

Implementing Cooperative Diversity Antenna Arrays with Commodity Hardware

Aggelos Bletsas and Andy Lippman, Massachusetts Institute of Technology

ABSTRACT

Cooperation among single-antenna transceivers and formation of distributed antenna arrays has recently attracted considerable interest. Such distributed antenna arrays are envisioned to provide resistance to slow wireless fading and improve performance of point-to-point wireless communication across various dimensions. Despite the plethora of recently proposed theoretical approaches that promise gains due to diversity at the physical layer through cooperation (*cooperative diversity*), there is not much work in the implementation of cooperative antenna arrays with existing wireless transceivers. In this article we summarize the main challenges in implementation of cooperative diversity antenna arrays for realistic wireless networks. We then present the basic building blocks of a cooperative diversity demonstration realized in the lab, utilizing commodity radio hardware. Our work sheds light onto the synergies needed between the physical, link, and routing layers that significantly simplify the overall network operation and decrease the transceiver complexity in cooperative diversity antenna arrays, making feasible the utilization of (existing) commodity radio hardware.

INTRODUCTION

A tremendous interest in *cooperative diversity* wireless systems has been observed over the last years [1–4]. In the simplest form, a set of cooperating nodes relay information transmitted from a single source towards a destination (Fig. 1). The main advantage explored in cooperative diversity is the redundancy offered by the availability of several paths between source and destination, through a set of relays; when the direct path between source and destination is in deep fade or blocked by an obstacle, reliable communication might be feasible through the available relays. Given that more than one transmitting nodes are utilized, cooperative diversity schemes are fundamentally different than conventional multihop communication, since they attempt to mimic the behavior of multiantenna links and, as a result, cooperative schemes are usually referred to as “virtual antenna” schemes.

Despite the multiplying numbers of published works on the general theme of cooperative diversity during the last three years, there has not been much work on implementation examples or demonstrations for realistic wireless applications. In this article we first describe the challenges the communication engineer faces during the construction of cooperative diversity schemes and we underline their inherent cross-layer nature. We then present the basic building blocks of the cooperative diversity demonstration we implemented in the laboratory using commodity hardware and explain how we overcame the major challenges.

Our work provides a concrete example of implemented virtual antenna arrays and sheds light onto the interactions needed among the physical, data, and routing layers in cooperative diversity schemes. Hopefully, this work will spark interest within the research community in devising constructive ways that bridge theory with practice in the emerging field of cooperative wireless communications.

IMPLEMENTATION CHALLENGES IN COOPERATIVE DIVERSITY SYSTEMS

ACQUISITION OF NETWORK STATE INFORMATION (NSI) AND IMPLEMENTATION OF NETWORK COORDINATION

The main assumption in cooperative communications is that relay transmissions are always beneficial, compared to direct communication from source to destination. Communication analysis treats the distributed relay transmissions as a set of *perfectly coordinated* links where all relays know *if* and *when* to transmit.¹ Therefore, such analysis provides an upper bound of performance without quantifying the system resources spent to discover which relays are indeed useful, or the required overhead for coordination among the cooperating nodes.

A relay node might have extremely poor wireless channel conditions during a specific time interval, and thus any attempt at relaying through that node would be strictly harmful and wasteful in terms of transmission power and

This work was supported in part by the National Science Foundation under Grant no. CNS-0434816, in part by the Massachusetts Institute of Technology Media Laboratory Digital Life Program, and in part by a Nortel Networks graduate fellowship award.

¹ A discussion about how to transmit is provided in the following subsection.

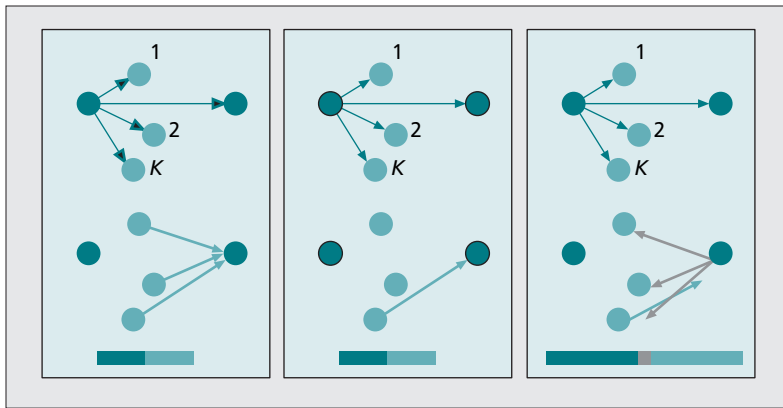


Figure 1. A set of K relays are overhearing the transmission of a single source. After the completion of source transmission, there are several options: a) all relays could utilize a distributed space-time code and simultaneously transmit in-band (left); b) relay selection among the relays can be performed, and a single relay can be utilized (center); c) a relay is utilized if and only if feedback from the final destination flags such necessity (right). The last two cases require distributed methods for efficient relay(s) selection.

