Stop-and-Go Waves in Traffic Flow Via Microscopic Simulation

R. Gibson R. Illner

September 27, 2005

Abstract

This paper explores causes for the formation of stop-and-go waves observed on multi-lane freeways near bottlenecks. A microscopic traffic flow model of a two-lane highway with a bottleneck at which vehicles must merge into a single lane is presented. We implement it through the use of a Java program simulation governed by our model. The simulation shows that our model gives a strong symptom of the formation of such stop-and-go waves.

Contents

1 Introduction

There are several different traffic scenarios which produce different types of stop-and-go waves. The simplest case would be found on a high density, multilane highway with aggressive drivers in which a bottleneck forces vehicles to merge into single lane traffic. By aggressive drivers, we mean that vehicles will fight for every bit of road they can advance; they will not slow down until they are forced to do so by the cars in front of them. Such simplistic behaviour means that drivers do not adjust their speeds when they see backed-up traffic ahead. Thus, cars pile up near the bottleneck and experience stop-and-go motion as the front car of their lane merges, followed by the next car moving ahead and then stopping to let the other lanes merge. A simulation of this behaviour and a cure can be found in Ref. [1].

Kinetic bottleneck simulations and real-life traffic studies have shown density "spikes" propagating backwards from the bottleneck, as in [2], and present to us a different type of stop-and-go wave formation (I'm speculating on the real-life studies existing... do they exist?). These suggest that vehicles will experience stop-and-go motion before even reaching the back of the bottlenecked traffic. What type of realistic driver behaviour produces such traffic flow? Obviously, the aggressive behaviour mentioned above will not produce this effect. What we need is more conservative drivers; vehicles must have the ability to look ahead at the situation on the road and adjust their speeds and/or change lanes if a traffic jam is ahead. In Section 2, we describe a traffic flow model used in an attempt to achieve this effect including the highway design, how and when cars accelerate and brake, when cars change lanes, and what cars do at a bottleneck. In Section 3, we present a simulation of traffic flow behaving according to this model. The implementation of the simulation is discussed along with the results which show symptoms of stop-and-go wave formation.

2 The Model

2.1 The Setup

Consider a straight stretch of road along the x-axis from $x = 0$ to $x = 850$: From $x = 0$ to $x = 800$, there are two lanes denoted lane 1 (top lane) and lane 2 (bottom lane); between $x = 760$ and $x = 800$, we have a bottleneck where cars must merge into one single lane; from $x = 800$ to $x = 850$, traffic continues in a single lane denoted lane 3 (see Figure 1). These lengths are all in pixels since our simulation will run directly on the computer screen.

Vehicles enter the system at $x = 0$ with initial speed v where v is a uniformly distributed random variable between 15 and 20 pixels per 200 milliseconds for cars entering lane 1, and is a uniformly distributed random variable between 10 and 20 pixels per 200 milliseconds for cars entering lane 2. The individual car's speed limit, v_{max}, is taken to be this initial speed. We do this since individual drivers have their own maximum speed which they feel comfortable driving at (we can consider 20 pixels per 200 milliseconds to be the legal speed limit on the road). Thus, we consider lane 2 to be the "slower lane" since slower traffic only enters into this lane (however can still change lanes), providing a realistic feature of highway traffic. Vehicles then continue travelling to the right until they pass $x = 850$, at which point they then leave the system. Cars enter lane 1 at "probability rate" ρ_1 and enter lane 2 at "probability rate" ρ_2 (see Section 3.1). We take our road to be initially empty.

Figure 1: The road for our model which is initially empty

2.2 Car Look-ahead Behaviour

Let us ignore lane changing for the moment and focus on how a car acts and reacts to the traffic situation ahead in its lane. Let R_i be a car on the road. We denote the car directly in front of R_i, if it exists, by R_{i-1}. We denote the positions and speeds of R_i and R_{i-1} by x_i and v_i, and x_{i-1} and v_{i-1} respectively.

We consider R_i to have a look-ahead distance in which it surveys the traffic situation ahead in its lane. The faster R_i travels, the further R_i must look to make a good decision on the road. Thus, we define:

$$LAD := \text{Look-ahead Distance} = k_1 v_i + (L + S)$$

Here, L is the length of a car (taken to be the same for every car on the road), S is a safety parameter, and k_1 is a constant. The effect of S, as will be seen later, is that R_i will always attempt to have at least S pixels between the front of itself and the back of R_{i-1}.

