

An Experimental Study of Acoustic Distributed Beamforming Using Round-Trip Carrier Synchronization

D. Richard Brown III, Boyang Zhang, Boris Svirchuk, and Min Ni

Abstract—This paper describes the development of an acoustic distributed beamforming system and presents experimental results for two-source and three-source acoustic distributed beamforming using the time-slotted round-trip carrier synchronization protocol. Each source node in the system was built using commercial off-the-shelf parts including a Texas Instruments floating-point digital signal processor, microphone, speaker, audio amplifier, and battery. The source node functionality, including phase locked loops and the logic associated with the time-slotted round-trip carrier synchronization protocol, was realized through real-time software independently running on each source node’s C6713 digital signal processor. Experimental results for two-source and three-source realizations of the acoustic distributed beamforming system in a room with multipath channels are presented. The two-source and three-source experimental results show mean power gains of approximately 97.7% and 90.7%, respectively, of an ideal beamformer.

I. INTRODUCTION

Distributed transmit beamforming has recently been proposed as a technique in which multiple individual single-antenna transmitters simultaneously transmit a common message and control the phase and frequency of their carriers so that their bandpass signals constructively combine at an intended destination. The transmitters in a distributed transmit beamformer form a *virtual antenna array* and, in principle, can achieve almost all of the gains of a conventional antenna array, i.e. increased range, rate, and/or energy efficiency, without the size, cost, and complexity of a conventional antenna array. Distributed transmit beamforming can also provide benefits in terms of security and interference reduction since less transmit power is scattered in unintended directions. As discussed in [1], a key challenge in realizing these benefits, however, is *precise carrier synchronization* such that the transmissions combine constructively at the intended destination.

In the last few years, several carrier phase and frequency synchronization techniques suitable for distributed transmit

beamforming have been proposed. A survey of these techniques is presented in [1]. This paper focuses on the time-slotted round-trip carrier synchronization technique first described in [2]. Time-slotted round-trip carrier synchronization is based on the equivalence of round-trip propagation delays through a multihop chain of source (single-antenna transmitter) nodes. A two-source round-trip system model is shown in Figure 1. The basic idea is that an unmodulated carrier transmitted by the destination node and “bounced” around the green (clockwise) circuit shown in Figure 1 will incur the same total phase shift as an unmodulated carrier transmitted by the destination node “bounced” around the blue (counterclockwise) circuit shown in Figure 1. In practice, the “bouncing” of carriers can be performed actively by having each source node track the signals received by other nodes using phase locked loops (PLLs) and then using the PLL in “holdover mode” to transmit a periodic extension of the signal received in a previous timeslot. Coherent combining is achieved since the destination is receiving the sum of two carriers, modulated by the common message, after they have propagated through circuits with identical phase shifts.




Fig. 1. Time-slotted round-trip carrier synchronization system with two source nodes.

D.R. Brown III is an Associate Professor with the Electrical and Computer Engineering Department, Worcester Polytechnic Institute, Worcester, MA 01609 USA. e-mail: drb@ece.wpi.edu.

B. Zhang is a System Integration Test Engineer at Azimuth Systems, Inc. in Acton, MA. e-mail: boyang.zhang@azimuth.net.

B. Svirchuk is a Research and Development Engineer and Philips Healthcare in Andover, MA. e-mail: boris.svirchuk@philips.com.

Min Ni is a Ph.D. candidate with the Electrical and Computer Engineering Department, Worcester Polytechnic Institute, Worcester, MA 01609 USA. e-mail: minni@ece.wpi.edu.

This work was supported by Texas Instruments and NSF award CCF-0447743.

While recent research has focused on the development and analysis of carrier phase and frequency synchronization techniques, relatively little has been published on prototypes and/or experimental studies of distributed transmit beamforming. In 2006, a prototype of a one-bit feedback carrier phase synchronization system described in [3] was built at the University of California, Berkeley, in collaboration with the University of California, Santa Barbara. In a bench-top experiment performed with three source nodes, the received

power at the destination node was measured to be better than 90% of ideal. The one-bit feedback system was also extended to include frequency synchronization as reported in [4].

This paper reports on the development of an acoustic proof-of-concept prototype for time-slotted round-trip carrier synchronization. The development of an acoustic system is motivated by the observation that one can easily replicate common electromagnetic radio-frequency (RF) carrier wavelengths acoustically by scaling all frequencies in the RF system by the ratio $340.3/(3 \cdot 10^8)$. This implies that results obtained through acoustic proof-of-concept prototypes can provide guidelines for the design and development of RF systems. The use of acoustic communications is also appealing due to the fact that acoustic transducers are simple and inexpensive and the inherently low data rates allow for real-time operation with low-cost hardware. The experimental study described in this paper was conducted using low-cost acoustic transducers with unmodulated 1021Hz carriers. Acoustic propagation at this frequency has the same wavelength as electromagnetic propagation at 900MHz. This paper provides a summary of our two-source and three-source experimental results showing mean power gains of approximately 97.7% and 90.7%, respectively, of an ideal beamformer.

II. TIMESLOTTED ROUND-TRIP PROTOCOL

This section describes the time-slotted round-trip carrier synchronization protocol in the context of the M -source system shown in Figure 2. For clarity of exposition, we begin with a detailed description of the protocol for the case with $M = 2$ sources and then describe the protocol for the case with more than two sources in Section II-C.

