Simulations in Wireless Sensor and Ad Hoc Networks: Matching and Advancing Models, Metrics, and Solutions

Ivan Stojmenovic, University of Birmingham and University of Ottawa

ABSTRACT

The objective of this article is to give advice for carrying out a proper and effective simulation activity for protocol design. It challenges some of the existing criticisms of simulation practices that emphasized *validation* aspects. This article advocates the use of *simple* models, matching assumptions and metrics in the problem statement and simulation to provide a basic "proof of concept," and *comparison* with truly competing solutions, which is possible only after a thorough and *critical literature review*. Then the complexity of the models can be increased (one parameter at a time), revising the algorithms themselves by adapting them to new assumptions, metrics, and the corresponding simulation environment. Selected *independent variables* should explain performance under a wide range of scenarios.

INTRODUCTION

Several recent studies have underlined the lack of rigor in simulation practices that threatens the credibility of the published claims. While agreeing with some of the criticism, this article argues against some others and advocates for a different overall view. In short, existing criticism is concentrated on simulation as an independent task, and advises better reporting (repeatability), more reliability, more realism, and more scenarios. This article leverages the views of different communities and benefits all researchers by presenting a common ground.

Our starting position is that simulation is merely a supporting factor for new ideas. Better simulations may always be added later. The main contribution of an article may be a novel valuable protocol, novel concepts, or theoretical analysis, with simulation (if needed at all) only used to provide minimal but necessary support for the claims made, provide justification for later deeper simulation, emulation, or implementation, and/or indicate weaknesses to be addressed by further analysis and research.

A single article will never solve all the problems on the road toward creating applications. For example, after a thousand articles on routing in ad hoc and sensor networks, it still remains unclear which protocol will perform well under a wide range of scenarios. It is our view that each article should be judged on its overall contribution, including the assumptions used, theory developed, new algorithms introduced, protocol details, simulation results, and relevance to an ultimate goal of staying on a path toward creating applications.

We begin with a literature review of existing criticism for simulation practices, and then discuss what we believe are the main issues. We take a more general view of simulation as a support for new ideas and theories, providing a platform for their comparison with truly competing existing solutions, not merely for their validation. We also discuss the proper meaning of a proof of concept, selection of independent variables to clearly represent a variety of scenarios, and simplicity of overall design that should provide for understanding and explanation.

EXISTING CRITICISMS OF SIMULATION PRACTICES

Takai, Martin, and Bagrodia [1] observed that typical simulations focus on specific higher-layer protocols that are being proposed, and tend to ignore details of models at other layers, particularly the interaction with the physical layer. They show that signal reception, path loss, fading, interference and noise computation, and preamble length have significant impact on performance. Furthermore, they reveal that different simulation tools yield quite different results even when they are configured with the same set of protocols, and such differences are mostly derived from different assumptions made at the physical layer. Some of the modeling factors of the physical layer can change the simulation results significantly, and can even change the relative ranking of routing protocols for specific scenarios [1]. Their findings are mainly based on NS-2 and GloMoSim simulators and Ad Hoc On Demand Vector (AODV) and DSR routing protocols for ad hoc networks. Therefore, when network layer protocols are designed based on a simple model such as an unit disk graph (UDG), and experimented on a simulator based on a

0163-6804/08/$25.00 © 2008 IEEE

(non-unique) realistic physical layer, the result is a list of suboptimal solutions whose performance cannot be fully explained because of too many variables being involved. For completeness, two nodes can communicate in UDG if and only if the distance between them is at most R, where R is the transmission radius, equal for all nodes.

Perrone, Yuan, and Nicol [2] state that countless experimental studies already published cannot be replicated because they do not fully report the conditions in which they were carried out. This situation is especially troublesome when the number of parameters in the simulation is large. It is not clear what the desired combinations are in parameter space. One should be aware that model parameters have default values hidden within the simulator. These defaults may not set the exact scenario the experimenter has in mind. For the sake of reproducibility, these parameters should be made clear.

