

Distributed Reception with Hard Decision Exchanges

D. Richard Brown III, *Senior Member, IEEE*, Upamanyu Madhoo, *Fellow, IEEE*, Min Ni, *Student Member, IEEE*, Matthew Rebholz, *Member IEEE*, and Patrick Bidigare, *Member, IEEE*

Abstract—This paper considers the problem of jointly processing messages received over a forward link from a single distant transmitter to a cooperative receive cluster connected by a local area network with finite available throughput. For N cooperating receivers, ideal distributed receive beamforming with direct exchange of unquantized observations leads to an N -fold gain in signal-to-noise ratio (SNR) for equal-gain additive white Gaussian noise channels, with significant additional gains over fading channels due to diversity. It is shown in this paper that a significant portion of these gains can be obtained simply by exchanging hard decisions among some or all of the nodes in the receive cluster. Mutual information computations and simulations of LDPC-coded systems show that optimal combining of hard decisions tends to perform within 0.5-2 dB of ideal receive beamforming. For the low per-node SNR regime of interest with large receive clusters, asymptotic analysis of a suboptimal combining technique termed “pseudo-beamforming” shows that distributed reception with hard decision exchanges performs within 1-2 dB of ideal receive beamforming.

Index Terms—Distributed reception, receiver cooperation, receive beamforming, cooperative communications, likelihood combining.

I. INTRODUCTION

WE consider the scenario shown in Figure 1 with a single transmitter and a cluster of N cooperative receive nodes connected by a wireless local area network (LAN) backhaul. The goal is to communicate common broadcast messages over the forward link from the distant transmitter to all of the receive nodes. As one example, the scenario in Figure 1 could correspond to a long-range downlink in which the receive cluster jointly processes messages from a distant base station. We consider the problem of “distributed reception” where some (or all) of the nodes in the receive cluster combine their observations to increase diversity and power gain and, consequently, improve the probability of successfully decoding noisy transmissions. Distributed reception can also result in increased communication range, increased




Fig. 1. Distributed reception scenario.

data rates, and/or decreased transmit power. The ideal performance in such a setting is obtained by receive beamforming, which could be implemented, for example, by each receiver broadcasting its unquantized (in practice, finely quantized) observations over a local area network (LAN) followed by maximal ratio combining at a fusion center or at each receiver. Direct implementation of receive beamforming can require unrealistic LAN throughput, however, even for modest forward link information rates. In this paper, we show that simply exchanging *hard decisions* among some or all of the nodes in the receive cluster provides a straightforward but powerful approach for fully distributed reception over a wireless LAN with limited capacity, with a degradation of only 0.5-2 dB relative to ideal receive beamforming.

The potential gains from our approach are best illustrated by an example. Suppose that the forward link employs 16-QAM modulation with an inform

Manuscript received September 24, 2013; revised January 14, 2014; accepted March 9, 2014. The associate editor coordinating the review of this paper and approving it for publication was R. Zhang.

This work was supported by the National Science Foundation awards CCF-1302104, CCF-1319458, and the NVIDIA Corporation through their donation of an NVIDIA Tesla K20 GPU used to accelerate the numerical results.

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

U. Madhoo is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara CA 93106 USA (e-mail: madhoo@ece.ucsb.edu).

M. Ni is with Oracle Corporation in Redwood City, CA 94065 (e-mail: min.ni@oracle.com).

M. Rebholz and P. Bidigare are with Raytheon BBN Technologies, Arlington VA 22209 USA (e-mail: matt.rebholz@bbn.com, bidigare@ieee.org).

Digital Object Identifier 10.1109/TWC.2014.051314.131763

to a subset of the nodes, depending on the channel conditions.

- We employ mutual information computations and simulations of LDPC-coded systems to demonstrate that optimal combining of hard decisions prior to decoding leads to performance within 0.5-2 dB of ideal receive beamforming. Simulations of a suboptimal but analytically tractable combining technique called “pseudo-beamforming” show that it performs within 1-2 dB of ideal receive beamforming. Soft decisions for the pseudo-beamformer are generated by applying a central limit theorem (CLT) based approximation, which is accurate even for a moderate number of receivers. These results demonstrate the efficacy of distributed reception with full and limited receiver participation.

- We analyze pseudo-beamforming in an asymptotic regime with a large number of receivers and low signal to noise ratio (SNR) per receiver, such that the SNR with ideal receive beamforming is finite and bounded away from zero. We show that the asymptotic degradation of performance relative to ideal receive beamforming is $2/\pi$ (about 2 dB) for QAM constellations (including QPSK) and $\pi/4$ (about 1 dB) for large PSK constellations (8PSK or higher). This also provides an asymptotic upper bound on the performance degradation of optimal combining relative to ideal receive beamforming. Our results for QAM can be viewed as a generalization of classical results on the hard decision penalty for binary communication in [1]–[3], while our results for larger PSK alphabets draw on the relatively recent low-SNR analysis in [4].

Prior work: There is a large body of related literature dating back more than three decades on the broad subjects of “multiterminal inference” [5]–[7], “distributed hypothesis testing” [8]–[10], and “distributed detection” [11]–[15]. In general, the setting in these problems is to have multiple agents forward quantized observations to a fusion center which then applies a detection or estimation rule according to a certain performance objective. Rate constraints limit the amount of information each agent can forward to the fusion center or decision maker. The problem we consider here is related to the classical model of distributed detection with a parallel fusion network [14] in that we assume communication over the wireless LAN backhaul among receive nodes is reliable. However, the key difference between the work reported in the present paper and the extensive literature on signal processing for distributed detection is that we focus on distributed detection for the specific purpose of improving the reliability of coded communication systems.

