Leveraging Physical Layer Cooperation For Energy Conservation

Christopher Hunter, *Member, IEEE,* Lin Zhong, *Senior Member, IEEE,* Ashutosh Sabharwal, *Senior Member, IEEE*

***Abstract*—Physical layer (PHY) cooperation is a technique for achieving MIMO-like performance improvements on small devices that cannot support antenna arrays. Devices in a network transmit on behalf of their neighbors to act as "virtual MIMO" antennas. Since small devices are typically battery constrained, PHY cooperation immediately leads to the following question related to the energy efficiency (bits-per-joule) of devices – is the performance improvement worth the extra energy costs of transmitting for others? Through an in-depth hardware testbed study, we find that PHY cooperation can improve energy efficiency by as much as 320% or it can *reduce* energy efficiency by as much as 25% depending upon topology. With this performance gap in mind, we propose the Distributed Energy-Conserving Cooperation (DECC) protocol. DECC tunes the amount of effort each device dedicates to providing cooperative assistance for others so that the energy each device spends on cooperation is commensurate with the personal benefits received by that device. With DECC, users can tune their level of cooperation with completely node-localized decision making. Thus, DECC allows nodes to tap into a large energy efficiency benefit, suffering only a bounded, preset loss when this benefit is not available.**

I. Introduction

As demand for mobile data increases, it is becoming increasingly important to improve physical layer spectral efficiency. The so-called "spectrum crunch" [1] has led to rethinking the physical layer beyond the traditional standpoint of establishing a reliable connection between only two devices (the transmitter and the receiver). One such concept is that of cooperation between devices at the signal level (see, e.g., [2]) to harness the antennas of multiple devices to increase rates, reliability and transmission range.

Physical layer (PHY) cooperation harnesses MIMO-like spatial-diversity benefits with only single-antenna devices and also exploits topological asymmetries where one device may have a better link quality than its neighbors.

Information and communication theoretic investigations of physical layer (PHY) cooperation have shown its potential for multifold network capacity improvement; see, e.g, book length exposition [3, and references therein]. Recently, experimental results using software-defined radio platforms have also been reported [4]. These results further demonstrate the capacity improvement potential of symbol-level cooperation at the PHY layer. However, a key barrier to the practical deployment of cooperation in mobile devices remains unsolved: when two mobile devices cooperate, the performance gain from cooperation may be uneven. Consider the simplest and most widely studied scheme in which one mobile device helps another by forwarding overheard packets from the latter. The device being aided is likely to gain in both link throughput (measured in bits per second) and energy efficiency (measured by bits per joule); on the other hand, the device giving assistance may well suffer in both throughput and efficiency. A significant roadblock to the widespread adoption of cooperative communications is that mobile devices have limited battery capacity – users will likely not want a technology where their mobile phone depletes its battery while attempting to assist others in the network.

The challenge is greatly simplified with a centralized approach, where a central entity like an access point gathers information about all the inter-device channels and then decides who should help whom in a mutually beneficial manner. However, such an approach is not scalable for large networks, as periodically collecting all link information and transmission of scheduling decisions will lead to very high measurement and management overhead.

We provide a scalable solution to the problem of energy conservation in the context of Wi-Fi-like devices by designing a completely device-localized protocol called Distributed Energy-Conserving Cooperation (DECC)[1]. DECC has three components. First, DECC estimates the extent of benefit that cooperation provides by measuring an energy efficiency reference by periodically disabling its cooperative mode. This allows the device to measure the amount of benefit it derives by cooperating with other devices. This estimation is performed completely asynchronously, without any assistance or packet exchanges with the Wi-Fi access point (AP) or other devices.

Second, based on the estimated benefit, the device scales its cooperative "effort," a dimensionless quantity, to control its level of cooperation. This parameter is adjusted by a proportional-integral control loop that tracks the ongoing energy efficiency calculations.

Third, the effort parameter is used to modify the MAC layer behavior by allowing the MAC to decide on a per-packet basis whether it will cooperatively transmit the last overheard packet or not. The determination of whether to cooperate on a packet also uses information about the instantaneous channel quality estimates of the relevant links. These estimates are derived from the last received packet via each link so no additional

C. Hunter, L. Zhong and A. Sabharwal are with the Department of Electrical and Computer Engineering, Rice University, Houston, TX, 77005 USA e-mail: [chunter, lzhong, ashu]@rice.edu. This work was partially funded by NSF grants CNS-0551692, CNS-0619767, CNS-0923479 and CNS-1012921.

[1]This paper presents a condensed version of the first author's Ph.D. dissertation [5].

overhead in channel estimation is incurred.

A key feature of DECC is that it is "association-free." That is, devices do not care whose packet they are helping but simply adapt their cooperative behavior to manage their own total energy efficiency loss or gain. This is an important aspect as it allows the device behavior to automatically adapt to topology changes due to device mobility and traffic changes.

We use the Wireless Open-Access Research Platform [6] (WARP) hardware testbed that employs a real-time OFDM PHY with a random access cooperative protocol [4] to investigate and determine two crucial findings. First, we show that there exist network situations where PHY cooperation is beneficial for *all* devices in the network. In these situations, even selfish devices will cooperate. Second, we show that in many cases, a little sacrifice by a helping device may produce disproportionally large gain for the device being helped. Even a small amount of altruism can lead to significant gain for others. Inspired by these observations, DECC allows each device to constrain its level of altruism to a user-selectable bound.

In addition to implementing DECC on WARP, we have also developed a custom ns-2 [7] simulator to show how DECC's behavior scales to larger network sizes and topologies. With DECC, devices can bound their maximum acceptable loss in energy efficiency to, for example, only 5%. But more importantly, DECC allows the devices to derive large energy efficiency benefits, if and when they are available, with completely localized decision making.

We envision many uses of DECC in actual devices. Wireless users are already growing accustomed to crowded airwaves. As such, these users could be inspired to enable a bounded altruism mode (e.g. they contribute no more than 5% energy efficiency as a worst case) in the hope that they will be helped in future by other devices as they move in the network and encounter poor connectivity locations. Since devices control their individual level of altruism, the amount that they are willing to help "for free" (5% in this example) can be changed by each user based on their specific use case scenario. If no one cooperates, DECC simply reduces to the base MAC and PHY, i.e, current Wi-Fi. However, as more devices participate in cooperation, they gain from the altruism of others and, in turn, dedicate more resources to providing cooperative aid. The gains and losses of cooperation change as users move through the network since that movement leads to topology changes and, therefore, different network conditions.

Methods to induce cooperation between nodes in a network have been extensively studied in multiple contexts. For example, in the context of ad hoc networks, there has been extensive work to promote cooperation so that all nodes participate in a fair manner to increase either the network capacity [8] or improve energy efficiency [9], [10]. The key mechanisms in this large body of work are inspired by game theoretic principles, where each node is considered rational and optimizes its utility function [11, and references therein]. The implicit assumption in game-theoretic development of cooperative (or competitive) networks is that each node is fully aware of whom it is playing this game against – which is akin to each node knowing about all the nodes in the network, their utility function and what policy they are using to maximize the utility. This information is critical for proving properties of the game (e.g. Pareto optimality, Nash equilibria). However, that level of information is not available in Wi-Fi networks, since the potential overhead of collecting, managing, and sending this information to all nodes would be very high and hence not scalable. In contrast, we focus on the design of network protocols where each node is completely agnostic to the network topology and has no access to other nodes' utilities or strategies. Despite the radically different approaches to the problem of energy efficient cooperative communications, many conclusions drawn from our work coincide nicely with the game theoretic literature. For example, the Nash bargaining solution from [10] results in a scenario where cooperation is employed among users when it is mutually beneficial to do so with all users. Otherwise, the system naturally falls back on non-cooperative communication when that mutual benefit cannot be achieved. This is not unlike the behavior achieved by our proposed DECC protocol, where users ratchet back the cooperative aid they provide others when they determine that cooperation is not beneficial to them.

A related subject that is widely studied is the problem of relay selection and power control for energy efficient cooperative communications [12]–[14]. In [12], the authors provide simple relay selection and power control heuristics from analysis of optimal strategies. Cooperative links that employ these heuristics are shown to have an improvement in energy efficiency versus non-cooperative links. In [13], the energy costs of CSI training are modeled, and the authors show that cooperative communications can still improve energy efficiency over non-cooperative communications when the number of relays selected is limited to the best few. Finally, the work in [14] provides a slightly different look at the problem in that the work does not rely on physical layer cooperation as described in our work or in the preceeding citations. Instead, a single relay is chosen through which to route the packet without any waveform-level combination for maximum ease of implementation. By selecting a single relay to forward a packet, a source can reduce the power of its transmissions and rely on the relay to send the packet the rest of the way to the destination. In all of these works, cooperation is shown to be capable of greatly improving energy efficiency for sources that use it. However, the relays studied in this work do not have traffic of their own, so the broader discussion of energy efficiency in the context of spending energy for one's own traffic alongside spending energy aiding others in the network is unexplored. In our work, we generalize the discussion of cooperative relaying to allow devices in the network with their own traffic to provide cooperative assistance to one another. Furthermore, relay selection requires setting up a link – a process the consumes time, creates overhead, and thus may not be applicable to high mobility scenarios. In contrast, DECC is "association-free" and can adapt to network changes without an explicit relay selection phase.

Finally, considerable study is available at higher layers in the networking stack such as cooperative routing. In these works, the use of cooperative links complicates the problem of routing because minimum distance routing decisions no longer

capture all the information about the capabilities of any given cooperative link. In [15], a strategy is devised to minimize power consumption along the route by exploiting cooperative communication links. The use of DECC to tune the energy efficiency of the individual cooperative links that make up these routes is certainly possible, and the study of this joint problem is an interesting area of future work.

The rest of the paper is organized as follows. Section II provides an overview of the energy efficiency challenge created by PHY cooperation and formally defines the key metric used in this work. Section III presents results from in-depth experimental characterization of network primitives with cooperating devices and distills the insights for the design of DECC, which is presented in Section IV. Section V describes our implementation of DECC on a hardware testbed and our evaluation of DECC using both prototype-based measurements and simulations of a larger network.

