Capture-Aware Protocols for Wireless Multihop Networks Using Multi-Beam Directional Antennas

Technical Report, March 2005

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—**We define a node to be "captured" when its receiver is engaged in receiving a packet that is not intended for that node. A unicast transmission using IEEE 802.11, although intended for a single destination node, can result in *capturing* several other nodes located within the interference footprint of the transmitter. Capture is mostly unproductive, because the physical layer of a captured node is kept occupied in receiving a packet which will eventually get dropped at the MAC layer. This can degrade network capacity because captured nodes refrain from scheduling useful communication during the period of capture. Antennas capable of beamforming may be able to overcome capture by spatially filtering directions from which signals lead to capture. As a result, nodes might be able to gain more free time to simultaneously schedule their own productive communication.**

This paper presents protocols that utilize multi-beam antennas to minimize the impact of capture. At the MAC layer, nodes characterize ongoing transmissions based on their directions, and use multi-beam antenna patterns to suppress unwanted signals – thus, the antenna can implement a spatial filter. The network layer assists the MAC layer by using a new *capture-aware routing metric* that prefers routes less prone to capture. Simulation results show that reducing the impact of capture can help improve performance in wireless multihop networks.

I. Introduction

Directional antennas have received increasing research attention in the recent years. Protocols that utilize directional antennas in wireless multihop networks, (such as [1],[2],[3],[4],[5],[6]) have been shown to outperform protocols that use omnidirectional antennas, especially in stationary environments. The improvements with directional antennas result primarily from higher spatial reuse, and longer communication range.

Although better than omnidirectional antennas, directional antenna protocols do not take maximal advantage of spatial reuse. This is because these protocols lend themselves to the problem of "capture", and thereby underutilize the network capacity.

- *Definition of Capture*

We define "capture" as a problem in which nodes get engaged in receiving an unproductive packet, and thereby cannot utilize that time duration for its own (useful) communication. Capture is a prominent problem with IEEE 802.11 – a wireless unicast, although intended for a single destination node, can result in *capturing* several other nodes located within the communication range of the transmitter. When using omnidirectional antennas, the motivation to mitigate capture is limited because captured nodes are anyway not expected to initiate simultaneous communications themselves.

In the wireless networking community, the term "capture" is typically associated to a physical layer phenomenon. The *effect of capture* refers to a purely physical layer phenomenon, namely, the successful demodulation of the stronger signal, when two signals collide at a common receiver [7]. As we later illustrate, our notion of "capture" applies more to a MAC layer phenomenon.

We argue that directional antennas can be utilized to mitigate the problem of capture. This is because when using directional antennas, a node has the ability to spatially filter out unwanted signals by choosing appropriate beamforms. In the past, spatial filtering has been exten-

Work reported here is supported in part by Motorola Center for Communications Fellowship.

sively studied at the physical layer. We propose to perform spatial filtering at the MAC layer with the aim to support more simultaneous communications. Previously proposed MAC protocols for directional antennas, such as those in [1],[2],[3],[4],[5],[6], are limited in the number of simultaneous communications they support, as illustrated next, using Figure 1.

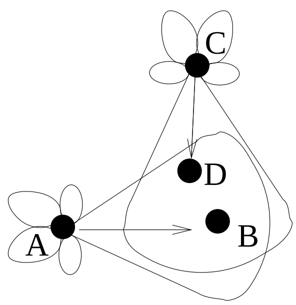

Fig. 1. Capture prevents parallel dialog between A-B and C-D

In Figure 1, assume that node A intends to transmit to node B, and node C to node D. When using a directional MAC protocol, such as DMAC [3] (or its variants), if A initiates a transmission to B before C can begin its transmission to D, node D will get captured by A's transmission. This happens because DMAC (and its variants) require node D to be in the omnidirectional mode, while waiting for signals to arrive. As a result, the protocols lend themselves to being captured by unproductive signals. Observe that if node C initiates transmission first, node B will get captured; in neither case will the two dialogs proceed simultaneously. Readers not familiar with DMAC may find the above explanation unclear. We will revisit these issues and explain them in more detail in later parts of the paper.

We propose to handle capture by observing that it might be possibile to discriminate directions of useful traffic from those of capture traffic. Based on such a discrimination, node D might choose to use a multi-beam antenna pattern that filters out signals from A. Clearly, if node B also chooses an appropriate antenna pattern, the two dialogs can progress simultaneously. However, node D cannot continue to use such an antenna pattern indefinitely, in view of the possibility that node A might wish to transmit to D in the future. At the MAC layer, a suitable schedule to avoid signals from A can free node D from getting captured over short time scales. At the network layer, selecting routes that are less prone to capture, may improve channel utilization over longer time scales (i.e., node A may prefer routes that do not affect communication between C and D). This paper proposes MAC and routing protocols that aim to achieve these improvements. The contributions of our work can be summarized as follows.

- We identify the problem of capture in some existing omnidirectional and directional antenna protocols. We show that capture can be alleviated using multi-beam directional antennas to improve performance.

- We propose a MAC protocol that implements a beam selection policy to reduce the impact of capture. In addition, we suggest an optimization called *Early Destination Detection*, wherein the *destination field* in a packet is detected early to avoid capture on small time scales.

- We propose a routing protocol that uses a *capture-aware* metric to select routes on which links are less prone to capture. This results in high-throughput routes since each route is less "distracted" by traffic from other routes. We show that capture-awareness at the MAC and network layer can complement each other, thereby improving on the performance.

II. Related Work

The possibility of using directional antennas in multihop networks has been investigated in the recent past. The benefits have been explored at several layers of the protocol stack. At the physical layer, signal processing algorithms have been proposed for spatial filtering with adaptive beamforming [8],[9]. Smart array antennas, that exploit these algorithms, are appearing in the commercial market [10], and offer orders of magnitude gain in terms of link quality. At the MAC layer, significant research has focused on designing protocols that exploit the directional capabilities [1],[2],[4],[3],[5],[11],[12]. Neighbor discovery, topology control, and multicast have been studied in [13],[14],[15],[16]. At the network layer, directional antennas have been shown to find fewer hop routes [17], energy-efficient routes [18] and interference-aware routes [19]. Several researchers have also investigated the use of multi-beam antennas [20],[16]. Some projects also propose to implement ad hoc networking testbeds using directional antennas [6],[21]. We discuss some of the existing work from this rich body of research, and explain how our contributions in this paper improve the state-of-the-art.

The early work on MAC protocols using directional antennas improved spatial reuse by enabling nearby communications to progress in parallel [1],[2]. Further investigation improved upon the possibilities of simultaneous communication, and introduced the notion of directional virtual carrier sensing through the directional NAV (DNAV) mechanism [4]. However, several tradeoffs were also identified, including new kinds of hidden terminal problems, and deafness [3], that limited the benefits of directional antennas in multihop scenarios. Subsequent protocols addressed some of these problems and raised the limits on achievable performance [22],[11],[12]. Another problem faced at the MAC layer was related to discovering neighbors with directional antennas. This problem has recently been addressed in several works including [14],[13],[23].

The impact of directional antennas on multihop routing has also emerged to be of interest. The longer transmission range of directional antennas holds promise of discovering shorter routes, in addition to the prospect of localizing queries only toward the direction of the destination [17]. Moreover, the ability to control the topology using directional antennas [15] made it a candidate for selecting less congested routes [18],[19]. From the perspective of energy awareness, directional antennas were also shown to be beneficial for routing, broadcast, and multicast traffic [16],[19],[23],[24],[25].

