Power-controlled Medium Access for Ad Hoc Networks with Directional Antennas

Aman Arora and Marwan Krunz
Department of Electrical and Computer Engineering
The University of Arizona, Tucson, AZ 85721
{amarora,krunz}@ece.arizona.edu

***Abstract*— Directional antennas have the potential to significantly improve the throughput of a wireless ad hoc network. At the same time, energy consumption can be considerably reduced if the network implements per-packet transmission power control. Typical MAC protocols for ad hoc networks (e.g., the IEEE 802.11 Ad Hoc mode) were designed for wireless devices with omnidirectional antennas. When used with directional antennas, such protocols suffer from several medium access problems, including interference from minor lobes and hidden-terminal problems, which prevent full exploitation of the potential of directional antennas. In this paper, we propose a power-controlled MAC protocol for directional antennas that ameliorates these problems. Our protocol allows for dynamic adjustment of the transmission power for both data and clear-to-send (CTS) packets to optimize energy consumption. It provides a mechanism for permitting interference-limited concurrent transmissions and choosing the appropriate tradeoff between throughput and energy consumption. The protocol enables nodes to implement load control in a distributed manner, whereby the total interference in the neighborhood of a receiver is upper-bounded. Simulation results demonstrate that the combined gain from concurrent transmissions using directional antennas and power control results in significant improvement in network throughput and considerable reduction in energy consumption.**

***Index Terms*— Ad hoc networks, medium access, directional antennas, power control, interference-limited transmissions.**

I. Introduction

A key issue in designing wireless mobile ad hoc networks (MANETs) is how to improve network throughput (i.e., increase the spatial reuse). Network throughput can be improved by employing directional antennas [13]. Besides their throughput gain, directional antennas provide wider coverage and lower power consumption. Because of these advantages, directional antennas have been adopted in IS-95 and 3G cellular systems [13]. For instance, sectoring provided by directional antennas enables a base station (BS) to serve more than one cell at a time, thus improving the capacity of the cellular network (in a typical configuration of 120° sectoring, i.e., 3 cells per BS, directional antennas provide a power gain of 18 dBi[1], which translates into increased coverage).

This work was supported in part by the National Science Foundation through grants ANI-0095626, ANI-0313234, and ANI-0325979. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

[1]The relative gain of an antenna system X in dBi is given by $10\log_{10}\frac{P_X}{P_i}$, where P_X is the power received from antenna X at a some reference point and P_i is the power received from an isotropic antenna at the same point, provided both antennas are fed with equal transmission powers.

Classic MAC protocols for MANETs (e.g., the IEEE 802.11b Ah Hoc scheme [2]) were not designed for use with directional antennas [19]. Such protocols assume that nodes have equal reception sensitivity and radiate equal powers in all directions. This is an underlying assumption in the request-to-send/clear-to-send (RTS/CTS) exchange mechanism that is used for collision avoidance in the IEEE 802.11b scheme. So if a node can cause interference at a receiver, then this node will likely hear the CTS from that receiver and will defer from transmitting. When directional antennas are used, the radiated power and reception sensitivity between two nodes are functions of these nodes' angular orientation. Thus, using the same power for both RTS/CTS and data packets can no longer prevent potential interferers from transmitting.

In this paper, we propose a novel MAC protocol called LCAP (*load-based concurrent access protocol*) for MANETs with directional antennas. LCAP's novelty lies in using an elaborate packet-based power control strategy that is aimed at increasing the channel's spatial reuse by allowing interference-limited, concurrent directional transmissions to take place in the same vicinity. By employing a separate control channel and by accounting for minor-lobe interference, LCAP alleviates many of the channel access problems that afflict previously proposed MAC protocols for MANETs with directional antennas.

According to LCAP, an idle node listens to the channel omnidirectionally and continuously measures the total interference-plus-noise power. A transmitting node sends an omnidirectional RTS at a fixed (maximum) power. Any node that receives this RTS estimates the path loss between itself and the transmitter along with the angle of arrival (AOA) of the transmitter's signal. If the intended receiver wishes to accept the data packet, it beamforms its antenna in the transmitter's direction and responds back with a directional CTS (DCTS). The transmission power of this DCTS is properly scaled to reach a set of *expected* interferers. Subsequently, the transmitter beamforms in the receiver's direction and sends its data packet at an adjustable power, whose value is computed by the intended receiver and included in the DCTS packet. In determining this value, the intended receiver uses *load control* to strike a balance between energy consumption and spatial reuse. The computed power is larger than the minimum power needed for correct packet reception; the difference is used as an *interference margin* to allow for future interfering transmissions to take place in the same vicinity.

The remainder of the paper is organized as follows. The next section discusses some of the fundamentals of directional antennas and directional virtual carrier sensing. Section III reviews previous work on MAC protocols for directional antennas. In Section IV we discuss the problems associated with these protocols. Section V presents the proposed LCAP. In Section VI we discuss how LCAP solves various channel access problems. Performance evaluation of LCAP is provided in Section VII. Finally, we conclude the paper in Section VIII.

II. PRELIMINARIES

A. Directional Antennas

In contrast to an isotropic antenna, which transmits the same amount of power is all directions, a directional antenna has preferred direction(s) for transmission and reception; while transmitting, the antenna concentrates the power in certain direction(s), and while receiving the antenna has a greater sensitivity for electromagnetic radiation in certain direction(s). The relative gain of a directional antenna is typically plotted as a function of direction (angle in azimuth or vertical plane). This plot is called the *radiation pattern* of the antenna. Figure 1 depicts the radiation pattern for a typical directional antenna (a six-element circular array) while beamformed in 90-degree azimuth. The peaks in the radiation pattern are the result of concentrating the power in certain directions. The peak with the maximum gain is called the *major (or main) lobe*. Peaks other than the major lobe are called *minor lobes* (they include *side* and *back lobes*). These minor lobes represent the power radiated/received in directions other than intended, and though undesirable, cannot be completely eliminated.

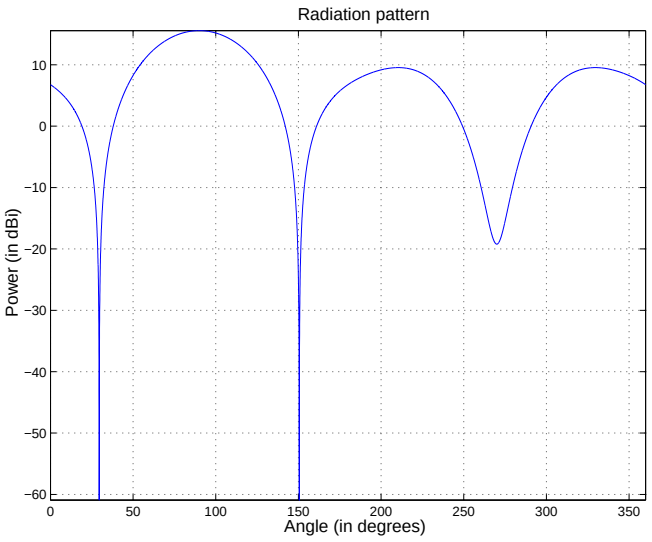

Fig. 1. Radiation pattern for a six-element circular array directional antenna.

B. Directional Virtual Carrier Sensing

Many channel access schemes, including the one used in the 802.11b standard, are based on carrier-sense multiple access with collision avoidance (CSMA/CA), with an optional virtual carrier sense (VCS) mechanism for large packets. In VCS, if a node wishes to send a data packet to another node, it first broadcasts an RTS packet. The intended receiver node replies back with a CTS packet, indicating its willingness to accept the sender's data packet. The data-packet transmission can then proceed, followed by an ACK transmission (if the protocol supports link-layer reliability). Upon overhearing an RTS or a CTS packet, a node in the vicinity of a transmitter and/or a receiver sets its NAV (network allocation vector) for the duration of the data and ACK transmissions. Nodes desist from transmitting until their NAVs expire.

Under omnidirectional antennas, the classic CSMA/CA with VCS mechanism can be quite conservative in its allocation of the channel's spatial capacity. Consider the example in Figure 2, where node C intends to transmit to node D. To reserve the channel, node C sends an RTS, which is received by nodes B and D. Node B sets its NAV to a value that corresponds to the time at which node D completes its ACK transmission. It is easy to see that the silencing of node B during the transmission of the data packet from C to D is unnecessary, as B's signal does not interfere with the received signal at node D. In the literature, this problem is known as the *exposed terminal problem*. Even during the transmission of the ACK from D to C, node B may, in principle, transmit concurrently at an adjusted power that does not degrade the signal-to-interference-and-noise ratio (SINR) at D below a given threshold SINR_{th}.

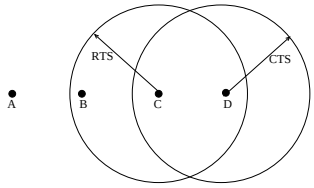

Fig. 2. Exposed terminal problem in MANETs with omnidirectional antennas.

