Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 51 (2007) 2594–2615
Computer Networks
www.elsevier.com/locate/comnet

Handling asymmetry in power heterogeneous ad hoc networks

Vasudev Shah, Ece Gelal *, Srikanth V. Krishnamurthy

Department of Computer Science and Engineering, University of California, Riverside, CA 92521, USA

Received 4 May 2006; received in revised form 19 September 2006; accepted 20 November 2006
Available online 20 December 2006

Responsible Editor: M. Venkatachalam

Abstract

Traditional MAC and routing protocols, which are primarily designed for homogeneous networks wherein all nodes transmit with the same power, suffer performance degradations when employed in power heterogeneous networks. The observed degradations are due to link asymmetry, which arises as high power nodes that do not *sense* the transmissions of low power nodes can potentially initiate transmissions that interfere with the low power communications. Link layer asymmetry in power heterogeneous networks not only disrupts the functioning of the routing protocol in use, but also results in unfairness in medium access. In this paper, we develop a cross-layer framework to effectively address the link asymmetry problem at both the MAC and the routing layers. At the MAC layer, the framework intelligently propagates low power control messages to the higher power nodes, so as to preclude them from initiating transmissions while there are low power communications in progress within their sensing range. At the routing layer, the framework facilitates the efficient use of unidirectional links. We perform extensive simulations to study the performance of our proposed framework in various settings, and show that the overall throughput in power heterogeneous networks is enhanced by as much as 25% over traditional layered approaches. In addition, we show that our schemes are also beneficial in power homogeneous settings, as they reduce the extent of false link failures that arise when the IEEE 802.11 MAC protocol is used. In summary, our framework offers a simple yet effective and viable approach for medium access control and for supporting routing in power heterogeneous ad hoc networks.

Keywords: Power control; Power-heterogeneous ad hoc networks; Link asymmetry; Cross-layer design; MAC; Routing

1. Introduction

Emerging ad hoc networks are likely to consist of devices with varying capabilities. One could envision low power sensor nodes and wireless hand-held devices integrated into a single network with higher power wireless devices such as laptops, wireless routers or wireless devices mounted on vehicles, powered by inside alternators. In such heterogeneous networks nodes are likely to transmit at different power levels, thereby causing communication links of varying range. In such networks *link asymmetry* is likely to be the norm. Link asymmetry may also appear in ad hoc networks when *power control* is employed, in order to reduce energy consumption or to enhance spatial re-use in the network (as in [11,20]). With link asymmetry, the transmission of

* Corresponding author. Tel.: +1 909 3798258.
E-mail address: ecegelal@cs.ucr.edu (E. Gelal).

1389-1286/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2006.11.016

a lower power node might not be received (or sensed) at a higher power node, while communication in the reverse direction could be feasible. As a result, traditional MAC and routing protocols that implicitly assume that links are *bi-directional* will either fail or perform poorly. At the MAC layer the hidden terminal problem is exacerbated [26]; routing becomes challenging due to presence of unidirectional links [19,24,3].

There have been a plurality of research efforts on alleviating the effects of link asymmetry in ad hoc networks. It has been shown in [26] that the performance of the IEEE 802.11 MAC protocol degrades in the presence of link asymmetry, since low power nodes cannot acquire the channel for sufficient durations. Link asymmetry at the MAC layer has also been studied in [4,18]; however, these studies are limited to certain scenarios with strict assumptions. In parallel, [19,24,3,2,31] propose methods for routing in the presence of unidirectional links; however, most of these studies ignore interactions with the MAC layer. Specifically, they either (i) employ the traditional IEEE 802.11 MAC protocol, which is inefficient with link asymmetry, or (ii) they assume an ideal MAC protocol which can only function effectively in the presence of unidirectional links. Finally, there are proposals to build link layer tunnels in order to hide unidirectional links from the higher layers [9,12]; again, MAC layer effects can impact these solutions.

The prior work in [26] considers the propagation of MAC layer control messages (in particular the IEEE 802.11 MAC protocol is studied) in order to alleviate the effects of asymmetry. However, the work considered variants of flooding for the above propagation and the overhead incurred was shown to be prohibitive; as a result, performance enhancements were not achieved. We first examine if reducing this overhead via some simple strategies can improve the MAC layer performance. In particular, we study the effect of (i) using an intelligent broadcasting scheme to quell unnecessary broadcasts, and, (ii) reserving the bandwidth for multiple data packets with a single RTS/CTS exchange (the multi-reservation scheme). Our simulations show that these extensions offer fairly limited performance enhancements (as compared to the legacy IEEE 802.11 protocol) in terms of the MAC layer throughput.

The above studies lead us to believe that a further reduction in the MAC layer control message propagation overhead could yield higher performance dividends. In order to do so, as opposed to relying on broadcasting strategies like before, we design a cross-layer solution, wherein, MAC layer control messages that are transmitted with low powers are *routed* to beyond the one-hop neighborhood of the low power transmitter. As before, the goal is to inform and silence the high power nodes that, being oblivious of low power communications within their sensing range, can potentially initiate transmissions that will collide with existing low power communications. However, with this approach the functionality is achieved with much lower overhead. In addition, our framework implicitly facilitates the *tunneling* of MAC and routing control packets via a *reverse path* (also referred to as the *inclusive cycle* [3]) that spans a unidirectional link. Tunneling supports the effective utilization of the unidirectional links at the routing layer. In summary, our framework effectively:

- Eliminates MAC layer inefficiencies in power heterogeneous ad hoc networks, thereby increasing throughput achieved at the MAC layer to a value that commensurates with that of a power homogeneous network.
- Identifies and effectively utilizes unidirectional links at the routing layer, thereby improving the performance in terms of throughput perceived at the higher layers (as compared to that with traditional routing protocols under power heterogeneous conditions).

We implement our framework in two steps. *First*, we modify the MAC layer to incorporate basic routing functionalities; this results in our **topology-aware CTS propagation (TACP) scheme**. With TACP, CTS messages are ''routed'' to the higher power nodes prior to any data exchange involving low power nodes. We combine TACP with the previously mentioned multi-reservation scheme. These two schemes are independent, and provide complementary improvements in performance. *Second*, we extend TACP to *tunnel* MAC and routing control packets in the reverse direction of a unidirectional link on a path that **spans** this link. (The structure that is formed for propagating the CTS message is used for this purpose.)

We perform extensive performance evaluations of our framework via simulations. First, we eliminate higher layer artifacts and examine the performance improvements exclusively at the MAC layer. In this stage, we do not tunnel packets or perform routing. We observe that the MAC throughput of the low

power nodes improves by approximately 24% as compared to the IEEE 802.11 MAC protocol. Next, we study the benefits offered by our framework as perceived at the transport layer (UDP). We show that the number of false link failures reported to the routing layer is reduced by up to 20%. In addition, the overall packet delivery ratio observed at the transport layer improves by up to 25% as compared with the traditional layered structure (with the IEEE 802.11 MAC protocol and AODV in place).

The rest of this paper is organized as follows: In Section 2, we describe related previous work. In Section 3, we briefly revisit the problems due to link asymmetry. In Section 4, we describe our methodologies in detail; in particular, we discuss: (a) specific MAC layer enhancements that were initially considered, and (b) our cross-layer framework. In Section 5, we present our simulation results and discuss the achieved performance enhancements. Our conclusions form Section 6.

2. Related previous work

In this section, we describe related work that address the effects of link asymmetry at various layers. Unlike most of these previous efforts, we consider a truly multifarious network wherein the nodes could have different maximum transmission power capabilities. Furthermore, our cross-layer framework enhances performance at both MAC and routing layers unlike almost all previous schemes which, in their design, ignore the effects of link asymmetry at other layers.

2.1. Link asymmetry due to power control at the MAC layer

Power control in MANETs has been explored in [5,11,20,26]. The authors in [11,20] propose power-controlled MAC protocols that incorporate the collision avoidance mechanism of the legacy IEEE 802.11 MAC protocol. With both protocols, the request-to-send (RTS) and the clear-to-send (CTS) frames are transmitted with a maximum preset power level, so that all nodes within the maximum range can hear the transmissions. Data frames are then communicated using the lowest power level that suffices for the intended communication to succeed. Both schemes avoid the effects of asymmetry by employing maximum-power transmission of control frames. These protocols are not able to exploit the spatial re-use gains that are potentially possible with power control. Furthermore, they are not applicable in networks consisting of multifarious nodes with different maximum transmission power capabilities.

A busy tone-based power control protocol is proposed in [33], where the 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 transmits only if its transmission does not add noise to the ongoing reception. This scheme as well, avoids the problems due to asymmetry, since nodes transmit busy tones at the preset maximum power levels; however, the aforementioned limitations with regards to spatial re-use are also valid for this protocol.

Muqattash and Krunz proposed both dual channel [21] and single channel [22] MAC protocols that enable power control in ad hoc networks. The approaches take into account the potential interferers while administering power control. However, as with the schemes in [11,20], these protocols require high power transmission of the control packets so that all potential interferers are reached. Such high power transmissions might not be possible in the multifarious, inherently power heterogeneous networks that we consider.

2.2. Link asymmetry due to topology control

Research on topology control in ad hoc networks via the deployment of multiple transmission power levels has particularly focused on (i) minimizing the overall energy consumption [14,6], (ii) bounding the node degree [17,15], and (iii) reducing the interference to increase effective network capacity [7]. The proposed solutions are based on adjusting the transmission power to avoid power-intensive communications [16,15,32]. Towards this, nodes choose a subset α of their neighbors to directly communicate with, and they reduce their transmission power (from the maximum) to the value sufficient to reach the farthest neighbor in α [16,15,32,6]. Ramanathan proposed a centralized topology control protocol to bound the energy consumption in the network [28]; with the proposed approach, nodes reduce their transmission powers to the extent possible without sacrificing the network connectivity. Similarly, the COMPOW protocol proposed in [23] requires the nodes to operate at the smallest power level at which the network is connected. This power level is computed by routing layer agents, and is used by all nodes in the network. Another approach for facilitating power control is clustering, as proposed in

[13]. With this approach, nodes adjust their power levels so as to communicate only with the clusterhead. The above efforts are geared towards topology control via power adaptations; they do not examine the effects of the created asymmetry on protocol performance at various layers.

2.3. Handling link asymmetry at the routing layer

Routing in the presence of unidirectional links has been studied in [2,3,19,27,31]. In [19], authors propose to bypass the unidirectional links and route the packets via only bi-directional links; this strategy may lead to the overloading of certain links, while under-utilizing others. In [3], reverse multihop paths are used to proactively discover and use unidirectional links. In [31], the authors extend the well-known zone routing protocol (ZRP) [10] to support unidirectional links. With these routing schemes, at the MAC layer either the traditional IEEE 802.11 MAC protocol is employed, or an ideal MAC protocol that can handle unidirectional links is assumed. While the former causes performance degradation due to unfairness (as mentioned before), the latter approach does not accurately reflect in the trends in the latter approach is unable to utilize system resources efficiently.

