Comparing the ns–3 Propagation Models

Mirko Stoffers
School of Electrical and
Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia, USA
Email: stoffers@gatech.edu

George Riley
School of Electrical and
Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia, USA
Email: riley@gatech.edu

***Abstract*—An important aspect of any network simulation that models wireless networks is the design and implementation of the *Propagation Loss Model*. The propagation loss model is used to determine the wireless signal strength at the set of receivers for any packet being transmitted by a single transmitter. There are a number of different ways to model this phenomenon, and these vary both in terms of computational complexity and in the measured performance of the wireless network being modeled. In fact, the *ns–3* simulator presently has 11 different loss models included in the simulator library. We performed a detailed study of these models, comparing their overall performance both in terms of the computational complexity of the algorithms, as well as the measured performance of the wireless network being simulated. The results of these simulation experiments are reported and discussed. Not surprisingly, we observed considerable variation in both metrics.**

I. Introduction

An important part of any wireless network simulation is the appropriate choice of the *Propagation Loss Model* to be used to model the performance of a wireless network channel or set of channels. These models are needed in order for the simulator to compute the signal strength of a wireless transmission at the receiving stations which in turn is required to determine whether or not each of the potential receivers can in fact receive the information without bit errors. There are a variety of such models, varying from abstract fixed loss models, to a simple exponential decay proportional to the distance between a transmitter and receiver, to models accounting for ground reflections, to models accounting for fast fading. Each of these models requires different amounts of computation to determine the relative signal strengths at each of the receivers, and correspondingly each of the models have differing levels of accuracy.

The *ns–3* simulation tool has 11 different loss models included in the distribution. We categorize those loss models into three groups:

1) Abstract propagation loss models do not model realistic propagation loss, but need to be configured to fit the given scenario.
2) Loss models in the second category model the deterministic path loss over the distance from sender to receiver.
3) The third category includes fading models. A stochastic fading process is intended to be applied on top of a path loss model in order to account for the non-deterministic effects caused by moving objects.

A short description of every loss model is provided in Section III. Each of these models will likely produce differing results from the propagation loss computation, and therefore lead to differing measured results in the wireless network being simulated. Further each of these models will utilize differing amounts of CPU time required for each propagation loss computation, resulting in varying computational complexity for the simulation itself.

The goal of this work is to categorize each of these *ns–3* models, both in terms of computational complexity and in terms of variations in measured results. We report on the relative computational complexity for each of the models (in terms of computation time per packet transmitted), and report on the variations we observed in measured results. We do not comment on relative *accuracy* of the measured results of these models, since different models are designed to model different environments. In addition, it has been shown by a number of field experiments that the actually measured value of receiver signal strength or receiver reception probability varies significantly depending on the environment and thus cannot be accurately predicted. Therefore, this paper is not about validating the implemented models. However, we hope that these results will assist *ns–3* users in choosing an appropriate model for their future wireless experiments.

The remainder of this paper is structured as follows: Section II discusses prior work in the field of propagation loss. We introduce the propagation loss models shipped with *ns–3* in Section III. Section IV discusses our experiment setup and Section V the measured results. Finally, we summarize our conclusions in Section VI.

II. Related Work

Accurate and efficient models for wireless data transmissions have been the subject of many research works over the last decade and longer. Rappaport [1] devotes several chapters in his popular textbook describing the mathematical formulae

This work is supported in part by the U. S. National Science Foundation under contract numbers CNS 0958015 and CNS 855110.

that can be used to model electromagnetic propagation between two points in a three dimensional space. Several of Rappaport's equations form the basis for some of the *ns–3* models. More common are published works that report on real-world experiments that measure and analyze the performance on an actual, deployed or temporary network.

While we understand that most of the works cited in the following paragraphs are not directly related to wireless channel modeling in simulations, we include the discussion here to reinforce the fact that electromagnetic propagation is highly variable, and is affected in many ways that are difficult or impossible for a modeler to incorporate.

