Experimental evaluation of wireless simulation assumptions

David Kotz, Calvin Newport, Robert S. Gray, Jason Liu, Yougu Yuan, and Chip Elliott

Abstract

All analytical and simulation research on ad hoc wireless networks must necessarily model radio propagation using simplifying assumptions. Although it is tempting to assume that all radios have circular range, have perfect coverage in that range, and travel on a two-dimensional plane, most researchers are increasingly aware of the need to represent more realistic features, including hills, obstacles, link asymmetries, and unpredictable fading. Although many have noted the complexity of real radio propagation, and some have quantified the effect of overly simple assumptions on the simulation of ad hoc network protocols, we provide a comprehensive review of six assumptions that are still part of many ad hoc network simulation studies. In particular, we use an extensive set of measurements from a large outdoor routing experiment to demonstrate the weakness of these assumptions, and show how these assumptions cause simulation results to differ significantly from experimental results. We close with a series of recommendations for researchers, whether they develop protocols, analytic models, or simulators for ad hoc 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 wireless networks, nearly all published MANET articles are buttressed with simulation results, and the simulations are based on common simplifying assumptions. Many such assumptions may

---

Note to readers who may have read the 2003 version of this paper as a TR [KNE03]: this revised version of the paper has an entirely new data set collected from a live ad hoc network experiment, a simulation study to demonstrate the impact of these axioms on three ad hoc routing protocols, and a new list of recommendations for routing protocol designers.

be too simple; a recent article in *IEEE Communications* warns that "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" [PJL02]. It then proceeded to survey 2200 published network simulation results to point out systemic flaws.

We recognize that the MANET research community is increasingly aware of the limitations of the common simplifying assumptions. Our goal in this paper is to make a constructive contribution to the MANET community by a) quantitatively demonstrating the weakness of these assumptions, b) comparing simulation results to experimental results to identify how simplistic radio models can lead to misleading results in ad hoc network research, c) contributing a real dataset that should be easy to incorporate into simulations, and d) listing recommendations for the designers of protocols, models, and simulators.

2 Radios in Theory and Practice

The top example in Figure 1 provides a simple model of radio propagation, one that is used in many simulations of ad hoc networks; contrast it to the bottom example of a real signal-propagation map, drawn at random from the web. Measurements of Berkeley Motes demonstrate a similar non-uniform non-circular behavior [GKW+02, ZHKS04]. The simple model is based on Cartesian distance in an X-Y plane. More realistic models take into account antenna height and orientation, terrain and obstacles, 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.

There are real risks to protocol designs based on overly simple models of radio propagation. First, "typical" network connectivity graphs look quite different in reality than they do on a Cartesian grid. An antenna placed top of a hill has direct connectivity with all other nearby radios, for example, an effect that 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., beacons) or "snooping" may therefore work significantly worse in reality than they do in simulation.

Figure 2 depicts one immediate drawback to the over-simplified model of radio propagation. The three different models in the figure, the Cartesian "Flat Earth" model, a three-dimensional model that includes a single hill, and a model that includes (absorptive) obstacles, all produce entirely different connectivity graphs, even though the nodes are in the same two-dimensional positions. As all the nodes move, the ways in which the connectivity graph changes 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 theoretical model (T) 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 a low rate of change in link connectivity. As such, this model makes it easy to determine when two nodes are, or are not, "neighbors" in the ad hoc network sense.

In the 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
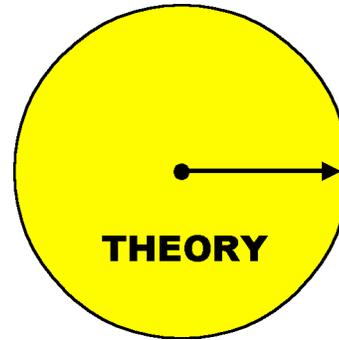
Typical theoretical model
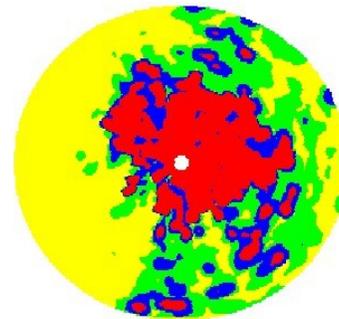
Source: Comgate Engineering
http://www.comgate.com/ntdsign/wireless.html

Figure 1: Real radios, such as the one at the bottom, are more complex than the common theoretical model at the top. Here different colors, or shades of gray, represent different signal qualities.

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 stationary 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 not considered 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 may fail.
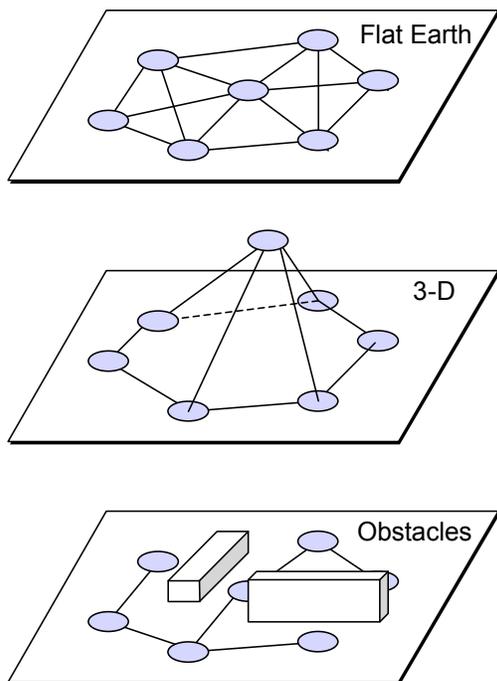
2

Figure 2: The Flat Earth model is overly simplistic.

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

Although the 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 (and link connectivity) varies much more rapidly in practice than in theory. 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, the propagation distances, and the timing of collisions can significantly alter the behavior of even the simplest flood-oriented network protocols [GKW+02]. Others [GC04, ZHKS04] have recently used two-node experiments to quantify specific characteristics of radio propagation, and used
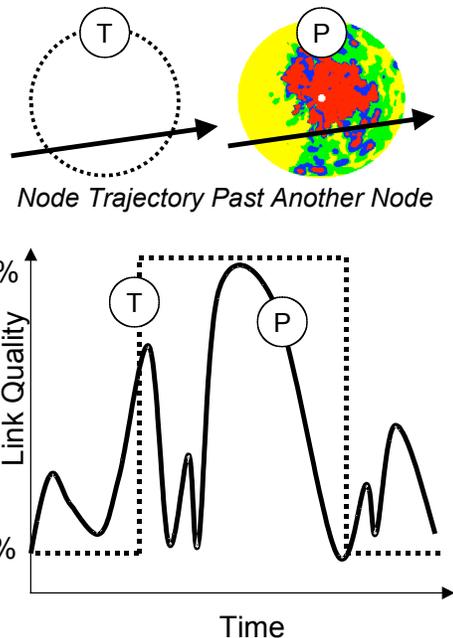
simulation to evaluate the impact of those characteristics on ad hoc routing protocols.

In summary,"good enough" radio models are 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   Models used in research

We surveyed a set of MobiCom and MobiHoc proceedings from 1995 through 2002. 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 results. Note that even in the best years, the Simple and Flat-Earth papers significantly outnumber the Good papers. A few [TMB01, JLW+96] deserve commendation for thoughtful channel models.

