

# Handling Asymmetry in Power Heterogeneous Ad Hoc Networks: A Cross Layer Approach

Vasudev Shah and Srikanth Krishnamurthy\*

Dept of Computer Science and Engineering, University of California, Riverside  
{vasu, krish}@cs.ucr.edu

## Abstract

*Power heterogeneous ad hoc networks are characterized by link layer asymmetry: the ability of lower power nodes to receive transmissions from higher power nodes but not vice versa. This not only poses challenges at the routing layer, but also results in an increased number of collisions at the MAC layer due to high power nodes initiating transmissions while low power communications are in progress. Previously proposed routing protocols for handling unidirectional links largely ignore MAC layer dependencies. In this paper, we propose a cross layer framework that effectively improves the performance of the MAC layer in power heterogeneous ad hoc networks. In addition, our approach seamlessly supports the identification and usage of unidirectional links at the routing layer. The framework is based on intelligently propagating low power MAC layer control messages to higher power nodes so as to preclude them from initiating transmissions while the low power communications are in progress within their sensing range. The integrated approach also constructs reverse tunnels to bridge unidirectional links thereby facilitating their effective usage at the routing layer. Extensive simulations are performed to study the proposed framework in various settings. The use of our framework improves the overall throughput of the power heterogeneous network by as much as 25 % over traditional layered approaches. In summary, our framework offers a simple, yet effective and viable approach for media access control and to support routing in power heterogeneous ad hoc networks.*

## 1 Introduction

As ad hoc networks gain popularity, one might expect emerging networks to consist of multifarious devices with differing capabilities. One could envision low power sen-

sor 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 are likely to have different power capabilities and thus, are likely to transmit with different power levels. This, in turn, leads to possible link asymmetry wherein the transmission of a high power node is received (or is sensed) by a lower power node whereas the high power node cannot sense the transmissions by the low power node. The effects of asymmetry also pose challenges when power control is to be employed in ad hoc networks [7, 14]. In the presence of such asymmetry, traditional protocols that are typically designed with an implicit assumption that links are bi-directional, either fail or perform poorly. Specifically, at the MAC layer, this leads to an exacerbation of the hidden terminal problem [5]. Routing becomes more complex due to the presence of such *unidirectional* links [15, 17, 19, 20].

While there has been a plurality of prior research efforts on performing routing in the presence of unidirectional links [15, 17, 19, 21, 22], to the best of our knowledge, most of these efforts ignore MAC layer dependencies. It has been shown that the performance of the IEEE 802.11 MAC layer protocol degrades in the presence of link asymmetry [5]. The inefficiency at the MAC layer affects the routing protocols primarily designed for ad hoc networks with bi-directional links. With these protocols, unidirectional links may not be identified and hence, may cause problems at the routing layer. The use of unidirectional links may effectively shorten the lengths of routes; this could in turn improve the throughput performance (over cases wherein such links are ignored and only bi-directional links are used) in power heterogeneous ad hoc networks. Furthermore, choosing longer paths can degrade the performance due to an overall increase in load (more transmissions). In the absence of appropriate MAC layer support, due to link level asymmetry, one may also expect to observe an increase in routing overhead: the aggravation of the hidden terminal problem increases the number of *false link failures* [11] thereby causing an increased frequency of

---

\*This work is supported in part by the NSF CAREER Grant No. 0237920 and the NSF NRT grant No. 0335302

route discovery attempts.

In our previous work [6], we proposed simple MAC layer mechanisms that alleviated the effects of asymmetry to a small extent. In this paper, we propose an integrated *cross-layer* framework that unites the MAC and the routing layers to achieve the following: (a) Eliminate, almost entirely, the MAC layer inefficiencies in power heterogeneous ad hoc networks and improve the performance in terms of achieved throughput to commensurate with that of a network in which nodes are homogeneous and, (b) Provide underlying support for identifying and effectively utilizing unidirectional links at the routing layer. This support enhances the performance in terms of throughput perceived at the higher layers, as compared to that of traditional routing protocols.

The key idea in constructing our framework is to *route* MAC layer control messages (specifically the clear to send or the CTS message used with the IEEE 802.11 MAC protocol) that are transmitted with low powers, to high power nodes, beyond the one-hop neighborhood of the low power transmitter, that can potentially initiate in the interim, new interfering transmissions. We first propose modifications to the MAC layer so as to enlist functionalities at the routing layer to achieve the above goal: this results in our **topology aware CTS propagation (TACP) scheme**. Next, to complete our framework, we extend TACP to *tunnel* MAC and routing layer control packets in the reverse direction of a unidirectional link on a path that spans the link.

We evaluate our cross layer framework, via extensive simulations, in two steps. We first eliminate higher layer artifacts and examine the performance improvements exclusively at the MAC layer. Later, we include higher layer protocols to study the performance at the transport layer (UDP). We observe that our framework improves the throughput by approximately 25 % as compared with the traditional layered protocol stack that includes the IEEE 802.11 MAC layer protocol and AODV. The performance with our schemes is equivalent to that of a network without power heterogeneity; in other words, the effects of asymmetry are completely handled with our framework.

The rest of this paper is organized as follows: In Section 2 we describe related previous work. In Section 3 we revisit the problems due to link asymmetry in brief to provide the basis for this work. In Section 4 we describe our cross layer framework in detail. In Section 5 we present our simulation results (using the two steps mentioned earlier) and discuss them at length. Our conclusions and a discussion of future work form Section 6.

