Towards Commercial Mobile Ad Hoc Network Applications: A Radio Dispatch System

Elgan Huang, Wenjun Hu, Jon Crowcroft, Ian Wassell
University of Cambridge Computer Laboratory
William Gates Building, Cambridge, CB3 OFD, United Kingdom
eh283@cam.ac.uk wenjun.hu@cl.cam.ac.uk Jon.Crowcroft@cl.cam.ac.uk ijw24@eng.cam.ac.uk

ABSTRACT
We propose a novel and plausibly realistic application scenario for mobile ad hoc networks in the form of a radio dispatch system. We evaluate the system from both financial and technical perspectives to gain a complete picture of its feasibility. Using a realistic mobility and propagation model drawn from real world data we investigate the effects of node density, connection times and traffic congestion on the network coverage. We discuss design considerations in the light of the results. These findings are not limited to this particular scenario but are applicable to any mobile ad hoc system operating in similar conditions.

Categories and Subject Descriptors
I.6 [**Simulation & Modelling**]: Applications

General Terms
Algorithms, Performance, Design, Economics

Keywords
Mobile ad hoc networks, application, performance

1. INTRODUCTION
In recent years there has been a growing interest in using mobile ad hoc networks (MANETs) for commercial purposes. Within the automotive industry there are already a number of efforts geared towards utilizing communications technology to improve automotive safety, provide passengers with information and entertainment, and achieve smooth traffic flow on the roads [2]. However, there has been little effort to tie business and technical aspects together in order to design solutions that are both technically and financially feasible; most research is currently focused primarily on certain technical issues. Standard scenarios commonly quoted for MANETs tend to be either unnecessary or too limited in scope.

In this paper, we propose a novel application scenario in the form of a MANET dispatch system for taxis. A major novel contribution of this paper is that we use realism in our models at all levels, ranging from RF propagation in a city, through the mobility of the nodes, up to the application and on to the business case for using a MANET rather than a regulated, centralised classical provider-based system.
In the following sections, we begin by laying out the vision for the dispatch system, analysing the value chain and value proposition for each player and evaluating the operating costs of existing dispatch technologies. This is similar to the tussle considerations in [9]. Applications for MANETs need to bring value to users, whilst at the same time provide business incentives for corporate participation. We then investigate the technical feasibility of the system, using a realistic mobility and propagation model to simulate the system's real world nature and scale. We analyse various factors and limitations that affect system performance and look into methods of improving desired performance metrics.

Despite the paper's focus on the application proposed, the main objective of this work is to highlight general principles and design considerations that are applicable to large-scale highly-mobile MANET applications operating in microcell environments.

2. CITY TAXIS SCENARIO
One of the most important components of a city taxi company's operations is its dispatch unit, which informs individual taxis about passenger pickups, assigns passengers to empty taxis, and passes on additional information such as directions and news about weather or traffic conditions.

Traditionally, taxi companies have relied on radio dispatchers for these tasks. Radio dispatch systems have variable QoS issues, and there are significant costs associated with installing, running and maintaining them. In many cases, radio licences need to be obtained to operate these systems, adding considerable cost and creating barriers to entry for small companies. Economies of scale mean that large companies are able to share the cost of radio dispatch amongst many drivers, whilst smaller players are unable to afford these systems.

However, the popularisation of mobile phones now means that smaller companies are able to enter the market by relying on public mobile networks as their dispatch systems. Unfortunately, this method has many drawbacks including a high average cost per call, and limited communication between drivers. Mobile phone calls are harder to coordinate and less time efficient; dispatchers can only communicate with one driver at a time. Voice relayed instructions are also subject to higher error rates.

Many larger companies have in turn upgraded their own systems to include smart software with GPS which allows tracking of taxis and efficient allocation of jobs, improving productivity and utilisation of resources. However, the hardware and software required is costly (costing over £630,000 for a company with 300 taxis [26]) and its operation involves transferring data wirelessly over public or private mobile radio networks, which is a significant recurring cost.

In our scenario, a MANET based dispatch system is used for communication. Each taxi is fitted with an ad hoc device and a central ad hoc server is located at dispatch headquarters. This is a hybrid scenario in that although the system is a fully mobile ad hoc

MobiHoc'05, May 25–27, 2005, Urbana-Champaign, Illinois, USA.

network, there is a degree of centralisation to it, with all dispatch information originating from one point, requiring all nodes to connect to it directly or indirectly to receive jobs.

A customer would call in to the taxi company's dispatch headquarters. The call would then be picked up by a dispatch handler who would enter the job details into the system. An automated request would then be sent via the ad hoc network to all free taxis in the request vicinity. In the case of ad hoc devices enabled with location information, the nearest free taxi to the pickup point could be automatically located and assigned the job, thereby decreasing waiting times for the customer and improving turnover rates for the taxi company. The ad hoc system would constantly be updating itself automatically with relevant job and road condition information for drivers, and drivers would also be able to communicate amongst themselves, empowering all parties concerned with valuable information.

Taxi stands throughout the city could also be equipped with ad hoc terminals that automatically route taxis to appropriate stands when requested. Should ad hoc devices become ubiquitous, it would even become possible for customers to simply order a taxi via their own standard ad hoc devices wherever they were and have the request sent directly to the nearest taxi. The taxi would then register the pickup and send the information through the network to update other taxi drivers and headquarters of its status. With the addition of an access point at dispatch headquarters, customers would also be able to book taxis over the Internet and have the request relayed immediately via the ad hoc network.

In all cases there would be a two way flow of information; once a particular taxi driver has accepted a job, details such as the taxi's license plate, its current location and the estimated time of arrival can be sent back via the ad hoc network to the user's phone, to the taxi kiosk, to the customer's ad hoc device or to the messaging service of the Internet customer respectively, reducing any potential misunderstandings that are common with voice relayed radio dispatching services and allowing better planning on the customer's part.

If the network were secure enough, credit card payment options could also be implemented using the ad hoc terminals, with encrypted verification information sent to and from the central node.

We envision the rollout of the ad hoc dispatch system in three distinct stages. In the first stage, a medium sized company is able to deploy a system as described in this paper, with taxi stands that double as stationary nodes to improve relay performance. In the second stage, multiple companies use the system and compatible standards are such that different companies' ad hoc networks can relay data across each other, with information protected by encryption. In the final stage, mobile ad hoc devices have become ubiquitous and data can be relayed through a multitude of devices.

3. FINANCIAL FEASIBILITY

3.1 Value Proposition for Primary Players

The value chain shown in Figure 1 is a map of the primary players within the taxi company radio dispatch industry. These players make contracts between each other to conduct exchanges of information, money, services or a combination of the three. Dotted lines indicate the flow of information or services, for example the information that is shared between taxi drivers and dispatchers, whereas solid lines indicate a monetary exchange, for example when a customer pays a taxi driver for the taxi ride.

Customers would now be able to make bookings with greater ease, convenience and efficiency through a variety of different media besides the phone. Also, information flow is now two-way and customers can use the information provided for better planning. With location specific job assignments and less information congestion on the dispatch side, taxis can get to customers much more quickly, reducing their overall waiting times. The chances of getting a taxi should also be improved due to the higher turnover rate and accuracy of the system. Finally, customers will be able to enjoy a quieter environment in taxis, without constant radio chatter.

Taxi drivers would also benefit, since faster and more efficient processing of jobs potentially means faster turnover rates and less empty cruising time, thereby increasing revenue and improving fuel efficiency. Furthermore, decreasing reliance on voice dispatching results in fewer misunderstandings between dispatchers and drivers. With location information and a hidden alarm switch, drivers in distress can immediately be located and helped, in the event of an accident, a mugging, a carjacking, etc [8]. Finally, taxi drivers can also enjoy a quieter environment in taxis.