The opportunity to benefit from eliminating capture has has not been exploited with existing directional antenna protocols. In this paper we utilize multi-beam directional antennas to design MAC and routing protocols to address this problem. Multi-beam antennas have also been proposed previously in the context of ad hoc networks [20],[2],[1], although none of them have attempted to address "capture" at the MAC or network layer.

III. Preliminaries

A. Antenna Model

We present an abstract antenna model here. The model is simplified for the convenience of explanation (the actual antenna model used for simulations are more realistic, as described later). In the abstract model, the antenna system comprises of N non-overlapping beams such that the union of all the beams covers the entire ($360°$) azimuth plane. The shapes of the beams can be envisioned as sectors, each sector having a beamwidth of $\frac{360}{N}$. Each antenna beam has an on-off switch, and signals from each of the beams are combined equally to feed a single transceiver. Using these switches, it is possible to turn on/off any subset of beams, resulting in various multi-beam patterns. A block diagram of the antenna system is presented in Figure 2, and example beam-patterns are shown in Figure 3. In Figure 3(b), a node only has a single beam turned on towards north. In all other directions, it has pointed a null implying that signals arriving in these directions are suppressed. In reality (and in beam-patterns used in our performance evaluation), nulls include low-gain sidelobes and backlobes, as illustrated in Figure 4.

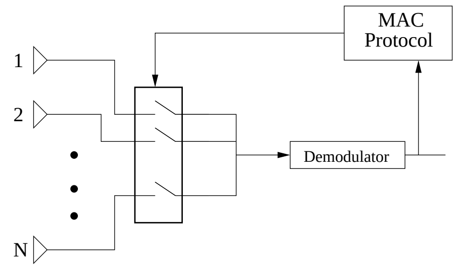

Fig. 2. Switched combining with multi-beam antennas

When all the beams of the antenna are turned on (Figure 3(a)), the beam-pattern is "omnidirectional". When transmitting in the omnidirectional mode, the antenna uses a fixed power P_{max}. When transmitting using a subset of its beams, the transmit power is appropriately reduced such that communication range always remains equal, irrespective of the beam-pattern being used.

We assume that the antenna system has the ability to detect the beam-of-arrival of a signal. In other words, 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.

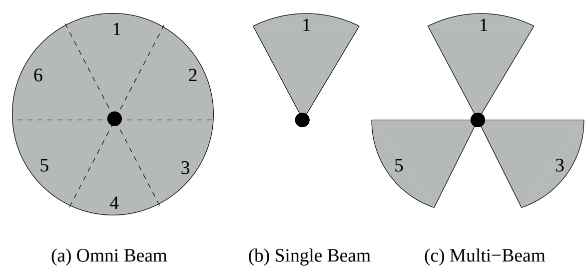

Fig. 3. An abstract antenna model

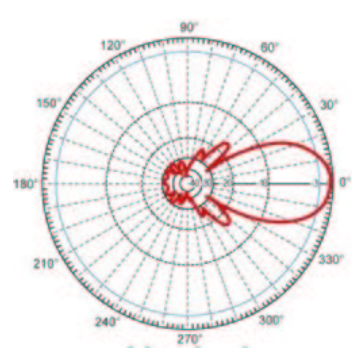

Fig. 4. A realistic radiation pattern including sidelobes.

Communication on antenna systems are governed by Friss equations [26]. Signals that impinge on the mainlobe of an active beam, are received with a higher gain, while those that arrive from other directions are suppressed by the lower gain of the antenna's sidelobes. One may view such a beam pattern as a spatial filter (a well-understood concept in the area of adaptive beamforming). Like frequency filters that filter out signals on the frequency domain, spatial filters can be made to filter signals in space. The abstract antenna model presented above can be realized in many ways, including diversity antennas and adaptive beamforming antennas [8]. Of course, such realizations will have beamforms with sidelobes (simulated in this paper for evaluating our proposed protocols).

B. IEEE 802.11, Directional MAC, and Deafness

We present a brief overview of 802.11, an existing directional MAC protocol (DMAC) [3], and the issue of deafness, previously identified as a problem with directional communication.

- *IEEE 802.11 DCF*

IEEE 802.11 is a CSMA/CA protocol that performs *physical carrier sense* before initiating transmission. Once the channel is sensed 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, S doubles its CW, counts down a newly chosen backoff interval before attempting retransmission. 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 the dialog, so that neighbors of the communicating nodes can defer their own transmissions – called *virtual carrier sensing*.

- *Directional MAC - DMAC*

DMAC [3] proposes modifications to 802.11 for adapting to directional antennas. The RTS is transmitted by pointing the transmitter's beam towards the intended receiver (called *directional mode*). The transmission of the RTS is preceded by a backoff phase (identical to 802.11) during which the transmitter remains in the omnidirectional mode. The receiver receives the RTS, and beamforms in the direction of the transmitter to reply with the CTS. The DATA and ACK packets are then communicated directionally. Once the dialog is over, both the transmitter and the receiver switch back to their omnidirectional mode, and wait for the next dialog. Idle nodes in the vicinity of the transmitter and receiver overhear the RTS/CTS packets, and defer communication for the proposed duration of the dialog. However, unlike 802.11, communication is deferred only along those directions in which the RTS/CTS packets arrived. This is termed as *directional virtual carrier sensing, DVCS*.

- *Deafness*

Deafness has been identified as a problem with directional communication [11]; briefly, deafness is caused when a node X intends to communicate to node Y, but is unaware that Y is pointing its beam toward a different direction. As a result, X might often retransmit multiple times to Y, and eventually drop the packet. Due to the exponential backing off algorithm in DMAC, deafness leads to degradation in throughput and fairness. Several authors from the past have addressed the problem of deafness [11],[22],[27],[28].

C. Motivation and Intuition

Many directional MAC protocols require an idle node to be in the omnidirectional mode [1],[2],[4],[3],[29],[11]. This is because when using CSMA/CA protocols like IEEE 802.11, an idle node is not expected to know the direction from which signals may arrive. Once a signal impinges on the receiver node, different protocols assume different techniques to beamform toward the signal. Once beamformed, the rest of the dialog is typically accomplished in the directional mode. Once the dialog is over, the node returns to the omnidirectional mode, waiting for the next packet. Observe that this model of operation lends itself to capture. *While an idle node is in the omnidirectional mode, it is susceptible to get captured by every (sufficiently strong) signal that impinges upon its transceiver.*

This observation motivates protocols that are "capture-aware". Beams at a node, that are prone to capture, but are not needed for productive communication, can be turned off for an appropriate time duration, say τ_{off}. Performance can potentially be improved, because nodes that would otherwise get captured may now be able to participate in productive communication for the τ_{off} duration. However, tradeoffs arise when traffic in the network changes within the time window of τ_{off}. In addition, turning off a subset of the beams does not eliminate the possibility of capture completely. A node will continue to get captured, if both productive and capturing signals arrive on the same beam.

These tradeoffs led us to explore the benefits of designing a capture-aware routing protocol. We observed that if a routing protocol is aware of the τ_{off} time window, then it might be possible to address the problem of changing traffic patterns. More importantly, we reasoned that if the cost of a route can account for *the "amount" of capture on the route*, then routes with lower capture can be preferred. Put differently, intermediate nodes on such routes are less likely be "distracted" by capture.

Motivated from these observations, we propose the design of capture-aware MAC and routing protocols. In Sections IV and V, we describe the protocols in detail, with some optimizations to improve on their performances.

