Evaluating the Benefits of Speed Harmonization with Lane Specific Speed Limits on a Motorway

Kevin Abraham Jacob, B.Tech

A Dissertation

Presented to the University of Dublin, Trinity College

in partial fulfilment of the requirements for the degree of

Master of Science in Computer Science (Future Networked Systems)

Supervisor: Vinny Cahill

August 2022

Declaration

I, the undersigned, declare that this work has not previously been submitted as an exercise for a degree at this, or any other University, and that unless otherwise stated, is my own work.

Kevin Abraham Jacob

August 19, 2022

Permission to Lend and/or Copy

I, the undersigned, agree that Trinity College Library may lend or copy this thesis upon request.

Kevin Abraham Jacob

August 19, 2022

Evaluating the Benefits of Speed Harmonization with Lane Specific Speed Limits on a Motorway

Kevin Abraham Jacob, Master of Science in Computer Science
University of Dublin, Trinity College, 2022

Supervisor: Vinny Cahill

Urban motorways are plagued by issues of congestion, reckless driving and accidents. A common cause to this is the uncoordinated acceleration and deceleration of car. To address this problem, this project proposes a strategy to use a constant speed for all vehicles and distribute these vehicles across the available lanes based on the speed they are travelling at. Through this model, the aim is to achieve a lower travel time on motorways and enable automobiles to travel at the maximum speed defined for the motorway while avoiding congestion.

The project analyses the potential of a lane specific- speed harmonised driving model. The system designed in this dissertation evaluates two scenarios. The first is a setup where Human Driven Vehicles are notified by a central control unit about the permissible speed ranges for each lane. The second uses Connected Autonomous Vehicles that communicate with a central controller and vehicles are assigned speeds, based on which they move to the designated lane. The baseline model used, is developed by Guériau and Dusparic (2020) and results from the simulations are compared to the baseline to evaluate its potential in a real world implementation.

In a three lane motorway, the Slow lane has a max speed of 50km/h, the Mid lane has a speed of 80km/h and the Fast lane has a speed of 100km/h. As the lane attains maximum capacity, the next lane's speed cap is increased and new vehicles join the lane. Through this process, all the vehicles on a highway travel at a speed close to the maximum speed limit. SUMO (Simulation of Urban MObility) is a tool used to run the simulations.

The experiments carried out yield the following results- average travel time is reduced by 21%. The average speed of vehicles is 80% of the maximum speed limit. Finally the effects of the model create a traffic flow pattern that promotes coordination between vehicles by lowering the number of lane changes by eliminating overtaking.

Acknowledgments

I would like to thank Professor Vinny Cahill for his feedback and inputs that helped me scope this project and helped me narrow down the outcomes I wanted to achieve through this dissertation. His advice has been helpful and his time greatly appreciated.

I thank the staff of the Trinity IT group, who got the required simulation software installed on the lab machines so that I could run the simulations.

Additionally, I would like to thank my friends Joanne and Aakash who helped me get in touch with a few people in the connected vehicles industry who took time to answer my questions as the domain was new to me when I began this project.

I am grateful to my family for their constant encouragement, support and instruction which helped me stay on track and bring me to this point after a year's work.

Finally, I am grateful to God for giving me this opportunity to study at Trinity and work on a project related to Connected Vehicle Infrastructure which is an avenue I enjoy exploring.

KEVIN ABRAHAM JACOB

University of Dublin, Trinity College
August 2022

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This managed motorway is a project that focuses on developing a strategy to manage traffic flow on highways. The flow is traffic is often hindered by congestion which is caused by drivers impulsively slowing down due to a trigger on the road such as diversions as explained in the the article (JAMS, 2022), which terms it as the butterfly effect. Another cause is vehicles switching lanes or in response to vehicles slowing down. The discussions and reflections from this dissertation are centred around how travel time can be reduced and speed of vehicles can be maximised while ensuring the vehicles can travel efficiently on a stretch of road with minimal accidents.

1.1 Motivation

The simple solution to traffic is to get humans to change their behaviour (CGPGrey, 2016). This is what CGPGrey says in his video on overcoming traffic congestion. However, this would only be achievable in an ideal world.

Human behaviour is often a result of a stimulus which triggers a response. On a motorway, if a driver notices a vehicle ahead slowing down, the drivers instinct is to slow down. The trigger here was the brake lights of the vehicle ahead and the behaviour was applying brakes. But as a human, one does not know what caused the driver ahead to slow down. In introducing a central control system and leveraging the use of vehicle to vehicle communication, messages and instructions can be circulated better thus having a single 'trigger' and a collective response.

This project attempts to have different speed limits defined for each lane. The technique of modifying the speed limits of a motorway based on conditions like weather, vehicle density, etc is called variable speed limits. The proposed model would incorporate speed harmonization, which means all the vehicles moving along a motorway- in this project, a lane will attempt to have the same speed. Differences in speed are often the cause of congestion- when vehicles ahead come to a sudden halt or sudden bursts in acceleration from a vehicle result in poor synchronization in the consequent vehicles. Studies show that traffic moves best at a steady pace as opposed to vehicles speeding up and slowing down (Liu et al., 2009). While this model can be implemented with CAVs, it proves to be more challenging with human driven vehicles. This is because humans can be impulsive in their driving behavior and reactions to braking, lane changes and signals have varying response time. In a system of completely connected vehicles, the braking distance, car following distance can be reduced and the speed limits can be increased. The coordination of autonomous vehicles would assist with uniform acceleration and then maintain a constant speed.

The concept of Variable Speed Limits (VSL) while not new, sees a new light when paired with speed harmonization. Some of the key factors for consideration within this project include-

- Distance to be maintained between vehicles
- Vehicle life spans and onboard equipment for CAVs and Driver experience for HDVs
- Attainable speeds
- Permissible speed limits on each lane
- Number of vehicles present on a motorway

1.2 Problem Description

Imagine a scenario while driving down a motorway and the vehicle ahead changes lanes abruptly. The vehicle behind sees this and reacts by slowing down. This leads to a domino effect of all following vehicles slowing down to some extend, and one vehicle somewhere along this chain comes to a complete halt. This start-stop motion of vehicles causes a phenomenon called Phantom Intersections, since it resembles vehicles stopping at a crossroad.

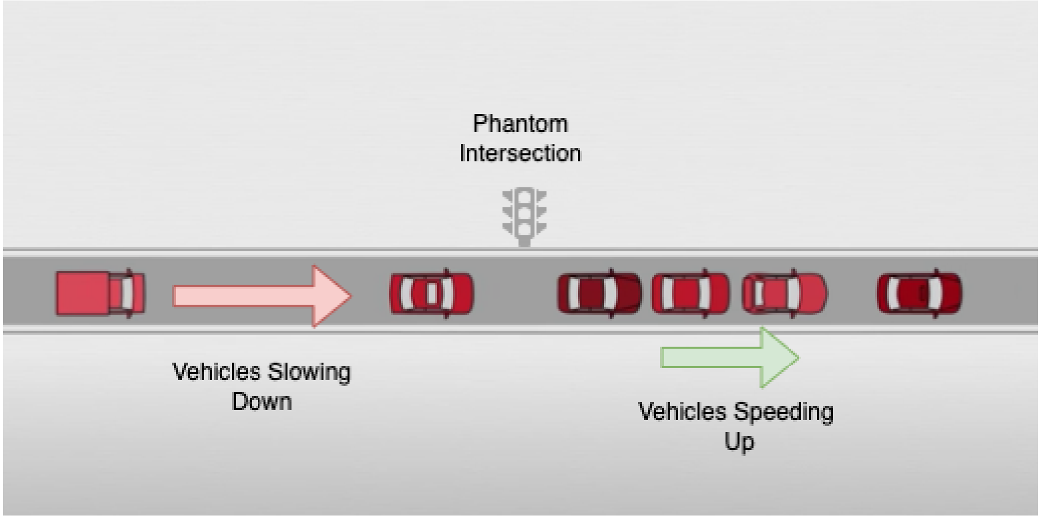

Figure 1.1: A Traffic Snake resulting from a phantom intersection.

Phantom intersections is one of the common issues on motorways. This results in congestion on the motorway and often leads to erratic lane changes by vehicles to avoid this jam, or accidents in more dire scenarios. It also causes the overall speed of vehicles to drop, which indirectly lowers travel time. (Ishikawa and Arai, 2016).

1.3 Project Overview

MODEL DESIGN

- **Human driver based model:** The short-term design would incorporate human drivers receiving messages via a central control for speed limits based on the lane and minimum distance to be maintained between each vehicle for the specific lane and instruments such a cruise control could be used. The decision to follow these instructions rests on each driver and as more people follow the limits, the highway could attain a structure that is the most efficient for all the vehicles.

- **Autonomous Model:** For a system designed using autonomous vehicles, the target is a long-term approach. The design envisioned for this project involves vehicles receiving instructions from a central control to adjust speed and acceleration based on the highway parameters like vehicle count and lane. The vehicles in a lane would maintain a fixed distance from the vehicle ahead of it and uniform acceleration and deceleration would be in effect when vehicles switch lanes.

  Another factor to be considered in this approach is, the middle and right lanes would have a higher top speed and a lower range of permitted speed, the leftmost lane would have the lowest top speed and a larger permitted range of speed to allow vehicles to exit and enter the motorway as well as cater to slow moving traffic. The values of the speed would have to be adjusted based on the use case.

Use Cases A few use cases we to address through this project are-

1.3.1 Co-operative Driving

Cooperative driving is a traffic behaviour where data is aggregated from multiple sources on the road network and the information processed is shared with the participants of the road network to attain a coordination of a specific type. This can be for vehicles to share the same cruising speed, or maintain a specific cruising distance. It constitutes a part of the Intelligent Transport Systems and improves the quality of information shared with drivers about the present environment, other vehicles on the road and any rules or alerts that need to be communicated.

Each vehicle measures the parameters to be shared and uploads the data to the roadway network. The required metrics are processed and sent to other vehicles to establish a coordination withing the system (Jose M. Armingol and Vicente Milanés, 2018).

Using the principles of coordinated driving, the aim is to achieve a synchronous flow from a point A on the highway to point B, coordinated lane changes and dynamic speed and acceleration variance.

1.3.2 Variable Speed Limits

The prospect of VSL forms the basis of our theory to coordinate driving. This serves as the key to managing the traffic in each lane, lowering the travel time, and reducing/ mitigating phantom intersections and issues that stem from unprecedented braking. In this project, the speed of each lane will vary based on conditions such as number of vehicles on the highway, presence of roadblocks, autonomous and HDV scenarios.

1.3.3 Speed harmonization

The introduction and promise of autonomous vehicles opens opportunities for a controlled traffic flow. While traditional vehicle acceleration was based on the decisions and intuitions of human drivers, autonomous vehicles could potentially be managed centrally to allow vehicles to speed up and slow down based on the situation on a road. Vehicles travelling on a single lane can have a constant speed defined or a specified range which will allow other vehicles to maintain this with fixed distances between consecutive vehicles and braking distances accounted for.

1.4 Research Question

Can a dynamically adjusted speed harmonisation strategy for each lane on a highway along with variable speed limits be more effective than present day driving schemes to reduce travel time and improve the flow of traffic to minimise traffic forming clusters on a highway?

1.4.1 Research Aims

The aim of the analysis is to evaluate how the flow of traffic in a controlled environment of speed and can result in an efficient and harmonized driving model. The results should ideally reduce the travel time when measured against exiting data and lower the rate of congestion. The specific objectives of this study are:

- Can a speed harmonisation model lead to speed limits on motorways being increased?
- How would travel times change by using speed harmonisation when compared to a system without it?
- Is a higher level of control better on a motorway as opposed to leaving the decisions to each driver?
- How long can vehicles on a motorway travel at the maximum speed limit?

1.4.2 Project Scope

This project is built to create a means to make all the vehicles on a lane travel at the same speed (Speed Harmonisation) and vary speed limits across each lane, to determine the benefits of such a strategy against one that follows traditional or driver based traffic flow. The project will be a simulated model using SUMO and the results cannot be interpreted as is for real world roadway system. It can only be used as a point of reference or ‘in theory’. This is because the testing of the system would be run using a simulation and not in a real world environment which would have other deterrents such as road conditions, weather which cannot be specified or be accurately defined within the simulation.

The actual means of vehicle to vehicle communication for the CAVs and HDVs is also beyond the scope of this project owing to which details about the techniques will not be addressed or explained in detail. The simulation does not address vehicular accidents as extensive research already exists in this sphere (Cheng et al., 2018).

In the aspects of security, a high level overview is discussed in brief with regard to threats, network overrides and physical issues that could occur due to faulty infrastructure. However, the security of the communication channels is not a domain of concern within this project.

1.4.3 Baseline and Roadway Model

The road network used for this simulation is the M50 motorway in Dublin. The primary model was developed as part of a network simulation for movement of traffic on the M50 motorway designed by Guériau and Dusparic (2020). This model serves as a baseline to which the results from the speed harmonisation model are compared.

Chapter 2

State of the Art

The chapter provides an overview of existing work in the area of smart highways and speed harmonisation. It draws the specifics to which this project draws parallels, analysis the current state of the art and describes in brief the existing pertinent work. The papers examined cover domains of smart highway architectures, lane changing models for connected vehicles, measures for coordinated driving and Variable speed limits in autonomous and human driven vehicles. The tool used for simulation (SUMO) is explained with reasoning for the selection of this tool.

2.1 State of the art

The state of the Art for Dynamic speed harmonisation is highly speculative. Variable speed limits and the their benefits have been exploited by systems across the globe. However, dynamically varying them based on control parameters with coordinated driving, HDVs and CAVs is still in its infancy with regard to research. This project focuses on VSL for each lane and harmonising the speed specific to a lane. This SOA research explored the existing papers on VSL, coordinated driving and its benefits along with the use cases highlighted specific to the papers.

The research papers have described the Impacts of Distributed Speed Harmonization (Goulet and Ayalew, 2020), Traffic Monitoring schemes (Villanueva et al., 2013), integration of smart vehicles and highways (McQueen, 1994), algorithms for coordination (Vrancken, 2005), Variable Speed Limits (Lahmiss and Khatory, 2020) and , motorway merging using VSL (Carlson et al., 2011).

2.1.1 Connected Autonomous Vehicles

A key feature of connected vehicles, as seen in most systems, is the interconnection within the network and to the infrastructure surrounding the network. The drawbacks of Autonomous Vehicles (AV's) is that they operate within a black box and the decision models do not have a clear explanation. Connected Vehicles (CV's) on the other hand depend extensively on the messages exchanged to make decisions and this causes failure within highly concentrated networks or lack of sufficient communication devices. Thus CAVs were introduced in order to address the issues mentioned above.

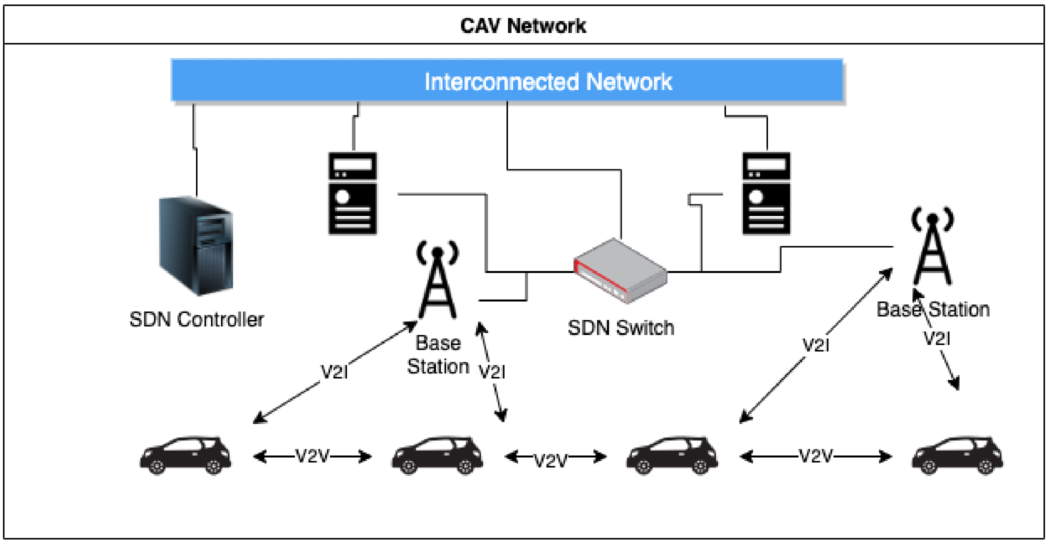

Figure 2.1: Cooperative CAV system

Role of V2X in connected vehicles: Connected vehicles make us of V2X communication to talk to other actors on a road. V2V is used for vehicle to vehicle, V2I is vehicle to infrastructure and these systems can send context aware messages (CAM) between vehicles. The information exchange can warn drivers of impending crashes. However V2X communication is still considerably unreliable and often issues due to latency or signal loss occurs when the concentration of devices in a small area is large.

