Noise-resilient scaling for wideband distributed beamforming

Muhammed Faruk Gencel[*], Maryam Eslami Rasekh[†], Upamanyu Madhow[‡]

Department of Electrical and Computer Engineering

University of California Santa Barbara

Santa Barbara, California 93106

Email: [*]gencel@umail.ucsb.edu, [†]merasekh@ece.ucsb.edu, [‡]madhow@ece.ucsb.edu

Abstract—We consider distributed transmit beamforming from a cluster of cooperating nodes towards a distant destination, over a wideband dispersive channel. We consider explicit aggregate feedback, with the destination broadcasting its feedback to all nodes in the transmit cluster. Explicit feedback allows for frequency division duplex (FDD) operation, since there is no reliance on channel reciprocity. Aggregate feedback enables scalability, since the destination is agnostic to the number and identity of the transmitters. There are two key contributions. First, it is shown in a narrowband setting that, unlike a well-known one bit aggregate feedback scheme, a training-based approach in which each transmitter learns its channel to the destination based on aggregate feedback, is resilient to noise, even when the feedback is heavily quantized. Second, extending the scheme to wideband settings by using OFDM, training on a subset of subcarriers, we show that interpolation across subcarriers using sparse time domain modeling provides accurate channel estimates. In particular, at low SNR per node, our approach, which exploits the correlation across subcarriers, outperforms independent per-subcarrier channel estimates in a hypothetical benchmark system in which training and feedback is applied to all subcarriers. Our simulation results show that the beamforming performance obtained by our approach is a substantial fraction of the ideal beamforming gain, even for typical received SNRs per node as low as -20 dB, and coarse 4-phase feedback quantization. In general, the method requires training time linearly proportional to the number of cooperating nodes and inversely proportional to the per-node SNR.

I. INTRODUCTION

We propose and evaluate an approach for distributed transmit beamforming over wideband dispersive channels. The goal is to form a virtual antenna array from a cluster of N cooperating transmitters, in order to focus power towards a distant receiver, over a wideband dispersive channel. Fixing the power at each transmitter, the net transmitted power scales with N, which we may term the *power pooling gain*. For a narrowband channel, if the carrier phases for the different transmitters align at the receiver, an additional N-fold *beamforming gain* is obtained due to coherent addition of the signals. Thus, narrowband distributed transmit beamforming leads to N^2 scaling of the received power. We adopt OFDM with per-subcarrier beamforming as a natural extension of this concept to wideband dispersive channels. We seek to achieve the following benchmark: each transmitter has an estimate of its channel to the receiver, and hence can undo the channel phase for each subcarrier (we do not consider optimization of

power allocation across subcarriers).

The key issues and constraints driving our proposed approach are as follows. First, the system must be able to operate in Frequency Division Duplex (FDD) mode, hence we cannot rely on implicit feedback via channel reciprocity. The receiver must therefore provide *explicit* feedback to the transmitters. Second, the system must scale at the protocol level: the receiver should be agnostic to the number and identity of the transmitters. This leads to the notion of *aggregate* feedback broadcast by the receiver to the transmit cluster. Third, the system must be able to scale physically: since the net received power scales as N^2, we should be able to scale down the received SNR per transmitter while operating at a desired range and spectral efficiency.

A well-known approach to distributed transmit beamforming with explicit aggregate feedback is stochastic ascent, with the transmitters using small random phase perturbations, and the receiver providing one bit of feedback per iteration [1]. While this approach is simple and has formed the basis for several prototypes [2], [3], it fails for low per-node SNRs [4], roughly speaking, because noise masks the effect of small phase perturbations. This motivates our approach, in which the transmitters employ large phase perturbations during a training phase, rather than attempting to make small adjustments while beamforming.

Approach and contributions: We first consider design of a training phase for a narrowband system, in which each transmitter sends an orthogonal (or quasi-orthogonal, depending on the level of coordination across transmitters) sequence. The receiver feeds back the corresponding sequence of received complex amplitudes. Each transmitter correlates the feedback sequence with its transmitted sequence to estimate its channel. Unlike the one-bit feedback algorithm, the approach scales to arbitrarily low per-node SNRs: measurement noise is averaged out by scaling the length of the training period inversely with the per-node SNR. The feedback can be drastically quantized, with a number of bits independent of the number of transmitters N: quantization noise (whose variance scales with N) is averaged out by scaling the training period with the number of transmitters. In particular, 1 bit per I/Q dimension provides 90% of the ideal beamforming gain, with the gap to unquantized feedback almost vanishing when the number of feedback bits per dimension is increased to 2.