The reduction in network throughput in the previous example can be addressed by using directional transmission and directional virtual carrier sensing (DVCS) [19]. DVCS is an extension of VCS, where upon receiving an RTS or a CTS, a node sets a directional NAV (DNAV), associating with it a duration, a direction, and an angular width. Figure 3 depicts an example of DVCS, where node B sets its DNAV after it overhears a CTS from node C (i.e., node C reserves a circular sector of a certain radius in B's direction). The angular width of this reserved sector depends on the width of the main lobe. Note that node B is free to transmit in the non-reserved directions.

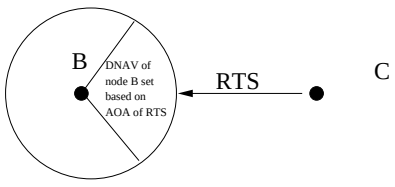

Fig. 3. DNAV of node B after overhearing a CTS from node C (based on the network topology in Figure 2) .

The following example demonstrates the potential improvement in spatial reuse that can be achieved using directional antennas. In Figure 4, node C is beamforming in the direction of node D and is sending an RTS. Due to the directional nature of the transmission, the RTS is not heard by node B. If node B hears a CTS from node D, it will set its DNAV in the direction of node D only. Node B can now transmit in node A's direction without disturbing the ongoing transmission between C and D (in this example, we ignore the effect of the minor lobes).

The advantages of using directional antennas in MANETs are evident, but their usage places additional requirements on the network and presents new channel access problems, as explained in Section IV.

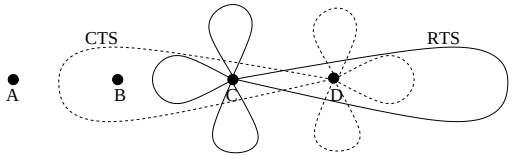

Fig. 4. Improving spatial reuse using directional antennas.

III. Related Work

The use of directional antennas in cellular and multihop packet radio networks has been extensively studied in the literature (see [18] [13] and the references therein). However, designing MAC protocols for MANETs that use directional antennas is a rather recent topic, and has been researched in only a handful of papers (e.g., [11] [24] [19] [5] [3] [15] [22] [4]). The authors in [11] proposed using location tables to keep track of directions via which a node can communicate with its neighbors. The RTS is sent directionally, successively in a circular manner to locate the intended receiver. In [19] the authors proposed a protocol that employs directional antennas and extends the concept of VCS to DVCS. AOA estimates are used to set the DNAV for a receiving node. The authors also demonstrated the adverse effects of minor lobes on the performance of their protocol. In [24] the authors proved that using directional transmission for control and data packets achieves the highest throughput among all combinations of omnidirectional and directional transmissions. However, they ignored the overhead associated with keeping track of neighbors when all control and data packets are sent directionally. Moreover, an oversimplified antenna pattern was used (side lobes were not considered). In [5] the authors proposed two MAC protocols. The first is somewhat similar to the protocol in [19]. The other protocol exploits the extended range of directional antennas, wherein a *multi-hop* RTS mechanism is used to beamform two far-off nodes in each other's direction before data transmission. In [16] the authors presented a scheme called UDAAN for ad hoc networks with directional antennas. UDAAN relies on an extra piece of hardware (called the Inertial Management Unit) to provide geo-position and orientation information. It implements an elaborate backoff procedure (called *forced idle*) for contention resolution following a collision, whereby the duration and window-adjustment mechanism of the forced idle period depend on the type of collision (control or data). Directional carrier sensing is used in UDAAN and is enabled by upper-layer position information (e.g., routing layer). A simple power control approach is adopted, according to which the transmission power for the RTS packet is successively increased upon each RTS retry. Field demonstrations of the protocol operation were performed using mobile vehicles equipped with custom radio boards.

Following the seminal work of Gupta and Kumar [9], which analytically demonstrated the potential of power control as a means of increasing network capacity, several researchers investigated the integration of power control in MAC protocols for MANETs (e.g., [14] [17] [20] [7]). These protocols, however, were designed for omnidirectional antennas. In [8] the authors proposed a power control technique for MANETs with directional antennas. According to this technique, SINR estimates are exchanged between nodes, and a "power reduction factor" is computed and updated iteratively. Although this protocol achieves some improvement in throughput, it does not allow for concurrent, interference-limited transmissions, which can significantly improve the spatial reuse. Concurrent, interference-limited transmissions were implemented in the PCDC protocol [14] for omnidirectional antennas. To the best of our knowledge, LCAP is the first protocol that uses power control to enable concurrent, interference-limited transmissions under directional antennas.

IV. Limitations of Previous MAC Protocols for Directional Antennas

In this section, we discuss through examples several channel access problems that arise when directional antennas are deployed in MANETs. Some of these problems can degrade the performance of previously proposed power-controlled MAC schemes for MANETs with directional antennas. Throughout the examples, we assume that idle listening is performed omnidirectionally.

A. Interference from Minor Lobes

1) Vulnerable Receiver: Previous MAC protocols for MANETs with directional antennas (e.g., [19], [11]) improve the spatial reuse by allowing concurrent transmissions to take place in the same vicinity. Such transmissions are permitted provided that nodes that intend to send data packets point their main lobes away from nodes that are already in the process of receiving data packets. Nodes keep track of the prohibited directions using various methods, e.g., by setting the DNAV [19], [3] or by location tables [11]. However, practical directional antennas have minor lobes, and the radiation from these minor lobes is significant. For example, for the six-element circular array shown in Figure 1, minor lobes have a peak gain of 10 dBi, i.e., the power radiated in the minor lobe direction is ten times greater than the power radiated from an isotropic antenna. Thus, a receiver that lies in the direction of a minor lobe will experience considerable interference, which may lead to packet collisions, as illustrated in Figure 5. In this figure, node A is sending data to node C after an RTS/CTS

exchange[2]. According to the protocols in [19], [3], [11], node D is free to transmit as long as it does not beamform in the direction of node C. So, node D sends a directional RTS to node B, causing interference at node C from the minor lobe radiation. This problem has not been identified in previous studies.

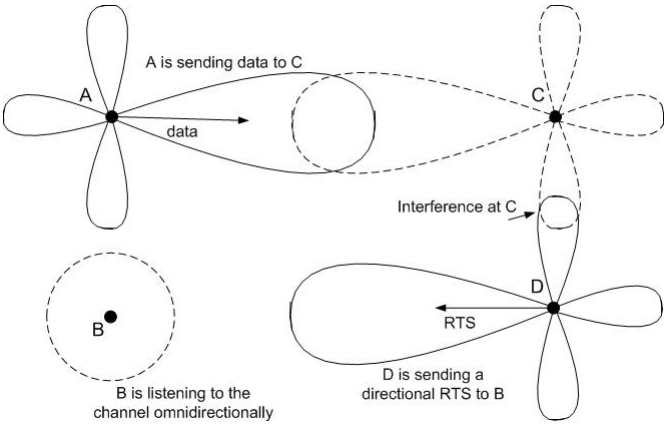

Fig. 5. Minor-lobe interference problem in existing MAC protocols for MANETs with directional antennas.

2) Vulnerable Transmitter: Consider the situation in Figure 6 (this topology will be used to illustrate several medium access problems). Suppose that node F sends a directional RTS (DRTS) to node H, which replies back with a directional CTS (DCTS). Upon receiving this DCTS, node F commences its data transmission. The DRTS/DCTS exchange is heard by node C, which sets its DNAV accordingly. Later on, suppose that node C wants to send a packet to node B. Since C's DNAV in the direction of B is not set, C sends a DRTS to node B and waits for a DCTS while beamformed in node B's direction. However, node F has its main lobe already pointed towards node C, and node C has some minor lobe gain in node F's direction. The interference from node F at node C may be significant. In this case, if node B replies with a DCTS, this DCTS may not be correctly received at node C. In other words, nodes in the vicinity of a transmitter may not be able to initiate data transmissions even in available directions (directions for which the DNAV has not been set) while a transmission is going on. A similar problem arises when a transmitter is waiting for an ACK packet.

B. Hidden Terminal Problems

In a MANET with omnidirectional transmissions, the solution for the hidden terminal problem lies in the exchange of the RTS/CTS packets prior to data transmission. MAC protocols for directional antennas (e.g., [5] [19] [3]) also use an RTS/CTS exchange but with the RTS and/or CTS packets sent directionally. In these protocols, when a node is idle, it listens to the channel omnidirectionally but uses directional transmission for data packets. This approach leads to the following hidden terminal problems [5]:

1) Unheard RTS/CTS Due to a Busy Node: In Figure 6, suppose that nodes A and B are beamformed in each other's direction and are communicating directionally. While this is taking place, nodes C and D proceed with an RTS/CTS exchange followed by a data packet from C to D. Due to the orientation of its antenna, node B is unaware of the ongoing communication between C and D. Suppose that after the end of the $A \leftrightarrow B$ transmission, node B wishes to transmit a packet to node D. It sends an RTS, which collides with the ongoing reception at D. Previous MAC protocols suffer from this problem.

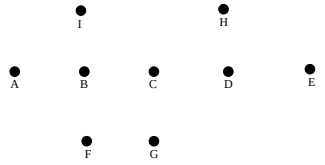

Fig. 6. Topology that is used to demonstrate several medium access problems in existing MAC protocols for MANETs with directional antennas.

