Modeling of Traffic Patterns on Highways

Jordan Hurley

April 6, 2006

Abstract

With the population shift to suburbs and the overall increase in population, the transit systems built as recently as 15 or 20 years ago have become problematic because they do not have enough space. Traffic jams form and dissolve, even when there are no accidents. The capability to model these situations would enable persons to predict these sorts of events, and thus the situations could be avoided, leading to a decrease in travel time even as overall volume increases.

1 Introduction

1.1 Purpose

The purpose of this project is to provide an effective simulation of highway traffic patterns. The first goal was to have the cars have a realistic speed. This involves them going faster in the left lanes than the right, and speeding up and slowing down based on the cars in front of them. They must also crash when they collide with each other. The cars also switch lanes similar to real life. The cars all have individual patience levels that will have them switch lanes if they are not satisfied with the lane they are in, most likely because of its relative speed. The simulation now effectively stays realistic over more situations, and the ultimate goal is standard behavior across the simulation and in all situations. The next goal is to have traffic going the other way too, and then to introduce "rubbernecking" that affects both sides of the road. The end goal is to have more than one highway, with congestion on one affecting the other.

1.2 Scope of Study

The study will only deal with highways and their traffic patterns, not city driving and traffic lights.

2 Background and Review of Literature

At this point, most research into traffic deals with it as fluid dynamics. This is good because of the similarity in terms of slowing the flow when the route narrows, but using just fluid dynamics doesn't work. The individual molecules in a pipe all act in the same way and according to the same rules, but every driver drives differently. They have accepted risk levels and make decisions based on their analysis of the situation. Fluids and their environments also don't evolve as rapidly as on-road situations.

3 Procedure, Development

3.1 Preliminary tasks

3.1.1 Finding an effective modeling tool

After trying RePast, NetLogo, and MASON, NetLogo was decided on for its easy navigation, simple language, and Model Library.

3.1.2 Familiarizing with NetLogo

A randomly moving population of default turtles was used to learn basic NetLogo commands.

3.2 Building the one-way traffic

3.2.1 Creating cars and moving them in a car-like manner

The shape of the turtle was changed to a car, and it was made to move horizontally in the direction it faced. Then multiple cars were added, with 4 lanes instead of the one.

3.2.2 Cars that don't loop and have individual characteristics

Cars were made to randomly generate at the left edge and disappear at the right, instead of looping around. They were given speeds based on what lane they were in and all had different colors.

3.2.3 More effective simulation of actual traffic

Crashing was added in, and so was a slider to choose the frequency at which cars appeared. A plot was made to demonstrate the relationship between number of cars, frequency of generation, and number of crashes. Cars were then enabled to change their speeds based on traffic in front of them. They were made not to generate when there was already a car on the spot where they were added in. This dramatically decreased the number of crashes to a more realistic level. After that, the cars were then made able to switch lanes when their speed relative to the lane limit is lower than their "patience". To make this more effective, it would be necessary to also compare it to speeds in nearby lanes. Finally, the cars now stop moving when they crash, instead of behaving like normal cars.

3.2.4 Continued work toward effective simulation

There was a little tinkering in terms of the variables with speeding up and slowing down, which helped out a lot. That made the number of crashes in the first 1/3 of the simulation more evident, though, and so that had to be fixed, in addition with the absurdly high frequency of crashes at startup. In order to fix this problem, more data was utilized. Each lane's average speed was calculated, and that number was used in the generation of new cars, instead of what previously had been the speed limit of each respective lane.

4 Results

4.1 end of quarter 2

At this point, crashes happen far too often. The last version before adding in lane changing and stopped crashes was pretty good at simulating that aspect. The cars do a good job of trying not to get into crashes, and the all show different characteristics, which is a good simulation of the actual road.

4.2 end of quarter 3

I guess the actual simulation part looks pretty good now. My plan going forward for the fourth quarter would be first of all to get traffic moving right to left in addition to left to right, even if they would only affect each other through rubbernecking. After that, who knows, maybe i'll try something crazy!

5 Discussion

6 Conclusion

Building an exact simulation of a road is probably impossible, given the time constraints, and the fact that I am working by myself. Limitations include environment size, because if it's too small then it won't show one pattern affecting another, and on the other side, the bigger it is, the higher number of cars it has to keep track of with their individual characteristics. In addition, limited computing power limits the size of the simulation.