

Experimental Evaluation of Wireless Simulation Assumptions

David Kotz
Dept. of Computer Science
Dartmouth College
Hanover, NH 03755, USA
dfk@cs.dartmouth.edu

Calvin Newport
Dept. of Computer Science
Dartmouth College
Hanover, NH 03755, USA
Calvin.Newport@alum.dartmouth.org

Robert S. Gray
Thayer School of Engineering
Dartmouth College
Hanover, NH 03755, USA
rgray@ists.dartmouth.edu

Jason Liu
Dept. of Mathematical and
Computer Sciences
Colorado School of Mines
Golden, CO 80401, USA
jasonliu@mines.edu

Yougu Yuan
Dept. of Computer Science
Dartmouth College
Hanover, NH 03755, USA
yuanyg@cs.dartmouth.edu

Chip Elliott
BBN Systems and
Technologies
Cambridge, MA 02138, USA
celliott@bbn.com

ABSTRACT

All analytical and simulation research on ad hoc wireless networks must necessarily model radio propagation using simplifying assumptions. We provide a comprehensive review of six assumptions that are still part of many ad hoc network simulation studies, despite increasing awareness of the need to represent more realistic features, including hills, obstacles, link asymmetries, and unpredictable fading. 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.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks; I.6 [Computing Methodologies]: Simulation and Modeling

General Terms

Measurement

Keywords

Wireless network, Wi-Fi, 802.11, ad hoc network, MANET, mobile computing, network simulation, experiment, measurement.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MSWiM'04, October 4-6, 2004, Venezia, Italy.

Copyright 2004 ACM 1-58113-953-5/04/0010 ...\$5.00.

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. Many such simulations may be based on overly simplistic assumptions, however; 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" [10]. It then proceeded to survey 2200 published network simulation results to point out systemic flaws.

Although 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, there are real risks to protocol designs based on overly simple models of radio propagation. As the nodes move in an ad hoc network, the connectivity graph changes over time, but these changes depend significantly on variations in the range due to antenna differences, elevation (extending range), and obstacles (restricting range).

We recognize that the MANET research community is increasingly aware of the limitations of its simplifying assumptions. Our goal in this paper is to make a constructive contribution to the MANET community by a) clearly identifying these assumptions and quantitatively demonstrating their weaknesses, 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.

Due to limited space, this paper presents only the main points. We present all of the details in an extended Technical Report [6].

2. MODELS USED IN RESEARCH

The simplest radio models are based on distance across flat terrain; radio communications are received perfectly within some circular "range" and not at all outside of that range. Real radios, including those used in the popular Berkeley Motes, demonstrate a strikingly non-uniform non-circular behavior [3, 13]. With real radios, spatial or temporal signal fluctuations may cause rapid changes

in network connectivity. Many algorithms and protocols may perform much more poorly under such dynamic conditions; some may even fail to converge and thus fail to work. More realistic models take into account antenna height and orientation, terrain and obstacles, surface reflection and absorption, and so forth. Simple radio models fail to explore these critical realities that can dramatically affect performance and correctness [3, 2, 13].

We surveyed a set of MobiCom and MobiHoc proceedings from 1995 through 2003. 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*. Even in the best years, the Simple and Flat-Earth papers significantly outnumbered the Good papers [6].

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 horizontal distance of node *B*.

Simple models are, almost without exception, ns-2 models using the CMU 802.11 radio model [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.

The two-ray ground-reflection model considers both the direct and ground-reflected propagation paths between transmitter and receiver; it is better, but not well suited to most MANET simulations.

More recently, ns-2 added a third, “shadowing” model [7], to account for indoor obstructions and outdoor shadowing via a probabilistic model [1]. The problem is that this model does not consider correlations: a real shadowing effect has strong correlations between two locations that are close to each other.

Although there are better radio models for ns-2 and other simulators (OpNet is one commercial example), most of the research literature uses ns-2 with simple models.

Good models have 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. 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?).

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 simple axioms to model their combined behavior. We take this network-layer point of view through the remainder of the paper. In the next two sections we show that 1) simple 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.

3. AXIOMS AND REALITY

We see six axioms in common use:

- 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.**

To study these axioms and their impact on simulation studies, we use data collected from a large MANET experiment in which 33 laptops with WiFi network cards roamed a field for over an hour while exchanging broadcast beacons and operating different ad hoc routing protocols [4]. Although our experiment represents just one environment, it serves our purpose to demonstrate that the axioms are untrue even in a simple environment, and that fairly sophisticated simulation models were necessary for reasonable accuracy. Indeed, our work represents the first detailed study of wireless assumptions that uses large-scale experimental data and corresponding simulation results for the same routing protocols.

