

Control Strategies with Human Factors for Active Steering

A Thesis Submitted to The University of Nottingham for the Degree of Doctor of Philosophy

JOSEPH S. HELLEWELL, MEng

December 2019

Abstract

Autonomous road vehicles are an increasingly popular research topic due to their potential to reduce the number of road deaths and road accidents while increasing energy efficiency of the vehicle and reducing congestion. To ensure a successful implementation of driverless cars, it is necessary that the vehicle's behaviour is acceptable to its occupants; that they feel as safe and comfortable as reasonably possible. What the resulting behaviour and driving style is like is different for different people and so it is necessary that an autonomous vehicle can change its performance accordingly.

This thesis presents an investigation into how existing human factor models can be incorporated into a control framework. This hierarchical Model Predictive Control (MPC) framework is evaluated in simulation for a constant speed, Active Steering Control (ASC) implementation. This aspect of autonomous behaviour is demonstrated as it represents a typical scenario, motorway driving, and it is illustrative of the effect that human factors can have on controller design and implementation. Published driving study data is used to create different percentile drivers, allowing for the evaluation of different driving styles. The resulting behaviour is analysed and evaluated quantifiably so that the effect that human preferences have on system performance can be understood.

Vehicle models are implemented in software and validated against established results. These models are used as both the system plant in the MPC formulation and as simulation models of the vehicle itself. It is demonstrated that the lower fidelity models have good agreement with more advanced, and complex, vehicle models. This allows for a simpler design process, as the vehicle-controller behaviour can be simulated much more

quickly.

A standard linear MPC formulation is then presented and its performance is demonstrated for the active steering control task. The ability to effectively constrain the vehicles lateral acceleration is shown, along with MPC's flexibility, by demonstrating All-Wheel Steering (AWS) performance. The necessity of providing suitable reference trajectories is demonstrated, which leads to the development of a hierarchical control framework, which creates reference trajectories for the ASC to track.

This developed framework is then used to implement human factor aware driving behaviour. Human factor parameters from the literature are included directly into the MPC formulation.

It is shown that incorporating human factor models of lateral acceleration preference, solely as constraints in the MPC steering module, does not provide a large differentiation in driving styles. This is because changes only occur at edge cases in driving scenarios and the human factor derived constraint is 'active'. Incorporating the models into the MPC cost function, in conjunction with constraints, increases the differences in performance by allowing for a trade-off between reference tracking of system states and lateral acceleration.

To allow for this trade-off for a range of longitudinal velocities and human acceleration tolerance, a penalty model is presented that determines the relative influence of the lateral acceleration penalty in the cost function. Although the method is attractive due to its ease of implementation, it is shown that its effectiveness is limited in scenarios where other concerns dominate, such as respecting road bound limits. Nevertheless, the models presented in this thesis, along with the results and conclusions obtained from them, provide a demonstration of how human factor issues can be included explicitly into a control framework for autonomous driving.

Acknowledgements

I would like to thank Professor Atanas Popov, for his support and guidance throughout this project. The time and knowledge that you shared has shaped this thesis. A thank you also to Professor Gary Burnett, your discussions on the far-reaching implications of human factors in the engineering problem provided the foundation on which the research was built.

Thanks, are also due to the Engineering and Physical Sciences Research Council (EPSRC), who provided funding for this research.

I also want to thank my family for their constant love and support that they have always shown. My final, but greatest, thanks and gratitude must go to Jess, whose love and support I am grateful for every day. Without it none of this would be possible.

Contents

Abstract i

Acknowledgements iii

Contents iv

Abbreviations vii

List of Figures ix

List of Tables xiii

List of Publications xv

1 Introduction 1

2 Literature Review 5

2.1 Introduction 5

2.2 History and Recent Developments 6

2.2.1 Autonomous Driving Behaviour 6

2.2.2 Human Factor Studies of Ride Comfort and Driving Behaviour 13

2.2.3 Control Methods in Automotive Systems 24

2.2.4 Vehicle Models 29

2.3 Thesis Aims and Objectives 31

2.4 Thesis Outline 32

3 Vehicle Modelling and Model Verification 35

3.1 Introduction 35

3.2 Model Verification and Validation 36

3.3 Vehicle Models 37

3.3.1 Tyre Models 37

3.3.2 The Bicycle Model 43

3.3.3 The Nonlinear Bicycle Model 48

3.3.4 Multi-Degree-of-Freedom Model 51

3.4 Displacement Tracking 63
3.5 Model Discretisation 67
3.5.1 Zero Order Hold 67
3.5.2 Euler Approximation 68
3.6 The Effect of Control Frequency on System Response 69
3.7 Conclusions 72
4 Model Predictive Control for Active Steering Control 75
4.1 Introduction 75
4.2 Formulation 77
4.2.1 Stability 80
4.2.2 Demonstrating the Control Performance 82
4.3 Conclusions 103
5 Hierarchical Model Predictive Control 105
5.1 Introduction 105
5.2 Hierarchical Structure 108
5.3 Geometric Path Generation 109
5.3.1 HMPC Coordinate Frame 110
5.3.2 Formulating the Geometric Path Generation Problem as a Quadratic Programming Problem 110
5.4 Path Generation and Tracking for a Double Lane Change Manoeuvre 114
5.5 Conclusions 131
6 Incorporating Human Factors 132
6.1 Introduction 132
6.2 Including Human Factors in the Control Problem 134
6.3 Human Preference of Vehicle Handling Behaviour 136
6.4 Results 139
6.4.1 Creating Reference Courses for Control Evaluation . 140
6.4.2 Evaluating Control Performance via RMS Error Tracking 140
6.4.3 Preliminary Results 141
6.4.4 Tracking the Minimum Acceleration Reference 145

6.5 Penalising Acceleration in the MPC Cost Function 149
6.5.1 Choosing the Acceleration Penalty Weight Value . . . 151
6.5.2 Results 155
6.6 Conclusions 162
7 Conclusions and Future Work **167**
7.1 Original Contributions to Knowledge 168
7.2 Future Work 172
Bibliography **176**
Appendix A Multi-Degree-Of-Freedom Model **191**
A.1 Equations of Motion 191
A.2 Model Verification 198
Appendix B Smooth Lateral Reference Trajectory **203**
Appendix C Linear MPC Formulated in QP form **205**

Abbreviations

ABS	Anti-lock Braking System
ACC	Adaptive Cruise Control
ASC	Active Steering Control
AWS	All-Wheel Steering
CAS	Collision Avoidance System
CAV	Connected and Automated Vehicle
DAE	Differential Algebraic Equation
DARE	Discrete Algebraic Riccati Equation
DARPA	Defense Advanced Research Projects Agency
DDT	Dynamic Driving Task
DOF	Degree-of-Freedom
ECOM	Extended Control Model
ESC	Electronic Stability Control
ESP	Electronic Stability Program
FOH	First Order Hold
HMPC	Hierarchical Model Predictive Control
IoT	Internet of Things
IPM	Interior-Point Methods

LQR	Linear Quadratic Regulator
LTV	Linear-Time Varying
MPC	Model Predictive Control
MTF	Magic Tyre Formula
ODE	Ordinary Differential Equation
PATH	Program on Advanced Technology for the Highway
PID	Proportional-Integral-Derivative
PROMETHEUS	PROgraM for European Traffic with Highest Efficiency and Unprecedented Safety
QP	Quadratic Programming
RMS	Root Mean Square
SDP	SemiDefinite Programming
SMMT	Society of Motor Manufacturers and Traders
SQP	Sequential Quadratic Programming
SUV	Sports Utility Vehicle
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
ZOH	Zero Order Hold

List of Figures

2.1 The mean lateral acceleration chosen by the drivers, in the Ritchie study 16
2.2 The Herrin exponential fitting, which modelled how a driver modifies their behaviour when driving a well-known route. . 17
2.3 Median, 95th and 99th percentile values from the modified-Levison criterion 19
2.4 Michon's mental model of the driving task 22

3.1 Concept of the tyre brush model. 39
3.2 The Magic Tyre Formula creates a highly-nonlinear force response 42
3.3 The Magic Tyre Formula handles combined slip conditions, where both longitudinal slip and slip angle are non-zero . . . 43
3.4 The linear bicycle model 44
3.5 System response of the bicycle model to a 0.75° step input at 80 km/h 47
3.6 System response of the nonlinear bicycle model to a 0.75° step input at 80 km/h (22.2 m/s) 50
3.7 System response of the Multi-Degree-of-Freedom (DOF) model, with tyre transience, to a 0.75° step input at 80 km/h (22.2 m/s) 57
3.8 System response of the Multi-Degree of Freedom model, with tyre transience, to a 0.75° step input at 80 km/h 58
3.9 Comparison of the lateral velocity, yaw rate and lateral acceleration during a sine-dwell manoeuvre at 80 km/h. 62
3.10 Linear biycle model with yaw and lateral displacement tracking 64

3.11 Lateral displacement of the bicycle model at 20 m/s with a 1° step-input 66
3.12 The lateral acceleration response of the two state bicycle model when subjected to a sine-dwell steering input applied at discrete time-instances 71

4.1 Vehicle state response to tracking a step change in lateral displacement 84
4.2 Control input and lateral acceleration of the vehicle with unconstrained MPC formulation 85
4.3 Vehicle state response to tracking a step change in lateral displacement with acceleration constraint 88
4.4 The linear lateral acceleration constraint successfully constrains the simulation vehicle's lateral acceleration 89
4.5 Vehicle state response with terminal penalty weight $\mathbf{Q}_f$ and preview horizon $N = 10$ 91
4.6 AWS vehicle respose to tracking a step function. 94
4.7 The linear lateral acceleration constraint successfully constrains the vehicle's lateral acceleration 95
4.8 Vehicle state response to tracking a smooth step function . . 98
4.9 Control sequence and corresponding lateral acceleration when tracking a smooth step reference. 99
4.10 Demonstration on the effect of both explicit stabilising techniques and reference trajectory on system stability 102

5.1 The double lane change manoeuvre 116
5.2 Demonstrating the effect of the Geometric path generation module 119
5.3 Vehicle response to tracking the reference path created by the geometric path generation module 122
5.4 Steering input and corresponding simulation vehicle acceleration when tracking static reference 123

5.5 The sequence of reference trajectories generated by the path generation module. In (b) the upper and lower road bounds, as functions of time, are shown as the bold lines, along with the optimisation problem bounds, which include an additional margin of 'safety' of 1 metre, shown as the dotted lines 126
5.6 Vehicle response to tracking the dynamic reference path created by the geometric path generation module 127
5.7 Control input and lateral acceleration response when tracking dynamic path planning trajectory 128

6.1 Median, 95th and 99.9th percentile values representing an upper bound on lateral acceleration 138
6.2 Reference tracking at 65 km/h, for the controller equiped with median, 95th and 99.9th driver profiles. The median profile can not track the reference trajectory as closely due its lower acceleration bounds 143
6.3 Lateral aceleration of the simulation vehicle with the the three acceleration profiles 144
6.4 State values for the minimum acceleration reference trajectory, shown for the median (50th), 95th and 99.9th percentile. 147
6.5 Comparison of the median (50th), 95th and 99.9th acceleration profiles for the minimum acceleration profile curves . . 148
6.6 Norm of the acceleration term added to the optimisation problem's quadratic term 152
6.7 Using u^3 to derive α restores the quadratic relationship between lateral acceleration and longitudinal velocity 154
6.8 Norm of quadratic matrix term for common road speeds . . 155
6.9 Reference tracking for the acceleration penalty profiles . . . 157
6.10 Lateral acceleration for the four penalty profiles 158
6.11 State variables for vehicle controlled with MPC controller with cost function 160

6.12 Lateral acceleration and saturation of the HMPC controller with acceleration penalty 161

A.1 Inertial acceleration of the vehicle body 199
A.2 Cardan angles of the Multi-DOF vehicle model 200
A.3 Body-fixed velocity components of the Multi-DOF model's unsprung mass 201
A.4 Angular acceleration of the wheels 202

List of Tables

2.1 Summary of levels of autonomy 7

3.1 Magic Tyre Formula (MTF) parameter values used in the thesis 41
3.2 Parameter values for the linear bicycle model. 45
3.3 Parameter values for the Multi-Degree of freedom model, from [106]. 53
3.4 Gain values for the driveline torque to provide horizontal speed control. 56
3.5 Eigenvalues of the bicycle model state matrix 69

4.1 Parameter values for the linear bicycle model. 83
4.2 MPC parameter values. 83

5.1 Course dimensions for the double lane change manoeuvre . . 116
5.2 Preview horzion value 117
5.3 Weighting values for path velocity and acceleration use in path generation problem 118
5.4 The effect of path planning call ratio on tracking performance. The simulations are performed at 65 km/h 129

6.1 Acceleration bound parameters using the modified Levision model 139
6.2 Dynamic RMS tracking error for controller with constant longitudindal velocity of 65 km/h 142
6.3 Static RMS tracking error for controller with constant longitudindal velocity of 65 km/h 142

6.4 Summary of acceleration data for the profile run 143
6.5 Dynamic RMS tracking error for controller with minimum acceleration reference path for front-wheel steering vehicle with constant longitudinal velocity of 65 km/h 145
6.6 Static RMS tracking error for controller with minimum acceleration reference path for front-wheel steering vehicle with constant longitudinal velocity of 65 km/h 146
6.7 Summary of acceleration data for the three profiles when tracking the minimum acceleration reference at 65 km/h . . 146
6.8 Summary of weighting values and saturation values 157
6.9 RMS dynamic tracking error for controller with acceleration penalty 159
6.10 RMS static tracking error for controller with acceleration penalty 162
6.11 Summary of acceleration data for the profile run at 65 km/h 162

List of Publications

The author's previous publications are listed in chronological order below.

[1] J. S. Hellewell, A. A. Popov, G. E. Burnett, and J. W. Griffin. "Model predictive control for active drive-torque distribution during a slalom manouevre". In: *Advanced Vehicle Control AVEC - Proceedings of the 13th International Symposium on Advanced Vehicle Control*. Ed. by P. Pfeffer, E. J., and M. Plochl. Munich; Germany: CRC Press/Balkema, 2016, pp. 611–618. ISBN: 9781315265285

[2] J. S. Hellewell, A. A. Popov, and G. E. Burnett. "Hierarchical model predictive control with human factor constraints". In: *The Dynamics of Vehicles on Roads and Tracks*. Ed. by M. Spiryagin, T. J. Gordon, C. Cole, and T. McSweeney. Vol. 1. Rockhampton; Australia: CRC Press/Balkema, 2018, pp. 289–294. ISBN: 9781138035713

[3] J. S. Hellewell, A. A. Popov, and G. E. Burnett. "Hierarchical Control for Trajectory Generation and Tracking Via Active Front Steering". In: *Journal of Dynamic Systems, Measurement, and Control* 142.1 (2020), pp. 1–10. ISSN: 0022-0434. DOI: 10.1115/1.4044620

Chapter 1

Introduction

Road vehicles equipped with autonomous driving capability have been a much discussed, and a much researched, idea for several decades. They offer an opportunity to revolutionise human transportation with their anticipated safety benefits saving lives, and their ability to maximise energy efficiency needed more than ever in a world affected by climate change induced by human activity. Either issue alone would warrant the pursuit of driverless cars, and the effect on the UK's economy is forecast to be in excess of £60 billion by the year 2030 [1].

For any successful implementation of driverless vehicles, the autonomous behaviour must be acceptable to its human occupants, who are now, by definition, exclusively passengers rather than drivers. Therefore, the occupants must feel as safe and ideally as comfortable as they do currently. A driverless vehicle must therefore choose its behaviour to not just satisfy the physical limitations of the system, but the expectations of its passengers. This task is made all the more difficult, as there is no single driving style or behaviour that would be satisfactory to all passengers. Consequently, the autonomous vehicle must be able to adapt its driving style depending on human preference. From a design perspective, the control methodology implementing the autonomous behaviour must be flexible and robust to changes in the parameters that modify the driving behaviour. This allows for the necessary changes in vehicle behaviour without the need to design specific controller behaviour for different human preferences or driving

styles.

Automation of road vehicle behaviour began in the 1980s [2], although the focus was on improving the capabilities of individual subsystems, such as the development of the Anti-lock Braking System (ABS) and later Electronic Stability Program (ESP) (also known as Electronic Stability Control (ESC)). While at first such systems may not be thought of as autonomous in the common sense, they represent an important first step as the vehicle determines a control input, such as brake pressure to individual wheels, based on both the driver input and additional sensor measurements. Such was the effectiveness of these active safety systems, so called because they attempt to prevent rather than provide protection in the event of a collision, that they are mandatory on all cars manufactured for EU, US and other markets since the early 2010s [3, 4]. It has been reported that the introduction of ESP reduced single-vehicle accidents by 30-50% for cars and by 50-70% for Sports Utility Vehicles (SUVs) [5, 6]. These were later augmented by the development of more advanced techniques such as Adaptive Cruise Control (ACC) or Collision Avoidance System (CAS). CAS operates autonomously by applying the brakes if it detects that the driver has not taken action quickly enough to prevent a collision with an obstacle in front of the vehicle, detected using radar or lidar sensors along with cameras using object recognition techniques.

When considering how such active safety systems have been developed, the approach of traditional automotive manufacturers can be contrasted to those that have entered via a computational or software background. There, the attempt has been to develop an autonomous functionality that requires little intervention from the driver, and with frequent updates to control software provided by over-the-air updates [7–9]. The differences in research philosophy can perhaps be explained by their respective backgrounds; established automotive manufacturers wish to update their existing vehicle products incrementally, only introducing features when they are technologically mature and there is a clear economic case for doing so. In contrast, new competitors wish to disrupt the existing market by introduc-

ing transformative products, mitigating the lack of a strong manufacturing and supply chain experience. One goal is to develop an autonomous driving system that can be integrated into existing vehicle platforms, allowing the system developer to focus on their strengths of computer vision, object recognition and computer hardware [10, 11].

The development of both control theory and computer hardware have allowed for the implementation of advanced control technologies, in a manner that would not have been achievable in previous decades. This progression allows for the specification and regulation of advanced control laws, necessary for the complex behaviour required to implement driverless vehicles. The availability of cheap, powerful, computer hardware that is compact and energy efficient enough to be installed within vehicles, have made it possible to implement advanced driving technologies for the first time.

In trying to implement human aware driving behaviour in a control framework, it makes sense to study how human drivers change their behaviour in different driving conditions and how this behaviour can be characterised. By understanding how humans change their driving style, it is possible to implement autonomous behaviour that mimics specific behaviour. This human factor aware control framework would then be able to provide autonomous driving behaviour, characterised identified models parameterised by both system parameters such as vehicle speed, and those which identify human preference that can vary among the population. Even if the specific, numerical, parameter values differ for human passengers, rather than human drivers, a successful control implementation would be able to adapt to moderate changes in said values.

For the evaluation of control strategies for autonomous driving behaviour, simulation models offer a method of rapidly prototyping and assessing performance. In using models with a lower complexity compared to physical prototypes or even advanced multi-body models, it is possible to give a more general judgement on the control performance rather than for a single vehicle design. For this reason, such models are popular eval-

uation tools within the vehicle dynamics research community and form a solid basis for the design and verification of autonomous control behaviour.

In this thesis, control methods that can incorporate human driving preferences, explicitly in their design, are investigated. One such method, Model Predictive Control (MPC), is then used to develop an active steering controller that uses human factor models of acceleration to modify its driving behaviour. This controller is then evaluated, using vehicle simulation models that are implemented and verified, by navigating a standard evaluation course. The results are evaluated, and the changes in driving behaviour is quantified, allowing for the inclusion of human factors into the control problem to be objectively analysed.

Chapter 2

Literature Review

2.1 Introduction

To allow for the study of autonomous driving behaviour and considerations of how human factor preferences can be accommodated within a control strategy, this chapter presents the history, along with the current state-of-the-art technology, of autonomous functionality in the automotive research field. Alongside this, human factor models are reviewed to determine how they might be incorporated into a control framework.

The developments in control theory, particularly the emergence of Model Predictive Control (MPC), that allow for the implementation of the proposed control design are presented in this chapter. An emphasis on matters concerning the closed-loop stability of the system, imperative for any successful control specification, are discussed and the subsequent implications on the implementation of human factor issues are raised.

So that the proposed control methodology can be assessed effectively, potential vehicle models and modelling methods are reviewed.

This literature review is then used to determine the aims and objectives upon which this thesis is built. From these objectives, a thesis outline is presented with a short summary of its chapters.

2.2 History and Recent Developments

The advances made in autonomous behaviour are summarised and work concerning human driving performance is reviewed. Control techniques for autonomous driving behaviour implementation and vehicle models available for use as simulation models are discussed.

2.2.1 Autonomous Driving Behaviour

Before examples of autonomous driving research and implantations are discussed, it is beneficial to define what it meant by autonomous behaviour. In order to provide a consistent categorisation of autonomous capability, the professional organisation and standards body SAE International defined six levels of autonomy, summarised in Table 2.1 [12]. This includes level 0, in which no assistance is provided to the human driver. For levels 2 and below, the human driver is responsible for monitoring the driving environment, and being ready to immediately react to any changes in the Dynamic Driving Task (DDT). With the aforementioned legislation for road vehicles, all modern cars have at least level 1 autonomy, with many, particularly premium vehicles, offering level 2 features such as Adaptive Cruise Control (ACC) and Collision Avoidance System (CAS).

In 2017, the first production vehicles with level 3 autonomy were introduced [8, 13]. In a level 3 vehicle, it is responsible for controlling the vehicle and monitoring the environment in specific driving conditions. The driver must remain alert and be aware of their surroundings incase of the vehicle encountering a situation it cannot deal with. For a level 4 vehicle, human interaction is not necessary because the vehicle can always bring itself to a stop safely, even if it cannot proceed with the driving task. The final stage of full automation, level 5, requires that the vehicle does not need any human intervention for any driving scenario, and it is always able to determine suitable decision, regardless of its environment. While there are claims of level 4 or 5 vehicles being ready for the mass market in the near

Table 2.1: Summary of levels of autonomy, adapted from [12]

SAE level	Name	Definition	Fallback of Driving Task	System Capability
Human driver monitors the driving environment and performs part or all of the Dynamic Driving Task (DDT)				
0	No Automation	The full-time performance by the human driver of all aspects of the dynamic driving task, even when enhanced by warning or intervention systems.	Human Driver	N/A
1	Driver Assistance	The driving mode-specific execution by a driver assistance system of either steering or acceleration/deceleration using information about the driving environment and with the expectation that the human driver perform all remaining aspects of the dynamic driving task.	Human Driver	Some driving modes
2	Partial Automation	The driving mode-specific execution by one or more driver assistance systems of both steering and acceleration/deceleration using information about the driving environment and with the expectation that the human driver perform all remaining aspects of the dynamic driving task	Human Driver	Some driving modes
Automated driving system monitors the driving environment and performs the DDT while engaged				
3	Conditional Automation	The driving mode-specific performance by an automated driving system of all aspects of the dynamic driving task with the expectation that the human driver will respond appropriately to a request to intervene.	Human Driver	Some driving modes
4	High Automation	The driving mode-specific performance by an automated driving system of all aspects of the dynamic driving task, even if a human driver does not respond appropriately to a request to intervene.	System	Some driving modes
5	Full Automation	The full time performance by an automated driving system of all aspects of the dynamic driving task under all roadway and environment conditions that can be managed by a human driver.	System	All driving modes

future [14, 15], this belief is not held by all experts [16, 17]. In a recent interview of industry experts in the UK, the accountancy firm KPMG, on behalf of The Society of Motor Manufacturers and Traders (SMMT), reported that level 4 functionality was not anticipated until 2025, with level 5 vehicles being introduced sometime after 2030 [1]. Either way, it is clear that there will be a large research effort required for the foreseeable future, before the autonomous driving environment is fully matured.

While the study of how human drivers react to handover events, that is, when a driver must resume driving in a scenario that a level 3 autonomous vehicle cannot handle itself ,is worth of study in itself such as [18], it is not necessary for the successful study of incorporating human factors directly into the control problem. For this reason, in this thesis it is assumed that in the driving scenarios considered, the vehicle and controller is always capable of performing the driving task and so could be considered to be a level 4 or 5 vehicle.

The term 'driverless cars' is used colloquially when discussing autonomous driving. With reference to Table 2.1, driverless cars correspond to either a level 4 or 5 vehicle, i.e., vehicles that can complete a given driving scenario without a driver, and without a human monitoring the environment. In this thesis, the proposed autonomous driving behaviour will take place without the need for human intervention, and so the phrase 'driverless cars' may be used in reference to the simulation vehicle equipped with the human factor aware controller.

Autonomous Driving Trials

Along with the many theoretical advances in autonomous vehicle functionality and behaviour, there have been a number of technology demonstrators that implement some level of autonomous driving or functionality. Some of the more prominent examples are listed below in an approximate chronological order.

The PATH Program

The Program on Advanced Technology for the Highway (PATH) [19] was a Californian collaboration between researchers in the Institute of Transportation Studies at UC Berkeley and the California Department of Transportation. The project began in 1986, and had a number of objectives with the aim of reducing the traffic congestion issues that the Californian freeway network was suffering in the 1980s [20].

One notable aspect of PATH was its experimental vehicles studies, with a demonstration of automatic vehicle control, which used permanent magnets embedded in the roadway to detect vehicle position from measurements of the magnetic field. While this is an important demonstration of autonomous behaviour, the necessary investment in updating infrastructure renders this technique infeasible for large scale deployment.

More modern methods of road positioning utilise existing infrastructure, e.g. detection of lane markings with object recognition and computer vision such as [21], to minimise the cost of large scale deployment.

PROMETHEUS

The PROgraM for European Traffic with Highest Efficiency and Unprecedented Safety (PROMETHEUS) [22] was an industry-led program initiated in 1987. Like the PATH program, it had many objectives focussed around improving drivers situational awareness, developing active safety systems, cooperative driving and traffic management. The program aimed to place an emphasis on an integrated approach with its sub-programs claiming that earlier research had suffered due to the lack of consideration of overall system design.

PROMETHEUS resulted in the 'VaMoRs–P vehicle, developed by the Bundeswehr University Munich. A twin was developed, VITA_2, by Daimler-Benz [23]. The work was based on even earlier work involving a 5 ton van 'VaMoRs' equipped with computer vision for lane-keeping autonomous driving [24]. They demonstrated autonomous behaviour in typical motorway driving in France, Germany and Denmark in 1994 [25]. Equipped with

forward and rear facing bifocal cameras, it had a detection range of up to 100 m in both directions. It was reported that the vehicles were capable of high-speed (130 km/h) driving, and were able to change lane, if permission was given by the supervising driver. A 1600 km trip from Germany to Denmark was undertaken, in which it was reported that 95% of the distance could be driven automatically. Dickmanns et al. summarised that, thanks to advances in computer hardware and machine vision, real world problems were now being solved in real-time [25].

From a modern perspective, the technology demonstrators that were part of PROMETHEUS were remarkably advanced for their time. It is perhaps an indication of the complexity and unforeseen implementation issues in bringing, what is a level 3 autonomous feature, from development to market.

DARPA Grand Challenge

In an attempt to accelerate the progress into the development of autonomous vehicles, an American competition was launched by Defense Advanced Research Projects Agency (DARPA) in 2004. The goal was to develop autonomous vehicles that could navigate off-road dirt tracks at high speed, with a $1 million prize offered as an incentive. As none of the vehicles that entered were able to finish the course set, the competition was held again in 2005, this time with an increased prize of $2 million. This time, the competition was won by 'Stanley', a vehicle developed by a team based at Stanford University, who used their experience of the competition in the previous year to improve the robustness of their autonomous vehicle [26]. The researchers noted that as the route was provided to them before the start of the race, the task was primarily high-speed road finding along with obstacle detection and avoidance in desert terrain.

DARPA Urban Challenge

Following the success of the Grand Challenge in 2005, DARPA organised the Urban challenge, which was held in 2007. The vehicles had to navigate a closed urban environment without human interaction; obeying Californian

driving laws, with no longer than ten seconds of time to stop and consider a given environment. The competition was won by a team from Carnegie Mellon University. Their vehicle 'Boss' used model predictive trajectory generation to dynamically create feasible paths to track from a given starting point [27]. It was able to detect a number of driving contexts, such as encountering intersections, merge planning or in-lane driving, and identifying appropriate driving behaviour for a given context. Boss was also equipped with a number of error recovery modes, to account for specific errors in the motion planning or unaccounted for anomalous scenarios.

Mercedes-Benz INTELLIGENT DRIVE

In 2013, Mercedes-Benz completed a 100 km journey autonomously, across country, driving on rural roads through villages and through major cities [28]. The route was chosen to mark the 125th anniversary of the worlds first long-distance journey taken by Bertha Benz, wife of Carl Benz.

The Mercedes-Benz S-class INTELLIGENT DRIVE prototype vehicle (Bertha Benz) successfully completed the journey in an autonomous fashion. The vehicle exhibited level 4 autonomy, as by design there were situations it could not enter: it was not allowed to enter the opposite road lane by more than one metre. Therefore, when the vehicle encountered a parked car in its lane, it came to a complete stop. The supervising driver then had to navigate the parked car before the autonomous driving mode could restart. The 2013 S-class used the standard production sensors along with an additional six radar sensors, two wide-angle cameras and a modified stereo vision camera. In [28], the authors emphasise that, compared to successful competitors in the earlier DARPA events, the S-class made extensive use of computer-vision and radar for vehicle perception. The earlier vehicles had relied on lidar, which is a much more expensive system. Because of this, the INTELLIGENT DRIVE project represents a step towards a more cost-effective mass producible system.

For this thesis, what is of more interest than the perception aspects of autonomous driving, is the motion planning and trajectory control of

the vehicle. In the Bertha Benz vehicle, the top level of motion planning was behaviour generation. This module handled mapping perceived objects into geometric constraints. Trajectory planning is then performed by solving a nonlinear geometric problem via Sequential Quadratic Programming (SQP). One potential drawback of this method is that if the nonlinear optimisation problem is non-convex, then SQP has to solve a sequence of quadratic approximations iteratively, which makes it struggle to meet the real-time requirements needed for control [29]. Many nonlinear trajectory problems are non-convex, due to the nature of the obstacle constraints.

Once a trajectory (which penalised deviation from the centre of the feasible road space, speed deviation from the current road limit and excessive acceleration and jerk) was determined, it was passed to separate lateral and longitudinal control modules. Longitudinal control was handled by an MPC controller [30], while for lateral control a mixture of feedforward control is used for disturbance compensation, with proportional-based feedback used for lateral displacement stabilisation [28].

In summarising the performance of the vehicle, the authors in [28] noted that, by design, it had a defensive driving style in that it took preventative action to avoid entering potentially dangerous driving scenarios. It was also reported that some passengers reported the driving to be akin to a human learner driver taking lessons and that overall the vehicle behaviour was far inferior to an attentive human driver. Nevertheless, the route is an impressive demonstration of the current capabilities of state-of-the-art driverless cars.

Autonomous Driving Trial Conclusions

In these autonomous driving trials, a common finding from the researchers is the effort required and emphasis on a robust testing framework for the software implementation, and a software design that is ready to handle unexpected hardware or sensor failures [26, 27]. This highlights that while the simulations of autonomous controller behaviour and resulting vehicle

response are a useful tool for researchers, there are implementation issues that remain and therefore physical prototypes and testing will be required. This is echoed by Guanetti et al. in their review [31], who noted a lack of experimental validation for many proposed driving technologies. This was attributed to difficulties on testing on public roads, but the increased funding costs of developing prototype vehicles must also be a factor in the lack of experimental validation of proposed technology.

Connected and Automated Vehicles

Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I), which are examples of Connected and Automated Vehicle (CAV) implementations, are an exciting research avenue. Using upcoming technology, such as 5G wireless communication devices and Internet of Things (IoT) protocols, autonomous vehicles further improve their performance by sharing information about their current vehicle state and driving intent. In their review article, Guanetti et al. [31] highlight that CAVs have the potential to improve autonomous driving far beyond the implementation of driverless features on individual cars alone.

By sharing sensor data, and broadcasting their own intent, it is possible for CAVs to improve environmental awareness far beyond that capable than a single autonomous vehicle, or more importantly, a single human driver. These benefits only begin to be realised when there is a large uptake of CAVs so that they represent a significant proportion of vehicles on the road network. While CAVs have an important role in realising the full potential of driverless cars, they are a separate development concern to the primary research question of this thesis, and are so not considered in their own right for the remainder of this work.

2.2.2 Human Factor Studies of Ride Comfort and Driving Behaviour

In the study of human driving performance, the perception of vehicle behaviour and subsequent response in driver control action are critical when

evaluating the ride comfort and handling performance of the vehicle. Ride comfort regards the physical sensation and perception of a passenger in a moving vehicle. It is a subjective issue, as it regards the classification of vehicle ride as comfort or discomfort and is reliant on numerous complex and interdependent factors [32].

These factors are dependent on both environmental sources, such as the smoothness of the road surface and the vehicle itself, such as the quality of the suspension system and its ability to supress vibration induced by road surfaces. Vehicle vibration is one of the most well studied aspects of ride comfort, particularly with regards to the vehicles response. This is so that suspension systems, whether they be passive, active or semi-active, have optimal parameter values to minimise the effect of road-induced vibration.

In these papers, it is common to define a ride performance index and then perform a parameter study or optimisation routine to determine optimal suspension parameters. Simple vehicle models are used as they can provide insight into fundamental design trade-offs [33, 34]. For example, two common ride indexes are the root mean square body vertical acceleration and the root mean square dynamic tyre load parameter [33, 35]. The dynamic tyre load index is a measure of the road-holding capability of the vehicle and is predominantly a vehicle safety issue; the design task is to ensure that there is always a large normal force acting upon each tyre to maximise the lateral and longitudinal response. The body vertical acceleration is more relevant when considering ride comfort as it can be perceived by humans exposed to it; at higher frequencies this vertical acceleration is referred to as vehicle vibration.

Ride comfort measured as exposure to vibration of time, has been extensively studied in both academic and industrial environments. ISO 2631 [36] and BS 6841 [37] are two standard methods used to evaluate whole body exposure to vibration. These use a weighted frequency response, to account for differing human sensitivities to specific frequencies and direction of vibration, to determine a scalar value as a ride comfort index. These can then be compared to thresholds that have been derived from subjective

evaluation to determine if the vehicle is comfortable.

Along with exposure to whole body vibration, other dynamic factors, such as vehicle pitch and roll oscillation determine a persons perception as to whether a vehicle is comfortable or not. For example, Wyllie and Griffin [38] showed that people perceive low frequency (0.2 - 1.6 Hz) transient lateral acceleration and roll oscillation differently, and that while both components contribute to the inertial lateral acceleration of the vehicle, roll oscillation results in greater discomfort than pure lateral acceleration of the same magnitude.

Additionally jerk, being the derivative of acceleration with respect to time, is an important factor of ride comfort and indeed is viewed as a critical aspect of ride comfort when evaluation vehicle vibration [39, 40]. Studies have attempted to quantify jerk thresholds for discomfort [41], and have found values ranging from 0.5 - 0.9 m/s^3 [42, 43]. While these studies focussed on seated rail passengers it is reasonable to assume that similar thresholds exist in road vehicles, but additional testing would be welcome in this area as it is less well studied than the effect of acceleration. The reason for the relative paucity of jerk studies compared to acceleration studies when evaluating ride comfort is that it cannot be measured directly. Instead it must be approximated via the numerical differentiation of a noisy acceleration signal, which results in large errors in the jerk calculation, even for small noise values [44].

There is clearly a multitude of personal, environmental and vehicle factors that determine the perceived comfort and safety of a vehicle. For a full implementation of human factor aware autonomous vehicle behaviour, it would be beneficial to consider all the above aspects in a control framework. Such an implementation is beyond the scope of this thesis and so instead a focus is placed upon the lateral acceleration of the vehicle. One reason for this is that it has a larger body of published literature from which appropriate human factor models can be identified. Another critical factor is that, unlike roll comfort, changing the lateral acceleration behaviour of the vehicle changes its ability to track a reference path. Therefore, when

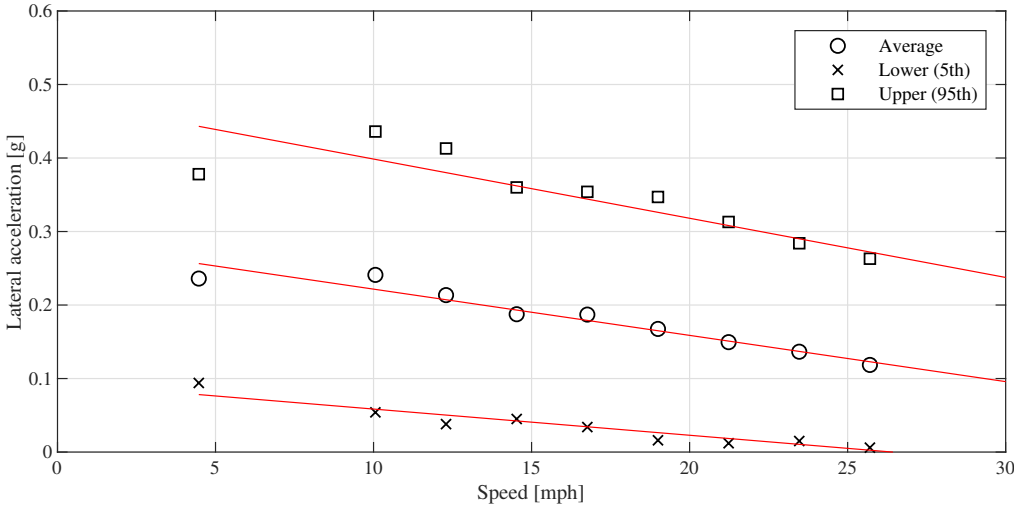


Figure 2.1: The mean lateral acceleration chosen by the drivers, in the Ritchie study [46]. The mean value, along with the 5^{th} and 95^{th} percentiles.

evaluating the effectiveness of incorporating human factor constraints directly into a control scheme, it is vital to characterise the consequent path tracking ability which is fundamental to a autonomous vehicles ability to safely navigate its surroundings.

Human behaviour in curve driving

The human factor discipline concerning drivers' driving style and tolerance to lateral acceleration, is referred to as curve driving. This may arise from the initial studies identifying the relationship between road curvature and longitudinal velocity. The first paper to study driver lateral performance when negotiating curves appears to be by Taragin and Leisch [45], before a more thorough investigation by Ritchie et al. [46, 47].

In [46], 50 drivers (25 men, 25 women) were asked to drive the same 110 miles of Ohio highway, with 227 curves. They were asked to drive 'normally' and their speed and lateral acceleration was recorded around each of the curves. They found a linear relationship between the vehicle's longitudinal velocity and lateral acceleration, with the drivers controlling the vehicle with a lower cornering acceleration when travelling at higher speeds.

In this experiment, no measure was made of the road curvature (presum-

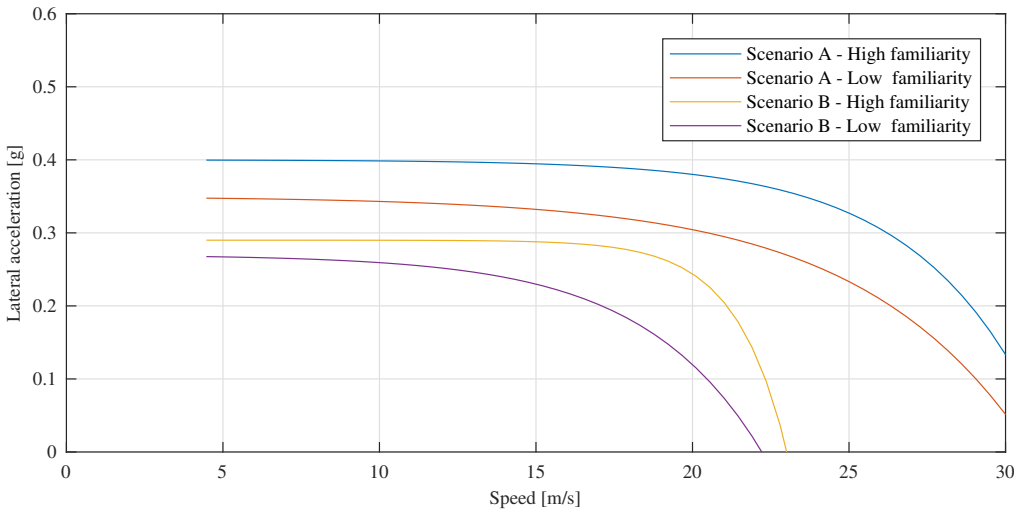


Figure 2.2: The Herrin exponential fitting [48], which modelled how a driver modifies their behaviour when driving a well-known route.

ably due to the technology available at the time) so instead this relationship between acceleration and speed is presented with no information on how 'tight' the curves are. There is the presumption that at higher speeds the curvature was lower, but there is no empirical link between curvature and a driver's acceleration tolerance.

The benefit of this simple empirical model is that it can give simple upper and lower bounds on the lateral acceleration based on the forward longitudinal speed. Therefore, the vehicle's lateral handling behaviour is adaptive to its longitudinal speed and the results provide an, albeit crude, set of parameter values based upon the mean chosen lateral acceleration and the standard deviation values that are provided. The reason that the proposed bounds might be viewed as crude is that the magnitude of the standard deviation values is large with respect to the mean value. Nevertheless, they do provide a large range of parameter values to investigate how these change driving behaviours.

This work by Ritchie was expanded upon by Herrin [48], which proposed an empirical model for a driver's negotiating 'horizontal curves'. In their paper, it is proposed that a desired longitudinal speed is an indication of expediency while lateral acceleration is a comfort or safety parameter. It is therefore theorised that a driver is continually balancing a desire to arrive

to their destination as quickly as possible and a desire to do so comfortably.

Naturally, this optimal compromise will vary from driver to driver, and also depending on a given situation. In [48], the 20 drivers were asked to drive one of two routes twice, once imagining that they were on a leisurely Sunday drive and in the other imagining that they were late for an important appointment.

The proposed trade-off model between acceleration and longitudinal velocity is:

$$a^{'} = e^{\beta u^{'}},$$
$$a^{'} = 1 - \frac{a}{a_{\mathrm{T}}}, \qquad u^{'} = u_{\mathrm{A}} - u,$$

therefore, $a^{'}$ is the fraction of remain allowable acceleration upto the tolerance accleration, a_{T}. $u^{'}$ is clearly the deviation from the aspiration velocity u_{A} and β is what Herrin refers to as the reflective degree of driver experience, which is constant for a given driving scenarion. This value has the following effect on a driver's behaviour:

- $\beta \ll 0$, the driver shows an unwillingess to slow down
- $\beta < 0$, a tradeoff between longitudinal speed and lateral acceleration
- $\beta \approx 0$, the behaviour follows the linear relation proposed by Ritchie in [46].

An appealing aspect of this model is that it extends the model proposed by Ritchie, while still providing the same behaviour in the appropriate regieme. Also, in its attempt to model driver expedience, it offers the potential to change a vehicle's behaviour depending on the autonomous vehicle's driver, but also on the scenario that said driver finds themselves in. It is to be expected that a passenger in an autonomous car would be grateful that the vehicle adapts it behaviour should it be aware that one is late for a meeting, for example. Much later, work by Bosetti proposed a much extended model of a human's control of a vehicle for lateral handling [49, 50]. A lot of the model focussed on existing proposed model for a

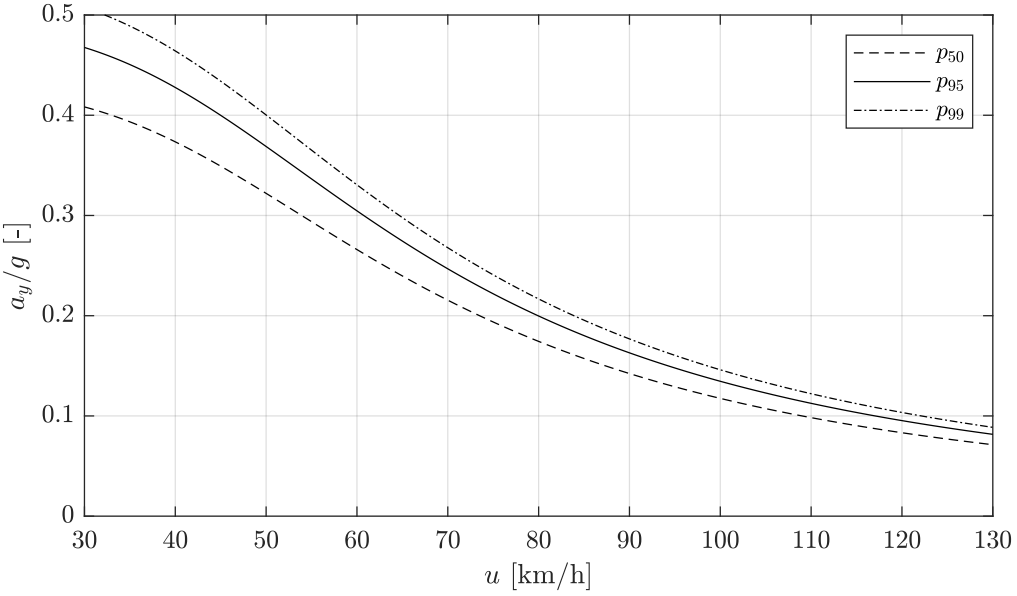


Figure 2.3: Median, 95th and 99th percentile values from the modified-Levison criterion

human's sensory-motor control, such as the so-called two-thirds power law, which relates a human's hand movements (velocity) to the traced path curvature, such a handwriting. It was proposed to use this to model the drivers steering inputs when driving, in which a larger steering input, and hence hand-speed, is needed to navigate a tighter bend. The two-thirds power law derives its name from the following experimentally observed correlation:

$$\omega = \alpha\kappa^{-2/3} \tag{2.2}$$

where ω is the rotational velocity of the hand, κ is the road curvature and α is a constant found for a given task. It is hypothesised that this law explains humans desire to control their movement via the minimum jerk principle, i.e. minimising changes in acceleration, to achieve a trade-off in motor control accuracy and speed depending on the task at hand. In [49], this law is used as a guiding principle on identifying a driver's trade-off between acceleration and velocity; as the human's motor control is the inner control loop that modulates the instantaneous tolerance to acceleration.

Reymond, et al [51] proposed a model of maximum speed choice and lateral acceleration behaviour of driver based upon dynamic safety margin. In

this, they assumed that the driver steers the vehicle to ensure that the lateral deviation from the path centreline was kept within safety limits. This results in a prediction of a quadratic decrease in peak lateral acceleration with increasing vehicle speed. Driving simulator results and experimental results from a test track suggests that both visual and acceleration cues play a role in driver behaviour.

Another model is the so-called modified-Levision criterion [49], in which they propose the following model for a driver's choice of lateral acceleration:

$$a_y = \frac{a_0}{\sqrt{\left(1 - \frac{v}{v_0}\right)^2 + 2\frac{v}{v_0}^2}} \tag{2.3}$$

which utilises two parameters, v_0 and a_0. By modifying a_0, different driver profiles can be considered, see Figure 2.3. The original research by Levison, et. al. [52, 53], was motivated by attempting to develop a combined human-vehicle model, for use as an analysis aid in projects sponsored by the US Department of Transport's Federal Highway Administration. In this research which focussed on the driving of 'rural highways' without interaction between other road vehicles, model data was obtained by two driving studies. One study was undertaken by Battelle Human Factors Transportation Center in Washington State and the other by Virginia Tech Transportation Institute.

It must be noted that while the work by Bosetti, using real-world driving data from the interactIVe project [54], supported the Levison criteria during non-motorway driving conditions, they noted that Eq. (2.3) did not hold for motorway driving. On the motorway, it was found that the acceleration did not reduce at higher speeds. Instead, the acceleration was bounded above by the induced centripetal acceleration for navigating a curve of a given radius at a given speed. They theorised that due to the wider road lanes on the motorway, the speed-acceleration trade-off is shifted to the former. This indicates that human drivers have a number of driving modes and change their driving style for specific driving scenarios. Therefore, it is reasonable to assume that autonomous driving behaviour would operate a

number of discrete driving modes. Of course, then the challenge would be to develop a robust 'overseer' to ensure that the relevant operating mode is selected for a given driving scenario.

Nevertheless, the modified-Levison criteria and corresponding driver percentile data in [50] identifies itself as a suitable, up-to-date model to use in the study of incorporating human factor models into control frameworks. It is a simple model, characterizable with few parameters required and so can provide insight for a generic vehicle implementation, yet the driving study it is derived from over 2500 kilometres of driver-mile data [49].

Driver Steering Control

One key area of research that incorporates human factor issues with control issues in the driving task is the development of models, both theoretical and computational, of a human driver [55, 56]. One of the identified benefits of accurate driver models used in steering control, for example, is that it allows for vehicle manufacturers to evaluate objectively how design changes modify the closed-loop driver-vehicle system [57].

In their 1985 review [58], Michon summarised the existing models of human driving behaviour, and what still needed to be developed in the future. In the review, Michon strongly stresses their belief that future human factor research should be based on cognitive methods, instead of behavioural observations. The (then recently) developing trend of cognitive research in behavioural sciences, particularly psychology, was based on a desire to understand the internal brain processes in a given situation, for example, during driving tasks. Michon's belief was based on the hierarchical cognitive control structure of human behaviour. Referencing their own earlier work [59, 60], and that of Janssen [61], Michon highlights the now well-known three-tier hierarchical model of human behaviour during the driving task. In the proposed framework, it is theorised that the human driver splits a driving task into three discrete levels: strategic, tactical and control, as shown in Figure 2.4.

The strategic level comprises of the high level, general, plans made for a

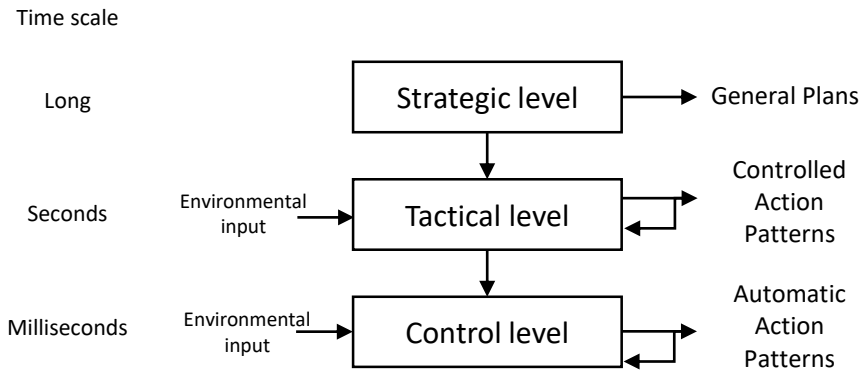


Figure 2.4: Michon's mental model of the driving task. Adapted from [58], it highlights how elements of the driving task are delegated and the time scales that they occur over

given journey, for example deciding on the route to take before the driving task actually begins. The tactical level focusses on shorter timescales, typically on the order of a few seconds, and deals with decisions that have to be made due to the driver's immediate environment. Examples of this type of behaviour are navigating parked vehicles in urban environments or overtaking a slower moving vehicle on the motorway.

The lowest stage in this hierarchical framework is the control level, which in this context refers to the human motor control of the vehicle inputs. This behaviour deals with how the driver achieves the tactical and strategic goals at a given instant, responding to physical feedback from the steering wheel, pedals and vibrations to ensure that the vehicle remains dynamically stable.

More recently, the Extended Control Model (ECOM) model was proposed [62]. In this, the number of hierarchical layers was increased to four levels, with Hollnagel splitting the cognitive task into targeting, monitoring, regulating and tracking.

Targeting directly corresponds to the strategic level, in that it is an open-loop process which is actively thought of over a relatively long-time scale. Monitoring may be thought of as covering high-level tactical processes, supervising external information such as road signs or traffic lights and updating the driving objectives from this information.

Regulating can also be thought of as similar to some aspects covered in the tactical level of Michons model. In the ECOM, regulating involves the determination of driving modes, prioritising certain driving objectives over others on short time-scales, to facilitate overtaking other road vehicles for example. This level of the model is also responsible to directing the lowest-level tracking loops and is itself a closed-loop process. Tracking itself refers to activities that, for the skilled drivers, occurs subconsciously, without active thought. These tasks, which are inherently part of closed-loop feedback action, are concerned with implementing the physical interaction with the system, to perform the immediate regulatory task.

In both the Michon and ECOM frameworks it is clear that, for the human driver, a hierarchical approach to driving is required. This separation of concerns allows for decisions to be addressed in the strategic or tactical spaces at the frequency at which they occur. The subconscious low-level control is then allowed to function as required. From the perspective of implementing autonomous driving behaviour, a hierarchical control framework also has a number of benefits; the modularity of individual sub-controllers allow for a simpler design and analysis process focussing on specific aims such as path generation or reference tracking. Discrete modules also allow for the use of individual controllers, implemented on separate hardware, to operate at different operating frequencies. This allows for a trade-off between operating frequency and control complexity to be made on a case-by-case basis.

The simulation theory of cognition [63] states that humans learn mental models of the plant they are to control (including their own body). This leads to the obvious motivation of using MPC to develop accurate driver models. MPC, which will be discussed in the next section, has been used by a number of authors to successfully model driving behaviour.

To surpass the current ability of human drivers, designers of autonomous driving systems must take full advantage of the capabilities of modern hardware, sensors and control theory. It is of course beneficial to understand how the driving task is perceived, processed and control implemented by

human drivers in the closed-loop human-driver system. Doing so gives an insight into how the human occupant responds to the dynamic behaviour of the vehicle.

While it is not necessary for a successful implementation of a human aware controller to develop an autonomous vehicle that mimics a human driver, it is beneficial to understand how a human driver might respond in a given situation.

2.2.3 Control Methods in Automotive Systems

In order to implement autonomous driving behaviour, it is necessary to make use of a suitable control strategy. Kalman's work on optimal control in the early 1960s introduced state-space models as a method of analysing stability in the time domain [64], in contrast to earlier 'classical' control techniques analysed in the frequency domain [65, 66], such as Proportional-Integral-Derivative (PID).

The Linear Quadratic Regulator (LQR) method [67] has been applied to the driver steering control task [68, 69], and other vehicle dynamic applications [70].

Similar methods of optimisation based control were developed, including Model Predictive Control (MPC), which has many theoretical and practical benefits that make it attractive as a control strategy for autonomous driving behaviour. This is the focus of the next section.

Model Predictive Control

Model Predictive Control (MPC), also known as Receding Horizon Control and historically other names such as Dynamic Matrix Control, has been around for more than three decades; it had a wide usage in industrial settings before academic rigour had been applied to its theory [71]. Initially it was used in the chemical processing industry, where due to the fact that the time scales of the processes were on the order of hours, there was time for the limited hardware of the day to calculate the appropriate control action, which was then implemented manually by the plant operators [72].

In the 1980s, the introduction of affordable microprocessors allowed for the introduction of constrained optimisation into the MPC algorithms. With MPC, a model of the system plant is used to predict the system behaviour in the near future. At each discrete time-step of the controller, a control sequence is computed by solving an optimal control problem, and the first control in this sequence applied to the system plant The attractiveness of MPC is that it can handle hard constraints on the system state inputs and outputs, as well as the control action.

