

Medium Access Control in a Network of Ad Hoc Mobile Nodes with Heterogeneous Power Capabilities

Neeraj Poojary⁺, Srikanth V. Krishnamurthy*, Son Dao*

⁺University of California, San Diego, CA 92091
E-Mail: npoojary@ucsd.edu

*HRL Laboratories, LLC, 3011, Malibu Canyon Road, Malibu, CA 90265
E-Mail: krish@hrl.com, skdao@hrl.com

Abstract—MAC layer protocols for wireless ad hoc networks typically assume that the network is homogeneous with respect to the transmit power capability of individual nodes in the network. The IEEE 802.11 MAC protocol has been popular for use in ad hoc networks. We investigate the performance of this protocol when it is used in a network with nodes that transmit at various power levels. We show that overall throughput is lower than the throughput of a network in which all nodes transmit at identical power levels. In addition, low power nodes have a disadvantage in accessing the medium due to higher levels of interference from the high power nodes. We consider propagating the control messages generated by a node wishing to initiate communication to distant nodes so that they may forbear transmissions for some time, thereby allowing clear access to the initiating node. We find that the overhead incurred due to the additional message transmissions outweighs the potential gain achieved by propagating these messages. This indicates that the signalling mechanism used in the IEEE 802.11 standard or the variants thereof are not sufficient to alleviate the loss in throughput and the lack of fairness engendered by networks that are heterogeneous with regard to the transmit power capabilities of individual nodes.

I. INTRODUCTION

A Mobile Ad Hoc Network (MANET) is defined as “an autonomous system of mobile routers (and associated hosts) connected by wireless links—the union of which form an arbitrary graph”[2]. Mobile ad hoc networks are primarily deployed in the military and in disaster relief operations. These networks need to be rapidly deployable, easily reconfigurable and are devoid of any centralized support infrastructure. This usually necessitates protocols that are distributed in nature for functions such as routing and medium access control. The mobility of nodes further complicates the design of such protocols in many ways.

The Medium Access Control (MAC) protocol is critical to achieving a statistically equitable distribution of the available capacity between contending users. This is also important for ensuring that the QoS requirements of different users are satisfied. The design of a good wireless MAC protocol has to address challenges raised by (i) *mobility* of the nodes and (ii) an unreliable, time-varying *channel*. Mobility affects the MAC protocol because the set of users competing for capacity on the medium keeps changing. This makes it difficult to allocate bandwidth in

an equitable fashion. Time-varying effects such as fading and interference also make it difficult to administer medium access control on the channel.




Fig. 1. An Example to illustrate the Hidden and Exposed Terminal Problems

The simplest medium access control mechanism would permit users to transmit any time they desire. The ALOHA protocol [3] operates based on this principle. While this method works well for light loads, at heavier loads it results in a high incidence of collisions among simultaneous transmissions. The “Carrier Sense Multiple Access” (CSMA) protocols [9] attempt to address this problem. When CSMA is deployed, nodes listen on the channel to sense the carrier due to another node’s transmission. If a carrier is detected, the sensing node refrains from transmitting. CSMA protocols do not, however, deal adequately with the *hidden terminal* problem [10] or the *exposed terminal* problem. The hidden terminal and exposed terminal problems are briefly explained by means of an example in the rest of this paragraph. In the scenario depicted in Figure 1, say node A is transmitting to node B. Just sensing the channel will not make node C aware of the transmission because it is beyond the range of node A. It may therefore attempt to transmit at the same time, thus causing a collision at node B. This is the hidden terminal problem. Now if node B is transmitting to node A, node C will sense the transmission and defer its transmission even though its range is not large enough to cause a collision at node A. Thus the

¹“This work was done at HRL Laboratories, LLC.”

²“2000, ©HRL Laboratories, LLC., All Rights Reserved”

channel utilization efficiency suffers in this case. This is known as the exposed terminal problem. Note that the collisions occur at the receiver and not at the transmitter.

