Towards a more accurate agent-based multi-lane highway simulation

Yi Luo and Ladislau Bölöni
School of Electrical Engineering and Computer Science
University of Central Florida
Orlando, Florida
yiluo@mail.ucf.edu, lboloni@eecs.ucf.edu

ABSTRACT

The next several years will see an acceleration of the adoption of intelligent driving aid devices. Studying the impact of such devices on the overall traffic performance and safety requires highly realistic microscopic traffic simulation models, which account not only for the overall traffic flow, but also for the details and variability of the individual driver's behavior. In this paper, we first describe a series of improvements on current state of the art multilane highway driving models. The general theme of these improvements is to make the models more agent-like, by considering the specific goals and limitations of the individual drivers and vehicles. We are also performing a more detailed simulation of some of the critical steps in multi-lane highway driving, such as the process of merging at highway entrances and changing lanes. We apply our model to a real world example of the busy commuter highway 408, which crosses Orlando. Through a series of experiments, we investigate the effects of individual driver behavior on the traffic flow.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Distributed Artificial Intelligence: Intelligent Agents

General Terms

Algorithms, Human Factors

Keywords

Agent-based simulation,Traffic models

1. INTRODUCTION

The upcoming years will see an unprecedented adoption of computing technologies in the traditionally conservative world of car manufacturing. We already see sensing systems such as rear and side view cameras and blind spot warning systems, as well as actuating systems such as "follow the preceeding car" cruise control systems, lane maintenance and lane departure warning systems, and automated parking aids. The upcoming standardization of vehicle-to-vehicle and vehicle-to-infrastructure communication around the 5.9GHz DSRC model will likely usher in a collection of new technologies, such as long range signaling of intentions, advance warning of traffic conditions, real time accident warning as well as negotiated long-lifetime highway convoys. Despite many technologies performing actuation as well as sensing, due to legal reasons, many of these technologies will be positioned as *intelligent driver aids* as opposed to *driving automation*. At least for the next decade, the human driver will still be (at least nominally) in control.

Will these devices allow for faster, safer, less congested traffic? This question is surprisingly hard, because the answer depends not only on the technological achievement of the device but also on the social mechanics of the driving. Even a safety device with a proven efficiency such as Anti-Lock Brakes (ABS) generated controversy about its overall effect. If drivers learn to drive on an ABS equipped car, they do not acquire certain skills such as "pumping the brakes", and can become overly confident of the ability of the car to stop on slippery roads. An even more subtle problem is that a driver of an ABS equipped car might overestimate the driving abilities of the other participants in the traffic.

The upcoming proliferation of intelligent driving aid devices will create a traffic with a heterogeneous mix of driving abilities. While some technologies such as ABS and stability control will be probably almost universal for new vehicles, other capabilities will be a source of differentiation for the individual manufacturers. And of course, the overall traffic mix will include older vehicles which possess none of these technologies.

We can conclude that the overall impact of the intelligent driver aids on the traffic will be a mix of clear improvement and unintended negative consequences. The improved driver control and extended amount of information available to the driver can be used for a more efficient driving and early detection of dangerous situations. On the other hand, a thus-equipped driver might be able to push the limits, drive faster, change lanes with only minimal space available, reduce car following distances - all this relying on the superior technology.

An even more problematic situation appears when the majority of drivers rely on intelligent driving aids. This can create traffic flow patterns which are safe and efficient for driver-aid-equipped vehicles (for instance, long and tight convoys relying of highly accurate speed synchronization), but which are dangerous for traffic participants without the necessary technology.

Except for post-facto studies of already deployed technologies, microscopic traffic simulation is virtually the only available technology which can study that impact of intelligent driver aids to the overall flow of the traffic. To achieve this, however, we need to simulate certain factors which currently are only marginally represented in the simulation literature. We need to consider the decision making process of the individual drivers, as well as the way in which the individual driver aids influence it. We need to also represent and model the psychological model of the human driver, its intentions, sensing and actuating limitations and strategies to cope with them (such as defensive driving), and the specifics of its awareness of the situation.

Finally, we need to model in more detail some of the specific actions taken by the drivers, such as lane changing, merging into the traffic, and exiting the highway, as these actions are the source of most traffic accidents.

Many of these requirements can be expressed by simply saying that the simulation of the traffic should be based on the simulation of independent agents. Significant previous work exists in traffic simulation based on statistical physics, particle systems, flow models and cellular automata [**?**, 1, 3]. For instance, dangerous situations arising during lane-changes have been investigated in [4]. We do believe that with increasing power of current computers, the performance increase due to the particle and cellular automata model becomes less important. We can perform real time simulations of long highway stretches representing the position of the cars without discretization, and even considering the gradual transition from one lane to another. Similarly, we can affort to model the individual decision making process of each driver in a relatively detailed way.

In this paper, we describe our work towards developing an agent based model for multilane highway driving, which reaches the sufficient level of realism for modeling the potential impact of intelligent driving aids.

We start with one of the most realistic microscopic traffic models. The simulation technologies which represent the starting point for our model are summarized in Section 2. Then we augment this model with a series of agent oriented extensions in Section 3. Section 4 describes an experimental study based on an accurately modeled stretch of highway (Highway 408, a busy commuter highway crossing Orlando in the East-West direction). We conclude in Section 5.

