A SURVEY OF IMPLEMENTATION EFFORTS AND EXPERIMENTAL DESIGN FOR COOPERATIVE COMMUNICATIONS

Glenn J. Bradford and J. Nicholas Laneman

Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556
{gbradfor, jnl}@nd.edu

ABSTRACT

Design and analysis of cooperative communication schemes based upon modeling and simulation exist in large quantities in the research literature. Despite this fact, there have been relatively few efforts directed toward implementing and experimentally evaluating such schemes. Cooperative protocols have many components that make them challenging to implement in real-world radio architectures, and their expected gains are highly dependent on the network topology and RF environment in which they operate. As such, experimental work will be crucial in the transition of such schemes from conceptual proposals to next-generation wireless standards. This paper motivates such practical work, surveys existing efforts in the area, and offers future direction for architectural and experimental design.

Index Terms— cooperative diversity, experimentation, implementation

1. INTRODUCTION

Since cooperation was first proposed for wireless networks, [1, 2], it has become a popular area of communications and networking research. The performance gains for cooperation stem from mitigating long-term path-loss and shadowing effects and short-term multipath fading effects. The combined path-loss of a two-hop transmission may be less than direct transmission, and RF absorbing objects in the environment can be circumvented to reduce shadowing. This is analogous to multi-hop routing and can be performed at either the network (NET) or medium access (MAC) layer, with greater efficiency at the latter [3]. Relaying can also provide a receiver with redundant messages that travel through spatially distinct paths, reducing the impact of multipath. This creation of spatial diversity, known as cooperative diversity, provides better average performance and can be achieved at either the MAC or physical (PHY) layer, with greater performance gains at the latter. The end result of path-loss and diversity gains provided through cooperation is improved network performance, realizable as various combinations of power savings, data rate increases, network coverage expansion, and interference reduction.

A number of characteristics of cooperative schemes make experimentation essential for transitioning proposed schemes to practical applications. Information flowing through a cooperative link visits multiple nodes within the network, suggesting cooperation will have implications for multiple layers of the protocol stack, including at a minimum the PHY, MAC, and NET layers. Decisions made at one layer will impact the efficiency of the others, increasing protocol complexity but providing many opportunities for cross-layer optimization. Performance metrics to evaluate such cross-layer schemes, such as throughput and delay at the MAC and NET layers, require large networks to obtain realistic estimates. Modeling and simulation can become unwieldy as the number of nodes within a network grows, making experimentation "one of the de facto approaches for benchmarking" [3]. Complex interactions between layers will also be impossible to completely predict, making full implementation necessary to identify conflicts that may degrade performance.

Cooperative gains determined by theoretical work and simulation vary significantly depending upon the models chosen for the network topology and the wireless channel. Real-world propagation environments are notoriously complicated to model; the most faithful models quickly become intractable for both theoretical work and simulation. Add to this the multi-layer simulation required by the cross-layer nature of cooperation, and experimentation becomes the only practical way to evaluate full protocols. Experimentation is an opportunity to obtain realistic quantifications of performance gains for proposed cooperative schemes in actual propagation environments and network topologies.

Finally, modeling and simulation often make idealistic assumptions about the capabilities of the physical radio platforms on which proposed protocols are to be run. Common limitations in real-world radios such as a lack of frequency and timing synchronization, imperfect channel estimation, and quantization errors have a large impact on the effectiveness of many cooperative schemes. Although a substantial amount of the literature addresses these issues directly, experimentation is the only way to identify other impediments and confirm devised solutions are indeed solutions.

This work has been funded in part by NSF Grant CNS06-26595

2. EXPERIMENT CLASSIFICATION

Current implementation efforts can be separated into two broad groups based on the architecture they employ:

- **Legacy-based** implementations use existing platforms and have the dual benefits of promoting quicker adoption and facilitating direct comparison with existing schemes. Their major drawback is that necessary modifications to allow full cooperation are often not possible.

- **Clean-slate** systems allow full cooperation but need to be designed from scratch. Because this can be an expensive and time consuming project, reconfigurable radio architectures, typically based upon Software-Defined Radio (SDR), have become popular for clean-slate implementational and experimental research.

