

Simplified Simulation Models for Indoor MANET Evaluation are not Robust

Andrés Lagar Cavilla, Gerard Baron[†], Thomas E. Hart, Lionel Litty and Eyal de Lara

Department of Computer Science

[†]Department of Electrical and Computer Engineering

University of Toronto

{andreslc, gsmbaron, tomhart, llitty, delara}@cs.toronto.edu

Abstract— We evaluate the robustness of simplified mobility and radio propagation models for indoor MANET simulations. A robust simplification allows researchers to extrapolate simulation results and reach reliable conclusions about the expected performance of protocols in real life. We show that common simplified mobility and radio propagation models are not robust. Experiments with DSR and DSDV, two representative MANET routing protocols, show that the simplifications affect the two protocols in very different manners. Even for a single protocol, the effects on perceived performance can vary erratically as parameters change. These results cast doubt on the soundness of evaluations of MANET routing protocols based on simplified mobility and radio propagation models, and expose the urgent need for more research on realistic MANET simulation.

I. INTRODUCTION

A multi-hop mobile ad hoc network (MANET) consists of a group of mobile wireless nodes that self-configure to operate without infrastructure support. Network peers communicate beyond their individual transmission ranges by routing packets through intermediate nodes [1]–[3].

Computer simulation is the most popular way to evaluate MANET routing protocols [4]–[6]. Simulation offers four important advantages: First, it enables experimentation with large networks. Second, it enables experimentation with configurations that may not be possible with existing technology. Third, it allows for rapid prototyping: by significantly abstracting the complexity of the real system, simulators enable the development and debugging of new protocols with reduced effort. Finally, it makes reproducible experiments in a controlled environment possible.

MANET protocol simulation presents challenging research problems. Besides having to simulate the networking stack and data traffic, MANET simulators also

need to incorporate models of node mobility and radio propagation. The mobility model determines how nodes choose destinations for their movement, the speed at which they move, and the physical paths they take. The radio propagation model determines whether communication between two given nodes is possible.

Random Waypoint (RWP) [4] and Free Space (FS) are the preeminent mobility and radio propagation models for MANET simulation. In RWP, a node picks a random destination inside a flat rectangular area, proceeds to it following a straight-line trajectory at a random speed, and pauses for a fixed time on arrival. The process then repeats itself. FS models propagation in an obstacle-free vacuum; signal strength degrades with the square of the distance between the transmitter and receiver. Several groups have extended these simple models with increasing levels of detail [7]–[12].

MANETs have been proposed for scenarios such as disaster relief, police, and military applications, which take place in complex obstacle-rich indoor environments. Therefore, it is important that MANET protocols be carefully evaluated in indoor conditions. Unfortunately, it is not clear that existing simulation models, which have been developed for outdoor environments, are appropriate for evaluating MANET protocols in indoor conditions.

This paper evaluates the *robustness* of simplified mobility and radio propagation simulation models for MANET simulations in indoor environments. A simplified simulation model is *robust* if the results obtained with the model for different protocols and simulation conditions are consistent (within a predictable error) with the results for the unsimplified model. A robust simplification allows researchers to extrapolate simulation results over different scenarios, and reach reliable conclusions about the expected performance of protocols in real life.

To determine the robustness of simplified models for indoor MANET simulation, we first introduce two detailed mobility and radio propagation models that take into account fine-grain obstacles and building materials. We then describe several simplifications to these detailed models that gradually decrease in sophistication. The least detailed models we consider correspond to the obstacle-free approaches provided by main-stream simulators (i.e., RWP and FS).

Experiments with DSDV [2] and DSR [1], two representative MANET routing protocols, show that simplifications to the mobility and radio propagation models are not robust, and have instead drastically different effects on the perceived performance of the two routing protocols. Whereas the performance of DSDV is virtually identical for all models, the performance of DSR varies widely between models. Moreover, even within DSR itself, the relative performance under the different models changes erratically as we vary experimental parameters. These findings raise troubling doubts over the soundness of MANET protocol evaluations based on simplified models, and expose the urgent need for more research on realistic MANET simulation models for indoor environments.

This paper makes two contributions: First, it shows that widely used simplified mobility and radio propagation models are not robust. We provide experimental evidence showing that the effects of simplifications of the simulation model are not uniform across protocols and evaluation conditions, leading to wrong conclusions about the performance of MANET protocols. Second, it provides the first evaluation of MANET routing protocols in indoor environments using detailed mobility and radio propagation models that account for fine-grain obstacles and building materials.

