

## Research Article

# On Building a Cooperative Communication System: Testbed Implementation and First Results

**Patrick Murphy, Ashutosh Sabharwal, and Behnaam Aazhang**

*Department of Electrical and Computer Engineering, Rice University, 6100 S. Main St., Houston, TX 77005, USA*

Correspondence should be addressed to Ashutosh Sabharwal, ashu@rice.edu

Received 1 November 2008; Revised 24 March 2009; Accepted 5 June 2009

Recommended by Alejandro Ribeiro

We present the results from over-the-air experiments of a complete implementation of an amplify and forward cooperative communications link. Our custom OFDM-based physical layer uses a distributed version of the Alamouti block code, where the relay sends one of the branches of Alamouti-encoded symbols. First we show analytically and experimentally that amplify and forward protocols are unaffected by carrier frequency offsets at the relay. This result allows us to use a conventional Alamouti receiver without change for the distributed relay system, thereby allowing cooperative systems to reuse components of current (noncooperative) systems. Our full system implementation shows gains of up to 5.5 dB in peak power constrained networks. Thus, we can conclusively state that even the simplest form of relaying can lead to significant gains in practical implementations.

Copyright © 2009 Patrick Murphy et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

## 1. Introduction

Cooperative communications [1, 2] (and references therein), has emerged as a significant concept to improve reliability and throughput in wireless systems. In cooperative communications, the resources of distributed nodes are effectively pooled for the collective benefit of all nodes. While cooperation can occur at different network layers (and hence at different time scales), physical layer cooperation at symbol time scales offers the largest benefit. However, symbol level cooperation is also potentially the hardest to implement due to significant challenges in enabling it in distributed systems. In this paper, we take the first significant steps in building and understanding the issues in implementing practical cooperative communication systems.

We focus on amplify and forward protocols [3] where the relay node simply amplifies and retransmits the analog waveform received from the source node. This simple protocol was shown to increase the diversity order [3], allowing single-antenna nodes to cooperate and achieve performance like a real MIMO system. However, most analyses to date have ignored the challenge of implementing such a distributed space-time scheme in the face of analog and digital distortions like carrier frequency offset, inaccurate synchroniza-

tion, and gain control for analog to digital conversion. All of these are significant parts of practical wireless systems which, if handled poorly, can cause significant performance degradation. Note that additive white Gaussian noise model with multiplicative fading is a highly oversimplified abstraction of an actual wireless link, which almost always has to deal with severe nonlinearities, finite precision, and lack of shared clock reference. In [4], the authors captured real wireless channels between multiple nodes to determine what rate various relay schemes could achieve. The measurement analysis in [4] clearly showed that relaying could potentially provide gains on real-world channels. An actual implementation, however, has to deal with many additional nonidealities due to automatic gain control, carrier offset estimation, channel estimation, and lack of perfect synchronization. We have built a real-time system to understand the gains of cooperative communication in an operational system in the presence of all channel and device imperfections.

Our contributions are threefold. First, we show analytically and experimentally that amplify and forward protocols are not affected by the carrier frequency offset of any relaying nodes. That is, the final received signal at the destination is only affected by the carrier offset between the source and destination, much like *relayless* system. This is significant finding




FIGURE 1: Nodes in basic amplify and forward system.

which shows that from the point of view of the destination, it can use a receiver built for a conventional multiple antenna transmissions without employing a multiuser-like front-end to handle simultaneous, noncoherent transmissions from multiple nodes.

The above finding leads to the second contribution which allows us to use a traditional Alamouti receiver *without* any change for the relay system. In fact, the destination can be potentially made agnostic to whether the transmission is  $1 \times 1$  (SISO),  $2 \times 1$  (MISO Alamouti), or  $1 \times 1 \times 1$  (relay system with distributed Alamouti). We note that the Alamouti receiver for a true  $2 \times 1$  system is not optimal in the current context due to different noise variances on the two paths (direct and one through relay). However, our experimental results show that the unmodified receiver provides significant gains over a noncooperative system, even with a suboptimal receiver. Our motivation for not changing the receiver is to enable reuse of current receivers and thus aid rapid adoption of cooperative techniques in deployed systems. Our future work will move toward building optimal receivers to quantify the gains from changing the current equipment.

Lastly, we build a fully operational amplify and forward system which assumes only time synchronization between source and relay to mimic packet synchronous systems like GSM or WiMAX. The system is built using the resources of the Rice University Wireless Open-Access Research Platform (WARP) [5] and implements a high-speed wireless link using an Alamouti-encoded OFDM physical layer. With a peak power constraint per node, relaying adds more power to the system leading to gains of up to 5.5 dB in BER performance for both BPSK and QPSK systems in *actual* wireless channels. With a total power constraint, where the total transmit power of the relaying system is the same as that of the point-to-point system, the relaying systems gains are still 2 dB or more. The gains can be attributed to a mix of diversity benefits and reduction in effective path loss due to relay location.

