

The mistaken axioms of wireless-network research

David Kotz, Calvin Newport, and Chip Elliott

Dartmouth College Computer Science Technical Report TR2003-467

July 18, 2003

Abstract

Most research on ad-hoc wireless networks makes simplifying assumptions about radio propagation. The “Flat Earth” model of the world is surprisingly popular: all radios have circular range, have perfect coverage in that range, and travel on a two-dimensional plane. CMU’s ns-2 radio models are better but still fail to represent many aspects of realistic radio networks, including hills, obstacles, link asymmetries, and unpredictable fading. We briefly argue that key “axioms” of these types of propagation models lead to simulation results that do not adequately reflect real behavior of ad-hoc networks, and hence to network protocols that may not work well (or at all) in reality. We then present a set of 802.11 measurements that clearly demonstrate that these “axioms” are contrary to fact. The broad chasm between simulation and reality calls into question many of results from prior papers, and we summarize with a series of recommendations for researchers considering analytic or simulation models of wireless networks.

1 Motivation

Mobile ad-hoc networking (MANET) has become a lively field within the past few years. Since it is difficult to conduct experiments with real mobile computers and real wireless networks in the real world, nearly all published MANET articles are buttressed with graphs produced by simulation, and these simulations are based on common simplifying assumptions.

It will come as no surprise that many reviewers and readers of such articles treat these simulation results with less than full respect. Indeed a recent article in *IEEE Communications* warned [PJL02]: “An opinion is spreading that one cannot rely on the majority of the published results on performance evaluation studies of telecommunication networks based on stochastic simulation, since they lack credibility.” It then proceeded to survey over 2200

published network simulation results to point out systemic flaws; the results make for interesting if depressing reading.

Our goals in this paper are somewhat different, and are in fact three-fold: a) to point out how simplistic radio models may lead to manifestly wrong results in ad-hoc network simulation; b) to note that most ad-hoc network researchers make common simplifying assumptions about the radio model, and to quantitatively demonstrate that these assumptions are far from realistic; and c) to make a modest contribution towards ameliorating this problem by contributing a real dataset that should be easy to incorporate into simulations.

2 Radios in Theory and Practice

The upper-left example in Figure 1 provides an all-too-familiar model of radio propagation, as used in many simulations of ad-hoc networks. This simple model stands in stark contrast to the three representative signal-propagation maps, drawn at random from the web, and to measurements from an ad-hoc network of Berkeley Motes [GKW⁺02]. The simple theory is based on Cartesian distance in an X-Y plane. More realistic models take into account antenna height and orientation, terrain, foliage, surface reflection and absorption, and so forth.

Of course, not every simulation study needs to use the most detailed radio model available, nor explore every variation in the wide parameter space afforded by a complex model. The level of detail necessary for a given analytic or simulation study depends on the characteristics of the study. The majority of results published to date use the simple models, however, with no examination of the sensitivity of results to the (often implicit) assumptions embedded in the model.

Impact of these overly simple assumptions. Two illustrative dangers loom for protocol and system designers who rely on overly simple models of radio propagation. First, “typical” network connectivity graphs look quite different in reality than they do on a Cartesian grid.

This research was supported by Dartmouth’s Center for Mobile Computing.

Every radio engineer knows that you erect your antenna on top of the tallest nearby hill so it has direct connectivity with all other nearby radios, or in short, high “fan in.” This effect cannot be observed in simulations that represent only flat plains. Second, it is often difficult in reality to estimate whether or not one has a functioning radio link between nodes, because signals fluctuate greatly due to mobility and fading as well as interference. Broadcasts are particularly hard-hit by this phenomenon as they are not acknowledged in typical radio systems. Protocols that rely on broadcasts (e.g., of beacons) or “snooping” may therefore work significantly worse in reality than they do in simulation. The following paragraphs expand on these remarks.


Figure 2 depicts one immediate drawback to the oversimplified model of radio propagation. The Cartesian (“Flat Earth”) approach links all network nodes as if they were on a flat plain. Contrast this with a simple three-dimensional model that includes some altitude, even a single hill, and note that the resulting network graph looks completely different from that of the Flat Earth model. Or consider a different simple model that includes obstacles (such as buildings or walls). Even if the obstacles are considered to be entirely absorptive (without reflections) the resultant connectivity graph again looks completely different from the Flat Earth model.

Now imagine all the nodes moving in these three scenarios. The ways in which connectivity changes depending on node locations— that is, the changes in graph edges over time— will be different in each scenario.


Figure 3 presents a further level of detail. At the top, we see a node’s trajectory past the theoretical (T) and practical (P) radio range of another node. Beneath, we sketch the kind of change in link quality we might expect under these two models. The Flat Earth (T) model gives a simple step function in connectivity: either one is connected or one is not. Given a long enough straight segment in a trajectory, this leads to quite a low rate of change in link connectivity. And such a model makes it easy to determine when two nodes are, or are not, “neighbors” in the ad-hoc routing sense.

In more realistic model (P) the quality of the link is likely to vary rapidly and unpredictably, even when two radios are nominally “in range.” In these more realistic cases, it is by no means easy to determine when two nodes have become neighbors, or when a link between two nodes is no longer usable and should be torn down. In the figure, suppose that a link quality of 50% or better is sufficient to consider the nodes to be neighbors. In the diagram, the practical model would lead to the nodes being neighbors, briefly, then dropping the link, then being neighbors again, then dropping the link.