Automotive Applications of MPC

Model Predictive Control has been investigated for many different automotive applications, such as engine and emission control, steering, suspensions and mechatronic actuators, among others [73]. Some work was done initially in the mid-1990s [74], before a flurry of work was published in the late 2000s. Advances in the theory of stability and robustness of MPC along with work being published on fast MPC methods [75], coupled with the ever advancing speed and availability of microprocessors made sub-second execution times possible for the controllers, and so the time-scales of the system dynamics could finally begin to be handled by MPC.

In [76], Cole et al. compared the performance of MPC to infinite and finite horizon LQR control strategies for the task of modelling the human driver completing the steering control task. They noted that for a long-enough (finite) preview horizon, the performance of LQR and MPC was similar. Due to the lack of additional stability methods, the MPC method they used was not stable for very short preview horizons. As this work was attempting to accurately model driver behaviour, rather than implement a robust controller, this is not necessarily an issue. As the MPC formulation does not utilise control or state constraints it is not a practical option for the general steering task, but it represents an early identification of the applicability of MPC for the active steering task.

In 2005, Borreli et al. [77] presented a nonlinear MPC [78] formulation for Active Steering Control (ASC), tracking smooth reference trajectories

for the lateral displacement and yaw angle, to navigate a double lane change manoeuvre in low grip (icy) conditions. The work is a demonstration of the suitability of MPC for control tasks, as it handles the nonlinear dynamics of the vehicle and tyre models. One of the key findings was that a larger preview horizon was required for closed-loop stability at high speeds. This introduces the complication that at higher speeds, when there is less time to react to developing situation, it is necessary to solve a more complex optimisation problem. For the nonlinear controller they presented, they expected, using 2005 hardware, to be able to experimentally validate the controller only at speeds of up to 10 m/s (36 km/h). The computational complexity was a limiting factor of control capability, but the paper demonstrates that MPC was becoming a viable control policy for the real-time control problem of active steering.

In 2007, the same authors presented experimental results of the active steering controller [79]. Due to the aforementioned complexity issues, the vehicle could only be tested at 7 m/s (25 km/h), due to instability occurring at 10 m/s (36 km/h).

To mitigate the control complexity, they then presented a Linear-Time Varying (LTV)-MPC formulation, based upon successive online linearization of the nonlinear model around the current vehicle state and control input. Because of this, the experimental vehicle could complete the double lane change manoeuvre at 21 m/s (75 km/h). This work is significant in that it details early experimental results for the use of MPC for steering control. It also highlights that while using nonlinear MPC is desirable as it allows for a more accurate modelling of the vehicle dynamics, there are implementation issues due to its higher model complexity that need to be addressed.

This framework was extended in the 2008 conference paper by the same authors, Falcone et al. [80]. They utilise the previous nonlinear ASC as a low-level subsystem of a hierarchical framework. In addition, a smooth reference trajectory is calculated offline, which is supplied to an MPC-based trajectory re-planner. This re-planner uses a point-mass model for

the vehicle dynamics and nonlinear constraints, to modify the initial path based on the vehicles current state. Due to the lower model complexity, the upper level module is able to optimise the vehicles future reference trajectory, over the next 2-3 seconds, beyond the horizon of the lower-level module. This allows for a shorter preview horizon in the low-level ASC module, while maintaining stability. The authors then demonstrated that a simulation model was able to navigate the double lane change manoeuvre when controlled by the HMPC framework, at higher speeds.

There are numerous other examples of active steering control, or lane change manoeuvring, such as Yoshida, et al. [81] who compared the performance of linear (bicycle) and four-wheel vehicle models. As the subject has become more developed, MPC for active steering has been incorporated with torque vectoring and path planning for collision avoidance [82], and as an assistance aid for human drivers [83]. Liu et al. considered stochastic constraints for an LTV-MPC model which imposed a lower bound on the probability of the vehicle remaining in its lane.

In [84], the authors present a two-tier hierarchical framework, for path planning with speed control and active steering in motorway scenarios. Due to the one-way nature of the traffic conditions, they were able to express the problem as a convex optimisation problem, which allows for much faster solution times. The effect of convexity is discussed in the later literature review section on computational effort. The results show that utilising low-complexity models, it is possible to successfully perform path planning and active steering control.

These papers, among others, show that MPC is a suitable control technique for the design, analysis and implementation of autonomous driving behaviour; particularly active steering control. The ability to specify constraints on control inputs and the dynamic behaviour, along with specifying optimal behaviour in the cost function, make it a clear choice for studying how human factor models of lateral acceleration can be incorporated into an autonomous driving control formulation. Indeed, Boyd recognises the utility of using optimisation in control, stating that it is "the PID control

of the 21st century" [85].

MPC Stability

Work on proving the stability of MPC schemes began during the late 1980s [71]. It was known that the infinite-horizon MPC law was closed loop stabilisable, so long as the initial state was stabilisable. Calculating the infinite horizon control law is not feasible in general however, as it requires up to an infinite number of dimensions in the optimisation problem at each time step. Methods were required therefore that maintain this stability with a finite, and practical preview horizon.

An effective stability method is the inclusion of a terminal state penalty matrix, which is appllied in the cost function to the state vector at the preview horizon [71, 86]. Additional stability methods, such as ellipsoidal or maximal output admissible terminal state constraint sets are not considered [71, 72, 87, 88]. Robust MPC practices [89–92] are likewise not considered, as the focus of the thesis does not require these additional techniques.

Computational effort

An area that was historically seen as a barrier for the practical implementation of control techniques such as MPC was the computational cost of the underlying optimisation problem. A number of advances in both MPC formulation and optimisation theory allow for the practical application of MPC when high operating frequency is required.

It was Rockafellar who noted that the distinction was between convexity and non-convexity, rather than linear and non-linear, when determining the relative difficulty of solving optimisation problems [93]. The work of Nesterov and Nemirovski [94], building on works by people like Karamakar [95], provided a detailed explanation of how Interior-Point Methodss (IPMs) could provide polynomial-time algorithms could solve convex optimisation problems. A polynomial time algorithm has computational complexity expressed, in big O notation, as $\mathcal{O}(n^c)$, for a problem of input size n, with a constant c being the leading power dependant on the problem and al-

gorithm. Polynomial-time algorithms are desirable as the solution time required does not increase unmanageably for larger optimisation problems.

Therefore, it is also desirable that MPC problems are formulated as convex optimisation problems. For linear systems, with linear state and control constraints, this is the case as they are expressed as Quadratic Programming (QP) problems. In fact, a wide range of MPC variations can be modelled as convex problems, specifically via SemiDefinite Programming (SDP) [96].

Along with improvements to the underlying theoretical framework, improvements were made to how MPC problems could be formulated. For linear systems, it is possible to remove the system states from the optimisation (as demonstrated in [76], for example), by implicit application of the state space model. This approach is referred to as *condensed form* [97].

In [75], Wang and Boyd outline a method of exploiting the structure of a linear MPC problem, that can outperform the condensed form at large enough preview horizons. In addition, in this full *regulatory form* [97], it is possible to enforce constraints on the system states, instead of just control inputs, and so is preferred for this thesis. In addition, using several methods such as *warm-starting*, a fixed barrier parameter in the IPM algorithm and a linear relationship between preview horizon and computational complexity, it is possible to solve MPC problems in milliseconds, showing that MPC is clearly a viable option for real-time control applications.

2.2.4 Vehicle Models

Analytic research into the dynamic response and stability of automobiles began in the 1940s, before a number of papers were published in the USA through the 1950s [98]. The researchers had a background in aerospace research and stressed the importance of utilising the expertise already available from this work [99]. Early research focussed on a linear two Degree-of-Freedom (DOF) model, known as the bicycle model, which is still popular to this day. The bicycle model assumed that there is no load transfer between the inner and outer wheels of the vehicle, and that the steer angle of

both wheels on an axle is the same. These assumptions allow the wheels to be doubled-up and so the front and rear tyre characteristics have been lumped together respectively. Although the model does not account for the roll motion of the vehicle, it still allows for an understanding of the cars dynamic behaviour [100]. The bicycle model assumes a constant longitudinal velocity and calculates the yaw acceleration and the lateral acceleration response to control inputs to the wheels' steering angle.

Although the bicycle model is simple, it allows for important stability analysis and accurately represents linear handling behaviour, which typically allows for manoeuvres up to 0.3 g of lateral acceleration [101]. This insight into vehicle handling behaviour makes it an appropriate starting point for validating more complex models, comparing their linear handling behaviour before the more complex models can be used for nonlinear analysis. The linear bicycle model can easily be extended to include longitudinal acceleration and nonlinear tyre models. This increases the degrees-of-freedom of the model, as the rotational accelerations of the wheel are required unless an external source can provide longitudinal slip values, such as a traction controller [102].

A slightly more complex model used for vehicle dynamics is the 7-DOF model which has seen use particularly for torque vectoring control research [103, 104]. This extends the bicycle model, reintroducing track width and modelling the rotational acceleration of each wheel. The roll, pitch and heave movement of the vehicles unsprung mass is not modelled however, and the suspension is omitted. Instead the normal force acting on each tyre is modelled as the static weight distribution plus the effect if the longitudinal and lateral inertial accelerations. These forces can additionally be taken as the steady state normal forces and a first order lag function can be used to approximate the suspension dynamics [105]. A 7-DOF models steady-state response matches those of a multi-body system [103], and so is a useful compromise between model fidelity and ease of verification of the simulation environment. Shim and Ghike developed an accurate 14-DOF model that incorporated vehicle suspension and vertical movement dynam-

ics [106]. They used the model to study the wheel 'lift-off' phenomenon, without having to implement complex multi-body models.

Tyre Models

Tyre modelling focusses on attempting to model the complex behaviour of the pneumatic tyre-road interaction. Despite their outwardly simple appearance, the tyre construction of 'rubber' compound, reinforcing steel belts and thermodynamic processes of the pressurised air within comprises some of the most complex dynamic behaviour of a vehicle's mechanical systems. The seminal work in developing practical tyre models for a wide range of applications was done by Pacejka, who proposed the semi-empirical Magic Tyre Formula (MTF) model [98, 107].

The model gained wide acceptance by the users in the automotive industry and academics in the vehicle dynamics research community alike. This was due to its hybrid design (hence being referred to as semi-empirical), which derived best-fit parameters from experimental data, but used fundamental physical parameters such as wheel slip and slip angle, as model inputs.

2.3 Thesis Aims and Objectives

The aim of this thesis is to develop methods and tools based upon vehicle control models which incorporate human driving preferences explicitly in the control formulation. To achieve this, the following objectives are specified:

1. Research control techniques that are available for implementing autonomous driving technology.
2. Research and identify examples of human factor models that are suitable to modify control performance.
3. Create and verify high-fidelity models of both the vehicle plant and

control implementations, which give an accurate and thorough representation of vehicle dynamics.

4. Investigate and understand the effect of human factor derived control laws on the handling of road vehicles.
5. Critically evaluate whether human factor derived control behaviour is beneficial and viable for implementation in the immediate or near future.
6. Identify further areas of research, based upon the findings of this thesis, which would expand the scope of the results.

In order to bound the scope of the thesis, the analysis will be performed at constant speed and focussing on a single simulation scenario, the double lane change manoeuvre, which is representative of overtaking on a motorway.

2.4 Thesis Outline

The thesis is split into seven chapters to organise it into the key areas that underly the work. They are summarised as follows:

Chapter 1: Introduction

The motivation for the research is presented along with a summary of research and results presented for a general audience.

Chapter 2: Literature Review

Existing literature on the key issues, including control techniques and human factor research is examined so that the state-of-the-art can be determined. From this, the key findings are used to define a basis for the thesis objectives. A more general history of control methodologies and vehicle dynamics research is also discussed to establish the background upon which the thesis is built upon.

Chapter 3: Vehicle Modelling and Model Verification

Vehicle models are a key component of the methodology in thesis. Models that provide an accurate representation of the vehicle behaviour and response to control input are presented and evaluated to ensure that the correct level of model fidelity is used to ensure that appropriate meaning can be taken from simulated results.

Model verification and validation is used, based upon results from dynamic theory to ensure that higher complexity models are correctly implemented and behave as expected in linear, low acceleration, conditions. Numerical considerations that affect the discretisation of the models from the fundamental continuous-time formulation are discussed, along with methods to mitigate implementation issues that arise.

Chapter 4: Model Predictive Control for Active Steering Control

The literature review identifies Model Predictive Control (MPC) as an appropriate control technique to fulfil the thesis objectives. Along with an introduction to the principles behind its operation, a standard formulation for linear models and constraints is presented, and stability issues are discussed along with typical methods that are available to overcome such issues.

MPC is then used to demonstrate its utility with an ASC example. MPC parameters and environmental factors that affect the closed loop system response are explicitly demonstrated in simulation to emphasise its utility as an effective control technique for complex situations such as autonomous driving scenarios.

Chapter 5: Hierarchical Model Predictive Control

Limitations in MPC, identified in Chapter 4, are addressed via the extension to a Hierarchical Model Predictive Control (HMPC) framework. The HMPC framework allows for dynamic reference trajectory planning, ensuring that MPC controllers have a feasible reference to track. The benefits of HMPC are demonstrated in simulation and key parameters are evaluated

to ensure that the proposed set-up is an appropriate method to overcome the limitations of MPC.

Chapter 6: Incorporating Human Factors

Human Factor models developed by experimental observation are reported and their applicability in control law design is discussed. Using the vehicle models developed in Chapter 3, and the HMPC framework presented in Chapters 4 & 5, Human Factor (HF) derived behaviour is evaluated in simulation by including HF behaviour explicitly in the control problem formulation.

Chapter 7: Conclusions and Future Work

The thesis objectives are recalled, and the major research findings are presented. From this, future research directions are discussed.

In the thesis, the following original contributions to science will be introduced:

- A demonstration through simulation that bicycle models are effective simulation tools for high-speed manoeuvres commonly used to evaluate motorway driving scenarios.
- An effective, linear, approximation of lateral acceleration is incorporated into an MPC scheme, allowing for the transient lateral acceleration of the vehicle to be accurately constrained.
- Human Factor derived acceleration models incorporated directly into an MPC scheme, allowing for vehicle behaviour to be changed by using existing models of human behaviour.
- A detailed study of the impact of human factor derived constraints on the dynamic tracking ability of a vehicle navigating a double lane change manoeuvre.
- A novel penalty function is developed to incorporate a dynamic penalty on lateral acceleration, to enhance effectiveness when using human factor models over a range of vehicle speeds.

Chapter 3

Vehicle Modelling and Model Verification

3.1 Introduction

To evaluate the performance of any controller acting upon a vehicle system, the closed-loop system of vehicle and controller is often simulated in an appropriate numerical environment, at least as a first stage of the design process, before prototypes are used to further investigate the control behaviour. In academic environments, it is often enough to evaluate a new controller in a simulated environment, as it allows for the studying of an archetypal controller, which exhibits the key behaviours of the system, without having to focus attention on application specific design and implementation issues.

To perform the numerical simulation, a mathematical representation of the vehicle is needed. Depending on the nature of the vehicle behaviour being studied, and the desired insight into vehicle response, there are many different formulations available. In this chapter, tyre models are first presented, as they are fundamental to the dynamic response of any automotive system. Then, linear and nonlinear vehicle models are detailed, and their implementation verified. Models are first verified against expected steady state values in the linear handling regime. The nonlinear, higher fidelity,

models are then used to observed system behaviour in more severe handling regimes. From this, the capability of the simpler vehicle models in these conditions can also be evaluated. These comparative tests are undertaken with a standardised sine-dwell manoeuvre, which is particularly effective at exciting the vehicles lateral response [108].

Next, displacement tracking is discussed, particularly the effect of coordinate linearisation, necessary for linear Model Predictive Control (MPC) implementation, which is the focus of the next chapter.

Two final aspects related to control implementation are then discussed, model discretisation and the effect of a discrete control input sequence on vehicle response. Two methods of discretisation are considered; Zero Order Hold (ZOH) and a first order Euler approximation, also known as First Order Hold (FOH).

3.2 Model Verification and Validation

The model verification used in this thesis is based on a sequence of verification steps; the lowest complexity models are first verified in the linear handling regime in which the assumptions made about their handling are valid. This allows for analytical verification of the steady state behaviour.

This provides a foundation on which more complex behaviour can be referenced against. The models used in this thesis are implementations of established work so that the focus is on the proposed control framework and behaviour. The use of established models means that the primary focus of the model verification is to ensure the absence of human errors in the software implementation.

Because of this, the higher order models are validated by observing the steady state behaviour and comparing it to the lower order models. The transient response is compared to published models and the physical behaviour we would expect form classical first and second order dynamic systems.

3.3 Vehicle Models

There is a wide variety of vehicle models available from modern literature. Many are specialised to try and understand specific dynamic behaviour, such as the quarter-car model to characterise a vehicle's ride quality [35, 109, 110] or as is the interest of this thesis, determining the lateral response of the vehicle.

Additionally a model of the tyre is required, to characterise the interaction between the vehicle and the road surface. Two key tyre models are outlined below.

3.3.1 Tyre Models

Despite a pneumatic tyres simple assembly, with no moving components or mechanisms, the vehicles tyre exhibits some of the most complex dynamic behaviour of the vehicle system. This is due to the composite construction of its rubber compound with steel reinforcement belts, along with the dynamics of the compressed air within. A further consideration is the interaction of the tyre with the ground itself, which can have friction characteristics that change quickly, affecting the tyres force generation capabilities.

Because of this, many models have been developed, to allow for the study of force generation within the tyre. The models vary in their complexity from simple linear models of the lateral and longitudinal tyre forces, to empirical nonlinear models obtained from physical experiments, that allow for the studying of the coupled lateral-longitudinal action. Complex multiphysics models, which use detailed Finite Element Models (FEM) are now feasible with modern computing hardware and software tools [98].

These models have an implementation complexity, greater than required for the successful completion of the thesis objectives, and so are not considered further. In the following, some key tyre characteristics are defined, and the two tyre models that are used in this thesis are introduced.

Tyre Slip

The longitudinal slip of a wheel, κ, is a measured of the relative longitudinal movement between the wheel hub centre and the tyre contact patch. There are a number of definitions for longitudinal slip, depending on the running conditions of the tyre, such as if the wheel is being accelerated or braked. In this thesis, it defined as follows:

$$\kappa = \frac{r\omega - u_{\mathrm{t}}}{|u_{\mathrm{t}}|}, \tag{3.1}$$

which accommodates both the acceleration and deacceleration conditions [98, Chapter 4]. ω is the rotational velocity of the wheel, r is the radius of the tyre, taken to be constant and u_{t} is the longitudinal velocity of the tyre at the wheel hub. It is often expressed as a percentage value.

Slip Angle

For both the brush model and the MTF, the lateral force is parameterised by the slip angle of the tyre. This is a measure of the deformation in the tyre, and is calculated from the ratio of the lateral to longitudinal velocity in the tyre:

$$\tan\alpha = \frac{v_{\mathrm{t}}}{u_{\mathrm{t}}} \tag{3.2}$$

where α is the slip angle of the tyre. The subscript t is to denote that the longitudinal velocity u and the lateral velocity v refers to velocity of an individual tyre itself, and not that of the vehicle centre of mass. This can change depending on the vehicle model, and so the specific formula will be specified when discussing a particular vehicle model.

The Brush Model

The simplest tyre model assumes that the tyre is a disc with bristles attached around its rim, hence the name [111, Chapter 13]. When an individual bristle comes into contact with the ground, it deforms linearly like a spring. The model assumes that the bristles tip stays in place on the contact patch until it is lifted from the ground due to the tyres rotation.

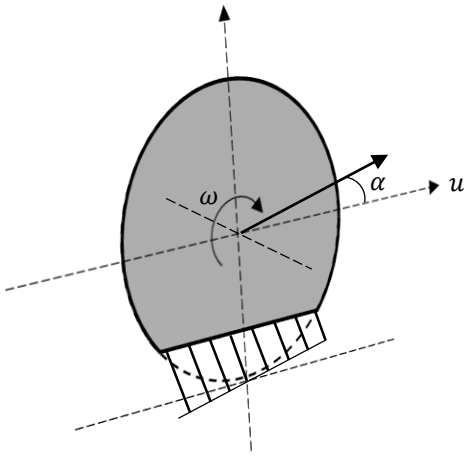


Figure 3.1: Concept of the tyre brush model.

This results in the following linear lateral tyre force, as a function of the tyres slip angle:

$$F_y = -C_\alpha \alpha, \tag{3.3}$$

where the negative sign indicates that it is a restoring force, resisting the lateral movement in the contact patch. The brush model is unbounded, in that its response does not saturate at large slip angles. For this reason, it is only valid for slip angles up to around 5 or 10 degrees. For this reason, it is only suitable when modelling low acceleration manoeuvres. Despite this apparent limitation, it provides an adequate approximation of tyre force response in many typical driving manoeuvres, particularly at higher vehicle speeds, and has the benefit of being easily incorporated into linear vehicle models.

The Magic Tyre Formula

The Magic Tyre Formula (MTF) [98, 107], is the most widely used non-linear tyre model in both academic and industry environments. It is a semi-empirical model in that model coefficients are fitted to slip terms determined from physical experiments.

As can be seen from Figure 3.2, which shows both the longitudinal and lateral tyre force in pure slip conditions, the MTF generates a tyre response that quickly peaks before tending towards a constant, lower limit. This occurs when the tyre saturates and is unable to provide further tractive or

cornering response.

The general form of the MTF used in this thesis is:

$$\begin{aligned} Y(\boldsymbol{x}) &= \mu F_z \cdot G \sin\left(C \tan^{-1}\left(Bx_1 - E\left(Bx - \tan^{-1}(Bx)\right)\right)\right) \\ G &= \cos\left(C_\mathrm{c} \tan^{-1}\left(B_\mathrm{c} x_2 - E_\mathrm{c}\left(B_\mathrm{c} x_2 - \tan^{-1}(B_\mathrm{c} x_2)\right)\right)\right) \end{aligned} \tag{3.4}$$

where μ is the road-tyre coefficient of friction, F_z is the normal tyre force acting upon the road. The MTF is designed so that it can be used for both longitudinal and lateral force calculations. As such, the term x_1 refers to the primary variable, while the x_2 refers to the secondary variable. For longitudinal tyre force calculation, κ is the primary variable and α is the secondary. For the lateral tyre force, these variables are reversed.

The G term determines the effect that the combined slip conditions have on the force generation. As can be seen in Figure 3.3, combined slip conditions reduce the maximum tyre force in the primary direction. In the example shown, when the longitudinal slip is held at 10%, the maximum lateral tyre force is only approximately 60% of the pure side-slip condition.

Using the MTF, it is possible to determine lateral tyre stiffness values that are used in the linear bicycle model. The slope of the MTF at the zero-slip condition provides the basis of the stiffness value:

$$C_\alpha = -2 \left.\frac{\partial F_\mathrm{c}}{\partial \alpha}\right|_{\alpha=0} \tag{3.5}$$

The factor of two is due to that, for the bicycle model, the tyre stiffness is doubled up to represent both wheels on the axle. The negative sign is to accommodate to the difference in sign convention.

In Fig 3.2 the MTF response is demonstrated, using the parameter values shown in Table 3.1. Fig 3.2 (a), it can be seen that the peak longitudinal tyre force is achieved as approximately ±10% (the tyre response is symmetric about the y-axis), before reducing to a value of 75% of the maximum as the slip value reaches 100%.

In Fig 3.2 (b), the maximum lateral force occurs at just before 10°, reducing to 90% of the maximum, as the slip angle reaches 45°.

Table 3.1: MTF parameter values used in the thesis

Parameter	Value
B_x	13.7
C_x	1.6
E_x	0.1
$B_{x,\mathrm{c}}$	13.0
$C_{x,\mathrm{c}}$	1.0
$E_{x,\mathrm{c}}$	-0.4
B_y	-11.5
C_y	1.3
E_y	-0.8
$B_{y,\mathrm{c}}$	10.6
$C_{y,\mathrm{c}}$	1.1
$E_{y,\mathrm{c}}$	0.3

For both the longitudinal and lateral tyre response, the point at which the maximum occurs and how quickly the tyre force falls away after saturation is dependent on an individual tyres design and composition, and the resulting fit of the MTF parameters.

Tyre Transience

The magic tyre formula models the steady state running conditions of the tyre, and a common assumption is that this value is achieved quickly enough that that simulation models assume that the steady state values are achieved instantaneously [98, Chapter 5]. When considering vehicle control scenarios as in this thesis, in which the steer angle can change rapidly, this assumption may no longer be valid [98, Chapter 8].

Then, a simple method to model the changing value of the slip angle is to assume that the slip is derived from a first-order system [98, Chapter 5]. Then, the slip angle is governed by the following Ordinary Differential Equation (ODE):

$$\dot{\alpha} = \frac{u}{\sigma} \left(\alpha_{ss} - \alpha \right), \tag{3.6}$$

where α is now the 'apparent' slip angle to be used in the tyre model and α_{ss} is the steady-state value, given by Eq. (3.2). σ is the relaxation length of the tyre, which can be approximated by the tyres radius [98], as is done in this thesis.

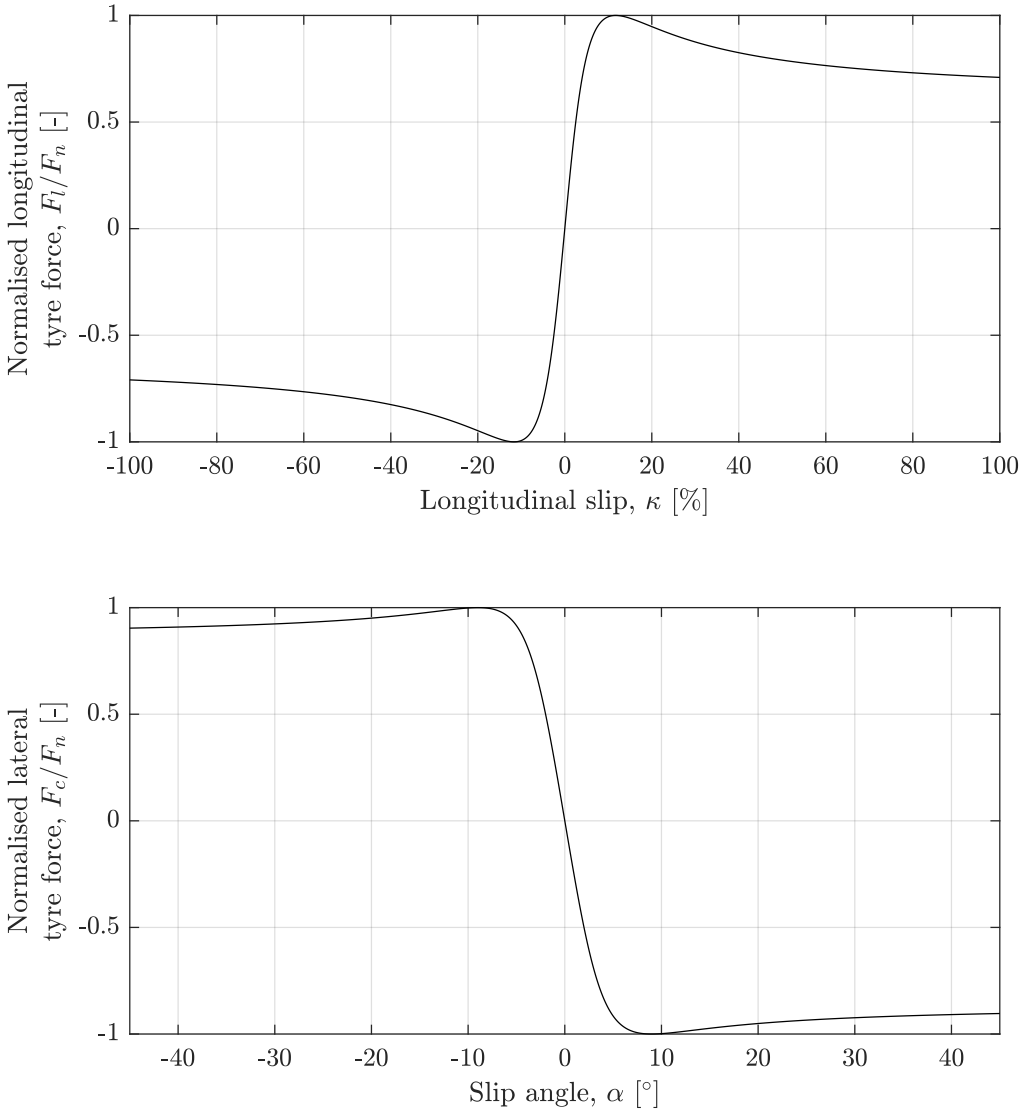


Figure 3.2: The Magic Tyre Formula creates a highly-nonlinear force response

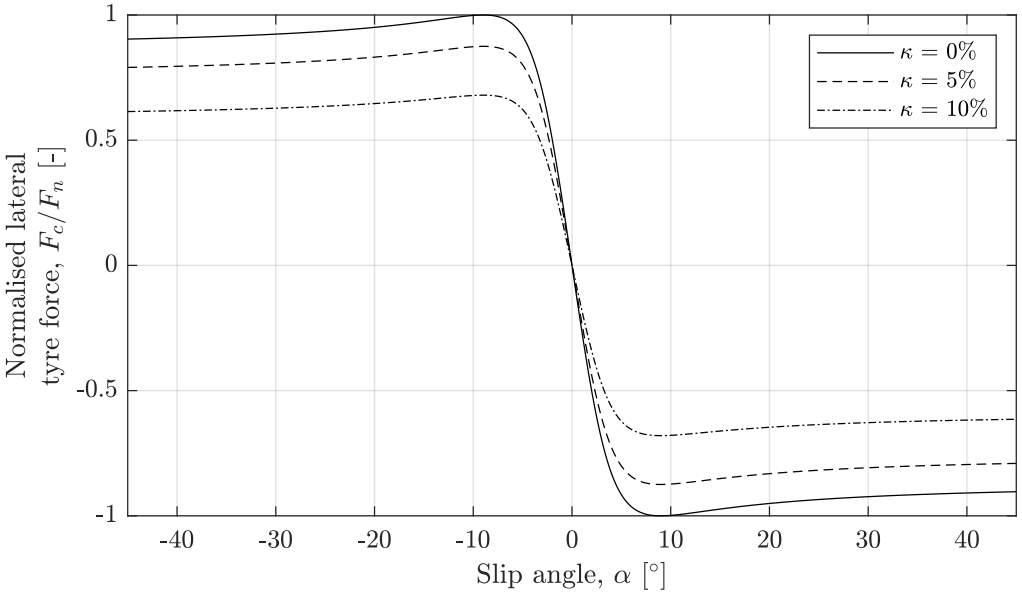


Figure 3.3: The Magic Tyre Formula handles combined slip conditions, where both longitudinal slip and slip angle are non-zero. The maximum lateral tyre force decreases for larger magnitudes of longitudinal slip

3.3.2 The Bicycle Model

Perhaps the most ubiquitous, and the most established, model for lateral vehicle dynamics, is the bicycle model. Despite being a simple model, it characterises the fundamental handling behaviour of the vehicle. This makes it an excellent starting point for many investigations into vehicle dynamics, and also makes it a good choice as a plant model in control applications, rather than just as a simulation tool.

The governing equations for the bicycle model are based on the following assumptions:

- The vehicle is travelling at a constant velocity u, travelling at a sufficiently high speed to prevent singularity in the slip angle term.
- The steer angle of the road wheels is small. Therefore, it is assumed that the steer angle of the inside and outside wheels is the same.
- The lateral acceleration of the vehicle is small, so that weight transfer between the inner and outer wheels can be neglected. This allows for the tyres on each axle to be treated as single tyres acting at the centre of each axle. This is where the model derives its name.

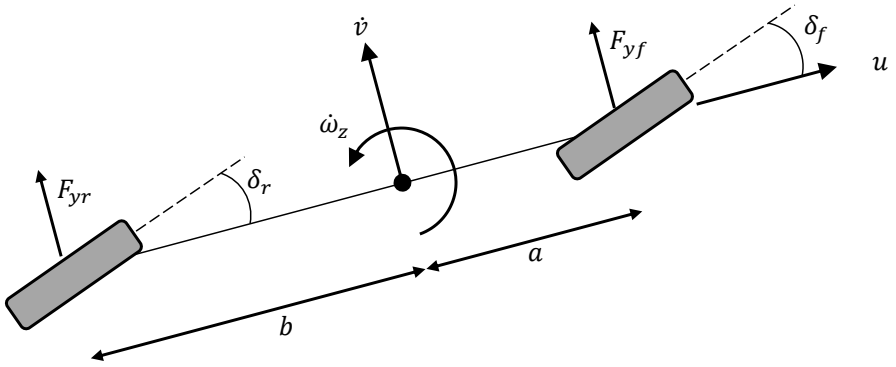


Figure 3.4: The linear bicycle model. Angles shown positive

- Due to the low acceleration and small steer angles, small angle approximations can be used, such that $\sin\theta \approx \theta$ and $\cos\theta \approx 1$.
- The tyre's lateral response is in the linear regime and not limited by the friction coefficient, so the brush model can be used to model the tyre force.

The resulting model is shown in Figure 3.4. From this, the equations of motion can be for the yaw angular acceleration and lateral acceleration can be derived from Newton's laws of motion:

$$\begin{aligned} m\left(\dot{v} + u\omega_z\right) &= F_{y\mathrm{f}} + F_{y\mathrm{r}} \\ I_z\dot{\omega}_z &= aF_{y\mathrm{f}} - bF_{y\mathrm{r}} \end{aligned} \tag{3.7}$$

As both the steer angles and lateral acceleration are small, the tyre forces are assumed to be linearly proportional to the tyre slip angle:

$$F_{y\mathrm{f}} = -C_{\alpha,\mathrm{f}}\alpha_{\mathrm{f}}, \qquad F_{y\mathrm{r}} = -C_{\alpha,\mathrm{r}}\alpha_{\mathrm{r}}. \tag{3.8}$$

Taking into account the position of the wheels, and the angular velocity of the vehicle and using a small slip approximation, the slip angles can be given by:

$$\alpha_{\mathrm{f}} = \frac{v + a\omega_z}{u} - \delta_{\mathrm{f}}, \qquad \alpha_{\mathrm{r}} = \frac{v - b\omega_z}{u} - \delta_{\mathrm{r}} \tag{3.9}$$

which is the application of Eq. (3.2). In most instances, the rear wheels

Table 3.2: Parameter values for the linear bicycle model.

Parameter	Value	Units
m	1,440	kg
I	2,000	kg·m^2
a	1.02	m
b	1.52	m

are not actively steered and so $\delta_r = 0$.

It is convenient to assemble the governing equations into state-space form:

$$\begin{pmatrix} \dot{v} \\ \dot{\omega}_z \end{pmatrix} = \begin{pmatrix} -\frac{C_\mathrm{f}+C_\mathrm{r}}{mu} & -\frac{aC_\mathrm{f}-bC_\mathrm{r}}{mu} - u \\ -\frac{aC_\mathrm{f}-bC_\mathrm{r}}{I_z u} & -\frac{a^2C_\mathrm{f}+b^2C_\mathrm{r}}{I_z u} \end{pmatrix} \begin{pmatrix} v \\ \omega_z \end{pmatrix} + \begin{pmatrix} \frac{C_\mathrm{f}}{m} & \frac{C_\mathrm{r}}{m} \\ \frac{aC_\mathrm{f}}{I} & -\frac{bC_\mathrm{r}}{I} \end{pmatrix} \begin{pmatrix} \delta_\mathrm{f} \\ \delta_\mathrm{r} \end{pmatrix} \tag{3.10}$$

In this form, the vehicle model is readily solved as a system of Ordinary Differential Equations (ODEs), and the system response is a function of both the system state and the instantaneous control input. Table 4.1 shows the parameter values used for the bicycle model in this thesis.

Verification

To verify the software implementation of the bicycle model, the numerical solution to the ODE system is compared to a steady-state, constant radius analysis. In this section, the verification is presented for front wheel steering only ($\delta_\mathrm{r} = 0$), although the analysis is readily extended to all-wheel-steering models. In steady state, the transient acceleration terms become zero ($\dot{v} = \dot{\omega}_z = 0$). This allows the steady state of the system in response to a step input to the control values by inverting the state matrix of the system:

$$\begin{pmatrix} v \\ \omega_z \end{pmatrix}_\mathrm{ss} = -\begin{pmatrix} -\frac{C_\mathrm{f}+C_\mathrm{r}}{mu} & -\frac{aC_\mathrm{f}-bC_\mathrm{r}}{mu} - u \\ -\frac{aC_\mathrm{f}-bC_\mathrm{r}}{I_z u} & -\frac{a^2C_\mathrm{f}+b^2C_\mathrm{r}}{I_z u} \end{pmatrix}^{-1} \begin{pmatrix} \frac{C_\mathrm{f}}{m} \\ \frac{aC_\mathrm{f}}{I} \end{pmatrix} \delta_\mathrm{f} \tag{3.11}$$

The bicycle models lateral velocity and yaw rate is simulated in response to a step input of 0.75° applied after 0.5 s, with a longitudinal velocity of

80 km/h (22.2 m/s). Along with the lateral velocity and the yaw rate of the system, the inertial lateral acceleration is also shown, which in steady state is equal to $u\omega_z$. Figure 3.5 shows that the system quickly reaches a steady-state, within 1 second of the steer input being applied. The lateral acceleration reaches a value of just over 0.25 g, which means that the assumptions made about low acceleration and large radius of curvature hold.

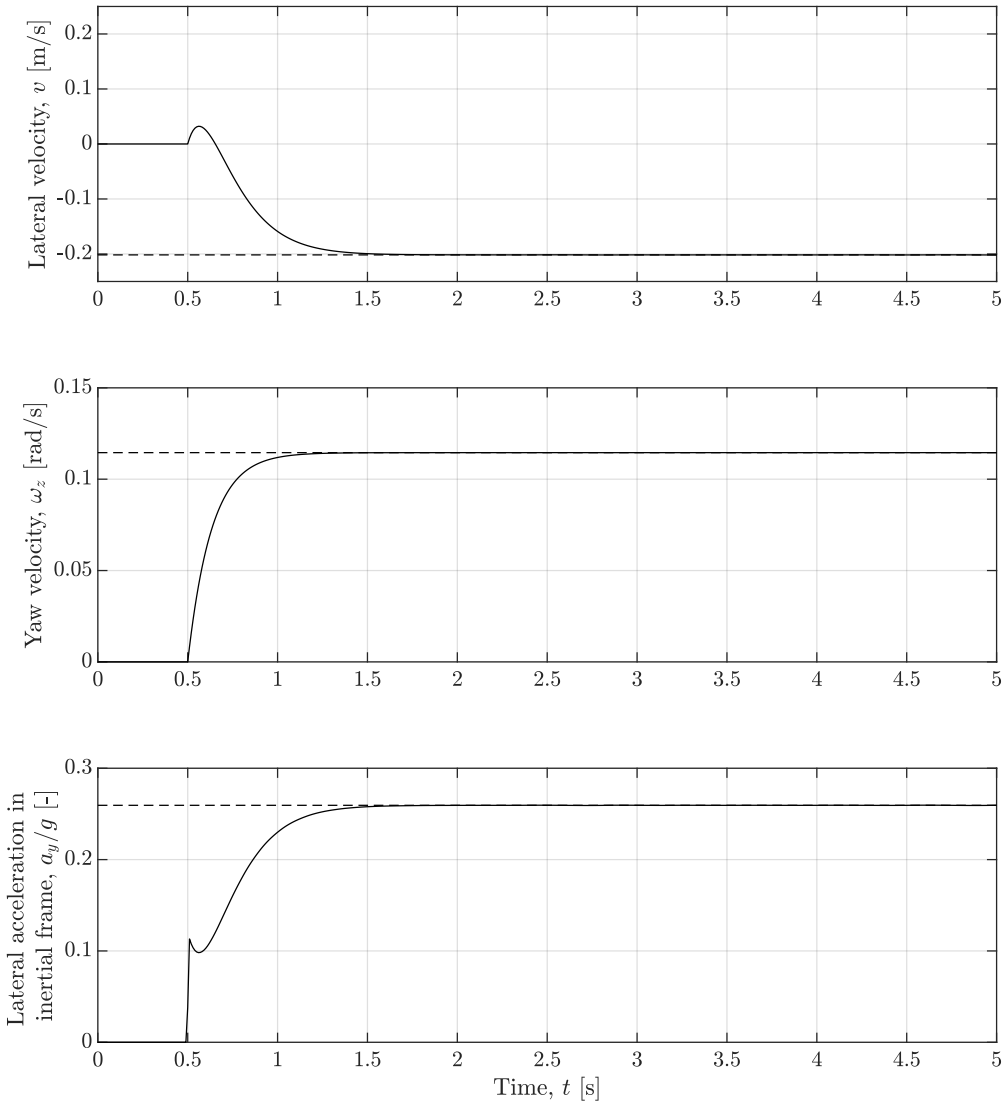


Figure 3.5: System response of the bicycle model to a 0.75° step input at 80 km/h (22.2 m/s). The steady-state values calculated from Eq. (3.11) are also shown. The system quickly approaches the predicted steady-state values

3.3.3 The Nonlinear Bicycle Model

While the linear bicycle model is a suitable vehicle model for many applications, it does have a number of limitations that mean that it is not always the most appropriate model. Due to the assumptions that were outlined previously, it is not representative of high-acceleration behaviour. When using the bicycle model in such scenarios, the system response will not match the physical system, as there is no tyre force saturation and therefore the resulting cornering force is unbounded.

Therefore, the linear bicycle model is extended by utilising the MTF instead of the brush model:

- The lateral tyre force is calculated using the MTF in pure lateral slip conditions ($\kappa = 0$).
- The slip angle is does not use small angle approximations to provide a more accurate model at larger steering angles.
- The normal force used in the MTF is calculated via the weight distribution of the vehicle.
- To allow for the evaluation of potentially high frequency control inputs, and to allow for the testing of the assumption that tyre transience is not needed in the linear model, the tyre transience is modelled via Eq. (3.6).

It should be noted that the resulting model still assumes that the vehicle dynamics can be well represented by a single-track vehicle. This is the case when the cornering radius of the vehicle is large, as the effect of the vehicles width on the inner and outer steer angles is negligible (i.e. $R_\mathrm{i} \approx R_\mathrm{o} \approx R$). Large radius of curvature cornering is often observed during high speed manoeuvres, such as motorway driving, which reinforces the single-track assumption as the load transfer across the axles is low due to the low cornering acceleration. Therefore, the normal force, which dictates the tyres lateral and longitudinal response, is approximately equal. In low-speed cornering manoeuvres a two-track model becomes necessary so that these effects can be incorporated into the simulated vehicle response.

The tyre slip angle is defined by the following equations:

$$\tan\alpha_{\mathrm{f}} = \frac{(v + a\omega_z)\cos\delta_{\mathrm{f}} - u\sin\delta_{\mathrm{f}}}{u\cos\delta_{\mathrm{f}} + (v + a\omega_z)\sin\delta_{\mathrm{f}}}, \quad \tan\alpha_{\mathrm{r}} = \frac{(v - b\omega_z)\cos\delta_{\mathrm{r}} - u\sin\delta_{\mathrm{r}}}{u\cos\delta_{\mathrm{r}} + (v - b\omega_z)\sin\delta_{\mathrm{r}}}, \tag{3.12}$$

rather than relying on the small angle approximation. The normal force acting upon a single tyre is given by:

$$F_{\mathrm{n,f}} = \frac{b}{2\,(a+b)}mg$$
$$F_{\mathrm{n,r}} = \frac{a}{2\,(a+b)}mg$$

As the MTF models are defined for a single tyre, Eq (3.7) has to be updated:

$$\begin{aligned} m\,(\dot{v} + u\omega_z) &= 2F_{y\mathrm{f}} + 2F_{y\mathrm{r}} \\ I_z\dot{\omega}_z &= 2aF_{y\mathrm{f}} - 2bF_{y\mathrm{r}} \end{aligned} \tag{3.13}$$

to account for the fact that there are two MTF tyres acting on the front and rear axles.

Verification

The nonlinear bicycle model is compared against the steady-state behaviour of the linear model given by Eq. (3.11). The relevant parameters are taken from Table 3.3. As with the linear bicycle model, a steer angle of 0.75° is applied to the system at 0.5 s. As expected, due to the fact that the system response is still well within the linear handling regime, the nonlinear model's results in Figure 3.6, are almost identical to the linear model's in Figure 3.5. The only observable difference is the initial lateral response at 0.5 s. Unlike the linear model, the acceleration response is not instantaneous due to the tyre transience present but the transient acceleration term, $\dot{v}$, rises very quickly. Despite this, the two model's behaviour is remarkably similar, which demonstrates that the common assumption that a tyre achieves its steady state slip angle instantaneously is valid.

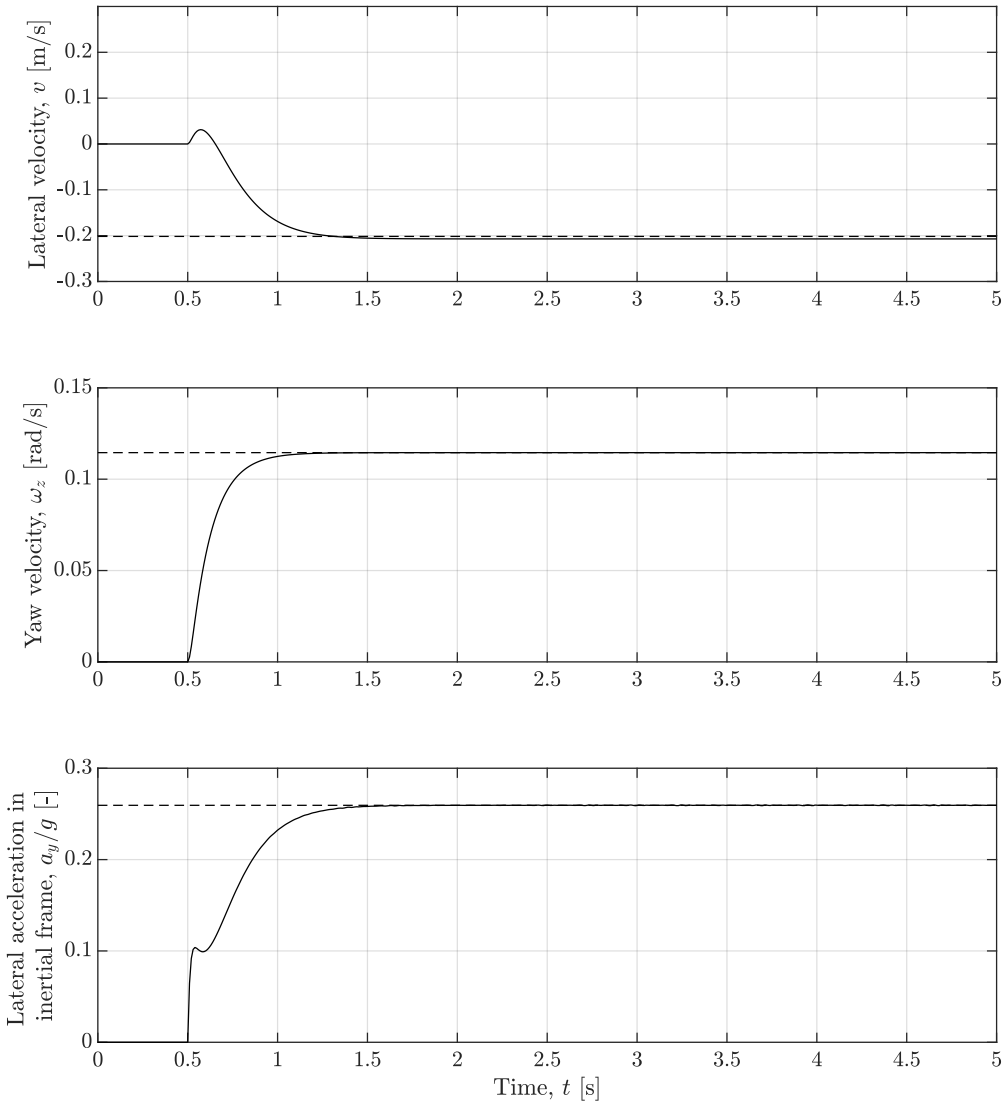


Figure 3.6: System response of the nonlinear bicycle model to a 0.75° step input at 80 km/h (22.2 m/s). The steady-state values calculated from Eq. (3.11) are also shown. The nonlinear model approaches the steady state values predicted by the linear model

3.3.4 Multi-Degree-of-Freedom Model

The models detailed so far give a reasonable approximation of vehicle behaviour, especially in steady state. However, there is certain dynamic behaviours that simpler models cannot simulate. As none of the models incorporate suspension into the vehicle design, they cannot represent manoeuvres that contain wheel lift-off, i.e. scenarios in which the tyre normal load of one or more of the wheels becomes zero. This behaviour may be seen in high-acceleration manoeuvre, which can occur during emergency scenarios such as collision avoidance [106]. In many instances, the controller is based on a simplification of the vehicle dynamics, such as the linear bicycle model, which are not capable of predicting this behaviour. The use of the fourteen degree-of-freedom model allows for the evaluation of such controllers in more extreme dynamic conditions, ensuring that the controllers response is adequate.

Although there are models of lower fidelity that provide some of the additional benefits, such as the seven DOF model [103] which along with the longitudinal, lateral and yaw acceleration of the vehicles centre of mass, models the rotation of the tyre, do not provide the additional benefits of modelling the dynamic normal force acting on the system due to the vehicles suspension system.

The equations defining the model, adapted from the work of Shim and Ghike [106], are summarised here due to their length and are instead presented in Appendix A. The main assumptions made for the model are as follows:

- The model consists of a sprung mass with six degrees-of-freedom, with the vertical displacement and wheel rotation also considered.
- The vehicle has two axles, with a track width (t_f and t_r for the front and rear track respectively) at which the tyre forces act.
- Each wheel is connected to the sprung mass via idealised spring and damper elements.
- The profile of the road surface is assumed to be perfectly flat.

- The role centres of the front and rear axles are assumed to be located at a fixed distance below the sprung mass, rather than being defined with reference to the ground.
- The roll centres are points about which the suspension forces act, rather than being kinematic constraints on the system.
- The tyre normal force and the spring and damper forces act perpendicular to the road surface and sprung mass respectively.

The model is presented in Appendix A. Due to the nature of the formulation (in particular calculating the vertical acceleration of the unsprung masses due to the damping and spring forces), the model cannot be represented as a system of ODEs ($\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, t)$) and instead is a system of Differential Algebraic Equations (DAEs) of the form $\boldsymbol{f}(\dot{\boldsymbol{x}}, \boldsymbol{x}, t) = \boldsymbol{0}$. There are commercial software packages to solve these equations, such as MATLAB, but the equations are in an implicit form, which means that certain cosiderations must be made when solving the equations to ensure numerical convergence [112]. Due to implicit form, vehicle models composed of Differential Algebraic Equations (DAEs) take much longer to solve than explicit models.

Table 3.3 shows the parameter values used for the Multi-Degree-of-Freedom (DOF), used for the numerical results in this thesis.

Driveline Model

The scenarios encountered in this thesis do not take place over a large speed range, and so it is assumed that when required, there is a torque value, up to a maximum, that is instantly applicable to the system. For vehicles with an Internal Combustion Engine (ICE), this assumption is acceptable so long as it can be assumed that the manoeuvre takes place over a small spped range in a single gear. This comes from the fact that if there is only a slight change in vehicle speed, the torque response will be relatively steady. This assumption fails if there were to be large fluctuations in torque response over a short time duration, as this would be unrealisable for an

Table 3.3: Parameter values for the Multi-Degree of freedom model, from [106].

Parameter	Value	Units
m	1,440	kg
I_x	900	kg·m^2
I_y	2,000	kg·m^2
I_z	2,000	kg·m^2
$m_{\mathrm{u,r}}$	80	kg
$m_{\mathrm{u,f}}$	80	kg
I_{w}	1	kg·m^2
a	1.02	m
b	1.52	m
t_f	1.50	m
t_r	1.50	m
h	0.75	m
$k_{\mathrm{s,f}}$	35,000	N/m
$k_{\mathrm{s,r}}$	30,000	N/m
$c_{\mathrm{s,f}}$	2,500	Ns/m
$c_{\mathrm{s,r}}$	2,000	Ns/m
$k_{\mathrm{t,f}}$	200,000	N/m
$k_{\mathrm{t,r}}$	200,000	N/m
$h_{\mathrm{rc,f}}$	0.65	m
$h_{\mathrm{rc,r}}$	0.60	m

ICE vehicle. It can be seen from the results shown, that this situation does not occur in the scenarios presented in this thesis.

For vehicles equipped with electric motors instead of an ICE, which is an increasingly prevalent situation, this assumption is even more valid. This is due to the quick response of the electric motor being able to supply a requested torque value, almost 'instantly' for the timescales under consideration.

Speed Control

Unlike the constant speed models presented thus far, the Multi-DOF model has a non-constant longitudinal (and vertical) velocity. This results in a negative longitudinal acceleration during cornering manoeuvres, which can be explained via observation of the acceleration terms for the Multi-DOF

model's unsprung mass:

$$\begin{aligned} a_x &= \dot{u} + \omega_y w - \omega_z v \\ a_y &= \dot{v} + \omega_z u - \omega_x w \\ a_z &= \dot{w} + \omega_x v - \omega_y u \end{aligned} \tag{3.14}$$

When the Multi-DOF model is being verified against steady state linear values or being used for constant speed manoeuvres, it is necessary to control the driveline torque to maintain a constant horizontal velocity.

Horizontal speed, u_{h}, being the magnitude of the u and v components of the vehicle's unsprung mass velocity is used rather simply the longitudinal velocity to match the common assumption of constant *speed* rather than constant longitudinal velocity. In practice, the performance of the controller is adequate with either value being used as the longitudinal velocity is typically a one or two of orders of magnitude higher than the lateral velocity, and so the difference between the longitudinal and horizontal velocity is negligible.

A Proportional-Integral (PI) controller is used to minimise the error between the horizontal speed u_{h} and the reference speed, which for verification and constant speed scenarios is taken to be the initial longitudinal velocity u_0. Two gain values are used to determine a driveline torque profile. For the integral action, the reference distance is taken to be the distance $u_0 t$ determined by the constant reference speed, while the measured value is taken from the magnitude of the X and Y displacements. X and Y are the vehicle's displacement in the global coordinate frame. This error term is then applied directly as the total driveline torque T_{d}, Table 3.4 shows the gain values used in this thesis.

