Local Interference Pricing for Distributed Beamforming in MIMO Networks

Changxin Shi, Randall A. Berry, and Michael L. Honig
Department of Electrical Engineering and Computer Science
Northwestern University, Evanston, Illinois 60208
Email: cshi@u.northwestern.edu, {rberry, mh}@eecs.northwestern.edu

***Abstract*—We study a distributed algorithm for adjusting beamforming vectors in a peer-to-peer wireless network with multiple-input multiple-output (MIMO) channels. Each transmitter precoding matrix has rank one, and a linear minimum mean squared error (MMSE) filter is applied at each receiver. Our objective is to maximize the total utility summed over all users, where each user's utility is a function of the received signal-to-interference-plus-noise ratio (SINR). Given all users' beamforming vectors and receive filters, each receiver announces an interference price, representing the marginal cost of interference from other users. A particular transmitter updates its beamforming vector to maximize its utility minus the interference cost to other users. We show that if the utility functions satisfy certain concavity conditions, then the total utility is non-decreasing with each update. We also present numerical results that illustrate the effect of ignoring interference prices from all but the closest users, and relaxing requirements on the frequency of beam and price updates.**

I. Introduction

Achieving high spectral efficiencies in multiuser wireless networks (e.g., cellular and wireless ad-hoc networks) depends critically on the application of interference mitigation techniques. When users in the network are equipped with multiple transmitters and receivers, the additional spatial degrees of freedom can be exploited to reduce interference. Joint optimization of power and spatial beams becomes especially challenging in peer-to-peer networks without centralized resource management, since an optimal allocation of resources at a particular transmitter (e.g., that maximizes total rate) generally requires information about other nodes.

We consider a peer-to-peer network with multiple-input multiple-output (MIMO) links. The performance of each transmitter-receiver pair is measured by a utility function, which is a function of the transmission rate and depends on the transmitter's precoding matrix and the received interference-plus-noise covariance matrix. Our objective is to maximize the total (sum) utility over all users by applying a distributed algorithm to adjust the users' precoding matrices. We assume that each user's precoding matrix has rank one, i.e., the precoding matrix is a beamforming vector. This case has practical significance, and simplifies the optimization problem.

The algorithm we study is motivated by the work in [1]–[3], which considers algorithms for power allocation in single-antenna wireless networks and beam assignments in multi-input single-output (MISO) networks. The *asynchronous distributed pricing (ADP)* algorithm was presented in [1] for single-antenna wireless networks. In that algorithm, each user announces an *interference price* to all other users, which is the user's marginal change in utility per unit interference power. Given the interference prices from all interfering users, each user updates his power and beam by optimizing his utility minus the interference cost to other users, which is determined from their announced interference prices. The ADP algorithm is extended to beamformer updates for MIMO channels in [3].

It is shown in [1] that the ADP algorithm converges for a suitable class of utility functions. The proof is based on relating the updates in the distributed algorithm to best response updates in a supermodular game. In [2], we present a different approach to the convergence analysis, which is based on the convexity of the utility functions with respect to the received interference power. Monotonic convergence of the sum utility is shown for a broader class of utility functions provided that interference prices are updated after every power/beam update. Similar distributed algorithms are studied for MIMO networks in [3], which have no restriction on the rank of the precoding matrix. These algorithms are observed to perform well numerically, although finding analytical conditions for convergence remains an open problem.

Here we again treat interference as noise (as in [1]–[3], and optimize the total utility by selecting an optimal beamforming vector for each transmitter, assuming a linear minimum mean squared error (MMSE) receive filter for each

This work was supported by ARO under grant W911NF-07-1-0028.

associated receiver. Unlike the MISO network considered in [2], the receiver MMSE filters must be iteratively updated along with beams. Based on the convergence analysis in [2], it is straightforward to establish convergence with an appropriate update rule.

We also numerically study the effect of limiting the information exchange. Specifically, we consider the scenarios in which users update their beamformers, ignoring interference prices and cross channel information from all but the closest (local) users. We also study the performance with infrequent price and receiver updates.

In terms of related work, the performance of the ADP algorithm and the effect of limiting the amount of information exchange for single-antenna links are studied through simulations in [4]. The Pareto-optimal rate-pair for a two-user MISO interference channel is characterized in [5]. For resource allocation techniques in MIMO networks, iterative waterfilling has been extensively studied (e.g., [6]–[8]), in which nodes implement best response updates without exchanging channel or interference information. However, iterative waterfilling is suboptimal in terms of the overall network performance.

