

Deafness: A MAC Problem in Ad Hoc Networks when using Directional Antennas

Technical Report
July 2003

Romit Roy Choudhury⁺ and Nitin H. Vaidya^{*}

⁺Dept. of Computer Science,
^{*}Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign
{croy, nhv}@uiuc.edu

Abstract—The benefit of using directional antennas in ad hoc networking has received increasing attention. By transmitting towards a desired direction, a node may be able to reduce wireless interference. *Signal to Interference Ratio (SIR)* may also improve if the receiver selectively beams in the direction of the *Signal of Interest*, thereby suppressing interferences that arrive from unwanted directions. However, selective beamforming introduces new problems in wireless medium access control (MAC). One such problem is *deafness*, caused when a node, C, is unable to communicate with another node, A, because A is beamformed in a direction away from C. Node C may misclassify the absence of a reply from node A as a sign of network congestion. Misclassification may affect the performance of MAC and, potentially, higher layer protocols. This paper evaluates the impact of *deafness* on wireless medium access control. We propose a tone-based directional MAC protocol (ToneDMAC) to address this problem. Simulation results show that our protocol performs better than existing directional MAC protocols.

I. INTRODUCTION

Advances in beamforming technology have motivated current research to revisit some of the problems in wireless networking. Greater spatial reuse and longer communication range are potential benefits of utilizing directional antennas. However, tradeoffs exist. New kinds of hidden terminal problems arise, that are absent when using omnidirectional antennas [1]. Another problem, termed *deafness*, also arises. Briefly, deafness is caused when a node, C, attempts to initiate dialog with a node, A, while A is engaged in communication with another node, B, as shown in Figure 1. Node A fails to receive signals from C since A

remains beamformed towards B over the duration of communication. Node C interprets the absence of a reply from A as indicative of a collision at A¹, and retransmits the packet. This can repeat multiple times, until node A has finished the dialog with B

². As we see later, deafness can become a serious issue when node B has multiple packets to transmit to A. In such cases, node C may repeatedly retransmit with a low probability of success, and finally drop the packet.

Several existing protocols attempt to maximize the benefits of beamforming antennas, and suffer from the problem of deafness as a tradeoff. This paper proposes ToneDMAC to address this tradeoff. ToneDMAC uses multiple tones to implicitly inform a node’s neighborhood of its activity. As detailed later, tones can reduce the overhead associated with transmitting explicit control packets, while serving as a notification signal to those that experience deafness. Unlike other proposals that use “busy tones”, ToneDMAC does not require simultaneous transmission of tones and data packets. We use the term “tones” to imply a form of control channel signaling. Our protocol assumes a single transceiver, having the capability to transmit or receive over multiple channels.

The rest of this paper is organized as follows. In Section II, we discuss related work along with a brief overview of

¹IEEE 802.11 and many other CSMA/CA protocols are designed based on this assumption.

²Deafness may appear even when directional antennas are not in use. For example, reconfigurable antennas may form different radiation patterns, which may not necessarily have narrow beamforms (for example, omnidirectional patterns in different planes). Deafness can occur in this case if a host can communicate with different sets of hosts using the different radiation patterns.

the IEEE 802.11 standard. Section III introduces preliminaries on antenna models and relevant terminology used in the rest of the paper. Section IV describes an existing directional MAC protocol, DMAC, and shows the impact of deafness in detail. Motivated by the observations from Section IV, we propose a tone-based directional MAC protocol, ToneDMAC, in Section V. We evaluate the performance of ToneDMAC in Section VI. Section VII concludes the paper with a brief discussion.

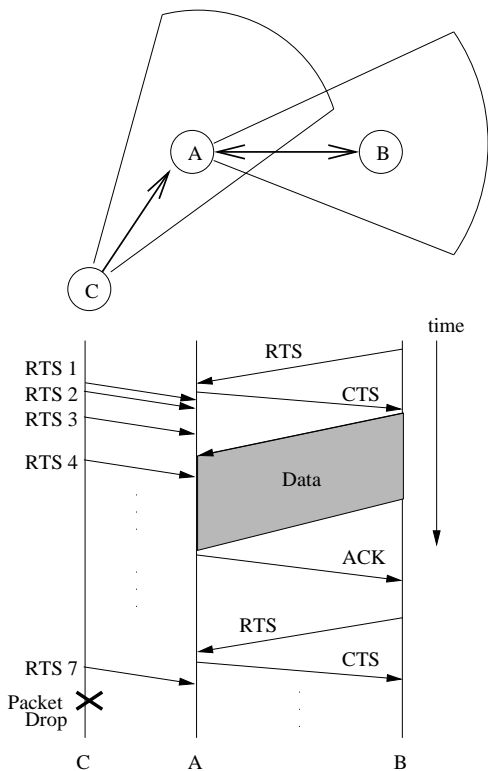


Fig. 1. A scenario illustrating the problem of deafness

II. RELATED WORK

The use of directional or beamforming antennas has been extensively studied in the context of broadband and cellular networks [2],[3],[4],[5]. Recently, attention has been focused on the possibility of using directional antennas for medium access control in multi-hop networks [6],[7],[8],[9],[10],[11],[12][13][14]. In principle, many of the proposed protocols are similar to IEEE 802.11, carefully adapted for use over directional antennas. We present a brief overview of IEEE 802.11, followed by a discussion of the existing protocols for directional medium access control.

A. IEEE 802.11 Distributed Coordinated Function (DCF)

In the IEEE 802.11 MAC protocol [15], an exchange of request to send(RTS)/clear to send(CTS) precedes DATA

communication. Both RTS and CTS packets contain the proposed duration of transmission. Nodes located in the vicinity of communicating nodes, which overhear either of these control packets, must themselves defer transmission for the proposed duration. This is called *Virtual Carrier Sensing* and is implemented through a mechanism called the *Network Allocation Vector (NAV)*. A node updates the value of the NAV with the duration field specified in the RTS or CTS. Thus the area covered by the transmission range of the sender and receiver is reserved for data transfer, to overcome the hidden terminal problem [16].

IEEE 802.11 is a CSMA/CA protocol that performs *physical carrier sense* before initiating transmission. Once the channel is sensed as idle for a DIFS (DCF interframe spacing) duration, 802.11 invokes a backoff mechanism for contention resolution. A node *S* chooses a random *backoff interval* from a range $[0, CW]$, where *CW* is called the *Contention Window*. *CW* is initialized to the value of CW_{min} . Node *S* then decrements the backoff counter once every idle “slot time”. When the backoff counter reaches 0, node *S* transmits the RTS packet. If the transmission from *S* collides with some other transmission (collision is detected by the absence of a CTS), *S* doubles its *CW*, counts down a newly chosen backoff interval, and attempts retransmission. The *Contention Window* is doubled on each collision until it reaches a maximum threshold, called CW_{max} . While in the backoff stage, if a node senses the channel as busy, it freezes its backoff counter. When the channel is once again idle for a duration called DIFS, the node continues counting down from its previous (frozen) value.

B. MAC using Directional Antennas

The design of IEEE 802.11 implicitly assumes an omnidirectional antenna at the physical layer. Although 802.11 may operate correctly when using directional antennas, performance may get affected. Recently, several MAC protocols have been proposed that suitably adapt 802.11 for beamforming antennas. Ko *et al.* [6] have proposed to transmit an RTS directionally, only if the RTS does not collide with other ongoing communications. Nasipuri *et al.* [17] proposes to reduce the interference in the wireless channel by communicating directionally. However, their proposals require the transmission of an omnidirectional CTS to inform the receiver’s neighborhood about the imminent dialog. This offsets spatial reuse – a key advantage of using directional antennas. In [10], Elbatt *et al.* propose an interesting idea – they use RTS/CTS to inform the neighborhood about the beam indices to be used

for the imminent communication. Based on this information, neighbors of the communicating nodes decide which beams may be used for initiating their own RTS packets. Bandyopadhyay *et al.* [18] present another MAC protocol that informs neighborhood nodes about ongoing communications through additional control messages. In addition, the protocol assumes knowledge of network traffic.

