Scaling wideband distributed transmit beamforming via aggregate feedback

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 investigate distributed beamforming from a cluster of N cooperating transmitters to a distant destination over a wideband dispersive channel. Feedback from the destination is critical for enabling this. In order to develop protocols that scale to arbitrarily large numbers of cooperating nodes, we restrict attention to *aggregate* feedback broadcast from the destination to the entire transmit cluster, rather than per-transmitter channel feedback as in conventional feedback-based MIMO systems. We first show that naive application of a one-bit feedback algorithm developed for narrowband channels to each subcarrier in an OFDM system does achieve beamforming gain on each subcarrier, but results in an effective channel at the destination with severe phase discontinuities across frequency, which is not amenable to standard receive channel estimation algorithms. We then show that it is possible to enforce smoothness of phase across frequency by augmenting the feedback to 2 bits per subcarrier, which enables modeling and estimation of the effective channel as sparse in the time domain. Our preliminary results show that, even when the SNR per node is well below the threshold for reliable demodulation, it is possible to bootstrap using the N-fold power pooling gain obtained from incoherent combining of the signals from multiple transmitters, and to attain a significant fraction of the N^2-fold beamforming gain using the proposed algorithm. We also discuss a number of open issues, recognizing that this is only a first step in developing scalable, wideband, distributed MIMO systems.**

I. Introduction

Distributed MIMO refers to a class of communication schemes in which a cluster of cooperating nodes emulates a virtual antenna array. Realization of this concept requires tight synchronization among the nodes, and there has been significant progress in recent years in developing and prototyping algorithms for achieving such distributed synchronization, with feedback from the receiver (either explicit, or implicitly obtained from channel reciprocity) playing a key role. In conventional feedback-based MIMO systems, explicit feedback is typically in the form of quantized estimates of the channel for each antenna. While this works for a relatively small number of antennas, for distributed MIMO systems, we would like to devise feedback strategies that scale to an arbitrary number of cooperating nodes. An interesting example of this is the one-bit feedback algorithm [1] for distributed transmit beamforming, in which cooperating transmitters employ a randomized ascent strategy for adapting their phases to achieve a beamforming solution based on a single bit of common feedback broadcast by the destination. Since the transmitters adapt their phases independently, this allows the protocol to scale to arbitrarily many cooperating nodes. It is of interest to explore, therefore, whether this model of *aggregate feedback* broadcast to the cooperative cluster is more generally applicable to devise scalable protocols for distributed MIMO.

Approach and Overview: In this paper, we explore the concept of aggregate feedback for distributed transmit beamforming over wideband dispersive channels. A natural approach is to employ the one-bit feedback algorithm over each subcarrier in an OFDM system, and this indeed works as expected, achieving coherent combining at the receiver of the transmitted signals for each subcarrier. However, because this is a randomized ascent algorithm, the received phase for different subcarriers are essentially unrelated. Thus, the effective channel at the receiver exhibits severe phase discontinuities across frequency. This is incompatible with standard OFDM channel estimation algorithms, which assume a level of smoothness in phase across frequency consistent with the time dispersion in the class of channels being designed for. In this paper, we show that by augmenting the algorithm with 1 bit of additional feedback per subcarrier, we can achieve smoothness in the phase of the effective channel at the receiver without any impact on the convergence of the basic 1-bit algorithm. We show that the effective channel thus obtained is smooth enough to be estimated by the receiver based on a small number of pilot subcarriers. The performance of the 2-bit algorithm is compared against two benchmarks. The first is the noncoherent combining benchmark, in which transmissions are loosely synchronized in time, but make no attempt to synchronize in phase at the receiver. This is the initial condition for bootstrapping our aggregate feedback scheme, and corresponds to a roughly N-fold "power pooling" gain, but the effective channel seen by the receiver still sees substantial frequency selectivity. The second is the ideal beamforming solution which provides roughly N^2-fold gain in received power, together with significant reduction in frequency selectivity due to diversity combining across transmitters. We observe that the 2-bit feedback algorithm approaches the beamforming solution when the SNR for the power pooling solution is "high enough" (e.g., beyond 5 dB), but that it has difficulty in progressing much beyond the power pooling initial condition when the power pooled SNR is low (e.g., 0 dB or below). We leave open the issue of whether the 2-bit algorithm can be effectively bootstrapped in such a noisy regime.