Link layer *tunneling* approaches to support routing in the presence of unidirectional links have been explored in [9,24]. These approaches hide the unidirectional nature of a link from higher layer protocols so as to facilitate their operations without any modifications. Tunneling is based on forming a reverse multi-hop path for each unidirectional link using the information gathered by the routing protocol. The reverse path is also sometimes referred to as the *inclusive cycle* [3]. A similar idea appears in [29], wherein a sub-layer beneath the routing layer is developed. The work also attests to the problems incurred in routing due to the lack of proper MAC layer protocols that handle asymmetry. There is also some work on using multi-hop acknowledgements to discover unidirectional links [2,24]. GPS-based approaches for enabling link level acknowledgements [12] over unidirectional links have also been proposed. The above link layer approaches however ignore the implications of link asymmetry at the MAC layer. In [19], the impact of unidirectional links on routing performance is studied and it is once again identified that more efficient MAC protocols are needed to handle unidirectional links.

2.4. Study of medium access in power heterogeneous networks

In [26], Poojary et al. quantify the inefficiencies in the use of the IEEE 802.11 MAC protocol in a network wherein nodes have heterogeneous power capabilities. A heterogeneous network with two types of nodes, operating at power levels of 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. This degradation was a consequence of the fact that the transmissions of low power nodes were often interfered with transmissions from high power nodes that were unable to hear the RTS/CTS exchange between the low power nodes. The authors proposed to solve this problem by means of a CTS propagation technique using a standard *flood-type* broadcast algorithm. The algorithm required that nodes that hear a CTS frame would propagate the frame further (up to a distance determined by the ratio of the range of high power nodes to the range of low power nodes). The propagated frames were called the bandwidth reservation (BW_RES) frames. The objective of this broadcast was to inform the high power nodes in the neighborhood about the ongoing communication so that they would inhibit their own transmissions for the duration specified in the BW_RES frame. It was found that such a *flood type* broadcast did not offer performance benefits; in fact, the scheme caused a further degradation in terms of throughput since the overhead incurred in propagating these BW_RES control frames outweighed the potential gains in terms of reducing the number of collisions. This work neglects the presence of "sensing range" with the IEEE 802.11 MAC protocol [11]; furthermore, the effects at higher layers have not been considered.

In [4], Bao et al. propose a MAC protocol for ad hoc networks with unidirectional links. The proposed approach schedules node transmissions based on the contention in their one and two-hop neighborhoods such that fairness is ensured. The scheme however, depends on the existence of accurate 2-hop neighborhood information at each node and requires perfect clock synchronization. Our proposed framework is designed to solve the problems due to link asymmetry at the MAC layer, without the requirements or constraints that limit the practicality of this solution.

In [18], Liu et al. propose a three phase approach for power heterogeneous ad hoc networks. The three phases are time division multiplexed. Link asymmetry is controlled by allowing the heterogeneous power transmissions only in one of the phases. The scheme imposes that within this phase, at any given time, only a single high power node can be active. This is because the MAC scheme does not utilize the CSMA/CA capability of the IEEE 802.11 MAC; it requires perfect scheduling between high power transmissions. The paper does not address these scheduling issues. This constraint may cause an under-utilization of the available bandwidth. In addition, the proposed scheme works efficiently only in scenarios with low mobility or long pause durations.

2.5. Summary

To summarize, the majority of the previous efforts only address the link asymmetry problem at specific layers. Most of the efforts propose solutions for routing via unidirectional links. There are a few approaches that study link asymmetry from the perspective of the MAC layer. However, they are limited in scope. In particular, the distinguishing contributions of our work in contrast with existing literature are:

- Most previous approaches that propose solutions for routing via unidirectional links ignore MAC layer effects. In particular they either assume an ideal MAC scheme or simply use the IEEE 802.11 MAC without modifications. Our scheme is the first to consider the interactions between the layers.
- Existing MAC layer solutions for handling link asymmetry are based on imposing time-schedules. These schemes require frame/clock synchronization and in some cases, the use of homogeneity within the network. Furthermore, the previous studies do not consider the impact at higher layers (routing in particular). Our cross-layer approach overcomes these limitations.
- Similarly, previous power control based media access schemes typically handle link asymmetry by assuming that all nodes can transmit with a fixed maximum transmission power level. Our approach is applicable in a truly heterogeneous setting where the above assumption does not hold.

3. Problem statement

In this section we revisit[1] the performance of the distributed coordination function (DCF) of the IEEE 802.11 MAC protocol in a power heterogeneous network setting. We discuss its deficiencies and highlight the resulting effects on the higher layers.

As alluded to earlier, the inefficiency of the IEEE 802.11 MAC protocol arises primarily in scenarios with link asymmetry. Link asymmetry causes lower power nodes to be *hidden* from higher power nodes. This, in turn, increases the number of collisions that are experienced by low power communications. This effect is depicted in Fig. 1. The RTS/CTS exchange between two low power nodes A and B is not overheard by node H since H is not within the sensing or interference range[2] of these communicating nodes. Thus, it is possible that while the data exchange between nodes A and B is in progress, node H could initiate another transmission, and cause a collision at node B.

A second problem that is manifested at the MAC layer occurs when a node fails to identify (and utilize) a unidirectional link. This effect is depicted in Fig. 2, where H_1 can reach L_1 but not vice versa. As a result, if L_1 responds to any frame (such as an RTS frame) sent by H_1, the response never gets to H_1. Similarly, any control frame initiation by L_1 (e.g. an RTS frame) would never reach H_1. Depending on the scenario, these problems could cause degradations due to wasteful control frame transmissions and back-offs at the MAC layer. The link asymmetry can also degrade the performance of traditional on-demand routing protocols[3] due to the loss of control packets. One such effect is depicted in Fig. 2 where node H_1 attempts to establish a route to H_2 through nodes L_1 and L_2. The

[1] The problems have been discussed in detail in [26,19,29,3].

[2] Typically two ranges are defined for the transmissions of a given node u viz., the transmission range and the interference range [1]. Nodes that are within the transmission range of u can decode the frames received from u. Nodes that are within the interference range but not within the transmission range of u cannot decode frames from u; however they still interfere with a transmission of u.

[3] Lost route update packets can lead to falsified routing tables when traditional proactive routing schemes are used. A low power node could wrongly assume that it can reach a high power node and a high power node may not know that it could reach a low power node. We omit the detailed discussion of these effects in this work since they are discussed in prior papers on routing on unidirectional links.

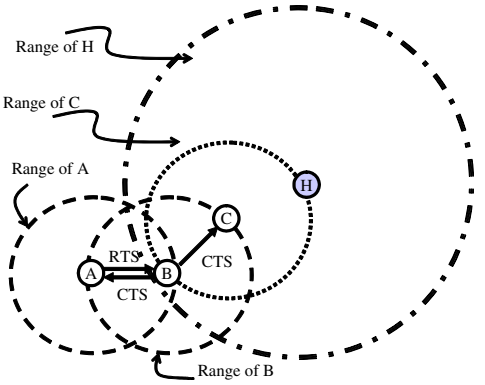

Fig. 1. Problem at the MAC layer due to link asymmetry in the power heterogeneous network.

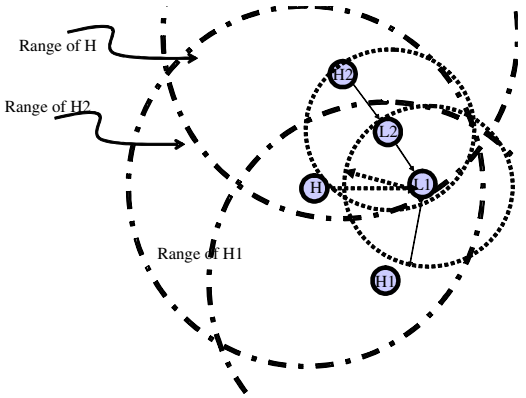

Fig. 2. Problem at the MAC/routing layers due to link asymmetry in the power heterogeneous network.

routing control packets from L_1 is not received by H_1 since it is outside the range of L1. Such effects could lead to repeated (albeit unsuccessful) route discovery attempts. We do not discuss these in great detail since prior work on unidirectional routing touch upon such problems [3,19].

4. Our cross-layer framework

In this section, we present our cross-layer framework for efficiently handling link asymmetry at the MAC and routing layers. The key idea of our approach is the selective forwarding of low power control packets to hidden high power nodes. In this section, we first introduce the components that aid the implementation of our cross-layer approach. Next, we study two extensions for a modular MAC layer solution, and discuss their performance enhancements. Finally, we present our cross-layer approach.

4.1. Preliminaries

In the following, we first revisit the Bandwidth Reservation (BW_RES) frame structure introduced in [26]; next, we compute the effective hop count up to which this frame should be forwarded such that asymmetry is eliminated.

The BW_RES frame: This frame is used to prevent high power nodes from initiating transmissions that may collide with ongoing lower power communications. The key idea is to have the nodes that hear a low power CTS frame broadcast a copy of the frame (the BW_RES frame) to reach nodes that are more than one low power hop away (Fig. 3). Reception of a BW_RES frame at a high power node indicates that a lower power communication is active in its vicinity, and that the node should inhibit its transmissions (for the time period specified in the BW_RES frame).

We modify the BW_RES and CTS frame structures to aid the realization of the proposed performance enhancements; the new frame structures and a timing diagram of the modified reservation scheme depicting the BW_RES transmissions are shown in Fig. 4a and b. The BW_RES frame format is similar to the RTS frame format except that the frame control field has a few additional attributes: (i) a sequence number, and (ii) the originator address field that contains the MAC address of the RTS sender. Our framework uses these fields for detecting duplicate BW_RES frames that may be received by third party nodes. A sequence number field is similarly added to the CTS frame structure. In the "frame control field" of CTS/BW_RES frames, we use the *To DS*, *From DS* and *MORE* bits

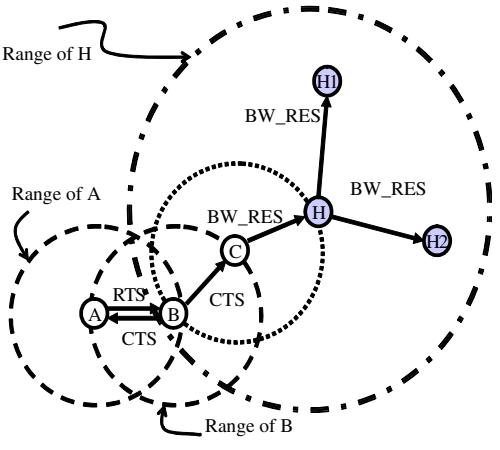

Fig. 3. CTS propagation scheme.

