A Site-specific Indoor Link Model for Realistic Wireless Network Simulations

Mustafa Al-Bado   Ruben Merz   Cigdem Sengul   Anja Feldmann

Deutsche Telekom Laboratories at TU-Berlin

{mustafa, anja}@net.t-labs.tu-berlin.de, {ruben.merz, cigdem.sengul}@telekom.de

ABSTRACT

Compared to testbeds, the efficiency and accuracy of wireless networking simulations are constantly questioned by the network community. It is widely accepted today that the current network simulators are not able to fully represent the real wireless characteristics, especially at the physical (PHY) layer. This affects the trustability of simulation-based performance evaluations. On the other hand, testbed experiments require taking a tedious and time-consuming implementation path. This path could be significantly reduced by using realistic network simulators as a first step to test novel algorithms or protocols. Therefore, we took on the challenge of representing the link characteristics of the indoor testbed of the Berlin Open Wireless Lab (BOWL) project in the ns-3 simulator. Our extensive measurements study of the link characteristics, namely received signal strength (RSS), frame detection ratio (FDR) and frame error ratio (FER), produced several guidelines for modeling our testbed with satisfying accuracy in the simulator. More importantly, the proposed empirical models take into account several crucial properties related to the radio hardware and the environment, which are shown to have a significant impact on the simulation accuracy. We validate our model against testbed results and show that, unlike the existing models in ns-3, our model shows high agreement with the measurement results for any pair of nodes in the testbed.

1. INTRODUCTION

Wireless network testbed deployments have become increasingly important for the development and evaluation of network protocols and algorithms. However, simulation is still necessary and extremely valuable. A study that looked at papers presented in the MobiHoc conference between 2001 and 2005 showed that more than 75% of them used a simulation tool for their evaluations [20]. This is not surprising as most testbeds require significant resources, a tedious and time-consuming system-level implementation, which can significantly impede their use. And the unpredictable physical

medium and complex interactions between various system components in testbeds is a real challenge, especially when it comes to troubleshooting them. In addition, testbeds carry the limitations of current hardware and deployment size. Thanks to simulators, these limitations can be overcome. Furthermore, simulation studies can also be used to debug testbeds. For instance, trace-driven simulation has been used to detect and identify network faults and to do a "what-if" analysis [26].

On the downside, the wide use of simulations for performance evaluation, in particular *packet-level* simulations, is a huge concern due to the lack of accuracy in wireless models, especially at the physical layer (PHY). The difficulty and high complexity to design PHY models stems from (1) the number of various factors that need to be considered (e.g., adjacent channel interference [9] or frame error rate with overlapping packet transmissions) and (2) the timescale difference between events at the physical layer and events at the medium access control (MAC) layer and above [25]. Another challenge is the representation of radio wave propagation. In the context of a packet-based wireless network simulator, a so-called *propagation model* is used to obtain, for each transmitted packet, the received signal strength (RSS) at any potential receiver node. Our measurement study (see Figure 4 in Section 5.1) and previous works [8, 12, 16, 18] show that channel propagation models can be one of the main reasons for the inaccuracy of wireless network simulation.

In this paper, we take on these challenges to create an accurate link model for a packet-level simulator of an IEEE 802.11 wireless network. Our underlying objective is to obtain a realistic representation of the BOWL indoor wireless network testbed [23] in ns-3, a well-known packet-level network simulator [3]. Therefore, contrary to the stochastic propagation models already available in ns-3, we develop a measurement-based *site-specific* link model [7]. The choice of a site-specific model is also motivated by the empirical validation of simulator propagation models (e.g., free-space and two-ray models, see Section 2 in [22]), which shows that although simple stochastic radio propagation models might predict the upper-layer routing performance, the results are sensitive to the model's parameters. Therefore, these models should be properly calibrated before they are used to represent a given topology [22, 16].

In addition to the propagation model, we also develop a site-specific frame detection model and frame error model. Frame detection is the first step for the correct reception of a packet, which consists of detecting the presence of a packet on the medium and acquiring the timing of the frame. A

frame needs to be properly detected before its payload can be decoded. Following detection, if the payload is successfully recovered, then the frame is handed over to the MAC layer. Otherwise, a frame error is declared. It must be noted that our frame error model is for single-user cases (i.e., only one transmitter is active at a time). This is essential to understand the effects stemming mainly from our environment and radio specifics. We leave the study of interference effects from multiple transmitters at a given time as future work.

Because of the large number of factors that can affect propagation and frame detection, for instance the carrier frequency or the data-rate of the transmitter, we ran an extensive measurement study covering these factors. This way, we obtained statistically meaningful measurements on a per-link basis. We also use this study to identify the factors that have a major effect on the output of the above-mentioned models. We present our learnings as guidelines for collecting measurements in a systematic fashion to feed as input to our simulation model.

Our work underlines the need for a more advanced channel propagation model that respects node deployment, environmental conditions and device specifications. For instance, our results confirmed that distance does not have a major influence on the communication between pair of nodes [24] but we also found that the carrier frequency selection has significant effects on frame detection. We summarize our contributions as:

- A novel, site-specific, link model which comprises a RSS model, a frame detection ratio (FDR) model and a frame error ratio (FER) model in an IEEE 802.11a wireless network testbed. As already explained, the FER model is single-user specific.

- A reproducible methodology to run, collect and import measurements to create a site-specific model.

- An extensive measurement study of propagation and frame reception in an IEEE 802.11a wireless network testbed that quantifies which are the most important factors for modeling the aforementioned quantities.

The rest of the paper is outlined as follows. In Section 2, we review the related work. We present our measurement and modeling methodology in Section 3. Section 4 describes our testbed and measurement setup, and in Section 5, the results from our extensive measurement study are used to identify the performance evaluation gaps between simulation and testbed environments. Based on this study, in Section 6, we propose a new site-specific propagation model and validate our model in Section 7. Finally, we conclude in Section 8.

2. RELATED WORK

There are a number of simulation tools for wireless networking, among which the most well-known are ns-2 [2], ns-3 [3], Opnet [4] and QualNet [5]. Several measurement studies [8, 12, 16] investigated the accuracy of typical channel propagation models in these simulators, and showed that these models do not accurately represent the reality of wireless network characteristics. In [22], it was shown that unrealistic radio propagation models might lead to overestimation of performance at the routing layer due to the unrealistic assumption that more wireless links are present than

reality. Furthermore, the comparison of two routing protocols may yield completely different results (i.e., protocol x outperforms protocol y in the simulator but not in reality). These studies encouraged the community to focus more on bridging this gap between reality and simulation models. In the following, we first discuss the existing propagation models in the aforementioned simulators and then, present the literature towards more realistic models.