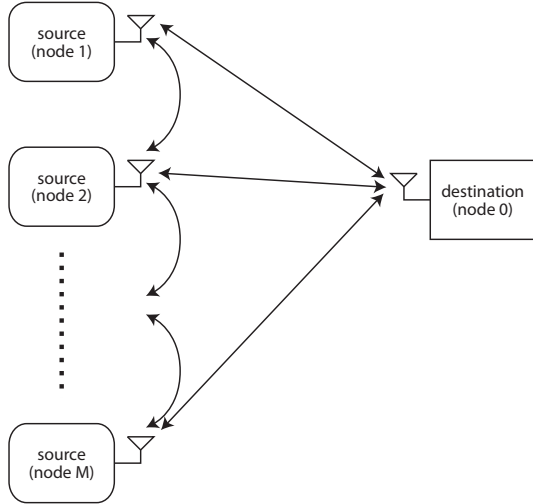


Fig. 2. M -source round-trip distributed beamforming system.

A. Two-Source Synchronization in Single-Path Channels

In the case with two sources, the time-slotted round-trip carrier synchronization protocol has a total of four timeslots: the first three timeslots are used for synchronization and the

final timeslot is used for beamforming. The activity in each timeslot is summarized below:

- TS⁽⁰⁾: The destination transmits the sinusoidal primary beacon to both sources. Both sources generate phase and frequency estimates from their local observations.
- TS⁽¹⁾: S_1 transmits a sinusoidal secondary beacon to S_2 . This secondary beacon is transmitted as a periodic extension of the beacon received in TS⁽⁰⁾. S_2 generates local phase and frequency estimates from this observation.
- TS⁽²⁾: S_2 transmits a sinusoidal secondary beacon to S_1 . This secondary beacon is transmitted as a periodic extension of the beacon received in TS⁽⁰⁾, with initial phase extrapolated from the phase and frequency estimates obtained by S_2 in TS⁽⁰⁾. S_1 generates local phase and frequency estimates from this observation.
- TS⁽³⁾: Both sources transmit simultaneously to the destination as a distributed beamformer. The carrier frequency of each source is based on both local frequency estimates obtained in the prior timeslots. The initial phase of the carrier at each source is extrapolated from the local phase and frequency estimates from the secondary beacon observation.

Figure 3 summarizes the time-slotted round-trip carrier synchronization protocol and shows how the protocol is repeated in order to avoid unacceptable phase drift between the sources during beamforming.




Fig. 3. Summary of the two-source time-slotted round-trip carrier synchronization protocol where PB and SB denote primary and secondary beacon synchronization timeslots, respectively.

Assuming temporarily that all of the channels in the system are single-path, it can easily be seen that the aggregate propagation times of the $D \rightarrow S_1 \rightarrow S_2 \rightarrow D$ and the $D \rightarrow S_2 \rightarrow S_1 \rightarrow D$ circuits are identical. As each source transmits periodic extensions of beacons it received in prior timeslots, each source is essentially “bouncing” the signal around the respective circuits. Beamforming is achieved since the destination is receiving the sum of two primary beacons after they have propagated through circuits with identical propagation times.

The time-slotted protocol begins in TS⁽⁰⁾ with the transmission of a unit-amplitude sinusoidal primary beacon of duration T_0 from the destination to both sources,

$$x_0(t) = \cos(\omega(t - t_0) + \phi_0) \quad t \in [t_0, t_0 + T_0]. \quad (1)$$

The signal received at S_i in TS⁽⁰⁾ can be written as

$$y_{0i}(t) = \alpha_{0i} \cos(\omega(t - (t_0 + \tau_{0i})) + \phi_0) + \eta_{0i}(t)$$

for $t \in [t_0 + \tau_{0i}, t_0 + \tau_{0i} + T_0]$ where $\eta_{0i}(t)$ denotes the noise in the $0 \rightarrow i$ channel and $i \in \{1, 2\}$. Each source tracks

the primary beacon from the destination using its first phase locked loop (PLL). Prior to the conclusion of the primary beacon, each source stops tracking and enters “holdover” mode on its PLL. If the PLLs are designed correctly, the transient response of the PLL will complete prior to entering holdover. This results in local frequency and phase estimates at each source, denoted by $\hat{\omega}_{0i}$ and $\hat{\phi}_{0i}$, respectively, at S_i for $i \in \{1, 2\}$. We use the usual convention that the phase estimate $\hat{\phi}_{0i}$ is an estimate of the phase of the received signal at the start of the observation at S_i , i.e. $\hat{\phi}_{0i}$ is an estimate of the phase of $y_{0i}(t)$ at time $t_0 + \tau_{0i}$.

Timeslot TS⁽¹⁾ begins immediately upon the conclusion of the primary beacon $y_{01}(t)$ at S_1 . At time $t_1 = t_0 + \tau_{01} + T_0$, S_1 uses its first PLL (running in holdover mode) to transmit a sinusoidal secondary beacon to S_2 that is a periodic extension of $y_{01}(t)$ (possibly with different amplitude). The secondary beacon transmitted by S_1 in TS⁽¹⁾ can then be written as

$$x_{12}(t) = a_{12} \cos\left(\hat{\omega}_{01}(t - t_1) + \hat{\phi}_1\right) \quad t \in [t_1, t_1 + T_1]$$

where

$$\hat{\phi}_1 = \hat{\phi}_{01} + \hat{\omega}_{01}(t_1 - (t_0 + \tau_{01})) = \hat{\phi}_{01} + \hat{\omega}_{01}T_0.$$

is the extrapolated phase of the first PLL at S_1 at time t_1 .