II. Energy Efficiency Using Signal-Scale Cooperation

By design, signal-scale cooperation is intended for small devices (e.g. smartphones or tablets) that are unable to directly achieve spatial diversity or multiplexing through the use of antenna arrays. These small devices share another fundamental trait: they are typically also constrained by limited battery life. As such, the *energy efficiency* of a device is extremely important. When this is taken into account, signal-scale cooperation exposes a clear tradeoff – in order to allow devices in the network to realize the significant rate and reliability performance gains that cooperation can provide, other devices must spend additional wireless transceiver energy to act as their cooperative transmitters.

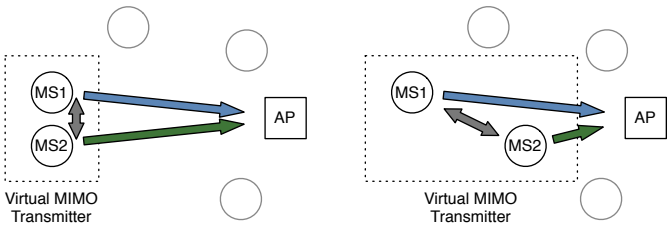

(a) Symmetric Topology. (b) Asymmetric Topology.

Fig. 1. Asymmetries in topology yield asymmetries in the performance/energy tradeoff.

The tradeoff between increased energy consumption and better performance is complicated by the fact that not everyone in a wireless network has the same amount to gain from cooperation nor do they have the same potential energy costs. Figure 1 illustrates a common uplink scenario that highlights this problem. In Figure 1(a), mobile stations MS1 and MS2 are equidistant from the access point AP and see similar distributions of fluctuations in the wireless channel. Due to the symmetry, the gains and costs of cooperation will be comparable for both devices. In Figure 1(b), MS2 is much closer to the AP than MS1 and therefore has a much higher SNR link to the AP on average. The asymmetry yields two important effects:

- MS2 has less to gain from cooperation than MS1 since it can rely more heavily on its direct link to the AP.
- MS2 has more energy to lose from cooperation since it can provide more substantial and more frequent cooperative assistance to MS1.

As above three-node example demonstrates, the "winners" and "losers" from cooperation are topologically dependent. In Section III, we perform an in-depth experimental characterization of topological dependence to demonstrate how severe the performance/energy tradeoff can be.

A. Key Metric

We use "energy efficiency" as our design metric. Specifically, we define energy efficiency to be

$$\gamma = \frac{\text{\# of bits the device successfully transmits + receives for itself}}{\text{total energy the device spends transmitting and receiving}},$$

where γ has a unit of bits/joule. Practically speaking, energy efficiency captures the amount of data that comes into and goes out of a device for a certain amount of battery capacity spent. The denominator of this expression captures the costs of the device's operation, even if that operation is in the cooperative service of other devices. The numerator, however, only captures the bits transferred that are relevant to the device performing the calculation of this metric. Any given cooperative transmission for another device is not directly beneficial to the device acting as the cooperative relay (i.e. the cooperative transmission will only increase the denominator of this metric). Cooperation can be of benefit on a longer time horizon than any single cooperative transmission – the presence of cooperative relays in the network may increase the number of bits a device can successfully transmit or receive (the numerator), even if that same device may also be increasing the energy it spends by helping others (the denominator). It is these two competing forces of cost and benefit of cooperation that is the focus of this work.

As a concrete example, suppose two devices (MS1 and MS2) each have a flow of traffic destined for a common access point (flows 1 and 2, respectively). Let X_i represent the number of bits delivered in flow i. Let E_{ij} represent the energy spent by device j for transmitting bits in flow i. In a non-cooperative system NC, the energy efficiency γ_j^{NC} of device j is

$$\gamma_{\text{MS1}}^{\mathsf{NC}} = \frac{X_1^{\mathsf{NC}}}{E_{11}^{\mathsf{NC}}}$$
$$\gamma_{\text{MS2}}^{\mathsf{NC}} = \frac{X_2^{\mathsf{NC}}}{E_{22}^{\mathsf{NC}}}.$$

However, for a cooperative system C where MS1 and MS2 aid one another, their energy efficiencies are

$$\gamma_{\text{MS1}}^{\mathsf{C}} = \frac{X_1^{\mathsf{C}}}{E_{11}^{\mathsf{C}} + E_{21}^{\mathsf{C}}}$$

$$\gamma_{\text{MS2}}^{\text{C}} = \frac{X_2^{\text{C}}}{E_{22}^{\text{C}} + E_{12}^{\text{C}}}.$$

From an energy efficiency standpoint, cooperation is only beneficial to device j if the number of additional bits delivered for its traffic flow offsets the increased energy burden of aiding others.

III. Experimental Characterization

To quantify the energy efficiency challenge for signal-scale cooperation, we study small, representative topologies using a custom hardware testbed implementation. In the subsequent experimental characterization, we make two key findings:

- Cooperation can be mutually beneficial to devices, but the regimes where mutual benefit occurs are topologically small and therefore unlikely in general.
- In regimes where one device suffers energy efficiency loss for the gain of another device, the loss to the former is smaller than the gain to the latter.

These findings provide the motivation for our protocol, DECC, in Section IV.

A. Hardware Platform Selection and Details

PHY cooperation is a challenging subject for empirical investigation because it requires novel techniques at multiple layers of the networking stack – including custom waveform construction. We choose Rice University's Wireless Open-Access Research Platform (WARP) [6] for our study as it allows for (a) custom physical layer behavior and (b) real-time operation to preserve the time scales over which multiple devices actually interact in the real world. As part of the open-source WARP project, our earlier works have made a decode-and-forward (DF) cooperative PHY [4], [16] and the Power-controlled Distributed On-demand Cooperation (PDOC) MAC [17] freely available in its repository. We use these designs as the starting point for our study.

Particularly relevant to our work is the PDOC MAC protocol. PDOC (and its prequel, DOC [18]) enables cooperation only in situations where MAC-level retransmissions would occur anyway. In this way, this protocol is "on-demand" and gracefully falls back on standard non-cooperative communication when cooperation is not needed. Devices employing PDOC first attempt to communicate using direct, non-cooperative communication much like the 802.11 distributed coordination function (DCF). In the event that the transmission fails, a cooperative retransmission occurs with simultaneous source and relay transmissions. The cooperative transmission is triggered by a receiving destination node through the explicit transmission of a negative acknowledgment control packet (NACK) when fading causes packet errors on the direct source-to-destination link. Significantly, this NACK carries feedback about the signal quality of the source-to-destination link. This feedback allows the relay to adjust its transmit power to a level that ensures that the link is reliable but uses no more power than is necessary. A detailed summary of this process is provided in Section IV-C.

For the experiments in this work, we have extended the PDOC design to implement the performance of maximal-ratio-combining (MRC) between the first source transmission and relay-only transmission. Because the DF cooperative PHY will only allow cooperative relay transmissions when the relay's reception passes a CRC32 checksum, this MRC combination cannot introduce errors to the system. Details of this extension are available in Appendix A. Throughout our study, we take an *event-driven* approach to power measurements. Offline, we calculate the energy costs of transmission and reception events and save these values to the WARP hardware. Then, the real-time FPGA implementation on WARP uses these energy cost values to calculate in real time how much power is being drawn. In this way, our design prototype effectively models the behavior of a final implementation where the design could request power draw information from an operating system that is already monitoring battery levels. The details of the offline power measurements that enable this study are available in Appendix B.

B. Experimental Results

Among distributed cooperative MAC protocols, PDOC is one of the best known energy conservers because it employs power-control at the relay. As such, PDOC is well-suited for the study of energy efficiency and serves as the baseline for this work. Using our WARP-based testbed, we measure the achieved throughput and energy efficiency of the cooperative PDOC protocol alongside a non-cooperative CSMA/CA protocol.

1) Topologies Studied: To study the impact of topological asymmetry on the gains and costs of signal-scale cooperation, we first study the atomic case of two flows in a three-node system, where mobile stations with backlogged queues are uplinking to a common access point. For all of the hardware experiments in this work, we use the Azimuth Systems ACE 400WB channel emulator [19]. All links undergo a single-tap Rayleigh fading channel with a velocity of 1.2km/hr to emulate channel coherence times consistent with walking-speed mobility.

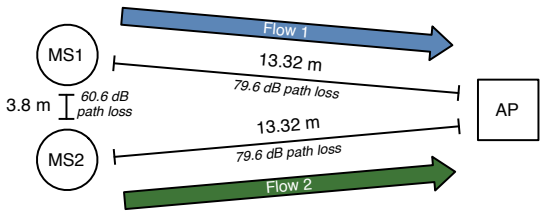

(a) Scenario A.

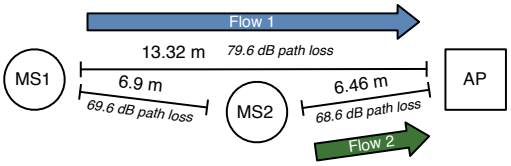

(b) Scenario B.

Fig. 2. Topologies under study.

We use the channel emulator to construct the two topologies shown in Figure 2. Using a path loss model, we convert the distances shown in the figure to average path losses using a path-loss exponent of 3.5. This is consistent with the urban environments in which non-line-of-sight Rayleigh fading occurs. We configure the emulator with the path loss values shown in Figure 2 and let the emulator then superimpose the Rayleigh channel fading on top of those values.

In Scenario A, the mobile stations are equidistant from the AP to create a symmetric topology. In Scenario B, MS2 is placed approximately halfway between MS1 and the AP. This introduces an asymmetry in the network where the flow sourced by MS2 undergoes less severe fading events.

