

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

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—This paper addresses *deafness* – a problem that appears when MAC protocols are designed using directional antennas. Briefly, *deafness* is caused when a transmitter fails to communicate to its intended receiver, because the receiver is beamformed towards a direction away from the transmitter. Existing CSMA/CA protocols rely on the assumption that *congestion* is the predominant cause of communication failure, and adopt backoff schemes to handle congestion. While this may be appropriate for omnidirectional antennas, for directional antennas, both *deafness* and *congestion* can be the reason for communication failures. An appropriate directional MAC protocol needs to classify the actual cause of failure, and react accordingly.

This paper quantifies the impact of deafness on directional medium access control, and proposes a tone-based mechanism as one way of addressing deafness. The tone-based mechanism, ToneDMAC, assumes *congestion* as the default reason for communication failures, and applies a corrective measure whenever the cause is deafness. Simulation results indicate that ToneDMAC can alleviate deafness, and perform 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. Recently, several MAC protocols have been proposed that aim to benefit from the ability to beamform in a desired direction. While higher spatial reuse and better link reliability can be achieved, tradeoffs appear. The problem of deafness is one such tradeoff. As discussed next, *deafness* is sufficiently serious to offset the advantages of beamforming, if left unaddressed. This motivates our investigation of *deafness*.

Consider Figure 1. Assume that node S intends to communicate to node D, using the route through X. When X receives a packet from S, it must beamform in the direction of D to forward the packet downstream to D. Observe that S is unaware of this communication because it does not overhear the directional signals between X and D. As a result, S initiates the next packet to X, but is unable to establish communication because X has its beam pointed towards D. Assuming congestion as the cause of failure, S backs off longer, before attempting retransmission. If data packets are large, then X remains engaged in communication to D for a long duration, during which S may attempt multiple retransmissions (each retransmission preceded by an increasing backoff duration). When X finally completes

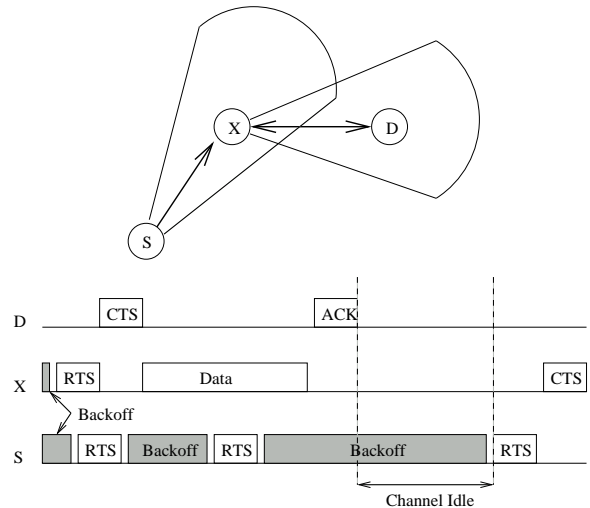


Fig. 1. A scenario illustrating the problem of deafness

packet delivery to D, and is ready to receive the next packet from S, it is highly likely that S would be counting down a large backoff counter. The channel remains idle over this entire duration – communication resumes only when the backoff period expires and S initiates the next transmission to X (as shown in Figure 1). Clearly, this is a serious wastage of channel resources, incurred for every packet communicated between X and D. Deafness does not arise when using omnidirectional antennas. Node S overhears the omnidirectional communication from X to D, and can thereby schedule its transmission to X only after X has finished forwarding the packet to D. This, however, is not possible when using directional antennas. As detailed later, the negative impacts of deafness can be cascading – leading to excessive packet drops, large delay variances, and unfairness. To higher-layer applications that rely on such metrics, deafness can be detrimental.

We propose ToneDMAC to address deafness – the main idea of our proposal can be summarized as follows. After the completion of every dialog, a communicating node transmits an out-of-band tone, omnidirectionally. Out-of-band tones act as feedback mechanisms to neighbors of the communicating node, and aids them in distinguishing deafness from congestion. When a

backlogged neighbor detects a tone, it tries to compare the originator of the tone to its intended receiver¹. If the tone-originator matches its intended receiver, then the neighbor deduces deafness as the cause of previous failures. The neighbor reduces its backoff interval, and proceeds to restart the pending transmission.

Observe that tones need not be demodulated, and therefore do not incur the large bandwidth overheads associated to demodulation and synchronization of explicit control packets. Also, tones need not be transmitted simultaneously while transmitting/receiving data packets – therefore, unlike proposals of concurrent “busy tones”, ToneDMAC relies only on a single transceiver.

The rest of the paper is organized as follows. We present related work and our assumed antenna model in Section II and III. Section IV details the impact of deafness, followed by ToneDMAC protocol description in Section V. The simulation results are presented in Section VI. We discuss some of the issues with ToneDMAC in Section VII, and conclude the paper in Section VIII.

II. RELATED WORK

The use of directional or beamforming antennas has been extensively studied in the context of single-hop communications [1],[2],[3]. Recently, attention has been focused on the possibility of using directional antennas for medium access control in multi-hop networks [4],[5],[6],[7],[8],[9],[10],[11][12][13]. While 802.11 may operate correctly in conjunction to a directional antenna, performance may get affected. As a result, many of the proposed protocols are similar to IEEE 802.11, carefully adapted for use over directional antennas. While earlier protocols proposed simple modifications, the opportunity to extract higher benefits have led newer protocols to make substantial changes to 802.11. We briefly present 802.11, followed by some of the existing directional proposals, and show how deafness is an unresolved problem.

