

Clique-Based Randomized Multiple Access for Energy-Efficient Wireless Ad Hoc Networks

Paul G. Flikkema and Brent West

Department of Electrical Engineering

Northern Arizona University

Flagstaff, AZ 86011-5600, USA

Email: {Paul.Flikkema, Brent.West}@nau.edu

Abstract—A fundamental tradeoff in MAC protocol design for wireless ad hoc networks is between proactive and reactive coordination, where the latter is used to resolve access conflicts whose severity is determined by the former. In this paper, we describe Clique-Based Randomized Multiple Access (CRMA), a distributed MAC protocol for wireless ad hoc network applications. Of the many objectives in MAC design for this application, CRMA places strongest emphasis on (i) energy efficiency and (ii) reliance only on local (one-hop) connectivity information. CRMA forms collections of nodes, or cliques, separated by one hop, and provides the proactive coordination required for clique members to synchronize their wake-sleep cycles. Each clique selects a slot in the clique’s frame pseudo-randomly, so that that no proactive coordination between cliques is required. To limit potential access conflicts, CRMA can exploit bandwidth via frequency hopping or spread spectrum coding; these also provide robustness to multiple-access interference, exogenous interference, and frequency selectivity. It also allows the use of multiple or multi-channel radios to increase performance. With a slight amount of additional proactive coordination, CRMA can also employ what we call predictive conflict resolution, wherein clique members predict access conflicts and resolve them ahead of time.

I. INTRODUCTION

WIRELESS sensor networks, as well as other ad-hoc systems that network energy-limited nodes, have different constraints than wired networks [1], and necessarily place more emphasis on some characteristics while simultaneously compromising in other areas. For example, unlike in wired networks, it is neither possible (nor, often, desirable) for a node to completely seize the channel resource. Wireless networking often involves nodes with significant energy consumption constraints. For this reason, a prominent difference between wired and wireless nets is that latency, throughput, and bandwidth efficiency are often traded for energy efficiency in the latter.

MAC (media access control) protocols for wireless ad hoc networks must balance a large number of conflicting goals. In addition to delivering sufficient energy and bandwidth efficiency, they should

- provide acceptable latency and throughput
- be scalable with network size
- admit energy-aware routing

This work was supported in part by the National Science Foundation under Grant EIA-0131691. The work of B. West was supported in part by the Merriam-Powell Center for Environmental Research.

- be robust to node failures and changes to the network topology
- be PHY layer aware, exploiting time, frequency, space, and angle diversity if available
- support robustness to interference at the PHY layer
- be scalable with node resources (e.g., multiple or multi-channel radios).

This list illustrates the recent re-awakening of interest in cross-layer interaction in wireless networking. Looking upward from the MAC layer, it is clear that the same connectivity information used in routing might improve MAC energy efficiency. In the other direction, the integrated design of MAC and PHY layers is well-established. For example, spread spectrum modulation is commonly used to mitigate the combination of multiple-access interference, frequency selectivity, and exogenous interference. One of the purposes of this paper is to describe a MAC protocol whose design is informed by both requirements and assets at neighboring network levels.

Conventional taxonomies place MAC protocols into two classes (hybrid strategies are also possible) [2]:

- 1) those that *avoid channel contention* using deterministic algorithms (e.g., TDMA) using either a fixed or dynamic assignment; and
- 2) those that *resolve contention* with random access techniques (e.g., CSMA protocols).

A second level of classification normally cited is whether the algorithm is centralized or decentralized. Since centralized protocols scale poorly in multi-hop networks and in geographic distance and number of nodes, we consider only distributed protocols here.

The fundamental MAC-layer design tradeoff in ad-hoc wireless networks is between *coordination* and *contention*. Routing algorithms can be classified into proactive or reactive [3]. Similarly, MAC-layer coordination can be further divided into *proactive coordination* during an *organizational* (or planning) phase and *reactive coordination* during the network’s *operational* phase. Greater levels of proactive coordination imply less contention and hence less need for reactive coordination.

A greater degree of proactive coordination also implies greater energy consumption in the organizational phase, while less implies more reactive coordination (and more energy consumption) in the operational phase. At one end of this continuum are ALOHA and carrier-sense methods (in, e.g., [4]),

which are reactive mechanisms. Slotted ALOHA improves over ALOHA in both throughput and energy efficiency by employing a modicum of proactive coordination: synchronization. Carrier-sense methods provide lightweight reactive coordination in that random backoffs explicitly acknowledge the presence of other users. Unfortunately, random access methods exhibit very poor energy efficiency, since they prevent nodes from sleeping—the best form of energy conservation.

