

Improving the MAC Layer Performance in Ad hoc Networks of Nodes with Heterogeneous Transmit Power Capabilities

Vasudev Shah*, Srikanth V. Krishnamurthy* and, Neeraj Poojary⁺

*University of California, Riverside, CA 92521

Email: {vasu, krish}@cs.ucr.edu

⁺Magis Networks, Inc.

12651 High Bluff Drive, San Diego, CA92130

Email: {npoojary@magisnetworks.com}

Abstract: - The performance of the IEEE 802.11 MAC protocol has been shown to degrade considerably in an ad hoc network with nodes that transmit at heterogeneous power levels. The main cause of this degradation is the potential inability of high power nodes to hear the RTS/CTS exchanges between nodes when at least one node involved in communication is a low power node. The propagation of the CTS message beyond the one-hop neighborhood of two communicating low power nodes was considered in our prior work in an attempt to alleviate this effect. However, this resulted in excessive overhead and further degraded the performance at the MAC layer.

In this paper we consider two techniques to reduce the overhead incurred due to the aforementioned propagation of the CTS message: (a) the use of an intelligent broadcast scheme and (b) the reservation of bandwidth for the sequential transmission of multiple data packets with a single RTS/CTS exchange (and propagation as needed). These techniques require changes only at the MAC layer. We find, by means of extensive simulations, that the techniques provide a significant improvement over the original 802.11 MAC protocol in the considered *power heterogeneous* ad hoc network. The overall throughput improves by as much as 12 % and the throughput of the low power nodes improves by up to 14 % as compared to the IEEE 802.11 MAC protocol. Furthermore, the schemes find applicability even in homogeneous networks as they reduce the number of *false link failures* that arise when the IEEE 802.11 MAC protocol is used, by about 20 %. We conclude that the schemes together offer a simple yet effective and viable means of performing medium access control in power heterogeneous ad hoc networks.

I. INTRODUCTION

Ad hoc networks are likely to consist of multifarious devices with differing capabilities. One could envision low power sensor nodes, wireless hand-held devices, laptops and bigger and more powerful wireless devices housed in vehicles, all integrated into a single network. In such a heterogeneous network, different nodes have different power capabilities, both in terms of their battery lives as well in terms of their achievable transmission range. In this work we consider a network in which the different nodes differ in terms of their achievable transmission range, i.e.,

the transmission range of the *lower power nodes* is smaller than that of the *higher power nodes*. We also use the terms *homogeneous* and *heterogeneous* to refer to networks in which all nodes have, respectively, identical or non-identical power capabilities in terms of transmission range. A central challenge in the design of the medium access control protocols is achieving the evenhanded distribution of the available capacity between contending nodes; the problem becomes even more challenging when nodes transmit at different powers. We have shown in our previous work [1] that the performance of IEEE 802.11 MAC protocol (popularly advocated for ad hoc networks), when it is used in a network with nodes that transmit at diverse power levels, degrades significantly when compared with its performance in a homogeneous network. The primary reason for this degradation was that high power nodes could not overhear the exchange of the low power Request to Send (RTS) and the Clear to Send (CTS) messages and as a result, the high power nodes initiated transmissions while low power nodes were in communication. This in turn caused an increase in the number of retransmission attempts by low power nodes, thereby increasing the effective traffic load in the network. Together, these effects caused the performance of the IEEE 802.11 MAC protocol to degrade significantly. In our previous work [1], we attempted to *broadcast* the CTS message within a neighborhood of a communicating low power node. Nodes that heard the CTS message propagated the message further in an attempt to notify high power nodes in the vicinity of the on-going communication. However, the schemes considered did not prove fruitful since they increased the amount of overhead generated due to the propagation of the CTS message¹. As a consequence the network throughput further degraded as compared to the legacy IEEE 802.11 MAC by about 20%. While the problems that arise with the use of the IEEE 802.11 MAC protocol were identified in [1], no viable solutions for alleviating the same were

This work is supported by DARPA FTN Grant no: F30602-01-2-0535 and NSF CAREER Grant no: ANI-0237920.

¹ The propagated CTS message beyond the first hop was referred to as the BW_RES message for Bandwidth Reservation message in [1].

suggested; the considered schemes did not clearly achieve the above objective.