We immediately note that our work has only scratched the surface in exploring the issues in implementing cooperative systems and studying their performance in real wireless environments. For example, we have only partially optimized the parameters in the receiver front-end (e.g., automatic gain control) and the choice of amplify and forward schemes. Further, our tests have been limited to indoor environments. Despite of these suboptimal elements, we show that cooperation can still lead to significant gains in real implementations with commercial grade components. As obvious extensions to this work, we will implement other forms of cooperation (decode and forward variants), study performance under different channel conditions and

network topologies, and gain a deeper understanding in energy-performance-complexity tradeoffs.

The rest of the paper is organized as follows. In Section 2, we review the amplify and forward scheme and show how it is unaffected by relay carrier frequency offset. Section 3 describes our complete implementation, experimental setup, and main results. We conclude in Section 4.

## 2. Amplify and Forward

Amplify and forward is the simplest class of cooperative communications schemes [6]. In amplify and forward systems, one node (the source) sends information to another (the destination). A third node (the relay) captures part of the source's transmission, amplifies it, and retransmits it without any further processing. The destination uses the combination of the source and relay's transmissions to decode the data, hopefully with fewer errors than if the source had transmitted alone. Figure 1 shows the basic configuration of these three nodes.

The underlying idea of amplify and forward can be applied in a wide variety of ways. Our goal is to construct a cooperative system based on one realization of amplify and forward, then to use this system to explore some of the issues which arise when building a cooperative communications system. We uncovered one particularly interesting property of amplify and forward systems, which we discuss in detail below.

**2.1. Carrier Frequency Offset.** In practice, wireless nodes generate radio-frequency carriers using phase-locked loops driven by a local frequency reference. The frequency of the generated carrier varies with the frequency of the local reference. When multiple nodes use independent local references, their RF carriers differ in frequency. In most hardware, this carrier frequency offset is large enough that it must be addressed by the wireless physical layer algorithms.

Carrier frequency offset is a well-studied problem; practical algorithms exist to mitigate CFO in a wide variety of wireless systems. However, the effects of CFO have largely been ignored in the development of cooperative communication algorithms. Some schemes have been proposed which attempt to synchronize the carriers of multiple transmitting nodes in hopes that their signals will constructively combine at the destination [7, 8]. These schemes rely on some kind of shared information among transmitting nodes, either in the form of communicated phase offsets or reception of a common beacon signal. In either case, the complexity of maintaining synchronization is nontrivial.

**2.2. Radio Transceiver Model.** Our first contribution in this paper is to explore the construction of an amplify and forward system which exploits a useful property of common radio hardware. We show how this property allows the destination node to simply ignore the carrier offset of an amplify and forward relay node.

The following analysis intentionally ignores many practical aspects of a wireless communications system, including

physical layer waveform design, gains, filters, analog/digital conversion, and channel effects. The goal of this derivation is solely to demonstrate the effects of carrier frequency offset in an amplify and forward link. In real communications systems, CFO is an analog (i.e., continuous time) problem, inherent in the local generation of RF carriers at each node. Thus, to trace the impact of CFO through a cooperative link, we consider only the analog baseband and RF signals in the following. The effects we omit here certainly play a part in constructing an actual cooperative link (as described in Section 3). However, in scenarios with little Doppler effect, carrier frequency offsets can be analyzed independent of these other impairments.

Figure 2 illustrates our models for the analog processes of RF upconversion and downconversion. These models reflect the inner workings of a direct conversion RF transceiver, where a common sinusoidal carrier is used for both the transmit and receive chains. The use of a common carrier reference for the transmit and receive paths at the relay node is a critical (but thankfully realistic) assumption in this analysis.

In the following, let  $\omega_c$  denote the frequency of the carrier, and let  $\text{LPF}(x)$  be a low-pass filter. Note that the baseband signals ( $X_{\text{BB}}$  below) are complex, but the RF signals ( $X_{\text{RF}}$  below) are allreal. This matches the implementation of wireless systems, where an RF signal is a single voltage, and complex baseband signals are represented with separate  $I$  and  $Q$  voltages.

First, we define the  $\text{Tx}()$  and  $\text{Rx}()$  functions which express the processes illustrated in Figure 2:

$$\begin{aligned} X_{\text{RF}} &= \text{Tx}(X_{\text{BB}}, \omega_c) \\ &= \text{Re}(X_{\text{BB}}) \cos(t\omega_c) - \text{Im}(X_{\text{BB}}) \sin(t\omega_c) \\ &= \frac{X_{\text{BB}} e^{jt\omega_c} + X_{\text{BB}}^* e^{-jt\omega_c}}{2}, \\ X_{\text{BB}} &= \text{Rx}(X_{\text{RF}}, \omega_c) \\ &= \text{LPF}(X_{\text{RF}} e^{jt\omega_c}). \end{aligned} \quad (1)$$

**2.3. CFO in Amplify and Forward.** We now apply these functions to trace the frequency offset of a signal as it propagates through an amplify and forward cooperative link. Figure 3 illustrates the nodes and signal names used in the following derivation.

