An Integrated Mobility and Traffic Model for Vehicular Wireless Networks

David R. Choffnes Fabián E. Bustamante
Department of Computer Science, Northwestern University, Evanston, IL 60201
{drchoffnes,fabianb}@cs.northwestern.edu

ABSTRACT

Ad-hoc wireless communication among highly dynamic, mobile nodes in a urban network is a critical capability for a wide range of important applications including automated vehicles, real-time traffic monitoring and vehicular safety applications. When evaluating application performance in simulation, a realistic mobility model for vehicular ad-hoc networks (VANETs) is critical for accurate results. This paper analyzes ad-hoc wireless network performance in a vehicular network in which nodes move according to a simplified vehicular traffic model on roads defined by real map data. We show that when nodes move according to our street mobility model, STRAW, network performance is significantly different from that of the commonly used random waypoint model. We also demonstrate that protocol performance varies with the type of urban environment. Finally, we use these results to argue for the development of integrated vehicular and network traffic simulators to evaluate vehicular ad-hoc network applications, particularly when the information passed through the network affects node mobility.

1. INTRODUCTION

The community is increasingly interested in developing network protocols and services for vehicular ad-hoc networks (VANETs). Due in part to the prohibitive cost of deploying and implementing such systems in the real world, most research in this area relies on simulation for evaluation. A key component of these simulations is a realistic vehicular mobility model that ensures conclusions drawn from such experiments will carry through to real deployments.

Unlike many other mobile ad-hoc environments where node movement occurs in an open field (such as conference rooms and cafés), vehicular nodes are constrained to streets often separated by buildings, trees or other objects. Street layouts and different obstructions increase the average distance between nodes and, in most cases, reduce the overall signal strength received at each node. We argue that a more realistic mobility model with the appropriate level of detail [10] for vehicular networks is critical for accurate network simulation results. We propose the STRAW (STreet RAndom Waypoint) mobility model, which constrains node movement to streets defined by map data for real US cities and limits their mobility according to vehicular congestion and simplified traffic control mechanisms.

We evaluate and compare ad-hoc routing performance for vehicular nodes in diverse urban environments using STRAW to that in an open field using the classical random waypoint (RWP) model. We show that the performance of wireless network protocols in urban environments is dramatically different than that in an open-field RWP model and, further, that the type of urban environment can significantly impact the performance of a protocol. In sum, this study clearly illustrates that a realistic mobility model and a variety of target environments are essential to ensure that VANET applications meet their performance criteria when deployed. Although trace-based evaluations provide accurate mobility in simulation, we argue for an integrated vehicular traffic and wireless network simulator when VANET applications are expected to influence the mobility of participating vehicles.

The remainder of this paper is structured as follows. In Section 2 we motivate the need for urban mobility models in ad-hoc networks. Section 3 introduces STRAW and details its implementation. We describe our simulation environment and present results in Section 4; we discuss these results in Section 5. We provide related work in Section 6 and conclude in Section 7 by suggesting some future directions for improving the accuracy of simulations for VANETs.

2. BACKGROUND

There is extensive literature on routing for ad-hoc networks in general and mobile ad-hoc networks (MANETs) in particular. Some of the protocols that have achieved prominence include topology-based ones (such as DSDV [27], DSR [13], AODV [26] and MRP [24]) that rely exclusively upon IP addresses to locate nodes and location-based protocols (e.g., DREAM [2] and GPSR [16]/GLS [20] [23]) that use geographical position information to determine routes.

The different proposed protocols are commonly analyzed and/or compared against competing or ideal ones in terms of metrics such as packet delivery ratio, throughput, latency and overhead. Given the cost, complexity and other limitations (primarily lack of repeatability) of real-world deployments, most reports in the literature today rely on simulators (e.g., [21, 37, 1]) for experimentation.

VANET'05, September 2, 2005, Cologne, Germany.

For these studies, researchers often adopt a common set of simulation parameters, such as:

- The number of nodes is small (i.e., ≤ 100).
- Nodes move in an open field.
- Nodes move according to a random waypoint model [36] or the Manhattan mobility model [9] with arbitrary pause times and often with arbitrarily uniform speed distributions between 0 and 20 m/s.
- Nodes transmit signals that propagate without error to other nodes within a radius of 250 m [18].

Such parameter settings are clearly inadequate for many MANETs, and particularly for VANETs. For example, in [18], the authors have shown that the relationship between distance and signal reception between two nodes is, at best, weakly correlated over large distances. Further, besides settings such as conventions in large conference halls, it is difficult to imagine many scenarios where nodes will move in an open field and/or in a way that can be accurately modeled by random waypoints. Specifically in VANETs, the number of nodes is generally large, the mobility of these nodes is constrained by roads and their velocities must be adjusted according to traffic control mechanisms (e.g. stop signs and traffic lights), speed limits and the level of congestion in the vehicular network.

The ad-hoc research community is increasingly aware of the limitations resulting from some of these simplifying assumptions [18, 15]. In the context of VANETs, various research groups are designing experiments that better model real vehicular traffic scenarios. For example, [17] studies the behavior of the MAC layer in a vehicular environment using arbitrary road plans while [35] and [33] use CORSIM, a proprietary vehicular traffic simulator, to provide mobility traces for the simulation. Our work contributes to this effort a new vehicular mobility model and its implementation as part of a publicly available wireless network simulator [1].