The PI controller has the following control law:

$$u\left(t\right) = K_P e\left(t\right) + K_I \int_0^t e\left(\tau\right)\ d\tau, \tag{3.15}$$

where e is the error term between the system state and reference value. K_P and K_I are the gain values for the proportional and integral action,

respectively. Using the above reference terms results in the following torque value for the driveline torque:

$$T_{\mathrm{d}} = K_P \left(u_{\mathrm{h}} - u_0\right) + K_I \left(\sqrt{X^2 + Y^2} - u_0 t\right), \tag{3.16}$$

To prevent excessively high torque values that could cause traction problems, a secondary driveline torque value is calculated, based upon the concept of a tyre friction circle, using the remaining tyre grip 'leftover' from the amount of grip used to handle any instantaneous lateral acceleration. In this way, the maximum amount of tyre grip is used for both the lateral and longitudinal control objectives. This second driveline torque can be expressed as:

$$T_{\mathrm{d}} = \begin{cases} T_{\max} \frac{\sqrt{\mu^2 g^2 - a_{\mathrm{y}}^2}}{\mu g}, & |a_{\mathrm{y}}| \leq \mu g \\ 0, & |a_{\mathrm{y}}| > \mu g \end{cases} \tag{3.17}$$

where $T_{\max}$ is a driveline parameter which dictates the maximum torque available; in straight-line running ($a_{\mathrm{y}} = 0$) the driveline torque reduces to this value. In situations in which the lateral acceleration exceeds the maximum tyre grip, i.e. a sliding scenario, T_{d} is taken to be zero. This formulation is adapted from Osborn and Shim's work [103]; having been modified to saturate at a specified torque value $T_{\max}$.

To ensure that the constraints for both methods of calculating a driveline torque are always satisfied, the final driveline torque supplied to the vehicle is taken as the smaller of the values calculated by the two methods:

$$T_{\mathrm{d}} = \min\left(T_{\mathrm{d},1},\ T_{\mathrm{d},2}\right), \tag{3.18}$$

where $T_{\mathrm{d},1}$ is the driveline torque calculated by Eq. (3.16) and $T_{\mathrm{d},2}$ is the torque value calculated by Eq. (3.17).

Table 3.4: Gain values for the driveline torque to provide horizontal speed control.

Gain	Value
K_{P}	500
K_{I}	75

Finally, the torque value at each individual wheel is given by [103]:

$$\begin{aligned} T_{\mathrm{fi}} &= r_{\mathrm{IO}} r_{\mathrm{FR}} T_{\mathrm{d}} \\ T_{\mathrm{fo}} &= (1 - r_{\mathrm{IO}})\, r_{\mathrm{FR}} T_{\mathrm{d}} \\ T_{\mathrm{ri}} &= r_{\mathrm{IO}}\, (1 - r_{\mathrm{FR}})\, T_{\mathrm{d}} \\ T_{\mathrm{ro}} &= (1 - r_{\mathrm{IO}})\, (1 - r_{\mathrm{FR}})\, T_{\mathrm{d}} \end{aligned} \tag{3.19}$$

.

r_{IO} is the torque distribution between the inner and outer wheels, defined for a positive steer angle turn, and r_{FR} is the distribution between the front and rear axles. For this thesis, the inner-outer distribution is split equally, with $r_{\mathrm{IO}} = 0.5$, and the front-rear distribution is set equal to the weight distribution of the vehicle, $r_{\mathrm{FR}} = 0.6$.

Steady-State Verification

Figure 3.7 shows the verification of the Multi-DOF model against the bicycle model's steady-state values (Eq. (3.11)) in the linear handling regime. Although the transient behaviour of the system is more unsettled, the system quickly reaches the steady state values predicted by the bicycle model. The lateral velocity exhibits damped oscillating behaviour due to the interaction of the sprung mass and the unsprung masses of the wheels interacting via the spring and dampers of the suspension system.

As the Multi-DOF model exhibits more complex behaviour, it is also necessary to validate other components or subsystems. In particular, the suspension model needs careful validation to ensure that it is correctly representing the system behaviour. In steady-state during a cornering radius turn, it is possible to calculate the vertical load of the sprung mass acting

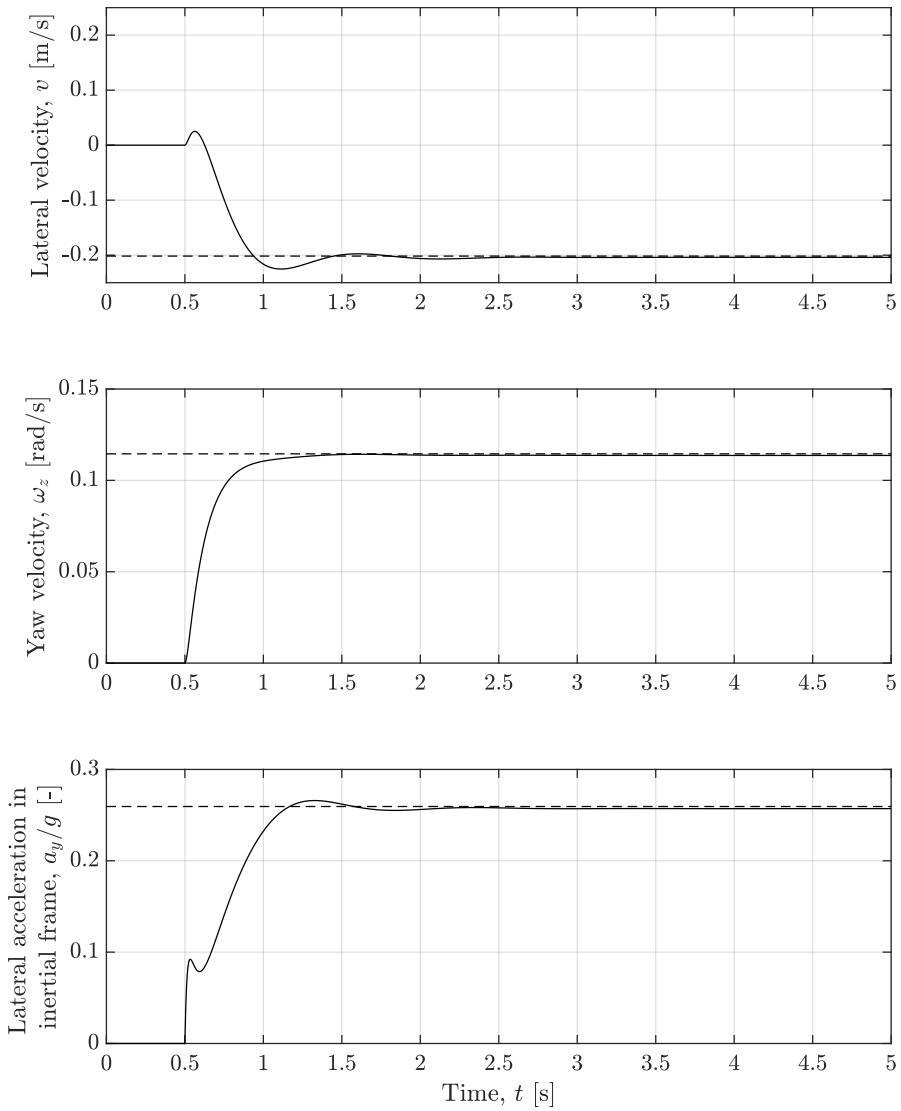


Figure 3.7: System response of the Multi-DOF Degree of Freedom model, with tyre transience, to a 0.75° step input at 80 km/h (22.2 m/s). The steady-state values calculated from Eq. (3.11) are also shown

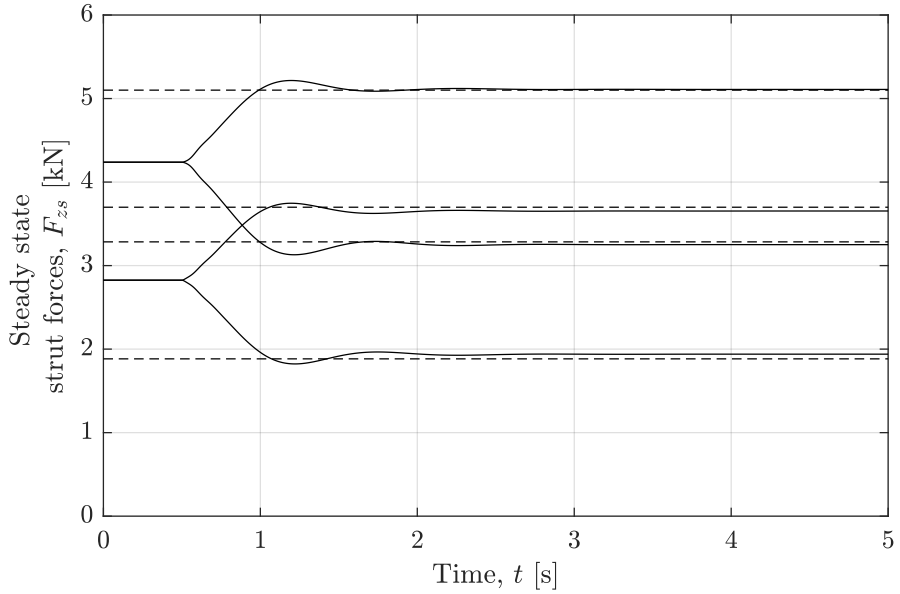


Figure 3.8: System response of the Multi-Degree of Freedom model, with tyre transience, to a 0.75° step input at 80 km/h

on each suspension strut by accounting for the static weight distribution with the additional load due to the inertial acceleration [103], and assuming that the front and rear roll moments are equal. This results in the following system of equations:

$$\begin{aligned} F_{z,\mathrm{sfi}} &= \frac{mb}{2\,(a+b)}g - \frac{mh}{2\,(a+b)}a_x - \frac{mh}{2t_\mathrm{f}}a_y \\ F_{z,\mathrm{sfo}} &= \frac{mb}{2\,(a+b)}g - \frac{mh}{2\,(a+b)}a_x + \frac{mh}{2t_\mathrm{f}}a_y \\ F_{z,\mathrm{sri}} &= \frac{mb}{2\,(a+b)}g + \frac{mh}{2\,(a+b)}a_x - \frac{mh}{2t_\mathrm{r}}a_y \\ F_{z,\mathrm{sro}} &= \frac{mb}{2\,(a+b)}g + \frac{mh}{2\,(a+b)}a_x + \frac{mh}{2t_\mathrm{r}}a_y \end{aligned} \tag{3.20}$$

Figure 3.8 shows the steady state values of the vertical forces acting upon the sprung mass when reacting to a 0.75° step steer input. The simulated strut forces respond smoothly to the steer input at 0.5 s and reach their steady state values approximately 1 s later. The steady state suspension forces strongly agree with the values determined by static weight plus inertial load, demonstrating that the vehicle is successfully modelling the force transfer through the suspension system.

Comparison of Model performance

It can be seen that the performance of the Multi-DOF model closely matches the steady state performance of the linear bicycle model. This shows that Multi-DOF model has been verified to correctly model the key steady-state behaviour. To observe the more complex dynamic behaviour of the system, a steering input is required that excites the system appropriately. The sine-dwell manoeuvre is commonly used in the automotive industry, in which a sinusoidal input is applied, before being held at the negative amplitude for a short time period.

It has been found that this is a particularly effective manoeuvre for disturbing the system, it having been shown theoretically, that from an optimal control perspective, it is the worst control input for the system to minimise state disturbances [108]. Clearly, for the analysis of system performance, this is the desired outcome.

BS ISO 19365:2016 [113] defines a sine-dwell steering input as follows:

$$\delta(t) = \begin{cases} A\sin\left(2\pi f t\right), & t \leq 0.75/f \\ -A, & 0.75/f < t \leq 0.75/f + t_d \\ A\sin\left(2\pi f\left(t - t_d\right)\right), & 0.75/f + t_d < t \leq 1/f + t_d \end{cases}$$

where A is the peak amplitude of the steering input, f is the steering frequency and t_d is the dwell time, the time at which the steering input is held at its negative peak amplitude. The steering input is zero for all times outside the defined window. The values used are 2.5°, 0.7 Hz and 0.5 s for the peak amplitude and steering frequency respectively. The steering angle is chosen to induce a large lateral acceleration in the vehicle, to study the nonlinear system response. Following the test laid out in the standard, the manoeuvre is carried out at 80 km/h (22.2 m/s).

Figure 3.9 shows the system performance of the Multi-DOF model compared to the linear and nonlinear bicycle models. It can be seen that the bicycle models represents the dynamic performance of the more complex model well. In particular, the yaw rate response is well-modelled, which in-

turn, and due to the high longitudinal speed, results in an accurate lateral acceleration response. The peak magnitude of the linear bicycle model's lateral acceleration is 0.83 g, while the Multi-DOF model has a peak acceleration of 0.76 g. The maximum error between the two acceleration profiles is 0.17 g, which is a 13% percentage error. The nonlinear bicycle model's maximum lateral acceleration is an even better match to the Multi-DOF response. The maximum error is 0.10 g, a 7.8% increase over the Multi-DOF model.

To give a more accurate representation of the Root Mean Square (RMS) error of the lateral acceleration, it is calculated between 1 and 4 seconds of the simulation. This is done to remove the steady state areas where there is zero error between the multi-DOF model and the respective bicycle models. The linear bicycle model has a RMS error of 0.08 g from the multi-DOF model. The nonlinear bicycles model is even smaller at 0.03 g.

There is a greater difference in the models transient lateral velocity. This can be ascribed to how the tyre force acts upon the system; in the bicycle model, the tyre force acts directly on the vehicle body itself. In the Multi-DOF model, the tyre forces act upon the vehicle body via the unsprung masses of the vehicle model.

Despite the added insight that a higher fidelity model can provide, this result shows that the bicycle models provide in this instance, an adequate modelling of the vehicle's lateral response. This is partly due to the test conditions, as the test is performed at a relatively high speed, the high lateral acceleration is achieved with a small slip angle. This means that the tyre force generated by the linear brush model agrees closely with the MTF model. At higher slip angles, the MTF model beings to saturate its output which does not occur with the brush model and the system responses would diverge.

The Multi-DOF has the disadvantage that it must be solved via implicit solvers. This is due to the inertial load acting upon the suspension struts. Implicit function solvers comparatively longer to simulate the system, making them less desirable during the design and analysis phase of controller

design. This, along with the comparative accuracy of the simulation model makes it less desirable for use as the vehicles simulation model.

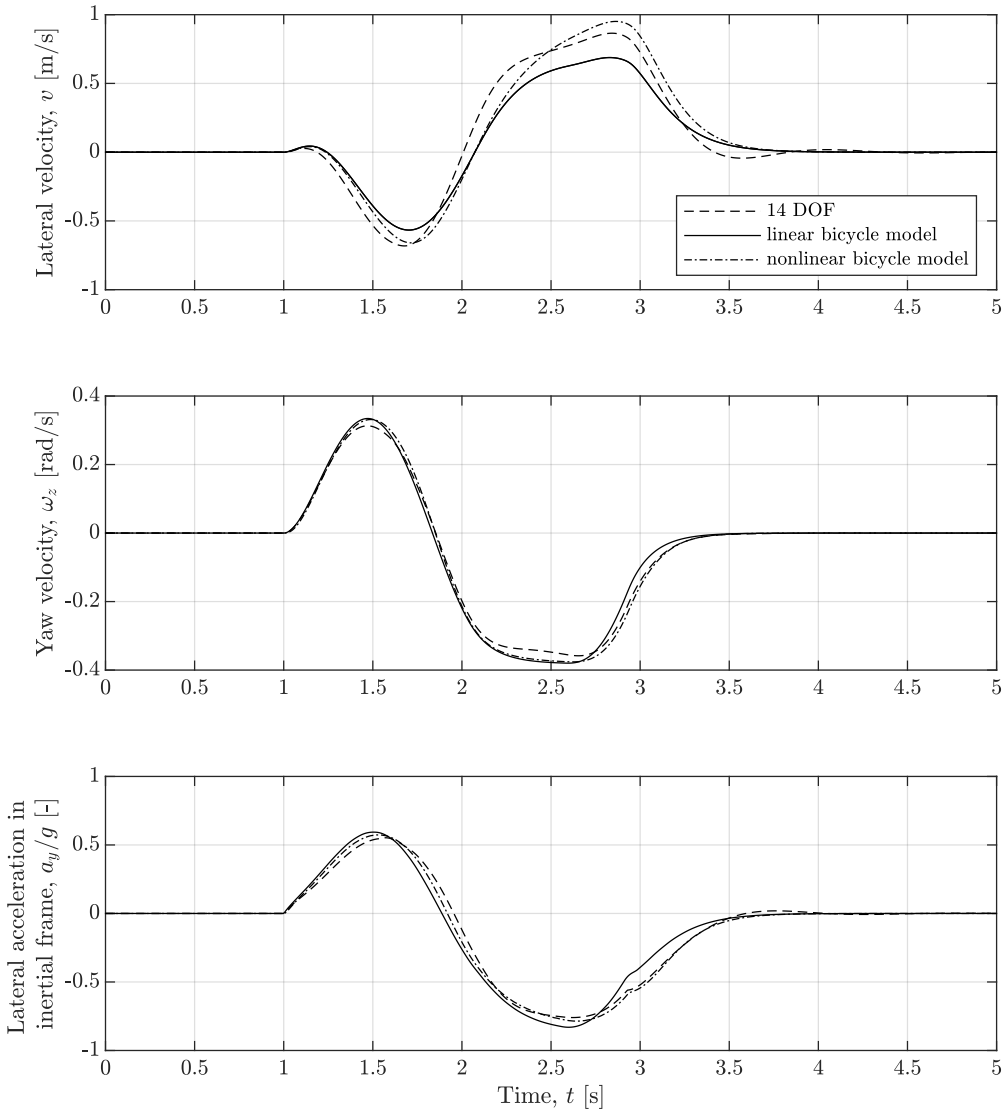


Figure 3.9: Comparison of the lateral velocity, yaw rate and lateral acceleration during a sine-dwell manoeuvre at 80 km/h. The peak amplitude is 2.5°, chosen to induce a large lateral acceleration response.

3.4 Displacement Tracking

The earliest research into vehicle dynamics focussed on understanding and characterising the stability of the vehicle. As such, the standard bicycle model was designed for the study of the lateral velocity and acceleration, along with the yaw rate. For autonomous driving applications, the models used need to be capable of modelling the displacement of the vehicle in its environment, and not just the distance travelled in the body-fixed frame of the vehicle.

The coordinate transform from the body-fixed frame to a Newtonian global reference is given by:

$$\begin{pmatrix} \dot{X} \\ \dot{Y} \end{pmatrix} = \begin{pmatrix} \cos\psi & -\sin\psi \\ \sin\psi & \cos\psi \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \tag{3.21}$$

and so the global position of the vehicle is given by the integration of Eq. (3.21) with respect to time. For the linear bicycle model, the following assumptions are made:

- $\dot{u} \approx 0$, constant longitudinal velocity
- $u \gg v$, the sideslip of the vehicle is small
- ψ is small so that the small angle approximations, $\sin\psi \approx \psi$ and $\cos\psi \approx 1$ are valid.

Which allows Eq. (3.21) to be simplified to:

$$\dot{X} = u + v\psi, \qquad \dot{Y} = v + u\psi \tag{3.22}$$

The coupled term $v\psi$ in the calculation of $\dot{X}$ prevents the inertial longitudinal displacement being expressed in state space form. However, based on the previous assumptions made, and the fact that the product of two small numbers (such as $v\psi$) is very small compared to the much larger value u

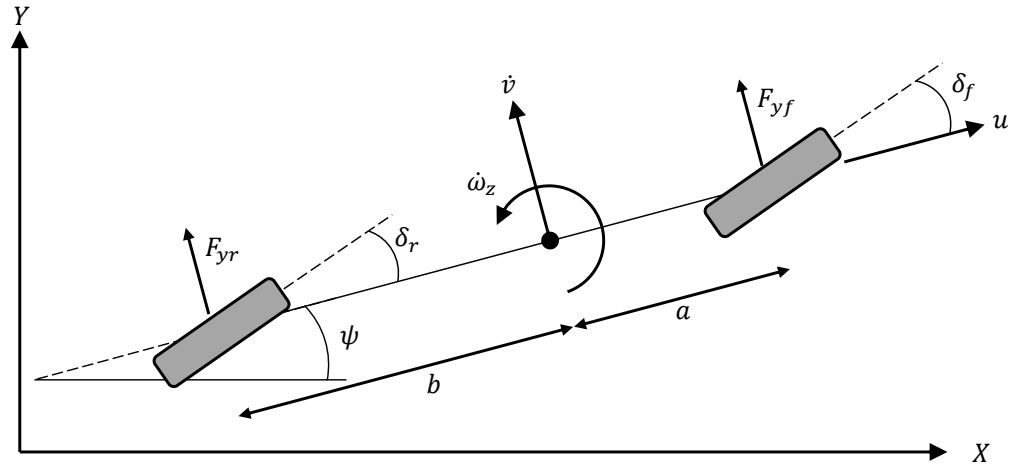


Figure 3.10: Linear biycle model with yaw and lateral displacement tracking

and can be neglected. The displacement terms are given by [76]:

$$\dot{X} = u, \qquad \dot{Y} = v + u\psi \tag{3.23}$$

$\dot{X}$ is no longer a state of the system, and so can be removed from the model. Also, the yaw angle of the vehicle of the vehicle is the integral of the yaw rate. This results in the following system of equations that allows for yaw and lateral displacement tracking:

$$\begin{pmatrix} \dot{v} \\ \dot{Y} \\ \dot{\omega}_z \\ \dot{\psi} \end{pmatrix} = \begin{pmatrix} -\frac{C_\mathrm{f}+C_\mathrm{r}}{mu} & -\frac{aC_\mathrm{f}-bC_\mathrm{r}}{mu} - u & 0 & 0 \\ 1 & 0 & u & 0 \\ -\frac{aC_\mathrm{f}-bC_\mathrm{r}}{I_z u} & 0 & -\frac{a^2C_\mathrm{f}+b^2C_\mathrm{r}}{I_z u} & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} v \\ Y \\ \omega_z \\ \psi \end{pmatrix} + \begin{pmatrix} C_\mathrm{f} \\ 0 \\ aC_\mathrm{f} \\ 0 \end{pmatrix} \delta_\mathrm{f} \tag{3.24}$$

Figure 3.11 shows the comparative performance of the nonlinear and linear terms. The global lateral displacement of the bicycle model, controlled by a step-input of 1° is shown. In Figure 3.11 (a), the lateral displacement is shown as a function of time, and in (b) it is shown as a function of longitudinal displacement. As expected, there is close agreement between the two methods, particularly over what is a short timescale. By the end of the five second simulation, the distance between the vehicles centre of mass as determined by the two methods is seven metres.

Over larger timescales, the error between the nonlinear and linear methods diverges. This is as the linear response is unbounded while the trigonometric terms in **??** bound the vehicles lateral response.

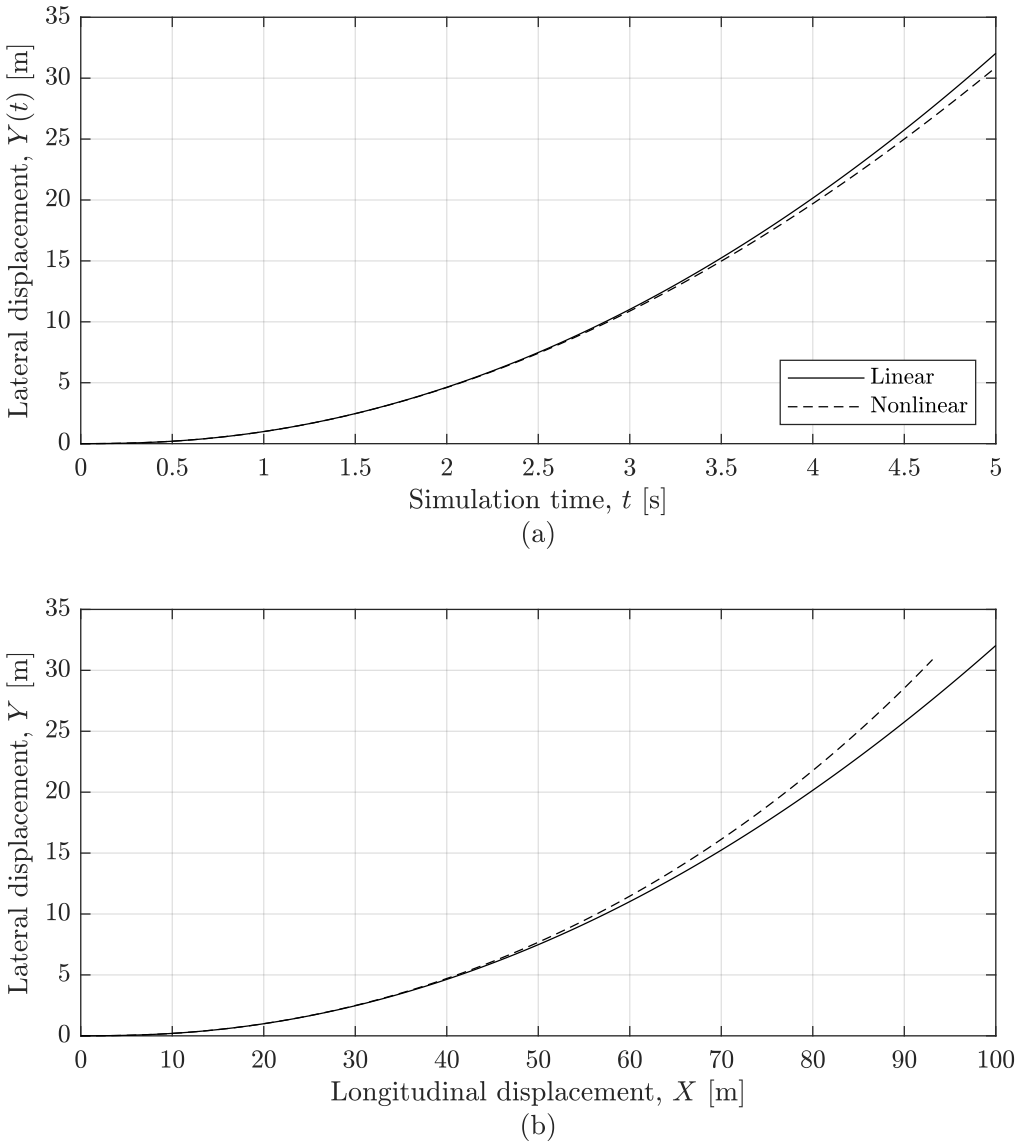


Figure 3.11: Lateral displacement of the bicycle model at 20 m/s with a 1° step-input, shown as both a function of time (above) and a plot of the vehicle in the XY plane.

3.5 Model Discretisation

In many control instances, the system dynamics are considered at discrete points in time rather than as continuous functions that depend on the rate of change of the system state. In these cases, a method of discretising the system is needed. With discretisation the continuous linear systems

$$\dot{\boldsymbol{x}} = \mathbf{A}_\mathrm{c}\boldsymbol{x} + \mathbf{B}_\mathrm{c}\boldsymbol{u} \tag{3.25}$$

is converted to the following form:

$$\boldsymbol{x}^+ = \mathbf{A}_\mathrm{d}\boldsymbol{x} + \mathbf{B}_\mathrm{d}\boldsymbol{u} \tag{3.26}$$

where $\boldsymbol{x}^+$ denotes the successor state of the system state $\boldsymbol{x}$ after a fixed time-step step Δt. In the following sections, two of the most common discretisation methods are outlined, along with the considerations to be made when deciding on a discretisation scheme.

3.5.1 Zero Order Hold

With Zero Order Hold (ZOH) discretisation, the sampled state is equal to the value of the continuous state at each interval Δt. For a continuous time system, such as Eq. (3.10), the discrete model can be calculated from the continuous model as [114, Chapter 4]:

$$\mathbf{A}_\mathrm{d} = e^{\Delta t \mathbf{A}_\mathrm{c}}, \qquad \mathbf{B}_\mathrm{d} = {\mathbf{A}_\mathrm{c}}^{-1}\left(\mathbf{A}_\mathrm{d} - \mathbf{I}\right)\mathbf{B}_\mathrm{c} \tag{3.27}$$

so long as $\mathbf{A}_\mathrm{c}$ is non-singular. ZOH discretisation ensures that the system state of discretised model maps directly to the continuous state at every integer multiple of the time-step Δt. Calculating the ZOH matrices is more computationally involved than other discretisation method, in particular due to the calculation of the matrix exponential. With the advances in modern day computer hardware, this is no longer an issue in most appli-

cations, especially when the state and control dimensions are in the tens or hundreds. For the models considered in this thesis, any computational differences are negligible.

3.5.2 Euler Approximation

A simpler discretisation is the Euler approximation, also known as FOH, which is based on the first order Taylor series expansion of Eq. (3.25), is leads to the following expressions for the discrete model:

$$\mathbf{A}_{\mathrm{d}} = \mathbf{I} + \Delta t \mathbf{A}_{\mathrm{c}}, \qquad \mathbf{B}_{\mathrm{d}} = \Delta t \mathbf{B}_{\mathrm{c}}. \tag{3.28}$$

The primary benefit of the Euler approximation is the ease with which it can be calculated, and that it preserves the physical intuition of the model. Apart from potentially on the main diagonal, the structure of the state matrix is preserved with this method, while ZOH discretisation does not preserve the continuous time structure. For large-scale systems with considerable structure, it would be beneficial to preserve the structure of the system to minimise the computation time of any calculation involving the state matrix.

The main drawback of the Euler approximation is that it does not preserve the stability of the continuous time system. For continuous time system, stability is obtained if the real components of all the eigenvalues of the state matrix are less than zero. For discrete time systems, stability is obtained when the absolute value of all the state matrix eigenvalues are less than one. Smaller values for the time-step typically decrease the magnitude of the eigenvalues, but it is not guaranteed that the system stability will be recovered for practical values of Δt.

Table 3.5 shows the effect of differing time-step values on the stability of the bicycle model. The first column shows the eigenvalues of the continuous time system. As they are both negative values, the system is stable. The remaining three columns show the absolute eigenvalues of the discretised system. With a time-step of 0.5 seconds, the discrete system is unstable,

Table 3.5: Eigenvalues of the bicycle model state matrix described by Eq. (3.10), and for the bicycle model discretised using the Euler approximation for three time-step values. Longitudinal velocity, u, 80 km/h

	$\mathrm{Re}(\lambda(\mathbf{A}))$	$\lvert(\lambda(\mathbf{I}+\Delta t\mathbf{A}))\rvert$		
		$\Delta t = 0.5$ s	$\Delta t = 0.25$ s	$\Delta t = 0.1$ s
1	-6.765	2.382	0.691	0.324
2	-7.542	2.771	0.885	0.246

but even with a time-step of 0.25 seconds the discrete system is stable.

In this thesis, a time-step of 0.1 seconds is used for discrete time controller designs. It can be seen in the final column that with this step size, the discrete system is also stable. For this reason, and the ease in which it can be implemented, the Euler method is the preferred discretisation method in this thesis. For the remainder of the text, when a state matrix is referred without qualification, i.e. $\mathbf{A}$, it shall refer to the discrete-time state matrix.

3.6 The Effect of Control Frequency on System Response

The use of digital microcontrollers to calculate control input necessitates the use of a discrete sequence, rather than a continuous spectrum, of control action. This discrete sequence is effectively a succession of step inputs of varying magnitudes, and the frequency of these steps determines whether the system response exhibits a measurable change in state over an infinitesimal timespan.

Such discrete behaviour is unlikely to be observed on a physical system, due to eventual realisation of the control input (such as a steering angle command applied to the front wheels), achieved through a sequence of electronic components and electric motors, operating on compliant mechanical components. The behaviour can be problematic when trying to characterise and control transient behaviour with reduced-order simulation

models. As an example, consider a predetermined, discrete sequence, of steer angle values applied to the nonlinear two state bicycle model detailed in Section 3.3.3, solved using a Runge-Kutta integration scheme using adaptive time-step control [115]. This ensures that the transient dynamics of the system are accurately modelled. We apply the control sequence to system at two frequencies to show the transient systems response dependency on it. The input frequencies where chosen based on the settling time of the system, approximated by calculating the eigenvalues of the linearised bicycle models state matrix.

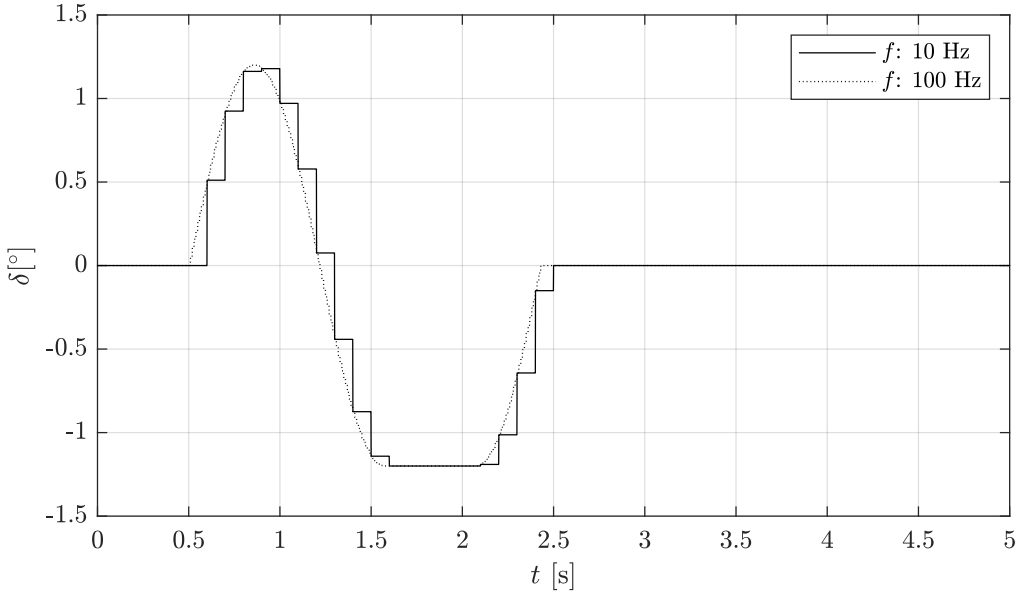
To demonstrate this behaviour, the sine-dwell steering function used earlier to study the dynamic system of the behaviour of the Multi-DOF model is used. The amplitude of the control input, 3 degrees, is chosen as specified in the standard to provide a steady-state lateral acceleration of 0.3 g. This continuous input can be discretised by substituting the continuous time variable, t, for its discretised output, τ. Figure 3.12 (a) shows the control sequence for the given frequencies. The discrete-time variable is calculated by

$$\tau = \frac{1}{f} \lfloor t \cdot f \rfloor$$

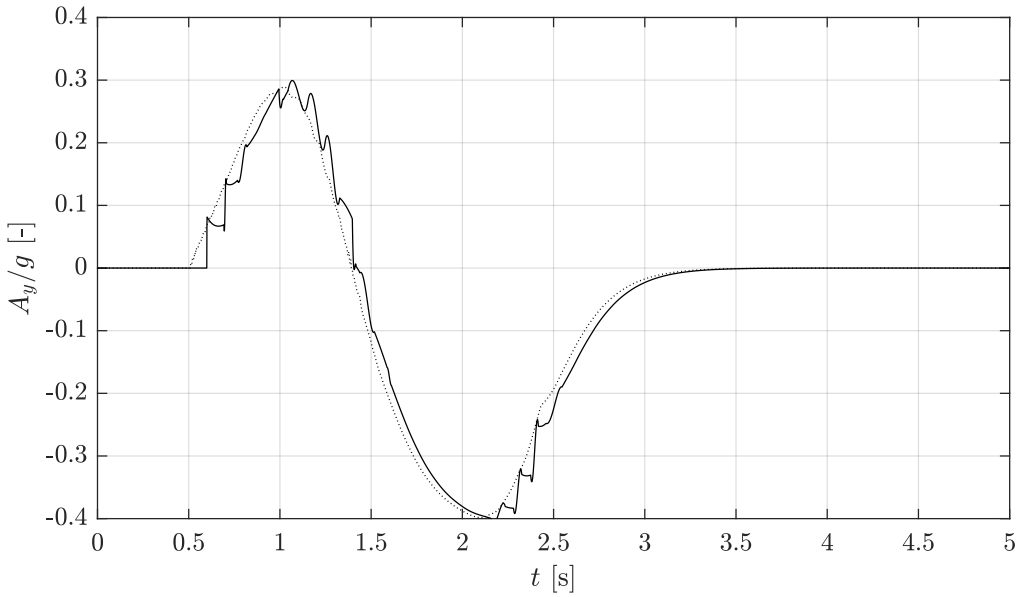
where $\lfloor \cdot \rfloor$ denotes the floor function.

Figure 3.12 (a) shows the discretised steering input compared to the continuous time steering input. Figure 3.12 (b) shows the lateral acceleration of the simulation model when controlled by the discrete steering inputs. With the 10 Hz control input, the sine dwell manoeuvre has an obviously stepped profile, and that the vehicle response is a sequence of peaks as the steering input changes.

At 100 Hz, the input profile looks smooth to the human eye, and the acceleration response is smooth as the control input value changes before the system has fully settled from the previous response. Despite the stepped acceleration response at 10 Hz, the lateral displacement response (not shown) is smooth. In addition, a 100 Hz steering frequency is too computationally prohibitive for implementation and so is not used. For the rest of the thesis, a 10 Hz operating frequency is used for the steering controller.



(a) The discrete sine-dwell steering input.



(b) The lateral acceleration response of the biycle model at 80 kph.

Figure 3.12: The lateral acceleration response of the two state bicycle model when subjected to a sine-dwell steering input applied at discrete time-instances. With a control input of 10 Hz, a stepped response in lateral acceleration can be observed, which is not observed at higher frequencies.

Despite the stepped acceleration response at 10 Hz, the lateral displacement response (not shown) is smooth and a 100 Hz steering frequency is to computationally prohibitive for implementation in this research.

3.7 Conclusions

In this chapter, a number of vehicle models have been presented and verified. They provide a framework for both the simulation of vehicle behaviour and as linear plant models for use in control design and implementation. In particular, it was shown that for first-order response behaviour, such as yaw rate response, the linear bicycle model has impressive agreement with a much more complex Multi-DOF vehicle model which includes suspension effects and a nonlinear tyre model. The maximum lateral acceleration error between the two models, during the sine-dwell manoeuvre, was 0.17 g, a 13.6% under-prediction by the linear bicycle model. Despite this, the RMS average error between the two lateral acceleration profiles is only 0.06 g over the entire manoeuvre. For the work presented in the remainder of the thesis, the focus is on developing and evaluating human factor aware vehicle behaviour, during the normal driving conditions. Because of this, the vehicles lateral acceleration will typically be less than 0.3 g, well within the valid range of linear bicycle model behaviour.

The bicycle model showed limitations with the modelling of integral terms. It was shown that it was a poor estimator of lateral displacement in the global XY-plane for any time-scales greater than a couple of seconds. Displacement tracking is a crucial element of autonomous driving and shows that careful consideration when implementing reference tracking of these terms.

For these reasons, the linear bicycle model will be used as the control plant as it allows for the MPC control problem to be solved quickly, efficiently and with a globally optimal solution. These issues will be discussed in the next chapter. For the simulation model, the nonlinear bicycle model,

with tyre transience, will be used. This allows for accurate displacement in the global coordinate frame. The tyre transience allows for the study of applying a discrete control sequence to the simulation vehicle at a relatively high frequency and helps to provide a more realistic system response by smoothing the acceleration profile.

The high-acceleration sine-dwell manoeuvre also showed that it is not necessary to make use of a more accurate, and correspondingly much more complex, simulation model to to complete the thesis objectives. By using the simpler nonlinear bicycle model, more time can be spent on the analysis of the controller behaviour, than a moderate increase in model fidelity with much longer simulation times. For a further and more complete analysis it would be necessary to use a higher-fidelity simulation model, so that other factors that influence ride comfort, such as body roll, can be taken into account.

Model discretisation was also studied, with a focus on the Zero Order Hold (ZOH) and Euler methods. ZOH offers a direct mapping between the continuous time model and the discrete model at the discrete time intervals. As it preserves system properties of continuous-time model, such as stability, it is an attractive possibility. The computational cost of calculating matrix exponentials in its formulation made it potentially prohibitive, but with the power of modern computer hardware, this would appear to be no longer an issue.

The Euler method by comparison is a much simpler model and is not guaranteed to preserve system stability for larger time-steps. It does preserve any structure, such as sparsity, of the system. For large scale systems, this could be of great benefit, but for the vehicle model applications of this thesis, this has no material effect. Its primary benefit is perhaps its simplicity and ease of implementation, especially as for practical time-step values considered in this thesis for autonomous driving, typically on the order of 0.1 s, the discrete models are still stable, as was demonstrated with the linear bicycle model. For the primary reason of simplicity, the Euler approximation will be used as the sole discretisation method for the rest of

the thesis.

The effect that controller discretisation was further considered when the effect of controller actuation frequency was studied for when a discrete control input is applied to a continuous-time plant. It was shown, that if the sampling time of the controller is too large, a step response behaviour is observed in the response.

Chapter 4

Model Predictive Control for Active Steering Control

4.1 Introduction

Autonomous driving behaviour, such as Active Steering Control (ASC), is a complex task and undertaken in an environment that cannot be fully determined ahead of time, with many unknown variables present, even in less challenging automotive scenarios, such as highway driving. This presents a distinct challenge, in which the operating regime of the autonomous vehicle cannot be exhaustively studied ahead of time.

A suitable technique, which can react to its current operating environment is needed. Model Predictive Control (MPC) is one such technique. In particular, the ability for MPC to handle state and control constraints explicitly in the control problem formulation makes it a particularly powerful framework for autonomous driving.

With MPC, a model of the system plant is used to predict the system behaviour in the near future. Optimisation of the control inputs is then performed to minimise the deviation of the predicted future state compared to reference values. Once the optimal control sequence has been determined, the first control input from the sequence is applied to the system.

At the next call to the controller, the system state is sampled again, and

the process is repeated. This method of repeatedly solving an optimisation problem online, with measurements of the system state updated every time-step, is how MPC provides feedback control, and prevents the accumulation of errors dues to model-system mismatch propagating through time.

The use of a predictive model also allows MPC to show proactive, rather than reactive, control. MPC anticipates future undesirable behaviour and takes steps to mitigate behaviour before it occurs, but it is the use of mathematical optimisation is that allows for what is perhaps MPC *raison d'être*: constraints. Linear MPC can be implemented as a Quadratic Programming (QP) problem, allowing for any linear combination of predicted system state and control action to be constrained. This allows, so long as there is a feasible solution, for a control designer to impose constraints upon the system, that ensure a system response that is within a predetermined operating regime. This feature of MPC will be exploited later to incorporate the preference of individual drivers explicitly into the control problem; allowing for a multitude of driving styles to be accommodated by a single, unified, control framework.

The remainder of this chapter is as follows: firstly, the formulation of linear MPC used for the remainder of the thesis is then presented, before considerations needed to ensure system stability are discussed. Some preliminary results of an MPC application are then presented, with ASC used to perform a lane change step manoeuvre. The flexibility of MPC with respect to system configuration and constraints is then demonstrated, with results compared for both front-wheel steering and all-wheel steering control. The state space model, with displacement reference tracking developed in Chapter 3 is used as the controller's plant model. The ASC controller is evaluated in simulation with the nonlinear bicycle model used as the simulation model.

Then, the effect of a standard stabilising feature; a terminal state penalty matrix is shown. The examples shown also demonstrate that the reference provided to an MPC controller can have affect performance and stability.

This enforces the notion that for autonomous driving behaviour, reference generation and reference tracking must be considered as two parts of a single problem. Any successful autonomous driving application must consider both aspects as equal challenges in the design process. This key observation is the foundation for the work in Chapter 5, where a hierarchical MPC framework is explored.

4.2 Formulation

This thesis utilises a linear MPC formulation, which is outlined below. Consider a discrete state space model, such as one discussed in Section 3.5. The vector $\boldsymbol{x} \in \mathbb{R}^n$ is the system state vector, and $\boldsymbol{u} \in \mathbb{R}^m$ is the control vector. The system also has a vector of outputs $\boldsymbol{y} \in \mathbb{R}^q$. The linear prediction model for the system is then given by:

$$\begin{aligned} \boldsymbol{x}^+ &= \mathbf{A}\boldsymbol{x} + \mathbf{B}\boldsymbol{u} \\ \boldsymbol{y} &= \mathbf{C}\boldsymbol{x} + \mathbf{D}\boldsymbol{u} \end{aligned} \tag{4.1}$$

The term $\boldsymbol{x}^+$ denotes the successor system state, one discrete time-step into the future. The matrices $\mathbf{A} \in \mathbb{R}^{n\times n}$ and $\mathbf{B} \in \mathbb{R}^{n\times m}$ are called the state matrix and control matrix respectively, while $\mathbf{C} \in \mathbb{R}^{q\times n}$ and $\mathbf{D} \in \mathbb{R}^{p\times m}$ are the output and feedforward matrices. In many cases, and for the remained of this thesis, the feedforward matrix is zero. For this reason, the term

$$\boldsymbol{y} = \mathbf{C}\boldsymbol{x} \tag{4.2}$$

is used, instead of Eq. (4.1), to state the system output. Indeed, for this thesis, full state output is assumed, such that $\boldsymbol{y} = \boldsymbol{x}$ ($\mathbf{C} = \mathbf{I}_n$), and the output is simply referred to as the system state.

With this model, it is possible to predict the future system state up to the preview horizon, which is N discrete time-steps in the future. Then

the following defines the regulatory MPC problem, $\mathbb{P}_N(\boldsymbol{x})$:

$$\underset{\boldsymbol{\mathcal{U}}}{\text{minimise}} \quad J = \sum_{i=1}^{N} \|\boldsymbol{y}_i - \boldsymbol{y}_i^{\text{r}}\|_{\mathbf{Q}}^2 + \|\boldsymbol{u}_{i-1}\|_{\mathbf{R}}^2 \tag{4.3a}$$

$$\text{subject to} \quad \boldsymbol{x}_0 = \tilde{\boldsymbol{x}}_0, \tag{4.3b}$$

$$\boldsymbol{x}_{i+1} = \mathbf{A}\boldsymbol{x}_i + \mathbf{B}\boldsymbol{u}_i, \tag{4.3c}$$

$$\boldsymbol{y}_i = \mathbf{C}\boldsymbol{x}_i, \tag{4.3d}$$

$$\boldsymbol{x}_i \in \mathcal{X} \qquad i = 1, ..., N, \tag{4.3e}$$

$$\boldsymbol{u}_i \in \mathcal{U} \qquad i = 0, ..., N-1, \tag{4.3f}$$

$$\boldsymbol{y}_i \in \mathcal{Y} \qquad i = 1, ..., N \tag{4.3g}$$

In this optimisation problem, the goal is to find the optimum control sequence $\boldsymbol{\mathcal{U}} := \{\boldsymbol{u}_i \mid i = 0, \ldots, N-1\}$ that minimises the cost function J. The vector $\boldsymbol{x}_i^r$ denotes the state reference value at the i-th time-step. $\|\boldsymbol{x}\|_{\mathbf{M}}^2 = \boldsymbol{x}^T\mathbf{M}\boldsymbol{x}$ denotes the square of the $\mathbf{M}$-weighted norm, where $\mathbf{M}$ is a positive-definite matrix such that $\boldsymbol{x}^T\mathbf{M}\boldsymbol{x} > 0$ for all values of $\boldsymbol{x}$. These weighted terms allow for the prioritisation of different system behaviours. The state penalty matrix is $\mathbf{Q} \in \mathbb{R}^{n\times n}$ and the control penalty matrix is $\mathbf{R} \in \mathbb{R}^{m\times m}$. It is common for the penalty matrix to be a diagonal matrix, in which case the individual state or control variables are penalised, rather than linear combinations of them.

Large values in the state penalty matrix encourages the controller to minimise the tracking error of that respective state and increases the feedback action of the controller. It is by the selection of these weights that the closed-loop action of the controller can be modified.

Large penalty weights on the control inputs discourage control action, which allows for both the prioritisation of certain control actuators over others, along with the possibility of allowing small steady-state tracking errors to prevent a hunting behaviour of the control inputs.

Equations (4.3b) to (4.3g) define the constraints on the system. Equation (4.3b) initialises the predicted state at the start of the control problem to the measured system state, denoted $\bar{\boldsymbol{x}}$. This is the mechanism that al-

lows for feedback control in MPC. Equations (4.3c) and (4.3d) force the system state and output to evolve as specified by the prediction model.

Equations (4.3e), (4.3f) and (4.3g) are the state and control constraints. Typical implementations for the constraints are as upper and lower bounds on each variable or constraining them via polytopic or ellipsoidal constraints. It is also common to constrain the maximum change in control variables between time-steps, commonly referred to as slew-rate constraints. For a linear MPC formulation, the constraint sets $\mathcal{X}$,$\mathcal{U}$ and $\mathcal{Y}$ can be represented by linear inequalities:

$$\begin{aligned} \mathcal{X} &:= \{\boldsymbol{x} \mid \mathbf{P}_{\boldsymbol{x}}\boldsymbol{x} \leq \boldsymbol{h}_{\boldsymbol{x}}\}, \\ \mathcal{U} &:= \{\boldsymbol{u} \mid \mathbf{P}_{\boldsymbol{u}}\boldsymbol{u} \leq \boldsymbol{h}_{\boldsymbol{u}}\}, \\ \mathcal{Y} &:= \{\boldsymbol{y} \mid \mathbf{P}_{\boldsymbol{y}}\boldsymbol{y} \leq \boldsymbol{h}_{\boldsymbol{y}}\}, \end{aligned} \tag{4.4}$$

where the matrices $\mathbf{P}$ and $\boldsymbol{h}$ for each constraint set are of appropriate dimensions.

Once an optimal solution to the MPC problem has been found, the control vector from the first time-step of the control sequence is then applied to the system. At the next call to the MPC controller, the system state is measured again, and the optimisation problem is solved starting from the updates system state. This process results in the implicit, nonlinear, MPC control law [71]:

$$\boldsymbol{\kappa}_N(\boldsymbol{x}_0) := \boldsymbol{u}_0 \tag{4.5}$$

with $\boldsymbol{u}_0$ being the first control input vector from the optimal control sequence $\boldsymbol{\mathcal{U}}$. In summary, at each discrete time-step, a control sequence is computed by solving an optimal control problem, and the first control in this sequence applied to the system plant.

This formulation, referred to as regulatory form [97], represents one approach, of many, for defining the MPC problem. Depending on the specific aims of the control design, there are other methods, such as the condensed form, removes the system states from the formulation through successive application of the state transition equation Eq. (4.3c), see for example, [76].

However, as $\mathbb{P}_N$ can be formulated as a standard QP (see Appendix C), efficient solution methods exist and so the condensed form is not necessarily quicker to solve [75], despite potentially having a significantly smaller size; which could benefit hardware constrained implementations. This formulation, which removes the ability to constrain system states, is inappropriate for the thesis aims and so is not considered further.

4.2.1 Stability

The MPC framework defined above allows for the implementation of a controller that can effectively track a reference trajectory while subject to state and control constraints. This basic formulation does not guarantee, a stable closed-loop control system. Depending on several system and controller parameters, the controller might not be able to maintain system stability without additional considerations in the MPC control formulation.

Some of the techniques used to provide stability will be discussed in this section. However, as the aims of this thesis are focussed on exploring how MPC as a strategy can implement effective personalised driving behaviour, the emphasis will be to highlight matters that need to be considered when implementing an MPC controller, and whether the techniques have any effect on task performance, rather than focussing on a specific stabilisation implementation for any of the proposed control strategies in this thesis.

The simplest stabilisation technique is to simply ensure that the prediction horizon is sufficiently long. This intuitively aids system stability; allowing the controller to consider more of the future states and determine a control input which stabilises the long-term future of the system. In general, for a stabilisable linear system with a fixed reference state, an infinite prediction horizon guarantees a stable system without any other stabilisation considerations [90]. In most cases however, it is not feasible to perform this infinite prediction and summation of future states. In practice, it is often enough to choose a sufficiently large finite value for the preview horizon. Unfortunately, there is no method to determine precisely how large a preview horizon needs to be, although certain heuristics do

exist, particularly with regards to the system state and control dimensions [116].

The temptation is to choose an overly conservative value, that far exceeds the necessary value, to ensure stability even when encountering unusual initial conditions or reference that were not considered when developing the MPC controller. The problem then is that the time required to calculate the optimal control action is tightly coupled to the preview horizon value. For the solution of a linear MPC algorithm, using a generic QP solver, the computational complexity is $O(N^3)$; proportional to the cube of the preview horizon [117]. Specialised QP solvers that exploit the unique structure of the MPC problem can reduce this proportionality to just the preview horizon value $O(N)$ [118], but there is still a limit to the maximum value that can be used for the preview horizon, before the computational cost is prohibitive.

For these reasons, in practice it is not enough to use the preview horizon as the lone stabilisation technique, except for simple or small-scale systems. It is, however, still a crucial aspect of controller design in other stabilisation techniques, as will be seen later. A common stabilisation technique is the use of a terminal state penalty. With this technique, the state penalty matrix, $\mathbf{Q}$, is replaced at the final prediction step of the optimisation routine, with the terminal weight matrix, $\mathbf{Q}_f$, resulting in a new cost function for the optimisation problem:

$$J = \|\boldsymbol{x}_N - \boldsymbol{x}_N^r\|_{\mathbf{Q}_f}^2 + \sum_{i=1}^{N-1} \|\boldsymbol{x}_i - \boldsymbol{x}_i^r\|_{\mathbf{Q}}^2 + \sum_{i=0}^{N-1} \|\boldsymbol{u}_i\|_{\mathbf{R}}^2 \tag{4.6}$$

One choice for $\mathbf{Q}_f$ is the solution to theDiscrete Algebraic Riccati Equation (DARE), as it was shown that for linear systems, this recovers the cost of the future states beyond the preview horizon, of the infinite-horizon optimal control problem by assuming that Linear Quadratic Regulator (LQR) optimal state feedback is used after the preview horizon [86, 117].

The utility of using this as the terminal state penalty matrix is clear, by recovering the remainder of the infinite-horizon cost, the controller is

stable, without having to increase the preview horizon. This, and the fact that the DARE can be solved *a priori*, allow a design that can be easily implemented with current hardware, as the computational complexity of the control problem can be minimised during the design phase.

In this thesis, a preview horizon of sufficient length is used to help stabilise the system for the software simulations used, but no proof of stability is presented. Therefore, additional stability measures, such as a terminal state constraint set discussed in Chapter 2, are not required for the completion of this thesis objectives as the system has adequate stability, and so are not implemented.

4.2.2 Demonstrating the Control Performance

To demonstrate the benefits of MPC in autonomous driving scenarios, Active Steering Control (ASC) is demonstrated, where MPC is used to calculate a sequence of steer angles to track a set of reference lateral displacement and yaw trajectories, along with the yaw velocity and lateral displacement. First, the controller will be used to track a step change in lateral displacement, with a constant, zero, yaw trajectory without state or control constraints. The simulation is then repeated with these constraints active to show their effectiveness. Then, a more realisable lateral displacement reference will be used, created by using an exponential sigmoid function. This allows for a reasonable yaw reference tracking value to be generated. Finally, the plant model will be changed to include rear steering, demonstrating how MPC allows for the control of multi-input, multi-output systems.