There is also by now a significant body of literature on distributed detection for communication. Information-theoretic studies on distributed reception for coded communication [16]–[18] typically focus on deriving achievable rates and capacity bounds for different relay forwarding strategies. Numerical quantizer optimization for information exchange in uncoded binary systems has been considered in [19], [20]. Distributed iterative message passing in coded *multiuser* binary systems is considered in [21], again using numerical optimization of quantizers. Constructive link layer iterative cooperation techniques coordinated by a cluster head are considered in [22], [23]. This body of literature focuses on a relatively small number of cooperating receivers, in contrast to the large receive cluster asymptotics that are the focus of the present paper. Furthermore, the papers with

constructive strategies [19]–[23] typically restrict attention to binary modulation, whereas our framework encompasses larger constellations (both QAM and PSK).

Note that, while we do not consider optimal quantization strategies in this paper (receivers in our system either broadcast their hard decisions over the LAN or are silent), such optimization could potentially further improve performance. For example, literature on quantization for *single receiver* coded systems [24], [25] shows that choosing quantization criteria such as mutual information can yield significant performance gains over standard mean squared error (MSE) based quantization.

Another set of closely related papers correspond to distributed reception methods with least-reliable and most-reliable bit exchanges [26]–[34]. In these methods, each receiver attempts to locally decode the message and then, in the case of least-reliable bit exchange, requests additional information from another receiver on a fraction of the least reliable outputs or, in the case of most-reliable bit exchange, broadcasts information about a fraction of its most reliable outputs. This process can be iteratively repeated to improve the coded bit error rate. Disadvantages of this class of approaches include the overhead of indexing the bits regarding which information is being exchanged (and the complexity of maintaining a memory of which bit indices have been used) as well as the latency caused by the iterative nature of the procedure. Both are avoided by our approach, in which a subset of nodes seeing the best channels simply broadcast *all* of their hard decisions, along with the associated channel magnitudes (assumed to be fixed over the duration of the message). Decoding is only attempted after all of the hard decisions have been received. To see the throughput advantage of our approach, consider a scenario with a BPSK forward link, a 10 node receive cluster, and a rate 1/2 block code with $(k, n) = (4050, 8100)$. Using (3), the hard decision combining technique considered in this paper requires a total LAN throughput of approximately 21 bits per forward link information bit. From [30], assuming three iterations, 10% most-reliable bits exchanged per iteration, and 5-bit reliability quantization, the total required LAN throughput is approximately 51 bits per forward link information bit. Even with 1-bit reliability quantization (hard decisions), the total required LAN throughput of most-reliable bit exchange is 39 bits per forward link information bit. This excess LAN throughput is largely due to the fact that each exchanged bit requires a 12-bit address to identify its bit index.

It is worth noting that there has been significant recent interest in using distributed reception techniques across multiple base stations in cellular uplinks e.g., [35]–[39]. These techniques are often called “coordinated multipoint” or CoMP. The focus of these papers is typically on mitigating interference, and high-fidelity information exchange between a small number of cooperating base stations via a high-speed wired or optical backhaul is assumed. This in contrast to our focus on a backhaul-constrained, single transmitter scenario with a large number of receivers.

An alternative approach to distributed reception, which sidesteps the need for a LAN backhaul with capacity increasing with the number of cooperating nodes, is for the cooperative receivers to act as amplify-forward relays, controlling their

phases so that their signals combine coherently in the air at a designated destination node. While promising experimental results have been reported recently with this approach [40], significant further effort and customized design is required to translate it into practice, in contrast to approaches utilizing explicit information exchanges such as the one in this paper, which can work with off-the-shelf radios.

The present paper represents a significant extension and generalization of prior conference papers [33], [34] involving a subset of the authors. These earlier papers contain numerical computations showing that distributed reception with hard decision exchanges for BPSK leads to a degradation of less than 2 dB relative to ideal receive beamforming. In the present paper, we provide a concise *analytical* characterization of the asymptotic performance degradation with respect to ideal receive beamforming for large QAM and PSK constellations. **Outline:** The system model and distributed reception protocol is described in Section II. Optimal and suboptimal combining rules are described in Section III. Section IV evaluates the performance of optimal hard decision combining using information-theoretic metrics, with numerical results showing that the performance attained is typically within 1 dB of ideal receive beamforming. Section V characterizes the asymptotic performance of the suboptimal pseudo-beamforming combining rule, thus also providing an asymptotic bound to the performance of optimal combining. Section VI provides numerical results demonstrating that we can indeed approach the performance of ideal receive beamforming for practical LDPC-coded systems. Section VII contains our conclusions.

II. SYSTEM MODEL

Referring to Fig. 1, we denote the total number of receivers as N and assume that messages from the distant transmitter to the receive cluster are (n, k) block coded where n and k correspond to the block length and the message length, both in bits, respectively. The forward link code rate is denoted as $r = k/n$. A mechanism for detecting a correctly decoded block (e.g., a CRC check) is also assumed at each receive node. The forward link channels are assumed to be constant over each block but may change from block to block. The forward link complex channel from the distant transmitter to receive node i for block m is denoted as $h_i[m]$ for $i = 1, \dots, N$. The vector channel for block m is denoted as $\mathbf{h}[m] = [h_1[m], \dots, h_N[m]]^T$.