The rest of this paper is organized as follows. Section II describes the main characteristics of indoor environments, and presents two detailed mobility and radio propagation models. Section III describes simplifications to the detailed mobility and radio propagation models. Section IV briefly describes DSR and DSDV, two representative MANET protocols that we use in our evaluation. Section V presents our experimental results. Finally, Section VI compares the paper to previous work on mobility and radio propagation models, and Section VII presents our conclusions and discusses avenues for future research.




Fig. 1. Mobility graph superimposed on floor plan.

II. INDOOR MANET SIMULATION

MANET simulation in indoor environments presents interesting challenges. Indoor environments tend to be much smaller than the outdoor scenarios traditionally considered in MANET research. Moreover, modern buildings can have irregular shapes and large numbers of obstacles, which affect both node mobility and radio propagation. Finally, buildings typically have multiple floors, which adds a three-dimensional aspect to the simulation.

For example, Figure 1 shows the blueprint of the 3rd floor of the Bahen Centre for Information Technology, an academic research building at the University of Toronto. The building stands on a 113 by 88 meter lot, and provides 5,400 square meters of living space. This area is two orders of magnitude smaller than what is usually considered in MANET simulations [4], [6]. The figure also portrays the irregular layout of this building, but fails to convey a sense of its architectural complexity: cement pillars, steel shafts, brick walls, and the pervasive presence of glass are just some of its relevant characteristics. Finally, movement between floors is possible using elevators and stairs, and given that the environment under consideration is not a ground floor, movement outside the floor plan is – for all practical purposes – impossible.

In the rest of this section, we describe our detailed node mobility and radio propagation models for MANET simulation in obstacle-rich indoor environments. In their current form, these models are limited to single-floor simulations. Extensions to multi-floor simulations are the subject of future work.

A. Constrained Mobility

Constrained Mobility (CM) is a mobility model for obstacle-rich indoor environments. CM uses a *mobility graph* to constrain node mobility according to the obstacles present in the environment. For example, Figure 1 shows a mobility graph superimposed over the floor plan of our building. Vertices represent possible destinations that nodes can visit, and edges correspond to physically-valid paths over which nodes can move toward their intended destinations. Movement from one destination to another is accomplished by traversing the edges that constitute the shortest path between the two corresponding vertices. Therefore, nodes move through doors and hallways to reach their destinations, instead of resorting to straight-line trajectories.

At present, we draw the mobility graph on top of the floor plan using a simple graphical editor we developed. We use existing AutoCAD drawings, so this is not a laborious task. Moreover, the mobility graph needs to be built only once for a given floor plan, and is then reused in a large number of simulations. Nevertheless, we plan to explore techniques to automate the generation of mobility graphs.

We follow a simple approach in choosing destinations: we limit the choice of destinations to the set of solid-colored vertices in the graph, situated in “interesting” locations such as offices, classrooms and conference rooms. Each node randomly chooses a vertex in this set, and moves toward it at a randomly selected speed. After reaching its destination, the node pauses for a fixed time period before resuming movement.

B. Attenuation Factor

Attenuation Factor (AF) [13]–[15] is an empirical radio propagation model for indoor environments that deterministically accounts for multiple obstacles. AF assumes a time-invariant (i.e. no fast fading) channel in which a primary ray follows a straight-line trajectory between the transmitter and receiver, and accounts for the majority of the signal strength at the receiver. AF models the attenuation of the transmitter signal strength as a function of the distance that separates the transmitter and receiver, and the effects of walls of different materials along the primary ray path. While AF does not explicitly account for propagation effects like reflection, diffraction and scattering, and only models obstacles after their material types but not their thickness, it has been shown to yield good accuracy and high computational efficiency [13]. To the best of our knowledge, this is the first application of AF to MANET simulation.

The AF model is given by eq. 1, where P_o is the power at some nearby reference distance r_o , n is the path loss exponent that determines the rate at which power decreases with distance r , m_i is the number of obstacles of material type i along the primary ray path, PF_i is the partition factor loss due to material type i , and σ is the number of distinguishable material types.

$$\bar{P}_{AF}(r) = P_o(r_o) - 10n \log \frac{r}{r_o} - \sum_{i=1}^{\sigma} m_i \cdot PF_i \quad (1)$$

We obtain the values for m_i from the floor plan, by counting the number of walls of various material types that intersect the path of the primary ray. The values for P_o , n and the PF_i s are site-specific empirical approximations derived from experimental measurements. We next describe the equipment used and methodology followed to derive these quantities.

