On the Feasibility of Distributed Beamforming in Wireless Networks

R. Mudumbai, *Student Member, IEEE,* G. Barriac, *Member, IEEE,* and U. Madhow, *Fellow, IEEE*

Abstract— Energy efficient communication is a fundamental problem in wireless ad-hoc and sensor networks. In this paper, we explore the feasibility of a distributed beamforming approach to this problem, with a cluster of distributed transmitters emulating a centralized antenna array so as to transmit a common message signal coherently to a distant Base Station. The potential SNR gains from beamforming are well-known. However, realizing these gains requires synchronization of the individual carrier signals in phase and frequency. In this paper we show that a large fraction of the beamforming gains can be realised even with imperfect synchronization corresponding to phase errors with moderately large variance. We present a master-slave architecture where a designated master transmitter coordinates the synchronization of other (slave) transmitters for beamforming. We observe that the transmitters can achieve distributed beamforming with minimal coordination with the Base Station using channel reciprocity. Thus, inexpensive local coordination with a master transmitter makes the expensive communication with a distant Base Station receiver more efficient. However, the duplexing constraints of the wireless channel place a fundamental limitation on the achievable accuracy of synchronization. We present a stochastic analysis that demonstrates the robustness of beamforming gains with imperfect synchronization, and demonstrate a tradeoff between synchronization overhead and beamforming gains. We also present simulation results for the phase errors that validate the analysis.

Index Terms— Distributed beamforming, synchronization, wireless networks, sensor networks, space-time communication.

I. INTRODUCTION

DISTRIBUTED beamforming has the potential of greatly enhancing energy efficiency in wireless networks. A group of cooperative transmitters can emulate an antenna array by transmitting a common message signal in such a way that the transmission is focused in the direction of the intended Base Station receiver (BS). For a narrowband message signal this can be arranged by adjusting the carrier phase of each transmitter in such a way that the individual transmissions combine coherently at the receiver. The energy efficiency gains from beamforming are well-known in the literature; if a single element antenna transmitting with power P_T achieves a received signal to noise ratio (SNR) of ρ_1, an array of N

R. Mudumbai and U. Madhow are with the Department of Electrical and Computer Engineering, University of California Santa Barbara, CA 93106 USA (e-mail: {raghu, madhow}@ece.ucsb.edu).

G. Barriac is with Qualcomm Inc., 9950 Barnes Canyon Rd, San Diego, CA 92121 USA (e-mail: gbarriac@qualcomm.com).
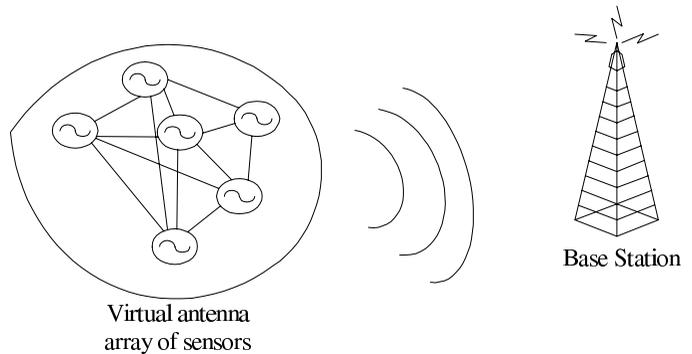
Fig. 1. Communication model for a sensor network.

identical antennas can use beamforming to achieve a SNR of $\rho_N = N\rho_1$ with the same total transmit power P_T, i.e. with each antenna transmitting with power $\frac{P_T}{N}$. Physically, this SNR increase arises from the increased directivity of the transmission from the antenna array; beamforming focuses N times more of the transmitted electromagnetic energy in the direction of the receiver. The biggest challenge in realizing these gains is the requirement of phase and frequency synchronization of the high-frequency carrier signals. In this paper, we present a master-slave protocol for synchronization, examine its performance using theory and simulation, and investigate the feasibility of distributed beamforming with imperfect synchronization.

While the ideas here may be of general applicability to different kinds of wireless networks, we focus on distributed beamforming in the context of a cluster of energy-constrained wireless sensor nodes communicating with a distant Base Station receiver (BS), as illustrated in Fig. 1. The main assumption is that local communication among the cooperating sensors is inexpensive compared to transmitting to the Base Station. Accordingly, we consider a *master-slave* architecture, where a designated *master* sensor coordinates the calibration and synchronization of the carrier signals of the other *slave* sensors. In this way, the sensors use cheap local communication between the *master* and the *slave* sensors to emulate a centralized antenna array, and to avoid the need for coordinating with the distant BS.

In a traditional (centralized) multi-antenna transmitter, one way to perform beamforming is by exploiting reciprocity to estimate the complex channel gains to each antenna element. These channel gains are computed in a centralized manner with reference to a RF carrier signal supplied by a local oscillator. However, in a distributed setting, each sensor has separate RF carrier signals supplied by separate local oscillator circuits. These carrier signals are not synchronized *a priori*. In the absence of carrier synchronization, it is not possible to estimate and pre-compensate the channel phase responses so as to assure phase coherence of all signals at the receiver.

It is encouraging to note that the achievable beamforming gains from imperfect phase synchronization are substantial. Consider the simple example of two equal amplitude signals from two transmitters combining at the BS with relative phase error of δ. The resulting signal amplitude is given by $\left|1 + e^{j\delta}\right| = 2\cos\left(\frac{\delta}{2}\right)$. Even a significant phase error of $\delta = 30°$ gives a signal amplitude of 1.93, which is 96% of the maximum possible amplitude of 2.0 corresponding to the zero phase error case. More generally, we show that it is possible to achieve SNR of up to 70% of the maximum with moderately large phase errors[1] on the order of $60°$.

The feasibility of distributed beamforming then depends on being able to keep the synchronization errors sufficiently small. We examine the different possible sources of phase error in detail. We observe that phase noise in practical oscillators causes them to drift out of synchronization, therefore, it is necessary for the master sensor to resynchronize the slaves periodically. This, combined with the duplexing constraints of the wireless channel (i.e. it is not possible to transmit and receive on the same frequency simultaneously), reveals a fundamental tradeoff between synchronization overhead and beamforming gain. We quantify this tradeoff using a stochastic model for the internal phase noise of oscillators.

There is now a growing body of research about cooperative transmission systems, including studies of distributed coding techniques for space-time diversity gains [1]. Diversity schemes do not offer average SNR gains, but rather reduce the probability of an outage event. Distributed diversity schemes are, therefore, of interest only in fading channels, and do not require coherent combining of signals. However, because typical diversity coding schemes [2] require a *baseband* channel that is constant at least over the length of a codeword, there is an implicit assumption of carrier *frequency* synchronization among the cooperating transmitters.

In contrast, beamforming offers SNR gains in both deterministic and fading channels, and in addition for fading channels reduces the probability of outage. However, distributed beamforming requires a globally consistent phase reference in addition to carrier frequency synchronization. Other authors have also independently considered the idea of using distributed transmitters as a *virtual antenna array* [3].