Consider the initial source to relay transmission, where  $\omega_{\text{CS}}$  and  $\omega_{\text{CR}}$  are the carrier frequencies of the source and relay, respectively. The sequence of operations and corresponding signals are

$$S_{\text{BB}} \longrightarrow \text{Tx}(S_{\text{BB}}, \omega_{\text{CS}}) \longrightarrow S_{\text{RF}} \longrightarrow \text{Rx}(S_{\text{RF}}, \omega_{\text{CR}}) \longrightarrow R_{\text{BB}}, \quad (2)$$

which when expanded gives the following, assuming  $\text{LPF}()$  is a linear filter with gain 2:

$$\begin{aligned} S_{\text{RF}} &= \text{Tx}(S_{\text{BB}}, \omega_{\text{CS}}) \\ &= \frac{S_{\text{BB}} e^{jt\omega_{\text{CS}}} + S_{\text{BB}}^* e^{-jt\omega_{\text{CS}}}}{2}, \\ R_{\text{BB}} &= \text{Rx}(S_{\text{RF}}, \omega_{\text{CR}}) \\ &= \text{LPF}(S_{\text{RF}} e^{jt\omega_{\text{CR}}}) \\ &= \text{LPF}\left(\frac{(S_{\text{BB}} e^{jt\omega_{\text{CS}}} + S_{\text{BB}}^* e^{-jt\omega_{\text{CS}}}) (e^{jt\omega_{\text{CR}}})}{2}\right) \\ &= \text{LPF}\left(\frac{S_{\text{BB}} e^{jt(\omega_{\text{CS}} - \omega_{\text{CR}})} + S_{\text{BB}}^* e^{-jt(\omega_{\text{CS}} + \omega_{\text{CR}})}}{2}\right) \\ &= S_{\text{BB}}(e^{jt(\omega_{\text{CS}} - \omega_{\text{CR}})}). \end{aligned} \quad (3)$$

As expected, the baseband signal received at the relay suffers a frequency offset due to the difference between the source and relay carrier frequencies.

Next, we trace the transmission from relay to destination:

$$R_{\text{BB}} \longrightarrow \text{Tx}(R_{\text{BB}}, \omega_{\text{CR}}) \longrightarrow R_{\text{RF}} \longrightarrow \text{Rx}(R_{\text{RF}}, \omega_{\text{CD}}) \longrightarrow D_{\text{BB}}, \quad (4)$$

$$\begin{aligned} R_{\text{RF}} &= \text{Tx}(R_{\text{BB}}, \omega_{\text{CR}}) \\ &= \frac{R_{\text{BB}} e^{jt\omega_{\text{CR}}} + R_{\text{BB}}^* e^{-jt\omega_{\text{CR}}}}{2}, \\ D_{\text{BB}} &= \text{Rx}(R_{\text{RF}}, \omega_{\text{CD}}) \\ &= \text{LPF}(R_{\text{RF}} e^{jt\omega_{\text{CD}}}) \\ &= \text{LPF}\left(\frac{(R_{\text{BB}} e^{jt\omega_{\text{CR}}} + R_{\text{BB}}^* e^{-jt\omega_{\text{CR}}}) (e^{jt\omega_{\text{CD}}})}{2}\right) \\ &= \text{LPF}\left(\frac{R_{\text{BB}} e^{jt(\omega_{\text{CR}} - \omega_{\text{CD}})} + R_{\text{BB}}^* e^{-jt(\omega_{\text{CR}} + \omega_{\text{CD}})}}{2}\right) \\ &= R_{\text{BB}}(e^{jt(\omega_{\text{CR}} - \omega_{\text{CD}})}). \end{aligned} \quad (5)$$

Finally, we substitute the previous expression for  $R_{\text{BB}}$ :

$$\begin{aligned} D_{\text{BB}} &= (S_{\text{BB}} e^{jt(\omega_{\text{CS}} - \omega_{\text{CR}})}) (e^{jt(\omega_{\text{CR}} - \omega_{\text{CD}})}) \\ &= S_{\text{BB}}(e^{jt(\omega_{\text{CS}} - \omega_{\text{CD}})}). \end{aligned} \quad (6)$$

Thus, the received baseband signal at the destination node suffers a frequency offset determined solely by the difference between  $\omega_{\text{CS}}$  and  $\omega_{\text{CD}}$ , independent of their respective offsets from  $\omega_{\text{CR}}$ . In other words, the relay's carrier frequency offset with respect to the source and destination nodes does not affect the final signal received at the destination.

**2.4. Empirical Verification.** In order to substantiate the preceding analysis and to verify the impact of its inherent

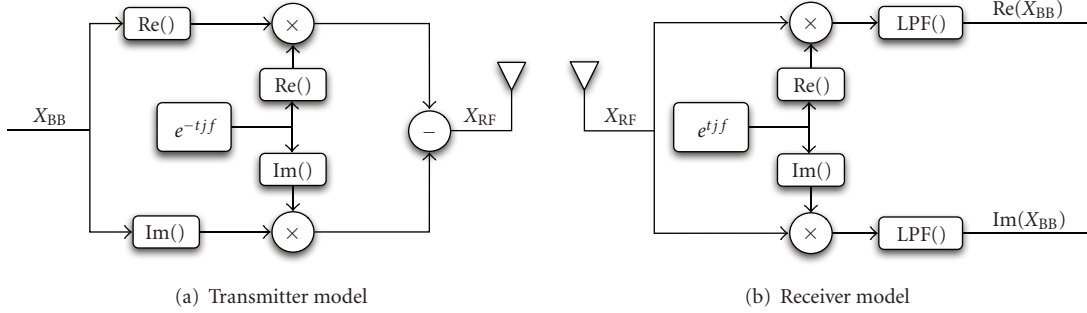


FIGURE 2: Tx and Rx models.




