Supporting Concurrent Transmissions in Multi-hop Wireless Networks

Arup Acharya
IBM TJ Watson Research Ctr
arup@us.ibm.com

Archan Misra
IBM TJ Watson Research Ctr
archan@us.ibm.com

Sorav Bansal
IBM India Research Labs
soravban@in.ibm.com

Abstract

Several studies have shown that multi-hop networks using IEEE 802.11 wireless LANs, based on a RTS/CTS MAC, exhibit significant throughput degradation. The IEEE 802.11 DCF MAC disallows any parallel transmission (or reception) or in the neighborhood of either a sender or a receiver (of an ongoing transmission). MACA-P is a set of enhancements to the 802.11 MAC that allow parallel transmissions in many situations when two neighboring nodes are either both receivers or transmitters, but a receiver and a transmitter are not neighbors. Like 802.11, MACA-P contains a contention-based reservation phase prior to data transmission. Unlike 802.11, the data transmission is delayed by a control *phase* interval, which allows multiple sender-receiver pairs to synchronize their data transfers thereby avoiding collisions.

1. INTRODUCTION

Wireless networks have received an inordinate degree of attention from the research community over the last 5-7 years. The bulk of this research can be classified into the following distinct categories:

1. *Improvements in MAC protocols for single-hop wireless LANs.*
   Enhancements such as QoS differentiation (e.g., [1]) and fair bandwidth sharing in IEEE 802.11, or the design of contention resolution schemes in the presence of uni-directional links, are designed primarily for the WLAN environment, where nodes attach to the wired backbone over a single wireless hop.
2. *Routing protocols for mobile ad-hoc networks.*
   The emphasis here has been on the development of protocols (e.g., AODV [2], DSR [3]) that establish traffic routes in environments where node mobility results in rapid changes in network topology. Although such protocols sometime utilize the broadcast nature of the wireless medium (e.g., route snooping in DSR), their primary focus is on rapid recovery from link failures and the avoidance of long-lived routing loops in mobile environments.

The research community has, however, largely ignored the problem of efficient data forwarding in multi-hop, wireless environments. For example, we have only recently seen some research publications (e.g., [4]) showing how TCP performs poorly in multi-hop 802.11 networks. In this paper, we present a set of techniques to allow better operation of *single-cell MAC protocols in a multi-cell/multi-hop network. Our principal objective is to significantly increase the number of parallel (or simultaneous) transmissions in such a multi-hop network by allowing the MAC layer to effectively exploit spatial diversity.*

Since 802.11 is the dominant technology for wireless LANs for now and the near future, we believe that the best techniques are those based on simple extensions to the basic 802.11 contention resolution mechanism. Our proposed protocol, MACA-P, is based on 802.11's basic 4-way handshake mechanism, and can achieve significant improvements in network throughput by simply relaxing some of the unnecessary restrictions in the 802.11 collision avoidance mechanism.

2. MEDIUM ACCESS CONTROL (MAC) FOR MULTI-CELL/ MULTI-HOP NETWORKS

The broadcast nature of the wireless medium implies that a transmitter and receiver node can communicate effectively as long as the MAC layer ensures adherence to the following fundamental constraint: *no receiving node can be within the reception range of more than one simultaneously transmitting node, since such concurrent transmissions will lead to collision and incorrect reception at the receiver (see Figure 1*).