Taxi companies would benefit in many ways. Streamlining the order taking, order processing and dispatching process result in significant capacity increases, improving both revenue streams and service levels while lowering costs. The system itself could bring about significant long term savings as operational costs are marginal compared to conventional methods. High customer and employee satisfaction could improve the taxi company's reputation and hence demand for its services. With faster information flow, management's decision-making capabilities could be enhanced, and the reaction time for a decision to filter down to the driver level would decrease. Location information allows taxi tracking in the case of emergencies, as well as facilitating better route forecasting and resource allocation.

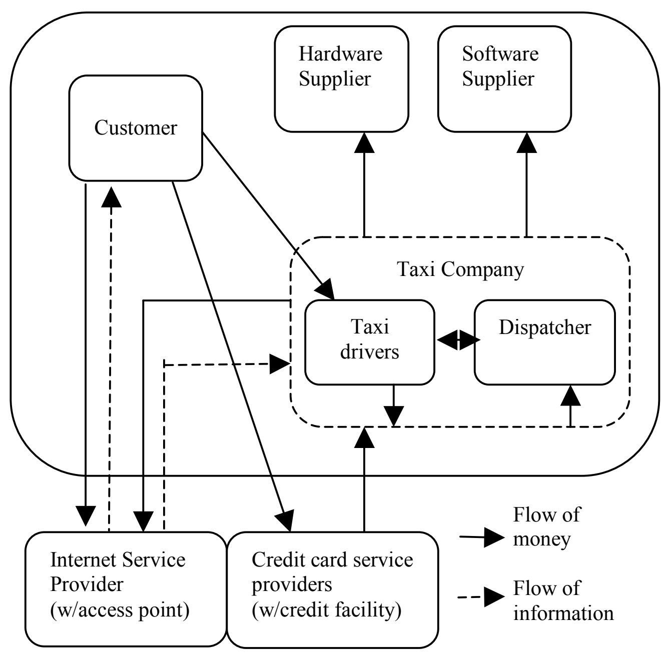

Figure 1: Value chain for a taxi company operation's primary players

Additionally, ad hoc systems are more robust than centralised systems, and network service is not dependent on the performance of service providers or subject to damaged infrastructure. Using decentralised ad hoc devices may also change the market structure by making it possible to offer off-the-shelf solutions to taxi companies, rather than forcing them to sign up with a service provider. Competition may then force service providers to offer more competitive prices.

Finally, as customers get used to alternatives to booking over the phone, the workload for ***taxi dispatchers*** will decrease. Reliable automated systems will improve the accuracy and speed with which orders are taken and passed on.

3.2. Comparison of Costs

Running a proprietary private radio network is very expensive, including infrastructure, spectrum licensing, maintenance and device costs. Consequently, this option is out of the question for most taxi companies.

The alternative for many smaller companies is the use of mobile phones. However, the cost of phone calls can mount up quickly. If each driver were to spend just one minute on each call with dispatch, and make 40 of these calls per day [17], at an average cost of 20p per minute (taking into account peak and non peak calls), it would cost £2,920 per taxi annually. A medium sized company with 300 taxis would therefore be paying £876,000 pounds a year on phone bills alone. Note that such a company could move to a system with better economies of scale, but the figure is useful for comparison purposes.

Subscribing to a public network radio service is generally a much cheaper option. A typical service would be a fixed charge, push-to-talk walkie-talkie type service, as provided by Dolphin Telecommunications [11] in the UK. With a monthly fixed charge of £25 per user plus VAT, a company with 300 taxis would need to spend £105,750 annually. Although much cheaper than using mobile phones, this is still a large recurring cost and is on top of hardware costs.

Another alternative is to use a Private Mobile Radio (PMR) system which can be operated by the taxi company or outsourced to a specialist service provider. In either case a license is required from Ofcom (Office of Communications) covering the intended service area. Through the use of a single radio communications mast a range of 15-20 miles can typically be covered at a relatively low cost. However this depends on the exact service used and features required. Panther Taxis [17] uses a service from Auriga Communications which costs £3,090 a year to operate (communication only), whereas a company like Cordic [10] might charge up to £124,800 a year (£8 per week per PDA rental).

In comparison, ad hoc systems cost nothing to run above the cost and maintenance of the hardware itself. In addition, both mobile phones and the push-to-talk service are voice only communication methods (data services would cost an additional premium), whereas an ad hoc system would be based on data communications and potentially also provide location information. According to [17], 95% of bookings are made within a 10-mile radius, which makes ad hoc networks feasible and potentially more scalable than licensed services.

MANETs are also likely to be much faster to roll out and cheaper in terms of hardware costs since they do not require any installation costs or infrastructure beyond the ad hoc units themselves.

4. TECHNICAL FEASIBILITY

Having laid out the business case for this conceptual system, we now turn to the technical feasibility. By nature, an ad hoc system's coverage depends on a suitable number and spread of nodes throughout the required service area. The mobility of nodes poses the danger of occasional service gaps when no taxis are acting as relays within a certain area, raising concerns of unpredictable QoS. Through simulations based on a realistic mobility and propagation model, we assess various performance metrics of the radio dispatch system to determine the extent of these problems.

4.1 Mobility model and assumptions

A number of mobility models have been proposed for MANET simulations, such as the Random Waypoint [6], Random Direction [23] and other Random Walk variants [7], amongst which the Random Waypoint is the most widely used. Despite its popularity, the Random Waypoint Model lacks both realism and desirable statistical properties [29]. Although there have been recent efforts to design more realistic models [18], the level of detail is insufficient for our feasibility evaluation. For this reason we did not use packages such as ns-2/GloMoSim, since their models were too simplistic. However, our simulation is compatible and our mobility traces can be fed into ns-2 using the *setdest* tool.

We model the city as a Manhattan style grid, as in [1], with a uniform block size across the simulation area. All streets are assumed to be two-way, with one lane in each direction. Taxi movements are constrained by these lanes, which are modelled using one-metre segments.

We assume that, under free flow conditions, each taxi's behaviour is homogeneous with respect to itself over time. A taxi is characterised by a preferred speed, a maximum acceleration and deceleration [3], a speed variation associated with the preferred speed at steady state, and a list of preferred destinations, i.e., the taxi stands. The taxis are randomly assigned one of three preferred speeds. All other parameters are set to be the same across all taxis, as they are from the same company.

At any instant, a taxi is i) carrying a passenger to a destination, assumed to be within the city, ii) heading for the taxi stand, or iii) roaming around until flagged by a passenger. The third case can be viewed as a taxi travelling to a particular location to collect a passenger. We combine into the roaming case the situation where a taxi collects a passenger according to the dispatch unit. Therefore, taxi journeys comprise of a sequence of movements of constant velocities, starting from the current position and ending at the destination. In the first two cases, the taxi pauses on reaching its destination; the roaming situation can comprise of several journey sequences with smooth transitions, depending on whether the taxi is to pause. The probability of a taxi picking up a passenger is estimated from a real taxi's daily empty cruising time.

