MAC-Layer Capture: A Problem in Wireless Mesh Networks using Beamforming Antennas

Romit Roy Choudhury
Dept. of ECE,
Duke University

Nitin Vaidya
Dept. of ECE,
University of Illinois

Abstract— **Beamforming antennas have been shown to improve spatial reuse in wireless networks. Protocols that aim to exploit beamforming antennas have leveraged benefits from directional transmission as well as directional reception. We argue that exploiting antenna characteristics for transmission and reception alone, is not sufficient. Controlling the antenna during the idle state of a node is necessary to attain further improvements. A node that does not control its antenna during the idle state will waste time in receiving packets not intended for it. These packets will eventually be dropped at the MAC layer, reducing channel utilization. With suitable beamforming, a node might be able to utilize this time for concurrent communication. We call this the problem of *MAC-layer Capture* in view of a node getting engaged in receiving unproductive packets. We present MAC and routing protocols to address MAC-layer capture. Simulation results show that improvements from avoiding MAC-layer capture can be substantial in static mesh networks.**

I. INTRODUCTION

In CSMA-based protocols (such as IEEE 802.11) the PHY layer does not know which packets are addressed to itself. The notion of "a packet's destination" is meaningful only at MAC and higher layers. Therefore, the PHY layer will attempt to decode every packet that it can, and pass them on to the higher layers. As a result, nodes may waste time in decoding unnecessary packets at the PHY layer, and later dropping them at the MAC layer. We call this the problem of *MAC-layer Capture* because the MAC layer remains captured in receiving unproductive packets. When using omnidirectional antennas, there is little incentive to select which packets should (or should not) be received. This is because, even if a node determines that decoding some packet is unnecessary, it cannot initiate a concurrent communication while that packet transmission is in progress. However, with suitable beamforming, the PHY layer can be made to avoid unnecessary packet reception, and the "free" time can be invested in useful, concurrent communication. To the best of our knowledge, existing protocols have not exploited the possibility of selectively receiving productive packets (and filtering out the unproductive ones). This is because existing protocols do not control the antenna during the idle state, as argued next.

Observe that transmission and reception can be associated to specific beam patterns. In contrast, there is no particular beam associated to the idle state of a node. Since a packet can arrive from any direction, most beamforming protocols require the idle PHY layer to listen in all directions (i.e., omnidirectional mode). As a result, the PHY layer gets "distracted" by every packet that originates in its vicinity. This can reduce spatial reuse, because the PHY layer might waste substantial time in receiving unnecessary packets. The following example illustrates the point.

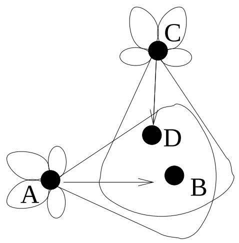

Fig. 1. Capture prevents parallel dialog between A-B and C-D

In Figure 1, assume that node A intends to transmit to node B, and node C to node D. Consider the point of time when all nodes are idle (and thus in the omnidirectional mode), and both A and C are preparing to initiate communication to their respective destinations. If A initiates transmission first, observe that both B and D will beamform toward A, and receive the entire packet. Concurrent communication between the two node pairs will not be possible, leading to underutilization. A similar situation arises even if C initiates the transmission first.

We propose to address capture at the MAC and network layer. The capture-aware directional MAC (CaD-MAC) protocol prescribes suitable beam-patterns (during the idle state) to filter out unproductive packets. The

Capture-aware Routing Protocol (CaRP) selects routes that further reduces capture between cross traffic. Together, CaDMAC and CaRP complement each other, enhancing network performance.

II. ASSUMPTIONS

Network Scenario: We envision stationary wireless mesh networks. We assume that the network nodes (access points) are mounted on rooftops or lamp-posts, and together form a multihop backbone. Each access point aggregates traffic from wireless clients associated to it, and routes the traffic to (and from) mesh gateways. We assume that the traffic between clients and access points is on an orthogonal frequency channel, and does not interfere with traffic on the backbone mesh network. Such an architecture has been considered in [1, 2]. We also assume that as a result of aggregation at access points (AP), the traffic is relatively persistent (i.e., routes do not vary in the time scale of few packets). Under this framework, this paper demonstrates the improvements possible when beamforming antennas are used by the APs.

Antenna Model: We use an antenna system that comprises of N beams, covering the (360°) azimuth plane. We assume that radiation pattern of the antenna can form multiple beams, where the number of beams can be controlled through signal processing or multi-mode switching. The maximum number of beams at any given time is limited by the number of elements in the antenna array. At any given time, the antenna can either transmit or receive a single packet. An abstract block diagram of a multibeam switched antenna is shown in Figure 2 – observe that subsets of the beams can be turned on/off using the switches, creating beam patterns like the ones shown in Figure 3. Of course, this is an overly simplistic representation, used for the convenience of explanation. Antenna patterns used in our simulation are taken from realistic antennas with sidelobes (not shown in the figure).

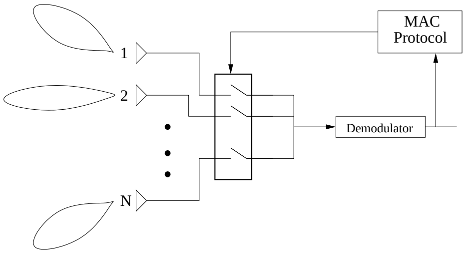

Fig. 2. Switched combining with multi-beam antennas

We assume that the antenna system has the ability to detect the beam-of-arrival of a signal. Briefly, when a signal arrives, the output from each of the beams can be compared to detect the beam on which the received signal strength was maximum. This information is made available to the higher layers, along with the received packet (details of direction-of-arrival estimation algorithms can be found in [3]).