In the next section, we present system model and the resource optimization problem. In Section III, we propose a distributed algorithm and consider some variations with reduced information exchange. Simulation results are presented in Section IV, and conclusions are given in Section V.

II. System Model

We consider a time-invariant wireless network with K pairs of transmitters and receivers, where each transmitter has N_T antennas and each receiver has N_R antennas. Given channel matrices $\mathbf{H}_{ik}$'s, representing the channel from transmitter k to receiver i, the received signal at receiver k is given by

$$\mathbf{y}_k = \mathbf{H}_{kk}\mathbf{x}_k + \sum_{i \neq k} \mathbf{H}_{ki}\mathbf{x}_i + \mathbf{n}_k \tag{1}$$

where $\mathbf{x}_k$ is the transmit signal vector for user k and $\mathbf{n}_k$ is additive complex Gaussian noise with covariance matrix $\mathbf{R}_{n_k}$. Assuming the interference is treated as additive Gaussian noise, the achievable rate for user k is given by [9],

$$R_k = \log\det\left(\mathbf{I} + \mathbf{H}_{kk}^{\mathrm{H}}\Big(\mathbf{R}_{n_k} + \sum_{i \neq k} \mathbf{H}_{ki}\mathbf{Q}_i\mathbf{H}_{ki}^{\mathrm{H}}\Big)^{-1}\mathbf{H}_{kk}\mathbf{Q}_k\right) \tag{2}$$

where $\mathbf{Q}_k = \mathrm{E}\left[\mathbf{x}_k\mathbf{x}_k^{\mathrm{H}}\right]$ is the transmit covariance matrix for user k and $(\cdot)^{\mathrm{H}}$ denotes Hermitian transpose. Generally, $\mathbf{Q}_k$ is determined by the precoding matrix $\mathbf{V}_k$ at the transmitter k. In this paper, we assume that each precoding matrix $\mathbf{V}_k$ is rank one, i.e., $\mathbf{V}_k = \mathbf{v}_k$ is a vector and $\mathbf{x}_k = \mathbf{v}_k x_k$ where x_k is a single transmitted symbol with $\mathrm{E}\left[|x_k|^2\right] = 1$.

With an optimized beamforming vector, the capacity in (2) can be achieved with a linear MMSE filter, given by

$$\mathbf{g}_k = \Big(\mathbf{R}_{n_k} + \sum_{i \neq k} \mathbf{H}_{ki}\mathbf{v}_i\mathbf{v}_i^{\mathrm{H}}\mathbf{H}_{ki}^{\mathrm{H}}\Big)^{-1}\mathbf{H}_{kk}\mathbf{v}_k. \tag{3}$$

The filtered signal is then

$$\hat{y}_k = \mathbf{g}_k^{\mathrm{H}}\Big(\mathbf{H}_{kk}\mathbf{v}_k x_k + \sum_{i \neq k} \mathbf{H}_{ki}\mathbf{v}_i x_i + \mathbf{n}_i\Big), \tag{4}$$

and the corresponding signal-to-interference-plus-noise ratio (SINR) for user k can be written as

$$\gamma_k = (\mathbf{H}_{kk}\mathbf{v}_k)^{\mathrm{H}}\Big(\mathbf{R}_{n_i} + \sum_{i \neq k} \mathbf{H}_{ki}\mathbf{v}_i\mathbf{v}_i^{\mathrm{H}}\mathbf{H}_{ki}^{\mathrm{H}}\Big)^{-1}\mathbf{H}_{kk}\mathbf{v}_k \tag{5}$$

The quality of service for each user is measured by a utility function of the transmission rate. Equivalently, we define the utility $u_k(\gamma_k)$ as a function of γ_k. As usual, we assume the utility function of SINR is monotonically increasing, concave and twice differentiable. Our objective is to choose the beamforming vector $\mathbf{u}_k$ for each user k to maximize the utility summed over all users, i.e.,

$$\begin{aligned} \max_{\mathbf{v}_1,\cdots,\mathbf{v}_K} \quad & \sum_{k=1}^{K} u_k(\gamma_k) \\ \text{s.t.} \quad & \|\mathbf{v}_k\|^2 \leq P_k^{max}, \quad k = 1, \ldots, K, \end{aligned} \tag{P}$$

where γ_k is determined by the beamforming vectors via (5), and P_k^{max} denotes the power constraint for user k. The nonlinear problem P is hard to solve in general. What is more challenging here is that we seek a distributed algorithm in which users do not know the entire network topology and other users' utility functions, yet the users' beamforming vectors are still guided to the optimal vectors with some limited information exchange among users.

