Technical Report

UCAM-CL-TR-631
ISSN 1476-2986

Number 631

UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

MIRRORS:
An integrated framework for capturing real world behaviour for models of ad hoc networks

Wenjun Hu, Jon Crowcroft

April 2005

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/TechReports/

ISSN 1476-2986

MIRRORS: An integrated framework for capturing real world behaviour for models of ad hoc networks

Wenjun Hu   Jon Crowcroft
University of Cambridge Computer Laboratory
{firstname.lastname}@cl.cam.ac.uk

Abstract

The simulation models used in mobile ad hoc network research have been criticised for lack of realism. While credited with ease of understanding and implementation, they are often based on theoretical models, rather than real world observations. Criticisms have centred on radio propagation *or* mobility models.

In this work, we take an integrated approach to modelling the real world that underlies a mobile ad hoc network. While pointing out the correlations between the space, radio propagation and mobility models, we use mobility as a focal point to propose a new framework, MIRRORS, that captures real world behaviour. We give the formulation of a specific model within the framework and present simulation results that reflect topology properties of the networks synthesised. Compared with the existing models studied, our model better represent real world topology properties and presents a wider spectrum of variation in the metrics examined, due to the model encapsulating more detailed dynamics. While the common approach is to focus on performance evaluation of existing protocols using these models, we discuss protocol design opportunities across layers in view of the simulation results.

1 Introduction

Mobile ad hoc networking has been an active research area over the past decade, producing a multitude of protocols across the stack. Yet there remains much not well known of the network dynamics. This inspires us to carefully study the network characteristics to gain insight for protocol design. Due to the nature of the wireless medium and mobility, however, mobile ad hoc networks are highly dynamic and complex. Modelling the real world underpinning these networks is therefore difficult.

Models for ad hoc networks encompass views of the real world, including space configuration, radio frequency propagation and mobility, as well as network operations, such as communication patterns. We focus on the real world views. Despite much previous effort across the above aspects, the current state of art is still unsatisfactory. Most notably, criticisms centre on the lack of realism. While earlier models are credited with ease of understanding and implementation, they are often based on theoretical models rather than real world observations. More recent works have addressed some criticisms, and yet there remains much scope to improve the level of details, to more accurately and precisely reflect the underlying network.

Moreover, our work is encouraged by the importance of the models in simulations. Performance evaluations of ad hoc network protocols generally rely on simulations. These require either real world traces or appropriate models. The former is difficult to obtain, and often describe a specific setting, from which it is hard to derive general network characteristics. For the latter, the accuracy of the models are critical to the credibility of simulation results. It has been shown that the models used could significantly affect protocol performance [2, 6, 14]. Therefore, an open question remains as to what effects those missing details might have. Our work considers these details.

The notion of space is embedded in most discussions on RF propagation or mobility, but we explicitly separate it out to emphasise its importance. Furthermore, criticisms and works have so far focused on individual propagation *or* mobility models, whereas the interplay between these models deserves much consideration. In this work, we propose *a new integrated framework for capturing real world behaviour* across different views and scenarios in the real world, MIRRORS: Mobility Integration of Radio Requirements in Real-world Simulations.

Although main discussions are in the context of simulations, we aim to model mobile ad hoc networks from a high level to deepen understanding of these networks. Therefore, instead of taking the common approach of focusing on evaluating existing protocols using these models, we use mobility as a focal point to study the topology characteristics of the networks synthesised. Our other contributions include: i) a new metric, neighbour occurrence count, to analyse the extent of intermittent connectivity in networks; ii) detailed distributions of the metrics studied; and iii) protocol design considerations highlighted from detailed modelling. An important observation is that these concerns could span several layers in the protocol stack. This could be of even greater interest and importance to cross-layer designs.

In the following sections, we first discuss existing approaches to modelling ad hoc networks in Section 2, and then introduce our own framework, MIRRORS, in Section 3. Section 4 studies homogeneous base cases within the framework in detail to illustrate the realism of our approach, followed by heterogeneous scenarios composed from the simple base cases. In Section 6, we present comparative simulation results derived from our approach and existing models. Section 7 discusses the use of the models to drive protocol design, and finally Section 8 concludes the paper.

2 Related work

2.1 The *de facto* standard

The set of models and parameters commonly used for performance analysis, especially of routing protocols, was originally proposed in [5]. The area is a flat unobstructed space of 1 km × 1 km. Radio propagation follows a two-ray ground reflection model, with a transmission range of 250 m, or 376 m for GloMoSim based simulations, over a perfect disc. Movement pattern is described by the Random Waypoint Model. The network size is often on the order of 50 nodes, and a typical simulation run lasts 900 s.

These settings certainly served as a good starting point for simulation based performance evaluations, though they do not match all scenarios. Unfortunately, they appear to have constrained later evaluations over years, which is probably unintended by the original authors.

2.2 RF propagation

Works on propagation models have largely involved link level measurements from testbed experiments, mainly as reported by Uppsala's APE [20], Dartmouth's experiments [10, 15], MIT Roofnet [1] and MSR's experiments [9]. Wide-ranging issues from connectivity to routing metrics reflecting link quality have been discussed, all highlighting the need for futher real world experiments. However, in contrast to the common alleged MANET applications in battlefields, disaster relief and conferences, most testbeds have been indoors and stationary. Roofnet is outdoors, but very high above the ground level. Dartmouth researchers recently conducted experiments of an outdoor mobile ad hoc network, but simulating Random Waypoint movements [10]. It is also a pity that none of the works above has proposed empirical models matching the measurements.

2.3 Mobility models and metrics

As an integral part of the simulation model for mobile ad hoc networks, mobility models have been under extensive studies. Other than the Random Waypoint Model, several earlier random walk variants were described in [6], such as the Random Direction Model, where nodes randomly choose an initial direction for movements, as opposed to a random destination in Random Waypoint. Also described in the survey, Gauss-Markov Model correlates successive movements, although it is based on cell networks. Smooth Mobility Model [3] includes detailed speed synthesis, especially for vehicles with high mobility, starting with a set of preferred speeds for each node. However, it is an improvement of Random Direction. The unconstraint environment is unlikely for the range of speeds, hence the types of nodes, studied. Mobility Vector Model [11] focuses on natural speed dynamics, but does not address the overall trajectories.

More recent efforts have involved restricted movement spaces. The City Section Model in [6] considers a Manhattan grid like street network, but the formulation only dealt with predefined paths and traffic laws. The Freeway and Manhattan Models were proposed in [2]. In both models, nodes were constrained within lanes, and by the nodes ahead in the same lane. An additional consideration for the Manhattan Model was the turning behaviour at a junction, where a probabilitic choice is taken. This practice of applying the same probabilities at all junctions is debatable. Also in [2], a number of metrics were proposed to evaluate the mobility models, such as the degrees of spatial and temporal dependence. Their definitions 'for nodes *not too far apart* in space/time' were vague, however. The effects of obstacles were examined in [14], where a Voronoi space graph was used to generate paths from the obstacles. We will argue in the next section that obstacles and paths should not be coupled. In all three works, mobility models were shown to affect the relative ranking of the protocol performance.