For wideband dispersive channels, we apply the preceding narrowband training to a subset of subcarriers in an OFDM system. We perform interpolation across subcarriers based on a sparse time domain modeling. The channel is modeled as a set of impulses, whose delays act as "frequencies" in the frequency domain, and can therefore be estimated using state-of-the-art frequency estimation algorithms [5]. At low SNR, the performance of our sparse estimation approach, which exploits channel correlation across subcarriers, is superior to that of benchmark per-subcarrier estimation in a hypothetical system with feedback for all subcarriers.

Related work: We do not attempt a comprehensive review of the extensive literature on distributed MIMO, but note that there are two broad categories of recent work. The first is based on coordination of infrastructure nodes (WiFi access points or cellular base stations) via a fast wired backhaul [6]–[8]. The second category consists of all-wireless schemes, with recent prototypes based on one-bit aggregate feedback [2], [3] or per-node feedback [9]. Our approach falls in the second category, but distinguishes itself in several ways: it applies to wideband systems, unlike the narrowband systems in [2], [3], [9]; it scales to low per-node SNR regimes, unlike the one-bit scheme; and it scales at the protocol level, unlike per-node feedback schemes such as [9].

II. SYSTEM MODEL

The system of interest is depicted in Fig. 1, with N cooperating transmitters aiming to send a common message to a distant receiver. The time domain channel from each transmitter to the receiver is modeled as a discrete set of paths, characterized by the (continuous-valued) complex gains and delays for each path. We consider an OFDM system with M subcarriers, with subcarrier spacing smaller than the channel's coherence bandwidth. The frequency response of the channel between the nth transmitter and the receiver at subcarrier m is denoted by $H_n(f_m), n = 1, \ldots, N$ and $m = 1, \ldots, M$.
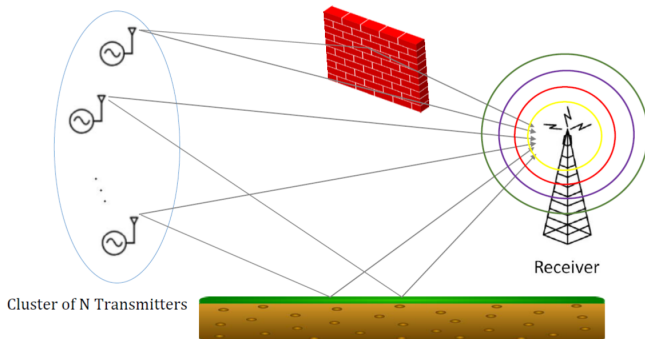

Fig. 1. N transmitters beamforming over frequency selective channels, using feedback broadcast from the receiver.

We assume that the transmitters are synchronized in carrier frequency and OFDM frame timing (this is relatively straightforward to achieve via local cooperation), hence the focus is on achieving phase alignment at the receiver for each subcarrier. Let $H_n(f)$ denote the channel from transmitter n

to the receiver. Each transmitter employs receiver feedback to estimate its channel. Denoting the estimate by $\hat{H}_n(f)$, transmitter n applies the beamformer $w_n(f_m) = e^{-j\angle \hat{H}_n(f_m)}$ to undo the channel phase on subcarrier m. Thus, the net channel seen by the receiver on this subcarrier is

$$y(f_m) = \sum_{n=1}^{N} H_n(f_m) w_n(f_m) = \sum_{n=1}^{N} |H_n(f_m)| e^{j\delta_{nm}}$$

where $\delta_{nm} = \angle H_n(f_m) - \angle \hat{H}_n(f_m)$ is the error in estimating the channel phase for transmitter n on subcarrier m.

In the next section, we describe channel estimation for a single subcarrier using aggregate receiver feedback. In Section IV, we show this approach, applied to a subset of OFDM subcarriers, can be used for wideband channel estimation using sparse time domain channel modeling.

