The Effect of Augmented Driver Behavior on Freeway Traffic Flow

Scott Chacon
Analyst, Investment Banking and Capital Markets
Wells Fargo Securities
Email: scottchacon@gmail.com
Phone: 626-497-8986

Professor Alain L. Kornhauser
Director, Transportation Program
Department of Operations Research & Financial Engineering
Princeton University
Email: alaink@princeton.edu
Phone: 609-258-4657

Word Count: 4,124
Figures: 14
Tables: 2

March 10, 2012

Abstract

This paper investigates the possible virtue of the modification of longitudinal and lane-change behaviors of drivers by intelligent cruise control systems that augment individual driver behavior by enforcing minimum separation between vehicles. Such systems would not only reduce collisions but may also improve traffic flow by reducing the frequency of bottlenecks on freeways. This hypothesis is tested using a modified microsimulation of a length of freeway in Los Angeles County. A transit-oriented minimum time headway controller is compared to a traditional minimum separation intelligent cruise controller. The results show that using a fixed distance policy to control the separation tends to keep the flow more stable during peak periods and reduces travel times.

1 Introduction

The modeling and microsimulation of inhomogeneous traffic flows has assisted city planners in predicting the effectiveness of infrastructure changes. Modeling has also given policy makers some sense of the value of new intelligent driver aids. This paper describes the modification and use of such modeling to explore alternative driver behaviors that could substantially reduce congestion. The focus is on simply changing drivers' headway (or following) distance on the freeway. Using a car-following microsimulation model, this paper first re-creates the downstream flow of a segment of a Los Angeles County freeway given the upstream, on-ramp, and off-ramp flows to prove the usefulness of the model. After justifying the accuracy of the simulation, the model is used to find the optimal minimum following distance for reducing traffic. The analysis shows that if vehicles aren't allowed to get within a few meters of the car in front of them (even in stop-and-go traffic), the frequency of the formation of traffic shock waves is significantly reduced, causing traffic flow during peak periods to remain high and reducing travel times. An intelligent cruise control system could easily be designed to augment individual driver behavior by enforcing this minimum separation between vehicles.

2 Traffic Flow Simulation

The simulation institutes Treiber et al. (2000)'s Intelligent Driver Model (IDM) longitudinal model and Kesting et al. (2007)'s Minimizing Overall Braking decelerations Induced by Lane-changes (MOBIL) model to determine driver behavior over a four-lane, $3,000$ meter length of roadway modeled after a stretch of freeway in Pasadena, California. The segment was chosen because of the number of loop detectors maintained by the Performance Measurement System (PeMS). See Chen et al. (2000) for an in-depth description. All that is important to know for the purposes of the simulation is the fact that the detectors report flow (number of vehicles per time period), occupancy (proportion of time a car is over the detector), and speed data for mainline, on-ramp, and off-ramp loop locations. The observed flows at each detector station allow for testing of the quality of the model and calibration of the parameters to simulate the traffic as realistically as possible.

The freeway being re-created is the 210 freeway heading east between Lake Avenue and N. Sierra Madre Blvd. There are mainline stations before the on-ramp at Lake, before the on-ramp at N. Hill Avenue, before the on-ramp at N. Allen Avenue, and before the on-ramp at Sierra Madre. There are also on-ramp stations at Lake, Hill, and Sierra Madre and off-ramp stations at Hill, N. Altadena Drive, and Sierra Madre. Altogether, there are a total of four in-flows and four out-flows across the segment. Figure 1 shows a satellite image of the area captured using Google Earth. In the figure, KML placemarks highlight the relative locations

Figure 1: Model Section of the 210 Freeway in Pasadena, CA

of each of the mainline stations. The stations are labeled with their unique six-digit station identifier per PeMS. The stations are labeled similarly in the Java model representation.

2.1 Pre-existing Java Applet

The Java simulation used is based off a web applet designed by Martin Treiber that is accessible on his website at http://www.traffic-simulation.de/. The applet implements the IDM and MOBIL models in various settings. In the different scenarios, the roads have two lanes traveling in the same direction, and the roads are either u-shaped or circular. The six scenarios are meant to show something different about driver interactions. The simple ringed road is intended to show how bottlenecks form under different vehicle densities. One of the u-shaped roads has an on-ramp to show how an additional in-flow affects mainline traffic. One of the other u-shaped roads begins as two lanes but then narrows to one near the end to model the effect of a lane closure. While all of these scenarios might induce lane-changes according to the MOBIL model, there is a scenario meant to study this aspect of driver interactions specifically - one circular road forces lane-changes through the placement of obstacles around the track.

Many of the scenarios allow the user to change the various model parameters to study how changing conditions affect the flow. The politeness factor, maximum acceleration, mainline in-flow, and on-ramp flow are a few of the parameters that the user may modify by using sliders located within the graphical user interface. The applet and Java code behind it provide a very good starting point for any type of modeling. Because there are so many different scenarios in Treiber's code, the simulation is easily tailored to model any number of possible real-life situations.

