WIRELESS COMMUNICATIONS AND MOBILE COMPUTING
Wirel. Commun. Mob. Comput. 2006; **6**:745–760

Understanding the key performance issues with MAC protocols for multi-hop wireless networks

Aravind Iyer[1] and Catherine Rosenberg[2*,†]

[1]*School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, U.S.A.*
[2]*Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, N2L 3G1, Canada*

Summary

Multi-hop wireless networks arise in the context of ad hoc networks, sensor networks, and mesh networks, and their performance depends critically on the underlying medium access control (MAC) protocol. Inspite of the large body of work devoted to MAC protocols and associated problems, the relative importance of these problems is still not well understood. This is because most of the previous work focuses on designing a protocol to solve a particular problem, or on identifying scenarios where a protocol will not work well. In addition, most of the work is also based on simplistic assumptions about the physical wireless medium, like fixed ranges for communication and interference, or concepts like capture threshold where the desired signal strength is compared with interference from a single node at a time, rather than cumulatively.

Our paper seeks to address these issues. We believe it is extremely critical that (i) we develop an understanding of the relative significance of the problems affecting MAC protocols, and that (ii) we use a realistic model for the physical channel for design and performance evaluation. Towards this end, we evaluate the performance of three currently proposed MAC protocols, IEEE 802.11 [1], RI-BTMA [2], and DUCHA [3] under a realistic channel model with additive interference. Since these protocols solve or suffer from different sets of problems, our evaluation provides a differential diagnosis of the severity of these problems. Based on our observations, we propose a simple and robust two channel MAC protocol (entitled 2CM) that is based on IEEE 802.11 augmented with a busy-tone channel. The 2CM protocol (i) mitigates the hidden node problem considerably, (ii) does not waste bandwidth in terms of logical control channels, and (iii) provides a reliable link layer acknowledgment. Through extensive simulations, we show that 2CM offers a consistently high throughput performance while not sacrificing link layer reliability in a variety of scenarios, thereby vindicating our approach. Copyright © 2006 John Wiley & Sons, Ltd.

KEY WORDS: multi-hop wireless networks; medium access control (MAC); busy-tone; additive interference model; performance evaluation

---

*Correspondence to: Catherine Rosenberg, Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, N2L 3G1, Canada.

†E-mail: cath@ece.uwaterloo.ca

Contract/grant sponsor: NSF; contract/grant number: 0087266.
Contract/grant sponsor: NSERC Discovery; contract/grant number: RGPIN42023-37200.
Contract/grant sponsor: Bell University Laboratories (BUL).

1. Introduction

Multi-hop wireless networks represent a key paradigm for the future of wireless networking, as they arise in the context of ad hoc networks, wireless sensor networks, and wireless mesh networks. The performance of these different types of networks depends critically on the design of the underlying medium access control (MAC) protocol. In order to better understand the operation of a typical MAC protocol, it is useful to think of MAC protocols as having two logical components. The first is a collision avoidance algorithm which uses physical carrier-sensing, protocol handshakes, busy tones, and so on, to 'reserve' the channel for the duration of the data transmission.The second is a contention resolution algorithm which uses mechanisms such as persistence and backoff, to 'regulate' the access to the channel.

For instance, the IEEE 802.11 protocol [1], one of the most widely used MAC protocols, uses a collision avoidance scheme called carrier-sensing, multiple access with collision avoidance (CSMA/CA) and a contention resolution scheme called the binary exponential backoff (BEB) algorithm. Collision avoidance and contention resolution work together to reduce data packet collisions. A badly designed collision avoidance mechanism can lead to poor throughput, congestion, and instability at the link layer, which can affect the operation of the other layers of the protocol stack. The contention resolution algorithm affects the fairness or how the wireless channel is shared among the different links, and also the throughput. Our work focuses on collision avoidance.

Research on MAC protocols has been active for a few decades now, and there is a wealth of literature devoted to MAC protocols and the problems associated with collision avoidance. Despite this, the understanding of the relative impact of these problems, is still quite limited. This is because most of the previous work focuses either on designing collision avoidance mechanisms to solve a particular problem, or on identifying scenarios where a particular MAC protocol, usually IEEE 802.11, does not work well. In addition, most of the work proposing new collision avoidance mechanisms, is evaluated based on simplistic assumptions about the physical wireless medium. Some works consider a fixed communication range and a fixed interference range, whereas others consider interference using the concept of a capture threshold.[‡] Although the capture threshold model is quite simplistic, it is still widely used, most notably in the *ns2* simulator [4] which is the most common simulation-based performance evaluation tool.

Our work seeks to address these issues. Our goal is to understand the key performance issues involved in designing MAC protocols for multi-hop wireless networks. We believe that in order to design an efficient and robust collision avoidance scheme, the following two factors are absolutely critical: first, an understanding of the relative significance of the problems affecting MAC protocols; and second, the use of a realistic model for the physical channel (i.e., one based on additive interference), for design and performance evaluation. We approach the problem from this standpoint. We start with a list of problems that affect MAC protocols in multi-hop wireless networks, *viz.,* the hidden node problem, the deaf node problem, the exposed node problem, and the link layer congestion problem. In order to understand their relative impact on performance, we evaluate the performance of three currently proposed MAC protocols, IEEE 802.11 [1], RI-BTMA [2], and DUCHA [3]. Since these protocols solve or suffer from different sets of problems, we expect our evaluation to provide a differential diagnosis of the severity of these problems.

Now, the evaluation of MAC protocols in a multi-hop setting is not a straightforward task. First owing to the lack of strong analytical results, and the lack of programmable hardware that can support busy-tones and control channels, we have to rely on simulations to understand the behavior of MAC protocols in large multi-hop settings. Second, the performance of a MAC protocol depends on how well its parameters such as backoff windows, power levels and so on, are configured. Clearly, when MAC protocols are compared based on their performance in a certain scenario, they are not necessarily compared on a fair ground since different protocols may perform better with different setting of parameters. Hence, we have to study a variety of scenarios, in order to identify the protocol that is the most robust and consistent. Once such a protocol is identified, there would still be much work required to tune it thoroughly. Finally, the importance of using a realistic model for the physical channel, cannot be overemphasized.

[‡] Under the capture threshold model, the signal strength of a packet transmission is compared with interference from a single interfering node at a time (rather than cumulatively as it happens in reality), and the transmission is successful if each signal-to-single-interference ratio is greater than the capture threshold.

We have implemented these protocols in an event-driven simulator which incorporates a realistic model for the physical channel based on additive interference and signal-to-interference-and-noise ratio (SINR) thresholding for correct packet reception. Our simulations bring several things to light. IEEE 802.11, the most widely deployed MAC protocol, is known to suffer from all of the aforementioned problem. RI-BTMA better addresses the hidden node problem by using a busy-tone channel, although it suffers from the link layer congestion problem. DUCHA solves all the above problems (under the capture threshold model), by employing a logical control channel, in addition to a busy-tone channel. However, under an additive model of interference, no protocol perfectly solves all problems.

Following are our main observations, based on an evaluation of these three protocols. All protocols use the BEB contention resolution algorithm used in IEEE 802.11

1. **Physical channel model:** Using a realistic physical channel model is extremely important. A simplistic model like the capture threshold model, can be *misleading* in the following ways:
   (a) Since the busy-tone is just an analog signal, nodes can sense sufficient power to detect a busy-tone without knowing which node (or combination of nodes) is (are) asserting the busy-tone. This can lead to problems if the busy-tone is used as a handshake signal, as in RI-BTMA where it is also used as a clear-to-send (CTS), and DUCHA where it is also used as a negative acknowledgment (NACK). Under the simplistic capture threshold based physical channel model, we *could not even observe* this problem.
   (b) Providing a **link layer acknowledgment** is indispensable, since under an additive interference model, DATA packet collisions are inevitable. However, under the *misleading* capture threshold model, RI-BTMA and DUCHA appear to *completely prevent* all DATA packet collisions, which is clearly not true in reality.
2. **Additional control channel:** The link layer congestion problem, in isolation, does not affect performance significantly, although in combination with the hidden node and associated problems, it can produce a lot of performance degradation. DUCHA uses an additional logical control channel to solve the link layer congestion problem, but this can result in a significant loss of bandwidth often outweighing the benefits.
3. **Busy-tone signal:** Although the busy-tone cannot be recommended to be used as a protocol handshake message, it can still be effectively used as an analog signal to block potential hidden nodes from disturbing an on-going data communication.