Flat Earth models are based on Cartesian X–Y proximity, that is, nodes A and B communicate if
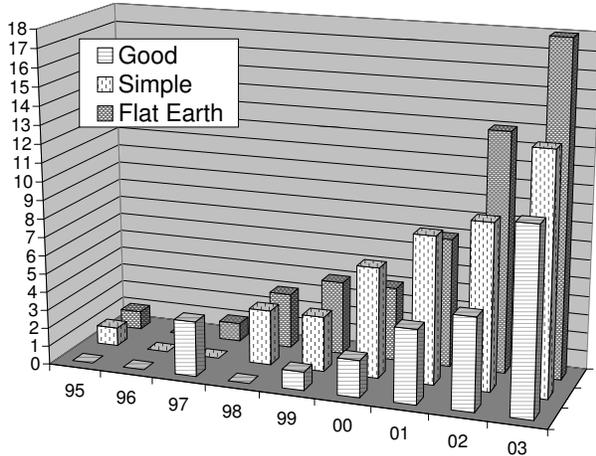
Figure 4: The number of papers in each year of Mobicom and MobiHoc that fall into each category.

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].[1] 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., 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 "Friis-free-space" model in the literature, and a *two-ray ground-reflection model*. Both are described in the `ns-2` document package [FV02].

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 sim-

ulations. 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, `ns-2` added a third channel model—the "shadowing" model described earlier by Lee [Lee82]—to account for indoor obstructions and outdoor shadowing via a probabilistic model [FV02]. The problem with `ns-2`'s shadowing model is that the model does not consider correlations: a real shadowing effect has strong correlations between two locations that are close to each other. More precisely, the shadow fading should be modeled as a two-dimensional log-normal random process with exponentially decaying spatial correlations (see [Gud91] for details). To our knowledge, only a few simulation studies include a valid shadowing model. For example, WiPPET considers using the correlated shadowing model to compute a gain matrix to describe radio propagation scenarios [KLM$^+$00]. WiPPET, however, only simulates cellular systems. The simulation model we later use for this study considers the shadowing effect as a random process that is temporally correlated; between each pair of nodes we use the same sample from the log-normal distribution if the two packets are transmitted within a pre-specified time period.[3]

Zhou et al. recently explored how signal strength varied with the angle between sender and receiver, between different (supposedly identical) senders, and with battery level. They developed a modifi-

---

[1]Other network simulators sometimes have better radio models. OpNet is one commercial example; see opnet.com. Most of the research literature, however, uses ns-2.

[2]See also Lundberg [Lun02], Sections 4.3.4–4.3.5, for additional remarks on the two-ray model's lack of realism.

[3]A recent study by Yuen et al. proposes a novel approach to modeling the correlation as a Gauss-Markov process [YLA02]. We are currently investigating this approach.

cation to path loss models that adds some random variation across angles and across senders, and then show how these better models lead to different simulation results than the original models. Different routing algorithms react differently to the more realistic radio model, leading a better understanding of each algorithm's strengths and weaknesses. Although they motivate their work with 2-node experiments, they do not have the ability to compare large-scale experiments with their simulation results as we do.

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 macrocells 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.

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, not all of which are orthogonal, 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.

There are many combinations of these axioms seen in the literature. In extreme cases, the combination of these axioms leads to a simple model like that in the top diagram in Figure 1. Some papers assume Axioms 0–4 and yet use a simple signal propagation model that expresses some fading with distance; a threshold on signal strength determines reception. Some papers assume Axioms 0–3 and add a reception probability to avoid Axiom 4.

In this paper we address the research community interested in ad hoc routing protocols and other distributed protocols at the network layer. The network layer rests on the physical and medium-access (MAC) layers, and its behavior is strongly influenced by their behavior. Indeed many MANET research projects consider the physical and medium-access layer as a single abstraction, and use the above axioms to model their combined behavior. We take this network-layer point of view through the remainder of the paper. Although we mention some of the individual physical- and MAC-layer effects that influence the behavior seen at the network layer, we do not attempt to identify precisely which effects cause which behaviors; such an exercise is beyond the scope of this paper. In the next two sections we show that 1) the above axioms do not adequately describe the network-layer's view of the world, and that 2) the use of these axioms leads simulations to results that differ radically from reality.

5 The Reality

Unfortunately, real wireless network devices are not nearly as simple as those considered by the axioms in

the preceding section. Although Gaertner and Cahill explicitly explore the relationship between link quality and radio characteristics or environmental conditions, they do so with only two nodes and with no evaluation of the impact on simulation or implemented routing protocols [GC04]. Similarly, Zhou et al. use two-node experiments to motivate their study of the impact of radio irregularity on simulation results [ZHKS04], but explore only that issue and do not validate their simulation study with experimental data.

In this section, we use data collected from a large MANET experiment in which forty laptops with WiFi and GPS capability roamed a field for over an hour while exchanging broadcast beacons. Although our experiment represents just one environment, it is not unlike that used in many simulation-based studies today (a flat square field with no obstacles and randomly moving nodes). For the purposes of this paper, it serves to demonstrate that the axioms are untrue even in a simple environment, and that fairly sophisticated simulation models were necessary for reasonable accuracy.

At different times during the field test, the laptops also tested the costs and capabilities of different routing algorithms. A companion paper [GKN+04] explores that experiment and compares four routing protocols, in what is to our knowledge the largest outdoor experiment with a mobile ad hoc wireless network.[4]

We begin with a description of the experimental conditions and the data collected.

5.1 Experimental data

The outdoor routing experiment took place on a rectangular athletic field measuring approximately 225 (north-south) by 365 (east-west) meters. This field can be roughly divided into four flat, equal-sized sections, three of which are at the same altitude, and one of which is approximately four to six meters lower. There was a short, steep slope between the upper and lower sections.

Each Linux laptop[5] had a wireless card[6] operating in peer-to-peer mode at 2 Mb/s. This fixed rate made it much easier to conduct the experiment, since we did not need to track (and later model) automatic changes to each card's transmission rate. Most current wireless cards are multi-rate, however, which could lead to **Axiom 6: Each packet is transmitted at the same bit rate.** We leave the effects of this axiom as an area for future work.

To reduce interference from our campus wireless network, we chose a field physically distant from campus, and we configured the cards to use wireless channel 9, for maximum possible separation from the standard channels (1, 6 and 11). In addition, we configured each laptop to collect signal-strength statistics for each received packet.[7] Finally, each laptop had a Garmin eTrex GPS unit attached via the serial port. These GPS units did not have differential GPS capabilities, but were accurate to within thirty feet during the experiment.

Each laptop recorded its current position (latitude, longitude and altitude) once per second, synchronizing its clock with the GPS clock to provide subsecond, albeit not millisecond, time synchronization. Every three seconds, the *beacon service program* on each laptop *broadcast* a beacon containing the current laptop position (as well as the last known positions of the other laptops). Each laptop that received such a beacon updated its internal position table, and sent a *unicast acknowledgment* to the beacon sender via UDP. Each laptop recorded all incoming and outgoing beacons and acknowledgments in another log file. The beacons allowed us to maintain a continuous picture of network connectivity, and, for-

---

[4]Lundgren et al. [LLN+02] briefly describes a slightly larger experiment, but indoors, with a limited mobility pattern, and with only a brief comparison of two routing algorithms.

[5]A Gateway Solo 9300 running Linux kernel version 2.2.19 with PCMCIA Card Manager version 3.2.4

[6]We used a Lucent (Orinoco) Wavelan Turbo Gold 802.11b. Although these cards can transmit at different bit rates and can auto-adjust this bit rate depending on the observed signal-to-noise ratio, we used an ad hoc mode in which the transmission rate was fixed at 2 Mb/s. Specifically we used firmware version 4.32 and the proprietary ad hoc "demo" mode originally developed by Lucent. Although the demo mode has been deprecated in favor of the IEEE 802.11b defined IBSS, we used it to ensure consistency with a series of ad hoc routing experiments of which this outdoor experiment was the culminating event. Our general results, which revolve around signal-strength measurements and beacon-reception probabilities, do not depend on a particular ad hoc mode.