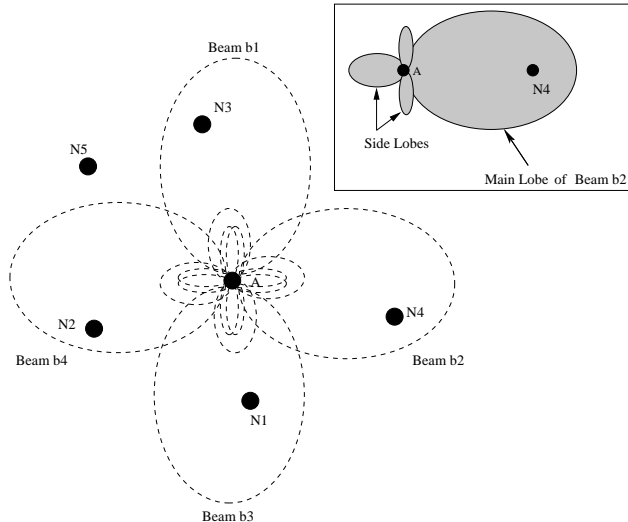


Fig. 2. The antenna radiation pattern of node A – inset shows node A beamformed in the direction of N4.

Takai *et al.* have described the DVCS mechanism in [19], where a node performs directional virtual carrier sensing. This is an useful optimization that increases the spatial reuse of the channel. Roy Choudhury *et al.* [1] have proposed protocols that aim to maximize the benefits of beamforming. They identify several new problems that arise as a tradeoff. Many of these problems have not been addressed in [1].

Some existing proposals have utilized “busy tones” for omnidirectional, and more recently, for directional antennas. In [20], Deng *et al.* have proposed DBTMA, that uses omnidirectional transmit and receive busy tones to avoid hidden terminal problems. Huang *et al.* [12] have extended this idea to the case of directional antennas. The protocol assumes multiple transceivers, capable of transmitting data packets as well as busy tones directionally. Moreover, the protocol suffers from the problem of deafness. Similar to DBTMA, and the protocol in [12], ToneDMAC uses an additional control channel for signaling. However, ToneDMAC requires a single transceiver. Tones are assigned on the control channel, and need not be transmitted in parallel with data packets. In this respect, tones are different from “busy tones” that are typically transmitted when a tone is “busy” with ongoing trans-

mission/reception. Korakis *et al.* [11] attempt to address the problems arising from directional antennas, including deafness. To inform neighbors of a communicating node pair, the authors propose to transmit directional RTS/CTS packets on every beam. Transmitting multiple RTS/CTS packets for each transmitted data packet can drastically degrade performance.

As detailed later, the problem of deafness does not appear when communicating nodes manage to inform their surrounding neighbors about the imminent communication. For example, in Figure 1, C may not attempt communication to A, if C was aware of the ongoing dialog between B and A. However, we later discuss why informing the neighborhood can be expensive, especially under heavy traffic. We argue that not informing neighbors explicitly of an imminent communication can favor performance improvement. This motivates the design of ToneDMAC.

III. PRELIMINARIES

A. Antenna Model

We assume a switched beam antenna system. The assumed antenna model is comprised of N beam patterns ($N = 4$ in Figure 2), sometimes referred to as radiation patterns. An example radiation pattern is shown in Figure 2 – a high gain main lobe points towards a certain direction, and low gain side lobes spread over other directions.

The antenna system offers two modes of operation: *Omni* and *Directional*. We assume that a node can operate only in one mode at a given time, but can toggle between modes with negligible latency. In *Omni* mode, after a signal is detected, the antenna determines the beam on which the received signal power is maximum. The rest of the packet is then received by using this beam. We assume that in omni mode, signals are received with a gain G^o . An idle node stays in the *Omni* mode. In *Directional* mode, a node can select *only one* of its beams and beamform with a main lobe gain of G^d , $G^d \geq G^o$. The expression that relates the main lobe transmit and receive gains (G_T and G_R), to the transmit and receive powers (P_T and P_R) is as follows [21]:

$$P_R = \frac{P_T G_T G_R}{L_p L_o}$$

$$L_p = \left(\frac{4\pi r}{\lambda}\right)^2$$

The term L_o is an additional path loss factor to account for atmospheric absorption, ohmic losses, etc. L_p is called the *free space loss*, caused due to spreading of transmitted waves. λ is the wavelength of the transmitted signals and

r is the physical distance between the transmitter and the receiver. Note that G_T and G_R are transmit and receive gains along the straight line joining the transmitter and receiver.

The physical distance over which two nodes can communicate is proportional to the product of the transmission and the reception gain. As a result, directional antennas offer the capability of range extension. Put differently, two nodes in omni mode may be out of communication range because the product of their omnidirectional transmit and receive gains, $(G^o \times G^o)$, is not large enough. However, if one of the nodes beamforms in the direction of the other, the new product, $(G^d \times G^o)$, may be sufficiently large to enable direct communication.

B. Deaf Zone

The coverage region, C_i , of a beam, b_i , is the region over which a node can communicate using beam b_i . We define a node's neighborhood region, \mathfrak{R} , as the union of the coverage regions, over all N beams. The deaf zone corresponding to beam b_i is defined as $\mathfrak{R} - C_i$. In other words, the portion of the neighborhood region from which a node does not receive signals, while beamformed using beam b_i , is the deaf zone associated with beam b_i . From Figure 2, nodes N1, N2, and N3 are located in the deaf zone of beam b_2 of node A. Node N5 lies outside the neighborhood region, and therefore outside the deaf zone.

IV. PROBLEM OF DEAFNESS

This section describes the problem of deafness, originally identified (but not further investigated) in [1]. We choose an existing directional MAC protocol, DMAC [1], and show how deafness can adversely affect protocol performance. We propose simple solutions to deafness, and expose how new problems can arise as a consequence. Our observations motivate the design of ToneDMAC.

A. Directional MAC (DMAC) protocol

In principle, DMAC is similar to IEEE 802.11, adapted for use over directional antennas. In describing DMAC, we refer to a sender node as S and the receiver node as R.

- *Physical Carrier Sensing and Backoff*

When using DMAC, physical carrier-sense is performed in the directional mode, using the same beam that must be used for the immediate communication. If the carrier is sensed idle for a DIFS duration, DMAC requires the node to choose a backoff interval, in a fashion similar to 802.11. The sender must remain in the directional mode

while counting down its backoff counter – this will be referred as directional backoff. If the carrier is sensed busy during the count-down, the node defers transmission for later and switches back to the omnidirectional mode. If the carrier remains idle, the sender node initiates channel reservation.

- *Channel Reservation*

Channel reservation in DMAC is performed using an RTS/CTS handshake, both being transmitted directionally. Sender S transmits a directional RTS (DRTS) meant for the receiver R. An idle node listens to the channel omnidirectionally. When it receives a signal arriving from a particular direction, it locks on to that signal and receives it. Assuming that R was idle, it receives the DRTS from S. Using a suitable beam to reply to S, node R transmits a directional CTS (DCTS).³

Nodes, other than S and R, that overhear the RTS and/or CTS, remember the directions from which the reservation messages arrived. These nodes defer their own transmissions in these directions, for the proposed duration of transmission between S and R. This is called directional *Virtual Carrier Sensing* [19],[1]. Directional NAV (DNAV) tables maintain the virtual carrier sensing information. To ensure that a new communication is not initiated in the direction of an ongoing communication, the transmitter must first consult the DNAV table. Only if the DNAV check is successful, will the node proceed to physical carrier-sense in that direction.