## 2 Related Previous Work

The IEEE 802.11 MAC protocol and derivatives thereof have been popularly considered for use in ad hoc networks

[1, 5, 7, 8]. Previous work on deploying multiple transmission power levels in ad hoc networks have primarily focused on achieving energy savings. There have been efforts geared towards topology control via power adaptations (as in [12]); however, the asymmetry effects due to such control on protocols (as we consider) are ignored in these efforts.

In [7] and [8], the use of power control with the IEEE 802.11 MAC protocol is examined. However, the proposed protocols require transmissions of the RTS and CTS messages with a maximum preset power level in order to reach all nodes that are within the maximum transmission range of the transmitter. By transmitting the control messages at the maximum power level, these schemes avoid the effects of asymmetry. Consequently, they are unable to exploit spatial re-use gains that are potentially possible with power control. Furthermore, they cannot be used in networks consisting of multifarious nodes with different maximum transmit power capabilities. In [13], a dual channel approach is proposed wherein, a busy tone is transmitted on a separate control channel at the maximum power level. Each neighbor estimates the channel gain from the busy tone and is allowed to transmit if its transmission does not interfere with the ongoing reception. The scheme once again avoids the problems due to asymmetry by assuming that nodes transmit busy tones at the preset maximum power levels. Furthermore, additional hardware complexity (two transceivers) is required. The effects of the protocol on routing are not studied. In [14], Muqattash and Krunz propose a dual channel dynamic power control protocol; the scheme supports frequency re-use. However, the proposal again necessitates the use of two radios. The problem of asymmetry is again avoided by using the control channel for transmission of the RTS and CTS messages using the maximum preset power level. The messages also carry additional information to specify the maximum power level that can be used for a newly initiated communication such that other ongoing communications are not affected. Unlike in our work, the protocol was not designed for the networks wherein nodes have differing transmit power capabilities. On the contrary, our schemes are applicable in homogeneous networks wherein power control is to be administered.

There has been significant work on routing in MANETs [9]. However, most of the protocols were primarily designed for networks that exclusively consist of bidirectional links (in other words, homogeneous networks). There has been prior work on routing in the presence of unidirectional links [15, 17, 19, 20, 22]. Link layer tunneling approaches to hide the unidirectional nature of links have been explored in [17, 18]. Tunneling is based on forming a reverse multi-hop path [19] for each unidirectional link using the information gathered by the routing protocol. A similar idea appears in [16]: here, a sub-layer beneath the routing layer is developed. There is also some work on us-

ing multihop acknowledgements to discover unidirectional links [17]. The above link layer approaches however ignore the MAC layer problems that exist due to link asymmetry. In [15], the authors propose to bypass the unidirectional links and route the packets via bi-directional links. However, this can lead to the choice of longer and possibly less efficient paths. Furthermore, the studies in [22] show that not accounting for the presence of unidirectional links can lead to erroneous routing decisions. Our framework provides support for unidirectional routing while taking MAC layer effects of asymmetry into account.

In our previous work [5], we consider a network in which nodes have heterogeneous power capabilities. The studies quantify the inefficiencies in the use of the IEEE 802.11 medium access control protocol in such networks. It was observed that low power transmissions were often interfered with, by transmissions from high power nodes that were unable to hear the RTS/CTS exchange between the low power nodes. In [5], we attempted to solve this problem by means of a CTS propagation technique using a standard *flood-type* broadcast algorithm to let the high-power nodes in the neighborhood 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. The propagated message was called the Bandwidth Reservation (BW\_RES) control message. With this method, the overhead incurred in propagating the BW\_RES control messages was seen to outweigh the potential gains achieved in terms of reducing the number of collisions. In our later work in [6], we considered two MAC layer enhancements for alleviating the effects due to asymmetry. First, we considered a smart broadcast scheme based on a counter based approach [3] in lieu of flooding to propagate the (BW\_RES) control message. Each node counts the number of rebroadcasts of a message that it overhears (due to its neighbors' rebroadcasts) and aborts its own broadcast if this count is higher than a preset threshold. This scheme has been shown to eliminate many of the unnecessary broadcasts that occur when flooding is used [3]. Second, we considered reserving the bandwidth for multiple data packets with a single RTS/CTS exchange/propagation (multireservation scheme). These enhancements provided small improvements to the performance observed at the MAC layer. In this paper, the new cross layer approach that we propose provides significant further improvements and the performance commensurates with that of a homogeneous network. We use the multi-reservation scheme in conjunction with our cross layer framework, since this is seen to provide benefits and is largely independent of the construction of our framework. For completeness, in the following paragraph, we describe the scheme in brief. As in [5] and [6], we refer to the control messages that are *routed* beyond the one-hop neighborhood of a low power communicating node as BW\_RES messages.

A single RTS/ CTS/ BW\_RES initiation is used to enable *multiple* sequential DATA/ACK exchanges with the *multi-reservation* scheme. The multiple DATA packets are in fact independent packets (they have their own fields including a separate checksum field for each packet). Before sending an RTS, a node checks its interface queue (between the network and the MAC layers) for other DATA packets with the same MAC layer destination address. If such packets are found, the duration field in the control messages is changed to account for multiple data transfers. We allow for a maximum of two DATA packets to be exchanged by means of a single control message exchange since this was shown to yield the best trade-offs via simulations in [6].