In this paper, we attempt to reduce the overhead incurred due to *the propagation* of the CTS message (as considered in [1]). We propose the use of two techniques: (a) the use of an intelligent broadcast scheme to quell unnecessary broadcasts and (b) reserving the bandwidth for multiple data packets with a single RTS/CTS exchange / propagation. The latter reduces the frequency with which CTS messages are generated and propagated and thereby might be expected to provide performance benefits. We find that these methods do in fact improve the performance tremendously and improve the medium access control throughput by approximately 12 % in total as compared to the legacy IEEE 802.11 MAC. Furthermore, they also improve the fairness in terms of providing a better share of the medium to low power nodes.

There have been a number of previously proposed techniques that help in reducing the overhead due to broadcast operations [2][3]. We simply borrow from a scheme proposed in [2], wherein a node is precluded from further broadcasting the CTS (or BW_RES²) message, if it overhears broadcasts of the message from a preset number of its neighbors.

In order to incorporate the second technique, if a node were to have multiple data packets for the same neighbor in its queue, it would send a single RTS message to reserve bandwidth for the sequential transmission of the multiple data packets. We evaluate the performance of our approaches by extensive simulations. The principal metrics of interest are the overall throughput, the throughput of the low power nodes, the data success rate of low power nodes and total number of link failures, which also includes false link failures. Our scheme improves the throughput of nodes transmitting at low power by up to 14%. We also show that our scheme reduces the number of link failures by 20%. We also show that the performance of the low power nodes in terms of successful data transmission after a successful RTS/CTS exchange improves by as much as 20%. Although the scenarios considered are limited, our results suggest that the proposed techniques offer a simple yet viable and effective option for use at the MAC layer in power heterogeneous wireless ad hoc networks.

The rest of the paper is organized as follows: In section II we describe the related previous work. In section III we describe our approach in detail. In section IV we present our simulation results and discuss them at length. Our conclusions form Section V.

II. RELATED WORK AND BACKGROUND

The IEEE 802.11 MAC protocols and possible derivatives have been popularly considered for use in ad hoc networks [1] [5] [6] [7] [8]. In our previous work [1] we investigated an ad hoc network wherein nodes had heterogeneous power capabilities and quantified the inefficiencies in the use of the IEEE 802.11 medium access control protocol in such networks. In our studies, we assumed that nodes transmitted with a constant power level, but this transmit power level was different for different devices in the network. The work is applicable in power managed ad hoc networks as well, wherein nodes choose temporary quasi-static power levels for transmissions over prolonged periods of time. In [1] a heterogeneous network, with nodes operating at two power levels (0.56 W and 0.14 W) was considered. It was shown that the low power nodes suffer up to 50% degradation in throughput as compared to the high power nodes under various conditions of network load. As mentioned earlier, this was primarily due to the fact that low power transmissions were often killed by transmissions from high power nodes that were unable to hear the RTS/CTS exchange between the low power nodes. This effect is depicted in Figure 1. The RTS/CTS exchange between nodes A and B is not overheard by node H. Thus, it is possible that while the data exchange between nodes A and B is in progress, node H could begin its own transmission, thereby causing a collision at node B.

Our previous attempt [1] to solve this problem was by means of a CTS propagation technique using a standard *flood-type* broadcast algorithm. Nodes that hear a CTS message propagate it further (up to a distance that is determined by the ratio of the range of the high power nodes to the range of the low power nodes [1]) in the form of Bandwidth Reservation (BW_RES) control packets (Figure 1). The objective of this broadcast was to let the high-power nodes in the neighborhood (such as H) know about the ongoing RTS/CTS exchange between the sender and the receiver so that they would in turn inhibit their own transmissions for the duration specified in the BW_RES packet. One might expect that this would have improved the performance at the medium access control layer. However, we show in [1] that such a broadcast does not work as expected since the overhead incurred in propagating the BW_RES control messages outweighed the potential gains achieved in terms of reducing the number of collisions. The scheme caused a further degradation in terms of throughput instead of an improvement. To the best of our knowledge our study of IEEE 802.11 MAC for a power heterogeneous ad hoc network is the first of its kind.

² To be consistent with the nomenclature used in [1], henceforth, we will refer to the propagated message as the BW_RES message.