2.2   Modifications to Borrowed Java Code

As mentioned before, Treiber makes all of his Java and class files available for download from his website. The available files were a good starting point for this paper's microsimulation because of its implementation of both the IDM and MOBIL behavior models. The applet is made up of fifty-five files, containing a few thousand lines of code. The re-creation of the stretch of freeway used in this study and depicted in Figure 1 leverages Treiber's code and expands it to include the implementation of four lanes of traffic, multiple on-ramps and off-ramps, and a minimum following distance policy. Figure 2 is a screenshot of the modified applet.

Starting at the top of the window, the four buttons control the program's operation. The applet may be stopped and restarted at any time using the "Start" and "Stop" buttons. Also, the user has the ability to switch between two different simulations. The first, denoted "Use Observed Ramp Data" is for comparing the performance of the IDM and MOBIL implementation as it relates to re-creating the observed PeMS flow data at the four mainline stations. The simulation takes time-varying observed data at the on-ramps, off-ramps, and upstream station and changes the in-flow and out-flow at each location according to the data. The simulated flows at the other three downstream stations are compared to the observed data to measure the functioning of the model. The quality of the model is based on its ability to accurately re-create the downstream flows. Being able to quantify its performance is important when using the model to analyze changes in flow after modifications to the behavioral model.

The button for the second simulation, which reads "Set All Variables Manually," creates a scenario where the user may modify any of eleven model parameters to test their effects on traffic flow. The interface presents eleven sliders corresponding to each of the parameters that may be modified. These include the upstream mainline in-flow, the in-flow at each of the three on-ramps, the safe deceleration for cars coming off the on-ramps, the lane-change politeness factor, the out-flow at each of the off-ramps, a "time warp factor" which controls the speed of the simulation, and a minimum following distance slider to control the minimum allowed headway distance.

In the top-left corner of the window, there is a timer displaying how long the simulation has been running in minutes and seconds. Immediately below, there are three statistics that appear whenever the user decides to institute a minimum following distance. The applet displays the total number of vehicles on the main road and the number and proportion of vehicles with actual headway less than the minimum following distance. Below these statistics and the sliders is the actual traffic visualization.

The stretch of the 210 freeway studied has four lanes and six ramps. In the simulation, each lane as well

as the ramp merging locations are separated by dashed white lines. The total road length is roughly the same as in real life, and the distances between mainline detector stations, on-ramps, and off-ramps is also as close to accurate as possible. The on-ramps have angled access roads beginning to the left of the merging location, and the off-ramps have angled exit roads to the right of where vehicles move off of the main road.

Below the gray strip representing the highway there are four updating plots that either show the time-varying flow at each of the mainline stations or the time-varying difference between the flow and the observed data. A data point is added to the plot every five minutes, and once the points reach the horizontal limit, the plot erases the left-most points and begins to redraw nearest the origin. The flow or flow difference is computed as an average over all of the lanes. The code could be changed easily to plot the flow or even speed at any single or combination of lanes. When testing the performance of the model, additional statistics are displayed above the station numbers. The applet shows current vehicle speed in each of the lanes as well as the average at the location and the observed speed for that time period.

2.3   New Java Classes and Additions to the IDM

Despite being a very good starting point for an implementation of the IDM and MOBIL behavior models, Treiber's publicly available code wasn't very expandable, and re-creating the segment of freeway on the 210 was very involved. Treiber's classes governing the lane-change behavior were written for a two-lane simulation and required massive changes in order to increase the size of the freeway to four. Similarly, the creation of multiple on-ramps required changes to many of the updating functions as well as additions to the private variables of the objects to remove global variables which now had more than one value (one for each on-ramp). Since the the logic was usually already in place, the lane and on-ramp augmentation required minimal work. In contrast, the creation of an "OffRamp" class was very difficult.

The exiting behavior of individual vehicles was modeled by assigning each new vehicle either one of three actual off-ramps or an imaginary distant one far down the road. This was done for each off-ramp according to the rate either specified by the user or by the PeMS data. The next off-ramp assigned a vehicle was the one that had the most backup of demand. Each vehicle object now had to have an added variable which stored the location along the road that was the beginning of their exit. This value was used when determining whether the vehicle should change lanes. The Distance-To-Exit value together with a variable scalar ensured that the vehicles were in the slow lane in time to merge onto the off-ramp. The Distance-To-Exit was given by:

$$\text{Distance-To-Exit} = \text{Beginning-of-Assigned-Exit} - x_M \tag{1}$$

where x_M is the position (in meters) of the vehicle. The scalar, referred to as the Right-Bias, increases the relative incentive a vehicle has to change lanes. In the case where either had passed its exit (circumstances prevented it from changing lanes in time) or was currently on an on-ramp, the scalar isn't affected by there being an instituted minimum following distance - it is simply determined by the default model parameter value (see Table 2). If this isn't the case and the vehicle is approaching its exit, then the Right-Bias is given by:

$$\text{Right-Bias} = \frac{\text{bias}^{\text{default}} - 100}{\text{Length-of-Road} \cdot \text{Distance-to-Exit}} + 100 \tag{2}$$

where 100 is the maximum Right-Bias. Therefore, the incentive that a vehicle has to change lanes is:

$$\text{Advantage} = a_M - a'_M + \text{Right-Bias} \cdot \left(1 - 2 \cdot 1_{\{i>j\}}\right) \tag{3}$$

where j is the current lane and i is the possible new lane (usually in traffic simulations the fast lane is lane zero).

