Complex Adaptive Systems: Energy, Information and Intelligence: Papers from the 2011 AAAI Fall Symposium (FS-11-03)

Using Agent-Based Simulation to Determine an Optimal Lane-Changing Strategy on a Multi-Lane Highway

Joseph Tuzo, John Seymour, Marie desJardins
University of Maryland Baltimore County
Department of Computer Science and Electrical Engineering
1000 Hilltop Circle, Baltimore MD, 21250
{tuzoj1, seymour1, mariedj}@umbc.edu

Abstract

Lane changing can increase or impede the flow of vehicular traffic, depending on traffic density and the lane-changing strategies used by individual drivers. We implement and extend the Nagel-Schreckenberg (N-S) traffic model as an agent-based model to investigate lane-changing behavior on a multi-lane roadway, with the goal of determining which lane changing strategies result in the greatest overall traffic flow. We show that in heavier traffic, an aggressive lane changing policy may be beneficial for overall traffic flow.

Introduction

When stuck in heavy traffic, many drivers attempt to gain headway by attempting to switch into whichever adjacent lane seems to be moving fastest at the moment. This greedy approach seems as though it would lead to a quicker egress from the jam, if only for the individual driver doing the weaving. However, lane changes by individual drivers alter the flow of traffic in the jam; moreover, which lane is moving fastest at a given time depends on many factors, including the influx of vehicles from other lanes. Therefore, excessive lane changing can also hinder both individual and overall progress through a jam.

In order to study the effects of lane changing on this complex system, we have created an agent-based simulation of multi-lane freeway traffic based on the model presented by Nagel and Schreckenberg (1992). Our simulation attempts to replicate stochastic driver behavior by providing each agent with an individual tendency to change lanes based on its situation.

We find that aggressive lane changing strategies do not result in a significant benefit for low ($< 40\%$) traffic densities, but may improve traffic flow as traffic becomes more congested.

Related Work

Much of the previous work on traffic flow analysis is based on Kerner's three-phase theory (Kerner 1998), in which traffic can be in one of three states, or "phases": free flow, in which vehicles move freely with little or no obstruction; synchronized flow, in which a roadway is somewhat congested and vehicles may be prevented from moving at their desired speed but may still move at a reasonable pace; and wide jam, in which a roadway is very congested and vehicles are slowed considerably or even stopped. Of the three phases, synchronized flow is of most interest to researchers because, as Heydecker and Addison (2011) show, it is most susceptible to perturbations that can quickly cause the system to change phases into a wide jam.

When studying the effects of lane changing on traffic flow, it is important to consider the actions of individual vehicles as well as properties of the overall flow. Driver behavior within the system of vehicular traffic affects the behavior of the system overall but depends on the comparatively small-scale analysis of one's immediate neighbors rather than the body of traffic as a whole.

Heydecker and Addison (2011) show that although they are correlated, vehicle speed generally does not influence the formation of traffic jams; rather, existing bottlenecks and jams are usually the cause of slow-moving traffic. Therefore, high congestion alone should not result in significant jams.

Naito and Nagatani (2011) have studied the effects of lane changing on the speed of vehicles in the destination lane, finding that the spacing between vehicles on a freeway plays an important role in preventing jams. This finding suggests that when spacing is tight, frequent or erratic lane changing may provide the perturbations required to cause a phase shift into a wide jam and may continue to exacerbate an existing jam.

Nagel and Schreckenberg (1992) introduced the seminal computer model of freeway traffic (N-S model), which is a cellular automaton model consisting of one-dimensional array of cells representing a single-lane roadway. Values in each cell represent the speed of the vehicle currently "occupying" the cell. Three rules are applied to all vehicles simultaneously to model traffic behavior:

1. Vehicles accelerate up to a predefined maximum velocity if there is room.
2. If there is not enough room, a vehicle modifies its speed (by coasting or braking) so that it does not collide with the vehicle ahead of it.
3. In order to simulate unpredictable driver behavior, some vehicles will slow down randomly.

