

On the Performance of Ad Hoc Networks with Beamforming Antennas*

Ram Ramanathan
Internetwork Research Department
BBN Technologies (A part of Verizon)
Cambridge, Massachusetts, USA
ramanath@bbn.com

ABSTRACT

Beamforming antennas have the potential to provide a fundamental breakthrough in ad hoc network capacity. We present a broad-based examination of this potential, focusing on exploiting the longer ranges as well as the reduced interference that beamforming antennas can provide. We consider a number of enhancements to a conventional ad hoc network system, and evaluate the impact of each enhancement using simulation. Such enhancements include “aggressive” and “conservative” channel access models for beamforming antennas, link power control, and directional neighbor discovery. Our simulations are based on detailed modeling of steered as well as switched beams using antenna patterns of varying gains, and a realistic radio and propagation model. For the scenarios studied, our results show that beamforming can yield a 28% to 118% (depending upon the density) improvement in throughput, and up to a factor-of-28 reduction in delay. Our study also tells us which mechanisms are likely to be more effective and under what conditions, which in turn identifies areas where future research is needed.

1. INTRODUCTION

Since the early days of ad hoc networking, researchers have strived to increase the capacity of ad hoc networks through a variety of innovative techniques. Along the way, it has been realized that there are fundamental limitations to how high one can push the capacity. Underlying this limitation is the inherent tension between the number of end-to-end packet hops and the spatial reuse per hop – that is, if one wishes to decrease the number of multihop transmissions end-to-end (and thereby reduce the demands on the network), then one has to increase the range per transmission, which results in a reduction in the number of simultaneous transmissions. This has been thoroughly examined in [1], who show that the throughput obtainable by each node is $\Theta(\frac{W}{\sqrt{(n \log n)}})$, where

W is the data rate, and n is the number of nodes. This limitation on capacity exists no matter what routing protocol or channel access scheme is used. Splitting the channel into subchannels does not change this result [1].

One of the chief contributors to this capacity limitation is the omni-directional nature of transmissions. Specifically, the distribution of energy in directions other than the intended direction not only causes unnecessary interference to other nodes, but also reduces the potential range of the transmission (due to lower signal strength and multipath components). On the other hand, with *directional* communications, both spatial reuse and range – the two contributors to the capacity limitation – would be simultaneously enhanced. Indeed, the capacity limit in [1] assumes omni-directional transmissions and acknowledges, without details, that beamforming will be “advantageous”. In recent years, beamforming technology has made great strides, and offers an unique and timely opportunity unshackle the capacity limitations of ad hoc networks.

While the use of directional communications in ad hoc networks promises to open new doors, a number of questions arise. How can we achieve directional communications? What kinds of antennas are available for the purpose? Are they too expensive and/or bulky to make sense? What implications does it have for *networking* – is it just a question of replacing the omni-directional antenna with a smart antenna or will networking protocols have to be changed to harness the potential? Which aspects of networking are important in harnessing this potential, and which are not? Does it depend upon the kind of antennas used? How much performance improvement can we get? Is it worth the trouble?

These questions are addressed in toto for the first time in this paper. We consider the applicability of *beamforming antennas* – antennas that have one or more steerable (pointable) directional beams. We evaluate, using modeling and simulation, the performance improvement that beamforming antennas can provide relative to omni-directional ones, and the factors upon which this depends. Our contributions include the following. We evaluate the effectiveness of a number of enhancements to a traditional ad hoc networking protocol stack, including two channel access approaches, link power control, and directional neighbor discovery. In particular, using the last enhancement, we give the first assessment of the benefits that longer-range directional links can bring to the robustness

*

(connectivity) and latency of an ad hoc network. We employ a detailed and thorough modeling of beamformed communications, using specific antenna patterns of different gains, and modeling of side as well as main lobes. Finally, we consider and compare steered (e.g. adaptive arrays) and switched (e.g. diversity) beamforming. The results of our study provide the motivation and methodology for using beamforming antennas as a “first class citizen” in ad hoc networking.

Many of the benefits of a beamforming antenna such as higher signal-to-noise ratio, multipath mitigation, etc. are relevant to and can be exploited at the *physical* layer. However, this is not the subject of this paper. Rather we ask: what *networking and medium access* advantages does beamforming provide, and how can they be exploited? In particular, we consider the exploitation of two advantages: the reduced interference due to the narrower beamwidths, and the extended range due to higher signal-to-noise ratio (by virtue of the higher gain and reduced multipath). Our focus is on examining the *relative* performance improvement over omni-directional antennas, and specifically on establishing a *lower bound*. As will become apparent later, this goal drives many of our assumptions and experimental methodology. While our work is in the context of a proactive ad hoc routing protocol, most of our results independent of the routing protocol per se.

2. RELATED WORK

While there has been a vast amount of research on physical layer issues related to beamforming antennas, only a small amount of work considers the medium access and network layer implications of beamforming antennas for ad hoc networking. Zander [2] has proposed the use of directional antennas in a slotted ALOHA multihop packet radio network. More recently [3] present several medium access control protocols for fixed directional antennas. These protocols utilize physical location information to implement collision avoidance. Subsequently, [4] suggest a protocol that uses signal strength information instead of position information. Beamforming antennas for single-hop packet radio networks have been studied in [5, 6] and other papers. The use of directional floods to limit the scope of route requests in an “on-demand” ad hoc routing protocol is suggested in [7], and explored in [8].

Our work is unique in several ways. First, most of the work cited above have restricted themselves to a specific aspect of utilizing directional antennas, such as medium access, whereas we model the effects of a number of mechanisms apart from medium access such as directional neighbor discovery, link power control, etc. and consider their interactions and effectiveness within a full-fledged ad hoc networking system. Second, none of the above works evaluate the impact of the longer-range links possible with directional antennas in improving the connectivity and reducing latency. Third, the only multihop specific works above consider multiple fixed directional antennas whereas we consider the support for more sophisticated beamforming such as steered beams. Finally, our study is based on a far more comprehensive antenna and propagation model with multiple antenna patterns, modeling of sidelobes, and the use of antenna transmit and receive gains in the bit-error-rate calculation at each node. In summary, this is the first broad-based, realistic analysis of how much beamforming will help, and what it requires of the MAC and network layers.

