COLiDeR: A Cross-Layer Protocol for Two-Path Relaying

Raphaël Naves
Thales SIX GTS France
raphael.naves@thalesgroup.com

Gentian Jakllari
IRIT-INPT/ENSEEIHT - France
gentian.jakllari@enseeiht.fr

Hicham Khalifé
Thales SIX GTS France
hicham.khalife@thalesgroup.com

Vania Conan
Thales SIX GTS France
vania.conan@thalesgroup.com

André-Luc Beylot
IRIT-INPT/ENSEEIHT - France
andre-luc.beylot@enseeiht.fr

ABSTRACT

In this work, we present COLiDeR, the first practical system for two-path relaying using off-the-shelf half-duplex radios. While two-path relaying has been mostly studied from a theoretical perspective in the literature, our solution addresses the challenges of making it practical by introducing two key contributions. First, using a measurement-driven approach, we identify the best approach for a radio to handle two overlapping signals. Just as important, we carefully quantify its limits and introduce the concept of the decoding areas. Second, we introduce a cross-layer protocol that seamlessly navigates the decoding areas with the objective of maximizing throughput while minimizing decoding failures. Experiments on a 4-USRP testbed show that COLiDeR delivers between 80-95% of the relaying performance of an ideal full-duplex radio while incurring negligible decoding failures.

ACM Reference Format:

Raphaël Naves, Gentian Jakllari, Hicham Khalifé, Vania Conan, and André-Luc Beylot. 2019. COLiDeR: A Cross-Layer Protocol for Two-Path Relaying. In *22nd Int'l ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '19), November 25–29, 2019, Miami Beach, FL, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3345768.3355941

1 INTRODUCTION

Two-path relaying was introduced as a solution to increase the throughput in the 4-node scenario depicted in Fig. 1. In this scenario, the source node, S, wants to send a batch of packets to the destination, D, which is too far to receive the packets via direct transmissions. Assuming a TDMA (Time-Division Multiple Access) channel access protocol for ease of presentation, a single-path relaying strategy can not deliver a end-to-end throughput higher than 0.5 frame per slot because of half-duplex hardware limitation. However, with a two-path relaying approach, the source, S, can transmit on every slot, alternating between two relays. As depicted in Fig. 1, in slot k, S transmits a packet to R_1. In slot $k+1$, R_1 forwards the packet to the destination, D, while S transmits a second packet, this time to R_2. If R_2 can successfully decode the packet from S, despite the interference from R_1, it will forward it to D in the following slot, while S will transmit a third packet to R_1. Thus, as long as one relay can decode a packet from the source while the other is transmitting a packet to the destination, two-path relaying can eliminate the multiplexing loss.

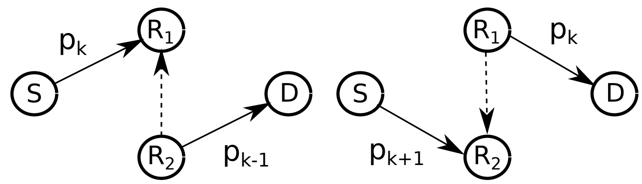

(a) Slot k (b) Slot k+1

Figure 1: Two-path relaying illustration

While two-path relaying has been studied theoretically [1, 2], realizing it in practice faces several challenges. First and foremost, two-path relaying requires a relay to successfully decode a transmission from the source when the other relay is transmitting. To this end, many theoretical works [1–3] rely on the well-known concept of Successive Interference Cancellation (SIC). Indeed, it is assumed that the inter-relay channel is very strong, allowing the interference to be decoded first and then eliminated from the received signal. However, in a practical system, there is no guarantee as to the inter-relay channel, especially if the nodes are mobile. In parallel, other works employ sophisticated amplify-and-forward coding techniques [4] while recent studies [5] propose to simply alternate between two-path relaying and a traditional one-path relaying depending on the packet decoding success rate at relays. The question, however, of how the stations would know which particular relaying approach to use was left open. More generally, to the best of our knowledge, there is no thorough comparative study, particularly using real-hardware implementations, on what is the best approach to be adopted by the relays.

