

Distributed Transmit Beamforming: Challenges and Recent Progress

Raghuraman Mudumbai, University of California at Santa Barbara

D. Richard Brown III, Worcester Polytechnic Institute

Upamanyu Madhow, University of California at Santa Barbara

H. Vincent Poor, Princeton University

ABSTRACT

Distributed transmit beamforming is a form of cooperative communication in which two or more information sources simultaneously transmit a common message and control the phase of their transmissions so that the signals constructively combine at an intended destination. Depending on the design objectives and constraints, the power gains of distributed beamforming can be translated into dramatic increases in range, rate, or energy efficiency. Distributed beamforming may also provide benefits in terms of security and interference reduction since less transmit power is scattered in unintended directions. Key challenges in realizing these benefits, however, include coordinating the sources for information sharing and timing synchronization and, most crucially, distributed carrier synchronization so that the transmissions combine constructively at the destination. This article reviews promising recent results in architectures, algorithms, and working prototypes which indicate that these challenges can be surmounted. Directions for future research needed to translate the potential of distributed beamforming into practice are also discussed.

INTRODUCTION

In wireless communication systems, transmit beamforming refers to a technique in which an information source transmits a radio frequency signal over two or more antennas and aligns the phases of the transmissions across the antennas such that, after propagation, the signals combine constructively at the destination. Fixing the power radiated by a given antenna element, ideal transmit beamforming with N antennas results in an N^2 -fold gain in received power. Compared to single-antenna transmission, transmit beamforming can therefore yield increased range (an N -fold increase for free space propagation), increased rate (an N^2 -fold increase in a power-limited regime), or increased power efficiency (an N -fold decrease in the net transmitted power for a fixed desired received power). In

addition, since more power is directed in the desired direction, less is scattered in undesired directions, resulting in reduced interference and increased security.

Given the many advantages of transmit beamforming, it is natural to ask whether it can be emulated in distributed fashion using a network of cooperating single-antenna sources. In order to operate as a “distributed transmit beamformer,” the sources must agree on a common message, transmit it at the “same time,” synchronize their carrier frequencies, and control their carrier phases so that their signals combine constructively at the destination. Hence, practical realization of this concept requires the development of implementable distributed techniques for information sharing, timing synchronization, and carrier synchronization. While these constitute a daunting set of challenges, recent results from several different research groups provide promising approaches for addressing them. The goal of this article is to take stock of the current state of the art, and to suggest directions for future research in the design and implementation of wireless networks that exploit distributed beamforming.

In addition to the N^2 -fold power gain from distributed beamforming, there is also a potential advantage in terms of wireless propagation. Consider the Friis formula for free-space propagation,

$$P_R = P_T G_T G_R \frac{\lambda^2}{16\pi^2 R^2},$$

where P_T and P_R are the transmit and receive powers, respectively, G_T and G_R are the directivity gains of the transmit and receive antennas, respectively, R is the range between the antennas, and λ is the carrier wavelength. For fixed antenna gains, the propagation loss

$$\frac{P_T}{P_R}$$

is smaller at longer wavelengths. However, antenna gains take the form


This research was supported in part by the U.S. National Science Foundation under grants CCF-04-47743, ANI-03-38807, CNS-06-25637, and CNS-0520335, as well as by UCSB's Institute for Collaborative Biotechnologies under grant DAAD19-03-D-0004 from the U.S. Army Research Office.

$$G = \frac{4\pi A}{\lambda^2},$$

where A is the effective area. Thus, in order to maintain a given directivity as wavelength increases, one must also scale the effective area of each antenna by λ^2 , which can make longer wavelengths unattractive. With distributed transmit beamforming, it is possible to have the best of both worlds: low propagation loss by operating at a long wavelength and high directivity by exploiting the natural spatial distribution of the cooperating nodes to emulate a large antenna array. While this argument is presented for free space propagation, longer wavelengths provide even more of an advantage in cluttered environments, since the radio waves are better able to diffract around obstacles.

As an example application, consider the scenario shown in Fig. 1, where a terrestrially deployed network of low-power single-antenna sensor nodes collects measurements and transmits these measurements to an overflying unmanned aerial vehicle (UAV) using a carrier frequency of 3 GHz and a bandwidth of 10 MHz. For a sensor transmit power of -10 dBm, the received power at an altitude of 3000 m (typical for intermediate range UAVs) is -110 dBm, assuming a sensor transmit antenna gain of $G_T = 2$ dBi and receive antenna gain $G_R = 10$ dBi for the aircraft. For a receiver noise figure of 6 dB, the noise power is -97 dBm (thermal noise at 300 Kelvin has a power spectral density of -173 dBm/Hz). Thus, the signal-to-noise ratio (SNR) for a single sensor transmission is -13 dB, making communication with reasonable spectral efficiency infeasible. On the other hand, the SNR increases to $+13$ dB if 20 sensor nodes form a distributed transmit beamformer. This could enable, for example, upload of image/video data or summaries of sensor data gathered over days or even months. Other interesting applications include reach-back using low-power soldier radios in battle-field communication, and collaboration between subscriber terminals for uplink transmission to a base station receiver, especially in rural or disaster recovery settings where longer range might be required.

As the preceding examples indicate, distributed transmit beamforming has the potential to enable fundamentally new functionalities in wireless communication and sensor networks. In the remainder of this article we discuss some of the technical issues that must be addressed in order to realize this potential. We review recent progress on the crucial distributed carrier synchronization problem in the next section and later describe two working prototypes that suggest this problem is solvable. We then discuss the characteristics of beam patterns realizable using distributed transmit beamforming with randomly placed sources. These results lay the foundation for “physical layer” feasibility of distributed transmit beamforming. In the following section we discuss the cross-layer design considerations for information sharing and coordination among sources. We end in the final section with a discussion of directions for future research.