2.1 Current Propagation Models

The typical propagation models found in simulators are the Friis free space model, the two-ray ground model and a shadowing model [13, 24, 27, 2]. In the free space model, the average signal strength fades with distance based on a power-law model. The two-ray ground model considers both the direct and ground-reflected propagation paths between the transmitter and the receiver. The shadowing model can be used to add a stochastic component to account for indoor abstractions and outdoor shadowing. However, this model ignores the correlations in terms of shadowing effects between two locations that are close to each other.

At the time of writing this paper, the ns-3 development branch contained also two log-distance models (LogDist and ThreeLogDist) applicable to indoor propagation modeling. The ThreeLogDist model is the same as the LogDist model except that it uses three distance fields: near, middle and far with different exponents. In addition, ns-3 supports a constant loss model, the COST-Hata model that models the urban area propagation loss (see Chapter 4 from [1]), a fast-fading Nakagami model for modeling multi-path effects [27] (which can also be used for Rayleigh fading), and the so-called Jakes propagation model for Rayleigh fading in mobile environments [17].

None of the above models are site-specific but stochastic models. Hence, these models are not tied to any environment but use statistical distributions of channel impulse responses. On the other hand, site-specific models allow predicting channel response for a specific transmitter location, receiver location and the environment. The commercial Qualnet simulator supports two well-known site-specific models: the irregular terrain model (Longley-Rice) and the commercially available The Terrain Integrated Rough Earth Model (TIREM). A popular method to predict the site-specific propagation characteristics is ray tracing. In [19], three parameters of the radio path: the path loss (dB), the Rician K factor, and the RMS delay spread are conjectured to be sufficient to capture the properties of a channel response. While the path loss and the Rician K factor determine the narrowband fading distribution, RMS delay spread measures the pulse dispersion of the channel. In contrast, in our study, in addition to channel propagation, we also take into account the hardware-specific properties and the effect of proprietary algorithms on frame detection and frame error rates.

2.2 Towards More Realistic Wireless Network Simulation

Several works proposed calibrating the simulator models to fit the real network behavior. In [16], the authors calibrated the shadowing model of ns-2. The results show that the calibrated model results deviated only 0.3% in terms of packet delivery ratio. Similar results were also obtained for network topology. However, the accuracy of the calibrated

model is limited as a fixed transmit rate is used, whereas the communication rate in the real network changes dynamically. In [18], two different methods are developed to determine the propagation behavior. In the first method, RSS values between each pair of nodes in the testbed are collected. In the second method, measurements are collected by moving two nodes to different locations in the building and recording the RSS values. The authors model propagation combining path loss and log-normal shadowing and test it under ns-2. Using least square linear regression on the collected RSS values, the path loss exponent is estimated. The shadowing model parameter is estimated by fitting a normal distribution on the error values in the regression. Using the same measurements, the authors also computed the probability that one node defers to the transmission of the another node. This deferral probability captures the effect of transmission and carrier sense range between the two nodes. Their results show that measurement-based models are able to achieve higher accuracy, whereas pure simulation models either underestimate or overestimate the capacity. However, the study does not take into account the effect of radio-specific properties, such as preamble detection, capture [30], adaptive noise immunity or weak signal detection mechanisms [15, 31]. These issues play an important role as shown in [8, 12, 30] and based on our own results in an indoor environment. In [29], the capture effects are also taken into account and *average* RSS for *successfully* received packets were used for two PHY models for reception and deferral. However, the information from packets that are received but errored are ignored. In [28, 11], even though authors proposed feasible models for rural environments, these models do not apply to indoor environments, where the conditions are completely different and more severe.

As a different effort to bring more detail to simulators, in [9], the need for a spectrum-aware channel modeling is argued. The authors calculated the Effective SNR (ESNR) that represents the performance of OFDM systems in which each sub-carrier experiences a different SNR and feed this value to the error model to infer whether a packet is in error or can be correctly received. Using the Friis propagation model, the authors show that their model decreases capacity errors when fading is considered in addition to path loss. There has also been work on the validation of MAC layer simulation models with testbed. A notable example is [10].

Several works mention channel, data rate and transmit power selections, as well as the environmental factors such as multi-path fading, as essential parameters of a realistic propagation model [14, 8]. To the best of our knowledge, there is yet no model that takes all these parameters into account. In this work, we try to address this problem through a measurement-based study, which incorporates measurement both from successfully received and errored packets. To this end, we first investigate the affect of these parameters in our indoor wireless testbed, and then propose a model based on the two parameters, channel and data-rate, that are found to have the highest impact.

3. MEASUREMENT AND MODELING METHODOLOGY

One of the goals of our research is to identify the influence from various factors, such as transmit rate, transmit power and channel selection on the characteristics of a wireless link. Understanding the effect of each factor allows us to better reduce the gap between simulation and real-world. Therefore, we follow a step-by-step approach and evaluate:

1. **The impact of radio-specific optimization algorithms.** In our network, we deployed Atheros radios. These radios use several optimization algorithms, such as adaptive noise immunity (ANI) [15, 31], to improve the device performance. Therefore, we first investigate the effects of the underlying radio properties and its effect on measurements.

2. **RSS variations in a stable environment.** In this step, the relationship between RSS and distance, and the distribution of RSS over time are investigated. For this particular step, in order to create a stable environment, all of the related measurements were run at night in a fixed network.

3. **The impact of channel selection.** We ran several measurements in different channels. To the best of our knowledge, none of the current simulation models take into account the performance diversity due to the use of different wireless channels.

4. **The impact of physical layer data-rate.** We study the relationship between different physical layer data rates and RSS, FDR and FER.

5. **The impact of a dynamic environment.** For this study, we followed two approaches. First, we extended our night measurement sessions to different days of the week and tried to capture how the changes during the day (e.g., due to moving of obstacles) affect communication during the night, when no such changes occur. Using these results, we were able to understand the impact of gradual environmental changes. Second, we conducted our experiments during the work days, where there are more dynamics in the environment.