The disadvantage to other drivers is computed as discussed earlier by subtracting the acceleration of the back vehicle in the new lane after a lane-change from the acceleration of the back vehicle without a lane-change. As long as the difference between the advantage to vehicle M and the disadvantage to others multiplied by the politeness factor is still greater than the threshold, then the lane-change will take place. The changing Right-Bias value gives greater incentive to vehicles to change lanes when approaching their off-ramps. Since the method that transfers the vehicles from the mainline lane onto the off-ramp simply pulls vehicles off the road if they have reached the beginning of their off-ramp, forcing each vehicle to be in the slow lane in between the front and back of their exit is crucial. If this doesn't happen, every vehicle would simply pass right by their off-ramp.

Acceleration was controlled under two different scenarios by making changes to the desired headway (s_0). If the bumper-to-bumper distance is increasing (the speed of the car in front is greater than the car behind), then desired headway remains unaffected. If the bumper-to-bumper distance is decreasing, meaning the rear vehicle is approaching the car in front of it, then s_0 is set as the minimum following distance. The velocity difference between the leading and following cars is given by $\Delta v_\alpha = v_\alpha - v_{\alpha-1}$ where the subscript α is the index of the rear car and $\alpha - 1$ is the car in front of it. Under the new minimum following distance policy,

the values for the desired headway are:

$$
s_0 = \begin{cases} s_0(\text{vehicle type}), & \text{if } \Delta v < 0 \\ \text{minimum following distance}, & \text{if } \Delta v \geq 0 \end{cases} \tag{4}
$$

where s_0(vehicle type) is the same across all vehicles of the same type in a simulation but is usually different for cars and trucks.

The reason there are two cases in setting s_0 is that the benefit of keeping vehicles further apart is undone if headway increases beyond the set distance at low speeds, which usually happens when the front of a platoon of cars begins to accelerate. Removing the artificial increase of s_0 causes the vehicles to accelerate more in unison.

The results showed that the modified model did a very good job of re-creating the downstream flow of the segment of the 210 freeway. For a detailed discussion of the results, please see Appendix A.

3 The Effect of Setting a Minimum Following Distance

Some Intelligent Cruise Control systems have been shown to affect freeways positively through modifications to human behavior. An Autonomous Intelligent Cruise Control (AICC) system discussed in Ioannou and Chien (1993) applies a constant time headway policy to achieve that end. Ioannou and Chien use a safety distance separation rule that is proportional to the vehicle velocity to eliminate traffic waves. Compared to three different human driver models, Ioannou and Chien showed that AICC has strong potential for smoothing traffic flows and increasing flow rates. The study's shortcoming is that it only looks at single-lane roads without allowing for lane-changes.

A more complicated type of AICC, called Cooperative Adaptive Cruise Control (CACC), has also been shown to positively affect freeway flows. The purpose of CACC is to coordinate the speeds of the vehicles in a platoon by allowing for communication between vehicles. Lead cars send location, speed, and directional information to following vehicles so that following distances and upstream speeds can be optimized for maximum flow. Arem et al. (2006) found that at low penetration rates ($< 40\%$), CACC has no effect on traffic flow throughput. In fact, there is some efficiency loss when only a few vehicles are equipped with CACC. At higher penetrations rates ($> 60\%$), the benefits depend heavily on the state of traffic flow. Because there are more vehicle interactions when traffic volume is high, flow improves significantly since CACC reduces time gaps and improves traffic stability. The overall conclusion was that the presence of

CACC increases freeway capacity, which during already high traffic volume could mean shorter commutes for those traveling at peak times.

While all of the previous studies involved setting minimum headway in terms of seconds, this section is intended to show how setting the length is better because it not only improves vehicle flow but travel times as well. Using the IDM and MOBIL models to simulate driver behavior, eight different minimum following distances were tested: 3,5,7,9,10,15,20, and 40 meters. Table 1 shows the summary statistics for the different trials. For each run, the initial inflows at the on-ramps and upstream mainline end and outflows at the off-ramps were set to ensure consistency. A control was also done with a minimum following distance of zero and then used to compare all of the data. The differences in flow were computed by subtracting the control flow from the flow with positive following distances.

The table shows that flow improvement increases with the headway distance until 10 meters, at which point it begins to decrease and eventually becomes negative at 40 meters. Figure 3 shows the points in a scatterplot with following distance along the horizontal axis and flow improvement on the vertical axis. Note the clear global maximum at 10 meters. Figure 4 plots the average proportion of policy violations versus following distance. As would be expected, the connected curve is upward sloping, approaching what appears to be an asymptote. This makes sense because it's difficult for vehicles to allow for greater following distances because of high initial velocities (the model assumes free flow upstream) and deceleration restrictions which prevent the vehicles from slowing down in time to obey even if they want to.