To start a sequence of movements, a taxi selects a destination according to the following distribution: probability of 0.5 uniformly shared between all taxi stand(s), and otherwise uniform within the remaining 0.5 probability for any other position in the simulation area. Therefore if there is only one stand, there is a 50% chance that a taxi will head towards it, otherwise there is an equal chance of it moving towards any other position on the map. We assume that these taxi stands are located in popular areas that see particularly high traffic such as train/coach stations to account for popular hot spots around the city. Adjustments to the selection procedure ensure that the destination is not the current position. The movement algorithm in principle computes a shortest path towards the destination. Since such a path is not unique in the grid layout, the probability of taking each direction at a junction is determined according to the distance components to the destination. At the borders the algorithm specifies 'bounce-back' behaviour. On reaching the destination, the taxi either pauses according to an exponentially distributed pause time plus a minimum 30 seconds, or continues roaming into the next movement sequence. An adjustment is made when the destination is a taxi stand, where the taxi joins a queue and waits for its turn to take the next passenger. The selection of destinations as opposed to initial directions [3, 23] better represents the true nature of real world movement towards a target.

Successive movements within a sequence are correlated by the taxi's velocities [3, 20]. The direction will be along the same lane until a possible turn at the nearest junction. The speed variation is bound by the acceleration or the deceleration [25], and if at steady state, it is specified by the steady state variation which simulates speed drifts. This also ensures that a taxi speeds up and slows down gradually, respecting physical laws [3, 16], at appropriate times. Further constraints on the taxi movements are imposed by the speed limit of the lane [3] (set to be 15m/s), previous taxis in the same lane [1] and traffic lights. We assume traffic lights are installed at each junction, with the green and red lights on for 60 seconds alternately. All lights along an entire street are consistent. For simplicity, traffic lights do not affect taxis about to turn into another street.

At the start of a simulation run, the taxis are placed according to their destination distribution, and their speeds set to the preferred speeds, with adjustments for those at junctions. We believe this approximates a steady state setup, and will allow the simulation to converge quickly [28]. Taxi positions are updated at one-second time steps to approximate continuous motion. Given the potential inter-dependency between movements of taxis, (e.g., two taxis turning into the same lane from different directions, one constraining the movements of the other), they are updated in a randomised order at each second.

4.2 Propagation Model

The simulation models a MANET based on the IEEE 802.11b standard, since it is well understood, readily available and adaptable to the application under consideration. IEEE 802.11b operates in the 2.4 GHz ISM band, and we assume a typical receiver sensitivity of -80dBm and a transmit power of 100mW (20dBm) (the legal limit for this band). Since our system operates using taxis' electrical systems as opposed to being powered by batteries, power consumption is not a big issue.

The results of the simulation will depend heavily on the environment that the nodes operate within. In this case, it is our intention to simulate harsh conditions in order to test the feasibility of the scenario. As such, we have used a Manhattan grid model where signal propagation is limited by blocks of buildings placed at regular intervals.

Since our omni-directional antennas will be mounted atop the roofs of taxis at a height of around 1.5m, propagation will be characterised most accurately by a microcell model. The dominant propagation mechanisms in such an environment are due to interactions between the direct path and paths reflected from buildings and the ground, as well as diffraction around the vertical edges of buildings. Given the low heights of our antennas, we ignore diffraction over rooftops.

We also ignore interference from other nodes in the model since the density of nodes is very low. Should node density increase, the distance from node to node would decrease and the transmit power could simply be lowered to shorten the propagation range, effectively maintaining or even increasing the bandwidth available. Coupled with the low heights of the antennas this would lead to very small reuse distances, thereby minimising problems from interference.

Our propagation model is a dual slope empirical model appropriate for path losses in microcells [24]. Essentially this model assumes a 2nd order loss (-20dB/dec) out to a 'break' distance and a 4th order loss (-40dB/dec) thereafter. For antenna separations greater than the 'break' distance, the plane earth propagation model yields,

$$\frac{P_r}{P_T} = \left(\frac{h_1 h_2}{d^2}\right)^2$$

where P_r is the received power (W), P_T is the transmit power (W), h_1 and h_2 are the node antenna heights (m) and d is the antenna separation (m). The breakpoint distance is approximately given by, $d_b = \frac{4\pi h_1 h_2}{\lambda}$, and is the point after which the ground reflection destructively interferes with the direct ray and reduces the field strength. For the simulation, the 'break' distance is set at 100m. Both nodes are assumed to be in line of sight.

Thus, before the breakpoint we have,

$$Gain(dB) = -20\log_{10} d - 32.9563 \quad (1)$$

and after the breakpoint,

$$Gain(dB) = 20\log_{10} h_1 + 20\log_{10} h_2 - 40\log_{10} d \quad (2)$$

To model the diffraction loss at the corner of buildings, we estimate that every time a corner is taken, a loss of 20dB is incurred in addition to the losses in equations 1 and 2. We also assume digital transmission via a narrowband channel. In other words, multipath propagation is assumed to be sufficiently small that the reception and detection of transmitted data will not be degraded by its presence. This is reasonable because the small coverage areas of microcells and line of sight paths lead directly to a reduction in multipath delay spread; fast fading data is most likely to conform to Rician rather than Rayleigh statistics [14].

Using the conditions specified, the maximum unobstructed transmit distance is 474m along a straight road and 150m if one corner is taken.

4.3 Network Operations

Operations in the system are based on simple message exchanges. Taxis regularly send status updates to dispatch, and dispatch sends job requests and other information to taxis. Messages are always sent to the nearest reachable hop in order to minimise transmit power and interference. Given that we know the location of dispatch and any taxi stands, as well as the taxis themselves, the system could use a location aware routing protocol such as location-aided routing (LAR), geographic distance routing, grid, zone-based two level routing [27] or a customised derivative. However, routing and MAC layer issues are beyond the scope of this paper, which is concerned mainly with the physical layer feasibility of such a mobile ad hoc network. It is assumed that upper layers are functioning appropriately; as long as a signal of sufficient strength is

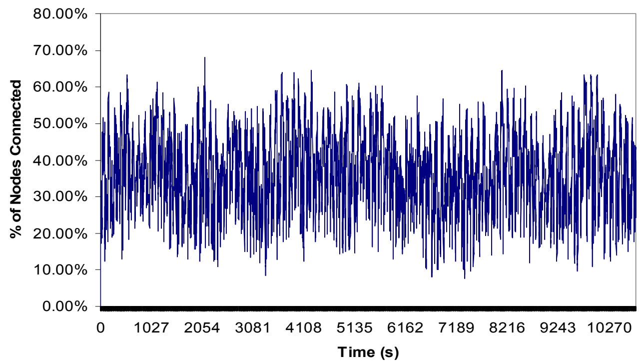

Figure 2: Typical Coverage Under Control Conditions

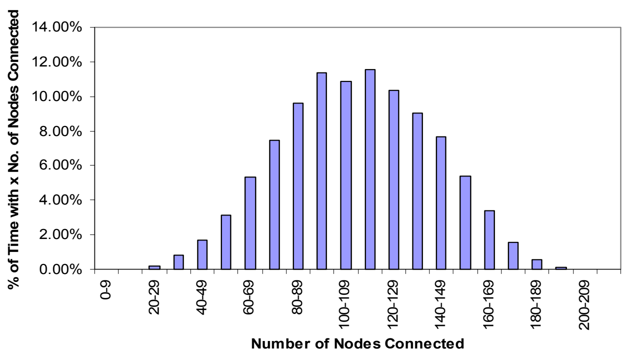

Figure 3: Typical Distribution of Coverage

received we assume that communication of acceptable QoS can be achieved. By taking a purely physical viewpoint, we are able to determine whether connectivity is physically achievable, independent of the performance of upper layers, which tend to vary considerably.

Being an ad hoc network, it is likely that taxis will wander in and out of coverage, the frequency and duration of which depend on the number of relay points, transmission range of each device, movement of vehicles, average distance between cars, etc.