■ Figure 1. Sensor network transmitting measurements to an overflying aircraft.

DISTRIBUTED CARRIER SYNCHRONIZATION

A key distinguishing feature of distributed transmit beamforming with respect to conventional beamforming is that each source node in a distributed beamformer has an independent local oscillator (LO). These LOs are typically generated by multiplying the frequency of a crystal oscillator up to a fixed nominal frequency. Carrier frequencies generated in this manner, however, typically exhibit variations on the order of 10–100 parts per million (ppm) with respect to the nominal. If uncorrected, these frequency variations among sources are catastrophic for transmit beamforming since the phases of the signals may drift out of alignment over the duration of the transmission and may even result in destructive combining at the destination. The first goal, therefore, is to synchronize the carrier frequencies for the different sources to minimize or eliminate frequency offset.

One approach to frequency synchronization is to employ a *master-slave* architecture [1, 2], where “slave” source nodes use phase-locked loops (PLLs) to lock to a reference carrier signal broadcast by a “master” source node. Alternatively, the destination node could broadcast a reference carrier to facilitate frequency synchronization among the source nodes [3–5]. A source node that estimates its frequency offset to be Δf can multiply its complex baseband transmitted signal by $e^{-j2\pi\Delta f t}$, where the operation can be implemented in a digital signal processor (DSP) prior to digital-to-analog conversion and carrier multiplication. Depending on the stability of the sources’ oscillators, the process of frequency synchronization may need to be repeated, and should be inherent to any networking protocol built around distributed beamforming.

Once frequency synchronization is achieved, the phase of the transmissions from the different

Regardless of the synchronization approach, it is known that beamforming gains are quite robust to moderate errors in phase alignment. For example, 90 percent of an ideal two-antenna beamforming power gain is attained even with phase offsets of the order of 30°.

sources must be synchronized to arrive with “reasonable” alignment at the destination. To understand why carrier phase synchronization is critical for distributed beamforming, consider first transmit beamforming using an N -element centralized array. To send a complex baseband message signal $s(t)$, the signal transmitted from antenna i is $w_i s(t)$, and the received signal is $\sum_i w_i h_i s(t)$, where h_i is the complex channel gain from antenna i to the receiver. The received SNR is therefore proportional to $|\sum_i w_i h_i|^2$. Given a constraint on the total transmitted power $\sum_i |w_i|^2$, it can be shown that the SNR is maximized by choosing $w_i \propto h_i^*$, i.e., $|w_i| \propto |h_i|$ and $\angle w_i = -\angle h_i$. Another option, appropriate for a peak power constraint per antenna element, is to use a fixed amplitude $|w_i| = w_{\max}$ and $\angle w_i = -\angle h_i$. When the channel gains are approximately equal in magnitude, both methods have similar performance: the received signal $|\sum_{i=1}^N w_i h_i| \propto N$, so that the received SNR scales as N^2 . In either case, the transmitter requires channel state information (CSI) regarding the $\{h_i\}$, with the phase $\angle h_i$ being the critical information required to obtain beamforming gains. Techniques for obtaining CSI at the transmitter fall into two broad categories: implicit feedback (e.g., using reciprocity in a time division duplexed [TDD] system), and explicit feedback, where the CSI is quantized and sent over a separate feedback channel. A detailed review of different beamforming techniques is given in [6].

In distributed beamforming scenarios the sources are assumed to be unsynchronized a priori. This lack of synchronization leads to ambiguous phase estimates at each source. To see this, consider first implicit channel feedback using reciprocity. Ignoring modulation and noise for simplicity, source node i receives the passband signal $\text{Re}(h_i e^{j2\pi f_c t})$. When this is down converted using the local oscillator (LO) at node i , using the quadrature carriers $\cos(2\pi f_c t + \theta_i)$ and $\sin(2\pi f_c t + \theta_i)$, the complex baseband channel estimate at node i will be $\hat{h}_i = h_i e^{-j\theta_i}$. Without carrier synchronization across nodes, the local oscillator phases $\{\theta_i\}$ may be modeled as independent and uniformly distributed over $(-\pi, \pi)$, which implies that the phase of the channel estimate contains no information about the actual channel phase. In other words, the channel phase cannot be disambiguated from the relative LO phase at node i with this approach. Now, suppose instead that the receiver measures the channel gains from each node, and feeds them back explicitly. When node i employs these explicit channel estimates, however, it must upconvert the baseband message using its LO, which means that it is effectively using the beamforming weight $w_i = h_i^* e^{j\theta_i}$. Again, the phase of the beamforming coefficient is essentially random without prior carrier synchronization across the nodes.

The preceding observations show that even under ideal timing synchronization across nodes, distributed beamforming is impossible without distributed carrier synchronization. Nonidealities in timing synchronization can also affect distributed beamforming, but the effects are easier to handle. Timing synchronization is required to ensure that all of the cooperating

nodes transmit the same symbol at a given time; timing errors between the nodes lead to misalignment between the symbols transmitted by each node, causing intersymbol interference (ISI) at the receiver. For relatively low data rates (say around 100 kb/s), the required level of timing synchronization can be obtained using well-known algorithms such as RBS [7]. These algorithms are capable of achieving accuracy on the order of 1 μ s with low complexity. For higher data rates, customized timing synchronization techniques might be needed to achieve the desired level of accuracy. Even with accurate timing synchronization, however, ISI can arise due to dispersive channels from each node to the receiver. A natural approach to handling this is multicarrier modulation: in this case, distributed beamforming would be performed separately for each subcarrier. For single-carrier modulation, we may wish to use transmit precoding to ensure that the same symbol sent by different transmitters appears at approximately the same time at the receiver. Thus, while timing synchronization does pose a challenge, it is not as fundamental a bottleneck as carrier synchronization; hence, we focus on the latter in this article.