Aguayo [2] reports on results from a series of measurement studies based on an existing network in Cambridge (MA, USA) around MIT called *RoofNet*. The RoofNet consists of 38 IEEE 802.11b base stations mounted on or near rooftops at various points around their campus. The study uses an active probing technique to measure packet reception probability at each of the potential receivers for a continuous burst of packets from a single transmitter. The reported results from this study are quite surprising, showing that conventional wisdom regarding reception probability (closer pairs should have higher reception probability) simply does not always hold. Indeed, in some of the Aguayo experiments, node pairs that were as far apart as two kilometers were in fact able to communicate about half the time, while other nodes as close as a few hundred meters got zero percent reception probability.

Kotz et. al [3] also perform a set of active measurement experiments, but their approach was to deploy a temporary network (graduate students and volunteers carrying laptops) with mobility to measure received signal strength and reception probability under controlled conditions. Using this approach, Kotz enumerates six common assumptions frequently used when designing path loss models, and subsequently shows that none of these six assumptions actually hold under experimental conditions.

Reddy et. al [4] report on a set of experiments similar to that of Kotz, but from a different perspective. That work shows that the measured signal strength as a function of distance can in fact be used to create a stochastic path loss model, and that the stochastic model can in fact produce simulated results that match reasonably well with the measured field experiments. However, such an approach depends on having access to the exact geographic location for the network being modeled, which is usually not possible.

More recently, Zheng and Nicol [5] describe a detailed experiment using an *Anechoic Chamber*, which is a large room with substantial radio signal shielding that essentially isolates the chamber from outside electromagnetic interference. Using this chamber, they measure and report on the received signal strength for various distances and antenna characteristics. As in other works reported above, even with the completely isolated chamber, the measurements often do not match those predicted by mathematical models.

The works mentioned above are just a small set of published works showing that mathematical models, regardless of the complexity, can rarely be expected to predict wireless network signal degradation with high accuracy.

This assumption is also proved by independent work by Abhayawardhana et al. [6] and Durgin et al. [7]. These works compare certain propagation loss models to the signal power degradation measured in real world scenarios. Additionally, they demonstrate high fluctuation in the received signal strength over time or while changing the distance only slightly.

We therefore argue that neither model can accurately predict the propagation loss in different environments. Instead, to retrieve really accurate results, measurements in the target environment are always necessary.

III. Propagation Models

In this section we briefly introduce the propagation models included in the most current release of *ns–3*, and categorize them into three groups:

1) Abstract propagation loss models:
   a) *Fixed Received Signal Strength.* Regardless of the distance the receive power is fixed to a predefined value.
   b) *Matrix Loss Model.* The propagation loss is fixed between each pair of nodes.
   c) *Maximal Range.* A maximal range determines how far the signal is retrieved. Within that range it is retrieved at the transmit power level.
   d) *Random Propagation Loss.* The propagation loss follows a random distribution.
2) Deterministic path loss models:
   a) *COST-Hata Model.* A model based on various experiments used to predict path loss in urban areas [8].
   b) *Friis Propagation Model.* The propagation model by Harald T. Friis [9] calculates quadratic path loss as it occurs in free space.
   c) *Log Distance Path Loss Model.* The log distance path loss model [10] assumes an exponential path loss over the distance from sender to receiver. It is designed for suburban scenarios.
   d) *Three Log Distance Model.* A variation of the log distance model. It applies different factors to the logarithmic path loss for different distance intervals.
   e) *Two Ray Ground Model.* This model was initially developed by Rappaport [1]. It assumes a radio propagation via two paths: One ray is received directly, the other one reflects on the ground.
3) Stochastic fading models:
   a) *Jakes Model.* The Jakes model [11] calculates the propagation loss by modeling a set of rays transmitted from the sender to the receiver via different paths.
   b) *Nakagami Model.* The Nakagami model [12] is similar to the Rayleigh model, but describes different fading equations for short-distance and long-distance transmissions.

