Networking with Cooperative Communications: Holistic Design and Realistic Evaluation

Justin Yackoski† Lu Zhang† Bo Gui† Chien-Chung Shen* Len Cimini*
*Department of Computer and Information Sciences
†Department of Electrical and Computer Engineering
University of Delaware, Newark, DE 19716, U. S. A.

Abstract—**Cooperative communications fundamentally changes the abstraction of a wireless link and offers significant potential advantages for wireless networks. However, much of the existing work on cooperative systems focuses on their physical layer capabilities. Such systems are evaluated with idealized scenarios where coordination, control signaling, and other costs are assumed to be negligible or even free. These costs are actually significant enough that the net gain of using cooperative communications in a networking context remains uncertain. Moreover, even if these costs appear small in simple network scenarios, they will increase significantly with the size and traffic level of the network, requiring careful evaluation, or even innovative design, of cooperative protocols to ensure their usefulness in realistic networks. In this article, taking into account the overheads incurred at the different protocol layers, we describe a realistic performance evaluation of cooperative communications in a networking context. The insights obtained from this evaluation help guide work on cooperative communications towards practical and potentially beneficial protocols.**

I. INTRODUCTION

The notion of cooperative communications[1] fundamentally changes the abstraction of a 'wireless link' and promises improved performance [1]. In contrast to single-node transmissions, cooperative communications may involve multiple relay nodes transmitting simultaneously to a receiver. To facilitate such simultaneous transmissions and receptions, many different overheads are incurred, some of which are not well understood, and are often ignored in most existing research. 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 when incorporating cooperative communications as a link abstraction into a networking context with multiple active flows and multi-hop relaying. For instance, relay selection and coordination requires control signaling at the MAC layer among potential relay nodes. Achieving maximum performance demands simultaneous consideration of routing, relay selection, power control, and other factors. Without a holistic evaluation within an end-to-end networking context which considers these costs, the true potential net benefit of cooperative communications cannot be explored.

We investigate both the overheads and net performance gain of incorporating cooperative communications in realistic networking scenarios by taking most of the necessary overheads into account in a network level study. As a minimal overhead baseline, we evaluate decentralized Distributed Space-Time Block Coding (Dis-STBC) in conjunction with a low-overhead, path-centric MAC protocol (which handles channel reservation as well as relay selection and coordination). Studying these mechanisms allows us to measure the cost in terms of bandwidth and energy consumption of the additional coordination required for cooperative communications. This work allows us to identify any "hidden" costs that are either not expected to impact the performance of cooperative communications or have been completely ignored by previous work. These costs identify areas where more consideration is needed, both by those measuring the theoretical benefit of new cooperative mechanisms, and by those designing cooperative communications-aware network protocols.

II. RELATED WORK

Cooperative communications, in terms of its physical layer capabilities and characteristics, have been extensively studied, but, only recently has there been work that explores the benefits of cooperation at different protocol layers. In this section we review some of the existing research in this area.

Several methods have been proposed to modify the IEEE 802.11 Distributed Coordination Function (DCF) protocol to enable nodes to cooperate to address poor quality at the link layer. In [2], [3], a single additional node is selected during the RTS/CTS exchange which can either replace

---

An earlier version of the paper was presented at IEEE Workshop on Cooperative Communications and Networking: Theory, Practice, and Applications (CoopNet), in conjunction with IEEE ICC, Beijing, China, May 23, 2008. 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.

[1] In this article, the terms cooperative communications and cooperative diversity are used interchangeably.

the original next-hop node in the route, or serve as an additional temporary hop between the current sender and the original next hop. Although these methods improve performance through cooperation, this cooperation is only used to temporarily modify the route; **no physical-layer cooperative diversity is employed**. In [4], these ideas are taken a step further, and the single additional node that is selected can transmit concurrently with the original node in the route using cooperative diversity.

In [5], another modification to the IEEE 802.11 DCF protocol is described where each *primary node* randomly selects a fixed number of additional nearby relay nodes. Each signal (RTS, CTS, DATA, ACK) is first sent non-cooperatively by the primary nodes. These signals are heard by the primary node's relay nodes, which then re-transmit the signal. Since this doubles the number of transmissions, the number of hops must be shortened by approximately a factor of two to achieve better performance than the conventional protocol. Unfortunately, the extent to which path shortening provides a benefit is strongly limited by the additional interference it causes. Due to this trade-off, as well as the random relay group selection and fixed transmit power schemes used, the achievable improvement of this design is limited.