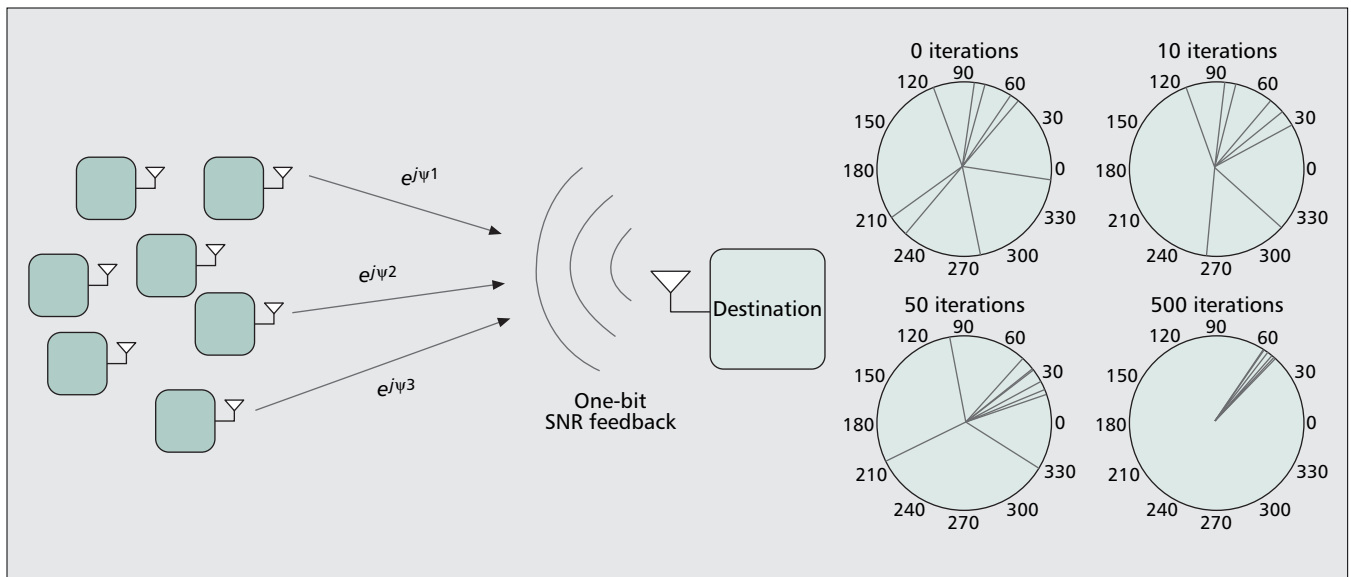
There are two basic approaches to phase synchronization distinguished by the interaction between the sources and the destination:

- **Closed-loop phase synchronization:** In closed-loop systems, the destination directly controls the phase alignment among the sources by measuring a function of the received phases of the source transmissions and then transmitting digital feedback signals to the sources to allow each source to compensate for its overall phase offset (LO and channel). Interaction among the sources can be minimal in closed-loop systems since the destination coordinates the synchronization process.
- **Open-loop phase synchronization:** In open-loop systems, the sources interact among themselves with only minimal signaling from the destination. Rather than providing feedback to be used for adapting the source phases, the destination may simply broadcast an unmodulated sinusoidal beacon to the sources. The sources use this beacon, as well as the signals from other source-source interactions, to achieve appropriate phase compensation for beamforming to the destination. The emphasis of open-loop systems is on using local interactions between the sources to minimize interaction with the distant destination.

Regardless of the synchronization approach, it is known that beamforming gains are quite robust to moderate errors in phase alignment. For example, 90 percent of an ideal two-antenna beamforming power gain is attained even with phase offsets on the order of 30° [1, 2].

FULL-FEEDBACK CLOSED-LOOP SYNCHRONIZATION

The first carrier synchronization scheme suitable for distributed beamforming is described in [3]. Carrier frequency synchronization is



■ **Figure 2.** One-bit feedback closed-loop carrier synchronization system.

achieved using a master-slave approach, with the intended destination acting as the master node. The unknown phase offset between the destination and the n th source node is corrected via a closed-loop protocol realized in the following steps:

- 1 The destination broadcasts a common master beacon to all source nodes.
- 2 Each source node “bounces” the master beacon back to the destination on a different frequency than the master beacon. The source nodes use distinct codes in a direct-sequence code-division multiple access (DS-CDMA) scheme in order to allow the destination to distinguish the received signals.
- 3 Upon reception of the bounced beacons, the destination estimates the received phase of each source relative to the originally transmitted master beacon. The destination quantizes these estimates, and then transmits the estimates via DS-CDMA to the source nodes in a “phase compensation message.” The phase compensation message may also contain clock correction information to facilitate symbol timing synchronization.
- 4 Each source receives the phase compensation message, extracts its own phase compensation estimate, and then adjusts its carrier phase accordingly.

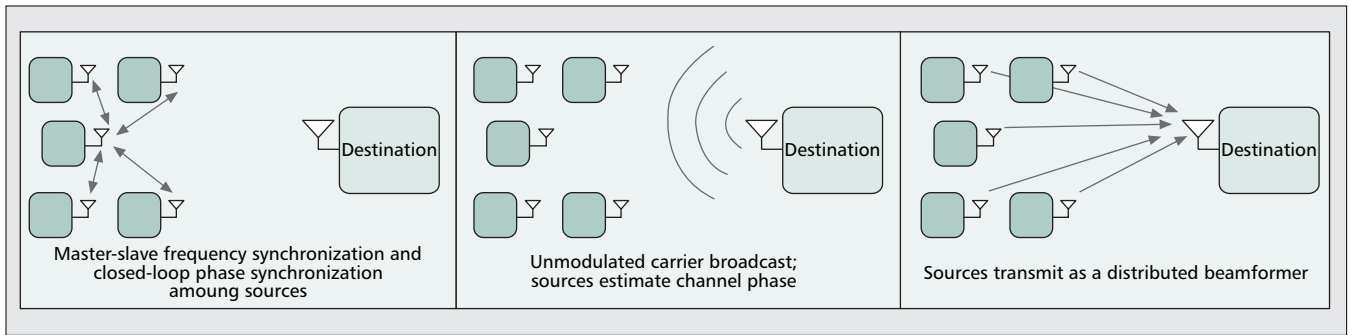
Assuming that the phase offsets have not changed significantly between the synchronization and beamforming intervals, the bandpass transmissions from each source will combine coherently when the sources transmit to the destination with compensated carrier phases. The effect of energy allocation between synchronization and information transmission on the error probability of digital signals transmitted by a distributed beamformer was also studied in [3]. The results showed that an optimal energy trade-off exists and that allocating too much or too little energy to carrier synchronization is inefficient.

