Time-Slotted Round-Trip Carrier Synchronization

Ipek Ozil and D. Richard Brown III
Electrical and Computer Engineering Department
Worcester Polytechnic Institute
Worcester, MA 01609
email: {ipek,drb}@wpi.edu

***Abstract*— We consider the problem of synchronizing the carriers of two sources in a wireless communication system with one destination. Carrier synchronization has been considered recently in cooperative communication systems where the sources wish to pool their antenna resources and transmit as a "distributed beamformer". Based on the concept of round-trip carrier synchronization first described in [1], we propose a new time-slotted round-trip carrier synchronization system and describe its implementation in systems with single-path or multipath channels. The performance of the time-slotted round-trip carrier synchronization system is investigated in terms of the phase offset at the destination and the expected beamforming time before resynchronization is required. Our results suggest that the synchronization overhead can be small with respect to the potential beamforming gains.**

I. Introduction

In multiuser wireless communication systems, the term "distributed beamforming" describes the case when two or more transmitters align the phase of their bandpass transmissions in order to emulate a centralized antenna array and focus a common transmission toward an intended destination [2]. The SNR gains of conventional beamforming are well documented in the literature. Distributed beamforming, however, is complicated by the fact that the transmitters have independent local oscillators. Transmitters in a distributed beamformer require some method to precisely synchronize their carrier signals so that they arrive with reasonable phase alignment at the intended destination. In addition to distributed beamforming, this sort of synchronization can also facilitate cooperative communication protocols that assume coherent combining of source and relay transmissions at the destination, e.g. the protocols described in [3] as well as the space-time cooperative protocols in [4].

A distributed beamforming scheme was proposed in [5] using a master synchronization beacon and knowledge of the relative source node locations. A synchronization scheme allowing for unknown transmitter locations was described in [6] where a beacon is used to estimate the phase delay between each transmitting node and the destination. Quantized versions of these estimates are transmitted to the source nodes for local phase pre-compensation. A round-trip carrier synchronization system using continuously transmitted beacons at different frequencies was proposed in [1]. While the continuously transmitted beacons allowed for high rates of source and/or destination mobility, the use of different frequencies for each beacon resulted in degraded performance in multipath channels.

This paper presents a new time-slotted round-trip carrier synchronization protocol that uses a single frequency for all beacons and satisfies the half-duplex constraint through time division of the channel. Implementation details are given for both single-path and multipath channels. We investigate the performance of the proposed protocol in terms of the phase offset during beamforming and the expected beamforming time before resynchronization is required. The performance of the time-slotted round-trip carrier synchronization protocol, unlike the protocol in [1], is not degraded in time-invariant multipath channels since the channels are accessed at the same frequency in both directions. Our results also show that the synchronization overhead can be small with respect to the potential beamforming gains in many cases.

This work was supported by NSF award CCF-0447743.

II. System Model

We consider the two-source one-destination system model shown in Figure 1. The destination (node 0) and both sources (nodes 1 and 2) are assumed to each possess a single isotropic antenna. The channel from node i to node j is modeled as an LTI system with impulse response $g_{ij}(t)$. Each channel in the system is assumed to be FIR with delay spread ν_{ij}. The impulse response of each channel in the system is assumed to be reciprocal in the forward and reverse directions, i.e., $g_{ij}(t) = g_{ji}(t)$, and the noise in each channel is assumed to be Gaussian and white with power spectral density $N_0/2$.

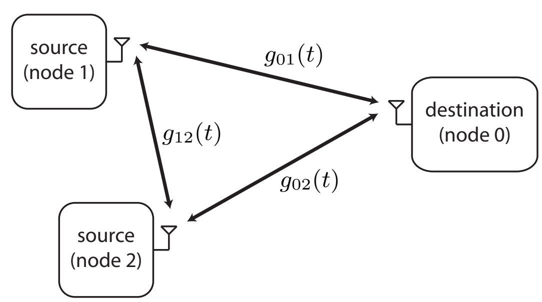

Fig. 1. Two-source distributed beamforming system model.

All nodes in the system are required to satisfy the half-duplex constraint. We also assume that both sources have identical information to transmit to the destination. Low-overhead techniques for disseminating information between source nodes in a distributed beamformer are outside of the scope of this paper but have been considered recently in [7].

III. Time-Slotted Synchronization Protocol

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:

TS0: The destination (D) transmits the sinusoidal primary beacon to both sources (S_1 and S_2). Both sources generate phase and frequency estimates from their local observations.

TS1: S_1 transmits a sinusoidal secondary beacon to S_2. This secondary beacon is transmitted at the frequency estimated by S_1 in TS0 and with initial phase extrapolated from the phase and frequency estimates in TS0. S_2 generates local phase and frequency estimates from this observation.