The forward link alphabet is denoted as $\mathcal{X} = \{x_1, \dots, x_M\}$. The ℓ^{th} symbol in block m is denoted as $X[m, \ell]$ and is assumed to be drawn equiprobably from the alphabet. The average energy per transmitted symbol is then

$$\mathcal{E}_s = \mathbb{E} |X[m, \ell]|^2 = \frac{1}{M} \sum_{i=1}^M |x_i|^2. \quad (1)$$

Given an additive white Gaussian noise channel with power spectral density $N_0/2$ in the real and imaginary dimensions, the phase-corrected complex baseband signal received at the i^{th} receive node for the ℓ^{th} symbol of block m can be written as

$$U_i[m, \ell] = |h_i[m]|X[m, \ell] + W_i[m, \ell] \quad (2)$$

where $W_i[m, \ell] \sim \mathcal{CN}(0, N_0)$ is spatially and temporally independent and identically distributed (i.i.d.) proper complex

Gaussian baseband noise. We assume here that channels magnitudes are scaled such that the noise variance is identical at each receive node. Since the average received energy per forward link symbol at node i is $|h_i[m]|^2 \mathcal{E}_s$, the quantity $\rho_i[m] = \frac{|h_i[m]|^2 \mathcal{E}_s}{N_0}$ corresponds to the signal-to-noise ratio (SNR) at receive node i for symbols received in block m .

We assume that the receive cluster has a wireless LAN backhaul that supports reliable broadcast from each node to all other nodes and that the throughput of this LAN exceeds the forward link information rate. To allow for uninterrupted forward link transmissions, the LAN and the forward link are assumed to operate on different frequencies so that the receive cluster can transmit/receive on the LAN while receiving signals from the distant transmitter over the forward link.

A. Distributed Reception Protocol

While our focus is on how to combine the information exchanged over the LAN, for concreteness, we specify in this section a particular distributed reception protocol for hard decision exchanges. After each node receives and locally demodulates a block, the following steps are performed by the receive cluster over the LAN:

- 1) To determine the set of participating¹ nodes $\mathcal{P} \subseteq \{1, \dots, N\}$ and also to enable combining of the hard decisions, all N nodes exchange estimates of their channel magnitudes $|h_i[m]|$.
- 2) The $|\mathcal{P}| = K \leq N$ nodes with the largest channel magnitudes participate by broadcasting all of their hard decisions, denoted as $V_j[m, \ell] \in \mathcal{X}$ for all $j \in \mathcal{P}$, over the LAN. As messages are received over the LAN, each receive node (including those that do not participate) combine this quantized information with their local unquantized information. This combined information is then provided as an input to a soft-input decoder at each receive node, including those that do not participate.
- 3) If any receive node successfully decodes the message, it broadcasts the decoded message over the LAN to the other receive nodes in the cluster. If two or more nodes successfully decode the message and attempt to broadcast the successfully decoded block, it is assumed the LAN has a mechanism for contention resolution.

Note that this example protocol has a fixed LAN throughput requirement which depends on the modulation order and the number of participating nodes, as discussed below, and a fixed latency since decoding only occurs once. Other protocols could be used to reduce the average LAN throughput at the expense of making the latency variable and potentially increasing the average latency. For example, each node could attempt to decode the message prior to step 1. If any node is successful, the successfully decoded message could be broadcast over the LAN to the other receive nodes in the cluster and the remaining steps could be skipped. As another example, the hard decisions could be broadcast by each receiver in order

¹Due to poor channel conditions or LAN capacity constraints, some nodes in the receive cluster may not broadcast hard decisions. A "participating" node is a node that broadcasts its hard decisions for the block over the LAN to the other nodes in the receive cluster. Non-participating nodes do not broadcast hard decisions but do receive messages from other nodes in the receive cluster, and combine this information with their local unquantized observations in attempting to decode the block.

for complex alphabets and

$$p_{V_j|X}(v_j|X=x) = \frac{1}{\sqrt{\pi N_0}} \exp - \frac{(v_j - |h_j|x)^2}{N_0/2}$$

for real alphabets. For $i \in \mathcal{P} \setminus j$, note that $p_{V_i|X}(v_i|X=x)$ are the channel transition probabilities of the discrete memoryless channel (DMC) induced by the hard decisions at node i . For example, with a BPSK forward link with alphabet $\mathcal{X} = \{x_1, x_2\}$, we have

$$p_{V_i|X}(v_i = x_m|X=x) = \begin{cases} 1 - p_i & m = \ell \\ p_i & m \neq \ell \end{cases}$$

with crossover probability

$$p_i = Q \left(|h_i| \sqrt{\frac{2\mathcal{E}_s}{N_0}} \right).$$

In general, the process of forming hard decisions results in a DMC with M inputs and M outputs. The DMC transition probabilities $p_{V_i|X}(v_i = x_m|X=x)$ for many typical modulation formats, e.g., BPSK, QPSK, M -PAM, and M^2 -QAM, with hard decisions can be exactly determined using standard analysis techniques. Transition probabilities for M -PSK with hard decisions and $M > 4$ require approximations or numerical evaluation.