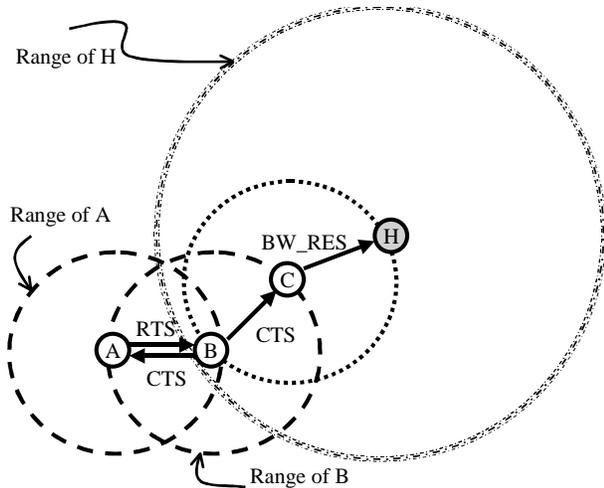


Figure 1: Propagation of the BW_RES message

There has been previous work on the use of power control in ad hoc networks. In [7] and [8], the authors examine the use of power control with the IEEE 802.11 MAC protocol. In both [7] and [8], it is assumed that the *maximum default power range* of each node in the network is the same i.e., the network is homogeneous. For communications the nodes might choose to use a power level that is lower than the maximum power range. However, both papers consider the transmissions of the RTS and CTS messages with the maximum power so as to reach all nodes that are within this maximum power range. Our work considers networks wherein the maximum default power range is different for different nodes, i.e., the network is heterogeneous.

In our previous work [1], we considered the propagation of the BW_RES message within a neighborhood of a low-power communication via a flood-type broadcast. Various broadcasting techniques have been studied and compared in [2]. Since, we wanted to implement a scheme that requires changes only at the MAC layer, we found that the counter based scheme proposed by [3] was the most suitable for our work. With this scheme, the MAC layer at a node requires no information with regards to the neighborhood of the node from the routing layer. We discuss this scheme with regards to our work in detail in section III. Due to space constraints we do not discuss the legacy IEEE 802.11 MAC protocol; the details may be found in [5].

III. THE MAC LAYER ENHANCEMENTS

In this section, first, we briefly revisit the BW_RES message (as in [1]) and discuss how the neighbors of a communicating node propagate this message. Later, we discuss the two techniques that we consider for improving the medium access efficiency. We also provide details with

regards to the specific changes needed to the existing IEEE 802.11 MAC packet format.

The BW_RES Message: As mentioned earlier, the key idea that we use in order to prevent high power nodes from initiating transmissions that can collide with a lower power communication is to have neighbors that hear a CTS message broadcast the message to reach nodes that further away (Figure 1). The new propagated message is called the BW_RES message. This BW_RES propagation notifies high power nodes that are further away that a low power communication is in progress in the vicinity and that they should inhibit their transmissions for the time period specified in the BW_RES message. The modified reservation scheme is shown in Figure 2.

Frame Format: The modified CTS and BW_RES packet formats are shown in Figure 3. The BW_RES message format is similar to the RTS message format except that that the frame control field has a few additional attributes: (a) a sequence number and (b) the originator address field which contains the MAC address of the node that initiated the communication (RTS sender). These fields are used for the detection of duplicate BW_RES messages that may be received. The sequence number field is also added to the CTS message for the same reason. We use the To DS, From DS and MORE bits in the frame control field [5] in the CTS and BW_RES messages to indicate a Time to Live (TTL) for the BW_RES packets. In our experiments we use TTL = 1 as we found that propagating the BW_RES messages to a distance of up to two hops resulted in the best overall throughput in the scenarios that we considered. In practice this is a system parameter that depends on the ratio of the maximum possible power range to the minimum possible power range [1].

Reserving bandwidth for sequential multiple transmissions: In order to reduce the overhead due to the broadcast of the BW_RES message, we attempt to reduce the frequency of such broadcasts. This may be achieved by using a *single* RTS/CTS/BW_RES initiation *for multiple* sequential DATA/ACK exchanges. The multiple DATA packets are in fact independent packets (they have their own fields including a separate checksum field for each packet) (Figure 2). Before sending an RTS, a node checks its interface queue (between and the network and the MAC layers) for other DATA packets with the same MAC layer destination address, and moves them to the MAC layer and buffers it along with the original packet to be sent. A node may also select multiple DATA packets already buffered in the MAC layer destined for the same neighbor. The RTS message now attempts to reserve the channel for these multiple (say N) data packets at the same time. When a node does not have additional packets for the destination (i.e., there is simply a single packet that can be transported at that given time), it reserves the channel for the single