The performance of distributed beamforming has been previously studied from an information theoretic perspective [4], [5]. In [6], the authors studied the scalability of ad-hoc networks using distributed beamforming with a "Listen and Transmit" protocol. Further, they considered the effects of imperfect synchronization and showed that the scalability results still hold in the presence of synchronization errors. In [7], the authors propose a synchronization protocol that is suitable for coherent transmission. However, this requires significant coordination with the distant BS, and does not scale for a large number of transmitters. In recent work, the authors in [8] propose a distributed phase synchronization protocol for two transmitters. Most work in the literature on clock synchronization, (e.g. [9]) focuses on network synchronization, and is unsuitable for distributed beamforming.

The authors in [10] also studied the performance of distributed beamforming in sensor networks. They model the sensor locations as random and evaluate the effect of location (and phase) uncertainty on the average beampattern. These results are consistent with and complementary to the results in this paper. While [10] examines the beampattern averaged over all possible sensor placements, we consider some fixed placement and examine the synchronization process in detail.

In summary, most prior work on cooperative communication [10], [4], [1], [3] takes synchronization as a given. To the best of our knowledge, the protocol presented in this paper (which minimizes coordination with the BS), and the complementary protocol in our related work [11] (which utilizes feedback from the BS to minimize coordination among the transmitters), are the first detailed studies of the synchronization mechanisms required for large-scale cooperative communication.

The rest of this paper is organized as follows. Section II presents a simple analysis of the effect of imperfect synchronization on beamforming gain. This shows the high tolerance for phase errors. We present a master-slave architecture for synchronization in Section III, and derive a distributed protocol for beamforming based on this architecture. The main contributor to the residual phase error is oscillator noise, and Section IV offers a stochastic analysis of this noise, and identifies a tradeoff between the synchronization overhead and beamforming gain. Section V concludes.

II. ANALYSIS OF BEAMFORMING GAIN

We consider a cluster of N sensors, communicating a common (baseband) message signal $m(t)$ to a distant Base Station receiver, by modulating $m(t)$ with a carrier signal at frequency f_c. Each sensor derives its carrier signal from a separate local oscillator, therefore, the carrier signals of the different sensors are not initially synchronized to each other. Therefore, an explicit synchronization process is necessary. Before we present our algorithm for carrier synchronization, we show using a simple analysis that beamforming gains are robust to moderately large phase errors. For this section we assume that the synchronization algorithm allows each sensor to obtain synchronized carrier signals at frequency f_c and an estimate of their own channel gain to the BS. Using this the sensors can cooperatively transmit the message $m(t)$ by beamforming, just like a centralized antenna array. The resulting received signal $r(t)$ is the superposition of the channel-attenuated transmissions of all the sensors and additive noise $n(t)$:

$$r(t) = \Re\left(m(t)e^{j2\pi f_c t}\sum_i |g_i h_i| e^{j\phi_i(t)}\right) + n(t) \qquad (1)$$

where g_i is the pre-amplification and $\phi_i(t)$ is the cumulative phase error from the synchronization process for slave i. Under a constraint on the total transmit power, the optimum $|g_i| \equiv |h_i|$.

The phase errors have two effects on the received signal: a reduction in the average SNR, and a time-dependent fluctuation of the received phase. The latter effect may cause limitations in the coherent demodulation of digital signals. However, there are several methods, e.g. differential modulation, available to deal with these fluctuations provided the

---

[1]The phase errors are the result of random noise, and are, therefore, random variables. In this paper we use the term "large phase error" to indicate a phase error distribution with a large root-mean squared error.

time-variations are not too rapid. In this paper, we concentrate on the first effect i.e. the reduction in average SNR. This is appropriate for power-limited sensor networks, where the feasible communication range is limited by SNR. For simplicity of notation, we suppress the time-dependence of $\phi_i(t)$ in this section.

We model the channel coefficients h_i, $i = 1...N$, as independent circularly symmetric complex normal random variables with zero mean and unit variance, as denoted by $h_i \sim CN(0, 1)$. This is an appropriate model for a non-LoS wireless channel. This allows us to evaluate the variation of beamforming gain in fading channels i.e. outage mitigation. We also assume that the phase errors ϕ_i are independent and identically distributed random variables for all the sensors i.

Equation (1) motivates as our figure of merit, the beamforming gain defined as the normalized received power P_R, given that the total transmit power is $P_T = 1$:

$$P_R = \frac{1}{N} \left\| \sum_{i=1}^{N} |h_i|^2 e^{j\phi_i} \right\|^2 \qquad (2)$$

Proposition 1: $\frac{1}{N}P_R \to \left(\beta_\phi\right)^2$ *a.s.* as $N \to \infty$, where $\beta_\phi = E[\cos\phi_i]$ and *a.s.* denotes almost sure convergence. In other words, when the total transmit power is kept a constant, the received signal power increases linearly with N as $N \to \infty$.

Note that when there are no phase errors, i.e. $f_\phi(\phi_i) = \delta(0)$, then $\frac{1}{N}P_R \to 1$ *a.s.*

Proposition 2: For finite N, $E[P_R] = 1 + (N-1)\left(\beta_\phi\right)^2$. Thus, even for finite N, the expected value of the received signal power increases linearly with N. (β_ϕ is defined as in Proposition 1, i.e. $\beta_\phi = E[\cos\phi_i]$.)

In the absence of phase errors, Proposition 2 gives that $E[P_R] = N$.

Proposition 3: When N is large enough for the central limit theorem to apply,

$$P_R \approx X_c^2 + X_s^2 \qquad (3)$$

where $X_c \sim N(m_c, \sigma_c^2)$, $X_s \sim N(0, \sigma_s^2)$, and the parameters m_c, σ_c^2, and σ_s^2, are given as follows:

$$\begin{aligned} m_c &= \sqrt{N}E[\cos(\phi_i)] \\ \sigma_c^2 &= 2E[\cos^2(\phi_i)] - E[\cos(\phi_i)]^2 \\ \sigma_s^2 &= 2E[\sin^2(\phi_i)] \end{aligned} \qquad (4)$$

The variance of the received signal power is then

$$var[P_R] = 4\sigma_c^2 m_c^2 + 2\sigma_c^4 + 2\sigma_s^4 \qquad (5)$$

which increases linearly with N.

When there are no phase errors, (5) reduces to $var[P_R] = 4N$.

(Refer to the Appendix for a proof of these results.)

Proposition 2 implies that as long as the distribution of phase errors is such that $\beta_\phi \equiv E\left(\cos\phi_i\right)$ is close to 1, large gains can still be realized using distributed beamforming.
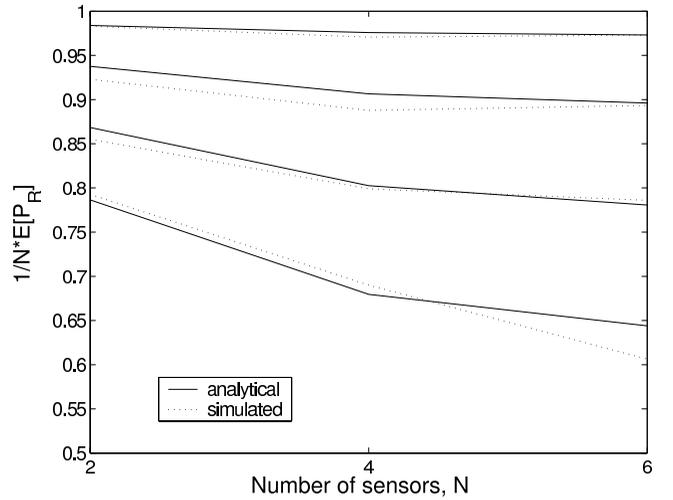
Fig. 2. $E[P_R]/N$ vs N, empirical and analytical results. The four sets of curves are for (top to bottom), $\Delta = 0.1 : 0.1 : 0.4$.

