

Distributed massive MIMO: algorithms, architectures and concept systems

U. Madhow, D.R. Brown III, S. Dasgupta, R. Mudumbai

Abstract—Making MIMO truly “massive” involves liberating it from the shackles of form factor constraints, by allowing arbitrarily large groups of neighboring nodes to opportunistically form virtual antenna arrays for both transmission and reception. Moving such distributed MIMO (DMIMO) systems from the realm of information theory to practice requires synchronization of the cooperating nodes at multiple levels. We are interested in all-wireless systems which are severely constrained in the amount of information that can be exchanged among the cooperating nodes, in contrast to recent proposals in massive MIMO (co-located antennas) or base station cooperation (which relies on a high-speed wired backhaul). The goal of this paper is to point out some of the research issues unique to scaling up such DMIMO systems. We briefly review the significant technical progress in design and demonstration over the past few years, and describe a research agenda for the next few years based on fundamental questions in attaining the “distributed coherence” required to realize concept systems such as DMIMO communication at large carrier wavelengths (e.g., white space frequencies for which standard antenna arrays are too bulky) and distributed 911 for emergency and rescue scenarios.

I. INTRODUCTION

The past two decades have seen tremendous progress in Multiple Input Multiple Output (MIMO) wireless systems, with significant performance breakthroughs achieved with multiple antenna techniques such as transmit/receive beamforming, spatial multiplexing, space division multiple access, and interference alignment. The theory and practice of MIMO communication has matured to the point where MIMO is now an integral component of several recent WiFi and cellular standards, such as 802.11n, 802.11ac, long-term evolution (LTE), WiMAX, and International Mobile Telecommunications (IMT)-Advanced. While the advantages of MIMO are significant, the applicability of MIMO is often limited by physical and economic constraints. For example, the form factor of handheld devices typically limits the number of antennas to only one or two. Even for infrastructure nodes such as access points and base stations, “massive” MIMO

Madhow is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106, madhow@ece.ucsb.edu. Brown is with the Department of Electrical and Computer Engineering, Worcester Polytechnic University, Worcester, MA 01609, drb@wpi.edu. Dasgupta and Mudumbai are with the Department of Electrical and Computer Engineering, University of Iowa, Iowa City, IA 52242. {rmudumbai, dasgupta}@engineering.uiowa.edu. This work is in part supported by US NSF grants EPS-1101284, ECCS-1150801, CNS-1329657 and CCF-1302456, CCF-1302114, CCF-1319458, ONR grant N00014-13-1-0202 and by the Institute for Collaborative Biotechnologies through the grant W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

transceivers with a very large number of antennas at current WiFi and cellular frequencies would perhaps be excessively bulky. And for lower, “white space” frequencies with carrier wavelengths as large as 6 meters, conventional MIMO is typically considered to be impractical.

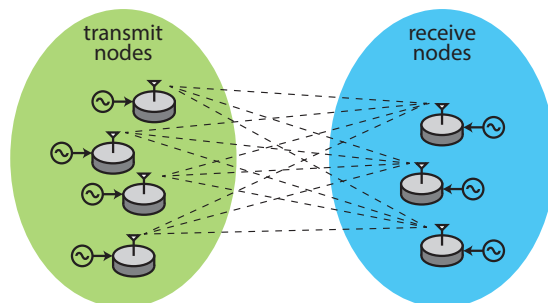


Fig. 1. Distributed MIMO system model with independent oscillators at each transmit and receive node.

One approach to unshackling MIMO from form factor constraints is to have multiple devices work together to form a *virtual* antenna array as shown in Fig. 1. In principle, such *distributed*, or *virtual* MIMO systems can emulate *any* MIMO communication technique feasible with centralized antenna arrays, thereby providing benefits such as improved transmit/receive directivity, reduced interference, increased degrees of freedom and spectral efficiency, and improved spatial diversity, while sidestepping the form factor and economic constraints limiting the applicability of conventional MIMO. In addition, if these distributed MIMO (DMIMO) techniques can be made scalable (i.e., insensitive to the number of cooperating nodes), then they provide arguably the most effective means of realizing truly massive MIMO systems. In order to make the leap from conventional networks of uncoordinated omnidirectional transceivers to intelligent networks of coordinated DMIMO nodes, however, we must address fundamental bottlenecks in synchronization and tracking of independent oscillator and kinematic dynamics and develop scalable, low-latency techniques for coherent distributed transmission and reception. While these are formidable challenges, in this paper, we review a selection of recent research results in this area that give us cause for hope, and sketch a research agenda for building on this progress. Our goal is to provide a roadmap for further exploration, and to motivate such exploration by pointing out why the technical problems involved are interesting from the point of view of both intellectual and technological impact.