III. NARROWBAND BEAMFORMING ALGORITHM

Transmitter n sends complex-valued training sequence $A_{l,n}$, $l = 1, ..., L$ over a training period of length L. We set the amplitudes $|A_{l,n}| = 1$, so that the training information is contained in the phases. We denote by $\mathbf{A}$ the corresponding $L \times N$ training matrix. We assume that the training sequences are orthogonal across transmitters (i.e., that $\mathbf{A}^H \mathbf{A}$ is diagonal), which requires a minimal level of local coordination. (Even this level of coordination can be dispensed with, at the cost of some performance loss, if we use quasi-orthogonal sequences.) The receiver is oblivious to the identity of transmitters and performs the same task at every iteration: it measure the complex amplitude and broadcasts its quantized value.

The observations at the receiver, collected over times $l = 1, ..., L$, can be written as the $L \times 1$ vector

$$\mathbf{y} = \mathbf{A}\mathbf{h} + \mathbf{w}$$

where $\mathbf{h}$ is the $N \times 1$ channel vector across different transmitters and $\mathbf{w} \sim CN(0, N_0 \mathbf{I})$ is the receiver noise.

For noiseless, unquantized feedback, the least squares estimate for the channel vector is given by

$$\widehat{\mathbf{h}} = \left(\mathbf{A}^H \mathbf{A}\right)^{-1} \mathbf{A}^H \mathbf{y} \qquad (1)$$

assuming that $L \geq N$ and $\mathbf{A}$ has rank N. Each node can thus obtain its channel estimate by taking the inner product of its corresponding row in the matrix $(\mathbf{A}^H \mathbf{A})^{-1} \mathbf{A}^H$ and the channel measurement feedback vector.

The Cramer-Rao lower bound on error covariance is $\mathbf{C}_h = N_0 \left(\mathbf{A}^H \mathbf{A}\right)^{-1}$. For each transmitter, the error covariance is bounded as

$$\text{Var}(\hat{\mathbf{h}}_n) \geq (N_0(\mathbf{A}^H \mathbf{A})^{-1})_{n,n} \geq \frac{N_0}{(\mathbf{A}^H \mathbf{A})_{n,n}}$$

with the bound attained for orthogonal training (diagonal $\mathbf{A}^H \mathbf{A}$). If this case, each node can estimate its channel by separately correlating the observations with its own training sequence:

$$\hat{\mathbf{h}}_n = \frac{1}{L}\mathbf{a}_n^H \mathbf{y} = \mathbf{h}_n + \frac{1}{L}\mathbf{a}_n^H \mathbf{w}$$

where $\mathbf{a}_n$ is the nth column of the training matrix. The estimation error covariance $\mathrm{Var}(\hat{\mathbf{h}}_n) = N_0/L$ can be made arbitrarily small by increasing the training interval L.

While it is possible to employ completely uncoordinated training, with each transmitter generating its training sequence pseudorandomly and independently, the coordination required for implementing truly orthogonal sequences (which provide the best possible performance for a given training duration and power) is minimal. There are many possible choices of training sequences, but we consider a DFT matrix in our results:

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & e^{-j2\pi/L} & \cdots & & e^{-j2\pi(N-1)/L} \\ \vdots & \vdots & \ddots & & \vdots \\ 1 & & & & \\ 1 & e^{-j2\pi(L-1)/L} & \cdots & & e^{-j2\pi(N-1)(L-1)/L} \end{bmatrix}$$

This is because DFT sequences are not only orthogonal, but they remain orthogonal when cyclically shifted by any amount. Assuming training is performed continuously by nodes repeating their training sequences (over a subset of subcarriers in an OFDM system, for example), a node can use any L-sized block of feedback and estimate its channel, without incurring interference from the sequences sent by the other transmitters. This permits nodes to join and leave the beamforming system at any time, assuming that their OFDM frames are aligned.

Feedback quantization:

In practice, the complex received signal amplitude measured at the receiver must be quantized to a limited number of bits and broadcast by the receiver. The variance of the received complex amplitude scales as N (the transmitted signals add up incoherently during the training period), hence a natural question is whether the quantization resolution also needs to be enhanced as N increases. Fortunately, the answer is no: as long as the receiver scales its quantizer step size Δ as $\sqrt{N}$ to accommodate the amplitudes it is seeing, we can use a fixed number of quantization bins, and average out the quantization noise across the training period.

