

DSP-CENTRIC ALGORITHMS FOR DISTRIBUTED TRANSMIT BEAMFORMING

Raghu Mudumbai* Upamanyu Madhow[‡] Rick Brown[◊] Patrick Bidigare[†]

* ECE Department, The University of Iowa, Iowa City IA 52242, rmudumbai@engineering.uiowa.edu

[‡] ECE Department, University of California, Santa Barbara, CA 93106, madhow@ece.ucsb.edu

[◊] ECE Department, Worcester Polytechnic Institute, Worcester, MA 01609, drb@wpi.edu

[†] Raytheon BBN Technologies, Arlington, VA 22207, bidigare@ieee.org

ABSTRACT

Distributed transmit beamforming is a means of increasing range and power efficiency via local collaboration among neighboring nodes in order to transmit a common message to a remote destination. While its basic feasibility has been established by recent analyses and prototypes, transitioning this concept to applications requires the development of protocols and architectures which can be implemented efficiently using digital signal processing (DSP). In this paper, we describe DSP-centric algorithms and their performance limits, and report on recent results from simulations and software-defined radio experiments.

Index Terms— distributed beamforming, synchronization, baseband algorithms

I. INTRODUCTION

Distributed transmit beamforming (DTB) is a wireless transmission technique where a group of transmitters organize into a *virtual antenna array* in order to cooperatively transmit a common message signal to a distant receiver. This technique is attractive because it allows nodes with simple omnidirectional antennas to collaboratively obtain the directivity (and associated energy efficiency benefits) of antenna arrays. For a cooperative array of N nodes with transmit power P per node, the received power at the destination is proportional to PN^2 . For free space transmission, a 10-node virtual array can provide range extension by a factor of 10, or, for fixed range, reduce power per node by a factor of 100, and total transmitted power by a factor of 10. This can have a dramatic effect on the size, weight and power (SWaP) of wireless devices and significantly extend the lifetime of systems operating off batteries.

The main challenge in realizing the large potential gains from distributed transmit beamforming is in precisely synchronizing the individual RF signals from each cooperating transmitter so that they are aligned in phase at the receiver. Since each node in a distributed array derives its RF carrier signal from separate local oscillators, it is necessary to estimate, track and compensate for offsets between these oscillators.

Related work: The problem of synchronizing transmitters for distributed transmit beamforming has attracted a great deal of attention over the last decade (see the survey article [1]). A number of synchronization techniques have been developed, including full-feedback closed-loop [2], 1-bit closed-loop [3], master-slave open-loop [4] synchronization, round-trip synchronization [5], and two-way synchronization [6], and the feasibility of these techniques has been demonstrated in several experimental prototypes [7], [8], [9], [10].

Thus, we now have a menu of synchronization techniques that represent different tradeoffs between overheads of coordination between the beamforming nodes, channel feedback from the receiver and implementation complexity as well as different sets of PHY assumptions such as Line-of-Sight channels and reciprocity. It is time, therefore, to focus research in this area towards the goal of realizing the benefits of distributed beamforming in practical wireless networks. To this end, it is important that the required RF synchronization functionalities be implementable on general-purpose signal processing platforms, without requiring custom RF hardware. It is also important that these synchronization functionalities be incorporated in network protocols with minimal overhead.

Contributions: In this paper, we observe that the synchronization functionalities for distributed beamforming can indeed be implemented in a DSP-centric baseband architecture, using off-the-shelf RF hardware. This is because, while RF signals at high frequencies cannot be directly digitized, the *relative frequency offsets* between two oscillators are small enough to process digitally. We discuss an illustrative DSP-centric architecture, based on a simple Kalman filtering framework for establishing and maintaining synchronization between oscillators in a robust manner with low overheads, using a simple phase estimation algorithm as a building block. We present simulation results, as well as the results of promising preliminary experiments on a software-defined radio testbed.

Outline: Section II motivates a DSP-centric model for synchronization using a simple dynamical model for phase drift between two oscillators. Section III illustrates the DSP-centric model with a simple Kalman filtering approach to

frequency and phase estimation and tracking. Simulation and experimental results showing the efficacy of this approach are presented in Section IV and Section V concludes.

II. BASEBAND SYNCHRONIZATION OF RF SIGNALS

We now introduce a baseband architecture for the synchronization of high frequency RF signals, which is suitable for DSP-centric implementation.