For concreteness, we focus on two canonical MIMO functions, beamforming and nullforming, and discuss what is entailed in achieving these in scalable fashion.

We are interested in DMIMO systems that do not rely on a high-speed backhaul connecting the cooperating nodes, and we are particularly interested in techniques that scale to a large number of nodes. Thus, our emphasis is quite different from DMIMO techniques proposed for base station cooperation, or “coordinated multipoint” (CoMP) [1], in cellular systems, which rely on a high-speed wired backhaul, and on explicit feedback from mobile nodes. A similar comment also applies to recent work on synthesizing virtual arrays using access points connected by a wired backhaul in a WiFi network [2], [3].

We organize the paper as follows. In Section II, we discuss the fundamental problem of synchronizing independently running oscillators, which is a building block for any coherent DMIMO technique. In Section III, we discuss the role of *aggregate* explicit feedback for scalable approaches to beamforming and nullforming. In Section IV, we discuss the alternative approach of *implicit* feedback via reciprocity, and the challenges in implementing it. Finally, while it is attractive to view DMIMO techniques as a direct generalization of centralized MIMO in any context, in practice, it must go hand in hand with synchronization-enabling protocols that support it. In Section VII, we discuss such protocol implications in the context of two *concept systems* in which DMIMO plays an essential role (i.e., centralized MIMO is simply not feasible): synthesizing a distributed base station for white space communication, and synthesizing a virtual array for reachback in emergency or disaster relief.

II. HOW TO SYNCHRONIZE OSCILLATORS

The first step to enabling a cluster of cooperating nodes cohere at radio frequency (RF) is to synchronize their oscillators in frequency and phase. A fairly generic model for achieving this (which lends itself to many different protocol implementations) is for the nodes in the cluster to synchronize to a common beacon. Such a beacon could be broadcast (typically intermittently) by a “master” node in the cluster, or by a distant node that the cluster is beamforming towards. Naturally, we would like to reduce the beacon overhead. A fundamental question, therefore, is how far we can reduce the duration and duty cycle of the beacon, without compromising synchronization performance. Before we answer this, we need a model for the oscillators whose frequency and phase are to be tracked.

Assume that a beacon of duration T_{est} is sent every T_{slot} (abbreviated as T_s henceforth) seconds, as shown in Figure 2. Let us consider the behavior of a particular node in the cluster. Let ϕ_t, ω_t denote the unwrapped phase and frequency of this node (relative to the beacon reference) at the time of the t th beacon. We model their evolution using the following standard linear state space model:

$$\begin{pmatrix} \phi_{t+1} \\ \omega_{t+1} \end{pmatrix} = \begin{pmatrix} 1 & T_s \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \phi_t \\ \omega_t \end{pmatrix} + \mathbf{p}_t \quad (1)$$

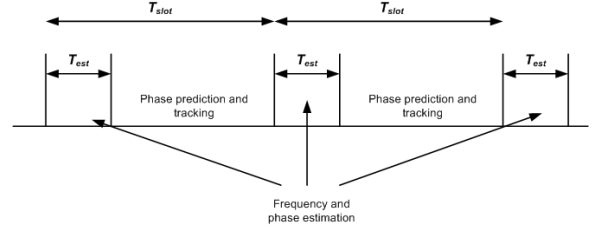


Fig. 2. We wish to track frequency and phase using intermittent beacons.

where \mathbf{p}_t is process noise modeled as Gaussian with zero mean and covariance matrix

$$\mathbf{Q} = \omega_c^2 q_1^2 \begin{pmatrix} T_s & 0 \\ 0 & 0 \end{pmatrix} + \omega_c^2 q_2^2 \begin{pmatrix} T_s^3/3 & T_s^2/2 \\ T_s^2/2 & T_s \end{pmatrix}$$

and where ω_c denotes the carrier frequency (in radians/second). The first term corresponds to phase drift, while the second corresponds to frequency drift. For any class of oscillators, we assume that nominal values for the parameters q_1 and q_2 governing these drifts are known (determined by measurements). The effect of mobility is easily incorporated within this framework, with the frequency state variable ω_t including Doppler shifts due to relative motion between the beacon emitter and the node implementing the tracking.

