Efficient Models of Fine-Grain Variations in Signal Strength[1]

John P. Mullen
Center for Stochastic Modeling
New Mexico State University
Las Cruces, NM.88001
E-mail: jomullen@nmsu.edu

Abstract

Presently, there are great differences between the performance of Mobile Ad Hoc Networks (MANETs) predicted by network simulators and actual field experience. One cause is that, although fine-grained variation in received signal strength greatly impacts the effectiveness of MANET routing protocols, it is almost never currently modeled. Although precise modeling of this variation is not feasible, it can be represented by rather simple stochastic models. This paper describes the impact of considering this variation in MANET models by means of modifications to the OPNET transceiver pipeline that portray fine-grained variation. This approach is highly efficient, requires little additional modeling information, and provides simulated data researchers may use to better predict actual field perfomance in order to better optimize OPNET routing protocols and their parameters.

Introduction

Recently, a number of researchers have attempted to implement *Mobile Ad Hoc Networks* (MANETs) in the field [1-6]. In general, performance differed greatly from that predicted by network simulations. These differences are very large and the result is that protocols must be heavily modified in order to perform at even a fraction of the predicted level. This gap in results means that simulation is not as good a tool to design MANET routing protocols as in could be. This presents a significant problem in that virtually all MANET protocol development is carried out on network simulators. In [7-9], the authors identified differences in physical layer modeling as a fundamental cause of differences between the performance of physical and simulated MANETs. Reducing this difference would improve the ability of network simulators to guide the design of physical tests and implementations.

This paper demonstrates the value of an improved wireless link model by means of a simple modification to OPNET ®[2] that vastly improves its ability to model physical MANETS. This modification does not add significantly to either modeling effort or execution time. While certainly not a perfect model, this increased fidelity improves the researcher's ability to predict field performance through simulation, improving the value of both approaches.

This paper is organized as follows. First, it describes fine-grained variations in signal strength and their impact on MANET routing protocol performance. Then, it develops a simple stochastic model of this variation and implements it in OPNET. Next, it validates the model through comparisons to field tests. After that, it uses this model to demonstrate some insights into basic MANET routing behavior. Next is a demonstration using a 25 node scenario. Finally, it summarizes, draws conclusions, and outlines future work.

The primary purpose of this work is to demonstrate the feasibility of this approach. For that reason, unless otherwise noted, default values were used in each example. While each protocol was checked for its main functions, no attempt was made to certify that the protocols, as implemented in OPNET, were exactly the same as those reported in other works. Therefore, while this paper outlines the need to consider fine grained signal variability and demonstrates the feasibility and utility of a method, it makes no claims for the strengths and weaknesses of any particular protocol.

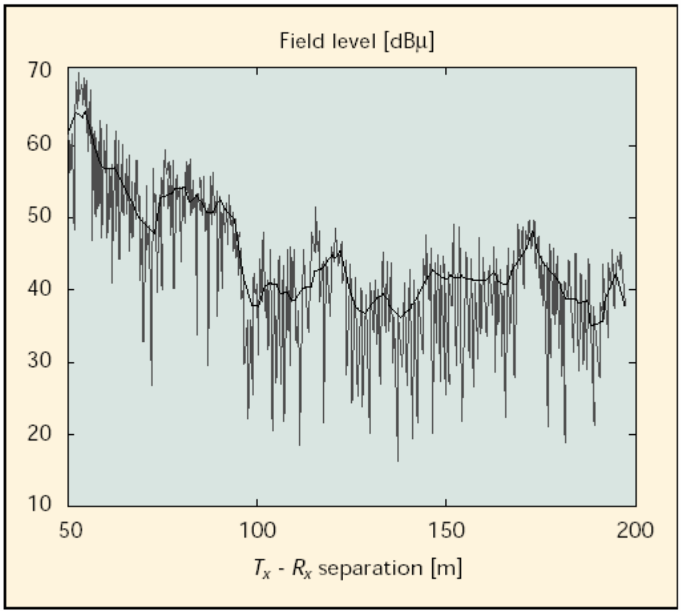

Figure 1: Actual Field Strength Measurements [10]

---

[1] This is an edited version of a paper presented at OPNETWORK 04, Washington, DC, August 30 – September 3, 2004. Permission is granted for reproduction for non-commercial purposes. Permission to copy granted provided that the copies are not made or distributed for commercial advantage and this notice is included.

[2] OPNET Technologies, Inc.