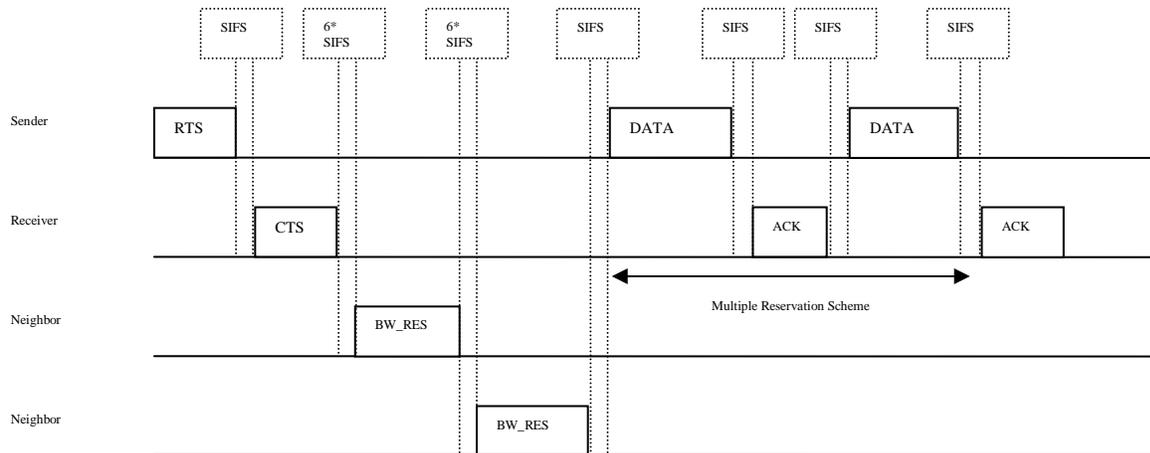


Figure 2 Modified Reservation Scheme

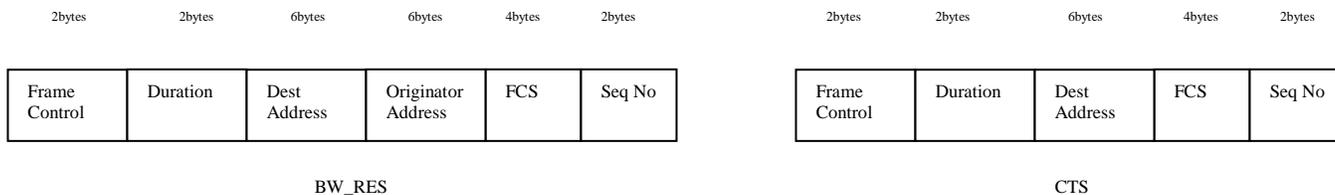


Figure 3 Modified Packet Structures

data packet; however, note that even in this case, neighbors propagate the BW_RES messages as discussed earlier. Clearly if a node is allowed to reserve bandwidth for a large number of sequential transmissions, it could lead to unfairness in the network. On the other hand, if this number is too small, there might not be a significant reduction in the number of generated BW_RES messages. We find by simulations that reserving bandwidth for just two sequential transmissions does provide the necessary benefits in terms of reducing the volume of BW_RES messages. In all our future discussions we refer to this technique as the **multi-reservation technique**.

Eliminating unnecessary broadcasts: In order to further reduce the number of broadcast BW_RES messages that are generated, we attempt to limit the number of such broadcasts by means of an intelligent broadcast technique that attempts to eliminate unnecessary broadcasts. The broadcast scheme that we employ is a derivative of a scheme proposed in [3] as mentioned earlier and is not entirely new. With this scheme, a node that has scheduled a BW_RES broadcast maintains a counter to record the number of BW_RES broadcast messages, corresponding to the same CTS message, that are overheard due to broadcasts from neighboring nodes. When the value indicated by the counter exceeds a preset fixed threshold T (a system parameter) it revokes its own BW_RES broadcast. The

idea behind implementing this scheme is that when a node overhears multiple copies of the broadcast message that it is scheduled to transmit; it would suggest that the node is in a highly dense neighborhood and the additional coverage that the node would achieve by performing its own broadcast is likely to be insignificant. The scheme reduces the number of unnecessary broadcasts to a significant extent [3]. We find by simulations that $T=3$ is a good value to choose; the coverage is maintained and a significant reduction in the number of unnecessary broadcasts is achieved. We refer to this technique as **the intelligent broadcast technique** in future discussions.