### 3 Problems due to Link Asymmetry

In this section, we revisit the performance of Distributed Coordinated Function (DCF) of the IEEE 802.11 MAC protocol in a power heterogeneous network setting. We briefly discuss its deficiencies and highlight the resulting effects on the higher layers [5].




Figure 1. Hidden terminal problem due to link asymmetry at the MAC layer.




Figure 2. Problem with identifying unidirectional links due to asymmetry.

As alluded to earlier, the inefficiencies of the IEEE 802.11 MAC protocol are primarily due to link asymmetry: in certain scenarios, high power transmissions can be sensed (or received) by low power transmissions but not vice versa. Thus, the low power nodes are *hidden* from high power nodes. This increases the number of collisions that are experienced by low power communications (for a more detailed discussion refer [5]). This effect is depicted in Figure 1. The RTS/CTS exchange between the two low power nodes A and B is not overheard by node H since node H is not within the sensing or interference range<sup>1</sup> of these communicating nodes. Thus, during the data exchange between nodes A and B, node H could possibly begin its own transmission, thereby causing a collision at node B.

The second problem that is manifested at the MAC layer is that a node fails to identify (and therefore utilize) a unidirectional link. This effect is depicted in Figure 2 where the node  $H_1$  can reach node  $L_1$  but not vice versa. As a result, if  $L_1$  responds to any message (such as an RTS message) sent by  $H_1$ , the response never reaches  $H_1$ . Similarly, any control message initiation by  $L_1$  ( $L_1$  could send an RTS message) would never reach  $H_1$ . Depending on the scenario, these problems could cause degradations due to wasteful control message transmissions and backoffs at the MAC layer. Furthermore, the link asymmetry can degrade the performance of traditional on-demand routing protocols<sup>2</sup> due to the loss of control messages. One such effect is depicted in Figure 2 where the node  $H_1$  is trying to establish a route to  $H_2$  through nodes  $L_1$  and  $L_2$ . The routing control packets (such as the RREP) from  $L_1$  are not received by  $H_1$  since it is outside the range of  $L_1$ . Such effects could lead to repeated (albeit unsuccessful) route discovery attempts. Prior papers on unidirectional routing describe these problems in detail [19], [15].

## 4 Framework for Handling Asymmetry

In this section, we present our cross layer approach and the interactions between the MAC and routing layers. We discuss the rationale behind the approach and explicate the possible benefits.

**Scope of Propagation for the BW\_RES Message:** As mentioned earlier, the key idea of our approach is to enable the intelligent propagation of the BW\_RES messages,

<sup>1</sup>Typically there are two ranges defined for MAC layer transmissions [1]. Neighborhood nodes that are 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; however, their receptions can be interfered with due to the transmission of the node.

<sup>2</sup>Similarly, route update messages can lead to falsified routing tables when traditional proactive routing schemes are used. The discussion/evaluation of such effects in detail is not included due to space constraints.


that are derived from the CTS message, to beyond the one hop neighborhood of a low power communication. The goal is to inform high power nodes that can potentially initiate transmissions while the low power communication is in progress, about the communication, so that they defer their transmissions until after the completion of the low power data transfer. The BW\_RES message is propagated with a TTL (Time to Live) value of 3, i.e., up to a distance of two hops. We provide a simple computation that provides a basis for using the above value for the TTL. We assume the channel model simply consists of attenuation with distance (as in popularly used models). Thus, if  $P_{trans}$  is the transmission power of a node, then at a distance  $d$ , the received power is defined by  $P_{recvd} = P_{trans} d^{-k}$ , where  $k$  is the path loss exponent ( $k=2$  for the free space propagation model [2]). Consider the power level of the high power nodes to be  $P_{max}$  and that of the low power nodes to be  $P_{min}$  (chosen as 0.56 W and 0.14 W respectively for our simulations) in our heterogeneous network with the two types of nodes<sup>3</sup>. Let  $d_{max}$  and  $d_{min}$  be the corresponding transmission ranges of the two types of nodes, respectively. At the periphery of a node's transmission range, we assume that the received power would be a fixed value (i.e., reaches a threshold just enough for the message to be decoded correctly). The received power threshold at the periphery of a high power node's transmission range and that at the periphery of a low power node's transmission range would be identical. The objective is then to ensure that a CTS message generated by a low power node reaches any high power node within whose range the low power node lies. In the worst case, the low power node could be on the periphery of the range of a high power node. We note that:

$$\frac{P_{max}}{P_{min}} = \frac{d_{max}^k}{d_{min}^k} \quad \text{and hence,} \quad \frac{d_{max}}{d_{min}} = \left(\frac{P_{max}}{P_{min}}\right)^{\frac{1}{k}} \quad (1)$$

For the chosen settings for  $P_{max}$  and  $P_{min}$  and with the free space propagation model, we obtain  $d_{max} = 2d_{min}$ . Thus, if the BW\_RES message of a low power node were to be propagated through two low power hops, the message would reach any high power node within whose transmission range the low power node lies. If we also take into account the sensing range, which is typically modelled to be twice the transmission range, we would need to propagate the message through 4 low power hops (or correspondingly 2 high power hops). Note that this simple analysis is extendable to cases where other channel models are assumed. Furthermore, we point out that in typical scenarios wherein asymmetry causes the highest degradations, a low power node typically would have high power nodes in its vicinity to forward the BW\_RES messages. Thus, choosing a TTL value of 2 or 3 can provide the desired benefits. It is

<sup>3</sup>A similar analysis is easily feasible for more generic cases.

important to note that, forwarding the BW\_RES messages to within the above determined distance *does not guarantee* that such messages would reach all of the high power nodes that can potentially initiate colliding transmissions; a path that is bounded by the above determined hop distance may not exist to some high power nodes. Via simulations, we ascertained that it is better to restrain the forwarding to a localized vicinity of the communication rather than to incur high overheads due to forwarding the message across a wider scope<sup>4</sup>.



