3-2003

# MACA-P: A MAC for Concurrent Transmissions in Multi-hop Wireless Networks

Arup ACHARYA
*IBM T. J. Watson Research Center*

Archan MISRA
*Singapore Management University*, archanm@smu.edu.sg

Sorav BANSAL
*Stanford University*

# MACA-P : A MAC for Concurrent Transmissions in Multi-hop Wireless Networks

Arup Acharya[1], Archan Misra[1], Sorav Bansal[2]

[1]IBM TJ Watson Research Center, [2]Stanford University

arup@us.ibm.com, archan@us.ibm.com, sbansal@stanford.edu

**Abstract**: *This paper presents the initial design and performance study of MACA-P, a RTS/CTS based MAC protocol that enables simultaneous transmissions in multi-hop ad-hoc wireless networks. Providing such low-cost multi-hop and high performance wireless access networks is an important enabler of pervasive computing. MACA-P is a set of enhancements to the 802.11 DCF that allows parallel transmissions in many situations when two neighboring nodes are either both receivers or both 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. However, the data transmission is delayed by a control phase interval, which allows multiple sender-receiver pairs to synchronize their data transfers, thereby avoiding collisions and improving system throughput.*

## I. INTRODUCTION

We believe that one of the key enabling technologies for pervasive computing will be the emergence of low-cost, high-performance all-wireless access networks. Although the data transmission rates associated with the 802.11 family of standards [1] are increasing rapidly (rates of up to 108 Mbps are under development), the actual throughput in 802.11-based multi-hop wireless networks remains a major performance bottleneck [2]. The 802.11 CSMA-CA (Carrier Sense Multiple Access with Collision Avoidance) mechanism[3,4] and its variants (e.g., [5]) for distributed access to the shared channel was principally designed for the single-hop wireless LAN scenario, where nodes typically formed a clique and multiple simultaneous transmissions are not possible.

Multi-hop wireless networks are however *spatially diverse*, with different nodes able to communicate directly with different sets of one-hop neighbors. In this paper, we present the basic design[1] and initial performance evaluation of an 802.11-based MAC protocol called MACA-P (Medium Access via Collision Avoidance with Enhanced Parallelism). MACA-P aims to improve the overall network utilization by exploiting spatial diversity to increase the number of concurrent transmissions in a multi-hop environment.

## II. PROBLEM DEFINITION

The 802.11 Distributed Coordination Function (DCF) uses a 4-way distributed handshake mechanism

---

[1] The complete design and evaluation can be found in [6].

(RTS/CTS/DATA/ACK) to resolve contentions between peers. We need to make a distinction between the sender and recipient of a particular packet, and the transmitter and receiver associated with a specific transmission activity: the terms *sender* and *recipient* refer to the RTS/CTS/DATA/ACK transaction as a whole, while *transmitter* and *receiver* refer to a specific transmission activity within such a transaction.

We first discuss why the 802.11 MAC does not permit two nodes to transmit simultaneously that are either neighbors or have a common neighboring node. The following observation **(SRS)** must be supported by any wireless MAC to avoid collisions at a receiver: If any node is currently a transmitter, there can be only one receiver node in the transmitter's 1-hop neighborhood. Conversely, if any node is a receiver, only one node in its 1-hop neighborhood is allowed to be a transmitter.



Fig. 1    Fig. 2a    Fig. 2b

Fig. 2c

Consider Fig.1 where the transmission from X (to Y) would interfere would P's transmission to Q, since Q is within range of both X and P. Therefore, the two transmissions cannot occur simultaneously. Now consider 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: B's transmission to A would collide with P's transmission at Q (case 3) and A's transmission to B would collide with Q's transmission at B (case 4).

Next consider the case when two receivers are neighbors: packet transfers A-to-B and P-to-Q, as shown in Fig. 2a (case 1 in Fig 2c). Since A's transmission range does not include Q and P's transmission range does not include B, the two transmissions should be allowed to proceed in parallel, according to observation SRS. However, the 802.11 MAC does not support such parallel transmissions: when B sends a CTS in response to A's RTS, Q is aware that B has reserved the channel for $T_{CTS}$ interval. If now P sends a RTS to Q, Q cannot respond with a CTS to P since it is aware of an existing channel reservation that would overlap with P's data transmission[2]. A similar situation exists for the scenario in Fig2b (and case 2 in Fig 2c): although B and Q should be able to transmit to A and P respectively at the same time, 802.11 does not permit such parallelism, as the transmission of the first RTS prohibits the 2nd sender from sending out any RTS during the entire interval $T_{RTS}$.