Other Details: The action taken upon transmission failure is as in the IEEE 802.11 standard; if the RTS sender does not receive a CTS response, it backs off and attempts to retransmit the packet after a back off period. As per the IEEE 802.11 MAC specifications, there is an upper bound on the number of retransmission attempts for RTS and DATA packets (4 and 7 respectively). After the pre-specified number of attempts the corresponding packets are dropped. If the sender does not receive an ACK for a specific DATA packet among the multiple sequential attempted transmissions, the node attempts to retransmit only the particular failed DATA packet at a later time. During the retransmission attempt, the sender again tries to reserve bandwidth for any additional packets (meant for the same MAC layer destination) in

addition to reserving bandwidth for the one being retransmitted.

IV. PERFORMANCE EVALUATION

Simulation Models: We study the performance of the proposed techniques using the event driven packet network simulator, ns2, specifically the version ns2.26 [4]. In order to be consistent with [1], we chose nodes that can transmit with either of two transmission power levels, 0.56 W or 0.14 W for simulating a heterogeneous ad hoc network; 50% of the chosen nodes were high power nodes and the other 50% were low power nodes. The physical layer is based on the IEEE 802.11 MAC specifications. All of our proposed enhancements require changes only at the MAC layer; any protocol may be used at the routing layer. In all of our experiments we use the AODV (Ad-hoc On-demand Distance Vector) protocol for performing routing [9]³. For all our experiments the packet size generated by the traffic agent is set to 1000 bytes. We use TCP as the transport layer protocol and the application is FTP⁴. The FTP agent generates packets at a rate of 1000 packets per second. Nodes move in accordance to a modified version of the *random waypoint* model with a constant speed of 6 mph between the chosen points and a pause time 0.1 seconds is assumed. The constant speed is chosen in light of the recent results that suggest that a choice of random speeds is inappropriate in terms of depicting mobility [11]. The simulated network is deployed in a square region whose area is varied so as to vary the geographical density of nodes in the network. The total number of nodes in the network is fixed at 40. In all our figures, the parameter along the abscissa indicates the length of the square grid.

The medium is assumed to be free of noise or any errors due to fading as in other previous work [1] [7] [8]. The channel bandwidth is set at 2 Mbps. All MAC control packets are transmitted at 1 Mbps and data is transmitted at 2 Mbps so as to conform to the IEEE 802.11 standards [5]. The radio specifications are based on the AT&T *Wave LAN* with only the transmit power being varied to form a heterogeneous ad hoc network. All the users in the network use the same frequency spectrum. For our multi-reservation scheme we choose to perform reservations for *two* packets with a single control message exchange⁵. Our metrics of interest are the data

³ We perform experiments with the Dynamic Source Routing protocol (DSR) [10] and the general behavior of the schemes that we implement remains the same.

⁴ We conducted extensive simulations using various load conditions, traffic types and numbers of connections; our results are consistent in terms of behavior for the various simulation environments. Due to space constraints we discuss the results for only specific scenarios.

⁵ We varied the limit on the number of data packets that can be transferred with a single RTS/CTS exchange. With an increased number the performance enhancements that we observe are not as significant. Longer transmissions lead to unfairness and hence we

restrict ourselves to at most reserving channel time for two packets with a single RTS/CTS exchange.

success rate and the throughput efficiency of each node. The throughput efficiency is defined as follows:

$$\text{Throughput Efficiency (\%)} = \left(\frac{\text{Total time spent in successfully transmitting data}}{\text{Total simulation time}} \right) \times 100.$$

We compare the performance of the following schemes: **Case (a)** the legacy IEEE 802.11 MAC protocol; **Case (b)** the MAC protocol with BW_RES propagation combined with the intelligent broadcast technique; **Case (c)** the MAC protocol with BW_RES propagation combined with the multi-reservation technique and, **Case (d)** the MAC protocol with BW_RES propagation combined with both the intelligent broadcast technique and the multi-reservation technique. Furthermore, we perform experiments wherein either: **Variation I:** only the low power nodes perform the BW_RES propagation and use the proposed techniques when initiating a communication or, **Variation II:** all the nodes regardless of whether they are high power nodes or low power nodes do so.