In each of these steps, we compared the measurement results with the simulation results using ns-3 to see whether the existing models match the real-world results. Based on the observed shortcomings, we proposed a new site-specific propagation model. In the following sections, we first describe our testbed and simulation environment and then present our measurement study in detail.

4. SYSTEM DESCRIPTION

We conducted our measurement study using the BOWL indoor testbed [23] at Deutsche Telekom Laboratories, in Berlin, Germany. Our indoor testbed consists currently of nine nodes, five of which are deployed in one floor, and the remaining four reside on the floor above (see Figure 1). The host names indicate the floor and the node numbers. For instance, `tel-16-2` is the second node on the 16th floor. Each node is built around a Gatework Avila GW2348-4 platform with 64 Mbyte of RAM, an Intel XScale IXP425 533 MHz processor (ARM architecture) and two Wistron CM9 miniPCI IEEE 802.11abg wireless network interface cards (NIC). Each wireless NIC is connected to an 8 dBi gain omnidirectional antenna (with a 2 dB loss because of cabling). The chipset on the wireless NIC is an Atheros

Figure 1: The BOWL indoor testbed spans two floors. There are five nodes on the 16th floor (top picture) and four nodes on the 17th floor (bottom picture). The host names are indicated with the convention: `tel-floor-node`.

AR5213A [1]. The operating system on our nodes is OpenWrt 8.09.2 with Linux kernel 2.6.26.8. The wireless driver is the OpenWrt maintained version, with revision number 3314 with HAL 20090508. Additionally, all nodes have a dedicated Ethernet management interface with at least a 100 Mbit/s capacity that is also used to collect measurement results on a central server.

In all our experiments, we used IEEE 802.11a [6] as the underlying physical layer. In addition, only one wireless interface is active at any given node. The transmitting nodes were set up in *ahdemo mode*[2], which disables the transmission of all management frames (e.g. beacon frames). This allows us to run more controlled experiments. The receiving nodes were set up in *monitor mode*[3], which allows us to gather the link layer information (e.g., RSS) using the so-called radiotap header[4]. Most of the measurements were conducted on channel 44 at 5.22 GHz except for the experiments of Section 5.3 where we looked at multiple channels. In any case, before each experiment, we always checked that the channel was free of other interfering transmitters.

In all experiments, there is always only one transmitting node at a given time. In all cases, each experiment comprises several measurement runs for each transmitting node. Because the accumulated duration of all runs for a given node can be on the order of an hour or more, a given transmitting node does not perform all runs in a row. After each measurement run, another transmitting node was chosen. All measurement runs for a given experiment have the same setup. Tcpdump was used to collect traces at each receiver. These traces were redirected to and stored on a central server using the Ethernet management interface.

For all experiments, CBR UDP broadcast traffic with a packet size of 512 B is generated using the Iperf tool. Because of the broadcast transmission, no RTS/CTS control

frames are transmitted. We validated that the packet size has no effect on the RSS and the FDR models. However, it is expected to have an effect on the FER model and an extension of the model for different packet sizes is planned but left for future work.

For simulations, we use ns-3 version 3.6[5]. The simulated topology represents our node deployment shown in Figure 1. As in the measurements, simulations also generate CBR UDP traffic with a payload of 512 B sent to the broadcast address. The packet generation rate is always set to be comparable with the Iperf settings. The transmission of beacon management frames is also disabled. At the physical layer, we use the `YansWifiPhy` with IEEE 802.11a settings. We set the noise floor value to the one captured in our measurement traces. In addition, we patched the ns-3 `PcapWriter` class to include the packets that do not pass the FCS check in the radiotap header trace, which the simulator can generate. We also corrected a bug in the noise floor calculation so that receiver noise figure was added correctly.

5. THE GAP BETWEEN REALITY AND SIMULATION

In this section, we follow the approach outlined in Section 3. We also compare our measurements to the typical propagation models in the ns-3 [3] network simulator (see Section 2.1). We use three metrics: RSS, frame detection ratio (FDR) and frame error ratio (FER). FDR represents the ratio of all the detected frames (by the radio signal detection unit i.e., including the frames with errors) to the transmitted packets. FER is the ratio of frames with errors to FDR. We decide that a frame has an error when it was received by the radio interface but did not pass the Cyclic Redundancy Check (CRC).

5.1 The Impact of Radio-Specific Optimization Algorithms

RSS measurement collection is driver and environment specific. Hence, as a first step, we investigate the relationship between RSS and the Atheros-specific optimization algorithms, and the current radio configuration. The optimizations we take into account are ANI (Adaptive Noise Immunity) and weak signal detection [15, 31]. To understand the impact of these optimizations, we ran the same scenario when these algorithms are turned on and when they are turned off. This scenario consists of 4 sessions. In each session, nodes take turn to broadcast packets with a rate of 4 Mb/s for 5 minutes. As mentioned before, there was always only one sender at a given time. In the experiments, the nodes were using the 6 Mb/s data rate and 17 dBm transmit power. We also ran the experiments at channels 44, 48, 52 and 56 and observed the same general behavior.

Figure 2 shows the FDR per RSS value (i.e. number of received frames at a particular RSS value divided by the total number of received packets) for two receivers on the 17th floor with and without optimizations. The total FDR between two nodes and the distance between these two nodes are given in the label. The results show that, for both cases when the optimizations were turned on and off, the range of RSS values of the received packets overlap. The main difference lies in the FDR per RSS and the total FDR. For
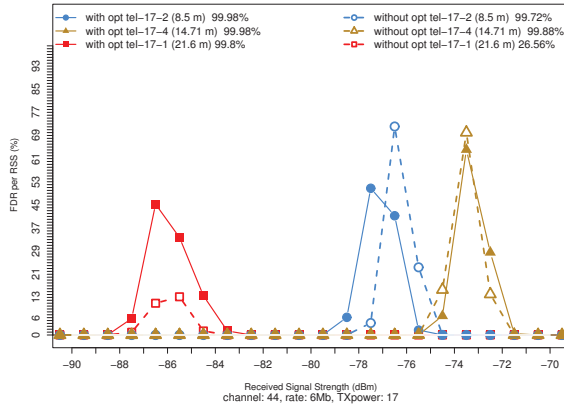
---