We propose a simple but robust two channel MAC protocol (entitled 2CM) which is based on the IEEE 802.11 protocol augmented with a busy-tone channel. The 2CM protocol takes into account the above observations, and successfully addresses the shortcomings of the other protocols. The principle of 2CM is quite simple. It uses an additional busy-tone channel in which nodes assert a busy-tone signal while receiving a DATA packet, so as to prevent hidden nodes from becoming active. Through extensive simulations, we show that 2CM is a very efficient MAC protocol which (i) makes effective use of the busy-tone signal in addition to the request-to-send/clear-to-send (RTS/CTS) handshake of IEEE 802.11, to mitigate the hidden node problem considerably, (ii) does not require a logical control channel, and (iii) provides a reliable link layer acknowledgment mechanism, just as in IEEE 802.11.

In what follows, we start by reviewing some related work in Section 2, followed by a brief discussion of physical channel models and MAC problems in Section 3. Section 4 looks at a few scenarios which provide the motivation for 2CM, which is introduced next, in Section 5. Section 6 provides a thorough performance evaluation of all the protocols, and finally Section 7 concludes the paper. A brief introduction to the protocols, RI-BTMA and DUCHA, is provided in Appendix A.

2. Related Work

There has been a considerable amount of research in the area of collision avoidance mechanisms for wireless networks. The hidden terminal problem was identified as early as 1975 by Tobagi and Kleinrock [5], in the context of infrastructure-based networks. In Reference [5], the authors proposed busy-tone multiple access (BTMA) wherein the central receiver protects itself from collisions due to hidden terminals, by asserting a busy-tone. The busy-tone is a narrowband, analog signal which is detected by measuring the power received in its band. BTMA was generalized to receiver-initiated BTMA (RI-BTMA) [2], by Wu and Li, for use in multi-hop wireless networks. Under the capture threshold model, RI-BTMA is a very effective collision avoidance strategy. It achieves perfect collision avoidance (i.e., it avoids data packet collisions completely), although at the cost of additional radio complexity. However, it does not provide a link layer acknowledgment, in addition to some other problems, as we shall see later.

The idea of using an in-band RTS-CTS handshake, in order to reduce radio complexity, was first proposed in multiple access collision avoidance (MACA) by Karn [6]. MACA was later refined in MACAW [7] by Bharghavan *et al.*, and in IEEE 802.11 [1]. IEEE 802.11 modified the basic RTS/CTS message exchange of MACA to a four-way RTS/CTS/DATA/ACK message exchange, while MACAW changed it to RTS/CTS/DS/DATA/ACK where the data-send (DS) message was used to indicate success/failure of the RTS/CTS dialog. In References [8] and [9], Fullmer and Garcia-Luna-Aceves proposed floor acquisition multiple access (FAMA) and FAMA-NCS (NCS stands for non-persistent carrier sensing). FAMA-NCS also avoids data packet collisions completely, but wastes way too much bandwidth in control message exchanges. IEEE 802.11 also employs virtual carrier sensing, using a network allocation vector (NAV) field in the RTS/CTS messages which is nothing but the time left for the completion of the on-going data transmission. Nodes that can successfully decode these messages are blocked from transmitting for the duration of the NAV.

In Dual Busy-Tone Multiple Access (DBTMA) [10] proposed by Haas and Deng, the idea of RI-BTMA [2] is used and additional protection is offered to RTS packets by means of a transmit busy-tone signal (in addition to the receive busy-tone of RI-BTMA) which is asserted when the RTS is being transmitted. In MAC-SCC [11], the idea of using a separate control channel to transmit control packets when the data channel is busy is proposed. MAC-SCC uses the additional control channel to avoid link layer congestion. In a more recent work entitled DUCHA [3], the authors propose a dual-channel MAC protocol for multi-hop wireless networks. DUCHA uses a separate control channel, and a busy-tone channel, in addition to the data channel. Under the capture threshold model, DUCHA not only achieves perfect collision avoidance, but also solves the link layer congestion problem. However, as we shall see later, DUCHA loses considerable bandwidth in terms of its additional control channel.

Several other works like [12–14] focus on IEEE 802.11, and identify several multi-hop scenarios under which it performs quite poorly. Recent works like PAMAS [15] and AIMRP [16] propose the use of duty-cycling as a means of saving power. This means that the radio module of a wireless node is shut off from time to time when the node is no longer communicating. This can lead to problems since the nodes which are asleep, cannot interpret handshake messages correctly. The problem of arbitrating fairness among the wireless links is considered in References [17,18]. Luo *et al.* in [17], consider a packet scheduling framework, and discuss how fairness between the links can be traded off, to achieve greater spatial reuse. Nandagopal *et al.* in [18], discuss how to design contention resolution algorithms to achieve a given model of fairness among the wireless links. Link layer fairness and related issues, are beyond the scope of this paper.

3. Physical Channel Model and Collision Avoidance Problems

In this section, we briefly describe the various problems that arise in designing collision avoidance strategies for multi-hop wireless networks. We start by introducing the physical channel model.

3.1. Modeling the Physical Channel

Before introducing our model for the physical channel, let us briefly discuss the need for a physical channel model. Modeling communications in a wireline network is very simple, since a node can communicate directly with only those nodes to which it is *wired*. By contrast, in a wireless network, there are three phenomena which need to be captured in order to model communications. First, the wireless medium is shared. In other words, the signal power from a wireless transmitter is radiated into space, rather than being confined within a wire. Second, wireless signals get attenuated considerably over distance. Thus, a node farther from a transmitter will perceive the signal at a lower power, than one closer to the transmitter. Third, a wireless signal can be successfully decoded only if the received signal power is sufficiently higher than the disturbance caused by noise and the *cumulative* interference due to other on-going signal transmissions.

Thus, a realistic model for wireless communications should take into account wireless signal propagation and the additive nature of interference. The choice of a physical channel model is important because our understanding of the properties of a MAC protocol, such as whether or not it successfully solves problems like the hidden node problem, etc., are model-dependent. Hence, the physical channel model cannot be too simplistic.

The physical channel model we use in this paper is based on additive interference, and SINR thresholding for correct packet reception. To be precise, denoting the set of simultaneously active wireless links (or transmitter–receiver pairs) by $\mathcal{L}$, and the transmit power and location of the transmitter of link $l \in \mathcal{L}$, as P_l and X_l respectively, the SINR perceived by the receiver of link m, γ_m, is given by:

$$\gamma_m = \frac{P_m \mathcal{A}(X_m, Y_m)}{\sum_{l \in \mathcal{L}\, l \neq m} P_l \mathcal{A}(X_l, Y_m) + \mathcal{N}_m f} \tag{1}$$

where Y_m denotes the location of the receiver of link m, $\mathcal{A}(x, y)$ denotes the channel attenuation from point x to point y, $\mathcal{N}_m$ denotes the power spectral density of the thermal noise at the receiver of link m, and f denotes the frequency bandwidth of the channel. The capacity of link m, C_m, depends on the modulation and coding scheme used at the physical layer. A packet reception at the data-rate C_m is successful, provided the SINR γ_m remains greater than some SINR threshold θ_m, corresponding to an acceptable bit error rate (BER), throughout the duration of the packet transmission. Otherwise, there is a collision. Also, a node located at Y will sense the channel busy if

$$\sum_{l \in \mathcal{L}} P_l \mathcal{A}(X_l, Y) + \mathcal{N}_m f > \theta_{\text{cs}} \tag{2}$$

where θ_{cs} is the carrier-sensing threshold. We will introduce actual numbers in Sections 4 and 6 when we look at simulation scenarios and results.

3.2. The Problem of Hidden, Exposed, and Deaf Terminals

A *hidden terminal* is defined in the context of a given transmitter–receiver pair. A node is said to be hidden from the transmitter if it can cause a collision (as defined above) at the receiver by transmitting, but cannot perceive any signal sent by the transmitter (according to Equation (2) above). For instance, nodes C and F are hidden nodes, for the transmitter–receiver pair, A and B, in Figure 1. While node A is transmitting to node B, if either node C or node F begins transmitting, there will be a collision at node B. The problem can be mitigated by requiring nodes A and B, to execute a handshake (e.g., RTS-CTS in IEEE 802.11), to reserve the channel, before every data transmission. However, the success of this handshake mechanism is severely hampered by two factors. First, a node such as node F that can cause a collision at node B, is unable to interpret the CTS sent out by node B since it is not close enough to node B to receive the CTS at an acceptable SINR (as defined above). Second, even a node such as node C which can successfully decode the CTS from node B could be *deaf* when node B transmits the CTS.

