Time-Slotted Round-Trip Carrier Synchronization for Distributed Beamforming

D. Richard Brown III, *Member, IEEE*, and H. Vincent Poor, *Fellow, IEEE*

***Abstract*—Transmit beamforming is an energy-efficient wireless communication technique that allows a transmitter with two or more antennas to focus its bandpass signal in an intended direction. Recently, the idea of transmit beamforming has been extended to networks of single-antenna cooperative transmitters that pool their antenna resources and behave as a "distributed beamformer." Unlike conventional transmit beamforming, however, the carriers of the single-antenna transmitters in a distributed beamformer are each synthesized from independent and imperfect local oscillators; carrier phase and frequency synchronization among the transmitters is necessary to ensure that a beam is aimed in the desired direction. In this paper, a new time-slotted round-trip carrier synchronization protocol is proposed to enable distributed beamforming in multiuser wireless communication systems. Numerical results for a distributed beamformer with two transmitters show that the parameters of the round-trip synchronization protocol can be selected such that a desired level of phase accuracy and reliability is achieved during beamforming and such that the synchronization overhead is small with respect to the amount of time that reliable beamforming is achieved. The impact of mobility on the performance of the round-trip carrier synchronization protocol is also shown to be small when the synchronization timeslots are short.**

***Index Terms*—Cooperative transmission, distributed beamforming, sensor networks, synchronization, virtual antenna arrays, wireless networks.**

I. Introduction

In multiuser wireless communication systems, the term "distributed beamforming" describes the situation in which two or more separate transmitters with common information work together to emulate an antenna array and focus their bandpass transmissions toward an intended destination [1]. Distributed beamforming has also been called "collaborative beamforming" [2], "virtual antenna arrays" [3], and has also been discussed in the context of coherent cooperative transmission [4] and cooperative multiple-input/multiple-output (MIMO) transmission [5]. In all of these systems, the basic principle is the same: individual sources with common information transmit with phase aligned carriers such that their bandpass transmissions combine constructively after propagation to the intended destination.

The advantages of conventional transmit beamforming are manifold and well-documented in the literature. By focusing the transmission toward the intended destination, for instance, less transmit power is needed to achieve a desired signal-to-noise ratio (SNR) target. This feature is particularly appealing in wireless communication systems with energy constrained nodes such as sensor networks. In these types of systems, nodes are typically too small to allow for the use of conventional antenna arrays. Distributed beamforming is a powerful technique that offers the potential power gains of conventional antenna arrays to wireless communication systems composed of multiple single-antenna users.

One thing that makes the distributed beamforming problem different from conventional beamforming is that each transmitter in a distributed beamformer has an independent local oscillator. Transmitters in a distributed beamformer require some method to synchronize their carrier signals so that the bandpass transmissions arrive with reasonable phase alignment at the intended destination. As discussed in [1], precise phase alignment is not critical for beamforming: a two-antenna beamformer with a 30 degree phase offset in the received carriers only suffers a loss with respect to ideal beamforming of approximately 7% of the power in the intended direction. The power gain of a beamformer becomes a power penalty, however, when the carriers arrive at the destination with more than 90 degrees of phase offset. Some form of carrier synchronization is required to ensure energy-efficient transmission to the destination and to ensure that the sources do not cancel each other's transmissions.

While there is a large amount of literature describing the applications and potential gains of distributed beamforming, there are relatively few papers that explicitly investigate the practical problem of multiuser carrier synchronization for distributed beamforming. To the best of our knowledge, the first multiuser carrier synchronization scheme suitable for distributed beamforming was proposed in [6]. In [6], a beacon is transmitted by the destination to each source node. The source nodes bounce the beacons back to the destination and the destination estimates and quantizes the phase delay in each channel. These quantized phase delays are then transmitted to the appropriate source nodes for local phase precompensation prior to beamforming. A hierarchical technique for multiuser carrier synchronization and distributed beamforming was recently proposed in [1]. Prior to beamforming, a master source node synchronizes the carriers of the sources using a technique similar to [6]. After the sources have been synchronized, the destination transmits a

Manuscript received December 3, 2007; revised May 1, 2008. First published June 13, 2008; current version published October 15, 2008. This work was supported by the National Science Foundation under Grants CCF-04-47743, ANI-03-38807, and CNS-06-25637. Portions of this paper were presented at the 41st Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, November 4–7, 2007. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Roberto Lopez-Valcarce.

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

H. V. Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: poor@princeton.edu).

Digital Object Identifier 10.1109/TSP.2008.927073

common beacon to the sources to facilitate channel estimation and phase precompensaton for beamforming. The same authors also recently proposed a randomized gradient search technique with one-bit feedback for carrier phase synchronization and distributed beamforming in [7].

A two-source "round-trip" carrier synchronization protocol was proposed in [8] and is distinguished from the aforementioned work by not requiring any digital signaling between nodes to establish synchronization; carrier phase and frequency synchronization is established through three continuously transmitted unmodulated beacons and a pair of phase locked loops (PLLs) at each source node. These beacons are transmitted on different frequencies (also different from the carrier frequency) to avoid simultaneous transmission and reception on the same frequency. This paper was the first to explicitly consider the effect of time-varying and multipath channels on carrier synchronization and showed that, while the continuously transmitted beacons allowed for high rates of source and/or destination mobility, the use of frequency division duplexing for the beacons and carriers resulted in non-reciprocal phase shifts and degraded performance in general multipath channels.

This paper presents a new M-source carrier synchronization protocol using the same "round-trip" approach as [8] except that a single frequency is used for all beacons and carriers. The proposed round-trip synchronization protocol uses time division duplexing for the beacons and carriers. The benefit of using a single frequency for the carrier and all beacons is that channel reciprocity is maintained in multipath propagation scenarios. As discussed in [1], however, the lack of continuous synchronization implies that periodic resynchronization is necessary to avoid unacceptable levels of phase drift during beamforming. We analyze the performance of the proposed "time-slotted round-trip" carrier synchronization protocol for a distributed beamformer with two sources in terms of the statistics of the carrier phase offset at the destination and the expected beamforming time before resynchronization is required in both time-invarant and time-varying channels. Our analysis also accounts for the effects of phase noise in the sources' local oscillators. We provide numerical results demonstrating that the parameters of the synchronization protocol can be selected such that a desired level of phase accuracy and reliability can be achieved during beamforming and such that the synchronization overhead is small with respect to the amount of time that reliable beamforming is achieved. We also show that the parameters of the synchronization protocol can be selected to make reliable beamforming possible even in systems with mobile sources.

II. System Model

We consider the M-source one-destination system model shown in Fig. 1. The destination (node 0) and the sources (nodes $1,\dots,M$) each possess a single isotropic antenna. The channel from node i to node j is modeled as a linear time-invariant (LTI) system[1] with impulse response $g_{i,j}(t)$. The noise in each channel is additive, white, and Gaussian and the impulse response of each channel in the system is assumed to be reciprocal in the forward and reverse directions, i.e., $g_{i,j}(t) = g_{j,i}(t)$.

[1]Time-varying channels are considered in Section V.

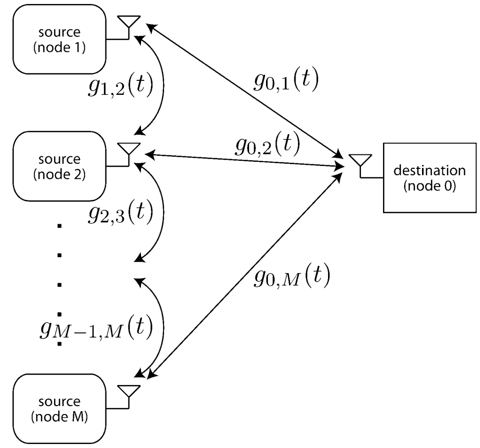

Fig. 1. M-source distributed beamforming system model.

All nodes in the system are required to satisfy the half-duplex constraint in the sense that they cannot transmit and receive signals simultaneously. We assume that all M sources have identical information to transmit to the destination. Low-overhead techniques for disseminating information among source nodes in a distributed beamformer are outside of the scope of this paper but have been considered recently in [9].

