

Formal Report
Motorway Madness - Cellular Automata

Sebastjan Meznarič

December 15, 2005

Contents

1	Background	3
2	Summary of the work done and results	3
2.1	Preliminary one-lane model	3
2.2	Inhomogeneous traffic	5
2.3	Multiple lanes	9
2.3.1	The description of the rules	9
2.3.2	A two-lane simulation	10
2.3.3	A three-lane simulation	10
2.3.4	The flux comparison	14
3	General discussion of results	16
4	Appendix	17

1 Background

Traffic problems are a very common occurrence in large cities. There also exist a number of consultant firms, which attempt to understand the traffic flow using the mathematical models and simulations. Let us first look at what approaches have been employed so far.

Since the 1950's people have been using the fluid-dynamical approaches to the traffic flow. These models seem to have shown the behaviour of start-stop waves for higher traffic densities, whereas for lower densities the smooth flow was observed. However, the granular approaches have also successfully been used to simulate the traffic. Similarly, they have also proven useful in simulating fluid flows [1].

The project concerns itself with a granular approach to the traffic model. The simulations have been run to model one-lane and multi-lane traffic, as well as traffic consisting of more than one vehicle-type (inhomogeneous traffic). Special behaviour can be observed when the traffic mixture is altered as well as when more lanes are added. All the simulations have used the traffic-in-the-circle approach¹.

The programming was conducted in *Mathematica*, v5.1.1.0. The language was chosen over the compiled languages because of the ease of visualising the results as well as the abundance of the mathematical functions readily available. It therefore allowed me to concentrate on the problem and its implementation rather than requiring me to spend my time on things such as the graphical interface. The short summary of the work done and results can be seen in the following sections of the report.

2 Summary of the work done and results

2.1 Preliminary one-lane model

The work began by programming a simple one-lane model containing only one type of vehicles. A grid was laid out such that each element in it would be 0 if empty and a list, containing the velocity of the vehicle and the distance to the vehicle in front, if occupied. While the distance is not an absolutely necessary component of the list, as it can always be calculated on the run, it is more convenient to apply the rules of traffic propagation if the distance element is present. The list was set up in a randomized manner² such that probability of a spot being occupied equals the average traffic density. The cars were set up such that their velocity $v \in \mathbb{Z}$ satisfied $0 \leq v \leq v_{max}$.

The rules of traffic propagation are the following[2]:

1. If the velocity v of the car is lower than v_{max} , and the distance³ to the next car ahead is larger than $v + 1$, the speed is increased by one.
2. If the distance between the cars $d \leq v$, then the speed is reduced to $d - 1$.
3. The velocity of each car is decreased by one with probability p_d , called dawdling probability.
4. Each vehicle is moved forward by v sites

This simple model already started to exhibit the sort of behaviour described in [1]. A graphical representation of a result can be seen on figure 1.

Several interesting things can be found from the simulation results. The first thing is to compare the average velocity on the lane with the traffic density. For this purpose the traffic setup function was redefined so that the traffic density was as required (the initial function's random nature meant certain amount of deviation from the specified traffic density). The plot of average velocity versus the traffic densities reveals that the average velocity is constant (with respect to density) until a certain critical density is reached. At that point the velocity begins to decrease (see figure 2).

¹A loop of traffic, like in a race.

²For the exact code please see the appendix, file Preliminary.nb and FirstSim.nb

³The distance was defined to be the number of empty cells between the cars




Figure 1: The results of a traffic simulation. The vertical axis represents the time increasing downwards and the horizontal axis is the position of the vehicles increasing to the right. The red vehicles are moving slowly, while the greener ones are faster.




Figure 2: The plot of average velocity versus the traffic density. It is seen that the velocity begins to decrease only after certain critical density, which appears to be approximately $\rho_c \approx 0.1$.

Another interesting quantity is traffic flux. Flux through a certain cell is defined to be $\phi(x) = \rho(x)v(x)$. Note that $\rho(x)$ is either 0 or 1 depending on whether the site is occupied or not. Average flux then is

$$\begin{aligned}\langle\phi\rangle &= \frac{1}{N} \sum_{x=1}^N v(x)\rho(x) \\ &\approx \frac{\langle v\rangle}{N} \sum_{x=1}^N \rho(x) \\ &= \rho \langle v\rangle,\end{aligned}$$

where ρ is the total (average) density. The graph of flux versus the average density can be seen on figure 3. We can also see that at a certain density the flux drops to zero, i.e. the traffic stops moving.