Current work on 802.11 (research [5] and standardization) however imposes a more rigorous constraint: the MAC layer effectively ensures that *no node that is a one-hop neighbor of either the sender or the receiver of a data packet may be engaged in any communication activity (either transmitting or receiving) during the entire 4-way (RTS-CTS-DATA-ACK) exchange*. To consider the differences between these two constraints, see Fig.2, where Q and B are one-hop neighbors, and A's transmission range does not include Q (and vice versa), and P's transmission range does not include B (and vice versa). It is clear that the transmission patterns shown in cases (3) and (4) shown in Fig.2c are not inherently feasible. In case (3), B's transmission to A would collide with P's transmission at Q, while in case (4), A's transmission (to B) would collide with Q's transmission (to P) at B. For case (1), however, since A's transmission range does not include Q and P's transmission range does not include B (Fig. 2a), the two transmissions can proceed in parallel; a similar argument applies to case (2) as well (Fig. 2b). The 802.11 MAC is unduly restrictive and prohibits cases (1) and (2) essentially because both the sender and the recipient of a data packet *revert between transmitting and receiving roles multiple times over a continuous interval* during the packet transfer. Since data packet recipient acts as a receiver during the RTS and DATA portions, and the sender acts as a receiver during the CTS and ACK portions, the entire neighborhood of *both* nodes is effectively silenced during the entire duration of the 4-way handshake.

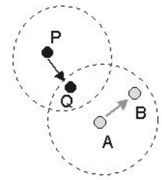

Fig. 1

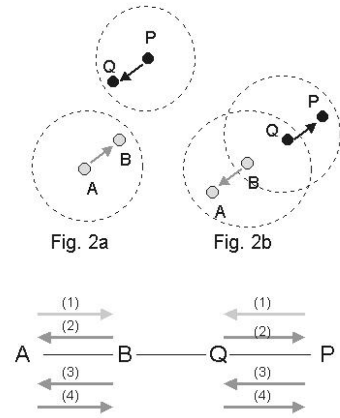

Fig. 2a Fig. 2b

Fig. 2c

Significant improvements to the overall system throughput of multi-hop networks can however be realized if this constraint on concurrent packet exchanges can be relaxed or modified. In fact, the initial papers on CSMA-CA (e.g., [5][6]) alluded to the possible exploitation of spatial diversity for parallel (concurrent) transmissions, but did not proceed with research in that direction. Recent attempts at improving the spatial reuse of the multi-hop network typically focus on two approaches, both of which fundamentally aim to *reduce the size of the one-hop neighborhood* and thus allow the network to be *partitioned* into a greater number of zones of concurrent transmissions:

a) Power control algorithms, e.g. [7]
b) Use of directional antennas, e.g. [8].

While simulation studies indicate that both approaches can significantly improve the aggregate channel capacity of multi-hop networks, they do suffer from certain drawbacks. Distributed versions of power-control protocols require nodes to include and decipher the transmission power levels in the header of MAC control packets. In real-life situations, where the interference range is larger than the actual packet reception range, a node can suffer interference effects from neighboring transmitters even though it cannot correctly receive their packets (and thus cannot correctly perform the appropriate power-level computations). Directional antennas, on the other hand, use sophisticated hardware and phase-modulation strategies, and may not prove to be cost-effective solutions for large-scale deployment, especially in pervasive and mobile devices. Indeed, the focus of these current approaches is on simply increasing the number of disjoint network segments that can proceed in parallel, rather than on fundamentally trying to relax 802.11's constraint.

We believe that a fresh look is needed for medium access control that is inherently targeted for wireless, multi-cell, concurrent operation. To that end, we propose a protocol called MACA-P (MACA with Parallel Transmissions), where the 802.11 constraint (of silent neighborhoods of both sender and receiver) is replaced with the fundamental constraint, i.e., a receiver cannot be not in the neighborhood of more than one transmitter. The key idea is to allow neighboring nodes to synchronize their reception periods, so that one-hop neighbors switch between transmitting and receiving roles in unison and thus avoid the problem of packet collisions. This objective can be achieved without any basic changes to the 802.11 4-way handshake, by introducing a *variable control gap* (a period of silence) between the RTS/CTS exchange and the DATA and ACK phases. One node generates a *master transmission* schedule and other neighboring nodes synchronize to that schedule (the interim control gap provides neighbors an opportunity to set up their individual transmissions).