[7]We used the `wvlan_cs`, rather than `orinoco_cs`, driver.

tunately, also represent network traffic that would be exchanged in many real MANET applications, such as our earlier work [Gra00] where soldiers must see the current locations of their fellows. Finally, every second each laptop queried the wireless driver to obtain the signal strength of the most recent packet *received* from every other laptop, and recorded this signal strength information in a third log.[8] Querying every second for all signal strengths was much more efficient than querying for individual signal strengths after each received packet.

These three logs provide all the data that we need to examine the axioms. Much more was going on in the experiment, however, since the overall goal was to compare the performance of four routing algorithms, APRL [KK98], AODV [PR99], ODMRP [LSG02], and STARA [GK97]. The laptops automatically ran each routing algorithm for 15 minutes, generating random UDP data traffic for thirteen out of the fifteen minutes, and pausing for two minutes between each algorithm to handle cleanup and setup chores. The traffic-generation parameters were set to produce the traffic volumes observed in our prototype situational-awareness applications [Gra00], approximately 423 outgoing bytes (including UDP, IP and Ethernet headers) per laptop per second, a relatively modest traffic volume. We do not describe the algorithms further here, since the routing and data traffic serves only as another source of collisions from the standpoint of the axioms. Note, however, that each transmitted packet was destined for only a single recipient, reducing ODMRP to the unicast case.

Finally, the laptops moved continuously. At the start of the experiment, the participants were divided into equal-sized groups of ten each, each participant given a laptop, and each group instructed to randomly disburse in one of the four sections of the field (three upper and one lower). The participants then walked continuously, always picking a section different than the one in which they were currently located, picking a random position within that section, walking to that position in a straight line, and then repeating. This approach was chosen since it

was simple, but still provided continuous movement to which the routing algorithms could react, as well as similar spatial distributions across each algorithm.

During the experiment, seven laptops generated no network traffic due to hardware and configuration issues, and an eighth laptop generated the position beacons only for the first half of the experiment. We use the data from the remaining thirty-two laptops to test the axioms, although later we simulate thirty-three laptops since only seven laptops generated no network traffic at all. In addition, STARA generated an overwhelming amount of control traffic, and we excluded the STARA portion of the experiment from our axiom tests. The final axiom dataset contains fifty-three contiguous minutes of beacons and acknowledgments for thirty-two laptops.

5.2 Axiom 0

The world is flat.

Common stochastic radio propagation models assume a flat earth, and yet clearly the Earth is not flat. Even at the short distances considered by most MANET research, hills and buildings present obstacles that dramatically affect wireless signal propagation. Furthermore, the wireless nodes themselves are not always at ground level; indeed, Gaertner and Cahill noted a significant change in link quality between ground-level and waist-level nodes [GC04].

Even where the ground is nearly flat, note that wireless nodes are often used in multi-story buildings. Indeed two nodes may be found at exactly the same x, y location, but on different floors. (This condition is common among the WiFi access points deployed on our campus.) 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.

We need no data to "disprove" this axiom. Ultimately, it is the burden of all MANET researchers to either a) use a detailed and realistic terrain model, accounting for the effects of terrain, or b) clearly condition their conclusions as being valid only on flat, obstacle-free terrain.

---

[8]For readers familiar with Linux wireless services, note that we increased the IWSPY limit from 8 to 64 nodes, so that we could capture signal-strength information for the full set of laptops.

5.3 Axioms 1 and 2

A radio's transmission area is circular.

All radios have equal range.

The real-world radio map of Figure 1 makes it clear that the signal coverage area of a radio is far from simple. Not only is it neither circular nor convex, it often is non-contiguous.

We combine the above two intuitive axioms into a more precise, testable axiom that corresponds to the way the axiom often appears (implicitly) in MANET research.

Testable Axiom 1. The success of a transmission from one radio to another depends only on the distance between radios.

Although it is true that successful communication usually becomes less likely with increasing distance, there are many other factors: (1) All radios are not identical. Although in our experiment we used "identical" WiFi cards, there are reasonable applications where the radios or antennas vary from node to node. (2) Antennas are not perfectly omnidirectional. Thus, the angle of the sender's antenna, the angle of the receiver's antenna, and their relative locations all matter. (3) Background noise varies with time and location. Finally, (4) there are hills and obstacles, including people, that block or reflect wireless signals (that is, Axiom 0 is false).

From the point of view of the network layer, these physical-layer effects are compounded by MAC-layer effects, notably, that collisions due to transmissions from other nodes in the ad hoc network (or from third parties outside the set of nodes forming the network) reduce the transmission success in ways that are unrelated to distance. In this section, we use our experimental data to examine the effect of antenna angle, sender location, and sender identity on the probability distribution of beacon reception over distance.

We first demonstrate that the probability of a beacon packet being received by nearby nodes depends strongly on the angle between sender and receiver antennas. In our experiments, we had each student carry their "node," a closed laptop, under their arm with the wireless interface (an 802.11b device in PC-card format) sticking out in front of them. By examining successive location observations for the node, we compute the orientation of the antenna (wireless card) at the time it sent or received a beacon. Then, we compute two angles for each beacon: the angle between the sender's antenna and the receiver's location, and the angle between the receiver's antenna and the sender's location. Figure 5 illustrates the first of these two angles, while the second is the same figure except with the labels Source and Destination transposed. Figure 6 shows how the beacon-reception probability varied with both angles.

To compute Figure 6, we consider all possible values of each of the two angles, each varying from $[-180, 180)$. We divide each range into buckets of 45 degrees, such that bucket 0 represents angles in $[0, 45)$, bucket 45 represents angles in $[45, 90)$, and so forth. Since we bucket both angles, we obtain the two-dimensional set of buckets shown in the figure. We use two counters for each bucket, one accounting for actual receptions, and the other for potential receptions (which includes actual receptions). Each time a node sends a beacon, every other laptop is a potential recipient. For every other laptop, therefore, we add one to the potential-reception count for the bucket representing the angles between the sender and the potential recipient. If we can find a received beacon in the potential recipient's beacon log that matches the transmitted beacon, we also add one to the actual-reception count for the appropriate count. The beacon reception ratio for a bucket is thus the number of actual receptions divided by the number of potential receptions. Each beacon-reception probability is calculated without regard to distance, and thus represents the reception probability across all distances. In addition, for all of our axiom analyses, we considered only the western half of the field, and incremented the counts only when both the sender and the (potential) recipient were in the western half. By considering only the western half, which is perfectly flat and does not include the lower-altitude section, we eliminate the most obvious terrain effects from our results. Overall, there were 40,894 beacons transmitted in the western half of the field, and after matching and filtering, we had 275,176 laptop pairs, in 121,250 of which the beacon was received, and in 153,926 of which the beacon was not received.

Figure 6 shows that the orientation of both antennas was a significant factor in beacon reception. Of course, there is a direct relationship between the an-
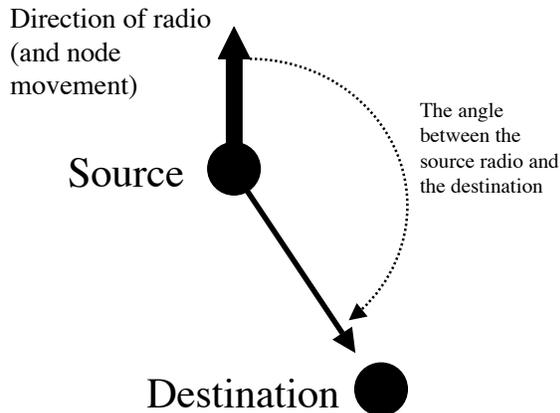
Figure 5: The angle between the sending laptop's antenna (wireless card) and the destination laptop. We express the angles on the scale of -180 to 180, rather than 0 to 360, to better capture the inherent symmetry. -180 and 180 both refer to the case where the sending antenna is pointing directly away from the intended destination.

tenna angles and whether the sender or receiver (human or laptop) is between the two antennas. With a sender angle of 180, for example, the receiver is directly behind the sender, and both the sender's body and laptop serves as an obstruction to the signal. A different kind of antenna, extending above the level of the participants' heads, would be needed to separate the angle effects into two categories, effects due to human or laptop obstruction, and effects due to the irregularity of the radio coverage area.