A *distributed pricing (MISO-DP)* algorithm is studied in [2] for MISO networks with guaranteed convergence for a wide class of utility functions. The key idea behind the MISO-DP algorithm is that users exchange *interference prices*, which reflect the marginal change in their utility per unit interference power. A particular user updates his beamforming vector to maximize his utility minus the interference cost to other users, which is determined from their announced interference prices. Furthermore, some similar numerical algorithms for MIMO networks are presented in [3] with the same beamformer update strategy. The differences between our algorithm and those in [3] are specified in the next section.

III. Distributed Pricing Algorithm

A. MIMO Distributed Pricing (MIMO-DP) Algorithm

From (5), it is not clear what is the "interference power", and how a change in interference will influence the corresponding SINR and utility. However, if we fix each receiver filter, the MIMO network reduces to the MISO case in which the channel vector from transmitter k to receiver i is $\mathbf{h}_{ik}^{\mathrm{H}} = \mathbf{g}_i^{\mathrm{H}}\mathbf{H}_{ik}$. Then, given a set of receive filters $\{\mathbf{g}_1, \mathbf{g}_2, \cdots, \mathbf{g}_K\}$, we can re-cast the MISO-DP algorithm in [2] to solve Problem P with

$$\gamma_k = \frac{|\mathbf{g}_k^{\mathrm{H}}\mathbf{H}_{kk}\mathbf{v}_k|^2}{|\mathbf{g}_k^{\mathrm{H}}\mathbf{n}_k|^2 + \sum_{i\neq k}|\mathbf{g}_k^{\mathrm{H}}\mathbf{H}_{ki}^{\mathrm{H}}\mathbf{v}_i|^2}. \tag{6}$$

Define the interference price for user k, given $\mathbf{g}_k$, as

$$\pi_k = -\frac{\partial u_k(\gamma_k)}{\partial I_k} \tag{7}$$

where $I_k = \sum_{i\neq k}|\mathbf{g}_k^{\mathrm{H}}\mathbf{H}_{ki}\mathbf{v}_i|^2$ is the interference power after the receive filter. Given fixed interference prices and the receive filters and beamforming vectors for the other users, transmitter k then updates its beamforming vector by solving the subproblem:

$$\begin{aligned} \max_{\mathbf{v}_k} \quad & u_k(\gamma_k(\mathbf{v}_k;\mathbf{v}_{-k})) - \sum_{i\neq k}\pi_i|\mathbf{g}_i^{\mathrm{H}}\mathbf{H}_{ik}\mathbf{v}_k|^2 \\ \text{s.t.} \quad & \|\mathbf{v}_k\|^2 \le P_k^{max}, \end{aligned} \tag{P_k}$$

where $\mathbf{v}_{-k}$ denotes the beamforming vectors of users other than user k.

If the users repeatedly update their beamforming vectors and interference prices according to P_k and 7, as in [2], then the MISO-DP algorithm will converge. Hence this gives (locally) optimal beamforming vectors, if we start with the optimal receive filters. However, the receiver updates must now be included in the analysis.

We formally state the distributed algorithm, referred to as the *MIMO distributed pricing (MIMO-DP)* algorithm, as follows:

1) *Initialization*: Each user k chooses an initial beamforming vector $\mathbf{v}_k$ satisfying the power constraint, and applies a corresponding MMSE filter at the receiver by (3).
2) *Price Update*: Using (7), each receiver k calculates the interference price π_k given the current beamforming vectors $\mathbf{v}_k$ and receive filters $\mathbf{g}_k$, and announces this price to every other user.
3) *Beamformer Update*: One random user k solves Problem P_k and updates his beamforming vector, given the interference prices $\{\pi_i\}_{i\neq k}$.
4) *Receive Filter Update*: Any set of users update their receive filter using (3), given beamforming vectors $\mathbf{u}_k$ and announce to other users.
5) Repeat from step 2).

