

Research Article

Reducing Congestion in Obstructed Highways with Traffic Data Dissemination Using Ad hoc Vehicular Networks

Thomas D. Hewer,^{1,2} Maziar Nekovee,^{1,3} and Peter V. Coveney²

¹Department of Computer Science, UCL, Gower Street, London WC1E 6BT, UK

²Centre for Computational Science, UCL, 20 Gordon Street, London WC1H 0AJ, UK

³Mobility Research Group, BT Research, Polaris 134, Adastral Park, Martlesham IP5 3RE, UK

Correspondence should be addressed to Thomas D. Hewer, t.hewer@cs.ucl.ac.uk

Received 1 December 2009; Revised 14 June 2010; Accepted 5 July 2010

Academic Editor: Shahrokh Valaee

Copyright © 2010 Thomas D. Hewer et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Vehicle-to-vehicle communications can be used effectively for intelligent transport systems (ITSs) and location-aware services. The ability to disseminate information in an *ad hoc* fashion allows pertinent information to propagate faster through a network. In the realm of ITS, the ability to spread warning information faster and further is of great advantage to receivers. In this paper we propose and present a message-dissemination procedure that uses vehicular wireless protocols to influence vehicular flow, reducing congestion in road networks. The computational experiments we present show how a car-following model and lane-change algorithm can be adapted to “react” to the reception of information. This model also illustrates the advantages of coupling together with vehicular flow modelling tools and network simulation tools.

1. Introduction

In the realm of vehicle-to-vehicle (V2V) communications there are several methods for the dissemination of data that are being actively researched. The use of satellite communication such as those linked to global positioning services (GPS) offers global communication but requires expensive equipment, large antennae (for two-ray transmission) and, due to the large distances the signal must travel, they will have a high latency. Cellular telephone networks offer a lower latency over large distances but are still slow when communicating with nearby vehicles, due to a centralised approach, and require cellular contracts to use the network. The scenarios presented in this paper require high-speed communication disseminated from source, which is difficult to achieve using either cellular or satellite communications.

Ad hoc networks offer a good method to spread information outwards from an origin quickly and efficiently. It has been shown by Nekovee [1] that in *ad hoc* networks worms spread in an epidemic pattern that can be modelled. Using such modelling techniques we can develop algorithms that allow for a change to be made to the speed, position and

route of a vehicle. A further advantage of *ad hoc* networking is the unlicensed use of the radio-spectrum and the recent reduction in cost for the equipment for communication. A separate strand of the research being undertaken on wireless fidelity (under IEEE standard 802.11) has been developed in the past few years specifically for vehicular *ad hoc* networks (VANETs). The 802.11p WAVE standard specifies network protocols which address the difficulties associated with vehicular networks. These difficulties involve short link times, delay tolerance, and the inefficiency of wired-network paradigms that have been inherited into the wireless standards.

The following simulation experiments and algorithms were developed with particular scenarios in mind. These scenarios operate in a dual-carriageway environment with no junctions together with an obstacle or danger that is present at some point in the field. The simulations use both vehicular flow modelling and message propagation to advise the vehicles of the obstacle at a greater distance than line-of-sight provides. The results show that by spreading information quickly and efficiently through the network we can develop systems to reduce congestion and limit negative vehicular

flow effects. By coupling the telecommunications and driver model in one tool we can perform these simulations in the same runtime. This method of simulation can also be useful for intelligent transport systems, where simulations run in parallel can sweep the parameter space for a desired outcome. The impact of applications such as this is to reduce the time spent in congested vehicular flow and also to increase the safety interval and reduce collisions. Many possible applications of VANET technology are designed to improve safety of life while driving, and in this work we show that the application of our novel algorithms will realise this goal.

The paper is organised as follows. Section 2 shows some related work in VANET and ITS applications. Section 3 presents the simulation system we have used for this work and the models that have been incorporated. Section 4 presents the algorithms for the intelligent driver model and the lane-changing decision, and also the epidemic propagation modelling for dissemination of information. We then present the simulation scenario in Section 5. Section 6 shows the studies we have explored and the results on vehicular flow through the simulated system. The paper ends with conclusions in Section 7.

2. Related Work

There are many examples of work using vehicular network simulation to test the applications of wireless technology. Yin et al. [2] use simulation to evaluate the performance of dedicated short-range communication (DSRC) and Eichler et al. [3] analyse the impact of V2V messaging on vehicular flow. Torrent-Moreno [4] has produced a great deal of work on safety applications for vehicular communications networks, providing empirical comparison of results under different network settings and providing mechanisms for the delivery and dissemination of vital data.

There are tools available for performing coupled simulations of vehicular mobility and network communication including VGrid [5] which has been designed for use mainly with V2I systems, and MobiReal [6] which produces scenario data for accurate mobile *ad hoc* network (MANET) simulations, similar to the work of Wellnitz et al. [7]. These systems are useful but we found that the ability to take the best from both aspects and combine them in a highly configurable and extensible simulation system, was best met by the vehicular flow application of Treiber. The work of Chen et al. [8] is similar to ours for a different vehicular flow and accident scenario. We explore our data deeply so that we can validate against this established work.

3. Simulation System

The simulation tool we use is adapted from the dynamic vehicular flow simulator by Treiber et al. [9]. This tool uses a simple model of a two-lane roadway, but contains an advanced driver model and lane-changing algorithm, MOBIL [10]. We have added telecommunications functionality to the original simulator for V2V communication. In this section we show the underlying models that operate within the simulation system.

3.1. Vehicular Modelling. To accurately simulate vehicular behaviour, there are several key components: a driver model to develop how real people will drive under certain circumstances, a lane-changing model to make realistic decisions on when it would be advantageous to change lane and a roadway with rules (e.g., drive on the left in the UK).

A car following algorithm will contain at least a desired velocity, a safe time separation when following other vehicles and acceleration and braking criteria [9]. At each simulation time step the acceleration is calculated for each vehicle. The parameters of these models can be changed to emulate more aggressive and more considerate drivers, as required.