Phil Karn proposed the “Multiple Access with Collision Avoidance” (MACA) protocol [8] in 1990, based on the Apple LocalTalk protocol. MACA does away with carrier sensing. Instead the initiator and intended receiver of a data transmission exchange control messages to gain access to the channel before commencing the transmission. The initiator sends a *Request-to-Send* (RTS) message to the intended receiver. The receiver responds with a *Clear-to-Send* (CTS) message. The initiator starts data transmission upon receipt of the CTS message. The initiator includes in the RTS message, the amount of data it intends to transmit. This information is also included in the CTS from the receiver. Nodes that overhear the RTS will defer their transmissions long enough for the CTS to be successfully received at the initiator. (Note that there is an assumption of symmetry here. If a node, say node X, can hear a second node Y, then node Y can also hear node X). Likewise, nodes that overhear a CTS message will defer their transmissions for a period long enough to ensure that the ensuing DATA packet is successfully received by the receiver.

MACA does not have link-level acknowledgements of data transmissions. If a data transmission fails, retransmission has to be initiated by the transport layer. This can cause significant delays in the transmission of data. MACAW [5] extends the RTS-CTS-DATA exchange by introducing a link level acknowledgement (ACK) from the receiver after the successful reception of data. The use of an ACK complicates the exposed terminal scenario. An exposed terminal can benefit from an opportunity to transmit only if it can hear the ensuing reply (a CTS or an ACK). For example, going back to Figure 1, say node B is transmitting to node A. If node C elects to transmit an RTS to another node at the same time, it may not successfully receive the CTS from its intended receiver due to a collision with the transmission from node B. Also the transmission from node C may itself cause a collision in node B’s reception of an ACK from node A, thus rendering node B’s data transmission futile. To address this issue, MACAW utilizes a *Data Sending* (DS) message from the initiator before the actual DATA transmission. The DS message announces to the neighbours of the initiator that there was a successful RTS/CTS dialog and a DATA transmission is about to follow. Nodes that hear this message will then defer their transmissions long enough for the initiator to transmit the DATA packet and successfully receive the ACK message from its intended receiver.

The IEEE 802.11 MAC protocol [1], [6] is derived from MACA. It uses both a physical and a *virtual* carrier sense mechanism to determine when the medium is busy. It uses an RTS-CTS-DATA-ACK dialogue to accomplish data transmission. Each message in the dialogue contain duration information for the remainder of the dialogue. The virtual carrier sense is implemented in the form of a *network allocation vector* (NAV) maintained by each node. The NAV at each node maintains a value which represents a time instant that indicates the duration

upto which the medium is going to be busy due to transmissions from other nodes. The NAV is updated based on the duration information advertised in messages overheard by the node.

Note that all the above protocols assume the presence of symmetric links. This is valid for a network in which all nodes transmit at the same power level. The rapid spread of multifarious “wireless network enabled” devices jeopardizes the assumption of homogeneous power capability. An ad hoc network may comprise low power transducers, PDAs, handheld computers and larger file servers. These devices will have different transmit power capabilities. Some of them may be “tethered” to a power supply at all times and others may be dependent on battery power for long durations of time. In any event, it will be critical to ensure that the MAC protocol in use does not unduly favour devices that can transmit at higher power levels. In the next section, we describe some of the issues associated with using the IEEE 802.11 MAC protocol in a network in which different nodes may transmit at different power levels. In Section III, we consider some modifications to the IEEE 802.11 MAC protocol in order to address these problems. In Section IV we provide the performance results of simulations of these modifications, interpret them and compare them to the performance of the standard protocol. In Section V we summarize our work and present our conclusions.

II. PROBLEM DESCRIPTION AND SIMULATION OVERVIEW

In this paper, we will concentrate on the IEEE 802.11 MAC protocol.³ We investigate a network with heterogeneous power capabilities and the inequities and inefficiencies in the use of the medium in such a network. In our study, each individual device is assumed to have a constant transmit power, but this transmit power may be different for different devices in the network. The term *power capability* will refer to the power level that a node is capable of using for transmissions. The terms *homogeneous network* and *heterogeneous network* will refer to networks in which all nodes have, respectively, identical or non-identical power capabilities. Simulations are performed using the *NS* network simulator.

