# Ad Hoc Networking with Directional Antennas: A Complete System Solution

Ram Ramanathan, Jason Redi, Cesar Santivanez, David Wiggins, and Stephen Polit

BBN Technologies, Cambridge, MA, USA
Email:{ramanath,redi,csantiva,dwiggins,spolit}@bbn.com

***Abstract*— In this paper, we present UDAAN ("Utilizing Directional Antennas for Ad Hoc Networking"), which is an interacting suite of modular network- and MAC-layer mechanisms for adaptive control of steered or switched antenna systems in an ad hoc network. UDAAN consists of several new mechanisms – a directional power-controlled MAC, neighbor discovery with beamforming, link characterization with directional antennas, proactive routing and forwarding – all working cohesively to provide the first complete systems solution. We describe the development of a real-life ad hoc network testbed using UDAAN with switched directional antennas, and we discuss the lessons learned during field trials. High fidelity simulation results, using the same networking code as in the prototype, are also presented. For the range of parameters studied, our results show that UDAAN can produce up to a factor-of-10 improvement in throughput over omni-directional communications.**

## I. INTRODUCTION

Directional antennas have a number of advantages over omni-directional antennas for ad hoc networking. By focusing energy only in the intended direction, directional antennas significantly increase the potential for spatial reuse. They provide a longer range and/or more stable links due to increased signal strength and reduced multipath components. Increased spatial reuse and longer ranges translate into higher ad hoc network capacity (more simultaneous transmissions and fewer hops), and longer ranges also provide richer connectivity. Further, since the spatial signature of the energy is reduced to a smaller area, chances of eavesdropping are reduced, and with "smart" antennas, steering of nulls allows suppression of unnecessary interference (such as jammers). In recent years, beamforming technology has made great strides, and offers a unique and timely opportunity to unshackle the limitations of current ad hoc networks.

Replacing an omni-directional antenna by a directional one in an ad hoc network is not by itself sufficient to exploit the offered potential. The antenna system needs to be appropriately controlled by each layer of the ad hoc networking protocol stack. Such control includes pointing in the right direction at the right time for transmitting and receiving, controlling the transmit power in accordance with the antenna gains, etc. Further, mechanisms that were designed with omni-directional communications in mind – for example, medium access, neighbor discovery and routing – have to be redesigned for directional antennas. Finally, modifications to such mechanisms interact with each other – for instance, medium access control may require knowledge of how to beamform for a particular neighbor discovered by the neighbor discovery mechanism. Managing these interactions so that the mechanisms work together as a system is a significant additional challenge.

We present the first such complete system for ad hoc networking using directional antennas, called UDAAN ("Utilizing Directional Antennas for Ad Hoc Networking"). While previous works have targeted specific problems, such as medium access for directional antennas, there has been no published work on designing, implementing and fielding a complete system that uses directional antennas. Our work has not only resulted in novel mechanisms for medium access, neighbor discovery, link characterization and routing, but also has solved the challenging problem of fielding these innovations as a complete system, and has provided insight into the dynamics of a real-life environment.

## II. RELATED WORK

Much of the work on medium access has been done in the context of extending CSMA/CA (in particular IEEE 802.11) to work with directional antennas. In [1], multiple fixed directional antennas are considered and the IEEE 802.11 protocol is executed on a per-antenna basis. Steered beams are considered in [2], where the main innovation is the use of a "short NAV", to exploit the increased opportunity for spatial reuse. A "directional NAV" was proposed in [3], and, independently in [4].

A small amount of work also exists in the area of TDMA using directional antennas [6], [7].

## III. THE UDAAN SYSTEM

In the next few sections we describe some of the individual UDAAN modules and algorithms in detail. Due to space restrictions we do not describe the routing protocol, except to mention that we adapted the Hazy-Sighted Link State protocol (HSLS) [8], which is a proactive link state routing protocol. HSLS has very good scalability properties [8]. The individual modules use *Linkprofiles* to express the various modes of communicating with a neighboring node. A linkprofile is a tuple of ($band$, $beamform$). Typically a single radio servicing a single "band", and a "beamform" indicates the antenna method used for exchanging packets as one of *No Beamforming* (N-BF), *Transmit Beamforming* (T-BF), or Transmit and Receive

Beamforming (TR-BF). The use of the term "beamforming" indicates the use of a directional antenna (either switched or steered). Therefore, "no beamforming" means that the transmitter and the receiver use omni-directional antennas, while transmit beamforming means that the transmitter uses a directional antenna, but the receiver uses an omnidirectional antenna.