However, as long as the period of outage is not too great, this may not necessarily impact performance significantly. Although real time communications require constant connectivity, a dispatch system does not need to transfer job data instantaneously and can operate on an asynchronous messaging basis. Since the amount of data required is very small (names, pickup and destination addresses, special instructions etc), data can be transmitted quickly whenever nodes are within coverage. Therefore as long as outage periods are short, the fact that connections are not constant should be of minimal concern to users and performance results should be interpreted with this in mind.

Even when the amount of data required is large, from similar types of studies, it is possible that in conditions of intermittent connectivity, a buffer or drive-thru proxy [12, 22] (based at dispatch, taxi stands and the individual nodes themselves) could be used to accumulate data and requests/responses from peers within the network and forward the information at the next opportunity without significantly sacrificing performance.

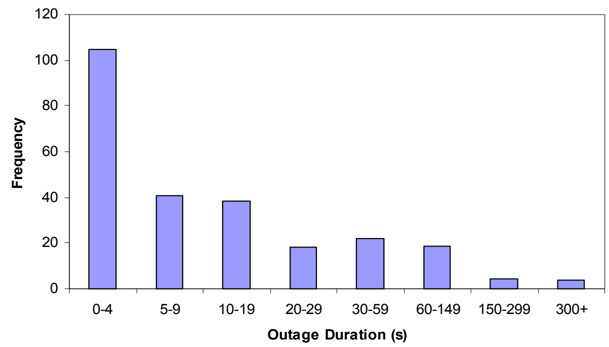

Figure 4: Typical Distribution of Outage Durations

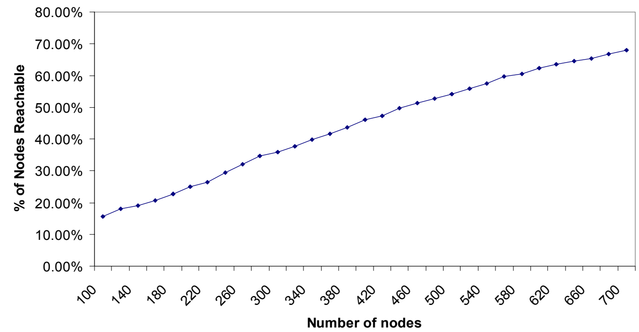

Figure 5: Percentage of Nodes Reachable as a Function of the No. of Nodes

5. SIMULATION RESULTS

Based on the model, we developed a statistical simulation to calculate the positions at 1 second granularity, from which connectivity information is derived. The control parameters for the simulation are as follows: We used a Manhattan grid of 5km x 5km, with a uniform block size of 100m x 100m. There is one central dispatch point located at the exact centre of the grid and this doubles as a taxi stand. Exact results vary according to the layout and distribution of taxi stands/popular destinations, however, the same general findings and system design considerations apply. Each simulation run lasted for 3 hours in simulation time after an initial warm-up period of 1,000 seconds. We chose 3 hours to give each taxi sufficient time to make several journeys to permit the speed and destination distributions to converge. Within the grid we placed a relatively conservative total of 300 ad hoc enabled taxis. A larger taxi company would obviously have more cars, and the number of nodes would also increase if different companies shared the same ad hoc system. Finally, a node is considered connected only if it is reachable from the dispatch point (either directly or via a number of intermediate hops) for three consecutive seconds or more, to ensure sufficient time for connection establishment and data transfer.

The following data were generated from up to 100 simulation runs each to ensure statistical significance. We define connectivity coverage as the percentage of taxis reachable from the central dispatch, and outage time as a continuous period of time when the taxi is not reachable from the central dispatch. Figure 2 shows a typical graph of coverage under control conditions. Clearly, the percentage of cars connected at any given point in time fluctuates, however there are always cars within coverage and even with only 300 nodes, the mean coverage over 100 runs is 107.7 cars or

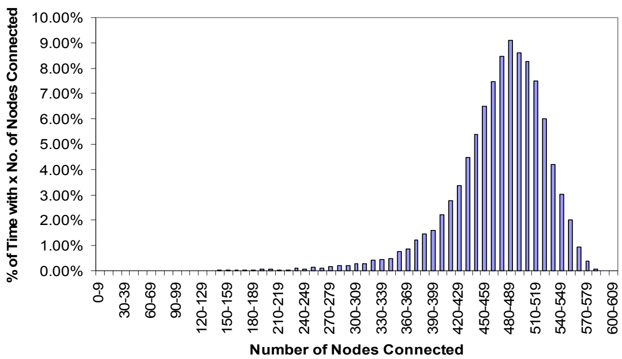

Figure 6: Typical Distribution of Coverage with 700 nodes

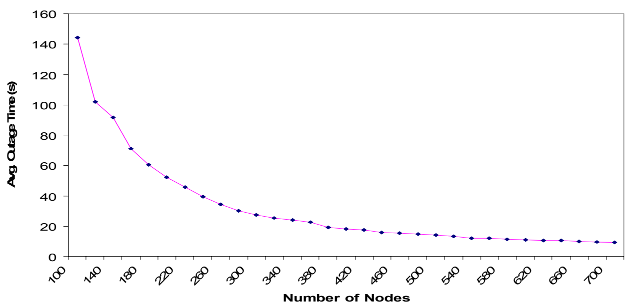

Figure 7: Average Outage Time as Function of No. of Nodes

35.91% (with a 95% confidence level of 0.12%). The median and the standard deviation of the mean is 35.84% and 0.6% respectively. The distribution of the coverage is fairly Gaussian in nature (Figure 3). Given the small error bounds, we omit error bars in the plots.

We calculated the length of outage durations to assess the extent of intermittent connectivity that the system needs to be able to handle. For control conditions, the average outage time was 28.47s, with a 95% confidence level of 0.15s, which should be acceptable where real time communication is not essential. However, given the complex relative motion between nodes, the maximum time that a node might be unreachable was considerably longer, averaging 696.49s (11 min), with a 95% confidence level of 3.79s. The longest time ever observed for a node being out of coverage was 2,785s (46 min), albeit an exceptional occurrence. As shown in Figure 4, the majority of outages are fairly short. We also need to consider the fact that with each hop across nodes there is an associated delay that could further increase the end to end latency. However, this is unlikely to be significant compared to the outage times of an unreachable node.

Therefore, the simulation seems to suggest that with 300 nodes in operation, real time communication is not feasible with only one third of taxis reachable at any given instant. Considering the absolute numbers of contactable taxis, however, the performance is adequate for the purposes of the dispatch system, given the likely number of requests at any typical point in time. Also, given the short time that each node spends out of coverage on average (28.47s), all taxis should be contactable within reasonable intervals. Conversely, on occasion certain nodes will be unreachable for a longer period (for example when taking passengers to far away destinations) and backup systems may be necessary.