In addition to spatial variations in signal quality, a radio’s signal quality varies over time, even for a station-




Typical theoretical model



Source: Comgate Engineering
<http://www.comgate.com/ntdsign/wireless.html>



Source: University of Stuttgart
http://www.ihf.uni-stuttgart.de/Winprop/Download/Poster_Models.pdf



Source: Midwest Radio Association
http://www.2meters.org/mra_cp.html

Figure 1: Real radios are more complex than the common theoretical model at the top. Here, color represents signal quality.




Figure 2: The Flat Earth model is overly simplistic.

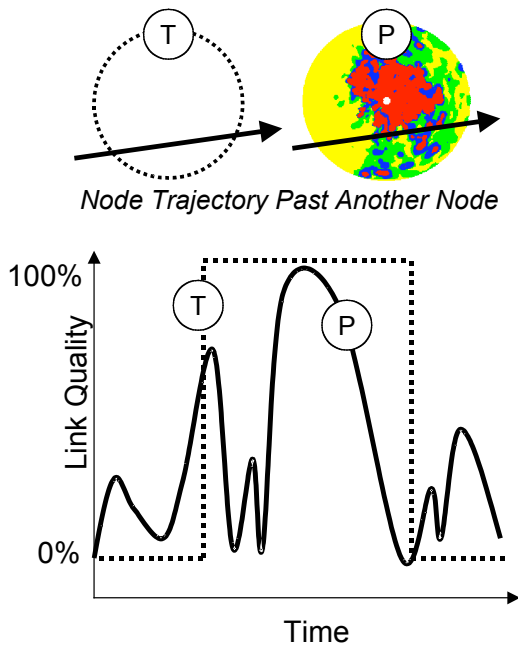


Figure 3: Difference between theory (T) and practice (P).

ary radio and receiver. Obstacles come and go: people and vehicles move about, leaves flutter, doors shut. Both short-term and long-term changes are common in reality, but ignored by most practical models. Some, but not all, of this variation can be masked by the physical or data-link layer of the network interface. Link connectivity can come and go; one packet may reach a neighbor successfully, and the next packet can fail.

Although the simple theoretical model may be easy to use when simulating ad-hoc networks, it leads to an incorrect sense of the way the network evolves over time. For example, in Figure 3 the link quality varies much more rapidly in practice than in theory, so the link connectivity may vary much more rapidly with a realistic model than with a simplistic model. Many algorithms and protocols may perform much more poorly under such dynamic conditions. In some, particularly if network connectivity changes rapidly with respect to the distributed progress of network-layer or application-layer protocols, the algorithm may fail due to race conditions or a failure to converge. Simple radio models fail to explore these critical realities that can dramatically affect performance and correctness. For example, Ganesan et al. measured a dense ad-hoc network of sensor nodes and found that small differences in the radios, in propagation distances, and the timing of collisions can significantly alter the behavior of even the simplest flood-oriented network protocols [GKW⁺02].

In summary, **“good enough” radio models are likely to be quite important in simulation of ad-hoc networks. The Flat Earth model, however, is by no means good enough.** In the following sections we make this argument more precise.

3 Is there a problem?

Yes, our community has a problem. To get some idea of the extent of the problem, we surveyed a nearby set of MobiCom proceedings from 1995 through 2002.¹ We inspected the simulation sections of every article in which RF modeling issues seemed relevant, and categorized the approach into one of three bins: *Flat Earth*, *Simple*, and *Good*. This categorization required a fair amount of value judgment on our part, and we omitted cases in which we could not determine these basic facts about the simulation runs.

Figure 4 presents the number of papers that fall into each category, year by year. Note that in good years (1999–2001) we count one Good simulation result per year. In most years the Flat Earth and Simple models run pretty much even. Two papers [JLW⁺96, TMB01] deserve commendation for their thoughtful channel models.

¹We used a full set of MobiCom with one volume missing (1997).




Figure 4: The number of papers in each year of Mobicom that fall into each category.

Flat Earth models are based on Cartesian X-Y proximity, that is, nodes A and B communicate if and only if node A is within some distance of node B .

Simple models are, almost without exception, ns-2 models using the CMU 802.11 radio model [FV02]. This model provides what has sometimes been termed a “realistic” radio propagation model. Indeed it is significantly more realistic than the “Flat Earth” model, e.g., because it models packet delay and loss caused by interference rather than assuming that all transmissions in range are received perfectly. We still call it a “simple” model, however, because it embodies many of the questionable axioms we detail below. In particular, the standard release of ns-2 provides a simple free-space model ($1/r^2$), which has often been termed a “Friss-free-space” model in the literature, and a *two-ray ground-reflection model*. Both are described in the ns-2 document package [FV02, Chapter 18].

The free-space model is similar to the “Flat Earth” model described above, as it does not include effects of terrain, obstacles, or fading. It does, however, model signal strength with somewhat finer detail than just “present” or “absent.”