We now present some numerical results comparing the above analytical model with Monte-Carlo simulations performed using SIMULINK. We assume that the sensors transmit a binary pulse train modulated by BPSK, with a bit-rate small compared to the carrier frequency:

$$m(t) = \sum_k p(t - kT)s_k \qquad (6)$$

where $\{s_k\}$ is the BPSK symbol stream, and $p(t)$ is the transmitted pulse. The average power of the pulse $p(t), t = 0..T$ is normalized to $\frac{1}{N}$ and $E[|s_k|^2] = 1$ so that the total power transmitted by all the sensors is $P_T = 1$. Further we assume that the phases ϕ_i are distributed uniformly in the range $(-\Delta\pi, \Delta\pi)$.

Figure 2 shows the variation of average beamforming gain normalized to the maximum possible: i.e. $\frac{E(P_R)}{N}$ against the phase error parameter Δ. We find that beamforming gains of more than 70% of the maximum are possible with phase errors as large as of $60°$. In other words, the term β_ϕ decreases very slowly with the parameter Δ, which leads to the key conclusion that the beamforming gains are robust to moderately large phase errors.

While Fig. 2 shows the *average* beamforming gain, the actual beamforming gain is a random variable. We now look at the variation of the SNR with the phase errors uniformly distributed as above. Histograms of P_R, calculated using the Normal approximations as in Proposition 3, are shown in Fig. 3 where $\Delta = 0.1$ and $N = 10 : 10 : 40$.

The histograms in Fig. 3 show increased averaging for larger numbers of transmitters. This is expressed quantitatively in Proposition 3, which shows that while the mean of P_R is proportional to N, the standard deviation is proportional to $\sqrt{N}$ i.e. the fractional deviation $\frac{\sqrt{var(P_R)}}{P_R}$ decreases with increasing N. This means that the probability of an outage event e.g. where the received SNR is smaller than 70% of its mean, decreases with increasing N, showing that beamforming has the effect of mitigating fading. This is true for perfect and imperfect synchronization. Of course, the existence of phase errors can only increase the variance over that of an ideal, error free system.

Fig. 3. Histograms of P_R, $\Delta = 0.1$.

Fig. 4. Master-Slave architecture for carrier synchronization.
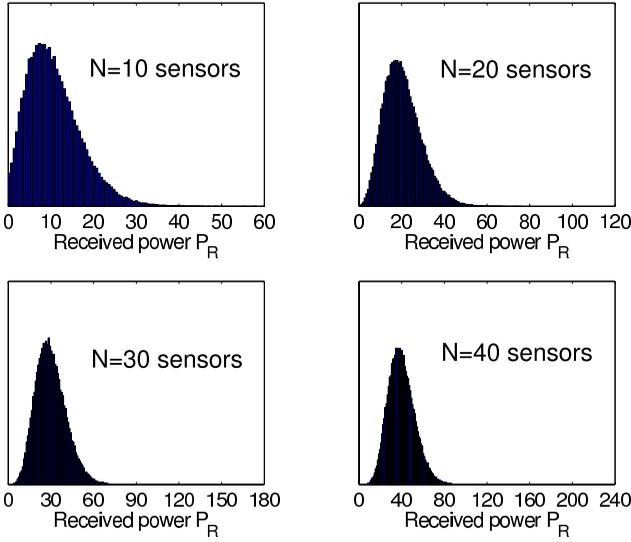
III. A MASTER-SLAVE ARCHITECTURE FOR BEAMFORMING

In this section, we present a protocol for achieving carrier phase synchronization based on a master-slave architecture. This is a multi-step process, and each step contributes to the overall phase error $\phi_i(t)$ that limits the beamforming gain. We now look at each step of the synchronization in detail.

The idea behind the protocol is illustrated in Fig. 4. The master sensor has a local oscillator which generates a sinusoid $c_0(t)$:

$$c_0(t) = \Re(\tilde{c}_0(t)), \text{ where } \tilde{c}_0(t) = e^{j(2\pi f_c t + \gamma_0)} \qquad (7)$$

that serves as the reference signal for the network. The master sensor broadcasts $c_0(t)$ to all the slaves. We assume that the local communication channel between master and slave sensors has a large SNR and ignore the receiver noise in this channel. After reception and amplification, the slave sensor i receives the signal broadcast by reception and amplification, the slave sensor i receives the signal broadcast by the master as:

$$c_{i,0}(t) = \Re(\tilde{c}_{i,0}(t)), \text{ where } \tilde{c}_{i,0}(t) = A_{i,0} e^{j(2\pi f_c t + \gamma_0 - \gamma_i)} \qquad (8)$$

where γ_i is the phase shift between the master and slave. $A_{i,0}$ is the amplitude of the received signal, its precise value is unimportant to the phase synchronization process (as the PLL is only sensitive to its phase). We simply set the term $A_{i,0}$ to unity, and the constant γ_0 to zero for simplicity.

The sensor i uses this signal $c_{i,0}(t)$ from (8) as input to a second-order phase locked-loop, driven by a VCO with a quiescent frequency close to f_c. From PLL theory [12], we can show that the steady-state phase error between VCO output and $c_{i,0}(t)$ is zero, and therefore, the steady-state VCO output can be used as a carrier signal consistent across all sensors - provided that the offset γ_i can be corrected for.

The phase offset γ_i is the total phase shift between the master sensors' reference oscillator signal $c_0(t)$, and the input signal at the slave sensors' PLL to which the slave VCO is synchronized in steady-state. One contribution to γ_i is from
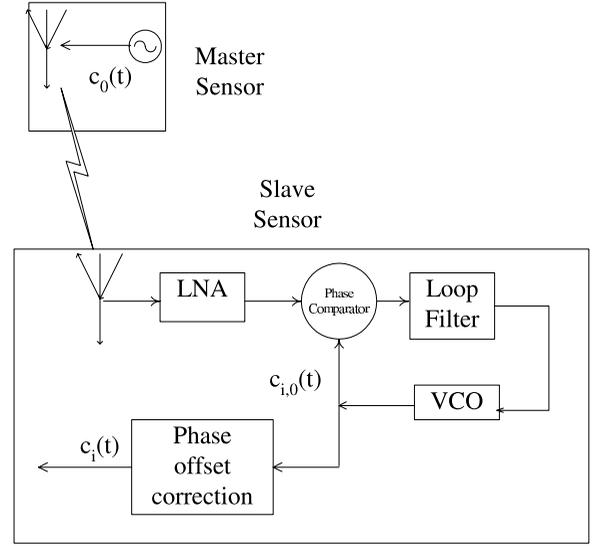
the phase response of the RF amplifiers at the master and slave sensor. These offsets are fixed and precisely known, and therefore, can be corrected for. However, the propagation delay of the wireless channel between master and slave also contributes to γ_i. This contribution can be characterized by an effective channel length d_i as $\gamma_i = \frac{2\pi f_c d_i}{c}$.