ONE-BIT FEEDBACK CLOSED-LOOP SYNCHRONIZATION

The rate of feedback necessary to establish and maintain reasonable phase alignment among the sources in the full-feedback closed-loop carrier synchronization system described in [3] may be prohibitive in some scenarios. Recently, a closed-loop carrier synchronization system was proposed using only one bit of feedback for all source nodes [8, 9]. The basic idea behind the one-bit feedback closed-loop synchronization system shown in Fig. 2 is as follows:

- 1 Each source node adjusts its carrier phase randomly.
- 2 The source nodes transmit to the destination simultaneously as a distributed beamformer.
- 3 The destination estimates the SNR of the received signal.
- 4 The destination broadcasts one bit of feedback to the sources indicating whether its SNR is better or worse than before the sources adjusted their phases. If it is better, all source nodes keep their latest phase adjustments; otherwise, all sources undo their latest phase adjustments.

These four steps form one iteration of the system. Since each source retains only those random phase adjustments that lead to performance improvement, the algorithm may be viewed as a randomized ascent procedure; hence, the number of iterations to achieve a desired degree of phase convergence is a random variable. The average number of iterations required to achieve phase convergence was shown to scale roughly linearly with N , where N is the number of source nodes [8]. Numerical and analytical results in [8] also showed that 75 percent of the ideal beamforming amplitude is achieved in roughly $5N$ iterations on average. Under mild conditions on the distribution of the source nodes’ random phase adjustments, the one-bit feedback system was also shown to converge to full phase coherence with probability one [9]. Figure 2 shows the



■ **Figure 3.** Master-slave open-loop carrier synchronization system.

evolution of the received phases from each source node in one instance of the algorithm with $N = 10$ nodes. In this case, after 500 iterations, all the received phases are between 45° and 60° (i.e., a spread of 15°). This is sufficient to achieve approximately 99 percent of the beamforming gains. Note that the precise value to which the received phases converge is irrelevant to the beamforming process; it only matters that the *differences* between the phases converge to zero to achieve coherent combining.

Like the full-feedback closed-loop synchronization system, the one-bit feedback closed-loop system corrects the overall phase offset for each source caused by both the LO and the channel. As such, the iterations can be continued indefinitely to track both channel time variations and oscillator drift. Moreover, while the system described in [8, 9] assumes that the sources are already synchronized in frequency (e.g., by using the master-slave approach), this approach can be extended to also explicitly include carrier frequency synchronization [10].

The simplicity and scalability of the one-bit feedback synchronization system make it an attractive candidate for practical implementation where closed-loop feedback from the destination is possible. Two experimental prototypes based on the one-bit feedback closed-loop approach are discussed later.

MASTER-SLAVE OPEN-LOOP SYNCHRONIZATION

In some applications, closed-loop feedback from the destination to the sources is undesirable due to the relatively high cost of communication over this link and the increased complexity incurred at the destination. Open-loop carrier synchronization systems minimize interaction between the source nodes and the destination by increasing the level of inter-source interactions. One open-loop approach inspired by master-slave frequency synchronization was described in [2] and is illustrated in Fig. 3.

In open-loop master-slave synchronization, one source node is designated as the master and the remaining source nodes are slaves. For frequency synchronization, the master source node broadcasts a sinusoidal signal to the slave nodes, and each slave node estimates and corrects its frequency offset. Phase synchronization among the source nodes is then achieved through a closed-loop method similar to [3] except that a TDD protocol is used between the master source node and the slave source nodes. The primary

difference in this case is that the feedback is from the master source node to the slave nodes and does not involve the destination.

Up to this point the synchronization process has been coordinated among the source nodes themselves without requiring any interaction with the destination node. In order for the sensors to beamform toward the destination, each source must estimate its channel response to the destination. This is achieved by having the destination broadcast a beacon (e.g., a sinusoidal signal at the carrier frequency) to the source nodes. Since the sources have already been synchronized, each source node can independently estimate its own complex channel gain to the destination using its frequency and phase-synchronized LO. The source nodes can then transmit as a distributed beamformer to the destination by applying the complex conjugate of these gains, typically at baseband, to their transmitted signals.

ROUND-TRIP OPEN-LOOP SYNCHRONIZATION

A different open-loop carrier synchronization system that eliminates the need for digital signaling during synchronization was proposed in [4, 5, 11]. The scheme is based on the equivalence of round-trip propagation delays through a multi-hop chain of source nodes and thus is called the *round-trip* carrier synchronization scheme. Like the open-loop master-slave synchronization system, the round-trip system requires minimal interaction between the source nodes and the destination.

A two-source round-trip system model is shown in Fig. 4 (round-trip carrier synchronization of more than two source nodes is discussed in [5]). The basic idea behind the round-trip synchronization system is that an unmodulated beacon “bounced” around the clockwise circuit shown in Fig. 4 will incur the same total phase shift as an unmodulated beacon “bounced” around the counterclockwise circuit shown in Fig. 4 when channels are reciprocal. The equivalence of the accumulated phase shifts for both round-trip circuits is the key feature of the round-trip carrier synchronization technique. The beacons are bounced around the circuit by having each source transmit periodic extensions of received beacons. Beamforming is achieved since the destination is essentially receiving the sum of two beacons, modulated by the common message, after they have propagated through circuits with identical phase shifts.