In [16], the authors investigated the effects of destination selections on node densities and success rate of route discoveries. It should be noted that the nodes were restricted to students on campus, who do not roam continuously. Furthermore, the destinations have significant dimensions, when compared to waypoints, and the route discovery success is related to co-locations, one element of context.

The survey [6] further describes several group mobility models, which deal with movement dynamics from different angles. It was pointed out in [17] that group formation requires due consideration. The social aspects thus implied may be highly scenario dependent, but we demonstrate in Section 5 that our framework allows for these situations.

A number of works have considered mobility models from a statistical perspective, studying their stochastic properties [3,4], pointing out the problems with the original Random Waypoint Model [21] and suggesting methodologies to ensure stationarity [22].

2.4 Overview

Where RF propagation is concerned, most models focus on path losses, and only [14] has investigated signal obstruction. Even so, many other aspects, such as multi-path fading, have not been discussed but could make a difference, as shown in [1,19].

In the case of mobility, constraints have been addressed in various models, but the fundamental movement mechanism is unclear, and the speed synthesis could be improved. Furthermore, none of the models incorporates all three streams of considerations.

It is worth noting that the geography is usually embedded in either RF propagation or mobility, as seen in the criticisms regarding flat free space. However, other aspects of the space geometry are left untouched but maybe significant.

Although mainly a simulation issue, the duration of a run is often limited to 900 s, except for [14]. The usual reason appears to be that the simulation will have converged to a steady state by then. However, it is questionable whether the full spectrum of behaviour is exhibited and whether stationarity is the only necessary concern. We point out in Section 4.2 that minimum pauses could be longer than 900 s.

Most notably, all the improved models address one issue or two at a time, regarding RF propagation *or* mobility. Naturally, questions arise as to correlations between these issues, whether these enhanced models are comparable to one another, and if so, how.

3 The MIRRORS framework

Our first observation is that scenarios vary considerably, and therefore one single model is unlikely to suit all situations. In general, the space constrains RF propagation and mobility. Typically, the mobile nodes in question fall into a few categories: pedestrian, city vehicular and highway vehicular. Within each category, node movements exhibit similar patterns. Complex scenarios, on the other hand, can be decomposed into simple units. Same considerations apply across categories, in the form of same parameters albeit different values, e.g., pedestrians and cars alike may be under constraints from the geography. In some cases, however, some of the issues may simplify to varying degrees, e.g., accelerations are not of great concern to pedestrians. We therefore propose an integrated framework to address common concerns of space, RF propagation and mobility—MIRRORS: Mobility Integration of Radio Requirements Of Real-world Simulations.

3.1 Outline

The MIRRORS framework consists of a set of base cases, from which complex scenarios can be derived through composition. In each base case, nodes have homogeneous mobility capabilities and identical probabilistic distributions. Space is considered as a standalone component, while RF propagation and mobility models are other components of a base case. For each component, we specify parameters, the typical values of which are discussed in the context of representative base cases corresponding to the above categories. Effectively we group all parameters into different sets for space, RF propagation and mobility respectively.

3.2 Base case—space

The space model specifies 3D geometries of the obstacles and paths, e.g., in the form of coordinates of obstacle vertices and control points for paths, and an appropriate projection onto a 2D area is often acceptable. The paths and obstacles are probably closely related, but not necessarily mutually exhaustive, considering, for example, lawns which neither favour movements nor block much signal transmission. This is a generalisation of the space graph in [14].

3.3 Base case—RF propagation

Given the common assumption of omni-directional antennas, propagation can usually be characterised by a microcell environment. The main issues are path loss, fading and shadowing models. According to [18], appropriate empirical models can be used to simulate real world scenarios to a reasonable accuracy, and this was confirmed in [10]. For line-of-sight (LOS) paths, assumed on a flat surface, path loss can usually be approximated using the two-way ground reflection model and Ricean Fading is suitable. For non-line-of-sight (NLOS) paths, we consider that signals will be blocked by buildings, but diffracted around the vertical edges of buildings and potentially over rooftops. Fast fading is more likely to conform to Rayleigh statistics. The signal power from the LOS path with respect to the power from NLOS paths can be controlled by the Ricean K Factor [19]. Shadowing is useful to model varying signal strength due to, e.g., leaves.

In most cases, it requires applying well known empirical models, and our emphasis is on considering all applicable components. RF propagation is highly complex, and therefore it is advisable to note both the merits and limits of simulation approaches.

3.4 Base case—mobility

For ad hoc networks, we model human movement behaviour, unlike for sensor networks for environment monitoring. Mobility is continuous, in contrast to mobile IP, cell network or WLAN situations, where mobility is discrete. Furthermore, we are concerned with microscopic behaviour, since movement dynamics is down to individual nodes. Therefore, we focus on entity mobility models. Compared with space and propagation models, mobility models are less well defined, and hence this is our focal point.

A number of general observations can be made. Nodes normally follow targets, and movements are under constraints. People prefer to save time, by travelling at a speed that is close to the possible maximum while ensuring a comfortable state for themselves.

The building blocks for a movement trace include movements and journeys. We define each movement as a period of motion at a constant velocity, and each journey as a sequence of movements from the last destination up to the current target. Pauses might be possible between journeys. Concatenating journeys yields the overall trajectory, while instantaneous velocity is the rate of change along the trajectory, as in standard physical definitions.

It is the trajectory that dictates a node's trace asymptotically, while the detailed speed variation determines the precise point along the trace at a particular instant. To that end, trajectories and speed variations can be derived separately.

3.4.1 Parameters

Each node is confined in a movement space, which may be an open campus for pedestrians, city streets or freeways for cars. In practice all nodes in the network are likely to be subject to the same movement space, and hence this could be a parameter across the whole simulation.

A node has a list of preferred destinations and other destinations, which together specify the destination distribution for a journey. In general, this distribution is dependent on the node's current location and time, and can be described by a Markov Chain.

Other parameters include characteristic speed(s), preferred/steady state speeds, associated speed drifts, speed limits, acceleration limits, altogether describing the speed dynamics. Compared with the Smooth Mobility Model, we assign only one preferred speed to each node, and leave other speeds synthesised through constraints.

Not all parameters are of concern in each scenario, e.g., in a pedestrian network it suffices to consider speeds without regard to accelerations. Furthermore, a pause time distribution is associated with each location and time. It can be short to represent roaming, or long to reflect the more common behaviour of travelling to a place and staying for a long time.

