The University of Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

Cellular Automata Model for Traffic Flow

Adam Howlett

4227558

University of Nottingham Physics & Astronomy

25th December 2017

Contents

Abstract

This document investigates a Monte Carlo simulation approach to a simple 'Cellular Automata' boolean traffic flow model. The model follows strict rules as laid out by the 1992 paper *'A Cellular Automaton Model For Freeway Traffic'* by Kai Nagel and Michael Schreckenberg, which successfully demonstrated a transition from laminar traffic flow to start-stop waves with increasing vehicle density. This report will, using the same model, demonstrate similar results to Nagel and Schreckenberg and also analyse the variation of certain parameters of the model and their relative effects on the flow of traffic along the road.

Introduction

It is becoming increasingly important to optimise traffic flow on the roads of the world. Minimising the amount of fuel used, a consequence thereof, would have not only economic but also profound environmental benefits, reducing the amount of money spent on fuel and reducing the amount of greenhouse gas emission into the atmosphere. One method of studying traffic flow in order to do this is the use of computer aided simulations. Results obtained from these simulations serve to increase our understanding of the ways traffic behaves under certain situations, as well as give insight into ways this might be implemented in the real world.

The Boolean model used in this project provides a valid model for one lane traffic flow on a road, with predictive capabilities. Based in nonlinear dynamics, the Monte Carlo simulation will be used to converge on several key results that demonstrate its usefulness, as well as demonstrate a phase transition from laminar flow to start-stop waves, analogous to behaviour in fluid dynamics. This is comparable to what happens in real traffic, as well as the maximisation of a traffic flow metric for a critical density point[1].

Method

The model upon which this paper will be based is a simple cellular model – The lane is represented by an array of Length L (the road length), with each cell either being occupied (indicating the presence of a vehicle) or not. The system starts with N cars at velocity v_i . The spacing will typically range between cell 1 and L, rounded to the nearest integer with a linspace function, although upper limits of 0.7L and 0.8L have also been used to generate some graphs.

For the purposes of this report, v_i will be taken exclusively to be 1, and will range in integer values between 0 and a given 'speed limit', v_{max}.

The system is then allowed to evolve over an integer number of steps, T. For each step, the fundamental rules that dictate the evolution of the road are, in order, as follows:

1. **Acceleration:** If the velocity v of a vehicle is lower than a given speed v_{max} and the distance to the next car ahead is larger than v+1, the vehicle will accelerate by a speed of 1
2. **Deceleration (to avoid collision):** If a vehicle in position i sees the vehicle in front of it at a distance j (& j ≤ v), it decelerates to speed j-1
3. **Deceleration (random):** With a given probability, p (to be varied), a given vehicle will decelerate by a speed of 1
4. **Car Motion:** Each vehicle will advance by (its respective) v number of cells

Other conditions will later be introduced, such as adding conditions for particularly slow drivers (or speeding drivers) by reducing/increasing v_{max} for a certain percentage of the cars on the road. Periodic boundary conditions are also applied, such that if a vehicle at site I has a speed v that would take it past the length of the road, it will instead advance to site i-L+v, 'wrapping' it back around to the start of the road. The main method of data representation of the road is in terms of a space time profile of the road, with t increasing downwards on the vertical axis. Hence a typical profile for a single car would look as in figure 1.

The car accelerates with each timestep, as indicated by the changing colours and increasing gaps initially between its position at t and t+1, until it reaches v_{max} (6 in this case) at which case it will continue to move at v=6, wrapping each time it reaches the end of the road, until t=T.

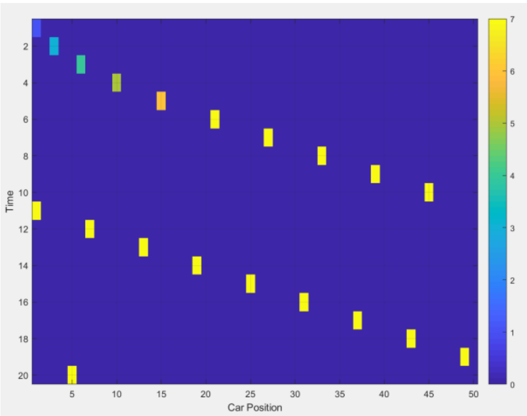

Figure 1: Typical space time profile for 1 car (p=0)