available bandwidth. At a different time interval, the same relay might become useful if the wireless channels towards source and destination are strong (e.g., the relay, due to mobility has acquired line-of-sight towards both source and destination). It is therefore imperative to devise cooperation schemes that discover how many relays are available and which of them are indeed useful, depending on their wireless channel conditions, as a function of time. Typically, the number of available nodes in a networked system is discovered at the routing layer (layer 3), while link (channel) state information is acquired at the link layer (layer 2). Therefore, cooperative systems need to exploit interactions between those two layers in a periodic fashion, so as to acquire time-dependent information about the useful relays in the system and their link (channel) states.

We refer to such information as network state information (NSI), addressing the question of *which* relays should participate in retransmissions. Network coordination addresses the issue of *when* relays should retransmit, given that the receiver should detect the transmitted information within a specified time interval. Relays could retransmit *synchronously*, assuming synchronization at the packet level, when all relays retransmit within the same time interval, or assuming synchronization at the network level, when the relays retransmit in a round-robin fashion, one after the other. In the latter case, coordination becomes even harder, given that when one relay transmits, all the others should be aware of that and avoid any type of in-band transmissions. Relays could also retransmit *asynchronously*, for example, one relay could retransmit as soon it has gathered enough signal to properly obtain the original message, independently of the other relays. Even such a scheme requires coordination, given that all relay retransmissions should occur within the specific time interval the receiver “listens”.

Network state information as well as coordination (and their associated overhead) are usual-

ly downplayed in the cooperative diversity literature, even though such issues are critical in realistic environments and thus, deserve special attention. Addressing such issues with distributed algorithms that do *not* rely on any type of global network knowledge (commonly referred to as *genie-aided* or server-based schemes) is a fertile research area that justifies further exploration.

REQUIREMENTS FOR DISTRIBUTED PHASED-ARRAYS (BEAMFORMING) VS. DISTRIBUTED SPACE-TIME CODING

Assuming that useful relays have been identified and coordination in the network has been established, the issue of *how* relays should transmit emerges. Cooperative diversity literature has addressed this problem with two distinct approaches.

The first approach is based on *coherence* and assumes that disconnected wireless transceivers have the ability to synchronize their carriers, in such ways that in-band transmissions from multiple, distributed transmitters *always* add constructively at the final receiver [2]. In that way, in-band transmissions do not cause interference, but instead enhance the transmitted signal. Such techniques are referred to as *beamforming* or *distributed phased-arrays* and they are commonly used to simplify theoretical analysis, since phase alignment at the carrier level reduces the treatment of baseband signals to algebraic sums of real numbers (instead of complex terms with different phases).

In practice, the implementation of distributed phased arrays is still an open area of research. Cost-effective ways to control carrier-level frequencies are required, which is a nontrivial task, especially in the high-frequency regime where modern radios operate. Moreover, some level of signaling from the final destination or a common controller is required, in order for distributed transmitters to estimate the required phase adjustments. In short, the radio hardware complexity and associated cost of distributed phased arrays is significant, and more research is required to keep it reasonable for practical applications. It is also true that there is no commodity radio hardware currently available that utilizes principles of distributed phased-arrays.

The second approach to the problem of cooperative transmission is based on distributed space-time coding techniques [5]. Such techniques do not require in principle any type of channel state information (CSI) or specialized radio hardware at the relays, but instead rely on intelligent coding to exploit the multiple available relay paths. The only “hidden” requirement is the need for linear RF-front ends at the receiver, since efficient space-time codes utilize in-band, simultaneous transmissions. The main challenge in conventional space-time coding research is to achieve good reliability, without sacrificing the achieved throughput rate [6]. The optimal space-time code is only known for two transmitters (the Alamouti scheme) and the research community strives to discover codes that achieve near-optimal trade-offs for an increased number of transmitters (e.g., see the work on lattice coding [7]). There is a long way

before results on space-time coding for the classic Multiple-Input Multiple-Output (MIMO) channel are mature enough to be extended to the distributed case of cooperative diversity, where the number of useful antennas is in general, unknown and time varying. There is also the challenge to quantify the performance of distributed space-time coding in realistic cooperative diversity networks and associate the observed benefits with the required overhead for network coordination (discussed in detail in the previous subsection).