Accident detection and mitigation (ADM) is a strategy that relies heavily on sensors and computationally demanding resources. Platooning is a strategy that's effective to ensure vehicles move at an effective speed while maintaining the desired spacing between consecutive vehicles within the road network. Initially radar based systems were used to determine the space to maintained from the vehicle in front. With V2X, this information can be carried out a lot more efficiently.

Connected autonomous vehicles are one of the actors used to test a scenario in this dissertation.

2.1.2 Speed Harmonization

Speed Harmonization is an effective mechanism to manage traffic congestion. The technique is gaining more momentum with the introduction on connected vehicles and autonomous driving. The concept of harmonising the speed of vehicles on a motorway allows more control if an unprecedented scenario occurs as it allows less scope for erratic behaviour. The connected vehicle infrastructure allows the participants to share real time data about speed, vehicles count on a motorway and for more granular networks, positioning data can be inferred. This information would be vital to decide the average speed for the motorway. The idea behind the technique is to improve the highway performance by mitigating vehicle break-downs and minimise bottlenecks on the road. It aims to uphold the Traffic Flow theory which states that the a break down of traffic at bottlenecks can be avoided by ensuring the speed of the upstream traffic matches that of the traffic at a bottleneck. This eases the flow of vehicles(Rios-Torres and Malikopoulos, 2017).

The paper goes on to describe a framework that allows vehicles to control their speed before entering a zone that is compromised due to a bottleneck. Speed harmonisation can be further categorised into

1. ***Reactive Approach:*** This implementation relies on human intervention to decide the speed limits to be imposed on a road based on conditions such as weather, traffic congestion, ongoing construction and road works. The advent of sensors has made this data faster to gather in real time. This has proved to be effective in improving the safety and and vehicular mobility in work-zones. However, the drawback of the reactive approach is that it relies on a negative scenario to occur in order to make a change.

2. ***Proactive Approach:*** This is a preventive approach designed to eliminate the shock waves that could occur with an accident on a motorway. It uses cost functions of Average travel time, traffic density in different sections and creates a predictive model with this data. It uses the time series data to make predictions. However, the performance of the model can be affected by the driving topology and driver behaviour. With CAVs, the second factor becomes more reliable.

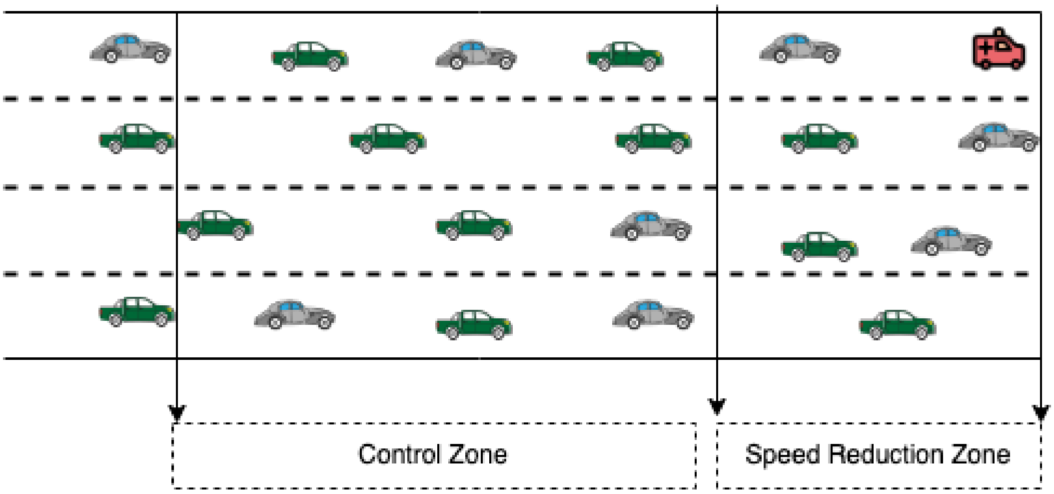

Figure 2.2: Vehicles within a control zone approaching a speed reduction zone.

The framework model has the following design to optimally manage speed, that consists of a speed reduction zone and a control zone. The control zone is the region where the vehicles can accelerate and decelerate in order to attain the appropriate speed before entering the speed reduction zone. The control zone ensures that vehicles are adequately prepared for the rules within the speed reduction zone and creates a virtual queue to manage traffic.

The idea and approaches discussed in this paper are borrowed and implemented in part within this dissertation. The speed harmonisation model builds on the control and speed reduction zone concept.

2.1.3 Variable Speed Limits

Variable Speed Limits are part of the intelligent transportation system (ITS) and functions as a tool to relieve congestion on roadways, improve the traffic flow and in-turn improve the road safety. The system functions by calibrating the speed limits on a stretch of road based on the real time traffic conditions, environmental factors (such as sleet, rainfall, snow) and the condition of the road infrastructure. This is often carried out by transmitting messages to a central control system that makes calculations based on the incoming data and displays the speed on highway speed boards, or sends the limits to the vehicles in Connected vehicular networks.

The main benefits of VSLs are described by the paper-

1. ***Safety Improvement:*** When road conditions are poor, it can adversely affect the rate of traffic flow. This also increases the risk of accidents. Couple poor infrastructure with environmental conditions like snowfall or rain and the impairment of

a drivers vision could result in dangerous accidents that could prove to be fatal. Experienced drivers would be able to react to these situations by driving at optimal speeds but this reliance on human behaviour can be unpredictable and is not reliable. Therefore Variable speed limits are used to fix the upper limit of a roadway and the data is updated periodically based on the conditions of the road network, traffic density.

An evaluation was carried out to gather results you the crash risk with and without VSL(Cheng et al., 2018). The results from the study showed that the risk of crashes are reduced considerably with VSL. The downside to this experiment showed that the trade-off for safety was increased travel time. It also proved a strong correspondence between vehicular speeds and crash potential. The highest impact of this is observed in high density situations like merging lanes and high traffic flow.

2. ***Environmental Impact:*** The use of variable speed limit has improved the rate of fuel consumption which reduces the pollution and adverse environmental impact due to harmful emissions. The study shows that higher vehicle speed and rapid acceleration and deceleration can cause a drop in fuel efficiency(Liu et al., 2012). The paper goes on to state that a speed limit of 70 mihour would reduce the carbon emission by a factor of 1 million.

3. ***Traffic Efficiency:*** VSL helps mitigate congestion with imposition of speed limitation on roadways. The incentives from congestion management can be categorised into benefits of VSL with speed modulation and breakdown sections of the road network that could cause a bottleneck. Reinforcement learning techniques are used to reduce the travel time by 21% and coordinate the movement of vehicles on a motorway.

2.1.4 Traffic Monitoring Schemes

The paper introduces smart traffic management as a pillar of a smart city. It goes on to propose an architectural framework to bring down the maintenance cost of highways by introducing a real time monitoring system. The monitoring encompasses the security fence, pollution levels and accident detection on highways, making use of Wireless Sensor Network (WSN) integrated with highway infrastructure. The highway is divided into different stretches which constitutes a portion of the road between two exits. The stretch can be a hundred meters to twenty kilometres in length. Each stretch of the highway is monitored. The parameters monitored are-

1. Temperature and humidity

2. Traffic
3. Pollution
4. Noise
5. Traffic Speed

The Network Topology design presents a challenge with the number of WSN nodes to be deployed within a highway. Estimating this is done by trial and error. Each WSN node has a temperature sensor along with humidity and pollution level indicators. The nodes periodically send data to a gateway with stores the information received and distributes it to the rest of the infrastructure within the network. Energy management for the deployed infrastructure is handled by connecting the nodes to the power lines within a city and solar panels or battery power backup us used for rural areas.

The architecture consists of 3 layers -

1. ***Event layer:*** It functions similar to the lower layers of the OSI model. The sensors connected in this system use the Secure MQTT protocol for communication. The MQTT protocol uses the publisher/subscriber model for messages and serves as the underlying framework for network communication between nodes. This information is relayed to the gateway infrastructure.

   The event layer also has multimedia sensors like cameras and microphones to identify anomalies in the road by capturing and forwarding media streams.

2. ***Management layer:*** This layer is responsible for data processing based on the information transmitted from the sensors. It provides a user a services for deployment and security. It is used in gateways and forms a mesh network

3. ***Presentation layer:*** The layer presents the information to the users such as Fire stations, Police etc based on the need, permission level and the immediate need of the system (accidents, Fire)

The project goes on to discuss how the future of highways would include collision detection, noise and traffic distribution (commercial to private vehicles, vehicle types) (Villanueva et al., 2013).

2.1.5 SUMO- Simulation of Urban Mobility

SUMO which stands for Simulation of Urban Mobility is a road network and traffic simulator available under the open source EPL 2.0 license. The suite provides tools to import road networks, create routes based on various sources and a simulator that can be run using the command line or graphical user interface.

SUMO can be used to evaluate the performance of traffic lights, evaluation of modern algorithms. It can be used to investigate choice of vehicle routes, SUMO is used extensively by the V2X community to provide real-time vehicle traces and to evaluate applications using a network simulation.

SUMO offers two solutions to interact with network objects in real-time. They are TraCI and LibSumo.

1. **TraCI** TraCI which stands for "Traffic Control Interface" is built on the TCP client-server architecture and provides access to the SUMO tools. SUMO functions as the server in this instance. TraCI uses SUMO in combination with communication networks to recreate vehicle to vehicle communication. TraCI creates client processes based on the scenario being simulated and supports a large number of languages such as python C++, .Net and Java. For this project simulation, python with the language of choice. Using TraCI, the transportation network can be modified dynamically. Aspects of the system such as vehicle states, routes taken by vehicles, traffic lights and types of objects on a road network can be changed at runtime. This provides dynamic control over the network since multi agent scenarios can be testing without spending time on modifying individual components using XML templates.

2. **LibSUMO:** LibSumo uses TraCI to interact with a running simulation which offers the flexibility of cross platform, and network interactions which SUMO acting as a server. LibSumo overcomes the drawbacks of TraCI by removing the need for socket based communications and offers a library that can be integrated with the programming environment. It has prebuilt binding for Python and supports other programming languages like JAVA. The interface relies on static functions and simple wrappers that can be bootstrapped to the client code base.

   Limitations of LibSUMO While libSumo has its benefits over TraCI, there are a few drawbacks that deter it from being the obvious choice.

   - Running LIBSUMO with sumo-gui does not work directly in Windows systems
   - Certain subscriptions like vehicle.getLeader requires additional arguments to be passed.
   - While using TraCI, every exception is generated as an error with stdout (Standard Output), however lib sumo does not generate this error message.

Chapter 3

Design

This chapter describes the design decisions used to create the model for speed harmonisation. It breaks down the components within the system and discusses the pertinent aspects of each.

In order to help the reader understand how the project was built, explanations are provided for the design choices. As stated in the previous chapters, the purpose of this project is to create a speed harmonisation model for motorways to lower the rate of congestion and streamline the flow of traffic minimising bottlenecks that are created due to lane changes. It also works toward attaining speeds close to the maximum speed defined for the highway. From the simulation runs and tests carried out, an evaluation is made to determine if the new model serves as a design can can be incorporated into the existing motorway architecture.

The entire project was developed using SUMO and interfacing this with TraCI using the python library. The communication between agents on the road network are constructed using the SUMO toolkit and simulated via commands and events within the code-base. The model is evaluated for the M50 motorway in Ireland which spans a distance of 49 km (Dowling, 2020).

The concepts used in this project are not entirely new, for instance the M50 motorway has been using Variable Speed Limits since October 2021 (Beehanl, 2021). The speed limits are monitored and controlled by the TII (Transport Infrastructure Ireland) motorway operations in Ireland. The original speed limit on the motorway used to be 120 km/h but this was altered to a maximum of 100 km/h to better control the traffic flow and mitigate accidents. Speed harmonisation is a technique that has also been around for some time. Adaptive Cruise Control is a subset of this, which uses lasers or radars to measure the distance to the vehicle ahead, and uses this data to maintain a safe distance. Incorporating these motorway management techniques into a single system provides an opportunity for dynamic control over the agents on the road network while also maintaining autonomy.

3.1 Requirements

- The system should function for all the types of vehicles that can travel on a motorway which includes most road vehicles.
- Roads on the motorway for which the rules apply should have a minimum of 3 lanes
- Vehicles should be able to merge into the lane without collisions
- Before changing lanes, vehicles must check if a lane change is possible.
- Unnecessary lane switches should be avoided. The lane change is restricted if a vehicle is close to the destination (5 km in this scenario)
- Traffic congestion should be avoided and the overall speed of the vehicles should tend toward the permissible maximum.

In order to achieve these requirements, a lot of challenges had to be overcome. Beginning with the model provided for the M50 (Guériau and Dusparic, 2020) each vehicle on the network had an individual definition within an XML file which contained over three hundred thousand vehicle definitions. These vehicles had to be modified based on the metrics specific to the harmonisation model and consolidated into a single type to negate redundancy.

3.2 Restrictions and Constraints

In order to complete this project within the stipulated time frame and few design decisions were made to create the model that does not account for every real world edge case. These constraints fall in with regard to the road network, vehicles, lane changing patterns and car following models.

- All the vehicles follow a predefined path with the source and destination configured.
- Vehicles do not change their destination midway through the journey or execution.
- Environmental conditions such as rain and snow are not factored into the model to modulate speed.

- Conditions such as roadworks, construction, diversions are eliminated in this scenario to test how the behaviour would be in ideal conditions.
- Traffic lights are avoided in this motorway scenario.
- Emergency Vehicles can exceed the speed threshold.
- Human Driven Vehicles have a permissible range of speeds.
- Connected Vehicles operate on a different scenario and have their speed at a fixed value with lower scope for variance.
- The car following model used allows for more freedom with vehicle separation and stopping distance
- Overriding the driving model is a complex task in SUMO therefore the fastest lane is configured as the left lane.
- The left lane is selected as the fast lane since the typical driving models have overtaking vehicles return to the original lane after the overtake.

3.3 System Implementation Overview

The overarching design for the system implementation is shown in Fig 3.1. The Speed harmonisation controller is written in python using Traci and interacts with the objects within the SUMO simulation. The command line interface is used to initialise the network. The *network configuration* file contains details about the roadway, number of intersections, the edges and junctions as well as the definition of the number of lanes on a stretch of road. This file is the first one to be primed and creates the foundational road network. Following this, the *vtypes* file is initialised which contains the vehicle type definition which is passenger vehicle or truck and assigns the car following model to be used. Based on the details within the *routes* file, each vehicular route is created with an identifier and the individual lanes that need to be traversed are added to the route in the order of traversal. This defines the path a vehicle would take during its journey. Finally the *emitters* file defines each vehicle with their point of entry, entry speed, time and the route to follow as defined by the identifier in the routes file.