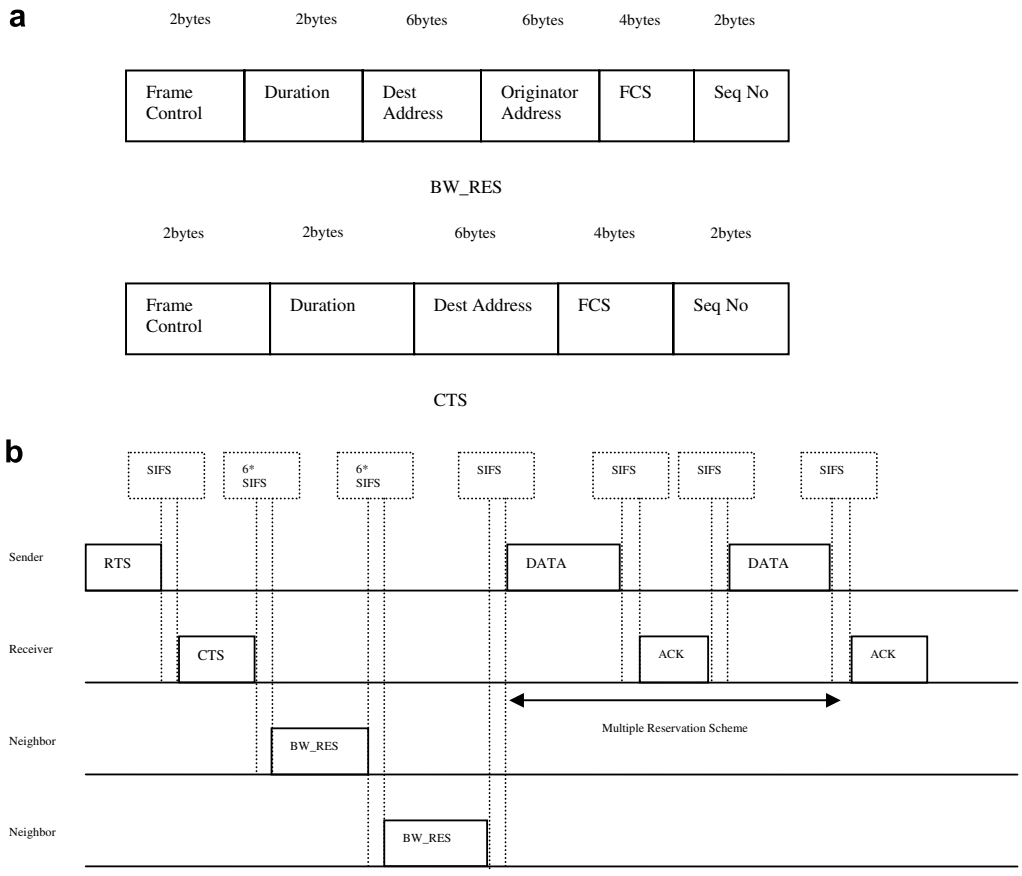

Fig. 4. Modifications on the IEEE 802.11 MAC protocol and on the scheme proposed in [26] to offer performance enhancements. (a) Modified frame structures. (b) Modified reservation scheme.

(used for signaling in infrastructured wireless networks in IEEE 802.11 [1]) to indicate a TTL value for the BW_RES frames. In our experiments we use TTL = 2; the choice for this value is justified in the following. With our cross-layer framework, the initiator of the CTS frame chooses the set of nodes (say L in number; L is typically small) that are to forward the BW_RES frame; the list of these L nodes is appended into the CTS and BW_RES frames. In our simulations we account for the overhead incurred due to this additional information in these frames.

Nodes that overhear the BW_RES frames set their network allocation vectors (NAVs) appropriately (the duration of the communication is indicated). If a node simply *senses* energy due to one or multiple BW_RES frames and is unable to decode a BW_RES frame, it simply sets its NAV to indicate that the medium is busy for an extended inter-frame space (EIFS) interval as with traditional IEEE 802.11.

We remark that the BW_RES frame is as small as the CTS control frame, except that it carries the IDs of the nodes that are to forward this frame. Hence, the collision of these frames has a small likelihood. Furthermore, nodes perform carrier sensing before transmitting BW_RESs; this further reduces the probability of a collision. However, if the BW_RESs collide, a subset of the high power nodes may not hear them; they will however sense the channel to be busy and therefore set their NAVs for an EIFS period.[4] Still, these nodes may initiate transmissions that collide with low power communications; in our simulations we find that these effects are not pronounced.

In the following, we determine the appropriate distance (in terms of hop count) up to which the

[4] This approach is similar to the IEEE 802.11 MAC protocol, where, upon the collision of RTS and/or CTS messages, nodes simply use physical carrier sensing to set their NAVs for an extended inter-frame space period.

BW_RES frames should be propagated. Specifically, we justify the selection of value 2 for the TTL of BW_RES frames.

At the physical layer, we adopt the path loss channel model from literature. With this model, if P_{tx} is the transmission power of a node u, then the received power at a distance d from u is given by $P_{rx} = P_{tx} \cdot d^{-\alpha}$, where α is the path loss exponent depending on the wireless transmission environment. We denote the power levels of high power and low power nodes as $P_{\max}$ and $P_{\min}$, and the corresponding transmission ranges as $d_{\max}$ and $d_{\min}$, respectively.

The ad hoc nodes define a unit disk graph (UDG), wherein a link exists between two nodes if and only if they are within a unit distance from each other, where the maximum transmission range is mapped to unit distance. While it is known that the transmission range of a node can be affected by wireless channel impairments such as multi-path fading and shadowing, we consider only the path loss to simplify our analysis. According to this model, the signal-to-noise ratio (SNR) of a frame received at the periphery of a unit disk is the *threshold* value that is necessary for the correct decoding of a received frame. This SNR threshold is the same at both high and low power nodes. Our objective is to ensure that the CTS frame of a low power node u reaches all high power nodes that have u within their transmission range. Given the UDG model, we note that:

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

For the chosen[5] values of $P_{\max}$, $P_{\min}$ and k we obtain $d_{\max} = 2 \cdot d_{\min}$. Thus, if the BW_RES frame of a low power node u is propagated through two low power hops, the frame can with *high probability* reach the high power nodes whose transmission is received at u. We may refine our model to distinguish between the sensing and the correct reception of a packet (in the former the packet is received but cannot be correctly decoded due to its low SNR value). In literature, the sensing range is typically modeled to be twice the transmission range; the frame must be propagated through four low power hops (or respectively two high power hops, from (1)). Choosing a TTL value of 2 is seen (in our simulations) to provide the requisite benefits; however, we remark that it *does not guarantee* that a BW_RES frame will reach *all* high power nodes that may potentially initiate colliding transmissions (high power nodes may be more than 2 hops away from a low power node). Yet, we support with our simulation results in various sce- narios that it is more effective to restrain the propagation to a localized vicinity of the communication as opposed to invoking a broadcast with a wider scope.

[5] In our simulations we chose $P_{\max} = 0.56$ W, $P_{\min} = 0.14$ W, and $k = 2$, representing propagation in free space. These values are parameters input to the simulations and can be altered for different scenarios. The behavioral results that we observed (via sample runs) with other scenarios were similar to the ones reported here.

4.2. Enhancements at the MAC layer

Before designing our cross-layer framework, we consider simple modifications to the IEEE 802.11 MAC protocol to enhance the MAC layer performance in power heterogeneous networks. In this subsection, we describe these modifications (namely the multi-reservation and the intelligent broadcasting techniques) and explain our intuitions for making the reported design decisions. The proposed approaches offer enhancements at the MAC layer, but introduce an overhead to the ad hoc network. Furthermore, these "*MAC layer enhancements*" do not provide support for routing over unidirectional links. We address these problems in the next subsection, where we propose our "*cross-layer framework*", which provides methods to handle asymmetry at the routing layer as well.

4.2.1. Reserving bandwidth for multiple sequential transmissions

As discussed earlier, [26] considered the scoped flooding of the BW_RES frame to preclude high power transmissions in the vicinity of ongoing low power transmissions; however, the overhead incurred was prohibitive. To limit the overhead due to broadcasts of BW_RES frames, we attempt to reduce their frequency. One way of achieving this is using a *single* RTS/CTS/BW_RES initiation for *multiple* sequential DATA/ACK exchanges. The multiple DATA frames can be *independent*, i.e. they may have separate fields, including the checksum. Before sending an RTS, a node checks its interface queue (between the network and the MAC layers)

for other DATA frames with the same destination address as that of the RTS. If such frames are found, they are moved to the MAC layer and buffered along with the original (RTS) frame to be sent. The RTS frame attempts to reserve the channel for a preset number (say N) of data frames at a time. When a node does not have additional frames for the destination (i.e., there is only one frame that can be transported at the given time), it reserves the channel for the single data frame.

Clearly, allowing a node to reserve bandwidth for a large number of sequential transmissions could lead to unfairness in the network. On the other hand, if only a small number of transmissions are reserved with a single control frame, a significant reduction in the number of generated BW_RES frames may not be feasible. We find by simulations that reserving the channel for 2 and 3 sequential transmissions provides a significant reduction in the volume of broadcast BW_RES frames. In our future discussions we refer to this technique as the **multi-reservation technique**. The sequence of frame transmissions due to the multi-reservation technique is depicted in Fig. 4b.

4.2.2. Intelligent broadcasting

To further reduce the volume of broadcast BW_RES frames we examine the benefits of excluding *unnecessary* re-broadcasts. Our approach is derived from a broadcast scheme proposed in [25], and it executes as follows. Each node u, upon receiving a BW_RES frame, sets a randomly chosen time-out in the future. It then records (in a counter) the number of subsequent BW_RES broadcasts (all of which correspond to the same CTS frame) that it overhears via broadcasts from its neighbors prior to the time-out. If the value indicated by the counter *exceeds* a preset fixed threshold T (a tunable system parameter), u revokes its own BW_RES broadcast. The idea behind this scheme is that overhearing multiple copies of the same broadcast frame suggests a node that it is located in a highly dense neighborhood, which implies that the additional coverage it would achieve by performing its own broadcast is likely to be insignificant. We find by simulations that a choice of $T = 3$ preserves coverage and significantly reduces the number of unnecessary broadcasts. We observe that at lower values of T there is a reduction in coverage (nodes quell their transmissions to a large extent), and for higher values there are broadcasts that do not offer an enhancement in coverage.

4.2.3. Failures and retransmissions

The actions taken upon transmission failure are the same as performed 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 with the IEEE 802.11 MAC specifications, we impose an upper bound on the number of retransmission attempts for RTS and DATA packets (7 and 4, respectively). Successive packets after these pre-specified number of attempts are dropped. If a sender does not receive an ACK for a specific DATA packet among the multiple sequentially attempted transmissions (with our multiple reservation scheme), it attempts to retransmit only the particular failed DATA packet after backing-off. During the retransmission of this packet, the sender will again attempt to reserve bandwidth for additional packets destined for the same node.