3.4.2 Algorithm

To start a new journey, a node selects a new destination according to the destination distribution at its current position and time. A movement algorithm then describes how the node travels to the destination. Along the trace, the velocity of the node is adjusted according to spatial, temporal and physical constraints and speed drifts, as well as any issue particular to the scenario. Node positions are updated accordingly. On reaching the destination, the node possibly pauses. Then the whole process is repeated.

3.5 Dependencies between models

From the viewpoint of mobility, the characteristic speed identifies the category of movement patterns, and hence the typical space. A key observation is that, for a particular scenario, the space is the underlying substrate for the network. This is in contrast to the wireline networking paradigm where space can often be abstracted away in the presence of wires, and it may be appropriate to abstract them further by capturing link metrics. Generally, the paths constrain line-of-sight RF propagation and mobility, whereas obstacles affect signal obstruction and diffraction. The introduction of an obstacle would therefore affect both the movement freedom and the signal propagation, as observed in [14]. Furthermore, approximations in RF propagation calculations may depend on the mobility scenario, especially when concerning fading and shadowing situations. On the other hand, effects of propagation models such as path loss calculations could be amplified by the mobility model.

Therefore, a scenario involves not only the mobility and the movement space, but implicitly the propagation. Studying the effects of a single component in isolating, be it space, mobility or RF propagation, may only reveal a partial picture. A consistent approach is necessary in modelling, starting with recognising the characteristic speed, as the case studies suggest.

3.6 Formulating an integrated model

To formulate an integrated model from MIRRORS for a scenario, we usually start by identifying a characteristic speed, which implies a typical terrain. The paths and obstacles in the space dictate RF propagation and movements, which are then considered respectively according to previous subsection. The parameters have been identified in the discussion in previous sections, and realistic values can be obtained from a modest amount of real world monitoring.

It is also to be noted that the framework covers deterministic models, which can be derived by setting the probability of e.g., selecting a particular destination to 1 and 0 for other destinations.

3.7 Aside: Simulation issues

A number of statistical issues are of concern to simulations. For example, statistical artifacts due to small area sizes have not received sufficient attention. This is one result of the common approach of modelling a closed system. Considering large systems, most are in fact self-contained, so a closed system can be a reasonable representation. To impose 'closedness' on a small area necessitates border rules, and previous studies have shown that their impacts are significant [3]. On the other hand, a small area limits the extent of dispersion, and therefore may not present the full spectrum of connectivity paradigms, from persistent, intermittent to transient, corresponding to a node in contact with the same neighbour continuously over a long time, or encountering the neighbour recurrently or occasionally. Also shown in various studies, the spatial distribution of nodes is normally non-uniform.

In order to reflect these, the configuration of the simulation area needs to comply with realistic vision. To model a metropolitan area, for example, the scale should be at least several kilometers on each side. In certain cases, we should allow nodes to depart and arrive temporarily, i.e., allowing the total number to vary but within a tolerance. Depending on the area size, the simulation duration should be sufficiently long to ensure convergence both speed- and destination-wise. The border behaviour, initialisation and position update procedures will also vary with the scenario. In summary, the movement scenario should determine the simulation scale.

4 Case studies

As mentioned earlier, mobile nodes are normally pedestrian, city vehicular or highway vehicular. Together with the case of stationary nodes, they form categories of ad hoc networks. We now discuss representative patterns, which implicitly emcompass space and RF propagation considerations. Through the description of the cases, we also demonstrate how to formulate an integrated model from MIRRORS.

4.1 Stationary network

We start with the characteristic speed, 0 in this case. Space-wise, this is usually a network of indoor devices. It may be generalised to cover fixed base stations on a wide area, although a standalone network of those is unlikely. RF propagation in an indoor environment is complicated and under extensive studies, but the absence of mobility reduces the network dynamicity considerably. A stationary network may be more concerned with mesh connectivity and capacity, but within this framework it still serves as a base case.

4.2 Pedestrian network

The characteristic speed is the maximum speed, 2 m/s. Typical pedestrian networks are in campus environments or metropolitan areas. In either case, the movement space is 2D with respect to the paths, and obstacles in the paths are 'points', mainly building and trees. This means that nodes are generally not under constraints from peers, unless the entire area is saturated. RF propagation involves path losses along the paths, obstruction by obstacles, diffractions around vertical edges of obstacles, fading and shadowing concerns.

In terms of mobility, the key is destination distribution, which is non-uniform across time and space. Different groups of people would form different cases. Consider students on campus, they tend to aim for residence halls, departments, libraries, canteens and so on, and the preferences vary according to time of the day. Professors would have different preferences, but similar considerations apply. This generalises [16], which is based on students' movements. Most likely the person would travel at a preferred speed he or she is comfortable with, and there is little variation beyond that. The pause time is associated with the purpose of the journey, such as meals in canteens and lectures in departments, hence the location, and is likely to be long, e.g., at least 20 minutes. Therefore these nodes do not roam. This can be contrasted with pedestrians in a shopping mall, where they tend to pause only briefly and roam for a considerable amount of time. Other issues or constraints, such as accelerations, are negligible.

4.3 City vehicular network

In this case the characteristic speed is the city-wide speed limit. The movement space consists of lanes in streets, and therefore movements are linear with respect to the paths. For RF propagation concerns, path losses are similar to those in the pedestrian cases, but fading/shadowing effects can often be ignored due to the high mobility.

Representative mobile nodes in these scenarios are buses, taxis and other traffic. Buses roam, but have fixed trajectories. Their destinations are always the next stops along the routes, where they pause very briefly or occasionally not at all, so the movements are mostly deterministic. Taxis also roam, but the traces are more random and pervasive. They tend to favour places like bus/train stations or other busy areas in town, but may also travel to other destinations to collect or drop off a passenger. The pause times tend to be short and above a non-zero minimum, but special considerations might apply at taxi stands, e.g., if the taxi joins a queue and waits for its turn to take the next passenger. Other vehicles often have transient journeys, potentially pausing for a significant length of time after a journey, and make fewer discrete journeys over a long period of time, on the order of days.

Speed-wise, in any of the three cases the driver usually prefers a speed close to the speed limit, and the actual speed drifts slightly around the preferred value under free flow conditions. The constraints include spatial—safety distance to vehicles ahead in the same lane and traffic lights, temporal—velocity correlations between successive movements in a journey, and physical—acceleration bounds in general and speed bounds for turning action and so on. Given the high mobility setting, accelerations are significant. Road policies are examples of additional concerns for this scenario, e.g., considering one-way streets, which could affect the initial movement in a new journey and turning actions.

4.4 Highway vehicular network

This setting is characterised by a higher speed limit than in built-up areas. The space is usually open and uncluttered, which simplifies both radio and mobility concerns. Movements are essentially continuous along the lanes, and the same speed variations as above apply. The main subtlety is in switching lanes, for which many driver behaviour models have been proposed for transport studies.