TS2: S_2 transmits a sinusoidal secondary beacon to S_1. This secondary beacon is transmitted at the frequency estimated by S_2 in TS0 and with initial phase extrapolated from the phase and frequency estimates in TS0. S_1 generates local phase and frequency estimates from this observation.

TS3: 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 obtained from the secondary beacon observation.

For long transmissions, the synchronization timeslots TS0-TS2 may need to be repeated in order to avoid unacceptable phase drift between the sources during beamforming. The following sections describe the protocol in more detail for the case of single-path and multipath time-invariant channels.

A. Single-Path Time-Invariant Channels

The time-slotted protocol begins in TS0 with the transmission of a unit-amplitude sinusoidal primary beacon of duration T_0 from the destination to both sources,

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

Under the assumption that all of the channels are single-path, i.e. $g_{ij}(t) = \alpha_{ij}\delta(t-\tau_{ij})\ \forall ij$, the signal received at S_j in TS0 can be written as

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

for $t \in [t_0+\tau_{0j}, t_0+\tau_{0j}+T_0)$ where $\eta_{0j}(t)$ denotes the AWGN in the $0 \rightarrow j$ channel and $j \in \{1,2\}$. Each source uses their noisy observation from the first timeslot to compute estimates of the received frequency and phase; these estimates are denoted as $\hat{\omega}_{0j}$ and $\hat{\phi}_{0j}$, respectively, at S_j for $j \in \{1,2\}$. We use the usual convention that the phase estimate $\hat{\phi}_{0j}$ is an estimate of the phase of the received signal at the start of the observation at S_j, i.e. $\hat{\phi}_{0j}$ is an estimate of the phase of $y_{0j}(t)$ at time $t_0+\tau_{0j}$.

Timeslot TS1 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 begins transmitting a sinusoidal secondary beacon to S_2 that is a periodic extension of $y_{01}(t)$ (possibly with different amplitude) using the phase and frequency estimates $\hat{\omega}_{01}$ and $\hat{\phi}_{01}$. To generate the periodic extension, the frequency estimate $\hat{\omega}_{01}$ is used to extrapolate the estimated phase of $y_{01}(t)$ at time $t_0+\tau_{01}$ to a phase estimate of $y_{01}(t)$ at time t_1. The extrapolated phase estimate at S_1 at time t_1 can be written as

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

The secondary beacon transmitted by S_1 in TS1 can be written as

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

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

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

for $t \in [t_1+\tau_{12}, t_1+\tau_{12}+T_1)$ where $\eta_{12}(t)$ denotes the AWGN in the $1 \rightarrow 2$ channel. From this noisy observation, S_2 generates estimates of the received frequency and phase; these estimates are denoted as $\hat{\omega}_{12}$ and $\hat{\phi}_{12}$, respectively.

Timeslot TS2 begins immediately upon the conclusion of $y_{12}(t)$ at S_2. At time $t_2 = t_1+\tau_{12}+T_1$, S_2 begins transmitting a sinusoidal secondary beacon to S_1 that is a periodic extension of $y_{02}(t)$ using the phase and frequency estimates $\hat{\omega}_{02}$ and $\hat{\phi}_{02}$. Note that S_2's secondary beacon is a periodic extension of the *primary* beacon it received in TS0 even though its transmission begins at the conclusion of the secondary beacon received in TS1. Here, S_2 extrapolates the phase estimate $\hat{\phi}_{02}$ obtained at time $t_0+\tau_{02}$ to time t_2 using the frequency estimate $\hat{\omega}_{02}$ in order to determine the appropriate initial phase of the secondary beacon. The extrapolated phase estimate at S_2 at time t_2 can be written as

$$\begin{aligned}\hat{\phi}_2 &= \hat{\phi}_{02}+\hat{\omega}_{02}(t_2-(t_0+\tau_{02})) \\ &= \hat{\phi}_{02}+\hat{\omega}_{02}(\tau_{01}+\tau_{12}-\tau_{02}+T_0+T_1).\end{aligned}$$

The secondary beacon transmitted by S_2 in TS2 can be written as

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

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

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

for $t \in [t_2+\tau_{12}, t_2+\tau_{12}+T_2)$ where $\eta_{21}(t)$ denotes the AWGN in the $1 \rightarrow 2$ channel and where we have applied the assumption that $\tau_{21} = \tau_{12}$ and $\alpha_{21} = \alpha_{12}$. From this noisy observation, S_1 generates estimates of the received frequency and phase; these estimates are denoted as $\hat{\omega}_{21}$ and $\hat{\phi}_{21}$, respectively.

