Application of robust nonlinear model predictive control to simulating the control behaviour of a racing driver

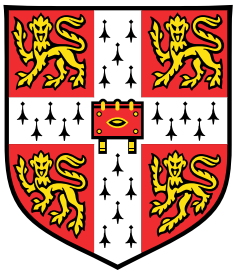

Giovanni Braghieri

Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

King's College March 2018

A L

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Giovanni Braghieri
March 2018

Acknowledgements

Firstly, I would like to thank my supervisor Dr. David Cole for his continuous guidance, support and understanding. His passion and knowledge have made this thesis possible.

I would also like to thank the UK Engineering and Physical Sciences Research Council and Renault Sport Racing Limited, for their financial support. I am thankful for the technical advice and insight of Daniele Casanova and Pierre Genon. Many thanks to Tim Johns for all the help with the simulator tests.

Over the past 4 years I learnt about vehicle dynamics, optimal control and experimental methods, but more importantly I have learnt a lot about myself.

I have to thank Dario, who's now better than me at cooking, for teaching me that Zeno is not the only source of wisdom.

I also have to thank Enrico who, despite losing to me around 150 foosball matches (and winning only 2) over the course of my PhD, has always helped me make the right choices.

Thanks to Reuben for the encouragement to accept my several genetic limitations and the discussions on social justice.

And Dan, for all the bromance.

A special mention goes to my friend Ginetto, my longest relationship to date.

Abstract

The work undertaken in this research aims to develop a mathematical model which can replicate the behaviour of a racing driver controlling a vehicle at its handling limit. Most of the models proposed in the literature assume a *perfect* driver. A formulation taking human limitations into account would serve as a design and simulation tool for the automotive sector.

A nonlinear vehicle model with five degrees of freedom under the action of external disturbances controlled by a Linear Quadratic Regulator (LQR) is first proposed to assess the validity of state variances as stability metrics. Comparison to existing stability and controllability criteria indicates that this novel metric can provide meaningful insights into vehicle performance. The LQR however, fails to stabilise the vehicle as tyres saturate.

The formulation is extended to improve its robustness. Full nonlinear optimisation with direct transcription is used to derive a controller that can stabilise a vehicle at the handling limit under the action of disturbances. The careful choice of discretisation method and track description allow for reduced computing times.

The performance of the controller is assessed using two vehicle configurations, Understeered and Oversteered, in scenarios characterised by increasing levels of nonlinearity and geometrical complexity. All tests confirm that vehicles can be stabilised at the handling limit. Parameter studies are also carried out to reveal key aspects of the driving strategy.

The driver model is validated against Driver In The Loop simulations for simple and complex manoeuvres. The analysis of experimental data led to the proposal of a novel driving strategy. Driver randomness is modelled as an external disturbance in the driver Neuromuscular System. The statistics of states and controls are found to be in good agreement. The prediction capabilities of the controller can be considered satisfactory.

Table of contents

List of figures

List of tables

List of symbols

$\mathbf{A}$	state-space state matrix
$\mathbf{B}$	state-space input matrix
$\mathbf{C}$	state-space output matrix
$\mathbf{D}$	state-space feedforward matrix
$\mathbf{F}$	state-space constant matrix
$\mathbf{E}$,$\mathbf{J}$	Prediction matrices
$\mathbf{G}$	state-space constant matrix
$\mathbf{Q}$	cost function output weight matrix
$\mathbf{R}$	cost function control effort weight matrix
$\mathbf{u}$	control or input vector
$\mathbf{w}$	additive disturbance
$\mathbf{x}$	state vector
$\mathbf{X}$	state prediction vector
$\mathbf{z}$	output vector
$\mathbf{Z}$	output prediction vector
$\mathbf{\Psi}, \mathbf{\Xi}, \mathbf{\Theta}, \mathbf{\Upsilon}$	prediction matrices
$\boldsymbol{\epsilon}$	error matrix
a	distance from CoG to front axle
A_x	frontal area
a_l,b_l	aerodynamic lift coefficients
a_x,b_x	aerodynamic drag coefficients
b	distance from CoG to rear axle
$\mathbf{b}$,$\mathbf{q}$	Parameter vector
B,D,E	Pacejka Magic Formula coefficients
b_f	brake balance
c_1, c_2	cornering stiffness parameters
C	cornering stiffness and Pacejka Magic Formula coefficient
C_α	normalised tyre slip stiffness

C_L	lift coefficient
C_x	drag coefficient
$\mathbf{d}$	search direction
dt	increment in time
D	tyre intercept
e	lateral displacement error between vehicle and track centreline trajectories
f_k	constraint function
F_{ax}	aerodynamics drag force
F_{az}	aerodynamics lift force
F_{exp}	experimental tyre force
F_{dist}	disturbance force
F_{sim}	simulated tyre force
F_{lon}	tyre longitudinal force
F_t	total force
F_x	longitudinal tyre force
F_y	lateral tyre force
F_z	vertical tyre load
$F_{z,t}$	total vertical tyre load
G_{sw}	steering wheel to road wheel angle ratio
g	acceleration due to gravity
$g_{nl}(\cdot)$	nonlinear function
h_k	constraint function
H_p	Hold time of holds of length p
$H(\cdot)$	Heaviside step function
I_f	front axle moment of inertia
I_r	rear axle moment of inertia
I_z	yaw moment of inertia
J	cost function
J_{lift}	drag parameter identification cost function
J_{lift}	lift parameter identification cost function
J_t	tyres identification cost function
J_{tl}	combined tyres identification cost function
J_p	parameter identification cost function
k	discrete time step
$l(\cdot)$	nonlinear function
L	wheelbase

L_t	total number of laps
M	vehicle mass
M_{dist}	moment disturbance
M_t	total moment
M_z	yaw moment with respect to the z axis
n	lateral offset
N_p	prediction horizon
N_i	intermittency period
O	quadratic penalty function
$p(\cdot)$	Lagrangian estimate
p_n	penalty parameter
p_t	number of track points
q_n	cost function weight
Q	tyre stiffness factor
r	tyre radius
R	Circle radius
r_n	random number
s	distance travelled
S	tyre shape factor
$S_{i,j}$	signal
$S_{\delta\delta}$	HWA MSSD
s_r	time to distance scalar
S_H	tyre force horizontal shift
s_x	normalised longitudinal tyre slip
s_y	normalised lateral tyre slip
t	time
T	road wheel torque
T_d	sampling interval
T_{run}	experimental run time
T_s	Prediction horizon
T_u	Update time
t_{SBH}	SBH length
t_{ZOH}	ZOH length
u	forward velocity of the vehicle
V	absolute velocity of the vehicle
$V_{(k)}$	cost function

v	lateral velocity of vehicle
x	longitudinal displacement
y	lateral displacement
X	vehicle horizontal position (cartesian coordinates)
Y	vehicle vertical position (cartesian coordinates)

Greek symbols

α	lateral tyre slip angle
α_l	Lomax distribution parameter
α_{max}	maximum lateral tyre slip angle
α_n	normalised slip
α_{max}	satuartion slip
β	vehicle body side-slip angle
Δ	time step
δ	road wheel angle
δ_{com}	commanded hand wheel angle
$\delta_{com,t}$	disturbed commanded hand wheel angle
δ_{sw}	actual hand wheel angle
ϵ	combined slip
ψ	heading direction
κ	longitudinal slip ratio
κ_{com}	commanded longitudinal slip ratio
κ_{max}	maximum longitudinal slip ratio
κ_n	normalised longitudinal slip ratio
κ_p	applied longitudinal slip ratio
κ_{sat}	saturation longitudinal slip
λ	Lagrangian multiplier
λ_l	Lomax distribution parameter
μ	coefficient of friction
μ_i	mean
ξ	angle between the vehicle and tangent to reference path
θ	absolute heading angle
θ_f	front axle angular displacement
θ_r	rear axle angular displacement
ρ	air density

σ	standard deviation
ξ_n	neuromuscular system damping ratio
ω	rate of change of yaw angle
ω_h	HWA frequency
ω_n	driver NMS natural frequency
Ω_z	track curvature
Ψ	wheel angular speed

Subscripts

0	arbitrary operating condition
app	approximate
c	continuous time
d	discrete time
global	global frame of reference
f	front
i	iteration step
r	rear
ref	reference
t	total
veh	vehicle variables

Superscripts

$*$	optimal

Acronyms

CAS	Collision Avoidance System
COP	Centre Of Pressure
COG	Centre Of Gravity
CPU	Central processing unit
HWA	Hand Wheel Angle
ICLOCS	Imperial College London Optimal Control Solver
IPOPT	Interior Point solver
LPE	Lateral Path Error
LTV	Linear Time Varying

MAD	Mean Average Distance
MPC	Model Predictive Control
MSSD	Mean Square Spectral Density
NMS	Neuromuscular Filter
NMPC	Nonlinear Model Predictive Control
PDF	Probability Density Function
PSI	Practical Stability Index
RAM	Random Access Memory
SBH	Serial Ballistic Hold
VSC	Vehicle Stability Control
ZOH	Zero Order Hold

Chapter 1

Introduction and literature review

Simulations are becoming more and more popular in the racing establishment as they provide a quick and inexpensive alternative to the now banned on-track testing. Formula One teams resort to simulations in the vehicle design stage, to train drivers and to find the optimal set of parameters for each racetrack [1].

The most common approach is based on quasi-steady-state assumptions, which approximates laps to a series of manoeuvres where the vehicle has no unbalanced forces or moments acting on it. This method allows complex vehicle models to be incorporated but does not consider the human driver and their role in feedback. A solution to this problem has been found by resorting to *driving simulators*. These machines replicate real driving conditions in the virtual environment; they can be thought as a glorified version of a racing video-game, where the response of the virtual vehicle is extremely realistic, the driver has all commands that are available in reality and motion cueing is provided by hydraulic or electric actuators. However, driving simulators, besides being very expensive to build and to maintain, do not guarantee repeatability as a human is present in the loop [2].

There is perhaps too much reliance on *perfect* driving models, such as [3–5], these not taking human limitations into account. Lately, an effort to include such limitations has been made [2, 6]; however, random factors that normally arise in racing circumstances have been barely taken into account. This work focuses on devising a driver model that can handle a variety of random factors at the handling limit. This implies that not only random external factors – such as road profile and wind speed – are considered, but also the inevitably imperfect cognitive process involved in the driving tasks. This translates to designing a robust nonlinear controller which can deal with arbitrary disturbances.

The aim of this research is to contribute to the development of a virtual driver model which is able to replicate with high fidelity the behaviour of a real human driver. Such a model could be used to optimise vehicle parameters without resorting to a driving simulator and in other areas of automotive engineering, such as collision avoidance.

Modelling a racing driver is a task that brings together various areas of automotive engineering. One of these is nonlinear vehicle dynamics, as the vehicle obviously plays a very important role in the driving task. Since racing involves very aggressive manoeuvres, modelling nonlinearities is very important; particular attention is given to modelling tyres, which are ultimately responsible for the forces acting on the vehicle. Stability metrics are also taken into consideration as they are important performance parameters. Research into modelling the perfect driver is also very useful as it provides the base for building more complex models. Lateral, longitudinal and combined control are considered. In order to be able to improve on existing driver models, advances in nonlinear robust control have to be considered, in particular tube-based nonlinear control. The literature review focuses on all these areas of knowledge to describe the state of art and highlight areas for further research.

1.1 Nonlinear vehicle dynamics

Tyres are the most significant source of nonlinearity in relation to vehicle dynamics [7].

Tyres have to bear the vertical load of the vehicle, bear the longitudinal accelerations governing the longitudinal dynamics and produce the lateral forces that allow the vehicle to steer [8]. They are therefore the most critical aspect of any vehicle model because they provide the forces which are responsible for the handling behaviour. Some tyre models will be reviewed and their benefits and drawbacks assessed.

The review will continue by looking at how these tyre models can be integrated into full vehicle models and how the modelling of the components of the vehicle influences the degree of non-linearity and the complexity of the final model. Finally, an overview of the work on stability and controllability metrics is outlined.

1.1.1 Tyre models

Tyres have to bear vertical load, absorb road deformations, and produce lateral and longitudinal forces for the vehicle to turn and accelerate. Tyre models can be steady state models, in which time is neglected, or transient models, which take time into

account. The latter are much more complex and intractable than the former, so they are only used if they are relevant to the problem studied.

The main characteristics of tyres' dynamic behaviour is captured by the tyre relaxation length; this measure has various definitions but the most intuitive is given by Cossalter *et al.* in [9]. It is defined as the distance a point on the tyre in contact with the road has to roll for the force to build up to 63% of its nominal value. Typical values for relaxation lengths range from 0.12 m to 0.45 m [9]. Another important parameter related to the relaxation length is the relaxation time constant, which determines the first order time lag of the tyre. It is defined in [10] as the relaxation length divided by the forward velocity.

For racing vehicles, whose tyres are characterised by short relaxation lengths, time constants are shorter than 5 ms for the majority of race conditions; given the overall transient dynamics of the vehicle, the relaxation time constant for a racing car is negligible. It is therefore sensible to consider steady state tyre models.

Pacejka and Sharp, [11], conducted a very comprehensive review of steady state tyre models, identifying three main approaches: physical, empirical and semi-empirical. While physical models are based on the mathematical representation of the tyres behaviour, empirical models use experimental observations to find tyre forces. Semi-empirical methods are a hybrid approach which assumes the tyre forces to be modelled by a certain function and try to fit the curves using experimental data. Physical models are reviewed first, followed by empirical and semi-empirical.

Wong [12] outlines a mathematical model which involves modelling the tread of the tyre as a stretched string which is restrained by lateral springs. The lateral force and the self-aligning moment are found assuming a linear lateral displacement of the string. The model is extended to include traction and braking. However, it is assumed that the combination of vertical load and lateral slip angle remains within the linear region and it does not consider high-slips which would lead to sliding. One of the best known physical models is Fiala model [13], also used in commercial modelling software. The tyre is modelled as two concentric rings, the inner rigid and the outer elastic, connected by springs. The contact area is approximated by a quadratic polynomial which allows the modelling of the non-linear tyre behaviour. Sakai *et al* [14] improves the Fiala model by adding additional terms to the self-aligning moment to give a better match between theoretical and experimental results. An important contribution to physical models has been given by Sharp who developed a computer model [15] to numerically calculate tyre forces by approximating the tyre as spokes attached to a rim.

All the physical models reviewed show common limitations: number of simplifications necessary to solve the system; approximation on the pressure distribution; inability to capture strong nonlinearities; computational demand and lack of a closed-form solution. The use of empirical methods can circumvent most of the aforementioned issues; these approaches are based on performing experiments and subsequently analysing the results. Data is either organised in tables, shown in graphs or used to fit equations. The latter, despite being more challenging, is the preferred method as it offers a lot of flexibility. The tyre response has been approximated using exponential, arctangent, polynomial, hyperbolic tangent curves and Fourier series. None of the aforementioned methods provide a satisfactory result apart from polynomials and series; however, these suffer from waviness [16]. This problem can be solved by using splines and lower order polynomials but the method is impractical.

Semi-empirical methods can be used to circumvent these limitations. Instead of simply fitting data to standard functions, experimental observations are used to generate functions that capture the tyre behaviour and experimental data is subsequently used to find the best fit. The most widely used formula is the *Magic formula*, initially proposed in [17], later refined in [16] and [7]. The formula is obtained by using functions that are able to capture the behaviour of the tyres observed in experiments. It was found that the requirements are met by using a function which includes sinusoidal and arctangent functions. This method has an excellent agreement with experiments, it is not computationally intensive and, most importantly, it has an expression in closed form, which is very important as it allows further manipulation. Combined slip characteristics can also be modelled by normalising the slip by factors proportional to the front and rear vertical loads.

1.1.2 Vehicle models

A vehicle model is a mathematical entity that has to capture the physical features of a real vehicle. For the purpose of this work, vehicle vibrations are not considered. Various vehicle models are analysed to assess their benefits and drawbacks.

The bicycle model is usually the first vehicle model to be considered because of its simplicity and the insights it gives into the dynamics of the vehicle. Pacejka [7] gives an excellent overview of this model. It is comprised of two wheels connected by a rigid rod representing the body of the vehicle; track width is therefore neglected. Longitudinal forces are not considered so the forward velocity is assumed to be constant; in this respect, the driving force required to keep the vehicle at this speed is assumed to be small with respect to the lateral forces. The steer and slip angles are restricted to

relatively small values; assuming linear tyres, the equations of motion are easily derived and can be expressed and solved in closed form, which is very convenient for vehicle handling analysis.

The bicycle model can also be coupled with nonlinear tyres, giving very good insights into the behaviour of the vehicle in limit conditions. A closed form solution cannot be obtained because of the nonlinear tyre forces. Various approaches have been developed to circumvent this problem. Cole *et al* [6] linearise the equations of motion at each time step about the current slip angle to design a controller to handle the vehicle in limit conditions. Mastinu [18] and Della Rossa [19] analysed a two degrees of freedom bicycle model using bifurcation analysis, hence phase planes, to predict the behaviour for ten different types of vehicles. They showed that the dynamic behaviour varies significantly and that both simple and complex bifurcations arise. The bicycle model with linear and nonlinear tyres is also used to derive stability criteria.

Longitudinal dynamics are very important as they allow throttle control. The simplest model is the trolley model, used for instance in [1]. It has three degrees of freedom, namely the front and rear wheel angular speed and the vehicle forward velocity. More complex models have been proposed in the context of nonlinear longitudinal control, for instance in [20]. A very comprehensive longitudinal vehicle model was employed by Majdoub *et al.* [21]. They take into account rolling resistance, the effect of temperature and the geometry of the vehicle.

Aerodynamic effects can easily be incorporated by adding or subtracting drag and lift forces. The equations of the longitudinal model can be used in conjunction with the bicycle model to have a simple vehicle model that allows the design of steering and throttle control. Longitudinal and lateral coupling of the tyres is obtained using the similarity method described in [7]. Timings chose this configuration in [1].

Even though this model captures well a number of features, it does not consider neither roll nor weight transfer. Furthermore, drifting manoeuvres cannot be modelled very accurately because the effect of the differential is neglected.

A four-wheels model can then be used to overcome these limitations; Velenis *et al.* [22] used this model to study a controller that stabilises steady-state drifting. They modelled the vehicle using four wheels, longitudinal and lateral weight transfer; they included a limited slip differential, which limits the left and right slip difference to the extent that the left and right drive torques are the same and they correspond to the driver's throttle input. The increased complexity and computational power are justified by a very effective resulting controller. Furthermore, the model can be complemented with various features to improve the quality of the simulation. Kelly

used a full dynamic car model in [23]. The 4 wheels model coupled with nonlinear tyres was complemented with several features of actual vehicles. A simple aerodynamics model was added to account for drag and lift. The engine was modelled using linear interpolation on existing Max/Min curves and the gearbox was modelled using an *atan* function. A differential in open and locked configurations was also considered.

1.1.3 Stability

Pacejka [7] examines the stability of steady state motion of a car in a straight line and on a curved path using the bicycle model with nonlinear tyres and handling diagrams. These are very powerful graphical methods for predicting the behaviour of the nonlinear bicycle model system. Plotting the ratio of the side force to the vertical load against the slip angle gives the normalised axles characteristics, which, subtracted horizontally from each other produce the handling curve.

The diagram is completed by including the graph that shows the relationship between lateral acceleration (in g units) and the relative path curvature for a series of speeds. This information can be used to find the hand wheel angle necessary to negotiate a turn of a given radius at a given forward velocity by finding the difference between a given point and the plot; furthermore, stability boundaries can be obtained.The analysis is however limited to steady state manoeuvres.

The stability metrics presented in Dixon [24] and Milliken and Milliken [25], are derived using the linearised bicycle model under steady state conditions. The theory behind these as well as their function is analytical and well understood. Milliken and Milliken [25] also present a quasi-steady state method for describing the behaviour of the car close to the tyre slip limit, called Force Moment Method. The authors graphically derive stability conditions and handling characteristics at the limit of adhesion of the tyres. One such metric is dC_N/dA_y, where C_N is the normalised yaw moment and A_y the lateral acceleration; it describes the understeer or oversteer behaviour of the car depending on the sign of the derivative. Ono [26] and Nguyen [27] describe stability using the phase plane method, where trajectories and states of the car are plotted for different sets of initial conditions and inputs. These plots are used to show stable or unstable equilibria and regions of operation as well as tracking the transient response of the car at a constant speed and steering input. Ono et al. [26] show that the phase plots bifurcate with changing operating conditions and inputs.

Skoog *et al.* [28] investigates the 'Frozen Time Eigenvalues' for time varying systems of the form $\dot{\mathbf{x}} = \mathbf{A}(\mathbf{t})\mathbf{x}$. Desoer [29] followed a previous derivation done by Rosenbrock [30] to show that for slowly varying systems, if $\sup ||\dot{A}(t)||$, where sup indicates the

supremum, is sufficiently small, the time varying system is asymptotically stable if the 'Frozen System' is stable; he also obtained explicit bounds. Skoog *et al.* [28] have additionally shown that if the eigenvalues of such slowly varying systems are away from the Imaginary axis, then the 'frozen time' system accurately reflects the stability of the system. If either condition is violated, however, it is shown that the relationship between the 'Frozen Time' and continuous systems may not follow.

Evangelou [31] used an eigenvalue analysis of the linearised, time varying state space model of a motorcycle to investigate its stability. Root loci of the poles of the motorcycle system were plotted as they varied with time, for different manoeuvres. Evangelou's conclusions drawn from his linearised frozen-time model were verified against nonlinear simulations to find that the frozen-time models could predict the behaviour of the motorbike model quite accurately. However, they stressed, that this method should be used with utmost care when inferring stability, especially when the eigenvalues cross the imaginary axis.

Limebeer *et al.* [32] studied frozen-time eigenvalues in the context of motion cueing. They focus on high-performance vehicle simulators where the handling characteristics are characterised by large accelerations. The stability and response characteristics for coupled longitudinal and lateral dynamics are characterised using frozen-time eigenvalue analysis. Motion cues are enhanced with lateral acceleration and yaw cueing filters based on the frozen-time eigenvalues analysis. These prove to be effective as they are tested by a professional race driver in a simulator receiving positive feedback.

Johansen *et al.* [33] argued that it is not sensible to evaluate the stability of a linearised system at an arbitrary, non-equilibrium point because it is by definition a transient state. However, Meijaard [34] considered the stability of systems in transient motion and proposed the concept of 'Linear Practical Stability' relating to perturbation growth. A metric called the Practical Stability Index was introduced and was found useful in analysing the dynamic behaviour of linearised systems.

Research on vehicle stability has been conducted in the context of Electronic Stability Control systems. They work by detecting when the car is about to become unstable or uncontrollable and they intervene using individual wheel braking action [35] or active steering [36] to bring the vehicle back to a stable operating condition; for the purpose of this project, the focus is on stability evaluation. It was found that most stability systems work by comparing measured car state data to reference states calculated in real time by a linear car model receiving the same driving inputs as the real car. Whenever the difference between the measured and reference states exceeded some threshold limit, the stability system intervened to stabilise the vehicle. This

threshold limit was calculated in different ways and details of such methods can be found in Gerdes *et al.* [37] and Youn *et al.* [36]. Their limitation is that real data from the car is needed to gauge how its handling is degraded in limit handling conditions.

The Aircraft Flying Qualities literature separates the stability of aircraft into static and dynamic stability. Static stability describes the tendency of the response of aircrafts to perturbations. Stability derivatives, relating moments about different axes of the aircraft to rotation about these axes, are used as metrics for static stability [38], [39].

The use of *stability and controllability derivatives* originally employed in the aircraft industry is advocated in [25]. If the total lateral force and yaw moment applied to the vehicle by the tyres are denoted by F_t and M_t then in the case of the linear operating regime they can be expressed in terms of the vehicle sideslip angle β, yaw rate $\dot{\psi}$, front road wheel steer angle δ and six partial derivatives:

$$F_t = \left(\frac{\partial F_t}{\partial \beta}\right)\beta + \left(\frac{\partial F_t}{\partial \dot{\psi}}\right)\dot{\psi} + \left(\frac{\partial F_t}{\partial \delta}\right)\delta \tag{1.1}$$

$$M_t = \left(\frac{\partial M_t}{\partial \beta}\right)\beta + \left(\frac{\partial M_t}{\partial \dot{\psi}}\right)\dot{\psi} + \left(\frac{\partial M_t}{\partial \delta}\right)\delta \tag{1.2}$$

Numerical values of the six partial derivatives can give insight into the handling behaviour of the vehicle.

A *handling diagram* is constructed in [7] to depict some aspects of the handling behaviour of a vehicle. Normalized characteristics of lateral force against slip angle for front and rear axles are used to construct the diagram. Equilibrium points and their stability can be determined for combinations of vehicle speed and steering angle. The portrait depicts the trajectories of two states of the vehicle, for example yaw rate and lateral velocity. Equilibrium points and their stability can be determined by observing convergence or divergence of the trajectories. A limitation of the phase portrait is that only a limited range of operating conditions can be depicted on an individual portrait.

Sideris [40] conducted an in-depth review of the methods outlined in the previous paragraphs highlighting their benefits and drawbacks. They went on to extend the practical stability index (PSI) approach by proposing the use of state variances as stability metrics, considering continuous disturbances instead of single perturbations. They used the covariance equation to calculate state variances of a disturbed vehicle. The algorithm was tested on various types of turns; however, it was found that the variances of lateral path and heading deviations were unbounded. Haslam [41], following Sideris' work, found that variances are bounded only if there is a compensatory controller acting on the disturbances. They implemented an ancillary LQR controller and showed

that state variances give insightful results for a 90° degree corner. The algorithm is validated against an iterative scheme. A parameter study was also conducted to demonstrate the potential of the method developed.

1.2 Modelling the perfect driver

The literature reviews done by Timings in [1] and Edelmann [42] are very relevant to this work; a similar structure is therefore followed. Work on control of vehicle dynamics is also reviewed as it is closely related to modelling the perfect driver.

1.2.1 Driver's characteristics

Driver models are mathematical representations of the control actions of a real driver; they are based on observing actual driving techniques and analysing the biological processes that occur when driving to determine the characteristics of the ideal driver. This section focuses on the processes involved in driving: gathering external information, elaborating on it, generating a command for the neuromuscular system and actuating it.

Research into the human visionary system has shown the driver relies mainly on vision for gathering information to determine the optimal path [43] and that the eye can elaborate positional and velocity information independently from a scene [44].

The first driver model based on preview information was developed by Sheridan [45, 46]; the path problem is expressed as a local optimal preview problem whereby the driver seeks to minimise the tracking error over a finite horizon. MacAdam [3] devised a driver model that minimises the lateral path error by considering the sum of the square difference between each preview point and the position of the vehicle over a finite horizon.

Another important characteristic of the driver is its ability to adapt to a varying plant, as it was shown in [47, 48]. Starting from those observations, McRuer *et al.* developed the well known cross-over model in [49, 50] which describes how humans adapt to a regulation task. It takes the form of a compensatory describing function that allows good tracking demand at low frequencies and less sensitivity at higher frequencies.

Research has also focussed on the neurological processes involved in the driving tasks; it has been proposed – for instance in [3] – that the brain creates a set of internal models which are selected according to sensory information to generate driving

commands. Various methods have been proposed to generate these internal models. Neural networks have been used in [4]; the internal models are described using the architecture associated with this type of control. Cheng and Fujioka use fuzzy logic to describe a hierarchical driver model comprised of four layers: decision making, task planning, manoeuvre and action [5, 51].

1.2.2 Lateral, longitudinal and combined control

A full driver model is not common because most studies focus on specific driver behaviours which may only require a specific type of control. The focus is generally on lateral control, longitudinal control or combined control. These will be reviewed separately.

The majority of driver models focus on the steering action needed to follow a specified trajectory, hence on lateral control only. Making use of concepts that emerged during the study of active vehicle suspension [52, 53], Sharp and Valtetsiotis [54] derived a preview steering controller by projecting an imaginary optical lever ahead of the vehicle and comparing lateral position points along this lever to the ideal path. These studies offer a number of interesting insights: firstly, preview control shows a diminishing return feature whereby preview information becomes decreasingly useful beyond a particular distance. Secondly, the preview gains reflect the lateral behaviour of the vehicle.

Ungoren and Peng [55] designed a Model Predictive Controller (MPC) controller to capture the ability of the driver to adapt to the changing dynamics of the vehicle. They assumed that the driver learns to invert the vehicle dynamics and to assess the mismatch between the actual and predicted response. The controller, however, rests on assumptions that do not fully reflect their ideas: they ignore driver speed control and restrict themselves to tests of short durations whereby there is no adaptation, which results in the driver changing the input only twice over the prediction horizon.

Cole *et al.* [56] compared Linear Quadratic (LQ) methods and predictive control theory designing path following controllers based on each method. The MPC controller was derived by solving an unconstrained predictive control problem following Maciejowski [57]. They highlighted the similarities between the two controllers, showing that for long previews, the model predictive controller and the linear regulator give an identical controllers.

In [54, 58], Keen and Cole designed a controller using multiple internal models, whereby the driver is assumed to have a set of internal models which are recalled following sensory information to determine the driving input. The number and scope

of internal models determines the level of skill of the driver. The MPC controller was derived by linearising the equations of motion about the predicted trajectory.

In the context of an autonomous vehicle, Falcone *et al.* [59] developed various controllers able to track vehicles on a slippery road. The path is assumed to be known; the first attempt at developing a controller was posed as a nonlinear MPC (NMPC) control problem; the computational burden prevented the controller from working for reasonable vehicle speeds in real time. The second attempt was a Linear Time Varying (LTV) controller, imposing constraints on the tyre curve to stabilise the vehicle at high speeds. Real-time implementation was made possible by reducing the number of internal models.

The development of driver safety support systems has encouraged research in vehicle longitudinal control. Sharp [60, 61] worked on speed tracking of road vehicles during acceleration and braking. They used a linear small-perturbation model whereby the equilibrium points were calculated using simulations, then deriving state-feedback and preview gains calculated. They found that, since the dynamics are simpler than in the lateral case, almost perfect tracking could be achieved. However, the use of small perturbations limited the tracking tests to profiles with small deviation from a nominal speed.

A similar approach was used by Fritz [62]. They improved the controller designed by Sharp by using a complex nonlinear longitudinal vehicle model and minimising an objective function together with the nonlinear vehicle model in closed-loop. The results were validated experimentally, showing that the controller can be used over a satisfactory range of operating conditions.

Even though considering lateral and longitudinal controls separately for a driving task is convenient, it is only with combined control models that a fuller understanding of driving can be achieved.

MacAdam [43, 63], used speed control to allow lateral steering control. They developed a driver model that includes internal vehicle dynamics models, prediction, steering/speed control and basic sensory and neuromuscular limitations. The information about the preview scene is used offline to generate a road centreline to be tracked. The novel aspect of the model is the ability of the driver to change the preview distances. A combined driver-vehicle system is obtained using simulations by coupling the driver model to a vehicle model developed by General Motors.

Falcone *et al.* [64, 65] detail a model that includes future vehicle speeds and positions information as well as lateral path and heading angle demands. Two active steering and torque controllers are developed, one based on NMPC and the other

based on LTV MPC with fixed model preview. Finally, the LTV MPC is simplified by substituting the internal vehicle model with a reduced order model linearised about the current operating conditions. Brake or throttle interventions are neglected for the linearisation.

The algorithms were tested by simulating a double lane change manoeuvre in icy conditions. The objective of the vehicle was to maintain a constant velocity throughout the manoeuvre, hence reducing the tractive force to a minimum.

The NMPC proves to be a good controller at a very high computational cost; the LTV MPC's performances are comparable to the NMPC at reduced computational cost. The reduced order LTV MPC controller shows good tracking abilities; however, the vehicle is only stable at low speeds. The major limitation in Falcone's work is the fact that the target trajectory is not generated.

Chang and Gordon [66, 67] relax this assumption in their work on Collision Avoidance Systems (CAS). They address the problem of path following through an infinite state representation of the desired path in the form of a reference vector field; however, this method proves too complex and impractical.

A simpler kinematic policy (KP) is proposed by Gordon and Magnuski [68]. Reference inputs are established about range and azimuth angles for a number of points that bound the available vehicle trajectory. The controller comprises three layers: upper, intermediate and lower. The upper layer can either be a vehicle stability control (VSC) mode which regulates the sideslip from a stability viewpoint or an autonomous vehicle controller which aids the driver. The intermediate layer seeks to minimise the predicted difference between the future and the desired responses of the system; the architecture of the controller is similar to [69, 58]. The lower layer involves a sliding mode control and active front steering based on fast slip control. The proposed controller is coupled to a CarSim model and simulated; the results suggest that the controller is comparable to an MPC formulation.

1.2.3 Control of high side-slip manoeuvres

The work done on high-sideslip manoeuvres is relevant to this project because modelling the vehicle during very nonlinear manoeuvres is one of the major challenges for high fidelity simulations. Various approaches have been proposed in recent years, each of which has benefits and drawbacks.

Voser *et al.* [70] considered various drifting equilibria and linearised the vehicle about these points. They used a simple bicycle model and analysed the dynamics of drifting manoeuvres assuming steady state conditions; they subsequently linearised

the vehicle around the equilibrium and designed a simple controller whose aim is maintaining the vehicle in steady state conditions. The benefits of this method are its simplicity and its ability to capture the drifting behaviour. However, its scope is too limited for any practical implementation.

Talvala [71] *et al.* (2011) designed a combined controller for autonomous racing. The work is based on studies on lane keeping. The vehicle model used is again a bicycle model with nonlinear tyres. The overarching idea relies on the minimisation of a potential field, which is found using an energy based approach. Both lateral and longitudinal dynamics are considered but they are not strongly coupled. The derived controller is stable and it has been successfully tested on a real autonomous vehicle. The main benefits of this approach are its robustness, simplicity and the ability to model the whole tyre curve. The drawbacks are its conservativeness and the poor performance at the tyre slip limit.

Keen and Cole [6] approached the problem of a vehicle in limit conditions by designing a MPC which stabilises a vehicle which is linearised at its current operating conditions. The model is an enhanced bicycle model with eight degrees of freedom to account for roll and weight transfer. Data from a real vehicle are used to fit the equations of motion. The front and rear tyre curves are linearised about several operating points to obtained an array of linearised vehicle models whose denseness and scope are meant to represent driver's skills. A LTV MPC controller is then designed. The benefits of this method are the accuracy in modelling the vehicle in the nonlinear region of the tyre curve, but before saturation, and the MPC computational efficiency; the only major drawback is the instability in the negative slope region of the tyre curve.

Tavernini [72] considered the problem of finding the optimum manoeuvre time for a vehicle cornering on different surfaces; the intent of the study is to compare different driving styles and understand why rally drivers perform very aggressive manoeuvres at very high slip angles. They used a bicycle model and formulated a constrained optimal control problem, finding the solution using an indirect method approach. Simulations are run for a variety of U-turns on surfaces characterised by different grip levels. Results show that high side-slip manoeuvres are more effective on gravel while low side-slip turns are better on high grip surfaces, which is what is observed in real racing. The benefits of Tavernini's approach benefits are the excellent modelling of the aggressive behaviour and the ability to model the entire tyre curve. The drawbacks include the lack of robustness, defined as the ability to drive the system to a stable state after a

perturbation, the lack of a model for the neuromuscular system and the computational burden, which prevents simulations from being run in real time.

Klomp [73] considered the optimal recovery of a vehicle from terminal under-steering. They considered a manoeuvre whereby the vehicle enters a corner too fast and needs to slow down. The manoeuvre is formulated as an optimal control problem which minimises the maximum off-tracking from the reference path. Two optimal control solutions are obtained: an analytical one for a friction-limited particle and the second a full numerical solution for a 4 wheels vehicle model; the agreement between the two methods is good. The benefits and drawbacks of this method are similar to those highlighted for Tavernini.

1.2.4 Minimum time manoeuvre

Timings [1] conducted an in-depth literature review on this topic as it was central to his PhD thesis. Since this project only makes use of Timing's results and does not focus on the generation of the nominal trajectory, the literature review for this section will be brief as it only seeks to contextualise the results in [1].

Minimum time manoeuvre problems consist of finding the minimum time a given vehicle can complete a lap in a track of given bounds. The first approaches to simulating racing laps were based on the quasi-steady-state (QSS) method [74], which approximates a lap as a sequence of steady state manoevures. The track is divided into several sections where the longitudinal acceleration of the vehicle is assumed to be constant. GG speed diagrams [75, 76] are used to determine the behaviour of the vehicle at the handling limits. In practice, the simulation works by calculating the velocity at the apex of each corner and extrapolating the dynamics backwards and forwards making use of the G-G diagrams. Even though this method is robust, gives sensible results and it is widely used in practice, the underling assumptions are inaccurate because the transient dynamics, which play an important role in racing, are neglected.

The problem of achieving a minimum laptime can be conveniently posed as one of optimal control. One of the first comprehensive studies was conducted by Hendrikx *et al.* [77]; they used a 3 DoF vehicle model with nonlinear tyre characteristics, load transfer, an engine torque curve and aerodynamic loads. In spite of the accuracy of the results, the model is not useful in practice because of excessive computational time.

Casanova studied the optimisation of a transient nonlinear race-car model over a complete lap of a circuit in [78]. The 7 DoF vehicle model features an engine map and a combined tyre slip model based on Pacejka's Magic Formula. The driver model was

replaced by a nonlinear optimiser which represents the skills of a very expert driver. The accuracy of the model resulted in a very computationally intensive algorithm.

Komatsu *et al.* in [79] are the first to apply time-varying linear optimal control theory to the vehicle path optimisation problem; they use a reference LQR controller coupled with an iterative method. It is found that the method is capable of determining an optimal path for limit cornering at a constant vehicle forward velocity. The computational cost is high.

Velenis and Tsiotras used different optimisation functions to analyse the driving styles of different racing drivers. In [80–82], they focused on the problem of minimum time manoeuvring for high-speed autonomous vehicles, focussing on real-time implementation.

One of the most accurate and comprehensive documented lap-time simulator to date is based on optimal control and was elaborated on by Kelly [23]. They consider transient response and tyre dynamics; their method does not attempt to model the driver but rather replaces the driver with a numerical nonlinear optimal controller with constraints as performance limits. The objective function is set up such that the method is able to simultaneously find the optimal racing line and driver control inputs. Kelly was also able to study the performance and stability of a vehicle running over smooth surfaces and simulated kerbing. In order to capture a variety of physical effects, such as frequency response over kerbs, each vehicle component was modelled separately and then as an assembly. These vehicle models are commonly referred to as *multibody models* as they assign independent degrees of freedom to different vehicle components. The final analysis conducted looks at the effects of a thermodynamic tyre model, characterised by temperature dependent parameters, on vehicle performance. However, their pursuit of robustness and accuracy has left the approach needing considerable computational resources, with single manoeuvre execution times typically being measured in number of GPU hours.

Gerdts *et al.* [83] used Nonlinear Model Predictive Control (NMPC) to generate local optimal trajectories which are combined by suitable continuity conditions to form paths of minimum time around a number of circuits. The vehicle model used is a single-track model with 3 DoF, aerodynamic drag, rolling resistance and a Magic Formula tyre model.

Perantoni *et al* [84] tackled the problem of computational inefficiency by using direct collocation methods which preserve the complexity of the model at a relatively low computational cost. The algorithm generates an optimal trajectory, driver inputs and optimises the vehicle set-up for a given track. The authors claim that, indicatively, a

simulation taking several hours would take 15 minutes on a standard desktop computer using their method. These improvements are the results of various refinements of previous models such as the use of curvilinear coordinates, model non-dimensionalisation and scaling, the avoidance of 'stiff' dynamics and more efficient numerical approaches.

Limebeer *et al.* [85] formulated an efficient optimal control problem to control hybrid vehicles to overcome problems associated to excessively long computational times which arise when modelling such systems. Hybrid vehicles are vehicles with at least two energy sources and they combine two modes of propulsion. The control problem is then more complex due to the increased complexity of the dynamics and the extra number of constraints. While traditional numerical techniques employed in the literature, such as direct multiple shooting, would take an unreasonable time to converge, the paper focuses on the orthogonal collocation method, which is only now being accepted as a sensible solution for a constrained control problems. An optimal control problem is formulated to find the optimal lap-time of a hybrid racing vehicle conforming with the F1 2016 regulations on energy recovery systems. The vehicle model is a four wheel vehicle with several features such as a locking differential, brake balance and weight transfer. The track is described using intrinsic coordinates. The cost function aim was lap time minimisation and fuel efficiency. The problem is then solved using a mesh refined orthogonal collocation method, which resulted in a real time solution for the SPA circuit in Belgium on a standard Intel i7 processor.

Lot *et al.* [86] solved the minimum lap time optimal control problem for a go-kart model using a similar approach. Equations of motion and the cost function Hessians are derived symbolically. The optimisation problem is then formulated and solved using indirect methods. The vehicle model used is a seven degrees of freedom vehicle model that captures all the characteristics of a racing go-kart, such as frame bending and wheel compliance. The model has been validated against actual track data. After confirming the reliability of the model, specific go-kart dynamics, such as tyre slippage and the role of the differential were studied.

Canerini *et al.* [87] studied the optimization of an innovative steering and roll systems for a four wheel vehicle. The system is a hybrid between a four and a two wheels vehicle: the driver operates the handlebars to steer and the throttle pedals to control speed. The driver can also control the roll movements through two pedals. A kinematic optimisation of all the components of the novel mechanism, i.e. handlebar, column, pinion, rack and track-rods to the wheels is performed to minimise undesired roll behaviour and provide appropriate outer/inner wheel steer angles.

Jardes *et al.* [88] focused on the evolution of modern motorsport limited slip differentials (LSD) which has allowed a high degree of flexibility in terms of tuning, effectively allowing an adjustment in the torque bias generated in the most salient phases of track manoeuvres. A particularly relevant task consists of finding the optimal torque bias profile under different braking and accelerating conditions. Since this task is prohibitively expensive in terms of computation time, a simple lane change manoeuvre is considered first. Vehicle stability and agility is evaluated using well known metrics and compared to open and closed differentials. The optimisation problem showed that the optimal LSD profile can give a performance gain over its locked differential counterpart where a quick direction change is required.

Velenis *et al.* [89] studied haptic steering support for a vehicle close to its handling limit. The device is aimed at promoting the driver's perception of the vehicle's behaviour by supplying torque cues on the steering wheel. Driving simulator and real vehicle tests are used to evaluate the proposed support; the scenario used is a skid-pad where drivers attempt to retain control of the vehicle in limit conditions. Both tests showed that the device helped drivers driving closer to the target path while reducing the cognitive load.

Similarly to Tavernini *et al.* [90], Velenis et al [91] investigated the minimum time manoeuvring problem using nonlinear optimal control techniques on different road surfaces. An optimal virtual driver is devised using nonlinear optimisation. Different transmission layouts and terrain types are considered and all the relative problems solved. Results showed that many of the driving strategies commonly observed in motorsports, such as hand braking or late cornering are in fact better for a specific vehicle configuration.

Even though optimisation has been widely used, other techniques have been employed for lap-time optimisation. Thommyppillai *et al.* [74, 92, 93] applied linear theory to the minimum manoeuvre time problem by considering it as one of accurate path tracking, which is well studied in the literature. They argue that the variation in racing line for different racing vehicles on a given track is negligible and that it is reasonable to assume that the optimal racing line is the one tuned by racing drivers over the years.

All the models considered so far assume that lap-time minimisation is independent of driver workload; cognitive effort is particularly high at the handling limit, an often unstable operating point. An important contribution of this work is accounting for driver workload in lap-time optimisation, in an attempt to provide a more realistic simulation tool. This is achieved by developing and validation a model that accounts for

driver's physical and cognitive limitations and can operate at the vehicle handling limit. Said model can then be used for lap-time minimisation. The overarching, long-term, purpose is shifting the focus from the vehicle properties to closed-loop driver behaviour.

1.3 Nonlinear robust control

Uncertainties play a very important role in lap time simulations because there are factors such as road roughness, wind speed or neuromuscular noise that cannot be modelled deterministically. Considering disturbances improves both the accuracy and the robustness of the model. Robust Model Predictive Control schemes (RMPC) are used to deal with disturbances; the most common RMPC schemes are: dynamic programming solution, feedback model predictive control, Min-Max predictive control and Tube-based model predictive control. Since the model developed in this thesis extends, among other things, the robust controller Timings developed in [1], the relevant literature is similar.

Dynamic Programming is a scheme developed by Bellman [94] and involves multi-stage optimisation. An equation, called the Bellman equation, is derived and if it is valid over the domain of the function that has to be optimised, the problem is easily solved using numerical algorithms. The limitation of this method is that it cannot be used in conjunction with MPC for a system that contains a random variable.

Feedback model predictive control approaches, such as those in [95] and [96], seek to calculate optimal control laws instead of control actions at each time step. This scheme modifies the system dynamics to reject disturbances. The resultant problem, however, is generally too complex to be amenable to a computer solution.

Min-Max model predictive controller aims at reducing the computational burden of feedback model predictive control by solving a nonlinear Min-Max optimal control problem at each time instant. The solution is a control input that satisfies all the constraints for all possible realisation of the uncertainty. Since the worst case scenario is taken into account, the scheme is conservative. The method has been applied successfully in [79, 97] but has proven to be dependent on the magnitude of the disturbances in [98].

Tube-based model predictive control is a method that is as efficient as the Min-Max but it is not as conservative: it is therefore the ideal candidate for this type of project. A review of tube-based model predictive control follows

Mayne and Kerrigan [99], building on a series of papers by Rakovic *et al.* [100, 101] *et al.*, developed a linear predictive control scheme based on the solution of two

MPC problems: one for the nominal trajectory with artificially contracted boundary conditions and one for rejecting disturbances. Feasibility is guaranteed if the tube is chosen in a way such that its outer limits stay within the original state constraints. In a more recent work, Mayne *et al.* [102] extend their work to nonlinear systems. The nominal controls are found first to determine the centre of the tube; these can be obtained either using any control policy or an heuristic routine. The ancillary controller is added to the system to reject disturbances. The deviation of the disturbed states from the nominal will be minimised in the cost function. Mayne *et al.* prove that the states are bound if the first instance of the control sequence is applied, so stability is not guaranteed if a longer sequence is applied.

Timings *et al* [1] apply tube based nonlinear control to a racing context. The nominal trajectory becomes the optimal vehicle path and speed profiles generated by the perfect driver model using and MPC based algorithm. Disturbances are then added and constraints redefined to account for the external disturbance forces. An ancillary LQR controller is added to reject disturbances; the statistics of the response are calculated in an iterative fashion and they determine the width of the *tube* within which the states are bounded. The main limitations of Timings work are the computational time required, the use of a linear anciliary controller and the incapability of dealing with the negative slope part of the tyre force curve. Sideris [40] and then Haslam [41] managed to improve the efficiency of the algorithm dramatically by calculating the statistics of the response using the covariance equation instead of an iterative method.

Johns [2] extended the work of Timings by considering the relationship between cognitive effort and laptime, hence refining Timings driver model. They conducted some experiments to validate the hypothesis that cognitive effort and lap-time are inversely related. They subsequently designed various controllers to account for the driver limited cognitive ability. It was found that an intermittent controller that utilises serial ballistic sequences of controls represents the experimental data better than other controllers. This controller was used to simulate a whole lap and it was found that an increase in the controller intermittency period could mimic a lap-time increase due to increasing cognitive load. Even though the controller is able to replicate the results from the experiments, its applicability to a vehicle driving at the limit of adhesion is questionable because it is based on linear assumptions. Robust model predictive control has also been used in the robotic literature [103] and [104].

1.4 Summary

A robust mathematical model for a driver is key to the development of low-cost design and tuning of race cars. The literature review has shown that the problem can be approached in different ways, each of which has benefits and drawbacks; it is clear that there is still room for development in various areas. The current state of knowledge is summarised in the next sections.

1.4.1 Nonlinear vehicle dynamics

The tyre model is the main source of nonlinearities. The most widely used are semi-empirical models which assume an analytical slip force relationship from trends observed in the experiments. Pacejka's Magic Formula stands out as the most accurate and convenient model. A nonlinear tyre model can be coupled with a vehicle model, the simplest being the bicycle model, which only models lateral dynamics. The longitudinal dynamics are usually captured through a trolley model, which allows for throttle control. It can be coupled with nonlinear tyres to gain insights into the behaviour of the vehicle through graphical means, such as handling diagrams.The trolley can be considered in conjunction with the bicycle model to have lateral and longitudinal controls.

Assuming a linear tyre model, a closed form solution to the bicycle model can be found and its stability assessed. Phase diagrams are used to represent the state space solution of the bicycle mode and identify stability boundaries. Frozen time eigenvalues are used to predict the disturbed behaviour of the vehicle. Stability and controllability derivatives are used to characterise the tendency to oversteer or understeer and the quantity of damping in the system. All the metrics so far described focus on the open-loop dynamics of the vehicle without taking the closed loop behaviour of the driver into account.

Nonlinear vehicle dynamics is a relatively mature field of research; however, there is the need for better stability and controllability metrics that can account for the driver's physical and cognitive limitations. This project aims at expanding the work of Sideris [40] and Haslam [41] to validate their results and further develop the proposed metrics.

1.4.2 Modelling the perfect driver

The problem of modelling a perfect driver is formulated as an optimal control problem, typically using an MPC or an LQR controller. These models, however, represent a

'perfect' driver and do not generally take into account human limitations. The most common and practical approach for minimum time manoeuvring relies on approximating the motion of the vehicle as a series of quasi-steady manoeuvres. The most successful attempts, however, have been those relying on nonlinear optimisation. These models are very accurate; nevertheless, they are very computationally intensive and do not take into account human limitations in a clear and efficient way.

The challenge for future research is the modelling of an *imperfect* driver, which takes into account the randomness in the neuromuscular system and random external disturbances.

1.4.3 Robust control

Two of the most common approaches to deal with uncertain control problems are Min-Max MPC, where a single sequence of control actions is used to minimise the worst case, or feedback min-max MPC where the worst case is minimised over a sequence of control laws. However, these methods are conservative because they take into account worst case scenarios. Tube-based model predictive control overcomes these limitations by finding the controls to generate an undisturbed nominal trajectory and superimposing the disturbance rejecting controls. Timings [1] applies tube based nonlinear control to a disturbed vehicle in a racing context by finding optimal controls and superimposing the action of an ancillary LQR controller. The statistics of the response are calculated in an iterative fashion which is very computationally intensive. The limitations of the work are the inability of dealing with the negative part of the tyre curve, the assumption that the driver has perfect knowledge of the vehicle, the lack of feedback constraints and the lack of control strategy variation. Sideris [40] and then Haslam [41] managed to improve the efficiency of the algorithm dramatically by calculating the statistics of the response using the covariance equation.

The shortcomings of Timings' algorithm can be overcome by applying nonlinear tube-based MPC to the lap-time problem. The controller has to be able to simulate the behaviour of a of a human driver in limit conditions at an acceptable computational power.

1.5 Research objectives

In order to circumvent some of the limitations in Timings' work and design an improved *imperfect* driver model, the following research objectives are proposed.

1. Complete the work of Sideris and Haslam in using the variance prediction equation to quantify the stability and controllability of a car performing a manoeuvre.

2. Extend the work of Timings and Johns to develop an MPC-based or other suitable compensatory controller that is appropriate for the nonlinear handling regime.

3. Implement and assess the effect of the driving strategies proposed by Johns.

4. Explore the performance of the controller, particularly the coupling of the lateral and longitudinal dynamics.

5. Devise and perform driving simulator experiments in order to identify driver behaviour and to validate the theoretical developments.

Chapter 2

Quantification of road vehicle handling quality using a compensatory steering controller

2.1 Introduction

Automotive engineers tasked with tuning the handling behaviour of a vehicle often use quantitative objective criteria to assess stability and controllability under various operating conditions. An ideal criterion would be predictable from the design parameters of the vehicle and correlate well with subjective assessments. Many criteria have been used in the past, but these have not taken account of the closed-loop behaviour of the driver-vehicle system.

Timings [105] adopted ideas from tube-based robust model predictive control. A continual random disturbance was applied to a simulated nonlinear vehicle negotiating a lap of a circuit. A simulated driver was assumed to perform closed-loop feedback control to compensate for the effect of the random disturbance and keep the vehicle close to the optimal path. The trade-off between the variance of the compensatory steering control (a measure of the driver's physical workload) and the lap time was quantified. A disadvantage of the method was that the variance of compensatory steering control was determined from an ensemble 1000 time domain simulations, which was computationally expensive.

The work described in the present Chapter extends the approach in [105] by calculating the variance of the simulated driver's compensatory control responses directly, rather than by simulating an ensemble of time domain responses. In addition, the way

in which the variance of the compensatory responses changes through the manoeuvre is examined, and compared against several existing stability and controllability criteria. It is anticipated that the variance of the driver's compensatory control action and the variance of the vehicle's lateral path displacement from the nominal path might relate to the driver's subjective assessment of the stability and controllability of the vehicle. In addition, by examining the way in which the variances change through the manoeuvre, it might be possible to tune the vehicle to behave in a desirable manner at various stages of the manoeuvre. The proposed new criteria differ from existing criteria by accounting for the closed-loop dynamics of the driver and vehicle, rather than considering only the open-loop dynamics of the vehicle without consideration of the driver.

Section 2.2 describes the nonlinear vehicle model and Section 2.3 summarises the algorithm used to calculate the nominal driver controls (drive/brake torque and steering angle) to negotiate a ninety-degree bend in minimum time. The compensatory steering controller is presented in Section 2.4, and the method for calculating the variance of its response is set out in Section 2.5. Results of the variance calculation for two different vehicles are presented in Section 2.6 and comparisons to existing stability and controllability criteria are discussed in Section 2.7.

The work documented in this Chapter is the result of the collaboration of multiple authors. The nominal controls are found using Timings MPC scheme [105]. Sideris [40] carried out preliminary work on stability and controllability metrics, focusing on comparing them. Haslam [41] developed the closed loop LQR controller and performed the variance calculations. The author of this Chapter reviewed the work, reproduced the results to ensure consistency and wrote a paper [106], which shares most of its elements with this Chapter.

2.2 Vehicle model

The vehicle model is similar to that employed in [105]. The lateral and yaw dynamics of the vehicle are represented by the familiar single-track 'bicycle' model with five degrees of freedom as shown in Fig. 2.1. The model complexity is appropriate for demonstrating the new stability and controllability criteria, but a more detailed model would likely be required to investigate the performance of a specific vehicle. The forces generated by left and right tyres on an axle are combined and throughout this Chapter all parameter values associated with the tyres relate to the combined left and right tires on the axle.

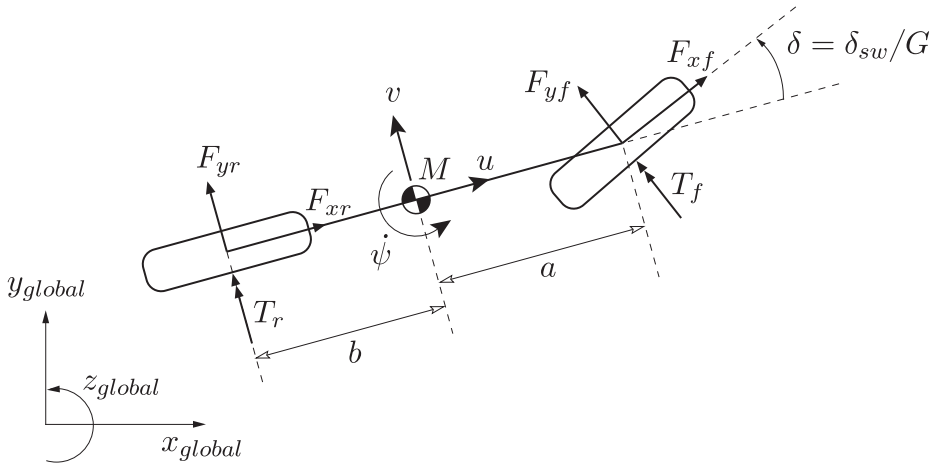

Figure 2.1: Vehicle model with associated forces and dimensions.

The equations of motion are:

$$M(\dot{v} + u\dot{\psi}) = F_{yf}\cos\delta + F_{xf}\sin\delta + F_{yr} + F_y \tag{2.1a}$$

$$I_z\ddot{\psi} = a(F_{yf}\cos\delta + F_{xf}\sin\delta) - bF_{yr} + M_z \tag{2.1b}$$

$$M(\dot{u} - v\dot{\psi}) = F_{xf}\cos\delta - F_{yf}\sin\delta + F_{xr} \tag{2.1c}$$

$$I_f\ddot{\theta}_f = T_f - F_{xf}r_f \tag{2.1d}$$

$$I_r\ddot{\theta}_r = T_r - F_{xr}r_r \tag{2.1e}$$

The axis sign convention for the vehicle model is included in Figure 2.1.

A baseline vehicle configuration which represents a standard saloon vehicle equipped with nonlinear tyres is used. Even though this work focuses on racing vehicles, there is no need to model one at this stage, as this is done for the validation in Chapters 5 and 6. The baseline configuration is US and not particularly challenging to handle. The vehicle is chosen so as to have a stable, predictable behaviour in the linear region.

The OS vehicle configuration is achieved by shifting the vehicle centre of mass rearwards, setting $a = 1.38$ m and $b = 0.92$ m, effectively increasing the cornering stiffness of the rear tyre as shown in [7]. Assuming effects causing a constant force are negligible, cornering stiffness is proportional to the vertical load on the tyre as per Eq (3.23). Shifting the centre of mass towards the back of the vehicle increases the load F_z on the rear tyre, hence the slope of the rear tyre curve. OS vehicles are unstable beyond critical speed, requiring closed loop control to remain stable.

The parameters and their nominal values are defined in Table 2.1. δ is the front road wheel steer angle, $T_{f,r}$ the axle torques and $\theta_{f,r}$ the wheel/axle angular displacements.

The subscripts f,r indicate front and rear and x,y indicate longitudinal and lateral. A lateral force F_y and a yaw moment M_z applied at the centre of mass of the vehicle act as disturbances, such as might arise from road roughness, friction variations or wind gusts. A single control torque T represents a combined drive action (if $T > 0$) and brake action (if $T < 0$). Assuming the vehicle is rear wheel drive, the torque T is distributed to the front and rear wheels such that

$$T_f = b_f(1 - H(T))T \tag{2.2}$$

$$T_r = T - T_f \tag{2.3}$$

where $H(\cdot)$ is the Heaviside step function and b_f is the brake balance (fraction of braking torque applied to the front axle).

The tyre forces are expressed as functions of the lateral and longitudinal slips. The effect of longitudinal weight transfer and camber angle have not been taken into account as racing vehicles are characterised by a low centre of mass. Furthermore, the purpose of this Chapter is to investigate the behaviour of the controller rather than achieving a high fidelity simulation. Concerning the sign convention, downwards vertical forces are considered positive and anticlockwise slip angles are considered positive.

The slips are defined as:

$$\alpha_f = \delta - \frac{v + \dot{\psi}a}{|u|}, \qquad \alpha_r = -\frac{v - \dot{\psi}b}{|u|} \tag{2.4}$$

$$\kappa_f = \frac{\dot{\theta}_f r_f - u}{|u|}, \qquad \kappa_r = \frac{\dot{\theta}_r r_r - u}{|u|} \tag{2.5}$$

where α_j is the lateral tyre slip, κ_j is the longitudinal tyre slip and j is f or r. A normalised slip vector is defined [107]:

$$\mathbf{s}_j = \begin{bmatrix} s_{xj} \\ s_{yj} \end{bmatrix} = \frac{C_{\alpha j}}{F_{pj}} \begin{bmatrix} \kappa_j \\ \tan\alpha_j \end{bmatrix} \tag{2.6}$$

where s is the normalised tyre slip and $C_{\alpha j}$ is the normalised tyre cornering coefficient.

The friction circle limits F_{pj} are functions of the static axle loads F_{zj} and of the form:

$$F_{pj} = \frac{F_{zj}}{1 + (\frac{2F_{zj}}{3Mg})^3} \tag{2.7}$$

where F_{zj} is the vertical force on the axle. For the purpose of the present study the height of the vehicle's centre of mass is set to zero, to minimise the number of

Table 2.1: Vehicle model parameters and nominal values. Tyre parameter values are per axle.

Parameter	Symbol	Value
Mass	M	1050 kg
Moment of inertia about z axis	I_z	1500 kgm^2
Front axle to CoM distance	a	0.92 m
Rear axle to CoM distance	b	1.38 m
Wheel radius	r_f, r_r	0.28 m
Wheel/axle moment of inertia	I_f, I_r	2 kgm^2
Front brake balance	b_f	0.6
NMS natural frequency	ω_n	18.85 rad/s
NMS damping factor	ξ_n	0.707
Steering gear ratio	G	17
Magic formula coefficient	B	1.03
Magic formula coefficient	C	1.60
Magic formula coefficient	D	1.36
Magic formula coefficient	E	0.00
Tyre coefficient	c_1	69 kN/rad
Tyre coefficient	c_2	1.4 kN
Gravitational constant	g	9.81 m/s^2
Discrete time step	T_d	0.02 s

parameters in the vehicle model and make the results as generic as possible. Thus only static vertical forces are considered so:

$$F_{zf} = \frac{b}{a+b}Mg, \quad F_{zr} = \frac{a}{a+b}Mg \tag{2.8}$$

The slip coefficient $C_{\alpha j}$ is defined as:

$$C_{\alpha j} = c_1 \left(1 - \exp\left(-\frac{F_{zj}}{c_2}\right)\right) \tag{2.9}$$

where c_1 and c_2 are tyre coefficients. Tyre forces can finally be expressed as:

$$\begin{bmatrix} F_{xj} \\ F_{yj} \end{bmatrix} = P(|\mathbf{s}_j|))\frac{F_{pj}}{|\mathbf{s}_j|}\begin{bmatrix} s_{xj} \\ s_{yj} \end{bmatrix} \tag{2.10}$$

where the function $P(s)$ is Pacejka's magic formula [7]:

$$P(s) = D\sin(C\arctan(Bs - E(Bs - \arctan(Bs)))) \tag{2.11}$$

The values of these coefficients, as well as c_1 and c_2, are shown in Table 2.1, and are identical for the front and rear axles. Lateral front tyre forces for a range of longitudinal slips are shown in Fig. 2.2.

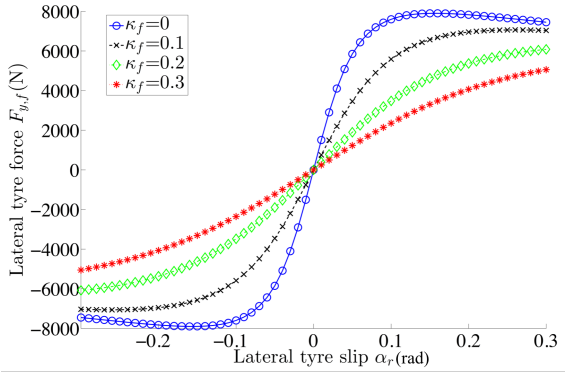

Figure 2.2: Lateral tyre force for one axle. The curves show four different levels of longitudinal slip κ. The vertical axle force F_{zj} is 6000 N.