The schemes described above are all variants of the IEEE 802.11 DCF MAC protocol, which is inefficient in terms of control signaling and bandwidth utilization. Although these approaches achieve better results than IEEE 802.11 DCF, IEEE 802.11 DCF suffers in general from poor performance due to high overheads, which cannot be addressed through the incorporation of cooperative diversity. As discussed in [6], the design problems of IEEE 802.11 DCF-based protocols demand a new integration of physical, link, and network layer designs to expand multi-hop wireless networks beyond their current limitations.

In [6], a path-oriented MAC scheme is described which reserves the channel for a burst of packets over multiple hops, eliminating the majority of the current per-hop and per-packet costs. A partial implementation was evaluated in [7]; however this work only reserves the channel for a single packet, instead of a burst, causing high control overheads.

III. COOPERATIVE COMMUNICATIONS

In this section we describe typical scenarios for a cooperative network which employs selective decode-and-forward (DF) relaying between randomly distributed single-antenna half-duplex nodes. Although several cooperative techniques have been proposed for use with selective DF relaying (for example, see [1], [8], [9]), to achieve a balance between performance and overhead, we focus on decentralized Dis-STBC. In [10], 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 [10], [11] 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 [1], [12]. In particular, [12] desribes a special discrete randomized Dis-STBC, termed m-group Dis-STBC, 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 the *All-Select* strategy where each node in the decoded set forwards the source's message[2].

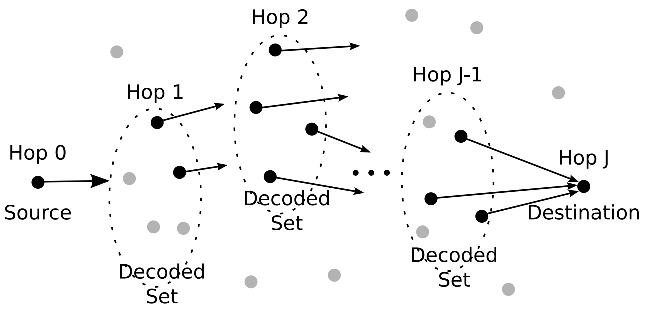

Fig. 1. A single flow J-hop selective DF cooperative network

Consider the single flow (i.e. single source transmitting to a single destination) J-hop ($J{>}2$) selective DF cooperative network is illustrated in Fig. 1. Except for hop 0, in which the source transmits, the remaining $J-1$ hops are called relaying hops. In selective DF relaying, a potential relay node is called a *decoded node* if it can correctly decode the source message. In the first hop, the source transmits while nodes in hop 1 listen. Then, the decoded nodes in hop 1 transmits while hop 2 listens, and so on until the source's message reaches the destination. We consider the case where the receiving nodes at each hop can only utilize the transmissions of the immediately previous relay group for decoding. Exploiting multi-hop diversity (e.g., by allowing nodes in hop 2 to collect and combine signals transmitted by hop 1 and hop 0) can increase the diversity gain but also incurs additional costs. In a DF multi-hop 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.

To incorporate cooperative communications into a realistic networking context, it is essential to consider the possibility of multiple concurrent flows attempting to transmit through the network. This scenario is the most complex

---

[2] We use *message* to refer to a piece of data from an end-to-end, application layer perspective, and *frame* to refer to a message combined with a link layer header.

since each flow competes for use of the channel and the potential relay nodes, requiring additional, and possibly significant, coordination overhead.

IV. EXAMPLE OF INCORPORATING COOPERATION WITH MINIMAL OVERHEAD

In this section we discuss the incurred overheads of, and the challenges emerging from, implementing cooperative communications in a realistic network. Specifically, we discuss the overhead incurred in our designed protocol, Cooperative PAC, as an example of that for an actual decentralized Dis-STBC system. Since we have made design decisions to minimize overhead and have ignored some overheads, this represents close to the minimum amount possible in a realistic system. The methodology described here, however, could also be used for evaluating other scenarios where additional overheads are present.

A. Physical Layer Overheads

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 training 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 training symbols. An example of these operations can be found in the IEEE 802.11g standard.

For the decentralized Dis-STBC transmissions described in Section III, the relays also need to employ decentralized Dis-STBC to implement the transmission of the training symbols. In our implementation, we assume the existing training 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). We use the All-Select strategy, described in the previous section.