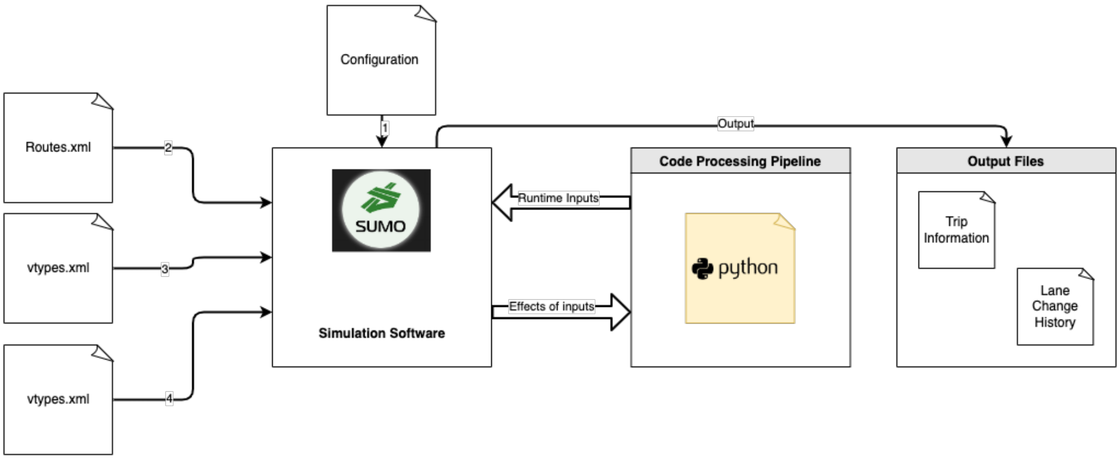

Figure 3.1: File System used to create the simulation

This data is loaded into SUMO on startup and as the execution runs the python script intercepts the vehicles to perform based on the definitions of the speed harmonisation model. Once the execution terminates, output files are generated. The files include-

1. A trip information file which details the number of vehicles, average journey time, average vehicular speed
2. A lane change file that displays every lane change that occurred on the motorway

3.4 System Architecture

The process for designing the Speed harmonisation system began by analysing systems like Cavnue (Nason, 2021) which implements a dedicated lane for autonomous vehicles. Various managed motorway strategies use VSL to determine motorway speeds, cruise control to maintain speeds but all these methods rely on human compliance. A combination of these techniques is what led to the development of the model that utilises VSL along with vehicular communication to achieve a constant speed while on a motorway.

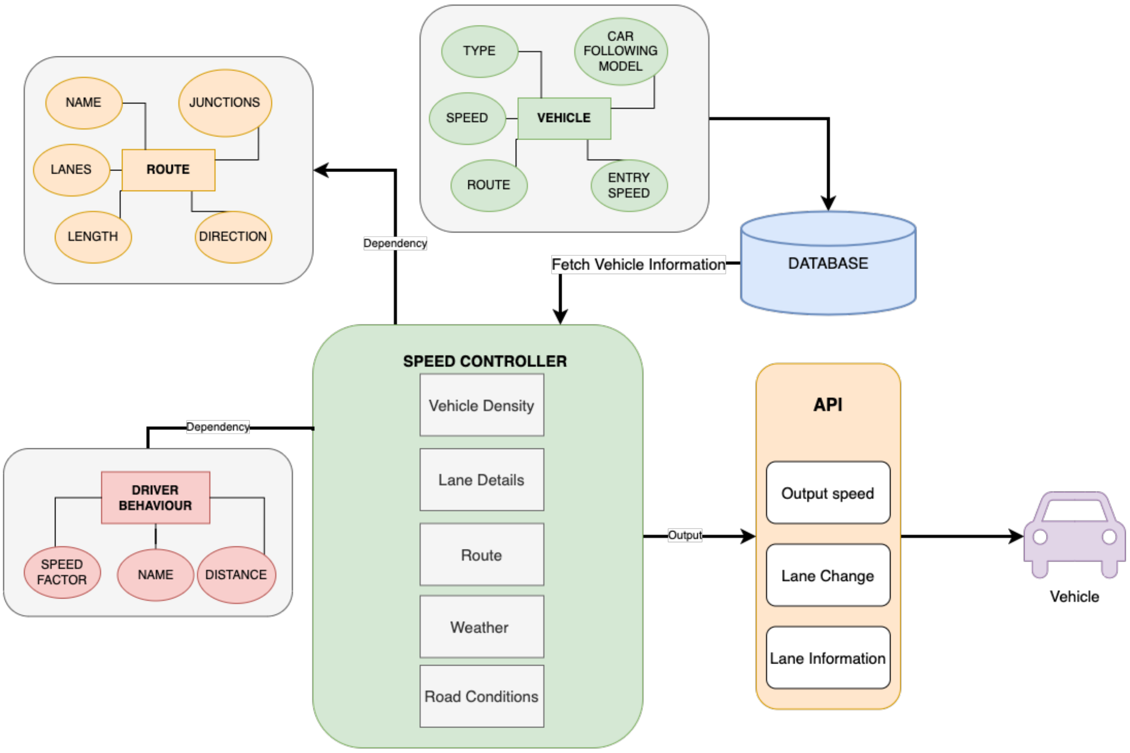

Figure 3.2: System Architecture for a real world implementation

The proposed real world model (Fig 3.2) accounts for vehicle registration based on the number-plate, make and model of the vehicle as well as maximum speed that can be achieved. In addition factors such as rain, sleet and snow are considered while deciding the speed limit. The final speed and lane switching information is projected to the user's dashboard with the connected vehicle architecture. The real world model would make use of an API to transmit the data to vehicles and a database to store the vehicular information.

The data output from the vehicles such as lane change information, average speed and destination or exit taken would be useful for processing and is discussed further in the evaluation section.

The simulated model architecture (Fig 3.3) does not include the environment considerations and vehicular speeds are displayed on the terminal as opposed to a dashboard.

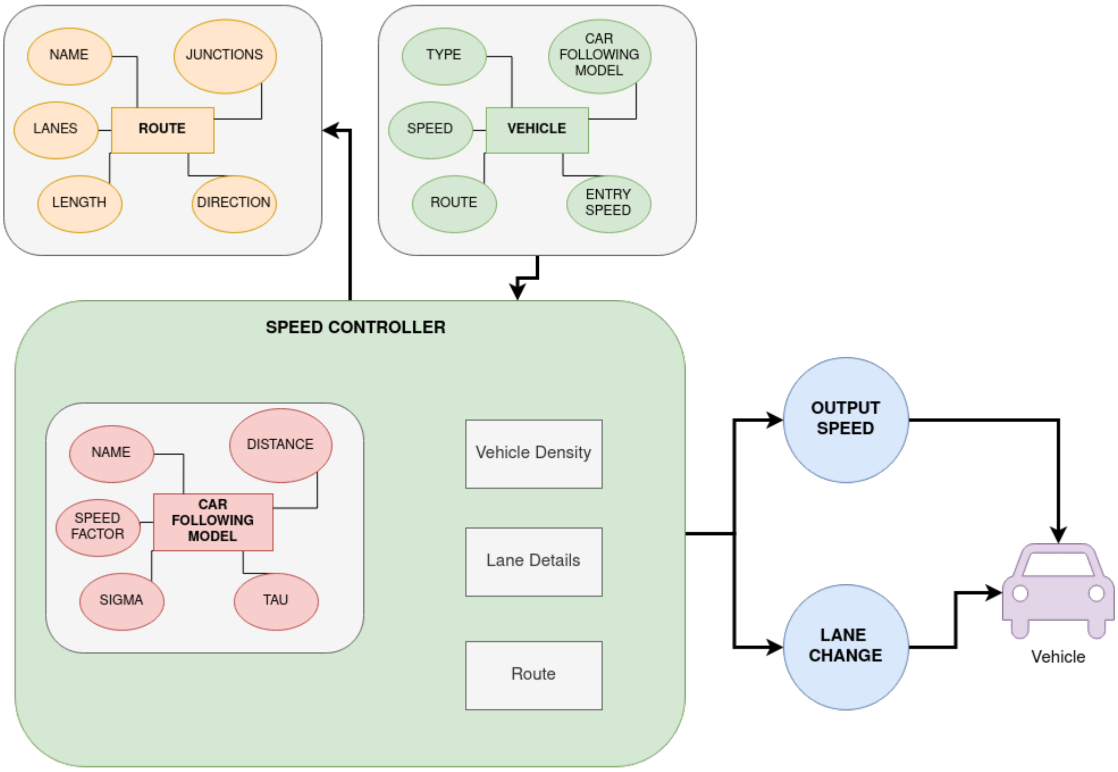

Figure 3.3: System Architecture-Simulation

The architecture of the system is as follows. The vehicle module defines the properties that can apply to a vehicle which include

1. Automobile type- Passenger vehicle
2. Maximum speed
3. Car following model
4. Entry Speed
5. Route to follow

The data captured from the car module is matched with the route module which contains details about-

1. Route Name
2. Length of the route
3. Number of lanes
4. Connecting junctions

5. Direction

The information from the routes module serves as a component upon which the vehicle plans its journey. The vehicle module communicates with the central speed controller which takes into account the various factors such as volume of vehicles within the network, current lane speeds and maximum speed factor.

A car following model makes up a part of the constraints applied to a vehicle traversing the highway. The properties of the car following model include -

1. Name
2. Minimum Distance between vehicles
3. Speed factor which defines the difference between the maximum permissible speed versus attainable speed for a vehicle in terms of a percentage
4. Distance a driver decides to maintain from the vehicle ahead
5. The divers rate of expertise

With the information provided, the speed harmonisation controller generates the ideal speed for each lane on the motorway and this data is communicated to the vehicles on the road for compliance. Based on the kind of vehicle (HDV or CVs) the compliance varies.

3.4.1 Human Driven Vehicles

The vehicles driven by humans follow a scenario where, the vehicle connects to the highway network and is added to the system. Each one is assigned an identifier and this is used to track the vehicle during its journey. Vehicles can have different points of entry as well as lanes unless they join in via a ramp. The vehicular system engages with a central control block that determines the lane of each vehicle based on their travel speed.

The algorithm progresses by capturing the speed of each vehicle. If the speed is greater than 80km/h, the vehicle is instructed to move to the rightmost lane as long as a lane switch is possible. A lane switch could fail to occur if the vehicle is already on the lane it is supposed to travel on, or if another vehicle is in its path which prevents lane switching.

If the vehicle is not approaching maximum speed, it is notified to move to the middle lane and constraints of lane switching are applied to ensure accidents are avoided and rules of the motorway are observed. A vehicle can stay on the middle lane while it maintains a speed between 50 to 75 km/h.

If the vehicle is slow moving, it traverses on the leftmost lane and can move along this lane while the speed is within the range of 20 to 50 km/h.

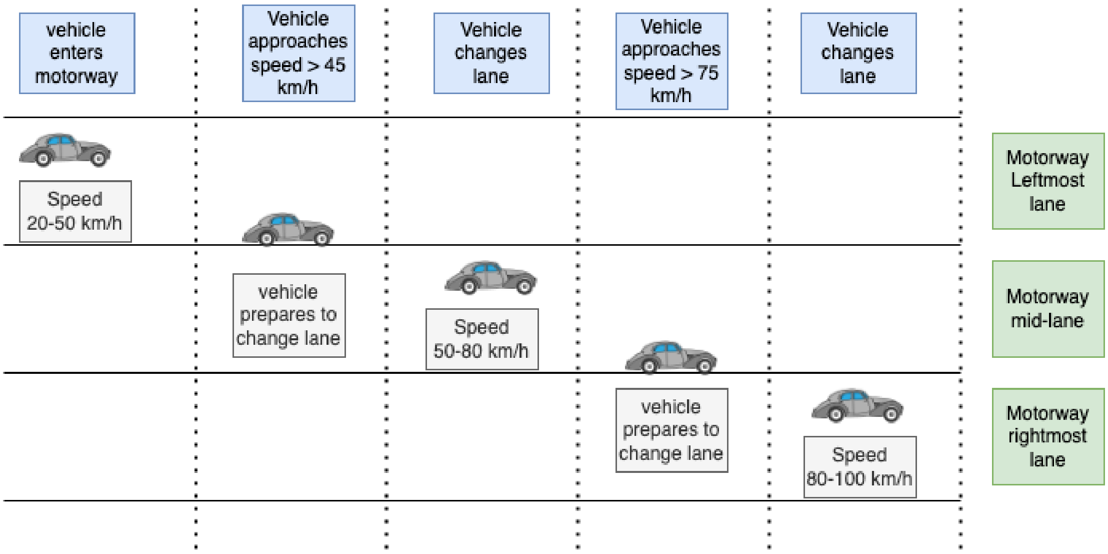

Figure 3.4: Representation of Human Driven Vehicles changing lanes

With this round of the simulation complete, the next phase focuses on objects on the motorway that have not switched lanes due to obstructions or have newly entered the motorway. The system fetches the current position of a travelling vehicle and checks its lane. If the car falls in the appropriate speed bracket, the vehicle can continue on the lane it is moving. If this check fails then the vehicle is instructed to move to lane defined by the Variable speed Limit set by the control unit, which is the speed controller. A slow moving vehicle on a middle or fast lane, moves to the left. A vehicle travelling between 50-80 km/h on the fast lane will be notified to move to the middle lane.

Table 3.1: Lane Speed Definitions

	Speed
Fast Lane	¿80km/h
Mid Lane	50-80km/h
Slow Lane	¡50km/h

This ensures distribution of traffic load and collisions due to speed are avoided. The flow of logic for a vehicle to stay or switch lanes is explained in Fig 3.5. The model continues to run in a cycle as log as vehicles are present on the motorway.