Since each receive node uses its local unquantized observation in combination with the hard decisions from the other receive nodes, the posterior probabilities $\text{Pr}(v_i = x_m|V = v)$ are different at each receive node. This may lead to a situation where some receive nodes can correctly decode the block while others cannot. The distributed reception protocol described in Section II-A allows any node that successfully decodes the message to broadcast the decoded message over the LAN to the full receive cluster. A block is unsuccessfully received only if all of the receive nodes are unable to decode.

B. Pseudo-beamforming

Recall that an ideal receive beamformer can be realized by scaling each continuous phase-corrected channel output U_j by its corresponding normalized channel magnitude and then summing, i.e.,

$$Y_{bf} \equiv Y_i = \sum_{j \in \mathcal{P}} \sqrt{\rho_i} U_j = \alpha \sum_{j \in \mathcal{P}} |h_j| U_j \quad (4)$$

where $\rho_i = \frac{|h_i|^2 \mathcal{E}_s}{N_0}$ and $\alpha = \frac{1}{\sqrt{\sum_{j \in \mathcal{P}} |h_j|^2}}$. Pseudo-beamforming is a simple but suboptimal combining technique where (4) is performed on the *hard decisions* from each node. Specifically, the pseudo-beamformer combiner output is

$$Y_{pbf} \equiv Y_i = \sum_{j \in \mathcal{P}} \sqrt{\rho_i} V_j = \alpha \sum_{j \in \mathcal{P}} |h_j| V_j \quad (5)$$

where $V_j \in \mathcal{X}$ for all j . Note that, since the unquantized local observation at each node is not used in the combiner, the pseudo-beamformer combiner output is the same at every receive node, so that all nodes either succeed or fail in decoding the block, under our assumption that the hard decision exchanges over the LAN are reliable.

Since the random variables $\{V_1, \dots, V_N\}$ are conditionally independent given the transmitted symbol, a CLT-based argument (which, in practice, tends to be a good approximation even for a moderate number of receivers) can be used to infer the conditional Gaussianity of the pseudo-beamformer output Y_{pbf} . To allow for different channel gains for different receivers, we apply the Lindeberg variant of the CLT, with the necessary conditions required to satisfy it in our context specified in the following Lemma.

Lemma 1. Let $\sigma_j^2(X) = \text{var}[\text{Re}\{V_j\} | X]$ and $\mu_j(X) = E[\text{Re}\{V_j\} | X]$. If, for all $j \in \mathcal{P}$,

$$\frac{|h_j|^2}{\sum_{j \in \mathcal{P}} |h_j|^2} \sigma_j^2(X) \rightarrow 0 \quad (6)$$

as $|\mathcal{P}| \rightarrow \infty$ then

$$A = \frac{\text{Re}\{Y_{pbf}\} - \alpha \sum_{j \in \mathcal{P}} |h_j| \mu_j(X)}{\alpha \sqrt{\sum_{j \in \mathcal{P}} |h_j|^2 \sigma_j^2(X)}} \xrightarrow{d} \mathcal{N}(0, 1) \quad (7)$$

when conditioned on X and \mathbf{h} where \xrightarrow{d} means convergence in distribution.

A proof of Lemma 1 is provided in Appendix A. The regularity condition in (6) can be thought of intuitively as requiring the channels to not vanish and for the local hard decisions to have some uncertainty as new nodes are added to the set of participating nodes.

An analogous application of Lemma 1 also implies the conditional asymptotic Gaussianity of the imaginary part of Y_{pbf} . Indeed, for the forward link alphabets considered in this paper, applying a two-dimensional version of the Lindeberg CLT implies the joint Gaussianity of the real and imaginary parts. The posterior likelihoods for each symbol (and bit) can therefore be easily computed once we specify the conditional second order statistics of the pseudo-beamformer output. In order to simplify these computations, we set the conditional covariances between the real and imaginary parts of the pseudo-beamformer output to zero (this holds asymptotically at low per-node SNRs), which amounts to approximating the effective noise at the output of the pseudo-beamformer as standard complex WGN.

The conditional means and variances of the pseudo-beamforming decision statistics are computed from the channel transition probabilities and known channel gains, i.e.,

$$\begin{aligned} E[Y_{pbf} | X = x_m] &= \alpha \sum_{j \in \mathcal{P}} |h_j| E[V_j | X = x_m] \\ &= \alpha \sum_{j \in \mathcal{P}} |h_j| x p_{m,j}^{(j)} \\ &= \mathbf{0} \end{aligned}$$

and

$$\begin{aligned} \text{var}[Y_{pbf} | X = x_m] &= \alpha^2 \sum_{j \in \mathcal{P}} |h_j|^2 \text{var}[V_j | X = x_m] \\ &= \alpha^2 \sum_{j \in \mathcal{P}} |h_j|^2 E[|V_j|^2 | X = x_m] - |E[V_j | X = x_m]|^2 \\ &= \alpha^2 \sum_{j \in \mathcal{P}} |h_j|^2 p_{m,j}^{(j)} |x|^2 - |E[V_j | X = x_m]|^2 \\ &= \mathbf{0} \end{aligned}$$

where $p_m^{(j)} = \Pr(\text{node } j \text{ decides } x \mid X = x_m)$ and where we have used the fact that $\{V_1, \dots, V_N\}$ are conditionally independent. Numerical results in Section VI show that this approach provides effective decoding performance, in agreement with the asymptotic performance predictions in Section V (within 1-2 dB of ideal receive beamforming), even for a small number of nodes.