A key assumption in this work is that the nodes in Fig. 1 do not possess a common time reference, at least not one with the sub-carrier-period accuracy required for distributed beamforming. The presence of an ideal common time reference significantly simplifies the distributed beamforming problem. All that is required in this case is for the destination to transmit a sinusoidal beacon with known initial phase to the sources. Upon reception of this beacon, each source can directly estimate the phase of its channel. Distributed beamforming can then be achieved simply by transmitting modulated carriers back to the destination with appropriate phase precompensation at each source. As long as the channels remain time invariant, no resynchronization is necessary since the sources can transmit with identical carrier frequencies derived from the common time reference.

The absence of a common time reference implies that each source keeps local time using its own independent local oscillator. We assume no *a priori* phase or frequency synchronization among these oscillators. Since none of the nodes in the system know the "true" time, they also do not know the "true" frequency or phase of their local oscillator. This implies that nodes cannot generate absolute phase or frequency estimates with respect to "true" time. Phase and frequency estimates obtained by node j must be considered relative to node j's local oscillator frequency and phase. It also implies that carrier frequency synchronization is necessary to maintain coherent combining during beamforming. Even if the channels are noiseless and the nodes remain at fixed positions, some amount of frequency synchronization error will be unavoidable between the sources due to phase noise and oscillator inaccuracy. Periodic resynchronization is necessary to avoid unacceptable phase drift during beamforming.

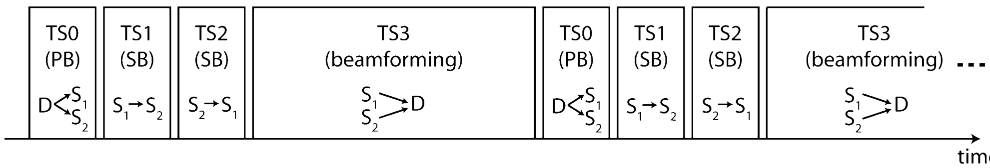

Fig. 2. Summary of the two-source time-slotted round-trip carrier synchronization protocol where PB and SB denote primary and secondary beacon synchronization timeslots, respectively.

Finally, we assume here that the nodes do not have knowledge of their precise position, either in an absolute sense or relative to each other. Relative position information could be used to enable distributed beamforming in systems with single-path channels since each source could, as in [10], use its position information to estimate its propagation delay to the destination and employ carrier phase precompensation to cancel the phase shift of the channel. While position estimates may be obtained via systems such as GPS, the amount of inaccuracy in these estimates is typically much greater than a carrier wavelength, making them impractical for accurate phase precompensation in distributed beamforming.

III. Time-Slotted Synchronization Protocol

We first develop the concept of time-slotted round-trip carrier synchronization in detail for the $M = 2$ source case in Section III-A and describe its implementation in multipath channels in Section III-B. The case with $M > 2$ sources is discussed in Section III-C. 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 here:

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.

Fig. 2 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.

A. Two-Source Synchronization in Single-Path Time-Invariant Channels

This section describes the implementation of the proposed time-slotted round-trip carrier synchronization protocol for a system with two sources under the assumption that all of the channels in Fig. 1 are single-path and time-invariant, i.e., $g_{i,k}(t) = \alpha_{ik}\delta(t - \tau_{ik}) \forall ik$. The basic intuition behind the round-trip protocol in this case is that the propagation times of the $\mathsf{D} \rightarrow \mathsf{S}_1 \rightarrow \mathsf{S}_2 \rightarrow \mathsf{D}$ and the $\mathsf{D} \rightarrow \mathsf{S}_2 \rightarrow \mathsf{S}_1 \rightarrow \mathsf{D}$ 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 essentially 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

$$x_0(t) = e^{j(\omega(t-t_0)+\phi_0)} \qquad t \in [t_0, t_0 + T_0). \tag{1}$$

The signal received at S_i in TS_0 can be written as

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

for $t \in [t_0 + \tau_{0i}, t_0 + \tau_{0i} + T_0)$ where $\eta_{0i}(t)$ denotes the additive white Gaussian noise (AWGN) in the $0 \rightarrow i$ channel and $i \in \{1, 2\}$. Each source uses its noisy observation from the first timeslot to compute estimates of the received frequency and phase; these estimates are 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 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 TS_1 can then be written as

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

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

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

for $t \in [t_1+\tau_{12}, t_1+\tau_{12}+T_1)$ where $\eta_{12}(t)$ denotes the AWGN in the $1 \to 2$ channel. From this noisy observation, S_2 generates estimates of the received frequency and phase; these estimates 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 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 TS_0 even though its transmission begins at the conclusion of the secondary beacon received in TS_1. 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

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

The secondary beacon transmitted by S_2 in TS_2 can then be written as

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

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

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

for $t \in [t_2+\tau_{12}, t_2+\tau_{12}+T_2)$ where $\eta_{21}(t)$ denotes the AWGN in the $2 \to 1$ 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 $\hat{\omega}_{21}$ and $\hat{\phi}_{21}$ of the received frequency and phase, respectively.

In timeslot TS_3, 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 the performance of the distributed beamformer is primarily affected by the phase offset between the carriers at the destination, we will 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

$$x_{i0}(t) = a_{i0} e^{j(\hat{\omega}_i(t-t_{3i})+\hat{\phi}_{3i})} \quad t \in [t_{3i}, t_{3i} + T_3) \tag{2}$$

where $\hat{\omega}_i$ is a frequency estimate at S_i that, as discussed in Section IV.B, may be a function of both $\hat{\omega}_{0i}$ and $\hat{\omega}_{ki}, k \neq i$. 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

$$\hat{\phi}_{31} = \hat{\phi}_{21} + \hat{\omega}_{21}(t_{31} - (t_2 + \tau_{12})) \quad \text{and} \tag{3}$$
$$\hat{\phi}_{32} = \hat{\phi}_{12} + \hat{\omega}_{12}(t_{32} - (t_1 + \tau_{12})) \tag{4}$$

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, and hence

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

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 begin its carrier transmission in timeslot TS_3 at time $t_{32} = t_2 + T_2 + \tau_{\text{delay}}$, where $\tau_{\text{delay}} = \tau_{01} + \tau_{12} - \tau_{02}$. By inspection of Fig. 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 TS_0 to the start of its secondary beacon observation in TS_1, i.e., $\tau_{\text{delay}} = (t_1 + \tau_{12}) - (t_0 + \tau_{02} + T_0)$.

The signal received at D in TS_3 can be written as

$$y_0(t) = \sum_{i=1}^{2} \alpha_{0i} a_{i0} e^{j(\hat{\omega}_i(t-t_3)+\hat{\phi}_{3i})} + \eta_0(t)$$

for $t \in [t_3, t_3+T_3)$ where $t_3 = t_{31}+\tau_{01} = t_{32}+\tau_{02}$ and where we have applied the assumption that $\tau_{0i} = \tau_{i0}$ and $\alpha_{0i} = \alpha_{i0}$ for $i \in \{1, 2\}$. Standard trigonometric identities can be applied to rewrite $y_0(t)$ as

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

where

$$\begin{aligned} a_{bf}(t) &:= \sqrt{a^2 + b^2 + 2ab\cos(\phi_\Delta(t))} \\ \phi_{bf}(t) &:= \hat{\omega}_1(t - t_3) + \hat{\phi}_{31} + \tan^{-1} \\ &\quad \times \left(\frac{b \sin(\phi_\Delta(t))}{a + b\cos(\phi_\Delta(t))} \right) \end{aligned}$$

with $a := \alpha_{01} a_{10}$ and $b := \alpha_{02} a_{20}$ and where we have defined the carrier phase offset

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

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

$$a_{bf}(t) = 2a \cos\left(\frac{\phi_\Delta(t)}{2}\right)$$
and
$$\phi_{bf}(t) = \frac{(\hat{\omega}_1 + \hat{\omega}_2)(t - t_3) + \hat{\phi}_{31} + \hat{\phi}_{32}}{2}.$$

B. Two-Source Synchronization in Multipath Time-Invariant Channels

Since the channels in Fig. 1 are accessed at the same frequency in both forward and reverse directions, the time-slotted round-trip synchronization protocol described in the previous section for single-path LTI channels can also be effective in communication systems with multipath LTI channels if minor modifications are made to the protocol to account for the transient effects of the channels. This section summarizes the necessary modifications and discusses the implications of multipath on the phase and frequency estimates at each source.