3. STRAW MOBILITY MODEL

STRAW incorporates a simple car-following model with traffic control to introduce vehicular congestion, which models real traffic conditions. The model also incorporates the notion of an *enabled vehicular penetration ratio*–the percentage of cars equipped with radios and actively communicating. The integration of a vehicular traffic mobility model makes possible to experiment with applications where the content of disseminated data (e.g., traffic information) is used to dynamically alter the routes taken by participating nodes during the course of a run.

STRAW relies on street plans to build a road map for the specified target region. For each *road segment* - the portion of a road between two intersections - STRAW maintains information such as the road class (e.g., residential road or divided highway), the start and end points of the segment, the name of the street and a list of points along the segment if it is not a straight line. STRAW also provides at least one lane in each direction on which vehicles can move. To determine the initial positions of vehicles on the field, we use a random street placement model that places a vehicle in a lane of random street just before an intersection. If another vehicle is already in that lane, the new vehicle is placed behind the existing one. All vehicles are initially stopped (i.e., assigned a speed of 0).

We currently use the freely available US Census Bureau's TIGER data files [22] as the source of street plans. The TIGER data is provided as packages, organized by state county, containing files that provide information about various geographic features, including locations and dimensions of schools, parks, roads and other landmarks. We use these files to extract the names, locations and shapes of roads, their corresponding street addresses and their road "classes", which can be used to approximate the speed limit and capacity of each road.

In the following subsections, we describe STRAW in detail and discuss its implementation.

3.1 Intra-segment Mobility

When the simulation starts, nodes move according to the *car-following* model [30] such that they will attempt to accelerate at a constant rate of up to 5 mph per second to move with a speed equal to the maximum speed for the current driver.[1] In [8], the authors report that observed speeds are normally distributed with a center at the posted speed limit. Unfortunately, we could not find a widely accepted standard deviation for this distribution. Thus, we set the maximum speed for a node to the speed limit for the current road plus a Gaussian distributed value with a zero mean and a 4 mph standard deviation. The car will alter its speed according to the following rules:

- *The car encounters an intersection and the next road segment on which it will travel is full.* In this case, the car stops before the intersection and remains stopped until there is room in the next road segment.
- *There is a car in front of the current car.* In this case, the node will slow down to the speed necessary to maintain a speed-based following distance between the current node and the node in front of it. We use the simple formula cited in [30]:

  $$S = \alpha + \beta V + \gamma V^2,$$

  where

  $S =$ the following distance
  $V =$ the current vehicle's speed
  $\alpha =$ the vehicle length
  $\beta =$ the reaction time (we use 0.75 seconds)
  $\gamma =$ the reciprocal of twice the maximum average deceleration of the following vehicle (we use the empirically-derived value, $0.0070104 s^2/m$ [30])

  If the car in front of the current car is moving faster than the current car, no speed adjustment is necessary.
- *The car encounters traffic control.* In this case, the car will slow down (at a uniform acceleration) before an intersection with a red stoplight or a stop sign; if the stoplight turns green, the car attempts to increase its speed if possible.

[1] We acknowledge that acceleration rates are hardly uniform in real life. Future iterations of the mobility model will include more accurate acceleration curves when such data becomes available.

- *The car turns onto a new street.* In this case, the car slows down before the intersection to make the turn at a reasonable speed (5 mph), then accelerates, if possible, to the highest speed it can attain given the other constraints.

Because vehicles for our experiments are constrained to roads in downtown urban environments and therefore exhibit average speeds no larger than 12 m/s (26.8 mph) for our experiments, we update each vehicle's position once per second using its current speed and direction. We intend to incorporate speed-based position updates in future iterations of STRAW.

3.2 Inter-segment Mobility

This section discusses the implementation of our inter-segment mobility model, i.e., vehicular behavior at an intersection. Our simulator supports two levels of admission control at an intersection. The first form of admission control simulates common traffic control mechanisms. Our simulator supports stop signs and timed traffic lights. We expect that future iterations of the model will include triggered lights and guarded turns. Note that because we do not currently support lane changing, we also do not consider a vehicle's current lane when it attempts to make a turn. The second form of admission control simply ensures that there is room for the vehicle on the next road segment before it crosses the intersection.

Because real-world, per-intersection traffic control information is unavailable, the simulator currently assigns traffic control according to the class of road segments at each intersection. For example, a stop sign controls access when two local/neighborhood roads meet; a timed stoplight controls access when a "secondary" road and a state highway intersect. The full details of our implementation are discussed in the technical report [6].

3.3 Route Management and Execution

The Route Management and Execution (RME) component determines the path taken by each vehicle for the duration of the simulation. We include two models: simple intersegment mobility (Simple STRAW) and mobility with origin-destination (OD) pairs (STRAW OD). In the former model, the next segment to which a vehicle will move is determined stochastically at each intersection. This model maintains a single value to select the next segment on which a vehicle will travel: the probability that it will turn at any given intersection. In the latter one, the decision is based on the precomputed shortest path between the vehicle's specified origin and destination. Although this model more accurately represents vehicular motion, it incurs more runtime overhead and requires knowledge of driving patterns (i.e., origins and destinations) in a particular scenario.

3.4 Implementation Details

STRAW is currently implemented as an extension to SWANS (Scalable Wireless Ad Hoc Network Simulator) [1], a Java-based, publicly available, scalable wireless network simulator. SWANS runs atop JiST (Java in Simulation Time), a high-performance discrete event simulation engine that features low memory consumption and fast run times. It supports large numbers of nodes (>1,000,000) and defines an extensible set of highly configurable simulation abstractions to model numerous real-world components for various levels of realism in simulation. In addition to providing high performance and ease-of-use, SWANS provides a bytecode rewriter that automates the porting of networked Java applications to the simulator. The STRAW implementation requires less than 5400 lines of code (not including comments).

