CHALMERS

Predictions and simulations of surrounding traffic for automated highway driving of long-combination vehicles

Master's thesis in Complex Adaptive Systems

BJÖRN PERSSON MATTSSON

Department of Applied Mechanics
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2017

MASTER'S THESIS IN COMPLEX ADAPTIVE SYSTEMS

Predictions and simulations of surrounding traffic for automated highway driving of long-combination vehicles

BJÖRN PERSSON MATTSSON

Department of Applied Mechanics
Division of Vehicle Engineering and Autonomous Systems
CHALMERS UNIVERSITY OF TECHNOLOGY

Göteborg, Sweden 2017

Predictions and simulations of surrounding traffic for automated highway driving of long-combination vehicles
BJÖRN PERSSON MATTSSON

Master's thesis 2017:73
ISSN 1652-8557
Department of Applied Mechanics
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone: +46 (0)31-772 1000

Cover:
Interstate 75/85 in Atlanta, Georgia, United States of America. Wikimedia Commons.

Chalmers Reproservice
Göteborg, Sweden 2017

Predictions and simulations of surrounding traffic for automated highway driving of long-combination vehicles
Master's thesis in Complex Adaptive Systems
BJÖRN PERSSON MATTSSON
Department of Applied Mechanics
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology

Abstract

Long combination vehicles (LCVs) are modular heavy trucks which can make the transport sector more effective. Due to the size and complex dynamics of these vehicles, automated driving (AD) functionality has the potential to improve traffic safety and prevent accidents. Lane changes on highways with very dense traffic is an example where traffic predictions are important in order for AD functionality to act safely. For LCV-sized vehicles, situations can occur where it simply is not possible to find a large enough gap in dense traffic which can accommodate the vehicle, and so communication with the surrounding traffic is necessary. This thesis examines the simulation of dense highway traffic situations where the traffic participants are able to react on intention signals such as turning indicators, as well as a method for predicting how the traffic situation will develop in the near future. This is done by introducing a concept of independent and dependent drivers in order to handle expected and emergency scenarios. Three highway traffic scenarios are identified for testing the functionality. It is shown that the system for automated driving becomes more risk-averse by considering the potential for emergency situations.

Keywords: Traffic prediction, traffic simulation, long combination vehicles, automated highway driving

Preface

This thesis was performed under the guidance of Volvo Group Trucks Technology in Gothenburg, Sweden, and parts of the work was done in their facilities. Supervisors at Volvo GTT was Peter Nilsson and Leo Laine.

As part of the work for the thesis, the author also took part in the second Grand Cooperative Driving Challenge (GCDC 2016), which was held in the Netherlands during the spring of 2016 [1], as a member in one of Chalmers two teams. GCDC 2016 was a competition where each participating team had to develop an autonomous control system for a vehicle, in order for it to take part in a number of different traffic scenarios. The primary goal of the competition was to further the development of autonomous and cooperative vehicles, and each team was scored according to how well their vehicle cooperated with those of the other participants. The author's main responsibility during the preparations for the competition was in regards to vehicle detection based on video data. Although Chalmers did not win the competition a great deal of work was accomplished, work which has been used as a basis for continued research at Chalmers.

Acknowledgements

Huge thanks to Peter Nilsson and Leo Laine at Volvo Group Trucks Technology for supervising me throughout the thesis. This project would not have been possible without your help.

I also want to thank Ola Benderius and Bengt Jacobson at Chalmers Division of Vehicle Engineering and Autonomous Systems for all the ideas and discussions we have had.

Finally, I want to express my appreciation to VTI, the Swedish National Road and Transport Research Institute (Statens väg- och transportforskningsinstitut), for making it possible to test the developed functionality in the Sim IV driving simulator at Lindholmen in Gothenburg.

Contents

1 Introduction

With the current environmental challenges that society stands before, there is an increased focus on reducing and optimising the energy consumption for the technology that is used today. In the European Union, the transport sector was responsible for almost a quarter of the total greenhouse gas emissions in 2012, and road transportation alone stood for more than one sixth of that total [2]. There is ongoing research for finding newer and more efficient fuel types, but that alone is not enough to reach the long-term environmental goals that have been set. Greatly decreasing the number of vehicles, in particular those that are fossil-fueled, on the roads is an obvious way to reduce emissions, but in a society that relies on constant transportation of goods and people to function such a step would require major changes in the way people live. However, it should be possible to optimise the transportation sector by outphasing the usage of small and inefficient vehicles, in favour of more efficient ones. This could involve for example car pooling, where several people share a vehicle for essentially the same fuel cost as if only one person used it. It could also be to use larger vehicles for transportation, and so increase the potential amount of transported goods with only a relatively small increase of air resistance. In order to increase the transportation and fuel efficiency even further, much research focuses on the benefits of using road trains and platooning vehicles [3], where a vehicle platoon is led by one driver while the rest of the platooning vehicles are automatically controlled to follow the leading one. This is hoped to lead to lower wind resistance and thus lower energy consumption, higher productivity of the non-leading drivers, and a better traffic flow.

A *long combination vehicle* (LCV) is a type of modular truck that is longer and potentially heavier than what is generally allowed on public roads. Thanks to the increased vehicle productivity and reduced environmental impact of LCVs compared to ordinary trucks, Sweden is now investigating a general introduction of LCVs on the roads. This could, for example, be used for more efficient transportation on highways between larger cities, where the size of smaller trucks (which might be useful on the small streets inside the actual cities) is less beneficial and they would be exposed to more air drag. However, in order for the introduction of LCVs to occur, it is important that issues such as traffic safety and road infrastructure are taken into account. One possible way of improving the safety of LCV driving is to automate some of the more tedious and dangerous tasks of the driver. This is often called *advanced driver assistance systems* (ADAS) and can include such functionalities as adaptive cruise control, parking assistance, emergency braking, or automated lane changes. In modern cars, cruise control and adaptive cruise control have become relatively common, and in even newer cars there can also be collision warning or emergency braking systems in use. Some car manufacturers, such as Tesla, Mercedes-Benz, and Volvo Cars, are also in full development of more complete autopilot systems, where many of the driver's tasks, for example lane changes, are automated. However, even though the problem of making a car fully automated is highly non-trivial, reaching the same level of functionality for a large truck or an even larger LCV is very complex. Due to the sheer size of a LCV, performing a lane change on a highway with dense traffic, which could be fully possible for a normal-sized car, might be impossible for a LCV without manipulating the surrounding traffic in some way, simply because there are no traffic gaps that are sufficiently large. By observing professional human LCV drivers, one can see that lane changes are possible by making use of the truck's turning indicators to signal their intentions to the surrounding drivers. The lateral lane positioning of the LCV can also be used to communicate the driver's desire to make a lane change.

In order for a LCV to handle the complex dynamics of surrounding highway traffic, it must be possible to predict how the traffic will behave to some degree. For human drivers, these sorts of predictions come intuitively, even though they might not always be completely correct. An automated system, however, does not have that same intuition when it comes to traffic predictions, and the resulting semi- or fully automated behaviour must fulfill certain criteria. First of all, the resulting driving behaviour must be safe for the ego vehicle and the surrounding traffic. Second, it must have a high driver acceptance and feel comfortable enough for the human drivers to not turn off the ADAS. This should include that the automated driving must follow traffic rules, but there are unfortunately also some situations in which a behaviour that does not strictly follow the law could result in higher driver acceptance or even safer driving. Third, an automated system should preferably generate efficient solutions in terms of the time and energy consumed during the driving manoeuvres. Another problem is that any developed method or system must be able to perform in real-time, while the vehicles are in movement. The raw sensor data has to be transformed into a useful format before the manoeuvre-generation system can be executed, and the whole process must take place in a time frame that allows a potential emergency behaviour, such as full braking, to be performed before a collision becomes inevitable. This situation imposes high demands on the efficiency of the algorithm and the time taken to complete it.

1.1 Problem description

The questions that this thesis aims to answer are:

- Can complex highway traffic situations be predicted in an approximate and safe way by considering expected and emergency situations?
- Is it possible to utilise these traffic predictions of emergency scenarios to achieve safer lane changes?

By answering these questions, the following implementation goals shall be reached:

- Identify highway scenarios that are relevant when testing autonomous functionality for LCVs in dense traffic.
- Implement a method for simulating the traffic surrounding the LCV, that takes into account intention signals such as turning indicators and lateral lane positioning.
- Propose a traffic-prediction method which allows safe manoeuvring of an autonomous LCV.

1.2 Limitations

Due to the complexity of the described problems, a general solution is not the goal for this thesis. First of all, no sensor noise is modelled in any way. This means that the controller system has access to perfect information of the environment, such as where the surrounding vehicles are, which velocities and accelerations they have, what the road ahead looks like, and so on. Secondly, the solution presented in this thesis is only intended to be used on highways and not in situations inside cities or where there is two-way traffic. Third, the highway road structure is assumed to be simple, with no appearing, disapperaing, or merging lanes. The complexity of the thesis problem lies mainly in the traffic situations, and not in road complexity.

2 Background

In the mid-twentieth century the need for transportation of goods grew notably in Europe. While the size and weight of heavy transports increased, the regulations of vehicle limits differed between countries. With the rise of international trade, a process of achieving a more unified legislation was initiated. In 1996, the modular-vehicle concept was presented in Europe, and a standardised regulation could finally go into effect, paving the way for modular vehicles such as LCVs in the European Union [4, 5]. As the demand grows for higher efficiency and lower environmental impact in the transport sector, LCVs are expected to become more and more common. In countries such as Australia, Canada, and parts of the USA, similar types of heavy transports are already widely in use. With larger or heavier vehicles it may become harder for human drivers to handle the vehicles safely, partly because a larger vehicle might be more difficult for a driver to monitor and partly because the vehicle dynamics are more complex [6]. In order to handle these issues, automated ADAS functionalities are being considered and are thus widely researched. With ADAS it should be possible to increase the safety for all nearby vehicles, at the same time as the strain on the driver can be relieved and the vehicle can be handled more efficiently.

In Sweden, one example of where LCVs could be applied to great effect is highway driving between larger cities, where these large transports would deliver goods to and from warehouses or depots located just outside of the cities. By then letting smaller transport vehicles take over and deliver the goods to the actual destinations, such as shops inside the city, the needs of redesigning the road infrastructure to handle the larger transport vehicles would be reduced. The active safety systems of LCVs would also benefit from letting the large vehicles drive primarily on highways, since the potential traffic scenarios that can occur on highways often are significantly less complex than those that can occur inside cities, where driving rules can be complicated, traffic dense, and pedestrians and other objects can appear on the road. By looking at the accident statistics of current heavy vehicles in Sweden, it can be seen that one in four accidents are rear-end crashes and around 16 % of the accidents lane-change crashes [6]. Most of the lane-change accidents occur on highways, often because a smaller vehicle is obscured by the truck's cabin when the truck driver initiates the lane change. For the larger LCVs, it is possible that the increased truck length makes it even harder for the truck drivers to keep watch on the surrounding traffic, and the ADAS of the truck must be able to handle these situations safely.

2.1 Examples of automated driving

There have been a large amount of work on automated driving of vehicles in traffic. Software companies like Waymo (previously Google) and automotive manufacturers such as Audi, Volvo, and Daimler all put much effort into developing commercial solutions for autonomous vehicles. Despite the lack of details regarding the commercial solutions, some examples of the approaches are recounted here.

Waymo's self-driving cars are still in development but can already handle many complex situations [7] [8]. By fusing sensor data from LIDARs, radars, cameras, and GPS, and combining it with Google's map database, the cars create a virtual environment in which a path is planned. When objects, such as traffic cones, cyclists, or pedestrians are detected, the system predicts where they will likely be located a short time into the future. This results in the car slowing down when a pedestrian is expected to cross the road, or swerving around parked cars to maintain a safe distance. The image processing is able to detect when e.g. a cyclist makes a turning signal, making sure to give way and avoid potential collisions.

Audi's piloted driving uses an Audi A7 for highly automated highway driving, which among other things is able to change lanes on its own to overtake slower vehicles [9]. The system is still being developed and the goal is that a human driver should only need to handle the beginning and end of the journey, while the intermediate highway driving is managed by the car.