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). Unlike the case with single-path channels, however, the primary beacon arrives at each source after propagation through multiple paths with differing delays and amplitudes. Denoting τ_{0i} as the delay of the shortest path and ν_{0i} as the finite delay spread of the channel $g_{0,i}(t)$, S_i observes the initial transient response of the multipath channel $g_{0,i}$ to the primary beacon over the interval $t \in [t_0+\tau_{0i}, t_0+\tau_{0i}+\nu_{0i})$. 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_i can then be written as

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

for $t \in [t_0+\tau_{0i}+\nu_{0i}, t_0+\tau_{0i}+T_0)$ where α_{0i} and θ_{0i} are the steady-state gain and phase response of the LTI channel, respectively, for $i \in \{1,2\}$. At the conclusion of the primary beacon, a final transient component is also received by S_i during $t \in [t_0+\tau_{0i}+T_0, t_0+\tau_{0i}+\nu_{0i}+T_0)$. The important thing here is that each source uses *only the steady-state portion* of its noisy observation in the first timeslot to compute 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 as periodic extensions of the steady state portion of the primary beacon observations.

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 only differences with respect to the single path case are that (i) the duration of each secondary beacon must exceed the delay spread $\nu_{12} = \nu_{21}$ 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. General M-Source Synchronization in Time-Invariant Channels

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 $\mathsf{TS}_0, \dots, \mathsf{TS}_{2M-1}$. The first $2M-1$ timeslots are used for the transmission of synchronization beacons and the final timeslot is used for beamforming. The basic concepts of two-source synchronization apply here with additional synchronization timeslots and minor modifications in the calculation of the transmission phase for source nodes $\mathsf{S}_2, \dots, \mathsf{S}_{M-1}$ shown in Fig. 1. The activity in each timeslot is summarized here:

1) In 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 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 TS_i is a periodic extension of the beacon received in TS_{i-1}. S_{i+1} generates local phase and frequency estimates from this observation.
3) In 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 TS_0, with initial phase extrapolated from the phase and frequency estimates obtained by S_M in TS_0. S_{M-1} generates local phase and frequency estimates from this observation.
4) In 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 TS_i is a periodic extension of the secondary beacon received in TS_{i-1}. S_{2M-i-1} generates local phase and frequency estimates from this observation.
5) In 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 $\mathsf{D} \rightarrow \mathsf{S}_1 \rightarrow \mathsf{S}_2 \rightarrow \cdots \rightarrow \mathsf{S}_M \rightarrow \mathsf{D}$ and the $\mathsf{D} \rightarrow \mathsf{S}_M \rightarrow \mathsf{S}_{M-1} \cdots \rightarrow \mathsf{S}_1 \rightarrow \mathsf{D}$ 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 III.A. The only difference is that the secondary beacons propagate between S_1 and S_M with multiple hops when $M > 2$, rather than via direct propagation in the two-source case. When $M > 2$, however, nodes $\mathsf{S}_2, \dots, \mathsf{S}_{M-1}$ must also derive appropriate transmission phases to participate in the distributed beamformer.

Under the assumption that all of the channels in Fig. 1 are time-invariant, and 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) of

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

where $\theta_{k,i}$ denotes the phase of the LTI single-path channel $g_{k,i}(t)$. Let $\mathcal{S}$ denote the set of source nodes S_m for $m \in \{2, \dots, M-1\}$. In order for source node $\mathsf{S}_m \in \mathcal{S}$ 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 $\theta^{\mathsf{rt}} - \theta_{m,0}$.

Source node $\mathsf{S}_m \in \mathcal{S}$ receives three transmissions during the synchronization phase of the protocol: a primary beacon in TS_0 at phase $\theta_{0,m} = \theta_{m,0}$, a secondary beacon during the counterclockwise[2] propagation of beacons in TS_{M-1} at phase $\theta_m^{\downarrow} = \theta_{0,1} + \theta_{1,2} + \cdots + \theta_{m-1,m}$ and another secondary beacon during the clockwise propagation of beacons in TS_{2M-m-1} at phase $\theta_m^{\uparrow} = \theta_{0,M} + \theta_{M,M-1} + \cdots + \theta_{m+1,m}$. Since each node in the system estimates the phase of received beacons relative to its own local time reference, absolute estimates of $\theta_m^{\uparrow}$ and $\theta_m^{\downarrow}$ 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 as

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

Since the unknown local phase offset has been canceled in the phase differences $\delta_m^{\uparrow}$ and $\delta_m^{\downarrow}$, 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

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

which is the desired phase for beamforming since $\theta_{0,m} = \theta_{m,0}$. After propagation through channel $g_{m,0}(t)$, the carrier from $\mathsf{S}_m \in \mathcal{S}$ arrives at the destination with phase θ^{rt} and constructively combines with the carriers from S_1 and S_M.

In the round-trip synchronization protocol for $M > 2$ sources, we note that the secondary beacons transmitted by S_m are only used by S_{m+1} or S_{M-1}, depending on the timeslot in which the beacon is transmitted. The other sources in the system (and the destination) ignore these transmissions. While more efficient protocols exploiting the broadcast nature of the wireless links between the sources and requiring less than $2M-1$ synchronization timeslots may be possible, the round-trip synchronization protocol possesses the property that the number of synchronization timeslots is linear in M.

D. Discussion

Although the events of the time-slotted round-trip synchronization protocol are described in terms of some notion of "true time" t, it is worth reiterating 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 $\mathsf{TS}_1, \dots, \mathsf{TS}_{M-1}$, each source transmission is simply a periodic extension of a beacon received in a previous timeslot. In the case of $M > 2$ sources, source nodes $\mathsf{S}_2, \dots, \mathsf{S}_{M-1}$ must also compute a phase offset for their carrier to arrive at the correct phase, but the carriers of these source nodes are each transmitted as a periodic extension of the primary beacon received in TS_0 with a phase shift computed from two local phase differences. No absolute notion of "time-zero" is needed since the phase of each source transmission is extrapolated from the estimated initial phase of the appropriate beacon observation in a previous timeslot. Moreover, each source transmission in timeslots $\mathsf{TS}_1, \dots, \mathsf{TS}_{M-1}$ is triggered by the conclusion of a beacon in a prior timeslot. The sources do not follow any schedule requiring knowledge of "true time."

Since each source transmission in the time-slotted round-trip carrier synchronization protocol is intended to be a periodic extension of a beacon received in a previous timeslot, we note that each source could realize its phase and frequency estimation functions during the synchronization timeslots by using PLLs with holdover circuits [11]. A PLL implementation of two-source round-trip carrier synchronization with continuously transmitted beacons is described in [8]. Holdover circuits are not required in [8] since the beacons are continuously transmitted and each PLL is able to constantly track the phase of a received beacon. Since beacons are transmitted for only finite durations in the time-slotted protocol, holdover circuits are necessary to prevent tracking of incorrect beacons or noise. Despite this difference, a PLL implementation for the time-slotted round-trip carrier synchronization protocol would be similar to that in [8] with the addition of the holdover circuits and their associated logic. Like [8], we note that each source requires as many PLLs as received beacons in order to "store" all of its local phase and frequency estimates. In the time-slotted synchronization protocol with $M > 2$, this implies that a source may require as many as three PLLs, depending on whether or not it is an "end node," i.e., S_1 or S_M, or an "interior node", i.e., $\mathsf{S}_2, \dots, \mathsf{S}_{M-1}$.

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 $\mathsf{TS}_1, \dots, \mathsf{TS}_{M-1}$ 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 for the case when $M = 2$ sources, the estimate of τ_{delay} at S_2 in TS_3 is not critical if the beamforming timeslot is sufficiently long. An inaccurate estimate of τ_{delay} only causes S_2's carrier to arrive slightly earlier or later than S_1's carrier at D; it does not affect the relative phase of the carriers during beamforming. The same intuition applies to the case of $M > 2$ sources.

As a final point, the indirect effect of gaps between the timeslots is that the extrapolated phase estimates will tend to become less accurate as the duration of the gaps increases. This is due

[2]In the context of Fig. 1, $\mathsf{S}_1 \rightarrow \mathsf{S}_2 \rightarrow \cdots \rightarrow \mathsf{S}_M$ is counterclockwise propagation and $\mathsf{S}_M \rightarrow \mathsf{S}_{M-1} \cdots \rightarrow \mathsf{S}_1$ is clockwise propagation around the circuit including D.

to the fact that the gaps cause more time to elapse between the observation and the start of the periodic extension transmission, magnifying the effect of frequency estimation error and/or phase noise. Hence, while the phase offset of the carriers at the destination during beamforming is not directly affected by the absolute starting and ending times of each timeslot, better performance can be achieved when any gaps between the timeslots are minimized.