**Figure 3. Propagating BW\_RES messages.**

**The Cross Layer Framework:** The MAC layer elicits assistance from the routing layer in order to determine a limited set of nodes to perform the BW\_RES re-broadcasts. In turn, the routing layer depends on the MAC layer for the discovery of unidirectional links and for assistance in the use of such links.

**Topology Aware CTS Propagation (TACP):** We propose a *routing assisted approach* using which a node that initiates communications, *multicasts* the BW\_RES message to the high power nodes in its vicinity. In order to facilitate this, broadly speaking, we require that each node maintain link-state information with regards to its two hop neighborhood. Towards this, each node broadcasts a *hello* message once every “*hello interval*” milliseconds. This is a message with a TTL of 1 i.e., it is only exchanged between neighbors. Before we discuss the contents of this message we define the following: the nodes from which a node can directly receive messages are called the *inbound neighbors* of the receiving node. In the *hello* message, at network instantiation, each node broadcasts a list of those in-bound neighbors from whom it is currently receiving *hello* messages. Each message also contains information that associates each neighbor with its corresponding maximum transmission power capability (in watts). Note that these messages are exchanged between the *network* layers of the nodes. After the initial phase, each node, with a collection of the received *hello* messages progressively constructs an *inbound tree*. This tree includes all of the neigh-

<sup>4</sup>We however do not include these results in this paper due to space constraints.


neighbors who can reach the node under discussion. As the network reaches a steady state (nodes construct stable in-bound trees), nodes begin transmitting *update hello* messages that are now modified to contain their inbound trees. Each node, then combines the *in-bound trees* reported by its neighbors with its own in-bound tree to progressively form a *localized graph* that depicts its local neighborhood. The update messages are then, further modified to include this localized neighborhood. Periodic transmissions of the update messages help refine this *localized graph*<sup>5</sup>. As information propagates, the localized graphs become more extensive. This would allow nodes to gather additional information (beyond their two hop neighborhoods). However, the size of the update messages grows considerably and there is a trade-off between the amount of information propagated and the extent of knowledge that is possessed by a node with regards to its vicinity. In our studies, nodes simply prune nodes that are beyond a certain number (*n*) of hops, from their localized graphs. The pruned version of the graph is included in the update messages. In each update, by transmitting *only the changes* to the localized graph that have occurred since the previous update, one may significantly constrain the size of the update messages. As mentioned earlier, a choice between 2 and 4 for *n* may be a reasonable value in general, for ensuring that most of the high power nodes that affect a given low power communication, are informed by means of BW\_RES messages. Our simulations suggest that setting *n* to 3 offers the best benefits<sup>6</sup>.

Using the localized graph, our objective is to have each node construct an *n-hop outbound Steiner tree* on which the BW\_RES messages will be **multicast** to the high power nodes. Note that a low power node does not initiate the propagation of the BW\_RES message if there are no high power nodes in its *n*-hop neighborhood. If such nodes exist, the low power node first identifies the minimum set of nodes in its one-hop neighborhood that can reach all other high power nodes in its two-hop neighborhood. We refer to these nodes as “*Candidate Nodes*” for relaying the BW\_RES message. The node then includes the identifiers (IPv4 addresses) of the candidate nodes in the CTS message. Since we want to minimize the number of BW\_RES rebroadcasts while maximally reducing the latency incurred during a MAC layer exchange, the candidate nodes are typically chosen to be high power nodes (if such nodes are available). The one-hop relays then perform a similar computation to identify the next set of relays (if needed) and so on. The identifiers of this next set of candidate relays are included in the BW\_RES message. If a node, upon the

<sup>5</sup>The periodicity of the update messages would depend on the mobility in a given scenario. If nodes are highly mobile, the update messages would have to be transmitted with high frequencies.

<sup>6</sup>This implies that all the high power nodes that can be reached via three low power hops from a low power communication are informed of the impending communication.

receipt of either a CTS or a BW\_RES message does not find its identifier in the message, it simply updates its NAV (network allocation vector) in accordance with the IEEE 802.11 MAC protocol and discards the message. The multi-reservation scheme from [6] is incorporated as well and if possible, a node reserves the channel for  $N$  (a system parameter discussed earlier) data packets destined for the same neighbor by means of a single RTS/CTS/BW\_RES initiation.