All results in this section come from 15 separate 120 second trials where all random number generators (at each node and the channel emulator) are reseeded. Error bars represent the standard deviation across those trials. All data transmissions are 1,300 byte packets modulated at 16-QAM in 10MHz of bandwidth. The data presented here represent the cumulative successful reception of over 28 billion bits, or 2.7 million data packets of 10,400 bits each.

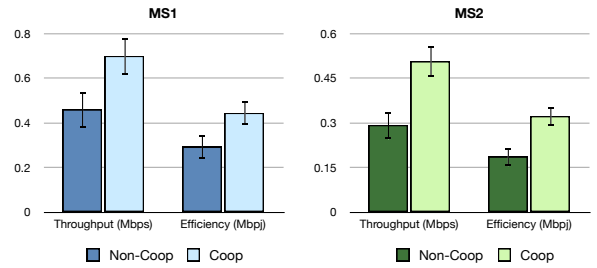

(a) Scenario A.

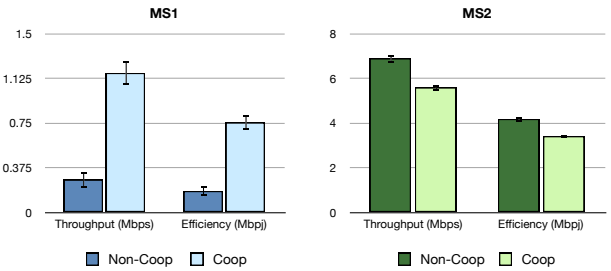

(b) Scenario B.

Fig. 3. Experimental results for Scenarios A and B.

2) Scenario A: Symmetric Topology: In the symmetric topology, both flows are in equally harsh fading environments. Figure 3(a) shows that cooperation is mutually beneficial to both flows. There are many cases where the instantaneous fades between each MS and the AP create link outages. By cooperating and achieving spatial diversity, the links between each MS and the AP can often be improved, resulting in an energy efficiency improvement in excess of 50% for each flow.

3) Scenario B: Asymmetric Topology: In the asymmetric topology, Flow 2 is considerably closer to the AP than Flow 1. As such, Flow 2 incurs fewer transmission failures due to adverse channel conditions. Figure 3(b) shows a significant result: while Flow 1 achieves a dramatic improvement in energy efficiency performance ($+320\%$), that benefit comes at the cost of Flow 2 (-18%). In the asymmetric topology, MS1 relies on MS2 to act as its relay far more often than the reverse case. MS2 spends considerable resources to help MS1 without meaningful reciprocation.

C. Establishing the Trend

While the study of Scenarios A and B in Sections III-B2 and III-B3 establishes the severity of improvement or degradation in energy efficiency, it does not establish how likely these two scenarios are. We examine the trends connecting these two regimes with another empirical study considering a linear topology.

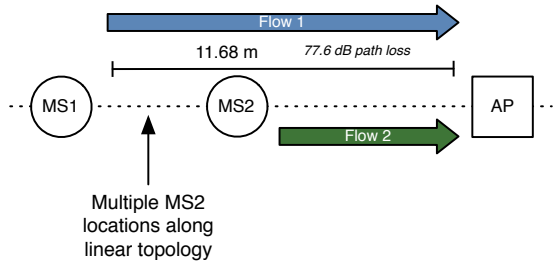

Fig. 4. Linear topology.

In the topology in Figure 4, MS2 is placed at a number of locations along the linear topology that connects MS1 and the AP. In this way, MS2 can be "swept" through the mutually beneficial and one-sided beneficial regimes that we previously established (in Scenarios A and B, respectively).

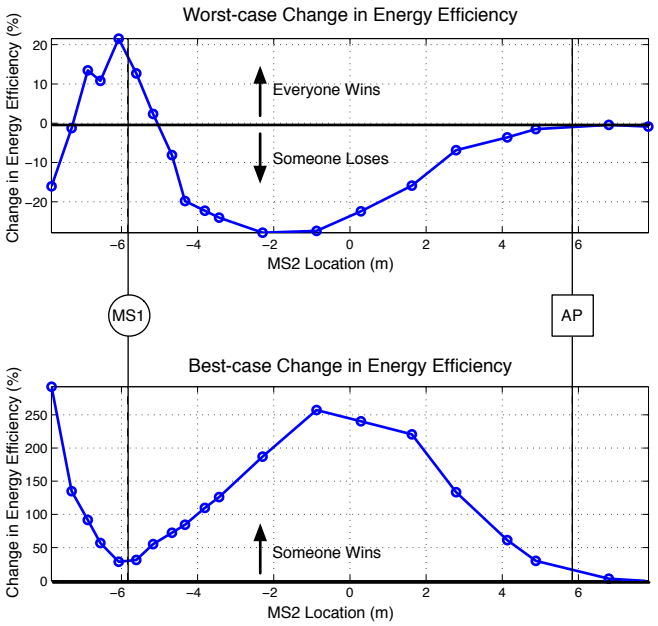

Fig. 5. Relative to non-cooperative CSMA/CA, cooperation yields a change in energy efficiency for both mobile stations. This change can be a improvement or a degradation depending on topology.

In Figure 5, we first examine the worst-case change in energy efficiency relative to non-cooperative communication. At every experiment point, we plot either the change in energy

efficiency of MS1 or MS2 (whichever is worse). When the worst-case change is larger than 0%, this describes a regime where *everyone wins*. Both MS1 and MS2 are better off with cooperation than without it. While this situation is clearly possible, it is not likely in a general network – there is a span of only approximately two meters near MS1 where mutually beneficial cooperation occurs. Thus, our first experimental finding is that cooperation can be mutually beneficial to devices, but the regimes where this occurs are topologically small and therefore unlikely in general.

We next examine the best-case change in energy efficiency. In the regimes where "someone loses," the other device does dramatically better. The sacrifice from one device yields tremendous improvements in energy efficiency to the other. Accordingly, our second experimental finding is that in regimes where one device suffers energy efficiency loss for the gain of another device, the loss to the former is smaller than the gain to the other. In this experiment, even when the results are at their most one-sided near MS2 location -1m, the energy efficiency loss of MS2 (approximately -28%) is more than exceeded by the gain in energy efficiency of MS1 (approximately +250%).

These two findings motivate the need to allow for *some* degree of altruism among devices. Not only are mutually beneficial regimes rare, but the gains that can be achieved for one device by allowing another device to suffer even a small amount can be very large. In this work, we present the Distributed Energy-Conserving Cooperation (DECC) system that can allow a device to provided a bounded amount of altruism in the cooperative aid of others in a network in a completely distributed manner.

IV. Protocol Description

In this section, we present the details of DECC. This design provides a completely distributed mechanism for devices to enforce *bounded altruism* in cooperative networks. With DECC, devices realize substantial improvements in energy efficiency or, at worst, they suffer a small and preset amount of degradation.

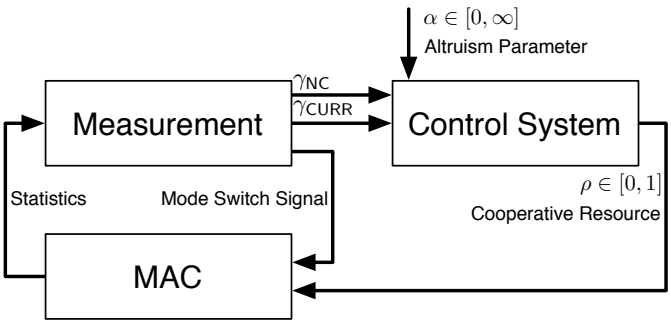

Fig. 6. DECC block diagram.

Figure 6 illustrates the block diagram for DECC. It is important to emphasize that this system is implemented at *each node in the network* and is completely distributed. There are three main subsystems of DECC: the measurement module, the control system module, and the MAC. Together, the measurement and control system modules determine the degree to which user cooperation is helping or hurting the device and then, accordingly, scale the amount of cooperation the device will employ by tuning the underlying MAC. The measurement and control system modules are analogous to a thermostat of an air conditioning system. They judge the "temperature" of the node's performance with cooperation and then make adjustments to the level of cooperation in which the underlying MAC takes part. We now provide detailed descriptions of each of these subsystems.

A. Measurement

The measurement module is responsible for determining whether or not cooperation is helping or hurting the energy efficiency of the device. It actively switches a device to and from a mode where cooperation is not allowed. During these mode switches, the measurement module directly measures the energy efficiency of the device (γ from Section II-A). Energy efficiency is measured by monitoring the statistics from the MAC. In a unit of time (e.g. once per second[2]), the measurement module asks the MAC for the number of bits that have successfully been transmitted or received. The measurement module also determines the amount of energy spent (in joules) over the same time. For example, if the DECC algorithm were running on a smartphone, the algorithm could ask the OS running on the smartphone for this energy expenditure information.[3] The ratio of successful bit transmissions and receptions to the total energy spent is the energy efficiency γ. The measurement module produces two outputs:

γ_{CURR}: the current energy efficiency of the device.

γ_{NC}: the energy efficiency of the device with cooperation disabled.

γ_{NC} provides ground truth for a device running DECC and allows the device to determine whether cooperation is currently helping it or not. The method for measuring non-cooperative energy efficiency requires each device to occasionally disable cooperation for a period of time when it senses changes to incoming or outgoing traffic load. During this explicit non-cooperative period, $\gamma_{\mathsf{NC}} = \gamma_{\mathsf{CURR}}$. The device can perform this switch by simply refusing to send any NACKs and refusing to act as a relay for any other device in the network. Additionally, the device will include a single bit flag inside any outgoing transmissions to notify any destinations that they should not send any NACK packets for that transmission. Including this flag will ensure that the device neither acts as a cooperative relay for others nor is aided by others on its transmissions. The frequency and

[2]In Section V-B3, we use our prototype implementation to study the duration of time needed to obtain an accurate measurement.

[3]In our implementation, a design on WARP lacks a mechanism to directly measure the amount of energy a node uses. Instead, our design keeps track of the amount of time a node has spent receiving and its number of packet transmissions with the amount of power such actions use to compute this energy expenditure.

duration of the non-cooperative periods are parameters that are exposed to higher layers. When a device senses changes to traffic patterns or mobility, these parameters can be adjusted at runtime accordingly. Each device asynchronously switches modes with no coordination with others. This allows each device to update its own γ measurement parameters.