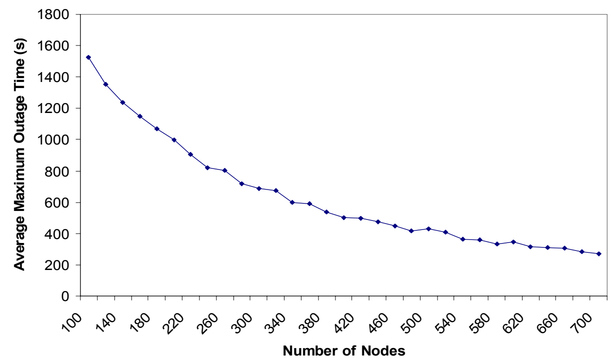

Figure 8: Average Maximum Outage Time as a Function of the No. of Nodes

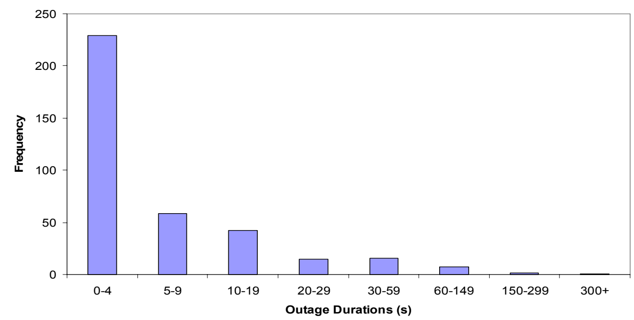

Figure 9: Typical Distribution of Outage Durations with 700 nodes

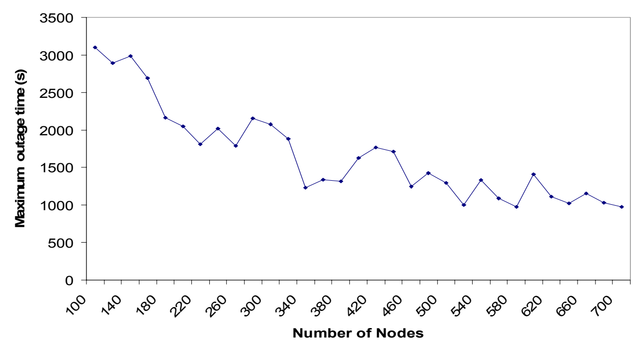

Figure 10: Maximum outage time

5.1 Node Density

To investigate the effect of node density, we varied the number of nodes from 100 to 700. Figure 5 shows the percentage of nodes reachable as a function of the number of nodes.

As expected, coverage improves as more nodes are available to act as relay points. With 700 nodes, almost 70% of taxis are reachable at any given time. By increasing the number of nodes, a larger taxi company or a number of small/medium companies that use compatible systems would be able to achieve proportionally good coverage. Figure 6 shows the distribution of coverage for 700 nodes. Compared to Figure 3, more nodes are connected more frequently resulting in the distribution skewing significantly.

Figures 7 and 8 illustrate the variation of average outage times and average maximum outage times with node density. The 'average maximum outage time' is defined as the average of all the maximum outage times of all the nodes over each simulation run.

Again, the trend shows clearly that as the number of nodes rises the average outage time decreases. However, the rate of decrease grows less pronounced with saturation effects becoming evident, such that the incremental benefit lessens. With 700 nodes, the average time a node spends out of coverage is only 9.4s. Likewise, the average maximum time that a node spends out of coverage decreases in a similar manner, with 270.1s outage time at 700 nodes (compared to almost 700s with 300 nodes).

Figure 9 shows the distribution of outage durations, and once again in comparison to Figure 4, with 700 nodes we observe a skew to the left, with most outages being of a very short duration.

Figure 10 shows the maximum time that a node was observed out of coverage as a function of the number of nodes. Due to the complex movements of the nodes, the plot exhibits a high degree of randomness owing to the considerable variation between simulation runs.

Overall, the decreasing trend is still clear, with an overall maximum outage time of 974s with 700 nodes, compared with well over 2,000s for 300 nodes. Figures 7 to 10 show that as the number of nodes increases, the time that a node is likely to spend out of coverage decreases significantly and the system grows closer to being able to support near real time communications. The dispatch system would be able to operate in such conditions with users largely unaware of the delay in receiving messages. However, longer outages could still occur, necessitating backup systems.

For the ubiquitous stage of deployment described at the end of Section II, results show that with 3,000 nodes, an average coverage of 95% can be achieved.

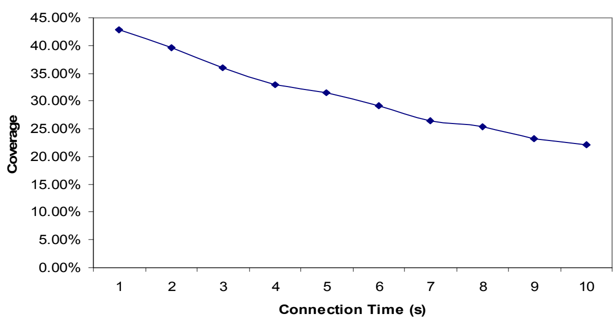

Figure 11: Coverage Variation with Connection Time

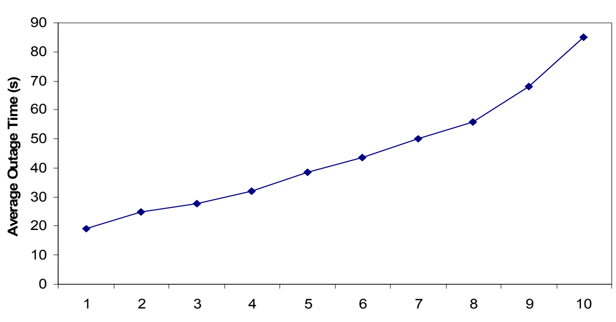

Figure 12: Average Outage Time vs. Connection Time

5.2 Connection Time

Another parameter that greatly affects the coverage statistics is the connection time, i.e., the time to set up links and transmit data between neighbouring nodes. The default connection time is set to 3s, however this depends on the exact protocols and hardware used. Keeping the other control parameters constant, we varied the connection time between 1 and 10s, with the results shown in Figures 11, 12 and 13.

With a connection time of 1s, 43% of nodes are within coverage on average, compared with 36% for a connection time of 3s. Given the potentially high relative mobility of the nodes, intuitively, the faster a connection is established and data transferred, the more likely it is that a passing node can stay within coverage for the time required. Likewise, it becomes less likely that a node is out of coverage for a long period of time.

The combined effects of larger network sizes and shorter connection times could mean significant performance gain, as presented in Figures 14 to 16. With 700 nodes and a 1s connection time the average coverage achieved was 77%, and the average outage time dropped to 8.8s. The majority of nodes are now connected most of the time, and generally experience very short outage durations.

Techniques that speed up connection times, such as making use of location awareness to predict handovers, may therefore result in significant improvements in connectivity.