Andel and Yasinsac [3] discuss a number of issues in existing simulations and give recommendations to improve simulation credibility. The first is the *lack of independent repeatability*. All settings should be properly documented so that the reader is able to repeat the experiments based on the provided information and obtain the same data. The next is the *lack of statistical validity*. One needs to determine the number of independent runs and sources of randomness, and to collect data only after eliminating transient values. This is also discussed in [4], which also gives statistics on the percentage of articles in top conferences that do not specify the simulator, transmission distance, number of simulation runs, confidence intervals, or traffic type, do not use mobility, and are not independently repeatable. Another one is the *lack of sensitivity analysis*. Analysis of variance should be applied to determine the significance of various factors (selected protocol, mobility setting etc.). Further criticism is about the use of *inappropriate radio models*, *unrealistic application traffic,* and *lack of real world implementation*. Finally, due to *improper precision*, [3] advises using simulations to provide proof of concept and general performance characteristics, not to directly compare multiple protocols with one another.

Kotz *et al*. [5] provide a comprehensive review of six assumptions that are still part of many ad hoc network simulation studies. In particular, they 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. They close with a series of recommendations for researchers, whether they develop protocols, analytic models, or simulators for ad hoc wireless networks. Despite the large impact of each individual factor, they also show that simulation data do not differ much from real world data under some combinations of these assumptions. Their best model is a stochastic model that captures radio signal attenuation as a combination of two effects: small-scale fading (e.g., Rayleigh and Ricean distributions) and large-scale fading (e.g., zero-mean log-normal distribution).

Conti and Giordano [6] critically discuss 10 years of research in ad hoc networking. There is a high risk of failure due to lack of realism regarding the objective and design, and assumptions such as the network being dense. They criticize the wide use of a random waypoint mobility model that does not match real users' behavior. The impact of transient phases is often neglected. The assumptions typically used (symmetric links, independence from ground height, constant radio links) are very far from reality.

Model/metric in simulations does not match the one in problem statement	25%
Simulation does not study one variable at a time	39%
Article does not give a thorough and critical literature review	41%
Simulation does not compare with truly competing solution(s)	59%
Average degree (density) is not an independent variable	69%
Article has no simulation section	22%

■ **Table 1.** *Observed practices in 45 ACM MobiHoc 2008 articles.*

The existing criticism appears to be based on simulation viewed as a validation tool for a single protocol, assuming that the protocol is a candidate for standardization and software embedding. Validation of a protocol is certainly desirable. However, it could be limited to the model and metrics used to design the protocol, and not necessarily run against a complex "realistic" physical layer that may not reflect the intended scope of research. Moreover, the goal is normally to present innovative winning protocols for certain simplified/theoretical communication models and assumptions. This is normally enforced since reality is too complex. In the context of evaluating a single article and its *contribution*, most of the criticism appears to be of limited significance, often counterproductive, and lacking one of the fundamental aspects in the design and analysis of algorithms: comparison.

CURRENT SIMULATION PRACTICES

We have studied all of the articles published at the 2008 ACM MobiHoc conference in order to derive current simulation practices with respect to items discussed in this article. The percentages of articles with the observed properties (with respect to 45 full-size studied articles) are listed in Table 1.

GOOD COMPARED TO WHAT?

We begin our discussion by opposing one of the pieces of advice given in [3]: to use simulations "to provide proof of concept and general performance characteristics, *not to directly compare multiple protocols against one another*" [3]. First, this sounds in itself contrary to other suggestions in the same articles, which ask for a thorough simulation with a variety of scenarios, and realistic modeling and testbeds. What is the value of such experimental data if we do not thoroughly compare the new protocol with the best com-

Too often some important measures are forgotten. For example, rarely can one see that the communication overhead to collect network information was measured by those that invent centralized protocols in wireless multi-hop networks.

petitors? What can be concluded about, say, 50 percent energy savings of a protocol if one does not (un)intentionally mention a competing protocol that yields, say, 90 percent savings?