The linear bicycle model, with displacement and yaw tracking, presented in Eq. (3.24) is used for the plant model terms $\mathbf{A}$ and $\mathbf{B}$. It is discretised using the forward Euler method, Eq. (3.28), using a time-step, Δt of 0.1 s. The parameter values used are the same as those in Chapter 3 Table 4.1.

The nonlinear bicycle model is used as the simulation model, again with the same parameter values shown in Table 3.1 and Table 3.3. The following

Table 4.1: Parameter values for the linear bicycle model.

Parameter	Value	Units
m	1440	kg
I	2000	kg·m^2
a	1.02	m
b	1.52	m

Table 4.2: MPC parameter values.

Parameter	Value	Units
N	20	-
$\boldsymbol{u}$	5	$\circ$
$\Delta\boldsymbol{u}$	10	$\circ$/s

terms were used as the penalty matrices:

$$\mathbf{Q} = \begin{bmatrix} 5 & 0 & 0 & 0 \\ 0 & 25 & 0 & 0 \\ 0 & 0 & 10 & 0 \\ 0 & 0 & 0 & 5 \end{bmatrix}, \qquad \mathbf{B} = 1, \tag{4.7}$$

chosen from experience to provide good tracking performance, with the lateral displacement prioritised to ensure good path following characteristics.

Step-change reference tracking

Figure 4.1 shows the system states of the simulation model, along with the corresponding reference values. Figure 4.1 (b) shows that, despite the nature of the reference, the MPC action results in a smooth lateral displacement profile with minimal overshoot. More interesting to note is the timing of the lane change manoeuvre; for the first two seconds of the simulation, the controller takes no action, as it is not aware of the change in reference profile. At $t = 3$ s, the Y reference at prediction step N becomes $\hat{Y}$, and so the controller begins to take a small control action to minimise the error. As time progresses, more of the lateral reference is equal to $\hat{Y}$ so the controller uses a larger steer angle to minimise this error. The resulting control profile can be observed in Figure 4.2 (a).

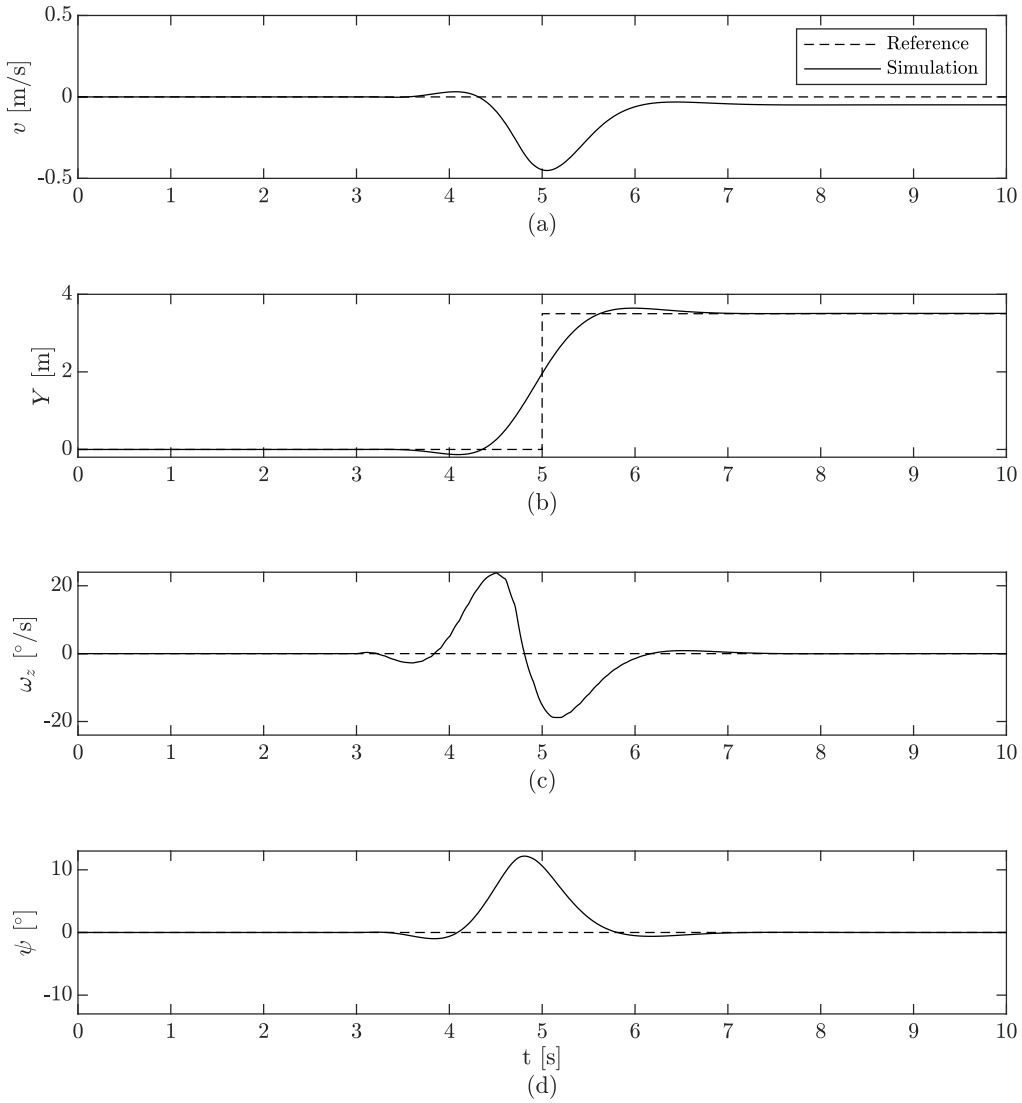
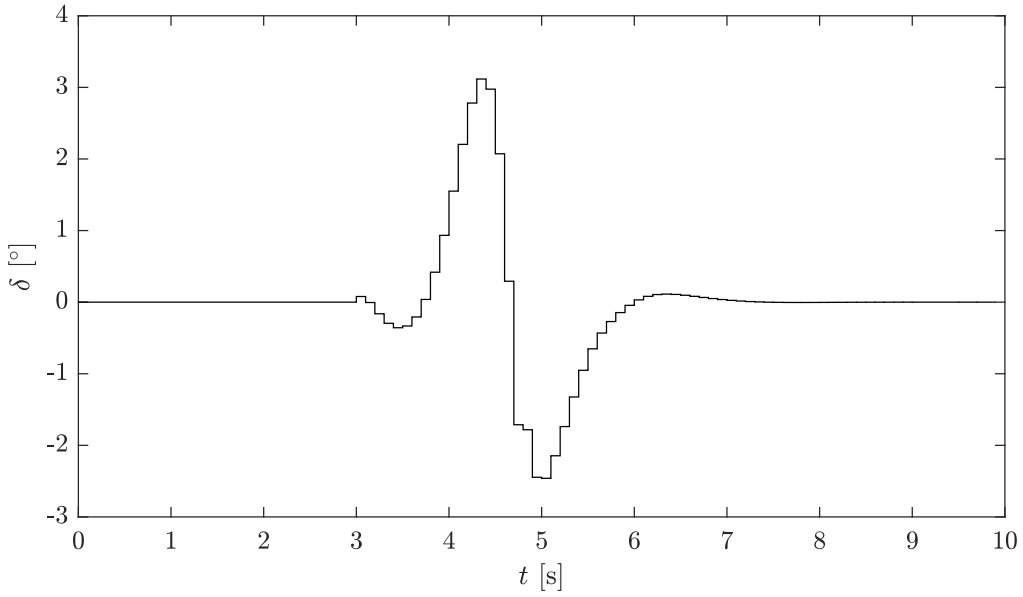
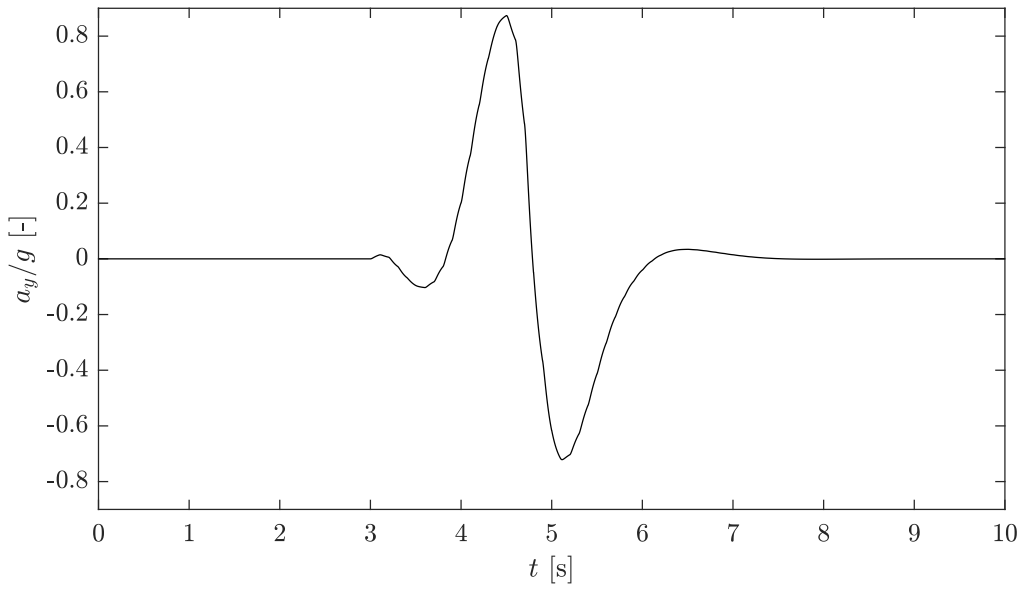


Figure 4.1: Vehicle state response to tracking a step change in lateral displacement. The reference trajectories for the other states are zero. The only stability consideration is ensuring a sufficiently long preview horizon ($N = 20$)



(a) The corresponding control input profile for an AFS controller



(b) Lateral acceleration of simulation vehicle

Figure 4.2: Control input and lateral acceleration of the vehicle with unconstrained MPC formulation. Without slew rate constraints on the control input, or state constraints, the MPC controller takes aggressive action to track the reference as closely as possible. The resulting lateral acceleration profile is not suitable for autonomous driving applications

As a result of using predicted future behaviour, rather than observed past behaviour, the controller begins to act before the step change, meaning that the vehicle has less response lag when tracking the reference trajectory.

Figure 4.2 (b) shows the lateral acceleration of the simulation vehicle. The peak value of over 0.8 g is exceptionally high for passenger vehicles and might not actually be realisable for physical systems. Even if so, it would be an unacceptable response for passengers. This highlights the need to properly constrain an MPC problem, ensuring that not only is the control action realisable for the system, but results in a vehicle response that is acceptable to its occupants.

Including lateral acceleration constraints

To demonstrate that the MPC controller can achieve these goals of ensuring a realisable and acceptable system response, simulation results for the same scenario with updated state and control constraints are presented. The first is to add slew-rate constraints, so that the maximum change in steer angle between time-steps is limited, resulting in a control sequence that can be readily applied by the control actuators. It also places an implicit constraint on $\dot{v}$, limiting any excessive change in lateral acceleration.

Rather than an implicit constraint on the acceleration, it would be preferable to be able to place an explicit constraint on a_y. The lateral acceleration is given by:

$$a_y = \dot{v} + u\omega_z. \tag{4.8}$$

In the MPC scheme, u is assumed to be constant and is therefore a system parameter. The yaw rate ω_z is a system state and so readily constrained. The transient, body fixed lateral acceleration, $\dot{v}$, is slightly more difficult to handle as the system rate vector is not considered in the discrete MPC scheme.

As obvious approximation is to use a finite-difference method for $\dot{v}$. For example, the first order, backward finite difference formula results in:

$$a_{y,i} \approx v_i - v_{i-1} + u\omega_z, \qquad \forall i = 1, \ldots, N \tag{4.9}$$

and allows for the measured lateral velocity at $i = 0$ to be used. This however results in a coupling of states between time steps that must be accounted for in the optimisation problem formulation. This formulation is addressed in Appendix C. In a more general form, by first denoting the first order approximation of the rate vector $\dot{\boldsymbol{x}}$ as $\boldsymbol{\Delta x}$, we can express the combined state-rate constraint as:

$$\mathbf{D}\boldsymbol{\Delta x}_i + \mathbf{E}\boldsymbol{x}_i \leq \boldsymbol{c}_i \tag{4.10}$$

where $\mathbf{D} \in \mathbb{R}^{p\times n}$, $\mathbf{E} \in \mathbb{R}^{p\times n}$, $\boldsymbol{c} \in \mathbb{R}^p$ and p is the number of constraints, $\boldsymbol{\Delta x}_i = \boldsymbol{x}_i - \boldsymbol{x}_{i-1}$. This constraint can include pure transient constraints, obtained by setting the elements of the relevant row(s) of $\mathbf{E}$ to zero.

Figure 4.3 shows the resulting system states when the lateral acceleration constraint is applied, with an upper and lower bound of 0.3 g applied to the system. Figure 4.3 (b) shows that the control still tracks the step change accurately, and that stability is maintained. The more obvious change in system behaviour can be observed in the yaw rate shown in Figure 4.3 (c) and the resulting yaw angle in Figure 4.3 (d). To limit the lateral acceleration, the controller saturates the yaw rate at 8 °/s for a longer period of time, compared to in the unconstrained case which peaked at 23 °/s for a shorter duration.

Figure 4.4 shows the control sequence and resulting lateral acceleration during the step change manoeuvre. At the lower limit, there is a small constraint violation of 0.5% of the acceleration bound. This can be attributed to either the solution tolerance of the MPC's underlying optimisation routine, or a difference in predicted and simulated lateral acceleration which could arise from the magic tyre formula, or tyre transience present in the simulation model. Nevertheless, this small violation is acceptable for the control application, and demonstrates that linear MPC is effective at constraining one of the more critical system behaviours.

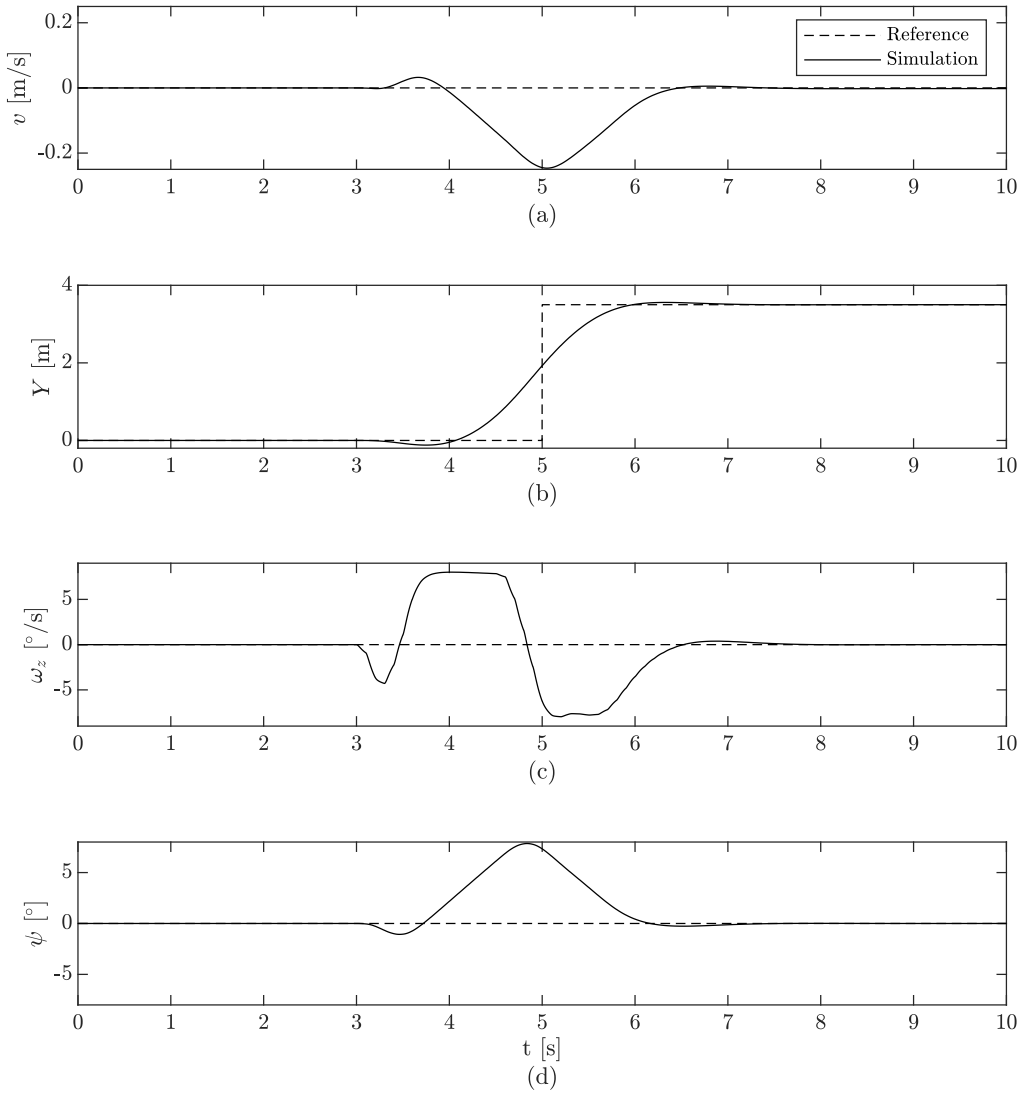
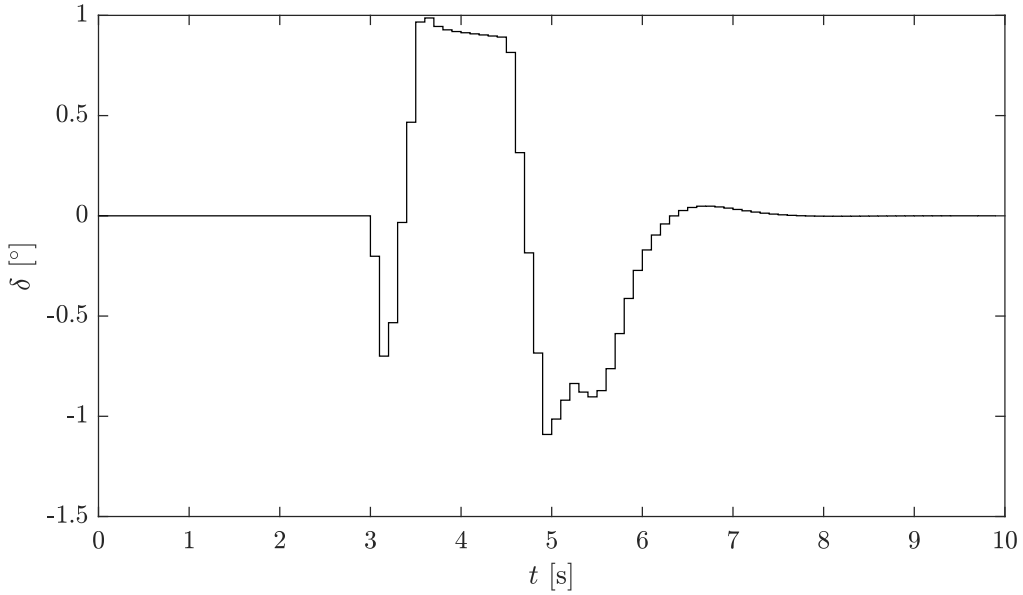
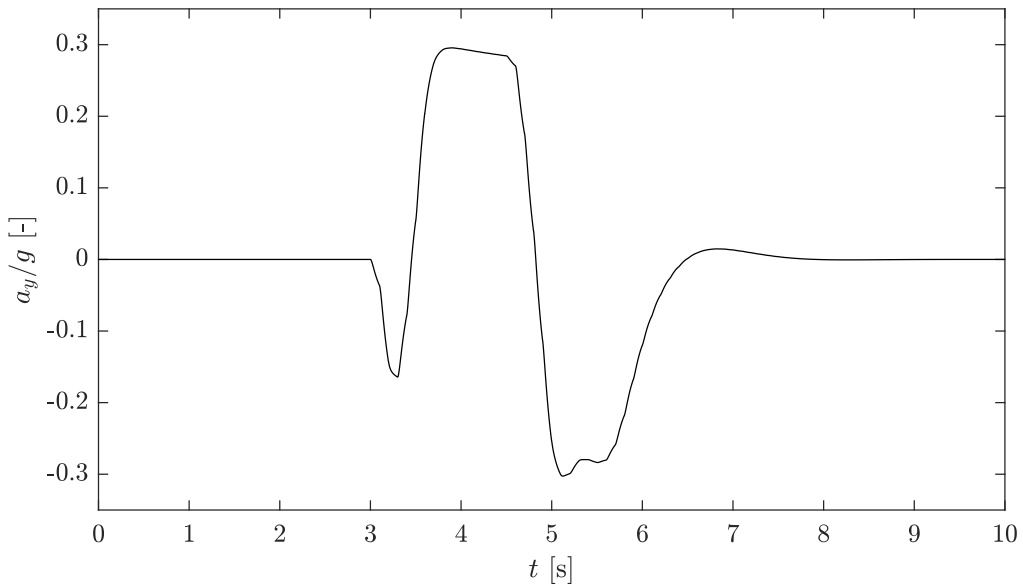


Figure 4.3: Vehicle state response to tracking a step change in lateral displacement with acceleration constraint. The system output, particularly the yaw rate, changes significantly in the presence of the state constraint imposed to limit lateral acceleration



(a) The corresponding control input profile for an AFS controller



(b) Lateral acceleration of simulation vehicle

Figure 4.4: The linear lateral acceleration constraint successfully constrains the simulation vehicle's lateral acceleration

Terminal state penalty matrix as a stabilising feature

It has already been shown in the results previously presented that for a sufficiently long preview horizon N confirms the expected trend of helping to stabilise the AFS controller. As discussed earlier in Section 4.2.1, this is not guaranteed and so often MPC formulations often include explicit features to improve or provide system stability.

The use of a stabilising state penalty matrix is presented as a simple example. The MPC controller has bound and slew-rate constraints on the front steer angle, so that the 'standard' formulation of MPC is unstable when $N = 10$, which equates to a time horizon T_p of 1 s. Figure 4.5 shows the system trajectory when $\mathbf{Q}_f$ is chosen to be the solution to the DARE. Even for what is a short preview horizon, it can be seen that the use of the terminal penalty weight allows for the controller to maintain stability throughout the manoeuvre.

In practice, it would not be advisable to use such a short preview horizon due to the presence of a terminal state penalty matrix and the purpose of the above example was merely to highlight the effect of $\mathbf{Q}_f$. As $\mathbf{Q}_f$ can be calculated offline during the design-phase of the controller, it is a highly valuable asset in the MPC formulation, particularly as it has minimal to no impact on optimisation complexity depending on the solver implementation.

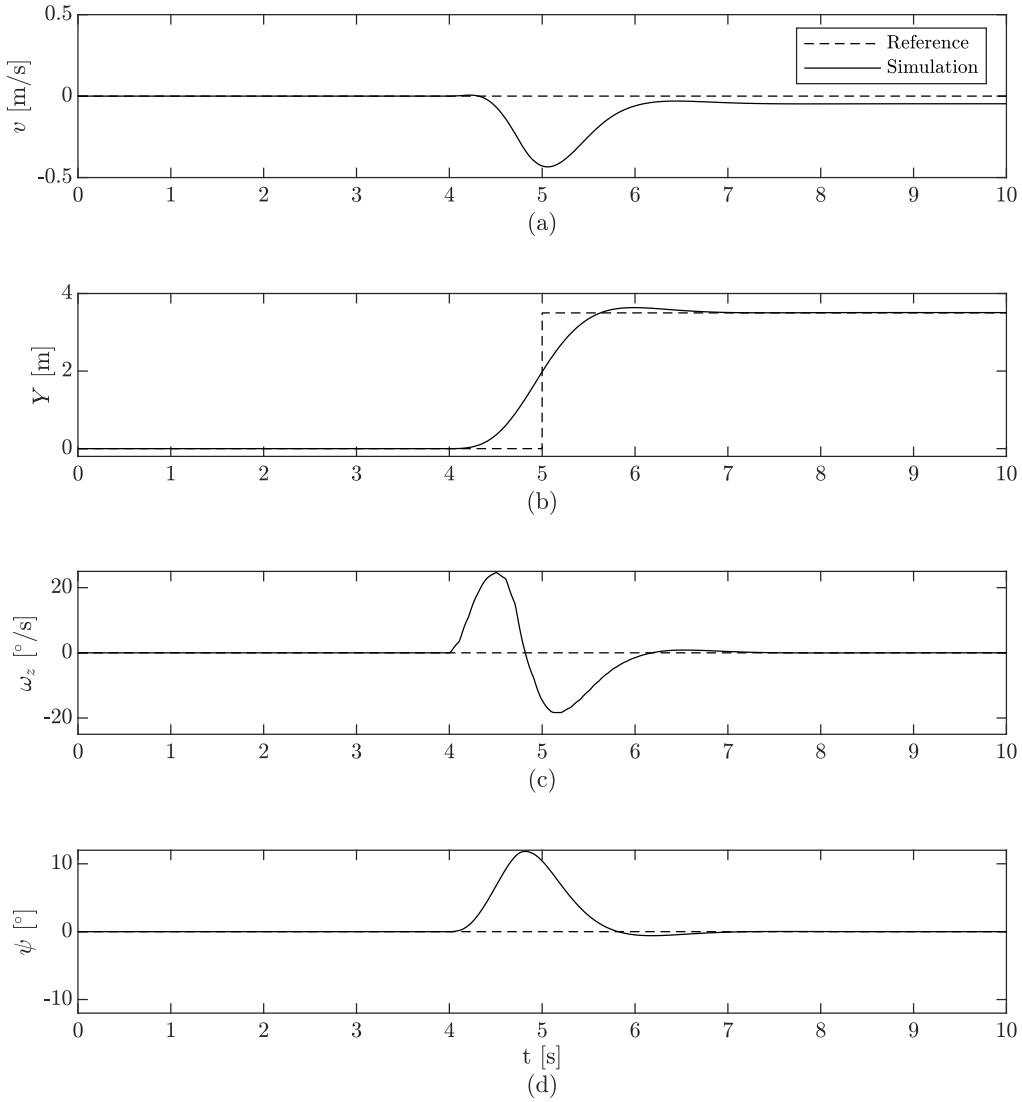


Figure 4.5: Vehicle state response with terminal penalty weight $\mathbf{Q}_f$ and preview horizon $N = 10$. The system maintains stability and has good reference tracking of the lateral step change

MPC for All-Wheel Steering

Along with the anticipated safety and efficiency benefits of successful autonomous driving, it also allows for the use of non-traditional control designs. An example of this is independent All-Wheel Steering (AWS). Although there are a few rare implementations of AWS in production vehicles, the steering angle of the rear axle is determined by the steering wheel angle chosen by the driver. This constrains the control authority of the steered axle, as a consistent system response must be provided, so that it meets the drivers expectations. By contrast, in a fully autonomous vehicle with independent steering, there is no need for there to be a consistent mapping between the front and rear control actions. Instead, the control can be chosen to optimise the system states, to provide a better passenger experience and possibility of further improving vehicle efficiency and safety.

A brief demonstration of MPCs ability to utilise additional control actuators. To do so, the vehicle models control matrix is updated to include the effects of the steered rear axle:

$$\mathbf{B} = \begin{pmatrix} \frac{C_\mathrm{f}}{m} & \frac{C_\mathrm{r}}{m} \\ \frac{aC_\mathrm{f}}{I} & -\frac{bC_\mathrm{r}}{I} \end{pmatrix}$$

with the control vector updated to include the rear steer angle, $\boldsymbol{u}^T = (\delta_f, \delta_r)$.

The steer change manoeuvre is repeated with this updated model, while the control penalty matrix updated to $\mathbf{R} = \mathbf{I}_2$, the two-by-two identity matrix, chosen so that the controller utilises both steering angles equally. The state penalty matrix is updated to

$$\mathbf{Q} = \begin{bmatrix} 5 & & & \\ & 5 & & \\ & & 250 & \\ & & & 50 \end{bmatrix}$$

to better prioritise the system states now that there are two control actuators. The yaw angle penalty is increased, to encourage the controller to exhibit crabbing behaviour to change lane, decreasing lateral acceleration.

Figure 4.6 shows the system states when tracking the step change ref-

erence. It can be seen in (d) that the peak yaw angle decreases from 5 to 12 degrees. In (b) it can be seen that it takes longer to complete the lane change, as the vehicle has to rely more on the lateral acceleration v, rather than the change in heading.

Figure 4.7 shows the control sequence, of the front and rear steer angle, and the resulting lateral acceleration. Even though the MPC controller has the same lateral acceleration constraint, in this scenario they are redundant, as the decreased vehicle yaw rate results in lower lateral acceleration manoeuvre.

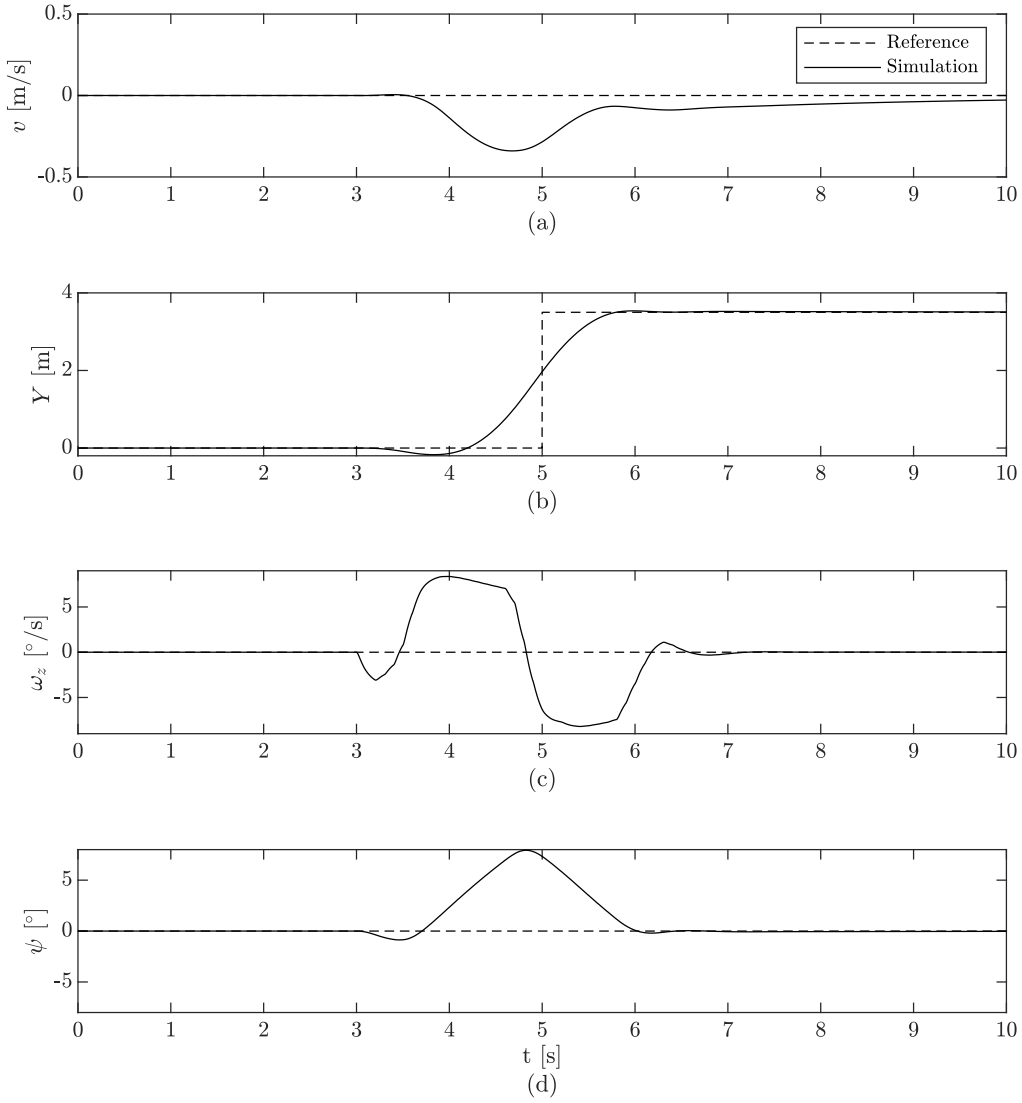
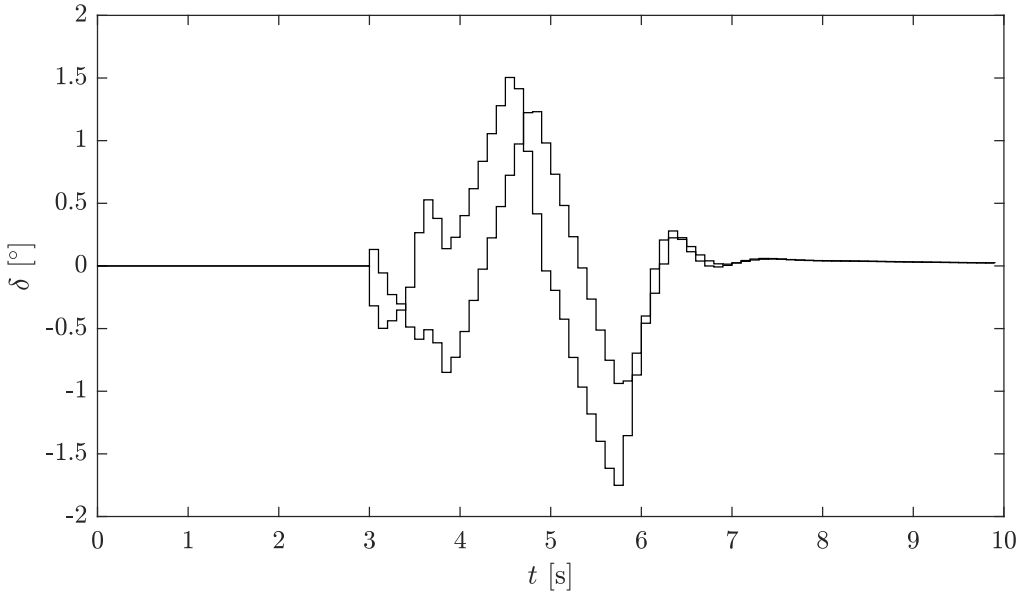
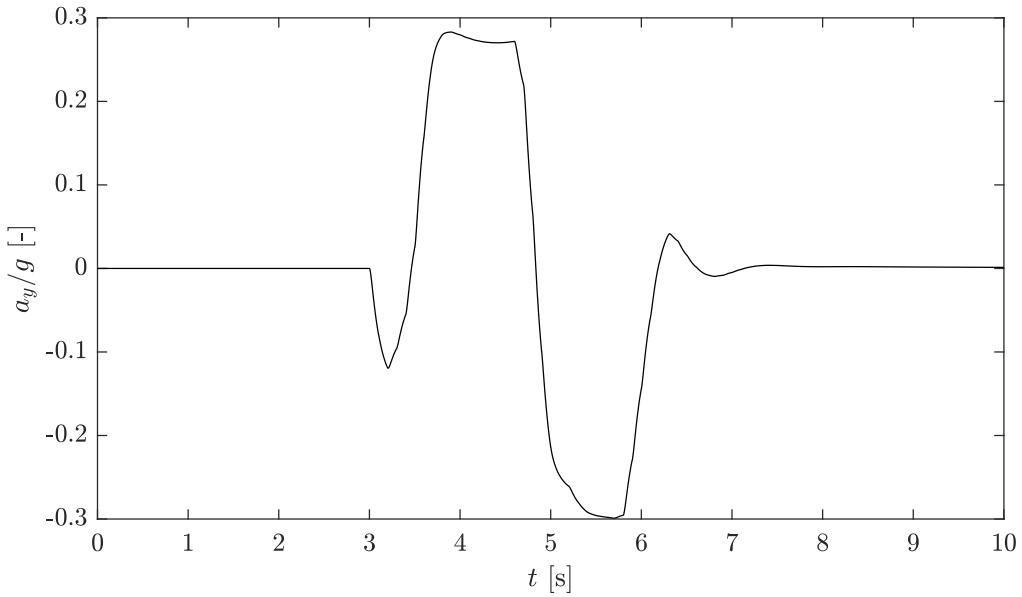


Figure 4.6: AWS vehicle respose to tracking a step function.



(a) The control sequence for the front and rear steer angles



(b) Lateral acceleration of simulation vehicle

Figure 4.7: The linear lateral acceleration constraint successfully constrains the vehicle's lateral acceleration

Smooth reference tracking

So far, the AFS controller has been tracking a step-change in lateral displacement. Clearly, this sudden change in reference is not physically realisable, nor is it desirable in practice. This motivates the development of a more gradual and smoother displacement trajectory, and the studying of the effect it has on the vehicles dynamic response.

A smooth, analytically defined reference also allows for the derivation of non-zero reference trajectories for the remaining system states. The full derivation of these can be found in Appendix B, but it is based on the following assumptions:

- The lateral displacement, Y, is the primary reference and the other variables can be derived from it.
- The lateral displacement step takes the form of a smooth, exponential, sigmoid function.
- The yaw angle is equal to the tangent angle of the lateral displacement's gradient dY/dX and the X-axis.
- The yaw rate is the time based derivative of the yaw angle.
- The lateral velocity in the body fixed frame v can be derived from the linearisation of the local-to-global coordinate transform.

This set of reference trajectories is used to demonstrate the effect that tracking has on both creating a system response that is desirable, and the effect that they can have on system stability.

Figure 4.8 again shows the system response of the simulation model, along with their respective reference values. The midpoint of the lateral displacement values occurs at five seconds to match the location of the previous step-change event. The MPC constraints used match the first scenario shown, i.e., on the control bounds only.

Figure 4.8 (b) shows that the controller can track the lateral displacement reference (which has the largest weight upon it) very closely. The vehicle also demonstrates good yaw angle and yaw rate tracking. In Figure 4.8 (a), the controller is not able to track the lateral velocity reference

accurately, but it does not have a material effect on overall performance. Figure 4.9 shows the secondary benefit of a smooth state reference. As there is no sudden change in error, the control sequence shown in Figure 4.9 (a) gradually adapts to minimise tracking errors. Figure 4.9 (b) shows that the lateral acceleration is much lower than the step change results. The maximum values still exceed the lateral acceleration constraint that was previously imposed. This highlights that while reference trajectories change behaviour, constraints are always an important part of the design problem.

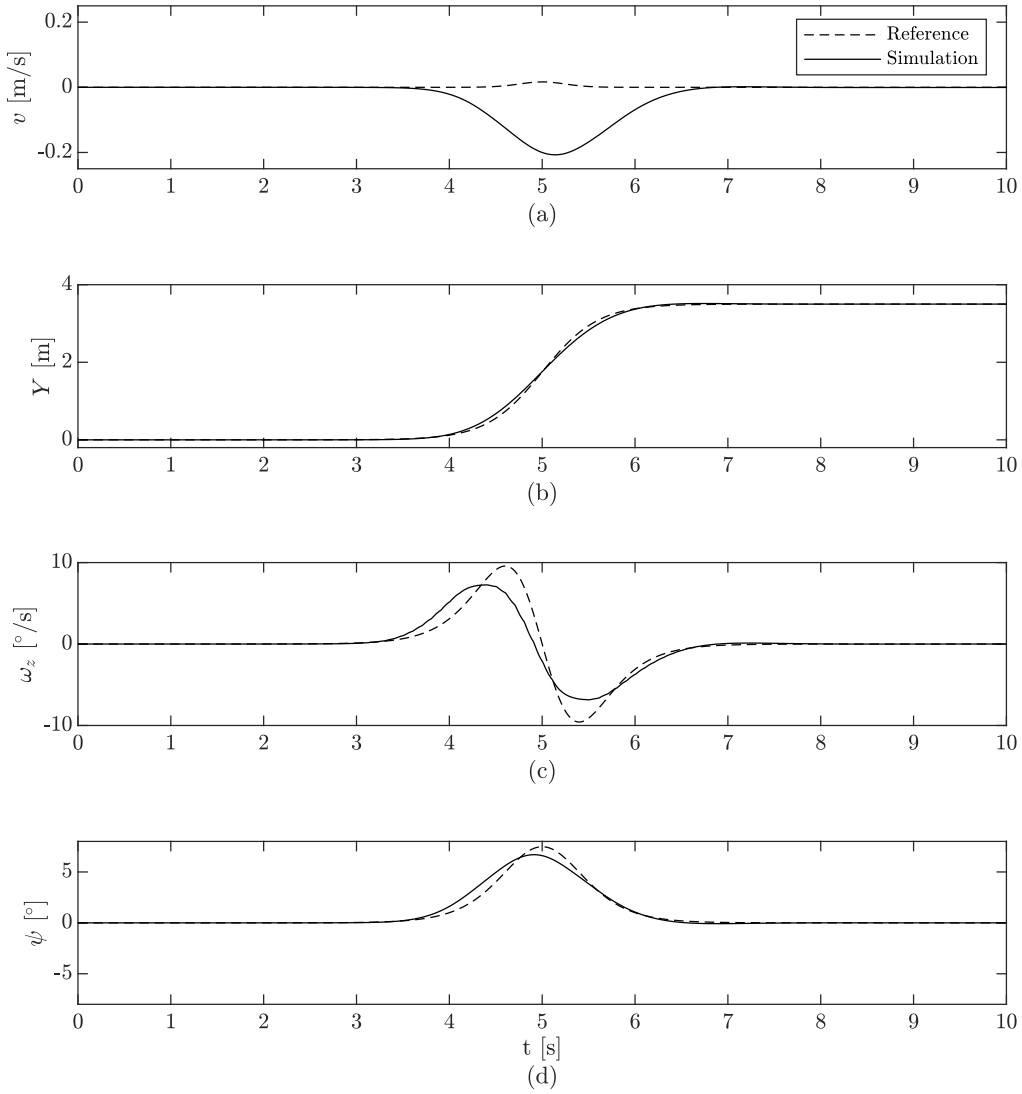
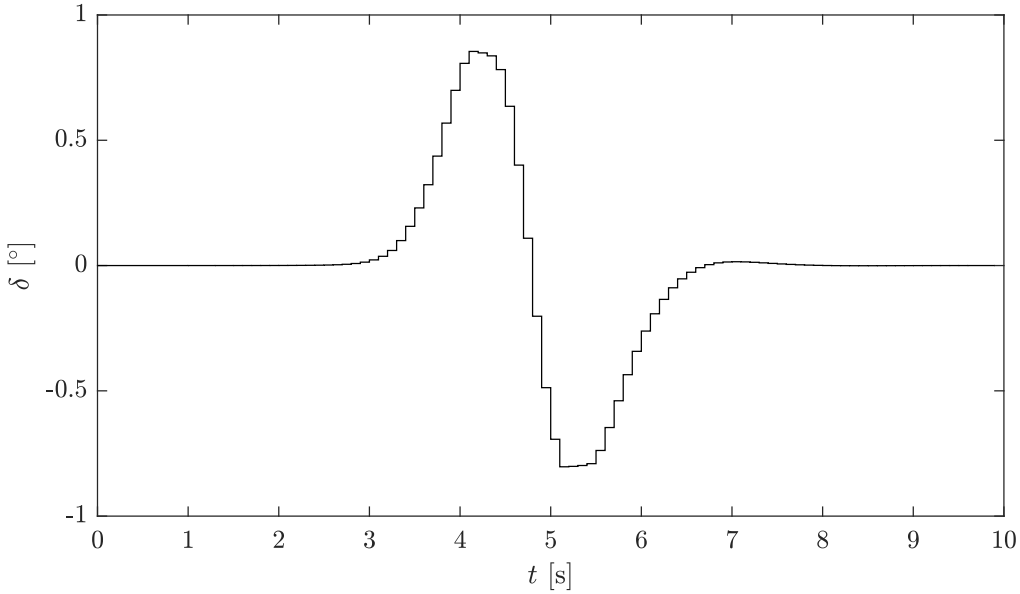
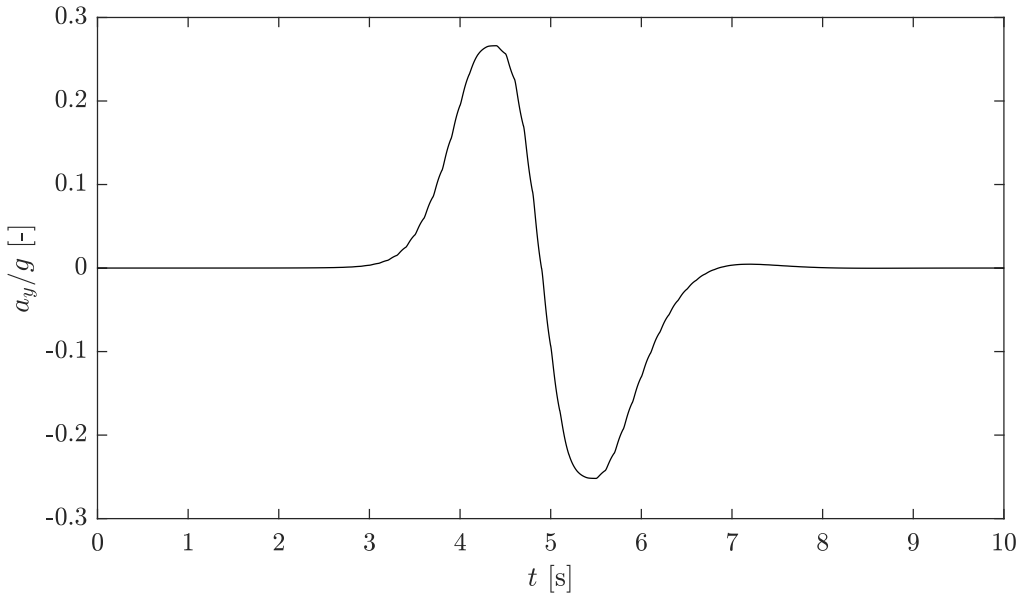


Figure 4.8: Vehicle state response to tracking a smooth step function. The only stability consideration is a sufficiently long preview horizon ($N = 20$)



(a) The corresponding control input profile for an AFS controller



(b) Lateral acceleration of simulation vehicle

Figure 4.9: Control sequence and corresponding lateral acceleration when tracking a smooth step reference. Due to the characteristics of the reference curve, the lateral acceleration stays below 0.3 g without an acceleration constraint

The influence of reference tracking on system stability

It has been shown that the state reference trajectory can change the system output substantially, and so it follows that it also has the potential to alter the stability of the system. Fortunately, the refence trajectory can be chosen to have a positive effect on system stability, as will now be shown. Three variations on the MPC controller for the AFS application are shown. In all three simulations, the only constraint is on the control bounds. This is only to emphasise the effects of the other stabilising measures, rather than it being a preferable control implementation. The preview horizon is reduced from 20 steps to 10, inducing system instability in the nominal case so that the stabilising measures can be analysed.

For the first and second simulations, the controller tracks the step-change in lateral displacement, while the final simulation uses the smooth state references. The second simulation makes use of a stabilising terminal state penalty matrix, while the other two do not.

Figure 4.10 (a) shows the lateral displacement of the three simulation runs. For the first, with no stabilising features and tracking a step reference, the system is unstable with an unbounded lateral output. Utilising $\mathbf{Q}_f$ in the second simulation run regains system stability, but there is significant oscillatory overshoot in the system response. This is further highlighted in the lateral acceleration results shown in Figure 4.10 (c), where the lateral acceleration oscillates in and out of the nonlinear handling regime (lateral acceleration above a magnitude of 0.3 g), as it attempts to find a control sequence that can track the sudden, and imminent, change in displacement reference. This behaviour, even if ultimately stable and short-lived, is unacceptable in passenger vehicles, except in the most dire emergency situations.

In contrast, the MPC scheme using the smooth state reference has no issues when the preview is shortened from 20 steps to 10. This is due to the fact that, at every advancing call to the MPC controller, using the updated reference, there is always a feasible state and control sequence, and furthermore, the optimal sequence are close to the initial starting point of

the optimisation solver.

These results taken together show, that for any successful implementation of autonomous driving behaviour utilising an MPC framework, it must utilise effective reference trajectories. This key observation will be used as the foundation for the work presented in the next chapter, which will investigate and present methods of implementing reference trajectories that are suitable for this role.

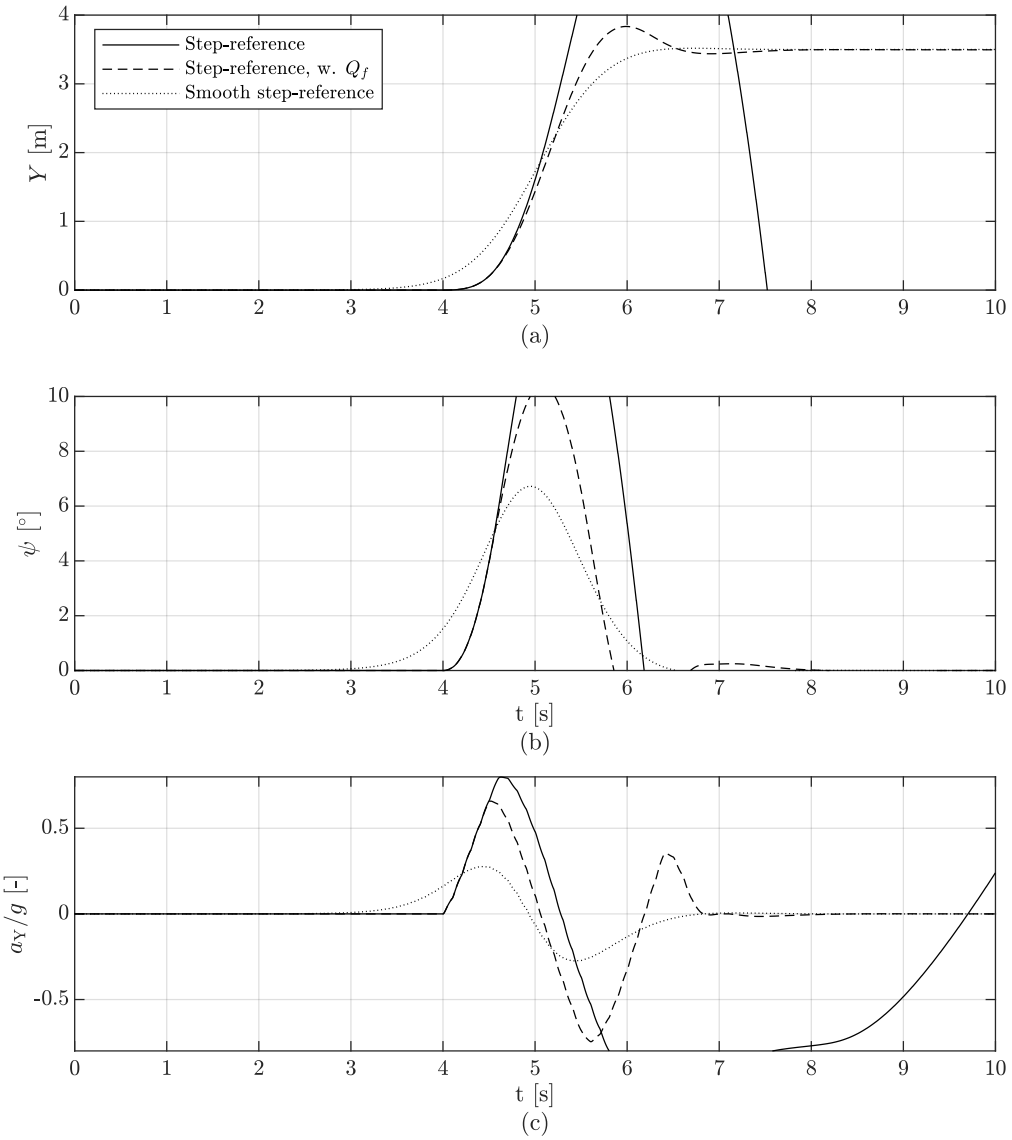


Figure 4.10: Demonstration on the effect of both explicit stabilising techniques and reference trajectory on system stability. In all cases, the preview horizon is 10. Without a stabilising terminal state penalty matrix, the vehicle is unstable when attempting to track a step change in lateral displacement and thus the system response increases overtime with unbounded error. When a stabilising state-penalty matrix is used ($\mathbf{Q}_f$), the vehicle regains stability but there is some overshoot and in tracking the displacement and the lateral acceleration in (c) can be seen to be quite large. By using a smooth reference trajectory, the vehicle can track lateral displacement change without any additional stabilising measures and the lateral acceleration is much lower than in the step-change scenarios

4.3 Conclusions

In this chapter, MPC has been shown to be an attractive control methodology due to its ability to prioritise specific system behaviours, track reference trajectories and is capable of handling constraints explicitly in the formulation. Additionally, the flexibility of MPC was demonstrated in Section 4.2.2 where the ASC formulation was changed to all-wheel steering by updating the controllers state matrix. The state penalty matrix also updated to make the best use of the new control authority and the maximum yaw angle of the vehicle was reduced by nearly 60%. To keep the focus on incorporating human factor issues, only front-wheel ASC shall be used for the remainder of this thesis.

The fact that the combined state constraints can be directly specified allows for the completion of one of the thesis objectives; to study if human factor derived constraints can be included explicitly in the control formulation. A method of including an approximation of the vehicles lateral acceleration was shown. The approximation is of sufficient quality to allow for an accurate representation of the vehicles lateral acceleration during standard acceleration manoeuvres. This allows for the study and application of human factor derived acceleration constraints using a linear MPC framework, which will be done in Chapter 6.

The effects of a standard stabilisation method on vehicle performance was studied. For example, it was shown that the terminal state penalty matrix, $\mathbf{Q}_f$, determined by the DARE, improves system stability without increasing the online computational complexity of the controller. In Section 4.2.2 it was demonstrated in the simulation for step-change reference tracking, $\mathbf{Q}_f$ could restore closed-loop stability when the preview horizon was reduced to 10 steps (1 second). Due to the benefits of this feature, it shall be used in all subsequent MPC implementations in this thesis.

It was shown that, along with the terminal stability matrix, the reference state trajectory profile has a major influence in both control quality and closed-loop stability of the vehicle. In Section 4.2.2 it was shown that

along with control specific issues, the reference trajectory provided to the controller was shown to be equally important, as using a smooth, instead of a step-change reference, also maintained system stability. The necessity of providing appropriate reference trajectories motivates finding a method to dynamically determine such a reference. This will be the focus of the next chapter.

Chapter 5

Hierarchical Model Predictive Control

5.1 Introduction

It has been shown that Model Predictive Control (MPC), with its ability to incorporate constraints explicitly into the problem formulation, is a promising control technique for autonomous driving applications. It was shown in Chapter 3, Section 3.4 that although linear approximations of coordinate tracking diverge quickly from the true value, MPC with a linear tracking model still performs well when tracking such a reference. This is due to the resampling process that takes place every time the controller is called, in which the error that accumulates through the prediction process is effectively reset. This enables the use of linear prediction models in the problem formulation, instead of resorting to more accurate, but more complex, nonlinear models to create a Nonlinear Model Predictive Control (NMPC) formulation.

