

Modeling VANET Deployment in Urban Settings

Atulya Mahajan^a, Niranjan Potnis^b, Kartik Gopalan^c, Andy Wang^b

^aLehman Brothers, New York, NY, USA

^bComputer Science, Florida State University, Tallahassee, FL, USA

^cComputer Science, State University of New York at Binghamton, Binghamton, NY, USA

Contact: kartik@cs.binghamton.edu, awang@cs.fsu.edu

ABSTRACT

The growing interest in wireless Vehicular Ad Hoc Networks (VANETs) has prompted greater research into simulation models that better reflect urban VANET deployments. Still, we lack a systematic understanding of the required level of simulation details in modeling various real-world urban constraints. In this work, we developed a series of simulation models that account for street layout, traffic rules, multi-lane roads, acceleration-deceleration, and RF attenuation due to obstacles. Using real and controlled synthetic maps, we evaluated the sensitivity of the simulation results toward these details. Our results indicate that the delivery ratio and packet delays in VANETs are more sensitive to the clustering effect of vehicles at intersections and their acceleration-deceleration. The VANET performance appears to be only marginally affected by the simulation of multiple lanes and careful synchronization at traffic signals. We also found that the performance in dense VANETs improves significantly when routing decisions are limited to a wireless backbone of mesh nodes, whereas in sparse VANETs, performance improves when vehicles also participate in ad hoc routing. Finally, through measurement and analysis of signal strengths around urban city blocks, we show that the effect of signal attenuation due to physical obstacles can potentially be parameterized in simulations. Our work provides a starting point for further understanding and development of more accurate VANET simulation models.

Categories and Subject Descriptors: C.2, I.6

General Terms: Measurement, Performance

1. INTRODUCTION

Wireless vehicular ad hoc networks (VANETs) consist of vehicles traveling on urban streets, capable of communicating with each other with/without the aid of a fixed infrastructure. VANETs have generated considerable research and commercial interest with promising applications, such as mobile communication, traffic monitoring, safety, and public utility management. The current simulation models used in popular wireless simulators such as NS-2 [15] tend

to ignore real-world constraints. For example, the widely used Random-Waypoint Model (RWM) [13] assumes that the nodes move in an open field without obstructions and ignores important factors such as street layouts, intersections with traffic signs, or inter-vehicle interactions. Similarly, the commonly used two-ray ground radio propagation model ignores effects such as RF attenuation due to buildings and other obstacles. Consequently, the simulation results are unlikely to reflect the protocol performance in the real world. In addition, a wide spectrum of VANET deployment scenarios are possible ranging from pure vehicle-to-vehicle (V2V) communication, routing through a fixed wireless mesh network (WMN) backbone, or even a hybrid deployment of V2V and mesh-based routing. These range of deployment possibilities are not adequately captured by current models.

To address these limitations, researchers have become interested in modeling 'realistic' VANET deployments [7, 19, 22, 4, 10, 8, 24, 3, 12, 9]. Although these studies capture different levels of simulation details and realism, existing research has shed little light on the actual level of details required and the sensitivity of simulation results to those details. Excessive details may prolong the running time of a simulation with little or no impact on the results, while too few details invariably lead to inaccurate results.

This paper addresses the following question: *what is the sensitivity of VANET simulation results toward individual modeling details?* We aim to identify factors that significantly affect the VANET performance, and also factors that marginally affect the performance and could potentially be ignored. We developed a series of simulation models that systematically capture various urban constraints at increasing levels of detail. Our purpose is not to advocate one model over another, but to use them as a means to gain better insights into the sensitivity of simulations results to different levels of realism.

- Our results indicate that the clustering effect of vehicles at intersections significantly affects the VANET performance. Increasing either the wait times at the intersections or the number of nodes lead to increased clustering. Consequently, increased clustering can produce different results depending upon whether the neighboring intersections are within or beyond each other's transmission range, the former leading to higher delivery ratios than the latter.
- We also find that VANET performance is sensitive to physical topology (block sizes and street layouts) and the acceleration and deceleration of vehicles.
- On the other hand, we find that adding additional

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MSWIM'07, October 22–26, 2007, Chania, Crete Island, Greece.

Copyright 2007 ACM 978-1-59593-851-0/07/0010 ...\$5.00.

complexity to the models, such as simulating multiple lanes and the synchronization of traffic lights, yields marginal impact on VANET performance.

- Comparing two mesh-enhanced VANET deployment strategies, we found that the delivery ratio in dense VANETs improves when routing decisions are limited to a wireless backbone of mesh nodes, whereas, in sparse VANETs, the performance improves when vehicles also participate in routing.
- We performed empirical signal strength measurements around urban city blocks and analyzed the data to show that the effect of signal attenuation due to obstacles can potentially be parameterized in simulations.

We now describe a series of enhancements to VANET simulation models, each successive model capturing increasing levels of detail.