[1] `http://www.wifi-stock.com/file/cm9_user_manual.pdf`

[2] `http://madwifi-project.org/wiki/UserDocs/AhdemoInterface`

[3] `http://madwifi-project.org/wiki/UserDocs/MonitorModeInterface`

[4] `http://madwifi-project.org/wiki/DevDocs/RadiotapHeader`

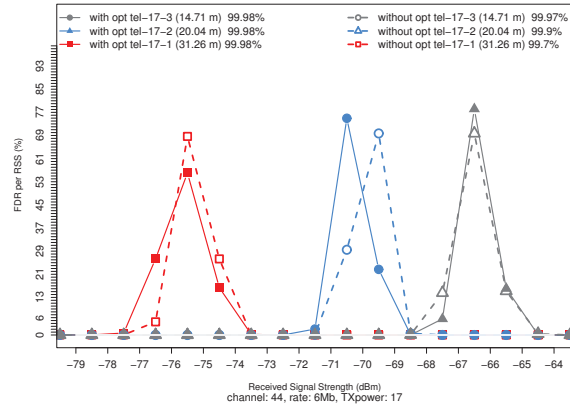[5]The latest release 3.10 does not contain any changes that affect our work.

(a) FDR per RSS from all senders at `tel-17-3`

(b) FDR per RSS from all senders at `tel-17-4`

Figure 2: FDR vs. RSS at two nodes in 17^{th} floor with and without using the Atheros-specific optimizations. For both cases, the range of RSS values overlap.
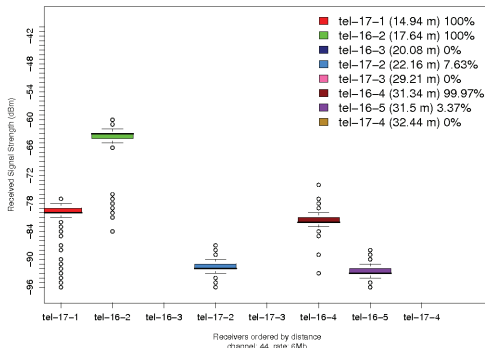

Figure 3: RSS variation at different nodes in a stable setting (sender). The results show no direct relationship with RSS and distance.
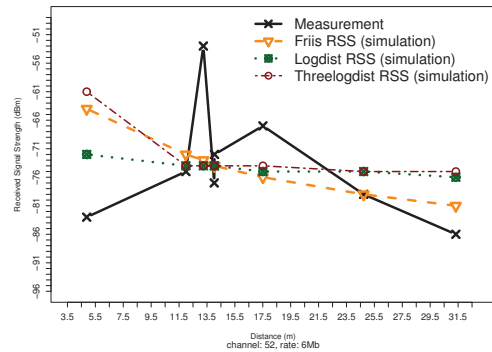

Figure 4: Comparison of RSS-distance relationship with simulator models. Calibration of models does not help capture the variability seen in practice.

instance, for the link between `tel-17-1` and `tel-17-3`, the total FDR is 26 % without optimization, and 99 % with optimization. Hence, the role of the optimization algorithms is to increase the probability of receiving a packet by adapting the noise thresholds, mitigating the interference or capturing the packet at a low SNR. In all of the following experiments, we turned these optimizations on.

5.2 RSS Variations in a Stable Environment

To see the behavior in a stable environment, we ran our first experiments at night when the lab is empty of people. We first investigated the relationship between RSS and distance (i.e. the length of a direct line between two points). To this end, we looked at the RSS values at all receivers for a particular sender. Figure 3 presents a boxplot for RSS values at different receivers (ordered by their distance) when the sender is `tel-16-1`. The figure shows that the variation of RSS and FDR (as seen in the labels in the graph) are not tied to distance. For instance, `tel-17-3` did not receive any packets, whereas `tel-16-4` received almost all packets (99.9% FDR) even though it is further away from `tel-16-1`.

One explanation might be that the sender and the receiver are residing in different floors. However, looking at `tel-16-5` and `tel-17-1`, we see that while the former had a FDR of 3.36%, the latter received almost all packets. On the other hand, in this case, `tel-17-1` is closer in distance to `tel-16-1`. Hence, it is not immediately clear how much affect can be attributed to these different factors.

In the next experiment, we look at only one of the two senders that was connected to all other nodes for the previous experiment (i.e., all receivers had a strictly positive FDR) and compare its measurement results with three typical propagation models found in ns-3: Friis, LogDist and ThreeLogDist. To represent the same conditions, we used nonlinear least squares (NLS) to find an appropriate fitting for each parameter of the simulation models based on our measurement data. Figure 4 shows that none of the simulation models is able to capture the measurement results. In contrast to earlier work that shows better fits with calibrated models, we conclude, in our case, NLS fitting of measurement data is not sufficient to represent RSS variations.
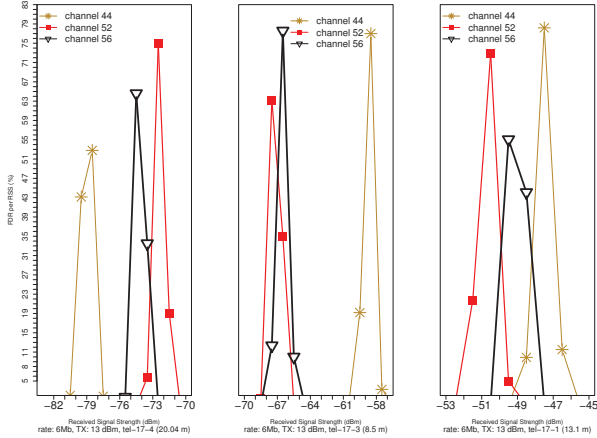
5.3 The Impact of Channel Selection

Figure 5: RSS values vary based on the wireless channel in use, requiring that each channel should be modeled separately in simulator.

Figure 6: RSS range does not change with rate.

Figure 7: FDR distribution varies for different data rates at low RSS values.

In this section, we try to answer the following questions: (1) Do all frequency channels have the same properties and performance? (2) If we assume that link l_1 has better performance on channel x compared to channel y, does this imply a better performance for link l_2?

To answer these questions, we ran our experiments on a single floor with one sender at night. The experiment consists of 4 sessions. Each session consists of 3 sub-sessions, in which the sender broadcasts packets at 4 Mb/s for 5 minutes using one of the following channels: 44, 52 and 56. Between sequential sub-sessions, the sender pauses for 10 seconds to change the channel. These three channels were chosen because both they and their overlapping channels were unoccupied. All of the chosen channels had the same noise level of -96 dBm. The nodes had the following setup: data rate was 6 Mb/s and transmit power was 17 dB.