IV. Capture-Aware MAC Protocols

We begin with a basic version of a capture-aware MAC protocol, CaMAC. The main aim of CaMAC is to implement a beam-selection policy, such that the selected beam can spatially filter out unproductive signals that lead to capture. We identify shortcomings of CaMAC, and adopt techniques from existing directional MAC protocols to improve its design. The improved protocol is termed CaDMAC. We discuss deficiencies of CaDMAC, and propose optimizations to further improve upon it.

A. Protocol Description: CaMAC

We divide time into recurring cycles of duration T_{cycle} as shown in Figure 5. Each cycle is further subdivided into an *ON-Duration* and an *OFF-Duration*, of lengths τ_{on} and τ_{off}, respectively. Each cycle begins with an ON-Duration, and ends with an OFF-Duration. For our proposed protocol, we assume that nodes are clock-synchronized. Nodes enter the ON and OFF durations synchronously, and the cycle repeats. The *ON-Duration* denotes the time period during which all nodes in the network have all their beams turned on, i.e., all nodes are in the omnidirectional mode. The *OFF-Duration* denotes the time period during which a subset of the beams may be turned off. Note that during the *OFF-Duration*, the transceiver is not switched off and the node continues communicating – only some of the antenna beams may be turned off (as detailed below).

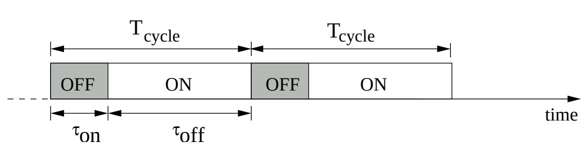

Fig. 5. Time divided into cycles of ON-Durations and OFF-Durations

While in the ON-Duration, the MAC layer at each node passively monitors the ongoing traffic on a per-beam basis. Monitoring ongoing traffic involves segregating incoming packets into 2 categories – (i) productive traffic and (ii) capture traffic. Productive traffic at a node includes unicast packets meant for that node, as well as broadcast packets. Capture traffic at a node includes overheard packets that are intended for other nodes. If a beam proves to be the receiver of only capture traffic[1], and if the beam is also not used for transmission over the ON-Duration, then the beam is "black-listed". At the end of the ON-Duration, the MAC layer decides to turn off all the 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 the OFF-Duration. As a result node i is left with more time to participate in productive communication, leading to a potential increase in network throughput.

In CaMAC, the basic channel access technique during the ON/OFF durations is very similar to 802.11, and is described briefly with the help of an example from Figure 6.

In Figure 6, assume that node A communicates with node C via node B, while node E communicates with node G via node F. While in the ON-Duration, nodes identify their respective beams that must be turned off. Figure 6(a) shows the beam patterns β^B_{off} and β^F_{off} used by B and F, respectively, during the OFF-Duration. Let us consider node B. When node B intends to transmit a packet to C, it proceeds through the same set of procedures as specified in 802.11. Node B senses the channel using beam β^B_{off}. If the channel remains idle over the entire DIFS and backoff periods, then node B initiates the dialog to node C. Once the dialog is over, node B returns to the idle state. While in the idle state, node B may receive a new packet from A, or may prepare to transmit another packet to C. Observe that while B communicates with nodes A and C, node F also communicates with E and G using the beamform β^F_{off}. The two flows progress in parallel, without being interrupted by capture.
Observe that the beam patterns shown in Figure 6 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.

As mentioned earlier, many directional antenna protocols proposed in the past [1],[2],[4],[3],[29],[11], are not designed to realize this form of spatial reuse. With these protocols, when E transmits packets to F, node B gets cap-

[1]Recall that each packet arrives at the MAC layer along with the beam-of-arrival information.

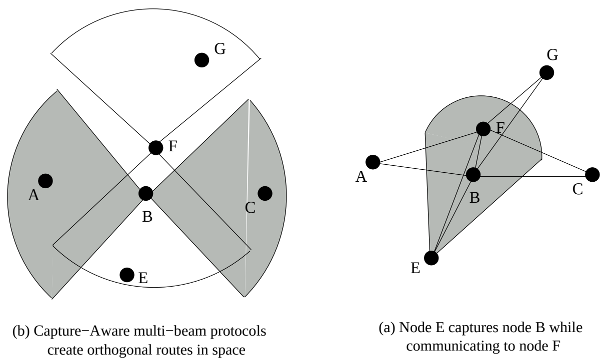

Fig. 6. Illustrating the efficacy of CaMAC in exploiting spatial reuse.

tured by the transmission, as shown in Figure 6(b). The problem occurs primarily because idle nodes remain in the omnidirectional mode, thereby lending themselves to the possibility of capture.

While our proposed protocol may be able to mitigate capture during the OFF-Duration, an important concern appears from the routing perspective. With reactive protocols such as DSR [30] or AODV [31], if a route discovery is initiated during the OFF-Duration, discovered routes may prove to be sub-optimal. In some cases, a route may not be discovered at all, even though a route exists. This can happen because, during the OFF-Duration, the topology of the network would be different from the original topology (recall that during the OFF-Duration, nodes may turn off a subset of their beams, and thereby remove links from the actual topology). We address this problem in a later subsection, titled “Dependence on Routing Protocol”.

B. Issues with CaMAC

While CaMAC appears to mitigate capture in the simple scenario in Figure 6, its benefits in larger networks proved to be limited in comaprison to DMAC. This was because the multi-beam patterns used by CaMAC interfered over a larger area, as opposed to single beams used by DMAC. A larger interference footprint caused each transmission to interfere with several other communications, offsetting the spatial reuse gained from mitigating capture. We present a simple scenario in Figure 7 to explain this point. In Figure 7, assume that node A communicates with C via B, while node E communicates with G via F. Now, observe that if node C also initiates a flow destined to node D, then the communication between C-D and A-B cannot occur simultaneously. This is because whenever node C communicates with node D using its multi-beam pattern, it captures node B. The vice versa happens when node B communicates to node A. In either case only one of the dialogs can progress at any given time. DMAC on the other hand, uses single beams for communication, and thereby allows for simultaneous communication between A-B and C-D (i.e., C points its conical mainlobe toward D while B points its conical mainlobe toward A). In such circumstances, DMAC improves spatial reuse from interfering over a smaller region, but continues to suffer from capture and deafness.

C. Enhancing CaMAC to CaDMAC

We enhance CaMAC to *Capture-aware directional MAC (CaDMAC)* with the aim to retain the benefits of both CaMAC and DMAC. To achieve this, we adopt techniques from single-beam communication protocols, and overlay them on CaMAC. With CaDMAC, a node always uses a single beam to communicate to its upstream and downstream nodes. Once a communication is over, an idle node i returns to the multi-beam pattern β^i_{off} (pre-selected for the ongoing OFF-Duration), and waits for the next dialog to begin. Referring to Figure 8, notice that CaDMAC can support all 3 communications simultaneously, namely communications between B-A, C-D, and F-G (or E-F). Such parallelism results from a combination of capture-awareness and reduced interference. DMAC [3] (and its variants) support any two out of these three dialogs in parallel, while omnidirectional protocols support only one.