**Figure 4. Reverse Route Construction to route MAC/routing control packets.**

**Construction of Reverse Routes for bridging Unidirectional Links:** Next, we discuss our approach for constructing reverse routes that span unidirectional links to assist routing. Exchanges of the previously discussed *hello* messages help detect the presence of unidirectional links in the network. When a node receives a *hello* message from a neighbor, and finds that it is excluded from the neighbor’s neighborhood list, the node infers that it is at the *tail* of a unidirectional link. The unidirectional links are also depicted in a node’s *pruned localized graph*. If a node happens to be at the tail of a unidirectional link, using the graph, it computes a reverse path to the node at the *head* of the unidirectional link. We illustrate this using the example in Figure 4. Note that the link between the nodes  $L_1$  and  $H_1$  is unidirectional i.e.,  $H_1$  can reach  $L_1$  *directly* but not vice versa. In order for the link to be utilized,  $L_1$  would need a reverse route to  $H_1$ . This is done as follows:  $L_1$  learns from its pruned localized graph (using the method explained above) that  $L_2$  can reach  $H_2$  which can in turn, reach  $H_1$ . Thus, the reverse path from  $L_1$  to the high power node  $H_1$  via nodes  $L_2$  and  $H_2$  is computed. We wish to point out here that if a reverse path of less than  $n$  hops does not exist from the node at the *tail* of the link to the node at the *head* of the link, the unidirectional link can neither be identified nor utilized. It may be possible to find longer reverse paths; however, discovering such paths would entail significant additional overhead and thus, could actually outweigh the gains incurred in utilizing the link.

Once the reverse path is computed, the node at the *tail*

of the unidirectional link simply embeds or encapsulates the control packets from the MAC and the routing layers into an IP packet and routes or *tunnels* the packet using the constructed reverse path. Note here that the above *proactive* form of routing is only used within the node’s  $n$  hop neighborhood and any traditional on-demand routing protocol can then be deployed for performing network-wide routing. With our scheme, in particular, we tunnel the following types of control packets:

- MAC layer packets: The CTS and the ACK packets.
- Routing control packets: The RREP (or route reply) packets [9].

The main motivation for tunneling the above packets is that the benefits are largely achieved *if the actual data is routed* on the unidirectional link (thereby potentially avoiding a longer alternate path). Thus, the unidirectional link is only found from the source to the destination. In our simulation scenarios we use UDP sources and this technique is therefore appropriate. For other applications, it may be beneficial to tunnel other control packets (such as the RTS or the RREQ packets) as well. Note here that if, either the reverse path or the unidirectional link were to fail, it would cause the *tunneled bi-directional link* to break. The network-wide on-demand routing protocol would then instigate a route error message. We finally point out that the node at the *tail* of the unidirectional link would also tunnel the network layer *hello* messages to the node at the *head* of the link in order to make the latter aware of the existence of the link.

In order to distinguish between tunneled MAC and network layer packets, the packet header at the the network layer is modified to support a *flag* indicating whether a packet is of the encapsulated type. At the network layer, this flag can be added as an option to the IP header (beyond the 20 byte standard header). The value of the flag would further indicate whether the encapsulated packet contains a MAC layer packet or a routing layer packet. Upon stripping the outer header, based on this value, the network layer either forwards the packet to the routing module or to the MAC layer.

## 5 Performance Evaluation

In this section we evaluate our proposed framework. As mentioned earlier, we sub-divide our evaluation studies into two steps. In our first step, our objective is to exclusively quantify the benefits observed at the MAC Layer. We eliminate the routing and transport layer artifacts in this step; this helps us evaluate the performance as seen at the MAC layer and compare the benefits with our proposed framework with the results reported in [5] and [6]. We later evaluate our integrated MAC/routing framework with the higher layers included.

## Performance Evaluation at the MAC layer:

**Simulation Models:** Simulations are performed using the event driven network simulator *ns2* (version 2.26). In order to decouple higher layer effects from our study of the MAC protocol, in our first step, we extend the *ns2* simulator to introduce a *poisson* traffic generation agent immediately above the MAC layer<sup>7</sup>. Each packet is 1000 bytes in size. When a data packet is generated at a node, it is randomly destined for one of the neighbors of the nodes. The average rate is varied to vary system load. In order to be consistent with previous experiments in [5] and [6], at this stage, we choose the number of nodes and the transmission power levels as shown in Table 1. 50% of the chosen nodes are high power nodes and the other 50% are low power nodes. The choice for using 50% low power nodes is based on the results in [15] which suggest that the maximum number of unidirectional links exist in the network for this value<sup>8</sup>. The physical layer is based on the IEEE 802.11 specifications. Nodes move in accordance to a modified version of the *random waypoint* model with a constant speed of 6 meters per second between the chosen end points of a trajectory instance<sup>9</sup>. A pause time of 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 [10]. Nodes generate *hello* messages every 0.5 to 1 seconds depending on the speed considered. The simulated network is deployed in a square region whose area is tuned 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, in which the nodes are deployed, in meters.*

The medium is assumed to be free of noise or any errors due to fading as in other previous work [5], [7], [8]. The only interference effects are due to simultaneous transmissions from multiple access users. All MAC control packets are transmitted at 1 Mbps and data is transmitted at 2 Mbps in order to conform to the IEEE 802.11 standards [1]. We find that reserving the bandwidth for a maximum of  $N=2$  (for high power nodes) and  $N=3$  (for low power nodes) sequential transmissions provide the necessary benefits<sup>10</sup>. As per our qualitative assessments in the previous section, we choose the TTL value for the BW\_RES messages to be  $n = 3$ . We have further done simulations that support our qualitative assessments. However, we omit these results due

<sup>7</sup>We use UDP traffic in a subsequent set of experiments.