When modelling vehicular networks over large areas (such as metropolitan areas) the flow of vehicles on a single road behaves as an incompressible fluid according to $Q = \rho V$, where ρ is the average density of vehicles (cars/km) and V is the average velocity on the road (km/h) [11]. At microscopic levels of simulation (across any field size) each vehicle is treated as an individual entity, which greatly increases the computational requirements of the system but provides a more realistic and component-based approach to modelling.

The Intelligent Driver Model (IDM) in the simulator follows the MOBIL model [10] developed by Treiber. MOBIL operates as a car-following model such that the acceleration and braking are defined by the distance from the car in front. The function of such an acceleration dv/dt is shown in

$$\frac{dv}{dt} = a \left[1 - \left(\frac{v}{v_0} \right)^\delta - \left(\frac{s^*}{s} \right)^2 \right], \quad (1)$$

where

$$s^* = s_0 + \left(vT + \frac{v\Delta v}{2\sqrt{ab}} \right) \quad (2)$$

for acceleration on an open road a , velocity v , desired velocity v_0 , distance s to front vehicle, desired dynamic distance to front vehicle s^* , velocity difference Δv , a safe time delay between vehicles T , a comfortable braking value b , a minimum distance between vehicles s_0 and finally an exponent δ which is adjusted in order to mimic real vehicular flow patterns.

Lane changing algorithms add a necessary level of complexity to the IDM. In order to decide whether to change lane or not, the current acceleration must be calculated for the current lane and the acceleration in the new lane (with regard to the car behind and in front in the original lane). If the acceleration in the new lane is greater than that in the current lane, there is an advantage to be gained by changing lane. Many models, including those in the original simulator, include a bias in the model for particular lanes, which simulates the real scenario of slow lanes and fast lanes.

3.2. MAC Layer Protocol. In simulating wireless fidelity networks, the majority of simulations use the IEEE 802.11 protocols [12], as these offer the best simulation of the MAC layer functionality. The IEEE 802.11 MAC layer uses a distributed coordination function (DCF), which has been simulated in [13], to ensure efficient communication on the

medium, and implements controls to reduce collisions. More recently the IEEE 802.11p standard has been tested in [2] specifically for intervehicular communications. This allows the foundation of underlying strengths in the 802.11 suite to be enhanced for vehicular networks.

The MAC layer in the simulator operates using an adapted version of IEEE 802.11 which removes the inter-frame spacing (IFS) model, enabling equal priority for all network traffic. Due to implementation, and the need for simplicity in the model, our implementation of 802.11 does not suspend the backoff timer when the medium is busy during that frame, as 802.11 does.

The network backoff when the medium is busy X operates as shown in [11]

$$X \in 2^n \times [B_{\min}, B_{\max}], \quad (3)$$

where n is the number of times it has previously had to back off in succession. B_{\min} and B_{\max} are the minimum and maximum possible back off time, respectively. B_{\min} is often set at 0. The medium is defined as busy if any car within the transmitters interference range, R_i , is currently broadcasting.

Every car within the transmission range, represented by R_c , (which is usually twice as small as the interference range) will receive the message with probability λ .

3.3. Radiowave Propagation. In the modified simulator the reception of messages is performed by a basic algorithm controlled by the simulation engine. In advanced network simulators the realistic reception of messages depends on the signal strength at the receiver. The basic propagation model is the Friis free-space calculation, which extends the ideal free-space propagation formula ($\text{Pr} \propto 1/d^2$) [14] where Pr is received power and d is the distance from transmitter) to incorporate the antenna gain(both transmit and receive). The Friis model, however, will only hold true with a clear line-of-sight (LOS) between transmitter and receiver, and assumes no level of multipath or shadowing of the radiowave (which becomes very apparent in urban and highway scenarios [15]). Friis operates as shown in

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

where Pr is the received power at distance d with respect to the antenna gain and height and the system loss L .

One method of altering the propagation of messages through the simulated network, is to change the transmitted power and therefore the transmission range. By transmitting information further the message has a greater probability of retransmission in sparse networks and also a faster dissemination through the system [16]. The main disadvantage of having a large transmission range is that information propagates further, eventually reaching a point where the information usefulness is low, and so takes up bandwidth and time for almost redundant data. In a system with limited time and bandwidth this can cause localised problems, where more pertinent (i.e., geographically closer) information is lost to redundant data.

4. Algorithms

This section examines the algorithms used in the simulation. These algorithms form the basis for the work we present here and have been designed specifically for vehicle-to-vehicle and VANET scenarios. We have added to the original simulator, by Treiber et al., including telecommunications functionality and adapting both the IDM and lane-changing model, based on reactions to information received from an obstacle or danger. From this, we are able to show that more efficient information dissemination through a network can increase vehicular flow and also reduce stop-and-go formations (presented in our results in Section 6).

4.1. Epidemic Message Passing Algorithm. The propagation of messages through a system requires an efficient delivery algorithm. In our simulation we use a probabilistic information dissemination protocol, which is fully defined in Nekovee and Bogason's work [11]. We assume that all vehicles know their location (via GPS technology) and that each message contains information about location and time of creation. To ensure that propagation does not extend to unnecessary distances from the source each message has a Time-To-Live (TTL) setting.

The algorithm allows for a reasonable amount of retransmission and dissemination through the network and balances the relevance of the information with the distance from the source. To this end, information can disseminate quickly and efficiently and also reduce information spreading to vehicles which do not require the information (as discussed previously, this can cause more pertinent information to be lost).