2. A BASELINE FRAMEWORK FOR MULTI-LANE HIGHWAY SIMULATION

We started our work by implementing in our simulation framework a collection of technologies which together represent a good snapshot of the state of the art in microscopic multi-lane highway simulation. We have chosen as starting point models which are relatively simple and can be augmented with behavioral considerations without a major disruption to the model.

Our baseline is a collection of three technologies: a time-continuous car following model, a lane change model and a human driver model. The first two form the base on which the augmentations described in the next section are applied. The human driver model has been superseeded by our model, but it has been highly influential on its design.

2.1 The car following model

Car following models describe the behavior of a car on a single lane highway. Most such models calculate the acceleration or deceleration of the car though a formula of the following general pattern:

$$\frac{dv_i(t)}{dt} = f(\Delta x_i, v_i, \Delta v_i) \tag{1}$$

where $\Delta x_i = x_{i+1}(t) - x_i(t)$ is the distance between the vehicle and its immediate leader, and $\Delta v_i = v_i(t) - v_{i+1}(t)$ is the approaching speed. The specific formula we choose to use is the one introduced by Treiber et al. [5]:

$$\frac{dv_i(t)}{dt} = a\left[1 - \left(\frac{v_i}{v_0}\right)^4 - \left(\frac{\delta(v_i, \Delta v_i)}{\Delta x_i}\right)^2\right] \tag{2}$$

where a is the maximum acceleration of the vehicle, v_0 is the desired speed, and $\delta(.)$ is the desired distance from the leading vehicle. This distance depends on a number of parameters through the following formula:

$$\delta(v_i, \Delta v_i) = \Delta x_{min} + v_i T + \frac{v_i \Delta v_i}{2\sqrt{ab}} \tag{3}$$

where Δx_{min} is the minimum distance in case of congestion ($v_i = 0$), T is the safe time headway which models the buffering time of the driver, and b is the comfortable deceleration, which couldn't be less than -9 m/s^2 on a dry road.

Let us now discuss the intuitions behind this formula. On a free road, the instant acceleration changes from the maximum acceleration a (when the vehicle is still $v_i = 0$) to 0 (when the vehicle reaches its desired speed $v_i = v_0$). If a vehicle follows a leader vehicle with a negligable approaching speed ($\Delta v_i \approx 0$), the term $v_i T$ in Equation 3 dominates such that the vehicle maintains a distance $v_i T$ from the leader.

In the situation that the vehicle approaches the leader with a high speed, the last term $v_i \Delta v_i / 2\sqrt{ab}$ dominates and the formula dictates a decceleration. The most extreme case is when the vehicle moves with its desired speed v_0 and observes a still obstacle at the distance of x_i. To avoid a collision, the vehicle must brake with deceleration $-b$ when it reaches a distance of $\Delta x_i = v_i^2/2b$. Indeed, this is exactly what the model predicts:

$$\frac{dv_i(t)}{dt} = -a\left(\frac{\delta}{\Delta x_i}\right)^2 = -a\frac{\left(\frac{v_i \Delta v_i}{2\sqrt{ab}}\right)^2}{\Delta x_i^2} = -\frac{v_i^4}{4b\Delta x_i^2} = -b \tag{4}$$

The car following model, defined in this way is considered *collision free*. Of course, this assumes that the drivers have perfect information about their environments. Collisions can still happen if, for instance, a static obstacle or a slow moving vehicle appears on the road at a distance where the vehicle can not come to a stop even if braking with maximum force. We also note that the model assumes that there is no delay between the moment when the acceleration is computed and actual application to the vehicle.

2.2 The lane changing model

Our baseline model extends the car following model with the lane change model described by Kesting et al. [2]. This model assumes that lane changes happen instantaneously: for a shift to the left lane, a vehicle which has been previously in the middle lane, at time t disappears from the middle lane and appears in the left lane. This opens the possibility that a car, coming from behind in the new lane with a higher speed can not break sufficiently quickly and collides with the lane changing car. The model assumes that it is the responsibility of the lane changing car to ensure that the rear left vehicle $j-1$ has sufficient buffer distance such that it can decelerate before hitting the lane changing vehicle

$$\hat{a}_{j-1}(t) \geq -b_{max} \tag{5}$$

If this condition is not satisfied, the vehicle concludes that it is not safe to change lanes.

The second feature of the lane changing model is the analysis of the motivations to change lanes, and the "politeness of the drivers". We assume that the goal of the vehicles is to achieve their desired speed, which implies a certain desired acceleration $\hat{a}_i$. The motivation of the driver to change lanes is such that it can achieve this acceleration (which, we assume, is not achievable in the current lane). However, the changing of lanes might also trigger accelerations in the other vehicles: for instance, it allows the current follower to accelerate, and it might force the new follower to brake.