### A. Medium Access Control

Our directional medium access control (D-MAC) protocol is based on the single-channel CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) approach. We modify traditional CSMA/CA to support and exploit directional communications. Such *directional MAC* protocols have been studied previously (for instance [1], [2], [3], [4]). Two novel features of our solution differentiate it from previous solution approaches. The first is the use of a backoff procedure where both the interval boundaries and the method of backoff depends upon the event (e.g., no CTS, no ACK, channel busy) that caused the backoff. The second is the tight integration of power control with direction control. As shown in [5] and other works, judicious power control is critical to exploiting the potential of directional antennas. Thus, for instance, our (directional) NAV table includes power values.

Our MAC protocol is designed to be simple enough to implement within a short time in a real-life prototype while harnessing most of the spatial-reuse benefits of directional antennas. Thus, rather than build in complex protocol features to eliminate all collisions, the protocol controls the collisions using judicious backoff schemes to achieve high throughput in practice.

We now describe the protocol briefly. Due to lack of space, many important details have been omitted. The reader is referred to [9] for a comprehensive description.

*1) The D-MAC Protocol:* The general control flow is illustrated in Figure 1.

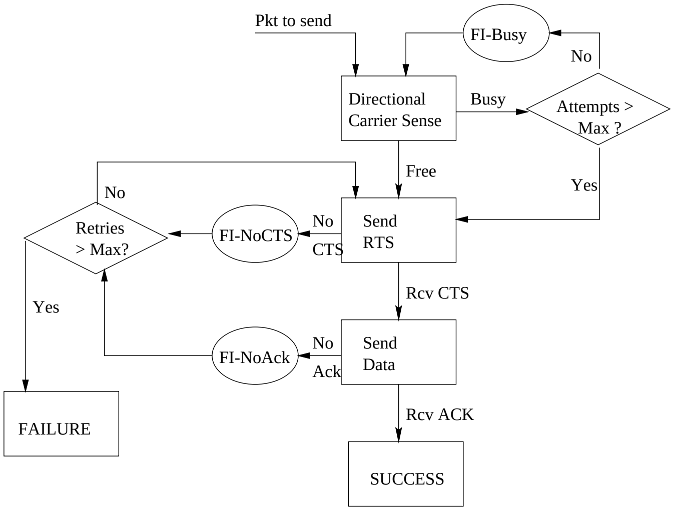

Fig. 1. High level flow chart of D-MAC.

When a node has a packet to send, it does directional carrier sensing (DCS) on the antenna corresponding to the destination for a randomly chosen period called the *DCS-period*. If the channel is free for the duration of the DCS-period, then the node sends an RTS. After the RTS is sent, the node listens on the same directional antenna and receives the CTS, sends the DATA and receives the ACK. All of this is done directionally.

The node goes into a *forced idle*[1] if one of the following happens: the channel is busy, a CTS is not received, and an ACK is not received. In forced idle the node is switched to the omni antenna and may receive packets, respond to RTS, etc. A forced idle period ends at the later of the completion of such communications and the expiration of a timer that is set when entering forced idle. A "contention window" is maintained and reset after each entry-exit sequence from the forced idle. The timer value is a random number within this contention window. As mentioned earlier, the way the contention window is reset depends upon the particular event. This is captured in table I, and explained below.

Two running variables HiFI and LowFI are maintained, along with several preconfigured constants, as explained in table I. The last column shows how the contention window changes upon each successive occurence of the event. Thus, if the channel is busy, the node backs off, tries again (no change in the window) and repeats this until a certain number of attempts is exceeded in which case the RTS is sent anyway (see figure 1). Similarly, if there is no CTS received, the node backs off using a "linear increase" and if there is no ACK received, the node implements an "exponential increase" and "exponential decrease" of the contention window. This is illustrated in the figure and the table. Further details may be found in [9].