The probability for rebroadcasting, P , as described by Nekovee [17] is obtained from

$$P = \begin{cases} 1 & \text{if } N_f \text{ or } N_b = 0, \\ 1 - \exp\left(-A \frac{|N_f - N_b|}{N_f + N_b}\right) & \text{otherwise,} \end{cases} \quad (5)$$

where N_f and N_b are the number of times the car has received that particular message from front and back, respectively, and A is a protocol parameter which controls redundant transmissions. In the case of directional message propagation, (5) is modified such that if a message is propagating in either direction it is only kept alive by nodes near the head or tail of the group. In this case the rebroadcasting probability, P , is computed from

$$P = \begin{cases} 1 & \text{if } N_k = 0, \\ 1 - \exp\left(-A \frac{N_k}{N_k + N_k^-}\right) & \text{otherwise,} \end{cases} \quad (6)$$

where N_k is the count of messages received from the direction of message propagation (e.g., if k is forward, N_k is the count of messages received from vehicles in front), and N_k^- is the count of messages received from the opposite direction.

4.2. Variable Speed Limit. The initial algorithm we introduce reduces the value of the desired velocity (v_0 in (5)) by a fixed amount when the vehicle has received the warning message. This achieves an overall slowdown in the network which can reduce the time delay between free-flow and gridlock (where $v = 0$) at the obstacle. This particular algorithm change has a transient effect on the network, such that the system will still become congested over time. The idea for this came from the London Orbital (M25) variable speed limit which operates on parts of the motorway.

The value we reduce v_0 by is of great importance. Initial tests showed that reducing the desired velocity too much (i.e., a reduction of over 10 m/s) when infected, the effect on the network was to cause congestion further back in the system, such that gridlock (i.e., all vehicles in the field of simulation are static) occurs much sooner. By reducing the value of v_0 by 2.7 m/s the network slowed well, and the time delay between free-flow and gridlock was increased without causing congestion further back in the network.

An important algorithmic change is the return to the normal value of v_0 once the obstacle has been passed geographically, otherwise the recovery from the obstacle will take a greater amount of time. We did test a proportional change in the desired velocity as the obstacle was approached, but this provided little observable effect at great distances and a highly-negative effect closer to the obstacle, as cars were congesting more smoothly but to a greater extent.

The results of this algorithm change were both transient and often detrimental to the overall velocity of vehicles in the system, so the change was dropped from the final algorithm described in Section 4.3. The reason for this negative effect is thought to be related to the size of field we are simulating. In future simulations we will explore a much larger field, where we expect the effect of the algorithm to become more pronounced as the distance from the obstacle increases.

4.3. Lane Changing Algorithm. The original lane change model, shown in (7), operates by determining an acceleration advantage (toward a goal velocity) to be gained by changing lane and then testing if a gain threshold is reached by the advantage. Early incarnations of our changes to the algorithm worked to forcibly increase the advantage if a message had been received and the vehicle was in the lane with the obstacle. This has some positive effects, but can cause problems when the message propagates a large distance back through the system. In the case of the message propagating beyond the reasonable extent of the need to change lane, this approach causes unnecessary congestion in the adjacent lane to the obstacle resulting in total congestion in a short time.

At the start of the simulation several static variables are applied to the model. A changing threshold is applied that indicates the increased acceleration the lane change will yield; this is set by default at 0.3 ms^{-2} for cars in the field. The other value is the ‘politeness factor’ which reduces the overall calculated advantage and which simulates the care drivers take when changing lane (i.e., the model may say it is advantageous to change lane, but the driver may be too polite or hesitant to do so).

The basic algorithm operates by calculating a value of advantage (A), the disadvantage this causes to other (B) and then calculates whether a function of these values reaches a changing threshold. If the threshold is reached the vehicle changes lane:

$$A = a_n - a_o + B, \quad (7)$$

where a_n is the acceleration in the new lane and a_o is the acceleration in the old lane. B refers to a weighting to keep the vehicle in the slow lane, as operates in reality:

$$B = a_{b(o)} - a_{b(n)}, \quad (8)$$

where $a_{b(o)}$ is the acceleration of the car behind in the old lane if the vehicles changes lane and $a_{b(n)}$ is the acceleration of the car behind in the new lane if I change lane. These values are then entered into the following equation to return true or false to changing lane:

$$(A - p) * B > T, \quad (9)$$

where p is the politeness factor and T is the changing threshold. The form of (9) is multiplicative so the values of A and B have a significant impact on each other. In the following equations, these values are only calculated if the vehicle has been infected with a message. If a vehicle is ignorant, it will continue to use the algorithm in (9). The initial change we made was to add a value to A in (10) as such

$$((A + V) - p) * B > T. \quad (10)$$

This is very much a brute force approach and as such does not truly represent real driving in a highway, where the value of V would increase as the vehicle approaches the obstacle and drop to zero after the obstacle has been passed. This proportional addition to A is shown in

$$\tau = \frac{P_o}{(P_m - P_o)}, \quad (11)$$

$$((A + \tau) - p) * B > T.$$

In this adaption of the original algorithm, the value τ is calculated as a function of the location of the obstacle and the vehicle’s distance to it. This value is capped at a maximum (currently 20) to prevent unrealistic behaviour (i.e., cutting in with zero safety headway), which means the effect is noticeable but quite subtle, when compared to the brute force method in (10). This means that as the vehicle approaches the obstacle the incentive to change lane becomes greater, reducing the appearance of congestion at the obstacle in the same lane.

5. Simulation Scenario

In order to deliver usable and realistic results, the simulated system must be an accurate representation of a real highway. The models for vehicular flow and network transmission have been introduced in previous sections, and the metrics

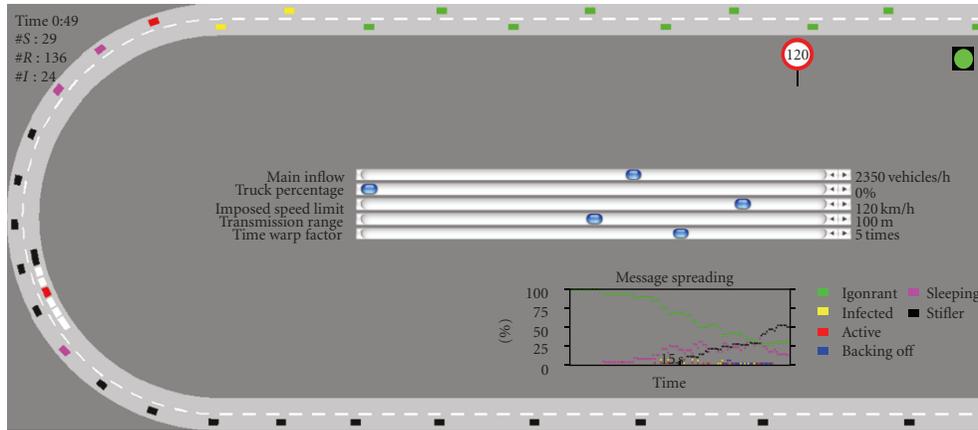


FIGURE 1: A graphic showing the simulator, with the obstacle in the lower quadrant of the curve to the left.