(Sidenote: The velocities on the colour bar are actually 1 value lower than displayed. This is because an actual value of 0 in matlab's 'imagesc' function would display as the background colour, resulting in confusing visuals)

Theory

The first quantity we wish to define in respect to the variables of the model is the density of the cars on the road, ρ, given simply by:

$$\rho = \frac{N}{L}$$

This variable is a constant for the road, as no cars will leave it due to the wrapping condition, and its length does not change. The density will turn out to be a critical variable when attempting to optimise the flow of traffic on the road, as we will see.

The flow will be taken to be the total average speed of all the cars on the road over the time length, T, and thus will be sum over all t of the flow at a given time:

$$Q = \frac{\sum_j v_j}{T},$$

Experimental data yields a surprisingly specific relationship between flow Q and density ρ, as shown in figure 2. There is a distinct change in the gradient dQ/dρ, at around the value of ρ=25%, where the flow reaches a maximum. This is the so-called 'critical density' at which the traffic is at its most optimal. Nagel and Schreckenberg reproduced this relationship in their paper, as shown by figure 3.

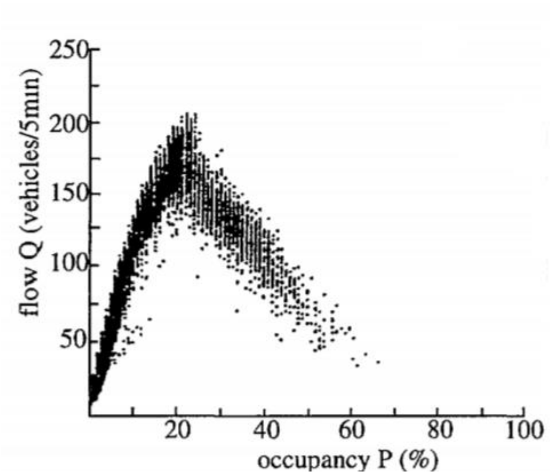

Figure 2: Experimental data of flow Q vs Occupancy taken by Japan Highway Public Corporation[2]

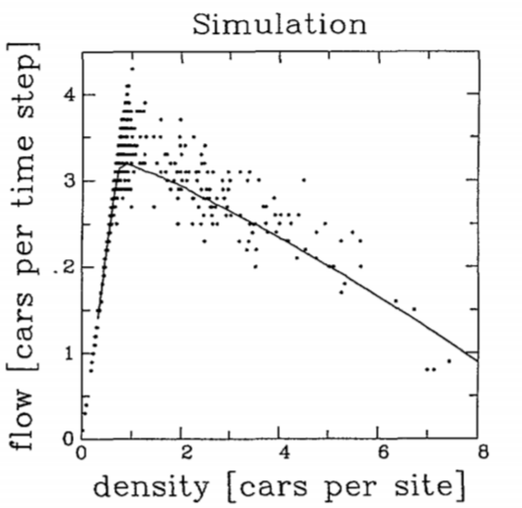

Figure 3: Simulation results by Nagel and Schreckenberg[1]

Simulation Results & Analysis

Initially, the slowdown probability p is taken to be 0. Without these random effects, we can still observe an interesting relationship between some of the core variables, as demonstrated by Nagel and Shreckenberg. If v_{max} is low enough, the motion will continue unhindered. However, if a relationship between the density and v_{max} is met the start/stop waves begin to form. This relationship comes from v_{max} attempting to define the spacing between the cars. If this is not allowed because there are too many cars to have a spacing of v_{max}, the cars that accelerate to v_{max} first will eventually reach cars that have not yet accelerated, and the start/stop waves will begin to generate, with a new, smaller spacing equal to the speed of the cars in the wave. Figures 4 and5 demonstrate how the increase in density affects the transition from laminar flow to stop/start waves, for $v_{max} = 5$, T=100 and L=10