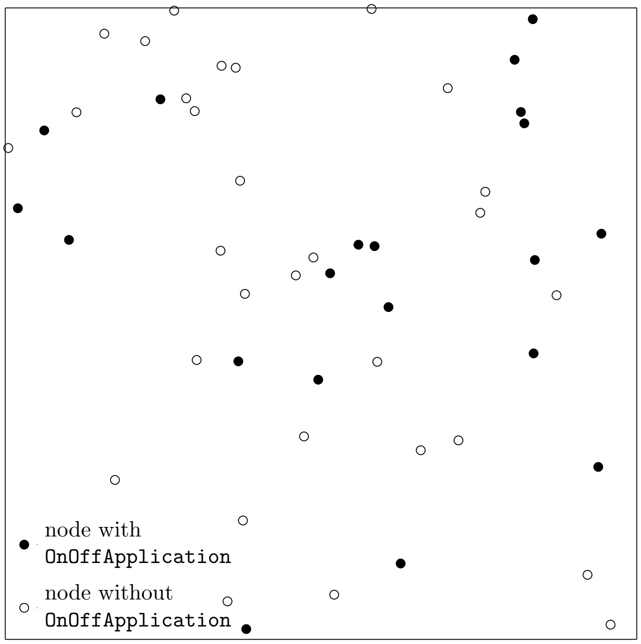

Fig. 1. 50 nodes are randomly placed in a 1 km by 1 km area. 20 nodes run an `OnOffApplication` sending *UDP* packets addressed to a randomly selected peer. This figure shows a possible placement of the nodes.

IV. Experiments

We designed an *ns–3* scenario using a simple wireless ad-hoc network that allows us to perform a comparative analysis of the measured network efficiency (total packets received at destinations divided by the total packets sent by sources). In these experiments 50 IEEE 802.11a wireless ad-hoc nodes are randomly placed on a 1 km by 1 km region. A subset of 20 nodes is formed. Each node in this subset chooses a random peer and generates a stream of *UDP* packets addressed to that peer using the *ns–3* `OnOffApplication` with a traffic intensity of 2 % and a data rate of 500 kb/s. A resulting network is shown in Figure 1. Three different routing protocols are used, specifically the *AODV*, *DSDV* and *OLSR* protocols. The transmission power is varied from 1 mW to 1 W.

Each of the abstract models, as well as each of the path loss models, are used in turn, configured with the *ns–3* default parameters shown in Table I. We emphasize that we chose those parameters since we believe that most *ns–3* users who are not experts in radio wave propagation will use the default parameters. However, since the default parameters do not yield reasonable results for four models, we ran additional experiments for those models with varied parameters. In particular, those models include the Fixed RSS model with a default receive signal strength of -150 dBm, and the Matrix model with a propagation loss of $1.8 \cdot 10^{308}$ dB. The Two Ray Ground model assumes two rays with one being reflected on the ground. By default all antennas reside directly on the ground, such that the model does not work and computes zero received power levels. Additionally, the random model by default draws a path loss of 1 dB every time. In order to actually utilize random values we ran an additional experiment featuring a normal distribution with a mean of 50 dB and a variance of 25 dB. We summarized the varied parameters in Table II.

Furthermore, we plug the Jakes and Nakagami fading models on top of the Friis path loss model. Finally, we end up with 15 different loss model configurations with 4 different transmission powers each for 3 different routing protocols. We executed each experiment 30 times for a duration of 5 minutes of simulated time resulting in a total of 5,400 experiments.