In timeslot TS3, both S_1 and S_2 transmit to the destination as a distributed beamformer with carriers generated as periodic extensions of the *secondary* beacons received at each source. Since our focus is on the performance of the distributed beamformer in terms of the phase difference of the received signals at the destination, we write the transmissions of S_1 and S_2 as unmodulated carriers. The unmodulated carrier transmitted by S_j during TS3 can be written as

$$x_{j0}(t) = a_{j0}\cos(\hat{\omega}_j(t-t_{3j})+\hat{\phi}_{3j}) \quad t \in [t_{3j}, t_{3j}+T_3) \quad (2)$$

where $\hat{\omega}_j$ is a frequency estimate at S_j that, as discussed in Section IV-B, is a function of both $\hat{\omega}_{0j}$ and $\hat{\omega}_{ij}$, $i \neq j$. The extrapolated phase estimates at times t_{31} and t_{32} are based on the phase and frequency estimates obtained from the secondary beacon observations and can be written as

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

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

respectively. As for the transmission start times t_{31} and t_{32}, S_1 begins transmitting its carrier immediately upon the conclusion of the secondary beacon from S_2, hence

$$t_{31} = t_2+\tau_{12}+T_2 = t_0+\tau_{01}+2\tau_{12}+T_0+T_1+T_2.$$

If S_2 begins transmitting immediately upon the conclusion of its secondary beacon transmission, its carrier will arrive at D earlier than S_1's carrier. To synchronize the arrivals of the carriers, S_2 should wait for $\tau_{delay} = \tau_{01}+\tau_{12}-\tau_{02}$ after the conclusion of the transmission of its secondary beacon before transmitting its carrier in timeslot TS3. This implies that

$$t_{32} = t_2+T_2+\tau_{delay} = t_0+2\tau_{01}+2\tau_{12}-\tau_{02}+T_0+T_1+T_2.$$

By inspection of Figure 1, we note that τ_{delay} must be non-negative. Moreover, S_2 can directly estimate τ_{delay} by observing the amount of time that elapses from the end of its primary beacon observation in TS0 to the start of its secondary beacon observation in TS1, i.e. $\tau_{delay} = (t_1+\tau_{12})-(t_0+\tau_{02}+T_0)$.

The signal received at D in TS3 can be written as the sum of both carrier transmissions after their respective channel delays, i.e.,

$$y_0(t) = \alpha_{01} a_{10} \cos(\hat{\omega}_1(t - t_3) + \hat{\phi}_{31}) + \alpha_{02} a_{20} \cos(\hat{\omega}_2(t - t_3) + \hat{\phi}_{32}) + \eta_0(t)$$

for $t \in [t_3, t_3 + T_3)$ where $t_3 = t_{31} + \tau_{01} = t_{32} + \tau_{02}$. Standard trigonometric identities can be applied to rewrite $y_0(t)$ as

$$y_0(t) = a_{bf}(t) \cos(\phi_{bf}(t)) + \eta_0(t) \quad t \in [t_3, t_3 + T_3) \quad (5)$$

where

$$a_{bf}(t) := \sqrt{(\alpha_{01} a_{10})^2 + (\alpha_{02} a_{20})^2 + 2\alpha_{01} a_{10} \alpha_{02} a_{20} \cos(\phi_\Delta(t))}$$
$$\phi_{bf}(t) := \hat{\omega}_1(t - t_3) + \hat{\phi}_{31} + \tan^{-1}\left[\frac{\alpha_{02} a_{20} \sin(\phi_\Delta(t))}{\alpha_{01} a_{10} + \alpha_{02} a_{20} \cos(\phi_\Delta(t))}\right]$$

and where we have defined the carrier phase offset

$$\phi_\Delta(t) := (\hat{\omega}_2 - \hat{\omega}_1)(t - t_3) + \hat{\phi}_{32} - \hat{\phi}_{31} \quad t \in [t_3, t_3 + T_3).$$

In the special case when carriers arrive at the destination with the same amplitude, i.e. $\alpha_{01} a_{10} = \alpha_{02} a_{20} = a$, the expressions for $a_{bf}(t)$ and $\phi_{bf}(t)$ simplify to

$$a_{bf}(t) = 2a \cos(\phi_\Delta(t)/2) \quad t \in [t_3, t_3 + T_3)$$
$$\phi_{bf}(t) = \left[(\hat{\omega}_1 + \hat{\omega}_2)(t - t_3) + \hat{\phi}_{31} + \hat{\phi}_{32}\right]/2$$

B. Multipath Time-Invariant Channels

The time-slotted round-trip synchronization protocol can also be effective in communication systems with multipath LTI channels if minor modifications are made to account for the transient effects of the channels. This section summarizes the necessary modifications.