Tesla disrupted the automotive market with their *Autopilot* system. Already in use by consumers on the road, and with continuos software updates, cars with the *Autopilot* engaged keeps their lane and a dynamical speed depending on the surrounding traffic [10]. The system is able to make lane changes, but will only do so if the human driver indicates that a lane change is desired. All cars in the Tesla fleet continuously communicates information to the cloud and uses machine learning algorithms to improve the performance for the whole fleet in subsequent software updates. Although the system can do much with its relatively inexpensive equipment, it is currently only suited for highway driving and the human driver is expected to always be alert and ready to interfer in emergency situations. These caveats are also a reason for the criticism that Tesla's *Autopilot* system has endured: the driver might become too relaxed by the seemingly safe functionality and simply not be prepared to take over in an emergency. On May 7 2016 this happened in Florida when a Tesla driver died in a crash with a semitrailer truck that made a turn and crossed the road in front of the Tesla. Although the *Autopilot* system was active during the crash, Tesla was cleared after an investigation showed that the truck driver failed to give right of way, the Tesla driver was unattentive, and the system was not designed to handle such a scenario [11]. This pinpoints the problematic transition towards autonomous vehicles where the driver feels secure but still needs to be on the alert because not all traffic situations are covered by the automated functionality.

In 2013, Volvo Cars announced the *Drive Me* project in which a fleet of XC90 cars will be able to drive autonomously on designated roads in Gothenburg, Sweden, starting in 2017 [12] [13]. The goal of the project is primarily to explore the potential benefits of self-driving cars when it comes to areas such as safety, environment, and efficiency. The cars use radars, LIDARs, ultrasonic sensors, and cameras to sense the surroundings, and combine the data with map information accessed via the cloud. In August 2016, Volvo also announced a cooperation with the taxi company Uber in order to develop self-driving cars [14].

Mercedes-Benz are currently developing a *Highway Pilot* system for autonomous trucks which is planned to be ready for deployment in 2025 [15] [16]. The system is meant to help and relieve the strain of the human driver in monotonous situations or congested traffic. When human input is needed, the system will inform the driver in advance to take over control. The *Highway Pilot* features radars and cameras for environment sensation, as well as a system for vehicle-to-vehicle communication in the vicinity of the truck.

Yutong announced in 2015 that they had succeeded with a trial drive of an autonomous bus on public roads [17]. The bus managed to handle traffic lights, busy traffic, lane changes, and overtakings without human intervention. In order to do this, sensors such as LIDARs, radars, and cameras positioned all around the vehicle were used.

2.2 Previous academic research

In the academic world, there is as well much research going on, particularly in regards to how to model traffic situations and predicting how they will develop. Previous work on decision making in highway traffic have considered methods for handling situations such as heavy congestion. A non-trivial problem in dense traffic is lane changing or lane merging where the drivers are forced to interact actively with each other, as opposed to sparse traffic in which drivers passively can wait for gaps to appear. In [18], data was gathered on how human drivers behaved while navigating through a highway intersection in California. The authors could see that, in congested traffic, two distinct types of merge behaviours were prevalent: courtesy merges (when the drivers slowed down and helped each other), and forced merges (when there was no time to act cooperatively or the merging driver grew impatient). Based on the gathered data, a model could be constructed based on utility-driven decision making. The model is structured as a flow diagram or state machine so that a driver that uses it switches between different states based on what the surrounding situation looks like.

In [19], a method is proposed which utilises risk-driven planning to control a simulated vehicle. This is done by generating a number of different trajectories and computing estimates of the risk and utility of the alternatives to determine the best one. A similar, and slightly more complex, procedure is shown in [20]. Here, potential near-future plans are generated by using rapidly-exploring random trees. These plans are evaluated against predictive risk maps that describe how the situation is likely to develop, to determine how risky they are. A plan with a minimised risk and maximised utility is then chosen.

In order to be able to accurately predict risk or utility, the traffic surrounding the ego vehicle must be sufficiently realistic. Much work has been done on simulating traffic on different levels of detail. In macroscopic models it is assumed that traffic as a whole behaves similarly to e.g. a fluid stream. On the contrary, when simulating traffic on the level of individual vehicles, one talks about microscopic traffic models. A discussion about microscopic models and the vehicle interactions that often are included can be found in [21]. Some examples of vehicle interactions are car following, lane changes, overtakings, and speed adaptation. It is also mentioned that if realism is an important aspect of the modelling, different types of drivers such as aggressive, fatigued, or even drunk drivers have to be simulated as well. In [22], a game-theoretic approach to traffic simulation is presented where the drivers are assumed to minimise their perceived effort. This can mean avoiding unnecessary lane changes or deviations from a desired speed, but it also takes minimising risk and maximising comfort and efficiency into account. Different sets of parameters, such as preferred driving speed, are used to model different driving styles. Based on the perceived cost of actions, such as modifying the current speed or changing lanes, the simulated driver acts in traffic, according to the authors, in a realistic way.

The *Intelligent Driver Model* (IDM) is a microscopic traffic model that describes the acceleration of a modelled vehicle based on a set of parameters [23]. The IDM is a continuous function that depends on the velocity v_α, the approach rate (difference in velocity) Δv_α to the vehicle in front, and the spatial distance s_α to the vehicle in front. Configurable parameters for the IDM are the desired velocity v_0, desired minimum standstill gap s_0, desired temporal headway T, maximum acceleration a, comfortable deceleration b, and an acceleration exponent δ. The equation describing the IDM (see Equation 2.1) can be divided into two parts: a free-road term and an interaction term. The free-road term dominates the function when there are no nearby vehicles, causing the modelled vehicle to converge to the desired velocity over time. On the other hand, when there is a nearby vehicle to the front, the interaction term comes into action and makes sure that the modelled vehicle keeps a safe distance to the front. With the IDM, a realistic single-lane behaviour can be achieved, and by varying the parameters it is possible to simulate vehicles with individual behaviours while still using the same model.

$$\dot{v}_\alpha = \frac{dv_\alpha}{dt} = a\left(1 - \left(\frac{v_\alpha}{v_0}\right)^\delta - \left(\frac{s_0 + v_\alpha T}{s_\alpha} + \frac{v_\alpha \Delta v_\alpha}{2 s_\alpha \sqrt{ab}}\right)^2\right) \tag{2.1}$$

In [19], a model called the *Foresighted Driver Model* is used to model driving in traffic. In order to do this there is a need to coarsely predict how the traffic situation will develop, including the most important behavioural alternatives of all relevant traffic participants (for example all vehicles that have a chance of interacting with the ego vehicle). This is done by simulating the main possible trajectories of the ego vehicle and the surrounding traffic participants, as well as their maneouvring alternatives (such as braking, accelerating, or turning). If a vehicle's path is unknown, possible trajectories can be derived by considering the road geometry. For example, if a car approaches a four-way intersection its possible trajectories could be turning onto either side road or continuing straight. From this discrete set of paths it is possible to consider the main behaviours along these paths. Even while concentrating on only a small set of alternatives (the authors suggest constant-speed and emergency-braking) it was possible to capture predictions of the major risks. When computing the risk of each path of the ego vehicle, it is also possible to weigh the different alternatives to capture that vehicles are more likely to keep a constant velocity than suddenly braking. The authors also state that in this approach a very limited subset of the main behavioural alternatives is considered, and there are alternatives (for example [24]) which instead make use of sampling in order to try to capture more situations.

2.3 Existing framework for automated driving

This section explains the *Traffic Situation Management* (TSM) and simulation framework that was supplied by Volvo GTT and used as a base for the implementation in this thesis project. The majority of the framework was created during the work of [25] and consists of a C++ module that is called upon from a Simulink simulation model.

2.3.1 Driver models

The driver model that is used to control the longitudinal and lateral movement of the LCV is designed to resemble the behaviour of a human driver in non-emergency highway situations as well as possible. To do this, the model makes use of aimpoints and optical flow theory, as explained in better detail in [26]. In this model, the concept of a perceived near point and a perceived far point is central, and the distances and angles to these points are important parameters. Figure 2.1 gives a visual example of the model.

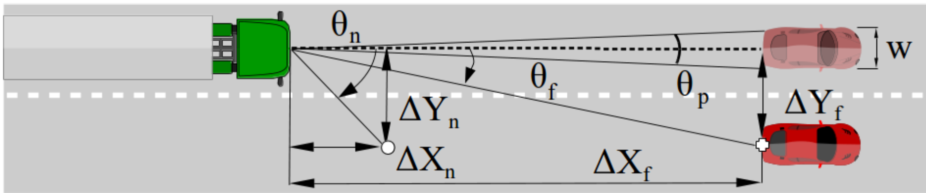

Figure 2.1. *Visualisation of the important parameters of the driver model. The two aimpoints, X_n and X_f, are shown as white circles. The parameters shown are the distances ΔX_n, ΔX_f, ΔY_n and ΔY_f to the aimpoints, the angles θ_n and θ_f to the aimpoints, the optical size θ_p, and the width W of the lead vehicle. Figure taken from [27].*

The longitudinal control is based on the time-to-collision and optical expansion rate of the lead vehicle as discussed by Lee in [28]. It is an iterative control model where the desired acceleration $a_{x,ref}$ is determined as follows:

$$a_{x,ref} = (1 + \dot{\tau}_m)\frac{\Delta v_x^2}{(\Delta X_L - v_{x,L} \cdot t_h)} \tag{2.2}$$

where $\dot{\tau}_m$ is the derivative of the time-to-collision, Δv_x is the velocity difference between the ego vehicle and the lead vehicle, ΔX_L is the distance from the first axle of the ego vehicle to the lead vehicle, $v_{x,L}$ is the velocity of the lead vehicle, and t_h is the desired temporal headway to the lead vehicle. The model is also subject to the constraints

$$\begin{aligned} \underline{a}_x \leq a_{x,ref} \leq \overline{a}_x \\ \underline{j}_x \leq j_x \leq \overline{j}_x \end{aligned} \tag{2.3}$$

that is, the acceleration $a_{x,ref}$ and longitudinal jerk j_x are to be kept within a predetermined interval.

The lateral control is based on the work of Salvucci and Gray [29] and utilises the two angles θ_n and θ_f (as shown in Figure 2.1) to accomplish the steering. The steering rate $\dot{\delta}_{ref}$ is formulated as

$$\dot{\delta}_{ref} = k_f \dot{\theta}_f + k_n \dot{\theta}_n + k_I \theta_n \tag{2.4}$$

This results in a lateral control algorithm that tries to maintain stable angles to the both aimpoints, while at the same time attempting to keep a zero-angle to the near point. This approach was used to good effect by Markkula et al [30] in their work on the modelling of human evasive driving behaviour.

2.3.2 Vehicle models

This section will describe the different vehicle models that were used in the project. For the ego vehicle, an A-double combination was used, as pictured in Figure 2.2. Due to the high complexity of the vehicle movement of an A-double LCV, the vehicle dynamics must be formulated and solved, something which was done in [31]. Below follows a brief explanation of how the prediction models for the ego LCV and surrounding vehicles function.

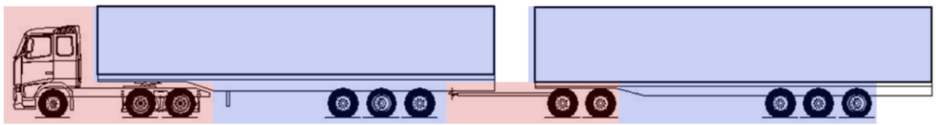

Figure 2.2. *Illustration of an A-double LCV combination. The different units of the combination are highlighted in different colors. Figure taken and adapted from [27].*

Ego vehicle prediction model

For the ego vehicle prediction, a one-track model of an A-double combination is used. Figure 2.3 shows the spatial parameters, motion variables and tyre forces of the vehicle model. Because a one-track model is used, the various forces affecting each axle are combined so that they can be modelled as only one single virtual tyre. When the motion of the ego vehicle has to be predicted for the tentative plans (as explained in Section 2.3.4), this simplified model is used as it is complex enough to capture the truck dynamics with sufficient accuracy, yet simple enough to be computed in real-time.