2) Unequal Gains in Omni and Directional Modes: The protocols in [19][3] use a fixed transmission power for data and control packets. This could cause a collision, as explained in the following example. Suppose that in Figure 6 node C sends a DRTS to node D, which replies back with a DCTS. The two nodes beamform in each other's direction, and node C starts transmitting to node D. Initially, node A is listening omnidirectionally and is distant from node D, so it does not hear the DCTS. While the $C \leftrightarrow D$ transmission is going on, suppose that node A now has a packet to send to node B, so it sends a DRTS. Because nodes A and D are now beamformed in each other's direction, it is quite possible that the DRTS from node A will interfere with node D . In other words, *any two nodes may be out of range when at least one of them is in the omnidirectional mode, but can come in range if they both beamform in each other's direction, potentially causing a collision*. Choudhary et. al [5] identified this problem but did not provide a solution to it.

V. PROPOSED PROTOCOL

A. Motivation

In contrast to previous protocols, the proposed LCAP allows transmissions to take place along already reserved directions, provided that the SINR at the receiving nodes remains above SINR_{th}. The idea is illustrated in Figure 7, where node A receives a DRTS from node B and responds with a DCTS . According to the protocols in [19] [5] [11], if node D overhears this DCTS, it refrains from sending to node C for the duration of the DNAV. Similarly, if node E overhears the DCTS through the *back lobe* of node A, it is not allowed to communicate with node F. In contrast, LCAP allows the three transmissions $A \leftrightarrow B$, $C \leftrightarrow D$, and $E \leftrightarrow F$ to proceed simultaneously (possibly with adjusted transmission powers),

[2]Under directional transmissions/receptions, it is no longer adequate to indicate a node's ability to correctly receive a packet by including this node within the range of the transmitting node. The reason is that the ability to correctly receive a packet now depends on both the transmitter-receiver distance as well as the relative orientation of the transmitter and receiver antennas.

provided that the SINR at the respective receivers is not below SINR_{th}.

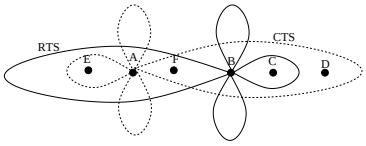

Fig. 7. Example that demonstrates the potential for concurrent, interference-limited transmissions in LCAP.

To enable interference-limited concurrent transmissions in the same neighborhood, a mechanism for allocating *link margins* (explained below) and power control is required. Note that the protocols in [19] [5] [11] do not implement power control; instead, data packets are sent at the maximum possible power, which may be well above the power required to achieve SINR_{th}.

B. Load Control

In power-controlled (interference-limited) wireless communications, in general, there exists a tradeoff between network capacity (throughput) and energy consumption. Energy considerations call for using the minimum transmission power that achieves SINR_{th}. However, by boosting the transmission power beyond this minimum, the receiver can tolerate more interference, which may permit more concurrent transmissions to take place (i.e., increase the spatial capacity)[3]. LCAP uses *load control* to manage this throughput/energy tradeoff.

The concept of load control has been used in CDMA cellular networks (e.g., UMTS) for connection admission purposes [10]. In principle, it can be performed on the basis of either interference or throughput. In interference-based load control, when a new user is to be admitted, the service provider estimates the expected total interference due to the addition of this user. The increase in interference depends on the user's QoS requirements (bit rate, required BER, etc.). The user is admitted only if the total expected interference is below a predefined threshold. In throughput-based load control, the admission decision is made based on the expected total throughput (sum of bit rates) normalized by the maximum allowable throughput. If the total normalized throughput following the admission of the prospective user is expected to exceed a predefined threshold, this user is not admitted.

In our work, load control is interference based. "Admission" is performed on a per-packet basis and is receiver dependent. Formally, let $\eta^{(j)}$ denote the *loading* at node j, defined as:

$$\eta^{(j)} \stackrel{\text{def}}{=} 1 - \frac{N_0}{I_{total}^{(j)}} \tag{1}$$

where N_0 is the thermal-noise power and $I_{total}^{(j)}$ is the total interference-plus-thermal-noise power seen by node j. Note that $0 \leq \eta^{(j)} < 1$. While listening to the channel omnidirectionally, an idle node j continuously measures $I_{total}^{(j)}$, which accounts for all sources of interference plus the thermal noise[4].

Load control ensures that for all nodes j, $\eta^{(j)}$ is not allowed to exceed a predefined *planned loading factor* L_p, where $0 < L_p < 1$. The value of L_p can be set a priori on the basis of the required communication range for the nodes.

In LCAP, rather than controlling the load of a given node, we control the load in the *neighborhood* of that node. The rationale is that an interferer not only affects a node, but also many of its neighbors. Hence, it would make sense to consider the impact of such an interferer on the whole neighborhood. To do that we revise (1), replacing $I_{total}^{(j)}$ with the average of all the $I_{total}^{(.)}$ values in the neighborhood of j, denoted by $I_{avg}^{(j)}$. The resulting loading is denoted by $\eta_{avg}^{(j)}$. Hence, we require that $\eta_{avg}^{(j)} \leq L_p$ for all nodes j. Note that because $I_{avg}^{(j)}$ is an average quantity, it is possible to have $\eta^{(i)} > L_p$ for some node i. However, this is not a major concern, as the purpose of load control is to limit the *average* energy consumption in the network.

Every time a node sends an RTS packet, it includes in it its latest $I_{total}^{(.)}$ value. To compute $\eta_{avg}^{(j)}$, node j maintains a cache of *active* neighbors[5] (nodes with recent RTS/CTS activity), denoted by $\mathcal{N}^{(j)}$. For each node $i \in \mathcal{N}^{(j)}$, node j also maintains $I_{total}^{(i)}$, the AOA of the signal received from node i, and the average path loss between i and j ($\Theta^{(ij)}$). Each entry in the cache is associated with a timer. Upon the expiration of this timer, the corresponding entry is flushed out. The value of $I_{avg}^{(j)}$ is the average of $I_{total}^{(i)}$ for $i \in \mathcal{N}^{(j)}$.

In cellular networks, base station controllers (BSCs) implement load control by adjusting the transmission powers of mobiles stations such that the loading in the network does not exceed L_p. In LCAP, load control must be implemented in a distributed manner, as described next.

C. Overview of LCAP Operation

We now give an overview of the operation of LCAP, explaining how load control is implemented, how concurrent transmissions in the same neighborhood are supported, and how the protocol protects ongoing receptions from future interfering transmissions. The main notation used in this and subsequent sections is summarized in Table I.

In designing LCAP, we make the following assumptions:

1) Separate channels are used for data and control packets. Besides being used for RTS/CTS packets, the control channel can also be used for route-discovery messages. As explained later, the use of a dual-channel solution in LCAP eliminates several of the channel access problems discussed in Section IV. Although such a solution involves more sophisticated hardware (but no additional

[3]Although the extra power causes additional interference at non-intended receivers, because of the locality of typical one-hop transmissions and the nonlinearity of the path loss phenomenon, the gain due to the extra power at the intended receiver overshadows its negative impact on non-intended receivers.

[4]By definition, $I_{total}^{(j)}$ must be below the *power-reception threshold* of node j; otherwise, it would initiate the capture-and-decode circuitry, putting node j in the *receive* mode.

[5]Node i is said to be a neighbor of node j if i's omnidirectional RTS can be correctly received at j in the absence of interference other than the thermal noise and while j is listening to the channel omnidirectionally.

cost in spectrum), the two channels are never used *simultaneously* for transmission. So while the data channel is being used for transmission, the control channel is in the receive (or listen) mode, but not in the transmit mode, and vice versa. As a consequence, there is no need for two antennas at a node; a duplexer is sufficient to implement both channels using one antenna.

2) A node can accurately estimate the AOA of a received signal (using, example, the techniques in [12], [21], [23]) as well as the total interference power.
3) Channel gain is reciprocal (i.e., $\Theta^{(ij)} = \Theta^{(ji)}$) and stationary for the duration of few control packets plus one data packet (i.e., the channel is slowly varying).
4) Each node can determine the antenna patterns of its neighbors, expressed as functions of the angle in the azimuth-vertical plane. This is trivial if nodes use identical antennas. If nodes use different types of directional antennas, then each antenna pattern can be assigned an integer identifier. For example, one byte is sufficient to encode the identifiers of 256 different antenna patterns. The one-byte identifier can be transmitted as part of the RTS packet.