As with single-path channels, the time-slotted synchronization protocol begins with the transmission of a sinusoidal primary beacon of duration T_0 from the destination as in (1). Since the beacon is of finite duration, the signals received by S_1 and S_2 will have an initial transient component, a steady state component, and a final transient component. It can be shown that the duration of the steady state component at S_j is equal to $T_0 - \nu_{0j}$, where ν_{0j} denotes the delay spread of channel $g_{0j}(t)$. In order to achieve a steady-state response at both S_1 and S_2, we require $T_0 > \max(\nu_{01}, \nu_{02})$. The steady-state portion of the beacon received at S_j can then be written as

$$y_{0j}(t) = \alpha_{0j} \cos(\omega t + \phi_0 + \theta_{0j}) + \eta_{0j}(t)$$

for $t \in [t_0 + \tau_{0j} + \nu_{0j}, t_0 + \tau_{0j} + T_0)$ and $j \in \{1, 2\}$. Each source uses *only the steady-state portion* of their noisy observation in the first timeslot to compute local estimates of the received frequency and phase. The transient portions of the observation are ignored.

The second and third timeslots are as described in the single-path case, with each source transmitting secondary beacons to the other source using the frequency and extrapolated phase estimates obtained from the first timeslot. The phase estimates at each source are extrapolated for transmission of the secondary beacons as periodic extensions of the *steady state* portion of the primary beacon observations. The only differences with respect to the single path case are that (i) the duration of each secondary beacon must exceed $\nu_{12} = \nu_{21}$ in order to ensure a steady-state observation and that (ii) the sources estimate the received frequency and phase of the secondary beacons using only the steady-state portion of the observations.

No other modifications to the synchronization protocol are necessary. In the final timeslot, both sources transmit as in (2). Assuming unmodulated carriers, the steady-state signal received at the destination during the final timeslot can be written in the same form as (5). The net effect of multipath on the synchronization protocol is that the beacons must be transmitted with durations exceeding the delay spread of the appropriate channels and that the duration of the steady-state observations used for phase and frequency estimation are reduced, with respect to single-path channels, by the delay spread of the multipath channels.

C. Discussion

Although the events of the time-slotted round-trip synchronization protocol described in Section III-A are described in terms of some notion of "true time" t, it is worth mentioning that the protocol does not assume that nodes share a common time reference. An essential feature of the protocol is that, in each of the timeslots TS1, TS2, and TS3, each source transmission is simply a periodic extension of a beacon received in a previous timeslot. No absolute notion of "time-zero" is needed since the phase of a source's transmission is extrapolated from the estimated initial phase of the appropriate beacon observation in a previous timeslot. Moreover, each source transmission in timeslots TS1, TS2, and TS3 is triggered by the conclusion of a beacon in a prior timeslot. The sources do not follow any schedule requiring knowledge of "true time".

The fact that sources have imperfect local oscillators also implies that local frequency estimates at each source are relative to the source's clock. Suppose, for example, a source's clock $t' = \gamma t$ runs at rate γ with respect to true time t and that the channels are noiseless. A beacon received at true frequency ω rad/sec will appear to this source to be received at frequency $\hat{\omega} = \omega/\gamma$. Nevertheless, when the source generates a periodic extension of this signal in a later timeslot, the frequency of the transmission will be equal to the product of the local estimate and the local relative clock rate, i.e. ω. Hence, the protocol does not require the sources to share a common time reference, either in terms of clock rate or phase.

We also point out that, as long as the half-duplex constraint is not violated, the absolute starting and ending times of each of the timeslots are not critical to the performance of the protocol. Since each source transmission in timeslots TS1, TS2, and TS3 is a periodic extension of a beacon received in a prior timeslot, gaps of arbitrary duration can be inserted between the timeslots without directly affecting the phase offset at the destination during beamforming. Gaps between the timeslots may be needed in practical systems, for example, to account for processing time at the sources and/or transient components of beacons received in multipath. In any case, these gaps do not directly affect the relative phase of the periodic extensions since they essentially delay the window in which the periodic extension is transmitted but do not change the phase or frequency of the periodic extension. As a consequence of this property, the estimate of τ_{delay} at S_2 in TS3 is not critical if the beamforming timeslot is sufficiently long. An inaccurate estimate of τ_{delay} only causes S_2's carrier to begin slightly earlier or later than S_1's carrier at D; it does not affect the relative phase of the carriers during beamforming.