IV. Two-Source Distributed Beamforming Performance in Time-Invariant Channels

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. We focus here on the two-source scenario and note that the analysis for this case can also be applied to the analysis of phase offsets in the $M > 2$ source scenario.

In an ideal two-antenna beamformer, the received signals add constructively at the destination and the resulting amplitude is $a_{bf}(t) = \alpha_{01}a_{10} + \alpha_{02}a_{20}$. The non-ideal nature of the distributed beamformer is captured in the carrier phase offset

$$\phi_\Delta(t) = \omega_\Delta(t - t_3) + \phi_\Delta + \chi_1(t) - \chi_2(t) \tag{6}$$

for $t \in [t_3, t_3 + T_3)$ where $\omega_\Delta := \hat{\omega}_2 - \hat{\omega}_1$ represents the linear phase drift during beamforming, $\phi_\Delta := \hat{\phi}_{32} - \hat{\phi}_{31}$ represents the initial phase offset at the start of beamforming, and $\chi_1(t) - \chi_2(t)$ represents the difference in the phase noise processes of the local oscillators of S_1 and S_2. Phase and frequency estimation errors at each source result in unavoidable initial carrier phase offset at the start of TS_3 as well as linear phase drift over the duration of TS_3. Moreover, the phase noise in each source's local oscillator causes each source's phase to randomly wander from the intended phase during beamforming, even in the absence of estimation error. The following section establishes a vector notation for the eight estimation errors in the two-source time-slotted round-trip carrier synchronization protocol. We then analyze the joint statistics of these estimation errors as well as the error due to phase noise in order to facilitate analysis of the carrier phase offset during beamforming. Our analysis assumes that the synchronization timeslots are sufficiently short such that the phase noise is negligible during synchronization and does not appreciably affect the accuracy of the phase and frequency estimates of the beacons received at each source.

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

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

where

$$\begin{aligned}\tilde{\omega}_{0j} &:= \hat{\omega}_{0j} - \omega \quad j \in \{1, 2\}\\ \tilde{\omega}_{ij} &:= \hat{\omega}_{ij} - \hat{\omega}_{0i} \quad ij \in \{12, 21\}\\ \tilde{\phi}_{0j} &:= \hat{\phi}_{0j} - \phi_0 \quad j \in \{1, 2\}\end{aligned}$$

and

$$\tilde{\phi}_{ij} := \hat{\phi}_{ij} - \hat{\phi}_i \quad ij \in \{12, 21\}.$$

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 D. 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 S_i.

To facilitate analysis, we assume that the frequency and phase estimates at both sources are unbiased, i.e., $\mathrm{E}\{\tilde{\boldsymbol{\theta}}\} = 0$. Denoting the covariance matrix of the estimation error vector as $\boldsymbol{\Theta} := \mathrm{E}\{\tilde{\boldsymbol{\theta}}\tilde{\boldsymbol{\theta}}^\top\}$, we note that almost all of the off-diagonal terms in this matrix are equal to zero since observations in different timeslots are affected by independent noise realizations and observations at S_1 and S_2 are affected by independent noise realizations. The frequency and phase estimates obtained from the same observation at a particular source, however, are not independent. Their covariance is denoted as $c_{ki} := \mathrm{cov}\,[\tilde{\omega}_{ki}\tilde{\phi}_{ki}]$.

It is possible to lower-bound the non-zero elements in $\boldsymbol{\Theta}$ with the Cramer-Rao bound (CRB). Given an N-sample observation of a complex exponential of amplitude a sampled at rate f_s, the CRBs for the variances and covariance of the frequency and phase estimates are given as [12]

$$\sigma_\omega^2 \geq \frac{12\sigma^2 f_s^2}{a^2 N(N^2 - 1)} \tag{7}$$

$$\sigma_\phi^2 \geq \frac{2\sigma^2(2N - 1)}{a^2 N(N + 1)} \tag{8}$$

and

$$\mathrm{cov}\{\tilde{\omega}, \tilde{\phi}\} \geq \frac{-6\sigma^2 f_s}{a^2 N(N + 1)} \tag{9}$$

where σ^2 is the variance of the uncorrelated real and imaginary components of the independent, identically distributed, zero-mean, complex Gaussian noise samples.

B. Carrier Frequency Offset During Beamforming

Since the carrier transmissions of each source in TS_3 are periodic extensions of the secondary beacons, it is reasonable to expect that S_1 and S_2 transmit their carriers in TS_3 at the secondary beacon frequency estimates $\hat{\omega}_{21}$ and $\hat{\omega}_{12}$, respectively. At the start of beamforming, however, each source actually possesses a pair of frequency estimates that could be used for beamforming: one from the primary beacon observation and one from the secondary beacon observation. Since both frequency estimates are unbiased, the source could simply select the estimate with the lower variance for the carrier transmission in TS_3. An even better approach is to generate the carrier at S_i from a linear combination of the local estimates, i.e.,

$$\hat{\omega}_i = \mu_i \hat{\omega}_{0i} + (1 - \mu_i)\hat{\omega}_{ki}.$$

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

$$\omega_\Delta = \boldsymbol{\Gamma}_1^\top \tilde{\boldsymbol{\theta}} \tag{10}$$

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

If the frequency estimates are unbiased and Gaussian distributed, the carrier frequency offset ω_Δ is Gaussian distributed with zero mean for any fixed μ_1 and μ_2. A good choice then for the linear combination parameters is one that minimizes $\mathrm{var}\{\omega_\Delta\}$ where

$$\begin{aligned}\mathrm{var}\{\omega_\Delta\} &= \mathbf{\Gamma}_1^\top \mathrm{E}[\tilde{\boldsymbol{\theta}}\tilde{\boldsymbol{\theta}}^\top]\mathbf{\Gamma}_1 \\ &= (1-\mu_1-\mu_2)^2\left(\sigma^2_{\omega_{01}}+\sigma^2_{\omega_{02}}\right) \\ &\quad + (1-\mu_2)^2\sigma^2_{\omega_{12}} + (1-\mu_1)^2\sigma^2_{\omega_{21}}.\end{aligned}$$

The linear combination parameters that minimize $\mathrm{var}\{\omega_\Delta\}$ can be determined using standard calculus techniques to be

$$\mu_1^* = \frac{1}{1+\frac{\sigma^2_{\omega_{12}}}{\sigma^2_{\omega_{21}}}\left(\frac{\sigma^2_{\omega_{01}}+\sigma^2_{\omega_{02}}}{\sigma^2_{\omega_{01}}+\sigma^2_{\omega_{02}}+\sigma^2_{\omega_{12}}}\right)}$$
$$\text{and}$$
$$\mu_2^* = \frac{1}{1+\frac{\sigma^2_{\omega_{21}}}{\sigma^2_{\omega_{12}}}\left(\frac{\sigma^2_{\omega_{01}}+\sigma^2_{\omega_{02}}}{\sigma^2_{\omega_{01}}+\sigma^2_{\omega_{02}}+\sigma^2_{\omega_{21}}}\right)}.$$

If each source is able to obtain high quality frequency estimates through, for example, properly designed PLLs or maximum likelihood estimators, the CRB result (7) can be used to compute the $\sigma^2_{\omega_{ki}}$ terms required for computation of μ_i. This computation, while simple, does however require knowledge of the signal to noise ratio $(a^2/2\sigma^2)$ for each channel of interest. If this information is not available but the durations of the secondary beacons are equal, i.e., $T_1 = T_2$, the reciprocal nature of $g_{1,2}(t)$ and $g_{2,1}(t)$ implies that $\sigma^2_{\omega_{12}} = \sigma^2_{\omega_{21}}$. A reasonable approach then could be to assume that the intrasource channels are of better quality than the source-destination channels, i.e.,

$$\sigma^2_{\omega_{12}} = \sigma^2_{\omega_{21}} \ll \sigma^2_{\omega_{01}} + \sigma^2_{\omega_{02}},$$

in which case $\mu_1^* = \mu_2^* = 1/2$.