TABLE I
Summary of *ns–3* Default Parameters

Propagation Model	Default Parameters
Fixed RSS	Receive Signal Strength: -150 dBm
Matrix	Loss: $1.8 \cdot 10^{308}$ dB
Range	Maximum Range: 250 m
Random	Loss: Constant(1 dB)
COST-Hata	Center Frequency: 2.3 GHz Base Station Antenna Height: 50 m Mobile Station Antenna Height: 3 m Minimum Distance: 0.5 m
Friis	Wave Length: 58.25 mm System Loss: 1 Minimum Distance: 0.5 m
Log Distance	Exponent: 3 Reference Distance: 1 m Reference Loss: 46.67 dB
Three Log Distance	Distances: 1 m, 200 m, 500 m Exponents: 1.9, 3.8, 3.8 Reference Loss at 1 m: 46.67 dB
Two Ray Ground	Wave Length: 58.25 mm System Loss: 1 Minimum Distance: 0.5 m Height above Z: 0 m
Jakes	Rays per Path: 1 Oscillators per Ray: 4 Doppler Frequency: 0 Hz Distribution: Constant(1)
Nakagami	Distances: 80 m, 200 m Exponents: 1.5, 0.75, 0.75

TABLE II
We modified certain parameters for additional experiments where the default parameters could not yield any reasonable results.

Propagation Model	Modified Parameters
Fixed RSS	Receive Signal Strength: -50 dBm
Matrix	Loss: 50 dB
Random	Loss: $\mathcal{N}(50\,\text{dB}, 25\,\text{dB})$
Two Ray Ground	Wave Length: 58.25 mm System Loss: 1 Minimum Distance: 0.5 m Height above Z: 1 m

The measured network metric is the network efficiency as mentioned above. The computational complexity of the propagation model was computed by measuring the execution time spent on the function to calculate the propagation loss normalized to the total number of packets transmitted. This normalization was needed to ensure a valid comparison, since the different experiments show significant variation in total packets transmitted, with the expected difference in execution time as a function of packet transmission events. We measured the computational effort by reading the RDTSC register before and after entering the `CalcRxPower` function responsible for calculating the propagation loss and the resulting receive signal strength.

We averaged the results over the 30 independent runs and calculated the 95 % confidence intervals. All simulations were run on a workstation with a 2.8 GHz quad-core Xeon CPU and 6 GB of RAM, running Ubuntu 11.10, and *ns–3*.12.1.

V. Experimental Results

The simulation experiments showed a number of interesting results regarding the overall effect of the chosen path loss and fading models on the measured network efficiency. Further, we were able to show considerable variation in the computational complexity of the various models. These results are described in detail below.

A. On-Demand Routing Efficiency

Figure 2a shows the efficiency of the given network with AODV routing and different propagation loss models applied. As already pointed out, the Fixed RSS, Matrix, and Two Ray Ground models feature inappropriate default values, resulting in almost no throughput. However, with modified parameters all models generate non-zero throughput.

Although the default parameter for the random propagation loss model does not comply with the expectations for a random model, it is well suited to gain considerable efficiency (about 20 %). However, there is a slight increase in the mean efficiency if we actually feed the model with random numbers. We explain this by the following assumption: If two (or more) data streams simultaneously arrive at one station with the same receive power, neither data stream can be decoded. The introduction of randomness, however, allows the receiver to successfully receive the data stream with higher receive power, such that only one data stream is discarded.

We summarize that the network features, independent of the transmission power, an efficiency of 20 % to 30 % when an abstract propagation model is applied.

For the path loss models we observe considerable differences on varying transmission power. This complies with our regular observations in real world. With COST-Hata model and Log Distance model applied, the network features very low efficiency for low transmission power rates. However, with increasing transmission power the efficiency increases up to 40 % for COST-Hata, and even 70 % for the Log Distance model. On the other hand, when the Friis model is applied, the network achieves considerable efficiency for low transmission rates which is not improved by increasing transmission power. We explain this by the environmental assumptions of the three models: COST-Hata and Log Distance assume urban scenarios while Friis assumes a free space scenario. This means that when using the Friis model with small transmission power several nodes are reached. This results in short paths, but also in high contention. With the urban models, more transmission power is required to reach nodes behind imaginary walls.