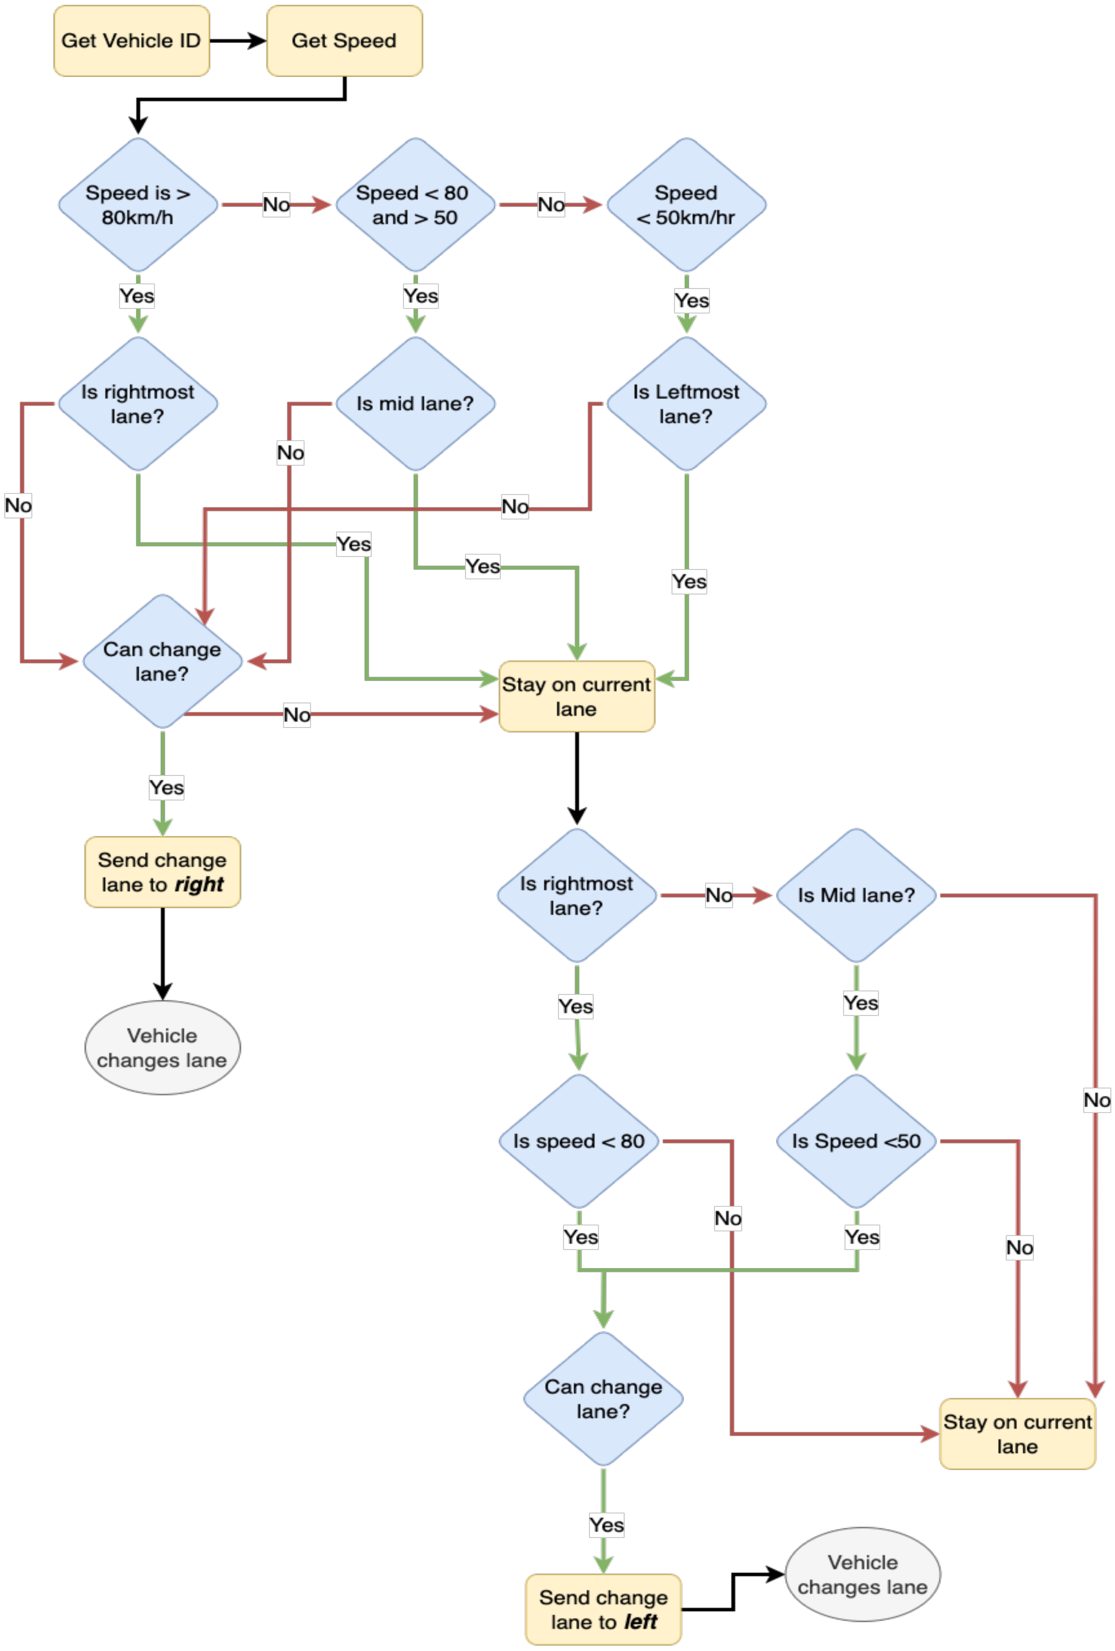

Figure 3.5: Control flow for Human Driven Vehicles

3.4.2 Connected Autonomous Vehicles

The scenario created for the autonomous vehicles offers a greater level of control to the Speed Controller or Central control unit, which would be a motorway authority in the real world.

Using the details gathered on the number of vehicles and the operational efficiency of the motorway, the speeds for each lane are defined by the central control. The range of speeds and duration of its validity vary dynamically as highway conditions change. For instance, the speed controller would assign higher speed limits at non-peak hours since the volume of traffic on the motorway is less than half or two-thirds of the maximum capacity. This would imply that cars have the luxury to move much faster with a greater stopping distance and following distance being maintained thus not compromising on security.

The speeds decided upon by the system relayed to the vehicles within the network

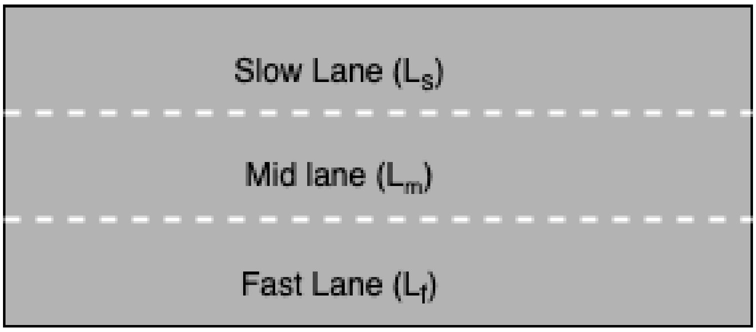

Figure 3.6: Lane Queues

The initialization of this model is similar to the HDVs. The system captures the identity of vehicles on the highway and runs a process to get the speed of each vehicle. Once a vehicle speed has been determined, the algorithm to harmonise speeds kicks in. In this scenario we consider a three lane motorway with vehicles entering the road network from one end. The controller checks the speed of the vehicle and then performs an operation by assigning the vehicle to a specific queue. The goal of the system is to have the most vehicles moving along the fastest lane without compromising the speed. In order to achieve this a vehicle that is recognised is assigned a speed and the vehicle accelerates to that specific speed. The object is also added to a queue that keeps track of the number of entities on the fastest lane. The vehicle then moves to the fast lane (L_f) and stays on that lane until it approaches the end of its journey. As more vehicles join the network, the process of assigning vehicles to the lane progresses. In order to balance the count of vehicles, the process assigns a random vehicle to the mid-lane(L_m) with a speed of 70 km/h while the count of vehicles in the fast lane is 10% greater than the other lanes.

As the fast lane($\mathrm{L_m}$) approaches maximum capacity, the speed limit of the mid lane is increased to match that of the fast lane. After this is complete, the vehicles joining the network are pushed to the mid lane and a separate queue is maintained for vehicles travelling along this lane. With the connected autonomous vehicles infrastructure, the system does not have to worry about drivers making decisions based on instincts or experience and a higher level of control can be established. With this design, vehicles tend towards a streamlined flow similar to that of a moving train with vehicles at the end of their journey looking like a single file of cars. The operation of identifying and assigning vehicles continues to run as a loop and at every interval the vehicles on the highway are monitored. To optimise the system within SUMO, an additional condition is assessed- if a vehicle is already assigned to a queue, then the control shifts to the next object so that the time of processing can be optimised. The entire control flow can be visualised using Fig 3.7

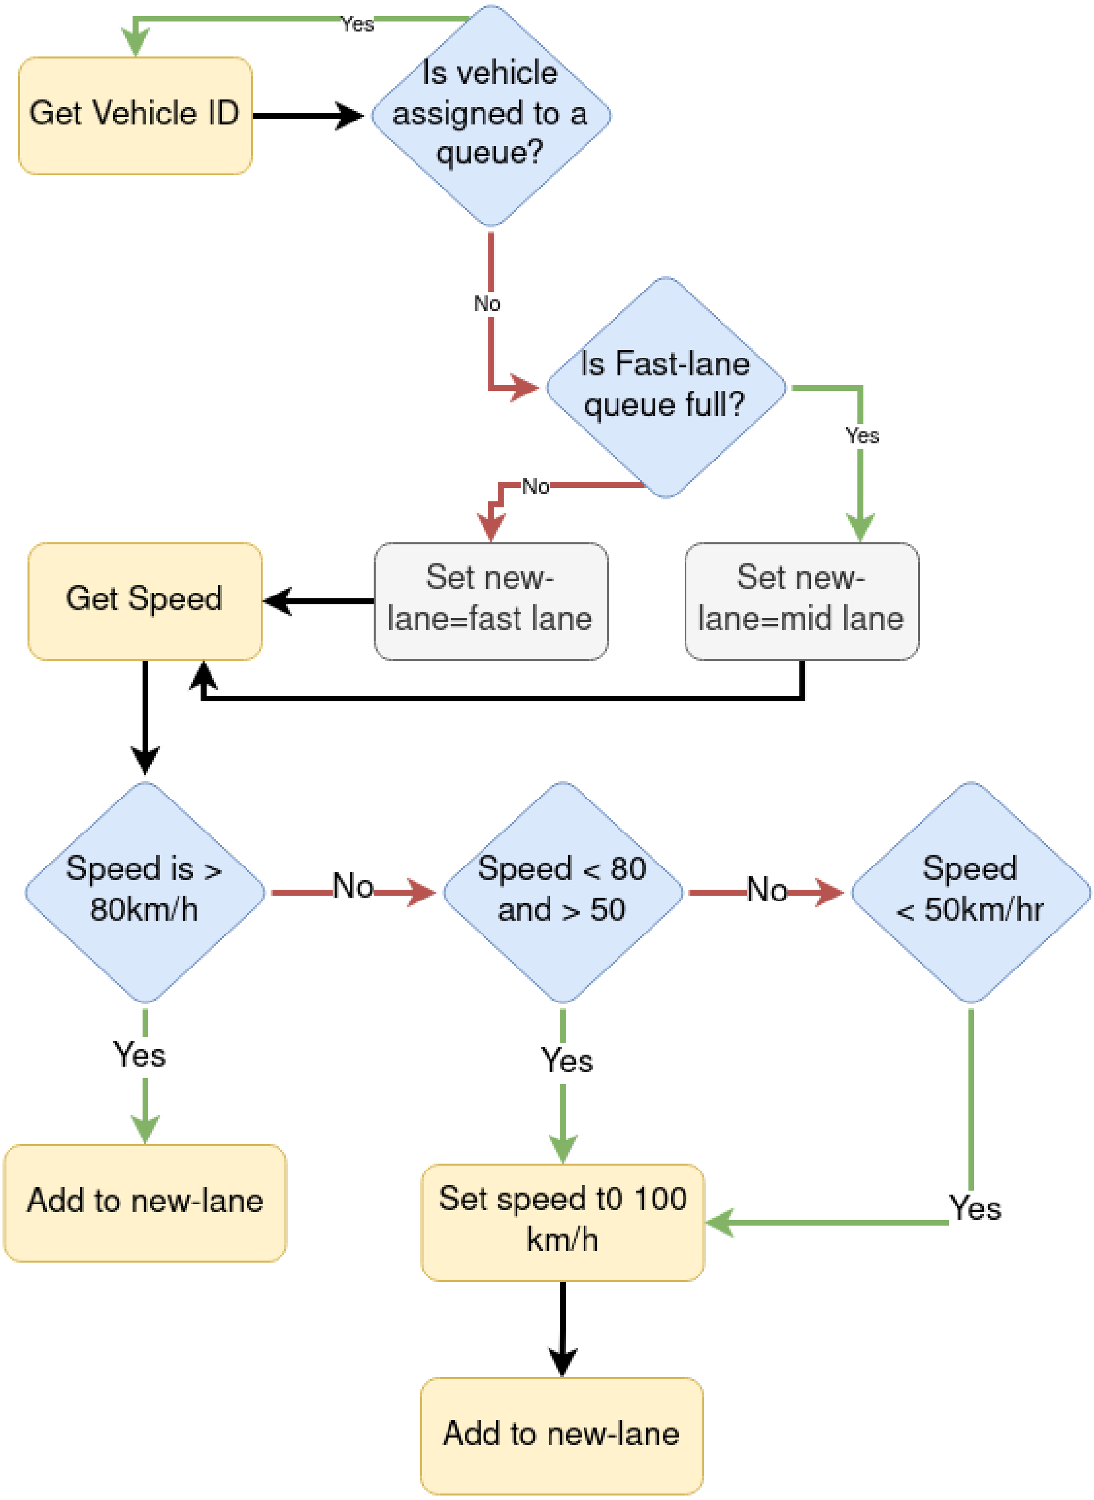

Figure 3.7: Control Flow for Connected Autonomous Vehicles

3.5 Speed Harmonisation Model Objects

The entities represented by the simulation are abstractions of real world objects. The ones considered for this project are vehicles, road and car following model. These objects form the foundational attributes upon which the simulation is built and the relationship and interactions between them influence how effective the road network is.

3.5.1 Vehicle

The automobiles on the motorway constitute the majority of the data and are the primary actors in this model. The vehicles used in this simulation are of two broad categories- one driven or influenced by human drivers and the decisions made by them. The second is an autonomous type that operates independently of human inputs and relies exclusively on communication from the central network or from other vehicles within the system. In this project, the vehicle to vehicle communications are not explicitly factored in and occur under the hood within the python script. In the real world, these interactions would occur between vehicles. The main attributes of a vehicle in the system are-

- **Identifier:** This is a unique value that is assigned to a vehicle to aid the system with referencing, distinguishing between various vehicles and creating operations specific to a vehicle.

- **Vehicle-class:** The abstract classes of vehicles determine how the object behaves on a roadway. A few examples of the values they can hold are passenger, emergency, truck, etc. The vehicle type explains certain aspects of the maximum speed, size of the vehicle and priority it gets for a specific lane. Certain types of vehicles may not be permitted on a roadway while others may have dedicated lanes.

- **MaxSpeed:** The maximum velocity that a road vehicle can travel at is defined by this metric. In SUMO this is measured in m/s. The default value is 55.55 m/s (200km/h)

- **Width:** The width of a vehicle is represented. It affects the view of the object within SUMO and would be useful in the real world when tunnels or single lane roads are used.

- **Height:** The height of a vehicle is represented in metres and is useful while the network has an underpass or tunnel.

- **Lane Change Model:** This defines the model to be used and forms the relationship between a vehicle and the Car following model.

- **Capacity:** The capacity of a vehicle can aid with differentiation with the number of people it can allow.
- **Origin:** The origin defines the point at which the vehicle enters the highway.
- **Destination:** The point of exit for a vehicle from the highway. The distance it traverses is the difference between origin and destination.

3.5.2 Route

The road network consists of a collection of edges and junctions that form a connected grid. Each road has a few properties that define how vehicles can behave on them. The road network forms the base layer for the vehicles to interact on. The road can be of different types such as single lane, multi-lane and entry and exit ramps as shown in Fig 3.8. The attributes that constitute the road are:

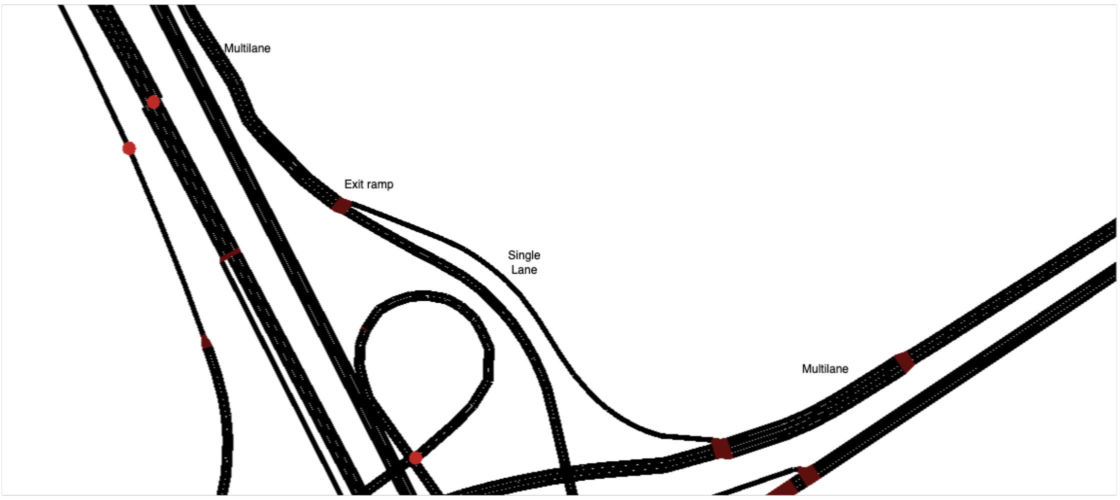

Figure 3.8: Types of routes within the road network

- **Identifier:** Identifier: A unique value to address a specific road to make references and run operations.
- **Name:** A human readable value that matches the real world name of the network.
- **Length**: The distance of the road from the starting point to the end.
- **Speed Limit:** A maximum speed that can be attained by any vehicle travelling on the highway. The speed limit defined by the road or lane overrides any other speed limit as the upper limit. This implies that no vehicle can overcome the speed cap defined by the road.