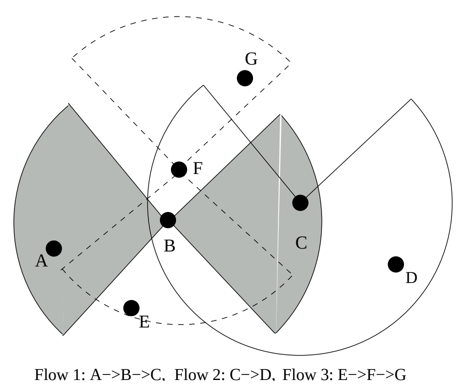

Fig. 7. An example showing the reduction in spatial reuse due to the larger interference footprint in CaMAC

• *Extending Directional Virtual Carrier Sensing (DVCS)* Observe that CaDMAC (as well as CaMAC) may not be able to eliminate capture completely. Beams that receive both productive traffic and capture traffic cannot be turned off during the OFF-Duration – these beams will therefore be prone to capture. However, CaDMAC can reduce the impact of such capture by utilizing directional virtual carrier sensing. On overhearing a RTS, CTS, or a DATA packet that is not meant for it, a node can turn off the corresponding beam for the proposed duration of that dialog[2]. During this time period, the node might be able to remain idle and thus might participate in useful communication. Directional protocols like DMAC [3], Rotational-MAC [21], Circular-MAC [22], and several others, do not incorporate this extension of DVCS. In these protocols, even if a node learns that it will be captured in the near future, it does not take measures to prevent the possibility[3]. As a result, in most cases the node gets captured by the imminent communication. Simulation results presented in Section VI reflect the impact of these factors.

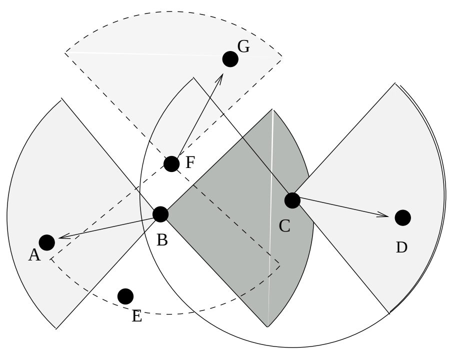

Fig. 8. The possibility of maximal spatial reuse with CaDMAC

D. Deafness in CaDMAC

Although CaDMAC can mitigate capture, and improve spatial reuse, it suffers from the problem of deafness. Deafness has been identified as a problem with directional communication [11]. Referring to Figure 8, while node F communicates with node G using a single conical beam pointed toward G, E will remain unaware of F's activity. In the meantime if node E initiates transmission to F, it will experience deafness because F will not receive these packets. Deafness has been addressed by several authors in the past [11],[22],[27],[28],etc. and ideas from these works could be used in conjunction with our CaDMAC protocol. However, since our focus here is on capture, for simplicity we do not incorporate the solutions for deafness in CaDMAC. Instead we intend to quantify the benefits of spatial reuse and capture, which together dominate the achievable performance.

[2]The proposed duration of a dialog is the remaining duration for which the dialog will continue, and is included appropriately in the RTS, CTS, and DATA packets.

[3]An idle node continues to remain in the omnidirectional mode. A node that has packets to send, does backoff in the omnidirectional mode to avoid the problem of deafness [11].

E. Early Destination Detection (EDD): An Optimization

With the above CaDMAC protocol, a node avoids capture in two time-scales – (i) at the time-scale of OFF-Durations, by choosing suitable multi-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)*. Observe that while a packet reception is in progress, if the MAC layer can quickly detect that the packet is not meant for it, then it might be possible to preempt reception of the rest of the packet. The beam-of-arrival for this packet can be turned off for the proposed duration of the packet (calculated using information from the packet header), thereby providing the node with even more idle time for its own useful communication. Quicker the *early destination detection*, longer is the idle time obtained from preempting the reception. Moreover, the benefits can be non-negligible when packet sizes are large. We incorporate this optimization into CaDMAC. For this optimization, we need to ensure that the destination node's identifier can be encoded within the initial bits of the packet header so that the decoder can decode this information and pass it to the MAC layer before the rest of the packet is received. This will require suitable encoding to be used for the packet header.

F. Dependence on Routing Protocol

We discussed earlier that using multi-beam patterns during the OFF-Duration impacts route discovery. A node that turns off a capture-prone beam during the OFF-Duration, essentially removes a link from the original network topology. As a result, routes discovered during the OFF-Duration may prove to be suboptimal. In some cases, a route may not be discovered at all, even though one exists in the original topology (during the ON-Duration).

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 needs to be 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 the next ON-Duration, it re-initiates route discovery to find an optimal route, and redirects traffic over it. Latency of route discovery may increase in this case, resulting in a delay-throughput tradeoff. However, if route discovery is not performed too frequently, then this issue may not prove to be serious.

Another way of handling the sub-optimal routing problem could be to turn off beams conservatively. For example, if turning a beam off significantly degrades network connectivity, then such a beam could be permanently kept on. Turning fewer beams off could clearly improve the quality of discovered routes, but at the expense of higher capture. We intend to study this tradeoff as a part of our future work.

In this paper, we adopt the approach of re-initiating a route discovery during the ON-Duration. A reactive routing protocol can be made compatible with our capture-aware MAC protocols, using this modification. However, this approach does not change the behavior of the routing protocol – the routing protocol remains capture-unaware. In the next section, we suggest a new *capture-aware routing protocol* that has the efficacy of choosing routes less affected by capture. Such a routing protocol, when used in conjunction with the proposed MAC protocols, can lead to further improvements in performance.

V. Capture-Aware Routing

In this section, we present a "capture-aware" routing metric, using which selected routes might be less prone to capture. Put differently, routes chosen with an awareness of capture can complement the benefits of CaDMAC. We illustrate this with a simple example in Figure 9. In Figure 9, assume that node A routes packets to node C via node B. 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 one of the following routes: S-A-B-G or S-M-F-G. Since route S-A-B-G shares intermediate nodes with the existing flow on route A-B-C, the capacity available to each of these routes will reduce. Even if the routing protocol chooses node-disjoint routes, like S-M-F-G, communications between nodes A-B and nodes M-F will still capture each other, thereby forcing them to again share the available channel capacity. In such a scenario, choosing route S-E-F-G proves to be significantly beneficial, because transmissions on routes S-E-F-G and A-B-C are free of mutual capture. We present a routing protocol that accounts for such mutual capture between links A-B and M-F, and prefers routes that do not include such capture-prone links.

A. Measuring the cost of links

We propose a routing metric that associates a cost to links that are prone to capture. The proposed metric also prefers routes that share fewer intermediate nodes with other active routes in the network. These two attributes are combined suitably with the hop count of a route to obtain the cost of using that route.

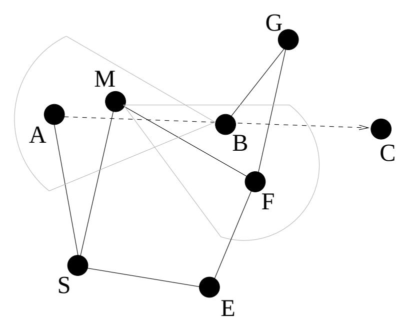

Fig. 9. A scenario in which minimum-hop routing can reduce throughput

- *Sharing active nodes*

We prefer routes that share fewer nodes with other active routes in the network. To achieve this, we require the routing module at a node i to maintain an *indicator variable*, I_i, that denotes whether that node participates in an ongoing communication. Nodes that act as source, intermediate nodes, or sink of a route, set $I = 1$, and are called "active nodes". Other nodes set $I = 0$, and are referred to as "inactive nodes". Indicator variables at active nodes are associated with timeouts. If an active node does not participate in a communication for a threshold period of time, $T_{participate}$, the node resets its indicator variable to 0.