Simulation Experiments and Discussion: To begin with, we examine the performance of our techniques (Case (d)) with both Variation I and Variation II. We consider 6⁶ connections established between arbitrarily chosen source destination pairs. We compare the performance with that of the legacy IEEE 802.11 MAC protocol (Case (a)). From Figure 4, we observe a significant improvement in terms of the data success rate of low power nodes with both Variation I and Variation II. Notice that the low power nodes see an overall improvement of up to 22% with Variation I (and up to 17 % with Variation II) with our schemes as compared to scenarios with the legacy IEEE 802.11 MAC protocol. Clearly, the low power nodes benefit more when only these nodes use the proposed techniques (Variation I); the data success rate for the low power nodes is better by about 5 % with Variation I than with Variation II. However note from Figure 5, that the overall data success rate however is better with Variation II than with Variation I. This is because the high power nodes can also benefit from our schemes. As described in prior work [5], each node has a *transmission range* and an *interference range*. Neighborhood nodes within the transmission range of a node can decode a received packet from the node whereas nodes that are within the interference range of the node but are not within the transmission range cannot decode the received packet but can be interfered with due to the transmission of the node. The presence of the interference range results in the failure of packet

⁶ Other Scenarios yielded similar results.

transmissions after a fixed number of RTS retries by the MAC layer [12][13], is deciphered as the failure of a particular link. In actuality, there is no link failure, i.e., the desired recipient node may still be within the transmission range of the node trying to retransmit; such link failures are generally termed as false link failures [12] [13]. Link failures instigate rediscovery of routes, TCP retransmissions and timeouts and can result in the wastage of wireless capacity and poor utilization efficiency. With Variant I, the high power nodes are able to alleviate some of the effects of the interference range in terms of reducing the false link failures. The improvements seen overall in terms of the data success rate with Variant I, are therefore higher than with Variant II as shown in Figure 5. This is primarily because, by using our schemes, we are able to reduce the number of false link failures by up to 20 % as shown in the Figure 6.

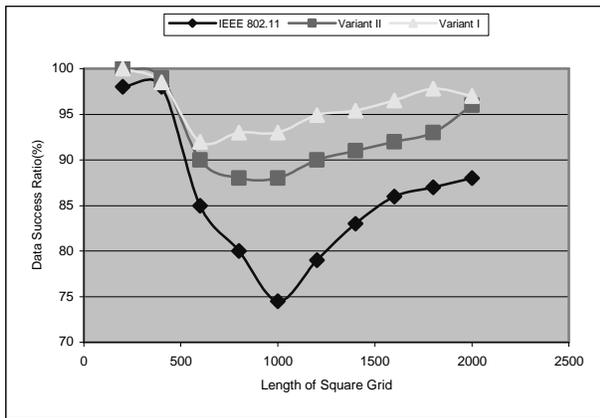


Figure 4 Data Success Rate (%) of Low Power nodes

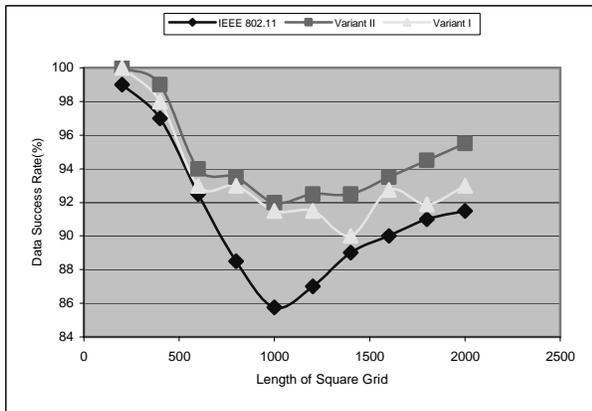


Figure 5: Overall Data Success Rate for a Heterogeneous Network

By using our schemes, the nodes that are within the interference range or beyond are now made aware of initiated data transfers at the MAC layer by means of the BW_RES message and this in turn reduces the false link

failures. In fact, we also deployed our schemes in a homogeneous network (with nodes transmitting at 0.28W) and we observed an overall increase in the network throughput of up to 10% and a reduction in the number of false link failures by about 15% as compared to the legacy IEEE 802.11 MAC.

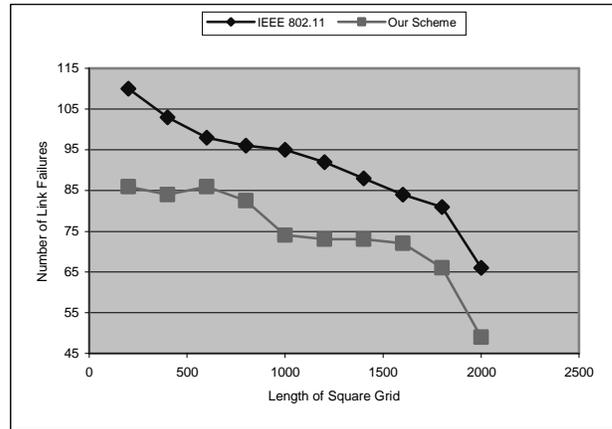


Figure 6: Link Failures for a Heterogeneous Network

Similar improvements were observed for other power levels (0.14 W and 0.56W). To summarize, the modifications that we propose seem to provide a generic framework for medium access control and provide benefits both in heterogeneous and homogeneous networks. Depending upon the number of low power nodes in the network, one might prefer to use Variant I or Variant II. The latter would provide a better performance if the fraction of low power nodes in the network is large. In general, since the number of low power nodes may differ from the scenarios that we consider, Variant II would be the preferred design specification.