In this work, we present COLiDeR, a novel PHY/MAC CrOss-Layer protocol for practical decode-and-forward Diamond[1] Relaying. Based on a testbed-driven evaluation and analysis of the best approach for a relay to decode a packet from the source node while the other relay is transmitting a packet to the destination, COLiDeR is capable of identifying the relaying strategy that maximizes throughput while minimizing packet losses due to decoding failures.

[1]Due to the geometrical shape of the Fig. 1 topology, two-path relaying is also known as a diamond relay network.

MSWiM '19, November 25–29, 2019, Miami Beach, FL, USA

ACM ISBN 978-1-4503-6904-6/19/11…$15.00
https://doi.org/10.1145/3345768.3355941

Figure 2: USRP testbed used for implementing and evaluating diamond-based relaying.

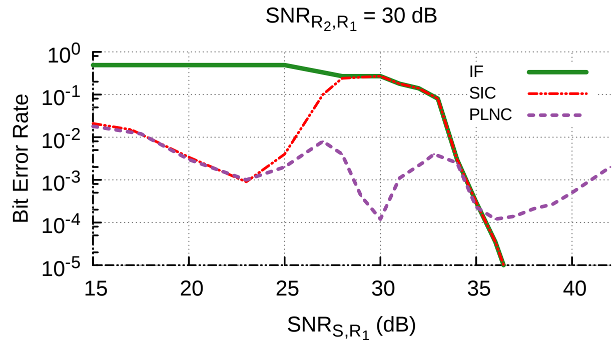

Figure 3: Error Rate at R_1 with different values of SNR_{S,R_1}

We implement COLiDeR on a 4-node USRP radio testbed with the GNU Radio framework [6], and evaluate its performance using over-the-air transmissions exclusively. The experimental results show that COLiDeR delivers between 80-95% of the relaying performance of an ideal full-duplex radio while incurring negligible decoding failures.

2 SIMULTANEOUS RECEPTION OF MULTIPLE PACKETS: A MEASUREMENT-DRIVEN ANALYSIS

Since the feasibility of two-path relaying closely depends on the ability of relays to decode the packet sent by the source while the other relay is also transmitting, we carry out in this section an experimental study to identify the best strategy for the relays to handle the simultaneous reception of two signals.

2.1 Problem formulation

To simplify the presentation, we consider one part of two-path relaying as depicted in Fig. 1(a). S transmits a packet, p_S, to R_1 while R_2 transmits a packet, p_{R_2}, to the destination, with R_1 facing the challenge of decoding p_S. All the results shown apply to the second part, taking place in the subsequent slot, as depicted in Fig. 1(b).

We consider that packets are sent using an OFDM coding structure with a QPSK (Quadrature Phase-Shift Keying) symbol modulation on each subcarrier. Denoting with $x_S^{m,k}$ and $x_{R_2}^{m,k}$ the symbol transmitted simultaneously on the m-th OFDM symbol and the k-th subcarrier by nodes S and R_2, respectively, the symbol $y^{m,k}$ received by R_1 can be expressed as follows:

$$y^{m,k} = h_{SR_1}^{m,k} x_S^{m,k} + h_{R_2R_1}^{m,k} x_{R_2}^{m,k} + n \quad (1)$$

where, $h_{XY}^{m,k}$ is the channel coefficient of the m-th OFDM symbol on the k-th subcarrier between X and Y nodes and n is the ambient noise.

Obviously, the symbol $x_S^{m,k}$ is the signal of interest for R_1.

2.2 The decoding candidates

We consider three decoding techniques that could potentially be applied at R_1.

1) Interference-Free (IF) equalization [3][5]: R_1 decodes directly the packet p_S considering the part $h_{R_2R_1}^{m,k} x_{R_2}^{m,k} + n$ of Eq. (1) as Gaussian noise. In particular, R_1 estimates the symbol $\hat{s}_S^{m,k}$ sent by S on the k-th subcarrier of the m-th OFDM symbol using the minimum distance criteria.