FIGURE 3: Nodes and signals in an amplify and forward link.

TABLE 1: Alamouti STBC encoding.

|          | $t_0$ | $t_1$    |
|----------|-------|----------|
| Stream A | $x_0$ | $-x_1^*$ |
| Stream B | $x_1$ | $x_0^*$  |

assumptions, we constructed an RF link which allows the direct observation of carrier frequency offsets. In this setup, one node acts as both the source and destination, while a second node acts as the relay. The source generates a constant valued baseband signal, which after upconversion results in the transmission of a sinusoid at exactly  $\omega_{CS}$ . The relay node receives this sinusoid, downconverts it with its local carrier, and saves the samples at baseband. If the analysis is correct, these samples should be of a sinusoidal signal with frequency  $(\omega_{CS} - \omega_{CR})$ . The relay then transmits the same samples back to the first node. If our assumptions and analysis hold, the first node should receive a constant valued signal at baseband, showing no frequency offset as a result of amplification and retransmission by the relay.

Figure 6 shows the results of this experiment. Two trials are depicted here. In the first, the relay transmits its received signal after a short delay, approximately 10 milliseconds. In the second, the relay waits two minutes before retransmitting. The transmission in both directions happens over a wire to eliminate any channel effects. The top plots depict the phase of the signal received at the relay. The phase of this signal is increasing linearly in time, corresponding to a received sinusoid. This sinusoid is the direct result of carrier frequency offset between the two nodes. The bottom plots depict the phase of the signal received at the source node, after it is buffered and retransmitted by the relay. The complete lack of the saw wave pattern clearly illustrates the relay canceling its own carrier offset during retransmission. In Figure 4(b), a very slight slope can be observed in the received signal's phase. This is the result of a minor drift in the node's local oscillator frequency. The WARP hardware

utilized in this experiment uses temperature-compensated crystal oscillators for the carrier reference, which accounts for the very minor drift, even after two minutes. Cheaper oscillators, like those used in low-end commercial wireless hardware, could exhibit larger drifts over time.

### 3. Building a Cooperative System

This section describes the construction of an amplify and forward cooperative communications system which relies on the properties described in Section 2. This system is implemented on WARP [5], making heavy use of the custom hardware, physical layer designs, and other support packages provided by the platform.

**3.1. Overview.** Our system is built on the idea of distributed space-time coding [3, 9], where multiple nodes cooperate to transmit a signal which approximates the transmission of a single, multiple-antenna node. In particular, we employ Alamouti's space-time block code (STBC) [10]. Figure 5 illustrates the classic  $2 \times 1$  STBC configuration which the proposed cooperative scheme imitates. The signal names here correspond to the two spatial streams generated by a two-antenna Alamouti transmitter; these signals play a key role in the proposed cooperative version of this link.

The Alamouti STBC encodes two data symbols across two symbol periods and two spatial streams. Given two data symbols  $x_0$  and  $x_1$ , the code outputs the signals shown in Table 1. In each symbol period at the receiver, the superposition of the two streams is received after each passes through separate channels; the signals received in two symbol periods are represented by  $r_0$  and  $r_1$ . The receiver uses local channel estimates and the following combining rules to recover the original data symbols:

$$\begin{aligned}
 r_0 &= h_A x_0 + h_B x_1 + n_0, \\
 r_1 &= -h_A x_1^* + h_B x_0^* + n_1, \\
 \tilde{x}_0 &= h_A^* r_0 + h_B r_1^*, \\
 \tilde{x}_1 &= h_B^* r_0 - h_A r_1^*.
 \end{aligned} \tag{7}$$

Much like other cooperative protocols for half-duplex radios, the proposed cooperative link operates in two time slots per packet. Figure 6 illustrates the activity of each node