5 Heterogeneous scenarios

All the representative patterns above depict very homogeneous networks, where nodes have the same mobility capabilities and identical probablistic distributions. They can serve as the building blocks for heterogenous situations.

5.1 Heterogeneous space

An interesting example of space composition is to combine indoor and outdoor environments over wide area, e.g., consider the vicinity of a campus building as well as the inside. More generally, non-uniform space layout could be derived from basic area units. Essentially, 'contituent states', positions or signal strengths, can be calculated within the respective space components, and care is needed for dealing with behaviour along component boundaries.

5.2 Heterogeneous radio models

Since the basic propagation models already embrace a number of issues, composition is not as distinct. Some are implied in space composition examples, e.g., indoor/outdoor environments would imply different path loss exponents, and therefore they should be applied as appropriate. It may however be possible to accommodate single-radio nodes with variable transmission ranges or multi-radio cases, by treating them as compositions of single radios with fixed ranges.

5.3 Heterogeneous mobility

If a group of nodes follow the same destination selection pattern, then group mobility could be observed. This further suggests that the choice of destinations could embody social aspects in the network. The notion of co-location in [17] is reflected in neighbouring relations, i.e., nodes within direct connectivity are likely to be close in location. Given the vast number of possibilities in resulting in group mobility, it can be more scenario dependent than entity cases.

Heterogeneous networks can be derived if the requirements of homogeneous node capabilities and idential probabilistic patterns are relaxed. One example is non-identical patterns within the same types of nodes, e.g., considering arts students and science students, whose destination distributions could be the same, except they involve different departments. Within the same category of mobility capability, we observe city traffic networks with buses, taxis and normal traffic, which have different behaviour. Crossing capability boundaries, we could derive the cases of mostly stationary networks with occasional mobility, as in many testbed experiments, and metropolitan area networks with both pedestrians and vehicles. Higher order compositions of composed models could generate even more complex patterns.

Despite the increasing degree of heterogeneity, the mobility states of nodes, such as destinations and current positions, can be obtained from respective mobility units, with additional adjustments as a result of interactions between different nodes. Where different movement space is involved, RF propagation calculations need potential corrections accordingly. The viability of such heterogeneous networks for the purposes of network services is beyond the scope of this discussion, however.

6 Comparative analysis

Having presented our formulation of the models, we now turn to effects of the details in these models on network connectivity, especially when compared to existing models. Our models are considerably more complex than existing approaches, in the hope of capturing fine dynamics. We focus on the city taxi scenario on a Manhattan grid for its complexity, and compare the resulting connectivity graphs with those derived from the Random Waypoint Model (RWP) and a variant of the Manhattan Model (MV). We note that these are not mobility models alone, but imply the terrain and radio propagation models. In fact, they represent typical combinations of those components. Given the models are not necessarily comparable, parameters are set to comparable values as far as possible. Moreover, we take a protocol independent approach to study the network characteristics.

6.1 Detailed models

Our Taxi Model is as outlined in Section 4.3 and detailed in [12], except for the destination distribution. We consider taxi movements under free flow conditions except when potentially subject to traffic lights at junctions. These lights cause significant speed variations over the entire simulation. Taxis travel to one of the three hotspot destinations at (0, 0), (2000, 2500) and (3000, 2000) 2/3 of times, roam to the central 36% of the area 1/6 of the times, and to the entire area in the remaining times. All probabilities are uniformly shared between the possible destinations. Time-varying aspects are ignored for now, given our focus on a short time with respect to the daily cycle. The movement algorithm is based on following shortest paths.

The original Manhattan Model was not precisely specified in [2] regarding speed variation and calculations of signal strength. Therefore, the speeds are synthesised as in the Taxi Model, with the exceptions of traffic light considerations, and signal strength is calculated in the same way in both simulations. Furthermore, [2] did not specify pause behaviour for the Manhattan Model, and therefore we set pause times to zero in all models for this comparative study. Standard RWP is used, with a non-zero minimum speed.

For the Taxi Model, the propagation mechanisms include path losses along lanes, following a two-ray ground reflection model [18], and diffraction around street corners [12]. The transmit power strength and receiver sensitivity are set to be 20 dBm, the legal limit, and -80 dBm respectively. The breakpoint distance is 100 m and the antenna height is 1.5 m. Attenuation of diffracted signals around each corner is estimated to be a further 20 dB loss. These together amount to a transmission range of 474 m along a straight line and 150 m if one corner is taken. The former is also used for the RWP calculations.

Since common simulators such as ns-2 and GloMoSim do not provide all the models necessary for the taxi scenario, we implemented the model as a statistical simulation to calculate positions and derive connectivity at 1-second granularity. The other two models were also implemented within the same system for compatibility. The traces can be formatted to be fed into ns-2 using the *setdest* tool.

500 nodes are distributed over an area of 5 km by 5 km, initially according to the destination distribution for the Taxi Model, and uniformly for RWP and MV. Additionally for MV and the Taxi Model, this area has a Manhattan grid layout with a uniform block size of 100 m by 100 m. These parameters are intended to minimise non-uniformity caused by the area, but mimicking the scale of a small city. There are two lanes in each street, one for each direction. This means there is a bypass enclosing the area. The minimum and maximum speeds for RWP are 1 m/s and 15 m/s, the latter of which is also the speed limit for the other models. For MV and the Taxi Model, nodes are assigned one of three preferred speeds 13, 14 and 15 m/s. The acceleration is set to be 3 m/s^2.

Each simulation run lasted 1 hour in simulation time, after an initial 1000 s warmup period. We first

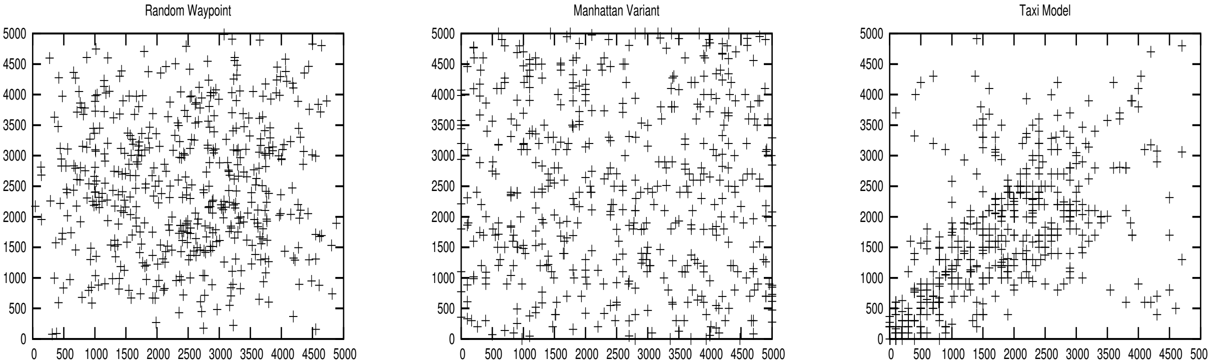

Figure 1: Spatial distribution of nodes: RWP, MV and Taxi Model. Axes indicate coordinates.

examine network wide snapshots of node distribution, connectivity characteristics and relative speed distribution at chosen times. Time-varying characteristics could be exhibited by comparing simulation results over models for different times. We then take a node's eye view to study neighbouring relations for those within direct connection. The results presented reflect individual runs, but several runs were executed for each case to confirm common behaviour. Our emphasis is on the overall trend.