3. BEAMFORMING ANTENNAS: CONCEPTS, TERMINOLOGY, RELEVANCE

In this section, we discuss some concepts related to beamforming antennas. This is not intended to cover all aspects of this technology, nor do we cover it precisely or formally. Rather, the idea is to give the basics in an informal and intuitive fashion to equip the reader unfamiliar with this topic with just enough knowledge to understand the remainder of this paper. Readers familiar with beamforming antennas may skip this section. Readers wishing to explore this field in detail are referred to [9] and the citations therein.

3.1 Antenna Concepts

Radio antennas couple energy from one medium to another. An *omnidirectional antenna* (sometimes known as an *isotropic antenna*) radiates or receives energy equally well in all directions¹. A *directional antenna* has certain preferred transmission and reception directions, that is, transmits/receives more energy in one direction compared to the other.

The *gain* of an antenna is an important concept, and is used to quantify the directionality of an antenna. The gain of an antenna in a particular direction $\vec{d} = (\theta, \phi)$ is given [9] by

$$G(\vec{d}) = \eta \frac{U(\vec{d})}{U_{ave}} \quad (1)$$

where $U(\vec{d})$ is the power density in the direction \vec{d} , U_{ave} is the average power density over all directions, and η is the efficiency of the antenna which accounts for losses. Informally, gain measures the relative power in one direction compared to an omnidirectional antenna. Thus, the higher the gain, the more directional is the antenna. The *peak gain* is the maximum gain taken over all directions. When a single value is given for the gain of an antenna, it usually refers to the peak gain. Gain is often measured in unitless decibels (dBi), that is, $G_{dBi} = 10 \cdot \log_{10}(G_{abs})$. An omni-directional antenna has a gain of 0 dBi.

An *antenna pattern* is the specification of the gain values in each direction in space, sometimes depicted as projections on the azimuthal and elevation planes. It typically has a *main lobe* of peak gain and (smaller gain) *side lobes*. As is common practice, we use the word *beam* as a synonym for “lobe”, especially when discussing antennas with multiple/controllable beams/lobes.

A related concept is the antenna *beamwidth*. Typically, this means the “3 dB beam width”, which refers to the angle subtended by the two directions on either side of the direction of peak gain that are 3 dB down in gain. Gain and beamwidth are related. Typically, the more directional the antenna, higher the gain and smaller the beamwidth. However, two antennas with the same gain could have different beamwidths – for instance, the antenna with the smaller main lobe width may have more or larger sidelobes.

¹In reality, no antenna is perfectly omnidirectional, but we use this term to represent any antenna that is not intentionally directional.

3.2 “Smart” Beamforming Antennas

The simplest way of improving the “intelligence” of antennas is to have multiple elements. The slight physical separation between elements, or *diversity* can be used counteract multipath effects. There are two well-known methods. In *switched diversity* the system continually switches between elements so as to always use the element with the best signal. While this reduces the negative effects of fading and multipath, there is no increase in gain. In *diversity combining*, the phase error of multipath signals is corrected and the power combined to both reduce multipath and fading, as well as increase the gain.

The next step in sophistication involves incorporating more control in the way the signals from multiple elements (the antenna *array*) are used to provide increased gain, more beams and beam agility. Again, there are two main classes of techniques, as described below.

In *switched beam* systems, multiple fixed beams are formed by shifting the phase of each element’s signal by a predetermined amount (this is done by a *beamforming network*). The transceiver can then choose between one or more beams for transmitting or receiving. While providing increased spatial reuse, switched beam systems cannot track moving nodes which therefore experience periods of lower gain as they move between beams.

In a *steered beam* system, the main lobe can be pointed virtually in any direction, and often automatically using the received signal from the target using sophisticated “direction-of-arrival” techniques. One may distinguish between two kinds of steered beam systems – *dynamic phased arrays* which maximizes the gain toward the target, and *adaptive arrays* which additionally minimize the gain (produce *nulls*) toward interfering sources.

In this paper, we consider switched beam and steered beam antennas, jointly referred to as a *beamforming antenna*.

3.3 Relevance for Ad Hoc Networks

When considering the use of beamforming antennas for ad hoc networks, a question is: *Aren’t beamforming antennas too expensive and/or too big for ad hoc networks?* In this section, we argue that there do exist antenna techniques with suitable price and form-factor combinations.

Applications for ad hoc networking may be classified broadly into three categories: military, commercial outdoor, and commercial indoor, each with its own distinctive profile, and able to accommodate different antenna technologies.

Military networks, which are by far the most prevalent application of *mobile* ad hoc networks, contain a significant number of large nodes (such as tanks, airplanes). The size of these platforms makes the form factor of most antennas quite irrelevant. Further, each platform by itself is so expensive that the cost of even the most sophisticated antenna is dwarfed by comparison. Thus, beamforming antennas are extremely relevant to military networks. An added bonus is that use of directional transmissions have better immunity to jammers and eavesdroppers.

Fixed ad hoc networks for commercial outdoor infrastructure extend the reach of base stations using wireless repeaters organized into a mesh network. Here steered beam approaches may be too expensive. However, switched beams using inexpensive beamforming networks such as the Butler matrix [10] are easily manufactured using inexpensive hybrid couplers etc. [9] making switched beamforming quite relevant.

The biggest deterrent to using beamforming antennas for networking small nodes such as PDAs and laptops within an indoor environment is the size. At 2.4 GHz, and the typical half-wavelength element spacing, an eight element cylindrical array would have a radius of about 8 cm, making it quite unwieldy. However, as the operating frequency continues to increase (already the IEEE 802.11a is working on wireless LANs in the 5 GHz band), the antenna sizes will shrink. At the 5.8 GHz ISM band, the 8-element cylindrical array will have a radius of only 3.3 cm, and at the 24 GHz ISM band, a mere 0.8 cm. Thus, the future looks bright for applying beamforming technology even to such applications.