If noisy measurements of ϕ_t and ω_t were available, a standard Kalman filtering/prediction formulation could be applied for tracking the frequency and phase. However, there are two important observations regarding the estimates we can derive from a beacon signal:

- 1) We can only estimate the phase ϕ_t modulo 2π , so that our measurement model is inherently nonlinear.
- 2) The errors in the phase and frequency estimates have very different characteristics. Specifically, the Cramer-Rao lower bounds on these estimation errors scale as follows as a function of the post-integration SNR and the length of the estimation interval [4]:

$$\sigma_\phi^2 \sim \frac{1}{SNR}, \quad \sigma_\omega^2 \sim \frac{1}{T_{est} \times SNR}$$

Thus, we can, in principle, shrink the estimation interval indefinitely without impacting the phase estimation accuracy, as long as we increase the beacon power so as to keep the post-integration SNR large enough. However, as we shrink T_{est} , the frequency estimation error becomes large. In the limit of very small estimation intervals, therefore, we can only rely on phase measurements. Of course, in such a regime, we become vulnerable to the problem of *frequency aliasing*: based on wrapped phase measurements at intervals of T_s , we cannot distinguish between two frequencies that are separated by $2\pi/T_s$ (radians/second).

These problems are far from insurmountable. Through simulations in [5] and an over-the-air demonstration of distributed beamforming in [6], we have shown that a standard Kalman filtering framework can be applied if we choose parameters that avoid frequency aliasing. This design employs

good oscillators and a start-up sequence of measurements that removes frequency ambiguity, and thereafter employs phase-only measurements based on very short measurement intervals to track both phase and frequency. An alternative approach, successfully demonstrated in our software-defined radio testbed [7], [8] (where the oscillator quality is poor), is to employ long enough measurement intervals to enable frequency measurements which are reliable enough to overcome frequency aliasing. Robustness to the nonlinearity of the phase measurement is obtained by using an extended Kalman filter: this limits the size of the innovations associated with phase measurements, and hence the effect of phase unwrapping errors.

Open issues: While the work done so far assures us that tracking oscillators is eminently feasible, the fundamental open question is how far we can drive down the overhead of doing so. For example, if we shrink the measurement interval, then at some point we must operate with phase-only measurements. And as we reduce the duty cycle, the frequency aliasing problem becomes more acute. How far can we go down this road if we abandon the comfort zone of the Kalman filter and consider nonlinear filtering techniques such as the particle filter (or so-called Rao-Blackwellized versions of it that exploit the underlying state space structure to the extent possible)? This is a topic of ongoing research, promising results from which are to be reported in forthcoming publications.

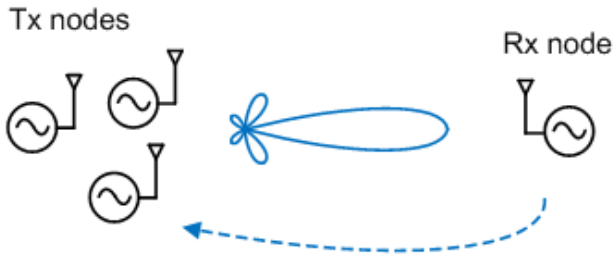


Fig. 3. Transmitters adapting based on aggregate feedback broadcast by the receiver.

III. THE ROLE OF EXPLICIT FEEDBACK

Let us first consider the problem of distributed transmit beamforming to a distant destination. In principle, if we know the locations of the transmitting nodes and that of the destination precisely, and if there is no multipath, then we can compute the beamforming weights to be used by each transmitter precisely. In practice, we never know the locations with sufficient precision, and the channel inevitably has multiple paths. Thus, some form of feedback regarding the channel is essential for transmit beamforming. In current centralized MIMO systems, explicit feedback to a transmitter with multiple antennas typically consists of quantized estimates of the channel from each antenna to the receiver. Such per-antenna explicit feedback does not scale well for DMIMO: the

overhead scales with the number of cooperating nodes, and the feedback protocol becomes dependent on the number of nodes. A far more attractive approach in such settings is *aggregate* feedback which can be broadcast from the destination to all cooperating transmitters, as shown in Figure 3.

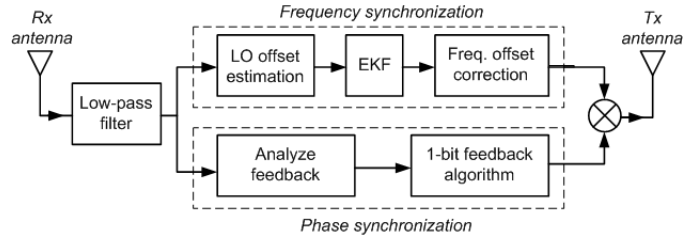


Fig. 4. The transmitters in [8] employ the one-bit feedback algorithm for phase alignment at the receiver, and an EKF for frequency synchronization.