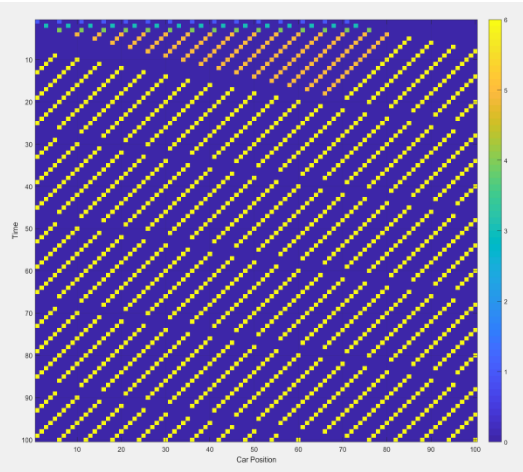

Figure 4: Laminar flow for ρ=0.15

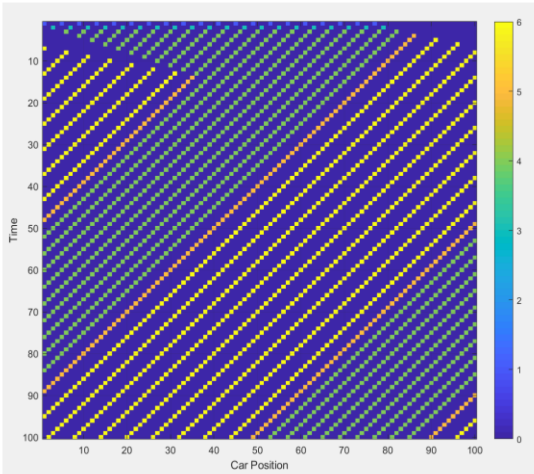

Figure 5: Start/stop waves forming at ρ=0.2

The transition to a steady solution here is not affected by the length of the road, but the time it takes the vehicles to initially accelerate. As long as they are sufficiently spaced, this is not an issue and predictable patterns typically emerge within the t=10 range. An increase in v_{max} will increase the width of the stop wave, up to a point. This is also relatively intuitive, though, as the vehicles leaving the wave will accelerate more quickly leaving it but also return to the back more quickly. After a certain point, however, the vehicles do not have enough space to reach v_{max},and the waves will reach a maximum width. The wave group travels with a speed of v=-1 in all cases. Interestingly, flow is maximised here for the situation in which v_{max} is such that the start stop waves have *just* begun to form. This is likely due to the emergence of them quickly outweighing the positive contribution of simply adding more cars to the road. However, this is not a particularly realistic example of how drivers actually perform on the road. While attempts to regulate spacing are made, human error amongst other issues consistently result in effectively random decelerations of specific cars. To this end, we allow p to take values between 0 and 1 for each of the cases in Figure 4 and 5. The results are shown in Figure 6, with the same densities, but over a larger time range and road length to better capture the resulting wave.

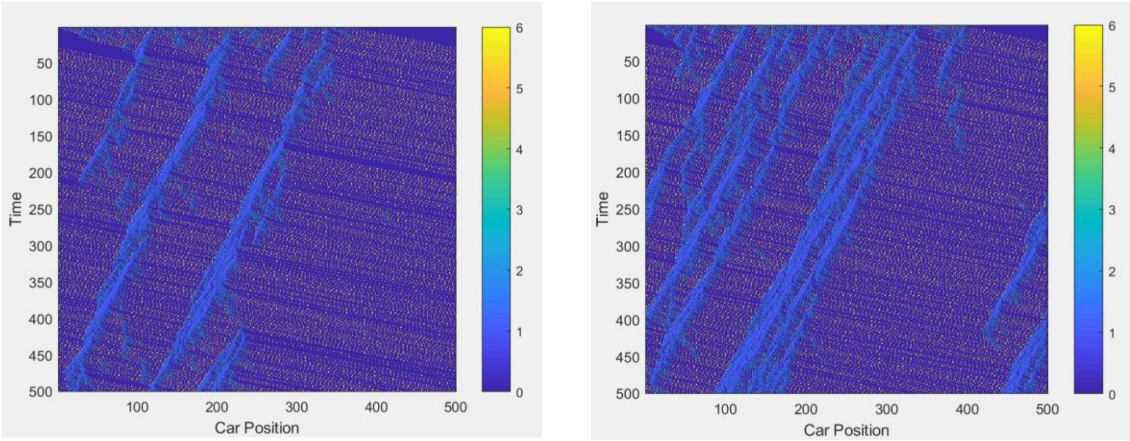

Figure 6: Jams forming for nonzero p at ρ=0.15 and ρ=0.2 respectively.