After propagation through the $1 \rightarrow 2$ channel, this secondary beacon is received by S_2 as

$$y_{12}(t) = \alpha_{12}a_{12} \cos\left(\hat{\omega}_{01}(t - (t_1 + \tau_{12})) + \hat{\phi}_1\right) + \eta_{12}(t)$$

for $t \in [t_1 + \tau_{12}, t_1 + \tau_{12} + T_1]$ where $\eta_{12}(t)$ denotes the noise in the $1 \rightarrow 2$ channel. S_2 uses its second PLL to track this beacon and enters holdover on the second PLL prior to the conclusion of this beacon. The frequency and phase estimates of the second PLL at S_2 are denoted by $\hat{\omega}_{12}$ and $\hat{\phi}_{12}$, respectively.

Timeslot TS⁽²⁾ begins immediately upon the conclusion of $y_{12}(t)$ at S_2 . At time $t_2 = t_1 + \tau_{12} + T_1$, S_2 uses its first PLL (running in holdover mode since the end to TS⁽⁰⁾) to transmit a sinusoidal secondary beacon to S_1 that is a periodic extension of $y_{02}(t)$. The secondary beacon transmitted by S_2 in TS⁽²⁾ can then be written as

$$x_{21}(t) = a_{21} \cos\left(\hat{\omega}_{02}(t - t_2) + \hat{\phi}_2\right) \quad t \in [t_2, t_2 + T_2]$$

where

$$\hat{\phi}_2 = \hat{\phi}_{02} + \hat{\omega}_{02}(t_2 - (t_0 + \tau_{02})).$$

is the extrapolated phase of the first PLL at S_2 at time t_2 .

After propagation through the $2 \rightarrow 1$ channel, this secondary beacon is received by S_1 as

$$y_{21}(t) = \alpha_{12}a_{21} \cos\left(\hat{\omega}_{02}(t - (t_2 + \tau_{21})) + \hat{\phi}_2\right) + \eta_{21}(t)$$

for $t \in [t_2 + \tau_{21}, t_2 + \tau_{21} + T_2]$ where $\eta_{21}(t)$ denotes the noise in the $2 \rightarrow 1$ channel and where we have used the fact that $\tau_{21} = \tau_{12}$ and $\alpha_{21} = \alpha_{12}$. S_1 uses its second PLL to track this beacon and enters holdover on the second PLL prior to the conclusion of this beacon. The frequency and phase

estimates of the second PLL at S_1 are denoted by $\hat{\omega}_{21}$ and $\hat{\phi}_{21}$, respectively.

In timeslot TS⁽³⁾, S_1 and S_2 each transmit to the destination as a distributed beamformer with carries generated using the second PLL (running in holdover mode). The carrier at each source is transmitted as a periodic extension of the *secondary* beacons received at each source. Since the performance of the distributed beamformer is primarily affected by the phase offset between the carriers at the destination, we can write the transmissions of S_1 and S_2 as unmodulated carriers. The unmodulated carrier transmitted by S_i during TS⁽³⁾ can be written as

$$x_{i0}(t) = a_{i0} \cos\left(\hat{\omega}_{ij}(t - t_{3i}) + \hat{\phi}_{3i}\right) \quad t \in [t_{3i}, t_{3i} + T_3] \quad (2)$$

for $j \neq i$ and where

$$\hat{\phi}_{31} = \hat{\phi}_{21} + \hat{\omega}_{21}(t_{31} - (t_2 + \tau_{12})) \quad \text{and} \quad (3)$$

$$\hat{\phi}_{32} = \hat{\phi}_{12} + \hat{\omega}_{12}(t_{32} - (t_1 + \tau_{12})), \quad (4)$$

are the extrapolated phases of the second PLLs at S_1 and S_2 , respectively, at time t_3 .

The signal received at D in TS⁽³⁾ can be written as

$$y_0(t) = \sum_{i=1}^2 \alpha_{0i}a_{i0} \cos\left(\hat{\omega}_{ij}(t - t_3) + \hat{\phi}_{3i}\right) + \eta_0(t)$$

for $t \in [t_3, t_3 + T_3]$ where $t_3 = t_{31} + \tau_{01} = t_{32} + \tau_{02}$ and where we have again used the fact that $\tau_{0i} = \tau_{i0}$ and $\alpha_{0i} = \alpha_{i0}$ for $i \in \{1, 2\}$.