The bandwidth-limiting effect of the driver's neuromuscular system (NMS) is represented by a second order low-pass filter acting on the hand wheel angle input δ_{sw} to the vehicle [108, 56], given by:

$$\ddot{\delta}_{sw} + 2\xi_n\omega_n\dot{\delta}_{sw} + \omega_n^2\delta_{sw} = \omega_n^2\delta_{com} \tag{2.12}$$

where δ_{com} is the commanded hand wheel angle, ξ_n and ω_n are the damping ratio and natural frequency of the NMS; values for these parameters were informed by [109]. The road wheel steering angle is given by

$$\delta = \frac{\delta_{sw}}{G} \tag{2.13}$$

The nonlinear vehicle dynamics equations can be expressed as

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}) \tag{2.14}$$

$$\mathbf{z} = g(\mathbf{x}) \tag{2.15}$$

with state vector

$$\mathbf{x} = [v \;\; \dot{\psi} \;\; \psi \;\; u \;\; \dot{\theta}_f \;\; \dot{\theta}_r \;\; \dot{\delta}_{sw} \;\; \delta_{sw}]^T \tag{2.16}$$

and control input

$$\mathbf{u} = [\delta_{com} \;\; T]^T \tag{2.17}$$

A linearised representation of (2.14) and (2.15) is achieved by performing a linearisation about an arbitrary, possibly non equilibrium point $(\mathbf{x}_0, \mathbf{u}_0)$ to give

$$\dot{\mathbf{x}}_c = \mathbf{A}_c\mathbf{x}_c + \mathbf{B}_c\mathbf{u}_c + \mathbf{F}_c \tag{2.18}$$

$$\mathbf{z}_c = \mathbf{C}_c\mathbf{x}_\mathbf{c} + \mathbf{G}_c \tag{2.19}$$

where $\mathbf{x}$ is the state vector and the matrices relate to the Jacobians through

$$\mathbf{A}_c = \left.\frac{\partial f(\mathbf{x},\mathbf{u})}{\partial \mathbf{x}}\right|_{\mathbf{x}_0,\mathbf{u}_0}, \quad \mathbf{B}_c = \left.\frac{\partial f(\mathbf{x},\mathbf{u})}{\partial \mathbf{u}}\right|_{\mathbf{x}_0,\mathbf{u}_0} \tag{2.20}$$

$$\mathbf{F}_c = f(\mathbf{x}_0,\mathbf{u}_0) - \mathbf{A}_c\mathbf{x}_0 - \mathbf{B}_c\mathbf{u}_0 \tag{2.21}$$

$$\mathbf{C}_c = \left.\frac{\partial g(\mathbf{x})}{\partial \mathbf{x}}\right|_{\mathbf{x}_0}, \quad \mathbf{G}_c = g(\mathbf{x}_0) - \mathbf{C}_c\mathbf{x}_0 \tag{2.22}$$

For more details of the linearisation see [105, 110]. Following common practice in the field of Model Predictive Control the system is then discretised. Zero order hold discretisation is used; the subscript k indicates that the object is in discrete time and evaluated at time $t = T_d k$, where T_d is the discrete time step and k is an integer. The system of equations takes the form

$$\mathbf{x}_{k+1} = \mathbf{A}_k\mathbf{x}_k + \mathbf{B}_k\mathbf{u}_k + \mathbf{F}_k \tag{2.23}$$

$$\mathbf{z}_k = \mathbf{C}_k\mathbf{x}_k + \mathbf{G}_k \tag{2.24}$$

2.3 Minimum manoeuvre time calculations

A single 90° bend shown in Fig. 2.3 is considered. The track is 10 m wide; it starts with a straight section of 360 m, followed by a 90° bend of length 100 m whose centreline has a constant radius of 63.7 m. A long straight section follows, with the simulation terminating 40 m into this. In order to minimise manoeuvre time, the driver has to maximise speed while attempting to minimise the distance travelled. The two requirements are clearly conflicting. Following [110] the minimum nominal manoeuvre time for the undisturbed vehicle described by (2.23) and (2.24) is obtained

using an efficient MPC formulation. The procedure rests on the idea of gathering information about the track through a receding horizon which mimics the line of sight of a real driver. At each time step the vehicle is linearised about the current operating conditions and future displacements are predicted over the whole prediction horizon. A quadratic cost function is then formulated to maximise the distance travelled over the time horizon, hence minimising laptime. The cost function takes the following form:

$$\begin{aligned} \min_{\theta_{(k)}} J_{(k)} = \ & \frac{1}{2}\theta_{(k)}^T H_{(k)} \theta_k + \eta_{(k)}^T \theta_{(k)}^T \\ \text{subject to} \quad & \lambda_{(k)}\theta_{(k)} = \gamma_{(k)} \\ \text{subject to} \quad & \Omega_{(k)}\theta_{(k)} \leq \omega_{(k)} \end{aligned}$$

where $J_{(k)}$ is the cost function, $\theta_{(k)}$ a vector with the input and controls for the system, $H_{(k)}$ a weight matrix, η_k, $\gamma_{(k)}$, $\omega_{(k)}$ and $\Omega_{(k)}$ matrices related to the system constraints.

The derivation of the above matrices is very involved so it will be omitted here as it is outside the scope of this Chapter. The interested reader can refer to [110] for further details.

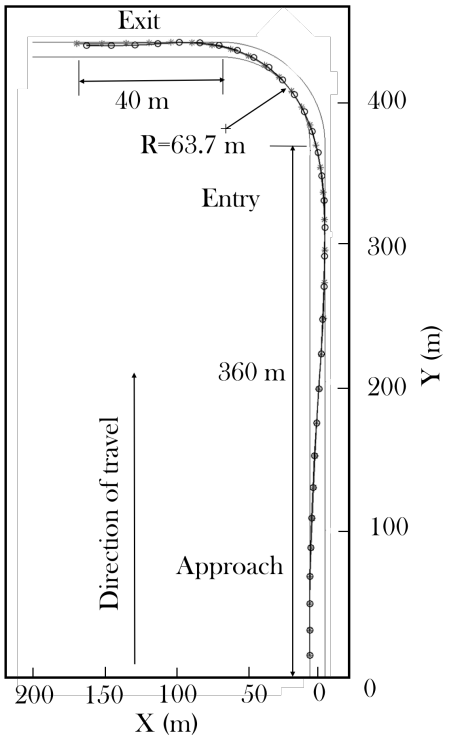

Figure 2.3: The road boundaries and the optimal paths of the US and OS vehicles.

Optimal controls δ_{sw} and T are calculated for two vehicle configurations. One is defined by the nominal parameter values in Table 2.1 and has a steady-state understeering (US) characteristic. The other is defined by the same parameter values except for $a = 1.38$ m and $b = 0.92$ m and has an over-steering characteristic (OS) due to the rearward position of the centre of mass (CoM). The brake balance b_f is the same for both vehicles. The vehicle begins on the left-hand boundary at the start of the first straight with speed 30 m/s. The optimisation involves some constraints: the vehicle is required to remain within the road boundaries; the maximum drive torque is 2 kNm; the maximum allowable tire slip at any time is that which gives 0.99 of the maximum feasible horizontal tyre force at that time.

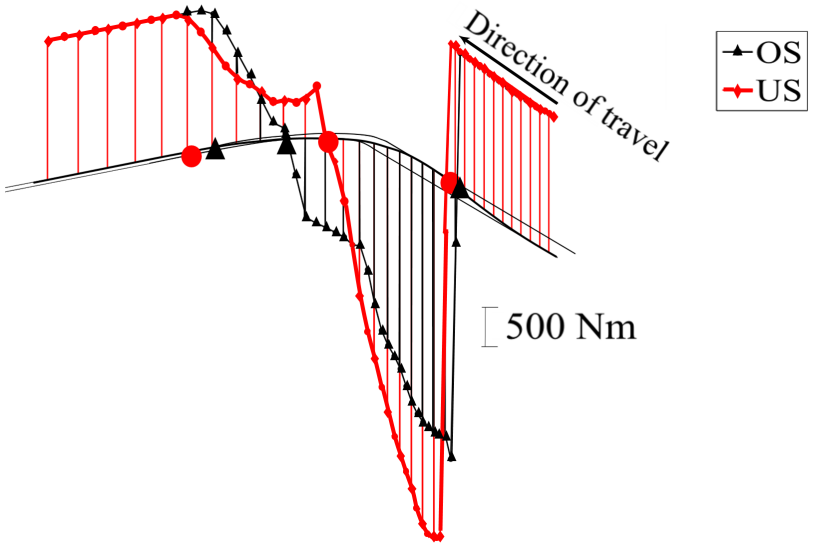

Figure 2.4: Optimal torque (T) through the manoeuvre for the US and OS vehicles. The torque is positive (accelerating) if the line is above road level and negative (braking) when the line is below road level. Both vehicles accelerate in the first part of the manoeuvre, then brake at corner entry and accelerate at the exit. The vertical lines correspond to data at time intervals of 0.4 s. The triangles and circles plotted at road level correspond to the three phases of the manoeuvre: braking on entry; transition from braking to accelerating at mid corner; and maximum drive torque at exit.

Fig. 2.4 shows the optimal torque control for the US and OS vehicles. The vertical lines correspond to data at time intervals of 0.4 s. The constraint on drive torque means that the torques are equal whilst accelerating on approach to the corner. However the braking point for the US vehicle is a little later than for the OS vehicle and the total braking torque is greater, due to the front-biased brake balance and forward CoM position of the US vehicle. After the apex of the corner the US vehicle begins

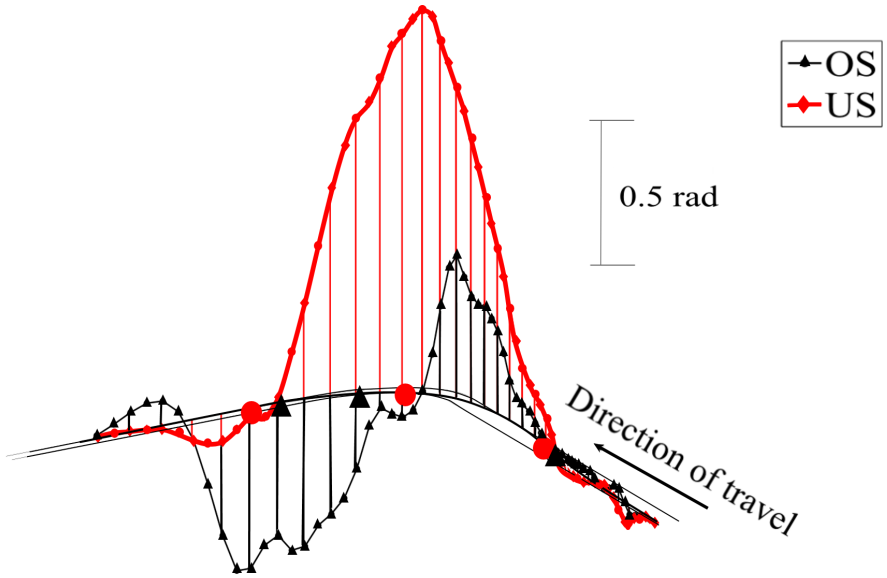

Figure 2.5: Optimal hand wheel angle (δ_{sw}) through the manoeuvre for the US and OS vehicles. The steering angle for the US vehicle is positive through most of the corner while it is positive at the beginning of the corner for the OS vehicle and negative (countersteering) towards the exit.

accelerating earlier than the OS vehicle, but the OS vehicle applies maximum drive torque earlier than the US vehicle. The triangles and circles plotted at road level correspond to these three phases of the manoeuvre: braking on entry; transition from braking to accelerating at mid corner; and maximum drive torque at exit. Fig. 2.5 shows the optimal hand wheel angle for the US and OS vehicles. Larger angles are required for the US vehicle compared to the OS vehicle. Another notable difference is that a significant countersteering action (negative angle) is applied to the OS vehicle after the apex of the corner. Comparing the torque and steering controls it can be observed that for both vehicles on entry to the bend the steering begins at about the same time as the switch from acceleration to braking. After this point the braking torque tends to reduce from its peak value as the steering angle increases. This ensures that the constraint on combined longitudinal and lateral slip is satisfied. For the US vehicle the peak steering angle occurs at about the same time that braking torque reaches zero, which is consistent with the slip at the front axle being limiting. Maximum drive torque is applied to the US vehicle at the point where the steering angle returns to near zero. The coordination of the steering and torque controls of the OS vehicle is different to the US vehicle and consistent with the slip at the rear axle being limiting. The countersteering action begins just before the switch from braking to acceleration and continues for a short while after maximum drive torque is reached.

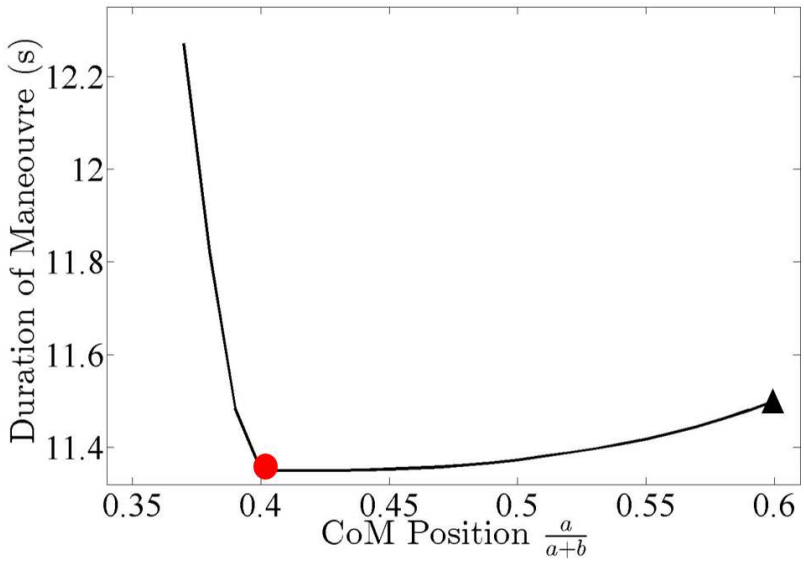

Figure 2.6: Relationship between manoeuvre time and centre of mass position. The optimal value to minimise manoeuvre time is 0.42.

The effect of CoM position, defined as $a/(a+b)$, on time to complete the manoeuvre is shown in Fig. 2.6. The brake balance is fixed to the value given in Table 2.1. Minimum time occurs at $a/(a+b) = 0.42$, which is close to the value of 0.4 specified for the US vehicle.

2.4 Response variance calculation

In the preceding section the optimal controls to achieve minimum manoeuvre time are calculated. These controls could in principle be applied in a feedforward, open-loop fashion to the vehicle, and in the absence of disturbances and other uncertainties, the nominal trajectory would be achieved. In practice the driver will be required to perform an additional feedback, closed-loop control in order to compensate for disturbances and other uncertainties, and to stabilise the vehicle if necessary. It is the proposition of this Chapter that the driver's compensatory control action and corresponding vehicle response provide a practical and objective way of quantifying the handling quality of the vehicle as it travels through the manoeuvre.

It will be assumed that the primary function of the driver's compensatory control is to minimise lateral deviation of the vehicle from the nominal optimum trajectory calculated in Section 2.3. This is an appropriate assumption for a racing driver, but drivers in other situations may have different objectives.

The state space equation (2.18) is augmented to include the lateral displacement as a ninth state.

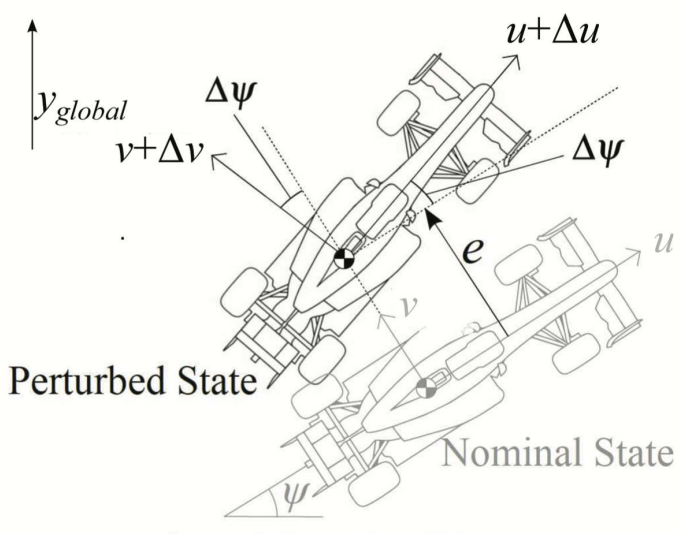

Figure 2.7: Vehicle in the nominal and perturbed state. The lateral path error e is shown together with nominal and perturbed longitudinal velocity u, lateral velocity v, and yaw angle ψ.

Referring to Fig. 2.7 the perturbed longitudinal and lateral velocities are $u + \Delta u$ and $v + \Delta v$ and the perturbed yaw angle is $\psi + \Delta\psi$, so that the time derivative of the lateral path error, neglecting second order terms, can be expressed as:

$$\begin{aligned} \dot{e} &= (v + \Delta v) + (u + \Delta u)\Delta\psi - v \\ &= \Delta v + u\Delta\psi \end{aligned}$$

where $\Delta\psi$ is assumed small. In order to include this expression in the state space equation an additional state is considered:

$$\dot{y} = v + u\psi \tag{2.25}$$

This equation is not valid for the large angles ψ that might arise in the nominal case. However, when small perturbations are considered about a linearisation point y_0, (2.25) reduces to

$$\begin{aligned} \dot{y}_0 + \Delta\dot{y} &= (v_0 + \Delta v) + (u_0 + \Delta u)(\psi_0 + \Delta\psi) \\ &= (v_0 + u_0\psi_0) + (\Delta v + u_0\Delta\psi + \psi_0\Delta u) \\ \Leftrightarrow \Delta\dot{y} &= \Delta v + u_0\Delta\psi + \psi_0\Delta u \end{aligned}$$

If Δy is to be equivalent to e then it is clear that the linearisation point must be taken about u_0, the current nominal longitudinal velocity. However the linearisation must also be about $\psi_0 = 0$ rather than the true current yaw angle, to ensure that the lateral path error is independent of the nominal heading angle in the global reference frame.

Disturbances $\mathbf{w}$ are considered to cause small perturbations $\Delta\mathbf{x}_k$ about the current nominal state $\bar{\mathbf{x}}_k$, with nominal control input $\bar{\mathbf{u}}_k$. The state space equation can be expressed as

$$\mathbf{x}_{k+1} = \bar{\mathbf{x}}_{k+1} + \Delta\mathbf{x}_{k+1} = \mathbf{A}_k(\bar{\mathbf{x}}_k + \Delta\mathbf{x}_k) + \mathbf{B}_k(\bar{\mathbf{u}}_k + \Delta\mathbf{u}_k) + \mathbf{H}_k\mathbf{w}_k + \mathbf{F}_k \tag{2.26}$$

where $\mathbf{w}_k$ are the disturbances and $\Delta\mathbf{u}_k$ is the compensatory control action added to the nominal control action.

Disturbances are considered to arise from random lateral force and yaw moment acting at the centre of mass, and random additional handwheel angle arising from neuromuscular noise. It is assumed that all three disturbances are zero mean, Gaussian and uncorrelated. The disturbance covariance matrix takes the following form

$$\mathrm{cov}(w_k) = \begin{bmatrix} \sigma^2_{\delta_{\mathrm{sw,dist}}} & 0 & 0 \\ 0 & \sigma^2_{\mathrm{F_{y,dist}}} & 0 \\ 0 & 0 & \sigma^2_{\mathrm{M_{z,dist}}} \end{bmatrix} \tag{2.27}$$

where σ denotes standard deviation and the subscripts denote the signal to which the standard deviation relates. Values for the variances of the random force and moment are derived from [111] and are $\sigma_{\mathrm{F_{y,dist}}}$ =730 N and $\sigma_{\mathrm{M_{z,dist}}}$ =360 Nm. Variance of the handwheel angle disturbance is derived from driving simulator experiments [108, 112] and is $\sigma_{\delta_{\mathrm{sw,dist}}}$ =0.1 rad.

The state-space equation (2.23) can be decomposed as the nominal disturbance-free dynamics:

$$\bar{\mathbf{x}}_{k+1} = \mathbf{A}_k\bar{\mathbf{x}}_k + \mathbf{B}_k\bar{\mathbf{u}}_k + \mathbf{F}_k \tag{2.28}$$

and the perturbation dynamics:

$$\Delta\mathbf{x}_{k+1} = \mathbf{A}_k\Delta\mathbf{x}_k + \mathbf{B}_k\Delta\mathbf{u}_k + \mathbf{H}_k\mathbf{w}_k \tag{2.29}$$

The human driver's compensatory control action is modelled as a Linear Quadratic Regulator (LQR) with full state feedback. Other control theories can be used to represent human control action, such as model predictive control and fuzzy control. An

LQR model has been shown to represent measured steering behaviour well [108, 112], and has been shown in some circumstances to be equivalent to model predictive control [56]. Future work will account for sensory and cognitive limitations [113, 114]. The objective of the compensatory controller is to provide a control action

$$\Delta \mathbf{u}_k = -\mathbf{K}_k \Delta \mathbf{x}_k \tag{2.30}$$

that minimises a quadratic cost function [115] comprising a weighted sum of mean square compensatory control action and mean square deviation from the nominal state:

$$J = \sum_{k=1}^{\infty} (\mathbf{x}_k^T \mathbf{Q} \mathbf{x}_k + \mathbf{u}_k^T \mathbf{R} \mathbf{u}_k) \tag{2.31}$$

where $\mathbf{Q}$ is the state cost matrix and $\mathbf{R}$ the input cost matrix. The two weighting matrices, $\mathbf{Q} = \text{diag}(q_i)$, where $i = 1:9$ and $\mathbf{R} = \text{diag}(r_l)$ where $l = 1:2$, are chosen to achieve an acceptable performance trade-off. To penalise steering action $\Delta\dot{\delta}_{sw}$ is weighted with $q_7 = 1$ $(\text{rad/s})^{-2}$ and $\Delta\delta_{sw}$ is weighted with $q_8 = 1$ rad^{-2}. Lateral path deviation Δe is weighted with $q_9 = 10$ m^{-2} and heading error $\Delta\psi$ with $q_3 = 1$ rad^{-2}. To discourage significant compensatory braking or acceleration control the weight on ΔT was set to $r_2 = 0.01$ $(\text{Nm})^{-2}$. All other states and controls were weighted with 10^{-6}. If measured driving response data are available it is possible to identify values of the cost function weights [108, 112].

2.5 Compensatory response calculation

Consider the covariance of (2.23):

$$\text{cov}(\mathbf{x}_{k+1}) = \text{cov}(\mathbf{A}_k \mathbf{x}_k + \mathbf{B}_k \mathbf{u}_k + \mathbf{H}_k \mathbf{w}_k + \mathbf{F}_k) \tag{2.32}$$

$\mathbf{x}_k$ can again be further split up into a nominal component $\bar{\mathbf{x}}_k$ with $\text{cov}(\bar{\mathbf{x}}_k)=0$, and a perturbation $\Delta\mathbf{x}_k$ about this with $\text{E}(\Delta\mathbf{x}_k)=0$. Similarly for $\mathbf{u}_k$ using (2.29) to give

$$\text{cov}(\mathbf{x}_{k+1}) = \text{cov}(\bar{\mathbf{x}}_{k+1} + \Delta\mathbf{x}_{k+1}) = \text{cov}(\mathbf{A}_k(\bar{\mathbf{x}}_k + \Delta\mathbf{x}_k) + \mathbf{B}_k(\bar{\mathbf{u}}_k + \Delta\mathbf{u}_k) + \mathbf{H}_k\mathbf{w}_k + \mathbf{F}_k) \tag{2.33}$$

The covariance of the nominal components is zero by definition. The variance of the constant is also zero, so (2.33) becomes

$$\text{cov}(\mathbf{x}_{k+1}) = \text{cov}(\Delta\mathbf{x}_{k+1}) = \text{cov}(\mathbf{A}_k\Delta\mathbf{x}_k + \mathbf{B}_k\Delta\mathbf{u}_k + \mathbf{H}_k\mathbf{w}_k) \tag{2.34}$$

The expression for $\Delta\mathbf{u}_k$ (2.30) can be substituted into (2.34). Noting that $\mathbf{w}_k$ is uncorrelated with the current vehicle state perturbation [116], the covariance equation becomes

$$\text{cov}(\mathbf{x}_{k+1}) = \text{cov}((\mathbf{A}_k - \mathbf{B}_k\mathbf{K}_k)\Delta\mathbf{x}_k) + \text{cov}(\mathbf{H}_k\mathbf{w}_k) \tag{2.35}$$

Exploiting the identity $\text{cov}(\mathbf{Ab}) = \mathbf{A}\text{cov}(\mathbf{b})\mathbf{A}^T$, and remembering that $\text{cov}(\mathbf{x}_k) = \text{cov}(\Delta\mathbf{x}_k)$, the covariance equation can finally be expressed as

$$\text{cov}(\mathbf{x}_{k+1}) = (\mathbf{A}_k - \mathbf{B}_k\mathbf{K}_k)\text{cov}(\mathbf{x}_k)(\mathbf{A}_k - \mathbf{B}_k\mathbf{K}_k)^T + \mathbf{H}_k\text{cov}(\mathbf{w}_k)\mathbf{H}_k^T \tag{2.36}$$

The variances predicted using (2.36) are compared with those determined numerically from an ensemble of 1000 time domain responses calculated using (2.29). Fig. 2.8 shows the time-varying standard deviations of the heading angle and lateral path deviations of the US vehicle calculated using the two methods. It is clear that the analytical and numerical results are consistent with each other.

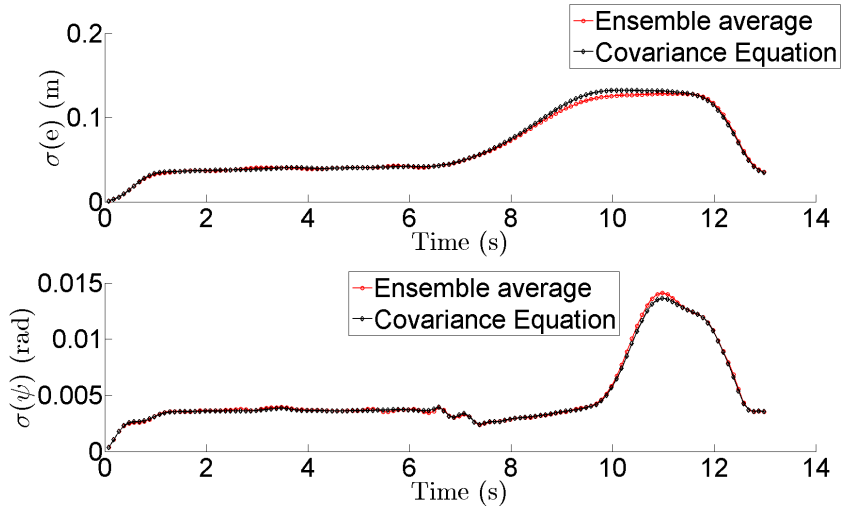

Figure 2.8: Standard deviation of the lateral path deviation and heading angle for the US vehicle, calculated using (2.36) and an ensemble of 1000 time domain responses.

In order to assess the controllability of the vehicle, the covariance calculation is extended to the control signals. The initial quantities of interest are the hand wheel rate and the drive/brake torque. The hand wheel rate is chosen over the hand wheel angle because it has been argued that the physical difficulty arises more from changing the angle than holding the angle constant [117]. This quantity can be calculated from (2.36) since $\dot{\delta}_{sw}$ is one of the vehicle states. However, this is not the case for the drive/brake torque. This requires computation of the variance of the vector $\mathbf{u}_k$. In

order to achieve this, $\mathbf{u}_k$ is split into nominal and perturbed components:

$$\text{cov}(\mathbf{u}_k) = \text{cov}(\bar{\mathbf{u}}_k + \Delta\mathbf{u}_k) = \text{cov}(\Delta\mathbf{u}_k) = \text{cov}(-\mathbf{K}_k\Delta\mathbf{x}_k) \tag{2.37}$$

Making again use of the identity $\text{cov}(\mathbf{Ab}) = \mathbf{A}\text{cov}(\mathbf{b})\mathbf{A}^T$, the following is obtained

$$\text{cov}(\Delta\mathbf{u}_k) = \mathbf{K}_k\text{cov}(\Delta\mathbf{x}_k)\mathbf{K}_k^T \tag{2.38}$$

Thus, once the state-covariance is computed at each time-step, the covariance matrix of the input vector can be found using (2.38). A comparison of the analytical and numerical results is shown in Fig. 2.9, which confirms the validity of the method.

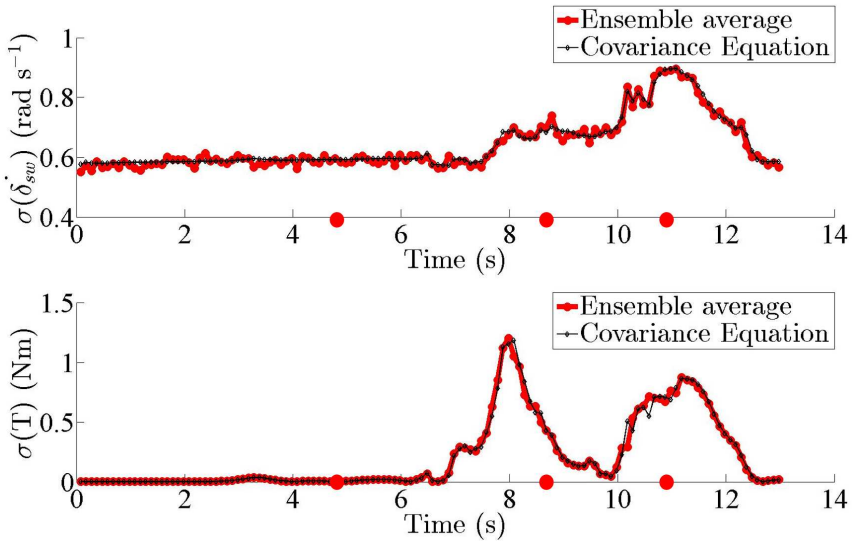

Figure 2.9: Standard deviation of the compensatory hand wheel rate calculated using (2.36) and drive/brake torque for the US vehicle calculated using (2.38) and an ensemble of 1000 time domain responses.

2.6 Results

Fig. 2.10 shows the standard deviation of the lateral path error as the two vehicles travel around the corner. On approach to the bend the error is about 0.04 m but the error increases as the vehicles pass the apex of the bend, reaching a maximum between the apex and exit of the bend. The two vehicles behave similarly, although the OS vehicle experiences a higher maximum value of path error standard deviation, about 0.15 m. Comparison to the nominal control actions in Fig. 2.4 and Fig. 2.5 shows

that the increase in path error begins some distance after the braking point, once the steering angle has reached about half its peak positive value. Once accelerating in a straightline after the exit of the bend the path error reduces to the about the value seen before braking on approach to the bend.

Fig. 2.11 shows the corresponding results for the heading error. The trends are similar, although heading error begins increasing later in the manoeuvre than lateral error. The OS vehicle exhibits two peaks, which occur either side of the switch from braking to accelerating.

Fig. 2.12 shows the standard deviation of the compensatory hand wheel angle. The increase in lateral and heading error is generally matched by increase in the compensatory hand wheel angle. The US vehicle has two peaks, located each side of the switch from braking to accelerating. The OS vehicle has a large peak just before the switch from braking to accelerating, and coincident with the first peak in heading error shown in Fig. 2.11

Further insight to the effect of the centre of mass position can be gained by plotting the standard deviations as a function of time and of CoM position, however in this format it is more difficult to relate feaures of the surface plot to locations on the bend. Fig. 2.13 shows the standard deviation of the compensatory steering control. The US and OS vehicles correspond to the edges of the surface, at $a/(a+b) = 0.4$ and $a/(a+b) = 0.6$.

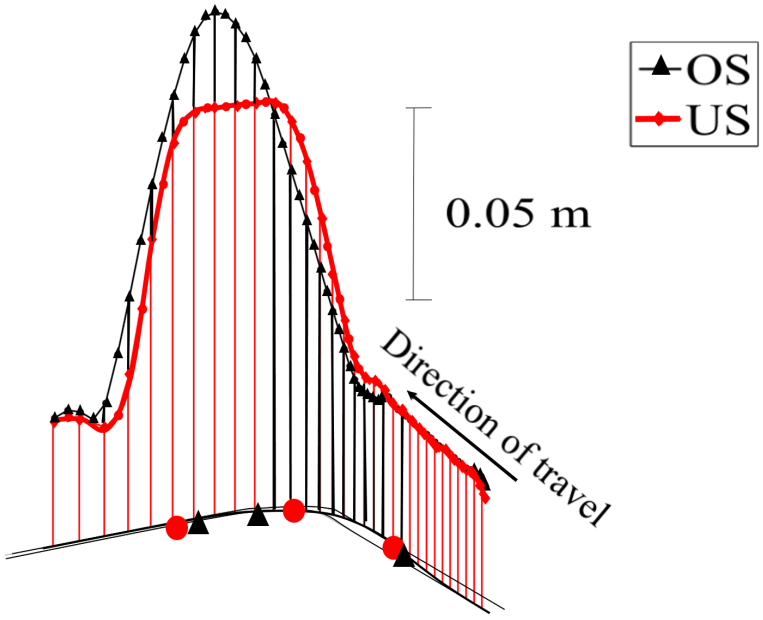

Figure 2.10: Standard deviation of lateral path error through the manoeuvre.

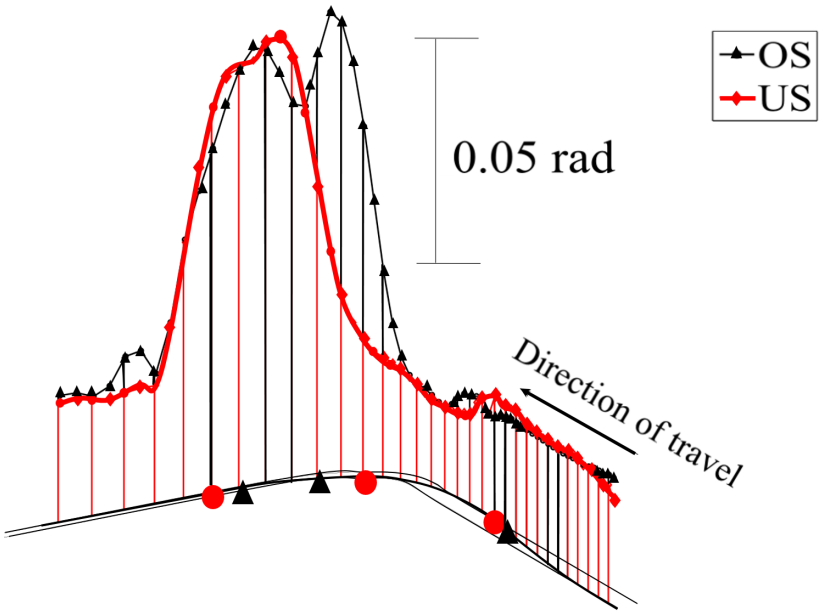

Figure 2.11: Standard deviation of heading error through the manoeuvre.

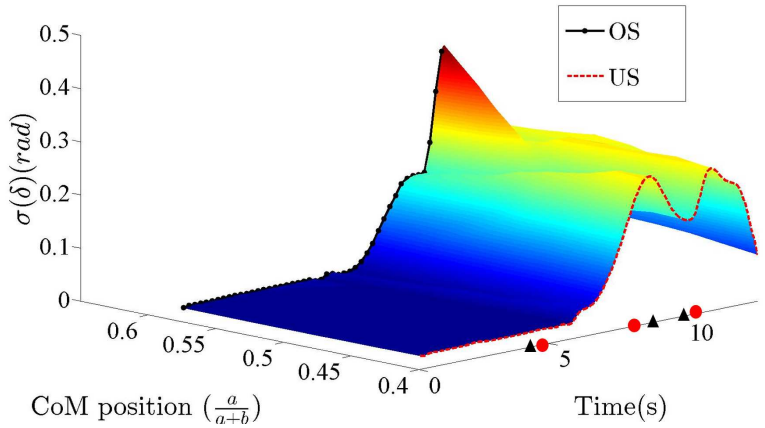

Figure 2.13: Surface plot showing how the standard deviation of the compensatory hand wheel angle varies with CoM position and time through the manoeuvre. The US and OS vehicles correspond to the boundaries of the surface.

2.7 Discussion

In this section the response variances of the compensatory steering controller are compared to some of the existing stability criteria reviewed in Section 1.1.3, specifically stability and controllability derivatives, and frozen-time eigenvalues. These existing

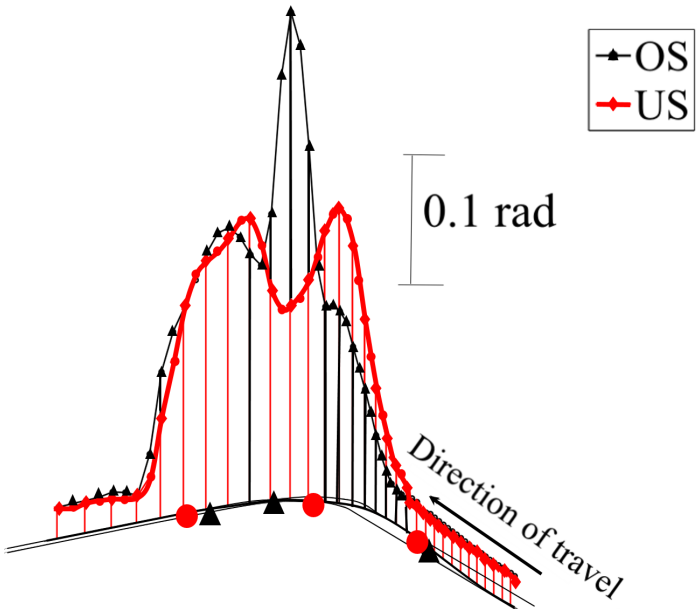

Figure 2.12: Standard deviation of compensatory hand wheel angle through the manoeuvre.

criteria are evaluated by considering only the linearised lateral-yaw dynamics of the vehicle (2.1a) and (2.1b) and linearising the combined slip tyre model at each point through the manoeuvre.

Fig. 2.14 shows the variation of the directional stability derivative $\partial M_t/\partial\beta$ as the US and OS vehicles travel through the curve. On approach to the curve the derivative is negative, which indicates a destabilising condition and arises because the traction force at the rear wheels decreases the effective cornering stiffness at the rear axle. After the braking point the derivative becomes positive, consistent with the stabilising effect of the front-biased brake balance. The subsequent changes of sign in the derivatives approximately follows the switches between drive and brake torque shown in Fig. 2.4, consistent with the combined slip behaviour of the tyres. Comparison with the responses of the compensatory dynamics in Figs. 2.10 to 2.12 reveals that the standard deviation of the compensatory control actions does not correlate well with the directional stability derivative: the control actions are large in regions of positive derivative and in regions of negative derivative.

Fig. 2.15 shows the variation of the yaw damping derivative $\partial M_t/\partial\dot{\psi}$ as the vehicles travel through the curve. The US and OS vehicles exhibit a similar variation. The derivative is negative throughout the manoeuvre, corresponding to a stabilising action. However the magnitude of the derivative reduces significantly in the middle of the manoeuvre, corresponding approximately to the peaks in the compensatory responses

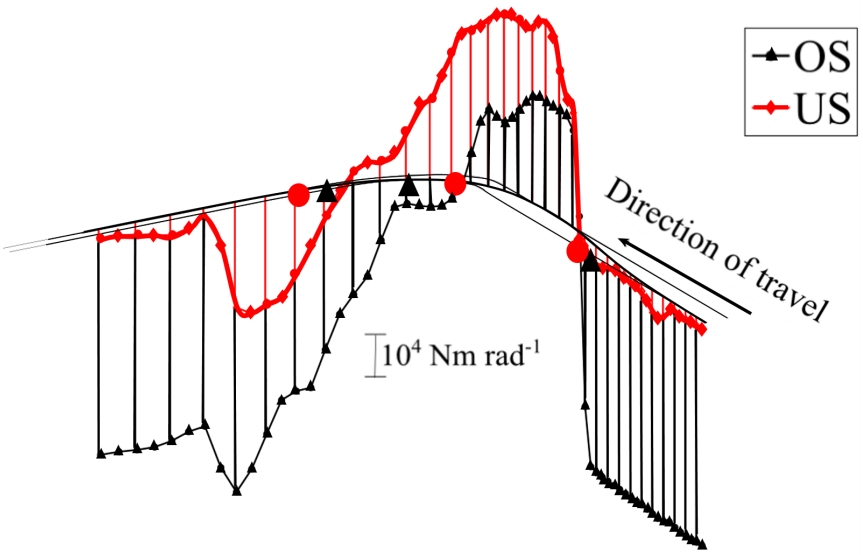

Figure 2.14: Directional stability derivative for the OS and US vehicles through the manoeuvre.

in Figs. 2.10 to 2.12. The reduction is due to a decrease in the cornering stiffness of the tyres as they approach saturation.

Fig. 2.16 shows the variation of the control moment derivative $\partial M_t/\partial \delta_{sw}$ as the vehicles travel through the curve. The derivative is high whilst the vehicle is accelerating on approach to the curve because there are no drive or braking torques on the front axle to reduce the cornering stiffness. The derivative reduces significantly once the brakes are applied, and reduces further as hand wheel angle is applied and the front tyre nears saturation. The US vehicle reaches a significantly lower value than the OS vehicle because of the forward CoM position of the US vehicle. The trough in the derivative corresponds approximately to the peak region of the compensatory hand wheel angle in Fig. 2.12.

Fig. 2.17 shows the real and imaginary parts of the eigenvalues for the lateral-yaw dynamics of the US vehicle through the manoeuvre. The circles plotted on the time axis correspond to the three phases of the manoeuvre: braking on entry; transition from braking to accelerating at mid corner; and maximum drive torque at exit. Positive real parts indicate instability; non-zero imaginary parts indicate oscillation. The vehicle is stable throughout the manoeuvre apart from the point where drive torque is applied at mid-corner until just after the point where maximum drive torque has been applied at exit. Fig. 2.18 shows the eigenvalues for the OS vehicle. In comparison to the US vehicle, it becomes unstable earlier in the manoeuvre (mid-way through the braking

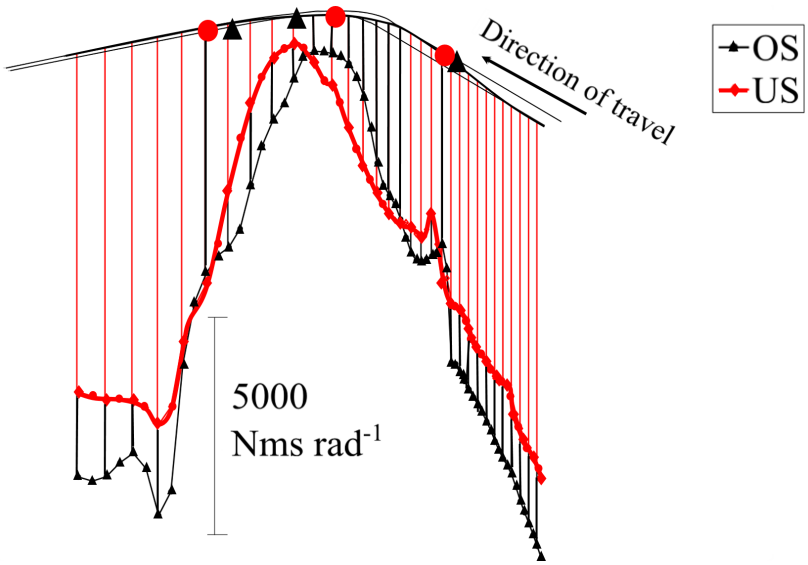

Figure 2.15: Yaw damping derivative for the OS and US vehicles through the manoeuvre

phase) but returns to stability at about the same time. Comparing this data to the response variances of the compensatory control, Figs. 2.10 to 2.12, does not show an obvious relationship to the features of the variance data, apart from a general increase in variance in the regions of instability.

It is clear from the comparisons made in this section that the response variances of the driver and vehicle compensatory dynamics give an additional view of the behaviour of a nonlinear vehicle as it travels through a manoeuvre near to the limit of adhesion. Further work is needed to understand precisely how the responses should be interpreted, and to investigate the effect of the neuromuscular properties (damping and natural frequency). A series of instrumented vehicle or driving simulator experiments with experienced drivers providing subjective assessments is an obvious next step. It is anticipated that such experimental data will allow potential benefits of the new criteria to be revealed. Other extensions to the work are planned: addition of human sensory and cognitive limitations to the compensatory control model; a nonlinear instead of linear compensatory controller; and incorporation of the compensatory response variances as constraints in the calculation of the nominal optimal controls.

2.8 Summary

The brief review of existing criteria in Section 1.1.3 for quantifying stability and controllability of road vehicles concluded that the *practical stability* approach was

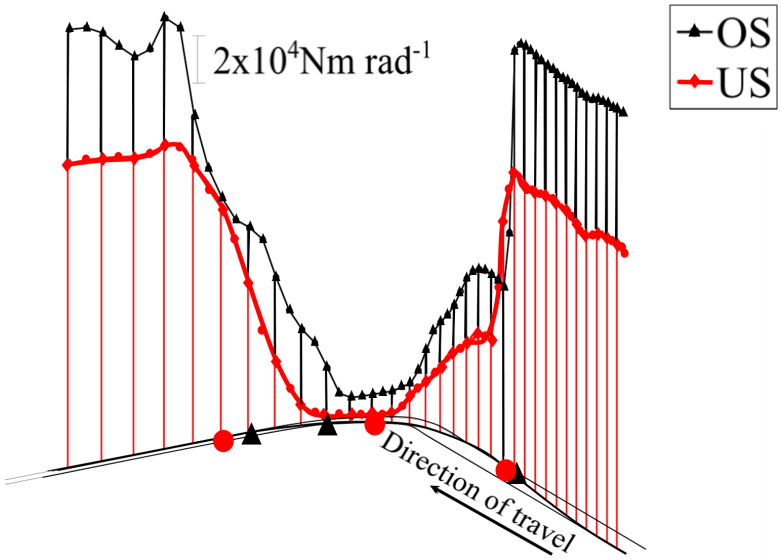

Figure 2.16: Control moment derivative for the US and OS vehicles through the manoeuvre

worthy of further development, with the aim of providing a closer relationship to subjective assessments by drivers. A minimum manoeuvre time calculation for a five DoF nonlinear vehicle model reveals that moving the CoM rearwards (without changing

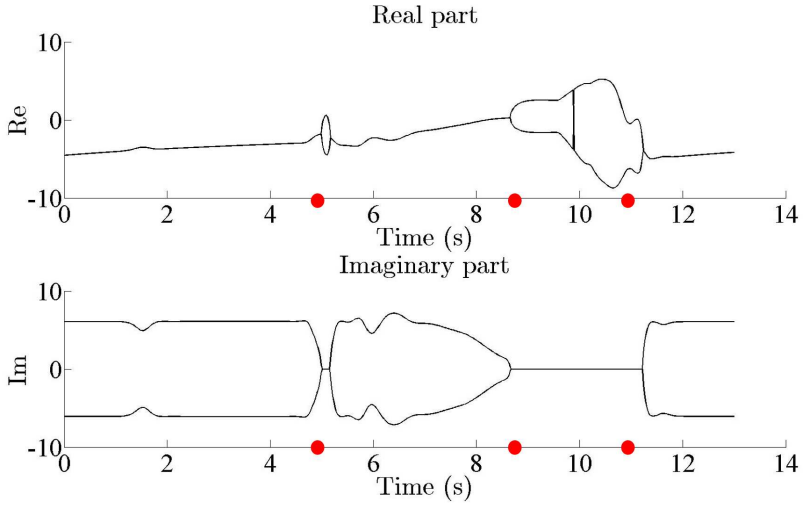

Figure 2.17: Eigenvalues for the US vehicle going through the manoeuvre described in Section 2.3. The circles plotted on the time axis correspond to the three phases of the manouvre: braking on entry; transition from braking to accelerating at mid corner; and maximum drive torque at exit. Positive real parts indicate instability; non-zero imaginary parts indicate oscillation.

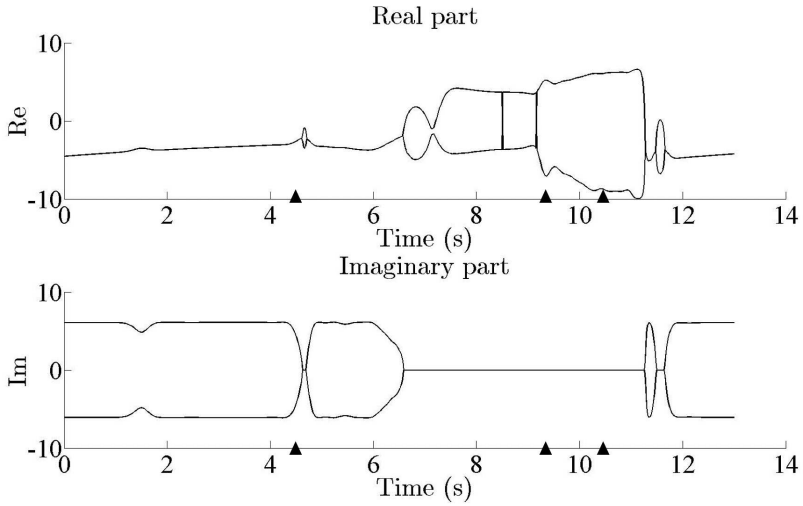

Figure 2.18: Eigenvalues for the OS vehicle going through the manoeuvre described in Section 2.3.

the brake balance) requires earlier braking on entry to the corner, countersteering between mid-corner and exit, and later acceleration out the corner. Calculation of the response variances of a compensatory LQR steering controller agrees closely with the variances calculated from an ensemble of time-domain responses, and thus enables efficient computation of the driver and vehicle response to random disturbances acting on the vehicle as it travels through the manoeuvre. The variance of lateral path error, heading error and compensatory hand wheel angle plotted as a function of distance through the corner reveal how the driver might perceive the change in vehicle operating point. The main effect of moving the centre of mass rearwards is a sharp increase in the variance of the hand wheel angle at the transition from braking to accelerating at mid-corner. Comparison of the variances to existing criteria (eigenvalues and stability derivatives) demonstrates that the new criteria present an additional perspective of the dynamic behaviour of a nonlinear vehicle through the manoeuvre, particularly concerning the driver's closed loop action. Metrics based on linear models, such as eigenvalues or phase portraits, have somewhat limited practical applications due to model simplifications. Stability and controllability metrics can be very insightful when designing a vehicle but do not provide any information about driver-vehicle interaction. The novel formulation can help predict a vehicle subjective handling qualities.

Different metrics serve for different purposes. Eigenvalues and state space are a useful tool for preliminary theoretical analysis; stability metrics can provide useful insights in the design phase while variance based metrics can be an aid to understand a vehicle's handling qualities subjective assessment.

Chapter 3

A NMPC compensatory controller formulation

3.1 Introduction

The results in [2] show that while the LQR controller performs well for a disturbed linear vehicle, it fails to control a disturbed nonlinear vehicle at the limit of adhesion. Racing drivers exploit the whole of the tyre friction circle to maximise tyre forces and minimise lap time. Non professional drivers are also likely to saturate tyres for extreme manoeuvres; when avoiding a crash for example. Being able to model the totality of the tyre curve and controlling the vehicle when operating in the negative slope region is therefore relevant to improving understanding of driving behaviour.

The problem of finding the minimum manoeuvre time for a vehicle controlled by an *ideal* driver has been widely studied in the Optimal Control (OC) literature. The main focus has been simulating reasonably detailed vehicle models over complex manoeuvres - such as a full lap on an international racetrack – in an efficient fashion. However, random factors – such as road roughness, wind speed or the noise in the driver neuromuscular system – significantly affect driving [111]. The objective of the proposed controller is to extend previous formulations to incorporate randomness and thus represent a human driver's ability to stabilise and control a nonlinear vehicle.

Nonlinear optimisation schemes are ideal for tackling such problems as they are characterised by a high degree of robustness, though this comes at the expense of a higher computational burden. The vast literature on control of nonlinear systems provides a starting point for the analysis. Some fundamental OC concepts are outlined to define the theoretical framework used to design the controller. Particular attention

is given to the development of numerical methods for digital computers and how integration schemes and optimisation algorithms are implemented.

While the subject is an expansive one, this Chapter discusses only concepts that are directly related to the implementation of the controller, without reviewing theories which are not directly related to the development of the controller. The interested reader can refer to [118] and [119] for more comprehensive theoretical explanations. Cognitive limitations and driving strategies are accounted for so as to have a better understanding of a real human driver's approach to the driving task. External random disturbances are added to the system to account for any non deterministic events present in the driving task. Due to the random nature of the system, a robust control scheme is the most obvious approach.

Section 3.2 focuses on the theoretical background necessary to formulate the problem, concluding with a brief overview on the algorithms which will be used to solve it. Section 3.3 details the derivation for the lateral controller while Section 3.4 derives the longitudinal controller. All the material covered in Section 3.2 is established, being reported only for completeness, and so does not constitute any new contribution to the subject. Section 3.3 and Section 3.4, however, introduce a novel formulation which, even though it rests on ideas from the literature, is an original contribution to knowledge.

3.2 Optimal control

Optimal control is a mathematical optimisation method aimed at deriving control actions that drive systems to a desired state [118]. The core of the process usually consists of finding the minima of a cost function subject to equality and inequality constraints. The complexity of the solution depends on the problem requirements. While low dimensional, linear, unconstrained problems can be solved using simple calculus techniques [119], more refined techniques – such as Langrangian multipliers – are required to tackle problems of practical interest. The most general frameworks that provide optimality conditions are the Pontryagin minimum and Bellman optimality principles [120]. Though the two formulations are mathematically different and have different conditions of existence for the solutions, they lead to essentially the same result for problems formulated in equivalent ways [121].

While problems are usually formulated in continuous time, imposing optimality conditions inevitably leads to equations that are too complex to be solved analytically. It is therefore necessary to discretise the problem and apply numerical methods to

solve it. Discretisation can happen either before or after applying optimality conditions. The former strategy is referred to as a *direct* approach while the latter is usually called an *indirect* approach [122]. The type of approach is probably the most critical decision in the solution strategy for an optimal control problem. Both approaches are viable and lead to the same solutions. However, one can be better suited than the other for a specific case. Numerical methods are always necessary to solve the resulting optimal problem. Parameter optimisation schemes are usually preferred over variational calculus based methods because digital computers made the implementation of parameter optimisation schemes efficient and reliable. Section 3.2.1 and Section 3.2.2 provide an overview of indirect and direct methods respectively, while Section 3.2.3 discusses the algorithms that can be used to solve the resulting optimal control problem.

Optimal control strategies are limited to deterministic systems. Since random disturbances are applied to the vehicle, the theory needs to be extended to include notions of robust control. Section 3.2.4 outlines the theoretical approach used to embed robustness into the system.

3.2.1 Indirect methods

Indirect methods impose optimality conditions on continuous problems. While calculus can be used to derive the optimality conditions for a set of time invariant parameters, calculus of variations serves the same purpose for a formulation in continuous time. Optimality conditions for unconstrained problems are obtained by finding a functional that minimises a cost function [123]. The theory behind this approach is established and will not be covered here. The interested reader can refer to [124] for a detailed explanation. Extending the fundamental variation principles to constrained problems leads to Pontryagin's minimum principle which states that, given a set of admissible controls, the control Hamiltonian must take an extreme value [124]. The control Hamiltonian is an expression that provides the necessary conditions for optimality. The derivation is again omitted as it is established. The interested reader can refer to [123] for a comprehensive overview of variational calculus and optimisation methods. Pontryagin's minimum principle generalises the condition of optimality to a two-point boundary-value problem, which is analytically intractable for most practical applications [119]. The boundary value problem is then solved using well established numerical methods, such as root-finding or Runge Kutta schemes [125].

3.2.2 Direct Methods

Direct methods discretise, or *transcribe*, the problem before applying optimality conditions. A well chosen transcription method is key to an efficient solution. Transcription techniques aim at describing a time varying problem using constant parameter vectors [121]. It is clear that difficulties arise when a dynamic system is also described by differential relationships. In general, the first step for a transcription process is to define a set of parameters that approximates the continuous functions describing the model and to then impose constraints that replicate the continuous time differential relationships. In practice, there are three different families of discretisation methods used to carry out the aforementioned steps [122].

- *Single shooting*: the state trajectories are found by placing the nodes on the control trajectory and solving the resulting state differential equations using appropriate numerical quadrature techniques. Solutions are found iteratively until convergence is achieved. The limitations of this method are the lack of robustness and the difficulty of handling highly nonlinear constraints.

- The lack of robustness can be overcome by using a *multiple shooting* method, which divides the trajectory into different sections and imposes boundary conditions to ensure continuity. The solution methods are then the same as single shooting techniques.

- The limit case of the multiple shooting technique consists of substeps of the control trajectory that can be solved using a one-step rule, referred to as *collocation methods*. In this method a quadrature rule is implemented at each node to formulate a parameter optimisation problem to find an array of parameters that characterise the system at every node.

The advantage of collocation methods is the decoupling of control and states, which adds robustness. Once a problem is transcribed, optimality conditions can be imposed and the problem solved, which is essentially equivalent to performing parameter optimisation on a large data set. An overview of parameter optimisation is given in Section 3.2.3, together with some of the numerical algorithms used to tackle the problem.

3.2.3 Parameter optimisation algorithms

Section 3.2.2 shows that optimality conditions can be imposed after discretisation. Since a discrete problem is described by time-invariant parameters, it can effectively

be treated as a parameter optimisation problem. Any time varying dynamic system can be described by a sufficiently large number of parameters using the appropriate discretisation techniques. The size of the resulting problem is usually significant so choosing the right algorithm is key to obtaining a fast and efficient solution. An appreciation of the numerical techniques used to solve such problems is essential to illustrate the operation of the robust controllers described in Section 3.3 and Section 3.4. The numerical techniques described in this Section can be found in any optimisation book, for instance in [119]. They are reported here for completeness.

The fundamental parameter optimisation problem consists of finding a vector $\boldsymbol{b}$ that minimises an objective function $J(\cdot)$. The problem is usually expressed as

$$\min_{b} J(\boldsymbol{b}) \tag{3.1}$$

The standard calculus result $\frac{\partial}{\partial \boldsymbol{b}} J(\boldsymbol{b}^*) = 0$ can be used to find stationary points and higher order derivatives to characterise them. $\boldsymbol{b}^*$ indicates a local minimum. In practice, unconstrained optimisation is quite rare. Most problems involve some sort of constraints; equality constraints being the most common. For a dynamic system evolving over a timespan, referred to as prediction horizon, the general form of an optimisation problem with equality constraints takes the following form

$$\begin{aligned} &\min_{b} J(\boldsymbol{b}) \\ \text{subject to} \quad &\mathrm{f_k}(\boldsymbol{b}) = 0 \qquad k = 1.....N_p \end{aligned} \tag{3.2}$$

where f_k is a nonlinear constraint function and N_p the number of time steps in the prediction horizon, also referred to as prediction horizon. The Lagrange Multiplier theorem, which can be derived using results from linear algebra, defines the conditions for optimality for a minimisation problem with an equality constraint. A new quantity, referred to as the Lagrangian, is introduced as a means of quantifying the sensitivity of the cost function around the constraints. It can be expressed as

$$L(\boldsymbol{b}, \boldsymbol{\lambda}) = J(\boldsymbol{b}) + \sum_{k=0}^{N_p} \lambda_k \mathrm{f_k}(\boldsymbol{b}) \tag{3.3}$$

where L is the Lagrangian and λ_k the Lagrangian multiplier at k. The condition for optimality can be found by minimising the Lagrangian with respect to the vector of

Lagrangian multipliers $\boldsymbol{\lambda}$ and the parameter vector $\boldsymbol{b}$:

$$\begin{aligned}\nabla_{\boldsymbol{b}} L(\boldsymbol{b}^*, \boldsymbol{\lambda}^*) &= 0 \\ \nabla_{\boldsymbol{\lambda}} L(\boldsymbol{b}^*, \boldsymbol{\lambda}^*) &= 0 \\ \boldsymbol{q}^T \nabla_{\boldsymbol{b}\boldsymbol{b}} L(\boldsymbol{b}^*) \boldsymbol{q} &\geq 0\end{aligned}$$

where $\boldsymbol{q}$ is any vector that belongs to the subspace of first-order feasible cost function variations. The interested reader can refer to [126] for the proof. The Lagrangian multipliers are not part of the solution as they are newly introduced variables but they are useful to understand the nature of the solution and of the constraints. Another class of constraints to be considered in optimisation is inequality constraints. They arise often in practical problems and impose bounds on the search domain. An optimisation problem including inequality constraints can be cast as

$$\begin{aligned} & \min_{b} \; J(\boldsymbol{b}) & \\ \text{subject to} \quad & \mathrm{f_k}(\boldsymbol{b}) = 0 \qquad & k = 1.....N_p \\ & \mathrm{h_k}(\boldsymbol{b}) \leq 0 \qquad & k = 1.....N_p \end{aligned}$$

where $\mathrm{h_k}(\boldsymbol{b})$ is a nonlinear constraint function. It is assumed that both equality and inequality constraints are applied over the whole prediction horizon. In this case, first order conditions for optimality can be found by augmenting the Lagrangian with another vector of multipliers and finding the optimal conditions based on that system. However, problems with a high number of constraints can make the Lagrangian prohibitively complex. Alternative solution strategies, which are not always analytic, have been developed to tackle this problem. *Penalty* methods, for example, use a series of unconstrained sub-problems to replace the original constrained problem. Constraints are enforced by adding an extra term – positive when the states violate the constraints and zero otherwise – to the cost function of each sub-problem. An example of such function is known as the *quadratic penalty function*, defined as

$$O(\boldsymbol{b}, p_n) = J(\boldsymbol{b}) + \frac{1}{2p_n} \sum_{k=0}^{N_p} \mathrm{f_k^2}(\boldsymbol{b}) \tag{3.4}$$

where p_n is a penalty parameter. The problem can be solved by iterating over p_n until a solution is reached. This method, however, suffers from ill conditioning which can be tackled by including the term $\sum_{k=0}^{N_p} p(\lambda_k)$ – where $p(\cdot)$ and λ_k are an estimate of the Lagrangian multiplier – to the cost function. This method also iterates over the

estimate of the Lagrangian multiplier and the barrier parameter until convergence is achieved. A more conservative and robust approach is the interior point barrier method, where a logarithmic penalty is added to the original cost function. The augmented cost function when only equality constraints are enforced is

$$O(\boldsymbol{b}, p_n) = J(\boldsymbol{b}) - p_n \sum_{k=0}^{N_p} \log \mathrm{f_k}(\boldsymbol{b}) \tag{3.5}$$

This approach effectively softens the constraints as it allows the solution to lie in their vicinity. Iteration is again necessary to find the parameters that satisfy (3.5). Another class of solvers is referred to as sequential quadratic programming methods. These methods exploit the properties of the Lagrangian to iteratively find the global search direction by solving a number of optimisation sub-problems. At each step a minimisation problem is solved to refine the search direction $\boldsymbol{d}_i$ until convergence is achieved. In the case of equality constraints only, the search direction $\boldsymbol{d}_i$ can be obtained by

$$\begin{aligned} \min_{d} \ & \frac{1}{2}\boldsymbol{d}_i^T \nabla^2_{bb} L(\boldsymbol{b}_i, \boldsymbol{\lambda}_i)\boldsymbol{d}_i + \nabla J(\boldsymbol{b}_i)^T \boldsymbol{d}_i \\ \text{subject to} \ & \mathrm{f_k}(\boldsymbol{b}_\mathrm{i}) + \nabla \mathrm{f_k}(\boldsymbol{b}_\mathrm{i})^\mathrm{T} \boldsymbol{d}_\mathrm{i} = 0 \end{aligned}$$

where $L(\boldsymbol{b}_i, \boldsymbol{\lambda}_i)$ is the Lagrangian at iterate $\boldsymbol{b}_i, \boldsymbol{\lambda}_i$. The value $\boldsymbol{d}_i$ is added to the states to update the search direction at every step of the iteration until convergence is achieved:

$$\boldsymbol{b}_{i+1} = \boldsymbol{b}_i + \boldsymbol{d}_i \tag{3.6}$$

3.2.4 Robust Control

The optimal control theory outlined so far allows the formulation of an optimal control problem for an arbitrarily complex nonlinear vehicle model without constraints on the geometry of the manoeuvre. These problems have been studied in great depth in the vehicle dynamics literature. However, the optimisation routines so far outlined are not enough to deal with uncertain systems, which is the aim of this work. Notions of robust control are therefore necessary to lay the theoretical foundations of a controller that can reject random disturbances. Robust control is the branch of control theory that deals with uncertain systems. For the purpose of this work a system that has additive, external disturbances is considered. Even though a number of approaches have been proposed to tackle this problem, there is no agreement on which is the most

suitable. A disturbed control problem in its most general form can be expressed as

$$\boldsymbol{x}_{(k+1)} = g_{nl}(\boldsymbol{x}_{(k)}, \boldsymbol{u}_{(k)}) + \boldsymbol{w} \tag{3.7}$$

where k indicates the current time step, $\boldsymbol{x}$ is the nominal state vector, $\boldsymbol{u}$ the control action, $g_{nl}(\cdot)$ a nonlinear function and $\boldsymbol{w}$ the disturbance vector. The two most common approaches to tackle this problem are known as the *min-max* approach and the *feedback* min-max MPC. The *min-max* approach uses a single open loop control series to minimise the worst case scenario while for the *feedback* min-max MPC, the worst case cost is minimised over a sequence of control laws. Numerous authors, for instance Rakovic [100], have shown that the open loop approach suffers from being overly conservative. Some *feedback* schemes tackle this problem at the expense of computational resources. Any attempt to improve process efficiency limited the practical use of the algorithms proposed to simple cases. Mayne tried to obviate this problem by sacrificing optimality for simplicity. In a work with Kerrigan [99], building on a series of papers by Rakovic [100, 101] et al., they developed *tube based* MPC. The fundamental idea is to use two controllers, where one controller – referred to as *nominal controller* – finds the optimal nominal trajectory in absence of disturbances and the other controller – referred to as *compensatory controller* – bounds the disturbed states in the vicinity of the nominal trajectory, effectively creating a *tube* whose width changes depend on the disturbances and the structure of the controller. Feasibility is guaranteed if the tube is chosen in a way such that its outer limits stay within the original state constraints. The constraints of the nominal controller may have to be tightened to account for the disturbances; this is usually done iteratively. In a more recent work, Mayne et al. [102] extended the idea of tube based MPC to nonlinear systems, showing that linear superposition is not required to obtain robust solutions.

