A Scalable Architecture for Distributed Transmit Beamforming with Commodity Radios: Design and Proof of Concept

François Quitin, Muhammad Mahboob Ur Rahman, Raghuraman Mudumbai, and Upamanyu Madhow

Abstract—We describe a fully-wireless prototype of distributed transmit beamforming on a software-defined radio platform. Distributed beamforming is a cooperative transmission technique that can achieve orders of magnitude increases in range or energy efficiency of wireless communication systems. However, this technique requires precise synchronization of the radio frequency signal from each transmitter. The significance of our prototype is in demonstrating that this requirement can be satisfied using digital signal processing methods on commodity hardware with low-quality oscillators. Our synchronization approach scales to large numbers of transmitters: each transmitter runs independent algorithms based on periodically transmitted feedback packets from the receiver. A key simplification is the decoupling of the algorithms for frequency locking and beamsteering at each transmitter, even though both processes use the same feedback packets. Frequency locking employs an Extended Kalman filter to track the local oscillator offset between a transmitter and the receiver, using frequency offset measurements based on the feedback packet *waveform*, while the phase adjustments for beamsteering are determined using a one-bit feedback algorithm based on the feedback packet *payload*. Our prototype demonstrates that distributed transmit beamforming can be incorporated into wireless networks without requiring hardware innovations, and provides open-source building blocks for future research and development.

Index Terms—Distributed MIMO, beamforming, synchronization, wireless networks.

I. INTRODUCTION

DISTRIBUTED transmit beamforming refers to a cooperative transmission technique for wireless networks under which several nodes calibrate their transmissions in such a way that their individual transmitted signals add up coherently at an intended receiver. In effect, under this scheme, the transmit nodes act as a "virtual antenna array" and direct a beam towards the receiver, enhancing communication range and/or power efficiency [1], [2], [3].

The key challenge in realizing the large potential gains from beamforming is in precisely synchronizing the RF signals from each transmitter so that they are aligned in phase at

the intended receiver. Each transmitter obtains its RF carrier signal from a separate local oscillator, and even when two oscillators are set to the same nominal frequency, because of manufacturing tolerances and temperature variations, they would in general have a non-zero frequency offset with respect to each other. In addition, all oscillators undergo random phase and frequency drifts over time. Finally, unlike a traditional phased array, a virtual array of collaborating wireless nodes does not have a regular and precisely known geometry. Thus, even if the destination's location were known precisely (which is not the case in practice), it is not possible to determine the phases that the transmitters in the virtual array must employ in order to direct energy towards the destination. Indeed, explicitly computing beamforming weights requires knowledge of the geometry of the transmitters and receiver to within a small fraction of the carrier wavelength, and standard localization techniques such as GPS fall far short of the accuracy necessary to overcome this geometric uncertainty. Thus, distributed beamforming requires a sophisticated synchronization architecture that accounts for all of the preceding uncertainties. The prototype described in this paper illustrates a robust and scalable feedback-based approach to this problem, and successfully demonstrates that distributed transmit beamforming is indeed feasible even with off-the-shelf hardware and low-quality oscillators.

Contributions: We present an architecture for distributed transmit beamforming that is suitable for packet wireless networks, and demonstrate it using a software-defined radio testbed. The objectives of our prototyping effort are two-fold: (a) a proof-of-concept implementation to show that distributed beamforming can be implemented on commodity hardware with minimal overhead using signaling similar to that in existing packet wireless networks, and (b) reusable open-source building blocks for further development, which we hope will lead to eventual incorporation of distributed beamforming in real-world WiFi, Zigbee and cellular wireless networks. Our main contributions are summarized as follows.

1) **Scalable synchronization architecture.** We propose a decentralized, feedback-based synchronization architecture in which each transmitter adapts its frequency and phase independently, based on feedback from the receiver. In principle, this allows the system to scale to an indefinitely large number of cooperating transmitters. At each transmitter, the problems of frequency locking and beamsteering are decoupled into two parallel algorithms.

From the point of view of implementation, a key contribution is the use of the same all-digital feedback packets for both parallel algorithms: the single bit of feedback used for phase adaptation for beamsteering is embedded in the payload of a GMSK modulated packet, and the modulated waveform in the packet is used to extract information for frequency locking.

2) **Extended Kalman filter for frequency locking.** Each transmitter employs an extended Kalman filter (EKF) to lock its frequency to that of the receiver. The EKF is driven by phase and frequency offset measurements made from the modulated waveforms in the feedback packets, but does not make use of the payload of the feedback packet (which carries information to be used for beamsteering). The EKF provides a robust mechanism for using wrapped phase offset measurements to track the unwrapped phase offset in the state evolution. We provide an approximate analysis of the EKF. This yields design guidelines for the required feedback rate which are validated through experiments.

3) **1-bit feedback based beamsteering.** The receiver broadcasts one bit indicating the change in its received signal strength in each feedback packet. Each transmitter uses this to make phase corrections (on top of the frequency/phase corrections for frequency locking based on the EKF) for beamsteering, using the randomized ascent algorithm first proposed in [4].

Related work: Many information-theoretic analyses, ranging from three decades back [5], [6] to the present [7], [8], [9], rely on the concept of distributed transmit beamforming, without addressing the synchronization requirements involved in realizing it. Over the last decade, however, there have been serious attempts to tackle these difficult synchronization problems, and a menu of synchronization techniques have been studied, including closed loop methods with explicit channel feedback for each transmitter [10], [11], one-bit aggregate feedback [4], [12], [13], implicit feedback using reciprocity [2], round-trip synchronization [14], [15], and two-way synchronization [16], [17]. These different synchronization techniques represent different sets of tradeoffs between overheads of coordination, channel feedback and complexity (see the survey article [3], and a discussion of more recent work in [18]). Distributed beamforming has also been recently studied in the context of "coordinated multipoint (CoMP)" capabilities for 4G-LTE cellular systems, where multiple base stations act as a distributed antenna array [19].

In addition to these theoretical investigations, this technique has also been demonstrated in several experimental prototypes [12], [20], [21], [22]. The 1-bit feedback algorithm, presented in [12], [13], has proved to be especially popular for experimental investigations because of its simplicity and scalability.

The results reported in this paper represents a significant advance over previous work in a number of important respects. Early experimental prototypes for distributed beamforming [12], [20], [21] all used wired side-channels for distributing common reference clock signals and/or channel feedback messages to the transmitters. Similarly the CoMP studies for 4G cellular systems assume that the cooperating base stations are all connected by dedicated backhaul links with high bandwidths and low latencies, as well as uninterrupted GPS satellite connections. These assumptions substantially simplify the synchronization process, and therefore these early prototypes do not fully resolve the question of the feasibility of distributed beamforming in practical ad-hoc wireless networks. Furthermore these prototypes all use custom-designed hardware which is not compatible with existing standards and devices.

Some previous work has addressed the problem of over-the-air frequency synchronization. In [23], a Kalman filter framework is presented to track the LO clock offset between any pair of nodes. However, it is assumed that the unwrapped LO phase offset could be measured, which is only possible if the LO frequency offset between any pair of nodes is sufficiently small to avoid any uncertainty in the number of 2π-phase rotations between two successive measurements. The early prototype in [22] uses analog signaling, and signal processing techniques based on analog circuits such as Costas loops, which are less efficient and require a dedicated node to serve as a master node, as well as a dedicated frequency band to transmit a continuous analog signal. If this signal is momentarily interrupted (due to a lag in the hardware or software), synchronization between the nodes is lost. Moreover, the Costas loop has two stable steady states that are 180° apart. Occasionally, phase noise may cause the Costas loop to switch from one stead-state to the other, resulting in a sharp drop in RSS of the beamforming signal.