Multipath Fading and its Impact on MANETs

There are a variety of mathematical models that may be used to predict received power as a function of distance, obstructions, radio carrier frequency, etc. [10, 11]. Such predictive models are the basis of almost all simulations of MANET routing protocols [7]. However, these models predict only the *expectation* of received power and do not take into account fine-grained variations in that signal strength due to shadowing, multipath fading, and other causes. These effects can easily cause the signal to vary from the predicted value by up to ±30dB [11, 12]. An example of this variation is illustrated in Figure 1. Clearly, this signal exhibits a very high level of fine granularity variability.

Key Impact: This high-magnitude fine-grained variation has two important effects on wireless digital systems.

1. Even at close range, a packet may be dropped, and
2. Even at great distances, a packet may be received.

In point-to-point and stub networks, Effect 1 is a serious consideration, but Effect 2 is usually a much less serious concern. However, this is not the case in MANETs. In this case, these effects can lead to two detrimental events:

1. Occasionally, packets are lost on a good link, and
2. Occasionally, packets are received on a poor link.

In a MANET, Event 1 is likely to trigger an unnecessary route search, significantly reducing performance. If Event 2 occurs during a route search, a poor link might appear to be reliable. If this happens in protocols that seek routes with fewer hops, the one-hop, but unreliable, link with be selected over a multihop, but reliable, alternative. When this route fails, a new search is initiated, reducing performance. Finally, whether a search is initiated by Event 1 or by an actual need to find a new route, it generates a chance for Event 2 to cause a poor route to be selected again. This has proven to be a very serious problem in the field [1, 2].

Let P_1 be the probability of Event 1 and P_2, that of Event 2. In stub networks, one can increase transmitter power, increase the number of retries, and use other methods to reduce P_1. In cellular networks, similar approaches can be used, though Event 2 can be a problem. However, P_2 is reduced by use of the basic seven-cell honeycomb pattern which places the distant users of the common frequency at least four cell radii apart. Also, since there is no multi-hop routing in either system, even if a packet is received over a poor link, its impact is fairly small.

In MANETs, however, all nodes share a common channel, at least for signaling and routing. “Distant” users may be barely further away than “close” ones. In addition, the occurrence of Event 2 during a route search can lead to a very significant reduction in performance. Thus, in a MANET, effectiveness of any strategy to reduce P_1 is limited by the fact that as P_1 decreases, it is likely that P_2 will increase. Unless this trade-off is considered in MANET modeling, the results can be significantly unrealistic and lead to poor decisions for parameter selection and protocol design.

Multipath Fading: While there are numerous reasons that received power may rapidly fluctuate, usually the most significant in MANETs is that due to multipath fading [11, 12]. This fading is the consequence of receiving multiple copies of the same transmission two or more paths of different lengths. The different copies can either reinforce or partially cancel each other out, depending on the difference in path lengths. As a consequence, the received signal may vary from the expectation by ±30dB [11, 12].

It is the nature of MANETs, with most transmissions from small, low antennas, especially in a reflection-rich settings, such as a convention center or urban environment, to have ample opportunities for multipath reception. It is the ability of the signal to be vastly *stronger* than expected that leads to Event two. Unless this multipath fading is considered, P_2 may be vastly underestimated and results will present an overly optimistic estimate of MANET performance.

A Simple Model of Multipath Fading

The main idea is to construct a practical model of a wireless link that takes more fully into account the random effects of multipath fading. Here, *practical*, means a model that:

1. predicts possible variations in received power,
2. without requiring inordinate modeling detail or excessive execution time, and
3. also makes available to the modeler such values as BER and SNR, so that methods to cope with this variability may also be modeled.

The model presented here satisfies all three requirements.

Detailed Fading Models: While methods exist to closely model multipath fading [13], these require an immense amount of detailed information about the site and, as a result, have limited general application. For a MANET, in which the nodes may be in a reflection-rich environment in which many elements are in motion, such an approach would consume immense amounts of CPU time and memory. In addition, if there's any change in the length of any of these paths, or if paths are either created or destroyed, the received power will change. Nearly complete cancellation can cause the signal to drop below the receiver's threshold at virtually any distance. On the other hand, nearly complete reinforcement can cause the received signal to be many times its predicted level. As a consequence, 1) in most situations, multipath fading has the greatest impact on received power at greater distances, and 2) for all practical purposes, it cannot be modeled explicitly in even a simple RF environment [10-12].

Stochastic Fading Models: An alternate approach is to characterize the fading as a stochastic process. This allows one to more faithfully model the nature of wireless signals without additional modeling detail or simulation effort. Although the result is not a precise prediction of every communications result, it does predict performance that is typical of the modeled situation. As shown below, this is usually adequate to predict MANET performance in the field, especially considering the fact that the ultimate aim is to design protocols that work in a wider variety of situations.