4. EVALUATION

In this section, we evaluate the impact of our mobility model on ad-hoc routing performance by comparing the performance of two well-known ad-hoc protocols, AODV [26] and DSR [13], when used with the random waypoint model to that using the STRAW mobility model. We first describe the experimental settings, then discuss our results.

4.1 Experimental Setup

SWANS provides implementations of Java standard network interfaces at the application layer, sockets at the network layer, UDP and TCP at the transport layer, AODV and DSR at the routing layer, 802.11 at the MAC layer and several path loss and fading models at the physical link layer. It also includes several node mobility models, including static, random waypoint and random walk, and two node placement models: random and grid.

In addition to incorporating our STRAW mobility model as described in Section 3, we have extended SWANS in multiple ways to evaluate routing performance in an urban environment. The set of network statistics was improved to include some absent useful metrics such as packet delivery ratio, latency and overhead. Several relatively minor bugs in the original AODV and DSR implementations were fixed.

We evaluate network performance using two reactive, address-based, point-to-point routing protocols: Ad-hoc On-Demand Distance Vector (AODV) and Dynamic Source Routing (DSR). We opted not to analyze proactive routing protocols (e.g., DSDV [27]) as they are unable to maintain an accurate view of the entire network in a highly mobile environment without initiating a "broadcast storm" that can prevent successful message transmission [32].

AODV is a next-hop routing protocol [26]. Route discovery occurs only when a source node transmits a packet to a destination for which no valid route is known and each node's routing table contains information only about the next hop on the path to a particular destination node.

We also evaluate network performance under DSR, another routing protocol [13] commonly evaluated in the literature. Similar to AODV routing, route discovery in DSR occurs only when a source node attempts to send a message to a destination for which there is no active valid route. Unlike AODV, DSR discovers source routes; i.e., upon successful route discovery, the originating node obtains an ordered list of node addresses along the path from origin to destination. More detailed descriptions of these protocols can be found in the associated references.

Our experiments were designed as follows. The simulator placed n nodes on streets contained in a rectangular region. Of these nodes, m nodes were assigned to be transmitters; these nodes transmit single UDP packets at a rate of s packets per minute. All of the nodes use an 802.11b MAC protocol operating at 2Mbps, share common radio properties typical of commodity wireless network cards and operate in an environment with a generic path loss model with shadowing [19], using exponent 2.8 and standard deviation 6.0,

Figure 1: Portion of downtown Chicago demonstrating a regular grid.

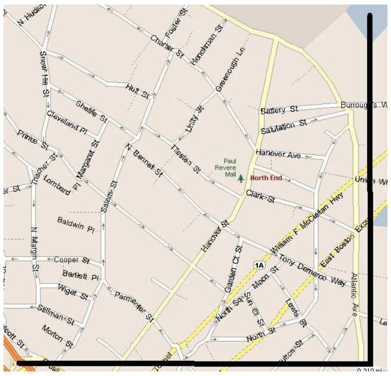

Figure 2: Portion of the North End of Boston, demonstrating a nonregular grid.

to determine signal strength at the receiver. Each simulation has a warm-up time of 60 seconds before any packets are sent and a resolution time of 60 seconds during which no packets are sent, to allow the last packets sent to reach their destinations. The total simulation time is 900 seconds, a common value in the literature. To determine the impact of street-constrained mobility on network performance, we conducted our simulations in three environments: an open field (no streets), a region of downtown Chicago (Figure 1) and a region of the North End of Boston (Figure 2). The downtown Chicago region is a regular, "Manhattan grid"; in contrast, Boston's North End is a nonregular road network.

Beyond the simplifying assumptions in our traffic model, we note that we did not incorporate a notion of buildings into our simulator, so signals are not blocked by obstructions. Also, the implementations of AODV and DSR do not include all of the optimizations recommended by the latest specifications. Thus, the results of our experiments should not be used to compare the performance between the two algorithms; rather, the results should be interpreted according to the relative performance for the same protocol when using different mobility models. Due to the low connectivity experienced by nodes using STRAW mobility, we believe that the missing protocol optimizations will not significantly impact the conclusions drawn from our results.

4.2 Results and Analysis

To determine the effects of street mobility on packet delivery ratio, we used the STRAW mobility model, resulting in a node speed that varied according to vehicular traffic conditions. To compare the results of STRAW to the random waypoint mobility model, we repeated each city experiment in an open field of the same size and allowed vehicle speeds to vary from 3 m/s to 15 m/s, with pause times of 3 seconds. We chose the speed range because it corresponds to the observed speed of uncongested traffic in our mobility model ($\approx$ 9 m/s) and the pause time approximates the time spent at a stop sign.

For all of our experiments, the field size is set to a square of 500 m sides. In our experience, using either AODV or DSR to route packets in a large test region led to abysmal performance when modeling commodity radios with realistic transmission ranges. We opted for the Simple STRAW configuration in which nodes turn at an intersection with some probability p and go straight otherwise. When a node decides to turn or cannot go straight, a new direction is chosen uniformly at random. We contend that for such a small region, the difference between mobility using OD pairs and that using the turn probability ratio is insignificant for the purposes of measuring routing performance. All of the experiments in this section use a value of 0.3 for p, providing a small but significant amount of entropy in a vehicle's path.