II-A. Signal model

Consider a distributed array with N nodes which seek to collaboratively transmit a common baseband message signal $m(t) = m_I(t) + jm_Q(t)$. Let the RF signal transmitted by node i be denoted by $x_i(t)$. Assuming narrowband signaling, the channel gain from transmitter i to the distant receiver can be represented as a complex scalar h_i . We further assume that the magnitude of the channel gain is the same for all transmitters¹; this gain can be set to unity without loss of generality so that we can write $h_i = e^{j\theta_i}$.

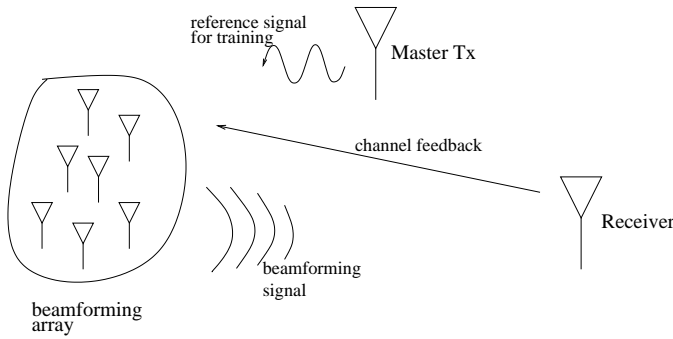


Fig. 1. General network model for distributed beamforming.

Our aim is to synchronize the oscillators on each of the nodes to a common reference signal $c_0(t) \doteq \cos(2\pi f_0 t)$ so that the RF signal transmitted by node i is $x_i(t) = \Re(m(t) \exp(j2\pi f_c t - j\theta_i))$, so that the overall received signal is (ignoring receiver noise for simplicity)

$$\begin{aligned} r(t) &= \sum_{i=1}^N \Re(h_i m(t) \exp(j2\pi f_c t - j\theta_i)) \\ &= N \Re(m(t) \exp(j2\pi f_c t)) \\ &\equiv N(m_I(t) \cos(2\pi f_c t) - m_Q(t) \sin(2\pi f_c t)). \end{aligned} \quad (1)$$

In (1), the factor of N represents the *beamforming gain* and shows the effect of constructive interference between the signals from each transmitter.

The challenge then is in constructing the signal $x_i(t)$ from each transmitter as above so that all the nodes are frequency

¹This is a reasonable assumption if we consider Line-of-Sight transmission from a cluster of transmitters to a distant receiver; however, the results of this paper do not depend on this assumption. We use it mainly to simplify the presentation.

locked and the phase of $x_i(t)$ is chosen to precisely cancel the effect of the channel phase θ_i . Note that in the above model, the transmitters may in general need to use the reference signal $c_0(t)$ at frequency f_0 to synthesize a carrier signal for beamforming at a completely different frequency f_c .

Different protocols for distributed beamforming have been developed that solve the above problem in different ways that represent various tradeoffs between in-network coordination, feedback from the receiver and so on. For instance, under beamforming schemes using a Master-Slave architecture [4], there is a designated master node that supplies the reference signal $c_0(t)$, whereas under round-trip synchronization schemes [5], the receiver itself implicitly provides the reference signal. A Master-Slave can also be used using a GPS signal as reference; in other words, a GPS satellite can act as an implicit Master node; however this requires uninterrupted access to a GPS signal and also additional hardware. We consider the general model of the beamforming process illustrated in Fig. 1.

Let us denote the RF carrier signals at the nominal frequency f_0 obtained from the free-running local oscillator on node i as $c_i(t) = \cos(2\pi(f_0 + \Delta f_i)t + \Delta\phi_i)$ and $s_i(t) = \sin(2\pi(f_0 + \Delta f_i)t + \Delta\phi_i)$ where Δf_i and $\Delta\phi_i$ are the frequency and phase offsets of node i 's oscillator with respect to $c_0(t)$.

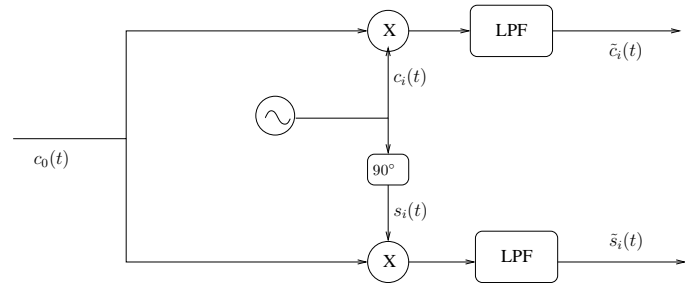


Fig. 2. Signal model for DSP-centric architecture.