We observe a different behavior on the Three Log Distance and the Two Ray Ground model: Too low transmission power here means that too few nodes are reached while too high transmission power induces too much contention. The best results are achieved by adapting the transmission power to a moderate level.

By combining the Friis model with the Jakes or Nakagami fading model we observe minor decrease or increase in the network efficiency. However, especially with lower transmission rates fading has a considerable negative impact on the throughput.

B. Proactive Routing Efficiency

When we compare the results for AODV to the results of DSDV and OLSR (see Figure 3a and Figure 4a), we make two general observations:

1) The results of DSDV and OLSR are very similar.
2) The results of AODV show significant difference to the results of the proactive routing protocols.

We explain this behavior by the general difference between on-demand routing and proactive routing. In AODV a route is only established if it is necessary in order to transmit a packet. In DSDV and OLSR the nodes initiate periodic routing table updates, and rely on the results for future conversation. Since in the *ns–3* implementation of the two routing protocols all nodes initiate those updates at the same time, heavy contention arises. This, however, results in incomplete routing tables (links are not discovered since the probe packets are lost) which results in suboptimal or missing routes, and finally in packet loss.

Since certain models introduce more interference than others do, this effect is more or less distinct when different propagation loss models are applied. For the abstract models Fixed RSS, Matrix, and Random the proactive routing protocols gain almost no throughput. With Fixed RSS or Matrix model applied, the proactive routing protocols suffer from the problem that all packets arrive at a receiver antenna at the same point in time with the same power level such that the receiver is unable to decode one of the packets. Also with the Random model applied, DSDV and OLSR cannot gain any significant throughput. We explain this by the fact that those routing protocols rely on established routes, but purely random propagation loss means links appear and disappear randomly.

On the other hand, DSDV and OLSR feature quite high

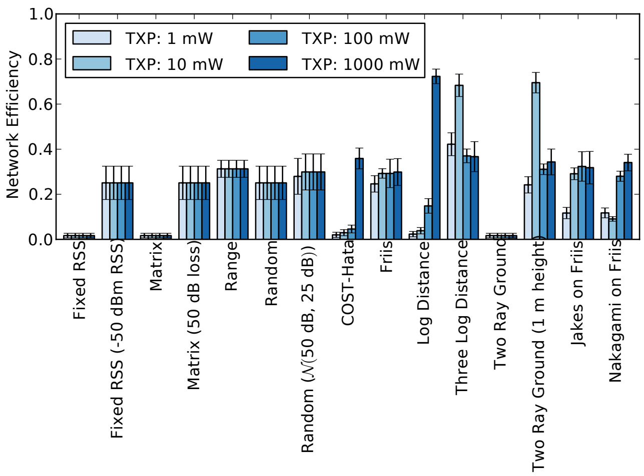

(a) Network efficiency (number of successfully received packets at destinations over the number of packets transmitted)

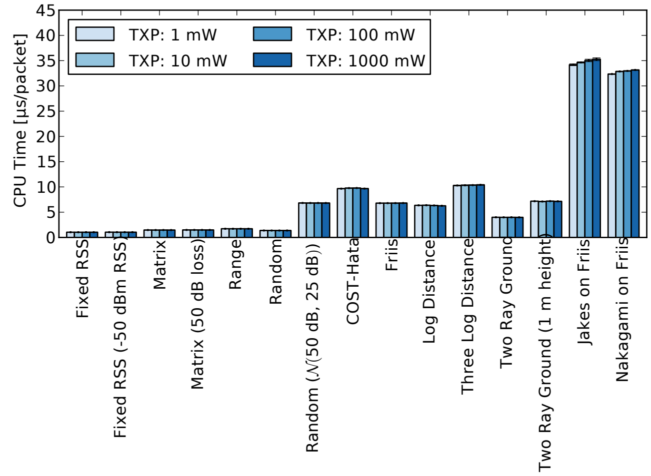

(b) Computational effort to compute the receive signal strength of one packet at all potential receivers by means of different propagation models