LIMITING TOTAL RECEPTION ENERGY

The *transmission* power (and consecutively, transmission energy) of any cooperative network is usually upper bounded. For example, in FCC part-15 Industrial Scientific Medical (ISM) bands, the total transmission power from multiple in-band *distributed* transmitters is fully specified. Therefore, the total transmission power the relays add into the network is (by definition) bounded. The challenge in cooperative relay schemes is to minimize the overall *reception* power, which is not bounded and depends on the total number of participating nodes in *listening* mode. Energy used for reception critically affects the battery lifetime, and hence requires special attention. This becomes particularly important if modern radios are used, since utilized forward error correction (FEC) has become energy expensive, and thus reception energy is comparable to transmission energy [8]. Even though cooperative diversity has been viewed as an energy-efficient alternative to direct, noncooperative communication, the critical issue of increased total reception energy, as a function of network size, is usually neglected in theoretical analysis and deserves additional attention.

CROSS-LAYER IMPLEMENTATION WITH COMMODITY HARDWARE

In an effort to demonstrate the benefits of cooperative diversity in a concrete way, we decided to implement a cooperative radio network and facilitate a wireless indoor application. The demonstration aimed to confirm the following:

- Improved reliability of cooperative data transfer compared to conventional (noncooperative) wireless communication
- Adaptation of the cooperative wireless network to the time varying wireless channel conditions, especially when people were moving inside the room
- Feasibility of cooperative diversity antenna arrays with commodity hardware

The latter was by far the most difficult task. The above sections discuss the cross-layer nature of cooperative diversity and, therefore, the requirement for flexible radios that provide access to all layers. Unfortunately, most development kits do not provide access to the detection techniques at the physical and link layers, mainly because detection is implemented in silicon and there is no provision to control the associated hardware. Additionally, there is limited access (if any) to the medium access control (MAC) layer which resides between the physical and the rout-

ing layer. Usually that may be programmable and an application programming interface (API) might be available to the communication engineer; however, it is not guaranteed that enough flexibility will be there, given that existing radio protocol stacks have been designed according to noncooperative principles.

Therefore, it was soon realized that we had to design our own transceiver architecture, providing flexible access to all necessary layers (physical, link, and routing). Furthermore, we chose to keep the design as simple as possible and incorporate off-the-shelf, commodity components to decrease the overall cost for the whole network and accelerate hardware development.

Finally, we set up a wireless application through our custom cooperative diversity antenna array system, with special attention to the input/output interface, so as to vividly demonstrate the benefits of cooperation, even to the nonspecialist.

DEMONSTRATION APPLICATION



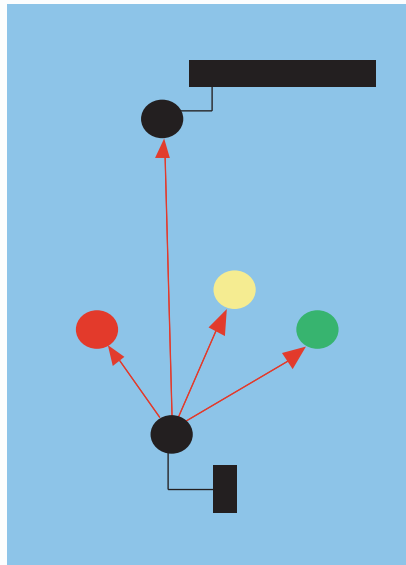
A handheld computer retrieves weather information from the World Wide Web and feeds it through the serial port to one of our custom transceivers. The information is transmitted towards a similar receiver connected to a large store display (Fig. 2). In that way, the received information is publicly displayed and can be readily viewed by humans. For that reason, new information is transmitted every approximately 3 s, which is the time needed for the text to scroll at the store display. Transmission power is decreased and distance between the two endpoints is maximized, so as to emulate poor wireless channel conditions, especially when many people are in the vicinity of the link. Therefore, the displayed text information at the destination includes errors that can be easily perceived.

Between source and destination, there is a set of fixed, immobile color-coded relays (“red,” “yellow,” “green”), packaged with a set of LEDs that show from distance whether the relays collaborate or not with the initial source and form a cooperative diversity antenna array (Fig. 2). Information about which relay is participating is also provided at the destination display via a simple color-coding scheme. Whenever a cooperative antenna array is formed, performance is improved (especially when people move inside the room), and that can be easily perceived from the quality of the text information received and displayed at the destination. More importantly, the network adapts to the time-varying wireless channel conditions and that is visible either through the LEDs at the relays or at the destination store display (Fig. 2). In that way, any spectator of the demonstration has a clear understanding that the network adapts to the dynamics of the wireless channel while cooperation benefits are simultaneously observed. In the following subsections, we describe our approach in detail.