C. Carrier Phase Offset at the Start of Beamforming

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

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

The frequency estimates $\hat{\omega}_{21}$ and $\hat{\omega}_{12}$ are generated from the secondary beacon observations at S_1 and S_2, respectively, and can be written as

$$\begin{aligned}\hat{\omega}_{21} &= \hat{\omega}_{02} + \tilde{\omega}_{21} = \omega + \tilde{\omega}_{02} + \tilde{\omega}_{21} \\ \hat{\omega}_{12} &= \hat{\omega}_{01} + \tilde{\omega}_{12} = \omega + \tilde{\omega}_{01} + \tilde{\omega}_{12}.\end{aligned}$$

From the protocol description in Section III, the phase estimate $\hat{\phi}_{12}$ can be written as

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

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 at S_2. Similarly, the phase estimate $\hat{\phi}_{21}$ can be written as

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

where $\Psi := t_{32} - (t_1+\tau_{12}) = \tau_{01}+\tau_{12}-\tau_{02}+T_1+T_2$. Since $t_{31}-(t_2+\tau_{12}) = T_2$, we can write

$$\phi_\Delta = \mathbf{\Gamma}_2^\top \tilde{\boldsymbol{\theta}} \tag{11}$$

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

D. Effect of Phase Noise During Beamforming

In addition to phase and frequency offsets that occur as a consequence of imperfect estimation, practical oscillators also exhibit phase noise. Oscillator phase noise can cause the phase of each carrier to randomly wander from the intended phase during beamforming and can establish a ceiling on the reliable beamforming time, i.e., the time at which the phase offset between the carriers drifts beyond an acceptable theshold, even in the absence of estimation error.

We model the oscillator phase noise $\chi_i(t)$ at S_i as a non-stationary Gaussian random process with zero mean and variance increasing linearly with time from the start of beamforming, i.e., $\sigma^2_{\chi_i}(t) = r(t-t_3)$ for $t \in [t_3, t_3+T_3)$, where the variance parameter r is a function of the physical properties of the oscillator including its natural frequency and physical type, e.g., Colpitts [13]. We assume that $\chi_1(t)$ and $\chi_2(t)$ share the same value of r but are independent phase noise processes. As discussed in [13], typical values of r for low-cost radio-frequency oscillators range from 1 to 20 $\mathrm{rad}^2\cdot$ Hz.

E. Statistics of the Phase Offset During Beamforming

Substituting (10) and (11) in (6), we can compactly express the phase offset during beamforming in terms of the estimation error vector and phase noise processes as

$$\phi_\Delta(t) = [(t-t_3)\mathbf{\Gamma}_1 + \mathbf{\Gamma}_2]^\top \tilde{\boldsymbol{\theta}} + \chi_1(t) - \chi_2(t)$$

for $t \in [t_3, t_3+T_3)$. Note that $\mathbf{\Gamma}(t-t_3) := (t-t_3)\mathbf{\Gamma}_1 + \mathbf{\Gamma}_2$ is deterministic. Under the assumption that the estimates are unbiased, Gaussian distributed, and independent of the mutually independent phase noise processes at each source, we can say that $\phi_\Delta(t) \sim \mathcal{N}(0, \sigma^2_{\phi_\Delta(t)})$ where

$$\sigma^2_{\phi_\Delta(t)} = \mathbf{\Gamma}^\top(t-t_3)\mathbf{\Theta}\mathbf{\Gamma}(t-t_3) + 2r(t-t_3) \tag{12}$$

at 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

$$\mathrm{Prob}[|\phi_\Delta(t)| < \lambda] = 1 - 2Q\left(\frac{\lambda}{\sigma_{\phi_\Delta}(t)}\right) \tag{13}$$

where $Q(x) = (1/\sqrt{2\pi})\int_x^\infty \exp\,(-t^2/2)\,dt$. The CRB can also be used to provide a lower bound on the variance of the phase offset during beamforming and, as such, an upper bound on the probability that the absolute carrier phase offset is less than a given threshold.

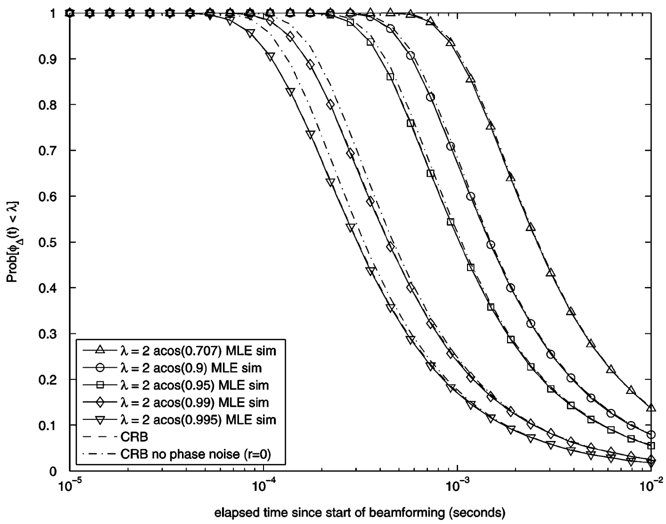

Fig. 3. $\mathrm{Prob}[|\phi_\Delta(t)| < \lambda]$ versus beamforming time for the case when $T_0 = 1\ \mu$s and $T_1 = T_2 = 2\ \mu$s.

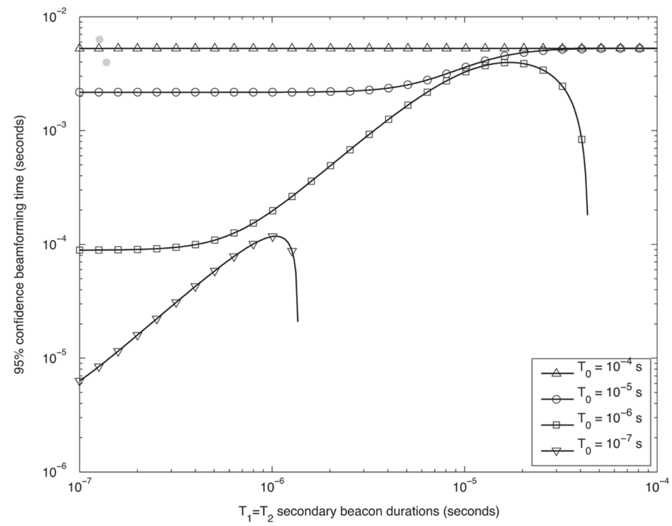

Fig. 4. 95% confidence beamforming time as a function of the primary beacon duration T_0 and secondary beacon durations $T_1 = T_2$ for a 90%-ideal beamforming quality threshold ($\lambda = 2\cos^{-1}(0.9)$).

F. Numerical Examples

This section presents numerical examples of the two-source time-slotted round-trip carrier synchronization protocol in single-path time-invariant channels. The examples in this section assume that all beacons are received at an SNR of $10\log_{10}((a^2/2\sigma^2)) = 20$ dB and that each channel has a random propagation delay. The primary beacon frequency is $\omega = 2\pi \cdot 900 \cdot 10^6$ rad/second and the oscillator phase noise variance parameter is assumed to be $r = 20\ \mathrm{rad}^2\cdot$ Hz.

The first example shows how the beamforming quality degrades as the duration of the beamforming interval TS_3 increases (for fixed beacon durations) and gives a sense of how often the sources will require resynchronization to maintain acceptable beamforming quality. Fig. 3 plots $\mathrm{Prob}[|\phi_\Delta(t)| < \lambda]$ versus elapsed time from the start of beamforming 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 carrier frequencies for beamforming with the optimum linear combining factors $\mu_1^* = \mu_2^* \approx 0.5152$ in this case. A Monte Carlo simulation with 50000 iterations was performed to obtain estimates of $\mathrm{Prob}[|\phi_\Delta(t)| < \lambda]$ where, for each new realization of the random parameters, both sources generate maximum likelihood phase and frequency estimates of their noisy observations and use these estimates to generate periodic extensions in the appropriate timeslots. The quality thresholds in Fig. 3 represent the ratio of received power of the distributed beamformer to that of an ideal beamformer in the sense that $\lambda = 2\cos^{-1}(u)$ and $|\phi_\Delta(t)| < \lambda$ implies that the received power of the distributed beamformer at time t is no worse than u^2 times that of an ideal beamformer, where $0 < u < 1$. The value of $u = (1/\sqrt{2})$ represents the "break-even" case where the distributed beamformer has the same power efficiency as orthogonal transmission, i.e., the carriers are received with 90 degrees of phase offset.