Fig. 2. Measurement results for the wireless network with AODV routing

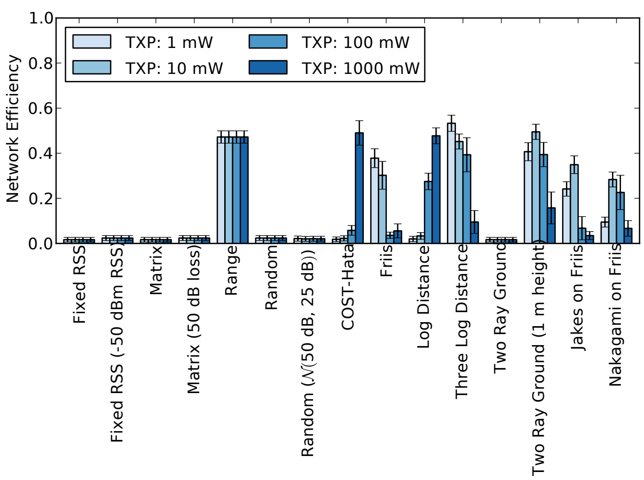

(a) Network efficiency (number of successfully received packets at destinations over the number of packets transmitted)

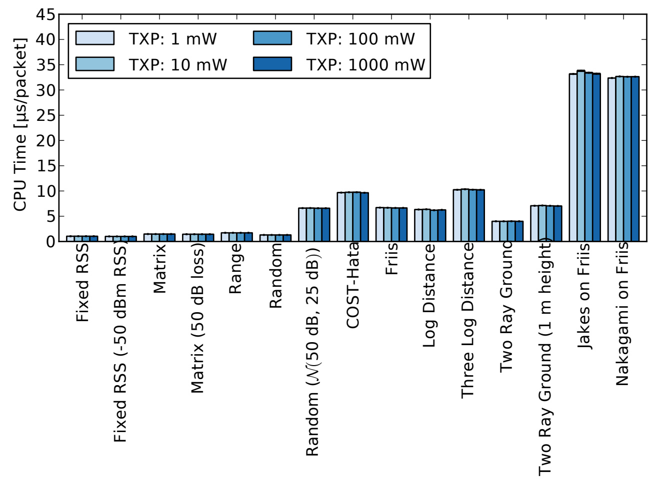

(b) Computational effort to compute the receive signal strength of one packet at all potential receivers by means of different propagation models

Fig. 3. Measurement results for the wireless network with DSDV routing

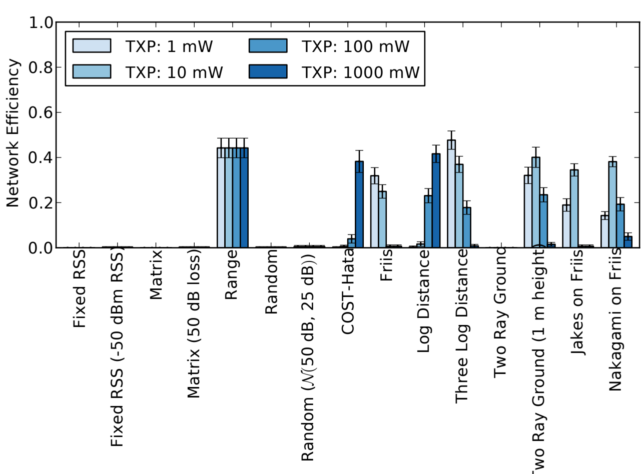

(a) Network efficiency (number of successfully received packets at destinations over the number of packets transmitted)

