On the Credibility of Manet Simulations

Todd R. Andel and Alec Yasinsac
Florida State University

Simulation is useful for evaluating protocol performance and operation. However, the lack of rigor with which it's applied threatens the credibility of the published research within the manet research community.

Mobile ad hoc networks (manets) allow rapid deployment because they don't depend on a fixed infrastructure. Manet nodes can participate as the source, the destination, or an intermediate router. This flexibility is attractive for military applications, disaster-response situations, and academic environments where fixed networking infrastructures might not be available.

Simulation has proven to be a valuable tool in many areas where analytical methods aren't applicable and experimentation isn't feasible. Researchers generally use simulation to analyze system performance prior to physical design or to compare multiple alternatives over a wide range of conditions. Unfortunately, errors in simulation models or improper data analysis often produce incorrect or misleading results.

The mainstream approach in the manet research community follows the development, simulation, and publish process. Manet publications typically include performance simulations and commonly compare routing protocols.

Simulation is a powerful tool, but it's fraught with potential pitfalls. We question this approach's validity and show how it can systematically produce misleading results. Although the problems aren't unique to manet simulations, we focus on issues affecting the manet research community. Rather than condemning simulation-based research, or simulation itself, we provide a cautionary tale for all who rely exclusively on simulation for validating technical concepts.

SIMULATION CHALLENGES

Simulation provides an attractive method for evaluating the performance of mobile ad hoc routing protocols because designers can use it to evaluate theoretical systems. However, if the simulation doesn't reflect an important aspect of reality, it can't give insight into the operating characteristics of the system the developers are studying. Generalization and lack of rigor can lead to inaccurate data, which can result in wrong conclusions or inappropriate implementation decisions.

Consider, for example, the inappropriate exercise of feature isolation. Investigators can use simulation to isolate target features by varying only those features, while holding all other parameters constant. Many manet simulations compare protocols by varying only the protocol in each simulation run.

Unfortunately, the constant factors can (passively) inject inaccuracies into the simulation. This type of unintended side effect can occur when, for example, investigators run a simulation with protocol A and subsequently run the identical simulation with protocol B. Intuitively, with all other settings the same, the protocol that performs the best head-to-head should have the better performance. However, this isn't necessarily true.

In reality, static features can dominate performance. For example, protocol A might outperform B at 10-megabits-per-second (mbps) bandwidth, but the converse might be true at 100 mbps. Worse yet, even if the simulation results reflect these two situations, they'd say little about performance in a 50-mbps or 500-mbps network. Omitting such intermediate-valued tests can lead to false conclusions. Determining which features to vary, how much to vary them, and in what combinations with which other features to simulate them is difficult. Often, the volume of interaction possibilities precludes thorough simulation, leading to less than rigorous valida-

Published by the IEEE Computer Society

tion. Consequently, the investigator can't attribute output differences solely to the varied parameter.

Network simulation model inconsistencies

Most research papers on manet routing protocols include a simulation to show a proposed solution's performance. Researchers typically conduct the simulation using only one simulation package. Network simulation packages are complex, and a researcher might have sufficient experience or research time for just one such package, whether it's commercial, open source, or independently developed.

David Cavin, Yoav Sasson, and Andre Schiper compare three commonly used network simulation software packages—Opnet, Network Simulator (NS-2), and Global Mobile Information Systems Simulation Library (GloMoSim)—by implementing a simple flooding protocol in each package.[1] If each simulation package reflects reality, the simulation results should agree. However, the results show large differences between the simulators. For example, with a simple well-understood flooding protocol, GloMoSim shows 100 percent delivery at a very low power range, while NS-2 and Opnet reflect much lower success rates at the same power setting. Conversely, Opnet and GloMoSim correlate much more closely toward 100 percent success as power increases, whereas delivery in NS-2 seems to flatten out at about 50 percent at the highest power threshold.

Only one, if any, of these simulations can be right—because they give three significantly distinct results, if one is right, the other two must be wrong. Worse yet, this divergence suggests that all three might be wrong.