The two-ray ground-reflection model, which considers both the direct and ground-reflected propagation path between transmitter and receiver, is better but not particularly well suited to most MANET simulations. It has been reasonably accurate for predicting large-scale signal strength over distances of several kilometers for cellular telephony systems using tall towers (heights above 50m), and also for line-of-sight micro-cell channels in urban environments. Neither is characteristic of typical MANET scenarios. In addition, while this propagation model does take into account antenna heights of the two

nodes, it assumes that the earth is flat (and there are otherwise no obstructions) between the nodes. This may be a plausible simplification when modeling cell towers, but not when modeling vehicular or handheld nodes because these are often surrounded by obstructions. Thus it too is a “Flat Earth” model, even more so if the modeler does not explicitly choose differing antenna heights as a node moves.²

More recently, Wei Ye of ISI added a third channel model to ns-2, called the “shadowing” model, which can account for indoor obstructions and outdoor shadowing via a probabilistic model [FV02, Chapter 18]. Although it does not appear to take antenna height or topography into account, it may provide more realistic propagation models than the older free-space or two-ray ground reflection models. To our knowledge, no MANET simulations to date have reported results using this shadowing model.

Good models have fairly plausible RF propagation treatment. In general, these models are used in papers coming from the cellular telephone community, and concentrate on the exact mechanics of RF propagation. To give a flavor of these “good” models, witness this quote from one such paper [ER00]:

In our simulations, we use a model for the path loss in the channel developed by Erceg et al. This model was developed based on extensive experimental data collected in a large number of existing macro-cells in several suburban areas in New Jersey and around Seattle, Chicago, Atlanta, and Dallas. . . [Equation follows with parameters for antenna location in 3-D, wavelength, and six experimentally determined parameters based on terrain and foliage types.] . . . In the results presented in this section, . . . the terrain was assumed to be either hilly with light tree density or flat with moderate-to-heavy tree density. [Detailed parameter values follow.]

Of course, the details of RF propagation are not always essential in good network simulations; most critical is the overall realism of connectivity and changes in connectivity (Are there hills? Are there walls?). Along these lines, we particularly liked the simulations of well-known routing algorithms presented by Johansson et al. [JLH+99], which used relatively detailed, realistic scenarios for a conference room, event coverage, and disaster area. Although this paper employed the ns-2 802.11 radio model, it was rounded out with realistic network obstacles and node mobility.

²See also [Lun02], Sections 4.3.4–5, for additional remarks on the two-ray model’s lack of realism.

4 Common MANET axioms

For the sake of clarity, let us be explicit about some basic “axioms” upon which most MANET research explicitly or implicitly relies. These axioms deeply shape how network protocols behave. We note that all of these axioms are contradicted by the actual measurements reported in the next section.

0: The world is flat.

1: A radio’s transmission area is circular.

2: All radios have equal range.

3: If I can hear you, you can hear me (symmetry).

4: If I can hear you at all, I can hear you perfectly.

5: Signal strength is a simple function of distance.

This last Axiom is not used in many MANET papers, because the prior axioms allow the protocol or algorithm to assume a simple model of connectivity and to ignore signal strength. We include Axiom 5 because it is often a core assumption of algorithms that use signal strength to estimate distance, e.g., to obtain radio position by triangulation.

5 The Reality

Unfortunately, real wireless network devices are not nearly as simple as those considered by the axioms in the preceding sections. Since we did not have at hand a large collection of devices in an ad-hoc wireless network, we set out to measure the characteristics of the radios in a production Wi-Fi network on the campus of Dartmouth College, which has a campus-wide network of over five hundred 802.11b access points. The full campus network is shown in Figure 5, but Figures 6 and 7 show the two regions of campus where we took measurements. Most access points are Cisco model 350, with a small number that are model 340.

Although the Dartmouth access points comprise a static infrastructure network rather than a mobile ad-hoc network, for the purpose of this study we treat them as a static set of wireless network radios. In the next subsection, we describe how we used a single mobile measurement device to consistently record the signal characteristics of the network and treat each access point as a “node” in an ad-hoc network.

5.1 Data-collection methods

We constructed a map and collected three data sets.

We obtained a detailed scale map of the campus from Dartmouth College as a AutoCAD file. The map shows the location of streets and the footprint (outline) of each building. We further obtained detailed scale floorplan drawings of each building, on which were marked the location of each access point. We scaled and rotated each

floor plan to fit inside the footprint of each building. The result is a map that places every access point on the map’s coordinate system, so we can compute relative distances with ease; although we do not know the conversion factor between map units and meters, the map units are sufficient for our purpose. Figures 6 and 7 are subsets of that map.

First data set: node radio coverage. We carried a palm-sized computer with Wi-Fi and GPS capability.³ We again used NetStumbler (Ministumbler version 0.3.23 (beta)) to collect the data, as we walked around outdoors near an isolated access point.

We then post-processed the data to account for noise in our measurements of GPS location, and small space and time variations in signal quality. GPS and NetStumbler report each observation to the nearest 0.0000001 degree of latitude and longitude. We rounded the location of each observation to the nearest 1/8000 degree (0.0001250 degree), and then computed the average value of all observations at the same rounded location. The result smoothes the fine variations of location and quality, and makes our maps easier to read by avoiding overlapping data points.