3. BUILDING BLOCKS OF MACA-P

Control Phase : The first enhancement we make to the RTS/CTS protocol is to add extra information in the RTS and CTS messages to explicitly delineate the intervals for both the data transfer phase and the ack phase, thereby allowing other nodes to know exactly when the two nodes associated with the transmission under consideration switch between tx and rx roles. This is done by explicitly introducing two time intervals to the RTS and CTS control messages:

- T_{DATA} : specified as a time interval after the reception of this control message, this indicates the start time of the data transmission.
- T_{ACK} : specified as a time interval after the reception of this control message, this indicates the start time of the ACK control message.

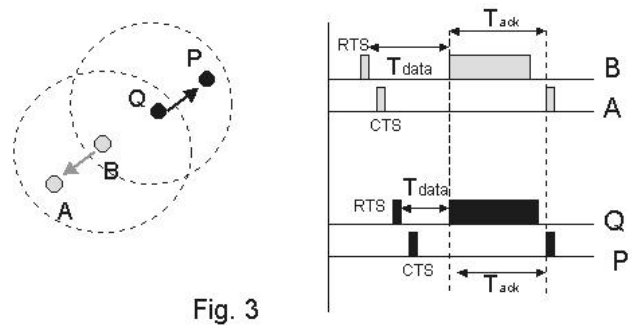

Fig. 3

This is a departure from the format of the RTS/CTS messages of the 802.11 MAC, which include a single time interval. By introducing the two time intervals, we are explicitly demarcating the control phase, the data transmission phase and the ACK phase. The time interval T_{DATA} informs all neighboring nodes of the duration of the control phase and allows for the possibility that an overlapping transmission may be scheduled aligned with the end of this interval. The second interval T_{ACK} is necessary to align the ACKs of overlapping transmissions. In Figure 3, Q overhears the RTS sent from Bto A, and becomes aware that a neighboring node is initiating a transmission. If it has a packet to transmit during this time, it will initiate a RTS whose T_{DATA} is aligned with the start time of Bs data transmission. Both RTS and CTS messages carry the two intervals so that nodes that are neighbors of either the sender or the recipient learn of scheduled data and ACK transmissions after T_{DATA} and T_{ACK} intervals.

Inflexible Bit in RTS : The RTS message is further enhanced to carry a bit which we call the **inflexible** bit. The purpose of this bit is to indicate to the receiver of the RTS message whether the transmission schedule proposed in the RTS message can be changed: if the bit is set, then this schedule cannot be changed.

Modification of T_{DATA} and T_{ACK} by CTS : When a node receives a RTS where the inflexible bit is not set, it may change the proposed schedule of the sender so that it can align the proposed data transmission by the sender with an existing scheduled reception in its neighborhood. This is done by modifying the T_{DATA} and T_{ACK} received on the RTS message, and sending back the modified values on the CTS message. Consider Figure 4, where B has overheard the CTS from Q and is aware of a scheduled reception in its neighborhood. Thus, when it receives a RTS from A with the inflexible bit unset, it responds with a modified T_{DATA} and T_{ACK} (shown as t1 and t2) so that B's reception of data from A overlaps with Q's reception.

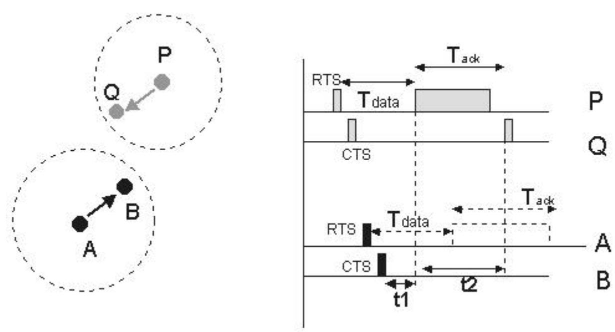

Fig. 4

RTS' control message : When a sender transmits a RTS message, neighbors of the sender have updated their respective NAVs based on overhearing the RTS. However, if the schedule proposed in the RTS is modified by the receiver of the RTS message (as discussed earlier), neighbors of the sender are not updated of the modified schedule. To avoid such a situation, the sender always sends a gratuitous RTS message containing the updated T_{DATA} and T_{ACK} that it received from the CTS. A second use of the RTS' message is to cancel a prior schedule made through the matching RTS, when the sender did not receive a CTS from the intended receiver (of the RTS).