Before continuing in this direction, let us first elaborate on a very frequent approach to doing research. This philosophy is to follow an "I am the best" approach. It is characterized by one or more of the listed flaws in doing or presenting research results. We note that the advice "not to directly compare multiple protocols against each another" [3] appears to provide direct support for this approach.

In the "I am the best" approach, the *literature review is incomplete*, since this is obviously the best way not to mention probably better existing solutions. Algorithmic descriptions are too often incomplete, vague, or given by pseudo-code. There are often no clear, concise, intuitive description of new ideas and no useful examples, which makes it difficult to correctly criticize them. The new solution is compared to some existing ones that are clearly inferior, often without even simulating the new one. It is often a solution that even uses different assumptions or different metrics. They give an unfair advantage to new protocols. For instance, many papers claim to be superior to AODV/DSR routing in terms of delay, power consumption, or other metrics. However, AODV/DSR/DD route discovery protocols merely use hop count as a metric and do not address congestion or power consumption issues.

Impressive simulation diagrams are presented, normally by choosing parameter values showing good performance. A lot of effort is made to use realistic scenarios, which complicates simulations. Generally great effort is exerted to show practical assumptions, scenarios, and modeling. However, too often some important measures are forgotten. For example, rarely can one see that the communication overhead to collect network information has been measured by those that invent centralized protocols in wireless multihop networks.

WHAT IS THE PROOF OF CONCEPT?

Combined with other suggestions, it is interesting to observe that proofs of concept appear possible only after thorough simulation and even testbeds (according to [3]). Our understanding of the meaning of proof of concept is quite different. We use it to mean a basic simulation, with a very simple model and scenario, matching the model with assumptions used to design the protocol, with the purpose of simply demonstrating that what is expected about the very basic performance is indeed true. Simulations simply replace theoretical proofs of performance (normally average case) because the latter are very difficult (often probably impossible) to derive, so basic claims and expectations are confirmed. In doing so, basic simulation setup provides for easy repeatability and independent verification, without any doubts and without any subtleties of any particular "well-known" simulator.

We now argue that the proof of concept as defined in [3] may in fact produce a "lack of concept" when too much realism is introduced. Moreover, it may be a declared "defeat" without any explanation because of introducing too much complexity in the quest for realism. Alternatively, it may not be clear which concept is under simulation, or the claim made for a particular concept may in fact end up being valid/invalid for another concept. Our example is the case of the georouting task in the next section.

CASE STUDY: GEOROUTING

Suppose that each node is aware of its own geographic position, and positions of all of its direct neighbors within the transmission radius, and the position of the destination (e.g., fixed sink in wireless sensor networks). What protocol should physical sensors run in real applications? Do we have a good protocol, or do we need to continue searching for it?

For the UDG model, in the *greedy routing* protocol (references on routing can be found in [7]) the node currently holding the packet will forward it to the neighbor that is closest to the destination. Simulations in the UDG model show that greedy routing is an elegant solution with *hop count* (number of transmissions) close to the optimal one, obtained from applying the shortest path algorithm. For dense networks it rarely fails (it can be augmented with a recovery procedure [7] to guarantee delivery). The shortest path algorithm, however, enjoys the advantage of centralized knowledge of all network nodes and edges available at each node, while greedy routing is a localized algorithm. Thus, the shortest path algorithm serves here as a kind of benchmark, the best possible performance for which one could hope.

Consider, for example, the *nearest neighbor* algorithm as an alternative solution. The node currently holding the packet will forward it to its closest neighbor among those that are closer to the destination than itself, and therefore provide an advance. Greedy routing certainly wins over the nearest neighbor on hop count and UDG based comparisons. What happens if a comparison is made on a simulator with a realistic physical layer?

The physical layer is quite complex. A reasonable metric would be the overall number of transmissions and acknowledgments sent between two nodes, which is called the *expected hop count* (EHC). Signal propagation loss reduces the reception probability with increased distance, and causes an increase in EHC between two nodes. The ideal routing protocol might be the shortest weighted path one, where the weight is the EHC between two nodes.

