Realistic Evaluation of Cooperative Relaying Networks Using Decentralized Distributed Space-Time Block Coding

Justin Yackoski†, Lu Zhang*, Bo Gui*, Chien-Chung Shen†, and Leonard J. Cimini, Jr.*

University of Delaware, Newark, DE 19716, U.S.A.
Email: yackoski,lzhang,bgui,cshen,cimini@eecis.udel.edu

***Abstract*—Information theoretic studies have shown the significant performance improvements of cooperative communications. However, these studies ignore both the overheads incurred in real implementations of the cooperative techniques at the physical layer and their interactions with higher layer protocols in a networking context. In this paper, we study the performance of realistic networking scenarios facilitated by cooperation by taking overheads incurred at the physical, MAC, and network layers into account. In particular, (1) we modify the physical layer model of the QualNet network simulator to incorporate decentralized distributed space-time block coding into all SINR calculations and to combine signals transmitted concurrently from multiple relays, (2) we implement a path-centric MAC protocol to both reserve a multihop path between source and destination nodes and coordination relay nodes, and (3) we modify the DSR protocol to support path reservation at the network layer. Preliminary simulation results demonstrate that significant performance improvement can be achieved by employing cooperation. We also demonstrate the overheads which challenge their effectiveness in real networks.**

I. INTRODUCTION

The notion of cooperative communications fundamentally changes the abstraction of a wireless link and promises improved performance [1-2]. However, much of the existing research focuses on information theoretic studies of the physical layer capabilities, which ignore both the overheads that may be incurred in realistic networking contexts and the interactions with higher layer protocols.

In contrast to single-node transmissions, cooperative communications may involve multiple relay nodes transmitting simultaneously to a receiver such as in space-time coded cooperation. To facilitate such simultaneous transmissions and receptions, many different overheads are incurred, some of which are not well understood, and are often ignored. In the physical layer, these overheads include those required for synchronization between the transmitter(s) and receiver(s), channel estimation to obtain the necessary channel state information (CSI), and power allocation. Additional overheads are incurred by incorporating cooperative communications as a link abstraction into a networking context with multiple source-destination pairs and multihop relaying. For instance, relay selection and coordination will require control signaling at the MAC layer among potential relay nodes, and performance optimization demands joint optimization of routing and relay selection. Without a holistic evaluation within an end-to-end networking context, the true potential of cooperative communications cannot be evaluated.

The objective of our work is to investigate the performance, and hence the benefits, of incorporating cooperative communications in realistic networking scenarios by taking most of the necessary overheads into account. Here, we consider only a few sources of overhead and focus on a cooperative network using *decentralized* Distributed Space-Time Block Coding (Dis-STBC). Specifically, we modify the physical layer model of the QualNet network simulator **(1)** to incorporate Dis-STBC into all signal-to-interference-plus-noise ratio (SINR) calculations and **(2)** to combine signals concurrently transmitted from multiple relays. At the MAC layer, we implement a path-centric scheme which reserves a multihop path between source and destination nodes, facilitates relay node selection, and coordinates cooperative tranamissions. At the network layer, we modify the standard Dynamic Source Routing (DSR) protocol so that the source route is provided to the MAC for path reservation. Networking performance is then evaluated for networking scenarios with and without cooperative communications.

Recently, there has been some work that explores the benefits of cooperative communications in the networking context at different protocol layers. The work in [3-4] explores a cross-layer framework where centralized Dis-STBC is used at the physical layer. In addition, the proposed MAC adopts a hop-by-hop approach, in contrast to the path-centric approach proposed here, a significant difference we will further describe in Section III.C. The work in [5] explores cooperative communication schemes at the MAC layer, which, however, is in sharp contrast to cooperative communications at the physical layer as investigated in our work.

The reminder of the paper is organized as follows. In the next section, we first review cooperative communications and describe the networking scenarios to be studied. In Section III, we enumerate many of the potential overheads, from

This material is based on research sponsored by the Air Force Research Laboratory*, under agreement number FA9550-06-1-0077, and by the National Science Foundation†, under grants CNS-0347460 and CNS-0721361. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

978-1-4244-2052-0/08/$25.00 ©2008 IEEE

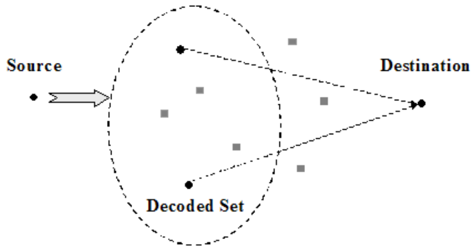

Fig. 1. A single-source two-hop selective DF cooperative network

the physical layer to the network layer, that may arise in a networking context. Section IV describes the simulation environment and discusses the simulation results. Section V concludes the paper with future research directions.