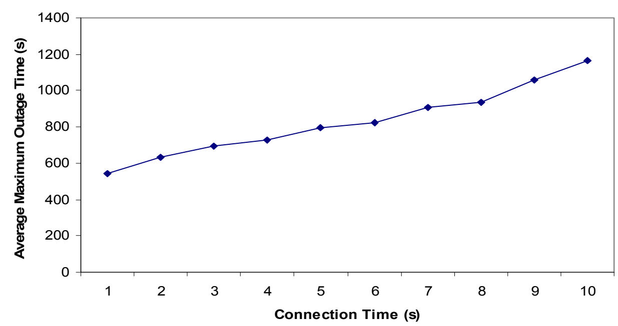

Figure 13: Average Maximum Outage Time vs. Connection Time

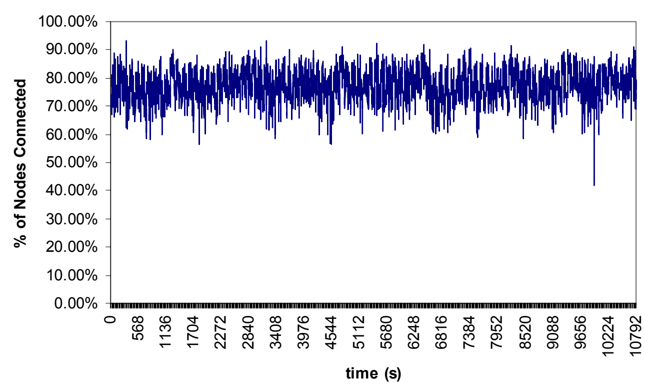

Figure 14: Typical Coverage with 700 Nodes and a 1s Connection Time

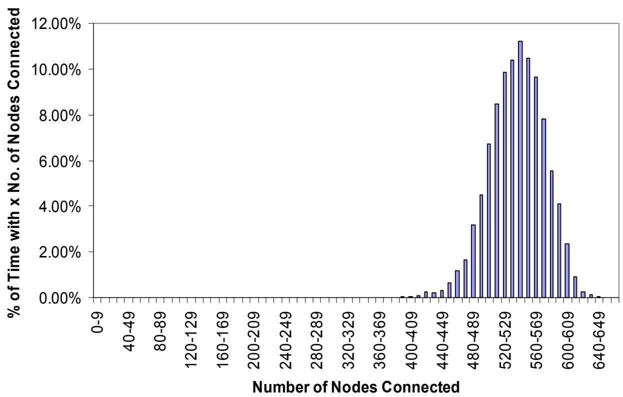

Figure 15: Typical Coverage with 700 Nodes and a 1s Connection Time

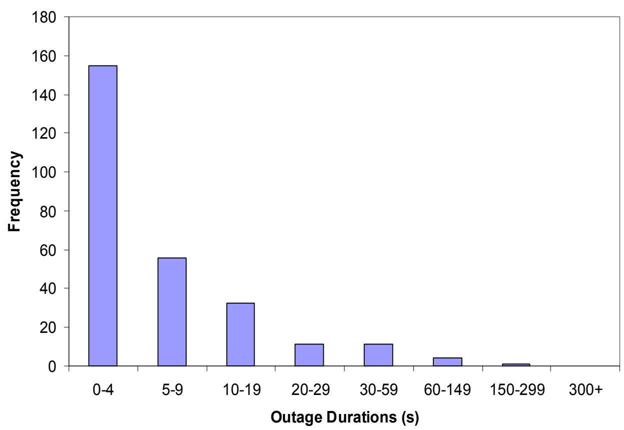

Figure 16: Typical Distribution of Outage Durations with 700 Nodes and a 1s Connection Time

5.3 Traffic Congestion

Finally, we investigated how traffic congestion would affect the results. To be feasible, the dispatch system has to be able to operate at all hours of the day regardless of traffic conditions. Traffic is introduced by inserting cars that behave the same as taxis, except that they do not have preferred destinations and do not pause between movement sequences. We therefore varied the number of cars from 300 (only taxis) up to 30,000 (300 taxis and 29,700 regular vehicles) to study how congestion affects node distribution, mobility and coverage. If all 30,000 cars were evenly distributed throughout the map, the average distance between vehicles would be 34m. The results are shown in Figures 17 to 24.

Our first observation is that the instantaneous variation in coverage is more noticeable with added congestion. Comparing Figure 17 with Figure 2, the peak to peak variations are larger, and the peaks and nulls in Figure 17 appear to be more regular and evenly spaced. Comparing Figure 18 with Figure 3 (30,000 cars vs. 300 cars), the distribution with congestion appears much more spread out, exhibiting bimodal characteristics. Although the mode and average coverage increase only slightly (38.4% vs. 36.3%), at any given instant, it is now much more likely that a higher number (over 140) of cars will be connected. On the other hand, it is also more likely that a lower number (below 70) are connected.

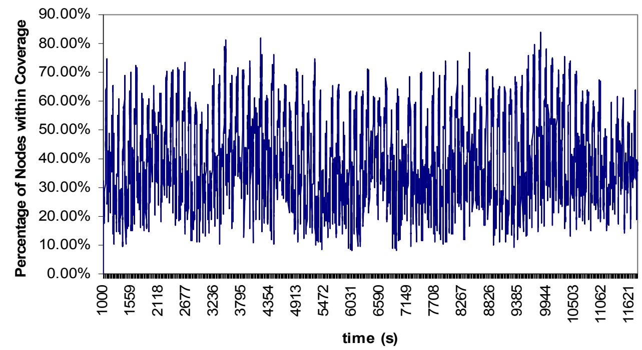

Figure 17: Typical Coverage with Congestion (30,000 cars)

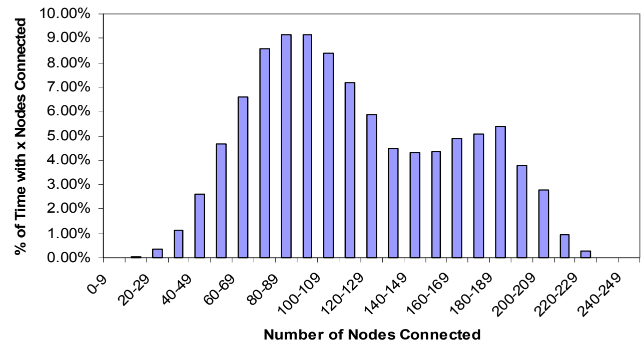

Figure 18: Typical Distribution of Coverage with Congestion (30,000 cars)

Indeed, in Figure 20 we can see that the standard deviation of coverage increases with congestion. However, after a steep rise it appears to flatten out at approximately 15,000 cars. Variability extends to outage times as well. Despite a nominal mean outage time change from 28.3s to 30.8s, a comparison of Figure 19 and Figure 4 shows that the distribution has spread out, making it is more likely that slightly longer periods of outage will be experienced.

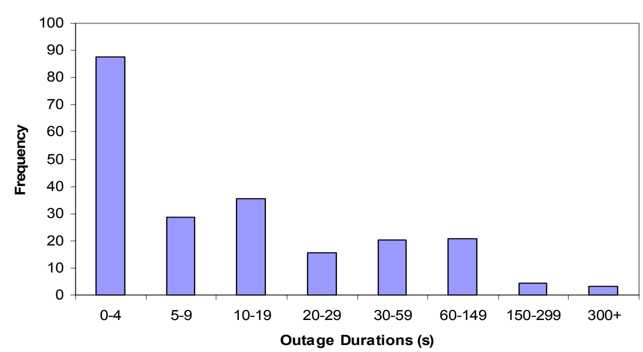

Figure 19: Typical Distribution of Outage Durations with Congestion (30,000 cars)

The effects of congestion appear erratic. However, the average coverage, as well as the average and maximum outage times, generally increases with congestion up to approximately 15,000 cars as shown in Figures 21 to 24. This is slightly counterintuitive; in previous cases when coverage has improved, outage time has decreased. This is due to the fact that increased congestion slows down the mobility of nodes in high traffic areas, such that they are more likely to be in contact for the required duration, hence improving coverage. At the same time, nodes that are out of coverage and caught in traffic are more likely to experience longer periods of outage.

Qualitative analysis suggests the results are attributable to non-uniform distribution of nodes and the mobility of nodes. Given that the destination selection mechanism for a movement sequence in our mobility model is an extension of that in the Random Waypoint Model, it is no surprise that the results exhibit Random Waypoint-like baseline behaviour. It has been shown that the asymptotic node density for the Random Waypoint Model is much higher at the centre of the simulation area but tends to zero at the borders, and the distribution is independent of node speeds [4, 5]. The placement of the taxi stand further contributes to the aggregation of the nodes. The non-uniform spatial distribution also applies to the non-taxi traffic, which mimics heavier traffic in the city centre. Meanwhile, mobility decreases under congested conditions, which generally reduces relative speeds. Therefore more nodes are likely to be close to the dispatch unit and hence reachable. Given the non-uniform distribution of congestion, however, nodes outside the highly congested city centre are able to move at higher speeds, and hence connectivity is poorer. On the other hand, the periodic variation is due to the density wave phenomenon observed for the Random Waypoint Model [23]. Similarly, our taxis are more likely to select paths through the central part of the simulation area, and therefore we can expect periodic dispersion of the nodes.