To see this, write the channel estimate at transmitter n with quantized feedback as

$$\hat{\mathbf{h}}_n = \frac{1}{L}\mathbf{a}_n^{\mathbf{H}}(\mathbf{y} + \mathbf{n}_q)$$

where $\mathbf{n}_q$ is the quantization noise vector. Assuming quantization noise is distributed uniformly over the span associated to each level, the variance of any element $n_q[l]$ of the quantization noise vector scales as $\mathrm{Var}(n_q[l]) = \frac{\Delta^2}{12} \sim N$. If the quantization noise values can be approximated as independent over time,

$$\mathrm{Var}(\frac{1}{L}\mathbf{a}_n^H \mathbf{n}_q) \sim \frac{NL}{L^2} = \frac{N}{L}$$

so that the effect of quantization noise on channel estimation can be made independent of N by scaling L linearly with N. Thus, we can use a fixed feedback rate even as we increase the size of the transmit cluster, as long as the length of the training period scales linearly with cluster size.

In practice, poorly chosen training sequences can lead to correlation of quantization noise across time, which leads to performance degradation. For example, while repeating N-length orthogonal sequences averages out measurement noise, it leads to periodicity in quantization noise. Such correlation is avoided, for example, by using L-length DFT sequences.

For drastic quantization (e.g., 1 bit for I and 1 bit for Q), measurement noise provides a natural dither that aids estimation, so that performance can actually degrade as we increase SNR per node. This is avoided by introducing dither at the receiver prior to quantization by adding an appropriate amount of artificial noise.
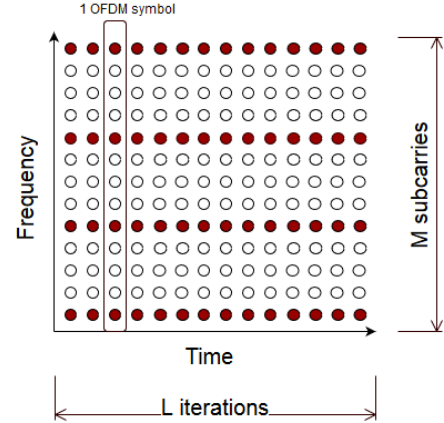

Fig. 2. Subset of OFDM subcarriers to be used in channel estimation

IV. EXTENSION TO FREQUENCY SELECTIVE CHANNELS

The preceding estimation algorithm is extended to OFDM over wideband channels by employing training on a subset of subcarriers, as shown in Fig. 2. The channel estimates over the training subcarriers are interpolated to all subcarriers by first performing a sparse time domain channel reconstruction. In particular, we assume that the time domain channel is modeled as a discrete multipath channel:

$$h_n(t) = \sum_{k=1}^{K} g_k \delta(t - \tau_k) \tag{2}$$

In the frequency domain, this channel is a mixture of sinusoids with "frequencies" equal to multipath delays:

$$H_n(f) = \sum_{k=1}^{K} g_k e^{-j2\pi f \tau_k} \tag{3}$$

The problem of reconstructing the channel from estimates at a subset of subcarriers is therefore equivalent to finding the frequency (τ_k) and amplitude (g_k) of each sinusoid in the mixture. In order to solve this problem, we use a dual of a frequency estimation algorithm developed in [5]. The inputs are the channel measurements $\hat{H}_n(f_i)$ on the training subcarriers. The outputs are estimates of the multipath delays and complex amplitudes. These estimates are then substituted
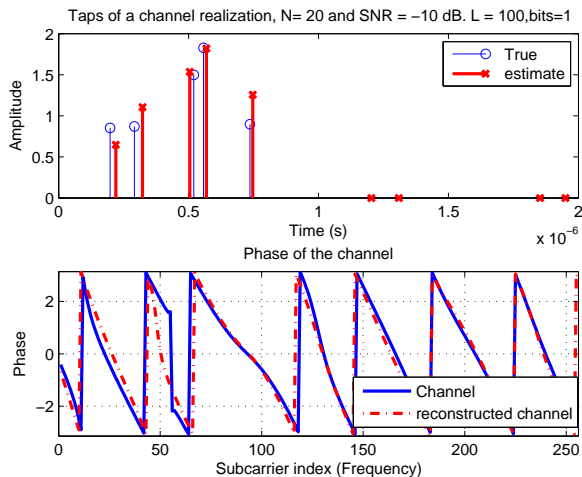
Fig. 3. A realization of time domain channel estimation algorithm using subset of subcarriers (32 out of 256) and reconstructed phase of the channel by time domain estimates

in (3) to obtain channel estimates (and hence beamforming weights) for all subcarriers.