IV. Performance Analysis

This section analyzes the performance of the time-slotted round-trip carrier synchronization protocol in terms of the carrier phase offset at the receiver during the beamforming timeslot. In an ideal beamformer, the amplitudes of the received signals add constructively at the destination and $a_{bf}(t) = \alpha_{01} a_{10} + \alpha_{02} a_{20}$. The non-ideal nature of the distributed beamformer is captured in the carrier phase offset

$$\phi_\Delta(t) = \omega_\Delta(t - t_3) + \phi_\Delta \qquad t \in [t_3, t_3 + T_3) \quad (6)$$

where $\omega_\Delta := \hat{\omega}_2 - \hat{\omega}_1$ represents the linear phase drift during beamforming and $\phi_\Delta := \hat{\phi}_{32} - \hat{\phi}_{31}$ represents the initial phase offset at the start of beamforming. Phase and frequency estimation errors at each source result in unavoidable initial carrier phase offset at the start of TS3 as well as linear phase drift over the duration of TS3. The following section establishes a vector notation for the eight estimation errors of the time-slotted round-trip carrier synchronization protocol and analyzes the joint statistics of these errors to facilitate analysis of the carrier phase offset during beamforming.

A. Statistics of the Frequency and Phase Estimation Errors

In the time-slotted round-trip carrier synchronization protocol, each source generates a pair of frequency estimates and a pair of phase estimates from the primary and secondary beacon observations. We define the estimation error vector

$$\tilde{\boldsymbol{\theta}} \quad := \quad [\tilde{\omega}_{01}, \tilde{\omega}_{02}, \tilde{\omega}_{12}, \tilde{\omega}_{21}, \tilde{\phi}_{01}, \tilde{\phi}_{02}, \tilde{\phi}_{12}, \tilde{\phi}_{21}]^\top$$

where $\tilde{\omega}_{0j} := \hat{\omega}_{0j} - \omega$, $\tilde{\omega}_{ij} := \hat{\omega}_{ij} - \hat{\omega}_{0i}$, $\tilde{\phi}_{0j} := \hat{\phi}_{0j} - \phi_0$, and $\tilde{\phi}_{ij} := \hat{\phi}_{ij} - \hat{\phi}_i$ for $j \in \{1,2\}$, $i \in \{1,2\}$, and $i \neq j$. Note that the frequency and phase estimation errors $\tilde{\omega}_{0j}$ and $\tilde{\phi}_{0j}$ are defined with respect to the primary beacon frequency and phase transmitted by $\mathsf{D} \rightarrow \mathsf{S}_j$. The frequency and phase estimation errors $\tilde{\omega}_{ij}$ and $\tilde{\phi}_{ij}$ are defined with respect to the secondary beacon frequency and phase transmitted by $\mathsf{S}_i \rightarrow \mathsf{S}_j$.

To facilitate analysis, we assume the estimation error vector is Gaussian distributed with zero mean and covariance matrix $\boldsymbol{\Theta} := \mathrm{E}\left[\tilde{\boldsymbol{\theta}}\tilde{\boldsymbol{\theta}}^\top\right]$. We note that the frequency estimation errors are all independent since (i) observations in different timeslots are affected by independent noise realizations and (ii) observations at S_1 and S_2 are affected by independent noise realizations. This is also true of the phase estimates. The frequency and phase estimates obtained from the same observation, however, are not independent. Hence, all of the off-diagonal elements of the covariance matrix are equal to zero except for the terms $\mathrm{cov}[\tilde{\omega}_{ij}\tilde{\phi}_{ij}]$ for $i,j \in \{1,2\}$.

The variances on the diagonal of Θ and the covariances on the off-diagonals of Θ can be lower bounded by the Cramer-Rao bound (CRB). Given a sinusoid of amplitude a in white noise with PSD $\frac{N_0}{2}$, the variances and covariance of the frequency and phase estimates can be lower bounded by [8]

$$\sigma_\omega^2 \quad \geq \quad 12N_0/(a^2T^3) \tag{7}$$
$$\sigma_\phi^2 \quad \geq \quad 4N_0/(a^2T) \tag{8}$$
$$\mathrm{cov}[\tilde{\omega}, \tilde{\phi}] \quad \geq \quad -6N_0/(a^2T^2) \tag{9}$$

where T is the duration of the observation.