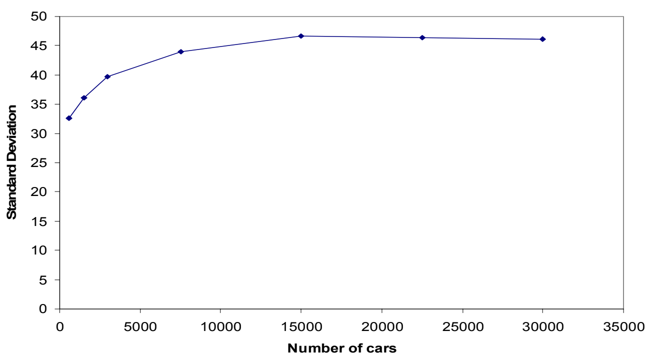

Figure 20: Standard Deviation of the coverage as Function of Congestion

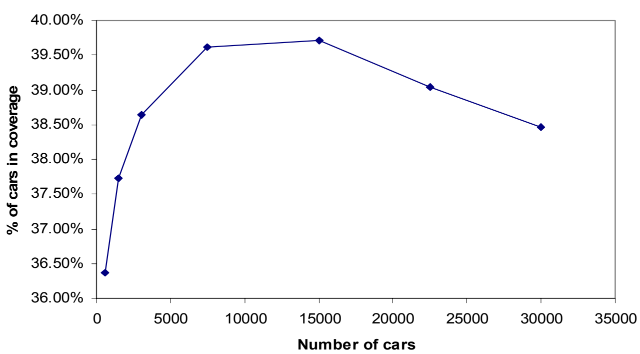

Figure 21: Coverage as Function of Congestion

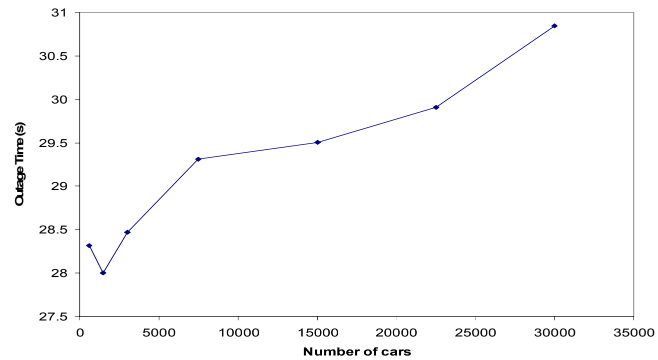

Figure 22: Average Outage Time

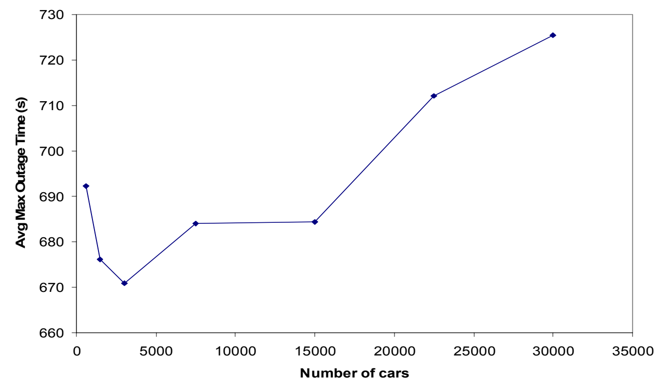

Figure 23: Average Maximum Outage Time

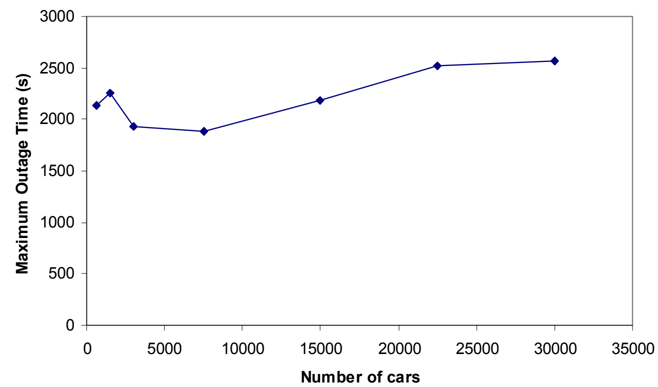

Figure 24: Maximum Outage Times

Above 15,000 cars, congestion is spread increasingly throughout the entire map such that all areas are now highly congested rather than just central, high traffic areas, and hence nodes are now more likely to be stuck in outlying areas, out of coverage. This subsequently results in a slight dip in average coverage and an increase in outage times as can be seen in the figures. Further investigations of the phenomena are under way to derive quantitative results.

These results suggest that with light to medium congestion the dispatch system will be able to operate on average as well or better than in zero congestion conditions. However, in conditions of very high congestion, coverage may suffer slightly and in any case congestion increases the risk of longer outages. In future work we intend to investigate the effect of events on mobility and coverage, such as football matches where many taxis will be sent to a particular high traffic area.

6. RELATED WORK

Projects such as [2] and [21] have explored various technological aspects of vehicular based mobile ad hoc networks, however there have been few specific commercial applications proposed to date. Most research seems to be centered on upper layers such as routing protocols and do not consider physical layer limitations. Simulations such as [19] have used real life bus traces to model general vehicular mobility. However, bus movements are regular and periodic, and only cover certain parts of the city, whereas taxis and other vehicles tend to have somewhat random destinations and pervasive trajectories. Furthermore, the physical layer was not considered, and it was assumed that signals had a 1.5 km radius range despite the city block environment, which may be overly optimistic.

Unlike related work, our model is based on a truly ad hoc scenario in its mobility, as opposed to scheduled movement like buses, or purely random movement as in many simulation studies, and our simulation is large scale and realistic.

7. CONCLUSION

In this paper we have investigated the concept of a MANET based radio communications dispatch system. Unlike conventional networks, the taxi network has a central origin of messages although communication is achieved in an ad hoc fashion. We looked at how the scenario might work and why it would be desirable as well as the risks and disadvantages of the system. By evaluating the system from both financial and technical perspectives, we are able to form a more complete picture of the systems feasibility.

Such a system has the potential to bring value to all players involved, increasing the revenue per taxi, decreasing the cost and time to fulfill jobs and improving profitability per customer as well as customer satisfaction. The system would also be fast to deploy and scalable in size.

To evaluate the technical feasibility, we developed a realistic mobility and propagation model to establish the effects of node density, connection time and traffic congestion on various metrics of performance.

From the results, we deduce that whilst such a system would not be usable for real time communications, performance would be satisfactory for the purposes of the dispatch application. Since parameters and conditions in the simulation were meant to simulate harsh conditions, it is possible that in less difficult real world environments performance might be better than expected. However it is to be noted that performance also depends on factors such as routing protocols which are assumed to work as desired in this model. Given the particular information being transmitted and the environment, fit-for-purpose protocols and optimisations could be designed to suit the application. However, the focus of this paper has been to consider the physical layer effects independently of upper layers, which has allowed us to effectively identify an upper physical bound for connectivity in this network setting.