The third graph we shall look at is the graph of average velocity versus the dawdling probability (figure 4). For small p_d it is not difficult to see that the average velocity is, on average, decreased by p_d . This is because p_d represents the probability of a vehicle decreasing its velocity by 1. If at the same time density is low ($p_d, \rho \ll 1$), then

$$v \approx v_{max} - p_d.$$

This can be seen on the left side of the figure 4, indicated by the red line.

2.2 Inhomogeneous traffic

The inhomogeneous traffic (i.e. such that the maximum velocity is allowed to differ from vehicle to vehicle) is an extension of the model described in section 2.1. The implementation required adding an extra parameter to the list describing the vehicles, in addition to the velocity and the distance to the nearest vehicle. The updating subroutine then had to be modified so that instead of reading the maximum velocity from a global variable, it would read it from the vehicle's parameter⁴.

The same sorts of simulations were run for the inhomogeneous traffic as for the homogeneous one. The densities of lorries (the slower vehicles) were varied while the car (faster vehicle) density was kept constant. This allowed for the investigation of average velocity and flux in relation to the lorry density.

It was found that because there is only one lane, the inhomogeneous traffic effectively moves as if the maximum velocity was global and equal to the lowest maximum velocity parameter of all vehicles. So the slowest vehicle is slowing everyone else down, as the vehicles cannot overtake. The time-space traffic diagram can be seen on figure 5.

The velocity and flux graphs show the behaviour familiar from the homogeneous simulations (figures 6 and 7). The difference is that there is not a constant velocity section of the graph. This is because when the lorry density is 0, there is still some traffic on the road (car density is not zero), so the total density is beyond the critical density in the beginning of the simulation. We saw above that the critical density was about 0.1. This is exactly the car density, which means we should not expect a straight line section of the graph. It is also seen that all the graphs' curves for lorries and cars are virtually indistinguishable. The traffic behaves as if the maximum velocity was that of lorries.

The other interesting thing is the relation between the average velocity and the dawdling probability (figure 8). The figure shows two lines, one for lorries and another for cars. As on the velocity-density graph, they are barely distinguishable from one another. Also note that the line does *not* have two different slopes as opposed to the one we obtained for the homogeneous traffic. This is because the lower maximum velocity causes congestion right away and not at some later value of p_d .

⁴The code can be seen in the appendix, file inhomogeneous.nb




Figure 3: The plot of flux versus density. The flux peaks at the same critical density where the velocity starts to decrease. Again, this is found to be at about $\rho_c \approx 0.1$.




Figure 4: The plot of average velocity versus the dawdling probability. We see the linear behaviour in the beginning. However, after a certain value, the slope increases. This can be understood in terms of congestion forming. Major congestion areas begin to form when the slope of the graph increases. The initial slope is -1 , because $\langle v \rangle = v_{max} - p_d$ for low p_d .




Figure 5: The time-space diagram. The vertical axis is time, decreasing downwards and the horizontal axis is space, increasing from left to right. The gray dots are (say) lorries and the white dots are cars. It can be seen how the faster cars are stopped by the slower lorries.




Figure 6: The plot of average velocity versus the lorry density. It is seen that the characteristic constant velocity section is not present




Figure 7: The plot of average flux against the lorry density. The maximum flux is reached at about $\rho \approx 0.2$.




Figure 8: The graph of average velocity versus the dawdling probability for inhomogeneous traffic. The dashed line represents the slower vehicles, while the full line represents the faster ones.