At the other extreme of the coordination-contention continuum are reservation-based methods such as TDMA [5] which avoid contention as much as possible by using a great deal of proactive coordination. The appropriate amount of proactive coordination is driven by the dynamics of the physical network topology and the traffic statistics. For example, high dynamics and low traffic levels imply that bandwidth and energy consumed during the planning phase would be largely wasted.

II. RELATED WORK

Most MAC algorithms for wireless ad-hoc networks attempt to strike a balance between pure proaction or pure reaction in order to allow nodes to sleep while still limiting the amount of proactive coordination. In the receiver-oriented protocol of [6], nodes advertise when they will be listening, and other nodes coordinate their transmissions accordingly. If multiple nodes wish to send to a node, they must resolve contention via RTS/CTS. The work of [7] uses a greater degree of proactive coordination to avoid contention. Nodes invite other nodes to form links, and linked nodes then cooperate to choose locally contention-free time slots. Some form of diversity (e.g., frequency hopping) can then be used to soften contention from nearby, interfering nodes.

A useful MAC strategy is *local proactive coordination*: nodes within one hop of each other proactively coordinate to some degree, while a reactive mechanism handles contention from other nodes. One advantage of this is scalability. In this paper, we introduce Clique-Based Randomized Multiple Access (CRMA), a MAC algorithm that relies on both local proactive coordination and reactive coordination. A primary objective of CRMA is to rely on simple, one-hop information, yet provide extremely good energy efficiency. CRMA employs a type of randomized (not random) access that simplifies proactive coordination and provides robustness to channel degradations such as frequency selectivity and interference. However, within local groups of nodes called cliques, access is deterministically scheduled.

In comparison with the protocols of Sohrabi and Pottie [7] and Ye *et al.* [6], our approach uses more proactive coordination than [6], but less than [7]. By using less coordination, CRMA has more potential contention, but this can be reduced by making more bandwidth available. Like CRMA, [7] has low bandwidth efficiency, as mentioned in [6]. However, sensor network applications typically require only low data rate links, enabling the aggressive trade of bandwidth efficiency for energy efficiency.

Three major features distinguish CRMA from previous works. The first is the concept of *cliques*: a node's cliques

are subsets of its set of one-hop neighbors that it may wish to communicate with. A clique determines meeting times much as people do: the member nodes agree to a pseudo-randomly determined future time interval (and, optionally, a frequency and/or a direct-sequence spread spectrum code, both also pseudo-randomly determined) at which to communicate. Cliques generalize the notion of links, allowing multicasts and broadcasts in addition to unicasts. Unlike receiver-oriented protocols, the use of cliques in CRMA implies that transmitters and receivers are coordinated. This enables extremely high energy efficiency: a node awakens (and hence consumes energy) for communication during a given slot only if it is time for the clique to meet. The protocol of [6] also forms groups of nodes called virtual clusters, but only for time synchronization.

Nodes using the SEEDEx protocol [8] also publish the outputs of a pseudo-random number generator to implicitly reserve slots. These slots are for listening only. SEEDEx was developed with the goal of improving per-node throughput in large networks, but can also be used to improve energy efficiency. SEEDEx does not incorporate cliques, so transmitters must contend for a node's listening slot. CRMA also employs the additional coordination infrastructure of frames, so that per-hop latency can be bounded.

The second distinguishing feature of CRMA is its ability to cleanly exploit multiple communication resources. The PAMAS protocol [9] is designed for an air interface consisting of a data channel and a dedicated control channel. This form of reactive control underpins a simple messaging structure to inform neighbors when they can sleep, but uses additional energy. The protocol of [10] uses two radios per node and can work at reduced capacity if one fails, but is based on local polling. CRMA efficiently uses an arbitrary number of multichannel radios per node and admits graceful degradation in cases of failure.

The third feature is what we call *predictive conflict resolution*. We use the term conflict resolution, rather than collision resolution, to indicate that clique meetings are rescheduled so that no collisions (actually, conflicts of a certain type) occur. Section V describes an approach that employs a small amount of additional proactive coordination to predict and avoid most conflicts.

III. THE BASIC CRMA PROTOCOL

We assume that each node in the network (perhaps after route discovery has completed) contains a list of its one-hop neighbors, and that each node is synchronized with its neighbors, so that time can be slotted.