Node i demodulates the reference signal $c_0(t)$ using $c_i(t)$, $s_i(t)$ as shown in Fig. 2 to get the pair of in-phase and quadrature baseband signals $\tilde{c}_i(t) \doteq \cos(\phi_i(t))$, $\tilde{s}_i(t) \doteq \sin(\phi_i(t))$, where $\phi_i(t) = 2\pi\Delta f_i t + \Delta\phi_i$ and $\Delta\phi_i$, Δf_i respectively represent the phase and frequency offsets of the local oscillator with respect to the reference signal. The offsets $\Delta\phi_i$, Δf_i can in general vary in time and we need to track these variations.

We use the time-slotted model for synchronization illustrated in Fig. 3. Training signals - including the reference signal $c_0(t)$ and possibly other feedback signals carrying information about the channel phase - are periodically retransmitted every T_{slot} time units in short bursts of duration T_{est} . The *duty cycle* of the training process $\eta \doteq \frac{T_{est}}{T_{slot}}$ can be quite small. The transmitters use the training signal in each

time-slot to update their estimate of the phase and frequency offsets $\Delta\phi_i$, Δf_i .




Fig. 3. Time-slotting model for synchronization.

II-B. Statistics of phase drift between two oscillators

The key idea behind the DSP-centric architecture is that, while the RF signals transmitted by the beamforming nodes are themselves not suitable for digital processing, the clock offsets between oscillators that are nominally set to the same frequency are relatively small. For instance, inexpensive crystal oscillators [11] have worst-case frequency deviations on the order of ± 10 parts per million of the nominal center frequency, which corresponds to offsets of no greater than ~ 10 kHz at frequencies ~ 1 GHz. Furthermore while these offsets vary in time, the time-variations are usually slow; typically the offsets remain roughly constant over time intervals on the order of several seconds. Thus, as long as we are working with *relative offsets* between two oscillators, the signal bandwidths are small enough and their time variations slow enough that they can be tracked and compensated in software with relatively small overheads. In other words, the baseband signals $\tilde{c}_i(t)$ and $\tilde{s}_i(t)$ from Fig. 2 have a small enough bandwidth to process digitally.

We now describe a simple model that is effective in capturing the statistics of time-variations in the clock offsets and is widely used in the literature on synchronization. This model assumes that the pairwise instantaneous frequency and phase both undergo independent Brownian motion increments. The frequency increments capture the effects of higher order correlations in the phase drifts and also motion-related Doppler effects.

Let $\mathbf{x}_i[k] \doteq [\Delta\phi_i \ 2\pi\Delta f_i]^T$ denote the offsets at a transmitter node i at time-slot k .

$$\mathbf{x}_i[k+1] = F\mathbf{x}_i[k] + w[k], \quad (2)$$

where $F = \begin{bmatrix} 1 & T_{slot} \\ 0 & 1 \end{bmatrix}$ and $w[k] \sim N(0, C)$ is a process noise whose covariance is modeled as

$$C = q_1^2 T_{slot} \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + q_2^2 T_{slot} \begin{bmatrix} \frac{T_{slot}^2}{3} & \frac{T_{slot}}{2} \\ \frac{T_{slot}}{2} & 1 \end{bmatrix},$$

where q_1 , q_2 represent the standard deviations of Brownian motion processes in phase and frequency respectively. In practice, these parameters can be estimated from measurements of Allan variances over different time-scales. Note that q_1 , q_2 have dimensions of $\text{sec}^{-\frac{1}{2}}$ and sec^{-1} respectively.

III. DSP-CENTRIC SYNCHRONIZATION

We now present a low-complexity approach to synchronization. We first present a simple algorithm to generate “one-shot” phase estimates i.e. estimates based on the reference signal in a single training epoch of duration T_{est} . These estimates can then be used as inputs to a simple Kalman filter based on the preceding state-space model to accurately estimate and correct the frequency and phase offsets of each node with respect to the reference signal $c_0(t)$.