The waves appear much more sporadic and randomised, due to the probabilistic nature of their formation. For higher densities (or indeed, v_{max}, again up to a certain threshold), the number of these jams is extended. For higher probabilities, the width of the jams is increased, leading to less efficient travel and a lower net flow, as shown in Figure 7:

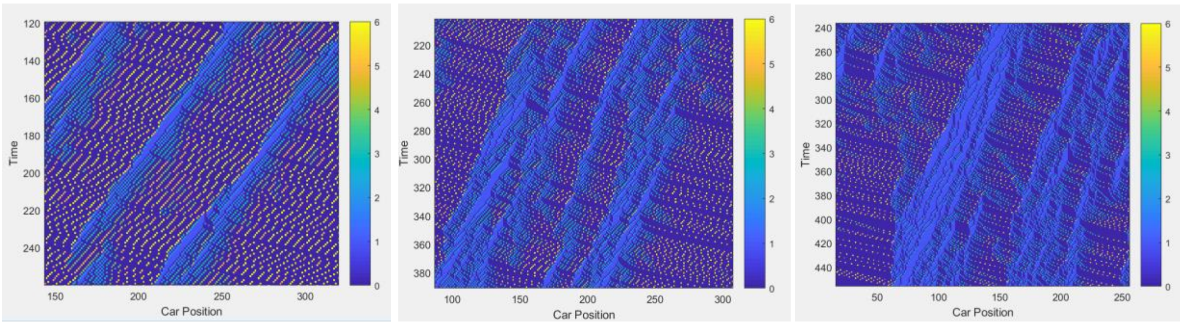

Figure 7: ρ=0.2 jams for p=0.1, p=0.3 and p=0.5

The formation of these jams is a clear inefficiency issue, and the resulting flow is reduced from the ideal value for these starting conditions (average speed of 5) to about 2.5 for the p=0.5 case.

Next, the flow-density relationship is observed, and similar results to Nagel & Shreckenberg are obtained. By running the program over many density values between 0 and 1 with the same v_{max} of 5 and (initially) a value of p=0, a graph of this nature is obtained (Figure 8):

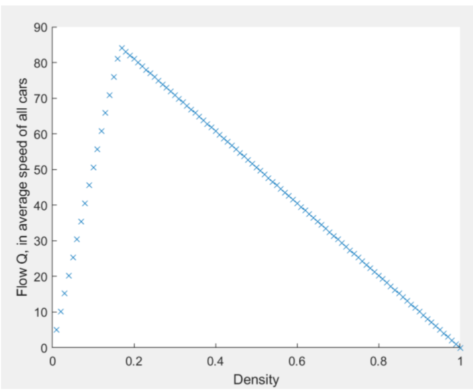

Figure 8: Density/Flow relationship for p=0, T=100 and L=100

As expected, the correct relationship is observed. Optimal traffic flow occurs at a density of around 0.18. After a density of around this value is reached, a saturation point of cars occurs and the transition to start/stop waves is made. The emergence of these waves quickly outweighs the contribution of more cars on the road in terms of the flow rate, and an effectively linear relationship between an increasing density and flow Q is observed.

A more unpredictable scatter is observed for nonzero values of p. However, these are shown to somewhat converge to their corresponding relationships if the average over a large number of samples are taken. This is demonstrated below (Figure 9).