Second data set: node-to-node measurements. We chose two reasonably self-contained sections of campus, the engineering school (Figure 6) and the medical school (Figure 7). We used a tablet computer⁴ equipped with a wireless 802.11b card⁵ to run the NetStumbler software (version 0.3.30). We carried the wireless tablet to the location of 92 unique Cisco access points⁶ distributed in 19 buildings. At each access point, we aligned our wireless card at a consistent (waist level) height and at a consistent angle to the antenna (we held the card parallel to the broad side of the paddle antenna) and recorded 10 seconds of signal strength data using the NetStumbler software. For access points not visible to us (e.g., in locked closets) we determined the proper alignment by choosing the angle that provided the maximum signal-strength reading for the local access point, before recording the data. After collecting all the data, we computed for each access point the average signal strength for each of the other access points (or zero for those that were not heard at that location). The average strength was determined by averaging all readings taken over that 10-second recording in-

³An iPAQ 3870 with a single PCMCIA card expansion pack, loaded with a Lucent Orinoco Gold card and a 5 dBi antenna mounted on a 5 foot pole. The antenna had 10 feet of low-loss cable, a type N connector, and a pigtail to plug into the Lucent card. The iPAQ serial port connected to a Garmin eTrex Vista GPS, which we carried in its cradle during the data collection process.

⁴A Fujitsu C-500, running Windows 2000, Service Pack 1.

⁵Dell TrueMobile 1150 running version 1.6.22.10 of the bundled Dell driver, and firmware version 4.04. No external antenna.

⁶All were a model 350 access point operating at power level 100. One used a dipole antenna; all others used a Cisco paddle-style antenna.



Figure 5: A map of all access points on the Dartmouth College campus. Readers using Acrobat Reader can zoom and pan to see full detail. Note, however, that this two-dimensional map does not clearly show all APs, because some multi-story buildings place APs at the same location on multiple floors.

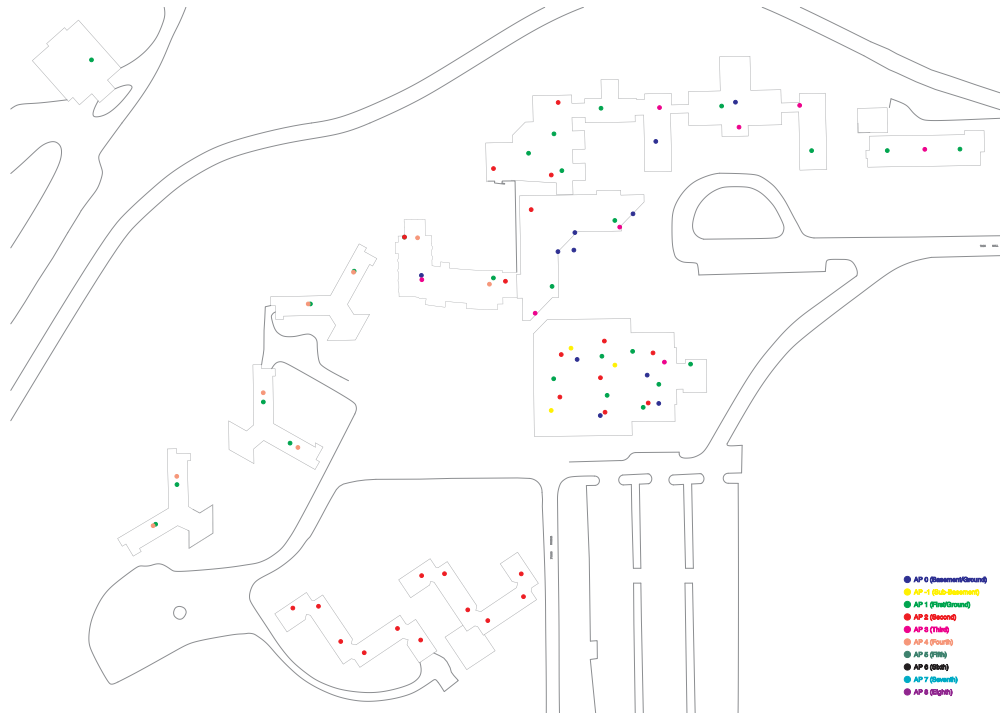


Figure 6: A map of access points in and around the engineering school of Dartmouth College. Readers using Acrobat Reader can zoom and pan to see full detail. Note, however, that this two-dimensional map does not clearly show all APs, because some multi-story buildings place APs at the same location on multiple floors.