The notion of *politeness* models the fact that the driver might consider the accelerations of the other vehicles as well when taking a decision to change the lane. The politeness parameter p specifies how much does the vehicle discount the other vehicles' desired acceleration compared to its own. A value $p = 0$ indicates an impolite, fully selfish driver which does not care about other drivers (however, it still considers the safety criteria). The vehicle i will decide to change the lane if the following inequality is verified:

$$(\hat{a}_i + p \cdot (\hat{a}_{j-1} + \hat{a}_{i-1}) - (a_i + p \cdot (a_{i-1} + a_{j-1})) \geq \Delta p_{th} \tag{6}$$

Δp_{th} is the politeness threshold. The left hand side is the difference between the new accelerations $\hat{a}_i$, $\hat{a}_{j-1}$ and $\hat{a}_{i-1}$ if the vehicle i successfully changes into the target lane and the old accelerations a_{i-1} and a_{j-1} if it doesn't change lane. The intuition is that the vehicle favors to change lane only when the advantage of the action is greater than the disadvantage it exerts to its neighboring vehicles. However, because the vehicle i can not obtain the parameters (T, v_0, a, b) for its successor $i-1$ and $j-1$, the utility of lane change can only be calculated by vehicle i's own parameters.

2.3 Human driver model

The models described until now still maintain some elements of the physical models used for traffic simulation starting from the 1950's. Once we move beyond the global picture, many details of the traffic are due to the specific behavior, psychology, cognitive skills and limitations of the *human driver*. There is a general tendency to extend traffic models towards the more accurate modeling of the human driver. Many aspects of our work are in this direction.

A human driver is in some aspects "less capable", but in other aspects "more capable" than the abstract driver envisioned in the models considered up to this point.

It is less capable, because the human driver will inevitably spend some time reasoning about the traffic situation, which delays action. There is also an effect of the cognitive load - humans can make only a limited number of independent decisions per unit of time. Finally, humans occasionally make mistakes, either by taking the wrong decision, not investigating the environment (for instance, by missing a car in the blind spot), or by actuating incorrectly (pushing the wrong pedal or not keeping the car in the lane).

On the other hand, humans usually form a more complete picture of the specific traffic, involving cars around the vehicle, instead of just the vehicle they are currently following. The driver scans ahead several vehicles, and also, occasionally scans the position of vehicles behind it on its own and neighboring lanes. Drivers can also make predictions about the movement patterns of the vehicles. These predictions are facilitated by the similarity of the human drivers' minds. Drivers can recognize the intention of another vehicle to change lanes in the early phases of the action, or even during preparation. In addition, there are standard inter-vehicle signaling methods, such as brake lights and turn signals.

Current microscopic traffic models implement a subset of these human behaviors. Our baseline model is based on Treiber et al. [6] and implements the following aspects.

First, we consider the fact that humans can not perform an indefinite number of decisions per unit of time. This is modeled by considering a time step Δt. At every time step Δt the drivers observe the traffic and make a decision about acceleration. This acceleration value will remain constant for the next interval Δt:

$$\begin{aligned} v_i(t + \triangle t) &= v_i(t) + \dot{v_i}(t)\triangle t \\ x_i(t + \triangle t) &= x_i(t) + v_i(t)\triangle t + \frac{1}{2}\dot{v_i}(t)\triangle t^2 \end{aligned} \tag{7}$$

Another aspect of the human behavior modeled is the reaction time T' necessary to reason about the traffic situation and make decisions accordingly. This can be achieved by substituting in Equation 1 the current state (Δx_i, v_i, Δv_i) at time $t - T'$. If $t - T'$ falls between two simulation steps, then it will be adjusted as:

$$x(t - T') = \beta x_{t-n-1} + (1 - \beta) x_{t-n} \tag{8}$$

3. AGENT-ORIENTED EXTENSIONS

The model presented in the previous section creates realistic multi-lane highway behavior in terms of global numbers and overall look of the traffic. If we are considering it, however, from the point of view of the individual driver, many of the individually important details are "simplified away". The system is calibrated for assuming no accidents (and also no near misses). It assumes that cars are driving on an indefinitely long highway with no concept of source and destination. The cars do not have a preference for a particular lane. There is no modeling of the maneuvers necessary to merge into the traffic from an entrance, neither the lane changes necessary to leave the highway at a particular exit. By assuming that lane changes occur in a single simulation step, it ignores the complex preparation and perspective change needed to change lanes.

The main goal of our model is to extend the realism of the model for a number of significant events in the course of driving. In the following we describe our contributions which extend or replace components of the baseline model.

3.1 Sensing and actuation delays

One of the important real world factors in driving is that there are some psychologically, technologically and physically motivated delays between the information based on which a decision is made, the time when the decision is actually taken and the moment when the consequences of a decision (such as applying a certain acceleration or decceleration or changing lanes) is actually enacted.