In general, there may be multiple solutions to Problem P_k, which gives the beamformer update. In that case, one of the solutions can be randomly chosen assuming the previous beam is not a solution. Otherwise, the previous beam is kept. The MIMO-DP algorithm assumes that each user k knows the product $\mathbf{H}_{ik}^{\mathrm{H}}\mathbf{g}_i$ for all i, in addition to local information, namely, his own utility function and the interference-plus-noise power at the output of $\mathbf{g}_k$. Hence, receiver k must announce $\mathbf{g}_k$ after each update. Later we show results illustrating the performance if each transmitter uses the initial values of $\mathbf{g}_k$'s in P_k, with no successive update or infrequent updates. Note that each transmitter does not need to know the other users' beamforming vectors or the channel matrices $\mathbf{H}_{ij}$ for $j \neq k$ and all i.

B. Convergence

Although the idea of the beamformer update in the MIMO-DP algorithm is the same as that in the algorithms in [3], they are still evidently different. Specifically, those algorithms in [3] involve beams and associated powers adaptation, while here we assume the precoding matrix are rank one.

As in [2], it can be shown that if the utility function satisfies that $-\frac{u''(\gamma)\gamma}{u'(\gamma)} \in [0,2]$, then it is convex with respect to the received interference power. If all interference prices are current (i.e., have been updated since the last beamformer update), it can then be shown that a subsequent beamformer update cannot decrease the total utility. Also, updating the receive filter $\mathbf{g}_k$ further increases user k's SINR and utility without changing others' utilities. Hence the total utility must monotonically converge to a limit with beam and receiver updates. This is summarized in the following proposition

Proposition 1: If for each user k, $-\frac{u_k''(\gamma_k)\gamma_k}{u_k'(\gamma_k)} \in [0,2]$ for all feasible γ_k, then the MIMO-DP algorithm converges to a stationary point, which satisfies the Karush-Kuhn-Tucker (KKT) conditions of Problem P with MMSE receive filters fixed.

C. Solving P_k

In the MIMO-DP algorithm, each user must solve the subproblem P_k repeatedly, which is a nonlinear optimization problem. Since this may be time-consuming for large networks, we present an efficient numerical algorithm for solving Problem P_k.

The KKT conditions for Problem P_k for user k are

$$\underbrace{\left[a_k(\mathbf{v}_k)\mathbf{H}_{kk}^{\mathrm{H}}\mathbf{g}_k\mathbf{g}_k^{\mathrm{H}}\mathbf{H}_{kk} - \sum_{i\neq k}\pi_i\mathbf{H}_{ik}^{\mathrm{H}}\mathbf{g}_i\mathbf{g}_i^{\mathrm{H}}\mathbf{H}_{ik}\right]}_{\mathbf{X}_k}\mathbf{v}_k = \lambda_k\mathbf{v}_k, \tag{8}$$

where

$$a_k(\mathbf{v}_k) = \frac{u'_k(\gamma_k)}{|\mathbf{g}_k^{\mathrm{H}}\mathbf{n}_k|^2 + \sum_{i\neq k}|\mathbf{g}_k^{\mathrm{H}}\mathbf{H}_{ki}\mathbf{v}_i|^2},$$

and λ_k is the Lagrange multiplier associated with the power constraint. In general, a_k is a function of $\mathbf{v}_k$. However, if the utility function is linear in the SINR, then the optimal beamforming vector $\mathbf{v}_k^*$ is either the eigenvector of the matrix $\mathbf{X}_k$ in (8) corresponding to its largest positive eigenvalue with an appropriate scale factor if the largest eigenvalue is positive, or is the zero vector if all eigenvalues are negative.

With this observation, we start with any feasible beamforming vector $\mathbf{v}_k^0$, and treat the coefficient a_k as a constant. We then compute $\mathbf{v}_k^1$ by finding the appropriate eigenvector of $\mathbf{X}_k$. Iterating we get the sequence $\{\mathbf{v}_k^0, \mathbf{v}_k^1, \cdots, \mathbf{v}_k^n, \cdots\}$. If it converges, then the limit point satisfies the KKT conditions. Although convergence is not guaranteed, the algorithm works well in our simulations.

D. Localized Distributed Pricing Algorithm

The MIMO-DP algorithm provides an efficient approach to optimize network performance. However, each user requires knowledge of cross channels and interference prices, which might incur large overhead. In addition, this overhead may provide little benefit for relatively weak (far away) users. On the other hand, to *guarantee* the convergence, we require users to follow the update rules. In what follows, we study the following modifications to the MIMO-DP algorithm.