2. STOP SIGN MODEL

In the Stop Sign Model, every street at an intersection has a stop sign. Any vehicle approaching the intersection must stop at the signal for a specified time, which is configurable. On the road, each vehicle’s motion is constrained by the vehicle in front of it. That is – a vehicle cannot move further than the vehicle in front of it, unless it is a multi-lane road and the vehicles are allowed to overtake each other. When vehicles follow each other to a stop sign, they form a per-street queue at the intersection. Each vehicle waits for at least the required wait time once it gets to the head of the intersection after other vehicles ahead in the queue clear up. Vehicle crossings at the intersection are not coordinated among different directions. Although in practice, an urban layout is unlikely to have stop signs at every intersection, this model does serve as a useful first step to understand the effect of mobility dynamics on VANET performance.

3. PROBABILISTIC TRAFFIC SIGN MODEL

Next, we refined SSM by replacing stop signs with traffic signals. Although it is possible to simulate the detailed coordination of traffic lights from various directions, we did not do so at this stage. The rationale was to understand whether such levels of detail would produce a significant impact on routing performance. As an intermediate step, we developed the Probabilistic Traffic Sign Model (PTSM), which approximates the operation of traffic signs by *not coordinating* among different directions. When a node reaches an intersection with an empty queue, it stops at the signal with a probability p and crosses the signal with a probability $(1 - p)$. If it decides to wait, the amount of wait time is randomly chosen between 0 and w seconds. Any node that arrives later at a non-empty queue will have to wait for the remaining wait time of the previous node plus the startup delay between queued cars. Whenever the signal turns green, the vehicles begin to cross the signal at the delay interval, until the queue becomes empty. The next vehicle to arrive at the head of an empty queue again makes a decision to stop with a probability p and so on. Like SSM, there is no coordination among vehicles crossing an intersection from different directions. This model reduces excessive stopping, while approximating the behavior of traffic lights.

4. TRAFFIC LIGHT MODEL

SSM and PTSM coarsely model vehicular mobility. In order to understand which other level of detail besides street

topology is absolutely essential, we further refined PTSM with incremental levels of mobility details. We call this model and its variants, the Traffic Light Model (TLM).

Coordinated traffic lights: Under TLM, traffic lights at each intersection are coordinated. First, consider an intersection with an even number of roads with single-lane opposing traffic. The lights turn green for one pair of opposing sides at a time to cross the intersection simultaneously, while the remaining lights stay red. Vehicles that need to turn follow the free turn rule once they reach the head of the queue. After a fixed period, green signals are rotated to another pair of roads with opposing traffic. The case of an odd number of roads meeting at an intersection (e.g., a T intersection) is modeled by permitting one of the roads to periodically have a green light by itself.

Acceleration and Deceleration: The next level of detail we added to TLM was the acceleration and deceleration of vehicles. With this feature, vehicles do not change their state between rest and the top speed abruptly. Instead, they accelerate gradually from rest up to the maximum speed. Similarly, when approaching a stop sign or red light, they decelerate gradually to a stop.

Multiple Lanes: Another feature of the TLM is the introduction of multiple lanes. For real maps, the number of lanes can be determined by the type of the road specified in the street database. When a vehicle enters a new road, e.g. when crossing or turning at an intersection, it selects the lane with the least number of vehicles.

Generating Variants of TLM: To study the sensitivity of VANET performance to mobility details, various features of TLM can be independently enabled or disabled to obtain different TLM variants. In particular, four variants of TLM can be obtained by enabling or disabling the acceleration-deceleration and multi-lane features. Hence, the basic TLM with neither features still has one enhancement over PTSM, namely the coordinated traffic lights.

5. MESH-ENHANCED DEPLOYMENT

We also implemented two alternative VANET deployment models in urban settings, where a wireless multi-hop network of stationary mesh nodes enables or supplements the network connectivity among mobile nodes. Mesh nodes can be strategically positioned at a subset of street intersections. Such a wireless backbone of mesh nodes can potentially reduce churn in network connectivity and increase route stability when the static mesh nodes participate in the ad hoc routing protocol. We implemented and evaluated two deployment scenarios: (1) **mesh-enhanced peer-to-peer routing (MEPPR)** where both the mobile nodes and static mesh nodes participate in routing, and (2) **mesh-enhanced infrastructural routing (MEIR)** where only the static mesh nodes route packets generated by the mobile nodes. For both MEPPR and MEIR, the TLM mobility pattern generator was altered to designate a subset of nodes as mesh nodes that are positioned at street intersections and remain stationary throughout the simulation.

6. MODELING OBSTACLES

Besides confining the vehicular movements to streets, physical obstacles also affect radio signal propagation through attenuation, reflection, diffraction, and refraction. This is in addition to the free-space attenuation of radio signals with distance. Since a receiver needs a minimum signal-to-noise




Figure 1: Signal strengths (dBm) around a downtown Tallahassee block. AP is the access point.

ratio to receive data, accounting for the obstacles is important when evaluating VANETs through simulations.

Traditional analytical models have limitations in capturing complex real-world factors that influence radio signal strengths. Our approach is to use empirically measured data from real urban settings to capture the impact of different factors on radio signals in a few simulation parameters.