These contrasting architectural options raise the question of how best to implement cooperation. PHY layer cooperation produces the greatest performance improvements but also requires the most sophisticated radio architecture. Generally, as one moves up the network stack, cooperation becomes easier to implement but with a decrease in the performance gains. Which relaying mechanisms are realizable in a system is also highly dependent on the functionality of the underlying radio architecture.

Experimentation with cooperative protocols is very much in its infancy, with most completed work only verifying simple results from theoretical work. Several features are useful in classifying experimental work:

- **Network Size** has an impact on how realistic the performance measures made during experimentation are for real-world settings. Larger networks are more useful, especially for the metrics of higher protocol layers.

- **Network Topology** is important as cooperation will have a greater impact in some topologies than others. Cooperative schemes need to be examined in a number of topologies to ensure both situational and average performance improvement.

- **Propagation Environment** also impacts how general the measures of performance are, requiring the collection of data for multiple RF environments.

- **Induced Fading** provides multiple channel realizations to accurately characterize diversity gain for a given topology. It can be obtained through node mobility or environmental alteration, but care must be taken to isolate spatial diversity gains from temporal.

With a few exceptions, experimental setups have employed simple networks with relatively few nodes and limited induced fading in a relatively small subset of RF propagation environments. The next section describes these experimental efforts with the various projects summarized in Table 1.

3. EXISTING COOPERATIVE EFFORTS

An FPGA-based SDR platform that has been used in a number of cooperative implementations is the Wireless Open Access Research (WARP) board [4] developed at Rice University. Its use of a high-end FPGA with embedded processors makes it possible to design sophisticated PHY, MAC, and NET layers that approach performance levels of commercially available systems. Cooperative experimentation at Rice [5] was done with a distributed Alamouti space-time code (DSTC) using OFDM modulation and amplify-and-forward (AF) processing at the relay. Experiments with the setup were done using three nodes with the relay halfway between the source and destination. The experiment was conducted indoors with line-of-sight propagation in the 2.4 GHz ISM band with no induced fading. Results show bit-error rates for cooperation are superior to those of direct transmission, but without node movement the diversity gain is not quantifiable.

A number of cooperative efforts have been undertaken at Polytechnic Institute ([3] and references therein) including a legacy-based effort which modified Linux WiFi card drivers to implement a proposed cooperative MAC. The scheme, called CoopMAC, enables multi-hop transmission at the link layer. Nodes within the network keep a table of the sustainable transmission rates between their neighbors; if a node determines transmission to a destination will be faster through an intermediate relay, it will send its packet via this two-hop route. Experiments with the testbed were done using 10 to 20 laptop nodes running the modified WiFi drivers in the 2.4 GHz band. Throughput, delay, and jitter were measured for multiple network configurations. Multiple topologies and a reasonable network size make the measured performance a good estimate of the scheme's capabilities. By adapting a legacy system, this approach allows for a full MAC implementation, easy large scale network deployment, and direct comparison with existing protocols. The major disadvantage is all PHY and time critical MAC functions are implemented in proprietary firmware, making PHY cooperation for spatial diversity impossible and creating inefficiencies at the MAC.

To eliminate the MAC layer inefficiencies, a full version of CoopMAC was implemented using WARP operating at 2.4 GHz. The ability to modify time critical MAC functions made it possible to eliminate inefficiencies revolving around control and acknowledgement packets. Experiments with this platform only used three nodes indoors in a line topology. Transmission rates on all three links were manually set irrespective of the actual channel quality. Throughput, average delay, and packet error rate were measured for standard WiFi, CoopMAC, and an implementation mimicking the legacy system CoopMAC. The full implementation outperformed all others, but the small network size makes it difficult to extrapolate performance. Although this work uses a clean-slate architecture, its focus is still legacy systems since PHY cooperation is not used.