gathered with the results obtained are shown in Section 6. The system we used for our study is shown in Figure 1. The road length is approx. 1400 m and has two lanes with vehicles moving in the same direction and an obstacle at 720 m, in the left-hand lane. Vehicles originate in the top-right of the road and travel around to the departure point in the bottom-right. The flow of vehicles entering the road can be varied from 0 to 4000 vehicles per hour, split across the two lanes. The desired speed (v_0 in (1)) can be varied between 20 and 140 km/h, which will directly affect the vehicular density as greater velocity will require greater separation, according to the IDM.

The transmission of data between vehicles is controlled by them operating in five distinct states: ignorant, where the vehicle has received no data and is not transmitting, informed, when the vehicle receives data from the network, transmitting, to resend received data through the network, backing off, when the vehicle wishes to transmit but the medium is busy, sleeping, when the vehicle has forwarded data but is waiting to make sure all neighbours have received that data and stifling, when the vehicle has performed all the reception/transmission required and therefore is no longer active. The movement through these states, from ignorant to stifler, can be seen as a loop. The propagation methods described in Section 6.3 control how the vehicle progresses through the states in different ways. The transmission range of the vehicles is also set globally for the simulation, between 0 and 200 m (it is possible to extend the transmission range of vehicular networks beyond 200 m but only by using very high transmission power and assuming little interference). The vehicle that is causing the obstacle (surrounded by several car-lengths of safety barriers) constantly transmits without moving to the sleeping or stifler states, so that any vehicle coming within range will receive the data, and then begin the process of receive-and-forward.

In the next section we present the findings of our work.

6. Simulation Studies

In this section the collection of metrics is discussed with diagrams and presentation of the results. All the simulations

were run with varying settings, so that an appreciation of all the situations that may occur (that we can control in the simulator) can be established. It is important to note, as previously mentioned, that in the initial studies we assumed that all vehicles in the simulation are equipped with the technology for message propagation. In reality this market penetration will take many years to achieve, but many car manufacturers are working on supplying this technology soon [18].

In later studies, as shown in Section 6.5, we have explored different rates of equipped vehicles in the simulation. This requires that all parameters remain static whilst testing the equipped vehicle rates, such that we can rule out any influence from transmission range, vehicular density and vehicle velocity. For each run at different parameters we tested equipped vehicle rates from 0% to 100% in increments of 10% and performed 20 runs to average out the data and remove any transient or artifactual data.

6.1. Velocity Experiments. These experiments test the effectiveness of the algorithm as vehicle velocity changes, to see if the algorithms are suited to an urban (slow) or highway (fast) environment. When ignorant (no messages have been received), the cars will still attempt to change lane to avoid the obstacle, but only as part of the original lane-changing algorithm, and so congestion builds up in a short amount of time, for most simulations. Below a certain network load the road will never become congested, so the vehicular load of the experiment was varied for each experiment. The vehicular load on the road was also set low enough so that the algorithm could affect the flow of cars as, at high loads, this would not be possible. To this end there is, in any system, a critical value of vehicular load after which no action can prevent or reduce congestion.

Some early simulations with low vehicular loads showed that it is sometimes more efficient to be ignorant of the obstacle, and this must be taken into account, as in this case the best course of action is to drive normally, using the normal algorithm.

Figures 2 and 3 show results from a sample of the experiments we ran to test this theory. By varying all

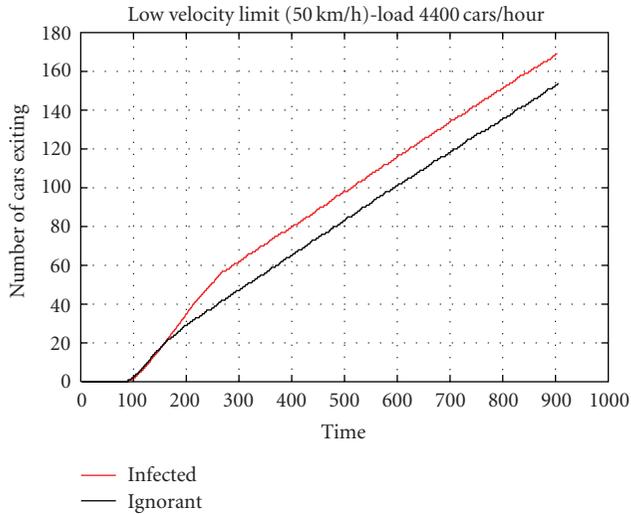


FIGURE 2: Comparison of exit aggregate for infected and ignorant simulations at urban velocities (below 50 km/h).