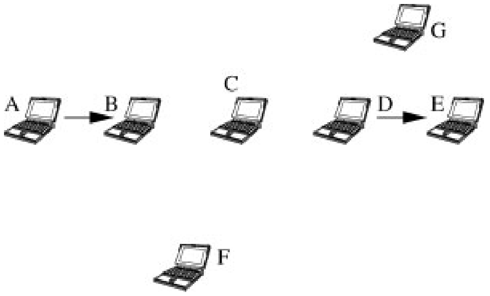

Fig. 1. Illustrating the hidden, exposed, and deaf terminal problems.

A *deaf* terminal is the one that is unable to interpret the handshake messages from a transmitter–receiver pair in its neighborhood. In Figure 1, node C will not be able to interpret the CTS from node B to node A, if there is an on-going data transmission from node D to node E. Thus, the purpose of the CTS from node B is defeated. Indeed, once the transmission from node D to node E is completed, node C could send out an RTS and cause a collision at node B. For a given transmitter–receiver pair, an *exposed terminal* is a node which can perceive signals sent by the transmitter, but cannot cause a collision at the receiver. In Figure 1, nodes C is an exposed node for the transmitter–receiver pair, D and E. Since exposed nodes are also deaf nodes, they cause the same problem. Note that all deaf nodes need not be exposed. Node G in Figure 1 is not exposed, but will be deaf when node D transmits to node E.

3.3. Link Layer Congestion

The RTS-CTS handshake coupled with virtual carrier sensing (using a NAV) suffers from the problem of link layer congestion. Consider Figure 2. When node B is transmitting data to node A, following an RTS-CTS exchange, node C is blocked from transmitting for the duration of the data transmission. Now, if node D sends an RTS to node C, then node C will be unable to reply. This results in two problems. First, although D did not acquire the channel successfully, node E will respect the NAV of node D's RTS, and would be blocked. Hence, an RTS from node F to node E will elicit no response from node E, thereby cascading the process. Thus, node E is *falsely blocked* [7,13] and node F sees *false contention*. Second, node D will interpret the lack

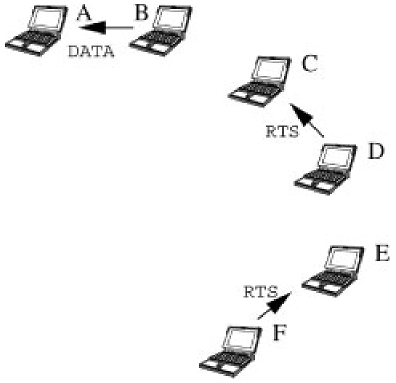

Fig. 2. Illustration of link layer congestion in a network employing IEEE 802.11 MAC.

of a reply from node C, as contention and will enter contention resolution. Now, in case of the BEB algorithm used in IEEE 802.11, this will result in node D (and later node F, etc.,) doubling its backoff window unnecessarily. If this happens often enough, nodes D, F, etc., will retry an RTS transmission several times, and eventually discard their packets. When used in conjunction with an on-demand routing protocol like DSR [19] or AODV [20], this false contention could trigger a *link failure*. The reason for this is that a non-response to an RTS is interpreted as an RTS collision. In other words, the protocol has no means of distinguishing between (i) an RTS collision which indicates contention and (ii) a receiver being blocked.

These problems often lead to significant performance degradation, most notably in IEEE 802.11. In order to tackle these problems, there have been numerous proposals in the literature, like RI-BTMA [2], DBTMA [10], MAC-SCC [11], and DUCHA [3]. In the next section, we try to understand the relative impact of these problems, by evaluating some of these proposed protocols.

4. Preliminary Results on Simple Topologies

In this section, we study the impact of the problems with collision avoidance, identified in the previous section, as well as the impact of an additive interference-based physical channel model. Our methodology for doing so, is the following. We simulate three currently used or proposed MAC protocols viz., IEEE 802.11 [1], RI-BTMA [2], and DUCHA [3], in three different scenarios. We assume that the reader is familiar with the protocol handshakes of IEEE 802.11. The reader is directed to Appendix A for a quick introduction to RI-BTMA and DUCHA. The reasons for choosing these particular protocols is the following. As mentioned earlier in the Introduction, IEEE 802.11 is known to suffer from all the problems identified in Section 3. RI-BTMA, by making use of an additional busy-tone channel, is able to better address the hidden terminal and associated problems, but it provides no link layer acknowledgment mechanism. DUCHA addresses both the hidden terminal problem and the link layer congestion problem, by making use of a busy-tone channel and an additional control channel. However, in order to understand both the relative impact of these problems on performance, and the tradeoffs involved in the design choices made by the aforementioned protocols, it is essential to objectively compare these protocols with one another.

We start by providing some details about the simulation parameters. The physical layer parameters used for the simulations presented here, are summarized in Table I. For the protocol DUCHA which uses a separate control channel for RTS-CTS, and a data channel for DATA packets, the overall bandwidth needs to be split between the two channels. In Table I, we show the proportion of bandwidth used for the control channel. The rest is used for the data channel[§]. According to the IEEE 802.11 standard [1], the control packets RTS and CTS, are always transmitted at the base data-rate of 1 Mbps, while the DATA and ACK packets can be transmitted at higher data-rates. In simulating the other protocols, we adopt the same rule, using only the base modulation and coding scheme to transmit the RTS and CTS packets. Thus, in the case of DUCHA, in a network operating at the 2 Mbps data-rate, the control channel data-rate is 300 kbps (30% of the base data-rate, 1 Mbps), while the data channel data-rate is 1.4 Mbps (70% of the high data-rate, 2 Mbps). The busy-tone channel bandwidth is taken to be 11 kHz. The transmit power and noise power levels are scaled in proportion to the bandwidth of the channel, and the carrier-sensing or detection threshold is set to 6 dB over the noise power in the channel. We do not simulate any rate adaptation algorithm, but we merely operate the entire network at some fixed data-rate (1 Mbps, 2 Mbps, or 11 Mbps). Finally, we do not simulate any congestion control mechanism, instead

---

[§] In the paper on DUCHA [3], the authors use about 22% of the bandwidth for the control channel for a network operating at 1 Mbps. We chose the values in Table I as they resulted in a better performance for DUCHA, in the scenarios we considered.

Table I. Physical channel parameters.

Overall channel bandwidth	22 MHz
Overall transmit power	0 dBm (1 mW)
Signal attenuation (in dB) as a function of distance	−40 dB − 40 log (distance)
Overall noise power	−100 dBm
Carrier-sensing threshold	6 dB over noise power
SINR threshold for 1 Mbps	12 dB
SINR threshold for 2 Mbps	15 dB
SINR threshold for 11 Mbps	24 dB
Busy-tone channel bandwidth	11 kHz
Control channel bandwidth in DUCHA (1–2 Mbps)	6.6 Mhz (30%)
Control channel bandwidth in DUCHA (11 Mbps)	11 Mhz (50%)

allowing the data sources to generate packets at any fixed rate.

Before we present the results of our simulation study, let us briefly discuss what should be the desirable behavior of a good MAC protocol, and introduce the metrics we use for performance evaluation. Figure 3 shows the qualitative throughput performance (throughput vs. arrival rate or load) of a good and a bad MAC protocol. A good MAC protocol achieves a throughput that increases with load up to a certain point, and then remains constant with increasing load. The higher the point at which the throughput saturates, the better. A bad MAC protocol, on the other hand, may become congested even for moderately high arrival rates, and its throughput may actually decrease with increasing load. Note that the behavior of the MAC protocol to the right of the point of maximum throughput is of secondary importance because arguably, a congestion control mechanism at a higher layer, will try to maintain the network to the left of this point. The first metric we use for comparison, is the average end-to-end throughput defined as the average amount of the data payload bits that are delivered from the source to the destination (end-to-end) in unit time.

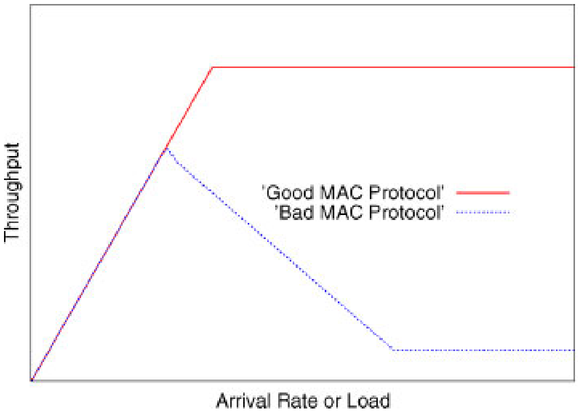

Fig. 3. Qualitative throughput of a MAC protocol.