Despite this inherent robustness to linearisation errors, it was seen in Chapter 4 that the reference trajectory supplied to the controller has a sizeable effect on its performance. Not only does it affect the nature of the system response, in that a smooth reference is more likely to result in a smooth system response, it influences system stability. It was shown that

using a smooth reference trajectory improved system stability, as the lack of sudden changes in the reference trajectory ensured that there was always a feasible control sequence that satisfied the system constraints.

What is needed therefore, to ensure that autonomous driving behaviour with MPC is successful, is that at all times, the reference trajectory provided to it is feasible, and the nonlinear aspects of global coordinate tracking can be handled. One such method of providing these conditions is to treat the reference trajectory and its generation, as a tightly coupled elements of the control process.

This concept, of treating reference generation and reference tracking as parts of a single whole, leads to the idea of a hierarchical control framework. Hierarchical multilayer control [119] has been used in several scenarios related to vehicle dynamics and autonomous driving functionality [80, 120]. There are several reasons for why it is an attractive control methodology for complex tasks such as path planning and active steering control, namely it allows for the logical splitting of the control task into separate, discrete modules. Additionally, the hierarchical control framework is analogous to those that exist to model human driver behaviour, such as [55, 121].

This allows for a simpler design process, as distinct aspects of route planning and control can be considered individually. Path planning and optimisation can focus on specific targets; be that calculating the most energy efficient route, minimising the total distance travelled, or maintaining a desired speed. The vehicle controller can then track the reference trajectory while focussing on vehicle specific constraints, or aspects of driving relating to human comfort.

Hierarchical control easily allows for the accommodation of dynamics with different timescales and controller frequencies [119]. The upper level controls the slower dynamics of the system in consideration; specifying a trajectory for the lower levels to follow by using a large value for the preview horizon. The lower level operates at a higher frequency to control the faster acting dynamics within the system, by using a shorter preview horizon than the upper levels; the low-level controller can calculate the

optimal control action over a shorter timescale.

The method of utilising different prediction horizon lengths aids the stability of the control problem. The high-level control has an awareness of potential events long before they come into consideration by the low-level controller. This allows for the upper level reference generator to provide a revised trajectory, which mitigates difficult scenarios, such as sudden step changes to road bounds, for the lower level to track. This is done before the lower level controller, which incorporates dynamics that are more complex, must optimise the control variables to navigate the scenario.

In this framework, the reference trajectory component sits at the top of the hierarchy and, by taking cues from the physical environment, creates a reference trajectory for the MPC controller to track. As with MPC, the trajectory generation is called again at a further time step with a resampled and updated system state. This hierarchical concept confers the same advantages of MPC to the path generation process, in that the reference trajectory is updated overtime, allowing it to adapt to changes in the vehicles environment and that as this process is performed online using current information of the system state, accumulated errors which occur due to modelling errors presented during trajectory generation are periodically removed.

For successful, stable tracking, the reference trajectory should have the following properties:

- The reference trajectory starts from the current position of the vehicle. This ensures that tracking errors do not accumulate and result in the vehicle drifting away from the reference. If this happens, then the control problem can become infeasible as the controller cannot guide the vehicle to a safe terminal region in the time up to the preview horizon.
- Any reference trajectory created spans a period of time greater than the prediction horizon of the MPC controller. MPC performance was shown to be good and the system to be stable when there was no sudden change in trajectory. Using a larger prediction horizon

for path planning helps avoid drastic changes in reference trajectory provided to the steering controller.

Therefore, the controller that is developed in this chapter, to supply the reference trajectory to the steering controller, must satisfy these conditions by its design. In the examples shown in the previous chapter, the MPC controller was tracking a predetermined, static, reference path. For generic driving scenarios this approach is not feasible, and the hierarchical controller needs to be generating a path for it to track. The path created should account for the current environment conditions, while possessing a number of properties desirable to either the vehicle passengers or the MPC problem formulation. In particular, it should be smooth as that was shown earlier to have a significant effect on the closed loop stability of the system.

In this chapter, the structure of the hierarchical framework used for the thesis is presented. Then the geometric path generation problem is given, along with how its formulation as a Quadratic Programming (QP) allows for its efficient solution in real-time, making it implementable as part of an Active Steering Control (ASC) system. Implementation issues, such as the ratio of calls between the upper and lower framework levels, are discussed. Simulation of the proposed Hierarchical Model Predictive Control (HMPC) are presented, and the effect of said implementation issues are explored.

5.2 Hierarchical Structure

It has already been shown that MPC is a promising control technique for autonomous driving manoeuvres, so long as it is provided with an appropriate reference trajectory. Therefore, the hierarchical controller must be equipped with a method of creating a valid reference path. The structure of the hierarchical controller outlined in this chapter therefore consists of two components. A lower level controller calculating an optimal front steer angle, and an upper level controller operating at a lower frequency that provides a reference trajectory for the lower level controller. This con-

troller creates a geometric path in the spatial domain, and it is shown that for constant speed situations, with driving scenarios similar to motorway driving, this geometric path can be used directly as a time-based reference for steering controller.

5.3 Geometric Path Generation

The first task is to determine a geometric path in the global (X, Y) coordinate frame. This allows for the determination of a path that avoids both static and dynamic obstacles. By considering a fixed X-grid with constant longitudinal step, the optimisation problem is to determine the appropriate Y coordinate. In doing so, the geometric path generation problem can be formulated as a convex optimisation problem allowing for solutions to be found rapidly and reliably, which is necessary for real-time applications.

For constant, high-speed scenarios, this geometric path is sufficient to use as the reference path for the MPC steering controller. This is as the lateral velocity is much smaller than the longitudinal velocity, the longitudinal velocity is a good approximation of the vehicle's constant speed. If the longitudinal spacing of optimisation problem, ΔX, is chosen to match the distance travelled by vehicle over a single controller time-step, that is:

$$\Delta X = u\Delta t,$$

then the resulting vector of lateral coordinates given as the optimal solution of the path generation problem at timestep k can be used directly as the MPC controller's reference. In doing so, the complexity of the optimisation problem is reduced as only the Y coordinates need to be determined, rather than the X and Y components.

For a non-constant speed scenario, it is possible to create a coordinate reference that depends on time, $X(k), Y(k)$. This can be achieved by calculating the arc-length of the reference path and then creating an interpolation vector of distance along the path [122].

For this optimisation problem, it is necessary for the autonomous vehicle to be equipped with a system of sensors that can detect and determine road bounds and obstacles. In this thesis, it is assumed that the simulation vehicle is equipped with such functionality and can perceive its environment without error or uncertainty.

5.3.1 HMPC Coordinate Frame

In this thesis it is assumed, without loss of generality, that the principal direction of the road path is aligned with the global X-axis. For scenarios where there is not the case, it is necessary to define a second set of ground-fixed axes $X'Y'$ [76]. The X'-axis is chosen so that it is aligned with the nominal road direction so that the vehicles lateral, Y, and angular, ψ, displacements are small. This ensures that the bicycle models tracking assumptions in Section 3.4 are valid. Once this second coordinate frame has been determined, road and object data, expressed as lateral displacement constraints, can be expressed in the $X'Y'$ frame. This frame can be updated as necessary, either at fixed time intervals, or when the vehicles displacement in the $X'Y'$ frame becomes too large for the small angle approximations to be valid.

5.3.2 Formulating the Geometric Path Generation Problem as a Quadratic Programming Problem

To determine a smooth path, the optimisation is formulated to minimise the magnitude of the paths first and second derivatives, which are approximated by finite differences. In this formulation, the optimisation problem is to choose the vector of lateral displacement values ($\mathbf{Y}^{\mathrm{r}} \in \mathbb{R}^n$), which are equally spaced along a fixed (for any single solution of the optimisation problem) vector of X values. Here, n refers to the size of the path planning optimisation problem, and not the size of the vehicle's state vector.

Informally, the problem is to choose the lateral coordinate values that minimise the total length of the path (by minimising the magnitude its gra-

dient) and minimises how quickly the path changes direction (by minimising the second derivative of the path), while respecting the road bounds. This can be achieved by choosing suitable penalty matrices that are equivalent to the squared Euclidean norm values. This then enables the geometric path generation problem to be solved as a convex QP [123].

The penalty matrices can be formed as follows, let $\mathbf{V} \in \mathbb{R}^{n-1\times n}$ be the matrix that represents the first-order finite difference coefficient matrix that is analogous to the lateral velocity of a point mass travelling along the path. $\mathbf{A} \in \mathbb{R}^{n-2\times n}$ is the second-order finite difference coefficient matrix, representing the acceleration of the point mass along the path. If a forward finite difference method is used, with first order accuracy is used, $\mathbf{V}$ and $\mathbf{A}$ can be given by:

$$\mathbf{V} = \begin{bmatrix} 1, & -1, & & & \\ & 1, & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \end{bmatrix}, \qquad \mathbf{A} = \begin{bmatrix} -1, & 2, & -1, & & \\ & -1, & 2, & -1 & \\ & & \ddots & \ddots & \ddots \\ & & & -1 & 2 & -1 \end{bmatrix} \tag{5.1}$$

The penalty matrices $\mathbf{Q}_1 = \mathbf{V}^T\mathbf{V}$ and $\mathbf{Q}_2 = \mathbf{A}^T\mathbf{A}$ can be combined as a weighted sum $\mathbf{Q} = w_1\mathbf{Q}_1 + w_2\mathbf{Q}_2$ to define the quadratic matrix in the quadratic programming problem. The positive scalar weights w_1 and w_2 are chosen appropriately to prioritise a path of shorter distance or a smoother path. Via a similar construction, it is possible to construct the jerk matrix $\mathbf{J} \in \mathbb{R}^{n-3\times n}$ and include it into the penalty matrix $\mathbf{Q}$ to minimise the rate of change of the path acceleration, but this is not considered in this thesis.

The path generation problem is constrained by upper and lower bounds on each Y coordinate. Large values of the paths first and second derivatives are discouraged by the penalty matrix in the cost function. One consideration that does have to be taken into account is ensuring that the step between the vehicle's location and the first reference coordinate is not too large. If this was not done, the first value in the lateral displacement reference trajectory could 'jump' to any point within the road boundaries. Then, it would be difficult for the steering controller to track and the stability benefit of dynamic path planning would be lost.

A simple way to do this is to introduce an additional optimisation variable at the start of the optimisation vector and then constrain it to be equal to the vehicle's current position. In doing this, the reference path is anchored to the current vehicle's location and so the generated path has a higher likelihood of being a feasible reference path compared to a static reference path. This increases the size of the optimisation problem by one ($n \rightarrow n+1$), but the change in complexity is negligible.

To improve the continuity of the path, the second decision variable (and the first that would be returned as a reference to the low-level module) is also constrained. As the yaw angle of the vehicle is approximately equal to the gradient of the path, then the following bound constraints can be derived by using a first order forward finite difference

$$\frac{dY}{dX} \approx \psi \implies \boldsymbol{Y}_{\text{lb},1} = \boldsymbol{Y}_{\text{ub},1} = \tilde{Y}_0 + u\Delta t\tilde{\psi}_0$$

where $\tilde{\psi}_0$ is the measured yaw angle of the vehicle and $\tilde{Y}_0$ is the position of the vehicle the The resulting formulation can be expressed as the following QP problem:

$$\begin{aligned}
\underset{\boldsymbol{Y}^{\text{r}}}{\text{minimise}} \quad & \boldsymbol{Y}^{\text{r}T}\mathbf{Q}\boldsymbol{Y}^{\text{r}} && && \text{(5.2a)}\\
\text{subject to} \quad & \boldsymbol{Y}_{\text{lb},1} = \tilde{Y}_0, && && \text{(5.2b)}\\
& \boldsymbol{Y}_{\text{ub},1} = \tilde{Y}_0, && && \text{(5.2c)}\\
& \boldsymbol{Y}_{\text{lb},2} = \tilde{Y}_0 + u\Delta t\tilde{\psi}_0, && && \text{(5.2d)}\\
& \boldsymbol{Y}_{\text{ub},2} = \tilde{Y}_0 + u\Delta t\tilde{\psi}_0, && && \text{(5.2e)}\\
& \boldsymbol{Y}_{\text{lb,i}} \leq \boldsymbol{Y}_i^{\text{r}} \leq \boldsymbol{Y}_{\text{lb,i}}, && i = 3,\ldots,n && \text{(5.2f)}
\end{aligned}$$

The optimal path coordinates from time-steps 2 to n can then be used as the reference trajectory for the steering controller.

Deriving Secondary Reference Variables from the Geometric Path

Once the optimal path has been determined, the secondary variables can be derived. To keep the computational complexity to a minimum, these are

derived solely from the lateral reference path. The difference between the lateral reference as a primary reference variable and the derived secondary variables represent a separation in how and why they are created. The lateral displacement values characterise a path that exists in the global XY coordinate plane. It is desirable for it to be smooth as it results in a path that does not change direction too rapidly. This enables for the path to be tracked by a vehicle. The derived variables, however, represent the desired state of the vehicle as it travels along the path to its destination. For example, for a vehicle with both front and rear steered axles at high-speed it would be desirable for the yaw-rate of the vehicle to be very low and for the vehicle to 'crab' as it tracks the lateral displacement reference, as this improves stability at high speeds. In this instance, the yaw reference would be different, even though the reference Y path, which can be thought to be the desired position of the vehicle's centre of mass, would be the same. The terms are created using Eq. (B.6) in Appendix B. The gradient terms are calculated via a finite-difference approximation of the reference trajectory determined by the path generation problem.

Call Frequency of the Path Generation Problem

In a hierarchical controller, it is possible for the individual modules of the system to have differing operating frequencies. This allows for the individual aspects of the control task to be addressed at a rate that is most appropriate. Where parts of the system have fast dynamics [119], they can be controlled by a controller that has an operating frequency high enough to control the transient behaviour of the system. Here, fast dynamics refers to system state variables in the inner loop whose transient dynamics have fallen close to zero by the time the outer loop controller has been called again. For components of the system that do not change as rapidly, a slower controller can be used. This reduces redundant computation of unnecessary or unused control action.

The path generation component of the hierarchical controller generates a reference path from the current position of the vehicle and responds to

changes in the system environment, specifically in changes to the road bounds that constrain the feasible search space for the reference path. If the lower-level controller achieves good reference tracking, at least over the short term, then the primary changes to the validity of a given reference path are the bound constraints. These are less likely to change over the time interval of the lower level controller, so it is a reasonable assumption to call the higher-level controller with a lower frequency.

For comparison of the relative frequency of two controllers in a hierarchical scheme, we define a call ratio, which determines how many times a lower-level controller is called for a single call to its parent controller,

$$\rho = \frac{\Delta t_p}{\Delta t_c}$$

where Δt_p, the time-step of the parent controller and Δt_c , the time-step of the lower-level child controller, are chosen so that the call ratio ρ is an integer value. The effect that this parameter has on system performance is investigated later in the Results section.

5.4 Path Generation and Tracking for a Double Lane Change Manoeuvre

The hierarchical controller is presented and evaluated by navigating a double lane change course. The double lane change manoeuvre is a typical driving scenario, as it is representative of both overtaking manoeuvres and emergency obstacle avoidance. These situations can occur at both low speed and high speed, and over different length scales. Therefore, despite being a simple manoeuvre to characterise, the double lane manoeuvre captures a broad range of driving scenarios and is an effective tool to evaluate the performance of the HMPC design.

First, the manoeuvre is demonstrated with the path generation carried out as a single static task at the start of the run to provide a baseline which

later HMPC runs can be compared against.

Displacement Constraints in the Steering Controller

In Chapter 4, it was demonstrated that MPC was capable of including constraints on the vehicles state variables. To ensure that the vehicle remains within the roadway throughout the simulation, the lateral is constrained above and below at every time-step. This is achieved by modifying the existing constraint terms in Eq. (4.10):

$$\mathbf{D} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \qquad \mathbf{E} = \begin{bmatrix} 0 & 0 & -u & 0 \\ 0 & 0 & u & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \tag{5.3}$$

where the first two rows are the existing acceleration constraints and the last two rows are the lower and upper road bounds, which can vary at each time-step. In this thesis, it is assumed that the HMPC is provided with an accurate set of upper and lower displacement bounds at all times during the simulation.

Simulation Setup

The course is defined by an upper and lower bound on the global lateral value Y, see Figure 5.1, with the specific course dimensions shown in Table 5.1. It is not possible to determine a feasible control sequence that satisfies both the lateral displacement and acceleration constraints at 80 km/h, which was used for the results presented in Chapter 4. For this reason, and to keep the vehicle within the linear handling regime, u is reduced to 70 km/h for these simulation results. Equally, the lateral acceleration bound could be relaxed to regain feasibility. In that case, the resulting behaviour is unlikely to be desirable to the vehicles occupants, and so this option was not pursued, and the lateral acceleration constraint is kept at ± 0.3 g. The upper and lower bounds, used in the constraint vector term $\boldsymbol{c}_i$, used in Eq. (4.10), are obtained from the course bounds (see Table 5.1) by determining the longitudinal displacement X of the ve-

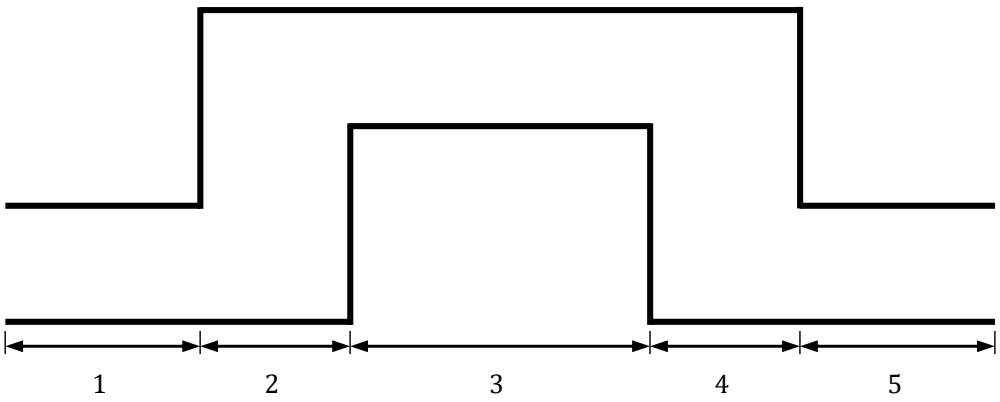


Figure 5.1: The double lane change manoeuvre

Table 5.1: Course dimensions for the double lane change manoeuvre

Section	1	2	3	4	5
Length [m]	15	40	25	25	55
Width [m]	3.5	6.5	3.5	6.5	3.5
Offset [m]	0	0	3	0	0

hicle at a given controllers discrete time-step, and using the corresponding lateral displacement value. As the path planning module does not consider the vehicle dimensions, a tighter set of bounds is used in the optimisation routine. The physical road bounds are reduced by a safety margin of 1 metre, which constrains the vehicles centre of mass to a narrower region. This safety margin was chosen conservatively to ensure that the entirety of the vehicle remains within the safe road bounds (although this was not explicitly modelled).

The preview horizon values for both controllers are shown in Table 5.2. The lower-level ASC preview horizon results in a time horizon of one second. Consequently, at the start of the simulation, the ASC module is not 'aware' of the sudden change in road bounds that are occur. As was shown earlier, the step change in reference trajectory would result in system instability, as it would not be able to take the required action in time.

The effect can therefore be seen of the path planning module, with its greater preview horizon; before the lower level of the controller is even aware of future obstacles, the path planning module is allocating a path to

Table 5.2: Preview horzion value

Controller	Parameter	Value
Path generation	n	40
Steering	N	10

avoid them.

Results

Demonstrating the effectiveness of the hierarchical control framework is done in the following stages:

- Static path: The geometric path generated by the path generation problem is assessed by creating a single reference path at the start of the simulation for the steering controller to track. This is done to demonstrate that the path generated is suitable as a reference, and that the derived secondary variables do not reduce the stability or performance of the closed-loop system.
- Dynamic path re-planning: The hierarchical system is evaluated with the path being dynamically re-planned during the simulation. This being the intended normal running of the controller, demonstrates the effect of having a dynamic path generated from the current vehicle position.
- The effect of re-planning frequency on system performance and tracking accuracy. The principle motivation for path re-planning to be undertaken in the control loop is the ability to adapt to sudden changes in the vehicle environment. How differing call ratios between the path planning and steering modules effect the tracking capability of the controller is investigated.

Along with evaluating the trajectory profiles qualitatively for the simulations runs, the tracking performance is also measured by calculating the Root Mean Square (RMS) error between the simulated vehicle states and the reference trajectories. The system state is sampled at every time-step

Table 5.3: Weighting values for path velocity and acceleration use in path generation problem

Parameter	Value
w_1	1
w_2	20

that the steering controller is called and the corresponding time-step from the most recent reference trajectory generated by the path generation problem is used.

Static Reference Path

To evaluate the viability of the geometric path as time-based reference for the constant speed steering controller; the controller is called once at the start of the simulation, with the prediction horizon chosen so that the entire simulation duration is covered.

Figure 5.2 shows the resulting reference trajectories for the lateral velocity, displacement, yaw velocity and yaw angle. In Figure 5.2 (b), which shows the lateral displacement reference, the road bounds are shows as bold black lines. The optimisation bounds, which include a safety margin to account for the vehicles width, are shown as dashed lines. Creating a path that minimises lateral velocity and acceleration, results in a smooth path, which touches the apex points of all corners present in the manoeuvre. When returning to the first lane, the path drifts below the centreline to reduce the lateral acceleration by decreasing the path curvature.

For the remaining references, it can be seen in Figure 5.2 (a) that the path generation problem successfully creates a path with minimal lateral velocity. Figure 5.2 (d) shows that the finite difference approach for determining the numerical gradient of the lateral displacement with respect to the longitudinal displacement results in a smooth yaw reference.

For the yaw rate reference, while a continuous reference is created, there are four sharp peaks that correspond to the four corner points of the optimisation bounds. This can be attributed to the fact that the derived yaw rate reference is proportional to the second derivative of the reference path

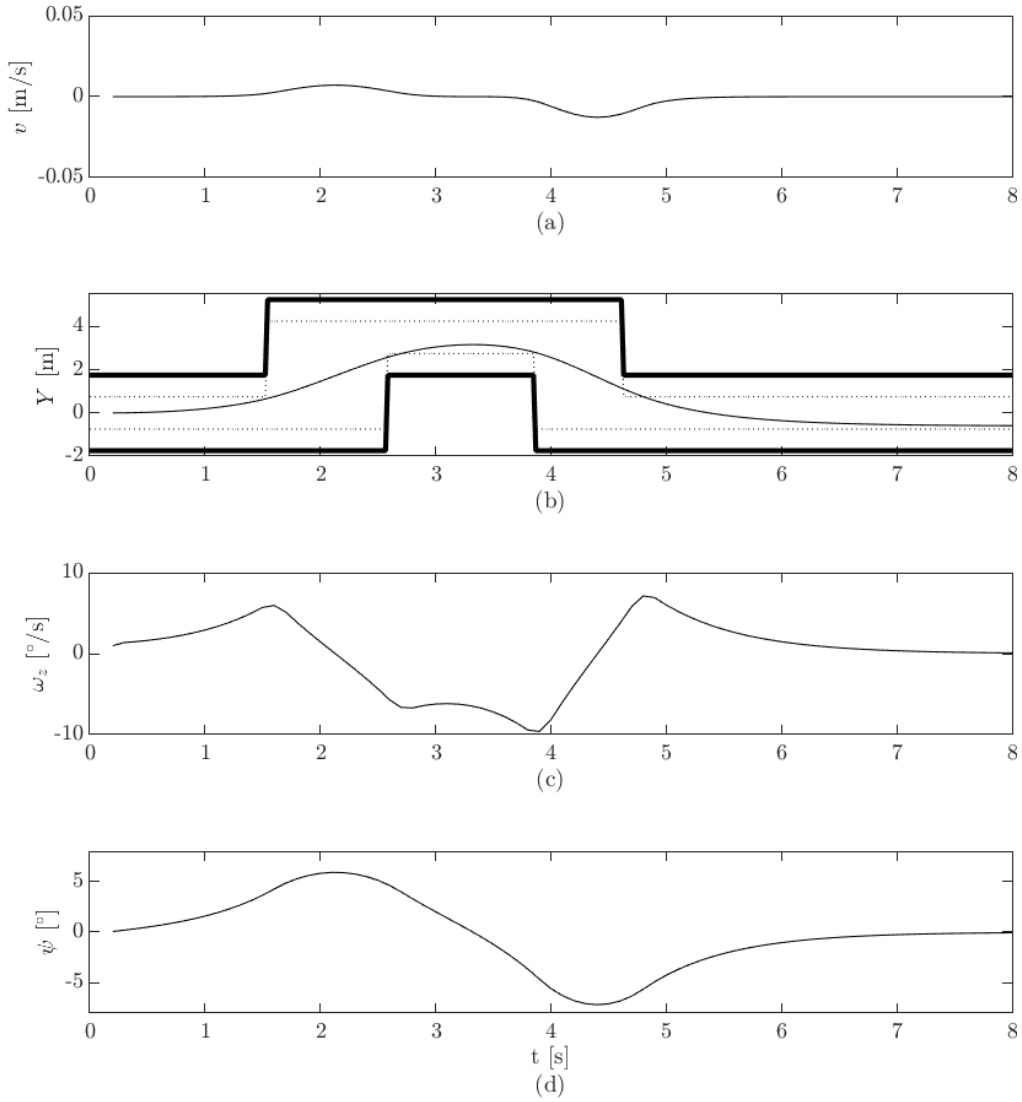


Figure 5.2: Demonstrating the effect of the Geometric path generation module. The preview horizon of the module was chosen to cover the entire course (10 seconds). A smooth reference is determined for the lateral displacement reference. In (b) the upper and lower road bounds, as functions of time, are shown as the bold lines, along with the optimisation problem bounds, which include an additional margin of 'safety' of 1 metre, shown as the dotted lines

returned by the optimisation routine. As the third derivative of the path is not minimised in the optimisation scheme, there is no effort to stop the yaw rate rapidly changing value. If this is undesirable for a given implementation, it can be accounted for by adding an appropriate penalty matrix and weighting value like those in Eq. (1.1).

Figure 5.3 shows the tracking performance or the simulation vehicle. In Figure 5.3 (a), the simulated lateral velocity of the vehicle is much higher than the reference trajectory, with a maximum value of 0.13 m/s compared to 0.01 m/s. This indicates that the method used to derive the secondary variables do not provide an accurate representation of the vehicle behaviour, resulting in poor tracking performance. The steering controller is able to track the reference trajectory closely and matches the reference in clipping the four apex points of the manoeuvre. Figures 5.3 (c) and (d) show similarly good tracking performance, with the yaw rate tracking in Figure 5.3 (c) indicating that the poor lateral velocity tracking performance is not due to the use of numerical second derivatives, but rather, that the method to derive the lateral velocity reference (in Appendix B) is the cause, and is not representative of the observed vehicle behaviour of the simulation vehicle.

In Figure 5.3 (c), the yaw velocity can be observed hunting around its reference. This has a corresponding effect on the yaw reference tracking, but the controller is still able to track the yaw reference well, as seen in 5.3 (d).

The reasons for this tracking behaviour can be seen in Figure 5.4. In (a), the control sequence determined by the steering controller is shown, with the simulation vehicles lateral acceleration and controller bounds (as dashed lines) presented in (b). The acceleration bound is saturated briefly around the five second mark, as the vehicle returns to the initial lane. The acceleration bound comes close to saturation earlier in the simulation, which suggests the reason why they yaw velocity does not have a smooth profile. This is due to in the calculation of lateral acceleration term; the yaw velocity is multiplied by the longitudinal velocity. Therefore, at higher

speeds, changing the yaw velocity is the controllers primary method of modulating the lateral acceleration. This is collaborated by the observation that if the simulation manoeuvre is performed at lower speeds, the yaw and yaw velocity tracking is greatly improved, with smooth response profiles.

Figure 5.3 shows the sequence of control inputs determined by the MPC scheme and the resulting lateral acceleration of the vehicle. The lateral acceleration saturates briefly between 5 and 6 seconds, as the vehicle returns to the inside lane. If the simulation results were dominated by the saturation of the acceleration bounds, the results would have been more focussed on the stability of the MPC steering controller, which was demonstrated in Chapter 4.

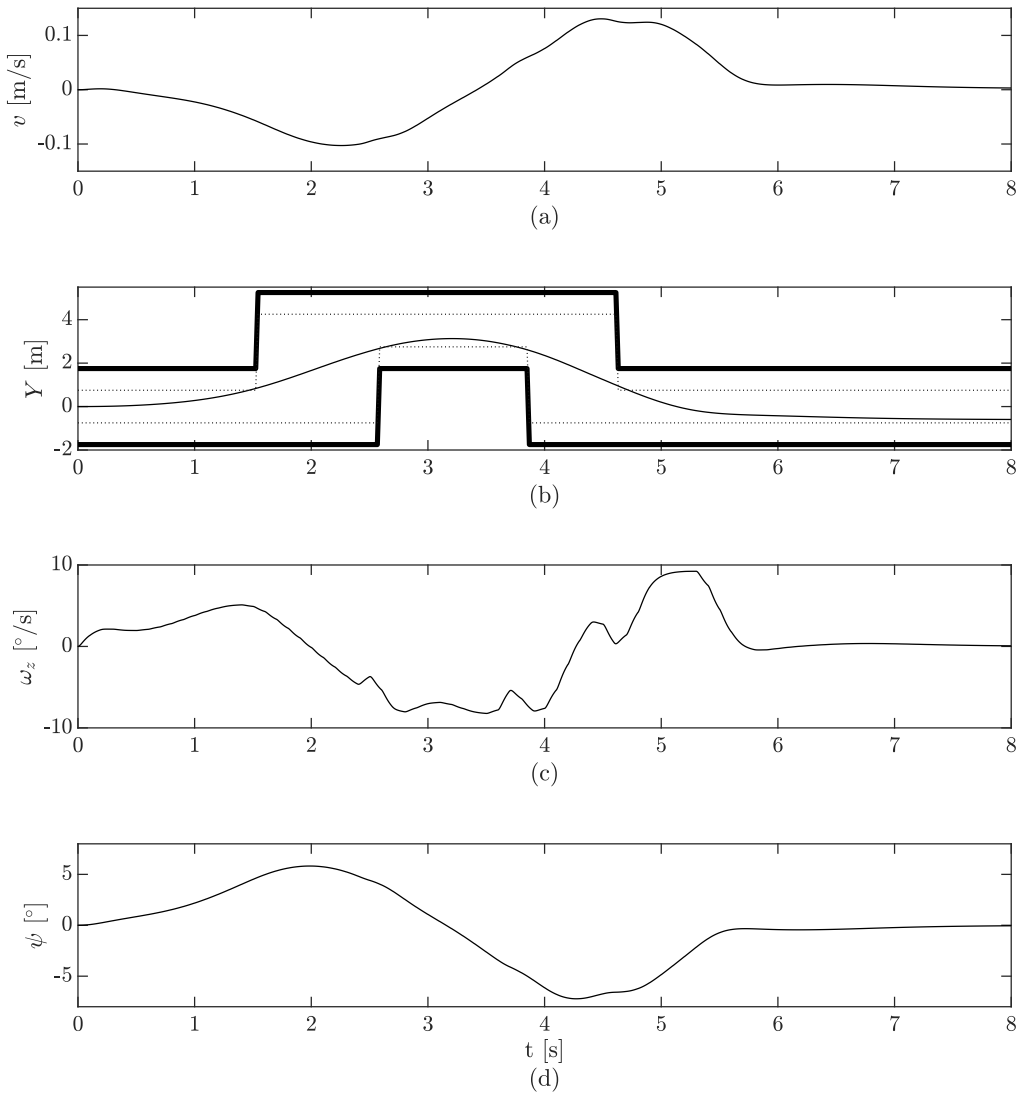
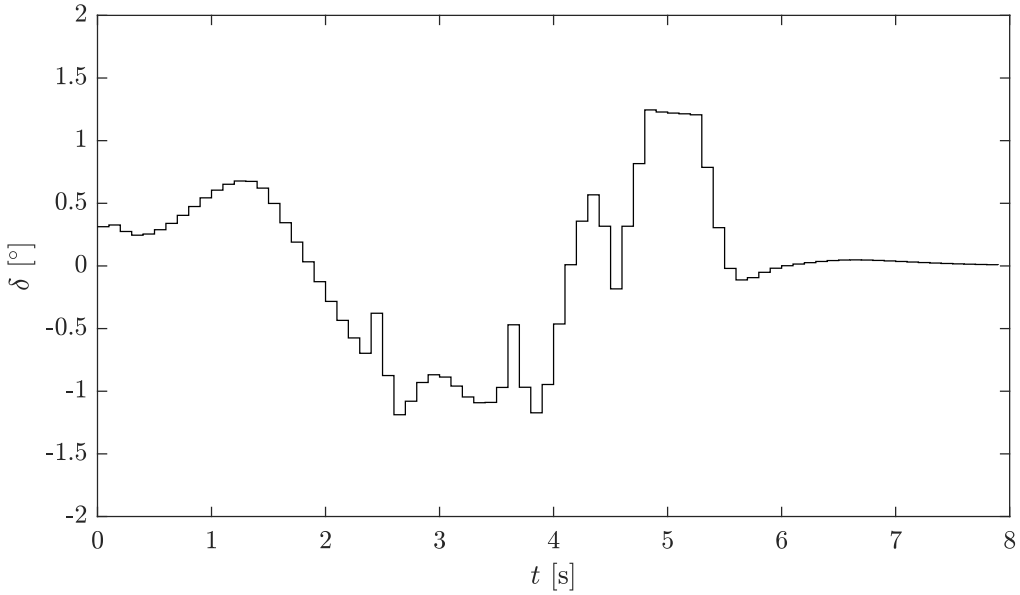
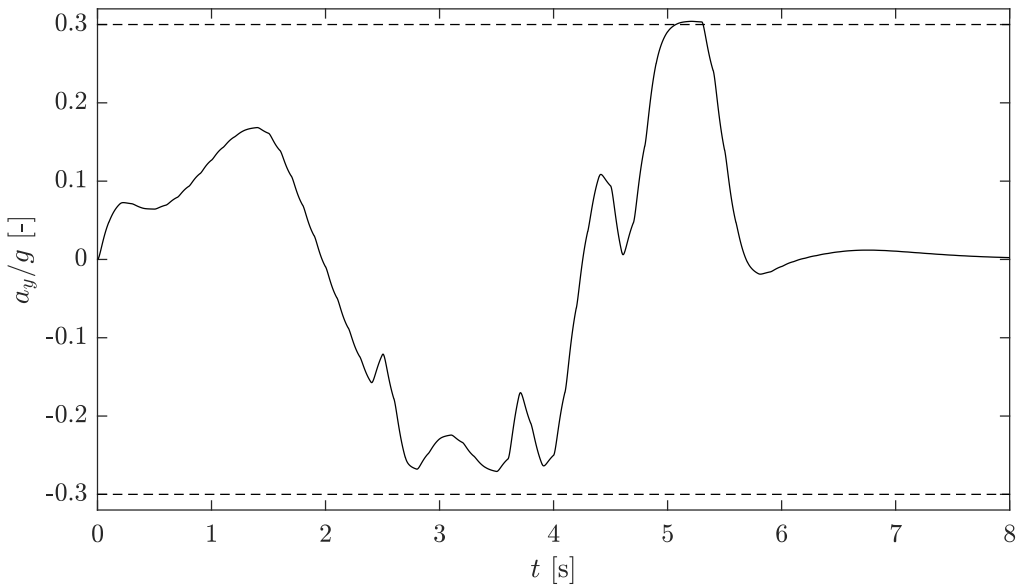


Figure 5.3: Vehicle response to tracking the reference path created by the geometric path generation module. In (b) the upper and lower road bounds, as functions of time, are shown as the bold lines, along with the optimisation problem bounds, which include an additional margin of 'safety' of 1 metre, shown as the dotted lines



(a) Steering input profile for an AFS controller



(b) Lateral acceleration of simulation vehicle

Figure 5.4: Steering input and corresponding simulation vehicle acceleration when tracking static reference

Dynamic Re-planning

Now that the viability of using the geometric path problem as a method of creating reference trajectories has been demonstrated, the effect of using it dynamically in the closed-loop system is analysed. Here, the purpose is to demonstrate that by using even a simple path planning method dynamically, so that the vehicles change in position is accounted for, results in the stabilising features of a smooth reference trajectory to be provided to the MPC controller.

In the following simulation, the steering controller, as previously, has a prediction horizon of ten steps. The geometric path planning module has a prediction horizon of 40, resulting in a look-ahead time of four seconds. The call ratio between the two modules is two. The effect of using different call ratios will be demonstrated in the next section. By using the larger horizon for the path generation module, the steering controller can track a reference that avoids an obstacle, before the steering module is even aware of it.

Figure 5.5 show the sequence of reference trajectories provided for each individual state, throughout the simulation. Figure 5.5 (b) shows the lateral displacement trajectory. This is the only value that is explicitly generated by the path planning module. It can be seen, by the fact that the trajectories generated over subsequent time-steps all aligning closely to one another, that the path planning module generates a feasible reference path that the steering module is able to track comfortably. The fact that the reference trajectories have large agreement with one another can be attributed to two reasons; as the path planner is attempting to generate a path that minimises the lateral acceleration, it is always feasible for the AFS to track. The second reason is that the reference path is constrained to start at the instantaneous location of the vehicle, therefore so long as the steering controller achieves good tracking of the lateral displacement, the new optimal path will always lie close to the previous one.

In Figure 5.5 (d), it can be seen that the derived yaw reference also align closely to one another between reference calls. This is due to the first

coordinate being constrained to align in the direction of the vehicles yaw (or heading) angle, at the time the path is called. Figure 5.5 (c) shows that the yaw rate reference does not appear to be particular consistent, particularly at the times when there is a large change in Y. An attempt to constrain the third lateral displacement variable to agree with the vehicles measured yaw-rate was not successful (as was done for yaw and lateral displacement) and so is not shown in the thesis. As the yaw rate has a much smaller associated time-constant compared to the yaw angle and displacement, it changes much more rapidly between discrete time intervals in the ASC module. This is also responsible for the oscillation observed in the results. When a smaller time step is used for the lower level controller, this oscillation becomes smaller, as the discretised controller becomes a better approximation of the yaw rate.

Figure 5.6 shows the system states when tracking the sequence of reference trajectories provided by the path generation module. It can be seen in (b) that the vehicle has a smooth lateral displacement trajectory as it navigates the course.

The yaw angle and lateral velocity trajectories are similar to the vehicle tracking the static reference trajectory shown in Figure 5.3. The same is true for the yaw rate in Figure 5.6 (c), except that it tends to zero five seconds after the start of the simulation much more slowly than the static reference tracking run. This occurs due to the path planner being able to create a steadier return to the initial lane.

Figure 5.7 shows the steering sequence and the lateral acceleration response of the vehicle. Unlike in the static reference trajectory run, the acceleration does not reach its constraint limit. This shows that by creating a reference path from the vehicles current location and heading, which tries to minimise the paths acceleration, allows for a reduction in vehicle acceleration. This occurs as for a given call to the path generation module, its optimal reference is a function of the vehicles current state whereas for the static case, the vehicle is being driven to a predetermined reference, regardless of the tracking performance.

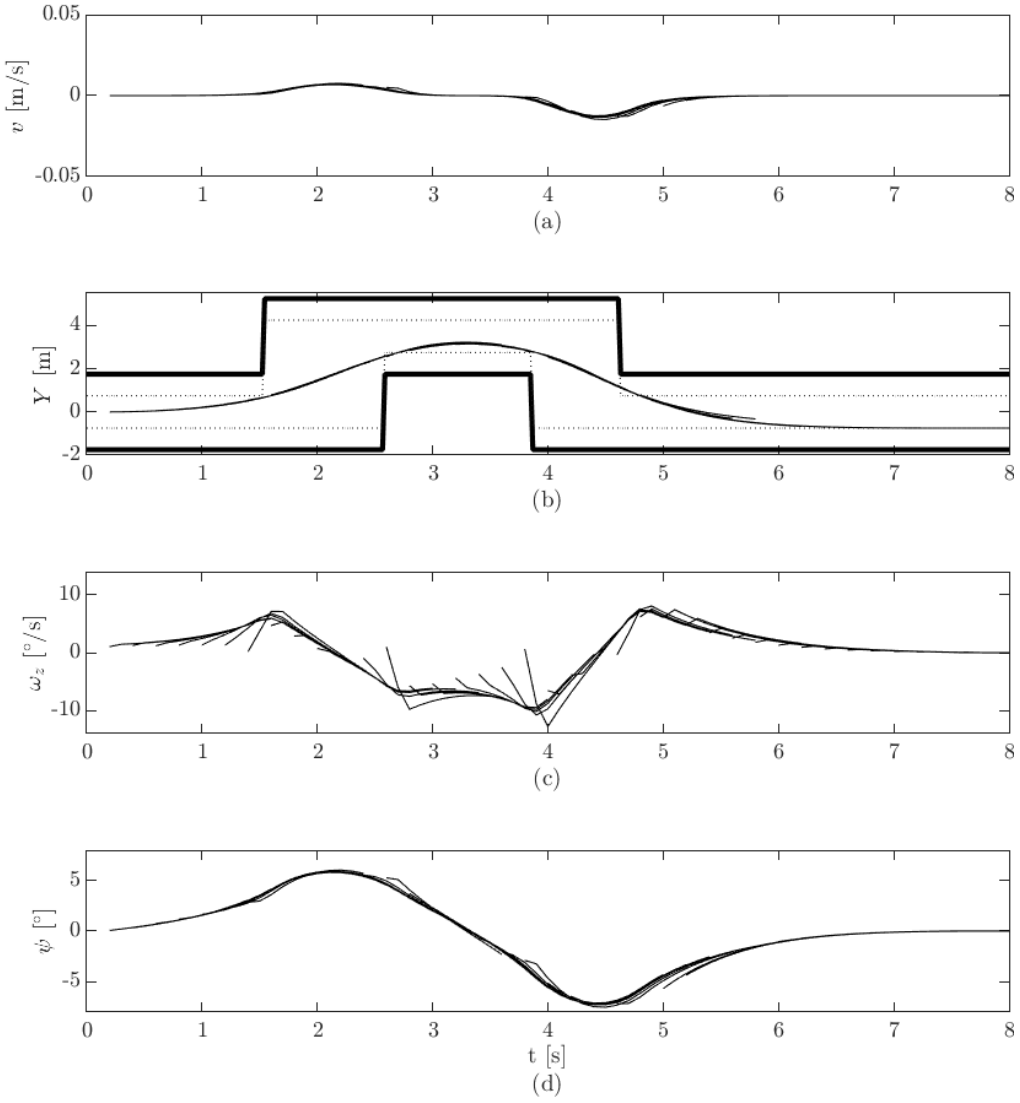


Figure 5.5: The sequence of reference trajectories generated by the path generation module. In (b) the upper and lower road bounds, as functions of time, are shown as the bold lines, along with the optimisation problem bounds, which include an additional margin of 'safety' of 1 metre, shown as the dotted lines

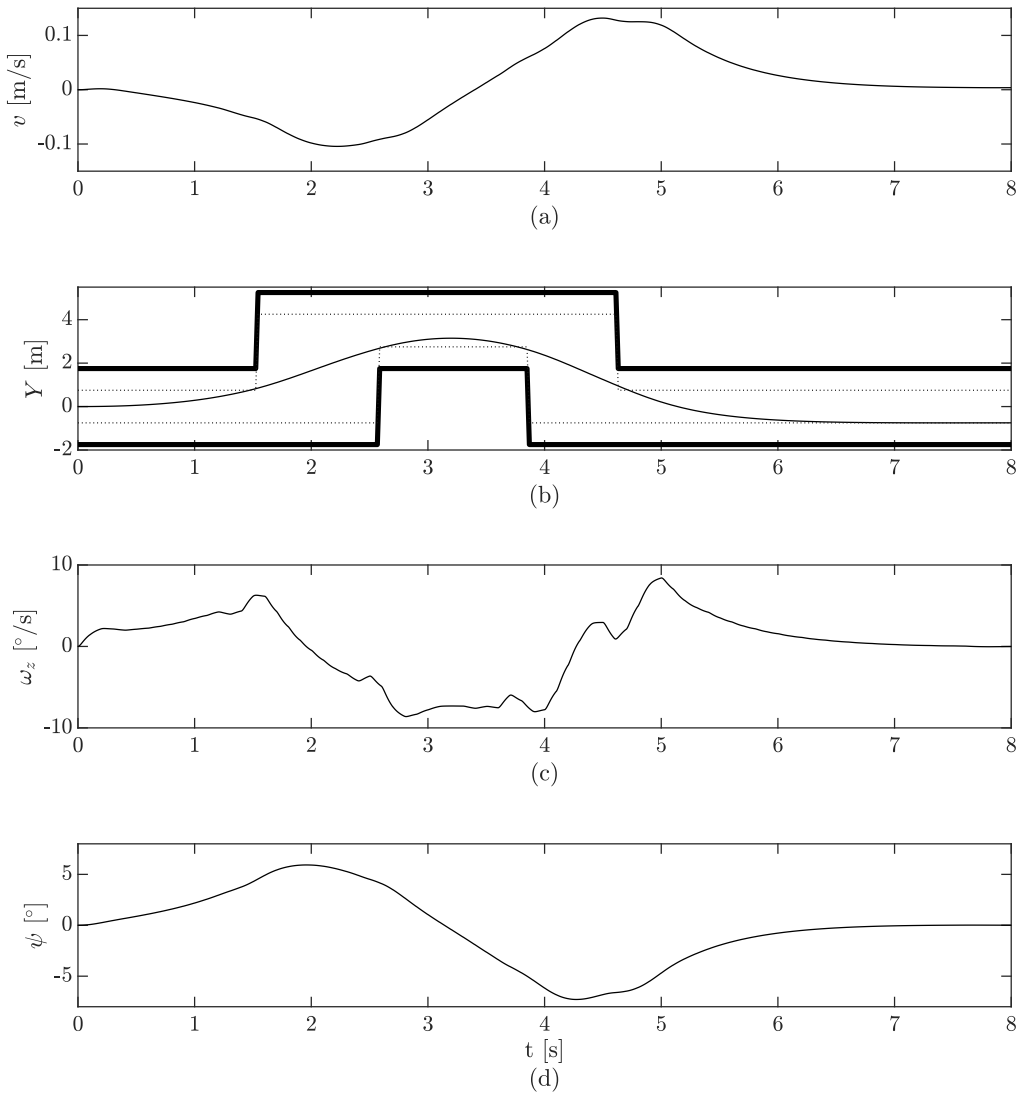
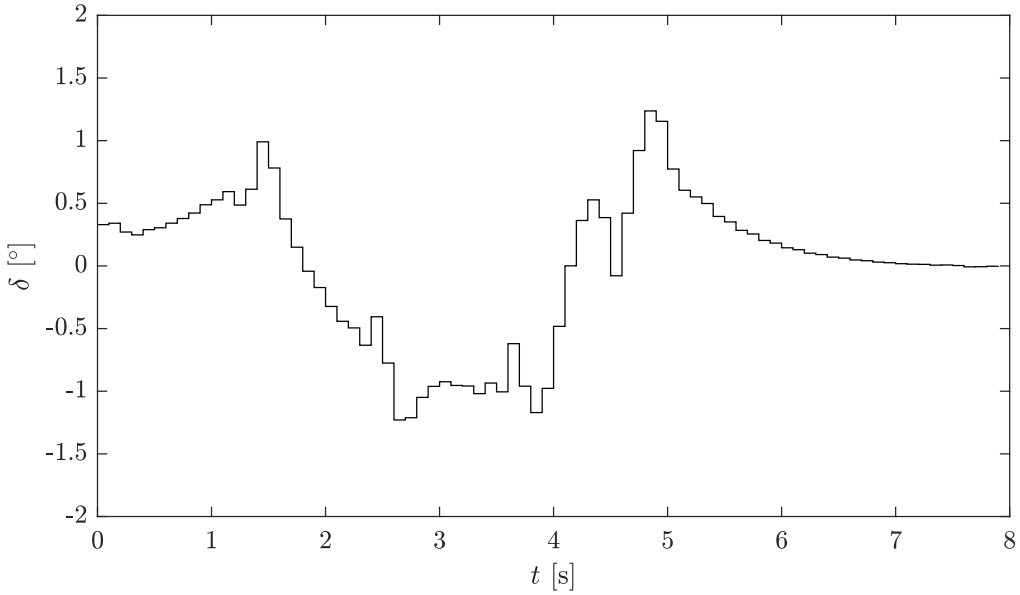
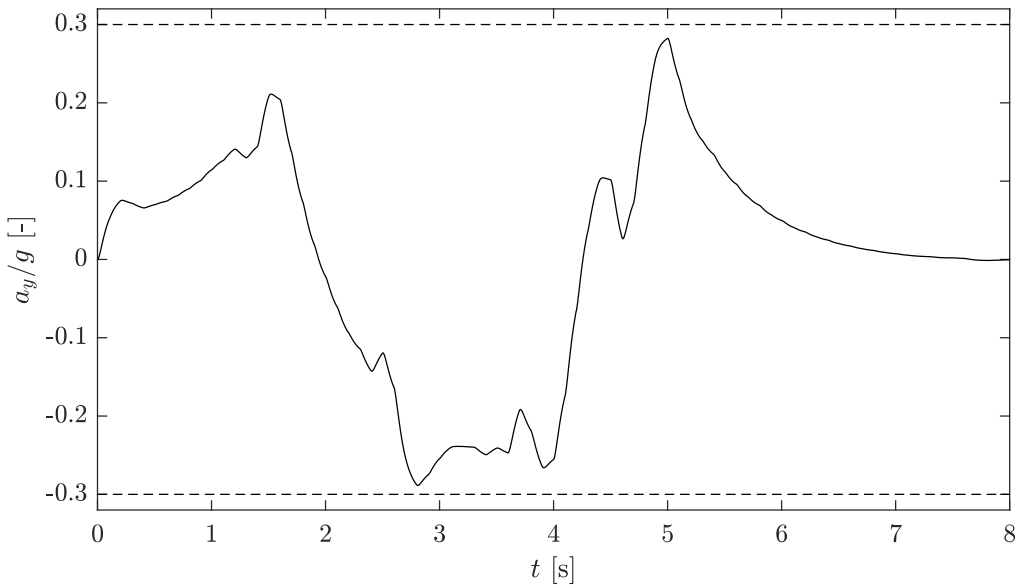


Figure 5.6: Vehicle response to tracking the dynamic reference path created by the geometric path generation module. In (b) the upper and lower road bounds, as functions of time, are shown as the bold lines, along with the optimisation problem bounds, which include an additional margin of 'safety' of 1 metre, shown as the dotted lines



(a) The corresponding control input profile for an AFS controller



(b) Lateral acceleration of simulation vehicle

Figure 5.7: Control input and lateral acceleration response when tracking dynamic path planning trajectory

The Effect of Call Frequency on Reference Tracking Performance

As the beneficial effect of using a hierarchical controller has been shown, one of the remaining questions regards the effect that the call ratio has on system performance. By recalculating the reference path and state trajectories more often, the system uses the most up-to-date environmental information and eliminates tracking errors as the path is redefined from the vehicles location.

The drawback for using a smaller call ratio is that, depending on the path planning algorithm being used, it may not be possible to update the reference path in time. Even if its possible to do so, it would be an unnecessary use of computational resource to continuously recalculate a path if it is similar to the one that it is replacing. This could occur if the vehicles physical environment does not change dramatically over the time scale of the inner loop.

To investigate the effect, the double lane change simulation was performed again, with seven different call ratios used between the MPC steering controller and the geometric path generation module. Otherwise, all system parameters are kept the same. The ratios used are shown in Table 5.4, along with the RMS tracking errors for the reference trajectories. It can be seen, as expected, that the lateral displacement tracking error decreases as the path is recalculated more frequently. This is by the design of the path generation problem, every time the planner is called, the path starts from the vehicles current position, allowing for the tracking

Table 5.4: The effect of path planning call ratio on tracking performance. The simulations are performed at 65 km/h

ρ	v [m/s]	Y [m]	ω_z [°/s]	ψ [°]
1	0.0684	0.0045	2.1166	0.3112
2	0.0682	0.0111	1.7735	0.3963
3	0.0685	0.0201	1.8117	0.4290
4	0.0686	0.0248	1.9305	0.4325
5	0.0683	0.0223	1.5457	0.3641
10	0.0674	0.0268	1.3901	0.3638
20	0.0668	0.0409	1.6117	0.3880

error to be (near) zero for larger parts of the manoeuvre, when called more frequently.

The lateral velocity RMS error stays approximately constant throughout the different simulation runs, even decreasing slightly when larger call ratios are used. This appears to be due to the fact that, as discussed in Section 5.4, the derived lateral velocity reference does not accurately represent a trackable reference for the steering controller.

More interestingly, is the yaw reference, and the subsequently derived yaw rate reference. Tracking the yaw reference, derived from the numerical gradient, was shown to have satisfactory performance in the previous result sections. This was partly attributed to the constraints placed upon the second reference lateral displacement value, to constrain the path to align with the vehicles attitude at the start of the optimal path. In the RMS tracking error results, the yaw error decreases from 0.146 degrees when the path is re-planned every timestep, to 0.076 degrees when the path is re-planned every 20 time-steps. This is due to the fact that, at low call ratios, when the path is being re-planned more frequently, the reference trajectory is more constrained due to the constraints Eqs. (4.3b) and (4.3c) being applied to the updated path more regularly. The results in the steering controller struggling to track all of the references accurately. At higher call ratios, the steering controller has more time to optimise the system states, to minimise the tracking error, as the system enters a steady state.

For all the tracking variables, the change in RMS error values are small and for all the call ratio results shown, the vehicle has good tracking performance. The results show that over a period of 0.1 to 2.0 seconds, the effect that the system call ratio has on tracking performance is not a significant factor in determining the value to use for the hierarchical system. Instead, the value should be chosen based upon the other factors outlined earlier. For the remainder of this thesis, the call ratio is kept at two, providing some reduction in the computational load but allowing for the system to react rapidly to any changes to the vehicles environment.

5.5 Conclusions

Hierarchical control has been added to the control framework to address some of the limitations of the linear MPC framework. The concept of a hierarchical framework along with its theoretical advantages have been discussed. A two-tier application has been decided on, incorporating a geometric path generation module to dynamically determine reference trajectories for the steering controller. It is shown that, for driving scenarios of sufficiently high-speed, the geometric reference path is suitable for constant speed driving scenarios.