6.2 Spatial distribution of nodes

Prior works have examined the spatial distribution of nodes for various mobility models. We present in Figure 1 observations at the end of each simulation.

As reported in [4], the long run node distribution of RWP is non-uniform, and at it most dense at the area centre. Our results confirm this even when RWP is run on a Manhattan grid. MV, on the other hand, results in rather uniform distribution of nodes. This is no surprise, when contrasted with the observed node distribution for the Random Direction model [3], since MV is essentially a variant of the latter. For the Taxi Model, nodes tend to aggregate around the hotspot destinations and the region enclosed by them, although the exact distribution varies according to the hotspot locations. These can be compared with findings in [16], which reported non-uniform spatial distribution of nodes observed from real world traces.

Furthermore, we have observed that the non-uniformity property is independent of the speeds, although all the speed considerations determine the exact distribution at a given instant. Even when pauses and different speed ranges were taken into account, the general trends of node distribution remained the same for RWP, as reported in [4], and the Taxi Model. This supports the separate formulation of trajectories and speed variations.

We therefore conclude that the destination selection mechanism, as opposed to the selection of an initial direction, gives rise to non-uniform spatial distribution of nodes asymptotically. As a key component of a realistic mobility model, the destination distribution varies with the scenario. This implies that, inherently, there are both highly congested and blank-out regions. It is also a case for a realistic scale for simulation, in order to study the full spectrum of variation.

If non-uniform street layout were to be considered, its effects would depend on its relative position within the entire simulation area. Essentially it creates non-uniform street capacity, thus affecting the overall connectivity more at the favoured regions than elsewhere.

6.3 Instantaneous connectivity

In graph theoretic terms, the nodes may form one or more connected components at any instant, based on potentially multi-hop communications. The number and sizes of network partitions are closely related to the spatial distribution of nodes. We performed a breadth-first search on nodes to identify reachability, and hence network partitions, at randomly chosen times. Figure 2 illustrates the frequency distributions of partition sizes at the end of simulation runs.

As expected, the plots match those of spatial node distributions. In RWP, nodes cluster around the centre, and hence there is a large central connected network component. This clustering is split around multiple spots in the Taxi Model, hence the existence of several major network partitions, but of smaller sizes. In MV, however, nodes are more evenly spread out, resulting in many more and small partitions.

Along the other end of the scale, the numbers of partitions of a node or two indicate the extent of disconnection. The value of this peak is related to the relative location and dimension of the sparse region within the entire area. For RWP, it is around the area borders and of small dimensions, while for the Taxi Model, it is the large area surrounding the hotspot region. Sparse regions are more difficult to identify in MV, as the spread of nodes is more erratic.

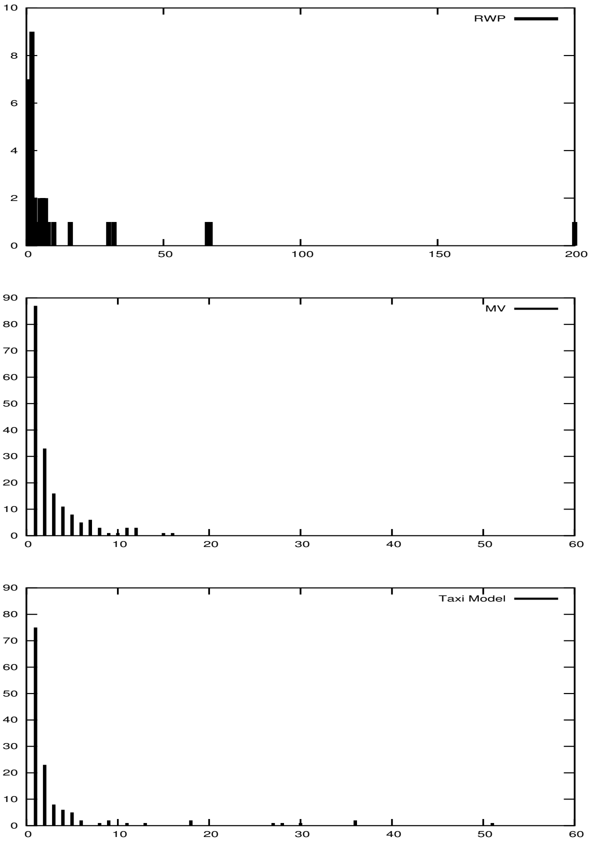

Figure 2: Frequency distribution of network partition sizes: RWP, MV and Taxi Model. Horizontal axes indicate the partition sizes, and on the vertical axes are observed frequency counts of individual sizes.

6.4 Relative speed distribution

Relative speeds and their average over all node pairs were metrics within the framework proposed in [2]. Instead of taking the average, we examine the distribution. Moreover, we focus on node pairs within direct connections. Characteristics of these speeds implicitly affect the durations of the connections, hence connectivity over the whole network.

Snapshots of relative speeds were taken at random points of a simulation run, and the representative distributions, taken again at the end of simulation runs, shown in Figure 3.

This is an example where formulations of space, signal propagation and mobility all play a part. Generally, the determining factors are speed range, path orientation and signal propagation, and the results are independent of the exactly block size. RWP completely fails to reflect relative mobility for inappropriate representation of these factors. Its speed variation uniformly spans a large range, unlike the discrete ranges in real life. Even when this range is reduced to 12 to 15 m/s as in the second plot, it still allows movements and signal propagation in all directions. MV presents a more realistic view, firstly for its path layout, as well as other formulation borrowed from the Taxi Model, and the peaks correspond to angles of 0, ±90 and 180 degrees between velocities.