We measured the signal strength variation from a commodity access point around two city blocks in downtown Tallahassee – including a 100m x 100m block with several three-story buildings and a 200m x 50m block with one-story buildings. We placed an 802.11b Linksys wireless access point at a corner of the block being measured. We then used the Wavemon [23] tool running on a Linux laptop equipped with a wireless PCI card to take signal strength measurements at various locations around the block. The empirical data were composed of the distances from the access point and the associated signal strength. A logarithmic transformation was performed on collected distances before a linear regression was applied on the signal strength S (in decibels/milliwatts or dBm), as a function of distance d (in meters) [11]. Logarithmic linear regressions yielded the following formulas, with R^2 (coefficient of determination) of 0.6836 and 0.9698, indicating that 68% and 97% of the variances in data are explained by these equations respectively.

$$\text{Block1} : S = -25.809 - 29.773 * \log(d) \quad (1)$$

$$\text{Block2} : S = -20.089 - 33.012 * \log(d) \quad (2)$$

From the structure of Equations 1 and 2, we can derive a simplified parameterization of the received signal strength.

$$P_r = P_t + A - B \log(d) \quad (3)$$

P_r and P_t are the signal strengths (in dBm) at the receiver and the sender respectively; d is the distance between the two in meters, and A and B are tunable parameters (whose significance we will explore below). Intriguingly, the propagation models used in NS2 can also be represented in the form of Equation 3. The NS2 propagation models are the Friis model for free space and the two-ray ground model that accounts for multipath reflection from the ground.

$$\text{Friis} : P_r(d) = \frac{P_t G_t G_r \lambda^2}{(4\pi)^2 d^2 L} \quad (4)$$

$$\text{Two ray ground} : P_r(d) = \frac{P_t G_t G_r h_t^2 h_r^2}{d^4 L} \quad (5)$$

$P_r(d)$ is the received power (watts) at distance d , P_t is the transmit power (watts), G_t and G_r are antenna gains for transmitter and receiver respectively, h_t and h_r are the antenna heights for the transmitter and the receiver respectively, L is the system loss, and λ is the wavelength. Both

Equations 4 and 5 could be represented in the form of Equation 3, after converting watts into dBm. The default values of A and B in NS2 for the Friis model are $A = -31$ and $B = 20$; for the two-ray ground model, $A = 7.5$ and $B = 40$.

Assuming that the received signal strength largely depends on the presence of obstacles and the distance from the sender, we can interpret A and B as follows. Parameter A captures the constant factor reduction in signal strength due to the presence of obstacles in a particular terrain. Parameter B captures the order of magnitude reduction in the signal strength with the distance from sender, the order of magnitude being determined by nature of the obstacles. We will refer to A as the *constant factor* and B as the *distance factor* in the remainder of the paper.

Of course, the actual values of A and B would be quite different for various urban settings, and even across different regions within a single urban setting. Regardless, these two parameters give us convenient knobs to capture and explore the effect of obstacles in VANET simulations.

7. IMPLEMENTATION

The mobility models are implemented as independent C++ programs that output mobility files, which serve as input to the NS2 simulations. The initial vehicle positions and their destinations are uniformly random. Each vehicle follows the shortest path to its destination. Upon reaching a destination, the vehicle begins its journey to another random destination. Each model takes a time parameter (in seconds). For SSM, this parameter denotes the duration each vehicle stops at intersections; for PTSM it denotes the maximum duration that each vehicle stops at the head of empty intersections; for TLM, it denotes the green light duration for each opposing pair of roads at an intersection. PTSM and TLM have a 1s vehicle startup delay.

The street topology is specified in a file that stores the road identifiers and the starting and ending road coordinates. All roads are modeled as bidirectional roads. SSM and PTSM assume a single lane in each direction of every road, whereas TLM provides the option for modeling multiple lanes. In SSM and PTSM, vehicles always travel within 5 miles/hour of the street speed limit. TLM has a slightly different mechanism with vehicles accelerating from rest to reach the speed limit, and then decelerating to stops. The acceleration and deceleration rates were 3 meter/second².

Implementations of MEPPR, MEIR, and obstacle models used the TLM mobility model. MEPPR and MEIR involved enhancing the NS2 simulator. The identity of the mesh nodes and their positions are specified from a separate file. Mesh nodes were positioned at a subset of street intersections chosen using the uniform random distribution. The implementation of MEPPR reuses AODV's implementation in NS2, which allows every node to participate in routing decisions. However, the MEIR deployment required changes in AODV to ensure that only the mesh nodes participate in routing and forwarding, whereas mobile nodes act as sources and destinations. For modeling obstacles, NS2's radio propagation model is modified according to Equation 3, such that parameters A and B can be configured.