Specifically, we will introduce two delay factors:

- The **sensing delay** t_{sense} is the delay between the time when a certain situation happens and the time when this situation becomes available to the driver for reflex action or conscious decision making. The sensing delay will exist even for non-human driver assistant systems: for instance a radar-based car following model

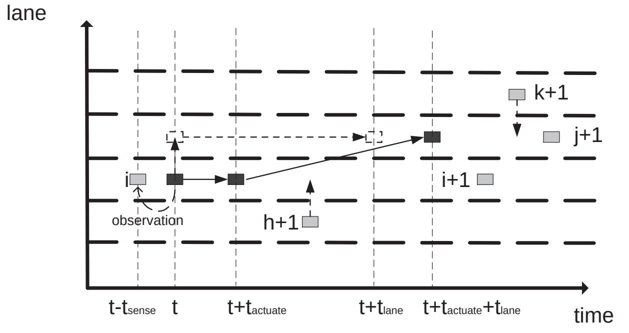

Figure 1: The sensing and actuation delays for a driver agent making a decision at time t

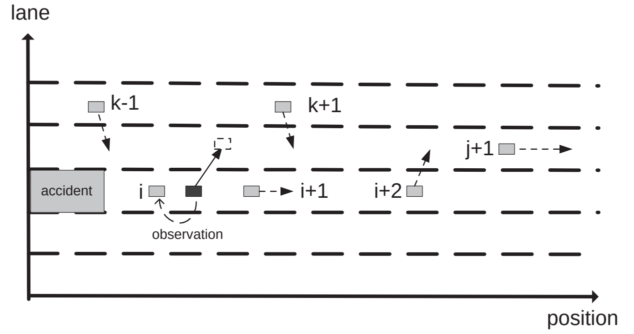

Figure 2: At time t, agent i tries to evaluate the decision of left change.

will have its own delay establishing the distance and speed of the preceding car.

- The **actuation delay** $t_{actuation}$ is the time between the decision making and the actual enactment of the acceleration in the motor vehicle. This includes several factors: first there is the actual decision making time and the motor enactment of the human driver. This is the component which is extended if the driver is under the influence of alcohol or certain medication types. Second, there is the time which is required to enact the control on the car's user interface: for instance, moving the feet from the acceleration to brake pedal and depressing it. Finally, the vehicle's actuation time is the time it takes between receiving a certain command, such as braking, to the moment when the vehicle actually decelerates with the specified value.

In summary, if a decision is made at time t, it is made based on the state of the world at time $t - t_{sense}$ and the decision is enacted at time $t + t_{actuation}$ (see Figure 1).

3.2 Increasing the realism of lane changes

Our baseline model assumes that the lane changes are instantaneous: a vehicle appears in the neighboring lane and disappears from the current lane without any warning. The following car will need to hit the brakes at the very instant when this movement happens. In practice, cars change lanes along a diagonal lane, over a period of time t_{lane} (see Figure 1). At the moment when the lane change starts, the cars in the back on all the lanes will know the intention of the driver, and they will react accordingly. To drive safely, the followers on the destination lane need to act as if the lane change has been completed as soon as it starts. On the other hand, the followers on the source lane act as if the car is still on the target lane until the completion of the maneuver. This caution is justified by the fact that the cars can, indeed, abandon a line change in the middle of the maneuver. We can model this pessimistic reaction by making the assumption that during the lane change maneuver the vehicle occupies both lanes.

This model has implications for the car following and lane-changing model. In the Figure 2, at time t, agent i tries to evaluate the decision of the left change based on the observation at $t - t_{sense}$. Behind itself, the agent observes an accident so it considers no follower in the original lane. On the left, the agent finds no vehicle, but a vehicle two lanes left is changing to the right. So the new follower in the target lane should be vehicle $k - 1$. There are several possible new leaders in the target lane: vehicle $k+1$, $j+1$ or $i+2$. In general, the agent should consider all the vehicles in the target lane, as well as all the vehicles which are moving towards the target lane. In this example, the new leader should be vehicle $k + 1$ as it is the nearest vehicle in the target lane.

To model the cognitive limitations of human drivers, we do not allow vehicles to decide on a second lane change during the time it is engaged in the first. During the lane change, the agent can only control the acceleration of the vehicle. The new accelerations are calculated based on the predicted leader on the destination lane.

3.3 Abandoning lane changes

Let us consider a situation when two vehicles, driving in parallel on the left and right lanes of a three-lane highway simultaneously make the decision to move to the middle lane. This leads to an accident under all the models discussed previously[1]. In real life, however, such situations rarely lead to accidents, because the drivers will become aware of the other driver's intentions either through the index signals or by the movement of the car. Seeing the dangerous situation developing, one or both drivers abandon the lane change and remain in (or return to) their previous lanes.

We implement this behavior in the model. When the two cars initiate the lane change, they will be perceived by other cars as occupying two lanes. Their accelerations will be, however, based on the target lane. If in their predicted position in the common target lane either the accident checker predicts and accident or the agent calculates an acceleration which exceeds the safety limit, the agent will abandon the lane change. If the vehicles are not aligned the vehicle which will abandon the lane will be the one behind. For aligned vehicles both vehicles will abandon the lane change.

3.4 Visibility constraints and defensive driving

Drivers on the highway can only see and consider a subset of the vehicles ahead of them. This is normally expressed through a distance value indicating the farthest distance ahead where a driver can locate a vehicle. Our model

[1]The accident can be avoided under the highly artificial assumptions that the lane change is instantaneous, there is no sensing and actuation delay and the vehicles decision and actions are performed in a strict, non-overlapping sequence.