As noted by Das (2011), the N-S model constitutes a "minimal" simulation, because it cannot be simplified without ceasing to resemble real-world traffic. A number of other researchers have built upon this model to implement more accurate simulations, introducing features such as heterogeneous populations (Zhu, Ge, and Dai 2007; Lárraga and Alvarez-Icaza 2010), slow-to-start behavior (Schadschneider and Schreckenberg 1997), rolling "trickle-flow" jams (as opposed to complete stops) (Das 2011), "intertial" driver behaviors (such as reaction time, which causes delayed response to stimuli) (Ding and Huang 2010), and multi-lane dynamics (Tanaka, Nagatani, and Masukura 2008; Alperovich and Sopasakis 2008).[1] These models report greater accuracy in replicating observed real-world traffic phenomena than the base N-S model.

The measurements generally used to characterize traffic include traffic density and traffic flow. Density, as proposed by Nagel and Schreckenberg, is defined as the ratio of time steps in which a specified cell in a discrete-spaced roadway is occupied to the total number of time steps. Similarly, flow is defined as the ratio of time steps during which a vehicle is moving over a specified cell to the total number of time steps.

Model Design

Our approach extends the original N-S model by casting it as an agent-based model and adding a heterogeneous population, multiple lanes, and lane changing behaviors. The model is implemented using the freely available NetLogo software (Wilensky 1999). We chose to extend the original N-S model, rather than a derivative model, due to its ease of implementation and minor overhead, which allowed simpler analysis.

The model environment consists of a canvas representing a multi-lane roadway. Traffic traverses the canvas in one direction and "loops around" from one end to the other, forming a closed, circular roadway. Each horizontal row of the canvas represents a single lane of travel, while the length of the canvas determines the "circumference" and thus the maximum capacity of the road. Time and locations are measured as discrete, integral units.

Each NetLogo agent, or "turtle," represents a driver-vehicle entity, with state variables describing the vehicle's current and maximum speeds as well as the driver's tendency to change lanes stochastically (*tendency*), current desire to change lanes (*desire*), and a "frustration" rate that increments the driver's desire the longer a vehicle is prevented from moving at its maximum speed. The values of the driver's variables determine whether an agent will employ a cautious or aggressive lane changing policy.

Speed and *max-speed* describe the current and maximum speeds of a vehicle, respectively. *Speed* determines how many cells a vehicle will attempt to move forward at each time step (as in the N-S model), and *max-speed* places an upper limit on *speed*.

Tendency is a fixed real value that determines the rate at which a driver will change lanes without a sufficient "desire," and represents stochastic driver behavior in which a driver does not have a readily apparent reason for changing lanes.

[1]These final two sources analyze and characterize complexity in multi-lane traffic but do not use this data to propose new strategies.

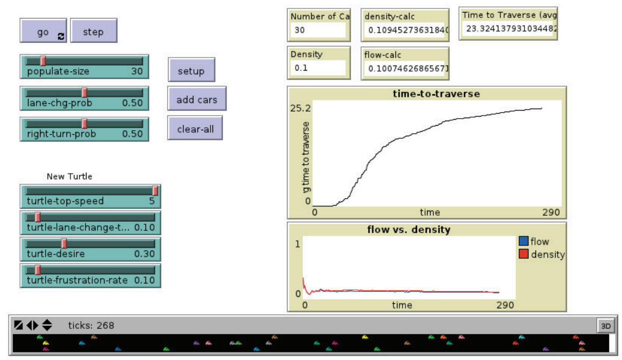

Figure 1: Simulation environment. The sliders on the left control the input parameters, while the readouts on the right display and graph statistics about the traffic flow. The simulated roadway runs along the bottom.