B. Two-Source Synchronization in Multipath Channels

The assumption of single-path channels in the prior development of the two-source round-trip carrier synchronization protocol was used for clarity of exposition but is not necessary to enable beamforming. Note that the beacons exchanged in the round-trip carrier synchronization system are all at the same frequency as the carrier. Hence, each bidirectional channel between a pair of source nodes (and between individual source nodes and the destination) is a time-division-duplex (TDD) channel that is reciprocal in both directions. The principles developed in the case of single-path channels can then be applied to the case with multipath channels with the difference being that it is now the phase shift, rather than the propagation delay, of each channel that is identical in both directions. Denoting the phase of the channel between node i and j as $\theta_{i,j}$, it is easy to see that the aggregate round trip phase shifts of the $D \rightarrow S_1 \rightarrow S_2 \rightarrow D$ circuit and the $D \rightarrow S_2 \rightarrow S_1 \rightarrow D$ circuit are identical and equal to

$$\theta^{\text{rt}} = \theta_{0,1} + \theta_{1,2} + \theta_{2,0}.$$

Although the steady-state phase shift of each channel is identical in both directions, the multipath channels also cause the finite-duration beacons received by S_1 and S_2 to have transient components that must be accounted for in the protocol. In a system with multipath channels, each source node should delay tracking the beacon with the appropriate PLL until the transient effects of the channel become negligible. The source

then tracks the beacon with the appropriate PLL during the steady-state portion of the beacon observation and puts the PLL into holdover mode prior to the conclusion of the steady-state portion of the beacon. This is summarized in Figure 4.

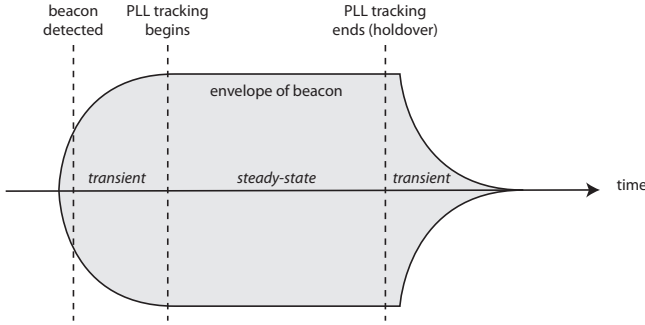


Fig. 4. Effect of a beacon received in multipath on PLL tracking and holdover.

The important thing here is that each source uses *only the steady-state portion* of its noisy observation in each timeslot for PLL tracking and subsequent computation of local estimates of the received frequency and phase. The initial and final transient portions of the observation are ignored. As with single-path channels, the phase estimates at each source are extrapolated for transmission of the secondary beacons and carriers as periodic extensions of the steady state portion of the primary beacon observations.

In order to ensure that some portion of the observation is steady-state observation, the duration of each beacon must exceed the delay spread of the channel in which the beacon is transmitted. Guard times may also be added between timeslots to allow for the transients in a previous timeslot to vanish before a new beacon is transmitted. No other modifications to the synchronization protocol are necessary. In the final timeslot, both sources transmit as in (2).

C. General M -Source Synchronization

In a distributed beamforming system with $M > 2$ sources, the time-slotted round-trip carrier synchronization protocol has a total of $2M$ timeslots denoted as $\text{TS}^{(0)}, \dots, \text{TS}^{(2M-1)}$. The activity in each timeslot is summarized below:

- 1) In $\text{TS}^{(0)}$ the destination transmits the sinusoidal primary beacon to all M sources. Each source generates local phase and frequency estimates from its observation.
- 2) In $\text{TS}^{(i)}$ for $i = 1, \dots, M-1$, S_i transmits a sinusoidal secondary beacon to S_{i+1} . The secondary beacon transmitted by S_i in $\text{TS}^{(i)}$ is a periodic extension of the beacon received in $\text{TS}^{(i-1)}$. S_{i+1} generates local phase and frequency estimates from this observation.
- 3) In $\text{TS}^{(M)}$, S_M transmits a sinusoidal secondary beacon to S_{M-1} . This secondary beacon is transmitted as a periodic extension of the primary beacon received by S_M in $\text{TS}^{(0)}$, with initial phase extrapolated from the phase and frequency estimates obtained by S_M in $\text{TS}^{(0)}$. S_{M-1} generates local phase and frequency estimates from this observation.

- 4) In $\text{TS}^{(i)}$ for $i = M+1, \dots, 2M-2$, S_{2M-i} transmits a sinusoidal secondary beacon to S_{2M-i-1} . The secondary beacon transmitted by S_{2M-i} in $\text{TS}^{(i)}$ is a periodic extension of the secondary beacon received in $\text{TS}^{(i-1)}$. S_{2M-i-1} generates local phase and frequency estimates from this observation.
- 5) In $\text{TS}^{(2M-1)}$, all M sources transmit simultaneously to the destination as a distributed beamformer. The frequency and initial phase of the carrier transmitted by each source is based only on the local phase and frequency estimates obtained in the prior timeslots.

Since, like the two-source case, the total phase shift of the $D \rightarrow S_1 \rightarrow S_2 \rightarrow \dots \rightarrow S_M \rightarrow D$ and the $D \rightarrow S_M \rightarrow S_{M-1} \dots \rightarrow S_1 \rightarrow D$ circuits are identical, distributed beamforming between source nodes S_1 and S_M can be achieved by following the round-trip protocol and transmitting secondary beacons as periodic extensions of previously received beacons in exactly the same manner as described in Section II-A. When $M > 2$, however, nodes S_2, \dots, S_{M-1} must also derive appropriate transmission phases to participate in the distributed beamformer.