Although the western half of our test field was flat, we observed that the beacon-reception probability distribution varied in different areas. We subdivided the western half into four equal-sized quadrants (northwest, northeast, southeast, southwest), and computed a separate reception probability distribution for beacons sent from each quadrant. Figure 7 shows that the distribution of beacon-reception probability was different for each quadrant, by about 10–15 percent for each distance. We bucketed the laptop pairs according to the distance between the sender and the (intended) destination—the leftmost bar in the graph, for example, is the reception probability for laptop pairs whose separation was in the range $[0, 25)$. Although there are
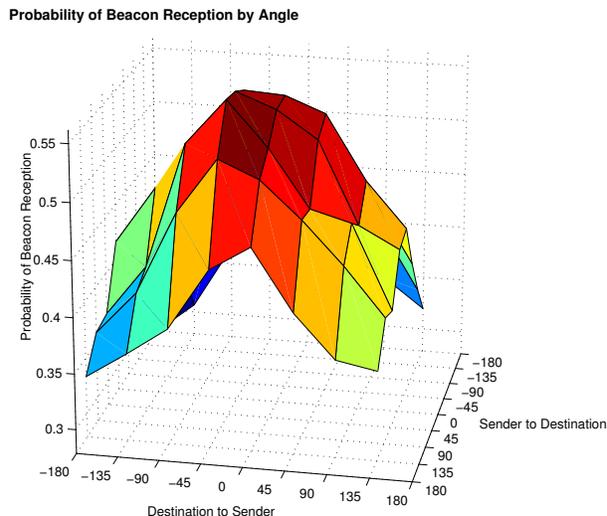

Figure 6: The probability of beacon reception (over all distances) as a function of the two angles, the angle between the sender's antenna orientation and the receiver's location, and the angle between the receiver's antenna orientation and the sender's location. In this plot, we divide the angles into buckets of 45 degrees each, and include only data from the western half of the field.

many possible explanations for this quadrant-based variation, whether physical terrain, external noise, or time-varying conditions, the difference between distributions is enough to make it clear that the location of the sender is not to be ignored.

The beacon-reception probability in the western half of the field also varied according to the identity of the sender. Although all equipment used in every node was an identical model purchased in the same lot and configured identically, the distribution was different for each sender. Figure 8 shows the mean and standard deviation of beacon-reception probability computed across all sending nodes, for each bucket between 0 and 300 meters. The buckets between 250 and 300 meters were nearly empty. Although the mean across nodes, depicted by the boxes, is steadily decreasing, there also is substantial variation across nodes, depicted by the standard-deviation bars on each bucket. This variation cannot be explained entirely by manufacturing variations within the antennas, and likely includes terrain, noise and other factors, even on our space of flat, open ground. It also is important to note, however, that there are
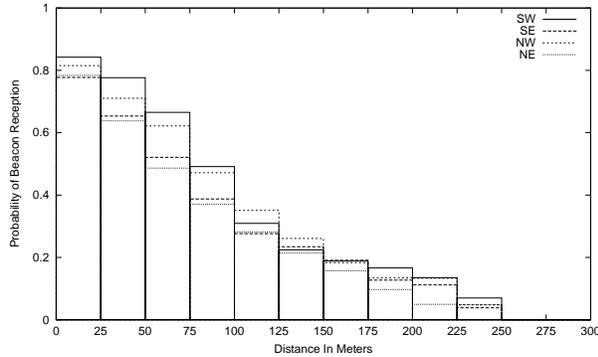
Figure 7: The probability of beacon reception varied from quadrant to quadrant within the western half of the field.
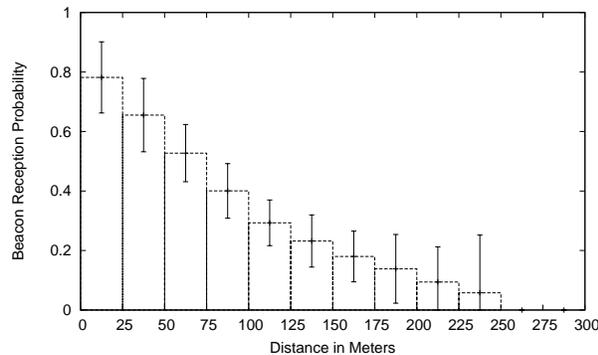

Figure 8: The average and standard deviation of reception probability across all nodes, again for the western half of the field.

only 500-1000 data points for each (laptop, destination bucket) pair. With this number of data points, statistical-significance issues come into play. In particular, if a laptop is moving away from most other laptops, we might cover only a small portion of the possible angles, leading to markedly different results than for other laptops. Overall, the effect of identity on transmission behavior bears further study with experiments specifically designed to test it.

In other work, 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].[9] The resulting contour map is not circular, nor convex, nor even monotonically decreasing with distance. Indeed, since the coverage

---

[9]The Berkeley mote is currently the most common research platform for real experiments with ad hoc sensor networks.

area of a radio is not circular, it is difficult to even define the "range" of a radio.

Zhou et al. [ZHKS04] also note that signal strength varies with the angle between sender and receiver, angle between receiver and sender, and sender identity, using two-node experiments.

5.4 Axiom 3

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

More precisely,

Testable Axiom 3: If an unacknowledged message from A to B succeeds, an immediate reply from B to A succeeds.

This wording adds a sense of time, since it is clearly impossible (in most MANET technologies) for A and B to transmit at the same time and result in a successful message, and since A and B may be moving, it is important to consider symmetry over a brief time period so that A and B have not moved apart.

There are many factors affecting symmetry, from the point of view of the network layer, including the physical effects mentioned above (terrain, obstacles, relative antenna angles) as well as MAC-layer collisions. It is worth noting that the 802.11 MAC layer includes an internal acknowledgment feature, and a limited amount of re-transmission attempts until successful acknowledgment. Thus, the network layer does not perceive a frame as successfully delivered unless symmetric reception was possible. Thus, for the purposes of this axiom, we chose to examine the broadcast beacons from our experimental dataset, since the 802.11 MAC has no internal acknowledgment for broadcast frames. Since all of our nodes sent a beacon every three seconds, we were able to identify symmetry as follows: whenever a node B received a beacon from node A, we checked to see whether B's next beacon was also received by node A.