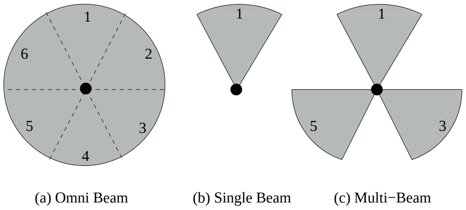

(a) Omni Beam (b) Single Beam (c) Multi−Beam

Fig. 3. An abstract antenna model

Multipath: The wireless channel is characterized by multipath propagation characteristics. In view of this, an antenna pattern that maximizes signal to interference ratio (SIR) may not have its beam pointed in the line-of-sight (LoS) of the transmitter/receiver. Moreover, the beam may change over time due to the time-varying nature of the channel. In this paper, we do not assume LoS communication. However, we assume that the channel variation is not so drastic, that an established beam must change in the time scale of few packets.

III. RELATED WORK

Early work on directional MAC protocols improved spatial reuse by enabling nearby communications to progress in parallel [4, 5]. Further investigation introduced the notion of directional virtual carrier sensing through the directional NAV (DNAV) mechanism [6]. However, several tradeoffs were also identified, including new hidden terminal problems, and deafness [7], that limit the benefits of beamforming. Subsequent protocols addressed some of these problems and raised the limits on achievable performance [8–11]. However, further improvements can still be obtained by avoiding MAC-layer capture, and receiving only those packets that are productive.

The impact of beamforming antennas on multihop routing has also emerged to be of interest [12], [13]. In [14], the authors considered the possibility of *zone disjoint routing*, wherein, routes are chosen such that they are spatially disjoint. A route is preferred if it passes through a zone in which there is less cross-traffic. However, when two flows are forced to pass through the same zone, they are unable to uphold concurrent communication. Our approach in this paper is motivated from the observation that nearby flows in a network should be able to communicate in parallel, as long as their directions of communications are non-overlapping. Our routing protocol chooses routes that are maximally non-overlapping in direction. As will be clear from Section IV and V, our protocol encompasses the benefits available from *zone disjoint routing*, and enhances it with capture-awareness.

A. Motivation and Intuition

In the absence of scheduled communication, existing directional MAC protocols require an idle node to be in the omnidirectional mode [4, 6, 7, 9, 15]. Once a signal impinges on the receiver, different protocols assume different techniques to beamform toward the signal. Once beamformed, the rest of the dialog is typically accomplished in the beamformed mode. When the dialog is over, the receiver returns to the omnidirectional mode, waiting for the next packet. The waiting time is not small because the transmitter will ensure a free channel, backoff, and then initiate the next dialog. During this entire duration, *the idle receiver is susceptible to get distracted by every unproductive signal that impinges on its transceiver*, leading to frequent occurrences of capture.

This observation motivates protocols that are "capture-aware". If a protocol can estimate directions[1] of unproductive traffic, it might be able to employ an appropriate beam-pattern *a priori* (like a spatial filter) that filters out unproductive signals. Nodes may now be free to participate in productive communication while the spatial filter is active. To respect the possibility that productive signals may arrive from the filtered-out directions, a node must periodically refresh the beam pattern, based on changes in local traffic patterns. To further reduce the effect of capture, route directions can be carefully chosen such that it is less captured by network cross-traffic. We detail these ideas in Sections IV and V.

IV. CAPTURE-AWARE DIRECTIONAL MAC (CADMAC)

We divide time into recurring cycles of duration T_{cycle} as shown in Figure 4. Each cycle is further subdivided into an *ON-Duration* and an *OFF-Duration*, of lengths τ_{on} and τ_{off}, respectively. For our proposed protocol, we assume that nodes are clock-synchronized (we show later that synchronization need not be tight, and can be within the limits of current accuracy levels). Nodes enter the ON and OFF durations synchronously, and the cycle repeats. In the *ON-Duration*, all nodes use the omnidirectional mode when they are idle. In the *OFF-Duration*, each node uses a suitably chosen multi-beam pattern when it is idle.

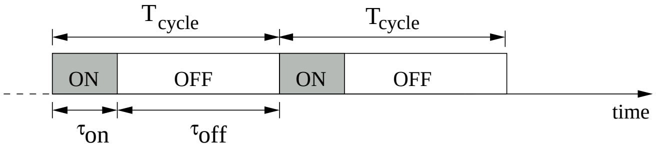

Fig. 4. Time divided into cycles of ON and OFF-Durations

During the ON-Duration, the MAC layer at each node records every received packet, and the beam used to receive it. The packets are then segregated into 2 categories – (i) productive traffic and (ii) capture traffic. Productive traffic includes unicast packets meant for that node, as well as broadcast packets. Capture traffic includes overheard packets that are intended for other nodes. If a beam proves to be the receiver of only capture traffic, then the beam is "black-listed". At the end of the ON-Duration, the MAC layer decides to turn off all black-listed beams for the next OFF-Duration. The resulting beam-pattern used by node i, during the ensuing OFF-Duration, is denoted by β^i_{off}. The beam β^i_{off} can be viewed as a spatial filter that makes node i less susceptible to capture during OFF-Durations. The example in Figure 5 illustrates this simple scheme.