<sup>8</sup>Various other parameter settings were studied. However the results observed and the interpretations thereof were similar to those reported and are thus not included.

<sup>9</sup>Other speeds were considered (of up to 10 meters per second), however the results were observed to be similar and are hence not presented.

<sup>10</sup>Note that simulation results suggested these values were the best in order to ensure good levels of multi-reservation while at the same time maintaining fairness [6]. We omit these results due to space constraints.

to space limitations.


**Table 1. Simulation Environment for Evaluations at MAC Layers**

Simulator	ns2 (version 2.26)
Total number of nodes	40
Power levels used	0.56 watts and 0.14 watts
Number of high power nodes	20
Number of low power nodes	20
Packet Generation Rate	1000 packets/second
Traffic model	Poisson
Mobility model	random way point
Pause time	0.1 seconds
Speed	6 meters per sec
Length of square grid	varied between 300 to 2000 meters


**Metrics:** Our primary metrics of interest are the Data Success Rate and the Throughput Efficiency of each node. In order to quantify channel usage, we define throughput efficiency at each node as the ratio of the time spent by a node in successfully transmitting data to the total simulation time. The Data Success Rate (%) is defined as the percentage of DATA transmissions that succeed after a successful RTS/CTS exchange between the two communicating nodes.

**Performance of the Topology Aware CTS Propagation scheme (TACP):** We first evaluate our Topology Aware CTS Propagation (TACP) to evaluate the benefits exclusively at the MAC layer. Since the multi-reservation scheme proposed in [6] can offer significant benefits and since TACP can be used in conjunction with this scheme, we use a combination of the two schemes in our experiments. We compare the performance of the following: **Case (a)** the legacy IEEE 802.11 MAC protocol, **Case (b)** the IEEE 802.11 MAC protocol with BW\_RES propagation by means of TACP, **Case (c)** the IEEE 802.11 MAC protocol with the multi-reservation technique and, **Case (d)** the IEEE 802.11 MAC protocol with BW\_RES propagation using the TACP technique combined with multi-reservation technique. Furthermore, we perform experiments wherein all the nodes, regardless of whether they are high power nodes or low power nodes, adopt our proposed scheme. We choose this strategy since, as seen previously in [6], high power nodes also benefit from BW\_RES propagations in terms of alleviating the effects of hidden terminals and reducing false link failures.

To begin with, we examine the performance of the combination of TACP with the multi-reservation technique (Case (d)). We compare the performance of the rendition in Case (d) with that in Case (a). In Figure 5, we observe a significant improvement in the data success rate of low power nodes, with the use of our schemes, in comparison to that observed with the legacy IEEE 802.11 MAC protocol. Notice that the low power nodes see an overall improvement of up to 20% as compared with the legacy IEEE 802.11 MAC protocol and the performance of the low power nodes is almost as good as the performance of the high power nodes in the heterogeneous network. **Furthermore, we observe**



**Figure 5. Data success rate improves significantly with our Framework.**




**Figure 6. Our Framework significantly reduces overhead due to BW\_RES propagation.**


that the overall performance is very close to that observed in a homogeneous network (This was never previously achieved). We also observe that the number of false link failures decreases by about 28% as compared to the legacy IEEE 802.11 and by about 8% as compared to when our intelligent broadcasting scheme is being used. This further improvement is due to the added intelligence at the MAC layer resulting in a significant reduction in the overhead, and as a consequence, the contention for the channel. In order to elucidate this further, we compare the number of BW\_RES messages generated per CTS instantiation with (a) the standard flooding scheme, (b) the intelligent counter-based broadcast of BW\_RES messages considered in [6] and (c) TACP. The results are shown in Figure 6. Note that the number of BW\_RES messages broadcasted per CTS message reduces by about 50 % with TACP as compared with the intelligent broadcast scheme. However, TACP does require the transmission of *hello* messages unlike the intelligent broadcast scheme. The overall improvements (relative to the IEEE 802.11 MAC protocol) in terms of the data success rate and throughput efficiency as observed, are higher with TACP than with the intelligent broadcast scheme *even*

with the overhead incurred due to *hello* messages. We observe that with the reduction in interference from the high power nodes, the throughput efficiency of the low power nodes increases by up to 24% (Figure 7). Our scheme also aids the high power nodes as it alleviates the false link failures to some extent; this is reflected by an improvement in their throughput efficiency by up to 12%.

We also conducted experiments by introducing additional heterogeneity by increasing the number of possible power levels. Although different scenarios were considered, we only provide sample results due to space limitations. We observe from Figure 8 that our scheme shows an overall improvement in network throughput by up to 18% as compared to that with the IEEE 802.11 MAC protocol. Again, the improvements are more significant as we increase the grid size (the density of nodes is chosen to be moderate as opposed to being set to extremely high) since the asymmetry is higher in these scenarios as described earlier. In this particular experiment, the fraction of nodes that belong to each class is made approximately equal.