Three common stochastic models of multipath fading are the Rayleigh, Ricean, and Nakagami [10-12]. In this paper, the Rayleigh model is used because 1) it is sufficient to illustrate the basic concepts, 2) it is simple to implement, and 3) the results can be compared to field results reported in [1] in which Rayleigh fading seemed to be in effect.

The Rayleigh fading model specifies that $P(d)$, the random received power at distance d has the following distribution:

$$\Pr\{P(d) \le p\} = 1 - \exp\{-[p / m_P(d)]\} \tag{1}$$

where $m_p(d)$ is the expectation of received power at distance d. The model used to predict $m_P(d)$ may be any one of several predictive models and may be arbitrarily detailed to include antenna gain or loss, bandwidth, RF frequency, obstructions, etc. Use of the *inverse transformation method* [14] leads to:

$$P(d) = -m_P(d) \ln(1 - r) \tag{2}$$

where $P(d)$ is the pseudorandom received power level at distance d and r is a pseudorandom number uniformly distributed on the open interval (0,1).

Implementation in OPNET is straight-forward. First, the line computing **rcvd_power** in the **wylan_power** stage of the OPNET transceiver pipeline is modified so that it computes **mean_rcvd_power**. Then, it is immediately followed by:

```
rcvd_power = op_dist_exponential(mean_rcvd_power);
```

The modified code was saved as a fading version of **wylan_power** stage and modules using this model similarly renamed to allow comparisons of the default (non-fading) and modified (fading) models. With this modification, subsequent code in this stage, as well as following stages in the pipeline, react to this random power level, more realistically portraying link behavior, as well as multipath effects on interference. This also satisfies Objective 3 by making the random power levels available to modules that compute SNR, BER, etc.

From limited testing in OPNET 9.1, 10.1, and 10.5, it appears that the modified modules can replace their equivalent unmodified forms without any other changes. Also, in most scenarios, there was little or no direct impact on execution speed. In those that did have execution time differences, they were moderate and should be considered in light of the higher fidelity. That is, without the fading model, the only alternative would be field tests, which would be far more expensive and time consuming.

Figure 2: First Scenario Layout

Validation: The first check was to employ a simple test that compared simulated received power predicted by default and fading models to the trace noted in Figure 1. As depicted in Figure 2, A mobile transmitter sent one 1024 bit packet each second to a stationary receiver. The mobile node traveled along a straight line at a rate of 0.5 m/s along a path that ran from 5 to 100m away from the receiver. In both runs, default OPNET settings were used, except that transmit power was set to 5μW and the *Short Retry Limit* (SRL) was set to one.

The smooth black trace in Figure 3 is $m_P(d)$, as predicted by the unmodified **wylan_power** stage in OPNET. The red erratic trace is an instance of simulated received power using the fading version. Note the similarity of the variation in $P(d)$ to that of the field measurements in Figure 1.

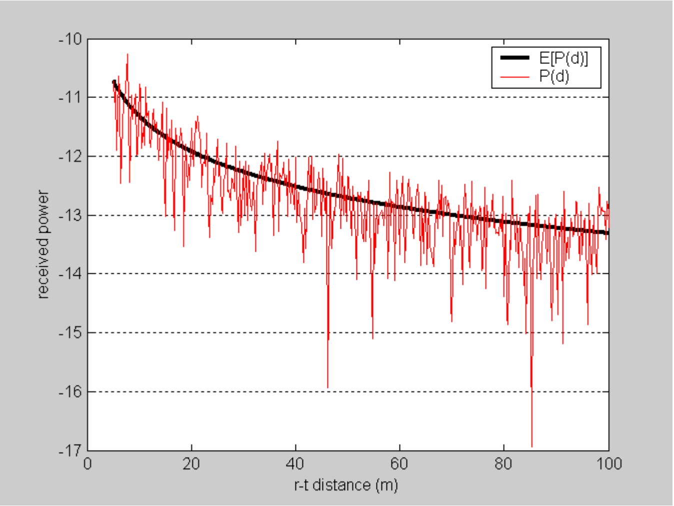

Figure 3: Simulated Received Power

Scenario 2: A second layout, adopted from [1], was designed to measure the basic ability of a MANET protocol to select and manage multihop routes. The layout is depicted in Figure 4. In this figure and the following description, all distances are in meters. All message traffic is generated by the mobile node (MH2) and is addressed to MH 1, which is stationary, as are. potential relay nodes N1 through N3. MH 2 starts at (0,40)[3] and moves to the right at a rate of 0.5m/s until it reaches (100,40). The transmit power is 5μW, yielding an effective range of 39m. Thus, at all times, there is a viable route from

---

[3] In OPNET, the positive y direction is downward.