8. PERFORMANCE EVALUATIONS

Table 1 summarizes the default values of parameters using NS-2 simulator [15]. We compared SSM, PTSM, TLM, the Random Waypoint Model (RWM) [13] and the Rice University Model (RUM) [19]. RWM captures mobility in an

Parameter	Default Value(s)
Simulation Time	900s (plus 450s warmup)
Routing Protocol	AODV
Transmit Range	250m
Number of Nodes	100
CBR Sources	4 pkts/sec and 64 byte pkt
Mobility Models	RWM, RUM, SSM, PTSM, TLM
Topologies	1200 × 1200m Grid, Real Map
Max. Wait Time	SSM-3s, PTSM-30s($p = 0.5$), TLM-30s
Max. Node Speed	35 mph
Accel./Decel.	3 meters/sec ² for TLM

Table 1: NS2 Wireless Simulation Parameters

open field with no obstacles, roads, or intersections. RUM simulates roads in a real map, but vehicles do not stop at intersections. For controlled experiments, we varied the block sizes in a grid topology over a 1200m × 1200m area. We also used real world street maps extracted from the US Census Bureau TIGER [21] database, which specifies road type, speed limit, and number of lanes. Each simulation run lasted 900s, preceded by 450s warmup period. Each data point is averaged over 5 to 10 different mobility patterns and node placements to attain a 95% confidence interval.

8.1 Number of Mobile Nodes

This section compares various mobility models with varying numbers of mobile nodes in a 1200m × 1200m grid topology with a block size of 200m × 50m. Figures 2 and 3 compare the delivery ratio and end-to-end delay among all mobility models. SSM had a wait time of 3 seconds. PTSM had a maximum wait time of 30 seconds. TLM switched signals with a periodicity of 30 seconds and used two lanes in each direction with acceleration-deceleration of vehicles enabled.

The results indicated that the RWM yields the lowest delivery ratio and the maximum end-to-end delay, for this particular topology. The range of performance variation across various models highlights our point regarding the importance of fidelity of mobility models in VANET simulations.

We observe that the delivery ratio increases with the number of nodes, up to 100 nodes, as the connectivity of the communication graph increases. Then the delivery ratio starts decreasing as the number of nodes increases further. This behavior is due to the increased channel contention as the large number of nodes leads to a flood of control messages in the network. The end-to-end delay in Figure 3 displays the opposite trend: it first decreases as the number of nodes increases, and then there is a sharp increase thereafter. We also experienced certain NS2 constraints as the number of nodes increased. The simulation time became a concern because we needed to explore a large parameter space. Also, the resource requirements of memory and storage (for output traces) became prohibitive. Additionally, with a large number of nodes, the confidence intervals of performance numbers widen significantly, further requiring more repetitions to reduce the variance of the results. Unless specified, we used 100 nodes in the remaining evaluations.

To understand the sensitivity of various mobility features (e.g. multi-lane roads and acceleration-deceleration of vehicles), we repeated the same experiment on TLM, with combinations of features enabled/disabled (Figure 4). The results indicate that acceleration-deceleration led to a significant increase in the delivery ratio because this feature reduces the average speed of vehicles. Thus, network routes are more stable. Additionally, the performance difference between the single-lane and multilane models is not noticeable below 100 nodes. However, with acceleration-deceleration disabled, it becomes noticeable beyond 100 nodes as the

channel contention begins to rise. It is interesting to note that, once the acceleration-deceleration is enabled, the difference between the single-lane and multi-lane models becomes negligible. At first glance, multiple lanes without acceleration and deceleration differ from single lanes without acceleration and deceleration. However, the confidence interval is rather wide. After checking the average vehicle speed (Figure 10), and the average percentage of nodes moving at a given time (Figure 8), the two models appear to have the same clustering effects at intersections with cars moving at a similar average speed. Therefore, we believe that the difference is largely within margins of statistical errors. Thus, with our experimental settings, the additional complexity of modeling multiple lanes does not seem to significantly affect performance.

8.2 Maximum Speed and Wait Times

We varied the maximum speed limit of vehicles in Figure 5 and observed that mobility models that place greater restrictions on node movement yield higher delivery ratios. One such restriction is the wait time at intersections. To further understand this effect, we varied the maximum wait time of nodes at intersections (Figure 6). The results brought out an interesting aspect of this study. As expected, the RUM model yields the lowest delivery ratio due to its highly dynamic pattern of mobility. However, in contrast to our earlier experiments, SSM yields a higher delivery ratio compared to PTSM. The reason is that SSM results in a more static network than PTSM does, where nodes are forced to stop at all intersections. On the other hand, PTSM nodes at intersections decide with a 50% probability to wait. The single-lane, no acceleration-deceleration TLM displays a marginally lower delivery ratio than PTSM for the same wait time because coordinated traffic lights provide a slightly higher rate of churn compared to PTSM. However, the addition of multiple lanes and acceleration-deceleration to TLM yields the highest delivery ratio among these models. This result, combined with our earlier observation about negligible impact of modeling multiple lanes, suggests that the introduction of acceleration-deceleration effectively slows down the vehicle speeds most significantly and dampens the changes in the network topology. However, these results are also dependent upon other factors, such as block sizes, which we will consider next.