- *Capture-awareness*

A routing protocol that aims to achieve higher throughput, needs to prefer routes on which links are less likely to get captured. In view of this, we associate each link with a cost that indicates the "amount of capture" on that link. Now, observe that each link can be represented by a pair of beams, namely, the beams used by the nodes that form the link. Extending this idea, each route can be represented as a sequence of beams. We quantify the *amount of capture* incident on each of these beams, thereby quantifying the "capture-awareness" of a route. We now describe a simple metric for measuring capture.

Assume that each active node informs its neighbors about the set of beams that it uses for productive communication[4]. If node x receives this information from node y, on beam B_{xy}, then node x infers that any new route that uses beam B_{xy} in future, will be prone to capture from existing transmissions from node y. Thus, node x increases the cost of using beam B_{xy} by some quantity $C_{B_{xy}}$. Notice that node x might use the same beam, B_{xy}, to communicate to another neighbor, z. For example, this may be the case if z and y are closely located. In such a case, B_{xy} and B_{xz} refer to the same beam. The cost of communication between x and z will therefore be affected by the communication between x and y. In general, the *capture-cost* of a link ij, κ_{ij}, can be calculated as

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

where $C_{B_{ij}}$ is the capture-cost incurred by node i for using the beam B_{ij}, and $C_{B_{ji}}$ is the capture-cost incurred by node j for using the beam B_{ji}. Similarly, the cost of a route, R, can be calculated as

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

When using such a metric, links on the least-cost route will typically be less prone to capture from ongoing cross-traffic.

Ideally, the cost of using a beam must be proportional to the amount of capture traffic incident on that beam. In this paper, we use a simpler metric that counts the number of transmitters from which transmitted packets were received on this beam. Clearly, in presence of different sending rates, the cost of capture can be different. As a part of our future work, we intend to study more refined ways of measuring the cost of capture on a beam. For example, while informing neighbors about its active beams, a transmitter may also include information about the total number of packets transmitted on each active beam, within a past window of time. The number of transmitted packets can be a more accurate measure of the cost of capture.

- Hop-Count

It is possible that a route is low in terms of capture and node-sharing, but takes a longer path to reach the destination. Since hop-count of a route also impacts the achievable throughput, we account for it in our metric. To achieve this, every link ij used in a route is associated with a cost H_{ij}. In our simulations, we assign H_{ij} to be 1.

[4]This information need not be explicit messages – we describe later how the information about beams can be included in the route request (RREQ) messages.

B. A combined routing metric

We discussed three factors based on which we intend to choose our routes – active-node-sharing, capture-awareness, and hop-count. A unified routing metric accounts for all these factors. With our unified routing metric, the cost of a route, C_{route}, is formulated as a weighted average of the three factors, as follows.

$$C_{route} = \sum_{ij \in R} (\omega_I (I_i + I_j) + \omega_\kappa \kappa_{ij} + \omega_H H_{ij}) \tag{3}$$

where ω_I, ω_κ, and ω_H are weights associated to node-sharing, capture-awareness, and hop-count, respectively. In our simulations, we use identical weights. However, more sophisticated ways of choosing the weights could potentially lead to larger benefits.

C. Route Discovery

We describe the route discovery mechanisms in this subsection. Our capture-aware routing protocol (CaRP) is a source-initiated routing protocol, similar to DSR [30]. The source node initiates a route request (RREQ) packet which is flooded across the network. Each node uses its current beam to forward the RREQ packet. Observe that the current beam may not be omnidirectional if the network happens to be in the OFF duration at that instant of time.

Consider the case in which a node X receives a RREQ initiated 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. If nodes use N beams, then a tuple of N separate costs are included in the RREQ. For instance, the cost corresponding to beam 3, B_3, at node X, is calculated as $\kappa_P + C_{B_3}$, where P denotes the partial route from node S to X over which this RREQ has arrived.

(2) The cumulative *number of active nodes* on the route along which this RREQ has arrived. The number of active nodes is 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.

When node Y receives the RREQ, it deduces the capture-cost on the partial route from from S to Y, using the information available in the RREQ. For example, if node X uses beam 5 to communicate to Y, then Y selects the 5^{th} 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[5]. The result represents the capture-cost for the partial route from S to Y.
If node Y happens to be already participating on another active route, then it increments the active-node-sharing cost, i.e., A_{sy} = $A_{sx} + 1$. Then, node Y calculates the hop count, say h, 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, C_{sy}, using equation (3). Node Y remembers the cost of this partial route, for reasons explained later. Node Y then updates the RREQ with the required information, and forwards it.
When the destination node receives the RREQ, it calculates the cost of the route using equation (3). 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. While this may be reasonable for discovering minimum-hop routes, it may not be applicable for our purposes. With our proposed 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.

When implementing DSR and CaRP over CaDMAC, we did not use caching, and other optimizations (such as ring-zero search, etc.). As a result, when performing route discovery, intermediate nodes did not transmit gratuitous replies from their caches. Moreover, when a route was successfully discovered during the ON-Duration, we forced DSR and CaRP to continue using it for the entire duration of that flow. If a route was discovered during the OFF-Duration, then route discovery was re-initiated during the next ON-Duration, and thereafter the new route was used for the rest of the durations. In rare cases, route errors (RERR) were triggered by intermediate nodes if a link experienced short term fading. For simplicity of implementation, we overlooked RERR messages, and persisted with the chosen routes – since nodes are not mobile, the occurrence of route errors were rare. Our ongoing work is currently focused on a more general and accurate implementation.

[5]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.

D. *Issues with capture-aware routing*

Since several source nodes may initiate (pending) route discoveries right at the beginning of an ON-Duration, they can get synchronized. Observe that synchronization can affect the quality of selected routes. We reduce synchronization using jitters – the pending route discovery is initiated at a time, t, chosen randomly from the ON-Duration. Of course, t is chosen conservatively, so that the initiated route discovery can finish before the node switches to the OFF-Duration.

VI. Performance Evaluation

We use the Qualnet simulator [32], version 3.6, for simulating our proposed protocols, and comparing them with 802.11 and another directional antenna protocol (DMAC [3] with DVCS [4]). 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 antenna beams used for our simulations are characterized with sidelobes and backlobes, obtained from realistic antenna patterns [9]. We evaluate the performance of our protocol for varying beamwidth. For our *Early Destination Detection (EDD)* optimization, we assume that the destination field of the packet is available to the MAC layer at the time when the PLCP and the physical layer header has been completely received.

A. *Simple Scenarios*

We begin by using two simple scenarios shown in Figure 10. In Figure 10(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, and the beamwidth is specified to be

90°. We measure the end to end throughput for different CBR sending rates. The results are shown in Figure 11.

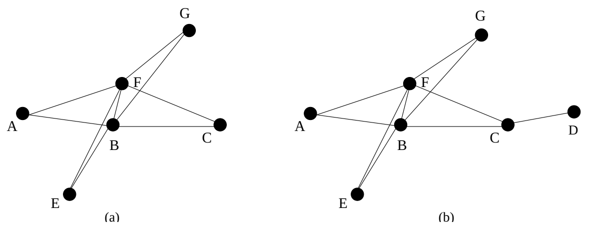

Fig. 10. Simple example scenarios.