Figure 9 shows the conditional probability of symmetric beacon reception. If the physical and MAC layer behavior was truly symmetric, this probability would be 1.0 across all distances. In reality, the probability was never much more than 0.8, most likely due to MAC-layer collisions between beacons. Since
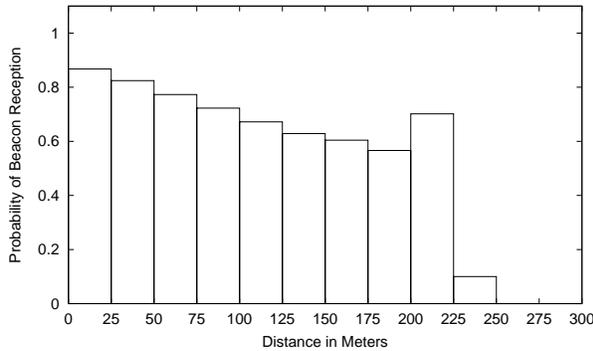
Figure 9: The conditional probability of symmetric beacon reception as it varied with the distance between two nodes, again for the western half of the field.
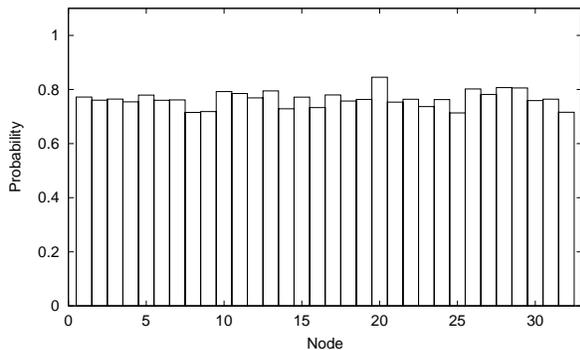

Figure 10: The conditional probability of symmetric beacon reception as it varied across individual nodes, again for the western half of the field.

this graph depends on the joint probability of a beacon arriving from A to B and then another from B to A, the lower reception probability of higher distances leads to a lower joint probability and a lower conditional probability. Figure 10 shows how the conditional probability varied across all the nodes in the experiment. The probability was consistently close to its mean 0.76, but did vary from node to node with a standard deviation of 0.029 (or 3.9%). Similarly, when calculated for each of the four quadrants (not shown), the probability also was consistently close to its mean 0.76, but did have a standard deviation of 0.033 (or 4.3%).

In other work, 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 asym-

metric 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.]

Zhou et al. also found through simulation that the use of angular variations in signal strength naturally led to asymmetric links in simulation, and that some protocols were unable to adapt gracefully to asymmetry [ZHKS04].

Overall, it is clear that reception is far from symmetric. Nonetheless, many researchers assume this axiom is true, and that all network links are bidirectional. Some do acknowledge that real links may be unidirectional, and usually discard those links so that the resulting network has only bidirectional links. In a network with mobile nodes or in a dynamic environment, however, link quality can vary frequently and rapidly, so a bidirectional link may become unidirectional at any time. It is best to develop protocols that do not assume symmetry.

5.5 Axiom 4

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

Testable Axiom 4: The reception probability distribution over distance exhibits a sharp cliff; that is, under some threshold distance (the "range") the reception probability is 1 and beyond that threshold the reception probability is 0.

Looking back at Figure 8, we see that the beacon-reception probability does indeed fade with the distance between the sender and the receiver, rather than remaining near 1 out to some clearly defined "range" and then dropping to zero. There is no visible "cliff." The common `ns-2` model, however, assumes that frame transmission is perfect, within the range of a radio, and as long as there are no collisions. Although `ns-2` 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 project[10] that are being used to explore how sophisticated channel models affect simulation outcomes.

Takai examines the effect of channel models on simulation outcomes [TBTG01], and also concluded

---

[10]http://www.scalable-networks.com/pdf/mobihocpreso.pdf

that different physical layer models can have dramatically different effect on the simulated performance of protocols [TMB01], but lack of data prevented them from further validating simulation results against real-world experiment results, which they left as future work. Zhou et al. also did not validate their simulation results against real-world experiment results. We compared the simulation results with data collected from a real-world experiment, and recommend below that simple models of radio propagation should be avoided whenever comparing or verifying protocols, unless that model is known to specifically reflect the target environment.

5.6 Axiom 5

Signal strength is a simple function of distance.

Rappaport [Rap96] notes that the average signal strength should fade with distance according to a power-law model. While this is true, one should not underestimate the variations in a real environment caused by obstruction, reflection, refraction, and scattering. In this section, we show that there is significant variation for individual transmissions.

Testable Axiom 5: We can find a good fit between a simple function and a set of (distance, signal strength) observations.

To examine this axiom, we consider only received beacons, and use the recipient's signal log to obtain the signal strength associated with that beacon. More specifically, the signal log actually contains per-second entries, where each entry contains the single strength of the most recent packet received from each laptop. If a data or routing packet arrives immediately after a beacon, the signal-log entry actually will contain the signal strength of that second packet. We do not check for this situation, since the signal information for the second packet is just as valid as the signal information for the beacon. It is best, however, to view our signal values as those observed within one second of beacon transmission, rather than the values associated with the beacons themselves.

As a starting point, Figure 11 shows the *mean* beacon signal strength observed during the experiment as a function of distance, as well as best-fit linear and
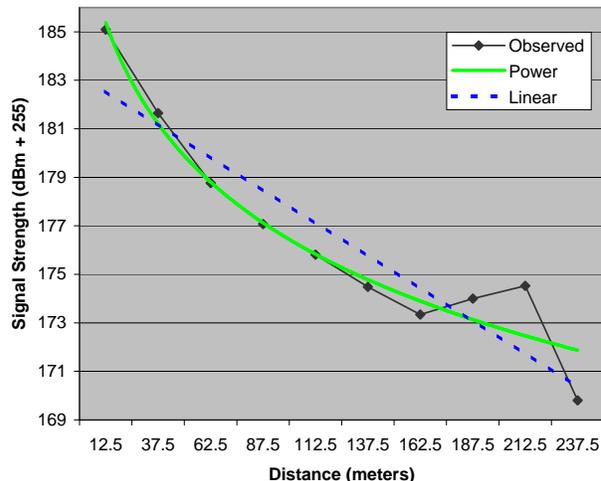

Figure 11: Linear and power-curve fits for the mean signal strength observed in the western half of the field. Note that we show the signal strength as reported by our wireless cards (which is dBm scaled to a positive range by adding 255), and we plot the mean value for each distance bucket at the midpoint of that bucket.

power curves. The power curve is a good fit and validates Rappaport's observation. When we turn our attention to the signal strength of individual beacons, however, as shown in Figures 12 and 13, there clearly is no simple (non-probabilistic) function that will adequately predict the signal strength of an individual beacon based on distance alone.

The reason for this difficulty is clear: our environment, although simple, is full of obstacles and other terrain features that attenuate or reflect the signal, and the cards themselves do not necessarily radiate with equal power in all directions. In our case, the most common obstacles were the people and laptops themselves, and in fact, we initially expected to discover that the signal strength was better behaved across a specific angle range (per Figure 6) than across all angles. Even for the seemingly good case of both source and destination angles between 0 and 45 degrees (i.e., the sender and receiver roughly facing each other), we obtain a distribution (not shown) remarkably similar to Figure 12. Other angle ranges also show the same distribution as Figure 12.

Overall, noise-free, reflection-free, obstruction-free, uniformly-radiating environments are simply not real, and signal strength of individual transmis-
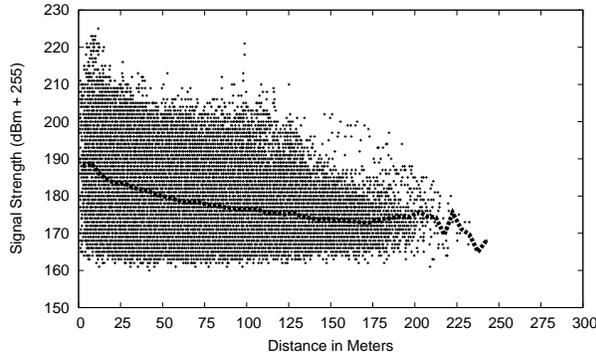
12

Figure 12: A scatter plot demonstrating the poor correlation between signal strength and distance. We restrict the plot to beacons both sent and received on the western half of the field, and show the mean signal strength as a heavy dotted line.
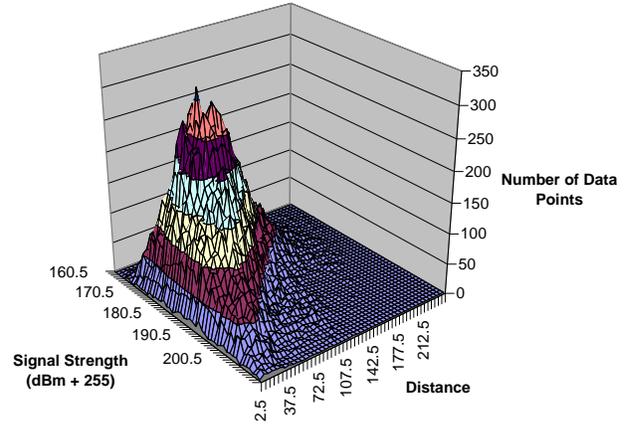

Figure 13: Same as Figure 12 except that it shows the *number* of observed data points as a function of distance and signal strength. There is significant weight relatively far away from the mean value.

sions will never be a simple function of distance. Researchers must be careful to consider how sensitive their simulation results are to signal variations, since their algorithms will encounter significant variation once deployed.