A. IEEE 802.11 Distributed Coordinated Function (DCF)

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 duration, 802.11 chooses a random *backoff interval* from a range $[0, CW]$, and decrements the backoff counter once every idle “slot time”. When the backoff counter reaches 0, the transmitter transmits the request to send (RTS) packet, to which the intended receiver responds with a clear to send (CTS) packet. If the transmission from S collides (collision is detected by the absence of a CTS), S doubles its CW, counts down a newly chosen backoff interval, and attempts retransmission – called the binary exponential backoff (BEB) algorithm. Once the RTS/CTS handshake is successful, the transmitter transmits the DATA packet, to which the receiver

responds with an ACK. The RTS/CTS packets contain the proposed duration of DATA/ACK transmission, so that neighbors of the communicating nodes can defer their own transmissions – called *virtual carrier sensing*.

B. Directional MAC Protocols

Early proposals on directional MAC protocols [4],[14] proposed directional transmission only when possible, and relied on omnidirectional RTS/CTS packets to inform neighborhoods of imminent communication. Although novel, the observations were conservative – omnidirectional communications traded off spatial reuse for deafness, allowing the gains to be limited. Takai *et al.* described the DVCS mechanism in [15], where a node performs directional virtual carrier sensing. This is a useful optimization that increases the spatial reuse of the channel. Roy Choudhury *et al.* [16] proposed DMAC that aimed to maximize the spatial reuse, and suggested that all MAC-layer operations be performed directionally. Combined with the DVCS mechanism, DMAC maximized spatial reuse. However, several tradeoffs appeared, including deafness, that were left unaddressed.

In [9], Elbatt *et al.* propose an interesting idea – they use omnidirectional 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. However, in order to avoid interfering with ongoing communications, a node may not be able to always transmit an omnidirectional RTS/CTS packet. Transmitting directionally would however lead to deafness, which remains a problem in [9].

Bandyopadhyay *et al.* [17] propose to inform a node’s traffic pattern to its neighbors. They suggest proactive control messages that may incur heavy overhead. In addition, the protocol assumes knowledge of network traffic.

Some existing proposals have utilized “busy tones” for omnidirectional, and more recently, for directional antennas. In [18], Deng *et al.* have proposed DBTMA, that uses omnidirectional transmit and receive busy tones to avoid hidden terminal problems. Huang *et al.* [11] 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 [11], ToneDMAC uses an additional control channel for signaling. However, as mentioned earlier, tones transmitted in our protocol are different from “busy tones” in the sense that “busy tones” are transmitted in parallel with data packets, and thus requires an additional transceiver. With ToneDMAC, we use a single transceiver.

Korakis *et al.* [10] propose Circular-DMAC and 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 sequentially on every beam. As an example, if nodes in the network are equipped with 6-beam-antennas, then each data

¹Recall that tones do not contain decodable information. We present a separate scheme, based on the frequency and duration of the tone, that allows a node to deduce the tone-originator. The details of the scheme is presented later.

packet is preceded by 6 RTS and/or 6 CTS packets. As evident from simulation results, transmitting multiple RTS/CTS packets for each transmitted data packet can drastically degrade performance.




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

III. ANTENNA MODEL

We assume a switched beam antenna system. The assumed antenna model is comprised of N beam patterns ($N = 4$ in Figure 2), also called radiation patterns. An example radiation pattern is shown in the inset of 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* (imagine two passive antennas, attached to a single transceiver). A node resides in the *Omni* mode when it does not know the direction from which a signal might arrive. Once a signal is detected, the antenna begins to receive the signal with an omnidirectional gain, G^o . While the signal is being received, the antenna performs an azimuthal scan in order to select the beam on which the impinging signal power is maximum. This beam index is cached at the receiver, and may be used in the near future for communicating back to the transmitter. Observe that when using directional beams for communicating (transmitting or receiving), the main lobe gain in the direction of the beam is G^d . Typically, $G^d \geq G^o$. The Friss Equation represents how transmit and receive gains (G_T and G_R) are related to the transmit and receive powers (P_T and P_R) [2]:

$$P_R = \frac{P_T G_T G_R}{K r^\alpha}$$

The term K is a constant that accounts for atmospheric absorption, ohmic losses, etc., r is the distance between the transmitter and the receiver, and α is the path-loss index. Note that G_T and G_R are transmit and receive gains along the straight line joining the transmitter and receiver.

Observe from the equation that the maximum distance, r , over which two nodes can communicate, increases with increase in transmit and the receive gains. 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.

Relating to 802.11, a transmitter can transmit the RTS directionally, assuming that it is aware of the direction of its intended receiver. Since the receiver does not know *a priori* who the transmitter might be, it cannot receive the RTS directionally. The receiver receives the RTS omnidirectionally and performs the azimuthal scan parallelly to obtain the best beam for this transmitter. Thereafter, the CTS, Data, and ACK packets can be transmitted and received directionally. However, since we assume that the RTS must be received omnidirectionally, the distance between two communicating neighbors is bounded by $G^d \times G^o$.