Related work: While idealized distributed MIMO has long been the subject of information-theoretic investigation, over the past few years, there has been significant progress in algorithms and experimental demonstrations that address the difficult synchronization problems in realizing these concepts. Demonstrations of base station or access point cooperation based on a fast wired backhaul in cellular and WiFi contexts include [2], [3]. All-wireless demonstrations of distributed beamforming include [4]–[6]. In particular, the demonstrations in [4], [5] are based on the one-bit aggregate feedback algorithm in [1]. However, this prior work on aggregate feedback, both theoretical and experimental, has focused on narrowband systems [7]. To the best of our knowledge, this is the first paper to investigate aggregate feedback for wideband dispersive channels.

Feedback-based centralized MIMO schemes are typically based on sending quantized estimates of the spatial channel for each antenna [8], and hence do not scale beyond a small number of antennas. The recent activity in massive MIMO [9] focuses on centralized arrays with a large number of antennas, and is typically based on the assumption of time division duplex (TDD), with channel estimates derived *implicitly* using reciprocity. In principle, reciprocity-based techniques should also yield scalable distributed MIMO systems employing TDD. However, in order to scale FDD-based distributed MIMO (for which reciprocity does not apply) to a large number of nodes, aggregate feedback is an attractive approach, since the receiver does not need to be cognizant of the number or identity of cooperating nodes.

The remainder of this paper is organized as follows; in section II, the system model is presented and the dispersive channel considered in the evaluation is described. The iterative feedback algorithm used for phase synchronization is explained in section III and its performance is analyzed in section IV. Conclusions and areas for further analysis of the proposed methods are reviewed in section V.

II. System Model

As depicted in Fig. 1, we consider a cluster of M cooperating transmitters that wish to communicate a common message to a distant destination. The channel between the mth transmitter and the destination is denoted by $H_m(f)$, $m = 1, ..., M$. In our simulations, we model these channels as independent realizations from a bandwith dependent tapped delay line model with exponential power delay profile [10] and all realizations are normalized to power of 0 dB. The channel model is

$$h(t) = \sum_{i=1}^{k} \alpha_i \delta(t - \frac{i}{W})$$

where W is the bandwith of the transmitted signal and $k = \tau_d W$ assuming there are no multipath components after the delay spread τ_d and amplitudes of the taps are

$$\alpha_i \sim CN(0, ab^i),$$

where $b = exp(-\frac{1}{W\tau_{rms}})$ and $a = 1 - b$. Fig. 2 shows a typical channel realization with root mean squared delay $\tau_{rms} = 1\mu s$ as a typical RMS delay spread value in an urban area [11], $W = 10$ MHz, $\tau_d = 4.7\mu s$ and carrier frequency $f_c = 2.5$ GHz. We note that that there are several significant fades.

We consider OFDM with N subcarriers, with subcarrier spacing smaller than the channel coherence bandwidth. We assume that the transmitters are synchronized in terms of clock and carrier frequency (this can be achieved by a number of mechanisms, including using a master-slave architecture), but have timing and phase offsets that are *a priori* unknown. Assuming that the transmitters are coarsely synchronized in timing

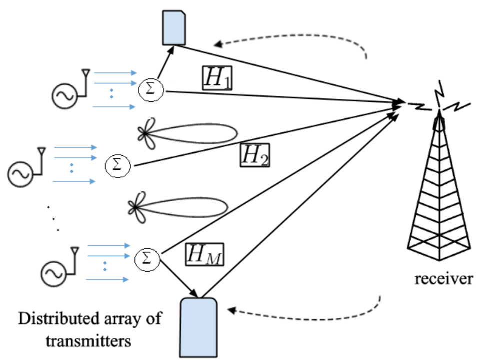

Fig. 1: System Model with M transmitters transmitting on N subcarriers on multipath channels