Unfortunately, making all of the simulators agree for a given scenario won't resolve this problem. Protocol simulation aims to produce results that represent real implementations. If a simulator is valid, real-life performance should correlate with the simulated performance. Agreement among multiple simulators doesn't lend credibility unless investigators independently validate the simulators and corroborate their application for each instance. Unfortunately, little in the literature documents rigorous validation for any simulation package routinely used in manet research. Thus, investigators should corroborate any simulation used for manet research against experimentation with a real implementation or through analytical results.

Protocol stack interactions. Figure 1 represents the manet protocol stack commonly used in simulation. In this view, the network layer is typically the layer under test because routing protocol development, testing, and comparison occur at this level.

Researchers typically combine the levels above the network layer into a single application layer and model application traffic as constant-bit-rate traffic generators, with a set packet size and interarrival time. Traffic-generation levels profoundly affect the simulation's outcome

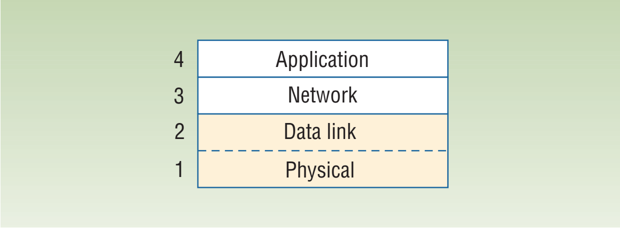

Figure 1. Modified protocol stack. Most testing in manet simulations takes place at the network layer, with the levels above the network layer combined into a single application layer.

by determining how a given routing protocol will operate under various traffic loads (that is, heavy congestion).

In 2004, Himabindu Pucha, Saumitra M. Das, and Y. Charlie Hu suggested that increasing the number of traffic connections in a single node reduces the delivery rate for networks with the same average traffic load.[2] They reported that transmitting sources in NS-2 manet simulations have on average one to two simultaneous connections. They attributed the tendency for higher connection rates to running multiple applications and supporting network services such as authentication. As they increased the number of connections per source, they simultaneously decreased each connection's traffic rate to maintain the same average traffic-generation rate. Their results showed that, for the same generation rate, packet-delivery rates decreased as the number of connections increased.

Their study questioned prior results based on simple traffic patterns that don't support many simultaneous connections per source. Unfortunately, traffic-generation parameters depend on specific applications and have a wide range of possibilities, such as packet size, variable versus constant bit-rate generation, and application or usage profiles. Without a universal or driving application, the best you can do is properly document the simulated settings so independent researchers can duplicate the results.

The data link layer—which includes the medium access control, or MAC, sublayer—and physical layer are highly correlated, and network simulation packages commonly treat them as a combined layer. For instance, Opnet, NS-2, and GloMoSim can simulate the combined IEEE 802.11 physical and data link.[1,3]

Physical layer interactions are an important part of manet simulation configurations, and investigators should recognize and document them as a controlled variable. In their study of the effects of different wireless physical layer models,[3] Mineo Takai, Jay Martin, and Rajive Bagrodia listed the fading, path loss, and signal-reception models used by Opnet (version not stated), NS-2 v.2.1b, and GloMoSim v.2.02. GloMoSim included both Rayleigh (used to model non-line-of-sight) and Ricean (used to model line-of-sight) fading models, whereas NS-2 and Opnet offered neither choice.

GloMoSim and NS-2 included two-ray path-loss models, which account for reflected ground transmission interference, and free-space path-loss models, which assume no interference and perfect channel conditions; Opnet included only free-space models. Opnet calculated signal reception based on bit error rate (BER), NS-2 used signal-to-noise ratio threshold (SNRT), and GloMoSim included both options. Free-space and line-of-sight modeling consider ideal conditions that might not be applicable in actual implementations.