Indeed, estimating how many people will be abiding by the policy is important because it greatly impacts the viability of the result if applied in the real world. In the theoretical model, every vehicle is assumed to have some intelligent systems that implements the policy, which is unrealistic because short of going door-to-door and handing out the hardware for free, this is an unattainable goal in the short-term.

As claimed in the beginning of the section, minimum following distance also improves travel time. For all following distances, the simulation computed the amount of time it took each vehicle to travel from the first to the last station and the results are shown in Figure 5 in the form of cumulative distribution functions. The control is in black, and from the plots it is clear that all following distances improve travel time, even 40 meters. The shapes of the various curves also supports that 10 meters is the optimal following distance. While the policy leads to slightly fewer travel times between 100 and 250 seconds (over 2,700 meters), it makes up for it by having even fewer travel times that are greater than 250 seconds.

The question that remains is: through what mechanism the following distance has this effect? From time series plots of the flow with the 10-meter following distance compared to the control, it appears that using

Table 1: Results Statistics for Various Following Distances

Headway	Model Statistic	Value
3 meters	Mean of flow differences	75.3
	Median of flow differences	0
	Standard deviation of flow differences	1,044
	Average vehicle violations	5.22%
5 meters	Mean of flow differences	88.2
	Median of flow differences	120
	Standard deviation of flow differences	1,359
	Average vehicle violations	15.96%
7 meters	Mean of flow differences	405.1
	Median of flow differences	360
	Standard deviation of flow differences	1,424
	Average vehicle violations	30.39%
9 meters	Mean of flow differences	975.2
	Median of flow differences	1,020
	Standard deviation of flow differences	1,151
	Average vehicle violations	37.50%
10 meters	Mean of flow differences	1,005.3
	Median of flow differences	960
	Standard deviation of flow differences	1,291
	Average vehicle violations	48.06%
15 meters	Mean of flow differences	832.6
	Median of flow differences	720
	Standard deviation of flow differences	1,341
	Average vehicle violations	61.39%
20 meters	Mean of flow differences	91.5
	Median of flow differences	0
	Standard deviation of flow differences	1,241
	Average vehicle violations	68.06%
40 meters	Mean of flow differences	-991.9
	Median of flow differences	-960
	Standard deviation of flow differences	1,221
	Average vehicle violations	71.16%

distance rather than time to control headway prevents bottlenecks from forming. During the simulations, as time progressed and the freeway became more crowded, bottlenecks formed multiple times during the control, but Figure 6 shows that this didn't happen under the 10-meter forced following distance policy. In the beginning, flow steadily increased during both the control and experiment. It wasn't until about 300 seconds in that the curves diverge. This was when the bottleneck began to form in the control. The position of the blue curve above the red one from then on shows that the minimum following distance prevented the freeway from reaching its critical level. By keeping cars further apart and affecting their movement as a platoon, the instituted policy maintained the capacity of the freeway at near-maximum levels throughout the simulation.

Figure 7 shows the histogram of the flow improvements for the 10-meters following distance. The differences seem to be normally distributed around a mean of 1,005 vehicles per hour and with a standard deviation of 1,291 (see Table 1). Using a 95% confidence interval, the null hypothesis that the flow improvements are random samples from a normal distribution with mean zero is rejected. The results of the t-test with 959 degrees of freedom are as follows:

1. Value of the t-statistic = 24.16

2. p-value (probability, under the null hypothesis, of observing a t-statistic as extreme or more extreme) = 0

3. 95% confidence interval = (923,1086)

Unfortunately, the results of the test can't be given as much weight as one would like. The value of the t-statistic and p-value would imply that instituting a minimum following distance emphatically improves traffic flow. That conclusion is incorrect because the flow differences compiled and used to compute the statistic aren't actually independent, which is a requirement for the t-test - the flows at the upstream stations affect all of the downstream stations at some other time. One would need to do many simulations, record the instantaneous flow difference at the same station and at the same time in each and then aggregate and perform the statistical test. This might seem like bad news, but it doesn't change the result that at least initially a minimum following distance improves both flow and travel time over freeway segments.

4   Conclusion

Using flow data, real arrivals at merges and real diverges were incorporated into a simulation re-creating a segment of freeway. The study concluded that based on this initial analysis, the behavior models used, the IDM for longitudinal motion and MOBIL lane-changes, were very good at re-creating observed traffic. Once tested, the models were then used to test the effects of various driver behaviors on traffic flows. The results show that there is an optimal following distance (measured in meters) that improves flow and travel times during peak demand. This was accomplished by modifications to the driver behavioral model. Specifically, an added headway control element prevented vehicles from getting within 10 meters of the vehicle ahead of them while an acceleration control forced vehicles to maintain no more than that separation if in an accelerating platoon of cars.