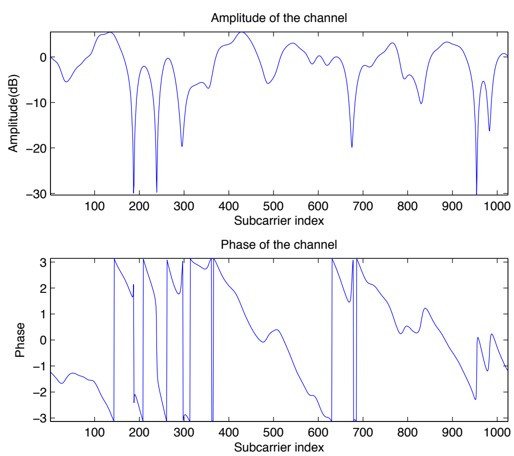

Fig. 2: Frequency response of realization of a wideband channel with $\tau_d = 4.7\mu s$, over $N = 1024$ subcarriers, with a total bandwidth of $W = 10$ MHz and $f_c = 2.5$ GHz

such that residual offsets are significantly smaller than the channel time dispersion, we can absorb these timing offsets within the OFDM cyclic prefix. Thus, the key problem that we focus on here is the problem of coherent phase combining for each subcarrier

We are interested in regimes in which the received SNR corresponding to any given transmitter is too low to permit reliable communication at the desired rates. If the (loosely synchronized) transmitters emit a common message at an agreed upon time, then the effective channel seen by the receiver is

$$G_{nc}(f) = \sum_{m=1}^{M} H_m(f) \tag{1}$$

where the subscript denotes that the channels from different transmitters are being combined noncoherently. Thus, while the net received power increases (assuming each transmitter sends at a fixed power), we still see significant frequency selective behavior, as shown in Fig 3(a), which shows a typical effective channel obtained from such noncoherent power pooling.

If each transmitter knows its channel to the receiver, the optimal strategy, subject to a per-transmitter power constraint, is to employ waterfilling. In this paper, however, we consider a suboptimal strategy in which each transmitter simply adjusts its phase at each subcarrier to compensate for the channel, while keeping its power constant across subcarriers. Specifically, if the mth transmitter applies a precoder of the form

$$P_m(f) = e^{-j\angle H_m(f)}$$

then the net channel obtained is

$$G_c(f) = \sum_{m=1}^{M} P_m(f) H_m(f) = \sum_{m=1}^{M} |H_m(f)| \tag{2}$$

we term this solution the ideal beamforming solution. Fig 3.a benchmark plot shows a typical effective channel obtained from ideal beamforming. In addition to power pooling and beamforming gains at each frequency, we also notice a diversity gain resulting in a significant reduction in frequency selectivity.

As we shall see, the one-bit feedback algorithm applied independently to each subcarrier leads to a solution of the form

$$G_{1bit}(f) = e^{j\phi(f)} \sum_{m=1}^{M} |H_m(f)| \tag{3}$$

so that we are achieving beamforming gain at each frequency, but with arbitrary frequency-dependent phase shifts $\phi(f)$ because we are not coordinating across subcarriers. Thus, the smoothness of phase across frequency that is a characteristic of natural channels, and is relied upon by receiver estimation algorithms, gets destroyed.

We show that a simple modification which runs in parallel to the one-bit feedback algorithm, in which the receiver sends back 1 additional bit (per subcarrier) to help enforce phase continuity, can be used to approach the ideal beamforming solution up to a constant phase offset. That is, the effective channel approaches

$G_{eff}(f) = G_c(f)e^{j\phi}$, where ϕ does not depend on frequency. This allows the receiver to use standard channel estimation algorithms that exploit smoothness of phase across frequency once it switches to decision-directed or pilot-based estimation.