IV. PROBLEM OF DEAFNESS

While exploring the space of directional MAC protocols, we gained insight into the impact of deafness. Careful analysis and simulation results indicated that deafness can be a serious problem, offsetting some of the key benefits of beamforming. While some approaches appeared to resolve deafness, the penalties they incurred deemed them impractical. This section presents our findings, and motivates our tone-based proposal of ToneD-MAC.

We consider directional MAC protocols that are motivated from the basic structure of 802.11. While deafness may not be a problem for TDMA-based protocols, we confine ourselves to the regime of 802.11-like CSMA/CA protocols only. We consider the design choices for the following protocol steps.

- *Carrier Sensing*
- *Backing Off*
- *RTS/CTS Handshake*
- *Data/ACK Exchange*

For each of the above steps, we discuss the ramifications of using the directional versus omnidirectional modes of operation. We identify how existing protocols adopt some of these design choices, and argue why they may not be sufficient to address deafness. We will refer to the scenario in Figure 3 for examples and explanations. In the figure, we will assume that node B intends to transmit data to node A.




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

- *Carrier Sensing*
- Omnidirectional carrier sensing does not seem to be appropriate, when directional antennas are available. This is because

node B (in Figure 3) would sense a transmission from X to Z, and defer its own transmission to A, even though the two transmissions are in different directions, and are not in conflict with each other. Clearly, this is unnecessary. However, if B sensed the carrier directionally, it would not sense the signal from X to Z, and could simultaneously begin communication to A. Simultaneous transmissions between B-A and X-Z improves network throughput.

- *Backing Off*

Based on the arguments on carrier sensing, it appears that a node must remain in the directional mode, while backing off as well. Remaining in the directional mode can prevent a node from getting unnecessarily captured by surrounding communications. However, an acute problem can arise. A node B, that backs off in the directional mode (and communicates subsequent data packets directionally as well) would perpetually remain in the directional mode while it has packets queued for transmission. Put differently, if each packet requires node B to carrier-sense, backoff, and communicate directionally, then multiple back-to-back packets would require B to be in the directional mode for a continued period of time. This implies that a neighbor of this node B, say X, intending to communicate to B, would experience complete deafness over this entire duration. Notice that in attempting multiple transmissions, X would also remain in the directional mode as long as it has more packets to transmit. A neighbor of X, say Y, intending to communicate to X, would similarly experience acute deafness from X. A chain is possible in which none of the nodes, except the first one, can communicate successfully – a “deadlock”. DMAC [16] and the protocol proposed in [17] propose directional backoff and fail to identify the possibility of “deadlocks”.




Fig. 4. Omnidirectional transmissions affect spatial reuse.

- *RTS/CTS Handshake*

Many variations have been proposed for exchanging the RTS/CTS packets with directional antennas. We discuss the ramifications of these variations in view of the deafness problem.

(i) *Omnidirectional RTS/CTS* was proposed by early protocols, and directional capabilities were utilized only for Data/ACK packets. Transmitting omnidirectional RTS/CTS limits the communication range of directional antennas. Node B may no longer communicate directly to node A, otherwise possible

when using directional RTS/CTS transmissions. In addition, omnidirectional RTS/CTS packets may reduce the benefits of spatial reuse. Referring to Figure 4, assume that nodes B and F perform omni RTS/CTS and begin communicating (the circles depict omni communication range). To protect this ongoing communication, Node Z must not initiate transmissions that may interfere at the receiver F. Notice that an omnidirectional communication from Z (meant for X) may interfere with F’s reception, and is therefore not appropriate.




Fig. 5. An example scenario in multi-hop ad hoc networks.

(ii) *Directional RTS/CTS* [9][11],[16],[17] allows for long-range, simultaneous transmissions, but suffer the adversities of deafness. Referring to Figure 4, if node X initiates communication to B while B is communicating to A, then X would experience deafness, as explained earlier. The cumulative impact of deafness would be clear when we discuss the issues related to directional Data/ACK packets.

(iii) *Multiple Directional RTS/CTS*: Another option could be to transmit multiple DRTS/DCTS packets (one on each beam) so that all neighbors of the transmitter and receiver are informed of the imminent dialog [10]. This strategy can eliminate deafness, while offering spatial reuse and range extension. However, the overheads can be extremely high if each communication needs to be preceded by N control packets². Circular-DMAC [10] is a proposal that adopts this design choice as a remedy to deafness, but incurs large overhead, as evident from our simulations in Section VI.

The above discussion led us to believe that single directional RTS/CTS is probably the appropriate design choice, even though deafness appears as a tradeoff to spatial reuse and range extension. ToneDMAC aims to explicitly resolve deafness, thereby retaining the benefits of reuse and extension.

- *Data/ACK Transmission*

The tradeoffs associated to RTS/CTS, are associated to Data/ACK packets as well. Omnidirectional Data/ACK transmissions reduce communication range and spatial reuse. Sacrificing these benefits will probably be inappropriate, and therefore proposals in [4],[9],[10],[14],[16], almost unanimously adopt directional Data/ACK transmissions. However, choosing to transmit the Data/ACK packets directionally leads to deafness as well.