Almost a decade ago [9], we showed that a particularly simple form of aggregate feedback (one bit per iteration) can be remarkably effective. Each cooperating transmitter independently perturbs its carrier phase in each iteration. The receiver broadcasts a single bit saying whether the resulting received signal strength is “better” or “worse” than before. If “better,” each transmitter keeps the phase perturbation; if “worse,” then each transmitter undoes its phase perturbation. This algorithm was shown to converge [10] with probability one, and forms the basis for distributed transmit beamforming for demonstrations on multiple testbeds, including recent demos on software-defined radios [7], [8]. A block diagram of a transmitter is shown in Figure 4: the feedback packets do double duty, with the waveforms in the packets providing the measurements of phase and frequency driving the state space models discussed in Section II for synchronizing the oscillators of the cooperating transmitters, and the payload (a single bit) to adapt the phases so as to form a beam at the destination. A typical prototype result [8] showing 3 transmit nodes cohering at a destination (i.e., with amplitudes adding up in phase) is shown in Figure 5.

Open issues: While the success of the one-bit algorithm is startling, such parsimonious feedback restricts the speed of convergence. How much can we gain if we feed back more information (e.g., the complex amplitude of the received signal)? And what is the most elegant generalization to dispersive, wideband channels? For the latter, a natural approach is to parallelize using OFDM, with explicit feedback for each subcarrier. However, it remains an open question as to how best to track and adapt to oscillator dynamics and channel time variations, especially when the time scales in the two settings are different.

IV. PRE-SYNCHRONIZATION AND IMPLICIT FEEDBACK

For time division duplex communication, *implicit feedback* from channel reciprocity can be used instead of, or in addition to, explicit feedback. Retrodirective antenna arrays [11], [12] reflect or retransmit an incident wave from a source directly

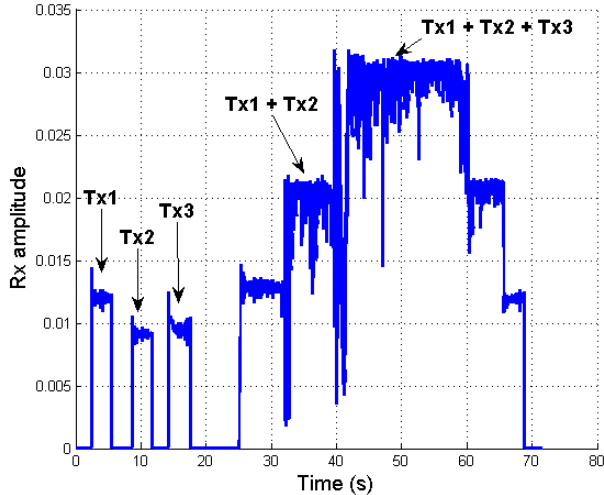


Fig. 5. The transmitters in [8] employ the one-bit feedback algorithm for phase alignment at the receiver, and an EKF for frequency synchronization.

back to the source without any prior knowledge of the source’s location; well-known examples the passive corner reflector, the Van Atta array [13] and the Pon array [14]. For distributed transmit beamforming, the transmit nodes could, in principle, estimate the phase of a beacon from the intended receiver, and form a beam back to the destination with no additional overhead. This approach is particularly attractive in highly mobile settings, where explicit feedback strategies may have difficulty keeping up. Our prior work in this direction includes theoretical studies in [15]–[18], with experimental results for acoustic signals reported in [19]. However, there are two major challenges in realizing the gains from implicit feedback at RF. First, the cooperative transmitters must make the channel measurements at “the same time,” with timing errors within a small fraction of a carrier cycle. This requirement is *orders of magnitude* more stringent than that achievable through off-the-shelf techniques such as GPS or NTP (network time protocol). Second, mismatch in the transmit and receive chains at each cooperative node must be systematically accounted for, in order to truly realize the benefits of reciprocity.

Recent progress in timing synchronization: We have some good news to report regarding the first challenge. Bounds by Weiss and Weinstein [21], [22] show that, in addition to the obvious improvement of timing accuracy with SNR and time-bandwidth product, accuracy also improves with carrier frequency if the SNR is high enough. These bounds indicate that precision of the order of *picoseconds* can be attained for reasonable values of system parameters (e.g., 10 dB SNR, 50 MHz bandwidth, duration 10 μ s, $f_c = 1$ GHz). In recent results [20], we have shown that we can attain the Weiss-Weinstein bound using fairly simple algorithms; see Figure 6. These algorithms have also been successfully demonstrated to attain accuracies of the order of picoseconds using a hardware prototype [23]. It is worth commenting on

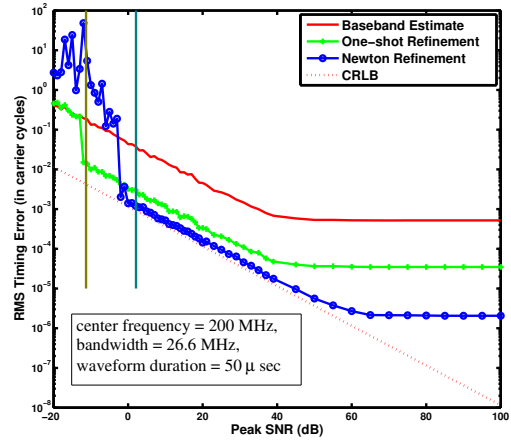


Fig. 6. The timing synchronization algorithm in [20] follows the Weiss-Weinstein bound.