Unfortunately, for the high-frequency RF carriers typical of wireless networks, even a small uncertainty in channel length d_i causes substantial phase uncertainty e.g. at $f_c = 1.0$ GHz, the wavelength of the transmission is 30 cm, and an uncertainty of 15 cm in the channel length causes an uncertainty of $180°$ in γ_i. If left uncorrected this is disastrous for distributed beamforming, because a $180°$ offset would change constructive interference between transmitters into destructive interference. In centralized antenna arrays, the array elements are arranged in a known geometry, and therefore, the offset for each element can be precisely computed. This is not a reasonable assumption for ad-hoc and sensor networks considered in this paper. Thus it is necessary to develop methods to explicitly measure and correct for this unknown offset. Fortunately, if the sensors are not moving relative to each other, this offset stays roughly constant for significant time intervals, and therefore, frequent recalibration is not required. In Section III-A, we describe a protocol for performing this calibration, based on each slave sensor transmitting their frequency-locked carrier signal $c_{i,0}(t)$ back to the master sensor. We now sketch the process of channel estimation, and the algorithm for distributed beamforming assuming that slave i has an estimate $\hat{\gamma}_i = \gamma_i + \phi_i^e$ of its phase offset, where ϕ_i^e is the estimation error in the phase calibration. Slave i then has the calibrated carrier signal $c_i(t)$, which it uses to perform channel estimation and beamforming:

$$c_i(t) = \Re(\tilde{c}_i(t))$$
$$\text{where } \tilde{c}_i(t) = \tilde{c}_{i,0}(t) e^{j\hat{\gamma}_i}$$
$$= e^{j2\pi f_c t + j\phi_i^e} \qquad (9)$$

So far the synchronization process has been coordinated within the sensor network by the master sensor without
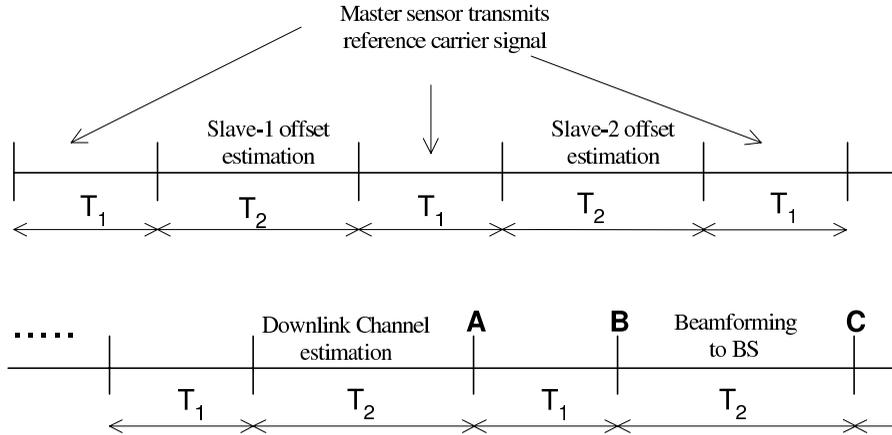
Fig. 5.   The Time-Division Duplexing constraint.
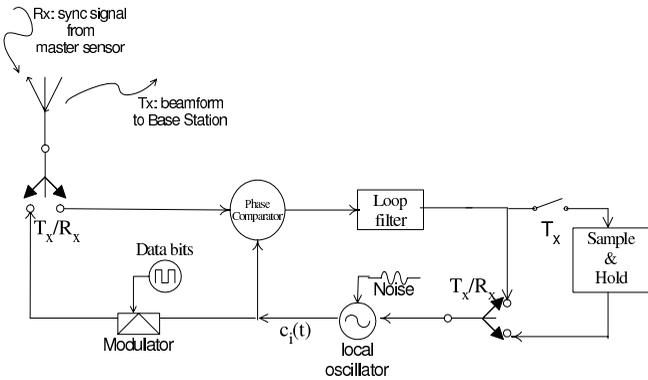
Fig. 6.   Schematic of a slave sensor.

requiring any interaction with the BS. In order for the sensors to beamform towards the BS, some information about the direction of the BS, or more precisely the channel response to the BS is required. Using channel reciprocity allows us to achieve this with only a minimum interaction with the BS. Specifically the BS broadcasts an unmodulated carrier signal $g(t)$:

$$g(t) = \Re(\tilde{g}(t)) = \Re(e^{j2\pi f_{c,0}t+\phi_0}) \qquad (10)$$

Each sensor independently demodulates its received signal $g_i(t) = \Re(h_i \tilde{g}(t))$ using $c_i(t)$ to obtain an estimate $\hat{h}_i$ of its own complex channel gain h_i to the receiver (for a narrowband message signal, the linear time-invariant channel to BS is represented as a scalar complex gain). More precisely, the channel estimate $\hat{h}_i$ is obtained by the sensor i by demodulating the received carrier signal $g_i(t)$ using $c_i(t)$, and sampling the result at some fixed time t_h. Note that while the sensor nodes have a mutually consistent carrier signal, the BS's carrier has not been explicitly synchronized to the master sensor's reference carrier, and therefore, would not be at the same frequency as the sensors. Letting $f_{c,0} = f_c + \Delta f$ we have:

$$\hat{h}_i = h_i \cdot e^{j(\phi_0 - \phi_i^e + \phi_h)} \qquad (11)$$

where $\phi_h = 2\pi \Delta f \, t_h$. We observe that the term ϕ_0 is just a constant scaling term and adds no *relative* phase errors between sensors. Similarly the term ϕ_h adds no relative phase error so long as the sampling term t_h is identical for all sensors. If the sampling times are off due to timing errors

τ_i, we get an effective phase noise: $\phi_i^h = 2\pi \Delta f \tau_i$. Therefore, we rewrite (11) as:

$$\hat{h}_i = C \cdot h_i \cdot e^{j(-\phi_i^e + \phi_i^h)} \qquad (12)$$

where C is a (complex) scaling constant that has no impact on the beamforming process. For simplicity, we take $C = 1$.

