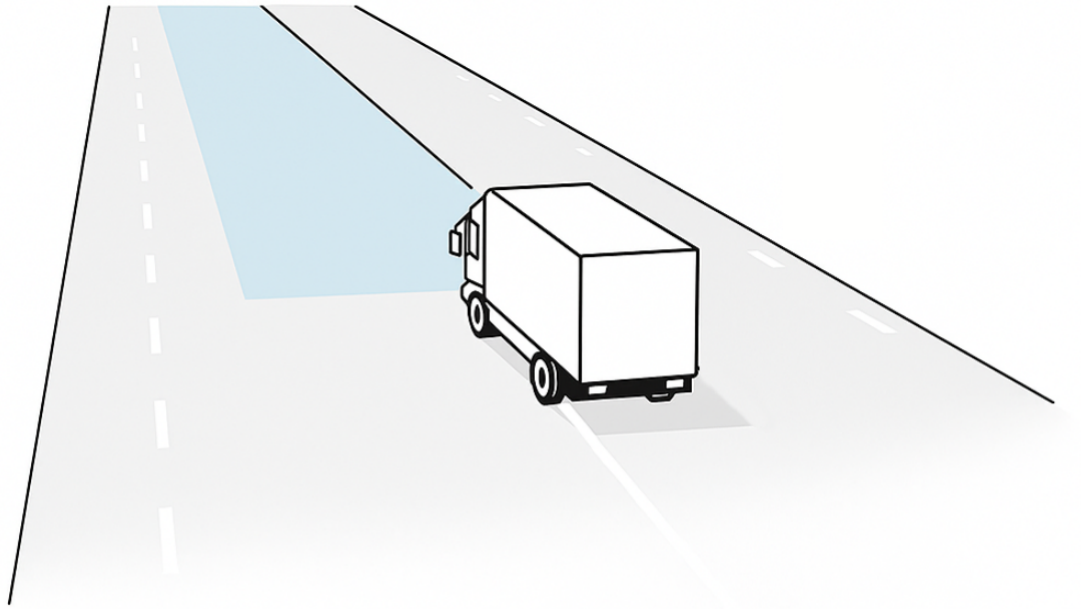

Improving heavy duty vehicle stability during split-μ braking situations with rear axle steering

OLLE ANDERSSON, ÅKE WILÉN

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis 2026

Improving heavy duty vehicle stability during split-μ braking situations with rear axle steering

OLLE ANDERSSON
ÅKE WILÉN

CHALMERS
UNIVERSITY OF TECHNOLOGY

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
Gothenburg, Sweden 2026

Improving heavy duty vehicle stability during split-μ braking situations with rear axle steering
Olle Andersson, Åke Wilèn

Supervisor: Mats Jonasson, Department of Mechanics and Maritime Sciences Chalmers University of Technology
Supervisor: Dawn Joy, Volvo Trucks TTI
Examiner: Mats Jonasson, Department of Mechanics and Maritime Sciences Chalmers University of Technology

Master's Thesis 2026
Department of Mechanics and Maritime Sciences
Division of Vehicle Engineering and Autonomous Systems
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Split-μ braking (Figure 3.3 converted to cartoon theme using AI)

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Improving Heavy Duty Vehicle Stability During Split-μ Braking Situations with Rear Axle Steering
OLLE ANDERSSON
ÅKE WILÉN
Department of Mechanics and Maritime Sciences
Chalmers University of Technology

Abstract

This thesis investigates the use of rear axle steering (RAS) to improve the stability of heavy-duty vehicles during split-μ braking scenarios. Split-μ braking, where the friction coefficient differs between the left and right wheel paths, generates asymmetric longitudinal forces and induces yaw moments that can compromise vehicle stability, control and safety.

Controllers for the RAS actuator were developed and evaluated, including PID control, gain-scheduled Linear Quadratic Regulator (LQR), robust $\mathcal{H}_\infty$ control, adaptive Model Predictive Control (MPC), and a brake pressure feed-forward approach. The controllers were implemented and assessed in a co-simulation environment using MATLAB/Simulink and IPG TruckMaker. The stability properties of the controllers were analysed using an induced-norm-based metric to assess their influence on the driver's intended vehicle behaviour. Performance was evaluated across multiple scenarios, including straight-line braking, curved-road braking, and lane-change manoeuvres under various slit-μ conditions. Performance metrics included yaw rate, yaw angle, lateral deviation, braking distance, and driver steering effort.