the common structure of the bounds and the algorithms used to attain them. In estimating a continuous-valued parameter, once we are “close enough,” the performance is predicted well by the Cramer-Rao Bound (CRB). However, in order to get close enough, we typically must do hypothesis testing across coarsely quantized bins. The Ziv-Zakai bound (ZZB), which Weiss and Weinstein adapted for the specific scenario of timing estimation using passband signals in AWGN, accounts for errors in such hypothesis testing, and tends to the CRB when the SNR is high enough for the probability of choosing the wrong bin is small enough. A natural approach to designing an algorithm for attaining these bounds is to follow this logic, first doing hypothesis testing to determine the bin, and then refining the estimate within the bin (which should work well if the SNR is high enough). In [20], we show that Newton-Raphson based approaches to the latter refinement work well. The additional twist, however, is that, if this refinement gets us to within the right carrier cycle (which happens at high enough SNR), then an additional refinement based on the carrier phase can get the timing error to within a small fraction of a carrier cycle. Thus, by exchanging messages among nodes within a cluster of cooperating nodes, there is every reason to hope that the pre-synchronization accuracy required for acquiring accurate implicit feedback can be achieved, since the SNR of such local communication is expected to be high.

Open issues: Scaling pre-synchronization protocols and algorithms to large node clusters is an open design issue, as is the problem of accounting for any channel dispersion in the local channels connecting the cooperating nodes. Putting together one-shot estimates based on local message exchanges into a framework including state space tracking is also an open problem. In addition, a systematic framework must be developed for learning and compensating for the mismatch in the transmit and receive chains at each cooperating node.

V. DISTRIBUTED RECEIVE BEAMFORMING

Distributed reception has a much longer history than distributed transmit beamforming. It is already used in synthetic aperture radio astronomy, e.g. the Very Large Array, where signals recorded at multiple receivers are forwarded to a central collection point for coherent combination. Such systems use high-speed wired/fiber-optic backhauls with high throughput and have no real-time latency constraints. Direct application of such approaches is not suited to a real-time all-wireless DMIMO system, since the local network capacity needed for pooling observations scales with the number of cooperating receivers.

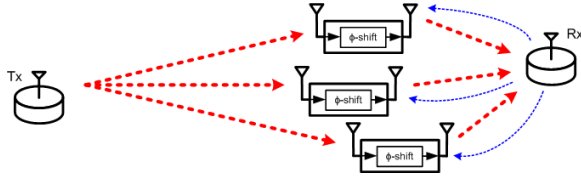


Fig. 7. Scalable distributed reception with over the air combining.

Scalable distributed reception via amplify-forward: Instead of pooling observations from the cooperating receivers after transporting them to a common processor node, we can linearly process these observations “in the air,” with each receiver relaying (potentially a transformed version of) its observations to a common processor node, as shown in Figure 7. Consider the problem of distributed receive beamforming. In a receive cluster of $N + 1$ nodes receiving a message from a distant transmitter, N nodes are designated as relays and one node is designated as the receiver, or processor. The key idea is that all of the relays forward the signal received on the long-range forward link over short-range local links to the processor *at the same time*, filtering and adjusting the phases of the forwarded signals so that they coherently combine at the processor. By using simultaneous coherent relaying, we obtain a scalable architecture in which the number of degrees of freedom for local communication need not increase with the number of cooperating nodes. We have thus transformed the problem of distributed reception on the long-range link into that of distributed transmission on the short-range links. A rudimentary prototype of distributed receive beamforming, utilizing one-bit feedback from the processor node to the relays to enable distributed transmit beamforming on the short-range links, has been recently reported in [24]. An interesting observation is that, for TDD operation on the long and short links, the relay nodes do not have to be synchronized in frequency: their frequency offsets with respect to the transmitter and the processor nodes “cancel out.” Distributed reception using amplify-forward relays can also be used to create spatial degrees of freedom, as observed in a paper at this workshop last year [25].

Open issues: “Over the air” combining for distributed reception restricts us to *linear* collaborative processing, and it is important to determine how effectively such an architecture

can handle complications such as interference and channel dispersion, in terms of the feedback structure required and the performance achievable. Also, while frequency offsets at the relays cancel out, the accumulated phase noise due to the delay between reception on the long link and transmission on the short link has an adverse impact on relay adaptation which needs to be quantified. Finally, when the number of cooperating nodes is not too large, explicit message exchanges with moderate backhaul requirements may be feasible, especially in light of recent results showing that performance close to that of ideal receive beamforming is possible with heavily quantized observations [26].