We provide a full description of the experimental conditions and the data collected in a companion paper [4] and our Technical Report [6]. Briefly, each laptop periodically recorded its position (latitude, longitude and altitude) according to an attached GPS device, and *broadcast* a beacon containing its current position (as well as the last known positions of the other laptops). Every laptop receiving a beacon sent a *unicast acknowledgment* to the beacon sender via UDP. We use extensive log files, containing positions and data about beacons sent and received, to examine the axioms.

Axiom 0: The world is flat.

Common stochastic radio propagation models assume a flat earth, and yet clearly the Earth is not flat. We need no data to “disprove” this axiom. 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 [2]. Furthermore, it is not uncommon to see two nodes in a multi-story building deployed at the same *x, y* location, but on different floors.

Axiom 1: A radio’s transmission area is circular.

Axiom 2: All radios have equal range.

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. For example, 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. Figure 1 shows how the beacon-reception probability varied with angle. In each bucket, the reception probability is the fraction of beacons received out of beacons sent to nodes within that range of angles, regardless of distance.

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

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. Since *A* and *B* may be moving, it is important to consider symmetry over a brief time period

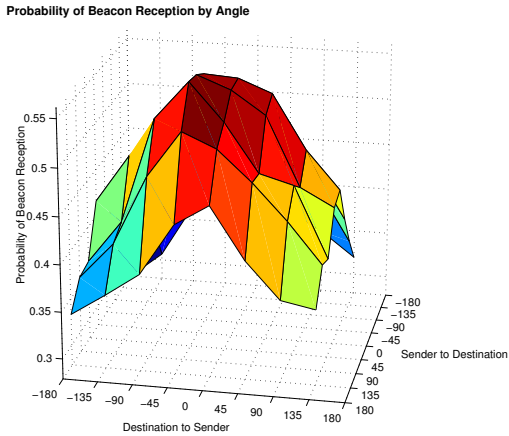


Figure 1: The probability of beacon reception as a function of the two angles, the angle between the sender’s antenna orientation and the receiver’s location, and the angle between the receiver’s antenna orientation and the sender’s location. In this plot, we divide the angles into buckets of 45 degrees each.

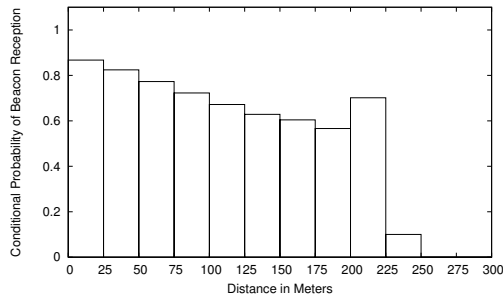


Figure 2: The conditional probability of symmetric beacon reception as it varied with the distance between two nodes.

before A and B have not moved apart. Our nodes sent a beacon every three seconds; thus whenever a node B received a beacon from node A , we checked to see whether B ’s beacon acknowledgement was also received by node A ; if so, we said the link was symmetric.

Figure 2 shows the conditional probability of symmetric beacon reception. The probability was never much more than 0.8, most likely due to MAC-layer collisions between beacons. At higher distances, a lower reception probability leads to a lower joint probability (of a beacon arriving from A to B and then another from B to A) and thus a lower conditional probability.

Significant asymmetry can arise from differences in transmission power (purposeful or otherwise). Marina and Das, for example, consider three different variable transmission-power models; they determine that using the resulting unidirectional links offers little routing benefit, and that those links should be removed from consideration when constructing routes [9].

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 at Figure 3, we see that the beacon-reception probability does indeed fade with the distance between the sender and

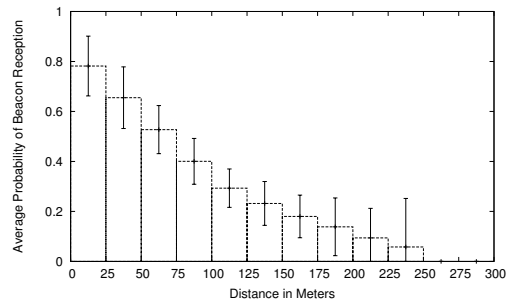


Figure 3: The average and standard deviation of reception probability across all nodes.

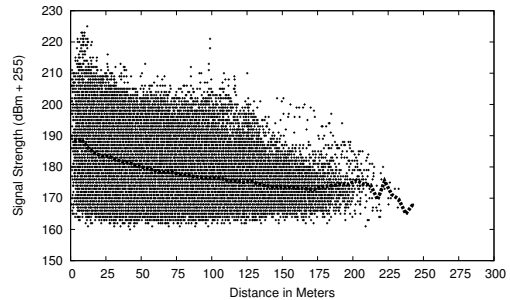


Figure 4: Signal strength and distance were poorly correlated. We show the mean signal strength as a heavy dotted line.