**Figure 7. Our Framework improves throughput efficiency.**



**Figure 8. Our Framework shows improvements with more classes of heterogeneity.**

In order to quantify the impact of our TACP technique and our multi-reservation technique in isolation on the over-



**Figure 9. Improvements seen by TACP and Multi-Reservation scheme in isolation.**

all improvement in performance in terms of the throughput efficiency, we performed two distinct sets of experiments as specified by Case (b) and Case (c) i.e., we consider one scheme at a time (Figure 9). We found that merely deploying the TACP provides improvements of as much as 16% and multi-reservation scheme provides improvements of as much as 8%. By using the schemes in conjunction the benefits due to one can complement the benefits due to the other.

**Evaluating Higher Layer Effects:**

In this second step of our evaluations, our objective is to incorporate the routing and transport layers atop our MAC layer. We now consider UDP traffic (Constant bit rate: CBR) as in many previous studies on ad hoc networks [7]. We evaluate the performance of two popular on-demand routing protocols (AODV and DSR [9]) in conjunction with our cross-layer framework. The simulation setup parameters that we use in this step of our evaluations are listed in Table 2. The performance metrics that we compute are now observed at the UDP layer and are listed below:


- *Packet Delivery Ratio:* Ratio of number of packets delivered to the number of packets generated.
- *Route Search Attempts:* Number of new route discovery attempts that are initiated.
- *Route Search Failures:* Number of times that a source node does not find a path to a destination<sup>11</sup>.
- *Average end-to-end delay:* Mean end-to-end delay that the packets experience.

We reiterate here that TACP as considered in our previous experiments is now supported with our methods to compute reverse routes that bridge unidirectional links. When

<sup>11</sup>Note here that this can be due to one of two reasons (a) the source and destination belong to different network partitions and (b) due to the presence of a unidirectional link on the path from the source to the destination. While in the first case none of the schemes can do anything to compute a route, in the second case, our framework can discover routes while the traditional methods fail.

**Table 2. Simulation Environment for Evaluations at Higher Layers**


Simulator	ns2 (version 2.26)
Bit rate of channel	2 Mbps
Radio model	Lucent WaveLAN
MAC protocols	IEEE 802.11 MAC DCF IEEE 802.11 with MAC Layer Enhancements IEEE 802.11 with Cross Layer Framework
Routing protocols	AODV DSR AODV and DSR with Cross-Layer Framework
Number of nodes	80
Heterogeneity	10%-50% low power nodes
Power Levels	0.56W and 0.14 W
Traffic	CBR
Simulation Area	Varied
Packet Size	512 bytes
Source sending rate	5 -50 packets/sec
Sources	Varied between 10-15




**Figure 10. Our Framework improves packet delivery ratio of nodes.**

we refer to *MAC layer enhancements* we allude to a combination of the multi-reservation scheme with the intelligent broadcast scheme considered in [6].

In Figure 10, we observe that the packet delivery ratio improves with either AODV or with DSR using our cross-layer framework due to reduced level of interference from high power nodes (thereby reducing false link failures) and the added ability of identifying and utilizing unidirectional links. Note that the improvements are higher with AODV than with DSR. This is primarily because AODV, in its traditional settings, does not support asymmetric links. DSR on the other hand, (without our framework) has features that allow a destination to invoke a separate route discovery in order to discover the source (upon the receipt of a route request) in the presence of unidirectional links. Thus, AODV in fact, has more to gain when deployed over our framework. DSR too benefits since these *reverse* route discovery floods are no longer needed. The improvements in packet delivery ratio (of about 25 % as compared to 12 % with the MAC layer enhancements) over the traditional IEEE 802.11 MAC are seen due to the overall reduction in contention in the network due to decrease in routing and MAC layer control packet overhead. This is evident from Figure 11. **The**




**Figure 11. Our Framework reduces the total number of route discoveries attempts.**



**Figure 12. Our Framework reduces the total number of route search failures.**

results in this figure indicate that the number of route discovery attempts are dramatically reduced (by about 35% for AODV and by about 25% for DSR) with our framework. The reason for this significant reduction is that, with our framework, the nodes are now able to easily discover and thereby use unidirectional links. Without our framework, these links were either rendered useless (in the case of AODV) or were discovered by expending high overhead (with DSR). Our results also show a reduction in the percentage of route search failures by up to 25% (Figure 12) for both the routing protocols under consideration.



**Figure 13. Nominal increase in end-to-end delay with our Framework.**

In Figure 13, we plot the mean delay experienced by packets versus the *level of heterogeneity*. The level of heterogeneity refers to the percentage of low power nodes in the network<sup>12</sup>. We note that the mean delay that is experienced is only marginally increased by our cross-layer framework as compared with the traditional layered protocols. We wish to point out that there are several conflicting factors at work. Since our cross-layer framework requires the additional transmission of BW\_RES messages at the MAC layer, each MAC layer transmission usually takes longer than when using traditional schemes. However, this is offset to some extent, by the use of multi-reservations. However, with the traditional schemes, route failures and consequently, route discovery attempts are observed to occur with greater frequency. During these periods, data packets simply wait in the source queue. We find that overall, the delay with our cross-layer framework is only marginally larger than that with traditional protocols due to these conflicting factors. We believe that this slight increase in delay is acceptable considering that, with our framework, we observe significant gains in the overall throughput efficiency and packet delivery ratio.

Our proactive maintenance results in an increase in the overhead incurred. However, we notice that this only causes a slight increase in the overall overhead incurred (by about 10 %) and the superior performance with our framework is in spite of this overhead.