Remark: The pseudo-beamformer output (5) is a particular linear combination of the hard decisions $\{V_1, \dots, V_N\}$ using the coefficients $\{|h_1|, \dots, |h_N|\}$. In principle, one could improve performance by optimizing these coefficients, e.g., to maximize the output SNR, still motivated by the CLT-based approximation for the output. The resulting coefficients would depend on the particular channel realizations in a complicated fashion, however. We do not consider this approach further, since it offers little gain relative to the simple combining rule (5), which is shown to provide excellent performance in the numerical results in Section VI. In addition, it is worth noting that (5) can actually be shown to be *asymptotically* optimal for the low per-node SNR regime considered in Section V.

IV. INFORMATION-THEORETIC PERFORMANCE ANALYSIS FOR OPTIMAL COMBINING

In this section, we develop an expression for the mutual information of distributed reception with optimal hard decision combining and provide numerical results demonstrating its performance gap relative to ideal receive beamforming.

Consider optimal hard decision combining at receive node j . Given equiprobable channel inputs X drawn from \mathcal{X} , the channel realization \mathbf{h} , the vector channel output $\mathbf{V} \in \mathcal{V}$ with elements arbitrarily quantized or unquantized the mutual information $I_{\mathbf{h}}(X; \mathbf{V})$ can be expressed as shown in (8) at the top of the next page, where $p(v|m) = p_{\mathbf{V}|X}(v|X = x_m)$ [41]. Note that all distributions in (8) are conditioned on \mathbf{h} and the conditional expectation is over the vector channel output \mathbf{V} given a scalar channel input $X = x_m$. Since the elements of \mathbf{V} are conditionally independent, we can write

$$\begin{aligned} p(\mathbf{v}|m) &= p_{\mathbf{V}|X}(\mathbf{v}|X = x_m) \\ &= p_{V_j|X}(v_j|X = x_m) \prod_{i \in \mathcal{P} \setminus j} p_{V_i|X}(v_i|X = x_m) \end{aligned}$$

and it follows that

$$\begin{aligned} &\frac{\prod_{\ell=1}^{M-1} p(\mathbf{v}|\ell)}{\prod_{\ell=1}^{M-1} p(\mathbf{v}|m)} \\ &= \frac{\prod_{\ell=1}^{M-1} p_{V_j|X}(v_j|X = x_\ell) \prod_{i \in \mathcal{P} \setminus j} p_{V_i|X}(v_i|X = x_\ell)}{\prod_{\ell=1}^{M-1} p_{V_j|X}(v_j|X = x_m) \prod_{i \in \mathcal{P} \setminus j} p_{V_i|X}(v_i|X = x_m)} \\ &= \frac{\prod_{\ell=1}^{M-1} p_{V_j|X}(v_j|X = x_\ell)}{\prod_{\ell=1}^{M-1} p_{V_j|X}(v_j|X = x_m)} \prod_{i \in \mathcal{P} \setminus j} \frac{p_{V_i|X}(v_i|X = x_\ell)}{p_{V_i|X}(v_i|X = x_m)}. \end{aligned}$$

At node j , since the marginal observation V_j is unquantized, the expectation in (8) must be approximated numerically, either by numerical integration or by Monte-Carlo simulation.

Figure 2 shows an example of the mutual information for distributed reception with BPSK and 16-QAM forward link modulation, $N = 10$ receive nodes, and fixed channels $\mathbf{h} = [1, \dots, 1]^T$. These results were obtained through Monte-Carlo simulation of (8) where 10^4 i.i.d. noise realizations




Fig. 2. Mutual information for a distributed reception system with BPSK and 16-QAM forward links, $N = 10$ receive nodes, full participation, and $\mathbf{h} = [1, \dots, 1]^T$. Distributed reception (a) is the case when the hard decisions are optimally combined with local unquantized observations. Distributed reception (b) is the case when the hard decisions are optimally combined with local hard decisions.

were generated at each receive node. All receive nodes are assumed to participate in the distributed reception protocol. Since the forward link channels to each receive node are the same in this example, the performance of distributed reception with hard decision exchanges is the same for all receive nodes (this is not the case for general \mathbf{h} , however). This example shows that distributed reception with optimal hard decision combining can provide significant capacity gains with respect to single-receiver processing and that simply exchanging hard decisions among the nodes in the receive cluster can result in performance within approximately 1.8 dB of ideal receive beamforming for a BPSK forward link and within approximately 0.8 dB of ideal receive beamforming for a 16-QAM forward link with fixed, equal-gain channels.

The results in Figure 2 also compare “distributed reception (a)” where the hard decisions from other nodes are optimally combined with local unquantized observations versus “distributed reception (b)” where the hard decisions from other nodes are optimally combined with local hard decisions. While the latter approach is clearly suboptimal, it may be simpler in practice to optimally combine local hard decisions with hard decisions from other nodes since the local posterior likelihoods can be computed in the same manner as the posterior likelihoods of the hard decisions received over the LAN. The performance loss of using local hard decisions is relatively minor in this example due to the fact that forward link channels are all identical.