An important contribution of the prototype described here, relative to prior work, is to show that all-wireless distributed beamforming can be achieved with commodity hardware with modest overhead, without requiring external (e.g., GPS-based) synchronization or wired side channels. Our EKF synchronization framework uses *wrapped* LO phase offset measurements to run the frequency synchronization, without making any assumptions on the unwrapped phase. The results in this paper significantly improve upon our prior efforts [22], [24] towards this goal. The present prototype is an evolution of an earlier implementation reported in our recent conference paper [24]: while this also employed digital feedback messages, it employed separate training signals for frequency locking and beamsteering. The present prototype does not require manual calibration as in [22], and improves upon [24] by using the same digital feedback packets (with standard modulation formats) for both purposes. A recent live demo of the prototype was presented at [25]. We go beyond these conference publications by providing detailed description and analysis of our architecture, including an analysis of the EKF which yields design guidelines for setting the feedback rate.

As with prior prototypes, we demonstrate coherent addition of a carrier tone at the receiver, which requires frequency locking and phase alignment. For distributed beamforming with modulated waveforms, an additional requirement is baseband timing alignment (or more generally, channel dispersion compensation). While the latter is not addressed here, frequency synchronization and narrowband phase alignment as implemented here could be used as building blocks for multi-carrier techniques for distributed beamforming over wideband dispersive channels.

**Outline:** The remainder of the paper is organized as follows.

Fig. 1.   Distributed beamforming general architecture.

Fig. 2.   Time-slotting model for frequency and phase estimation.

Section II presents the general architecture of our distributed transmit beamforming setup, and discusses the design implications of fundamental limits on frequency/phase synchronization. Section III focuses on frequency synchronization. The extended Kalman filtering framework is presented, together with approximate modeling and performance analysis which yield design guidelines. Finally, Section IV describes implementation of our architecture on software-defined radios, and presents experimental results from our testbed.

II. ARCHITECTURE FOR DISTRIBUTED BEAMFORMING

We first outline a system model motivating our implementation. The underlying approach is quite general, and is broadly applicable to WiFi, Zigbee and other packet wireless networks.

A. System model

Consider a distributed array with N nodes which seek to collaboratively transmit a common complex baseband message signal $m(t)$ to a receiver. Let the RF signal transmitted by a transmit node be denoted by $x(t)$ and θ the phase gain of the channel from the transmitter to the receiver (note that θ is different for each transmitter). In order to achieve beamforming, each transmit node needs to synthesize and transmit an RF signal $x(t) = \Re\left(m(t)\exp(j2\pi f_c t - j\theta)\right)$.

However, each transmitter has a local oscillator (LO) that derives its RF signal from its own separate crystal reference which in general have small but non-zero frequency offsets with each other. Let α denote the fractional offset of the reference signal of the transmit node from some standard universal reference. These offsets are usually specified in parts per million (ppm); thus, a 10 ppm offset corresponds to $\alpha = 10e - 6$. The resulting frequency offset of the LO signal of the transmit node at the nominal frequency f_c is $\Delta f \equiv \alpha f_c$. Thus the node's LO signal can be written as $e(t) = \exp\left(j(2\pi(f_c + \Delta f)t + \Delta\phi)\right)$, where $\Delta\phi$ is the (unknown) phase offset. The problem of beamforming is to synthesize the signal $x(t)$ from $e(t)$.

Figures 1 shows a schematic representation of our system. As shown in the figure, the receiver regularly broadcasts feedback messages; the transmitter nodes use these feedback messages to estimate and correct for the offsets Δf, $\Delta\phi$ and to construct the desired beamforming signal $x(t)$.

Figure 2 shows the the time-slotting model for the transmit nodes. Every time a feedback message is received from the rece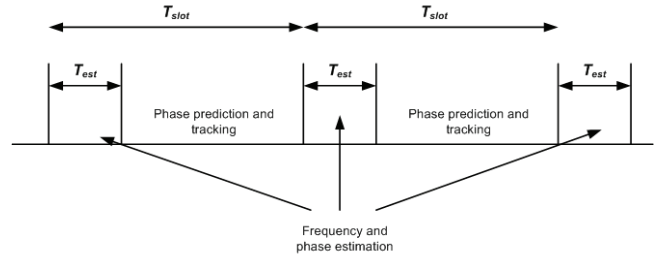iver, the transmit node will use the feedback message to make an estimation of the frequency and phase offset Δf and $\Delta\phi$. These estimations will be used by the transmit nodes to predict and compensate for the LO offset until the next feedback message is received.

The feedback messages can be transmitted over the same frequency band as the beamforming signal using a medium access control mechanism enabling time sharing, or we can employ frequency division multiplexing, with the feedback and beamforming signals sent over different frequency bands. While our architecture applies to both scenarios, we employ frequency division multiplexing in our prototype. Thus, the beamforming nodes employ a reference signal at one frequency to synthesize a synchronized RF signal at a different frequency. To do this, we assume that both carrier frequencies are synthesized from the same crystal oscillator, hence there is a known multiplicative relationship between them which also applies to the frequency offsets between two nodes that we wish to estimate and correct for. Specifically, the offset Δf_2 at a frequency $f_{c,2}$ can be obtained using the measured offset Δf_1 at another frequency $f_{c,1}$ as $\Delta f_2 = \alpha f_{c,2} \equiv \Delta f_1 \frac{f_{c,2}}{f_{c,1}}$.

B. Two synchronization sub-processes

A key feature of our implementation is that the beamforming process is decoupled into two sub-processes that run independently and concurrently. Roughly speaking, the first sub-process compensates for the frequency offset Δf, and the second sub-process compensates for the unknown phases $\Delta\phi$ and θ.

1) **Frequency locking.** In this sub-process, each transmitter locks its oscillator on to a shared reference signal, which in our case is the set of feedback packets from the receiver. The purpose of this sub-process is to ensure that the transmitters all have RF signals with the same frequency and a fixed (but unknown) phase relationship with each other. The LO frequency offsets that can occur in typical software-defined radios can range up to several kHz, making the frequency synchronization of the transmit nodes especially challenging.

2) **Beamsteering.** This sub-process adjusts the phase relationship between the transmitters in such a way that their transmitted signals add up coherently at the intended receiver. In our case, the feedback packets from the receiver contains the one bit of SNR feedback in its payload that is used to steer the beam using the 1-bit algorithm.

Fig. 3.   Block diagram of a transmit node.

Fig. 4.   Block diagram of a receive node.

The main motivation for this decoupled design is simplicity: the 1-bit algorithm is easy to implement and has low overhead. While it can be modified to provide both frequency and phase synchronization [20], it cannot handle the significant frequency drifts that we encounter in our prototype, especially given the large latencies in the feedback channel. The frequency locking process estimates and eliminates these frequency offsets and allows the simple 1-bit algorithm to achieve and maintain coherence.