The sensors now use the synchronized carrier signal $c_i(t)$, and the channel estimate $\hat{h}_i$ to modulate the message signal for beamforming. The slave sensors obtain their carrier signal from the VCO that is synchronized to the reference signal from the master sensor, however, it is not possible for the slave sensors to receive a synchronization signal from the master sensor, while they are transmitting. Therefore, the VCOs of the slave sensors need to operate in an open-loop mode as shown in Fig. 6, while the slave sensors are transmitting. While in the open-loop mode, the slave's carrier signals obtained from the VCO undergoes uncompensated phase drift because of internal oscillator noise, and over time, the different slave carriers drift out of phase. This motivates the time-division duplexed mode of operation shown in Fig. 5, where the master sensor periodically transmits a reference carrier signal to resynchronize the slave carriers, to keep the total phase error bounded. The phase noise can be considered as a cyclostationary random process with period $T = T_1 + T_2$, and we analyze it in detail in Section IV. The noisy carrier signal used by the slave sensor i for modulation can be written as:

$$c_i^o(t) = \Re(\tilde{c}_i^o(t))$$
$$\text{where } \tilde{c}_i^o(t) = \tilde{c}_i(t)e^{j\phi_i^d(t)}$$
$$= e^{j2\pi f_c t+j\phi_i^e+j\phi_i^d(t)} \qquad (13)$$

$\phi_i^d(t)$ represents the uncompensated VCO drift when slave i is transmitting. After modulation by the carrier signal $c_i^o(t)$, slave sensor i applies a complex amplification $\hat{h}_i^*$ to compensate for the channel, and transmits the signal:

$$s_i(t) = \Re(\tilde{s}_i(t))$$
$$\text{where } \tilde{s}_i(t) = \hat{h}_i^* m(t)\tilde{c}_i^o(t) \qquad (14)$$

The received signal at the BS is then given by

$$r(t) = \Re\Big(\sum_i h_i \tilde{s}_i(t) + n(t)\Big)$$

$$= \Re\Big(m(t) \sum_i h_i \hat{h}_i^* \tilde{c}_i^o(t)\Big)$$

$$= \Re\Big(m(t) \sum_i |h_i|^2 e^{j2\pi f_c t - j\phi_i^h + j2\phi_i^e + j\phi_i^d(t)}\Big). \quad (15)$$

Comparing (15) with (1), we have for the total carrier phase error

$$\phi_i(t) = -\phi_i^h + 2\phi_i^e + \phi_i^d(t) \quad (16)$$

Equation (16) shows the different contributions to the total phase error in the received signal at the BS. In Section IV we look at the phase error in detail; we argue that the dominant component is the drift term $\phi_i^d(t)$, and show quantitatively how it affects the total beamforming gains.

A. Closed-Loop Method for Carrier Phase Calibration

In this section, we propose a flexible method for carrier phase calibration, where the master sensor measures the round-trip phase offset, and uses it to estimate the unknown phase offset γ_i from (8) for each slave, assuming symmetry in the forward and reverse channels to the slave nodes. The flexibility of this method comes at the price of complexity, and the necessity of synchronizing each of the slaves individually. However, the calibration process has to be repeated only when the RF channel between the master and slave sensor changes, therefore, the overhead from this process is small.

Remark: In the ideal case where the relative positions of the master and slave sensors as well as any multi-path scatterers do not change, the calibration process has to be performed only once (at startup time). In practice, wireless channels are not perfectly static: mobile scatterers and physical changes in the medium may change the channel phase response even when the sensors are stationary. Therefore, it makes sense to recalibrate the slave sensors periodically to track the channel changes. Fortunately, the channel variations are slow compared to the channel transmission times, and the robustness benefits of this periodic recalibration (e.g. every 100 seconds) outweigh the small extra overhead.

Fig. 7 illustrates the process of round-trip phase offset estimation. The basic idea is for the slave sensor i to transmit back to the master sensor the (uncompensated) VCO signal $c_{i,0}(t)$ represented in (8). The symmetry of the forward and reverse master-slave channels imply that the signal $c_{i,1}$ at the master sensor can be written as:

$$c_{i,1}(t) = A_{i,1} \cdot \Re\big(e^{j(2\pi f_l t + \gamma_0 - 2\gamma_i)}\big) \quad (17)$$

where $A_{i,1}$ is the received signal amplitude at the master sensor ($A_{i,1}$ is equal to $A_{i,0}$ by symmetry, but the actual value is not relevant to the phase noise, therefore, $A_{i,1}$ is set to unity for the discussion). Estimating the phase difference between $c_{i,1}(t)$ from (17) and $c_0(t)$ from (7) gives:

$$\Delta\phi_i = \big(2\gamma_i \bmod 2\pi\big) \quad (18)$$

Given a measured value of $\Delta\phi$, we have the estimated value of the offset γ_i:

$$\hat{\gamma}_i = \frac{\Delta\phi}{2} \quad (19)$$
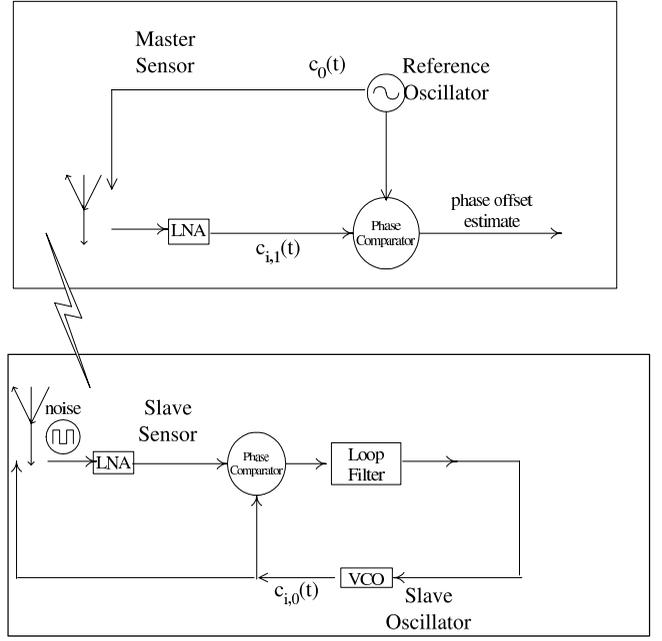
Fig. 7. Round-trip phase calibration.

Remark: There is one subtlety that needs to be noted here: the round-trip measurement of phase offset as in (19) leaves a 180° ambiguity in γ_i. In other words, by measuring $\Delta\phi_i$ we cannot distinguish between γ_i and $\gamma_i + 180°$. While it is possible to resolve this ambiguity by exchanging another set of messages between master and slave i, it turns out that a 180° phase difference does not affect the beamforming process. The reason is that the same carrier $c_i(t)$ is used by slave i for both channel estimation and distributed beamforming, and as (16) shows, the two ambiguities cancel each other.

B. Discussion

The time-division duplexing requirement for the master-slave link is the most important constraint of the synchronization protocol of Section III. Other authors [8] have considered using two frequencies to avoid this problem, with one frequency f_1 reserved for the master-slave link and the slave sensors beamforming to the BS on a completely different frequency f_2. In such schemes, the slave sensors use a frequency dividing PLL to obtain a carrier signal at frequency f_2 as $f_2 = \frac{m}{n}f_1$, where m and n are integers. Under this scheme the slave PLLs do not need to be open-loop while transmitting, therefore, an interesting question is whether such a frequency division duplexed (FDD) architecture can eliminate the problem of uncompensated carrier drift.

Unfortunately, the frequency divider introduces a phase ambiguity of integer multiples of $\frac{2\pi}{n}$ in the derived carrier signal. While it may appear that a constant phase ambiguity can be estimated and corrected for in a one-time calibration process, closer analysis shows that such phase ambiguities may also occur during the dynamical operation of the PLL, e.g. due to cycle slips [12]. Therefore, periodic recalibration is still necessary even with a FDD architecture, and we conjecture that a tradeoff between the synchronization overhead and the achievable beamforming gain still applies in this case.