Master Transmission Schedules: A master transmission schedule is one that allows neighbors of either the sender or recipient to schedule their own *data transmission* overlapping with the master. A set of aligned data transmissions is referred to as a transmission set. The data sender of a master transmission will be referred to as a master sender, and similarly the data recipient will be referred to as a master recipient. In general, we will simply use the term "master" when it is clear from the context whether it refers to a sender or a recipient. We now state a key requirement whether it is possible for a sender/recipient pair to schedule a transmission depending on the number of masters in their neighborhood:

> *A sender/recipient pair can schedule a data transmission only if there is at most one master transmission in the sender's neighborhood or at most one master reception in the recipient's neighborhood, but not both.*

To see why we mandate such a requirement, consider Figure 5. In Figure 5a, Y is neighbor of B and Q, but B is not a neighbor Q. The two transmissions A-to-B and P-to-Q have been scheduled, i.e. Y has two masters, B and Q. X then sends a RTS to Y. If Y has to fit in this transmission, it must align X's data transmission with P's data transmission (Q's reception) and stretch out its (Y's) ACK to X to align with B's ACK to A. In general, if a node has more than 1 master, it has to align the proposed data transmission with that of the master with earliest data transmission and align the ACK with that of the master with the latest ACK. First, this adds complexity to our solution. Second, all master (recipient) nodes other than the master with the latest ACK, are blocked from scheduling any further receptions till the master transmission with the latest ACK, completes. In the figure, this means Q cannot schedule any further reception (from P, say) before Y sends its ACK to X (aligned with the ACK from B to A). Otherwise, a subsequent CTS from Q could interfere with Y's reception of data. For MACA-P, we take a conservative approach and disallow a node with more than one master from participating in a parallel transmission/reception. Figure 5b shows the analogue for a node with more than one master sender.

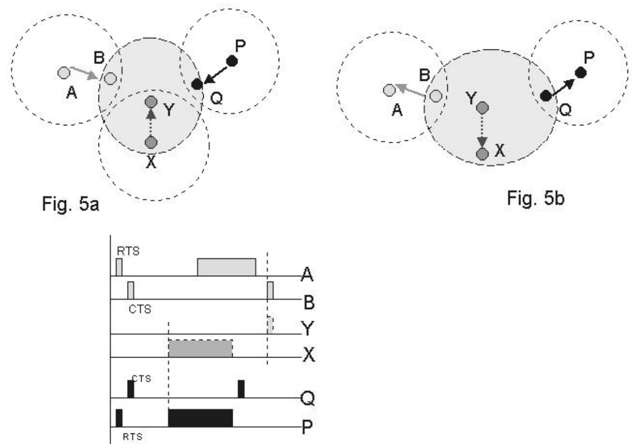

Next, consider Figure 6, where both the sender and the recipient of a proposed transmission schedule, has a master each. In Fig 6, assume that both B-to-A and P-to-Q transmissions have been scheduled. B is a master of Y while Q is a master of X. Y sends a RTS with the inflexible bit *set*. Since X has a master already, it must align any reception of its own with master Q. However, Y's proposed transmission schedule is aligned with its (Y's) master B. In this case, X will not be able to respond to Y's RTS, since the flexible bit indicates that Y is not willing to accept any re-alignment of its proposed schedule. What this example shows is that if both the sender and recipient have a master transmission and reception scheduled in their respective neighborhoods, then it is not possible for the sender/recipient pair to have a data transfer without collision.