8.3 Effect of Block Sizes

The block sizes in the topology play an important role in determining the performance of VANETs. With large block sizes, vehicles spend more time in traversing between intersections; thus, nodes are mobile more often. This increased mobility leads to a weakened connectivity in the network, and a corresponding drop in the delivery ratio. To validate this hypothesis, we conducted experiments with varying block sizes in a 1200m × 1200m area. Figure 7 largely confirmed our hypothesis – as the block size increases, the delivery ratio decreases. The RUM model is not sensitive to block sizes, since nodes do not stop at intersections. With the largest evaluated block, SSM outperforms PTSM due to a lower churn rate of routes, illustrating the interplay between block sizes and wait times in VANET simulations.

8.4 Analysis of Increased Mobility

The results of our experiments showed a distinct trend between the performance of various mobility models -TLM resulted in the highest delivery ratios, and the performance




Figure 2: Delivery ratio vs. number of mobile nodes.




Figure 3: End-to-end delay vs. number of mobile nodes.




Figure 4: Delivery ratio vs. number of mobile nodes for TLM variants.




Figure 5: Delivery ratio vs. maximum speed of vehicles.




Figure 6: Delivery ratio vs. maximum wait time at intersections.




Figure 7: Delivery ratio vs. increase in block size.

did not degrade considerably with an increase in the number of mobile nodes; PTSM showed a higher delivery ratio than SSM, and the throughput obtained through use of these models was considerably higher than RUM. This brings into context our hypothesis that varying the degree of mobility (node speed) within these networks is the reason for differing performance. In SSM, each node is forced to stop at each intersection. On the other hand, PTSM nodes stop only at non-empty intersections and some of the empty intersections. However, the default wait times for PTSM are higher as compared to SSM. This leads to a network that is effectively more static when compared to SSM, with better connectivity and corresponding performance improvements. TLM eliminates the probabilistic behavior of traffic lights and introduces acceleration and deceleration of vehicles, which leads to an even more stable network. To gain a detailed understanding, we identified metrics that measured the mobility of the nodes and the clustering of vehicles at intersections. The first metric provided us with a measure of the fraction of nodes we expected to actually be mobile at any given instant. The second metric was the extent of clustering at intersections. The number of clusters of vehicles could be treated as an effective number of nodes in the network, since all the nodes in a cluster displayed similar connectivity to nodes outside the cluster. The third metric measured the average speed.

Average Number of Mobile Nodes: To compute this metric we determined the number of nodes that are not waiting in a queue at any intersection. We took samples each second, averaged them over the simulated lifetime, and represented the result as a percentage of total nodes.

The first observation is that for the same wait time, varying the number of nodes does not appear to affect the percentage of mobile nodes significantly. This implies that the topology and wait time are more influential to the percentage of moving nodes compared to the number of nodes, up to 400 nodes within a $1200m \times 1200m$ area. Under similar conditions of wait time and topology, SSM is less mo-

bile when compared to PTSM as expected. The introduction of acceleration-deceleration of vehicles to TLM increases the percentage of moving nodes in the network significantly, as slower average speeds reduce the chance of nodes being queued at intersections. To illustrate the effect of the wait time, we also evaluated both PTSM and TLM with a similar value of the wait time. The plots indicate that for the same wait time of 10 seconds, PTSM is more mobile than SSM, with PTSM having an average of 85% of the nodes moving at any time compared to 68% for SSM.

Average Number of Clusters: Stopping of nodes at intersections effectively creates many clusters all over the network. Connectivity among the nodes within a cluster is near perfect (minus the network contention effects). On the other hand, if one node in the cluster cannot reach a distant node outside the cluster, then most likely all nodes in the cluster are unable to reach the same distant node. The number of such clusters can be treated as the effective number of (logical) nodes in the VANET at any time. Thus, we postulate that clustering has an effect similar to decreasing the number of nodes in the network.

To estimate the number of effective nodes, we divided the topology into $60m \times 60m$ regions, counted the number of regions containing at least a node each second, and took the average. Figure 9 shows that when the number of nodes increases, the number of effective nodes grows sub-linearly as more nodes are clustered at intersections. TLM resulted in a marginally greater number of effective nodes as compared to PTSM, for a similar wait duration of 30 seconds. This indicates a reduced clustering effect in TLM – a consequence of the reduced average speeds of the vehicles. We also studied the variation of this effect with the maximum wait time at intersections. With a wait time of 10 seconds, we observed that SSM with a wait time of 3 seconds resulted in a similar value as that of PTSM, which is consistent with our findings in Figure 8. Interestingly, we observed that acceleration-deceleration and multiple lanes do not significantly impact the difference in clustering level between the




Figure 8: Percentage of mobile nodes at a given time.




Figure 9: Clustering Effect: Decreasing slope of plots indicates increased clustering.




Figure 10: Average speed for various models.

various versions of TLM. This indicates that the difference across TLM variants is mainly due to average speed.

Average Speed of Vehicles: We computed the average speed for each vehicle as the ratio of the entire distance it travels during the simulation and the simulated time (Figure 10). We observed that PTSM results in lower average speeds compared to SSM, because of the longer wait times involved at intersections. For TLM variants, the addition of acceleration-deceleration leads to a significant decrease in average speed, which translates into higher delivery ratios. Also observe that TLM with multiple lanes does not noticeably affect the average speed compared to single lane.