2) Successive Interference Cancellation (SIC) [3][5][2]: This decoding technique is used to decode many packets received simultaneously with different power levels. The principle is to decode the strongest one with an IF equalization and then to cancel its contribution in equation (1), repeating the operation until the signal of interest is decoded [7].

When $SNR_{R_2,R_1} > SNR_{S,R_1}$[2], SIC equalization becomes a two-step decoding process. R_1 estimates the symbol $\hat{s}_{R_2}^{m,k}$ sent by R_2, removes its contribution from the the received signal and finally it estimates $\hat{s}_S^{m,k}$ with second IF equalization.

3) Physical-Layer Network Coding (PLNC): Never before associated with two-path relaying, we make it the third candidate because it is designed to benefit from two interfering signals. In its simplest form, PLNC [8] consists of decoding a linear combination (typically a XOR combination) of two interfering packets rather than the two individual packets. In the considered scenario for instance, R_1 aims to decode a linear combination of p_S and p_{R_2} rather than p_S only. To do so, we apply the optimal decoding technique in the sense of the minimum distance criteria introduced in [9]. Note that, with PLNC, R_1 decodes a linear combination of p_S and p_{R_2} rather than the individual packet p_S. The destination, using previously received packets (the very first is a single native packet) can perform the PLNC decoding.

2.3 IF vs. SIC vs. PLNC

Experimental platform: To compare the performance of the three decoding techniques, we use the testbed depicted in Fig. 2. It consists of 4 Ettus Research Universal Software Radio Peripherals (USRPs) N210 equipped with an SBX daughterboard and GPS antennas. Communications take place over the 1.8 GHz frequency band with a sampling rate of 400kB/s and the signal/MAC processing is done

[2] $SNR_{X,Y}$ denotes the Signal-To-Noise Ratio of a signal sent by node X measured at node Y.

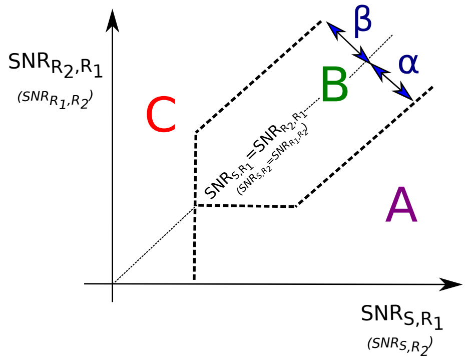

Figure 4: Three decoding strategies at R_1 depending on the intra-diamond channel (IDC)

in C++ and Python within the GnuRadio framework [6]. A slotted TDMA channel access is built thanks to the GPS synchronization.

Implementation details: Packets are 1280 bit-long and each OFDM symbol includes 64 subcarriers: 40 subcarriers for data symbols, 12 guard subcarriers and 12 pilot subcarriers. To estimate the channel coefficients $\hat{h}^{m,k}_{R_2R_1}$ and $\hat{h}^{m,k}_{SR_1}$ each interfering payload transmission is preceded by two interference-free preambles sent by S and R_2 successively [10].

Experiment description: To evaluate R_1's capability to decode the signal from the source under diverse intra-diamond channel conditions, we keep the inter-relay link stable and vary the source transmission power. Specifically, we start by modulating the transmission power of R_2 so as to obtain $SNR_{R_2,R_1} = 30\ dB$. Leaving R_2's transmission power unchanged for the rest of the experiment, we vary the transmission power of node S and evaluate the Bit Error Rate (BER) at R_1. When PLNC decoding is applied, the BER is obtained by comparing the decoded packet to $p_S \oplus p_{R_2}$ or $p_S \oplus Rot(p_{R_2})$, depending on the θ value obtained at R_1.