After the ACK is received, the node goes into a forced idle to give other nodes a chance. This forced idle is set as per table I(row FI-Ack).

*a) Directional NAV and Power Control:* The first RTS for a given packet is sent at the power indicated in the radio profile (or maximum power if none is specified). The RTS contains the power $P$ and its current receive threshold ($T$). Then, the CTS is sent with a power equal to $P$ - ($R$ - $T$) + MARGIN-POWER, where R is the received signal strength indication (RSSI) of the received RTS packet, and MARGIN-POWER is a margin to account for fades etc. The DATA and ACK are power-adjusted in a similar manner.

The NAV table contains, in addition to the duration field, the antenna number[2], and the *allowed power*. This field indicates the power above which interference will occur. It may be used to transmit if it is deemed (see below) that the intended transmission is sufficiently low power so as to not bother the busy nodes.

---

[1] In most respects, the "forced idle" concept is similar to backoff. However, the term better captures the semantics in a directional antenna setting where being in idle also requires certain rules to be followed (e.g. switch to omni).

[2] We say antenna *number* because we are using a switched antenna model – but as mentioned at the beginning of this section, steered model can be mapped into this. Alternatively, the field can contain the direction, as in [3], [4]

|  | LowFI | HiFI | CWMin | CWMax | Comments |
|---|---|---|---|---|---|
| FI-Busy | No change | No change | $B_{low}$ | $B_{high}$ | Constant |
| FI-NoCTS | No change | No change | $C_{low}$*retries | $C_{high}$*retries | Linear increase |
| FI-NoAck | 0 | min(HiFi*2, $A_{max}$) | LowFI | HiFi | Exponential Increase |
| FI-Ack | $A_{init}$ | max(HiFi/2, $A_{min}$) | LowFI | HiFi | Exponential Decrease |

TABLE I

TABLE OF FORCED IDLE CONTROL. THE ROWS INDICATE THE EVENTS FOR WHICH THE FORCED IDLE IS EXECUTED. THE COLUMNS INDICATE HOW THE VALUES OF RUNNING VARIABLES LOWFI AND HIGHFI ARE RESET, AND THE CONTENTION WINDOW (CWMIN AND CWMAX) USED FOR THAT EVENT. FOR THE FI-BUSY, THE CONTENTION WINDOW ONLY APPLIES IF THE ATTEMPT NUMBER IS LESS THAN A CONFIGURED THRESHOLD, OTHERWISE THE NODE BECOMES PERSISTENT (SEE FIGURE 1). THE $A_{min}$, $A_{max}$, $C_{low}$, $C_{high}$, $B_{low}$, $B_{high}$, AND $A_{init}$ ARE CONFIGURATION PARAMETERS

The *allowed power* field is set as follows. Only RTS/CTS received when the node is in idle (omni) are processed for NAV. When an RTS (CTS) is received, the *allowed power* is set as the smaller of the current *allowed power* and $P$ - ($R$ - $T$) - *VCSMarginPower*, where $P$ is the transmitted power of the RTS (CTS) (indicated in the RTS (CTS)), $R$ is the received power, $T$ is the current receive threshold of the sender of the RTS (CTS), and *VCSMarginPower* is a configured parameter to account for fades etc.

When a packet is to be transmitted with a power $P_{xmit}$, the following deference procedure is employed. If $P_{xmit} <$ (*allowedPower* - ($G_{xmit}$ - $G_{omni}$)), then the packet is allowed to be transmitted, regardless of busy indication of the antenna. Here, $G_{xmit}$ is the gain of the antenna that is intended to be used for transmission, and $G_{omni}$ is the gain of the omnidirectional antennna. This is required because the RTS and CTS are received on omni. Otherwise, the node defers for a period indicated by the duration field in the NAV table.

## *B. Neighbor Discovery*

In order to exploit the longer-range advantage of directional antennas, UDAAN incorporates *directional neighbor discovery*, that is, the ability to discovery neighbors that can only be reached if one or both of the nodes use beamforming. As mentioned earlier, UDAAN has three kinds of links/neighbors in each band: N-BF (without beamforming), T-BF (using transmit-only beamforming), and TR-BF (using transmit and receive beamforming).