III-A. Limits of frequency and phase estimation

We now consider the process of obtaining one-shot frequency and phase estimates using a noise-corrupted reference signal received by transmitter i over the training epoch of duration T_{est} in one time-slot. Let $a_i(t) = A_i c_0(t) + n_i(t)$, $t \in [0, T_{est}]$. The post-integration SNR of this signal is $\text{SNR} \equiv \frac{A_i^2 T_{est}}{2N_0}$, where N_0 is the power spectral density of the white noise process $n_i(t)$. After downconversion as shown in Fig. 2, we get the complex baseband signal that can be written using the notation introduced in Section II-A as

$$\begin{aligned} \tilde{a}_i(t) &\doteq A_i \left(\tilde{c}_i(t) + j\tilde{s}_i(t) \right) + \tilde{n}_i(t) \\ &= A_i \exp(j2\pi\Delta f_i t + j\Delta\phi_i) + \tilde{n}_i(t), \end{aligned} \quad (3)$$

where $\tilde{n}_i(t)$ denotes the baseband complex Gaussian noise. Given $\tilde{a}_i(t)$, $t \in [0, T_{est}]$, we want to estimate the frequency and phase offset Δf_i , $\Delta\phi_i$ of the local oscillator of transmitter i with respect to the reference signal $c_0(t)$.

The Cramer-Rao lower bounds for this one-shot phase and frequency estimation process are well-known in the literature [12], [13]. If ϕ_{err} and f_{err} respectively denote the one-shot phase and frequency estimation errors, we have

$$\begin{aligned} \sigma_\phi^2 &\doteq E[\phi_{err}^2] \geq \frac{2}{\text{SNR}} \\ \sigma_f^2 &\doteq E[f_{err}^2] \geq \frac{3}{2\pi^2 T_{est}^2 \text{SNR}}. \end{aligned} \quad (4)$$

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 $\phi_i(t) - \hat{\phi}_i(t)$ between the predicted phase offset $\hat{\phi}_i(t)$ and actual phase offset $\phi_i(t)$ of transmitter i with the reference signal grows with time and its value at the end of the time-slot can be written as

$$\begin{aligned} E\left[(\phi_i(t) - \hat{\phi}_i(t))^2 \right]_{t=T_{slot}} &= \sigma_\phi^2 + T_{slot}^2 (2\pi\sigma_f)^2 \\ &\geq \frac{2}{\text{SNR}} \left(1 + \frac{3}{\eta^2} \right). \end{aligned} \quad (5)$$

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 (5) dominates; this reflects the fact that *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

$$\text{var}(\tilde{f}) = \frac{2\sigma_\phi^2}{(2\pi T_{slot})^2} \geq \frac{1}{\pi^2 \text{SNR} T_{slot}^2}, \quad (6)$$

and this variance can be significantly smaller than the one-shot frequency variance σ_f^2 in (4). This suggests that we can dispense with one-shot frequency estimates altogether, and rely on averaging phase estimates over multiple time-slots to get good frequency estimates. Furthermore, when the estimation interval T_{est} is small, we can neglect the effects of the frequency offset on the training waveform itself; specifically when $\Delta f_i T_{est} \ll 1$, we can rewrite (3) as

$$a_i(t) = A_i \exp(j2\pi\Delta f_i t + j\Delta\phi_i) + \tilde{n}_i(t) \quad (7)$$

$$\approx A_i \exp(j\Delta\phi_i) + \tilde{n}_i(t). \quad (8)$$

The approximation in (8) allows us to obtain a very simple ML phase estimate

$$\hat{\phi}_i \doteq \angle \left(\int_0^{T_{est}} a_i(t) dt \right). \quad (9)$$

We can extend this idea of averaging one-shot phase estimates using a Kalman filtering approach which we describe next.

III-B. A two-state Kalman filter for phase prediction and tracking

The analysis in Section III-A motivates the following simple measurement model:

$$y[k] = H\mathbf{x}_i[k] + \nu[k], \quad (10)$$

where $y[k]$ is the one-shot phase estimate obtained in time-slot k , $H = [1 \ 0]$ and $\nu[k] \sim N(0, R)$ is a measurement noise representing estimation error in the phase estimation algorithm. If we assume the phase estimation achieves the Cramer-Rao lower bound in (4), we have $R = \frac{2}{\text{SNR}}$.