Results: Figure 3 delivers two main lessons. First and probably most surprisingly, contrary to what is recommended by almost all literature on the subject [2, 3, 5], SIC is a poor enabler of two-path relaying. Instead, the data shows that on off-the-shelf hardware, using PLNC yields an acceptable BER ($< 10^{-2}$)) for far more intra-diamond channel values.

The second lesson is that depending on the intra-diamond channel state, there are three distinct areas relevant to two-path relaying. When $SNR_{S,R_1} > SNR_{R_2,R_1}$ by more than 5dB (right part of the curve), R_1 can easily decode p_s by using IF, as expected. Second, when the difference in SNR between the two signals is between ($5dB$, $-15dB$) – a wide region – using PLNC enables R_1 to decode with a reasonable BER. Finally, when the desired signal becomes too weak ($< 15dB$), the BER starts reaching unacceptable levels.

Then, from a higher perspective, the empirical study showed that the performance of two-path relaying will highly depend on the intra-diamond channel (IDC). Thus, we identify three IDC-level decoding areas, depicted in Fig. 4, which could lead, as shown in § 3, to different decoding strategies at relays:

- **Area A** – the desired signal is stronger than the interfering signal. Using interference-free equalization (IF), the relay, say R_1 (resp. R_2), can decode the packet from the source while the other relay, R_2 (resp. R_1), is transmitting a packet to the destination. Therefore, the source and R_2 (resp. R_1) can transmit simultaneously.
- **Area B** – the desired signal is similar in strength to the interfering signal. Using PLNC-based decoding, the relay, say R_1 (resp. R_2), can decode the xor-ed packet containing the packet from the source while the other relay, R_2 (resp. R_1), is transmitting a packet to the destination. Therefore, the source and R_2 (resp. R_1) can transmit simultaneously.
- **Area C** – the desired signal is significantly weaker than the interfering signal. In this case, there is no known approach for one of the relays, say R_1 (resp. R_2), to decode the desired signal coming from the source, if the other relay, R_2 (resp. R_1) transmits to the destination at the same time. These transmissions have to be orthogonalized.

3 COLIDER

In this section, we introduce COLiDeR, a CrOss-Layer Diamond Relaying protocol, whose design is driven by the measurement study presented in § 2. The goal is to achieve two complementary objectives: maximizing the received throughput at the destination while minimizing the number of packet losses due to decoding failures.

3.1 IDC-aware states and scheduling

As our measurement study in Sec. 2 demonstrated, the performance of two-path relaying depends on the intra-diamond channel (IDC). To successfully navigate the state of the diamond, COLiDeR identifies 5 states: for three of them scheduling is trivial while for the other two a new scheduling approach is introduced. Each state is defined by the the ability of R_1 and R_2 to handle interfering signals.

3.1.1 ***States 0 and 0'.*** None of the relays can decode a packet transmitted from the source, S, while the other relay is transmitting – area C in Fig. 4 for the two relays. In this case, COLiDeR relays packets through R_1 only (state 0) or R_2 only (state 0'), following the traditional one-path relaying approach. As a result, the source can send at most 1 packet every 2 time slots, for a 0.5 frames/slot maximum throughput.

3.1.2 ***States 1 and 1'.*** R_1 can handle the reception of two signals while R_2 cannot (state 1) or R_2 can handle the reception of two signals while R_1 cannot (state 1') – areas A or B in Fig. 4 for one of the relays and area C for the other. In this situation, COLiDeR introduces the concept of 2/3-capacity two-path relaying. To illustrate it, let us consider the case in which R_1 can handle the reception of two signals while R_2 cannot (State 1), as depicted in Fig. 5(b). In slot 1, the source, S, transmits an interference-free packet to R_2. In slot 2, R_2 relays the packet to the destination while at the same time S transmits a second packet to R_1. Finally, in slot 3, R_1 transmits the second packet, interference-free, to the destination, which ends up receiving two packets over 3 time slots (0.66 packets/slot) – a 2/3-capacity relaying.