Takai, Martin, and Bagrodia performed experiments to show the results of various physical layer factor settings within GloMoSim v.2.02 for the Ad Hoc On-Demand Distance Vector (AODV) and Dynamic Source Routing (DSR) protocols.[3] Their results show that fading, path-loss, and signal-reception factor settings significantly affect simulation results in GloMoSim. When the path-loss model factor is set to two-ray, the results significantly differ depending on the signal reception model (BER or SNRT) and fading model (none, Ricean, or Rayleigh). In this scenario, AODV and DSR packet-delivery rate and end-to-end delay vary significantly for the various combinations.

Omitting detail or oversimplifying the model can lead to ambiguous or erroneous outcomes.

More disturbing is that the relative rankings between AODV and DSR change. With no fading model and the Ricean model, AODV's end-to-end delay was much smaller than DSR's. With Rayleigh fading, DSR had the smallest delays. When Takai, Martin, and Bagrodia performed the same experiment with the free-space path-loss model, output varied little between all signal-reception and fading-model combinations for AODV and DSR.

These two simple examples show that differences in comparative analysis between routing protocols can be due to underlying (and possibly undocumented) parameter settings and not the protocols being compared.

MAC settings, which belong to the data link layer, include the ability to enable or disable the request-to-send/clear-to-send and to use the Address Resolution Protocol (ARP).[1,3] With RTS/CTS enabled, simulations can model the 802.11 standard's distributed coordination function, which has an obvious impact on packet collision. Furthermore, Luiz Felipe Perrone, Yougu Yuan, and David Nicol discuss how simple modeling of ARP network-to-physical address translations can impact end-to-end delay in manet simulations.[4] In their study, including ARP increased end-to-end delay by as much as 16 percent.

Obviously, settings at all protocol stack layers significantly impact simulation outcome. For the simulation to be constructive, investigators must clearly understand and document their setting choices.

Effects of detail. It's impossible to simulate all aspects of a manet. Numerous tradeoffs are related to the effects of abstraction or level of detail reflected. Excessive detail can cause problems in simulation development and execution. As a simulation's detail increases, development time and simulation execution time also increase. Additionally, bugs that could affect results are less likely to be detected as the modeling detail increases.

Unfortunately, omitting detail or oversimplifying the model can lead to ambiguous or erroneous outcomes. John Heidemann and his colleagues compared energy consumption among four manet routing protocols: AODV, DSR, Destination-Sequenced Distance-Vector (DSDV), and the Temporally Ordered Routing Algorithm (TORA).[5] When they ignored idle-time energy consumption, the on-demand AODV and DSR protocols outperformed the table-driven DSDV and TORA protocols. However, when they considered idle-time energy consumption, the results were approximately the same for all routing protocols.

The problem lies in interpreting these inconsistent results. In the Heidemann case, using the results that include the idle-time energy consumption would seem to provide the most realistic representation for a protocol's real-world energy requirements. The energy model used in the simulation represented the WaveLAN wireless radio interface, which has idle-time energy requirements. Although the energy model's abstraction within the simulation can never give a completely equivalent result, it should provide a realistic outcome. This reinforces the call for validating a simulation model against a real-world implementation because using the incorrect model for the given environment would produce incorrect results.

Model validation. Kwan-Wu Chin and colleagues' real-world manet implementation study brings into question the results of manet studies based solely on simulation.[6] They tested both AODV and DSDV on a small four-node network using Lucent WaveLAN IEEE 802.11b-compliant network cards. Their main goal wasn't to determine these routing protocols' performance, but to answer a more basic question: Do manet routing protocols work? Their results showed that both protocols had unreliable routes for communication outside a source's one-hop distance. Further investigation revealed that these protocols improperly chose neighbors at the transmission footprint's edges, making subsequent communication rely on a direct route to each node instead of using an intermediate note as a relaying hop. Consequently, two routing protocols deemed feasible via simulation didn't work properly as designed when tested in an actual implementation.