Frustration-rate is a fixed real value that is used to modify *desire* at each time step according to:

$$desire = \begin{cases} desire * (1 + frustration\text{-}rate) \\ \quad \text{if } speed < max\text{-}speed \\ desire / (1 + frustration\text{-}rate) \\ \quad \text{if } speed = max\text{-}speed \\ desire/2 \\ \quad \text{if the agent has just changed lanes} \end{cases} \tag{1}$$

This formula attempts to model building frustration as a driver is stuck in slow traffic, decreasing frustration as they move freely, and sudden relief felt after changing lanes into (presumably) faster traffic.

The state variables for each agent are initialized randomly within certain ranges, using a uniform distribution. Agents may be placed at any unoccupied location on the grid, facing forward (right). Each agent's *tendency*, *desire*, *frustration-rate*, and maximum speed values are set randomly within a range specified by the user. For the first three variables, real values in the range [0, 1] are used, whereas the speed is set as an integer value in the range [0, 5]. All random values are generated by NetLogo's internal random number generator, which uses an implementation of the Mersenne Twister algorithm.

Additional global parameters that can be set by the user control the population size/density, global lane change tendency (*lane-chg-prob*), and proportion of right-merges to left-merges when both are available. These values affect the traffic at an overall rather than an individual level.

Agents in the model behave according to the following rules:

1. If there is room to accelerate, do so according to the N-S specification.
2. With probability *tendency*, attempt to change lanes (randomly) by moving sideways.

3. If there is room to accelerate:
   (a) With probability *lane-chg-prob*:
      i. With probability *desire*, attempt to change lanes (with reason) by moving sideways.
      ii. With probability $(1 - desire)$, slow down.
   (b) With probability $(1 - \textit{lane-chg-prob})$, slow down.
4. If a lane change was unsuccessful, slow down.
5. Update *desire* according to Equation 1.
6. Move vehicle forward.

This rule set is applied to each vehicle in the simulation atomically, so that any changes in the traffic configuration will be reflected immediately after a vehicle moves instead at the end of a time step.[2]

Experimental Design

Three attributes are used to characterize traffic in the simulation: traffic density, traffic flow, and traversal time.

Traffic density is measured, as in Nagel and Schreckenberg (1992), by calculating the ratio of the number of time steps at which a given cell is "occupied" by a vehicle to the total number of time steps. In our simulation, a line of cells to be used for measurement is established across the three lanes of the roadway, forming a "lap" line of sorts.

Traffic flow, again per Nagel and Schreckenberg, is calculated similarly to density but only considers vehicles that are moving ($speed > 0$) at the time they are counted. We modify the calculation from the N-S model to reflect the average flow per lane in multi-lane situations.

Traversal time is the number of time steps it takes a vehicle to travel one lap of the roadway. Each agent tracks its time-to-traverse (*TTT*) for the current lap, for all laps, and its average *TTT*. The global average time-to-traverse (*avg-TTT*) is taken as the average of average *TTT*'s over all agents.

The experimental setup used for testing consists of three lanes of length 100. For each agent, *tendency* is set to a random real value between 0 and 0.1, *desire* is set to a random real value between 0 and 0.3, and *frustration-rate* is set to a random real value between 0 and 0.1. Vehicles are created with a maximum speed of 5 spaces per time step, and are initialized with a random integer speed between 1 and 5.

Results

We conducted trials to determine the effects of lane changing on traffic flow at varying traffic densities. We tested population density in 10% increments from 10% to 80% (early experiments showed that 90% and (trivially) 100% densities resulted in a complete deadlock), and lane-chg-prob in 10% increments from 0 to 1. Each permutation of parameters {population density, lane-chg-prob} was run 10 times through a simulation of 10,000 steps. Table 1 shows the 10-run average *avg-TTT* results for the test runs, while Table 2 gives the corresponding flow values.