However, as we shall see later, it is possible for a MAC protocol, to have a good throughput performance, but discard an unduly high number of packets at the link layer close to saturation. Thus, in addition to having a good throughput performance, a good MAC protocol also needs to be *reliable*. In order to capture the reliability of MAC protocols, we use a second metric for comparison, *viz.*, the average link layer packet loss defined as the average amount of the data payload bits that are dropped before reaching the destination. Note that we do simulate all the protocols with the same retry mechanisms as in IEEE 802.11, so that the link layer packet loss is the amount of packets that are dropped after the maximum number of retries. In the case of RI-BTMA, since there is no link layer acknowledgment, a packet can be retried if the RTS collides, but it is lost if the DATA packet collides. To summarize, a good MAC protocol achieves a throughput performance like that shown in Figure 3 while maintaining a low rate of link layer packet losses.

4.1. Scenario 1

We now move on to our simulation study. Our first scenario is based on a simple topology (see Figure 4) to illustrate the potential impact of hidden nodes, and link layer congestion. The distances between the nodes are as indicated. We operate the network at a data-rate of 2 Mbps. Based on the parameters chosen (see Table I), node 1 is a hidden node for the transmitter–receiver pair of node 4 and node 3, and node 4 is a hidden node for the pair of node 1 and node 2. We consider two one-hop flows, labeled flow 1 and flow 2, as shown. The simulation experiment consists of simulating the three MAC protocols for this topology, with Poisson packet arrivals with equal rates for both flow 1 and flow 2. Each simulation run is 1000 simulation seconds long. The total packet arrival rate is varied from 0 to about 2500 kbps. Each protocol is simulated with the BEB contention resolution algorithm with a minimum backoff window

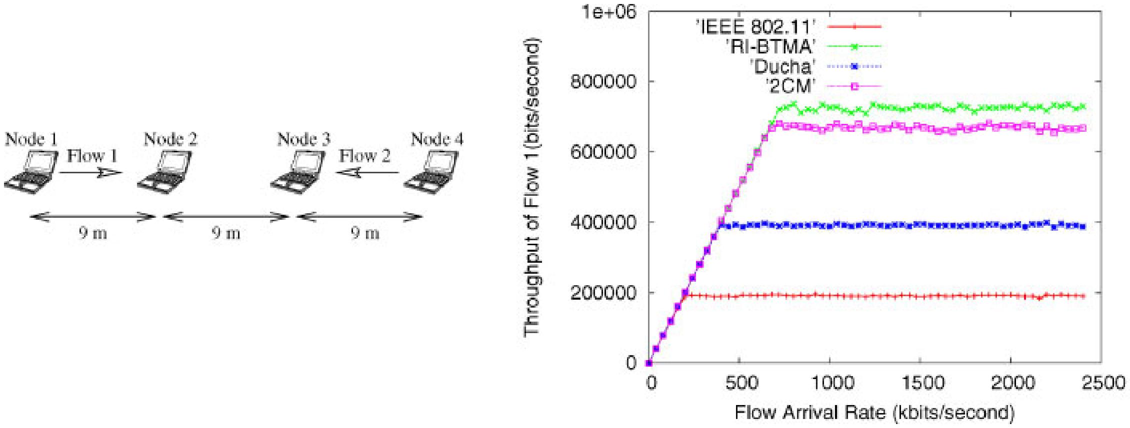

Fig. 4. Scenario 1: Topology (left); throughput of flow 1 for various protocols (right).

of 32 slots and a maximum of 1024 slots. The throughput achieved by flow 1 (the throughput for flow 2 is almost identical), as a function of its arrival rate, for different protocols is shown in Figure 4. Figure 4 also shows the throughput for the 2CM protocol which will be introduced in the next section.

As seen in Figure 4, IEEE 802.11 shows a very poor throughput performance, achieving a total throughput of less than 30% of that achieved by RI-BTMA. The reason for this can be explained by considering the following sequence of events. Consider Figure 4 and let flow 1 be active following an RTS/CTS dialog. Thus, node 2 is receiving a DATA packet from node 1. Now, node 4 is a hidden node for flow 1, and further it cannot decode any CTS messages from node 2, and could cause a collision at node 2, by sending an RTS request to node 3. Not only does this cause a collision at node 2, but since node 3 is exposed to the transmission of node 1 which is active, it cannot interpret the RTS message. This leads node 4 to reattempt multiple RTS transmissions, till the data transmission from node 1 is concluded. This is the link layer congestion problem explained in Section 3. RI-BTMA and DUCHA are unaffected since the busy-tones successfully prevent hidden nodes from transmitting, and hence exhibit a good throughput behavior (refer to Figure 3). DUCHA achieves only about 60% of the throughput of RI-BTMA, because of the bandwidth penalty involved in using an additional control channel. We have not shown the link layer packet loss experienced by the protocols, since in this scenario, the packet losses are insignificant.

4.2. Scenario 2

Now let us consider a similar but more interesting scenario to illustrate the potential impact of the link layer congestion problem (Figure 5). The flows are named as indicated. Packets arrive at the source of these flows according to a Poisson distribution. The

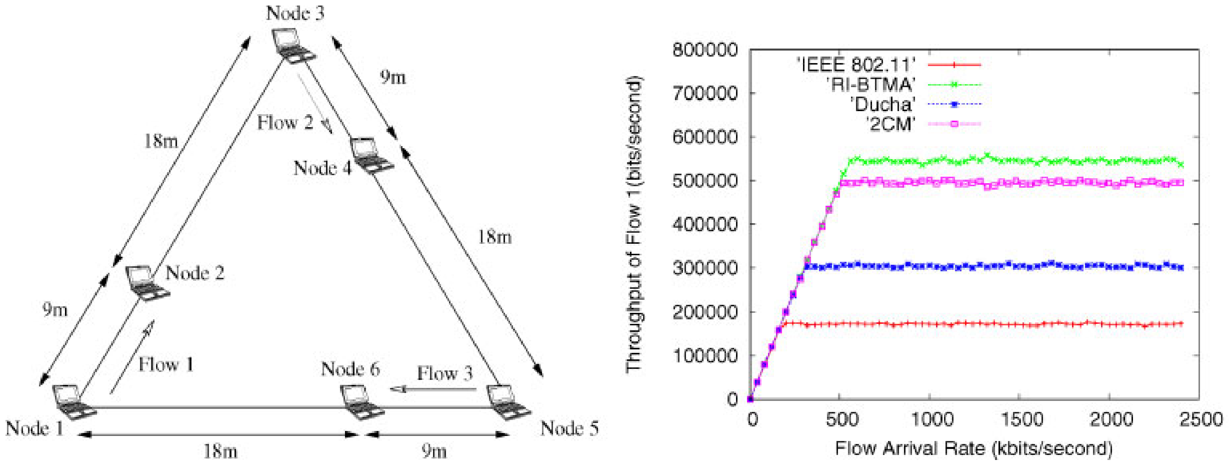

Fig. 5. Scenario 2 Topology (left); throughput of flow 1 for various protocols (right).

network is operated at 2 Mbps. Again, the arrival rates of the three flows are taken to be equal and varied from 0 to 2500 kbps. Each protocol is simulated at a given arrival rate, for 1000 simulation seconds, and with the BEB algorithm as earlier. The throughput curves from this experiment are depicted in Figure 5. Again, we show only the throughput achieved by flow 1, since the throughput achieved by the other two flows is almost identical. Note the similarity of Figure 5 to Figure 4. The explanation for the poor performance of IEEE 802.11 is the same as discussed earlier.

In the case of RI-BTMA, despite its good throughput performance, our simulations reveal the following problem. Consider RI-BTMA, and consider flow 1 to be active. Thus, node 1 is transmitting a data message, and node 2 is asserting a busy-tone signal. This busy-tone is perceived by node 3, but cannot be perceived by node 5. Hence, node 5 could initiate a communication by sending a request packet to node 6. However, since node 6 is exposed to the data transmission from node 1, it cannot receive the request packet correctly. This leads node 5 to reattempt the request multiple times, till the conclusion of the data transmission of flow 1, frequently leading to packet discards. Thus, in this scenario, RI-BTMA suffers from the link layer congestion problem. Although, we do observe this phenomenon repeatedly in the simulation traces, it still does not bring about any perceptible degradation in throughput. Again, we have not shown the link layer packet loss experienced by the protocols, since the packet losses, in this scenario, are insignificant.