Figure 3 shows an example of the outage probability for distributed reception with 16-QAM forward link modulation, full participation, and i.i.d. Rayleigh fading channels with $h_i \stackrel{i.i.d.}{\sim} \mathcal{CN}(0, 1)$. These results were also obtained through Monte-Carlo simulation of (8) with 5000 i.i.d. channel realizations for each node and 1000 i.i.d. noise realizations per channel realization. This example shows that distributed reception with optimal hard decision combining and local hard decisions (“distributed reception (b)”) performs within




Fig. 4. Block error rate of distributed reception with hard decision exchanges versus energy per symbol for a BPSK forward link with full participation.

are equivalent to the results obtained for 2-PAM, Gray-coded QPSK, and Gray-coded 4-QAM. We see that the gap between ideal receive beamforming and pseudo-beamforming is approximately 2 dB, even for smaller values of N , and that optimal hard decision combining tends to perform closer to ideal receive beamforming, especially at smaller values of N , due to the use of unquantized local information and exact posterior likelihood calculations.

Figure 5 shows the block error rate of distributed reception versus \mathcal{E}_S/N_0 for Gray-coded 16-QAM and 16-PSK forward links with $N = 1, 2, 5, 10$ and full participation ($K = N$). Here we see that the gap between ideal receive beamforming and pseudo-beamforming is approximately 1 dB for both M -PSK (consistent with Corollary 2) and 16-QAM (better than the 2 dB performance loss predicted by Corollary 1). Optimal combining tends to perform within approximately 0.5 dB of ideal receive beamforming for the settings shown in Figure 5. These trends closely match the information-theoretic outage probability results in Figure 3, except for a 1-2 dB shift in the curves that can be attributed to the gap between the LDPC code and the Shannon limit for the equivalent binary input channel that it sees.

Figure 6 shows the block error rate of distributed reception versus \mathcal{E}_S/N_0 for a BPSK forward link with partial participation. The set of participating receive nodes is selected as the K receive nodes with the largest channel magnitudes from the total pool of N receive nodes for $(K, N) \in \{(5, 10), (5, 20), (10, 10), (10, 20)\}$. Only results for optimal hard decision combining are plotted, but results for pseudo-beamforming exhibit the same trends. We see that partial participation can lead to significant reduction in LAN throughput requirements, while incurring a modest performance loss; for example, using the best $K = 5$ receive nodes from a $N = 10$ node pool with respect to full participation ($K = N = 10$) incurs only about 1 dB loss while cutting the required LAN throughput approximately in half. Furthermore, for fixed K (e.g., based on the available LAN throughput), the block error rate improves as N increases because of the added selection diversity.

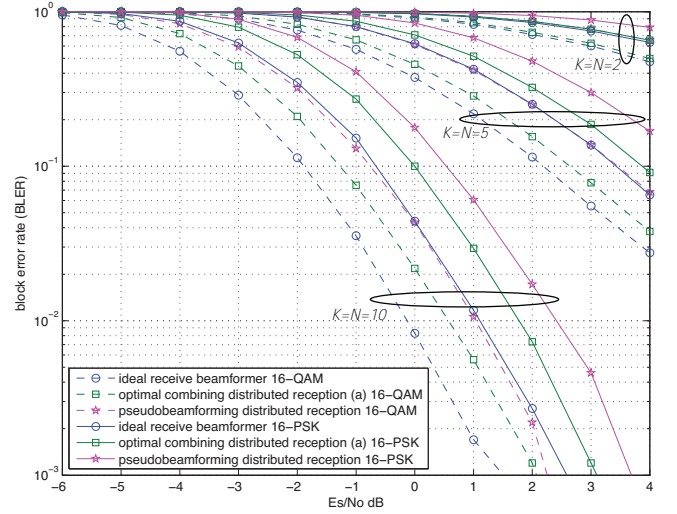


Fig. 5. Block error rate of distributed reception with hard decision exchanges versus energy per symbol for Gray-coded 16-QAM and 16-PSK forward links with full participation.




Fig. 6. Block error rate of distributed reception with hard decision exchanges and optimal combining versus energy per symbol for a BPSK forward link with partial participation. Approximate normalized LAN throughputs according to (3) are also shown.

VII. CONCLUSIONS

We have shown, using information-theoretic computations, simulations of LDPC-coded systems, and asymptotic analysis, that distributed reception with hard decision exchanges suffers a relatively small penalty relative to ideal receive beamforming. From a practical perspective, this implies that excellent performance can be achieved with off-the-shelf hardware (e.g., a receive cluster connected via WiFi), with a significant reduction in LAN throughput requirements relative to sharing lightly quantized observations.

The results reported here open up a number of important questions for future research. First, it is natural to explore distributed compression strategies for further reducing LAN throughput requirements without increasing latency and while exploiting the unique features of our problem; unlike conventional distributed compression, where the goal is to reduce

$$\mathbb{E}[\operatorname{Re}(V_j) | X = x] \approx \frac{1}{2} - \frac{(2\ell - 1| - 1)\rho_j}{\sqrt{2\pi}} x_1 + \sum_{m=2}^{M-1} \frac{2\rho_j}{\sqrt{2\pi}} x_m + \frac{1}{2} - \frac{(2\ell - M| - 1)\rho_j}{\sqrt{2\pi}} x_M \quad (15)$$

$$= \frac{2(2 - M - 1)\rho_j}{\sqrt{2\pi}} x_M \quad (16)$$

$$= \frac{2(M - 1)\rho_j}{\sqrt{2\pi}} x \quad (17)$$

$$\operatorname{var}[\operatorname{Re}(V_j) | X = x] = \mathbb{E}[\operatorname{Re}(V_j)^2 | X = x] - (\mathbb{E}[\operatorname{Re}(V_j) | X = x])^2 \quad (18)$$