IV. ANALYSIS OF PHASE ERROR

So far, we have described a protocol for carrier synchronization and beamforming, while enumerating the different sources of phase errors ϕ_i^e, ϕ_i^h and $\phi_i^d(t)$. Of the three different sources of error, ϕ_i^e and ϕ_i^h are constant calibration errors, whereas $\phi_i^d(t)$ is a time-varying noise term that arises from oscillator drift. Theoretically, we could perform carrier phase calibration and channel estimation several times independently and reduce the error terms ϕ_i^e and ϕ_i^h to arbitrarily small levels. However, the drift term $\phi_i^d(t)$ represents an irreducible phase error. Therefore, we consider this as the dominant cause of performance degradation and we now develop a stochastic model to characterize it.

The previous discussion in Section III motivated the time-division duplexed (TDD) mode of operation as shown in Fig. 5, where the slave sensors alternate between *sync* and *transmit* timeslots. The timeslots T_1 where the slave sensor synchronize to the master is a synchronization overhead, therefore, it is desirable to keep it small relative to the useful timeslots T_2. T_1 is determined by the settling time of the slave PLL, and T_2 is determined by the maximum admissible phase error, and the statistics of oscillator phase noise. By tolerating larger phase error, we are able to make T_2 higher and thereby reduce the synchronization overhead. We show in Section II that the SNR gains from beamforming are robust to moderately large phase errors. In the remainder of this section, we offer a quantitative analysis of this tradeoff using a stochastic model for oscillator phase noise.

Consider the PLL of the slave sensor as shown in Fig. 6. We use a loop-filter with one pole to obtain a second-order PLL with the closed-loop transfer function[12]:

$$H(s) = \frac{s^2}{s^2 + 2\xi\omega_n s + \omega_n^2} \quad (20)$$

where ω_n is the *natural frequency* and ξ is the *damping ratio* of the loop. By standard PLL theory, the steady state phase error of a second-order PLL is zero, and if we require a $90°$ phase margin, then we need a damping ratio of at least $\xi = 1.0$, and the settling then we need a damping ratio of at least $\xi = 1.0$, and the settling time (defined as the time required for the phase error to decrease to less than a given small fraction, say $\rho = 1\%$ of the initial error) is $T_s \approx \frac{4}{\omega_n}$. Since the synchronization timeslot $T_1 \geq T_s$, in order to minimize overhead we want to make T_s as small as possible. However, we observe that the loop has a low-pass frequency response with approximate bandwidth of ω_n, therefore, increasing ω_n also increases the phase noise. At the end of the *sync* timeslot, the loop-filter output is sampled and the VCO input is held to this value for the duration of the *transmit* timeslot. The phase error process in the transmit timeslot determines the achievable beamforming gain.

A. Stochastic model for the phase noise process

In order to study this more quantitatively, we assume that the PLL input signal from the master sensor (in the synchronization timeslot) is noiseless, and the only source of phase error is internal phase noise $\phi_i^d(t)$ in the slave sensors'
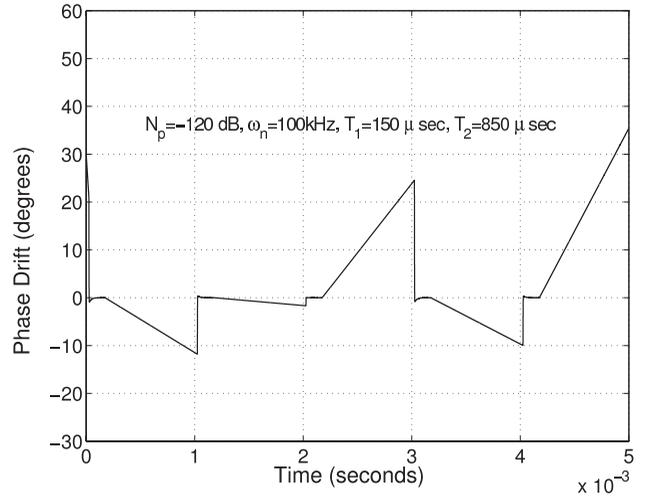
Fig. 8. Simulation of oscillator phase drift.

local oscillator signal:

$$c_i(t) = \Re\big(e^{j2\pi f_c t + j\phi_i^d(t)}\big) \quad (21)$$

(We also assume that the PLL phase drift is always small enough to allow the use of a linearized model.)

The traditional way to measure phase noise is by specifying its root-mean squared frequency deviation and Allan variance [12]. However, these measures are most useful if the noise process is stationary in time. In our case the drift process $\phi_i^d(t)$ is not stationary; in the *transmit* timeslot, the dominant phase noise contribution is from a random residual frequency offset that causes the phase error to increase linearly in time until the next *sync* timeslot (see Fig. 8). Therefore, the statistics are more appropriately modeled as *cyclostationary* with the period $T = T_1 + T_2$. We use a more fundamental approach to model this process.

In our model, the phase error in the oscillator in closed-loop (i.e. in the *sync* timeslot) consists of two components: a decaying transient of the initial phase offset, and a phase noise internal to the oscillator. The phase error in the free-running oscillator (i.e. in the *transmit* timeslot) has those two components and an additional linear phase drift. The linear drift arises because the VCO frequency set by the sample-and-hold (see Fig. 6) may have a small but non-zero offset from the reference frequency f_l. The oscillator internal phase noise is modeled as a wideband (white) Gaussian noise process with spectral density N_p. While phase noise in practical oscillators may also have other types of spectral densities e.g. flicker noise and random-walk noise, white Gaussian phase noise represents a worst case in terms of large instantaneous frequency deviations, because of the power in the high frequencies. Let N_p be the normalized spectral density defined such that the total power of the phase noise is $N_p\omega_n$. In other words, a white Gaussian phase noise with spectral density N_p will have the same power in a system of bandwidth ω_n as the oscillator's total internal phase noise.

Since the phase error process $\phi_i^d(t)$ is a zero-mean Gaussian process at all times, therefore, we characterize its statistics by computing its variance at the key time instants labelled A, B

and C in Fig. 5. Let the random phase values at these instants be denoted as ϕ_A, ϕ_B and ϕ_C, and their standard deviations as σ_A, σ_B and σ_C respectively. By the cyclostationarity of $\phi_i^d(t)$, $\sigma_C \equiv \sigma_A$. Using the linearity of the PLL's phase response, we can write the phase at time B as the superposition of the deterministic decay of the initial error ϕ_A, to a small fraction ρ of its starting value (ial error ϕ_A, to a small fraction ρ of its starting value (e.g. $\rho = 1\%$), and a noise term:

$$\phi_B = \rho\phi_A + \psi_1 \tag{22}$$

ϕ_B is small by design, and its variance can be written as:

$$\sigma_B^2 = \rho^2\sigma_A^2 + N_p\omega_n \tag{23}$$

In addition to the small phase error, at time instant B, the VCO input is sampled to set the VCO frequency for the *transmit* timeslot. The sampled value has a random offset Δf from the reference carrier frequency, and this offset consists of a transient term and a noise term:

$$\Delta f = \rho\omega_n\phi_A + \omega_n\psi_3$$
$$\text{Therefore } \sigma_f^2 = \rho^2\omega_n^2\sigma_A^2 + \omega_n^3 N_p \tag{24}$$