Table 1. Summary of Experimental Cooperative Projects

Project	Architecture	Layers	Nodes	Topology	Band	Data Rate	Fading
Rice	clean-slate (FPGA)	PHY	3	line	2.4 GHz	15 Mbps	no
CoopMac WiFi	legacy	MAC	10-20	random	2.4 GHz	11 Mbps	no
CoopMac WARP	clean-slate (FPGA)	MAC	3	line	2.4 GHz	24 Mbps	no
Polytechnic PHY	clean-slate (FPGA)	PHY	3	line	2.4 GHz	10 Mbps	no
KTH	clean-slate (DSP)	PHY	4	random	1.77 GHz	19.2 kbps	no
ETH	RF front-end	N/A	10	random	5.25 GHz	N/A	yes
Notre Dame	clean-slate (GPP)	PHY	3	triangle	400 MHz	50 kbps	yes

Another effort at Polytechnic incorporates PHY cooperation at 2.4 GHz, using WARP boards to capture waveforms which are then sent to a computer for batch processing. Source and relay transmissions are orthogonal with the relay performing decode-and-forward (DF). Two schemes were implemented, including maximum ratio combining (MRC) and a convolutional code scheme with parity bits generated by the relay. The lack of a real-time PHY makes large scale experiments and cross-layer work impractical, but work is proceeding to move the PHY into the FPGA. Three nodes were used, each positioned by trial and error to give link qualities favorable to cooperation. The parity check code with soft decisions outperformed hard decisions, followed by DF and direct transmission. Although this setup shows topologies exist for which PHY cooperation is beneficial, biasing the network link performances in favor of cooperation and not inducing fading means the diversity gain of the system could not be quantified.

A DSP-based testbed for cooperative communications at the Royal Institute of Technology (KTH) and its experimental results are described in [6]. The testbed consists of four nodes, making it possible to have two relays, which operate at low data rates with RF boards for the 1.77 GHz band. All baseband processing is done on a TI DSP board. A number of cooperative protocols were tested, including AF, DF, cooperative MRC (CMRC), a DSTC, and selection based on the strongest relay. Measurements for the testbed were taken indoors using the four nodes and multiple network topologies but with no induced fading for a given topology. Performance was analyzed using packet error rate and a novel metric called implementation loss. Complexity, in terms of the number of DSP clock cycles, and overhead, in terms of the number of symbols not used to convey data, are reported for each scheme. Their inclusion is a useful detail not found in other works.

Radio Access with Cooperating Nodes (RACooN) Lab [7] is a testbed at ETH-Zurich comprised of ten single antenna, half-duplex radio nodes. The radios operate in the 5.25 GHz band with 60 MHz of user bandwidth and are battery powered for portability. Nodes consist of an RF unit and a baseband processing and data storage unit. Like the batch processing WARP setup, transmit waveforms must be predefined and received waveforms must be processed off-line. These nodes were used to obtain real-world channel measurements for cooperative simulations focusing on joint cooperative diversity and scheduling. The setup consisted of two sources, four relays, and two destinations. Measurements were taken in an indoor lab with the source and destination nodes being moved during data collection to obtain multiple channel realizations. Because the testbed was only used to obtain channel transfer functions, implementation issues were not addressed.

Last but not least, an experimental effort at the University of Notre Dame based on an SDR performing baseband processing on a general purpose processor (GPP) is reported in [8]. Nodes use a combination of the Universal Software Radio Peripheral (USRP) [9] as an RF front-end and a host computer running GNU Radio [10] for baseband processing. There is considerable delay between packet reception and processing, making turnaround times quite large for reaction to channel conditions and the generation of acknowledgements. These delays make realistic MAC layer implementations difficult, but performing baseband processing in high-level software makes implementing PHY protocols much easier. Experimentation in the 400 MHz band used a three node network with source and relay transmissions orthogonal in time and the relay performing DF. Data was collected with the nodes arranged equidistant from one another, normalizing average link performance and putting the relay link at a disadvantage. Additionally, the nodes were fixed to a moveable frame so that fading could be induced to quantify the diversity gain via bit-error curves. Results show a diversity gain for DF, but only when the relay selectively forwards packets received without error, determined through the use of a cyclic redundancy check. The lack of a MAC implementation means cross-layer issues were not explored.

4. DISCUSSION

By far, the most extensive experimentation in terms of number of nodes and topologies used has been done with the legacy implementation of CoopMac. It has displayed the benefits of cooperation at the MAC layer and some approaches needed for incorporating cooperation into existing standards. The remainder of the experimentation outlined above has focused on PHY cooperation in simple setups with no more than three to four nodes, limited induced fading, and little MAC integration. This is primarily because much of the focus has been on developing frameworks and testbeds with which to experiment. They have been useful in confirming a number of important basic theoretical predictions and demonstrating that currently available hardware and software solutions are sufficient for developing cooperative networks. With the development of the basic infrastructure in terms of network nodes, more elaborate implementations and experiments are expected in the near future.