B. Control System

The control system is responsible for scaling the effort the device dedicates to acting as a cooperative relay for others such that the effort is sufficiently offset by the performance gain from cooperation. The control system uses the outputs of the measurement module as inputs to a feedback loop that controls a normalized value corresponding to cooperative effort.

The output of the control system is a cooperative resource $\rho \in [0, 1]$. When $\rho = 0$, a device never provides cooperative assistance to others, while $\rho = 1$ represents the case where the device provides as much assistance as it is able to.

In the control system, we introduce a new parameter α that represents an *altruism factor*. This parameter serves as an input to the control system.

$$\alpha \in [0 \text{ —— } 1 \text{ ———— } \infty]$$

Selfless | Selfish

When $\alpha \to \infty$, the control system will disallow cooperation with others at all times. In other words, the control system will drive the output $\rho \to 0$. When $\alpha = 1$, the control system will disallow cooperation if there is *any* loss to energy efficiency. When $\alpha = 0$, the control system is effectively disabled and will not attempt to enforce a minimum energy efficiency. In other words, the control system will drive the output $\rho \to 1$. In this scenario, the device is completely selfless and willing to cooperate with others as much as possible. The α factor allows a user to scale how altruistic their device is.

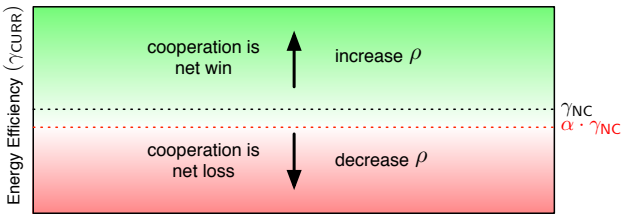

Fig. 7. Relationship between cooperative resource ρ and energy efficiency γ.

Figure 7 illustrates the goal of the control system. When cooperation is a "net win" for the device ($\gamma_{CURR} \geq \alpha\gamma_{NC}$), cooperative resources are increased. When cooperation is a "net loss" ($\gamma_{CURR} < \alpha\gamma_{NC}$), cooperative resources are decreased. The idea is that, if a device is performing better with cooperation than without, the only reason is because other devices in the network are aiding it by cooperatively sending its packets and thus saving the device resources. If devices start rolling back cooperative resources when they are "winning," the gain can disappear for everyone. It is only when a device knows that cooperation is actively harming its energy efficiency that it should lower the amount that it is willing to help others.

In determining a value for α, setting $\alpha = 1$ will not ensure that everyone is helped by cooperation; it will only ensure that no-one is hurt. If every device in a network operates under the rule that it will tolerate *no* negative deviation from non-cooperative performance, then the steady-state operation of that network can easily devolve to no device ever cooperating. As an example, imagine a general network where all cooperation is currently disabled. In other words, $\gamma_{CURR} = \gamma_{NC}$. If the devices in the network employ an altruism parameter $\alpha = 1$, then cooperation stays disabled *even if* some users in the network could benefit at no cost to others.

The α altruism parameter, when set below 1 even by a slight amount, can serve to bias the network in the direction of providing cooperative resources in case there is mutual benefit to be had. Returning to the previous example, if $\alpha = 1-\epsilon$ (where ϵ is a very small number), then the cooperative resources of devices will start to ramp up among devices where cooperation is a "net win." As devices start to see gains from cooperation they will correspondingly cooperate more, resulting in further gains.

For our implementation in Section V, we evaluate a value for α that is *prima facie* tolerable to everyone ($\alpha = .95$). An excellent avenue for the future extension of this work is the creation of policies that govern the selection of α. As an intuitive example, one might imagine that a user's ability to tolerate loss of energy efficiency may be related the remaining capacity of the battery in their mobile device. As a battery depletes its charge, a policy might govern that α increases to protect remaining energy stores.

As a control system, we employ a classic design known as a proportional-integral (PI) controller. These controllers date back to the 1890s [20] and are used heavily to this day.

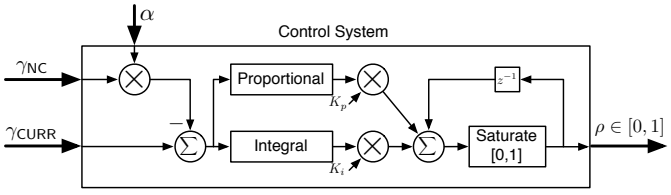

Fig. 8. Proportional-integral control loop.

Figure 8 shows the PI controller in DECC. The controller reacts to two terms:

- **Proportional:** The proportional branch of the controller reacts to a current measurement of error. The current error describes how far the system is away from some established baseline at this very moment.
- **Integral:** The integral branch of the controller reacts to the history of errors.

The two filter parameters K_p and K_i can be adjusted to set how quickly the control loop reacts to changes in its inputs as well as the damping response of the filter.

C. MAC

Given ρ (i.e., the output from the control system given, among other inputs, some α), the MAC should scale its *effort* as acting as a cooperative relay. We use the PDOC MAC from the investigation in Section III such that, when DECC lowers ρ, the MAC avoids cooperating on the most energy-expensive packets *first*. This exploits the fact that not all transmissions cost the same amount of energy.

To describe this process, we first present a detailed summary of the PDOC protocol that was originally presented in [17]. PDOC is a cooperative MAC protocol that accomplishes two tasks:

- A PDOC cooperative relay only transmits when doing so is likely to result in a successfully delivered packet.
- A PDOC cooperative relay will also adjust downward each transmission's output power such that successful decoding of that transmission by the destination is still likely.

PDOC is able to accomplish these tasks by exploiting feedback from the receiving destination node in the event of packet losses.

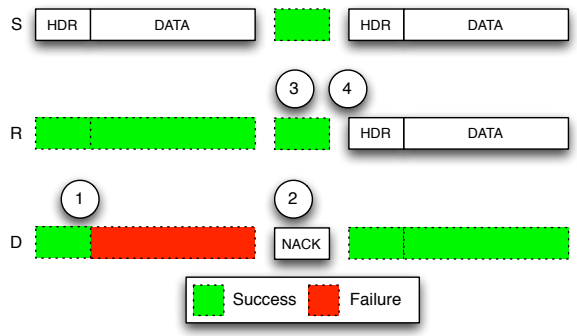

Fig. 9. PDOC protocol timeline.

Figure 9 illustrates a timeline of events that occur with the PDOC protocol. The figure shows a failed direct packet exchange between a source of traffic (S) to a destination for that traffic (D). The destination then triggers a cooperative transmission by actively sending a NACK to both the source (S) and a cooperative relay (R). The numbered circles in Figure 9 correspond to the following events:

1) The destination estimates the channel quality between source and destination when receiving the first transmission. This estimation occurs in any coherent communication system with channel-state-information-at-the-receiver (CSIR). Rather than throw away the estimate after using it to attempt to detect and decode the packet, PDOC receivers save the value into memory. For the implementation in this paper, this estimate is simply a 10-bit received signal strength indication (RSSI).
2) A failure event in decoding the payload of the source's transmission causes the destination to generate a NACK and broadcast it to both the source and relay. PDOC includes the channel quality measurement from the previously described event (event 1) inside the NACK packet.
3) Upon receiving the NACK, the relay directly estimates the channel between the relay and the destination using exactly the same CSIR mechanism as event 1. This mechanism relies on channel power reciprocity, which can be applied to time-division duplexing systems if the use of the return channel is within a channel coherence time [21]. Additionally, the mechanism reads the source-destination channel estimate out of the NACK given by the destination, giving it access to both the source-destination and relay-destination channels.
4) Finally, the relay decides whether or not to transmit and at what power level. This determination is made as a function of the SD and RD channel values. For our implementation, this function is determined via calibration measurements that map SD and RD channel powers to packet error rate. The details of these measurements are available in our prior publication [17].

With PDOC, relay transmissions powers will occur in a range spanning up to the maximum allowed transmission power. In the prior work, this maximum value was the maximum transmission power through the radio. In this work, we recognize that this maximum transmission power can be lowered by the ρ cooperative resource parameter from the control system. As ρ falls and the control system demands fewer cooperative transmissions, the PDOC MAC is instructed to allow maximum transmission powers at decreasing values.

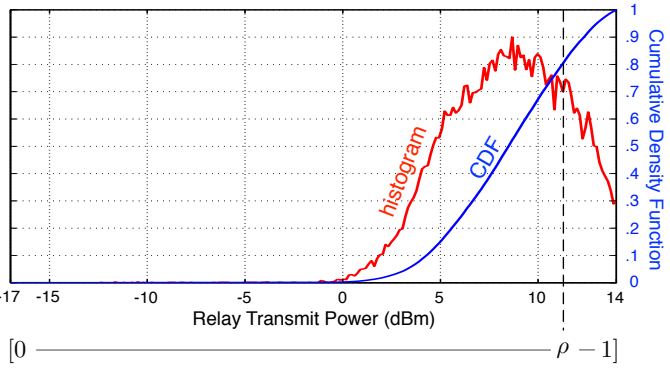

Fig. 10. Using ρ to cap relay transmission power.

Figure 10 shows data from the MS2 device from the experiment in Section III-B3. This data shows the PDOC protocol in action. Over the range of possible transmission powers (-17 to $+14$ dBm), MS2 selects a transmission power such that the cooperative transmission is likely to succeed. Notice that the majority of transmissions do not occur at maximum transmission power. By using the ρ input from DECC, the MAC caps its maximum relay transmission power and ensures that the cooperative exchanges that are skipped are the ones that draw the most power from the device. As illustrated in Figure 10, a $\rho \approx .79$ would eliminate 20% of relay transmissions. The eliminated relay transmissions would be the highest power and therefore the most costly in terms of energy efficiency.