In IEEE 802.11, the problems of link layer congestion and hidden nodes, compound each other and bring about severe throughput degradation, as seen in both the above scenarios. However, as seen above in scenario 2, the link layer congestion problem may have a limited impact, in a protocol which mitigates the hidden node problem considerably (in this case, RI-BTMA).

4.3. Scenario 3: Impact of Physical Channel Model

Now we consider a scenario to illustrate the impact of the physical channel model (Figure 6). This is again a three flow scenario, with distances, node labels, and flow labels, as indicated. Again, the network is operated at 2 Mbps. We carry out exactly the same simulation experiment as in the previous two scenarios, and observe the throughput achieved by the three flows. Based on the combination of distances chosen, the following facts will help the reader understand the scenario better. Node 2 is completely immune to collisions due to its close proximity to node 1. Node 3 can receive an RTS packet from node 4, but not a DATA packet when node 1 is transmitting. This is because a DATA packet transmitted at 2 Mbps needs a higher SINR for correct reception than an RTS packet transmitted at 1 Mbps. When both node 1 and node 6 transmit, node 3 cannot receive any transmission from node 4. Node 5 can receive an RTS packet from node 6, in all circumstances, but not a DATA packet if node 4 is transmitting, again because a DATA packet requires a greater SINR for reception. In case of busy-tone-based protocols, node 4 cannot perceive the busy-tone asserted by node 2 or node 5 individually, but it can perceive the busy-tone if both node 2 and node 5 assert it simultaneously.

To illustrate the impact of the physical channel model, we perform the same simulation experiment with a capture threshold model at the physical layer. We use the same parameters as in Table I, except for the following. Note that the capture threshold model uses a signal-to-noise threshold (commonly called `RxThresh`) and a capture or signal-to-interference threshold (commonly called `CpThresh`), instead of an SINR threshold. We take `RxThresh` and `CpThresh` in the capture threshold model, to be equal to the SINR threshold in the additive interference model. To be precise, if the SINR threshold is 15 dB (say), then the received signal power has to be greater than the noise by 15 dB, and greater than the interference from each *individual* node by 15 dB, to be correctly received under the capture threshold model, whereas it has to be greater than the *sum* of the noise and all the interference by 15 dB, to be correctly received under our additive interference model. One would expect the capture threshold model to overestimate the achievable throughput. Indeed, the total throughput achieved by the different protocols, as a function of the total arrival rate is shown in Figure 7. As can immediately be noted, the capture

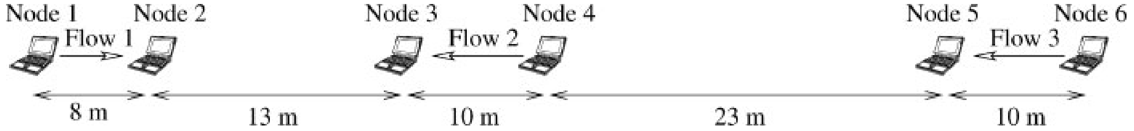

Fig. 6. Scenario 3.

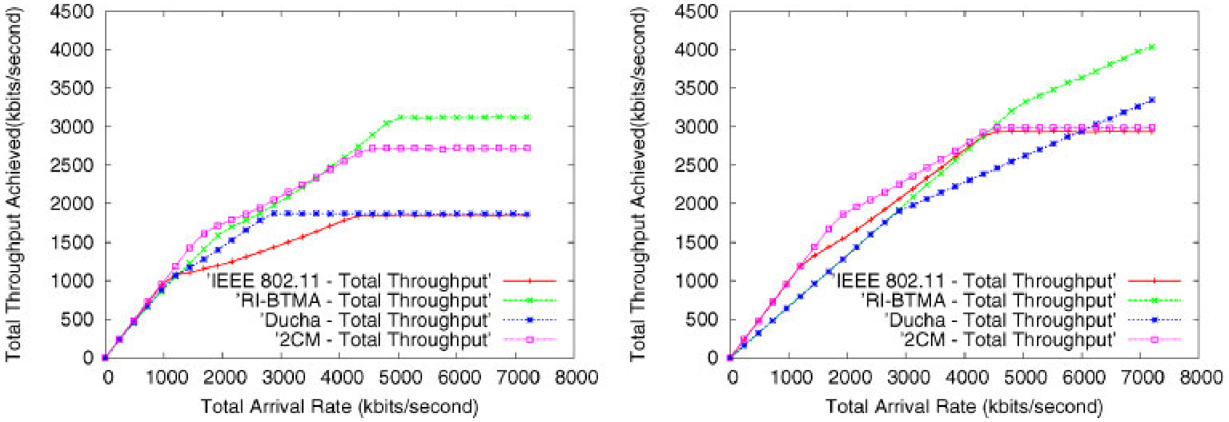

Fig. 7. Scenario 3: total throughput achieved; additive model (left); capture threshold model (right).

threshold model not only overestimates the throughput performance of all the protocols, but, even worse, it also gives different qualitative results.

There are several other noteworthy points to be explained here. First, note the penalty in bandwidth paid by DUCHA, to employ the additional control channel. The total throughput it achieves is only about as high as that of IEEE 802.11, under both the physical layer models, except under very high load with the capture threshold model. This is consistent with what can also be observed in Figures 4 and 5, where again DUCHA achieves a significantly lower throughput than RI-BTMA and 2CM. We also observed that DATA packet collisions are *inevitable* under an additive model of interference, whereas they are *completely prevented* by RI-BTMA, DUCHA, and 2CM, under the capture threshold model. An important consequence of this is that since RI-BTMA does not have a link layer acknowledgment mechanism, the collided packets are *directly lost* at the link layer, under the additive interference model. Although the proportion of DATA packet collisions was not found to be significant, in RI-BTMA, DUCHA, and 2CM, in these simple scenarios, we expect it to aggravate in larger scenarios (as we shall see later, in Section 6).

Finally, the problem of DATA packet collisions, is further compounded in RI-BTMA due to the following problem which can *only be observed* under an additive model of interference. Refer to Figure 6, and consider flow 1 to be active. Although node 2 is asserting a busy-tone, it cannot be perceived at node 4. Now suppose node 4 and node 6 transmit an RTS to node 3 and node 5, respectively, at about the same time. Since node 3 is exposed to the transmission from node 1, it does not receive the RTS and makes no reply. However, node 5 receives the RTS successfully and asserts a busy-tone. Due to the collective power of the busy-tones from node 2 and node 5, node 4 begins perceiving a busy-tone (analogous to a CTS in RI-BTMA) and begins transmitting the DATA packet. This results in a collision at node 3 and also at node 5, and both the DATA packets are *lost immediately* without even being retried, since RI-BTMA has no link layer acknowledgement mechanism. Note that we have not shown the link layer packet loss experienced by the protocols, since the packet losses are not significant, in this scenario.

So far, we have looked at the performance of IEEE 802.11, RI-BTMA, and DUCHA in three scenarios, and made certain observations which provide the motivation for our two channel MAC protocol (2CM). Specifically, the link layer congestion problem is not significant enough to warrant the use of an additional control channel to solve it. Second, since busy-tones are analog signals, it is quite risky to use them as protocol handshake messages, although they can still be effectively used to silence potential hidden terminals. Finally, DATA packet collisions are inevitable. For quick retransmission of the collided packets, a link layer acknowledgment mechanism is a must. This is in contrast with the capture threshold model, where RI-BTMA, DUCHA, and 2CM completely prevent DATA packet collisions. Even though DATA packet collisions did not have a statistically significant impact, in these simple scenarios, we expect them to be important in larger network settings, as we shall see later in Section 6.

5. 2CM: A Two Channel MAC Protocol

In this section, we provide a detailed description of our protocol, entitled 2CM which takes into account all the observations made at the end of the previous section. Although the 2CM protocol is actually nothing but IEEE 802.11 augmented with a busy-tone channel, we provide a detailed description for completeness. To deploy 2CM, we need wireless nodes equipped with two independent half-duplex radio channels. In particular, a node can simultaneously use both the channels, but on each one of them, it cannot transmit and receive at the same time. One of these two channels is a logical channel (a 'bit-pipe'), while the second one is an analog narrowband busy-tone channel. 2CM uses a combination of virtual carrier sensing and busy-tone based reservation to protect DATA packets from hidden nodes.