While B and A (in Figure 5) are engaged in directional

²Observe that with 802.11b, the time to transmit RTS or CTS packets is around one-third of the total time for sending 512-Byte data packets. When number of beams, N , are of the order of six or more, overheads can be intolerable.

Data/ACK, node X may attempt multiple transmissions to B, each transmission preceded by an exponentially longer backoff duration. Now, once B finishes transmitting a packet to A, it must choose a new backoff value in preparation of transmitting the next queued packet. Since X has exponentially increased its backoff duration multiple times, it is highly probable that the backoff value B chooses for the next transmission, is smaller than X's remaining backoff duration. As a result, B can finish counting its backoff earlier than X, and proceed with the next directional Data/ACK communication. Observe that when X finally finishes counting down, B has initiated a new dialog (DRTS/DCTS/DData/DACK), causing X to fail again. After a *threshold* number of unsuccessful attempts, X drops the packet, and informs the upper layer of a link failure. Higher layer protocols will initiate elaborate mechanisms to reroute packets, and trigger congestion control algorithms. Performance can deteriorate significantly.

Consider another consequence of deafness. Assume that node X intends to send packets to A using B as an intermediate node (Figure 5). When packets are large, B would remain engaged in transmission to A for a longer duration, during which X may attempt multiple retransmissions (each transmission preceded by an exponentially increasing backoff duration). When B finally completes packet delivery to A, and is ready to receive the next packet from X, it is likely that X would be counting down a large backoff counter. The channel remains idle over this entire duration (as shown earlier in Figure 1) – communication resumes only when the backoff period expires and X initiates the next transmission to B.

V. TONEDMAC – PROTOCOL DESCRIPTION

To be able to handle deafness, a directional MAC protocol seems to require two main features.

1. The ability to communicate directionally (in the interest of higher spatial reuse)
2. The ability to inform all surrounding neighbors about a communication (in the interest of addressing deafness).

In addition, the above goals need to be achieved without a significant performance penalty.

Our proposal of ToneDMAC aims to achieve these design goals. In the beginning, a node performs directional communication (including DRTS/DCTS), and does not try to inform its neighbors of its intended communication. However, once the dialog is over, the node transmits an out-of-band tone to inform its neighbors that it had been engaged in communication over the recent past. Tones are transmitted omnidirectionally. On detecting the tone, a neighbor that had been failing to establish communication can deduce that the cause of previous failures was deafness (and not congestion). The node cancels its remaining backoff, chooses a new, relatively smaller, random backoff value, and reattempts transmission. The channel utilization increases, packet drops and transmission failure decreases, while spatial reuse and range extension is still maintained.

However, observe that tones do not contain information bits, and thus are not demodulated – they are only detected

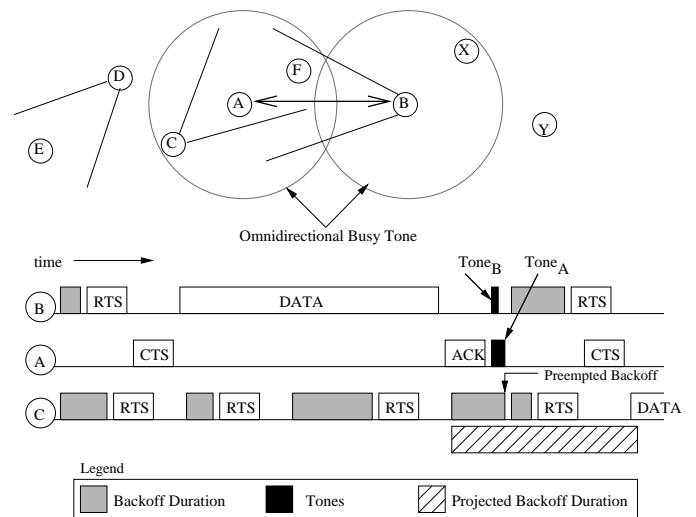


Fig. 6. An example scenario and timeline of the ToneDMAC protocol.

(through energy estimation) on the corresponding narrow frequency band. This implies that a neighbor may not be able to determine the sender of a tone, unless nodes are associated to specific tones. ToneDMAC proposes to hash node identifiers to tone-frequencies using a simple static hash function. The details of ToneDMAC are presented next.

• ToneDMAC Carrier Sensing and Backoff

When using ToneDMAC, a transmitter (B) beamforms in the direction of its intended receiver (A) and carrier senses in the directional mode. If the channel is idle, the transmitter selects a backoff duration from an interval of $[0, CW_{min}]$ (similar to 802.11). Now, ToneDMAC requires B to switch the antenna to the omnidirectional mode while performing the backoff countdown. However, a node does not cancel its backoff on sensing any arbitrary signal. Instead, when the node senses a signal, it performs the azimuthal beam scan to determine the direction of the arriving signal. If this direction is in a different direction from that of A, then B continues with the countdown; otherwise B cancels the timer until the channel becomes idle. Observe that if the arriving signal is strong enough for B to receive it, then B cancels the backoff timer, and decodes the packet. If the packet is meant for B, then B can reply to the packet, alleviating the possibility of deadlock and deafness.