Thus, while at first glance it may seem that ad hoc networks and beamforming antennas are not compatible, a more careful examination opens up a number of possibilities.

4. PROTOCOL MODELS

In this section, we describe the details of each key protocol used in the simulation system. The simulation system was developed using the OPNET modeling and simulation tool, version 6.0. Compared to more commonly used ad hoc network simulation tools such as ns-2 and GloMoSim, OPNET offers better support for directional communications. Specifically, OPNET allows the creation of arbitrary antenna gain patterns in 3-D, allows the pointing of the main lobe toward an arbitrary position in 3-D and automatically computes the energy received at every node in the system taking into account the different gains in different directions. It also supports a detailed modeling of path loss, SNR, SIR, and bit-error-rate computations. These features allow for a very realistic modeling of directional radio communications.

Our simulation system contains some *abstract* models. An *abstract* model of a protocol is one where some of the protocol dynamics are replaced with “short-cut” equivalents. For instance, instead of modeling the generation, transmission and reception of a control message over the channel, the message could be sent through a “back-door”, that is, a single event that delivers the necessary contents to the relevant module. Another example is when multiple short messages are replaced with a single long message (or a “packet train”).

Abstraction is a way of trading off fidelity for shorter running times and quicker development. It is appropriate when the purpose is to determine the rough-order-of-magnitude performance figures, or when comparing mechanisms where the abstraction tends to affect each mechanism more or less equally, so that the *relative* performance is not affected as much². Both these cases apply to our work, as well as the ubiquitous

²For instance, we have not modeled the explicit sending of RTS/CTS messages. An equal number of such messages have to be sent in the omni-directional as well as directional case, and hence have little impact when *relative* performance is considered.

need for short running times. Therefore, some of our models are abstract.

We now discuss the protocol models. We begin with the beam forming and pointing model, which, although not a “protocol”, forms a key component of any ad hoc networking study with beamforming antennas.

4.1 Antenna Patterns and Beam Control

Each node in the simulation system has an antenna that can be associated with one of a number of predefined antenna patterns. Beam steering is modeled by orienting the pattern so that the center of the main lobe points toward the target node. We first describe our antenna patterns and then the pointing techniques.

4.1.1 Antenna Patterns

Since modeling a real antenna with precise values for main and sidelobes is difficult and deviates from the focus of this paper, we use an approximate antenna pattern as follows. Given a gain value g_m , the antenna pattern for this gain consists of a main lobe of beamwidth θ_m , and a sidelobe of gain g_s of beamwidth $(2\pi - \theta_m)$. That is, the main lobe is a cone of uniform gain, and the sidelobes are aggregated to a single “bulb” at the base of the cone. Figure 1 illustrates one of the patterns used in our simulations.

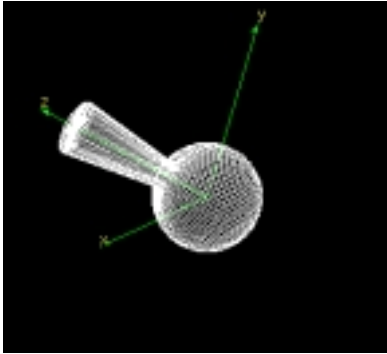


Figure 1: Antenna pattern 20 dBi

As discussed in section 3, a directional antenna merely redistributes the energy. Thus, the choice of θ_m and g_s must be chosen to provide a realistic model of an antenna beam, while at the same time not breaking any laws of nature. Specifically, the question is: in order to model an antenna as in figure 1 with a given (main lobe) gain of g_m , what should the its beamwidth θ_m and sidelobe gain g_s be, so that the total energy is the same. This question will arise in any work that wishes to use a quick, yet reasonable model for directional antennas. We derive below a formula for computing θ_m and g_s , given a gain value g_m , and hope that it might be of use to other researchers.

First we derive the approximate maximum beamwidth θ_{max} possible for g_m . Consider a sphere of some radius r (no loss of generality). The surface area A on the sphere for a beamwidth of θ_m can be approximated as a circle of radius $r \tan(\theta_{max}/2)$. Let S be the surface area of the sphere, and P the emanated power. By definition of gain, we have

Gain (g_m)	Beamwidth (θ_m)	Sidelobe gain (g_s)
10 dBi	60 deg	-7.4 dBi
14 dBi	40 deg	-7.6 dBi
20 dBi	20 deg	-6.5 dBi
26 dBi	10 deg	-4.0 dBi

Table 1: Antenna pattern parameters used in our study

$$g_m = \frac{P/A}{P/S} = \frac{4\pi r^2}{\pi \cdot (r^2 \cdot \tan^2(\theta_{max}/2))} \quad (2)$$

Solving for θ_{max} ,

$$\theta_{max} = 2 \cdot \tan^{-1} \sqrt{\frac{4}{g_m}} \quad (3)$$

We then pick a $\theta_m < \theta_{max}$. The difference leaves energy that can be used for sidelobes, and is dependent upon how much sidelobe gain we want to model. In this study, we simply choose θ_m to be the largest multiple of 10 that is less than θ_{max} , for purely pragmatic reasons (OPNET antenna pattern granularity).

Given θ_m and g_m , we now derive g_s to complete the model, as follows. By definition,

$$g_m \cdot U_{av} \cdot A + g_s \cdot U_{av} \cdot (S - A) = \eta \cdot P \quad (4)$$

where U_{av} is the average power density, and is given by P/S . Letting Δ represent S/A , substituting, and simplifying, we have

$$g_s = \frac{\eta \cdot \Delta - g_m}{\Delta - 1} \quad (5)$$

where $\Delta = \frac{4}{\tan^2 \frac{\theta_m}{2}}$ (using expressions for S and A , similar to equation 2). As an aside, Δ is an often used quantity called the *directivity* of the antenna.

Using this formula, we have generated and used four antenna patterns with parameters illustrated in the table 1.

4.1.2 Beam Steering and Switching

Beam steering is modeled as follows. Assume that a pattern P corresponding to gain g has been associated with a node S . Suppose that S wants to send a packet using a steerable beam to another node R .