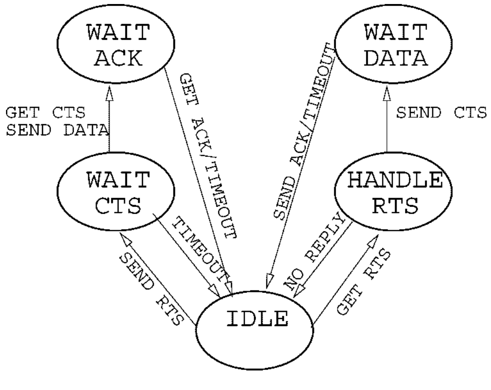

Fig. 8. The finite state machine of the 2CM protocol.

5.1. Protocol Messages

2CM uses four basic messages (like IEEE 802.11), *viz.*, RTS, CTS, DATA, and ACK, in addition to an analog busy-tone signal which will be referred to as BT. These are described below:

1. Request-to-Send (RTS): The RTS message is used for initiating a data packet transmission. It contains the MAC addresses of the source and the destination nodes, and a NAV field, exactly as in IEEE 802.11
2. Clear-to-Send (CTS): The CTS message is used by the receiver node to indicate that it can receive the data packet transmission, and it also contains the NAV field.
3. Busy-tone Signal (BT): The busy-tone signal is used by the receiver node to reserve the data channel and protect itself from potential data transmissions by hidden nodes which could collide with the current data transmission. The busy-tone is asserted at the same time as the receiver node starts receiving a DATA message.
4. Data (DATA): The data message contains the actual data and includes all the higher layer protocol headers and payload.
5. Acknowledgment (ACK): The ACK message is used by the receiver to confirm the correct reception of the data packet. It is used as a mechanism to trigger fast retransmissions of collided DATA packets.

5.2. Collision Avoidance Scheme

A simplified finite state machine for 2CM is shown in Figure 8. A node following the 2CM protocol can be in any of the states shown in Figure 8. The actions producing the state transitions are indicated on the arrows that represent them. The collision avoidance scheme described below specifies these state transitions in detail.

A node following the 2CM protocol is by default in the `IDLE` state. It maintains its own copy of the NAV for virtual carrier sensing. The NAV keeps track of whether the current node is potentially hidden/exposed with respect to an on-going communication. Whenever a node receives a protocol message meant for another node, it infers that it could possibly be a hidden or exposed node for that communication. The node also copies the value of the NAV into its own NAV and remains silent till the NAV expires. The behavior of the node in all the protocol states is described below:

`IDLE` **State:** Upon receiving data from the higher layers of the protocol stack, the node attempts to get access to the channel by following the contention resolution algorithm. In order to initiate a transmission, the node needs to have a zero value of the NAV, in addition to sensing no activity on the busy-tone channel. Then, it needs to sense the main channel idle for DIFS long (as in IEEE 802.11) before accessing the channel through an RTS transmission. On transmitting an RTS message, the node is now in the `WAIT CTS` state (Figure 8). In the `IDLE` state, if a node receives an RTS message intended for itself, then it enters the `HANDLE RTS` state.

`WAIT CTS` **State:** If the node receives a CTS reply, then it proceeds to transmit the DATA packet, going into the `WAIT ACK` state. If the node receives no reply for more than SIFS long (as in IEEE 802.11), or if there is a collision while receiving the CTS reply, then the state times out to `IDLE`, where the node enters contention resolution again.

`WAIT ACK` **State:** If the node receives an ACK reply, then it returns to the `IDLE` state and dequeues the current packet. If the node receives no reply for more than SIFS long, or if there is a collision while

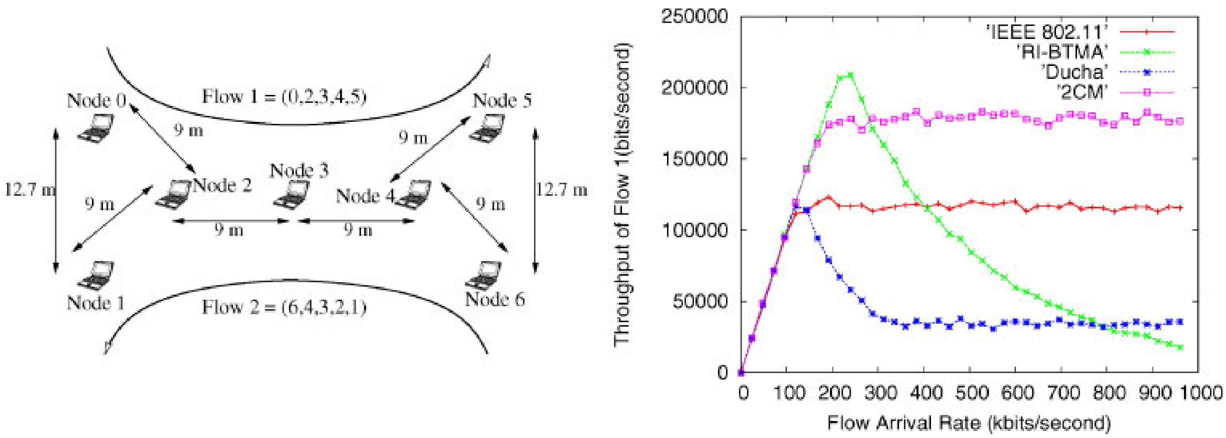

Fig. 9. Topology (left) and throughput of flow 1 for various protocols (right).

receiving the ACK reply, the state times out to `IDLE`, and it attempts to retransmit the packet according to the contention resolution algorithm.

`HANDLE RTS` **State:** If after SIFS long, the node has a non-zero NAV, or perceives some activity in the busy-tone or main channel, it makes no reply, and returns to the `IDLE` state. Otherwise, it replies with a CTS message, and enters the `WAIT DATA` state.

`WAIT DATA` **State:** If the node begins receiving the DATA packet before SIFS long, it asserts the busy-tone until the completion of the DATA packet. On correct reception, the node sends an ACK message and passes the packet to higher layers. If there are any errors or a collision, or if the node receives no reply for more than SIFS long, then it merely returns to the `IDLE` state.

This specifies the collision avoidance scheme used in 2CM.

As already substantiated by the results in Figures 4, 5, and 7, 2CM is a very efficient protocol. In the next section, we carry out a more thorough performance comparison of 2CM with IEEE 802.11, RI-BTMA, and DUCHA, in a variety of realistic multi-hop scenarios, to better illustrate its advantages.

6. Performance Evaluation

By means of Sections 4 and 5, we have understood the relative impact of the problems affecting collision avoidance, and recognized the importance of using an accurate physical channel model. We have also motivated and introduced the 2CM protocol, which performs very well in the scenarios discussed in Section 4. Now, we look at larger and more complicated multi-hop network scenarios, to thoroughly evaluate all the MAC protocols under consideration. Again the simulations are carried out under the additive model introduced in Section 3, with the parameter values from Table I in Section 4.

6.1. Cost of an Additional Control Channel

This scenario illustrates the potential cost of using an additional control channel, as in DUCHA. Consider the simulation topology shown in Figure 9. As indicated, we consider two multi-hop flows. Flow 1 is a flow from node 0 to node 5, and is routed via nodes 2, 3, and 4, in that order, while flow 2 is from node 6 to node 1, via nodes 4, 3, and 2, in that order. The network is operated at 2 Mbps with the BEB algorithm, and as earlier, the simulation experiment consists of simulating each MAC protocol, with Poisson packet arrivals with equal rates for both flow 1 and flow 2. The flow arrival rate is varied from 0 to about 1000 kbps, and for each given arrival rate, the network is operated for 1000 simulation seconds.