III. Phase Synchronization with Periodic Feedback

A. Narrowband channels

The 1-bit feedback algorithm is a simple iterative procedure that synchronizes the received signals corresponding to all the transmitters without attempting to explicitly estimate channel states. In this procedure, time is divided into slots and each transmitter m applies a beamforming weight of $1e^{j\theta_m(n)}$ to its transmitted signal in slot n. In each slot, each transmitter applies a random perturbation to its phase. At the end of the slot the receiver broadcasts one bit of feedback indicating whether or not the RSS has improved compared with the previous iteration. Upon receiving the feedback, transmitters adopt the latest beamforming phase if RSS has improved and undo the perturbations if it has degraded. This procedure is repeated until coherence is achieved. The resulting beamforming phases compensate for the phase of the channel response as well as the phase offsets of the local oscillators. The feedback from the receiver based on the received power observation in iteration n is formulated as:

$$F_n = \begin{cases} 1 & r[n] > r_{best}[n] \\ 0 & r[n] < r_{best}[n] \end{cases} \tag{4}$$

$$r_{best}[n] = \max_{t<n} r[t]$$

Each transmitter m keeps track of the best beamforming phase so far, $\theta^m_{best}(n)$. At the beginning of each slot, a random perturbation $\delta^m(n)$ is added to this phase and the beamforming phase used in slot n is equal to:

$$\theta^m(n) = \theta^m_{best}(n) + \delta^m(n)$$

Based on the feedback from the receiver, each transmitter updates its best beamforming phase to:

$$\theta^m_{best}(n+1) = \begin{cases} \theta^m_{best}(n) & F_n = 0 \\ \theta^m(n) = \theta^m_{best}(n) + \delta^m(n) & F_n = 1 \end{cases} \tag{5}$$

B. Extension to wideband channels: the 2-bit feedback algorithm

A wideband frequency selective channel is parallelized into flat fading subcarriers using OFDM, hence we can directly apply the narrowband one-bit algorithm However, the beamforming phases evolve independently for different subcarriers, which results in a lack of continuity in the received signal phase across subcarriers. As shown in Fig. 3, the noncoherent benchmark for 10 transmitters has a smooth phase response at the receiver, but the phase response after adaptation based on parallel 1-bit feedback algorithms exhibits severe discontinuities. This can be problematic when using a standard OFDM receiver, where typically a number of subcarriers are reserved as pilots for channel estimation and no data is transmitted on them. The receiver interpolates the channel response between these pilots to obtain estimates for the channel response of all other subcarriers. These channel estimates are then used to decode the symbols sent on the data carrying subcarriers. After the initial training phase, our goal is to treat cluster of nodes as a single transmitter at the receiver side. We therefore modify the 1-bit algorithm to provide an effective channel that is smooth across subcarriers.

In order to obtain channel continuity, an additional bit of feedback is sent by the receiver for each subcarrier demanding all transmitters to either increase the phase of their beamforming weight by a predefined increment of γ or remain at the same phase for that subcarrier (on top of their individual random perturbations for the one-bit algorithm). In order to achieve phase continuity at the channel seen by receiver, the receiver compares the phase of each subcarrier with the average phase of all subcarriers and send feedback that will bring this phase closer to the average. For iteration n and subcarrier frequency f_i, the feedback decision for the second bit at the receiver is formulated by defining $g[f_i, n]$ such that:

$$g[f_i, n] = r^*[f_i, n] \sum_{j \neq i}^{N} r[f_j, n]$$

for $i = 1, .., N$ and $r[f_i, n]$ is the received signal at i^{th} subcarrier. Then, feedback bit for each subcarrier is decided by comparing $\angle g[f_i, n]$ with the predefined constant α and incrementing beamforming gains accordingly as