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 [13], 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 at the time of deployment, 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. 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 transmissions by multiple relays positioned at different locations, the different propagation delays from the relays to the destination (or the next relay group) result in inherent time offsets at the destination, even if all the relays are perfectly synchronized. These time offsets 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. Frequency offsets 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 [14]. 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 [15].

In general, some additional symbols are required for synchronization. To account for this bandwidth overhead in our implementation, we require that each individual data transmission within a burst (as described in Section IV-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 and assume correctable time and/or frequency offsets are always perfectly corrected.

B. Network Layer and Routing Overheads

Cooperative PAC uses DSR 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*. 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 message to be transmitted. Cooperative PAC therefore incurs a penalty in terms of sub-optimal routing but as a trade-off incurs little control signal overhead. We have modified DSR so that the source route is provided to the MAC protocol, for use as described in the next subsection.

C. Link Layer and Medium Access Control Overheads

As detailed in [6], the traditional CSMA mechanisms incur significant signaling and other overheads which limit performance. With cooperative communications, each transmission is made by a group of nodes, instead of a single node, so decisions must be made on a group basis. This requires some form of additional information collection and/or distribution within each potential relay group. A *potential relay group* is a group of *potential relay nodes*, nodes which participate in the cooperative transmission if they are decoded nodes.

As just one example of these problems, when a DATA frame is sent from one potential relay group to the next, the receiving potential relay group should acknowledge reception of the frame. By design, the decoded nodes may include only a subset of the potential relay group. A single node within the potential relay group, regardless of whether it decoded the frame, cannot determine whether the frame should be acknowledged, since the node is not aware of whether enough other potential relay nodes decoded the frame. Compared to the already high overheads of regular IEEE 802.11 DCF, the cooperative diversity MAC protocol in [5] sub-optimally addresses the acknowledgement problem by designating one node in the potential relay group as the leader who determines, based solely on its own decoded state, whether the frame should be acknowledged.

Given the high overhead of IEEE 802.11 DCF based solutions, we design a MAC protocol, termed ‘Cooperative PAC’ that is closely based on Path Access Control (PAC) [6], with two exceptions. First, the path is not broken into multiple segments. Second and more importantly, although PAC suggests incorporating cooperative diversity, no working method of doing so is described in [6]. PAC is a CSMA-based scheme where, before a burst of DATA frames are transmitted, the source sends an 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 in Cooperative PAC is illustrated by Fig. 2. 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 [5] 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. Since [6] traditional CSMA mechanisms already incur significant signaling and other overheads [6], we seek to reduce these overheads instead of increasing them (e.g. [5] increases control signaling by a factor of two).

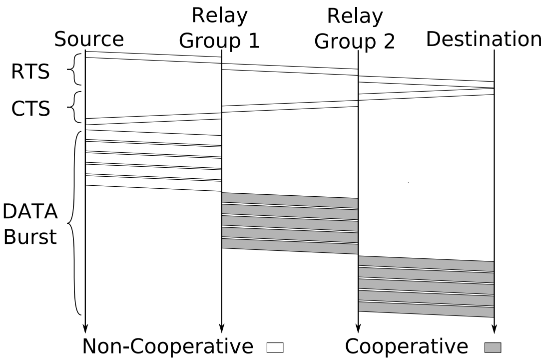

Fig. 2. Timing diagram of our Cooperative PAC implementation for a single-source three-hop network with transmissions from the source (S) to the destination (D) forwarded via relay groups 1 and 2.

The RTS-CTS exchange reserves the channel for the entire duration of the transmissions by all hops, which has little impact on performance in small networks ($\leq$ 3 hops) where spatial reuse is not possible. In a larger network, coordinating the simultaneous transmissions by sufficiently distant relay groups requires additional coordination but may be necessary to improve throughput. Even when only one flow is active, with longer routes ($\geq$ 4 hops) it becomes possible for the transmissions to be “pipelined” to increase throughput. For example, the source may transmit to the first hop while the third hop transmits to the fourth hop. In our implementation, nodes may pipeline their transmissions using knowledge of frame transmission durations to maintain a fixed, globally pre-set number of hops separation between simultaneously transmitting hops within the same flow. This simple pipelining scheme allows for intra-flow spatial reuse with no additional overhead. We further discuss the effectiveness of such pipelining in Section V-B.