Our implementation of the two synchronization sub-processes are shown in the block diagram of a transmit node in Figure 3. A key feature of our design is that the same set of periodically transmitted feedback packets from the receiver (depicted in Figure 4) are used for both synchronization sub-processes: the payload of the feedback packet contains the one bit of SNR information to drive the beamforming algorithm, and the frequency and phase offset of the transmitter's oscillator with respect to the receiver is implicitly contained within the preamble and header symbols of the packet.

Note that the phase synchronization with the one-bit feedback is useless as long as frequency synchronization has not been attained. However, since the frequency synchronization is based only on the waveform of the feedback packet, it is not slowed down by the one-bit feedback algorithm. The one-bit feedback algorithm will only become useful once the frequencies of the different transmit nodes are synchronized.

It should also be noted that the use of multiple-bit feedback has the potential to improve the performances of the one-bit feedback algorithm. A preliminary study in [26] investigates how rich feedback can significantly enhance convergence speed, while still maintaining scalability by designing the feedback such that the same aggregate feedback can be used by all nodes in the network (as opposed to feedback customized for each transmitter, which does not scale with the number of cooperating nodes). This is especially critical for larger networks, where the slow convergence speed of the one-bit feedback algorithm becomes a problem. For the sake of simplicity, however, the prototype presented in this paper requires only a single bit of feedback for the beamsteering.

C. Fundamental limits of frequency and phase estimation

The performance limits and convergence properties of the 1-bit beamforming algorithm (assuming ideal frequency synchronization) have been studied in detail [13]. We therefore focus here on understanding the frequency synchronization sub-process, which estimates frequency and phase using training signals transmitted in short periodic bursts of duration T_{est}, and then smooths these estimates using an EKF. The time between these bursts is denoted by T_{slot}. In this section, we discuss what insights the Cramer-Rao Lower Bound (CRLB) for one-shot frequency/phase estimation provides regarding the desirable regime of operation for the frequency synchronization sub-process. These insights are then verified by simulations and experiments quantifying EKF performance.

Consider the process of obtaining one-shot frequency and phase estimates using a noise-corrupted reference signal received by a transmitter over the training epoch of duration T_{est} in one time-slot. Let $a(t) = A \exp(j\phi(t)) + n(t), t \in [0, T_{est}]$ which is the complex baseband waveform corresponding to one feedback packet upon demodulation using the LO signal of the transmit node. The post-integration SNR of this signal is defined as $\text{SNR} \equiv \frac{A^2 T_{est}}{2N_0}$, where N_0 is the power spectral density of the white noise process $n(t)$. The CRLBs for this one-shot phase and frequency estimation process are well-known in the literature [27], [28]: if ϕ_{err} and f_{err} respectively denote the one-shot phase and frequency estimation errors, we have

$$\sigma_\phi^2 \doteq E\left[\phi_{err}^2\right] \geq \frac{2}{\text{SNR}}$$
$$\sigma_f^2 \doteq E\left[f_{err}^2\right] \geq \frac{3}{2\pi^2 T_{est}^2 \text{SNR}} \quad (1)$$

Consider now the phase error that results when transmitters use one-shot frequency and phase estimates from the training interval to predict and correct for the frequency and phase offsets of their oscillators over the subsequent time slot. The variance of the resulting error $\phi(t) - \hat{\phi}(t)$ between the predicted phase offset $\hat{\phi}(t)$ and actual phase offset $\phi(t)$ of the transmitter with the reference signal grows with time and its value at the end of the time-slot can be written as

$$E\left[\left(\phi(t) - \hat{\phi}(t)\right)^2\right]_{t=T_{slot}} = \sigma_\phi^2 + T_{slot}^2 (2\pi\sigma_f)^2$$
$$\geq \frac{2}{SNR}\left(1 + \frac{3}{\eta^2}\right). \quad (2)$$

When the duty cycle of the estimation process is small i.e. $\eta \equiv \frac{T_{est}}{T_{slot}} \ll 1$, then the second term in (2) dominates; in this setting, *one-shot frequency estimates are highly unreliable* as compared to the phase estimate.

Now consider an alternative approach to the frequency estimation problem. Instead of doing one-shot frequency estimates, we can also estimate frequency by using two one-shot phase estimates in two successive training epochs T_{slot} seconds apart. In other words, we consider the frequency estimate $\tilde{f} \doteq \frac{\hat{\phi}(T_{slot}) - \hat{\phi}(0)}{2\pi T_{slot}}$. This estimate has the variance

$$\text{var}(\tilde{f}) = \frac{2\sigma_\phi^2}{(2\pi T_{slot})^2} \geq \frac{1}{\pi^2 \text{SNR} T_{slot}^2}, \quad (3)$$

and this variance can be significantly smaller than the one-shot frequency variance σ_f^2 in (1). This suggests that we might be better off dispensing with one-shot frequency estimates altogether, and rely on averaging phase estimates over multiple time slots to get good frequency estimates. Indeed, this approach, implemented using a Kalman filter, is what is employed in [23]. However, using phase estimates alone for both phase and frequency tracking requires access to *unwrapped* phase estimates. This in turn requires that the frequency error in our estimate is small enough that 2π ambiguities in phase do not appear over the slot duration T_{slot} between successive training bursts. For the low-quality oscillators in our software-defined radios, the frequency drift is severe enough that satisfying this assumption would require excessive overhead.

In order to circumvent the preceding phase unwrapping problem, we employ crude one-shot frequency estimates to complement the phase estimates. These one-shot frequency estimates need only be good enough to avoid phase unwrapping errors over a single time-slot; in other words, we want $\sigma_f T_{slot}$ that is not too much larger than unity. Plugging this into (1), we obtain the following rule of thumb.

CRLB-based rule of thumb: $\frac{T_{slot}}{T_{est}} \approx k\sqrt{SNR}$, where $k = \sqrt{\frac{2}{3}}\pi$. Interestingly, this requirement only applies to the ratio $\frac{T_{slot}}{T_{est}}$ or equivalently to the duty-cycle of the training signal, not individually to T_{slot} or T_{est}. Note that this requirement is only meant to provide very rough guidance. More detailed design insights are obtained via numerical simulations and experiments in Section III-B.

III. FREQUENCY SYNCHRONIZATION

The frequency synchronization sub-process is divided into three stages, as shown in Figure 3. In the first stage, the transmit node, upon receipt of each feedback packet, makes a measurement of its LO frequency and (wrapped) phase offset relative to the receiver using a blind estimation algorithm. In the second stage, the EKF uses these LO frequency and wrapped LO phase offset measurements to keep track of the unwrapped LO phase offset. In the third stage, the transmit node compensates for the LO offset based on the latest LO frequency and phase offset values as predicted by the EKF.

The blind estimation algorithm used in the first stage depends on the modulation format used for the feedback message. For most classical modulation formats (PSK, QAM, GMSK etc.), these algorithm transform the feedback message into a pilot tone, whose frequency can easily be estimated with classical frequency estimation theory. After compensating the feedback message for the LO frequency offset, the LO phase offset can easily be measured by correlating the feedback message with the (known) message header. This will yield a *wrapped* measurement for the LO phase offset. The estimation algorithm used for our implementation applies to GMSK feedback messages [29], [30], and is described in detail in Appendix A.