B. Carrier Frequency Offset

Since each source has a pair of unbiased frequency estimates prior to the start of beamforming, we can reduce the phase drift during beamforming by generating the carrier at S_j from a linear combination of the local estimates, i.e.,

$$\hat{\omega}_j \quad = \quad \mu_j \hat{\omega}_{0j} + (1-\mu_j)\hat{\omega}_{ij}.$$

In this case, the carrier frequency offset during beamforming can be written as

$$\omega_\Delta := \hat{\omega}_2 - \hat{\omega}_1 = \boldsymbol{\Gamma}_1^\top \tilde{\boldsymbol{\theta}} \tag{10}$$

where

$$\boldsymbol{\Gamma}_1 = [1-\mu_1-\mu_2, -(1-\mu_1-\mu_2), 1-\mu_2, -(1-\mu_1), 0, 0, 0, 0]^\top.$$

It can be shown that the carrier frequency offset ω_Δ is Gaussian distributed with zero mean for any choice of μ_1 and μ_2 when the frequency estimates are unbiased and Gaussian distributed. A good choice then for the linear combination parameters μ_1 and μ_2 is one that minimizes $\mathrm{var}[\omega_\Delta]$. For $j \in \{1,2\}$ and $i \neq j$, the linear combination parameters that minimize the variance can be determined using standard calculus techniques to be

$$\mu_j^* \quad = \quad \frac{1}{1 + \frac{\sigma^2_{\omega_{ji}}}{\sigma^2_{\omega_{ij}}}\left(\frac{\sigma^2_{\omega_{01}} + \sigma^2_{\omega_{02}}}{\sigma^2_{\omega_{01}} + \sigma^2_{\omega_{02}} + \sigma^2_{\omega_{ji}}}\right)}$$

C. Initial Carrier Phase Offset

From (3) and (4), the carrier phase offset at the start of beamforming can be written as

$$\phi_\Delta = \left[\hat{\phi}_{12} + \hat{\omega}_{12}(t_{32} - (t_1 + \tau_{12}))\right] - \left[\hat{\phi}_{21} + \hat{\omega}_{21}(t_{31} - (t_2 + \tau_{12}))\right]. \tag{11}$$

The secondary beacon frequency estimates $\hat{\omega}_{21}$ and $\hat{\omega}_{12}$ can be written as

$$\hat{\omega}_{21} \quad = \quad \hat{\omega}_{02} + \tilde{\omega}_{21} \; = \; \omega + \tilde{\omega}_{02} + \tilde{\omega}_{21} \text{ and}$$
$$\hat{\omega}_{12} \quad = \quad \hat{\omega}_{01} + \tilde{\omega}_{12} \; = \; \omega + \tilde{\omega}_{01} + \tilde{\omega}_{12}$$

and the phase estimate $\hat{\phi}_{12}$ can be written as

$$\hat{\phi}_{12} = \hat{\omega}_{01}T_0 + \hat{\phi}_{01} + \tilde{\phi}_{12} = (\omega + \tilde{\omega}_{01})T_0 + \phi_0 + \tilde{\phi}_{01} + \tilde{\phi}_{12}$$

where $\tilde{\phi}_{01}$ is the primary beacon phase estimation error at S_1 and $\tilde{\phi}_{12}$ is the secondary beacon phase estimation error error at S_2. Similarly, the phase estimate $\hat{\phi}_{21}$ can be written as

$$\begin{aligned}\hat{\phi}_{21} &= \hat{\omega}_{02}(\tau_{01} + \tau_{12} - \tau_{02} + T_0 + T_1) + \hat{\phi}_{02} + \tilde{\phi}_{21} \\ &= (\omega + \tilde{\omega}_{02})(\tau_{01} + \tau_{12} - \tau_{02} + T_0 + T_1) + \phi_0 + \tilde{\phi}_{02} + \tilde{\phi}_{21}.\end{aligned}$$

Letting $\Psi := t_{32} - (t_1 + \tau_{12}) = \tau_{01} + \tau_{12} - \tau_{02} + T_1 + T_2$ and noting that $t_{31} - (t_2 + \tau_{12}) = T_2$, we can plug these results into (11) to get

$$\phi_\Delta = \quad = \quad \boldsymbol{\Gamma}_2^\top \tilde{\boldsymbol{\theta}} \tag{12}$$

where $\boldsymbol{\Gamma}_2 = [\Psi + T_0, -(\Psi + T_0), \Psi, -T_2, 1, -1, 1, -1]^\top$.