An important goal of STRAW is to model vehicle interaction when roads become congested. The regions that we chose required approximately 300 to 400 vehicles to significantly lower the average speed of vehicles in the network. Unfortunately, the high density of nodes along the road segments leads to crippling network congestion for AODV and DSR when every node is equipped with a radio. To generate results with vehicular congestion, we varied the penetration ratio such that the number of participating vehicles remained constant while varying the number of vehicles traveling in the road network.

We chose to offer a modest packet traffic load in these experiments to ensure the observed performance was not dependent on the routing protocol's capacity to handle heavy loads. We used 10 transmitters, each sending a 512-byte UDP packet every two seconds, to a randomly chosen destination. We varied the number of vehicles in the road network from 50 to 400, in increments of 50. For each of these scenarios, we evaluated performance when 50 nodes were equipped with radios and, for scenarios with more than 50 vehicles, we also equipped 100 nodes with radios. Figure 3 plots the average vehicle speed as a function of the number of vehicles in the region, both for Chicago and for Boston. Not surprisingly, the average speed of vehicles in a fixed-size region decreases as the density of vehicles in the region increases.

The first experiment using the STRAW mobility model compares the performance of DSR on a regular section of downtown Chicago, IL to an urban region of similar size in the North End neighborhood of Boston, MA. The North End features irregular road patterns, which decreases the average distance between nodes. Figure 4 shows the performance when 50 nodes are equipped with radios and Fig. 5 shows the performance with 100 radio-equipped vehicles. The random waypoint data points represent the average of nine runs with a standard deviation of 3.9% for 50 nodes and 4.7% for 100 nodes. The STRAW data points represent the aver-

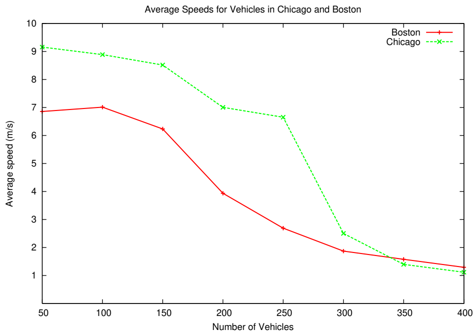

Figure 3: Average speeds of vehicles in the Chicago and Boston regions, as a function of the number of vehicles in the region.

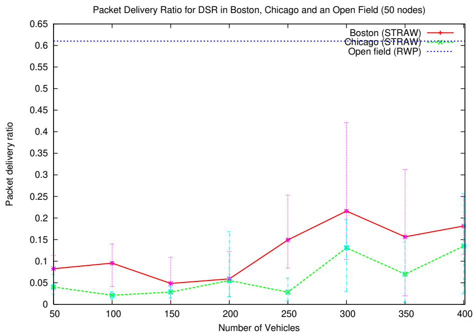

Figure 4: Comparison of DSR packet delivery ratio between mobility in downtown Chicago, Boston and the random waypoint model using 50 radio-equipped nodes.

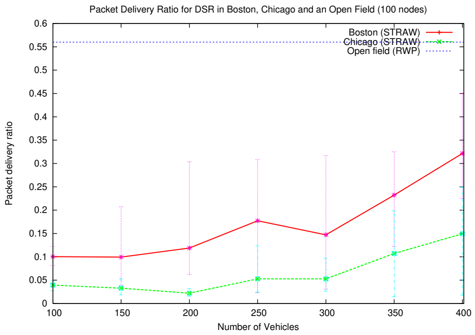

Figure 5: Comparison of DSR packet delivery ratio between mobility in downtown Chicago, Boston and the random waypoint model using 100 radio-equipped nodes.

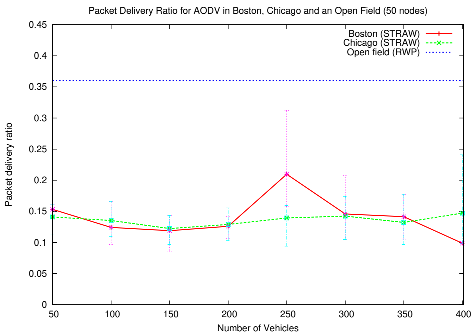

Figure 6: Comparison of AODV packet delivery ratio between mobility in downtown Chicago, Boston and the random waypoint model using 50 radio-equipped nodes.

age of five runs and the error bars represent the range of observed packet delivery ratios. Note that the range increases as the average vehicular speed decreases, indicating that congested vehicular networks are more sensitive to the initial placement of vehicles.

These results clearly illustrate the impact of more realistic mobility models in terms of delivery ratio. Constraining vehicle mobility based on city street maps instead of the more simplistic open field model result in a significantly lower (43% to 55%) mean delivery ratio. Despite the large variance in the STRAW data points, indicating unstable network connectivity, none of the data points come within one standard deviation of the RWP data points. In this experiment, the most significant cause for dropped packets using STRAW is due to attempts to send packets from intermediate nodes to nodes that were not within range of the transmitters. In this case, it appears that a combination of increased average internode distance (caused by the large city blocks) and the average vehicular speed (9.1 m/s or 20.4 miles per hour), leading to stale routing data, were the primary reasons for this effect.

Note that the packet delivery ratio in Boston is consistently higher (up to 17%) than that of Chicago. We believe that this is largely due to the decreased average internode distance caused by smaller "blocks". Further, the vehicles in this simulation experienced a slightly slower average speed, caused by the larger number of intersections (leading to increased traffic control). Finally, we note that the system experienced a higher packet delivery ratio with 100 radio-equipped nodes than with 50 under STRAW, indicating that a higher density of nodes can lead to increased performance (as long as the nodes do not increase the number of messages lost to radio interference).