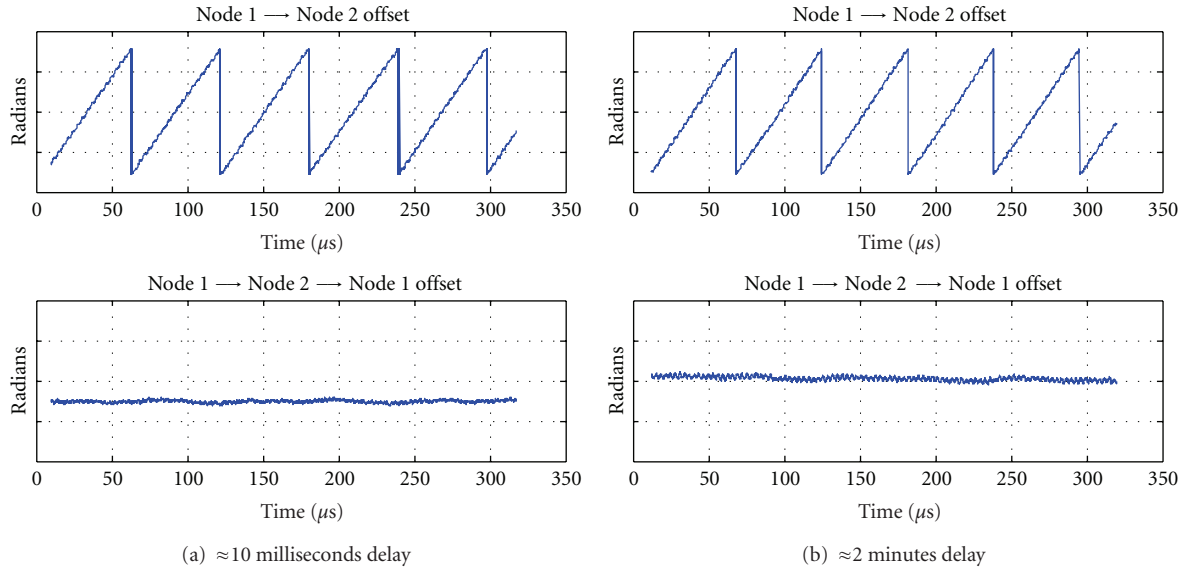


FIGURE 4: Experimental observation of carrier offset in a relay system.




FIGURE 5: Equivalent  $2 \times 1$  Alamouti setup.

in our scheme’s two time slots. In the first slot, the source node transmits the full packet, encoded, using the Alamouti space-time block code. This signal matches that which would be sent from one antenna in a true two-antenna Alamouti transmission. The relay node receives this transmission and stores the raw samples in a buffer. In the second time slot, the source node transmits the other half of the Alamouti-encoded sequence, and the relay transmits its stored copy of the first transmission. The destination node receives the superposition of these simultaneous transmissions. From the perspective of the destination, it receives a standard Alamouti-encoded packet, where each of the spatial components was exposed to an independent channel.

**3.2. Physical Layer Design.** In the proposed amplify and forward scheme, the timing of the two transmissions in the second time slot cannot be perfectly guaranteed. The offset between the arrival times of the source and relay’s transmissions can be modeled as multipath. This is analogous to the signals sent from a standard two-antenna Alamouti transmitter arriving at slightly different times at the receiver after passing through different channels.

In order to cleanly handle this potential impairment, we chose OFDM as the underlying physical layer for our coop-

erative system. OFDM’s inherent immunity to multipath makes it an ideal PHY for an amplify and forward system, as a delayed transmission which is treated as just another reflection in the channel.

The details of the physical layer design are described below.

**3.2.1. Frame Format.** Our cooperative physical layer uses the following frame format, partially inspired by IEEE 802.11a [11]. The transmissions are composed of four components:

- (i) short training symbols (STS): 10 16-sample sequences, used for AGC convergence,
- (ii) long training symbols (LTS): 2.5 64-sample sequences, used for fine symbol timing,
- (iii) channel training symbols ( $Tr_A/Tr_B$ ): 80-sample channel training symbols,
- (iv) spatial streams: Payload data encoded with Alamouti’s space-time code.

The composition of the transmissions in each time slot is illustrated in Figure 7.

In the first time slot, the source node transmits a frame designed to trigger packet detection at the relay but avoids packet detection at the destination. This is achieved by omitting the long training symbols. Our OFDM receiver uses a correlator to search for the LTS within a fixed window after an energy detection event. If this correlation fails, the receiver assumes a false packet detection and resets. The relay node does not perform this check.

In the second time slot, the LTS must be included to allow the destination to properly detect the packet and synchronize the receiver. It is critical as both the source and relay nodes send the LTS so that the destination can detect the packet based on either node’s transmission. The source node transmits the full STS/LTS preamble in this slot. The




FIGURE 6: Node configurations and activity in the amplify and forward system.

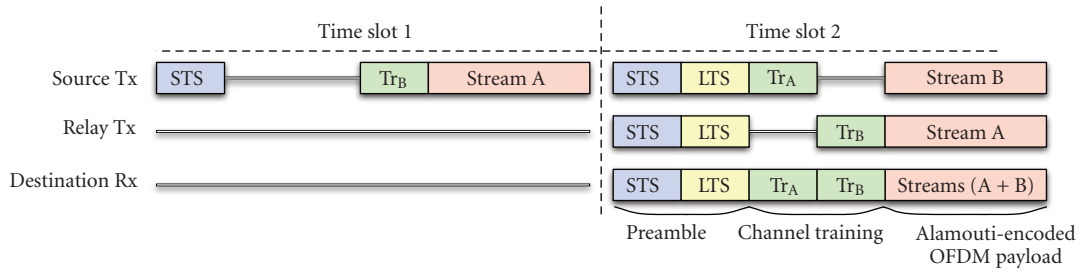


FIGURE 7: Physical layer frame formats.

relay node also sends a full preamble in the second time slot. The relay stores this preamble in a lookup table and sends it in place of the STS captured in the first slot.