Each node then forms memberships in a number of cliques (subsets of one-hop neighbors). We will describe this in detail in Section IV; for the moment it is sufficient to consider the set of two-node cliques (each consisting of the node in question and a neighbor).

Time for each clique is divided into frames which contain N fixed time-duration slots. In general, a frame consists of N_t time slots, N_f frequency, and N_c code division multiple access (CDMA) slots, so that $N = N_t N_f N_c$. These slots can be made orthogonal, meaning that a collision occurs only when two or

more nodes transmit at the same time, frequency, and with the same code. In many cases, due to cost or bandwidth constraints, $N_f = 1$ or $N_c = 1$.

Each node then contacts the members of each clique q , and they cooperate to distribute a common package of information, consisting of (1) a pseudo-random number generator (PRNG), (2) an initial seed, and (3) a start-of-first-frame time some interval in the future. The members of the clique update their PRNG's each frame, thus agreeing on a common slot n_q in the next frame.

IV. CLIQUE SELECTION

Since, by definition, each node may only communicate with its one-hop neighbors, the most obvious arrangement for a node is to form a clique with each of its one-hop neighbors. Therefore, each clique consists of only two nodes as shown in Figure 1, and is logically a point-to-point (unicast) link. If a node has n one-hop neighbors, then it can join up to n unicast cliques. To ensure contention-free access during the clique's meeting time only one of the clique's members may transmit at a time. This can be accomplished by defining two distinct cliques for each pair of nodes or by simply dividing the meeting time in half and allowing each node to transmit only in their respective half.

In multicast cliques (cliques with > 2 members) a master-slave arrangement is more suitable, whereby the clique has only one master who controls access to the channel—slaves speak only when spoken to. If the application requires broadcast support in addition to the point-to-point links the number of cliques that each node belongs to increases to $2n + 1$ since now each node must keep track of a point-to-point clique and broadcast clique for each of its neighbors plus a broadcast clique for itself. For example, in the case of Figure 1, node A would become a member of 7 cliques: three point-to-point cliques with nodes B, C, and D, three broadcast cliques during which it must wakeup and listen for broadcasts from B, C, and D, and a broadcast clique during which it transmits to nodes B, C, and D simultaneously. Some applications may not require dense connectivity among nodes in which case the number of cliques each node belongs to may be substantially reduced through a method of selective attachment.

V. PREDICTIVE CONFLICT RESOLUTION

Since the number of slots per frame is finite, two cliques will occasionally schedule a conflicting slot. There are two types of conflicts that can occur in CRMA. To analyze this we start with a definition: two cliques are said to intersect if they share one or more nodes. Thus all cliques that a node belongs to intersect. This is a logical intersection; non-intersecting cliques can still interfere physically. *Hard conflicts* occur at a node when the number of intersecting cliques reserving the same slot exceeds the number of available radios. For the purposes of this paper, we define a radio as an air interface instrument that can transmit or receive (but not both) one information stream, and stores packets, meaning it implements PHY, MAC, and LLC processing. (Note that the use of orthogonal codes such as Walsh codes




Fig. 1. Example ad hoc network; point-to-point cliques shown.

would permit multiple transmit streams per radio.) *Soft conflicts* occur when neighboring, non-intersecting cliques use the same slot. While soft conflicts result in collisions, they may not be catastrophic since packets from two geographically dispersed nodes still may not result in fratricide at the intended receiver(s) due to channel losses (propagation, shadowing, and fading) and the capture effect.

Hard conflicts can be prevented using *predictive conflict resolution*. In PCR, all nodes in a clique individually look ahead an arbitrary number of frames to predict if conflicts will occur. If a conflict is found by any node in the clique, it (1) substitutes a *conflict resolution seed* for the frame in which the hard conflict occurs, updates the PRNG, and checks if the conflict is resolved. If not, it keeps updating the PRNG and checking until it is. Then it notifies the other members of the clique who perform the same seed substitution and number of updates. This notification is performed far enough in advance for all the other nodes in the clique to positively acknowledge the change.

For example, suppose a node predicts a conflict between two of its cliques. It then chooses one of the two cliques and employs the conflict resolution seed in the PRNG until the conflict is resolved. When it is, it notifies the other members of that clique of the (future) frame in which the conflict occurs. All members then use the conflict resolution seed for determination of a new slot in that frame. The updated conflict resolution seed is then stored for later use (when the next conflict is discovered), and the default seed is recalled for use.