3.1.3 ***State 2.*** Both relays can decode packets transmitted by the source while the other is transmitting – areas A or B in Fig. 4 for the two relays. In this case, COLiDeR switches to full-capacity two-path relaying, as depicted in Fig. 5(c). The source transmits 1 packet

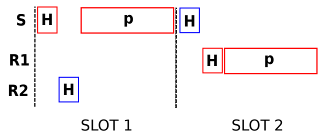

(a) State 0 (half-capacity relaying)

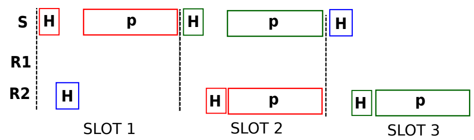

(b) State 1 (2/3-capacity relaying)

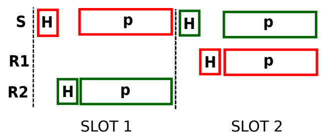

(c) State 2 (full-capacity relaying)

Figure 5: COLiDeR introduces 3 scheduling policies: the traditional half-capacity one-path relaying (5(a)); 2/3-capacity two-path relaying (5(b)); and full-capacity two-path relaying (5(c)).

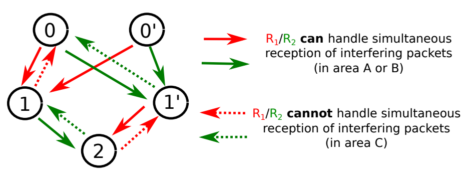

Figure 6: Source rate adaptive protocol state machine

per slot while either R_1 or R_2 relays a packet to the destination, which receives 1 packet/slot.

3.2 Distributed state transition

In this section, we describe how COLiDeR addresses the challenge of identifying the proper two-path relaying state out of the 5 possible introduced in § 3.1. What makes this task challenging is that the state depends on the intra-diamond state. Not only this information is not known in advance by any of the stations, but it can be subject to change due to the dynamic nature of the wireless channel.

At the high level, COLiDeR's approach is to have the source drive the transitions, with input from relays which are in the best position to estimate the best decoding strategy. The information necessary for the coordination is shared via the transmission of orthogonal control packets (H packets in Fig. 5). We use the finite-state machine, depicted in Fig. 6, to model the selection of the two-path relaying state. The initial state, as well as the conditions for each transition, are as follows:

Initial state: To bootstrap, COLiDeR uses the traditional single-relay, half-capacity forwarding of packets. Either R_1 (state 0) or R_2 (state 0') exclusively forwards the packets from the source to the destination.

Transitioning to a higher state: The source and two relays periodically transmit control packets, as depicted in Fig. 5. Leveraging the reception of these packets, R_1 (resp. R_2), can estimate SNR_{S,R_1} and SNR_{R_2,R_1} (resp. SNR_{S,R_2} and SNR_{R_1,R_2}), necessary for deciding in which decoding area of Fig. 4 it most likely is. If R_1 estimates that decoding two signals is possible, i.e. area A or B, it notifies the source in the next control packet it transmits. The source, based on the the control packets it receives, decides whether to transition to a higher state. To avoid overreacting to what could be a momentary change, the source only moves to a higher state if it receives a given number of consecutive packets notifying that a relay is in area A or B. Before transitioning, the source notifies the relays to immediately update their adopted state. This information is included within source's control packets.

Transitioning to a lower state: The source will decide to shift to a lower state if one of the two events occur: a) it receives a given number of control packets from a relay notifying that it is in area C (the relay can not receive two simultaneous packets), or b) it receives no control packets that pass the CRC check for a given number of cycles. Should that happen, the source executes a state change by modifying the adopted state field in the control packets it sends out, similarly to the transition to higher state.

To account for the fact that packet losses are inevitable, the relays will transition by themselves to the initial state if they receive no control packets from the source passing the CRC check for a given number of cycles.