To explain the working of LCAP, consider Figure 8. Suppose that node i wants to transmit a data packet to node j . It first sends an RTS packet *omnidirectionally* at a fixed, known power ($P_{\max}$) on the control channel. Every neighbor of i, say j, that hears the RTS updates its information about $I_{total}^{(i)}$ and the AOA of i's signal, and adds an entry for node i in $\mathcal{N}^{(j)}$. The neighbor also calculates the path loss between i and j: $\Theta^{(ij)} = P_{\max}/P_{recvd}^{(ij)}$, where $P_{recvd}^{(ij)}$ is the received power of the RTS. If j is the intended receiver, it computes the *maximum additional interference* ($P_{allowed}^{(j)}$) it can allow in its neighborhood that would not cause $\eta_{avg}^{(j)}$ to exceed L_p. This power can be calculated by setting $L_p = 1 - N_0/(I_{avg}^{(j)} + P_{allowed}^{(j)})$ and solving for $P_{allowed}^{(j)}$. Thus,

$$P_{allowed}^{(j)} = \frac{N_0 - I_{avg}^{(j)}(1 - L_p)}{(1 - L_p)}. \tag{2}$$

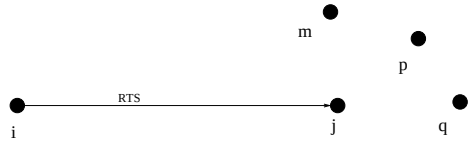

Fig. 8. Example topology that demonstrates the working of LCAP.

It is important to notice that $P_{allowed}^{(j)}$ is a *neighborhood quantity*, and is computed prior to the $i \to j$ data transmission. This quantity must be treated as a resource to be shared by receiver j itself[6] and its neighboring nodes (nodes m, p, and q in Figure 8).

To explain how $P_{allowed}^{(j)}$ is shared, we first consider the impact of the $i \to j$ transmission on the interference (load) in the neighborhood of j. Let $P_{min}^{(ij)}$ be the minimum transmission power required to achieve SINR_{th} at receiver j:

$$P_{min}^{(ij)} = \frac{\text{SINR}_{th} \cdot I_{total}^{(j)} \cdot \Theta^{(ij)}}{G_M^2}. \tag{3}$$

where G_M is the antenna gain in the direction of the main lobe. In (3) the factor G_M^2 appears because during the $i \to j$ transmission, nodes i and j are beamformed in each other's direction.

The power $P_{min}^{(ij)}$ is computed from the transmitter's point of view. To reflect its impact on receiver j's neighborhood, we "translate" it into another quantity, $\hat{P}_{min}^{(ij)}$, that is on par with $P_{allowed}^{(j)}$. The details of such translation are given in Section V-D. After accounting for $\hat{P}_{min}^{(ij)}$, the leftover interference margin $P_{margin}^{(j)} \stackrel{\text{def}}{=} P_{allowed}^{(j)} - \hat{P}_{min}^{(ij)}$ is split into two parts. The first part, denoted by $P_{other}^{(j)}$, represents the maximum interference power that j's neighbors other than i are allowed to generate in the future without disturbing the $i \to j$ transmission. The second part is used to combat the additional neighborhood interference caused by scaling up the transmission power of i

$I_{total}^{(j)}$	Total interference-plus-noise at node j.
L_p	Planned loading factor.
$I_{avg}^{(j)}$	Average of $I_{total}^{(.)}$ values in j's neighborhood.
$\eta^{(j)}$	Loading at node j.
$\eta_{avg}^{(j)}$	Average loading in j's neighborhood.
G_M	Main lobe gain.
N_0	Thermal noise at each node.
$\Theta^{(ij)}$	Path loss between nodes i and j.
$P_{recvd}^{(ij)}$	Received power at j for a transmission from i.
$P_{min}^{(ij)}$	Minimum required transmission power at node i to achieve SINR_{th} at receiver j.
$\hat{P}_{min}^{(ij)}$	Interference power associated with $P_{min}^{(ij)}$ when measured at the centroid of j's neighborhood.
$P_{allowed}^{(j)}$	Maximum allowable future interference in the neighborhood of node j.
$P_{margin}^{(j)}$	Allowed interference in the neighborhood of node j if node i starts transmission for j at $P_{min}^{(ij)}$.
$P_{self}^{(j)}$	Power required at the receiver j over and above the minimum power needed for coherent reception.
$P_{other}^{(j)}$	Maximum interference power that node j can tolerate from all future interferers while receiving a data packet.
$I_p^{(j)}$	Interference power allowed per future interferer at node j.
$G^{(ij^*)}$	Node i's antenna gain in the direction of centroid j^*.
$\mathcal{N}^{(j)}$	Set of active nodes in the neighborhood of node j.
$P_{\text{CTS}}^{(j)}$	Transmission power at which node j is required to send a CTS to reach all nodes in $\mathcal{N}^{(j)}$.
$P_{req}^{(ij)}$	Transmission power at which node i is required to send a data packet to node j that ensures capture if interference $\leq P_{other}^{(j)}$.
$P_{\text{DNAV}}^{(ik)}$	Maximum power that node i can send in the direction of receiver k listed in i's DNAV table.

TABLE I
NOTATION USED IN THE PAPER.

[6]Although j is the intended receiver for i's transmission, from the standpoint of j's neighbors such a transmission is considered as interference, and must thus be accounted for in $P_{allowed}^{(j)}$.

beyond $P_{min}^{(ij)}$. In order for other transmissions to take place in the neighborhood of j while j is receiving its packet from i, node i must transmit its data packet at power $P_{req}^{(ij)}$ that is greater than $P_{min}^{(ij)}$:

$$P_{req}^{(ij)} \stackrel{\text{def}}{=} P_{min}^{(ij)} + \left(\frac{P_{self}^{(j)}}{G_M}\right)\Theta^{(ij)} \tag{4}$$

where $P_{self}^{(j)}$ is the additional power used by i after the transmitting antenna gain (hence the gain factor G_M in the denominator), *as measured at the receiver* j.

Note that $P_{self}^{(j)}$ and $P_{other}^{(j)}$ are receiver-based quantities, and thus are not comparable with $P_{allowed}^{(j)}$ and $P_{margin}^{(j)}$, which are neighborhood quantities.

From (3), $P_{min}^{(ij)}$ is determined by SINR_{th} and $\Theta^{(ij)}$, both of which cannot be controlled. However, $P_{self}^{(j)}$ and $P_{other}^{(j)}$ can be set such that if future transmissions in j's neighborhood produce $P_{other}^{(j)}$ worth of interference power at receiver j, then $P_{self}^{(j)}$ should be large enough to ensure that the SINR at node j does not fall below SINR_{th}. The determination of $P_{self}^{(j)}$ and $P_{other}^{(j)}$ is provided later in Section V-E.

To distribute $P_{other}^{(j)}$ among future interfering transmitters, node j must estimate the number of such transmitters (denoted by $M^{(j)}$). It does that on the basis of the number of RTS messages it has overheard over a fixed time period, not counting the RTS messages that are intended for node j itself. Accordingly, $M^{(j)}$ is the size of the set $\mathcal{N}^{(j)}$. Node j then calculates the allowed interference power per future transmission: $I_p^{(j)} = P_{other}^{(j)}/M^{(j)}$. Thereafter, node j sets the power of its CTS (denoted by $P_{\text{CTS}}^{(j)}$) such that this CTS reaches all future potential interferers. The details of calculating $P_{\text{CTS}}^{(j)}$ is given in Section VI-A. In its CTS, node j also includes the values of $P_{req}^{(ij)}$, $I_p^{(j)}$, and $P_{\text{CTS}}^{(j)}$. The CTS packet is sent directionally on the control channel. Upon receiving this CTS, node i beamforms in the direction from which the CTS has been received and starts sending the data packet at power $P_{req}^{(ij)}$.

Upon overhearing the CTS, an idle neighbor[7] sets its DNAV and computes the channel gain between itself and the sender of the CTS. This is illustrated in the example in Figure 9, where node m overhears j's CTS and accordingly computes $\Theta^{(jm)} = P_{\text{CTS}}^{(j)}/P_{recvd}^{(jm)}$, where $P_{recvd}^{(jm)}$ is now the power at which j's CTS is received at node m. Node m then calculates the maximum permitted power $P_{\text{DNAV}}^{(mj)}$ that m can use if it is to transmit a packet in the direction of node j during j's activity. This power is given by $P_{\text{DNAV}}^{(mj)} = \Theta^{(mj)} \cdot I_p^{(j)}$. Thereafter, node m stores $P_{\text{DNAV}}^{(mj)}$ along with its expiration time, which is obtained from the timing information provided by j's CTS. The gain $\Theta^{(mj)}$ includes the path loss from node m to node j as well as the receiving antenna gain of node j in the direction of node m (since j's CTS is sent to i directionally). In other words, $P_{\text{DNAV}}^{(mj)}$ represents the maximum allowable transmission power that node m can transmit with an antenna gain of one while node j is receiving a data packet directionally. Table II depicts an example of the DNAV table with two entries that represent two power-limited directions. Note that the DNAV table is not the same as the cache table; unlike the DNAV table, the cache table may contain neighbors that are currently inactive but that were recently active and have not expired according to the cache expiration period. Such nodes are considered potential interferers (see Section VI-A for details).

[7]In our terminology, an idle node is one that is not waiting for a CTS, a data, or an ACK packet. Such a node listens to the channel omnidirectionally.