The throughput curves for the four protocols with the BEB algorithm are plotted in Figure 9. The first point to note is that surprisingly, IEEE 802.11 achieves a better performance than DUCHA. Although the throughput curves of RI-BTMA and DUCHA are of similar nature, DUCHA suffers a huge throughput penalty owing to the use of its additional control channel. The throughput, in the case of RI-BTMA and DUCHA rises smoothly with the arrival rate until a certain point, beyond which the throughput rapidly decreases. The reason for this 'bad' behavior (as in Figure 3) is the following. Consider the neighborhood of node 2 in Figure 9. We find that the link from node 0 to node 2, achieves greater throughput at the expense of the other neighboring links (node 2 to node 1, node 2 to node

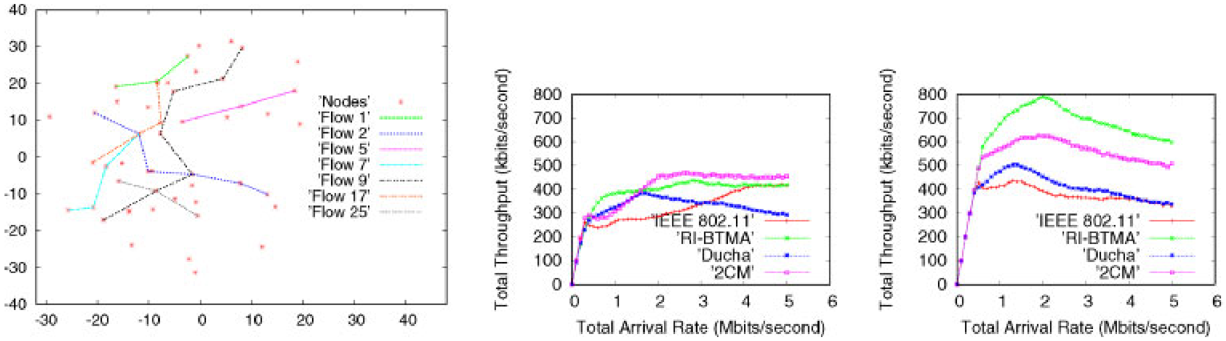

Fig. 10. Topology and a few flows (left); total throughput—additive model (center); total throughput—capture threshold model (right).

3, etc.), thereby driving the overall flow throughput down. This is an undesirable consequence of the BEB contention resolution algorithm. In the case of IEEE 802.11 and 2CM, the throughput rises up to a certain point and then stays constant with increasing load. Both IEEE 802.11 and 2CM are able to avoid any problems due to the unfairness of BEB, in this scenario. As can be seen, 2CM combines the handshake mechanism of IEEE 802.11 with the busy-tone employed by the other two protocols, to achieve the best of both approaches.

6.2. Impact of Physical Layer Model

Next we consider a 50 node random topology shown in Figure 10, to further illustrate the impact of the physical layer model. In Section 4, we used a simple three flow scenario to show that (i) DATA packet collisions are inevitable, and (ii) protocol handshakes based on the busy-tone are prone to misinterpretation. However, we noted that the proportion of DATA packets lost consequently, was not significant. In this scenario, these problems come to light more emphatically. In the node topology shown in Figure 10, end-to-end flows are formed by connecting 25 of the 50 nodes to the other 25, randomly. In order to route the packets from the source to the destination of each flow, we employ a static shortest path routing algorithm. We have not implemented a routing protocol, but we simply store the relevant route information in the nodes, at the beginning of the simulation. Figure 10 indicates a few of the 25 flows. Again, we operate the network at 2 Mbps. All protocols use the BEB contention resolution algorithm. Each flow receives Poisson packet arrivals with equal rates. The total arrival rate is varied from 0 to about 5 Mbps, and the network is operated for 1000 simulation seconds for a given arrival rate.

The throughput curves for the four protocols are plotted in Figure 10. Compare the throughput curves in the center for the additive model, with the ones on the right for the capture threshold model. It can immediately be noted that in this scenario, not only is the throughput overestimated considerably by the capture threshold model, but also there are qualitative differences too. In particular, RI-BTMA achieves almost twice the throughput with the capture threshold model, as compared with the additive model. Another important aspect of the differences in the physical channel models, is illustrated in Figure 11 which shows

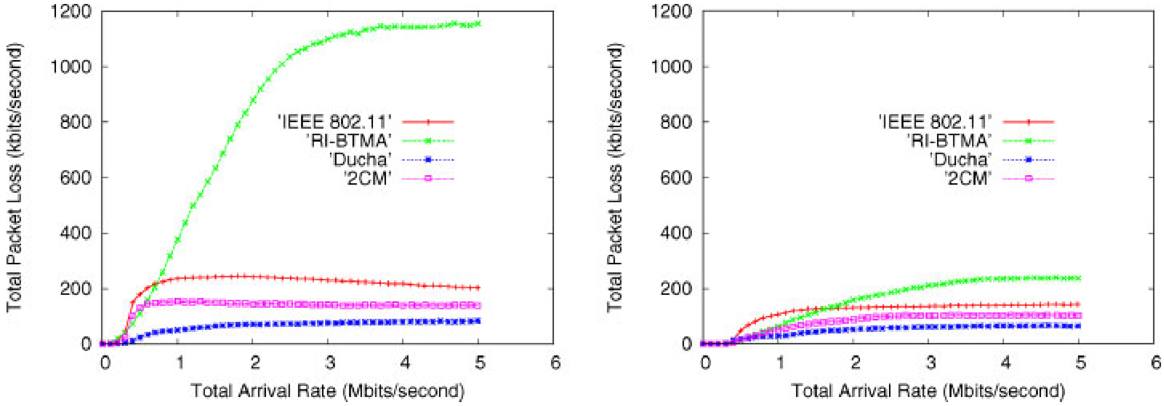

Fig. 11. Total packet loss: additive model (left); capture threshold model (right).

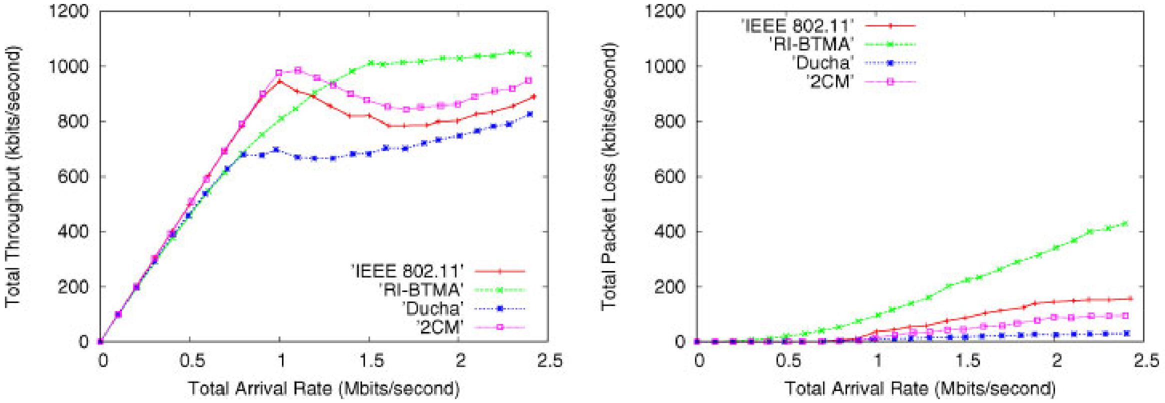

Fig. 12. Ad hoc network (11 Mbps): throughput (left) and packet loss (right).

the total packet loss experienced by the protocols as a function of the arrival rate. Comparing the curves for the additive model on the left, with those for the capture threshold model on the right, we see that RI-BTMA experiences a *tremendous* amount of packet loss at the link layer, which is underestimated by a factor of almost 6 by the capture threshold model. The problems we identified with RI-BTMA in Section 4 get aggravated *drastically* in a larger multi-hop setting with more concurrent transmissions. Due to several concurrent transmissions, not only do DATA packet collisions become inevitable, but also transmitters repeatedly infer a busy-tone CTS incorrectly, and transmit DATA packets even though the receiver is not actually ready. RI-BTMA is simply not an acceptable alternative for such scenarios. Among the others, 2CM achieves the overall best throughput performance, with reasonably low packet losses.

6.3. Effect of Higher Data-Rates

Figure 12 depicts the throughput and loss characteristics of a 50 node 11 Mbps random ad hoc network (its topology is different from that shown in Figure 10). Again, 25 flows are chosen by randomly pairing the 50 nodes, and the routing is carried out statically as earlier. The simulation experiment is exactly as before. Figure 12 confirms what we have seen so far. DUCHA experiences an overall low throughput due to its control channel, while RI-BTMA experiences a high amount of packet losses at the link layer. 2CM again offers a good throughput performance without sacrificing link layer reliability, unlike RI-BTMA.

The simulation results of the three scenarios presented in this section, confirm the observations we made in Section 4. Further, we clearly see that the 2CM protocol is a very efficient protocol with the following desirable features:

1. 2CM uses the busy-tone signal in conjunction with the RTS/CTS of IEEE 802.11, to considerably mitigate the hidden node problem.
2. It does not waste any bandwidth in terms of additional control channels.
3. Finally, it provides a reliable packet-based link layer acknowledgment mechanism.

We also note that the importance of a realistic physical channel model, cannot be overemphasized. With a capture threshold model, RI-BTMA appears to be a very effective protocol. However, under closer evaluation with a more realistic physical channel model, this conclusion breaks down. Indeed, phenomena like DATA packet collisions, or misinterpretation of busy-tone based protocol handshakes, which can happen in reality, can be observed only under an additive interference based channel model. As the deployment scenarios for wireless networks become more complicated, the details of the physical layer can no longer be ignored by the network designer.