Notice from Figures 4 and 5 that at high densities (grid size is small), the performance improvements are not significant. This is because, even with the smaller power level (0.14W) the transmission range is about 205m; this implies that all the nodes are typically within the hearing distance of one another in the network most of the time i.e., no asymmetry exists. At lower densities, the asymmetry increases and our schemes tend to provide benefits. If we further increase the grid size, the network tends to become sparse; the possibility of collisions now reduces and therefore the benefits seen due to our schemes decrease. In Figure 7, we show the improvements seen by the nodes in terms of throughput, with the use of both the proposed techniques; Variant II is used. The throughput of the low power nodes improves by as much as 14% as compared with the IEEE 802.11 scheme. Notice that our scheme also benefits high power nodes because (a) the reduction in number of retransmissions by low power nodes reduces the overall

contention for wireless medium and (b) as mentioned earlier, the effects of false link failures are alleviated. We observe an overall improvement in network throughput of up to 12% as compared to the IEEE 802.11 MAC protocol.

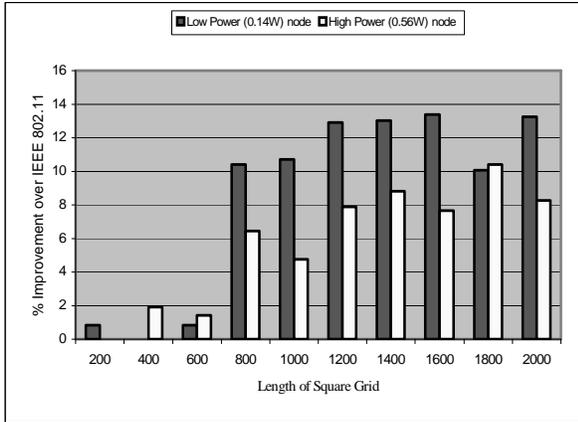


Figure 7: Percentage Improvement in Throughput over IEEE 802.11 MAC

We also examine the benefits that are achieved by using our schemes (Variant II) at various loads. We vary the packet generation rate of the FTP agent in order to vary load. We observe that the frequency of multi-reservations increases with load as one might expect. When a node wishes to initiate a transmission, with higher loads, it is more likely to find multiple packets destined for the same neighbor. Thus, the efficiency of our schemes improves as compared with the IEEE 802.11 MAC protocol with load (results not shown in detail due to space constraints). At a load of 500 packets/second the improvement in throughput is 5%, whereas with a load of 1000 packets/second, the improvement in throughput is 12 %.

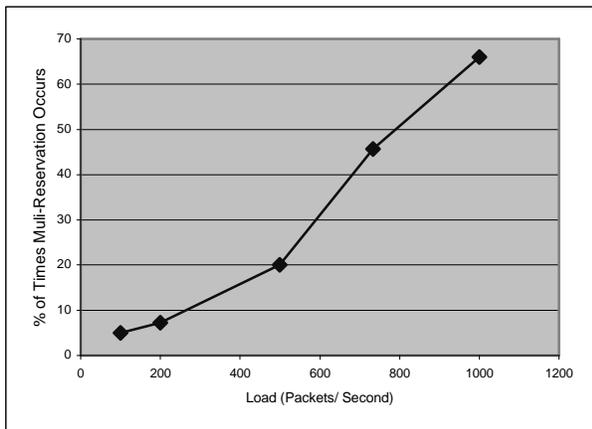


Figure 8: Frequency of Multi-Reservations

In order to quantify the impact of our intelligent broadcasting technique and our multi-reservation technique in isolation on the overall improvement in

performance in terms of the throughput efficiency, we performed two distinct set of experiments as specified by Case (b) and Case (c) i.e., we consider one scheme at a time. We found that merely deploying the intelligent broadcasting technique does not give a significant improvement in the overall network throughput (Figure 9). Our multi-reservation scheme does provide improvements of as much as 9 %. However, since the two schemes are independent they can be used in conjunction with each other and the benefits due to one can supplement the benefits due to the other.