The actual implementation of a round-trip distributed beamformer is complicated, however, by the constraint that wireless transceivers may not transmit and receive on the same frequency at the same time. One approach is to use continuously transmitted beacons with distinct frequencies (also distinct from the carrier frequency). This approach, called the frequency-synthesis round-trip carrier synchronization system, was considered in [4] where each source employed a pair of frequency-synthesis PLLs in order to generate appropriate frequency-scaled periodic extensions of the beacons it received. An audio-frequency prototype of the frequency synthesis round-trip carrier synchronization system is discussed later. While the continuously transmitted beacons allowed for high rates of source and/or destination mobility, the use of distinct frequencies for the beacons and carriers resulted in non-reciprocal phase shifts and degraded performance in general multipath channels.

To ensure channel reciprocity in general multipath channels, a single-frequency time slotted round-trip carrier synchronization technique was proposed for a two-source distributed beamformer in [11] and extended to $N > 2$ sources in [5]. A total of $2N - 1$ synchronization time slots are needed to synchronize the sources prior to beamforming. The protocol is also repeated in order to avoid unacceptable phase drift, resulting from frequency estimation errors as well as phase noise and/or mobility, between the sources during beamforming. Long duration synchronization time slots tend to result in low estimation error but increased drift due to phase noise and/or mobility. Short duration time slots reduce the effects of phase noise and mobility, but lead to increased drift from low-quality frequency and phase estimates. Guidelines for achieving an efficient trade-off with low synchronization overhead are discussed in [5].

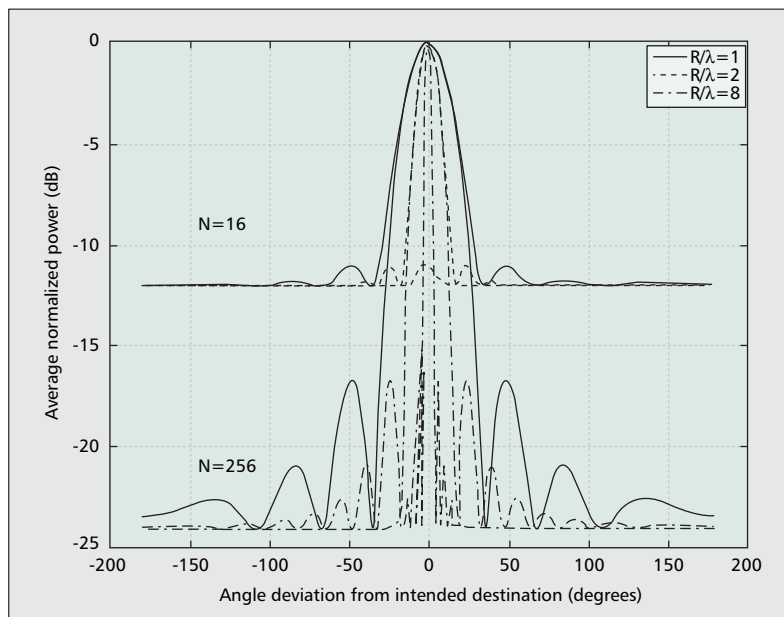
BEAMPATTERNS FOR RANDOMLY PLACED SOURCES

The carrier synchronization techniques described previously are necessary to ensure that the directional gain of the distributed beamformer is close to that of an ideal conventional beamformer. Given a particular antenna geometry and the sources' carrier phases, it is also possible to use standard techniques to compute the beamwidth and sidelobe characteristics of a distributed beamformer. These characteristics may be of interest in applications where, for example, security or interference is important. Since the "antenna geometry" of a distributed beamformer may be random, however, a statistical characterization of the beam-pattern is necessary. This section summarizes recent work in this area.

The probability distribution of the far-field beam pattern of a distributed beamformer with node locations uniformly distributed on a two-dimensional disk of radius R was analyzed in [12]. The average far-field beam pattern for a N -source distributed beamformer was shown to be



■ Figure 4. Round-trip open-loop carrier synchronization system, two source nodes.

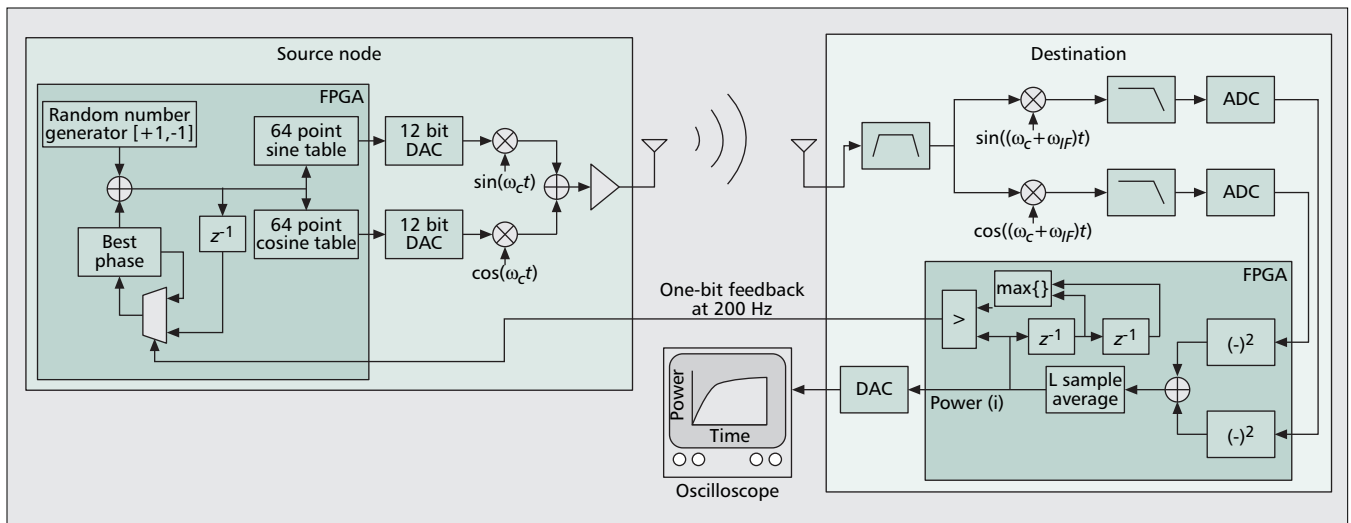


■ Figure 5. Average beam pattern of a distributed beamformer with randomly placed nodes.