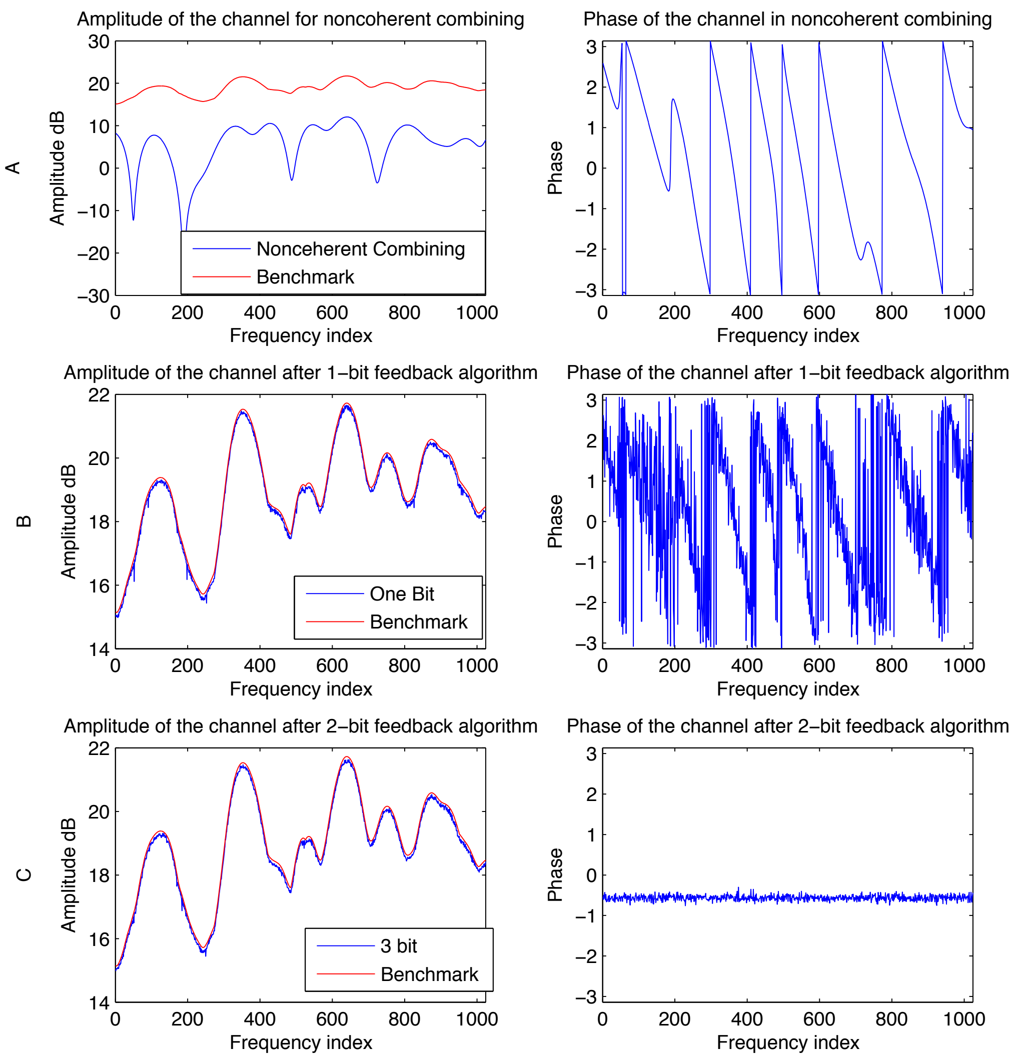

Fig. 3: Amplitude and phase of net channel response of 10 transmitter for A) non-coherent combinations, B) coherent combination after convergence of 1-bit feedback, C) coherent combination and smoothing with 2-bit feedback

$$\theta^m_{best}[f_i, n+1] = \begin{cases} \theta^m_{best}[f_i, n] + \gamma & \angle g[f_i, n] > \alpha \\ \theta^m_{best}[f_i, n] & otherwise \end{cases} \tag{6}$$

That is, while the one-bit feedback algorithm is adapting the phases at each subcarrier to achieve beamforming gain, we are running a consensus-style algorithm on the received phases to attain phase continuity across subcarriers. However, instead of directly comparing phases, we work with complex amplitudes in order to avoid phase wrapping issues, and to provide a soft averaging mechanism in which subcarriers with larger received amplitudes have larger weight. The resulting net channel response of the 2-bit feedback algorithm is shown in Fig. 3(c) for 10 transmitters beamforming over a dispersive channel with delay spread of $4.7\mu s$ using 1024-subcarrier OFDM with frequency spacing of 9.76kHz. It can be seen that the channel phase response is smooth. The amplitude remains equal to the amplitude provided by beamforming (i.e. equal to the sum of the channel amplitudes of all transmitters), since the one-bit feedback algorithm operates in parallel to the phase smoothing mechanism using the second bit. The diversity provided by pooling the power of 10 transmitters significantly decreases frequency selectivity relative to both a single transmitter and noncoherent power pooling