- *DATA/ACK exchange*

The exchange of DATA/ACK packets is similar to 802.11, except that they are transmitted directionally. After transmitting the RTS, the sender S waits for the CTS, using the beam that it had used to transmit the RTS. If the CTS arrives within a *CTS-timeout* duration, S transmits the DATA packet directionally. R acknowledges successful data reception with a directional ACK. If the CTS does not arrive within the specified timeout duration, S chooses a new backoff value from a doubled *Contention Window*, counts down this backoff value, and retransmits the RTS (similar to 802.11). Once the exchange of DATA/ACK packets are over, both S and R switch back to the omnidirectional mode.

B. Deafness in DMAC

To show the impact of deafness on DMAC, let us consider Figure 3. Assume that all nodes are idle and that B

³DMAC [1] assumes that a higher layer is responsible for directional neighbor discovery.

intends to transmit a data packet to A. DMAC requires B to beamform in the direction of A, and detect if the channel is idle for a DIFS duration. If the channel is found idle, B proceeds to the backoff phase and counts down the backoff counter while still beamformed towards A. Observe that while B is counting down its backoff counter, node X may intend to communicate with B. If X completes its own backoff before B and transmits an RTS to B, B would not receive the RTS. In the absence of a reply from B, X would repeatedly backoff and retransmit the RTS, until the dialog between B and A is over. Unproductive retransmissions is an outcome of deafness.

Now consider the case where B has multiple packets to transmit to A. Once B has finished transmitting the first packet, it immediately prepares to transmit the next packet by beamforming in the direction of A, and then repeating the sequence of DMAC operations, namely, carrier-sense, backoff, DRTS, DCTS, etc. Note that unless B is in the omni mode, X would not be able to communicate with B. When using DMAC, B would not switch to the omni mode until it has finished transmitting all its queued packets for A. If B remains backlogged for a long time, X may end up dropping the packets meant for B. A scenario is possible where node Y intends to communicate with node X. If X has multiple packets queued for B, it would remain engaged either in directional backoff or in transmitting a directional RTS. Y would experience prolonged deafness, until X has dropped all its packets. A chain is possible in which none of the nodes communicate successfully – a “deadlock”. This is a serious problem, caused when the intended receiver of a node is itself a transmitter.




Fig. 3. An example scenario in multi-hop ad hoc networks.

Once the directional backoff phase has been completed, DMAC requires B and A to exchange the DRTS/DCTS packets. From Figure 3, observe that nodes X and C (in the neighborhood region of B and A, respectively) do not receive either of these packets, and remain unaware of the ongoing communication. Deafness arises if C intends to initiate communication with A. Notice that the intended receiver of C (i.e., A), is itself a receiver of another communication. C may attempt multiple transmissions while A is engaged in receiving DATA packets from B. Because

of the longer transmission-duration of DATA packets, the possibility of deafness is proportionally higher. Repeated failure to communicate to A, causes C’s contention window to grow exponentially.

Now consider a point of time when the dialog between B and A is over. Due to the exponential growth of C’s contention window, it is likely that node C has chosen a large backoff value. C continues to count down its backoff duration, although node A has switched back to the omnidirectional mode, and is no longer deaf to C. In the meantime, if B has more packets queued for A, it chooses a backoff duration from the minimum backoff interval, $[0, CW_{min}]$. With a high probability, this chosen backoff would be less than C’s remaining backoff duration. Therefore, B counts down its backoff duration first, and initiates successful communication once again. Later, when C retransmits the DRTS, it once again receives no reply from A. Retransmissions may continue several times until C has reached its RTS-retry-limit. At this time, C is forced to drop the DATA packet. If C has more packets to send to A, it chooses a new (and potentially smaller) backoff duration from the interval $[0, CW_{min}]$ and begins counting down. However, choosing a small backoff does not ensure channel access. Observe that if B is communicating with A while C has counted down the small backoff, the entire process may repeat. Simulations show that in a simple 3 node scenario, C may drop several packets before it successfully transmits a packet to A. Interestingly, once C succeeds in communicating with A, the problem of deafness appears at B. B is now located in the deaf zone of A. B drops multiple packets before it gets fortunate enough to steal channel access back from C. Multiple packet drops at the source node, without actual congestion or link failure, can adversely affect performance. Higher layer protocols, that use packet drops as indicators of the network condition, may be misled. End-to-end throughput and latency can degrade. We quantify these impacts in Section VII.

Unfairness is also an outcome of deafness. When multiple nodes (say, C and F in Figure 3) attempt to communicate with node A, the node that wins channel contention retains the privilege to access the channel for a long time. Although the receiver remains busy almost all the time, the transmitter nodes experience short-term unfairness. As we show later, the variance of the end-to-end delay increases, which in turn affects throughput.

C. Addressing Deafness

A rather simple strategy to alleviate (or avoid) deafness could be to transmit the RTS/CTS omnidirectionally. Om-

unidirectional neighbors of the communicating nodes would be informed of the ongoing dialog, and therefore, would refrain from initiating communication. Consider an example scenario in Figure 4. Nodes X, F, and C receive the omni-RTS and/or CTS packets from B and A, and defer their own transmissions towards B and A for the proposed duration. Deafness does not appear.




Fig. 4. A scenario to illustrate the problems when using omnidirectional mode for DMAC protocols.

However, several tradeoffs arise.

1. Using omnidirectional RTS/CTS limits the range extension capability of directional antennas. Node B may no longer communicate directly with node Y, otherwise possible when using DRTS/DCTS. Networks may get partitioned. Routes may grow longer because of shorter links, affecting throughput and delay.
2. Even if range extension is sacrificed by transmitting omnidirectional RTS/CTS, problems still remain. Consider the case where node C receives an omni-CTS from A, and records the direction towards which it must not initiate communication. Now, if C intends to communicate to node D, it must transmit an omni-RTS itself. DNAV restrictions at C inhibits any transmission towards A – therefore an omnidirectional transmission must be inhibited. This leaves two choices to node C – (i) to defer transmission until the dialog between B and A is over, or (ii) to transmit a directional RTS towards D in order to respect the DNAV restrictions. The first option sacrifices spatial reuse, a key advantage of directional antennas. Communications that are possible simultaneously get serialized over time – a waste of network capacity. The second option suffers from deafness again. Observe that node F does not receive the DRTS transmission from C to D, and thus, can attempt communication to C. Since C would be beamformed towards D, it would fail to receive F’s RTS. F would continue reattempting without success. Deafness remains unresolved.

V. TONEDMAC – PROTOCOL DESCRIPTION

The key contribution of ToneDMAC lies in the use of multiple tones to alleviate deafness. We observed that deafness occurs primarily because a transmitter is unaware of the activities of its intended receiver. The main idea of our protocol is to inform a node’s neighborhood of its activity, through omnidirectional tone-notification. Nodes that experience deafness may use these notifications to suitably schedule transmissions to its intended receiver. Previous proposals in directional MAC have used “busy tones” to replace RTS/CTS, or to alleviate hidden terminal problems [12],[22],[23]. We use tones for a different purpose. We detail the ToneDMAC protocol in this section. The scenario in Figure 5 has been used for protocol description.

In ToneDMAC, the common channel is split into two sub channels: a data channel and a narrow control channel. RTS, CTS, DATA, and ACK packets are transmitted on the data channel. The tones (essentially sinusoids with sufficient spectral separation) are assigned on the control channel. We assume that an idle transceiver is capable of tuning to any tone/signal that arrives either on the control or data channel.

