Quantize-Map-Forward (QMF) Relaying: An Experimental Study

Melissa Duarte
École Polytechinque Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
melissa.duartegelvez@epfl.ch

Ayan Sengupta
École Polytechinque Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
ayan.sengupta@epfl.ch

Siddhartha Brahma
École Polytechinque Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
siddhartha.brahma@epfl.ch

Christina Fragouli
École Polytechinque Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
christina.fragouli@epfl.ch

Suhas Diggavi
University of California Los Angeles (UCLA)
Los Angeles, CA 90095
USA
suhas@ee.ucla.edu

ABSTRACT

We present the design and experimental evaluation of a wireless system that exploits relaying in the context of WiFi. We opt for WiFi given its popularity and wide spread use for a number of applications, such as smart homes. Our testbed consists of three nodes, a source, a relay and a destination, that operate using the physical layer procedures of IEEE802.11. We deploy three main competing strategies that have been proposed for relaying, Decode-and-Forward (DF), Amplify-and-Forward (AF) and Quantize-Map-Forward (QMF). QMF is the most recently introduced of the three, and although it was shown in theory to approximately achieve the capacity of arbitrary wireless networks, its performance in practice had not been evaluated. We present in this work experimental results—to the best of our knowledge, the first ones—that compare QMF, AF and DF in a realistic indoor setting. We find that QMF is a competitive scheme to the other two, offering in some cases up to 12% throughput benefits and up to 60% improvement in frame error-rates over the next best scheme.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*wireless communication*; E.4 [**Coding and Information Theory**]: Error control codes.

Keywords

Cooperative Communication; Relaying; 802.11; Quantize-Map-Forward; Software Defined Radio; WARP.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MobiHoc'13, July 29–August 1, 2013, Bangalore, India.
Copyright 2013 ACM 978-1-4503-2193-8/13/07 ...$15.00.

1. INTRODUCTION

In theory, it is well established that relaying can extend the range of a wireless network and increase its bandwidth efficiency. In practice, although relaying is increasingly becoming part of standards, it is not clear which of the many proposed relaying strategies performs best, and when. We present in this paper an experimental study of state-of-the-art relaying strategies, and compare their performance for several indoor 802.11 (WiFi) topologies.

WiFi is what most homes already have: if we can enhance its performance, we pave the road for a number of interesting applications. Ease of installation, low-cost, interoperability and ubiquity have made WiFi the most popular WLAN choice with homeowners. Smart homes seem a tangible possibility thanks to this wide penetration of WiFi, as many appliance manufacturers are looking into WiFi-enabled machines that communicate with each other. Green home solutions as well rely on WiFi connectivity for controlling energy efficiency. High rate wireless data, home support for impaired people, entertainment, the list of potential applications is long. Yet, the reality is that the indoor WiFi links are not always reliable and cannot support high enough rates, both because of physical obstruction (*e.g.,* an elevator) and interference (*e.g.,* a cordless phone)—leading to a disappointing performance for a number of interesting applications. WiFi is thus a good example where relaying can have high impact.

We focus in this paper on the most basic topology: a source sends data to a destination with the help of a relay, as depicted in Fig. 1. In our home, the source could be a WiFi router that streams video to a tablet with the help of a WiFi enabled appliance that boosts the network performance by acting as a relay. We select this setup for several reasons: the simpler the experiment the easier to interpret results; it can serve as a building block towards more complex configurations; this is commonly accepted as the topology that would be used in practice, at least in the immediate future, thus making it more meaningful for performance comparisons; and finally, surprisingly, even for this small topology, few experimental results have been published.

In information theory, relaying has been a topic of research for the last forty years, and many approaches have been proposed and analyzed. Perhaps the two most studied are Amplify and Forward (AF) and Decode and Forward (DF). DF in particular has been shown to be within one bit from the capacity of the single relay network we are focusing on. Recently, a third strategy, Quantize-Map-Forward (QMF), was also shown to achieve the capacity of the single relay network within one bit, and moreover, to achieve within a constant gap the capacity of *arbitrary* wireless networks.

QMF has in theory a number of attractive attributes. It dictates a simple relay operation: the relay captures the signal, quantizes it, performs an appropriate mapping and forwards it to the destination, which is much simpler than decoding for instance. It only requires receive channel information at each relay. It is scalable, in the sense that adding or removing relays does not require to change what the operating relays do. Most importantly, it is proved that it approximately achieves the capacity of *arbitrary* wireless networks.

However, it is not clear how many of these advantages translate into practical systems. The information-theoretic analysis assumes infinite length coding and no complexity constraints on the transceivers; even then there is a possible gap to information-theoretic capacity; and of course the analysis assumes perfect estimates of the noise and the channel coefficients for decoding. It is well possible that for networks at moderate SNR these advantages disappear – this is especially so, if we want to operate with backwards compatible coding schemes at the encoder, such as those that 802.11 supports. Moreover, for small networks, such as the one we examine, DF is also (approximately) optimal and may perform better in practice. As far as we know, there has been no experimental evaluation of QMF relaying (before our work).

We present in this paper the first system design, deployment and experimental comparison of QMF, AF and DF relaying. We compare these results to a baseline direct transmission (DT) which does not employ a relay. Our system design emulates the 802.11 physical layer procedures, such as the frame structure, use of Orthogonal Frequency Division Multiplexing (OFDM) and use of standardized LDPC encoders. We implement all schemes and deploy them simultaneously on the same network to offer a fair comparison. We find that relaying schemes significantly outperform the baseline DT. We also find that QMF is a competitive scheme to DF and AF, offering in some cases up to 12% throughput benefits and up to 60% improvement in frame error-rates over the next best scheme. However, we also see that in a single-relay network, there are scenarios where DF does better than QMF.

Our main contributions are:

- We design and deploy the first QMF relaying system, as well as *coded* AF and DF relaying systems.
- We report a number of experiments that compare the three schemes, and evaluate the utility of relaying in the context of WiFi.

The rest of the paper is organized as follows. Section 2 presents our physical layer signaling as well as the network operation; Section 3 describes the relaying schemes and our designs for WiFi; Section 4 details the implementation of our system; and finally Section 5 provides and discusses our experimental results.

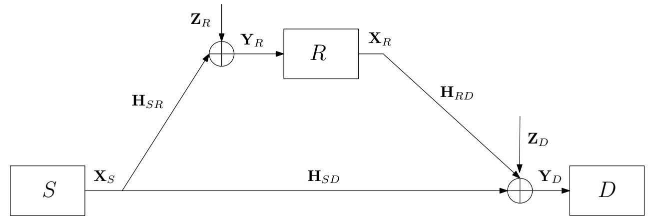

Figure 1: The 3-node network considered in this paper: the source S communicates with the destination D with the help of the relay R.

2. BACKGROUND AND MODEL

We consider a network that consists of three wireless nodes: a source S communicates with a destination D with the help of a relay R (see Fig. 1). The relay R is half-duplex in nature, i.e., it cannot simultaneously listen and transmit over the same frequencies. Our focus is on physical layer communication schemes: we are interested in comparing the performance of QMF with DF and AF, and more generally, evaluate how much physical layer cooperation through relaying can improve the performance as compared to routing. Our performance metrics are frame error rate and throughput, and are described in Section 5.