II. Cooperative Networking Scenarios

In this section, we will describe three typical scenarios for a cooperative network which employs selective decode-and-forward (DF) relaying: single-source two-hop, single-source multihop, and multi-source multihop. In particular, we focus on a network which has several randomly distributed single-antenna half-duplex nodes. The wireless channel includes the effects of path loss and quasi-static flat Rayleigh fading. Several specific cooperative techniques have been proposed for use with selective DF relaying (for example, see [1-2, 6-8]). To achieve a good balance between performance and overhead, we focus on decentralized Dis-STBC [9,10] which will be described in the next subsection.

A. *Single-Source, Two-Hop Cooperative Network*

A single-source, two-hop relaying scenario, as depicted in Fig. 1, has been used in most proposed cooperative systems. In the first hop, the source transmits and all of the other nodes listen. In the second hop, the relays cooperate to retransmit the source message to the destination. In particular, consider a network with M uniformly distributed nodes. When a source-destination (s,d) pair is active, all of the remaining $M-2$ nodes can serve as potential relays. In selective DF relaying, a potential relay node is called a *decoded node* if it can correctly decode the source message. The decoded set is defined as the set of N ($N \leq M-2$) decoded nodes. K ($K \leq N$) decoded nodes are then selected to relay the source message.

In [6], a Dis-STBC was proposed in which each relay transmits one unique column of the underlying STBC matrix. To enable each relay to know which column to transmit, most Dis-STBC schemes [6,11-12] require a central control unit or full inter-node negotiations, which incurs significant overhead. In contrast, several *decentralized* Dis-STBC schemes have been proposed to implement code assignment at the relays without control signaling [9,10]. In particular, in [9], a special discrete randomized Dis-STBC, termed m-group Dis-STBC, was presented where each relay randomly and independently chooses one column from the given underlying STBC matrix. For decentralized Dis-STBC, one approach to implementing relay selection without control signaling is for each node in the decoded set to forward the source message; this is called the *All-Select* strategy, where $K = N$.

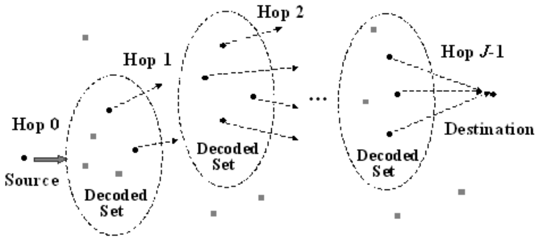

Fig. 2. A single-source J-hop (J>2) selective DF cooperative network

Although a single-source, two-hop cooperative network is the simplest scenario, significant overhead might be required for synchronization, channel estimation, power allocation, relay selection, and link-layer coordination. More detailed discussions of the overhead are presented in Section III.

B. *Single-Source, Multihop Cooperative Network*

A single-source J-hop (J>2) selective DF cooperative network is illustrated in Fig. 2. Except for hop 0, in which the source transmits, the remaining $J-1$ hops are called relaying hops. Although exploiting multihop diversity [13] can improve the energy efficiency, we do not consider that here. Instead, the receiving nodes at each hop can only utilize the transmissions of the immediately previous relay group for decoding.

In a DF multihop network, the destination will not receive the message correctly if any intermediate hop fails to decode the message. Thus, the end-to-end performance is determined by the performance of each individual hop. To improve the performance with minimal overhead, decentralized Dis-STBC can be applied to each relaying hop. The presence of multiple hops complicates this scenario since potential relay nodes must be assigned in some way to an appropriate hop, and each hop must avoid interfering with both the source and the other hops.

C. *Multi-Source, Multihop Cooperative Network*

To incorporate cooperative communications into the networking context, it is essential to consider the possibility of multiple source nodes. For instance, in a multihop selective DF cooperative network, assume that there are U ($U \geq 2$) active sources. All of the U sources may have U distinct destinations, resulting in multiple (s,d) pairs. This scenario is the most complex since each pair competes for use of the channel and the potential relay nodes, requiring additional, and possibly significant, coordination overhead.

III. Overhead Considerations and Challenges

In this section we discuss the incurred overheads of, and the challenges emerging from, implementing decentralized Dis-STBC in a realistic cooperative system. Specifically, we discuss the overhead incurred in our simulation implementation as an example of that for an actual decentralized Dis-STBC system. Since we have made design decisions to minimize overhead and have ignored many other overheads, this represents close to the minimum amount possible in a realistic system. The methodology described here, however, could be used for even more realistic considerations.