TRANSCEIVER IMPLEMENTATION

We designed and implemented a low-cost, embedded, software defined radio (SDR), as shown in Fig. 3. Our board includes a radio frequency (RF) module, directly interfaced to

There is a long way before results on space-time coding for the classic MIMO channel are mature enough to be extended to the distributed case of cooperative diversity, where the number of useful antennas is general, unknown, and time varying.



Red Relay

The red Relay is blocked. Information is relayed through the yellow relay.



Yellow Relay

The yellow Relay is blocked. Information is relayed through the red relay.

Figure 2. The setup of the laboratory demonstration is depicted. A single source transmits to a single destination. Three colored relays in the room can assist the communication. The source is connected to a handheld computer, and the destination is connected to a large store display, which outputs the received text information. The relays are equipped with LEDs that demonstrate activity. While people are moving inside the room, the network in a distributed and dynamic (non-static) way discovers which relays should be utilized, and relay information is depicted at the destination display.

an 8-bit microcontroller unit (MCU). The technical specifications of the utilized hardware are depicted in Table 1. The MCU is based on the *pushpin* computing architecture [9], originally used for distributed sensing and computation research. For this project, we specifically used the microcontroller to fully control the radio. All necessary functionality, including frame transmission, frame synchronization, frame reception and data detection, cyclic redundancy check (CRC), as well as upper-level link access and routing, are provided in software by the microcontroller and were developed from first principles. Special care was given to fully utilize the available code space without the need of external components and keep the overall hardware design as simple as possible. The microcontroller incorporated a serial port interface, and therefore it was straightforward to interface our transceivers with external handheld computers, or other devices, through the serial port.

PROTOCOL

Approach — The nonlinearity of our radio module front-end precluded the utilization of in-band simultaneous transmissions from multiple relays. Therefore, space-time coding techniques were not an option, with the particular radio design. Nevertheless, the availability of received signal strength indication (RSSI) allowed us to utilize each transceiver and consecutively the entire relay network, as a distributed sensor of the wireless channel and thus, exploit the richness of wireless RF propagation, even with our low-complexity, low-cost radios.

Using the relays as a distributed sensor of RF propagation is the main theme of “opportunistic relaying” [4, 10, 11]: the relays sample the wireless channel in a distributed and periodic fashion and manage to elect the best available single relay path, among a collection of several possible candidates. Specifically, the relays overhear pilot signals transmitted from source (e.g., ready-to-

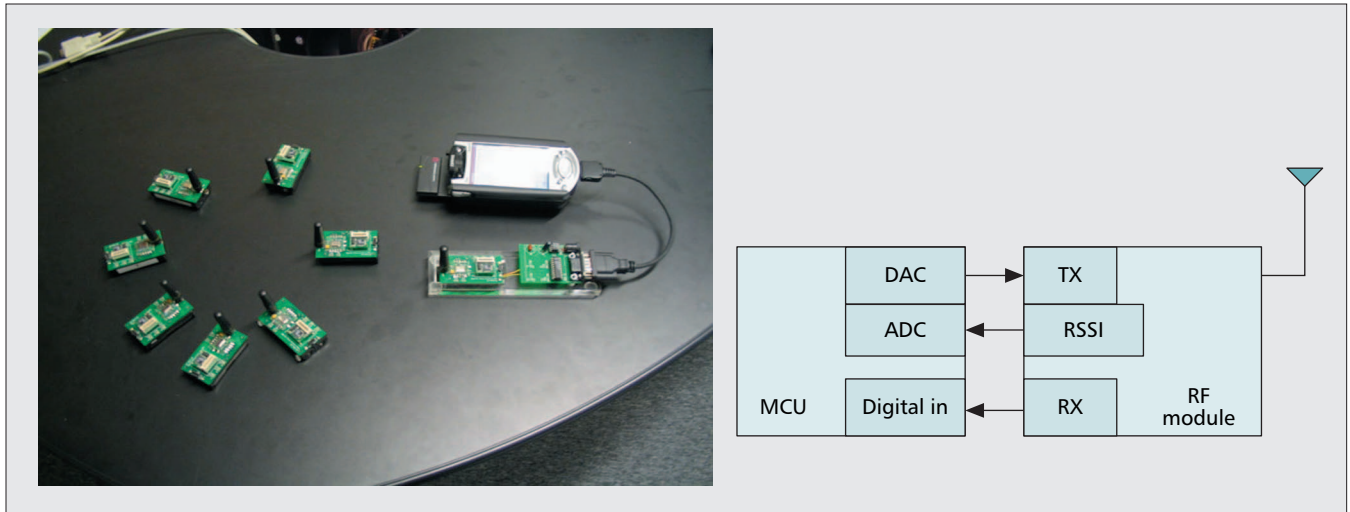


Figure 3. Low-cost embedded, software defined radios (eSDRs) were created in order to ensure full access to the physical, link (access), and routing layers. A microcontroller unit (MCU) was interfaced directly to a 916.5 MHz on-off keying radio. All necessary functions for transmission/reception, synchronization, detection, and access were implemented in software at the 8051 MCU. Serial port interfaces also allowed connection to handheld computers and external devices. The hardware cost for each eSDR was on the order of \$30 total (in quantities of 10).