## 6 Conclusions and Future Work

In this paper, our key contribution is the development of a unified framework that marries the MAC and the routing layers in order to deal with link level asymmetry in power heterogeneous ad hoc networks. There are no prior solutions that handle asymmetry effectively at the MAC layer. Furthermore, previously *unidirectional* routing schemes have ignored MAC layer dependencies. We argue that a tightly coupled MAC/routing framework is needed in order to effectively overcome the effects of asymmetry. In our framework, the MAC layer elicits assistance from the routing layer to identify link asymmetry. Low power nodes then route MAC layer control packets to high power nodes that are beyond their transmission range, to inhibit them from performing transmissions while they are in the process of communicating with other nodes. At the same time, the framework allows for the identification and usage of unidirectional links at the routing layer. This, in turn, leads to shorter routes and consequently to improved perfor-

<sup>12</sup>If one were to instead consider the percentage of high power nodes as the level of heterogeneity, the observed results are almost identical. Note again that the level of heterogeneity is maximized when the fraction of the low power nodes is almost equal to the fraction of high power nodes in the network [15].

mance. We study the performance exclusively at the MAC layer and at the higher layers. Our cross layer framework can improve the transport layer throughput of low power nodes by up to 25% and alleviate the unfairness caused by the legacy IEEE 802.11 MAC. We also show a significant reduction (by 20%) in total number of interference related *false link failures* in the network. Our observations leads us to argue that our integrated MAC/Routing layer framework offers a simple yet viable and effective solution for handling asymmetry in power heterogeneous ad hoc networks.

## References

- [1] IEEE STD 802.11 1999. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999.
- [2] T.S.Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall, 2003.
- [3] S. Ni, Y. Tseng, Y. Chen and 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] N. Poojary, S.V. Krishnamurthy and S. Dao, "Medium Access Control in a Network of Ad Hoc Nodes with Heterogeneous Transmit Power Capabilities", *Proceedings of ICC 2001*.
- [6] V.Shah, S.V.Krishnamurthy and N.Poojary, "Improving MAC Layer Performance in Ad Hoc Networks with Heterogeneous Transmit Power Capabilities", *Proceedings of ICC 2004*.
- [7] E. Jung and N.H.Vaidya, "A Power Controlled MAC Protocol for Ad Hoc Networks", *ACM/IEEE MOBICOM 2002*.
- [8] J. Monks and V.Bharghavan, "A Power Controlled Multiple Access Protocol for Wireless Packet Networks", *IEEE INFOCOM 2001*.
- [9] E.Royer and C-K.Toh, "A Review of Current Routing Protocols for Ad hoc Wireless Networks", *IEEE Personal Communications Magazine*, April 1999.
- [10] J. Yoon, M. Liu and B. Noble, "Random Waypoint Considered Harmful", *IEEE INFOCOM*, 2003.
- [11] T. Saadawi and S. Xu, "Performance Evaluation of TCP Algorithms in Multi-hop Wireless Packet Networks", *Journal of Wireless Communications and Mobile Computing* 2002.
- [12] R. Ramanathan and R. Rosales-Hain, "Topology Control of multihop wireless networks using transmit power adjustment", In the proceedings of IEEE INFOCOM, volume 2, pages 404-413, 2000.
- [13] S. L. Wu, Y. C. Tseng and J. P. Sheu, "Intelligent medium access for mobile ad hoc networks with busy tones and power control", *IEEE Journal on Selected Areas in Communications*, 18(9): 1647-1657, 2000.
- [14] A. Muqattash and M. Krunz, "Power controlled dual channel (PCDC)Medium Access Protocol for Wireless Ad hoc networks", In the proceedings of IEEE INFOCOM 2003.
- [15] M. K. Marina and S. R. Das, "Routing performance in the Presence of Unidirectional Links in Multihop Wireless Networks", *ACM Mobihoc* 2002.
- [16] V. Ramasubramanian, R. Chandra and D. Mosse, "Providing Bidirectional Abstraction for Unidirectional Ad Hoc Networks", In Proceedings of IEEE INFOCOM, June 2002.
- [17] S. Nesargi and R. Prakash, "A Tunneling Approach to Routing with Unidirectional Links in Mobile Ad Hoc Networks", In Proceedings of International Conference on Computer Communications and Networks (IC3N), pages 522-527, 2000.
- [18] E. Duros, W. Dabbous, H. Izumiyama, N. Fujii and Y. Zhang, "A Link-Layer tunneling mechanism for Unidirectional links", RFC 3077, 2001.
- [19] L. Bao and J. J. Garcia-Luna-Aceves, "Link-state Routing in Networks with Unidirectional Links", In Proceedings of International Conference on Computer Communications and Networks (IC3N), pages 358-363, 1999.
- [20] P. Sinha, S. V. Krishnamurthy and S. Dao, "Scalable Unidirectional Routing using ZRP extensions for Wireless Networks", In Proceedings of IEEE WCNC, 2000.
- [21] R. Prakash, "A Routing Algorithm for Wireless Ad Hoc Networks with Unidirectional Links", *ACM/Kluwer Wireless Networks*, 7(6):617-625, 2001.
- [22] R. Prakash, "Unidirectional Links Prove Costly in Wireless Ad Hoc Networks", In Proceedings of the Workshop on the Discrete Algorithms and Methods for Mobile Computing and Communication (DIALM'99), pages 15-22, August 1999.