D. Statistics of the Phase Offset During Beamforming

Plugging (10) and (12) into (6), we can compactly express the phase offset during beamforming in terms of the estimation error vector as

$$\phi_\Delta(t) \quad = \quad [(t - t_3)\boldsymbol{\Gamma}_1 + \boldsymbol{\Gamma}_2]^\top \tilde{\boldsymbol{\theta}} \qquad t \in [t_3, t_3 + T_3).$$

Since $(t - t_3)\boldsymbol{\Gamma}_1$ and $\boldsymbol{\Gamma}_2$ are deterministic and the estimation error vector is assumed to be Gaussian distributed with zero mean, we can say that $\mathrm{E}[\phi_\Delta(t)] \sim \mathcal{N}(0, \sigma^2_{\phi_\Delta(t)})$ at any $t \in [t_3, t_3 + T_3)$. The variance of the phase offset can be written as

$$\sigma^2_{\phi_\Delta(t)} \quad = \quad [(t - t_3)\boldsymbol{\Gamma}_1 + \boldsymbol{\Gamma}_2]^\top \boldsymbol{\Theta} \left[(t - t_3)\boldsymbol{\Gamma}_1 + \boldsymbol{\Gamma}_2\right] \tag{13}$$

for any $t \in [t_3, t_3 + T_3)$. This result can be used to quantify the amount of time that the distributed beamformer provides an acceptable level of carrier phase alignment with a certain level of confidence. At any time $t \in [t_3, t_3 + T_3)$, the probability that the absolute carrier phase offset is less than a given threshold λ can be written as

$$\mathrm{Prob}\left[|\phi_\Delta(t)| < \lambda\right] \quad = \quad 1 - 2Q\left(\frac{\lambda}{\sigma_{\phi_\Delta}(t)}\right) \tag{14}$$

where $Q(x) = \frac{1}{\sqrt{2\pi}}\int_x^\infty \exp(-t^2/2)\,dt$. The CRB results (7), (8), and (9) can be used to provide a lower bound on the variance of the phase offset during beamforming and, as such, an upper bound on $\mathrm{Prob}\left[|\phi_\Delta(t)| < \lambda\right]$.

E. Numerical Results

In this section, we present numerical examples of the time-slotted round-trip carrier synchronization protocol in single-path time-invariant channels. All beacons are transmitted with unit amplitude and each channel is assumed to have unit gain, random propagation delay, and an AWGN PSD of $N_0 = 2.25 \cdot 10^{-12}$ W/Hz. The primary beacon frequency is $\omega = 2\pi \cdot 900 \cdot 10^6$ radians/second. Both sources generate carrier frequencies with the optimum linear combining factors μ_1^* and μ_2^* derived in Section IV-B.

Figure 2 plots Prob $[|\phi_\Delta(t)| < \lambda]$ versus beamforming duration when the primary beacon duration is fixed at $T_0 = 1\mu$s and the secondary beacon durations are fixed at $T_1 = T_2 = 2\mu$s. Both sources generate maximum likelihood phase and frequency estimates of the primary and secondary beacon observations. In addition to the experimental results, Figure 2 also plots the theoretical predictions for Prob $[|\phi_\Delta(t)| < \lambda]$ using (13), (14), and the CRB.

The results in Figure 2 demonstrate that it is possible for the distributed beamformer to provide near-ideal performance with high confidence for a duration much longer than the amount of time spent synchronizing the sources. For example, in the case shown in Figure 2, the beamformer can be expected to maintain an amplitude within 90% of ideal with 95% confidence, i.e. Prob$[\cos(\phi_\Delta(t)/2) \geq 0.9] \geq 0.95$, for $t - t_3 > 300\mu$s. The total time spent synchronizing the carriers in this example was only 5μs, however. This result suggests that the gain obtained by beamforming (in terms of rate improvement or energy savings) is likely to far outweigh the synchronization costs (rate and/or energy loss) in this case.

To understand the effect of the beacon durations on the performance of the beamformer, Figure 3 plots the 95% confidence beamforming time given a 90%-ideal beamforming quality threshold ($\lambda = 2\cos^{-1}(0.9)$) using the CRB analytical predictions. All other parameters are identical to Figure 2. The results in Figure 3 show that the 95% confidence beamforming times are approximately flat when the secondary beacon durations are significantly shorter than the primary beacon duration. When the secondary beacon durations begin to exceed the primary beacon duration, the 95% confidence beamforming times increase at a rate proportional to the secondary beacon durations. If the secondary beacon durations become too long, however, the 95% confidence beamforming time quickly drops to zero. This is due to the fact that the extrapolated phase estimates from the primary beacon become increasingly inaccurate for longer secondary beacon durations. Hence, for a fixed primary beacon duration, these results suggest that the best performance is achieved when the secondary beacon durations are selected to exceed the primary beacon duration, but not by too much.