MH 2 to MH 1, but not always a one-hop route. This test requires routing for one through four hops, as well as effective detection of broken links and searches for new routes. In this scenario, the transmission rate was fixed at ten 1024 packets per second.

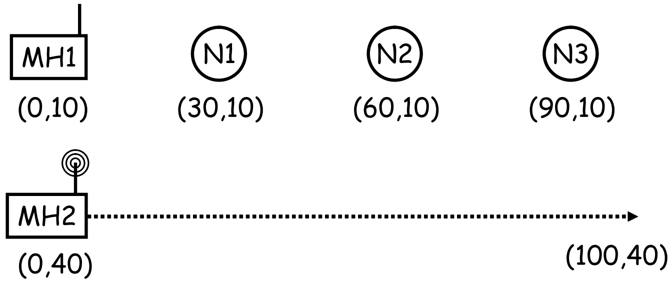

Figure 4: Second Scenario Layout

Figure 5 illustrates the packet throughput in two simulations. The red trace is that predicted by the standard, non-fading, model and the blue is that predicted by the Rayleigh fading model. Note that the two sets of results are radically different. The results with the fading model are consistent with the observations in [1] and [5] that AODV could not reliably maintain a path more than two hops long. The intermediate performance is also similar to that in [15] where only paths up to two hops in length were considered. Thus, it appears that in this test, the fading model more realistically predicts the real-world performance of AODV.

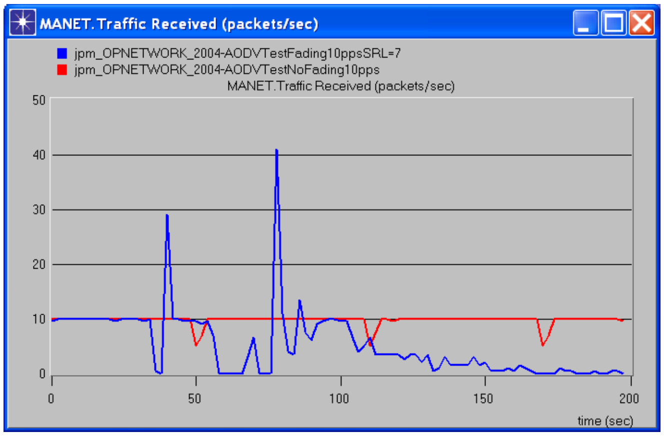

Figure 5: Simulated AODV Throughput

Demonstrations in Simple Scenarios

The models in this section employ either Scenario 1 or Scenario 2. Although these tests are very simple, they isolate the main effects of multipath fading on link throughput and routing. The next section demonstrates the use of the model in a larger, more complicated, model.

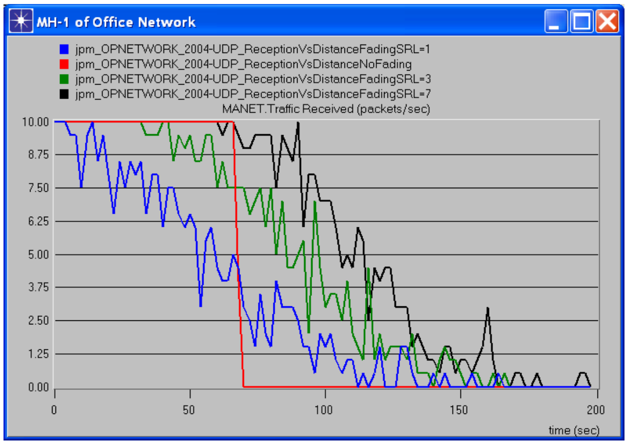

Figure 6: Effect of SRL on Throughput in Scenario 1.

Impact of Short Retry Limit (SRL): Figure 6 shows the packet throughput in Scenario 1 with the *Short Retry Limit* (SRL) set to 1 (blue), 3 (green), and 7 (black). The red trace is the result in the non-fading model, which is unaffected by the setting of SRL. In the fading model, as the SRL is increased, throughput within the nominal range, (39m) improves significantly. However, at the same time, the probability of Event 2 also increases. This would imply that large values of SRL could be counter-productive in a MANET. This point was explored further in [16].

AODV vs. DSR: In [6, 17, 18] the authors found in field tests significant differences between AODV and DSR. Although the non-fading model predicts no differences in either Scenario 1 or 2, the fading model does so. Figure 7 shows that the fading model predicted significant differences between the throughput of DSR and AODV in Scenario 2. Once again, the fading model presents results that are more consistent with field experience than the default model

.