Given the unpredictable performance of the system, backup systems (such as mobile phones) may need to be in place for nodes which are stuck outside of coverage for an unusually long time.

It appears that QoS can be improved to suit individual specifications by increasing the number of nodes or decreasing the connection time or both. One option would be to introduce additional stationary nodes that double as taxi stands. This would also help to reduce the variability of coverage and outage times.

Although a MANET based dispatch system appears to be both technically and financially feasible, there are a number of risks associated with it, in addition to the switching costs incurred by incumbent taxi companies. Given the incremental short term benefits, taxi companies may well choose not to adopt the new technology despite the potential long term cost savings. This could change depending on a variety of factors including technology and regulation. As spectrum pricing is currently under evaluation, it is possible that new regulations could make licensed bands more expensive, hence increasing the cost advantage for ad hoc networks. This could provide the necessary impetus to switch.

New operators would be able to adopt this system without worrying about switching costs, however the system will not be suitable for small companies with only a few cars until the system is standardised and competing companies can share the same relay system. A caveat is that standardising systems brings in additional issues to be resolved such as cooperation incentives.

The main alternatives to ad hoc systems are public or private radio networks. Research is also being performed to link conventional wireless LAN systems with seamless handoffs. The primary advantages that ad hoc networks have over these other forms of communication are low cost and ease of deployment. Should the alternatives advance sufficiently in performance or decrease substantially in price, they could render ad hoc systems inadequate in comparison for commercial uses.

On the technical side, security issues may raise concerns. Transmission must be secure to prevent eavesdropping of sensitive information. In the case of a standardised ad hoc network shared between different taxi companies, data encryption would avoid potential problems where companies attempt to poach each other's customers. Moreover, encryption is vital if credit card processing facilities are to be offered to customers. Situations may also arise where competitors try to perform denial of service attacks or otherwise disrupt communications to gain unfair competitive advantages.

Scalability is also an issue which should be considered. Both the software and the underlying radio technology must allow for significant growth in the number of nodes. With restricted bandwidth in a given area, there is a potential limit to the number of users that can be supported due to interference. This is unlikely to be an issue for taxi companies alone, but should the devices be sharing spectrum with other devices, problems could occur. One way to mitigate these problems would be to lower transmit powers or use adaptive/cognitive radios. Furthermore, opening up additional unlicensed bands could provide room for growth.

Regardless of whether the application is implemented in the proposed form, the main objective of this paper is to highlight issues concerning the design of MANET based systems in city environments with high mobility nodes. The technical results discussed are applicable to any MANET application operating under similar conditions, as are many of the business concerns. An important conclusion from this work is that unless a critical mass of nodes can be deployed at the same time, a MANET based system could suffer from poor performance due to low node density. The expected number of nodes in operation needs to be considered when designing a system. Moreover, it is recommended that these systems be able to handle intermittent connectivity in light of the short but frequent outages which will be experienced [13].

8. REFERENCES

[1] F. Bai, N. Sadagopan, A. Helmy, "IMPORTANT: A framework to systematically analyze the Impact of Mobility on Performance of Routing protocols for Adhoc NeTworks." IEEE INFOCOM, March 2003

[2] M.Bechler, W.J. Franz and L.Wolf, February 2003, "Mobile Internet Access in Fleetnet", Verteilten Systemen KiVS 2003

[3] C. Bettstetter, "Smooth is Better than Sharp: A Random Mobility Model for Simulation of Wireless Networks", MSWiM'01, July 2001

[4] C. Bettstetter, H. Hartenstein, X.Pérez-Costa, "Stochastic Properties of the Random Waypoint Mobility Model", in ACM/Kluwer Wireless Networks, Vol. 10, No. 5, Sept 2004

[5] D.M. Blough, G. Resta , P. Santi, "A Statistical Analysis of the Long-Run Node Spatial Distribution in Mobile Ad Hoc Networks", in MSWiM'02, Sept 2002

[6] J. Broch, D. A. Maltz, D. B. Johnson, Y. C. Hu, J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols", in ACM/IEEE MobiCom, 1998

[7] T. Camp, J. Boleng, V. Davies, "A Survey of Mobility Models for Ad Hoc Network Research", in WCMC: Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications, vol.2, no.5, 2002

[8] Y. Cheng, "Singapore's Smart Taxis", Asian Technology Information Program (ATIP), March 1996

[9] D.D.Clark, J. Wroclawski, K.R.Sollins, R.Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", in ACM SIGCOMM, Aug 02

[10] Cordic Communications, www.cordic.com

[11] Dolphin Telecommunications, http://www.dolphin-telecom.co.uk

[12] M. Esbjornsson, O. Juhlin and M. Ostergren, "The Hocman Prototype-fast motor bikers and ad hoc networking", International Conference on Mobile and Ubiquitous Multimedia Dec 2002

[13] K. Fall, "A Delay-Tolerant Network Architecture for Challenged Internets", in ACM SIGCOMM, August 2003

[14] E Green, "Radio link design for microcellular systems", British Telecom Technology Journal Vol 8 No 1 Jan 1990

[15] I. Gruber, O. Knauf, H. Li "Performance of Ad Hoc Routing Protocols in Urban Environments", *European Wireless 2004*, Feb 2004

[16] X. Hong, T.J. Kwon, M. Gerla, D. Gu and G. Pei, "A mobility framework for ad hoc wireless networks", in ACM International Conference on Mobile Data Management (MDM), Jan 2001

[17] Interview with director of Panther Taxis, Cambridge

[18] A. Jardosh, E. M. Belding-Royer, K.C. Almeroth, S.Suri, "Towards Realistic Mobility Models For Mobile Ad hoc Networks", ACM/IEEE MobiCom, Sept 2003

[19] J. Jetcheva, Y.C. Hu, S. PalChaudhuri, A. Saha, D. Johnson, "Design and Evaluation of a Metropolitan Area Multitier Wireless Ad Hoc Network Architecture", in Mobile Computing Systems and Applications, 2003

[20] B. Liang and Z.J. Haas, "Predictive Distance-Based Mobility Management for PCS Networks", in IEEE INFOCOM, March 1999

[21] R. Morris, J. Jannotti, F. Kaashoek, J. Li, D. Decouto, "CarNet: A Scalable Ad Hoc Wireless Network System", ACM SIGOPS, Sept 2000

[22] J. Ott,D. Kutscher, "The "Drive-thru" architecture: WLAN-based internet access on the road", IEEE Vehicular Technology Conference May 2004

[23] E.M. Royer, P. M. Melliar-Smith, L.E. Moser, "An analysis of the optimum node density for ad hoc mobile networks", IEEE ICC, June 2001

[24] S. R. Saunders, "Antennas and propagation for wireless communication systems" WILEY 1999

[25] C. Schindelhauer, T. Lukovszki, S. Rührup and K. Volbert, "Worst Case Mobility in Ad Hoc Networks", in SPAA'03, June 2003

[26] J.R. Stone, G. Gilbert, and A. Nalevanko, "Assessment of Computer Dispatch Technology in the Paratransit Industry", U.S. Department of Transportation (monetary figures converted to current values using http://eh.net/hmit/ppowerusd/)

[27] Y.C. Tseng, S.L. Wu, W.H. Liao, C.M. Chao, "Location Awareness in Ad Hoc Wireless Mobile Networks" Computer, Vol: 34/6, June 2001

[28] J. Yoon, M. Liu and B. Noble, "Sound Mobility Models", in ACM/IEEE MobiCom, Sept 2003

[29] J. Yoon, M. Liu, B. Noble, "Random Waypoint Considered Harmful", in IEEE Infocom, April 2003