We have for the evolution of the phase between time instants B and C:

$$\phi_C = \Delta f\, T_2 + \phi_B + \psi_2 \tag{25}$$

Of the three terms in (25), the frequency offset is the dominant term because it causes a phase drift that grows with time. The phase ϕ_B is small by design, and ψ_2 represents a stationary term, and we can safely neglect both terms compared to the linear drift. This is also illustrated in the simulation shown in Fig. 8. Thus we have:

$$\sigma_C^2 \equiv \sigma_A^2 = \sigma_f^2 T_2^2 \tag{26}$$

Combining (24), (23) and (26), we get:

$$\sigma_A^2 = \frac{N_p\omega_n^3 T_2^2}{1 - \rho^2\omega_n^2 T_2^2} \tag{27}$$

Fig. 8 shows a simulation of the phase error over time with $T_1 = 150\mu sec$, $T_2 = 0.85$ ms, $\omega_n = 100$ kHz, $\rho = 1\%$ and $N_p = 7 \times 10^{-11} Hz^{-1}$ or -101 dBc/Hz. The VCO in this simulation has a quiescent frequency that is 1 kHz offset from the reference carrier signal. The spectral density of phase noise is chosen conservatively compared to typical numbers reported e.g. -110 dBc/Hz in [13]. For these numbers, we get $\sigma_A \approx 24°$ from (27). Since the phase error is a Gaussian variable with standard deviation smaller than $24°$ at all times, $\beta_\phi = E(\cos\phi_i) \geq 0.91$, and by Proposition 1, we can see that average beamforming gains of at least 91% are achievable. This is an average number and occasionally, phase errors larger than this can occur as seen in Fig. 8, where phase error becomes almost $35°$ at one point. Even with this large phase error, the resulting beamforming gain is 81% of the maximum. This confirms the results of Section II, that beamforming gain is robust to phase errors and demonstrates the basic feasibility of the distributed beamforming algorithm.

B. Comparison with fundamental Cramer-Rao bounds

So far in this analysis we have limited ourselves arbitrarily to a second-order PLL because it is the most commonly used device in practice. However, we can also derive fundamental limits on the size of the frequency and phase offsets, by viewing the PLL as a frequency and phase estimator. The PLL uses the (noisy) oscillator signal in the *sync* timeslot to form estimates $\hat{f}_l$ and $\hat{\phi}$. It uses the estimate $\hat{\phi}$ to drive the phase difference with the PLL input to zero, and $\hat{f}_l$ to tune the VCO's input to the frequency of the reference, and the sample-and-hold element keeps the VCO tuned to that estimate in the *transmit* timeslot. The Cramer-Rao lower bound for the variance of these offsets has been computed in previous work on frequency estimation [14]:

$$\hat{\sigma}_f^2 \doteq var(\hat{f}_l) = \frac{3N_p}{\pi^2 T_1^3}$$
$$\hat{\sigma}_\phi^2 \doteq var(\hat{\phi}) = \frac{2N_p}{T_1} \tag{28}$$

Using the same values used in Fig. 8, we find $\hat{\sigma}_f = 2.5$ Hz, and $\hat{\sigma}_\phi \ll 1°$. Since $\hat{\sigma}_f = 2.5$ Hz, and $\hat{\sigma}_\phi \ll 1°$. Since $\hat{\sigma}_f$ is substantially smaller than the PLL's root-mean squared frequency offset $\sigma_f = 418$ Hz, we conclude that there is significant suboptimality in using an analog PLL, therefore, performance can be further improved by using optimal digital processing.

V. CONCLUSION

We have investigated a master-slave architecture for achieving the carrier synchronization necessary for distributed transmit beamforming. There are several sources of synchronization error in this procedure, and the aggregate phase errors limit the achievable SNR gains from beamforming. We identified the dominant source of errors as VCO drift arising from time-division duplexed operation of the synchronization protocol. We examined the phase noise process of a second order analog PLL by simulation and analysis, and calculated the resulting beamforming gains. We also compared the performance of the PLL with the fundamental Cramer-Rao estimation bounds. Our results show that even with the suboptimal analog PLL, and with phase errors on the order of $60°$, it is possible to achieve SNR gains of 70% of the maximum. In summary, our investigation indicates that implementation of distributed beamforming at RF frequencies is challenging but potentially feasible.

One way to improve the beamforming performance is to use an optimal frequency and phase estimator instead of an analog PLL. This would require a more sophisticated digital implementation, and a detailed performance study of such a scheme is a topic for future work.

In this paper, we studied a reciprocity based approach to channel estimation designed to minimize coordination with the BS. An alternative methodology for distributed beamforming is to use feedback from the receiver (i.e. the BS) for distributed phase synchronization. We explored this approach in [11], [15] and our results are promising. It is also possible to combine these ideas with multi-hop routing schemes for wireless networks. These issues are not considered in this present work and are interesting areas for further inquiry.

<div style="text-align:center">APPENDIX I
PROOFS:</div>

Proof of Proposition 1: We can rewrite (2) as follows:

$$P_R = N \left\| \frac{1}{N} \sum_{i=1}^{N} |h_i|^2 e^{j\phi_i} \right\|^2 \tag{29}$$

Invoking the law of large numbers, and the fact that the $\{|h_i|^2\}$, are i.i.d. exponential random variables which are independent from the i.i.d $\{\phi_i\}$, we have

$$\frac{1}{N} \sum_{i=1}^{N} |h_i|^2 e^{j\phi_i} \to E\left[|h_i|^2 (\cos\phi_i + j\sin\phi_i)\right] \quad a.s. \tag{30}$$

The expectation on the RHS of (30) simplifies as follows

$$
\begin{aligned}
E\left[|h_i|^2 (\cos\phi_i + j\sin\phi_i)\right] &= E[|h_i|^2]E[\cos\phi_i] \\
&= E[\cos\phi_i] \tag{31}
\end{aligned}
$$

We have assumed that ϕ_i is symmetrically distributed around 0, and hence $E[\sin\phi_i] = 0$. Equation (31) results because $h_i \sim CN(0,1)$ and hence $|h_i|^2$ is exponential with unit mean. We thus have that

$$\left\| \frac{1}{N} \sum_{i=1}^{N} |h_i|^2 e^{j\phi_i} \right\|^2 \to (E[\cos(\phi_i)])^2 a.s. \tag{32}$$

since continuous functions of variables which are converging almost surely also converge almost surely, and the desired result follows. $\quad\square$