What happens with the evaluation of greedy routing? This elegant solution is penalized in simulations based on realistic physical layers because long hops lead to low reception probability and high EHC. Simulators then may declare this protocol inefficient. The nearest neighbor protocol may have overall even better EHC because of higher reception probabilities between neighbors. The proof of concept turns out to be a demonstration of the lack of concept. Does it mean that the greedy protocol was really inefficient? We believe not. It was simply evaluated in a different model from the one in which it was designed. The protocol, the metric, and the model did not advance in parallel. How should this be fixed?

A general cost-to-progress ratio framework was proposed in [8]. The best neighbor to forward to is the one that minimizes the ratio of cost over progress, where progress is difference in current and new distances to the destination, while cost is matching the desired metric and model. Greedy routing is a special case of this framework with the hop count metric used in UDG. The metric could be EHC, with a properly defined realistic physical layer, delay when transport and MAC layers are added, and so on. One can observe that the protocol does not really change; only the metric and models change. But often good protocols in one model need to be adjusted when the model changes.

We generally advocate advancing models, protocols, and simulations in parallel, on a road from UDG to applications. At any stage in the process, simulations with both matching and realistic models could be used to derive the next important factor that needs to be addressed and used to improve the model. For instance, the packet reception probability function that depends on distance was used to derive EHC [7] as a more accurate model than UDG. It is important to identify the variable from a realistic physical model that shows the biggest impact, then design a simplified model and protocol to address it in a tractable but useful way.

SELECTING INDEPENDENT VARIABLES

The choice of parameter is often such that real investigation and analysis is restrictive. A particular sound example is the choice of transmission radius as the independent variable when simulating routing protocols. Note that [3, 4] even demand reporting the selected choices. However, all distances in the communication graphs and the transmission radius itself can be multiplied by the same factor, and one gets the very same graph with different transmission radii.

What is the variable that will immediately give us some clue about the expected performance of some routing protocols? In [9] a shift in the selection of simulating position-based routing protocols was proposed: to use the average density (the average number of neighbors per node) as the independent variable itself. There is a straightforward relation between average density and transmission radius once the size of the area containing nodes is fixed. It turns out that the performance of a variety of geo-routing protocols is excellent for dense networks (density over 10), good for medium densities (6–10), and poor for sparse networks (densities below 6). This can be seen even before performing the simulations, by looking at what average choices a given node has to make progress in forwarding the traffic. Such an observation is impossible with the transmission radius as the independent variable. It is easy then to use only transmission radii that correspond to dense networks, not discuss this fact, and claim excellent performance.

SIMPLICITY AND THE STEPWISE APPROACH

The discussion in this section is primarily oriented toward simulation of network layer protocols. Oversimplifying the model can lead to erroneous outcomes [3]. This may or may not be true. In the context of the other suggestions from [3] it may be mostly disputed. Outcomes are normally erroneous with respect to various realistic physical layer models, not with respect to the simple model where the protocol is defined.

"While all models are unrealistic, some of them are useful" (George Box). "Everything should be made as simple as possible, but not simpler" (A. Einstein). Obviously, it may be counterproductive to oversimplify. For example, the assumption of an ideal medium access control (MAC) layer needs to be reexamined as it may impact the conclusions made. For instance, flooding appears to be an ideal method for route discovery. However, in dense networks many nodes receive the message at about the same time and will contend for retransmitting at the same time, leading to message losses and propagation failures since there are no acknowledgments or multiple trials. Excessive messaging is generally counterproductive for any protocol, causing a similar effect. A simple MAC layer with collisions is advisable to show the advantage of methods relaying on smaller communication overhead.

The very first questions are *what* to simulate and *whether* to simulate. A thorough literature review of this "hot" field is necessary to find existing relevant solutions. They should be clearly classified into those that do not solve exactly the same problem or do not use the same assumptions, with reasons given for exclusion from simulation. The main ideas of competing ones should be briefly described, which is a step toward explaining their simulation results. The main task during simulation is to identify scenarios and assumptions under which a new protocol is better than its competitors. The next important task is to also identify scenarios where the new protocol is inferior, and find reasons for such behavior. Understanding advantages and drawbacks of compared solutions is the first step toward finding an even better solution.