The performance of the hierarchical controller is verified, and it is shown that the hypothetical benefits of the hierarchical framework does allow for the usage of MPC in the proposed driving scenarios. The frameworks dependency on its implementation parameters is also presented, and it was demonstrated that system performance was robust to hierarchical control parameters, such as call ratios. Simulation results demonstrated that while larger call ratios decreased state tracking performance, as measured by RMS, the decrease was not sufficient enough to be a concern. This allows for a simple implementation of the hierarchical controller, with minimal parameter tuning, before using it for the research of MPC-specific implementations.

Because of this, a flexible and generic framework has been developed, which forms a strong foundation for the study of the incorporation of human factor informed control strategy and subsequent analysis.

Chapter 6

Incorporating Human Factors

6.1 Introduction

For the successful adoption of autonomous road vehicles, the behaviour must be perceived by the vehicle occupant to be at least as safe as current vehicle designs, being controlled by a human driver. Chapters 4 and 5 have shown that Model Predictive Control (MPC), particularly when part of a hierarchical framework, can provide a control law that is suited to the dynamic environment encountered in autonomous driving scenarios. In this chapter, human factors pertaining to the autonomous driving problem are studied and incorporated into the hierarchical control framework. The effects that this has on vehicle performance is investigated, and it is shown that human factor concerns can be included explicitly into the control formulation. This then allows for dynamic vehicle behaviour that can, depending on the preference of the human drivers, while maintaining a single control implementation.

In a hierarchical framework, many aspects of the driving task play a role in the autonomous behaviour, including the path generation module and steering controller. The focus will be on the steering controller originally outlined in Chapter 5 by considering how to include human factor constraints.

When considering the lateral performance of the vehicle, the lateral ac-

celeration behaviour is critical for the human acceptance of the autonomous driving behaviour. The vehicle occupants must feel that the vehicle is always operating safely, and that the vehicle behaviour is comfortable. Because of this, the maximum desired acceleration in normal driving scenarios, is far below what is achievable by the vehicle, which is ultimately bounded by the friction between the road and the vehicles tyres.

Therefore, the autonomous vehicle controller must provide a control sequence that minimises or avoids high lateral acceleration manoeuvres. As shown previously, by using an online optimisation scheme, MPC provides methods for both penalising unwanted behaviour, by use of the cost function, or constricting the feasible state space, by use of inequality constraints, so that the vehicle behaves within human factor derived limits.

In this chapter, the relative merits of these methods are discussed, along with the implications on the underlying control problems. Some human factor models of acceleration and curve negotiation are recalled and various ways of including them into the Hierarchical Model Predictive Control (HMPC) framework are considered.

Simulation results are then presented, with the HMPC scheme developed in Chapter 5 modified to utilise the human factor models as acceleration constraints, being used to control the nonlinear bicycle model, with tyre transience, through a double lane change manoeuvre.

Along with an evaluation of controller performance and the effect of the human factor models, the influence of the reference trajectories on the effectiveness of the proposed framework is shown.

Based on these insights, the HMPC framework is again updated, to show how with the flexibility of MPC, it is possible to adapt the system behaviour in response to human preferences that have been established in existing models.

6.2 Including Human Factors in the Control Problem

Before human factor models are discussed, a brief consideration is given to how they might be included in the MPC formulation. One of the key features of MPC is how it can handle constraints on the system states. In what follows, some potential methods of enforcing the constraints are discussed. Using human factor derived constraints restricts the feasible state space when determining the optimal control action. There are numerous scenarios, such as collision avoidance arising in emergency situations, where the occupants are not interested in comfort constraints and instead wish for the full vehicle capacity to be used. It is therefore necessary that the controller can adapt to these situations, for which different methods to accommodate this are outlined.

One implementation is to include the human factor constraints into the MPC optimisation problem as 'soft' constraints, which are constraints that should be obeyed whenever possible, but can be violated to ensure that the remaining constraints, now termed hard constraints, are satisfied. For example, ensuring that the lateral acceleration constraint is observed unless there is a lateral displacement bound constraint representing an obstacle which requires a higher acceleration manoeuvre to avoid it. For a typical quadratic programming problem, which linear MPC problems can be formulated as (see Appendix C), soft constraints are typically implemented by heavily penalising the soft constraint violation in the cost function [116]. This problem can be written as follows:

$$\underset{\boldsymbol{z},\, s}{\text{minimise}} \quad \boldsymbol{z}^T \mathbf{H} \boldsymbol{z} + \boldsymbol{f}^T \boldsymbol{z} + \rho \cdot s \tag{6.1a}$$

$$\text{subject to} \quad \mathbf{F}\boldsymbol{z} = \boldsymbol{b}, \tag{6.1b}$$

$$\mathbf{P}\boldsymbol{z} \leq \boldsymbol{h}, \tag{6.1c}$$

$$\bar{\mathbf{P}}\boldsymbol{z} \leq \bar{\boldsymbol{h}} + s, \tag{6.1d}$$

where $\bar{\mathbf{P}}$ and $\bar{\boldsymbol{h}}$ define the soft linear inequality constraints and s is the slack variable that allows for constraint violation. ρ is the positive scalar penalty weight for the slack variable, which must be chosen carefully so that it dominates the error tracking terms that make up the remainder of the cost function. The parameter must be chosen such that any small violation of the soft constraint results in a much larger cost than that occurred for even large state tracking errors. In doing so, the soft constraints are only violated to ensure that the hard constraints are satisfied. It should be noted that if in Eq. (6.1d), s is bounded above by the difference between the physical acceleration constraint terms and the soft constraint, the system of hard constraints can be reduced while still incorporating the friction derived constraints, and the increase in computational complexity of the problem is minimised.

An alternative method of including the human factor constraints, which avoids the need of determining the penalty weight ρ, is to solve the feasibility problem associated with the MPC constraints. In this, the human factor constraints are treated as hard constraints (or simply, constraints), and included in the linear inequality terms. The feasibility problem for the constraint set is then expressed as:

$$\underset{\boldsymbol{z},\, s}{\text{minimise}} \quad s \tag{6.2a}$$

$$\text{subject to} \quad \mathbf{F}\boldsymbol{z} = \boldsymbol{b}, \tag{6.2b}$$

$$\mathbf{P}\boldsymbol{z} \leq \boldsymbol{h} + s, \tag{6.2c}$$

where s is now the feasibility parameter associated with the constraint set. When non-positive, it indicates that there is at least one feasible solution to the MPC problem. Large negative values indicate that there is spare capacity in the constraint set, meaning that the feasible set is in some way quite large, and that the control sequence can be chosen that provides a good value for the MPC cost function.

For the human factor constrained problem, the more interesting scenario is when the feasibility parameter is positive, which reveals that there is no

admissible state and control sequence satisfying the constraints. In these conditions, there are a number of ways to proceed; the current constraint set could be enlarged, for example by deriving the lateral acceleration constraint from physical conditions rather than human preference.

As an alternative, or in addition, to using an enlarged constraint set, infeasibility detection could be used to trigger a number of protocols, such as emergency braking, differential braking, or specific emergency hierarchical frameworks to provide provision avoidance.

In practice, it might be prudent to use both of these methods in a combined form, by using the slack variable implementation to obey the human factor constraints wherever possible. In the event of a slack variable being activated, this could be used to trigger the necessary protocols, as in the case of the infeasibility detection. While this event is being processed, the MPC controller with soft constraints has already determined the best possible steering action, given the conditions, attempting to avoid the obstacle in the friction constrained scenario.

In this thesis, it is assumed that the vehicle is operating in conditions in which an emergency will not occur, i.e., there is a feasible solution to the MPC problem. This allows for the focus on how human factors change the behaviour of the controller. As such, the human factor derived constraints are simply implemented as standard linear inequality constraints.

6.3 Human Preference of Vehicle Handling Behaviour

So far, the discussion has been focussed on how human factors, particularly preference or tolerance to lateral acceleration, can be incorporated into the MPC framework. It assumes that there are clear numeric values for the human factor bound on the lateral acceleration value. The determination of correct values for these bounds is a non-trivial task, as it will depend on a specific individual preference. A complete industrial implementation,

would require large-scale prototyping and trial with people from different socio-economic groups, spanning a range of ages, to gather insight into differing reactions to a specific vehicle or vehicle platform.

The goal would be to create a driving profile that is satisfactory to as much of the population as possible. This would tend to a conservative constraint, and so, as is the case with numerous electronically controlled components (suspension damping, gearbox shifting speed, electronic throttle control), the manufacturer can provide various driving modes that change the dynamic response of the vehicle. By incorporating a broader range of dynamic driving styles, the vehicle then appeals to a larger proportion of the population.

Fortunately, by incorporating these human factor preferences as constraints, the task is to determine dynamic regions to avoid (such as high acceleration), rather than having to determine optimal lateral acceleration references for the controller to track, which would have to cover a wide variety of driving scenarios.

A primary objective of this thesis is to investigate how human factor preference towards autonomous driving behaviour can be incorporated into the control strategy and the subsequent effect it has on the control response. As such, the focus is not on finding acceleration bounds that accurately reflect specific values to use in an automotive application. Instead, it is desirable to show, using representative values for the acceleration bounds in a range of driving scenarios. These values, which cover a spectrum of acceleration tolerance bounds, prove the utility of MPC in the application of human factor driven constraints.

Therefore, data obtained from driving studies, which observed the driving behaviour of many drivers, across a range of driving conditions. From this, the researchers were able to determine, that among many influences, the longitudinal speed of the vehicle dominates in determining driver tolerance to lateral acceleration.

As discussed in the literature review in Chapter 2, the modified-Levison criterion is one such model to investigate, which has the following expres-

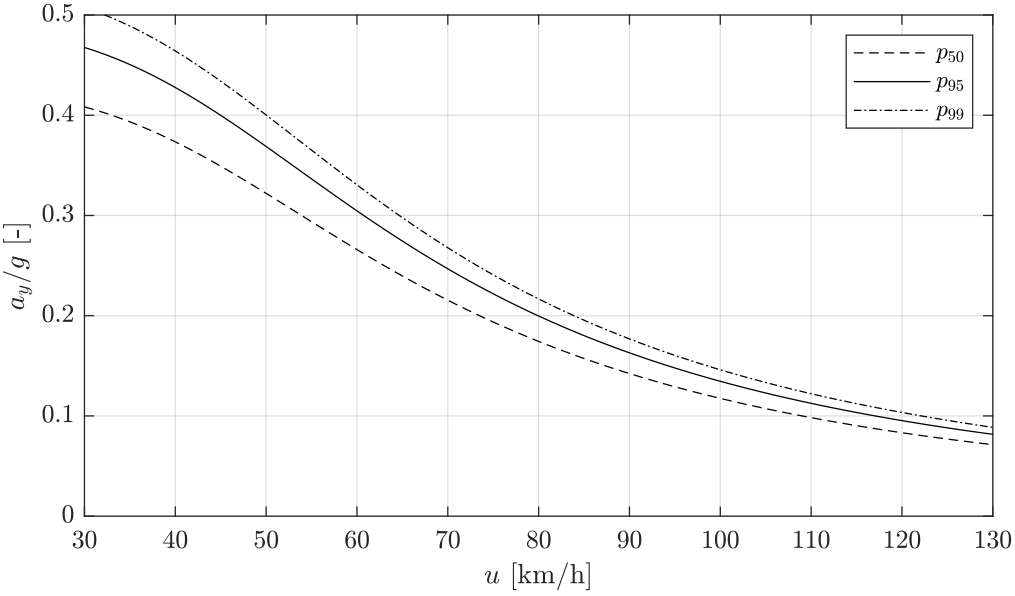


Figure 6.1: Median, 95th and 99.9th percentile values representing an upper bound on lateral acceleration

sion to model an upper limit for lateral acceleration [49]:

$$\bar{a}_y = \frac{a_0}{\sqrt{\left(1 - \frac{u}{u_0}\right)^2 + 2\frac{u}{u_0}^2}}, \tag{6.3}$$

where u_0 is the reference velocity at which the results were determined for, and more interestingly, a_0 is the parameter which determines the driver's lateral acceleration tolerance. The values in Table 6.1 show representative values for the median, 95th and 99.9th percentile of studied drivers.

Figure 6.1 shows the acceleration profiles for median, 95th and 99.9th percentile drivers, using the data obtained from the driver studies in [49]. It shows the decreasing tolerance to lateral acceleration is shared by all drivers. At motorways speeds, around 110-120 km/h, the difference becomes negligible with a 0.02 g difference between the 99.9th and median driver profiles at 120 km/h. The characteristic curve was adapted from an earlier empirical study [53], which had the limitation of being unbounded at low-speeds. Figure 6.1 shows that this is not the case anymore with the median drivers profile tolerance just exceeding 0.4 g at 30 km/h, while the 99.9th percentile profile exceeds 0.5 g.

Table 6.1: Acceleration bound parameters using the modified Levision model, with reference velocity $u_0 = 14.84$ m/s, taken from [50]

	a_0	Units
median	4.20	m/s^2
95th	4.81	m/s^2
99.9th	5.22	m/s^2

It should be noted that in [49], the authors identify a second behaviour mode in pure motorway driving conditions. In this, the lateral acceleration was bounded by the centripetal acceleration required to navigate a corner of fixed curvature, i.e. $\bar{a} = \kappa u^2$, where κ is the road curvature. This indicates that during motorway driving, certain drivers prefer to prioritise their driving speed, rather than slowing to reduce cornering acceleration. This behaviour was also observed in the earlier driving study by Herrin and Neuhardt [48].

In the following results, this expression will be used as a proxy for tolerance or acceptance, which is valid assuming that the drivers in the study successfully navigated corners as they expected.

By using the median and percentile value shown in Table 6.1 to derive bound limits in the linear MPC controller, the effect this has on controller feasibility and dynamic closed-loop system response can be studied. The three different profiles will be used to study how driving behaviour changes across the driver preference spectrum.

6.4 Results

The linear MPC controller is evaluated for three different simulation speeds for the median 95th and 99.9th percentile driver profiles. To ensure the valuation of the controller provides meaningful results the cost mentioned in the simulation speed must be chosen to demonstrate the dynamic conditions that could occur in standard driving situations.

For the simulation model, the nonlinear bicycle model is used, while

the control plant is the linear bicycle model discretised with the Euler approximation. As throughout the entirety of the thesis, the controller has an operating frequency of 10 Hz.

6.4.1 Creating Reference Courses for Control Evaluation

As stated previously the results profile the results profile vehicles and standard (non-emergency) driving scenarios, which means that the driving scenario should be feasible for all acceleration profiles and speed combinations. The second consideration is to ensure that the driving scenario is challenging enough. By this, it is meant that if the constraints are relaxed enough such that the optimal control sequence is the same for all acceleration profiles, then there would be no change in driving behaviour. While this consistent behaviour is acceptable or even desirable for autonomous driving vehicles, the area of interest from a research perspective is establishing the conditions in which this behaviour diverges.

6.4.2 Evaluating Control Performance via RMS Error Tracking

It is beneficial to have a scalar quantitative measure to give a single objective metric to evaluate controller performance. One choice is to calculate the Root Mean Square (RMS) error between the simulated system state and the reference trajectories. This then allows for a direct comparison between for controllers assessed on the same course. A limitation with using the dynamically generated reference path is that it does not give a complete representation of the vehicle tracking behaviour over time. This is because, the reference trajectory provided to the lower level of the HMPC controller changes in part due to the tracking performance of the lower HMPC level. It is hard therefore to identify changes in controller behaviour over time, such as small but consistent drifts in tracking due to active constraints like an acceleration bound.

To mitigate against this issue, a static RMS tracking error is also defined, with the previous term referred to as the dynamic RMS tracking error. The static RMS error, attempts to account for changes in tracking performance over time, compared to what might have been determined to be the optimal reference path ahead of time. To achieve this, a static reference trajectory is determined for a given road scenario. Unlike the reference planning undertaken in the HMPC framework, the preview horizon is chosen to cover the entire course. Calculating the error term at each time-step is then simple as the simulated state vector at each time-step can be used with the corresponding reference vector from the static path to determine the error. It is emphasised that this calculation is done for analysis purposes after the simulation is complete and that this is not the error term that the MPC controller is trying to minimise.

In the case of the dynamic RMS tracking error, there exists a sequence of reference trajectories provided to the lower-level MPC controller. In this case for a given time-step, the most recent reference trajectory is used to calculate the error vector.

6.4.3 Preliminary Results

First, the simulation results using the three acceleration constraint profiles for the 50th, 95th and 99.9th Levison profiles is shown using the same parameter values for the reference trajectory planning module shown in the previous chapter. Then, the simulations performed again with the reference path providing the minimum acceleration path, in which it is shown that the reference path is a considerable influence on the closed loop response. For the road conditions presented in Section 5.3, 65 km/h was shown to be an acceptable longitudinal value. The results, for the three acceleration bound profiles are shown in Figure 6.2. For all three profiles it can be seen in Figure 6.2 (a) the lateral velocity is broadly the same. It is instead the yaw rate that changes to limit the lateral acceleration as observed in Figure 6.2 (c). The resulting acceleration profiles are shown in Figure 6.3 (a) with all the profiles obeying the respective limits. In

Table 6.2: Dynamic RMS tracking error for controller with constant longitudindal velocity of 65 km/h

Driver percentile	v [m/s]	Y [cm]	ω_z [°/s]	ψ [°]
50th	0.0493	0.1255	1.8402	0.3698
95th	0.0474	0.1268	1.6674	0.3683
99.9th	0.0471	0.1297	1.7018	0.3759

Table 6.3: Static RMS tracking error for controller with constant longitudindal velocity of 65 km/h

Driver percentile	v [m/s]	Y [cm]	ω_z [°/s]	ψ [°]
50th	0.0479	3.5454	2.2665	0.8180
95th	0.0465	1.6181	1.7127	0.4583
99.9th	0.0463	1.2480	1.5922	0.3825

Figure 6.3 (b) acceleration saturation (the ratio $a_y/\bar{a}$) limit is shown and the median profile can be observed to fully saturate while returning the vehicle to the inside lane.

Ultimately, this results in the increased lateral displacement tracking error observed in Figure 6.2 (b). The 95th percentile saturates momentarily and so has similar tracking performance to the 99.9th percentile acceleration profiles acceleration profile. The median profile however saturates at a lower magnitude, the lateral acceleration takes longer to return to zero at the end of the simulation manoeuvre.

Tables 6.2 and 6.3 shows a summary of the dynamic and static RMS error tracking. The dynamic tracking errors have similar values, across the state variable, for all three runs. This is due to the capabilities of the path planning module of the hierarchical framework, creating a reference that the controller can accurately track, regardless of the acceleration bound applied to the Active Steering Control (ASC) module. The static RMS errors show how the reference has evolved over time, in particular, the lateral displacement error for the 50th percentile is nearly three times higher than the 99.9th percentile error. For the 50th percentile, the yaw tracking static error is more than double the dynamic error, a consequence of the much higher lateral displacement error. This is not observed for the other simulation results.

Table 6.4: Summary of acceleration data for the profile run

Driver percentile	$a_{y,\ \mathrm{RMS}}/g$ [-]	$a_{y,\ \mathrm{RMS}}/\bar{a}$ [%]	$\|a_y\|_\infty/\bar{a}$ [%]
50th	0.140	58.3	102.4
95th	0.142	51.7	101.8
99.9th	0.142	47.6	100.3

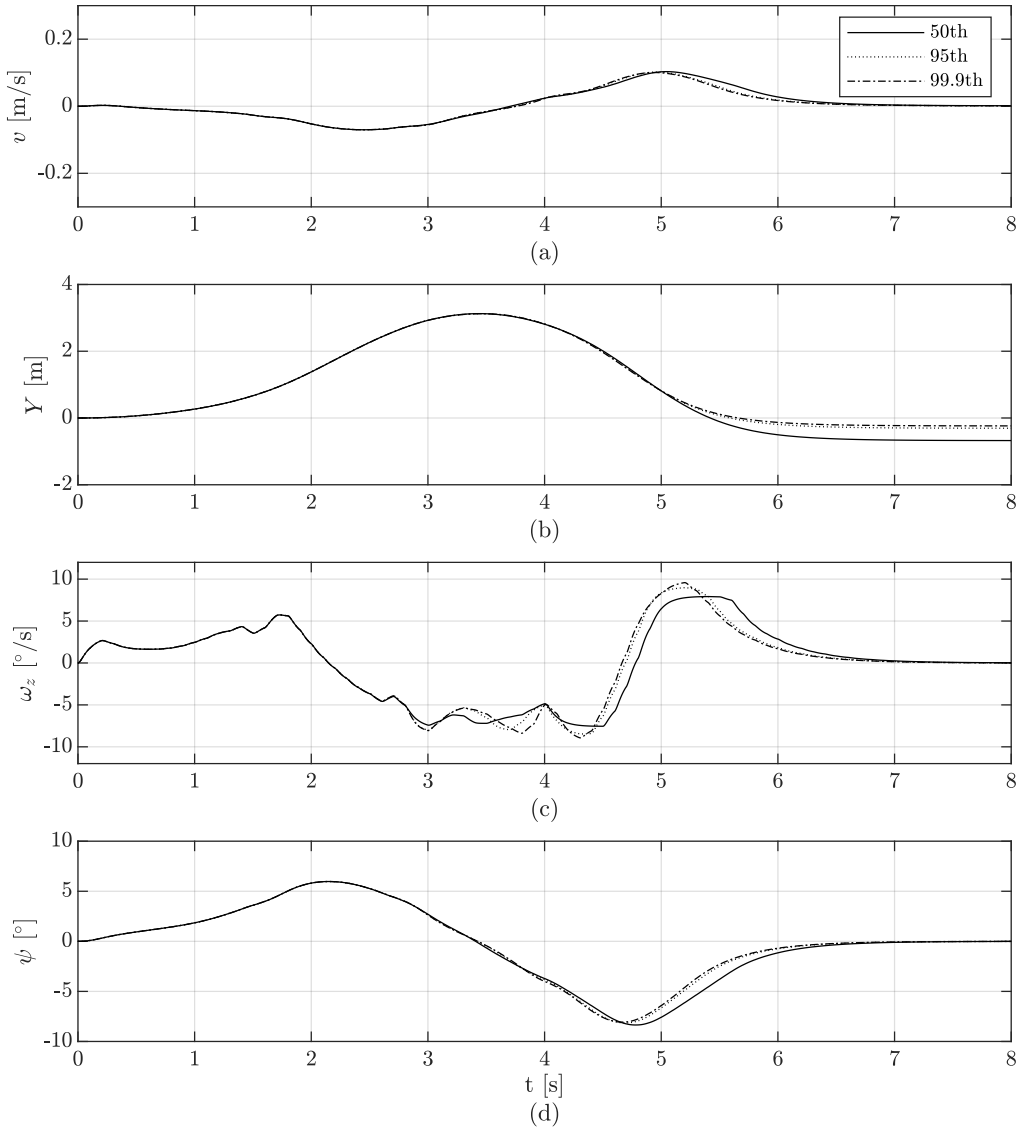


Figure 6.2: Reference tracking at 65 km/h, for the controller equiped with median, 95th and 99.9th driver profiles. The median profile can not track the reference trajectory as closely due its lower acceleration bounds

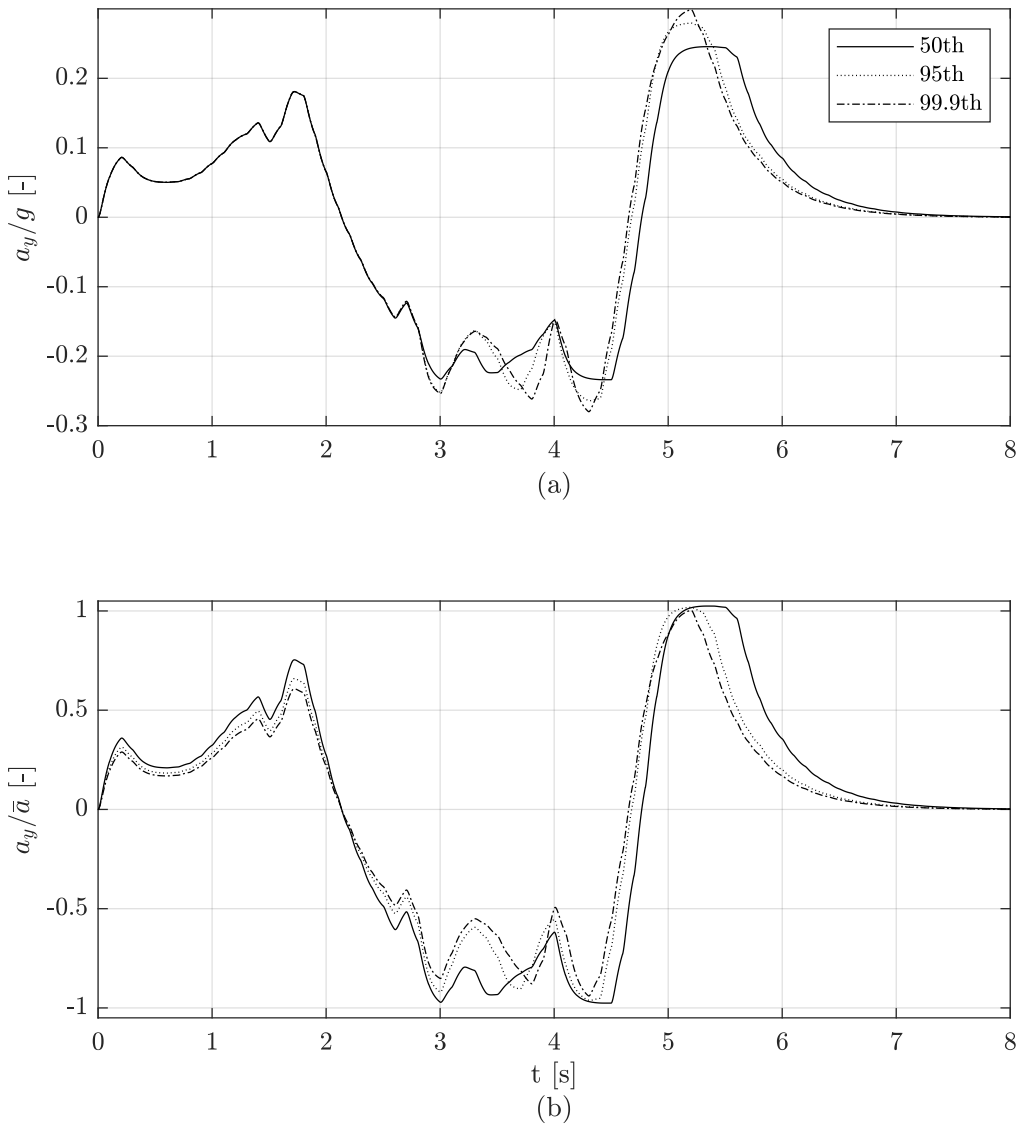


Figure 6.3: Lateral aceleration of the simulation vehicle with the the three acceleration profiles using the modified Levison criterion. The saturation of the respective bounds is also shown, and it can be observed that the median profile spends a larger proportion of its time saturated. The 99.9th percentile profile is never fully saturated

6.4.4 Tracking the Minimum Acceleration Reference

The results so far have shown the controller response when tracking the provided, smooth, trajectory. In this section the previous reference path is replaced with a profile that minimises the second derivative of the lateral displacement reference. In doing so the path can be thought of as an approximation of the optimal minimum acceleration path for vehicle tracking it.

Figure 6.4 shows the state trajectories for the acceleration profiles. In contrast to those shown in Figure 6.2, the state projectors are exceptionally similar for three different acceleration bounds. This can be explained by considering, that when tracking the minimum acceleration profile, the optimum state trajectories is feasible to all three control instances with the differing acceleration bounds. The only observable difference between the state trajectories occurs in the yaw velocity, around four seconds into the simulation. This change happens to prevent the respective lateral acceleration constraints being violated as the vehicle returns to the initial lane. It can be observed in Figure 6.5 (b) that the median profile in particular is close to saturation, which causes the change is yaw velocity.

Tables 6.5 and 6.6 show the vehicles dynamic and static tracking errors respectively. Compared to the initial path tracking results just shown, there is less change between the dynamic and static measure. This is particularly evident when comparing the lateral displacement errors. When tracking the minimum acceleration trajectory, the static RMS error for Y is approximately three times larger than the dynamic RMS value. In contrast, when tracking the original trajectories, the static Y RMS error range from

Table 6.5: Dynamic RMS tracking error for controller with minimum acceleration reference path for front-wheel steering vehicle with constant longitudinal velocity of 65 km/h

Driver percentile	v [m/s]	Y [cm]	ω_z [°/s]	ψ [°]
50th	0.0498	0.0983	1.4414	0.3416
95th	0.0500	0.0940	1.3450	0.3229
99.9th	0.0500	0.0931	1.3363	0.3209

Table 6.6: Static RMS tracking error for controller with minimum acceleration reference path for front-wheel steering vehicle with constant longitudinal velocity of 65 km/h

Driver percentile	v [m/s]	Y [cm]	ω_z [°/s]	ψ [°]
50th	0.0494	0.3067	0.9205	0.2347
95th	0.0496	0.2630	0.8578	0.2090
99.9th	0.0496	0.2554	0.8468	0.2042

Table 6.7: Summary of acceleration data for the three profiles when tracking the minimum acceleration reference at 65 km/h

Driver percentile	$a_{y,\,\mathrm{RMS}}/g$ [-]	$a_{y,\,\mathrm{RMS}}/\bar{a}$ [%]	$\|a_y\|_\infty/\bar{a}$ [%]
50th	0.133	55.7	100.8
95th	0.134	48.8	95.4
99.9th	0.134	45.0	90.1

ten times larger for the 99.9th percentile run to 28 times larger for the median profiles.

As expected, this can be explained by the vehicles lateral acceleration in the initial simulation. In Table 6.7, the vehicles acceleration data is shown. It can be seen that only the 50th percentile data has a saturated limit. The 0.8% overshoot can be attributed to the model mismatch between controller and vehicle simulation. For the 95th and 99.9th percentile manoeuvres, there is a 5% and 10% capacity remaining respectively.

The median profiles static RMS Y error is 17-20% greater than the 95th or 99.9th percentile runs, suggesting that, at 65 km/h, the median acceleration bound vehicle is close to its saturation limit, and if the simulation was performed at a higher speed, the controller would not be able to determine a feasible control sequence that satisfies all of the constraints imposed upon it.

It can be seen, that by comparing Figures 6.2 and 6.4 that in the median profile, the controller changes the vehicles lateral displacement output in the standard reference trajectory run to be similar to the minimum acceleration references to maintain feasibility when tracking non-minimum acceleration references.

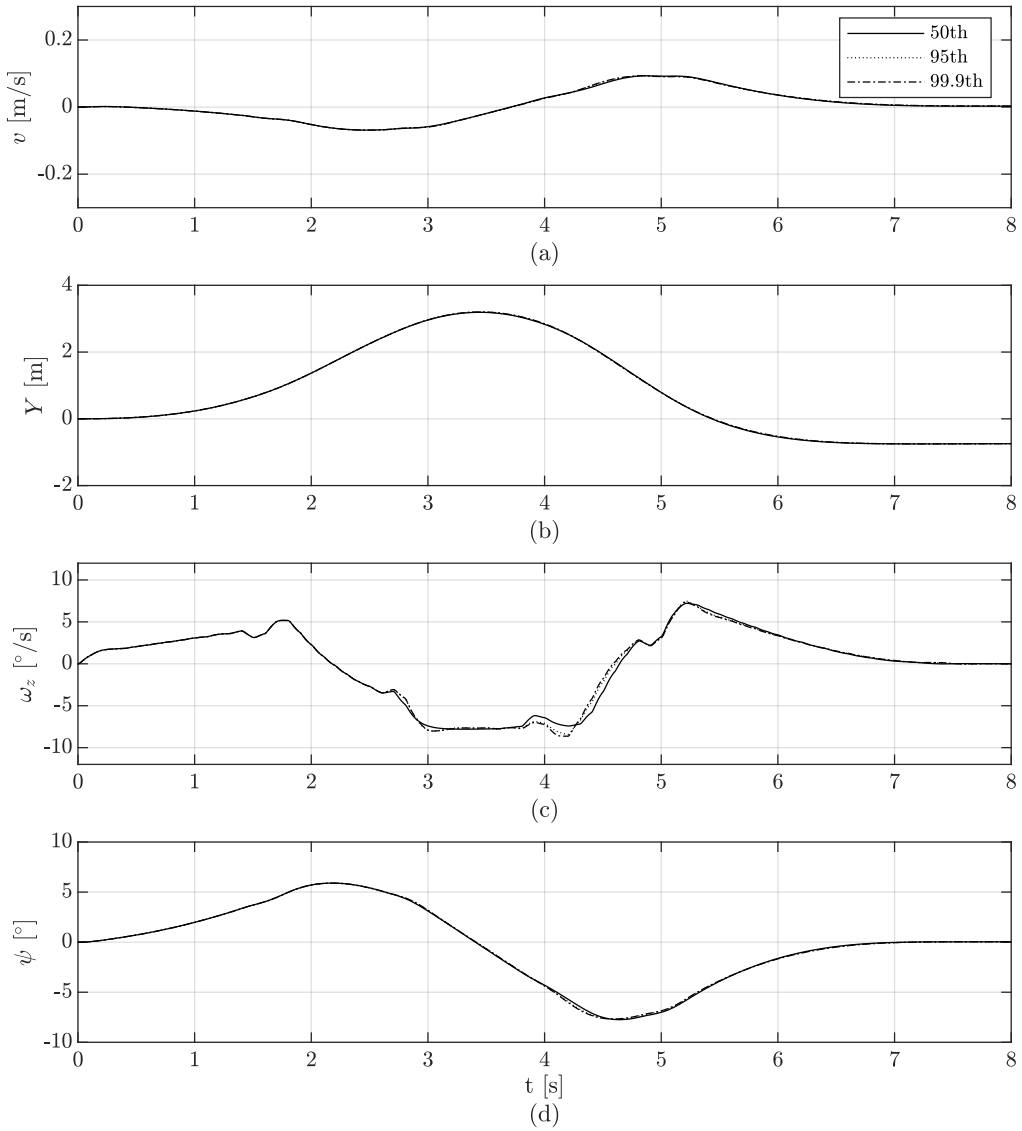


Figure 6.4: State values for the minimum acceleration reference trajectory, shown for the median (50th), 95th and 99.9th percentile. As the minimum acceleration trajectory is feasible to all tolerance profiles, the vehicle behaviour is the same for all.

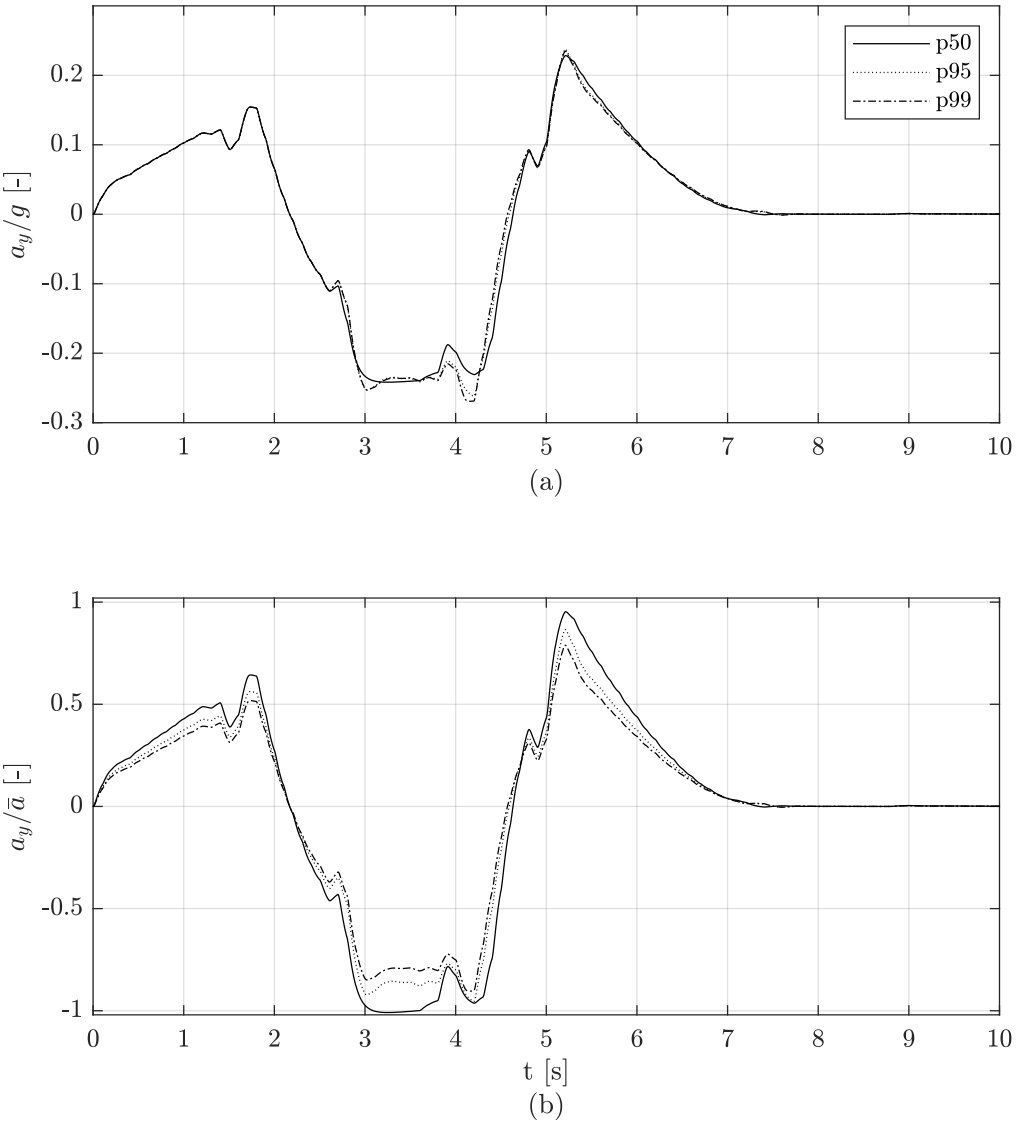


Figure 6.5: Comparison of the median (50th), 95th and 99.9th acceleration profiles for the minimum acceleration profile curves. As the state trajectory tracking is the same for all profiles, so is the lateral acceleration (a). In (b), it can be seen that this results in a greater saturation of the acceleration bound for the median driver

6.5 Penalising Acceleration in the MPC Cost Function

The results so far have shown that the changes in vehicle performance can be subtle. When a minimum acceleration curve is used as the reference trajectory, the differences become negligible. This is to be expected, if one considers the underlying optimisation problem for the MPC formulation. For a feasible solution, the only concern is minimising the cost function, which for MPC is governed by minimising system tracking error and control input. For the minimum acceleration reference, MPC tracks the trajectory until at a sufficiently high longitudinal speed, u, there is no longer a feasible solution that satisfies all the constraints. At this point, the standard MPC formulation is not capable of determining an appropriate control sequence.

For non-minimum acceleration curves, such as those shown in Section 6.4.4, the same is also true in that the controller is attempting to track its reference and that there is a speed at which it is no longer possible. The difference is that in this instance there is still a set of feasible routes to choose from. By using an acceleration bound like the modified Levison criteria, which is dependent on u, the speed at which the controller cannot track the reference changes and results in the different control profiles and vehicle behaviour.

From this, it can be concluded that while there are reasons why it is desirable to use Human Factors as constraints, other factors in the overall control framework, can render them ineffective as differentiators of autonomous driving behaviour.

If this lack of variation is problematic, it leads us to consider ways of modifying autonomous driving behaviour more markedly in response to a given persons acceleration percentile. In doing so, the controller can be encouraged to avoid high acceleration values when possible.

The most obvious extension is to incorporate into the cost function a penalty on the magnitude of the lateral acceleration. To do so is straight-

forward, as the lateral acceleration is modelled in the MPC formulation as a linear function of system state. This allows for the existing cost function terms to be augmented to include an acceleration penalty rather than introducing additional states or variables in the formulation.

The augmented penalty terms can be most concisely expressed by considering the formulation in Quadratic Programming (QP) form (see Appendix C for its presentation). In doing so, we can note that the predicted acceleration vector, $\boldsymbol{a}_y \in R^N$, can be expressed as a linear function of the optimisation vector, $\boldsymbol{z} \in R^{N(m+n)}$:

$$\boldsymbol{a}_y = \mathbf{A}\boldsymbol{z} + \boldsymbol{v}, \tag{6.4}$$

where

$$\mathbf{A} = \begin{bmatrix} \mathbf{0}, & \mathbf{D}+\mathbf{E} & \mathbf{0} & \mathbf{0} & \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} & \mathbf{D}+\mathbf{E} & \\ & \ddots & \ddots & \ddots & \\ & & -\mathbf{D} & \mathbf{0} & \mathbf{D}+\mathbf{E} \end{bmatrix} \in \mathbb{R}^{N \times N(n+m)},$$
$$\boldsymbol{v}^T = \begin{bmatrix} (\mathbf{D}\boldsymbol{x})^T, & \mathbf{0}, & \ldots, & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{N(n+m)}.$$

The matrix terms $\mathbf{D} = [1,\ 0,\ 0,\ 0]$ and $\mathbf{E} = [0,\ 0,\ u,\ 0]$ are used to calculate the lateral acceleration values. Then the cost function terms can be modified to incorporate a cost on the magnitude of the acceleration vector by writing them as:

$$\begin{aligned} \tilde{\mathbf{H}} &= \mathbf{H} + \alpha \mathbf{A}^T \mathbf{A} \\ \tilde{\boldsymbol{f}} &= \boldsymbol{f} + 2\alpha \mathbf{A}^T \boldsymbol{v} \end{aligned} \tag{6.5}$$

where α is the penalty weight for the lateral acceleration to be chosen. It should be noted that if a terminal stabilising weight matrix, $\mathbf{Q}_f$, was used in the original problem this must be recalculated by using the modified $\mathbf{Q}$ matrix penalty in the LQR problem [86], and the $\mathbf{H}$ and $\boldsymbol{f}$ terms must be changed accordingly.

An interesting observation to note is that, by inspection of Eq. (6.5), it can be seen that there is now an explicit relationship between longitudinal

speed u and the cost function terms. Previously, u was only present in the MPC constraints, and so u affected the feasible states, rather than the desirable states. Of course, if so desired it would be possible to choose α as a function of u to remove this dependency (or even augment it). This is investigatged in the next section.

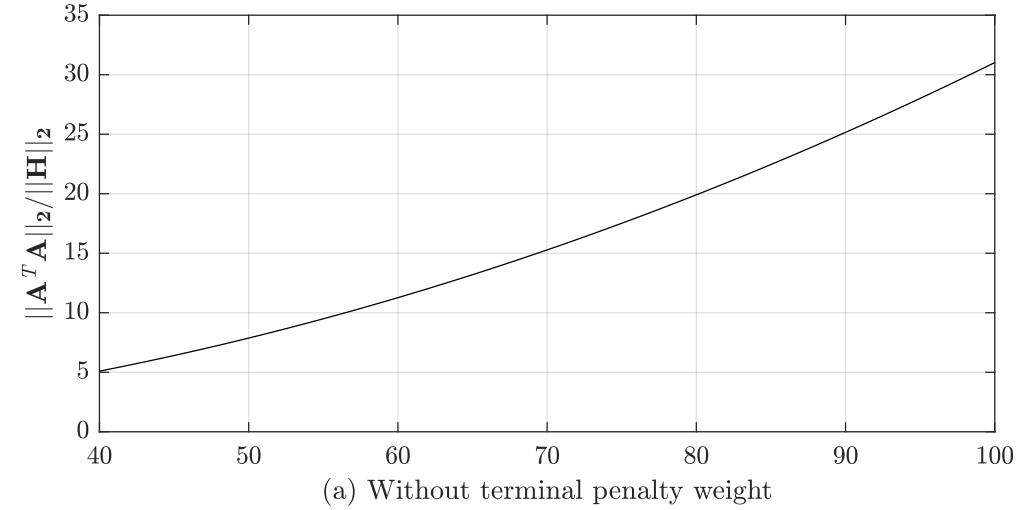
6.5.1 Choosing the Acceleration Penalty Weight Value

Now that a method of incorporating a penalty on the lateral acceleration, it is necessary to determine an appropriate value for α to update the matrix terms in Eq. (6.5). For this thesis, in which the aim is to demonstrate how human factor issues can be incorporated into the control problem formulation, it is desirable to derive the penalty value α from the model data chosen. One obvious choice is to make use of the a_0 values from [49], and assume that the a driver that has a higher limit on lateral acceleration would naturally tolerate a larger average acceleration.

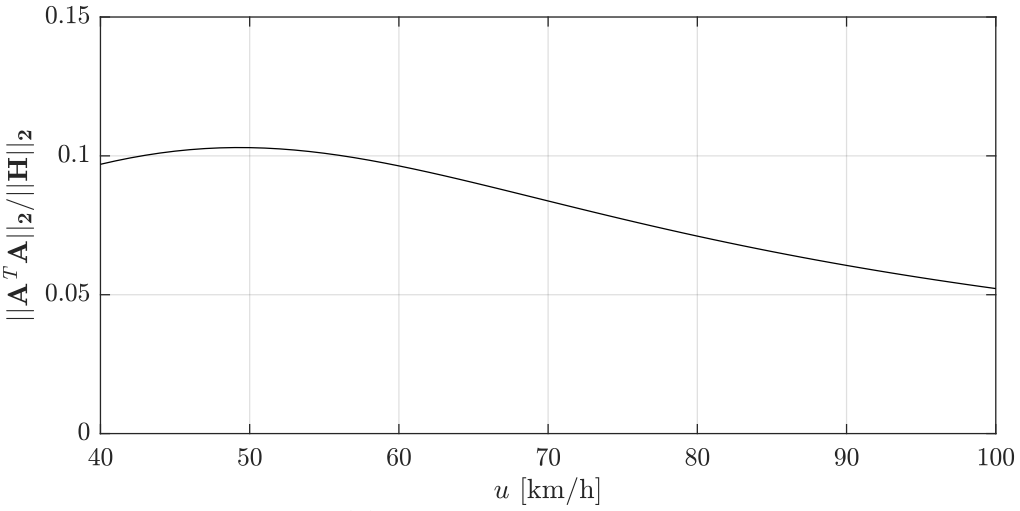
A small parameter study was conducted to find appropriate values for α. First, the influence of the acceleration term in the QP matrix $\tilde{\mathbf{H}}$ was studied.

In Figure 6.6 (a), the Euclidean norm of the acceleration term (i.e. $\left\|\mathbf{A}^T\mathbf{A}\right\|_2$) is shown for a range of longitudinal velocity values. The graph is scaled by the norm of the original $\mathbf{H}$ term. The Euclidean term is useful as it presents a scalar value that can be thought of as the maximum gain of the matrix and represents the relative influence that the acceleration penalty has compared to the state terms. A ratio greater than unity indicates that the acceleration term dominates the tracking term. In these results, a preview horizon of ten time-steps was chosen with $\mathbf{Q} = \text{diag}(5,\ 25,\ 10, 5)$ and $\mathbf{R} = 1$.

It can be seen that the acceleration term dominates the standard quadratic matrix, even at the relatively low speed of 40 km/h, with its norm value being 5 times greater. As expected, this increases with a quadratic dependence on longitudinal velocity, u, indicating that a small value is required to scale α to ensure that acceptable path tracking is achieved while penalising



(a) Without terminal penalty weight



(b) With terminal penalty weight

Figure 6.6: Norm of the acceleration term added to the optimisation problem's quadratic term. It can be seen that including the terminal penalty matrix $\mathbf{Q}_f$ (in (b)) has a large effect on the comparative weight of the matrix terms, with a large change in relative importance that can be seen by the change in scale of the y-axis between (a) and (b).

the lateral acceleration. That u has a positive correlation the relative importance of acceleration penalty is beneficial, as it helps achieve the design goal of matching the behaviour of the Levison curve.

An interesting variation of this is to consider replacing the final block in $\mathbf{H}$ with $\mathbf{Q}_f$ as discussed in Chapter 4, Section 4.2.1. In this, $\mathbf{Q}_f$ must be determined for every value of u.

Figure 6.6 (b) shows the relative of influence of the acceleration term compared to the standard quadratic matrix term, when $\mathbf{Q}_f$ is included. $\mathbf{Q}_f$ changes drastically for u increases, and that the norm value $\left\|\mathbf{Q}_f\right\|$ increases as u gets larger an (elevenfold increase from 40 km/h to 100km/h). Using the $\mathbf{Q}_f$ term in the optimisation problem has been done for all previous results, and means that without scaling α accordingly, the relative penalty placed on lateral acceleration decreases at higher speeds. This is directly opposed to the findings that led to the modified Levison curve characterisation. It is clear therefore that α should be chosen, in the instance that a terminal penalty matrix is chosen, so that the $\mathbf{A}^T\mathbf{A}$ term becomes stronger as u increases.

It was found that be determining α as a function of u^3 restored the influence of longitudinal speed on the ratio of the two matrix norms. The can be observed in Figure 6.7, where the weighting values was given by:

$$\alpha = \frac{u}{u_0}^3, \tag{6.6}$$

where u_0 is a reference velocity to normalise the term and to determine when the weighting value becomes greater than one, i.e. at u_0. For the results shown in the thesis, $u_0 = 25$ km/h.

From preliminary studies, it was seen that penalising acceleration from a_0 values with a homogenous polynomial of a_0 (e.g. a_0^2, a_0^3) was insufficient at differentiating vehicle behaviour. Instead it was found that an exponential function of a_0 was more appropriate, resulting in the following term for α:

$$\alpha = \beta \cdot e^{-ka_0} \left(\frac{u}{u_0}\right)^3, \tag{6.7}$$

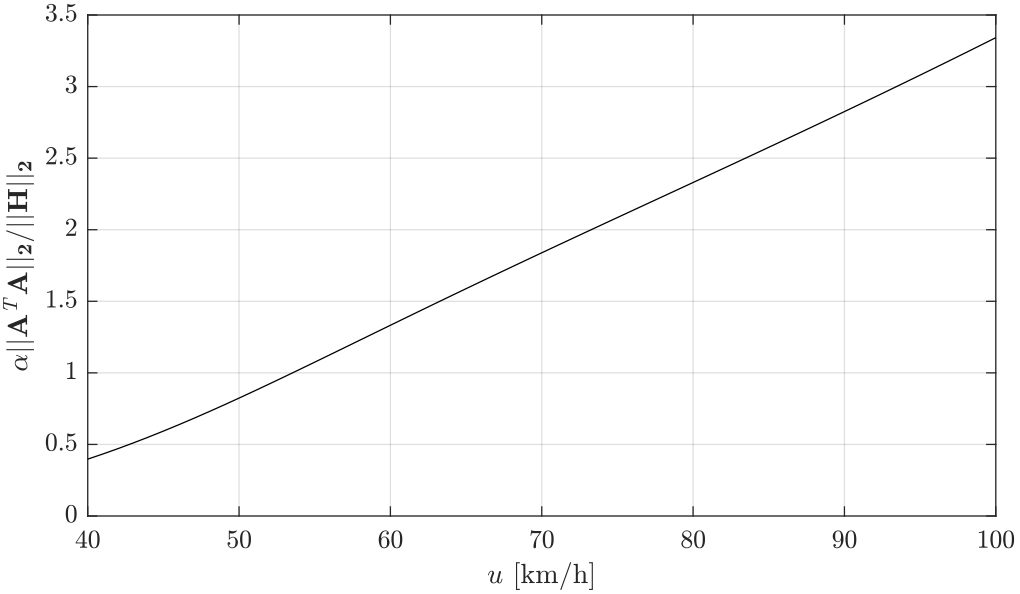


Figure 6.7: Using u^3 to derive α restores the positive relationship between lateral acceleration penalty and longitudinal velocity

where $\beta,\ \kappa > 0$ are additional constants to be determined. To find accurate values corresponding to human preferences, experimental studies would have to be undertaken in driving simulators or a suitable equipped test vehicle. For this thesis, representative values are chosen to demonstrate how a simulation vehicles behaviour changes when equipped with this human factor aware HMPC framework. Simulation results for this augmented controller are shown in the next section.

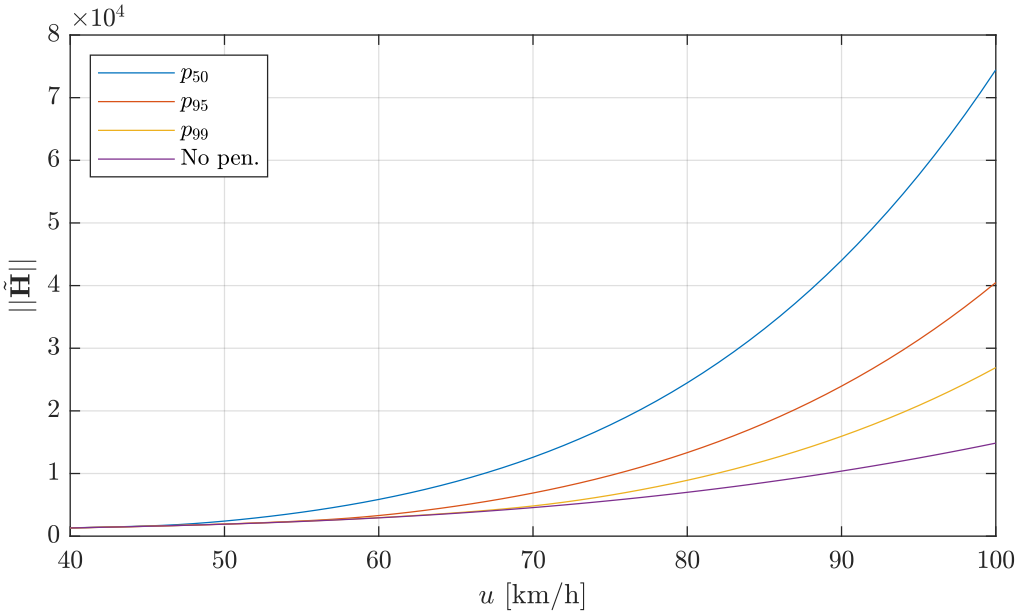


Figure 6.8: Norm of quadratic matrix term for common road speeds, using the penalty term Eq. (6.7). Until 50 km/h, the three profiles have a very similar matrix norm values to that without the acceleration penalty term

6.5.2 Results

The augmented cost function has a fundamental impact on the objective; not only is the controller attempting to minimise the error between predicted and reference states, it is attempting to minimise the magnitude of a linear combination of these states. To verify the change in formulation and to allow for the study of the resulting vehicle behaviour, the ASC behaviour is first investigated without the hierarchical framework. Instead, a simpler linear step in Y displacement is used, as in Chapter 4.

Four sets of results are presented, with a focus on the lateral acceleration behaviour. The MPC controller, with acceleration penalty weighting, is shown for the median, 95th and 99.9th percentile profiles along with the standard MPC controller with acceleration bounds for the median driver, as a reference. For this thesis, $\beta = 100$ and $u_0 = 25$ km/h. All other parameters are kept the same as the previous simulation run, to allow for comparison between the two sets of results.