4.3. Integrated MAC/routing framework

In the following, we describe our **cross-layer** framework for supporting medium access control and routing in power heterogeneous ad hoc networks. Specifically, our cross-layer framework improves the enhancement achieved at the MAC layer by the mechanisms introduced in the previous subsection; in addition, it introduces methods to handle the asymmetry at the routing layer. In our framework, the MAC layer solicits assistance from the routing layer in determining a small subset of nodes that will perform the BW_RES re-broadcasts. In return, the routing layer depends on the MAC layer for the discovery and use of unidirectional links. Thus, the key functionalities of our framework are the intelligent propagation of BW_RES messages and the construction of reverse routes for bridging unidirectional links.

4.3.1. Topology-aware CTS propagation (TACP)

We propose a *routing-assisted approach* for reducing the number of propagated BW_RES frames. With this approach, nodes *multicast* the BW_RES frame to the high power nodes in their vicinity, as opposed to broadcasting this frame.

To facilitate this, we require that each node maintains link-state information with regards to its two-hop neighborhood [8]. This information is collected via *Hello* messages as follows. At network

instantiation, every node broadcasts[6] a list of its one-hop neighbors (we call these one-hop neighbors of a node u, as the *in-bound neighbors* of u) that it is currently receiving *Hello* messages from. These *Hello* messages also contain the corresponding maximum transmission power (in watts) for each inbound neighbor included. After the initial phase each node constructs an *inbound tree* with the *Hello* messages it has received. The inbound tree of a node u includes all neighbors that can reach node u. As the network reaches a steady state, nodes begin transmitting *Update Hello* messages that are now modified to contain their inbound trees. These messages are broadcast every "Hello interval" milliseconds. Each node then combines the *in-bound trees* reported by its neighbors with that of its own, and forms a *localized graph* that depicts its local neighborhood. The *Update Hello* messages are further modified to include this localized neighborhood. Periodic transmissions of the *Update Hello* messages help refine this *localized graph*.[7] As information propagates, the localized graphs become more extensive. This allows nodes to gather additional information (beyond their two hop neighborhoods); however, the sizes of *Update Hello* messages grow considerably. 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; and this pruned graph is included in the *Update Hello* messages. In each update, by transmitting *only the changes* to the localized graph as compared with the previous update, one may significantly constrain the size of the *Update Hello* messages. As mentioned earlier, a choice of n between 2 and 4 ensures to a great extent that most of the high power nodes that affect a given low power communication are informed by means of BW_RES frames. Our simulations suggest that setting $n = 3$ offers the best benefits[8]; our sample studies suggest that with $n = 2$, a significant fraction of the high power nodes are missed by the BW_RES broadcasts due to collision effects while $n = 4$ does not offer extended coverage benefits. We justify the chosen value for this parameter, based on our simulation results in Section 5.

Using the localized graph, our objective is then to have each node construct an n-hop *outbound* Steiner tree on which the BW_RES frames will be **multicast** to the high power nodes. Note that a low power node u does not initiate the propagation of the BW_RES frame if there are no high power nodes in its n-hop neighborhood. If high power nodes exist within the n-hop neighborhood of u, u 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 frame. Node u includes the IDs of these nodes in its CTS frame. As our goal is to minimize the number of BW_RES re-broadcasts and to reduce (to the extent possible) 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). The IDs of this next set of candidate relays are included in the BW_RES frame. If a node, upon receipt of either a CTS or a BW_RES frame does not find its ID in the frame, it simply updates its NAV (network allocation vector) in accordance with the IEEE 802.11 MAC protocol and discards the frame. The multi-reservation scheme is incorporated as well; if possible, a node reserves the channel for N data frames destined for the same neighbor by means of a single RTS/CTS/BW_RES initiation. Since the BW_RES frame is multicast to the optimum set of one-hop neighbors along the node's Steiner tree (as opposed to simple broadcasting), the number of propagated BW_RES frames does not increase with increasing local density. In fact, a better set of forwarding nodes could be viable at a higher density and this would further decrease the BW_RES overhead. We provide an algorithmic representation of the topology-aware propagation scheme in Fig. 5.

4.3.2. Construction of reverse routes for bridging unidirectional links

Our framework also proposes a cross-layer approach, to assist routing in the existence of link asymmetry. The key idea is to construct reverse routes that span the *unidirectional* links. In the

[6] *Hello* messages have a TTL (time to live) value of 1, i.e. they are only exchanged between one-hop neighbors.

[7] The periodicity of the *Update Hello* messages would depend on the mobility in a given scenario. If nodes are highly mobile, the *Update Hello* messages must be transmitted with higher frequencies.

[8] This implies that all high power nodes that can be reached via three low power hops from a low power communication are informed of the impending communication.

In order to determine the candidate nodes we assume that the following sets have already been computed by a node using periodic hello messages. These sets are inputs to the procedure Compute_Potential_Candidate which return a List ***F consisting of candidate nodes****.*

H_{x1} = *Set of all high power nodes in one hop neighborhood of x.*
H_{x2} = *Set of all the high power nodes that are 2 hop neighbors of x and not in* H_{x1}.
L_{x1} = *Set of all low power nodes in the one hop neighborhood of x.*
L_{Hx2} = *Set of low power nodes in* L_{x1} *that have high power nodes in their neighborhood.*

A node executing the procedure Compute_Potential_Candidates should satisfy following conditions

1. It should have its identifier in the received CTS or the BW_RES message.
2. The TTL field in the received message should be greater than zero.

If above conditions are not met then the node just updates its NAV as specified by the standard legacy IEEE 802.11 MAC and does not take any further action. Let S be the nodes that hear the CTS or BW_RES message and let the indicated TTL be higher than zero.

```
Procedure Compute_Potential_Candidates (Hx1 , Hx2 , Lx1 , LHx2)
{
  For each node x in S
  Do
        Add  Hx1 to F
        If
              Hx1 = nil and LHx2! =nil
        Then add LHx2 to F
         Else
              If ( LHx2! = nil )
              Then add to F those nodes of LHx2 that can reach high power nodes not in
              Hx1 and not reachable by nodes in Hx1.
              Else
                   if( LHx2 = nil and Hx1 = nil)
                        F = nil (if node is replying to a RTS then just broadcast CTS
                        message as in 802.11 MAC signaling else do not propagate CTS
                        or BW_RES further)
                        update NAV
return F
}
```

Fig. 5. Algorithm to select potential candidates.

following, we describe how these reverse routes are formed.

The exchange of the previously discussed *Hello* messages help in detecting the unidirectional links in the network. A node detects that it is at the *tail* of a unidirectional link, when it receives a *Hello* message from a neighbor and finds that it is excluded from this neighbor's ''neighborhood list''. Unidirectional links are also depicted in nodes' *pruned localized graphs*. Using this graph, a node that is at the tail of a unidirectional link can compute a reverse path to the node at the *head* of this unidirectional link. We illustrate this process in Fig. 6. In this figure, the link between nodes L_1 and H_1 is unidirectional (H_1 can reach L_1 but L_1 cannot reach H_1). To utilize this link, L_1 constructs

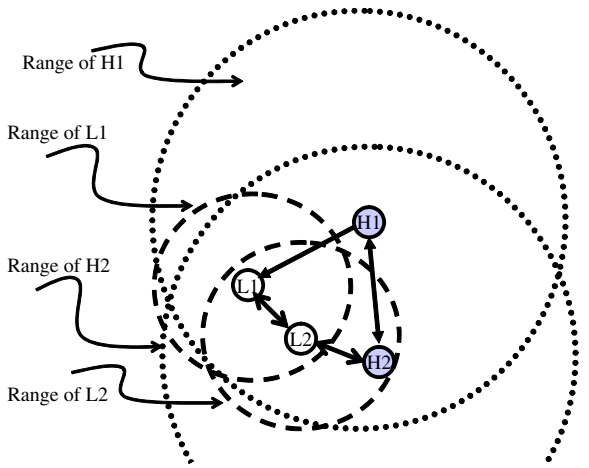

Fig. 6. Reverse route construction to route MAC/routing control packets.

a reverse route to H_1. L_1 learns from L_2's inbound tree (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 is via nodes L_2 and H_2. If a reverse path from L_1 to H_1 of less than n hops does not exist, the unidirectional link (H_1, L_1) will not be 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 found, L_1 encapsulates the control frames from the MAC and routing layers into an IP packet and routes (or *tunnels*) this packet using the constructed reverse path. This *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 network-wide routing. If either the reverse path or the unidirectional link were to fail, the *tunneled bi-directional link* would break. The network-wide on-demand routing protocol would then instigate a route error packet.

With our scheme, *Hello* messages are tunneled to the node at the *head* of the unidirectional link, in order to inform this node of the existence of the link. We also tunnel the MAC layer control frames (namely the CTS and ACK frames) and the routing layer control packets (namely the RREP packet, as proposed in [30]). The motivation for tunneling these frames/packets is, that significant benefits are achieved when the *unidirectional link* is utilized (thereby potentially avoiding a longer alternate path).

In order to distinguish between tunneled MAC layer frames and network layer packets, the packet header at 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 frame 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.

Finally, we note that our tunneling scheme is similar in spirit to those proposed in [9,12]. The key difference is, that our integrated framework (TACP) for handling MAC layer asymmetry provides us with a simple and seamless way of discovering reverse routes to bridge unidirectional links.

5. Performance evaluation

We have performed simulations using the event driven network simulator *ns2* (v.26). We divide the evaluation of our proposed enhancements into two parts. *First*, we exclusively measure the enhancements at the MAC layer (as proposed in Section 4.3) and quantify the benefits as compared to the performance of the IEEE 802.11 MAC protocol and also the performance reported in [26]. *Second*, we evaluate the performance enhancements offered by the implementation of our framework at the higher layers. We consider various scenarios and discuss the observed results.

5.1. Performance evaluation at the MAC layer

5.1.1. Simulation setup

Towards our first goal, to decouple the effects of routing and transport layer artifacts from the MAC-layer throughput in this study, we introduce a *Poisson* traffic generation agent above the MAC layer. The traffic generation rate is 1000 packets/s (unless stated otherwise), each data packet being 1000 bytes in size. We also vary the average packet generation rate in order to observe the effects with different system loads. When a data packet is generated at a node, it is randomly destined for one of the node's neighbors.

The network consists of 40 nodes deployed in a square region. We vary the area of this region in order to consider different network densities. The ratio of high power to low power nodes is 1, i.e. 50% of the nodes are high power nodes and 50% are low power nodes. This choice is based on the results in [19], which suggest that in the network a maximal number of unidirectional links exist for this value. We also studied various other parameter settings; however the observed results and the interpretations thereof are similar to those reported and thus, are not included.

The physical layer is based on the IEEE 802.11 specifications. Nodes move in accordance to a modified version of the *random waypoint* mobility model with a constant speed of 6 mps, and pause for 0.1 s at each destination point on the plane. This constant speed is chosen in light of the recent results suggesting that a choice of random speeds is not appropriate for a realistic modeling of mobility [34].