A. Extended Kalman filter state-space model

We use the following discrete state-space model for the LO offsets of each transmit node relative to the receiver.

$$\mathbf{x}_{k+1} = \mathbf{F}\mathbf{x}_k + \mathbf{w}_k \qquad (4)$$

where $\mathbf{x}_k = [\phi_k, \omega_k]^T$ is the LO phase and angular frequency offset of the transmit node with respect to the receive node at time-slot k (where $\omega_k = 2\pi\Delta f_k$). The state update matrix $\mathbf{F}$ is defined by

$$\mathbf{F} = \begin{bmatrix} 1 & T_{slot} \\ 0 & 1 \end{bmatrix}$$

and T_{slot} is the period of the feedback messages. Note that if aperiodic feedback messages are considered, T_{slot} is not fixed and the state update matrix $\mathbf{F}$ is allowed to be time-varying. The process noise vector $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{Q}(T_{slot}))$ is the noise that causes the LO phase and frequency offset to deviate from their nominal value[1].

We use the following measurement model for the blind LO offset estimation algorithm which provides the inputs for the EKF.

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{v}_k \qquad (5)$$

where

$$\mathbf{h}(\mathbf{x}_k) = \begin{bmatrix} \cos(\phi_k) \\ \sin(\phi_k) \\ \omega_k \end{bmatrix}$$

and $\mathbf{v}_k \sim \mathcal{N}(0, \mathbf{R})$ is the additive white Gaussian measurement noise. Note that (5) defines a *non-linear* measurement model reflecting the fact that the blind estimation algorithm yields only an estimate of the *wrapped* phase offset.

The equations that determine the EKF evolution are split into two stages: an update phase and a prediction phase. The update phase corrects the current state estimate given the last measurement $\mathbf{z}_k$, and is mathematically defined as

$$\mathbf{y}_k = \mathbf{z}_k - \mathbf{h}(\mathbf{x}_{k|k-1}) \qquad (6a)$$

$$\mathbf{S}_k = \mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^T + \mathbf{R} \qquad (6b)$$

$$\mathbf{K}_k = \mathbf{P}_{k|k-1} \mathbf{H}_k^T \mathbf{S}_k^{-1} \qquad (6c)$$

$$\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + \mathbf{K}_k \mathbf{y}_k \qquad (6d)$$

$$\mathbf{P}_{k|k} = (\mathbf{I}_2 - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_{k|k-1} \qquad (6e)$$

The matrix $\mathbf{H}_k$ is the Jacobian of the function $\mathbf{h}$:

$$\mathbf{H}_k = \left. \frac{\partial \mathbf{h}}{\partial \mathbf{x}} \right|_{\mathbf{x}_{k|k-1}}$$

The prediction phase gives an estimation of the future state $\mathbf{x}_{k+1|k}$ to be used in the update phase of next EKF cycle:

$$\mathbf{x}_{k+1|k} = \mathbf{F}\mathbf{x}_{k|k} \qquad (7a)$$

$$\mathbf{P}_{k+1|k} = \mathbf{F}\mathbf{P}_{k|k}\mathbf{F}^T + \mathbf{Q} \qquad (7b)$$

For each EKF cycle, the values contained in the vector $\mathbf{x}_{k|k}$ give a filtered estimate for the unwrapped LO phase offset and LO angular frequency offset. The output of the EKF in Figure

---

[1]The model described in (4) is accurate when modeling static scenarios, where the nodes do not move. In the case of mobile scenarios, a three-state model can be used to include the effects due to kinematics (which produces Doppler shift) in the LO model [23].

3 will be the filtered unwrapped LO phase offset $\widetilde{\phi}_k$ and LO frequency offset $\widetilde{\omega}_k$.

The interplay between LO phase offset and LO frequency offset in equations (6)-(7) can be intuitively understood by considering the elements of $\mathbf{y}_k$. It is first important to observe that the phase terms of $\mathbf{y}_k$ (the first two elements of $\mathbf{y}_k$) cannot exceed 2, whereas the frequency term of $\mathbf{y}_k$ (the third element of $\mathbf{y}_k$) can be arbitrarily large. In the early cycles of the EKF, the differences between the estimated and measured LO frequency offsets are often large. As a results, the phase terms of $\mathbf{y}_k$ will be negligible compared to the frequency term of $\mathbf{y}_k$, and the LO frequency offset will be the main driving element of the EKF. Once the estimated LO frequency offset approaches its measured values, the phase terms of $\mathbf{y}_k$ will no longer be negligible compared to the frequency term of $\mathbf{y}_k$. In this regime, the previously predicted LO frequency offset is used to determine the number of 2π-phase wraps that has occurred between the previous cycle and the current one. The current LO phase and frequency measurement are then used to adjust the previously predicted LO phase and frequency offset.

B. EKF convergence

It was argued in Section II-C that in order for our estimation framework to successfully handle phase unwrapping ambiguities, $\sigma_f T_{slot}$ should not be too large. The value of $\sigma_f T_{slot}$ depends on T_{slot}, T_{est} and the SNR of the feedback link. In this section, we will investigate how these parameters affect the performances of the EKF, both with simulations and experiments.

Our simulations are based on a model for the process noise covariance matrix $\mathbf{Q}$ borrowed from [31], [23]. The state-space noise covariance matrix is defined by

$$\mathbf{Q}\left(T_s\right) = \omega_c^2 q_1^2 \begin{bmatrix} T_s & 0 \\ 0 & 0 \end{bmatrix} + \omega_c^2 q_2^2 \begin{bmatrix} \frac{T_s^3}{3} & \frac{T_s^2}{2} \\ \frac{T_s^2}{2} & T_s \end{bmatrix} \quad (8)$$

where ω_c is the carrier frequency and T_s is the sample period. The parameters q_1^2 and q_2^2 are the process noise parameters that correspond to white frequency noise and random walk frequency noise, respectively. For a class of oscillators, these two parameters can be obtained by using the Allan variance.

The Allan variance is a tool to characterize the frequency stability of an oscillator, under the presence of various noise sources. It is mathematically defined as

$$\sigma_y^2\left(\tau\right) = \frac{1}{2\tau^2} \left\langle \left(\phi\left(t + 2\tau\right) - 2\phi\left(t + \tau\right) + \phi\left(t\right)\right)^2 \right\rangle_t \quad (9)$$

where $\phi\left(t\right)$ is the LO phase offset at time instant t with respect to some absolute reference. By applying equation (9) to the state-space model (4) and the noise model (8), it is shown in [31] that the following theoretical model can be obtained for the Allan variance:

$$\sigma_y^2\left(\tau\right) = \frac{q_1^2}{\tau} + \frac{q_2^2 \tau}{3} \quad (10)$$

The Allan variance can also be measured experimentally be sending a pilot tone with a transmitter, and by recording the received pilot tone (which will contain a certain LO clock offset). By entering the unwrapped phase of the received


Fig. 5. Average phase error of the EKF for various values of T_{est} and T_{slot} obtained with simulations. The colors intensity represents the phase error (in °), and the black line represents the CRLB.

pilot tone in (9) for various values of τ, it is possible to obtain an experimental curve for the Allan variance. By fitting experimental Allan variance measurements to the theoretical model (10), it is possible to obtain values for q_1^2 and q_2^2. For the software-defined radios used in the setup described in this paper, the obtained parameters are $q_1^2 = 8.47 \times 10^{-22}$ and $q_1^2 = 5.51 \times 10^{-18}$.