4 GENERAL EVALUATION

To evaluate the reactivity and the overall performance of COLiDeR in different channel conditions, we implement a scenario incorporating four representative configurations. Initially, configuration I, the 4 USRP radios depicted in Fig. 2 are placed in the classic diamond topology and the transmission power levels are selected such that only one relay can handle two simultaneous transmissions with a reasonable bit error rate. At time $t_0 = 100$ slots, relay R_2 is switched off, creating configuration t_0 aimed at emulating a relay loss due to mobility. At $t_1 = 200$ slots, relay R_2 is switched back on, emulating the selection of a new relay. In the final configuration, at $t_2 = 300$ slots, the source transmission power is increased, emulating a scenario in which the source moves closer to the relays.

Comparison: Our goal is to evaluate how COLiDeR navigates the trade-off between reliability and throughput, its principle design consideration. Note that COLiDeR could optimize reliability by permanently staying in state 0, the equivalent of traditional one-path relaying. Thus, this approach is included in the study. On the opposite side, having COLiDeR permanently in state 2 may maximize throughput although it causes severe packet losses. To see why this is true, consider the case in which R_1 and R_2 can not handle the reception of two interfering signals. While R_1 is transmitting a packet to the destination, the source has two choices. It can transmit to R_2, the maximalist approach, leading to a packet loss. Or it does not transmit, which is what COLiDeR would have done since it would have transitioned to state 0. In terms of throughput, assuming layer two will retransmit the same packet in the following slot, the maximalist approach did no worse than COLiDeR even though it transmitted unnecessarily, wasting energy and augmenting interference. COLiDeR, however, may sometimes do worse than the maximalist approach in terms of throughput, either because it is designed to wait for a few rounds before it transitions to the next best state, or because it chooses to be in a lower state to avoid packet losses. Thus, we also compare COLiDeR to this maximalist approach, when used with a static PLNC decoding at relays.

Results: Figure 7 plots the end-to-end throughput (Fig. 7(a)) and percentage of lost packets (Fig. 7(b)) against time. Each value is obtained by computing the throughput and loss ratio over the 30 previous slots.

In the initial configuration, I, R_1 can handle the reception of two signals with a reasonable bit error rate whereas R_2 cannot. COLiDeR adopts scheduling state 1 (§ 3.1), resulting in 0.66 packets/slot throughput and no packet losses – the best joint throughput-reliability performance. When the source adopts the maximalist approach and transmits a packet every slot, a slightly higher (around 8%) throughput is achieved, because occasionally, R_2 successfully decodes the interfering packets. It comes, however, with a high penalty in terms of packet losses, around 25%. Furthermore, COLiDeR performs just as well as one-path relaying in terms of packet losses while outperforming it by 25% in terms of throughput.

At $t_0 = 100$ slots R_2 is switched off, leaving a single path to the destination. Obviously, as Fig. 7(a) shows, the maximum throughput in this case, realized by all schemes, is 0.5 packet/slot. Having the source sending 1 packet/slot when only R_1 is active, however, leads to an inevitable 50% rate loss. As described in § 3.2, COLiDeR, after receiving no control packets from R_2, transitions from state 1 to state 0, limiting the packet losses.

At $t_1 = 200$ slots, relay R_2 is switched back on and is eventually integrated back to two-path relaying, enabling the two-path relaying schemes to realize the performance observed in the period 0-100 slots. Finally, at $t_2 = 300$ slots the source transmission power is increased, enabling R_2 to handle the reception of two simultaneous signals with a reasonable bit error rate. COLiDeR takes advantage of the newly created opportunity and transitions to state 2, achieving close to channel-capacity performance in terms of throughput with almost zero packet losses. The maximalist approach, while also transmitting 1 packet/slot, only includes one decoding technique (PLNC), explaining its inferior performance when compared to COLiDeR.