1) *Measurement Equipment*: Our experimental equipment consisted of two laptops running Linux with Wireless Extensions [16] enabled. Each had a Lucent Orinoco 802.11b based PCMCIA network interface card configured in ad hoc mode, and attached to a special external omni-directional antenna [17] that provided a gain of 9 dBi. The antennas also provided a horizontally-shallow radiation pattern that minimized the effects of reflection on the floor and ceiling. At 2 Mbps, the Orinoco network interface has a nominal transmit power of 15 dBm, and a receive sensitivity (Bit Error Rate $< 10^{-5}$) of -91 dBm [18]. With a cumulative gain of approximately 17 dB (two 9 dBi antennas minus pigtail losses), the setup was capable of recording receive signal strength values of -108 dBm for equivalent unity gain (0 dBi) antennas.

2) *Site-specific Parameterization*: We recorded over 250 measurements of signal strength over the floor plan illustrated in Figure 1. Each trial involved three steps. First, the two laptops were randomly positioned on different locations corresponding to vertices of the mobility graph. Second, an attempt was made to establish communication between the two laptops. If successful, both laptops were configured to ping each other; otherwise, a new pair of vertices was chosen. Finally, both laptops simultaneously recorded signal strength values over a period of one minute; on average, each made 30 measurements. We set the signal strength to the average of the measurements from both laptops.

Given stationary measurements and the symmetry of our experimental setup, we expected both laptops to record roughly the same signal strength values per




Fig. 2. Signal strength measurements, and AF and PL fits.




Fig. 3. AF generated visualization of signal strength.

trial because of the electromagnetic principle of *reciprocity* [19]. We did not anticipate the large effect the movement of other people would have on the assumption of a time-invariant channel. To achieve reciprocity, measurements had to be taken late at night.

Figure 2 plots the signal strength measurements in decibels (dBm) as a function of distance from the transmitter. To obtain the site-specific values for P_o , n and the PF_i s we ran a regression test in MATLAB. For each measurement point k we provided MATLAB with the signal strength \bar{P}_k , the distance r_k from the transmitter, and the number of walls of each type m_{ik} between the transmitter and the receiver.

We could distinguish seven material types in our AutoCAD floor plan: exterior walls, interior walls, exterior glass, interior glass, steel, concrete, and wood. However,

the best fit to the empirical measurements involved only four materials ($\sigma=4$). The interior walls and wood were combined into one material ($PF_1=2.479$ dB), metal and steel into another ($PF_2=4.7727$ dB), interior and exterior glass into a third ($PF_3=3.11104$ dB); exterior walls were our fourth material ($PF_4=6.50076$ dB). The effect of furniture and smaller obstacles was accounted for by n and P_o , which were fit to 1.9665 and -31.4627 dBm, respectively. r_o was nominally set to one meter. The resulting AF parameterization (the circles in Figure 2) comes within 8.9% of the experimental data. Figure 3 shows an AF-generated visualization of the signal strength of a transmitter placed in the center of the floor plan, and demonstrates the dramatic effect of wall attenuations on signal strength.

In its present state, AF shares many of the simplifying assumptions of other propagation models, such as a two-dimensional topology, omni-directional antennas and a time-invariant channel. Overcoming these limitations is a subject for future work.

III. SIMPLIFICATIONS

We next consider several simplifications of the node mobility and radio propagation models presented in the previous section. In Section V, we determine the extent to which these simplifications affect the evaluation of MANET protocols in indoor environments.

A. Mobility

The CM model described in Section II-A takes both internal and external walls into account. An initial simplification, *Shell*, discards the internal walls of the building and the mobility graph; instead, nodes select destinations randomly within the area outlined by the external walls of the building, and follow straight-line trajectories to their destinations. Shell thus increases the number of possible destinations, and distributes them uniformly. However, choice of destinations is constrained to locations that will not force nodes to step outside the floor plan perimeter.

Discarding the external walls from the Shell model yields the Random Waypoint (RWP) [4] model. We consider two variants of RWP. In the RWP_S (*small*), nodes move within a square with 73.5 meter sides; this area is equivalent to that of the Shell and CM models. In RWP_L (*large*), nodes move in a rectangle of 113 by 88 meters, the area of the lot over which the building stands.

B. Radio Propagation

A natural simplification to the AF model is to remove the explicit consideration of obstacles. The Free Space (FS) model usually employed in MANET simulations is an extreme example of this approach which assumes that signals propagate through a vacuum. This is an inappropriate assumption for our indoor environment, as any single node will obtain full radio coverage of the network (default FS radio range is 250 meter). Therefore, to obtain a realistic basis for comparison with AF, we need to scale down the radio range of FS. To do so, we employ the Log-Distance Path Loss (PL) function given by eq. 2. Here we assume an arbitrary homogeneous medium characterized by a path loss exponent n . The P_o and r_o components have the same meaning as in eq. 1.

$$\bar{P}_{PL}(r) = P_o(r_o) - 10n \log \frac{r}{r_o} \quad (2)$$

We used MATLAB to fit the PL equation to the empirical measurements taken in Section II-B.2. The obtained fit ($n = 4.0602$ and $P_o = -19.2464$ dBm, with $r_o = 1$ m) is plotted in Figure 2 as a dashed line; it yields a 14.85% relative error and offers a reasonable set of radio ranges for comparison.