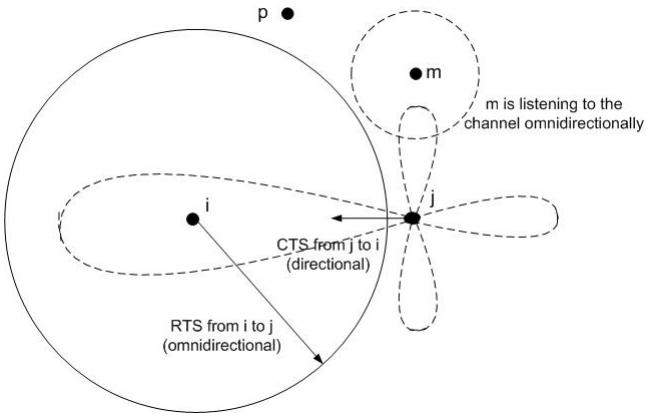

Fig. 9. Setting of the DNAV in LCAP.

k	Start	End	$P_{\text{DNAV}}^{(mk)}$ (Watts)	Expiration (μs)
1	240°	300°	6.4E−2	2000
2	150°	210°	7.9E−1	4000

TABLE II
EXAMPLE OF THE DNAV TABLE AT NODE m.

In the above example, if node m intends to send a data packet to node p, it not only checks the DNAV in the direction of p, but it also inspects every direction in m's DNAV table. Specifically, node m calculates the power it would radiate in the direction of each receiver listed in m's DNAV table if m were to beamform in that direction. To do that, node m requires the antenna gain along that direction. Let such gain be denoted by G_{mk} for each receiver k in m's DNAV table. Since the antenna pattern is known, the angular offset of the DNAV with respect to the boresight can be calculated by node m, and the gain in any direction can be determined. So after receiving the CTS from node p, node m executes the algorithm in Figure 10.

Basically, for every receiver k in m's DNAV table, node m determines how much power it would radiate in the direction of k. If that power is larger than $P_{\text{DNAV}}^{(mk)}$, no transmission takes place. This strategy solves the vulnerable receiver problem mentioned in Section IV-A.1, as nodes consider the interference resulting from minor lobes before deciding to transmit.

Figure 11 illustrates how various quantities mentioned previously are calculated in LCAP.

D. Neighborhood Centroid and Computation of $\hat{P}_{min}^{(ij)}$

In the previous section, we discussed how the maximum allowable interference at node j ($P_{allowed}^{(j)}$) is computed. If

CHECK-DNAV($P_{req}^{(mp)}, m$)
1 **for** each node k in m's DNAV table, **do**:
2 Find the antenna gain in k's direction (G_{mk}).
3 Determine the power that would be radiated in k's direction
4 $P_{mk} = G_{mk} P_{req}^{(mp)}$
5 **if** $P_{\text{DNAV}}^{(mk)} < P_{mk}$
6 Abort-transmission and **terminate** CHECK-DNAV.
7 **end-for**
8 Start-transmission.
9 **end** CHECK-DNAV

Fig. 10. Algorithm for checking the DNAV table at a prospective transmitter m that is beamformed in the direction of a receiver p.

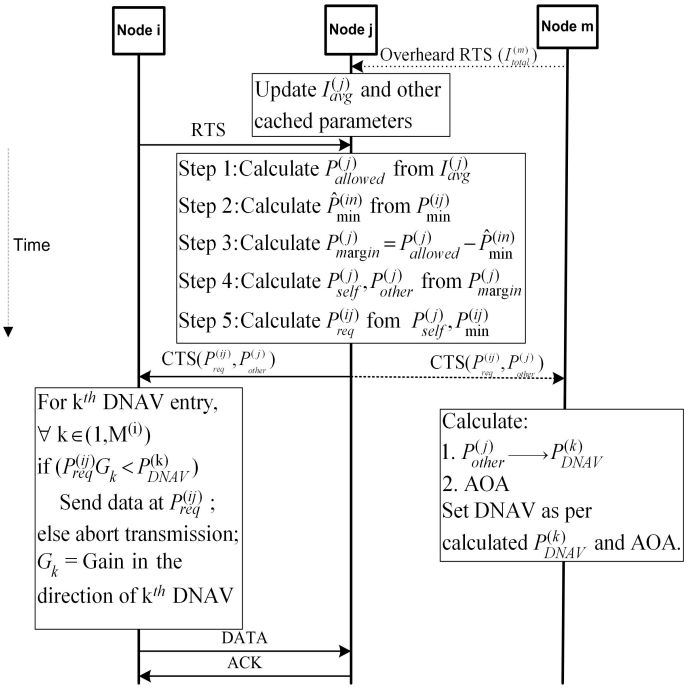

Fig. 11. Diagram of the relative timing of various computations in LCAP.

node j intends to receive a data packet from another node i, it splits $P_{allowed}^{(j)}$ into two parts: $P_{margin}^{(j)}$ and $\hat{P}_{min}^{(ij)}$. In this section, we are concerned with the computation of $\hat{P}_{min}^{(ij)}$. This quantity reflects the impact (interference) of the minimum transmission power $P_{min}^{(ij)}$ on node j's neighborhood. To compute $\hat{P}_{min}^{(ij)}$, we abstract node j's neighborhood by its "centroid", j^*, which is a fictitious point in the 2D space that represents the average distance (equivalently, path loss) and orientation of j's *active* neighbors relative to j. To illustrate, consider Figure 8. Let $\Omega = \{m, p, q\}$ be the set of currently active neighbors of node j. Let $\overrightarrow{jk}$ be a vector in the 2D space that points from j in the direction of k, $k \in \Omega$. The amplitude of $\overrightarrow{jk}$ is $||\overrightarrow{jk}|| = \Theta^{(jk)}$; the path loss between j and k. Node j determines the vector $\overrightarrow{jk}$ using AOA and path-loss estimates obtained through previously transmitted RTS packets from node k. Node j can then compute the path-loss vector $\overrightarrow{jj^*}$ using

$$\overrightarrow{jj^*} = \frac{1}{N} \cdot \sum_k \overrightarrow{jk} \tag{5}$$

where N is the number of active neighbors of node j. In other words, the x- and y-coordinates of the vector $\overrightarrow{jj^*}$ are, respectively, the means of the x- and y-coordinates of the vectors $\overrightarrow{jk}$, $k \in \Omega$. Note that $||\overrightarrow{jj^*}|| = \Theta^{(jj^*)}$.

To calculate $\hat{P}_{min}^{(ij)}$, node j needs the path loss between node i and the centroid j^* ($\Theta^{(ij^*)}$) as well as i's antenna gain in the direction of j^* when i is beamformed in the direction of node j (G_{ij^*}). The path loss $\Theta^{(ij^*)}$ can be computed from $\Theta^{(ij^*)} = \|\overrightarrow{ij^*}\| = \|\overrightarrow{jj^*} - \overrightarrow{ji}\|$, where $\overrightarrow{ji}$ is computed at j from the AOA and path-loss estimate of i's RTS packet. Node j can also compute the angle ϕ_{ij^*} between $\overrightarrow{ij^*}$ and $\overrightarrow{ji}$:

$$\phi_{ij^*} = \cos^{-1}\left\{ \frac{\overrightarrow{ji} \cdot \overrightarrow{ij^*}}{\|\overrightarrow{ji}\| \|\overrightarrow{ij^*}\|} \right\}. \tag{6}$$

Note that ϕ_{ij^*} represents the angular offset of the direction $\overrightarrow{ij^*}$ with respect to the boresight $\overrightarrow{ji}$ (since i is beamformed in the direction of j). From ϕ_{ij^*}, $\Theta^{(ij^*)}$, and the antenna pattern, node j can easily compute G_{ij^*}.

Finally, node j can now calculate $\hat{P}_{min}^{(ij)}$:

$$\hat{P}_{min}^{(ij)} = \frac{G_{ij^*} \cdot P_{min}^{(ij)}}{\|\overrightarrow{ij^*}\|^\alpha} \tag{7}$$

where α is the path loss exponent.

E. Computation of $P_{self}^{(j)}$ and $P_{other}^{(j)}$

As indicated before, besides the $i \rightarrow j$ transmission, additional (future) transmissions are to be allowed in the neighborhood of node j. For such transmissions to take place, node j must request that node i uses a power level $P_{req}^{(ij)}$ greater than $P_{min}^{(ij)}$. The extra *received* power, which we denoted by $P_{self}^{(j)}$, will consume a part of $P_{margin}^{(j)}$, whereas the leftover power (denoted by $P_{other}^{(j)}$) represents the maximum future interference that nodes in the neighborhood of j can afflict upon node j. We now explain how $P_{self}^{(j)}$ and $P_{other}^{(j)}$ are determined by node j.