The results in Fig. 3 show that the carrier phases are closely aligned at the destination with high probability up to $t - t_3 \approx 50\ \mu$s. By $t - t_3 = 3$ms, however, the probability of having carrier phase alignment such that beamforming is more power efficient than orthogonal transmission is less than $1/2$. Hence, depending on the quality threshold and the confidence in which the threshold must be satisfied, these results show that the distributed beamformer must be periodically resynchronized in order to maintain an acceptable level of performance with high confidence.

In addition to the Monte Carlo simulations, Fig. 3 also plots an upper bound for $\mathrm{Prob}[|\phi_\Delta(t)| < \lambda]$ using the CRB results in Section IV.A to lower bound (12). The close match of the experimental and analytical results shows that, as might be expected (see, e.g., [14]), the CRB can be used to efficiently predict the performance of the two-source time-slotted round-trip carrier synchronization system when the sources use maximum likelihood phase and frequency estimation. The CRB without phase noise, i.e., $r = 0$, is plotted to also show the effect of phase noise on the expected beamforming time. In these results, since the synchronization timeslots are short, the phase drift due to the phase and frequency estimation errors tends to dominate the effect of the phase noise.

The second example considers the effect of the beacon durations on how long acceptable beamforming quality can be maintained at a desired level of confidence. Fig. 4 plots the 95% confidence beamforming time given a 90%-ideal beamforming quality threshold ($\lambda = 2\cos^{-1}(0.9)$ and $\mathrm{Prob}\ [|\phi_\Delta(t)| < \lambda] = 0.95$) using the CRB analytical predictions. All other parameters are identical to Fig. 3. The results in Fig. 4 show two regimes of operation. In the short-beacon regime, the phase and frequency estimation errors dominate the phase noise. In this regime, the 95% confidence beamforming times are approximately flat with respect to the secondary beacon duration 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

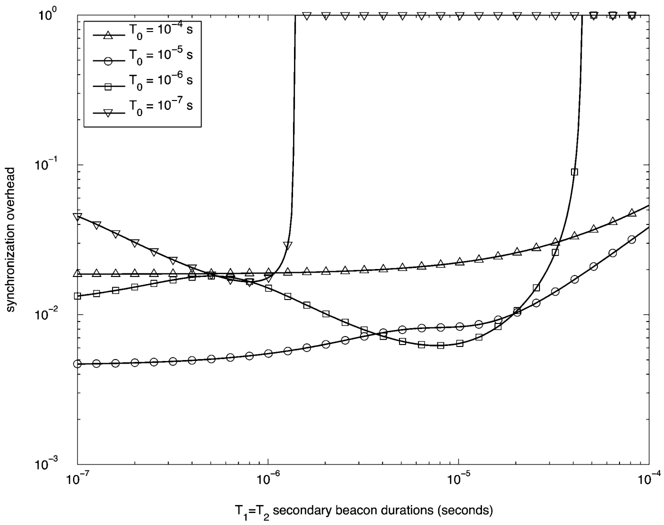

Fig. 5. Synchronization overhead as a function of the primary beacon duration T_0 and secondary beacon durations $T_1 = T_2$ for a 90%-ideal beamforming quality threshold at 95% confidence ($\lambda = 2\cos^{-1}(0.9)$ and $\mathrm{Prob}[|\phi_\Delta(t)| < \lambda] = 0.95$).

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.

In the long-beacon regime, the oscillator phase noise dominates the phase and frequency estimation errors. When operating in this regime, the 95% confidence beamforming time is flat with respect to the primary and secondary beacon durations. In Fig. 4, the phase noise establishes a 95% confidence beamforming time ceiling at approximately 5 ms. This ceiling is achieved when $T_0 = 10^{-4}$ or when $T_0 = 10^{-5}$ and $T_1 = T_2 > 2 \cdot 10^{-5}$. Since there is nothing to gain from increasing beacon durations when operating in this regime, these results intuitively suggest that beacon durations should be selected to be long enough such that the estimation error does not dominate but also short enough such that phase noise does not dominate.

Fig. 5 plots the synchronization overhead of the two-source round-trip protocol under the same conditions as in Fig. 4. We define synchronization overhead as the fraction of total time spent synchronizing the sources ignoring propagation delays, i.e.

$$\text{synchronization overhead} := \frac{T_0 + T_1 + T_2}{T_0 + T_1 + T_2 + T_3}.$$

The beacon durations T_0, T_1, and T_2 are design parameters, but the beamforming duration T_3 is computed from the CRB analytical predictions assuming a a 90%-ideal beamforming quality threshold ($\lambda = 2\cos^{-1}(0.9)$) at 95% confidence. The results in Fig. 5 show that a synchronization overhead of less than 1% can be achieved by selecting beacon durations such as $T_0 = 10^{-5}$ and $T_1 = T_2 = 10^{-7}$ or $T_0 = 10^{-6}$ and $T_1 = T_2 = 8 \cdot 10^{-6}$. In both cases, the beacon durations are short enough to avoid operation in the phase-noise dominated regime but also long enough to avoid excessive estimation error. An important conclusion from these results is that the synchronization overhead can be small when the beacon durations are selected such that the phase drift due to estimation error is balanced with the phase noise.

V. Effect of Mobility

In this section, we consider the question of how mobility affects the performance of the two-source time-slotted round-trip carrier synchronization system. Since each of the three channels in the system are accessed in both directions at different times, mobility may lead to phase and frequency offset at the destination during beamforming even in the absence of phase and frequency estimation error at the sources. While our focus here is on the scenario with $M = 2$ sources, the ideas developed in this section also apply to the $M > 2$ source scenario.

In order to isolate the effect of mobility on the two-source distributed beamformer, we consider a system with noiseless single-path channels and perfect oscillators, i.e., no phase noise. In this scenario, the sources are able to perfectly estimate the phase and frequency of their primary and secondary beacon observations. We assume a constant velocity mobility model where, in D's reference frame, S_1 and S_2 are moving with constant velocity $v_i \leq v$ for $i \in \{1, 2\}$, where v is the maximum source velocity and $v/c \ll 1$. While this constant bounded velocity assumption may appear to be somewhat restrictive, it is reasonable in the regime where the synchronization timeslots are short. Any acceleration that occurs in a timeslot will result in a relatively small velocity change if the timeslot is short.

A. Initial Carrier Phase Offset Due to Mobility

Under the notation established in Section III, S_1, and S_2 begin transmission of their carriers in TS_3 at times t_{31} and t_{32}, respectively. The resulting arrival time difference at the destination can be written as

$$\begin{aligned}\tau_\Delta &= t_{31} + \tau_{10}(t_{31}) - (t_{32} + \tau_{20}(t_{32})) \\ &= \tau_{10}(t_{31}) - \tau_{01}(t_0 + T_0) + \tau_{21}(t_2 + T_2) \\ &\quad - \tau_{12}(t_1) + \tau_{02}(t_0 + T_0) - \tau_{20}(t_{32})\end{aligned}$$

where $\tau_{ki}(t)$ denotes the propagation time of an impulse emitted from transmitter k to receiver i at time t. If the carriers from S_1 and S_2 are both transmitted at frequency ω, the arrival time difference results in a carrier phase offset at the start of beamforming of $\phi_\Delta = \omega\tau_\Delta$. Applying the bound for $|\tau_\Delta|$ developed in the Appendix, we can bound the initial carrier phase offset as

$$|\phi_\Delta| \leq \omega\beta[9t_p + 8T_1] \tag{14}$$

where $\beta := v/c \ll 1$, t_p denotes the maximum propagation delay over all links in the system, and where we have assumed that the secondary beacon durations are identical ($T_1 = T_2$).

As a numerical example of the bound, consider a system with a 900 MHz carrier frequency and secondary beacon durations $T_1 = T_2 = 2\,\mu$s. Suppose the maximum velocity of each source is set to $v = 100$ meters/sec and the maximum propagation time is set to $t_p = 1\ \mu$s, corresponding to a maximum distance between any two nodes of 300 meters. From (14), the maximum initial carrier phase offset can be computed to be $|\phi_\Delta| \leq 0.0471$ radians, or about 2.7 degrees.