$$P_{\text{av}}(\phi) = \frac{1}{N} + \left(1 - \frac{1}{N}\right) \left| \frac{J_1(4\pi\tilde{R}\sin(\phi/2))}{2\pi\tilde{R}\sin(\phi/2)} \right|^2,$$

where $\tilde{R} = R/\lambda$ is the radius of the disk normalized by the wavelength of the transmission, ϕ is the angle with respect to the intended destination, and $J_1(x)$ is the first-order Bessel function of the first kind. The average far-field beam pattern is shown in Fig. 5 for two different values of N and three different values of \tilde{R} .

For sufficiently large N and $\tilde{R} \gg 1$, the average 3 dB beamwidth of the main lobe was shown to be inversely proportional to \tilde{R} by numerically solving Eq. 2 for the case $P_{\text{av}}(f) = 1/2$. The dependence of the average 3 dB beamwidth on \tilde{R} is also evident in Fig. 5. These results suggest that very narrow beamwidths can be achieved in typical sensor network applications. For example, a sensor network with $N = 10$ randomly placed source nodes on a disk of radius 25 m



■ **Figure 6.** Block diagram of the one-bit feedback closed-loop carrier synchronization prototype described in [15].

will, on average, achieve a 3 dB beam width of less than half a degree if the sources transmit with 900 MHz carriers. The average sidelobe power of a distributed beamformer with N randomly placed source nodes was also shown to be on the order of $1/N$, plus some margin for sidelobe peaks near the main beam in [12]. The dependence of the sidelobe power on N is also evident in Fig. 5.

INFORMATION SHARING AMONG BEAMFORMING NODES

In conventional transmit beamforming, a common message is transmitted across all antennas in the array. The transmitted signal at each antenna element is simply a complex weighted version of the common message with weights selected to achieve a desired beam pattern. In distributed beamforming systems, nodes must share information prior to beamforming. When the links between cooperating nodes in a distributed beamformer are short with respect to the link to the destination, it is reasonable to assume that energy required for information sharing prior to beamforming is negligible with respect to the energy required to transmit to the intended destination. The time required for information sharing may not be negligible in some cases, however, and depends to some extent on the architecture of the network.

The problem of information sharing in distributed beamforming systems has primarily been studied for the case of heterogeneous networks with K master source nodes, each with distinct information to convey to a distinct destination [13]. These master source nodes share a pool of N “non-master” source nodes (or relays) that can transmit as a distributed beamformer. A straightforward approach in this scenario is to use time sharing: one master source node broadcasts its message to the relay pool. The relays then transmit this message, including any noise in the received signals, as a beamformer to the destination. This broadcast beamforming cycle is performed one master source

node at a time; hence, the throughput per master source node for this scheme is inversely proportional to the number of master source nodes.

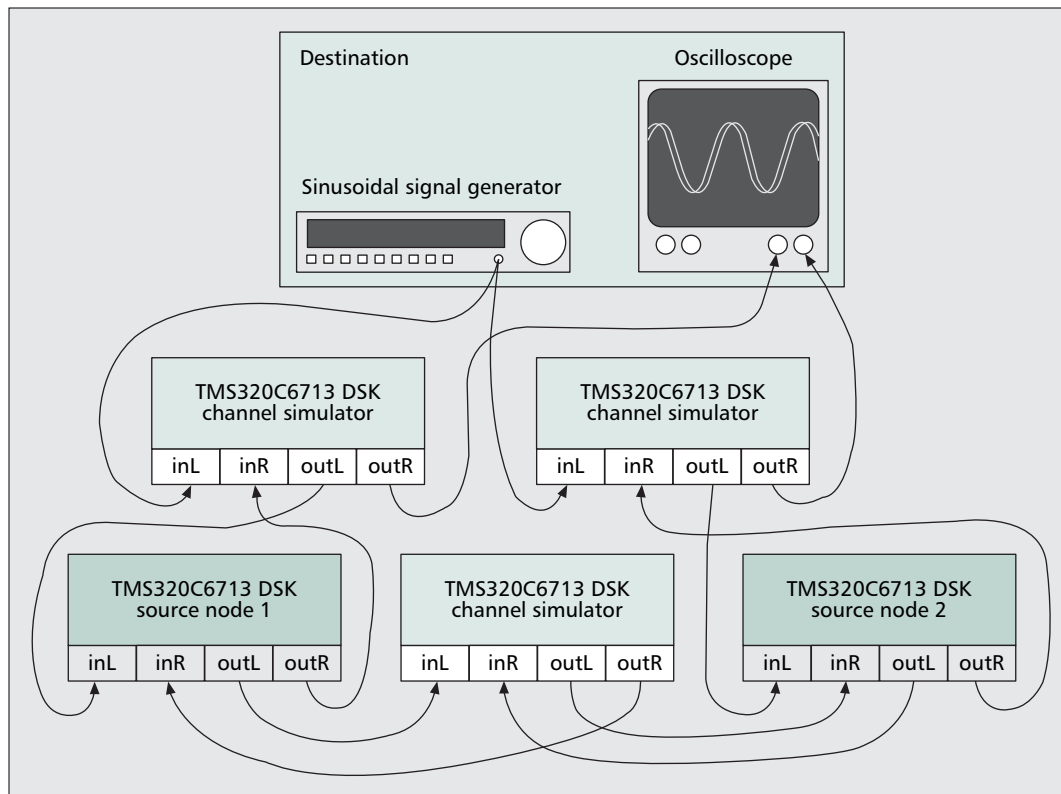
A higher-throughput space-division information sharing strategy is proposed in [14] in which all K master source nodes simultaneously broadcast their independent information to the pool of non-master source nodes. The N non-master source nodes then simultaneously beamform to the K destination nodes so that the message of the k th master source node combines coherently at the k th destination. While the throughput of this space-division approach is clearly better than that of time sharing, the simultaneous broadcast of messages by the master source nodes and simultaneous beamforming by the non-master source nodes may result in interference in the signals received at each destination.