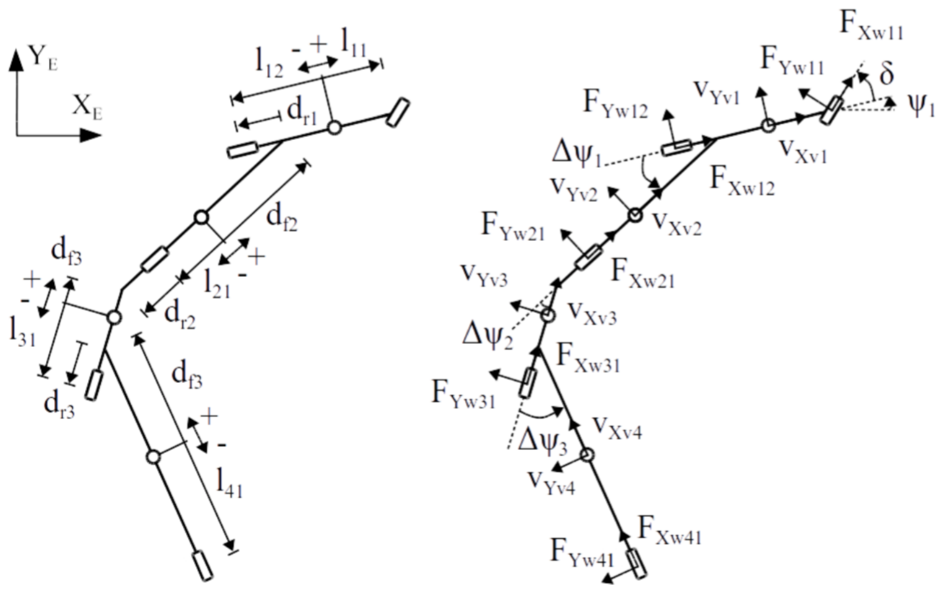

Figure 2.3. *The left part visualises the spatial parameters of the vehicle model, while the right part shows the included motion variables and tyre forces in the model. Figure taken from [27].*

The differential equations used in the ego vehicle model, as derived in [32], are formulated as follows:

$$
\begin{aligned}
\dot{z}_1 = 47.0 \cdot \delta - z_{10} \cdot z_3 + 1.9 \cdot z_4 + 0.9 \cdot z_6 - 0.002 \cdot z_8 + \\
\left(\frac{-70.7 \cdot z_1 + 9.7 \cdot z_3 + 21.7 \cdot z_5 + 4.5 \cdot z_7 - 0.02 \cdot z_9}{z_{10}} \right)
\end{aligned} \tag{2.5}
$$

$$
\dot{z}_2 = z_3 - \kappa_{R,1} \cdot (z_{10} \cdot \cos(z_2) - (z_1 + 1.5 \cdot z_3) \cdot \sin(z_2)) \tag{2.6}
$$

$$
\begin{aligned}
\dot{z}_3 = 25.0 \cdot \delta - 1.9 \cdot z_4 - 0.8 \cdot z_6 + 0.002 \cdot z_8 + \\
\left(\frac{27.6 \cdot z_1 - 174.2 \cdot z_3 - 20.8 \cdot z_5 - 4.3 \cdot z_7 + 0.02 \cdot z_9}{z_{10}} \right)
\end{aligned} \tag{2.7}
$$

$$
\dot{z}_4 = z_5 \tag{2.8}
$$

$$
\begin{aligned}
\dot{z}_5 = -25.5 \cdot \delta - 4.0 \cdot z_4 + 2.5 \cdot z_6 - 0.007 \cdot z_8 + \\
\left(\frac{-36.5 \cdot z_1 + 165.4 \cdot z_3 - 10.9 \cdot z_5 + 13.0 \cdot z_7 - 0.05 \cdot z_9}{z_{10}} \right)
\end{aligned} \tag{2.9}
$$

$$
\dot{z}_6 = z_7 \tag{2.10}
$$

$$
\begin{aligned}
\dot{z}_7 = 0.6 \cdot \delta + 2.3 \cdot z_4 - 22.9 \cdot z_6 - 0.9 \cdot z_8 + \\
\left(\frac{19.9 \cdot z_1 - 216.8 \cdot z_3 - 169.7 \cdot z_5 - 125.8 \cdot z_7 - 7.2 \cdot z_9}{z_{10}} \right)
\end{aligned} \tag{2.11}
$$

$$
\dot{z}_8 = z_{10} \tag{2.12}
$$

$$
\begin{aligned}
\dot{z}_9 = -0.19 \cdot \delta + 5.1 \cdot z_4 + 22.7 \cdot z_6 - 7.1 \cdot z_8 + \\
\left(\frac{-12.5 \cdot z_1 - 195.8 \cdot z_3 + 168.6 \cdot z_5 + 68.2 \cdot z_7 - 54.7 \cdot z_9}{z_{10}} \right)
\end{aligned} \tag{2.13}
$$

$$
\dot{z}_{10} = a_{x,1} \tag{2.14}
$$

$$
\dot{z}_{11} = \frac{a_{x,1,des} - a_{x,1}}{\tau} \tag{2.15}
$$

$$
\dot{z}_{12} = \frac{z_{10} \cdot \cos(z_2) - (z_1 + 1.5 \cdot z_3) \cdot \sin(z_2)}{1 - \kappa_{R,1} \cdot z_{13}} \tag{2.16}
$$

$$
\dot{z}_{13} = z_{10} \cdot \sin(z_2) + (z_1 + 1.5 \cdot z_3) \cdot \cos(z_2) \tag{2.17}
$$

$$\dot{z}_{14} = \frac{1}{(1 - \kappa_{R,4} \cdot z_{15})} \cdot \\ (-24.6 \cdot z_3 - 22.7 \cdot z_5 - 12.3 \cdot z_7 - 7.7 \cdot z_9 - z_4 \cdot z_{10} - z_6 \cdot z_{10} - z_8 \cdot z_{10} + z_1) \cdot \\ - \sin(z_2 + z_4 + z_6 + z_8 - \theta_{R,4} + \theta_{R,1}) + z_{10} \cdot \cos(z_2 + z_4 + z_6 + z_8 - \theta_{R,4} + \theta_{R,1}) \tag{2.18}$$

$$\dot{z}_{15} = (-24.6 \cdot z_3 - 22.7 \cdot z_5 - 12.3 \cdot z_7 - 7.7 \cdot z_9 - z_4 \cdot z_{10} - z_6 \cdot z_{10} - z_8 \cdot z_{10} + z_1) \cdot \\ \cos(z_2 + z_4 + z_6 + z_8 - \theta_{R,4} + \theta_{R,1}) + z_{10} \cdot \sin(z_2 + z_4 + z_6 + z_8 - \theta_{R,4} + \theta_{R,1}) \tag{2.19}$$

$$\dot{z}_{16} = \dot{\delta} \tag{2.20}$$

$$z = [\dot{y}_1, \phi, \dot{\phi}, \theta_1, \dot{\theta}_1, \theta_2, \dot{\theta}_2, \theta_3, \dot{\theta}_3, v_{x,1}, a_{x,1}, s_1, e_1, s_{11}, e_{11}, \dot{\delta}] \tag{2.21}$$

$$u = [a_{x,1,des}, \delta] \tag{2.22}$$

Here, $\dot{y}_1$ corresponds to the lateral velocity, while ϕ and $\dot{\phi}$ corresponds to the yaw angle and yaw angle rate of the vehicle's first axle. $\theta_1, \theta_2, \theta_3$ and $\dot{\theta}_1, \dot{\theta}_2, \dot{\theta}_3$ are the towed units' articulation angles and their corresponding rates, respectively. The first axle's longitudinal velocity and acceleration is modelled by $v_{x,1}$ and $a_{x,1}$. The distances and perpendicular distances of the first and last axles, as projected on the geometry of the road, are represented by s_1, e_1, s_{11}, e_{11}. The curvature and heading of the road, based on the first and last vehicle axles, are modelled by $\kappa_{R,1}, \kappa_{R,4}, \theta_{R,1}, \theta_{R,4}$. τ represents a time constant used for the longitudinal dynamics. The input to the vehicle model consists of u, that is the first axle's longitudinal acceleration $a_{x,1,des}$ and the steering angle δ in relation to the road.

Surrounding vehicle models

The surrounding vehicles are modelled using point-mass models in the longitudinal direction, along the road. Laterally, they are restricted to always stay in the exact center of the lane they are in. The vehicle model used in the simulation plant is configured to follow a prespecified acceleration profile, which can be set differently for each vehicle. For traffic predictions, on the other hand, a constant-velocity model (based on the vehicle's current velocity) is used for the range of the prediction horizon. The longitudinal movement can be expressed using the following equation

$$\frac{d}{dt}\begin{bmatrix} s \\ \dot{s} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} s \\ \dot{s} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \cdot \ddot{s} \tag{2.23}$$

where s is the vehicle's longitudinal position, $\dot{s}$ is the velocity, and $\ddot{s}$ is the acceleration along the road.

Ego vehicle plant model

In order to model the ego vehicle in the simulation plant, a high-fidelity two-track model is used. An illustration of the spatial parameters and motion variables can be seen in Figure 2.4. Comparing Figures 2.3 and 2.4 shows the difference in complexity between a one-track and a two-track model of the A-double. The simulation plant uses a library of proprietary high-fidelity vehicle models that is developed by Volvo. The models in this library take into account the detailed dynamics regarding many of the components of the vehicles, for example brakes, chassis, steering system, suspension, and tyres.

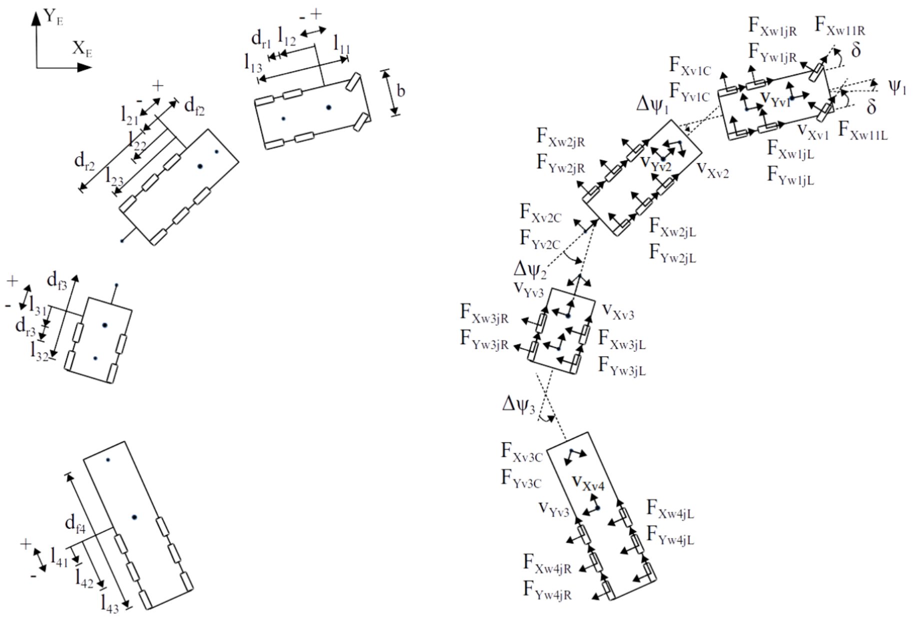

Figure 2.4. *Illustration of the spatial parameters and motion variables of a two-track model of an A-double LCV. Figure taken from [31].*

2.3.3 Road modelling

In order to simulate the vehicles in traffic and evaluate their behaviour, the road must be modelled in a clearly defined way. For this purpose, the road is represented as a number of clothoid segments. Each segment consists of a length as well as a constant curvature. By combining multiple segments of different lengths and curvatures, it is possible to achieve a smooth and flexible model of a road. Figure 2.5 illustrates the clothoid modelling of road segments.