Table I shows a set of AF thresholds and the corresponding down-scaled FS radio ranges. We will refer to this site-specific down-scaled FS model as FS'.

AF Threshold	-51	-61	-71	-81	-91
FS' Range	6.05	10.67	18.82	33.19	58.51

TABLE I

FS' EFFECTIVE RANGES (M) FOR DIFFERENT AF SENSITIVITY THRESHOLDS (DBM).

IV. MANET ROUTING PROTOCOLS

MANET routing protocols fall into two broad categories: reactive and proactive. Reactive, or on-demand routing protocols only update routes when packets need to be transmitted along them, while proactive protocols try to keep up-to-date routing tables at all times. We next describe two representative MANET routing protocols: DSR [1] (on-demand) and DSDV [2] (proactive).

A. DSR

The key feature of DSR is the use of *source routing*. Each packet carries in its header the full route to its destination. Intermediate nodes just forward the packet to the next hop in the source route. Routes are kept in a *route cache*. The cache is filled with routes the node discovers on demand, or that it overhears from

packets placed in the channel. When a packet needs a route and the cache offers no alternatives, DSR sends a *route request* broadcast message with an empty source route. Upon receipt of a route request, a node attempts to answer it with a cached route, or appends itself to the source route of the request and rebroadcasts it. Eventually the request will reach the destination, and a unicast *route reply* message will be sent using the route constructed by this process – note that many replies with different routes may be generated. Whenever a packet fails to be sent to its next hop, DSR assumes the link is broken, cleanses its cache of routes using the link, and sends a unicast *route error* message to the originator of the packet, who will try to retransmit the packet using another route.

B. DSDV

DSDV is a table-driven proactive routing protocol, that builds on the Bellman-Ford distance-vector routing algorithm. Every node keeps a routing table with entries for all other nodes in the network. A routing entry includes the destination's address, the next hop to the destination, a metric (usually the path length), and a sequence number to indicate the freshness of the information, and to ensure loop-freedom in the formation of routes. Nodes exchange route entries periodically (typically every 15 seconds). A node modifies its routing table when it learns of a fresher route or a current route that is shorter.

V. EXPERIMENTAL EVALUATION

In this section, we evaluate the robustness of simplifications to the mobility and radio propagation models for indoor MANET simulation. We consider five combinations of the mobility and radio propagation models described in Sections II and III. We evaluate the robustness of Constrained Mobility under site-specific Free Space (CM-FS') as a simplification of Constrained Mobility under Attenuation Factor (CM-AF), and of Shell and Random Waypoint under site-specific Free-Space (Shell-FS' and RWP-FS') as simplifications of CM-FS'. For RWP, we consider both RWP_S-FS' and RWP_L-FS'.

A. Simulation Environment

We ran our experiments using the *ns2* [20] network simulator version 2.26, augmented with the CMU wireless extensions [21]. We extended *ns2* with the AF propagation model. Using eq. 1 and the empirical parameters derived in Section II-B, our AF implementation can compute signal strength at the receiver for any pair of nodes

arbitrarily positioned inside the modeled floor plan. We determine the number of walls in the primary ray path, and their material types, by computing the intersections between the primary ray and the building’s geometry captured in the AutoCAD floor plan. To generate CM and Shell mobility patterns, we developed extensions to the CMU *setdest* program that incorporate the mobility graph and floor plan shape specifications as additional inputs.

We report results for networks of 20, 30, 40, 50 and 60 nodes. All experiments ran for 1200 seconds of simulated time. The channel’s capacity and frequency were 2 Mbps and 2.4 GHz, respectively. We used the ns2 implementations of the 802.11 DCF MAC protocol, and DSR and DSDV routing protocols. In all simulations, nodes choose a speed uniformly distributed between 0.5 and 3 m/s, which we regard as the range of human walking speeds in an indoor environment; in particular, we chose a non-zero minimum speed to avoid the average speed decay phenomenon analyzed in [22]. To stress the network, we set the pause time to 0.

We modeled network traffic using Constant Bit Rate (CBR) sources. In each experiment, half the nodes in the network are CBR sources, and each source transmits 64-byte packets at a rate of 4 per second. We experimented with other sending rates, packet sizes and number of sources. We omit those results, as they show similar trends.