After S has obtained access to the channel (see section 4.2 below), it sets its pattern to P and the main beam is pointed in the direction of R . The packet is then sent for transmission. Immediately after the packet is transmitted, S sets its pattern back to omni-directional. Thus, we only use transmit directionality, all receives are omni-directional.

Beam switching is modeled as follows. Suppose that the antenna contains K switched beams. We assume that the beams are identical (have a pattern P_g), and are equally spaced. Without loss of generality, the direction of each beam is fixed upon startup, with direction of antenna i is given by $D_i = (i - 1) * 360/K$ degrees relative to due east. When K not divide 360, the directions are *approximately* given by the above equation.

When a node S wants to transmit a packet to a node R , it determines the direction D that R is relative to itself. It then determines D_j such that $|D - D_j|$ is the minimum over all D_i , $i = 1$ through K . That is, it finds the beam of minimum angular separation with the intended direction. After S has obtained access to the channel (see section 4.2 below), it sets its pattern to P and the main beam is pointed in the direction D_j . The packet is then sent for transmission. Immediately after the packet is transmitted, S sets its pattern back to an omni-directional pattern. Thus, we only use transmit directionality, all receives are omni-directional. We note that switched beam modeling does not require construction of a K -lobed pattern or K beams.

In both steered and switched beam models, it is necessary to find the relative direction of a neighboring node. This could be done either using position information or direction-or-arrival techniques within the antenna controller. We assume that the latter is used, and hence do not model the position information dissemination³.

4.2 Channel Access

Our channel access protocol model is based on the CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) approach. Traditional CSMA/CA works as follows. When the carrier is free, a sending node sends a Request-to-Send (RTS) packet to the intended receiver. The receiving node replies with a Clear-to-Send (CTS). This exchange acquires the channel “floor” for the communication by having any node that hears this exchange desist from transmitting, and hence avoids collision. Several protocols published in the literature are representatives of this approach, including [11], and the the standardized wireless LAN mechanism IEEE 802.11 [12].

The goal of CSMA/CA is to have a node transmit a packet if and only if there will be no collisions. In particular, not only should it be blocked if there will be a collision, but it should also *not* be blocked if there will be no collision. CSMA/CA meets this goal admirably in ad hoc networks when omni-directional antennas are used. However, with directional transmissions, this is not true whether one uses omni-directional or directional RTS/CTS. One example of each is illustrated in figure 2.

In figure 2(a), A wants to send a packet to B, and C wants to send a packet to D. The RTS from A is sent omni-directionally or directionally to B, but is heard by C and C is inhibited from sending even though it can do so without interfering with the transmission from A to B.

³If one wishes to use position information, then the predicted performance improvement should be downgraded slightly to reflect the overhead of conveying this information, since traditional omni-directional communications does not require this.




Figure 2: Examples when traditional RTS/CTS is insufficient when beamforming is employed.

In Figure 2(b), suppose A is sending a packet to B after having initiated an RTS-CTS exchange. Neither the RTS nor the CTS is heard by C, which proceeds to initiate a transmission to D. The RTS-CTS exchange between C and D is directional and not heard by A, which, after completing its transmission to B, now initiates a transmission to E. The RTS from A to E interferes with the data being received by D from C.

Thus, with directional communications, the receipt of an RTS or CTS does not imply that you must be blocked (refer figure 2(a)) and the non-receipt of an RTS or CTS does not imply that you can transmit (refer figure 2(b)). More generally, there is a spectrum of tradeoffs between having better spatial reuse with parallel transmissions (but more collisions), and having lesser number of collisions (but less spatial reuse). A good solution to the problem should be parametrized to allow for a selectable tradeoff in this spectrum.

This paper does not attempt to *solve* this general problem. Rather, we consider the two ends of the spectrum, and have modeled two protocols called *aggressive CA* and *conservative CA*, described briefly below. Both are over and above a basic CSMA protocol, and operate when the packet has passed the carrier sensing. In both cases, the RTS and CTS are transmitted omni-directionally and are assumed to have a range at least that of the directional antenna⁴. Both are abstract models, that is, the RTS and CTS are not actually sent over the air – this helped us reduce the running time by an order of magnitude.

- *Aggressive Collision Avoidance Model.* This models a protocol in which a node is *never* blocked upon receiving an RTS or CTS. The handshake is used only for ensuring that the receiver is not already busy sending or receiving. In particular, no attention is paid to the receive status of other nodes. Thus, this could potentially cause collision at nodes other than the intended receiver.
- *Conservative Collision Avoidance Model.* This models a

⁴This can be done using larger power, higher processing gain or a lower frequency for the RTS and CTS.

protocol in which a node is *always* blocked upon receiving an RTS or CTS. This is similar to the traditional collision avoidance approach. In particular, a transmission takes place only if none of the nodes in its range are busy. When used with directional antennas, this mode sometimes passes up on collision-free transmission opportunities.

We emphasize that we do not suggest these as a solution for the beamforming MAC problem. Rather, by analyzing the performance for these two extremes, we can place a lower bound on the performance of a future solution to this problem, presumably one that combines the best characteristics of each. Working toward a lower bound on the performance improvement with beamforming antennas is consistent with the goals of this paper (see the last paragraph in section 1). As will be seen in section 5, even with such a trivial protocol, one can get huge performance gains.

4.3 Link Power Control

The system can be configured to do link power control on a per-packet basis. The idea here is to control the transmit power so that it is just about sufficient to activate the link. This reduces interference, as well as battery consumption.

A straightforward way of doing this within the RTS/CTS framework is as follows. The RTS is sent at a predetermined power (for instance, the maximum power). The receiver determines the difference δ between the received power for the RTS and its receive threshold. The value of δ is sent along with the CTS. When transmitting the DATA, the sender uses a power that is δ less than the power used for the RTS. Optionally, it may add a small “margin” to the value to account for fading and mobility. This mechanism is widely implemented in military radios. A variant of this idea is described in [13].

In our model, there are no RTS/CTS packets actually sent, and therefore, we use an abstract power control model which monitors the received power at the target node using a “backdoor” mechanism, and adjusts its transmit power accordingly.