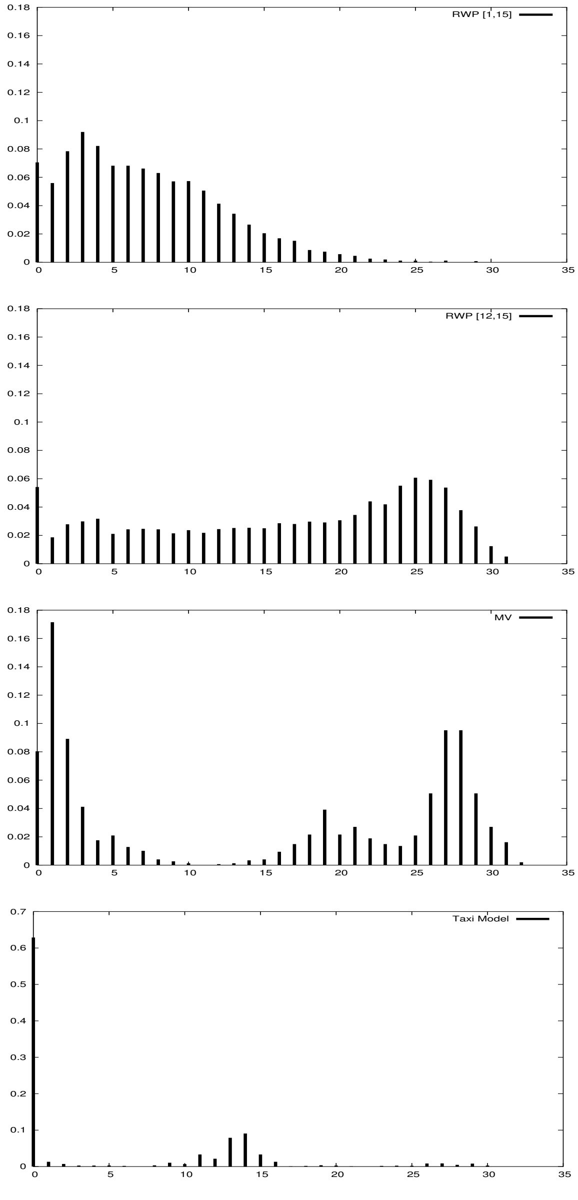

Figure 3: Probability distribution of relative speeds: RWP with speed range [1,15], RWP with speed range [12,15], MV and Taxi Model. Horizontal axes are relative speeds in m/s, and vertical axes indicate probabilities.

An interesting observation is that traffic light considerations lead to nodes stopping in the Taxi Model, hence contributions to the lower speed peaks from larger inter-velocity angles. Results show that pauses make further pronouced contributions to this shift, by introducing larger speed variations and higher probabilities of low speeds. Therefore, detailed speed synthesis would also affect relative mobility, with implications on link durations.

For the Taxi Model, a related observation is cycles in the proportions of relative speeds between nodes along the same, opposite and orthogonal directions. The percentages of same/opposite-direction nodes in connection remain mostly comparable, while the proportions of orthogonal-direction node pairs increase and decrease continuously from around 16% to 33% then back to 16% or lower, and the 'period' is usually around 60 seconds, corresponding to the duration of red lights.

6.5 Link durations

As the relative speed affects the lifetime of a link, we next investigate link durations. These links are between nodes directly within the transmission ranges. Statistics were run of link durations over all node pairs throughout the entire simulation, and the frequency distribution of different durations shown in Figure 4. Despite the randomness in the exact frequency counts, the same distribution shapes have been observed for all links involving the same sample node. In order to emphasise the absolute numbers, the plots were not normalised.

The shapes can be explained with reference to the distributions of relative speeds. The peaks and humps in Figure 4 can be accounted for by those in corresponding plots of Figure 3. Naturally, the higher the relative speed, the shorter the link duration.

Again, the distribution shapes for MV and the Taxi Model are more similar to each other than to RWP. Due to more random directions at junctions for MV, more short links are observed. In the meantime, the randomisation effectively smoothes out differences, and the absolute numbers of short links are much smaller than for the Taxi Model.

6.6 Inter-contact durations

Inter-contact durations are closely related to link durations. An inter-contact interval is defined as the duration between two consecutive direct connection between the same pair of nodes. We collect statistics over complete simulation runs, and plot the normalised frequency distributions of these inter-contact durations in Figure 5.

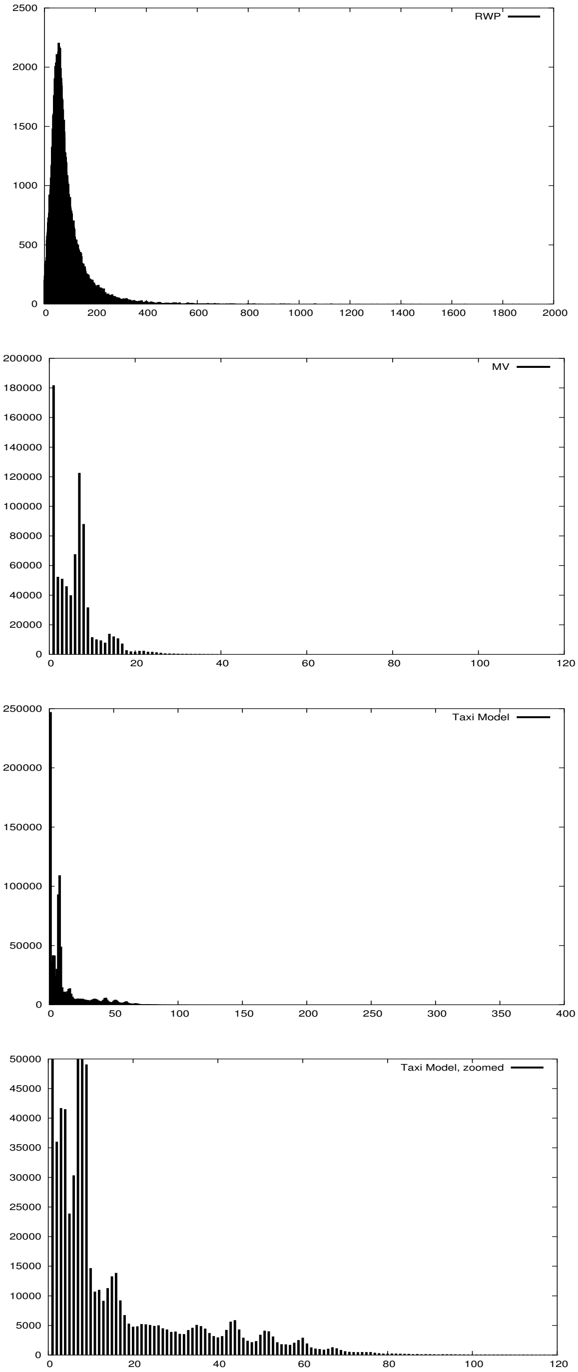

Figure 4: Distribution of link durations: RWP, MV, Taxi Model and Taxi Model zoomed in. Horizontal axes indicate link duration in seconds, and vertical axes report frequency counts.

These can be compared to [7] where power-law distributions for inter-contact durations were observed from real world mobility traces. RWP fails to represent the power-law characteristics. Initial results suggest that the speed variation has a notable effect on the distribution for short inter-contact times, which also accounts for the agreement between the

MV plot and that for the Taxi Model. On the other hand, the destination distribution affects the distribution of long inter-contact durations, i.e., those of 1000 s or more, and this can be seen from the similarities between the RWP and the Taxi Model plots. We are investigating this further.