At a high level, the communication over our network depends on three choices: (i) the physical layer signaling; we emulate the physical layer operation prescribed by 802.11. We highlight the parts we need to describe our relaying schemes in Section 2.1 and give specific details in Section 4. (ii) the network operation; that is, whether we use the relay and in what network topology. For example, we may use the relay only when direct transmissions fail, or we may use it to form a two-hop path that connects the source to the destination. We describe our choices in Section 2.2. (iii) the relaying scheme; that is, if we do use the relay, which relaying strategy we deploy? We discuss several strategies and how we optimize their system implementation in Section 3.

2.1 Signal Model

The communication happens in frames. We work with coded systems, i.e., to improve the error resiliency, the source encodes the information bits using the 802.11 standard LDPC codes (we give more details on the codes we implement in Section 5). After encoding, the codeword bits are first mapped to constellation symbols (we use QAM constellations supported in WiFi), and then modulated using OFDM. We employ OFDM modulation as specified in the WiFi physical layer to combat channel frequency selectivity. We denote by $\mathbf{X} = [X_1, X_2, \ldots, X_m]^T$ an input vector to an OFDM modulator and by $\mathbf{Y} = [Y_1, Y_2, \ldots, Y_m]^T$ the output vector of the corresponding OFDM demodulator, where the subscript denotes the inputs and outputs at subcarrier $l \in \{1, \ldots, m\}$, with m being the total number of subcarriers (we use $m = 64$). The inputs X_l are complex numbers corresponding to the signal constellation used. With OFDM, the point-to-point signal model [24] for each subcarrier is:

$$Y_l = H_l X_l + Z_l, \tag{1}$$

where Z_l is the additive Gaussian noise and H_l is the channel for subcarrier $l \in \{1, \dots, m\}$.

In our network, both the source and the relay may transmit at the same time. From the superposition property of wireless channels, if the two transmitters S, R each emit an OFDM symbol $\mathbf{X}_i = [X_{i1}, X_{i2}, \dots, X_{im}]^T$, where $i \in \{S, R\}$, then the demodulated OFDM symbol at node D, denoted by $\mathbf{Y}_D = [Y_{D1}, Y_{D2}, \dots, Y_{Dm}]^T$, is given by [24]

$$Y_{Dl} = H_{SD,l}X_{Sl} + H_{RD,l}X_{Rl} + Z_{Dl}, \tag{2}$$

where $H_{SD,l}$ is the channel from S to destination D and $H_{RD,l}$ is the channel from R to destination D.

2.2 Network Operation

Even in a small network with a single relay, there are several possibilities for the *network operation*, *i.e.,* which links of the network we use and when we use them. We explore the following four choices in our experimentation.

1. Direct Transmission (DT): In the DT mode, we do not use the relay, i.e., the source S and destination D communicate directly using only the S–D direct link. The failure recovery is simply through retransmission of the same frame. We deploy the DT mode in our experiments to assess how useful the relay and physical layer cooperation are in our scenario.

2. Direct Transmission at Half Rate (DTHR): This scheme also uses the S–D direct link only, but incorporates a failure recovery mechanism through rate adaptation as advocated in the WiFi standard [1]. More specifically, we first attempt a direct transmission using a 16 QAM constellation. If this fails, we retransmit the same frame at half the rate. This is implemented by using a constellation of half the rate for the second transmission, i.e., 4 QAM[1]. As a result, the time duration of the second transmission doubles. Note that we need to account for the change in the rate when comparing the throughput performance with the other schemes. DTHR offers a recovery mechanism well suited to a direct link that is in general strong, but may infrequently experience deep fades.

3. Link Switching (LS): In LS we employ all three links: in a first attempt we employ broadcasting from the source, i.e., the S–D and the S–R links; in a second attempt we employ only the R–D link. First, the source broadcasts a frame that is intended for the destination. The relay R also overhears the source transmission. If the destination successfully decodes the frame through the direct link, we declare success and proceed with the next frame. However, if the first attempt fails and the destination cannot decode, then in a second attempt, the relay retransmits the frame it has overheard from the source, while the source remains silent. The operation for LS is summarized in Table 1, and the corresponding signal model is:

$$\begin{aligned} \mathbf{Y}_R^{(1)} &= \mathbf{H}_{SR}^{(1)}\mathbf{X}_S^{(1)} + \mathbf{Z}_R^{(1)} \\ \mathbf{Y}_D^{(1)} &= \mathbf{H}_{SD}^{(1)}\mathbf{X}_S^{(1)} + \mathbf{Z}_D^{(1)} \\ \mathbf{Y}_D^{(2)} &= \mathbf{H}_{RD}^{(2)}\mathbf{X}_R^{(2)} + \mathbf{Z}_D^{(2)}, \end{aligned} \tag{3}$$

where the superscript (k), $(k \in \{1, 2\})$, denotes the attempt, and $\mathbf{H}_{ij}^{(k)} = \text{diag}(H_{ij,1}^{(k)}, \dots, H_{ij,m}^{(k)})$ denotes the (sub-carrier) channels from node i to node j at attempt k. Note that since the transmissions from the source and the relay are orthogonalized in time, LS does not require synchronized transmissions.

Attempt	Source	Relay	Destination
T1	Transmit	Receive	Receive
T2	Silent	Transmit	Receive

Table 1: Schedule for Link Switching

LS can be interpreted as trying to utilize two paths: first the direct S–D link and then, if this fails, the S–R–D path. When the relay uses DF (Decode-Forward is described in Section 3) the LS mode of failure recovery is the single-relay analog of *routing* over larger networks.

4. Link Cooperation (LC): In LC as well, we employ all three links. Similar to LS, in a first attempt we employ broadcasting from the source, i.e., the S–D and the S–R links; however, if we fail, we then simultaneously employ both the S–D and R–D links. More specifically, if the destination fails to decode the first direct transmission of the source, in the second attempt both the source and relay cooperatively transmit. The schedule for LC is summarized in Table 2, and the corresponding signal model is:

$$\begin{aligned} \mathbf{Y}_R^{(1)} &= \mathbf{H}_{SR}^{(1)}\mathbf{X}_S^{(1)} + \mathbf{Z}_R^{(1)} \\ \mathbf{Y}_D^{(1)} &= \mathbf{H}_{SD}^{(1)}\mathbf{X}_S^{(1)} + \mathbf{Z}_D^{(1)} \\ \mathbf{Y}_D^{(2)} &= \mathbf{H}_{RD}^{(2)}\mathbf{X}_R^{(2)} + \mathbf{H}_{SD}^{(2)}\mathbf{X}_S^{(2)} + \mathbf{Z}_D^{(2)}, \end{aligned} \tag{4}$$

where the notation is similar to (3).