4.4 Neighbor Discovery

Neighbor discovery is done using a Hello protocol. Each node periodically (with a small random jitter) transmits a Hello packet. Received Hellos are tracked over a history (sliding) window of size N . A neighbor is deemed up if the number of Hellos received from that neighbor over the current window is greater or equal to a configured value K_{up} . If the number of Hellos received is lower than a configured value K_{down} the neighbor is deemed unreachable and the link is erased.

While neighbor discovery is fairly straightforward with omnidirectional antennas, a number of interesting research problems arise when beamforming is considered. For instance, consider two nodes A and B who are out of each other’s range when omnidirectional transmissions are used but can communicate when transmit and/or receive beamforming is used. The problem is that A and B must both determine *independently* – without communicating with each other – where to point and when to point so that they are pointing

at each other⁵.

In this paper, we consider two cases – *omni-directional neighbor discovery*, and *directional neighbor discovery*. In the first, all Hellos are sent omnidirectionally (but data packets may be sent directionally). Thus, the routing topology is exactly the same as one would get when beamforming does not exist. In the second case, which we only use with switched beams⁶, Hellos are sent directionally. In particular, a Hello is sent on each of the beams. Since these Hellos travel further, the potential topology is richer than in the first case. This is a detailed model, and incorporates all control messages with high fidelity.

4.5 Routing

We use the well-known link-state routing protocol. Although we are aware of its scalability limitations, it is not of concern to us in this study as we consider only 40 node networks with no mobility. Further, the overhead induced in most of our experiments (the ones based on omnidirectional neighbors) is the same whether or not beamforming is used for data, and does not affect the *relative* performance.

Briefly, link-state updates are triggered by a node when one of its links goes up or down, as deemed by neighbor discovery. Updates are flooded throughout the network, updating the forwarding table. Route generation is done using Dijkstra’s shortest-path algorithm. For ad hoc networks, scalable versions of this “proactive” or “table-driven” approach abound, such as [14]. This is a detailed protocol model and incorporates all control messages with high fidelity.

5. EXPERIMENTAL RESULTS

We first describe our simulation environment, and then discuss the results. Due to the abstract nature of some of the models, we caution the reader that while the *relative* performance improvements are predicted with some degree of accuracy, the *absolute* performance numbers are not.

5.1 Simulation Environment

OPNET has a detailed propagation model where the bit error rate on a packet is computed based on the signal to interference and signal to noise ratios. The radio model is based on a direct sequence spread spectrum radio, with QPSK modulation and a data rate of 1.6 Mbps. The propagation model uses the following equation, from [15], to compute the received power.

$$P_r(d) = \frac{P_t \cdot \lambda^2 \cdot d_{ref}^2 \cdot G_t \cdot G_r}{4\pi^2 \cdot d^4} \quad (6)$$

where P_t and P_r are the transmit and received powers respectively, λ is the wavelength, G_t and G_r are the transmit antenna and receive antenna gains respectively, d is the distance between the nodes, and d_{ref} is a reference distance

⁵Alternatively, one transmit omnidirectionally but using a higher processing gain, thereby trading data rate for longer range. We do not consider this or other methods in this paper.

⁶Directional neighbor discovery using steered beams can be more powerful, but is much more challenging and beyond the scope of this paper.

given by $d_{ref} = 2D/\lambda^2$, where D is the maximum antenna dimension.

All simulations are with a 40 node ad hoc network. The nodes are placed randomly in a 2-dimensional square area of varying size depending upon the density parameter. We study the dependance of the system on the following parameter ranges.

- Densities 4, 8, 16, 32, 48, 80, and 112 nodes/sq mile. These correspond to average node degrees (number of neighbors in the network topology, with omni-directional neighbor discovery) of approximately 3, 5, 9, 16, 22, 30, and 38 respectively.
- Gains 10, 14, 20, and 26 dBi. These correspond to beamwidths of 60, 40, 20, and 10 degrees respectively (refer table 1). The baseline case is an omni-directional antenna, represented by a beam with gain 0 in the plots.
- For the switched beamforming case, antennas with 4,8, and 12 beams are considered.

We consider two performance metrics: *throughput*, the percentage of packets sent by any source that was successfully received at the intended destination; and *delay*, the average time elapsed, for all successful packets, between the packet being sent by a source and it being received. Each result is an average over 3 runs with random seeds.

For all of the results presented in the network, 20 streams are used, between randomly chosen source-destination pairs. Each stream consists of packets of size 1800 bits and the inter-arrival time is uniformly distributed around a mean rate of 90 kbps per stream. Nodes have buffers of 30 packets, and packets that find the queue full are dropped. This represents a high but not excessive load on the network. The use of buffering implies that capacity improvements are often reflected as reduced delay rather than increased throughput.

In all cases, receiving is done omni-directionally. Further, in all steered beam cases, control messages are sent omni-directionally. An implication of this is that the network topology used does not change as the gain is increased. One set of results reported here has control packets being sent directionally with switched beam antennas in which case the topology gets richer with increasing gain.

Only stationary networks are simulated. The main reason for this is that the advantages we wish to evaluate – spatial reuse, and longer range – are largely orthogonal to the mobility of the nodes. In other words, since the *same* routing algorithm is used for both omni-directional and directional experiments, and the MAC layer is mobility impervious, the numbers for the *relative* performance between directional and omni-directional communications will not change significantly for a mobile scenario. Given this, it makes sense to use the available simulation resources (CPU time) to model a wider range of parameter settings and a larger variety of capabilities such as power control and switched beams, which, as we shall see, affect the relative performance more significantly.

5.2 Simulation Results

We begin by considering the performance of an ad hoc network with CSMA channel access, no power control, steered beams, and omni-directional neighbor discovery. We then progressively change each of these parameters to discern the effect that each one has on the performance.