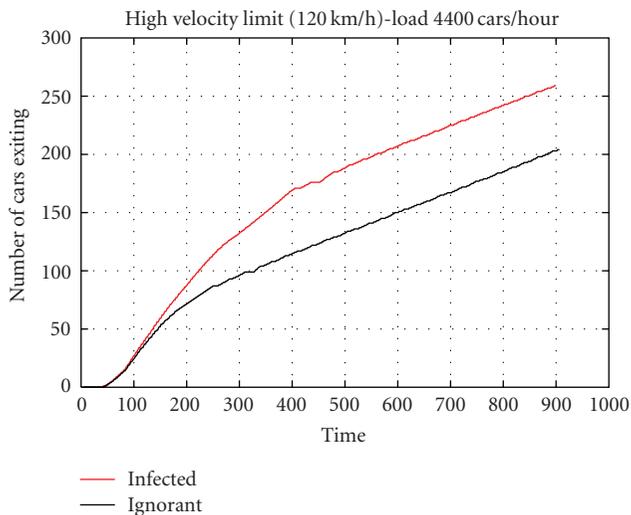


FIGURE 3: Comparison of exit aggregate for infected and ignorant simulations for motorway velocities.

the parameters available we found that certain velocities, transmission ranges, and vehicular load had different levels of effectiveness to the overall congestion in the system. In the main the results showed that our algorithm always produced a positive effect. In order to produce valid data we ran several simulations with the same parameters and then averaged this data, to show results for a typical case.

Figures 2 and 3 show the number of cars exiting the field in a simulation run as an aggregate over time. Both show an advantage for infected cars using the advanced lane-changing algorithm, but the advantage is greater at higher velocities, where the vehicles have more distance between them for the same vehicular load, meaning they can more easily change lane.

We found, and Figures 2 and 3 corroborate this, that the advanced algorithm increased the time before congestion

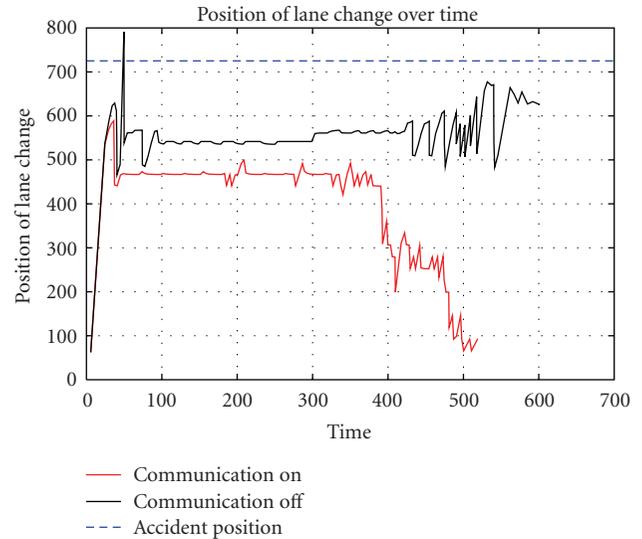


FIGURE 4: Chart showing location on the field of lane changes with and without communication.

began to build up and then once congested, the infected cars still moved through the system more efficiently. By running the simulation for 15 minutes we can monitor the development of congestion in the system and how the flow is affected by the changed algorithm.

The results show interesting behaviour, beyond the reduction of congestion in the system. By analysing the first 5–7 minutes of simulation time, we can see that the development of congestion is also slower once vehicles do start to slow down. This is because of the algorithm moving vehicles into the opposite lane to the obstacle, reducing the load on the lane with the obstacle and therefore reducing the number of stopped vehicles behind the obstacle which, when changing lane, cause a dramatic slowdown in the new lane. This reduction in stop-and-go vehicular formation is also seen elsewhere in the field when the cars are infected with the warning message and switch to our adapted algorithm.

6.2. Position of Lane Change. A factor that affects the build-up of congestion in the system is related to the location of the lane change. The following results show where the lane change occurs with no communication and then using our enhanced lane change algorithm with communication active. The simulation settings were set at 2200 cars/hour load, speed limit of 120 km/h and a transmission range of 100 m.

During the simulation the message propagates backwards towards position 0 and, with the advanced algorithm, the location of the lane-change also reduces. When the system starts to slow and vehicular density increases, the lane-change moves right back, causing a slower build-up of congestion and a greater amount of free vehicular flow, as shown by Figure 4. This does place greater load on the opposite lane to the obstacle, but the reduction of stop-and-go behaviour negates this. In Figure 4 the rapid reduction in position of most lane changes between 395–405 seconds and again at

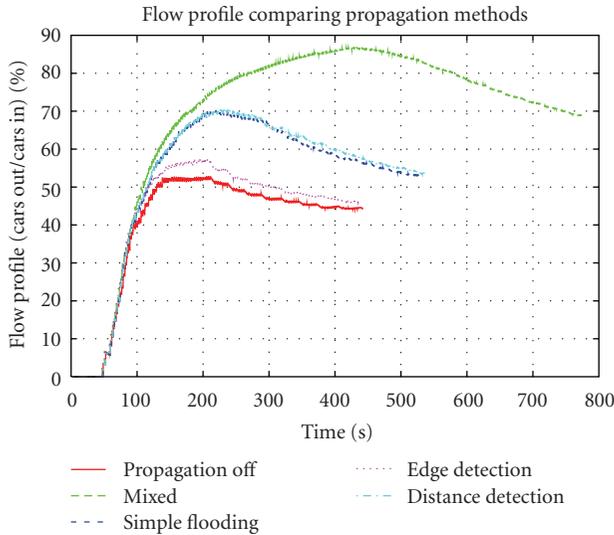


FIGURE 5: Comparison of message propagation methodologies.

475–500 seconds represents a period when the propagation of the message is continuous, and the periods of little change (of lane change position) are due to reduced propagation of the warning message. The initial peak of lane change position between 0 and 35 seconds represents the initialisation of the system, that cars can change lane very close to the obstacle due to the road being less loaded. The data to produce this result came from a single typical simulation, with parameters set as per the previous paragraph.

6.3. Transmission Method Experiments. In this experiment the available transmission methods are tested with the new algorithm, to see how they affect the overall congestion in the system. Figure 5 shows a simulation run until congestion is present at the origin of the field (i.e., position = 0). The simulation is stopped when the congestion reaches the origin as after this point the algorithm is not affecting the vehicular flow. The values represent the proportion of cars leaving the field in relation to the cars entering (Cars Exiting/Cars Arriving). This indicates how vehicles are flowing through the system, where an increase of the gradient represents free flow and a decrease represents congestion.