Ignoring estimation errors to ease exposition, the round-trip nature of the protocol and the transmission of periodic extensions implies that the destination will receive carriers from S_1 and S_M at a phase (relative to the phase of the primary beacon transmitted in $\text{TS}^{(0)}$) of

$$\theta^{\text{rt}} = \theta_{0,1} + \theta_{1,2} + \dots + \theta_{M-1,M} + \theta_{M,0}$$

where $\theta_{k,i} = \theta_{i,k}$ denotes the phase of the channel between node i and node k . Let \mathcal{S} denote the set of source nodes S_m for $m \in \{2, \dots, M-1\}$. In order for source node $S_m \in \mathcal{S}$ to transmit a carrier that arrives at the destination with the same phase as S_1 and S_M , S_m must transmit its carrier with phase $\theta^{\text{rt}} - \theta_{m,0}$.

Source node $S_m \in \mathcal{S}$ receives three transmissions during the synchronization phase of the protocol: a primary beacon in $\text{TS}^{(0)}$ at phase $\theta_{0,m} = \theta_{m,0}$, a secondary beacon during the counterclockwise¹ propagation of beacons in $\text{TS}^{(m-1)}$ at phase $\theta_m^{\downarrow} = \theta_{0,1} + \theta_{1,2} + \dots + \theta_{m-1,m}$ and another secondary beacon during the clockwise propagation of beacons in $\text{TS}^{(2M-m-1)}$ at phase $\theta_m^{\uparrow} = \theta_{0,M} + \theta_{M,M-1} + \dots + \theta_{m+1,m}$. Since each node in the system estimates the phase of received beacons relative to its own local time reference, absolute estimates of θ_m^{\downarrow} and θ_m^{\uparrow} at S_m will both have an unknown phase offset that depends on the phase of the local time reference at S_m . To avoid the problem of determining this unknown phase offset, S_m can calculate the phase *difference* between any two phases that were measured under the same local time reference and effectively cancel the offsets. Accordingly, S_m can calculate the phase difference between each secondary beacon phase estimate and the primary beacon phase estimate

¹In the context of Figure 2, $S_1 \rightarrow S_2 \rightarrow \dots \rightarrow S_M$ is counterclockwise propagation and $S_M \rightarrow S_{M-1} \dots \rightarrow S_1$ is clockwise propagation around the circuit including D.

as

$$\begin{aligned}\delta_m^\uparrow &= \theta_m^\uparrow - \theta_{0,m} \\ \delta_m^\downarrow &= \theta_m^\downarrow - \theta_{0,m}.\end{aligned}$$

Since the unknown local phase offset has been canceled in the phase differences δ_m^\uparrow and δ_m^\downarrow , the sum of these terms will also not have any unknown phase offset. Hence, if S_m transmits its carrier as a periodic extension of the *primary* beacon received in $TS^{(0)}$ with an additional phase shift of $\delta_m^\uparrow + \delta_m^\downarrow$, the carrier phase of S_m can be written as

$$\phi_m = \theta_{0,m} + \delta_m^\uparrow + \delta_m^\downarrow = \theta^{rt} - \theta_{0,m}$$

which is the desired phase for beamforming.

III. INDEPENDENT LOCAL OSCILLATORS IN DISTRIBUTED BEAMFORMING

Unlike a conventional transmit beamformer in which each antenna element is driven by the same local oscillator, a distributed transmit beamformer is realized by cooperative transmission of multiple single-antenna nodes, each with their own independent local oscillator. Distributed transmit beamforming requires precise carrier synchronization in order to appropriately align the frequency and phase of each node's transmission so that the bandpass signals coherently combine at the intended destination. Estimation errors incurred during synchronization as well as independent phase noise in each local oscillator all lead to some loss of performance with respect to an ideal conventional beamformer. At time t , the power of the aggregate received signal at the destination from the source nodes can be expressed as

$$|y_0(t)|^2 = \sum_m a_{0,m}^2 + \sum_m \sum_{n \neq m} a_{0,m} a_{0,n} \cos(\delta_{m,n}(t)) \quad (5)$$

where $a_{0,m}$ is the amplitude of the channel between source-node m and the destination and the non-ideal nature of the distributed beamformer is captured in the carrier offset terms

$$\delta_{m,n}(t) := (\hat{\omega}_m - \hat{\omega}_n)t + (\hat{\phi}_m - \hat{\phi}_n) + \chi_m(t) - \chi_n(t) \quad (6)$$

between source nodes m and n where $\hat{\omega}_m$, $\hat{\phi}_m$, and $\chi_m(t)$ represent the estimated carrier frequency, carrier phase, and local oscillator phase noise for source node m . Note that (6) is composed of three components: carrier frequency offset, initial carrier phase offset at $t = 0$, and phase noise. The effect of each of these components is illustrated in Figure 5.

The statistical properties of each of the carrier offset components were analyzed in [2] for the two-source case in additive white Gaussian noise channels. The theoretical predictions were based on a Cramer-Rao lower bound analysis under the assumption that the source nodes used maximum likelihood estimators. The theoretical results showed that, even with low-cost oscillators, it was possible to achieve near-ideal beamforming performance with little synchronization overhead.

As shown in (5), any carrier offset term $\delta_{m,n}(t)$ not equal to zero results in some loss of received power at the destination node with respect to the ideal beamforming power when $\delta_{m,n}(t) \equiv 0$ for all m and n .