5.2.1 CSMA, No Power Control, Steered Beams, Omni-directional Neighbors

As illustrated in figure 3 there is no improvement in throughput when beamforming antennas are employed. This is because CSMA is even more unsuitable with beamforming – the “hidden terminal” problem is exacerbated due to reduced gains on all but the main lobe. Thus, a node often ends up transmitting to a neighbor that is busy transmitting or receiving. This seems to negate any gains due to slightly increased spatial reuse.

As illustrated in figure 4 the delay with beamforming antennas is a fair bit lower, especially at high densities. This is because, due to directional transmissions, nodes back off less and are able access the channel sooner. Although many of these packets collide, nodes nonetheless get their packets out quicker. The delay difference is a factor of about 17 in case of density 112, between omni-directional and the 20 dBi beam.

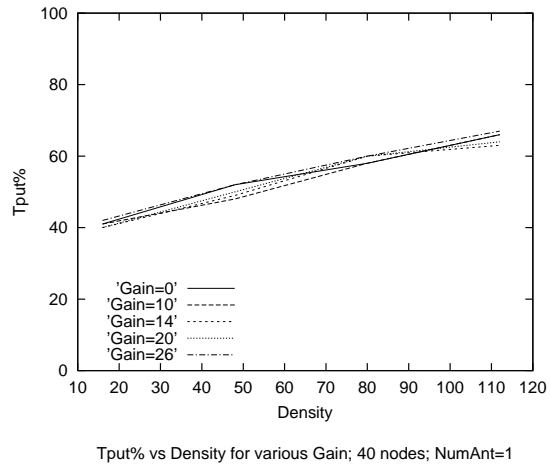


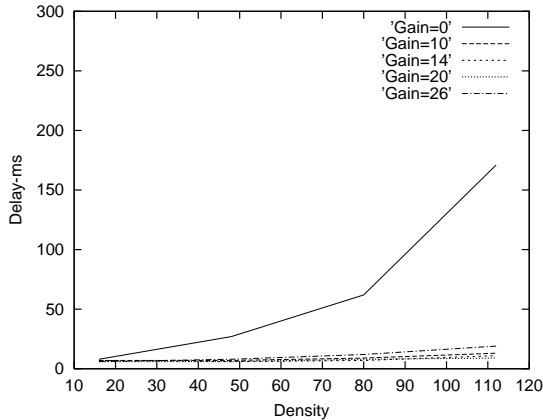
Figure 3: Throughput: Steered beams, CSMA, no power control

Thus, even with one of the simplest MAC protocol and a typical ad hoc networking algorithm, there are some gains in performance. However, they are probably not enough to justify deploying beamforming antennas, and hence motivate the enhancements considered below.

5.2.2 Adding Aggressive Collision Avoidance

We consider the performance with *aggressive collision avoidance* (refer section 4.2), no power control, steered beams, and omni-directional neighbors. The results are in figure 5, and 6.

When CSMA is augmented by an aggressive collision avoidance protocol, there is a marked difference between the throughput of directional and omni-directional antennas. The new MAC protocol ensures that a transmitter only sends a packet




Delay-ms vs Density for various Gain; 40 nodes; NumAnt=1

Figure 4: Delay: Steered beams, CSMA, no power control

when the receiver is not busy, thereby drastically increasing the number of successful transmissions. Throughput is increased by 15% (20dBi at 112 nodes/sq mile).


The delay of both directional and omni-directional antennas increases, when compared with using CSMA, due to the additional waiting for collision avoidance. Beamforming antennas have the same or less delay than omni-directional antennas, with the difference increasing at higher densities when omni-directional antennas start suffering from a greater decrease in spatial reuse.



Tput% vs Density for various Gain; 40 nodes; NumAnt=1

Figure 5: Throughput: Steered beams, Aggressive CA, No power control

We found that the Conservative Collision Avoidance described in section 4.2 had consistently less performance than the Aggressive Collision Avoidance, even for the omni-directional case. To understand this, recall that we used a finite queue size of 30 packets at each node. Packets arriving at a node find the queue full with a greater probability compared to Aggressive Collision Avoidance, and so more packets are dropped than would have collided. We noticed that increasing the buffer size does increase the throughput somewhat, but produces excessive delay. Therefore, we did not consider the



Delay-ms vs Density for various Gain; 40 nodes; NumAnt=1

Figure 6: Delay: Steered beams, Aggressive CA, No power control

Conservative Collision Avoidance further.

5.2.3 Adding Link Power Control

We consider the performance with aggressive collision avoidance, *link power control* (refer section 4.3), steered beams, and omni-directional neighbors. The results are in figure 7, and 8.

Adding power control to the steered, aggressive CA case improves the performance for both directional as well as omni-directional beams. However, the improvement is greater when beamforming is used. The throughput (see figure 7) with beamforming antennas is significantly higher than with omni-directional antennas. Interestingly, the greatest difference occurs at middle densities – e.g., at density of 48, using a 26 dBi antenna gives 28% better throughput than the omni directional antenna. This is a result of counteracting forces – the shorter number of hops, whose beneficial effects increase with increasing density, versus the detrimental effect of side-lobes, which also increases with increasing density, resulting in a peak at the middle densities. The difference between the various gains is much less - to within 6% in most cases, and about 10% in the density 16 case.

The delay (see figure 8) for steered beams with aggressive CA and power control is dramatically lower than with omni-directional antennas (also using aggressive CA and power control). When the density is 112 nodes/sq mile, there is a reduction by a factor of about 28 in the delay when the 26 dBi antenna is used. The difference is less at lower densities (about a factor of 2-5 at density 16). The delay of both omni-directional as well as beamforming antennas is higher at lower densities. This is due mainly to the increased number of hops at lower densities.