Table 1
Simulation setup for evaluation at the MAC layer

Simulator	ns2 (version 2.26)
Number of nodes (high power/low power)	40 (20/20)
Length of square grid	Varied from 300 to 2000 m
Power levels	0.56 W and 0.14 W
Traffic model	Poisson, 1000 packets/s
Mobility model	Modified random waypoint
Node speed, Pause time	6 m/s, 0.1 s

Depending on the speed simulated, nodes choose a frequency between 0.3 and 0.5 s, and generate *Hello* messages with this period. The parameters[9] used for this first set of simulations are tabulated in Table 1.

We assume that the communication channel is obstacle-free and that signal degradation occurs only due to path loss (as in other previous work [26,11,20]). We also assume that the channel is symmetric, and asymmetry occurs due to differences in transmission power levels. The channel bandwidth is set to 2 Mbps. All MAC control frames are transmitted at 1 Mbps and data at 2 Mbps, so as 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 significant benefits.

5.1.2. Parameters

We vary the **node density** and the **traffic load**, and we observe the effects on the performance, in terms of our metrics that we define below.

5.1.3. Metrics

The primary metrics of our interest in this first part of simulations are:

- *Data success rate*. The percentage of successful data frame transmissions *given* the RTS/CTS exchange between the pair of communicating nodes is successful.
- *Throughput efficiency per node*. We define this metric to quantify channel usage; it is the ratio of the time spent by a node in successfully transmitting data, to the total simulation time.

Using these metrics, we measure and compare the performance of the following:

1. Intelligent broadcasting (IB) and multi-reservations (MR)
   - *Case (a)*. The legacy IEEE 802.11 MAC protocol.,
   - *Case (b)*. 802.11 MAC protocol with BW_RES propagation using IB,
   - *Case (c)*. 802.11 MAC protocol using MR,
   - *Case (d)*. 802.11 MAC protocol with BW_RES propagation using IB *and* MR.

   To observe the performance of the above four cases on the low power and high power nodes, we define **Variant I** and **Variant II**, which stand for (i) only the low power nodes perform the intelligent BW_RES propagation and multi-reservations when initiating a communication, and (ii) all nodes, regardless of their power level, use these enhancements, respectively.
2. Topology-aware CTS propagation scheme (TACP)
   - *Case (e)*. The IEEE 802.11 MAC protocol with BW_RES propagation via TACP,
   - *Case (f)*. The IEEE 802.11 MAC protocol with BW_RES propagation using TACP *and* MR.

5.1.4. Simulation results

First, we examine the performance of the simple MAC layer enhancements (Case (d)) with both Variant I and Variant II, and compare their performances with that of the legacy IEEE 802.11 MAC protocol (Case (a)). The data success rate of low power nodes with both variants is depicted in Fig. 7a. As compared to scenarios with the legacy IEEE 802.11 MAC protocol, the low power nodes see an overall improvement of up to 14% with Variant I (and up to 12% with Variant II) with our schemes. Clearly, low power nodes benefit more when only these nodes use the proposed techniques in the network (Variant I); the data success rate for the low power nodes is better by about 2% with Variant I than with Variant II. However, note from Fig. 7b that the overall data success rate is better with Variant II than with Variant I. This is because with Variant II, high power nodes can also benefit from our schemes; specifically, the

[9] To draw a fair comparison we try to be consistent with the scenario in [26]; this dictates our choice of number of nodes and transmission power levels, as shown in Table 1.

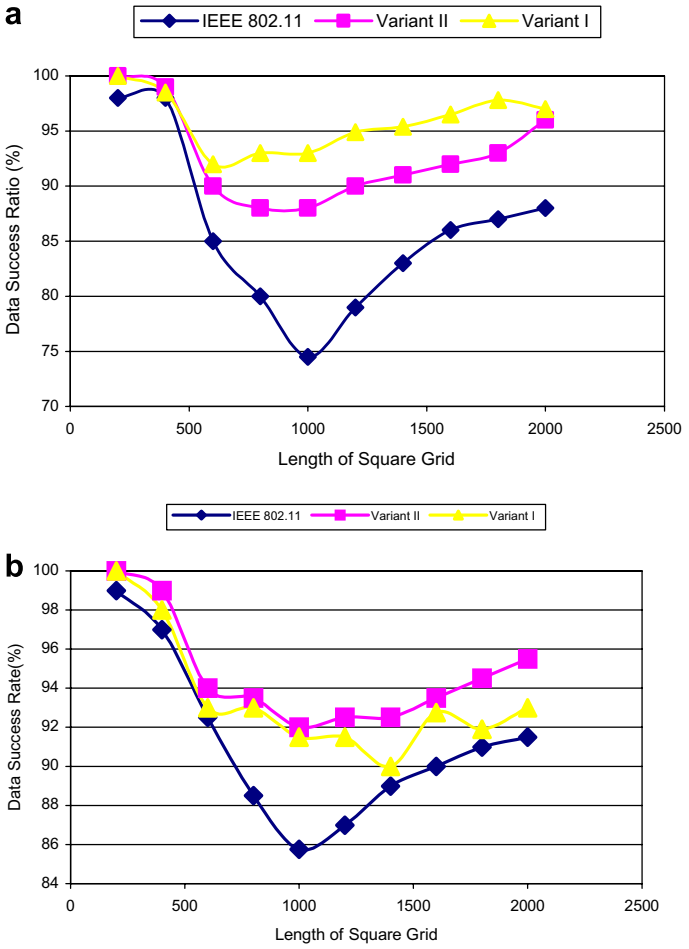

Fig. 7. Data success rate in the heterogeneous network with MAC layer enhancements. (a) Data success rate of low power nodes. (b) Overall data success rate.

effects of *false link failures*[10] are alleviated. As shown in Fig. 8, our proposed framework is able to reduce the number of false link failures by up to 20%.

We observe in Fig. 7a and b that the data success rates with the raw IEEE 802.11 MAC protocol and with the MAC layer enhancements are high, both at low and at high densities; however, there is a degradataion observed at moderate node densities. At low densities, the probability of a low power communication being interfered by a high power communication is small (the probability of having a high power node in its vicinity is small). At high densities, the nodes are close to each other and hence the effects of link asymmetry are not severe. Therefore, the performance in these two extreme

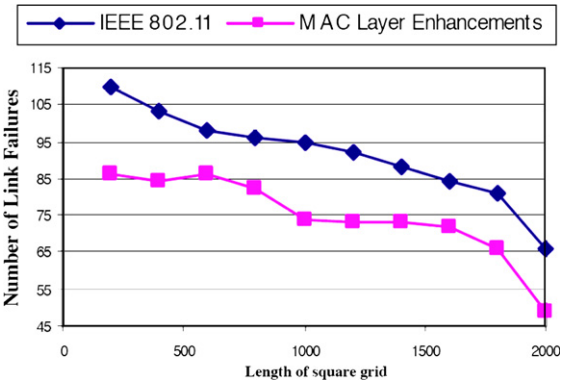

Fig. 8. Link failures in the heterogeneous network.

[10] False link failures occur if a node deems that a link has failed as a consequence of successive failed RTS transmission attempts. This happens if the receiver is being interfered with some other transmission.

cases is fairly good. At moderate densities, the link asymmetry effects are more pronounced, as the probability of a low power communication being in the vicinity of a high power transmitter is not insignificantly low. Hence the data success rate is low in this regime. The MAC layer enhancements do alleviate the degradation but cannot completely eliminate the effect.

In the same setting we also deployed a homogeneous network where all nodes have the same transmission power of 0.28 W; we observed an increase of up to 10% in network throughput, and a reduction in the number of false link failures by about 15% as compared with the IEEE 802.11 MAC protocol. In sum, the modifications that we consider are shown to provide benefits at the MAC layer in both power heterogeneous and power homogeneous networks. Depending on the percentage of low power nodes in the network one might prefer to use Variant I or Variant II. The former would provide better performance if the fraction of low power nodes in the network is large. Since the number of low power nodes may typically vary in the scenarios considered, Variant II would be the preferred design specification.

The throughput efficiency offered by Case (d) – Variant *I* is shown in Fig. 9 for different node densities. We observe that the throughput of low power nodes improves by as much as 14% as compared to the IEEE 802.11 protocol. High power nodes also benefit with our enhancements, since reducing the number of retransmissions of low power nodes decreases the overall contention for the wireless medium. Furthermore, the aforementioned effects of false link failures are alleviated. In sum, we observe an overall improvement in the network throughput of up to 12% as compared to the IEEE 802.11 MAC protocol. Fig. 9 indicates that at high node densities (when the area of deployment is small) the performance improvement with our scheme is not significant. As alluded to earlier, this is because, even with a low power (of 0.14 W) the transmission range is about 205 m, which implies that all nodes in the deployment area are within the sensing range of each other, i.e., no asymmetry exists. At lower densities, however, link asymmetry in the network increases and it is when the benefits of our schemes become significant. The network becomes highly sparse when the size of the deployment area is further increased. This reduces the possibility of collisions; this case does not pose a significant challenge on the MAC throughput and therefore the gains of our scheme are again not visible.

We observe that the frequency of multi-reservations increases with load (Fig. 10), since at higher loads when a node wishes to initiate a transmission, it is more likely that the node finds multiple packets destined for the same neighbor. Thus the efficiency of the scheme improves with load: the improvement in throughput is about 5% at a load of 500 packets/s, whereas with a load of 1000 packets/s, it is about 12%.

In the following experiments, we evaluate the performance of TACP exclusively (Case (e)) and in conjunction with the MR scheme (Case (f)), and compare this with the performance with Case (b) and Case (d).

We first compare the data success rate with Case (f), with that of the legacy IEEE 802.11 MAC protocol (Case (a)) for low power and high power nodes at different node densities. As shown in Fig. 11, in the power heterogeneous scenario the low power nodes see a significant improvement (of

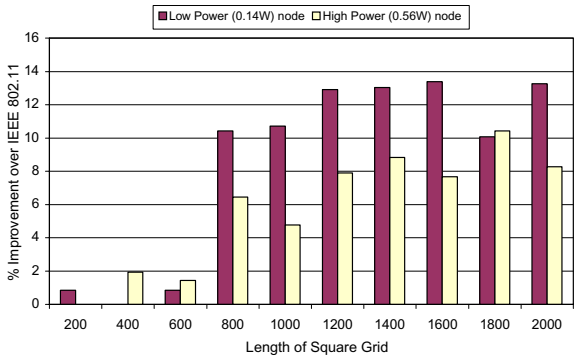

Fig. 9. Increase in average throughput with the MAC layer enhancements.