A. Carrier Sensing and Backoff

Consider Figure 5. If node B intends to initiate a transmission to node A, it beamforms in the direction of A and performs directional physical carrier sensing. If the channel is idle, B selects a backoff duration, uniformly chosen from an interval of $[0, CW_{min}]$. Unlike DMAC, ToneDMAC requires a node to switch back to the omnidirectional mode while performing the backoff countdown. While backing off in the omnidirectional mode, a node senses the channel as busy only if a signal arrives from the direction in which the node intends to transmit. However, if a RTS or CTS arrives from other directions, a node will be capable of receiving them. This mitigates the “deadlock” problem arising from directional backoff, discussed in detail in the previous section. When the backoff count successfully reaches zero, B proceeds to channel reservation.

B. Channel Reservation, Data Communication

To exploit range extension, ToneDMAC requires node B to transmit a DRTS, to which A replies with a DCTS. The channel reservation and data communication phases are similar to DMAC. DATA and ACK packets are exchanged directionally. Nodes like F, that overhear the DRTS and/or DCTS, update their DNAV tables suitably.

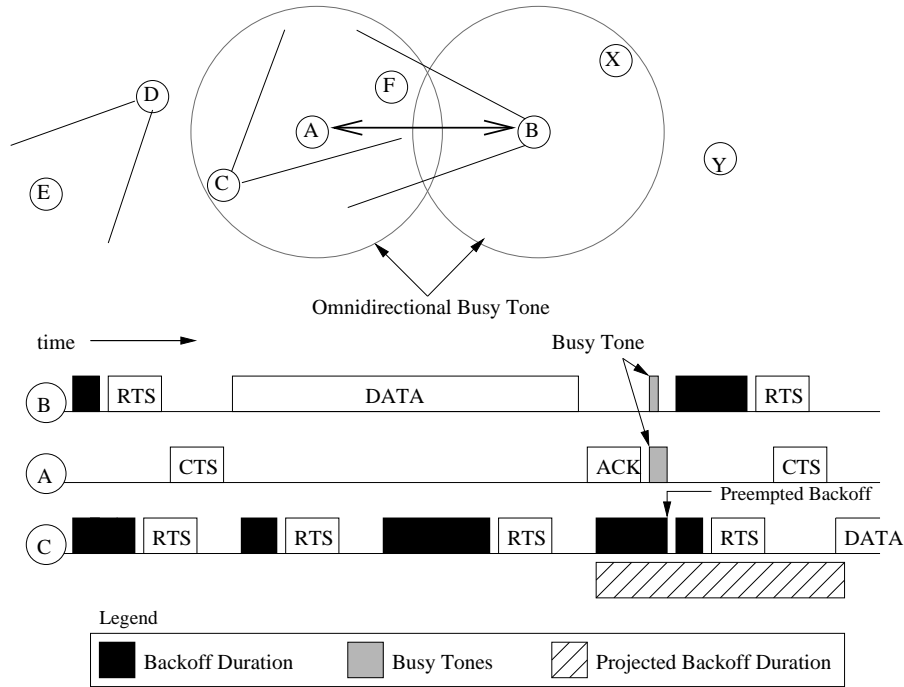


Fig. 5. An example scenario and timeline of the ToneDMAC protocol.

Observe that nodes like X and C, located away from the beamformed direction of B and A respectively, do not receive either of DRTS or DCTS packets, and remain unaware of the communication between B and A. Once directional DATA/ACK exchange is accomplished, both B and A switch back to the omnidirectional mode.

C. Tone Transmission

After exchanging the DATA/ACK, and switching back to the omnidirectional mode, nodes B and A initiate tone transmission. Tones serve as indicators that nodes B and A were recently engaged in communication. A neighboring node, C, unable to communicate with A in the meanwhile, can utilize the tone from A, as an indication of recent deafness. Realizing deafness, C can suitably modify its retransmission strategy. The motivation of using tones (instead of explicit control packets) is as follows. Control packets, however small, are prepended with a physical layer PLCP preamble and header for receiver demodulation. Together, the preamble and the header consume multiple time slots for every transmitted packet. This leads to excessive wastage in channel bandwidth. Tones overcome this problem. However, a tradeoff arises. Since tones do not contain information, a node may not be able to identify the sender of a tone. However, to determine whether its intended receiver is deaf, a node needs to correctly identify this sender. As an example, C must ignore tones trans-

mitted by D, and react to tones that arrive only from A⁴. Clearly, correct identification may not be possible if a single tone is used by all the nodes. Ambiguities may arise, leading to misclassification.

In ToneDMAC, we use a group of tones and different transmission-durations, to reduce the probability of misclassification. A node i chooses a tone τ_i from a set of K tones, and an integer time duration t_i from an interval $[1, T]$. The tuple (τ_i, t_i) serves as the signature of node i . The values of τ_i and t_i are static hash functions of the node's identifier. We assume that a higher layer is capable of assigning consecutive identifier's to nodes, so that the tuple (τ_i, t_i) can be uniformly distributed over the node identifier space. A simple hash function is used to assign (τ_i, t_i) to a node i .

$$\begin{aligned}\tau_i &= (i \bmod K) + 1 \\ t_i &= (i \bmod T) + 1\end{aligned}$$

Nodes B and A now transmit their corresponding tones in the omnidirectional mode. To cover the neighborhood region, the transmit power of the tones are suitably increased. Node B transmits τ_B for t_B slots, node A transmits τ_A for t_A slots. Observe that DNAV restrictions are not applied to the control channel.

⁴Observe that nodes D and A might be engaged in distinct dialogs – D to E and A to B. Tones transmitted by both these nodes would arrive at C.

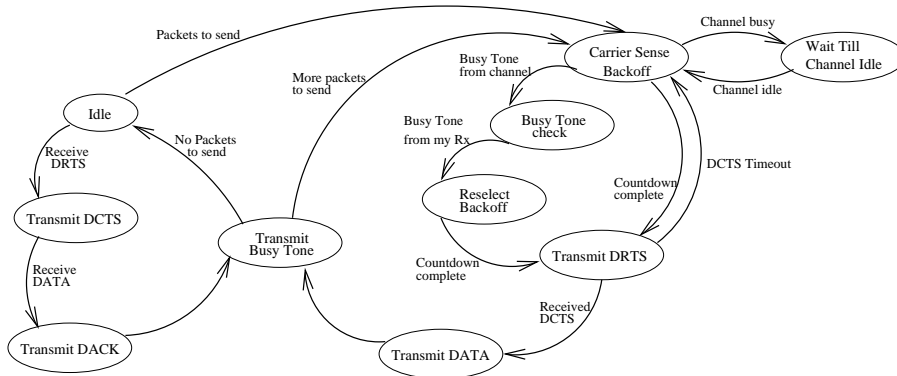


Fig. 6. A finite state machine for ToneDMAC.

D. Tone Check

On receiving a tone, a backlogged node needs to check whether the tone originated from its intended receiver. Based on the tone frequency and the transmission durations, it might be possible to deduce the node identifier of the source, with reasonable accuracy. False alarms may occur when (i) two or more neighbors of a node are hashed to identical signatures (τ, t) , or (ii) two neighbors, with signatures differing only in the values of t , transmit tones that overlap in time. ToneDMAC utilizes antenna capabilities to reduce false alarms. In Figure 5, assume that the identifiers of nodes A and D map to identical values of (τ, t) , say (τ_q, t_q) . Also assume, for this example, that A is the sender of the tone. On receiving tone τ_q for t_q slots, C first determines the beam on which the received signal strength is maximum - the beam facing north-east in this case. Assuming that node C is aware of its neighbors on each beam, it can immediately eliminate D, as the potential sender of the tone. Of course, if both A and D were located on the same beam, accurate classification would not be possible. Simulation results have shown that under reasonable node density, with $K = 4$, $T = 3$, and $beamwidth = 60^\circ$, probability of misclassification is small.