The theory behind this approach is illustrated as it is at the core of the design of the controllers in Section 3.3 and Section 3.4. Firstly, the nominal controls are found. Considering the system in Eq 3.7, the centre of the tube can also be found using a standard MPC controller

$$\boldsymbol{x}^*_{nom(k+1)} = g_{nl}(\boldsymbol{x}^*_{nom(k)}, \boldsymbol{u}^*_{nom(k)}) \tag{3.8}$$

The optimal nominal state, $\boldsymbol{x}^*_{nom}$ and control $\boldsymbol{u}^*_{nom}$ can be also be found in different fashions, for instance by assuming constant controls and states or applying different control strategies.

A second controller, referred to as ancillary model predictive controller, is added to the system to keep the state of the system close to the reference trajectory in presence of disturbances. The aim of the controller is to minimise the deviation of the actual trajectory $\boldsymbol{x}_{tot}$ from the optimal trajectory $\boldsymbol{x}^*_{nom}$ previously calculated. The cost function to be minimised can therefore be expressed as

$$J = \sum_{k=0}^{N_p-1} l(\boldsymbol{x}^*_{nom(k)} - \boldsymbol{x}_{tot(k)}, \boldsymbol{u}^*_{nom(k)} - \boldsymbol{u}_{tot(k)}) \tag{3.9}$$

where $l(\cdot)$ is a function that can take any form. The optimal problem yields $\boldsymbol{u}^*_{tot(k+1)} = \min_J$ which is the signal that is applied to the plant. It contains the nominal and disturbed signals. This is a significant difference from the approach used in Chapter 2 as linear superposition is no longer needed. The increase in computational cost is justified by the improved robustness. Mayne and Kerrigan prove in [102] that the uncertainty set $\mathbf{S}_K$ that contains the disturbed trajectories is bounded only if the first instance of the input sequence $\boldsymbol{u}^*_{tot(k)}$ is applied in a classical MPC fashion. The tube is not *necessarily* bounded if more than one instance is applied. This is an important point to investigate as prolonged control actions from a single optimisation are an important aspect of the driver model proposed.

In order to implement tube based MPC on a vehicle, the current states and controls of the vehicle $\boldsymbol{x}_{tot(k)}$ are measured first. For the purpose of this work, it is assumed that the driver has perfect knowledge of the states of the vehicle so no predictive filters are necessary. The optimal undisturbed control sequences $\boldsymbol{x}^*_{nom(k)}$ and $\boldsymbol{u}^*_{nom(k)}$ are then determined. These include the nominal racing line, all the vehicle states and the lateral and longitudinal inputs necessary to stabilise the vehicle. The ancillary problem is then solved to obtain $\boldsymbol{u}^*_{tot(k+1)}$, which is then fed to the disturbed system to integrate it over time to obtain a measurement for the following set of states $\boldsymbol{x}_{tot(k+1)}$.

3.3 Lateral controller

Section 3.2 provides the theoretical background for the formulation of the driver model. A lateral controller is considered first. It is assumed that the vehicle travels at constant speed and that the driver can only control the hand wheel. Since the speed is constant, accelerating and braking are not considered. This assumption simplifies the vehicle model but more importantly the tyre model, as only lateral forces have to be considered. The following sections set out the problem in its general form and then move onto

describing all its components; namely a track model, a vehicle model, and a tyre model. Constraints, cognitive limitations and the solution methodology are also considered. The concepts outlined to develop the lateral controller will be extended to include the longitudinal control in Section 3.4.

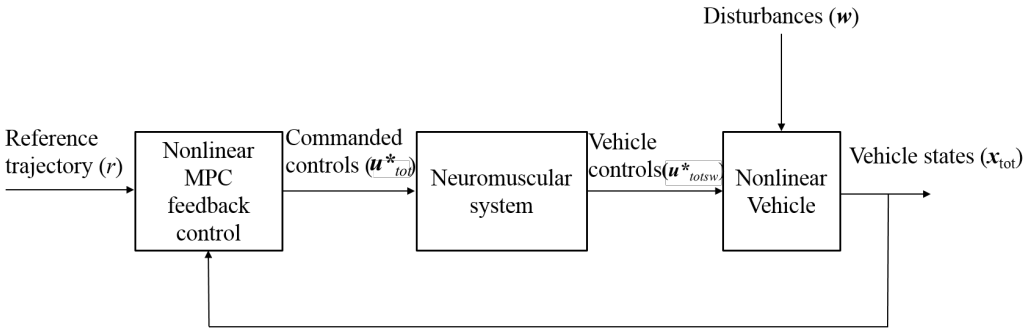

Figure 3.1: Robust controller structure highlighting the main components. The nonlinear formulation reduces significantly the number of blocks. The subscript $_{sw}$ indicates steering wheel – i.e. applied to the vehicle.

3.3.1 Problem formulation

Figure 3.1 shows the main blocks of the proposed control approach described in Section 3.2.4. The reference trajectory, which is assumed to be known, is fed to the nonlinear MPC controller together with the vehicle states. The controller finds the optimal control input $\boldsymbol{u}^*_{tot}$ which, after being passed through the neuromuscular filter, is input to the disturbed nonlinear vehicle. The resulting total states are fed back to the MPC controller. For the purpose of this work, the nominal state trajectory is assumed to be known. A sufficiently general formulation for the problem to be solved to find $\boldsymbol{u}^*_{tot}$ is

$$\begin{aligned} \min_{\boldsymbol{x}_{tot},\boldsymbol{u}_{tot}} \quad & J(\boldsymbol{x}_{tot},\boldsymbol{u}_{tot},\boldsymbol{x}^*_{nom},\boldsymbol{u}^*_{nom}) \\ \text{subject to} \quad & \mathrm{f}_k(\boldsymbol{x}_{tot},\boldsymbol{u}_{tot},\dot{\boldsymbol{x}}_{tot}) = 0 \qquad k = 1...N_p \\ & \mathrm{h}_k(\boldsymbol{x}_{tot},\boldsymbol{u}_{tot}) \geq 0 \qquad k = 1...N_p \end{aligned} \tag{3.10}$$

In traditional MPC only the first instance of the control sequence $\boldsymbol{u}$ is fed to the disturbed plant. The formulation allows for a longer section of the control sequence to be fed to the plant to account for the driver's cognitive limitations. Section 3.3.2 describes the track model and Section 3.3.3 the vehicle and tyre model, completing the description of all terms in Eq. 3.10. Section 3.3.4 elaborates on how driver cognitive limitations are incorporated in the controller.

3.3.2 Track description

The first aspect of the solution to be considered is the track description as it determines the frame of reference, which in turn affects the formulation. The Lateral Path Error (LPE), defined as the perpendicular distance between the centre of mass of the vehicle and the reference trajectory, is the most important quantity to consider when defining the track, as it needs to be minimised to ensure the vehicle stays within a tube. A Cartesian frame of reference would allow easy expression of the coordinates of the nominal trajectory with X_{ref} and Y_{ref} and the vehicle position using X and Y. Expressing the LPE in cartesian coordinates would not be straightforward. The problem can be tackled as one of finding the minimum distance between a point and a line, which has a well known solution. The proximity strategy proposed by Timings in [1] could be implemented. Figure 3.2 shows the main quantities involved in the derivation. The true position of the vehicle X,Y, which is available from the previous prediction, and its approximate position on track X_{app}, Y_{app}, assuming the distance travelled is Tds, are compared. An arbitrary number p_t of track points in front and behind the approximate position on track X_{app}, Y_{app} of the vehicle is considered. The distance between the vehicle true position and these points on the track is calculated and trigonometry is then used to determine the track data point whose normal goes through the vehicle position, which is the true minimum distance.

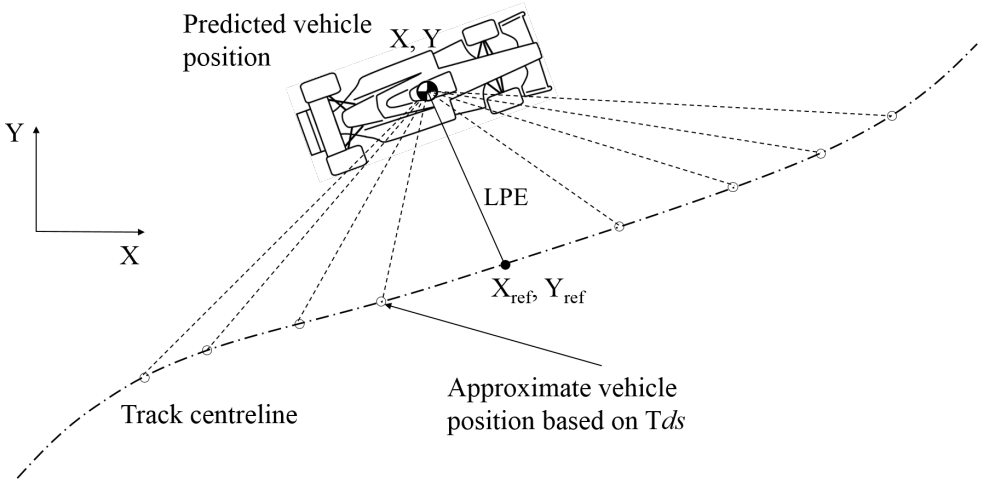

Figure 3.2: Strategy to evaluate the points on the nominal line such that they are perpendicular to the track and pass through the vehicle centre of mass

This approach is problematic as it involves expressing the LPE with trigonometric expressions, which eventually appear in the cost function, increasing significantly the numerical complexity of the problem. A more efficient track formulation, where an

expression for the LPE would be readily available and could be expressed in a simple form, would greatly improve the tractability of the problem.

An intrinsic track formulation allows us to express the LPE as a state of the system, making it readily available in the cost function. The track model described here is the same used by Lot [127] and Perantoni [128]. Curvilinear coordinates are used to express the vehicle position. Referring to Figure 3.3, the vehicle position is described in terms of distance travelled s and lateral offset $n(s)$. The nominal line is described by the

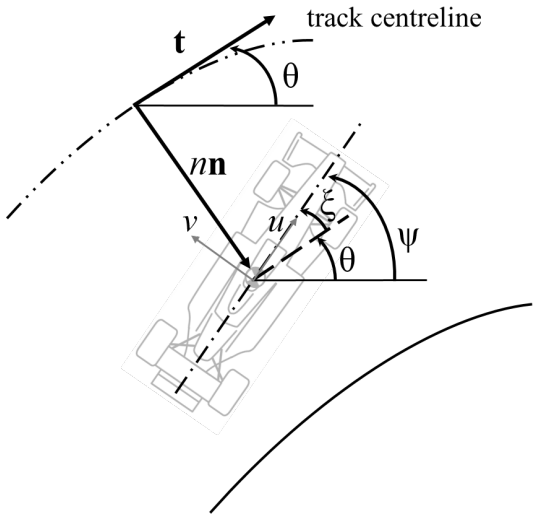

Figure 3.3: Diagram showing the intrinsic curvilinear track description.

distance travelled s and normal $\mathbf{n}$ and tangent $\mathbf{t}$ unit vectors. The heading direction of the vehicle is given by ψ and the angle between the vehicle and the reference is given by ξ, which implies that $\psi = \theta + \xi$, where θ is the angle between $\mathbf{t}$ and the global axis. The key advantage of this track formulation is that the position of the vehicle and its orientation are included among the states of the vehicle without resorting to numerically complex functions. A relationship among the different quantities mentioned before needs to be found to describe the vehicle position in time. The track curvature Ω can be found exploiting simple geometry concepts:

$$\Omega = \frac{d}{ds}\left(\arctan\frac{dy}{dx}\right) \tag{3.11}$$

where x and y are the intrinsic local Cartesian coordinates and s is the total distance travelled. Since only differential quantities are involved, a global rectangular frame of reference does not have to be defined. Eq. 3.11 provides a convenient way to relate rectangular and curvilinear coordinates. Resolving velocities in the $\mathbf{t}$ direction provides

a relationship between the distance travelled and the orientation of the vehicle.

$$\dot{s} - n\dot{\theta} = u\cos\xi - v\sin\xi \tag{3.12}$$

where n is the lateral path error, θ the absolute heading angle, u the forward vehicle speed, v the lateral vehicle speed and ξ the angle between the vehicle longitudinal axis and the reference trajectory. Differentiating θ and rearranging gives

$$\dot{s} = \frac{u\cos\xi - v\sin\xi}{1 - n\Omega} \tag{3.13}$$

The Lateral Path Error and the heading angle, calculated with respect to the track, are also included in the formulation

$$\dot{n} = u\sin\xi + v\cos\xi \tag{3.14}$$

$$\dot{\xi} = \dot{\psi} - \Omega\dot{s} \tag{3.15}$$

The track description shows that most quantities are better expressed as a function of distance rather than time. This can be achieved by multiplying the equations of motion by a scalar, derived by Casanova in [78] for rectangular coordinates and Perantoni [128] for intrinsic coordinates, which converts time increments to distance increments. It can be expressed as:

$$s_r = \left(\frac{ds}{dt}\right)^{-1} = \frac{1 - n\Omega}{u\cos\xi - v\sin\xi} \tag{3.16}$$

3.3.3 Vehicle and tyre model

The vehicle model is the bicycle model as derived in Pacejka [7], where the details of the derivation can be found. The model is similar to the one used in Section 2.2, where the formulation, however, included both lateral and longitudinal dynamics. The equations for the lateral only case are therefore described. They read

$$M(\dot{v} + u\omega) = F_{yf} + F_{yr} + F_{y,dist} \tag{3.17}$$

$$I_z\dot{\omega} = aF_{yf} - bF_{yr} + M_{dist} \tag{3.18}$$

where M is the total mass, I_z the total inertia, v the lateral velocity, u the forward velocity, F is force, $_y$ indicates that it is lateral, $_f$ and $_r$ are front and rear respectively,

$_{dist}$ indicates that it is a disturbance, ω is the angular yaw rate and a and b the distance of the centre of mass from the front and rear axle respectively.

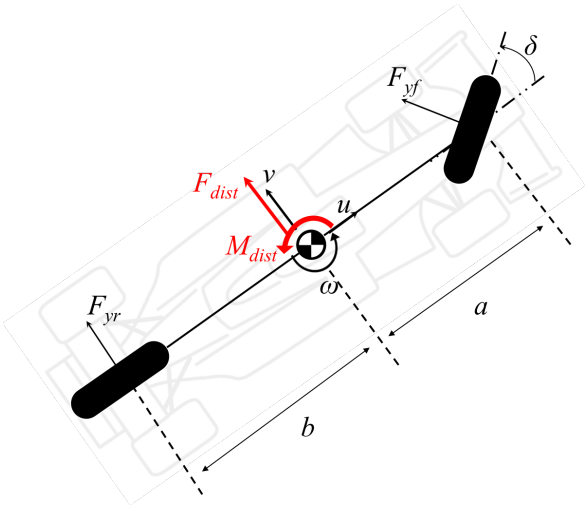

Figure 3.4: Vehicle model used to model the lateral dynamics. Force and moment disturbances are shown in red.

The lateral force is a function of the lateral slip of the tyres, which is in turn a function of the lateral velocity and the yaw rate. The slip can be expressed as

$$\alpha_f = \delta - \frac{v + \omega a}{|u|} \tag{3.19}$$

$$\alpha_r = -\frac{v - \omega b}{|u|} \tag{3.20}$$

where δ is the steering angle at the tyres. Nonlinear tyres are the obvious choice as the aim of the project is to investigate the behaviour of the vehicle close to tyre saturation. A well established model is Pacejka tyre model, described by Eq. (2.11). Even though the model is established and allows the description of almost any possible road friction interaction, its numerical complexity makes the optimal control problem difficult to solve. Whilst for the lateral controller the resulting optimisation problem is numerically tractable, the lateral and longitudinal coupling that occurs when longitudinal forces are considered makes convergence difficult to achieve. Even though it is not strictly necessary for the lateral case, Equation 2.11 needs to be simplified to ensure consistency with the combined controller. This particular work focus on track racing, so modelling different ground surfaces, such as gravel, is not required. Trading control of the shape of the curve against numerical simplicity does not affect the overall performance of the controller. The following tyre model has been adapted from Kelly [23]. Defining the

coefficient of friction as

$$\mu_y = \mu_{ymax} \sin(Q_y \arctan(S_y \alpha_n)) \tag{3.21}$$

where μ_{ymax} is the maximum coefficient of friction, $\alpha_n = \alpha/\alpha_{max}$, where α is the slip angle, α_{max} the maximum slip angle, and S_y is defined as

$$S_y = \frac{\pi}{2 \arctan(Q_y)} \tag{3.22}$$

where Q_y is the stiffness factor. The lateral force can be expressed as

$$F_y = \mu_y F_z \alpha \tag{3.23}$$

where F_z is the vertical force.

Figure 3.5 shows the lateral force characteristics for the front tyre of an Understeering and an Oversteering vehicle. Tyre parameters are $\mu_{ymax} = 1.27$, $Q_y = 1.61$, $\alpha_{max} = 0.12$. These parameters are just indicative and are used for illustration purposes only. The saturation region starts at 0.12 rad.

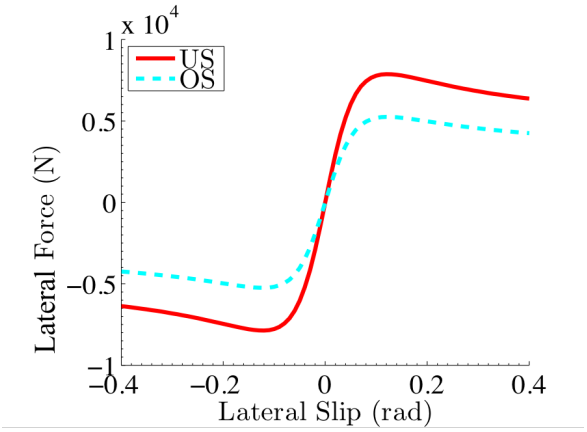

Figure 3.5: Front tyre force slip characteristics for an Understeered and an Oversteered vehicle. The curve shows saturation and the negative slope region.

Equation 3.23 gives the force for equations 3.17 and 3.18 as a function of the slip, which are expressed in terms of v and ω in (3.19).

A second order filter acting on the steering input is added to the system to take limitations of the neuromuscular system (NMS) into account [6]. The filter takes the following form

$$\ddot{\delta}_{sw} + 2\zeta_n \omega_n \dot{\delta}_{sw} + \omega_n^2 \delta_{sw} = \omega_n^2 \delta_{com} \tag{3.24}$$

where ω_n is the NMS natural frequency, ξ_n the NMS damping factor, $_{sw}$ indicates that the signal is applied at the hand wheel and $_{com}$ indicates that the signal is commanded from the brain.

A simple aerodynamic model is also added to account for downforce. As vehicle's speed is constant for a lateral controller, aerodynamic forces will not vary, resulting in a constant increase of tyre vertical load. Furthermore, drag is not considered as longitudinal forces are neglected. Even though computational cost could be reduced by modelling these effects as an increase in vehicle mass and a CoG shift, they are included for further developments. Lift is assumed to be governed by standard quadratic functions

$$F_{az} = \frac{1}{2} C_l \rho A_f u^2 \tag{3.25}$$

where C_l is the lift coefficient, ρ is air density, A_f the vehicle front area and u the vehicle speed. F_{az} is assumed to be positive downwards. Assuming that the distances of the Centre of Pressure (CoP) from the front and rear axles are denoted by a_{CoP} and b_{CoP}, the vertical forces can be calculated in the following fashion.

$$F_{zf} = \frac{b}{L} Mg + F_{az} \frac{b_{CoP}}{L} \tag{3.26}$$

$$F_{zr} = \frac{a}{L} Mg + F_{az} \frac{a_{CoP}}{L} \tag{3.27}$$

where $L = a + b$ is the wheelbase of the vehicle and all other terms have been previously defined.

Disturbances are added to the system to account for the randomness of a real driving scenario. A lateral disturbance acting at the centre of mass perpendicularly to the vehicle's longitudinal axis and a yaw moment disturbance are considered; as Figure 3.4 shows. The lateral and moment disturbance terms F_{dist} and M_{dist} represent various factors that can affect driving such as road roughness or wind. Even though such measurements could be obtained from track data, the disturbance profile used in Chapter 2 – white Gaussian noise with zero mean is adopted. This assumption allows testing of the controller: for the validation work, a different disturbance profile is chosen and it will be discussed in Chapter 5.

3.3.4 Cognitive limitations

The literature shows that real drivers have cognitive limitations, which cannot be accurately represented by an optimal controller. Young and Stanton [129] define

cognitive workload as the level of attentional resources required to meet certain performance criteria. Young and Stanton's research key outcome is that cognitive workload for a given task is related to the driver information processing capacity. Patten [130] shows that human information processing capacity is done through temporary memory – a sort of human Random Access Memory (RAM) – where data can easily be overwritten. This theory suggests that the driver cognitive processing capacity is limited, which implies that online complex optimisations are unlikely to happen. Johns [2] performed theoretical and experimental work to assess the extent to which cognitive limitations affect driving. He enriches existing driver models to account for these limitations and validates the proposed improvements experimentally.

Johns further developed and validated the idea of intermittent control. In traditional MPC, a control sequence is calculated at every time step and the first term of the control sequence applied. Following Patten [130], Johns argues that it is unlikely for the driver to be able to perform an optimisation for frequencies higher than 3 Hz. He therefore advances the idea of intermittent control, which hypothesises that the driver applies a longer sequence of the control inputs from the optimisation rather than only the first one. This approach inevitably leads to a suboptimal solution, and the robustness and accuracy of this scheme are investigated in this thesis. The time over which controls are applied will be referred to as the *intermittency period*, which depends on the human *refractory period*; defined as the period of time over which a nerve or a muscle is unresponsive to further stimuli after being stimulated. The refractory period length is closely related to cerebral activity. The steering signal applied over the intermittency period will be referred to as a *hold*. Johns identifies a typical hold length to be approximately 0.5 s and investigated holds further to determine their nature. Two hold types are identified: Zero Order Hold (ZOH) – when the driver holds the first control action calculated constant over the *intermittency period* – and Serial Ballistic Hold (SBH) – when instances of the calculated control sequence are applied up to the *intermittency period*. Figure 3.6 illustrates the characteristics of a controller incorporating human limitations. The controller samples the states, which are assumed to be known, at time k. The optimal controls, shown by the dashed line, is calculated over the prediction horizon, N_p. In traditional MPC only the first instance of the control sequence is applied to the system. In this case it would be up to $k+1$. In order to account for the driver cognitive limitations, the controls will be applied over the intermittency period, up to $k+N_i$, where N_i is the intermittency period. The ZOH control is shown as the solid red line; the signal is calculated assuming it is constant over the intermittency period N_i and the first instance is applied. In the case of SBH,

shown by the dotted line, the actual control sequence is applied up to $k + N_i$. The controls are suboptimal after $k+1$ if the system is perturbed. Figure 3.6 also shows the predicted vehicle states, dotted black line, and actual vehicle states, solid black line. If the optimal control sequence is applied in absence of disturbances in a traditional MPC fashion, the vehicle states are the same as the predicted states from the optimisation. In the presence of disturbances, as in this case, the actual vehicle states will not be the same as the predicted ones. The controller intermittency period needs to be short enough for the controller to keep the disturbed states bounded.

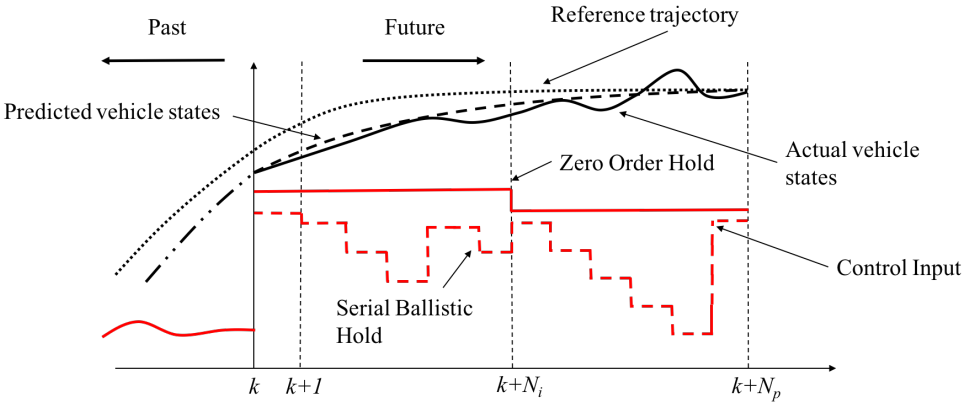

Figure 3.6: Sketch illustrating the core characteristics of the controller.

Johns showed that drivers are more likely to adopt an SBH approach. In order to include the driver cognitive limitations so far described, SBH intermittent control will be assumed for this work.

3.3.5 Cost function

The cost function $J(\boldsymbol{x}, \boldsymbol{u})$ needs to minimise the lateral offset and ensure the bandwidth of the input signal to the neuromuscular system is computable by a human. The proposed cost function is

$$J = \boldsymbol{n}^T q_1 \boldsymbol{n} + (\boldsymbol{\xi} - \boldsymbol{v}_{nom}/\boldsymbol{u}_{nom})^T q_2 (\boldsymbol{\xi} - \boldsymbol{v}_{nom}/\boldsymbol{u}_{nom}) + \dot{\boldsymbol{\delta}}_{com}^T q_3 \dot{\boldsymbol{\delta}}_{com} \tag{3.28}$$

where q_n are the cost function weights, bold italic indicates a vector over the whole prediction horizon, the subscript $_{nom}$ the nominal value of the signal and all other terms have been previously defined. $\boldsymbol{u}$ indicates the vector of longitudinal vehicle velocities. A quadratic cost function has been chosen because of its numerical tractability. The term $\boldsymbol{n}^T q_1 \boldsymbol{n}$ minimises the lateral offset. This term of the cost function shows the

clear advantage of the intrinsic formulation: the lateral offset can be expressed as state of the system and enters the cost function as a simple quadratic term. The term $(\boldsymbol{\xi} - \boldsymbol{v}_{nom}/\boldsymbol{u}_{nom})^T q_2 (\boldsymbol{\xi} - \boldsymbol{v}_{nom}/\boldsymbol{u}_{nom})$ penalises the heading of the vehicle, effectively targeting the the rate of change of the lateral offset. The last term $\dot{\boldsymbol{\delta}}_{com}^T q_3 \dot{\boldsymbol{\delta}}_{com}$ penalises the rate of the change of the commanded steering angle to the neuromuscular system, avoiding unrealistic frequency content. This term penalises the signal coming out of the brain, hence serving a different purpose to the NMS filter.

3.3.6 Constraints

A number of constraints are needed to complete the description of the problem. The dynamics of the vehicle define the most important constraint, the equality constraint which can be conveniently expressed in state space $\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u})$ from Eq. 3.10. The states of the vehicle are obtained by its dynamics, described in Section 3.3.3, the driver neuromuscular system, Eq. 3.24 and the track, Eqs. 3.13 3.14 and 3.15. All the equations are multiplied by s_r, defined in Eq 3.16 to change the independent variable from time to distance. There are no inequality constraints on the system as it only has to follow the nominal line. However, a constraint is put on the tyre slip to limit the search space to realistic tyre slip. This is done purely to make the numerical problem more tractable. The slip constraint is at twice the saturation slip to ensure the controller can explore the totality of the tyre curve. Hence

$$\alpha < 2\alpha_{sat} \tag{3.29}$$

where α is the lateral slip and α_{sat} is the saturation slip.

All terms in Eq. 3.10 have been defined. The states of the vehicle by Eqs. 3.13, 3.14, 3.15, 3.17 and 3.24 which result in the following system of equations:

$$\dot{v} = s_r \left(\frac{1}{M} \left(F_{yf} + F_{yr} + F_{ydist} \right) - u\omega \right) \tag{3.30}$$

$$\dot{\omega} = s_r \left(\frac{1}{I_z} \left(aF_{yf} - bF_{yr} + M_{dist} \right) \right) \tag{3.31}$$

$$\dot{\xi} = s_r \omega - \Omega_z \tag{3.32}$$

$$\dot{n} = s_r \left(u \sin(\xi) + v \cos(\xi) \right) \tag{3.33}$$

$$\ddot{\delta}_{sw} = s_r \left(-2\zeta\omega_n \dot{\delta}_{sw} - \omega_n^2 \delta_{sw} + \omega_n^2 \delta_{com} \right) \tag{3.34}$$

The cost function is defined in Eq. 3.28.

3.4 Combined controller

The controller framework so far developed can be easily extended to include the longitudinal dynamics as the problem formulation is very similar; relying, as it does, on the same robust control ideas. In fact, the general formulation cast in Section 3.3.1 is valid for a longitudinal scheme as well. The vehicle, driver and tyre model have to be adapted to include the effect of varying speed to redefine the various terms of Eq. 3.10.

3.4.1 Vehicle model

The vehicle longitudinal dynamics are captured by a simple trolley model, which allows the vehicle to change its speed. Wheel dynamics are neglected, as spinning wheels are characterised by fast dynamics which would significantly increase the numerical complexity of the problem without any great improvement in accuracy [84]. This is due to the small inertia of the wheel compared to the overall vehicle. The control inputs for the longitudinal dynamics are therefore the front and rear longitudinal slip, which enter the equation through the tyre model. The two are treated as independent control inputs. A longitudinal disturbance is also added to the system as a force acting longitudinally through the centre of mass.

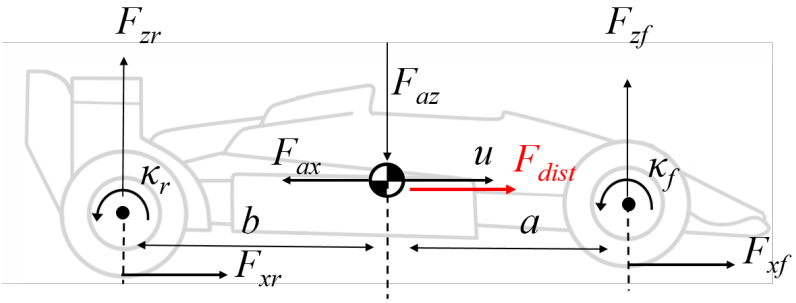

Figure 3.7: Trolley model showing the forces and the longitudinal force disturbance

The equation of motion of the vehicle are different from 2.1 because wheel dynamics are not accounted for. They can therefore be expressed as

$$M(\dot{v} + u\omega) = F_{yf}\cos\delta + F_{xf}\sin\delta + F_{yr} + F_{y,dist} \tag{3.35}$$

$$I_z\dot{\omega} = a(F_{yf}\cos\delta + F_{xf}\sin\delta) - bF_{yr} + M_{dist} \tag{3.36}$$

$$M(\dot{u} - v\omega) = F_{xf}\cos\delta - F_{yf}\sin\delta + F_{xr} - F_{ax} + F_{x,dist} \tag{3.37}$$

where δ is the steering angle at the tyres, M is the total mass, I_z the total yaw moment of inertia about the centre of mass, v the lateral velocity, u the forward velocity, F is force, $_y$ indicates that it is lateral,$_x$ that is longitudinal, $_f$ and $_r$ are front and rear respectively, $F_{y,dist}$ the lateral disturbance,M_{dist} the moment disturbance, $F_{x,dist}$ the longitudinal disturbance, ω is the angular yaw rate and a and b the distance of the centre of mass from the front and rear axle respectively. Small angles have not been assumed because the longitudinal component of the front tyre force has an effect on the speed, so (3.35) and (3.36) include trigonometric expressions while (3.17) and (3.18) do not. An aerodynamic force is introduced by the term F_{ax}. Aerodynamic drag is assumed to take the simple form

$$F_{ax} = \frac{1}{2} C_x \rho A_x u^2 \tag{3.38}$$

where C_x is the drag coefficient and all other terms have already been defined.

3.4.2 Tyre model

The tyre model has to take into account the lateral, longitudinal forces and the effect of the friction circle. The tyre model used by Kelly in [23] is adopted. The model is a parameter-reduced version of Pacejka's, the same used in Section 3.3.3 for lateral dynamics, sacrificing control over the curve shape for numerical simplicity. The ability to model very specific – and uncommon – tyre-surface characteristics is not relevant as the aim of this work is to evaluate the controller performance at the limit of handling and not to replicate very specific tyre force-slip characteristics. Furthermore, all the experimental work is done on asphalt, which can be reproduced by the proposed tyre model. The lateral slips are defined in Eq. 3.19 and Eq. 3.20. Pacejka's convention is used to define the longitudinal slip

$$\kappa = -\frac{u - r\Psi}{u} \tag{3.39}$$

where u is the forward speed, r is the wheel radius and Ψ is the wheel angular speed. The combined slip is defined as

$$\epsilon = \sqrt{\alpha_n^2 + \kappa_n^2} \tag{3.40}$$

where α_n and κ_n are the normalised slip which are defined as

$$\alpha_n = \frac{\alpha}{\alpha_{max}} \tag{3.41}$$

$$\kappa_n = \frac{\kappa}{\kappa_{max}} \tag{3.42}$$

where α_{max} and κ_{max} are the maximum lateral and longitudinal slip and are parameters of the tyre model. The maximum force in the lateral and longitudinal direction is assumed to be a function of the vertical force and the coefficient of friction which is defined as

$$\mu_x = \mu_{xmax} \sin(Q_x \arctan(S_x \epsilon)) \tag{3.43}$$

$$\mu_y = \mu_{ymax} \sin(Q_y \arctan(S_y \epsilon)) \tag{3.44}$$

where S_x and S_y are defined as

$$S_x = \frac{\pi}{2 \arctan(Q_x)} \tag{3.45}$$

$$S_y = \frac{\pi}{2 \arctan(Q_y)} \tag{3.46}$$

where all terms have been previously defined. Finally, the lateral and longitudinal forces can be expressed as

$$F_x = \mu_x F_z \frac{\kappa_n}{\epsilon} \tag{3.47}$$

$$F_y = \mu_y F_z \frac{\alpha_n}{\epsilon} \tag{3.48}$$

Fig. 3.8 shows the force slip characteristics for a tyre with the following parameters: $\mu_{max,y}$=1.2734, $\mu_{max,x}$=1.2736, Q_y=1.5978, Q_x=1.6092, α_{max}=0.0897, κ_{max} =0.0902. These parameters have been obtained by the author while testing the tyre fitting algorithm used in Chapter 5 and Chapter 6. Saturation is present in the both the lateral and longitudinal direction, with only a slight asymmetry present.

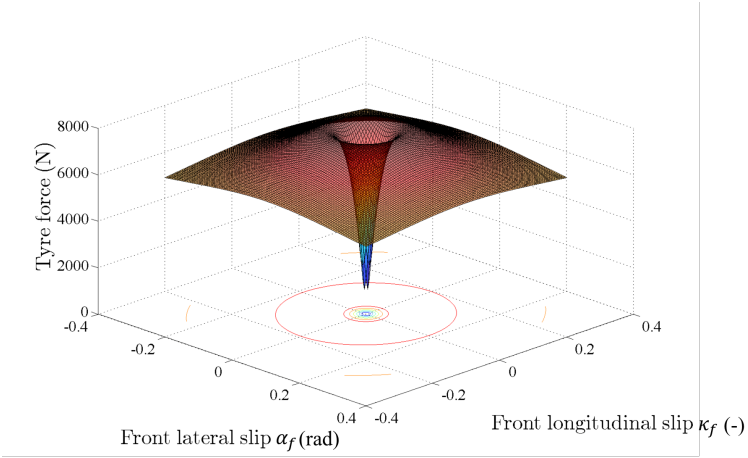

Figure 3.8: Tyre force surface as a function of lateral and longitudinal slip

3.4.3 Cost function

The cost function needs to be expanded to account for longitudinal dynamics. An expression for cost function is given and all terms explained

$$J = \boldsymbol{n}^T q_1 \boldsymbol{n} + (\boldsymbol{\xi} - \boldsymbol{v}_{nom}/\boldsymbol{u}_{nom})^T q_2 (\boldsymbol{\xi} - \boldsymbol{v}_{nom}/\boldsymbol{u}_{nom}) + \dot{\boldsymbol{\delta}}_{com}^T q_3 \dot{\boldsymbol{\delta}}_{com} + (\boldsymbol{V} - \boldsymbol{V}_{nom})^T q_4 (\boldsymbol{V} - \boldsymbol{V}_{nom}) + \dot{\boldsymbol{\kappa}}_f^T q_5 \dot{\boldsymbol{\kappa}}_f + \dot{\boldsymbol{\kappa}}_r^T q_6 \dot{\boldsymbol{\kappa}}_r + (\boldsymbol{\kappa}_f + \boldsymbol{\kappa}_r)^T q_7 (\boldsymbol{\kappa}_f + \boldsymbol{\kappa}_r) \tag{3.49}$$

where all vectors are up to prediction horizon,$\boldsymbol{u}$ is the vector of forward velocities, and $_{nom}$ indicates nominal. The first three terms are the same as for the lateral controller: minimising the lateral path error, the heading of the vehicle and $\dot{\delta}$. The difference between the total velocity $V = \sqrt{v^2 + u^2}$ and the total reference speed $V_{nom} = \sqrt{v_{nom}^2 + u_{nom}^2}$ only minimises the difference in speed magnitude but not in direction. However, $(\boldsymbol{\xi} - \boldsymbol{v}_{ref}/\boldsymbol{u}_{ref})^T q_2 (\boldsymbol{\xi} - \boldsymbol{v}_{ref}/\boldsymbol{u}_{ref})$ ensures that the heading is correct. $\dot{\kappa}_f$ and $\dot{\kappa}_r$ are also minimised to avoid unrealistic frequency content in the longitudinal slip controls. Finally, the sum of the squares of front and rear longitudinal slip is also minimised with a very low weight q_7 to encourage the driver not to accelerate and brake at the same time.

3.4.4 Constraints and filter

A filter is also added to the longitudinal dynamics to replicate the neuromuscular behaviour of the driver. The same filter is added for the front and rear slip and takes

the same form as the NMS filter imposed on lateral dynamics.

$$\ddot{\kappa}_p + 2\zeta_n\omega_n\dot{\kappa}_p + \omega_n^2\kappa_p = \omega_n^2\kappa_{com} \tag{3.50}$$

where $_p$ indicates the applied signal and $_{com}$ commanded. Lateral and longitudinal tyre slips are also constrained to avoid unrealistic values of slip. Again, the constraint is put at twice the maximum slip angle the vehicle can achieve.

$$\alpha < 2\alpha_{sat} \tag{3.51}$$

$$\kappa < 2\kappa_{sat} \tag{3.52}$$

where κ denotes the longitudinal slip for front and rear tyres. The front slip is also constrained to be less than zero, so it can only be used in braking

$$\kappa_f < 0 \tag{3.53}$$

The model is fully characterised. The complete system of equations for the longitudinal case is:

$$\dot{v} = s_r\left(\frac{1}{M}(F_{yft} + F_{yrt} + F_{ydist}\right) - u\omega \tag{3.54}$$

$$\dot{\omega} = s_r\left(\frac{1}{I_z}(aF_{yft} - bF_{yrt} + M_{dist})\right) \tag{3.55}$$

$$\dot{u} = s_r\left(\frac{1}{M}(F_{xft} + F_{xrt} + F_{xdist}) + v\omega\right) \tag{3.56}$$

$$\dot{\xi} = s_r\omega - \Omega_z \tag{3.57}$$

$$\dot{n} = s_r(u\sin(\xi) + v\cos(\xi)) \tag{3.58}$$

$$\ddot{\delta}_{sw} = s_r(-2\zeta\omega_n\dot{\delta}_{sw} - \omega_n^2\delta_{sw} + \omega_n^2\delta_{com}) \tag{3.59}$$

$$\ddot{\kappa}_p = s_r(-2\zeta_n\omega_n\dot{\kappa}_p - \omega_n^2\kappa_p + \omega_n^2\kappa_{com}) \tag{3.60}$$

The parameter for the NMS filter are assumed to be the same for lateral and longitudinal controls.

The total forces can be expressed as

$$F_{yft} = F_{xf} \sin \frac{\delta_{sw}}{G_{sw}} + F_{yf} \cos \frac{\delta_{sw}}{G_{sw}} \tag{3.61}$$

$$F_{xft} = F_{xf} \cos \frac{\delta_{sw}}{G_{sw}} - F_{yf} \sin \frac{\delta_{sw}}{G_{sw}} \tag{3.62}$$

$$F_{yrt} = F_{yr} \tag{3.63}$$

$$F_{xrt} = F_{xr} \tag{3.64}$$

Vertical forces are described by Eq. 3.26. Small angle approximations have not been adopted because the longitudinal force of the front tyre in steering is relevant to speed control. More specifically, approximating to 0 the term $\sin \frac{\delta_{sw}}{G_{sw}}$ in Eq (3.62) would lead to a significant error in the modelling of the longitudinal dynamics because when F_{yf} is large, the braking action is not negligible.

3.5 Solution routine

A direct method is preferred because the conditions for optimality do not have to be directly imposed on the continuous system, allowing for easy incorporation of different driving strategies, such as SBH and ZOH. ICLOCS, a open-source optimisation suite from Imperial College [131], has been chosen to implement the optimisation. The problem is first transcribed using a multiple shooting approach to increase the robustness of the solution. The CVODES package is used to determine the solution to the initial value problem. Data structures for the partial derivatives of the stage cost, the constraints and the system dynamics are found numerically. The resulting discrete problem is then solved using IPOPT, an interior point algorithm implemented in C++ which interfaces with Matlab through a Mex file. IPOPT is a particularly efficient algorithm, which eliminates the inequality constraints introducing barrier logarithmic functions. IPOPT uses the Fortran based MA57 matrix solver to deal with the large, sparse matrices that are obtained when evaluating inequality constraints. Further details on the architecture of IPOPT can be found in [131]. The solver parameters are summarised in Table 3.1. Results are shown in Chapter 4.

3.6 Summary

This Chapter details the derivation of the lateral and longitudinal controller and the relevant theoretical background.

Table 3.1: ICLOCS parameters

Parameter	Function	Value
options.derivatives	Derivatives calculation	Numeric
options.transcription	Transcription method	Hermite
options.hessianFD	Hessian matrix calculation method	central
options.NLP solve	Non linear problem solver	ipopt
options.ipopt.tol	Ipopt tolerance	1e-10
options.ipopt.mu	Logarithmic barrier strategy	adaptive

Optimal control and robust control theory is described first. The problem for the lateral controller is then cast and the track and vehicle models described. The model is then enriched with various features to mimic actual driving behaviour and subsequently extended to include longitudinal dynamics. An overview of the solution methodology is also presented.

The contribution to knowledge of this Chapter is the design of a controller than can mimic the action of a human driver stabilising a vehicle at the handling limit under the action of random disturbances. Several published controllers can achieve stability for a vehicle with saturating tyres, however, none considers disturbances and human limitations in such context.

Chapter 4

Controller performance assessment

4.1 Introduction

Optimal and robust control have been used in Chapter 3 to propose a controller which can stabilise vehicles at the handling limit under the action of disturbances. Two separate formulations for lateral and combined controls have been considered. The aim of this Chapter is to gain insight into the performance of the controllers as well as assessing the effect the model parameters have on the output. A variety of driving scenarios are selected to highlight the features of the controller.

The performance of the lateral controller is assessed first. Scenarios with increasing levels of non-linearity are considered to determine if the control action can stabilise the vehicle at the handling limit. The chosen target path is a circle, as it allows assessment of the vehicle's behaviour about a steady operating point. Two vehicle configurations are considered to evaluate the different driving strategies: Under Steering (US) and Over Steering (OS). A parameter study is then carried out to assess the effect of the main controller parameters. Since this work focuses on mimicking a human driver, the following driver parameters are considered: the NeuroMuscular System (NMS) cut-off frequency (ω_n), the update time T_u and the penalty on the hand wheel angle steering rate q_3.

The combined controller is then considered. Adding longitudinal controls significantly increases complexity, particularly for highly nonlinear manoeuvres where the coupling between lateral and longitudinal dynamics is significant. The controller is first tested on a circular path at constant speed in a highly nonlinear scenario to show that the lateral and longitudinal controls can stabilise the vehicle at the limit of handling. A more realistic driving scenario, a 90 degree manoeuvre is then considered to show that the controller can track and stabilise the vehicle as speed and target path curvature

vary. Finally, a parameter study is carried out, not only to assess the effect of driver parameters on the controller performance, but also to investigate the coupling between the lateral and longitudinal dynamics.

Section 4.2 details the baseline parameters used for both the lateral and combined controller. Section 4.3 includes results for the lateral controller navigating a circle at constant speed. The action of the combined controller is explored in Section 4.4.

4.2 Vehicle and driver model parameter values

An understanding of the vehicle's and controller's parameters is fundamental for attempts to replicate the behaviour of a human driver. In order to structure the analysis, they are categorised in the following way: vehicle parameters and driver parameters.

4.2.1 Vehicle parameters

These parameters define the vehicle characteristics and its handling properties. They only affect driver responses indirectly so they will not be considered in detail in the parameter studies. The fundamental objective of the controller, which is to stabilise a vehicle at its limit of adhesion, is tested together with the trade-off between path error and control action. Baseline vehicle parameters are shown in Table 4.1. The baseline values are the same as the vehicle model used in Chapter 2. Some of the parameters are missing because of the differences in vehicle models discussed in Chapter 3.

The OS vehicle configuration is achieved by shifting the vehicle centre of mass rearwards, setting $a = 1.38$ m and $b = 0.92$ m, effectively increasing the cornering stiffness of the rear tyre as shown in [7].

Tyre parameters determine the slip force characteristics of the tyres and also significantly affect the response of the vehicle, however, not being directly related to the driver response, will not be investigated in depth. They are chosen to represent a standard nonlinear tyre on asphalt, which saturates above a certain value of slip; values, taken from [105], are shown in Table 4.4.

4.2.2 Driver parameters

Driver parameters chosen to explore the behaviour of the controller rather than to attempt the behaviour of a human driver, which will explored in Chapters 5 and 6. The NMS values are chosen to replicate the bandwidth of an expert driver; they

Table 4.1: Vehicle parameters from [105]

Parameter	Symbol	Value
Vehicle mass	M_t	1050 kg
CoG to front axle distance	a	0.92 m
CoG to rear axle distance	b	1.38 m
Steer to road wheel angle ratio	G_{sw}	17
Centre of Pressure (CoP)	c_p	0.5
Cross section area	A_x	2 m^2
Lift coefficient	C_l	0.26
Drag coefficient	C_x	0.35
Air density	ρ	1.2 kg/m^3

are adopted from Timings [105]. Optimal solutions are unaffected by changes in the prediction horizon if this is long; running simulations for various prediction horizon lengths showed that the threshold is around 3 seconds. A 4 seconds prediction horizon is chosen to ensure it does not affect the behaviour of the controller. Johns [2] suggests an update time of 0.5 s. If random disturbances are applied to the system, only the first instance of the control sequence guarantees stability [102]. Even though the calculated control sequence is optimal, if disturbances are applied the longer the sequence used, the farther the system deviates from optimality. Error propagation depends on the nature of the system [34]; an intrinsically stable US vehicle allows for longer update times even when tyres are close to saturation while an unstable OS vehicle needs higher frequency feedback [7]. Various tests showed that a 0.5 s update time does not guarantee stability for OS vehicle. A brief heuristic study showed in fact that a 0.05 s update time -which corresponds to only applying the optimal input- is necessary to guarantee stability at the handling limit for an OS vehicle. Even though such as short update time would only be necessary for the OS vehicle at the limit of adhesion, the same value is also used for the other runs to make them comparable and facilitate the analysis. The update time is not representative of a human driver, but at this stage the focus is on the performance of the controller rather than its accuracy in replicating human driving actions, which will be considered for the validation in Chapters 5 and 6. Driver parameters, shown in Table 4.2, are taken from [93] and they represent a skilled driver [93].

Cost function parameters also significantly affect driver's response as they determine the overall level of cognitive effort and how it is allocated. They are not directly related

Table 4.2: Driver parameters. Lateral and longitudinal NMS parameters are assumed to be the same.

Parameter	Symbol	Value
NMS damping ratio	ζ_n	0.707
NMS natural frequency	ω_n	18.85 rad/s
Update time	T_u	0.05 s
Prediction horizon	T_s	4 s

to physical quantities, hence choosing baseline values is inevitably heuristic. The controller is only tested for robustness, so the cost function weights values are only indicative and they are chosen to yield reasonable results. Firstly, penalty on LPE is set to 1 and used as a reference; penalty on heading is also set to 1 so that it is penalised less than the LPE, as the magnitude of the former is typically more than one order of magnitude smaller than the latter. The frequency content of the HWA is penalised with a smaller weight to ensure the driver can reject high frequency disturbances. The penalty on the speed deviation, is the same as the one on the LPE, so the driver allocates the same cognitive resources to speed and path tracking. Again, the choice is arbitrary and a real driver may opt for a different strategy but this will be discussed in details in Chapter 5 and Chapter 6. The weights on $\dot{\kappa}_f$ and $\dot{\kappa}_r$ are chosen to yield a prompt longitudinal response. Finally, q_7 is chosen to avoid braking and accelerating on the straight without affecting the longitudinal controls significantly. It is found that given that a value of 0.0001, does not affect the vehicle behaviour yet ensures that accelerating and braking do not happen simultaneously on the straight. Table 4.3 shows the baseline cost function values.

Table 4.3: Cost function parameters

Parameter	Symbol	Value
Penalty on deviation from nominal path	q_1	1
Penalty on deviation from nominal heading	q_2	1
Penalty on hand wheel angle speed	q_3	0.0001
Penalty on deviation from nominal speed	q_4	1
Penalty on front longitudinal slip rate	q_5	0.001
Penalty on rear longitudinal slip rate	q_6	0.001
Penalty on longitudinal slip sum	q_7	0.0001

Tyre parameters are chosen to represent symmetrical, standard tyres and are taken from [105]. All the baseline parameters are shown in Tables 4.1, 4.2, 4.3 and 4.4. The same values were used for the bicycle model in Chapter 2 but are also reported here for completeness.

Table 4.4: Nominal tyre parameters. The tyre is symmetric in the x and y direction.

Parameter	Symbol	Value
Lateral stiffness factor	Q_y	1.6
Maximum lateral slip	α_{max}	0.12
Lateral friction coefficient	μ_y	1.3
Longitudinal stiffness factor	Q_x	1.6
Maximum longitudinal slip	κ_{max}	0.12
Longitudinal friction coefficient	μ_x	1.3

4.2.3 Disturbances

Disturbances are added to assess how the controller stabilises the vehicle as it deviates from the reference trajectory. A random disturbance force is added to each of the three degrees of freedom of the vehicle. A lateral force, $F_{y,dist}$ acting perpendicularly to the centre of mass (CoM), disturbs the vehicle laterally, a longitudinal force, $F_{x,dist}$ acting parallel to the vehicle centreline disturbs the vehicle longitudinally and a yaw moment, $M_{z,dist}$ acting at the centre of mass (CoM). These disturbances are described in detail in Section 2.4. The nature of the disturbances is arbitrary and can be changed to investigate the behaviour of the controller under different conditions. Such random disturbances represent not only the unpredictable factors that characterise driving – such as road roughness and wind speed – but also mismatches between predicted model and real model. The nominal value of the standard deviations of the forces and moments is taken from Ulsoy *et al.* [111] and included in Table 4.5.

For the parameter study in Section 4.4.3, a pulse disturbance is also introduced to understand how the driver model recovers from a loss of grip.

Table 4.5: Disturbance values

Parameter	Symbol	Value
$F_{y,dist}$ standard deviation	$\sigma_{F,y}$	730 N
$F_{x,dist}$ standard deviation	$\sigma_{F,x}$	1430 N
$M_{z,dist}$ standard deviation	$\sigma_{M,z}$	360 Nm

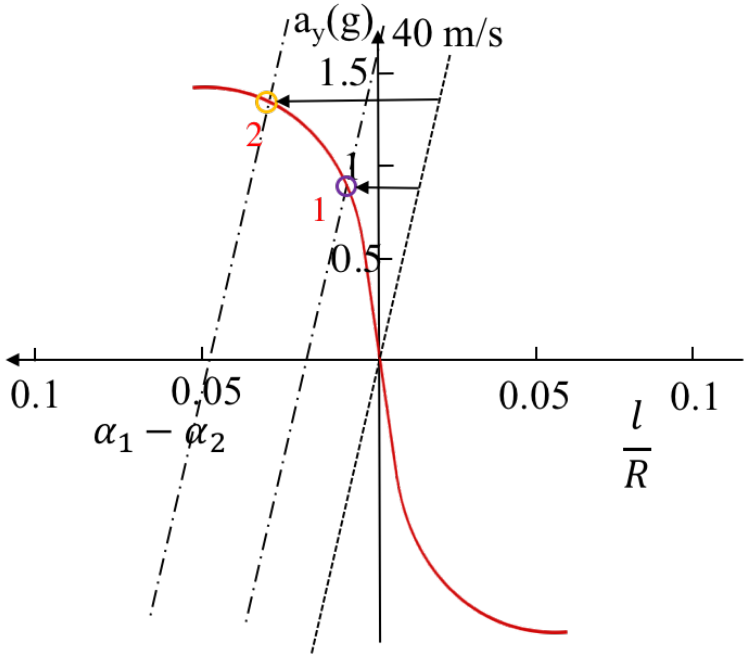

Figure 4.1: Handling diagrams for the US vehicle. Operating points for both US cases, are shown. Stability margins are not shown because US vehicles are always stable.

4.3 Lateral controller results

The lateral controller is first tested in a simple scenario where the vehicle navigates a target circle at constant speed. The nominal solution is clearly a constant steering angle, whose magnitude depends on the radius of the circle, the speed and the vehicle configuration. Since the nominal solution is constant, the effect of disturbances is apparent. This set up readily allows for changes in the slip level, so as to explore the totality of the tyre curve. The aim of this simulation is to demonstrate that the lateral controller can compensate for disturbances acting on the vehicle when the tyres are in the nonlinear part of the slip-force curve, particularly the saturation region.

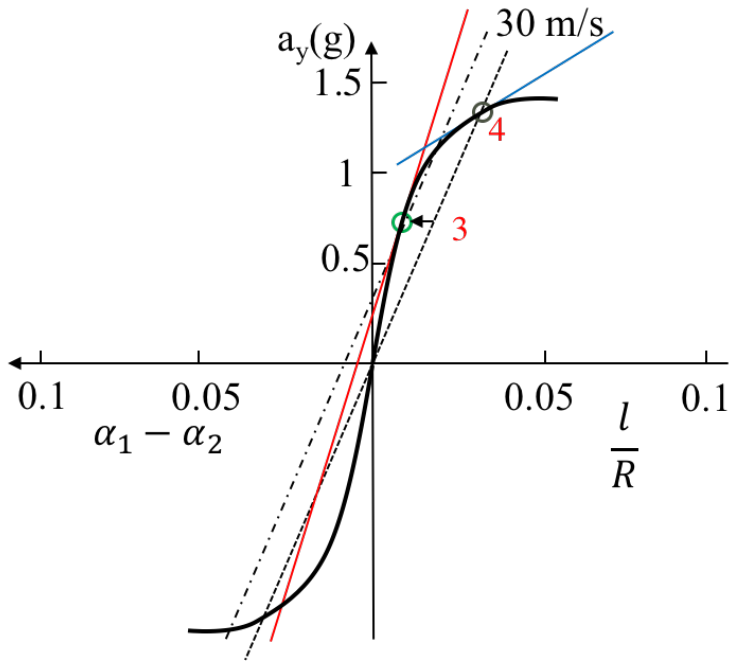

Figure 4.2: Handling diagrams for OS vehicle. The operating points for cases 3 and 4 are shown. The red and light blue solid lines tangent to the handling diagram represent respectively the critical speeds for points 3 and 4. Point 3 is stable because the critical speed at that point, represented by the tangent to the curve, is higher than 30 m/s. Point 4 is unstable because the critical speed is lower than the vehicle speed.

Two vehicle configurations, Under-Steered (US) and Over-Steered (OS), are considered at different slip levels for a total of four cases; these are described in Table 4.6. An US and an OS vehicle have been chosen because, whilst the former is easier to handle and highlights the compensatory action of the controller, the latter is unstable at the speed and radius it is tested; it is therefore an ideal test for the robustness of the controller. Two linear cases, Case 1 and Case 3, are considered to provide predictable solutions against which the nonlinear ones can be compared. Furthermore, they serve as a preliminary validation; observing trends such as decreasing HWA for similar slip levels confirms that the controller behaves as expected. Both US and OS vehicle cases are tested at the handling limit – Case 2 and Case 4 – to show that saturating tyres do not cause instability. For each case, the HWA, LPE, slips and side-slip angle are shown.

The states of vehicles in all four tests converge to a steady state solution. Handling diagrams are an useful way to graphically represent non-linear steady cornering solutions. The method is established and reported in details in [7] so only an overview is provided

here. The following force balance can be derived for the steady state solution of the bicycle model

$$\frac{F_{y,f}}{F_{z,f}} = \frac{F_{y,r}}{F_{z,r}} = \frac{a_y}{g} \tag{4.1}$$

where a_y is the lateral acceleration. Eq (4.1) states that the ratio of the normalised tyre loads is equivalent to the normalised lateral acceleration. Considering the kinematic relationship

$$\delta - (\alpha_f - \alpha_r) = \frac{w_b}{R} \tag{4.2}$$

where w_b is the wheelbase and R the radius of the circle, a plot of $\alpha_f - \alpha_r$ against a_y/g, usually referred to as *handling diagram*, can be obtained by subtracting the normalised tyre characteristics. The diagram is completed by adding the straight line that shows the relationship between lateral acceleration and the relative path curvature at constant speed.

The handling diagram can be used to find the equilibrium conditions for a vehicle with nonlinear tyres navigating a circle at constant speed by shifting the speed line to the left over a distance equal the steering angle. The handling diagram can also be used to assess the stability of the vehicle. The critical speed, which only applies to OS vehicles, can be found by finding the tangent to the handling diagram at the operating point. If the vehicle speed is greater than the critical speed, the vehicle is unstable.

Figure 4.1 shows the handling diagram, solid red line, for the US vehicle and the speed line for 40 m/s, dashed line. This line is shifted to find the two operating equilibrium points, 1 and 2. The shifted lines are shown as dashed-dotted lines. These are for the undisturbed vehicle. The US vehicle is always stable so stability margins need not to be investigated.

In Figure 4.2 the solid black thick lines shows the handling diagram for the OS vehicle. The speed line for 30 m/s is shown as a dashed line. The line is shifted to find the equilibrium for point 3; the steering necessary to achieve it is small. This point is stable as the tangent to the curve, solid red line, is steeper than the speed line for 30 m/s. Point 4 is essentially on the handling diagram, in fact the steering necessary to turn the vehicle is essentially 0. The tangent has lower gradient than the vehicle speed line, indicating that the vehicle is unstable.