• ToneDMAC DRTS/DCTS/DData/DACK

The RTS/CTS handshake is carried out directionally when using ToneDMAC. Once the handshake is established successfully, Data and ACK packets are transmitted directionally as well. Nodes that overhear the DRTS, DCTS, or DData packets, update their records suitably to avoid interfering with communications in those directions. Both B and A switch back to the omnidirectional mode once the DData/DACK exchange has been accomplished.

• ToneDMAC Tone Transmission

In ToneDMAC, the 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. Once DData/DACK is over, ToneDMAC requires B and A to transmit tones omnidirectionally. The frequency of the tone, and the duration for which the tone is transmitted depends on the transmitter of the tone. A node i chooses a tone of frequency f_i from a set of K frequencies, and an integer time duration t_i from an interval $[t_{min}, T]$. The tuple (f_i, t_i) serves as the signature of node i . This allows a neighbor of node i , that detects the tone, to deduce the transmitter of the tone with some degree of accuracy. The values of f_i and t_i are static hash functions of the node's identifier. We assume that a higher layer is capable of assigning consecutive identifiers to nodes, so that the tuple (f_i, t_i) can be uniformly distributed over the node identifier space. A simple hash function is used to assign (f_i, t_i) to a node i .

$$f_i = (i \bmod K) + 1$$

$$t_i = (i \bmod (T - 1)) + 2$$

Nodes B and A transmit their corresponding tones omnidirectionally. Since transmission in the omnidirectional mode may not reach all the directional neighbors, the transmit power of the tones are suitably increased to equal the range of directional transmissions. This ensures that nodes that may potentially be encountering deafness from B and A, are within transmission range of the tones.

- *ToneDMAC Tone Reception*

A neighbor of A, say C in Figure 6, that has been not been able to communicate to A because of A's communication to B, would receive the tone transmitted by A. To be able to conclude deafness as the cause of communication failure, C must first check whether the originator of the tone matches its intended receiver. Therefore C applies the static hash function on its intended receiver, calculates the (f, t) signature, and compares it with the frequency and duration of the received tone. If the comparison is positive, C concludes that the reason for communication failure was indeed deafness. In order to avoid backing off unnecessarily, it preempts its current backoff timer, reselects a new smaller backoff from $[0, CW_{min}]$ (as in 802.11), and reattempts transmission. Figure 6 shows the sequence of events where C successfully initiates transmission before B can initiate the second DRTS to A. Observe in the figure that, without the tones, C's projected backoff could be longer than B's second backoff, and therefore C would experienced deafness repeatedly.

- *Tone Aliasing*

Observe that the hashing mechanism proposed for ToneDMAC does not guarantee that the transmitter of the tone will always be deduced accurately. Aliasing occurs when two neighbors of a node are hashed to the same (f, t) signature (e.g., $f_D = f_A$, and $t_D = t_A$). A node like C, that has both A and D as its neighbors, will not be able to deduce the tone-sender correctly. We alleviate this problem by utilizing the direction in which the tones are received. By azimuthal beam scanning, C determines the direction in which the tone arrives. In cases where the aliased nodes are on different beams, the direction of the tones would be sufficient to deduce the tone-transmitter. For example, in Figure 6, if

C determines that the tone arrived from north-east, then assuming that C knows the direction of its neighbors, C would be able to deduce that the tone came from A and not from D. Of course, if both A and D were located on the same beam, accurate classification will not be possible.

Another case that can cause incorrect deduction is when two nodes, hashed to the same tone-frequency, transmit tones that overlap in time. For example, assume that $f_D = f_A = f$, and $t_D \neq t_A$. It is possible that C would receive a tone at frequency f for a duration, τ , such that $\tau > t_D + t_A$. In such cases, incorrect deduction occurs.

Although incorrect deduction is possible, we argue that the consequences may not be detrimental. Deafness may arise when two neighbors of a node finish their dialogs at nearly the same time – the tones would therefore overlap. However, because of the randomization in backoff counters, the subsequent communications of the same neighbors will probably get separated in time, preventing repeated overlap of tones. Moreover, an incorrect deduction would result in a node not resetting its backoff counter when necessary, or the vice versa. This may occasionally cause slight wastage of channel resources. Simulation results have shown that under reasonable node density, with $K = 4$, $T = 4$, and $beamwidth = 60^\circ$, inaccurate deductions is not a serious problem.

VI. PERFORMANCE EVALUATION

We use the Qualnet simulator [19], version 3.1, for simulating ToneDMAC and comparing its performance with DMAC [16] and circular-DMAC [10]. We consider several scenarios and systematically analyze the impact of deafness on directional medium access control. The transmit power is assigned such that communication range of directional transmissions is approximately 300 meters. The assumed data rate is 11 Mbps. Due to the allocation of bandwidth to the control channel (in ToneDMAC), available data rate for data communication is 10.5 Mbps. We used both UDP and TCP traffic in our experiments, with packet sizes of 512 Bytes. Sources are always backlogged unless mentioned otherwise. We do not consider mobility.

We begin by using three simple scenarios in order to understand the individual benefits of each design choice made in ToneDMAC. The scenarios are shown in Figure 7 – dotted lines indicate flows and solid lines indicate links.