Once C infers that the tone is from node A, it realizes that A was deaf in the recent past. If C does not intend to communicate with A, it ignores the tone completely. Otherwise, C enters the *Reselect Backoff* phase.

E. Reselect Backoff

Once a backlogged node infers that a overheard tone originated from its intended receiver, it preempts its current backoff countdown, resets its contention window to its minimum value, CW_{min} , and selects a new backoff from the interval $[0, CW_{min}]$. The node now enters the carrier-sense/backoff phase, as shown in the timeline in Figure 5.

Clearly, based on the tone notification from node A, node C now has a fair chance to win channel contention – assuming B has more packets to send to A, both B and C choose backoff values from the same contention window $[0, CW_{min}]$. Observe in Figure 5 the “Projected Backoff Duration” of node C. The “Projected Backoff Duration” shows the duration over which C would have continued backing off in the absence of the tones. As a result, node B would have won channel contention yet again. Clearly, ToneDMAC ensures a higher degree of fairness. Moreover, packet-loss probability reduces through explicit tone-based notifications from deaf nodes. Figure 6 shows a partial finite state machine, summarizing ToneDMAC.

VI. PERFORMANCE EVALUATION

We use the Qualnet simulator [24], version 3.1, for simulating ToneDMAC. We compare our protocol mainly with DMAC [1], and quantify some of the key impacts of deafness on medium access control. The transmit power is assigned in a way such that the communication range of directional transmissions is approximately 300 meters. The assumed data rate is 11 Mbps. We have used both UDP and TCP traffic in our experiments. Sources are always backlogged unless mentioned otherwise. We do not consider mobility.




Fig. 7. Scenarios affected by deafness.

A. Simulation Results

Deafness may arise in two types of scenarios (Figure 7) – (i) when multiple senders intend to transmit simultane-

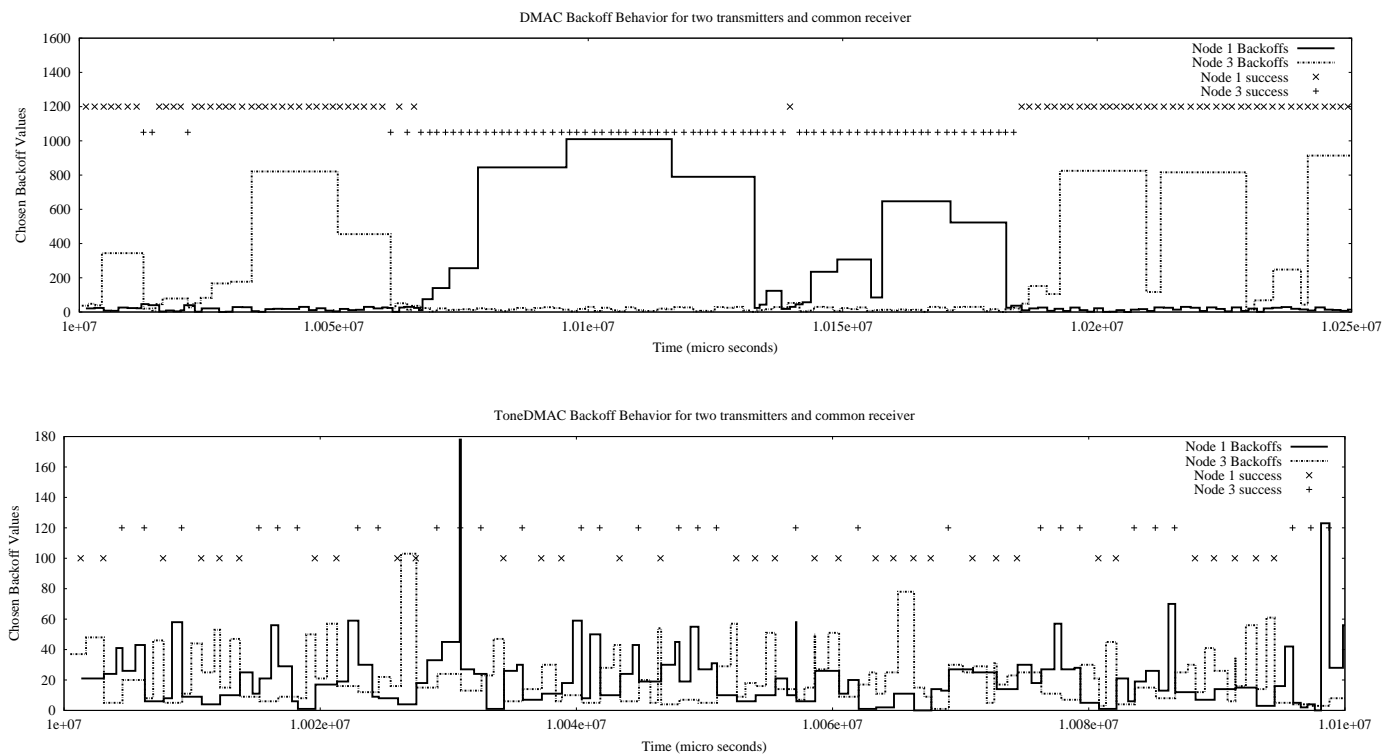


Fig. 8. Behavior of the backoff values against time, for scenario (i). Successful packet transmissions have also been marked.

ously, to a common receiver, and (ii) when the intended receiver of a transmission is itself busy transmitting in some other direction. The arrows in Figure 7 indicate single hop flows. We compare the performance of DMAC and ToneDMAC for these simple scenarios, and then show the overall effects in larger multi-hop networks.

As discussed previously, ToneDMAC differs from DMAC primarily in two ways – (i) Omnidirectional Backoff and (ii) Tone Feedbacks. To understand the individual benefits of each of these factors, we simulated a protocol called ZeroToneDMAC. As explained later, ZeroToneDMAC is identical to ToneDMAC, except that it does not use tones. Put differently, ZeroToneDMAC can be viewed as DMAC with omnidirectional backoff.

We begin evaluation of ToneDMAC, with an analysis of scenario (i). Later we will discuss scenario (ii), and point out important differences in performance.

- Scenario (i)

When using DMAC, Figure 8 shows the variation of backoff values against time, for nodes 1 and 3. (Please note that the graphs in Figure 8 are plotted using a step function. For example, if backoff values B1 and B2 are cho-

sen at times t_1 and t_2 , then a horizontal line at a height of B1 extends from time t_1 to t_2 followed by a jump to a height B2, at time t_2). Evident from Figure 8, the backoff value of one of the nodes remains small for a long interval of time, while that of the other continues to grow. The node with smaller backoff communicates successfully to the common receiver, while the other node experiences prolonged deafness. After multiple packet transmissions, the situation reverses. Figure 8 captures this behavior of DMAC. Observe that node 3 loses channel contention initially, and is forced to choose increasingly larger backoffs. Node 1 repeatedly wins channel contention, and transmits packets back to back – shown by the sequence of “cross” marks on the graph. Later, node 3 gets fortunate enough to complete counting down earlier than node 1, and steals channel access. Node 1 now suffers from a growing backoff, while node 3 transmits multiple packets in sequence (shown by “plus” marks on the graph). The alternation continues. A trend is visible whereby each node transmits multiple packets once it grabs the channel. Clearly, short-term unfairness is an outcome of DMAC. Figure 8 shows the variation of backoff values when ToneDMAC is used. As evident, backoff values remain low and channel access is performed with reasonable fairness. Also, as discussed below, packet drops are fewer when using ToneDMAC.