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 often unused. More sophisticated models do exist, particularly those developed by the GloMoSim project (and commercialized by QualNet), and can be used to explore how sophisticated channel models affect simulation outcomes.

Axiom 5: Signal strength is a simple function of distance.

Rappaport [11] 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.

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. We found that the power-law model was a good fit for the mean beacon signal strength observed during the experiment as a function of distance [6], validating Rappaport’s observation. When we turn our attention to the signal strength of individual beacons, however, as shown in Figure 4, there clearly is no simple (non-probabilistic) function that will adequately predict the signal strength of an individual beacon based on distance alone.

4. 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. We thus quantify the impact of the axioms on the simulated behavior of routing protocols.

While others have used simulation to explore the impact of different radio propagation models [12, 13], we used the same protocol implementation in both the experiment and our simulator [8], and used a large number of nodes in the outdoor experiment [4]. By extending the simulator to read the node mobility and application-level data logs generated by the real experiment, we were able to reproduce the same traffic in simulation as in the real experiment, and compare the results directly.

Our best model. The signal propagation model is a stochastic model that captures radio signal attenuation as a combination of small-scale fading and large-scale fading. 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. For the propagation model, we chose 2.8 as the distance path-loss exponent and 6 dB as the shadow fading 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.) To this we add a detailed 802.11 model, with parameters chosen to match the settings of our real wireless cards.

Note that our “best” model assumes some of the same axioms (namely flat earth, omni-directional radio propagation length, and symmetry) we discount in the preceding section! Because this model sufficiently matches the experiment results, it is clear that the stochastic model with proper parameters can provide a close approximation to the radio propagation environment of the environment we used for our outdoor experiment. Nonetheless, this model is sufficient for the purposes of this paper, because we can still demonstrate how the other axioms may affect performance.

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. In this study, we chose to use the two-ray ground reflection model since its signal travel distance matches observations from the real experiment. This weaker model assumes Axiom 4: “If I can hear you at all, I can hear you perfectly.” Without variations in the channel, all signals travel the same distance, and successful reception is subject only to the state of interference at the receiver.

Finally, we consider a third model that further weakens the simulator by assuming that the radio propagation channel is *perfect*. 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.

The Results. Consider the reception ratio of the beacon messages, shown in Figure 5. Compared with the two simple models, our best model is a better fit for the real experiment results. The sharp cliff in the beacon reception ratio curve, seen in both of the simpler models, simply does not exist in reality; the better propagation model was a critical feature.

Now consider the effect of different simulation models on the overall performance of the AODV routing algorithm. Figure 6 shows the packet delivery ratio, as we varied the application traffic intensity by adjusting the average packet inter-arrival time at each node. The real experiment’s result is represented by a single point (42.3% rather than 46.8% in simulation with our best model). We saw similar behavior for the APRL and ODMRP routing algo-

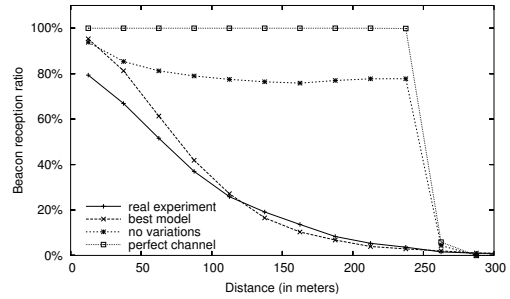


Figure 5: The beacon reception ratio at different distances between the sender and the receiver. The probability for each distance bucket is plotted as a point at the midpoint of its bucket.

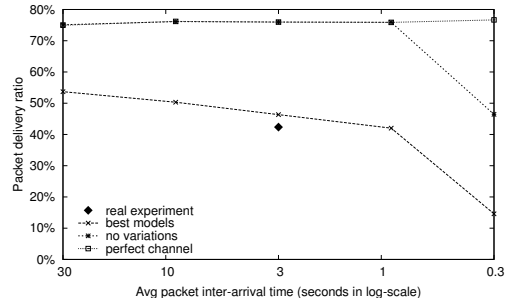


Figure 6: Packet delivery ratios for AODV.

ritms [6]; in both cases, the best simulation came even closer to the experimental delivery ratio. Clearly, routing performance was significantly exaggerated by the simpler models. 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 routes to their intended destinations.

Earlier, others studied the effect of detail in radio propagation models on wireless network simulations [5]. In particular, they found that simple radio models can sometimes be effective in cases where either the applications or the metrics describing the applications are insensitive to the abstractions in lower layers. They use radio-based localization and robot following problems as case studies to support this claim.