The 802.11 MAC thus precludes the possibility of parallel communication by two neighboring nodes that are either both senders or both recipients. The key reason for this restrictive behavior is that a node reverts between a transmitter (tx) and receiver (rx) roles multiple times during a packet transfer without a precise, explicit knowledge of when these role reversals take place. Additionally, the 802.11 4-way handshake mechanism is effectively contiguous—once a node pair initiates a packet transfer, neighboring nodes cannot assume the role of a transmitter until the original 4-way handshake is complete. Thus, for a MAC to support parallel transmissions, we see that (1) two neighboring nodes must either both be transmitters or both be receivers, and (2) a gap ("control" gap), between the RTS/CTS exchange and the subsequent DATA/ACK exchange must exist to allow (a) other neighboring pairs to exchange RTS/CTS messages within the control phase gap of the first pair, and (b) subsequent pairs to align their DATA/ACK transmission phases with that of the first pair.

Note that the control gap is put in place by the first pair (A-to-B). A subsequent RTS/CTS exchange by a neighboring pair (P-to-Q) does not redefine the gap; subsequent pairs instead use the remaining portion of the control gap to align their data transmission with the first pair. MACA-P's principal goal is the enhancement of the 4-way handshake to allow parallel communication in cases 1 and 2 of Fig. 2c. As described in the next section, this is achieved through

---

[2] The data structure at each node that records channel reservations that it knows about, is called a NAV (Network Allocation Vector), as per the 802.11 MAC specifications.

enhancements to the basic IEEE 802.11 RTS/CTS mechanism to introduce an appropriate control gap.

## III. DESCRIPTION OF MACA-P

**Control Gap**: We add extra information in the RTS and CTS messages to explicitly delineate the intervals for both the DATA and ACK transmissions, thereby allowing neighboring nodes to know exactly when the two nodes associated with the DATA/ACK switch between tx and rx roles. The RTS and CTS control messages now specify two instants, both specified as the end of time intervals, relative to the time of receiving the associated control packet:

- $T_D$: the start time of DATA transmission occurs at the end of the interval $T_{DATA}$.

- $T_A$: the start time of ACK transmission occurs at the end of interval $T_{ACK}$.



Fig. 3

In figure 3 above, Q overhears the RTS sent from A to B. If Q has a packet to transmit, it will initiate a RTS whose $T_D$ is aligned with the start time of B's 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 the DATA and ACK transmission schedules.

**State of neighboring nodes**: Each node maintains the state of its neighboring nodes by overhearing the RTS/CTS exchanges from its neighbors. Consider Fig 3, where B initiates a RTS/CTS exchange with A. Since Q hears the RTS from B, it will update its NAV to indicate that B has scheduled a transmission to A. For each neighbor from which a RTS or a CTS has been overheard, a node maintains an entry in the NAV consisting of the neighbor's MAC address, neighbor's state (sender, recipient or idle), $T_D$ and $T_A$ time instants. If a node wishes to send a data packet, it must check that no entry in its NAV is marked as a recipient. (Otherwise, it would violate the SRS observation made earlier). Similarly, a node receiving an RTS cannot respond with a CTS if any entry in its NAV is marked as a transmitter. In addition to this basic test, nodes use the NAV to schedule an overlapping data transmission of its own. For example, in Fig.3, Q updates its NAV on

overhearing B's RTS to A, and then uses this information to schedule an overlapping transmission of its own (such that the $T_D$ and $T_A$ of the two transfers are respectively aligned).

**Inflexible Bit in RTS**: The RTS message is further enhanced to carry a bit which we call the inflexible bit, which indicates to the RTS receiver whether the transmission schedule proposed in the RTS message can be changed: if the bit is set, then this schedule cannot be changed. A sender attempting to align with a pre-existing $T_D$ must set this bit.

**Modification of $T_{DATA}$ and $T_{ACK}$ by CTS**: When a node receives a RTS where the inflexible bit unset, it may change the proposed schedule by modifying the $T_{DATA}$ and $T_{ACK}$ of the RTS, and sending back the modified values on the CTS. 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 $T_{ACK}$ interval (shown as t1 and t2) so that B's reception of data from A overlaps with Q's reception.



Fig. 4

**RTS' message**: Nodes update their respective NAVs on overhearing a RTS. If the schedule proposed in the RTS is modified by the RTS receiver (as discussed earlier), neighbors of the RTS sender would have an incorrect view of the transmission schedule unless notified of this change. To avoid this situation, a RTS sender always sends a gratuitous-RTS message (RTS') with updated $T_{DATA}$ and $T_{ACK}$ (received from the CTS), which informs all neighbors of the RTS sender of the updated schedule. A second use of RTS' is to cancel a prior schedule made through a matching RTS, when the sender did not receive a CTS from the intended receiver (of the RTS).

**Master Transmission Schedules**: The notion of aligning with an existing DATA/ACK transmission schedule leads to the concept of a master transmission. A master transmission schedule is one which is used by neighbors of either the sender or recipient to synchronize their own DATA/ACK transmissions. We

now state a key requirement for scheduling parallel transmissions via MACA-P:

> 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.