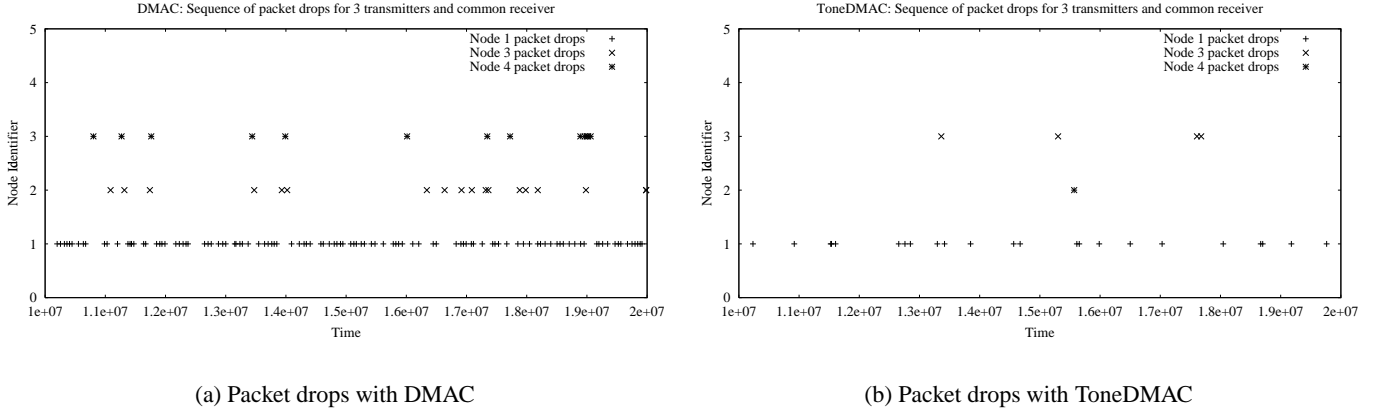


Fig. 9. Packet drop due to deafness in DMAC and ToneDMAC for scenario (i), using UDP traffic

The impact of deafness is accentuated when more than two transmitters intend to transmit to a common receiver. In such scenarios packet drops increase because a transmitter must attempt many more times before it can win channel contention. Put differently, a node must be luckier than all its contenders to be able to successfully initiate transmission. We measured packet drops for a scenario similar to scenario (i), but with 3 transmitters transmitting to a common receiver. The results are presented in Figure 9. As in IEEE 802.11, we used $CW_{min} = 31$, $CW_{max} = 1023$, and $RTS-Retry-Limit = 7$. ToneDMAC drops fewer packets in comparison to DMAC. When using equal sending rates for all the flows, we observed approximately 2.5% packet drops in DMAC, compared to 0.25% in ToneDMAC. When different sending rates are assigned, performance degrades further. If one of the transmitters is always backlogged and another is not, packet drop rate for the latter increases significantly in DMAC. When using ToneDMAC, the degradation is negligible.

We now evaluate throughput. Figure 10(a) shows the results of simulating UDP traffic for scenario (i). The sending rates are equal for both transmitters. With all the adverse effects of deafness, DMAC achieves comparable performance with ToneDMAC⁵. This happens because of the following reason. Over a reasonably long interval of time, the aggregate throughput is determined by the fraction of time the “common receiver” remains busy. Observe that for both DMAC and ToneDMAC, the receiver receives packets at similar rates. The difference lies in the sequence in which channel is accessed and the rate at which packets

⁵DMAC and ZetoToneDMAC performs almost identically, because the advantages of omnidirectional backoff is not exploited in scenario (i).

are dropped. UDP traffic simulations do not reflect these effects.

In terms of average end-to-end delay, the performances of DMAC and ToneDMAC are similar. Again, when using DMAC, packets that remain queued for a long time due to deafness, are transmitted in quick succession once the node wins channel contention. The latency gets amortized over multiple packets. The end-to-end delay, averaged over all packets, is therefore comparable to that of ToneDMAC.

- Scenario (ii)

We now consider scenario (ii), as shown in Figure 7(b). Figure 10(b) shows aggregate throughput obtained by DMAC, ToneDMAC, and ZeroToneDMAC. Unlike scenario (i), ToneDMAC and ZeroToneDMAC outperform DMAC under heavy traffic. The reason follows from our previous discussions. Under heavy traffic, transmitters are always backlogged. When using DMAC, node 3 remains beamformed towards its intended receiver, almost all the time (recall that carrier-sense, backoff, and packet transmissions, are all performed in the directional mode in DMAC). As a result, node 2 repeatedly fails to establish communication with node 3. Moreover, node 2 also remains in the directional mode, since it is either backing off or retransmitting a DRTS. Clearly, node 1 can never initiate successful communication to 2. While node 3 communicates, nodes 1 and 2 drop large number of packets. If node 3 finishes transmitting all the packets, node 2 gets a chance. Only after node 2 finishes transmitting all its queued packets to node 3, does node 1 acquire channel access. Clearly, deafness can cause large number of packet drops in such scenarios. Figure 10(c) indicates that percentage of packet drops increase drastically with increase

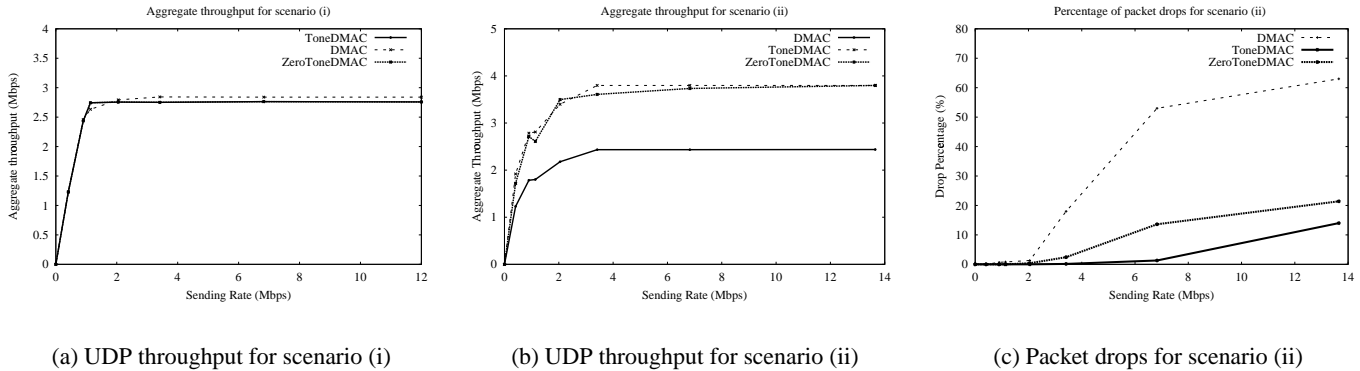


Fig. 10. Performance comparison for DMAC and ToneDMAC for scenarios (i) and (ii), using UDP traffic

in sending rate at the source node. Furthermore, spatial reuse of the channel reduces – the 3 flows in scenario (ii) are serialized over time. When using ToneDMAC and ZeroToneDMAC, recall that nodes remain in the omnidirectional mode during the backoff stage. As a result, two communications take place simultaneously – one initiated by node 1 and the other by node 3. This enhances spatial reuse, and reduces packet drops. ToneDMAC clearly outperforms DMAC in terms of aggregate throughput and packet drop rate. End-to-end delay is also lower with ToneDMAC. ZeroToneDMAC performs comparably with ToneDMAC because the benefit of tones is not conspicuous in a topology like scenario (ii). Since both ZeroToneDMAC and ToneDMAC uses omnidirectional backoff, they both alleviate the “deadlock” problem discussed earlier. We evaluate large multihop networks next, and show how the higher improvements can be achieved when using ToneDMAC.