6   Impact

We demonstrate above that the axioms are untrue, but a key question remains: what is the effect of these axioms on the quality of simulation results? In this section, we begin by comparing the results of our outdoor experiment with the results of a best-effort simulation model, and then progressively weaken the model by assuming some of the axioms. The purpose of this study is not to claim that our simulator can accurately model the real network environment, but instead to show quantitatively the impact of the axioms on the simulated behavior of routing protocols.

Clearly, analytical or simulation research in wireless networking must work with an abstraction of reality, modeling the behavior of the wireless network below the layer of interest. Unfortunately, overly simplistic assumptions can lead to misleading or incorrect conclusions. Our results provide a counterexample to the notion that these axioms are sufficient for research on ad hoc routing algorithms. We do not claim to validate, or invalidate, the results of any other published study. Indeed, our point is that the burden is on the authors of past and future studies to a) clearly lay out their assumptions, b) demonstrate

whether those assumptions are reasonable within the context of their study, and c) clearly identify any limitations in the conclusions they draw.

While others have used simulation to explore the impact of different radio propagation models [TMB01, ZHKS04], we use the identical implementation of the routing protocol in both the simulator and the experiment [LYN+04], use a large number of nodes in an outdoor experiment [GKN+04], and are able to compare our simulation results with the actual experiment.

6.1   Our simulator

Our SWAN simulator for wireless ad hoc networks provides an integrated, configurable, and flexible environment for evaluating ad hoc routing protocols, especially for large-scale network scenarios. SWAN contains a detailed model of the IEEE 802.11 wireless LAN protocol and a stochastic radio channel model, both of which were used in this study.

We used SWAN's direct-execution simulation techniques to execute within the simulator the *same* routing code that was used in the experiments from the previous section [LYN+04]. We modified the real routing code only slightly to allow multiple instances of a routing protocol implementation to run simultaneously in the simulator's single address space. We extended the simulator to read the node mobility and application-level data logs generated by

13

the real experiment. In this way, we were able to re-produce the same network scenario in simulation as in the real experiment. Moreover, by directly running the routing protocols and the beacon service program, the simulator generated the same types of logs as in the real experiment.

In the next few sections, we describe three simulation models with progressively unrealistic assumptions, and then present results to show the impact.

6.2 Our best model

We begin by comparing the results of the outdoor experiment with the simulation results obtained with our best signal propagation model and a detailed 802.11 protocol model. The best signal propagation model is a stochastic model that captures radio signal attenuation as a combination of two effects: small-scale fading and large-scale fading. Small-scale fading describes the rapid fluctuation in the envelope of a transmitted radio signal over a short period of time or a small distance, and primarily is caused by multipath effects. Although small-scale fading is in general hard to predict, wireless researchers over the years have proposed several successful statistical models for small-scale fading, such as the Rayleigh and Ricean distributions. Large-scale fading describes the slowly varying signal-power level over a long time interval or a large distance, and has two major contributing factors: distance path-loss and shadow fading. The distance path-loss models the average signal power loss as a function of distance: the receiving signal strength is proportional to the distance between the transmitter and the receiver raised to a given exponent. Both the free-space model and the two-ray ground reflection model mentioned earlier can be classified as distance path-loss models. The shadow fading describes the variations in the receiving signal power due to scattering; it can be modeled as a zero-mean log-normal distribution. Rappaport [Rap96] provides a detailed discussion of these and other models.

For our simulation, given the light traffic used in the real experiment, we used a simple SNR threshold approach instead of a more computational intensive BER approach. Under heavier traffic, this choice might have substantial impact [TMB01]. For the propagation model, we chose 2.8 as the distance path-loss exponent and 6 dB as the shadow fading

	Experiment	Simulation	Error
AODV	42.3%	46.8%	10.5%
APRL	17.5%	17.7%	1.1%
ODMRP	62.6%	56.9%	-9.2%

Table 1: Comparing packet delivery ratios between real experiment and simulation.

log normal standard. These values, which must be different for different types of terrain, produce signal propagation distances consistent with our observations from the real network. Finally, for the 802.11 model, we chose parameters that match the settings of our real wireless cards. We then conducted the simulation of the wireless network with 40 nodes, of which 7 did not generate any network traffic, but were available for selection as potential packet destinations. This duplicated the 7 crashed nodes from the real experiment, and allowed us to reproduce the same traffic pattern.

Table 1 shows the difference in the overall packet delivery ratio (PDR)—which is the total number of packets received by the application layer divided by the total number of packets sent—between the real experiment and the simulation. The simple propagation model produced relatively good results: the relative errors in predicted PDR were within 10% for all three routing protocols tested. We caution, however, that one cannot expect consistent results when generalizing the simple stochastic radio propagation model to deal with all network scenarios. After all, this model assumes some of the axioms we have identified, including flat earth, omni-directional radio propagation length, and symmetry. Thus this model, our best, nonetheless assumes some of the same axioms we discount in the preceding section! This ironic situation is testimony to the difficulty of detailed radio and environment modeling; in situations where such assumptions are clearly invalid—for example, in an urban area—we should expect the model to deviate further from reality. On the other hand, this approximation is sufficient for the purposes of this paper, because we can still demonstrate how the other axioms may affect performance.

On the other hand, since the model produced good results amenable to our particular outdoor experiment scenario, we use it in this study as the base-

line to quantify the effect of the axioms on simulation studies. As we show, these assumptions can significantly undermine the validity of the simulation results.

6.3 Simpler models

Next we weakened our simulator by introducing a simpler signal propagation model. We used the distance path-loss component from the previous model, but disabled the variations in the signal receiving power introduced by the stochastic processes. Note that these variations are a result of two distinct random distributions: one for small-scale fading and the other for shadow fading. The free-space model, the two-ray ground reflection model, and the generic distance path-loss model with a given exponent—all used commonly by wireless network researchers—differ primarily in the maximum distance that a signal can travel. For example, if we assume that the signal transmission power is 15 dBm and the receiving threshold is -81 dBm, the free-space model has a maximum range of 604 meters, the two-ray ground reflection model a range of 251 meters, and the generic path-loss model (with an exponent of 2.8) a range of only 97 meters. Indeed, the SWAN authors also noted that the receiving range plays an important role in ad hoc routing: longer distance shortens the data path and can drastically change the routing maintenance cost [LYN+04].

In this study, we chose to use the two-ray ground reflection model since its signal travel distance matches observations from the real experiment.[11] This weaker model assumes Axiom 4: "If I can hear you at all, I can hear you perfectly," and specifically the testable axiom "The reception probability distribution over distance exhibits a sharp cliff." Without variations in the radio channel, all signals travel the same distance, and successful reception is subject only to the state of interference at the receiver. In other words, the signals can be received successfully with probability 1 as long as no collision occurs during reception.