- **Shape:** The geometry of the road is communicated by this value It cannot be empty and has a minimum of two positions (Straight line).

3.5.3 Lane Changing Model

The framework used to determine how vehicles function on the motorway and their idiosyncrasies while interacting with other motorway objects are governed by the rules dictated by the Lane changing model. Within the simulation, the car following model is representative of a driver's behaviour on a road. It provides a template to determine optimal lane changes, following distances and safety considerations within a lane. In the simulation environment, all vehicles are abstractions of the drivers who make the decisions. The attributes of the lane changing model are-

- **Reaction time:** individuals react to stimuli differently and the rate at which they respond has an impact on the outcome of a traffic interaction. The reaction time sets a threshold for how long a vehicle takes to react to a change that occurs on the motorway.

- **Tau:** This represents the minimum time the driver maintains between the vehicle ahead of it. Tau is a factor of time as opposed to distance. As the tau value increases, the driving behaviour moves towards a free driving style. The values have to be greater than or equal to 0.

- **Sigma:** It represents the imperfection in the driving skill of the driver. Sigma allows the simulation to be more realistic with random changes occurring with higher sigma values. 0 denotes a perfecting driving style

- **Name:** Name: The name denotes the model used. SUMO applies the Krauss Model by default. There are multiple types that can be followed which varies the metrics of how the car behaves on the road. A good car following model assesses the current state a vehicle is in and uses that to make decisions for future states. A frequently used configuration constitutes three driving styles (LAZAR, 2019):

    1. *Free driving:* These vehicular flows are unconstrained and often achieve the speed they desire without any obstruction. It represents free flowing traffic.
    2. *Normal following:* This is the typical traffic flow system where a vehicle behind modulates its speed according to how the vehicle ahead is moving.
    3. *Emergency deceleration:* The priority of this system is to avoid collisions.

Getting into more detail about the car following models, they define the propagation of vehicles in the forward direction and replicate how a driver would behave when faced with situations and encounters on a road. The model chosen in this implementation is the Wiedemann 99 and the justification for the choice can be made by understanding how a few other commonly used models operate.

- **IDM or Intelligent Driver Model:** This model was created by Treiber (Treiber et al., 2000) and is used in the model created by Ivana and Maxine (Guériau and Dusparic, 2020). It is known for its simplicity in generating realistic profiles for vehicular acceleration which are also accurate. The acceleration of the model is calculated as a product of its velocity, distance to the preceding vehicle and step difference in velocity between the two vehicles (Pourabdollah et al., 2017).

  $$a_{\alpha}(t) = a_{\max}(1 - \left(\frac{v_{(X}(t)}{v_{des}(t)})^{\delta} - (\frac{s^{*}(.)}{s_{\alpha}(t)})^{2}\right), \tag{1}$$

  The distance to the vehicle ahead is large on a free road and the acceleration of the vehicle is governed by the limits of the road.

- **Krauss Model:** The SUMO default model is Krauss. It was developed by Stephan Krauss and calculates the velocity of the vehicle without using the acceleration to derive it. The velocity of a vehicle is calculated based on a term called safe speed which is a function of the time to the vehicle ahead, the distance between the two vehicles, the reaction time of the driver and maximum rate of deceleration. Using these parameters, a safe speed is arrived at. The Krauss model is more conservative than the rest and hence was not selected for the purpose of this evaluation as it failed to achieve the desired speeds in the time taken for other models.

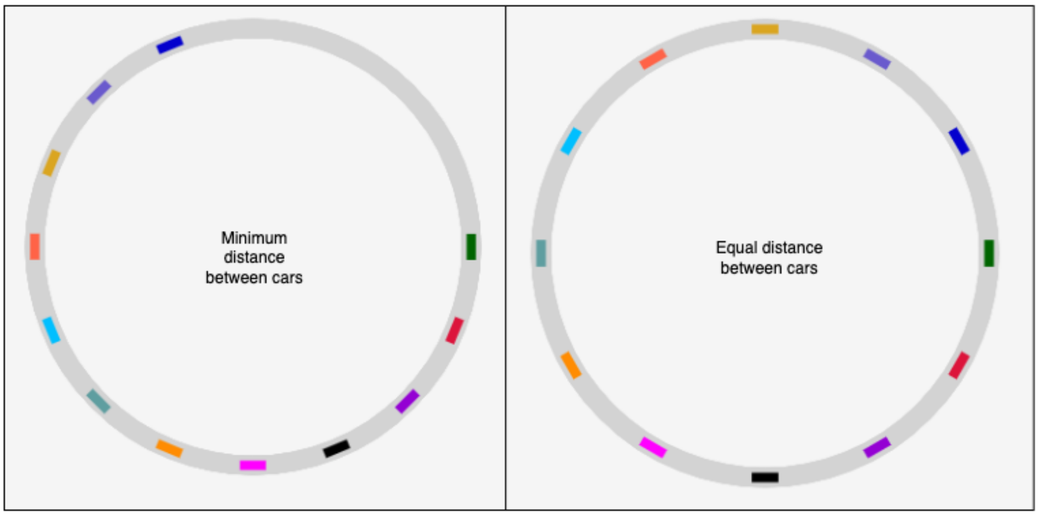

Figure 3.9: W99-Minimal Spacing vs Equidistant Car Following

- **Weidemann Model:** Introduced by Rainer Wiedeman in 1974, this model is widely used in simulations. The model has two overarching designs, one that focuses on equidistant vehicles and the other that tends toward minimising the distance to the vehicles ahead, and this can be seen in Fig 3.9. For this dissertation, a variant of the Weidermann called W99 is used as it offers aspects of granular control with regard to spacing time to the vehicle ahead, car following distance, etc (Center, 2016)

Chapter 4

Implementation

4.1 The Speed Harmonisation Model

Revisiting the discussion in the introduction chapter, the goal of the project is to create a framework for vehicles to travel on a motorway by utilising the lanes most effectively and maximising their speed. The solution proposed is to have a specific speed cap set for each lane on a motorway and every vehicle travelling on it can have a speed that is less than the preset maximum. This ties in with the concept of VSL. In addition to this, a second layer is added where the lanes function in ascending order of maximum speed while moving from the left to right. At this point the concept of speed harmonisation plays an important role, because every vehicle on the motorway is moved to a specific lane in order for it to attain the highest velocity. Since velocity is a derivation of distance and time-

$$velocity = \frac{Distance}{Time}$$

The velocity of the vehicle moving to a lane is determined by the distance it has to cover to reach its destination and the time it will stay on the new lane. As mentioned earlier, the project observes two scenarios, one that focuses on vehicles with a human driver and another that allows autonomous vehicles controlled by a central unit.

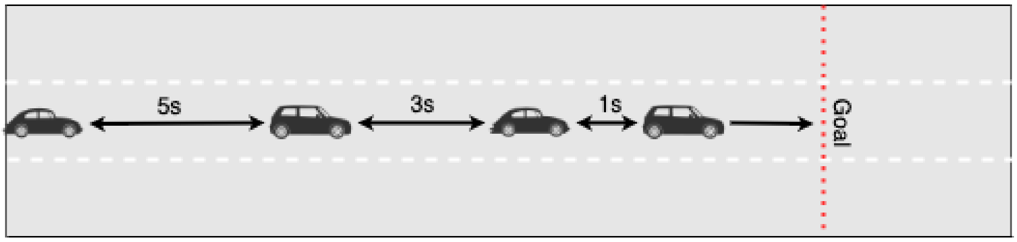

Figure 4.1: Traditional driving effect on a road way

The final visualisation or design that is achieved is similar to the way a train functions. When a train accelerates, the engine determines the speed and every carriage attached to the engine travels at the same velocity. This effectively creates a continuous flow of objects with a dependency constraint on a single entity. In traditional driving schemes and teaching, a driver is taught to observe the vehicle ahead and respond accordingly and as drivers gain more experience, they are able to observe at most 2 preceding vehicles without compromising safety. This then forms a chain reaction where the entity behind reacts to the one ahead and as traffic progresses, the image resembles Fig 4.1. But if vehicles could move synchronously, the movement would resemble Fig 4.2.

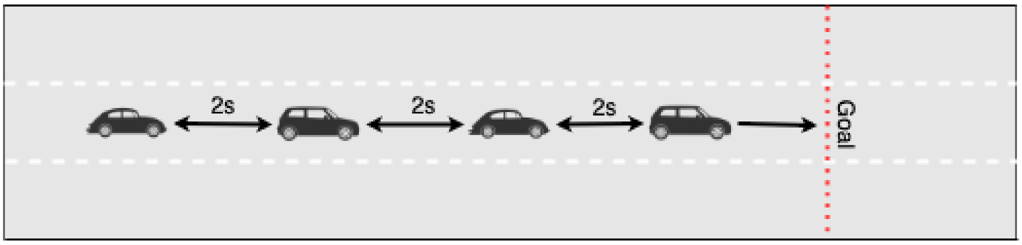

Figure 4.2: Driving with uniform acceleration and constant speed

In Fig 4.1, if every vehicle is spaced 2 seconds apart let us consider a travel time scenario. If vehicle V_1 takes 10s to reach the destination, V_2 should take 12s, V_3-14s and so on. But realistically there is a reaction time that comes into play in this system. If every driver takes 1 second to react, then the new times become V_1:12s, V_2:15s, V_3:21s. By this example we can see how the time scales up. Now not all drivers react the same way, and often vehicle types affect the rate of acceleration. Considering these factors we can say that a vehicle that is following another starts moving or continues to move with a delay that ranges between 1-4s. The effect produced by these events results in what is called a traffic snake (CGPGrey, 2016). The scenario discussed takes a single lane into account; if lane changes, careless driving and other real world scenarios are factored into this equation, the time to travel only grows exponentially.

This is an issue that can be overcome using the technique of speed harmonisation by allowing vehicles to accelerate uniformly by utilising a means of communication such that when the first car moves, the other vehicles on the lane are notified and they begin to accelerate simultaneously. The resulting motion has all the cars moving with a constant velocity along a lane. This is the targeted effect that this dissertation proposed and works toward.

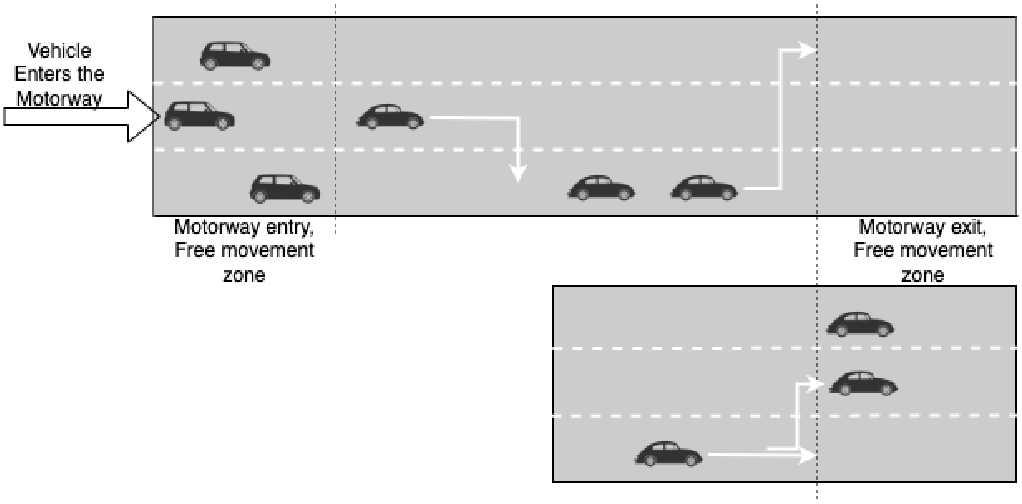

Figure 4.3: Control flow for vehicles on the motorway from entry to exit

The implementation of this model begins with taking the M50 motorway and adding an element derived from the paper by (Rios-Torres and Malikopoulos, 2017) where the entry and exit points have a control zone to ensure even distribution of traffic. The lane speeds at this point are evenly distributed without variance and as the vehicles approach the exit, they spread out based on the exit path. This is demonstrated in Fig 4.3. The region in between serves as the speed harmonisation zone where the core traffic management strategy is applied and vehicles move to respective lanes based on their travel speed.

The python script governing the simulation kicks in the moment the execution begins and for every 10 step counts (steps are the timer function in SUMO which act as points in time) the number of active vehicles in the simulation are read. The vehicle IDs are captured and stored in a temporary memory. This would translate to a database in the real world to log the number of vehicles entering the highway. For the simulation the maximum number of vehicles that can exist on the highway at any given time is 150,000. This serves as the upper limit. On portions of the motorway where the number of lanes are greater or equal to 3, when at maximum capacity, the vehicles are evenly distributed.

As the vehicle exits the control space and joins the speed harmonised motorway, the control unit, which in this simulation is the python script, tracks the vehicle and gets the current speed. Vehicles can enter the motorway at different speeds which are defined in the emitters.xml file specific to each vehicle. The system records the speed of the vehicle and assigns it to a queue based on the lane. The fast lane 3.6 has a speed range between 80-100 km/h. The mid lane has a range of 50-80km/h and the slow lane is for vehicles travelling at speeds less than 50km/h. In the HDV scenario the vehicles switch lanes as they accelerate but with connected vehicles, a speed is assigned and the vehicle is moved to the respective lane.

Within the simulation, each vehicle is assigned the maximum speed that is permitted on the motorway and is moved to the fast lane (L_f). During this process it is added to a queue with an ID of 2. Before the vehicle switches lanes, a conditional check is performed where the distance to the destination is calculated. If the distance is less than 10km, the vehicle does not change lanes. If it is greater than 10km, the time to reach the destination is calculated based on the max speed and distance and the vehicle is moved to the fast lane for that period of time. As the lane continues to fill and concentration of vehicles increases, the mid lane(L_m) is promoted to the same speed as the fast lane (L_f) and new vehicles are assigned to the mid lane with ID of 1 for the queue. The vehicles continue to progress in this manner and as they approach the end of the journey, the movement resembles soldiers in a single file Fig 4.4.

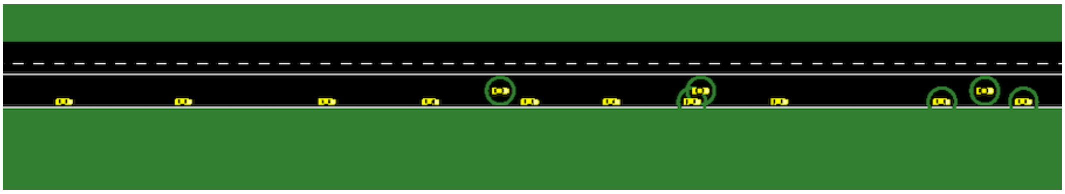

Figure 4.4: Effects of speed harmonization for a lane resulting in streamlined traffic flow.

When the vehicle draws closer to the exit, the control region once again has speeds that are V_{max} across all the lanes and the cars distribute themselves as depicted in Fig 4.3. This method prevents bottlenecks from forming on the highway as the automobiles tend to have higher concentration on the lanes with max speed as they approach this point and other lanes may be free. With this model, the overall speed vehicles attain is more than the current average. An additional benefit of this approach is certain lanes have a significantly lower volume of vehicles, this allows emergency vehicles and other priority vehicles to use the free lanes and since these vehicles have a higher speed factor, they can exceed the speed limits that otherwise apply to regular vehicles.

The baseline followed is the model created by Guériau and Dusparic (2020) where the car following model used is IDM. The IDM model is a safe model for vehicles and does not permit short following distances and lower values of spacing distance qualify as *tailgating*.