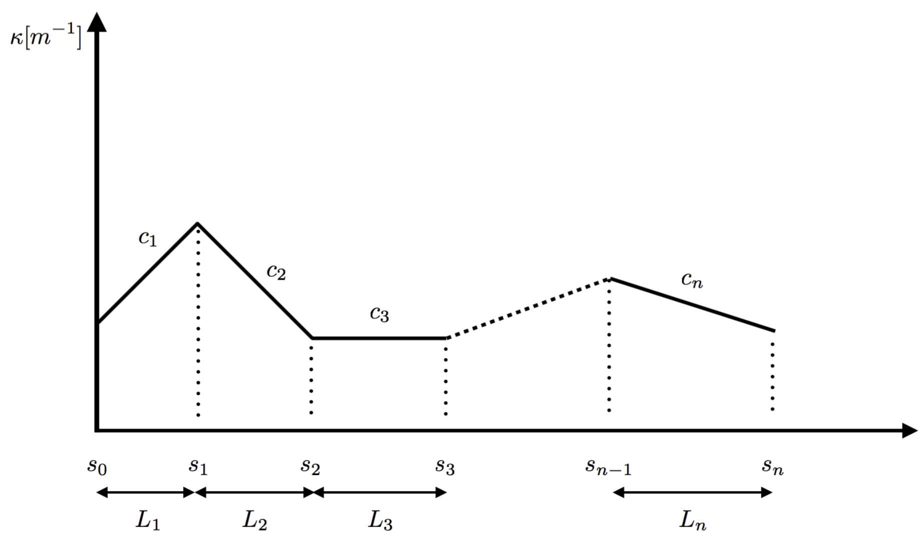

Figure 2.5. *An illustration of the clothoid-based road modelling. The road segments are shown on the horizontal axis, while each segment's curvature is shown on the vertical axis. Figure taken from [33].*

In order to use a coordinate system that is easily represented relative to the road, origo is set to the point along the road which is perpendicular to a fixed point on the truck. The orientation of the coordinate system is then set so that the longitudinal axis points along the road segment at origo and the lateral axis is perpendicular to the road, which implies that the previously mentioned fixed point on the truck always has a longitudinal position of zero. Lanes are modelled simply as lateral offsets along the road profile. A lane holds information, such as width, length, curvature, lane index, and adjacent lanes, and each modelled surrounding vehicle is aware of the lane it is driving in. In the simulations, the road was only modelled in two dimensions, where height was excluded.

2.3.4 Traffic situation predictions

Using the components described above, it is now possible to form an understanding of the traffic situation which is complete enough to form *traffic situation predictions* (TSPs). In the simulation framework, TSPs are formed for the current and directly adjacent lanes, resulting in three TSPs to be evaluated. Each prediction is simulated for a time horizon of $t = 3.5$ s, and during this time the surrounding traffic is simply predicted as moving along their lanes with a constant velocity. The integration of each surrounding vehicle's position is done using forward-Euler integration with a time step of $h = 0.05$ s. With the information of where the surrounding vehicles are predicted to be located for each time instant during the TSP, it is possible to form a plan for the ego vehicle. At each integration step, the vehicle movement plan is tested against a set of constraints to make sure that the vehicle is not exceeding any limits related to vehicle dynamics as well as respectecting the lane boundaries.

$$\underline{v}_x \leq v_{x,1} \leq \overline{v}_x \tag{2.24}$$

$$\underline{a}_x \leq a_{x,1} \leq \overline{a}_x \tag{2.25}$$

$$\underline{a}_x \leq a_{x,4} \leq \overline{a}_x \tag{2.26}$$

$$\underline{d} \leq d_1 \leq \overline{d} \tag{2.27}$$

$$\underline{d} \leq d_4 \leq \overline{d} \tag{2.28}$$

Here, $\underline{v}_x$ and $\overline{v}_x$ correspond to the minimum and maximum speed limits while $v_{x,1}$ represents the velocity of the LCV's first axle. The limits on the lateral acceleration is similarly expressed as $\underline{a}_x$ and $\overline{a}_x$, while $a_{x,1}$ and $a_{x,4}$ correspond to the actual lateral acceleration of the first and the fourth axles. $\underline{d}$ and $\overline{d}$ correspond to the limits in lateral positioning based on an offset from the center lane, and d_1 and d_4 correspond to the lateral position of the first and the fourth axles.

In addition to these five constraints there is also one for collision detections, in order to invalidate a tentative plan which causes the ego vehicle to collide with the surrounding traffic. This collision detection is based on creating bounding boxes around the LCV and the surrounding vehicles, and then checking if there are any intersections between the bounding boxes. If any of the constraints are violated during the closed-loop generation of the tentative plans, the plan is marked as infeasible [27].

The input to the LCV's driver model control is based on the state of the ego vehicle as well as the observations of the traffic environment. In addition, it is possible for an external signal (e.g. a human driver) to request a lane change in either direction. Based on this, the different TSPs are set up and evaluated by the driver model control. The tentative plans are evaluated for the predictions' whole time horizon. Based on the goals of the ego vehicle (e.g. if a lane change should be performed) and potential constraint violations of the TSPs, a tentative plan is selected to follow. From this tentative plan, the first iteration's longitudinal acceleration ($a_{x,des}$) and steering wheel rate ($\dot{\delta}$) are used as input to the vehicle motion management, which uses them to step the actual simulation of the vehicle plant. The model that is used in the simulation plant can either be the two-track high-fidelity model mentioned in Section 2.3.2, or a simpler one-track model. A high-level illustration of the ego vehicle's control hierarchy is shown in Figure 2.6.

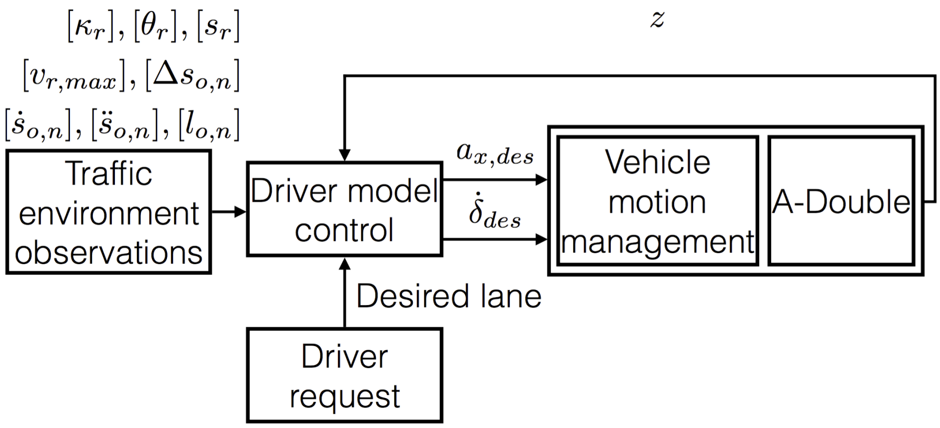

Figure 2.6. *A high-level illustration of the design of the automated driving framework. The driver model control receives inputs from its own state, the traffic situation, and potentially and external lane change signal. From the evaluated plans, values for longitudinal acceleration and steering wheel rate are extracted and sent to the vehicle motion management, which steps the simulation framework. Figure taken from [33].*

3 Method

The functionality described in this thesis is implemented and tested using the simulation framework described in Section 2.3. In this section, the modifications made to the Traffic Situation Management (TSM) module is explained, as well as the behaviour of the surrounding traffic used in the simulation environment and the traffic-prediction model utilised by the ego vehicle. Finally, the scenarios that are used to test the functionality of this thesis are presented and explained.

3.1 Modifications to the Traffic Situation Management

When the ego vehicle signals to the surrounding traffic that it wants to do a lane change, the turning indicators are the main means of communication. A second, more aggressive, way of signalling the intention of making a lane change is to move closer to the target lane without actually crossing the line between the two lanes. A driver who might have missed or even ignored the turning signals should be more likely to react if they see that a large vehicle is actually moving closer to them. This functionality was implemented in the TSM so that the LCV activates its turning indicators when it wants to change lane. After the turning indicators have been active for five seconds the LCV adds a lateral offset to the far aimpoint, so that the LCV gradually moves closer to the target lane without moving across the line. As explained in Chapter 2.3, the simulation framework has previously based its lane-changing decisions mainly on the available spatial and temporal gaps between the leading and trailing vehicles in the target lane. In order to test the traffic situations that this thesis explores, lane changes in dense highway traffic, the restrictions on temporal and spatial gaps that had previously been used in the TSM were removed. Instead, the LCV is allowed to change lanes as soon as the traffic predictions (explained in more detail in Section 3.3) say that it is possible without causing a collision.

As is discussed in Section 3.4, a test scenario including a lane change on a limited distance is performed. In order to accomplish this, a speed-regulating mechanism had to be introduced since there was no decision-making logic already implemented in the TSM that could modify the speed while performing lane changes. This speed-regulation feature can easily be turned on and off. The motivation behind the mechanism is that lane changes in highway traffic often take a certain time to complete, normally between 5-12 seconds [34]. By counting the number of lane changes that are needed to get to the desired lane and estimating the time it takes to perform one lane change, it is possible to compute the maximum velocity that can be kept in order to have time to perform all lane changes before reaching a certain position, such as a highway exit ramp. The speed-regulating mechanism will therefore reduce the velocity of the LCV as it gets closer to a point where it should be located in a certain lane, giving the controller a way to perform lane changes even in dense traffic. The estimated maximum velocity $\overline{v}$ is described by Equation 3.1.

$$\overline{v} = \min(\overline{v}^{(d)}, \overline{v}^{(r)}) \tag{3.1}$$

where $\overline{v}^{(r)}$ is the maximum allowed speed on road segment r, and

$$\overline{v}^{(d)} = \frac{d_r}{t_{LC} \times k} \tag{3.2}$$

where, in turn, d_r is the distance that remains until the lane change must be completed, t_{LC} is an estimation of the upper limit of the time a lane change takes, and k is the number of lane changes that must be performed before the ego vehicle is located in the target lane. In the implementation, t_{LC} was set to 20 seconds. This gave reasonable results and did not cause the controller to decrease the speed too suddenly.

3.2 Traffic simulation

The behaviour of the simulated surrounding traffic is implemented in the Simulink part of the simulation framework. In order to achieve a more reactive and realistic behaviour than the prespecified acceleration profiles that were used previously, the simulated traffic was changed to use a customised version of the Intelligent Driver Model (which was explained in Section 2.2). The modifications are mainly used to let the simulated traffic react on the intention signals of the ego LCV, for example another vehicle's turning indicators or lateral position on the road. The longitudinal acceleration of a simulated vehicle is controlled by the two terms of the IDM (see Equation 2.1), so that the freeway term (which aims to reach the desired speed) is active at all times while the interaction term (which keeps the distance to other vehicles) is active only when there is another vehicle in front. In the traffic simulation, it is also possible to configure individual vehicles to behave cooperatively and try to help other vehicles complete their lane changes. A longer explanation of cooperativeness is found below. This was done so that if a simulated vehicle is set to be cooperative, the interaction term of the IDM also gets activated when the simulated vehicle detects another vehicle that wants to merge into the lane. An example of this would be when the ego LCV uses its turning indicators to show that it intends to perform a lane change. When the interaction term is activated by a detected lane change, the IDM uses the rear end of the merging vehicle as the point which to relate to (see Figure 3.1).