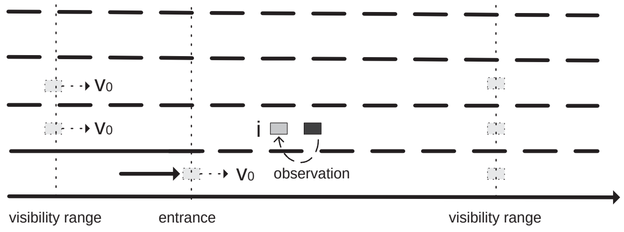

Figure 3: Modeling visibility constraints and defensive driver.

represents the visibility constraints by limiting the vehicles considered in the equations to those which are inside the visibility range.

Environmental conditions such as fog, rain or night affect the visibility range of all the vehicles. However, the visibility can change from vehicle to vehicle. Visibility can be reduced by the driver's vision problems or vehicle dependent factors such as broken windshield wipers. On the other hand, technologies such as fog lights, infrared nighttime vision, or, in the future, inter-vehicle notification can extend the visibility range. To allow such modeling, our simulator allows us to individually set the visibility on a vehicle-by-vehicle basis.

Low visibility situations can trigger accidents because if a vehicle suddenly appears at the edge of the visibility range, the driver might not be able to brake in time. Human drivers respond to this through *defensive driving*, by setting a lower speed under low visibility.

We model defensive driving by adding virtual, non-moving vehicles (obstacles) at the end of the visibility range (see Figure 3). This obstacle appears only in the calculations of acceleration, and will move with the vehicle. The overall result on all lanes is that the vehicles will set their speed in such a way as to be able to stop at the visibility edge. Defensive driving vehicles will not reach their nominal desired speed even on an empty highway even if the highway is empty.

3.5 Modeling the source and destination

Many previous simulations only modeled linear stretches of highways. Our goal is to create a more realistic system which also models the lane changes, the entrances, exits, and lane shifts on the highway. In a typical highway crossing an urban area, there is a high density of entrances, which are usually, but not always, paired with exits. Most exits are right exits, but occasional left exits exist. The number of lanes in a highway changes. A typical pattern for a highway crossing a city is that the number of lanes increases from 2 outside the city, to 3, 4 or 5 lanes in downtown, followed by a gradual decrease as the highway moves out from the city. In addition, a frequent feature of the highways is the gradual shift of the lanes to the right. For many exits the rightmost lane becomes an exit lane and a new lane is added to the left.

For many highways, relatively good statistics are available for the inflow rate $In(i)$ and outflow rate $Out(i)$, at a specific entrance or exit i. This value is easily measurable by simply counting the cars at these exits and entrances. There is, however, no information about the respective source and destinations of the cars, because this would require a full trace of every vehicle throughout the length of the highway.

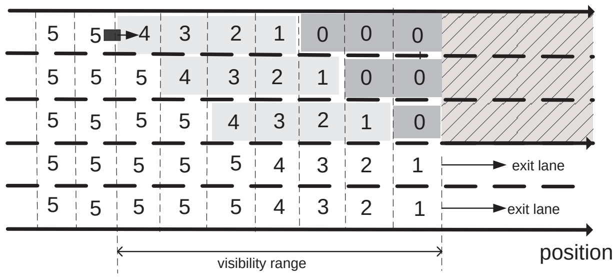

Figure 4: The evaluation of lane preferences of a vehicle approaching the destination exit. Note that the preferences are comparable only within a single column.

As a note, such a tracing *is* possible for vehicles which are using automatic payment transponders at toll roads but this information is not publicly available, and it does not include cars paying cash.

We are using a probabilistic model for the entrance and exit of the vehicles. We assume that the vehicles arrive at entrances with an average rate $In(i)$ following a Poisson distribution. Vehicles have a merging lane of limited length, if this is full, they are queued at the entrance. For each vehicle, we assume that its destination is one of the highway exits downstream. The actual destination is chosen stochastically, with the probability that the vehicle entering at entrance i will have a destination at exit j being:

$$Pr(j) = \frac{Out(j)}{Out(j) + \sum_{k>j} Out(k) - \sum_{l>j} In(l)} \tag{9}$$

where $In(l)$ is the inflow rate of entrance with label l, and $Out(k)$ is the outflow rate of exit with label k. The denominator in the equation 9 is the total number of vehicles which will pass or exit the location. However, the selection probability is calculated in condition that the vehicle doesn't exit before j, so we need to normalize them as

$$Pr(i, j) = \prod_{i<m<j} (1 - Pr(m))Pr(j) \tag{10}$$

3.6 Modeling the preference of vehicles for specific lanes on the highway

Our baseline model assumes that the only reason for a vehicle to change lane is to be able to achieve an acceleration which will bring it closer to the desired speed. In real life, however, drivers also consider other factors. Most drivers will prefer not to drive on the leftmost lane because of the distractions and slowdowns created by entering and exiting cars. On the other hand, when vehicles are approaching their destination exit, they need to gradually get closer to the exit lane. If drivers are notified that a lane ends or it is blocked by an accident, they will try to move away from the lane as soon as possible. We model all these aspects by introducing a preference value for each lane, by each driver, for every point of the road. The introduction of preferences introduces two problems: how are the preference values chosen and how are the preference values affecting the driving.