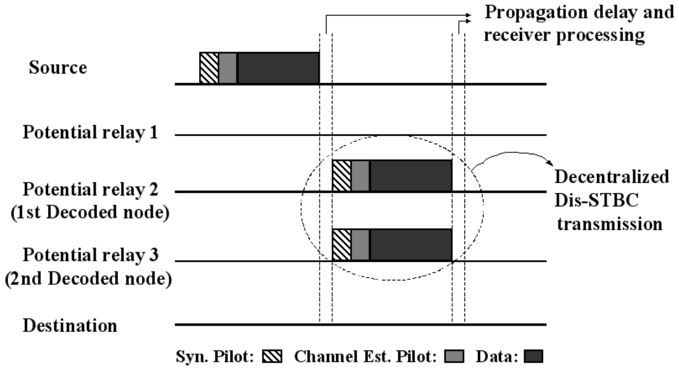

Fig. 3. The physical-layer packet transmission process in a single-source two-hop selective DF cooperative network using All-Select relay selection

A. *Physical Layer*

According to the three scenarios described in Section II, at the physical layer, the operations and the types of incurred overhead are inherently the same in all scenarios. Thus, here, we describe the physical-layer overheads and challenges by using a single-source, two-hop network as an example.

***Synchronization and Channel Estimation*:** In a realistic system, before transmitting data, the transmitter(s) and the receiver(s) must first be synchronized, in both the time and frequency domains. This is generally done by periodically padding a group of known pilot symbols to the beginning of the transmitted signal. Also, to achieve the maximum benefit, each receiver needs to obtain CSI. This could also be estimated by the receiver(s) using the known pilot symbols. An example of these operations can be found in the IEEE 802.11g standard [14].

For the decentralized Dis-STBC transmissions described in Section II, the relays also need to employ decentralized Dis-STBC to implement the transmission of the pilot symbols. In our implementation, we assume the existing pilot symbols in IEEE 802.11 can be re-used for these purposes. To incur minimal overhead, we also assume the estimated CSI at the receiver(s) are not fed back to the transmitter(s).

The physical-layer packet transmission process is illustrated in Fig. 3 for a single-source, two-hop network. In this example, there are three potential relays and the number of nodes in the decoded set is two. In addition, the All-Select strategy is used, and we omit the link-layer control signals (discussed later in Section III-C). The duration of the pilot symbols depends on the duration of the transmitted data symbols.

***Power Allocation*:** For decentralized Dis-STBC, one approach to implementing power allocation without control signaling is for all nodes to use the same fixed transmit power; we call this the uniform-power strategy and use it in our evaluation. In [15], two ad hoc, yet more efficient, power allocation strategies have been proposed for decentralized Dis-STBC. One is the open-loop strategy where the power used by each relay is equal to the power used by the source divided by the number of columns, L, in the underlying STBC matrix. Since the information about L is known to all the nodes *a priori*, this strategy can be implemented without any control signaling. The other is the feedback-assisted strategy where the power used by each relay is equal to the power used by the source divided by the number of relays, K ($K \leq N$). This strategy requires control signaling to enable each relay to obtain the knowledge about K. The uniform-power and open-loop strategies do not incur any overhead; however, the implementation of the feedback-assisted strategy does.

***Time and Frequency Offsets*:** For simultaneous retransmissions by multiple relays positioned at different locations, the different propagation delays from the relays to the destination (or the next relay group) will result in inherent time offsets at the destination, even if all the relays are perfectly synchronized. These time offsets will make the original non-dispersive (flat) relay-to-destination channels time-dispersive (frequency selective), resulting in inter-symbol interference. Further, in a network with mobile nodes, if the impact of Doppler spread is non-negligible, different frequency shifts of different relays will make the pilot-assisted frequency estimation and compensation ineffective, resulting in frequency offsets at the destination. Synchronization errors could also cause time and frequency offsets. The frequency offsets will make the original quasi-static (time-invariant) relay-to-destination channels time-varying. Achieving the theoretic diversity benefit of STBC generally requires that the channels be constant within an encoding block [16]. Obviously, the time-variation of the channels will result in non-negligible performance degradation. Thus, combating the effects of time offsets and possible frequency offsets becomes an important challenge. Fortunately, some effective and practical countermeasures have been proposed, including equalization [17], OFDM [18], and time-reversal STBC (TR-STBC) [18-19].

In general, some additional symbols (such as a guard interval between the neighboring blocks in TR-STBC) are required. To account for this bandwidth overhead in our implementation, we require that each individual data transmission within a burst (as described in Section III-C) has a separate preamble and is separated by a single inter-frame space (SIFS). We assume that if one of the above techniques were used, the required additional symbols can be substituted for the preamble and SIFS, resulting in a similar amount of overhead. We do not consider the processing overhead.

B. *Network Layer and Routing*

We use DSR [20] as the routing protocol. In DSR, when a source node requires a route, a request is flooded through the network to locate the destination. The destination then sends, to the source, one or more unicast responses containing the discovered route(s). The shortest (fewest hops) route is then selected. We term the intermediate nodes in the route *primary nodes*, as opposed to *potential relay nodes* which are selected later. The primary nodes are used until the MAC protocol reports that the route is unusable, at which point a new set is discovered.