send, RTS) and the destination (e.g., clear-to-send, CTS) and use them to estimate the channel conditions towards source and destination. A timing method has been proposed, so as the network discovers the relay with the best end-to-end channel conditions, without requiring global CSI information in a central controller or anywhere else in the network [4]. This is accomplished by an intelligent relay access scheme: as soon as each relay receives the pilot signal from the destination (CTS), it initiates a timer with an initial value inversely proportional to the quality of its own end-to-end channel conditions towards source and destination. The timer of the relay with the best channel conditions expires first and consecutively, that node notifies destination as well as the rest of the network for its availability, with a flag packet. The destination could further notify the rest of the network about the discovery of a useful relay.

Two functions of received signal strength (or, equivalently, signal-to-noise ratios) for each relay j about the path from source to relay γ_{Sj} and relay to destination γ_{jD} (which is the same for the path between destination to relay, due to reciprocity) have been proposed [4]:

$$\gamma = \min(\gamma_{Sj}, \gamma_{jD}), \quad (1)$$

$$\gamma = 2(\gamma_{Sj} \gamma_{jD}) / (\gamma_{Sj} + \gamma_{jD}) \quad (2)$$

The first seems more appropriate in regenerative (decode-and-forward) relay networks, while the second, which is a smoother function of the relay link strengths compared to the first, is more appropriate for amplify-and-forward relays [10].

The intelligent channel access scheme achieves selection of the relay that maximizes γ across all relays, without requirement for global CSI anywhere in the network. The intuition is simple: in order to find out the tallest student in a classroom, you do not need to measure the height of each and everyone in the room, but

instead you can invite all students to stand up and ask the tallest member to observe the class and raise her hand. As every channel access scheme, there is a nonzero probability for two or more relays to access the channel, within the same time interval. Probabilistic analysis of such an event for various wireless channel models with incorporation of practical limitations, such as the nonzero radio switch time from listening-to-transmit mode or propagation delay differences among the several links in the network, has been detailed elsewhere [4]. We note that even at the case where two or more relays have similar end-to-end channel quality, relay selection is still feasible by using randomized algorithms.

In our implementation, we exercised function (1) as a relay path quality metric. A 16-bit timer was used for each relay, and RTS/CTS packets transmitted from source and destination respectively, allowed the estimation of γ at every relay. The CTS reception initiated the distributed relay selection and the selected relay was used for a specific period of time, smaller than the coherence time of the channel. For 916.5 MHz carrier frequency and mobility of approximately 1 m/sec (corresponding to walking people), the channel coherence time becomes approximately 300 msec.² The measured indoor channel coherence time often revealed values close to 800 msec for a 916.5 MHz carrier and provided an approximate repetition rate for relay selection.

We note that our relay selection scheme is performed *proactively*, before the source transmits the message, in contrast to prior art that has focused on *reactive* schemes, where selection could be performed among relays that have correctly decoded the message. Such a design choice was intentional, since we attempted to minimize the total reception power, given that relays which are not selected, could enter an *idle* mode and avoid any reception for a specific period of time. In contrast, reactive schemes have *all*

² Channel coherence time is inversely proportional to Doppler shift, which depends on mobility speed and carrier frequency.

MCU	RF module
Pushpin MCU	RF-Monolithics
Architecture: 8051 (8-bit)	Frequency: 916.5 MHz
Clock speed: 22.1184 MHz	Baud rate: 115 kHz
12-bit ADC, 10-bit DAC	Modulation: on-off keying
Voltage: 3 V (2 AA batteries)	

■ **Table 1.** Specification details of the utilized hardware.

relays listen in order to receive the message, even though a subset of them eventually forwards the message and thus, total reception energy increases with network size. This might be a serious limitation of reactive schemes in battery operated networks.


Finally, we note that our protocol does not require any type of in-band transmissions, and thus any low-complexity radio transceivers can be employed.

Signaling and Receiver Structure — Information was sent periodically, in blocks corresponding to 16 characters of information, since that was the selected message length that could be displayed at the receiver display. The message would scroll from left to right with duration of approximately 3 s. Therefore, messages of 16 characters were sent within that period.