Attempt	Source	Relay	Destination
T1	Transmit	Receive	Receive
T2	Transmit	Transmit	Receive

Table 2: Schedule for Link Cooperation

LC is the mode of operation that is the most promising: it puts in use all the network resources–both broadcasting and transmit cooperation to forward the information and thus, if we use a "good" relaying scheme, it has the potential to offer the best performance no matter which of the network links are stronger. The question is whether one of the relaying schemes (that we describe in the next Section) makes this possible.

Discussion. The focus of our work is in evaluating the potential of PHY-layer cooperation; we leave open what a well-matched MAC-layer protocol would be. However, we do not think that heavy MAC-layer redesign would be required; there exist schemes in the literature that already implement the simple functionalities we would need. For instance, the modes that use a relay (LS and LC) operate *on-demand*, i.e., the relay is invoked only when the first attempt to communicate with the destination has failed. MAC-layer protocols for activating on-demand relaying in 802.11-like networks have been discussed in [26, 11, 9].

3. RELAYING SCHEMES

In this section we describe the three relaying schemes we will compare: Decode-Forward (DF), Amplify-Forward (AF) and Quantize-Map-Forward (QMF). We first give the

[1]Clearly, adaptation to different rates is possible; we use this to get an indicative performance benchmark and because it ties well with 802.11 rate adaptation mechanisms.

high level operation and then describe our system design choices. The relaying schemes are only relevant for the LS and LC network operations (since in DT and DTHR we do not employ the relay at all).

3.1 Relay Operation

Decode Forward (DF): Once the relay receives a frame, it attempts to decode it and retrieve the information bits. If decoding is successful (i.e., the CRC passes *after* running the LDPC decoder), the relay can re-encode the decoded information, remodulate and forward the frame to the destination. If correct decoding is possible, it is an optimal relay operation, as it removes the S–R link noise. If decoding at the relay fails however, the relay remains silent and cannot cooperate for that frame.

Amplify Forward (AF): The relay does not attempt to decode; rather it simply amplifies its received signal to the maximum transmit power of its radio (by multiplying with an appropriate amplification factor A) and retransmits it. This is advantageous in cases where the relay cannot decode, but the destination can do so with the help of both the source and relay transmissions.

Quantize-Map-Forward (QMF): The relay quantizes the received symbols, collects a sequence of quantized values, and operates on the entire sequence to produce a transmit sequence. This is distinct from DF since the relay does not decode and is distinct from AF since it operates on a *sequence* of quantized symbols.

3.2 System Design

3.2.1 Cooperative Transmissions

In the LC mode, effectively we have the source and relay cooperating like a *distributed* 2×1 transmit antenna system to the destination. For point-to-point links with 2 transmit and 1 receive antennas, the WiFi standard recommends an Alamouti space-time code [3, 1], as it gives the best rate-reliability tradeoff for the 2×1 MISO channel, asymptotically in SNR (*i.e.,* it is diversity-multiplexing-tradeoff optimal). Accordingly, to optimize the performance of AF and DF, we implement a *distributed* Alamouti code in the LC mode. QMF cannot take advantage of Alamouti coding, because of the random mapping at the relay. We describe in the following, the transmissions for all schemes.

Decode Forward (DF): In the LC mode, if the relay successfully decodes the source frame, then we can exactly create a 2×1 multiple transmit antenna situation. That is, we can implement a distributed Alamouti scheme by using the decoded symbols at the relay in cooperation with the source. If S_1 and S_2 are two QAM symbols transmitted on a particular subcarrier l over two OFDM symbols, then the transmission scheme for the distributed Alamouti scheme is given in Table 3. Using the OFDM signal model of (2), the demodulated symbols at the destination D for that subcarrier l are

$$\begin{aligned} Y_{Dl}^{(2,1)} &= H_{RD,l}^{(2)} S_1 + H_{SD,l}^{(2)} S_2 + Z_{Dl}^{(2,1)} \\ Y_{Dl}^{(2,2)} &= -H_{RD,l}^{(2)} S_2^* + H_{SD,l}^{(2)} S_1^* + Z_{Dl}^{(2,2)}, \end{aligned} \tag{5}$$

where $Y_{Dl}^{(2,k)}$ is the demodulated signal on subcarrier l at D across two OFDM symbols, $k = 1, 2$ and $H_{RD,l}^{(2)}$, $H_{SD,l}^{(2)}$ denote the channels on subcarrier l in attempt 2. Standard Alamouti combining at D [3] results in the following effective point-to-point channels per subcarrier:

$$\begin{aligned} \tilde{Y}_{Dl}^{(2,1)} &= S_1 \sqrt{|H_{RD,l}^{(2)}|^2 + |H_{SD,l}^{(2)}|^2} + \tilde{Z}_{Dl}^{(2,1)} \\ \tilde{Y}_{Dl}^{(2,2)} &= S_2 \sqrt{|H_{RD,l}^{(2)}|^2 + |H_{SD,l}^{(2)}|^2} + \tilde{Z}_{Dl}^{(2,2)}, \end{aligned} \tag{6}$$

where $\tilde{Z}_{Dl}^{(2,k)}$, $k = 1, 2$ is Gaussian noise of the same variance as the noise in (5).

	OFDM Symbol 1	OFDM Symbol 2
Relay	S_1	$-S_2^*$
Source	S_2	S_1^*

Table 3: Transmitted signals from source and relay per subcarrier for DFLC

Amplify Forward (AF): For AF, since the relay does not decode, the distributed Alamouti scheme of Table 3 is modified by using the amplified received signal at subcarrier l, instead of the decoded symbols as in DF (see Table 4).

	OFDM Symbol 1	OFDM Symbol 2
Relay	$AY_{Rl}^{(1,1)}$	$AY_{Rl}^{(1,2)}$
Source	S_2	S_1^*

Table 4: Transmitted signals from source and relay per subcarrier for AFLC

where

$$\begin{aligned} Y_{Rl}^{(1,1)} &= H_{SR,l}^{(1)} S_1 + Z_R^{(1,1)} \\ Y_{Rl}^{(1,2)} &= -H_{SR,l}^{(1)} S_2^* + Z_R^{(1,2)} \end{aligned}$$

The received signals are modified from (5) as:

$$\begin{aligned} Y_{Dl}^{(2,1)} &= H_{RD,l}^{(2)\prime} S_1 + H_{SD,l}^{(2)} S_2 + Z_{Dl}^{(2,1)\prime} \\ Y_{Dl}^{(2,2)} &= -H_{RD,l}^{(2)\prime} S_2^* + H_{SD,l}^{(2)} S_1^* + Z_{Dl}^{(2,2)\prime} \end{aligned} \tag{7}$$

where $H_{RD,l}^{(2)\prime} = AH_{RD,l}^{(2)} H_{SR,l}^{(1)}$. Note that the noise-variances corresponding to $Z_{Dl}^{(2,k)\prime}$ are larger than $Z_{Dl}^{(2,k)}$ in (5) since we are forwarding noise in AF. Using the same Alamouti combining at D as in (6), we get

$$\begin{aligned} \breve{Y}_{Dl}^{(2,1)} &= S_1 \sqrt{|H_{RD,l}^{(2)\prime}|^2 + |H_{SD,l}^{(2)}|^2} + \breve{Z}_{Dl}^{(2,1)} \\ \breve{Y}_{Dl}^{(2,2)} &= S_2 \sqrt{|H_{RD,l}^{(2)\prime}|^2 + |H_{SD,l}^{(2)}|^2} + \breve{Z}_{Dl}^{(2,2)}, \end{aligned} \tag{8}$$

where $\breve{Z}_{Dl}^{(2,k)}$, $k = 1, 2$ is still Gaussian noise with same (larger) variance as $Z_{Dl}^{(2,k)\prime}$ in (7).