For R_i to drive safely on the road, it needs to keep a safe distance from R_{i-1}. This safe distance needs to increase as v_i increases to allow R_i enough time to brake if R_{i-1} happens to slam on its brakes. This is the philosophy of the generally accepted road rule in British Columbia called the "2-second rule" [3]. So we define:

$$OD := \text{Optimal Distance} = k_2 v_i + (L + S)$$

Again, k_2 is a constant. Thus, R_i tries to stay further back from R_{i-1} as its speed increases.

We now come to the ordinary differential equations which define R_i's acceleration. Let $D = |x_{i-1} - x_i|$.

Case 1: $D < OD$

$$\dot{v}_i = \begin{cases} max\left\{-b_{max}, -c_1\left(\frac{1}{D} - \frac{1}{OD}\right)v_i^2 + \epsilon\right\};\ v_{i-1} \leq v_i \\ max\left\{-b_{max}, -b_{slight}\left(\frac{1}{D} - \frac{1}{OD}\right)v_i^2 + \epsilon\right\};\ v_{i-1} > v_i \end{cases}$$

Here, b_{max} is the maximum amount of braking that can be applied by R_i (since it is only physically possible to decelerate so fast), c_1 and b_{slight} are dimensionless constants, and ϵ is a small random error, taken to be a normally distributed random variable with mean 0 and standard deviation 0.2. The error accounts for the driver's inability to perfectly judge the speeds and distances required to make the decision. We will take c_1 to be larger in magnitude than b_{slight} since R_i must brake harder if R_{i-1} is travelling slower than R_i. However, no matter what the speed of R_{i-1}, R_i brakes since it is closer to R_{i-1} than desired.

Case 2: $D = OD$

$$\dot{v}_i = \epsilon$$

Case 3: $D > OD$ *and* $D < LAD$

$$\dot{v}_i = \begin{cases} max\left\{-b_{max}, -c_2\left(\frac{1}{OD} - \frac{1}{D}\right)(v_i - v_{i-1})^2 + \epsilon\right\};\ v_{i-1} < v_i \\ min\left\{a_{max}, c_3 \cdot sgn(v_{max} - v_i)\left(\frac{1}{OD} - \frac{1}{D}\right)max\{(v_{max} - v_i)^2,\ (v_{i-1} - v_i)^2\} + \epsilon\right\};v_{i-1} > v_i \\ \epsilon;\ v_{i-1} = v_i \end{cases}$$

Similar to b_{max}, a_{max} is the maximum acceleration of R_i, and c_2 and c_3 are dimensionless constants. So when R_i is further than the optimal distance from R_{i-1} (and within the look-ahead distance), it brakes according to the top equation when R_{i-1} is going slower and accelerates according to the middle equation when R_{i-1} is going faster. We check the sign of $(v_{max} - v_i)$ to ensure that R_i decelerates if travelling faster than its speed limit.

Case 4: $D \geq LAD$

$$\dot{v}_i = min\left\{a_{max}, c_4(v_{max} - v_i) + \epsilon\right\}$$

Note that c_4 is another constant. Here, R_{i-1} is outside the look-ahead distance of R_i, so R_i just accelerates up to its speed limit. This acceleration term is the only dimensionally inconsistent expression; however, the behaviour of R_i for this case is the least important to our investigation and as long as R_i accelerates towards its speed limit, our model will behave properly.

The key to our desired conservative driving behaviour lies in Case 3. As noted above, R_i brakes when $D > OD$ and $v_{i-1} < v_i$. Also notice that we multiply by the square of the relative velocity, $v_i - v_{i-1}$, rather than the square of the velocity of R_i, v_i, as was done in Case 1. By doing this, R_i will adjust to the flow of traffic ahead of it given that the magnitude of c_2 is large enough. So if a jam is ahead, the flow of traffic will be slow and hence R_i will slow down to R_{i-1}'s speed before reaching the back of the jam. This will help produce stop-and-go waves.

2.3 Lane Changing

Stop-and-go waves seem to appear only in macroscopic models which incorporate bifurcated fundamental diagrams (see [2]). To produce this, we include lane changing terms in the model. To be consistent with the macroscopic models, we have allowed vehicles in lanes 1 and 2 to change lanes in our microscopic model when realistic conditions for doing so are satisfied. The lane changing decision of R_i is dependent on, among other factors, the positions of several other cars on the road. We will continue to denote R_{i-1} as the car directly in front of R_i and in the same lane as R_i. We will now denote the car closest to and ahead of R_i in the opposite lane as R_{ahead} and the car closest to and behind R_i in the opposite lane as R_{behind} (see Figure 2). All of the following conditions must be satisfied for R_i to change lanes:

1. $x_i < 760$
2. $\dot{v}_i \leq a_{min}$
3. $|x_i - x_{behind}| \geq (M+1)L$
4. $v_i < f v_{max}$
5. a) $|x_{ahead} - x_i| \geq LAD$; or b) $|x_{ahead} - x_i| \geq OD^*$ and $v_{ahead} \geq v_{i-1}$

Here, a_{min} is a small constant so that R_i will change lanes only when decelerating or travelling at approximately constant speed; M is a constant denoting the minimum number of car lengths between the back of R_i and the front of R_{behind} needed to safely change lanes; f is taken to be a positive value < 1 so that R_i will only change lanes when travelling some fraction of its speed limit or slower; OD^* is the optimal distance R_i would want to be from R_{ahead} if R_i does change lanes (i.e. $OD^* = OD$ computed as if R_i is in the other lane). If all five conditions are satisfied at any given moment, then R_i changes lanes "with probability p" (see Section 3.1).

Figure 2: The approriate references R_i must consider when deciding whether or not to change lanes

2.4 Merging at the Bottleneck

Now we are concerned with the look-ahead behaviour of cars near the bottleneck. Once $x_i = 760$, R_i must begin to merge into lane 3. We then take the car directly behind R_i and in the opposite lane (R_{behind}) to look ahead at R_i. In other words, R_i becomes the car directly in front of R_{behind} for the purpose of calculating the acceleration of R_{behind} as in Section 2.2. It is possible here for R_{behind} to be very close to R_i at the moment that R_i begins to merge. This makes it important for the parameters c_1 and b_{slight} to be large enough so that once both x_i and x_{behind} are > 800, $|x_i - x_{behind}| > (L+S)$; ie. the cars do not hit each other as they complete their merge.

3 The Simulation

3.1 Implementation

A Java animation program is used to observe traffic flow under the microscopic model described in Section 2. The differential equations for the cars' accelerations are advanced using Euler's method with time step $dt = 0.025$, where each time step produces a frame of animation.

As described in the model, cars enter lane 1 at probability rate ρ_1 and enter lane 2 at probability rate ρ_2. We interpret this in the program by inserting a car at $x = 0$ in lane 1 with probability ρ_1 at each integration time step, and similarly in lane 2 with probability ρ_2. However, we may only insert a car R_i at $x_i = 0$ when $x_{back} > k_2 v_i + (L + S)$, where R_{back} is the car at the back of the lane. We impose this restriction so that there is sufficient room for R_i to enter the road without running into R_{back}. Also, cars are coloured according to their randomly generated speed limits, v_{max}: If $10 \leq v_{max} < 13$ pixels per 200 milliseconds, then the car is red; if $13 \leq v_{max} < 17$ pixels per 200 milliseconds, then the car is yellow; if $17 \leq v_{max} < 20$ pixels per 200 milliseconds, then the car is cyan.

The lane changing conditions in Section 2 are checked for every car at each time step. When all these conditions are satisfied for a particular car at a given time step, it immediately changes lanes with probability p. For simplicity, lane changing occurs instantaneously. Since a car will often have its lane changing conditions satisfied for a lengthy duration of time rather than just at one time step, our cars will almost always change lanes when able to if p is not too small. This means that when a group of cars all see the same jam ahead, they will all change lanes at nearly the same time. This is unrealistic behaviour; we expect cars to change lanes one at a time rather than as a group. For this reason, we choose a small value for p.

The numerical values used in the simulation for the various parameters are listed in Table 1.

3.2 Results

Since our road is initially empty, it takes some time for traffic to build up and to observe interesting behaviour. Figures 3 and 4 show the states of traffic at an early and a later time in one particular simulation respectively.

In Figure 3, we see that no jam has developed yet at the bottleneck.

Parameter	Value	Parameter	Value
b_{min}	1	a_{max}	0.75
b_{slight}	2	k_1	30
b_{max}	10	k_2	1.5
c_1	15	S	5
c_2	3	a_{min}	0.05
c_3	10	f	0.95
c_4	10	ρ_1	0.0029
L	6	ρ_2	0.0058
p	0.005	M	2

Table 1: Data for the model parameters used in the simulation

Thus, cars travel along at safe distances from the cars in front of them and change lanes when appropriate. There is only a very small degree of stop-and-go waves forming at this early stage caused by instances where a slower car changes lanes in front of a faster travelling car. This forces the faster car behind to immediately slow down while the slower car in front begins to speed up after having made more room in front of itself by changing lanes. After a short while, the car behind can then accelerate since there is a larger, safer distance between the two cars. Thus, the car behind experiences a slow down, speed up behaviour (a very small symptom of stop-and-go waves).