4.3.1 Under Steering (US) vehicle

Figure 4.3 shows the results for Case § 1, where an US vehicle navigates a 200 m radius circle at 40 m/s. Tyres are operating beyond the linear region but before saturation,

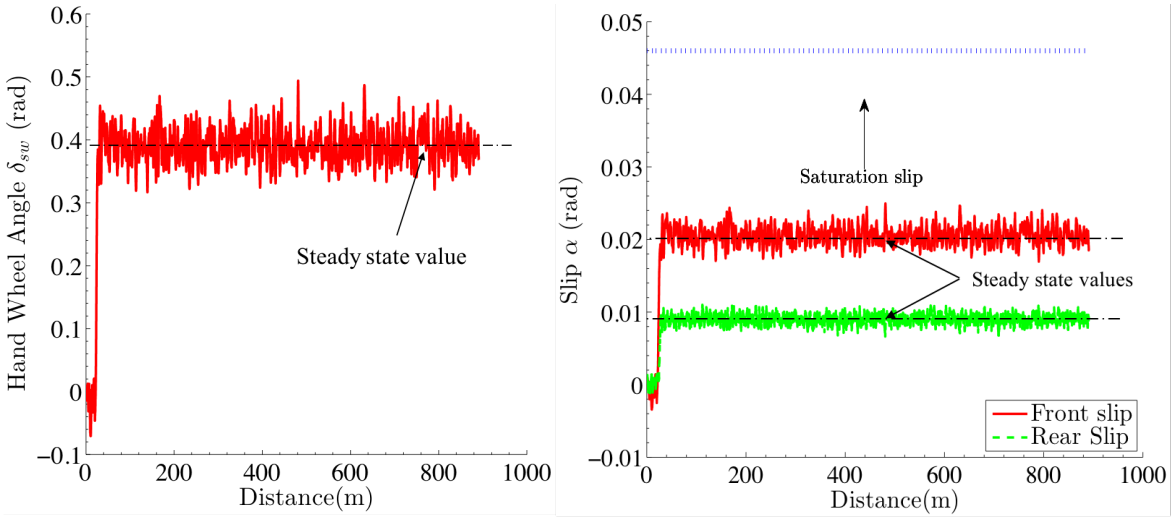

(a) Hand Wheel Angle

(b) Front and rear slip

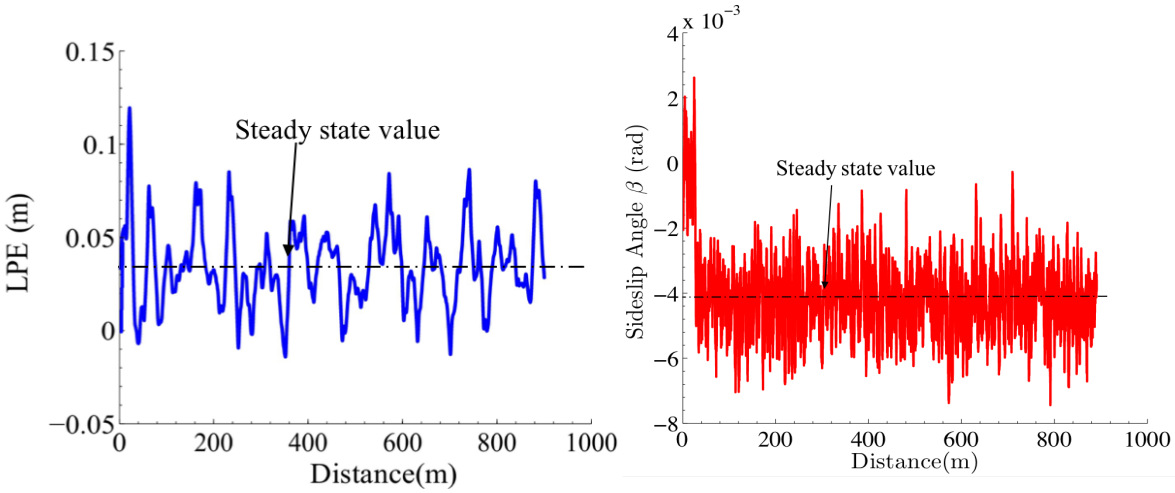

(c) Lateral Path Error

(d) Side-slip angle

Figure 4.3: US vehicle navigating a 200 m radius circle at 40 m/s under the action of disturbances. Tyres are just out of the linear region but far from saturation. The controller is successful at rejecting disturbances and bounding the states. In order to judge the frequency content, the distance travelled can simply divided by 40 to obtain the corresponding time into the simulation.

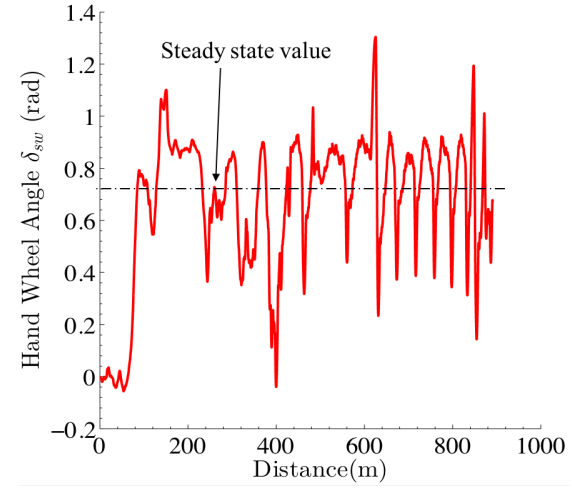

(a) Hand Wheel Angle

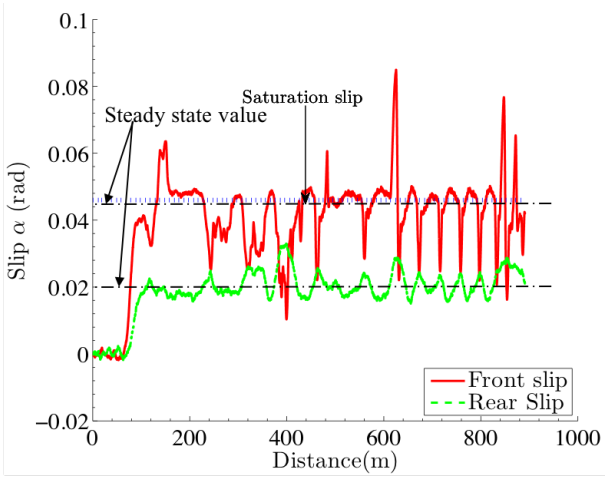

(b) Front and rear slip

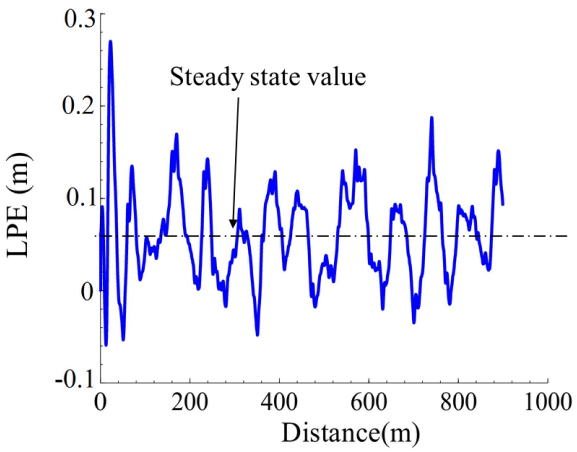

(c) Lateral Path Error

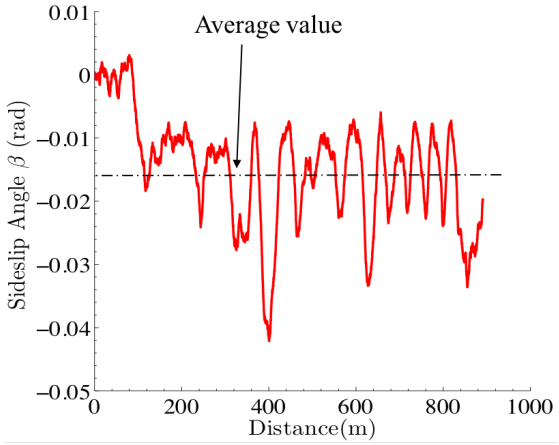

(d) Side-slip angle

Figure 4.4: US vehicle navigating a 110 m radius circle at 40 m/s under the action of disturbances. Even though tyres are saturating the controller is successful at rejecting disturbances and bounding the states. In order to judge the frequency content, the distance travelled can simply divided by 40 to obtain the corresponding time into the simulation.

Table 4.6: Testing conditions for the lateral controller. Four cases are considered to explore the capabilities of the controller and asses its ability to reject disturbances. Different velocities have been used to show that the controller is able to stabilise fast as well as slower vehicles.

Case	Nonlinearity	Vehicle configuration	Radius (m)	Speed (m/s)
§ 1	Mild	US	200	40
§ 2	High	US	110	40
§ 3	Mild	OS	120	30
§ 4	High	OS	80	30

so the action of the controller is focussed only on rejecting disturbances. Figure 4.3a shows that the Hand Wheel Angle (HWA) is characterised by a high frequency response. Since disturbances have zero mean, the HWA mean is very close to the steady state value. The effect of tyre nonlinearity is small because for this test deviations from the steady state value are small. The frequency of the oscillations would be lower if the update time were higher, so the results are not representative of a real human driver. Again, the validation in Chapter 5 and 6 will consider a more accurate comparison with an actual human driver. Figure 4.3b shows that the front slip is greater than the rear slip, as expected for a US vehicle. Both slips are far from saturation, indicated by the dotted line. Again, oscillations around the nominal value are present due to disturbances. Figures 4.3c and 4.3d show that the controller successfully manages to keep the Lateral Path Error LPE and the side-slip angle bounded. The mean LPE is not zero because integral action is not present. This applies to all cases. Case §2 provides a more interesting set of results, shown in Figure 4.4. The radius of the circle is reduced to 110m to push the tyres into the saturation region. Figure 4.4a shows that the controller operates in a radically different fashion under these conditions. Steering corrections are characterised by high amplitude and low frequency. The tyre force-slope has a low gradient, which implies that large changes in HWA are necessary to change the tyre forces to reject disturbances. Since changes are large, the frequency is lower because $\dot{\delta}$ is limited by the cost function. As shown in Figure 4.4b, the front slip is significantly higher than the rear, and even though it goes well beyond the saturation point, the controller successfully stabilises the vehicle. The magnitude of the slip increases suddenly and significantly as it crosses the saturation point. This is due to the disturbances pushing the tyres into the negative slope region. Since the force disturbance is normally distributed, peaks of 4σ can have a significant effect when the vehicle is at the limit of adhesion. Figure 4.4c confirms that the LPE is

effectively bounded and the vehicle manages to follow the target path. Finally, Figure 4.4d shows that the side-slip angle follows a pattern which is very similar to the steering, characterised by the superposition of high and low frequency dynamics. Again, the controller successfully manages to bound the side-slip angle despite the saturating tyres, ensuring good path tracking.

4.3.2 Over Steering (OS) vehicle

Results for Case §3 are shown in Figure 4.3. Even though speed and radius for Case §1 and Case §2 are not the same as for Case §3 and Case §4, the slip level is similar, making a qualitative comparison reasonable. The compensatory action of the HWA, shown in Figure 4.5a, around the nominal value, which is approximately 0.14 rad, successfully stabilising the vehicle. At equivalent slip levels, the OS vehicle HWA in Figure 4.5a is considerably lower than the US vehicle HWA in Figure 4.3a, confirming that the model behaves as expected. This follows from the fact that the understeer coefficient – a linear concept which can extended to a nonlinear vehicle operating at constant slip – decreases as the centre of mass is moved towards the rear of the vehicle. A more comprehensive discussion of this standard result in vehicle dynamics can be found in [7]. The Lateral Path Error (LPE), Figure 4.5c is bounded and small but has a higher absolute value than its US counterpart. The difference is attributable to the steering characteristics and the open loop response of the different vehicle configurations [25].

The radius is reduced in Case §4 to assess the performance of the controller for a saturating vehicle. In fact, Figure 4.6a shows that the mean of the HWA is close to zero. The vehicle speed is higher than its critical speed at that operating point, as indicated by the handling diagram in Figure 4.2, which means that the vehicle is beyond its stability limit. The steering activity is characterised by a high frequency and high amplitude response, necessary to avoid instability. The rear slip steady state mean is again significantly larger than the front, as per the nature of the vehicle. The variation in front slip angle is greater than the rear because it is a function of the HWA activity. The LPE, Figure 4.6c, is characterised by a steady state error and oscillates around that value. The constant offset is caused by the vehicle being just above the limit of adhesion, with tyres that cannot provide enough force to keep it on the nominal line and the absence of an integral term in the cost function. Finally, the sideslip angle, 4.6d, is the largest of all the four runs so far analysed, the reason having been outlined in the previous paragraph.

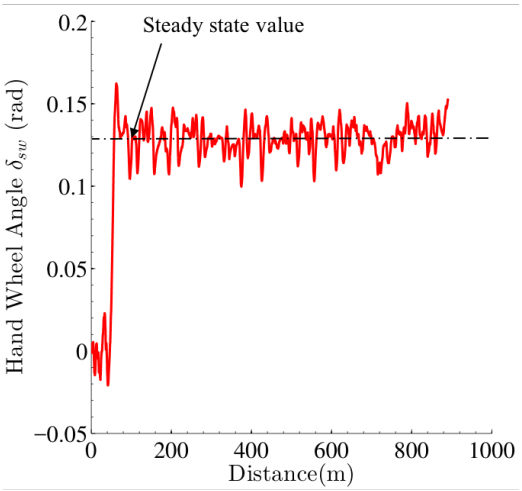

(a) Hand Wheel Angle

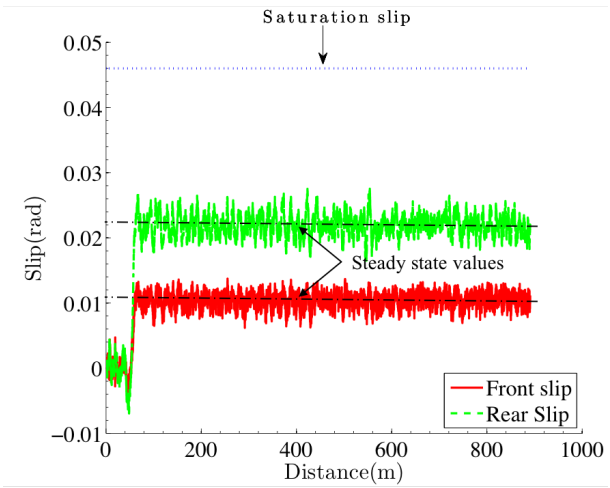

(b) Front and rear slip

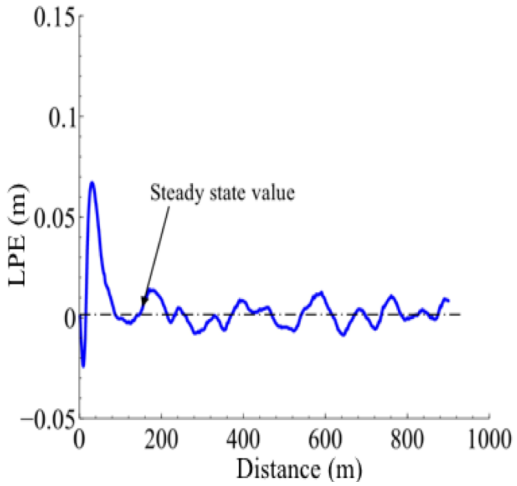

(c) Lateral Path Error

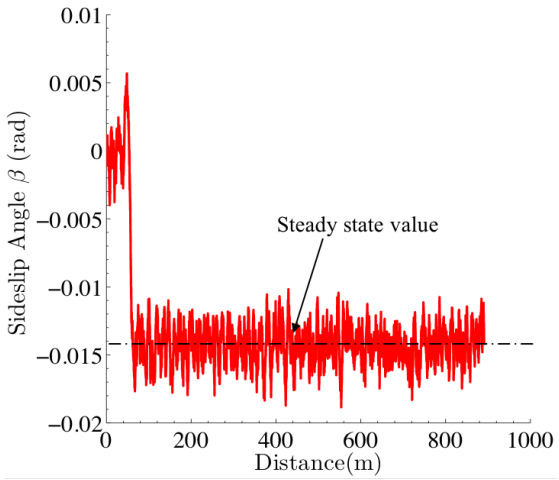

(d) Side-slip angle

Figure 4.5: OS vehicle navigating a 120 m radius circle at 30 m/s under the action of disturbances. Tyres are just out of the linear region but far from saturation. The controller is successful at rejecting disturbances and bounding the states. In order to judge the frequency content, the distance travelled can simply divided by 30 to obtain the corresponding time into the simulation.

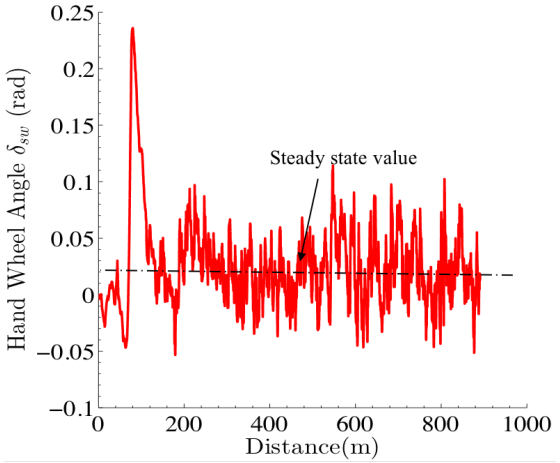

(a) Hand Wheel Angle

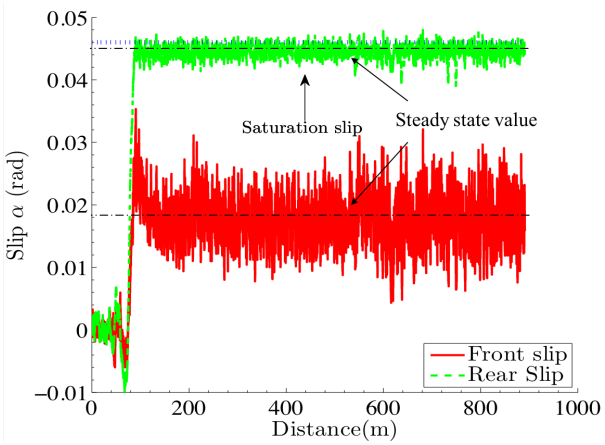

(b) Front and rear slip

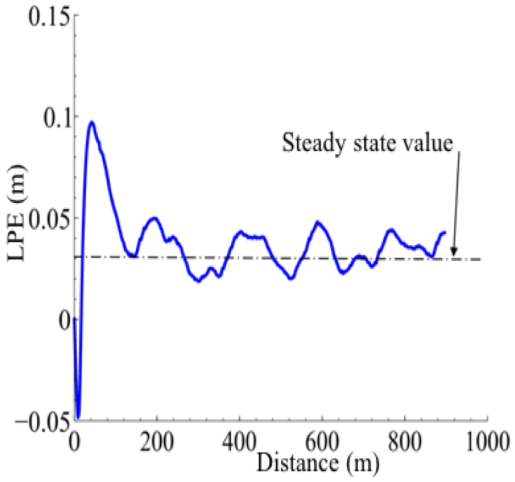

(c) Lateral Path Error

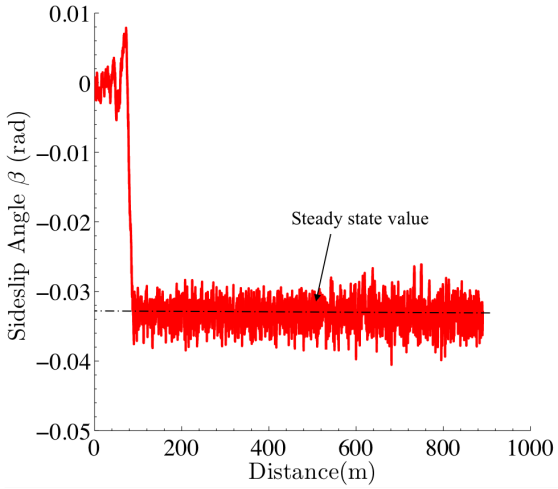

(d) Side-slip angle

Figure 4.6: OS vehicle navigating an 80 m radius circle at 30 m/s under the action of disturbances with tyres saturating. The controller is successful at rejecting disturbances and bounding the states. The vehicle is clearly very difficult to stabilise – in fact significant steady state LPE is present. In order to judge the frequency content, the distance travelled can simply divided by 30 to obtain the corresponding time into the simulation.

4.3.3 Lateral controller parameter study

The controller in its baseline configuration is able to stabilise the vehicle for highly nonlinear manoeuvres and unstable configurations, confirming the robustness of the proposed formulation. The effect of the key parameters can now be considered to gain a better understanding of the operation of the controller and to further confirm that it behaves as expected. A parameter study on all the model parameters would obviously be infeasible given the large number of parameters involved. The effect of changing the centre of mass position has already been explored in Section 4.3.2, establishing that the controller behaves as expected. The focus of this section is to explore driver behaviour and cognitive load. As such, those parameters affecting driver behaviour are considered in most depth. These are:

- The NMS cut-off frequency, which affects the bandwidth of the compensatory response control of the driver. The bandwidth characteristic of the driver influences the closed loop behaviour significantly. [6].
- The penalty on the hand wheel angle speed q_3, which is a primary determinant of the steering action characteristics [105].
- The update time, which determines how frequently the driver performs optimisations to calculate control sequences.

The US configuration described in Section 4.2 is not representative of a real driver; some of the parameters – such as update time – are chosen to ensure the controller can stabilise vehicles at the stability limit. The main objective of this parameter study is to understand the influence of the parameters of the driver model. It is therefore necessary to modify some of the baseline parameter values described in Section 4.2 to achieve a more suitable baseline configuration. Firstly, the update time is set to 0.5 s to represent a real human driver. Secondly, the NMS natural frequency is set to 1 Hz, in the middle of the human-achievable bandwidth. Moreover, tyre stiffness is decreased by 40 % to accentuate the effect of the disturbances and increase the variation in controls as parameters are changed, making the results of the parameter study easier to interpret.

The radius is of the circle is 160 m and the vehicle is travelling at 40 m/s. Figure 4.7 shows the equilibrium position on a handling diagram. Tyres are beyond the linear region but not saturating. Extreme handling conditions are avoided because they would make the effect of each parameter harder to isolate. The combined controller

parameter study – Section 4.4.3 – investigates the effect of saturating tyres in details. The disturbances are described in Section 4.2.3.

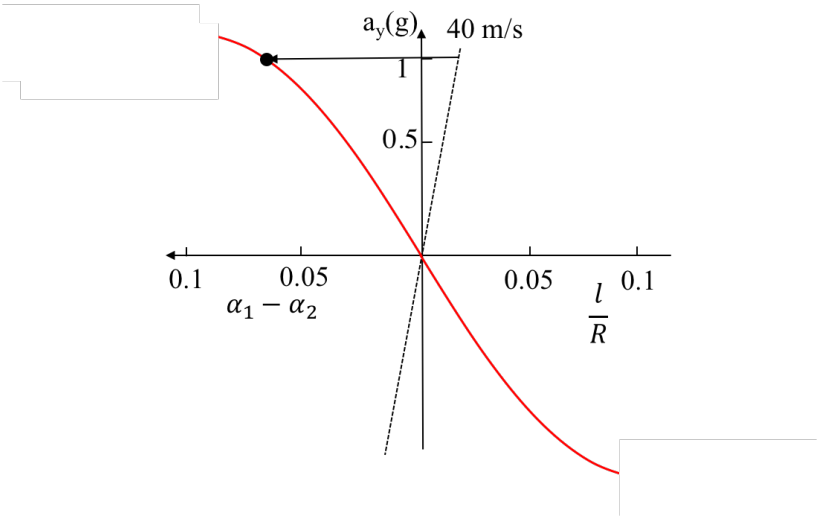

Figure 4.7: Handling diagrams for the parameter study; the speed is 40 m/s and the radius is 160m.

A few terms are clarified. When referring to the Mean Square Spectral Density (MSSD), the roll-off frequency indicates the boundary at which the signal starts being attenuated. The roll off indicates the slope after the roll-off frequency. When referring to the LPE, the time history is considered. An element of the LPE time history is referred to as LPE_k. The LPE standard deviation is defined as

$$\sigma_{LPE} = \sqrt{\frac{\sum_{k=1}^{LPE_{num}} (LPE_k - \mu_{LPE})^2}{LPE_{num}}} \tag{4.3}$$

where LPE_{num} is the number of elements in the LPE time history and μ_{LPE} is the LPE mean. The LPE Root Mean Square (rms) is defined as

$$rms_{LPE} = \sqrt{\frac{\sum_{k=1}^{LPE_{num}} LPE_k^2}{LPE_{num}}} \tag{4.4}$$

The LPE standard deviation is the time history probability density function standard deviation and the LPE rms is the Root Mean Square of the time signal.

The first parameter to be considered is the NMS natural frequency. The baseline configuration is 1 Hz, which is increased to 3 Hz and lowered to 0.2 Hz.

Figure 4.8a shows that, as the NMS natural frequency decreases, the standard deviation of the HWA signal decreases, while the mean remains the same. Figure 4.8b

shows the HWA MSSD. While the roll-off is the same for the three values, the roll-off frequency for the 0.2 Hz case is significantly lower than for 1 and 3 Hz. This explains the decrease in standard deviation 4.8a as the bandwidth of the NMS filter limits the high frequency steering activity.

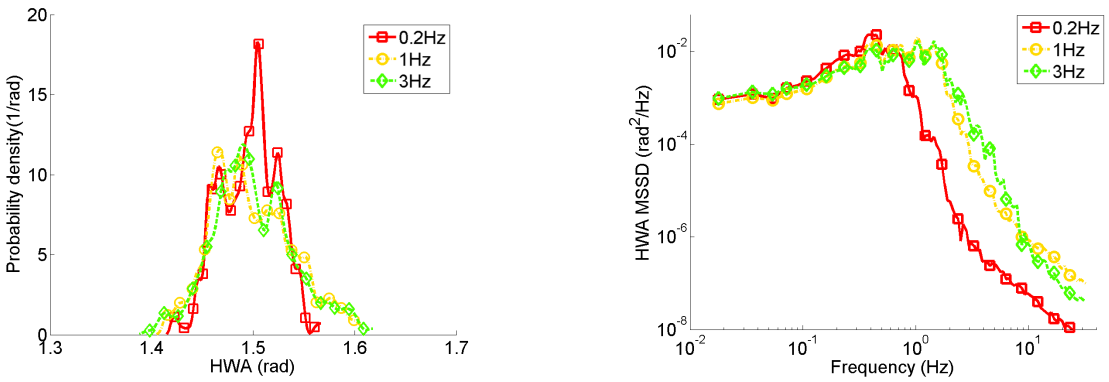

(a) Hand Wheel Angle probability density function

(b) Hand Wheel Angle Mean Square Spectral Density

Figure 4.8: HWA probability density function and HWA Mean Square Spectral Density for a US vehicle navigating a circle at 40 m/s with tyres in the nonlinear region but not saturating at three levels of NMS frequency.

The difference between 1 Hz and 3 Hz is very small compared to the 0.2 Hz result. The reason for this behaviour is due to the penalty q_3, which reduces the frequency content of the HWA signal as shown in Figures 4.9a. Since the HWA MSSD is already limited by the weight on the $\dot{\delta}$, increasing the NMS natural frequency will not have any effect beyond a certain threshold. This result is very important for the validation, as it shows that these two parameters are closely linked and have a similar effect on the frequency response of the driver.

Figure 4.9a shows the significant effect changing q_3 has on the HWA MSSD, affecting both the roll-off frequency and the roll off. The lower q_3, the higher the cutoff frequency and the steeper the roll-off, resulting in improved path tracking capabilities due to the increase in NMS bandwidth. It is also interesting to explore the idea that the penalty is somewhat correlated to the effort the driver puts into driving the vehicle. In order to show this, the LPE Root Mean Square (rms) and HWA speed rms are plotted as q_3 is varied. Figure 4.9b clearly shows that there is inverse correlation between HWA speed and LPE. Increasing q_3, which corresponds to limiting the "effort" [105] provided by the driver, causes the LPE to increase.

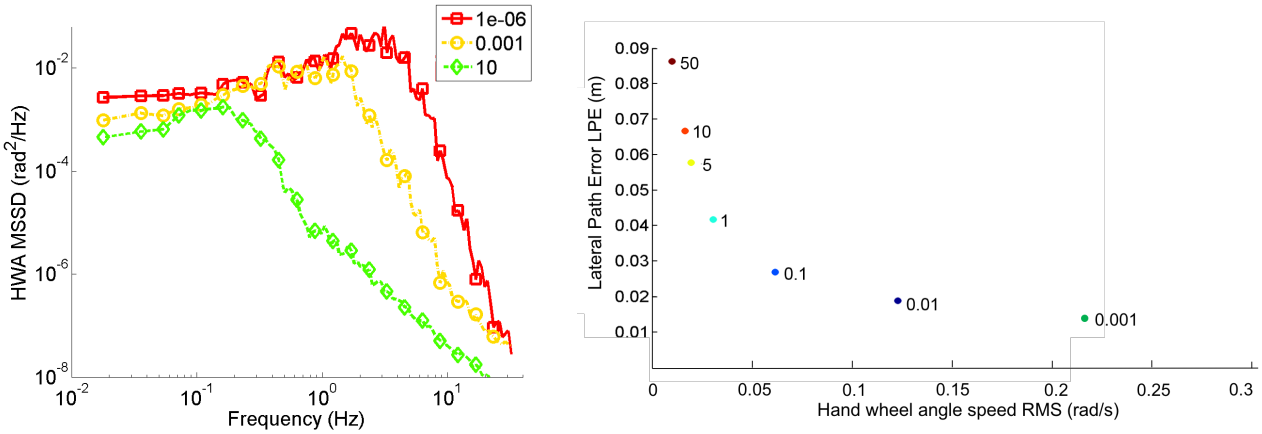

(a) Hand Wheel Angle Mean Square Spectral Density

(b) Tradeoff between HWA speed and LPE

Figure 4.9: HWA MSSD and tradeoff between HWA speed and LPE for a US vehicle navigating a circle at 40 m/s with tyres in the nonlinear region but not saturating for varying q_3.

Finally, the effect of the update time is considered. Figure 4.10a shows that a shorter update time decreases the LPE standard deviation significantly. This is expected, as increasing the update time increases the degree of suboptimality of the solution. The HWA MSSD, Figure 4.10b is not affected significantly, showing that the frequency content is not necessarily correlated to the LPE.

4.4 Combined results

The lateral controller can successfully stabilise a nonlinear vehicle at the limit of handling at constant speed as shown in Section 4.3. Longitudinal dynamics are added to the system to account for speed changes, making the controller more representative of actual racing. The changes to the model and the formulation are detailed in Chapter 3. The aim of this Section is to confirm that the controller can stabilise a vehicle at the limit of adhesion using lateral and longitudinal controls. A particularly challenging manoeuvre, a circle whose radius of curvature is smaller than a physically attainable one, is analysed in Section 4.4.1 to test the effectiveness of the compensatory action of the controller. A 90 degree corner where the nominal trajectory was generated by an MPC controller is considered in Section 4.4.2 to evaluate the response in an actual driving scenario.

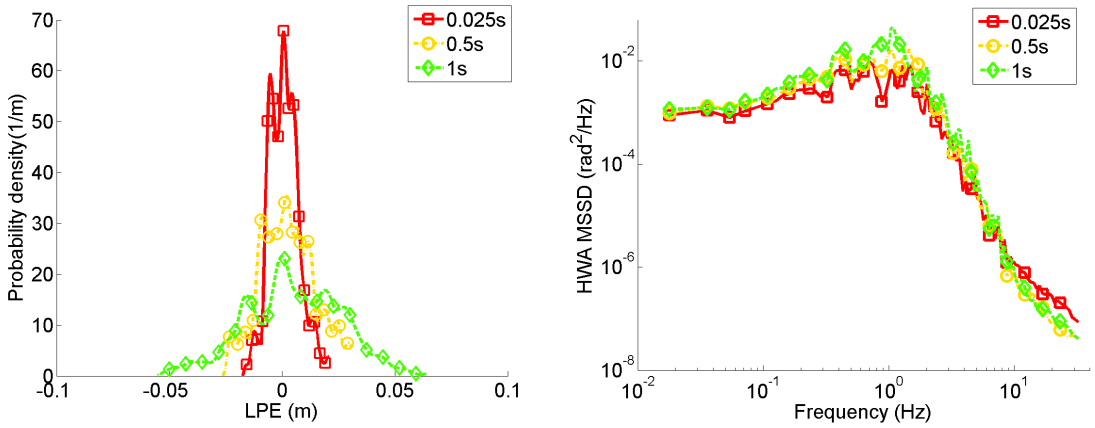

(a) Lateral Path Error probability density function

(b) HWA Mean Square Spectral Density

Figure 4.10: LPE probability density function and HWA Mean Square Spectral Density for a US vehicle navigating a circle at 40 m/s with tyres in the nonlinear region but not saturating for varying update time.

4.4.1 Circle

Various tests in Section 4.3 showed that the lateral controller can stabilise the vehicle as it navigates a circle at constant speed under the action of disturbances, highlighting the differences between OS and US vehicles. The same tests could be repeated for the combined controller, however, a more challenging manoeuvre is chosen to better highlight the characteristics of the controller.

The vehicle navigates target circle at constant speed whose radius is smaller than the physically achievable one for the given tyre characteristics, speed and configuration. For this particular test, the tyre parameters are given in Table 4.4, the speed is set to 30 m/s, the vehicle is a US one and the radius is set at 55 m. The vehicle is subjected to the lateral and moment disturbances described in Table 4.5 and Section 4.2.3. Longitudinal disturbances are not considered so as to isolate the effect of the lateral disturbances and gain an understanding of how they can affect longitudinal controls. The constraint on the lateral and longitudinal slip is removed to better explore the behaviour of the controller. The parameter study in Section 4.4.3 considers the effect of longitudinal disturbances.

If the lateral controller was tested under the same conditions, the vehicle would find its equilibrium on a circle whose radius is larger than the target one, at constant speed. Evidence of this behaviour was present in Case §4 of Section 4.3.2, where the

vehicle was right at the limit of adhesion and the LPE showed a significant steady state error. Even though no integral action was present, the steady state error was significantly larger than the other cases so it is attributable to saturating tyres. The test is not repeated for the lateral controller because, since speed would be constant, comparing vehicle states would not be insightful.

This set-up forces the tyre to be in the highly nonlinear region of the curve. Clearly, a trade-off between speed error and LPE is necessary because the tyres cannot provide enough force to match the target path and speed. The nominal solution, without disturbances, is superimposed on all Figures 4.11, indicating that the first section of the time history is the transient response arising from the initial conditions. Steady state is achieved in the second section, where the disturbed signals show the response to random disturbances.

Figure 4.11a shows the HWA. The standard deviation of the signal is very significant because of the highly nonlinear nature of the tyre curve at the operating point. The HWA speed $\dot{\delta}_{sw}$ would be impossible to achieve for a human driver. The literature, for instance [132], shows that the maximum achievable $\dot{\delta}$ is about 10 rad/s while the simulated $\dot{\delta}_{sw}$ reaches values of 15 to 20 rad/s multiple times. The graph is not shown because it would not add any further relevant information. The aim of the test is to show that the controller can stabilise the vehicle at the very limit of handling, which the extreme values demonstrate.

Figure 4.11b shows the LPE that, again, is far from zero, as the vehicle cannot physically follow the nominal path. In absence of disturbances the lateral path error would settle to a constant value such that the vehicle would end up on a concentric circle with a larger radius. Due to the action of random disturbances, the vehicle oscillates around the nominal solution. Figure 4.12 shows the path of the vehicle. After the initial transient, the vehicle settles around a circle which has the same centre as the reference and a larger radius.

Figure 4.11c shows the speed of the vehicle. The average speed is lower than the target as the radius of the circle is too small for the vehicle to match both the target speed and path.

The LPE and Speed Error relative magnitude depends on the cost function weights, as there is an inevitable trade off between the two: tyre load can be decreased by decreasing the velocity or increasing the LPE. The cost function weights q_1 and q_4 determine the nature of the trade-off.

Figure 4.11d shows the longitudinal slip of the front and rear tyres of the vehicle. The front slip is zero for most of the manoeuvre as the vehicle is rear wheel drive;

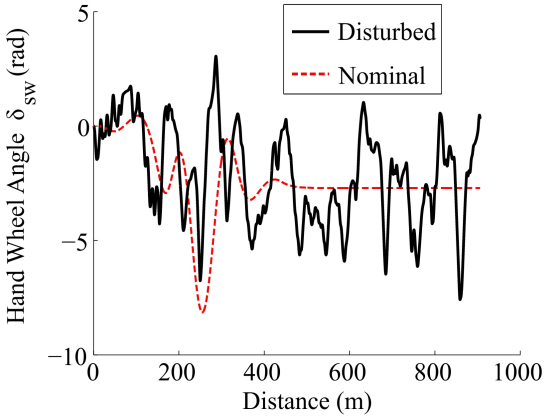

(a) Hand Wheel Angle

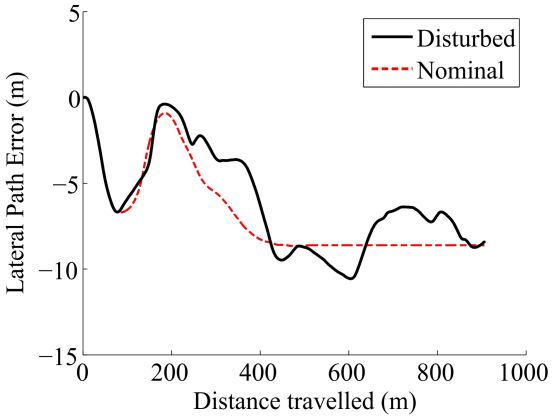

(b) Lateral Path Error

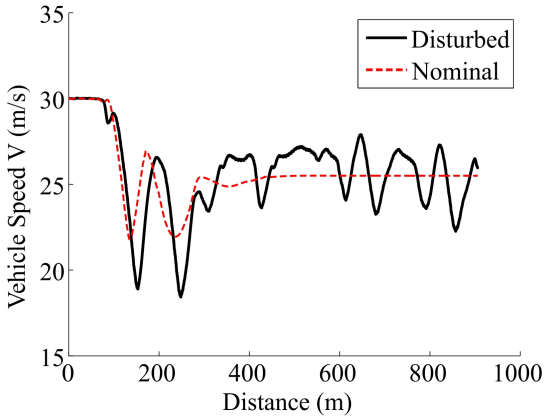

(c) Vehicle Speed

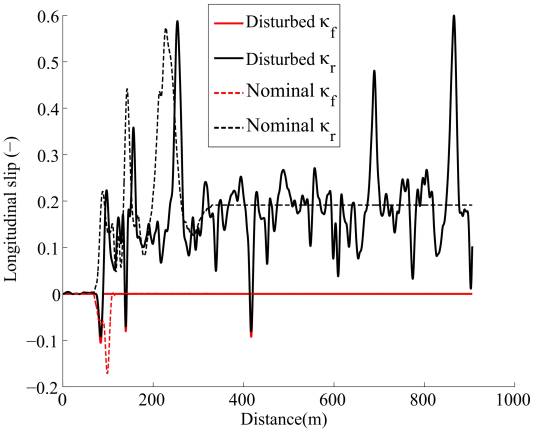

(d) Front and rear longitudinal slip

Figure 4.11: Vehicle navigating a 55 m radius circle at 30 m/s under the action of disturbances. The radius of the circle is too small for the vehicle to track the trajectory, so there is a steady state error in both speed and lateral path error. Even though there are instances where both tyres are on the saturation limit, the controller stabilises it successfully.

significantly negative values are observed in 3 instances where the controller finds necessary to apply braking torque at the front. This can be surprising because the vehicle is below its target speed. However, the nominal speed and radius the controller is tracking are different from the required ones due to them being beyond the vehicle physical limits. The speed therefore oscillates around a lower value, around 25 m/s. Even though the overall longitudinal torque is positive to overcome the longitudinal component of the front tyre, the magnitude of the oscillations is large enough for the controller to require a small braking torque in some instances. The rear slip average value is positive because it needs to provide the torque to overcome air resistance and overcome the front wheel longitudinal force component, which is slowing the vehicle down. Since the front wheel steer angle values are large, this component is significant.

Figures 4.13a and 4.13b show that both the front and rear slip are well beyond the saturation limit while the vehicle is still stable. This shows that the controller is successful at stabilising the vehicle when the tyres are beyond saturation. Figure 4.14 shows the front and rear tyre saturation as a percentage of the maximum combined slip throughout the manoeuvre. The initial transient causes both the front and the rear to go well beyond saturation. Front and rear slip decrease in the subsequent part of the manoeuvre.

Figure 4.15 shows the slips time histories with reduced vertical range. While the front slip stays close to 100 %, sometimes lower, the rear is beyond saturation for most of the manoeuvre. While the longitudinal action is needed to overcome the

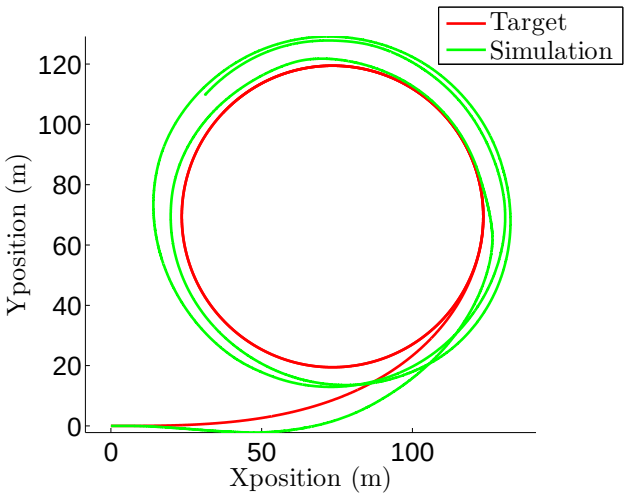

Figure 4.12: Vehicle path for a vehicle navigating a 55 m radius circle at 30 m/s under the action of disturbances. The radius of the circle is too small for the vehicle to track the trajectory, so a constant LPE is present.

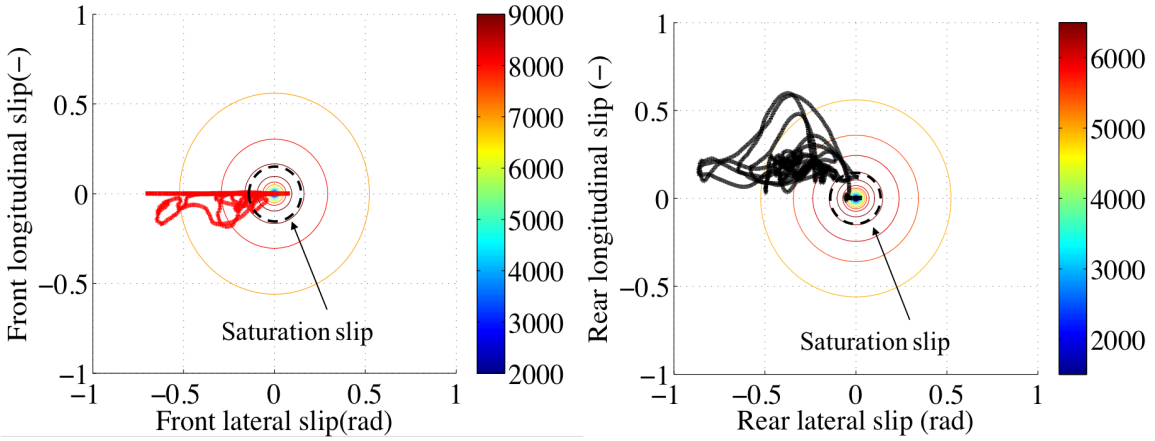

(a) Front tyre slip (b) Rear tyre slip

Figure 4.13: Front and rear slip for a vehicle navigating a 55 m radius circle at 30 m/s under the action of disturbances. The radius of the circle is too small for the vehicle to track the trajectory, so tyres saturate. The colorbar indicates the tyre force (N)

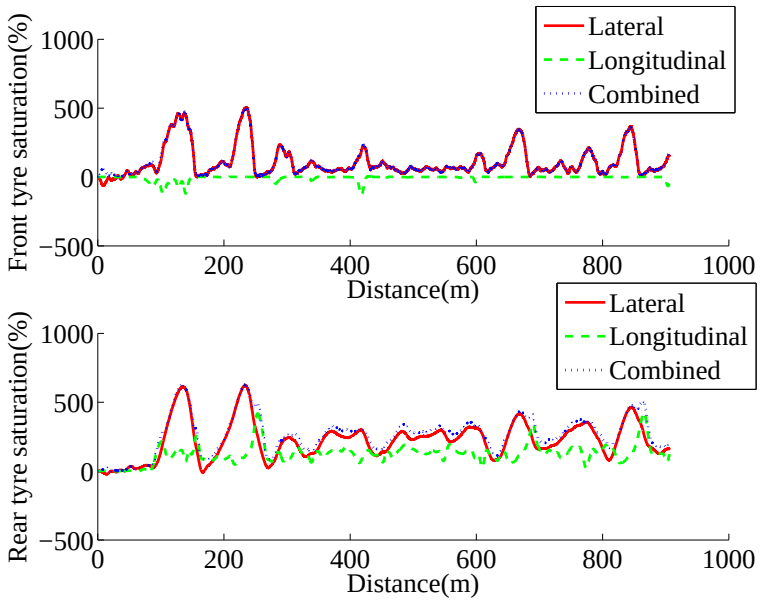

Figure 4.14: Slip time history showing that both front and rear tyres saturate for most of the manoeuvre.

front tyre longitudinal force, the large rear slip level also suggests that the controller accelerates to cause the rear tyre to saturate, effectively increasing its compliance and

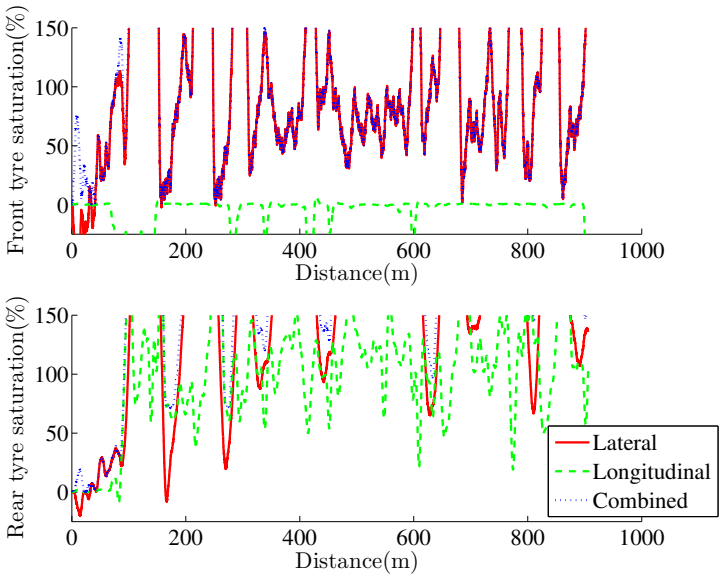

Figure 4.15: Slip time history with reduced vertical range to show the slip around saturation. The rear saturates far more than the front.

making the vehicle easier to steer. This behaviour is difficult to isolate in this test; the lateral/longitudinal coupling is investigated in more depth in Section 4.4.3.

4.4.2 Comparison with MPC controller

Nominal path and speed profiles used so far lead to steady state solutions in absence of disturbances. While they are useful to gain a better understanding into the behaviour of the controller, a more complex manoeuvre needs to be considered to assess the behaviour of the controller in a more realistic scenario where speed varies. The 90° bend used in Chapter 2, Figure 2.3, is considered. The path and speed trajectories that minimise manoeuvring time are found using an efficient MPC framework [133] and are used as target trajectories for the NMPC controller. They are described in Section 2.3.

The first set of result examines how well the nonlinear controller tracks the nominal path and speed in absence of disturbances. The driver and vehicle model used in [133] is essentially equivalent to the one used for the nonlinear controller; there is however one subtle difference worth mentioning. The wheel motion is neglected in the nonlinear controller but considered in the MPC. The maximum absolute value of the acceleration in the target speed trajectory happens in braking and it is about 1g, which, using the wheel parameters in [133], leads to a total wheel inertial torque of 20 Nm, which is around 1.5 % of the total torque. The effect is therefore negligible. The driver and

vehicle parameter values used are the same as [133] to ensure results are comparable. The cost function weights are those described in Section 4.2.2.

Given that the vehicle is at its handling limit, the task is challenging and it also serves as a preliminary validation. Figure 4.16a shows the MPC and the nonlinear controller HWA. The two driving strategies are very similar, with the nonlinear controller putting in slightly less steering at the apex. Also, fewer oscillations are present at the beginning of the corner and the manoeuvre is generally smoother. The different approach is due to the difference in formulation. Firstly, the MPC linearises the system at every time step whereas the nonlinear controller uses a direct collocation method. Secondly, the NMPC cost function, which is described in Section 3.4.3, is different. While the MPC minimises manoeuvre time, the nonlinear formulation tracks optimal states, effectively using the MPC controller output as the target.

Figure 4.16b shows that the LPE is zero at the beginning of the manoeuvre and in the order of 10^{-2} as the vehicle negotiates the corner. Figure 4.16c shows a good match of the speed profiles with the vehicle accelerating, braking into the corner, hitting the apex at constant speed and then accelerating to exit the corner. A closer look at the difference in speed, Figure 4.16d, shows that the error is around 0.1 % for corner entry to increase to about 1 % at corner exits, where tyres are close to saturation and both lateral and longitudinal forces are needed. Figures 4.17a and 4.17b show that the whole of the tyre curve is utilised, both laterally and longitudinally. Both tyres saturate in both the lateral and the longitudinal direction, and the controller is able to stabilise the vehicle in both braking and turning. The MPC controller, despite the difference in δ_{sw}, exploits the tyre force in the same way.

After establishing that the controller can successfully track the optimal signals, external disturbances are added to confirm that the stabilising action is not compromised by random external forces. Disturbances are added to all the three degrees of freedom of the vehicle and they are described in Section 4.2.3 and Table 4.5. The controller manages to track the reference states as Figure 4.18 shows. A full set of results is not shown here because it would be redundant and would not add any new or interesting insights. Figure 4.18a shows that the LPE is higher in the middle of the corner, which is where the tyres saturate laterally. Figure 4.17b shows that the maximum errors in speed occur at corner entry and corner exit; that is at the highest moment of longitudinal activity. For the disturbed system, the speed error, Figure 4.18b is on average 30 % higher than the speed error for the undisturbed case, shown in Figure 4.16d, in the first section of the manoeuvre. At entry and exit, it is three times and twice as large respectively.

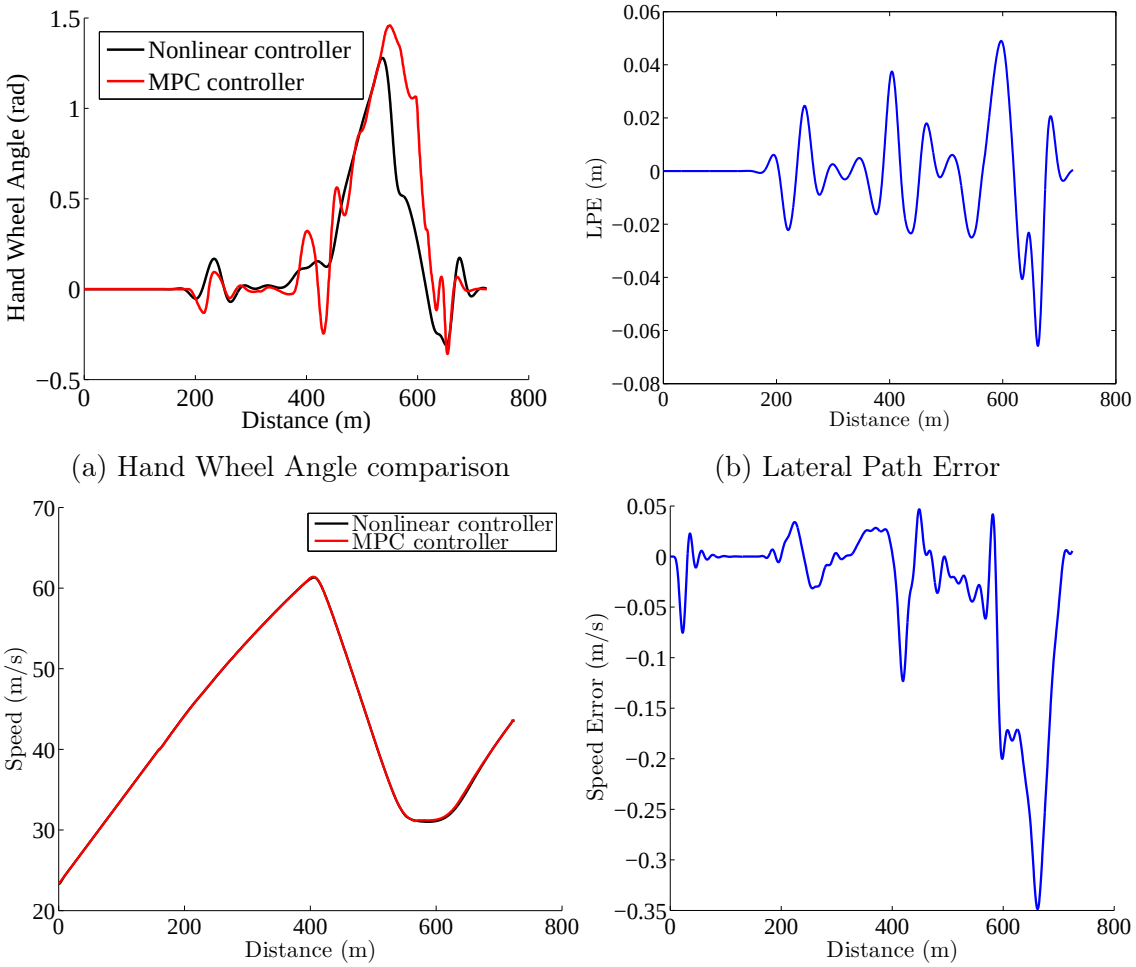

(a) Hand Wheel Angle comparison

(b) Lateral Path Error

(c) Vehicle Speed

(d) Speed error relative to nominal speed

Figure 4.16: Comparison between the MPC and the nonlinear controller for the 90 degrees corner shown in Figure 2.3.

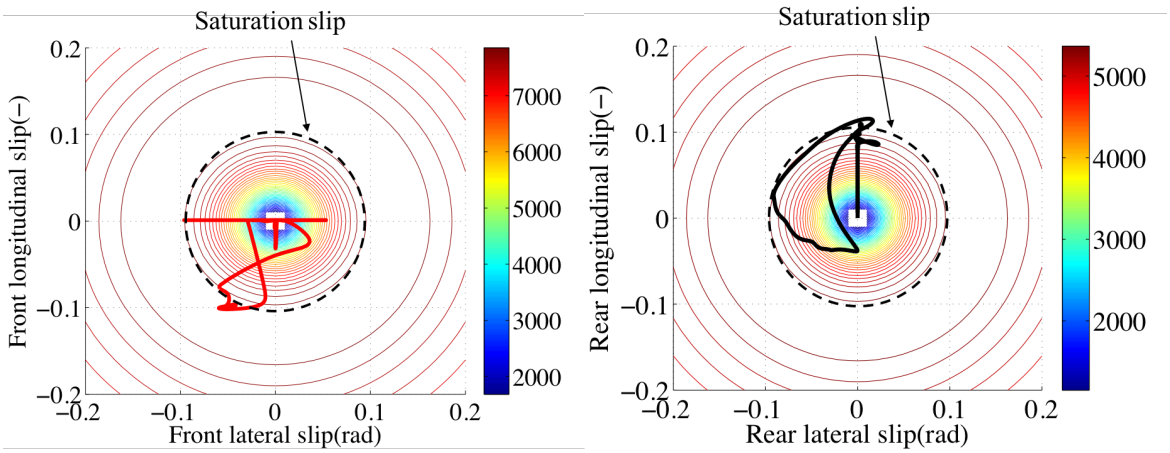

(a) Front slip (b) Rear slip

Figure 4.17: Nonlinear controller slips for the 90 degrees corner shown in Figure 2.3. Both front and rear tyres are saturating. The colorbar indicates the tyre force (N).

Disturbed results could be compared to the results from the LQR controller described in Chapter 2. One of the main research objectives of the thesis is improving the linear controller as it is deemed unsatisfactory. The comparison could be done in a number of ways, such as comparing time histories or statistics. However, drawing quantitative, fact-based conclusions on which approach is better would only be speculative work because of the lack of an objective, experimental benchmark. Observing that both the LPE and the speed error increase as tyres saturate is consistent with the findings in Chapter 2, for instance in Figure 2.10. More specific conclusions are difficult to obtain. The fact that the NMPC manages to stabilise the vehicle as tyres saturate confirms that the research objective is met, as the LQR would be unstable for the negative slope region section of the tyre curve. An in-depth assessment of the controller accuracy is performed in Chapters 5 and Chapter 6, where an experimental benchmark is available. The LQR is not tested as it would not ensure stability for the tests performed to obtain the experimental data.

4.4.3 Combined controller parameter study

The parameter study described in Section 4.3.3 is extended to the longitudinal case. The main objective is to investigate the cross-coupling between steering control and speed errors, and between wheel torque control and lateral path errors. Two scenarios are considered to meet these objectives:

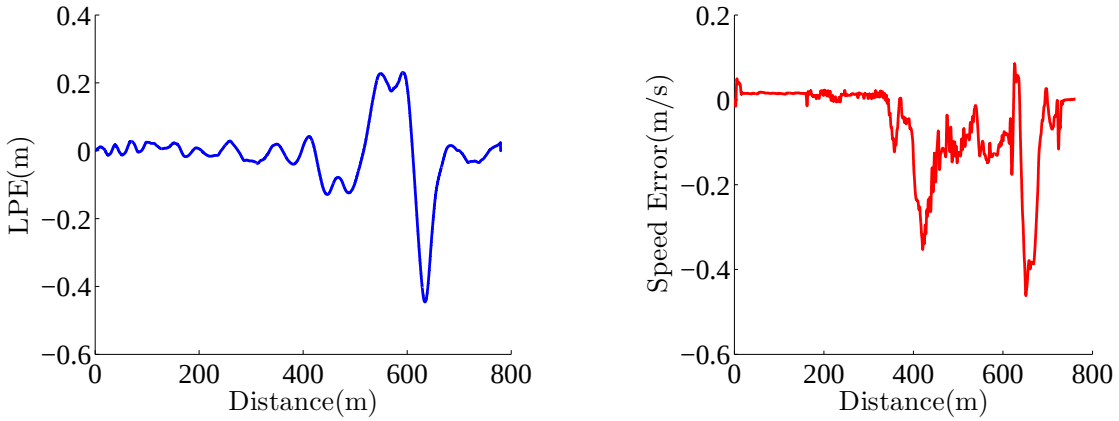

(a) Lateral path error

(b) Speed error

Figure 4.18: Lateral Path Error and Speed error for the 90 degrees corner shown in Figure 2.3 with nominal disturbances acting on the vehicle.

- An pulse in lateral force, longitudinal force or yaw moment is applied to the vehicle at the centre of mass while the vehicle navigating a circle.
- Continuous random disturbances are applied to the vehicle.

Pulse disturbance

The control actions required to recover the vehicle from a disturbed set of states are particularly sensitive to model parameters, making the scenario ideal for a parameter study. These conditions are replicated by applying an pulse disturbance to the baseline US vehicle navigating a circle at a constant speed of 30 m/s. The radius is changed to achieve different slip conditions. Pulse disturbances are constant external forces or moments acting on the vehicle over a single time step, essentially a Dirac delta function shifted in time to allow the vehicle to achieve steady state, as shown in Figure 4.19. Pulse disturbances are introduced on all three degrees of freedom.

The lateral pulse disturbance is a 1230 N force perpendicular to the centre of mass pointing to the outside of the circle acting after the vehicle has travelled 200m along its target path. The pulse magnitude has been chosen to increase the slip enough to cause a significant corrective action, without causing the vehicle to spin. The value has been found iteratively by trial and error. The yaw disturbance is a 960 Nm moment acting at the centre of mass 200 m into the simulation. Finally, the longitudinal disturbance is a 2240 N force acting longitudinally at the same point in distance.

Disturbances are not active for all tests. Different configurations in terms of disturbances, variable parameters and constraints are used to meet the parameter

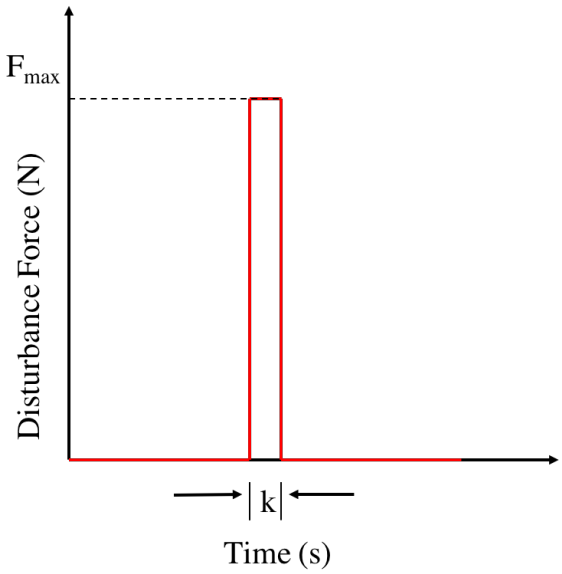

Figure 4.19: Diagram showing the nature of the pulse disturbance. The force at the beginning of the simulation is zero and then increases to F_{max} over a single time-step k to then return to zero. The moment disturbance takes the same form.

study objectives. Table 4.7 shows the four Cases considered. The radius determines the level of nonlinearity: mild for a 120 m radius and high for a 70 m radius. The parameters that are varied are the cost function weights on the lateral and longitudinal control activity. These are:

- The cost function weight q_1 determines the penalty on the deviation from the nominal trajectory.
- The cost function weight q_3 determines the penalty on $\dot{\delta}$ magnitude.
- The cost function weight q_4 determines the penalty on the deviation from the nominal speed.
- The cost function weight q_6 determines the penalty on $\dot{\kappa}_r$ magnitude.

Each Case is run for three values of the study parameter, namely, 1E-7, 0.025 and 1, unless otherwise stated. Longitudinal disturbances denote a longitudinal force only, while lateral disturbances denote a lateral force and a moment. The front longitudinal slip is constrained to 0 in all cases to isolate the rear action. Each case focus on a different aspect of the controller:

- Case 1 focuses on longitudinal controls. There are only longitudinal disturbances and lateral controls are constrained so the longitudinal action is isolated.
- Case 2 focuses on lateral controls. There are only lateral disturbances and longitudinal controls are constrained so the lateral action is isolated.
- Cases 3 and 4 focus on the lateral/longitudinal interaction. Only lateral disturbances are present but both lateral and longitudinal controls are active.