Typically MAC layer throughput is affected by routing and transport layer artifacts. For instance, TCP retransmissions and acknowledgements make it difficult to model the input load to the MAC layer accurately. Also, the use of stale routing information may manifest inappropriately as transmission failures at the MAC layer. In order to decouple these effects from our study of the MAC protocol, we extended the *NS* simulator to introduce a traffic generation agent immediately above the MAC layer. This agent has perfect information about the node’s neighbours at every instant. Every time a data packet is generated, it will be randomly destined for one of the nodes that are neighbours at that instant. The data packets are fixed-size packets of 1000 bytes each. The traffic model at each node has exponentially distributed packet inter-arrival times with the average rate λ being varied to vary system load. The mobility model

³“Specifically, we study the IEEE 802.11 MAC protocol in the framework of the Distributed Coordination Function (DCF)[1]”

is a *random waypoint* model with constant speed of 6mph between points and pause time 0.1 seconds. In other words, each node chooses a random direction in which it moves at a constant speed of 6mph for a random time. After this time, the node pauses for 0.1 seconds, then chooses a new direction at random and repeats the process.⁴ The simulation network is assumed to be deployed in a square region whose area may be varied so as to vary the geographical density of nodes in the network. The medium is assumed to be free of noise and any errors due to fading or interference other than the interference from other users in the network. The channel bandwidth is set at 2Mbps. The radio specifications are based on the AT&T *WaveLAN* with only the transmit power being varied. The same frequency band is used by all users in the network. Two or more simultaneous transmissions received by a node either result in a collision or *capture*. A transmission captures the medium when the received power due to that transmission exceeds ten times the received power due to any other simultaneous transmission. In order to quantify channel usage, we define throughput efficiency at each node as follows:

$$\eta = \frac{\text{total time spent in successfully transmitting data}}{\text{total simulation time}}$$




Fig. 2. MAC layer throughput for a homogeneous network




Fig. 3. MAC layer throughput for a heterogeneous network

In Figure 2, we show the throughput of two homogeneous networks at power levels of 0.14W and 0.56W and an average

⁴“Note that this traffic model has been chosen for simplicity and any generic traffic model may be expected to result in similar performance”

offered load (λ) of 1000 packets/second at each node. The total number of nodes in the network is fixed at 40 and the node density is varied by varying the area of the square grid used in the simulations. The parameter along the X axis indicates the length of the square grid. We note that at very high densities (grid length < 500m), both networks perform virtually identically. This is because the nominal transmit range at the smallest power level (0.14W) is about 205m, which implies that all the nodes are sharing a single channel almost all the time. The same is true for the network operating at higher transmit power level. However as the grid area increases, we notice that the network in which nodes transmit with the lower transmit power does much better than the other network. This is because the lower transmit power increases network capacity by increasing spatial reuse of the spectrum. This is in marked contrast to what happens in the operation of a heterogeneous network with nodes operating at two transmit power levels. Figure 3 depicts a network of 40 nodes with half of the nodes transmitting at 0.14W and the other half transmitting at 0.56W. We note that the low power nodes suffer a 50% degradation in throughput efficiency in comparison with their performance in a homogeneous network in which there are no high power nodes. A similar trend was discerned at lighter traffic loads of 100, 50 and 10 packets per second per node. Clearly the low power nodes are being overwhelmed by the higher power nodes in accessing and using the channel successfully.

The IEEE802.11 MAC protocol uses a reservation scheme based on the exchange of Request to Send (RTS) and Clear to Send (CTS) messages between a source and destination as explained in Section I. For a homogeneous network, on average this mechanism works satisfactorily in ensuring a fair allocation of the channel. But in a heterogeneous network, when a low-power node attempts to reserve the channel for a subsequent data transmission, it may not be heard by high-power nodes that are potentially close enough to disrupt its data exchange. For instance in Figure 4, node A is a high-power node, node B is a low-power node and node C is another high-power node. Node C may potentially interfere with the reception of data at node B in spite of the RTS/CTS exchange between nodes A and B since it is unable to hear the CTS message from node B.