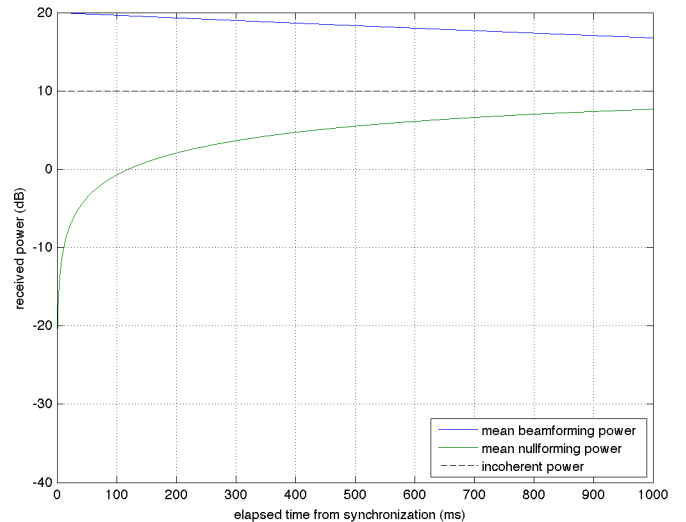


Fig. 8. Nulls deteriorate much faster than beams. The figure shows mean beamforming and nullforming powers as a function of elapsed time from synchronization for 10 cooperating nodes with independently drifting oscillators independently drifting with a standard deviation of 62 ps in one second. The carrier frequency is 2.4 GHz.

VI. DISTRIBUTED NULLFORMING

While much recent attention in DMIMO has been focused on beamforming, it is important to also develop a robust design framework for distributed nullforming. The latter is critical for interference suppression and avoidance, and when used in conjunction with distributed beamforming, enables simple linear precoding techniques for SDMA. Figure 8 illustrates that nulls are far more sensitive to phase errors than beams, so that it is expected to be challenging to achieve deep nulls in a DMIMO system. On the other hand, deep nulls can be achieved with a small number of cooperating nodes (e.g., we need only 2 nodes to achieve a single null), which may lead to interesting DMIMO system concepts that do not require a large number of cooperating nodes. For example, we may have a few powerful transmitters using uncoordinated broadcast to serve a region (even if the transmitters are not synchronized in phase, they can provide power pooling gains), while coordinating to steer nulls towards certain protected receivers.

Recent progress in distributed nullforming: There have been several recent studies of distributed nullforming algorithms. In [27]–[29], explicit feedback-based distributed nullforming systems accounting for channel state uncertainty were studied with numerical results demonstrating that nulls more than 20 dB below the incoherent power level could be achieved with low channel measurement overhead, infrequent measurement intervals, and significant feedback latency. A simplifying assumption in [27], [28] was that each receiver individually tracked the channels from the transmit array. This approach was suboptimal since it does not exploit the statistical coupling of the pairwise phase and frequency offsets across all of the receive nodes. The performance of distributed nullforming with global tracking was studied in [30] where it was shown that significant gains in null depth can be achieved at the expense of the additional overhead of exchanging observations. In [31], random vector quantization methods were studied for nullforming with very low rate feedback. Analytical results showed that, for a system with K simultaneous nulls and random codebooks with $N = 2^B$ precoding vectors, the mean received power in the nulls was upper bounded by $N^{-1/K} = 2^{-B/K}$ even without coordination among the receivers.

While beamforming requires only that the phases of the received signals match at the beam location, nullforming requires a much more intricate combination of phases and magnitudes. Accordingly, [27], [28] both require that each transmitter know the channel state seen by all the transmitters. The preliminary work reported in [29] relaxes this requirement by proposing a decentralized nullforming algorithm that requires each transmitter to only know its channel state to the null location. The algorithm exploits an aggregate feedback from the receiver to the transmitter array; specifically, the total baseband signal at the receiver. To achieve unsynchronized power pooling gains away from the null location, each transmitter adapts only its transmitted phase. The algorithm involves the gradient descent minimization of the total received signal power, and is provably convergent, even though the cost function being minimized is nonconvex under phase only adaptation.

Open issues and ongoing work: The phase only adaptation algorithm of [29], makes the assumption that the received signal power from each individual transmitter is the same. Ongoing work relaxes this assumption. Open issues include extending the results of [29] to achieve nulls at multiple locations. A related issue is to go beyond phase only adaptation to adapt signal magnitudes as well. This brings forth the additional challenge of avoiding adaptation strategies that drive the transmitted signals to zero. A possible avenue is to impose additional constraints such as requiring a prescribed beam power at another location or a precise baseband signal at this location. The former constraint leads to a nonconvex problem, as suggested in the literature surrounding the so called constant modulus adaptive algorithm, [32]. The latter, though convex, sacrifices potentially useful degrees of freedom.