During the preamble, the source and relay are transmitting the same signal simultaneously. This opens the possibility of unintended beamforming, where the identical signals constructively or destructively interfere at the destination. However, we did not observe this effect during our experiments. We attribute this to a number of phase variations between the source and relay transmissions. First, the relay generates its preamble locally, so its transmission occurs with a carrier frequency offset relative to the source. This offset does not exist for the buffered and retransmitted channel training and payload symbols, for the reasons discussed above. Second, the source and relay have independent phase noise characteristics, as each node's carrier is generated from a local PLL. Finally, the source and relay use separate sampling clocks, introducing another source of independent jitter (phase noise) which will tend to blur their combined beam pattern.

After sending their preambles, the nodes send channel training symbols in alternate symbol periods. This orthogonal-in-time training allows the destination to estimate the two channels separately. Following the training symbols, both nodes transmit payload symbols simultaneously. These transmissions are entirely nonorthogonal, overlapping in frequency and time. The source transmits the second Alamouti spatial stream while the relay transmits the version of the first stream it captured during the first time slot. These simultaneous transmissions utilize the Alamouti space-time block code to avoid interfering with one another and exploit the delay spread tolerance of OFDM to combat small synchronization differences between the overlapping transmissions.

We designed this scheme to allow successful packet detection and synchronization at the destination even if it

receives just one of the two transmissions in the second time slot. Both transmissions contain everything the destination needs to receive the packet—preamble, channel training symbols and the full payload. In an intuitive sense, this scheme preserves full diversity as it will fail only if two (presumably) independent channels are simultaneously in deep fades.

**3.2.2. Synchronization.** In our setup, the relay node uses a wired synchronization signal from the source to initiate its buffering and retransmission processes. The packet lengths are also fixed throughout our experiments and are known ahead of time by every node. This kind of synchronization mimics what is used in scheduled access systems like GSM or WiMAX.

The destination node implements autonomous packet detection. This system uses the RSSI (received signal strength indicator) signal from the RF transceiver to detect a spike in received energy indicating the start of a new packet. The timing of the packet is refined in the PHY by cross-correlation against the LTS in the packet's preamble. This is the same approach to packet detection and timing used in a noncooperative random access system. If the uncertainty of packet arrival times at the destination was eliminated, as in slotted systems like GSM or WiMAX, we expect that the system performance would improve.

Every node has independent sampling and radio reference clocks. Given the relatively short packets, we ignore sampling frequency offsets throughout. Offsets among the radio reference clocks result in carrier frequency offsets, the effects of which we explored in Section 2.1.

**3.2.3. Gain Control.** Both the relay and destination nodes implement automatic gain control, which executes with each packet detection. The AGC algorithm sets the gains for

the receive amplifiers in the RF transceiver in the first 2-3 microseconds after packet detection, well within the STS section of the preamble.

The relay node amplifies its received signal in both the analog and digital domains. The relay's RF transceiver uses low-noise amplifiers to boost the analog RF and baseband signals in the receive path. The gain settings for these amplifiers are chosen for each packet by the AGC system. The result of this amplification is an analog signal whose amplitude is independent of the received power. This signal is sampled by the relay's ADC and buffered in the FPGA. During the second time slot, the relay multiplies these stored samples by a constant before driving them into the DAC. The radio board's RF transceiver and power amplifier apply the final stages of gain before transmission. The digital gain value is fixed, as it is determined solely by the difference in the ADC and DAC dynamic ranges and does not depend on the RF transceiver's gain settings.

**3.2.4. Channel Estimation.** The destination must estimate two channels in order to properly combine the Alamouti-encoded symbols, analogous to the two channels in a classic  $2 \times 1$  Alamouti configuration. In our setup, however, one of these channels is actually the combination of two physical channels: source-to-relay and relay-to-destination. Only the relay's retransmitted signal experiences this compound channel. The destination node uses a training symbol originally embedded by the source, then retransmitted by the relay, to estimate the compound channel. The second channel the destination must estimate is the source-destination channel using a training symbol embedded in the source's transmission in the second time slot. The source node constructs its transmissions so that in the second time slot, the two training symbols do not overlap, allowing independent estimates at the destination node.

**3.2.5. OFDM.** The source and destination nodes implement identical, full Alamouti OFDM transceivers. This PHY was originally implemented for use in a standard  $2 \times 1$  Alamouti OFDM link. Due to the structure of our amplify and forward configuration, the same receiver design works as it is, without modification, in the cooperative system. The transmitter design requires minor modifications to enable the back-to-back transmissions of the spatial streams from a single antenna. The universality of the receiver design which functions without modification in  $1 \times 1$ ,  $2 \times 1$ , and  $1 \times 1 \times 1$  configurations is a significant benefit of amplify and forward systems.