Properly validating simulation models against the intended or real-world implementation and environment can mitigate many of the problems of simulation package differences, incorrect parameter settings, and improper level of detail. As we discussed earlier, a simulation model must make abstractions from the concrete representation. But how do you know the abstraction level is correct? The basic problem here is that we can't know what we don't know—that is, we can never be sure that we've accounted for all aspects that could affect a simulation model's ability to provide meaningful results. For example, consider an AODV simulation performed and validated against a baseline real-world case of 50 laptops roaming an open field in flat terrain with humid conditions. Will the same simulation model provide realistic results for 100 laptops in hilly terrain in a dry desert?

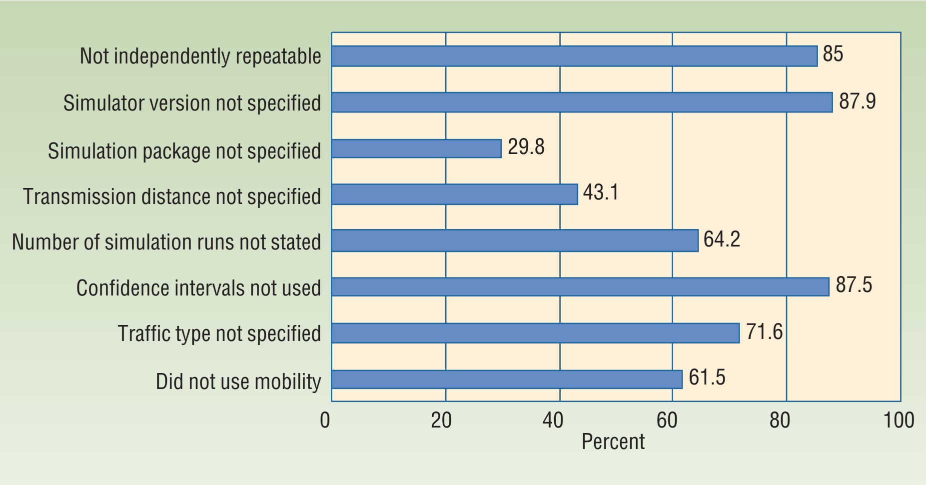

Figure 2. Manet simulation issues. The omissions cited affect the research's integrity and credibility within the manet research community.

David Kotz and colleagues demonstrated how to validate simulation models against an actual manet protocol implementation.[7] They compared a simple beaconing protocol implemented on 33 laptops against the same protocol simulated in the simulator for wireless ad hoc networks (SWAN) package for three radio models. In the real-world implementation, they roamed an open field, collecting location and application data. They transformed this information into simulation input for traffic and mobility to ensure an equal comparison when running the protocol in SWAN. Their results show that their "best model" most closely represents the real experiment because it accounts for radio signal shadowing and fading. The "perfect channel" (similar to NS-2 free space) and "no variation" (similar to NS-2 two-ray) radio models didn't produce real-world results and would provide unrealistic conclusions if used for simulations in this environment.

Often, a newly developed protocol doesn't have an actual implementation or testbed to serve as a baseline. In this case, investigators can validate a baseline against the protocol specifications or mathematical calculations. However, such a validation's reliability would obviously be lower than validation against an actual implementation because including environmental conditions and channel contention issues would be difficult.

Another option is to use area experts to validate simulation models. Josh Broch and his colleagues used experts to validate radio propagation and 802.11 MAC implementations during the development of their NS-2 manet extensions.[8] They also asked the original protocol authors to validate their routing protocols. Using experts can increase confidence when real-world implementations are unavailable. However, because this process validates each model of the entire simulation (that is, physical layer, MAC layer, and given protocol) separately, investigators can't account for interactions between various aspects of the simulation.

Improper simulation practices

In addition to parameter settings within the protocol stack, poor experimental procedures can bias manet simulation studies. To increase such studies' credibility, investigators should follow a rigorous scientific procedure that provides for repeatability and statistical validity and uses appropriately justified assumptions.

Repeatable simulation. Researchers need documentation to understand published results. It's also vital to the credibility of manet studies. Documentation lets peers appropriately review and understand the simulation environment, and full disclosure of all underlying parameters allows fair protocol comparisons. Independent researchers must be able to repeat simulation studies to ensure their credibility.