We experimented with a variety of sensitivity thresholds for the AF model, ranging from the default -91 dBm up to -51 dBm, with a step of 10 dBm. Based on the mappings from Table I, we experimented with FS’ effective ranges between 5 and 60 meters. All the results presented are averages of five runs over different randomly generated mobility patterns. No significant variance was observed among different runs for the same scenario; standard deviation values were consistently smaller than 10% of the corresponding average.

B. Experimental Results

We evaluate the robustness of simplifications of the mobility and radio propagation models by comparing the Packet Delivery Rate (PDR) of the protocols under consideration. Figures 4 and 5 show the PDR for networks of 30, 40 and 50 nodes for DSDV and DSR, respectively. The results for networks of 20 and 60 nodes are not shown for space considerations. The results for 20 nodes are virtually identical to the 30 node graph for both protocols. In the 60 nodes network, congestion effects dominate the DSR simulation and packet delivery

rate for all models drops below 30%. The DSDV graph for 60 nodes is similar to the 50 nodes graph.

A comparison of Figures 4 and 5 provides a first indication that simplifications of the mobility and radio propagation models may not be robust; the effects of the simplifications on performance are not uniform across the two protocols. Successive simplifications do not alter the perceived performance of DSDV, but the performance of DSR changes dramatically across models. This indicates that conclusions reached about the relevance of detail in the evaluation of one MANET protocol may not carry over between protocols. Assuming otherwise will likely produce misleading results.

For example, assume that based on the DSDV results we were to (wrongly) conclude that RWP_L-FS' , a simple model which assumes no obstacles for mobility or radio propagation, is a good approximation for the more complex CM-AF. The similar performance curves of both models, and the fact that the results obtained with RWP_L-FS' are within a bounded and consistent error from the results of the sophisticated model, would seem to bolster this assertion. Unfortunately, if we evaluated DSR using RWP_L-FS' , we would reach the erroneous conclusion that DSR outperforms DSDV in this environment. Note that the exact opposite occurs with the more detailed model.

Figure 5 provides further evidence that simplified models are not robust; the relationship between the observed performance of DSR under different models changes dramatically as we increase the number of nodes. For example, CM-FS’ outperforms RWP_L-FS' for 30 nodes, but the opposite is true for 50 nodes. The implication is that observations about the relevance of detail do not necessarily carry over even within the evaluation of the same protocol as we change the experimental parameters.

For example, assume that based on the DSR results for the 30 nodes network (Figure 5(a)), we were to (wrongly) conclude that RWP_L-FS' is a good approximation of CM-FS’. Further assume that we are evaluating an energy-aware routing protocol that increases battery life by reducing transmitter power at all nodes to the same lower level. Evaluation of this optimization under RWP_L-FS' would lead us to the (false) conclusion that for the 50 nodes network an effective transmission range of 45 meters (roughly -85dBm in our environment) achieves a delivery rate close to 100%. In contrast, experiments with CM-FS’ show that at this transmission range, delivery rate is closer to 50% (Figure 5(c)); therefore, the power adaptation policy would not be effective for this network.

Based on the DSR results we make the following




Fig. 4. DSDV packet delivery rate.




Fig. 5. DSR packet delivery rate.

three observations: (i) there is a drastic difference in the performance of DSR between CM-AF and CM-FS', a strong indication that the sophistication of the radio propagation model can affect the results of the evaluation; (ii) there are important differences in the delivery rate of DSR under CM-FS' and Shell-FS', which indicates that internal walls even when they only limit mobility (and do not affect radio propagation) can affect significantly the results of the evaluation; and (iii) it appears that for DSR, at least under our simulation conditions, Shell-FS' and RWP_S-FS' are equivalent, and that therefore external walls have much less of an effect on the simulation results than the internal walls. The fact that interior walls turned out to be relevant, whereas exterior walls did not, proved contrary to our initial expectations and is further evidence of the difficulty of

foreseeing the effects of simplifications on the evaluation of MANET protocols.

C. Performance Breakdown

The difference in the performance of DSDV and DSR can be traced to the protocols routing overhead under the various mobility and radio propagation models. Figures 6(a) and 6(b) show the normalized routing load – the fraction of routing packets transmitted per application-layer delivered packet – of DSDV and DSR for networks of different numbers of nodes, with a FS' transmission range of 35 meters and an AF sensitivity threshold of -81 dBm. For DSR there is a large difference in routing overhead between the models, with a significant increase in overhead as the network grows in size (notice the logarithmic scale used in the graph). In contrast, for DSDV there is little variation between the models, and

the overhead increases modestly with the number of nodes.