The main problem with existing analysis and criticism of simulation practices is that simulation is treated as a step on its own. However, simulation is rather a tool in the search for an optimal protocol, which does not need to be among those currently compared. The algorithm, new solution, new idea, or contribution is what matters most. Theory (mathematical analysis) and simulation (and associated protocol engineering) help, but need to be introduced and added in healthy doses.

To "cure" the mentioned problems with simulation, we further advocate a stepwise approach, adding one variable at a time, *solving one problem at a time*, and providing a full explanation for the performance of protocols on simulators that match the assumptions. *Simplicity* is the key to understanding. Suppose that a new routing

The choice of parameter is often in itself such that real investigation and analysis is restrictive. A particular sound example is the choice of the transmission radius as the independent variable when simulating routing protocols.

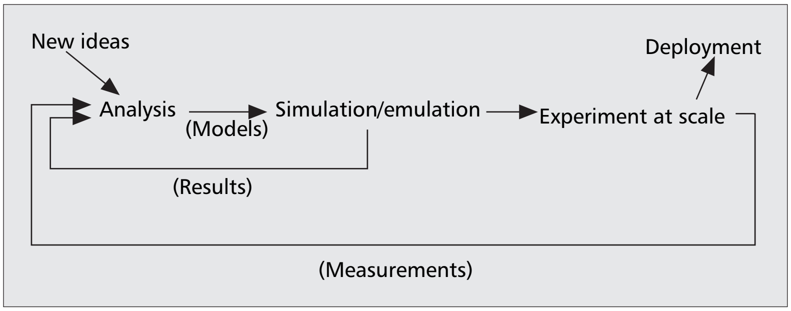

Figure 1. *Role of simulation.*

protocol is being proposed based on the UDG model and ideal MAC layer. Does it have some basic desirable properties? For example, does it create loops? If mathematical analysis finds loops, no simulation is needed. Instead, go back to the protocol design and fix it to avoid loops (e.g., [9]). If not, and it looks promising, perform *basic simulation*, comparing it with the proper competitors. Basic simulation normally means the UDG model, an ideal MAC layer (no message collisions), and one routing task at a time, matching the assumptions made when designing the protocol. Simplified simulations provide the best grounds for extracting major properties quickly.

In other words, we advise "turning off" the features not used in higher and lower layers, if possible, or alternatively building new simulators that will be incrementally improved as new features from other layers are added. Based on our experience, we find it extremely useful to limit our analysis and advances to one parameter/variable at a time, test it on the simplest possible simulator, and gain full understanding of it before making the next move. Understanding is rather challenging if the simulators contain "magic boxes" or complex modeling.

CONCLUSIONS

The existing and new recommendations on simulation practices are summarized in Table 2. While the overall recommendations are certainly good practice, they should be taken with caution. There is no single recipe for solving all problems. Different problems require different approaches, and it is difficult to give general recommendations. We believe that the arguments presented in this article provide researchers with necessary and important new views to balance currently prevailing ones.

It is clear to all researchers that the UDG model is not sufficiently realistic, although it is a good initial approximation. A typical response to address the issue is to evaluate protocols, designed for the UDG model, using popular simulators that implement realistic physical layers and scenarios. These simulations revealed some problems not seen in the UDG model. However, the experiments with UDG based protocols on simulators implementing a realistic physical layer do not explain the performance and do not provide sufficient insight toward developing better protocols. Even worse, these simulators may in fact "defeat" novel valuable protocols because of differences in the models and the impact such differences have.

Simulations done with realistic physical layers normally lead to investigating phenomena with too many variables, too many puzzles, leading to too few explanations, and too few hints for future progress. Although such simulations certainly provide some useful insight about the performance, it remains inconclusive. It is difficult to even compare two protocols properly, because some changes in physical constants may change the order of protocols, without any insight.