1) Limited Information Exchange: In practice, we can modify the MIMO-DP algorithm by taking into account the interference cost to a limited number of interfering users in Problem P_k (those that are relatively close). The criterion for making this reduction depends on the cost of exchanging prices. Obviously, keeping all channel and price information gives the best performance. In Section IV, we examine the performance with different scenarios.

2) Localized Update Rule: The convergence proof for the MIMO-DP algorithm relies on the fact that whenever one user updates his beamformer, he has current interference prices from every user in the network. To reduce the overhead required to broadcast these prices, it is desirable to relax this restriction and allow multiple users to update their beamformers before new prices are announced, i.e., update the interference prices infrequently.

In [2] it is observed that simultaneous power updates can cause oscillations in some users' powers. Hence we relax the preceding restriction by allowing some users to update simultaneously provided that those users are far enough away, i.e., the update rule is satisfied locally. In other words, within a threshold distance from any updating user, there can be no other user updating his beam. Although there is no analytical proof to guarantee convergence, the MIMO-DP algorithm still converges in all cases we simulated for the rate utility.

IV. Simulation Results

In this section, we first show a typical convergence plot for the MIMO-DP algorithm in Section III-A. Then, the performance of algorithms with the variations presented in Section III-D is shown and discussed.

A. MIMO Distributed Pricing Algorithm

We consider a MIMO network of 5 users with 3 transmitter antennas and 3 receive antennas for each user, randomly placed in a square of 100m x 100m. The direct and cross channel matrices have *iid* complex Gaussian entries with each variance determined by the distance attenuation. Namely, the average gain (received power) of each single link is $\sigma^2(d) = \sigma_0^2(\frac{d}{d_0})^{-4}$, where $d_0 = 10$m is the reference distance and also the minimum distance allowed between any transmitter and receiver, and σ_0^2 is the reference power. The maximum power, noise, and σ_0 are selected so that the expected received signal-to-noise ratio (SNR) of a single link with a separation of 10m is 100 (20 dB). The rate utility function, $u_k(\gamma_k) = \log(1 + \gamma_k)$ is assumed for all users.

Fig. 1 shows sum rate versus number of iterations for a particular model realization, achieved by applying the MIMO-DP algorithm. We also examine the scenario in which receive filter updates are announced only every n iterations (even though the filters are updated after every beam update). In other words, the transmitter uses obsolete information for the receive filters when updating its beamformers. The algorithm starts from an arbitrary selection of beamformers, and converges to a stationary point, predicted by Proposition 1. To check whether the limit point is indeed optimal, we use MATLAB to solve the global optimization problem P directly, starting from the limit point achieved by the MIMO-DP algorithm. The total utility obtained in this way is indicated by the dash-dot line in Fig. 1, which matches the limit point of the MIMO-DP algorithm. Admittedly, the more frequently the receive filter updates are announced, the faster the algorithm converges.

B. Effects of Limited Information Exchange

Here, we consider the scenario where users ignore prices announced by remote users.

Similar to the previous subsection, now we consider a MIMO network of N users with 3 transmitter antennas and 3 receive antennas for each user, in which each transmitter and the associated receiver are randomly placed at nodes in a grid of 200m x 200m formed from minimum squares

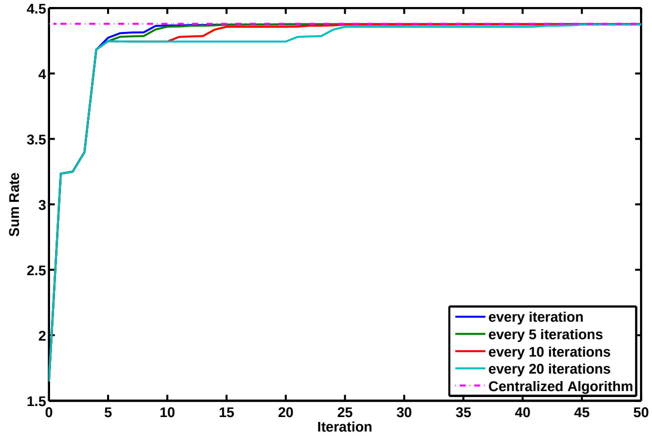

Fig. 1. Illustration of the convergence of the MIMO-DP algorithm with different frequencies for announcing receive filter updates.