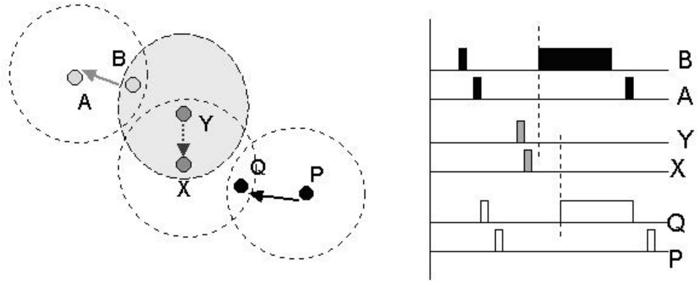

Fig. 6

4. IMPLEMENTATION AND PERFORMANCE EVALUATION

We have implemented MACA-P (with additional enhancements) as an additional MAC-layer module on the ns-2 simulator. Initial experiments with various multi-hop topologies show that MACA-P can lead to significant improvements (often as much as 150%) in throughput, simply by increasing the number of concurrent transmissions possible. Our simulation studies also led us to the design of several other enhancements, including an adaptive learning algorithm in the MAC layer, which are necessary for obtaining significantly higher throughput in real wireless environments (where two nodes may cause interference even if they are outside each other's reception range). For example, Fig. 7 shows the total network throughput with UDP packets in a circular "ring" topology as the number of nodes is varied.

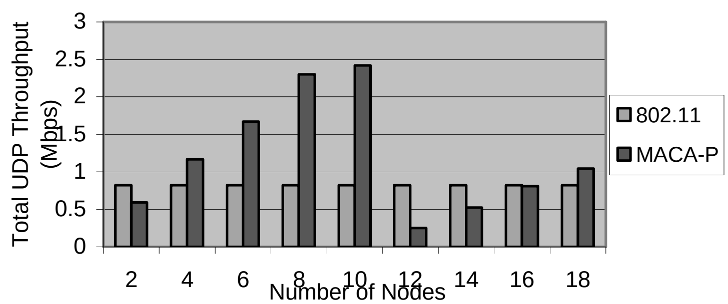

Fig. 7: Total Network Throuhput (802.11 vs. MACA-P)

REFERENCES

[1] IEEE 802.11 WG. Draft Supplement to Standard for Telecommunications and Information Exchange between System, Medium Access Control (MAC) Enhancements for Quality of Service (QoS), *IEEE 802.11e/D2.0*, November, 2001.

[2] C. Perkins, E. M. Belding-Royer, and S. Das. Ad Hoc On Demand Distance Vector (AODV) Routing. *IETF Internet draft, draft-ietf-manet-aodv-10.txt*, March 2002 (Work in Progress).

[3] D. B. Johnson and D. A. Maltz. Dynamic Source Routing in Ad Hoc Wireless Networks. In *Mobile Computing, Chapter 5, pages 153-181*, Kluwer Academic Publishers, 1996.

[4]S. Xu and T. Saadawi. Does the IEEE 802.11 MAC Protocol work well in multihop wireless ad hoc networks? *IEEE Communications Magazine*, June 2001.

[5] V. Bhargavan, A. Demers, S. Shenker and L. Zhang. MACAW : A Media Access Protocol for Wireless LAN's. *Proceedings of ACM SIGCOMM* '94.

[6] J. Weinmuller, H. Woesner, J. Ebert, A. Wolisz. Analyzing the RTS/CTS Mechanism in the DFWMAC Media Access Protocol for Wireless LANs. *IFIP TC6 Workshop on Personal Wireless Communications*, April '95.

[7] J. P . Monks, V. Bhargavan and W. Hwu. A Power Controlled Multiple Access Protocol for Wireless Packet Networks. *Proceedings of IEEE Infocom 2001*.

[8] Y.Ko, V. Shankarkumar and N.H. Vaidya, Medium Access Control Protocols using Directional Antennas in Ad Hoc Networks. *IEEE INFOCOM'2000*, March 2000.