Before every message transmission, “best” relay selection would be performed, according to the described algorithm. Then, 16 frames were transmitted from the source, corresponding to

the 16 characters of the message. Each frame (out of those 16 frames) was repeated from the best relay, provided that it had been correctly decoded. That is why the measured signal structure, acquired with a digital oscilloscope and shown in Fig. 4 (second row, second picture), has empty slots destined for transmission from the selected relay. Each frame included the necessary synchronization preamble, followed by 4 bytes (32 bits) that included header information (source id, destination id, sequence id), data information, as well as a cyclic redundancy check (CRC) for error detection purposes (Fig. 4, upper-right figure). CRC information was required so that the relay could find out whether it had correctly decoded the message. The destination received information from the source as well as information from the best relay and decided about the original message. Even though we could use a maximum ratio combiner (MRC), we chose to further simplify the receiver structure: the receiver decoded both messages and kept the one with the correct message (assertion made with the help of the CRC field).

The signal structure (depicted in Fig. 4 from captured oscilloscope traces) is a specific example of how opportunistic relaying can be used in cooperative diversity contexts. It should be viewed as a concrete example for a specific application, built for demonstration purposes. Additional optimization could be performed if that was necessary. For example, the time required for “best” relay selection could be further reduced. We did not perform such optimization, since there was no such need in our slow bit-rate and low duty cycle demonstration. However, we have studied such optimization and have shown that relay selection can be efficiently performed within a time interval that is two to



■ **Figure 4.** Measured traces at the RX (receive) pin of the destination, using a digital oscilloscope. Upper left: The CTS packet from destination is followed by the Flag packet from the “best” (selected) relay. Then 32 frames follow (16 from source and 16 from selected relay). Lower left: Direct and selected relay frame transmissions are interlaced. Lower right: Direct communication when no relay retransmits. Upper right: The structure of each frame. A preamble is used for frame synchronization, followed by 32 on-off bits that include the message (8 bits), as well as CRC and protocol information (source address, destination address).

three orders of magnitude smaller than the channel coherence time, in slow fading environments [4]. Therefore, sampling of space can be repeated with small overhead more often, within intervals smaller than the coherence interval. Such sampling would be sufficient even for cases where the wireless channel fluctuated in a discontinuous and abrupt fashion.

Additionally, our embedded radios did not have much computation power, given the 8 bit processor structure. More complex receiver structures, like a (MRC) receiver or an advanced error correcting code combiner receiver require more powerful computation and could be used in conjunction with a powerful microprocessor for each embedded radio. Note, however, that increased complexity at each receiver increases the necessary required reception energy, having a significant impact on the overall energy budget [8]. We chose to keep the individual nodes as simple as possible and rather exploit distributed intelligence at the network level.

Practical Considerations — One of our main concerns during implementation was the limited resolution of the analog-to-digital converter (ADC) at each relay radio, during the evaluation of the signal strength path, towards source and destination. Fortunately, the microcontroller's 12 bit ADC was proven adequate in practice. A slight movement in space could easily result to a factor of 10 in strength fluctuation, as we experimentally observed and, therefore, crude digitization of such variation is sufficient.³

A second concern during implementation was about specific channel estimation algorithms at the relays, given the 8-bit architecture of each micro-processor which resulted to limited computation performance: we had to reduce all floating-point calculations in order to improve accuracy and speed. That was the main limitation from a hardware perspective and could be resolved by using more advanced microcontrollers.

A third concern involved the case when a collision among relays did occur. There are several possibilities with regard to what the relays should do after a collision. One solution could be to have source or destination notify the relays that they have collided, especially when the relays cannot listen to each other; in that case, one of the relays could back off. This is easy to implement, since the relays switch between receive and transmit mode periodically in order to receive the information from the source. Therefore, a control-bit indicating collision and transmitted by the source (or the destination) is straightforward. Another, even simpler solution could be to have the relays that indeed participate in the retransmission randomly avoid retransmitting information and wait to see if other relays are retransmitting. This is a valid approach when there is a path between source and destination, and the additional path via the best relay is used to increase reliability. In the case when there is no direct path between source and destination, that solution is clearly suboptimal. For our room-size demonstration, where a path between source and destination was available (although with variable quality), that approach was followed.