No ACK or ARQ mechanism is used in Cooperative PAC in the interest of minimizing overhead. If some messages 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 multi-hop network. Omitting ARQ allows us to measure peak throughput, but has several drawbacks discussed in Section V.

D. Relay Selection and Coordination Overheads

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 loss probability). In practice, this selection is difficult both due to the computational complexity of choosing the optimal potential relays if the channel

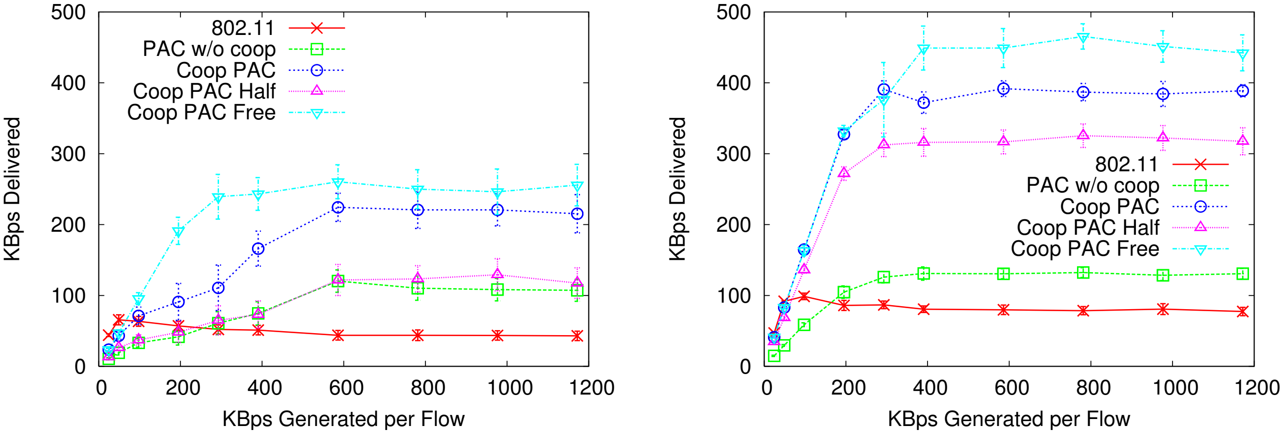

(a)

(b)

Fig. 3. Throughput in the **(a)** sparse and **(b)** dense two-hop, two-flow scenarios

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 of Cooperative PAC, potential relay nodes are selected during the RTS-CTS exchange to reserve the channel. RTS and 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 RTSs and CTSs. 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 control signal transmission rate. However, this is acceptable since the node will not act on such decisions unless it has decoded a DATA frame transmission, and since the DATA frames are relayed by multiple nodes at a higher effective power.

To become a member of potential relay group i, a node must overhear either the RTS or CTS of both the $(i-1)$th and $(i+1)$th primary nodes (or the source/destination for the first/last potential relay group). If pipelining is used, nodes may also meet this criteria by similarly hearing any of the first few DATA frames in the burst. Each RTS, CTS, and DATA frame 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 also contains a 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 group i of this flow, and will cooperatively forward the DATA frame with other decoded nodes.

In a given 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 when pipelining is not used, a potential relay in a given 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 group will forward the DATA frame.

V. REALISTIC EVALUATION AND SIMULATION RESULTS

In this section we evaluate Cooperative PAC and describe some of the problems these results identify by incorporating the true cost of all control signals used for coordination. We evaluate IEEE 802.11 DCF with and without the RTS-CTS mechanism enabled, but only show the enabled case (*802.11*) which has better results in the multi-hop networks we consider. To determine whether the improvement of Cooperative PAC is due to the path-based channel reservation or cooperative communications, we show both regular Cooperative PAC (*Coop PAC*) and *PAC w/o coop* where only the primary nodes relay. In *Coop PAC Free*, control signals consume no bandwidth or energy. Finally, *Coop PAC Half* is a simple power allocation scheme where relays use half the normal transmit power when transmitting cooperatively.