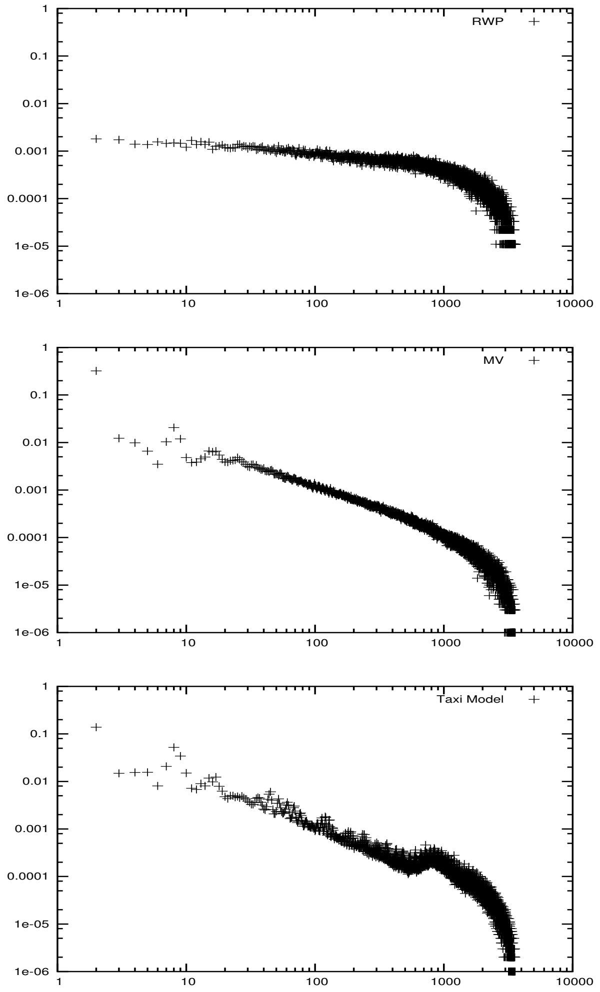

Figure 5: Distribution of inter-contact durations: RWP, MV and Taxi Model. Horizontal axes represent the lengths of the durations, and vertical axes report probabilities.

6.7 Neighbour statistics

When switching to a nodal view for detailed analysis of links, each node appears to encounter most other nodes at least once, and often repeatedly, but the neighbouring durations and intervals fluctuate considerably. We define a neighbour occurrence as the event of a continuous period of time within which a direct connection is possible with a particular node. Frequent occurrences could also imply shorter link durations during each occurrence. Statistics were run over the occurrences of all neighbours for the same node. We select nodes at random, and plot the frequency distribution of these occurrences over the entire simulations (Figure 6).

The shapes of the plots are attributable to destination preferences, relative speeds and hence link durations. Compared with RWP, the Taxi Model gives significant preferences to the same small number of hotspot destinations, but produces higher peaks at large relative speeds. Compared with MV, our Model again creates more opportunity for repeated encounters.

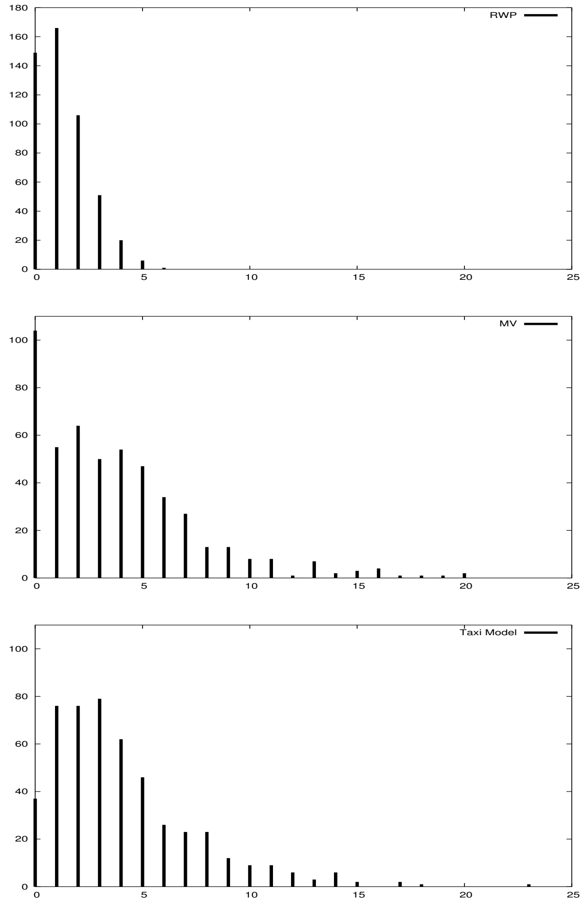

Figure 6: Distribution of neighbour occurrence frequencies: RWP, MV and Taxi Model. Horizontal axes represent the number of occurrences, and vertical axes report frequency counts.

6.8 Model complexity

Compared with RWP, the Taxi model is significantly more complex. Space-wise, our model considers a Manhattan grid style street layout that places restrictions on both movements and signal propagation. The propagation component in the Taxi

Model takes into account defraction in addition to path losses, while the mobility component synthesises more detailed destination selection and speed variation. MV shares much with the Taxi Model, but the initial direction is determined with a simpler probablistic choice.

Overall, the Taxi Model as a representative of the MIRRORS framework involves a larger set of variables than existing models, and it is these parameters that capture fine details of the dynamics. We are currently investigating the effects of individual variables.

6.9 Discussion

Our analysis do not capture effects at the physical, MAC and network layers, as we hope to understand fundamentals of the networks independent of protocols. Furthermore, this set of protocol independent observations suggest the optimal performance of the network, hence potentially guiding protocol design.

The node number was low enough to ensure network partitioning, hence the issues of accurately representing link durations and the amount of intermittent connectivity to handle. Path durations were not calculated, since most paths are very short-lived. If the network is hardly partitioned, local node densities would become the major concern as discussed earlier. Even at a high density, intermittent connectivity may still be observed, in the form of frequent neighbour changes, or link changes. Furthermore, our studies considered a 2D projection of the space for movements, but the results can be generalised to 3D. Signal strengths, in particular, need to be calculated in 3D, as reflected in the antenna heights for taxis.

As with neighbour statistics, some network states recur over time, but aperiodically. This also casts doubt on the stationarity requirement in simulations, as noted in [16].

It should also be pointed out that we have simplified the study to binary state connectivity, given most metrics concern instantaneous snapshots or short link durations. When considering physical layer issues, exact signal strengths need to be obtained, which highlights the importance of spatial dimensions.

On the whole, models such as RWP and MV both overestimate and underestimate some part of the variations, so the combined effects could be complex. While for the MIRRORS framework, even though the specific model might vary with the exact parameter values, the general observations about topology characteristics hold for the type of networks investigated.