The preference values depend on the source and destination of the vehicle, the events on the road (such as obstacles and accidents) as well as the preferences of the human driver.

They do not depend on the other vehicles on the road, which are already modeled by the baseline model.

Figure 4 shows the evolution of the preferences of a vehicle driving on a five lane highway which is planning to exit at an exit where the two rightmost lanes are exit lanes. While initially the preferences are identical across all the lanes, the preferences gradually shift towards the right. Right before the exit only the two exit lanes have non-zero preferences.

The preference values are used as a weight in the line change decision process. To allow this we reinterpret the Equation 6 in the following way: $U_{current} = \Delta p_{th}$ is the utility of the current lane, while $U_i = (\hat{a}_i + p \cdot (\hat{a}_{j-1} + \hat{a}_{i-1}) - (a_i + p \cdot (a_{i-1} + a_{j-1}))$ is the utility of the neighboring lane i. The vehicle will move from the current lane i to the neighboring lane j only if $Pr(i) \cdot U_{current} < Pr(j) \cdot U_j$. Notice that the lane change decision is still subject to the safety conditions.

Let us now consider some of the implications of this model. A vehicle travels on the fastest lane. When its destination exit approaches, the preference weights will start to gently favor the right lanes. However, if the other lanes are slower and the desired speed is high, the vehicle will still remain in the leftmost lane, until its preference drops to zero. At this moment, the vehicle will definitely *want* to move to a lane on the right, but it might not be able to do it safely for a while, due to the lane busy with cars. It is possible that the vehicle will not be able to make it to the exit lane in time and it will miss the exit. Such an occurrence is more likely if the vehicle has a fast decrease of the priorities over a short distance span, in contrast with vehicles which adapt their priority long distance ahead of the exit.

3.7 Modeling accidents

Despite the best efforts of the participants in traffic, dangerous situations and accidents do occur in real world highway driving. It should not be the goal of a simulation to model a hypothetical, collision free model. Also, we do not want to introduce accidents for reasons related to the internal working of the simulation.

We want to model the real causes of accidents - whether that be human error coupled with risky driving behavior, increased reaction time due to alcohol or medication, slippery roads or low visibility. Although we strive to model the dangerous and accident situations from first principles, inevitably, some calibration based on real world data will be necessary.

In addition to accidents, which are rare events, we also model and detect situations which are *dangerous*. We also model the aftermath of the accidents, such as lane closures and the congestions triggered by them.

We are considering three specific situations:

- **Dangerous situation:** are cases when an accident did not occur, but only because of a "lucky" set of choices made by other cars. One such example are accidents avoided through the abandoning of the lane changes.

- **Minor accident:** is a collision at relative speeds of < 10 km/h, or a lateral collision when changing lanes. For a light collision, we assume that the vehicles are damaged, they are not continuing to travel, but they will not constitute a road block (they are able to move to the side of the road).

Parameter	Symbol	Value
simulation step	Δt	$0.1s$
maximum deceleration	b_{max}	$5.0m/s^2$
car length	x_{length}	$4m$
minimum distance	Δx_{min}	$2m$
acceleration	a	$1.5m/s^2$
desired deceleration	b	$2.0m/s^2$
headway time	T	$1.5s$
desired speed	v_0	$105 \pm r\%$
politeness	p	$0.5 \pm r\%$
politeness threshold	Δp_{th}	$0.2 \pm r\%$
visibility range	$x_{visibility}$	$200m \pm r\%$
sensing time	t_{sense}	$0.1s$
actuating time	$t_{actuation}$	$0.4s$
lane change time	t_{lane}	$2.0s$
heterogeneity range	r	$0\% \dots 50\%$

Table 1: Simulation parameters

- **Major accident:** the vehicles collide with a relative speed of > 10 km/h. We assume that the vehicles will occupy a highway lane for a period $t_{accident}$. Even if the collision is a result of two vehicles colliding, we assume that only one highway lane will be blocked.

4. EXPERIMENTAL RESULTS

4.1 Experimental settings

In this section we will describe a series of initial experiments performed with out simulator, which implements the baseline model in Section 2 and the agent oriented enhancements in Section 3.

We create a heterogeneous population of vehicles by choosing the vehicles' individual parameters from a heterogeneity range r. The simulation parameters are summarized in Table 1.

4.2 High traffic simulation

For the first set of simulations, we use a hypothetical model of a 4km-long stretch of highway depicted in Figure 5. The originally 3-lane highway has an entrance at 1000m, and it is temporarily extended with a merging lane. At 2500m an exit lane starts which ends at an exit at 3000m. The simulation assumes a very heavy flow of vehicles (1000 vehicle/hour/lane). Figure 6 shows the number of dangerous situations versus the average politeness factor. As expected, more polite driving reduces the number of dangerous situations. Figure 7 shows the number of lane changes function of the politeness. As most of the lane changes done by a car impact negatively the neighboring car's utility, the number of lane changes decreases with the increase in politeness.

4.3 Modeling a real highway

In the second part of our simulation we are modeling a stretch of Highway 408 between UCF and Goldenrod Road (see Figure 8). The inflow and outflow rates for each entrance and exit were taken from official statistics [7].