Fig. 5. An illustration of the components of the carrier offset terms $\delta_{m,n}(t)$ as a function of time.

IV. EXPERIMENTAL METHODOLOGY

The acoustic source nodes used in the experimental study were developed as part of an Acoustic Cooperative Communication Experimental Network Testbed (ACCENT). All of the hardware components of the ACCENT node are low-cost off-the-shelf parts. Figure 6 shows a block diagram of the major components of the source node including a Texas Instruments TMS320C6713DSK floating point DSP starter kit, microphone, power amplifier, speaker, and battery. As shown in Figure 7, the components are mounted in an plastic enclosure with the microphone and speaker placed in close proximity to approximate a single transducer. Note that each source node operates independently using its own local oscillator; there are no wires or signals shared among the source nodes other than the acoustic signals generated during the round-trip protocol.




Fig. 6. Block diagram of an ACCENT acoustic source node.



Fig. 7. ACCENT acoustic source node hardware in a plastic enclosure.

Photographs of the test environment configured for an acoustic experiment with two and three source nodes are shown in Figures 8 and 9, respectively. The room in which the acoustic experiments were performed was a typical carpeted conference room with dimensions approximately 7.5 meters

by 7.5 meters. In the two-source experiments, the nodes were placed in an approximately equilateral triangle configuration with approximately 4 meters of separation between the source nodes and between each source node and the destination. In the three-source test, the approximate node separations are given in Table I.




Fig. 8. Two-source acoustic distributed beamforming test configuration.

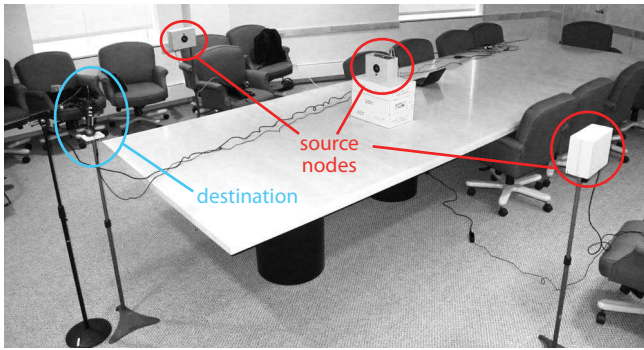


Fig. 9. Three-source acoustic distributed beamforming test configuration.

TABLE I

THREE-SOURCE TEST APPROXIMATE NODE SEPARATIONS IN METERS.

	dest	S ₁	S ₂	S ₃
dest	0	2.4	2.1	2.4
S ₁	2.4	0	2.1	3.8
S ₂	2.1	2.1	0	1.9
S ₃	2.4	3.8	1.9	0

The “destination node” was realized by using a portable CD player and a self-amplified loudspeaker for primary beacon generation, as well as a microphone and a Marantz digital recorder for recording of the signals. An oscilloscope was also connected to the output of the Marantz digital recorder for real-time monitoring.

Each acoustic experiment consisted of $N = 100$ “tests” where a test is a complete execution of the $2M - 1$ timeslots of the round-trip protocol. Upon initialization, each node enters into a state where it listens for a primary beacon from the destination node. When the start of the primary beacon is detected, the nodes execute the round-trip protocol according to the schedule in Table II where each node keeps time

by counting samples received from the codec onboard the TMS320C6713DSK sampling at a rate of 44.1 kHz. Note that Table II corresponds to the timing for a two-source test; the three-source tests have similar timing but require more timeslots to exchange the beacons as discussed in Section II-C. In all of the tests reported in this paper, the duration of each beacon was one second with a 0.25 second guard time between timeslots. After the final beacon, a guard time of 0.3 seconds occurs before beamforming. The experiments were automated by creating a compact disk with the one second primary beacon signal repeating every 7 seconds for the two-source tests and every 10 seconds for the three-source tests.

TABLE II

TWO-SOURCE ROUND-TRIP SYNCHRONIZATION PROTOCOL TIMING.

time	S ₁	S ₂
0.00s	detect primary beacon	detect primary beacon
0.00-0.10s	wait	wait
0.10-0.60s	track PLL1	track PLL1
0.60-1.00s	holdover PLL1	holdover PLL1
1.00-1.25s	holdover PLL1	holdover PLL1
1.25-1.35s	transmit secondary beacon using holdover PLL1	holdover PLL1
1.35-1.85s	transmit secondary beacon using holdover PLL1	holdover PLL1 and track PLL2
1.85-2.25s	transmit secondary beacon using holdover PLL1	holdover PLL1 and holdover PLL2
2.25-2.50s	wait	holdover PLL1 and holdover PLL2
2.50-2.60s	wait	transmit secondary beacon using holdover PLL1; also holdover PLL2
2.60-3.10s	track PLL2	transmit secondary beacon using holdover PLL1; also holdover PLL2
3.10-3.50s	holdover PLL2	transmit secondary beacon using holdover PLL1; also holdover PLL2
3.50-3.80s	holdover PLL2	holdover PLL2
3.80-5.80s	transmit carrier using holdover PLL2	transmit carrier holdover PLL2
5.80-6.80s	clear state	clear state
6.80s	re-arm primary beacon detector	re-arm primary beacon detector

A. Source Node Functionality and PLLs

The primary beacon detection and round-trip protocol functionalities were implemented by programming the TMS320C6713DSKs in C using Texas Instrument’s Code Composer Studio integrated development environment. Each source node runs identical software and determines its identity by polling a bank of DIP switches upon initialization. In order to reduce the likelihood of false detection of the primary beacon caused by room noise, a second-order IIR filter with peak frequency 1021 Hz and bandwidth of 100 Hz is used to filter all of the signals prior to subsequent processing.