The rationale is as follows. 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. Master schedules can clearly be cascaded, with neighbors of a newly aligned transmission schedule treating this as their own master schedule.
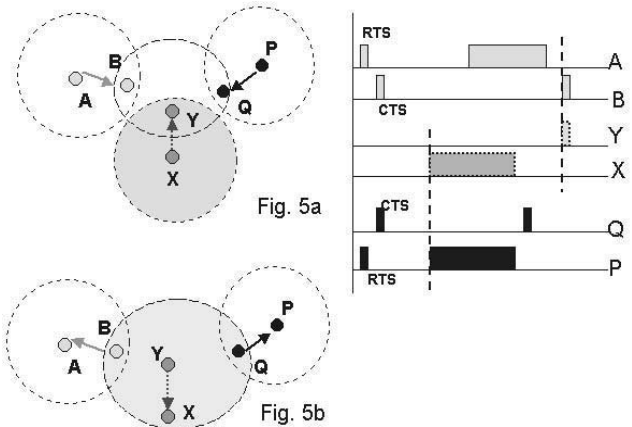


Fig. 5a



Fig. 5b

**Packet size for a master transmission**: An important implication of a transmission being a master is that all overlapping transmissions must transfer data packets whose size is less than that of the master transmission. Otherwise, the DATA of an overlapping transfer will

interfere with the ACK of the master. Therefore, a sender with no existing neighboring master applies MACA-P on a packet only if the packet size is greater than a certain threshold. Thus MACA-P's control gap is introduced only for "large" packets; for smaller packets, a master node uses the standard 802.11. Such an approach also implies that the larger MACA-P control packets do not cause unacceptably high overhead for "small" data packets.

## IV.   PERFORMANCE STUDY

We implemented MACA-P by extending the 802.11 DCF MAC available in the ns-2 simulator. The RTS/CTS/RTS' exchange was implemented with an extra 2-byte field $T_{data}$ (see Fig 3) in the MAC header.



**Figure 6: Concentric Ring Topology**

For an initial study of the performance gain of MACA-P, we present simulation results on a concentric ring scenario (illustrated in Figure 6), consisting of an equal number of nodes, placed in inner and outer concentric circles. While all the inner nodes form a clique, the outer nodes are radially aligned with the inner nodes. While 802.11 does not allow more than one transmission at any given time for the concentric ring, the number of simultaneous transmissions for MACA-P can be as high as N/2 (for N nodes). We measured the cumulative throughput for a traffic pattern with Inner Senders, with traffic going from each inner node to its corresponding outer node.



Figure 7: Base MACA-P/802.11 on Concentric Ring

Figure 7 shows the relative performance of 802.11 and MACA-P for packet sizes of 1536 bytes, and a MACA-

P control gap of 640 Bytes. While throughput in 802.11 never exceeds the channel capacity, MACA-P can give a performance improvement of almost 200% in some scenarios. We can also see a sharp drop in MACA-P throughput when N exceeds 10. This occurs because the higher node density precludes slave transmissions from occurring in parallel with scheduled master transmissions—unlike 802.11, sender nodes attempt concurrent transmissions by sending RTS and then suffer from timer backoff. This drawback can be rectified via an adaptive learning mechanism described in greater detail in [6].

## V.   CONCLUSIONS AND FUTURE WORK

This paper first showed how the limited support for concurrent transmissions in the 802.11 MAC acts as a key bottleneck in high performance packet forwarding in multi-hop wireless networks. MACA-P's design seeks to increase the feasible set of concurrent transmissions by introducing a control gap between the RTS/CTS and DATA/ACK phases. This gap allows two neighboring senders (recipients) to synchronize the start of DATA transmission (reception) and ACK reception (transmission). We described the use of additional control packets (such as RTS' and RTS-NACK) in MACA-P, and showed how these control packets could be used to align secondary transfers with a master transmission schedule. Preliminary simulation results presented here validate the operational correctness of MACA-P and show the potential for significant throughput improvement (at least in selected topologies). In subsequent work, we have identified and rectified additional performance drawbacks in MACA-P through the use of adaptive learning strategies and better physical layer capabilities.

## REFERENCES

[1] B. Crow, I. Widjaja, J. G. Kim and P. T. Sakai. IEEE 802.11 Wireless Local Area Networks. IEEE Communications Magazine, Sept '99.

[2] 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.

[3] P. Karn. MACA – a new channel access method for packet radio. ARRL/CRRL Amateur Radio 9th Computer Networking Conference, 1990.

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

[5] R. Rozovsky and P. R. Kumar. SEEDEX: A MAC protocol for ad hoc networks. Proceedings of ACM Mobihoc 2001.

[6] A.Acharya, A. Mishra and S.Bansal. MACA-P : A MAC for Concurrent Transmissions in Multi-Hop Wireless Networks. IBM Research Report RC22528.