Finally, we consider a third model that further

---

[11]When we consider the full experiment field, which provides possible reception ranges of over 500 meters, we see almost no receptions beyond 250 meters. The 251-meter range of the 2-ray model is computed from a well-known formula, using a fixed transmit power (15 dBm) and antenna height (1.0 meter).
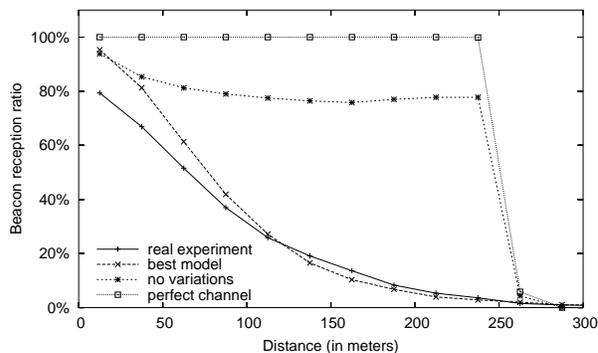

Figure 14: The beacon reception ratio at different distances between the sender and the receiver. The probability for each distance bucket is plotted as a point at the midpoint of its bucket; this format is easier to read than the boxes used in earlier plots.

weakens the simulator by assuming that the radio propagation channel is *perfect*. That is, if the distance between the sender and the recipient is below a certain threshold, the signal is received successfully with probability 1; otherwise the signal is always lost. The perfect-channel model represents an extreme case where the wireless network model introduces no packet loss from interference or collision, and the reception decision is based solely on distance. To simulate this effect, we bypassed the IEEE 802.11 protocol layer within each node and replaced it with a simple protocol layer that calculates signal reception based only on the transmission distance.

6.4 The Results

First, we look at the reception ratio of the beacon messages, which were periodically sent via broadcasts by the beacon service program on each node. We calculate the reception ratio by inspecting the entries in the beacon logs, just as we did for the real experiment. Figure 14 plots the beacon reception ratios during the execution of the AODV routing protocol. The choice of routing protocol is unimportant in this study since we are comparing the results between the real experiment and simulations. We understand that the control messages used by the routing protocol may slightly skew the beacon reception ratio due to the competition at the wireless channel.
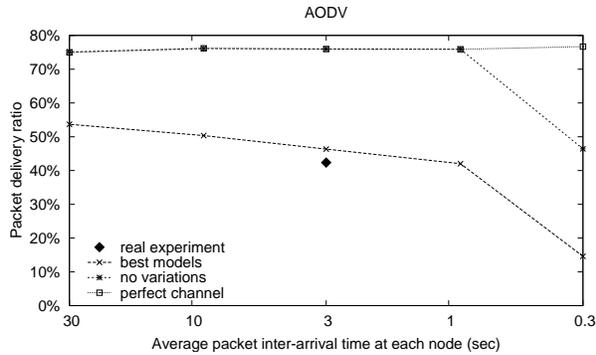
Compared with the two simple models, our best

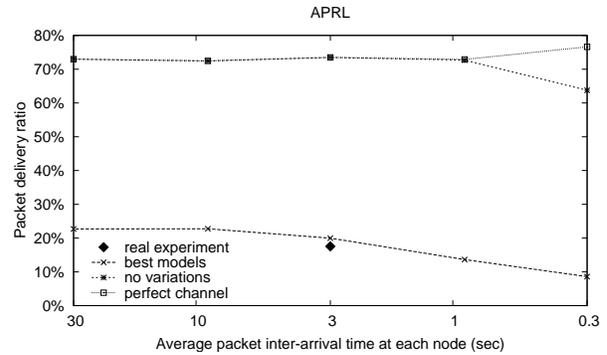Figure 15: Packet delivery ratios for AODV.

Figure 16: Packet delivery ratios for APRL.

Figure 17: Packet delivery ratios for ODMRP.

model is a better fit for the real experiment results. It does, however, slightly inflate the reception ratios at shorter distances and underestimate them at longer distances. More important for this study is the dramatic difference we saw when signal power variations were not included in the propagation model. The figure shows a sharp cliff in the beacon reception ratio curve: the quality of the radio channel changed abruptly from relatively good reception to zero reception as soon as the distance threshold was crossed. The phenomenon is more prominent for the perfect channel model. Since the model had no interference and collision effects, the reception ratio was 100% within the propagation range.

Next, we examine the effect of different simulation models on the overall performance of the routing protocols. Figures 15–17 show the packet delivery ratios, for the three ad hoc routing algorithms, as we varied the application traffic intensity by adjusting the average packet inter-arrival time at each node. Note the logarithmic scale for the x-axes in the plots. The real experiment's result is represented by a single point in each plot.

Figures 15–17 show that the performance of routing algorithms predicted by different simulation models varied dramatically. For AODV and APRL, both simple models exaggerated the packet delivery ratio *significantly*. In those models, the simulated wireless channel was much more resilient to errors than the real network, since there were no spatial or temporal fluctuations in signal power. Without variations, the signals had a much higher chance to be successfully received, and in turn, there were fewer route invalidations, and more packets were able to
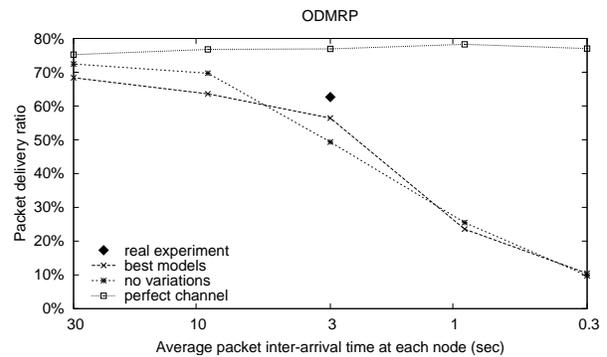
find routers to their intended destinations. The performance of the perfect-channel model remained insensitive to the traffic load since the model did not include collision and interference calculations at the receiver, explaining the divergence of the two simple models as the traffic load increases. For ODMRP, we cannot make a clear distinction between the performance of the best model and of the no-variation model. One possible cause is that ODMRP is a multicast algorithm and has a more stringent bandwidth demand than the strictly unicast protocols. A route invalidation in ODMRP triggers an aggressive route rediscovery process, and could cause significant packet loss under any of the models.

In summary, the assumptions embedded inside the wireless network model have a great effect on the simulation results. On the one hand, our best wireless network model assumes some of the axioms, yet the results do not differ significantly from the real experiment results. On the other hand, one must be extremely careful when assuming some of the ax-

16

ioms. If we had held our experiment in an environment with more hills or obstacles, the simulation results would not have matched as well. Even in this relatively flat environment, our study shows that proper modeling of the lossy characteristics of the radio channel has a significant impact on the routing protocol behaviors. For example, using our best model, one can conclude from Figure 15 and Figure 17 that ODMRP performed better than AODV with light traffic load (consistent with real experiment), but that their performance was comparable when the traffic was heavy. If we use the model without variations, however, one might arrive at the opposite conclusion, that AODV performed consistently no worse than ODMRP. The ODMRP results are interesting by themselves, since the packet-delivery degradation as the traffic load increases is more than might be expected for an algorithm designed to find redundant paths (through the formation of appropriate forwarding groups). Bae has shown, however, that significant degradation can occur as intermediate nodes move, paths to targets are lost, and route rediscovery competes with other traffic [BLG00]. In addition, the node density was high enough that each forwarding group could have included a significant fraction of the nodes, leading to many transmitted copies of each data packet. An exploration of this issue is left for future work.