The discrete-time phase locked loops in each source node are implemented in software. Depending on the number of nodes and the node number, as many as three independent phase locked loops are implemented on a source node. The

PLL loop filter is realized by following the analog active-PI loop filter design procedure in [5] with 3 dB bandwidth of approximately 13 Hz and then using the bilinear transform to convert the analog loop filter to discrete time. The PLL’s “voltage controlled oscillator” is implemented in software as a numerically controlled oscillator (NCO) centered at the nominal frequency of 1021Hz. All processing is performed on the DSP in floating point.

The phase detector in each PLL is implemented in two stages. In the “rough acquisition” stage (the first 0.34 seconds of tracking), the PLL uses a phase-frequency detector (PFD) [5]. The PFD is used for two reasons. First, unlike most other phase detectors, e.g. the multiplier, the PFD does not possess any unstable equilibria and convergence times are predictable. Second, the PFD output is independent of the input amplitude. Hence, the PLL can perform rough acquisition without automatic gain control. The PFD output after convergence, however, has occasional transients that are not fully suppressed by the loop filter which can lead to inconsistent beamforming performance. Hence, after rough acquisition, the PLL switches to “fine acquisition” for the remainder of the tracking period by changing the phase detector to a standard multiplier. The multiplier phase detector does not have output transients like the PFD after convergence, but is not suitable for rough acquisition due to its sensitivity to input amplitude and unpredictable convergence times caused by the presence of unstable equilibria. During fine acquisition, the input signal is normalized to unity amplitude by using a local estimate of the signal amplitude obtained during rough acquisition so that the PFD and the standard multiplier can share the same phase detector gain. This is done to ensure the phase detector gain is consistent between the PFD and multiplier phase detectors. Inconsistent phase detector gain may result slow convergence for the multiplier phase detector.

B. Data Analysis Methodology

At the conclusion of an experiment consisting of N acoustic beamforming tests, the uncompressed .wav recording of the experiment was transferred to a PC and analyzed in MATLAB to generate the statistical results presented in Section V. To quantify the efficacy of the distributed beamformer, the “power ratio” ρ of the beamformer is calculated by estimating the power received during beamforming and computing its ratio with respect to the ideal beamforming gain when the carriers are received in perfect phase alignment. A power ratio of one corresponds to an ideal beamformer with perfect phase alignment. A power ratio of zero corresponds to the case where the carriers completely cancel at the destination.

To understand how power ratio is computed from the recordings, Figure 10 shows a figurative example of a typical recording for a two-source round-trip beamforming test. Since the secondary beacons in $TS^{(1)}$ and $TS^{(2)}$ are transmitted at the same amplitude as the carriers in $TS^{(3)}$, the power ratio of the n^{th} test can be computed by estimating the amplitudes of the signals recorded in timeslots $TS^{(1)}$, $TS^{(2)}$, and $TS^{(3)}$

and calculating

$$\rho[n] = \left(\frac{\hat{a}_{bf}[n]}{\hat{a}_{10}[n] + \hat{a}_{20}[n]} \right)^2.$$

The amplitude estimates in each test are obtained via the MLE FFT technique described in [6] using an 0.2 second window of the steady state portion of each signal in timeslots $TS^{(1)}$, $TS^{(2)}$, and $TS^{(3)}$.

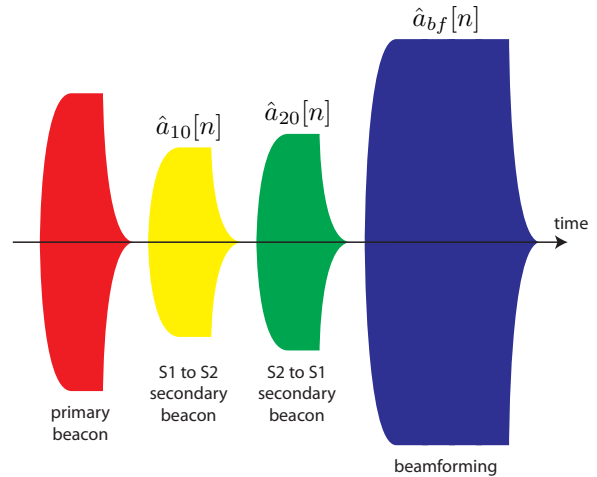


Fig. 10. Amplitude estimation in the n^{th} test of a two-source round-trip distributed beamforming experiment.

In the general M -source case, the power ratio of the n^{th} test can be calculated as

$$\rho[n] = \left(\frac{\hat{a}_{bf}[n]}{\hat{a}_{10}[n] + \dots + \hat{a}_{M0}[n]} \right)^2.$$

As with two nodes, a power ratio of one corresponds to an ideal beamformer with perfect phase alignment.