To understand the differences in routing overhead, we look at the effect the models have on the network topology. Figures 7(a) and 7(b) show the normalized neighbor density and the normalized link changes count for a network of 40 nodes, respectively. Neighbor density measures the average number of nodes within transmission range from each other. We normalize this quantity by $n - 1$, where n is the number of nodes in the network. The link changes count is a measure of the average number of connectivity changes between each node pair over the length of the simulation. We normalize this quantity by the total number of links ($n * (n - 1) / 2$). Because of normalization, the figures for networks of 20, 30, 50, and 60 nodes are very similar and are therefore not shown. Note that both of these topology metrics are independent of the routing protocol.

The neighbor densities of CM-AF and CM-FS' differ significantly. The neighbor density values of CM-FS' are initially lower than those reported for CM-AF. When no obstacles block the primary ray between transmitter and receiver, AF propagation can actually reach farther in obstacle-free areas such as hallways and large conference rooms. For ranges greater than 25 meters, the situation reverses as the effects of transmission through multiple walls become the limiting factor in AF propagation. While FS' coverage is effectively a disk, the attenuations induced by multiple obstacles in AF propagation render a non-circular coverage zone, as Figure 3 shows. Thus, for the upper band of sensitivity thresholds, individual nodes in FS' reach an almost complete network coverage (90%), but coverage of the network with AF propagation is restricted to 62%. The lower neighbor density of CM-AF results in fewer single-hop and longer multi-hop routes between communicating nodes, potentially degrading the routing performance.

Perhaps not surprising, CM-FS', Shell-FS' and RWP_S-FS' have similar neighbor densities. This result seems to suggest that under FS', at least in our environment, neighbor density is mostly dependent on the effective movement area, as opposed to the specific paths taken by nodes. RWP_L-FS' models a larger space and consequently has lower neighbor density.

Figure 7(b) is much more informative. It shows that there is a very large difference in the number of link changes between CM-AF and CM-FS' (note the logarithmic scale). In CM-AF, the radio connectivity between two nodes suffers abrupt changes as nodes move behind obstacles, resulting in numerous short disconnections. In

contrast, with CM-FS' connectivity degrades slowly and smoothly as nodes move away. Albeit smaller, there is a significant difference in link changes between CM-FS' and the other FS' models. In CM-FS' nodes move away from one another much more quickly by traversing graph edges – walking through hallways – in opposite directions; this in turn causes a higher rate of link changes.


DSR is seriously affected by the much higher rate of link changes in CM-AF and CM-FS'. The protocol reacts by purging from the cache all routes involving the broken link, even if the link breakage is short-lived. Route error messages are generated, and packets are either salvaged with cached routes – probably stale due to the high number of link breakages –, or a new route discovery cycle is triggered. In either case, the network is clogged by the additional traffic, leading to a congestion breakdown. Beyond diminishing the availability of the shared channel, congestion misleads DSR into believing that links have been broken, when communication has actually been prevented by collisions and interference. As the number of nodes in the network increases, the higher number of link changes and neighboring nodes exacerbate the congestion effect. This situation is particularly acute in CM-AF, due to the frequent short disconnections previously discussed.

DSDV, on the other hand, reacts in a very controlled and regular manner to events such as link breakages or congestion. For instance, a link is diagnosed as broken only after three periodic updates from the corresponding node have not been received, and several timing constraints prevent unrestricted propagation of updates. Therefore the protocol's simulated performance is mostly immune to the problems that negatively affect DSR.


D. Discussion

While the CM-AF model is significantly more detailed than the simple models used in mainstream simulators (i.e., RWP_L-FS'), the model makes several simplifying assumptions, such as not modeling multiple floors, or assuming a time-invariant channel with no small scale fading. Therefore, the robustness of CM-AF as a simplification of real-life conditions is not guaranteed.

In this light, one should be careful not to view the results we present for DSDV and DSR as *realistic* predictors of the expected protocol performance in real-life indoor environments. Instead, in this paper we consider the CM-AF results only as a benchmark against which to compare the simplified models. Given that the simplified models are not robust simplifications of CM-




(a) DSDV




(b) DSR

Fig. 6. Routing overhead. FS' transmission range is 35m, AF sensitivity threshold is -91 dBm.



(a) Normalized neighbor density.



(b) Normalized link changes.

Fig. 7. Topology Metrics.

AF in our environment, we hypothesize that it is very unlikely that they are robust simplifications of real-world environments.

VI. RELATED WORK

Related work in improved mobility modeling falls under three areas: behavioral modeling, and outdoor and indoor simulation models. Research into more realistic radio propagation models is also very relevant to our work.