At low SNR per node and severe quantization, the frequency estimation algorithm used to detect delay components can overfit by producing spurious low-amplitude taps with delays larger than the channel delay spread. We find that channel estimates are improved by a simple denoising procedure which excludes these components. One realization of a 5-tap time domain channel is shown in Fig. 3 along with the estimated taps using channel estimates at 32 out of 256 subcarriers. The channel estimates were obtained using the aggregate feedback method described in Section III in a 20 element array over training time of 100 iterations. The phase of the reconstructed channel, consisting only of the first 5 estimated taps with estimated delays smaller than the channel delay spread of $1/\Delta f = 1\mu s$, is shown in Fig. 3 alongside the correct channel phase. It can be seen that the channel phase profile across frequency is estimated accurately even at a low per-node SNR of -10 dB and drastic quantization of 1 bit per I/Q dimension (or equivalently, four-phase quantization).

V. SIMULATION RESULTS

We first evaluate the performance of the narrowband algorithm in Section III. Fig. 4 plots the fraction of ideal beamforming gain as a function of per-node SNR for different quantization levels. These values were obtained using Monte Carlo simulation of the algorithm in a 20 element distributed array over $L = 160$ training steps using DFT sequences. This figure demonstrates that, even with heavily quantized feedback of 1 bit per real and imaginary dimension, 90% of beamforming gain can be achieved at -10 dB SNR per node. We see that when we increase the feedback to 2 bits per dimension, the performance is close to that of unquantized feedback, and that even this small gap disappears when feedback is increased to 3 bits per dimension. A significant advantage
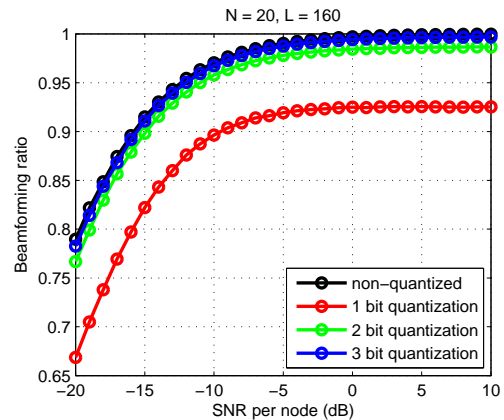

Fig. 4. Effect of quantization bits (per I&Q components) on the beamforming on a narrowband channel

of this method is its resilience to noise: the channel can be estimated to desired accuracy at arbitrarily low SNR simply by noise averaging over a long enough training period. At higher SNR, we add dither to prevent performance degradation due to drastic quantization (details omitted due to lack of space).
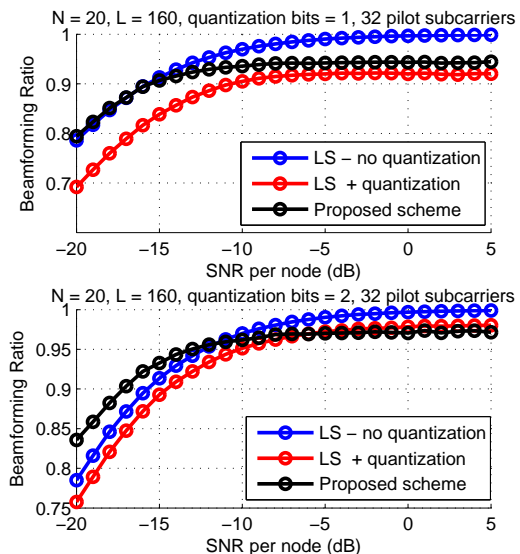

Fig. 5. System model of N transmitters beamformig with feedback broadcast from the receiver over frequency selective channels

Wideband beamforming performance is depicted in Fig. 5. In these simulations, beamforming is performed on a 20 node array over 160 iterations using OFDM with 256 subcarriers. The fraction of ideal beamforming gain averaged over subcarriers is plotted as a function of per-node SNR. The blue and red curves represent results for independent least squares channel estimation using the algorithm of Section III for each subcarrier, while the black curves are obtained by performing the least squares channel estimation for a subset of 32 subcarriers and sparse time domain reconstruction for interpolation across subcarriers, as described in Section IV.
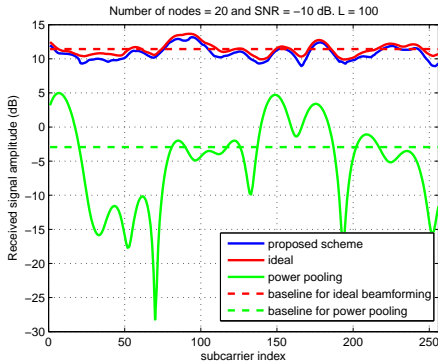
Fig. 6. Amplitude and phase of the channel seen by the receiver

We see that, especially in low SNR regimes, the sparse reconstruction method actually outperforms per-subcarrier channel estimation in the hypothetical system with feedback for all subcarriers. This is because the imposition of time domain sparsity provides implicit noise resilience, exploiting channel correlation across subcarriers.