Figure 2: Screenshot of User Interface of Java Microsimulation
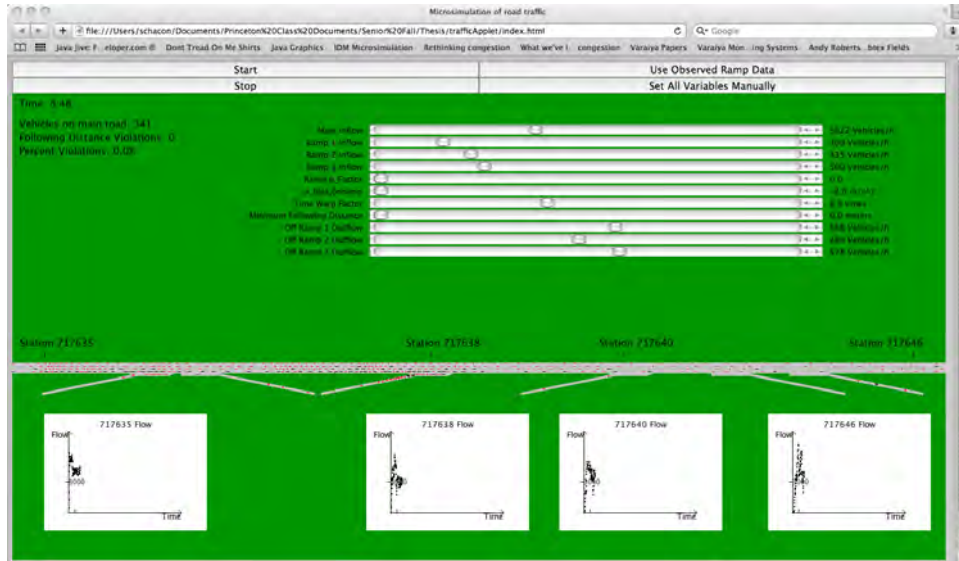
Figure 3: Average Flow Improvement versus Minimum Following Distance
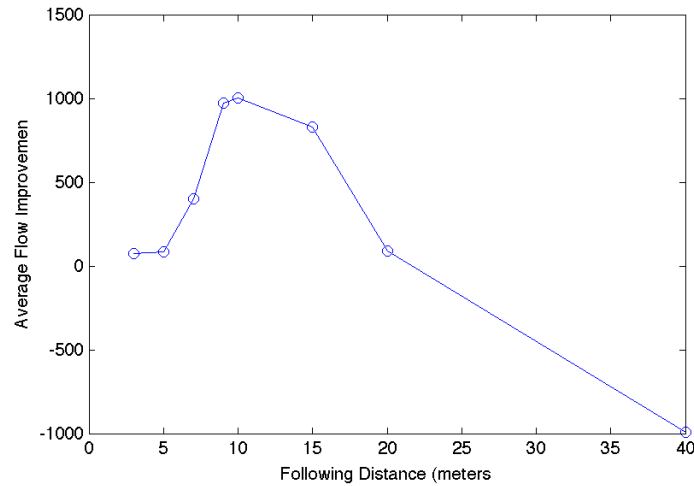
Figure 4: Average Percent of Policy Violations versus Minimum Following Distance
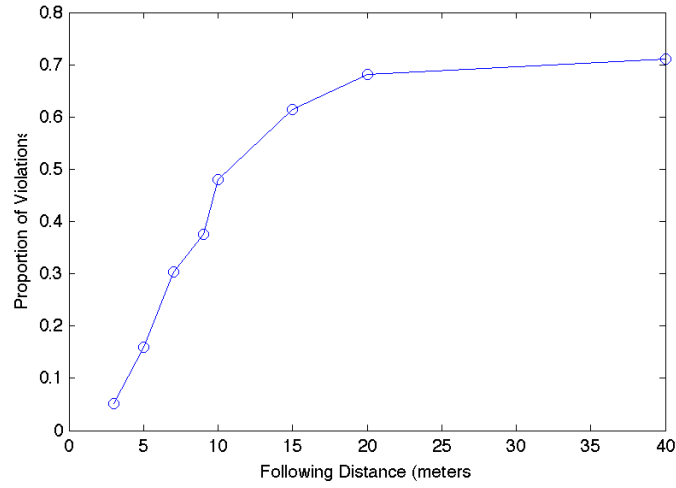
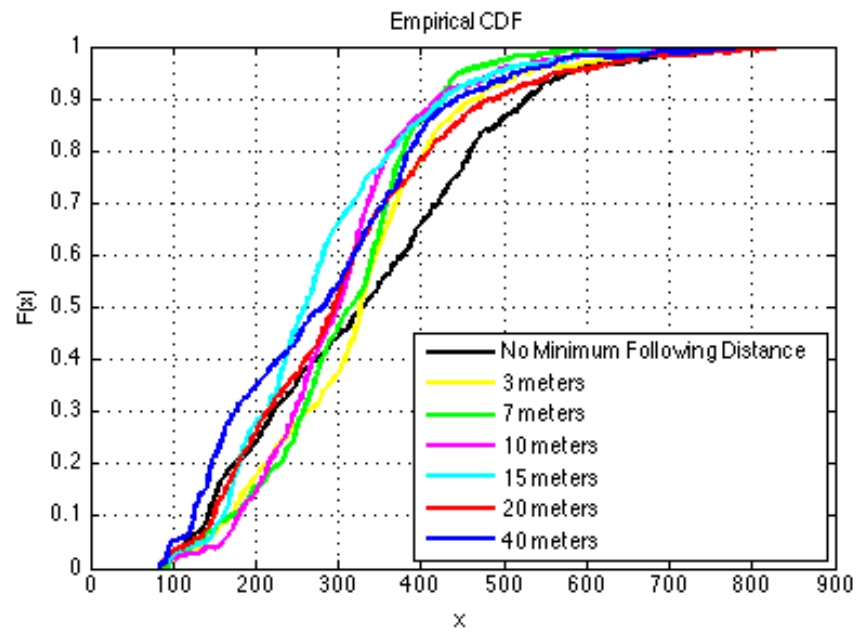Figure 5: Cumulative Distributions of Travel Times for Minimum Following Distances Tested