Figure 9 plots the accident rate function of the actuation time. The main source of the accidents are situations where the vehicles could not abort a lane change in time. As expected, the accident rate increases with the actuation time, with a remarkable take-off after the actuation rate is

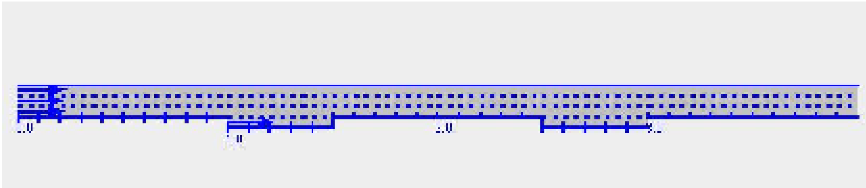

Figure 5: The sample scenario used to calibrate our model

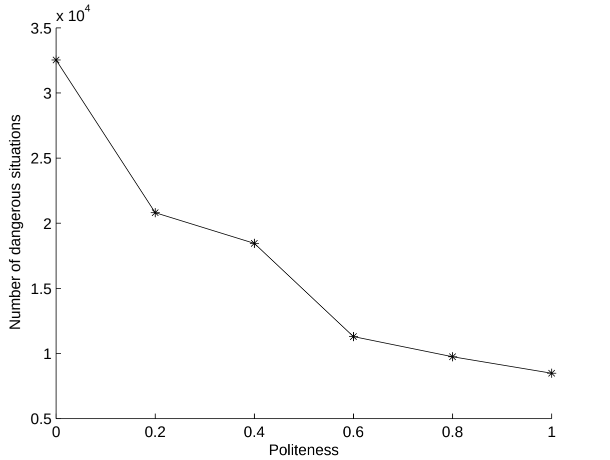

Figure 6: Dangerous situations function of politeness

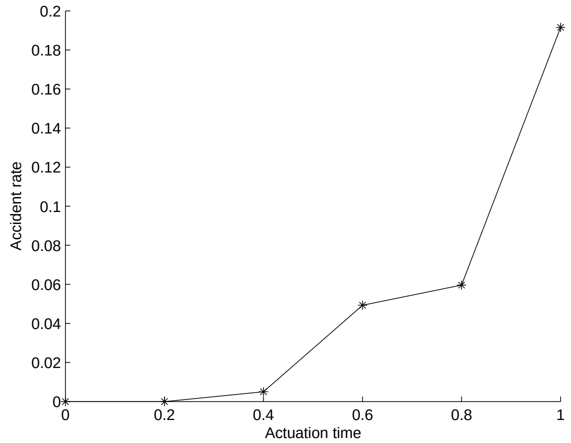

Figure 9: Accident rate function of the actuation time on highway 408

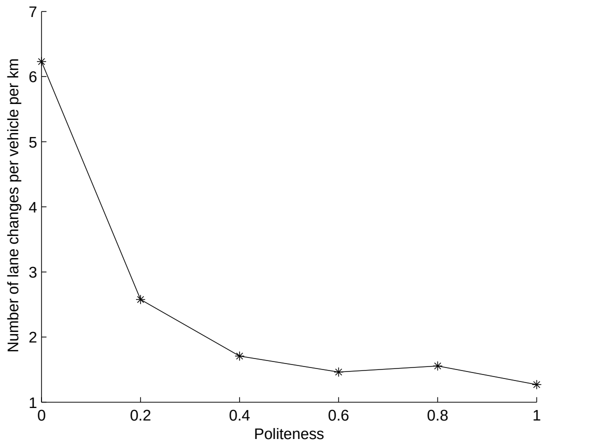

Figure 7: Lane changes function of politeness

higher than 0.4. (Such a universally high actuation time corresponds to the unlikely situation that all the drivers are DUI).

In the last two experiments we study the impact of the heterogeneity of the drivers on the driving performance. Figure 10 plots the number of lane changes versus the heterogeneity range r. As expected, the increase in the heterogeneity of the drivers triggers higher number of lane changes as the drivers with faster desired speed overtake the slower ones.

Figure 11 plots the average relative speed of the drivers over their complete trip from source to destination. For instance, if a driver's desired speed was 100 km/h but it averaged only 70 km/h from source to destination, the relative speed is 0.7. Although this appears low, it also includes the time spent merging in traffic, as well as the time spent sitting in congested traffic. We note that the heterogeneity of the drivers decreases the average relative speed, by creating more traffic congestion and forcing more lane changes.

5. CONCLUSIONS AND FUTURE WORK

In this paper we described a simulation framework which allows a more accurate modeling of individual drivers behavior in multi-lane highway driving, as well as the more specific simulation of specific events such as merging into traffic, preparation for an exit, avoiding accident cars and so on. This simulator opens wide avenues for future studies concerning the implications of particular human driving styles, vehicle heterogeneity and intelligent driving aids.

6. REFERENCES

[1] D. Helbing. Traffic and related self-driven many-particle systems. *Reviews of modern physics*, 73(4):1067–1141, 2001.

[2] A. Kesting, M. Treiber, and D. Helbing. General lane-changing model MOBIL for car-following models. *Transportation Research Record: Journal of the Transportation Research Board*, 1999(-1):86–94, 2007.