The MAC can react to *instantaneous* changes in channels; even with a small ρ, the MAC may still opt to act as a

cooperative relay if the necessary power cost for that action is sufficiently small.

V. Evaluation Using Hardware Implementations

In this section, we evaluate DECC by presenting a custom real-time prototype on the WARP testbed. We then revisit the two topologies from Sections III-B2 and III-B3 and characterize DECC's performance. Finally, we use an implementation of DECC in a custom simulator to evaluate larger network scenarios (since our hardware testbed size is limited).

A. Implementation Details

We have implemented DECC on WARP entirely in the built-in PowerPC without any changes to the FPGA fabric needed to implement our prior cooperative PHY [4], [16] and MAC [17].

The MAC keeps a cumulative counter of the number of bits successfully transmitted (i.e. ACKed bits) and the number of bits successfully received. The MAC also keeps a cumulative measurement of the amount of energy spent according to the description provided in Appendix A. These values are provided to the DECC software module via a periodic 1 Hz timer. Every second, the DECC software module calculates the energy efficiency of the previous second of traffic and feeds this into the proportional path of the PI control system. Additionally, it updates a cumulative energy efficiency calculation over the entire runtime of the experiment. This cumulative energy efficiency calculation forms the integral path of the PI control system. The ρ cooperative resource is updated at the conclusion of each timer callback.

As described in Section IV-C, the MAC layer uses the value of ρ to decide whether or not any given potential relay packet should actually be transmitted. Specifically, ρ is used as a scale factor on the maximum allowed transmission power. With the WARP radios and OFDM PHY, our maximum possible transmission power is 14 dBm, so the maximum allowed transmission power is simply

$$P_T \leq \rho \cdot 14 \text{ dBm}. \quad (1)$$

Any potential relayed packet that exceeds this power threshold is cancelled and nothing is transmitted. We use our WARPnet Measurement Framework [22]–[24] to allow each device to send packet traces over Ethernet to a host computer for analysis.

For simplicity of visualization we have first opted for a "one-shot" non-cooperative mode performed at the beginning of the experiment for the devices to measure γ_{NC}, the energy efficiency of a non-cooperative link. We then demonstrate the full asynchronous mode-switching functionality of DECC's measurement module in Section V-B4.

Table I shows the DECC parameters used for the experiments in this section. We do not claim optimality in the selection of these parameters. In fact, we have intentionally left the control loop slightly underdamped as it makes the visualization of the input-output relationship of the DECC control system easier to discuss.

Figure 11 shows an example timeline of events taken from an oscilloscope capture of a trial. One can see MS1 and MS2 transmitting their own data packets as well as a cooperative transmission from MS2 to assist MS1 for a packet that was NACKed. In order for this cooperative transmission to have occurred, the required relay transmission power from MS2 was less than $\rho \cdot 14$ dBm.

TABLE I
DECC Parameters

K_p	10^{-8}
K_i	10^{-7}
α	$0, .95$
Direct Mode Duration	45 s
Trial Duration	500 s

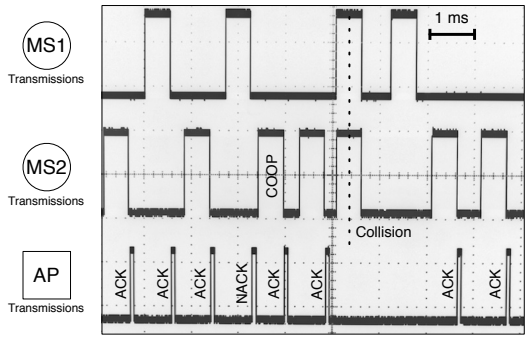

Fig. 11. Oscilloscope-captured timeline of events.

B. Experimental Evaluation

In Section III, we demonstrated that the gains or losses in energy efficiency due to signal-scale cooperation are severely topologically dependent. First, in Sections V-B1 and V-B2, we revisit two of the same topologies and show that DECC can (a) preserve mutual cooperative energy efficiency gains when they occur and (b) protect a device from losses in energy efficiency due to cooperation.

Second, in Section V-B3, we investigate the DECC parameters that yielded the experimental results from the prior sections. Specifically, we empirically study the amount of time DECC needs to spend in a non-cooperative mode such that the measurement of γ_{NC} in the non-cooperative phase is representative of what γ_{NC} would be during the cooperative phase.

Finally, in Section V-B4, we demonstrate the full asynchronous mode-switching capability of DECC.

1) Revisiting Scenario A: Symmetric Topology: In Section III-B2, we observed that devices in a symmetric topology can provide equal amounts of cooperative assistance to one another and, therefore, provide mutual benefit. With DECC activated, devices should be able to recognize that they are each being aided by cooperation and do nothing to reduce cooperative resources; DECC should set $\rho = 1$ indefinitely.

Figure 12(a) shows a 500 second timeline of a trial for both devices (MS1 and MS2). During the first 45 seconds, both MS1 and MS2 refuse to act as cooperative relays for one another. During this period, γ_{NC} (and therefore, $\alpha\gamma_{\mathsf{NC}}$) are measured. For this trial we set $\alpha = .95$ so that devices would

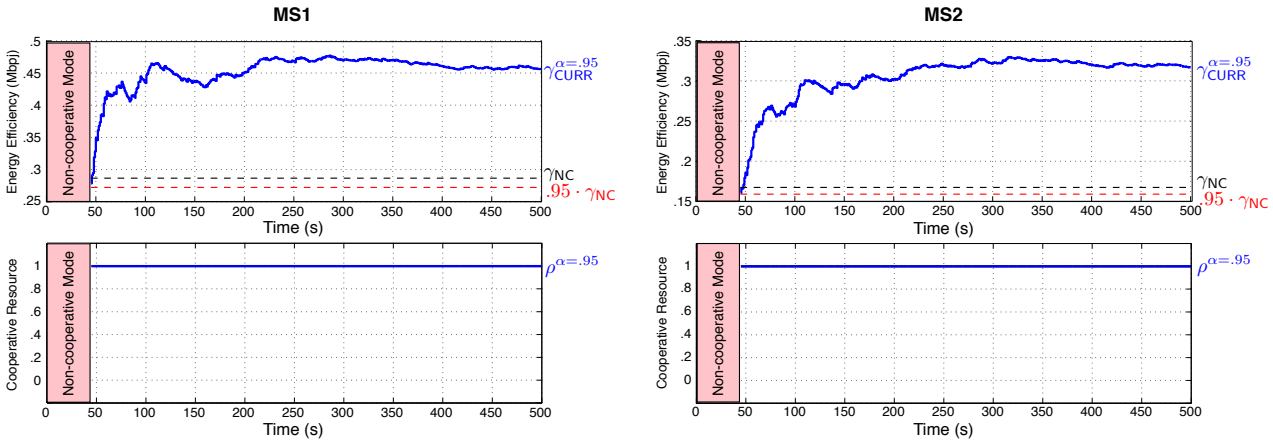

(a) Scenario A: Symmetric Topology.

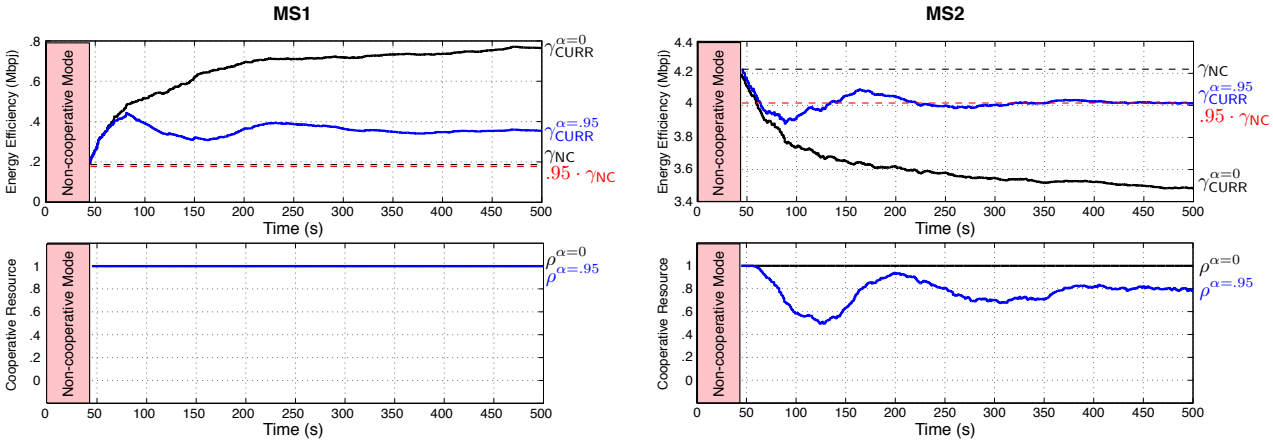

(b) Scenario B: Asymmetric Topology.

Fig. 12. Energy efficiency with DECC.

be willing to tolerate as much as a 5% reduction in energy efficiency. Because of the symmetric nature of the topology, both devices perform significantly better in the cooperative mode versus the non-cooperative mode. Throughout the entire experiment, DECC leaves the cooperative resource $\rho = 1$ because both devices are better off with cooperation than without it. The $50\%+$ improvement in energy efficiency that was measured in Section III-B2 is retained.

2) Revisiting Scenario B: Asymmetric Topology: In Section III-B3, the asymmetric topology caused MS1's energy efficiency to improve by over 320% and MS2's energy efficiency to degrade by 18%. With DECC activated, we expect MS2 to reduce the amount of effort it puts into cooperation in order to halt this loss.

Figure 12(b) shows a timeline for both devices in the asymmetric topology. For reference, we have included a trial for $\alpha = 0$ where the DECC control loop is effectively disabled. In this case, we can see that MS2 maintains a $\rho = 1$ despite performing significantly worse with cooperation than without; MS2 is completely selfless for $\alpha = 0$.