• Multi-hop Networks

Figure 11 shows protocol performance when multi-hop UDP traffic is simulated, in a network of 30 nodes. The nodes were placed randomly in a region of $1500 \times 1500 m^2$. Random source destination pairs were chosen for 5 flows, and minimum-hop routes were assigned statically. We used $K = 4$ and $T = 3$ (i.e., 4 tones and maximum 3 slots) for ToneDMAC. The simulation results are averaged over 25 runs. Figure 11 indicates that, ToneDMAC attains higher aggregate throughput compared to DMAC⁶. The reason is attributed to the frequent occurrence of scenarios (i) and (ii) in large networks. The result also implies that when using ToneDMAC, tone misclassification and bandwidth wastage (due to control channel signaling) does not degrade the performance of ToneDMAC.

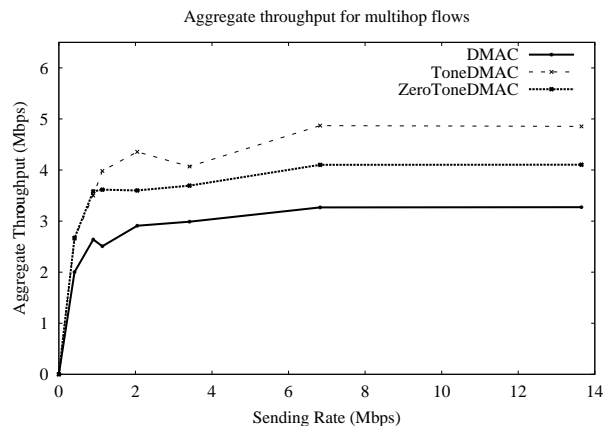


Fig. 11. Throughput comparison for multi-hop flows in a random network.

The improvement of ToneDMAC over DMAC may be attributed to two main modifications – (1) notification from tones and (2) omnidirectional backoff. To evaluate the individual benefits of each modification, we simulated ToneDMAC, without the tones (i.e., $K = 0$ and $T = 0$) – called ZeroToneDMAC. The performance of ZeroToneDMAC is indicative of the benefits derived simply from omnidirectional backoff. The differential of ZeroToneDMAC and ToneDMAC is, approximately, indicative of the gains due to implementing tones. Figure 11 shows the throughput of ZeroToneDMAC, when simulated in multihop networks. Clearly, wherever cases like scenario (ii) appeared in the network topology, both ZeroToneDMAC and ToneDMAC benefitted due to omnidirectional backoff. Whenever cases similar to scenario (i) occurred, ToneDMAC outperformed ZeroToneDMAC (and DMAC) due to the additional tone feedback. Moreover, cases where packets were forwarded over multiple hops, ToneDMAC proved to be beneficial over the other protocols. To illustrate this, consider a case where source

⁶DMAC outperforms IEEE 802.11 in similar scenarios, shown in [1].

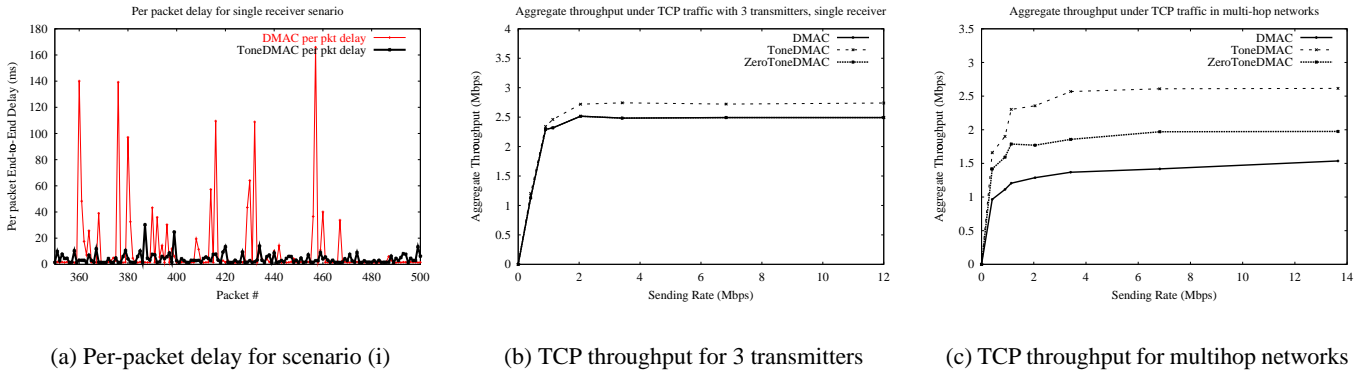


Fig. 12. Per packet delay under scenario (i) and throughput of DMAC and ToneDMAC under TCP traffic

node S intends to communicate with destination node D , through an intermediate node I . Consider ZeroToneDMAC or DMAC in this situation. While node I forwards packets to node D , observe that node S would continue to retransmit to I . As explained earlier, node S would exponentially increase its contention window upon each failure. Since S 's backoff window can quickly grow to a large value, S would backoff even while node I has finished forwarding the packet to D . After a comparatively longer time (when countdown has finished) S initiates communication with I . Clearly, slots are wasted in unnecessary backoff. When using ToneDMAC, S receives a tone from I and therefore can initiate communication sooner. Consequently, higher throughput can be attained using tone-notifications. This is evident from Figure 11. Fairness among the flows were also observed to be higher in ToneDMAC compared to ZeroToneDMAC and DMAC.

The outcome of mitigating deafness may provide additional advantages when network layer issues are considered. For example, fewer packet-drops may trigger fewer route errors. Suboptimal route selection, caused due to unheard route-discovery packets, affects end-to-end throughput and delay [13]. ToneDMAC may alleviate such problems. The results presented in this paper do not reflect such potential performance gains. Evaluating the impact of ToneDMAC on the network layer is outside the scope of this paper, and is a part of our future work. However we believe that when a suitable routing protocol is executed, ToneDMAC would outperform DMAC even further.

- TCP in single and multi-hop networks

For the common receiver scenario, the average end-to-end delay for ToneDMAC was observed to be comparable to DMAC. As discussed earlier, packets that were queued for a long time due to deafness, were transmitted in quick suc-

cession once the node won channel contention. As a result, latency got amortized, leading to comparable average end-to-end delay. However, in addition to the ‘‘average’’, the ‘‘variance’’ of per-packet delay is also an important metric that needs to be compared. Figure 12(a) compares per-packet end-to-end delay. (Graphs for ZeroToneDMAC has not been plotted because it is almost identical to DMAC). The end-to-end delay is calculated as the latency from the time a packet is dispatched at the source application, till the time the packet is received at the destination application. Figure 12(a) shows how DMAC’s delay is characterized by large fluctuations. When using ToneDMAC, the per-packet delays vary within a small range – tone notifications schedule nodes to access the channel with increased regularity. As a result, the variance of ToneDMAC is smaller in comparison to DMAC. To emphasize the importance of smaller variance in end-to-end delay, we now compare the two protocols under TCP traffic.

Figure 12(b) shows results when 3 transmitters transmit to a common receiver. ToneDMAC achieves higher throughput than DMAC, even for the ‘‘common receiver’’ case. This happens because the round trip time (RTT) estimation of TCP depends on the variance in end-to-end transmission latency. When DMAC is used at the MAC layer, the higher variance and recurring packet drops degrade system performance. ZeroToneDMAC, which also suffers from the same problems as DMAC, shows almost identical throughput variation as DMAC. Figure 12(c) shows the results of simulation when a larger network is used. The nodes were placed randomly in a region of $1500 \times 1500 m^2$. Random source destination pairs were chosen for 4 flows, and minimum-hop routes were assigned statically. The aggregate throughput achieved by ToneDMAC is higher than ZeroToneDMAC, which in turn is higher than DMAC, indicating that the impact of deafness is acute under TCP

traffic. The improvement can be attributed to frequent occurrences of scenarios (i) and (ii), and ToneDMAC's efficacy to alleviate deafness in both these scenarios. Also observe that when using multihop flows, the benefit from tone notifications is even more pronounced. Recall the scenario where source node S intends to communicate with destination node D , through an intermediate node I . When using DMAC or ZeroToneDMAC, while node I forwards packets to node D , node S would continue to retransmit to I , and increase its contention window upon each failure. Soon S 's contention window grows to a large value, leading to large backoff durations. Therefore, S would continue to backoff even while node I has finished forwarding the packet to D . Clearly, delay increases, which in turn affects the RTT, and thus the throughput. While we discussed an identical scenario for UDP flows, the impact is more pronounced when using TCP. When using ToneDMAC, S receives a tone from I and therefore can initiate communication sooner. As evident from Figure 12(c), performance improves over ZeroToneDMAC, and further over DMAC. In other words, ToneDMAC evidently outperforms ZeroToneDMAC and DMAC, in multi-hop ad hoc networks.