In Figure 5, assume that node A communicates with node C via node B, while node E communicates with node G via node F. Nodes identify the unproductive beams during the ON-duration, and turn them off during the next OFF-Duration. As shown in Figure 5(a), node B turns off its north and south pointing beams, while node F turns off the east and west pointing beams. As a result, dialogs between nodes A, B, or C, do not capture F, and dialogs between nodes E, F, and G, do not capture B. Spatial reuse increases in this simple scenario. Observe that the beam patterns shown in Figure 5(a) will have sidelobes and backlobes. However, parallel communications will still be achievable, provided F and B are not too close to interfere each other significantly, along the direction of low-gain side lobes.

The above benefit of spatial reuse in CaDMAC arises from a node's ability to use beam-pattern β_{off} while idle. However, using the same beam-pattern for transmission/reception will be inefficient. Therefore, with

---

[1] The term "direction" is used loosely, and it does not imply the direction of line-of-sight.

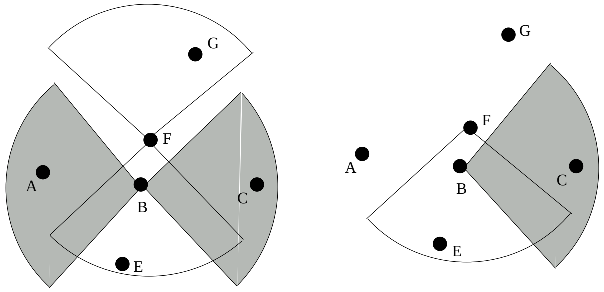

(a) Capture−aware beam patterns of nodes B and F during idle state

(b) Nodes B and F transmitting using a single beam

Fig. 5. CaDMAC enables parallel communication by avoiding mutual capture between nearby nodes.

CaDMAC, a node always uses (both in ON and OFF durations) a suitably chosen single-beam to transmit or receive packets. The single beams used by nodes B and F, to transmit to C and E respectively, are shown in Figure 5(b). Once a dialog is over, an idle node i returns to the beam-pattern β^i_{off} (pre-selected for that ongoing OFF-Duration), and waits for the next dialog to begin. Using a single-beam for transmission and reception increases the spatial reuse in the network, as advocated in most of the existing directional MAC protocols [6–8]. CaDMAC achieves these benefits in addition to the benefits of avoiding capture during the OFF-Durations.

Directional virtual carrier sensing (DVCS) is a useful mechanism that requires a node to refrain from initiating transmissions toward a direction from which it has overheard an RTS, CTS, or a Data packet [6]. However, observe that the DVCS mechanism is not designed to address the problem of capture, because, even while implementing DVCS, a node will get captured by the RTS packet in the first place. More importantly, a node that overhears an RTS from a specific direction does not attempt to filter out the unwanted data packet that is likely to arrive from the same direction in the immediate future. When the unwanted data packet arrives, the node gets captured again, and thereby cannot schedule its own useful communication.

We enhance the DVCS mechanism to make it capture-aware. When a node, X, overhears a RTS or a CTS packet on a particular beam, CaDMAC recommends the physical layer to turn off that beam for the proposed duration of the imminent dialog. As a result, X can remain idle over the duration of the dialog, and can potentially initiate (or respond to) a new communication. Since the duration of dialogs can be long, the benefits from avoiding capture at such time scales can be non-marginal.

A. *Early Destination Detection (EDD): An Optimization*

CaDMAC avoids capture in two time-scales; (i) at the time-scale of OFF-Durations, by choosing suitable beam patterns, and (ii) at the time-scale of "RTS/CTS/DATA/ACK" dialogs, by extending the DVCS mechanism. However, capture is still possible at the time-scale of single packets, e.g., a node may get captured by a single RTS or CTS packet. More importantly, when RTS/CTS packets are not used, a node may be captured by the entire DATA packet. To minimize capture on the time-scale of single packets, we present an optimization that we call *Early Destination Detection (EDD)*. The main idea in EDD is to decode the destination address of the packet (before the packet's payload is received), and turn off appropriate beams if the packet is addressed to a different node. Quicker the *destination detection*, higher is the benefit (especially when packets are large). For this optimization to be effective, we need to ensure 2 conditions. (1) A packet's destination address is included within the initial bits of the packet header[2]. (2) The antenna system will need to have a small beamforming delay to ensure that it can remain capture-free for a large fraction of the packet transmission time. With inexpensive antennas, such time-scales may not be achievable, and hence, EDD may not be applicable in such cases.

B. *Routing over CaDMAC*

When using CaDMAC, nodes turn off a subset of beams during the OFF-Duration, resulting in a less connected topology. Therefore, a route discovery initiated during the OFF-Duration can result in the discovery of sub-optimal routes. We propose simple modifications to the routing protocol to address this issue. The routing protocol could be made aware of the ON and OFF durations. Route discoveries initiated during the ON-Duration require no modification. However, if the route-discovery is initiated during the OFF-Duration, then the routing protocol sets a flag to remember this action. If a route is discovered during the OFF-Duration, the routing protocol begins sending packets on this route (even though the route might be sub-optimal). If a route is not discovered, it buffers the arriving packets and waits for the next ON-Duration. If the flag is set at the beginning of an ON-Duration, it re-initiates route discovery to find a better route, and redirects traffic over it. Of course, latency of finding the optimal route increases on average. However, the gains from avoiding capture over subsequent OFF-Durations can compensate adequately for the

---

[2] Support for such capabilities is becoming available in a transceiver prototype from Lucent [16].

losses from initial sub-optimality. Moreover, if the network is dense, the sub-optimal routes may not be too bad in comparison to the optimal. To validate this, we evaluate the impact of ON and OFF durations on network performance in Section VI.