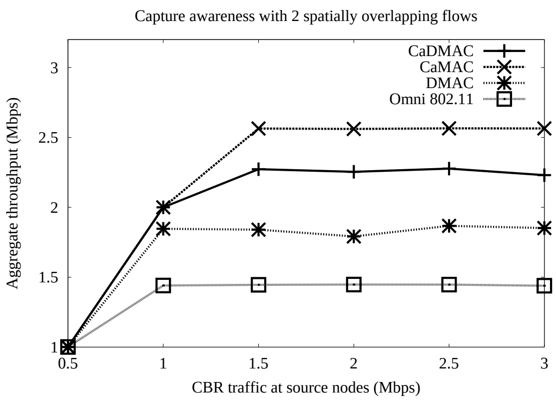

Fig. 11. Spatially filtering capture-causing signals lead to higher spatial reuse.

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. This is because node B often gets captured by communications between nodes F-G or nodes E-F. Observe that node F also gets frequently captured by transmissions from nodes A and C. Capture limits the available spatial reuse in DMAC, which translates into lower aggregate throughput. Observe that both CaDMAC and DMAC suffer from deafness, implying that the difference between the CaDMAC and DMAC graphs is an indicator of the benefits from capture avoidance. CaMAC outperforms CaDMAC only in this simple scenario. As discussed earlier, this is because, in the absence of other flows in the vicinity, CaMAC's large interference footprint does not prove to be harmful. In fact, use of the multi-beam pattern aids the performance of CaMAC by eliminating the impact of deafness. Since CaDMAC is affected by deafness, CaMAC outperforms CaDMAC for this simple scenario.

We evaluate the scenario in 10(b), in which a single hop flow is added between node C and node D. This scenario is a simple instance that reflects the deficiencies of CaMAC. The results for this experiment are shown in Figure 12.

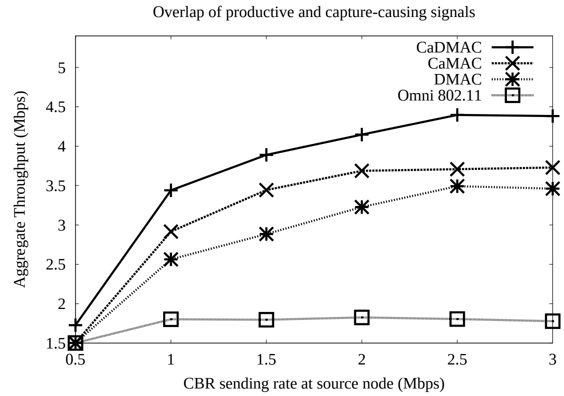

Fig. 12. Larger interference footprint of CaMAC reduces throughput in comparison to CaDMAC.

Since CaMAC uses a fixed multi-beam pattern over the entire OFF-Duration, node B gets captured by the communications between C and D. Observe that with CaMAC, node B cannot avoid this capture because both the capture-causing signals and the productive signals arrive on the same beam (i.e., on the beam toward node C). Spatial reuse reduces as a result, which is reflected in lower aggregate throughput. However, CaMAC can still outperform DMAC implying that the benefits of avoiding capture is significant. CaDMAC outperforms CaMAC because it can often support 3 simultaneous communications, namely communications between A-B, C-D, and E-F (or F-G).

B. Large Random Scenarios

We simulated large networks, with 50 nodes randomly placed in a square region of side 1500m. We selected 15 random pairs of source and destination nodes, 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 13 reflect only the benefits achievable from a capture-aware MAC protocol. 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.

As evident from Figure 13, CaDMAC outperforms CaMAC, DMAC and 802.11. The EDD optimization with CaDMAC improves performance slightly. This is because the packet size used for this measurement was 512 Bytes. Recall that the benefits of EDD is higher for larger packets. In our simulation of the same experiments with 1024 Byte packets (results not included in this paper), we observed that the EDD optimization improves on CaDMAC by a larger margin (approximately 18%). Although CaDMAC clearly outperforms DMAC, CaMAC is only marginally better than DMAC. In networks where several flows are

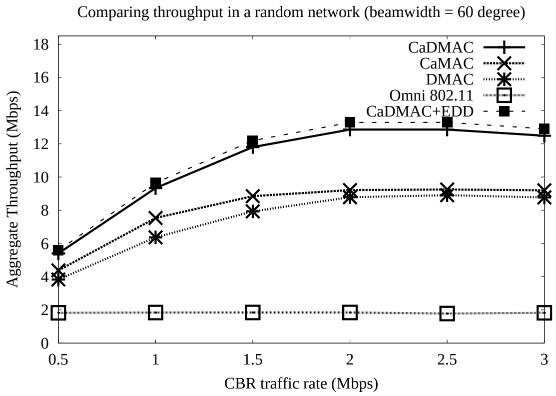

Fig. 13. Throughput benefits from capture-aware MAC protocols in large random networks, with beamwidth of 60°

contending for space, the large interference footprint of CaMAC reduces spatial reuse, offsetting the benefits of capture-awareness. IEEE 802.11 with omnidirectional antennas evidently achieves the least throughput.

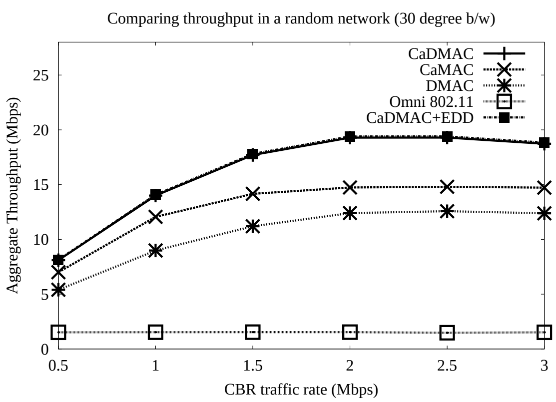

Fig. 14. Throughput benefits from capture-aware MAC protocols in large random networks, with beamwidth of 30°

Figure 14 shows the results of repeating the above experiments, but with a beamwidth of 30°. CaMAC noticeably shows improvement because with a smaller beamwidth it is able to filter out capture with better granularity. More precisely, with higher beamwidths, a productive beam is more likely to be susceptible to capture-causing signals. Since these beams cannot be turned off, nodes lose spatial reuse. As the beamwidths become narrower, capture can be filtered out with more efficacy. While DMAC also uses a smaller beamwidth, its inability to avoid capture limits the improvement it can obtain. CaDMAC obtains benefits from both smaller beamwidth and capture-awareness, thereby achieving higher aggregate throughput in networks with multi-hop flows. However, at lower beamwidths, benefits from EDD are negligible – the graph for EDD is not visible due to the overlap with CaDMAC.

Delay in Random Networks

Figure 15 presents the improvements in 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 with our proposed capture-aware MAC protocols, as evident from the graphs.

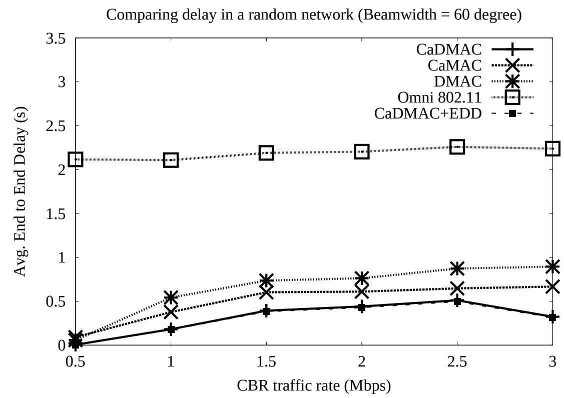

Fig. 15. Latency benefits from capture-aware MAC protocols in large random networks, with beamwidth of 60°

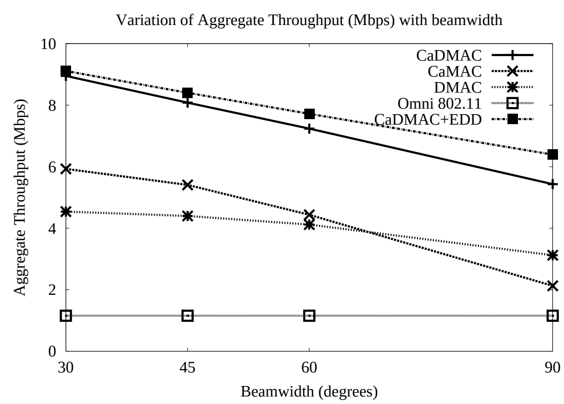

Fig. 16. Throughput variation with variation in beamwidth.