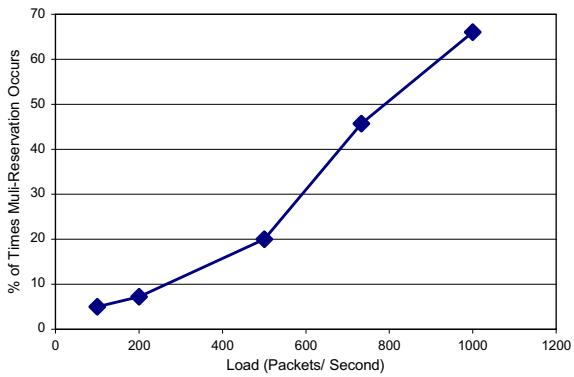

Fig. 10. Frequency of multi-reservations at various traffic load.

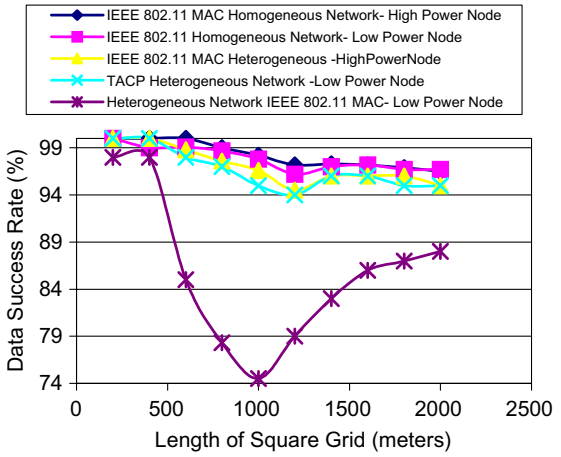

Fig. 11. Data success rate in the heterogeneous network with TACP scheme.

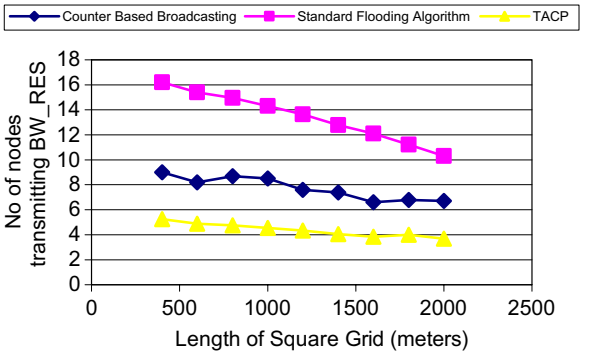

Fig. 12. BW_RES overhead with flooding, intelligent broadcasting and TACP schemes.

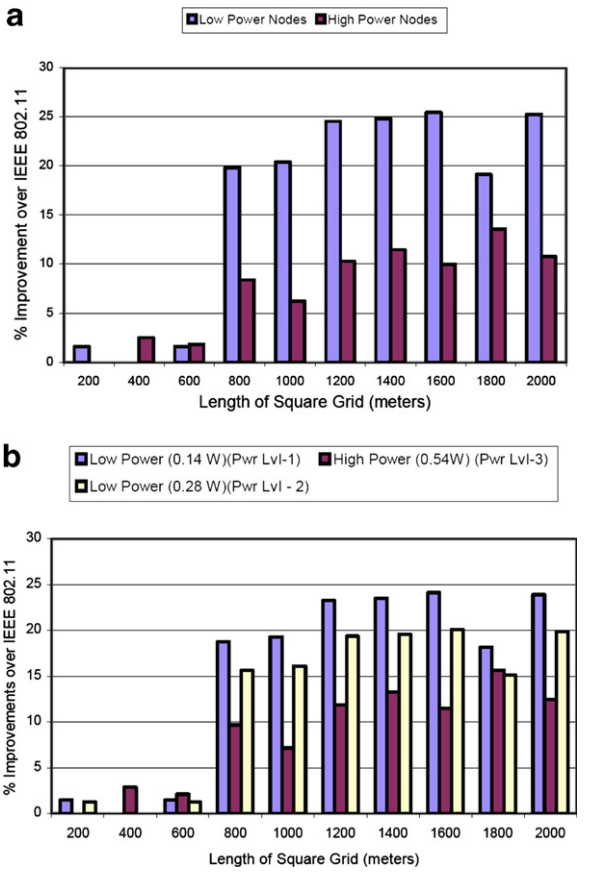

Fig. 13. Increase in average throughput with TACP in power heterogeneous settings. (a) Throughput of nodes in a heterogeneous network. (b) Throughput of nodes in a network with three power levels.

up to 20%) with Case (f) as compared to the scenario with Case (a); their performance is almost as good as that of the high power nodes. In addition, the number of false link failures decreases with TACP by about 28% compared to the IEEE 802.11, and by about 8% compared to when Intelligent Broadcasting is used (Fig. 7). The latter enhancement is due to the additional intelligence at the MAC layer, which significantly reduces the overhead traffic in the network and the contention for channel access. To elucidate this further, we compare the number of BW_RES frames generated per CTS instantiation with the standard flooding scheme, Intelligent Broadcasting and TACP. The number of BW_RES frames broadcast per CTS frame reduces by about 50% with TACP even as compared to that with IB (Fig. 12). In addition, the overall improvement (relative to the IEEE 802.11 MAC protocol) in terms of data success rate and throughput are higher with TACP than with IB despite the overhead incurred due to the *Hello* messages.[11] We observe that with this improvement in data success rate (Fig. 11) and the reduction in interference from high power nodes, the throughput efficiency of the low power nodes increases by up to 24% (Fig. 13a). TACP also alleviates false link failures at high power nodes, this is reflected by the improvement of up to 12% in their throughput efficiency.

Up to this point, we assumed two power levels – high and low – in our simulations. Next, we introduce additional heterogeneity by incorporating multiple possible power levels. The fraction of nodes belonging to each power level is (almost with negligible variation) equal. In Fig. 13b, we depict the throughput efficiency improvement for the nodes

[11] TACP does require the transmission of *Hello* messages unlike the Intelligent Broadcasting scheme.

with the three power levels in the network. We observe that our scheme enables the lower power nodes to have a fair share of channel bandwidth. Overall improvement in network throughput as compared to the IEEE 802.11 MAC protocol is up to 18%. Again, the improvements are more visible as we increase the size of the deployment area, since the asymmetry is emphasized at moderate densities as opposed to extremely high densities as described earlier.

We also quantify the impact on overall MAC throughput, of TACP and IB techniques in isolation from the MR scheme. Specifically, we simulate Case (b), Case (c), Case (e) and Case (f); Fig. 14a compares Case (b) and Case (c), and Fig. 14b compares cases (e) and (f). While the mere deployment of IB does not result in drastic improvements in the overall network throughput, MR scheme by itself offers an improvement of up to 9%. However, merely deploying TACP improves performance as much as 16% and MR provides an additional improvement of up to 8%. In sum, TACP and IB function independent from MR and can be used in conjunction to offer supplementary benefits.

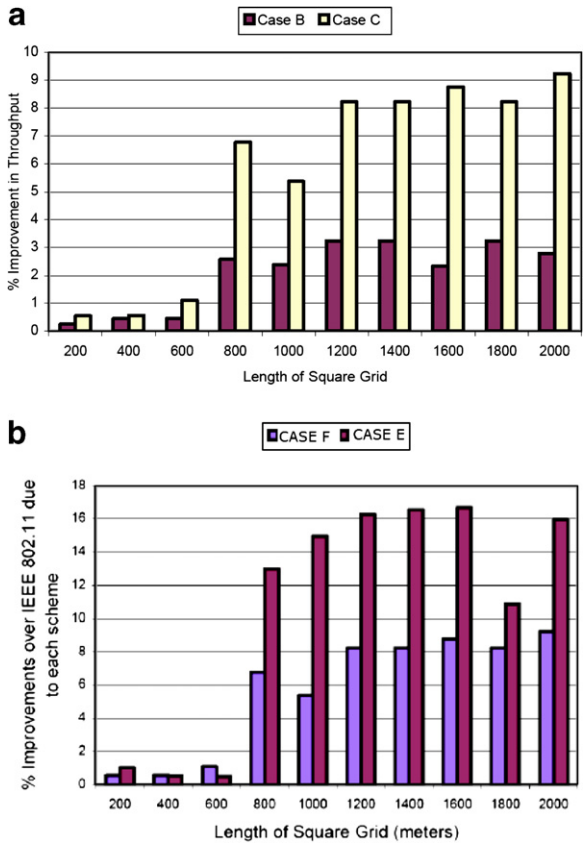

Fig. 14. Throughput improvement using IB and TACP along with multi-reservation scheme, Case (b), Case (c), Case (e) and Case (f). (a) Improvements with the intelligent broadcasting. (b) Improvements with the TACP.

5.2. Performance evaluation at the higher layers

In this second part of our simulations, our objective is to incorporate the routing and transport layers atop our MAC layer, and quantify the performance at higher layers.

We now consider UDP traffic with constant bit rate (CBR), and evaluate the performance of two of the most commonly used on-demand routing protocols AODV and DSR. We assume that each node is represented by its IPv4 address (32 bits), and that the power level of a node is represented by a single bit to indicate whether the relevant node is a high power node or a low power node.[12] The simulation parameters used in these experiments are listed in Table 2.

We are interested in quantifying the performance in terms of the following metrics:

- *Packet delivery ratio (PDR)*: Ratio of the number of higher layer packets delivered, to the number of such packets generated.
- *Average end-to-end delay*: Mean end to end delay experienced by the packets.
- *Route search attempts*: Number of initiated route discovery attempts.
- *Route search failures*: Number of times that a source node fails to find a path to its destination.[13]

Fig. 15 depicts the performance improvement at the routing layer by using our cross-layer framework at different *levels of heterogeneity*, as compared to using the 802.11 MAC protocol. The level of heterogeneity refers to the percentage of low power nodes in the network.[14] We first study

[12] If multiple power levels are to be used, the number of bits that specify the power level need to be increased in a logarithmic proportion.

[13] A route discovery may fail due to two reasons: (i) the network is disconnected and the source and destination belong to different network partitions; (ii) there are unidirectional links on the path from the source to the destination. While in the first case none of the schemes can compute a route, in the second case our framework can discover routes while the traditional methods fail.

[14] This value is maximized when the fraction of low power nodes is almost equal to that of high power nodes in the network [19].