The results depicted in Figure 5 answer the first question and show that different channels have significantly different RSS performance. Figure 5 also shows the RSS range for the different links for the chosen channels (when the sender was `tel-17-2`). This answers our second question: we observe that there is no relative relation between channels. For instance, while channel 44 is the best channel for `tel-17-3` (i.e, the range of RSS values at this node is the highest for this channel), this corresponds to the worst channel for `tel-17-4`. Therefore, to improve the realism of propagation models, it is necessary to model each channel separately.

5.4 The Impact of Data Rates

Similar to the channel experiments in Section 5.3, these experiments were also run on the 17th floor at night. Three sessions were run, where each session consists of eight sub-sessions. In each sub-session, the sender broadcasts packets at 4 Mb/s for 5 minutes using one of the IEEE 802.11a data rates. Between sequential sub-sessions, the sender pauses for 10 seconds to change to a different data rate. The nodes have the following setup: channel is 44 and transmit power is set to 13 dBm. The transmit power was chosen to be 13 dBm because the Atheros calibration table states that this is the highest power the node can transmit at channel 44 with the highest data rate of 54 Mb/s.
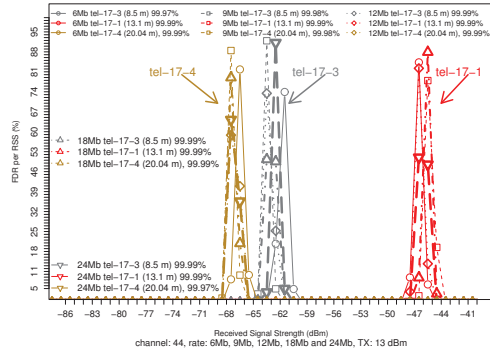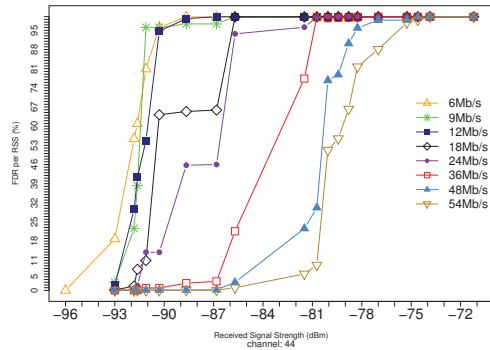
Figure 6 shows that the RSS range and FDR do not change with rate (only 5 rates were depicted for the clarity of the presentation). However, note that the FDR is high for all measurements. In some measurements, we observed that this does not always hold for low RSS values. Therefore, we ran a second experiment, which is explained next.

In the next experiment, we chose the two nodes which have full connectivity to others as senders. The experiment was run during the weekend, and hence, conditions can be considered stable. Each session consisted of four sub-sessions. To create varying RSS conditions, the sender broadcasts at six different transmit powers: 0, 5, 7, 9, 11 and 13 dBm. At each transmit power the sender broadcasts packets for four minutes for each of the data rates. All nodes were again at channel 44. The sender tried to saturate the medium based on the current data rate. However, the maximum rate we could achieve was 18 Mb/s because of the limitation in capturing packets in monitor mode with our hardware.

The experiment was able to produce different RSS ranges. Figure 7 shows that, at a given RSS value range, the FDR differs based on the data rate. We have repeatedly observed this behavior with Atheros radios, and plan to do more controlled tests with a signal attenuator and oscilloscope to validate our observations. Figure 8 shows the relationship between RSS and FER for each data rate. As expected, FER stays high for higher data rates until a threshold RSS value. These results are used in Section 7 to validate the standard FER calculation in the simulator.
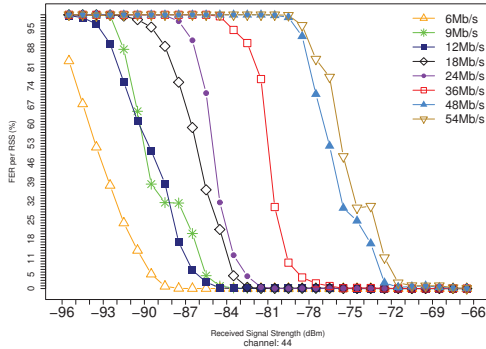
Figure 8: FER distribution based on data rate. FER stays high for higher data rates until a threshold RSS value.

Figure 10: The FDR at a given receiver on a given day, when `tel-16-1` is the sender. At certain days the FDR shows dramatic changes, as also hinted by Figure 9.
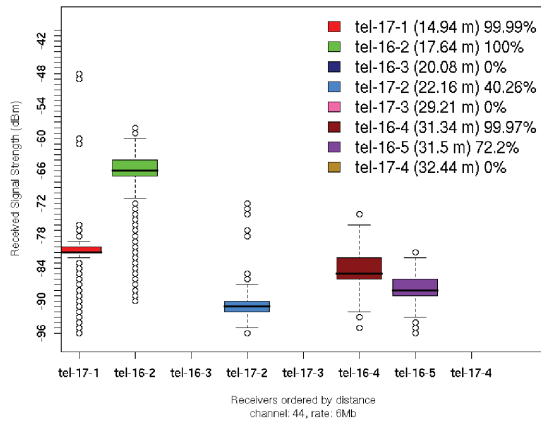

Figure 9: RSS variation at different nodes for a week, when `tel-16-1` is the sender. While some receivers experienced no change in RSS range, others show high variations.

Figure 11: RSS variation at different nodes during day time, when `tel-16-1` is the sender. RSS and FDR show higher variability - either RSS range increases, or the average FDR decreases.

5.5 The Impact of a Dynamic Environment

To understand the impact of changes during the day on our night experiments, we ran the first scenario in Section 5.1 for one week. Figure 9 shows a boxplot for RSS at different receivers when `tel-16-1` is the sender. We notice that compared to the results depicted in Figure 3, some links did not exhibit any changes either in their RSS range or FDR (e.g., `tel-17-1` and `tel-16-3`). For others, either the RSS value range or the FDR has changed, (e.g. the cases of `tel-16-2` and `tel-17-2`). To see things more clearly, we plot the FDR for each links in Figure 10 for each day. We notice that at certain days there are dramatic variations. For instance, FDR at `tel-16-5` on the second and the sixth days fluctuated from almost 100% to 3%. However, note that the RSS range of this node is also close to the noise level. These results show that while weak links need to be measured for a long term to really understand their behavior, links that are strong in terms of RSS exhibit high stability and might require less number of measurements.