8.5 Real Map Results

Having the insights into the various factors affecting VANET performance in grid topologies, we conducted experiments using real maps extracted from the TIGER database. We performed a set of experiments using a smaller section of the map used by RUM [19]. The original map was $2400m \times 2400m$, but the NS2 simulations at this size do not scale due to the large number of nodes required (or conversely, one needs to set unrealistic transmission ranges) to maintain meaningful delivery ratios. To address this problem, RUM [19] used a transmission range of 500 meters, which we considered to be too large for our settings. Hence, we decided to maintain the original default NS-2 setting of 250 meters transmission range, with a truncated map size of $1200m \times 1200m$. In our experiments using this map, we observed that the delivery ratio for each model increased with the number of nodes up to 100 nodes, followed by a rapid degradation in performance thereafter. However, the performance using TLM remained constant up to almost 200 nodes. These results reconfirm our understanding regarding the correlation between topology and mobility, and between the mobility and performance.

For another experiment, we extracted a map of Tallahassee, over an area of $2000m \times 2000m$. The results in this case were different from what we had seen so far, owing to a much larger area as compared to the first map. In this experiment, we were able to observe the effect of network partitioning due to the large area and the initial low density of nodes. This effect was also strengthened due to the stoppages enforced by our mobility models – once a node is in the waiting state at an intersection, it is highly likely to communicate with other nodes in other intersections due to the large size of the map. The delivery ratios were initially very low with a small number of nodes, and the performance actually improved as the number of nodes was increased up to 200. This was in contrast to the results obtained with the smaller map, where the performance went down with an increase in the number of nodes, perceivably due to network

saturation. This reinforces our observation that these simulation results must be analyzed with the topology in mind. However, our basic understanding remains valid. The delivery ratio with TLM still remains higher than that with PTSM, SSM, and RUM due to a lower network churn.

8.6 Mesh-Enhanced VANET

In this section, we try to understand the performance of MEPPR and MEIR, without considering obstacles. Figure 11 shows the effects of varying the number of mobile nodes on the delivery ratios and end-to-end delay in the MEPPR deployment scenario. The two plots correspond to fixed number of mesh nodes at 171 mesh nodes (one per intersection) and 40 (approximately 23% of intersections). For 171 mesh nodes, as the number of nodes participating in the routing process increases and the network becomes dense, the resulting channel contention increases. Consequently, the delivery ratio of MEPPR deployment degrades. On the other hand, with only 40 mesh nodes, the MEPPR deployment maintains a high delivery ratio (and low end-to-end delay) with the addition of mobile nodes. To rule out the performance degradation due to only the total number of mobile and static nodes, we also extended the number of mobile nodes to 170 (not shown in the plot) and observed no significant performance degradation with only 40 mesh nodes. These results confirm our hypothesis that the channel access contention generated by the number of nodes participating in the routing process is an important factor in VANET’s performance.

Figure 12 show the effects of increasing the number of mobile nodes on the delivery ratio under the MEIR deployment. We again used 171 and 40 mesh nodes. The plot shows that the delivery ratio does not vary significantly for either setting as the number of mobile nodes increases. Since the mobile nodes do not participate in the routing process (they are merely sources and sinks for data packets), and the mesh nodes that participate in routing decisions are stationary, the resulting routes are much more stable than in MEPPR routes. MEIR routes change only when the mobile endpoints move out of range of their immediate mesh node.

Interestingly, in the sparser case of 40 mesh nodes in MEIR, since the mobile nodes do not participate in the routing process, the resulting network coverage and connectivity is poorer than the case of 171 MEIR mesh nodes (in Figure 12) and 40 MEPPR mesh nodes (in Figure 11). Consequently, the delivery ratio is also lower (around 90%).

Figure 13 shows the effect of varying the number of stationary mesh nodes, with fixed the number of mobile nodes (30 or 35). Since both mesh nodes and mobile nodes perform routing, a relatively small number of mobile nodes combined

with mesh nodes can achieve good routing coverage and delivery ratio. On the other hand, too many mesh nodes severely limit the number of mobile nodes due to channel contention. This is seen in the case of 35 mobile nodes and 171 mesh nodes in the above graphs.

Figure 14 shows how the number of mesh nodes affects the delivery ratios with MEIR. Clearly, a sufficient number of mesh nodes are needed to achieve good routing coverage and delivery ratio. However, MEIR scales better when compared to MEPPR because the routes are more stable, resulting in fewer route breakages, route discovery, and recovery events.

Thus, in a dense network, where the total number of nodes is high, MEPPR deployment can lead to decreased performance as a result of channel contention. In this case, the MEIR deployment scenario is preferable. In addition, MEIR deployment can scale better with increased network loads. On the other hand, in a sparse network with a smaller number of nodes, MEPPR deployment provides better routing coverage and higher connectivity.