Quantize-Map-Forward (QMF): We implement it as follows: the relays employ a scalar quantizer (symbol-by-symbol on a QAM symbol level) to quantize the received signals from the source (as proposed in [22, 4]). Instead of the random mapping operation proposed in [5], we use a permutation mapping (randomly chosen bit-interleaver) on the quantized bits corresponding to each codeword in the frame. This mapping was shown to perform well in [22], and it also facilitates a simpler decoder operation. After mapping, the relay re-modulates the frame as per the 802.11 specifications, and if the direct transmission in the first time slot fails, forwards it to the destination.

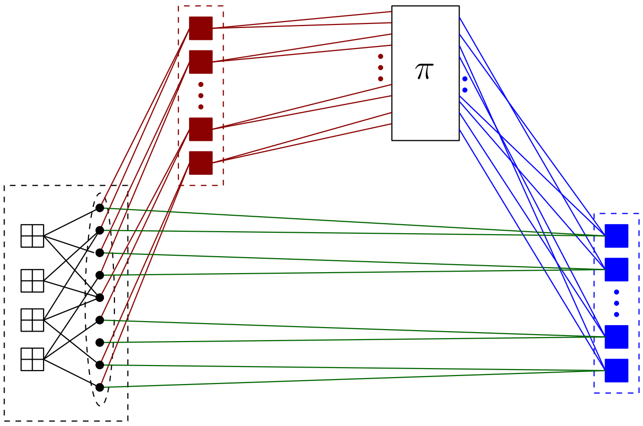

Figure 2: LDPC-based joint decoder for QMF

3.2.2 *Decoding at the Receiver*

In attempt 2 of the LC mode, as we already described in (2), the destination receives a superposition of the source and relay transmissions

$$Y_{Dl} = H_{SD,l}X_{Sl} + H_{RD,l}X_{Rl} + Z_{Dl}.$$

For AF and DF, we have already seen that effectively we have point-to-point channels with parameters that take into account the operations at the relay. Thus, standard point-to-point LDPC decoders are sufficient for decoding AF and DF frames.

QMF, on the other hand, uses non-linear operations (quantization) at the relay and hence a point-to-point decoder does not suffice for decoding. As an example, Fig. 2 outlines the graphical structure of the decoder for one codeword with QMF LC. This is *not* a standard structure for iterative (LDPC) codes due to two reasons: (i) the quantization at the relay which takes the received signal Y_{Rl} and quantizes it to $\hat{Y}_{Rl}$. (ii) the superposition of the two streams X_{Sl}, X_{Rl}.

For the QMF iterative decoder, we introduce two novel features: (i) We use the stochastic quantizer check nodes (red boxes in Fig. 2) for the functions $p(\hat{Y}_{Rl}|X_{Sl})$ that capture the quantization operation at the relay. We use the multiple-access check nodes (blue boxes in Fig. 2) to capture the signal superposition at the destination.(ii) We use a graphical structure *across* the nodes of the network and run our decoding iterations over the graph. That is, the iterative decoder runs across the graphical structure of Fig. 2, with the bit level reliabilities (represented by log-likelihood ratios, LLRs) flowing along the edges in the graph.

In addition, to reduce the occurrence of low-weight error events, we use a further *enhancement layer* for the LDPC decoders of all three schemes, QMF, AF and DF, that treats LLR magnitudes above a certain pre-determined threshold as correct, and others as *erasures* (lost bits) after a fixed number of iterations. Then another erasure correction round of iterations is performed to clean up the residual errors.

4. SYSTEM IMPLEMENTATION

Overall system. The source, relay, and destination nodes were implemented using the WARP SDR platform [2]. We used the WARPLab framework, which allows interaction with the WARP hardware via a host PC running MATLAB (see Fig. 3). The host PC was used to control real-time over-the-air transmission and reception.

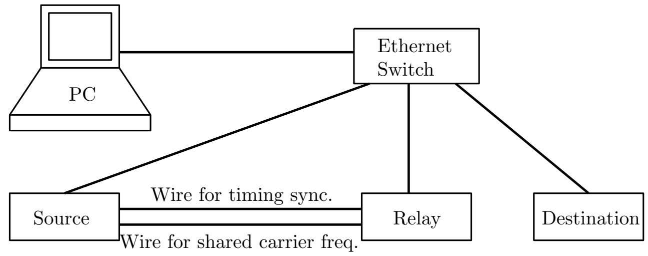

Figure 3: Node and host PC configuration

We implemented the four network operation methods described in Section 2.2 coupled with the three relaying strategies defined in Section 3.1. Table 5 summarizes the implemented schemes when the relay is active; additionally, we implemented Direct Transmission (DT) and Direct Transmission Half Rate (DTHR).

4.1 Frame Structure and Operation

We use a frame structure specified in the 802.11 standard [1], which stipulates that we first transmit training for Automatic Gain Control (TAGC), followed by training for timing synchronization (TSYNC), training for channel estimation (TCHE) and then the payload. We use the signal that the 802.11 standard defines as long training symbols for timing synchronization and channel estimation. Our experiments correspond to a 20 MHz bandwidth system with 64 subcarriers. During the payload transmission, for each OFDM symbol transmitted, 48 subcarriers are data subcarriers, 4 subcarriers carry pilots, and the rest of the subcarriers are unused, as per the 802.11 payload structure. The signals transmitted over-the-air were centered at a frequency of 2.4 GHz in a 20MHz bandwidth. Next, we give more specific details of the frame structures for each scheme for the attempts over the two time-slots.

Time-slot operation: To make a fair comparison between schemes with and without cooperation, we allow a maximum of two attempts for every source information frame to be successful, as described in Section 2.2. We send each attempt over a *time slot*, which is one frame in length. If the communication is not successful after two attempts then the frame is declared in error and the source moves on to the first attempt for the next information frame.

***First Time Slot*:** For all schemes, the first attempt is always only via the direct link from source to destination. The source transmits using the standard frame structure we previously described; the frame structure is shown in Fig. 4. If the destination is able to decode successfully, we proceed with the first attempt of a new frame. If the transmission attempt fails, then we move to the next attempt in the second time-slot described below.