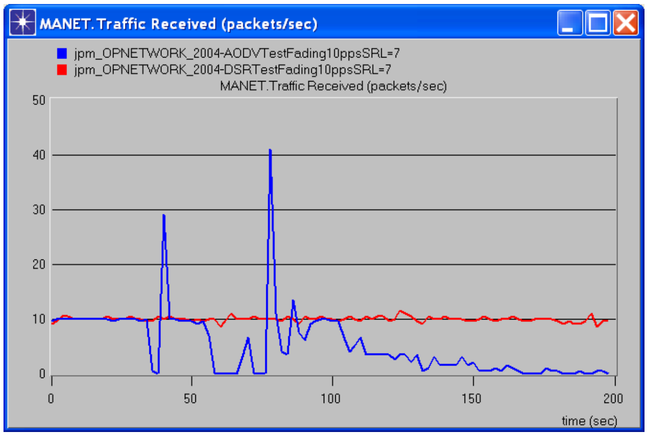

Figure 7: Simulated AODV and DSR Throughput

While this does not mean that DSR is superior to AODV in this situation, it does imply that 1) the non-fading model is not adequate to predict this difference in performance and 2) at the least, AODV parameters need to be adjusted.

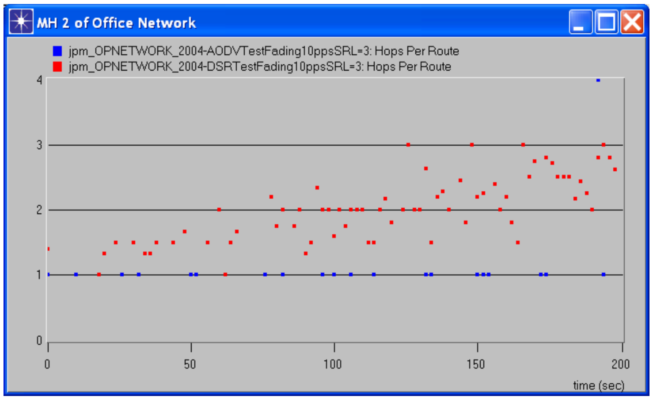

Figure 8: Number of Hops pre Route in Scenario 2.

Figure 8 illustrates a probable explanation for the difference in behavior. This is a plot of the number of hops per route for AODV (blue) and DSR (red) with SRL set to three. Note that as MH 2 moves further to the right, DSR attempts to use an appropriate number of hops, but AODV repeatedly attempts to use a one-hop route.

A Demonstration in a Larger Scenario

While the demonstrations above show differences at the link level and in a specific routing problem, the question arises, "How much difference does this make in a more realistic scenario." In addition, a load of one 1024 bit packet per second is a very light load on a 1Mb/s channel. The difficulty due to Event 2 is present at all loads. So, another question is, "Would Event 2 still be a problem at higher loads?" This section attempts to answer that question.

The scenarios used in this section are slightly modified versions of the DSR and AODV **station_mobility** scenarios the OPNET MANET example model. Each basic scenario was run in three versions:

1. Using the existing non-fading power stage, unmodified,
2. Using the fading power stage in the existing model, and
3. Using the fading power stage in the existing model, but with the range for transmitter exclusion set to 5km.

The first set of simulations established a baseline for the current non-fading model. The second simulated a situation in which receivers are GPS enabled and exclude transmissions from transmitters more than one km away and the third a situation in which receivers do not use GPS data.

In addition, tests were repeated with the mean packet generation rate set to 1, 4, and 10 packets per node per second. Since there are 25 transmitting nodes, these loads are far higher than the ones of Scenario one and two. Also, using three levels provides insight into how the protocols behave at different loads. Of course, this leads to differences in the non-fading, as well as the fading, models. For that reason, the non-fading case is shown in each figure, as a baseline. Finally, rather than traces of network behavior, the results in this section are the average packet throughput and delay.

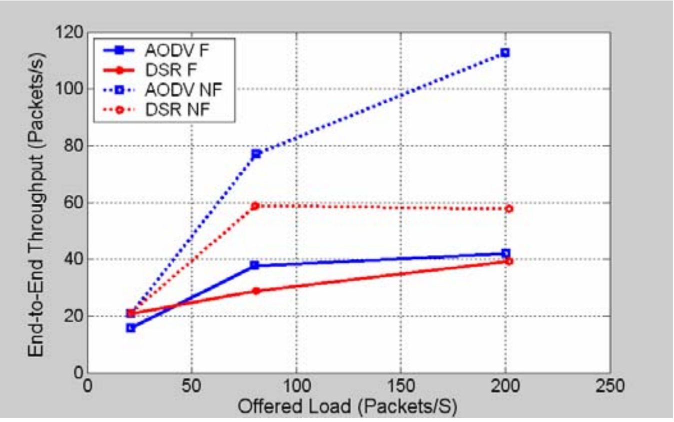

Figure 9: AODV vs. DSR: Mean Throughput, No GPS

AODV vs. DSR, no GPS: Figure 9 illustrates the mean throughput as a function of offered load in data packets per second (P/s), with GPS not enabled. Note that the non-fading models are very optimistic. Also, while the non-fading models show AODV is uniformly superior to DSR, the fading model indicates a more mixed situation.