Figure 6: Time Series of Simulated Flows at Station 717640 with and Without 10-meter Following Distance Constraint
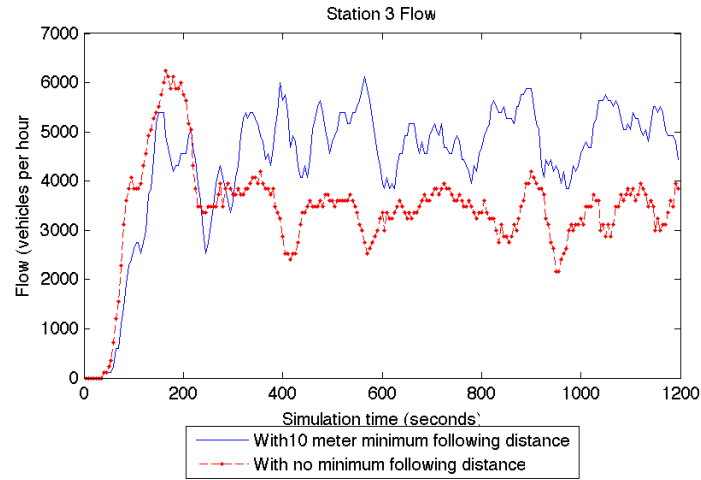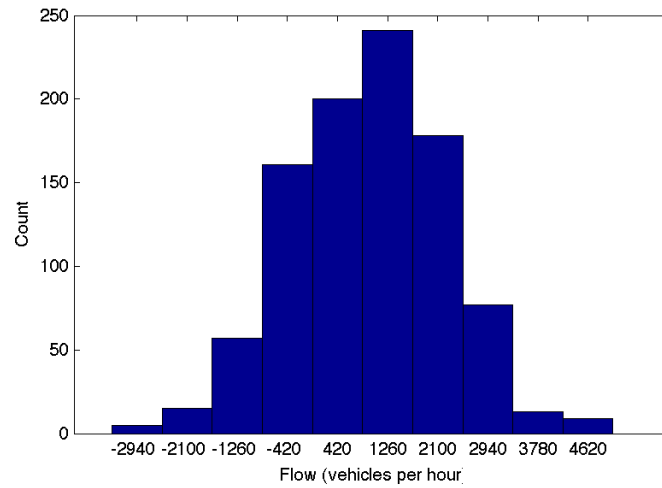

Figure 7: Histogram of Flow Differences with 10-meter Following Distance Constraint

References

Arem, B. V., van Driel, C. J., and Visser, R. (2006). The impact of cooperative adaptive cruise control on traffic-flow characteristics. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):429–436.

Chacon, S. (2009). The analysis, characterization, and visualization of freeway traffic data and the effect of driver behaviors on traffic flows. Unpublished undergraduate thesis, Princeton University, Princeton.

Chen, C., Petty, K., Skabardonis, A., and Varaiya, P. (2000). Freeway performance measurement system: Mining loop detector data. Retrieved November 11, 2008 from `http://paleale.eecs.berkeley.edu/ṽaraiya/papers_ps.dir/pems_paperf.pdf`.

Ioannou, P. and Chien, C. (1993). Autonomous intelligent cruise control. *IEEE Transactions on Vehicular Technology*, 42(4):657–672.

Kesting, A., Treiber, M., and Helbing, D. (2007). General lane-changing model mobil for car-following models. *Transportation Research Record: Journal of the Transportation Research Board*, 1999:86–94.

Treiber, M., Hennecke, A., and Helbing, D. (2000). Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62:1805.

Appendix

A   Model Calibration and Performance Testing

In order to evaluate the realism of the microsimulation based on the IDM and MOBIL driver behavior models, the simulated traffic flows at each station were compared to 5-minute data from September 22, 2008 provided by PeMS. The error in the simulated downstream flows at the last three mainline stations were then computed for a 24 hour period from the early morning to midnight. The testing period covered all 24 hours to ensure that there weren't any errors from seeding problems. If this hadn't been done and the simulation had reached a peak period before the mainline lanes were correctly saturated, the results would've been meaningless since the actual conditions wouldn't have been re-created properly. A situation like this would've led to gross underestimates of travel time and incorrect flows. The model parameters used in the simulation appear in Table 2.