The results demonstrate that RAS can significantly improve lateral stability and reduce driver effort without notably increasing braking distance. Among the evaluated methods, the LQR and PID controllers provided the best overall performance in terms of stability and driver workload reduction, with the LQR offering a balance between performance and control effort. The $\mathcal{H}_\infty$ controller showed robustness, while the MPC showed less consistent performance for short-duration, high-dynamics manoeuvres. The feed-forward approach proved effective in reducing initial yaw disturbances when combined with feedback control. Furthermore, it is shown that a more aggressive ABS strategy can be applied without compromising stability when combined with RAS control, thereby improving both braking performance and stability. The results also show that performance is primarily limited by the steering rate of the RAS actuator rather than by its maximum angle. Additionally, the Linear Parametric Varying (LPV) reference model used for control design can be further extended to improve controller performance and expand the operational capability of the RAS system.

Overall, the thesis confirms that rear axle steering is a viable approach for enhancing vehicle stability during critical split-μ braking scenarios and can be implemented using existing vehicle signals without requiring predictive sensing technologies.

Keywords: rear axle steering, split-μ braking, heavy-duty vehicles, yaw control, yaw stability, vehicle stability, vehicle dynamics

Acknowledgements

We would like to express our thanks to our supervisors at Volvo Trucks TTI, Dawn Joy and Emma Heiwall, for their continuous support, guidance, and valuable insights throughout this thesis. Their contributions have been essential to the completion of this work.

We would also like to thank the entire Advanced Steering Systems team, as well as other teams at Volvo Trucks TTI, for their support, collaboration, and knowledge sharing. Their engagement and willingness to assist created a supportive environment that enabled the results presented in this thesis.

In addition, we want to thank our examiner at Chalmers University of Technology, Professor Mats Jonasson, for his valuable expertise and guidance throughout this thesis.

Olle Andersson, Åke Wilèn, Gothenburg, May 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

ABS	Anti-lock Braking System
BP_{FF}	Brake Pressure Feed Forward
CAN	Controller Area Network
CoG	Center of Gravity
CUSUM	Cumulative Sum
ECU	Electronic Control Unit
ESC	Electronic Stability System
FB	Feedback
FF	Feedforward
ISO	International Organisation for Standardisation
LPV	Linear Parameter Varying
LQR	Linear Quadratic Regulator
LTVKF	Linear Time-Varying Kalman Filter
MD	Measured Disturbance
MPC	Model Predictive Control
MV	Manipulated Variable
PID	Proportional-Integral-Derivative
RAS	Rear Axle Steering
RMS	Root Mean Squared
RWA	Rear Wheel Angle
SMO	Sliding Mode Observer

Nomenclature

Below is the nomenclature of indices, sets, parameters, and variables that have been used throughout this thesis.

A, B, C, D	: State-space system matrices
BP_{tot}	: Total applied wheel braking torque
$BP_{\text{Lwhl},i}$	: Braking torque on left side on axle i
$BP_{\text{Rwhl},i}$	: Braking torque on right side on axle i
C_i, C_{eq}	: Cornering stiffness (for axle i or equivalent 2-axle system)
$\mathcal{C}$	: Controllability matrix
$e(t)$	: Error signal
F_x	: Longitudinal force acting on the vehicle
F_y	: Lateral tire force
F_z	: Vertical wheel load
$G(s)$	: Transfer function matrix
$\|G(s)\|_\infty$	: $\mathcal{H}_\infty$ norm / peak induced 2-norm
I_z	: Total yaw inertia
J	: Optimization cost function
K	: LQR optimal feedback gain matrix
K_{fx}	: Longitudinal force gain with respect to brake pressure
K_z, K_λ	: Weighting factors for wheel deceleration and relative slip
l_f	: Distance from front axle to CoG
l_m	: Distance from middle axle to CoG
l_r	: Distance from rear axle to CoG
l_{eq}	: Equivalent wheelbase distance
L_k	: Time-varying Kalman gain
$\mathcal{L}$	: Set of wheels on the left side of the vehicle
m	: Total vehicle mass
M_z	: Yaw moment acting on the vehicle