B. Frequency Offset Due to Mobility

In addition to an initial carrier phase offset ϕ_Δ at the start of beamforming, mobility may also result in frequency offset ω_Δ during beamforming due to Doppler shifts [15] in each link. Before analyzing the behavior of the round-trip system, we first consider an exchange of short sinusoidal beacons between two nodes denoted by A and B. Suppose B transmits the first beacon. In A's reference frame, B is moving at constant velocity v_B and is separating from A at the rate $v_B \cos\theta_{BA}$ where θ_{BA} is the angle of B's motion with respect to the line from A to B. When a beacon is transmitted from A to B, in B's reference frame, the geometry is the same as before except that now A is moving at constant velocity $v_A = v_B$ and is separating from B at the rate $v_A \cos\theta_{AB}$ where θ_{AB} is defined in the same way as θ_{BA}. We can write the ratio of the Doppler shifts as

$$\rho := \frac{\text{Doppler}(\mathsf{B} \to \mathsf{A})}{\text{Doppler}(\mathsf{A} \to \mathsf{B})} = \frac{1 + \frac{v_B}{c}\cos\theta_{BA}}{1 + \frac{v_A}{c}\cos\theta_{AB}}$$

Since the beacons between A and B were not transmitted simultaneously, it cannot be assumed that $\theta_{AB} = \theta_{BA}$. Nevertheless, when the elapsed time between transmissions is short, we can expect any difference in the angles to be small. Suppose that we can bound the difference of the angles such that $|\theta_{BA} - \theta_{AB}| \le \epsilon$. Then we can bound the range of ρ as

$$\begin{aligned}\rho &\in \left[\frac{1 - \frac{v}{c}\sin(\epsilon/2)}{1 + \frac{v}{c}\sin(\epsilon/2)}, \frac{1 + \frac{v}{c}\sin(\epsilon/2)}{1 - \frac{v}{c}\sin(\epsilon/2)}\right] \\ &\subset \left[\frac{1 - \beta\epsilon/2}{1 + \beta\epsilon/2}, \frac{1 + \beta\epsilon/2}{1 - \beta\epsilon/2}\right] \approx [1 - \beta\epsilon, 1 + \beta\epsilon]\end{aligned}$$

where v is the maximum source velocity, $\beta := v/c \ll 1$, and the approximation results from discarding the insignificant higher order terms in the series representation of the ratio.

In the two-source time-slotted round-trip carrier synchronization protocol, beacons are transmitted around the $\mathsf{D} \to \mathsf{S}_1 \to \mathsf{S}_2 \to \mathsf{D}$ and $\mathsf{D} \to \mathsf{S}_2 \to \mathsf{S}_1 \to \mathsf{D}$ circuits. We denote the Doppler shift, i.e., the ratio of the output to input frequency, in the $g_{k,i}(t)$ channel as $\gamma_{k,i}$. Focusing on propagation through the first circuit, the frequency of the primary beacon received by S_1 can be written as $\omega_{01} = \gamma_{0,1}\omega$. Since the secondary beacon transmitted by S_1 is a periodic extension of this primary beacon, the secondary beacon transmitted by S_1 will be received by S_2 at frequency $\omega_{12} = \gamma_{1,2}\omega_{01} = \gamma_{1,2}\gamma_{0,1}\omega$. And finally, since S_2 will transmit its carrier as a periodic extension of this secondary beacon, the carrier from S_2 will be received by D at frequency $\omega_{20} = \gamma_{2,0}\gamma_{1,2}\gamma_{0,1}\omega$. A similar analysis can be applied to the second circuit.

When the sources are moving with constant velocity, each circuit incurs a *composite* Doppler shift that is the product of the individual Doppler shifts in each link. If we bound the absolute difference of the angles of all of the forward and reverse links in the round-trip synchronization system by ϵ, the previous result implies that the ratio of the composite Doppler shifts, i.e.

$$\begin{aligned}\rho_c := &(\text{Doppler}(\mathsf{D} \to \mathsf{S}_1 \to \mathsf{S}_2 \to \mathsf{D})/ \\ &\text{Doppler}(\mathsf{D} \to \mathsf{S}_2 \to \mathsf{S}_1 \to \mathsf{D}))\end{aligned}$$

is bounded by

$$\begin{aligned}\rho_c \in &[(1-\beta\epsilon)^2(1-2\beta\epsilon), (1+\beta\epsilon)^2(1+2\beta\epsilon)] \\ &\approx [1 - 4\beta\epsilon, 1 + 4\beta\epsilon]\end{aligned}$$

where we have again discarded the higher order terms in the approximation. Note that the $(1 \pm 2\beta\epsilon)$ terms correspond to the ratio of the Doppler shifts for the secondary beacons exchanged between the sources since a source's velocity relative to the other source is bounded by $2v$. This result implies that the carrier frequency offset during beamforming can be bounded by

$$|\omega_\Delta| \le 4\beta\epsilon\omega.$$

As a numerical example of the bound, consider a system with a 900–MHz carrier frequency, maximum velocity $v = 100$ m/s, and an angle bound $\epsilon = 0.0175$ radians, corresponding to 1 degree of maximum angle difference between all forward and reverse transmissions. In this case, the maximum carrier frequency offset at the destination can be computed to be $|\omega_\Delta| \le$ 21 Hz. This amount of carrier frequency offset is likely to be insignificant in most cases when compared to the frequency offset caused by estimation error as well as the phase noise of typical low-cost radio frequency oscillators.

C. Discussion

The results in this section demonstrate that the round-trip carrier synchronization protocol can be designed such that even significant levels of mobility have little effect on the performance of the distributed beamformer in single-path channels. For a given carrier frequency, the initial carrier phase offset due to mobility can be reduced to a desired level by limiting the range of the links (and hence limiting the maximum propagation time t_p), limiting the maximum source velocity, and/or by transmitting short secondary beacons. Similarly, carrier frequency offset during beamforming due to mobility can be reduced by limiting the change in the velocity angles between forward and reverse accesses of a channel. This can also be achieved by limiting the maximum velocity and/or by using short beacon durations.

In systems with mobility, the selection of beacon durations involves a tradeoff between the carrier phase and frequency offsets caused by mobility and the carrier phase and frequency offsets caused by estimation error at each source. Selecting very short beacon durations can minimize the effects of mobility but lead to high levels of estimation error and unnecessary synchronization overhead. On the other hand, very long beacon durations may result in precise phase and frequency estimates at each source, but significant carrier phase and frequency offsets due to mobility as well as operation in the phase-noise dominant regime. As was the case without mobility, beacon durations should be selected to balance the effect of these impairments.

As a final point on the implementation of round-trip carrier synchronization in systems with mobility, we note that the optimum linear combining parameters derived in Section IV-B to minimize $\text{var}\{\omega_\Delta\}$ are, in general, not optimum in systems with Doppler shift. The combining factors μ_1^* and μ_2^* were derived for the case when both frequency estimates $\hat{\omega}_{0i}$ and $\hat{\omega}_{ki}$ at S_i are unbiased estimates of ω, as is the case for systems with LTI channels. In systems with time-varying channels and Doppler shifts, however, $\hat{\omega}_{0i}$ and $\hat{\omega}_{ki}$ may have different biases with respect to the primary beacon frequency ω. If these biases are small with respect to the standard deviation of the frequency estimation errors, the optimum combining factors derived in Section IV-B may still be used with little loss of optimality. On the other hand, if the biases due to Doppler effects are large with respect to the standard deviation of the frequency estimation errors, use of the combining factors derived in Section IV-B may significantly degrade the performance of the beamformer. In this case, the carriers should be generated as periodic extensions of the secondary beacons in order to ensure that the ratio of the composite Doppler shifts in each circuit is close to one.

VI. Conclusion

In this paper, we have proposed a time-slotted round-trip carrier synchronization protocol to enable distributed beamforming in multiuser wireless communication systems. We have provided a detailed description of the half-duplex protocol for a system with two sources and single-path time-invariant channels and have also described how the protocol can be applied to systems with more than two sources, multipath channels, and mobility. We analyzed the performance of the time-slotted round-trip carrier synchronization protocol in the two-source scenario in terms of the statistics of the carrier phase offset at the destination and the expected beamforming time before resynchronization is required. Our numerical results demonstrated that the parameters of the synchronization protocol can be selected such that a desired level of phase accuracy and reliability can be achieved during beamforming and such that the synchronization overhead can be small with respect to the amount of time that reliable beamforming is achieved. We have also analyzed the effect of time-varying channels on the performance of the protocol and demonstrated that the parameters of the synchronization protocol can be selected to make reliable beamforming possible even in systems with mobile transmitters.