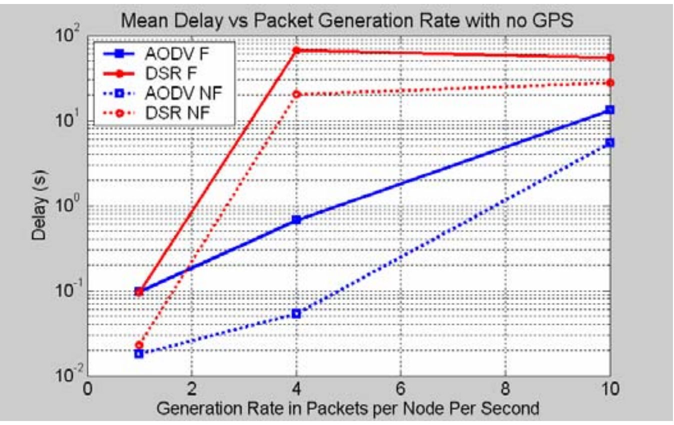

Figure 10: AODV vs. DSR: Mean Delay, No GPS

Figure 10 shows the mean delay for both protocols as a function of the mean number of packets sent per node per second. Here, a logarithmic scale is used because of the great differences in values. Note the delay predicted by the fading and non-fading models for DSR can differ by a ratio of 20:1 while those for AODV can differ by up to two orders of magnitude.

AODV vs. DSR, with GPS: Figure 11 displays the mean throughput with GPS enabled. Note that the non-fading throughput is identical to that of the first case, while the fading model indicates a significant improvement for both protocols.

Also, it appears that by ignoring the long, unreliable, routes, AODV's performance is very much improved by this strategy.

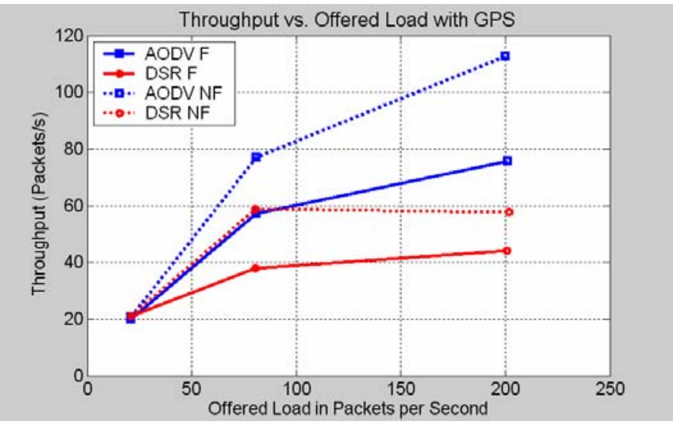

Figure 11: AODV vs. DSR: Mean Throughput, with GPS

Figure 12 shows the mean delay for both protocols with GPS enabled. Again, the non-fading model predicts no improvement, while the fading model predicts a significant decrease in delay. However, it appears that DSR can still have some very long delays.

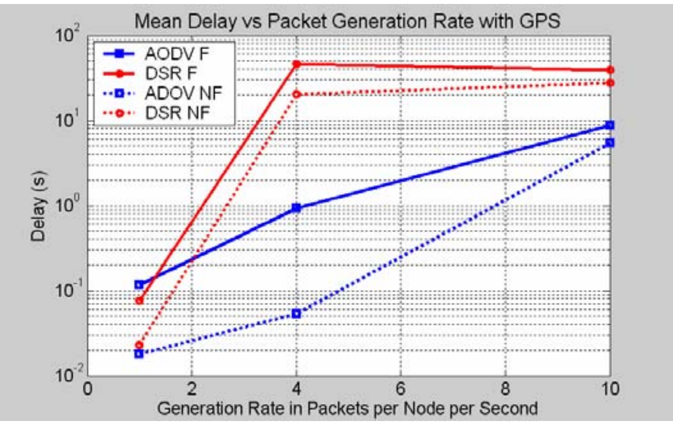

Figure 12: AODV vs. DSR: Mean Delay, with GPS

Impact of SRL on DSR: There is some indication that reducing the short retry limit (SRL) can be beneficial in DSR [16]. Figure 13 shows the mean throughput of DSR with SRL set to 7 (the default) and to 3. For reference, the figure also shows the prediction of the non-fading model, which is the same for both SRL values. Note that throughput is greatly improved and even approaches that predicted by the non-fading model at the highest offered load.