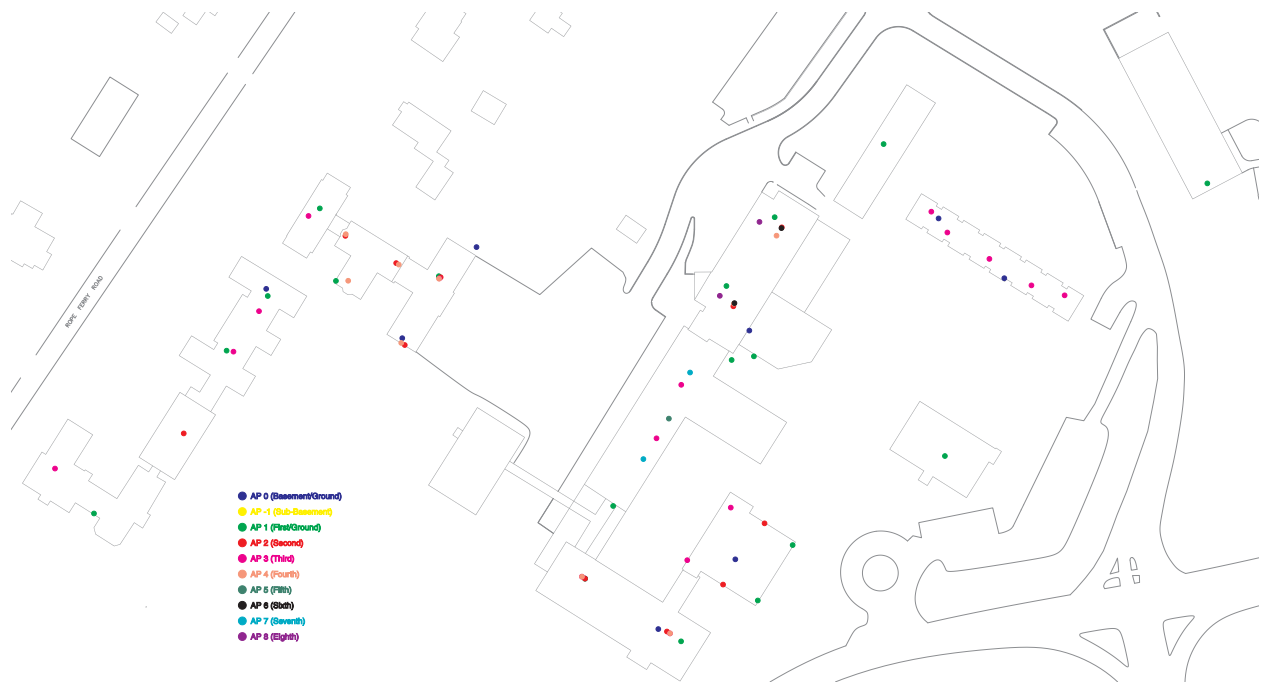


Figure 7: A map of access points in and around the medical school of Dartmouth College. Readers using Acrobat Reader can zoom and pan to see full detail. Note, however, that this two-dimensional map does not clearly show all APs, because some multi-story buildings place APs at the same location on multiple floors.

terval. To ensure that the resulting network is completely self-contained, we removed reference to any access points from whose location we did not record data. Because we used two reasonably self-contained sections of campus, there were few such fringe readings. The resulting data provides the edges of a directed node-connectivity graph, with a weight (average signal strength) on each edge.

Using the map, we know the X,Y coordinates for each access point, along with the floor number. We compute distance in this “two and a half” dimensional space by converting the floor number to an elevation, assuming that each floor is 10 feet below the next floor, using an empirically derived constant to convert feet into map units. Then we simply use the three-dimensional Cartesian distance formula.

Third data set. For Axiom 4, we wanted to measure how well one radio can hear another radio, and vice versa. We used our extensive network of access points and our large collection of mobile users as a source of information. We used SNMP to poll 149 of our access points for a day, collecting information about the number of frames sent and received by each client, and the number of frames that had errors. We selected 15 mobile users at random from those that moved about extensively during the trace; these 15 users were active in 5 different locations.

We now use this data to explore each of the axioms.

5.2 Axiom 0

The world is flat.

Clearly, the Earth is not Flat. The core area of our campus is mostly flat, although there are a few outlying nodes down by the river, about one hundred feet below the main campus. Even in a world that is nearly flat, like our campus, note that wireless nodes are often used in multi-story buildings. Figure 8 demonstrates the wide range of node elevations, expressed simply as the floor number on which the access point is mounted (using American-style numbering, floor zero is the basement, and floor one is the ground floor). In many tall buildings, two nodes may be found at exactly the same X,Y location, but on different floors. Any Flat Earth model would assume that they are in the same location, and yet they are not. In some tall buildings, we found it was impossible for a node on the fourth floor to hear a node in the basement, at the same X,Y location.

As we noted above, the common $ns-2$ model essentially assumes a flat earth.

A local researcher using Berkeley “motes” for sensor-network research notes the critical impact of elevation and ground-reflection effects:

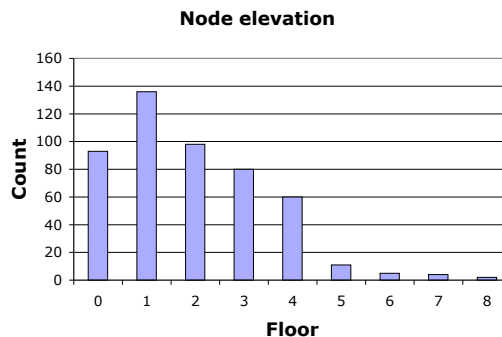


Figure 8: The distribution of nodes across building floors in campus buildings.

In our current experiments we just bought 60 plastic flower pots to raise the motes off the ground because we found that putting the motes on the ground drastically reduces their transmit range (though not the receive range). Raising them a few inches makes a big difference.

5.3 Axiom 1

A radio’s transmission area is circular.

The radio maps of Figure 1, calculated by other researchers with a variety of propagation modeling tools, make it clear that the signal coverage area of a radio is far from simple. Not only is it not circular, nor convex; the coverage region is often non-contiguous.