Table 4.7: Test cases for parameter study with pulse disturbances. The nonlinearity level is adjusted by changing the radius. The weights on the LPE and u, i.e. speed, indicate which state is penalised.

Case	Radius	Disturbances	Parameter	Constraints	q_1	q_4
1	120	Longitudinal	q_6	$\dot{\delta}$ & $\dot{\kappa}_f$	1E-3	10
2	120	Lateral and yaw	q_3	$\dot{\kappa}_r$ & $\dot{\kappa}_f$	10	1E-3
3	120	Lateral and yaw	q_3	$\dot{\kappa}_f$	10	1E-3
4	70	Lateral and yaw	q_3	$\dot{\kappa}_f$	10	1E-3

In the following analysis, the term heading angle refers to the angle ξ, the angle between the vehicle longitudinal axis and the target path described in Section 3.3.2.

In Case 1 the vehicle navigates the circle and a forward pulse longitudinal force is applied at its centre of mass. The simulation is repeated for three values of q_6, the penalty on $\dot{\kappa}$, 1E-7, 0.025 and 1. The lateral controls are not active and they do not intervene, making the case straightforward to interpret. Figure 4.20a shows that, as q_6 increases, the longitudinal action decreases and the speed error increases , Figure 4.20b. This confirms that the longitudinal action and speed error are inversely proportional.

Case 2 focuses on confirming that lateral control action and LPE are inversely proportional. The vehicle is disturbed laterally, q_1 is high,10, and the test is repeated for three values of q_3, 1E-7, 0.025, 1. The longitudinal rear slip is constrained to zero to isolate the lateral dynamics. Figure 4.21a shows that as q_3 increases the control action decreases. Figure 4.21b confirms that the less prompt the control action causes an increase in the LPE. Figure 4.21c shows that the vehicle heading oscillates if the control action is not prompt, causing a larger path error. It is also interesting to note that the LPE for $q_3 = 1$ is much larger than the other two cases, while the heading is very similar. This shows that the LPE is more sensitive to the lateral force disturbance than the yaw moment disturbance.

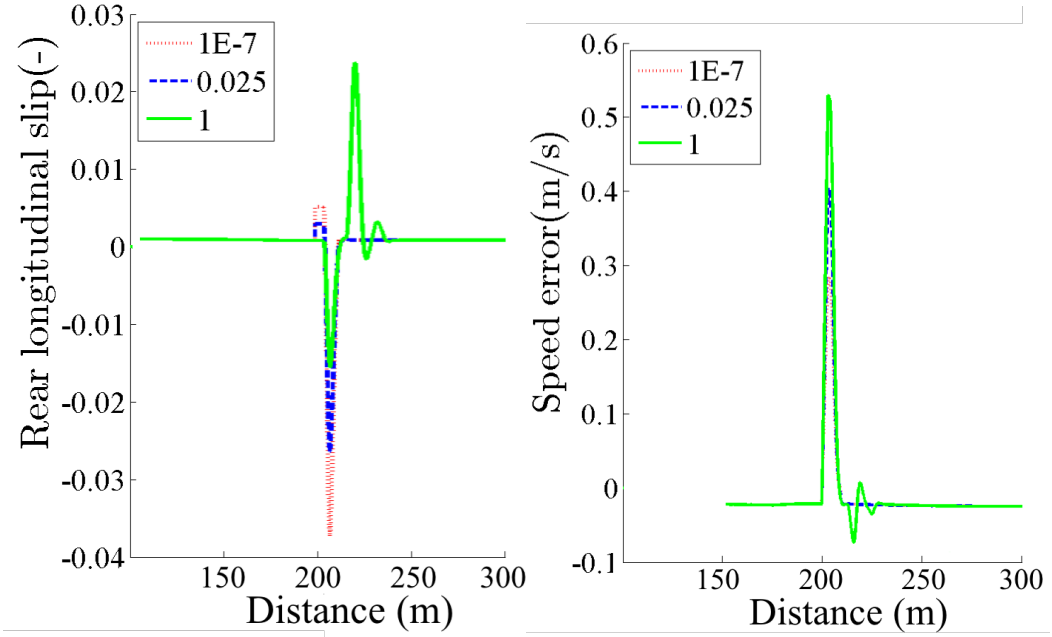

(a) Rear Longitudinal Slip (b) Speed error

Figure 4.20: Rear longitudinal slip and speed error for Case 1 showing the effect of changing the longitudinal controls' weight. The legend indicates q_6 values.

Case 3 is characterised by the same conditions as Case 2 but the longitudinal rear slip is not constrained, with the aim of exploring the coupling between lateral and longitudinal dynamics.

The HWA for Case 3, shown in Figure 4.22a, is the same as Case 2 for $q_3 = 1E-7$ and $q_3 = 0.025$. Only the corrective action for $q_3 = 1$ is smaller, with a peak to peak value of 0.1 rad vs 0.15 rad. Since all conditions, but the longitudinal action, are the same, this difference confirms that the lateral controls are affected by the longitudinal action of the controller. The LPE, Figure 4.22b, follows a similar pattern: the response for $q_3 = 1E-7$ and $q_3 = 0.025$ is the same as Case 2 but the maximum amplitude of the deviation for $q_3 = 1$ is smaller, -0.08 m vs -0.12 m, indicating the mitigating effect of the longitudinal action. The heading error, shown in Figure 4.22c, is also reduced by about 30 % for $q_3 = 1$. Figure 4.22d shows the HWA speed, which decreases as the penalty increases.

Figure 4.23a shows the longitudinal action of the controller and the vehicle speed error. Longitudinal controls intervene when the HWA action is not prompt. For the first two values of q_3, 1E-7 and 0.025, rear controls do not intervene as the LPE is

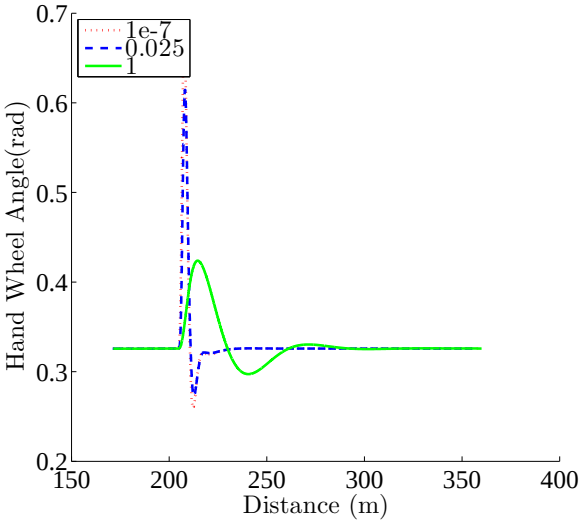

(a) HWA

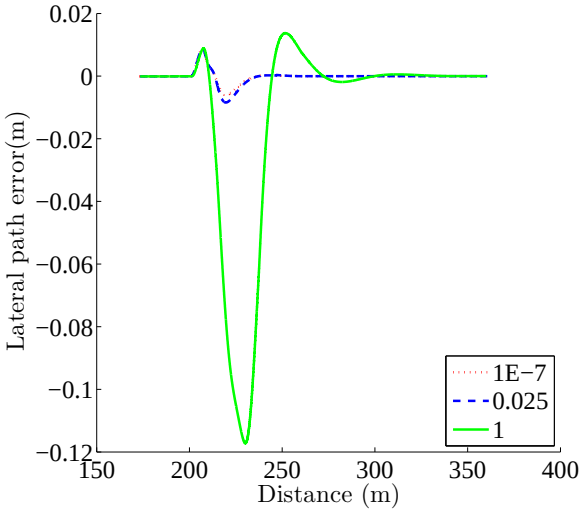

(b) Lateral Path Error

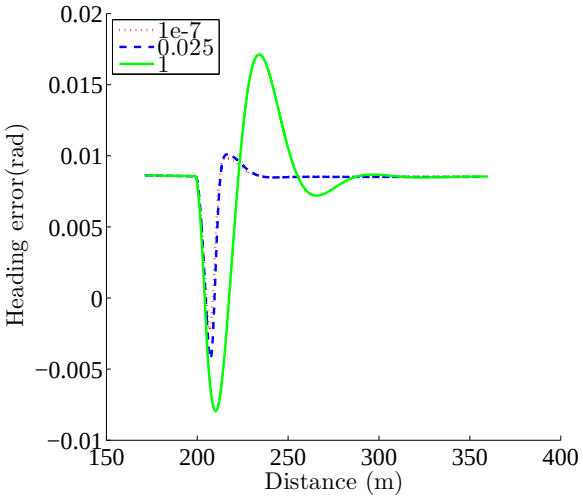

(c) Heading Error

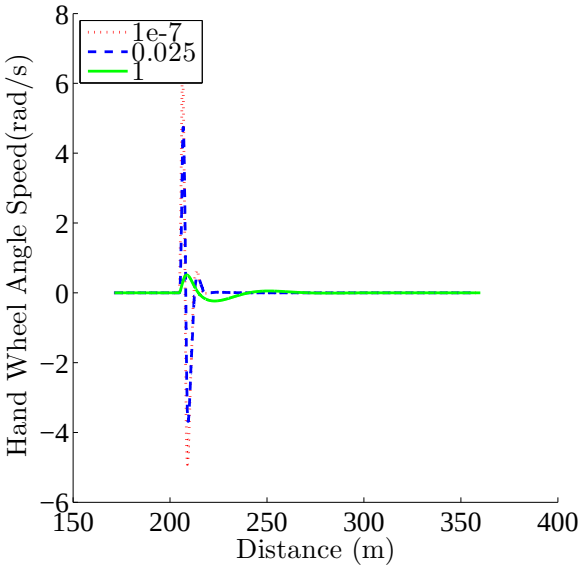

(d) HWA speed

Figure 4.21: Key states and controls for Case 2 showing the effect of changing the lateral controls' weight. The legend indicates q_3 values.

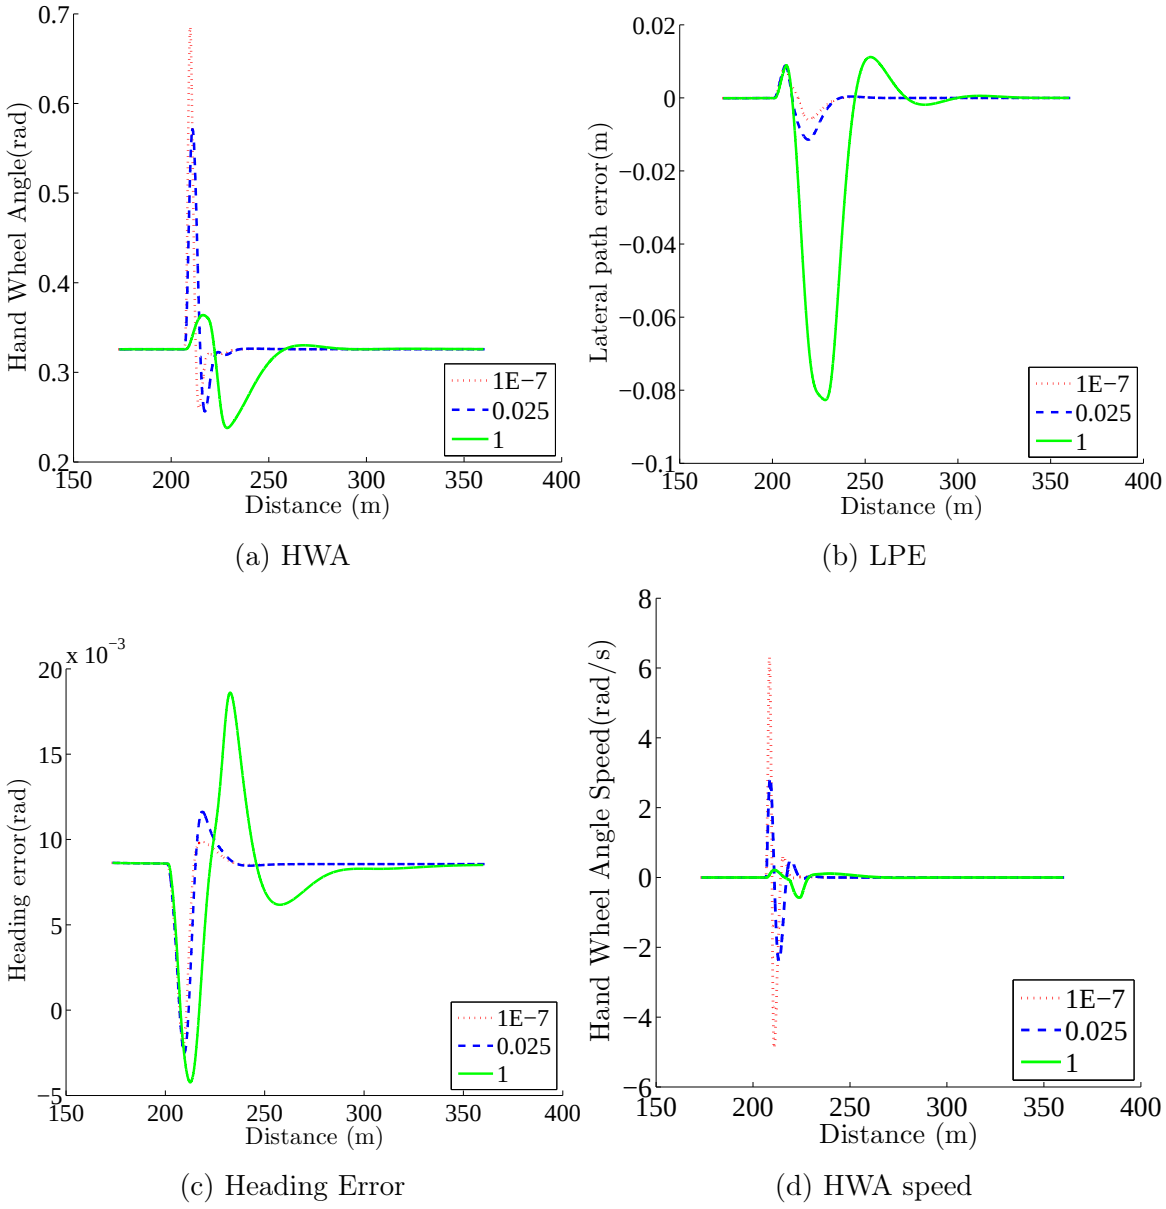

(a) HWA (b) LPE

(c) Heading Error (d) HWA speed

Figure 4.22: Key states and controls for Case 3 showing the effect of changing the lateral controls' weight and the interaction between lateral and longitudinal controls. The legend indicates q_3 values.

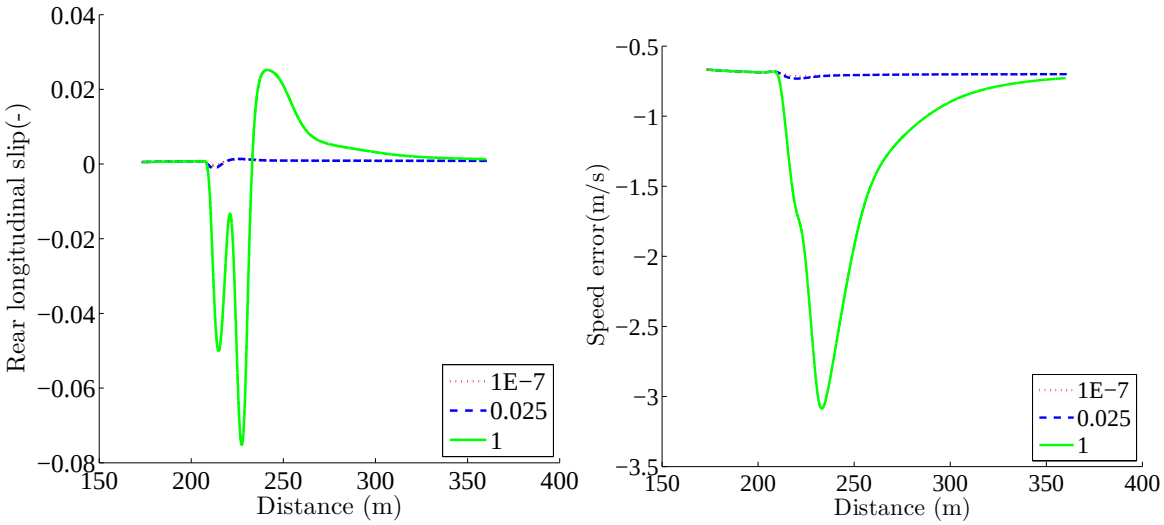

(a) Rear longitudinal slip (b) Speed error

Figure 4.23: Rear longitudinal slip and Speed error for Case 3 showing the effect of changing the lateral control weights and the interaction between lateral and longitudinal controls. The legend indicates q_3 values.

small but, as the HWA action becomes less responsive, negative slip is generated. The braking action reduces the cornering stiffness of the rear tyre, increasing the turning moment generated by the front force. The controller accelerates later to compensate for the velocity error. Figure 4.23b shows that the speed error is small for the first two cases because the longitudinal action is not significant, while it is significant in the third case. The improvement in path tracking comes at the expense of a large speed error; this is the result of the high q_1/q_4 ratio.

Case 4 replicates the conditions in Case 3, but at the limit of adhesion with the front tyre saturating. This allows assessment of how the increase in nonlinearity affects the rear longitudinal action. Figure 4.24a shows that the HWA decreases as the penalty on q_3 is increased. The control action variation is more marked than Case 3, because of the low tyre stiffness near saturation. Thus, a large change in slip, and hence HWA, is required to correct the lateral force. Figure 4.24b shows that the LPE increases as q_3 increases. The LPE is significantly larger than in Case 3 because of the limited force tyres can provide for a corrective action when close to saturation. Saturating tyres have a significant effect on the heading error as well; Figure 4.24c in fact shows that it is not very sensitive to changes in lateral control action. HWA speed follows the expected pattern, decreasing as q_3 increases, as Figure 4.24d shows.

Figure 4.25a shows that longitudinal controls intervene in all cases because the LPE is large. This results in the speed error shown in Figure 4.25b but more importantly

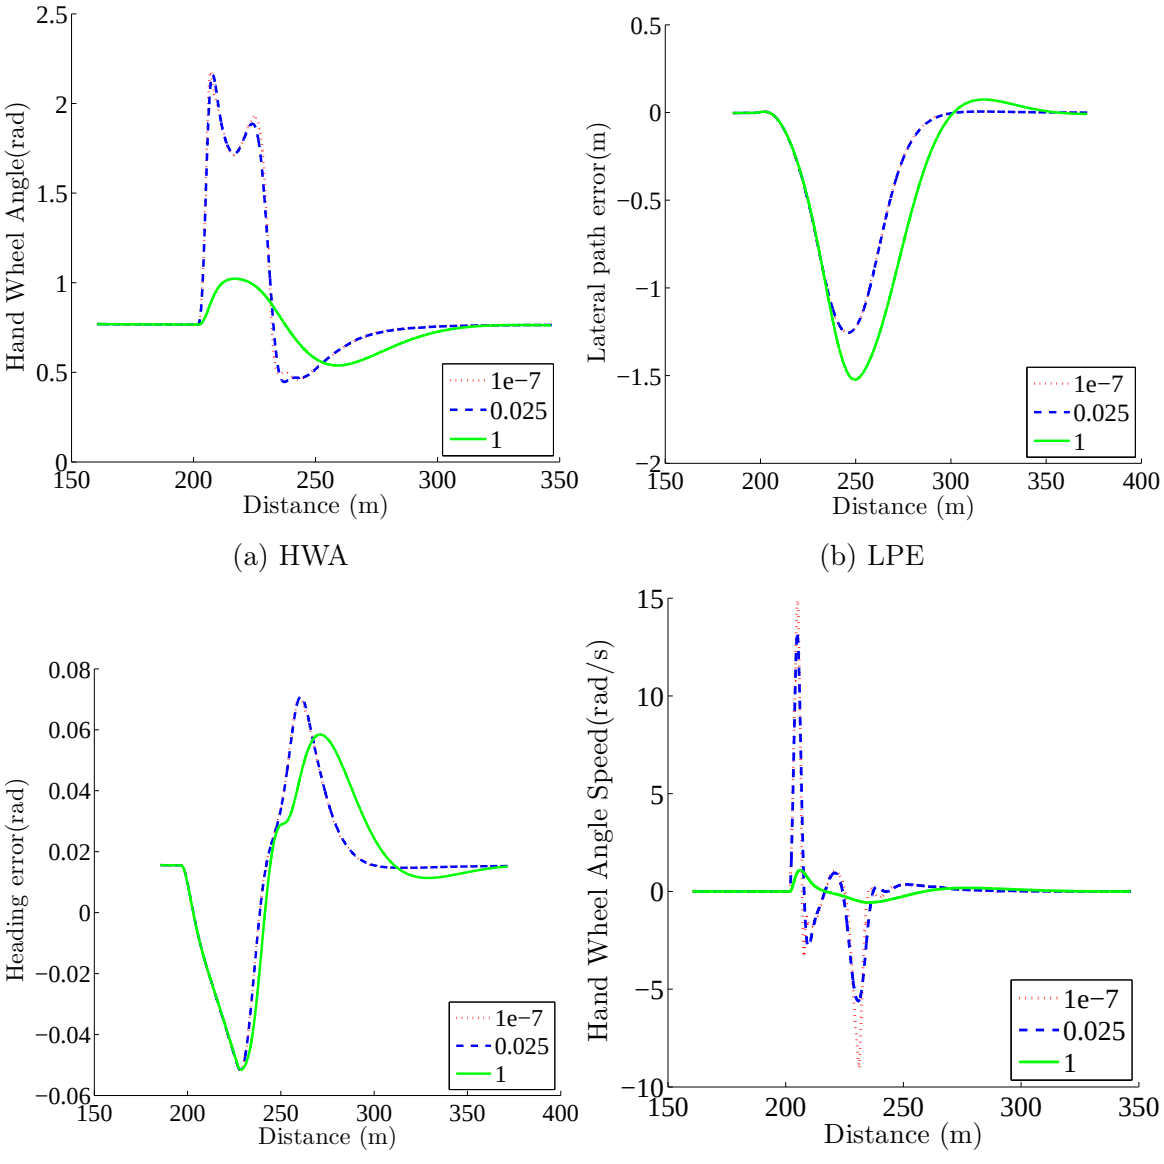

(a) HWA (b) LPE

(c) Heading Error (d) HWA speed

Figure 4.24: Key states and controls for Case 4, showing the effect of changing the lateral controls weight and the interaction between lateral and longitudinal controls when the vehicle is saturating.

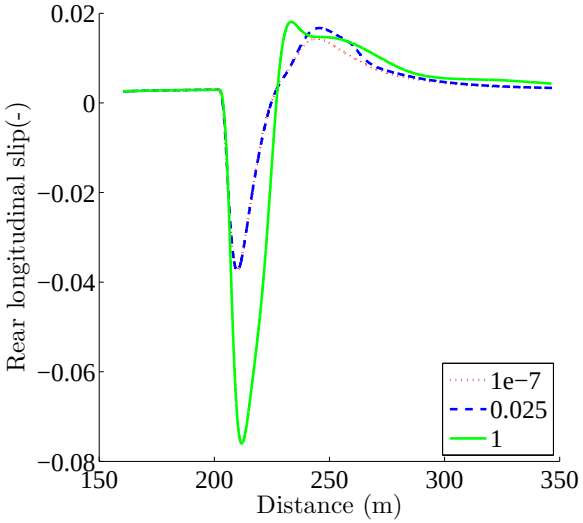

(a) Rear longitudinal slip

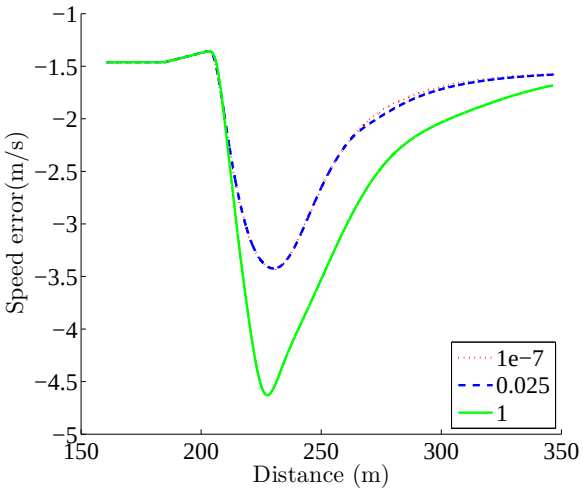

(b) Speed error

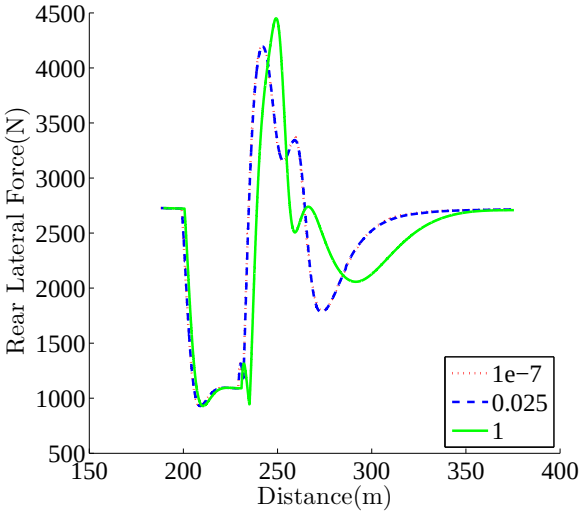

(c) Rear lateral force

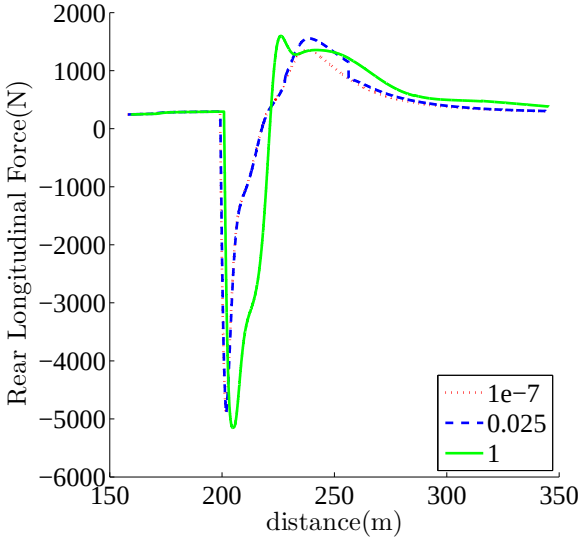

(d) Rear longitudinal force

Figure 4.25: Rear longitudinal slip and speed error for Case 4 showing the effect of changing the lateral control weights and the interaction between the lateral and longitudinal controls when the vehicle is saturating. Rear lateral and longitudinal forces are also included showing how the controller increses the longitudinal slip to decrease the cornering stiffness of the tyres. The legend indicates q_3 values.

in a decrease of the rear lateral force, shown in Figures 4.25c, again increasing the turning moment of the front tyre. The longitudinal slip, Figure 4.25d, also decreases significantly as a result of the braking action. Since the aim of the controller is decreasing the rear tyre lateral stiffness, it could either increase or reduce the speed. The latter is chosen because the heading of the vehicle is pointing away from the reference path, hence covering more distance in the wrong direction would increase the LPE.

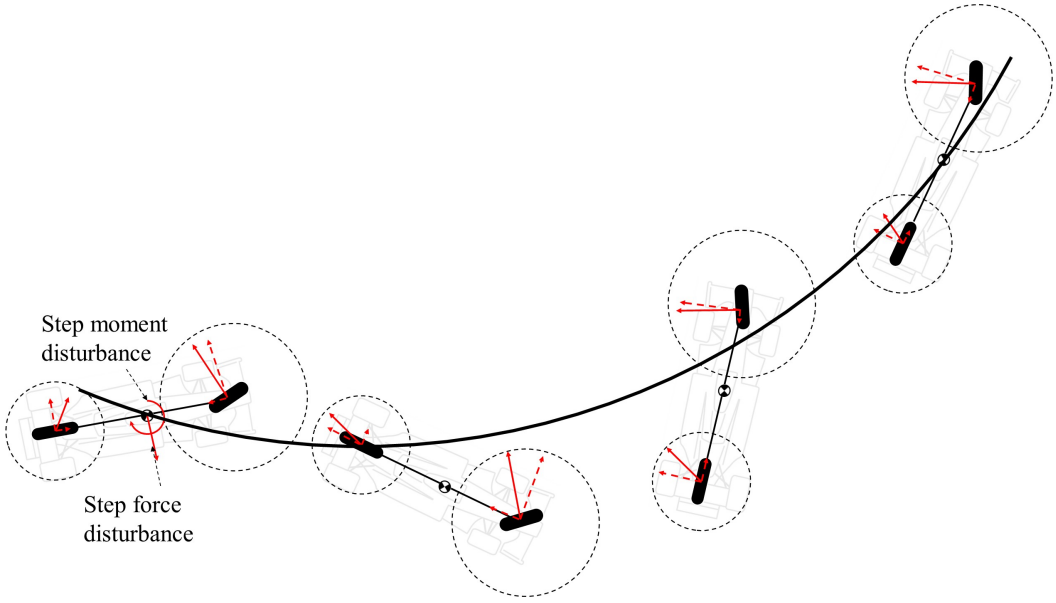

Figure 4.26: Sketch of the vehicle for Case 4 as it is disturbed. Both the heading angle and wheel angles in the sketch are not to scale to emphasise the behaviour of the vehicle. Figure 4.25 shows the actual values. Rear and front tyre circles are shown with arrows indicating the magnitude and direction of the forces. The vehicle is shown at 200 m, 225 m, 275 m and 375 m.

Figure 4.26 shows four snapshots of Case 4 for $q_3 = 0.025$ to illustrate the vehicle time evolution as it is disturbed. The drawing is purely illustrative and does not introduce any new quantity. Geometries have been exaggerated to show the behaviour of the vehicle. The dotted circles around the tyres indicate the maximum available force and the red solid arrows the actual forces. Force components parallel and perpendicular to the vehicle longitudinal axis are also shown by the dotted arrows.

The first frame, taken at 200m, shows that as the vehicle is navigating the circle in steady state before the pulse disturbance is applied. The lateral and moment pulse disturbances are applied immediately after the frame. The vehicle is displaced laterally and turned towards the outside of the circle. In the second frame, corresponding to 225m, it travels towards the outside of the circle. The lateral and longitudinal coupling

becomes apparent here: the rear longitudinal slip increases, increasing the longitudinal force, thus decreasing the effective cornering stiffness of the tyres, and making the vehicle easier to turn towards the inside of the circle. While the LPE is still negative – Figure 4.24c at 275 m – the heading error overshoots the equilibrium position in the third frame. The rear slip becomes positive to accelerate the vehicle and to provide the lateral force to turn the vehicle back to its equilibrium position. The fourth frame shows the vehicle back to its equilibrium state.

Random disturbance

The response of the controller under the action of random disturbances is also considered. The US vehicle in its baseline configuration navigates a 120 m radius circle at 30 m/s. Tyres in the nonlinear region but far from saturation. The nature of the disturbances is described in Section 2.4 white Gaussian noise with zero mean and an arbitrary standard deviation. Figures 3.4 and 3.7 shows the model with the disturbances acting on it. Given that the mean is zero, the direction of the force or moment is irrelevant. The standard deviations values are included in Table 4.5.

Two cases are considered. In Case a1 a force and moment disturbance act on the vehicle and q_3 is varied logarithmically from 10E-6 to 10E0. In the second case, Case a2, a longitudinal disturbance acts on the vehicle and q_5 and q_6 are varied logarithmically from 1E-6 to 1E0. Disturbances are described in 4.2.3.

Figure 4.27 shows the relationship between the LPE and $\dot{\delta}$ for Case a1. As the weight on $\dot{\delta}$ increases, its RMS decreases and the LPE increases as expected, showing an inverse relationship between control effort and path tracking accuracy.

Figure 4.28, again Case a1, shows that as $\dot{\delta}$ decreases, $\dot{\kappa}$ increases slightly. Even though the variation is not dramatic, this test suggests a degree of lateral and longitudinal coupling. In this case, the increase in longitudinal action does not affect tyre cornering stiffness as tyres are not saturating so it is likely to be used to adjust the speed, reducing it when the vehicle heading is not optimal. The overall effect on the LPE is however negligible as Figure 4.27 shows that $\dot{\delta}$ is the main factor influencing the LPE.

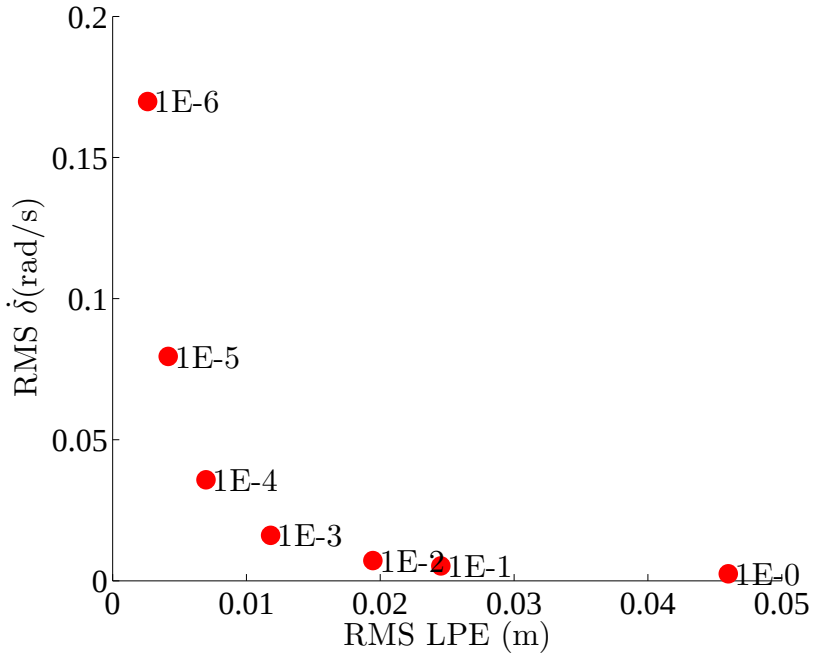

Figure 4.27: LPE RMS vs $\dot{\delta}$ RMS for a US vehicle navigating a circle at constant speed under the action of lateral disturbances – Case a1. Labels indicate the value of q_3.

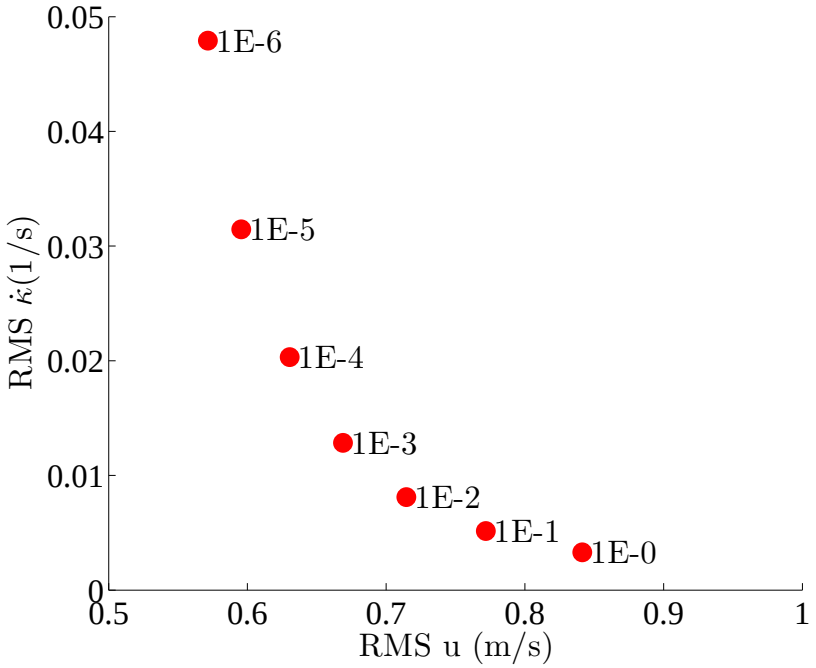

Figure 4.29: $\dot{\kappa}$ RMS vs speed error RMS for a US vehicle navigating a circle at constant speed under the action of longitudinal disturbances – Case a2. Labels indicate the value of q_5 and q_6.

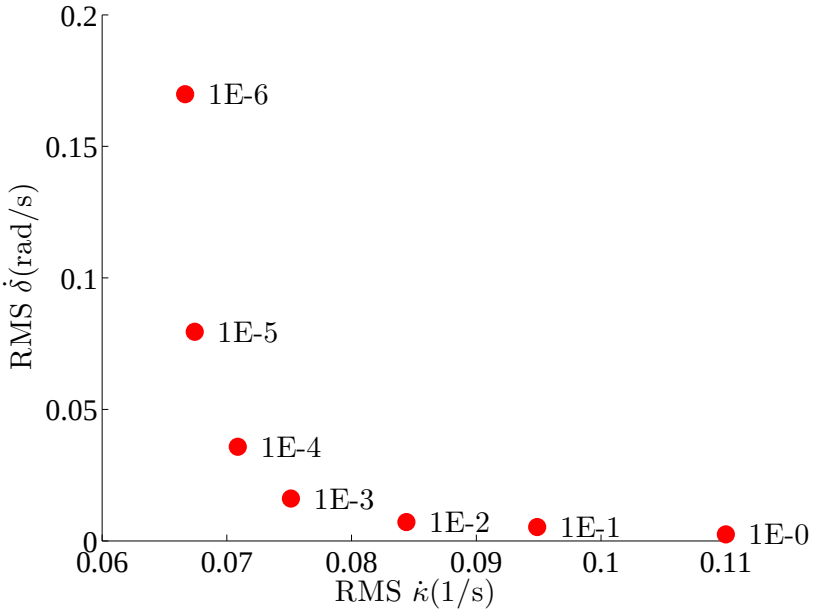

Figure 4.28: $\dot{\delta}$ RMS vs $\dot{\kappa}$ RMS for a US vehicle navigating a circle at constant speed under the action of lateral disturbances – Case a1. Labels indicate the value of q_3.

Case a2 focuses on the longitudinal controls. The longitudinal slip weights q_5 and q_6 are both varied between 1E-6 and 1E0 simultaneously. Figure 4.29 shows that there is an inverse relationship between longitudinal controls and speed error, again highlighting the inverse relationship between control effort and speed tracking accuracy. Finally, Figure 4.30 shows that the coupling between the lateral and longitudinal controls is weaker when longitudinal disturbances are applied. The $\dot{\delta}$ range is in the order of 1E-3 while the variation in Case a1 was in the order of 1E-1. Despite its magnitude, there is a clear trend that shows an increase in $\dot{\delta}$ for increasing $\dot{\kappa}$. Interestingly, the LPE range, not shown as it would not add any relevant insight into the behaviour of the controller, is 1E-6 so it is not the cause for the $\dot{\delta}$ increase. The vehicle travels slightly slower than the reference speed. As $\dot{\kappa}$ increases, the speed error reduces and the vehicle travels slightly faster, which increases the mean δ necessary to navigate the circle as the vehicle is US. The tyre curve, despite not saturating, is still nonlinear so an increase in the operating point can decrease the local stiffness, hence an increase in steering activity. This effect is almost negligible but relevant to explain the $\dot{\delta}$ increase.

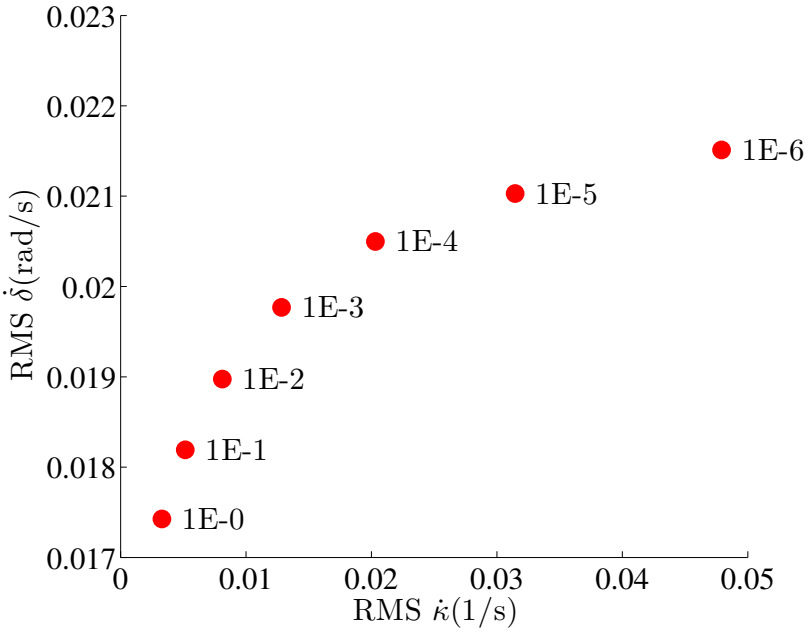

Figure 4.30: $\dot{\delta}$ RMS vs $\dot{\kappa}$ RMS for a US vehicle navigating a circle at constant speed under the action of longitudinal disturbances – Case a2. Labels indicate the value of q_5 and q_6.

4.5 Summary

The aim of this Chapter was to assess the validity of the formulation proposed in Chapter 3 and to gain insight into the operation of the controller. A lateral only and a combined controller are considered. Driving scenarios of increasing difficulty are considered. The lateral controller is tested on a circular path at constant speed under the action of disturbances. The controller manages to stabilise a US and OS vehicle at the handling limit, confirming the robustness of the formulation. The combined controller is tested in a similar yet more challenging scenario where the target path and speed are physically unachievable for the controller. Again, the controller minimises the deviation from the target, stabilising the vehicle. The controller has also been successfully tested in more demanding driving conditions – a 90 degree corner at the limit of handling – managing to stabilise the vehicle.

Parameter studies on the lateral and longitudinal controller are also included in the Chapter. Different methodologies have been used to confirm a correct understanding of the effect of the parameters on the states and the controls. It was found that the weight q_3 on the hand wheel angle rate has a significant effect on the HWA and that it is coupled with the NMS natural frequency. It was also shown that the update time

does not affect the frequency response of the steering significantly but does have a significant effect on the LPE. The combined parameter study further confirmed the direct relationship between control effort and path tracking accuracy. It also showed that longitudinal controls can be used to drive lateral states to equilibrium.

There are two main original contributions to knowledge in this Chapter. Firstly, it has been confirmed that the formulation in Chapter 3 can stabilise a vehicle under the action of disturbances in a variety of scenarios. Secondly, it has been shown that the controller can use longitudinal controls to drive lateral states to equilibrium, confirming the predicted coupling.

Chapter 5

Lateral controller validation

5.1 Introduction

Chapter 3 detailed the derivation of the lateral and longitudinal controllers while Chapter 4 investigated their performance in different driving conditions and assessed the effect of key parameters. The results obtained are encouraging, as the controller manages to stabilise the vehicle in very challenging driving scenarios. The nonlinear formulation successfully tracks the optimal nominal states generated by an MPC minimising manoeuvre time for a 90 degree corner.

A more structured validation is necessary to confirm how well the controller can mimic a real human driver. To this end experimental data is collected from a Driver In the Loop (DIL) driving simulator. Even though one might argue that comparing the simulation results with actual track data would be more insightful, performing the experiments using a driving simulator has several advantages:

- *Ability to vary track geometry and car setups*

  Performing experiments in a virtual environment allows one to choose arbitrary track geometries and car set-ups with minimal effort, facilitating exploration of controller performance in very specific and controlled operating conditions. Data from track sessions would allow for only a limited range of operating conditions.

- *Data availability*

  Even though racing cars are fitted with several sensors, certain quantities are difficult to measure directly and have to be estimated from other measurements. Performing experiments in a driving simulator gives easy access to all variables in the environment.

- *Driver feedback*

  Performing experiments in the driver simulator allows subjective feedback from the driver to be collected easily, which would not be possible if track data from previous races was used. Even though the driver of a real car could also provide feedback, performing the experiments on a real vehicle to gather feedback would be overly expensive. Since the overarching aim of the study is to enhance understanding of how racing drivers handle vehicles at the handling limit, having qualitative feedback is useful as it allows for comparison with the quantitative results.

Despite the aforementioned advantages, it is important to question the validity of the driving simulator data as the whole experimental validation rests on it. Formally quantifying driving simulators performance would be problematic because of the analysis intrinsic subjectivity; however, they are ubiquitous in the racing industry. Formula 1 teams have been using them for several years as substitute for testing and for driver training purposes. This reliance on simulators show that DIL simulations can replicate actual driving to a satisfactory degree of accuracy. If the proposed controller could match DIL simulations fidelity, Formula 1 teams could use it, saving a considerable amount of resources every year.

Section 5.2 describes the experimental setup, Section 5.3 analyses the data collected from the driving simulator, Section 5.4 outlines the procedure used for model identification and Section 5.6 shows the comparison between experimental and simulated data. This Chapter focuses on the validation of the lateral controller while Chapter 6 examines the combined controller. The two are considered separately because they require different driving experiments.

5.2 Driving experiment

5.2.1 Experimental setup

The experiment was performed in a driving simulator consisting of the body of a car mounted on a moving platform allowing mainly lateral and yaw motion. Electric actuators are used for motion cueing. An offboard, 8 meters diameter, 270-degree viewing angle screen is mounted in front of the platform for visual cueing. Images are projected by 6 projectors with a 120 Hz update frequency and 2560x1600 px resolution. The steering system has a force feedback mechanism for advanced cueing. More information about the simulator hardware and software, developed in-house,

could not be obtained due to confidentiality reasons. The vehicle model used in the driving simulator is a full nonlinear version of a 2014 Formula One car in the baseline configuration for the Circuit de Barcelona-Catalunya. The vehicle model, tyre model and aerodynamic map were not disclosed to the author for confidentiality reasons.

5.2.2 Driving task

A professional driver with extensive racing experience was asked to follow circular paths of different radii at constant speed under the action of disturbances. Engineers set up the digital environment with the desired vehicle configuration and a circular track. The simulator would then start and the vehicle would be brought to steady state at the target speed before disturbances started acting on the vehicle. The driver would then try to follow the target path under the action of the disturbances for at least 3 km, before decelerating and ending the session. This procedure, referred to as a *run*, is repeated several times with different parameters values to explore the behaviour of the vehicle. Vehicle speed is held constant and the driver only controls the steering.

Figure 5.1 shows the vehicle navigating the circle and the main associated quantities. Table 5.1 details the conditions for all the runs. Since simulation time is expensive, the experiment was not exclusively performed for this validation so specific changes in vehicle configuration are also due to the requirements of other experiments that were performed in the same session. The first column indicates the run number, the second the radius of the circle and the third the constant speed of the vehicle in m/s. Set up changes are indicated in the fourth column. The Baseline configuration refers to the vehicle parameters in Table 5.2. Finally, the steady state lateral acceleration, expressed as a percentage of the maximum lateral acceleration, is also provided. The value is the ratio between the lateral acceleration of the given vehicle in the given conditions and the maximum lateral acceleration that the tyres can physically allow for the given speed. Lateral acceleration is a convenient metric to define how close to the limit the vehicle is as it considers both vehicle speed and radius of curvature. The value provided refers to steady state conditions, so the actual disturbed vehicle is likely to be closer to saturation.

The vehicle was subjected to an external random force acting 0.67 m in front of the vehicle CoG pointing towards the outside of the circle; the yawing motion of the vehicle is therefore disturbed as well. While the disturbance profile used in Chapter 4 was random Gaussian noise, the experimental one is different. Figure 5.2 shows the disturbance MSSD. It is characterised by a 1 Hz cutoff frequency. The mean force is 57 N [49 N, 66 N at 95 % confidence interval] and the standard deviation is 717 N [712 N,

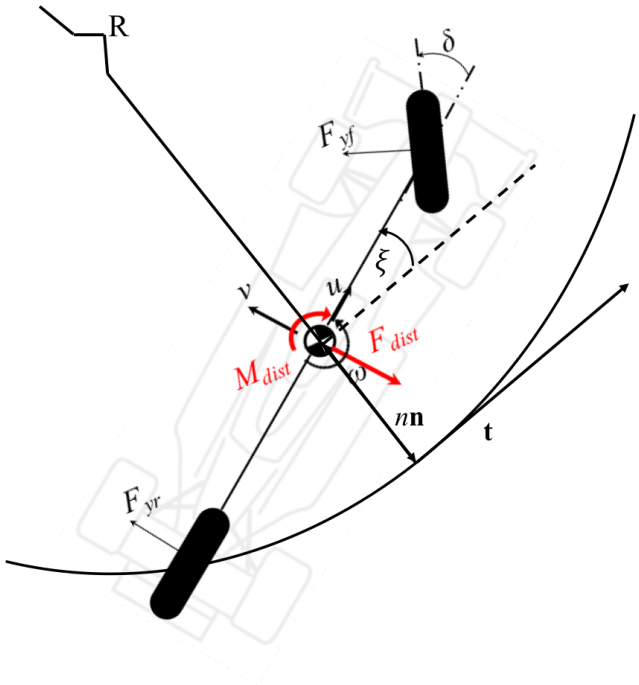

Figure 5.1: Setup for the experiment. The vehicle navigates a circle while a random force and moment disturbances are applied. R indicates the radius of the circular path. The force and moment are resolved at the centre of mass. Their distribution is almost symmetrical about zero so the direction in the diagram is arbitrary.

724 N at 95 % confidence interval]. These values are used for all the runs apart from 6,7,8 and 9, where they are halved. The mean is not zero as in Chapter 4; however, this is not a concern for the validation as the exact same disturbances are replicated in the simulation environment. Furthermore, the mean is 7 % of the standard deviation, suggesting that its contribution is minor. Kolmogorov-Smirnov test indicates that the force probability density function can be approximated by a normal distribution with a 5 % confidence level. The pdf and the normal distribution are shown in Figure 5.3.

Table 5.1: Description of the experimental conditions for all runs. The percentage of the maximum acceleration is the ratio between the vehicle lateral acceleration and the maximum lateral acceleration tyres can physically allow for the given speed rounded to 5% . CoG indicates the centre of gravity and CoP centre of pressure.

Run	Radius R (m)	Speed u (m/s)	Set-up changes	% of max lat. acc'n a_y
1	80	40.8	Baseline	85
2	80	40.8	CoP moved 4 % forward	85
3	80	40.8	CoG moved 2% forward	85
4	80	40.8	Vehicle yaw inertia increased by 20%	85
5	80	31.7	Baseline	60
6	80	44.4	Disturbance halved	90
7	40	26.9	Disturbance halved	90
8	115	60.3	Disturbance halved	90
9	115	55.6	Disturbance halved	85
10	165	67.5	Baseline	75
11	115	54.2	CoP moved 4 % forward	85
12	165	67.5	CoP moved 4 % forward	75

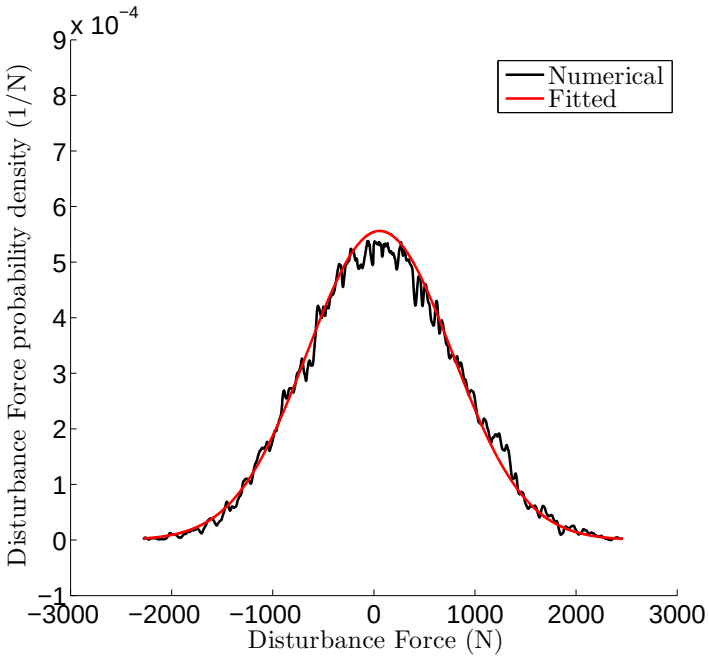

Figure 5.3: Disturbance force probability density function and fitted probability distribution. Kolmogorov-Smirnov test indicates to a 5 % confidence interval that a normal distribution fits the data .

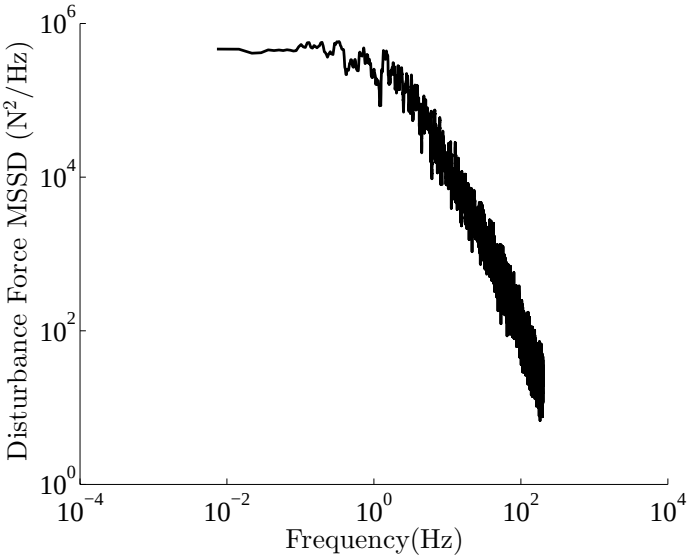

Figure 5.2: Mean Square Spectral Density for force disturbances showing the cutoff frequency at 1 Hz

Most of the simulation parameters, namely the tyre model, driver characteristics, aero-map and vehicle model, are not known, so need to be estimated. Data for a bicycle model approximating the simulator model was provided by the company. The aero-model does not have to be fitted because only steady state operations are considered. The identification procedure for driver characteristics and tyre model is described in Section 5.4. Parts of the reported runs have not been considered because of the driver losing control of the vehicle or hardware malfunction.

5.3 Data analysis

The data is collected from the driving simulator at a sampling frequency of 200 Hz and fully describes the motion of the vehicle and the tyre forces. The travelled path is given in cartesian coordinates. The total distance travelled and the vehicle speed are also provided separately. The hand wheel angle is provided, together with the tyre slips and the tyre forces, allowing calculation of all the vehicle states and tyre forces.

In order to identify the experimental model parameters and perform the validation, the lateral path error, the steering rate and the tyre forces for the bicycle model described in Section 3.3.3 are required. Since the vehicle travels along a circle centred at $X = 0$ and $Y = 0$ with deviations on either side, the lateral path error will simply be the difference between the magnitude of the position vector and the radius of the

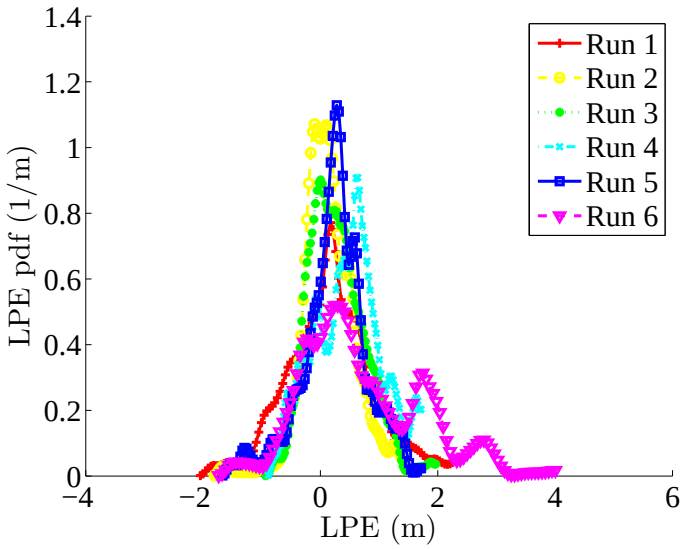

Figure 5.4: LPE pdf for the first 6 runs. The probability density function shows that the distribution resembles a normal distribution with a fairly large standard deviation. Only the first 6 runs have been shown to avoid clutter.

circle.

$$n_{exp} = R - \sqrt{X^2 + Y^2} \tag{5.1}$$

where X and Y are the Cartesian coordinates of the vehicle. The steering rate is obtained by numerically differentiating δ_{sw}.

5.3.1 Lateral Path Error

Figure 5.4 shows the LPE pdfs for runs 1 to 6 and Figure 5.5 for runs 7 to 12. The LPE is approximately normally distributed with a large standard deviation and nonzero mean in most cases. The fact that the driver may tend slightly to the outside instead of the inside could be attributable to the nature of the disturbance and the tyre saturation. The following qualitative discussion attempts to explain the phenomenon which deserves more attention. Section 5.3.2 shows that disturbances have a positive mean, resulting in a mean force towards the outside. Even though the driver should be able to compensate for such force, it may still have an effect. Secondly, an increase in steering action would increase the front axle saturation level while decreasing δ would decrease it, causing the vehicle to be harder to control when steered towards the inside. Statistically, the driver will therefore be better at rejecting disturbances pushing the vehicle towards the inside of the circle, this resulting in a positive LPE.

Such asymmetry is confirmed by Figure 5.6, which shows that the mean LPE values for all runs are positive. The mean of all mean values, 0.34 m, is half of the mean standard deviation, 0.78 m, indicating that the steady state deviation is not negligible. Furthermore, the highest mean values, correspond to high levels of tyres saturation, confirming that slip level affects the compensatory action.

5.3.2 Driving strategy

The experimental data is quite insightful, as it reveals interesting features of the driving strategy.

The hand wheel angle time history shows that the driver does not continually adjust δ_{sw}, instead holding it constant for time spans from 0.1 to more than 1.0 s. Figure 5.7 clearly shows ZOHs across runs and hence across operating conditions. Hold strategies have been described in detail in Section 3.3.4.

For a ZOH $\dot{\delta}_{sw}$ is theoretically zero and practically very close to zero. Quantifying the ZOH gives an important insight into the driver's behaviour. The simplest and most direct approach is to calculate the proportion of time over which $\dot{\delta}_{sw}$ is close to zero. The integration boundaries are somewhat arbitrary; for the purpose of this work $|\dot{\delta}_{sw}|<0.005$ rad/s is considered close to zero. Integrating $\dot{\delta}_{sw}$ pdf around zero shows that δ is flat on average 60 % of the time, quantitatively confirming the presence of

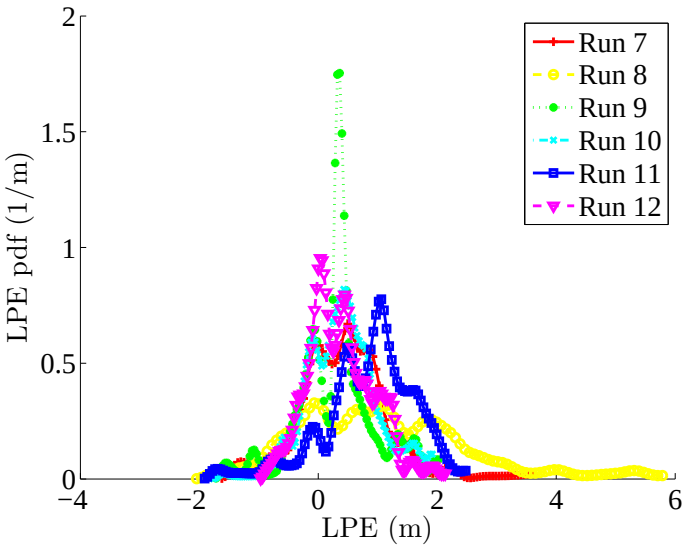

Figure 5.5: LPE pdf for runs 7 to 12. The probabilities density functions are similar to the first 6 runs.

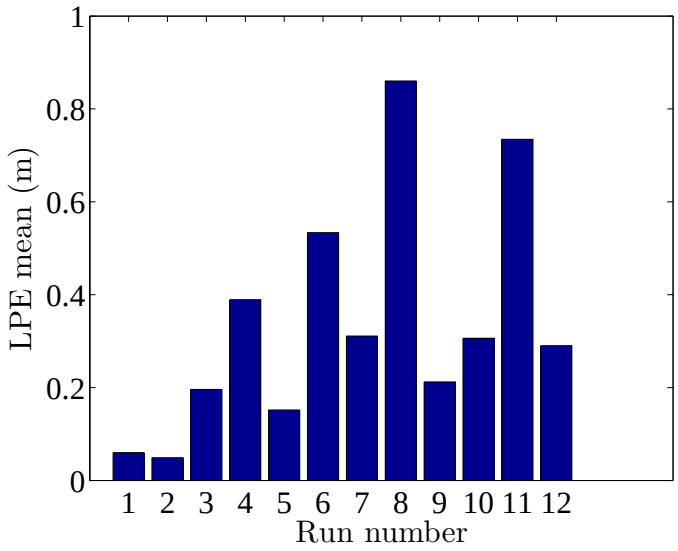

Figure 5.6: LPE mean values for all 12 runs. The chart shows that the values are always positive, implying a consistent bias towards the outside of the circle. The highest mean values correspond to runs where the vehicle is closer to saturation.

hold behaviour. This metric however only shows the total length of the holds for each run but it does not provide any insight into the length of each hold.