All processing in the PHY is implemented in the WARP FPGA and executes in real-time. Carrier frequency offset estimation, symbol timing estimation, phase noise tracking, equalization, and detection are all implemented in fixed point in the FPGA. The physical layer operates in a 12.5 MHz bandwidth with a raw data rate of 7.5 or 15 Mbps by transmitting BPSK or QPSK symbols in 48 of 64 subcarriers. One training symbol is used per channel, and 4 pilot subcarriers are used to track phase noise and residual carrier offset.

**3.3. Experiment Design.** Our experiments were conducted in a three-node setup, each implementing a single antenna half-duplex transceiver. The nodes were built with WARP hardware, each with one FPGA board and one radio board [5] connected to omnidirectional antennas. The physical locations of the nodes and their antennas were fixed throughout the experiment. The source and destination antennas were separated by 4.5 m; the relay's antenna was approximately half-way in between. All transmissions were over-the-air in the 2.4 GHz ISM band. A channel not used by any other wireless devices was chosen to minimize the impact of interference on our BER results. All experiments were conducted indoors. As a result, there was very little mobility in the channel during each experiment, and each node had a clear line-of-sight to the other two.

Given that the nodes' locations are fixed throughout, we used the transmission power of the source and relay nodes as a proxy for SNR and as the independent variables in the results below.

The transmit power and received gains are adjusted inside the WARP radio board's RF transceiver. The various gain stages are applied to the analog and RF signals by low-noise amplifiers. The analog signals at the ADCs and DACs are always the same amplitude, so the contribution of quantization to the overall performance is fixed and independent of a node's transmit or receive power.

**3.4. Results.** From the plots in Figure 8, it is immediately clear that the relay node significantly improves the BER performance in the cooperative link.

The top curve shows the performance of the noncooperative link. For these tests, the source and destination nodes operate exactly as described above, but the relay is switched off. A copy of this curve shifted left 3 dB is also included. This shifted curve illustrates the best possible performance the destination could achieve if it performed maximal ratio combining (MRC) on the two copies of each packet it receives, instead of simply ignoring the energy it received in the first time slot.

A second observation we can make from these results is whether adding a relay helps even if the system's total transmit power was artificially constrained. To make this comparison, we first choose a point along the  $X$ -axis, determine the total transmit power (source + relay), then find the point on the  $X$ -axis of equal power. The comparison of relay aided versus norelay BER values at these two points reveals whether fixed total power is better allocated to the relay or source.

Figure 9 shows a region from the BER plot in more detail and illustrates this comparison. It can be seen that allocating some power to the relay outperforms the comparable norelay configuration by at least 5 dB. This gain is heavily dependent on the network topology and channel conditions. An exhaustive study of networks would be required to state this result more generally. However, this example still clearly demonstrates that given a total power constraint, the tradeoff between source and relay transmit power can favor allocating power to the relay in some realistic situations.

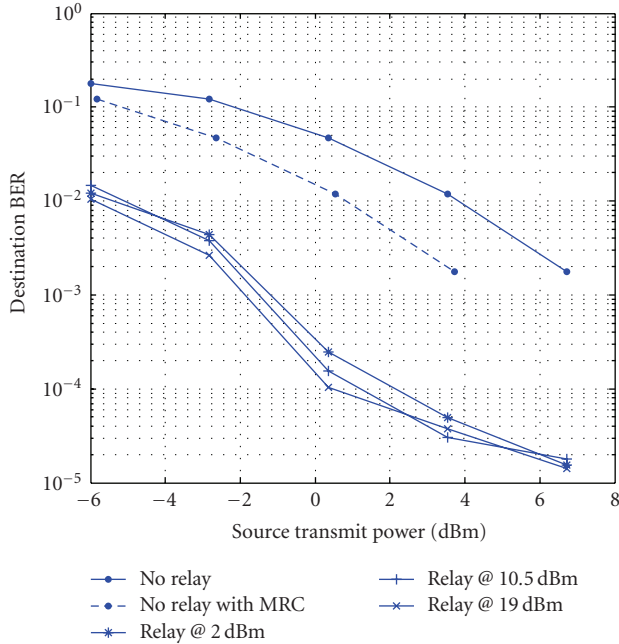


FIGURE 8: BER results for relay versus relayless links.




FIGURE 9: BER results with equivalent power comparison.

A final point to observe in these results is the relatively minor performance improvement which results from extra transmission power at the relay. This strongly indicates that the source-relay link dominates the overall performance. This fits the intuitive notion that if the relay receives a poor signal in the first time slot, it will spend most of its power retransmitting noise with little benefit to the destination.

The second set of results in Figure 10 compares the performance of two modulation schemes. For these tests, we




FIGURE 10: BER results for BPSK and QPSK modulations.

configured the source and destination to use either BPSK or QPSK on the same 48 (of 64) subcarriers used in the previous trials. For BPSK, half as many bytes were sent per packet in order to maintain the same source/relay timing parameters used for QPSK transmissions. As expected, BPSK performs better than QPSK, and both schemes benefit when using the relay. This plot offers another interesting comparison. For large payloads, the total transmission duration for relay aided QPSK and relayless (i.e., single time slot per packet) BPSK is roughly equivalent. Thus, the source can choose between the two schemes without significantly affecting medium access control. Given this choice, the results in Figure 10 indicate that at low to moderate SNR, the source should choose relay aided QPSK over relayless BPSK.