Figure 5 shows the simulated average phase error at the output of the EKF for various values of T_{slot} and T_{est}. Figure 5 also shows the CRLB-based rule of thumb deduced in Section II-C. It can be observed from the simulation results that there is a steep transition between the area where the EKF works perfectly, with phase errors close to zero, and the area where the EKF diverges, with phase errors that become arbitrarily large. It can also be observed that, for various SNR levels of the feedback link, the convergence/divergence transition area of the EKF depends on $\frac{T_{slot}}{T_{est}}$. For SNR levels above 20 dB, it can be seen that the convergence/divergence transition area of the EKF matches fairly well with the CRLB-based rule of thumb. Note that the rule of thumb does not represent a hard threshold, but rather a transition zone between the EKF convergence and EKF divergence area.

The experimental setup (described in Section IV) with one transmit node was also used to determine the convergence/divergence of the EKF. Note that in order to obtain a measurement of the LO frequency offset, the GMSK feedback packet is transformed in a pilot tone with a non-linear transformation, as explained in Appendix A. The theory developed before for pilot tone thus also applies to our setup, which uses GMSK feedback packets. Figure 6 shows the phase of the signal received by the receive node when only one of the transmit node is transmitting (when the transmit node is synchronizing on the receive node). In this figure, T_{slot} is varied while T_{est} is fixed to 5.1 ms. It can be seen that for high values of T_{slot} (feedback rates below 20 Hz), synchronization is not achieved and the frequency of the transmitted signal is not synchronized with the receiver, hence the phase variations. For feedback rates of 20 Hz and above, the received signal has a fairly constant phase indicating EKF convergence.

Fig. 6.   Phase of the received signal as a function of feedback rate, obtained with our experimental setup when using only one transmit node.

Fig. 7.   EKF convergence/divergence as function of T_{est} and T_{slot}, both for simulations and measurements. The blue/red area represents the $\{T_{slot}, T_{est}\}$-pairs for which the EKF converges/diverges, based on the simulation results. The circles/crosses represents the $\{T_{slot}, T_{est}\}$-pairs for which the EKF converges/diverges with the experimental setup. The black line show the CRLB.

Figure 7 shows the convergence/divergence of the EKF for various values of T_{slot} and T_{est}, both with simulations and with our experimental setup. The blue/red area represents the simulation results: the red area represents the operating points where the phase error at the output of the EKF was larger than $15°$, the blue area represents the operating points where the phase error at the output of the EKF was smaller than $15°$. The CRLB of $\sigma_f T_{est}$ is again plotted for comparison. The LO phase offset error can not be determined with our experimental setup since the real LO phase offset cannot be measured; it can only be observed if the EKF of the transmit node is synchronized on the receive node or not (similarly to what can be observed in Figure 6). The circles represents values of $\{T_{slot}, T_{est}\}$ for which the transmit node was able to synchronize to the receive node, the crosses represent values of $\{T_{slot}, T_{est}\}$ for which the transmit node could not synchronize to the receive node. It can be seen that the experimental results match well with the simulations, and that both match quite well with the CRLB-based rule of thumb.

C. Performance of the EKF

The error on the filtered LO phase and frequency estimate $\widetilde{\phi}_k$ and $\widetilde{\omega}_k$ can be recovered from the state error covariance matrix $\mathbf{P}_{k|k}$. The diagonal elements of this matrix correspond to the LO phase estimation error $\sigma_{\widetilde{\phi}}^2$ and the LO frequency estimation error $\sigma_{\widetilde{\omega}}^2$. For a stable system, this matrix converges to a stable value, which can be computed recursively with equations (6b), (6c), (6e) and (7b). For an EKF, this value cannot be computed independently of the state, since the Jacobian $\mathbf{H}_k$ depend on the actual state $\mathbf{x}_{k|k}$. For that reason, an equivalent linear Kalman filter is defined with the state-space model from (4) and the following measurement model:

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}'_k \qquad (11)$$
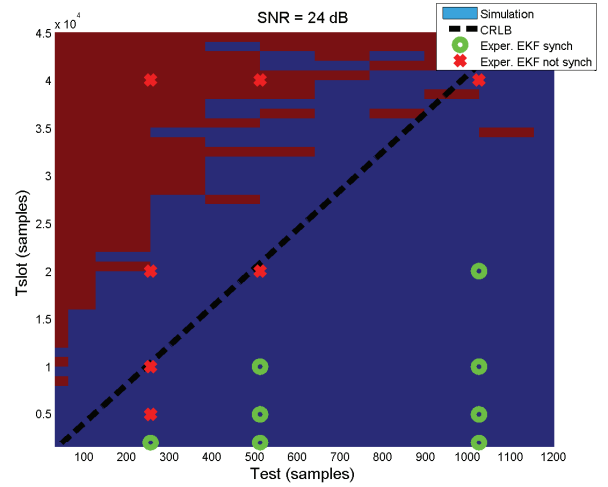
where $\mathbf{H} = \mathbf{I}_2$ is an identity matrix and $\mathbf{v}'_k \sim \mathcal{N}(0, \mathbf{R}')$ is the additive white Gaussian measurement noise. For this linear Kalman filter, the measurement matrix $\mathbf{H}_k = \mathbf{I}_2$ is independent of the state, and $\mathbf{P}_{k|k}$ can be computed independently of the state. If the EKF is in the convergence area determined in Section III-B, it can be assumed that the number of 2π-phase wraps is measured correctly, and the performance of the EKF is expected to be identical to the performance of the linear Kalman filter.

The measurement noise matrix corresponding to our setup is $\mathbf{R} = \begin{bmatrix} 0.05 \cdot \pi/180 & 0 \\ 0 & 1.5 \cdot 2\pi \end{bmatrix}$. By entering this and the process noise parameters in equations (6b), (6c), (6e), (7b) and (8), it can be computed that the state error covariance matrix converges to a stable value after a few cycles. The error on the frequency estimate at the output of the Kalman filter $\sigma_{\widetilde{\omega}}^2$ is given by element $(2, 2)$ of $\mathbf{P}_{k|k}$.

The frequency offset correction block in Figure 3 will use the latest phase and frequency estimates until new ones are obtained from the EKF. When computing the error on the filtered LO phase and frequency estimates $\sigma_{\widetilde{\phi}}$ and $\sigma_{\widetilde{\omega}}$, it can be determined that the LO phase estimation error is very low. It is mainly the LO frequency estimation error that will deviate the phase of the beamformed signal until the next state estimate is obtained from the EKF. The standard deviation of the maximum phase error then corresponds to $T_{slot}\sigma_{\widetilde{\omega}}$. The standard deviation of the maximum phase error as a function of the feedback rate $1/T_{slot}$ is shown in Figure 8. In [2] it was determined that as long as the phase error stays below $15°$, the beamforming gain stayed above 95% of the maximum beamforming gain. The $15°$ treshold is shown in Figure 8. For feedback rates above 15 Hz, the phase error shows that the beamforming gain can be maintained. The standard deviation of the maximum phase error was also measured experimentally with our experimental setup (described in Section IV). It can be seen that there is a reasonable match between the theoretical
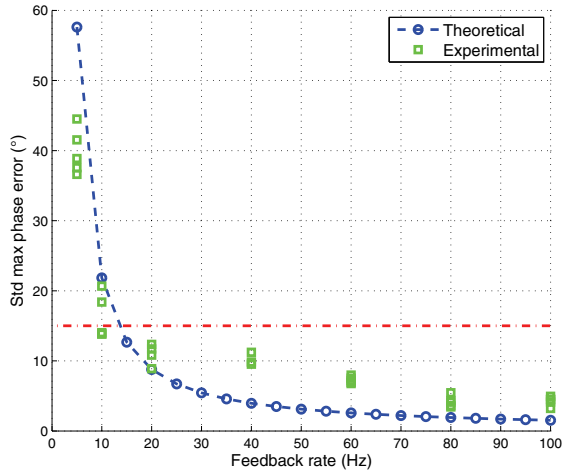
Fig. 8.   Standard deviation of the maximum phase error as a function of $1/T_{slot}$, both theoretically and experimentally. The dash-dotted line represents the $15°$ treshold below which 95% of the maximum beamforming gain is achieved.

prediction and the experimental result.