The net channel seen by the receiver over the frequency bandwidth of the signal after beamforming using the proposed scheme is shown in Fig. 6 at -10 dB SNR per node with $N = 20$ nodes. We note that, in addition to beamforming gain, the spatial diversity provided by coherent combining across different transmitters significantly reduces the frequency selectivity of the net channel. On the other hand, incoherent combining in power pooling mode continues to result in significant frequency selective fading: at each subcarrier, central limit theorem arguments imply that the net complex gain in power pooling mode is well modeled as zero mean complex Gaussian, corresponding to Rayleigh fading.

VI. CONCLUSION

Our proposed channel estimation (and associated distributed beamforming) scheme is scalable at both the protocol and the physical level. The receiver can be oblivious to the number of cooperating transmitters in its own channel estimation and feedback generation mechanisms, and can employ quantization as drastic as 1 bit per dimension. The SNR per node can scale down indefinitely, as long as we are willing to increase the length of the training period correspondingly. Furthermore, in the context of OFDM, a subset of subcarriers can provide continuous training, given the demonstrated efficacy of interpolation via sparse time domain reconstruction.

An important topic for future work is to incorporate these ideas and algorithms into a complete system design and protocol specification, and to demonstrate them via prototypes. To this end, we must explicitly address issues such as frequency and OFDM frame timing synchronization within the transmit cluster, which we have abstracted out in the present paper. The transmitters must also account for the receiver's frequency offset, again abstracted out here, jointly with their channel estimates. Incorporating the effect of mobility is also an important consideration in certain applications.

REFERENCES

[1] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Distributed transmit beamforming using feedback control," *Information Theory, IEEE Transactions on*, vol. 56, no. 1, pp. 411–426, 2010.

[2] F. Quitin, M. M. U. Rahman, R. Mudumbai, and U. Madhow, "A scalable architecture for distributed transmit beamforming with commodity radios: Design and proof of concept," *Wireless Communications, IEEE Transactions on*, vol. 12, no. 3, pp. 1418–1428, 2013.

[3] F. Quitin, A. Irish, and U. Madhow, "Distributed receive beamforming: A scalable architecture and its proof of concept," in *Vehicular Technology Conference (VTC Spring), 2013 IEEE 77th*, June 2013, pp. 1–5.

[4] M. F. Gencel, M. E. Rasekh, and U. Madhow, "Distributed transmit beamforming with one bit feedback revisited: How noise limits scaling," in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015, pp. 2041–2045.

[5] B. Mamandipoor, D. Ramasamy, and U. Madhow, "Newtonized orthogonal matching pursuit: Frequency estimation over the continuum," *arXiv preprint arXiv:1509.01942*, 2015.

[6] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H.-P. Mayer, L. Thiele, and V. Jungnickel, "Coordinated multipoint: Concepts, performance, and field trial results," *Communications Magazine, IEEE*, vol. 49, no. 2, pp. 102–111, 2011.

[7] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Airsync: Enabling distributed multiuser mimo with full spatial multiplexing," *IEEE/ACM Transactions on Networking (TON)*, vol. 21, no. 6, pp. 1681–1695, 2013.

[8] H. Rahul, S. Kumar, and D. Katabi, "MegaMIMO: Scaling Wireless Capacity with User Demands," in *ACM SIGCOMM 2012*, Helsinki, Finland, August 2012.

[9] P. Bidigare, M. Oyarzyn, D. Raeman, D. Chang, D. Cousins, R. O'Donnell, C. Obranovich, and D. Brown, "Implementation and demonstration of receiver-coordinated distributed transmit beamforming across an ad-hoc radio network," in *Signals, Systems and Computers (ASILOMAR), 2012 Conference Record of the Forty Sixth Asilomar Conference on*. IEEE, 2012, pp. 222–226.