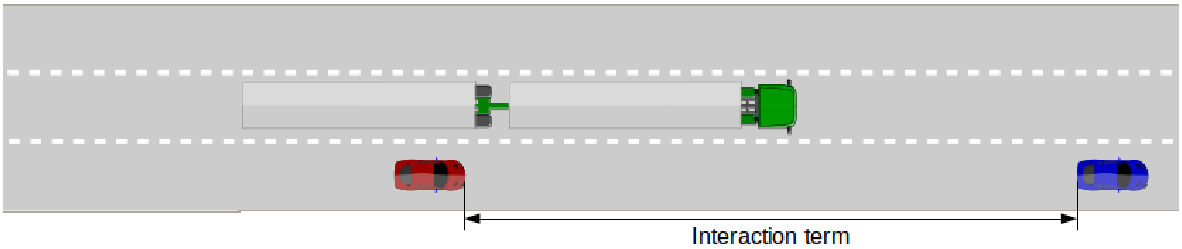

(a) *During normal driving, the interaction term of the red vehicle will depend only on the closest vehicle to the front in the same lane, which in the displayed situation is the blue car.*

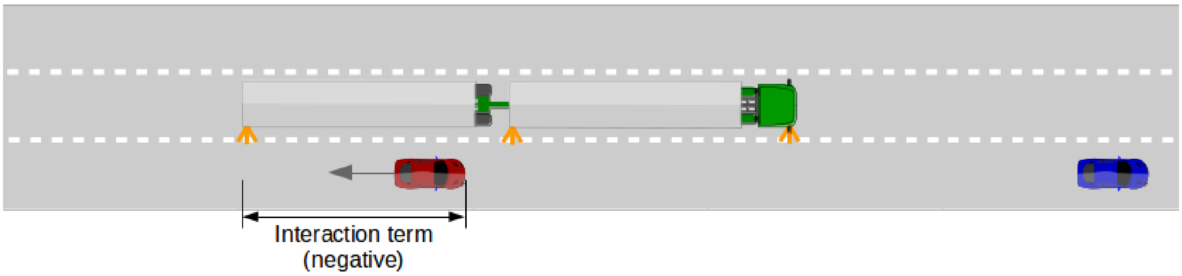

(b) *When the red vehicle has detected that another vehicle (the LCV) wants to change lane, the interaction term instead is set to depend on the end of the merging vehicle. Since the distance in the interaction term in this case becomes negative, a strong deceleration force causes the red vehicle to brake and let the LCV enter the lane.*

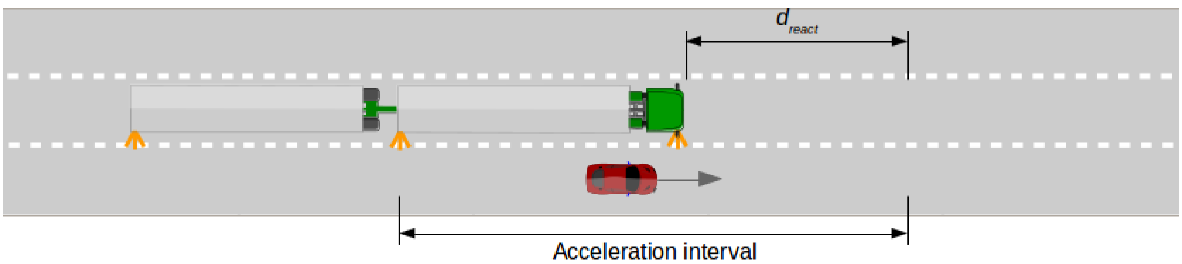

(c) *When the red vehicle is further to the front than half the merging LCV, it accelerates to the front of the LCV to make space.*

Figure 3.1. *Modifications made to the IDM model to enable reactions on intention signals.*

In order to achieve a more visually realistic behaviour of the simulated vehicles that give way to a merging vehicle, the model described above was also modified so that if a vehicle was further to the front than half the merging vehicle it would instead speed up and move forward instead of backward. This can be exemplified by a vehicle positioned beside the tractor of the LCV (as shown in Figure 3.1c). Most human drivers would in this situation not brake and let the LCV merge into the lane in front of them, but instead accelerate to try and give the LCV space behind them. This was implemented so that if the simulated vehicle is in the affected area (denoted *acceleration interval* in Figure 3.1c), an extra acceleration is used instead of the IDM's interaction term.

The extra acceleration a_e forward is determined by:

$$a_e = a_f \left(1 - \min\left(\max\left(\frac{s_v - s_{ego}}{d_f}, 0\right), 1\right)\right) \tag{3.3}$$

Here, a_f is the maximum possible acceleration forward, s_v is the longitudinal position of the simulated vehicle, s_{ego} is the longitudinal position of the merging vehicle (in our case the LCV), and d_f is the distance from the front of the merging vehicle to the front of the acceleration interval. The resulting behaviour is that $a_e = a_f$ when the simulated vehicle is behind the merging vehicle, and that a_e decreases linearly to 0 as the simulated vehicle gets closer to the front end of the acceleration interval. In the implementation, the constants $a_f = 2[m/s^2]$ and $d_f = 13[m]$ were used.

The model for reacting on turning signals is relatively simple in the current implementation. Depending on the chosen simulator settings, each simulated vehicle is either cooperative or non-cooperative. A non-cooperative vehicle will simply ignore any other vehicle that is not in its own lane. On the other hand, a vehicle that is cooperative may react to a signalling vehicle in the adjacent lane after it has detected the signals. The detection of turning signals can be made as complex as one would want, but for the purpose of this thesis, a cooperative vehicle is said to have detected a turning signal if it has been active for two seconds. As described above, a vehicle that reacts will try to move out of the way to make place for the merging vehicle, while non-cooperative vehicles will ignore any merge requests and just continue with business as usual.

3.3 Traffic predictions

The problem of predicting traffic in the near future is a very complex task. It is possible for a large amount of potential scenarios to play out, and any implemented model must be abstracted and simplified to some extent. Because of how the trajectory generation in the simulation framework functions, as described in Section 2.3.4, there are limitatations on how the traffic predictions can be modelled. The simulation framework requires two specified locations in order to generate a trajectory, and this means that the predicted traffic cannot be modelled in a purely probabilistic fashion, as is done in some of the research described in Section 2.2. Instead, for each generated trajectory, there must be a corresponding determined and non-probabilistic traffic prediction. This means that in order to consider various potential scenarios, each corresponding traffic prediction must be described and determined beforehand. In turn, this means that there are limitations on how uncertainty can be modelled in the traffic predictions, and also on the number of potential scenarios that can be considered since even a small scenario variation would require one more relatively time-consuming evaluation.

In order to partly counter this, the traffic predictions are modelled using a concept of dependent and independent drivers. The term *independent* is used for a driver that is positioned far enough, spatially and temporally, from another vehicle that their behaviour is assumed to not be significantly affected by the actions of the other vehicle. On the contrary, a *dependent* driver is one that is close enough for their actions to be influenced by those of another vehicle. This distance will hereafter be denoted as the *dependency distance*. An illustration of independent and dependent drivers can be seen in Figure 3.2. In this project, a dependent driver is assumed to follow the Intelligent Driver Model described in Section 2.2, while an independent driver is simulated by constant-acceleration models. The prediction simulations are implemented so that if an independent driver moves into dependency distance of another vehicle, they are regarded as a dependent driver for the rest of the prediction simulation. This way, a potential collision between two vehicles, that at the start of the prediction were simulated with constant-acceleration models, will be much more unlikely since the IDM will try to keep the vehicles from colliding.

In a prediction simulation, an independent driver is modelled by two constant-acceleration models: It can either keep its current acceleration, which we call the *expected scenario*, or it can brake heavily, which we in turn call the *emergency scenario*. For the emergency scenario, two different types of braking are considered in this thesis. The first is a constant braking, for example -4 m/s^2, which is predicted in the emergency scenario no matter what the vehicle's current acceleration is. The second braking type is a relative braking, where the predicted emergency scenario considers the vehicle's current acceleration minus a constant braking deceleration. The difference between the two is that with a constant braking of say -4 m/s^2 and a current acceleration of -1 m/s^2, the emergency scenario would condider a deceleration of -4 m/s^2. On the other hand, with a relative braking of -4 m/s^2 and a current acceleration of -1 m/s^2, the emergency scenario would use a deceleration of $-1-4=-5$ m/s^2 (see Figure 3.3). Based on the assumption that a vehicle's acceleration at a specific moment normally is similar to the acceleration it had the moment before, the second model could be more realistic. It is, however, not confirmed whether this is actually the case. Even though one might argue that it could be dangerous to consider a relative emergency braking when it is known that a vehicle has the ability to brake harder than it, this problem should be mitigated by the fact that these traffic predictions are executed so often in the simulation framework that any sudden braking in one time step would quickly be accounted for in the next time step's prediction.

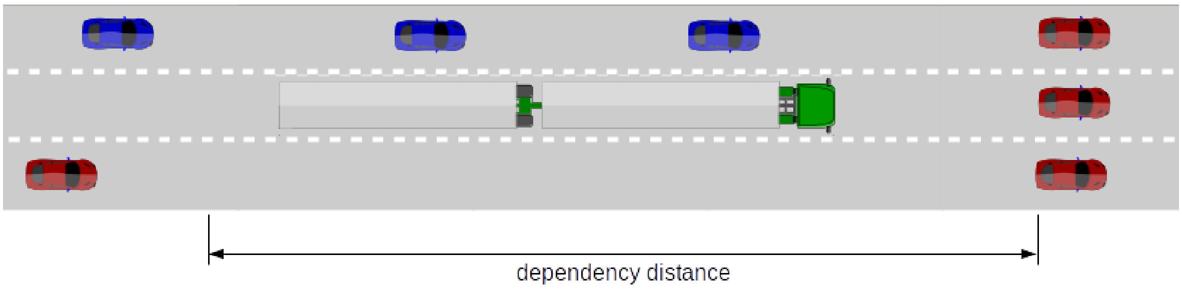

Figure 3.2. *An illustration of dependent and independent drivers. A red car represents an independent driver, while a blue car represents a dependent driver.*

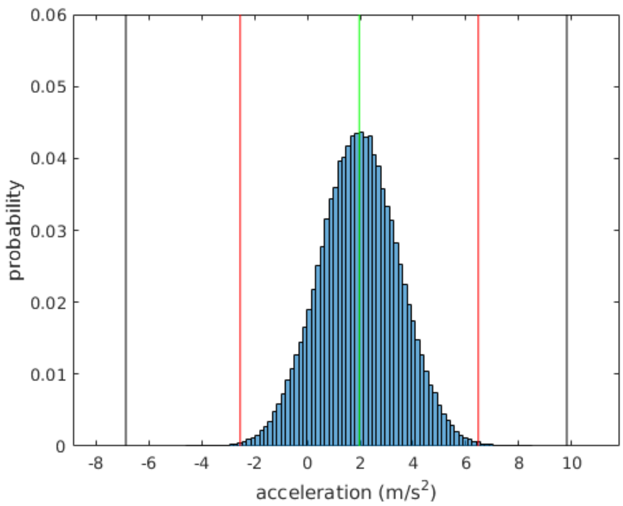

Figure 3.3. *An example of a hypothesised model for how the acceleration of a vehicle might change in a certain moment. It is not yet confirmed whether such a model is realistic or not. By considering an unlikely, but not impossible, acceleration (left-most red line) as what could happen in a predicted emergency situation, it is believed that the subsequent decisions would become safer. The green line represents the vehicles current acceleration in a given instant (in this example $2m/s^2$), the red lines correspond to the 99-percent confidence interval (three standard deviations), and the grey lines would be the vehicle's maximum and minimum acceleration limits.*

If we were to consider an expected and an emergency scenario for each vehicle in the vicinity of our ego vehicle and look at all the possible combinations that could occur, we would have a problem with exponentially increasing complexity. For N vehicles there would be 2^N possible combinations, and when looking at situations with dense traffic, this would quickly escalate to a problem where it is extremely hard to evaluate all possible predictions in real-time. For the relatively simple situation shown in Figure 3.2, one would have to evaluate $2^7 = 128$ different prediction possibilities, something that would not be possible to accomplish here in real-time. However, by only looking at the two different possibilities for the independent drivers and letting the dependent drivers be controlled by a deterministic model like the IDM, it is possible to simplify the problem somewhat. For I independent and D dependent drivers, the number of possible combinations would be $2^I \times 1^D = 2^I$. Once again, considering Figure 3.2 with the displayed number of dependent and independent drivers, we would have to evaluate $2^4 = 16$ different prediction possibilities. Even though the problem has been simplified greatly at this stage, it would still not be possible to perform this number of evaluations in real-time in the simulation framework. Assuming a dependency distance that is not too small, a worst-case situation could imply three independent drivers to the front, three to the rear, and one to either side. In this worst-case situation, there would be $2^8 = 256$ different combinations, even with the simplifications described above. With an improved trajectory generation algorithm and better hardware, this might however be possible in the future. In this thesis, one more simplification is needed to reduce the number of possible traffic predictions and thus allow for real-time evaluations. By not considering different combinations and instead only looking at one expected scenario for all vehicles and one emergency scenario for all vehicles, it is possible to reduce the problem to only two traffic predictions no matter how many surrounding vehicles there are. The implications of these simplifications are discussed in Section 5.