7. Conclusion

The performance of multi-hop wireless networks depends on the efficiency and robustness of the underlying protocols, most notably, the MAC protocol. In order to select an efficient and robust MAC protocol, it is critical that we understand the various problems affecting performance, and their relative significance. Moreover, as deployment scenarios for these network become more and more complicated, it is essential to use accurate models of the physical layer, to prevent being misled. We believe our work is a significant step in this direction. Based on a careful study of some of the currently proposed MAC protocols, we are able to understand the relative importance of some of the well-known problems affecting MAC protocols. The success of our approach is best illustrated through the good performance of our proposed 2CM protocol, in a variety of scenarios.

While our work is an important step, it also raises several questions. The performance of a MAC protocol depends not just on the collision avoidance mechanism used, but also on several other factors such as the network topology, the routing, the contention resolution algorithm, and the tuning of the protocol parameters such as the busy-tone power level. Clearly through careful tuning, the performance of a MAC protocol can be improved on a scenario-to-scenario basis. However, we feel that a protocol such as 2CM has just the right combination of features. Coupled with an understanding of how the topology, routing, and contention resolution impact performance, a protocol based on the framework of 2CM can be expected to perform consistently well in any scenario.

Appendix: Additional Material

RI-BTMA

RI-BTMA or received-initiated BTMA [2] is a generalization of BTMA for multi-hop wireless networks. The principle of this protocol is very simple. A node wishing to initiate a transmission, sends a short control message (similar to RTS in IEEE 802.11) to its intended recipient, provided it cannot hear any busy-tone signal. If the receiver node can hear the request message without collisions, it responds by asserting a busy-tone signal. The busy-tone signal serves the dual purpose of letting the transmitter know the readiness of the receiver (like a CTS in IEEE 802.11), and of silencing nodes which may cause a collision at the receiver. The transmitter node responds by sending the data message, and the receiver node continues to assert the busy-tone for the entire duration of the data message. Upon conclusion of the data message, the receiver stops the busy-tone. Since the busy-tone is asserted for the entire duration of the data message, it is more efficient than a CTS message as in IEEE 802.11, which could be 'missed,' as discussed in Section 3. RI-BTMA is susceptible to link layer congestion as explained in Section 4, and also does not provide a link layer acknowledgment for a data transmission.

DUCHA

DUCHA [3] proposes to use a total of three channels for communication. The first is a data channel which is used only for data communication, the second is a control channel which is used for exchanging RTS, CTS, and NCTS (not-clear-to-send) messages, and the third is a narrowband, busy-tone channel which is used to assert a busy-tone. To initiate a communication, the transmitter node sends an RTS message in the control channel, provided it cannot perceive any busy-tone power. If the RTS is received collision-free, the receiver responds with a CTS message, provided it does not perceive any activity in the data channel. If the data channel is perceived active at the receiver, it responds with an NCTS message which informs the transmitter an estimate of when the receiver might become available. After a successful RTS-CTS dialog, the transmitter sends the data message, during the entire duration of which the receiver node asserts a busy-tone signal. Incorrect reception of the data message is notified to the transmitter, by a NACK message which is nothing but the busy-tone signal of the receiver that continues to be asserted after the data transmission is over. Since DUCHA uses an additional logical control channel, the overall channel bandwidth needs to be split appropriately between the data and the control channels. The separation of the control channel from the data channel, and the use of the NCTS message, mitigate the link layer congestion problem considerably, by alerting the transmitter to retry later.

Acknowledgments

This work was supported in part by a research grant from NSF (contract no. 0087266), by an NSERC Discovery Grant (contract no. RGPIN42023-37200), and by a Bell University Laboratory grant.

References

1. IEEE 802 LAN/MAN Standards Committee. Wireless lan medium access control (MAC) and physical layer (PHY) specifications. *IEEE Std. 802.11-1999 Edition*, 1999.
2. Wu C, Li VOK. Receiver-initiated busy-tone multiple access in packet radio networks. In *Proceedings of the ACM SIGCOMM'87*, 1987, pp. 336–342.
3. Zhai H, Wang J, Fang Y, Wu D. A dual-channel mac protocol for mobile ad hoc networks. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM 2004), Workshop on Wireless Ad Hoc and Sensor Networks*, November 2004.
4. The *ns2* network simulator. Available at http://www.isi.edu/nsnam/ns/

5. Tobagi F, Kleinrock L. Packet switching in radio channels: Part ii—the hidden terminal problem in carrier sense multiple-access and the busy-tone solution. *IEEE Transactions on Communications*, COM-23 1975; **12**: 1417–1433.
6. Karn P. Maca—a new channel access method for packet radio. In *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, 1990, pp. 134–140.
7. Bharghavan V, Demers A, Shenker S, Zhang L. Macaw: a media access protocol for wireless lans. In *Proceedings of the ACM SIGCOMM'94*, 1994, pp. 212–225.
8. Fullmer CL, Garcia-Luna-Aceves JJ. Floor acquisition multiple access (fama) for packet radio networks. In *Proceedings of the ACM SIGCOMM'95*, 1995, pp. 262–273.
9. Fullmer CL, Garcia-Luna-Aceves JJ. Solutions to hidden terminal problems in wireless networks. In *Proceedings of the ACM SIGCOMM'97*, 1997, pp. 39–49.
10. Haas Z, Deng J. Dual busy tone multiple access (dbtma)—a multiple access control scheme for ad hoc networks. *IEEE Transactions on Communication* 2002; **50**(6): 975–985.
11. Li Y, Wu H, Perkins D, Tzeng N, Bayoumi M. Mac-scc: medium access control with a separate control channel for multi-hop wireless networks. In *Proceedings of the 23rd International Conference on Distributed Computing Systems Workshops (ICDCSW'03)*, May 2003.
12. Dhoutaut D, Guerin-Lassous I. Impact of heavy traffic beyond communication range in multi-hop ad hoc networks. In *Proceedings of the 3rd International Network Conference (INC 2002)*, July 2002.
13. Ray S, Carruthers J, Starobinski D. Rts/cts-induced congestion in ad hoc wireless lans. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC 2003)*, 2003, pp. 1516–1521.
14. Xu S, Saadawi T. Does the ieee 802.11 mac protocol work well in multi-hop wireless ad hoc networks? *IEEE Communications Magazine* June 2001, pp. 130–137.
15. Singh S, Raghavendra CS. Pamas: power aware multi-access protocol with signaling for ad hoc networks. *ACM Computer Communication Review* 1998; **28**(3): 5–26.
16. Kulkarni SS, Iyer A, Rosenberg C. An address-light, integrated mac and routing protocol for wireless sensor networks. *IEEE/ACM Transactions on Networking*, August 2006, to appear.
17. Luo H, Lu S, Bharghavan V. A new model for packet scheduling in multi-hop wireless networks. In *Proceedings of the ACM/IEEE MobiCom 2000*, August 2000, pp. 76–86.
18. Nandagopal T, T-E. Kim, Gao X, Bharghavan V. Achieving mac layer fairness in wireless packet networks. In *Proceedings of the ACM/IEEE MobiCom 2000*, August 2000, pp. 87–98.
19. Johnson DB, Maltz D, Broch J. The dynamic source routing protocol for multi-hop wireless ad hoc networks. *Ad Hoc Networking*, 2001, pp. 139–172.
20. Perkins CE, Royer EM. Ad hoc on-demand distance vector routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, February 1999, pp. 90–100.

Authors' Biographies

Aravind Iyer received his B.Tech. and M.Tech. degrees in Electrical Engineering, from the Indian Institute of Technology (IIT) Bombay in August 2002. He is currently pursuing his Ph.D. degree at the School of Electrical and Computer Engineering at Purdue University. His research interests include the design, modeling, optimization, and performance evaluation of wireless sensor networks and ad hoc networks. His email address is iyerav@ecn.purdue.edu.

Catherine Rosenberg received her M.Sc. degree from UCLA and her Doctorat en Sciences from University of Paris, Orsay, France. She is currently the Chair of the Department of Electrical and Computer Engineering at the University of Waterloo, Canada, where she also holds a University Research Chair. She has authored over 70 papers and has been awarded six patents in the USA. Dr. Rosenberg is currently an Associate Editor for IEEE Transactions on Mobile Computing. Her email address is cath@ece.uwaterloo.ca.