The route selection in DSR is certainly different from the idealized case where the best route (e.g., in terms of the signal strength of each relay group) is re-calculated for each packet.

We therefore incur a penalty in terms of sub-optimal routing but as a trade-off incur little control signal overhead. We have modified DSR so that the source route is provided to the MAC protocol with each packet received from the network layer, for use as described in the next subsection.

C. Link-Layer and Medium Access Control

The MAC protocol that we employ here is closely based on Path Access Control (PAC) [21], with two exceptions. First, we do not break the path into multiple segments. Second and more importantly, although PAC suggests incorporating cooperative diversity, no working method of doing so is described in [21]. We describe such a method in Section III.D.

PAC is a CSMA-based scheme where, before a burst of DATA frames are transmitted, the source sends a RTS-like control frame via each primary node in the route to the destination. The destination then responds with a CTS-like control frame sent in reverse along the primary nodes. This exchange reserves the channel along the entire route for a specified period of time, allowing a burst of DATA frames to then be transmitted. Each DATA frame contains a separate preamble and checksum and is separated in time by a SIFS. The data burst is then cooperatively relayed by each relay group in the route.

This process is illustrated in Fig. 4, where each “DATA” includes the “Syn. Pilot, Channel Est. Pilot, and Data” from Fig. 3. We reserve the channel for the entire duration of the transmissions by each hop, which has little impact on performance in small networks ($\leq$ 4 hops). In a larger network, coordinating the simultaneous transmissions by sufficiently distant relay groups requires additional coordination but may be necessary to improve throughput. In our simulation implementation, the source route is included on the RTS, ensuring no ambiguity in the primary nodes to be used. Our protocol differs from [4] in that the RTS/CTS exchange is sent non-cooperatively, and we perform a full-path reservation for a transmission burst instead of a hop-by-hop reservation for each DATA frame. As detailed in [21], the traditional CSMA mechanisms incur significant signaling and other overheads which limit performance. The MAC protocol in [4] doubles the amount of signals over IEEE 802.11 DCF, a substantial cost which we seek to avoid.

No ACK or ARQ mechanism is used here for three reasons. First, if some frames successfully travel only partway to the destination, control signaling must be sent to inform the source and/or relays whether to re-transmit. Coordinating partial re-transmission by relay nodes is difficult and may result in significant overhead in a multihop network. Second, the CTS serves as a response from the destination that the route is working, and a field in the CTS could be used by the destination to indicate the success or failure of the previous burst, allowing the route to be changed accordingly. Third, we assume that cooperative relaying improves the packet loss, so that any loss is significantly less frequent. Thus, we assume any re-transmissions must originate at the source, making the peak throughput independent of the use of ARQ.

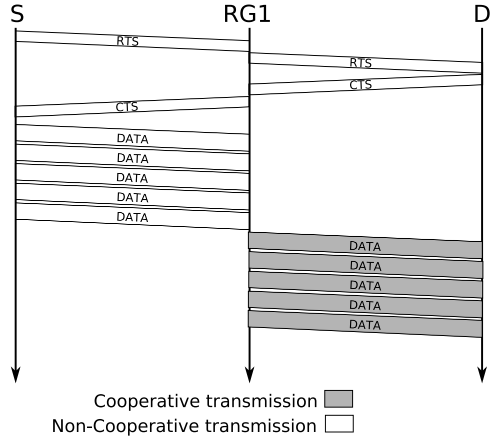

Fig. 4. Timing diagram of our cooperative PAC implementation for a single-source two-hop network with transmissions from the source (S) to the destination (D) forwarded by a single relay group (RG1)

D. Relay Selection and Coordination

In the ideal case, the potential relays for each relay group are selected to create the best end-to-end channel gain (and, therefore, the lowest packet loss probability). In practice, this selection is difficult both due to the computational complexity of choosing the optimal potential relays if the channel conditions are known, and due to the difficulty of predicting future channel conditions based on past history, especially in a mobile environment. The collection of CSI might also require additional control signaling. As a compromise, we utilize a simple, yet effective, selection scheme.

In our implementation, potential relay nodes are selected during the RTS/CTS exchange to reserve the channel. RTS/CTS frames, as well as broadcast route requests, are sent at a lower rate than that for DATA frames. This allows longer links to be selected by the routing protocol and increases the number of nodes near the primary route which can overhear the RTS/CTS. Whether a node hears an RTS or CTS serves as an estimate of the channel conditions at both the node and the transmitter. This is an optimistic estimate due to the lower transmission rate. However, this is acceptable since the node will not act on such decisions unless it is part of the decoded set for a transmission, and since the DATA frames are relayed by multiple nodes at a higher effective power.