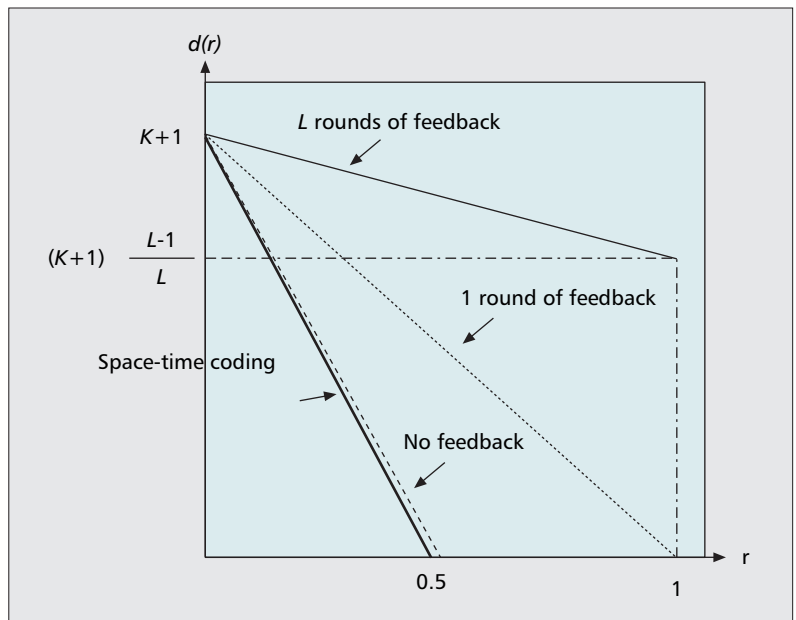


Figure 5. The diversity-multiplexing gain trade-off (DMT) for the implemented protocol (thick line). Diversity $d(r)$ provides a measure of reliability, while multiplexing gain (or degrees of freedom) r provides an indication of the achievable rate (b/s/Hz). Relay selection provides the same DMT as space-time coding for K relays. If one round of feedback is utilized, DMT performance is improved. Additional rounds (total L rounds of feedback) of feedback enhance the performance. Intelligent relay selection offers cooperation benefits without simultaneous in-band transmission and allows for utilization of low-complexity radio hardware.

Finally, we need to emphasize the fact that the utilized cooperative diversity technique is about increasing reliability in slowly fading environments. Sampling of space, in the form of pilot signals transmitted from source or destination, needs to be periodically repeated. Emphasis on this work was given in minimizing the overhead time required for best relay selection and no assumptions were made regarding smoothness of the wireless channel fluctuations. Future work could focus on dynamic channel access measurement, modeling, and prediction so as to minimize the overhead for pilot signals and channel estimation.

DISCUSSION

We established a demonstration of cooperative diversity using low-complexity, commodity radio and attempted to address all challenges. The distributed nature of cooperative relaying was by far the most intriguing difficulty. The introduction of an intelligent channel access scheme at the link layer (layer 2) with characteristics of adaptive routing (layer 3) provided distributed ways for acquisition of network state information and coordination. Such intelligence at layers 2 and 3 allowed simplification of the physical layer (layer 1), and thus utilization of low-complexity radios was made feasible. Furthermore, the proactive nature of relay selection reduced the total network reception energy.

Information theoretic analysis of our protocol, for both amplify-and-forward as well as decode-and-forward relays, revealed maximum

³ We note, however, that the timing protocol used for relay selection is benefited by a fine resolution at the ADC.

Our work has provided a concrete implementation of cooperative diversity antenna arrays using commodity hardware and hopefully will spark interest in the research community to study cooperation in all layers and become adventurous enough to implement them in custom, experimental test-beds.

diversity order on the number of participating nodes in the system, even though a single relay transmits [4]. Moreover, the diversity-multiplexing gain trade-off⁴ (DMT) in opportunistic relaying was the same as that in distributed space-time coding schemes, where in-band, simultaneous transmissions and optimal processing are assumed [4]. If one (or several rounds of feedback) from destination towards the selected relay is available, then the DMT performance can be further enhanced [11], offering improved reliability without sacrifice in terms of the achieved rate. In that way, the implemented cooperative diversity technique resembles the benefits of optimal centralized antenna arrays, without in-band, simultaneous relay transmissions (Fig. 5). Subsequent theoretical analysis has shown that under an aggregate transmission power constraint, the implemented technique outperforms space-time coding techniques at the finite-SNR regime [10]. Those findings suggest that the approach followed in this work not only allows implementation with commodity hardware, but also outperforms other techniques found in the literature.

Our work demonstrates that the benefits of cooperative diversity do not necessarily arise from in-band, simultaneous transmissions, but instead emerge from:

- The existence of several potential relay paths between source and destination and
- The dynamic discovery of the most useful of them, by means of distributed and adaptive techniques

Future work should extend our results to the wideband regime, possibly through the use of multicarrier modulation (OFDM) with implementation of our algorithms for each subband. More work is also needed to extend our scheme to multihop environments and quantify end-to-end performance.