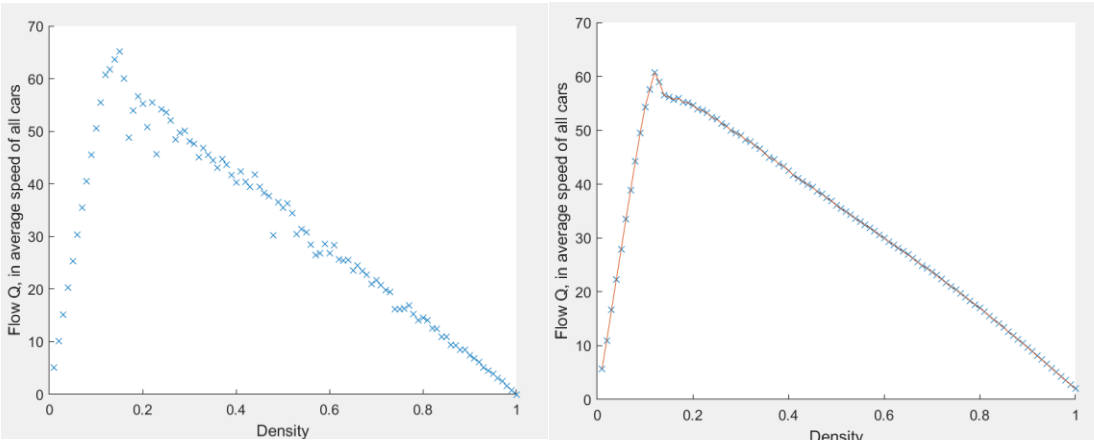

Figure 9: Density/Flow relationship for p=0.2, with T=50, L=100. The left shows 1 sample of data and the right demonstrates an average over 100 samples.

There are a few things to note here. The first is that the behaviour around the peak is decidedly different from that of the p=0 case, although the general shape is still converging to the correct result. The beginning of the start/stop waves forming initially behaves differently to the p=0 case due to the introduction of the random element generating more waves than would usually expected, resulting in a marked drop in flow rate initially. Another thing to note is the flow rate for every point has been reduced, as well as the peak moving slightly to the left. Both of these are also explained by the overall reduction in flow rate and thicker nature of the stop waves that are generated. We would expect to see the last two observations amplified for increasing p, and indeed this is demonstrated in Figure 10.

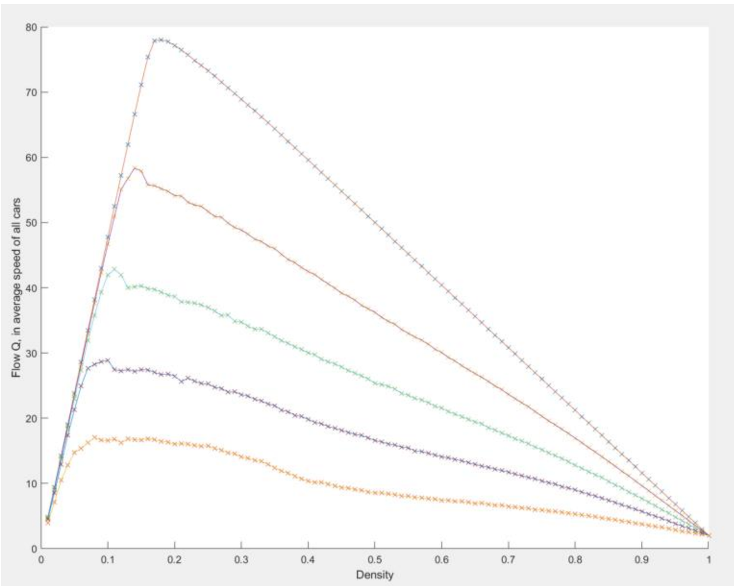

Figure 9: Density/Flow relationship for p=0, 0.2, 0.4, 0.6 & 0.8, with T=50, L=100. All are averaged over 100 samples.

Next, we should investigate what happens if we allow v_{max} to vary. We again run the program over many samples, fixing this time the probability p at 0.2.