We therefore argue that simulators should match the assumptions made in the protocol, and advance both protocols and assumptions in parallel from the UDG toward realistic physical layers. Existing simulators also offer a platform for the proposed path. It suffices to simply switch off features not used in assumptions. Then one could add one feature at a time, observe the impact, measure, and show it. After gaining sufficient insight, the next step could be made, by designing a new model and new protocol.

The comments made here are applicable for completing a single article. The ultimate goal is to design capable protocols for use in real equipment, which may require hundreds of articles,

Problem	Current advice	New advice
Proof of concept	Thorough evaluation, testbed	Basic simulation using assumptions in designed protocol
Comparison	Not to directly compare multiple protocols against one another	Comparison with truly competing existing protocols
Literature review	Partial	Thorough and critical
Number of variables	Multiple variables of a realistic scenario	Solving one problem at a time, study one variable before adding another
Model complexity	Complex realistic modeling	Simplicity to preserve tractability
Simulation progress	Immediate simulation with complex modeling	Parallel advance of useful modeling and protocol design

Table 2. *Summary of existing and new recommendations on simulation practices.*

each solving a specific task and making specific progress. Basic simulations in one article, along the ideas and analysis, may provide insight for making further advances in the next one. In the process, the need for addressing issues at other layers and for more realism in simulation, along with the appropriate parallel changes in the modeling, metric, and protocol will arise. Arguably, this corresponds to the role of simulation as described in [10] (Fig. 1).

ACKNOWLEDGMENTS

This research is supported by EU FP7-CSA ProSense Project, UK Royal Society Wolfson Research Merit Award and Discovery and Strategic grants from NSERC Canada. The author greatly appreciates comments received by reviewers, which have improved the presentation.

REFERENCES

[1] M. Takai, J. Martin, and R. Bagrodia, "Effects of Wireless Physical Layer Modeling in Mobile Ad Hoc Networks," *ACM MobiHoc*, 2001.

[2] L. Perrone, Y. Yuan, and D. Nicol, "Modeling and Simulation Best Practices for Wireless Ad Hoc Networks," *Proc. Winter Simul. Conf.*, 2003.

[3] T. R. Andel and A. Yasinsac, "On the Credibility of MANET Simulations," *IEEE Comp*., July 2006, pp. 48–54.

[4] S. Kurkowski, T. Camp, and M. Colagrosso, "MANET Simulation Studies: The Incredibles," *SIGMobile MCCR*, vol. 9, no. 4, 2005, pp. 50–61.

[5] D. Kotz *et al*., "Experimental Evaluation of Wireless Simulation Assumptions," *ACM MWSIM*, 2004.

[6] M. Conti and S. Giordano, "Multihop Ad Hoc Networking: The Theory," *IEEE Commun. Mag*., Apr. 2007, pp. 78–86.

[7] H. Frey and I. Stojmenovic, "Geographic and Energy Aware Routing in Sensor Networks," in *Handbook of Sensor Networks: Algorithms and Architectures*, Wiley, 2005, pp. 381–415.

[8] I. Stojmenovic, "Localized Network Layer Protocols in Sensor Networks Based on Optimizing Cost Over Progress Ratio," *IEEE Netw*ork, vol. 20, no. 1, Jan./Feb. 2006, pp. 21–27.

[9] I. Stojmenovic and Xu Lin, "Loop-free Hybrid Single-Path/Flooding Routing Algorithms with Guaranteed Delivery for Wireless Networks," *IEEE Trans. Parallel Distrib. Sys*., vol. 12, no. 10, 2001, pp. 1023–32.

[10] D. Goodman, "Keynote Speech" *IEEE PIMRC*, Sept. 2008; www.geni.net

BIOGRAPHY

Ivan Stojmenovic [F'08] (Stojmenovic@storm.ca) received all degrees in mathematics. He has published over 200 papers and edited four books. He is editor of a dozen and founder of three journals. One of his articles was recognized as the Fast Breaking Paper for October 2003 (the only one for all of computer science) by ISI. He has chaired and/or organized over 30 workshops and conferences.