Using this measurement model together with the state-space model of the clock offsets presented in Section II-A, we can use a standard Kalman filter to obtain estimates of the state \mathbf{x} i.e. the frequency and phase offsets Δf_i , $\Delta\phi_i$. This estimate is updated every time-slot at the end of the training interval T_{est} , using the new one-shot phase estimate obtained during this time-slot.

During the rest of the time-slot, the state estimate from the previous training interval is used to *predict* the phase offset. During this phase, the prediction error has a component increasing linearly with time arising from non-zero frequency estimation error as in (5).

IV. RESULTS

We now present some results from simulation and experiment that show that the DSP-centric approach described in this paper is effective in achieving synchronization of RF carrier signals for beamforming.




Fig. 4. Phase and frequency estimation error over time.

Fig. 4 shows the phase and frequency estimation error from a simulation filter using a one-shot phase observation per time-slot as described in Section III. For this simulation, we used post-integration SNR of 20 dB to generate the one-shot phase estimates, an estimation interval of 10 ms, time-slots of 250 ms (and thus a training duty cycle of 4%), frequency offset of 0.1 ppm at a center frequency of 900 MHz, and oscillator phase noise parameters $q_1 = 5 \times 10^{-22} \text{ sec}^{-\frac{1}{2}}$ and $q_2 = 2 \times 10^{-23} \text{ sec}^{-1}$. It can be seen that in the first few time-slots when the frequency estimate from the Kalman filter has a large error, the phase prediction error becomes quite large over the time-slot. However, after about 2 seconds (corresponding to 8 time-slots, the Kalman filter estimate is sufficiently accurate to keep the phase error below 20° or so.

We have also implemented distributed beamforming based on the DSP-centric architecture for synchronization described in this paper on a software-defined radio platform. In this implementation we used baseband algorithms to lock oscillators at a set of slave transmitters to a reference

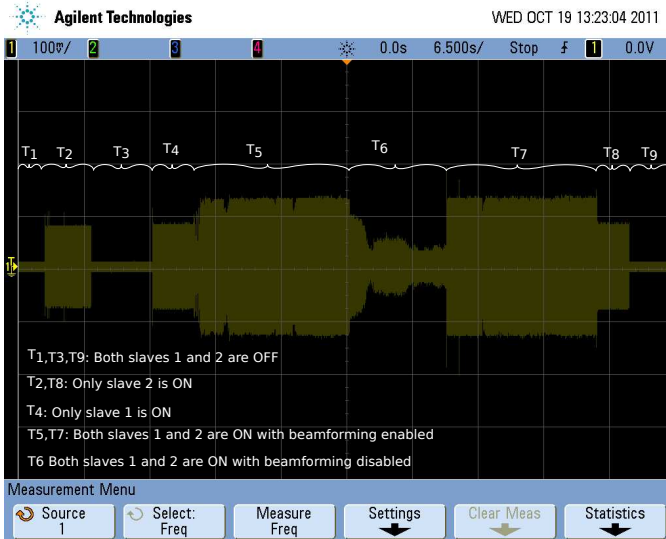


Fig. 5. Distributed beamforming with software-defined radios.

sinusoidal tone broadcast by a master transmitter, and used 1-bit SNR feedback [3] to steer the transmission towards the receiver. Fig. 5 shows an oscilloscope trace of the received signal strength from this setup with two slave transmitters; it can be seen that amplitude of the received signal when both slaves are transmitting is close to the sum of their individual received amplitudes, which shows coherence. These results show the practical feasibility of the DSP-centric approach i.e. it is possible to synchronize high frequency RF signals digitally in baseband and achieve beamforming gains on commodity platforms without any hardware modifications. Details of this experimental work will be reported in a forthcoming publication.

V. CONCLUSIONS

Modern communication transceivers leverage Moore’s law by employing DSP-centric architectures to perform sophisticated baseband operations prior to upconversion and after downconversion. As discussed in this paper, the synchronization required for distributed beamforming can also be implemented within such a framework. We have provided one example of DSP-centric design (more details can be found in [14], while fundamental limits are explored in [15]), in which explicit feedback based on phase measurements is used to drive frequency/phase Kalman tracking. While phase measurements must be individually performed for each transmitting node in the scheme considered here, alternative feedback schemes based on aggregate measurements also fit within a DSP-centric framework. The one bit feedback scheme in [3] is an example of such a technique, but richer forms of aggregate feedback can provide considerably better performance [16], especially in terms of tracking frequency

offsets as well as aligning phases. In addition, recent results in timing synchronization [17] show that timing accuracies within a fraction of a carrier cycle are achievable at relatively low complexity, which implies that it may be possible to achieve distributed beamforming even without such explicit feedback.