Similar experiments were run employing AODV for routing. Figures 6 and 7 show the packet delivery ratios for 50 and 100 radio-equipped nodes, respectively, in the same region. The data points represent the average of ten runs with a standard deviation of 6.9% for 50 nodes and 4.4% for 100 nodes. As with the previous experiments, packet delivery ratios for the city environment are significantly lower than in the open field using random waypoint. Unlike the previous experiment, the packet delivery ratio is not significantly affected by the speed of the nodes; however, the packet delivery ratio does improve with increased node density. Note that for AODV, the packet delivery ratio does not differ significantly between the Chicago and Boston environments and the variance of those data points is not quite as large. This indicates that the impact of road plans on routing performance is not uniform across different routing protocols.

In the following experiments, we evaluate the effect of send rate on routing performance for AODV (Fig. 8) and DSR (Fig. 9). Note that the send rate from the previous figures, 30 packets per minute per transmitter, is not in-

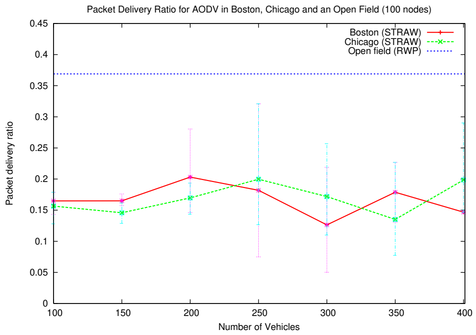

Figure 7: Comparison of AODV packet delivery ratio between mobility in downtown Chicago, Boston and the random waypoint model using 100 radio-equipped nodes.

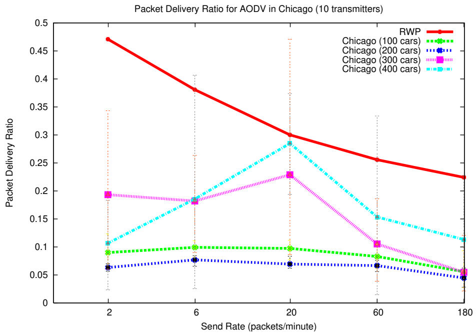

Figure 8: Comparison of AODV packet delivery ratio between mobility in downtown Chicago and the random waypoint model using 50 radio-equipped nodes and varying send rates.

cluded in these graphs. Both figures demonstrate that for relatively uncongested vehicular networks (i.e., 100-200 vehicles) the number of packets sent per minute has little effect on the packet delivery ratio, indicating poor connectivity in the system. As vehicles slow down due to vehicular congestion, (i.e., 300-400 vehicles), packet delivery ratio improves as nodes are better able to establish and maintain routes. In all cases, the packet delivery ratio curves for STRAW vary significantly from that of the RWP model.

For experiments with smaller numbers of vehicles (i.e., 100-200), vehicles move quickly and packet delivery ratio suffers due to unstable routes and limited connectivity. The connectivity is so poor that the nodes are relatively unaffected by the offered send rate. For experiments with larger numbers of vehicles, (i.e., 300-400), vehicular congestion slows node movement, leading to more stable routes and less routing overhead. We believe that the reactive nature of these protocols leads to lower performance when the send rate is low, as routes change significantly between successive data packets. As the send rate increases, the nodes are able to better maintain routes, leading to increased routing performance. Eventually, however, the send rate becomes so high ($>$ 20 packets/minute in these experiments) that the system is overwhelmed by interfering message broadcasts that inhibit successful route establishment and packet delivery. For all experiments, including the open field with

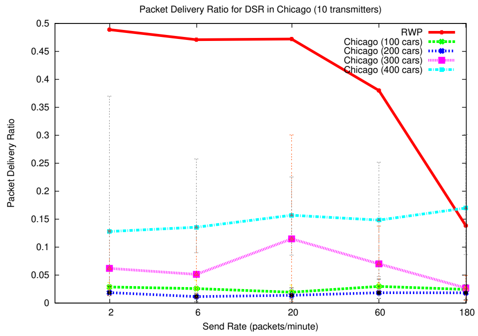

Figure 9: Comparison of DSR packet delivery ratio between mobility in downtown Chicago and the random waypoint model using 50 radio-equipped nodes and varying send rates.

random waypoint mobility, the packet delivery ratio suffers from the interference caused by higher send rates.

It is important to note that in both urban environments, the mean packet delivery ratio was significantly lower than that in an open field. Further, the range of packet delivery ratios for STRAW mobility never exceeded those in an open field for comparable average speeds (i.e., 100-200 nodes in Chicago). We attribute the low performance to the higher average internode distance imposed by street-constrained motion and the increased interference near intersections.

We opted to focus our analysis of the implication of mobility models employing two representative ad-hoc protocols, AODV and DSR. However, we believe that the trends observed will persist independent of the communication protocol employed due to poor connectivity. Early results with ZRP and GPSR implementations confirm our intuition.

5. DISCUSSION

Our results clearly indicate that the packet delivery ratio for common topology-based ad-hoc routing algorithms varies significantly between an environment using a model of vehicular movement confined to real roads and one using the random waypoint model. Although location-based routing protocols may perform better than AODV and DSR, the lower packet delivery ratio is indicative of poor connectivity in the network due to the increased average distance imposed by constraining movement to streets, and by the increased interference due to node clustering at intersections. Further, we demonstrated that the type of street plan can also impact the packet delivery ratio.

These results contradict those reported in Saha and Johnson [31], where the authors state that a random waypoint model is sufficiently similar to the street mobility in terms of network connectivity. The authors reach this conclusion using a 500 m transmission range and an unspecified path loss model. Under more realistic settings, network performance when nodes are constrained to streets is clearly significantly worse than when nodes move according to the random waypoint model.