To become a potential relay in relay group i, a node must overhear either the RTS or CTS of both the $(i-1)$th and $(i+1)$th primary node (or the source/destination for the first/last relay group). The RTS/CTS specifies the source and destination address as well as the combined duration of the remainder of the RTS/CTS exchange and the burst of DATA transmissions. Each RTS, CTS, and DATA frame contains a relay group header field, which is incremented (decremented in the case of CTS) before transmission by each primary node. Over the duration of the burst, when a potential relay which belongs to group i correctly decodes a DATA frame (i.e., becomes a decoded node), the node becomes a *selected forwarding relay node* in relay group i of this (s,d) pair, and will cooperatively forward the DATA frame with other decoded nodes. That is to say, from the physical-layer point of view, the All-Select strategy is used at each hop.

In a given relay group, when each potential relay correctly decodes either all or none of the DATA frames in the burst, the selected forwarding relay nodes can easily synchronize their transmissions by always sending a fixed time after the end of the last DATA frame in the burst. However, a potential relay in a given relay group might only be a decoded node for some subset of the DATA frames in the burst, making coordination more difficult. Ideally, a potential relay should participate in forwarding any DATA frames that can be correctly decoded by the node. To facilitate this in our implementation, the source sets header fields on each DATA frame indicating the duration of the entire burst and the offset (in μs) within the burst at which the particular DATA frame is transmitted. This allows a potential relay node which receives an arbitrary DATA frame within a burst to calculate when its relay group will forward the DATA frame.

IV. SIMULATION EXAMPLE

A. *Simulation Setup*

In this section we describe simulations using our implementation in QualNet 3.7 [22] as described above. We use a transmit power which results in a transmission range (defined as 95% success) of 225 m for 11 Mbps and 355 m for 1 Mbps. The m-group Dis-STBC with two columns in the underlying STBC matrix is used as the decentralized Dis-STBC scheme. The sources generate constant-bit-rate (CBR) data of size 1000 bytes at regular intervals starting at a random time, and we use a minimum burst size of 10 (with the exception of route replies which may be sent individually) and a maximum burst size of 100 DATA frames. We use IEEE 802.11b with our Dis-STBC modifications at the physical layer, and base our PAC implementation on IEEE 802.11 DCF. All results shown are the average of 25 random network topologies, with each topology simulated for 100 s. Error bars show 95% confidence intervals (calculated as 1.96 × the standard deviation). We simulate cooperative PAC as described in the previous section. To determine whether the performance gain is due to the burst transmission mechanism of PAC or the physical-layer cooperation, we also simulate non-cooperative PAC where only the primary nodes transmit. We also compare against IEEE 802.11 DCF both with the RTS/CTS exchange enabled and without (i.e. direct access mode).

To eliminate the impact of node density and transmission rates, we use four scenarios for each protocol variation. For PAC, we evaluate both a broadcast/control rate of 1 Mbps with a data rate of 11 Mbps, and a fixed rate of 11 Mbps for all transmissions. For IEEE 802.11 DCF, we evaluate both automatic rate adjustment (with broadcasts at 1 Mbps and data transmissions varying automatically between 1 Mbps and 11 Mbps), and a fixed rate of 11 Mbps for all transmissions. We randomly place nodes in either a sparse or dense configuration. For brevity, we only show the results for the scenario under which each of the four protocol variations performs best. Intuitively, cooperative PAC performs best in the dense network with a lower broadcast rate since additional relays are beneficial. Non-cooperative PAC and IEEE 802.11 with RTS/CTS enabled perform best in the dense network with a fixed rate, since the RTS/CTS protection mechanism is effective, long links are undesirable due to their unreliability, and, for IEEE 802.11, the rate adjustment mechanism is overly conservative. Without the RTS/CTS enabled, IEEE 802.11 performs best in the sparse network with a fixed rate, since, without protection, additional nodes more easily interfere.

B. *Decentralized Dis-STBC in QualNet*

For the physical channel, we have modified QualNet 3.7 to incorporate decentralized Dis-STBC into all SINR calculations since QualNet does not normally support this. We use m-group Dis-STBC; thus, a column is randomly chosen for each transmission by each selected forwarding relay. We store the complex Gaussian channel values used in QualNet's existing path loss and Rayleigh fading calculations for each transmission. QualNet's fading calculation has been updated so that the path loss is the variance of the fading, whereas previously the path loss and fading were independent. In QualNet, fading varies on a per-packet basis.