***Second Time Slot*:** The source re-transmits at lower rate in DTHR mode. In LS mode, only the relay transmits, using a standard frame structure. In the LC mode, both the source and the relay transmit. The different relaying operations for each of the modes are described in Section 3.1. The frame structures for the second attempt are shown in Fig. 4. As noted in Fig. 4, in AFLC/AFLS and QMFLC/QMFLS, the relay sends one more OFDM symbol than in DFLC/DFLS. This extra OFDM symbol is only sent by the relay and it is

	Amplify-Forward (AF)	Decode-Forward (DF)	Quantize-Map-Forward (QMF)
Link Switching (LS)	AFLS	DFLS	QMFLS
Link Cooperation (LC)	AFLC	DFLC	QMFLC

Table 5: Implemented schemes when the relay is active. The network modes (rows) are described in Section 2.2 and the relaying strategies (columns) in Section 3.1.

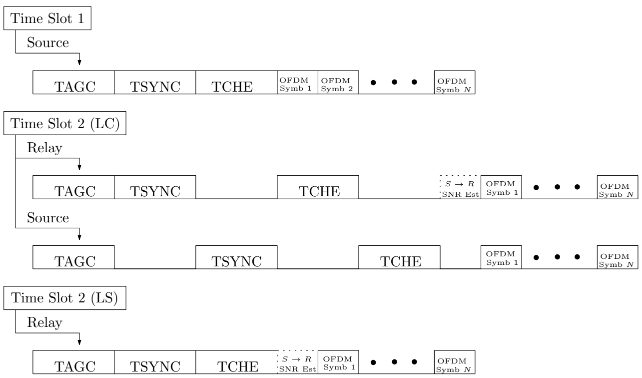

Figure 4: Time-slot frame structure for LS and LC.

used to forward an estimate of the source to relay SNR. This estimate is used by the iterative decoder at the destination. The SNR estimate is not forwarded by DFLC/DFLS because the bits are decoded at the relay before forwarding them.

4.2 Timing and Carrier Synchronization

The WARPLab framework allows a coarse timing synchronization of the nodes involved in communication. However, for 20MHz OFDM time scales this coarse synchronization is not accurate enough; more elaborate mechanisms that enable the required timing and carrier frequency synchronization are required both at the receivers and distributed transmitters.

Operation at a receiver: A node in receiver mode can solve timing synchronization by exploiting the autocorrelation properties of the training sent for timing synchronization. The signal used for channel estimation, labeled as TCHE in Fig. 4, is also used for time-domain estimation and correction of the carrier frequency offset (CFO). The time domain correction is applied to the training for channel estimation and to the OFDM symbols before performing the FFT for conversion to the frequency domain. To correct for residual phase noise present, another level of estimation and correction of CFO and a phase noise correction is applied in the frequency domain. The frequency domain CFO and phase noise are estimated based on the four known pilot subcarriers that are included in each OFDM symbol. Since DT/DTHR/LS modes effectively create single-transmitter, single-receiver (point-to-point) channels, we distinguish its operation from the LC modes.

Link Cooperation (LC): In order for the destination to be able to set the AGC correctly for the reception of the sum signal, the source and relay send training for AGC simultaneously. The waveform for AGC sent by the source is similar to the one sent by the relay, with the only difference being that it is cyclically shifted in order to avoid accidental nulling. The source and relay both send training for timing synchronization and channel estimation and these training signals are orthogonal in time. This orthogonality ensures that the two node diversity is also present in the timing synchronization phase, since the destination can solve timing synchronization from any of the two copies it receives. The orthogonal training for channel estimation is needed in order to compute clean channel estimates from each of the links; these two channel estimates are needed for decoding in the LC mode. To implement the orthogonality of training/timing in link cooperation mode, the four pilot subcarriers are split and alternated between the source and the relay in the following way. In odd OFDM symbols the relay transmits pilots 1 and 3 and the source transmits pilots 2 and 4. In even OFDM symbols the relay transmits pilots 2 and 4 and the source transmits pilots 1 and 3. This is the same method for pilot assignment proposed in [17, 21].

The mechanism described above for CFO correction enables a receiver to solve for only one CFO. However, in LC mode, at the destination there would be two CFOs to correct for–one due to the R–D link and another due to the S–D link. In [18, 21] this issue of two CFOs is solved by having the relay lock on to the carrier of the source by estimating its CFO with respect to the source and applying a time domain correction so that the destination observes only one CFO. Although feasible to implement, this was not the focus of our study. Consequently, for our implementation, we used a wire to share the carrier frequency clock of the source between the source and the relay, as shown in Fig. 3.

Distributed synchronization: In LC, in the second time slot, the relay and the source transmit synchronously to ensure required orthogonality in the training phase and avoid intersymbol interference when receiving the payload. Thus, the synchronization mechanism must be accurate enough to ensure that at the destination, the time of arrival of the signals from the source and the relay is within the cyclic prefix of the OFDM symbols. In our implementation, the duration of cyclic prefix is 0.8 μs, which is the exact same duration of the cyclic prefix specified in 802.11 for a 20 MHz bandwidth system. Recent work [18, 21, 6] has demonstrated that one can design protocols that achieve accurate timing synchronization (between 20 ns and 100 ns accuracy) for distributed OFDM communication enabled by implementing a large part of the distributed timing synchronization mechanisms in real time in the FPGA. Incorporating this fast turnaround time to the WARPLab framework, although feasible, was not the focus of our study. Consequently, we synchronized the time of transmission of distributed transmitters via a wired connection between source and relay, as shown in Fig. 3, for all schemes.

4.3 Estimation of Effective Noise

To decode, all our schemes require an estimate of the effective noise variance in the digital (sampled) domain. This estimate is used to compute the log likelihood probabili-

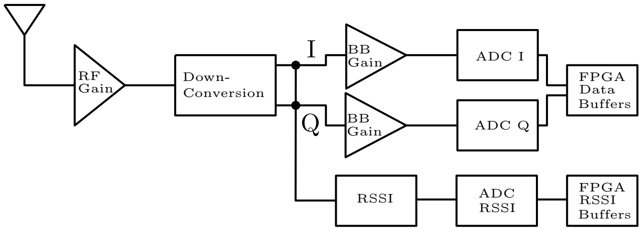

Figure 5: Radio board receiver path

ties that serve as the input to the decoders, and thus an accurate estimation is important for the iterative (LDPC) decoder performance. However, accurate estimation of the noise variance can be a challenging and complex task, since there are many sources of noise in a hardware implementation. For example, there is additive noise, quantization noise at ADC and DAC, IQ imbalance, channel estimation errors, carrier frequency offset errors and timing synchronization inaccuracy. Also, the noise added by the radios is a function of the gains set by the AGC and the received signal power. To balance the need for an accurate estimate and the system need for low complexity, we developed the following simple algorithm that builds on four components: (i) the (analog) RSSI measurement provided by the radios (ii) the approximate (calibrated) noise level of the radios we use (iii) the channel estimates per subcarrier (iv) the model for the received signal. We describe how we put these together in the following two steps.