7   Conclusions, recommendations

In recent years, dozens of Mobicom and Mobihoc 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.

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 enumerated the set of common assumptions used in MANET research, and presented a real-world experiment that strongly contradicts these "axioms." The results cast doubt on published simulation results that implicitly rely on these assumptions, e.g., by assuming how well broadcasts are received, or whether "hello" propagation is symmetric.

We conclude with a series of recommendations, ...for the MANET research community:

1. Choose your target environment carefully, clearly list your assumptions about that environment, choose simulation models and conditions that match those assumptions, and report the results of the simulation in the context of those assumptions and conditions.

2. Use a realistic stochastic model when verifying a protocol, or comparing a protocol to existing protocols. Furthermore, any simulation should explore a range of model parameters since the effect of these parameters is not uniform across different protocols. Simple models are still useful for the initial exploration of a broad range of design options, due to their efficiency.

3. Consider three-dimensional terrain, with moderate hills and valleys, and corresponding radio propagation effects. It would be helpful if the community agreed on a few standard terrains for comparison purposes.

4. 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.

5. Use real data as input to simulators, where possible. For example, 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. The dataset also may be helpful in the development of new, more realistic radio models.

17

...for simulation and model designers:

1. Allow protocol designers to run the same code in the simulator as they do in a real system [LYN$^+$04], making it easier to compare experimental and simulation results.

2. Develop a simulation infrastructure that encourages the exploration of a range of model parameters.

3. Develop a range of propagation models that suit different environments, and clearly define the assumptions underlying each model. Models encompassing both physical and data-link layer need to be especially careful.

4. Support the development of standard terrain and mobility models, and formats for importing real terrain data or mobility traces into the simulation.

...for protocol designers:

1. Consider carefully your assumptions of lower layers. In our experimental results, we found that the success of a transmission between radios depends on many factors (ground cover, antenna angles, human and physical obstructions, background noise, and competition from other nodes), most of which cannot be accurately modeled, predicted or detected at the speed necessary to make per-packet routing decisions. A routing protocol that relies on an acknowledgement quickly making it from target or source over the reverse path, that assumes that beacons or other broadcast traffic can be reliably received by most or all transmission-range neighbors, or that uses an instantaneous measure of link quality to make significant future decisions, is likely to function significantly differently outdoors than under simulation or indoor tests.

2. Develop protocols that adapt to environmental conditions. In our simulation results, we found that the relative performance of two algorithms (such as AODV and ODMRP) can change significantly, and even reverse, as simulation assumptions or model parameters change. Although some assumptions may not significantly affect the agreement between the experimental and simulation results, others may introduce radical disagreement. For similar reasons, a routing proto-

col tested indoors may work very differently outdoors. Designers should consider developing protocols that make few assumptions about their environment, or are able to adapt automatically to different environmental conditions.

3. Explore the costs and benefits of control traffic. Both our experimental and simulation results hint that there is a tension between the control traffic needed to identify and use redundant paths and the interference that this extra traffic introduces when the ad hoc routing algorithm is trying to react to a change in node topology. The importance of reducing interference versus identifying redundant paths (or reacting quickly to a path loss) might appear significantly different in real experiments than under simple simulations, and protocol designers must consider carefully whether extra control traffic is worth the interference price.

Availability. We will make our simulator and our dataset available to the research community upon completion of the camera-ready version of this paper. The dataset, including the actual position and connectivity measurements, would be valuable as input to future simulation experiments. The simulator contains several radio-propagation models.

Acknowledgements

References

[BLG00]   Sang Ho Bae, Sung-Ju Lee, and Mario Gerla. Unicast performance analysis of the ODMRP in a mobile ad-hoc network testbed. In *ICCCN 2000*, October 2000.

[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.

[GC04]   Gregor Gaertner and Vinny Cahill. Understanding link quality in 802.11 mobile ad hoc networks. *IEEE Internet Computing*, Jan/Feb 2004.

[GK97]   P. Gupta and P. R. Kumar. A system and traffic dependent adaptive routing algorithm for ad hoc networks. In *Proc. of the 36th IEEE Conference on Decision and Control*, pages 2375–2380, December 1997.

[GKN+04] Robert S. Gray, David Kotz, Calvin Newport, Nikita Dubrovsky, Aaron Fiske, Jason Liu, Christopher Masone, Susan McGrath, and Yougu Yuan. Outdoor experimental comparison of four ad hoc routing algorithms. Technical Report TR2004-511, Dept. of Computer Science, Dartmouth College, 2004.

[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.

[Gra00]   Robert S. Gray. Soldiers, agents and wireless networks: a report on a military application. In *Proceedings of the 5th International Conference and Exhibition on the Practical Application of Intelligent Agents and Multi-Agents*, April 2000.

[Gud91]   M. Gudmundson. Correlation model for shadow fading in mobile radio systems. *Electronics Letters*, 27(3):2145–2146, November 1991.

[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 Shnivakumar. 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.

[KK98]   B. Karp and H. T. Kung. Dynamic neighbor discovery and loopfree, multi-hop routing for wireless, mobile networks. Harvard University, May 1998.

[KLM+00] O. E. Kelly, J. Lai, N. B. Mandayam, A. T. Ogielski, J. Panchal, and R. D. Yates. Scalable parallel simulations of wireless networks with WiPPET: Modeling of radio propagation, mobility and protocols. *ACM/Kluwer MONET*, 5(3):199–208, September 2000.

[KNE03]   David Kotz, Calvin Newport, and Chip Elliott. The mistaken axioms of wireless-network research. Technical Report TR2003-467, Dept. of Computer Science, Dartmouth College, July 2003.

[Lee82]   W. .C. Y. Lee. *Mobile communications engineering*. McGraw-Hill, 1982.

[LLN+02] Henrik Lundgren, David Lundberg, Johan Nielsen, Erik Nordström, and Christian Tschudin. A large-scale testbed for reproducible ad hoc protocol evaluations. In *Proceedings of the Third Annual IEEE Wireless Communications and Networking Conference (WCNC)*, pages 412–418. IEEE Computer Society Press, March 2002.

[LSG02]   S.-J. Lee, W. Su, and M. Gerla. On-demand multicast routing protocol in multihop wireless mobile networks. *ACM/Kluwer MONET*, 7(6):441–453, December 2002.

[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.

[LYN+04] Jason Liu, Yougu Yuan, David M. Nicol, Robert S. Gray, Calvin C. Newport, David Kotz, and Luiz Felipe Perrone. Simulation validation using direct execution of wireless ad-hoc routing protocols. In *Proceedings of the Workshop on Parallel and Distributed Simulation (PADS)*, May 2004. Nominated for Best Paper award.

[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.

[PR99]   C. E. Perkins and E. M. Royer. Ad hoc on-demand distance vector routing. In *WMCSA*, pages 90–100, February 1999.

[Rap96]   T. S. Rappaport. *Wireless Communications, Principles and Practice*. Prentice Hall, Upper Saddle River, New Jersey, 1996.

[TBTG01] M. Takai, R. Bagrodia, K. Tang, and M. Gerla. Efficient wireless network simulations with detailed propagation models. *Wireless Networks*, 7(3):297–305, May 2001.

[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.

[YLA02]   W. H. Yuen, H. Lee, and T. Andersen. A simple and effective cross layer networking system for mobile ad hoc networks. In *13th IEEE Int'l Symposium on Personal, Indoor and Mobile Radio Communications*, September 2002.

[ZHKS04] Gang Zhou, Tian He, Sudha Krishnamurthy, and John A. Stankovic. Impact of radio irregularity on wireless sensor networks. In *Proceedings of the 2004 International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 125–138, Boston, MA, June 2004. ACM Press.