For the remainder of the paper, T_{slot} was fixed to 50 ms (a feedback rate of 20 Hz) and T_{est} was fixed to 5.1 ms, so that the setup would work even for lower SNR values, while maintaining the phase error of the beamformed signal below $15°$.

IV. IMPLEMENTATION ON SOFTWARE-DEFINED RADIOS AND RESULTS

A. Hardware implementation

The previously described architecture was implemented on a software-defined radio testbed using USRP RF and baseband boards [32]. Our setup used a mix of USRP-2 and USRP-N200 baseband boards, and WBX 50-2200 MHz RF daughter-boards; however our implementation is completely portable to any compatible SDR hardware. Our implementation is freely available for download online [33].

Three of the USRPs were used as transmit nodes, and one USRP was used as a receiver node. As explained in Section II, the different transmit nodes use the feedback messages from the receiver to synchronize both their frequency and their phase. The distance between the transmit nodes and the receive node was approximately 2 m, and the distance between the several transmit nodes was about 20 cm. We plan to port this testbed to outdoor scenarios with larger distances between the different nodes.

The estimation of the LO frequency and phase offset, using the GMSK feedback message, is explained in Appendix A. This operation also gives the precise time of arrival of the packet, and this measured value is then used for T_{slot} in the state-space matrix of the EKF.

In our setup, the feedback messages from the receiver to the beamforming nodes are sent at a frequency of 964 MHz and the beamforming signal itself is transmitted at a different frequency 892 MHz. Thus the EKF tracks the LO offset at the 964 MHz frequency, and translates this to the 892

MHz frequency using the method described in Section II, to compensate for the LO offsets in the beamforming signal.

Once frequency synchronization is achieved between the different transmit nodes, the absolute phase of the different nodes must be adjusted so that the different transmitted signal add up constructively at the receiver. The one-bit feedback algorithm is used to achieve this [12], and works as follows. At each time slot, all the transmitters add a random phase perturbation to their current phase. The receiver monitors the aggregate received signal strength (RSS) of the beamforming signal and periodically broadcasts a packet with one bit of feedback indicating increase or decrease of the RSS. The transmitters retain those perturbations that lead to RSS increases and discard the others.

The main challenge for the implementation of this algorithm on the SDR testbed is the large global round-trip latency of the system. The round-trip latencies measured with the USRPs ranged from 10 ms to 30 ms. The chosen feedback rate of 20 Hz (or 50 ms between packets) ensures that the transmitters have the time to add their new random phase perturbation before the receiver starts to measure the RSS. In our implementation the random phase perturbations are chosen to be either $+15°$ or $-15°$ which offers a good trade-off between convergence speed and residual signal fluctuations [12].

B. Experimental results

Figure 9 shows the received power at the receiver node when two nodes are cooperating to beamform towards the receiver. When the second node is turned on, it can be seen that it takes some time for the second transmit node to synchronize its frequency and to synchronize its phase. Once this is achieved, it can be seen that the received amplitude is the sum of amplitudes of the two individual transmitters. Figure 10 zooms in on the synchronization part in Figure 9. It can be seen that at first, the second transmit node is still synchronizing its frequency: the total signal is the sum of two sinusoids of unequal frequency. Once frequency synchronization is attained, the one-bit feedback algorithm will adapt the phases of the two transmit nodes until the resulting signal is the sum of the two individual signals. The time of the frequency synchronization can be reduced by entering a more precise value for the initial LO frequency offset in the EKF. The convergence speed of the synchronization with our setup equals the sum of the time needed for frequency and phase locking. With a proper initial estimate for the LO frequency offset, the frequency synchronization can be brought down to as low as a few (< 10) cycles[2]. The convergence speed of the one-bit feedback algorithm for the phase synchronization depend on the number of transmit nodes, and has been studied in [13]. As a rule of thumb, with three transmit node the one-bit feedback algorithm typically converges in less than 20 cycles, corresponding to 1 s in our setup.

The algorithm described in this paper is scalable to larger network sizes. Figure 11 shows an example with three transmit

---

[2]In our implementation, the initial frequency estimate was set manually. A better way of choosing a proper initial frequency estimate would be to use the frequency offset measurement of the first feedback message as an initial value for the EKF.

Fig. 9. Received power when two transmit nodes are beamforming. The first two bursts are the two individual nodes transmitting.
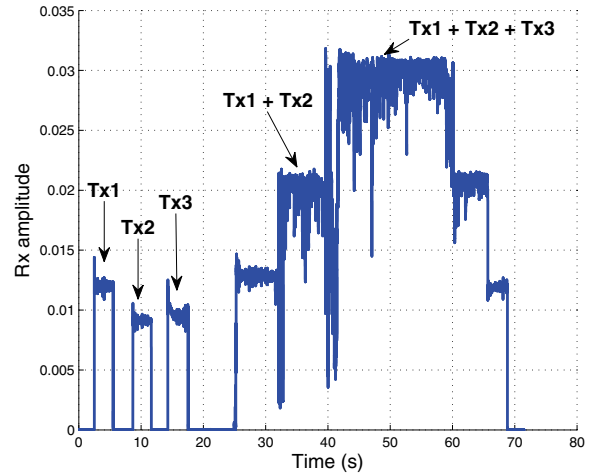

Fig. 11. Received power when three transmit nodes are beamforming.


Fig. 10. Synchronization of the two transmit nodes. At first, the two nodes are synchronizing their frequency. Then, the phase of the two nodes is adapted so that their signals add up.

its host laptop. In such cases, it takes a few cycles for the one-bit feedback algorithm to restore the phase of all the transmit nodes to their optimal value.

V. CONCLUSION

We have presented a scalable feedback-based synchronization architecture for distributed transmit beamforming, and have demonstrated that theoretical beamforming gains can be attained with commodity SDR hardware with moderate overhead. The EKF provides a robust frequency locking mechanism, using measurements of the frequency and wrapped phase LO offsets to track the unwrapped phase offset. The theoretical predictions from our analysis of the EKF match experiments, and indicate that further reductions in feedback rate can be obtained with better quality LOs. The decoupling of the one-bit feedback based phase adjustments for beamforming from EKF-based locking, while using the same set of feedback packets for both algorithms, greatly simplifies the implementation. We have shown that it is possible to use advanced modulation formats for the feedback, which opens the way for incorporating our synchronization architecture into prevalent wireless standards. The open-source implementation of our prototype is publicly available, and hopefully provides a good starting point for further experimental research on distributed MIMO.