Fig. 4. Failure of RTS/CTS in heterogeneous power environment

As the use of ad hoc wireless networks becomes more ubiquitous, the assumption of uniform transmit power capability will be less and less valid. *Network-enabled* devices with disparate

power capabilities will be pressed into service and some of them may not be able to operate satisfactorily in the network due to unfairness in the MAC protocol. Therefore MAC protocols will need to be designed to be more sensitive to the different transmit power capabilities of devices.

III. MODIFYING THE RTS/CTS RESERVATION SCHEME




Fig. 5. Success rate for DATA transmissions

As shown in Figure 4, a successful RTS/CTS exchange will not guarantee successful transmission of data in a heterogeneous network. This is borne out by the graph in Figure 5 which shows the percentage of successful DATA transmissions after a successful RTS/CTS exchange for homogeneous networks operating at different power levels as well as for a heterogeneous network comprising two types of nodes. We see that in the heterogeneous network, the degradation in DATA transmission success rate for the low power nodes can be as high as 30%. As the network density goes down, there are fewer neighbours that can interfere with the DATA transmission. Hence the success rate of DATA transmissions from low power nodes improves, but it is still far below the success rate for high power nodes.




Fig. 6. Using BW_RES message to propagate CTS

A possible solution to prevent this degradation is to extend the reach of the RTS/CTS reservation so that all high power nodes that could potentially interfere with the DATA transmission are made aware of the reservation. One way to extend the reach of the RTS/CTS reservation mechanism without boosting transmit power is for nodes that hear the CTS message to propagate it again. For instance, let us revisit the earlier scenario in Figure 6,

now with an additional node D in the picture. Say node D broadcasts the CTS it hears from node B. The CTS from node B could not reach node C, but the broadcast message from node D will reach node C and node C will then defer its own transmissions during the ensuing DATA/ACK sequence between nodes A and B. Note that in most cases a single broadcast of the CTS will not suffice. At the same time, we obviously do not want to keep broadcasting the CTS ad infinitum. We need to propagate the CTS a reasonable number of times to ensure adequate reach for the reservation without causing too much overhead. Adequate reach means covering a radius equal to the transmission radius of the highest power node in the network. Assuming that the network is not partitioned, and transmit ranges are normalized such that the lowest power node in the network has range 1 unit and the highest power node in the network has range N units, we have the following result:

Lemma 1: With the nodes distributed along a straight line such that the distance between any two neighbours is less than one unit (no partitioning, in some sense) and each transmission having a range of one unit, and assuming that among the nodes that hear a transmission, the node that is furthest from the transmitting node will retransmit the message, we have:

A message needs to be propagated $2N - 1$ times to ensure that it is heard at a distance N from the originator of the message.

Proof: Let *transmission* refer to both, the original transmission of a message or subsequent retransmissions by nodes that hear the message. Say the originator of the message is at the origin and transmissions occur along the positive X axis. The first transmission covers one unit. Say m transmissions of the message are needed to ensure a reach of k units, specifically a distance $k + \delta$. If there is a node in $[k, k + \delta]$, its transmission of the message will cover distance $k + 1$. In this case we need $m + 1$ transmissions to cover $k + 1$ units. Suppose now that there is no node in $[k, k + \delta]$. Then \exists a node n_1 in $(k + \delta, k + 1]$ (else there will be two neighbours with the distance between them being greater than 1 unit) and a corresponding node n_2 in $(k - 1 + \delta, k]$ such that $d(n_1, n_2) \leq 1$. Then one transmission from n_2 followed by one transmission from n_1 will be required to cover distance $k + 1$ units. In this case we require $m + 2$ transmissions to ensure that the message covers a distance of $k + 1$ units. Thus the original transmission covers one unit and for each additional unit of coverage, two additional transmissions are required. Therefore, the minimum number of transmissions required to ensure coverage of N units is $2N - 1$.

An illustration of the lemma for $N = 3$ is provided in Figure 7.