$$\approx 2 \frac{1}{2} - \frac{(2\ell - 1| - 1)\rho_j}{\sqrt{2\pi}} x_1^2 + 2 \sum_{m=2}^{M/2-1} \frac{2\rho_j}{\sqrt{2\pi}} x_m^2 - \frac{2(M - 1)\rho_j}{\sqrt{2\pi}} x^2 \quad (19)$$

$$\approx (M - 1)^2 a^2 \quad (20)$$

$$\mathbb{E}[V_j | X = x_1] = \sum_{m=1}^M x_m p_{m,1} \quad (21)$$

$$\approx \sum_{m=1}^M a e^{j 2(m-1)/M} \frac{1}{M} + \frac{1}{\sqrt{\pi}} \cos \frac{2\pi(m-1)}{M} \sin \frac{\pi}{M} \rho_j \quad (22)$$

$$= \frac{2a\rho_j \sin(\pi/M)}{\sqrt{\pi}} \sum_{m=1}^{M/2} \cos^2 \frac{2\pi(m-1)}{M} \quad (23)$$

$$= \frac{M\rho_j \sin(\pi/M)}{2\sqrt{\pi}} x_1 \quad (24)$$

APPENDIX C PROOF OF LEMMA 3

Proof: An M -PSK constellation has a complex-valued alphabet given as $\mathcal{X} = \{x_1, \dots, x_M\} = a, a e^{j 2/M}, a e^{j 4/M}, \dots, a e^{j (M-1)2/M}$ with $\mathcal{E}_s = a^2$. Since the constellation is symmetric, we focus on $X = x_1$. The probability of deciding $V_j = x_m$ given $X = x_1$ can be expressed as

$$p_{m,1} = \int_{(2m-3)\pi/M}^{(2m-1)\pi/M} f_{|X}(\theta | X = x_1) d\theta$$

for $m \in \{1, \dots, M\}$ with the conditional phase distribution given as [4]

$$f_{|X}(\theta | X = x_1) = \frac{1}{2\pi} e^{-\frac{\rho_j}{\sqrt{\pi}} \cos(\theta)} e^{-\frac{\rho_j}{\sqrt{\pi}} \sin^2(\theta)} \frac{1 - Q\left(\frac{\rho_j \cos(\theta)}{2\rho_j^2 \cos^2(\theta)}\right)}{2\rho_j^2 \cos^2(\theta)}$$

where $\rho_j^2 := \frac{|h_j|^2 a^2}{N_0}$. In the low per-node SNR regime, we can calculate a first-order Taylor series expansion of $p_{m,1}$ at $\rho_j = 0$ by computing

$$p_{m,1} \Big|_{\rho_j=0} = \int_{(2m-3)\pi/M}^{(2m-1)\pi/M} f_{|X}(\theta | X = x_m) d\theta = \frac{1}{M}$$

and

$$\begin{aligned} \frac{\partial}{\partial \rho_j} p_{m,1} \Big|_{\rho_j=0} &= \int_{(2m-3)\pi/M}^{(2m-1)\pi/M} \frac{\partial}{\partial \rho_j} f_{|X}(\theta | X = x_1) d\theta \\ &= \int_{(2m-3)\pi/M}^{(2m-1)\pi/M} \frac{\cos(\theta)}{2\sqrt{\pi}} d\theta \\ &= \frac{1}{2\sqrt{\pi}} \sin \frac{(2m-1)\pi}{M} - \sin \frac{(2m-3)\pi}{M} \\ &= \frac{1}{\sqrt{\pi}} \cos \frac{2\pi(m-1)}{M} \sin \frac{\pi}{M} \end{aligned}$$

Hence, in the low per-node SNR regime with ρ_j small, we have

$$p_{m,1} \approx \frac{1}{M} + \frac{1}{\sqrt{\pi}} \cos \frac{2\pi(m-1)}{M} \sin \frac{\pi}{M} \rho_j.$$

Under the assumption that $M \geq 4$ is even, we can now compute the conditional expectation shown in (21) — (24). By symmetry, it is easy to see that $\mathbb{E}[V_j | X = x] = \frac{M \rho_j \sin(\pi/M)}{2\sqrt{\pi}} x$.

The conditional variance can be computed similarly as

$$\begin{aligned} \operatorname{var}[V_j | X = x] &= \mathbb{E}[|V_j|^2 | X = x] - |\mathbb{E}[V_j | X = x]|^2 \\ &\approx a^2 - \frac{M\rho_j \sin(\pi/M)}{2\sqrt{\pi}} x^2 \\ &\approx a^2 \end{aligned}$$

for all $\ell \in \{1, \dots, M\}$ where we have discarded the term with ρ_j^2 in the final approximation since ρ_j is small under the low per-node SNR assumption