• *Step 1 - Computation of the Analog SNR:* We compute the SNR in the analog domain as

$$\text{SNR-analog } (dB) = \text{PayloadRSSI } (dBm) - N_0 \ (dBm) \quad (9)$$

The analog thermal noise value N_0 is initialized with the calibrated value for the WARP boards. The receiver radio path is shown in Fig. 5. The RSSI reading is made after downconversion to baseband and the radios provide the value of the receiver RF Gain and the mapping of RSSI readings to RSSI at the receiver antenna. Hence, from the RSSI reading we actually get the RSSI of the signal at the receiver antenna. This is the PayloadRSSI that we use to compute the SNR-analog in (9). The analog RSSI values are reported every 0.1μsec by the radios or 40 RSSI readings per OFDM symbol. Hence, for a payload of L OFDM symbols we have $40L$ RSSI readings that are averaged to get the payload RSSI represented in (9).

• *Step 2 - Solve for the effective digital noise variance:* In the digital domain, we model the received signal as shown in (1). Since we have m subcarriers, the SNR in the digital domain can be estimated as

$$\text{SNR-digital} = \frac{\sum_{l=1}^{m} |H_l|^2}{m\sigma^2}, \quad (10)$$

where σ^2 is effective (digital) noise power per subcarrier (the signal power is normalized to 1). We estimate that the SNR computed in the digital domain is the same as the SNR computed in the analog domain minus an estimate Δ of effects like imperfect channel estimation and imperfect CFO correction.

$$\text{SNR-digital } (dB) = \text{SNR-analog } (dB) - \Delta \ (dB) \quad (11)$$

From (9), (10) and (11) we solve for the effective noise variance σ^2, which is used in the iterative decoder. Guided by the WARP radio calibrations, we use values of $N_0 = -95\ dBm$ and $\Delta = 5\ dB$.

As for any estimator, the method for σ^2 estimation described in the two steps above can have errors that affect the decoder performance. To the best of our knowledge, there is currently no work on the comparison of DF, AF and QMF schemes under errors in σ^2. We leave the development and implementation of more robust estimators for future work.

For AF and QMF, the decoding at the destination needs the SNR-digital computed for the source to relay link. The relay computes this information and forwards it to the destination using one extra OFDM symbol, as shown in Fig. 4. To do so, we first quantize the SNR estimate to one out of 40 possible values ranging from −10 to 30 dB (in steps of 1 dB). We can describe these 40 values using 6 bits. We repeat these 6 bits 8 times, modulate them with BPSK and allocate them to the 48 data subcarriers in the OFDM symbol used to forward the SNR information.

5. EXPERIMENTATION

Setup. Our testbed covers a rectangular area of 175 m^2, and spans across three rooms. We report the results for three scenarios, which we describe in terms of the average Received Signal Strength Indicator (RSSI) in dBm. We note that the RSSI values not only depend on the distance between devices, but also on the multipath effects in our indoor environment. Moreover, in addition to the relative positioning, we also needed to adjust the transmit power of the radios in some cases, in order to achieve the RSSI values for our scenarios that are typical for 802.11 systems and recommended for the modulation we employ. The RSSI values for all the scenarios is given in Fig. 7.

- **Scenario 1** captures the case where the source and the relay are close: we deployed the relay at a distance of 1.5m from the source as depicted in Fig. 6(a). As a result, the S–R link is much stronger than the S–D and R–D links, which is reflected in the average RSSI values we measure at our receivers.

- **Scenario 2** captures the case where all three links have approximately equal strength. The node placement is shown in Fig. 6(b).

- **Scenario 3** captures the case where the S–R and R–D links are much stronger than the S–D direct link. The nodal arrangement is depicted in Fig. 6(c).

In all our experiments we used 16 QAM modulation, an LDPC code of rate 3/4 with codeword length of 1944 bits and the parity matrix defined in the 802.11 standard. Each frame encapsulated four codewords, leading to 5832 information bits (payload) per frame. For each experiment we transmitted at least 1600 frames, which corresponds to at least 9331200 information bits transmitted per experiment.
Metrics: We use two performance metrics:
(i) *Frame Error Rate (FER)* quantifies the error probability our frame transmissions have with each scheme.
(iii) *Throughput* quantifies the amount of successfully decoded information bits that reach the destination, normalized by the the total over-the-air transmission time and

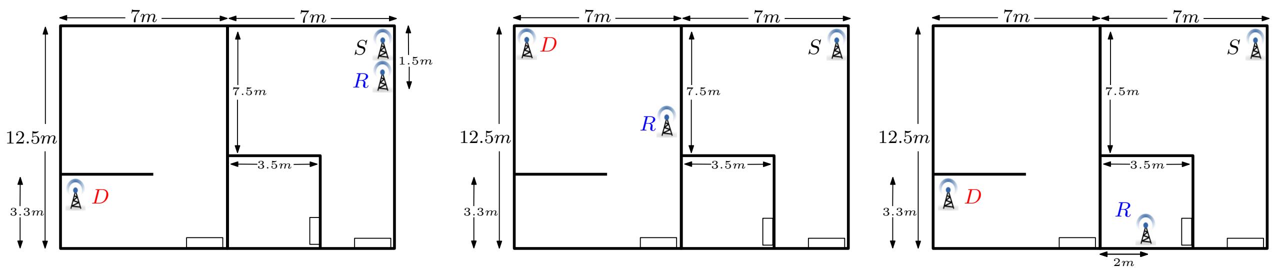

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3

Figure 6: Node placement in our testbeds

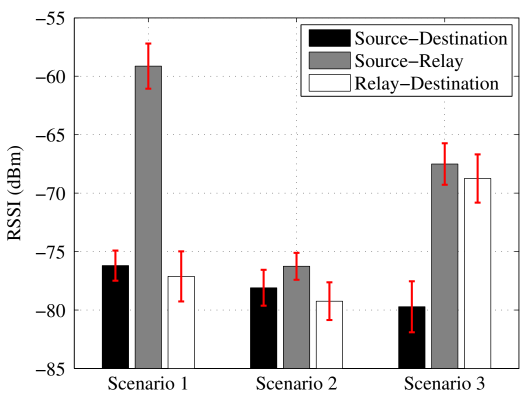

Figure 7: Average measured RSSI values per scenario.

bandwidth used. Thus, we count the amount of decoded information bits per channel use in bps/Hz.

Discussion of results. The following observations are based on the FER results shown in Fig. 8 and the throughput results shown in Fig. 9.

(1) *Relaying helps over DT:* In all three scenarios, LC schemes significantly outperform DT. The FER gains of relay transmission ranged from a factor of 10 in Scenarios 1 and 2, to nearly a factor of 100 in Scenario 3, where the direct path was very weak. The throughput gains were 20-40% in Scenarios 1 & 2, and was nearly a factor of 15 in Scenario 3. DTHR[2] can improve the throughput performance with respect to DT, for example in Scenarios 2 and 3. However, in all the scenarios the performance of DTHR was always worse than all of the LC schemes. This is because the reliability benefit of transmitting at a lower rate does not offset the diversity gain obtained with link cooperation.