The hard problem in directional neighbor discovery, such as for T-BF and TR-BF linkprofiles, is in determining *where to point*, and *when to point* the antenna for transmit and/or receive. For instance, TR-BF neighbors can be discovered only if both the transmitter and the receivers point toward each other at the exact time a heartbeat is sent.

We have developed two methods for this problem: *informed discovery* and *blind discovery*. In informed discovery, a node $A$ has available some form of information about a non-neighbor $B$ (for instance, from the routing information) that will enable pointing. In blind discovery, a node $A$ is not even aware of the existence of a node $B$. Conventional (N-BF) neighbor discovery is a form of blind discovery. Blind discovery for T-BF and TR-BF links is, however, far more challenging, yet the only approach when the network is disconnected, and we need to create links across the partition(s).

UDAAN neighbor discovery is based on sending and scoring *heartbeats*, which are periodic control messages broadcast by each node. For informed T-BF neighbor discovery, the hearbeats are sent directionally toward a potential neighbor. A detailed description of the neighbor discovery mechanism can be found in [9]. Here we describe the main ideas behind our novel blind TR-BF mechanism.

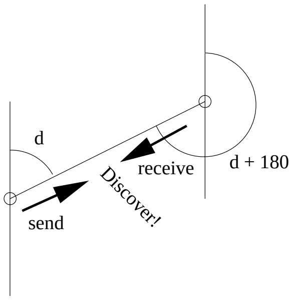

Fig. 2. Blind TR-BF discovery.

The main challenge with blind TR-BF discovery is to get two nodes that do not know of each other's existence to beamform toward each other simultaneously. To accomplish this, we require that the clocks on all of the nodes are synchronized with each other, as might be the case if the common clock source is GPS. Periodically, all nodes engaged in blind TR-BF discovery do the following (at the same time). A direction is chosen based on the time - imagine one of the hands of an analog clock. Each node alternates randomly between sending heartbeats in that direction and listening in the opposite direction for such heartbeats. For example, at a certain point in time that corresponds to 1 o'clock in the imaginary analog clock, all nodes point toward 30 degrees east of due north to send heartbeats and 30 degrees west of due south to receive heartbeats. As seen in Figure 2, for any two nodes, when the direction matches the angle between the nodes, this scheme will cause the transmit beamforming of one node to align with the receive beamforming of the other create an opportunity for the nodes to transfer a heartbeat.

It is easy to see that after only one such cycle all TR-BF neighbors that are within range will be discovered. In contrast to randomly scanning, this is a remarkable improvement without any more complexity, other than the requirement of approximately synchronized time. Although the current mechanism is restricted to 2-D, the same principle can be extended for 3-D operation.

### C. Link Characterization

The Link Characterization (Linkchar) module is designed to take information from the link layer interface and summarize it into a set of metrics for use by other modules. With the exception of occasional requested probes, Linkchar does not send any traffic to test the link, but relies on the traffic sent or received by other modules and the IP applications. Linkchar will send an occational ping-like packet to test a link if it has not received any packets from a neighbor within a pre-defined amount of time.

The design of Linkchar defines a variety of metrics. In field testing, however, we only used energy and utilization metrics, so we only describe those here.

The *minimum required energy* for transmitting over a link is a very useful metric because it gives insight into the error rate, stability, and the probability of detection when using that particular linkprofile. Because of this, the routing module uses a quantized version of this energy metric for each link when creating paths in the network (e.g. it does not necessarily use the shortest number of hops). Additionally, this active determination of minimum required energy is used by the routing module for packet-by-packet power control of the MAC.

If we assume generally symmetric pathloss between two nodes $A$ and $B$, then we can use the packets received at $A$ from $B$ to determine the amount of energy needed by $A$ to send a packet back to $B$. To do this we need to know the transmit power, the received power, and the antenna gains for that packet (assuming code, data and FEC rates stay the same).