In order to generate the traffic predictions, the actions of the ego vehicle must also be considered since the vehicles that are predicted must be able to interact with it. This creates a cyclic dependency, since the ego vehicle's tractory generation depends on a generated traffic prediction, and the generation of traffic predictions requires a trajectory for the ego vehicle. This was solved by using the previous iteration's ego vehicle trajectory when generating the traffic predictions, and then generate the current iteration's ego vehicle trajectory from those predictions. For the first iteration when we have no computed ego vehicle trajectory, we instead use a constant-velocity trajectory for the ego vehicle while generating the traffic predictions.

3.4 Test scenarios

The situation that will be used in order to test the findings in this thesis involves a LCV driving in dense highway traffic. From the initial state, which is presented in Figure 3.4, three different scenarios will be run where the surrounding traffic behaves differently in each scenario. The main goal with these scenarios is to simulate a lane change to the right, performed by the LCV, and to safely be able to either complete a lane change, or abort it if the surrounding vehicles are not cooperative enough to allow the LCV to switch lane. In all scenarios the road speed limit is set to 80 km/h, and the vehicles (both surrounding traffic and the ego vehicle) move at approximately this speed.

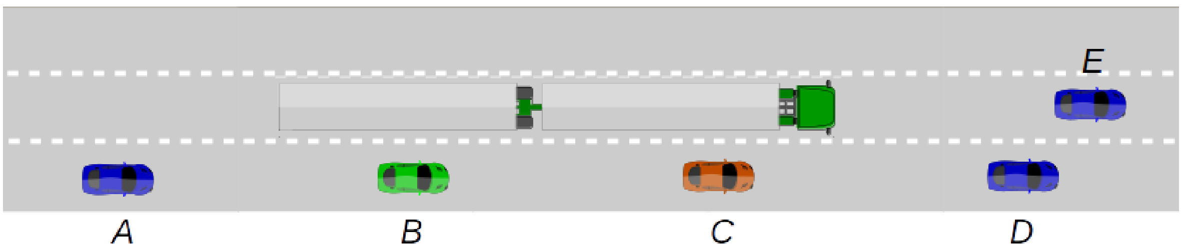

Figure 3.4. *The initial traffic situation where the LCV is faced by dense traffic on the right side.*

3.4.1 Scenario 1

The first scenario will start with the driver of a LCV requesting a lane change to the right by turning on the right turning signal. The LCV will then search for a possible way to merge into the right lane, and after some time also strongly indicate that it wants to change lane by moving itself closer to the edge of its current lane. The surrounding vehicles are, in this first scenario, modelled to be cooperative and create a gap for the LCV when they see the signal. They will react to the turning signal and the lateral offset it creates. Figure 3.5a shows one possible interaction between the LCV and vehicle B and C. Vehicle B decreases the speed and vehicle C continues driving, leading to a gap opening for the LCV. Other possibilities are that both vehicle B and C decrease or increase their speed creating the desired gap for the LCV. The scenario is intended to demonstrate the system's request at performing a lane change and indicating it to the surrounding traffic.

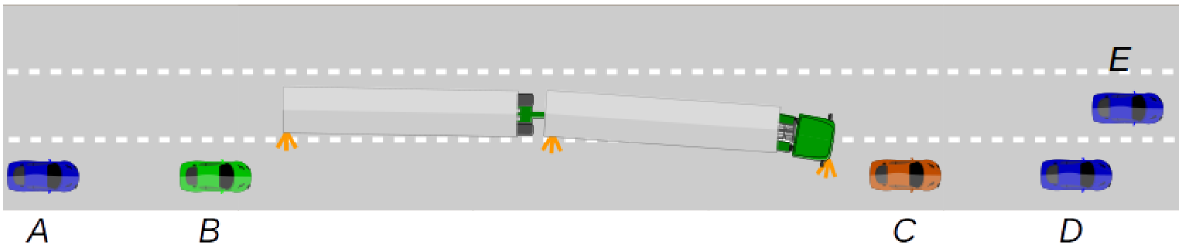

(a) *Possibility 1: Cars B and C have split up to allow for a lane change.*

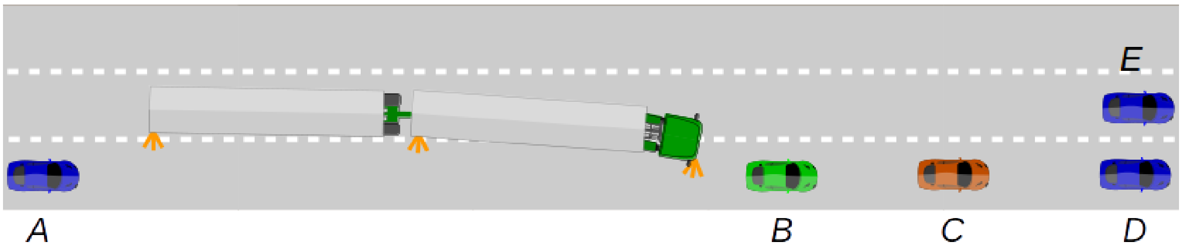

(b) *Possibility 2: Cars B and C have moved forward to make way for a lane change.*

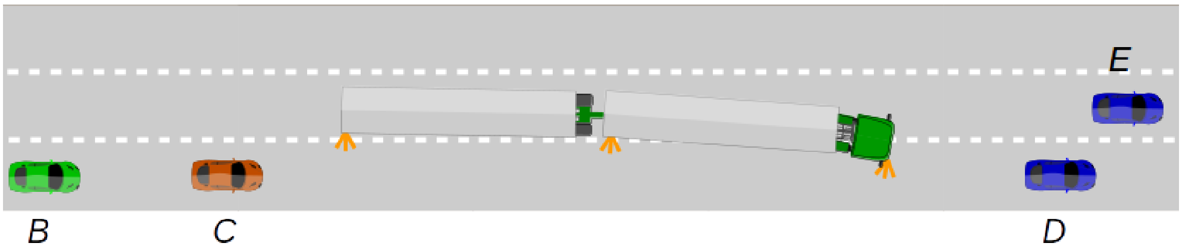

(c) *Possibility 3: Cars B and C have moved backward to make way for a lane change.*

Figure 3.5. *Scenario 1: The surrounding vehicles are cooperative and make way for the LCV.*

3.4.2 Scenario 2

This scenario also starts with a driver requesting a lane change to the right and the LCV behaving in the same way as explained in Section 3.4.1. However, this time the vehicles will be non-cooperative and no feasible gap will be created for the LCV, as in Figure 3.6. The result is that the LCV will have to abort any attempt to change lane since one or more vehicles will be blocking the manoeuvre. This scenario is intended to test the functionality of aborting a dangerous manoeuvre in favor of a safe one.

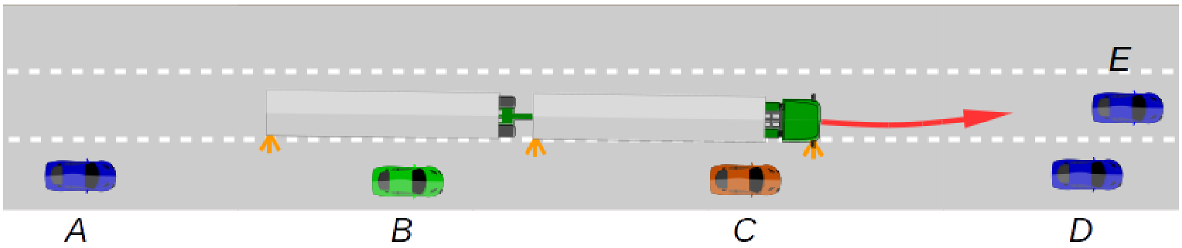

Figure 3.6. *Scenario 2: The LCV has started a lane change, but since the adjacent cars have not moved out of the way a safe abortion maneuver has to be performed.*

3.4.3 Scenario 3

The last scenario includes a hard constraint in that the LCV aims to exit the highway using a specific exit ramp, and thus the LCV has to complete the lane change within a given spatial frame. In this scenario, the longitudinal spatial limit d_{exit} at which the lane change must be completed is set to 1000 meters. The surrounding vehicles should be modelled in a realistic way so that they are self preserving and try to avoid collisions, even though they may not be completely cooperative and want to help the LCV in every situation. The intention with the last scenario is to try out the system in a more realistic environment, where surrounding traffic is not strictly cooperative or non-cooperative, and where a lane change might have to be performed under pressure.

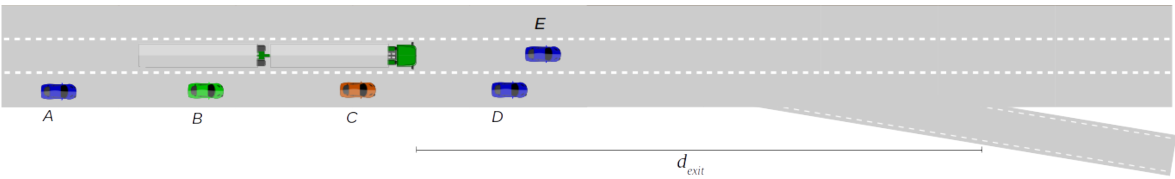

Figure 3.7. *Scenario 3: The goal of the LCV is now to exit the highway from a specific exit ramp, and must thus perform the lane change within a certain distance d_{exit}.*

4 Results

This chapter contains the results from running the various traffic scenarios presented in Section 3.4 in the simulation framework.

4.1 Scenario 1

When simulating scenario 1, it was noted that the way in which vehicles move forward to create a gap for the LCV (see explanation of method in Section 3.2) regularly caused the LCV to initiate a lane change only to abort it a short while later. Because of this problem, the relative emergency deceleration discussed in Section 3.3 was switched to a model where a constant emergency deceleration was used instead. In the moment when the controller initiated a lane change, the temporal headway to the front vehicle in the target lane was recorded. The front temporal headway is plotted against the success or failure (abort) of the lane change in Figure 4.1. To increase the readability of the graph, all data points have been moved a random distance along the x axis to spread them out.

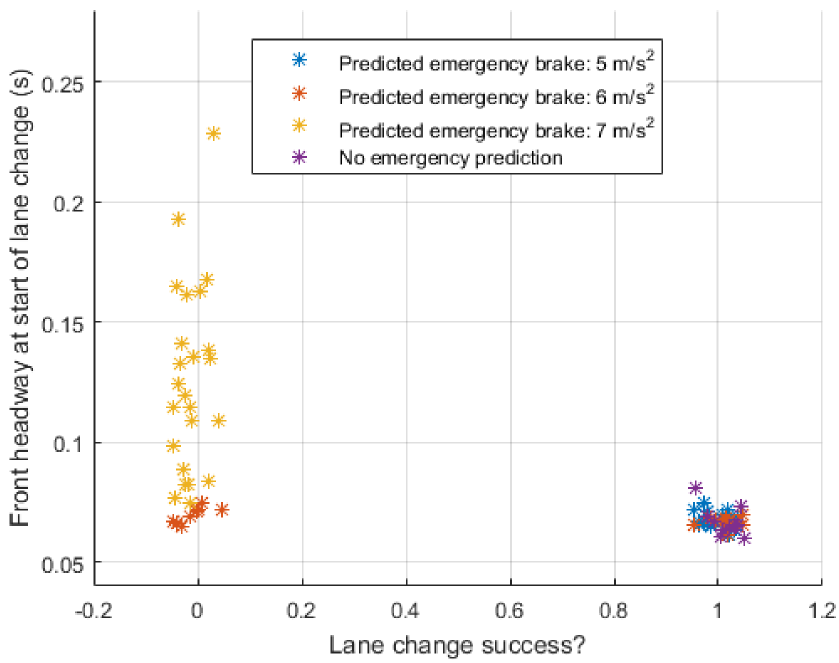

Figure 4.1. *The temporal headway to the front vehicle in the target lane at the initiation of the lane change plotted against whether the lane change was a success or not. A successful lane change is denoted by a 1, while an initiated and aborted lane change is denoted by a 0. In order to improve readability, all data points have been translated a short random distance* $[-0.05, 0.05]$ *along the x axis.*

4.2 Scenario 2

Independent of parameter settings, there was no case in which the simulated LCV managed to complete a lane change when the surrounding traffic was non-cooperative, nor were any collisions with the surrounding traffic recorded. When the lane-changing initiation signal was received by the LCV it activated its turning indicators, and after five seconds moved closer laterally to the target lane. Since the surrounding traffic never created a gap, the controller never took the decision to initiate a lane change. The decision-making system did not take any other action, such as decreasing the speed of the LCV, in order to create a gap on its own.