- *Scenario (i)*

Scenario (i) shows a two-hop flow between node 1 to 3 via node 2. As discussed earlier, node 1 experiences deafness when node 2 forwards the packet to node 3, which in turn results in long durations of an idle channel. Figure 8 shows the variation of end to end throughput for scenario (i). Note that scenario (i) measures only the losses suffered from deafness – we do not expect potential benefits of directional communication because there is no scope for spatial reuse. Therefore, IEEE 802.11 (simulated with omnidirectional antennas) performs better than all other directional protocols because of its ability to completely avoid deafness. ToneDMAC's throughput is slightly lower than

802.11 because of the control channel bandwidth allocated to out-of-band tones.




Fig. 7. Scenarios affected by deafness.

Between the directional protocols, ToneDMAC outperforms the others, except at low sending rates, where the performances are comparable. This is because multi-hop deafness does not arise at low sending rates – i.e., node 2 completes forwarding a packet to node 3 by the time node 1 has generated the next packet. At higher loads, the improvement of ToneDMAC over the other protocols is pronounced. Although Circular-DMAC alleviates deafness, the heavy performance penalty associated to multiple RTS packets, degrades its performance.




Fig. 8. Throughput comparison for scenario (i).

- *Scenario (ii)*

Scenario (ii) is designed to measure the cascading effects of deafness in terms of packet losses, delay variances, etc. Both nodes 1 and 3 intend to communicate to 2. Now when one of the nodes (say node 3) initiates successful communication, the other would experience deafness. When 3 finishes its first transmission, 1 would probably be counting down a large backoff interval. For its next transmission, 3 would choose a small backoff value, which has a high probability of being smaller than 1's remaining backoff duration. As a result 3 initiates the second packet successfully, and when 1 finishes counting down its backoff, it will experience deafness again. This can continue many times until 3 gets lucky enough to steal channel access from 1 (possibly after dropping many packets). Observe that once 3 gets successful, 1 suffers from sustained deafness. The alternation continues – the winning node transmits multiple packets once it grabs the channel, while the losing node drops multiple packets. Figure 9 shows the variation of backoff windows (and successful transmissions) with DMAC and ToneDMAC.

Clearly, backoff values remain low, and channel access fairness is reasonably maintained with ToneDMAC.

Although the UDP throughput for scenario (ii) is comparable for ToneDMAC and DMAC³, the high rate of packet drops can cause higher layers to suffer. The network layer may interpret a link failure and unnecessarily trigger route-repair algorithms, all of which incur overhead. We have not investigated the adverse effects of deafness on higher layer protocols, and it is a part of our future work. Figure 10 shows the rate of packet drops for scenario (ii) – clearly, packet drops are lower when using ToneDMAC.

Due to the alternation of channel access, delays experienced by nodes 1 and 3 can vary significantly in scenario (ii). A node that loses channel access to another node, might have to wait for a long duration before it can re-establish successful communication – a case of short-termed unfairness. Although the “average” delay over all the packets is comparable, “variance” in packet delays is often a metric of interest (to transport protocols like TCP, for example). Figure 11(a) compares the per-packet delay experienced by node 1 when using DMAC and ToneDMAC. While DMAC's delay is characterized by large fluctuations, ToneDMAC maintains a smaller range of per-packet delay. This reduces the “variance”, which in turn improves TCP throughput, as evident from Figure 11(b).

- *Scenario (iii)*

Scenario (iii) aims to evaluate the “deadlock” problem discussed earlier in Section IV. Recall that “deadlock” occurs if the directional MAC protocol chooses to backoff in the directional mode, as proposed in DMAC [16]. As explained earlier, all nodes on a chain may fail to communicate except for the first node on the chain (node 3 in this case). However, when using ToneDMAC, nodes would apply the combination of *omnidirectional backoff* and *out-of-band tones* to handle deafness. To observe the individual benefits from each of these design decisions, we have simulated a variation of ToneDMAC, called ZeroToneDMAC, that does not use tones at all. The performance of ZeroToneDMAC will reflect the benefits of *omnidirectional backoff* only. Figure 11(c) compares the percentage of aggregate packet drops for DMAC, ToneDMAC, and ZeroToneDMAC. Due to “deadlock”, DMAC experiences severe rate of packet drops at nodes 1 and 2. ZeroToneDMAC performs significantly better than DMAC. This is because nodes backoff in the omnidirectional mode, and therefore simultaneous communication between links 1-2 and 3-4 can progress simultaneously. However, ToneDMAC outperforms both of the other protocols since in addition to allowing simultaneous communication, channel wastage between two consecutive transmissions are also reduced using the out-of-band tones. Fairness among flows were also observed to be higher when using ToneDMAC.

- *Multi-hop Networks*

With reasonable understanding of how deafness impacts the

³This is because the common receiver remains occupied in communicating to either of the transmitters for both DMAC and ToneDMAC. With Circular-DMAC, the throughput is however much lower, and is therefore omitted for this comparison.