Within each packet the MAC header indicates the transmit antenna used as well as the angle relative to North that the antenna was pointing. The header also includes the location of the sending node and the power used for transmission. At a receiver node, this information is used to determine the angles between each node's antennas, which can be used to determine the exact transmit and receive gains given preloaded 3D tables of all gains for all antennas in use.

Routing uses a quantized version of the minimum energy to build most of its paths. This allows the selection of links with a higher SNR, but by quantizing to just a few values we do not take as many additional hops as we would with strict minimum energy routing.

Linkchar additionally provides a *utilization* metric for biasing against otherwise nearly equal paths which are currently carrying large amounts of traffic. This is a node-based, not link-based metric because outgoing transmissions and incoming receptions affect all links since we only have one transceiver per band.

The idea is to attempt to determine the amount of time that a particular radio is already being used for sending or receiving. We take a particular window of time $W$ and then subtract the amount of time used for transmitting and receiving packets to us. For each packet sent or received we also subtract out the time for sending or receiving MAC control messages, performing retransmissions, or in timeout state (such as the forced idle time). The amount left over within $W$ is the residual transmit time ($T_{res}$). A higher $T_{res}$ can be considered "better" than a low $T_{res}$.

## IV. FIELD DEMONSTRATIONS

Six months after beginning the research into the UDAAN protocol design, we were given a deadline of five months to build, integrate, and test the protocols in a government defined set of field experiments. The intention of the tests was to measure the quantitative gains of using directional antennas over existing omnidirectional ad hoc networks. Specifics of all the 3 hour tests including the exact location, vehicle routes and speeds, and traffic loads, were defined. Each vehicle in the test was given a government laptop that sourced and sunk traffic, and each vehicle was driven by a government representative to insure that the systems were truly operating autonomously and motion was unbiased. The system was graded on the throughput, packet loss rates, and delay of individual traffic flows through the network.

The tests were conducted in a sparsely used 4km by 3km area which provided little other vehicle traffic to contend with. The area consisted of around one third pine tree forest corridors, one third sparsely clustered 2-4 story buildings, and one third wide open airfield. The vehicle trails varied from paved roads over flat areas to rarely used dirt paths over hilly terrain.

The first field test utilized 20 ground vehicles outfitted with a single 2.4 Ghz radio each. The government provided a baseline performance to beat by running their own 20 node test using NRL's version of OLSR, 6W power amplifiers and 6dBi antennas.

Inside each of our vehicles, the UDAAN protocol software (neighbor discovery, link characterization, routing and forwarding) was hosted on a laptop, while the D-MAC software ran on a x86 processor which was part of a custom radio board which utilized a standard CDMA 802.11b 2.4Ghz RF front end. The output of the radio was fed to a 6W 2.4Ghz power amp whose output then went to an antenna switch which fed multiple antennas. Lines were run from the x86 to the antenna switch so the MAC could switch between antennas on demand. Multiple switched antennas were used instead of an adaptive array or other steered antenna due to time and cost limitations. For this first demonstration we used a single omni antenna with 6dBi gain plus four directional antennas which had 10dBi gain at boresight and 6 dBi gain at the 45 degree cross-over points. Antennas were mounted with one pointing front, one pointing back, and one on each side.

Fig. 3. Row of demo 1 vehicles. Mounted on top of each vehicle are four directional antennas, an omni antenna, and a mast holding the GPS device away from the vehicle.

All the antenna patterns were measured at zero degrees elevation so we could use the exact antenna patterns in the simulation and well as have the Linkchar module back out receive and transmit antenna gains for use in estimating link quality.

The test performance came remarkably close to the simulation results, in large part due to the meticulous mapping of the foliage, buildings and paths through the terrain. Our system soundly beat the omnidirectional OLSR network tested by the government, particularly in terms of delay, but also in terms of throughput capacity.

The second demonstration added the requirements of a helicopter, the addition of high-band (38Ghz) at six of the nodes, multicast requirements, and ToS-based metrics. This demonstration also consisted of 20 nodes, but they were more scattered throughout the area.

The following are a few of the lessons we learned from our demonstration test experiences.