4.3 Scenario 3

When a lane change was finished, the distance that remained to the position where the lane change had to be finished was recorded. The lowest velocity recorded during the lane change was also noted. These results can be seen as histograms in Figure 4.2, where Figure 4.2a shows the remaining distance and Figure 4.2b shows the lowest velocity.

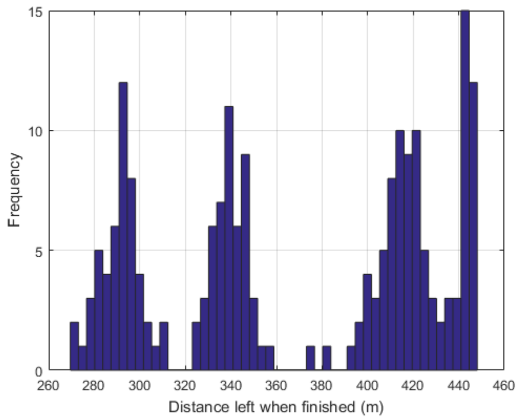

(a) *Remaining distance after lane change is completed.*

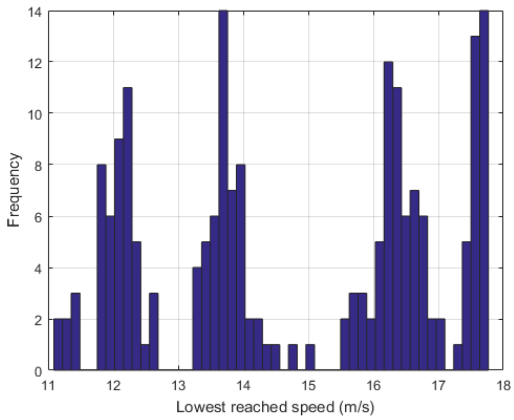

(b) *Lowest velocity reached during the lane change.*

Figure 4.2. *Histograms of the lowest velocity reached during the lane change, and the remaining distance to the virtual exit ramp after the lane change has completed.*

The result after plotting the lowest recorded speed against the remaining distance can be seen in Figure 4.3. The colours represent the different parameter settings that were used for the fixed emergency braking. By changing the colours to only represent the usage and the non-usage of the emergency-braking scenario, the result in Figure 4.4 is achieved.

Due to the highly linear relationship between the lowest recorded speed and the remaining distance, it is possible to visualise the data in an easier way by plotting the emergency brake setting against the lowest recorded speed, without losing much information. This is displayed in Figure 4.5. Here, not using the emergency scenario is shown as having an emergency braking of 0 m/s^2. The data points that make up the cluster visible in the top right of Figure 4.3 (where lowest recorded velocity was higher than 17.1 m/s), can be broken up in simulations that considered emergency predictions vs those that did not. The result is shown in Table 4.1.

Table 4.1. Breakdown of the simulation runs based on whether they utilised emergency scenario predictions or not, and whether they were part of the cluster in the top right of Figure 4.3 or not.

	Emergency prediction	No emergency prediction
Lowest recorded speed > 17.1 m/s	23	10
Lowest recorded speed ≤ 17.1 m/s	142	21
Total count	165	31
Part exceeding 17.1 m/s	13.94 %	32.26 %

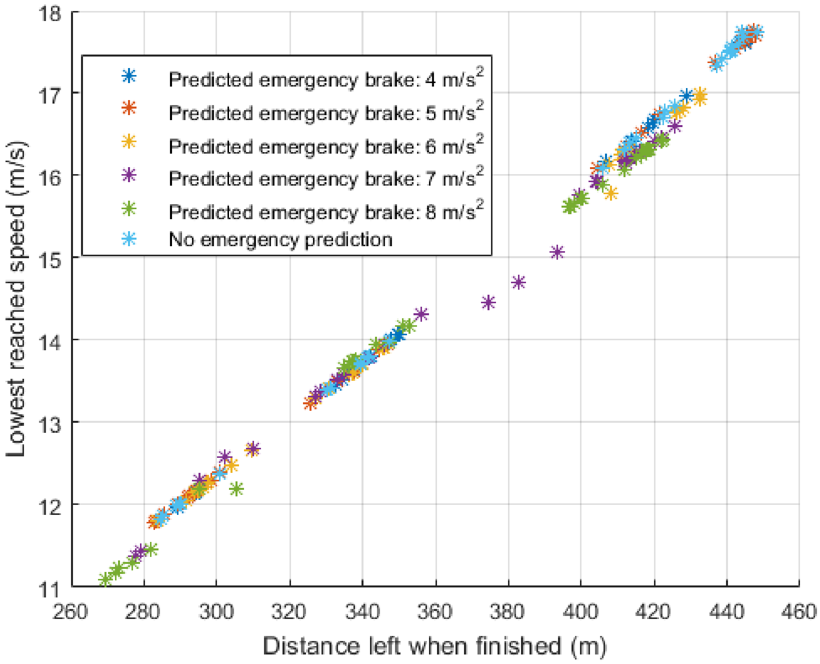

Figure 4.3. *The content of the histograms in Figure 4.2 plotted in relation to each other. Each parameter setting of the emergency brake deceleration is plotted in its own colour.*

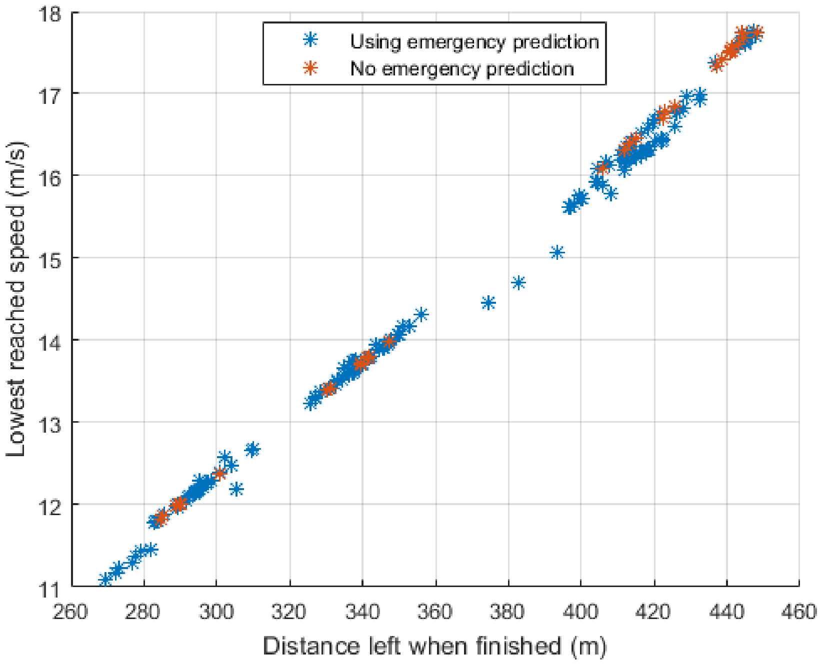

Figure 4.4. *The content of the histograms in Figure 4.2 plotted in relation to each other. Every simulation where an emergency brake scenario was considered is plotted in blue, while the red points denote the simulations where no emergency scenario was contemplated.*

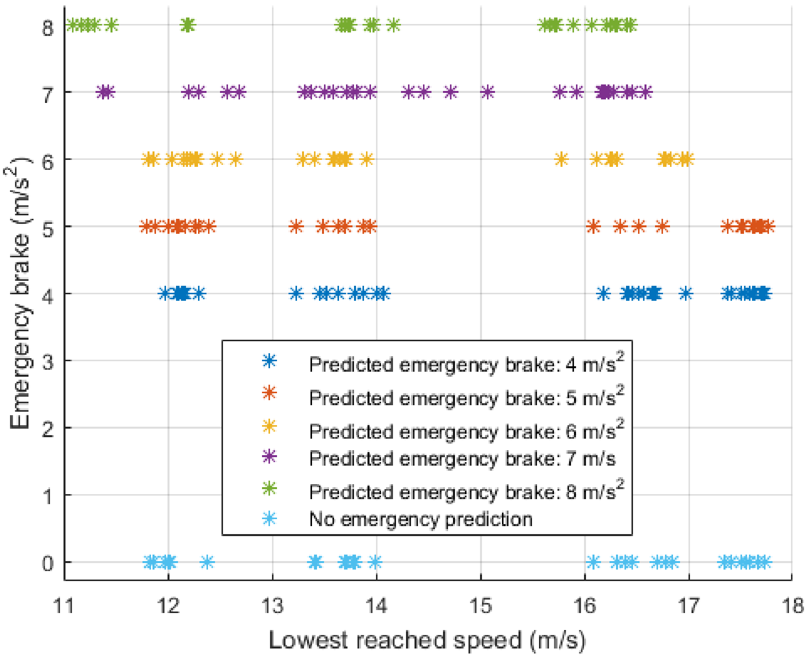

Figure 4.5. *Emergency brake parameter setting plotted against the lowest speed reached during the lane change. Simulations where no emergency braking scenario was considered are depicted as having an emergency brake deceleration of 0 m/s*^{2}*.*

5 Discussion

While running the simulation framework to gather results, it was noted that parts of the simulation model that was used for the surrounding vehicles were too simple: When cooperative vehicles accelerated in order to create a gap for the LCV, an extra acceleration term (described by Equation 3.3) was simply used instead of the interaction term in the IDM model. Since this extra acceleration term is activated only when the cooperative vehicle is located in a certain interval (the *acceleration interval* in Figure 3.1c) and the IDM interaction term is used otherwise, the fact that those two terms can be very dissimilar causes the cooperative vehicle to behave in an unnatural way near the front end of the *acceleration interval*. Consider for example the situation where a simulated cooperative vehicle tries to create a gap but there is another vehicle to its front. While inside the *acceleration interval* the vehicle will be affected by the extra acceleration term causing it to move closer to the front vehicle, but as soon as the interval is exited the IDM interaction term takes over and causes a deceleration because the front vehicle is closer than desired. This deceleration causes the cooperative vehicle to enter the *acceleration interval* again, and the process starts all over again, entering a loop (until the traffic situation changes). Since the LCV makes use of traffic predictions that are based on the surrounding vehicles' accelerations, the sudden jumps between acceleration and deceleration causes the LCV to deem a plan as feasible in one iteration and infeasible in the next, which in turn affects the behaviour of the LCV and hence also the results.

In addition, the decision-making algorithm in the existing framework (Traffic Situation Management) did not use any logic for e.g. decreasing the speed in order to merge into the target lane behind a lead vehicle. Instead, the LCV kept its velocity, waiting for a gap to appear beside it instead of actively searching for, or trying to create, a suitable traffic gap. This meant that it was hard to directly see a relation between the temporal headway to the lead vehicle and the emergency braking deceleration, as was first the intention. As the LCV did not modify its speed to increase the temporal headway and find an available gap, another measure instead was used: whether a lane change was successful or not, as seen in Figure 4.1. This binary measure is not as informative as e.g. the temporal headway of the completed lane changes would be, but gives at least a hint of how different values of emergency deceleration affects the driving.

Figure 4.1 shows that in scenario 1, the lane changes suceeded to a higher degree and were initiated with a very low front headway if the emergency-brake parameter was using lower values (or if the emergency-brake scenario was not considered). Only for values greater than around 6 m/s^2 did the lane changes fail (they were first initiated and then aborted), and only for the simulations that used 7 m/s^2 was the front headway larger than around 0.08 seconds. This indicates that for low values of the emergency-brake parameter, the LCV does not deem the emergency brake scenario to be a critical situation and starts the lane change as soon as there is a positive temporal headway. Larger values of the emergency-brake parameter (that is, more critical emergency scenarios) caused the LCV to wait for a slightly larger headway before initiating the lane change, and more often caused the LCV to start the lane change before reconsidering the situation and aborting it. The results from scenario 2 were as expected: the LCV never tried any lane changes because there were surrounding traffic in the way.