Unfortunately, lack of documentation is a widespread problem in the manet research community. Stuart Kurkowski, Tracy Camp, and Michael Colagrosso studied 114 peer-reviewed manet research papers published between 2000 and 2005 in a highly respected venue.[9] The results are startling.

Figure 2 summarizes some of the issues raised in the study. The percentages aren't a direct comparison from the same baseline; rather, the authors drew them from different populations with the intent of visualizing some significant issues in the manet research community.[9]

The survey concluded that 85 percent of the research papers aren't independently repeatable because of the lack of documentation. Missing documentation included simple environment issues. For example, of the 58 papers that specified a publicly available simulation package, 87.9 percent didn't list the version. This omission might seem innocuous, but investigators can't repeat a simulation without this information. The survey also reported that 29.8 percent of the papers didn't even identify the simulation package that was used.

Additionally, the papers often omitted simulation input parameters, such as transmission range and traffic type. For example, the survey noted that 43.1 percent of the simulation papers didn't specify the transmission distance used. This basic information ensures the research's integrity and credibility within the manet community.

Statistical validity. Statistical validity is another important aspect in simulation. Because simulation is inherently imprecise, a single simulation run is rarely credible. Rather, multiple simulation runs—for example, with different pseudorandom number generator (PRNG) seeds evaluated for convergence, deviation, modal values, and other statistically significant metrics—provide credibility and insight. In contrast, insufficient statistical analysis of independent simulation runs and improper data collection techniques can produce ambiguous or inaccurate conclusions. For example, if investigators run a simulation 30 times, varying only the PRNG seed with no clear convergence and wide result deviation, any single run is meaningless, and the simulation itself must be analyzed for flaws. Moreover, such a result can mean that you can't accurately estimate the target property with simulation (that is, the simulation is invalid).

Insufficient statistical analysis of independent simulation runs and improper data collection techniques can produce ambiguous or inaccurate conclusions.

Raj Jain provides numerous techniques to guide computer simulations for performance studies.[10] A crucial aspect of any study is proper simulation output analysis. For example, multiple simulation runs using different PRNG seeds ensure that investigators can use independent results to calculate average results, and they can use the central limit theorem to calculate a confidence interval. As the number of trials increases, variance should decrease, allowing better mean value estimation and providing a more precise range or deviation. In addition to providing an expected range, confidence intervals can help determine if two data sets are statistically equivalent. If the mean value for a simulated measurement falls within the confidence interval of the same measurement from the real experiment and vice versa, the simulation value for that data point would be statistically equivalent to the real experiment.

The central limit theorem can also determine the number of simulation runs required. An investigator starts by performing a small number of independent runs. Next, the theorem uses the resulting mean and standard deviation to determine the number of simulation runs necessary, as given by $n = ((100 \times z \times s)/r \times \bar{x}))^2$, where n is the number of replications, $\bar{x}$ is the sample mean, r is +/– precision level, z is the normal variate (that is, 1.645 constant for 90 percent confidence interval), and s is the standard deviation.[10]

Data collection can also be the source of incorrect simulation analysis. Investigators must collect data when a simulation is in steady state to protect against transient, or start-up, data injecting bias into the results. Research estimates that the average relative error for end-to-end delay can be as high as 30 percent when data analysis includes transient values.[4] You can remove transients by properly initializing state variables, such as preloading queues and buffers with application traffic, and presetting routing table assignments before executing the simulation. If presetting variables isn't feasible, performing initial data deletion can remove simulation initialization bias. Jain provides various statistical methods, such as tracking data rate of change, to identify the length of the initialization period from which you must discard collections.[10]

Unfortunately, manet community researchers often ignore or neglect these issues. Kurkowski, Camp, and Colagrosso's survey showed that 64.2 percent of the analyzed papers didn't state the number of simulation runs, 87.5 percent didn't include confidence intervals, 93 percent didn't remark on transient removal, and none addressed using a PRNG.[9] All of these omissions profoundly affect the reported results' statistical validity.