The RWP model and its properties have been subject to extensive research [22], [23]. Betstetter [23] and Camp et al. [8] review several variations of RWP that increase

realism by applying statistical distributions to the choice of node destinations and speed. Camp et al. [8] also reviews models of group behavior, where nodes are distributed around reference points. GEMM [9] provides a set of constructs for designing behavioral models. Research on the behavioral aspects of node mobility is complementary to the work we have presented in this paper.

Several models simulate outdoor movement along grids or graphs. The City Section Mobility Model [8] uses a grid to model vehicle movement on city streets in a coarse-grained manner. The Obstacle Mobility Model

(OM) [10] uses automatically-generated Voronoi graphs to model building-to-building movement in a campus. Tian et al. [11] restrict movement to an arbitrary user-defined graph over a large outdoor environment. Jetcheva et al. [24] use actual traces of city buses to model vehicular movement over a major city.

There is little previous work on indoor MANET simulation. Johansson et al. [5] consider conference, event coverage, and disaster area scenarios with a few simple obstacles. CAD-HOC [25] is a tool for designing indoor mobility patterns. CM improves over these efforts by modeling fine-grain obstacles and significantly more complex indoor environments.

Several groups have explored the use of various radio propagation models for MANET protocol simulation. *ns2* [20] comes bundled with the Two-Ray Ground and Shadowing models. The former is a variation of the FS propagation model that considers a second ray reflecting off the ground for long distances. The latter augments FS propagation with a probabilistic fluctuation. Jardosh et al. [10], experimented with the Line-of-Sight (LOS) model, where Two-Ray propagation is preempted if an obstacle blocks the primary ray between transmitter and receiver. We experimented with LOS in conjunction with CM, and found it is too coarse-grained to produce any useful results: node communication was mostly preempted, unless an intermediate node was conveniently placed at a hallway intersection. We omit these results for space considerations. Dricot et al. [12] have attempted Ray-Tracing, but have been limited to very small floor plans due to its high computational cost. Finally, a simpler variant of the AF model has been implemented in RADAR [15], but for the purposes of location tracking, has not been previously used in MANET simulations.

In contrast to these deterministic models, probabilistic propagation models use stochastic methods to reproduce the variations in signal strength induced by obstacles. SIRCIM [13] is a typical example, simulating both large scale and small scale fading at the MAC and physical layers. Takai et al. [26] also use probabilistic Ricean and Rayleigh distributions for modeling small scale fading.

Studies on simplified simulation models have focused on the limitations of either the mobility or propagation model individually [7], [26], [27]. To the best of our knowledge, we are the first group to consider detailed propagation and mobility models in conjunction, and to identify and evaluate the robustness of simplified models in complex indoor environments.

VII. CONCLUSIONS AND FUTURE WORK

We addressed the robustness of simplifications of the mobility and radio propagation models for indoor simulation of MANET routing protocols. A simplification of a mobility or radio propagation model is robust if the results obtained with the simplification for different protocols and simulation conditions are within a predictable error of the expected result for the unsimplified model. Robust simplifications allow researchers to extrapolate simulation results and reach reliable conclusions.

Experimental results showed that simplifications of the mobility or radio propagation model are not robust (at least) for indoor environments. The simplifications we considered had drastically different effects on the perceived performance of the two protocols we evaluated. Even for the same protocol, the effects on perceived performance varied erratically for different simulation configurations.

These results cast serious doubt on the validity of simulation-based MANET evaluations using simplified models. Even if a simplified model appears to be a good approximation for evaluating a specific MANET protocol, there are no assurances that the model will be valid for other routing protocols, or even the same protocol under different experimental conditions. This troubling conclusion is a compelling indication of the importance of further research on the development and validation of realistic models for indoor MANET simulation.

In the future, we will extend our evaluation to other buildings, and will experiment with mobility and radio propagation models that extend to multiple floors and take into consideration smaller obstacles such as furniture. We also intend to explore the effect of different movement patterns where destinations are not chosen randomly. In the long run, we want to relax the assumption of a time-invariant radio channel and model the effect of human activity, which our empirical measurements showed to be significant.

ACKNOWLEDGMENT

We would like to thank Angela Demke, Baochun Li, Ben Liang, Peter Marbach, the students of the systems group, and the anonymous reviewers for their helpful suggestions. We would also like to thank Parham Aarabi for providing access to the computing cluster where we ran our simulations. This research was partially funded by the NSERC Discovery Grant 261545-03. Andrés Lagar Cavilla is partially funded by a Wolfond Scholarship.