The quantities $P_{self}^{(j)}$ and $P_{other}^{(j)}$ can be set such that if in the future, nodes in the neighborhood of node j generate $P_{other}^{(j)}$ worth of interference at j, then $P_{self}^{(j)}$ should be such that the SINR does not drop below SINR$_{th}$ at node j. This can be expressed by the following condition:

$$\frac{P_{min}^{(ij)} \cdot G_M^2 / \Theta^{(ij)} + P_{self}^{(j)} \cdot G_M}{I_{total}^{(j)} + P_{other}^{(j)}} \geq \text{SINR}_{th}. \tag{8}$$

For the optimal case when no power more than required is used, i.e., SINR threshold is exactly met, (8) becomes an equality. Substituting the value of $P_{min}^{(ij)}$ from (3) in this equality, we arrive at following relationship between $P_{self}^{(j)}$ and $P_{other}^{(j)}$:

$$\frac{P_{self}^{(j)}}{P_{other}^{(j)}} = \frac{\text{SINR}_{th}}{G_M}. \tag{9}$$

By definition, $P_{margin}^{(j)}$ is a neighborhood resource, and $P_{self}^{(j)}$ and $P_{other}^{(j)}$ (which are power margins at receiver j) are derived from this resource. To receive the extra $P_{self}^{(j)}$, receiver j requests transmitter i to use an extra power of $P_{self}^{(j)}/\Theta^{(ij)}$. This extra power consumes a part of $P_{margin}^{(j)}$ at the centroid j^* that is given by:

$$(P_{self}^{(j)}/\Theta^{(ij)}) \cdot G_{ij^*} \cdot \Theta^{(ij^*)} = P_{self}^{(j)} \cdot \gamma \tag{10}$$

where $\gamma \stackrel{\text{def}}{=} G_{ij^*} \cdot \Theta^{(ij^*)}/\Theta^{(ij)}$. The part of $P_{margin}^{(j)}$ consumed by future transmissions is given by: $P_{other}^{(j)} \cdot \Theta^{(j^*j)}$. Note that $\Theta^{(j^*j)}$ is known to receiver j from the path loss of previously received RTS packets. Accordingly,

$$P_{other}^{(j)} \cdot \Theta^{(j^*j)} + P_{self}^{(j)} \cdot \gamma = P_{margin}^{(j)}. \tag{11}$$

By solving (11) and (9), we determine the values of $P_{self}^{(j)}$ and $P_{other}^{(j)}$:

$$P_{self}^{(j)} = \text{SINR}_{th} \cdot \frac{P_{margin}^{(j)}}{\Theta^{(j^*j)} G_M + \text{SINR}_{th}\gamma} \tag{12}$$

$$P_{other}^{(j)} = G_M \cdot \frac{P_{margin}^{(j)}}{\Theta^{(j^*j)} G_M + \text{SINR}_{th}\gamma}. \tag{13}$$

From $P_{self}^{(j)}$ and $P_{min}^{(ij)}$, node j can compute $P_{req}^{(ij)}$ using (4). If $P_{req}^{(ij)}$ is greater than $P_{\max}$, then the transmission is aborted and the DCTS is not sent.

VI. SOLUTIONS FOR CHANNEL ACCESS PROBLEMS

In this section, we discuss how LCAP addresses the channel access problems that afflict other MAC protocols for directional antennas.

A. Adaptive Power Scaling of DCTS Transmissions

As explained in Sections IV-B.2, the unequal gains of the omni and directional modes increase the likelihood of collisions. The essence of this problem is that in previous protocols for MANETs with directional antennas, the DCTS does not reach all potential interferers. To address this problem, LCAP allows a receiver j to adaptively adjust the transmission power $P_{\text{CTS}}^{(j)}$ of its DCTS packet such that this packet reaches all potential interfers in the neighborhood of j. We now explain how node j computes $P_{\text{CTS}}^{(j)}$. Clearly, this computation requires determining the set of potential interferers $\mathcal{N}^{(j)}$. As will become clear shortly, $\mathcal{N}^{(j)}$ is actually a superset of the DNAV table at node j. Recall that each entry in the DNAV table has an expiration time, which reflects the time by which the corresponding data transmission will be completed. In contrast, each entry in $\mathcal{N}^{(j)}$ is associated with a larger expiration time, called the *activity expiration time* (AET). This time reflects the duration over which the corresponding neighbor is expected to stay active. Whenever an entry is added to the DNAV table, the same entry is also added to $\mathcal{N}^{(j)}$. However, the AET should be set to a value that is larger than the transmission time of the corresponding data packet if a neighbor is expected to be active for more than one packet time (which is often the case). Note that by increasing the value of the AET, more *potential* interferers are considered, which may in turn reduce the likelihood of packet collisions and hence increase the network throughput. However, a larger set of potential interferers also means less interference margin per interferer (the total interference margin at node j is split equally among potential interferers). If some of these potential interferers are no longer active, then their interference margins have been unnecessarily allocated, thereby reducing the number of concurrent transmissions and negatively impacting the spatial channel reuse. As a compromise between the two trends, we heuristically set the AET to twice the timeout value of the corresponding entry in the DNAV table. The simple rationale behind this choice is that the activity of a node at the link layer (LL) is a highly autocorrelated process. So if a node transmits a LL packet at the present time, there is a good chance it will continue to send more packets in the near future. It is also possible to customize the setting of the AET so that it depends on the observed load of each neighbor (i.e., the higher the observed load, the smaller the value of the AET), but this increases the complexity of the protocol.

Note that the definition of $\mathcal{N}^{(j)}$ does not take into account a node, say m, that starts its transmission after a time period greater than the AET, which could cause collisions. However, the transmission of an RTS/DCTS message by node m will make other nodes aware of node m's activity. Thus, in the future, node m will be treated as a potential interferer during the calculation and distribution of the interference margin.

When a node j broadcasts its DCTS, it scales its power such that this packet can be heard by all nodes listed in $\mathcal{N}^{(j)}$ and are not located in the direction of j's antenna nulls (nodes located in the direction of antennas nulls do not cause interference). The value of $P_{\text{CTS}}^{(j)}$ is calculated from the path loss of the previously heard RTS packet of a potential interferer and from the antenna gain in the direction of that interferer, assuming that node j beamforms in the direction of the given transmitter:

$$P_{\text{CTS}}^{(j)} = \max_{k \in \mathcal{N}^{(j)}} \left\{ \frac{\text{SINR}_{th} \cdot \Theta^{(jk)} \cdot N_0}{G_{jk}} \right\}. \tag{14}$$

By adjusting the transmission power of the DCTS, node j can inform potential interferers of the maximum interference power they can send in node j's direction for a given amount of time. Thus, collisions at a receiver due to asymmetric gains and minor lobes can be significantly reduced.

It should be noted that the aim of LCAP is to improve the spatial reuse and decrease the energy consumption, and not to extend the communication range. In fact, by sending the RTS packet omnidirectionally, LCAP does not realize the full potential of directional antennas as a means of extending the transmission range. When coupled with power control, a reduced transmission range translates into significant energy saving.

B. Segregation of Data and Control Channels

Collisions due to unheard RTS/CTS messages (Section IV-B.1) and to a vulnerable transmitter (Section IV-A.2) are

avoided by using separate data and control channels. Consider, for example, the unheard RTS/CTS problem discussed in Section IV-B.1 in relation to Figure 6. While node B is transmitting data to node A, node C sends an RTS to node D. Upon receiving the RTS, node D replies with a DCTS. Although nodes A and node B are busy sending/receiving data on the data channel, their control channels are idle, and thus, they are able to hear the DCTS from node D. Accordingly, both nodes set their DNAVs. This resolves the unheard RTS/CTS problem mentioned in Section IV-B.1. As for the exposed transmitter problem, separating the control and data channels eliminates entirely the possibility of a collision between data and control messages. Collisions can occur only when two control messages reach a node simultaneously. However, this has a low probability of occurrence due to the small size of the control messages. Consequently, the vulnerable transmitter problem (Section IV-A.2) is addressed. Furthermore, there is no need for nodes to acquire location information of neighboring nodes from higher layers [5] or via AOA caching [19].

VII. PROTOCOL EVALUATION

A. Simulation Setup

In this section, we evaluate the performance of LCAP and contrast it with three other MAC protocols. Our simulation programs are written in Csim; a process-oriented discrete-event programming environment that is based on the C language [1]. Compared with the popular ns-2 simulator, our code provides a more accurate account of interference. Specifically, when computing I_{total}, our programs *simultaneously* account for all sources of interference, including those that are very far away from the intended receiver. The total interference is taken as the sum of the received powers of *all* interferers. In contrast, ns-2 considers interferers one at a time; it compares the received power of a given interferer with the power of the intended signal and accordingly decides whether or not *that* interferer is enough to cause packet collision (it does not add the powers of interferers). This often results in an overly optimistic assessment of packet collisions.

We compare LCAP with three other protocols: RMAC (*Reference MAC*), *omnidirectional LCAP* (O-LCAP), and the IEEE 802.11b Ad Hoc scheme. RMAC is similar to previously proposed protocols for directional antennas [19] [11] [5] except that it uses two channels and hence does not suffer from the vulnerable transmitter problem. It does not implement power control nor does it account for the effect of minor-lobe interference when setting the DNAV table. The objective of comparing RMAC with LCAP is to demonstrate the performance gain due to power control and interference-limited concurrent transmissions (RMAC is equivalent to an instance of LCAP operating at $L_p = 0$). O-LCAP is a version of LCAP for MANETs with omnidirectional antennas. Like LCAP, O-LCAP implements power control, uses two channels, and allows for concurrent transmissions. The study of O-LCAP is meant to demonstrate the virtues of directional transmissions in the context of power-controlled MAC protocols for MANETs. The 802.11 scheme is the simplest of the four protocols. It uses omnidirectional transmissions over a common channel with no power control. Its performance serves as a lower-bound reference point for other protocols. All four protocols have the same *total* channel bandwidth. Note that using directional transmissions over a single channel creates several hidden-terminal problems, so we did not consider this option. Also, there is no point in studying a two-channel solution for omnidirectional MANETs without power control (the separation of the control and data channels is mainly intended to address problems that arise in the context of directional transmission). In our simulations for LCAP, RMAC, and O-LCAP, we implement energy-oriented routing based on the connectivity set (CONSET) protocol [14]. Such a routing strategy results in paths with more hops and a smaller distance per hop than shortest-hop routing.