- [30] A. Nayagam, J. Shea, and T. Wong, "Collaborative decoding in bandwidth-constrained environments," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 2, pp. 434–446, Feb. 2007.
- [31] X. Li, T. Wong, and J. Shea, "Performance analysis for collaborative decoding with least-reliable-bits exchange on AWGN channels," *IEEE Trans. Commun.*, vol. 56, no. 1, pp. 58–69, Jan. 2008.
- [32] Y. Zhou, O. Takyu, and Y. Umeda, "A CSI reduction method of block cooperative decoding system in OFDM wireless communication scheme," in *Proc. 2010 IEEE International Symposium on Personal, Indoor and Mobile Radio Communications Workshops*, pp. 461–466.
- [33] D. R. Brown III, M. Ni, U. Madhow, and P. Bidigare, "Distributed reception with coarsely-quantized observation exchanges," in *Proc. 2013 Conference on Information Sciences and Systems*.
- [34] R. Wang, D.R. Brown III, M. N. U. Madhow, and P. Bidigare, "Outage probability analysis of distributed reception with hard decision exchanges," in *Proc. 2013 Asilomar Conference on Signals, Systems, and Computers*.
- [35] E. Aktas, J. Evans, and S. Hanly, "Distributed decoding in a cellular multiple-access channel," *IEEE Trans. Wireless Commun.*, vol. 7, no. 1, pp. 241–250, Jan. 2008.
- [36] D. Gesbert, S. Hanly, H. Huang, S. Shamai Shitz, O. Simeone, and W. Yu, "Multi-cell MIMO cooperative networks: a new look at interference," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 9, pp. 1380–1408, 2010.
- [37] I.-H. Wang and D. Tse, "Interference mitigation through limited receiver cooperation: symmetric case," in *Proc. 2009 IEEE Information Theory Workshop*, pp. 579–583.
- [38] P. Marsch and G. P. Fettweis, Eds., *Coordinated Multi-Point in Mobile Communications: From Theory to Practice*. Cambridge University Press, 2011.
- [39] C. Yang, S. Han, X. Hou, and A. Molisch, "How do we design comp to achieve its promised potential?" *IEEE Wireless Commun.*, vol. 20, no. 1, pp. 67–74, 2013.
- [40] F. Quitin, A. Irish, and U. Madhow, "Distributed receive beamforming: a scalable architecture and its proof of concept," in *Proc. 2013 IEEE Vehicular Technology Conference — Spring*, pp. 1–5.
- [41] G. Ungerboeck, "Channel coding with multilevel/phase signals," *IEEE Trans. Inf. Theory*, vol. 28, no. 1, pp. 55–67, 1982.
- [42] Y. Xiao and K. Kim, "Alternative good LDPC codes for DVB-S2," in *Proc. 2008 International Conference on Signal Processing*, pp. 1959–1962.
- [43] X. Huang, J. Fan, Y. Xiao, and K. Kim, "Girth 4 and low minimum weights' problems of LDPC codes in DVB-S2 and solutions," in *Proc. 2009 International Conference on Wireless Communications, Networking and Mobile Computing*, pp. 1–4.
- [44] D. Love and P. Bidigare, "Receive spatial coding for distributed diversity," in *Proc. 2013 Asilomar Conference on Signals, Systems, and Computers*.
- [45] A. Ozgur, R. Johari, D. Tse, and O. Leveque, "Information-theoretic operating regimes of large wireless networks," *IEEE Trans. Inf. Theory*, vol. 56, no. 1, pp. 427–437, Jan. 2010.
- [46] P. Billingsly, *Probability and Measure*, 3rd ed. Wiley, 1995.



D. Richard Brown III (S'97–M'00–SM'09) received the B.S. and M.S. degrees in Electrical Engineering from The University of Connecticut in 1992 and 1996, respectively, and received the Ph.D. degree in Electrical Engineering from Cornell University in 2000. From 1992–1997, he was with General Electric Electrical Distribution and Control. He joined the faculty at Worcester Polytechnic Institute (WPI) in Worcester, Massachusetts in 2000 and currently is an Associate Professor. He also held an appointment as a Visiting Associate Professor at

Princeton University from August 2007 to June 2008. His research interests are currently in coordinated wireless transmission and reception, synchronization, distributed computing, and game-theoretic analysis of communication networks.



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 on distributed 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*, published by Cambridge University Press in 2008.

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*, published by Cambridge University Press in 2008.

Min Ni (S'07–S'13) received the B.S. degree in Electrical Engineering from Sichuan University, Chengdu, China, in 2007, and received the M.S. degree and Ph.D. degree in Electrical and Computer Engineering from Worcester Polytechnic Institute in 2011 and 2014, respectively. She joined Oracle corporation in Redwood city, California in 2014 and currently is a member of technical staff. Her research interests are currently in coordinated wireless transmission and reception and synchronization of communication networks.



Matthew Rebolz (M'13) received his B.S. and M.S. degrees in Electrical Engineering from Virginia Tech in and 2013, respectively. From 2011 to the present, he has been with Raytheon BBN Technologies as a Staff Scientist.

T. Patrick Bidigare received the B.S.E.E. degree in computer engineering and the and the Ph.D. degree in mathematics from the University of Michigan, Ann Arbor, in 1992 and 1997 respectively. His doctoral work was in the area of algebraic combinatorics. He was with the Environmental Institute of Michigan (now General Dynamics) from 1991 to 2007 where he served as principal investigator on programs in areas of protection, radar responsive tagging and secure communications. Since 2007 he has been a technical director at BBN Technologies (now Raytheon) in Arlington, VA where he leads programs in distributed RF sensing and communications. His research interests include hardware and algorithm development for implementing distributed clock synchronization, network self-localization and distributed beamforming.