REFERENCES

- [1] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*, T. Imielinski and H. Korth, Eds. Kluwer Academic Publishers, 1996, ch. 5, pp. 153–181.
- [2] C. Perkins and P. Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers," in *ACM SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*, London, UK, Aug. 1994.
- [3] C. E. Perkins and E. M. Royer, "Ad hoc on demand distance vector routing," in *2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, New Orleans, LA, Feb. 1999.
- [4] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols," in *4th ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Dallas, TX, Oct. 1998.
- [5] P. Johansson, T. Larsson, N. Hedman, B. Mileczarek, and M. Degermark, "Scenario-based performance analysis of routing protocols for mobile ad-hoc networks," in *5th ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Seattle, WA, Aug. 1999.
- [6] S. R. Das, C. E. Perkins, and E. E. Royer, "Performance comparison of two on-demand routing protocols for ad hoc networks," in *19th Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Tel Aviv, Israel, Mar. 2000.
- [7] J. Liu, Y. Yuan, D. M. Nicol, R. S. Gray, C. C. Newport, D. Kotz, and L. F. Perrone, "Simulation validation using direct execution of wireless ad-hoc routing protocols," in *8th Workshop on Parallel and Distributed Simulation*, Kufstein, Austria, May 2004.
- [8] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Communications and Mobile Computing (WCMC): Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, vol. 2, no. 5, pp. 483–502, 2002.
- [9] M. J. Feeley, N. Hutchinson, and S. Ray, "Realistic mobility for mobile ad hoc network simulation," in *3rd International Conference on AD-HOC Networks and Wireless (ADHOC-NOW)*, Vancouver, BC, July 2004.
- [10] A. Jardosh, E. M. Belding-Royer, K. C. Almeroth, and S. Suri, "Towards realistic mobility models for mobile ad hoc networks," in *9th International Conference on Mobile Computing and Networking (MobiCom)*, San Diego, CA, Sept. 2003.
- [11] J. Tian, J. Hahner, C. Becker, I. Stepanov, and K. Rothermel, "Graph-based mobility model for mobile ad hoc network simulation," in *35th Annual Simulation Symposium*, San Diego, CA, Apr. 2002.
- [12] J.-M. Dricot and P. D. Doncker, "High accuracy physical layer models for wireless network simulations in ns-2," in *International Workshop on Wireless Ad-hoc Networks (IWWAN)*, Oulu, Finland, May 2003.
- [13] T. S. Rappaport, *Wireless Communications: Principles & Practice*. Prentice Hall, 2002.
- [14] S. Y. Seidel and T. S. Rappaport, "Site-specific propagation prediction for wireless in-building personal communication system design," *IEEE Transactions on Vehicular Technology*, vol. 43, no. 4, pp. 879–891, Apr. 1994.
- [15] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *19th Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Tel Aviv, Israel, Mar. 2000.
- [16] Wireless Tools for Linux. [Online]. Available: http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/Tools.html
- [17] Superpass SPDG80-D3 antenna specifications sheet. [Online]. Available: <http://www.superpass.com/SPDG80-D3.html>
- [18] Enterasys RoamAbout wireless high-rate PC card. [Online]. Available: <http://www.enterasys.com/products/wireless/CSIxD-AA/>
- [19] C. S. Balanis, *Advanced Engineering Electromagnetics*. Wiley Text Books, 1989.
- [20] N. Drakos and R. Moore, *ns2 - The Manual (formerly Notes and Documentation)*, 1999. [Online]. Available: <http://www.isi.edu/nsnam/ns/doc/node215.html>
- [21] The Monarch Project. [Online]. Available: <http://www.monarch.cs.rice.edu>
- [22] J. Yoon, M. Liu, and B. Noble, "Random waypoint considered harmful," in *22nd Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, San Francisco, CA, Mar. 2003.
- [23] C. Bettstetter, "Mobility modeling in wireless networks: categorization, smooth movement, and border effects," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 3, pp. 55–66, July 2001.
- [24] J. G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A. K. Saha, and D. B. Johnson, "Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture," in *5th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, Monterey, CA, Oct. 2003.
- [25] S. Shah, E. Hernandez, and A. Helal, "CAD-HOC: A CAD like tool for generating mobility benchmarks in ad-hoc networks," in *2nd IEEE/IPSJ Symposium on Applications and the Internet (SAINT)*, Nara, Japan, Jan. 2002.
- [26] M. Takai, J. Martin, and R. Bagrodia, "Effects of wireless physical layer modeling in mobile ad hoc networks," in *International Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC)*, Long Beach, CA, Oct. 2001.
- [27] J. Heidemann, N. Bulusu, J. Elson, C. Intanagonwiwat, K. Lan, Y. Xu, W. Ye, D. Estrin, and R. Govindan, "Effects of detail in wireless network simulation," in *SCS Multiconference on Distributed Simulation*, Phoenix, AZ, Jan. 2001.