Internally, we maintain a list of the signals each node is currently receiving. Whenever a new signal arrives at a node which is not currently locked onto any signal, we check whether the new signal can constructively combine with any concurrently received signals. If so, the effective power of the combined signals is used to determine whether the node can lock onto the combined Dis-STBC signal. To constructively combine, signals must contain identical bytes, have the same duration and data rate, and arrive within a $2\mu s$ window (although a smaller window is needed for IEEE 802.11 to operate properly, we assume that in an outdoor deployment a suitable protocol will be used, and use IEEE 802.11 only as an example). We use this information along with the stored channel values and column information to calculate the effective SINR. We want to emphasize that, although QualNet separately tracks each signal sent by each of the simultaneously transmitting relays, those signals which would in reality combine into a single Dis-STBC signal are treated precisely as such due to our modifications. The SINR (in particular, the interference) is re-calculated whenever other interfering signals arrive or end during the middle of the desired signal's reception. The corresponding BERs for each bit in the desired signal are then used to determine the probability of successful decoding.

C. *Results in a Two-Hop, Two-Pair Network*

To evaluate Dis-STBC under a basic scenario, we simulate a 265 m × 265 m network, with four nodes placed at the extreme corners and either 3 or 10 additional nodes placed randomly. CBR traffic is sent by each of two sources in the top corners to the corresponding destinations in the opposite bottom corner, and must travel through at least one intermediate hop.

As shown in Fig. 5, PAC has two to four times higher peak throughput than non-cooperative PAC and IEEE 802.11,

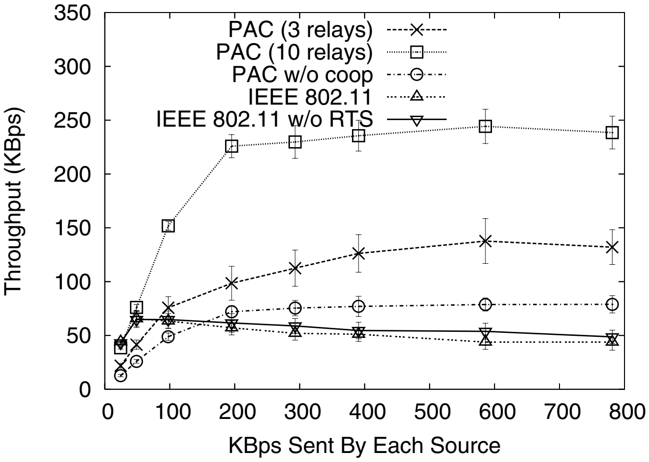

Fig. 5. Comparison of throughput achieved by each protocol in the two (s,d) pair, two-hop case

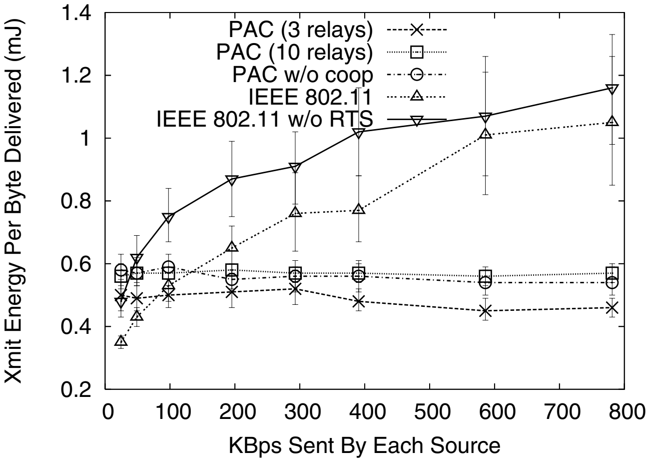

Fig. 6. Comparison of energy efficiency of each protocol in the two (s,d) pair, two-hop case

achieving the highest performance in the dense (10 relay) network since cooperation allows a higher SINR and fewer failed transmissions to more efficiently utilize bandwidth. When the load is low, PAC performs slightly worse than IEEE 802.11 since PAC does not incorporate an ARQ scheme. Non-cooperative PAC has a slightly higher throughput than IEEE 802.11 under high load, but it is clear that cooperative relaying significantly improves performance.

In terms of energy efficiency, as shown in Fig. 6, PAC uses about half the energy per byte delivered under high load compared to IEEE 802.11. However, two significant observations must be made. First, with 3 relays present, PAC is only slightly more energy efficient than non-cooperative PAC. With 10 relays, the case where PAC has significantly higher throughput, there is no statistically significant difference compared to non-cooperative PAC because the additional nodes present result in higher SINR at the cost of additional relay transmissions. Second, when IEEE 802.11 is operating at peak throughput (approximately 50 KBps/source), IEEE 802.11 has similar energy efficiency to PAC when PAC is operating at peak throughput (approximately 600 KBps/source). Thus, the difference in energy efficiency between PAC and IEEE 802.11 is only a result of IEEE 802.11's medium access mechanism operating poorly under much higher load than can be supported. Further, PAC's energy efficiency under high load is primarily a result of the medium acquisition costs being amortized over the multiple packets in each DATA burst, since PAC and non-cooperative PAC show similar trends.