We used our first data set, collected outdoors using GPS and a tall antenna, to measure the signal strength (SS) of nodes (access points) in our campus network. If the coverage area of a radio were circular, we would see consistently high SS readings within the range of the radio, and low SS outside that range. In Figure 9 we see an aerial photograph of a portion of our campus, with measurements of a single node superimposed.⁷ Concentric circles at 100m intervals provide a guide to discover the “range” of this radio. The black dots and white, gray, and black circles indicate locations where we measured the strength of the signal from this radio. Although we could not measure every location in this mixed terrain of playing fields,

⁷This access point was a ruggedized Cisco model 350 series bridge, attached to a 5.2 dBi Omnidirectional antenna (AIR-ANT2506) mounted at the peak of the roof. The antenna was connected through 20 feet of Cisco Low-Loss cable (AIR-CAB020LL-R), which loses 3.5 dB in the run. There was also a Lightning Arrestor (AIR-ACC3354) at every installation. The AP operated at 100 mW.

low buildings, and forested hillside, it is clear that the SS varies from nonexistent, to poor, to excellent even within the three innermost circles. For example, consider the second innermost circle, representing range from 100–200 meters. This region contains points of every color, from no service (black dots), through white, gray, and black circles. Indeed, so does the next region, 200–300 meters. We have similar maps for other access points on campus.

Although we do not have sufficient data to draw the coverage map of this access point, it is clear that it is not circular.

Ganesan et al. used a network of Berkeley “motes” to measure signal strength of a mote’s radio throughout a mesh of mote nodes [GKW+02]. [The Berkeley mote is currently the most common research platform for real experiments with ad-hoc sensor networks.] The resulting contour map is not circular, nor convex, nor even monotonically decreasing with distance.

5.4 Axiom 2

All radios have equal range.

Since the coverage area of a radio is not circular, as we showed when dispelling Axiom 1, it is difficult to even define the “range” of a radio. Nonetheless, our data makes it clear that the radios have differing ranges: some close pairs of radios could not hear each other, while some distant pairs could.

In their study, Ganesan et al. [GKW+02] defined *connectivity radius* (range) as the distance beyond which the probability of successfully receiving a radio’s transmissions drops below a given threshold (65%). They computed this probability-vs-distance relationship over all node pairs in a dense, rectangular grid of nodes. In this probabilistic approach they do not distinguish individual node ranges. Without attempting to compute the range of a node, we can use our data to show that each node has a different range.

We used our second data set, in which we measure the indoor signal quality of each node from the location of each other node. We define $SS(i, j)$ as the signal strength of node i observed at the location of node j (which is zero when i is not heard at j). Note that $SS(j, i)$ is typically different than $SS(i, j)$ (see Axiom 3).

Since we also know the position of each node, we can easily compute the distance between two nodes. If each node’s radio has the same circular range, per Axioms 1 and 2, then all nodes “in range” of a given node would hear that node, and no nodes “out of range” would hear that node.

In Figure 10 we consider every node pair (i, j) (for $i \neq j$), and divide them into 10-unit distance buckets. Notice that (i, j) and (j, i) are considered two node pairs




Figure 10: A histogram that depicts the number of node pairs (i, j) in a given 10-unit distance bucket in which j can hear i ($SS(i, j) > 0$). The left-most bar represents nodes 10–20 units apart.

in this computation, albeit at the same distance. Because the nodes are not uniformly distributed across our campus, the inter-node distances are not uniformly distributed; the number of node pairs in each bucket varies. We therefore plot the *fraction* of node pairs (i, j) in each distance bucket in which j can hear i (that is, $SS(i, j)$ was non-zero). If the axiom were true, Figure 10 would be level at value 1.0 out to the range of the radio, and then level at 0.0 thereafter.

Our data shows that the radios had different ranges. Our nodes cannot reliably hear each other unless they are extremely close together, because there are some node pairs 20–30 units apart with zero signal strength. On the other hand, some distant nodes can hear each other, because there are some node pairs quite far apart that can hear each other. The curve in Figure 10 falls off gradually, indicating that the each node has a different range.

5.5 Axiom 3

If I can hear you, you can hear me (symmetry).

Clearly, not all node pairs can hear each other. Even in a symmetric relationship, where i can hear j and j can hear i , the amount of symmetry can vary widely. We define the *signal-strength symmetry* (SSS) of that pair to be

$$SSS(i, j) = \min[SS(i, j)/SS(j, i), SS(j, i)/SS(i, j)]$$

except where both $SS(i, j) = 0$ and $SS(j, i) = 0$, in which case $SSS(i, j)$ is undefined. The $\min()$ forces SSS to the range [0:1] and to zero when one of the nodes can-




Figure 9: An aerial photograph of one corner of our campus, with an access point node in the center. The photo is aligned with the top to the North. White circles indicate SS from 0 to 25, light gray circles indicate SS from 26 to 50, dark gray circles indicate SS from 51 to 75, and black circles indicate SS over 75. Black dots indicate places where we could not hear the node.




Figure 11: A histogram of signal-strength symmetry (SSS) for all pairs (i, j) where $i < j$, where defined.




Figure 12: The cumulative distribution function (CDF) of signal-strength symmetry for all pairs (i, j) where $i < j$, where defined.

not hear the other.⁸ Figure 11 makes it clear that there are about as many wholly asymmetric relationships (SSS=0) as there are perfectly symmetric relationships (SSS=1), and a wide range of asymmetry in between. Indeed, we were surprised by the large number of wholly asymmetric relationships on our campus. Figure 12 demonstrates the distribution of symmetry values as a CDF.

Ganesan et al. [GKW⁺02] noted that about 5–15% of the links in their ad-hoc sensor network were asymmetric. In that paper, an asymmetric link had a “good” link in one direction (with high probability of message reception) and a “bad” link in the other direction (with a low probability of message reception). [They do not have a name for a link with a “mediocre” link in either direction.] Although we measure signal strength rather than message reception probability, and we present the degree of symmetry rather than a thresholded definition of symmetry, the conclusion is the same: the two directions of a link can be very different.