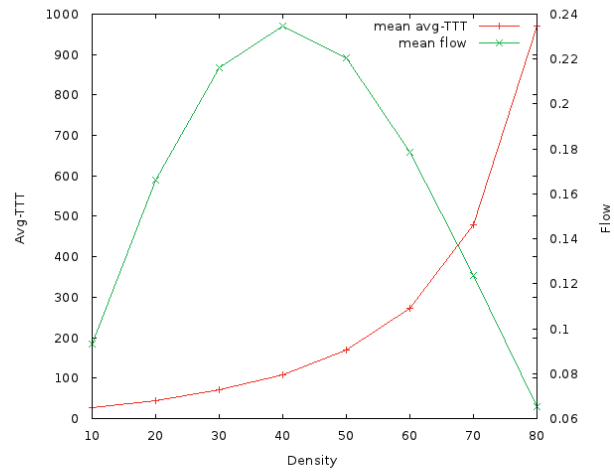

Figure 2: Graphs of *flow* (green "X") and *avg-TTT* (red "+") for $\textit{lane-chg-prob} = 0.2$. *Avg-TTT* increases exponentially with increasing density, while *flow* follows a parabolic arc. Graphs of all other values of *lane-chg-prob* were similar in shape.

In general, as traffic density increased, *avg-TTT* rose exponentially, while *flow* followed a downward-facing parabola (see Figure 2). This is intuitively sensible, since when the traffic density is low, individual agents experience high flow, but there are relatively few of them; conversely, when the density is high, there is high congestion, and many agents experience low flow.

The highest flow values overall were observed at a 40% traffic density; at this density, reported flow was highest at 40% lane changing frequency, presenting a global maximum for our given lane-changing problem. 40% density flow values were higher than those for any other density, suggesting that 40% is an ideal density.

The effect of *lane-chg-prob* varied with the traffic density. At low densities, fluctuations in *lane-chg-prob* were statistically significant; however, the effects were minimal, as shown in Figure 3. At high densities (above 40%), *lane-chg-prob* was positively correlated with flow rate, as shown in Figure 4.

Additionally, fluctuations at low values of *lane-chg-prob* proved to be more significant than those at high values (see Table 3): a jump from 0.1 to 0.2 is likely to be more significant than from 0.9 to 1.0.

Conclusions and Future Work

Our model was able to generalize the Nagel-Schreckenberg paradigm to include stochastic and deterministic lane-changing behavior.

We found that in general, a driver's lane changing strat-

[2]Our model differs significantly in implementation from the N-S model in this regard. N-S updates vehicle positions in parallel, and its rule set is contrived to avoid collisions; since our rule set has to account for the condition of an agent's neighbors at the time the agent makes a decision to move, updates must take into account the current state of the roadway, and not the state at the beginning of the time step. Note that this is an implementation detail, and should not adversely affect the results of the simulation.

	lane-chg-prob										
Density	**0.0**	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**
10%	27.65	27.51	27.38	27.52	27.48	27.64	27.47	27.50	27.51	27.70	27.62
20%	45.34	45.39	45.01	45.08	45.09	45.15	45.23	45.06	45.04	45.17	45.21
30%	71.06	70.45	70.11	70.10	70.15	69.91	69.89	69.75	69.98	69.76	69.91
40%	109.56	108.19	107.41	107.00	106.44	106.23	106.38	106.27	106.03	106.41	106.50
50%	171.42	169.69	168.31	166.88	165.58	164.80	163.95	163.13	163.63	162.99	163.01
60%	278.01	275.76	273.72	272.36	269.22	267.30	264.78	264.17	261.63	260.71	260.45
70%	484.96	481.80	479.72	476.40	474.31	469.08	467.16	463.95	460.24	457.97	455.39
80%	981.03	973.66	971.44	966.84	962.13	963.21	955.71	953.36	954.41	957.91	956.64

Table 1: Effects of varying *lane-chg-prob* (columns) on average time to traverse (*avg-TTT*) at different traffic densities (rows). *Avg-TTT* is measured in units of time steps.