The models shown in Figure 5 are simple flooding, where the message is rebroadcast just once, edge detection which is explained in Section 3.1, and distance detection, which is a different probabilistic method and a mixture of edge and distance detection. These models are all running in the simulation, but the mixed (edge detection and distance detection) offers the best simulation of a real epidemic protocol.

As can be seen in Figure 5 the edge detection method alone offers little improvement over no propagation, and the simple flooding and distance detection methods offer a good initial advantage (0–200 seconds) but then suffer very fast congestion build-up. The mixture of edge and distance detection algorithm, with our changes to the lane change model offers excellent results keeping near free flow until

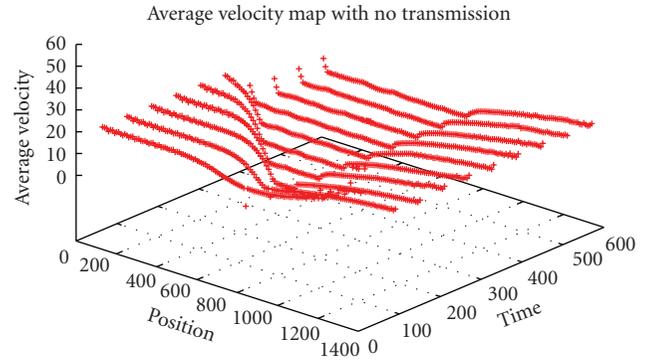


FIGURE 6: Average velocities across the simulated highway without transmission.

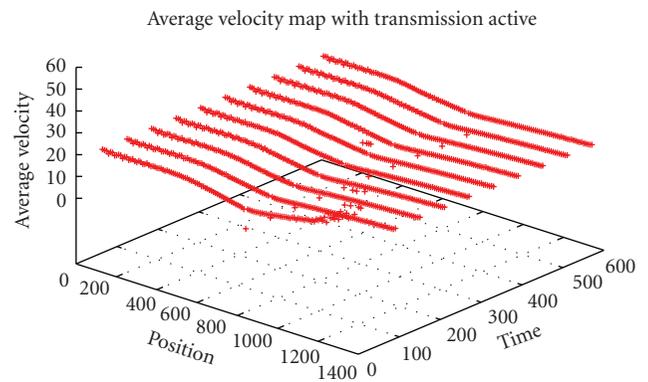


FIGURE 7: Average velocities across the simulated highway with active transmission

ca. 420 seconds, when the network then slows and starts to congest, but this takes longer (ca. 250 seconds from the first slowdown) than the other algorithms.

The addition of this propagation method and the changed algorithm prevent several congestion-causing situations to occur. The main situation avoided is when vehicles are unable to change lane to avoid the obstacle and begin to slow down, but then do change lane causing the cars behind to slow. This cause has been seen to initiate the build-up of congestion by earlier warning of the obstacle so that the cars can change lane at a high velocity. Another behaviour of the system is that when the congestion is initiated, the cars will fill up behind the obstacle, unable to change lane. With the adapted algorithm the extra incentive to change lane means that the opposite lane fills first and so vehicles can still move, increasing the time before the whole system becomes congested.

6.4. System Velocity. The average velocity through a system is of great importance. If a higher average velocity can be achieved the number of vehicles passing through the area of congestion will be higher than if there is much slowing of vehicles. Figures 6 and 7 show the average velocity calculated for intervals of 10 metres on the x axis and an interval of 30 seconds across the y axis. Each point represents the average velocity at that time/position interval. In both figures there