Finally, we ran experiments during day time from 8am to 8pm - a total of 13 measurement sessions of 50 minutes. Fig-
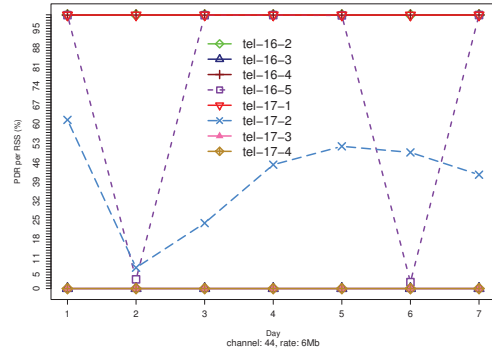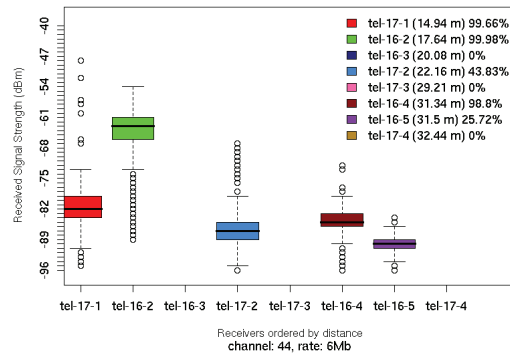
ure 11 shows, in this case, the RSS range becomes wider and the FDR values decrease compared to night experiments for both the one-night and week-long cases. Based on the experiment results, we classified the collected traces to three categories: First category is *relatively stable environment*, which occurs early morning and at the end of the work day. Figure 12 shows that if we look at only the first two measurements, we see that the RSS variation is much more stable (i.e., the RSS range is less than 5 dB). Second category is *medium dynamic environment*, where the distribution of RSS is between 5 and 10 dB. Finally, the third category is *highly dynamic environment*. In the last category, the RSS range reaches up to 15 dB variations. To the best of our knowledge, current simulation models do not take into account such time-of-the-day effects, which might lead to high inaccuracy in performance evaluation.

6. A SITE-SPECIFIC INDOOR LINK MODEL FOR NS-3

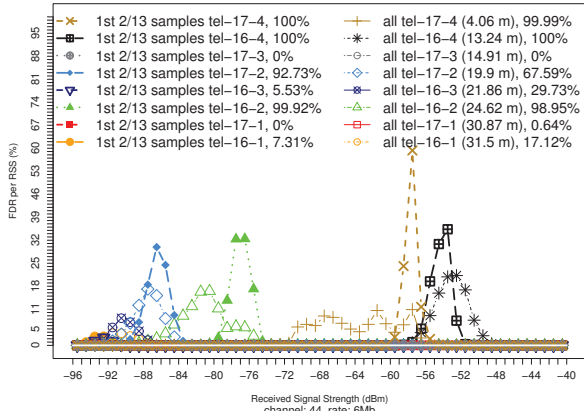In this section, we use the results of our measurement

Figure 12: The frame detection ratio vs. RSS for daytime measurements, when `tel-16-5` is the sender. When we look at the first two measurements out of 13, we see RSS variation is more stable. This represents a stable environment at early morning.

study to build the BOWL indoor model (BIM)[6], a measurement-based and site-specific model of the BOWL indoor network for the ns-3 simulator. BIM comprises a radio propagation model, a FDR model and a FER model. BIM is added to ns-3 by writing a new propagation model and adding three minor modifications to the `YansWifiPhy` class which simulates the IEEE 802.11a physical layer in ns-3 [21].

The modifications added to `YansWifiPhy` are:

1. Ability to feed measurement-based RSS distributions.

2. Per-rate FDR threshold that models the radio signal detection unit behavior.

3. A FER model that takes a particular frame RSS and modulation as parameters.

4. Transmit power limitations that represent Atheros hardware specific limitations of the transmission power for high data rates (see Section 5.4).

5. Recording of dropped frames in the radiotap trace.

The new radio propagation model is based on per-link RSS distributions. For each link and for each combination of data rate and channel frequency, we process our measurements to build an empirical cumulative distribution function (ECDF) of the RSS. Then, each ECDF is fed to a file-based database which is read by the ns-3 simulator at the beginning of the simulation. The creation of the per-link ECDF is performed in two steps:

• For each combination of link, data rate and channel, we first parse the RSS value of each recorded frame. Note that our data includes only the RSS of detected frames (i.e., excludes frames that were dropped by the radio signal detection unit). To represent the effect of the dropped frames in each RSS distribution, we post-process our traces as follows. Because we know exactly how many frames are transmitted for a measurement

---

http://code.nsnam.org/rmerz/ns-3-bim-simutools2011

Table 1: FDR threshold and α per data rate

Data rate Mbit/s	6	9	12	18	24	36	48	54
α dB	0	1	3	4	4	6	8	9
FDR thresh. dBm	-96	-95	-93	-92	-92	-90	-88	-87

run, we easily deduce the number of frames dropped by the radio signal detection unit. We then add "shadow" frames to our measurement traces with an RSS value equal to the noise floor plus a *data-rate specific* correction factor α. The sum of the noise floor plus α is equal to the FDR threshold (i.e., the threshold RSS value where FDR per RSS is greater than 0) for a particular data rate. The values of the FDR threshold for each data rate are given in Table 1.(The values of α were obtained with the measurements represented in Figure 7).

• The ECDF is then created based on the RSS values using a bin size of one and written to a file. This ECDF file is then added to the database that contains RSS distributions for all combination of links, data rate and channel.

In the simulator, the RSS value for a particular received frame is obtained simply by sampling the RSS distribution of the given link with the corresponding data rate and channel used for this particular frame.