For $\alpha = .95$, however, we can see that MS2 lowers the cooperative resource ρ once γ_{CURR} falls below $.95\gamma_{\mathsf{NC}}$. When this happens, γ_{CURR} at MS2 begins to increase back towards the non-cooperative performance while γ_{CURR} at MS1 begins to decrease since fewer cooperative resources are being provided by MS2. After some underdamped oscillations around the $.95\gamma_{\mathsf{NC}}$ line, DECC's selection of ρ ensures that MS2 suffers no more than $1 - \alpha = 5\%$ loss of energy efficiency rather than the 18% loss suffered without DECC.

Despite the reduction in cooperative resources provided by MS2 when DECC is activated, MS1 still sees nearly a $2\times$ improvement in energy efficiency. Substantial cooperative gains can be achieved in return for small, bounded losses in efficiency for a subset of nodes even in severely uneven scenarios.

3) Measurement Deviation Analysis: In the previous experiment, we chose a non-cooperative duration of 45 seconds and witnessed that the DECC control system at MS2 converged on a 5% loss in energy efficiency. In this section, we investigate this selection of duration.

Figure 13 illustrates a mode switch performed by the measurement module of DECC. During the T1 period, DECC measures the energy efficiency of a non-cooperative mode, γ_{NC}. When cooperation is allowed during the rest of T2, DECC implicitly relies on the measurement of γ_{NC} during T1 being representative of what γ_{NC} would be *if it were* measured throughout T2. To investigate the durations of T1 and T2 required for such reliance to be appropriate, we

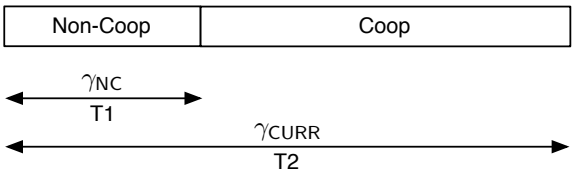

Fig. 13. Ideally, the measurement of γ_{NC} during the T1 period should be representative of what γ_{NC} would be if measured throughout the T2 period.

construct an experiment to measure the percentage change in energy efficiency measurements between periods T1 and T2. Formally, we measure the deviation[4]

$$d = \frac{|\gamma_{\text{T1}} - \gamma_{\text{T2}}|}{\gamma_{\text{T1}}}, \tag{2}$$

where γ_{T1} is the energy efficiency of the T1 period and γ_{T2} is the energy efficiency of the T2 period.

The goal of the following experiment is to determine the deviation between the two periods if non-cooperative links were used in both. As such, we use our non-cooperative CSMA/CA implementation on WARP. We test the asymmetric topology described in Section III-B3. Furthermore, for ease of visualization, we consider the case where $\text{T2} = 4 \cdot \text{T1}$. In other words, we consider the $\frac{1}{4}$ duty cycle case where, for example, a measurement in one second is expected to be valid for the next three seconds.

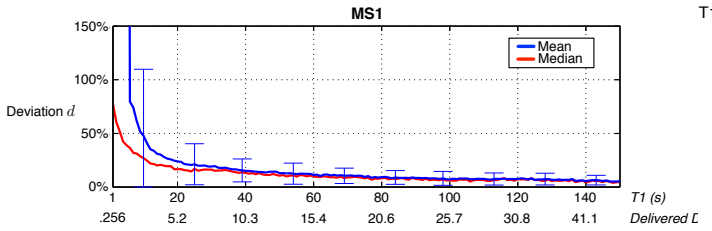

Fig. 14. Because the energy efficiency of MS1 is so small, the percentage deviation can be quite large absent substantial measurement time.

Figure 14 shows the percent deviation between T1 and T2 for MS1. Recall that MS1 is the device that is a long distance away from the AP. In any given second, the achieved throughput of the device only ensures that, on average, 256 kilobits are successfully delivered. With the 1300 byte payloads being sent, this is an average of only just over 24 successful packet transmissions per second. As such, the measurement from one second to the next can deviate significantly. The 45 seconds of non-cooperative measurement in Section V-B2 resulted in a deviation of approximately 20% on average, which is far below the nearly 100% improvement measured with DECC set to $\alpha = .95$ and even further below the 320% improvement seen with $\alpha = 0$. As such, DECC is unlikely to confuse cooperative gain for cooperative loss despite the difficulty in measuring energy efficiency.

Figure 15 shows the percent deviation between T1 and T2 for MS2. Recall that MS2 was the device where the 5% loss

[4]We note that the metric d in Equation 2 can be dominated by outliers if γ_{T1} is very small or zero. As such, we provide the median across all trials as well as the mean and standard deviation error bars.

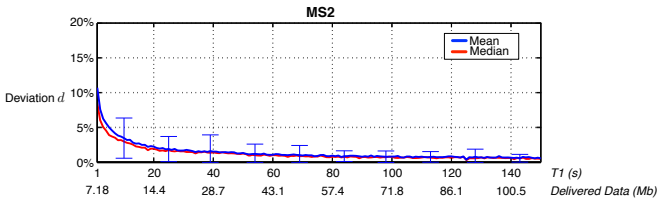

Fig. 15. Because MS2 delivers far more data per unit time than MS1, its percentage deviation between T1 and T2 is far superior to that of MS1.

(for $\alpha = .95$) was actually a binding constraint. As such, MS2's ability to measure γ_{NC} to well within 5% is paramount. With the 45 seconds of non-cooperative measurements, MS2 measures γ_{NC} to within an approximately 2% deviation. Of course, if MS2 is less strict about how bounded its loss is (e.g. $\alpha = .85$ or lower), then even shorter non-cooperative durations can be tolerated since the deviation between measurements becomes small relative to the bound.

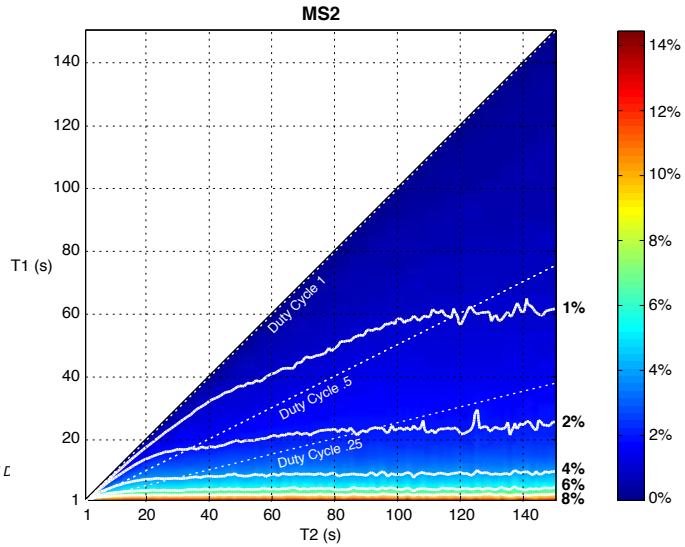

Fig. 16. MS2 measurement deviation is represented in color and in contour lines.

In the previous analysis, a duty cycle of $\frac{1}{4}$ was assumed. To characterize the impact of duty cycle selection on energy efficiency measurements, we generalize our study in Figure 16. In this figure, measurement deviation between T1 and T2 is represented by color and contour lines along the [1%, 2%, 4%, 6%, 8%] ridges. When $\text{T1} = \text{T2}$, the measurement deviation between T1 and T2 is zero because the measurements occur over exactly the same regions of time. Of course, a 100% duty cycle also means that cooperation can never be enabled to take advantage of the perfectly accurate measurement. The data in Figure 15 is, in fact, a cross section of Figure 16 along the line labelled "Duty Cycle .25."

We make no claims as to the universality of these specific T1 and T2 durations for general networks. That said, the topology studied is very likely to occur in general networks: some users are much closer to base stations or access points than others. The deviation from one measurement to the next is highly dependent on channel coherence times, traffic

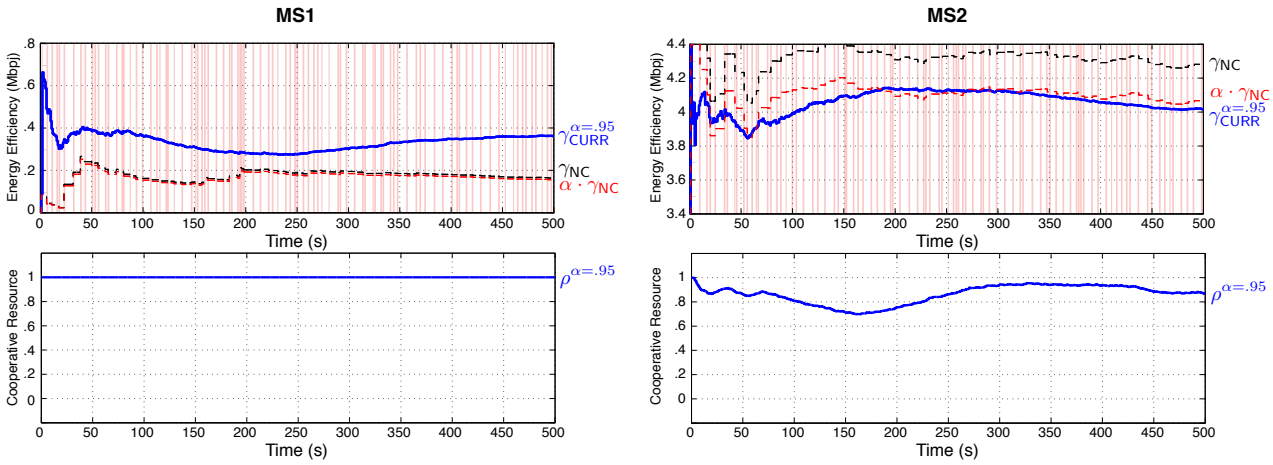

Fig. 17. With DECC, each device can asynchronously switch between cooperative and non-cooperative modes.

patterns, device mobility, and network size. However, the implementation presented in this work can be used for a future user study where large, mobile networks with actual user-generated traffic patterns affect these measurements. In this experiment, MS1 observed considerably higher measurement deviation between T1 and T2 than MS2 because MS1 delivers far fewer bits than MS2 in the same amount of time. Consequently, for DECC, devices that stand to lose energy efficiency due to cooperation will tend to be able to measure non-cooperative energy efficiency more easily. This tendency is well-aligned with DECC's structure – by design, DECC treats *any* improvement in efficiency with cooperation as a reason to cooperate more. It is only when the device's energy efficiency is degraded by cooperation that accurate measurement of non-cooperative energy efficiency is very important. Said another way, the higher the deviation between measurements, the less such deviation matters for the operation of DECC.