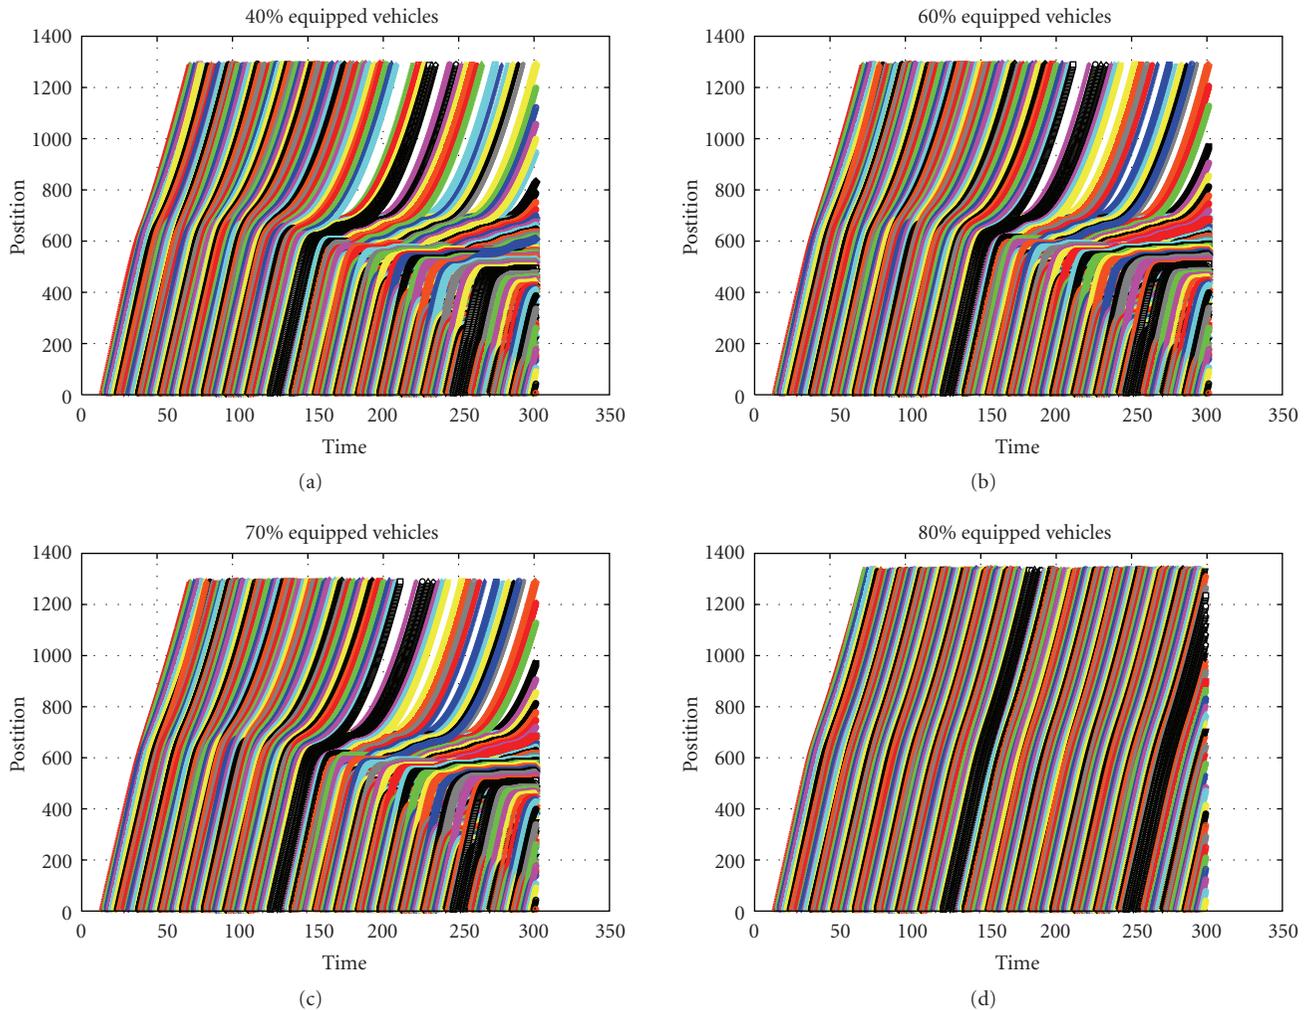


FIGURE 8: Vehicle trajectories through the system for (a) 40% equipped vehicles, where we see little change from 0% equipped vehicles, (b) 60% where some improvement is shown but congestion still builds up, (c) 70% where the improvement is noticeable and (d) 80% where the trajectory is only slightly affected by the obstacle. Each coloured trajectory represents a single vehicle as it enters and leaves the simulated highway.

is a noticeable slowdown as the vehicles pass the obstacle. This can be accounted for by the IDM attempting to retain a minimum safe distance between vehicles.

Figure 6 shows that after an initial even velocity through the field, congestion begins to build up at the position of the obstacle between 100 and 200 seconds, which causes a slowdown further back to position 0. By time of 330 seconds the congestion has reached position 0, and the average velocity has fallen from $20\text{--}25\text{ ms}^{-1}$ to $0\text{--}5\text{ ms}^{-1}$.

As can be seen in Figure 7, there is a uniform average velocity before and after the obstacle during the whole period of the simulation (10 minutes). This reinforces the results from the other simulations and proves that there is a better flow of vehicles through the network, as well as a reduction in the build-up of congestion, when effective transmission of the road condition occurs.

We note that for both figures there is a spike in velocities (between time 0–10 and position 1200–1400). This is the period when the first cars are leaving the field. As they have

no cars in front they can accelerate up to the full speed limit unhindered. To remove this artefact future simulations will have a “warm-up” period, with low vehicular load, that initialises the field.

6.5. Varying Penetration Rate of Radio-Equipped Vehicles.

In all of the simulation studies and results presented in Section 5, we have assumed that 100% of the vehicles in the simulated field are radio equipped and therefore able to receive and retransmit (where necessary) the data about the upcoming incident or congestion. This is a somewhat idealised view of how deep uptake of radio-equipped vehicles will be. More likely, as vehicle manufacturers and radio equipment providers begin releasing the products to operate and support a VANET, the ratio of equipped vehicles will increase. Due to the critical safety-of-life applications of VANET, it is important to study the efficiency and resilience of simulated applications, to study the point at which they become usable and reliable. This simulation study involved

the same system as in previous experiments, but here the proportion of equipped vehicles is increased from 10% to 90%. The plots in Figure 8 show the flow trajectories of vehicles in the system.

We ran each simulation several times, to ensure no artifactual or erroneous data was included in our results, but the results shown here represent only one simulation. As can be seen at low penetration rates (10%–40%) the vehicles reach the obstruction and begin to congest, but between 60% and 70% there is a noticeable improvement, or smoothing, of the vehicles velocity through the field of study. The figures show that above 70%, where there is still some congestion build-up around the obstruction, the propagation of information to vehicles further away, and the change to the lane-changing algorithm this triggers, allows for more steady driving through the field of study, and less congestion.

7. Conclusions

This paper has presented a simulation of a specific road condition, that of an accident blocking a lane on a dual carriageway. The simulation uses the coupled approach to mobility modelling and network simulation in a single process. The models and algorithms that represent the vehicular flow and network traffic have been implemented according to well-known and standard models. We have shown that adapting the algorithms when information about the accident is received via wireless transmission, we can reduce the build-up of congestion and increase the flow of vehicles passing the obstacle, both in terms of quantity of cars and average velocity.

The various algorithms we tested achieved overall improvement in the majority of cases. In some simulations the improvement was not only in the prevention of congestion overall, but by also keeping a consistent average velocity through the network, which helps to reduce the effects of stop-and-go vehicular flow and smooth possible congestion “waves” that emanate from the source backwards. The driver model and lane-changing algorithms come from well validated sources, and so the adaptations we have made are highly realistic and can show the effect of even simple changes (as in the brute force addition of a value to A). In order to fully test our algorithms we ran numerous tests with a wide variety of parameters, to test for any transient or artifactual effects. From these repetitive runs we established that most effects were long-lasting and that where those were transient, this was only due to a vehicular load on the road where congestion could no be prevented (as corroborated by control tests).

In the common situation where the obstacle is temporary (i.e., a vehicle malfunction), any reduction in congestion build-up allows for the obstacle to be removed, before the velocity of all vehicles behind the obstacle drops to zero. In more complex roadways the advanced warning could also lead to a change of route, so drivers can avoid the section of road where the obstacle exists.