While to the best of our knowledge, there is no FDR model for ns-3, there exists an analytical model for FER [21]. In this paper, we propose measurement-based FDR and FER models. In contrast to our per-link propagation model, the FDR and the FER models are obtained by aggregating measurement traces for all links in the network. This implies we use the same FDR and FER model for each node in the simulation. The FDR model is a simple threshold that is applied to each frame at a receiver (see Table 1). Frames with RSS below the FDR threshold are dropped and considered as interference. Frames equal to or higher than the threshold are considered for potential reception and passed to the FER model. In this model, for a particular frame, the RSS and data rate of this frame is used to look up the corresponding frame error probability p. This probability is then compared to a uniformly sampled random variable q ($0 \leq q \leq 1$). If $q > p$, the frame is successfully received. Else, the frame contains an error and is dropped.

7. VALIDATION OF BIM WITH NS-3

In this section, we validate the components of our site-specific link model, namely propagation, FDR and FER models, with measurements. The simulation scenario is as described in Section 4. For the measurements, we use the so-called stable environment scenario from Section 5.2. In particular, we use a transmission power of 13 dBm, channel 44 and a data rate of 6 or 54 Mbit/s. We observed similar results for other rates and channel selections.

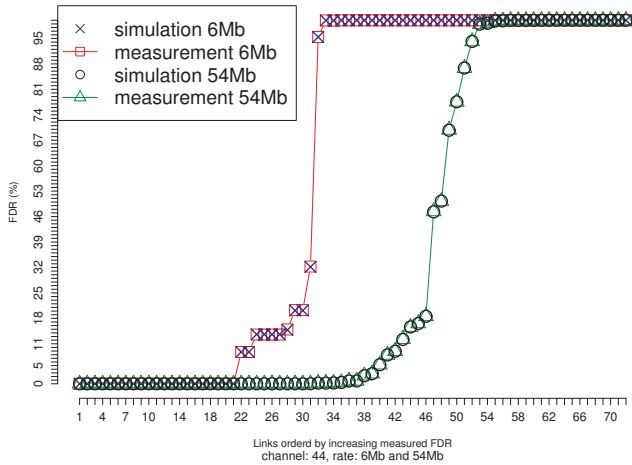To validate the propagation model, we compare, for each

185

Figure 13: Comparison of the measurement and simulation-based FDR at 54 Mb/s for all links. Simulations are performed with the BIM model only as there is no FDR model in ns-3. The FDR model shows an excellent agreement. Indeed even with a single FDR threshold per data-rate, the RSS distribution per link and per data-rate embeds the varying FDR behaviors.
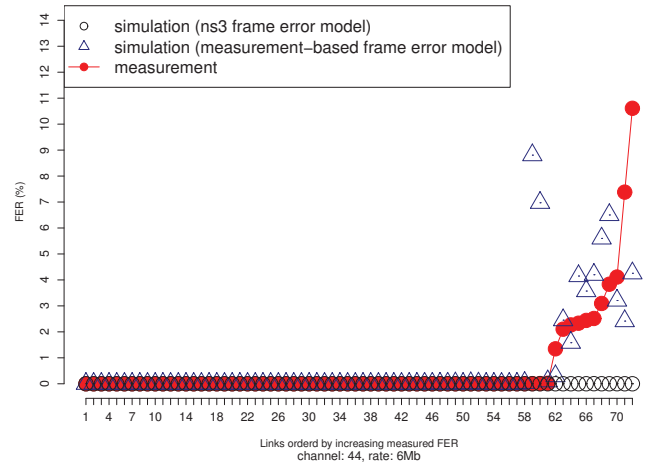

Figure 14: Comparison of the measurement and simulation-based FER at 6 Mb/s for all links. The simulation results show reasonable agreement. The root mean squared error (RMSE) between the measurements and the BIM-based results is equal to 1.73. The RMSE for the ns-3 model is equal to 1.78.

of the 72 links, the RSS distribution obtained from the simulation with the one obtained from the measurements. We use the Kullback-Leibler (KL) divergence to assess the distance between distributions. For 6 Mbit/s, the average of the KL divergence is 2.73×10^{-6} with confidence interval $[2.35 \times 10^{-6}, 3.10 \times 10^{-6}]$. Similar numerical values were obtained for 54 Mbit/s.

For the validation of the FDR model, Figure 13 compares measurement and simulation-based FDR at 6 and 54 Mb/s for all links. Simulations are performed with the BIM model only as there is no model for frame detection in ns-3. The FDR model shows excellent agreement with the measurement results even if we use a single FDR threshold per data rate for all links. Note that the RSS distributions are per link, and hence, this way, we are able to expose varying FDR behavior for each link.

For the validation of the FER, Figures 14 and 15 compare measurement results, simulation results with the BIM model and simulation results with the ns-3 default FER model. We consider a data-rate of 6 and 54 Mbit/s. For 6 Mbit/s, the maximum variation between the measurements and the BIM-based simulation results is around 10% and is observed only for a few links. The root mean squared error (RMSE) between the measurements and the BIM-based results is equal to 1.73. The RMSE for the ns-3 model is equal to 1.78. We observe higher variance for 54 Mbit/s (see Figure 15). Nevertheless, the BIM model exhibits a better agreement with RMSE equal to 6.60, whereas the RMSE for the ns-3 model is equal to 12.73. These higher variations are a result of using a network-wide FER model. Better accuracy can be achieved by switching to a node-based FER model. However, it remains to be seen whether higher accuracy in terms of FER will be necessary to simulate the higher-layer behavior accurately.

8. CONCLUSION

In this work, we have used an extensive measurement study to develop a novel site-specific link model of a wireless indoor network for the ns-3 simulator. In particular, our model comprises a link specific propagation model along with a network-wide frame detection ratio model and a frame error ratio model. To the best of our knowledge, the effect of signal detection has not been considered so far in wireless network simulation.

Our measurement study also allows us to draw a few guidelines for the simulation of wireless networks. First, per channel modeling is necessary. Then hardware specific effects have to be taken into account. For instance, the effect of transmission power calibration for different data rates is significant. Furthermore, weak links need special attention as more measurements are necessary to get a concrete picture of their performance. Finally, the time of the day can have an important effect on the overall network performance.

For future work, we plan to run more controlled simulations, especially with controlled channels by using attenuators or making use of environments such as the CMU wireless emulator[7]. We will also start investigating the effect of our modeling on the upper layers. Finally, we will extend our models to include time and space correlations.