N_p, N_c	: Prediction and control horizons
P	: Riccati equation solution matrix / Error covariance matrix
Q, R	: State/output and control weighting matrices
r, r_{ref}	: Yaw rate and reference yaw rate
$\mathcal{R}$	: Set of wheels on the right side of the vehicle
s	: Laplace variable
SK	: ABS trigger criterion
t_w	: Vehicle track width
T_{min}	: Minimum total wheel torque threshold
T_{whl}	: Wheel torque
u	: Control input vector / manipulated variable
v_x	: Longitudinal vehicle velocity
v_y	: Lateral vehicle velocity
$\dot{v}_y$	: Lateral vehicle acceleration
$w_i(\rho)$	: Weighting factor for convex combination of local LTI systems
$W_{(\cdot)}$	: Frequency-domain weighting functions
x	: State vector
x_e	: State error vector
y	: Output vector
z_R	: Wheel deceleration
α_i	: Tire slip angle at axle i
δ_f	: Front wheel angle
δ_r	: Rear wheel angle
ΔBP_{whl}	: Braking torque difference between left and right sides
λ_{rel}	: Relative wheel slip
μ	: Friction coefficient
ρ	: LPV scheduling parameter vector
ϕ	: Roll
θ	: Pitch
ψ	: Yaw
$\omega_z, \dot{\psi}$	: Yaw rate

Contents

List of Figures

List of Tables

1
Introduction

This thesis was conducted in collaboration with and within the Volvo Trucks organisation, specifically within the Advanced Steering Systems team. The focus of this thesis is on the control of the rear-axle steering actuator, which is currently implemented in certain heavy-duty trucks. While the system is primarily used to improve low-speed manoeuvrability, this thesis investigates its application in enhancing vehicle stability under critical braking conditions, such as split-μ scenarios. The following sections present the background and the problem addressed in this thesis, including relevant previous research, the overall purpose of the work, the research questions that guide the thesis, as well as its limitations.

1.1 Background

According to the National Highway Traffic Safety Administration (NHTSA), 5,472 people were killed in traffic crashes involving large trucks in 2023, representing an 8% decrease compared to the 5,969 fatalities reported in 2022 [1]. Despite this reduction, the number of fatalities remains substantial, highlighting the safety challenges associated with heavy-duty vehicles. Furthermore, the Large Truck Crash Causation Study (LTCCS), conducted by the Federal Motor Carrier Safety Administration (FMCSA) in 2007, identified the most frequent critical event types in large-truck crashes. The second most common critical event was vehicle loss of control, accounting for 29% of the cases in the study [2]. This category includes incidents resulting from excessive speed relative to road conditions, cargo shift, vehicle system failures, and adverse roadway surface conditions.

One particularly critical scenario for heavy-duty vehicles occurs when braking on surfaces with different friction coefficients between the left and right wheels, a situation known as split-μ braking. In such situations, asymmetric longitudinal tyre forces generate an unwanted yaw moment, causing the truck to rotate toward the side with higher friction. Split-μ braking has been studied extensively due to its high risk of vehicle instability, increased stopping distances, and potential to cause rollover.

Rear-axle steering (RAS) is commonly used to improve turning radius and reduce tyre wear on the rear axles. However, it may also enhance vehicle stability under split-μ braking conditions. Vehicle stability is a critical factor for ensuring safety during various driving scenarios, particularly during braking. When the friction

coefficients differ between the left and right wheels, the resulting yaw moment can lead to loss of control and increased stopping distance, thereby increasing the risk of accidents. While modern vehicles are equipped with Anti-lock Braking Systems (ABS) to prevent wheel lock-up and Electronic Stability Control (ESC) to mitigate yaw instability during asymmetric braking, further improvements in vehicle stability under split-μ conditions may be achieved by integrating rear-axle steering. The Swedish Parliament introduced the Vision Zero policy (Nollvisionen) in 1997, aiming to eliminate deaths and serious injuries in road traffic and the transport system [9]. One potential contribution toward this goal is improving heavy-duty vehicle stability during critical braking scenarios, such as split-μ braking, which is the focus of this thesis.

1.2 Purpose

The purpose of this thesis is to investigate how a rear-axle steering system can be designed and controlled to improve the stability of heavy-duty vehicles during split-μ braking, given the physical feasibility constraints of a rear-axle steering system.

1.3 Research Questions

The following questions are the main research questions to be answered:

- How does rear axle steering influence yaw rate, lateral deviation, and vehicle controllability during split-μ braking for different friction ratios, speeds, and road geometries?
- How much can rear axle steering improve the yaw stability and braking performance of heavy-duty vehicles during split-μ braking compared to a baseline vehicle without rear axle steering, under realistic actuator and system constraints?
- Using commonly available vehicle signals, how can split-μ braking conditions be characterised, and what is necessary to support the activation and operation of rear-axle steering control?
- What rear axle steering actuator capabilities (maximum steering angle, steering rate, delay, etc.) are required to achieve meaningful stability and braking improvements during split-μ braking?