2.3 Multiple lanes

2.3.1 The description of the rules

In the same way as the inhomogeneous simulation is the extension of the homogeneous one, the multilane simulation is the extension of the single-lane inhomogeneous traffic. This results in a general simulation, allowing an arbitrary number of lanes as well as an arbitrary number of vehicle types, each with their own maximum velocity. However, this extension requires a bit more work and thought, as the lane-changing rules need to be defined and implemented. After some consideration, they were chosen to be:

1. The lane being changed to must be a neighbouring lane.
2. The vehicle changing the lanes must not be at rest ($v \neq 0$).
3. The same position on the new lane must be free.
4. The distance to the next vehicle in front of the position must be greater than the distance to the next vehicle on the current lane.
5. The first car behind the position is not allowed to brake for more than $\frac{v_{max}}{3}$.
6. If both left and right lanes satisfy the conditions 1–5, the vehicle changes to the lane with the greater distance to the next vehicle in front. If they are equal, the new lane is chosen randomly, with equal probability of choosing either of the lanes.
7. All the vehicles changing the lanes have their velocities reduced by one⁵.

The multiple lanes were created as a list of lanes, which are themselves lists⁶. The lane-changing was then implemented by writing a special subroutine tasked only with lane-changing. This subroutine further used another subroutine, which decided for every vehicle where to change, returning 1 for changing to the lane to the right, -1 for the lane to the left and 0 for no change. This resulted in a matrix that was then applied to the traffic by the lane-changing subroutine. After the lane changing has been completed, each lane would be updated separately as if there were no other lanes. The already existing inhomogeneous lane-updating function was reused for this purpose.

Two different simulations were run, one containing the rule 7 and one not⁷. In addition to that, a large number of simulations was run automatically to do a flux comparison between a two-lane and one-lane systems. Besides plotting the usual quantities, it also seemed interesting to plot the number of vehicles on different lanes as a function of time. The average velocity was also calculated for each lane separately as well as for all lanes put together, whereas the average maximum velocity on a particular lane made it easier to see whether the traffic was distributing itself through the lanes in any way.

⁵This rule was not implemented in all the simulations. When they are considered separately, it is specified whether the simulation used the rule or not.

⁶The code can be found in the file `multilane.nb`

⁷This was done to compare the flux with and without the rule.

2.3.2 A two-lane simulation

The very first multi-lane simulation contained two lanes, 100 cells each, and was run for 100 steps. The rule 7 was present in this simulation.

Let us first look at the number of vehicles on each lane (figure 9). The graph shows the number of vehicles oscillating up and down on both of the lanes. This can be understood by looking again at the lane-changing rules. If the number of vehicles is increased on a lane, the average distance between them will drop. This means more cars will be changing off the lane. Conversely, if the number of vehicles is decreased, more of them will be changing *to* the lane. The number of vehicles is therefore oscillating about a half of the total vehicles (for the case of N lanes, N -th of the total number of vehicles). On the graph of average velocity versus the time step (figure 10) it is seen that the average velocity oscillates more slowly with time than the number of vehicles. This seems to be the reflection of the reaction time required for the velocity to change when the number of vehicles changes, as vehicles' acceleration is limited to 1 unit of velocity per time step.

The average maximum velocity is also good to look at because it is a measure of the vehicle types on different lanes. If the average maximum velocity is high, it means there are more fast vehicles and vice-versa. It is particularly useful for a low number of vehicle types. As we have only two vehicle types, it is not difficult to deduce the distribution. In our particular simulation (figure 11) the right lane seems to be dominated by the fast vehicles in the last 20 time-steps. It is not, however, always the same lane that is dominated by the fast/slow vehicles. The roles may swap (and indeed often do), as can be seen in the time steps 70-80 on the graph.

2.3.3 A three-lane simulation

The next simulation was a three-lane one. Besides the usual graphs which do not show anything particularly new compared to the two-lane ones (figures 12, 13 and 14), I also plotted a visual representation of the traffic (figure 15). Each strip shows all three lanes at a certain time step. Each lane runs from left to right. Vehicles disappearing at the right side of the strip reemerge on the left side, consistent with the loop boundary conditions.

Even though the density was set at 0.1 for cars and 0.05 for lorries, the start-stop waves do not seem to be appearing.




Figure 9: The figure shows the oscillating number of vehicles on both lanes. It was plotted for the last 30 time-steps. The gray line represents the right lane and the black line the left lane.




Figure 10: The figure shows the graph of average velocity versus the time step for the last 30 time-steps. The gray line corresponds to the right lane and the black line to the left one.