Given the increasing interest in vehicular ad-hoc networks, we suggest that simulations for evaluating vehicular communication protocols incorporate motion constrained to roads in representative geographic regions. Also, if such protocols are meant for purposes such as traffic advisory and safety-

Mobility	100 nodes	800 nodes	1600 nodes
RWP	0.6038 s	1.2332 s	2.4658 s
Simple STRAW	2 s	7.4 s	14.6 s
STRAW with OD	15.2 s	96.0 s	99.0 s

Table 1: Runtime for mobility models using different numbers of nodes. Each simulation ran for 960 seconds. Numbers for STRAW were generated using a 5,000,000 square meter region of Chicago.

related applications, we posit that simulations should account for vehicles that change routes according to information contained in the data packets. In this situation, it is essential to use an integrated street mobility model.

5.1 The Cost of Detailed Mobility Models

For such a detailed mobility model to be practical, it must not contribute significant CPU and memory overhead to the wireless simulation. To evaluate STRAW's overhead, we isolate the STRAW component by disabling wireless communication in the simulator and record STRAW's performance in terms of the resulting runtime and memory consumption.

For the runtime performance evaluation, we summarize the results presented in the associated technical report [6] (see Table 1). The experiments were performed on a desktop computer containing a Pentium 4 2.4GHz processor with HyperThreading enabled.[2] The Simple STRAW mobility model incurs a small (approximately constant) factor of runtime overhead compared to the random waypoint model. The STRAW OD model requires a significantly longer execution time, due to the cost of computing shortest paths. It is important to note that runtimes for this mobility model eventually decrease as the number of nodes increase. This occurs because there is significant congestion in the network (i.e., a traffic jam), meaning that each node covers less distance per unit of simulation time and thus requires fewer shortest path searches. For example, with 2400 nodes in the specified region (not shown in the table), the simulation took approximately 91.7 seconds to run, which is faster than when the simulator models 1600 nodes in the same region.

Figure 10 demonstrates how the simulation's memory consumption varies according to the size of the region when using Simple STRAW. As shown in [6], memory consumption for STRAW OD is not significantly greater than for Simple STRAW. Table 2 shows how memory consumption varies with the number of nodes. The figures demonstrate that, although memory consumption can become significant, it is not a limiting factor for simulation with today's hardware. In fact, when loading map data for all of Cook County, IL, which contains the entire city of Chicago (230 square miles containing 157,120 road segments, not shown), memory consumption was approximately 92 MB. Although the size of the data structures supporting STRAW varies during execution, the 92 MB value yields approximately 58 bytes of memory per road segment object, on average.

These results demonstrate that, in general, one can successfully model large-scale realistic vehicular motion on commodity hardware. Although STRAW OD does not scale as

[2]The simulations ran on only one of the two logical processors, so it is difficult to determine how much more than 50% of the CPU's resources were made available to the simulator.

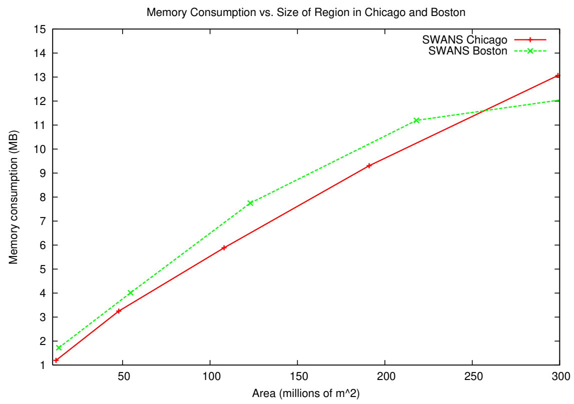

Figure 10: Effect of size of region on memory consumption for simple STRAW using 400 nodes.

Mobility	100 nodes	800 nodes	1600 nodes
RWP	797 KB	675 KB	1.09 MB
Simple STRAW	1.10 MB	1.52 MB	2.16 MB
STRAW with OD	2.03 MB	3.17 MB	4.90 MB

Table 2: Memory consumption for mobility models using different numbers of nodes. Each simulation ran for 960 seconds. Numbers for STRAW were generated using a 5,000,000 square meter region of Chicago.

well as other mobility models, its worst-case performance is bounded by the finite capacity of the underlying road plan. We expect that realistic scenarios using OD pairs will contain many nodes following the same path through a region during a simulation run, enabling STRAW to take advantage of cached routes to reduce the runtime overhead.

6. RELATED WORK

Mobility models have been the focus of significant research for MANETs. Several surveys have focused on mobility models [5] [38] and their effects on routing performance [9]. Due to the popularity of the random waypoint model for evaluating routing protocol performance, many researchers have focused on analyzing its characteristics [7] [29]. Others have attempted to more realistically simulate motion in RWP; in [3], for example, the authors introduce acceleration and deceleration near the waypoints.

Although RWP is commonly used as a general-purpose or worst-case mobility model for evaluating performance in MANETs, we agree with [25] that it neither conforms to any particular realistic mobility scenario nor does it actually capture the worst-case performance for all scenarios. Due to the effect of mobility models on network performance, many researchers have developed alternative mobility models that include features such as group mobility [11] and obstacles [12]. In [4], the authors present a random trip mobility model that integrates a number of models including random waypoint, random walk and city section. An important contribution of their work is the capacity to provide a "perfect simulation" that starts the simulation's mobility in a stationary state.