To understand the variation of throughput with increasing beamwidth, we simulated random scenarios using beamwidths of 30°, 45°, 60°, and 90°. The results from Figure 16 show that with increasing beamwidth, the performance degrades for all the protocols. This is a result of lower spatial reuse with higher beamwidths. However, observe that the degradation in CaMAC is faster than others – as discussed earlier, the large interference footprint of CaMAC limits its performance, and offsets its benefits of capture-awareness. It is also worth noting that at higher beamwidths, the benefits of the EDD optimization is more conspicuous, in comparison to CaDMAC. As explained earlier, when the beamwidths are higher, both productive

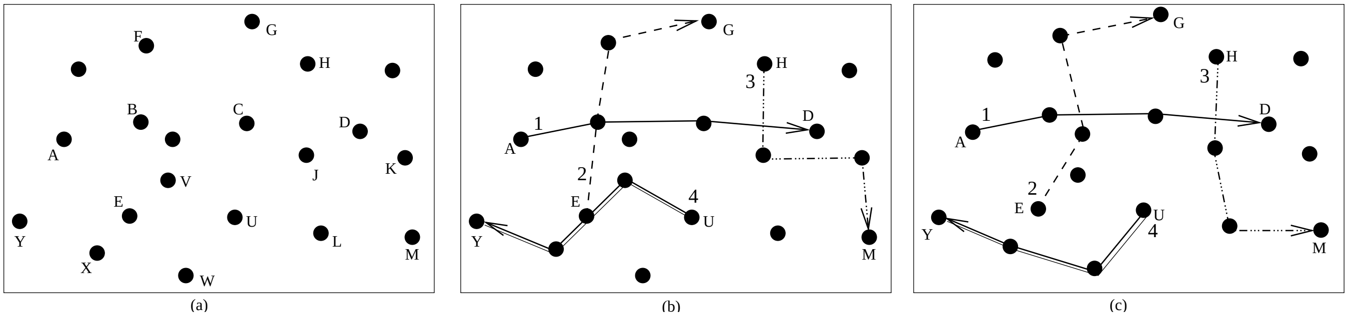

Fig. 17. Enumeration of routes when using DSR and our capture-aware routing protocol, on a small topology with 4 ordered flows.

and capture-causing signals are more likely to impinge on the same beam. As a result, nodes become more prone to capture. In such cases, the EDD optimization helps in freeing up nodes for longer time durations to continue with useful communication.

C. Capture-Aware Routing Protocol (CaRP)

The results presented till this point reflect the benefits of capture-awareness at the MAC layer. In the simulations so far, routes were selected arbitrarily, without attention to the capture-aware metric proposed in Section V. In this section, we present initial results on the combined performance of our MAC and routing protocols. We aim to show that when CaDMAC is used in conjunction with our capture-aware routing metric, performance can be superior.

We begin with a simple scenario and show how capture-awareness in routing can improve achievable throughput. We consider the scenario in Figure 17, where a few nodes are placed in a rectangle. In this simple topology, we introduce four flows with source and destination as the following. (1) From A to D, (2) from E to G, (3) from H to M, and (4) from U to Y. The starting time of the four flows are in the order in which they are enumerated above. However, all the flows end together at the end of the simulation experiment. To measure the best possible benefits, we forcibly initiated all the flows only during ON-Durations. More precisely, flow i is initiated in the i^{th} ON-Duration, so that the routing protocols can always choose the optimal routes.

Figure 17 (b) and (c) show the routes chosen by DSR and CaRP respectively (both use CaDMAC as the MAC protocol). Recall that CaRP is a source-initiated protocol similar to DSR, but enhanced with a capture-aware routing metric as discussed earlier in Section V. As evident from the figure, flows 2, 3, and 4 choose different routes from DSR. The aggregate throughput, as a result of using these routes, is shown in Figure 18. Clearly, CaRP outperforms DSR, indicating the advantages of choosing capture-aware routes, even in a moderately congested network.

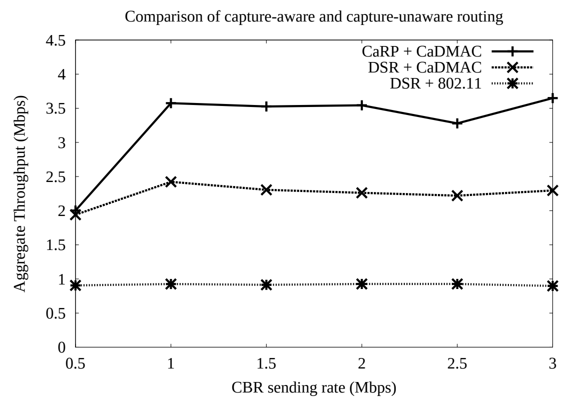

Fig. 18. Capture-aware routes, chosen by CaRP, offers higher aggregate throughput.

The properties of the chosen routes in Figure 17 explain the differences in achieved throughput. Observe that for flow 1, the routes chosen by DSR and CaRP are identical. This is natural, because in the absence of any other spatially overlapping traffic, the metrics of DSR and CaRP are influenced only by the hop count of the discovered route. However, when flow 2 is introduced, CaRP avoids choosing a route through node B, and thereby finds a route that is node-disjoint from route 1. Thus, CaRP can permit both flows 1 and 2 to proceed in parallel. On the other hand, node B becomes the bottleneck for DSR.

Later, when flow 3 is introduced, DSR chooses a route that goes through nodes J and K. Observe that the links J-K and link C-D cannot operate in parallel, and therefore they are prone to capture by each other's traffic. However, CaRP recognizes the possibility of capture on link J-K, and therefore chooses a route that goes through node L – note that link J-L does not share the channel with link C-D, and can operate parallelly. Spatial reuse increases with CaRP.

Finally, when flow 4 is introduced, CaRP again selects a route less prone to capture. As evident from Figure 17(c), 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-B. The benefits of discovering capture-aware routes lead to significant benefits with CaRP. The relative benefits can be larger when the network is even more congested.

The previous results were simulated for data flows that were initiated forcibly during the ON-Duration. This allowed us to measure the best-case benefits available from capture-aware routing. However, in reality, source nodes may generate traffic at any point of time – route discoveries may very well be initiated during the OFF durations as well. As discussed earlier, traffic initiated during the OFF-Duration may cause CaRP to choose suboptimal routes. The next experiment takes into account such possibilities, and initiates traffic at random points of time. We compare the performance of DSR with omnidirectional antennas, DSR with CaDMAC, and CaRP with CaDMAC.