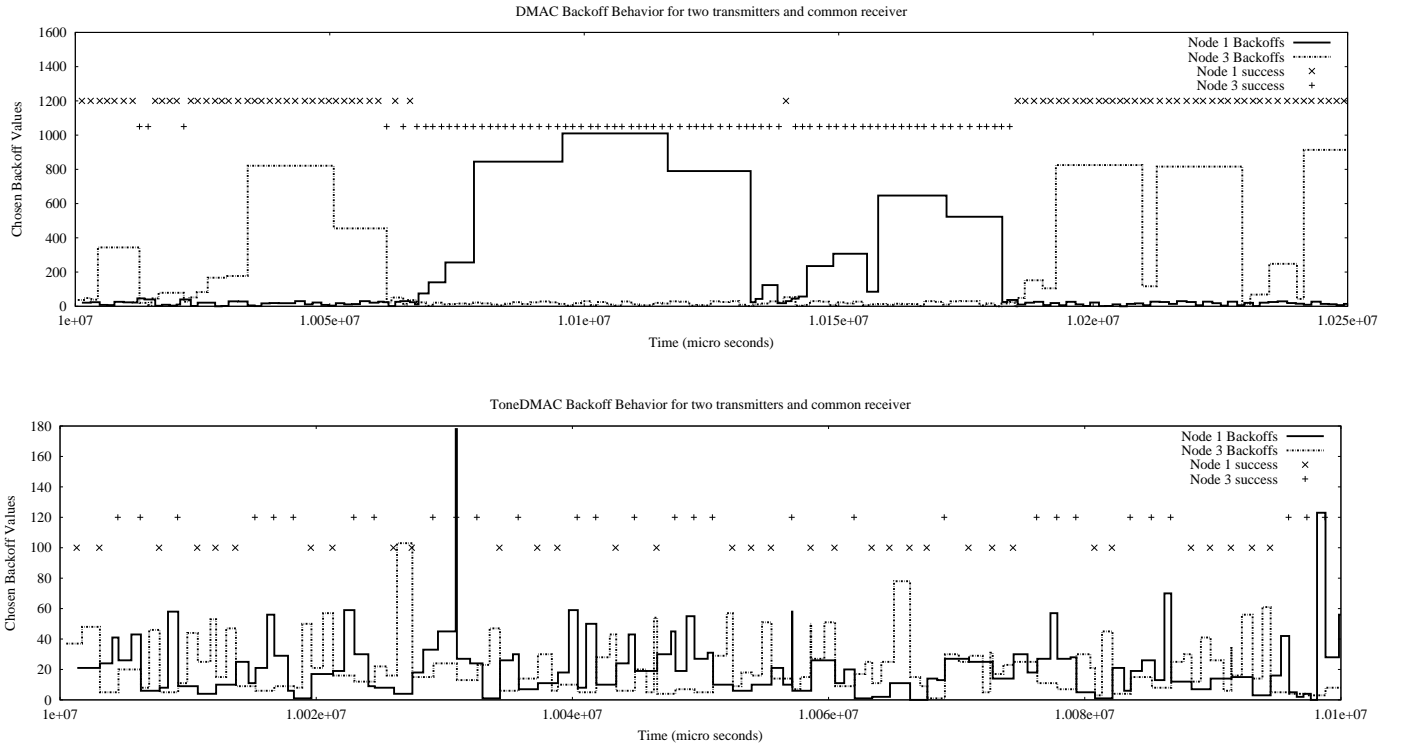


Fig. 9. Behavior of the backoff values against time, for scenario (ii). The backoff values are plotted using a step function. For example, if backoff values B_1 and B_2 are chosen consecutively at times t_1 and t_2 , then a horizontal line at a height of B_1 extends from time t_1 to t_2 , followed by a jump to a height B_2 , at time t_2 . Successful packet transmissions have also been marked. For visual clarity, the time durations over which the graphs are drawn, are different in length.




Fig. 10. Packet drop due to deafness in DMAC and ToneDMAC for scenario (i), using UDP traffic

simple scenarios, we simulated larger multihop networks, for various traffic conditions. 30 nodes were placed randomly in a region of $1500 \times 1500 m^2$. 5 flows were set up simultaneously between randomly chosen source and destination node-pairs. The routes were statically assigned. We used $K = 4$ and $T = 4$ (i.e., 4 tone-frequencies and maximum 4 slots for transmission) for ToneDMAC. The simulation results are averaged over 25 runs. As described earlier, we include results for ZeroToneDMAC in order to discern the benefits of *tones* from *omnidirectional backing off*.

Figure 12 shows the performance of ToneDMAC when UDP traffic is used. Evident from the graphs, ToneDMAC out-

performs both DMAC, ZeroToneDMAC, and Circular-DMAC mainly because of its ability to reduce channel-idle time and resolve deadlock. DMAC performs worse than ZeroToneDMAC and ToneDMAC because of its inability to address either of the above issues. The benefit of omnidirectional backing off is evident from the performance improvement of ZeroToneDMAC over DMAC. The benefit of tones, to reduce idle-channel time, is evident from the improvement of ToneDMAC over ZeroToneDMAC. Circular-DMAC achieves the least throughput not only because of the multiple RTS transmissions, but also because of deafness appears due to the transmission of a single directional CTS packet (as proposed in [10], to reduce overheads).