The lane changing attribute configured for the speed harmonisation system is W99 which allows really short distances to be maintained between two vehicles and this allows cars to switch lanes in situations that often are difficult. The configuration for the model used is described in Table . . . .

The results of the experiments performed yielded results that proved that speed harmonisation served as a better alternative but this comes with certain limitations.

4.1.1 Limitations of the Model

- Since vehicles have to overtake or merge into a lane, when a large stretch of cars are flowing with the intended design, new vehicles cannot join the fast lane in that area and this causes the vehicle to either speed up to overtake the leader or slow down to join the chain at the end. However, this causes inefficiencies when the length of the highway is small.
- With the car following model enforcing traditional driving rules, vehicles move back to the original lane after an overtake. This configuration requires thorough understanding of the W99 configuration state file and could not be achieved within the time frame of this dissertation. Therefore a workaround followed was to switch the slow and fast lane properties. Thus, once a vehicle overtakes, or is on a lane that is not the left lane, it merges into the vehicular flow and achieves the desired speed harmonisation. This can be visualised in Fig..

4.2 Security and Privacy Considerations

The connected vehicles infrastructure makes use of modern wireless technology including WiFi and mobile networks for communication between vehicles and infrastructure. This makes them susceptible to most types of flaws present in wireless networks, as the system is essentially a peer to peer network. The system of a connected vehicle has moving parts and typically a host operating system running within the vehicle, and a communication system that interacts with other vehicles as well as infrastructure on the road network (V2X). The security threats to the V2X systems are determined by the experience and capabilities of an attacker, motivation and could also be determined by software bugs which could lead to a malfunction. Often the incentives to destabilise the vehicular networks are vandalism, monetary gains and information/ data mining (Sui and Muehl, 2020).

The threats to vehicular systems can be categorised into active and passive attacks (Hasan et al., 2020). Active attacks involve the attacker interacting with the system to compromise it, while passive attacks do not have direct involvement of the attacker but information and data could be lost or stolen. Active attacks could include overriding the hardware of a system giving the hacker complete control over the system, or online attacks such as message tampering, Denial of service (DoS) attacks, Sybil attacks and data injection. Passive attacks could involve man in the middle schemes for message inspection, packet capture and eavesdropping. While the range of threats is large, for this paper, we shall consider 4 kinds of attacks, namely

1. Man in the middle- message inspection
2. Sybil attacks
3. False data injection

With dynamic speed harmonisation, the actors involved are the central control that sets the speed limits on the highway lanes based on the number of vehicles on the motorway. This information is sensitive and hence Man in the middle (MITM) threats as well as false data injection attacks have to be addressed.

Vehicles also communicate with each other while switching lanes in order to accelerate or decelerate to make room for incoming vehicles. This line of communication is essential and compromises here could lead to accidents and potentially detrimental situations if spoofing or sybil attacks occur.

4.2.1 A brief explanation of each attack

- **Man in the middle (MITM):** This type of attack occurs when a malicious party hijacks a session and acts as a proxy between the agents involved in communication. The attacker can choose to be an active or passive participant (Ahmad et al., 2020), but essentially intercepts every message and can eavesdrop on the data being exchanged. This allows them to disguise themselves as either party and control the type of data, as well as the information contained within the messages.

  With an attack of this kind, the eavesdropper could analyse patterns and signals being sent from the central control and find the IDs of the vehicles within the network. This would then allow them to impersonate any of the vehicles if they opt to do so. It gives them visibility over the type of speed ranges and identifies when the specified range is transmitted. With an active attack here, the proxy could alter the speed that's transmitted to the road network.

- **False Data Injection:** This type of attack involves an attacker intercepting the network traffic and adding false data to the data streams. In a V2V scenario (Bécsi et al., 2015), when a vehicle is switching lanes, a false actor could tamper with the lane switching indication and send the incorrect signal to the other vehicles. A worse scenario would be the hacker gaining access to the internal network of a vehicle and tampering with the messages sent within the system, which could cause untoward reactions.

- **Sybil Attacks:** A sybil attack is a type of identity impersonation where an attacker assumes multiple fake identities and inserts fake nodes into the network. By adding a large number of fake participants into the network, the malicious user can assume a larger attack surface by spamming the network with traffic resulting in a DDoS (Distributed denial of service) and false voting in cases of leader election. Within the autonomous motorway network, this could lead to false information being passed to the control centre and adversely affect the permissible speed, or worse, a DDoS attack could crash the vehicular network leading to accidents.

4.2.2 Addressing Security and Privacy

Protecting wireless V2X communication using cryptographic credentials is a first layer of security that can be applied. V2X protection mechanisms can be proactive or reactive. Based on the type of mechanism, the implementation schemes may vary. Proactive security would involve the use of enforcing a security policy with PKI (Private Key Infrastructure), digital signatures and certificates to verify the clients connecting to the motorway network.

When a vehicle joins the network, a physical ID can be scanned and a digital identity can be passed to identify the vehicle. A tag can be active while the vehicle is on the motorway and when it exits, the state changes. With authentication and identification, the chance of an unidentified exchange can be reduced and the identified would allow verification of messages. While these schemes are great for stationary networks, high speed communication on a motorway would require better memory management and the infrastructure would have to scale to support the high volume of data exchanged in a peer to peer network.

Reactive security would vary based on the type of data entering the system or the type of actor. With vehicular communication, a pattern can be established for the entities on the roadway and when a node begins to act out of sync regularly, a system can be in place to remove the bad actor or verify its identity. If the actor fails to confirm the identity, an alarm can be raised. Similarly, when looking at the data entering the system, the frequency and type of data can be analysed to identify imposter nodes or denial of service attacks.

False Data injection mitigation

The way to identify data injection attacks is to analyse response patterns of the nodes in the system (Hasan et al., 2020). Using short term private keys with each entry into the motorway and pre-authenticated signatures would allow the motorway control to identify intrusions. A mechanism called watchdog uses the principle of a vehicle trusting the neighbouring vehicle. Data packets will not be forwarded if a collision or attack is detected and vehicles would continue to work on the present instruction set. If a vehicle repeatedly drops packets based on a preset tolerance, the actor is considered malicious. The behaviour of a vehicle can also be assessed by the frequency and type of data it sends when it joins the network and two instances of verification can be carried out. The first would be comparing the incoming data to existing traffic models and ensuring the metrics do not stray from the established patterns. The second is to monitor the newly joined vehicle and if the type of messages change, an alert can be raised.

In addition, since vehicles are real world objects, when repeated offences are observed, the digital identity can be flagged and this can help identify a blacklisted identity before it joins the network.

Sybil Attack Mitigation

Based on the paper by Grover (Grover et al., 2011), the fake actors in a system would exist in the same space as the valid entities, especially true for vehicular networks. The scheme proposes a cooperative protocol that uses group signatures to maintain the privacy of vehicles and all mobility tracing. This is accomplished by vehicles around an imposter vehicle broadcasting warning messages with partial signatures and from this, a complete signature can be derived. The data that is collected from the sensors can be used to validate a node.

Another means of validation is to track the GPS positioning of vehicles within the network, and use the strength of the signal to determine its position. Often a bad actor would focus on certain regions of the motorway, since constant position switching would cause additional overhead for the attacker to maintain. This however has a flaw when a signal is jammed by the attacker or a sophisticated impersonation attack is used. Another detection scheme that can be employed has a centralised authority like an RSU to detect nodes. The scheme assigned metrics to the vehicles such as

- Claimer : The vehicle claiming an identity.
- Witness: a node that receives a signal and measures the proximity using the strength of the incoming signal.

- Verifier: the vehicle that collects the signal measurements to triangulate the claimed position of a vehicle.

The message can contain a timestamp, and if multiple messages come in with similar and periodic timestamps, the central authority establishes it as a valid node.

Another protocol enlists the use of session key-based certificates. Each vehicle has an identity registered with a global provider and the vehicles entering the motorway then generate anonymous IDs that are validated by a server within the network. Any recipient vehicle can verify the message by comparing it to the other vehicle's identity on the local server. This system can use PKI and the local certificate to compare messages received from the RSU. The principle is similar to Public and private keys used in data exchange.

Chapter 5

Evaluation

This chapter examines the effectiveness of the speed harmonisation model. The model is evaluated against a baseline that was created by Guériau and Dusparic (2020) and is compared to the output metrics of travel time and average speed to evaluate its benefits on a motorway. The constraints that are in place are explained in Section 3.2. The performance metrics that are evaluated through the experiments and simulation runs are-

- Average travel time across a motorway
- Average speed of vehicles on the motorway.
- Traffic flow patterns on the motorway to minimise bottlenecks and congestion.

5.1 Testing the designed system

To evaluate the performance of the system, 3 types of tests are carried out. The first test run uses Human Driven Vehicles and compares it to the baseline model. The second test run uses Connected Autonomous Vehicles and compares this to the HDV scenario and the baseline. The final test considers a subsection of the motorway and showcases the movement of vehicles along the highway for a smaller volume of traffic and draws out the details of how a vehicle changes lanes and the step by step progression along a highway to visualise the benefits of the speed harmonisation model.

The tests are carried out using SUMO and the metrics are captured using two output files and one log file. The first output file provides information about the lane change history of each vehicle that travelled through the motorway as well as the direction of the vehicle's journey. The second file provides information about the trip that each vehicle makes. This includes data such as the time the vehicle joined the network, the time it exited the network, the speed of arrival and departure. In case the vehicle encounters any collisions it records the time the vehicle stops. In an instance of traffic jams, it records the waiting time. The log file records the details of the motorway for each timestamp of execution. This is essential to create plots and compare the differences in each model as it provides details about the position of a vehicle at a specific execution timestamp as well as the speed it was travelling at that point in time. This data is captured as lane specific information, which also provides the count of vehicles on each lane which is useful for measuring the traffic density.

5.1.1 Testing Metrics

- **Average travel time across a motorway**
  Every vehicle on the motorway has a point of entry and a fixed destination. The travel time is the duration it takes a car to traverse from the starting point till the end of its journey which is its destination. The travel time of a vehicle is calculated by using the execution time within SUMO and converting that into a real world time. The time factor in SUMO translates every 1s in simulation time to 10s in real world time. The average travel time is calculated by calculating the median of all the vehicular travel times on the motorway for the specific execution.

- **Average speed of vehicles on the motorway**
  A vehicle travelling on a motorway does so with a specific speed. While entering the motorway, the speed of the vehicle can vary. Based on the vehicle type, the rate of acceleration can also change which is affected by parameters such as weight, size of the vehicle, number of passengers. The combination of these factors lead to an overall speed that the vehicle can attain. Within SUMO, the default acceleration of the vehicles is set to 2.6 m/s2 and the net-state file records the change in speed at each timestamp. Using this, the average speed of each vehicle can be determined. Taking a median of the average speed of all the vehicles on the motorway from the start to the exit is used to deduce the average speed of vehicles on the motorway.

- **Traffic flow patterns**
  A unique aspect of the speed harmonisation model is the shape the traffic flow takes as it conforms to the rules of the model. As stated in the previous chapter, each vehicle is assigned a lane based on the speed it's travelling at- for HDVs and for CAVs, the speed is assigned based on the available lane with the fastest permissible speed. When the traffic density is less, this leads to one lane having a large volume of traffic leaving other lanes free and as the volume of traffic increases, the lanes mimic the traffic flow pattern of the first lane. This leads to a sort of flow that

mimics an infantry troop marching with equal spacing and a pattern that follows the row in front. This specific model leads to a more streamlined vehicular movement that helps overcome phantom intersections, which is when a single vehicle creates a virtual intersection due to the stop start motion causing a similar effect on the vehicles behind. The vehicle at the end of this chain is forced to stop or slow to a near halt because it cannot move.

5.2 Results

The tests were conducted using 5-7 simulation runs of each scenario described in Section 4.1. The scale of the scenarios were for hundred and fifty thousand vehicles (150,000) and three hundred thousand vehicles (300,000) that spanned the motorway. For each of the tests run, the input metrics that are varied include the Driver reaction time, distance to be maintained from the vehicle ahead. Metrics such as the Car Following Model and lane speeds remain the same for the baseline which uses IDM and Speed Harmonisation which uses W99 respectively. The first two test use the entire stretch of the M50 Fig 5.1 to derive its results and the final test utilises a section of the motorway with only straights and no intersections or turns.

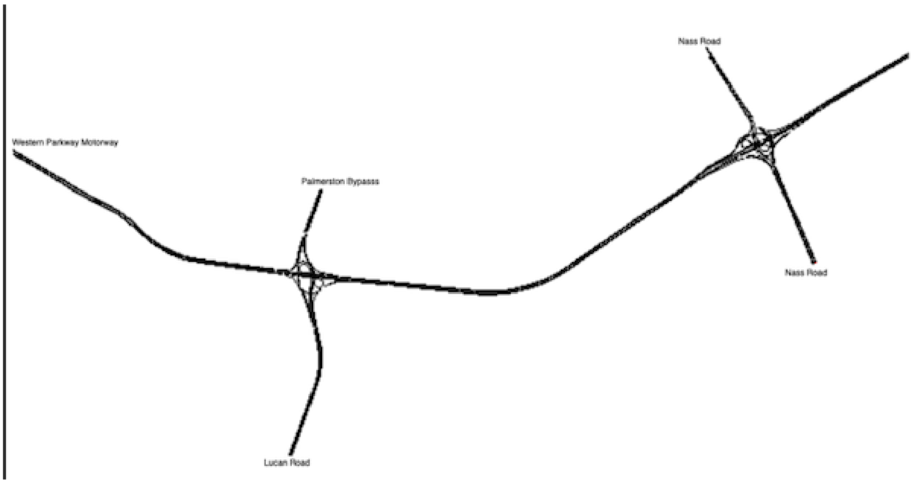

Figure 5.1: M50 Model

5.2.1 Average Travel Time

The model of the M50 covers a span of 116km of road along the longest stretch. From the simulations run the average time taken for vehicles to cover this distance proved to be lower using Speed Harmonisation. The current average time on the M50 is 56 mins (Carswell, 2017) and the new model lowers the travel time by 21%. Fig 5.2.

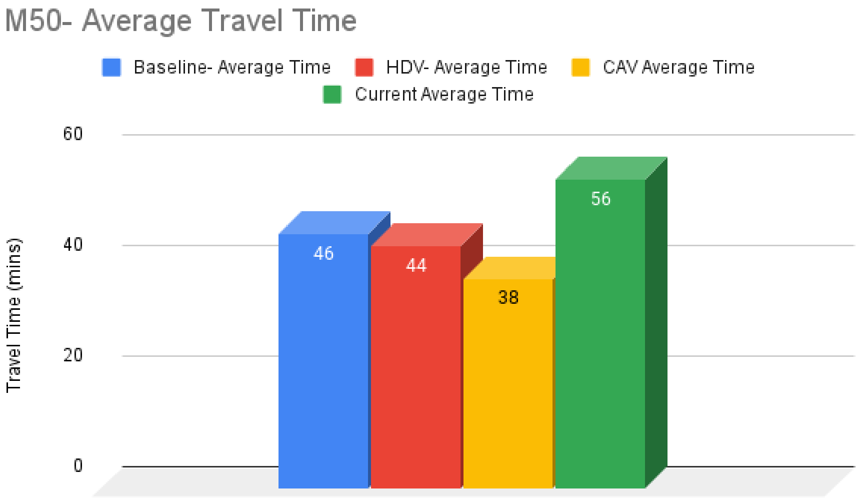

Figure 5.2: Average travel time comparison

This drop in travel time can be attributed to the lane changing model used, which allows vehicles to change lanes based on speed and then stick to the lane. Since vehicles do not overtake or shuffle between lanes, the impact of slowing down and speeding up of vehicles is significantly lower.

A drawback of the speed harmonisation model, however, is that it is ineffective on short roadways. A test carried out using 70 vehicles on a 3 lane highway required a minimum length of 17km to be effective for the vehicles to attain their maximum speed.

5.2.2 Average Speed

Using the data from the simulation, the average speed variance of the Speed Harmonisation model is compared to the baseline. The input metrics are kept constant except for the values defined in Table 4.1.