Appendix

For a time-slotted round-trip distributed beamformer with two sources, as shown in Section V-A, the arrival time difference of the start of the carriers at the destination can be written as

$$\tau_{\Delta} = \tau_{\Delta_{01}} + \tau_{\Delta_{21}} + \tau_{\Delta_{02}} \tag{15}$$

where $\tau_{\Delta_{01}} := \tau_{10}(t_{31}) - \tau_{01}(t_0 + T_0), \tau_{\Delta_{21}} := \tau_{21}(t_2 + T_2) - \tau_{12}(t_1), \tau_{\Delta_{02}} = \tau_{02}(t_0 + T_0) - \tau_{20}(t_{32})$ and where $\tau_{ki}(t)$ denotes the propagation time of an impulse emitted from transmitter k to receiver i at time t. An upper bound on the magnitude of τ_{Δ} can be written as

$$|\tau_{\Delta}| \leq |\tau_{\Delta_{01}}| + |\tau_{\Delta_{21}}| + |\tau_{\Delta_{02}}|.$$

Before we make this bound more explicit, we note that the required accuracy of the forward and reverse propagation times requires consideration of relativistic effects. For example, in a system with 900 MHz beacons/carriers, a error of 100 ps in a propagation time calculation corresponds to a phase error of more than 30 degrees. To maintain consistency and avoid unnecessary transformations between reference frames, we will take all lengths and times in our analysis to be in the destination's reference frame.

Since the speed of light is constant in any reference frame, an impulse emitted by a stationary or moving source in the destination's reference frame always propagates at velocity c from the perspective of the destination. To provide an example of the fundamental calculation that we use to derive the bound, suppose a receiver is distance d_0 from an emitter of an impulse at time t_0 and is moving at constant velocity v on a line directly away from the emitter (all lengths and times are in the destination's reference frame). The emitter can be stationary or moving in the destination's reference frame without affecting the analysis. Since the impulse propagates with velocity c and the receiver is moving directly away from the point at which the impulse was emitted at velocity v, the impulse emitted at t_0 arrives at the receiver when

$$c(t - t_0) = d_0 + v(t - t_0)$$

where the lengths and times are all in destination's reference frame. Hence, the propagation time in the destination's reference frame is

$$t_{\text{prop}} = t - t_0 = \frac{d_0}{c - v}.$$

This expression can also be used to calculate the propagation time of an impulse to a stationary receiver by setting $v = 0$, or to a receiver moving directly toward from the point at which the impulse was emitted under the convention that motion toward the emitter corresponds to negative velocity [16].

With this example in mind, the first term $\tau_{\Delta_{01}}$ in (15) represents the propagation time difference in the reverse and forward links, respectively, between D and S_1. Since $t_{31} > t_0 + T_0$, this term is maximized when S_1 moves away from D at maximum velocity and is minimized when S_1 moves toward D at maximum velocity. In the former case, we can write

$$\begin{aligned}\tau^{+}_{\Delta_{01}} &= \max[\tau_{10}(t_{31}) - \tau_{01}(t_0 + T_0)]\\ &= \frac{d_{01} + v(t_{31} - t_0)}{c} - \frac{d_{01} + vT_0}{c - v}\\ &= \frac{\beta(t_{31} - t_0 - T_0 - d_{01}/c) - \beta^2 t_{31}}{1 - \beta}\end{aligned}$$

where $\beta := v/c$, d_{01} is the distance from S_1 to D at time $t = t_0$. Similarly, when S_1 moves toward D at maximum velocity, we can write

$$\begin{aligned}\tau^{-}_{\Delta_{01}} &= \min[\tau_{10}(t_{31}) - \tau_{01}(t_0 + T_0)]\\ &= \frac{d_{01} - v(t_{31} - t_0)}{c} - \frac{d_{01} - vT_0}{c + v}\\ &= \frac{-\beta(t_{31} - t_0 - T_0 - d_{01}/c) - \beta^2 t_{31}}{1 + \beta}.\end{aligned}$$

It can be shown that $|\tau^-_{\Delta_{01}}| \geq |\tau^+_{\Delta_{01}}|$. Hence, we can say that

$$|\tau_{\Delta_{01}}| \leq \frac{\beta(t_{31} - t_0 - T_0 - d_{01}/c) + \beta^2 t_{31}}{1+\beta} \leq \frac{\beta(t_{31} - t_0 - T_0) + \beta^2 t_{31}}{1-\beta} \quad (16)$$

where the second inequality is applied to allow for a more convenient representation of the overall bound on $|\tau_\Delta|$.

The second term $\tau_{\Delta_{21}}$ in (15) represents the propagation time difference in the links between S_1 and S_2. Since $t_2 + T_2 > t_1$, this term is maximized when the sources are separating at maximum velocity and is minimized when the sources are approaching each other at maximum velocity. In the former case, we can write

$$\begin{aligned}\tau^+_{\Delta_{21}} &= \max[\tau_{21}(t_2+T_2) - \tau_{12}(t_1)] \\ &= \frac{d_{12} + 2v(t_2+T_2-t_0)}{c-v} - \frac{d_{12}+2v(t_1-t_0)}{c-v} \\ &= \frac{2\beta(t_2+T_2-t_1)}{1-\beta}\end{aligned}$$

where d_{12} is the distance from S_1 to S_2 at time $t = t_0$. Similarly, when the sources are approaching each other at maximum velocity, we can write

$$\begin{aligned}\tau^-_{\Delta_{21}} &= \min[\tau_{21}(t_2+T_2) - \tau_{12}(t_1)] \\ &= \frac{d_{12} - 2v(t_2+T_2-t_0)}{c+v} - \frac{d_{12}-2v(t_1-t_0)}{c+v} \\ &= \frac{-2\beta(t_2+T_2-t_1)}{1+\beta}.\end{aligned}$$

It can be shown that $|\tau^+_{\Delta_{21}}| \geq |\tau^-_{\Delta_{21}}|$. Hence, we can say that

$$|\tau_{\Delta_{21}}| \leq \frac{2\beta(t_2+T_2-t_1)}{1-\beta}. \quad (17)$$

The third term $\tau_{\Delta_{02}}$ in (15) represents the propagation time

[12] D. Rife and R. Boorstyn, "Single-tone parameter estimation from discrete-time observations," *IEEE Trans. Inf. Theory*, vol. IT-20, pp. 591–598, Sep. 1974.
[13] A. Demir, A. Mehrotra, and J. Roychowdhury, "Phase noise in oscillators: A unifying theory and numerical methods and characterization," *IEEE Trans. Circuits Syst. I, Fund. Theory Appl.*, vol. 47, pp. 655–674, May 2000.
[14] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York: Springer-Verlag, 1994.
[15] T. Gill, *The Doppler Effect.*. New York: Academic, 1965.
[16] R. Ferraro, *Einstein's Space-Time: An Introduction to Special and General Relativity*. New York: Springer, 2007.

D. Richard Brown III (S'97-M'00) received the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, in 2000.

He joined the faculty of Worcester Polytechnic Institute (WPI), Worcester, MA, in 2000 and is currently an Associate Professor. His research focus is on the development of energy-efficient communication systems, synchronization, and game-theoretic analysis of communication networks.

H. Vincent Poor (S'72–M'77–SM'82–F'87) received the Ph.D. degree in electrical engineering and computer science from Princeton University, Princeton, NJ, in 1977.

From 1977 to 1990, he was on the faculty of the University of Illinois at Urbana-Champaign. Since 1990 he has been with Princeton University, where he is the Michael Henry Strater University Professor of Electrical Engineering and Dean of the School of Engineering and Applied Science. His research interests are in the areas of stochastic analysis, statistical signal processing, and their applications in wireless networks and related fields. Among his publications in these areas are the recent book *MIMO Wireless Communications* (Cambridge, U.K.: Cambridge University Press, 2007), and the forthcoming book *Quickest Detection* (Cambridge, U.K.: Cambridge University Press, 2009).

Dr. Poor is a member of the National Academy of Engineering, a Fellow of the American Academy of Arts and Sciences, and a former Guggenheim Fellow. He is also a Fellow of the Institute of Mathematical Statistics, the Optical Society of America, and other organizations. In 1990, he served as President of the IEEE Information Theory Society, and in 2004–2007 he served as the Editor-in-Chief of the IEEE TRANSACTIONS ON INFORMATION THEORY. Recent recognition of his work includes the 2005 IEEE Education Medal, the 2007 IEEE Marconi Prize Paper Award, and the 2007 Technical Achievement Award of the IEEE Signal Processing Society.