These results are a centerpiece of this paper for two reasons: first, they show that *dramatic* improvements are possible even under conservative assumptions on new protocols; and second, they show in light of the previous section that unless power control is employed, such improvements are likely to be elusive. The latter makes sense because without power control, the total amount of energy injected into the network

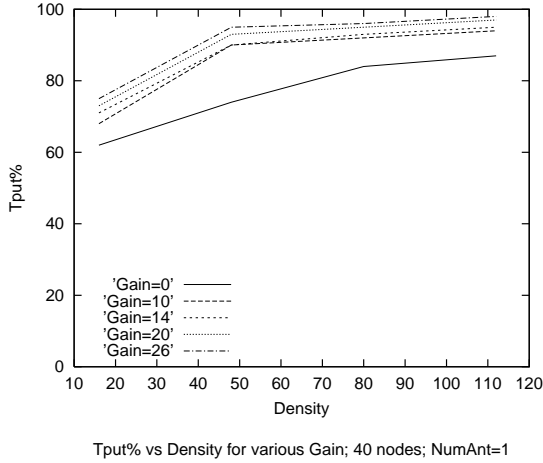


Figure 7: Throughput: Steered beams, Aggressive CA, Power Control

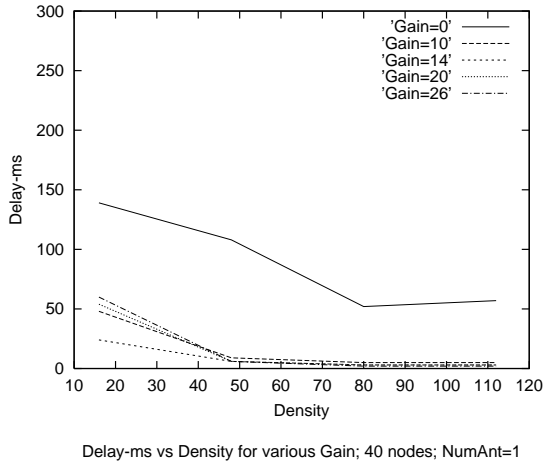


Figure 8: Delay: Steered beams, Aggressive CA, Power Control

in both omni-directional and directional communications is the same, causing roughly the same amount of interference.

5.2.4 Using Switched instead of Steered Beams

We consider the performance with aggressive collision avoidance, power control, *switched beams*, and omni-directional neighbors. The results are in figure 9, 10, and 11.

As seen in figure 9 and figure 10, the switched beam case behaves very similar to the steered case, both in terms of throughput and delay. An important contributor to this is the number of beams used – 16 in this case – which allows good coverage of the azimuthal plane. However, as seen in figure 11, performance at higher gain is reduced when only 8 beams are used. This is because the beamwidth of the 20 dBi antenna is only 20 degrees. The large amount of “gap” in the coverage (only 160 degrees out of 360, which is less than 50%) implies that some of the neighbors discovered omni-directionally are not reachable now, and negates any spatial reuse gains. In fact, we noticed in experiments not presented here (due to space limitations) that using 8 20 dBi beams

results in *less* throughput than the omni-directional case.




Figure 9: Throughput: Switched beams, Aggressive CA, Power Control




Figure 10: Delay: Switched beams, Aggressive CA, Power Control

Our results here indicate that one may consider switched beams as a less expensive alternative to fully adaptive beams, at least in terms of capacity gains from spatial reuse. However, a minimum number of beams, depending on the beamwidth, is crucial for good performance.

5.2.5 Using Directional Neighbor Discovery

We consider the performance with aggressive collision avoidance, power control, switched beams, and *directional neighbors*. Recall that with directional neighbor discovery, Hellos are sent out directionally on each beam and travel further, enabling a topology with longer-range links. We consider a lower density range for these experiments, as that is the more interesting range in this case. There are 12 beams per node. The results are in figure 13 and figure 14.

At low densities, the throughput with beamforming antennas is far higher. For density 8, using a 20 dBi switched beam antennas yields 118 % better throughput and a factor-of-20 reduction in delay. This is even more interesting when you



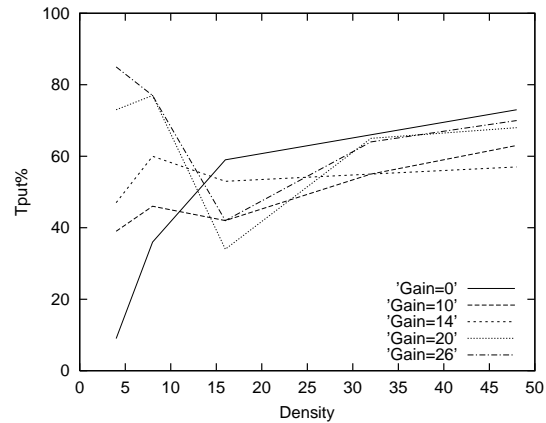


Figure 12: For very sparse deployments, use of omni-directional antennas leaves the network highly partitioned (left), whereas use of directional neighbor discovery with 10 dBi beams (middle) and 20 dBi beams (right) provides good connectivity and commensurate performance. We note that each link depicted is a directional link, and hence, unlike the omni-directional case, a high average degree is not bad.



Tput% vs Gain for various NumAnt; 40 nodes; Density=80

Figure 11: Throughput: Switched beams, Aggressive CA, Power Control



Tput% vs Density for various gains; 40 nodes; NumAnt=12

Figure 13: Throughput: Switched beams, Aggressive CA, Power Control, Directional Neighbors

consider that the beamwidth of the 20 dBi antenna is 20 degrees, and therefore the 12 switched beams do not cover the azimuthal plane completely.

For density 4 nodes/sq mile, the network is partitioned with omni-directional antennas but connected with beamforming antennas, as illustrated in figure 12.

The performance drops at middle densities before rising again. This reflects a playing-out of the interference versus range forces.

We also note that for higher densities, the throughput is about the same or worse with beamforming. This is due to the fact that directional neighbor discovery tends to use the longer links by virtue of the shortest path routing, which in turn causes more interference (recall the sidelobes). This bears out and extends to directional communications the conclusion in [1] that all things being equal, one should use the smallest power (shortest links) that provides a connected network. This motivates the use of novel topology control and routing algorithms that use shorter links even when longer directional links exist.

6. CONCLUDING REMARKS

Based on our simulation study of the performance of ad hoc networks with beamforming antennas, we arrive at the following conclusions.