VII. CONCEPT SYSTEMS

While DMIMO can be used to enhance conventional MIMO by increasing the effective number of antennas, perhaps the most compelling applications are to settings in which large centralized arrays are simply not possible. We briefly discuss two such scenarios, or “concept systems,” here. The first is a distributed base station (DBS) for rural broadband, where a number of low-cost component nodes (connected locally by, for example, WiFi at 5 GHz) are used to synthesize a large array for white space frequencies (50-800 MHz), used to communicate with distant subscribers at ranges of the order of 10s of kilometers. A centralized array with a large number of elements would be too bulky at such large carrier wavelengths (especially at the lower end of the white space spectrum), hence DMIMO is the only means of reaping the benefits of MIMO, such as enhanced spatial reuse via beamforming, SDMA, and interference reduction for primary users. The second concept system is based on an emergency scenario in which a cooperative cluster of mobile devices (e.g., belonging to a group of lost hikers) is communicating with a distant base station (or a search and rescue vehicle) that none of the devices can reach on its own.

Open issues: Even as DMIMO algorithms mature, realizing such concept systems requires detailed design and evaluation of synchronization-enabled protocols crafted for the system at hand. It is fair to say that every aspect of such designs is wide open.

VIII. CONCLUSIONS

We believe that the time has come for a concerted effort by the research community to transform DMIMO from a collection of information-theoretic constructs to a suite of algorithms, architectures and protocols that can truly enable us to build massive MIMO systems that sidestep form factor constraints for individual nodes. We hope that this paper provides a sense of the exciting opportunities in this area, as well as a starting point for further exploration.

Acknowledgements

We would like to acknowledge current and past collaborators on the work we have reviewed in this paper. Current and former graduate students and postdoctoral researchers whose prior and ongoing research we have discussed include Andrew Irish, Amy Kumar, Dr. Francois Quitin, Dr. Mahboob Rahman, Maryam Eslami Rasekh. We also acknowledge our highly fruitful collaborations with the BBN/Raytheon team led by Dr. Patrick Bidigare.

REFERENCES

- [1] P. Marsch and G. P. Fettweis, Eds., *Coordinated Multi-Point in Mobile Communications: From Theory to Practice*. Cambridge University Press, 2011.
- [2] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, “Achieving high data rates in a distributed mimo system,” in *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*, ser. Mobicom ’12. New York, NY, USA: ACM, 2012, pp. 41–52. [Online]. Available: <http://doi.acm.org/10.1145/2348543.2348552>