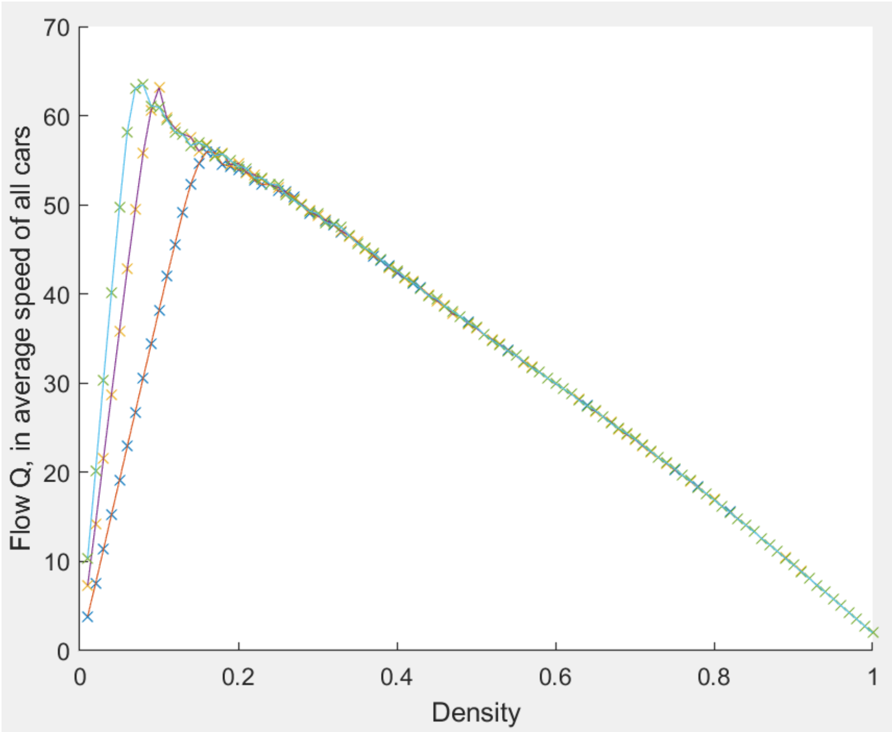

Figure 10: Density/Flow relationship for p=0.2, with T=50, L=100. The graphs correspond to (in increasing order of peak height) v_{max}=4, 8 & 12

The behaviour here is also interesting. For an increasing v_{max}, we see an increase in the height of the peak and an increasing steepness before the phase transition begins. These are both relatively natural consequences of the flow being allowed to be higher before saturation becomes an interference. Afterwards, all v_{max} values converge to the same relationship, which is also to be expected as at this point of saturation, none of the vehicles will be able to reach their respective maximum speeds, and the relationship will be the same regardless of it. Another thing to note here is an increasing steepness of the dropoff behaviour exhibited just after the peak. Due to the wrapping nature of the road, cars accelerating to a higher speed will take less time to reach the front of the next stop wave, and so the flow rate decreases more and more quickly once the critical density has been reached.

Finally, we will allow there to be a probability of particularly slow/fast drivers, by altering their respective v_{max} by a given probability, p_{shift}. We will vary this probability between 0 and 0.5, as well as the extent to which the respective v_{max} values are altered to attempt to note a significant change in the flow/density relationship.

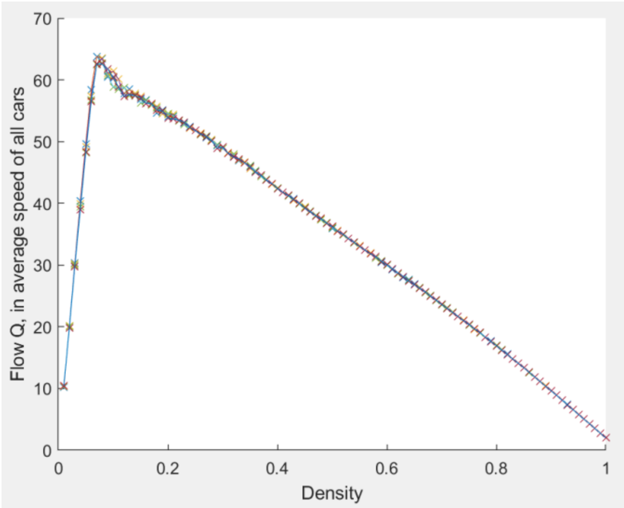

Figure 11: Results for p_{speed} =0, 0.2, 0.4 & 0.6 respectively for v_{max}=12

These results are relatively underwhelming. For different values of p_{speed} we see effectively no change in the trend of the flow/density relation. This could potentially be explained by the fact that once critical density is reached, and the phase change introduced, speeding drivers will have to move at the speed of the wave regardless – they have no room to speed up as a result of the congestion at this point. At densities lower than the critical value, flow rates deviate slightly but negligibly. Another thought is that both of the speeding/slower drivers have been implemented simultaneously, each one with a probability of 0.5 to occur if p_{speed} is met.