Fig. 7. An example of five transmissions required when $N = 3$




Fig. 8. The BW_RES frame

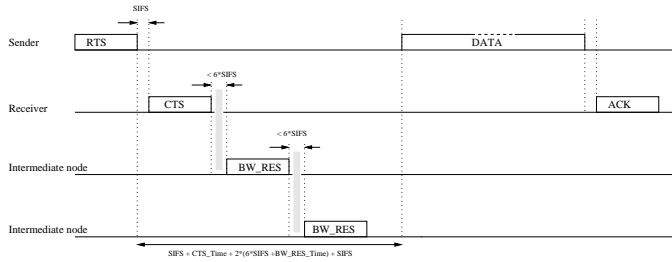


Fig. 9. The modified reservation scheme

We extend the RTS/CTS mechanism by adding another message called BW_RES which is essentially a broadcast propagation of the CTS message. The BW_RES message format is shown in Figure 8. It is similar to the RTS message format except that the frame control field has an additional attribute called *seqno*. This is a sequence number intended for the detection of duplicate BW_RES messages that may be received when a standard flooding algorithm is used to propagate these messages. A similar *seqno* attribute is added to the frame control field of the CTS frame. For our simulations, the To DS, From DS and More Frag bits of the frame control field (see [1]) were overloaded in the CTS and BW_RES messages to indicate a 'Time to Live' (or *ttl*) for the message. The *ttl* is initially set to $2N - 1$ when the CTS is sent out and then decremented appropriately by each node that retransmits the message (in the form of a BW_RES), depending on the transmit power level of the node. Each node that hears the CTS and determines that it needs to send a BW_RES, waits a random number (between 0 and 6) of *short interframe space* (SIFS) [1] units before transmitting the BW_RES message. This is to minimize collisions due to multiple simultaneous BW_RES transmissions from neighbours that hear the same CTS message. The complete RTS-based data transfer sequence with the extended scheme is depicted in Figure 9.

IV. EFFECTS OF THE MODIFICATIONS

We performed simulations with the modified RTS/CTS mechanism for a two-level heterogeneous network of 40 nodes with half the nodes operating at a transmit power level of 0.14W and the other half at 0.56W. Results were obtained at different network densities and for different offered loads. Owing to space constraints, graphical results for all offered loads that were tested are not presented here. Initially, the system was configured such that the CTS messages would be rebroadcast 2 times (since $N < 2$). The results for an average offered load of 1000 packets per second per node are shown in Figure 10 and Figure 11. Figure 10 shows the throughput of the heterogeneous system along with that of the homogeneous networks operating at each power level. We see that overall system performance has




Fig. 10. Throughput for modified reservation scheme





Fig. 11. DATA transmission success ratio for extended reservation scheme

actually degraded significantly. Though the extended RTS/CTS mechanism brings about fairness in the sense that the difference in throughput between the high power and low power nodes is not as high, the additional message overhead probably exceeds the benefit accrued in propagating the CTS messages. We also note from Figure 11 that while there is no significant change in the DATA transmission success rate for low-power nodes, the rate for high-power nodes goes down by about 5 to 10 percent.

Fig. 12. Heterogeneous network throughput comparisons for different reservation schemes ($\lambda = 1000$)




Fig. 13. Performance of modified RTS/CTS mechanism at low load ($\lambda = 20$)

In order to reduce the overhead involved in sending the BW_RES messages, we considered a location information system like GPS to optimize the transmission of BW_RES messages. We made changes to the simulation software so that the MAC layer at each node was aware of the node's position. This position information was then incorporated into the CTS and BW_RES messages. Three optimizations were performed based on the location information now contained in each message in conjunction with the node's awareness of its own location:

- Among nodes that overhear a CTS message, nodes that are further from the sender of the CTS are more likely to transmit a BW_RES first.
- If a node that receives a BW_RES message is already as far from the sender of the CTS as the range of the strongest node in the network, it will not propagate the BW_RES further even if the message has a non-zero time-to-live.
- If a node that receives a BW_RES determines that the message is being propagated back toward the sender of the CTS (i.e. if the receiver of the BW_RES is closer to the sender of the CTS than the sender of the BW_RES), it will not propagate the BW_RES message further even if the message has a non-zero time-to-live.