IV. Performance

Fig. 4 shows the progress of net received signal power (summed over subcarriers) with time for 10 transmitters. The corresponding symbol error rates for each scenario (analytically calculated based on the received SNR, assuming QPSK modulation) are shown in Fig. 5. We vary the SNR per user at the receiver across curves, keeping the transmit power the same. The initial condition for the algorithm is non coherent power pooling, and the progress towards the ideal beamforming solution depends on the noise level. We see that the 2-bit algorithm is fairly robust to noise, and enables reliable operation in regimes where a single transmitter would not be able to close the link without going to very low spectral efficiencies. However, there is some noise threshold beyond which the algorithm breaks down. For example, even when the SNR per user is as low as -5 dB, we do attain a significant fraction of the beamforming gain, but when the SNR per user dips to as low as -10 dB, we barely progress beyond noncoherent power pooling. We conjecture that this threshold effect is based on how large the power-pooled SNR that we bootstrap with is, but more analysis and performance evaluation, as well as potential optimization of the two-bit algorithm, are needed to come up with quantitative guidelines.

The benchmark for perfect beamforming gain is shown in Fig. 3(c) along with the net channel amplitude after beamforming for comparison. This value is calculated in (7) for subcarrier frequency f_k as:

$$|H_{opt}(f_k)| = \sum_{m=1}^{M} |H_m(f_k)| \tag{7}$$

where M is the number of transmitters and m is the transmitter index. In the proposed scheme, the process of channel smoothing is performed independently from beamforming, i.e. based on the smoothing feedback all transmitters apply the same phase rotation in each subcarrier causing a rotation in the net channel phasor but leaving its amplitude unchanged and beamforming gain unaffected. Consequently, the performance of 2-bit beamforming with OFDM modulation is identical to the basic 1-bit feedback procedure in a narrowband link, hence the speed of convergence, as seen in Fig. 4, is similar to the predictions of [12] for the basic 1-bit feedback algorithm. For noiseless feedback, using optimum phase perturbation increment size of 9 degrees (obtained from simulations), 75% of beamforming gain is obtained after $5M = 50$ iterations. This step size, however, may not be optimal in the noisy settings considered here, and basic analysis of the one-bit feedback algorithm under noise is an open issue.

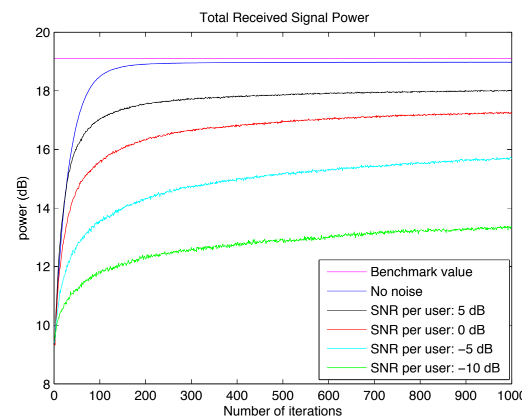

Fig. 4: Convergence of 2-bit feedback with time

Receiver channel estimation: The resulting channel after convergence is relatively smooth and can be estimated by the receiver from a subset of subcarriers. In the LTE standard, channel estimation is performed at the receiver by interpolating the channel measured from a number of pilot subcarriers. One in six subcar-