4) Asynchronous Switching: As noted earlier, the previous DECC evaluation hinged on a single beginning-of-time non-cooperative measurement period of 45 seconds. In this section, we demonstrate that the asynchronous switching functionality of DECC specified in Section IV-A holds. This functionality allows a device to effectively spread non-cooperative training over a longer window to allow the control system to begin adapting even if the underlying measurements are noisy. In this experiment, we consider a $1/4$ duty cycle where, on average, one second of non-cooperative measurement supports the next three seconds where cooperation is allowed. To demonstrate that mode switches are not coordinated between MS1 and MS2, we allow each device to enter a one second non-cooperative mode followed by a random cooperative duration chosen with a uniform random variable between zero and six seconds. MS1 and MS2 separately decide to disable cooperation and their decisions need not align.

In Figure 17, vertical red bars represent non-cooperative periods for MS1 and MS2. We can see that, over time, γ_{NC} is refined as cumulatively more time is spent in a non-cooperative mode. By 180 seconds into the experiment, MS2 has spent approximately $\frac{180}{4} = 45$ seconds measuring non-cooperative energy efficiency – equivalent to the "one-shot" experiments in Sections V-B1 and V-B2. We can see that the steady-state behavior of the asynchronously switched DECC is similar to the behavior observed in Figure 12(b). The only difference is the DECC control system must track a moving target of a continually updating measurement of γ_{NC} rather than a static measurement – a task that DECC's PI control system is adept at handling.

C. Larger Network Simulation

The previously described experiments present a targeted investigation of atomic topological relationships. In this section, we evaluate the scalability of DECC by considering a larger network. In this experiment, each device randomly chooses to enter its non-cooperative test mode independently from the other devices in the network. Furthermore, to demonstrate the distributed nature of DECC, we allow any device in the network to provide cooperative relaying assistance to any other device. Specifically, if any device (a) decodes the original transmission from the source and (b) decodes the NACK packet from the AP, it can act as a cooperative relay on that particular packet. If more than one relay transmits on a given packet, we assume they will not interfere with one another as they can employ conjugate beamforming to align the phase of their transmission to the channel estimate garnered from the reception of the NACK.[5]

The channel emulator we have employed in earlier sections of this work is a 4×4 MIMO channel emulator and lacks a sufficient number of emulated paths to test networks larger than three devices. Instead, we have opted to implement DECC in the well-known ns-2 network simulator [7].

Figure 18 shows the topology under consideration. In Section III-C, we established regimes of mutual benefit and one-sided benefit. In this simulation we consider a variety of devices at different distances away from an AP in order

[5]We recognize that this kind of beamforming relies on channel reciprocity between the reception of the NACK and transmission of the cooperative packet. Until recently, phase reciprocity was assumed to be impossible in Wi-Fi-like radios that switch between transmission and reception modes. However, modern (and ongoing) work has demonstrated that offsets between transmissions and receptions can be calibrated away [25].

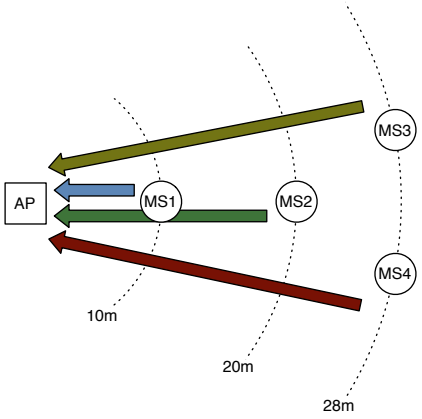

Fig. 18. Network topology.

to sample points in these various regimes. The purpose of this topology is to showcase scenarios where devices may be simultaneously helped by some devices in the network and harmed by others. For example, from our study so far we expect that MS2 will be able to derive benefit from MS1. However, the existence of MS3 and MS4 on the fringes of the network will impact those gains in until-now unknown ways. In this simulation, we perform the same asynchronous mode switch of one second of non-cooperative mode to an average of three seconds of cooperation. Furthermore, to allow devices' measurements of γ_{NC} to stabilize, we have performed the simulation over a very long time of 2000 seconds. By the end of this simulation timeline, virtually no deviation in the cumulative measurement of γ_{NC} from one update to the next is observed.

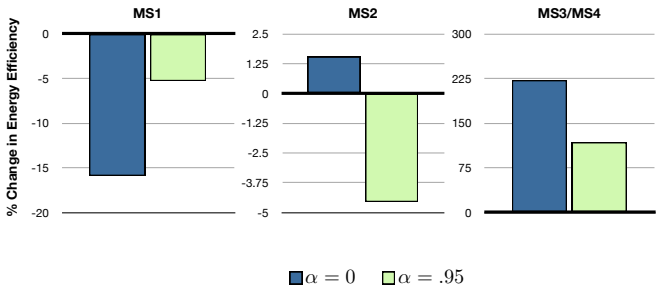

Fig. 19. Results of larger network simulation.

Figure 19 shows the percentage change of energy efficiency for each device in the network. Specifically, we plot $(\gamma_{\mathsf{CURR}} - \gamma_{\mathsf{NC}})/\gamma_{\mathsf{NC}}$ for every device. Furthermore, we consider the case of completely selfless devices ($\alpha = 0$) and the case of devices employing DECC for bounded altruism ($\alpha = .95$).

When $\alpha = 0$, MS1 acts as a cooperative relay completely selflessly. Since it is in a location that can benefit all of the other nodes, MS1 sees a significant degradation in energy efficiency of 15%. By enabling DECC with $\alpha = .95$, that loss is reduced to the specified 5%. This is consistent with our earlier investigations in Section V-B.

MS2, on the other hand, shows a surprising implication of DECC in large networks. When every device acts completely selflessly with $\alpha = 0$, the gains to MS2 from MS1 are enough to offset the costs of helping MS3 and MS4. Overall, MS2 sees a modest increase in energy efficiency. However, when $\alpha = .95$, the amount of gain that MS1 provides MS2 is reduced accordingly. This decrease in improvement for MS2 means that MS2 is now more burdened with acting as a cooperative relay; MS2 starts to suffer an energy efficiency loss. However, since MS2 is also implementing DECC, DECC ensures that MS2 does not suffer more than 5%. In this way, DECC shifts the burden of cooperation off of one user and distributes it to others.

The net effect for MS3 and MS4 is that DECC will lower the amount of improvement they see, but the gains are still considerable ($2\times$ the energy efficiency) considering the amelioration of harm on the other devices in the network.

VI. Conclusion

In a general network, the mobility of devices means that sometimes they may be close to an AP and subject to potential losses in energy efficiency due to cooperation. Other times, the devices will be on the outskirts of the network and enjoy tremendous improvements in energy efficiency. With DECC, we have created a way for devices to bound their potential losses to a tolerable level and still allow the achievement of large gains elsewhere in the network.

Appendix A
MRC Emulation Calibration

The cooperative PHY from [4] employs a 2×1 Alamouti STBC to achieve spatial diversity. For architectural simplification, the PHY ignores the first "broadcast" time slot of cooperation and requires the source to retransmit its message alongside the relay during the second time slot. Equivalently, the PHY could combine the source's transmission from the first broadcast phase with a relay-only transmission during the relay phase via a technique known as maximal-ratio combining (MRC) and save the source from having to retransmit again. Adding MRC capability to the PHY would require re-architecting beyond the scope of this work. Instead, we use the 2×1 Alamouti mode of the PHY to characterize the performance of an equivalent MRC-capable physical layer. We then use these calibration values to emulate its performance on the current PHY with relay-only cooperative transmissions.

Figure 20 shows the packet error rate (PER) of the destination node as a function of the powers of the source-destination and relay-destination links. The PER characterization was created using the same 16-QAM modulation used throughout the rest of this work.

In our emulated MRC implementation, when a device sends a NACK packet it saves the received signal strength (RSSI) of the data packet it just received. Then, after receiving the RSSI of a relay-only transmission, the destination draws a random number according to a Bernoulli distribution whose parameter is derived from Figure 20. If that random draw is one, the packet is discarded. If the draw is zero, the packet is deemed successful. In this way, we encode the physical layer with the

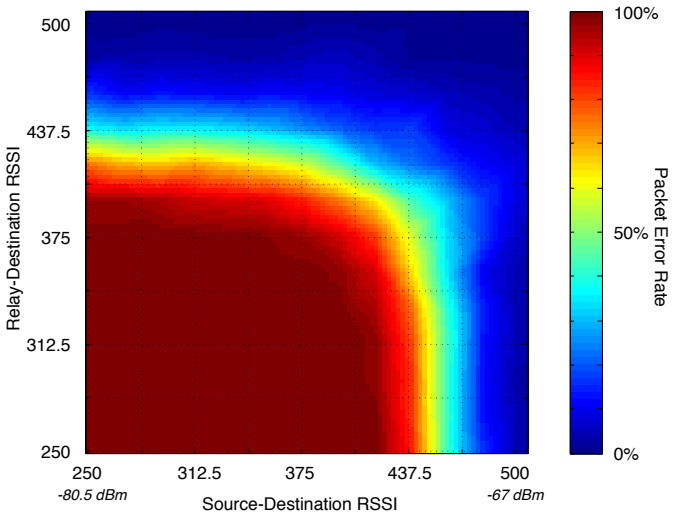

Fig. 20. PER of 2×1 Alamouti PHY.

performance of an MRC mode without making the physical layer actually perform the task.