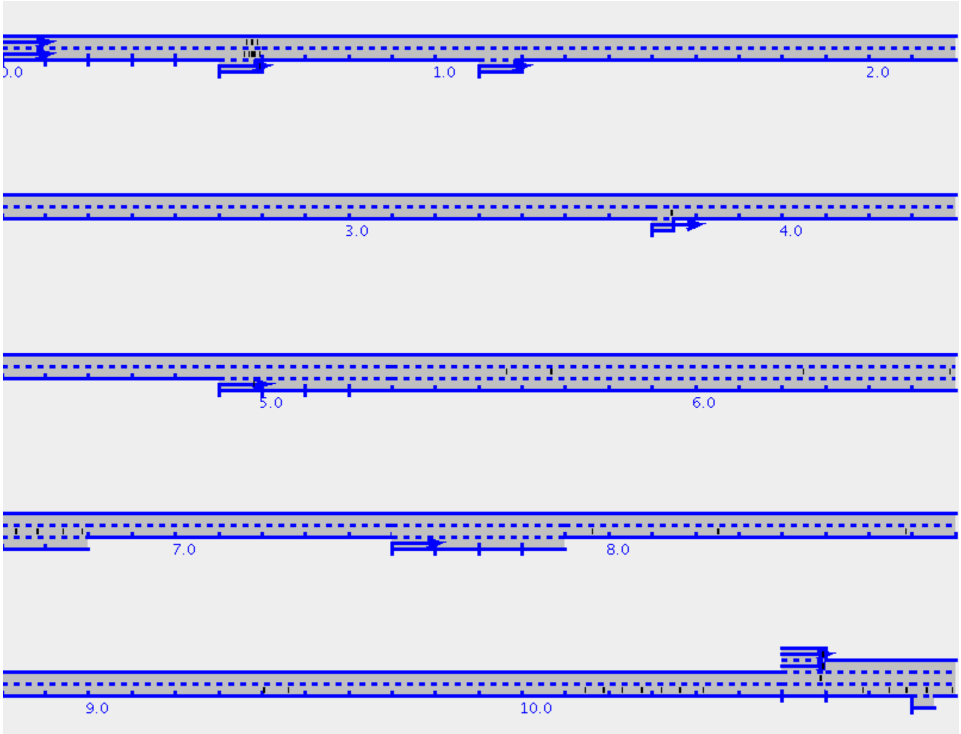

Figure 8: The stretch of Florida Highway 408 from UCF to Goldenrod Rd, with the entrances, exits and evolution of the lanes accurately represented. The distances are in miles.

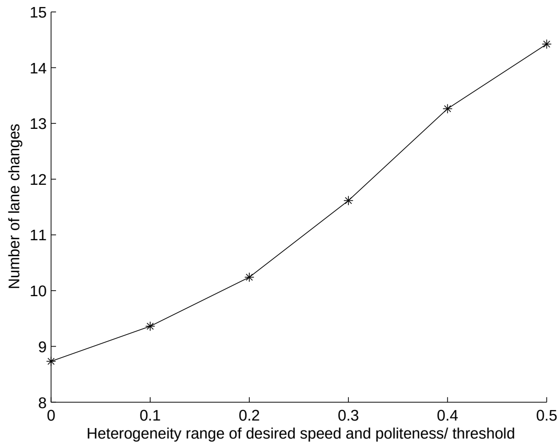

Figure 10: Lane changes function of the heterogeneity range on highway 408

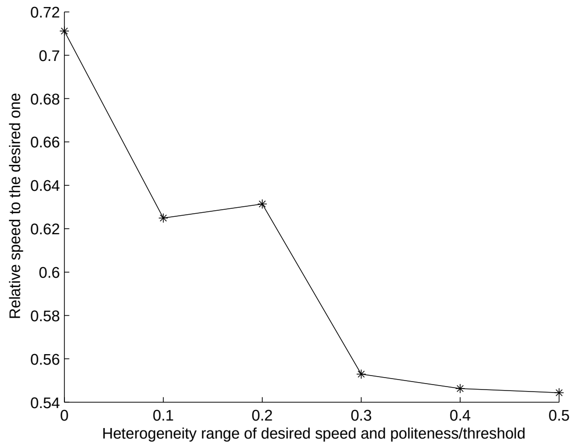

Figure 11: Relative speed function of the heterogeneity range on highway 408

[3] S. Maerivoet and B. De Moor. Cellular automata models of road traffic. *Phys.Rep.*, (419):1–64, 2005.

[4] N. Moussa. Dangerous Situations in Two-Lane Traffic Flow Models. *International Journal of Modern Physics C*, 16:1133–1148, 2005.

[5] M. Treiber, A. Hennecke, and D. Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2):1805–1824, 2000.

[6] M. Treiber, A. Kesting, and D. Helbing. Delays, inaccuracies and anticipation in microscopic traffic models. *Physica A: Statistical Mechanics and its Applications*, 360(1):71–88, 2006.

[7] Expressway history. URL `http://www.expresswayauthority.com/Corporate/aboutStatistics/HistoricalTraffic.aspx`.