Table 5.1: Input metrics for model

	HDV	CAV	Baseline
Following Distance	0.5	0.3	0.5
Driver In-expertise	1	0	0.7
Lane Changing Model	W99	W99	IDM

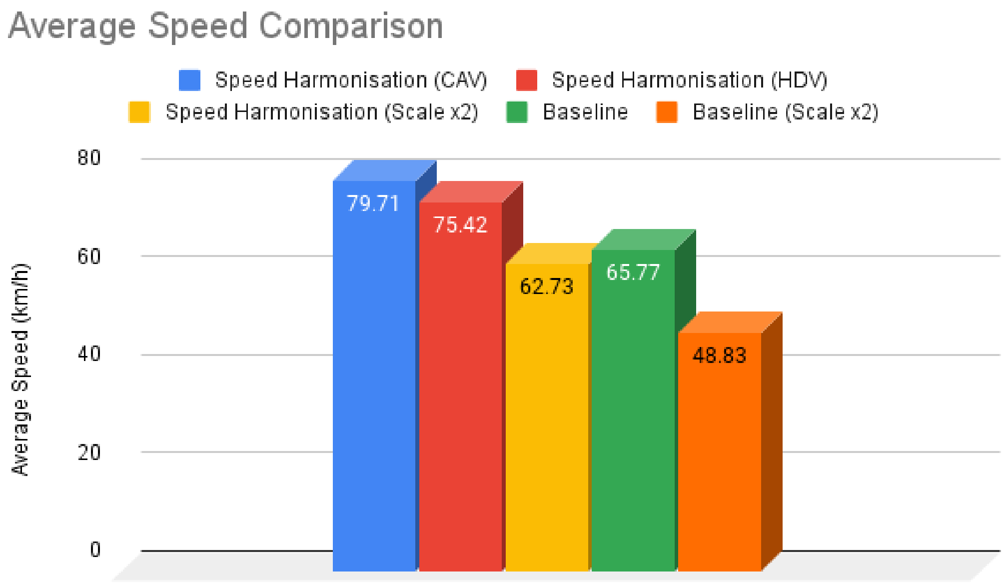

Figure 5.3: Average Speed of vehicles based on the different models

The simulation is run for 15000 steps or 4.5 hours real world time. While comparing the results obtained from the simulation runs, the models have significant variance in how vehicles change their speeds over time. The speed harmonised model stabilises with very small deviations in speed after one-fourth of the simulation time elapses and attains an average speed between 75-80 km/h. The baseline however has a much lower average speed of 65 km/h. This is because the model has vehicles distributed across lanes based on driving preferences of the driver and congestion tends to occur on the roadway near junctions and roundabouts. The graph generated from SUMO to show the average speed distribution proves this- Fig 5.5.

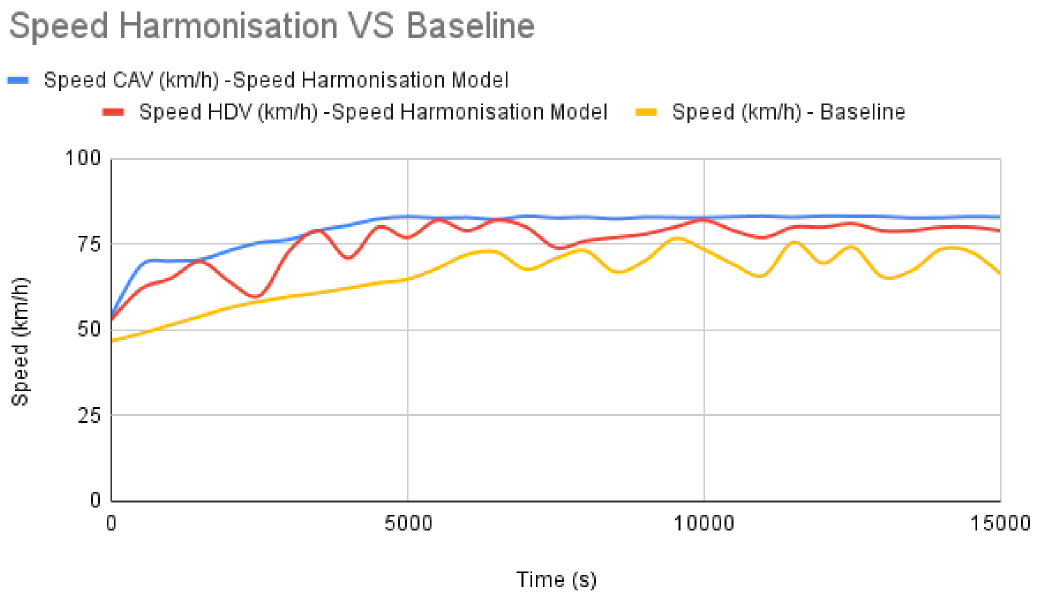

Figure 5.4: Speed Harmonisation scenario speeds against the baseline

While comparing the speeds of each scenario in Fig 5.4, it is evident that the HDVs have a larger variance in speed at the beginning because vehicles entering the highway from the multi-lane entry points have to move the lane with the highest speed based on the lane change algorithm. However, the distribution of vehicles across the lanes causes cars on the neighbouring lanes to be blocked and hence cannot move to the fast lane during the initial iterations of the simulation. When a vehicle attempts to change lanes, it notices another car blocking its path and has to wait for the next iteration to attempt a switch in lanes. This is the cause for the deviation during the beginning of the HDV scenario.

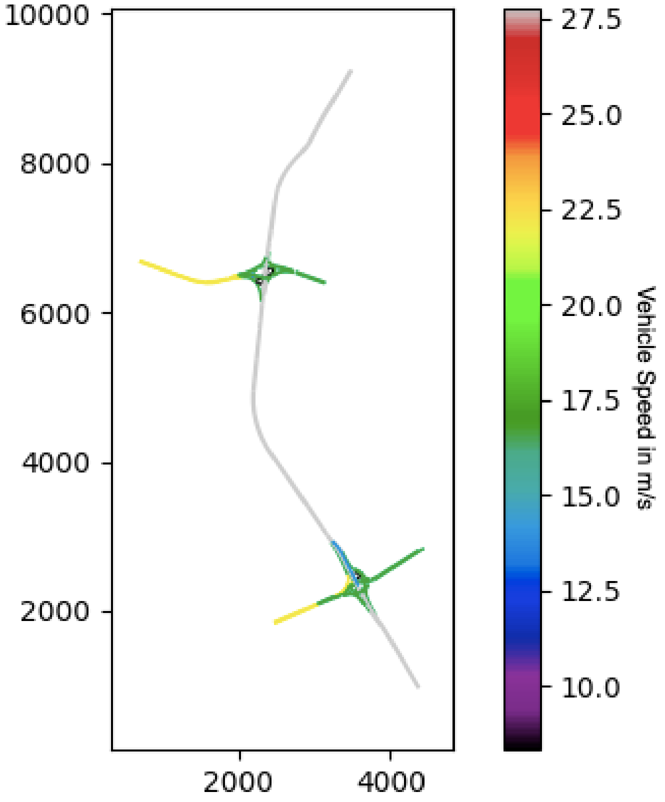

Figure 5.5: M50 speed distribution

The baseline on the other hand sees a drop in speed as the simulation enters its latter half because the volume of vehicles increases. This leads to congestion and vehicles slowing down as they approach the M50 roundabout. This forms a vehicle buildup at these points which has an effect on the overall average speed.

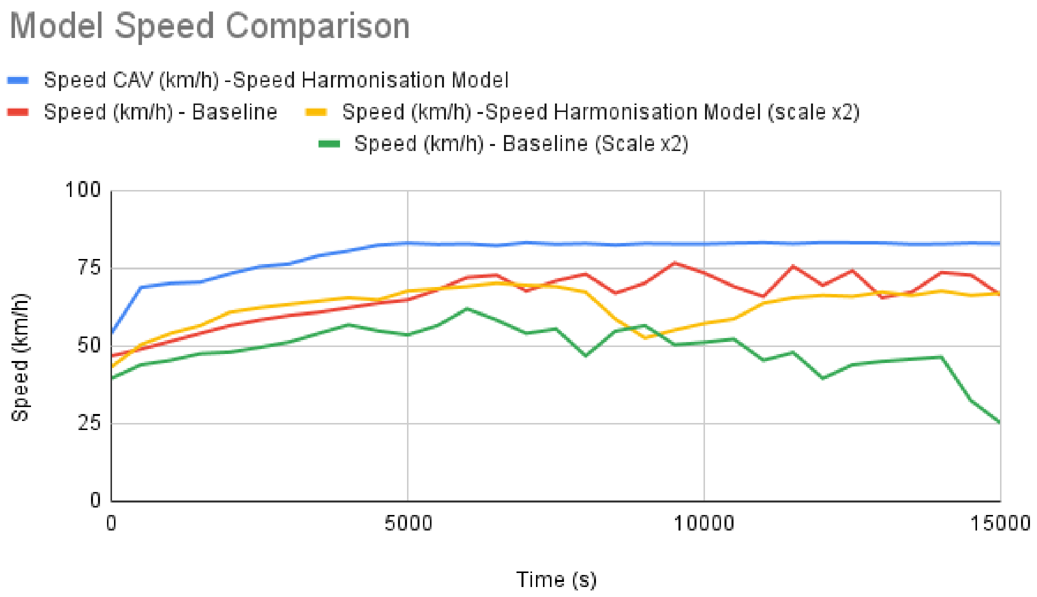

Figure 5.6: Comparison of speeds based on traffic density

Looking at Fig 5.6, the yellow and green lines represent the CAV and Baseline scenarios at twice the traffic density. The algorithm is designed to allow only 150,000 vehicles as defined by the M50 authorities as the maximum capacity. Since the load increases, the high speed movement of vehicles leads to collisions early on (observe time 8000) and after the collision occurs a small percentage of vehicles come to a halt and a large volume slows down. This leads to an overall drop in the average speed. Taking a look at the baseline at double capacity, more collisions begin to occur towards the end, the delay being due to a lower overall speed, but once collisions begin, they continue to occur, as driver behaviour is erratic and impulsive. This causes a severe drop in the overall highway speed towards the end of the simulation.

Further experimentation, using the Speed Harmonisation model, was conducted by breaking down the M50 into subsections, taking the straight line roads as one part and the curves as another section. Running simulations on these miniature highways resulted in a key observation. The model running the straight line simulation attained the maximum highway speed of 100km/h 60% faster than the baseline model. On the contrary, the curved roads either matched the baseline or performed worse because vehicles could not sustain high speeds along the curves and turns. Based on this, it is possible to infer that the high speeds on the M50 are attained from long unwinding stretches, which makes the speed harmonisation model designed in this dissertation effective for straight roads. Fig 5.7 shows a small comparison of how vehicles distribute themselves on curves vs straight roads.

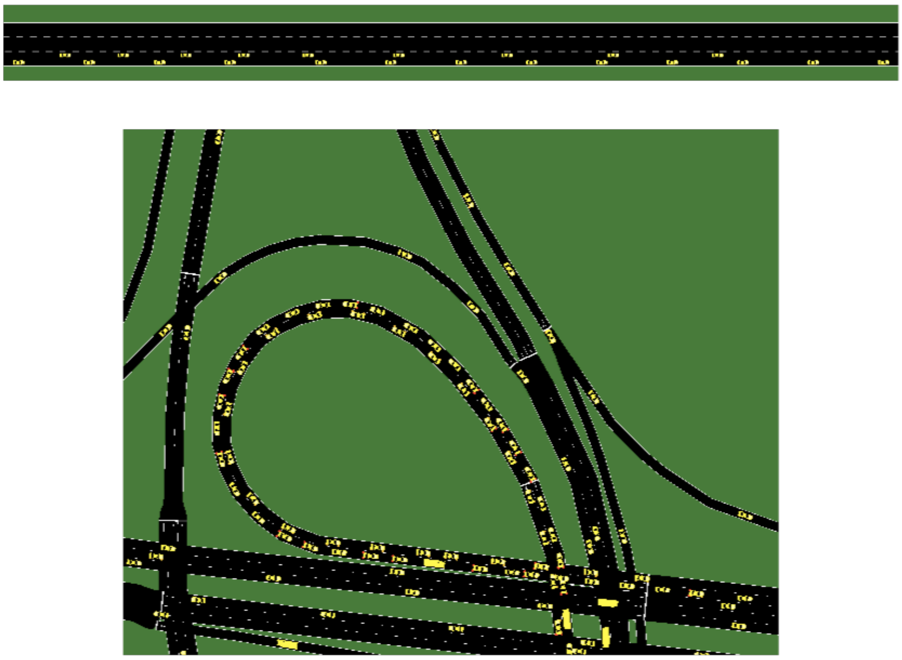

Figure 5.7: Distribution of vehicles on curves and straight roads

5.2.3 Traffic Flow Pattern

Traditional driving techniques follow a driving model where vehicles on a highway move to a right lane, overtake and return to the left lane. Sometimes the vehicles distribute themselves on different lanes based on the speed they are travelling at where faster moving vehicles travel on one lane and slow vehicles on another. However, there is no rule that states specifically how this distribution should be. Drivers also use their own discretion and impulses to drive along a certain lane. This causes issues because certain vehicles would cut across lanes to turn at an intersection, or when a car changes lanes, the vehicle behind would have to slow down to accommodate the change. This is a traffic flow pattern.

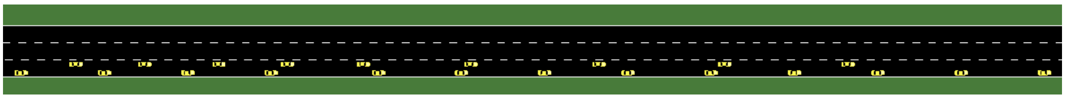

Figure 5.8: Vehicles preparing to merge to a fast lane

Using the speed harmonisation model, a new pattern is established on the motorway. Since overtaking is not a feature considered, the issue with changing and returning to a lane is overcome. Furthermore, a vehicle does not change lanes if it is close to its destination to avoid cutting lanes. And finally, cars driving on a lane form a sequential flow mimicking the movement of a train and its carriages or an infantry troop marching. This method of movement helps overcome issues like phantom intersections and confusion in regard to a driver's intent. The model developed has vehicles nearing the end of their journey, distribute themselves as shown in Fig 4.3.

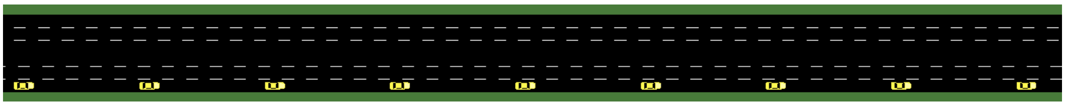

Figure 5.9: The result of speed harmonisation with a streamlined flow of vehicles on a lane

The cars merge (Fig 5.8) into the designated Fast lane when an opening is found and stay on the lane till the end of the journey Fig 5.9. This approach to traffic helps the system alleviate bottlenecks on the road network, streamlines the flow of vehicles at the exit and creates a more systematic way for cars to move. An additional advantage presented by the model is, lanes that have a lower density of automobiles can be used by emergency vehicles to travel from one point to another without causing other vehicles to lower their speed or move aside.

Chapter 6

Conclusions & Future Work

Based on the simulations run and the results observed from the tests, the speed harmonisation model presents clear benefits for implementation on a motorway. The average travel time is reduced by 21% using a constant speed specific to a lane. The vehicles traversing through the motorway were also able to travel at speeds close to 100km/h which was 12-14km/h higher than the baseline speed on the motorway. While the model is built under ideal conditions, it still shows more promise than the existing highway travel times. Based on the data recorded by Carswell (2017), the current average time taken for a car to travel across the M50 motorway is 56 mins. This time is reduced to 40 mins with speed harmonisation, and once CAVs can be introduced onto motorways, a higher level of control means this time can be reduced further.