A given routing protocol can be executed over CaDMAC with the above modifications. However, we next illustrate that incorporating capture-awareness in the routing protocol can further improve performance. We begin with a simple motivating example.

V. CAPTURE-AWARE ROUTING

In Figure 6, assume that node B routes packets to node A. While such a flow exists, assume that node S intends to discover a route to node G. Routing protocols that use minimum hop-count metrics for selecting a route, may very well choose the route $\{$S-M-F-G$\}$. When using route $\{$S-M-F-G$\}$, observe that CaDMAC at node M must ensure that its east-pointing beam is turned on, in order to communicate with F. However, this east-pointing beam at M is susceptible to capture by transmissions from node B to A. As a consequence, routes $\{$B-A$\}$ and $\{$S-M-F-G$\}$ will be forced to share the available channel capacity. However, choosing route $\{$S-M-E-G$\}$ proves to be beneficial because transmissions on link M-E and B-A are free of mutual capture. *Intuitively, in the presence of more flows, more routes can traverse through this region without having to share the channel between them, as long as there is an available direction that is capture-free from others.* Our capture-aware routing metric is designed to pick out all such capture-free directions (whenever possible), and route traffic along these directions.

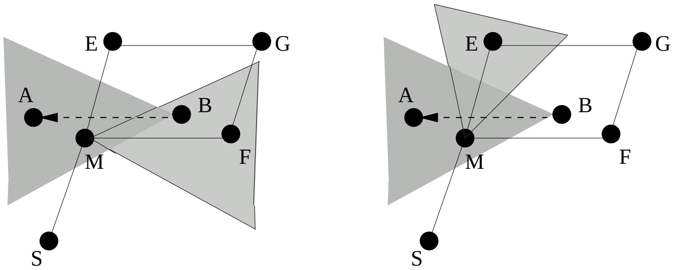

Fig. 6. Communication on link A-B captures link M-F. Routes that avoid using link M-F (and chooses link M-E instead) extracts greater spatial reuse from the channel

A. Measuring Route Cost

Our goal is to associate each route with a cost that indicates the "amount of capture" on that route. To achieve this, we determine the capture-cost of each link on the route. Since each link can be represented by a pair of directional beams, quantifying the *amount of capture* on each beam can determine the capture-cost of the end-to-end route.

To be precise, consider an end to end route, R, between nodes S and D. Let nodes i and j be two adjacent nodes on this route. Let B_{ij} be the beam that i uses to communicate to j, and let B_{ji} be the beam that j uses to communicate to i. Let $C_{B_{ij}}$ represent the "amount of capture" on beam B_{ij}. Then, the capture cost of link ij, denoted as κ_{ij}, can be calculated as

$$\kappa_{ij} = C_{B_{ij}} + C_{B_{ji}} \tag{1}$$

Consequently, the capture-cost of route R, denoted by κ_R, can be calculated as

$$\kappa_R = \sum_{ij \in R} \kappa_{ij} \tag{2}$$

We now describe a simple mechanism to calculate capture-cost of a beam. We use the specific example in Figure 6 for our description, and then generalize the mechanism. In Figure 6, when there is only a single flow between nodes B and A, beam B_{MB} is affected by capture. Now observe that beam B_{MB} is the very same beam as beam B_{MF}, i.e., the east-pointing beam. While estimating the capture cost of this east-pointing beam, the capture-cost of all links that utilize this beam needs to be taken into account. For this specific example, the updated cost of beam B_{MF}, $C'_{B_{MF}}$, is equal to $C_{B_{MB}}$. In general,

$$C'_{B_{ij}} = \{\sum_k C_{B_{ik}} :\ B_{ik}\ is\ same\ as\ B_{ij}\} \tag{3}$$

Using this mechanism, observe that the capture-cost of beam B_{ME} is zero because it is not captured by already existing traffic in the network. Thus, when comparing the cost of alternate routes from S to G, the cost of route $\{$S-M-F-G$\}$ is higher than $\{$S-M-E-G$\}$, making $\{$S-M-E-G$\}$ the preferred route with our capture-aware routing metric.

Incorporating Hop-Count: Since capture-cost of a beam can be zero, the cost of a route does not increase monotonically with increasing number of hops. This can lead to routing loops. More importantly, selected routes may become very long in order to avoid capture. Since longer routes affect end-to-end latency, we incorporate hop count into our metric. Each link ij on a route is associated with a cost H_{ij}. Existing metrics such as ETT, that account for link quality, can be used to update the value of H_{ij}.

Sharing Active Nodes: Routes that are equivalent in terms of capture and hop-count, may not offer equal throughput. This is because intermediate nodes in one route may participate in other communications, and thereby be less available to a newly admitted flow. To favor routes on which nodes are more available, we require the routing module at node i to maintain a *participation variable*, P_i, that reflects the node's participation in existing traffic. Nodes that act as source, intermediate nodes, or sink are assigned a positive value for P_i, and are called "active nodes". Nodes that do not participate in any ongoing traffic are assigned $P_i = 0$, and are referred to as "inactive nodes". Participation variables at active nodes are associated with timeouts. If an active node does not participate in any communication for a threshold period of time, the node resets P_i to 0.