While PAC shows definite promise in terms of throughput, it is clear that if energy efficiency is a motivating factor, the relay selection protocol must intelligently adjust the number and/or transmit power of relay nodes. Such coordination must be carefully designed to maintain high SINR without wasting significant bandwidth on additional control signals.

D. *Results in a Large Four-Pair Network*

We evaluate Dis-STBC in a rectangular network of width 300 m and varying length. We use four (s,d) pairs, randomly placing two sources and two destinations within 100 m of the left end of the network, and the corresponding destinations and sources within 100 m of the right end, resulting in traffic flowing in both directions. For the 450 m length, we randomly place 8 additional nodes in the sparse case, and 16 nodes in the dense case. As the network length increases, nodes are added to maintain a constant density. We evaluate each protocol under various traffic levels and show the peak throughput achieved in Fig. 7.

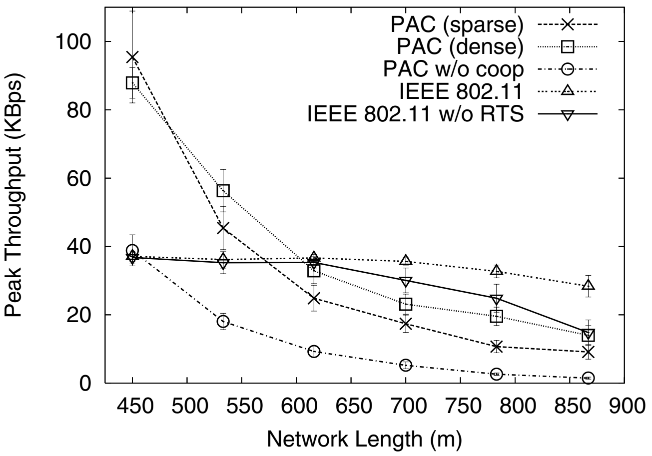

Fig. 7. Comparison of peak throughput achieved by each protocol in the four (s,d) pair, large network case

When the network is small (below 600 m long), PAC achieves much higher throughput than IEEE 802.11 and non-cooperative PAC, as observed in Section IV-C. However, as the network grows, PAC's performance decays rapidly while IEEE 802.11 with the RTS/CTS enabled maintains its level of performance. PAC's poor performance is somewhat unexpected due to the larger number of possible relay combinations, but several key overheads affect the results.

First, non-cooperative PAC shows a drop in performance similar to PAC because both protocols suffer from poor spatial channel re-use since the entire path between the source and destination (effectively the majority of the network) is reserved for a single (s,d) pair at any given time; IEEE 802.11, on the other hand, opportunistically allows multiple distant nodes to transmit simultaneously. This is further compounded by the cooperative relaying in PAC, which causes a larger amount of interference over a wider region than non-cooperative relaying and is somewhat contrary to the concept of increasing channel re-use through aggressive power control.

Second, the routes chosen by the routing protocol are imperfect and are sometimes chosen such that there are few potential relay nodes. Predicting the quantity or quality of potential relay nodes a given route will have is difficult but, even in a narrow network where routing is “easy,” some consideration of appropriately determining potential relay nodes during route selection must be made.

Third, DSR, like many efficient routing protocols, relies on overheard control signaling to cache and maintain route information. In a cooperatively relayed transmission, each individual sender cannot identify itself to its neighbors, otherwise each signal would contain different and conflicting bytes. Route replies and other routing traffic are critical and, therefore, should be sent cooperatively to increase their delivery, but, without the identification of each transmitter, their information is less useful when overheard. Thus, more overhead-intensive route discovery schemes than those used in non-cooperative networks may be necessary.

V. CONCLUSION

The great promise of cooperative communications can only be demonstrated by taking this new link abstraction into networking contexts to evaluate realistic networking scenarios with the consideration of all of the overheads incurred across protocol layers. In this paper, we take the first steps in analyzing the potential overheads incurred when incorporating cooperative communications into networking contexts, and in performing a holistic evaluation of cooperative networks using decentralized Dis-STBC. Simulation results in terms of network throughput and energy efficiency demonstrate the superior performance of networks employing cooperative communications, as well as the challenges faced.

Our results show that while cooperative communications is promising, several problems must be addressed before its practical use in general networks can be realized. In particular, careful and joint consideration of relay selection, power control, and routing is necessary to achieve the full benefits. Further, it is vital that the evaluation of such protocols be conducted in a realistic environment, where all of the overheads are considered, in order for their true performance to be known. In the future, we plan to extend this work to design a systematic methodology of evaluating cooperative communications enabled wireless networks that may employ different relay selection protocols, different distributed space-time coding schemes, and different MAC and routing protocols, while taking all of the necessary overheads into account.