(2) *Link Cooperation (LC) helps:* In Scenarios 1 and 2, link cooperation gives a factor of 10–40 gain in FER and a 40–70% gain in throughput over link switching. In Scenario 3, link switching performs marginally better than link cooperation. There are two main reasons why in Scenario 3 LC does not improve performance over LS. First, at an average RSSI of −80 dBm, the S–D link in Scenario 3 is 13 dB weaker than the R–D link. Hence, the contribution from the S-D link is minimal. Second, in LC mode, the pilots for frequency domain CFO and phase noise corrections are split between the source and the relay as was explained in Section 4.2. However, since the S–D link is very weak, the pilots assigned to the source are received with very weak power and hence do not help towards CFO or phase noise correction. Consequently, in Scenario 3, the last hop in LC is effectively a single R–D link, reminiscent of LS, with the added disadvantage of having only half the pilots as compared to LS.

(3) *Universality of QMF:* QMFLC shows the most competitive performance across all the scenarios. QMF outperforms AF because the frames it sends in the second time-slot are less noisy since the quantization removes some of the noise. In Scenarios 1 and 3 there is very little difference between the performance of QMF and DF. In these two scenarios, since the S–R link is at a medium-to-high RSSI (−60 dBm in Scenario 1 and −67 dBm in Scenario 3), DF can decode most of the time and hence exhibits good performance. In Scenario 2, the S–R link is at a low RSSI (−78 dBm); hence decoding at the relay fails more often, affecting the performance of DF negatively. As a result, the gains from QMF are most pronounced in Scenario 2, as the relay always transmits a reasonable quantized signal that the destination can exploit in decoding. Scenario 2 shows a 12% throughput benefit for QMFLC and up to 60% improvement in FER over the next best scheme (DFLC).

Note: Although we observe that cooperation helps and QMFLC outperforms the other schemes in Scenarios 1 and 2, we neither believe nor expect that this is true for all possible scenarios; for example, if the direct link S–D is much stronger than the rest, we expect DT to do best; and if this link essentially does not exist (effectively creating a line network as seen in Scenario 3) DFLS (*i.e.,* routing) does better than all others schemes. In this paper we focused on scenarios where the outcome of the comparison is not obvious, and thus more interesting.

Why frames are lost: We grouped frame failures in categories: failed due to timing, or a certain number of bit errors between $[a, b]$ remained in the decoded frame (and thus the CRC check failed). We plot an example of this analysis for Scenario 2 in Fig. 10. We observe that for DF and QMF, frame errors primarily occur when a *large* fraction of the bits are in error, indicating that relay processing reduces low-weight error patterns. In contrast, in AF we see frames failing even with a small number of bit errors: a plausible explanation being that since AF does not process its received signal at the relay (DF decodes, QMF quantizes and remaps), the amplified noisy signals propagate to the destination resulting in a (relatively) less sharp threshold behavior in the LDPC decoder.

[2]For DTHR, we only compared throughput since the transmission rate was lower than the other schemes. This implied that the FER was not directly comparable.

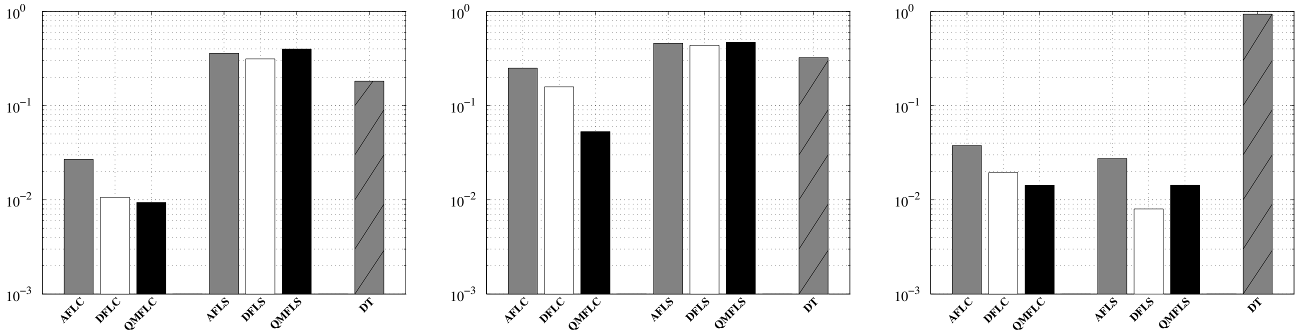

(a) FER Scenario 1 (b) FER Scenario 2 (c) FER Scenario 3

Figure 8: Frame error rates.

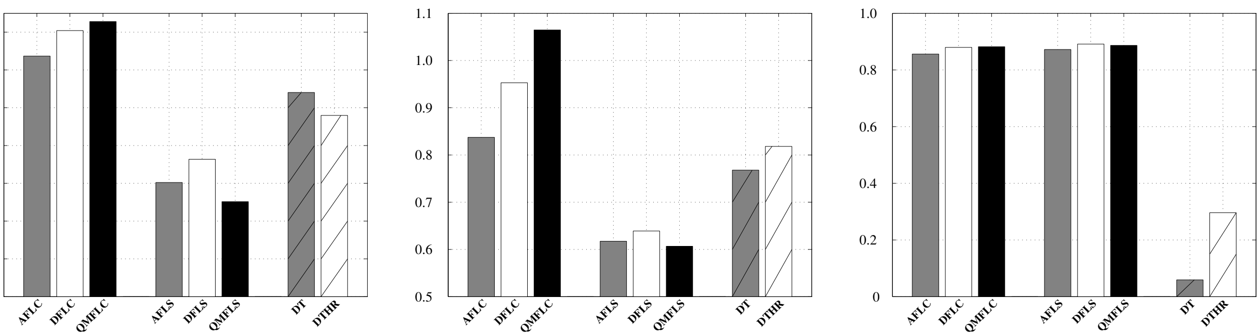

(a) Throughput Scenario 1 (b) Throughput Scenario 2 (c) Throughput Scenario 3

Figure 9: Throughput in bps/Hz.

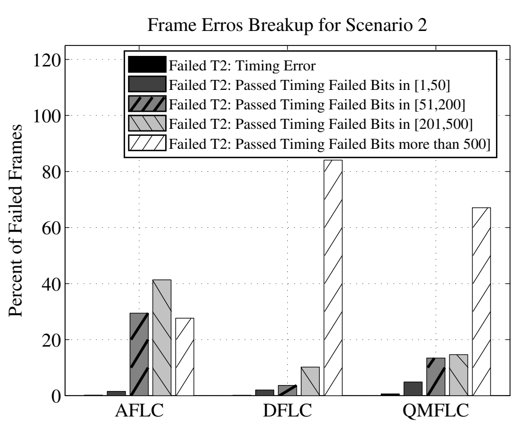

Figure 10: Distribution of percentage of frames that fail in $T2$.

6. RELATED WORK

Surveys covering the field of testbed implementations of physical layer cooperative communications are found in [13], [8]. In [7], a testbed implementation of PHY layer cooperation using uncoded DF on a single relay network is investigated. A WARP radio testbed implementation of cooperative relaying is presented in [12] where again DF based schemes are implemented. Both [7, 12] assume orthogonal transmissions from the source and relay. In [25] uncoded DF protocols with one and two relays on the Universal Software Radio Peripheral (USRP) platform are studied. In all these studies advanced error correction or broadband OFDM modulation are not implemented as done in our paper. The work [18] is perhaps closest to our paper. They implement both (uncoded) AF and DF relaying strategies along with distributed Alamouti-based transmission. It is important to note however, that this implementation also did not use error correcting codes at any point in the network. Other testbed implementations that specifically focus on tackling the synchronization problem for multiple simultaneous transmissions are found in [21] and [26].