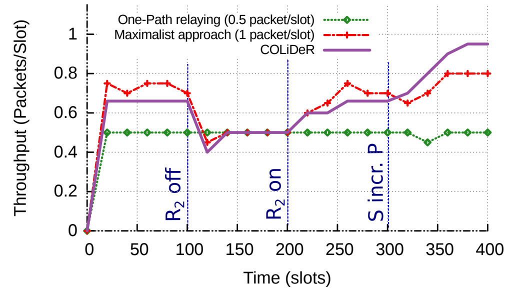

(a) Received Throughput

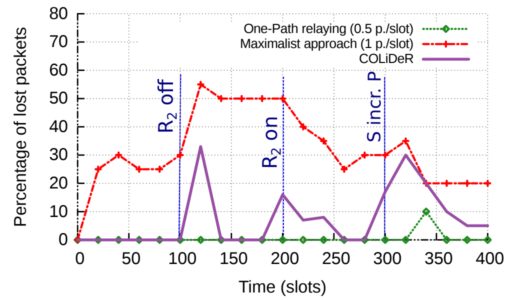

(b) Lost packets

Figure 7: COLiDeR performance in a general evaluation scenario

5 CONCLUSION AND FUTURE WORK

We have presented COLiDeR a PHY/MAC cross-layer solution for adaptive two-path relaying. Experimentally driven, our approach combines interference-free equalization and Physical-Layer Network Coding (PLNC) for decoding at relays, and relies on a source rate adaptation mechanism to respond to channel state and topology changes. Experimental results have shown that COLiDeR ensures smooth and efficient sailing between the defined protocol states, achieving between 80-95% of the relaying performance of the ideal full-duplex radio while incurring negligible decoding failures.

While this work was focused on minimizing the relaying multiplexing loss, as future work we aim to integrate and evaluate COLiDeR in large-scale wireless networks. Therefore, two complementary questions are of major interest. First, how to identify diamonds in a network with multiple flows? Then, how should we adapt the TDMA framework to fit most of current ad-hoc networks relying on CSMA-based channel reservation methods?

ACKNOWLEDGMENTS

This work was supported in part by the Agence Nationale de la Recherche under the ANR JCJC CiTADEL grant.

REFERENCES

[1] C. Zhai, W. Zhang, and P. C. Ching, “Cooperative spectrum sharing based on two-path successive relaying,” *IEEE Transactions on Communications, Vol. 61, No. 6*, 2013.
[2] R. Simoni, V. Jamali, N. Zlatanov, R. Schober, L. Pierucci, and R. Fantacci, “Buffer-aided diamond relay network with block fading and inter-relay interference,” *IEEE Transactions on Wireless Communications, Vol. 15, No. 11*, 2016.
[3] Y. Fan, C. Wang, J. Thompson, and H. V. Poor, “Recovering multiplexing loss through successive relaying using repetition coding,” *IEEE Transactions on Wireless Communications, Vol. 6, No. 12*, 2007.
[4] H. Wicaksana, S. H. Ting, C. K. Ho, W. H. Chin, and Y. L. Guan, “Af two-path half duplex relaying with inter-relay self interference cancellation: diversity analysis and its improvement,” *IEEE Transactions on Wireless Communications*, vol. 8, no. 9, pp. 4720–4729, September 2009.
[5] F. Tian, W. Zhang, W.-K. Ma, P. Ching, and H. V. Poor, “An effective distributed space-time code for two-path successive relay network,” *IEEE Transactions on Communications, Vol. 59, No. 8*, 2016.
[6] “Gnuradio overview available at http://www.gnuradio.org/.”
[7] D. Halperin, T. Anderson, and D. Wetherall, “Taking the sting out of carrier sense: Interference cancellation for wireless lans,” in *Proc. IEEE MOBICOM Conference*, 2008.
[8] S. Zhang, S. Liew, and P. Lam, “Hot topic: physical-layer network coding,” in *Proc. ACM MobiCom Conference*, 2006.
[9] V. N. V. T. Muralidharan and B. Rajan, “Wireless network-coded bidirectional relaying using latin squares for m-psk modulation,” *IEEE Transactions On Information Theory, Vol. 59, No. 10*, 2013.
[10] L. You, S. Liew, and L. Lu, “Network-coded multiple access ii: Toward real-time operation with improved performance,” *IEEE Journal On Selected Areas in Communications, Vol. 33, No. 2*, 2015.