Evaluation is performed using QualNet, with our modifications to accurately simulate cooperative communications in Cooperative PAC with two column m-group Dis-STBC. We use a transmit power which results in a transmission range (defined as 95% success rate) of 225 m for 11 Mbps and 355 m for 1 Mbps. We use a 2-μs window within which cooperatively transmitted signals must start being received to constructively combine[3]. Each source generates 1000 byte CBR packets and we use a minimum and maximum burst size in Cooperative PAC of 10 and 100 packets, respectively. To eliminate the impact of transmission rates, we evaluate Cooperative PAC using a broadcast rate, for route requests and RTS-CTS frames, of 1 Mbps and a data rate of 11 Mbps (allowing more potential relay nodes to be selected). We evaluate *PAC w/o coop* and *802.11* using a fixed 11 Mbps rate for all transmissions, the best scenario for non-cooperative protocols since a lower broadcast rate results in the selection of long, poor-quality links being and IEEE 802.11's automatic rate adjustment is overly conservative. All results 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).

A. Two-Hop, Two-Flow Network

To evaluate Cooperative PAC under a basic scenario, we simulate a 265 m × 265 m network, with four nodes placed at the extreme corners and traffic generated between opposite corners. We place either 3 (for the sparse scenario) or 10 (for the dense scenario) additional nodes randomly. Cooperative PAC selects a subset of these nodes as potential relay nodes. Of those, generally 1-2 for the sparse scenario and 3-6 for the dense scenario actually relay. As shown in Fig. 3, some throughput gains over *802.11* are from using path-based reservation (*PAC w/o coop*). Adding cooperation, *Coop PAC*, particularly in the dense scenario, results in significant further improvement.

These results lead to several observations. First, in low traffic *802.11* achieves higher throughput because of the use of ARQ. For low traffic or cases where loss is not handled well by upper layers, ARQ may provide a net benefit to cooperative protocols despite the added complexity. Second, all variants of Cooperative PAC have higher throughput in the dense network due to the additional relay nodes. Thus to maximize performance it is important to be able to potentially utilize many available nodes, which may require coordination with the routing protocol. It may also be useful for nodes to dynamically determine whether the cost of a given cooperative mechanism is worthwhile depending on the node density. Third, even through Cooperative PAC transmits only two small control frames per hop, per burst, removing their bandwidth consumption in *Coop PAC Free* increases throughput by 10-15% over *Coop PAC*. The high cost of these few control frames illustrates the challenge of achieving a net gain in performance by sending additional control signals to gather information and make more informed decisions. Substantial performance gains and/or very efficient use of control frames may be needed to realize any useful benefits. Finally, in the sparse scenario *Coop PAC Half*'s throughput is half that of *Coop PAC* and nearly identical to *PAC w/o coop*, while in the dense case *Coop PAC Half*'s throughput is only about 15% lower than *Coop PAC*, which brings us to the question of energy efficiency.

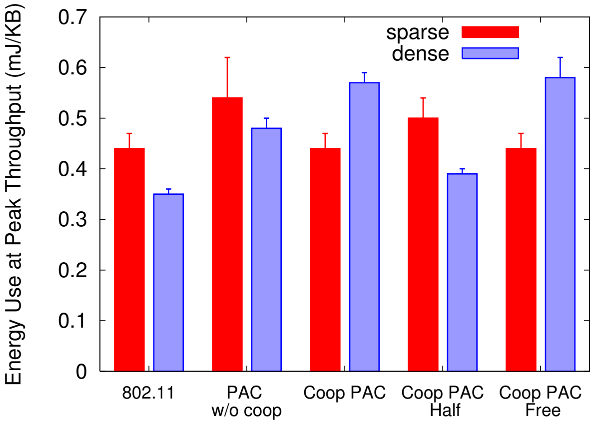

Fig. 4. Energy efficiency in the two-hop, two-flow scenario

Fig. 4 shows the energy efficiency in both the sparse and dense scenarios, with each protocol's energy use measured at its peak throughput from Fig. 3. We measure only energy used for transmission. For scenarios with devices which consume significantly more energy receiving than in an idle or sleep state, one must also consider the energy consumed both by adding idle nodes to the network and by increasing the number of nodes actually used to relay.