B. A Unified Routing Metric

The unified cost of a route, U_{route}, can now be formulated as a weighted sum of all the factors that contribute to the capacity of a route. Thus U_{route} is expressed as

$$U_{route} = \sum_{ij \in R} \omega_\kappa \kappa_{ij} + \omega_p P_i + H_{ij} \qquad (4)$$

where ω_κ and ω_p are weights used to control the influence of participation and capture, on the choice of routes. In our simulation of this routing metric, we assign $H_{ij} = 1$. For the capture-cost, we assign $C_{B_{ij}} = n$, where n is the number of nodes that capture node i on its beam $C_{B_{ij}}$. For active nodes, we assign $P_i = 1$ when a node is a source or destination of a flow. However, if an active node is an intermediate node of at least one flow, we assign $P_i = 2$, since it bears the onus of both receiving and forwarding.

We are aware that $C_{B_{ij}}$ and P_i are very course-grain representations of the amount of directional-capture and load on a node. However, they have the property of being insensitive to changing traffic patterns in the network, and thereby prevent oscillations. We intend to refine our cost models as a part of our future work, and incorporate them in our proposed framework.

C. Protocol Design

We present the design of a source-initiated capture-aware routing protocol (CaRP), similar to DSR [17]. The source initiates a route request (RREQ) flood in the network – each RREQ accumulates the cumulative cost over all the links it traverses. Observe that RREQ packets are always transmitted using omnidirectional beams. The mechanism of cost accumulation in RREQs is performed as follows. Consider the case in which a node X receives a RREQ originated by a source node, S. The RREQ forwarded by node X contains the following information.

(1) The cumulative *capture-cost* associated with the partial route from S to X, added to the *capture-cost* for each of X's beams. This implies that a tuple of N separate costs are included in the RREQ, where N is the number of beams. The i^{th} element of the tuple reflects the sum of the capture cost of all beams on the partial route $\{$S-X$\}$, added to the capture-cost of X's i^{th} beam. Figure 7 illustrates this with a diagram.

(2) The cumulative *participation cost* on the route along which this RREQ has arrived, denoted as A_{sx}.

(3) The intermediate nodes' identifiers (i.e., the partial source route) along which the RREQ has traversed to reach X from S. This indicates the hop-count.

When node Y receives the RREQ, it deduces the capture-cost on the partial route from S to Y, using the information available in the RREQ. For example, if node X uses beam 2 to communicate to Y, then Y selects the 2^{nd} value from the tuple included in the RREQ, and adds to it the capture-cost of the beam that Y used for receiving the RREQ from X[3]. The result represents the capture-cost for the partial route from S to Y.

Node Y then updates the participation cost as $A_{sy} = A_{sx} + P_Y$. Then, it calculates the hop count using the list of intermediate nodes included in the RREQ. Once the three factors are calculated individually, node Y calculates the unified cost of the partial route from S to Y, U_{SY}, using equation (4). Node Y remembers the cost of this partial route, for reasons explained later. It then updates the RREQ with the required information, and forwards it. Figure 7 illustrates the cost accumulation procedure in a propagating RREQ.

When the destination node receives the RREQ, it cal-

---

[3] We assume a neighbor discovery module that enables node Y to be aware of the beam used by X, to communicate to Y. In stationary networks like mesh networks, this can be achieved with low overheads [18].

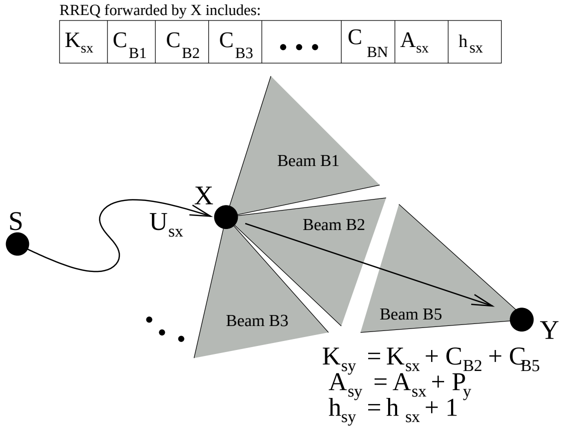

Fig. 7. Accumulation of route cost in the propagating RREQ

culates the cost of the route using equation (4). If the RREQ happens to be the first one for this route discovery phase, the destination node always responds with a route reply (RREP). For RREQs that arrive later, the destination responds with an RREP only if the cost of the received RREQ is lower than all of the previously seen RREQs.

In the basic DSR protocol, an intermediate node does not forward a duplicate RREQ under the assumption that the earlier arriving RREQ traverses the better path. However, with capture-aware routing, a longer path may be associated with a lower cost. To account for this possibility, we require intermediate nodes to remember the least cost RREQ that it has forwarded for a particular route discovery. If a lower-cost RREQ arrives later, then we require the intermediate node to forward this RREQ as well.

D. *CaRP Details*

Choosing ON and OFF durations: The length of ON and OFF durations lead to a performance tradeoff. Longer OFF-Durations is beneficial because higher spatial reuse can be obtained for longer durations. However, a sub-optimal route discovered in an OFF-Duration will also persist longer. While the optimal ratio between ON and OFF durations requires investigation, we observe that for a given optimal ratio, choosing the smallest possible ON duration is desirable. The smallest possible ON-Duration is a function of the time taken for a source node to discover a route. Knowing the diameter of a network, the length of the ON-Duration can be calculated (with necessary factors of safety). The OFF durations can in turn be calculated based on the ON-Duration.