As seen in the results from scenario 3 (Figures 4.3 and 4.4), the relationship is very linear between the lowest reached speed and the remaining distance when the lane change is completed. This is not surprising since, as explained in Section 3.1, the maximum allowed speed of the LCV is set to directly depend on the remaining distance. The reason that the data points tend to cluster together in Figures 4.3, 4.4 and 4.5, is because each cluster corresponds to a certain combination of cooperativeness for the different surrounding vehicles. The cluster at around 17.5 m/s, for example, appears when both the immediate neighbours of the LCV are cooperative and split up to create a gap, which causes the LCV to change lane quicker thus maintaining a higher speed. The other clusters correspond to situations where one of the two immediate neighbours are cooperative (lane change takes some more time to complete), when neither of the immediate neighbours are cooperative but the rear-most vehicle is cooperative, when none of the simulated surrounding vehicles are cooperative (causing the slowest lane changes), and so on. From Figure 4.5 one can also see that if the emergency-brake paramater is set to a low value, a larger proportion of the simulations result in a quick lane change (a data point in the highest-speed cluster). This implies that a higher value for the emergency-brake parameter causes the LCV to wait a little longer, and so reach a lower speed, before attempting a lane change. This is unsurprising because if a prediction concludes that a vehicle to our front may perform an emergency brake with a higher deceleration, more headway would naturally be needed before deeming any maneuvre as safe. So, while the LCV waits with the lane change because it is not seen as safe, the speed-decreasing mechanism described in Section 3.1 causes the LCV to lower the velocity which, if the vehicles are cooperative, causes a gap to appear in the traffic which is large enough to be deemed safe.

When it comes to potential improvements to the system, one thing would be the mechanism that decreases the speed of the LCV as it approaches the exit. The maximum speed is currently set based directly on the remaining distance, and this speed limitation is utilised until the middle point of the LCVs front axle has entered the target lane. This causes the LCV to keep decreasing the velocity even after the lane change has been started since it has not yet passed the line to the other lane, even though it is obvious for a human driver that no further speed decrease is needed. Changing this to a more advanced model that not just sets the maximum velocity based on the remaining distance would likely help in achieving a more human-like and realistic behaviour. The concept of a dependency distance could perhaps also be improved, since it currently is very dependent on tuning. To make it less affected by various velocities, it could be possible to change it into a *dependency headway*, that is, let it take the velocity into account and not only the distance. The emergency scenario currently makes the assumption that all independent vehicles are braking. This is not necessarily the worst-case scenario, since a more critical situation would be if vehicles behind the ego vehicle speed up while vehicles to the front perform an emergency brake. Although such a situation would be a worst-case scenario, it is not clear whether such a scenario is realistic and should be considered. Additional improvements could be made to the decision-making system in the Traffic Situtation Management. It is currently implemented as a finite state machine and changes state based on e.g. whether the tentative plans are deemed feasible or not. This sometimes causes problems when plans constantly hop between success or failure for every iteration, which happens just at the boundary for when a plan is seen as feasible or not. In order to avoid this hopping between states, some sort of memory or margins should be implemented in the decision making. Another enhancement that can be made concerns the simulation of the surrounding vehicles, to let them react on the lateral positioning of other vehicles. This would allow for the LCV to move closer to the boundaries between lanes in order to affect the behaviour of the simulated vehicles.

Although human drivers are very good at using various means, for example turning indicators, positioning, or even eye contact and sign language, to signal their intentions to the surroundings, this might run into unexpected problems in the future. Many modern car models make use of ACC systems that do not communicate with their surroundings, nor understand the signals that other road users employ. This could be problematic in the future if a majority of the cars on roads use semi-automatic functionality that does not react on intention signals, not only for the concepts examined in this thesis but also for human drivers that need to navigate in dense traffic. Research into systems for vehicle-to-vehicle communication might solve some of these problems but more research into human-machine interaction (HMI) is also important, and not only for the in-car graphical user interfaces but also for how the autonomous vehicles communicate with other road-users, including pedestrians.

6 Conclusion

This thesis studies the problematic situations that can occur for LCVs when they try to change lanes in dense highway traffic. The study has been performed using an existing framework for automated driving that utilises a driver model for the navigation of the LCV. A Simulink-based simulation environment was used for running the framework, and the behaviour of the highway traffic in the environment was modified from the original one to more closely resemble the highway traffic encountered in real life. The control system of the LCV makes use of short-term traffic predictions in order to find a suitable path for the ego vehicle. The traffic predictions are generated based on the concept of *independent* and *dependent* drivers, whose actions can be seen as independent or dependent of the behaviour of other traffic participants, respectively. In addition, the traffic predictions also consider two different types of situations: an expected situation where the independent drivers move with a constant acceleration, and an emergency scenario where the independent drivers brake heavily.

The developed system was tested against three different lane-changing scenarios: one with cooperative traffic, one with non-cooperative traffic, and one with mixed cooperativeness and a distance limit which the lane change must have been completed within. The results show that with a larger presumed emergency deceleration the control system for automated driving becomes more risk-averse and avoids lane changes where the temporal headway to the vehicle in front is too low. Because there is no decision-making logic to handle cases of too low temporal headway, the LCV never brakes to increase the headway and make a safe lane change possible. When letting the surrounding traffic behave selfishly and not help the LCV find a gap in any way, the results are as one would expect: feasible traffic gaps never occur since the surrounding vehicles are so close to each other, and the LCV never manages to change lane. By introducing a velocity-decreasing mechanism for the LCV, by estimating the maximum velocity that can be kept in order to complete the required number of lane changes to be located in a specific lane at a specific position, the LCV manages to always change lane. This is ensured both by cooperative surrounding traffic and the finite traffic flow in the simulation: Sooner or later, the LCV will have decreased its speed to let any obstructing vehicles pass by, and the lane change can be performed.

Future work should be aimed at improving the decision-making capabilities of the system, making sure that the LCV is able to actively search for traffic gaps instead of waiting for them to appear beside it. Another focus area should be remedying the issues that were encountered during the implementation, such as the hopping between states for the simulated vehicles. With these issues fixed, the system should behave in a more natural way, both when it comes to the controller and the simulation, which in turn should lead to more informative and reliable results.

References

[1] *Grand Cooperative Driving Challenge*. http://www.gcdc.net. Accessed: 2016-01-12.

[2] *European Commission - Road transport: Reducing CO2 emissions from vehicles*. http://ec.europa.eu/clima/policies/transport/vehicles_en. Accessed: 2016-12-04.

[3] *The SARTRE Project*. http://www.sartre-project.eu/en/Sidor/default.aspx. Accessed: 2016-12-04.

[4] J. Aurell, T. Wadman, and Volvo Trucks. "Vehicle combinations based on the modular concept". *NVF-reports, Report* 1 (2007).

[5] H. Kyster-Hansen and J. Sjögren. *Roadmap High Capacity Transports on Road in Sweden*. Tech. rep. CLOSER, 2013.

[6] J. Sandin et al. "A Lane-Change Gap Acceptance Scenario Developed for Heavy Vehicle Active Safety Assessment: A Driving Simulator Study". *Future Active Safety Technology Towards zero traffic accidents, FAST Zero 2015, Göteborg, September 2015*. 2015, p. 7.

[7] *Google Self-Driving Car Project*. https://www.google.com/selfdrivingcar/. Accessed: 2016-08-22. 2016.

[8] A. S. Brown. "Google's autonomous car applies lessons learned from driverless races". *Mechanical Engineering-CIME* 133.2 (2011), pp. 31–32.

[9] *Audi Piloted Driving*. http://www.audi.com/com/brand/en/vorsprung_durch_technik/content/2014/10/piloted-driving.html. Accessed: 2016-08-25. 2016.

[10] *Model S Software Version 7.0*. https://www.tesla.com/presskit/autopilot. Accessed: 2016-08-25. 2016.

[11] *Read the Florida Highway Patrol's full investigation into the fatal Tesla crash*. http://www.theverge.com/2017/2/1/14458662/tesla-autopilot-crash-accident-florida-fatal-highway-patrol-report. Accessed: 2017-03-11. 2017.

[12] L. Laursen. "Volvo to test self-driving cars in traffic". *IEEE SPECTRUM, http://spectrum. ieee. org/tech-talk/green-tech/advanced-cars/volvo-to-test-selfdriving-cars-in-traffic#. Up-Vr1KXla k. facebook (last visited Mar. 16, 2014)* (2013).

[13] *Volvo Car Group initiates world unique Swedish pilot project with self-driving cars on public roads*. https://www.media.volvocars.com/global/en-gb/media/pressreleases/136182/volvo-car-group-initiates-world-unique-swedish-pilot-project-with-self-driving-cars-on-public-roads. Accessed: 2016-08-26. 2013.

[14] *Volvo Cars and Uber join forces to develop autonomous driving cars*. https://www.media.volvocars.com/global/en-gb/media/pressreleases/194795/volvo-cars-and-uber-join-forces-to-develop-autonomous-driving-cars. Accessed: 2016-08-26. 2016.

[15] J. Roberts. "'Driverless' trucks on the way." *Commercial Carrier Journal* 171.8 (2014), p. 26. ISSN: 15337502.

[16] *Future Truck 2025 and Highway Pilot*. https://roadstars.mercedes-benz.com/en_GB/events/2015/october/future-truck-2025.html. Accessed: 2016-07-19. 2015.

[17] *Yutong completes world's first trial operation of unmanned bus*. http://en.yutong.com/pressmedia/yutongnews/2015/2015IBKCFbteUf.html. Accessed: 2016-08-26. 2015.

[18] C. Choudhury, V. Ramanujam, and M. Ben-Akiva. "Modeling acceleration decisions for freeway merges". *Transportation Research Record: Journal of the Transportation Research Board* 2124 (2009), pp. 45–57.

[19] E. Julian and F. Damerow. "Complex Lane Change Behavior in the Foresighted Driver Model". *Intelligent Transportation Systems (ITSC), 2015 IEEE 18th International Conference on*. IEEE. 2015, pp. 1747–1754.

[20] F. Damerow and J. Eggert. "Balancing risk against utility: Behavior planning using predictive risk maps". *Intelligent Vehicles Symposium (IV), 2015 IEEE*. IEEE. 2015, pp. 857–864.

[21] J. Olstam. "Simulation of surrounding vehicles in driving simulators" (2009).

[22] S. P. Hoogendoorn and P. Bovy. "Generic driving behavior modeling by differential game theory". *Traffic and Granular Flow'07*. Springer, 2009, pp. 321–331.

[23] M. Treiber, A. Hennecke, and D. Helbing. "Congested traffic states in empirical observations and microscopic simulations". *Physical Review E* 62.2 (2000), p. 1805.

[24] F. Damerow and J. Eggert. “Predictive risk maps”. *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*. IEEE. 2014, pp. 703–710.
[25] N. J. van Duijkeren. “Real-time receding horizon trajectory generation for long heavy vehicle combinations on highways”. PhD thesis. TU Delft, Delft University of Technology, 2014.
[26] P. Nilsson et al. “A driver model using optic information for longitudinal and lateral control of a long vehicle combination”. *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE. 2014, pp. 1456–1461.
[27] P. Nilsson et al. “Driver model based automated driving of long vehicle combinations in emulated highway traffic”. *Intelligent Transportation Systems (ITSC), 2015 IEEE 18th International Conference on*. IEEE. 2015, pp. 361–368.
[28] D. N. Lee. “A theory of visual control of braking based on information about time-to-collision”. *Perception* 5.4 (1976), pp. 437–459.
[29] D. D. Salvucci and R. Gray. “A two-point visual control model of steering”. *Perception* 33.10 (2004), pp. 1233–1248.
[30] G. Markkula, O. Benderius, and M. Wahde. “Comparing and validating models of driver steering behaviour in collision avoidance and vehicle stabilisation”. *Vehicle System Dynamics* 52.12 (2014), pp. 1658–1680.
[31] P. Nilsson. “On Traffic Situation Predictions for Automated Driving of Long Vehicle Combinations” (2015).
[32] P. Nilsson and K. Tagesson. *Single-track models of an A-double heavy vehicle combination*. Tech. rep. Chalmers University of Technology, 2014.
[33] I. Batkovic. *Optimization of driver model parameters for Long Combination Vehicles*. 2016.
[34] P. Nilsson, L. Laine, and B. Jacobson. “A Simulator Study Comparing Characteristics of Manual and Automated Driving During Lane Changes of Long Combination Vehicles”. IEEE. 2016.