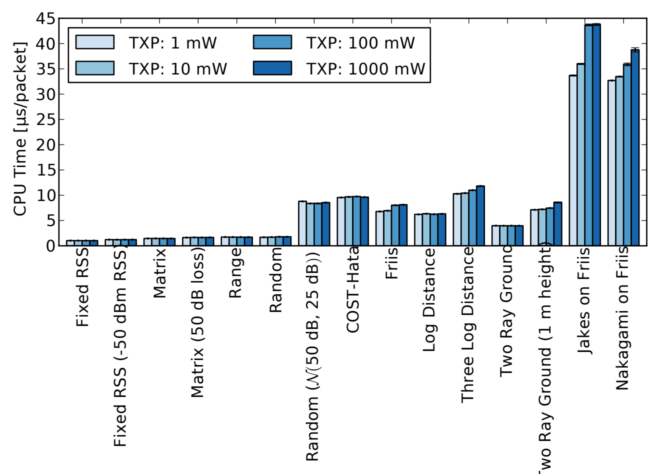

(b) Computational effort to compute the receive signal strength of one packet at all potential receivers by means of different propagation models

Fig. 4. Measurement results for the wireless network with OLSR routing

efficiency (40 %) when the Maximal Range model is used since only a subset of nodes receives the routing packets which reduces the contention. Furthermore, the links are stable since the maximal range does not vary.

The results for all three routing protocols are qualitatively equal when the COST-Hata, Log Distance, or Two Ray Ground model is applied. There are differences in the quantitative values. However, we cannot determine a superior routing protocol.

On switching from on-demand routing to proactive routing with Friis or Three Log Distance propagation loss model, we observe a shift in the optimal transmit power towards lower values. Again we explain this by the high contention induced by periodic routing table updates. A reduction of transmission power reduces the contention.

The fading models again appear to decrease the signal quality on average. However, with too high transmission power the proactive routing protocols benefit from this decrease since the contention is reduced.

C. Computational Effort

Figure 2b to Figure 4b illustrate the computational effort induced by the different propagation models. We observe almost no differences in the effort per packet on varying transmission power or routing protocol. We expected this behavior since the calculation of path loss and fading does not depend on either the transmission power or the packet content. Also the packet size is not considered by the fading calculations in *ns–3*. However, in fact there are slight differences in the complexity, in particular in the complexity of the calculations of the Jakes fading coefficients. We explain this by the implementation in *ns–3* which advantages certain traffic patterns over others by maintaining a list of senders. This list is accessed in order to enable stateful fading, i. e. correlation in the fading coefficients of subsequent transmissions. Furthermore, caching effects can influence the CPU times.

In general we observe the following distribution of CPU times: The abstract models determine the propagation loss within 1 µs to 2 µs. The only exception is the random model which takes about 5 µs to draw the pseudo random numbers while it finishes within 1 µs when a constant distribution is selected.

The path loss models take about 5 µs to 10 µs per packet to compute the path loss coefficients and are therefore as complex as the Random model. However, fading computations are heavily complex tasks, such that the Friis model in combination with Nakagami or Jakes fading takes more than 30 µs per packet on our evaluation system.

D. Routing Model Comparison

If we compare the different routing models, we observe that the best routing model depends on the propagation model. While measurements with most of the abstract propagation models result in the observation that AODV is the best routing model, DSDV performs best if COST-Hata is used. OLSR can benefit from fading effects on certain power levels, and achieves therefore the most throughput if configured adequately.

E. Summary of Experimental Results

We observed that the application of different propagation models results in different network efficiencies. Abstract simulation models do not lead to highly efficient networks as long as the receive power is independent on the distance between sender and receiver. This behavior is induced by high contention since every packet is spread over the entire simulation area. Further, modifying the transmission power for all nodes does not improve the situation since this adapts both signal and interference.

Applying an *ns–3* path loss model allows for highly efficient networks if the transmit power is adapted in such a way that a reasonable subset of nodes is reached by each packet. However, simply increasing of transmission power does not always improve the situation due to the consequently arising increase of interference for the other transmissions. If fading models are applied, additional loss in the signal strength has to be considered.