These results have several implications for cooperative communications. First, *802.11* uses energy most efficiently in both scenarios. While cooperative diversity increases throughput and in theory could reduce the energy needed for transmissions, the lack of power control wastes energy. Note that since *Coop PAC* and *Coop Pac Free* show similar results, the energy to transmit control signals has minimal impact. Second, *Coop PAC Half* in the dense scenario has both higher throughput and lower energy usage per byte than *Coop PAC* in the sparse scenario. In other words, given that *Coop PAC* in the dense scenario has desirable throughput but undesirable energy usage, if the goal is to reduce energy usage with the least impact on throughput, it is better to adjust the transmit power used by relays (i.e. *Coop PAC Half* in the dense scenario) than to reduce the number of relays (i.e. *Coop PAC* in the sparse scenario). This is intuitive since having more relays leads to a fuller realization of the diversity gains, motivating the use of power control by varying the transmit power at each node, instead of only varying the number of relays as in current cooperative protocols. Finally, with *Coop PAC*, less energy is used per byte in the sparse scenario than in the dense scenario, but with *Coop PAC Half* the opposite is true. This is because there are few relay nodes in the sparse

---

[3] We use IEEE 802.11b which in reality requires a larger window. We assume a suitable encoding will be used in practice and use IEEE 802.11b as an example.

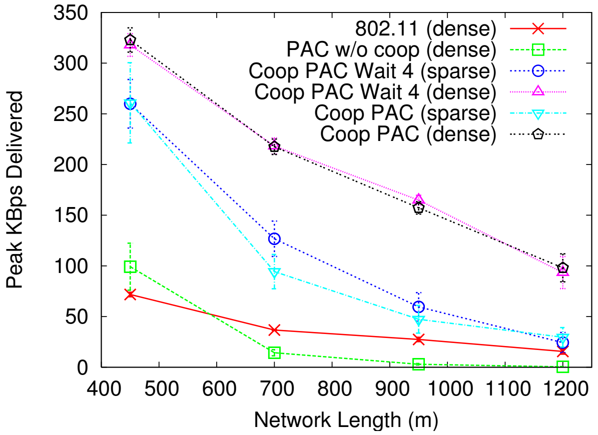

Fig. 5. Throughput in the four-flow, large network scenario

scenario. Thus lowering the relay transmit power in *Coop PAC Half* results in a weak signal and, while less total power is consumed, a significant amount of this power is wasted since many transmissions are too weak to be decoded. In the dense scenario, many potential relays are present and often the combined signal is unnecessarily strong when relay transmit at full power. Thus, we conclude that power control should be considered during the relay node selection (and route selection) process, and must adapt to the relay node availability in order to avoid both possible causes of wasted energy.

B. *Large, Four-Flow Network*

We consider a rectangular network of width 300 m and varying length. Source and destination nodes are placed randomly near opposite ends of the network, with two flows in each direction. For the 450 m length, we randomly place 8 additional nodes in the sparse scenario, and 16 nodes in the dense scenario. 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. 5. We omit results for *802.11* and *PAC w/o coop* in the sparse scenario as they are nearly identical to the dense scenario.

Cooperative PAC achieves much higher throughput than *802.11* and *PAC w/o coop* when the network is small, but degrades rapidly as the length increases. The decrease is somewhat unexpected due to the larger number of possible relay combinations (resulting in a stronger combined signal and less loss), but several issues limit performance. First, *PAC w/o coop* shows a drop in performance because the path-based reservation, while efficient in terms of using minimal control signaling, suffers from poor spatial channel reuse since the entire path between the source and destination is reserved (sometimes unnecessarily) for a single flow at any given time. *802.11*, on the other hand, opportunistically allows multiple distant nodes to transmit simultaneously since channel reservations are made on a smaller scale. The reservation problem is further compounded by cooperative relaying, which causes a larger amount of interference over a wider region than non-cooperative relaying. The RTS-CTS exchange, which is sent non-cooperatively, is thus subject to more interference. 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 for a given route is difficult but, even in this narrow rectangular network where routing is "easy," considering some estimate of potential relay nodes during route selection may be useful. Third, DSR, like many routing protocols, relies on overheard control signaling to cache and maintain route information. Since this overhearing is not possible with cooperative transmissions (each cooperatively transmitting node does not identify itself), the routing protocol has less information available when making routing decisions. Finally, Cooperative PAC does not use ARQ and so the loss probabilities at each hop accumulates. An efficient method is needed to coordinate recovery from loss, particularly when a packet is lost in the middle of a route.