REFERENCES

[1] A. Scaglione, D. Goeckel and J. Laneman, ”Cooperative communications in mobile ad hoc networks: rethinking the link abstraction,” *IEEE Sig. Proc. Magazine*, pp. 18-29, Sept. 2006.

[2] F. Fitzek and M. Katz, *Cooperation in Wireless Networks: Principles and Applications - Real Egoistic Behavior is to Cooperate!*, Springer, 2006.

[3] G. Jakllari, S. Krishnamurthy, M. Faloutsos, P. Krishnamurthy, and O. Ercetin, ”A framework for distributed spatio-temporal communications in mobile ad hoc networks,” in *Proc. of Infocom 2006*, May 2006.

[4] G. Jakllari, S. Krishnamurthy, M. Faloutsos, P. Krishnamurthy, and O. Ercetin, ”A cross-layer framework for exploiting virtual MISO links in mobile ad hoc networks,” *IEEE Trans. on Mobile Computing*, vol. 6, no. 6, pp. 579-594, June. 2007.

[5] Z. Tao, T. Korakis, Y. Slutskiy, S. Panwar, and L. Tassiulas, ”Cooperation and directionality: a co-opdirectional MAC for wireless ad hoc networks,” in *Proc. of WiOpt 2007*, April 2007.

[6] J. Laneman and G. Wornell, ”Distributed space-time-coded protocols for exploiting cooperative diversity in wireless networks,” *IEEE Trans. on Inform. Theory*, vol. 49, no. 10, pp. 2415-2525, Oct. 2003.

[7] J. Luo, R. Blum, L. Cimini, L. Greenstein, and A. Haimovich, ”New approaches for cooperative use of multiple antennas in ad hoc wireless networks,” in *Proc. of VTC 2004-Fall*, pp. 2769-2773, Sept. 2004.

[8] A. Stefanov and E. Erkip, ”Cooperative coding for wireless networks,” *IEEE Trans. on Comm.*, vol. 52, no. 9, pp. 1470-1476, Sept. 2004.

[9] J. Luo, R. Blum, L. Greenstein, L. Cimini, and A. Haimovich, ”Link-failure probabilities for practical cooperative relay networks,” in *Proc. of VTC 2005-Spring*, pp. 1489-1493, May 2005.

[10] B. Sirkeci-Mergen and A. Scaglione, ”Randomized space-time coding for distributed cooperative communication,” *IEEE Trans. on Sig. Proc.*, vol. 55, no. 10, pp. 5003-5017, Oct. 2007.

[11] S. Barbarossa and G. Scutari, ”Distributed space-time coding for regenerative relay networks,” *IEEE Trans. on Wireless Comm.*, vol. 3, no. 5, pp. 2387-2399, Sept. 2005.

[12] P. Anghel, G. Leus, and M. Kaveh, ”Distributed space-time cooperative systems with regenerative relays,” *IEEE Trans. on Wireless Comm.*, vol. 5, no. 11, pp. 3130-3141, Nov. 2006.

[13] J. Boyer, D. Falconer, and H. Yanikomeroglu, ”Multihop diversity in wireless relaying channels,” *IEEE Trans. on Comm.*, vol. 52, no. 10, pp. 1820-1830, Oct. 2004.

[14] IEEE802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Aug. 1999.

[15] L. Zhang and L. Cimini, ”Efficient power allocation for decentralized distributed space-time block coding,” submitted to *IEEE Trans. on Wireless Comm.*, 2007. (http://www.eecis.udel.edu/ lzhang/publication.html)

[16] E. Larsson and P. Stoica, *Space-Time Block Coding for Wireless Communications*, Cambridge University Press, 2003.

[17] J. Mietzner and P. Hoeher, ”Distributed space-time codes for cooperative wireless networks in the presence of different propagation delays and path losses,” in *Proc. of IEEE Wrksp. on Sensor Array and Multi channel Sig. Proc.*, July 2004.

[18] Y. Mei, Y. Hua, A. Swami, and B. Daneshrad, ”Combating synchronization errors in cooperative relays,” in *Proc. of IEEE ICASSP 2005*, March 2005.

[19] M. Sharp, A. Scaglione, and S. Galli, ”Distributed randomized space-time coding for HF transmission,” in *Proc. of Milcom 2006*, Oct. 2006.

[20] D. Johnson, D. Maltz, and J. Broch, *DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks*, in *Ad Hoc Networking*, edited by Charles E. Perkins, Chapter 5, pp. 139-172, Addison-Wesley, 2001.

[21] R. Ramanathan, ”Challenges: A radically new architecture for next generation mobile ad hoc networks,” in *Proc. of MobiCom 2005*, Aug. 2005.

[22] Scalable Network Technologies, Inc., QualNet Simulator, http://www.scalable-networks.com.