We investigate the effect of loading on the throughput-energy tradeoff. Our performance metrics include the total network throughput, the one-hop throughput, and energy consumption per transmitted packet. We set the size of a data packet to 2000 bytes. Simulations for random, uniformly distributed packet sizes (not reported) were also performed, and they indicated similar trends to the ones observed for the fixed-size case. We use a random-way point mobility model with speeds that are uniformly distributed between 0 and 2 meters/sec. The radiation pattern for the directional antennas is the one in Figure 1. Other parameters used in the simulations are given in Table III. The combined data rate of the control and data channels in LCAP is chosen to be equal to the IEEE 802.11b data rate (2 Mbps). Typical values are used for the antenna gains of a standard six-element circular array directional antenna. The chosen reception and carrier sense thresholds are also typical for contemporary WLAN cards.

Directional antenna type	6-element circular array
Omnidirectional antenna gain	2.2 dBi
Main lobe gain	15 dBi
Beamwidth	60^o
Data packet size	2 KB
LCAP data-channel rate	1.6 Mbps
LCAP control-channel rate	400 Kbps
SNR threshold	6 dB
Reception Threshold	-94 dBm
Carrier sense threshold	-108 dBm
Thermal noise	-169 dBm/Hz
Maximum EIRP	35 dBm

TABLE III
PARAMETERS USED IN THE SIMULATIONS.

Our simulations are performed for two types of topologies: *random grid* and *clustered*. For both types, 64 nodes are placed within a square area of dimensions 500x500 m^2. In the case of a random-grid topology, the square area is split into 64 smaller squares, and a node is placed at a random location within one of these smaller squares. During the initialization phase, sending nodes randomly choose their destination nodes. Thereafter, each node generates packets according to a Poisson process with rate λ (same for all nodes)[8]. For the random-grid

[8]Although more sophisticated traffic models can be used in the simulations, e.g., ON/OFF model, such models are not expected to make a qualitative difference in the relative performance of the tested protocols.

topology, we set $P_{\max} = 10$ micro-watts (recall that $P_{\max}$ is the power used to transmit the RTS packet omnidirectionally). Given a reception threshold of -94 dBm and assuming a 2-ray propagation model, this $P_{\max}$ corresponds to a maximum (omni-directional) reception range of 106 meters.

For the clustered topology, the 64 nodes are divided into four clusters, and each cluster is placed in a 50x50 m^2 square in one of the corners of the larger area. A source node selects a destination from within its own cluster with probability $1-p$ and out of its cluster with probability p. We set $P_{\max}$ to 30 milli-watts (reception range $=$ 700 meters), so that all transmissions are one-hop.

B. Results

We first consider random-grid topologies. Figure 12 demonstrates the effect of L_p on the total network throughput and on the energy consumed per packet in LCAP. It can be observed that increasing L_p improves the network throughput but at the cost of increased energy consumption per packet. This is expected from (2), since increasing L_p results in receiving nodes choosing a greater value of $P^{(j)}_{allowed}$ and, in turn, $P^{(j)}_{margin}$, providing a greater interference margin for future transmissions. As a result, more interference-limited transmissions are allowed to take place. Also from (12), we know that by increasing L_p, the link margin $P^{(j)}_{self}$ also increases. As receiving nodes control the powers of the transmitting nodes, by increasing $P^{(j)}_{self}$, receiving nodes request higher transmission power $P^{(ij)}_{req}$ (see (4)). By adjusting the value of L_p, LCAP provides a mechanism for tuning the tradeoff between network throughput and energy consumption.

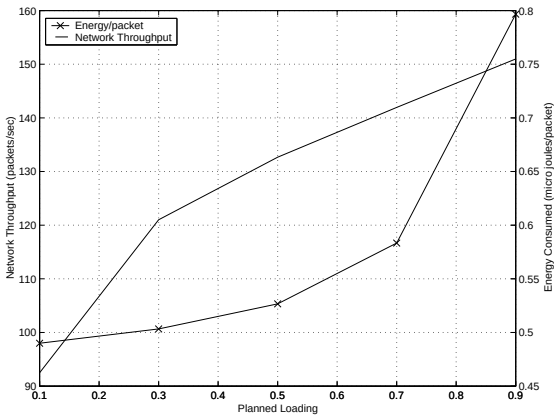

Fig. 12. Network throughput and energy consumed per packet versus L_p for the LCAP protocol (random-grid topology, λ = 25 packets/sec).

To quantify the channel's spatial reuse, we use the *probability of concurrent transmissions*, defined as the fraction of transmissions that cause interference to ongoing receptions but that are allowed to take place in LCAP (because the interference is within the acceptable limit). As shown in Figure 13, as L_p increases, the allowed interference per future transmission also increases, and so is the probability of concurrent transmissions. Such an increase explains the increase in network throughput in Figure 12.

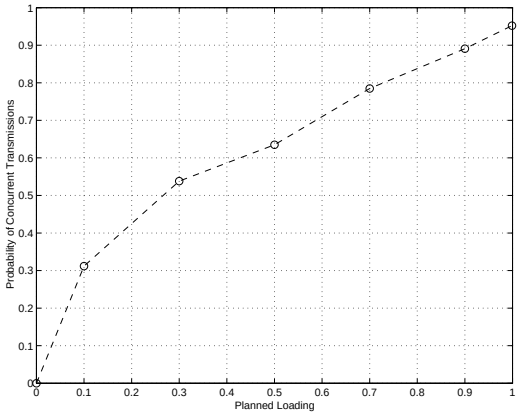

Fig. 13. Probability of concurrent transmissions versus L_p in LCAP.

Figure 14 depicts network throughput versus λ under the random-grid topology. At light loads ($\lambda \leq 1$), there is not much difference among the tested protocols. As λ increases, the differences become more significant, with LCAP displaying the highest throughput. At $\lambda = 5$, LCAP achieves more than 140% improvement in throughput over the IEEE 802.11b scheme and about 71% improvement over RMAC. Recall that RMAC uses directional antennas, but does not implement power control, does not allow for concurrent interference-limited transmissions, and does not account for side-lobe interference in channel access. Interestingly, the figure shows that the omnidirectional version of LCAP (O-LCAP) achieves higher throughput than RMAC.

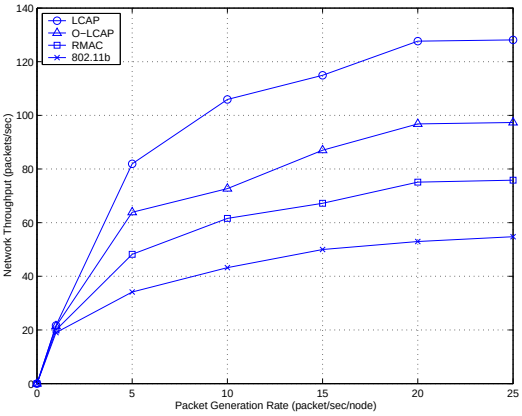

Fig. 14. End-to-end network throughput versus λ for the random grid topology ($L_p = 0.9$ for LCAP and O-LCAP).

For the same topology, Figure 15 depicts the one-hop throughput versus λ. At $L_p = 0.9$, the probability of concurrent transmissions is high (about 0.9, as indicated in Figure 13). Unlike RMAC and the IEEE 802.11b scheme, 90% of all transmission requests in LCAP are immediately accepted without requiring the nodes to go into backoff. This explains the considerable increase in the one-hop throughput. Another factor that contributes to the higher throughput of LCAP is its consideration of side-lobe interference. By accounting for such interference in setting the DNAV, LCAP reduces the likelihood

of collisions and subsequent backoffs. This is shown in Figure 16, which depicts the probability of a packet collision due to minor lobe interference versus λ for both LCAP and RMAC.

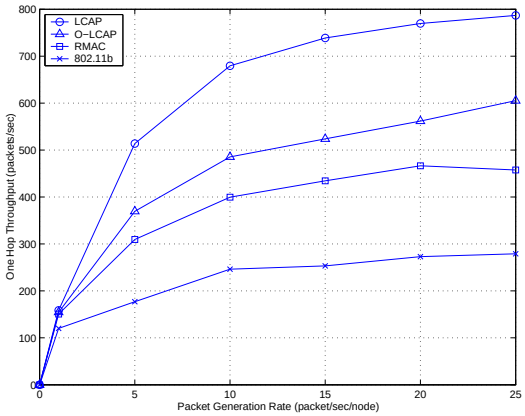

Fig. 15. One-hop throughput versus λ for LCAP ($L_p = 0.9$), O-LCAP, RMAC, and the IEEE 802.11b scheme under the random-grid topology.