Figure 3: The traffic state of one particular simulation at time $t = 102.2$

In Figure 4, a small pile-up of cars are merging at the bottleneck. The cars in lanes 1 and 2 directly behind this pile-up "see" the traffic jam ahead and slow down as designed in our model in Section 2. In turn, these cars create what appears to be a minor jam-up to the drivers further back down the road. This means that they too slow down and wait for this minor jam to clear before speeding up again. Notice that this creates pockets of high density regions where the jams occur on the road and pockets of low density regions of gaps between the groups of cars. Our microscopic model thus shows a type of driver behaviour that could explain density "spikes" found in macroscopic models, as in Ref. [2]. Once the pile-up at the bottleneck clears, the next minor jam of cars then speed up, followed by the cars behind them speeding up as that minor jam clears. The new front group of cars then reach the bottleneck and again create a jam while they merge into lane 3. The group of cars behind them again slow down and wait for this new jam to clear before once again speeding up towards the merging point. Thus, many individual cars experience a slow down, speed up, slow down, speed up cycle. These are not pure stop-and-go waves since cars behind the bottleneck rarely ever come to a full stop, but our model certainly provides a strong symptom of such wave formation in traffic flow.

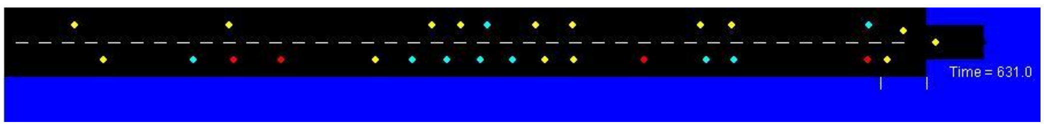

Figure 4: The traffic state for the same simulation at time $t = 631.0$

Other values for ρ_1 and ρ_2 were also simulated. We found that for $\rho_1 = 0.001$ and $\rho_2 = 0.002$, the car insertion rates were too low to produce the stop-and-go waves found in Figure 4. For $\rho_1 = 0.005$ and $\rho_2 = 0.01$, traffic became very dense after just a short while in the simulation and the amplitudes of the stop-and-go waves appeared to flatten out. As was expected, there seems to be a specific density regime where the formation of good stop-and-go waves occur; outside of this regime, they either are very small in amplitude or don't exist at all.

The most significant cause of the stop-and-go waves is the adjustment to the flow of traffic made by R_i in Case 3 of Section 2.2 (braking scenario). We ran a simulation with $c_2 = 0$ and found that stop-and-go waves disappeared completely. This is understandable since $c_2 = 0$ implies that R_i will not brake until it passes its optimal distance from R_{i-1}, but since R_i can only brake so fast it will end up much closer to R_{i-1} than desired. This behaviour is much too aggressive. In addition, we set $p = 0$ for one simulation (no lane changing). Although the stop-and-go waves did not vanish from the traffic flow entirely, they certainly were much less magnified compared to our model where p is nonzero. We also adjusted the rate parameters ρ_1 and ρ_2 and found that the density regime for stop-and-go wave formation was smaller with $p = 0$. Finally, a simulation was run with each car given the same initial speeds and the same speed limits. This resulted in no change to the stop-and-go waves.

In previous versions of our model, R_i observed the density of traffic ahead and would attempt to stay further back from R_{i-1} as the density ahead increased. To this end, we had $OD = (k_2 + k_3\sqrt{n-1})v_i + (L + S)$, where n was the number of cars ahead of R_i within its look-ahead distance and k_3 was a constant. This actually had little effect on stop-and-go wave formation. If anything, it did more harm than good because some cars were becoming too spaced out and were not creating good enough jams to produce nice propagating stop-and-go waves. Thus, we have not included such terms since they are unneccessary and only further complicate the model.

In summary, it appears that the key ingredients of stop-and-go wave formation are a bottleneck where cars must merge from a multilane highway to a single lane road and driver behaviour that promotes drivers to adjust to the flow of traffic ahead of them. Lane changing also helps to widen the density regime where stop-and-go waves are found.

References

[1] W. Beaty. (n.d.). The Physics Behind Traffic Jams. Retrieved July 21, 2005, from http://www.smartmotorist.com/wav/wav.htm

[2] M. Herty, R. Illner, A. Klar, and V. Panferov. Qualitative Properties of Solutions of Systems of Fokker-Planck Equations in Multilane Traffic Flow. (I don't know the rest of the information where this paper can be found).

[3] Insurance Corporation of British Columbia. (2004, November). ICBC - Road Safety - Safe Driving Tips - Everyday Driving Tips - Handling Your Vehicle. Retrieved July 26, 2005 from http://www.icbc.com/Road_Safety/roadsafety_tips_daily_handle.html