Nonetheless, many researchers assume this axiom is true, and thus all network links are bidirectional. Some acknowledge that real links may be unidirectional (i cannot hear j even though j can hear i), and usually discard those links so that the resulting network has only bidirectional links. Nonetheless, in most such cases the protocol does not prevent the use of the unidirectional links, that is, the protocol does not reliably detect and discard asymmetric links. For example, if i sends a packet to j and j receives it, j typically will use it without testing whether a return packet from j to i would have arrived. In a network with mobile nodes, or a dynamic environment, link quality can vary frequently and rapidly, so a bidirectional

⁸In an asymmetric relationship, either $SS(i, j) = 0$ or $SS(j, i) = 0$, and $SSS = \min(\infty, 0) = 0$.

link may become unidirectional at any time. It is best to develop protocols that do not assume symmetry.

5.6 Axiom 4

If I can hear you at all, I can hear you perfectly.

We used our third data set to examine the reliability of frame transmission. The data demonstrates that, although the mobile clients are within range of the access point, a non-zero number of frames are lost due to transmission errors (collisions or noise). Six of the clients had no errors, two had errors in frames sent from the mobile client to the AP, and eight had a small number of errors in frames sent from the AP to the mobile client. Overall, there was an average 0.44% error rate in frames sent by the client, and 0.56% error rate in frames sent by the AP to the client. In addition, there was one notable outlier client, in which 36% of all frames sent by the AP were in error.

In short, the frame error rate in our well-provisioned network is small, but in some cases decidedly non-zero.

The common ns-2 model assumes that frame transmission, within the range of a radio, is perfect. Although it provides hooks to add a bit-error-rate (BER) model, these hooks are unused. More sophisticated models do exist, particularly those developed by Qualnet and the GloMoSim project⁹. We particularly commend their efforts to show how their sophisticated channel models affect the outcome of network simulations.

5.7 Axiom 5

Signal strength is a simple function of distance.

⁹<http://www.scalable-networks.com/pdf/mobihocpres0.pdf>




Figure 13: A scatter plot of the signal strength readings at various distances, showing only weak correlation.

Rappaport [Rap96] notes that the signal strength should fade with distance according to a power-law model. In this section we show the results of our measurements, which indicate only a weak correlation.

Recall that our second data set provides an observation at the location of each node, giving the average signal strength of every other node that can be heard at that location. We use the base map to compute the three-dimensional distance between the node and the observation point, and plot the relationship between signal strength and distance in Figure 13. The signal-strength units are those reported by NetStumbler, and the distance units are those used by the X,Y coordinate system of our base map.¹⁰

There is no clearly visible correlation in Figure 13. We used the SPSS statistical modeling package to fit each of the common distribution functions to this data; Figure 14 shows the resulting fits. The R-squared value for these fits is never better than 0.260, which implies a poor fit (0 is no correlation, 1 is perfect correlation). Some modelers assume that signal strength drops as the square or cube of the distance, which makes physical sense. Our quadratic and cubic fits each had R-squared of only 0.187, however, implying that quadratic or cubic functions are poor models of signal strength over distance.

More generally, our power-law fit attempts to fit the data to the general equation $SS = ad^b$ where d is the distance and a and b are constants. The best fit, with R-squared 0.258, is a weak correlation at best:

$$SS = 129.350d^{(-0.1403)}$$

Note that 0.14 is far from the exponents commonly used

¹⁰Neither are physically meaningful units, but both are “to scale” and sufficient for our purpose.




Figure 14: Curves fit to the preceding scatter plot.

in models (typically, in the range 2–5).

The S model gives the best fit (R-squared of 0.260), but we doubt this model is particularly meaningful.

$$SS = exp(2.1916 + 138.950/d)$$

The reason for the poor fits is clear: our environment is full of obstacles that attenuate or reflect the signals. An empty-space, noise-free environment is simply not real.

6 Summary and recommendations

Over the past seven years, dozens of Mobicom papers have presented simulation results for mobile ad-hoc networks. The great majority of these papers rely on overly simplistic assumptions of how radios work. Both widely used radio models— “flat earth” and ns-2 “802.11” models— embody the following set of axioms: the world is two dimensional; a radio’s transmission area is roughly circular; all radios have equal range; if I can hear you, you can hear me; if I can hear you at all, I can hear you perfectly; and signal strength is a simple function of distance.

In this paper, we present a real-world dataset that strongly contradicts all these “axioms.” It thus casts doubt on published simulation results that may implicitly rely on these assumptions, e.g., by assuming how well broadcasts are received, or whether “hello” propagation is symmetric.

We have the following recommendations for the MANET research community.

1. Always state your assumptions explicitly. Wherever possible, avoid assuming these axioms are true.

2. All simulations should run in three dimensions, e.g., on terrain with moderate hills and valleys, with corresponding radio propagation. It would be helpful if the community agreed on a few standard terrains for comparison purposes.
3. All simulations should include some fraction of asymmetric links (e.g., where A can hear B but not vice versa) and some time-varying fluctuations in whether A 's packets can be received by B or not. Here the ns-2 “shadowing” model may prove a good starting point.
4. In the meantime, use real data (such as our dataset) as input to simulators. Using our data as a static “snapshot” of a realistic ad-hoc wireless network with significant link asymmetries, packet loss, elevated nodes with high fan-in, and so forth. Researchers should verify whether their protocols form networks as expected, even in the absence of mobility.