The streamlining of traffic flow by separating cars into different lanes, reworks how vehicles presently operate on a motorway. The freeing up of certain lanes and allowing traffic to move at high speeds on other lanes results in a better strategy for rerouting traffic in an emergency and allows priority vehicles like ambulances or fire trucks to move through the motorway without having a huge impact on the overall travelling speed and time of other vehicles.

The solution however, presents a few issues and challenges that have to be addressed. The first issue is that the higher vehicular speeds can result in more serious accidents if a driver is not careful or is inexperienced and this would cause the lane to be blocked. Secondly, closing a lane at either extreme does not cause an impact to the traffic flow model, but if any of the mid lanes are closed, the lane switching system is impacted and vehicles would have to traverse in a disconnected flow or the lane switching strategy would have to be employed at the point of diversion. The final challenge presented by this model is that it is only effective for road networks with a large stretch of straight roads. The straight line speeds that the vehicles can attain is the highlight of this model and when the roadway has multiple turns or deviations, it causes the algorithm to perform poorly and vehicles cannot corner at the same straight line velocity. This would cause a traffic buildup at these points.

6.1 Future Work

Autonomous Vehicles are expected to change the landscape of motorways in the near future. The speed harmonisation strategy used here can present a lot of benefits with enhancements to the model.

The data captured in the lane-change file can be used to plot the path of a vehicle and is useful to link a vehicle's identity in instances where authorities need to track a vehicle. The lane change information is also useful for accident analysis. If multiple accidents occur on a motorway, the data from this file is useful to analyse how changes in lane take place at a specific junction and improvements can then be made either to the model or the highway infrastructure. For instance, an exit ramp along a highway has vehicles moving at speeds close to 100km/h and sudden changes in vehicles lanes are causing widespread breaking, the data would show the rapid change of speed and lane switches. This can then be used to update the algorithm by creating a new rule set for that section of the motorway.

With the rapid onset of industries moving towards connected autonomous vehicles, the speed harmonisation model can be used to provide a granular level of control. Used in conjunction with a digital twin, the model can equip a motorway for proactive speed changes. Regular users of the motorway would serve as input data and VSL patterns can be established for different times of the day to allow traffic to move even faster. In addition, with the higher level of control in CAVs, the speed limit on the motorway can be raised without a higher risk of accidents.

If the past decade has taught us anything, it's that security can no longer take a back seat and needs to be a co-passenger with the efforts being put into developing the systems. While this paper addresses a few threats that could affect the speed harmonisation and communication channels specific to my dissertation, the reaches of privacy and security exceed the bounds we have currently discussed. With new laws being laid down for peer to peer networks (GDPR rules), evidence indicates that the security risks of CAVs will not be overlooked. While the overlying issue of adversarial AI looms, the immediate threats can be controlled.

Bibliography

Ahmad, F., Kurugollu, F., Adnane, A., Hussain, R., and Hussain, F. (2020). Marine: Man-in-the-middle attack resistant trust model in connected vehicles. *IEEE Internet of Things Journal*, 7(4):3310–3322.

Beehanl (2021). What are the new variable speed limits on the m50? aa ireland.

Bécsi, T., Aradi, S., and Gáspár, P. (2015). Security issues and vulnerabilities in connected car systems. In *2015 International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, pages 477–482.

Carlson, R. C., Ragias, A., Papamichail, I., and Papageorgiou, M. (2011). Mainstream traffic flow control of merging motorways using variable speed limits. In *2011 19th Mediterranean Conference on Control Automation (MED)*, pages 674–681.

Carswell, S. (2017). M50 blues: Ireland's busiest road, dublin's biggest car park.

Center, G. A. (2016). Vehicles and routes.

CGPGrey, C. (2016). The simple solution to traffic.

Cheng, Z., Lu, J., and Li, Y. (2018). Freeway crash risks evaluation by variable speed limit strategy using real-world traffic flow data. *Accident Analysis Prevention*, 119:176–187.

Dowling, J. (2020). National primary route 50.

Goulet, N. and Ayalew, B. (2020). Impacts of distributed speed harmonization and optimal maneuver planning on multi-lane roads. In *2020 IEEE Conference on Control Technology and Applications (CCTA)*, pages 305–311.

Grover, J., Gaur, M. S., Laxmi, V., and Prajapati, N. K. (2011). A sybil attack detection approach using neighboring vehicles in vanet. In *Proceedings of the 4th International Conference on Security of Information and Networks*, SIN '11, page 151–158, New York, NY, USA. Association for Computing Machinery.

Guériau, M. and Dusparic, I. (2020). Quantifying the impact of connected and autonomous vehicles on traffic efficiency and safety in mixed traffic. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8.

Hasan, M., Mohan, S., Shimizu, T., and Lu, H. (2020). Securing vehicle-to-everything (v2x) communication platforms. *IEEE Transactions on Intelligent Vehicles*, 5(4):693–713.

Ishikawa, S. and Arai, S. (2016). Cooperative learning of a driving strategy to suppress phantom traffic jams. In *2016 IEEE International Conference on Agents (ICA)*, pages 90–93.

JAMS, W. C. T. (2022). What causes traffic jams.

Jose M. Armingol, Cristina Olaverri-Monreal, F. and Vicente Milanés, D. M. (2018). Cooperative systems for autonomous vehicles. *Journal of Advanced Transportation*, 2018(Article ID 8714549):1.

Lahmiss, H. and Khatory, A. (2020). Variable speed limit (vsl) system applications in motorways. In *2020 IEEE 13th International Colloquium of Logistics and Supply Chain Management (LOGISTIQUA)*, pages 1–5.

LAZAR, H. (2019). Comparison of microscopic car following models. In *2019 International Conference on Systems of Collaboration Big Data, Internet of Things Security (SysCoBIoTS)*, pages 1–6.

Liu, B., Ghosal, D., Chuah, C.-N., and Zhang, H. M. (2012). Reducing greenhouse effects via fuel consumption-aware variable speed limit (fc-vsl). *IEEE Transactions on Vehicular Technology*, 61(1):111–122.

Liu, X., Dai, S., and Lu, S. (2009). Macroscopic model for interrupted traffic flow of signal controlled intersection. In *2009 ISECS International Colloquium on Computing, Communication, Control, and Management*, volume 3, pages 78–81.

McQueen, B. (1994). Smart vehicles and intelligent highways-an integrated system? In *Proceedings of VNIS'94 - 1994 Vehicle Navigation and Information Systems Conference*, pages A/7–.

Nason, N. (2021). An intelligent deal for infrastructure.

Pourabdollah, M., Bjärkvik, E., Fürer, F., Lindenberg, B., and Burgdorf, K. (2017). Calibration and evaluation of car following models using real-world driving data. In *2017*

IEEE 20th International Conference on Intelligent Transportation Systems (ITSC), pages 1–6.

Rios-Torres, J. and Malikopoulos, A. A. (2017). A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps. *IEEE Transactions on Intelligent Transportation Systems*, 18(5):1066–1077.

Sui, A. and Muehl, G. (2020). Security for autonomous vehicle networks. In *2020 IEEE 3rd International Conference on Electronic Information and Communication Technology (ICEICT)*, pages 67–69.

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

Villanueva, F. J., Albusac, J., Jiménez, L., Villa, D., and López, J. C. (2013). Architecture for smart highway real time monitoring. In *2013 27th International Conference on Advanced Information Networking and Applications Workshops*, pages 1277–1282.

Vrancken, J. (2005). A coordination algorithm for road traffic control measures. In *Proceedings. 2005 IEEE Networking, Sensing and Control, 2005.*, pages 34–43.

Appendix

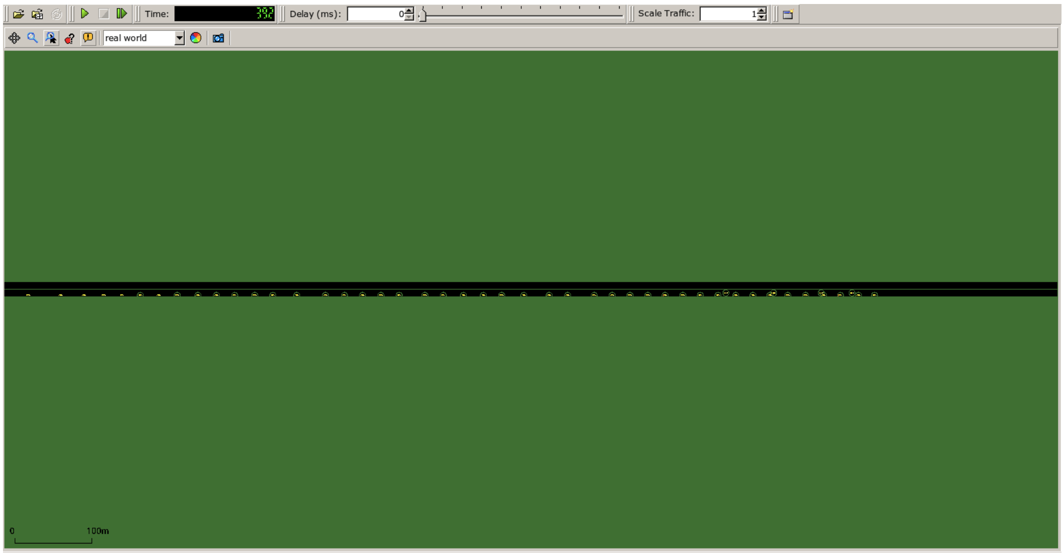
Time: 392
Delay (ms): 0
Scale Traffic: 1
real world

Figure 1: The result of speed harmonisation at the end of the motorway

Figure 2: A zoomed in view of the vehicles following the synchronous traffic flow

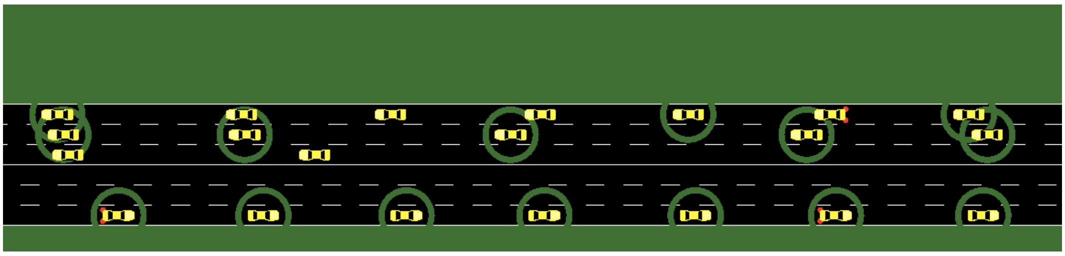

Figure 3: Vehicles on the lanes preparing to merge into the Fast Lane

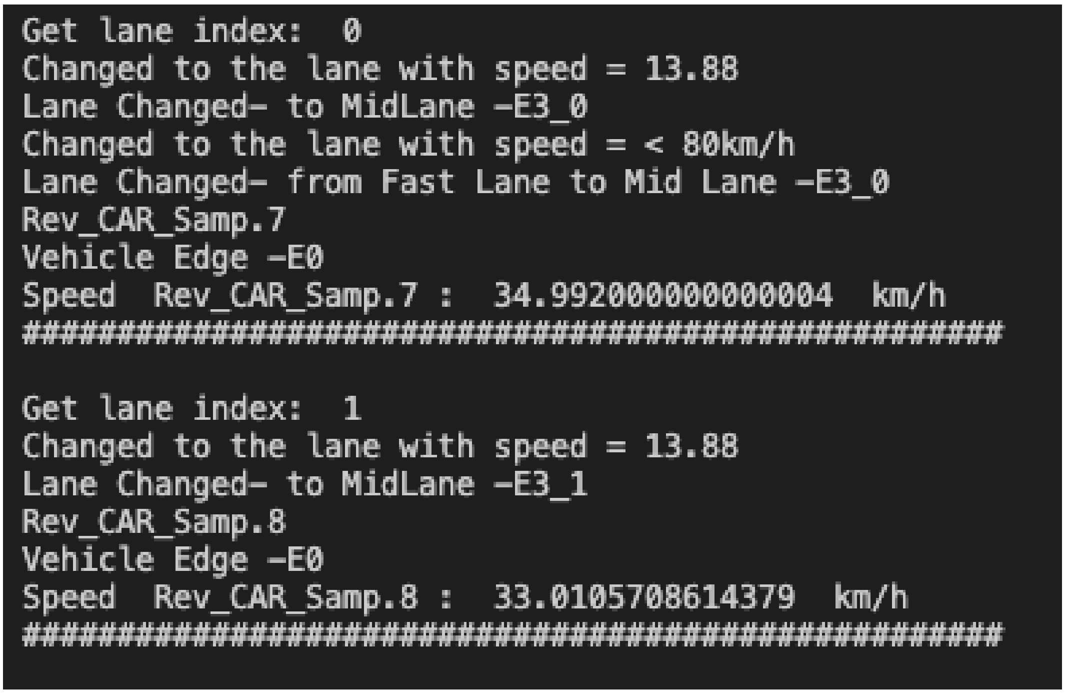

```
Get lane index:  0
Changed to the lane with speed = 13.88
Lane Changed- to MidLane -E3_0
Changed to the lane with speed = < 80km/h
Lane Changed- from Fast Lane to Mid Lane -E3_0
Rev_CAR_Samp.7
Vehicle Edge -E0
Speed  Rev_CAR_Samp.7 :  34.992000000000004  km/h
####################################################

Get lane index:  1
Changed to the lane with speed = 13.88
Lane Changed- to MidLane -E3_1
Rev_CAR_Samp.8
Vehicle Edge -E0
Speed  Rev_CAR_Samp.8 :  33.0105708614379  km/h
####################################################
```

Figure 4: Sample Output of the simulation to show a change to the Mid lane.

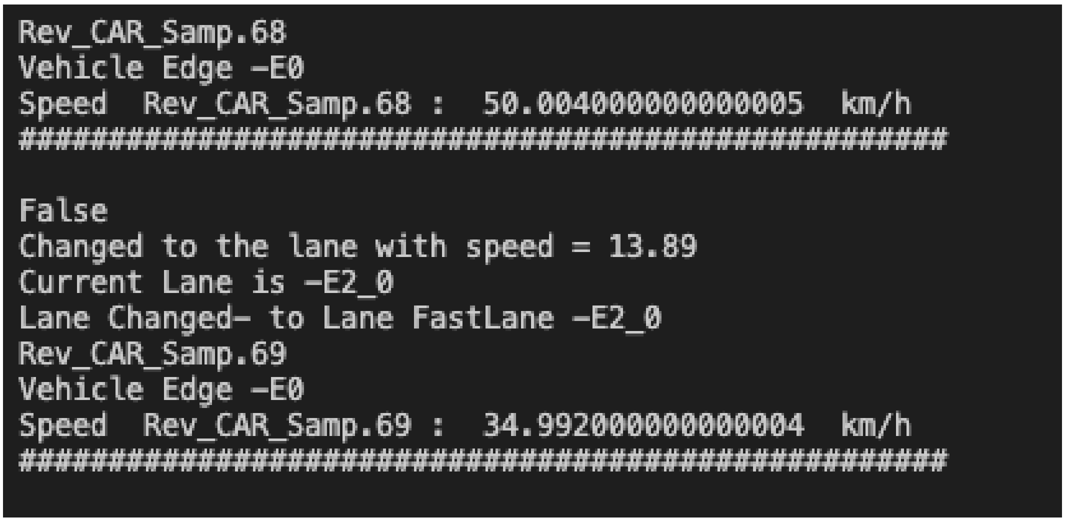

```
Rev_CAR_Samp.68
Vehicle Edge -E0
Speed  Rev_CAR_Samp.68 :  50.004000000000005  km/h
####################################################

False
Changed to the lane with speed = 13.89
Current Lane is -E2_0
Lane Changed- to Lane FastLane -E2_0
Rev_CAR_Samp.69
Vehicle Edge -E0
Speed  Rev_CAR_Samp.69 :  34.992000000000004  km/h
####################################################
```

Figure 5: Sample Output of the simulation to show a change to the Fast lane.