9. REFERENCES

[1] Cost action 231, final report. `http://www.lx.it.pt/cost231/final_report.htm`.
[2] ns-2, `http://www.isi.edu/nsnam/ns/ns-documentation.html`.
[3] ns3, `http://www.nsnam.org/`.
[4] Opnet, `http://www.opnet.com/`.

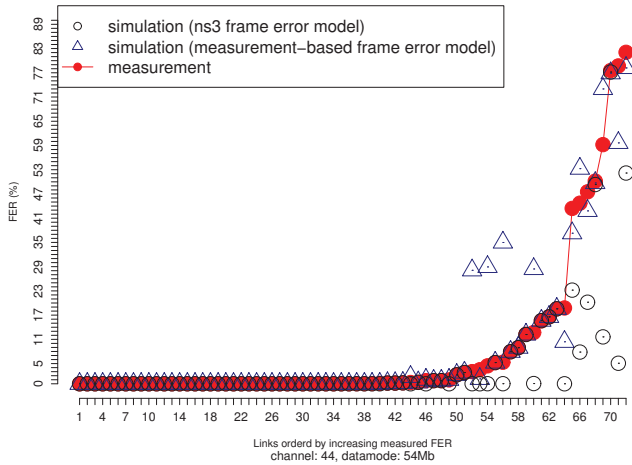[7] `http://www.cs.cmu.edu/~emulator/`

Figure 15: Comparison of the measurement and simulation-based FER at 54 **Mb/s for all links. There exists higher variance for** 54 **Mbit/s. The BIM model exhibits a better agreement with RMSE equal to** 6.60**, whereas the RMSE for the ns-3 model is equal to** 12.73**.**

[5] Qualnet. `http://www.scalable-networks.com/products/qualnet/`.

[6] IEEE standard for information technology-telecommunications and information exchange between systems-local and metropolitan area networks-specific requirements - Part 11: wireless lan medium access control (MAC) and physical layer (PHY) specifications, June 2007.

[7] IEEE 802.16m evaluation methodology document (EMD), IEEE 802.16m-08/004r5. `http://www.ieee802.org/16/tgm/docs/80216m-08_004r5.zip`, June 2009.

[8] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. *SIGCOMM Comput. Commun. Rev.*, 34(4):121–132, 2004.

[9] N. Baldo and M. Miozzo. Spectrum-aware channel and PHY layer modeling for ns3. In *Intl. Workshop on Network Simulation Tools (NSTOOLS)*, Oct. 2009.

[10] N. Baldo, M. Requena-Esteso, J. Nú nez Martínez, M. Portolès-Comeras, J. Nin-Guerrero, P. Dini, and J. Mangues-Bafalluy. Validation of the IEEE 802.11 MAC model in the ns3 simulator using the EXTREME testbed. In *Proc. of SIMUTools*, 2010.

[11] P. Barsocchi, G. Oli, and F. Potorti. Measurement-based frame error model for simulating outdoor Wi-Fi networks. *IEEE Trans. Wireless. Comm.*, 8(3):1154–1158, 2009.

[12] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom*, pages 31–42, 2005.

[13] I. Gruber, O. Knauf, and H. Li. Performance of ad hoc routing protocols in urban environments. In *European Wireless*, 2004.

[14] D. Halperin, W. Hu, A. Sheth, and D. Wetherall.

Predictable 802.11 packet delivery from wireless channel measurements. In *ACM SIGCOMM*, 2010.

[15] P. J. Husted, Y. Huanchun, and S. Aman. Adaptive interference immunity control. US Patent 7349503, March 2008.

[16] S. Ivanov, A. Herms, and G. Lukas. Experimental validation of the ns-2 wireless model using simulation, emulation, and real network. In *WMAN*, pages 433–444, 2007.

[17] William C. Jakes and Donald C. Cox, editors. *Microwave Mobile Communications*. Wiley-IEEE Press, 1994.

[18] A. Kashyap, S. Ganguly, and S. R. Das. Measurement-based approaches for accurate simulation of 802.11-based wireless networks. In *Proc. of MSWIM*, pages 54–59, 2008.

[19] A. O. Kaya, L. J. Greenstein, and W. Trappe. Characterizing indoor wireless channels via ray tracing combined with stochastic modeling. *Trans. Wireless. Comm.*, 8(8):4165–4175, 2009.

[20] S. Kurkowski, T. Camp, and M. Colagrosso. Manet simulation studies: The incredibles. *SIGMOBILE Mob. Comput. Commun. Rev.*, 9:50—61, 2005.

[21] M. Lacage and T. R. Henderson. Yet another network simulator. WNS2 '06. ACM, 2006.

[22] J. Liu, Y. Yuan, D. M. Nicol, R. S. Gray, C. C. Newport, D. Kotz, and L. F. Perrone. Empirical validation of wireless models in simulations of ad hoc routing protocols. *Simulation*, 81:307–323, April 2005.

[23] R. Merz, H. Schiöberg, and C. Sengul. Design of a configurable wireless network testbed with live traffic. In *Proc. of TridentCom*, May 2010.

[24] C. Newport, D. Kotz, Y. Yuan, R. S. Gray, J. Liu, and C. Elliott. Experimental evaluation of wireless simulation assumptions. *SIMULATION*, 83(9):643–661, September 2007.

[25] S. Papanastasiou, J. Mittag, E. G. Strom, and H. Hartenstein. Bridging the gap between physical layer emulation and network simulation. In *IEEE WCNC*, pages 1–6, Apr. 2010.

[26] L. Qiu, P. Bahl, A. Rao, and L. Zhou. Troubleshooting wireless mesh networks. *SIGCOMM Comput. Commun. Rev.*, 36(5):17–28, 2006.

[27] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2001.

[28] D. Reddy and G. Riley. Measurement based physical layer modeling for wireless network simulations. In *MASCOTS*, pages 46–53, 2007.

[29] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *ACM SIGCOMM*, pages 51–62, 2006.

[30] J. Ryu, J. Lee, S.-J. Lee, and T. Kwon. Revamping the IEEE 802.11a PHY simulation models. In *Proc. of MSWIM*, pages 28–36, 2008.

[31] I. Tinnirello, D. Giustiniano, L. Scalia, and G. Bianchi. On the side-effects of proprietary solutions for fading and interference mitigation IEEE 802.11b/g outdoor links. *Computer Networks*, Feb. 2009.