Cooperative diversity is, by nature, a cross-layer approach and hence requires exploitation of the physical, link, and routing layers which have been traditionally addressed assuming non-cooperative communication. Our work has provided a concrete implementation of cooperative diversity antenna arrays using commodity hardware and hopefully will spark interest in the research community to study cooperation in all layers and become adventurous enough to implement them in custom, experimental test-beds.

ACKNOWLEDGMENTS

The authors would like to thank the members of the Viral Communications Group at the MIT Media Lab for their invaluable help throughout this work. They are also grateful to Ashish Khisti, Josh Lifton, Jamey Cooley, Dean Christakos, Iliia Mirkin, and Marios Michalakis for their kind assistance and advice during various stages of the demonstration development.

⁴ The DMT provides a common tool for characterization of reliability versus degrees of freedom, and has been recently adopted in classic multi-antenna systems.

REFERENCES

- [1] J. H. Winters, "On the Capacity of Radio Communication Systems with Diversity in Rayleigh Fading Environment," *IEEE JSAC*, vol. 5, no. 5, June 1987, pp. 871–78.
- [2] A. Sendonaris, E. Erkip, and B. Aazhang, "User Cooperation Diversity, Part I and Part II," *IEEE Trans. Commun.*, vol. 51, no. 11, Nov. 2003, pp. 1927–48.
- [3] 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*, vol. 50, no. 12, Dec. 2004, pp. 3062–80.
- [4] A. Bletsas *et al.*, "A Simple Cooperative Diversity Method Based on Network Path Selection," *IEEE JSAC*, vol. 24, no. 3, Mar. 2006, pp. 659–72.
- [5] J. N. Laneman and G. W. Wornell, "Distributed Space-Time Coded Protocols for Exploiting Cooperative Diversity in Wireless Networks," *IEEE Trans. Info. Theory*, vol. 59, Oct. 2003, pp. 2415–2525.
- [6] K. Lu, S. Fu, and X. G. Xia, "Closed Form Designs of Complex Orthogonal Space-Time Block Codes of Rates $(k + 1)/(2k)$ for $2k - 1$ or $2k$ Transmit Antennas," *IEEE Trans. Info. Theory*, vol. 51, no. 12, Dec. 2005, pp. 4340–47.
- [7] H. E. Gamal, G. Caire, and M. O. Damen, "Lattice Coding and Decoding Achieve the Optimal Diversity-Multiplexing Trade-Off of MIMO Channels," *IEEE Trans. Info. Theory*, vol. 50, no. 6, June 2004, pp. 968–85.
- [8] R. Min *et al.*, "Energy-Centric Enabling Technologies for Wireless Sensor Networks," *IEEE Wireless Commun.*, vol. 9, no. 4, 2002, pp. 28–39.
- [9] J. Lifton, M. Broxton and J. A. Paradiso, "Experiences and Directions in Pushpin Computing," *Proc. 2005 Symp. Info. Processing in Sensor Networks*, Los Angeles, CA, Apr. 25–27, 2005, pp. 416–21.
- [10] A. Bletsas *et al.*, "Cooperative Diversity with Opportunistic Relaying," *IEEE WCNC '06*, Las Vegas, NV, Apr. 2006.
- [11] A. Bletsas, A. Khisti, and M. Z. Win, "Low Complexity Virtual Antenna Arrays with Cooperative Relay Selection," *ACM Int'l. Conf. Wireless Commun. and Mobile Computing*, Vancouver, Canada, July 2006.

BIOGRAPHIES

AGGELOS BLETSAS (aggelos@media.mit.edu) received his diploma degree (with excellence) in electrical and computer engineering from Aristotle University of Thessaloniki, Greece, in 1998. In 1999 he joined the Massachusetts Institute of Technology (MIT) Media Laboratory, where he earned M.S. and Ph.D. degrees in 2001 and 2005, respectively. His research interests span the broad area of scalable wireless communication and networking, with emphasis on relay techniques, signal processing for communication, radio hardware/software implementations for wireless transceivers and low-cost sensor networks, time/frequency metrology, and nanotechnology. He is currently researching and developing detection and signal processing techniques for the long-range backscatter wireless channel.

ANDREW LIPPMAN (lip@media.mit.edu) received B.S. and M.S. degrees in electrical engineering from MIT, and a Ph.D. degree from the Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland, in 1995. His history at MIT spans more than 30 years. His work at the Media Laboratory has ranged from wearable computers to global digital television. Currently, he heads the Laboratory's Viral Communications program, which examines scalable real-time networks whose capacity increases with the number of members. In addition, he codirects MIT's interdisciplinary Communications Futures program, and established and directs the Digital Life consortium, which focuses on both technical invention and human understanding, and works to create a networked world where communication becomes fully embedded in our daily lives.