Table 2
Simulation setup for evaluation at the higher layers

Simulator	ns2 (version 2.26)
Channel bit rate	2 Mbps
Radio model	Lucent WaveLAN
Number of nodes	80
Heterogeneity	10–50% low power nodes
Power levels	0.56 W and 0.14 W
Traffic Model	CBR, 5–50 packets/s, 512 bytes/packet
Number of source nodes	Varied between 10 and 15
MAC protocols	IEEE 802.11 MAC DCF IEEE 802.11 w/MAC enhancements IEEE 802.11 w/cross-layer framework
Routing protocols	AODV, DSR, AODV/DSR w/cross-layer framework

the performance in terms of the first metric, PDR. Fig. 15a depicts that our MAC layer enhancements register PDR improvements of up to 12% over the traditional IEEE 802.11 MAC protocol. The PDR further increases with our cross-layer framework with both AODV and DSR, due to the ability to identify and utilize unidirectional links. The improvement is higher with AODV than with DSR at higher levels of heterogenity. This is because AODV, with its traditional settings, does not support asymmetric links. DSR, on the other hand, allows a destination (upon the receipt of a route request) to invoke its own route discovery to discover the source in the presence of unidirectional links. Therefore, AODV has more to gain when deployed over our framework. We remark that DSR benefits from our framework as well, since its *reverse route discovery* floods are no longer needed. These improvements are owing to the overall reduction in network contention with decreasing MAC and routing layer control packet overhead. This is borne out by Fig. 15b, which depicts that the number of route discovery attempts are significantly reduced (by about 35% for AODV and by about 25% for DSR) with our framework. This reduction is a consequence of our framework enabling nodes to easily discover and use unidirectional links. Without our framework these links were either rendered useless (in the case of AODV) or were discovered with a high overhead (with DSR). Our results also show that the percentage of route search failures are reduced by up to 25%; Fig. 15c depicts the performance for both routing protocols under consideration.

Finally, we study the mean end-to-end delay at the routing layer, experienced by packets under scenarios with different levels of heterogeneity. Fig. 16a

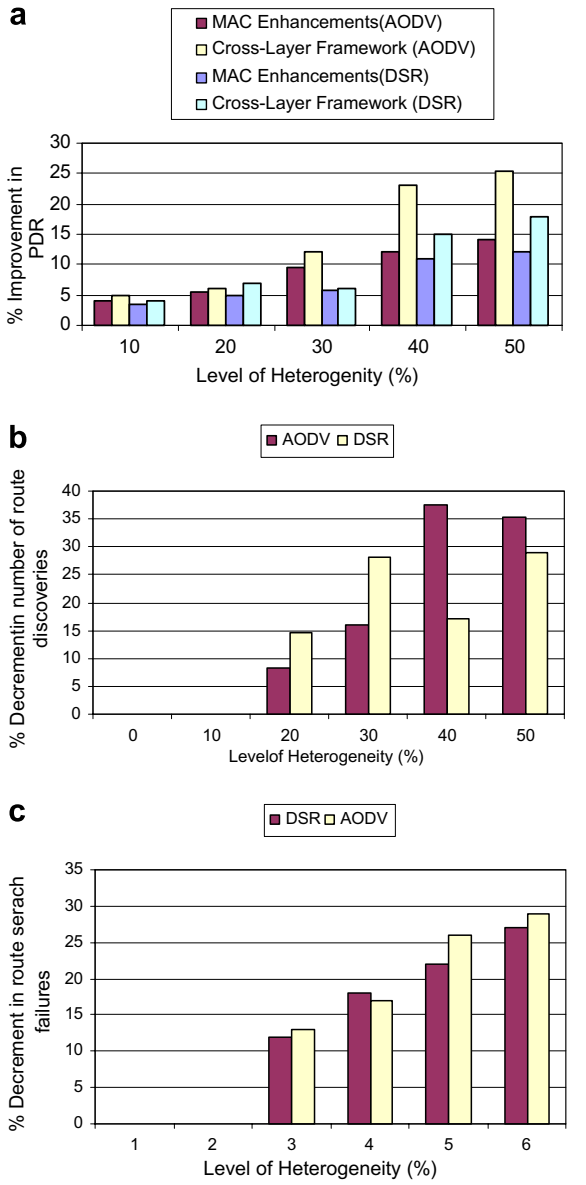

Fig. 15. Performance improvements at the routing layer with our cross-layer framework. (a) Increase in packet delivery ratio (PDR) of nodes in the heterogeneous network. (b) Percentage decrement in route discovery attempts. (c) Percentage decrement in route search failures.

shows that the mean delay experienced is only marginally increased by our cross-layer framework. Conflicting factors contribute towards these results. On the one hand, our framework requires the transmission of additional BW_RES frames at the MAC layer and thus, a MAC layer transmission takes a longer time than with the traditional schemes. On the other hand, this effect is somewhat offset via the use of multi-reservations. With the traditional

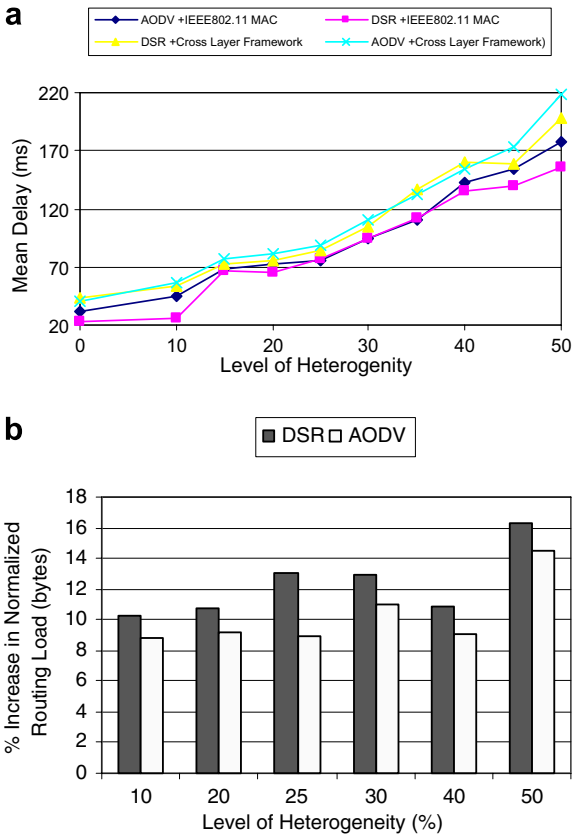

Fig. 16. End-to-end delay and routing overhead with our framework. (a) End-to-end delay. (b) Routing overhead with our cross-layer framework in the heterogeneous network.

schemes, route failures and the consequent route discovery attempts occur with greater frequency, and during these periods data packets simply wait in the source queue. At higher levels of heterogeneity, the overall delays experienced at the MAC layer increases, due to increased collisions and retransmissions caused by asymmetry (Fig. 16b). However we still find that the mean delay with our cross-layer framework is marginally larger than that with traditional protocols. 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.

5.3. Choice of system parameters

In this subsection we study the sensitivity of our simulation results to the system parameters that were used. We have earlier presented intuitive justifications for our selected values of the system parameters; in the following, we provide experimental results that corroborate our previous discussions.

Choice of T: T is a parameter associated with the Intelligent Broadcasting scheme; if a node overhears broadcasts from at least T neighbors, it quells its own. As depicted in Fig. 17c, when T is 1 or 2, nodes are over-aggressive in quelling their broadcasts. As a result, the coverage in terms of reaching all the potential interfering high power nodes is small.

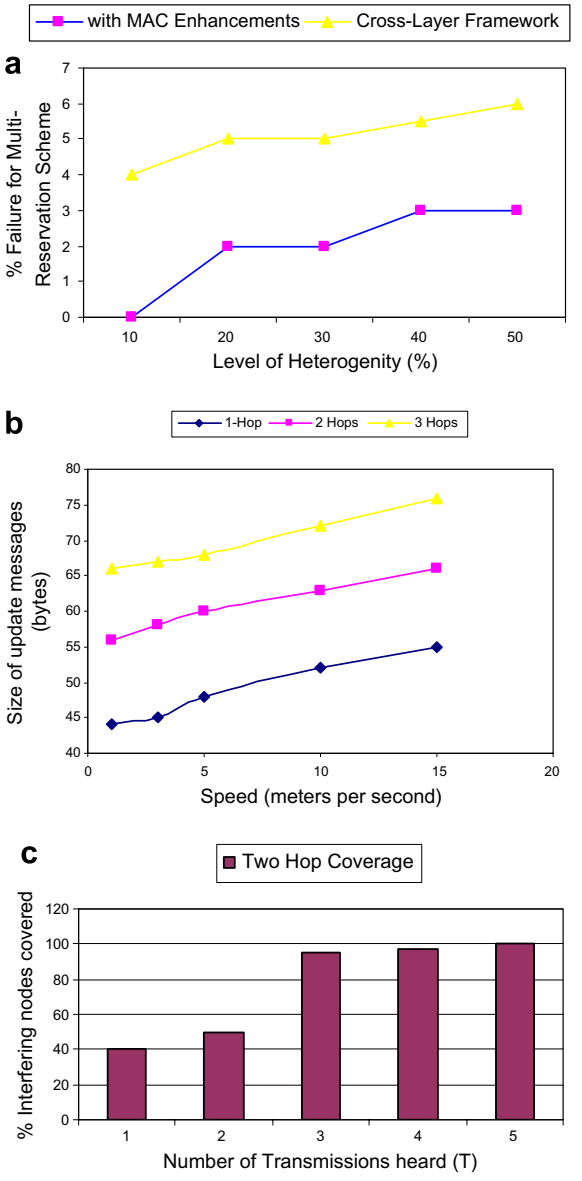

Fig. 17. System parameters with our scheme. (a) Packet failures with the multi-reservation scheme when $N = 3$. (b) Size of update *Hello* messages at different levels of mobility. (c) Choice of parameter T.

For values of T strictly greater than 3, nodes end up performing more re-broadcasts than needed, thus rendering the additional coverage benefits negligible. Thus, we set $T = 3$.

Choice of N: To recap, N is the number of packets (intended for sequential transmissions), for which a node can simultaneously reserve the channel with a single RTS/CTS/BW_RES instantiation using the multi-reservation scheme. In Fig. 17a we plot the percentage of times when a packet fails while using the multi-reservation scheme with various levels of heterogenity, when the value of N is fixed to 3. We plot, (a) the performance using the MAC layer enhancements with our Poisson traffic agent at the MAC layer and (b) the performance of our cross-layer framework. The percentage of failures observed with the latter is higher, due to the additional overhead of routing. Failures occur either due mobility or because our schemes are unable to reach all potential interferer high power nodes (unavailability of paths). The latter effect also increases with the level of heterogeneity (increased asymmetry) as borne out by the figure. Failure rates are nominal with $N = 3$ (between 3% and 6%). At higher values of N we observe that failure rates increase considerably (we do not report results here due to space constraints). This is because a higher number of back-offs[15] occur for larger values of N. In addition, larger values of N also lead to the dominance of the channel by certain nodes thereby increasing the unfairness in channel access. Thus, in our simulations we use a small value of N, specifically, 2 for high power nodes and 3 for low power nodes. We provide a larger value of N for low power nodes in order to improve fairness for these nodes in their channel access procedure.