Precision. Unlike mathematical proof systems, which measure results by orders of magnitude, simulation models can show percentage improvements between manet protocols. However, simulation is naturally imprecise and is subject to errors injected by inaccurate parameters or false assumptions. Simulation assumptions always affect research outcomes.

In most simulations, precision is difficult to attain. Manet research uniformly makes many imprecise assumptions. For example:

- Transmission range is a critical factor in many manet protocols, but its characteristics are not precisely defined. Rather, investigators generally represent transmission distance as a circle's radius.
- Researchers commonly model node distribution as uniform or random. In reality, roads, trees, water, and other obstacles affect node distribution.
- Interference models are typically based on SNRs or BERs. This neglects interference based on increasing traffic or unpredictable background noise.
- Researchers typically assume that node communication is bidirectional. However, unlike wired implementations, wireless communication doesn't guarantee

Table 1. Recommendations to improve simulation credibility.

Problem	Solution
Lack of independent repeatability	Properly document all settings. Publication venues have limited space, so typically include only major settings (such as transmission distance and bit rate). Provide all settings as external references to research Web pages, which should include freely available code/models and applicable data sets.
Lack of statistical validity	Determine the number of required independent runs. Address sources of randomness (such as pseudorandom number generators) to ensure simulation run independence. Collect data only after deleting transient values or eliminating it by preloading routing tables and traffic queues.
Use of inappropriate radio models	Free-space radio models are sufficient during early model development, but two-ray and shadow models provide a more realistic environment during data collection and analysis. Tune settings against an actual implementation when available. Improve radio model abstractions as more implementations and experimental manet testbeds become available.
Improper/nonexistent validation	Validate the complete simulation (developed protocol, traffic, radio model, and scenario) against a real-world implementation. When this isn't possible (such as during early concept development), validate the simulation against analytical models or protocol specifications. The latter will be less precise, but you can further refine it as implementations are realized.
Unrealistic application traffic	Simple constant-bit-rate traffic might be unrealistic. Base traffic generation on intended applications.
Improper precision	Use manet simulations to provide proof of concept and general performance characteristics, not to directly compare multiple protocols against one another.
Lack of sensitivity analysis	Sensitivity analysis can identify a chosen factor's significance (parameter settings that change in a study). For example, if you're testing two routing protocols (such as Ad Hoc On-Demand Distance Vector and Dynamic Source Routing) against three mobility speeds, using the analysis of variance (ANOVA) technique can determine if the output changes are due to the routing protocol, the mobility setting, both, or neither. Raj Jain's book[10] lists procedures for performing ANOVA calculations.

signal transmission, and reception distances are equivalent. Manet nodes might have different power reserves available for transmission.

- Researchers commonly model node mobility as random, but it rarely is. For instance, a group of soldiers will commonly follow a preplanned path, or at least travel in the same general direction. Individuals rarely travel in random directions, pause for random times, and then embark in a completely different direction. Instead, they usually follow some pattern.
- Simulations typically model a square or rectangular network area. Although convenient, this rarely reflects reality. When a node reaches the network edge, does it abruptly turn or continue moving and fall out of the network area? If the latter, the node moving outside the network area would still be in transmission and interference range to nodes on the simulation area's edges.

Investigators might not fully understand the model abstraction and intended operating environments. Although investigators can't eliminate many assumptions, they must understand their effects on the overall outcome. It would be improper to compare two protocols with constant parameter settings without fully understanding the true sources of performance differences.

RECOMMENDATIONS

Our observations and the issues we've raised might seem obvious or trivial. Unfortunately, however, much of the current published research in this area doesn't follow these basic principles. These issues aren't new. In fact, in 1999, the US Defense Advanced Research Projects Agency and the National Institute of Standards and Technology hosted a network simulation validation workshop to discuss similar problems.[11] Nevertheless, investigators aren't following many of the workshop recommendations.