Testbed implementations focussing on the benefits of MAC-level cooperative communication over WiFi was presented in [15], using (DF-based) multi-hop transmission. Other approaches to increase the throughput of WiFi include intelligently selecting access points based on channel utilization and adaptive switching [23], and recommendations to improve the way in which rate adaptation is done in WiFi ([16] and references therein).

Many of the fundamental information theoretic ideas of relaying were laid in the seminal work by Cover and El-Gamal [10], where the DF strategy was formalized. They also proposed Compress-Forward (CF). which is closely related to QMF proposed in [5]. The distinction is that the CF needs complete network state at the relays and therefore does not allow distributed operation. QMF also extends seamlessly for multicast, which is not possible with CF [5]. It was shown that QMF approximately achieves network capacity (within a constant gap, independent of the channel gains and SNR). Similar schemes using lattice-based codes were developed in [20] and QMF was extended to discrete memoryless networks in [14]. Practical coding schemes for QMF relaying with LDPC codes were proposed for an orthogonalized version of the half-duplex single relay network in [19] and for the full-duplex diamond network in [22], where interleavers as relay maps were also proposed. The use of demodulation instead of quantization was used in [22] and [4]. None of these papers had an experimental evaluation of performance.

7. ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their comments that improved the presentation of the work. We also acknowledge funding from the European Research Council grant NOWIRE ERC-2009-StG-240317, the EU project CONECT FP7-ICT-2009-257616, the NSF award 1136174 and MURI award AFOSR FA9550-09-064.

8. REFERENCES

[1] Local and metropolitan area networks–specific requirements part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. *IEEE Std 802.11-2012*.

[2] WARP Project, http://warpproject.org.

[3] Alamouti, S. A simple transmit diversity technique for wireless communications. *IEEE Journal on Select Areas in Communications 16*, 8 (October 1998), 1451–1458.

[4] Atsan, E., Knopp, R., Diggavi, S., and Fragouli, C. Towards integrating quantize-map-forward relaying into LTE. In *Proceedings of the IEEE Information Theory Workshop* (September 2012), pp. 212–216.

[5] Avestimehr, A. S., Diggavi, S. N., and Tse, D. N. C. Wireless network information flow: A deterministic approach. *IEEE Transactions on Information Theory 57*, 4 (April 2011), 1872–1905.

[6] Balan, H. V., Rogalin, R., Michaloliakos, A., Psounis, K., and Caire, G. AirSync: Enabling distributed multiuser MIMO with full spatial multiplexing. *IEEE/ACM Transactions on Networking*, 99 (2013).

[7] Bradford, G., and Laneman, J. N. An experimental framework for the evaluation of cooperative diversity. In *Proceedings of the IEEE CISS* (March 2009), pp. 641–645.

[8] Bradford, G., and Laneman, J. N. A survey of implementation efforts and experimental design for cooperative communications. In *Proceedings of the IEEE ICASSP* (2010), pp. 5602–5605.

[9] Chang, K. et al. Relay operation in 802.11ad. In *IEEE 802.11ad TGad 1-/0494r1* (2010).

[10] Cover, T. M., and Gamal, A. E. Capacity theorems for the relay channel. *IEEE Transactions on Information Theory 25*, 5 (September 1979), 572–584.

[11] Hunter, C., Murphy, P., and Sabharwal, A. Real-time testbed implementation of a distributed cooperative MAC and PHY. In *Proceedings of the IEEE CISS* (March 2010), pp. 1–6.

[12] Knox, M., and Erkip, E. Implementation of cooperative communications using software defined radios. In *Proceedings of the IEEE ICASSP* (March 2010), pp. 5618–5621.

[13] Korakis, T., Knox, M., Erkip, E., and Panwar, S. Cooperative network implementation using open-source platforms. *IEEE Communications Magazine 47*, 2 (2009), 134–141.

[14] Lim, S. H., Kim, Y.-H., Gamal, A. E., and Chung, S.-Y. Noisy network coding. *IEEE Transactions on Information Theory 57*, 5 (May 2011), 3132–3152.

[15] Liu, P., Tao, Z., Narayanan, S., Korakis, T., and Panwar, S. CoopMAC: A cooperative MAC for wireless LANs. *IEEE Journal on Selected Areas in Communications 25*, 2 (February 2007), 340–354.

[16] Loiacono, M., Rosca, J., and Trappe, W. The snowball effect: Detailing performance anomalies of 802.11 rate adaptation. In *Proceedings of the IEEE GLOBECOM* (November 2007), pp. 5117–5122.

[17] Murphy, P. *Design, Implementation, and Characterization of a Cooperative Communications System*. PhD thesis, Rice University, 2010.

[18] Murphy, P., and Sabharwal, A. Design, implementation, and characterization of a cooperative communications system. *IEEE Transactions on Vehicular Technology 60*, 6 (July 2011), 2534–2544.

[19] Nagpal, V., Wang, I.-H., Jorgovanovic, M., Tse, D., and Nikolić, B. Coding and system design for quantize-map-and-forward relaying. *IEEE Journal on Selected Areas in Communications* (August 2013). See also arXiv:1209.4679 [cs.IT].

[20] Özgür, A., and Diggavi, S. Approximately achieving Gaussian relay network capacity with lattice codes. In *Proceedings of the IEEE International Symposium on Information Theory* (June 2010), pp. 669–673. See also arXiv:1005.1284 [cs.IT].

[21] Rahul, H., Hassanieh, H., and Katabi, D. SourceSync: a distributed wireless architecture for exploiting sender diversity. In *Proceedings of the ACM SIGCOMM* (August 2010), pp. 171–182.

[22] Sengupta, A., Brahma, S., Ozgur, A., Fragouli, C., and Diggavi, S. Graph-based codes for quantize-map-and-forward relaying. In *Proceedings of the IEEE Information Theory Workshop* (October 2011), pp. 140–144.

[23] Sun, T., Zhang, Y., and Trappe, W. Improving access point association protocols through channel utilization and adaptive switching. In *Proceedings of the 8th IEEE International Conference on Mobile Ad-Hoc and Sensor Systems* (October 2011), pp. 155–157.

[24] Tse, D., and Viswanath, P. *Fundamentals of Wireless Communication*. Cambridge University Press, May 2005.

[25] Zhang, J., Jia, J., Zhang, Q., and Lo, E. M. K. Implementation and evaluation of cooperative communication schemes in software-defined radio testbed. In *Proceedings of the IEEE INFOCOM* (March 2010), pp. 1307–1315.

[26] Zhang, X., and Shin, K. G. DAC: Distributed asynchronous cooperation for wireless relay networks. In *Proceedings of the IEEE INFOCOM* (March 2010), pp. 1064–1072.