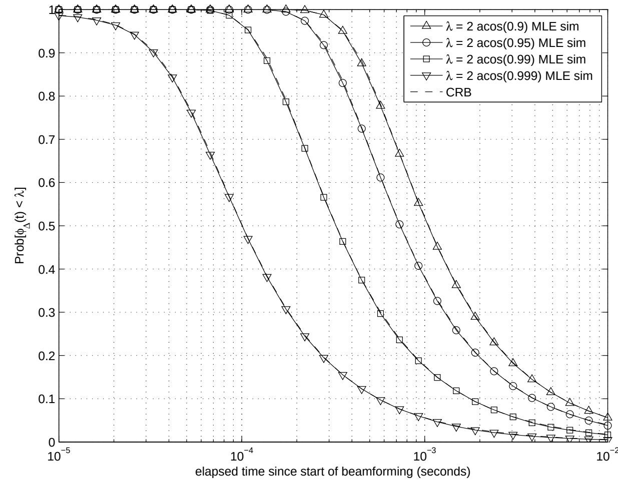

Fig. 2. Prob $[|\phi_\Delta(t)| < \lambda]$ vs. $t - t_3$ when $T_0 = 1\mu$s and $T_1 = T_2 = 2\mu$s.

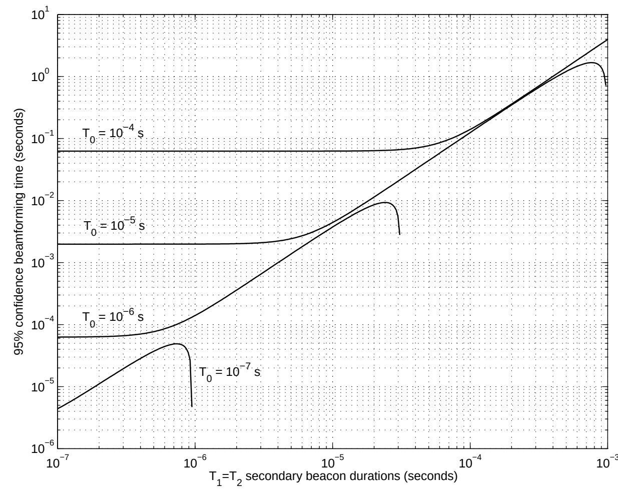

Fig. 3. 95% confidence beamforming time as a function of the primary beacon duration T_0 and secondary beacon durations $T_1 = T_2$ for a 90%-ideal beamforming quality threshold ($\lambda = 2\cos^{-1}(0.9)$).

V. Conclusions

This paper describes an explicit method for synchronizing the carriers of two sources in a cooperative communication system with one destination. The performance of the proposed synchronization system is investigated in terms of the phase offset of the distributed beamformer at the intended destination and the expected beamforming time before resynchronization is required. Our results suggest that the overhead required to synchronize the carriers of the sources can be small with respect to the potential beamforming gains.

References

[1] D.R. Brown III, G. Prince, and J. McNeill, "A method for carrier frequency and phase synchronization of two autonomous cooperative transmitters," in *Proc. IEEE 6th Workshop on Signal Processing Advances in Wireless Communications (SPAWC 2005)*, New York, NY, June 5-8, 2005, pp. 260–264.

[2] R. Mudumbai, G. Barriac, and U. Madhow, "On the Feasibility of Distributed Beamforming in Wireless Networks," *IEEE Transactions on Wireless Communications*, vol. 6, no. 5, pp. 1754–1763, May 2007.

[3] A. Sendonaris, E. Erkip, and B. Aazhang, "User cooperation diversity — part I: System description," and "User cooperation diversity — part II: Implementation aspects and performance analysis," *IEEE Trans. on Communications*, vol. 51, no. 11, pp. 1927–1948, Nov. 2003.

[4] J.N. Laneman and G.W. Wornell, "Distributed space–time-coded protocols for exploiting cooperative diversity in wireless networks," *IEEE Trans. on Information Theory*, vol. 49, no. 10, pp. 2415–2425, Oct. 2003.

[5] G. Barriac, R. Mudumbai, and U. Madhow, "Distributed beamforming for information transfer in sensor networks," in *Inf. Proc. in Sensor Networks (IPSN), Third Intl. Workshop*, Berkeley, CA, April 26-27 2004.

[6] Y.S. Tu and G.J. Pottie, "Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through AWGN channels," in *Proceedings of the IEEE Vehicular Technology Conference (VTC)*, Spring 2002, vol. 1, pp. 130–134.

[7] L. Dong, P. Petropulu, and H. Poor, "A cross-layer approach to collaborative beamforming for wireless ad hoc networks," *IEEE Trans. on Signal Processing*, in review.

[8] D.C. Rife and R.R. Boorstyn, "Single-tone parameter estimation from discrete-time observations," *IEEE Trans. on Information Theory*, vol. IT-20, no. 5, pp. 591–598, Sep. 1974.