6.1 Data availability

We will make our data available to interested parties. Our first data set, useful for simulation of static networks, includes:

Map. The $2\frac{1}{2}$ -dimensional location of every access point on campus (X, Y location plus floor number).

Signal strength. For two regions of campus, a set of observations from every access point in that region noting the strength of each other access point “audible” at that location.

Graph. The result is a static network graph in which access points are nodes, and directed edges indicate node pairs with non-zero signal strength. Each edge has two weights: signal strength and physical distance.

Our second data set, useful for more detailed analysis of network coverage, and possibly for limited simulations involving mobile nodes, includes

Map. An aerial photo of the area, registered to latitude and longitude.

Observation points. Numerous observation points, each tagged with GPS position (latitude and longitude), recorded continuously while driving or walking in the area of the map.

Signal measurements. Signal, noise, and signal-to-noise ratio at each observation point, for each and any access points “audible” at that point at that time.

A recommendation for Wi-Fi manufacturers. In the course of our data-collection efforts we also arrived at a recommendation for the manufacturers of Wi-Fi access points:

- Most access points make a wealth of information available through SNMP, including the list of connected clients and the quality of their signal. Presumably each access point can also “hear” other nearby access points. If each access point records the signal quality and other observations of nearby access points, then SNMP-based tools can be used to map and monitor the signal quality of the network. Such a tool would be invaluable for monitoring and maintaining an infrastructure network. It would also be beneficial to research projects like ours.

6.2 Contributions

Others have noted that real radios and ad-hoc networks are much more complex than the simple models used by most researchers [PJL02], and that these complexities have a significant impact on the behavior of MANET protocols and algorithms [GKW⁺02]. In this paper, we enumerate the set of common assumptions used in MANET research, provide data demonstrating that these assumptions are not usually correct, and recommend critical actions for our community. Our data can be used in network simulations, as one example of a real-world network environment, and may be helpful in the development of new, more realistic radio models.

6.3 Future work

We plan to collect another set of data from a set of 50 mobile nodes in an active ad-hoc network. We will continuously record GPS position, signal measurements, and protocol-level information from a variety of common ad-hoc routing protocols.

We would like to see research exploring this issue further, in particular, examining the effect of detail in the radio model on the behavior and performance of ad-hoc routing algorithms.

Acknowledgements

We are extremely grateful to the many people that helped make this project possible.

Jim Baker supplied a floorplan for every building with AP locations marked. Gurcharan Khanna, James Pike, and the FO&M department supplied the campus base map. Erik Curtis, a Dartmouth undergraduate, painstakingly mapped each floorplan to the campus base map.

Zach Berke and Chris Lentz, Dartmouth students, and Jim Christy, a Cisco engineer, helped us understand the capabilities of our Cisco AP 350 access points. Chris, in particular, stomped around campus with snowshoes to collect signal-strength readings that support our analysis of axiom 1.

Qun Li, Jason Liu, Ron Peterson, and Felipe Perrone all provided invaluable feedback on early versions of this paper.

Dr. Jason Redi loaned us his collection of Mobicom proceedings.

This project was supported in part by a grant from the Cisco Systems University Research Program, and by the Dartmouth Center for Mobile Computing.

References

- [ER00] Moncef Elaoud and Parameswaran Ramanathan. Adaptive allocation of CDMA resources for network-level QoS assurances. In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking*, pages 191–199. ACM Press, 2000.
- [FV02] Kevin Fall and Kannan Varadhan. *The ns Manual*, April 14 2002. www.isi.edu/nsnam/ns/ns-documentation.html.
- [GKW⁺02] Deepak Ganesan, Bhaskar Krishnamachari, Alec Woo, David Culler, Deborah Estrin, and Stephen Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical Report UCLA/CSD-TR 02-0013, UCLA Computer Science, 2002.
- [JLH⁺99] Per Johansson, Tony Larsson, Nicklas Hedman, Bartosz Mielczarek, and Mikael Degermark. Scenario-based performance analysis of routing protocols for mobile ad-hoc networks. In *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networking*, pages 195–205. ACM Press, 1999.
- [JLW⁺96] Jan Jannink, Derek Lam, Jennifer Widom, Donald C. Cox, and Narayanan Shrivakumar. Efficient and flexible location management techniques for wireless communication systems. In *Proceedings of the Second Annual International Conference on Mobile Computing and Networking*, pages 38–49. ACM Press, 1996.
- [Lun02] David Lundberg. Ad hoc protocol evaluation and experiences of real world ad hoc networking. Master’s thesis, Department of Information Technology, Uppsala University, Sweden, 2002.
- [PJL02] K. Pawlikowski, H.-D.J Jeong, and J.-S.R. Lee. On credibility of simulation studies of telecommunication networks. *IEEE Communications*, 40(1):132–139, January 2002.
- [Rap96] T. S. Rappaport. *Wireless Communications, Principles and Practice*. Prentice Hall, Upper Saddle River, New Jersey, 1996.
- [TMB01] Mineo Takai, Jay Martin, and Rajive Bagrodia. Effects of wireless physical layer modeling in mobile ad hoc networks. In *Proceedings of the 2001 ACM International Symposium on Mobile Ad-Hoc Networking and Computing (MobiHoc 2001)*, pages 87–94. ACM Press, 2001.