V. EXPERIMENTAL RESULTS

Figure 11 shows a histogram of the two-source power ratios over $N = 100$ tests of the time-slotted round-trip carrier synchronization protocol with one second beacons and 0.25 second guard times. The mean power ratio of the distributed beamformer was computed to be 0.977 and the standard deviation was computed to be approximately 0.003.

Figure 12 shows a histogram of the three-source power ratios over 100 tests of the time-slotted round-trip carrier synchronization protocol with one second beacons and 0.25 second guard times. The mean power ratio of the distributed beamformer was computed to be 0.907 and the standard deviation was computed to be approximately 0.006.

These results show that the time-slotted round-trip carrier synchronization protocol was consistently effective at synchronizing the phase of the carriers of the ACCENT nodes and that the synchronization errors lead to only a small loss in performance with respect to the ideal beamforming gain. The following section discusses the factors that contribute to the non-ideal performance observed in the acoustic experiments.

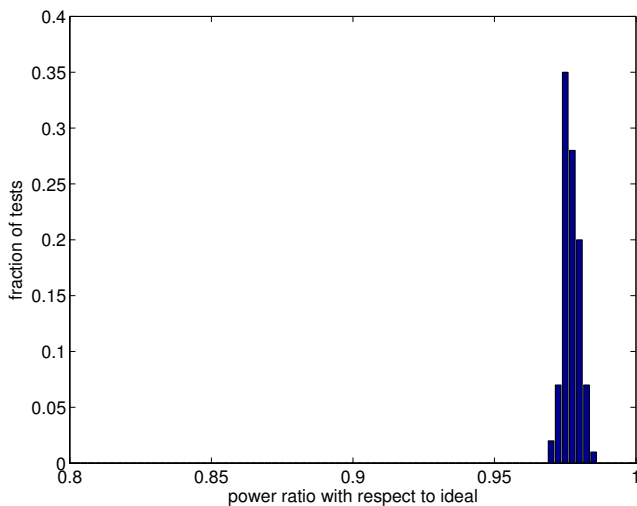


Fig. 11. Two-source power ratio distribution for an experiment with $N = 100$ tests.




Fig. 12. Three-source power ratio distribution for an experiment with $N = 100$ tests.

A. Discussion

Prior to performing the acoustic experiments, the real-time implementation of the round-trip protocol was tested over *wired* channels by connecting the line-level outputs of the DSKs to an audio mixer and connecting the line-level inputs of the DSKs to the output of the mixer. The CD player used for primary beacon generation was also connected to the audio mixer. Several wired-channel experiments were performed and these experiments consistently resulted in power ratios greater than 0.98 and standard deviations on the order of 10^{-3} . Hence, the wired-channel experiments confirmed that the round-trip carrier synchronization protocol can consistently offer near-ideal performance over “perfect” channels.

The two-source acoustic results shown in Figure 11 are very similar to the results obtained over wired channels. The average power ratio of the three-source results shown

in Figure 12, however, is somewhat lower than the average power ratios observed in the wired experiments. One important factor in the acoustic experiments is that the microphone and speaker at each source node (and at the destination) are separate transducers in slightly different locations with different radiation patterns. Hence, the channel reciprocity between each pair of nodes required by the round-trip protocol is only approximate. Also, the mean power ratio results were sensitive to the position of the microphone at the destination. The best results were obtained when the microphone was placed such that the amplitudes of the secondary beacons were similar. Room reverberation and background noise (primarily caused by air conditioning) also affect the PLLs as well as the accuracy of the amplitude estimates generated in the analysis of the results.

It is worth emphasizing that the two-source and three-source power ratio results over acoustic channels are nevertheless consistent with the wired channel results in that the standard deviation of the acoustic experiments is similar to the standard deviation of the wired channel results. The consistency of these results confirms that the source node PLLs are converging consistently and that the round-trip protocol can be used to realize a distributed beamformer with near-ideal performance and low computational complexity even in noisy multipath channels.

VI. CONCLUSION

This paper presents the first experimental results for time-slotted round-trip carrier synchronization. Two-source and three-source wireless acoustic distributed beamforming systems were built and tested in a room with noisy multipath channels. The 1021 Hz acoustic signals used for the beacons and carriers had a wavelength equivalent to 900MHz electromagnetic propagation. The results in this paper confirm that a distributed beamformer using time-slotted round-trip carrier synchronization can consistently achieve a large fraction of the power gains of an ideal conventional beamformer.

REFERENCES

- [1] R. Mudumbai, D.R. Brown III, U. Madhow, and H.V. Poor, “Distributed transmit beamforming: Challenges and recent progress,” *IEEE Communications Magazine*, vol. 47, no. 2, pp. 102–110, February 2009.
- [2] 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.
- [3] R. Mudumbai, B. Wild, U. Madhow, and K. Ramchandran, “Distributed beamforming using 1 bit feedback: from concept to realization,” in *44th Allerton Conf. on Comm., Control, and Computing*, Monticello, IL, Sep. 2006, pp. 1020 – 1027.
- [4] M. Seo, M. Rodwell, and U. Madhow, “A feedback-based distributed phased array technique and its application to 60-ghz wireless sensor network,” in *IEEE MTT-S International Microwave Symposium Digest*, Atlanta, GA, June 15-20, 2008, pp. 683–686.
- [5] R. Best, *Phase-Locked Loops : Design, Simulation, and Applications*. New York: McGraw-Hill, 2003.
- [6] 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.