1. Beamforming antennas have tremendous potential within ad hoc networks. For a typical 40 node stationary ad hoc network, and moderate-to-high load, we observe an improvement of up to 118% in throughput and a up to factor of 28 reduction in end-to-end delay.
2. Even with simple channel access techniques, the performance improvement is drastic. In the tradeoff spectrum between parallel transmissions and collisions, leaning toward more collisions but more parallel transmissions seems to pay off. This may be the appropriate operating point for real-time traffic.
3. Link power control is essential in exploiting the benefits of beamforming antennas to their fullest.
4. For ad hoc networks, when just spatial reuse is considered, switched beams are nearly as good as steered beams (and a lot less expensive).

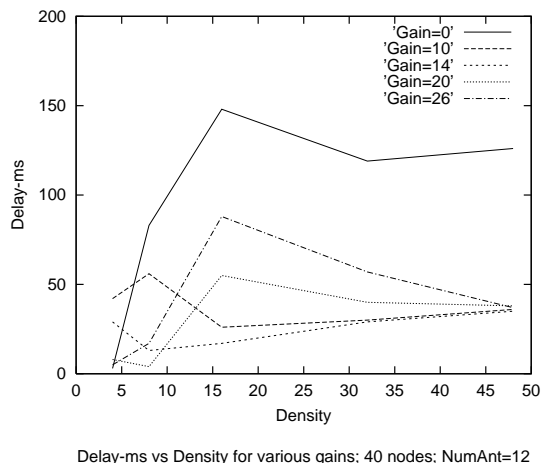


Figure 14: Delay: Switched beams, Aggressive CA, Power Control, Directional Neighbors

- The marginal utility of both steered and switched antennas decreases with increasing gain. Considering the cost and larger form-factor of higher gain antennas, this is good news.
- Directional neighbor discovery works wonders at low densities where it makes all the difference between a connected and a partitioned network.

Our preliminary results point to a dense mesh of low cost transceivers equipped with switched beam antennas as the best candidate for a wireless extension of gigabit networks. Needless to say, much work remains to be done to realize the true potential of beamforming antennas. Our work points to a number of exciting areas at the MAC and network layers that require research, including:

- Channel access protocols for steered and switched beams that allow for a tradeoff between spatial reuse and collisions (for real and non-real-time traffic).
- Techniques for exploiting the larger range of directional communications for better connectivity and lower latency. This in turn would require neighbor discovery using steered beams such as efficient “scanning” using transmit and/or receive beams, or use of higher processing gain, etc.
- New analytical and simulation models that do not assume that the the range is equal in all directions (no more “unit disks”).
- New techniques for characterizing directional links and their use in supporting quality-of-service.
- For slowly steerable antennas, topology control using beamforming in a manner similar to topology control using power [16].

From a practical viewpoint, there is a need for defining standardized interfaces between the link and the physical layer

for networking protocols to be able to use smart antennas. Finally, often-used simulation tools such as NS-2 need to be modified to incorporate antenna patterns and beam steering.

7. ACKNOWLEDGEMENTS

We thank Cesar Santivanez for his fixes to the OPNET pipeline stages which resulted in more accurate simulation. We also thank Dr. Nitin Vaidya for suggesting the example in figure 3(b), and Dr. Martha Steenstrup for her valuable comments.

8. REFERENCES

- P. Gupta and P.R. Kumar, “The Capacity of Wireless Networks”, *IEEE Transactions on Information Theory*, vol. IT-46, no. 2, pp. 388-404, March 2000.
- J. Zander, “Slotted ALOHA multihop packet radio networks with directional antennas,” *Electronic Letters*, vol. 26, no. 25, 1990.
- Y.B. Ko and N.H. Vaidya, “Medium Access Control Protocols Using Directional Antennas in Ad Hoc Networks,” *Proceedings of IEEE Infocom*, March 2000.
- A. Nasipuri, S. Ye, and R.E. Hiromoto, “A MAC Protocol for Mobile Ad Hoc Networks Using Directional Antennas,” *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2000.
- N. Pronios, “Performance considerations for slotted spread-spectrum random access networks with directional antennas,” in *Proc. of IEEE GLOBECOM*, November 1989.
- J. Ward and R. Compton, “High throughput slotted ALOHA packet radio networks with adaptive arrays,” *IEEE Transactions on Communications*, vol. 41, pp. 460-470, Mar 1993.
- Y-B. Ko and N.H. Vaidya, “Location-Aided Routing (LAR) in Mobile Ad Hoc Networks”, *ACM/Baltzer Wireless Networks (WINET)*, Vol. 6, No. 4, pp. 307-322.
- A. Nasipuri, J. Mandava, H. Manchala, and R.E. Hiromoto, “On-Demand Routing Using Directional Antennas in Mobile Ad Hoc Networks,” in *Proceedings of the IEEE International Conference on Computer Communication and Networks (ICCCN2000)*, October, 2000, Las Vegas.
- J.C. Liberti and T.S. Rappaport, “Smart Antennas for Wireless Communications,” Prentice-Hall PTR, 1999.
- J. Bulter and R. Lowe, “Beamforming Matrix Simplifies Design of Electronically Scanned Antennas,” *Electronic Design*, Apr 1961.
- V. Bhargavan, A. Demeers, S. Shenker, and L. Zhang, “MACAW - A media access protocol for wireless LANs,” in *Proc. of ACM SIGCOMM '94*, September 1994.
- “Wireless LAN medium access control (MAC) and physical layer (PHY) specifications,” 1997. Draft Standard IEEE 802.11, P802.11/D1: The editors of IEEE 802.11
- J. Monks, V. Bhargavan and W. Hwu, “Transmission Power Control for Multiple Access Wireless Packet Networks.” IEEE Conference on Local Computer Networks (LCN), Tampa, FL, November 2000
- P. Jacquet, P. Muhlethaler, and A. Quayyum, “Optimized link state routing protocol”, IETF MANET Working Group Internet-Draft, Work in Progress
- T. S. Rappaport, *Wireless Communications, Principles and Practice*, Prentice-Hall, 1996.
- R. Ramanathan and R. Hain, “Topology Control of Multihop Radio Networks using Transmit Power Adjustment,” in *Proceedings of IEEE Infocom*, Tel Aviv, Israel, 2000.