Other simulations have used mobility traces to evaluate performance in MANETs [14]. Although traces provide perfectly real mobility, they are not generalizable and cannot be used to close the feedback loop in applications such as "traffic advisory," where participating nodes may alter their routes based on traffic conditions.

In the context of vehicular networks, a small number of researchers have accounted for street-constrained motion using real road plans. In [33], the authors use CORSIM to provide a highly accurate model of vehicular movement that has been validated against observed traffic patterns for the target region. Besides the limitation that the CORSIM software is not free or open source, which can inhibit research and development, CORSIM is detached from the wireless network simulator, making it difficult to close the aforementioned feedback loop. Although STRAW has not been validated against observed traffic data, it takes negligible time to configure and uses nonproprietary software and data.

Xu and Barth [34] use the proprietary PARAMICS [28] vehicular traffic simulator to provide node mobility in the NS-2 network simulator. Although this configuration permits closing the feedback loop mentioned above, the solution is limited by NS-2's poor scalability and the overhead of synchronizing the two simulations. This setup also hides the road plan from the wireless simulator, making it more difficult to incorporate street and building information into communication protocols and wireless signal propagation models.

In a closely related work, Saha and Johnson [31] incorporate real map data into the NS-2 network simulator. A limitation of their mobility model, however, is that cars do not interact with one another and there is no notion of traffic control, so each car consistently moves at or near the estimated speed limit. A more realistic model significantly impacts the location distribution and average speeds for each node during a simulation run, which, as we showed, leads to different levels of ad-hoc networking performance. In addition, the authors chose to employ less realistic radio properties (such as a 500 m transmission range) in order to improve network performance.

7. SUMMARY AND FUTURE DIRECTIONS

We introduced STRAW, a new mobility model for vehicular networks in which nodes move according to a simplified vehicular traffic model on roads defined by real map data. We analyzed the implications of mobility models in the performance of ad-hoc wireless routing protocols by contrasting the performance of two well-known protocols using both the commonly employed Random Waypoint Model and STRAW. This study makes the case for a more realistic mobility model integrated with the wireless network simulator, particularly when the information passed through the network affects node mobility.

There are a number of unanswered questions and open issues we would like to address. We are currently examining how the performance of geographic routing protocols changes when evaluated using STRAW. We also plan to investigate more realistic scenarios for our mobility models, such as flows of vehicles with common origins and destinations and vehicles that may change their participation in the simulation (e.g., by arriving at a destination, parking and turning off the car) at any time. Another interesting scenario involves simulating car crashes or other accidents and examining their effects on ad-hoc networking performance.

We intend to refine our simple traffic model to incorporate more realistic elements, such as timed and untimed stoplights, nodes traveling from origins to destinations with varying start times, obstructions and more accurate support for multiple lanes in each direction. To improve the accuracy of simulation, we plan to incorporate real path loss measurements taken in a city environment to augment our simulator's path loss model. Note that none of these changes will affect the main conclusion of this paper; i.e., that simulated routing performance using more realistic vehicular mobility models (like STRAW) is significantly different from that of the popular random waypoint model, and that differences among target regions can have a significant effect on the measured results. The STRAW source code is publicly available for download from *http://aqualab.cs.northwestern.edu/projects/C3.html*.

STRAW was designed in support of the C3 (Car-to-car cooperation) project, the goal of which is to build high-level services following a cooperative model that depends solely on the contribution of participating vehicles. To that end, we are investigating the effectiveness of various data aggregation and directional routing protocols for example applications such as an infrastructureless traffic advisory system.

8. REFERENCES

[1] Barr, R. *An efficient, unifying approach to simulation using virtual machines*. PhD thesis, Cornell University, 2004.

[2] Basagni, S., Chlamtac, I., Syrotiuk, V. R., and Woodward, B. A. A distance routing effect algorithm for mobility (dream). In *Proc. of ACM/IEEE MobiCom* (1998).

[3] Bettstetter, C. Smooth is better than sharp: A random mobility model for simulation of wireless networks. In *Proceedings of the Fourth ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems* (2001), ACM Press.

[4] Boudec, J.-Y. L., and Vojnovic, M. Perfect simulation and stationarity of a class of mobility models. In *IEEE INFOCOM* (2005).

[5] Camp, T., Boleng, J., and Davies, V. A survey of mobility models for ad hoc network research. *Wireless Communications & Mobile Computing (WCMC): Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications 2*, 5 (2002), 483–502.

[6] Choffnes, D., and Bustamante, F. E. Modeling vehicular traffic and mobility for vehicular wireless networks. Tech. Rep. NWU-CS-05-03, Department of Computer Science, Northwestern University, Evanston, IL, July 2005.

[7] Christian Bettstetter, H. H., and Perez-Costa, X. Stochastic properties of the random waypoint mobility model: Epoch length, direction distribution, and cell change rate. In *Proceedings of the Fifth ACM International Workshop on Modeling Analysis and Simulation of Wireless and Mobile Systems* (2002), ACM Press.

[8] Clarke, R. The distribution of deviance and exceeding the speed limit. *Br J Criminol 36*, 2 (1996),

169–181.

[9] F. Bai, N. Sadagopan, A. H. Important: A framework to systematically analyze the impact of mobility on performance of routing protocols for adhoc networks. In *Proc. of IEEE INFOCOM* (2003).

[10] Heidemann, J., Bulusu, N., Elson, J., Intanagonwiwat, C., chan Lan, K., Xu, Y., Ye, W., Estrin, D., and Govindan, R. Effects of detail in wireless network simulation. In *Proc. of SCS Multiconference* (2001).