Future implementation work needs to encompass a wide range of characteristics lacking in the current work. First, although the suitability of current radio architectures for basic cooperation has been demonstrated, more complicated cooperative schemes, such as distributed beamforming, have stricter hardware requirements; experimentation that focuses solely on the PHY should concentrate on determining appropriate radio architectures to enable these advanced techniques. Second, one notices from Table 1 that the outlined projects have all dealt with one isolated protocol layer, either the PHY or MAC, without exploring the plethora of complex interactions that arise if cross-layer optimization is attempted. Many cross-layer topics of great interest in the literature, such as network resource allocation or joint partner selection and routing, have not yet been studied from an experimental viewpoint. Addressing these areas will also require that larger networks be constructed to obtain realistic measures of higher layer performance.

Additionally, experiments should be done in a more diverse set of network topologies and propagation environments with the inclusion of induced fading to gauge cooperative performance under a representative sample of real-world conditions. Most of the current efforts only quantify benefits resulting from path-loss gains as opposed to the equally important gains that come from the creation of spatial diversity. Experimental design to obtain meaningful insight into cooperation is no small task, and will become increasingly difficult as larger networks are used. The ultimate goal of experimentation is to ensure that cooperative techniques are adopted into industry standards in an intelligent manner that maximizes performance. Wireless technologies continue to be used in an ever growing number of disparate applications, so experimentation also needs to be carried out for various network architectures. Along with the ad hoc networks described, infrastructure oriented experiments should be created to determine which schemes are appropriate for next generation cellular standards under development, such as LTE Advanced.

Continuing implementation efforts at Notre Dame are working to transform the relatively simple cooperative architecture developed in [8] to cover a number of the areas suggested for improvement in this section. This entails the implementation of a wider range of relaying strategies and the modification of the underlying SDR system to allow for more realist MAC functionality. Although MAC performance on par with commercial broadband standards is not expected, it will be possible to achieve levels sufficient to extrapolate performance to the typical ad hoc network. Once these improvements are complete, a number of experiments investigating partner selection, relaying strategy selection, and cross-layer interaction are planned. These experiments will have roughly 10-20 nodes and will be performed in numerous indoor and outdoor environments. The core of this development plan is only one of many possible, and it is hoped that others will continue their efforts to investigate cooperative communications experimentally. Timing is somewhat critical, as a number of advanced wireless standards are quickly approaching a stage of finalization.

5. REFERENCES

[1] A. Sendonaris, E. Erkip, and B. Aazhang, "User cooperation diversity – Part I: System description," *IEEE Trans. Comm.*, Nov. 2003.

[2] J.N. Laneman, D.N.C. Tse, and G.W. Wornell, "Cooperative diversity in wireless networks: Efficient protocols and outage behavior," *IEEE Trans. Info. Theory*, Dec. 2004.

[3] T. Korakis, M. Knox, E. Erkip, and S. Panwar, "Cooperative network implementation using open source platforms," *IEEE Comm. Magazine*, Feb. 2009.

[4] "Wireless Open Access Research Platform (WARP)," `http://warp.rice.edu`.

[5] P. Murphy, A. Sabharwal, and B. Aazhang, "On building a cooperative communication system: Testbed implementation and first results," *EURASIP Journal on Wireless Comm. and Networking*, 2009.

[6] P. Zetterberg, C. Mavrokefalidis, A. Lalos, and E. Matigakis, "Experimental investigation of cooperative schemes on a real-time DSP-based testbed," *EURASIP Journal on Wireless Comm. and Networking*, 2009.

[7] I. Hammerstroem, J. Zhao, S. Berger, and A. Wittneben, "Experimental performance evaluation of joint cooperative diversity and scheduling," in *Vehicular Technology Conference*. Sept. 2005.

[8] G. Bradford and J.N. Laneman, "An experimental framework for evaluating cooperative diversity," in *Conf. on Info. Sciences and Systems*, Mar. 2009.

[9] "Ettus Research USRP," `http://www.ettus.com`.

[10] "GNU Radio," `http://gnuradio.org/trac`.