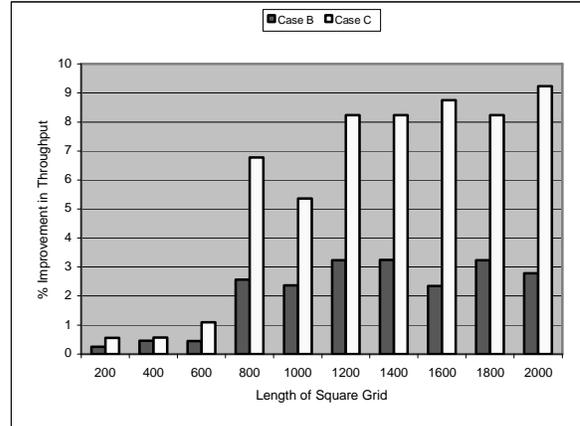


Figure 9: Improvements due to each scheme above IEEE 802.11

V. CONCLUSIONS

In this paper we consider ad hoc networks that consist of devices that have differing transmit power capabilities. Future networks are likely to consist of an integration of such multifarious devices. Prior work has shown that the IEEE 802.11 MAC protocol performs poorly in such networks. The primary reason for this is the fact that high power nodes could potentially be beyond the range of low power communicating nodes and might therefore initiate transmissions that collide with low power packets. One way to alleviate this is to propagate the CTS message beyond the one hop neighborhood of low power nodes. We propose effective ways to reduce the overhead due to such propagations. These techniques are based on (a) making a single reservation for multiple packet transmissions and (b) using an intelligent broadcast scheme to propagate the CTS message. We show that using our schemes we can improve the throughput of low power nodes by up to 14% and alleviate the unfairness caused by the legacy IEEE 802.11 MAC. We also show a significant reduction (by 20 %) in total number of link failures caused in the network due to interference from neighboring nodes. We conclude that our MAC layer framework offers a simple yet viable and effective option for medium access control in power heterogeneous ad hoc networks.

VI. ACKNOWLEDGEMENTS

We are indebted to Daniel Berger, graduate student in department of Computer Science, University of California–Riverside for his help and valuable comments with regards to implementation of our scheme in ns2.

VII. REFERENCES

- [1] Neeraj Poojary, Krishnamurthy, S.V. and Dao, S, “Medium Access Control in a Network of Ad Hoc Nodes with Heterogeneous Transmit Power Capabilities”, *Proceedings of ICC 2001*.
- [2] Brad Williams, Tracy Camp, “Comparison of Broadcasting Techniques for Mobile Ad Hoc Networks”, *ACM MOBIHOC 2002*.
- [3] S. Ni, Y. Tseng, Y. Chen, J. Sheu, “The broadcast storm problem in a mobile ad hoc network”, *ACM/IEEE MOBICOM 1999*.
- [4] NS2: Network Simulator.
<http://www.isi.edu/nsnam/ns/>
- [5] IEEE STD 802.11 1999. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999.
- [6] Bhargavan, V., Demers, A., Shenker, S. and Zhang, L., “MACAW: A media access protocol for wireless LANs”, *Proceedings of ACM SIGCOMM 94*.
- [7] E.Jung and N.H.Vaidya, “A Power Controlled MAC Protocol for Ad Hoc Networks”, *ACM/IEEE MOBICOM 2002*.
- [8] J. Monks, V.Bharghavan, “A Power Controlled Multiple Access Protocol for Wireless Packet Networks”, *IEEE INFOCOM 2001*.
- [9] C.E.Perkins and E.M.Royer, “Ad hoc on-demand distance vector routing,” in *Proceedings of IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, 1999.
- [10] D. Johnson and D. Maltz, “Dynamic Source Routing in Ad hoc Wireless Networks,” in *Mobile Computing*, edited by T. Imielinski and H. Korth, pp 153-181, 1996.
- [11] J. Yoon, M. Liu and B. Noble, “Random Waypoint Considered Harmful”, *IEEE INFOCOM*, 2003.
- [12] S.Kopparty, S.V.Krishnamurthy, M.Faloutsos and S.K.Tripathi, “Split TCP for Ad Hoc Networks”, *IEEE GLOBECOM 2002*.
- [13] F.Klemm, S.V.Krishnamurthy and S.K.Tripathi, “Alleviating Effects of Mobility on TCP Performance in Ad Hoc Networks using Signal Strength based Link Management”, *IFIP Personal and Wireless Communications (PWC)*, 2003.