Figures 8, 9, and 10 are time series plots of the flows at the different locations. Note that the flows around midnight and midnight are close to zero, which ensured proper seeding. The first figure shows that flow takes the same general shape at each of the on-ramps - flow is low in the morning and then rises quickly during the morning peak period. Throughout the middle of the day, flow fluctuates but stays relatively high, making the graphs appear to have a plateau. Then later at night, after the evening rush hour, flow drops off again to close to zero.

The plots of the off-ramps have similar characteristics. The interesting difference to note is that the plots reach a global maximum during the morning peak period. That would imply that most travel to these locations is focused heavily during the morning. There isn't as much business traffic during the middle of the day comparatively. Also, the first off-ramp generates the most demand - roughly two and a half times that of the off-ramps.

The mainline stations show much less variation in comparison. They all have almost exactly the same

Table 2: IDM Model Parameters used in the Java Simualtion

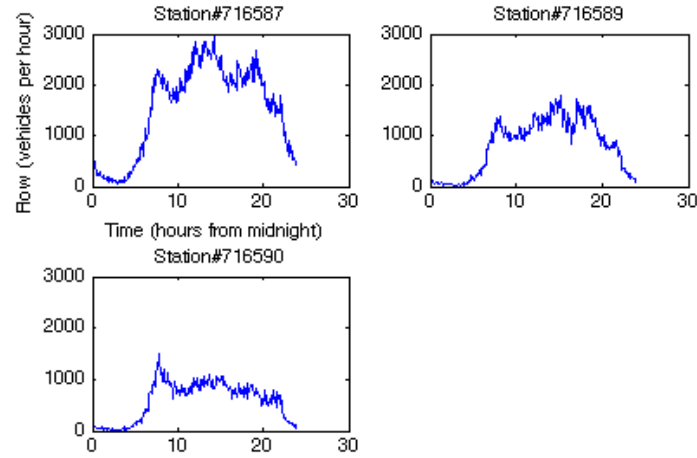| Parameter	Car Value	Truck Value
Desired velocity (v_0)	140 km/h	100 km/h
Time headway T	1.5 seconds	1.7 seconds
Minimum gap s_0	2.0 m	2.0 m
Acceleration (a^{max})	0.3 m/s^2	0.3 m/s^2
Deceleration (b)	3.0 m/s^2	2.0 m/s^2

Figure 8: Observed On-ramp Flows

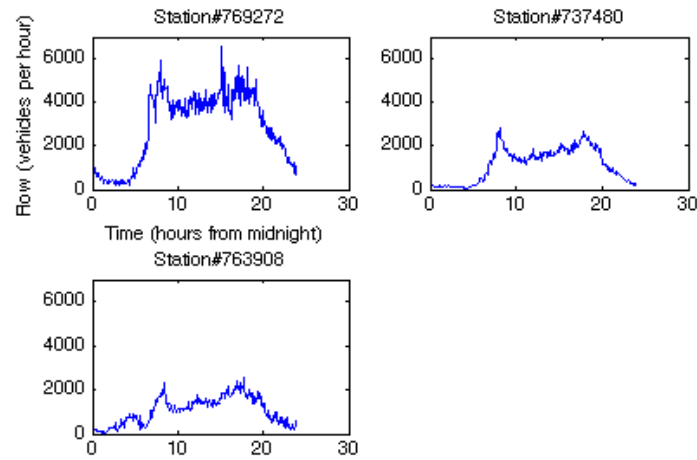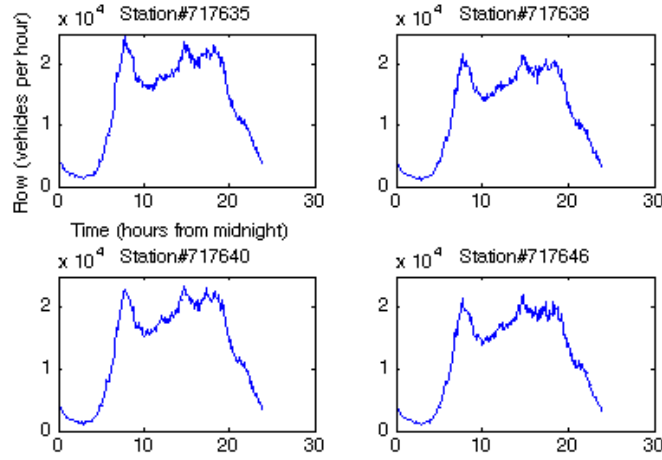Figure 9: Observed Off-ramp Flows

Figure 10: Observed Mainline Flows

shape and have values very close to one another at every single data point. The plots are so similar that it's difficult to tell them apart at first glance. In comparing the simulated flows to the mainline plots, the Java model showed that the IDM did a decent job at modeling driver behavior. Figure 11 shows the simulated flows at the first station. The plots show that the errors are very tightly spread around zero except during the first few hours. This is peculiar because at this time the flows shouldn't have been affected by any bottlenecks or other traffic patterns that could've caused a disruption in the flow generation. The flow rate should've directly reflected the rate that vehicles were added to the simulation, which should've come from the data. The only explanation is that the simulation's computation of flow isn't as accurate when there are very few cars on the road.