Table 1 provides some recommendations for credible simulation. While not an all-inclusive list, it represents a solid proposal of practices to increase research credibility. Depending on the development stage, complexities, and available resources, incorporating all recommendations might not be possible. The greater the number of practices a study can adapt, the greater the study's scientific rigor and resulting higher credibility. These recommendations can be applied to any area that relies on simulation-based research.

Without proper documentation, the additional recommendations are a moot point. Our suggestions are similar to and echo recommendations made elsewhere.[1,2,4,6-11]

The manet community must take it upon itself to ensure the integrity and credibility of published research. The responsibility lies with the researcher publishing results, the conference and journal reviewers, and these publications' readers. It's time for a wake-up call to the research community: Police our science or lose our credibility. ■

Acknowledgments

We based this material on work supported in part by the US Army Research Laboratory and the US Army Research Office under grant number W91NF-04-1-0415. The views expressed in this article are those of the authors and don't reflect the official policy or position of the US Air Force, US Department of Defense, or US government.

References

1. D. Cavin, Y. Sasson, and A. Schiper, "On the Accuracy of Manet Simulators," *Proc. 2nd ACM Int'l Workshop Principles of Mobile Computing*, ACM Press, 2002, pp. 38-43.
2. H. Pucha, S. Das, and Y. Hu, "The Performance Impact of Traffic Patterns on Routing Protocols in Mobile Ad Hoc Networks," *Proc. 7th ACM Int'l Symp. Modeling, Analysis, and Simulation of Wireless and Mobile Systems*, ACM Press, 2004, pp. 211-219.
3. M. Takai, J. Martin, and R. Bagrodia, "Effects of Wireless Physical Layer Modeling in Mobile Ad Hoc Networks," *Proc. 2nd ACM Int'l Symp. Mobile Ad Hoc Networking and Computing*, 2001, ACM Press, pp. 87-94.
4. L.F. Perrone, Y. Yuan, and D.M. Nicol, "Modeling and Simulation Best Practices for Wireless Ad Hoc Networks," *Proc. 2003 Winter Simulation Conf.*, vol. 1, IEEE Press, 2003, pp. 685-693.
5. J. Heidemann et al., "Effects of Detail in Wireless Network Simulation," *Proc. Soc. for Computer Simulation Multiconf. Distributed Simulation*, Soc. for Computer Simulation, 2001, pp. 3-11.
6. K. Chin et al., "Implementation Experience with Manet Routing Protocols," *SIGMobile Mobile Computing Comm. Rev.*, vol. 32, no. 5, 2002, pp. 49-59.
7. D. Kotz et al., "Experimental Evaluation of Wireless Simulation Assumptions," *Proc. 7th ACM Int'l Symp. Modeling, Analysis, and Simulation of Wireless and Mobile Systems*, ACM Press, 2004, pp. 78-82.
8. J. Broch et al., "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," *Proc. 4th Ann. ACM Int'l Conf. Mobile Computing and Networking*, ACM Press, 1998, pp. 85-97.
9. S. Kurkowski, T. Camp, and M. Colagrosso, "Manet Simulation Studies: The Incredibles," *SIGMobile Mobile Computing Comm. Rev.*, vol. 9, no. 4, 2005, pp. 50-61.
10. R. Jain, *The Art of Computer Systems Performance Analysis*, John Wiley & Sons, 1991.
11. J. Heidemann, K. Mills, and S. Kumar, "Expanding Confidence in Network Simulations," *IEEE Network*, vol. 15, no. 5, 2001, pp. 58-63.

***Todd R. Andel** is a PhD student in the Department of Computer Science at Florida State University. His research interests include wireless security protocols, simulation, and formal methods. He received an MS in computer engineering from the Air Force Institute of Technology. Contact him at andel@cs.fsu.edu.*

***Alec Yasinsac** is an associate professor of computer science at Florida State University. His research interests include application software, managing and installing systems software on mainframe systems, and engineering tactical and permanent data networks. He received a PhD in computer science from the University of Virginia. He is a Senior Member of the IEEE and a member of the IEEE Computer Society and the ACM. Contact him at yasinsac@cs.fsu.edu.*