The newly chosen slot may, with a very low probability, also conflict with slots selected by other non-intersecting cliques. When this occurs, the node that finds such a conflict will then initiate the same process with any of the conflicting cliques. Clearly, if the network load is light and one looks far enough in advance, most predicted collisions can be resolved.

Implementation of PCR involves negligible additional com-

putational load. The future slots that must be computed using the PRNG's would have been computed later. The additional load comes from the search for conflicts and—when a conflict is discovered—the additional updates of the PRNG's with the conflict resolution seed. PCR also requires allocation of additional memory for tables containing the future slots. Hence PCR requires additional proactively allocated resources (primarily memory), as well as more reactive computation to resolve predicted conflicts.

There are a number of possibilities that require less reactive coordination. The simplest approach is for all conflicting cliques but one to go to sleep during the conflicting slot. This saves energy, but also causes additional latency. Another possibility is a one-shot attempt at collision avoidance: rather than use a conflict resolution seed, the nodes can simply choose to claim the previous or next slot in that frame: if N is large, then this will usually resolve the conflict. However, neither of these approaches can guarantee resolution of a conflict.

VI. DYNAMIC RESERVATIONS

The basic CRMA protocol provides a method for a clique to claim a slot (or, by simple extension, a fixed number of slots) per frame. The computational and memory infrastructure used for PCR can also be exploited for demand-based claims on a variable number of slots. If a node needs to send traffic at a higher rate to a given clique, it can claim additional slots in a future frame as follows. It first chooses the slots (perhaps randomly), then checks for hard conflicts. If the new slots do not conflict with other cliques in that frame, the node informs the other members of the clique of the new plan. The newly claimed slots may conflict with slots already claimed by cliques that intersect at the clique members of the originating node. The effect of this conflict propagation will be worse than for PCR, since more (rather than just different) slots are claimed. For this reason, the protocol can include a veto by clique members if desired.

VII. PERFORMANCE ANALYSIS

For a fixed number of cliques in an area, the probability of conflicts goes down as the number of slots increases. Let c_j be the number of cliques intersecting at node j with r_j radios. Assume that the set of all PRNG's generates a set of mutually independent processes that are also temporally independent. Assume also that the number of slots N is common across all cliques. Let the probability that all c_j cliques intersecting at node j communicate free of hard conflicts in one frame of N slots after k PCR updates be denoted by

$$p_{\text{free},H}^{(j)}(c_j, r_j, N, k).$$

For the case $r_j = 1$ and no PCR, this probability is

$$p_{\text{free},H}^{(j)}(c_j, 1, N, 0) = \frac{N!}{(N - c_j)!N^{c_j}}.$$

When $r_j > 1$, the number of radios is not the only limitation, since no conflict occurs only if the number of cliques claiming the slot is $\leq r_j$ and the time-frequency-code slot is free. For $N_f N_c \rightarrow \infty$, we have for $r_j = 2$ that

$$\begin{aligned} p_{\text{free},H}^{(j)}(c_j, 2, N, 0) &= p_{\text{free},H}^{(j)}(c_j, 1, N, 0) \\ &+ \frac{1}{N^{c_j}} \sum_{i=1}^{\lfloor \frac{c_j}{2} \rfloor} \overbrace{(2, 2, \dots, 2)}^i \overbrace{(1, 1, \dots, 1)}^{c_j-2i} (i, c_j - 2i, N - c_j + i)! \end{aligned} \quad (1)$$

This is an upper bound when $N_f N_c$ is finite. Additionally, if PCR is employed, then $1 - p_{\text{free},H}^{(j)}(c_j, r_j, N, 0)$ is the probability that PCR will be invoked. The sum in (1) is driven by the first multinomial factor, which defines an equivalence class of outcomes; the extension to cases when $r_j > 2$ is straightforward, but lengthy.