PROOF-OF-CONCEPT PROTOTYPES

As theoretical research on distributed transmit beamforming has advanced, experimental prototypes have recently been constructed to confirm theoretical predictions and to better understand the inherent nonidealities in practical realizations. This section describes two such prototypes and summarizes the results of the laboratory experiments.

ONE-BIT FEEDBACK CLOSED-LOOP SYNCHRONIZATION PROTOTYPE

In 2006 a prototype of the one-bit feedback closed-loop synchronization system described earlier was built at the University of California at Berkeley in collaboration with the University of California at Santa Barbara [15]. A block diagram of a single source node and the destination node is shown in Fig. 6. Carrier frequency synchronization was achieved by distributing a common clock to the source nodes, which was multiplied up in frequency by each source node separately using a PLL. The one-bit feedback was conveyed from the destination to the source nodes via separate wired links. Using an FPGA-based power estimator, the destination fed back



■ **Figure 7.** Block diagram of the two-source round-trip open-loop frequency-synthesis carrier synchronization prototype described in [16].

The next step is to investigate and demonstrate distributed beamforming in a networked context, with a detailed design that spans information sharing, timing synchronization, carrier frequency synchronization, and carrier phase alignment.

a value of 1 to the source nodes when the current received power was greater than the averaged power estimates from each of the last L iterations ($L = 4$ for the results reported in [15]).

In a bench-top experiment performed with three source nodes, the measured received power was better than 90 percent of ideal. Convergence took approximately 60 iterations, which for a 200 Hz feedback rate corresponds to a convergence time of approximately 300 ms. The experiment was performed with unmodulated carriers as well as binary phase shift keying (BPSK) modulated carriers; as expected, data modulation did not affect convergence time or beamforming gain.

The one-bit algorithm has also been extended to provide distributed frequency and phase synchronization; this was demonstrated in 2007–2008 for a millimeter-wave sensor network testbed at the University of California at Santa Barbara [10].

TWO-SOURCE FREQUENCY-SYNTHESIS ROUND-TRIP SYNCHRONIZATION PROTOTYPE

In 2005–2006 a two-source distributed transmit beamformer using the round-trip open-loop carrier synchronization technique described in [4] was built and tested at Worcester Polytechnic Institute, Massachusetts [16]. The source nodes were realized by using Texas Instruments TMS320C6713 digital signal processing starter kits (DSKs) [17]. All synchronization functionality was realized in software running in real time on the 225 MHz floating point digital signal pro-

cessor. By using audio carrier frequencies and exploiting the built-in AIC23 stereo codec, a real-time proof of concept was built without custom hardware development.

A block diagram of the round-trip open-loop carrier synchronization prototype is shown in Fig. 7. A total of five TMS320C6713 DSKs were used to realize the system. Three of the DSKs were programmed to work as single-path channel simulators to facilitate repeatable simulation of time-invariant or time-varying channels. The remaining two DSKs were programmed to work as source nodes. Each source node simultaneously ran two PLLs, one for each analog channel.

Several distributed transmit beamforming experiments with unmodulated carriers and time-invariant and time-varying channels are reported in [16]. For time-invariant channels, convergence typically occurred in less than 5000 carrier cycles (at a frequency of 5.4 kHz), with a received power almost 99 percent that of an ideal beamformer. For single-path time-varying channels, sources moving at constant velocity suffer no performance loss, as long as the PLL filters are of at least second order.

DISCUSSION AND CONCLUSIONS

The results reviewed in this article indicate that distributed transmit beamforming is on the cusp of feasibility. The prototypes reported in the literature thus far have focused on demonstrating that the critical task of aligning carrier phases at the intended destination is feasible. The next step is to investigate and demonstrate distribut-

Protocols must be designed both for local coordination among the sources and for communication between the sources and the destination. The gains from distributed beamforming must, of course, be traded off against the overhead required to implement it.

ed beamforming in a networked context, with a detailed design that spans information sharing, timing synchronization, carrier frequency synchronization, and carrier phase alignment. Protocols must be designed for both local coordination among the sources and communication between the sources and the destination. The gains from distributed beamforming must, of course, be traded off against the overhead required to implement it.

The N^2 -fold power gain provided by distributed transmit beamforming with N collaborating sources can be exploited in different ways, depending on the needs of the application. If each source is constrained in transmit power, collaboration can be used to increase the range beyond what is attainable by a single source, which can be exploited for extending network access in rural settings, for example. If the link budget is sufficient for a single source to communicate with the destination, collaboration can be used to significantly increase the rate of communication, assuming that the system operates in a power-limited rather than bandwidth-limited regime. This could dramatically increase the upload rate from a network of sensors or soldier radios. On the other hand, if a single source can already communicate with the intended destination at the desired rate and range, distributed beamforming can be employed to reduce the transmit power per source by a factor of N^2 and reduce the energy radiated in undesired directions, which can be exploited for energy efficiency in sensor networks or low-probability-of-intercept communication in military applications. While each application may require a different cross-layer protocol and physical layer design, we hope that this article has conveyed the fundamental issues that must be addressed by such a design.