[11] Hong, X., Gerla, M., Pei, G., and Chiang, C.-C. A group mobility model for ad hoc wireless networks. In *MSWiM '99: Proceedings of the 2nd ACM international workshop on Modeling, analysis and simulation of wireless and mobile systems* (New York, NY, USA, 1999), ACM Press, pp. 53–60.

[12] Jardosh, A., Belding-Royer, E. M., Almeroth, K. C., and Suri, S. Towards realistic mobility models for mobile ad hoc networks. In *Proc. of ACM/IEEE MobiCom* (New York, NY, USA, 2003), ACM Press.

[13] Johnson, D. B., and Maltz, D. A. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, Imielinski and Korth, Eds., vol. 353. Kluwer Academic Publishers, 1996.

[14] Jorjeta G. Jetcheva, Yih-Chun Hu, S. P. A. K. S., and Johnson, D. B. Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture. In *Proceedings of the Fifth IEEE Workshop on Mobile Computing Systems and Applications (WMCSA'03)* (2003), IEEE Computer Society.

[15] Judd, G., and Steenkiste, P. Using emulation to understand and improve wireless networks and applications. In *Proc. of the 2nd Symposium on Networked Systmes Design & Implementation* (May 2005).

[16] Karp, B., and Kung, H. T. Gpsr: greedy perimeter stateless routing for wireless networks. In *Proc. of ACM/IEEE MobiCom* (2000).

[17] Korkmaz, G., Ekici, E., Özgüner, F., and Özgüner, U. Urban multi-hop broadcast protocol for inter-vehicle communication systems. In *Proc. of ACM VANET* (2004).

[18] Kotz, D., Newport, C., Gray, R. S., Liu, J., Yuan, Y., and Elliott, C. Experimental evaluation of wireless simulation assumptions. In *Proc. of ACM MSWIM* (2004).

[19] Lee, W. C. Y. *Mobile Communications Engineering.* McGraw-Hill, 1982.

[20] Li, J., Jannotti, J., Couto, D. S. J. D., Karger, D. R., and Morris, R. A scalable location service for geographic ad hoc routing. In *Proc. of ACM/IEEE MobiCom* (2000).

[21] McCanne, S., and Floyd, S. ns Network Simulator. http://www.isi.edu/nsnam/ns/.

[22] Miller, C. L. UA census 2000 TIGER/Line files technical documentation, April 2002.

[23] Morris, R., Jannotti, J., Kaashoek, F., Li, J., and Decouto, D. Carnet: a scalable ad hoc wireless network system. In *Proc. of ACM SIGOPS European Workshop* (2000).

[24] Nain, D., Petigara, N., and Balakrishnan, H. Integrated routing and storage for messaging applications in mobile ad hoc networks. *Mob. Netw. Appl. 9*, 6 (2004), 595–604.

[25] P. Johansson, T. Larsson, N. H. B. M., and Degermark, M. Scenario-based performance analysis of routing protocols for mobile ad-hoc networks. In *International Conference on Mobile Computing and Networking (MobiCom99)* (1999), pp. 195–206.

[26] Perkins, C. Ad hoc on demand distance vector (aodv) routing, 1997.

[27] Perkins, C., and Bhagwat, P. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *Proc. of ACM SIGCOMM* (1994).

[28] Quadstone, Inc. The PARAMICS Transportation Modeling Suite. *http://www.paramicsonline.com/* 2004.

[29] Resta, G., and Santi, P. An analysis of the node spatial distribution of the random waypoint mobility model for ad hoc networks. In *POMC '02: Proceedings of the second ACM international workshop on Principles of mobile computing* (New York, NY, USA, 2002), ACM Press, pp. 44–50.

[30] Rothery, R. W. Car following models. In *Trac Flow Theory* (1992), Transportation Research Board, Special Report 165.

[31] Saha, A. K., and Johnson, D. B. Modeling mobility for vehicular ad-hoc networks. In *Proc. of ACM VANET* (2004).

[32] Tseng, Y.-C., Ni, S.-Y., Chen, Y.-S., and Sheu, J.-P. The broadcast storm problem in a mobile ad hoc network. *Wirel. Netw. 8*, 2/3 (2002), 153–167.

[33] Wu, H., Fujimoto, R., Guensler, R., and Hunter, M. Mddv: a mobility-centric data dissemination algorithm for vehicular networks. In *Proc. of ACM VANET* (2004).

[34] Xu, H., and Barth, M. A transmission-interval and power-level modulation methodology for optimizing inter-vehicle communications. In *VANET '04: Proceedings of the first ACM workshop on Vehicular ad hoc networks* (New York, NY, USA, 2004), ACM Press, pp. 97–98.

[35] Yin, J., ElBatt, T., Yeung, G., Ryu, B., Habermas, S., Krishnan, H., and Talty, T. Performance evaluation of safety applications over dsrc vehicular ad hoc networks. In *Proc. of ACM VANET* (2004).

[36] Yoon, J., Liu, M., and Noble, B. Sound mobility models. In *Proc. of ACM/IEEE MobiCom* (2003).

[37] Zeng, X., Bagrodia, R., and Gerla, M. Glomosim: a library for parallel simulation of large-scale wireless networks. In *Proc. of PADS* (1998).

[38] Zheng, Q., Hong, X., and Ray, S. Recent advances in mobility modeling for mobile ad hoc network research. In *ACM-SE 42: Proceedings of the 42nd annual Southeast regional conference* (New York, NY, USA, 2004), ACM Press, pp. 70–75.