- Real antenna patterns are far more complex than “pie-slices” or “cone plus ball” models. Real antennas often have significant side and back lobes that affect performance of a directional ad hoc network due to the additional noise. Increases in performance are not necessarily gained from higher gain antennas unless the average front-to-back and front-to-side gain ratios are high.
- When switched beams are used in conjunction with an omni-directional beam, sufficient overlap must exist between the directional beams so that the choice of a directional antenna never results in a poorer gain than an omni-directional one.
- Using position information to select the right beam mostly works, but not always. Sometimes, multipath and other factors might make an antenna that points in a different direction a better choice

## V. EXPERIMENTAL RESULTS

We have used OPNET for the simulation results presented here. We have used the same networking code (except for MAC) in our simulations, as in the testbed. Thus, the simulations are of very high fidelity. The downside is that simulations take very long to run, thereby reducing the number of data points that can be generated in a given time.

All simulations are with a 20 node ad hoc network. The nodes are placed randomly in a 2-dimensional square area of varying size depending upon the density parameter. Mobility is similar to the random waypoint model, but nodes do not stop at waypoints.

For all of the results presented in the network, 20 streams are originated, one per node, with the destination chosen randomly. Each stream consists of packets of size 8192 bits and the inter-arrival time is uniformly distributed around a mean rate of 10 pps per stream for some experiments and 50 pps for some others. The raw bit rate of the radio for the random scenarios is 11 Mbps.

Figures 4 - 6 illustrates the throughput performance of UDAAN over the random scenarios. The throughput is computed as the total number of bits delivered successfully at the destination.

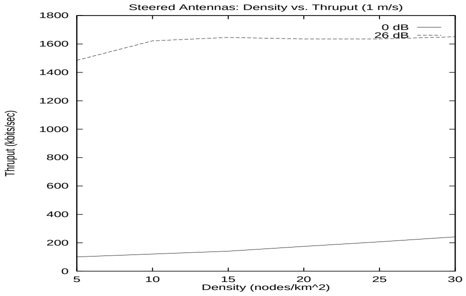

Fig. 4. Throughput dependence on density

Figure 4 shows that with a steered antenna of gain 26 dBi, the throughput is increased by a factor of about 8 (at density 30) to a factor of about 10 (at density 5). Both throughputs increase with increasing density, although the directional case reaches maximum at density 10. The main reason for the large throughput difference here is due to the network being disconnected when omni-directional antennas are used, but well connected with directional antennas (thanks to the longer links provided by directional neighbor discovery).

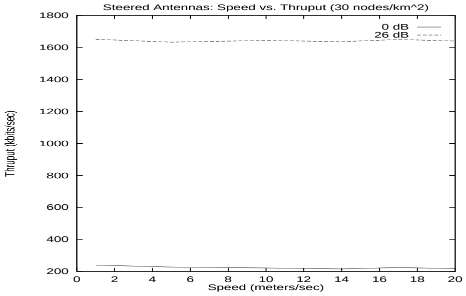

Fig. 5. Dependence on speed

The throughput versus speed plot (figure 5) also shows about a factor of 8-10 improvement in throughput with the directional antenna. The throughput in each case is largely unaffected by speed, because for the range of speeds tested, the hazy sighted link state routing was able to adapt quickly enough.

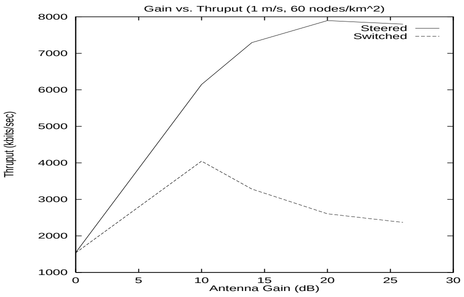

Fig. 6. Effect of antenna gain

Figure 6 compares throughput using switched and steered antennas as a function of antenna gain. An important fact to note here for the switched antenna curve is that the number of antennas is kept constant. Thus, when gain is increased (and beamwidth is decreased), the azimuthal "coverage" (the fraction of the 360 degree plane with at least 0 dBi gain) goes down. Interestingly, as gain increases, the performance of switched antenna still increases upto a certain gain value (10 dBi) indicating that the increased gain makes up for the lost coverage. But after that the decreased azimuthal coverage starts to hurt us. For steered antennas, this issue does not arise as the beam can be pointed wherever desired. Thus, the performance continues to increase with gain. Predictably, steered antenna provides better throughput than switched antennas, but the difference is not high at low gains (and probably won't be at higher gains too, if the number of switched antennas is increased commensurately).