REFERENCES

- [1] G. Barriac, R. Mudumbai, and U. Madhow, "Distributed Beamforming for Information Transfer in Sensor Networks," *Proc. Info. Process. Sensor Net., 3rd Int'l. Wksp.*, Berkeley, CA, Apr. 26–27, 2004.
- [2] R. Mudumbai, G. Barriac, and U. Madhow, "On the Feasibility of Distributed Beamforming in Wireless Networks," *IEEE Trans. Wireless Commun.*, vol. 6, May 2007, pp. 1754–63.
- [3] Y. Tu and G. Pottie, "Coherent Cooperative Transmission from Multiple Adjacent Antennas to a Distant Stationary Antenna Through AWGN Channels," *Proc. IEEE VTC*, vol. 1, Birmingham, AL, Spring 2002, pp. 130–34.
- [4] D. R. Brown III, G. Prince, and J. McNeill, "A Method for Carrier Frequency and Phase Synchronization of Two Autonomous Cooperative Transmitters," *Proc. 5th IEEE Signal Proc. Advances Wireless Commun.*, New York, NY, June 5–8, 2005, pp. 278–82.
- [5] D. R. Brown III and H. V. Poor, "Time-Slotted Round-Trip Carrier Synchronization for Distributed Beamforming," *IEEE Trans. Signal Proc.*, vol. 56, Nov. 2008, pp. 5630–43.
- [6] L. C. Godara, "Application of Antenna Arrays to Mobile Communications. II. Beam-Forming and Direction-of-Arrival Considerations," *Proc. IEEE*, vol. 85, Aug. 1997, pp. 1195–1245.

- [7] J. Elson, L. Girod, and D. Estrin, "Fine-Grained Network Time Synchronization Using Reference Broadcasts," *SIGOPS Op. Sys. Rev.*, vol. 36, no. SI, 2002, pp. 147–63.
- [8] R. Mudumbai et al., "Scalable Feedback Control for Distributed Beamforming in Sensor Networks," *Proc. IEEE Int'l. Symp. Info. Theory*, Adelaide, Australia, Sept. 2005, pp. 137–41.
- [9] R. Mudumbai et al., "Distributed Transmit Beamforming Using Feedback Control," *IEEE Trans. Info. Theory*, in review.
- [10] M. Seo, M. Rodwell, and U. Madhow, "A Feedback-Based Distributed Phased Array Technique and its Application to 60-ghz Wireless Sensor Network," *IEEE MTT-S Int'l. Microwave Symp. Digest*, Atlanta, GA, June 15–20, 2008, pp. 683–86.
- [11] I. Ozil and D.R. Brown III, "Time-Slotted Round-Trip Carrier Synchronization," *Proc. 41st Asilomar Conf Signals, Sys., Comp.*, Pacific Grove, CA, Nov. 4–7, 2007, pp. 1781–85.
- [12] H. Ochiai et al., "Collaborative Beamforming for Distributed Wireless Ad Hoc Sensor Networks," *IEEE Trans. Signal Proc.*, vol. 53, Nov. 2005, pp. 4110–24.
- [13] A. Dana and B. Hassibi, "On the Power Efficiency of Sensory and Ad Hoc Wireless Networks," *IEEE Trans. Info. Theory*, vol. 52, July 2006, pp. 2890–2914.
- [14] L. Dong, P. Petropulu, and H.V. Poor, "A Cross-Layer Approach to Collaborative Beamforming for Wireless Ad Hoc Networks," *IEEE Trans. Signal Proc.*, vol. 56, July 2008, pp. 2981–93.
- [15] R. Mudumbai et al., "Distributed beamforming using 1 bit feedback: from concept to realization," *Proc. 44th Allerton Conf. Commun., Control, Comp.*, Monticello, IL, Sept. 2006, pp. 1020–27.
- [16] J. McGinley, "Real-Time Software-Defined-Radio Implementation of a Two-Source Distributed Beamformer," Master's thesis, WPI, Worcester, MA, Dec. 2006.
- [17] R. Chassaing and D. Reay, *Digital Signal Processing and Applications with the TMS320C6713 and TMS320C6416 DSK*, 2nd ed., Wiley, 2008.

BIOGRAPHIES

RAGHURAMAN MUDUMBAI is currently a post-doctoral researcher at the University of California, Santa Barbara (UCSB). He received a B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, in 1998 and an M.S. degree in electrical engineering from Polytechnic University, Brooklyn, New York, in 2000. He worked for two years at Ericsson Telephone Co. and returned to graduate school at UCSB in 2003, where he received a Ph.D. in electrical and computer engineering in 2007.

D. RICHARD BROWN III is an associate professor of electrical and computer engineering at Worcester Polytechnic Institute (WPI), Massachusetts. He is a recipient of the NSF CAREER award, and his research interests are currently in the development of energy-efficient cooperative communication systems, synchronization, distributed computing, and game-theoretic analysis of communication networks.

UPAMANYU MADHOW [F] is a professor of electrical and computer engineering at UCSB. His current research interests include wireless communication and sensor networks. He is a recipient of the NSF CAREER award. He has served as Associate Editor for several IEEE journals, and is the author of the textbook *Fundamentals of Digital Communication* (Cambridge University Press, 2008).

H. VINCENT POOR [F] is the Dean of Engineering and Applied Science at Princeton University, where he is also the Michael Henry Strater University Professor of Electrical Engineering. His interests lie in the area of statistical signal processing with applications in wireless networks and related fields. Among his publications in this area is the recent book *MIMO Wireless Communications* (Cambridge University Press, 2007). He is a member of the National Academy of Engineering.