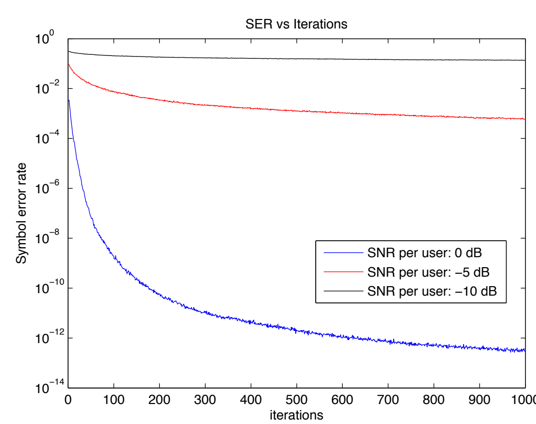

Fig. 5: Symbol error rate at every iteration

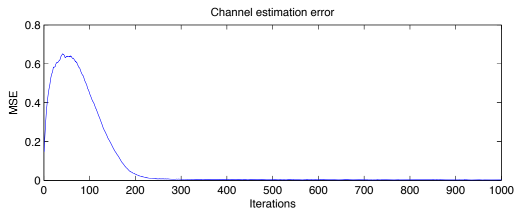

Fig. 6: MSE of interpolation of channel estimate from subset of subcarriers

riers are allocated for pilot transmission and channel sounding and the remaining subcarriers are used for data transmission. The net channel after convergence of 2-bit feedback is smooth in phase and relatively smooth in amplitude therefore channel estimations obtained by interpolation between one every seven subcarriers is low. The variation of interpolation error with iteration time is shown in Fig. 6.

V. Conclusions

We have shown that distributed transmit beamforming with aggregate feedback can be extended to wideband frequency selective channels while preserving the smoothness of the channel seen by the receiver. The concept of aggregate feedback allows us to bootstrap the noncoherent power pooling solution (already a factor of N better than for a single transmitter) towards the beamforming solution.

We view this work as a first step in developing a framework for scalable wideband distributed MIMO. Even for our specific problem of distributed transmit beamforming, there are a number of open issues. For example, is it possible to bootstrap without any training data at all? Our initial noncoherent power pooling solution is smooth across subcarriers. Can we use the 2-bit algorithm to maintain this smoothness, while making decisions and generating decision-directed feedback using differentially coherent modulation? Another important issue is to determine the fundamental limits of this approach, in terms of the lowest SNR per user, or power-pooled SNR, at which we can operate. Answering this question might require revisiting and optimizing the one-bit feedback algorithm in the presence of noise. Finally, we would like to model and mitigate the effects of clock drifts (which affect our assumptions of frequency synchronization and loose timing synchronization across the transmit cluster), as well as mobility. Our overall goal is to come up with cross-layer designs in which minimal modifications are required to existing cellular and WLAN architectures to accommodate node clusters acting as a single node.

Acknowledgments

This research was supported in part by the National Science Foundation under grant CCF-1302114, and by the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office.

References

[1] R. Mudumbai, B. Wild, U. Madhow, and K. Ramchandran, "Distributed beamforming using 1 bit feedback: from concept to realization," in *Proceedings of the 44th Allerton conference on communication, control and computation*, 2006, pp. 1020–1027.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] R. Mudumbai, U. Madhow, R. Brown, and P. Bidigare, "Dsp-centric algorithms for distributed transmit beamforming," in *Signals, Systems and Computers (ASILOMAR), 2011 Conference Record of the Forty Fifth Asilomar Conference on*, Nov 2011, pp. 93–98.

[7] R. Mudumbai, D. Brown, U. Madhow, and H. Poor, "Distributed transmit beamforming: challenges and recent progress," *Communications Magazine, IEEE*, vol. 47, no. 2, pp. 102–110, 2009.

[8] Y.-S. Tu and G. Pottie, "Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through awgn channels," in *Vehicular Technology Conference, 2002. VTC Spring 2002. IEEE 55th*, vol. 1, 2002, pp. 130–134 vol.1.

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

[10] U. Madhow, *Fundamentals of digital communication*. Cambridge University Press, 2008.

[11] S. Saunders and A. Aragón-Zavala, *Antennas and propagation for wireless communication systems*. John Wiley & Sons, 2007.

[12] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Scalable feedback control for distributed beamforming in sensor networks," in *Information Theory, 2005. ISIT 2005. Proceedings. International Symposium on*. IEEE, 2005, pp. 137–141.