An important next step is to extend our framework to transmission of modulated messages over wideband dispersive channels, which require baseband timing alignment as well as channel dispersion compensation. This involves design of signal processing architectures and algorithms, and their incorporation into suitably defined cross-layer protocols. Experimental investigation of the effect of mobility is another important topic.

nodes. When running the prototype, it could be observed that the received signal power sometimes drops. This can happen when the EKF convergence is momentarily lost due to a software or hardware lags in the setup. However, it could be observed that the system was able to recover quite quickly from such errors.

Note that in Figures 9-11, there are some drops in the received power. This is not due to a loss of synchrony between the transmit nodes; it is due to the self-interference of the receiver node. When the receiver is transmitting its feedback message, imperfections in the RF design of the USRP radios cause the received power on the forward link to drop, despite the different frequencies for forward and feedback link. This can be solved by moving towards a TDMA-oriented design, where forward message and feedback message are not overlapping. Occasionally, a larger drop in signal power can be observed (e.g. around 28 s in Figure 9). This can happen when the phase is lost, due to a software lag between a USRP and

APPENDIX A
LO FREQUENCY AND PHASE OFFSET ESTIMATION WITH A GMSK SIGNAL

We first describe the blind estimation of the LO frequency offset. The complex envelope of a GMSK signal can be written as

$$s(t) = \exp\left(j\frac{\pi}{2}\sum_i a_i q(t - iT)\right) \quad (12)$$

where a_i is the symbol stream, T is the symbol period and $q(t)$ is the phase pulse (which is a Gaussian phase pulse in the case of GMSK). Note that for an over-sampling rate of P, the symbol period is $T = PT_s$ where T_s is the sample period of the system.

If transmitted over a AWGN channel, the complex envelope of the received signal in the presence of LO frequency offset can be modeled as

$$r(kT_s) = \exp\left(j(\omega kT_s + \phi)\right) \cdot$$
$$\exp\left(j\frac{\pi}{2}\sum_i a_i q(t - ikT_s)\right) + v(kT_s) \quad (13)$$

where ω and ϕ are the LO frequency and phase offset, respectively, and $v(kT_s)$ is the noise term.

We define a non-linear transformation of the signal as follows:

$$z[k] = (-1)^k r^2(kPT_s) \quad (14)$$

It is shown in [29], [30] that $z[k]$ can be approximated as a discrete time sine wave with a frequency that is twice the LO frequency offset of the original signal:

$$z[k] \approx A\exp\left(j(2\omega kPT_s + 2\phi)\right) + w[k] \quad (15)$$

where $w[k]$ is the noise term.

The estimation of the LO frequency offset turns to the estimation of a sine wave's frequency, which can be solved by several means. In our implementation, the frequency of $z[k]$ is estimated by evaluating the frequency that maximizes the amplitude of the DFT of $z[k]$:

$$\hat{f_0'} = \frac{1}{NPT_s}\arg\max_{-N/2 \leq m \leq N/2}\{|Z[m]|\} \quad (16)$$

where $Z[m]$ is the DFT of $z[k]$, and N is the number of points of $z[k]$ used for the frequency estimation. This frequency estimate is further refined by using an interpolated estimate for our frequency estimate:

$$\hat{f_1'} = \frac{1}{NPT_s}\left[\hat{m} + I^+\frac{|Z[\hat{m}+1]|}{|Z[\hat{m}+1]|+|Z[\hat{m}]|} - \right.$$
$$\left. (1-I^+)\frac{|Z[\hat{m}-1]|}{|Z[\hat{m}-1]|+|Z[\hat{m}]|}\right] \quad (17)$$

where $\hat{m}$ is the index obtained in (16), and I^+ is defined as

$$I^+ = \begin{cases} 1, & \text{if } |Z[\hat{m}+1]| > |Z[\hat{m}-1]| \\ 0, & \text{if } |Z[\hat{m}+1]| < |Z[\hat{m}-1]| \end{cases} \quad (18)$$

The LO phase offset is then estimated as follows. If $s[k]$ is the sampled GMSK message sent by the transmitter over an AWGN channel (which contains the message header $t[k]$ of length N), the receiver will receive the following sampled message:

$$r[k] = \exp\left(j(\omega kT_s + \phi)\right)s[k] + v[k] \quad (19)$$

where $v[k]$ is the noise term. The receiver first estimates the LO frequency offset as described above, which will yield the estimate of the LO frequency offset $\hat{f_1'}$. The receiver can then compensate the received message for the LO frequency offset as follows:

$$r'[k] = \exp\left(-j\left(2\pi\hat{f_1'}kT_s\right)\right)r[k] \quad (20a)$$
$$\approx \exp(j\phi)s[k] + v'[k] \quad (20b)$$

where $v'[k]$ is the noise term. The receiver then correlates the frequency-compensated received signal r' with the (known) message header $t[k]$:

$$\gamma[k] = \sum_{n=1}^{N} r'[k-n]t[n] \quad (21)$$

The amplitude of $\gamma[k]$ will be maximum when $r'[k]$ is aligned with the message header. The phase offset is then also given by the phase of $\gamma[k]$:

$$k_0 = \arg\max\{(|\gamma[k]|\} \quad (22a)$$
$$\hat{\phi} = \angle\{\gamma[k_0]\} \quad (22b)$$

REFERENCES

[1] H. Ochiai, P. Mitran, H. Poor, and V. Tarokh, "Collaborative beamforming for distributed wireless ad hoc sensor networks," *IEEE Trans. Signal Process.*, vol. 53, no. 11, pp. 4110–4124, Nov. 2005.
[2] R. Mudumbai, G. Barriac, and U. Madhow, "On the feasibility of distributed beamforming in wireless networks," *IEEE Trans. Wireless Commun.*, vol. 6, no. 5, pp. 1754–1763, May 2007.
[3] R. Mudumbai, D. Brown, U. Madhow, and H. Poor, "Distributed transmit beamforming: challenges and recent progress," *IEEE Commun. Mag.*, vol. 47, no. 2, pp. 102–110, Feb. 2009.
[4] R. Mudumbai, J. Hespanha, and U. Madhow, "Scalable feedback control for distributed beamforming in sensor networks," in *Proc. 2005 IEEE International Symp. Inf. Theory*.
[5] T. Cover and A. El Gamal, "Capacity theorems for the relay channel," *IEEE Trans. Inf. Theory*, vol. IT-25, no. 5, pp. 572–584, 1979.
[6] A. El Gamal and T. Cover, "Multiple user information theory," *Proc. IEEE*, vol. 68, no. 12, pp. 1466–1483, 1980.
[7] S. Borade, L. Zheng, and R. Gallager, "Amplify-and-forward in wireless relay networks: rate, diversity, and network size," *IEEE Trans. Inf. Theory*, vol. 53, no. 10, pp. 3302–3318, Oct. 2007.
[8] S. Jin, M. McKay, C. Zhong, and K.-K. Wong, "Ergodic capacity analysis of amplify-and-forward MIMO dual-hop systems," *IEEE Trans. Inf. Theory*, vol. 56, no. 5, pp. 2204–2224, May 2010.
[9] S. Gharan, A. Bayesteh, and A. Khandani, "Asymptotic analysis of amplify and forward relaying in a parallel MIMO relay network," *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 2070–2082, Apr. 2011.
[10] Y.-S. Tu and G. Pottie, "Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through AWGN channels," in *Proc. 2002 IEEE Veh. Technol. Conf. – Spring*, vol. 1, pp. 130–134.
[11] Z. Ding, W. Chen, and K. Leung, "Distributed beamforming and power allocation for cooperative networks," *IEEE Trans. Wireless Commun.*, vol. 7, no. 5, pp. 1817–1822, 2008.