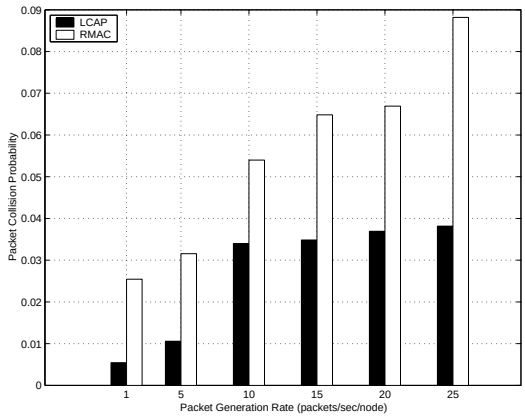

Fig. 16. Probability of a packet collision due to side-lobe interference vs. λ ($L_p = 0.9$ for LCAP; random-grid topologies).

Because of power control, the throughput improvement in LCAP does not come at the cost of increased energy consumption, as shown in Figure 17. The energy consumption per packet is the same for RMAC and the IEEE 802.11b scheme since in both protocols data packets are sent at the maximum power. The simulations show that with $L_p = 0.9$, the average energy consumption per packet in LCAP is about 16% of that of RMAC and the IEEE 802.11b scheme. This significant decrease is attributed to the combined gain from using directional antennas and transmission power control. Note that the results in Figure 17 are for the energy consumed in transmission, not reception. The energy consumed per packet during reception is identical for LCAP, RMAC and IEEE 802.11b. However, for RMAC and IEEE 802.11b, the average energy consumption during transmission is 66.67% greater than energy consumption during reception [6].

Figure 18 shows the total network throughput versus λ for the clustered topology with $p = 0.25$. At a moderate traffic load of $\lambda = 10$, LCAP achieves about 116% increase in

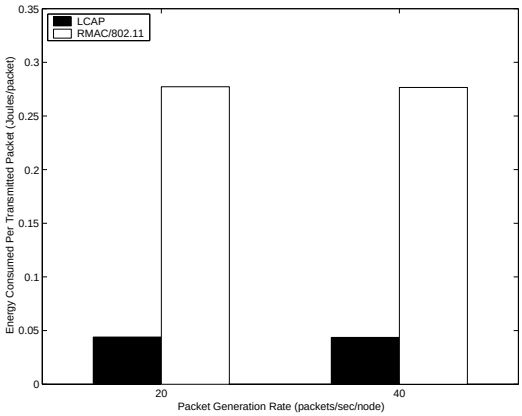

Fig. 17. Energy consumed per transmitted packet versus λ for LCAP ($L_p = 0.9$) and RMAC/IEEE 802.11b under the random-grid topology.

throughput over IEEE 802.11b and 27% over RMAC. At a heavier load of $\lambda = 20$, the relative throughput gains jumps to 192% and 43%, respectively. Note that the percentage increase in throughput of LCAP over IEEE 802.11b for the clustered topology is greater than for the random-grid topology (compare Figures 14 and 18). This is because in the clustered topology, all nodes within a cluster are in the carrier-sense range of each other. Thus, for the IEEE 802.11b scheme, most often only one transmission takes place at a time in a cluster.

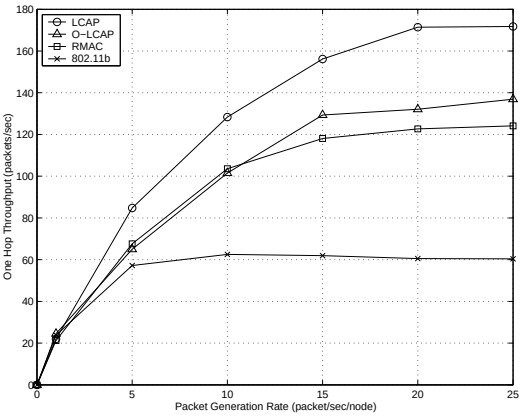

Fig. 18. Network throughput versus λ for clustered topologies ($p = 0.25$).

VIII. CONCLUSIONS

In this paper, we proposed a power-controlled MAC protocol for MANETs with directional antennas. The proposed protocol improves spatial reuse by allowing interference-limited concurrent transmissions. It also provides a planned loading mechanism for setting the desired tradeoff between network throughput and energy consumption. Nodes employ load control in a distributed fashion to upper-bound the interference in their neighborhoods. LCAP extends the concept of DNAV by associating a maximum permitted power value with each reserved direction. It ensures that the radiated power remains below this value not only in the direction of the main lobe but in all directions, thereby resolving the minor lobe radiation

problem. Simulations show that when compared to a two-channel version of a previously proposed protocol (the RMAC protocol), LCAP improves the network throughput by up to 71% and, at the same time, achieves about 89% reduction in the transmission-energy consumption per packet.

REFERENCES

[1] Mesquite Corporation. Online at http://www.mesquite.com.

[2] *IEEE Std 802.11b Part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications*, 1999.

[3] S. Bandyopadhyay, K. Hasuike, S. Horisawa, and S. Tawara. An adaptive MAC protocol for wireless ad hoc community network (WACNet) using electronically steerable passive array radiator antenna. In *Proceedings of the IEEE GlobeCom Conference*, volume 5, pages 2896–2900, 2001.

[4] L. Bao and J. Garcia-Luna-Aceves. Transmission scheduling in adhoc networks with directional antennas. In *Proceedings of the ACM MobiCom Conference*, pages 48–58, 2002.

[5] R. R. Choudhary, X. Yang, R. Ramanathan, and N. H. Vaidya. Using directional antennas for media access control in ad-hoc networks. In *Proceedings of the ACM MobiCom Conference*, pages 59–70, 2002.

[6] Cisco Systems. *Cisco Aironet 350 Series Client Adapters Datasheet*.

[7] T. A. ElBatt, S. V. Krishnamurthy, D. Connors, and S. Dao. Power management for throughput enhancement in wireless ad-hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, volume 3, pages 1506–1513, 2000.

[8] N. S. Fahmy, T. D. Todd, and V. Kezys. Distributed power control for ad hoc networks with smart antennas. In *Proceedings of the IEEE Vehicular Tech. Conference*, volume 4, pages 2141–2144, 2001.

[9] P. Gupta and P. Kumar. The capacity of wireless networks. In *IEEE Transactions on Information Theory*, volume 46, pages 388–404, 2000.

[10] H. Holma and A. Toskala. *WCDMA for UMTS: Radio Access for Third Generation Mobile Communications*. John Wiley and Sons Ltd, 2002.

[11] T. Korakis, G. Jakllari, and L. Tassiulas. A MAC protocol for full exploitation of directional antennas in ad-hoc wireless networks. In *Proceedings of the ACM MobiHoc Conference*, pages 95–105, 2003.

[12] J. Li and R. Compton. Maximum likelihood angle estimation for signals with known waveforms. In *IEEE Transactions on Signal Processing*, volume 41, pages 2850–2862, 1993.

[13] J. C. Liberti and T. S. Rappaport. *Smart antennas for wireless communication: IS-95 and Third generation CDMA Applications*. Prentice Hall, 1999.

[14] A. Muqattash and M. Krunz. Power controlled dual channel (PCDC) medium access protocol for wireless ad hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, volume 1, pages 470–480, 2003.

[15] R. Ramanathan. On the performance of ad hoc networks with beam forming antennas. In *Proceedings of the IEEE GlobeCom Conference*, pages 95–105, 2001.

[16] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit. Ad hoc networking with directional antennas: A complete system solution. *IEEE Journal on Selected Areas in Communications*, 23(3):496–506, Mar. 2005.

[17] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proceedings of the IEEE INFOCOM Conference*, volume 2, pages 404–413, 2000.

[18] M. Sanchez, T. Giles, and J. Zander. CSMA/CA with beam forming antennas in multihop packet radio. In *Proceedings of Swedish Workshop on Wireless Ad Hoc Networks*, pages 63–69, 2001.

[19] M. Takai, J. Martin, A. Ren, and R. Bagrodia. Directional virtual carrier sensing for directional antennas in mobile ad hoc networks. In *Proceedings of the ACM MobiHoc Conference*, pages 59–70, 2002.

[20] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, volume 3, pages 1388–1397, 2001.

[21] G. Xu and T. Kailath. Direction of arrival estimation via exploration of cyclostationarity - a combination of temporal and spatial processing. In *IEEE Transactions on Signal Processing*, volume 40, pages 1175–1185, 1992.

[22] S. Yi, Y. Pei, and S. Kalyanaraman. On the capacity improvements of adhoc networks using directional antennas. In *Proceedings of the ACM MobiHoc Conference*, pages 108–116, 2003.

[23] J. Yli-Hietanen, K. Kalliojarvi, and J. Astola. Low-complexity angle of arrival estimation of wideband signals using small arrays. In *IEEE Signal Processing Workshop on Statistical Signal and Array Processing (SSAP '96)*, 1996.

[24] W. Yu and J. Garcia-Luna-Aceves. Collision avoidance in single-channel ad hoc networks using directional antennas. In *Proceedings of the International Conference on Distributed Computing Systems*, pages 640–649, 2003.