The hold behaviour would be better quantified by a metric that takes into account the ratio between hold time and total time. Given that data sampling frequency is 200 Hz, the hold length space can be discretised with a 0.005 s time-step. H_p indicates the total hold time of holds of length p, where p is a multiple of 0.005 s.

$$H_p = p * N_{t,p} \tag{5.2}$$

where $N_{t,p}$ is the total number of holds of length p. For instance, $H_{0.1}$ indicates the total hold time for holds of 0.1 s.

$$H_{0.1} = 0.1 * N_{t,0.1} \tag{5.3}$$

where $N_{t,0.1}$ is the total number of holds of length 0.1 s. Referring to the longest hold length as $H_{p,max}$, the hold length space is divided into $f_{hold} = H_{p,max}/0.005$ units. A metric capturing the hold behaviour would be the ratio between the cumulative hold time for a specific hold and the total hold time. The ratio for hold lengths p can be

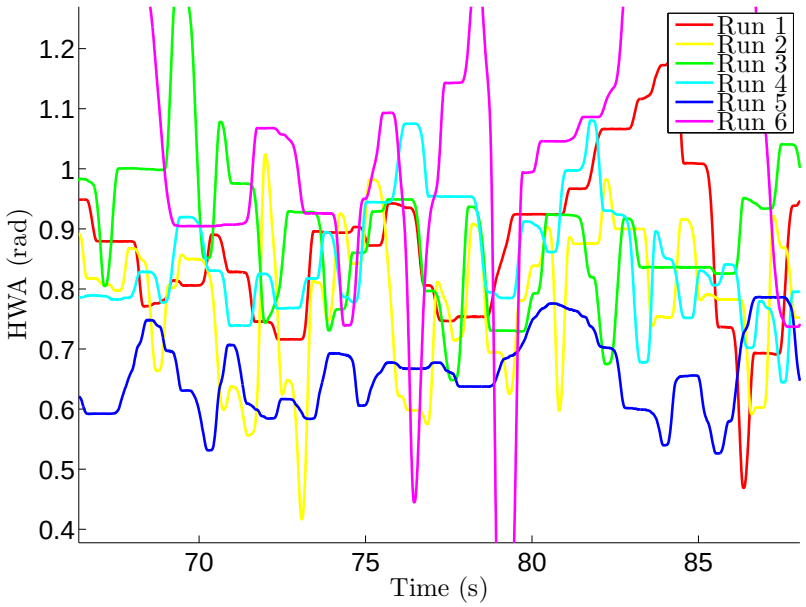

Figure 5.7: Time histories for the HWA for 6 runs. The hold behaviour is evident in all the runs.

defined as

$$H_p^{per} = \frac{\sum_{j=p}^{f_{hold}} H_j}{\sum_{j=1}^{f} H_j} * 100 \tag{5.4}$$

Clearly, for $p = 1$, $H_p^{per} = 100\%$. Plotting H_p^{per} $\forall p$ against p, Figure 5.8 gives an excellent insight into the hold behaviour. Firstly, it shows that the hold behaviour is consistent over different runs; run 5 deviates slightly from the mean for short holds but the all other runs show a remarkable consistency. At 0.005 s, the cumulative time is equal to the total hold time. H_p^{per} decreases slightly up to 0.1 s, indicating that a large proportion of holds are shorter than 0.1 s. It then drops off significantly between 0.2 and 0.4 s, which corresponds to a racing driver update time range [2]. Only few holds are longer than 1 s.

Figure 5.8 enables understanding of how holds are distributed but it does not provide any information about how often they are used. A new metric, similar to H_p^{per} is therefore introduced to account for the total time. Referring to the total run time as T_{run},

$$H_p^{pert} = \frac{\sum_{j=p}^{f_{hold}} H_j}{T_{run}} * 100 \tag{5.5}$$

Figure 5.9 shows that the total hold time is between 60 % and 70% for the first 6 runs and it is significant across hold lengths.

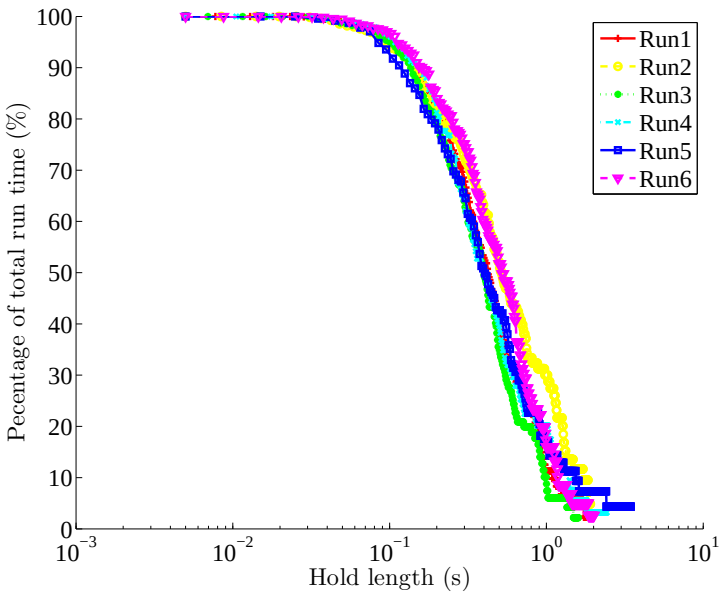

Figure 5.8: H_p^{per} showing a marked decrease between 0.2 and 0.4 s.

It is now clear that ZOH needs to be accounted for in the model. It has so far been assumed that the driving strategy relies on SBHs and hence on the driver applying the calculated sequence of controls over the refractory time. This assumption needs to be modified to account for the new findings discussed.

The driving strategy can be modelled as pure ZOH, where hold transitions are smooth because they are filtered by the NMS. In order to check this hypothesis, a pure ZOH has been superimposed to the experimental one assuming that δ_{sw} takes the value of the following hold as soon as each one finishes. The signal is then filtered by the neuromuscular filter and superimposed. Figure 5.10 shows the three signals: experimental, ZOH and filtered ZOH.

From now on, ZOH present in the experimental signal will be referred to as ZOH_{exp} to distinguish them from the superimposed ZOH. The time interval is divided into five sections. The experimental signal exhibits three ZOHs_{exp}, section 1; at $\delta_{sw,a}$; 3; at $\delta_{sw,b}$; and 5; at $\delta_{sw,c}$, connected by smooth signals. A pure ZOH signal is superimposed at $\delta_{sw,a}$ over section 1. The pure ZOH then increases to $\delta_{sw,b}$ over sections 2 and 3 where the new ZOH_{exp} begins to finally decrease to $\delta_{sw,c}$ in section 4 and 5. The filtered ZOH is the pure ZOH passed through a NMS. Its actual shape depends on the nature of the filter. This assumption is tested by superimposing the pure and filtered ZOH on two sections of Run 1. Figure 5.11 show that the difference between the experimental δ_{sw} and the filtered ZOH is acceptable for small changes between holds but becomes

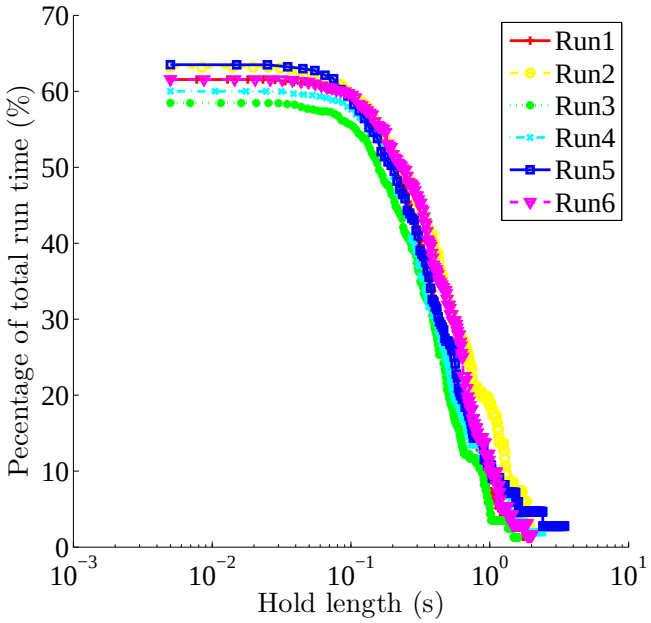

Figure 5.9: H_p^{pert} showing that the driver holds δ_{sw} for more than 60 % of each run.

significant for larger changes. More importantly, the experimental δ_{sw} gradient is different from the filtered ZOH, showing that the driving strategy is different.

The idea of a filtered ZOH is not the best approach. An alternative way to model the steering action is to assume that it is a mix of SBH and ZOH. One can assume that the driver applies a series of steering actions and then, once they feel that the vehicle is in a stable configuration, they keep the δ_{sw} constant, essentially moving from

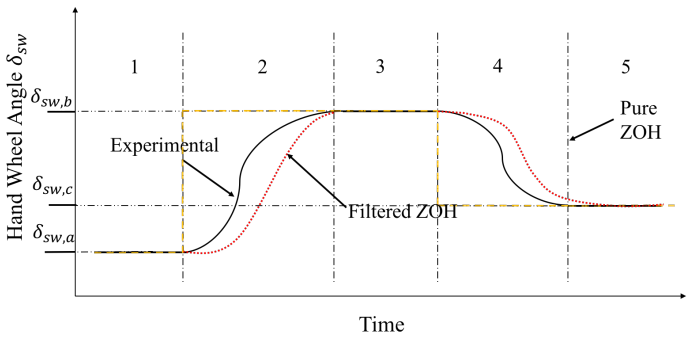

Figure 5.10: Sketch showing the superposition of the ZOH and the filtered ZOH signal over the experimental one. Labels do not report units because the sketch is for illustrational purposes only.

one stable configuration to the other. Figure 5.12 sketches the proposed steering signal, formed by a SBH sequence, lasting T_{SBH}, at the beginning followed by a ZOH, lasting T_{ZOH}. Both T_{SBH} and T_{ZOH} are allowed to vary to replicate experimental results. The total update T_u time will be

$$T_u = T_{ZOH} + T_{SBH} \tag{5.6}$$

5.4 Vehicle and tyre model characterisation

The model used in the simulator is a fully nonlinear version of the 2014 car in the Barcelona configuration, with full details not disclosed to the author for confidentiality reasons. Bicycle model parameters have to be identified to give best fit to the simulator vehicle model.

Before exploring the procedure that has been used to fit parameters, a new convention is introduced for the sake of clarity. From now on, *experimental* data/model refers to the data from the driving simulator and *simulation* data/model refers to the results of the simulations run with the nonlinear controller. Since both data come from a digital environment, the distinction is made clear. In order to replicate the experimental data three groups of parameters need to be considered: the vehicle, tyres and driver.

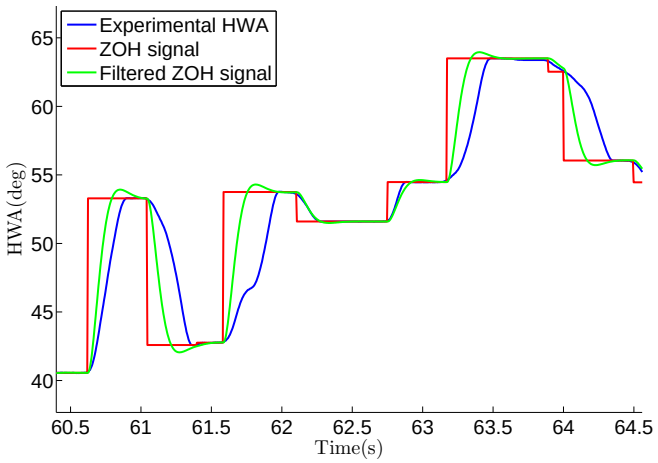

Figure 5.11: Comparison between experimental data, Zero Order Hold approximation and filtered ZOH for a section of Run 1 that shows poor agreement.

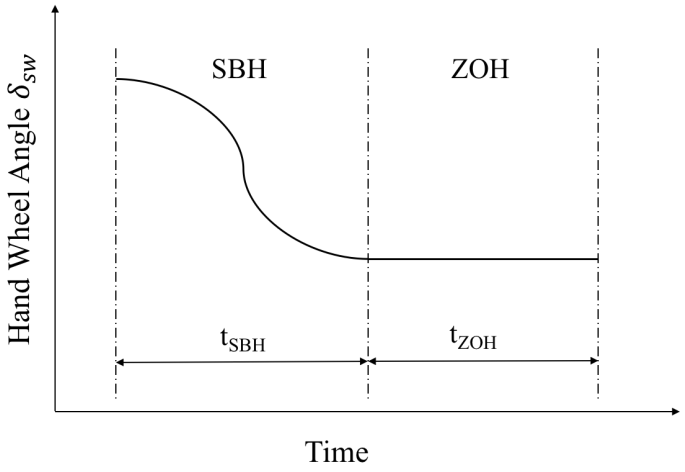

Figure 5.12: Proposed driving strategy, a Serial Ballistic Hold followed by a ZOH.

5.4.1 Vehicle parameters

Vehicle parameters for an equivalent bicycle model were provided by Renault and are reported in Table 5.2. The model includes the centre of pressure for the aerodynamic model without specifying the aerodynamic map. However, the map is not necessary for this experiment as the speed is constant, meaning the aerodynamic effects are static and are embodied in the vertical tyre forces.

Table 5.2: Bicycle model parameters

Parameter	Symbol	Value
Mass	M	735 kg
Inertia	I_z	1100 kgm^2
Front distance to CoG	a	1.89 m
Rear distance to CoG	b	1.61 m
Front distance to CoP	a_{COP}	1.93 m
Rear distance to CoP	b_{COP}	1.57 m
Steering ratio	G_{sw}	11

5.4.2 Tyre model identification

The tyre model is not provided for confidentiality reasons. The time history of tyre slip and the tyre forces are available, so a tyre model can be fitted to the experimental data. While the experimental vehicle model is four wheeled, simulations are run with a bicycle model, which only has two wheels. In order to accommodate this difference, the total axle force is obtained by adding the right and left wheel forces

$$F_{yit} = F_{yil} + F_{yir} \tag{5.7}$$

where the subscript $_y$ refers to the force direction, $_i$ the axle, front or rear, and $_l$ or $_r$ are abbreviations for left and right. The slip values for the left and right wheels, on the other hand, are averaged.

$$\alpha_t = \frac{\alpha_f + \alpha_r}{2} \tag{5.8}$$

where α_t is the total slip.

An established parameter identification technique consists of minimising the difference between the simulated and experimental output for the same experimental input [134]. In the case of tyre parameters identification, this translates to minimising the difference between the experimental tyre force and the simulated tyre force for the experimental slip time history. The cost function to be minimised is therefore

$$J_t = \sum_{k=1}^{k=P_v} (F_{exp,k}(\alpha_{exp,k}) - F_{sim,k}(\alpha_{exp,k}))^2 \tag{5.9}$$

where P_v is the vector length. There are a number of well established algorithms to perform this operation, from brute force to heuristic routines [135]. The Nelder Mead algorithm has been chosen because it is robust to noise and does not rely on gradient [134], making it ideal for non-smooth functions such as experimental data.

Nelder Mead makes use of polytopes formed by n + 1 vertices in n-dimensional space to numerically locate the minimum value of a function [134]. The function values are calculated at the vertices of the simplex and the worst point is replaced by one generated by reflecting across the centroid. Comparing the current best point with the previous best point indicates whether the simplex has to expand towards the newly reflected point or contract towards the previous best point. The iteration process continues until the diameter of the simplex is smaller than a specified tolerance value. The method can be implemented using the Matlab function *fminsearch*.

The parameters that have to be fitted for Equation (3.23) are: Q_i, μ_i, α_i.

Table 5.3: Tyre model coefficients fit for all runs. The subscript $_f$ denotes front and $_r$ denotes the rear.

Run	μ_f	Q_f	α_f	$feval_f$	μ_r	Q_r	α_r	$feval_r$
1	2.4858	1.6905	0.1600	2.67E-8	2.3564	1.7667	0.1597	5.98E-8
2	2.3564	2.5667	0.1597	4.56E-8	2.5439	1.6077	0.1565	4.16E-8
3	2.4438	1.6613	0.1246	3.57E-8	2.8926	1.0288	0.1343	4.2E-8
4	3.2446	0.7301	0.1101	3.72E-8	2.6126	1.2136	0.1318	3.71E-8
5	3.6432	0.6559	0.1132	3.90E-8	2.5544	1.9422	0.1432	3.43E-8
6	1.5749	3.6936	0.1495	5.47E-8	2.7386	1.7638	0.1813	5.47E-8
7	2.6259	2.6849	0.1807	9.70E-8	2.0802	1.4678	0.1332	2.99E-8
8	1.9846	1.3362	0.0921	2.77E-8	3.3103	2.1019	0.2102	4.04E-8
9	3.2751	2.6538	0.1859	6.16E-8	3.1584	1.6257	0.1241	8.76E-8
10	2.4858	1.6905	0.1435	1.10E-8	2.2432	2.4523	0.2334	3.61E-8
11	2.6374	1.8001	0.1495	7.49E-8	3.8414	0.8276	0.1394	6.13E-8
12	2.9898	1.5798	0.1241	2.82E-8	2.4136	1.6176	0.1662	4.32E-8

Tyres are fitted to a good degree of accuracy in all runs with normalised residues, $feval$, in the region of 1E-8, which is widely considered more than acceptable. The standard value in Matlab is 1E-4.

Table 5.3 shows that there are significant differences between the coefficients for different runs. In other words, the values of the coefficients that allow a good fit for a certain run are not the same for all other runs. Even though this is expected as conditions are different for each run, the extent of the discrepancies, up to 35 %, needs to be discussed. There may be two main causes of such differences: the assumptions made and the actual tyre model. The main assumption, made by adding tyre forces and averaging slips, can introduce an error. However, since the experiment is done in steady state conditions, this is unlikely to be significant because load transfer is accounted for in the tyre coefficients. Steady state conditions also eliminate variations in vertical force due to aerodynamic effects. Differences between the experimental and simulation tyre model are the most likely cause of discrepancies among runs. The simulation tyre model does not take into account thermal effects, wear and constant offsets. Even though the experimental model is not disclosed, it is very likely to include if not all, most of these effects. While the simulation tyre model may capture very well the slip force characteristics for one run, it may be problematic to extend this to more runs.

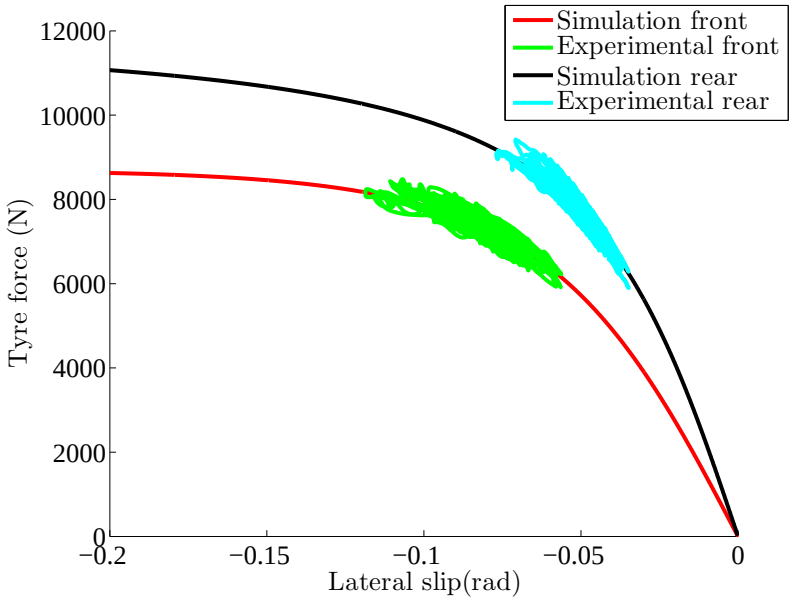

Figure 5.13: Tyre model fit for run 4

Figure 5.13 shows the fit for run number 4. The data fits the experimental results well. The experimental slip region does not extend beyond saturation so the slip force characteristics for large slips need to be extrapolated.

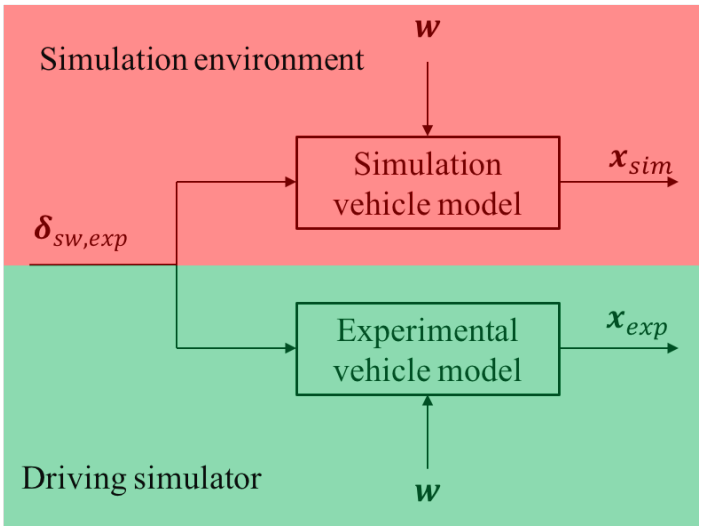

Figure 5.14: Diagram showing the vehicle model validation process. The experimental δ_{sw} and the same set of disturbances are fed to the experimental and simulated vehicle and tyre models and the outputs compared.

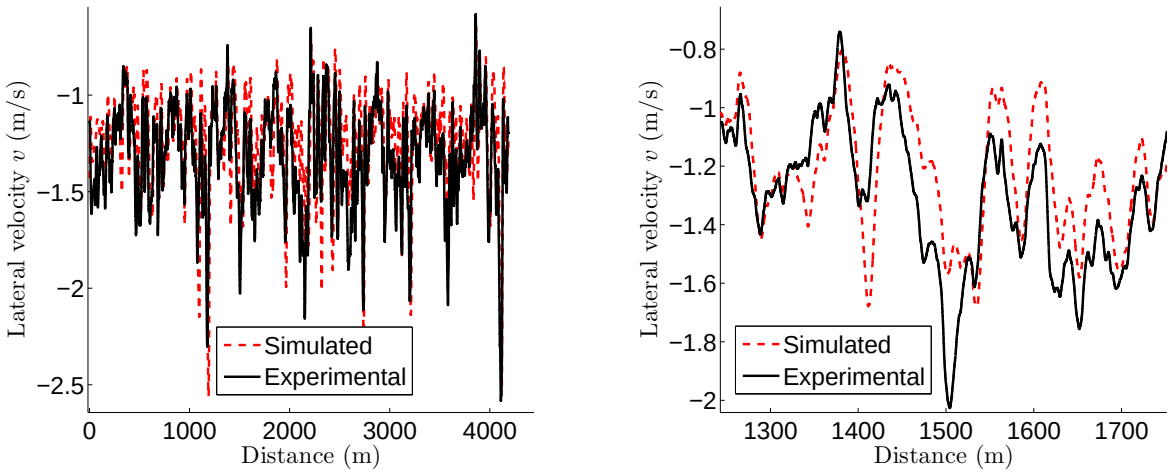

(a) Full time history

(b) 400 m time history section

Figure 5.15: Comparison between the experimental and simulated model lateral velocity.

5.4.3 Vehicle and tyre model validation

In this section the term *vehicle* model will refer to the *vehicle* and *tyre* model. The model accuracy is quantified to provide an indication for its contribution to the differences between the experimental and simulated signals in Section 5.6. The measured δ_{sw} is fed to the simulated vehicle model and its outputs compared with the experimental vehicle model. The process is detailed in Figure 5.14.

The output $\mathbf{x_{sim}}$ is then compared to $\mathbf{x_{exp}}$. Only v and ω are compared as they fully characterise the lateral vehicle motion. Data from run 1 is used for the comparison.

Given that the same disturbance time history is used, the systems are not random. Figure 5.15 shows that the experimental and simulation lateral velocities are reasonably close. The differences in the simulated and experimental signals is quantified using the mean of the relative error, $\mathbf{e}_{r,x}$. For any signal $\mathbf{x}$, the relative error is

$$\mathbf{e}_{r,x} = \left| \frac{\mathbf{x}_{sim} - \mathbf{x}_{exp}}{\mathbf{x}_{exp}} \right| * 100 \tag{5.10}$$

The mean $\mathbf{e}_{r,v}$, is 12 % with a 9% standard deviation. Figure 5.15b shows the signals in more details, confirming that the match is satisfactory. The difference in for the yaw rate is lower than the average velocity: the mean error is 3% with a 2% standard deviation. Given the assumptions and the parameter identification procedure, the match is satisfactory. Since the assumptions made for the bicycle model such as

lack of lateral load transfer and lack of differential, do not introduce a significant error for steady state driving, differences are mainly attributable to the tyre model. The tyre model identification inevitably leads to an error as the experimental tyre model is likely to account for several features that are not present in the simulation one.

The statistics for the first 6 runs are summarised in Table 5.4.

Table 5.4: Vehicle model validation errors from the first 6 runs rounded to the nearest percentage point.

Run	$\mathbf{e}_{r,v}(\%)$	$\mathbf{e}_{r,\omega}(\%)$
1	9	3
2	9	4
3	8	3
4	12	5
5	11	4
6	14	6

5.5 Driver model parameters fitting

The driver model parameters to be fitted are: the NMS natural frequency ω_n, the NMS damping factor, ξ_n, Eq 3.50, the weights q_1, q_2 and q_3, Eq 3.28, the driver update time T_u, holds length T_{SBH} and T_{ZOH} and the prediction horizon T_s, all described in

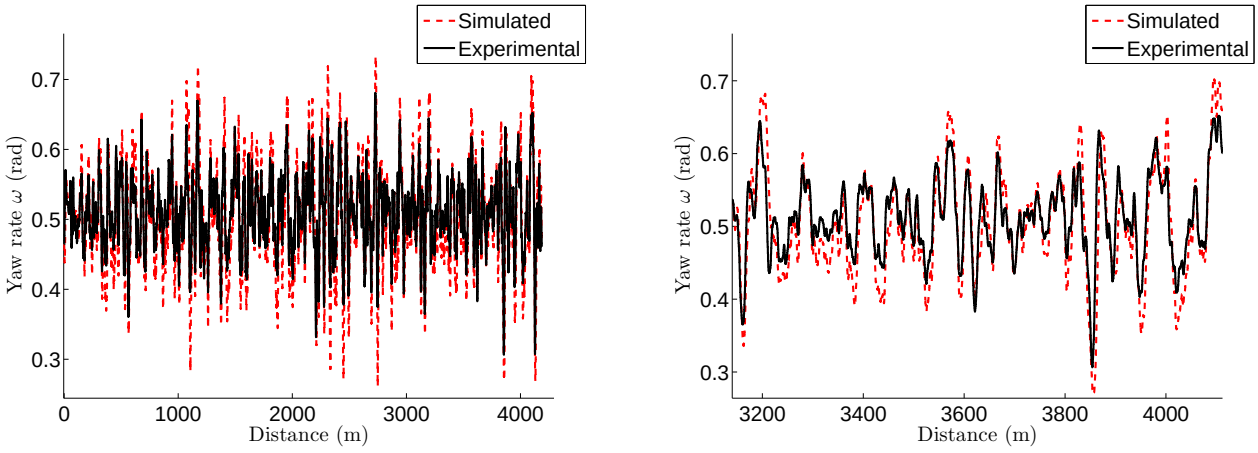

(a) Full time history

(b) 1000 m time history section

Figure 5.16: Comparison between the experimental and simulated model.

Section 3.3.4. Given the large number of parameters to fit, some assumptions have to be made to avoid over-fitting [135].

Since it is only the relative values of the weights of the cost function that are important, one of the three weights can be set to 1. The choice is arbitrary; however, preliminary parameter studies, not shown in this work, indicated that the system is less sensitive to changes in q_2/q_1 or q_2/q_3 than q_1/q_3 , so q_2 it will be set to 1 to improve the identification algorithm efficiency.

A reliable measure of the SBH and ZOH holds would be difficult to capture as they are not constant. The identification procedure would only yield constant values instead of a probability distribution. The lengths of the SBH and ZOH are therefore identified by measuring them directly from the experimental data; they then enter the identification procedure as known parameters. An algorithm has been devised to identify holds; these are taken as sections of δ_{sw} where the speed is close to zero, with the limit being set at 0.01 rad/s. The HWA is assumed to be a sequence of SBH and ZOH holds, each of length $t_{SBH,h}$ and $t_{ZOH,h}$, where h indicates the hold number. The update time will therefore be $T_{u,h} = T_{SBH,h} + T_{ZOH,h}$. The algorithm finds the signal section where $\dot{\delta}$ <0.01 rad/s. These are the ZOHs; the SBHs are the sections of the signal connecting the ZOHs. Figure 5.17 shows a portion of the time history for run 1 where holds are identified. The horizontal line shows the ZOH as solid lines and SBH as dashed lines. The number of holds prior to 50 s is referred to as n. The first SBH in Figure 5.17 has a length of $T_{SBH,n+1}$, which is followed by a ZOH of length $T_{ZOH,n+1}$. The SBH and ZOH holds that follow will last $T_{SBH,n+2}$ and $T_{ZOH,n+2}$ and so on. It is clear that the solid sections on the horizontal line correspond to the ZOH, showing that the algorithm successfully captures the ZOH and SBH.

The majority of ZOHs are below, with only a few beyond 1 s. The pdfs for the first 6 runs is shown in Figure 5.18, indicating an exponentially decreasing trend for all runs. A Pareto distribution would best describe the pdfs. For ZOHs shorter than 0.1 s, the ratio T_{SBH}/T_{ZOH} is uniformly distributed between 0.5 and 2, indicating that the SBH section of T_u ranges between half and twice T_{ZOH}. For longer ZOHs, T_{SBHs} is always less than 0.4*T_{ZOH}. This results in T_u never being longer than 2 seconds.

Furthermore, it has been shown in the literature – for instance in [136] – that the effect of the prediction horizon on the output becomes insignificant beyond a certain value, which varies depending on the driver model and the driving scenario. Running a quick parameter study showed that the prediction horizon has very little effect if it is beyond 6 s. In order to further simplify the parameter fitting, the prediction horizon

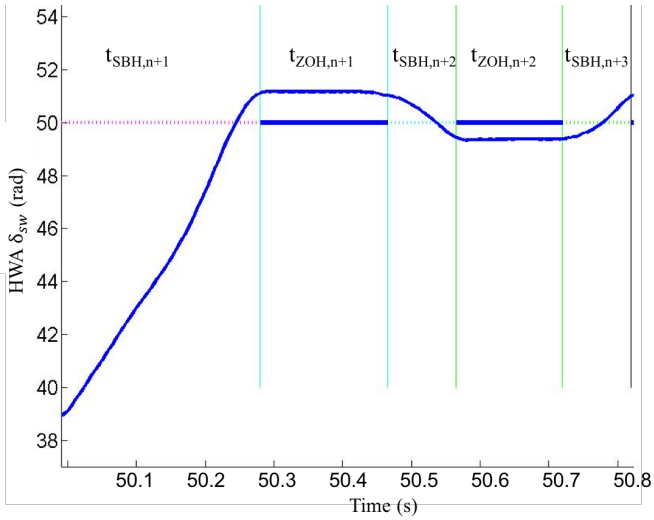

Figure 5.17: Section of the δ_{sw} signal for Run 1 showing the identified sequence of SBHs and ZOHs and their lengths.

is set to 6 s. The prediction horizon is three times longer than the longest update, so there is not coupling between the two.

Finally, experimental disturbance time histories have been used instead of generating random time histories using the statistics of the normal distribution fitted in Figure 5.3.

After all the aforementioned simplifications, the parameters that have to be fitted are therefore the NMS natural frequency ω_n, the NMS damping factor, ξ_n, the weight q_3 on the steering velocity and the weight q_1 on LPE. Due to the highly nonlinear nature of the problem, the choice of parameter identification algorithms is quite restricted. The main classes are [134]:

- *Derivative approximation methods*

  These methods approximate derivatives by finite differences and then use linear regression to find the parameters. However, this would defeat the purpose of the nonlinear system. Also, due to the iterative nature of the problem, it would be problematic to implement.

- *Bayesian methods*

  This family of methods makes use of probability based higher order systems such as *Markov Chain* and *Monte Carlo* to estimate parameters. This family of methods would require an prohibitively expensive computational power and also

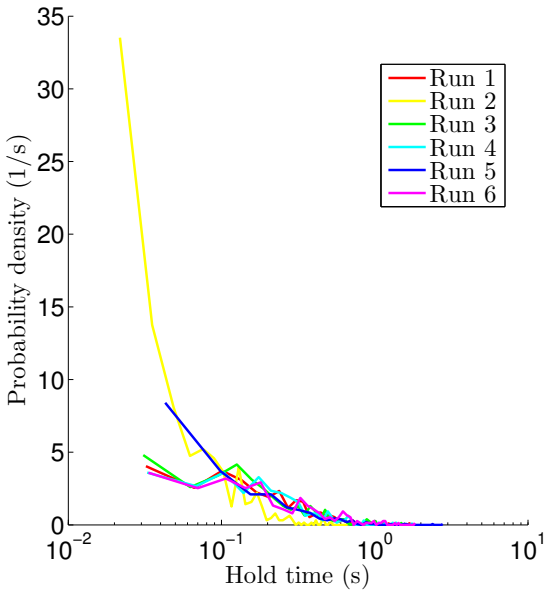

Figure 5.18: ZOH probability density function for the first 6 Runs.

would not be very useful as parameter identification does not involve random variables.

- *Optimisation methods*

  This family of methods attempts to minimise the difference between the simulation and experimental data by varying the model parameters; the algorithm used for the controller actually belongs to this class. The choice of cost function is therefore crucial.

Optimisation methods are the most suitable option because they offer a good trade off between accuracy and computational time. Derivative free methods are particularly indicated for experimental identification, where intricate computations are needed to find derivatives efficiently and accurately [137]. Nelder Mead, which has also been used to calculate the parameters for the tyre model, is a very fast algorithm and can achieve convergence significantly faster than other methods [134]. In order to fit the parameters, a cost function involving a metric that captures the essence of the model has to be set up. Since the noise in the driver NMS system introduces randomness in the system, it is necessary to introduce quantities that capture the statistics and frequency content of the signals to quantify the goodness of fit.

For the purpose of this analysis, the Mean Square Spectral Density (MSSD) of the steering input is considered. The MSSD is particularly useful, as it describes how the

variance is distributed over the frequency domain. The magnitude of the difference between the experimental and simulation values multiplied by the frequency window is therefore used to compare the two areas. An algebraic description is provided by Eqs (5.11) and (5.12) describes the The difference in time histories would not be a reliable metric because of the randomness in the driver NMS.

Even though it could be argued that an additional term accounting for LPE would be necessary, minimising the difference in steering MSSD is enough to match the simulation and experimental models. Since the process is stationary, the MSSD will just be the Fourier Transform of the autocorrelation function. In discrete time, considering a finite window of $1 \leq g \leq G$ with the signal sampled at discrete times δ_g=$\delta(g\Delta t)$ the MSSD becomes

$$\mathrm{S}_{\delta\delta} = \frac{(\Delta \mathrm{t})^2}{\mathrm{T}_{\mathrm{run}}} \left| \sum_{\mathrm{g}=1}^{\mathrm{G}} \delta_{\mathrm{g}} \mathrm{e}^{-\mathrm{i}\omega_{\mathrm{f}}\mathrm{g}} \right|^2 \tag{5.11}$$

where T_{run} is the total time of the experiment, ω_f is the normalised frequency and t time.

The cost function therefore becomes

$$J_p = \sum_{1}^{G} (\mathrm{S}_{\delta\delta}^{\mathrm{exp}} - \mathrm{S}_{\delta\delta}^{\mathrm{sim}}) \mathrm{d}\omega_{\mathrm{h}} \tag{5.12}$$

Eq. 5.12 is minimised using Nelder Mead algorithm. Some of the data is therefore needed for the fitting and some for the actual comparison. Parameters are also likely to vary between runs because of the human intrinsic randomness and the different vehicle operating conditions so it is necessary to perform a dedicated fitting for each run.

Augmented Dickey-Fuller tests [138] showed that in every run any window of at least 1.5 km can be considered stationary, implying that signal statistics are constant for any sample longer than 1.5 km. Since each run is about 4 km long, half of the run can be used for the comparison and half for the fitting. However, Nelder Mead is very sensitive to initial conditions [135]. Given that it is an heuristic algorithm, it does not guarantee a global optimum, allowing a variety of solutions. The interaction between the driver natural frequency and the penalty on δ_{sw} shown in Chapter 4 causes further complication. In fact, if the initial conditions for the NMS filter and weight q_3 in the cost function are chosen randomly, some numerically sound yet physically unrealistic combinations arise. The initial optimisation parameters need to be chosen in the vicinity of the solution for reasonably fast convergence.

Making use of the results in the parameter studies in Chapter 4, a set of initial optimisation parameters for all 12 runs is found. For this first set of initial optimisation parameters, convergence was only achieved for 8 runs. For the remaining 4, the initial optimisation parameters were iteratively tuned by plotting the cost function. Convergence was then achieved for all runs. The first 2 km of each run are used to identify the parameters. Results are shown in Tables 5.5 and 5.6.

Table 5.5: Simulation parameters runs 1 to 6 using 2 km of data. J_p^* indicates the minimum value of the cost function.

	1	2	3	4	5	6
NMS Natural Frequency (ω_n)(rad/s)	15.9	14.9	16.1	16.1	14.1	15.2
NMS Damping Factor (ξ_n)	0.71	0.69	0.70	0.69	0.68	0.67
LPE rate Weight (q_1) ($m^{-0.5}$)	3.52E-3	6.23E-3	2.83E-3	4.24E-3	5.61E-3	6.22E-3
$\dot{\delta}_{sw}$ Weight (q_3) ($rad/s^{-0.5}$)	2.12E-4	1.21E-4	6.34E-4	3.92E-4	5.16E-4	6.31E-4
J_p^* (rad^2)	1.12E-5	2.31E-4	5.42E-3	9.92E-2	9.20E-5	4.54E-2

Table 5.6: Simulation parameters for runs 7 to 12. J_p^* indicates the minimum value of the cost function.

	7	8	9	10	11	12
NMS Natural Frequency (ω_n)(rad/s)	16.9	14.7	14.9	15.3	15.4	16.1
NMS Damping Factor (ξ_n)	0.65	0.68	0.68	0.67	0.67	0.68
LPE Weight (q_1) ($m^-0.5$)	1.11E-3	4.24E-3	9.23E-3	6.87E-3	2.45E-3	9.46E-4
$\dot{\delta}_{sw}$ Weight (q_3) ($rad/s^-0.5$)	3.23E-4	7.21E-4	1.05E-3	6.18E-4	6.62E-4	4.29E-4
J_p^* (rad^2)	4.56E-5	1.92E-4	3.21E-5	7.62E-5	3.49E-5	6.21E-4

The identified values are consistent with the observed behaviour. Both the NMS Natural Frequency (ω_n), which varies between 15 and 17 rad/s, and the NMS damping factor, which varies between 0.65 and 0.7, are within an acceptable range for a racing driver [136]. Cost function weights are more difficult to consider as they do not have an immediate physical meaning. However, they can be compared to those used in Chapter 4; even though the model, driver and tyre parameters, are different, a qualitative comparison can indicate if the identified q_1 and q_3 are reasonable. Section 4.3.3 shows that the LPE has a 10^{-2} m order of magnitude for $q_1 = 1$; the experimental LPEs and the identified q_1 are in the order of 1 m and 10^{-3} respectively. A decrease of 3 orders of magnitude in the cost function weight resulted in an increase of 2 orders of magnitude in the signal, showing that the trend is reasonable. Finally, the value of q_3 and the experimental $\dot{\delta}$ are consistent with the set of results shown in Section 4.3.3 .

This first qualitative assessment is only indicative and preliminary, as a more insightful comparison is shown in Section 5.6.

Rather than using an average set of parameters from Table 5.5 and 5.6, each run is simulated with its respective fitted parameters. Since driving conditions change significantly, the NMS parameters and cost function weights vary as well, hence using the fitted parameters for each run leads to a better match. If an average set of parameters had to be used, the model should be extended to account for variable driver parameters.

5.6 Data comparison

After identifying vehicle, tyre and driver parameters, experimental and simulation outputs can be compared. The aim of the proposed controller is to replicate what the driver does. Given the noise in the driver NMS, comparing time histories – despite being indicative – would not be a reliable measure of the experimental and simulated signals differences, so statistics are considered.

The system has one input, δ_{com}, and four output states, namely δ_{sw}, v, ω and the LPE. Comparing all of them would not be practical so only δ_{sw} and the LPE are used, as they are the most indicative measures of the model performance.

The MSSD of the δ_{sw} is compared, the reasons for this choice being outlined in Section 5.5. Instead of comparing the MSSD of the total δ_{sw} signals, only the compensatory activity – obtained by subtracting the signal mean – is considered to highlight the corrective action of the controller. The signal mean is calculated by finding the average of the δ_{sw} time history. An increase or decrease in the mean could be attributable to the action of the controller compensating the disturbances due to the nonlinearity of the vehicle. However, the mean is subtracted for both the simulated and experimental signals, effectively taking this effect into account if present. A snapshot of the δ_{sw} time histories is also included to provide a more intuitive visual comparison.

The LPE probability density function (pdf) is also compared to assess how effective the model is at keeping the vehicle close to the reference path. Three runs: namely 1, 3, 4 and 6, are analysed as they provide enough information to gain an insight into the performance of the controller in different slip conditions. In Run 1 the vehicle is in its baseline configuration at 85 % of the max lateral acceleration, so not too close to saturation. Vehicle parameters are changed in run 3 and it is instructive to compare the performance of the controller as vehicle parameters are changed. Run 4, where differences between experimental and simulated data are the largest among runs, is

included to highlight the fact that the controller cannot replicate all driving conditions with the same level of fidelity. Finally, run 6 is considered because tyres are close to saturation.

5.6.1 Run 1

Figure 5.19 shows a snapshot of δ_{sw} time history for run 1. Firstly, it is clear that both time histories show similar statistics in terms of mean and standard deviation. The ZOH behaviour is present in both signals, confirming that the mixed SBH and ZOH approach helps mimic the behaviour of the driver. The experimental and simulated signal velocity, $\dot{\delta}_{sw}$ can have opposite sign, for instance at 1600 m. Since T_u is measured, a better match could be expected. It is reasonable to expect a difference between the driver and controller action for every optimisation. As the simulation progresses, differences in control actions will cause the states of the vehicle to be different, resulting in different time histories. The statistics, however, are unaffected because the mean and standard deviation of the signal measure the overall compensatory action rather than single instances.

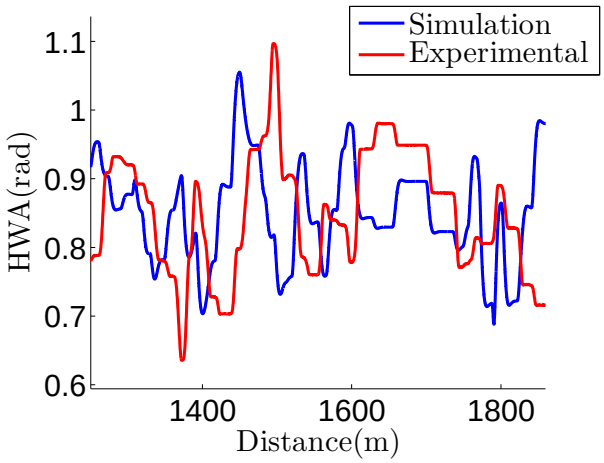

Figure 5.19: Experimental and simulated HWA time history for Run 1. The mean is not subtracted to provide a more informative figure.

Figure 5.20 shows the δ_{sw} MSSD. The match here is excellent; the power distribution over the frequency spectrum is very similar, showing that both mean and standard deviation are close across the whole frequency spectrum. The only noticeable difference is a "bump" at very low frequency, around 0.1 Hz, indicating a control action that happens over 10 seconds – approximately over a whole lap. This behaviour is present in most, but not all, runs. The controller does not account for any control action increase

at such low frequency. This difference is likely caused by a minor change in the mean signal over the laps, which is thought to be the result of the change in cognitive effort the driver devotes to the task. The human attention capacity is limited [139] and the controller does not account for any cognitive deterioration over time, which is likely to happen especially for such a repetitive task. This behaviour should be investigated further to confirm the claim.

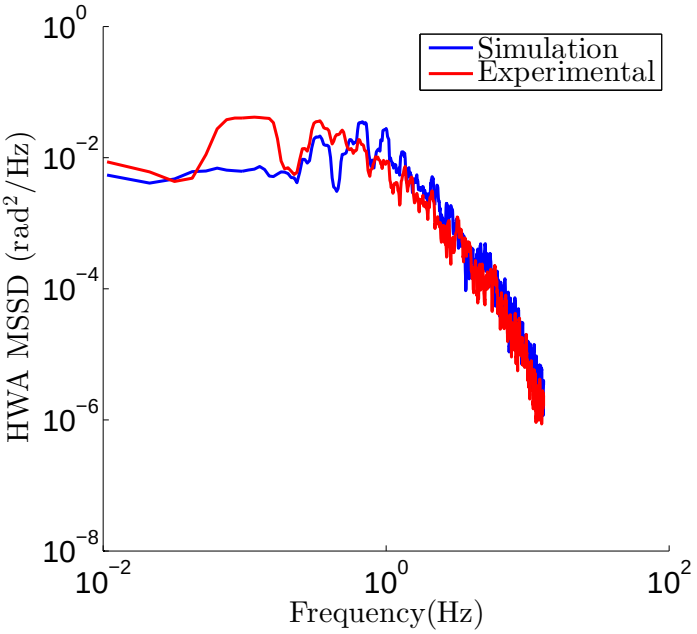

Figure 5.20: Experimental and simulated HWA MSSD for Run 1

Finally, Figure 5.21 shows the pdf of the LPE. Since, the LPE is a combination of two vehicle states, v and ξ, the difference between the experimental and simulated signals is amplified. Despite the differences, the match is satisfactory as the mean and standard deviation are comparable. As such, given that the intent is to replicate the behaviour of a human, the lateral path error match can be considered satisfactory.

5.6.2 Run 3

In this run the CoG is moved 2 % forward.

Figure 5.22 shows the time history of δ_{sw}. The mean values of this run, 0.89 rad and 0.90 rad for experimental and simulated respectively, are slightly larger than the means of run 1, 0.84 rad and 0.85 rad for simulated and experimental respectively. As the CoG is moved forward, under-steering is increased hence a larger steering angle is necessary. This effect is mitigated by the large aerodynamic forces that decrease the

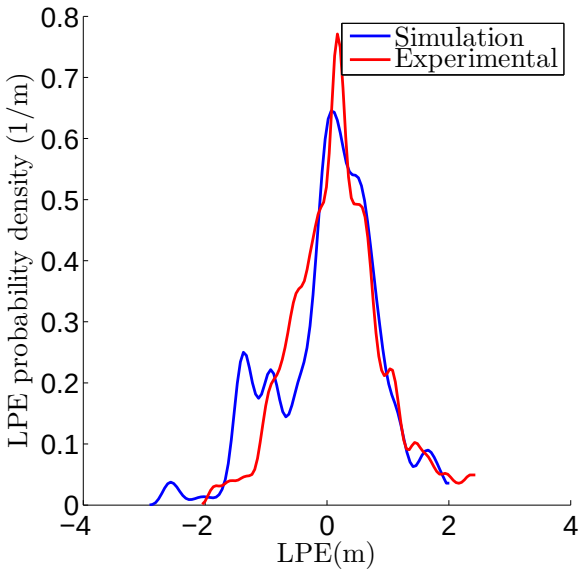

Figure 5.21: Experimental and simulated LPE pdf for Run 1

effect of the change in CoG on the tyre vertical load. After the test, the driver also mentioned that, though he noticed a difference, it was not major. The time history again shows a good agreement between the experimental and the simulated results.

The frequency response – shown in Figure 5.23 again shows an excellent agreement, especially at higher frequencies. The frequency increase around 0.1 Hz is still present.

The LPE, Figure 5.21 shows a very similar mean but a slightly different standard deviation. The experimental results are clearly not symmetric, while the simulation is. This was considered in Section 5.3 where it was observed that it is easier to straighten the vehicle rather than turning it in. While the controller does not require additional "resources" to control the vehicle as tyres are saturating, the driver will try to reduce time spent at the limit to avoid mistakes and reduce the cognitive load required. Furthermore, limb lateralisation [140] could also introduce a further bias in the LPE. The difference in standard deviations is also attributable to discrepancies between the simulation and experimental vehicle models and the noise in the driver's NMS.

5.6.3 Run 4

The vehicle inertia is increased by 20 %.

It is immediately clear from Figure 5.25 that the driving strategy is different, with the simulation requiring much larger corrective actions to achieve stability for a section of the time history, between 1500 m and 1800 m. This is also reflected in the δ_{sw} MSSD, shown in Figure 5.29, where the frequency content of the simulated signal is

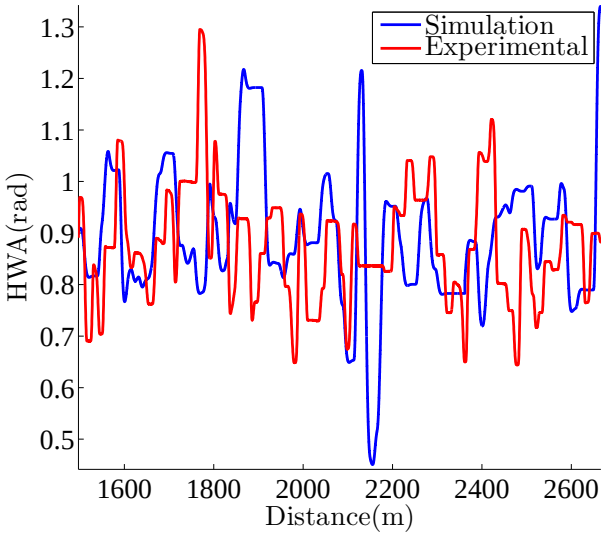

Figure 5.22: Experimental and simulated HWA time history for Run 3. This run was 6 km long so 3 km of data was used for identification and 3 km for data comparison. The mean is not subtracted to provide a more informative figure.

higher. The simulated LPE also has as larger standard deviation, indicating that path tracking capabilities are reduced.

The cause of the discrepancy is attributable to the increase in steering activity between 1500 and 1800 m, rather than to a steady and constant mismatch. This excessive corrective action, which also leads to an increase in the LPE, is caused by 4 σ disturbance realisation. While the driver does not react promptly, allowing the heading error to increase to avoid a destabilising action, the controller tries to correct it, increasing the LPE and the steering activity. This is an example where the cost function weights should be varied; with a lower q_2, the response would be more similar to the driver. This development can be considered in further work.

5.6.4 Run 6

Run 6 is the most challenging to reproduce because the vehicle is closer to the limit. Even though the maximum acceleration is only increased by 5 %, the nonlinear nature of the tyre model makes the increase significant. The δ_{sw} time history shown in Figure 5.28 indicates that the match is not as good as for the other runs.

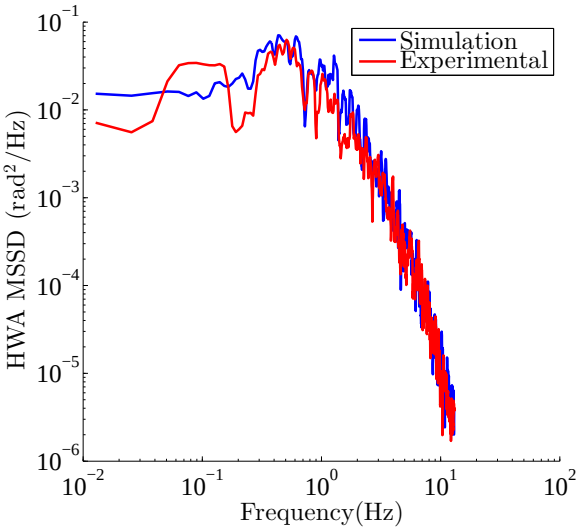

Figure 5.23: Experimental and simulated HWA MSSD for Run 3

This is confirmed by Figure 5.29, which shows that the frequency responses do not overlap as well as before. However, given that the intent is modelling a human driver, the match can still be considered more than satisfactory.

Figure 5.30 shows that the LPE pdf match is acceptable, but the difference is marked. The controller outperforms the driver considerably, with the latter allowing for LPE of up to 4 m. As the vehicle approaches the limit and tracking becomes more erratic, the difference between the experimental and simulated LPE increases. As the vehicle approaches the handling limit, the degree of nonlinearity increases significantly; differences in the vehicle and driver model are therefore amplified. Furthermore, the tyre model around saturation is only approximated due to the lack of experimental data points; this introduces further error.

5.7 Summary

This Chapter validated the lateral controller by comparing its output to data from a driver-in-the-loop driving simulator. Analysing experimental data has shown that the SBH assumption is not the best approach for this scenario, so a mixed ZOH and SBH driving strategy has been proposed. Driver and tyre parameters have been chosen to make experimental and simulation conditions comparable.

Comparing the statistics of the simulated and experimental signals has shown that the controller can mimic the controls of a real human driver to a good degree of accuracy. As vehicle limits are approached, biological biases and differences in vehicle

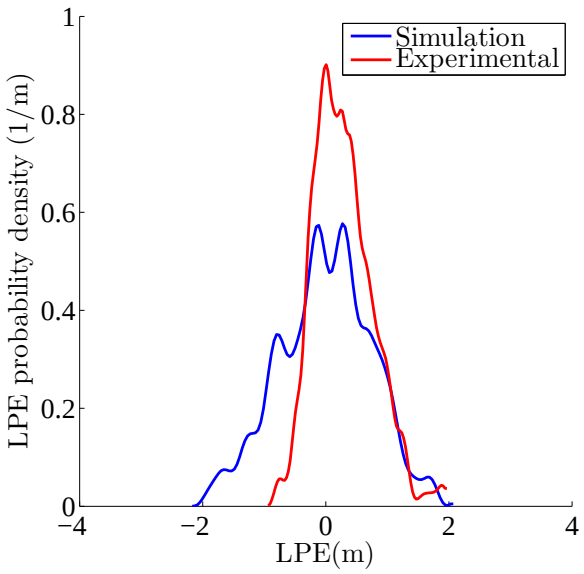

Figure 5.24: Experimental and simulated LPE pdf for Run 3

models become more significant. The fact that the difference between experimental and simulated LPE was more marked than the difference between simulated and experimental steering input shows that the tyre and vehicle model differences contributed to the overall error. Other discrepancies, such as the low frequency increase in the HWA MSSD and the difference in mean LPE, have been justified, but further work is needed to corroborate the claims made.

The primary original intellectual contribution in this Chapter is the formulation and validation of a novel driving strategy that includes both SBH and ZOH. The validation of the lateral controller under the action of disturbances is also a relevant contribution.

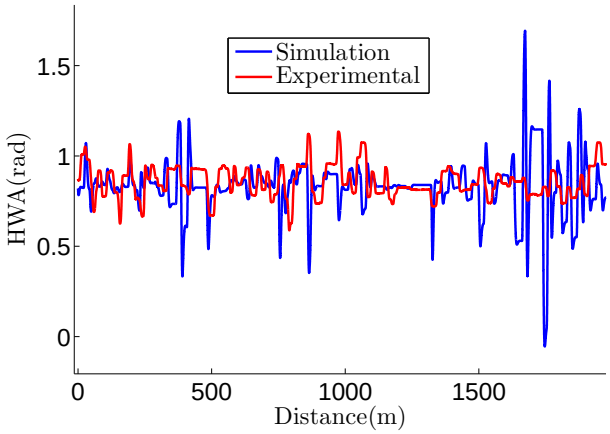

Figure 5.25: Experimental and simulated HWA time history for Run 4. This specific run was 6 km long so 3 km of data was used for identification and 3 km for data comparison. The mean is not subtracted to provide a more informative figure.

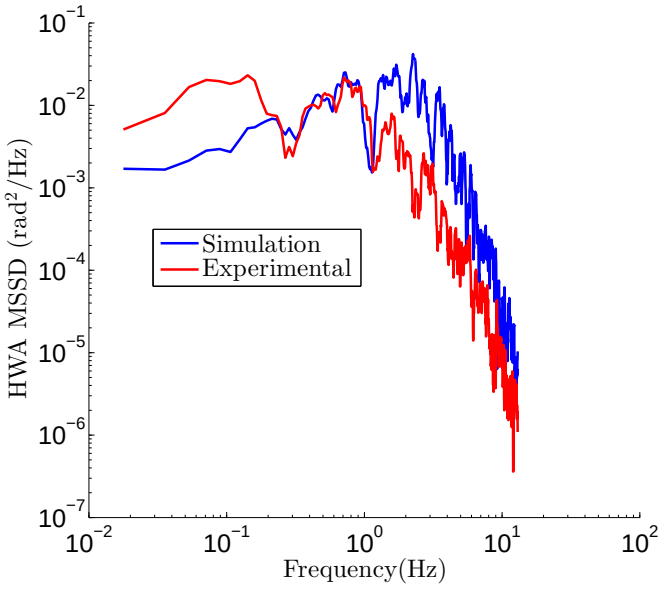

Figure 5.26: Experimental and simulated HWA MSSD for Run 4

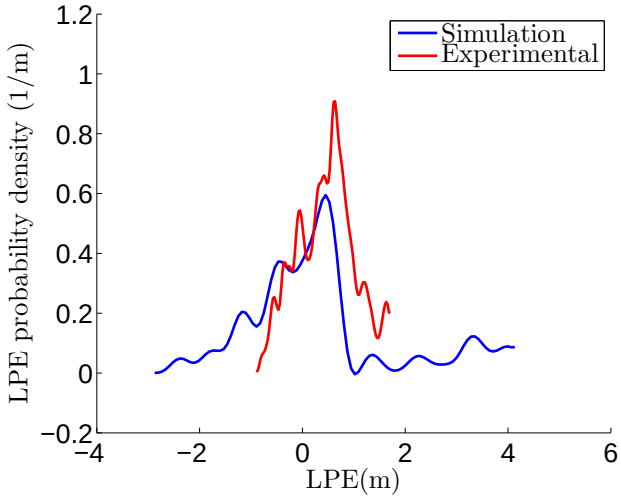

Figure 5.27: Experimental and simulated LPE pdf for Run 4

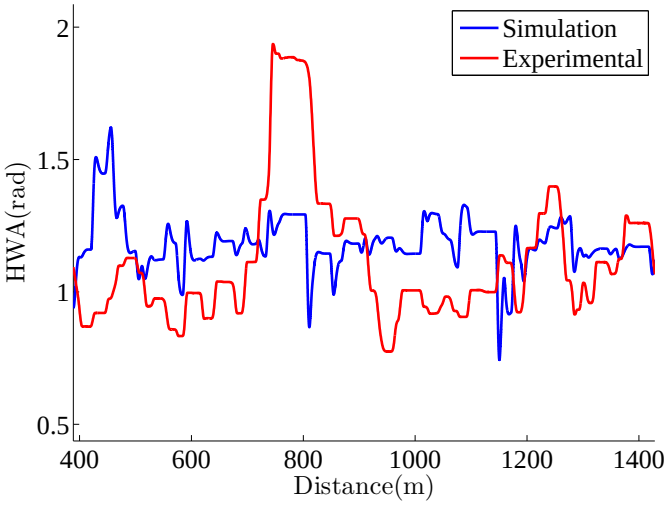

Figure 5.28: Experimental and simulated HWA time history for Run 6

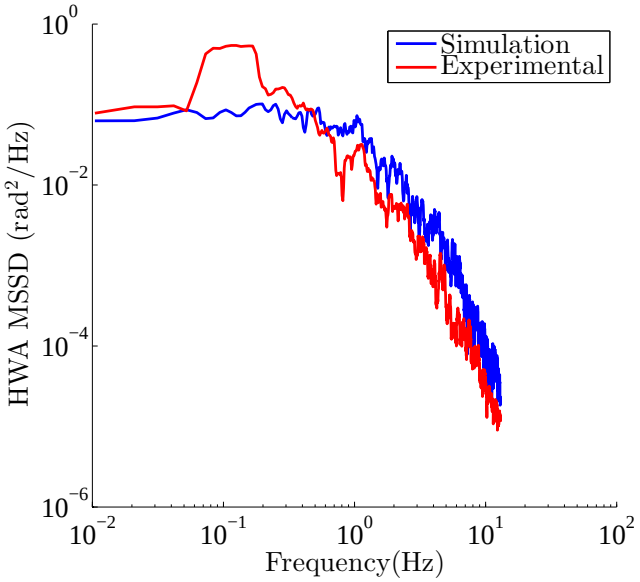

Figure 5.29: Experimental and simulated HWA MSSD for Run 6. The mean is not subtracted to provide a more informative figure.

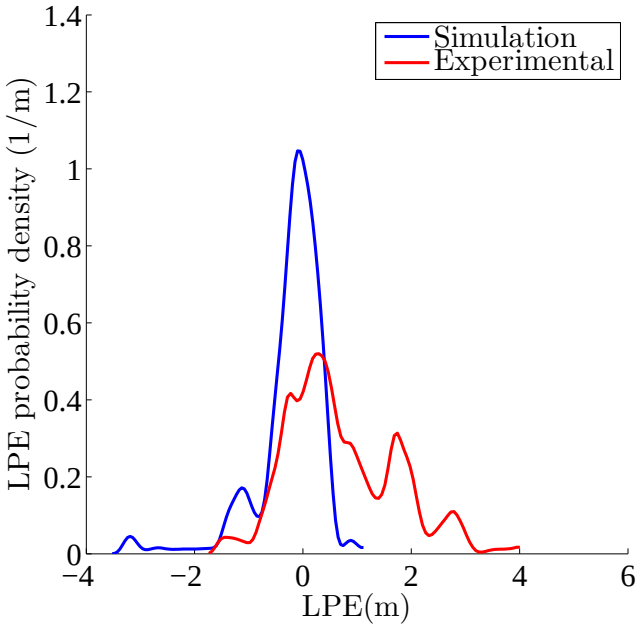

Figure 5.30: Experimental and simulated LPE pdf for Run 6

Chapter 6

Combined controller validation

6.1 Introduction

The lateral controller was validated in chapter 5, where it was shown that it can replicate the behaviour of a racing driver to a good degree of accuracy. The experimental conditions were limited to a constant curvature manoeuvre, which is not representative of actual driving. This chapter extends the validation to the combined controller, using data from a simulator practice session on an international racetrack as a benchmark.

The driver completed 30 laps on the Circuit de Catalunya racetrack on the same high fidelity simulator described in Section 5.2. Experimental data analysis confirms a number of driving behaviours, such as Zero Order Holds (ZOHs), observed in the constant curvature experiment in Section 5.3.2. The tyre and driver model parameters are identified from the experimental data. Simulated and experimental signals are then compared for a manoeuvre where tyres saturate both laterally and longitudinally.

Section 6.2 analyses the experimental data while Section 6.3 details the process used to fit the tyre and driver models and Section 6.4 the disturbance modelling the sensorimotor noise. Sections 6.5 describes the comparison between the simulated and experimental data.

6.2 Experimental data

Data is collected as a professional driver laps around the Montmelo racetrack, shown in Figure 6.1 [141]. While in the constant curvature experiment an external lateral random force was acting on the vehicle, external disturbances are not considered in this test. The main advantage of this approach is isolating the effect of the driver

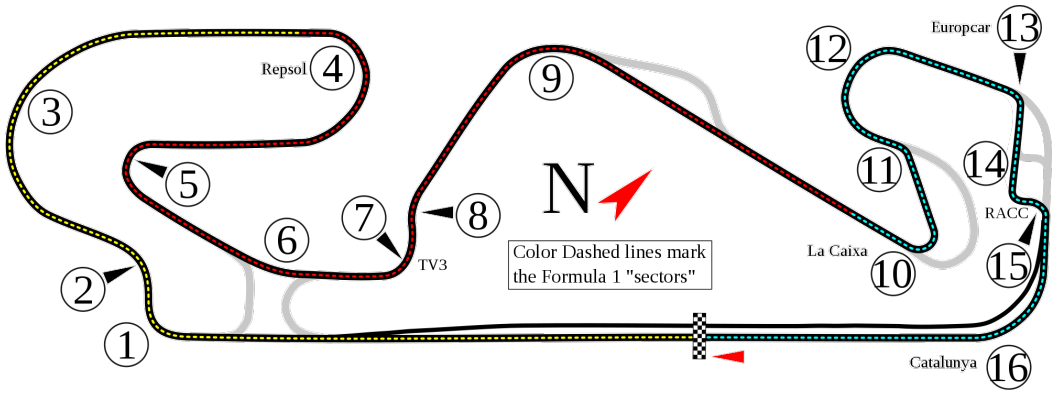

Figure 6.1: Montmelo circuit in the configuration used for the experiments. The same numbering convention will be used throughout the chapter.

intrinsic noise. The vehicle is in its baseline configuration and it is described in Section 5.4.1. Driver mistakes and spins are excluded from the analysis. The time histories, despite the absence of disturbances, display a varying degree or randomness caused by the noise in the driver sensorimotor system.