Clock synchronization: At the beginning of an ON duration, we propose to initiate route discovery after a random delay, t. The delay t can be viewed as a guard band, that relaxes nodes from the need to be accurately clock synchronized. Since the accuracy of clock synchronization can be fairly high (less than 10 μs [19]), the guard band remains reasonably small compared to the length of the ON-Duration. Of course, t is chosen such that the initiated route discovery can finish before the node switches to the OFF-Duration.

VI. PERFORMANCE EVALUATION

We use the Qualnet simulator, version 3.6, for simulating our proposed protocols. We consider several scenarios, and systematically analyze the factors that impact protocol performance. We use a transmission data rate of 11 Mbps. We use a two-ray ground reflection model with Rayleigh fading. We use constant bit rate (CBR) traffic at each source node. The simulated antenna beams are characterized with sidelobes and backlobes, obtained from realistic antenna patterns [20]. We compare the performance of CaDMAC and CaRP with existing MAC and routing protocol, namely DMAC [7] and DSR [17].

A. *Simple Scenarios*

We begin by using the simple scenario from Figure 5(a) – node A communicates with node C via node B, while node E communicates with node G via node F. Routes are specified statically for this simple scenario. The ON and OFF durations are assigned as 1 and 3 seconds respectively (for reasons discussed later), and the beamwidth is specified to be $90°$. We compare the end to end throughput of CaDMAC against DMAC [7], Circular-MAC [8], and Omnidirectional 802.11.

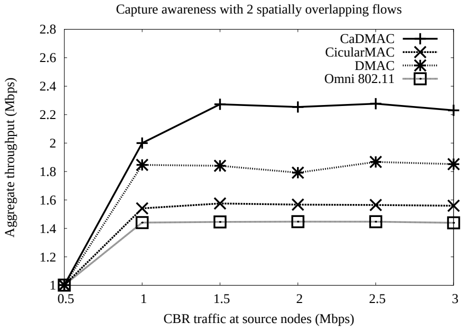

Fig. 8. CaDMAC improves throughput for the simple scenario

At low sending rates, the benefits of capture-awareness are not evident. As the sending rates increase, the network begins to saturate earlier with DMAC and Circular-MAC. This is because node B and F frequently get distracted (i.e., captured) by the ongoing cross traffic. Filtering out such distractions during the OFF periods enable CaDMAC to achieve better throughput.

B. Large Random Scenarios

We simulated large networks, with 50 nodes randomly placed in a square region of side 1500m. We selected 15 random source-destination pairs, and chose arbitrary multihop routes between them. The routes are *not* selected based on our *capture-aware routing metric*. As a result, the graphs in Figure 9 reflect the MAC-layer benefits alone. The ON and OFF durations are specified to be 1 and 3 seconds respectively, while the beamwidth is 60°. The results are an average of 25 topologies.

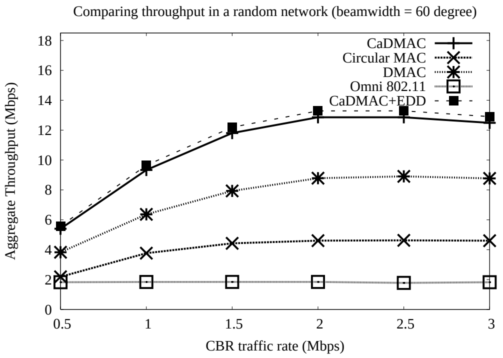

Fig. 9. Throughput benefits from capture-aware MAC protocols in large random networks, with beamwidth of 60°

As evident from Figure 9, the relative benefits from CaDMAC increases for larger, congested networks. In congested networks, the channel capacity saturates later for CaDMAC, leading to throughput improvements. The EDD optimization with CaDMAC offers marginal improvements because of using RTS/CTS packets, as well as small data packets (512 Bytes). With 1024 Byte packets, and absence of RTS/CTS, we observed an improvement of 18% (not reported here for the interest of space).

Delay in Random Networks

Figure 10 presents the improvements in average end to end delay available from capture-awareness. Nodes that are kept less occupied by capture, can accomplish productive message deliveries, quicker. Thus average end to end delay also improves in comparison to other protocols that have to postpone communication while captured.

C. Capture-Aware Routing Protocol (CaRP)

The results presented till this point reflect the benefits of capture-awareness at the MAC layer. In this section, we aim to show that when CaDMAC is used in conjunction with CaRP, performance can be superior.

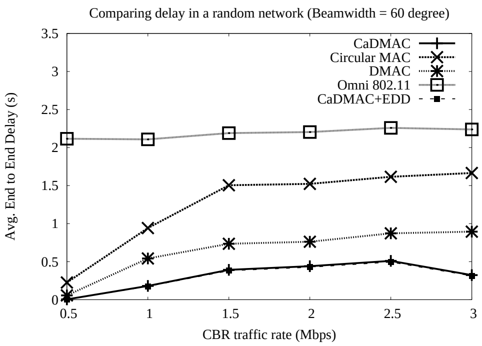

Fig. 10. Latency benefits from capture-aware MAC protocols in large random networks, with beamwidth of 60°

We begin with a small network, laying out the detailed impact of choosing capture-aware routes. Consider Figure 11, where a few nodes are placed in a rectangle. In this small network, we introduce four flows in the following order: (1) From A to D, (2) from E to G, (3) from H to M, and (4) from U to Y. More precisely, flow (i) is initiated in the i^{th} ON-Duration, allowing CaRP to choose its optimal route. Figure 11 (b) and (c) show the routes chosen by DSR and CaRP respectively (both use CaDMAC as the MAC protocol). The aggregate throughput (in Mbps), as a result of using these routes, are as follows: (CaRP+CaDMAC) = 3.64, (DSR+CaDMAC) = 2.38, (DSR+802.11) = 0.98.