Proof of Proposition 2: The expected value of P_R can be written as

$$
\begin{aligned}
E[P_R] &= \frac{1}{N} E\left[\sum_{i=1}^{N} |h_i|^2 e^{j\phi_i} \sum_{l=1}^{N} |h_l|^2 e^{-j\phi_l} \right] \\
&= \frac{1}{N}(N + \frac{N(N-1)}{2} E[|h_1|^2 |h_2|^2 \\
&\quad .2\Re(e^{j(\phi_1-\phi_2)})] \tag{33} \\
&= \frac{1}{N}\left(N + \frac{N(N-1)}{2} 2E[\cos(\phi_1-\phi_2)]\right) \\
&= 1 + (N-1)E[\cos(\phi_1-\phi_2)] \\
&= 1 + (N-1)E[\cos(\phi_1)\cos(\phi_2) \\
&\quad - \sin(\phi_1)\sin(\phi_2)] \\
&= 1 + (N-1)E[\cos(\phi_i)]^2 \tag{34}
\end{aligned}
$$

where we have used the fact that the $\{h_i\}, \{\phi_i\}$ are i.i.d. and that the $\{\phi_i\}$ are symmetrically distributed around 0. $\quad\square$

Proof of Proposition 3: We once again begin with the definition for P_R.

$$
\begin{aligned}
P_R &= \frac{1}{\sqrt{N}} \left\| \sum_{i=1}^{N} |h_i|^2 \cos(\phi_i) + j|h_i|^2 \sin(\phi_i) \right\|^2 \\
&= \left\| \frac{1}{\sqrt{N}} \sum_{i=1}^{N} (|h_i|^2 \cos(\phi_i) - \alpha) \right. \\
&\quad \left. + j\frac{1}{\sqrt{N}} \sum_{i=1}^{N} |h_i|^2 \sin(\phi_i) + \sqrt{N}\alpha \right\|^2 \tag{35}
\end{aligned}
$$

where $\alpha = E[|h_i|^2 \cos(\phi_i)] = E[\cos(\phi_i)]$. Invoking the central limit theorem, as N gets large, the first term in (35) tends to a Gaussian random variable with mean 0 and variance $\sigma_c^2 \equiv var[|h_i|^2 \cos(\phi_i)]$. Similarly, the second term tends to a Gaussian random variable with mean 0 and variance $\sigma_s^2 \equiv var[|h_i|^2 \sin(\phi_i)]$. Since the last term in (35) is real and constant, it only shifts the mean of the first Gaussian random variable, so we can write

$$P_R \approx |X_c + jX_s|^2 \tag{36}$$

where $X_c \sim N(\sqrt{N}\alpha, \sigma_c^2)$, and $X_s \sim N(0, \sigma_s^2)$. Making use of the fact that $|h_i|^2$ is a unit mean exponential random variable,

$$
\begin{aligned}
\sigma_c^2 &= var[|h_i|^2 \cos(\phi_i)] \\
&= E[|h_i|^4 \cos^2(\phi_i)] - E[|h_i|^2 \cos(\phi_i)]^2 \\
&= 2E[\cos^2(\phi_i)] - E[\cos(\phi_i)]^2 \\
&\text{and similarly, } \sigma_s^2 = 2E[\sin^2(\phi_i)] \tag{37}
\end{aligned}
$$

Letting $m_c \equiv \sqrt{N}\alpha$, we have that $P_R = X_c^2 + X_s^2$, as given. The variance of P_R follows from standard calculations for moments of Gaussian random variables. $\quad\square$

REFERENCES

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

[2] S. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, Oct. 1998.

[3] M. Dohler, J. Dominguez, and H. Aghvami, "Link capacity analysis for virtual antenna arrays," in *Proc. 56th IEEE Vehicular Technology Conference 2002*, vol. 1, pp. 440–443.

[4] A. swol Hu and S. Servetto, "Optimal detection for a distributed transmission array," in *Proc. IEEE International Symposium on Information Theory 2003*, pp. 200–200.

[5] O. Oyman, R. Nabar, H. Bolcskei, and A. Paulraj, "Characterizing the statistical properties of mutual information in MIMO channels," *IEEE Trans. Signal Processing*, vol. 51, no. 11, pp. 2784–2795, Nov. 2003.

[6] B. Hassibi and A. Dana, "On the power efficiency of sensory and ad-hoc wireless networks," in *Proc. IEEE International Symposium on Information Theory 2003*, pp. 412–412.

[7] Y.-S. Tu and G. Pottie, "Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through AWGN channels," in *Proc. 55th IEEE Vehicular Technology Conference Spring 2002*, vol. 1, pp. 130–134.

[8] D. Brown, *A method for carrier frequency and phase synchronization of two autonomous cooperative transmitters.* [Online]. Available: http://spinlab.wpi.edu/publications/conferences/Brown_SPAWC_2005.pdf

[9] J. Elson, L. Girod, and D. Estrin, "Fine-grained network time synchronization using reference broadcasts," *SIGOPS Oper. Syst. Rev.*, vol. 36, no. SI, pp. 147–163, 2002.

[10] H. Ochiai, P. Mitran, H. Poor, and V. Tarokh, "Collaborative beamforming for distributed wireless ad hoc sensor networks," *IEEE Trans. Signal Process. (See also IEEE Trans. Acoustics, Speech, Signal Processing*, vol. 53, no. 1053-1058, pp. 4110–4124, 2005.)

[11] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Scalable feedback control for distributed beamforming in sensor networks," in *Proc. IEEE Intl. Symp. on Inform. Theory (ISIT) 2005*.

[12] H. Meyr and G. Ascheid, *Synchronization in Digital Communications*. New York: John Wiley and Sons, 1990.

[13] B. Razavi, "A study of phase noise in CMOS oscillators," *IEEE J. Solid-State Circuits*, vol. 31, no. 3, pp. 331–343, 1996.

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

[15] R. Mudumbai, B. Wild, U. Madhow, and K. Ramchandran, "Distributed beamforming using 1 bit feedback: from concept to realization," in *Proc. 44th Allerton Conference on Communication Control and Computing*, Sept. 2006.

Raghuraman Mudumbai received the B. Tech degree in electrical engineering from the Indian Institute of Technology, Madras, India in 1998, and the MS degree in electrical engineering from Polytechnic University, Brooklyn, USA in 2000. He is currently working for his Ph.D at the University of California Santa Barbara.

Gwen Barriac received the B.S. degree in electrical engineering from Princeton University, Princeton, NJ, in 1999, and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Santa Barbara in 2004. She is currently with Qualcomm, San Diego, CA.

Upamanyu Madhow (S 86 M 90 SM 96 F 05) received the bachelor's degree in electrical engineering from the Indian Institute of Technology, Kanpur, India, in 1985, and the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois, Urbana-Champaign, in 1987 and 1990, respectively. From 1990 to 1991, he was a Visiting Assistant Professor at the University of Illinois. From 1991 to 1994, he was a Research Scientist with Bell Communications Research, Morristown, NJ. From 1994 to 1999, he was on the faculty of the Department of Electrical and Computer Engineering, University of Illinois, Urbana-Champaign. Since December 1999, he has been with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, where he is currently a Professor. His research interests are in communication systems and networking, with current emphasis on wireless communication, sensor networks, and data hiding. Dr. Madhow is a recipient of the NSF CAREER award. He has served as Associate Editor for Spread Spectrum for the *IEEE Transactions on Communications*, and as Associate Editor for Detection and Estimation for the *IEEE Transactions on Information Theory*.