[12] R. Mudumbai, B. Wild, U. Madhow, and K. Ramch, "Distributed beamforming using 1 bit feedback: from concept to realization," in *2006 Allerton Conf. Commun., Control, Comput.*

[13] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Distributed transmit beamforming using feedback control," *IEEE Trans. Inf. Theory*, vol. 56, no. 1, pp. 411–426, Jan. 2010.

[14] I. Ozil and D. Brown, "Time-slotted round-trip carrier synchronization," in *Proc. 2007 Asilomar Conf. Signals, Syst. Comput.*, pp. 1781–1785.

[15] D. R. Brown III and H. Poor, "Time-slotted round-trip carrier synchronization for distributed beamforming," *IEEE Trans. Signal Process.*, vol. 56, no. 11, pp. 5630–5643, Nov. 2008.

[16] R. Preuss and D. R. Brown III, "Retrodirective distributed transmit beamforming with two-way source synchronization," in *Proc. 2010 Conf. Inf. Sciences Syst.*.

[17] ——, "Two-way synchronization for coordinated multicell retrodirective downlink beamforming," *IEEE Trans. Signal Process.*, vol. 59, no. 11, pp. 5415–5427, 2011.

[18] R. Mudumbai, U. Madhow, D. Brown, and P. Bidigare, "DSP-centric algorithms for distributed transmit beamforming," in *2011 IEEE Asilomar Conf. Signals, Syst. Comput.*.

[19] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H.-P. Mayer, L. Thiele, and V. Jungnickel, "Coordinated multipoint: concepts, performance and field trial results," *IEEE Commun. Mag.*, vol. 49, no. 2, pp. 102–111, 2011.

[20] S. Munkyo, M. Rodwell, and U. Madhow, "A feedback-based distributed phased array technique and its application to 60-GHz wireless sensor network," in *Proc. 2008 IEEE MTT-S International Microwave Symp.*, pp. 683–686.

[21] S. Sigg and M. Beigl, "Algorithms for closed-loop feedback based distributed adaptive beamforming in wireless sensor networks," in *Proc. 2009 International Conf. Intelligent Sensors, Sensor Netw. Inf. Process.*, pp. 25–30.

[22] M. M. U. Rahman, H. E. Baidoo-Williams, S. Dasgupta, and R. Mudumbai, "Fully wireless implementation of distributed beamforming on a software-defined radio platform," in *2012 ACM/IEEE Conf. Inf. Process. Sensor Netw.*

[23] D. R. Brown III, P. Bidigare, and U. Madhow, "Receiver-coordinated distributed transmit beamforming with kinematic tracking," in *Proc. 2012 International Conf. Acoustics, Speech, Signal Process.*, to appear.

[24] F. Quitin, M. M. U. Rahman, R. Mudumbai, and U. Madhow, "Distributed beamforming with software-defined radios: frequency synchronization and digital feedback," in *Proc. 2012 IEEE Global Communications Conference.*

[25] F. Quitin, U. Madhow, M. M. U. Rahman, and R. Mudumbai, "Demonstrating distributed transmit beamforming with software-defined radios," in *Proc. 2012 International Symp. World Wireless, Mobile Multimedia Netw.*

[26] R. Mudumbai, P. Bidigare, S. Pruessing, S. Dasgupta, M. Oyarzun, and D. Raeman, "Scalable feedback algorithms for distributed transmit beamforming in wireless networks," in *Proc. 2012 IEEE International Conf. Acoustics, Speech, Signal Process.*, pp. 5213–5216.

[27] D. Rife and R. Boorstyn, "Single tone parameter estimation from discrete-time observations," *IEEE Trans. Inf. Theory*, vol. 20, no. 5, pp. 591–598, Sept. 1974.

[28] S. Kay, "A fast and accurate single frequency estimator," *IEEE Trans. Acoustics, Speech, Signal Process.*, vol. 37, no. 12, pp. 1987–1990, Dec 1989.

[29] M. Morelli and U. Mengali, "Feedforward carrier frequency estimation with MSK-type signals," *IEEE Commun. Lett.*, vol. 2, no. 8, pp. 235–237, Aug. 1998.

[30] H. Peng, J. Li, and L. Ge, "Non-data-aided carrier frequency offset estimation of GMSK signals in burst mode transmission," in *Proc. 2003 IEEE International Conf. Acoustics, Speech, Signal Process.*, vol. 4, pp. 576–579.

[31] C. Zucca and P. Tavella, "The clock model and its relationship with the Allan and related variances," *IEEE Trans. Ultrason., Ferroelect, Freq. Contr.*, vol. 52, no.2, pp. 289–296, 2005.

[32] USRP products. Available: http://www.ettus.com, 2012.

[33] Available: http://ece.ucsb.edu/~fquitin/.

François Quitin received the Ph.D. degree in Electrical Engineering from the Université Libre de Bruxelles (ULB), Brussels, Belgium and from the Université catholique de Louvain (UCL), Louvain-La-Neuve, Belgium in 2011. He is currently working as a post-doctoral researcher at the Wireless Communications and Sensornets Labs at the University of California, Santa Barbara (UCSB). He is the recipient of the 2012 Alcatel-Lucent Bell Scientific Award . His research interests focus on experimental and prototyping aspects in wireless communications, taking advanced theoretical ideas all the way to practice.

Muhammad Mahboob Ur Rahman (S'12) received his B.Sc. degree in Electrical engineering from the University of Engineering and Technology (UET), Lahore, Pakistan, in 2007. Since 2009, he is a Ph.D. student at University of Iowa, USA. Before joining University of Iowa, he has worked as a Lecturer/Lab engineer at Faisalabad campus of UET Lahore for two years (2007-2009).

Raghuraman Mudumbai is currently an assistant professor of electrical and computer engineering at the University of Iowa. He received a B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, in 1998 and an M.S. degree in electrical engineering from Polytechnic University, Brooklyn, New York, in 2000. He worked for two years at Ericsson Telephone Co. and returned to graduate school at UCSB in 2003, where he received a Ph.D. in electrical and computer engineering in 2007.

Upamanyu Madhow is Professor of Electrical and Computer Engineering at the University of California, Santa Barbara. His research interests broadly span communications, signal processing and networking, with current emphasis on millimeter wave communication and bio-inspired approaches to networking and inference. He received his bachelor's degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1985, and his Ph.D. degree in electrical engineering from the University of Illinois, Urbana-Champaign in 1990. He has worked as a research scientist at Bell Communications Research, Morristown, NJ, and as a faculty at the University of Illinois, Urbana-Champaign. Dr. Madhow is a recipient of the 1996 NSF CAREER award, and co-recipient of the 2012 IEEE Marconi prize paper award in wireless communications. He has served as Associate Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE TRANSACTIONS ON INFORMATION THEORY, and the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. He is the author of the textbook *Fundamentals of Digital Communication* (Cambridge University Press, 2008).