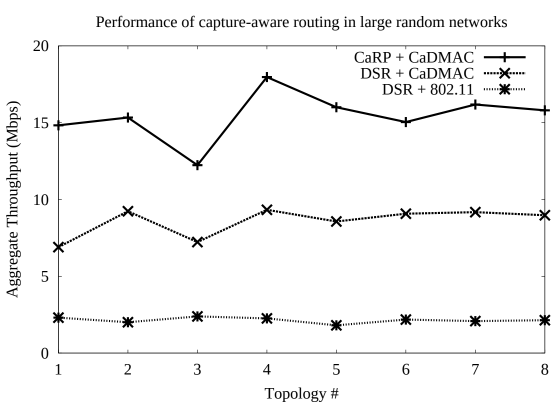

Fig. 19. Capture-aware routing in large random networks.

For this experiment, we randomly placed 50 nodes within a square of size 2000m, and initiated CBR traffic between randomly chosen source and destination nodes. We generated 15 multihop flows, and the starting time of the flows were randomly chosen as well (without accounting for the ON and OFF durations). We simulated our experiments for various topologies. Observe that when using omnidirectional antennas, the routes 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 thereby suboptimal routes are used for the corresponding OFF-Duration. Route discovery is re-initiated in the next ON-Duration, and traffic is redirected if a different route proves to be better. The ON and OFF durations for these experiments are assigned to be 1 and 3 seconds respectively. The beamwidth for this experiment was 60°.

Figure 19 shows the results for 8 different topologies. Observe that even though both (DSR+CaDMAC) and (CaRP+CaDMAC) outperform (DSR+802.11), in some topologies the margin of performance benefits is slightly lower. This is attributed to the use of a suboptimal route for a large part of the OFF-Duration. However, the achieved throughput in all the random topologies are always higher than (DSR+802.11). The benefits with CaRP are particularly encouraging. This is because, even though CaRP may use sub-optimal routes, it compensates for the losses by choosing a new route that has the capability to support higher throughput. Moreover, in moderately dense networks, chosen routes may not be severely suboptimal because of the availability of multiple nodes that may serve as alternates. Although our evaluations in this case are preliminary, it is indicative of possible benefits of capture-aware routing. With narrower beamwidth, we believe that the improvements can be higher.

D. *Ongoing Work*

Our evaluation of CaRP is preliminary. We are currently evaluating the performance of CaRP more exhaustively, in various scenarios and with varying parameters. In particular, we are evaluating the impact of ON and OFF durations. Clearly, longer OFF durations can reduce the impact of capture, but at the cost of using suboptimal routes. The behavior of this tradeoff is not evident from our current results. Insights into this behavior will also gain us insights into the latency trends experienced by CaRP.

We are also evaluating the impact of beam-width on the performance of CaRP. Narrower beamwidths are beneficial for our MAC protocol, CaDMAC. We believe that with CaRP, narrower beamwidth may prove to be even more beneficial.

VII. Discussion

With our capture-aware routing protocol, a route discovered by node S depends on other active routes in the network, present at that point of time. In other words, the cost of a route is a function of the traffic pattern in the network. If the traffic pattern changes between time instants t_1 and t_2, a route selected before t_1 may not be the best-possible route after time t_2. As a result of this dependence, spatial reuse of the channel may not be maximal, even with CaRP. One way to address this could be to have nodes periodically rediscover routes during the ON-Duration. We intend to investigate this issue in detail, and evaluate possibilities to further improve the achievable performance.

Our proposal on *Early Destination Detection (EDD)* may be applicable to existing (capture-unaware) MAC protocols as well. For example, in DMAC, if an idle node detects that a packet is not meant for it, the node can form a null toward the direction-of-arrival of the signal. In such cases, benefits can be obtained at the MAC layer, without incorporating ON/OFF durations. The benefits from EDD can be higher when used with DMAC than with CaDMAC, because in DMAC idle nodes will be in the omnidirectional mode, and thus would be frequently susceptible to capture.

In mobile ad hoc networks, DSR utilizes the option of overhearing route reply packets, thereby learning routes to other nodes without having to discover them. Capture-aware protocols reduce the opportunity to overhear packets, reducing the benefits of promiscuous listening. However, in applications such as wireless mesh networks, nodes remain almost stationary – the need for promiscuous listening may not be severe in such cases.

VIII. Conclusion

We identified the opportunity to improve performance by mitigating the phenomenon of capture. We use multi-beam antennas, and presented a basic capture-aware MAC protocol that spatially filters directions from which signals cause capture. We enhanced the basic protocol by having transmitters and receivers use only single beams during communication, and switch to multi-beams when they are idle. To further reduce capture, we enhanced our protocol with an *Early Destination Detection* optimization, with which nodes quickly decode the destination of a packet, and preempts reception if the packet is not destined for it. To assist the MAC layer on a larger time scale, we also proposed a routing protocol. By using a capture-aware routing metric, our protocol aims to choose routes on which links are less prone to capture. We show that network performance can be reasonably improved, when a capture-aware MAC and routing protocol are used in conjunction.

References

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] K. Sundaresan, R. Sivakumar, M. Ingram, and T. Chang, "Medium access control in ad hoc networks with mimo links: Optimization considerations and algorithms," Tech. Rep., 2004.

[6] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit, "Ad hoc networking with directional antennas: A complete system solution," in *To Appear in Journal of Selected Areas in Communications*, Jan 2005.

[7] A. Kochut, A. Vasan, A. UdayaShankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," .

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

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

[10] Motia. Inc, "Motia's javelin enhances wifi," Tech. Rep., http://www.motia.com/javelin.htm, 2004.

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

[12] 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.

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

[14] M. Streenstrup, "Neighbor discovery among mobile nodes equipped with smart antennas," in *Porceedings of ADHOC Networks*, May 2003.

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

[16] A. Spyropoulos and C. Raghavendra, "Energy efficient communications in ad hoc networks using directional antennas," in *Proceedings of Infocom*, 2002.

[17] 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.

[18] A. Nasipuri, J. Mandava, H. Manchala, and R. Hiromoto, "On demand routing using directional antennas in mobile ad hoc networks," in *Proceedings of IEEE ICCCN*, September 2000.

[19] 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.

[20] L. Bao and J. Garcia-Luna-Aceves, "Transmission scheduling in ad hoc networks with directional antennas," in *Proceedings of Mobicom*, 2002.

[21] 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.

[22] 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.

[23] R. Ramanathan, "On the performance of ad hoc networks with

beamforming antennas," in *Proceedings of ACM MobiHoc*, October 2001, pp. 95–105.

[24] K. Wang, J. Proackis, and R. Rao, "Energy-efficient routing algorithms using directional antennas for mobile ad hoc networks," April 2002.

[25] J. Cartigny, D. Simplot, and I. Stojmenovic, "Localized energy efficient broadcast for wireless networks with directional antennas," Tech. Rep., Univerity of Sciences and Technologies de Lille, July 2002.

[26] T. Rappaport, "Wireless communications principles and practice," *Prentice Hall*, 2002.

[27] K. Sudaresan and R. Sivakumar, "On the medium access control problem in ad-hoc networks with smart antennas," in *Proceedings of MobiHoc*, June 2003.

[28] H. Gossain, C. Cordeiro, D. Cavalcanti, and D. Agrawal, "The deafness problems and solutions in wireless ad hoc networks using directional antennas," in *Proceedings of Globecom*, 2004 November.

[29] 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.

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

[31] S. Das, C. Perkins, and E. Royer, "Ad hoc on demand distance vector (AODV) routing," Internet-Draft Version 4, IETF, October 1999.

[32] Scalable Network Technologies, "Qualnet simulator version 2.6.1," www.scalable-networks.com.