Figure 14 shows the mean delay for DSR at the two levels of SRL, as well as with the non-fading model. Note that when SRL is equal to 3, the actually delay can be less than that predicted by the non fading model. However, it should be borne in mind that this reduction in delay occurs because the older packets are being dropped. In a given situation, this may or may not be as beneficial as this plot implies. Nevertheless, it does seem that reducing SRL to 3 has a significant impact on DSR performance that the fading model would not predict.

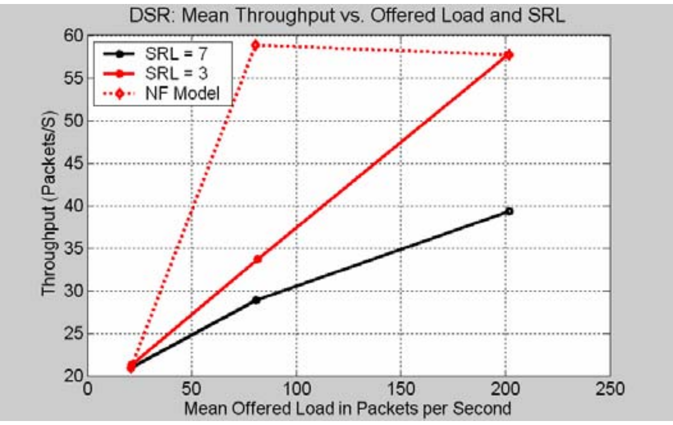

Figure 13: DSR: Throughput vs. SRL

While changing SRL to 3 in the DSR model can make the fading and non-fading models predict similar values of throughput and delay at high loads, because the non-fading model is insensitive to this parameter, the non-fading model could not be used to reach this conclusion. Also, since SRL has a much smaller effect on AODV, selecting SRL to have a default value of 3 is not likely to improve the fidelity of all simulations with the non-fading model.

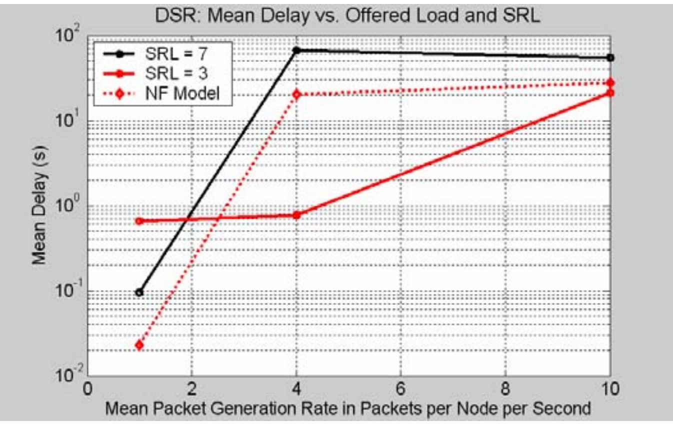

Figure 14: DSR: Delay vs. SRL

Impact on Execution Time: In Scenarios 1 and 2, there was no significant difference in execution times. However, in Scenario 3, there were. Table 1 shows the execution times for the models described above. In some cases, the execution times were two to four times as great. However, bear in mind that 1) these results may be far more consistent with field results than current ones and 2) even a 400% increase in execution time is a small price to pay when the alternative is a field test.

Table 1: Execution times in seconds

GPS?	Load (Pkts/s)	Fading		No Fading	
		AODV	DSR	AODV	DSR
Yes	1	180	493	144	351
	4	502	1833	345	1670
	10	681	2160	702	2303
No	1	638	726		
	4	931	2696		
	10	1312	3221		

Note: Because GPS had no effect on the results in the non-fading model, those models were not run and no times are available.

Summary and Conclusions

This paper demonstrates that multipath fading causes significant effects in MANETs that might not arise in other types of wireless networks. It also presents a simple stochastic model of multipath fading and demonstrates its use in OPNET. After validating the simulated results against those of several field tests, the paper presents number of examples in three different settings. Finally, the paper explores the impact of protocol, GPS elimination, and Short Retry Limit. In each case, the default non-fading model predicted little or no difference between the options, while the fading model showed that significant differences in performance exist.