Data is sampled at 200 Hz, so signals are not equally spaced along track length because velocity varies. The model independent variable is distance and not time, therefore the distance domain must be equispace. This is achieved by finding the intersection between the vehicle trajectory and the normal to the centreline at 0.1 m increments. Interpolation is used to find the exact value of the signal, thus ensuring consistency in space for all signals.

The Lateral Path Error (LPE) is not provided as it can be defined in a number of ways. For the purpose of this validation, the LPE is defined as the minimum perpendicular distance between the position of the vehicle and the nominal line, the same definition used in Section 3.3.2. Instead of simply calculating the magnitude, this approach records which side of the nominal line the vehicle is. The sign convention used is consistent with the one described in Section 3.3.2.

The signals lap to lap variation is of interest as it gives insight into driver randomness; the standard deviation at a given position on the track is used to quantify it. For any signal $S_{i,j}$, where $_i$ indicates the position on the track and j the lap number, the mean at position i will be

$$\mu_i = \frac{1}{L_t} \sum_{j=1}^{L_t} S_{i,j} \tag{6.1}$$

where L_t is the total number of laps. The standard deviation can then be calculated as

$$\sigma_i = \sqrt{\frac{1}{L_t}\sum_{j=1}^{L_p}(S_{i,j} - \mu_i)^2} \tag{6.2}$$

Figure 6.2 shows δ_{sw} time histories for all 30 laps and their average; several features, such as holds and high frequency actions, are lost in the averaging process. The absolute value of the δ_{sw} reaches 3 rad, which corresponds to about half a turn of the Hand Wheel Angle.

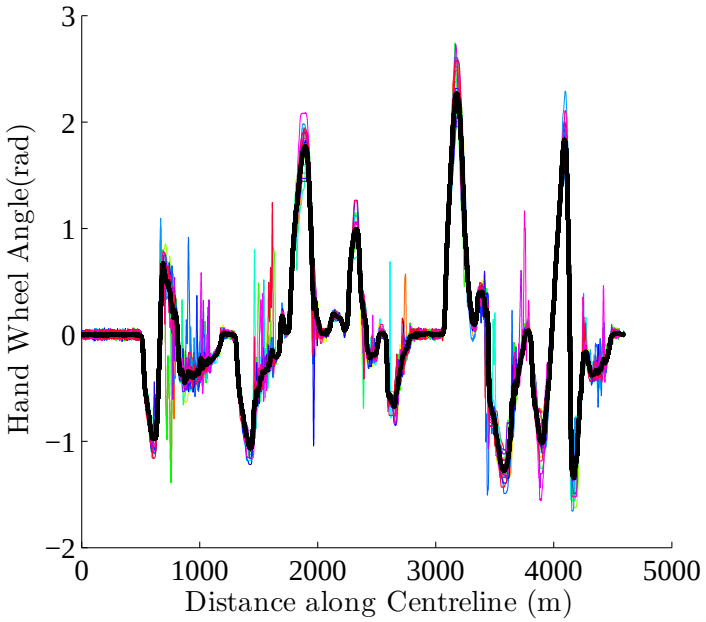

Figure 6.2: Hand Wheel Angle time histories from the driving simulator. Thirty laps are recorded. The thick black line shows the average.

The signal standard deviation is considered in Figure 6.3. Figure 6.4 shows that the standard deviation is significant in sections of the track characterised by high curvature, which correspond to high levels of tyre saturation. The increase in standard deviation in areas of high curvature, and hence higher steering activity, suggests that the noise in the driver NMS system is roughly proportional to the absolute value of δ_{sw}. Even though the data presented does not provide enough evidence to infer a causal relationship, this observation is relevant for the data comparison in Section 6.5.1. While it is not clear for the averaged signal, there is strong evidence of hold behaviour as $\dot{\delta}$ is consistently close to zero. Section 6.2.1 further analyses the hold behaviour.

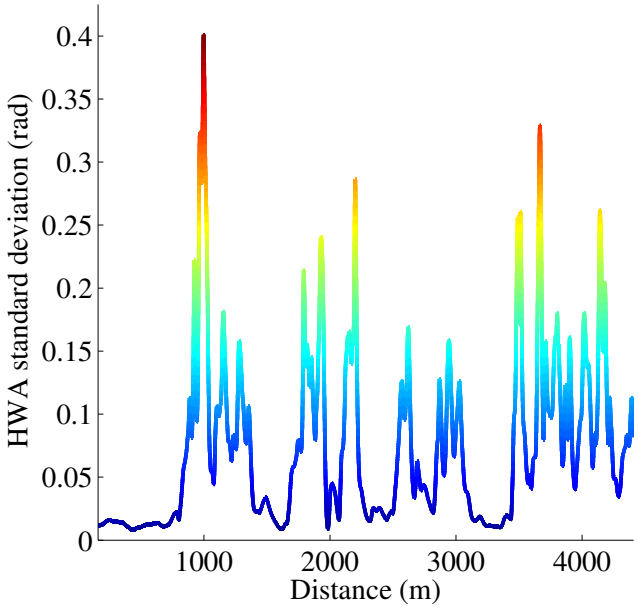

Figure 6.3: HWA standard deviation against distance travelled. The magnitude is approximately 0.01 rad on straight sections, to increase up to 0.4 rad.

The HWA standard deviation is related, yet does not translate directly, to the standard deviation in LPE. The LPE time histories are shown in Figure 6.5. Firstly, the LPE is not zero on the straight, implying that the driver does not always follow the same line. This behaviour is counter-intuitive, as the controller would manage to drive the LPE to zero on a straight line in absence of disturbances. This discrepancy will

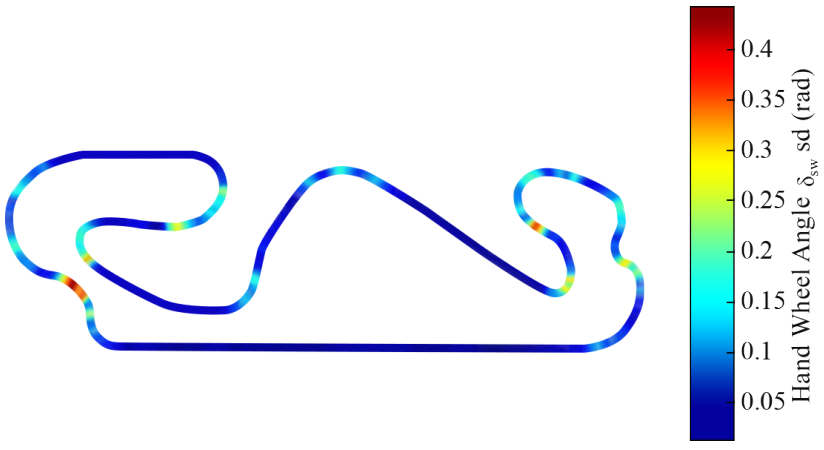

Figure 6.4: HWA standard deviation along the track. The value is generally highest as the vehicle is navigating corners with tyres close to saturation. The evidence provided, however, is not enough to determine a causal relationship.

be further discussed in Chapter 7. Figures 6.6 and 6.7 show that the LPE is highest either midcorner or at corner exit. These driving scenarios are again characterised by significant steering activity and tyre saturation. Finally, a marked increase in standard deviation is apparent in turn 3. The corner is long and characterised by decreasing curvature, which allows drivers to choose different trajectories, trading speed for distance travelled. Furthermore, it follows from corners 1 and 2, so any difference in exit position and speed affects the driving strategy significantly.

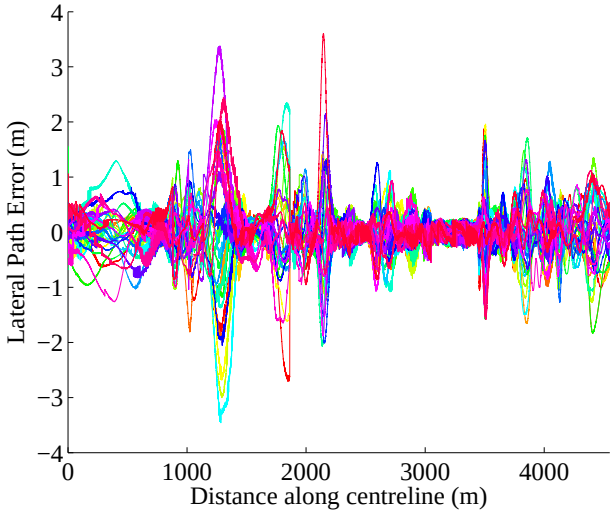

Figure 6.5: Lateral path error time histories ensemble for 30 laps from the driving simulator. Large standard deviations are evident in certain parts of the track.

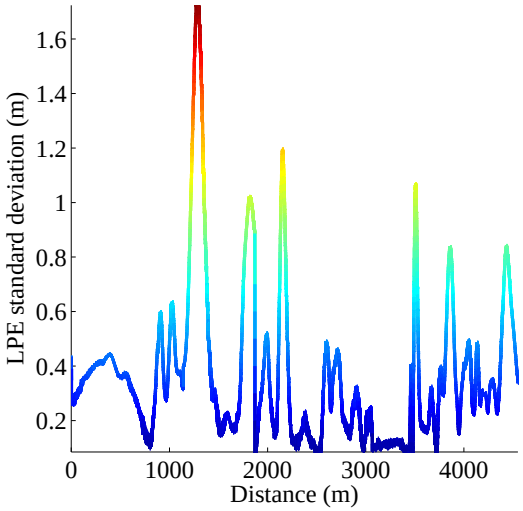

Figure 6.6: Lateral Path Error standard deviation against distance travelled. The highest value of 1.65 m corresponds to corner 3 where the curvature is low.

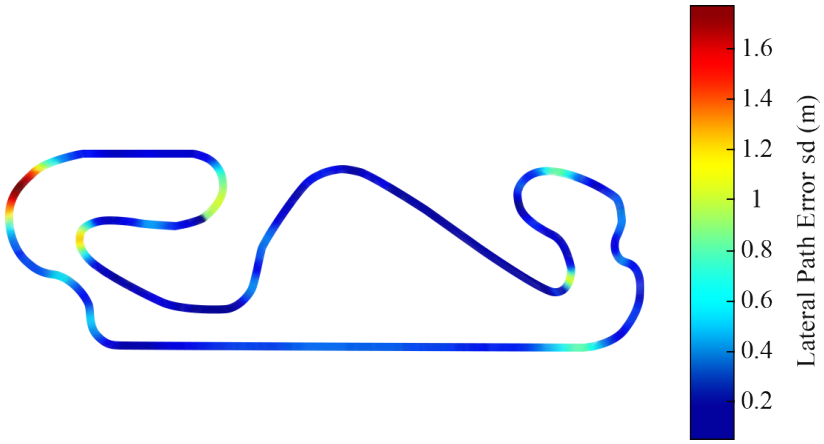

Figure 6.7: Lateral path error standard deviation along the track. Corner entries are characterised by a low standard deviation while it increases from mid-corner until the end of corner. One would usually expect the magnitude to be lowest at the apex. However, given the high power of the vehicle, position accuracy is sacrificed for a faster exit.

Gerdes *et al.* performed a similar data collection exercise in [142] where they considered the deviation of the vehicle from the nominal line. It is instructive to compare the collected data to Gerdes *et al.* work; even though the vehicles are different, a degree of similarity in the responses is expected. The deviation from path is taken as the absolute value of the local deviation from the median, referred to as MAD, which can be compared to the standard deviation of the LPE. The LPE patterns are

mostly similar. The magnitude ranges from 10 cm to about 1.5 m for both. The most significant difference is in large curvature corners where the MAD is low while the LPE standard deviation is higher. This is particularly noticeable in corner 3 of the Barcelona racetrack. The difference is thought to be due to the nature of the vehicle and the driver skill level.

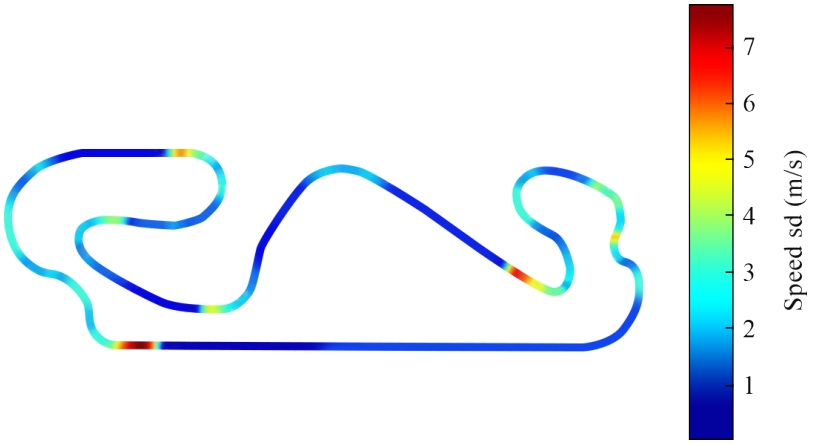

Figure 6.10: Speed standard deviation. The highest values correspond to braking points, where small differences cause a large speed error.

Figure 6.8 shows the speed profile, which is generally consistent across laps. Braking points and corner 3 are the only sections where the signal deviates significantly from

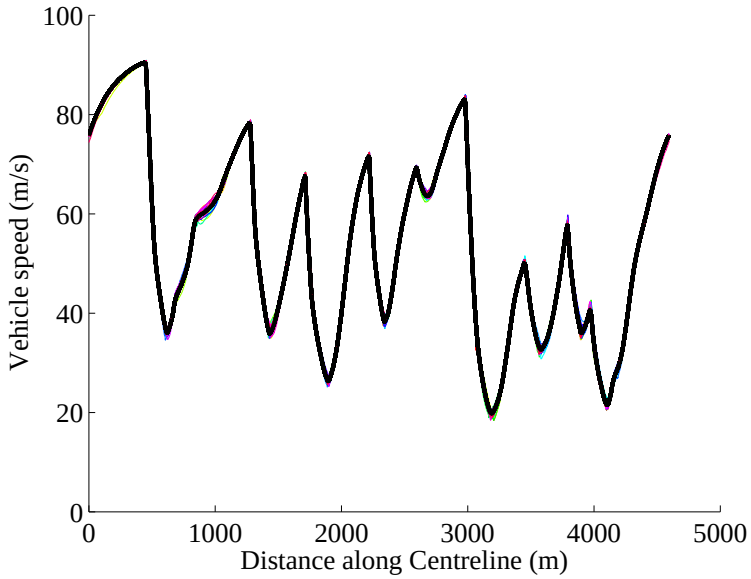

Figure 6.8: Vehicle speed time histories ensemble for 30 laps from the driving simulator. The thick black line shows the average.

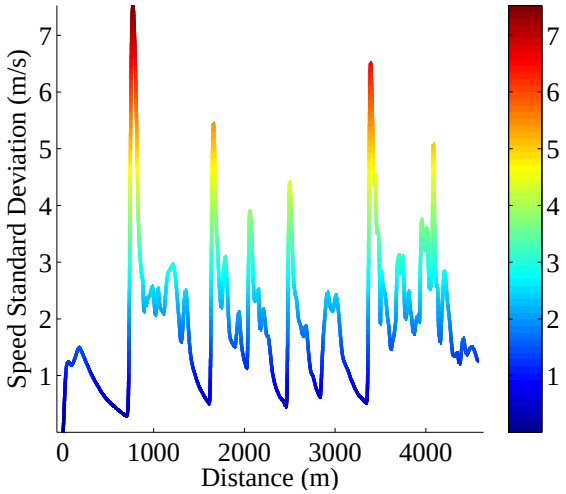

Figure 6.9: Speed standard deviation against distance travelled. The highest value is 7 m/s at the first braking point. Apart from braking points, the standard deviation is always less than 3 m/s.

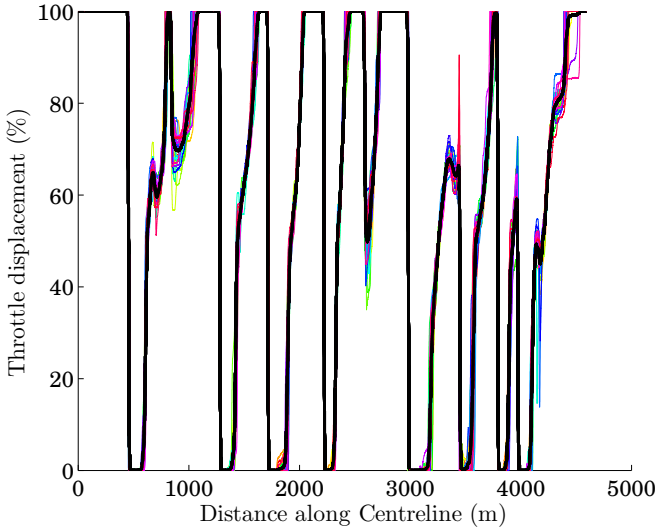

Figure 6.11: Throttle time histories ensemble for 30 laps from the driving simulator. The *on-off* behaviour is clear.

the nominal. Figures 6.9 and 6.10 show the standard deviation for the speed. The highest values are clearly in braking. Interestingly, the standard deviation along corners is independent of curvature. The lowest values of standard deviation are found on the straight. Gerdes *et al.* [142]only reports the velocity profiles for two corners, where the deviations are comparable to the ones in this work.

The throttle time history, shown in Figure 6.11, is consistent. The black line superimposes well on the time histories and the standard deviation is generally small. Figures 6.12 and 6.13 show the rear longitudinal slip standard deviation so as to determine the extent to which the longitudinal controls vary, especially traction forces. Clearly, the highest levels of standard deviation happen at corner exits where the acceleration is higher. Different braking points, shown in Figure 6.14 for corner 1, while resulting in significant speed standard deviation, do not translate to an increase in rear slip standard deviation. Around the braking point, the speed decreases slightly, then increases to finally drop. This behaviour is unexpected; speed signals from [142] do not show this "step", which suggests that it is likely to be a feature of the simulator vehicle model that is not considered by the simulation vehicle model.

The HWA, speed profile, the LPE and the throttle profile provide good insight into the driving strategy, highlighting some expected and unexpected features. Figure 6.2 shows evidence of hold behaviour, so is worth investigating to yield a more comprehensive picture of the driver's behaviour.

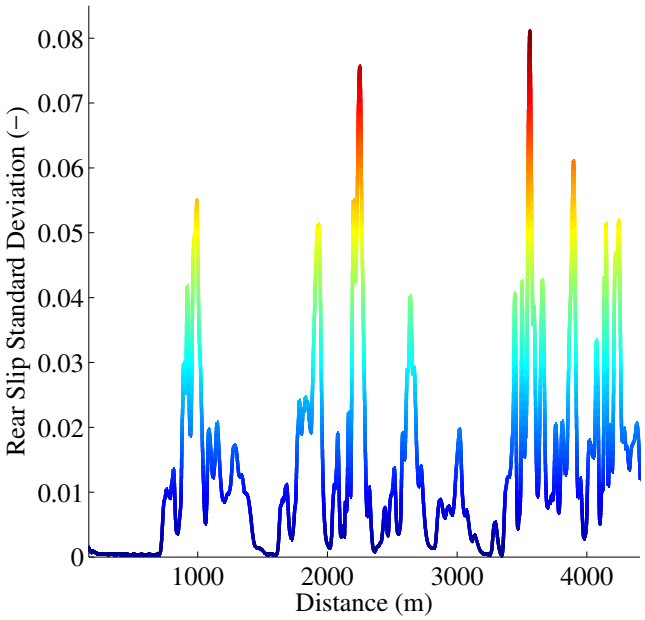

Figure 6.12: Rear slip time history. It is difficult to readily quantify the extent of the variation because slip does not have an intuitive physical meaning. The maximum value is 0.08 while it is close to 0 on the straight sections.

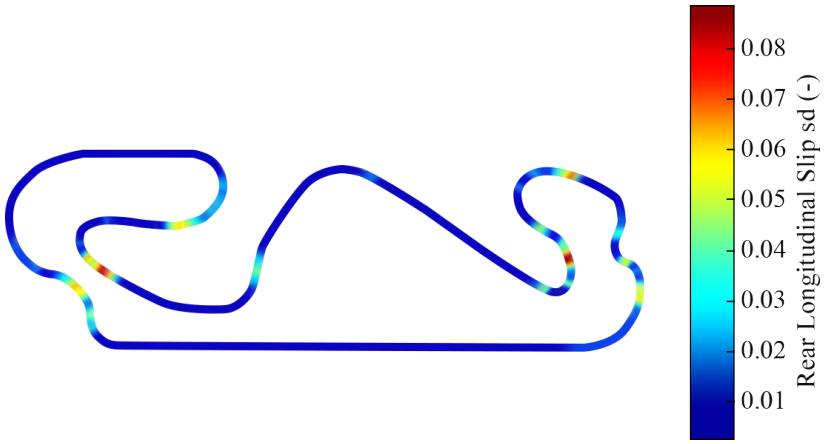

Figure 6.13: Rear longitudinal slip standard deviation along the track. The magnitude generally increases at corner exit, where the longitudinal activity is highest.

6.2.1 Hold strategy

The Hand Wheel Angle for a vehicle navigating a circle at constant speed is constant [7]. Constant δ_{sw}, hence ZOHs, would therefore be expected for corners characterised by constant curvature and constant speed in the absence of disturbance. It is therefore instructive to look at how the holds relate not only to the curvature (Ω_z), but also to the rate of change of curvature ($d\Omega_z/ds$). Holds are found using the algorithm described in Section 5.3.2; the lower bound being set to 0.005 s for a hold to be considered as such.

Figure 6.15 shows the relationship between curvature and hold lengths. Each dot corresponds to an instance of a hold of given length at a given curvature. The same hold can span a range of curvature, hence the short vertical lines. The curvature range over which holds happen is wide. Figure 6.16 confirms that the driver actually makes extensive use of ZOH as holds up to 0.5 s are consistently present over the whole range of $d\Omega/ds$. This result has very important consequences for the data comparison; it suggests that the driver makes extensive use of ZOH in all sections of the track, hence the ZOH+SBH approach can be extended for the combined controller as well.

H_p^{per} and H_p^{pert}, defined in defined in Section 5.3.2, are considered to gain insight into the driving strategy and to compare them to the hold statistics of the constant radius experiment described in Chapter 5. The two straight sections, between corners 16 and 1 and 9 and 10, are excluded because the curvature is 0.

Figure 6.17 shows that H_p^{pert} is constant for hold lengths shorter than 0.1 s, where it starts rolling off. The percentage of lap time where δ_{sw} was held constant varied

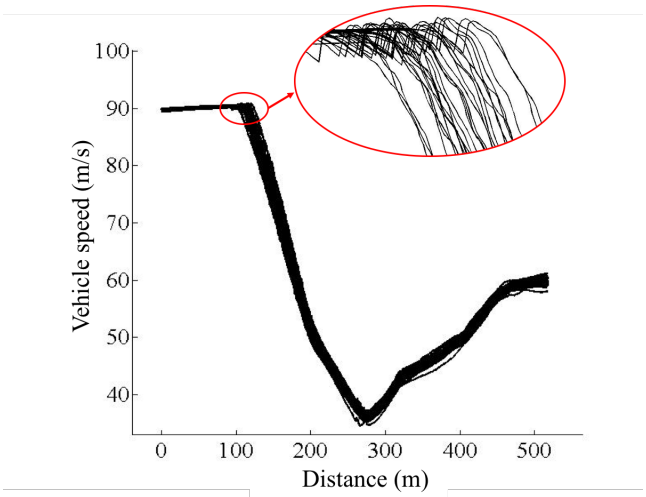

Figure 6.14: Change in speed at Corner 1 braking point. The variation in braking point is highlighted top right.

between 25 % and 35 %, indicating a marked hold strategy lap to lap variation. A clear and consistent roll off is observed around 0.1 s, which suggests that the hold activity is higher for shorter holds. Figure 6.18 shows that, while the hold behaviour is consistent across laps for short holds, differences in percentage of total hold time arise as holds get longer. The figure also indicates that the SBHs are between 65 % and 75 % of the total lap time, with a mean of 70%. For ZOHs up to 0.5 s, hence over the majority of the ZOH domain, T_{ZOH}/T_{SBH} is uniformly distributed with a 0.78 mean, indicating that the SBH holds are longer than the ZOHs. The update time statistics can also be inferred from Figure 6.17; the hold time where the roll off starts, short of 0.1 s, indicates that most ZOH are shorter. This suggests that if $T_{ZOH}/T_{SBH} = 0.78$, the T_u average will be in the region of 0.18 s.

Figure 6.19 highlights the differences in roll off. While it initiates at the same hold length, just short of 0.1s, H_p^{per} slope for the combined controls experiment is higher than the lateral control experiment. This indicates that the former is characterised by shorter holds, again a consequence of varying Ω_z

Figure 6.20 shows the comparison between the mean values of H_p^{pert} for the constant and varying curvature scenarios. It is immediately clear that the driver holds δ_{sw} significantly more, 62 % vs 30% for the constant curvature experiment, which is what one would expect. This confirms that even though ZOHs are found as Ω_z is varying, ZOHs are more likely to be observed for low values of $\mathrm{d}\Omega_z/ds$.

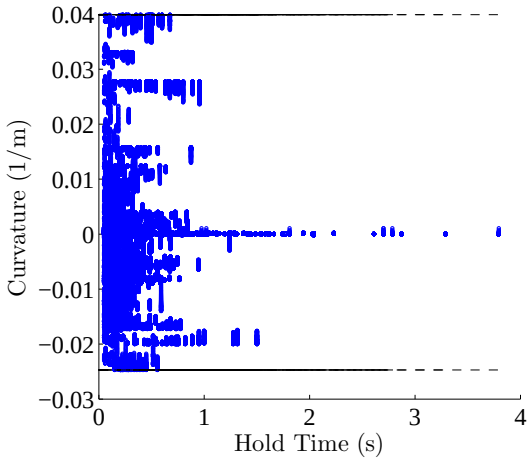

Figure 6.15: Hold times for all 30 laps for various Ω_z along the track. The horizontal lines at 0.04 and 0.025 m^{-1} indicate the maximum and minimum values of $d\Omega/ds$. The focus of the figure is the presence of holds over the curvature range rather than their distribution.

6.3 Model identification

After analysing the experimental data, the simulation model must be characterised.

The vehicle model is the same as the one used in Chapter 5, characterised by the parameters in Table 5.2. However, for the constant speed experiment, drag was not considered and the constant lift forces simply affected the weight distribution of the vehicle. These assumptions are no longer valid for the variable speed experiment; models for drag and lift must therefore be specified.

Tyre model parameters must also be identified; varying vertical forces, lateral and longitudinal components must be taken into account. Values from Table 5.3 are not used because the lateral/longitudinal coupling and aerodynamic effects affect tyres lateral behaviour so the fitting is repeated.

The driver model and the cost function parameters are the final part of the identification process.

6.3.1 Aerodynamic parameters identification

Both lift and drag models are to be identified.

The aerodynamic map was not disclosed, so a model for lift needs to be proposed. The vehicle model used to derive the controller assumes that the lift force is proportional to the square of the velocity. Aerodynamic forces for F1 cars are modelled by complex

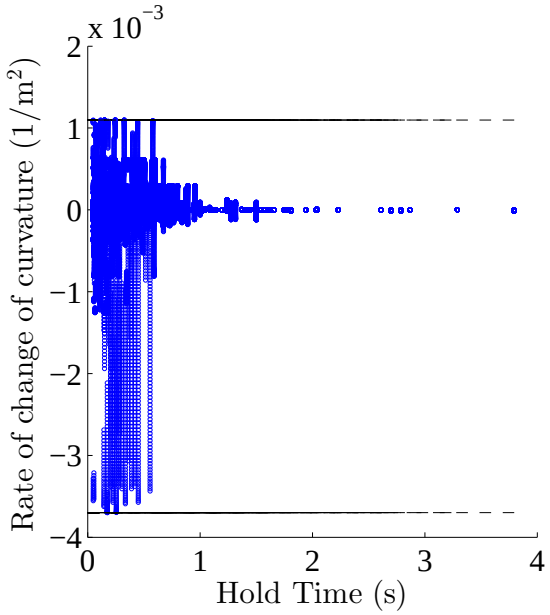

Figure 6.16: Hold times for all 30 laps for various Ω_z/ds levels along the track. The horizontal lines at 1E-3 and -4E-4 $1/\mathrm{m}^2$ indicate the maximum and minimum curvature. The focus of the figure is the presence of holds over $d\Omega/ds$ range rather than their distribution.

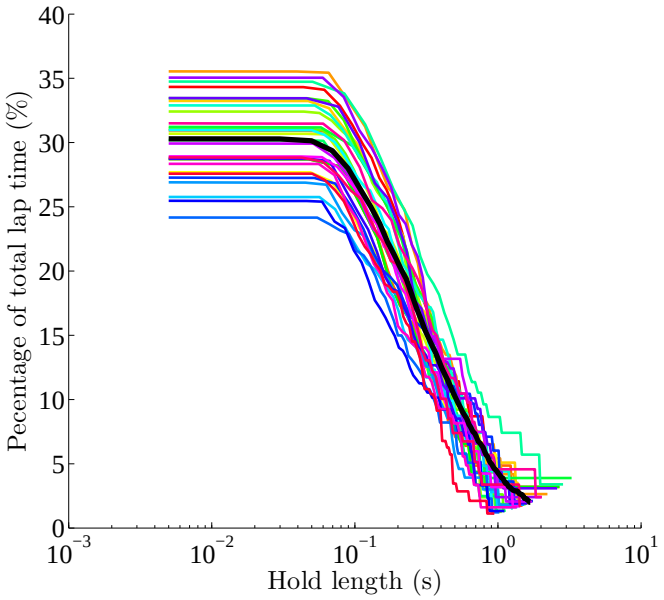

Figure 6.17: H_p^{pert} showing that the total hold time is between 25 % and 35 % of the total lap time, with an average of 30 %. A marked decrease in percentage of total lap time is observed for holds longer than 0.1 s.

empirical aerodynamic maps [143], it is therefore unrealistic to assume that the lift coefficient is constant. A better approximation can be found following Kutz [144], who shows that linear dependence on speed can capture the aerodynamic effects to a satisfactory degree of accuracy. The lift force, assuming it is positive downwards, can therefore be expressed as

$$F_{az} = \frac{1}{2} C_L A_f \rho v^2 \tag{6.3}$$

where all terms have been previously defined and C_L can be expressed as

$$C_L = a_L v + b_L \tag{6.4}$$

where a_L and b_L are constant coefficients. Since the force is assumed to be positive downwards, C_L is positive.

The total vertical force balance is considered so lateral and longitudinal load transfer does not have to be taken into consideration.

$$F_{zt} = Mg + F_{az} \tag{6.5}$$

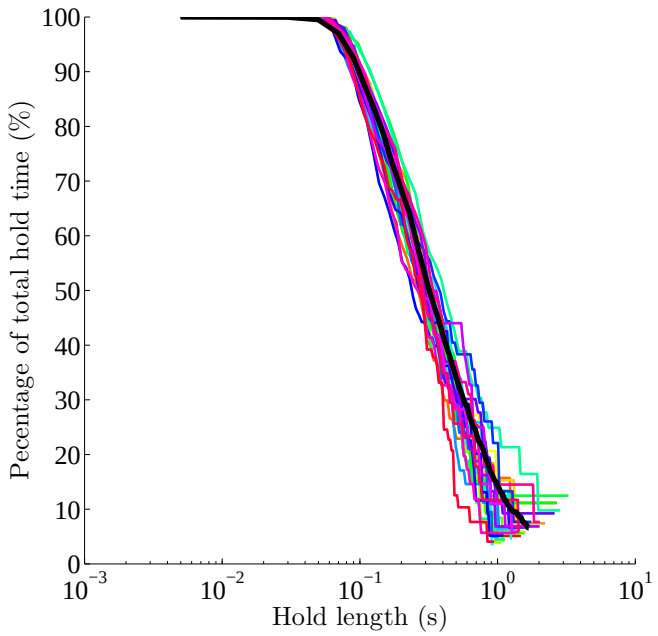

Figure 6.18: H_p^{per} showing a marked decrease in percentage of total lap time is observed for holds longer than 0.1 s. The lower bound for the holds is set to 0.005 s, which corresponds to the simulator sampling rate.

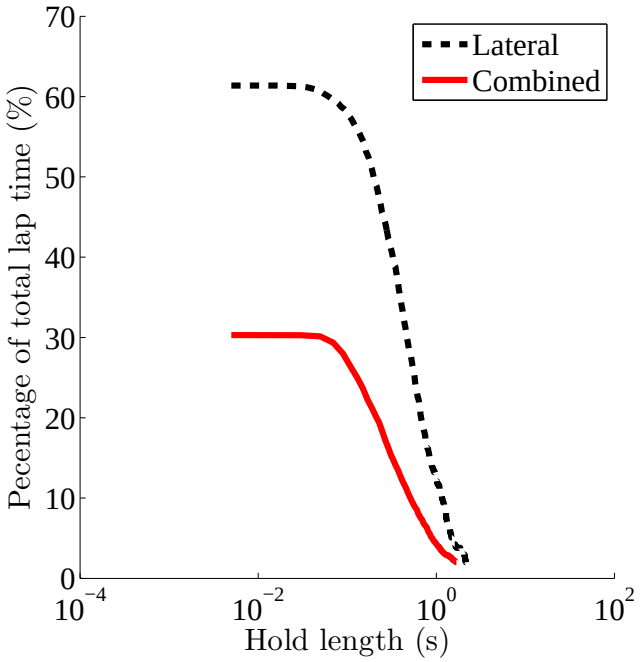

Figure 6.19: H_p^{pert} comparison between the lateral and combined experiment for the whole time showing a marked difference.

where all terms have been previously defined. Since both F_{zt} and F_{az} are a function of velocity, a_L and b_L can be found by minimising the following cost function:

$$J_{lift} = \sum_{k=1}^{k=N_l} (F_{zt(k)}^{sim} - F_{zt(k)}^{exp})^2 \tag{6.6}$$

where $F_{zt(k)}^{sim}$ and $F_{zt(k)}^{exp}$ are the simulated and experimental vertical forces at distance step k and N_l is the total number of distance steps over the 30 laps. Nelder Mead was used to minimise the cost function, obtaining $a_L = -0.033$ and $b_L = 5.026$. Data for $F_{zt(k)}^{exp}$ was provided by the industrial partner.

The drag force also needs to be considered. Drag force is not available so it can only be estimated. Assuming that

$$F_{ax}^{exp} = F_{Lon} - M_t \dot{u} \tag{6.7}$$

where F_{Lon} is the total tyre longitudinal force.

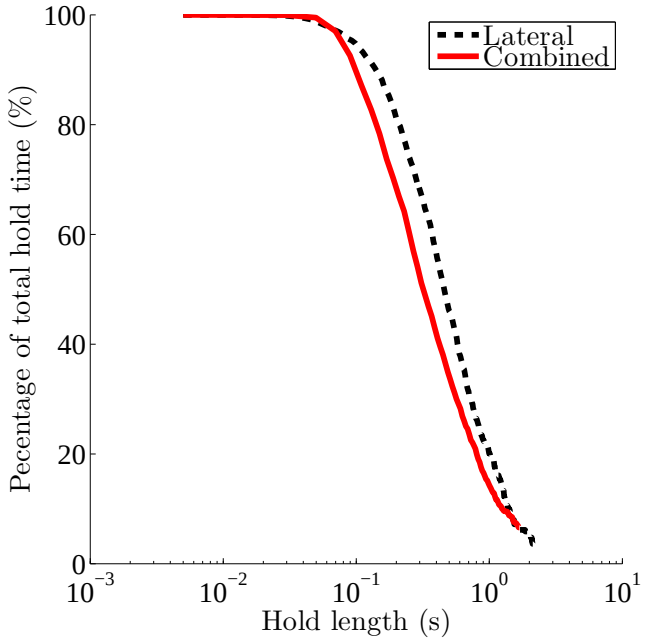

Figure 6.20: H_p^{per} comparison between the lateral and combined experiment.

Assuming the drag force and coefficient are modelled in the same fashion as the lift in (6.3) and (6.4) , the simulated drag force can be obtained by

$$F_{ax}^{sim} = \frac{1}{2} C_x A_f \rho v^2 \tag{6.8}$$

where all terms have been previously defined and C_d can be expressed as

$$C_x = a_x v + b_x \tag{6.9}$$

where a_x and b_x are constant coefficients.

The difference between the experimental and simulated drag force can be minimised to find the a_d and b_d The cost function can be expressed as

$$J_d = \sum_{k=1}^{k=N_l} (F_{(d,k)}^{exp} - F_{(d,k)}^{sim})^2 \tag{6.10}$$

where $F_{(d,k)}^{exp}$ and $F_{(d,k)}^{sim}$ are the experimental and simulated drag force at distance step k and N_l is the total number of distance steps over the 30 laps. Nelder Mead is used to minimise (6.10) to give $a_d = -0.0023$ and $b_d = 1.72$.

6.3.2 Tyre model identification

The tyre model was not disclosed for confidentiality reasons, so parameters need to be chosen to make the simulation and experimental models comparable. The time histories for the slip and tyre forces are provided for the four wheel vehicle. Left and right wheel forces are added to find the total front and rear axle forces, while left and right slips are averaged to find average slips. The procedure follows from the one described in Section 5.4.2. However, fitting the tyre model for a vehicle with combined controls is more challenging than fitting the model for a lateral only controller due to two factors:

- *Varying vertical force*

  Since speed is changing, the effect of the aerodynamic forces is not constant and cannot be taken into account by the tyre model.

- *Combined lateral and longitudinal effect*

  As lateral and longitudinal forces are present, the model is inevitably more complex. While only the magnitude of the force error was minimised in the lateral case, both magnitude and direction need to be considered in the combined case.

The variation in vertical forces is considered in Section 6.3.1. Tyres are characterised by three parameters per direction, namely Q, μ and α_{max} or κ_{max}. There are therefore a total number of six parameters to be fitted per axle.

The cost function that needs to include both longitudinal and lateral forces to optimise for the magnitude as well as for the direction of the force. It can be expressed as

$$J_{tl} = \sum_{k=1}^{k=N_l} |\boldsymbol{F}_{(k)}^{exp} - \boldsymbol{F}_{(k)}^{sim}|^2 \tag{6.11}$$

where $\boldsymbol{F}_{(k)}^{exp}$ and $\boldsymbol{F}_{(k)}^{sim}$ are the horizontal tyre force vectors at distance step k and N_l is the total number of distance steps over the 30 laps.

The procedure is applied to both front and rear tyres separately. The resulting parameters are reported in Table 6.1.

Figure 6.21 shows that the force magnitude and direction of the front tyre are a close match, with an average error of only 2 % for the force and 6 % for the direction. Only one lap is reported for clarity but the full set of 30 laps has been used to generate the data in Table 6.1.

Table 6.1: Tyre model coefficients fit for front and rear tyres.

Tyre	μ_y	Q_y	α_{max}	μ_x	Q_x	κ_{max}
Front	1.7909	1.7061	0.1564	1.4554	1.5922	0.1419
Rear	1.7406	1.5770	0.1795	1.6303	1.6077	0.1364

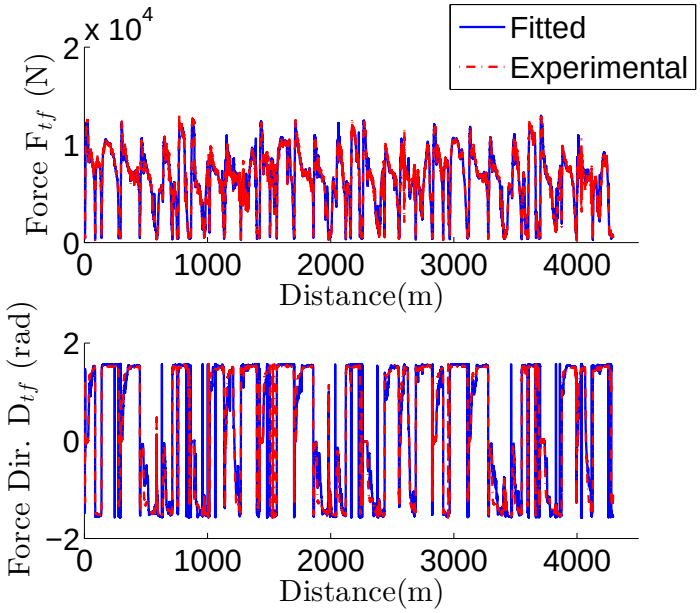

Figure 6.21: Comparison between the total front tyre force magnitude and direction for the fitted and experimental results.

Figure 6.22 shows the contour plot of the slip-force surface. The advantage of plotting the contour instead of the surface is the fact that it is constant over the whole range of vehicle speeds as aerodynamic forces simply scale the surface vertically. This implies that the projection of the force slip characteristics on the $\alpha - \kappa$ plane gives an immediate graphical representation of the tyre slip characteristics for the lap and for the whole speed range. Aerodynamic forces only scale the tyre curve, limiting their effect to the $F - \alpha$ and $F - \kappa$ planes. The shape of tyre curve is asymmetric, which shows that the experimental tyre model has different characteristics in the lateral and longitudinal direction. Load transfer effects, which are not accounted for in the bicycle model, are captured by the tyre model and they could contribute to tyre asymmetry. Red lines show the measured slip, which exceeds the 0.15 radius friction circle several times both laterally and longitudinally.

Figures 6.22 show that the magnitude and direction of the rear force also match closely.

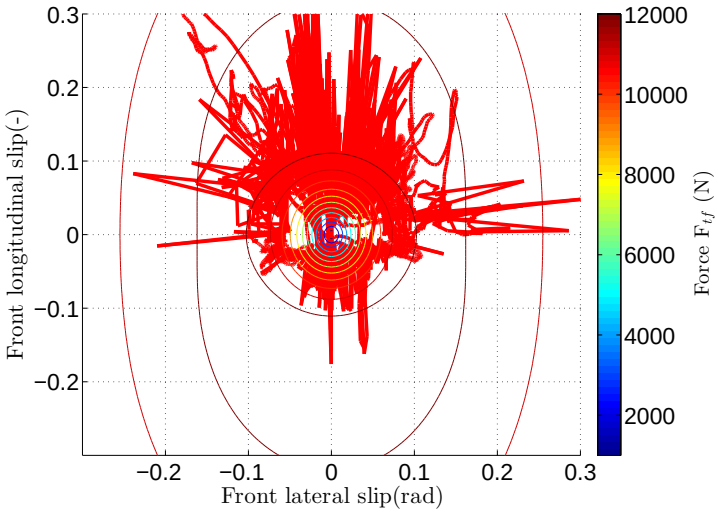

Figure 6.22: Fitted front tyre curve and slip force characteristics. The front tyre curve is plotted for 40 m/s. This, however, cannot be inferred from the figure as only the $\alpha - \kappa$ plane is showed.

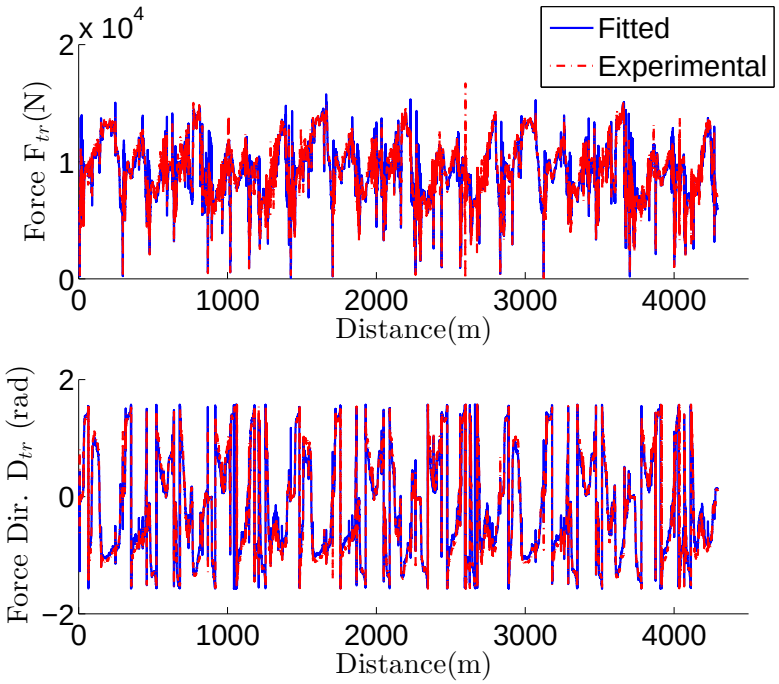

Figure 6.23: Comparison between the total rear tyre force magnitude and direction for the fitted and experimental results.

6.3.3 Driver and cost function parameters identification

Tyre parameters have been successfully characterised, so driver model parameters should now be considered. Following the identification procedure in Section 5.5, the difference between experimental and simulated δ_{sw} MSSDs could be minimised. While comparing the simulated output with the experimental one is still a valid strategy, external disturbances are not acting on the vehicle so the lap to lap variation is due to driver internal disturbances, which are unknown. Given the number of variables in the model and its complexity, trying to identify all the parameters with an unknown disturbance in the system would be computationally too expensive. An alternative approach needs to be considered.

The identification process is simplified by neglecting the system randomness. Rather than comparing the disturbed signals, the experimental averaged signals and the nominal, undisturbed simulated signals are compared. This assumption does not introduce a significant error for two reasons:

- *Lap to lap variation magnitude*

  The lap to lap variation is small, typically around 5 %, of the actual signal magnitude. While it is significant for the data comparison, it can be neglected in the parameter identification.

- *Driver intrinsic randomness*

  Cost function weights attempt to model driver behaviour. In this work, they are assumed to be constant. However, it is clear from the experimental results in Section 6.2 that they are not only variable, but they are also characterised a certain degree of randomness. Cost function weights are likely to vary with track conditions, such as curvature and speed, but also lap to lap due to the driver intrinsic randomness. A similar argument can be made for the NMS filter: both the natural frequency and damping attempt to model an intrinsically random system. Since the deterministic, constant driver model only approximates the random nature of the human, an estimate obtained from comparing nominal signals is sensible. The approach to include driver randomness into the system is discussed in 6.5.1. Systematic disturbances, such as road roughness at same point on the track, have not been taken into consideration because they were not present in the experimental model.

Before specifying the cost function and detailing the identification, a number of assumptions is made to reduce the search space of the algorithm. Given that the

same driver performed both the constant speed and variable speed experiments, the lateral dynamics NMS filter parameters are assumed to be the average of those found in Section 5.5. The longitudinal filter dynamics are assumed to the same as the lateral as the muscles and the nervous system are the same. The effect of the filters is affected by the cost function weights q_3, q_5 and q_6 so these should account for any discrepancy that arises from the aforementioned assumptions.

The update time is taken from measured data. As per Section 5.3.2, the signal is assumed to be a sequence of holds which are composed of ZOHs and SBHs, the sum of the two being the update time. These are measured using the algorithm described in Section 5.3.2.

Finally, the cost function weight q_7 is assumed to be 1E-4 because several tests performed by the author showed that this value ensures that the driver does not apply accelerating and braking torque simultaneously on the straight for any cost function parameter space used in this work.

The cost function for the identification can now be set up. Parameters need to be adjusted to track the nominal speed and path, while replicating the behaviour of the driver. All controls are included in the cost function, which can be expressed as

$$J_{sl} = (\boldsymbol{\delta}_{sw,sim} - \boldsymbol{\delta}_{sw,exp})^T(\boldsymbol{\delta}_{sw,sim} - \boldsymbol{\delta}_{sw,exp}) \tag{6.12}$$

$$+(\boldsymbol{\kappa}_{f,sim} - \boldsymbol{\kappa}_{f,exp})^T(\boldsymbol{\kappa}_{f,sim} - \boldsymbol{\kappa}_{f,exp}) \tag{6.13}$$

$$+(\boldsymbol{\kappa}_{r,sim} - \boldsymbol{\kappa}_{r,exp})^T(\boldsymbol{\kappa}_{r,sim} - \boldsymbol{\kappa}_{r,exp}) \tag{6.14}$$

J_{sl} is minimised using the Nelder Mead algorithm. The parameter studies performed in Chapter 4 and the validation in Chapter 5 provide a good baseline for an initial guess.

Since disturbances are not present, optimality is guaranteed over the whole prediction horizon. An arbitrary long update time can therefore be used, decreasing computational time considerably. The resulting values are shown in Table 6.2, specifying the values that have been identified and those assumed from the previous experiment or measured. The average states and controls are calculated using Eq (6.1).

Figure 6.24 shows the comparison between the experimental and the simulated δ_{sw}. The match is excellent with the biggest difference being at δ_{sw} peaks, hence at the limit of adhesions. The metric used for comparison is the relative error, calculated using Eq 5.10. The maximum relative error is around 20%, which happens at 2200 m into the simulation. The average error is much lower, around 3% for the whole lap. Neglecting δ_{sw} peaks at 2100 m and 3500 m, a t-test shows that the difference between

Table 6.2: Parameters for the combined validation

Parameter	Symbol	Value	Derivation
Steering NMS damping ratio	ζ_n	0.68	Assumed from Ch. 5
Steering NMS natural frequency	ω_n	15.47 rad/s	Assumed from Ch. 5
Penalty on deviation from nominal path	q_1	0.97	Identified
Penalty on deviation from nominal heading	q_2	1.03	Identified
Penalty on hand wheel angle speed	q_3	0.0101	Identified
Penalty on deviation from nominal speed	q_4	0.504	Identified
Penalty on front longitudinal slip speed	q_5	0.0015	Identified
Penalty on rear longitudinal slip speed	q_6	0.00097	Identified
Penalty on longitudinal speed sum	q_7	0.0001	Assumed from parameter study
Update time	T_u	Variable	Measured

experimental and simulation signals is not significant at 5%. The match is excellent in some sections of the track. Even for corner number 3, which is difficult to interpret, a good match is observed. This is a very encouraging result because the approximations made for driver, vehicle and tyre models do not result in a substantial difference in nominal driving strategy.

The match between the lateral forces is also excellent, as shown in Figure 6.25. The error in this case is always below 4%. Such a good match would not be expected due to approximate modelling of the aerodynamic map and the load transfer. However, the tyre model also captures aerodynamics effects, increasing the fidelity of the simulation.

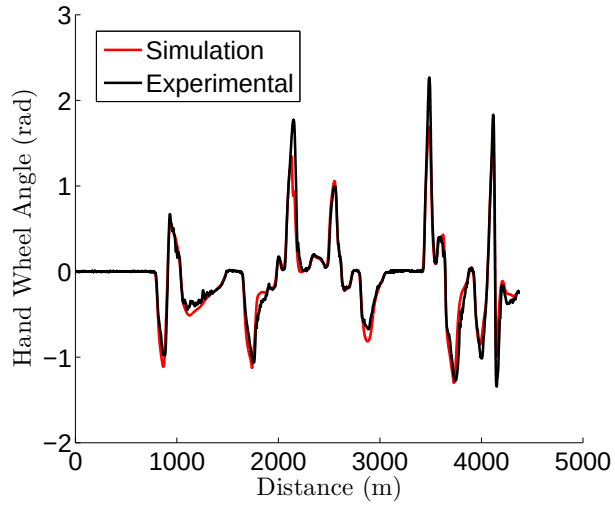

Figure 6.24: HWA for experimental and simulated lap. The match is excellent, with an average error of 3 %.

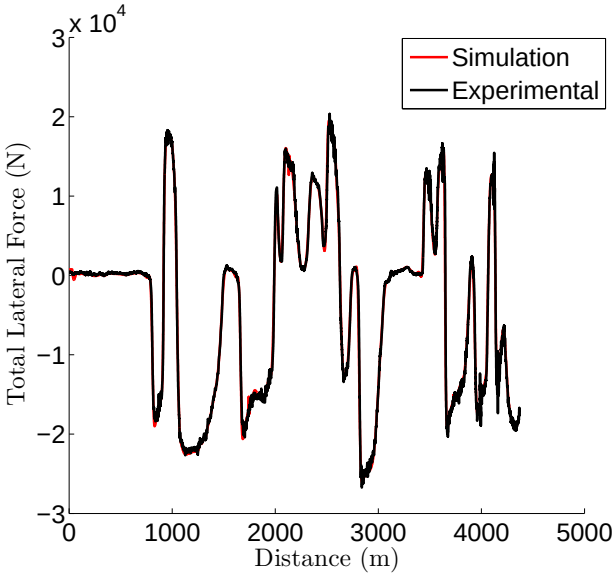

Figure 6.25: Sum of front and rear, hence total, lateral force for experimental and simulated lap.

Figure 6.26 gives a measure of the front tyre saturation for the simulated vehicle. Plotting the tyre surface for different velocities would not be insightful, also in light of the fact that the maximum available force is speed dependent. Instead, the difference between the maximum available slip, which is independent of the vertical load, and the total slip is considered.

$$\Delta\epsilon = \epsilon_{max} - \epsilon \tag{6.15}$$

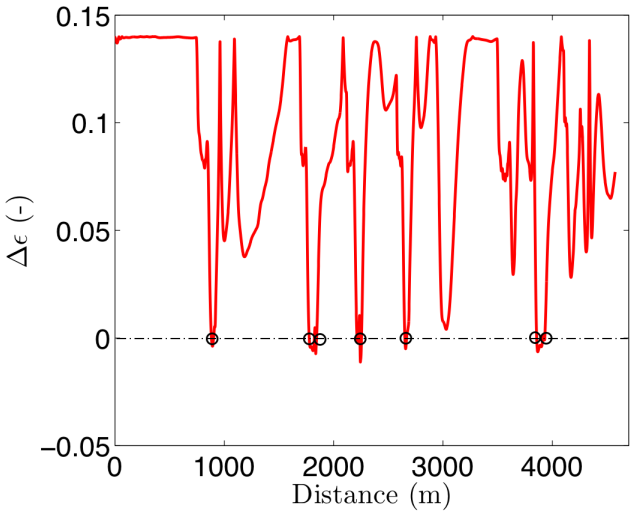

Figure 6.26: Difference between total slip and maximum available slip for the front tyre of the simulated vehicle. The front saturates several times over the course of the lap.

Clearly, when $\Delta\epsilon < 0$, tyres are saturating. Since the tyre model is only linearly dependent on F_{az}, the slip at which the horizontal force reaches its maximum value for a given vertical force, which is about 0.14, is constant. Instead of using $\Delta\epsilon$ as a metric, the total slip could simply be plotted, indicating its saturation value. The proposed method, however, can be extended to tyre curves that vary arbitrarily with speed so it is a valuable addition to the thesis.

On the main straight, between corners 16 and 1, the total slip is approximately zero, so $\epsilon = 0.14$. As the vehicle brakes, the slip increases, causing the tyres to saturate just before 1000 m. This happens various times over the lap, showing that the controller can stabilise the vehicle even when tyres are saturating. It is interesting to notice that the slip changes quite abruptly throughout the lap apart from a section around 1250 m. This corresponds to corner 3, where the slip decreases gradually, making it one of the hardest corners to interpret.

The behaviour of the rear tyre is similar, it saturates in correspondence of corners and braking points. The only difference is that slip never decreases to zero as a certain amount of traction force is always present.

Figure 6.27 shows that the speed profiles are very close, confirming that the longitudinal dynamics are modelled well. Figure 6.28 better quantifies the speed error by considering the difference in speed. The error is small and never exceeds 4 %. The error is negative for most of the lap, which means that the simulated vehicle is slower than the experimental one. The speed tracking error for the nominal signal does not

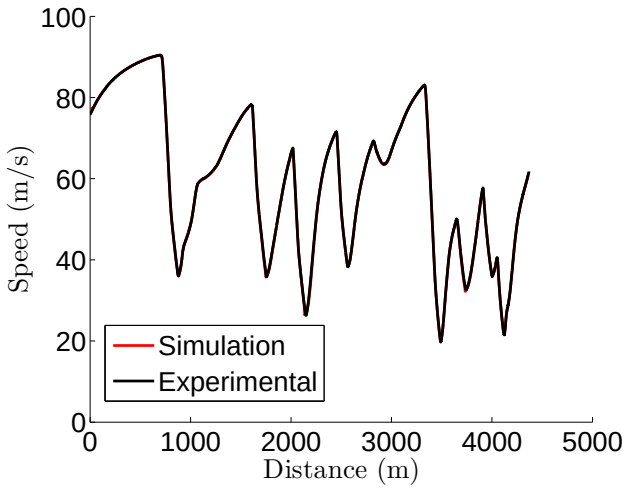

Figure 6.27: Vehicle speed for experimental and simulated lap

take into account the difference in braking points and other features of the signals which are lost in the averaging process. These will be further explored in Section 6.5.1.

Figure 6.29 shows the total longitudinal force. The match is excellent, but not as good as for the lateral case. Poor modelling of drag is likely to be the most significant contributing factor, given the complexity of the aerodynamic map and the assumptions made. In order to investigate this further, the aerodynamic map is needed.

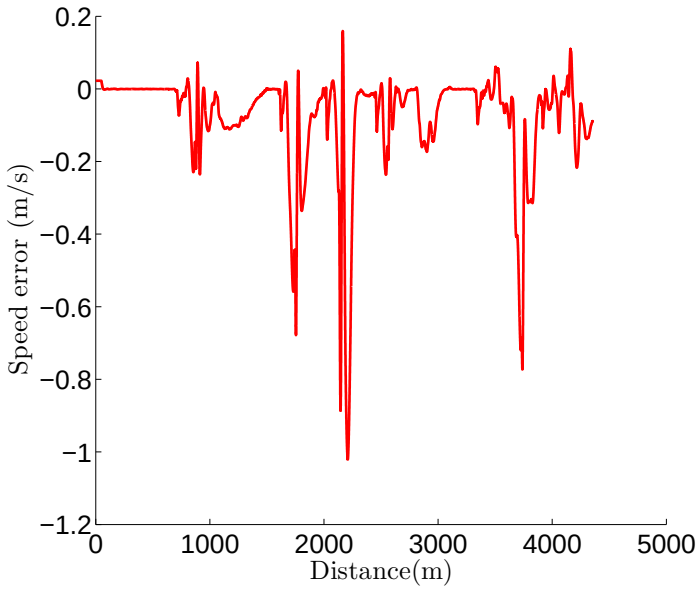

Figure 6.28: Speed error for experimental and simulated lap. A negative values indicates that the simulated vehicle is slower than the experimental one.

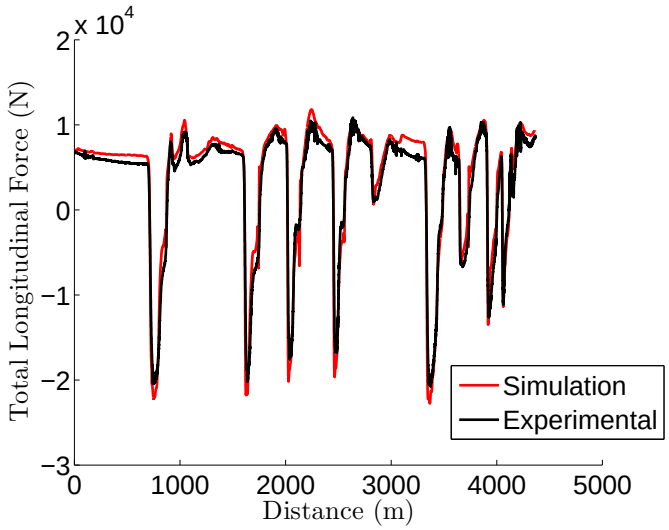

Figure 6.29: Total vehicle longitudinal force. The front and rear tyre forces are added.

6.4 Sensorimotor system noise

All model parameters have now been identified so experimental and simulated data can be compared. However, Section 6.2 shows that there is a statistically significant lap to lap standard deviation for all signals in absence of external disturbances, suggesting that a certain degree of randomness is also present in the driver. This needs to be accounted for to replicate the experimental results in the simulation environment.

Wolpert *et al.* [145] show that both the sensory system and the neuromuscular system are characterised by the presence of a certain degree of unpredictability. They confirm that noise contributes significantly to trial-to-trial variability. This suggests that confirming that the lap to lap variation may be attributable to said noise. Further randomness is also introduced by an imperfect knowledge of the dynamics of the vehicle and grip levels. These phenomena can be modelled as random disturbances in the driver controls.

Figure 6.4 shows that sections of high δ_{sw} standard deviation correspond to high curvature, hence high steering activity, suggesting that the disturbance is proportional to the signal magnitude. Nash *et al.*[114] show that noise in the neuromuscular system is signal-dependent. The disturbance is assumed to be white noise with zero mean and a standard deviation proportional to the signal magnitude. It is applied directly on δ_{com}. The new commanded signal $\delta_{com,t}$ becomes

$$\delta_{com,t} = (1 + r_n)\delta_{com} \tag{6.16}$$

where $\delta_{com,t}$ is the total steering angle and r_n is a random number drawn from a Gaussian distribution with zero mean and standard deviation Ξ, where Ξ is a constant. The value of Ξ is found heuristically and it is set to 0.6.

The HWA at the steering wheel will then be $\delta_{com,t}$ filtered by the neuromuscular system

$$\ddot{\delta}_{sw} + 2\zeta_n\omega_n\dot{\delta}_{sw} + \omega_n^2\delta_{sw} = \omega_n^2\delta_{com,t} \tag{6.17}$$

Figure 6.11 shows the *on-off* nature of longitudinal signals and Figure 6.10 shows that high standard deviation values correspond to braking points and that subsequent deviations from the mean are caused by the lap to lap variation in braking points rather by than a continuous disturbance. This suggests that the longitudinal controls are affected by continuous random sensorimotor noise and randomness in the timing of control events, with the latter being more significant. In order to simplify the analysis, only continuous longitudinal noise is taken into account. Even though it is not as significant as the timing of control events, it is similar to all the disturbances used for the model so far. Introducing a different type of disturbance would require a careful reconsideration of the controller architecture. The continuous disturbance is assumed to be a random gaussian noise with zero mean and a 0.02 standard deviation, which is chosen heuristically. There is no evidence of signal dependent noise for the longitudinal controls in the literature so the disturbance enters the equation of motion unaltered.

6.5 Data comparison

The controller was then tested on a single manoeuvre, Corner 1 and 2, multiple times to replicate the standard deviation naturally present in the driving task. The test is performed in absence of external random forces on the vehicle.

