int. ITG Workshop on Smart Antennas*, Feb. 2008, pp. 285-290.
[5] C. Shi, R. A. Berry, and M. L. Honig, "Distributed Interference Pricing with MISO Channels," in *Proc. Allerton Conference*, Sept. 2008.
[6] M. F. Demirkol and M. A. Ingram, "Power-Controlled Capacity for Interfering MIMO Links," *Vehicular Technology Conference*, 2001. VTC 2001 Fall, vol 1, pp. 187–191.
[7] S. Ye and R. S. Blum, "Optimized Signaling for MIMO Interference Systems with Feedback," *IEEE Transactions on Signal Processing*, Volume 51, Issue 11, Nov 2003, pp. 2839-2848.
[8] G. Scutari, D. P. Palomar and S. Barbarossa, "Competitive Design of Multiuser MIMO Systems based on Game Theory: A Unified View," *IEEE J. Sel. Areas Commun.*, Sept. 2008.
[9] S. A. Jafar and M. J. Fakhereddin, "Degrees of Freedom for the MIMO Interference Channel," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, July 2006.
[10] K. Gomadam, V. R. Cadambe and S. A. Jafar, "Approaching the Capacity of Wireless Networks through Distributed Interference Alignment," in *arXiv:0803.3816v1*, Mar. 2008.
[11] E. Telatar, "Capacity of Multi-Antenna Gaussian Channels," *Eur. Trans. Telecomm. ETT*, vol. 10, no. 6, pp. 585-596, Nov. 1999.
[12] C. Shi, R. A. Berry, and M. L. Honig, "Distributed Interference Pricing for OFDM Wireless Networks with Non-Separable Utilities," in *Proc. Conference on Information Sciences and Systems*, March 2008.
[13] D. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.