Recommendations: This technique has promise, but to use it with confidence, more validation would be required. If it is used in conjunction with field tests, it has the potential to improve the efficiency of those tests while also improving its own validity. However, by whatever means, if simulation is to continue to be a useful tool in the design of MANET protocols, this and others methods of significantly improving fidelity, even if inexact, need to be explored.

Future work: While Rayleigh fading is a good first approximation to received power, Ricean, and Nakagami fading may be more appropriate. Hence, these models, together with a suitable interface to select the type of fading and its parameters would be needed. In addition, this approach ignores some factors, such as packet size, which affect both P_1 and P_2. If these effects play a significant role in analysis, other changes may be needed in the transceiver pipeline. Finally, since this approach opens up more possible modes of comparison, it may be advisable to augment the existing statistics package and other modules to consider more factors.

Acknowledgements

This work is possible, in part, due to a software grant from OPNET Technologies, Bethesda, MD. The author would also like to thank Dr. Tim Matis, for his continued support through the Center for Stochastic Modeling at New Mexico State University. He would also like to thank Dr. Steve Horan and Dr. Hong Huang, both of the Klipsch School of Electrical and Computer Engineering at NMSU for their inspiration and very helpful suggestions.

References

[1] K.-W. Chin, J. Judge, A. Williams, and R. Kermode, "Implementation experience with MANET routing protocols," *ACM SIGCOMM Computer Communications Review*, vol. 32, pp. 49 - 59, 2002.

[2] I. D. Chakeres and E. M. Belding-Royer, "The Utility of Hello messages for determining link connectivity," *Wireless Personal Multimedia Communications*, vol. 2, pp. 504 - 508, 2003.

[3] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A High Throughput Path Metric for MultiHop Wireless Routing," presented at MobiCom '03, San Diego, California, USA, 2003.

[4] H. Lundgren, E. Nordström, and C. Tschudin, "Coping with Communication Gray Zones in IEEE 802.11b based Ad hoc Networks," presented at WoWMoM'02, Atlanta, GA, 2002.

[5] S. Desilva and S. Das, "Experimental evaluation of a wireless ad hoc network," 2000.

[6] D. A. Maltz, J. Broch, and D. B. Johnson, "Lessons from a full-scale multihop wireless ad hoc network testbed," *IEEE Personal Communications*, vol. 8, 2001.

[7] D. Cavin, Y. Sasson, and A. Schiper, "On the Accuracy of MANET Simulators," presented at POMC'02, Toulouse, France, 2002.

[8] Y. Zhang and W. Li, "An Integrated Environment for Testing Mobile Ad-Hoc Networks," presented at MOBIHOC'02, EPFL Lausanne, Switzerland, 2002.

[9] M. Takai, J. Martin, and R. Bagrodia, "Effects of Wireless Physical Layer Modeling in Mobile Ad Hoc Networks," presented at MOBILEHOC, 2003.

[10] A. Neskovic, N. Neskovic, and G. Paunovic, "Modern Approaches in Modeling of Mobile Radio Systems Propagation Environment," *IEEE Communications Surveys*, pp. 2 - 12, 2000.

[11] T. S. Rappaport, *Wireless Communications: Principles and Practice*, 2 ed. Upper Saddle River, NJ: Prentice-Hall, Inc., 2002.

[12] J.-P. Linnartz, *Narrowband Land-Mobile Radio Networks*. Boston, MA: Artech House, 1993.

[13] M. Nidd, S. Mann, and J. Black, "Using Ray Tracing for Site-Specific Indoor Radio Signal Strength Analysis," Computer Science Department, University of Waterloo, Ontario, Canada 1995.

[14] A. M. Law and W. D. Kelton, *Simulation & Modeling Analysis*, Second ed. New York, NY: McGraw-Hill, 1991.

[15] I. D. Chakeres and E. M. Belding-Royer, "AODV-UCSB Implementation and Ad hoc Network Field Results," vol. 2004. San Francisco, CA: MobSys, Poster, May 2003, 2003, pp. Poster.

[16] J. P. Mullen, T. Matis, K. Adams, and S. Rangan, "Achieving Robust Protocols for Mobile Ad Hoc Networks," presented at 2004 IERC, Houston, TX, 2004.

[17] D. A. Maltz, J. Broch, and D. B. Johnson, "Experiences Designing and Building a MultiHop Wireless Ad Hoc Network Testbed," School of Computer Science, Carnegie Mellon University, Pittsburgh, PA CMU-CS-99-116, March 5, 1999 1999.

[18] D. A. Maltz, J. Broch, and D. B. Johnson, "Quantitative Lessons From a Full-Scale Multi-Hop Wireless Ad Hoc Network Testbed," Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 2000.