Surprisingly, we found that these modifications only give marginal benefit at high densities and actually degrade performance further at low densities. We surmise that one reason for the degradation is that adding location information to the CTS and BW_RES messages increases the respective packet sizes by almost fifty percent.

We also considered a scenario in which a CTS message originating at a low-power node is propagated only once by nodes that hear it. Figure 12 provides a comparison of results for the various modifications at saturation load. Figure 13 provides a similar comparison at a lighter offered load. We note that for saturation load, throughput for the heterogeneous network keeps worsening as the overhead in the form of BW_RES messages increases. However at relatively lighter loading, throughput for a heterogeneous network in which a CTS message is propagated twice (i.e. 2 BW_RES) is sometimes better than and never worse than that for a heterogeneous network in which a CTS message is propagated only once. However for every modification, the throughput is still worse than the throughput for the standard protocol.

Thus the modifications to the IEEE 802.11 protocol to extend the reachability of the CTS messages by means of flood-

ing actually degrade the performance of the protocol. We are considering intelligent dissemination mechanisms whereby the gain achieved in avoiding collisions actually outweighs the loss incurred in terms of overhead.

V. CONCLUSIONS

We have shown, in the context of the IEEE 802.11 protocol, that heterogeneous networks suffer significant degradation in performance in comparison with homogeneous networks. This degradation is primarily caused by poor medium access for the low-power nodes in the network. It is clear that the MAC protocol has to be changed to make medium access more efficient in a heterogeneous network. We have investigated the feasibility of one such mechanism. This involves extending the RTS/CTS mechanism by adding another message type, to ensure that the reservation information is propagated a greater distance than before. We have found that the overhead due to the additional messages outweighs the potential benefits of the greater reach of the reservation mechanism. Hence other mechanisms, possibly involving a different kind of reservation scheme, will need to be investigated.

REFERENCES

- [1] "Draft International Standard ISO/IEC 8802-11, IEEE P802.11/D10, Jan 1999", *LAN MAN Standards Committee of the IEEE Computer Society*.
- [2] "<http://www.ietf.org/html.charters/manet-charter.html>" *IETF MANET Working Group*.
- [3] Abramson, N., "The ALOHA System - Another Alternative for Computer Communications", *Proceedings of the Fall Joint Computer Conference*, pp. 281-85, 1970.
- [4] Bertsekas, D. and Gallager, R., "*Data Networks*", Prentice-Hall Inc., 1987 (2nd Ed. 1992).
- [5] Bhargavan, V., Demers, A., Shenker, S. and Zhang, L., "MACAW: A media access protocol for wireless LANs", *Proceedings of ACM SIGCOMM '94*, pp. 212-25, ACM, 1994.
- [6] Biba, K., "A hybrid wireless MAC protocol supporting asynchronous and synchronous MSDU delivery services", *Tech. Rep. Paper 802.11/91-92*, IEEE 802.11 Working Group, 1992.
- [7] Haas, Z. J. and Tabrizi, S., "On some challenges and design choices in ad-hoc communications.", *IEEE Military Communications Conference. Proceedings. MILCOM 98, vol. 1*, 1998.
- [8] Karn, P. "MACA - a new channel access method for packet radio", *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, pp. 134-40, ARRL, 1990.
- [9] Kleinrock, L. and Tobagi, F. A., "Packet switching in radio channels: Part I - Carrier Sense Multiple-Access modes and their throughput-delay characteristics", *IEEE Trans. Commun.*, vol. COM-23, no. 12, pp. 1400-1416, 1975.
- [10] Tobagi, F. A. and Kleinrock, L., "Packet switching in radio channels: Part I - The hidden terminal problem in carrier sense multiple-access modes and the busy-tone solution", *IEEE Trans. Commun.*, vol. COM-23, no. 12, pp. 1417-1433, 1975.