6.5.1 Comparison – measured time parameters

The whole lap was not considered because of the high computational cost. Corners 1 and 2 were chosen because the manoeuvre involves a representative variety of driving conditions. That is braking after the straight, with tyres saturating longitudinally, and two corners where tyres saturate laterally. The simulation is repeated fifteen times, which, even though it is not thirty, is enough to ensure that the results are statistically relevant. This is confirmed by running a Bayesian bootstrap on the data.

A prominent feature of the signal is the presence of ZOH, as Figure 6.15 and 6.16 show, which is consistent with the mixed SBH-ZOH hypothesis proposed in Section

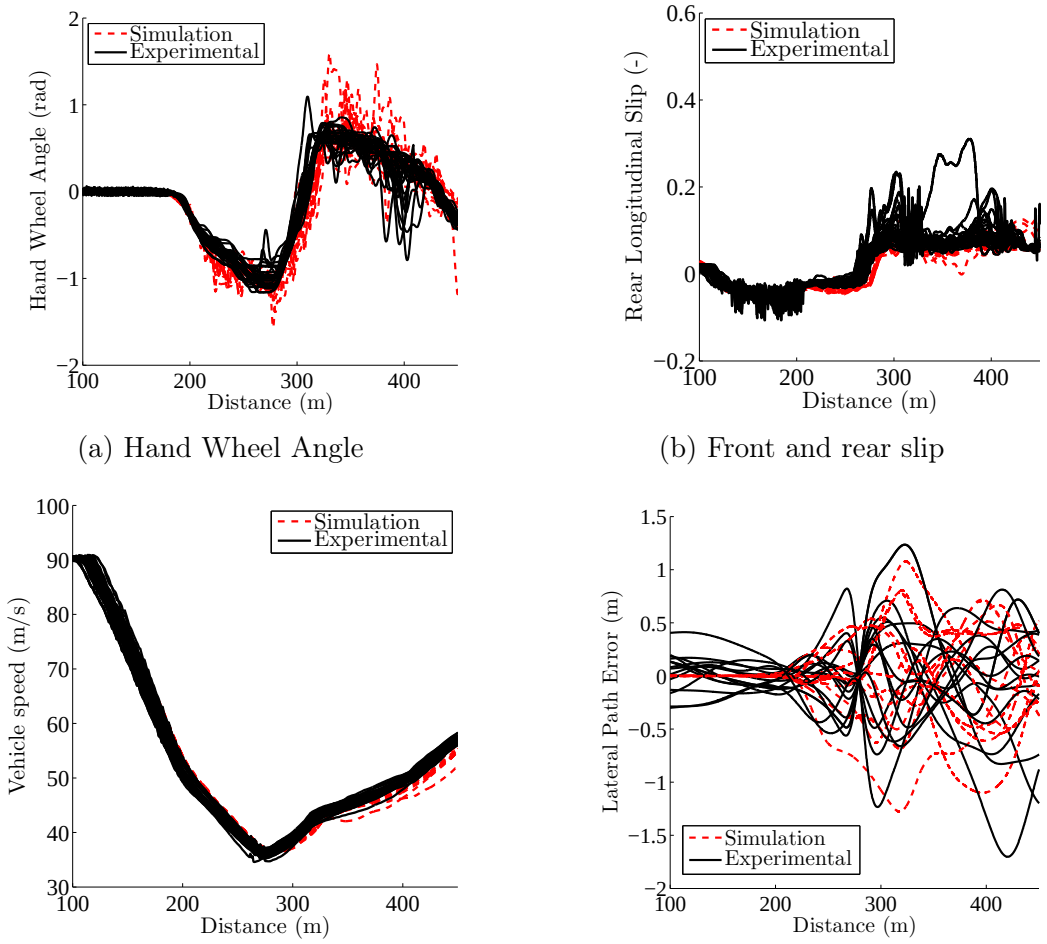

(a) Hand Wheel Angle

(b) Front and rear slip

(c) Vehicle Speed

(d) Lateral Path Error

Figure 6.30: Comparison between the disturbed experimental and simulated states for corners number 1 and 2.

5.3.2 . This strategy introduces further randomness into the system and it is replicated by assuming that each hold is an SBH followed by a ZOH, with the sum the two being the update time. The SBH and ZOH pattern for the the first fifteen laps from the driving simulator experiment was measured and replicated in the simulation environment, introducing a variable update time as a consequence.

Figure 6.30 shows the time histories of the inputs and the key states, the speed and the LPE. Figure 6.30a confirms that the controller tracks the nominal signal well, as the difference between the experimental and simulated means is negligible. The variance of both signals is also comparable, suggesting that the SBH and ZOH approach and the magnitude of the disturbance are reasonable.

The means of the experimental and simulated longitudinal rear slip, Figure 6.30b, do not match as closely as for the lateral controls. Longitudinal aerodynamics effects are difficult to capture and cause a discrepancy.

Figure 6.30c shows that even though the speed match is reasonably good, the controller fails to capture the differences in braking points, effectively affecting the rest of the signal time history. The mean and the variance, however, are reasonably close.

Figure 6.30d shows that the deviation from the LPE error is again comparable as the driver enters the corners. On the straight, the simulated path error is zero, which is consistent with what one would expect. However, the experimental data shows that there actually is path error. In the absence of external disturbances, the LPE on a straight line should theoretically be zero. The real driver, however, does not put significant emphasis on the trajectory on the straight, as this does not affect lap-time; resulting in LPE on the straight. Cost function weights are assumed to be constant, so the controller cannot predict such behaviour, which is limited to a driving scenario only affecting the lap-time marginally.Random errors are more significant than in other cases; this is due to the fact that on the straight the driver does not consistently follow the same path.

The experimental and simulated LPE and δ_{sw} standard deviation are compared, as they are good metrics to quantify the accuracy of the model. The longitudinal states and controls are not considered because longitudinal disturbances have not been applied to the vehicle.

The experimental and simulated HWA, Figure 6.31a, do not match well at the beginning of the manoeuvre. This is the result of the higher LPE deviation on the straight, which the controller fails to replicate. After the 200 m mark, the magnitudes of the standard deviations are comparable; while their time histories do not match perfectly, their mean values differ by 20%, which is reasonable given that the model is attempting to replicate a very complex system.

The LPE sd, Figure 6.31b, confirms that the controller can replicate the LPE to a very satisfying degree of accuracy. While the difference is marked in the first section of the manoeuvre, as curvature increases the sd time histories follow the same pattern, the simulation lagging slightly in the first corner and the experimental lagging in the second corner. The peak magnitudes have very similar magnitudes, the first being almost identical and the second only 10 % different.

The MSSD of the inputs are also considered, as they contain information on mean, variance and frequency content all in one signal.

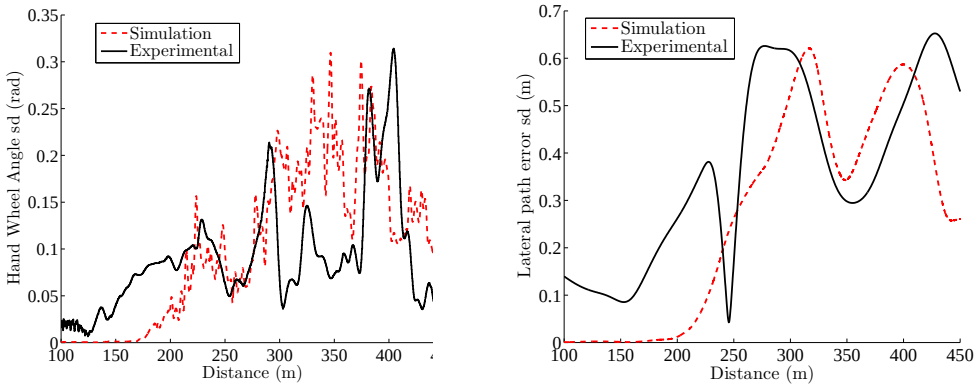

(a) Hand Wheel Angle standard deviation comparison.

(b) Lateral Path Error standard deviation comparison.

Figure 6.31: Comparison between signals standard deviations for corners number 1 and 2.

Figure 6.32a shows the average experimental and simulated δ_{sw} MSSD for 15 runs . At low frequencies the match is excellent, with an average 4 % error, which is consistent with the fact that the nominal HWAs match well. As frequency increases, the match in the frequency range of interest is still satisfactory, with errors never exceeding 6 %, showing that the δ_{sw} standard deviation is distributed equally over the frequency content. Some discrepancies are noticeable at higher frequencies, but these have very low magnitude, as they are beyond the limit of the driver's bandwidth.

The compensatory action of the controller can also provide good insights into its performance. For any signal S, the compensatory action can be calculated as

$$S_{com} = S - \mu_S \tag{6.18}$$

where μ_S is the signal mean. Figure 6.33a shows δ_{sw} compensatory action MSSD. The simulated activity is 5 % to 10 % higher over the whole frequency range, indicating that either the disturbance level is too high or the match is still satisfactory.

Figure 6.32b shows that, for low frequencies, the match for the longitudinal rear slip is good – with errors below 5 %-, confirming that the longitudinal dynamics can be captured to a good level of accuracy. The match is again not as good as for the δ_{sw} because assumptions made for modelling aerodynamic forces, drag in particular, have a more significant effect. At higher frequencies the only noticeable difference is the effect of the filter around 3 Hz, whose impact is not as clear for the experimental results. The interaction between the filter dynamics and the penalty of the signal in

the cost function is strong. In this case, the filter effect is more noticeable, though the difference is small and insignificant for the overall driving strategy given the low magnitude of the signal at which it happens. The compensatory κ_r MSSD, Figure 6.33b, shows a similar pattern: a very good match at low frequencies and a higher difference at higher frequencies.

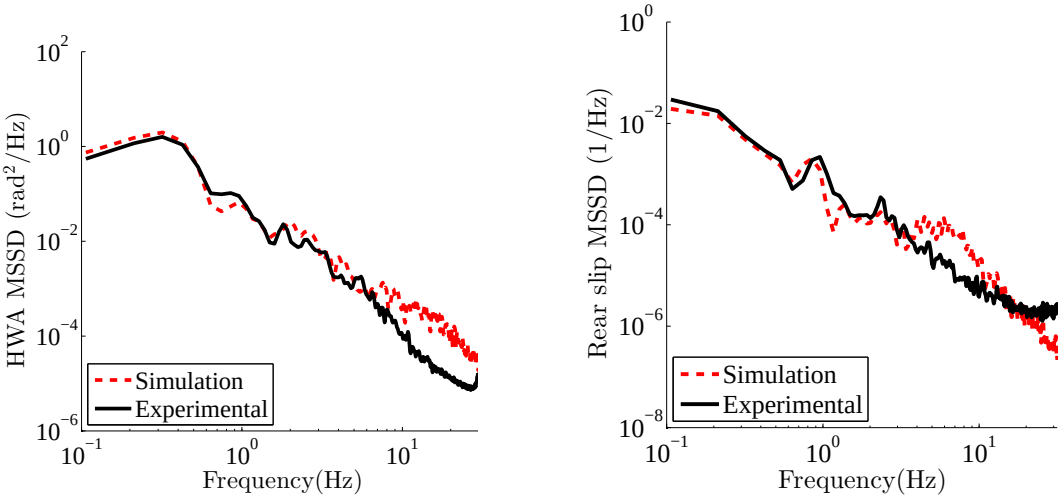

(a) Hand Wheel Angle MSSD

(b) Rear slip MSSD

Figure 6.32: Comparison between the disturbed experimental and simulated MSSD for corners number 1 and 2.

6.5.2 Comparison - Estimated time parameters

Section 6.5.1 shows that the proposed controller can replicate to a reasonable degree of accuracy the behaviour of a human driver. In the identification procedure, the ZOHs, the SBHs lengths, hence the update time, are measured. This approach does not allow to reproduce driver's behaviour in absence of experimental data, which partially defeats the purpose of the simulation. In order to circumvent this limitation, time parameters can be estimated from the measured data. Firstly, the minimum update time is set to 0.01 s. A shorter minimum update time, 0.005 s, was used to calculate H_p^{per} and H_p^{pert}; however, research on human refractory times [2] suggests that the lower bound is 0.01 s.

In order to have an estimate of the update time, the ZOH and SBH lengths for each control action need to be estimated. A joint probability density function could therefore be derived. However, the probability distribution of the ratio between

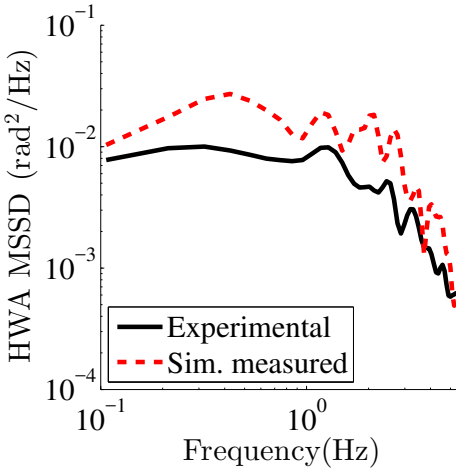

(a) Compensatory Hand Wheel Angle MSSD

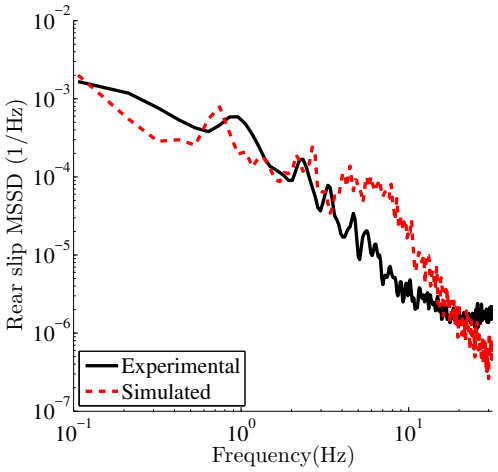

(b) Compensatory Rear slip MSSD

Figure 6.33: Comparison between the compensatory experimental and simulated MSSD for corners 1 and 2.

the ZOH length and the update time, T_{ZOH}/T_u, can be approximated to a uniform distribution with a 0.56 mean, as outlined in Section 6.2.1. It is therefore safe to assume that $T_u = 1/0.56 * T_{ZOH}$, eliminating the need for a joint probability function. A pdf for T_{ZOH} can be found by fitting the experimental data. Figure 6.34 shows the experimental pdf in black. The clear exponential nature of the curve suggests that a Pareto distribution would best fit the data. Since the asymptote is not zero, a Pareto type II, also known as Lomax distribution, is the best candidate.

$$p(t_{ZOH}) = \frac{\alpha_l}{\lambda}\left(1 + \frac{t_{ZOH}}{\lambda_l}\right)^{-(\alpha+1)} \tag{6.19}$$

where α_l and λ_l are constant parameters. The line of best fit, shown in Figure 6.34, is found using the Nelder-Mead algorithm, which yielded $\alpha_l = 7.9$ and $\lambda_l = 0.22$. The data set was obtained from a single driver, so little can be said about the universality of the result. However, even though the distribution parameters are likely to vary for different drivers, results in Chapter 5 and Chapter 6 suggest that a Pareto distribution could be used to model the hold behaviour of other drivers as well.

The time parameters are obtained by generating a random T_{ZOH} from $p(t_{ZOH})$ and multiplying by $1/T_{uz,avg}$ to obtain T_u.

The comparison between simulated and experimental data is repeated using the estimated time parameters, keeping everything else constant. The longitudinal controls

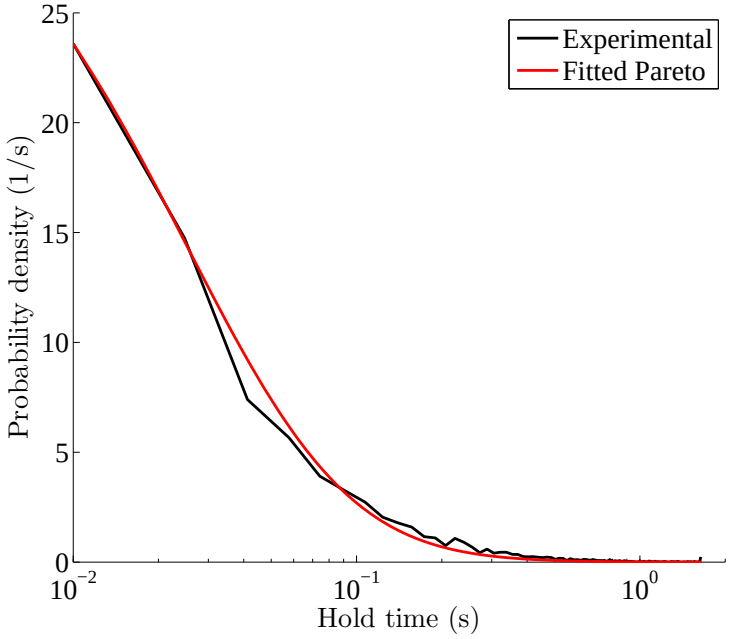

Figure 6.34: Average experimental hold length pdf and the fitted Pareto distribution

and states are not considered as they are not directly affected and do not change significantly. Results in the following Figures are labelled in the following way:

- Simulated results using estimated parameters - *Sim. estimated*
- Simulated results using measured parameters - *Sim. measured*
- Experimental results - *Experimental.*

The δ_{sw} time histories, Figure 6.35a show that estimating time parameters does not introduce a significant error. The HWA standard deviation of the simulation with estimated parameters, Figure 6.35b, follows the same pattern as the simulation with measured parameters; the magnitude is larger in places but the difference is generally less than 10 %.

The LPE time histories and standard deviation, Figures 6.36a and 6.36b, indicate that estimating time parameters increases the path tracking error by no more than 10 %, without changing the fundamental behaviour of the controller significantly.

Figure 6.37 shows the comparison between the MSSDs of the simulated δ_{sw} with estimated time parameters, $\text{MSSD}_{sim,e}$, the experimental one MSSD_{exp} and the simulated one with measured time parameters $\text{MSSD}_{sim,m}$. The frequency response is consistent

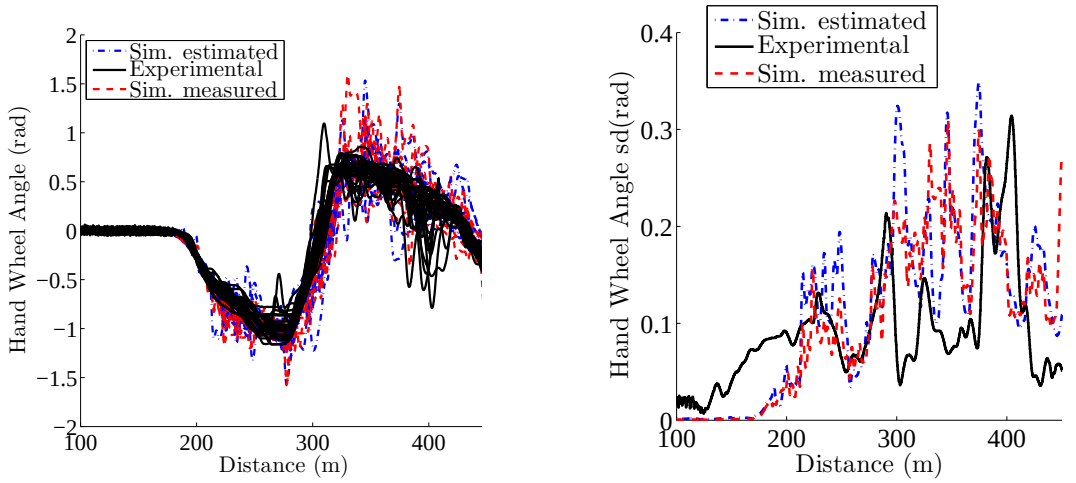

(a) Hand Wheel Angle time history

(b) Hand Wheel Angle standard deviation

Figure 6.35: Comparison between the disturbed experimental and simulated δ_{sw} time history and standard deviation for corners 1 and 2.

with what was previously observed; $\text{MSSD}_{sim,e}$ and $\text{MSSD}_{sim,m}$ are characterised by very similar responses, with variations well below 3 % for the frequency range of interest. The match with the experimental signal is still acceptable, which indicates that the controller can be used to replicate the behaviour of a human driver. Figure 6.38, shows that estimating parameters increases the compensatory action at higher frequencies by no more than 15 %.

6.6 Summary

This Chapter has focused on the analysis of experimental data from a driving simulator and the validation of the combined controller.

The statistics of the inputs and states have been considered and most expected behaviours confirmed. As tyres saturate, the driver becomes more erratic, introducing significant randomness in all signals. Long sections of the track characterised by decreasing curvature, such as corner 3, are the most difficult to interpret, as a high LPE and variance in throttle position indicate. While the longitudinal controls are consistent across laps, the braking point, especially at the end of the straight, varies considerably. The analysis of experimental data also revealed that ZOH are common in sections of the track where the curvature is not constant, highlighting the importance of this strategy.

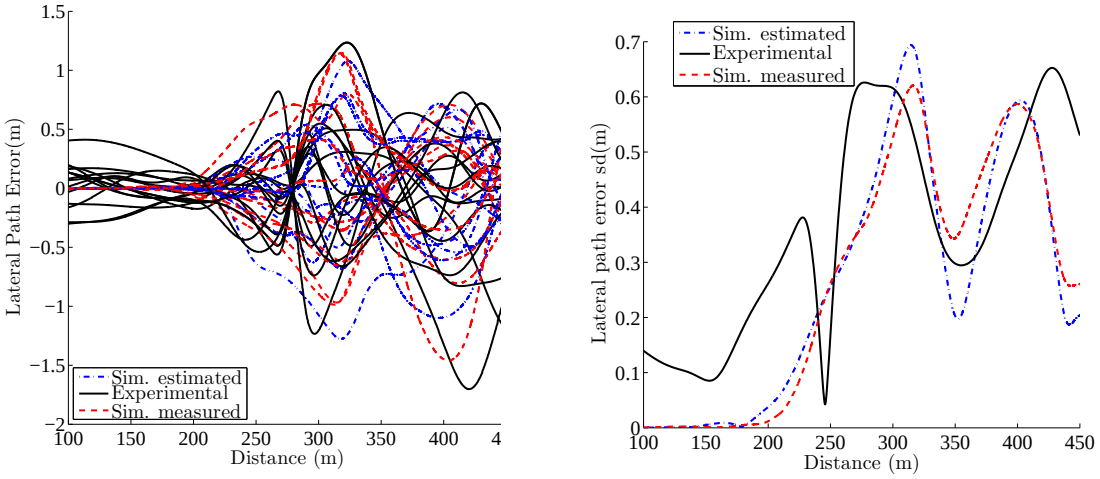

(a) Lateral Path Error time history

(b) Lateral Path Error standard deviation

Figure 6.36: Comparison between the disturbed experimental and simulated δ_{sw} time history and standard deviation for corners 1 and 2.

The tyres have been fitted by minimising the difference between the time history of the forces. A simple aerodynamic model has also been devised. The controller has then been used to track the nominal data to tune the model parameter with a very good match obtained. Simulated and experimental data were compared for corners 1 and 2, showing that the controller can predict the randomness of the driver to a good degree of accuracy. Tyres saturate both laterally and longitudinally as the vehicle navigates the two corners, making the manoeuvre representative of most track driving conditions. The result can be expected to be generalised to any geometry without any significant loss of accuracy. Data from different drivers would be needed to formally quantify the universality of the model; however, tuning cost function weights and NMS parameters can gather for a variety of driving skills.

This Chapter has shown that the controller can track the nominal states of a vehicle at the limit of adhesion in standard racing conditions and predict driver behaviour to a good degree of accuracy. The proposed model can capture the system's intrinsic randomness.

The main intellectual contribution made in this Chapter is the design of the procedure to compare experimental and simulated data to confirm that the controller can achieve the research objective.

Experimental data also revealed some areas for improvement such as variable cost function weights. These are explored in more detail in Chapter 7.

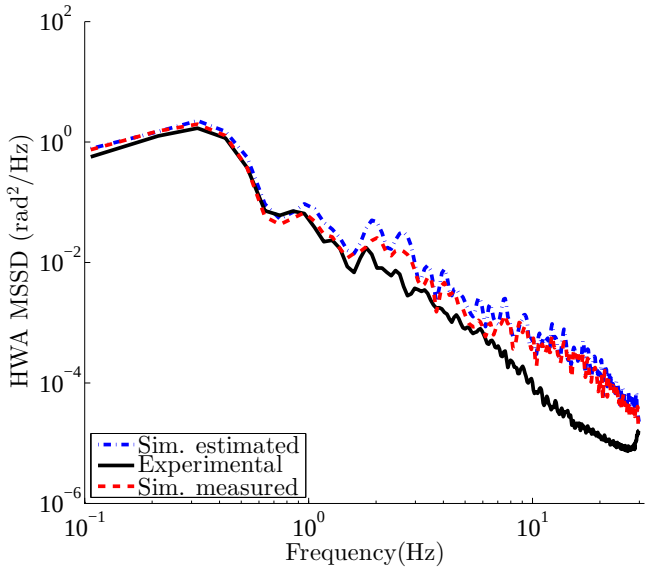

Figure 6.37: Comparison between δ_{sw} MSSD$_{sim,e}$, δ_{sw} MSSD$_{exp}$ and δ_{sw} MSSD$_{sim,m}$ for corners 1 and 2.

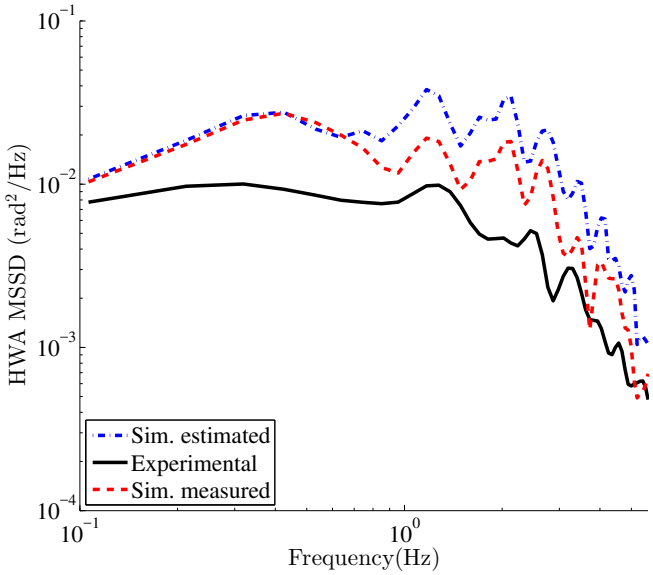

Figure 6.38: Comparison between the compensatory δ_{sw} MSSD$_{sim,e}$, δ_{sw} MSSD$_{exp}$ and δ_{sw} MSSD$_{sim,m}$ for corners 1 and 2.

Chapter 7

Conclusions and further work

The work undertaken in this research was aimed at developing a mathematical model which can mimic the behaviour of a racing driver controlling a vehicle at the limit of adhesion. Reproducing a human driver in a virtual environment could reduce the need for a driver in DIL simulators as well as providing an extremely useful design tool for all types of vehicles.

Section 7.1 summarises the main conclusions from each Chapter, specifying all contributions to knowledge while Section 7.2 recommends directions for further work.

7.1 Conclusions

7.1.1 Literature Review

The literature review considered developments in nonlinear vehicle dynamics; more precisely the tyre, vehicle and driver models aimed at modelling a car at the handling limit. This review highlighted that the vast amount of work done on lap-time minimisation led to the development of very sophisticated tools to minimise lap time in a deterministic environment.

It was also shown that very few attempts, however, have been made at trying to mimic the actual response of a human driver and little attention has been given to dealing with the uncertainties that characterise the real driving task. Moreover, most of the work done in this direction assumes a certain degree of linearity; the operational range of the tyres is often limited, thus making the models unsuitable for handling the vehicle at the limit of adhesion.

Control techniques were therefore reviewed to lay the foundations for a robust scheme that could stabilise the vehicle at the handling limit under the action of random

disturbances and attempt to mimic a human driver. It was therefore established that devising a controller that could mimic a human driver handling a vehicle at its limit could be an important contribution to knowledge.

7.1.2 Stability Metrics

Preliminary work included concluding a project aimed at developing novel stability metrics based on robust control ideas.

An extensive literature review on stability and controllability metrics is performed by Sideris [40]. A bicycle model augmented with a filter to account for the driver's neuromuscular system is used to find an ideal nominal, undisturbed, time-optimising trajectory using Timings MPC scheme [105]. Haslam [41] added disturbances to the system and superimposed a Linear Quadratic Regulator (LQR) to mimic the action of a human driver. He then computed state variances in an efficient fashion exploiting the linearity of the system. Calculations are validated against an iterative method.

Other stability metrics are also computed for comparison. The proposed stability and controllability metrics prove to be superior as:

- They relate to the response of the closed-loop driver-vehicle system and are thus likely to represent the driver's subjective assessment of the vehicle better than metrics derived from the vehicle dynamics alone.
- They are very fast to compute.

These stability metrics are used to analyse a 90 degree corner, yielding intuitive results confirming that the vehicle is harder to control as tyres are close to saturation. The original contribution to this Chapter was reviewing the work, independently reproducing the results to ensure consistency and representing them in an intuitive fashion. The work, including a refined version of the review by Sideris, resulted in a paper [106].

7.1.3 Driver Model Derivation

The linear driver model developed in Chapter 3 cannot guarantee stability for a vehicle at the limit of adhesion. Rather, a robust formulation able to deal with non-linearities has to be implemented. Firstly, notions of optimal and robust control are reviewed. These are then applied to derive a robust, nonlinear controller based on the concept of *tube based* MPC.

An MPC controller is used to find a nominal trajectory that is tracked by an ancillary controller as the vehicle is subjected to disturbances. Nonlinear optimisation is implemented for the ancillary controller. A single input, lateral control is considered first.

A bicycle model with nonlinear tyres, augmented with a filter to account for the driver NMS, is used for the derivation. An intrinsic track model is preferred, as it greatly simplifies the cost function. The model is enriched with features, such as update time and different hold types, which allow it to mimic the behaviour of a human driver.

The formulation is extended to include longitudinal dynamics, modifying the vehicle and tyre models accordingly. The work done in this Chapter shows how to apply the robust control techniques developed in [102] to a physical system. The authors of [102] highlight that most of the advances in the nonlinear robust control field are theoretically sound but of limited practical use. Controlling a vehicle under the action of disturbances is a practical application that reveals various features of the the controller.

The contribution to knowledge of this Chapter is the derivation of a driver model able to stabilise a vehicle under the action of disturbances. Optimal control formulations present in the literature, for instance [83–86], despite the very complete vehicle models, are deterministic and do not take randomness into account.

7.1.4 Driver Model Testing

The validity of the formulation proposed in Chapter 3 is tested on increasingly challenging driving scenarios. The lateral controller manages to stabilise US and OS vehicles at the handling limit under the action of disturbances, confirming the robustness of the formulation.

The combined controller deals with an even more challenging scenario, where the target path and speed are physically unachievable. Here, the model manages to keep the vehicle at the limit of adhesion, minimising both the deviation from the nominal path and the speed error. The controller is also tested in more realistic driving conditions – a 90 degree corner, showing that robustness is independent of track geometry.

Parameter studies are an important section of the Chapter. The effect of the parameters on the states and controls is investigated using a variety of approaches. It is concluded that both the weight q_3 on $\dot{\delta}$ and the NMS filter have a significant effect on δ_{sw} frequency content. The update time is found not to affect the frequency response of the steering significantly. The effect on the LPE is evident, however. The parameter study on the combined lateral and longitudinal controller further confirms two important features of the model: the direct relationship between control effort

and path tracking accuracy and the coupling between longitudinal controls and lateral dynamics.

This Chapter confirms that the controller can stabilise a vehicle with saturating tyres under the action of random forces, extending the work in [105]. The variety of driving scenarios tested demonstrates the controller effectiveness.

Two original contributions to knowledge are made in this Chapter, namely confirming that the formulation can stabilise a vehicle when at the limit of adhesion and finally investigating the effect of longitudinal controls on lateral states.

7.1.5 Driver Model Validation – Lateral Controls

While the controller achieves stability under a number of driving conditions, comparing the controls to those of an actual driver can give important insights into how well the model can mimic a human operator. Data for a driver navigating a circle under the action of disturbances from a driver-in-the-loop simulator is analysed. The SBH assumption is refined by adding a ZOH at the end, effectively proposing a mixed ZOH and SBH driving strategy.

The driver and tyre models' parameters are fitted to experimental data, rendering experimental and simulation conditions comparable. The statistics of the simulated and experimental signals show an excellent match. Despite the assumptions underlying the model, such as the lack of load transfer, the controller can mimic the controls of the human driver to a good degree of accuracy. Biological biases and differences in vehicle models become more significant as the vehicle's limits are approached. The difference between experimental and simulated LPE was more marked than the difference between controls, suggesting that the tyre and vehicle model discrepancies contributed to the overall error.

The formulation of a novel driving strategy that includes SBH and ZOH is the most important original contribution in this Chapter. This extends the work in [2] on driving strategy. Confirming the lateral controller soundness and the ability to mimic a real human driver is another significant contribution.

7.1.6 Driver Model Validation – Combined Controls

The validation is extended to the combined controller. Experimental data of a driver lapping an international racetrack is considered. Disturbances are not applied to the vehicle, so the driver's intrinsic randomness is accounted for by disturbances in the neuromuscular system. Experimental data is analysed, confirming most expected

behaviours. As tyres saturate, the driver becomes more erratic, introducing significant randomness in all signals. Longitudinal controls are generally consistent across laps, but braking points are characterised by considerable randomness. The randomness in lateral controls appears to be proportional to the signal amplitude. Finally, a careful consideration of the driving strategy shows that ZOHs are used extensively, regardless of track curvature.

Model parameters are identified from experimental data and a simple aerodynamic model is proposed. Cost function parameters are tuned by tracking the nominal signal. Simulated and experimental data is compared for corners 1 and 2, showing that the controller can predict the randomness of the driver with a satisfactory degree of accuracy.

There are two main contributions to knowledge in this Chapter. Firstly, data from a human driver is analysed to characterise the driver's behaviour, in a similar fashion to [142] but for a Formula 1 vehicle. Secondly, simulated and experimental data is compared to confirm that the controller can simulate the control actions of the human test subject with statistical properties that agree to a useful level of accuracy.

This Chapter has also highlighted various areas for improvement and potential for further work, expanded upon in Section 7.2.

7.2 Further work

Even though the research objectives have been met, the controller developed would still not be suitable for practical applications. However, extending the work done could make the formulation implementable in a practical context. Suggestions for further work follow.

7.2.1 Model extensions

Though the vehicle model used here is able to capture to a good degree of accuracy the behaviour of a real vehicle, refining the dynamic formulation would clearly lead to a model more applicable to actual racing conditions.

Firstly, the bicycle model could be extended to a four-wheels model to account for load transfer and a limited slip differential. Longitudinal load transfer, despite being almost negligible, could be added. Engine maps, gear shifts and KERS would also be relevant, especially when considering longitudinal dynamics. An empirical aerodynamic map could be added to model both drag and down-force. Very strong non-linearities,

such as driving on curbs, wheels lifting or locking, would be important contributions to the model if real racing scenarios are to be reproduced. Finally, the tyre model could be enhanced to include thermal effects and wear.

7.2.2 Numerical efficiency

One of the biggest hurdles to the practical implementation of this scheme is the computational burden associated with it. The efficiency of the proposed formulation could be improved in a number of ways.

Firstly, symbolic rather than numerical derivatives could be used, avoiding very expensive differentiation routines. Secondly, the whole algorithm could be coded in a lower level programming language, such as C. Finally, and perhaps most importantly, parallelisation could be exploited. The dynamics, the constraints and the cost function can be evaluated independently on separate processors due to the full state and control discretisation.

7.2.3 Variable update time

The hold behaviour and variable update time are clearly governed by a number of feedback inputs, including the distance from the ideal line, the orientation of the vehicle and the grip level. It would be instructive to find the correlation among the feedback signals and the update time variation.

A threshold strategy - which assumes a perfect correlation between update time and LPE – has been implemented and tested in the context of the validation in Chapter 5. Comparison with experimental data has shown that the assumption is wrong so it has not been included in the thesis. A cross correlation study could be carried out on a number of parameters to understand those that have a significant effect on the update time. Slip level is likely to be an important factor. Hand wheel torque feedback is in fact very widely used by racing drivers to adjust their feedback action.

7.2.4 Parametric uncertainty

Real drivers do not have perfect knowledge of the vehicle dynamics, so internal models are not the same as actual full scale ones. Uncertainties might involve time variation of parameters such as the tyre-road friction coefficient. The current work assumes that the driver has either perfect knowledge of the vehicle or that parametric uncertainty is included into the driver's sensorimotor system noise.

An more rigorous approach to model the discrepancies – for instance using predictive filters – would be an significant improvement for the controller.

7.2.5 Control strategy variation

Experimental data suggests that the driver varies their compensatory control strategy throughout the lap. For instance, the LPE is not zero on the straight, even in absence of disturbances. Furthermore, depending on vehicle states, the driver may opt for different ways of allocating their cognitive resources. These differences in control strategy might be accommodated by the variation in cost function weights.

In order to formulate a more realistic driving strategy, it would be necessary to identify which factors - such as speed or curvature – affect driving strategies and then adjust the cost function weights accordingly.

Another aspect to take into account in this context is driver fatigue. The total cognitive capacity assigned to the driving task should somehow be correlated with the length and intensity of the driving task, making nominal state tracking increasingly more erratic for long and challenging scenarios. An important step in all these developments will be experimental validation of the driver's feedforward and feedback control strategies and of the relationship between objective and subjective responses, extending the work in 2.

7.2.6 Online controls retrieval

It is very unlikely that a human can perform online optimisations, even for longer update times. It could be hypothesised that the driver retrieves signals previously learned and applies them depending on the states of the vehicle. This idea follows from observing the learning process [146].

An inexperienced driver explores the vehicle performance boundaries and then replicates the motions learned. Racing drivers may have an internal control library generated through experience, from which they would choose controls depending on the situation. It is also possible that the actual driving signals are generated from very few *primitives* [146] which are then scaled or modified to fit the current conditions.

An experience driver can push a vehicle to the limit on a track they have never driven on in less than 50 laps. Evidence can be found in the lap-times for races on new racetracks. In the inaugural Sochi GP in 2015, drivers were able to consistently lap within 3 % of their personal best lap in Free Practice 2, so after approximately 25 laps. Even though they all practised on the simulators before, this suggests that signals may

be stored and then adapted to the circumstances. It would be instructive to explore this idea by creating libraries of controls and applying them at each update time.

References

[1] J.T.Timings. *Application of Convex Optimisation and Robust Control to Dynamic Lap-Time Simulation*. PhD thesis, Cambridge University Engineering Department, 2011.

[2] T. Johns. *The effect of cognitive workload on a racingdrivere's steering and speed control*. PhD thesis, University of Cambridge, 2013.

[3] C.C. MacAdam. Application of an optimal preview control for simulation of closed-loop automobile driving. In *IEEE Trans. on Systems Man and Cybernetics*, volume 11, pages 393–399, 1981.

[4] S. J. Rutherford. *Modelling driver nonlinear steering control*. PhD thesis, Department of Engineering, University of Cambridge, 2007.

[5] B. Cheng and T. Fujioka. Driver model by fuzzy logic control for obstacle avoidance. *JSAE Review*, 18(2):201–210, 1997.

[6] S. Keen and D. Cole. Application of time-variant predictive control to modelling driver steering skill. *Vehicle System Dynamics*, 49(4):527–559, 2011.

[7] H. Pacejka. *Tyre and vehicle dynamics*. Butterworth Heineman, Oxford, 2002.

[8] U. Kiencke and L. Nielsen. *Automotive control system for engine,driveline, and vehicle*. Springer, 2005.

[9] V.Cossalter. *Motorcycle Dynamics*. Wiley, 2006.

[10] R.Rajamani. *Vehicle Dynamics and Control*. Springer, 2012.

[11] H.B. Pacejka and R.S. Sharp. Shear force development by pneumatic tyres in steady-state conditions: A review of modelling aspects. *Vehicle System Dynamics*, 20:121–176, 1991.

[12] Wong J.Y. *Theory of ground vehicles*. John Wiley & Sons, 2001.

[13] E. Fiala. Seitenctyres am rollenden luftreifen. *VDI-Zeitschrift*, 96:973–979, 1954.

[14] H. Sakai. Theoretical and experimental studies on the dynamic properties of tyres, part 1: review of theories of rubber friction. *International journal of vehicle design*, 2(1):78–110, 1981.

[15] R.S. Sharp and M.A. El-Nashar. A generally applicable digital computer based mathematical model for the generation of shear fforce by pneumatic tyres. *Vehicle System Dynamics*, 40(2):621–643, 1986.

[16] H.B.Pacejka and E. Bakker. The magic formula tyre model. *Tyre models for vehicle dynamics analysis, Supplement to vehicle system dynamics*, 21:1–18, 1993.

[17] E. Bakker and H.B. Tyre modelling for use in vehicle dynamics studies. Technical Report 870421, SAE, 1987.

[18] G. Mastinu and R.Cesarini. Stabilita' e guidabilita' nel piano orizzontale. *ATA*, 1991.

[19] F. Della Rossa, G. Mastinu, and C. Piccardi. Bifurcation analysis of an automobile model negotiating a curve. *Vehicle System Dynamics*, 50(10):1539–1562, October 2012.

[20] P.A. Ioannou and C.C. Chien. Autonomous intelligent cruise control. *IEEE Transactions on vehicular technology*, 42(4):657–672, 1993.

[21] K.E. Majdoub, F. Giri, H. Ouadi, L. Dugard, and F.Z. Chaoui. Vehicle longitudinal motion modelling for nonlinear control. *Control Engineering Practice*, 20:69–81, 2011.

[22] E.Velenis, D. Katzourakis, E. Frazzoli, P. Tsiotras, and R. Happee. Steady-state drifting stabilization of rwd vehicles. *Control Engineering Practice*, 19(11):1363–1376, November 2011.

[23] D. P. Kelly. *Lap Time Simulation with Transient Vehicle and Tyre Dynamics*. PhD thesis, Cranfield University School of Engineering, 2008.

[24] J.C. Dixon. *Tyres, suspensions and handling*. SAE International, 1996.

[25] W. Milliken and D. Milliken. *Race car vehicle dynamics*. Society of Automotive Engineers, 1995.

[26] E. Ono, S. Hosoe, H.D. Tuan, and S. Doi. Bifurcation in vehicle dynamics and robust front wheel steering control. *IEEE Transactions on Control Systems Technology*, 6(3):412 – 420, May 1998.

[27] V. Nguyen. Vehicle handling, stability and bifurcation analsys for nonlinear vehicles models.

[28] R. Skoog and C. Lau. Instability of slowly varying systems. *IEEE Transactions on Automatic Control,*, 17(1):86–92, 1972.

[29] C. Desoer. Slowly varying system. *Automatic Control, IEEE Transactions on*, 14(6):780–781, Dec 1969.

[30] H.H. Rosenbrock. The stability of linear time-dependant control systems. *Journal of Electronics and Control*, 15(1):73–78, 1963.

[31] S. Evangelou. *The Control and Stability Analysis of Two-Wheeled Road Vehicles*. Phd, Imperial College London, 2003.

[32] I.G. Salisbury and D.J.N. Limebeer. Motion cueing in high-performance vehicle simulators. *International Journal of Vehicle Mechanics and Mobility*, 55(6), 2017.

[33] T.A. Johansen, K.J. Hunt, P.J. Gawthrop, and H. Fritz. Off-equilibrium linearisation and design of gain-scheduled control with application to vehicle speed control. *Control Engineering Practice*, 6(2):167–180, 1998.

[34] J.P. Meijaard and A. Popov. Practical stability analysis for transient system dynamics. *Proceedings of the Royal Society of London*, 463(2084):2123–2135, 2008.

[35] Bo-Chiuan Chen, Cheng-Chi Yu, Wei-Feng Hsu, and Min-Fang Lo. Design of electronic stability control for rollover prevention using sliding mode control. *International Journal of Vehicle Design*, 56(1):224–245, 2011.

[36] W.Y. Youn and J.B.Song. Improvement of vehicle directional stability in cornering based on yaw moment control. *KSME International Journal*, 14(8):836–844, 2000.

[37] M. Gerdts, S. Karremberg, B. Muller-BeBler, and G. Stock. Generating locally optimal trajectories for an automatically driven car. *Optimization and Engineering*, 10(4):439–463, 2009.

[38] H.V. de Castro. *Flying and handling qualities of a fly-by-wire blended-wing-body civil transport aircraft*. Phd, Cranfield University, 2003.

[39] R.F. Stengel. *Notes on aircraft stability*. Princeton University Press, 2004.

[40] M. Sideris. F1 car stability analysis. Master's thesis, University of Cambridge, 2011.

[41] A. Haslam. Stability and controllability of a formula one racing car. Master's thesis, University of Cambridge, 2012.

[42] M. Plochl and J. Edelmann. Driver models in automobile dynamics application. *Vehicle System Dynamics*, 45(7-8):699–741, 2007.

[43] C. C. MacAdam. Understanding and modeling the human driver. *Vehicle System Dynamics*, 40:101–134, 2003.

[44] L. Kaufman and J. P. Thomas. *Handbook of Perception and Human Performance: Sensory Processes and Perception.* John Wiley, 1986.

[45] T. Sheridan. Control models of creatures which look ahead. In *Proceedings of the Fifth National Symposium on Human Factors in Electronics*, 1964.

[46] T. Sheridan. Three models of preview control. *IEEE Transactions on Human Factors in Electronics*, 7:91–102, 1966.

[47] D. McRuer, L. Hofmann, H. Jex, G. Moore, A. Phatak, D. Weir, and J. Wolkovitch. New approaches to human-pilot/vehicle dynamic analysis. Technical report, Systems Technology Inc., 1968.

[48] A. Phatak and G. Bekey. Model of the adaptive behavior of the human operator in response to a sudden change in the control situation. *IEEE Transactions on Man-Machine Systems*, 10:72–80, 1969.

[49] D. McRuer and E. Krendel. The man-machine system concept. In *Proceedings of the IRE*, 1962.

[50] I. Ashkenas and D. McRuer. A theory of handling qualities derived from pilot-vehicle system considerations. *Aerospace Engineering*, 21:60–61, 1962.

[51] B. Cheng and T. Fujioka. A hierarchical driver model. In B. Cheng and T. Fujioka, editors, *Proceedings of IEEE Conference on Intelligent Transportation System*, Proceedings of IEEE Conference on Intelligent Transportation System, 1997.

[52] R. Sharp. Preview control of active suspensions. *Smart Vehicles*, Swets & Zeitlinger, Lisse, Netherlands, 1995.

[53] N. Louam, D. Wilson, and R. Sharp. Optimal control of vehicle suspension incorporating the time delay between front and rear wheel inputs. *Vehicle System Dynamics*, 17(6):317–366, 1988.

[54] R. S. Sharp and V. Valtetsiotis. Optimal preview car steering control. *Supplement to Vehicle System Dynamics*, 35:101–117, 2001.

[55] A. Ungoren and H. Peng. An adaptive lateral preview driver model. *Vehicle System Dynamics*, 43:245–259, 2005.

[56] D. J. Cole, A.J. Pick, and A. M. C. Odhams. Predictive and linear quadratic methods for potential application to modelling driver steering control. *Vehicle System Dynamics*, 44(3):259–284, 2006.

[57] J.Maciejowski. *Predictive control with constraints.* Pearson Education, 2002.

[58] S. D. Keen and D. J. Cole. Steering control using model predictive control and multiple internal models. In *Proceedings of The 8th International Symposium on Advanced Vehicle Control*, 2006.

[59] P. Falcone, F. Borrelli, J. Asgari, Hong, and D. Hrovat. Predictive active steering control for autonomous vehicle systems. In *IEEE Transactions on Control Systems Technology*, volume 15, pages 566–580, 2007.

[60] R. S. Sharp. Application of optimal preview control to speed-tracking of road vehicles. In *Proceedings of IMechE, Part C: Journal of Mechanical Engineering Science*, volume 221, pages 1571–1578, 2007.

[61] R. S. Sharp. Optimal preview speed-tracking control for motorcycles. *Multi-body System Dynamics*, 18:397–411, 2007.

[62] H. Fritz. Neural speed control for autonomous road vehicles. *Control Engineering Practice*, 4(4):507–512, 1996.

[63] C. C. MacAdam. Development of a driver model for near / at-limit vehicle handling. Technical report, University of Michigan, 2001.

[64] P. Falcone. *Nonlinear Model Predictive Control for Autonomous Vehicles*. PhD thesis, Department of Engineering, University of Sannio,Benevento, 2007.

[65] P. Falcone, H. E. Tseng, F. Borrelli, J. Asgari, and D. Hrovat. Mpc-based yaw and lateral stabilisation via active front steering and braking. *Vehicle System Dynamics*, 46:611–628, 2008.

[66] S. Chang and T. Gordon. Model-based predictive control of vehicle dynamics. *International Journal of Vehicle Autonomous Systems*, 5:3–27, 2007.

[67] S. Chang and T. J. Gordon. A flexible hierarchical model-based control methodology for vehicle active safety systems. *Vehicle System Dynamics*, 46:63–75, 2008.

[68] T. J. Gordon and N. Magnuski. Modeling normal driving as a collision avoidance process. In Taipei, editor, *Proceedings of The 8th International Symposium on Advanced Vehicle Control*, 2006.

[69] S. D. Keen. *Modeling Driver Steering Behaviour using Multiple-Model Predictive Control*. PhD thesis, Department of Engineering, University of Cambridge, 2008.

[70] C. Voser, R. Hindiyeh, and J. Gerdes. Analysis and control of high sideslip manoeuvres. *Vehicle System Dynamics*, 48:317–336, 2010.

[71] K. Talvava, K. Kritayakirana, and J. Gerdes. Pushing the limits: From lanekeeping to autonomous racing. *Annual Reviews in Control*, 35(1):137–148, 2011.

[72] D. Tavernini, M. Massaro, E. Velenis, D. I. Katzourakis, and R. Lot. Minimum time cornering: the effect of road surface and car transmission layout. *Vehicle System Dynamics*, 51(10):1533–1547, 2013.

[73] M. Klomp, M. Lidberg, and T.J.Gordon. On optimal recovery from terminal understeer. In *Proceedings of the Institution of Mechanical Engineers,Part D: Journal of Automobile Engineering*, volume 228, pages 412–425, 2012.

[74] B. Siegler, A. Deakin, and D. Crolla. Lap time simulation: Comparison of steady state, quasi-static and transient racing car cornering strategies. In *Proceedings of 2000 SAE Motorsports Engineering Conference and Exposition*, 2001.

[75] J. Blasco-Figueroa. Minimum time manoeuvre based in the gg-speed envelope. Master's thesis, School of Engineering, Cranfield University, 2000.

[76] M. Gadola, D. Vetturi, D. Cambiaghi, and L. Manzo. A tool for lap time simulation. In *Proceedings of the SAE Motorsport Engineering Conference and Exposition*, Dearborn, Michigan, USA, 1996.

[77] J. Hendrikx, T. Meijlink, and R. Kriens. Application of optimal control theory to inverse simulation of car handling. *Vehicle System Dynamics*, 26:449–461, 1996.

[78] D. Casanova. *On Minimum Time Vehicle Manoeuvring: The Theoretical Optimal Lap*. PhD thesis, School of Mechanical Engineering, Cranfield University, 2000.

[79] A. Komatsu, T. Gordon, and M. Best. Vehicle path optimisation using a time-variant linear optimal reference control. In *Proceedings of The 8th International Symposium on Advanced Vehicle Control*, 2002.

[80] E. Velenis and P. Tsiotras. Minimum time vs maximum exit velocity path optimization during cornering. In *Proceedings of the 2005 IEEE International Symposium on Industrial Electronics*, 2005.

[81] E. Velenis and P. Tsiotras. Minimum-time travel for a vehicle with acceleration limits: Theoretical analysis and receding horizon implementation. *Journal of Optimization Theory and Applications*, 138(2):275–296, 2008.

[82] E. Velenis and P. Tsiotras. Optimal velocity profile generation for given acceleration limits: theoretical analysis. In *Proceedings of the 2005 American Control Conference*, 2005.

[83] M. Gerdts, S. Karrenberg, B. Muller-BeBler, and G. Stock. Generating locally optimal trajectories for an automatically driven car. *Optimization and Engineering*, 2008.

[84] D. J. N. Limebeer and G. Perantoni. Optimal control for a formula one car with variable parameters. *Vehicle System Dynamics: International Journal of Vehicle Mechanics and Mobility*, 52(5):653–678, 2014.

[85] David J.N. Limebeer and Anil V. Rao. Faster, higher, and greener: Vehicular optimal control. *IEEE control systems*, 35(2):36–56, April 2015.

[86] Roberto Lot and Nicola Del Bianco. Lap time optimisation of a racing go-kart. *International Journal of Vehicle Mechanics and Mobility*, 54(2):210–230, 2016.

[87] M. Caneri, R. Lot, V. Cossalter, and M. Massaro. Optimisation of the roll and steer systems of a four-wheeled tilting vehicle. In *Proceedings of the ASME 2014 12th Biennial Conference on Engineering Systems Design and Analysis ESDA2014*, 2014.

[88] A.J. Tremlett, M. Massaro, D. J. Purdy, E. Velenis, F. Assadian, A. P. Moore, and M. Halley. Optimal control of motorsport differentials. *International Journal of Vehicle Mechanics and Mobility*, 53(12):1772–1794, 2015.

[89] D.I. Katzourakis, E. Velenis, E. Holweg, and R. Happee. Haptic steering support for driving near the vehicle's handling limits; skid-pad case. *International Journal of Automotive Technology*, 15(1):151–163, 2014.

[90] D. Tavernini, E. Velenis, R. Lot, and M. Massaro. The optimality of the handbrake cornering technique. *Journal of Dynamic Systems, Measurement, and Control*, 136(4), 2014.

[91] D Tavernini, M Massaro, E Velenis, D I Katzourakis, and R. Lot. Minimum time cornering: the effect of road surface and car transmission layout. *International Journal of Vehicle Mechanics and Mobility*, 51(10), 2013.

[92] M. Thommyppillai, S. Evangelou, and R. S. Sharp. Advances in the development of a virtual car driver. *Multibody System Dynamics*, 22:245–267, 2009.

[93] M. Thommyppillai. *Optimal path-tracking of virtual race-cars using gain-scheduled preview control*. PhD thesis, Department of Mechanical, Electrical and Electronic Engineering, Imperial College London, 2010.

[94] R. E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.

[95] E. Bertolazzi, F. Biral, and M. Dalio. Symbolic-numeric efficient solution of optimal control problems for multibody systems. *Journal of computational methods in science and engineering*, 2(3):1–3, 2003.

[96] D. Li, Y. Xi, and P. Zheng. Constrained robust feedback model predictive control for uncertain systems with polytopic description. *International Journal of Control*, 82:1267–1274, 2009.

[97] H. Genceli and M. Nikolaou. Robust stability analysis of constrained l1-norm model predictive control. *AIChE Journal*, 39:1954–1965, 1993.

[98] J. H. Lee and Z. Yu. Worst-case formulations of model predictive control for systems with bounded parameters. *Automatica*, 33:763–781, 1997.

[99] D. Mayne and E. Kerrigan. Tube-based robust nonlinear model predictive control. In *Proceedings of the Ninth IFAC World Congress*, 2007.

[100] K. Kouramas S. Rakovic, E. Kerrigan and D. Mayne. Approximation of the minimal robustly positively invariant set for discrete-time lti systems with persistent state disturbances. In *Proceedings of The 42nd IEEE Conference on Decision and Control*, 2003.

[101] S. Rakovic, E. Kerrigan, K. Kouramas, and D. Mayne. Invariant approximations of the minimal robust positively invariant set. *IEEE Trans. Automatic Control*, 50(3):406–410, 2005.

[102] D.Q. Mayne, E.C. Kerrigan, E.J. van Wyk, and P.Falugi. Tube-based robust nonlinear model predictive control. *International Journal of Robust and Nonlinear Control*, 21(11):1341–1353, July 2011.

[103] M. Farrokhsiar and H. Najjaran. A robust probing motion planning scheme: a tube-based mpc approach. In *American Control Conference*, 2013.

[104] N.E.D. Toit and J.W.Burdick. Robot motion planning in dynamic uncertain environments. *IEEE transactions on robotics*, 28(1):101–115, 2012.

[105] J. Timings and D. Cole. Robust lap-time simulation. *Journal of Automobile Engineering*, 1(17), March 2014.

[106] G. Braghieri, A. Haslam, M. Sideris, J. T., and D. Cole. Quantification of road vehicle handling quality using a compensatory steering controller. *Journal of Dynamic Systems, Measurement, and Control*, 139(3), 2016.

[107] M.C. Best and T.J. Gordon. On the synthesis of inputs for the simulation of closed-loop handling manoeuvres. *International Journal of Vehicle Design*, 40(1/2/3):52–76, 2006.

[108] A. M. C. Odhams and D. J. Cole. Identification of preview steering control models using data from a driving simulator and a randomly curved road path. *International Journal of Autonomous Vehicle Systems*, 12(1):44–64, 2014.

[109] D. J. Cole and A. J. Pick. Dynamic properties of a driver's arms holding a steering wheel. *Proceedings of the Institution of Mechanical Engineers, Part D*, 221:1475–1486, 2007.

[110] J. P. Timings and D. J. Cole. Minimum manoeuvre time calculation using convex optimisation. *ASME Journal of Dynamic Systems, Measurement and Control*, 135(031015), May 2013.

[111] A. G. Ulsoy, C. Lin, and D. J. LeBlanc. Vehicle dynamics and external disturbance estimation for vehicle path prediction. *Control Systems Technology, IEEE Transactions*, 8:508–518, May 2000.

[112] S. D. Keen and D. J. Cole. Bias-free identification of a linear model predictive steering controller from measured driver steering behavior. *IEEE Transactions on Systems, Man and Cybernetics, Part B*, 42(2):434–443, 2012.

[113] T. A. Johns and D. J. Cole. Measurement and mathematical model of a driver's intermittent compensatory steering control. *Vehicle System Dynamics*, 53(12):1811–1829, 2015.

[114] C. J. Nash and D. J. Cole. Development of a novel model of driver-vehicle steering control incorporating sensory dynamics. In *Proceedings of IAVSD Symposium on Vehicles on Roads and Tracks, Graz, Austria*, August 2015.

[115] R.C. Dorf. *Modern Control Systems*. Addison-Wesley Series in Electrical Engineering. Addison-Wesley, 1980.

[116] M. Grewal and A.P. Andrews. *Kalman Filtering: Theory and Practice Using MATLAB*. Wiley-IEEE Press, 2008.

[117] M. Apetaur and F. Opicka. Assessment of the driver's effort in typical driving manoeuvres for different vehicle configurations and managments. *International Journal of Vehicle Mechanics and Mobility*, 20:42–56, 1992.

[118] D. P. Bertsekas. *Dynamic Programming and Optimal Control Volume I*. Athena Scientifi, 1995.

[119] A. E. Bryson and Yu-Chi Ho. *Applied Optimal Control: Optimization, Estimation, and Control*. Taylor & Francis, 1975.

[120] R. E. Bellman and S. E. Dreyfus. *Applied Dynamic Programming*. Princeton University Press, 1962.

[121] E. Bertolazzi, F. Biral, and M. Da Lio. Symbolic-numeric indirect method for solving optimal control problems for large multibody systems. *Multibody System Dynamics*, 13(233), 2005.

[122] J. T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. SIAM, second edition, 2010.

[123] G. A. Bliss. *Lectures on the Calculus of Variations*. 1946.

[124] D. Liberzon. *Calculus of Variations and Optimal Control Theory*. Princeton University Press, 2012.

[125] Erwin Fehlberg. Low-order classical runge-kutta formulas with step size control and their application to some heat transfer problems. Technical report, NASA, 1969.

[126] D.P. Bertsekas. *Nonlinear programming*. Athena scientific, 2001.

[127] M. Massaro and R. Lot. A virtual rider for two-wheeled vehicles. In *49th IEEE Conference on Decision and Control (CDC)*, pages 5586–5591, 2010.

[128] G. Perantoni and D. J. N. Limebeer. Optimal control of a formula one car on a three-dimensional track part 1. *Journal of Dynamic Systems, Measurement and Control*, 137(5), 2015.

[129] M.S. Young and N.A. Stanton. It is all relative: Defining mental workload in the light of annett's paper. *Ergonomics*, 45(14):1018–1020, 2002.

[130] CJD Patten. *Cognitive Workload and the Driver*. PhD thesis, Stockholm University Department of Psychology, 2007.

[131] P. Falugi, E.C. Kerrigan, and E.J. van Wyk. Imperial college london optimal control software user guide. *Imperial College London, London UK*, 2010.

[132] G.J. Forkenbrock and D. Elsasser. An assessment of human driver steering capability. Technical Report DOT HS 809 875, US National Highway Traffic Safety Administration, 2005.

[133] J.Timings and D. Cole. Vehicle trajectory linearisation to enable efficient optimisation of the constant speed racing line. *Vehicle System Dynamics*, 50(6):883–901, June 2012.

[134] B. Chen, Y. Zhu, J. Hu, and J. Principe. *System Parameter Identification*. Elsevier, 2001.

[135] R. Aster, B. Borchers, and C.H. Thurber. *Parameter estimation and inverse problems*. Elsevier, 2013.

[136] J.T.Timings. *Application of Convex Optimisation and Robust Control to Dynamic Lap-Time Simulation*. PhD thesis, Cambridge University Engineering Department, 2011.

[137] S.Korkel, H. Qu, G. Rucker, and S Sager. *Current Trends in High Performance Computing and Its Applications*. Springer, 2005.

[138] R. Schumway. *Time series analysis and applications*. Springer, 2000.

[139] J. B. Carroll. *Human cognitive abilities: a survey of factor analytic studies*. Cambridge Univeristy Press, 1993.

[140] D. Csermely and L. Regolin. *Behaviour lateralization in vertebrates: two sides of the same coin*. Springer Science, 2012.

[141] Will Pittenger. Circuit de barcelona-catalunya.

[142] John C. Kegelman, Lene K. Harbott, and J. Christian Gerdes. Insights into vehicle trajectories at the handling limits: analysing open data from race car drivers. *Vehicle System Dynamics*, 55(2):191–207, 2017.

[143] Pi Research. Aerodynamics application note. Technical Report 34324, Cosworth, 2003.

[144] W.F. Philips. *Eshbach handbook of engineering fundamentals*, volume Chapter 10 Aerodynamics of wings. Wiley, 5th edition, 2009.

[145] Faisal AA, Selen LPJ, and Wolpert DM. Noise in the nervous system. *Nature reviews Neuroscience*, 4(9):292–303, 2008.

[146] F.L. Moro, N.G. Tsagarakis, and D.G. Caldwell. On kinematic motion primitives - theory and application. *Frontiers in Neurorobotics*, 6(10), 2012.