7 Applications of the analysis

With insight from the above topology analysis, we could design protocols more prescriptively. Although stationary and mobile networks aim at different goals, i.e., increasing capacity vs coping with dynamicity, issues such as node density apply to both.

The spatial distribution of nodes relates to contention or lack of connectivity. Even for stationary networks, e.g., in office environments, node distribution is likely to be non-uniform. In dense areas, one should expect higher levels of interference and contention at the physical and MAC layers, and congestion at the network and transport layers. Most prominently, the protocols concerned span several layers! However, most comparative performance analysis of routing protocols, possibly over different mobility models, implicitly ignore effects at lower layers. Meanwhile, direct communications could favour nearest or nearby hops to reduce both power consumption and interference. In sparse areas, existing routing protocols that rely on full connectivity are likely to suffer. An explicit location aware scheme could be adopted for optimisation. It may also be possible to infer location information from the connectivity or contention level.

Relative mobility could be useful for predicting link durations. Links would be classified as 'persistent', intermittent or transient. Preferences could be attached to the link, either to stabilise the route or forward information opportunistically.

In view of network partitions and neighbour statistics, asynchronous routing, e.g., using a store-and-forward paradigm, could prove essential at times. Context awareness bears huge relevance at middleware and application levels, and could be further exploited to optimise operations at the network layer or below. Related previous works suggest the use of neighbour context based on such conjectures as 'frequently seen neighbours are likely to be seen again' [8], and our observations provide a supporting case. This context might be applicable to delay tolerant networking, when scheduled information is involved. For a more classical reactive routing protocol, co-location context could determine route discovery success.

Inter-contact durations are of particular importance to sparse networks and delay tolerant networking [13], for the distributions determine the viability of the forwarding algorithm, as shown in [7].

The representative cases can also form a test suite to identify the applicability of a general-purpose protocol, or compare a set of protocols.

8 Conclusions

We have proposed a new integrated framework, MIRRORS, that models real world ad hoc networks. This integration can be seen as across space, RF propagation and mobility, across nodes with different mobility capabilities and with consideration to both protocol design and evaluation. By presenting comparative results, we highlight fine variations within the network dynamics captured by the details in our model.

Parameters have been identified for models within our framework, and appropriate values could be extracted from a modest amount of real world traces to obtain the precise models. To that extent, our framework provides a mere necessary condition for realism.

Despite the microscopic focus, macroscopic behaviour of the networks could be derived from the models within the framework. This could facilitate fit-for-purpose protocol design or prescriptive optimisations of existing protocols. We call for a synergy between realistic modelling and protocol design, noting that a real world scenaro impacts *all* layers in different ways. The set of representative cases in the framework could also serve as a testsuite to evaluate the performance of a general protocol.

Furthermore, given the extreme flexibility of mobile ad hoc networks, these models could compute the exact node states in terms of positions and signal strengths. These could even be used in cell networks or WLANs to derive the likely associations with base stations, although simpler models targeting those networks are likely to be more computationally efficient.

References

[1] D Aguayo, J Bicket, S Biswas, G Judd, and R Morris. Link-level measurements from an 802.11b mesh network. In *Proceedings of ACM SIGCOMM*, 2004.

[2] F Bai, N Sadagopan, and A Helmy. IMPORTANT: A framework to systematically analyze the Impact of Mobility on Performance of RouTing protocols for Ad hoc NeTworks. In *Proceedings of IEEE INFOCOM*, 2003.

[3] C Bettstetter. Mobility modeling in wireless networks: Categorization, smooth movement, and border effects. *ACM Mobile Computing and Communications Review*, 5(3), 2001.

[4] C Bettstetter, H Hartenstein, and X Pérez-Costa. Stochastic properties of the random waypoint mobility model. *ACM/Kluwer Wireless Networks, Special Issue on Modeling and Analysis of Mobile Networks*, 10(5), 2004.

[5] J Broch, D Maltz, D Johnson, Y-C Hu, and J Jetcheva. A performance comparison of multihop wireless ad hoc network routing protocols. In *Proceedings of ACM/IEEE MobiCom*, 1998.

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

[7] A Chaintreau, P Hui, C Diot, R Gass, J Scott, and J Crowcroft. Pocket switched networks. Technical Report UCAM-CL-TR-617, University of Cambridge Computer Laboratory, 2005.

[8] X Chen and A Murphy. Enabling disconnected transitive communication in mobile ad hoc networks. In *Proceedings of ACM POMC*, 2001.

[9] R Draves, J Padhye, and B Zill. Comparison of routing metrics for static multi-hop wireless networks. In *Proceedings of ACM SIGCOMM*, 2004.

[10] R Gray, D Kotz, C Newport, N Dubrovsky, A Fiske, J Liu, C Masone, S McGrath, and Y Yuan. Outdoor experimental comparison of four ad hoc routing algorithms. In *Proceedings of ACM MSWiM*, 2004.

[11] X Hong, T Kwon, M Gerla, L Gu, and G Pei. A mobility framework for ad hoc wireless networks. In *Proceedings of ACM MDM*, 2001.

[12] E Huang, W Hu, J Crowcroft, and I Wassell. Towards commercial mobile ad hoc network applications: a radio dispatch system. In *Proceedings of ACM MobiHOC*, 2005.

[13] S Jain, K Fall, and R Patra. Routing in a delay tolerant network. In *Proceedings of ACM SIGCOMM*, 2004.

[14] A Jardosh, E Belding-Royer, K Almeroth, and S Suri. Towards realistic mobility models for ad hoc networks. In *Proceedings of ACM/IEEE MobiCom*, 2003.

[15] D Kotz, C Newport, and C Elliott. The mistaken axioms of wireless-network research. Technical Report TR2003-467, Dartmouth College Computer Science, 2003.

[16] K Merchant, W Hsu, H Shu, C Hsu, and A Helmy. Weighted way mobility model and its impact on ad hoc networks (poster). In *Proceedings of ACM/IEEE MobiCom*, 2004.

[17] M Musolesi, S Hailes, and C Mascolo. An ad hoc mobility model founded on social network theory. In *Proceedings of ACM MSWiM*, 2004.

[18] S R Saunders. *Antennas and propagation for wireless communication systems*. Wiley, 1999.

[19] M Takai, J Martin, and R Bagrodia. Effects of wireless physical layer modeling in mobile ad hoc networks. In *Proceedings of ACM MobiHOC*, 2001.

[20] C Tschudin, H Lundgren, and E Nordström. Embedding MANETs in the real world. In *Proceedings of IFIP-TC6 PWC*, 2003.

[21] J Yoon, M Liu, and B Noble. Random waypoint considered harmful. In *Proceedings of IEEE INFOCOM*, 2003.

[22] J Yoon, M Liu, and B Noble. Sound mobility models. In *Proceedings of ACM/IEEE MobiCom*, 2003.