	lane-chg-prob										
Density	**0.0**	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**
10%	0.0926	0.0929	0.0933	0.0937	0.0938	0.0940	0.0939	0.0923	0.0937	0.0931	0.0934
20%	0.1667	0.1650	0.1661	0.1648	0.1659	0.1642	0.1659	0.1652	0.1659	0.1669	0.1652
30%	0.2143	0.2145	0.2162	0.2155	0.2155	0.2160	0.2169	0.2152	0.2157	0.2161	0.2159
40%	*0.2315*	*0.2334*	*0.2345*	*0.2372*	*0.2386*	*0.2381*	*0.2370*	*0.2378*	*0.2382*	*0.2373*	*0.2380*
50%	0.2157	0.2183	0.2203	0.2226	0.2249	0.2265	0.2265	0.2280	0.2293	0.2294	0.2294
60%	0.1751	0.1758	0.1783	0.1795	0.1828	0.1861	0.1881	0.1896	0.1917	0.1918	0.1951
70%	0.1215	0.1227	0.1236	0.1253	0.1271	0.1290	0.1313	0.1343	0.1362	0.1409	0.1455
80%	0.0637	0.0646	0.0652	0.0661	0.0671	0.0677	0.0690	0.0709	0.0726	0.0759	0.0863

Table 2: Effects of varying *lane-chg-prob* (columns) on average flow at different traffic densities (rows). Flow is measured as the number of cars passing the “lap” line per time step. Maximum flow (in italics) was observed at 40% density for all values of *lane-chg-prob*.

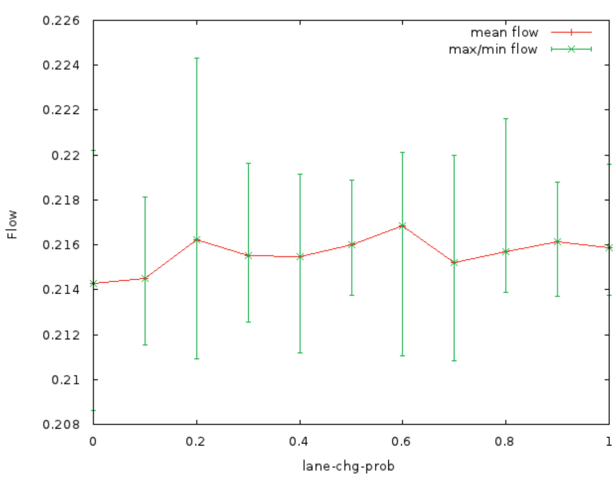

Figure 3: Traffic flow at population 90, or 30% density. Lane changing rate is not strongly correlated with flow, likely due in part to the high variance of the sample.

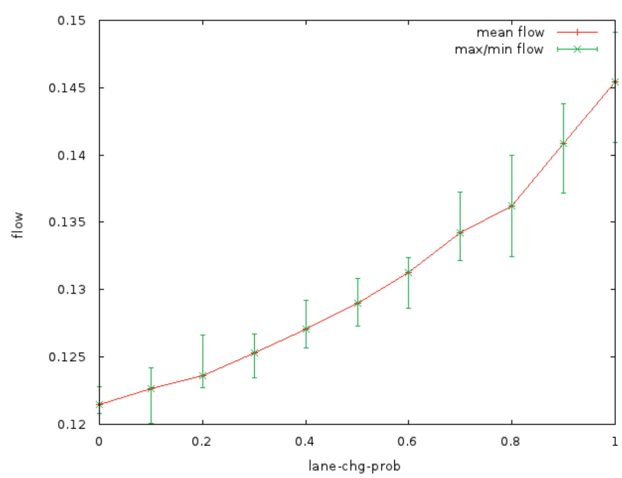

Figure 4: Traffic flow at population 210, or 70% density. Lane changing rate is strongly correlated with flow, although the flow values are visibly lower than in Figure 3.