8.7 Obstacle Representation

Figure 15 shows how obstacle factor A in Equation 3 affects network performance for both routing configurations. An increasingly negative value of obstacle factor A should lead to a decrease in signal strength at receivers and lower performance. This is observed when $A < -35$ for both routing configurations. The default value of A in the NS2 propagation model is -31 , which corresponds to a total absence of obstacles. However, it is interesting to note that in the MEPPR deployment scenarios, when $A > -15$, performance degrades. For such a negative value of obstacle factor A , the signal strength at the receiver is high enough to cause unwanted reception and interference among these receptions. This is not observed for MEIR because the static mesh nodes maintain a fixed distance from one another throughout the simulations. Figure 16 shows how the distance factor B in Equation 3 affects network performance for both mesh-enhanced routing configurations. A more positive value of distance factor B should reduce signal strength at receivers and decrease performance. This is observed in cases of values of $B > 21$ for both routing configurations. However, MEPPR deployment performs better as compared to MEIR deployment. This is the result of the mobile nodes' participation in routing to enhance connectivity and coverage. For a high value of distance factor, network connectivity is still better in the MEPPR deployment as more nodes are reachable through the mobile nodes. Thus, we see that obstacles could potentially be represented in network simulation through a few key parameters. MEPPR improves coverage and network performance in the presence of obstacles.

9. RELATED WORK

To date, studies in the fields of mobility modeling, obstacle modeling, and mesh networks have been largely performed in isolation. To the best of our knowledge, our work is the first attempt to synthesize and systematically evaluate the individual/combined effects of various factors on urban VANET deployment. In contrast to earlier works, the focus of our paper is not to recommend any one model over another but to understand and evaluate the performance impact and significance of various factors on VANET simulations.

The most used mobility model in literature is the Random Waypoint Model (RWM) [13]. Every node selects a random destination and speed and then moves to that destination

with the chosen speed, pauses, and then moves again to another random destination. Other similar open-field models include the Random Walk, Random Direction Model and the Boundless Simulation Area Model [2]. Camp [2] observed that the spatial distribution of nodes in such models is toward the center of the simulation area. [1, 16] have attempted to improve RWM to make it more realistic, though not within the context of VANETs. Bai [7] introduced the Freeway and Manhattan mobility models on roads specified through maps. A vehicle's path from an intersection was decided using a fixed probability. The vehicles did not pause, stop, or queue up at intersections. Saha et al. [19] modeled mobility of vehicles on real street maps, obtained from the TIGER database [21]. Their model, which we call RUM in this paper, does not enforce any traffic rules with the resulting performance similar to RWM. We used RUM as one of the base cases for comparison. Choffnes and Bustamante [4] recently introduced a vehicular mobility model for urban environments. With their simulators configured to generate delivery ratios between 0.05 and 0.3, they observed that the network performance in such a network was significantly different from the RWM. Our evaluations confirmed their findings regarding the correlation between network performance and the simulated topology. Additionally, our evaluations in this paper used reasonable parameter settings to generate delivery ratios over 90% that are within the usable range. Because the model was implemented in the SWANS simulator, it was difficult to evaluate it without significant porting effort to NS2. [10] proposed several theoretical models like city area, area zone, and unit street. [6] points out that these models lack specific details for actual node movement calculations. Proprietary tools for modeling transportation systems are also available [17, 24, 5], though their details are not fully understood.

While wireless mesh networks have been studied extensively in isolation, we evaluate their impact on urban VANETs using realistic mobility models. Microsoft [20] has proposed self-organizing community mesh networks. RoofNet [18] provides broadband Internet access to residential subscribers. Motorola has proposed mesh-based network solutions for intelligent transport systems and communication [14].

The effect of obstacles on wireless networks is relatively little studied. [3] points out that commonly used radio propagation models for indoor MANET evaluations are highly inaccurate and relative protocol performance varies highly depending upon the model. [12] models a terrain by specifying the shapes and sizes of obstacles. The effect of obstacles on signal propagation is determined by a static table based on the type of obstacle. [9] presented the principles of a WCDMA radio network simulator that accounts for path-loss, shadowing, and fast fading effects in radio signal propagation. In contrast to earlier efforts, our work uses empirical measurements around urban city blocks to derive a parameterized model that can be tuned to study the effect of obstacles on radio signal propagation in VANETs.

10. CONCLUSIONS

Simulation models play a critical role in the evaluation of Vehicular Ad Hoc Networks (VANETs). In this paper, we have systematically evaluated the sensitivity of various modeling details on VANETs through step-by-step implementation of details in the mobility, radio propagation, and deployment models in urban settings. We proposed a series




Figure 11: Delivery ratio vs. number of mobile nodes in MEPPR.




Figure 12: Delivery ratio vs. number of mobile nodes in MEIR.




Figure 13: Delivery ratio vs. number of mesh nodes in MEPPR.




Figure 14: Delivery ratio vs. number of mesh nodes in MEIR.




Figure 15: Delivery ratio vs. obstacle factor A in Equation 3.