Figures 12 and 13 show that the model did well re-creating the flows further down the freeway segment as well. The errors seem to be normally distributed around zero with very similar tight spreads. The errors for the last station are the only ones that seem to imply a breakdown in the model. Figure 14 shows that the flow wasn't as good at re-creating the flow that far down the freeway. The mean error seems to be slightly greater than zero, and the spread has also increased. Given that in between the mainline flow origination and the last station there are three on-ramps and three off-ramps, one would expect the spread to have increased.

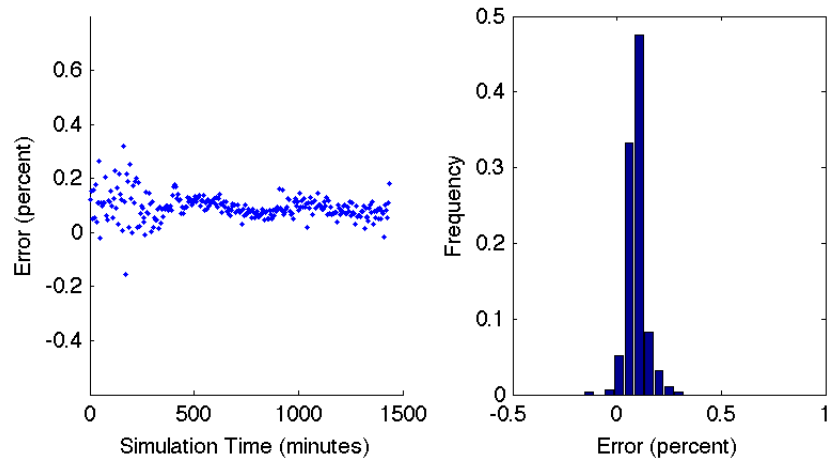Figure 11: Percent Error Between Simulated and Observed Flow at Station 717635

Figure 12: Percent Error Between Simulated and Observed Flow at Station 717638
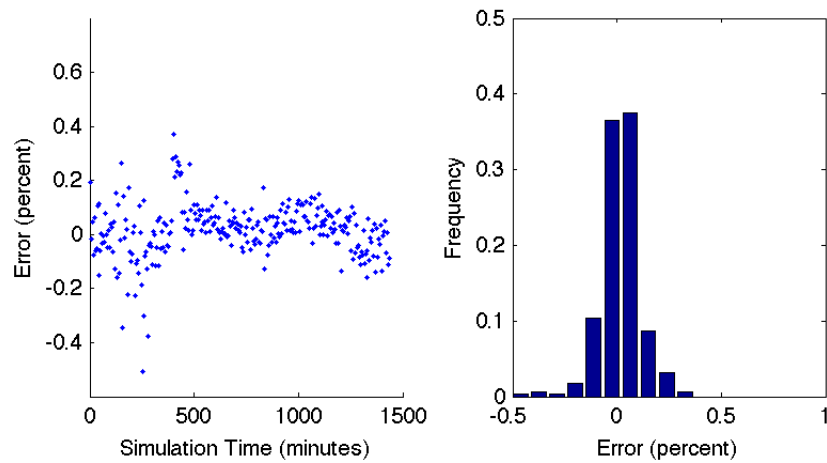
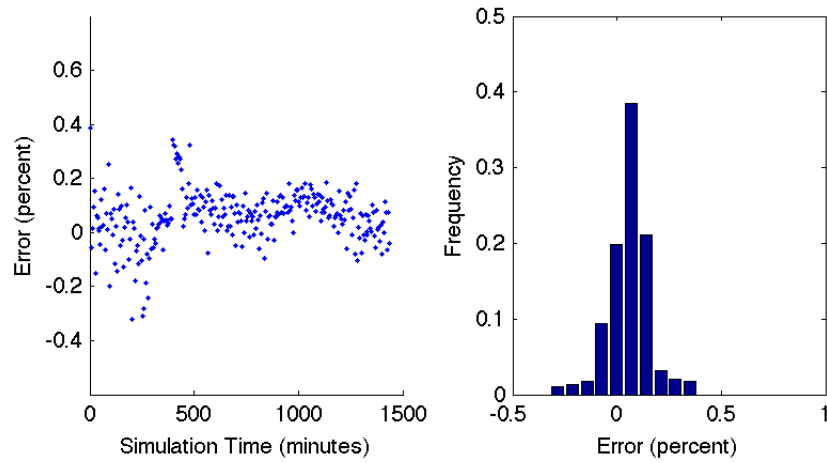Figure 13: Percent Error Between Simulated and Observed Flow at Station 717640

Figure 14: Percent Error Between Simulated and Observed Flow at Station 717646