Choice of n: The choice of n governs the ability of nodes to find reverse routes for bridging a unidirectional link. As mentioned earlier, n refers to the distance in terms of hop count, up to which a node collects and disseminates local neighborhood information. We know from our earlier intuitive discussions that n should be at least 2 in order that a low power node can identify the high power nodes that are beyond its own transmission range. We observe that the size of the n-hop neighborhood is relatively unaffected by node mobility, as we use constant speeds in our model. Fig. 17b depicts that the size of the route update packets increases with speed and with n. Since with higher mobility the localized neighborhood of a node is likely to change more often, the size of these update *Hello* messages increases with mobility. A higher value of n causes an increase in the size of the update messages since nodes have larger localized graphs. We remark that an update message contains only the changes in the localized graph incurred since the previous update message. Since, the average increase in the size of the update messages with an increase in n from 2 to 3 is nominal (increases by about 10 bytes[16]) we use $n = 3$. We find that this value improves the performance, since a large number of unidirectional links are now bridged and a sufficiently high number of interfering high power nodes are prevented from initiating transmissions when a low power communication is in progress.

6. Conclusions and future work

In this paper our key contribution is the development of a unified framework that offers a coupling between the MAC and the routing layers to deal with link level asymmetries in power heterogeneous ad hoc networks. While there had been no prior solutions that handle asymmetry effectively at the MAC layer, previous *unidirectional* routing schemes had ignored the MAC layer dependencies. In our framework the MAC layer solicits assistance from the routing layer to identify link asymmetry. Low power nodes then route MAC layer control frames 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 also allows for the identification and usage of unidirectional links at the routing layer. This in turn leads to shorter routes and consequently to improved performance. We also considered two techniques that are exclusively deployed at the MAC layer based on (a) the use of an intelligent broadcast scheme to quell unnecessary broadcasts, and (b) reserving the bandwidth for multiple data frames with a single RTS/CTS exchange. We find that while these schemes can also provide improvements in performance, our cross-layer approach does significantly better. We study the performance exclusively at the MAC layer and at the higher layers. At the transport layer, our cross-layer framework can improve throughput at

[15] A transmission failure causes a node to back-off.

[16] Note that the IEEE 802.11 data frame is typically of the order of 2 K Bytes.

the low power nodes by up to 25%, alleviating the unfairness introduced with the legacy IEEE 802.11 MAC. We also show a significant reduction (by 20%) in the total number of false link failures caused in the network due to interference from neighboring nodes. As a result of this comprehensive simulation-based analysis, our integrated MAC/Routing layer framework is shown to offer a simple yet viable and effective solution for handling asymmetry in power heterogeneous ad hoc networks.

In this work, we propose the use of traditional on-demand routing protocols (DSR and AODV) atop our integrated framework. While this is a preferable approach for ensuring backward compatibility with possible previous implementations of these protocols, one might also explore future optimizations to network wide routing in power heterogeneous networks. In addition, in this work we have simulated fixed power levels and therefore coarse-grained power heterogeneity. Our schemes need to be examined when nodes deploy fine-tunable power levels, i.e. in the presence of power control. There is a trade-off between exploiting spatial re-use with power control and dealing with the consequent link level asymmetry. We are planning to investigate these aspects in our future efforts.

References

[1] IEEE STD 802.11 1999. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999.

[2] Y. Afek, E. Gafni, Distributed algorithms for unidirectional networks, SIAM Journal of Computing 23 (6) (1994) 1152–1178.

[3] L. Bao, J.J. Garcia-Luna-Aceves, Link-state routing in networks with unidirectional links, in: IEEE International Conference on Computer Communications and Networks (IC3N), 1999.

[4] L. Bao, J.J. Garcia-Luna-Aceves, Channel access scheduling in ad hoc networks with unidirectional links, in: ACM Worskshop on the Discrete Algorithms and Methods for Mobile Computing and Communication (DIAL-M), 2001.

[5] V. Bhargavan, A. Demers, S. Shenker, L. Zhang, Macaw: a media access protocol for wireless lans, in: ACM SIGCOMM, 1994.

[6] D.M. Blough, M. Leoncini, G. Resta, P. Santi, The lit K-neigh protocol for symmetric topology control in ad hoc networks, in: ACM MOBIHOC, 2003.

[7] M. Burkhart, P.v. Rickenbach, R. Wattenhofer, A. Zollinger, Does topology control reduce interference? in: ACM MOBIHOC, 2004.

[8] G. Calinescu, Computing 2-hop neighborhoods in ad hoc wireless networks, in: ADHOC-NOW, 2003.

[9] E. Duros, W. Dabbous, H. Izumiyama, N. Fujii, Y. Zhang, A link-layer tunneling mechanism for unidirectional links, RFC 3077, 2001.

[10] Z.J. Haas, M.R. Pearlman, The zone routing protocol (ZRP) for ad hoc networks, Internet draft-Mobile Ad hoc networking (MANET) Working Group of the Internet Engineering Task Force (IETF), March 2001.

[11] Eun-Sun Jung, Nitin H. Vaidya, A power control mac protocol for ad hoc networks, in: ACM MOBICOM, 2002, pp. 36–47.

[12] D. Kim, C.K. Toh, Y. Choi, GAHA and GAPA: two link-level approaches for supporting link asymmetry in mobile ad hoc networks, IEICE Transaction on Communications E-86B (4) (2003) 1297–1306.

[13] T.J. Kwon, M. Gerla, Clustering with power control, in: IEEE MILCOM, vol. 2, 1997, pp. 919–923.

[14] E.L. Li, J.Y. Halpern, P. Bahl, Y.M. Wang, R. Wattenhofer, Analysis of a cone-based distributed topology control algorithm for wireless multi-hop networks, in: ACM PODC, 2001.

[15] N. Li, J.C. Hou, L. Sha, Design and analysis of an mst-based topology control algorithm, in: IEEE INFOCOM, 2003.

[16] X.Y. Li, W. Song, W. Wang, A unified energy-efficient topology for unicast and broadcast, in: ACM MOBICOM, 2005.

[17] X.Y. Li, P.J. Wan, Y. Wang, O. Frieder, Sparse power efficient topology for wireless networks, in: IEEE HICSS, 2002.

[18] W. Liu, Y. Zhang, W. Lou, Y. Fang, DELAR: Device/energy/load aware relaying in heterogenous wireless ad hoc networks, in: Military Communications Conference (MILCOM), vol. 3, 2004, pp. 1303–1309.

[19] Mahesh K. Marina, Samir Ranjan Das, Routing performance in the presence of unidirectional links in multihop wireless networks, in: ACM MOBIHOC, 2002.

[20] J.P. Monks, V. Bharghavan, W.-M.W. Hwu, A power controlled multiple access protocol for wireless packet networks, in: IEEE INFOCOM, 2001, pp. 219–228.

[21] Alaa Muqattash, Marwan Krunz, Power controlled dual channel (pcdc) medium access protocol for wireless ad hoc networks, in: IEEE INFOCOM, 2003.

[22] Alaa Muqattash, Marwan Krunz, A single-channel solution for transmission power control in wireless ad hoc networks, in: ACM MOBIHOC, 2004, pp. 210–221.

[23] S. Narayanaswamy, V. Kawadia, R.S. Sreenivas, P.R. Kumar, Power control in ad hoc networks: theory, architecture, algorithm and implementation of the compow protocol, in: European Wireless Conference, 2002, pp. 156–162.

[24] S. Nesargi, R. Prakash, A tunneling approach to routing with unidirectional links in mobile ad hoc networks, in: IEEE International Conference on Computer Communications and Networks (IC3N), 2000, pp. 522–527.

[25] S. Ni, Y. Tseng, Y. Chen, J. Sheu, The broadcast storm problem in a mobile ad hoc network, in: ACM MOBICOM, 1999.

[26] N. Poojary, S.V. Krishnamurthy, S. Dao, Medium access control in a network of ad hoc nodes with heterogeneous transmit power capabilities, in: IEEE ICC, 2001.

[27] R. Prakash, A routing algorithm for wireless ad hoc networks with unidirectional links, ACM/Kluwer Wireless Networks 7 (6) (2001) 617–625.

[28] Ram Ramanathan, Regina Hain, Topology control of multihop wireless networks using transmit power adjustment, in: IEEE INFOCOM, 2000, pp. 404–413.
[29] Venugopalan Ramasubramanian, Ranveer Chandra, Daniel Mossé, Providing a bidirectional abstraction for unidirectional adhoc networks, in: IEEE INFOCOM, 2002.
[30] E. Royer, C-K. Toh, A review of current routing protocols for ad hoc wireless networks, IEEE Personal Communications Magazine 5 (1999).
[31] P. Sinha, S.V. Krishnamurthy, S. Dao, Scalable unidirectional routing using zrp extensions for wireless networks, in: IEEE WCNC, 2000.
[32] W. Song, Y. Wang, X. Li, Localized algorithms for energy efficient topology in wireless ad hoc networks, in: ACM MOBIHOC, 2004.
[33] S.L. Wu, Y.C. Tseng, J.P. Sheu, Intelligent medium access for mobile ad hoc networks with busy tones and power control, IEEE Journal on Selected Areas in Communications (JSAC) 18 (9) (2000) 1647–1657.
[34] Jungkeun Yoon, Mingyan Liu, Brian Noble, Random waypoint considered harmful, in: IEEE INFOCOM, 2003.

Vasudev Shah received his B.S. in Electrical Engineering at the University of Bombay, India in 2000, M.Sc. in Computer Science at the University of California, Riverside in 2004 and he is a Ph.D. Candidate at the University of California-Riverside. Currently he is a Chief Technology Officer in First Network Solutions, a Los Angeles based company that specializes in the design and implementation of WAN solutions, Voice and Video over IP solutions for Medium and Large businesses. His research interests are the design of medium access control (MAC) and routing protocols, and power control mechanisms in wireless ad hoc networks.

Ece Gelal received her B.Sc. in Telecommunications from Sabanci University, Turkey and her M.Sc. in Computer Science and Engineering from the University of California, Riverside in 2004 and 2006, respectively. Currently she is a Ph.D. candidate in Computer Science and Engineering in the University of California, Riverside. Her research interests are cross-layer design approaches for mobile ad hoc networks (MANETs), power control, topology control in ad hoc networks, and effective use of smart antennas in ad hoc networks via properly designed MAC and routing protocols.

Srikanth V. Krishnamurthy received his Ph.D. degree in electrical and computer engineering from the University of California at San Diego in 1997. From 1998 to 2000, he was a Research Staff Scientist at the Information Sciences Laboratory, HRL Laboratories, LLC, Malibu, CA. Currently, he is an Associate Professor of Computer Science at the University of California, Riverside. His research interests span CDMA and TDMA technologies, medium access control protocols for satellite and wireless networks, routing and multicasting in wireless networks, power control, the use of smart antennas and security in wireless networks. He has been a PI or a project lead on projects from various DARPA programs including the Fault Tolerant Networks program, the Next Generation Internet program and the Small Unit Operations program. He is the recipient of the NSF CAREER Award from ANI in 2003. He has also co-edited the book ''Ad Hoc Networks: Technologies and Protocols'' published by Springer Verlag in 2005. He has served on the program committees of INFOCOM, MOBIHOC and ICC and is the associate editor-in-chief for ACM MC2R.