Figure 6.10 (a) shows the lateral acceleration of the four controller configurations. Only the profile without an acceleration penalty fully saturates

its acceleration bound at any point. In doing so, it is able to minimise the state tracking errors. In contrast, all other profiles stay within their respective bounds, and so have a smoother acceleration profile, rather than the flattened profile seen in the results without the acceleration penalty. Despite none of the other results saturating their respective bounds, it can be seen that the maximum acceleration for the 95th and 99.9th profiles exceeds the median, unweighted, profile.

One important aspect of the acceleration penalty design is that it should encourage the controller equipped with a median profile to minimise the acceleration more that it does for the 99.9th percentile profile. This is valid if it is a reasonable assumption that drivers with a higher acceleration bound tolerate a larger average acceleration.

To inspect this behaviour, the acceleration saturation is also shown, in Figure 6.10 (b). The acceleration saturation is defined as the ratio of the lateral acceleration of the vehicle to the controllers acceleration bound. This allows for a direct comparison between the controllers, even though they have different bound values. The unweighted results have the highest saturation values. This is to be expected as there is no penalty placed on having a consistently high acceleration that is within the acceleration bound constraint.

Table 6.8 supports these findings, with the maximum acceleration relative to the acceleration bounds being 92%, 96% and 98% for the median, 95th and 99.9th profiles respectively. The median profile, without acceleration penalty, exceeds its constraint by 3%. This small violation is due to the mismatch between the simulation and controller vehicle models.

The RMS average acceleration saturation for the median control run is lower, as expected, at 0.238, than the 95th and 99.9th profile run values at 0.249 and 0.258, respectively. These values are significantly lower than the median reference profile, which has an RMS saturation value of 0.499, due to the lack of acceleration penalty. The median profile also has the lowest value for the RMS average of absolute acceleration. This shows that in this sense, the acceleration penalty, and particularly the penalty weight α given

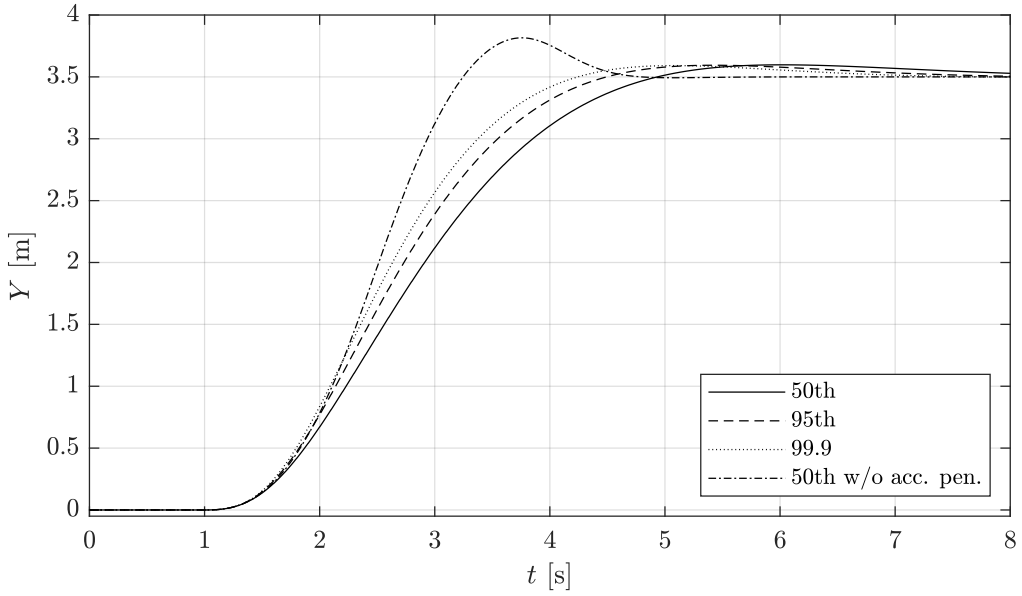


Figure 6.9: Reference tracking for the acceleration penalty profiles

Table 6.8: Summary of weighting values and saturation values. The median profile without acceleration profile is the last entry ($\beta = 0$)

Profile	α [-]	$\|\|a_y/\bar{a}\|\|_\infty$ [%]	RMS($a_y/\bar{a}$) [-]	RMS(a_y/g) [-]
50th	26.36	91.5%	0.238	0.057
95th	14.32	95.7%	0.249	0.068
99.9th	9.50	98.0%	0.258	0.077
50th ($\beta = 0$)	0.00	102.7%	0.499	0.120

by Eq. (6.7), have been effective at changing the vehicles driving style in response to the human factor parameters derived in [49].

Finally, Figure 6.9 shows the lateral displacement, Y, for the four simulation runs. The behaviour is broadly similar for all, but the reference result, without penalty, rises the last but rises the fastest, to minimise displacement tracking error.

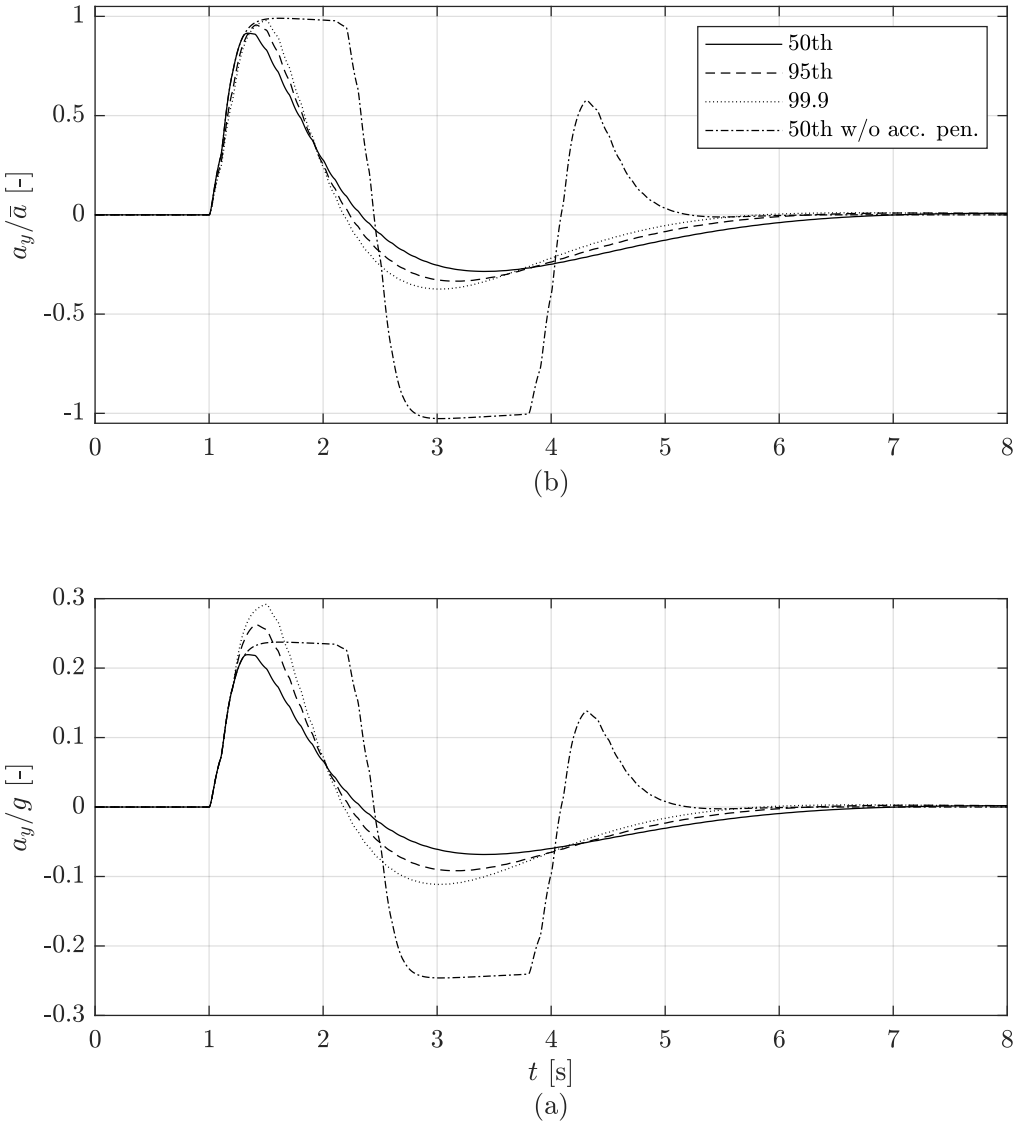


Figure 6.10: Lateral acceleration for the four penalty profiles, with the acceleration saturation shown below

HMPC Results

Using the parameter values just used, and the acceleration penalty model developed, the double lane change manoeuvre is simulated again. All parameters are kept the same, and the manoeuvre is performed at 65 km/h. The parameter values from Table 5.3 are used for path generation, rather than the minimum acceleration trajectory parameters.

Figure 6.11 shows the simulation results for the 50th, 95th and 99.9th percentile runs. Interestingly, there is less differentiation between the three lateral displacement trajectories compared to the results shown in Figure 6.2. This is due to the acceleration penalty included in the cost function 'pushing the vehicle lateral response away from the reference, in an effort to minimise acceleration. The median profile vehicle cannot adjust its trajectory in any meaningful way as it is already saturating its acceleration bound, but the other two vehicles displacement straightens as they return to the initial lane.

Table 6.9: RMS dynamic tracking error for controller with acceleration penalty for front-wheel steering vehicle with constant longitudinal velocity of 65 km/h

Driver percentile	v [m/s]	Y [cm]	ω_z [°/s]	ψ [°]
50th	0.0504	0.0946	1.2514	0.2336
95th	0.0494	0.0859	0.8994	0.2004
99.9th	0.0491	0.0870	0.7784	0.1950

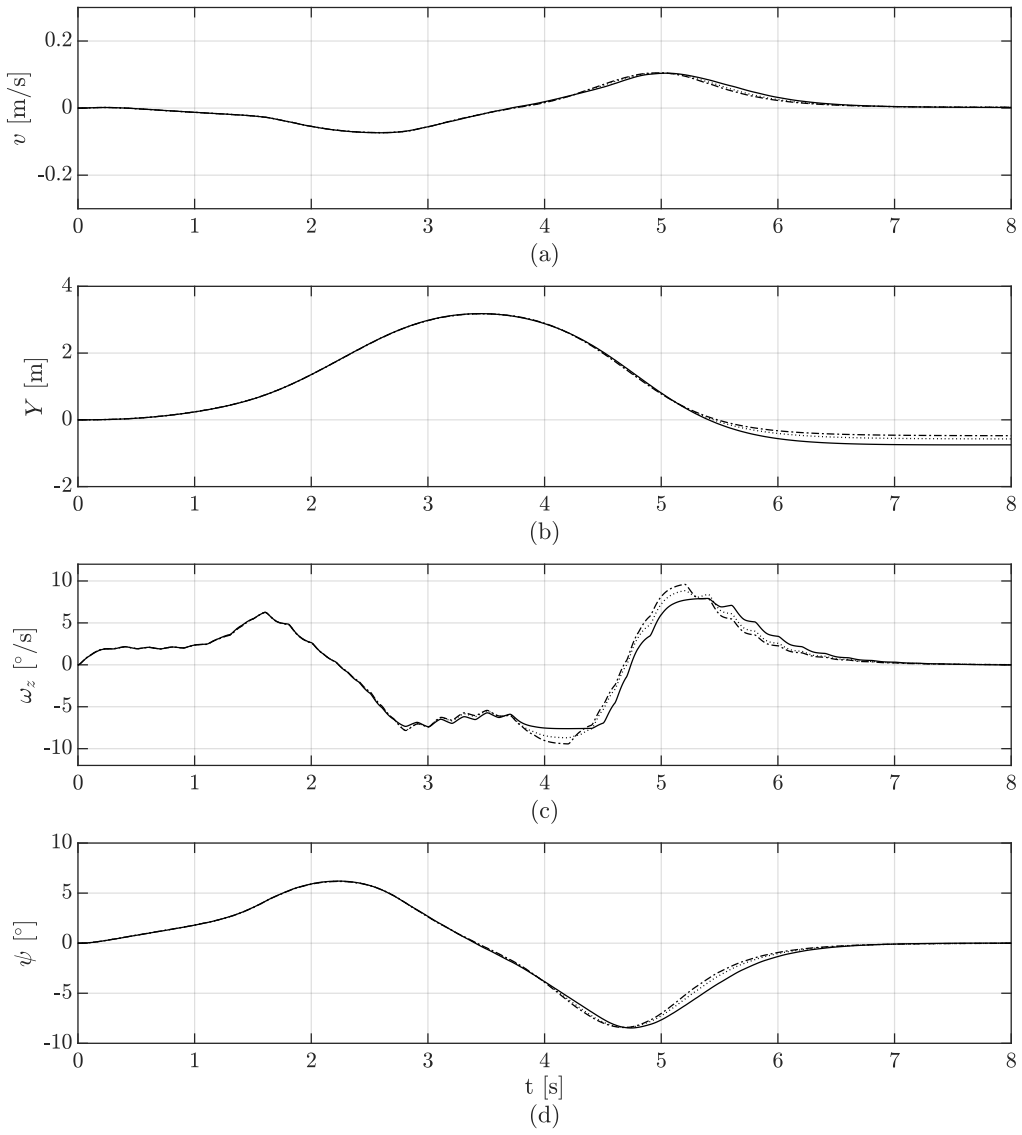


Figure 6.11: State variables for vehicle controlled with MPC controller with acceleration penaltyt in cost function

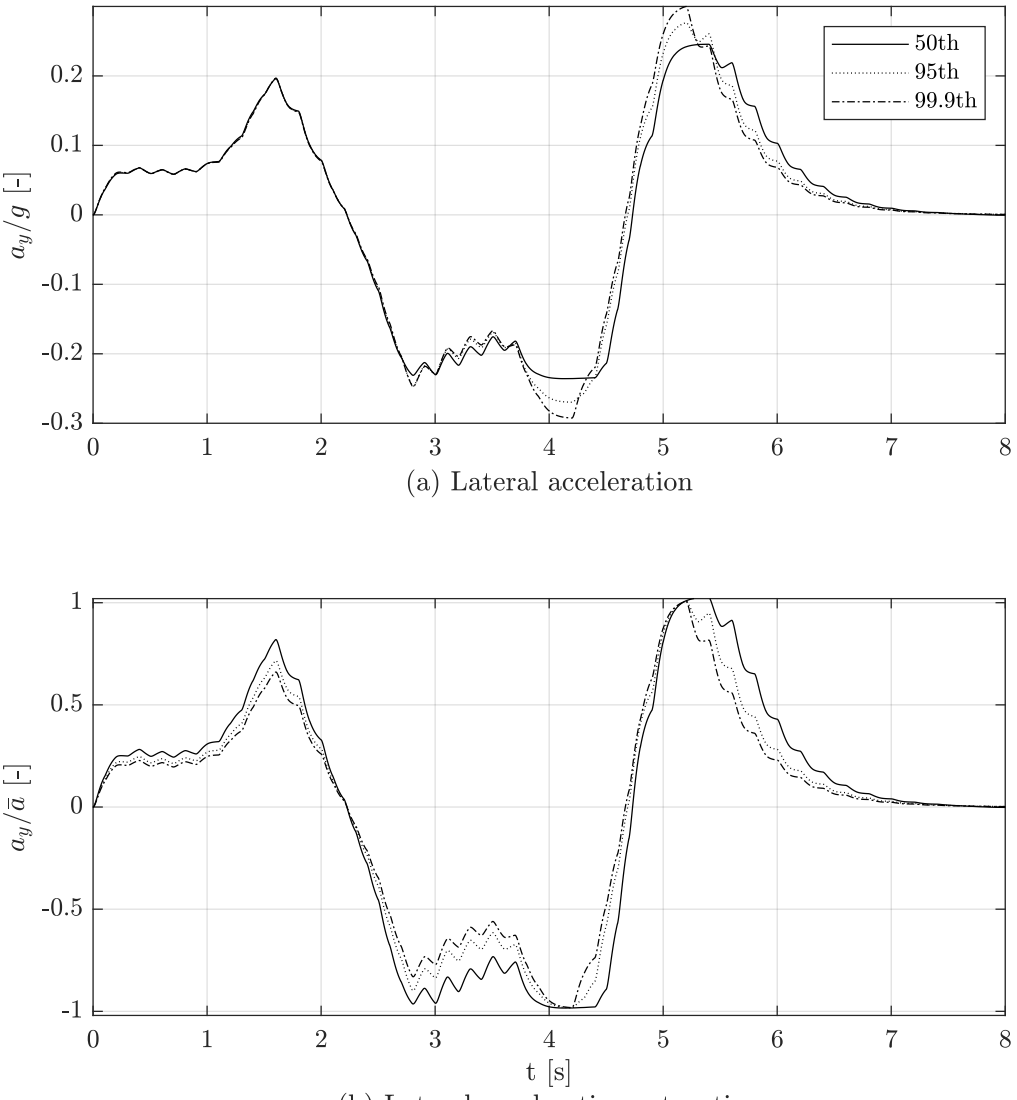


Figure 6.12: Lateral acceleration and saturation of the HMPC controller with acceleration penalty

Table 6.10: RMS static tracking error for controller with acceleration penalty for front-wheel steering vehicle with constant longitudinal velocity of 65 km/h

Driver percentile	v [m/s]	Y [cm]	ω_z [°/s]	ψ [°]
50th	0.0489	3.6899	1.8876	0.7885
95th	0.0483	2.5638	1.3942	0.5567
99.9th	0.0481	2.1126	1.1853	0.4695

Table 6.11: Summary of acceleration data for the profile run at 65 km/h

Driver percentile	$a_{y,\,\mathrm{RMS}}/g$ [-]	$a_{y,\,\mathrm{RMS}}/\bar{a}$ [%]	$\|a_y\|_\infty/\bar{a}$ [%]
50th	0.141	58.8	102.4
95th	0.142	52.0	100.4
99.9th	0.143	48.2	99.9

6.6 Conclusions

The results in this chapter have shown that it is simple to include human preference derived acceleration parameters directly into the model predictive control problem formulation, to change autonomous driving behaviour.

Firstly, the lateral acceleration constraints, derived from driving study data, were included. It was shown that the vehicle equipped with the 50th percentile acceleration constraint, had a larger static RMS error than the other profiles. Here, the static RMS error was introduced to measure how the vehicle drifted from its reference over time. The tracking performance for the median acceleration percentile was worse than the 95th and 99.9th percentiles as its acceleration bound, lower than the others, saturated during the presented double lane change manoeuvre. For example, the static RMS Y error was 2.8 times larger than the 99.9th percentile results. This provides an example of when it is beneficial to include, as a constraint, a human factor derived limit on acceleration. It showed that the controller, being able to detect that a constraint was to be violated, traded reference tracking performance for preferred human behaviour, i.e. lower acceleration.

This method of incorporating human factors as constraints does have potential limitations however. If, for a given driving scenario, the controller

can navigate without saturating the acceleration constraint, then the vehicle performance would be the same for all acceleration percentiles. This scenario would have occurred if the double lane change manoeuvre was performed at a lower speed, for example.

The lack of differentiation in driving styles occurs because of how constraints are handled in an optimisation problem; if the constraints are satisfied, then no further consideration is given to them, as 'just' satisfying a constraint is as preferable to a constraint with a lot of capacity or slack. Therefore, in the above scenario, the optimal control sequence would be the same, regardless of the acceleration profile.

Another potential issue of including human factor issues as constraints only was then highlighted, by creating a minimum acceleration path in the HMPC framework. It was shown then that tracking performance was similar for all acceleration profiles, with the RMS static error for lateral displacement being only 1.2 times larger than for the 50th percentile results compared to the 99.9th percentile. This was due to the optimal reference path also being the lowest acceleration path. Therefore, the acceleration constraint does not affect driving behaviour, only control feasibility. Regardless of acceleration limit, the controller would track the reference until it is able to do so. In the presented double lane change scenario, the effect of the numerical value of the acceleration constraint was to change the longitudinal speed at which the manoeuvre becomes infeasible.

In order to provide a more varied driving behaviour for different driver profiles, the MPC formulation was modified so that a penalty was placed on the magnitude of lateral acceleration. Using the linear approximation of lateral acceleration developed in Chapter 4, as a constraint, was an apparent extension.

By formulation as a quadratic penalty, it was shown in Eq. (6.5) that it is easily incorporated into the QP problem underlying the MPC formulation. The next step was to determine a scalar weighting term so that the acceleration term has an appropriate influence on vehicle performance. Noting, studies on human driver performance, such as the European driv-

ing study by Bosetti et al. [49, 50], that there is a decrease in lateral acceleration tolerance at higher speeds and so it was decided to derive the penalty weight that increases with longitudinal velocity. By nature of the penalty formulation, there is an inherent quadratic penalty for increasing longitudinal velocity, u, as shown in Figure 6.6 (a). However, a common stability measure for MPC, the terminal state penalty matrix, had an adverse effect on increasing the acceleration penalty with increasing speed. This was shown in Figure 6.6 (b), where it was demonstrated that the relative influence of acceleration in the MPC cost function decreased at higher speeds.

To overcome this, the penalty function Eq. (6.7) was developed. It utilised a cubic penalty on longitudinal velocity and an exponential function on the modified-Levison parameter, a_0. This allowed for the HMPC framework to increase the acceleration penalty for all driver profiles, while allowing for drivers with a higher tolerance to acceleration to have a relatively lower penalty value.

The effectiveness of the acceleration penalty weighting function, and including acceleration in the MPC cost function itself, was demonstrated in tracking a lateral step change for the median, 95th and 99.9th percentile driver profiles. All three showed a reduction in peak acceleration, with none of them saturating their acceleration bound, which does occur without the penalty term in the cost function. The penalty function Eq. (6.7) also allowed for the 99.9th percentile to have a higher maximum acceleration, at 98% of its limit compared to the median profile which had a 92% saturation. This demonstrates that a design goal of the penalty function was successfully met; not only did the higher percentile profiles have a higher absolute acceleration, they also had a higher average acceleration saturation. This allows the HMPC controller to provide better tracking of state variables to the occupants with a higher tolerance of lateral acceleration in a driving scenario.

The complete HMPC performance, with dynamic path generation, was then investigated with the double lane change manoeuvre. Compared to

step change tracking, it was shown that the acceleration penalty was less effective in this scenario. This was due to the issues discussed earlier; as the given course constraints are 'just' feasible when the longitudinal velocity is 65 km/h, the controller has very little capability to minimise lateral acceleration while maintaining feasibility.

Overall, due to the ease of inclusion in the problem, including human factor preferences into the MPC cost function provide a convenient method of directly modifying the vehicle performance, when the given driving conditions allow for it.

Incorporating the human factor conditions as constraints provides for an effective method of preventing behaviour that vehicle occupants wish to avoid. Implementing the behaviour as hard constraints might be overly restrictive however, as it makes avoiding acceleration that is safe from a tyre grip perspective as important as other constraints such as remaining within the defined road bounds.

An additional drawback of enforcing the human factor derived behaviour as hard constraints arises when attempting to study the resulting vehicle behaviour: feasibility. As the assumption that the control problem remains feasible has to be made, care must be taken when designing the driving scenario to do ensure this assumption holds true. This was done in this chapter by careful selection of longitudinal velocity.

Using hard constraints also results in the counter-intuitive consequence that for the median profile controller, which implements a more conservative, and therefore assumed to be safer, driving style it is harder to guarantee feasibility from a control theory perspective.

As discussed earlier, when studying the effect of a minimum acceleration reference trajectory, the path generation problem has a large influence on the resulting vehicle behaviour and the ability of the ASC module to influence performance. It is apparent therefore, that the path generation problem should also be studied for its suitability as a method to include human factors in the HMPC framework. In doing so, the driver preferences become a property of the path as well as a property of the autonomous ve-

hicle. This is discussed of the Future Work section of the next chapter.

Chapter 7

Conclusions and Future Work

The research in this thesis was based on the question 'Is it possible to include human factor preferences directly in the control problem for autonomous driving functionality?'. This naturally leads to further questions; namely what control techniques would be suitable for such a task? What effect does it have on vehicle performance for a given driving task? And, how does the inclusion of such human factors change the control specification and implementation.

This thesis provides an investigation into these questions, by focussing on a Hierarchical Model Predictive Control (HMPC) implementation of Active Steering Control (ASC) with dynamic path planning, which included human factor models of acceleration explicitly in the HMPC framework. This proposed framework was evaluated in simulation with a double lane change manoeuvre at motorway speeds used to observe the effect these human factor constraints had on vehicle tracking performance. In this chapter, conclusions are drawn from the results and the original contributions to the state-of-the-art knowledge are raised. Finally, based on this presented research, recommendations for future work are discussed.

7.1 Original Contributions to Knowledge

This thesis presents a control scheme that incorporates established human factor models explicitly into an advanced Model Predictive Control (MPC) formulation, demonstrating a convenient and effective method of changing autonomous driving behaviour in response to individual human preference.

In particular, the work developed in the thesis that, in the author's opinion, represent original contributions to knowledge is in the following areas:

- A demonstration through simulation that bicycle models are effective simulation tools for high-speed manoeuvres commonly used to evaluate motorway driving scenarios.
- An effective, linear, approximation of lateral acceleration is incorporated into an MPC scheme, allowing for the transient lateral acceleration of the vehicle to be accurately constrained.
- Human Factor derived acceleration models incorporated directly into an MPC scheme, allowing for vehicle behaviour to be changed by using existing models of human behaviour.
- A detailed study of the impact of human factor derived constraints on the dynamic tracking ability of a vehicle navigating a double lane change manoeuvre.
- A novel penalty function is developed to incorporate a dynamic penalty on lateral acceleration, to enhance effectiveness when using human factor models over a range of vehicle speeds.

These contributions are discussed in turn below.

Bicycle models are effective simulation tools for high-speed manoeuvres

It was shown that for moderately-high (0.5 g) acceleration manoeuvres, a non-linear bicycle model provides comparable performance as a simulation

model to much more complex models, such as the multi-DOF model presented in Chapter 3 and Appendix A. In Section 3.3.4, it was shown that the Root Mean Square (RMS) error between the non-linear bicycle and the multi-Degree-of-Freedom (DOF) was 0.03 g, while maximum lateral acceleration error during a sine-dwell manoeuvre was less than 8%.

By incorporating nonlinearities through the tyre model, it is possible to model the saturation in cornering force that occurs in high acceleration manoeuvres, limiting the vehicles lateral response and subsequent path tracking capability. The nonlinear bicycle model is a much faster model to use in simulation, due to the lack of implicit terms that arise in multi-DOF model, which require it to be solved as a system of Differential Algebraic Equations (DAEs).

This enables researchers and control designers to implement effective and time-efficient simulation tools for the modelling of motorway-speed manoeuvres.

An effective, linear, approximation of lateral acceleration in MPC

This allows for constraints on a key aspect of vehicle dynamic behaviour to be specified in a linear MPC formulation. In Section 4.2.2, it was shown that during a step change manoeuvre, the maximum acceleration of the nonlinear bicycle model was successfully controlled via MPC using the linear approximation of lateral acceleration. The simulated lateral acceleration only exceeded the constraint once during the lane change manoeuvre, when the constraint was violated by 0.5%. For typical driving scenarios this is an acceptable violation, and the constraint bound can be reduced accordingly if necessary.

As real-time considerations can be a limiting factor for implementation, this method of constraining a dominating factor of system response (acceleration) allows for QP representation and dramatically decreases the cost, by avoiding computational performance barriers for effective implementation in hardware.

Human Factor derived acceleration models incorporated directly into an MPC scheme

By doing so, this thesis demonstrates the potential to leverage large amounts of existing work in human modelling directly into the control scheme of autonomous behaviour. Instead of existing work that attempts to model the human driver as a part of the closed-loop control system, this work shows how human factor models can be used to shape autonomous driving behaviour and without attempting to replicate how humans perform the driving task. This allows for the controller to operate in a way that is not constrained by inherent limitations of human behaviour.

In Chapter 6, the modified-Levison criterion, presented in [49, 50], was used to demonstrate the effect of including human-factor derived constraints. It was used to generate three different driver profiles (50th, 95th and 99.9th percentile), each with a different acceleration bound. These were used to investigate the effect this had on a vehicle navigating a double lane change manoeuvre.

The impact of human factor derived constraints on the dynamic tracking ability of a vehicle was investigated for a double lane change manoeuvre

A systematic investigation into the effect of lateral acceleration was presented. The study focused on the behaviour of the vehicle at the controllers transition regime between feasible and infeasible control sequences, which represents the critical behaviour when evaluating performance.

It was shown that the vehicle equipped with the 50th percentile acceleration constraint, obtained by using Eq. (6.3), had a larger static RMS error than the other profiles. The tracking performance for the median acceleration percentile was worse than the 95th and 99.9th percentiles as its acceleration bound, lower than the others, saturated during the presented double lane change manoeuvre, resulting in a lower curvature path further from the reference trajectory. The static RMS Y error of the median profile was 2.8 times larger than the 99.9th percentile results.

Despite these difference in lateral displacement tracking, the effectiveness in changing behaviour was limited. At lower speeds, vehicle behaviour is unchanged for a different acceleration percentile. It was also shown that when tracking certain types of state-reference trajectories, such as for minimum path acceleration, specifying human preferable behaviours as constraints, in the low-level ASC module of the HMPC framework, is ineffective. This is because the constraint in this instance only changes the speed at which infeasibility occurs (for different acceleration profiles).

These conclusions show that while the control framework can easily incorporate existing models of human behaviour, MPC constraints do not provide a complete solution to changing vehicle behaviour as they are only active at boundaries of vehicle behaviour. In normal driving conditions, vehicle behaviour is not constantly held at a constraining boundary, and control designers must consider other methods to ensure the autonomous driving behaviour is acceptable. This motivates the next development.

Novel penalty function developed to incorporate a dynamic penalty on lateral acceleration

The influence of incorporating a penalty on lateral acceleration, in addition to a constraint, into the MPC problem was investigated. The effect that this has on the properties of the Quadratic Programming (QP) underlying MPC was shown.

From this, a penalty model was created, so that the active steering controllers behaviour adapted to both vehicle speed and human tolerance of acceleration. It was shown that this allowed for occupants with a greater acceleration tolerance to obtain better tracking of the state reference trajectories. At the same time, the approach is also reducing the peak acceleration, relative to the acceleration constraint limit, for people with a lower acceleration tolerance. In Section 6.5.2, this was demonstrated first by tracking a step change manoeuvre, in which it was shown that the 99.9th percentile profile had an acceleration saturation value 7% higher than the 50th percentile profile.

For manoeuvres that are limited by other constraints, such as road bounds, the acceleration constraint is again limited in its ability to change driving behaviour. It was shown that during a double lane change manoeuvre, the lateral acceleration was effectively the same for the three profiles.

The penalty was designed to ensure that the underlying QP formulation was preserved, allowing for the benefits with regards to complexity and performance to be kept. This allows for existing implementations to be modified to change their performance without requiring substantial modifications to the underlying framework or design paradigm.

7.2 Future Work

The final objective of this thesis is to identify further areas of research, to expand upon the work already presented. This it done by organising future work into the following categories:

- further study of how the human factor derived preferences can be included into the HMPC formulation
- extension of the control model to non-constant speed scenarios
- incorporation of different human factors, other than lateral acceleration, into the framework
- MPC robustness and stability guarantees when including model uncertainty.

It was shown that enforcing human factor considerations as hard constraints in the MPC formulation can cause issues when creating test scenarios and evaluating the simulation results. This was because if the manoeuvre was too 'easy', then regardless of acceleration parameter, the vehicle would be able to track the reference trajectory and the driving behaviour is the same. If the constructed manoeuvre is too challenging, then it will not be possible for one or more of the controllers to complete it and therefore comparison is not possible.

It also has the drawback that it makes a vehicle with a lower human tolerance value less capable than others. From the perspective of design intent, this is not ideal as it could trigger additional safety measures for one profile and not the other, resulting in a fundamentally different control approach due to a small change in human acceleration preference.

In Section 6.2, some different methods were outlined for including human factors. These methods, among others, should be investigated using the same analysis that was presented in the thesis, so that researchers and control designers have a more complete overview of the effect of the implementation issues available to them.

A second issue with regards to the proposed HMPC framework, is that the application of human factor design goals was not considered in the path generation problem. From the work of Michon and MacAdam [55, 59], among others, it is clear that the human driver, and therefore an individuals specific preference, is involved in the short term trajectory planning of the vehicle. It was also shown in Section 6.4.4 that the reference trajectory provided to the low-level ASC module had a large influence on the framework's ability to provide different driving styles. Therefore, performing a parameter study on the effect of changing the weighting values and the resulting acceleration response of the vehicle are an obvious extension. One possible method of changing vehicle behaviour is modification of the parameter weights in the path generation problem, which control the emphasis that is put on path velocity (which influences the shortest distance) and acceleration.

In the path generation problem presented, the reference path is only constrained by the road bounds (apart from the initial coordinates so that the path matches the vehicles initial heading). By incorporating curvature constraints on the optimal reference path, it is possible to, for a given vehicle speed, place a constraint on the paths acceleration. A natural starting point, based on the work presented in this thesis, would be to impose a curvature bound derived from the modified-Levison criterion. Again, as with the ASC module, the constraints raise issues of feasibility for individ-

ual human preferences in certain driving scenarios. As the path generation module has a larger preview horizon than the steering controller, it is possible that path infeasibility could be used as a way of determining that the vehicle speed needs to be reduced.

In the work of Bosetti et al. [49], building on the earlier work of Levison [53] and Herrin [48], the observation was made that human drivers perform a speed-acceleration trade-off to navigate bends. Due to the constant speed analysis presented here, it is not possible to investigate the effect that this has on the effect this trade-off has on control performance.

To investigate this, the HMPC framework should be extended to handle non-constant longitudinal velocity, u. Using a nonlinear bicycle model, with additional states u and X, it is possible study this (and allows for the full nonlinear coordinate transform from xy to XY). This has added complexity and implementations issues, particularly as the model is non-convex. Methods such as successive linearisation to pose the problem in Linear-Time Varying (LTV)-MPC form are problem are promising, along with the low-complexity approach proposed in [79].

Extending the framework to the non-constant speed case requires modification or replacement of the path-planning module. Due to the general structure of path-planning in 2D space, the resulting optimisation problem is non-convex due to the presence of obstacles and road boundary constraints. Sequential Quadratic Programming (SQP) is often used to determine the path, although the result can be sub-optimal. Solvers exist that take advantage of the geometric properties of the path planning problem to obtain solutions quickly [124, 125], allowing for effective real-time implementation. Because of these technologies, there is an existing methodology to investigate the effect of incorporating human factors in non-constant speed scenarios.

In addition, using a method such as this allows for a wider range of scenarios to be considered. For the work presented in this thesis, the standard double lane change manoeuvre was sufficient to demonstrate how the proposed control framework changed vehicle behaviour. For a more complete

understanding of the controllers performance across a vehicles operating envelope, it is necessary to use different driving scenarios. These could include lower speed driving scenarios to simulate urban driving or path following curves at higher speeds to study how the controller affects vehicle behaviour during rural road driving.

Along with acceleration constraints, other ride comfort factors should be included into the HMPC framework. For example, it is known that along with acceleration, jerk is an important aspect of comfort [39]. In the same way that the lateral acceleration was incorporated as both a constraint and as part of the cost function in Chapter 6, the MPC controller could be extended to constrain or minimise the maximum lateral jerk. Other human factor issues, such as roll oscillation [38], would require the use of a vehicle plant that can model the relevant phenomena.

In the work presented in this thesis, it was only necessary to consider the effect the human factor issues had on the terminal state penalty matrix. For the physical implementation of the control framework in prototype hardware, additional stability methods [71, 86], such as tube based methods [126] and the effect that the human factor terms have on them need to be characterised.

Furthermore, in the MPC formulation, it was assumed that both the system plant state was known perfectly; that is, that there is no uncertainty in the controller. This cannot be achieved in a real word implementation, and so the effects of this uncertainty on the proposed human factor based control scheme should be studied. Fortunately, there is already a significant body of work how uncertainty can be characterised and controlled via tractable, robust, MPC solutions and their underlying optimisation routine [89, 96, 127–129]. These methods should be evaluated to determine their applicability to the human factor based controller.

Bibliography

[1] *Connected and Autonomous Vehicles*. Tech. rep. The Society of Motor Manufacturers and Traders, 2019.

[2] T. Gordon and M. Lidberg. "Automated driving and autonomous functions on road vehicles". In: *Vehicle System Dynamics* 53.7 (2015), pp. 958–994. ISSN: 0042-3114. DOI: 10.1080/00423114.2015.1037774.

[3] C. Wallbank, L. Lloyd, J. Scoons, M. Muirhead, M. McCarthy, J. Carroll, and K. McRae-McKee. *Vehicle Safety Design Features and Future Safety Benefits in London*. Tech. rep. Transport Research Laboratory, 2015.

[4] "REGULATION (EC) No 661/2009 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 13 July 2009 concerning type-approval requirements for the general safety of motor vehicles, their trailers and systems, components and separate technical units intended therefor". In: *Official Journal of the European Union* L 200 (2009), pp. 1–24.

[5] S. A. Ferguson. "The effectiveness of electronic stability control in reducing real-world crashes: A literature review". In: *Traffic Injury Prevention* 8.4 (2007), pp. 329–338. ISSN: 15389588. DOI: 10.1080/15389580701588949.

[6] D. Piyabongkarn, J. Y. Lew, R. Rajamani, and J. A. Grogg. "Active driveline torque-management systems". In: *IEEE Control Systems Magazine* 30.4 (Aug. 2010), pp. 86–102. DOI: 10.1109/MCS.2010.937005.

[7] *Software Updates*. URL: https://www.tesla.com/en%7B%5C_%7DGB/support/software-updates (visited on 11/04/2019).

[8] P. E. Ross. *Tesla Model 3 to Driver: Look Me in the Eye*. 2017. URL: https://spectrum.ieee.org/cars-that-think/transportation/self-driving/tesla-model-3-to-driver-look-me-in-the-eye (visited on 11/04/2019).

[9] E. Woollacott. *Could a hacker hijack your connected car?* 2017. URL: https://www.bbc.co.uk/news/business-41367214 (visited on 11/04/2019).

[10] G. Topham. *Jaguar to supply 20,000 cars to Google's self-driving spin-off Waymo*. 2018. URL: https://www.theguardian.com/technology/2018/mar/27/waymo-self-driving-taxis-jaguar-land-rover (visited on 11/04/2019).

[11] J. M. Gitlin. *A cab ride in a Jaguar? Waymo will go classy for 20,000 electric robotaxis*. 2018. URL: https://arstechnica.com/cars/2018/03/a-cab-ride-in-a-jaguar-waymo-will-go-classy-for-20000-electric-robotaxis/ (visited on 11/04/2019).

[12] *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. Tech. rep. SAE International, 2018. DOI: https://doi.org/10.4271/J3016_201806.

[13] P. E. Ross. *The Audi A8: the World's First Production Car to Achieve Level 3 Autonomy*. 2017. URL: https://spectrum.ieee.org/cars-that-think/transportation/self-driving/the-audi-a8-the-worlds-first-production-car-to-achieve-level-3-autonomy (visited on 11/04/2019).

[14] J. M. Gitlin. *Tesla announces $35,000 Model 3, is closing its stores to pay for it*. 2019. URL: https://arstechnica.com/cars/2019/02/tesla-announces-35000-model-3-is-closing-its-stores-to-pay-for-it/ (visited on 11/06/2019).

[15] T. B. Lee. *Elon Musk says Tesla is "vastly ahead" on self-driving*. 2019. URL: https://arstechnica.com/cars/2019/04/elon-musk-says-tesla-is-vastly-ahead-on-self-driving/?comments=1%7B%5C&%7Dpost=37200121 (visited on 11/06/2019).

[16] K. Naughton. *Ford CEO Tamps Down Expectations for First Autonomous Vehicles*. 2019. URL: https://www.bloomberg.com/news/articles/2019-04-09/ford-ceo-tamps-down-expectations-for-first-autonomous-vehicles (visited on 11/06/2019).

[17] J. M. Gitlin. *Musk's remarks at conference imply Tesla has huge autonomous car advantage*. 2016. URL: https://arstechnica.com/cars/2016/06/musks-remarks-at-conference-imply-tesla-has-huge-autonomous-car-advantage/ (visited on 11/06/2019).

[18] R. E. Llaneras, J. Salinger, and C. a. Green. "Human factors issues associated with limited ability autonomous driving systems: Drivers' allocation of visual attention to the forward roadway". In: *PROCEEDINGS of the Seventh International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle Design* (2013), pp. 92–98. ISSN: 1057-8153.

[19] S. E. Shladover, C. A. Desoer, J. K. Hedrick, M. Tomizuka, J. Walrand, W. B. Zhang, D. H. McMahon, H. Peng, S. Sheikholeslam, and N. McKeown. "Automatic Vehicle Control Developments in the PATH Program". In: *IEEE Transactions on Vehicular Technology* 40.1 (1991), pp. 114–130. ISSN: 19399359. DOI: 10.1109/25.69979.

[20] S. E. Shladover. "PATH at 20 - History and major milestones". In: *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC* 8.4 (2006), pp. 1–8. ISSN: 15249050. DOI: 10.1109/TITS.2007.903052.

[21] A. Joshi and M. R. James. "Generation of accurate lane-level maps from coarse prior maps and lidar". In: *IEEE Intelligent Transportation Systems Magazine* 7.1 (2015), pp. 19–29. ISSN: 19391390. DOI: 10.1109/MITS.2014.2364081.

[22] J. Hellåker. "PROMETHEUS - Strategy". In: *Vehicle Electronics in the 90's: Proceedings of the International Congress on Transportation Electronics*. Oct. 1990, pp. 195–199. DOI: 10.1109/ICTE.1990.713012.

[23] E. D. Dickmanns, R. Behringer, D. Dickmanns, T. Hildebrandt, M. Maurer, F. Thomanek, and J. Schiehlen. "Seeing passenger car 'VaMoRs-P'". In: *Proceedings of the Intelligent Vehicles '94 Symposium*. 1994, pp. 68–73. DOI: 10.1109/IVS.1994.639472.

[24] E. Dickmanns and A. Zapp. "Autonomous High Speed Road Vehicle Guidance by Computer Vision". In: *IFAC W.* Vol. 20. 5. Elsevier, 1987, pp. 221–226. DOI: 10.1016/s1474-6670(17)55320-3.

[25] E. D. Dickmanns. "Vehicles capable of dynamic vision". In: *Artificial Intelligence* (1998), pp. 49–76. ISSN: 10450823.

[26] S. Thrun, M. Montemerlo, and A. Aron. "Probabilistic Terrain Analysis For High-Speed Desert Driving". In: *Proceedings of the Robotics Science and Systems Conference*. Ed. by G. Sukhatme, S. Schaal, W. Burgard, and D. Fox. Philadelphia, 2006. DOI: 10.15607/RSS.2006.II.021.

[27] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. N. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. M. Howard, S. Kolski, A. Kelly, M. Likhachev, M. McNaughton, N. Miller, K. Peterson, B. Pilnick, R. Rajkumar, P. Rybski, B. Salesky, Y.-W. Seo, S. Singh, J. Snider, A. Stentz, W. ". Whittaker, Z. Wolkowicki, J. Ziglar, H. Bae, T. Brown, D. Demitrish, B. Litkouhi, J. Nickolaou, V. Sadekar, W. Zhang, J. Struble, M. Taylor, M. Darms, and D. Ferguson. "Autonomous driving in urban environments: Boss and the Urban Challenge". In: *Journal of Field Robotics* 25.8 (2008), pp. 425–466. DOI: 10.1002/rob.20255.

[28] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, T. Dang, U. Franke, N. Appenrodt, C. G. Keller, E. Kaus, R. G. Herrtwich, C. Rabe, D. Pfeiffer, F. Lindner, F. Stein, F. Erbs, M. Enzweiler, C. Knoppel, J. Hipp, M. Haueis, M. Trepte, C. Brenk, A. Tamke, M. Ghanaat, M. Braun, A. Joos, H. Fritz, H. Mock, M. Hein, and E. Zeeb. "Making bertha drive-an autonomous journey on a historic route". In: *IEEE Intelligent Transportation Systems Magazine* 6.2 (2014), pp. 8–20. ISSN: 15249050. DOI: 10.1109/MITS.2014.2306552.

[29] C. Liu, C.-Y. Lin, and M. Tomizuka. "The Convex Feasible Set Algorithm for Real Time Optimization in Motion Planning". In: *SIAM Journal on Control Optimisation* 56.4 (2018), pp. 2712–2733.

[30] A. Geiger, M. Lauer, F. Moosmann, B. Ranft, H. Rapp, C. Stiller, and J. Ziegler. "Team AnnieWAY's Entry to the 2011 Grand Cooperative Driving Challenge". In: *IEEE Transactions on Intelligent Transportation Systems* 13.3 (2012), pp. 1008–1017. ISSN: 1524-9050. DOI: 10.1109/tits.2012.2189882.

[31] J. Guanetti, Y. Kim, and F. Borrelli. "Control of Connected and Automated Vehicles: State of the Art and Future Challenges". In: *Annual Reviews in Control* March (2018), pp. 0–1. ISSN: 13675788. DOI: 10.1016/j.arcontrol.2018.04.011. arXiv: 1804.03757.

[32] M. J. Griffin. *Handbook of Human Vibration*. First. London: Academic Press, 1990.

[33] F. Scheibe and M. C. Smith. "Analytical solutions for optimal ride comfort and tyre grip for passive vehicle suspensions". In: *Vehicle System Dynamics* 47.10 (2009), pp. 1229–1252. ISSN: 00423114. DOI: 10.1080/00423110802588323.

[34] J. K. Hedrick and T. Butsuen. "Invariant properties of automotive suspensions". In: *Proceedings of the Institution of Mechanical Engineers. Part D, Transport engineering* 204.1 (1990), pp. 21–27. ISSN: 02651904. DOI: 10.1243/pime_proc_1990_204_128_02.

[35] M. C. Smith and F. U. Wang. "Performance benefits in passive vehicle suspensions employing inerters". In: *Vehicle System Dynamics* 42.4 (2004), pp. 235–257. ISSN: 00423114. DOI: 10.1080/00423110412331289871.

[36] BSI. *BS ISO 2631-1:1997 Mechanical vibration and shock - Evaluation of human exposure to whole-body vibration - Part 1: General requirements.* 1997.

[37] BSI. *BS 6841:1987 Guide to measurement and evaluation of human exposure to whole-body mechanical vibration and repeated shock*. 1987.

[38] I. H. Wyllie and M. J. Griffin. "Discomfort from sinusoidal oscillation in the roll and lateral axes at frequencies between 0.2 and 1.6Hz". In: *The Journal of the Acoustical Society of America* 121.5 (2007), pp. 2644–2654. ISSN: 0001-4966. DOI: 10.1121/1.2715654.

[39] K. Strandemar. "On Objective Measures for Ride Comfort Evaluation". Technical Licentiate. KTH Royal Institute of Technology, 2005. ISBN: 9171782044.

[40] Q. Huang and H. Wang. "Fundamental study of jerk: Evaluation of shift quality and ride comfort". In: *SAE Technical Papers* (2004). ISSN: 01487191. DOI: 10.4271/2004-01-2065.

[41] J. Eriksson and L. Svensson. "Tuning for Ride Quality in Autonomous Vehicle: Application to Linear Quadratic Path Planning Algorithm". Master's thesis. Uppsala University, 2015.

[42] J. Förstberg. "Ride comfort and motion sickness in tilting trains: Human responses to motion environments in train and simulator experiments". Doctoral Thesis. KTH Royal Institute of Technology, 2000.

[43] A. S. Kilinç and T. Baybura. "Determination of Minimum Horizontal Curve Radius Used in the Design of Transportation Structures, Depending on the Limit Value of Comfort Criterion Lateral Jerk". In: *FIG Working Week 2012*. May. Rome, 2012, pp. 6–10.

[44] F. Jauberteau and J. L. Jauberteau. "Numerical differentiation with noisy signal". In: *Applied Mathematics and Computation* 215.6 (2009), pp. 2283–2297. ISSN: 00963003. DOI: 10.1016/j.amc.2009.08.042.

[45] A. Taragin and L. Leisch. "Driver Performance on Horizontal Curves". In: *Proceedings of the 33rd Annual Meeting of the Highway Research Board*. Washington, D.C., 1954, pp. 446–466.

[46] M. L. Ritchie, W. K. Mccoy, and W. L. Welde. "A Study of the Relation between Forward Velocity and Lateral Acceleration in Curves During Normal Driving". In: *Human Factors: The Journal of Human Factors and Ergonomics Society* 10.3 (1968), pp. 255–258. ISSN: 00187208. DOI: 10.1177/001872086801000307.

[47] M. L. Ritchie. "Choice of speed in driving through curves as a function of advisory speed and curve signs". In: *Human Factors* 14.6 (1972), pp. 533–538. ISSN: 00187208. DOI: 10.1177/001872087201400604.

[48] G. D. Herrin and J. B. Neuhardt. "An Empirical Model for Automobile Driver Horizontal Curve Negotiation". In: *Human Factors: The Journal of Human Factors and Ergonomics Society* 16.2 (1974), pp. 129–133. ISSN: 00187208. DOI: 10.1177/001872087401600203.

[49] P. Bosetti, M. Da Lio, and A. Saroldi. "On the human control of vehicles: An experimental study of acceleration". In: *European Transport Research Review* 6.2 (2014), pp. 157–170. ISSN: 18668887. DOI: 10.1007/s12544-013-0120-2.

[50] P. Bosetti, M. Da Lio, and A. Saroldi. "On Curve Negotiation: From Driver Support to Automation". In: *IEEE Transactions on Intelligent Transportation Systems* 16.4 (2015), pp. 2082–2093. ISSN: 15249050. DOI: 10.1109/TITS.2015.2395819.

[51] G. Reymond, A. Kemeny, J. Droulez, and A. Berthoz. "Role of Lateral Acceleration in Curve Driving: Driver Model and Experiments on a Real Vehicle and a Driving Simulator". In: *Human Factors: The Journal of the Human Factors and Ergonomics Society* 43.3 (2001), pp. 483–495. ISSN: 0018-7208. DOI: 10.1518/001872001775898188.

[52] W. H. Levison, A. C. Bittner, J. L. Campbell, and C. Schreiner. "Modification and partial validation of the Driver/Vehicle Module". In: *Transportation Research Record* 1803 (2002), pp. 52–58. ISSN: 03611981. DOI: 10.3141/1803-08.

[53] W. H. Levison. *Development of a Driver Vehicle Module for the Interactive Highway Safety Design Model*. Tech. rep. November. U.S. Department of Transportation, 2007.

[54] *interactIVe - Accident avoidance by active intervention for Inteligent Vehicles*. URL: http://interactive-ip.eu (visited on 11/18/2019).

[55] C. C. MacAdam. "Understanding and Modeling the Human Driver". In: *Vehicle System Dynamics* 40:1-3.January 2014 (2003), pp. 101–134. ISSN: 0042-3114. DOI: 10.1076/vesd.40.1.101.15875.

[56] M. Plöchl and J. Edelmann. "Driver models in automobile dynamics application". In: *User Modeling and User-Adapted Interaction* 45.7-8 (2007), pp. 699–741. ISSN: 17445159. DOI: 10.1080/00423110701432482.

[57] C. J. Nash and D. J. Cole. "Identification and validation of a driver steering control model incorporating human sensory dynamics". In: *Vehicle System Dynamics* 0.0 (2019), pp. 1–23. ISSN: 17445159. DOI: 10.1080/00423114.2019.1589536.

[58] J. A. Michon. "A critical view of driver behavior models: what do we know, what should we do?" In: *Human behavior and traffic safety*. Ed. by L. Evans and R. C. Schwing. New York: Plenum Press, 1985, pp. 485–520. ISBN: 978-1-4612-9280-7. DOI: 10.1007/978-1-4613-2173-6.

[59] J. A. Michon. "The Mutual Impacts of Transportation and Human Behaviour". In: *Transportation Planning for a Better Environment*. Ed. by P. Stringer and H. Wenzel. Boston, MA: Springer US, 1976, pp. 221–235. ISBN: 978-1-4615-8861-0. DOI: 10.1007/978-1-4615-8861-0_18.

[60] J. A. Michon. *Dealing With Danger*. Tech. rep. Gronigen, The Netherlands: Traffic Research Centre of the University of Groningen, 1978, pp. 157–171. DOI: 10.1177/101269099034002005.

[61] W. H. Janssen. *Route planning and guidance: A literature study (in Dutch)*. Tech. rep. Soesterberg, The Netherlands: Institute for Perception TNO, 1979.

[62] E. Hollnagel. "Cognition as control: A pragmatic approach to the modelling of joint cognitive systems". In: *Control* 9.October 2002 (2002), pp. 1–23.

[63] G. Hesslow. "The current status of the simulation theory of cognition". In: *Brain Research* 1428 (2012), pp. 71–79. ISSN: 00068993. DOI: 10.1016/j.brainres.2011.06.026.

[64] R. E. Kalman. "A New Approach to Linear Filtering and Prediction Problems". In: *Control Theory: Twenty-Five Seminal Papers (1932-1981)*. Ed. by T. Basar. Piscataway, NJ: IEEE Press, 2001, pp. 167–179.

[65] K. J. Aström and P. R. Kumar. "Control: A perspective". In: *Automatica* 50.1 (2014), pp. 3–43. ISSN: 00051098. DOI: 10.1016/j.automatica.2013.10.012.

[66] S. W. Sung, J. Lee, and I.-B. Lee. *Process Identification and PID Control*. Vol. 36. 8. John Wiley & Sons, Ltd., 2009. ISBN: 9780470824108. DOI: 10.1016/b978-1-85573-589-7.50015-1.

[67] D. Liberzon. "The Linear Quadratic Regulator". In: *Calculus of Variations and Optimal Control Theory - A Concise Introduction*. Princeton University Press, 2012. Chap. 6, pp. 180–199. ISBN: 978-0-691-15187-8.

[68] R. S. Sharp. "Driver steering control and a new perspective on car handling qualities". In: *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 219.10 (2005), pp. 1041–1051. ISSN: 09544062. DOI: 10.1243/095440605X31896.

[69] R. S. Sharp. "Optimal linear time-invariant preview steering control for motorcycles". In: *Vehicle System Dynamics* 44.SUPPL. 1 (2006), pp. 329–340. ISSN: 00423114. DOI: 10.1080/00423110600871509.

[70] R. S. Sharp and H. Peng. "Vehicle dynamics applications of optimal control theory". In: *Vehicle System Dynamics* 49.7 (2011), pp. 1073–1111. ISSN: 0042-3114. DOI: 10.1080/00423114.2011.586707.

[71] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert. "Constrained model predictive control: Stability and optimality". In: *Automatica* 36.6 (2000), pp. 789–814. ISSN: 00051098. DOI: 10.1016/S0005-1098(99)00214-9.

[72] J. H. Lee. "Model predictive control: Review of the three decades of development". In: *International Journal of Control, Automation and Systems* 9.3 (2011), pp. 415–424. ISSN: 15986446. DOI: 10.1007/s12555-011-0300-6.