To perform a first-order analysis of the effect of PCR on hard conflicts, we assume that the conflict resolution seed generates an independent slot, so that we are concerned with the probability that it will conflict with the other cliques intersecting at that node. We also assume that the clique's conflict is the only hard conflict at that node. Since there are N slots, the probability of conflict after a single PRNG update with the conflict resolution seed is $\frac{c_j-1}{N}$ if $r_j = 1$. Hence the probability that a clique can communicate free of hard conflicts at node j with k updates in this case is

$$\begin{aligned} p_{\text{free},H}^{(j)}(c_j, 1, N, k) &= \\ &1 - \left[1 - p_{\text{free},H}^{(j)}(c_j, 1, N, 0) \right] \left(\frac{c_j - 1}{N} \right)^k. \end{aligned} \quad (2)$$

Soft conflicts depend on the number of slots and the physical topology via the number of cliques within interfering range, or earshot. They also depend on the number of radios per node, r . Somewhat paradoxically, increasing r can increase the probability of soft conflicts, since bandwidth utilization increases. The probability of soft conflicts could be reduced by the use of connectivity information beyond one hop at the sacrifice of simplicity.

The energy efficiency of a MAC can be defined as the ratio of successfully used energy to the total energy consumed. The authors of [6] identify several sources of energy waste in wireless ad hoc nets: collisions, overhearing, control packet overhead, and idle listening. CRMA addresses these by its use of cliques and the strategy of only awakening for communication when the clique meets. Organization into cliques is a form of proactive coordination that allows efficient scheduling of synchronized wake-sleep cycles. Increasing either or both of N and the number of PCR k rounds can reduce the number of hard collisions to a negligible level. As mentioned earlier, soft conflicts may result in collisions, due to our choice of limiting the information exchange to single hops. The probability of soft conflicts can be reduced by the use of a large number of slots (N). A generalized analytical quantification of the throughput

and energy efficiency of CRMA and similar MAC approaches is underway.

VIII. CONCLUDING REMARKS

The design of MAC protocols for energy-constrained ad hoc wireless networks must balance a number of conflicting goals. Perhaps of paramount importance is where the MAC protocol lies in the joint parameter space of capacity, energy efficiency, and degree of proactive coordination. The CRMA protocol emphasizes the latter two characteristics, and employs cross-layer interaction by using and/or enabling capabilities found at neighboring network architecture layers. The capacity of wireless ad hoc nets is under intense investigation; in particular, [11] considers capacity within the context of a particular class of TDMA MAC protocols that assumes a high degree of proactive coordination. The development of analytical frameworks that explicitly capture coordination and energy efficiency in addition to capacity would be of great interest.

IX. ACKNOWLEDGMENT

The authors thank B. Ruggeri for helpful discussions on this topic.

REFERENCES

- [1] K. Sohrabi *et al.*, "Protocols for self-organization of a wireless sensor network," *IEEE Personal Communications*, pp. 16–27, 2000.
- [2] F. A. Tobagi, "Multiaccess link control," in *Computer Network Architectures and Protocols* (C. A. Sunshine, ed.), pp. 139–189, Plenum, second ed., 1989.
- [3] Z. J. Haas and S. Tabrizi, "On some challenges and choices in ad-hoc communications," in *Proc. IEEE MILCOM 1998*, pp. 187–192, 1998.
- [4] E. Callaway *et al.*, "Home networking with IEEE 802.15.4: A developing standard for low-rate wireless personal area networks," *IEEE Communications Magazine*, vol. 40, no. 8, pp. 70–77, 2002.
- [5] I. Chlamtac and S. Pinter, "Distributed nodes organization algorithm for channel access in multihop dynamic radio network," *IEEE Transactions on Computers*, vol. C-36, pp. 728–737, June 1987.
- [6] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," in *Proc. IEEE INFOCOM 2002*, 2002.
- [7] K. Sohrabi and G. J. Pottie, "Performance of a novel self-organization protocol for wireless ad hoc sensor networks," in *Proc. of the IEEE 50th Vehic. Tech Conf.*, pp. 1222–1226, 1999.
- [8] R. Rozovsky and P. R. Kumar, "SEEDEX: a MAC protocol for adhoc networks," in *Proc. MOBIHOC 2001*, pp. 67–75, Oct. 1998.
- [9] S. Singh and C. S. Raghavendra, "PAMAS: power-aware multi-access protocol with signalling for ad hoc networks," *ACM Computer Communication Review*, vol. 28, pp. 5–26, July 98.
- [10] Z. J. Haas, "On the performance of a medium access control scheme for the reconfigurable wireless networks," in *Proc. IEEE MILCOM 1997*, 1997.
- [11] S. Toumpis and A. J. Goldsmith, "Capacity regions for wireless ad hoc networks." (<http://wsl.stanford.edu/Publications/Stavros/twc01.pdf>; submitted for publication).