The properties of the chosen routes (in Figure 11) explain the differences in achieved throughput. Observe that for flow 1, the routes chosen by DSR and CaRP are identical – in the absence of other traffic in the network, the metrics of DSR and CaRP degenerate to hop-count. However, when flow 2 is introduced, unlike DSR, CaRP avoids using link Q-R because it is prone to capture from link B-C. Instead CaRP uses link Q-F, which is capture-free from the traffic on route $\{$A-B-C-D$\}$. Similar trends are visible in other cases as well – for flow 3, DSR chooses link J-K that is susceptible to capture from C-D, whereas CaRP chooses J-L. Finally, when flow 4 is introduced, CaRP utilizes the uncongested region of the network by choosing the route through node W. DSR chooses a route through link V-X, which is prone to capture from traffic over link E-Q. *Observe that there is still more capacity available (in this scenario), since all the available directions have not been saturated.* In a more congested network, CaRP will still be able to “squeeze in” more flows that do not share the bandwidth between themselves.

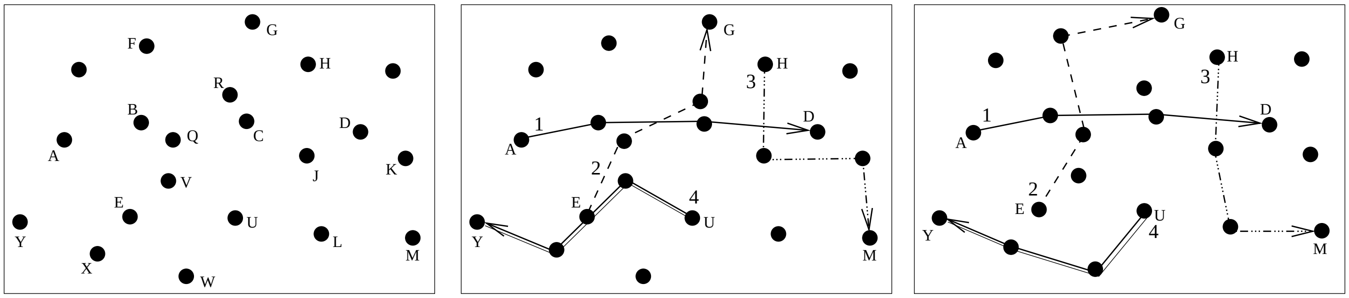

(a) A small network
(b) Routes chosen by DSR+CaDMAC
(c) Routes chosen by CaRP+CaDMAC

Fig. 11. Enumeration of routes when using DSR and our capture-aware routing protocol, on a small topology with 4 ordered flows.

To this point, all data flows were forcibly initiated during ON-Durations. This allowed us to measure the optimistic benefits from capture-aware routing. However, in reality, route discoveries may very well be initiated during the OFF durations. The next experiment initiates traffic at random points of time. We compare the performance of DSR with omnidirectional antennas, DSR with CaDMAC, and CaRP with CaDMAC.

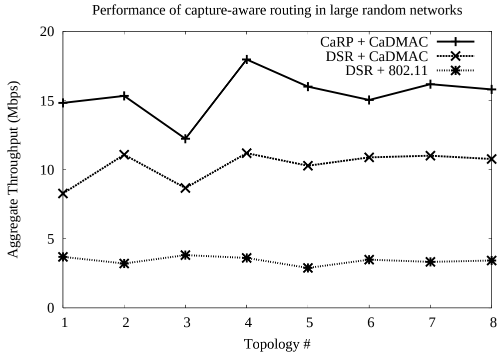

Fig. 12. Capture-aware routing in large random networks.

We randomly placed 50 nodes within a square of size 2000m, and initiated 20 CBR flows between randomly chosen source and destination nodes. The starting times were chosen randomly as well (without attention to the ON and OFF durations). We simulated our experiments for various topologies. Observe that when using omnidirectional antennas, the route discovery process is always performed on the unaltered topology – DSR chooses the minimum-hop routes in this case. However, when DSR and CaRP are executed over CaDMAC, the route discovery can often be initiated during the OFF-duration, and then re-initiated in the subsequent ON-Duration.

Figure 12 shows the results for 8 different topologies. Although many route discoveries were initiated in OFF-Durations, the benefits of CaRP remain reasonably higher than DSR. This is because CaRP compensates for the sub-optimal routes by choosing a high throughput route in the immediate next ON-Duration. Moreover, in moderately dense networks, the route chosen first may not be severely suboptimal because of the availability of multiple nodes that may serve as alternates. With narrower beamwidth, we believe that the improvements can be higher.