Such significant recent progress on DSP-centric synchronization sets the stage for further research on synchronization-enabled network protocols which allow clusters of nodes to cooperate not only to perform distributed transmit beamforming, but also distributed versions of other MIMO techniques such as nulling, SDMA and interference alignment. An important part of such a research agenda is to demonstrate the feasibility of such ideas using prototypes based on general-purpose platforms such as software-defined radios.

VI. REFERENCES

- [1] R. Mudumbai, D.R. Brown III, U. Madhow, and H.V. Poor, “Distributed transmit beamforming: Challenges and recent progress,” *IEEE Communications Magazine*, vol. 47, no. 2, pp. 102–110, February 2009.
- [2] Y.S. Tu and G.J. Pottie, “Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through AWGN channels,” in *IEEE Vehicular Technology Conf. (VTC)*, Birmingham, AL, Spring 2002, vol. 1, pp. 130–134.
- [3] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, “Scalable feedback control for distributed beamforming in sensor networks,” in *IEEE International Symp. on Information Theory (ISIT)*, Adelaide, Australia, September 2005, pp. 137–141.
- [4] R. Mudumbai, G. Barriac, and U. Madhow, “On the feasibility of distributed beamforming in wireless networks,” *IEEE Trans. on Wireless Communications*, vol. 6, no. 5, pp. 1754–1763, May 2007.
- [5] Ipek Ozil and D.R. Brown III, “Time-slotted round-trip carrier synchronization,” in *Proceedings of the 41st Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, November 4-7, 2007, pp. 1781 – 1785.
- [6] R.D. Preuss and D.R. Brown, “Retrodirective distributed transmit beamforming with two-way source synchronization,” in *Information Sciences and Systems (CISS), 2010 44th Annual Conference on*, march 2010, pp. 1 –6.
- [7] D.R. Brown, Boyang Zhang, B. Svirchuk, and Min Ni, “An experimental study of acoustic distributed beamforming using round-trip carrier synchronization,” in *Phased Array Systems and Technology (ARRAY), 2010 IEEE International Symposium on*, oct. 2010, pp. 316 –323.
- [8] R. Mudumbai, B. Wild, U. Madhow, and K. Ramchandran, “Distributed beamforming using 1 bit feedback:

from concept to realization,” in *44th Allerton Conf. on Comm., Control, and Computing*, Monticello, IL, Sep. 2006, pp. 1020 – 1027.

- [9] S. Sigg and M. Beigl, “Algorithms for closed-loop feedback based distributed adaptive beamforming in wireless sensor networks,” in *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2009 5th International Conference on*. IEEE, 2009, pp. 25–30.
- [10] Munkyo Seo, M. Rodwell, and U. Madhow, “A feedback-based distributed phased array technique and its application to 60-ghz wireless sensor network,” in *Microwave Symposium Digest, 2008 IEEE MTT-S International*, june 2008, pp. 683 –686.
- [11] “Datasheet for EC2620ETTS-64.000M oscillator,” <http://www.ecliptek.com/SpecSheetGenerator/specific.aspx?PartNumber=EC2620ETTS-64.000M>, 2011.
- [12] D. Rife and R. Boorstyn, “Single tone parameter estimation from discrete-time observations,” *Information Theory, IEEE Transactions on*, vol. 20, no. 5, pp. 591 – 598, sep 1974.
- [13] S. Kay, “A fast and accurate single frequency estimator,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 37, no. 12, pp. 1987 –1990, dec 1989.
- [14] D. R. Brown, P. Bidigare, and U. Madhow, “Receiver-coordinated distributed transmit beamforming with kinematic tracking,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012, to appear.
- [15] D. R. Brown III, R. Mudumbai, and S. Dasgupta, “Fundamental limits on phase and frequency tracking and estimation in drifting oscillators,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012, to appear.
- [16] R. Mudumbai, P. Bidigare, S. Pruessing, S. Dasgupta, M. Oyarzun, and D. Raeman, “Scalable feedback algorithms for distributed transmit beamforming in wireless networks,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012, to appear.
- [17] P. Bidigare, U. Madhow, R. Mudumbai, and D. Scherber, “Attaining fundamental bounds on timing synchronization,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012, to appear.