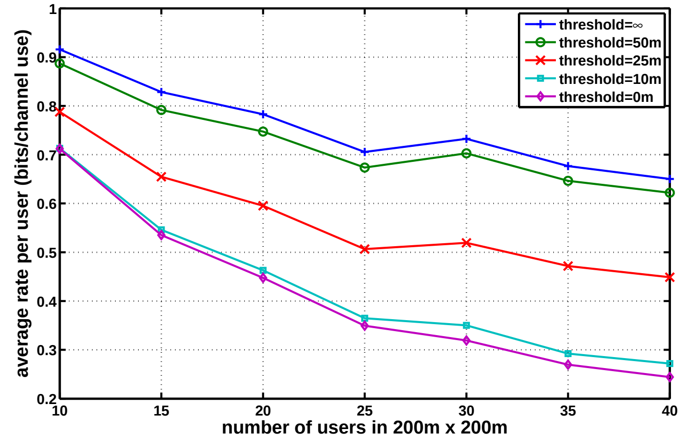

Fig. 2. Illustration of the effect of ignoring interference prices from remote users (receivers are uniformly distributed).

of 10m x 10m according to a uniform distribution. Channel matrices are configured the same way as before with the reference SNR at 10m set to be 30 dB.

We still assume each user k knows the channel matrices $\mathbf{H}_{ik}$ for all i. However, when each user updates his beamformer by solving Problem P_k, he ignores the interference cost to user i if $y_{ki} = \mathbf{1}^{\mathrm{H}}\mathbf{H}_{ik}^{\mathrm{H}}\mathbf{H}_{ik}\mathbf{1}$ is smaller than some threshold, where $\mathbf{1}$ is an all-one vector. The reason to choose this quantity as the criterion is as follows. Consider the feedback link from receiver i to transmitter k for broadcasting price π_i. If we assume the channel is reciprocal and receiver i uses the all-one vector as its beamforming vector, from (5), the received SINR at transmitter k is determined by y_{ki}, assuming the interference and noise covariance matrix $\mathbf{R}_n$ is a scaled identity matrix. Hence, this quantity should be sufficiently large to enable reliable reception of π_i.

Fig. 2 and Fig. 3 show the sum rate performance versus number of users (equivalently density) with different thresholds, averaged over 100 channel realizations. The difference is that all receivers are uniformly distributed as mentioned above in Fig. 2, while the associated receiver in Fig. 3 is randomly placed at some node within a 100m x 100m square centered around its transmitter. To interpret these results in terms of local interference, we convert the threshold with respect to channel gain into a threshold with respect to distance. Specifically, the distance threshold means a single-antenna link with this separation without considering Rayleigh fading will have the same received SNR as the MIMO link using the original threshold for y_{ki}. This comparison is reasonable because the reliability of broadcasting the price through such a MIMO link with the quantity of y_{ki} smaller than the original threshold is on average the same as that of receiving whatever information

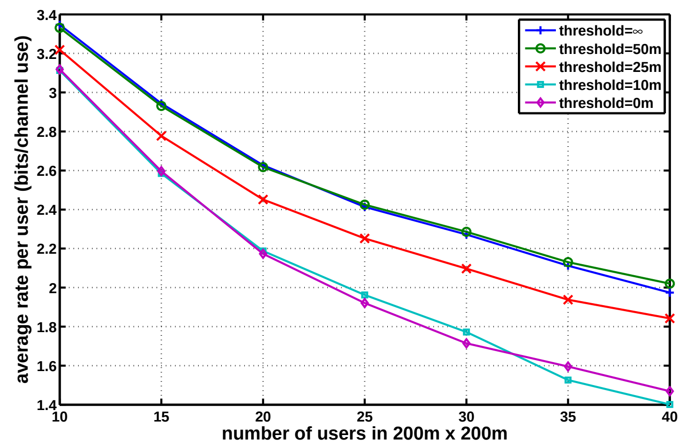

Fig. 3. Illustration of the effect of ignoring interference prices from remote users (each receiver is near its transmitter).

transmitted from a distance greater than the equivalent distance threshold through a single-antenna link. In these figures, "threshold $= \infty$" corresponds to the original MIMO-DP algorithm, in which all interfering users' prices are included, while "threshold $=$ 0m" corresponds to an algorithm without information exchange: each user maximizing his own utility without considering any interference cost. From the figures, it is shown that a reasonably small threshold (50m) does not hurt the performance significantly. Furthermore, although the average rate per user is higher when each receiver is located near its transmitter (Fig. 3), the gain of price exchanges compared to the myopic strategy without information exchange is more significant in the case that receivers are uniformly distributed (Fig. 2). This is because, when each receiver is located near its transmitter, interference becomes equivalently less dominant.