A large number of additional plots have been generated as part of the simulation study, but could not be included for lack of space.

## VI. SUMMARY

Directional antennas offer potential for an order-of-magnitude improvement in the capacity and connectivity of an ad hoc network. Translating this potential into reality requires support for antenna control at several layers of the protocol stack. While specific mechanisms have been developed in prior work - mostly at the MAC layer - the *complete* system design problem for an ad hoc network with directional antennas had been heretofore far unsolved.

In this paper we presented the first complete system solution for utilizing directional antennas in ad hoc networks (UDAAN). UDAAN supports both switched and steered antennas. UDAAN makes several contributions, some of which where described: a novel CSMA/CA based directional MAC (D-MAC) protocol that manages backoffs in a way that is suited to directional antennas, and fully integrates power control as part of its operation; a novel neighbor discovery mechanism that guarantees TR-BF discovery within one cycle time; and assignment of QoS metrics using link characterization when directional antennas are present.

Lacking in this paper is a comparison our solutions to any existing work. A chief reason for this is the paucity of prior work in all but medium access control. In particular, there has been no previous system described that we could model.

We believe there is considerable scope for further research in this area. In our opinion, the key research areas that will most benefit state of the art include: power-controlled directional (multichannel) MAC, integrated power/antenna topology control, reactive ad hoc routing based directional antenna system, new ways of getting pointing direction information (maybe without positions), and prototype development using small smart antennas.

## REFERENCES

[1] Y.B. Ko and N.H. Vaidya, "Medium Access Control Protocols Using Directional Antennas in Ad Hoc Networks," *Proc. of IEEE INFOCOM*, March 2000.

[2] N.S. Fahmy, T.D. Todd, V. Kezys, "Ad Hoc Networks with Smart Antennas Using IEEE 802.11-Based Protocols," *Proc. IEEE ICC*, 2002.

[3] M. Takai, J. Martin, A. Ren, R. Bagrodia, "Directional Virtual Carrier Sensing for Directional Antennas in Mobile Ad Hoc Networks," *Proc. ACM MOBIHOC*, Lausanne, Switzerland, June 2002.

[4] R.R. Choudhury, X. Yang, R. Ramanathan, N. Vaidya, "Using Directional Antennas for Medium Access Control in Ad Hoc Networks,"*Proc. ACM MOBICOM*, Atlanta, Georgia, September 2002.

[5] R. Ramanathan, "On the Performance of Ad Hoc Networks with Beamforming Antennas," *Proc. ACM MobiHoc*, Long Beach, CA, October 2001.

[6] K. Dyberg, L. Farman, F. Eklof, J. Gronkvist, U. Sterner, J. Rantakokko, "On the Performance of Antenna Arrays in Spatial Reuse TDMA Ad Hoc Networks," *Proc. IEEE MILCOM*, Anaheim, California, October 2002.

[7] L. Bao, J.J. Garcia-Luna-Aceves, "Transmission Scheduling in Ad Hoc Networks with Directional Antennas," *Proc. ACM MOBICOM*, Atlanta, Georgia, September 2002.

[8] C. Santivanez, S. Ramanathan, I. Stavrakakis, "Making Link State Routing Scale for Ad Hoc Networks," in *Proceedings of MobiHoc 2001* Long Beach, CA, Oct. 2001.

[9] "Utilizing Directional Antennas for Ad Hoc Networking – System Design Document", BBN Technical Report No. 8370, Cambridge, MA, 2003. http://www.ir.bbn.com/projects/udaan/tech-rep.pdf

## VII. ACKNOWLEDGEMENTS

The design and initial implementation of UDAAN was sponsored by DARPA/ATO and Army Research Laboratory under contract number DAAD19-01-C-0027. Content of this work does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

The integration and demonstrations of the system was performed as part of a team lead by Northrop Grumman Mission Systems (formerly TRW Systems) under contract number MDA972-01-9-0023.