5. CONCLUSIONS, RECOMMENDATIONS

Hundreds of research papers have presented simulation results for mobile ad hoc networks, but the great majority of these papers rely on overly simplistic assumptions of how radios work. Many widely used radio 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 are much more complex than the simple models used by most researchers [10], and that these complexities have a significant impact on the behavior of MANET protocols and algorithms [3, 13]. In this paper, we enumerate the set of common assumptions used in MANET research, and present an experiment that strongly contradicts these “axioms,” and demonstrate the impact of these assumptions on simulation results.

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

1. Choose the target environment carefully, clearly list the assumptions about that environment, choose a simulator that supports 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. The model should exhibit important properties of the wireless environment under study, including asymmetric links (e.g., where *A* can hear *B* but not vice versa) and some time-varying fluctuations in link quality.
3. Simulation should explore a range of model parameters since the effect of these parameters is not uniform across protocols.
4. Consider three-dimensional terrain, with moderate hills and valleys, and corresponding radio propagation effects.
5. Use real data as input to simulators, where possible. For example, using our data as a “snapshot” of a realistic ad hoc wireless network, researchers should verify whether their protocols form networks as expected.

...for simulation and model designers:

1. Allow protocol designers to run the same code in the simulator as they do in a real system, making it easier to compare experimental and simulation results. We did this in SWAN [8].
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 the assumptions of lower layers. In our experimental results, we found that the success of a transmission between radios depends on many factors that cannot be accurately modeled, predicted or detected at the speed necessary to make per-packet routing decisions.
2. Develop protocols that adapt to environmental conditions.
3. Explore the costs and benefits of control traffic; consider carefully whether extra control traffic is worth the interference price.

Availability. Our simulator and our dataset is available to the research community. 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

We are extremely grateful to the many people that helped to code our simulator and ad hoc routing protocols, who carried laptops around the field, and who commented on early versions of this paper, including the anonymous reviewers.

This project was supported in part by a grant from the Cisco Systems University Research Program, the Dartmouth Center for Mobile Computing, DARPA (contract N66001-96-C-8530), the Department of Justice (contract 2000-CX-K001), the Department of Defense (MURI AFOSR contract F49620-97-1-03821), and the Office for Domestic Preparedness, U.S. Department of Homeland Security (Award No. 2000-DT-CX-K001). Additional funding provided by the Dartmouth College in the form of a Presidential Scholar Undergraduate Research Grant and a Richter Senior Honors Thesis Research Grant. Points of view in this document are those of the author(s) and do not necessarily represent the official position of any of the sponsors.

6. REFERENCES

- [1] K. Fall and K. Varadhan. *The ns Manual*, April 14 2002. www.isi.edu/nsnam/ns/ns-documentation.html.
- [2] G. Gaertner and V. Cahill. Understanding link quality in 802.11 mobile ad hoc networks. *IEEE Internet Computing*, Jan/Feb 2004.
- [3] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. 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.
- [4] R. S. Gray, D. Kotz, C. Newport, N. Dubrovsky, A. Fiske, J. Liu, C. Masone, S. McGrath, and Y. Yuan. Outdoor experimental comparison of four ad hoc routing algorithms. In *Proceedings of the MSWiM 2004*, Oct. 2004.
- [5] J. Heidemann, N. Bulusu, J. Elson, C. Intanagonwivat, K. chan Lan, Y. Xu, W. Ye, D. Estrin, and R. Govindan. Effects of detail in wireless network simulation. In *Proceedings of the SCS Multiconference on Distributed Simulation*, pages 3–11, Phoenix, Arizona, January 2001. Society for Computer Simulation.
- [6] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. Technical Report TR2004-507, Dept. of Computer Science, Dartmouth College, June 2004.
- [7] W. C. Y. Lee. *Mobile communications engineering*. McGraw-Hill, 1982.
- [8] J. Liu, Y. Yuan, D. M. Nicol, R. S. Gray, C. C. Newport, D. Kotz, and L. F. Perrone. Simulation validation using direct execution of wireless ad-hoc routing protocols. In *PADS '04*, pages 7–16, May 2004.
- [9] M. K. Marina and S. R. Das. Routing performance in the presence of unidirectional links in multihop wireless networks. In *MobiHoc '02*, pages 12–23, Lausanne, Switzerland, June 2002.
- [10] K. Pawlikowski, H.-D. Jeong, and J.-S. Lee. On credibility of simulation studies of telecommunication networks. *IEEE Communications*, 40(1):132–139, January 2002.
- [11] T. S. Rappaport. *Wireless Communications, Principles and Practice*. Prentice Hall, 1996.
- [12] M. Takai, J. Martin, and R. Bagrodia. Effects of wireless physical layer modeling in mobile ad hoc networks. In *MobiHoc '01*, pages 87–94, Oct. 2001.
- [13] G. Zhou, T. He, S. Krishnamurthy, and J. A. Stankovic. Impact of radio irregularity on wireless sensor networks. In *MobiSys '04*, pages 125–138, June 2004.