Choosing ON/OFF Durations: We study the sensitivity of throughput and delay to the length of OFF-Durations. Figure 13 presents normalized performance (with respect to DSR+802.11) against the value of $\frac{\tau_{off}}{\tau_{on}}$. Observe that the normalized throughput increases with increase in the OFF-Duration, but after a point (around $\frac{\tau_{off}}{\tau_{on}} = 3$), the curve flattens. Since the duration of CBR flows are specified to be short, very long OFF-Durations force the initially discovered (suboptimal) route to take up a large fraction of the route lifetime, offsetting the gains from capture-awareness. The latency to discover the optimal route has been shown in Figure 13. Observe that for more persistent traffic, the benefits from longer OFF-Durations can be higher.

The normalized end to end delay also improves for small OFF durations – recall that even though sub-optimal routes are used during OFF-Durations, the capture-aware routes selected in the next ON-Duration compensates for the losses. However, the compensation is not sufficient when the OFF-Durations are too long.

Observe that Figure 13 can be used by a network designer for selecting the length of ON and OFF durations. For this paper, we used $\frac{\tau_{off}}{\tau_{on}} = 3$, since that region appeared to be reasonable both in terms of throughput and delay.

VII. DISCUSSION

Dependence On Traffic Pattern: With CaRP, the cost of a route is a function of the traffic pattern in the

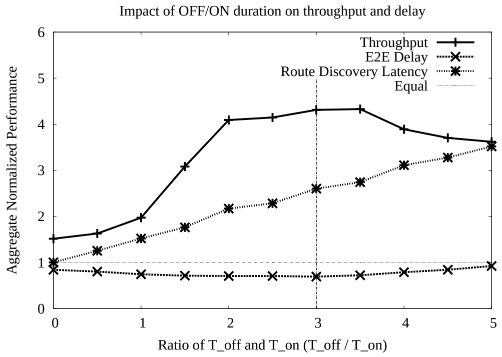

Fig. 13. Impact of ON and OFF durations on performance.

network. If the traffic pattern changes between time instants t_1 and t_2, a route selected before t_1 may be suboptimal after time t_2, affecting spatial reuse. One way to address this could be to have nodes periodically rediscover routes during the ON-Duration. We intend to investigate this issue as a part of our future work.

Incorrect Filtering: Due to multipath and time-varying nature of the wireless channel, it is possible that the chosen multi-beam pattern proves to be inappropriate during the OFF-Duration. In such cases, some communications will fail, leading to channel wastage. CaDMAC is unable to handle this issue efficiently. We intend to address this problem in our future work.

VIII. Conclusion

We have identified the need to exploit beamforming antennas even while a node is in the idle state – we show that remaining in the omnidirectional mode can cause a idle node to get frequently distracted by unproductive communication. We call this problem "MAC-Layer capture", and we believe that the problem will apply to a large body of existing directional antenna protocols. We provide one way of resolving capture at the MAC and the network layer. We believe that for stationary mesh networks, with relatively persistent route patterns, our protocols can improve spatial reuse. For networks with mobility and more dynamic route variations (such as in mobile ad hoc networks), the problem of capture needs to be further researched.

References

[1] S. Biswas and R. Morris, "Opportunistic routing in multi-hop wireless networks," in *ACM Sigcomm*, 2005.

[2] R. Karrer, A. Sabharwal, and E. Knightly, "Enabling large-scale wireless broadband: The case for taps," in *Proceedings of Hot Topics in Networks (HotNets II)*, 2003.

[3] J. Litva and T. Lo, "Digital beamforming in wireless communications," in *Artech House Publishers*, 1996.

[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] 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.

[6] 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.

[7] 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.

[8] 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.

[9] R. Roy Choudhury and N. H. Vaidya, "Deafness: A MAC problem in ad hoc networks when using directional antennas," in *Proceedings of the IEEE International Conference on Network Protocols (ICNP)*, October 2004.

[10] K. Sundaresan and R. Sivakumar, "A unified mac layer framework for ad-hoc networks with smart antennas," in *Proceesinngs of ACM International Symposium on Mobile Ad hoc Networking and Computing (MOBIHOC)*, May 2004.

[11] V. Navda, A. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "Mobisteer: Using steerable beam directional antenna for vehicular network access," in *ACM MobiSys*, 2007.

[12] C. Shen and Z. Huang, "Topology control with directional power intensity for ad hoc networks," *Proceedings of IEEE Computer Communications and Networks*, 2002.

[13] J. Tang, G. Xue, C. Chandler, and W. Zhang, "Interference-aware routing in multihop wireless networks using directional antennas," in *Proceedings of INFOCOM*, March 2005.

[14] S. Roy, D. Saha, S. Bandyopadhyay, T. Ueda, and S. Tanaka, "A network-aware MAC and routing protocol for effective load balancing in ad hoc wireless networks with directional antenna," in *Proceedings of ACM MobiHoc*, June 2003, pp. 88–89.

[15] 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.

[16] J. Boer, "Wireless lan with enhanced capture provision," in *U.S.Patent 5987033, Technical Report*, 1999.

[17] D. Johnson and D. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing, Kluwer Academic Publishers*, 1996.

[18] S. Vasudevan, J. Kurose, and D. Towsley, "On neighbor discovery in wireless networks with directional antennas," in *Proceedings of INFOCOM*, March 2005.

[19] J. Elson, L. Girod, and D. Estrin, "Fine-grained network time synchronization using reference broadcasts," in *Proceedings of Operating Systems Design and Implementation (OSDI)*, 2002.

[20] R. Schlub, "Practical realization of switched and adaptive parasitic monopole radiating structures," Tech. Rep., PhD Thesis, Griffith University, January 2004.