The real-world application of this algorithm is limited with current technology, as the recommendation for lane change can be ignored by a human driver. The current possi-

bility for this algorithm is for an advanced driver information system, where the details of the obstruction are suggested and displayed in-car. The modelling of penetration rates that we performed (Section 6.5) also studies the obedience of human drivers. In the future, when automated driving systems are in place and more relied on, this algorithm would form a part of the plethora of functions and processes used to drive a vehicle under computer control.

Büscher et al. [19] discuss the challenges of systems such as ours, from a social perspective. The authors work motivates the need for accurate human-behaviour modelling in vehicular network simulation and what impact the social aspects of mobility provide. The simulation system we designed to perform our studies uses accurate mobility modelling, but for a real implementation the social aspects of the driver and the context in which they are flowing becomes essential.

In the size of field simulated here we are between the microscopic and macroscopic scale of simulation, which is achieved seamlessly by the use of a coupled model of simulation. The tool we performed the simulations with was lightweight and so we could easily implement new algorithms and protocols. In order to run more complex networks we would require a more complex simulation engine for vehicle flow. With this increase in complexity the field size will increase and therefore a more powerful network simulator is required. The proof of concept that this paper provides will lead to further work in this field, including the use of parallel and distributed computational resource.

Acknowledgments

M. Nekovee acknowledges the Royal Society for supporting his work through an Industry Fellowship. T. Hewer acknowledges the support of British Telecom as part of a CASE Industrial PhD and funding from the UK Engineering and Physical Sciences Research Council.

References

- [1] M. Nekovee, “Modeling the spread of worm epidemics in vehicular ad hoc networks,” in *Proceedings of the 63rd IEEE Vehicular Technology Conference (VTC '06)*, vol. 2, pp. 841–845, May 2006.
- [2] J. Yin, T. Elbatt, G. Yeung et al., “Performance evaluation of safety applications over DSRC vehicular ad hoc networks,” in *Proceedings of the 1st ACM International Workshop on Vehicular Ad Hoc Networks (VANET '04)*, pp. 1–9, October 2004.
- [3] S. Eichler, B. Ostermaier, C. Schroth, and T. Kosch, “Simulation of car-to-car messaging: analyzing the impact on road traffic,” in *Proceedings of the 13th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS '05)*, pp. 507–510, September 2005.
- [4] M. Torrent-Moreno, *Inter vehicle communications, achieving safety in a distributed wireless environment: challenges, systems and protocols*, Ph.D. thesis, Universität Karlsruhe, 2007.
- [5] J. Anda, J. LeBrun, D. Ghosal, C.-N. Chuah, and M. Zhang, “VGrid: vehicular ad hoc networking and computing grid for

- intelligent traffic control,” Tech. Rep., University of California at Davis, 2005.
- [6] K. Maeda, T. Umedu, H. Yamaguchi, K. Yasumoto, and T. Higashino, “MobiREAL: scenario generation and toolset for MANET simulation with realistic node mobility,” in *Proceedings of the 7th International Conference on Mobile Data Management (MDM '06)*, p. 55, May 2006.
 - [7] O. Wellnitz, S. Lahde, E. Eden, and L. Wolf, “A scenario editor for mobile ad hoc networks,” Tech. Rep., IBR, Technische Universit at Braunschweig, 2006.
 - [8] A. Chen, B. Khorashadi, C.-N. Chuah, D. Ghosal, and M. Zhang, “Smoothing vehicular traffic flow using vehicular-based ad hoc networking & computing grid (VGrid),” in *Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC '06)*, pp. 349–354, September 2006.
 - [9] M. Treiber, A. Hennecke, and D. Helbing, “Congested traffic states in empirical observations and microscopic simulations,” *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, 2000.
 - [10] A. Kesting, M. Treiber, and D. Helbing, “General lane-changing model MOBIL for car-following models,” *Transportation Research Record*, vol. 1999, pp. 86–94, 2007.
 - [11] M. Nekovee and B. B. Bogason, “Reliable and efficient information dissemination in intermittently connected vehicular adhoc networks,” in *Proceedings of the 65th IEEE Vehicular Technology Conference (VTC '07)*, pp. 2486–2490, April 2007.
 - [12] IEEE, “IEEE standard for information technology-telecommunications and information exchange between systems-local and metropolitan area networks-specific requirements—part 11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications,” IEEE Std 802.11-2007, (Revision of IEEE Std 802.11-1999), pp. C1–1184, June 2007.
 - [13] J. Weinmiller, H. Woesner, and A. Wolisz, “Analyzing and improving the IEEE 802.11-MAC protocol for wireless LANs,” in *Proceedings of the 4th IEEE International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS '96)*, pp. 200–206, 1996.
 - [14] R. Struzak, “Radio-wave propagation basics,” Tech. Rep., ICTP-ITU-URSI School on Wireless Networking for Development, 2006.
 - [15] K. Laasonen, “Radio propagation modeling,” Tech. Rep., University of Helsinki, 2003.
 - [16] M. Torrent-Moreno, “Inter-vehicle communications: assessing information dissemination under safety constraints,” in *Proceedings of the 4th Annual Conference on Wireless on Demand Network Systems and Services (WONS '07)*, pp. 59–64, January 2007.
 - [17] M. Nekovee, “Epidemic algorithms for reliable and efficient information dissemination in vehicular ad hoc networks,” *IET Intelligent Transport Systems*, vol. 3, no. 2, pp. 104–110, 2009.
 - [18] D. Waters, “Connected cars ‘promise safer roads,” BBC News, 2007.
 - [19] M. Büscher, P. Coulton, C. Efstratiou, et al., “Intelligent mobility systems: some socio-technical challenges and opportunities,” in *Communications Infrastructure: Systems and Applications in Europe*, R. Mehmood, E. Cerqueira, R. Piesiewicz, and I. Chlamtac, Eds., pp. 140–152, 2009.