- [3] H. S. Rahul, S. Kumar, and D. Katabi, "Jmb: Scaling wireless capacity with user demands," in *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '12. New York, NY, USA: ACM, 2012, pp. 235–246. [Online]. Available: <http://doi.acm.org/10.1145/2342356.2342401>
- [4] D. Rife and R. Boorstyn, "Single-tone parameter estimation from discrete-time observations," *IEEE Trans. on Information Theory*, vol. 20, no. 5, pp. 591–598, September 1974.
- [5] D. Brown, P. Bidigare, and U. Madhow, "Receiver-coordinated distributed transmit beamforming with kinematic tracking," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, march 2012, pp. 5209–5212.
- [6] P. Bidigare, M. Oyarzun, D. Raeman, D. Cousins, D. Chang, R. O'Donnell, and D.R. Brown III, "Implementation and demonstration of receiver-coordinated distributed transmit beamforming across an ad-hoc radio network," in *Proceedings of the 46th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, November 4-7, 2012, p. tbd.
- [7] F. Quitin, U. Madhow, M. Rahman, and R. Mudumbai, "Demonstrating distributed transmit beamforming with software-defined radios," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2012 IEEE International Symposium on a*, june 2012, pp. 1–3.
- [8] 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, March 2013.
- [9] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Scalable feedback control for distributed beamforming in sensor networks," in *IEEE International Symp. on Information Theory (ISIT)*, Adelaide, Australia, September 2005, pp. 137–141.
- [10] —, "Distributed transmit beamforming using feedback control," *IEEE Trans. on Information Theory*, vol. 56, no. 1, pp. 411–426, January 2010.
- [11] R. Miyamoto and T. Itoh, "Retrodirective arrays for wireless communications," *Microwave Magazine, IEEE*, vol. 3, no. 1, pp. 71–79, mar 2002.
- [12] Y.-C. Guo, X.-W. Shi, and L. Chen, "Retrodirective array technology," *Progress In Electromagnetics Research B*, vol. 5, pp. 153–167, 2008.
- [13] E. D. Sharp and M. A. Diab, "Van Atta reflector array," *IRE Transactions on Antenna Propagation*, vol. 8, no. 4, pp. 436–438, 1960.
- [14] C. Pon, "Retrodirective array using the heterodyne technique," *IEEE Trans. on Antennas and Propagation*, vol. 12, no. 2, pp. 176–180, Mar 1964.
- [15] I. Ozil and D.R. Brown III, "Time-slotted round-trip carrier synchronization," in *Proceedings of the 41st Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, November 4-7 2007, pp. 1781–1785.
- [16] D.R. Brown III and H.V. Poor, "Time-slotted round-trip carrier synchronization for distributed beamforming," *IEEE Trans. on Signal Processing*, vol. 56, no. 11, pp. 5630–5643, November 2008.
- [17] R. Preuss and D.R. Brown III, "Retrodirective distributed transmit beamforming with two-way source synchronization," in *Proceedings of the Conference of Information Sciences and Systems (CISS 2010)*, Princeton, NJ, March 2010, pp. 1–6.
- [18] R. Preuss and D. Brown, "Two-way synchronization for coordinated multicell retrodirective downlink beamforming," *Signal Processing, IEEE Transactions on*, vol. 59, no. 11, pp. 5415–5427, nov. 2011.
- [19] D. Brown, B. Zhang, B. Svirchuk, and M. Ni, "An experimental study of acoustic distributed beamforming using round-trip carrier synchronization," in *Phased Array Systems and Technology (ARRAY), 2010 IEEE International Symposium on*, October 2010, pp. 316–323.
- [20] P. Bidigare, U. Madhow, R. Mudumbai, and D. Scherber, "Attaining fundamental bounds on timing synchronization," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, march 2012, pp. 5229–5232.
- [21] A. Weiss and E. Weinstein, "Fundamental limitations in passive time delay estimation—part i: Narrow-band systems," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 31, no. 2, pp. 472–486, apr 1983.
- [22] E. Weinstein and A. Weiss, "Fundamental limitations in passive time-delay estimation—part ii: Wide-band systems," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 32, no. 5, pp. 1064–1078, oct 1984.
- [23] P. Bidigare, S. Pruessing, D. Raeman, D. Scherber, U. Madhow, and R. Mudumbai, "Initial over-the-air performance assessment of ranging and clock synchronization using radio frequency signal exchange," in *Statistical Signal Processing Workshop (SSP), 2012 IEEE*, aug. 2012, pp. 273–276.
- [24] 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.
- [25] A. Irish, F. Quitin, U. Madhow, and M. Rodwell, "Sidestepping the rayleigh limit for los spatial multiplexing: A distributed architecture for long-range wireless fiber," in *Information Theory and Applications Workshop (ITA), 2013*, Feb 2013, pp. 1–6.
- [26] D. Brown, M. Ni, U. Madhow, and P. Bidigare, "Distributed reception with coarsely-quantized observation exchanges," in *Information Sciences and Systems (CISS), 2013 47th Annual Conference on*, March 2013, pp. 1–6.
- [27] D. R. Brown, P. Bidigare, S. Dasgupta, and U. Madhow, "Receiver-coordinated zero-forcing distributed transmit nullforming," in *Statistical Signal Processing Workshop (SSP), 2012 IEEE*, aug. 2012, pp. 269–272.
- [28] D. R. Brown, U. Madhow, P. Bidigare, and S. Dasgupta, "Receiver-coordinated distributed transmit nullforming with channel state uncertainty," in *Information Sciences and Systems (CISS), 2012 46th Annual Conference on*, march 2012, pp. 1–6.
- [29] M. M. Rahman, R. Mudumbai, and S. Dasgupta, "A scalable feedback-based approach to distributed nullforming," in *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, Springer Verlag, Berlin-Heidelberg*, sept. 2013, pp. 78–84.
- [30] D.R. Brown III and R. David, "Receiver-coordinated distributed transmit nullforming with local and unified tracking," in *Proceedings of the 2014 International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2014)*, Florence, Italy, May 4-9 2014.
- [31] D.R. Brown III and D. Love, "On the performance of MIMO nullforming with random vector quantization limited feedback," *IEEE Transactions on Wireless Communications*, 2014, accepted to appear.
- [32] C. R. Johnson, P. Schniter, T. J. Endres, J. D. Behm, D. R. Brown, and R. A. Casas, "Blind equalization using the constant modulus criterion: a review," *Proceedings of the IEEE*, vol. 86, no. 10, pp. 1927–1950, 1998.