p-values of T-test on Flow Data

lane-chg-prob	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.1370	**0.0087**	**0.0214**	**0.0118**	**0.0199**	**0.0108**	**0.0223**	**0.0140**	**0.0144**	**0.0208**
0.1		**0.0018**	**0.0073**	**0.0062**	**0.0152**	**0.0092**	**0.0214**	**0.0130**	**0.0145**	**0.0220**
0.2			0.1107	**0.0226**	**0.0392**	**0.0202**	**0.0445**	**0.0245**	**0.0269**	**0.0368**
0.3				**0.0106**	**0.0430**	**0.0280**	**0.0550**	**0.0277**	**0.0327**	**0.0436**
0.4					0.2080	0.0802	0.1168	**0.0462**	0.0555	0.0669
0.5						0.0746	0.1277	**0.0328**	0.0556	0.0711
0.6							0.3673	**0.0486**	0.0699	0.0878
0.7								**0.0011**	**0.0310**	0.0694
0.8									0.1921	0.1409
0.9										0.1583

Table 3: Results of a two-tailed, paired T-test on flow data. Significant values (below 0.05 confidence threshold, bold face) generally appear away from the diagonal, while insignificant values (> 0.05) appear along the diagonal and are concentrated above *lane-chg-prob* $= 0.7$.

egy does not affect traffic flow in cases of low traffic density. When traffic becomes highly congested, however, it may become advantageous for drivers to take opportunities for advancement.

Although our simulation is a realistic enough model of lane changing behavior to answer the question we posed in a mathematical context, the model is also quite basic and could be improved in a number of ways, including:

- Adaptive frustration level: implement a more complex frustration calculation that includes previous experience and may "learn" altruistic or more efficient behaviors.
- Smart lane selection: decide which lane to switch into based on the available lanes' apparent speeds, rather than the current method of using a (weighted) coin toss.
- Non-uniform driver distribution: implement classes of drivers (such as slow, careful drivers and highly aggressive ones) and determine how the ratio of these classes affects traffic flow.

Extending the model in these ways should provide a more accurate representation of actual driver motivations and more realistically model actual traffic.

References

Alperovich, T., and Sopasakis, A. 2008. Stochastic description of traffic flow. *J. Stat. Phys.* 133:1083–1105.

Das, S. 2011. Cellular automata based traffic model that allows the cars to move with a small velocity during congestion. *Chaos, Solitons & Fractals* 44:185–190.

Ding, J. X., and Huang, H. J. 2010. A cellular automata model of traffic flow with consideration of the intertial driving behavior. *International Journal of Modern Physics C* 21(4):549–557.

Heydecker, B. G., and Addison, J. D. 2011. Analysis and modeling of traffic flow under variable speed limits. *Transportation Research Part C* 19:206–217.

Kerner, B. S. 1998. Experimental features of self-organization in traffic flow. *Phys. Rev. Lett.* 81(17):3797–3800.

Lárraga, M. E., and Alvarez-Icaza, L. 2010. A cellular automaton model for traffic flow with safe driving policies. *Journal of Cellular Automata* 5:421–429.

Nagel, K., and Schreckenberg, M. 1992. A cellular automaton model for freeway traffic. *J. Phys. I* 2:2221–2229.

Naito, Y., and Nagatani, T. 2011. Safety-collision transition induced by lane changing in traffic flow. *Physics Letters A* 375:1319–1322.

Schadschneider, A., and Schreckenberg, M. 1997. Traffic flow models with 'slow-to-start' rules. *Ann. Physik* 6(541):1–15.

Tanaka, K.; Nagatani, T.; and Masukura, S. 2008. Fundamental diagram in traffic flow of mixed vehicles on multilane highway. *Physica A* 387:5583–5596.

Wilensky, U. 1999. *NetLogo*. Center for Connected Learning and Computer-Based Modeling, Northwestern University.

Zhu, H. B.; Ge, H. X.; and Dai, S. Q. 2007. A new cellular automaton model for traffic flow with different probability for drivers. *International Journal of Modern Physics C* 18(5):773–782.