Figure 16: Delivery ratio vs. distance factor B in Equation 3.

of simulation models that account for various urban constraints such as street layout, traffic rules, multi-lane roads, acceleration-deceleration, and RF attenuation due to obstacles. Our evaluations, using both real and controlled synthetic maps, provide many interesting insights. VANET performance is sensitive to the clustering of vehicles at intersections and acceleration-deceleration of vehicles. Simulation of multiple lanes and synchronization at traffic signals only marginally impact VANET performance. However, modeling of multiple lanes might be fundamental to applications such as Forward Collision Warning and Lane Change Assistance systems, which we do not consider. In dense VANETs, performance improves significantly when routing decisions are limited to a wireless backbone of mesh nodes compared, whereas in sparse VANETs, performance improves when vehicles also participate in ad hoc routing. Finally, we measured signal strengths around urban city blocks and showed that the effect of signal attenuation due to obstacles can potentially be parameterized via empirical real-world measurements. Our evaluations provide a starting point to understand and develop better VANET simulation models.

11. REFERENCES

- [1] C. Bettstetter. Smooth is better than sharp: a random mobility model for simulation of wireless networks. In *Proc. of MSWIM*, 2001.
- [2] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *Wireless Communications & Mobile Computing (WCMC): Special issue on Mobile Ad Hoc Networking*, 2(5):483–502, 2002.
- [3] A.L. Cavilla, G. Baron, T.E. Hart, L. Litty, and E. de Lara. Simplified simulation models for indoor MANET evaluation are not robust. In *Proc. of the IEEE SECON*, Oct. 2004.
- [4] D.R. Choffnes and F.E. Bustamante. An integrated mobility and traffic model for vehicular wireless networks. In *Proc. of VANET*, 2005.
- [5] CORSIM Microscopic Traffic Simulation Model. <http://mctrans.ce.ufl.edu/featured/tsis/version5/corsim.htm>.
- [6] V. Davies. Evaluating mobility models within an ad hoc network. Master's thesis, Colorado School of Mines, 2000.
- [7] F. Bai, N. Sadagopan, and A. Helmy. The IMPORTANT framework for analyzing the impact of mobility on performance of routing protocols for adhoc networks. *AdHoc Networks Journal*, 1:383–403, Nov 2003.
- [8] Z. Fu, B. Greenstein, X. Meng, and S. Lu. Design and implementation of a TCP-friendly transport protocol for ad hoc wireless networks. In *Proc. of ICNP*, Washington, DC, USA, 2002.
- [9] S. Hamalainen, H. Holma, and K. Sipil. Advanced WCDMA Radio Network Simulator. In *Proc. of the IEEE Intl. Symposium on Personal, Indoor and Mobile Radio, Osaka, Japan*, Sep. 1999.
- [10] D. Tsirikas, J. Markoulidakis, G. Lyberopoulos. Mobility modeling in third-generation mobile telecommunications systems. *IEEE Personal Commun.*, 4:41–56, August 1997.
- [11] R. Jain. *The Art of Computer Systems Performance Analysis*. John Wiley and Sons Inc. New York, 1991.
- [12] A. Jardosh, E.M. Belding-Royer, K.C. Almeroth, and S. Suri. Towards realistic mobility models for mobile ad hoc networks. In *Proc. of MobiCom*, 2003.
- [13] D.B. Johnson, D. Maltz, and J. Broch. *The Dynamic Source Routing Protocol for Multihop Wireless Ad Hoc Networks*, chapter 5, pages 139–172. Addison-Wesley, 2001.
- [14] Mesh-Enabled Solutions for Intelligent Transportation. <http://www.motorola.com/mesh/pages/applications/its.htm>.
- [15] NS2 Network Simulator. <http://www.isi.edu/nsnam/ns/>.
- [16] S. PalChaudhuri, J.-Y. Le Boudec, and M. Vojnovic. Perfect simulations for random trip mobility models. In *Proc. of Symposium on Simulation*, 2005.
- [17] Quadstone, Inc. The PARAMICS Transportation Modeling Suite. <http://paramics.quadstone.com/>.
- [18] Roofnet. <http://pdos.csail.mit.edu/roofnet/doku.php>.
- [19] A.K. Saha and D.B. Johnson. Modeling mobility for vehicular ad-hoc networks. In *Proc. of VANET*, 2004.
- [20] Self-Organizing Neighborhood Wireless Mesh Networks. <http://research.microsoft.com/mesh/>.
- [21] Topologically Integrated Geographic Encoding and Referencing. <http://www.census.gov/geo/www/tiger>.
- [22] T. Tugcu and C. Ersoy. How a new realistic mobility model can affect the relative performance of a mobile networking scheme. *Wireless Communications and Mobile Computing*, 4:383–394, June 2004.
- [23] Wavemon: Wireless Device Monitoring Application. <http://packages.debian.org/unstable/net/wavemon>.
- [24] H. Xu and M. Barth. A transmission-interval and power-level modulation methodology for optimizing inter-vehicle communications. In *Proc. of VANET'04*.