We attempt to address spatial reuse within a flow by allowing transmissions to be pipelined, shown by *Coop PAC Wait 4*, where the source may start transmitting another packet once the fourth relay group in the route has time to receive the previous packet, instead of the normal method used in *Coop PAC* (shown in Fig. 2). The non-uniformity in the distances between relay groups means that such a scheme applied on either a global or per-flow basis remains sub-optimal. Furthermore, without the use of power control, cooperation has the effect of increasing the average length of hops such that the 1200 m long network uses only 3 to 4 hops instead of the expected 5 to 6 (based on a maximum 225 m transmission range). This increase significantly expands the region in which a transmission produces interference and limits opportunities for spatial reuse of the channel, evidenced by the smaller gap between the two schemes in the dense scenario.

VI. CONCLUSION

By taking the first steps in analyzing the potential overheads incurred when incorporating cooperative communications into networking contexts, this article describes a realistic evaluation of cooperative networks. Using the insights obtained, the article shows the importance of a holistic design of cooperative networking protocols to reduce overhead and improve end-to-end performance. However, much work needs to be done to support the adoption of cooperative communications into multi-hop wireless networks.

REFERENCES

[1] A. Scaglione, D. L. Goeckel, and J. N. Laneman, "Cooperative communications in mobile ad hoc networks: rethinking the link abstraction," *IEEE Signal Processing Magazine*, vol. 23, pp. 18–29, September 2006.

[2] P. Liu, Z. Tao, S. Narayanan, T. Korakis, and S. Panwar, “A cooperative MAC protocol for wireless LANs,” *IEEE JSAC, Special Issue on Cooperative Communications and Networking*, vol. 25, Feb. 2007.

[3] J. Yackoski and C.-C. Shen, “Cross-layer inference-based fast link error recovery for MANETs,” in *Proc. of IEEE WCNC 2006*, April 2006.

[4] S. Moh, C. Yu, S. Park, H. Kim, and J. Park, “CD-MAC: cooperative diversity MAC for robust communication in wireless ad hoc networks communications,” in *Proc. of IEEE ICC 2007*, pp. 3636–3641, June 2007.

[5] G. Jakllari, S. V. Krishnamurthy, M. Faloutsos, P. V. Krishnamurthy, and O. Ercetin, “A cross-layer framework for exploiting virtual MISO links in mobile ad hoc networks,” *IEEE Transactions on Mobile Computing*, vol. 6, pp. 579–594, June 2007.

[6] R. Ramanathan, “Challenges: A radically new architecture for next generation mobile ad hoc networks,” in *Proceedings of the 11th Annual International Conference on Mobile Computing and Networking (MobiCom)*, (Cologne, Germany), pp. 132–139, August 2005.

[7] R. Ramanathan and F. Tchakountio, “Channel access over path segments for ultra low latency MANETs,” in *Proc. of IEEE Milocm 2007*, Oct. 2007.

[8] J. Luo, R. S. Blum, L. J. Cimini, L. J. Greenstein, and A. M. Haimovich, “New approaches for cooperative use of multiple antennas in ad hoc wireless networks,” in *Proceedings of the 60th IEEE Vehicular Technology Conference*, (Limassol, Cyprus), pp. 2769–2773, September 2004.

[9] A. Stefanov and E. Erkip, “Cooperative coding for wireless networks,” *IEEE Transactions on Communication*, vol. 52, pp. 1470–1476, September 2004.

[10] J. N. Laneman and G. W. Wornell, “Distributed space-time-coded protocols for exploiting cooperative diversity in wireless networks,” *IEEE Transactions on Information Theory*, vol. 49, pp. 2415–2525, October 2003.

[11] P. A. Anghel, G. Leus, and M. Kaveh, “Distributed space-time cooperative systems with regenerative relays,” *IEEE Transactions on Wireless Communications*, vol. 5, pp. 3130–3141, November 2006.

[12] J. Luo, R. S. Blum, L. J. Greenstein, L. J. Cimini, and A. M. Haimovich, “Link-failure probabilities for practical cooperative relay networks,” in *Proceedings of the 61st IEEE Vehicular Technology Conference*, (Stockholm, Sweden), pp. 1489–1493, May 2005.

[13] L. Zhang and L. J. Cimini, “Efficient power allocation for decentralized distributed space-time block coding,” *IEEE Transactions on Wireless Communications*, to appear.

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

[15] Y. Mei, Y. Hua, A. Swami, and B. Daneshrad, “Combating synchronization errors in cooperative relays,” in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, (Philadelphia, PA, USA), pp. 369–372, March 2005.