VII. FUTURE WORK

While proposing the ToneDMAC protocol, we modified IEEE 802.11 to suit a directional antenna system at the physical layer. However, it is not clear that modifying 802.11 is optimal in terms of performance. MAC protocols, designed specifically for directional antennas, may prove to be more efficient. For example, it is unclear whether CSMA/CA protocols are appropriate when using directional antennas. Time division multiple access (TDMA) schemes might prove to be more effective. Even if CSMA/CA principles are used, it is unclear whether RTS/CTS exchanges (as in 802.11) are necessary – with narrow beamwidths, bandwidth wastage due to RTS/CTSs might exceed the gains from channel reservation [7]. Directional carrier-sense is another mechanism that might not be meaningful when using directional antennas. Consider a scenario in which nodes A , B , and C are situated in a horizontal line, and B communicating to C . If A intends to communicate with C , and does not have a DNAV set towards C , it beamforms in the direction of C and carrier senses – this is done in DMAC, ToneDMAC and in many existing directional MAC protocols. Clearly, A would not sense B 's transmission to C , and might be led to believe that it can transmit an RTS to C . A collision is likely at C . Other scenarios may exist, where carrier sensing in the direction opposite to the direction of intended transmission, can be useful. Furthermore, binary exponential back-

off (BEB), as performed on packet collisions in 802.11, may not be a suitable policy when using directional antennas. BEB assumes that a CTS does not arrive from the intended receiver because of a collision. This is clearly not true in view of deafness, as discussed in this paper. We plan to take these considerations into account and explore MAC protocols, specifically designed for directional antenna systems. We intend to explore the possibilities of using tones more effectively, and analyse the effects of fading, interference, etc. This is a part of our future work.

VIII. CONCLUSION

This paper addresses deafness, an outcome of exploiting beamforming capabilities of directional antennas. A tone-based directional MAC protocol (ToneDMAC) has been proposed. The protocol uses an explicit notification mechanism to indicate the end of a dialog. Notification is implemented by transmitting a carefully chosen tone for a suitable number of time slots. Nodes, waiting to transmit, use the tones to alleviate the impacts of deafness. Simulation results indicate that under multi-hop UDP traffic, ToneDMAC performs better than DMAC. When TCP traffic is used, the performance benefits are even greater. When channel contention is high, ToneDMAC drops fewer packets in comparison to DMAC – in addition to enhancing TCP performance, this might be a desired metric for certain applications. In summary, ToneDMAC retains the benefits of beamforming while mitigating the adverse effects of deafness on MAC layer performance.

REFERENCES

- [1] R. Roy Choudhury, X. Yang, N. H. Vaidya, and R. Ramanathan, "Using directional antennas for medium access control in ad hoc networks," in *Proceedings of ACM MOBICOM*, September 2002.
- [2] C. A. Balanis, "Smart antennas for future reconfigurable wireless communication networks," Tech. Rep., Arizona State University, Annual Report, april 2000.
- [3] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. IT-46, no. 2, pp. 388–404, 2000.
- [4] M. Sanchez, T. Giles, and J. Zander, "CSMA/CA with beam forming antennas in multi-hop packet radio," in *Proceedings of Swedish Workshop on Wireless Ad Hoc Networks*, March 2001, pp. 63–69.
- [5] J. Ward and R. Crompton, "Improving the performance of slotted ALOHA packet radio networks with adaptive arrays," *IEEE Transactions on Communications*, vol. 40, no. 11, pp. 1716–1724, 1992.
- [6] Y. Ko, V. Shankarkumar, and N. H. Vaidya, "Medium access control protocols using directional antennas in ad hoc networks," in *Proceedings of IEEE INFOCOM*, March 2000, vol. 1.
- [7] R. Ramanathan, "On the performance of ad hoc networks with

- beamforming antennas,” in *Proceedings of ACM MobiHoc*, October 2001, pp. 95–105.
- [8] L. Bao and J. Garcia-Luna-Aceves, “Transmission scheduling in ad hoc networks with directional antennas,” in *Proceedings of Mobicom*, 2002.
- [9] S. Yi, Y. Pei, and S. Kalyanaraman, “On the capacity improvement of ad hoc wireless networks using directional antennas,” in *Proceedings of Mobihoc*, 2003.
- [10] T. Elbatt, T. Anderson, and B. Ryu, “Performance evaluation of multiple access protocols for ad hoc networks using directional antennas,” in *Proceedings of WCNC*, 2003.
- [11] T. Korakis, G. Jakllari, and L. Tassiulas, “A MAC protocol for full exploitation of directional antennas in ad hoc wireless networks,” in *Proceedings of Mobihoc*, 2003.
- [12] Z. Huang, C. Shen, C. Srisathapornphat, and C. Jaikaeo, “A busy-tone based directional MAC protocol for ad hoc networks,” in *Proceedings of Milcom*, 2002.
- [13] R. Roy Choudhury and N. H. Vaidya, “Impact of directional antennas on ad hoc routing,” in *Proceedings of Personal and Wireless Communication (PWC)*, September 2003.
- [14] K. Sudaesan and R. Sivakumar, “On the medium access control problem in ad-hoc networks with smart antennas,” in *Proceedings of MobiHoc*, June 2003.
- [15] IEEE 802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.
- [16] P. Karn, “MACA - a new channel access method for packet radio,” in *Proceedings of 9th Computer Networking Conference*, 1990.
- [17] A. Nasipuri, S. Ye, J. You, and R. Hiromoto, “A MAC protocol for mobile ad hoc networks using directional antennas,” in *Proceedings of IEEE Wireless Communication and Networking Conference*, September 2000, vol. 3, pp. 1214–1219.
- [18] S. Bandyopadhyay, K. Hausike, S. Horisawa, and S. Tawara, “An adaptive MAC and directional routing protocol for ad hoc wireless networks using ESPAR antenna,” in *Proceedings of ACM MobiHoc*, October 2001, pp. 243–246.
- [19] M. Takai, J. Martin, R. Bagrodia, and A. Ren, “Directional virtual carrier sensing for directional antennas in mobile ad hoc networks,” in *Proceedings of ACM MobiHoc*, June 2002, pp. 39–46.
- [20] J. Deng and Z. Haas, “Dual Busy Tone Multiple Access (DBTMA): A New Medium Access Control for Packet Radio Networks,” *IEEE Transactions on Communications*, 1998.
- [21] T. Rappaport, “Wireless communications principles and practice,” *Prentice Hall*, 2002.
- [22] T. S. Yum and K. W. Hung, “Design algorithms for multihop packet radio networks with multiple directional antennas,” *IEEE Transactions on Communications*, vol. 40, no. 11, pp. 1716–1724, 1992.
- [23] S. Wu, Y. Tseng, and J. Sheu, “Intelligent medium access for mobile ad hoc networks with busy tones and power control,” *IEEE Journal on Selected Area in Communications*, 2000.
- [24] Scalable Network Technologies, “Qualnet simulator version 2.6.1,” www.scalable-networks.com.