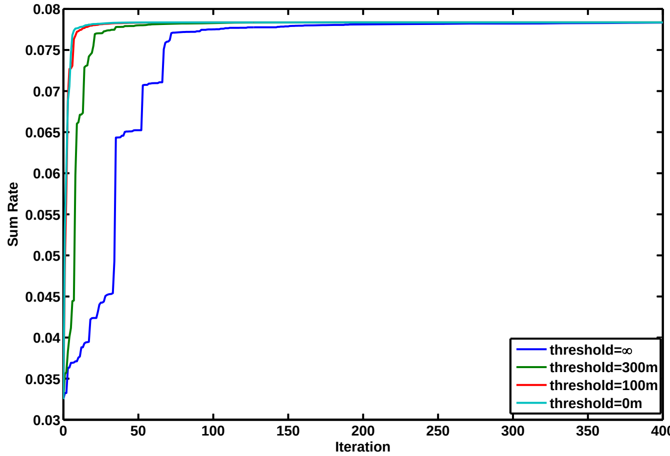

Fig. 4. Illustration of the effect of allowing faraway users to update beamformers simultaneously.

C. Localized MIMO-DP Algorithm

We now require that there are no simultaneous beamformer updates only within a distance smaller than some threshold. Hence multiple users can update at the same time as long as they are far enough away. We simulate a network of 20 users within a square of 500m x 500m in the same way as in Section IV-A. In all cases simulated, no matter how small the threshold is, we observed convergence of the distributed algorithm. Fig. 4 shows typical convergence performance with different distance thresholds. "Threshold $= \infty$" corresponds to the original MIMO-DP algorithm, while "threshold $=$ 0m" refers to the scenario where simultaneous beamformer updates are completely allowed. Furthermore, according to the numerical results above, the speed of convergence is faster when simultaneous updates are allowed.

V. Conclusions

We have presented a distributed algorithm for adjusting beamforming vectors in a peer-to-peer MIMO network to maximize the sum utility over all users, with rank-one precoding matrices. Convergence is established based on the result in [2] that each beamformer update increases the total utility, provided that all interference prices are current, and the observation that updating receive filters increases those users' utilities while others' utilities are unchanged. Examples were also presented, which indicate that ignoring some prices related to remote users, or relaxing the update rules does not compromise the performance of the MIMO-DP algorithm significantly.

Due to the spatial degrees of freedom in MIMO networks, it is not optimal in general to apply a rank-one precoding matrix for each user. Therefore, how to adjust a precoding matrix to approach an optimal solution in a MIMO network without rank estimation is still open. In addition, the convergence rate as a function of system parameters has not been considered.

References

[1] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 1074-1084, May 2006.

[2] C. Shi, R. A. Berry, and M. L. Honig, "Monotonic Convergence of Distributed Interference Pricing in Wireless Networks," to appear in *Proc. IEEE International Symposium on Information Theory*, June 2009.

[3] C. Shi, D. A. Schmidt, R. A. Berry, M. L. Honig, and W. Utschick, "Distributed Interference Pricing for the MIMO Interference Channel," to appear in *Proc. IEEE International Conference on Communications*, June 2009.

[4] J. Huang, R. A. Berry, and M. L. Honig, "Performance of Distributed Utility-Based Power Control for Wireless Ad Hoc Networks," in *Proc. Military Communications Conference, 2005*, pp. 2481-2487, Oct 2005.

[5] 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.

[6] M. F. Demirkol and M. A. Ingram, "Power-controlled capacity for interfering MIMO links," in *Proc. IEEE Vehicular Technology Conference,* vol. 1, Oct. 2001.

[7] S. Ye and R. S. Blum, "Optimized signalling for MIMO interference systems with feedback," *IEEE Trans. Signal Processing,* vol. 51, no. 11, pp. 2839-2848, Nov. 2003.

[8] G. Scutari, D. P. Palomar and S. Barbarossa, "Competitive design of multiuser MIMO systems based on game theory: A unified view," *IEEE Journal on Selected Areas in Communications: Special Issue on Game Theory*, vol. 26, no. 7, pp. 1089-1103, Sept. 2008.

[9] E. Telatar, "Capacity of Multi-Antenna Gaussian Channels," *Eur. Trans. Telecomm. ETT*, vol. 10, no. 6, pp. 585-596, Nov. 1999.