Again, we do not comment on the accuracy of those results since there is no correct value to compare to. We observed differences in the computational effort from 1 µs per packet for the computation of the receive signal strength in abstract propagation models to more than 30 µs per packet when path loss and fading models are applied.

VI. Conclusion

By means of the evaluation results we draw the following conclusions: Most of the abstract propagation loss models won't work at all with the default parameters of *ns–3*. We are working with the *ns–3* team in order to find more appropriate default parameters. However, adaption of the parameters can help to gain more realistic behavior. Nevertheless, high contention arises for abstract models when data streams are spread out over the whole simulation area without attenuation. We suggest using those models if sufficient certainty is available that the chosen parameters fit the given situation. For example, the Matrix model might be applied when the path loss between every pair of nodes has been measured and can be set as a parameter.

Nevertheless, for most purposes path loss models are the best choice. Although they are more computationally complex than the abstract models, the complexity is not that significant in the overall simulation. On the other hand, there are models accounting for the effects in different scenarios (urban to free-space) that feature more realistic behavior than abstract models do.

Fading models aim at increasing accuracy of the calculated propagation loss by taking into account the frequent changes to the communication environment. However, as depicted in the results this induces high computational effort. We therefore recommend the use of fading models only for scenarios where fading is important, such as in the evaluation of novel wireless technologies.

All in all there is no one-fits-all solution with respect to the proper choice of a propagation loss model. There are many different models, and the most appropriate model always depends on the research goal.

References

[1] T. S. Rappaport, *Wireless Communications, Principles and Practice*. Prentice Hall, 1996.

[2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level Measurements from an 802.11b Mesh Network," in *Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. ACM Press, 2004, pp. 121–132.

[3] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott, "Experimental Evaluation of Wireless Simulation Assumptions," in *Proceedings of the 7th ACM International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems*. ACM Press, 2004, pp. 78–82.

[4] D. Reddy, G. Riley, B. Larish, and Y. Chen, "Measuring and Explaining Differences in Wireless Simulation Models," in *Proceedings of the 14th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, 2006, pp. 275–282.

[5] Y. Zheng and D. Nicol, "Validation of Radio Channel Models Using an Anechoic Chamber," in *Proceedings of the 24th ACM/IEEE/SCS Workshop on Principles of Advanced and Distributed Simulation*, 2010, pp. 1–8.

[6] V. Abhayawardhana, I. Wassell, D. Crosby, M. Sellars, and M. Brown, "Comparison of Empirical Propagation Path Loss Models for Fixed Wireless Access Systems," in *Proceedings of the 61st IEEE Semiannual Vehicular Technology Conference*. IEEE, 2005, pp. 73–77.

[7] G. Durgin, T. S. Rappaport, and H. Xu, "Measurements and Models for Radio Path Loss and Penetration Loss In and Around Homes and Trees at 5.85 GHz," *IEEE Transactions on Communications*, vol. 46, no. 11, pp. 1484–1496, 1998.

[8] E. Damosso and L. M. Correia, *Digital Mobile Radio Towards Future Generation Systems: COST Action 231*. European Commission, Brussels, Belgium, 1999.

[9] H. T. Friis, "A Note on a Simple Transmission Formula," *Proceedings of the Institue of Radio Engineers*, vol. 34, no. 5, pp. 254–256, 1946.

[10] V. Erceg, L. Greenstein, S. Tjandra, S. Parkoff, A. Gupta, B. Kulic, A. Julius, and R. Bianchi, "An Empirically Based Path Loss Model for Wireless Channels in Suburban Environments," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 7, pp. 1205–1211, 1999.

[11] Y. R. Zheng and C. Xiao, "Simulation Models with Correct Statistical Properties for Rayleigh Fading Channels," *IEEE Transactions on Communications*, vol. 51, no. 6, pp. 920–928, 2003.

[12] M. Nakagami, "The m-Distribution, A General Formula of Intensity Distribution of Rapid Fading," *Statistical Methods in Radio Wave Propagation*, pp. 3–36, 1960.