We will allow also the v_{max} to drastically alter in another attempt to illustrate a deviation from the typical relationship, and will also only allow speeding drivers to exist, not the slower ones.

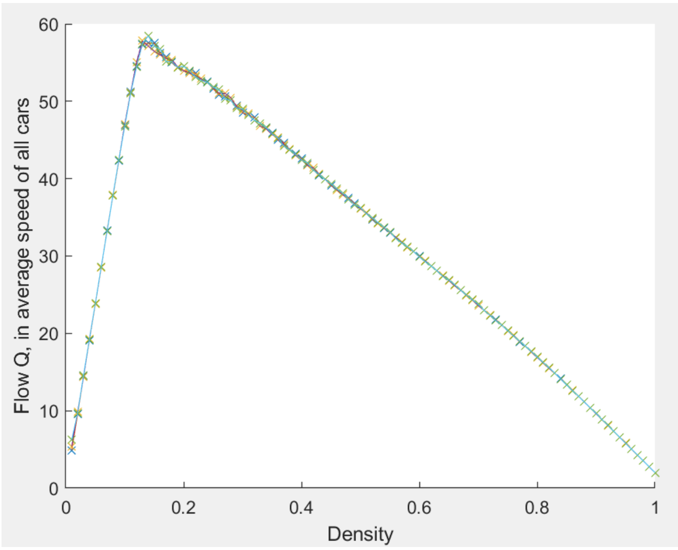

Figure 12: Results for p_{speed} =0.2, with v_{max} alterations of +1, +5 & +10 for v_{max} =5

Again, there is no significant change. The relationships are almost identical. The inference here is that the inclusion of speeding drivers has almost no effect on the net flow of the road. This is likely simply due to the fact that, even if they can speed up momentarily, they are still bound to the movement conditions of the cars directly in front and behind them, and must travel with the speed of the resulting wave.

Evaluation

With a successful model of traffic flow, found by cross-referencing results with both experimental data & results obtained by Nagel & Schreckenberg, we have demonstrated various conditions under which traffic flow is optimised for given conditions. The braking probability introduced to mimic real life deviations from 'perfect' driving behaviour were shown to have a significant impact on not only the maximum flow rate attainable but also the density at which that maximum can be achieved. The natural conclusion here is that drivers should be encouraged to stay as close to this behaviour as possible in order to maximise flow and ease congestion. Alternatively, the human element could be removed completely by the introduction of driverless cars to the road, which could theoretically mimic the behaviour seen in Figures 4 & 5. While increasing the maximum speed allowable did have a slight impact on the peak, it is effectively negligible compared to the random stopping probability's contribution, implying a slight need for possible 'tuning' weighed against the increased risks that would come with raising a speed limit on a given road. The introduction of speeding motorists resulted in no noticeable effect on any of the flow/density relationship, and thus is not particularly noteworthy in evaluating the model's usefulness for the optimisation of traffic flow.

Conclusion

In conclusion, a valid model has been created that effectively and accurately mimics traffic flow within certain restrictive conditions. The results from Nagel and Schreckenberg were extended to vary other parameters of the model and evaluate its effects on the traffic flow.

From this model comes suggestions of ways to optimise traffic on real roads, but limitations of the model are obvious. It is only valid, for example, on single lane traffic with no bottlenecks that are not the result of driver error. It also does not account for how junctions may affect traffic flow. However, quantitative generalisations to the real world *are* possible in the cases where it is valid, and the main result of attempting to keep the density of cars close to a critical value to maximise traffic flow is not a trivial one.

References

[1] – K Nagel, M. Shreckenberg., *"A cellular automaton model for freeway traffic"* - J. Phys. I France 2, pp.2221-2229, Dec 1992.

[2] – M Bando et al., *"Phenomenological Study of Dynamical Model of Traffic Flow"* -Journal de Physique I, EDP Sciences, 5 (11), pp.1389-1399., 1 Jan 1995.