Figure 11: The average maximum velocity versus the time step. The gray line corresponds to the right lane and the black line to the left one.




Figure 12: The number of vehicles as a function of time step on different lanes for the three-lane traffic.




Figure 13: The average velocity as a function of time step on different lanes for the three-lane traffic.




Figure 14: The average maximum velocity as a function of time step on different lanes for the three-lane traffic.




Figure 15: The graphical representation of the traffic motion for three-lanes. Gray dots are the slower vehicles and white dots are the faster ones.




Figure 16: The average velocity as a function of density. The dashed line corresponds to the two-lane traffic, while the full one represents the one-lane one. We see that the two-lane average velocity is *slightly* higher at low densities than the one-lane one. However, it later begins to lag behind.

2.3.4 The flux comparison

The final simulation was actually a set of simulations. It simulated the two-lane traffic with and without rule 7 for 18 different densities. The one-lane simulation was also run for the same densities. This allowed the investigation of the average velocity and average flux as a function of density for both two and one-lane traffic. The simulation was run in the same way as the one used for the inhomogeneous traffic flux calculation, which means the faster vehicles' density was kept constant and the slower vehicles' density was the one that was varied⁸.

First, I am going to present the results of the simulation with rule 7 included. The average velocity and flux comparisons can be seen on figures 16 and 17. It was expected before the simulations were run that the multi-lane set-up would have an advantage over a single lane one in that the faster vehicles could overtake the slower ones and would therefore not be stopped as much. However, comparing the average velocities, it seems that the advantage really only comes to play when densities are relatively low, since at higher densities the situation on the new lane may not be significantly better. In addition, the vehicles suffer a speed penalty upon changing the lanes which makes the total situation worse than with one lane traffic.

Because the rule 7 seems to be a big disadvantage to the multi-lane system, another simulation without it seemed like a good idea. Looking at the lower densities (figures 18 and 19) the difference between running the simulation with and without the rules is apparent. However, at the higher densities the rule does not make much of a difference. The conclusion from this would be that the multi-lane system has a special advantage over a single-lane one at *low* densities, although at the higher densities the advantage wears off.

To support the claim that the difference between the single and multi-lane system is about 0, the average difference between the fluxes was calculated and its standard deviations. It was found to be

$$\langle \Delta \phi \rangle = 0.03 \pm 0.1 \approx 0.$$

⁸In the following paragraphs the word density shall refer to the density of the slower vehicles.




Figure 17: The average flux as a function of density. The dashed line corresponds to the two-lane traffic, while the full one represents the single-lane one.




Figure 18: The average velocity as a function of density. The dashed line is the two lane traffic, while the full one is the single-lane one. The distinction between this graph and that of figure 16 is apparent at low densities. It also shows that at higher densities the multi-lane advantage begins to wear off.




Figure 19: The average flux as a function of density. As with other figures, the dashed line is the two lane traffic, while the full one is the one-lane one. Again, the flux advantage is apparent at low densities. At higher densities the flux still appears to be lower.

3 General discussion of results

The overall results of the simulations seem to be quite similar to the real traffic. Compare, for instance, figures 1 and 20. Both pictures at least qualitatively agree. The cars move with a velocity, given by the slope of the trajectories, which seems to be roughly constant outside the congestion areas. On the other hand, they come to a virtual standstill when hitting the congestion. The motion of the congestion area is in the direction opposite to the traffic motion, which sheds some additional credibility on the simulation.

The multi-lane model, as an extension of the standard one-lane one, makes a prediction that the multi-lane systems only have an advantage over the single-lane one in terms of flux when the traffic density is sufficiently low. However, it should be noted that this advantage comes in addition to the double capacity that such a system offers. This could be applied to the situations when it is necessary to extend the number of lanes on a road.




Figure 20: The space-time lines (trajectories) for cars from Aerial Photography. Each line represents the movement of one vehicle in the space-time domain (from [1]).

References

- [1] K. Nagel and M. Schreckenberg: *A cellular automaton model for freeway traffic*. J.Phys.I France, 2:2221-2229, 1992
- [2] Dr A H Harker: *Motorway Madness II - Cellular Automata*. The project specification script.

4 Appendix

The following pages contain the literal code print-out.