Fig. 11. Evaluation of deafness on simple scenarios of Figure 7

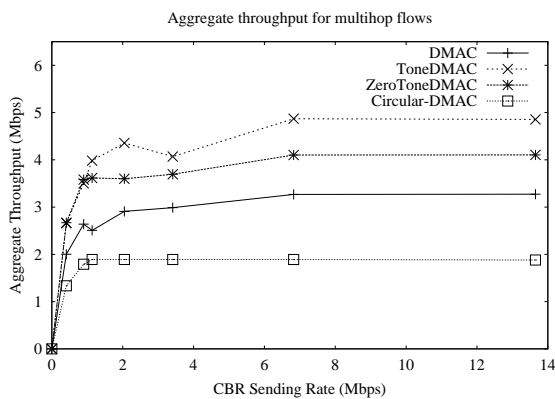


Fig. 12. Throughput comparison for multi-hop flows in a random network.

ToneDMAC	ZeroTone	DMAC	C-DMAC
2.64	1.82	1.41	1.23

TABLE I
AGGREGATE TCP MULTIHOP THROUGHPUT

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 undelivered route-discovery packets, affects end-to-end throughput and delay [12]. 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 a part of our future work. We believe that when a suitable routing protocol is executed, the performance gains of ToneDMAC would be even more pronounced.

While evaluating scenario (ii) in Figure 7, we observed large fluctuations in delay for DMAC transmissions. To evaluate the impact of high fluctuations (causing higher variance), we simulated flows with TCP traffic. As shown in Table I, the aggregate throughput achieved by ToneDMAC is higher than ZeroToneDMAC, which in turn is higher than DMAC and Circular-DMAC. This happens because the round trip time (RTT) estimation of TCP depends on the variance in end-to-end transmission latency.

When using DMAC, the higher variance and recurring packet-drops degrade system performance. Tone-based feedbacks reasonably limit delay-fluctuations, in addition to increasing channel utilization. The net benefits are encouraging.

VII. DISCUSSION

Most directional MAC protocols, including ToneDMAC, have adopted the CSMA/CA principles of IEEE 802.11. However, it is unclear that modifying 802.11 is the optimal approach to designing directional protocols. While binary exponential backoff (BEB) may be suitable for omnidirectional antennas, such conservative backing off may not be the best choice with directional antennas. We plan to take these considerations into account and explore new MAC protocols using directional antennas.

An issue with ToneDMAC pertains to successful tone detection, especially in the presence of fading and multipath. Since tones are transmitted on a narrow bandwidth channel, and the durations of transmissions are short, tones may be detected partially, or may not be detected at all. Multipath effects can also cause a tone to arrive from a different direction than the known direction of a neighbor. Clearly, both these effects can cause a node to misclassify the cause of transmission failure. However, notice that the consequences of misclassification may not be severe – a node may initiate its RTS too early or transmit its RTS too late. Performance penalties associated to these consequences may not be severe, unless misclassification is frequent.

Omnidirectional backing off, as opposed to directional, has a subtle tradeoff. While in the omnidirectional mode, a node, S , may not be able to overhear an RTS/CTS packet from a distant communication. If the node later beamforms in the direction of this communication, and initiates an RTS, then in some cases collisions may occur. However, such cases are rare since the receiver of the distant communication is not only outside the omnidirectional range of S , but also points its beam in the direction of its transmitter. As a result, the signal to interference plus noise ratio (SINR) at the receiver is likely to remain sufficiently high for successful reception.

With ToneDMAC, we increase the power of tones to equalize the communication range to that of a directional transmission. The increase in power is calculated from the difference in gain of the directional mainlobe and the omnidirectional beam. Even though tones are sent over a small duration, they can incur penalties in terms of power consumption at the node. We have not addressed the issues of power consumption in this paper.

VIII. CONCLUSION

We investigated the problem of deafness, appearing when MAC protocols are designed for directional antennas. We evaluated several cases in which deafness can lead to channel under-utilization. Moreover, we identified scenarios in which the impact can be cascading, leading to packet drops, delay-fluctuations, unfairness, etc. Our proposal of ToneDMAC uses a tone-based notification mechanism that allows neighbors of a node to classify congestion from deafness, and react appropriately. Simulation results show that ToneDMAC can be effective in mitigating the adverse effects of deafness, while retaining the benefits of directional beamforming.

REFERENCES

- [1] C. A. Balanis, "Smart antennas for future reconfigurable wireless communication networks," Tech. Rep., Arizona State University, Annual Report, april 2000.
- [2] T. Rappaport, "Wireless communications principles and practice," *Prentice Hall*, 2002.
- [3] 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.
- [4] 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.
- [5] R. Ramanathan, "On the performance of ad hoc networks with beamforming antennas," in *Proceedings of ACM MobiHoc*, October 2001, pp. 95–105.
- [6] L. Bao and J. Garcia-Luna-Aceves, "Transmission scheduling in ad hoc networks with directional antennas," in *Proceedings of Mobicom*, 2002.
- [7] S. Yi, Y. Pei, and S. Kalyanaraman, "On the capacity improvement of ad hoc wireless networks using directional antennas," in *Proceedings of Mobihoc*, 2003.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] K. Sudaesan and R. Sivakumar, "On the medium access control problem in ad-hoc networks with smart antennas," in *Proceedings of MobiHoc*, June 2003.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] Scalable Network Technologies, "Qualnet simulator version 2.6.1," www.scalable-networks.com.