APPENDIX B
RADIO POWER CONSUMPTION

The WARP hardware used in this work does not allow a direct measurement of power consumption at the granularity required for this work (i.e. power during receptions and transmissions at different output power levels). Instead, we use a method based on offline power consumption calculations. We used the WARP Radio Board v1.4 for all experiments in this work. There are four main components that are responsible for the bulk of the radio board's power consumption:

Digital-to-Analog Converter (DAC) and Analog-to-Digital (ADC): are responsible for converting baseband digital I and Q samples from the FPGA into analog waveforms and vice versa. Because of latencies involved with waking these devices from sleep, current WARP designs leave both the DAC and the ADC running at all times.

Transceiver (TRA): is responsible for upconverting to RF and downconverting to baseband during transmission and reception, respectively. Additionally, the transceiver contains a transmission variable gain amplifier (VGA) that covers a 32 dB range. The power draw from the transceiver depends on the output power of this VGA.

Power Amplifier (PA): The radio board has a power amplifier (PA) after the transceiver for boosting transmissions by another 32 dB. The power consumption of the power amplifier is a function of the power of the input signal from the transceiver. During reception, the PA is deactivated and its current draw is negligible.

[5]The WARP Radio Board uses a Sharp IRM046U7 power amplifier. Since this part does not provide supply current specifications as a function of transmission power, we have opted to substitute an equivalent power amplifier for this analysis [26].

TABLE II
WARP RADIO POWER CONSUMPTION

	Tx	Rx
DAC [27]	.418W	.418W
ADC [28]	.713W	.713W
Transceiver [29]	Equation (3)	.354W
PA[6]	Equation (4)	-

Table II summarizes the power consumption of the primary components on the WARP Radio Board. For the two parts that depend on transmission power (the transceiver and the PA), we have used their respective data sheets to perform a simple exponential curve fit. Their expressions are

$$P_D^{\text{PA}} = .0034e^{.2072P_T} + .2162e^{.0275P_T} \quad (3)$$

$$P_D^{\text{TRA}} = .2482e^{-.0084P_T} + .0922e^{.0264P_T}, \quad (4)$$

where P_T is the transmission output power (in dBm) and P_D^{PA}, P_D^{TRA} are the power dissipated by the PA and transceiver (in Watts), respectively.

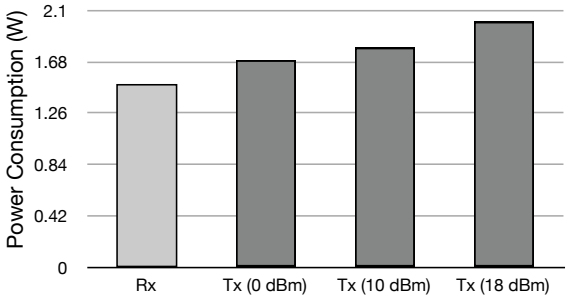

Fig. 21. Power consumption comparison of reception and transmission at a variety of output powers.

Figure 21 shows a comparison of reception and transmission power consumption.

In our analysis, we consider only the power consumption of the radio board since it would likely be the dominant energy sink in a final implementation that has been reduced to custom silicon. That said, our approach is general and can easily be extended to consider baseband processing costs as well as the energy costs of higher layer components such as the transport layer and applications. Updating DECC for additional power measurements would only require updating the offline power calculations. The architecture of the DECC implementation would remain unchanged.

REFERENCES

[1] D. Goldman, "Sorry, America: Your wireless airwaves are full." [Online]. Available: http://money.cnn.com/2012/02/21/technology/spectrum_crunch/?npt=NP1&hpt=hp_c2

[2] J. Laneman, D. Tse, and G. Wornell, "Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behavior," *IEEE Transactions on Information Theory*, vol. 50, no. 12, 2004.

[3] G. Kramer, I. Marić, and R. Yates, "Cooperative Communications," *Foundations and Trends in Networking*, vol. 1, no. 3, 2006.

[4] P. Murphy and A. Sabharwal, "Design, Implementation, and Characterization of a Cooperative Communications System," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 6, 2011.

[5] C. Hunter, "Distributed protocols for signal-scale cooperation," Ph.D. dissertation, Rice University, 2012. [Online]. Available: http://warp.rice.edu/w/HunterPhDThesis

[6] "Rice University WARP Project." [Online]. Available: http://warp.rice.edu/

[7] "The Network Simulator - ns-2." [Online]. Available: http://nsnam.isi.edu/nsnam/index.php/User_Information

[8] B. Wang, Z. Han, and K. Liu, "Distributed Relay Selection and Power Control for Multiuser Cooperative Communication Networks Using Buyer/Seller Game," in *Proceedings of IEEE INFOCOM*, 2007.

[9] V. Srinivasan, P. Nuggehalli, C. Chiasserini, and R. Rao, "Cooperation in Wireless Ad hoc Networks," in *Proceedings of IEEE INFOCOM*, 2003.

[10] M. Nokleby and B. Aazhang, "User cooperation for energy-efficient cellular communications," in *IEEE International Conference on Communications*, 2010.

[11] Z. Han, D. Niyato, W. Saad, T. Başar, and A. Hjørungnes, *Game Theory in Wireless and Communication Networks: Theory, Models, and Applications*. Cambridge University Press, 2011.

[12] L. Simic, S. M. Berber, and K. W. Sowerby, "Partner choice and power allocation for energy efficient cooperation in wireless sensor networks," in *IEEE International Conference on Communications*, 2008.

[13] R. Madan, N. Mehta, A. Molisch, and J. Zhang, "Energy-efficient cooperative relaying over fading channels with simple relay selection," *IEEE Transactions on Wireless Communications*, 2008.

[14] Z. Zhou, S. Zhou, J. Cui, and S. Cui, "Energy-Efficient Cooperative Communication Based on Power Control and Selective Single-Relay in Wireless Sensor Networks," *IEEE Transactions on Wireless Communications*, 2008.

[15] A. Ibrahim, Z. Han, and K. R. Liu, "Distributed energy-efficient cooperative routing in wireless networks," *IEEE Transactions on Wireless Communications*, 2008.

[16] P. Murphy, C. Hunter, and A. Sabharwal, "Design of a Cooperative OFDM Transceiver," in *Proceedings of IEEE Asilomar Conference on Signals, Systems, and Computers*, 2009.

[17] C. Hunter, M. Kanga, L. Zhong, and A. Sabharwal, "On-demand Cooperation with Power Control: Protocol and Experimental Results," in *Proceedings of IEEE Asilomar Conference on Signals, Systems, and Computers*, 2011.

[18] C. Hunter, P. Murphy, and A. Sabharwal, "Real-time testbed implementation of a distributed cooperative MAC and PHY," in *Proceedings of CISS*, 2010.

[19] "Azimuth Systems ACE 400WB." [Online]. Available: http://www.azimuthsystems.com/products/ace-channel-emulators/ace-400-wb-for-wifi/

[20] S. Bennett, *A History of Control Engineering, 1930-1955*. Peter Peregrinus Ltd, 1993, vol. 47.

[21] A. F. Molisch, *Wireless communications*. John Wiley & Sons, 2007.

[22] S. Gupta, C. Hunter, P. Murphy, and A. Sabharwal, "WARPnet: Clean Slate Research on Deployed Wireless Networks," in *Proceedings of ACM MobiHoc*, 2009.

[23] S. Gupta, P. Murphy, C. Hunter, and A. Sabharwal, *Cognitive Radio: System Design Perspective*. Springer, 2010, ch. WARPnet: A Platform for Deployed Cognitive Radio Experiments.

[24] "Rice University WARPnet Framework." [Online]. Available: http://warp.rice.edu/trac/wiki/WARPnet

[25] F. Kaltenberger, H. Jiang, M. Guillaud, and R. Knopp, "Relative Channel Reciprocity Calibration in MIMO/TDD systems," in *Proceedings of IEEE Future Network and Mobile Summit*, 2010.

[26] "Anadigics AWL6950 Datasheet." [Online]. Available: http://www.anadigics.com/products/view/awl6950

[27] "Analog Devices AD9777 Datasheet." [Online]. Available: http://www.analog.com/en/digital-to-analog-converters/da-converters/ad9777/products/product.html

[28] "Analog Devices AD9248/BCP-65 Datasheet." [Online]. Available: http://www.analog.com/en/analog-to-digital-converters/ad-converters/ad9248/products/product.html

[29] "Maxim MAX2829 Datasheet." [Online]. Available: http://www.maxim-ic.com/datasheet/index.mvp/id/4532

Christopher Hunter (S'04-M'12) received the B.S., M.S., and Ph.D. degrees from Rice University, Houston, TX, in 2006, 2008, and 2012 respectively. He is currently with Mango Communications in Houston, TX. His research interests include cooperative communications and the implementation of novel wireless systems.

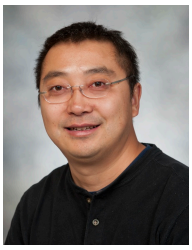

Lin Zhong (S'02-M'05-SM'12) is an Associate Professor at the Department of Electrical and Computer Engineering, Rice University. Prior to joining Rice, he received his Ph.D. from Princeton University in 2005, and his M.S. and B.S. from Tsinghua University in 2000 and 1998, respectively.

He is a recipient of the National Science Foundation CAREER award and the best paper awards from ACM MobiSys 2011, IEEE PerCom 2009, and ACM MobileHCI 2007. His research interests include mobile commuting, human-computer interaction, and nanoelectronics.

Ashutosh Sabharwal (S'91-M'99-SM'04) received the B. Tech. degrees from the Indian Institute of Technology, New Delhi, India, in 1993 and the M.S. and Ph.D. degrees from The Ohio State University, Columbus, in 1995 and 1999, respectively.

He is currently a Professor with the Department of Electrical and Computer Engineering, Rice University, Houston, TX. His research interests include information theory and communication algorithms for wireless systems.

Dr. Sabharwal received the 1998 Presidential Dissertation Fellowship Award.