**3.5. Comments.** Our observation of carrier offset cancellation at the relay is based on a few important but realistic assumptions. First, the magnitude of the relay's offset must be small relative to the signal's bandwidth. If the offset is too large, the resulting baseband spectrum at the relay will be shifted into the stop band of the transceiver's low-pass filters. The same problem would occur in a noncooperative system if the source-destination CFO was too large. As long as the wireless hardware uses oscillators of sufficient accuracy to allow noncooperative links, our CFO-free amplify and forward observation will hold. The second assumption is that the relative frequencies of the carrier frequencies at each node do not drift significantly between time slots. This is again a function of the quality of the system's oscillators. An oscillator's frequency stability over temperature and time is generally well specified by the manufacturer and can be used to determine the expected frequency drift. In practice, frequency changes on per-packet time scales are very small (as we demonstrate in Section 2.4).



OFDM is a natural PHY for an amplify and forward system, as it allows the retransmitted signal to be treated as just another multipath reflection at the destination. However, this approach reduces the overall delay spread tolerance of the OFDM PHY, as it uses some part of each OFDM symbol's cyclic prefix to account for inaccuracies in the timing of the relay's transmission. While the size of the cyclic prefix (1.28 microseconds in our case) in OFDM systems is generally conservative, especially for stationary indoor environments, this loss of delay spread tolerance could impact performance in more hostile channels.

Finally, we note that our results are only a first but important step toward studying cooperative communications in practice. We observed real performance gains when using amplify and forward, but the magnitude of these gains is certainly subject to many parameters, including network topology, channel conditions, and physical layer design.

#### 4. Conclusion

In this work, we have built a full cooperative communications system which operates in real-time over real wireless channels. This system puts to practice some ideas from existing work in cooperative communications. It also relies on our own results in understanding carrier frequency offset in amplify and forward systems. Our performance evaluation shows a clear benefit to using amplify and forward relays, demonstrating a significant BER improvement under realistic wireless conditions. It is clear that physical layer cooperation is largely uncharted territory, especially with regard to implementation in practical systems. To enable a deeper understanding, our implementation of an amplify and forward cooperative system will be available in the Rice WARP project's open-source repository [5], allowing the community to systematically evaluate different relaying protocols over real wireless channels with all practical considerations.

#### Acknowledgments

The authors would like to thank Dr. Chris Dick at Xilinx and the Xilinx University Program for their continuing support of the WARP project. This work was partially funded by NSF Grants CNS-0325971, CNS-0619767, and CNS-0551692.

#### References

- [1] A. Sendonaris, E. Erkip, and B. Aazhang, "User cooperation diversity—part I: system description," *IEEE Transactions on Communications*, vol. 51, no. 11, pp. 1927–1938, 2003.
- [2] G. Kramer, I. Marić, and R. D. Yates, "Cooperative communications," *Foundations and Trends in Networking*, vol. 1, no. 3-4, pp. 271–425, 2006.
- [3] J. N. Laneman and G. W. Wornell, "Distributed space-time-coded protocols for exploiting cooperative diversity in wireless networks," *IEEE Transactions on Information Theory*, vol. 49, no. 10, pp. 2415–2425, 2003.
- [4] S. Berger and A. Wittneben, "Experimental performance evaluation of multiuser zero forcing relaying in indoor scenarios,"

in *Proceedings of the 61st IEEE Vehicular Technology Conference (VTC '05)*, vol. 2, pp. 1101–1105, Stockholm, Sweden, May–June 2005.

- [5] "Rice University WARP Project," <http://warp.rice.edu>.
- [6] P. Herhold, E. Zimmermann, and G. Fettweis, "On the performance of cooperative amplify-and-forward relay networks," in *Proceedings of the 5th International ITG Conference on Source and Channel Coding (SCC '04)*, pp. 451–458, Erlangen, Germany, January 2004.
- [7] Y.-S. Tu and G. J. Pottie, "Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through AWGN channels," in *Proceedings of the 55th IEEE Vehicular Technology Conference (VTC '02)*, vol. 1, pp. 130–134, Birmingham, Ala, USA, May 2002.
- [8] D. R. Brown III, G. B. Prince, and J. A. McNeill, "A method for carrier frequency and phase synchronization of two autonomous cooperative transmitters," in *Proceedings of the 6th IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC '05)*, pp. 260–264, New York, NY, USA, June 2005.
- [9] P. A. Anghel, G. Leus, and M. Kaveh, "Multi-user space-time coding in cooperative networks," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '03)*, vol. 4, pp. 73–76, Hong Kong, April 2003.
- [10] S. M. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 8, pp. 1451–1458, 1998.
- [11] IEEE Std 802.11a-1999, "Part 11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications: high-speed physical layer in the 5 GHz band," 1999.