[73] D. Hrovat, S. Di Cairano, H. Tseng, and I. Kolmanovsky. "The development of Model Predictive Control in automotive industry: A survey". In: *2012 IEEE International Conference on Control Applications*. 2012, pp. 295–302. ISBN: 978-1-4673-4505-7. DOI: 10.1109/CCA.2012.6402735.

[74] D. Hrovat and J. Sun. "Models and control methodologies for IC engine idle speed control design". In: *Control Engineering Practice* 5.8 (1997), pp. 1093–1100. ISSN: 09670661. DOI: 10.1016/S0967-0661(97)00101-9.

[75] Y. Wang and S. Boyd. "Fast Model Predictive Control Using Online Optimization". In: *IEEE Transactions on Control Systems Technology* 18.2 (2010), pp. 267–278. ISSN: 0952-3499. DOI: 10.1109/TCST.2009.2017934.

[76] D. J. Cole, A. J. Pick, and A. M. C. Odhams. "Predictive and linear quadratic methods for potential application to modelling driver steering control". In: *Vehicle System Dynamics* 44.3 (2006), pp. 259–284. ISSN: 0042-3114. DOI: 10.1080/00423110500260159.

[77] F. Borrelli, P. Falcone, T. Keviczky, J. Asgari, and D. Hrovat. "MPC–based approach to active steering for autonomous vehicle systems". In: *International Journal of Vehicle Autonomous Systems* 3 (2005), pp. 265–291. DOI: 10.1504/IJVAS.2005.008237.

[78] G. Lars and J. Pannek. *Nonlinear model predictive control: theory and algorithms*. 2nd. Communications and control engineering. Springer International Publishing, 2017. ISBN: 9783319460246.

[79] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat. "Predictive Active Steering Control for Autonomous Vehicle Systems". In: *IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY* 15.3 (2007), pp. 566–580. ISSN: 1063-6536. DOI: 10.1109/TCST.2007.894653. arXiv: arXiv:1604.07446v1.

[80] P. Falcone, F. Borrelli, H. E. Tseng, J. Asgari, and D. Hrovat. "A hierarchical Model Predictive Control framework for autonomous ground vehicles". In: *2008 American Control Conference*. 2008, pp. 3719–3724. ISBN: 978-1-4244-2078-0. DOI: 10.1109/ACC.2008.4587072.

[81] H. Yoshida, S. Shinohara, and M. Nagai. "Lane change steering manoeuvre using model predictive control theory". In: *Vehicle System Dynamics* 46.SUPPL.1 (2008), pp. 669–681. ISSN: 00423114. DOI: 10.1080/00423110802033072.

[82] T. Shim, G. Adireddy, and H. Yuan. "Autonomous vehicle collision avoidance system using path planning and model-predictive-control-based active front steering and wheel torque control". In: *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering* 226.6 (2012), pp. 767–778. ISSN: 09544070. DOI: 10.1177/0954407011430275.

[83] X. Na and D. J. Cole. "Linear quadratic game and non-cooperative predictive methods for potential application to modelling driver-AFS interactive steering control". In: *Vehicle System Dynamics* 51.2 (2013), pp. 165–198. ISSN: 00423114. DOI: 10.1080/00423114.2012.715653.

[84] J. Nilsson, P. Falcone, M. Ali, and J. Sjöberg. "Receding horizon maneuver generation for automated highway driving". In: *Control Engineering Practice* 41 (2015), pp. 124–133. ISSN: 09670661. DOI: 10.1016/j.conengprac.2015.04.006.

[85] *Convex Optimization: An Overview by Stephen Boyd: The 3rd Wook Hyun Kwon Lecture*. 2018. URL: https://youtu.be/U41e7hKAAPQ?t=3318 (visited on 11/12/2019).

[86] J. B. Rawlings, D. Q. Mayne, and M. Diehl. *Model predictive control*. 2nd Editio. Nob Hill Publishing, 2017, pp. 1–401.

[87] E. G. Gilbert and K. T. Tan. "Linear Systems with State and Control Constraints: The Theory and Application of Maximal Output Admissible Sets". In: *IEEE Transactions on Automatic Control* 36.9 (1991), pp. 1008–1020. ISSN: 15582523. DOI: 10.1109/9.83532.

[88] M. N. Zeilinger, C. N. Jones, D. M. Raimondo, and M. Morari. "Real-time MPC - Stability through robust MPC design". In: *Joint Proceedings of the IEEE Conference on Decision and Control and Chinese Control Conference*. Shanghai, China: IEEE, 2009, pp. 3980–3986. ISBN: 9781424438716. DOI: 10.1109/CDC.2009.5400903.

[89] M. V. Kothare, V. Balakrishnan, and M. Morari. "Robust constrained model predictive control using linear matrix inequalities". In: *Automatica* 32.10 (1996), pp. 1361–1379. ISSN: 00051098. DOI: 10.1016/0005-1098(96)00063-5. arXiv: arXiv:1011.1669v3.

[90] A. Bemporad and M. Morari. "Robust model predictive control: A survey". In: *Robustness in identification and control* 245 (1999), pp. 207–226. ISSN: 13624962. DOI: 10.1007/BFb0109870.

[91] D. Q. Mayne, M. M. Seron, and S. V. Raković. "Robust model predictive control of constrained linear systems with bounded disturbances". In: *Automatica* 41.2 (2005), pp. 219–224. ISSN: 00051098. DOI: 10.1016/j.automatica.2004.08.019.

[92] D. Q. Mayne. "Model predictive control: Recent developments and future promise". In: *Automatica* 50.12 (2014), pp. 2967–2986. ISSN: 00051098. DOI: 10.1016/j.automatica.2014.10.128.

[93] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge: Cambridge University Press, 2009, pp. 487–487. ISBN: 9780521833783. DOI: 10.1080/10556781003625177. arXiv: 1111.6189v1.

[94] Y. Nesterov and A. Nemirovski. *Interior-Point Polynomial Methods in Convex Programming*. Philadelphia: Society for Industrial and Applied Mathematics, 1994. ISBN: 978-0-89871-319-0.

[95] N. Karmarkar. "A New Polynomial-Time Algorithm for Linear Programming". In: *Combinatoria* 4.4 (1984), pp. 373–395. DOI: 10.1145/800057.808695.

[96] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. Vol. 15. Philadelphia: Society for Industrial and Applied Mathematics, 1994, p. 203. ISBN: 089871334X. DOI: 10.1109/TAC.1997.557595. arXiv: arXiv:1011.1669v3.

[97] S. Richter, C. N. Jones, and M. Morari. "Computational complexity certification for real-time MPC with input constraints based on the fast gradient method". In: *IEEE Transactions on Automatic Control* 57.6 (2012), pp. 1391–1403. ISSN: 00189286. DOI: 10.1109/TAC.2011.2176389.

[98] H. Pacejka. *Tire and Vehicle Dynamics*. 3rd. Oxford: Butterworth-Heinemann, 2012. ISBN: 978-0-08-097016-5.

[99] L. Segel. "Theoretical Prediction and Experimental Substantiation of the Response of the Automobile to Steering Control". In: *Proceedings of the Institution of Mechanical Engineers: Automobile Division* 10.1 (1956), pp. 310–330. ISSN: 0367-8822. DOI: 10.1243/pime_auto_1956_000_032_02.

[100] D. W. Whitcomb and W. F. Milliken. *Design implications of a general theory of automobile stability and control*. 1956. URL: http://pad.sagepub.com/lookup/doi/10.1243/PIME%7B%5C_%7DAUTO%7B%5C_%7D1956%7B%5C_%7D000%7B%5C_%7D035%7B%5C_%7D02.

[101] R. S. Sharp. "Some contemporary problems in road vehicle dynamics". In: *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering* 214.1 (2000), pp. 137–148. ISSN: 09544062. DOI: 10.1243/0954406001522877.

[102] T. Keviczky, P. Falcone, F. Borrelli, J. Asgari, and D. Hrovat. "Predictive control approach to autonomous vehicle steering". In: *Proceedings of the American Control Conference* 2006 (2006), pp. 4670–4675. ISSN: 07431619. DOI: 10.1109/acc.2006.1657458.

[103] R. P. Osborn and T. Shim. "Independent control of all-wheel-drive torque distribution". In: *Vehicle System Dynamics* 44.7 (2006), pp. 529–546. ISSN: 00423114. DOI: 10.1080/00423110500485731.

[104] J. W. Griffin. "Influences of Drive Torque Distribution on Road Vehicle Handling and Efficiency". PhD. University of Nottingham, 2015.

[105] A. J. Tremlett, M. Massaro, D. J. Purdy, E. Velenis, F. Assadian, A. P. Moore, and M. Halley. "Optimal control of motorsport differentials". In: *Vehicle System Dynamics* 53.12 (2015), pp. 1772–1794. ISSN: 17445159. DOI: 10.1080/00423114.2015.1093150.

[106] T. Shim and C. Ghike. "Understanding the limitations of different vehicle models for roll dynamics studies". In: *Vehicle System Dynamics* 45.3 (2007), pp. 191–216. ISSN: 17445159. DOI: 10.1080/00423110600882449.

[107] H. B. Pacejka and E. Bakker. "The magic formula tyre model". In: *Vehicle System Dynamics* 21.sup1 (1992), pp. 1–18. ISSN: 00423114. DOI: 10.1080/00423119208969994.

[108] F. Bruzelius. "A theoretical justification of the sine with dwell manoeuvre". In: *Vehicle System Dynamics* 53.4 (2015), pp. 493–505. ISSN: 17445159. DOI: 10.1080/00423114.2014.1002794.

[109] D. Hrovat. "Applications of optimal control to advanced automotive suspension design". In: *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME* 115.2B (1993), pp. 328–342. ISSN: 15289028. DOI: 10.1115/1.2899073.

[110] D. Hrovat. "Survey of Advanced Suspension Developments and Related Optimal Control Applications". In: *Automatica* 33.10 (1997), pp. 1781–1817. ISSN: 00051098. DOI: 10.1016/S0005-1098(97)00101-5.

[111] R. Rajamani. *Vehicle Dynamics and Control*. Second. Boston, MA: Springer, 2012. ISBN: 9781461414322. DOI: https://doi.org/10.1007/978-1-4614-1433-9.

[112] L. F. Shampine. "Solving 0 = F(t, y(t), y'(t)) in Matlab". In: *Journal of Numerical Mathematics* 10.4 (2002), pp. 291–310. ISSN: 1570-2820. DOI: 10.1515/JNMA.2002.291.

[113] BSI. *Passenger cars — Validation of vehicle dynamic simulation — Sine with dwell stability control testing*. Tech. rep. British Standards, 2016.

[114] C.-T. Chen. *Linear System Theory and Design*. 3rd. New York: Oxford University Press, 1999.

[115] L. F. Shampine and M. W. Reichelt. "The MATLAB ODE Suite". In: *SIAM Journal on Scientific Computing* 18.1 (1997), pp. 1–22. ISSN: 1064-8275. DOI: 10.1137/S1064827594276424.

[116] J. M. Maciejowski. *Predictive Control: with Constraints*. London: Prentice Hall, 2002.

[117] C. V. Rao, S. J. Wright, and J. B. Rawlings. "Application of interior-point methods to model predictive control". In: *Journal of Optimization Theory and Applications* 99.3 (1998), pp. 723–757. ISSN: 00223239. DOI: 10.1023/A:1021711402723.

[118] Y. Wang and S. Boyd. "Fast model predictive control using online optimization". In: *IEEE Transactions on Control Systems Technology* 18.2 (2010), pp. 267–278. ISSN: 10636536. DOI: 10.1109/TCST.2009.2017934.

[119] R. Scattolini. "Architectures for distributed and hierarchical Model Predictive Control - A review". In: *Journal of Process Control* 19.5 (2009), pp. 723–731. ISSN: 09591524. DOI: 10.1016/j.jprocont.2009.02.003.

[120] S. Chang and T. J. Gordon. "A flexible hierarchical model-based control methodology for vehicle active safety systems". In: *Vehicle System Dynamics* 46.SUPPL.1 (2008), pp. 63–75. ISSN: 00423114. DOI: 10.1080/00423110701882306.

[121] T. A. Ranney. "Models of Driving Behavior: A Review of Their Evolution". In: *Accident Analysis and Prevention* 26.6 (1994), pp. 733–750.

[122] J. S. Hellewell, A. A. Popov, and G. E. Burnett. "Hierarchical Control for Trajectory Generation and Tracking Via Active Front Steering". In: *Journal of Dynamic Systems, Measurement, and Control* 142.1 (2020), pp. 1–10. ISSN: 0022-0434. DOI: 10.1115/1.4044620.

[123] C. Liu, C. Y. Lin, Y. Wang, and M. Tomizuka. "Convex feasible set algorithm for constrained trajectory smoothing". In: *Proceedings of the American Control Conference* (2017), pp. 4177–4182. ISSN: 07431619. DOI: 10.23919/ACC.2017.7963597.

[124] J. Liu, P. Jayakumar, J. L. Stein, and T. Ersal. "A study on model fidelity for model predictive control-based obstacle avoidance in high-speed autonomous ground vehicles". In: *Vehicle System Dynamics* 54.11

(2016), pp. 1629–1650. ISSN: 17445159. DOI: 10.1080/00423114.2016.1223863.

[125] C. Liu and M. Tomizuka. "Real time trajectory optimization for nonlinear robotic systems: Relaxation and convexification". In: *Systems and Control Letters* 108 (2017), pp. 56–63. ISSN: 01676911. DOI: 10.1016/j.sysconle.2017.08.004.

[126] Y. Gao. "Model Predictive Control for Autonomous and Semiautonomous Vehicles". PhD. University of California, Berkeley, 2014. ISBN: 9788578110796. DOI: 10.1017/CBO9781107415324.004. arXiv: arXiv:1011.1669v3.

[127] B. Pluymers, J. A. Rossiter, J. A. Suykens, and B. De Moor. "A simple algorithm for robust MPC". In: *IFAC Proceedings Volumes (IFAC-PapersOnline)*. Vol. 38. 1. Prague: IFAC, 2005, pp. 257–262. ISBN: 008045108X. DOI: 10.3182/20050703-6-CZ-1902.00987.

[128] B. Kouvaritakis, M. Cannon, S. V. Raković, and Q. Cheng. "Explicit use of probabilistic distributions in linear predictive control". In: *Automatica* 46.10 (2010), pp. 1719–1724. ISSN: 00051098. DOI: 10.1016/j.automatica.2010.06.034.

[129] A. Ben-Tal and A. Nemirovski. "Robust optimization - Methodology and applications". In: *Mathematical Programming* 92.3 (2002), pp. 453–480. ISSN: 00255610. DOI: 10.1007/s101070100286.

[130] S. Liu and G. Trenkler. "Hadamard, Khatri-Rao, Kronecker and other matrix products". In: *Int. J. Inf. Syst. Sci* 4.1 (2008), pp. 160–177.

[131] S. Eddins and L. Shure. *Matrix Indexing in MATLAB*. 2001. URL: https://uk.mathworks.com/company/newsletters/articles/matrix-indexing-in-matlab.html (visited on 12/03/2019).

Appendix A

Multi-Degree-Of-Freedom Model

The multi-DOF model is derived from the 14 DOF model presented in [106]. Its constituent equations are presented here for completeness and to account for any differences in implementation. Verification plots, in addition to those presented in Chapter 3, are also provided. When it is possible to represent the model for all four wheel/suspension components, e.g. there is no change in sign, the equation is presented for a single wheel subsystem with subscript notation dropped for brevity.

A.1 Equations of Motion

Preliminaries and Notation

To begin, the following vector terms are defined for the linear, $\boldsymbol{v}$, and rotational, $\boldsymbol{\omega}$ velocity of the sprung mass:

$$\boldsymbol{v} = \begin{pmatrix} u \\ v \\ w \end{pmatrix}, \qquad \boldsymbol{\omega} = \begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix}. \tag{A.1}$$

The tyre indicator set is defined as follows:

$$\mathcal{I} = \{\text{‘fi’},\ \text{‘fo’},\ \text{‘ri’},\ \text{‘ro’}\}, \tag{A.2}$$

which refers to the 'front inside', 'front outside', 'rear inside' and 'rear outside' tyres or suspension components. The terms 'inside' and 'outside' refer to the vehicle's orientation for a positive steer-angle curve.

It will be convenient later to make use of the matrix-vector form of the vector cross-product. Given two vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ in $\mathbb{R}^3$, the standard vector cross-product, $\boldsymbol{a} \times \boldsymbol{b}$ can be computed via [130]:

$$\boldsymbol{a} \times \boldsymbol{b} = [\boldsymbol{a}]_\times \boldsymbol{b} \tag{A.3}$$

where $[\boldsymbol{a}]_\times$ is defined to be the following skew-symmetric matrix:

$$[\boldsymbol{a}]_\times := \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix}$$

Sprung Mass

The sprung mass of the vehicle is assumed to be a rigid body, with body-fixed coordinates, xyz, fixed at the centre of mass and aligned along the principle axes. The velocities of the sprung mass in the body-fixed frame are given by u, v and w respectively.

The six equations of motion for the vehicle body, taken about the centre of mass, are given by:

$$m\left(\dot{u} + \omega_y w - \omega_z v\right) = \sum F_{x\mathrm{s,k}} + mg\sin\theta \tag{A.4}$$

$$m\left(\dot{v} + \omega_z u - \omega_x w\right) = \sum F_{y\mathrm{s,k}} - mg\sin\phi\cos\theta \tag{A.5}$$

$$m\left(\dot{w} + \omega_x v - \omega_y u\right) = \sum F_{z\mathrm{s,k}} + F_{d\mathrm{s,k}} - mg\cos\phi\cos\theta \tag{A.6}$$

$$I_x\dot{\omega}_x = \sum M_{x\mathrm{k}} + \frac{F_{z,\mathrm{sfi}} - F_{z,\mathrm{sfo}}}{2t_f} + \frac{F_{z,\mathrm{sri}} - F_{z,\mathrm{sro}}}{2t_r} \tag{A.7}$$

$$I_y\dot{\omega}_y = \sum M_{y\mathrm{k}} - a\left(F_{z,\mathrm{sfi}} + F_{z,\mathrm{sfo}}\right) + b\left(F_{z,\mathrm{sri}} + F_{z,\mathrm{sro}}\right) \tag{A.8}$$

$$I_z\dot{\omega}_z = \sum M_{z\mathrm{k}} + a\left(F_{y,\mathrm{sfi}} + F_{y,\mathrm{sfo}}\right) - b\left(F_{y,\mathrm{sri}} + F_{y,\mathrm{sro}}\right) \tag{A.9}$$

$$- \frac{F_{x,\mathrm{sfi}} - F_{x,\mathrm{sfo}}}{2t_f} - \frac{F_{x,\mathrm{sri}} - F_{x,\mathrm{sro}}}{2t_r}, \quad k \in \mathcal{I} \tag{A.10}$$

where m is the sprung mass. The moments acting about the sprung mass are given by:

$$\begin{aligned} M_x &= F_{y,\mathrm{s}} h_{rc} \\ M_y &= -\left(F_{x,\mathrm{sg}} r + F_{x,\mathrm{s}} l_{\mathrm{s,i}}\right) \\ M_z &= 0. \end{aligned} \tag{A.11}$$

Velocity transform

To determine terms in the suspension system equation, it is necessary to determine the velocity of the vehicle body at the four strut points. The general form is:

$$\boldsymbol{v}_{\mathrm{s},k} = \mathbf{M}_k \boldsymbol{\omega} + \boldsymbol{v}, \qquad k \in \mathcal{I} \tag{A.12}$$

where $\mathbf{M}_k$ is unique for each corner of the vehicle body. Using Eq. (A.3), the matrix-vector product in Eq. (A.12) can be given as $\boldsymbol{a}_k \times \boldsymbol{\omega} = [\boldsymbol{a}_k]_\times \boldsymbol{\omega}$ with:

$$\begin{aligned} \boldsymbol{a}_{\mathrm{fi}} &= (0, \ -t_f/2, \ -a)^T \\ \boldsymbol{a}_{\mathrm{fo}} &= (0, \quad t_f/2, \ -a)^T \\ \boldsymbol{a}_{\mathrm{ri}} &= (0, \ -t_r/2, \quad b)^T \\ \boldsymbol{a}_{\mathrm{ro}} &= (0, \quad t_r/2, \quad b)^T \end{aligned} \tag{A.13}$$

Tyre-Ground Contact Patch Transform

The lateral and longitudinal ground forces in the body fixed coordinate frame are obtained by resolving the longitudinal and cornering forces at the tyre contact patch as:

$$\begin{pmatrix} F_{x\mathrm{g,k}} \\ F_{y\mathrm{g,k}} \end{pmatrix} = \begin{bmatrix} \cos\delta_k & -\sin\delta_k \\ \sin\delta_k & \cos\delta_k \end{bmatrix} \begin{pmatrix} F_{x\mathrm{t,k}} \\ F_{y\mathrm{t,k}} \end{pmatrix}, \qquad k \in \mathcal{I} \tag{A.14}$$

where δ_k is the steer angle of the tyre. In this thesis, it is assumed that the steer angle of the front two tyres is equal, and that the steer angle of the rear two tyres is zero.

Contact Patch Velocities

The contact patch velocities are given by:

$$u_{\mathrm{g},k} = \cos\theta\left(u_{\mathrm{u},k} - \omega_y \cdot r_k\right) + \sin\theta\left(\cos\phi \cdot w_{\mathrm{u},k} + \sin\phi\left(v_{\mathrm{u},k} + \omega_x \cdot r_k\right)\right) \tag{A.15}$$

$$v_{\mathrm{g},k} = \cos\phi\left(v_{\mathrm{u},k} + \omega_x \cdot r_k\right) - \sin\phi \cdot w_{\mathrm{u},k}, \qquad k \in \mathcal{I} \tag{A.16}$$

Tyre slip angles

From the contact patch velocities, the tyres steady state slip angles are given by:

$$\tan\alpha_{\mathrm{ss},k} = \frac{v_k}{u_k}, \qquad k \in \mathcal{I}. \tag{A.17}$$

As in Section 3.3.1, a first order transience model, Eq. (3.6), is used to obtain the apparent tyre slip angle used to calculate the lateral and longitudinal tyre forces, using the Magic Tyre Formula (MTF), Eq (3.4), using the parameter values in Table 3.1.

Unsprung Mass Velocities

The longitudinal and lateral velocities of the unsprung mass, in the body-fixed coordinate frame, is given by:

$$u_{\mathrm{u},k} = u_{\mathrm{s},k} - l_{\mathrm{s},k} \cdot \omega_y, \qquad k \in \mathcal{I} \tag{A.18}$$

$$v_{\mathrm{u},k} = v_{\mathrm{s},k} + l_{\mathrm{s},k} \cdot \omega_x, \qquad k \in \mathcal{I} \tag{A.19}$$

The vertical velocity of the unsprung mass is given by the integration of the following equation:

$$\begin{aligned}\dot{w}_{\mathrm{u},k} = \frac{1}{m_{\mathrm{u}}}\big(&\cos\phi\left(\cos\theta\left(F_{z\mathrm{g},k} - m_{\mathrm{u},k}g\right) + \sin\theta F_{x\mathrm{g},k}\right) - \sin\phi F_{y\mathrm{g},k} - F_{\mathrm{d}z,k} \\ &- k_{\mathrm{s},k} \cdot x_{\mathrm{s},k} - c_{\mathrm{s},k} \cdot \dot{x}_{\mathrm{s},k} - m_{\mathrm{u}}\left(v_{\mathrm{u},k} \cdot \omega_x - u_{\mathrm{u},k} \cdot \omega_y\right)\big), \qquad k \in \mathcal{I}\end{aligned} \tag{A.20}$$

Instantaneous Spring Deflection Rate

The suspension spring deflection rate is given by the difference between the vertical velocity of the unsprung mass and the corresponding velocity of the vehicle body at the strut mounting point:

$$\dot{x}_{\mathrm{s},k} = w_{\mathrm{u},k} - w_{\mathrm{s},k}, \qquad k \in \mathcal{I} \tag{A.21}$$

Suspension Strut Length

The instantaneous strut length is given by:

$$l_{\mathrm{s},k} = l_{\mathrm{si},k} - \left(x_{\mathrm{s},k} - x_{\mathrm{si},k}\right), \qquad k \in \mathcal{I} \tag{A.22}$$

where l_{si} is the initial strut length and x_{si} is the initial suspension spring deflection. The initial spring length is given by:

$$l_{\mathrm{si},k} = h - \left(r_0 - x_{\mathrm{t},k}\right), \qquad k \in \mathcal{I} \tag{A.23}$$

The initial suspension spring compression, x_{si}, and initial tyre compression, x_{ti}, for the front and rear axles are given by:

$$x_{\mathrm{si,f}} = \frac{mb}{2\left(a+b\right)k_{\mathrm{s,f}}}, \qquad x_{\mathrm{si,r}} = \frac{ma}{2\left(a+b\right)k_{\mathrm{s,f}}} \tag{A.24}$$

$$x_{\mathrm{ti,f}} = \frac{\left(mb/2\left(a+b\right)\right) + m_{\mathrm{u,f}}}{k_{\mathrm{t,f}}}, \qquad x_{\mathrm{ti,fr}} = \frac{\left(ma/2\left(a+b\right)\right) + m_{\mathrm{u,r}}}{k_{\mathrm{t,f}}} \tag{A.25}$$

Unsprung Mass Forces

The longitudinal and lateral forces transmitted to the unsprung mass through the suspension strut are determined by subtracting the unsprung weight and inertial forces from the respective forces in the tyre contact patch. The vertical forces are determined by the forces of the suspension spring and damper:

$$F_{xs,k} = F_{xgs,k} - m_{\mathrm{u}}\left(-\sin\theta + \dot{u}_{\mathrm{u},k} - \omega_z v_{\mathrm{u},k} + \omega_y w_{\mathrm{u},k}\right) \tag{A.26}$$

$$F_{ys,k} = F_{ygs,k} - m_{\mathrm{u}}\left(\sin\phi\cos\theta + \dot{v}_{\mathrm{u},k} - \omega_x w_{\mathrm{u},k} + \omega_z u_{\mathrm{u},k}\right) \tag{A.27}$$

$$F_{zs,k} = k_{s,k} \cdot x_{s,k} + c_{s,k} \cdot \dot{x}_{s,k}, \qquad k \in \mathcal{I} \tag{A.28}$$

The presence of the $\dot{u}$ and $\dot{v}$ terms are what makes the model a system of DAEs, rather than Ordinary Differential Equations (ODEs).

Jacking Forces

The term 'jacking forces' refer to the load transfer due to the roll centre [106]:

$$F_{dz,fi} = -F_{dz,fo} = \frac{F_{ygs,fi} r_{fi} + F_{ys,fi} l_{s,fi} + F_{ygs,fo} r_{fo} + F_{ys,fo} l_{s,fo} - h_{rcf} \left(F_{ys,fi} + F_{ys,fo}\right)}{t_f} \tag{A.29}$$

$$F_{dz,ri} = -F_{dz,ro} = \frac{F_{ygs,ri} r_{ri} + F_{ys,ri} l_{s,ri} + F_{ygs,ro} r_{ro} + F_{ys,ro} l_{s,ro} - h_{rcr} \left(F_{ys,ri} + F_{ys,ro}\right)}{t_r} \tag{A.30}$$

Instantaneous Tyre Deflection Rate

The tyre deflection rate is given by the difference between the vertical velocity at the tyre contact patch and the vertical velocity of the wheel centre in the inertial coordinate frame. For this thesis, the vertical velocity at the tyre contact patch is zero, due to the assumption that the road surface is perfectly smooth and flat. The instantaneous tyre deflection rate is therefore:

$$\dot{x}_{t,k} = -\cos\theta \left(w_{u,k} \cos\phi + v_{u,k} \sin\phi\right) + u_{u,k} \sin\theta, \qquad k \in \mathcal{I} \tag{A.31}$$

Tyre Radius

The instantaneous radius for a given tyre is:

$$r = \begin{cases} r_0 & x_t \leq 0 \\ \frac{r_0 - x_t}{\cos\phi \cos\theta} & \end{cases} \tag{A.32}$$

where x_t is determined by integration of Eq. (A.31).

Normal tyre force in the contact patch

The normal tyre force, acting in the tyre contact patch in coordinate frame two is given by:

$$F_{\mathrm{t},k} = k_{\mathrm{t},k} x_{\mathrm{t},k}, \qquad k \in \mathcal{I} \tag{A.33}$$

where k_{t} is the tyre stiffness, and x_{t} is the tyre deflection rate, given by the integration of Eq. (A.31).

Angular Acceleration for Wheel

The angular acceleration of a wheel is given by:

$$\dot{\omega}_k = \frac{T_{\mathrm{d},k} - r_k F_{\mathrm{xt},k}}{I_{\mathrm{w}}}, \qquad k \in \mathcal{I} \tag{A.34}$$

The torque provided to each wheel is given by Eq. (3.19), in Chapter 3.

Cardan Angles

To obtain the orientation of the vehicle in the global inertial coordinate frame, XYZ, the vehicle's body fixed frame is rotated first through the roll, ϕ, and pitch, θ, angles before being rotated through the yaw angle, ψ. The cardan angles ϕ, θ, ψ are obtained from the numerical integration of the following equations:

$$\begin{aligned} \dot{\phi} &= \omega_x + \omega_y \sin\phi \tan\theta + \omega_z \cos\phi \tan\theta \\ \dot{\theta} &= \omega_y \cos\phi - \omega_z \sin\phi \\ \dot{\psi} &= \frac{\omega_y \sin\phi}{\cos\theta} + \frac{\omega_z \cos\phi}{\cos\theta}. \end{aligned} \tag{A.35}$$

The position of the vehicle's centre of mass in the global XYZ frame is given by referring through the yaw-angle, pitch-angle and then roll-angle. In matrix form, the overall rotation matrix is given by:

$$\mathbf{R} = \mathbf{R}_\phi \mathbf{R}_\theta \mathbf{R}_\psi \tag{A.36}$$

where $\mathbf{R}_\phi$ is the roll rotation matrix, $\mathbf{R}_\theta$ is the pitch rotation matrix and $\mathbf{R}_\psi$ is the yaw angle rotation:

$$\mathbf{R}_\phi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & -\sin\phi \\ 0 & \sin\phi & \cos\phi \end{bmatrix},$$

$$\mathbf{R}_\theta = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos\theta \end{bmatrix}$$

$$\mathbf{R}_\psi = \begin{bmatrix} \cos\psi & -\sin\psi & 0 \\ \sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Let $\boldsymbol{X} = (X, Y, Z)^T$ be the position of the vehicle in the global coordinate frame, then the global coordinate position is given by the integration of:

$$\dot{\boldsymbol{X}} = \mathbf{R}(t)\,\boldsymbol{v} \tag{A.37}$$

as it is assumed that at the start of the simulation, the vehicle is at the centre of the coordinate frame. The rotation matrix is a function of time as the cardan angles change as described in Eq. (A.35).

A.2 Model Verification

In addition to the figures presented in Section 3.3.4, the figures below are provided to demonstrate the verification of the Multi-Degree-Of-Freedom model. All model and simulation parameters are the same as those used in Chapter 3.

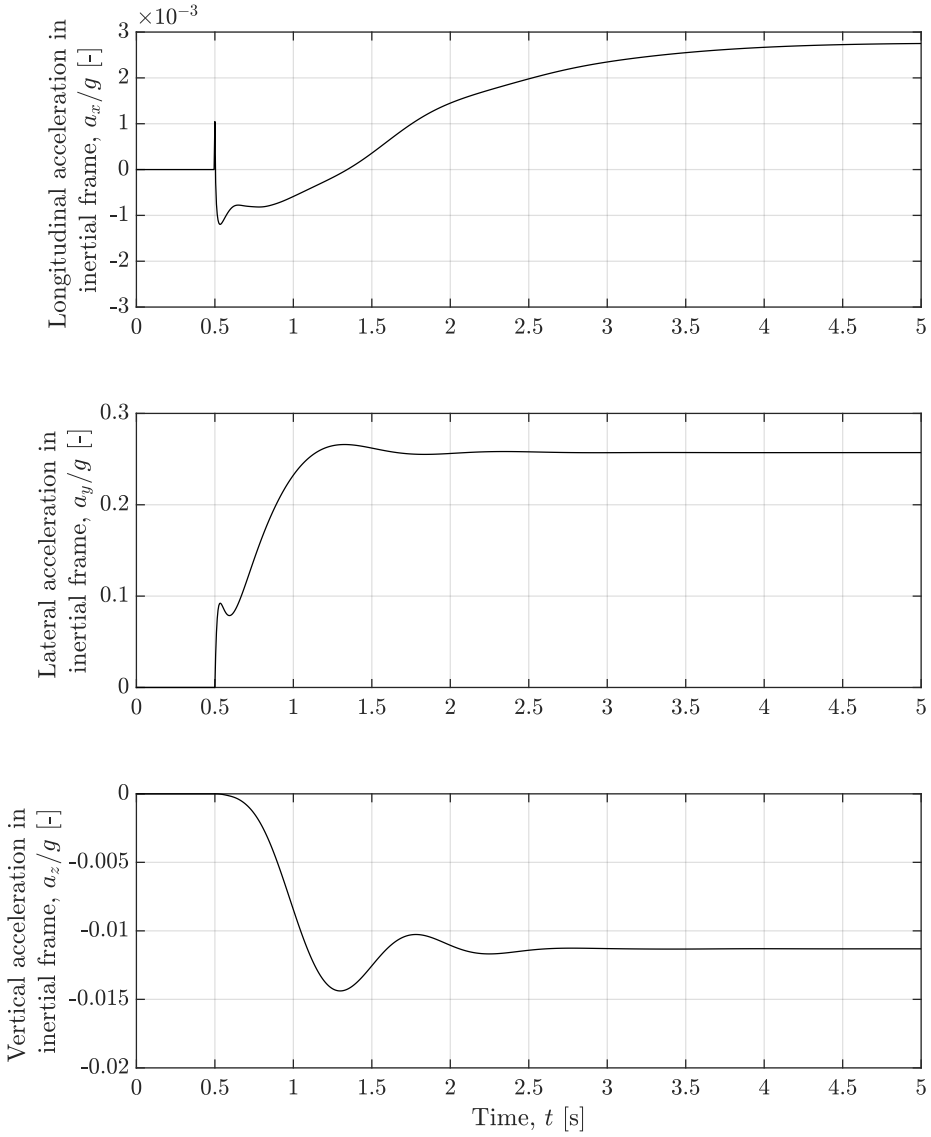


Figure A.1: Inertial acceleration of the vehicle body

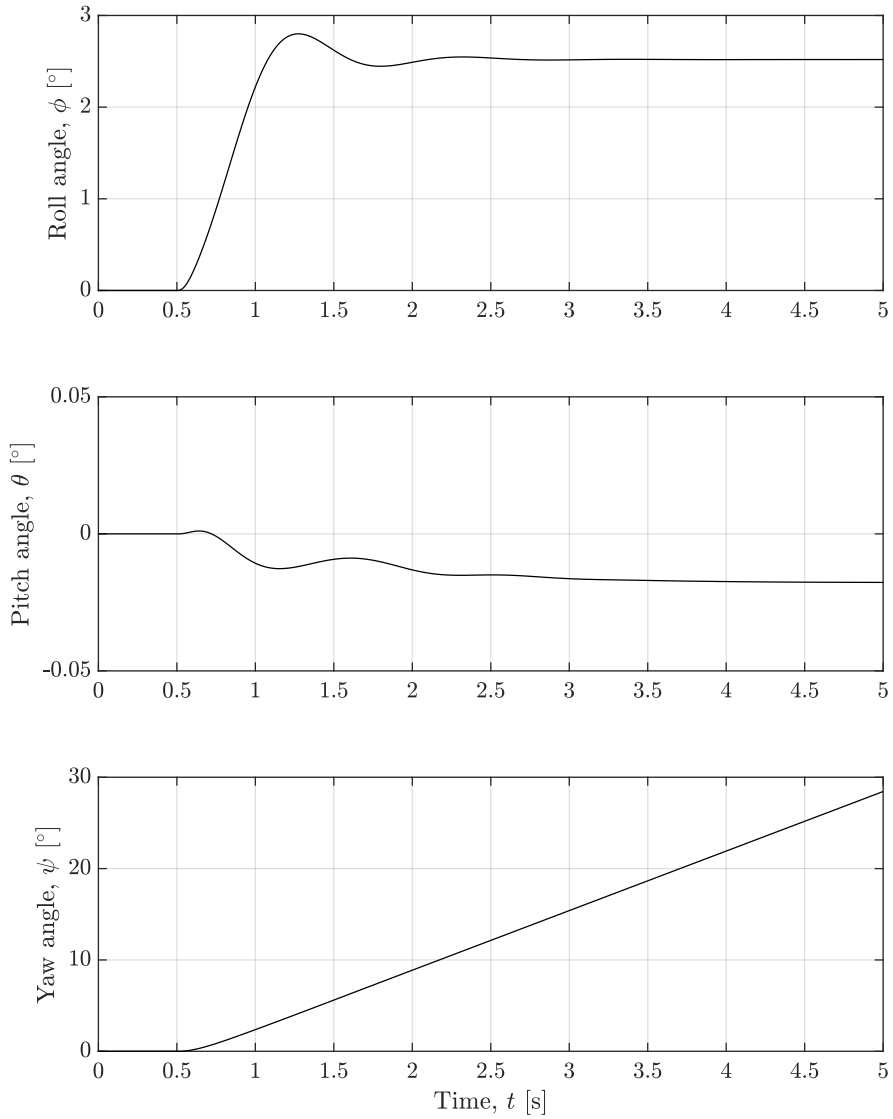


Figure A.2: Cardan angles of the Multi-DOF vehicle model. The roll angle of the vehicle settles to steady state after approximately 1 s after the step input, as the model enters the steady-state turn. The pitch angle of the vehicle is small due to the vehicle successfully tracking the reference speed

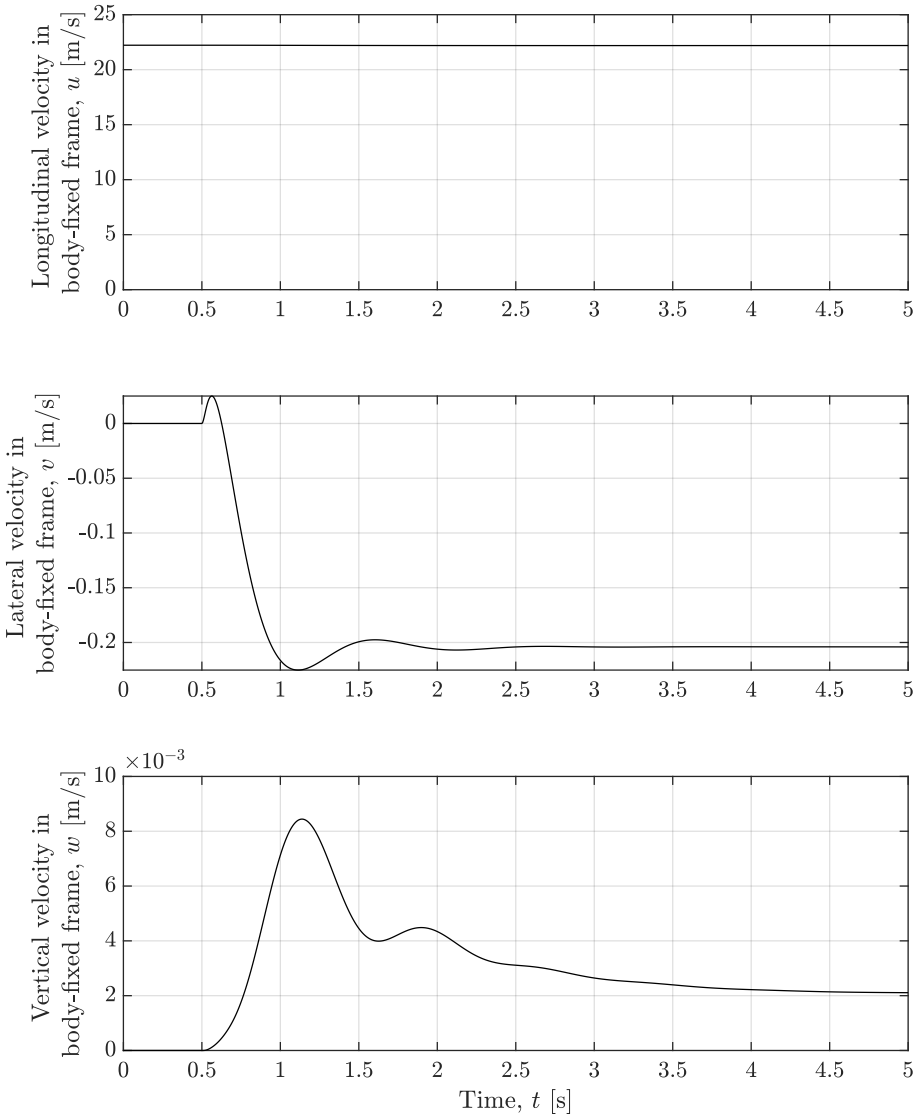


Figure A.3: Body-fixed velocity components of the Multi-DOF model's unsprung mass. The small vertical component is due to the non-zero roll and pitch angles of the vehicle body

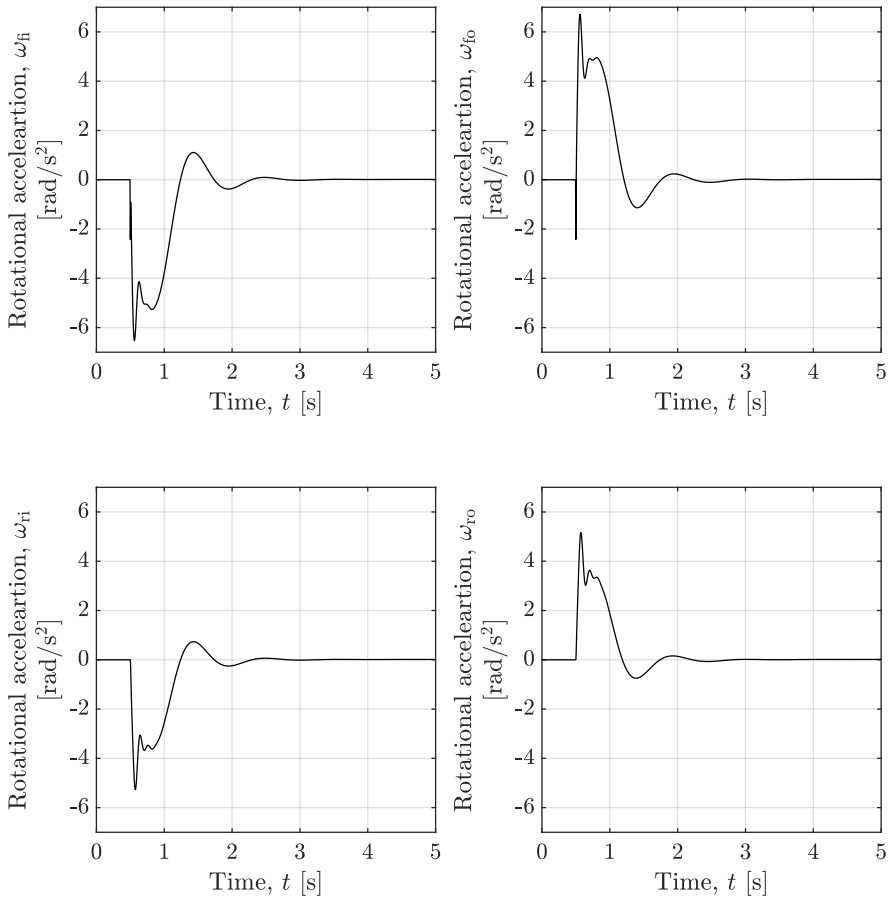


Figure A.4: Angular acceleration of the wheels

Appendix B

Smooth Lateral Reference Trajectory

To model a smooth reference trajectory for the lateral step manoeuvre, the following sigmoid curve is used:

$$Y(X) = f(X) = \frac{\hat{Y}}{1 + e^{-k(X - X_0)}} \tag{B.1}$$

where $\hat{Y}$ is the amplitude of the function, i.e., the height of the smooth step change, k is a positive sharpness value, that determines how quickly the function transitions to between the low and the high values of the function. X_0 is the midpoint of the curve ($Y(X_0) = \hat{Y}/2$).

For the four-state biycle model, along with with the lateral displacement it is desireable to create reference trajectories for the remaining states v, ω_z and ψ. To do this we note that the reference yaw angle is the angle between the slope of the lateral reference trajectory and the X-axis, which can be expressed via:

$$\psi = \tan^{-1} \frac{dY}{dX}. \tag{B.2}$$

The yaw rate is then the time-based derivative:

$$\omega_z = \frac{d}{dt} \tan^{-1} \frac{dY}{dx}. \tag{B.3}$$

For the lateral velocity, an expression can be found by rearranging the

equation for the linear approximation of lateral displacement:

$$\dot{Y} = v + u\psi \quad \rightarrow \quad v = \dot{Y} - u\psi \tag{B.4}$$

As the four-state bicycle model assumes constant longitudinal speed, which is much larger than the lateral velocity, the time based derivative of X is simply u:

$$\frac{dX}{dt} = u \tag{B.5}$$

From this, the following equations for the reference trajectories can be derived using standard calculus results:

$$\begin{aligned} Y &= f(X) \\ \psi &= \tan^{-1}\frac{dY}{dX} \\ \omega_z &= \frac{u\frac{d^2Y}{dX^2}}{1+\left(\frac{dY}{dX}\right)^2} \\ v &= u\left(\frac{dY}{dX} - \tan^{-1}\left(\frac{dY}{dX}\right)\right) \end{aligned} \tag{B.6}$$

where the first and second derivatives of the smooth step function are given by:

$$\begin{aligned} \frac{dY}{dX} &= -\frac{k\hat{Y}e^{-k(X-X_0)}}{\left(1+e^{-k(X-X_0)}\right)^2} \\ \frac{d^2Y}{dX^2} &= \frac{k^2\hat{Y}e^{-k(X-X_0)}\left(e^{-k(X-X_0)}-1\right)}{\left(1+e^{-k(X-X_0)}\right)^3} \end{aligned} \tag{B.7}$$

By substituting in the deriative terms of Eq.(B.7) into Eq. (B.6) the state references for the smooth-step are obtained for any value of X. The references can be found at time t by the simple relationship $X = ut$.

The terms in Eq. (B.6) for yaw angle, yaw rate and lateral velocity are used to derive secondary references from the path generation problem in Chapter 5. The derivative terms in that instance are calculated via finite difference, rather than Eq. (B.7).

Appendix C

Linear MPC Formulated as a Quadratic Programming Problem

The derivation of the standard MPC formulation represented in Chapter 4 is presented. The combined state and rate constraints are also included, and the constraint sets $\mathcal{X}, \mathcal{U}$ and $\mathcal{Y}$ are assumed to be represented by convex polytopes, and a terminal penalty matrix is used along with a seperate terminal state constraint set. The corresponding MPC definition is therefore:

$$\underset{\boldsymbol{\mathcal{U}}}{\text{minimise}} \quad J = \sum_{i=1}^{N} \|\boldsymbol{y}_i - \boldsymbol{y}_i^{\mathrm{r}}\|_{\mathbf{Q}}^2 + \|\boldsymbol{u}_{i-1}\|_{\mathbf{R}}^2 \tag{C.1a}$$

$$\text{subject to} \qquad \boldsymbol{x}_0 = \tilde{\boldsymbol{x}}_0, \tag{C.1b}$$

$$\boldsymbol{x}_{i+1} = \mathbf{A}\boldsymbol{x}_i + \mathbf{B}\boldsymbol{u}_i, \tag{C.1c}$$

$$\boldsymbol{y}_i = \mathbf{C}\boldsymbol{x}_i, \tag{C.1d}$$

$$\boldsymbol{\Delta u}_{\mathrm{lb}} \leq \boldsymbol{\Delta u}_i \leq \boldsymbol{\Delta u}_{\mathrm{ub}}, \qquad i = 0, ..., N-1, \tag{C.1e}$$

$$\mathbf{P}_u \boldsymbol{u}_i \leq \boldsymbol{h}_u, \qquad i = 0, ..., N-1, \tag{C.1f}$$

$$\mathbf{P}_y \boldsymbol{y}_i \leq \boldsymbol{h}_y, \qquad i = 1, ..., N, \tag{C.1g}$$

$$\mathbf{D}\boldsymbol{\Delta x}_i + \mathbf{E}\boldsymbol{x}_i \leq \boldsymbol{c}_i, \qquad i = 1, ..., N \tag{C.1h}$$

where $\boldsymbol{x} \in \mathbb{R}^n$ is the system state vector, $\boldsymbol{u} \in \mathbb{R}^m$ is the control vector and $\boldsymbol{y} \in \mathbb{R}^q$ are the system outputs. $\boldsymbol{\Delta x}$ represents the finite difference approximation of the continuous-time rate vector $\dot{\boldsymbol{x}}$, which in this instance is determined by a first-order backward difference approximation. The state constraint, $\mathbf{P}_x \boldsymbol{x}_i \leq \boldsymbol{h}_x,\ i = 1, \ldots, N$, and the output constraint, $\mathbf{P}_y \boldsymbol{y}_i \leq \boldsymbol{h}_y$, have been removed from the problem statement as it was redundant; any state or output constraint can be included in the generalised constraint Eq. (C1.h) by setting the corresponding row of $\mathbf{D}$ to zero.

To express the MPC formulation problem as a Quadratic Programming (QP) problem, first the optimisation variable is defined as:

$$\boldsymbol{z}^T = \left(\boldsymbol{u}_0^T,\ \boldsymbol{x}_1^T,\ \ldots,\ \boldsymbol{u}_{N-1}^T,\ \boldsymbol{x}_N^T\right) \quad \in \mathbb{R}^{N(m+n)},$$

with the corresponding set of state and control references denoted $\boldsymbol{z}^r$. The control references are assumed to be zero.

The QP can then be expressed as:

$$\begin{aligned} \underset{\boldsymbol{z}}{\text{minimize}} \quad & \boldsymbol{z}^T \mathbf{H} \boldsymbol{z} + \boldsymbol{f}^T \boldsymbol{z} \\ \text{subject to} \quad & \mathbf{F}\boldsymbol{z} = \boldsymbol{b}, \\ & \mathbf{P}\boldsymbol{z} \leq \boldsymbol{h}, \end{aligned} \tag{C.2}$$

with the problem terms defined as follows:

$$\mathbf{H} = \begin{bmatrix} \mathbf{R} & \mathbf{0} & & & & \\ \mathbf{0} & \mathbf{C}^T\mathbf{Q}\mathbf{C} & & & & \\ & & \ddots & & & \\ & & & \mathbf{C}^T\mathbf{Q}\mathbf{C} & \mathbf{0} & \mathbf{0} \\ & & & \mathbf{0} & \mathbf{R} & \mathbf{0} \\ & & & \mathbf{0} & \mathbf{0} & \mathbf{Q}_f \end{bmatrix} \in \mathbb{R}^{N(m+n) \times N(m+n)}$$

$$\boldsymbol{f} = -2\mathbf{H}\boldsymbol{z}^r \in \mathbb{R}^{N(m+n)}$$

$$\boldsymbol{z}^r = \left(\mathbf{0}^T,\ \boldsymbol{y}_1^{\mathrm{r}T},\ \ldots,\ \mathbf{0}^T,\ \boldsymbol{y}_N^{\mathrm{r}\ T}\right)^T \in \mathbb{R}^{N(m+n)}$$

$$\mathbf{F} = \begin{bmatrix} -\mathbf{B} & \mathbf{I}_n & \mathbf{0} & \mathbf{0} & & & & \\ \mathbf{0} & -\mathbf{A} & -\mathbf{B} & \mathbf{I}_n & & & & \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & -\mathbf{A} & & & & \\ & & & & \ddots & & & \\ & & & & & \mathbf{I}_n & \mathbf{0} & \mathbf{0} \\ & & & & & -\mathbf{A} & -\mathbf{B} & \mathbf{I}_n \end{bmatrix} \in \mathbb{R}^{Nn \times N(m+n)}$$

$$\boldsymbol{b} = \left((\mathbf{A}\boldsymbol{x}_0)^T,\ \mathbf{0}^T,\ \ldots,\ \mathbf{0}^T \right)^T \in \mathbb{R}^{Nn}$$

$$\begin{aligned} \boldsymbol{h} = \big(& -\boldsymbol{u}_{\text{lb}}^T,\ \boldsymbol{u}_{\text{ub}}^T, -\Delta\boldsymbol{u}_{\text{lb}}^T,\ \Delta\boldsymbol{u}_{\text{ub}}^T,\ \boldsymbol{c}_1^T, \ldots, \\ & -\boldsymbol{u}_{\text{lb}}^T,\ \boldsymbol{u}_{\text{ub}}^T, -\Delta\boldsymbol{u}_{\text{lb}}^T,\ \Delta\boldsymbol{u}_{\text{ub}}^T,\ \boldsymbol{c}_N^T \big)^T \in \mathbb{R}^{N(4m+p)} \end{aligned}$$

where blank entries within a matrix are zero and $\mathbf{I}_n$ is the $n \times n$ identity matrix. It is convienient to define the matrix inequality term, $\mathbf{P} \in \mathbb{R}^{N(4m+p)\times N(m+p)}$, using MATLAB indexing notation [131]:

$$\mathbf{P}\left(j_i,\ k_i\right) = \mathbf{M}, \qquad i = 2, \ldots, N \tag{C.3}$$

$$j_i = 2\left(m+p\right)\left(i-1\right) : 2\left(m+p\right) i \tag{C.4}$$

$$k_i = \left(m+n\right)\left(i-2\right) + 1 : 2\left(m+n\right) + \left(m+n\right)\left(i-1\right) \tag{C.5}$$

where $\mathbf{M} \in \mathbb{R}^{m_\mathbf{M} \times n_\mathbf{M}}$ and $m_\mathbf{M} = (4m+p)$, $n_\mathbf{M} = 2\left(m+n\right)$, is the block matrix:

$$\mathbf{M} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & -\mathbf{I}_m & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_m & \mathbf{0} \\ \mathbf{I}_m & \mathbf{0} & -\mathbf{I}_m & \mathbf{0} \\ -\mathbf{I}_m & \mathbf{0} & \mathbf{I}_m & \mathbf{0} \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} & \mathbf{D}+\mathbf{E} \end{bmatrix} \tag{C.6}$$

The initial block is given by:

$$\mathbf{P}\left(1 : 4m+p,\ 1 : m+n\right) = \begin{bmatrix} -\mathbf{I}_m & \mathbf{0} \\ \mathbf{I}_m & \mathbf{0} \\ -\mathbf{I}_m & \mathbf{0} \\ \mathbf{I}_m & \mathbf{0} \\ \mathbf{0} & \mathbf{D}+\mathbf{E} \end{bmatrix} \tag{C.7}$$