1.4 Related Work

This section provides an overview of existing research related to split-μ braking, vehicle stability control, and rear-axle steering systems.

1.4.1 Rear-Axle Steering in Split-μ Braking

Kharrazi et al. [3] conducted a study on extending the functionality of a heavy truck's rear-axle steering system from low-speed manoeuvrability to high-speed sta-

bility and split-μ braking scenarios. Their work is directly relevant to this thesis, as it investigates the same core hypothesis: whether RAS can improve yaw stability and braking performance. The authors developed separate rear-axle steering controllers for high-speed manoeuvres and split-μ braking scenarios. Their results showed that the high-speed controller reduced yaw-rate error by approximately 64% during dynamic lane-change manoeuvres, while the split-μ controller enabled a reduction in braking distance of about 10% when combined with a more aggressive ABS strategy. Although both approaches demonstrated significant improvements in yaw stability and braking performance, the two control strategies were treated independently. The authors explicitly state that the integration of the high-speed and split-μ controllers was left as a topic for future work. The study provides strong evidence that an active rear-axle steering system can significantly enhance vehicle stability and reduce braking distance in critical split-μ situations.

1.4.2 Braking on Split-μ Surfaces

Gao et al. [4] investigated a control strategy that coordinates differential braking with the driver's steering input, rather than relying on rear-axle steering. However, their work addresses the same fundamental challenge as Kharrazi et al. [3], maintaining vehicle stability during split-μ braking, but approaches it through a different actuator set. By analysing wheel-speed differences, brake pedal position, and steering-wheel angle, a potentially dangerous heavy-braking event on split-μ is detected. A 20% wheel-speed difference between the left and right wheels is used to trigger the stability controller. The control strategy uses a reference generator and the Modified Hamiltonian Algorithm (MHA) to implement a steering-adaptive function that coordinates brake and steering actions in real time. Simulation and experimental results demonstrate that the proposed controller reduces yaw rate and body side-slip angle compared to ABS-only braking. Furthermore, compared to a conventional electronically controlled braking system (EBS) with ESC, the MHA-based controller achieves comparable stopping distances while requiring less steering correction from the driver. This indicates improved vehicle controllability and reduced driver workload during split-μ braking. This work is particularly relevant to the thesis because it highlights the importance of early and reliable split-μ detection as a prerequisite for effective stability control. While Gao et al.'s main focus is on differential braking, the underlying detection principles and stability objectives are directly applicable to integrating a detection function into the rear-axle steering system controller investigated in this thesis.

1.4.3 Split-μ Detection

While the previously studied studies, including this thesis, focused on control strategies for maintaining stability during split-μ braking, the problem of reliably detecting a split-μ situation is equally critical. In [5], the detection problem is addressed directly by comparison of two different detection approaches using real vehicle tests. The first method, a kinematic detector, uses wheel speed sensors and GPS/GNSS (Global Positioning System/Global Navigation Satellite Systems) data to compute

front-wheel slip, flagging a split-μ situation when the slip difference between the left and right front wheels exceeds a threshold during braking. This detector proved very fast, with an average detection time of 0.14 seconds, but was susceptible to false detections, particularly during low-μ braking. The second, more complex dynamic detector uses a Sliding Mode Observer (SMO) based on a single-track model to estimate the "unexpected yaw torque" generated by the split-μ condition. The observer's error-injection signal is then monitored using a Cumulative Sum (CUSUM) change-detection algorithm. This detector was more robust, with a lower rate of false detections, but slower, 0.25 seconds average, due to filtering and required more complex tuning. Their simulation results showed that a steering controller could handle a split-μ situation even without a detector. The detector's role is therefore to improve performance by enabling a change in control strategy. Based on this, they concluded that detection speed is more important than robustness, since a false detection did not have dangerous consequences. This makes the faster kinematic detector more suitable for enhancing controller performance during split-μ braking, which is particularly relevant for integrating a detection function into the rear-axle steering system controller studied in this thesis.

1.5 Limitations

The main limitation of the thesis is that the rear-axle steering system is the only system the controller actively controls to achieve better stability. It may use signals from the other subsystems in the truck for calculations but only in a read-only mode. Trucks with trailers also lie outside the scope of this thesis and will not be considered, as the trailer adds a higher degree of complexity in dynamics, modelling, and control. The thesis is also limited to and will only consider the modelling and control of 6x2 trucks with steerable tag axles. This limitation is motivated by the fact that these trucks are the most commonly used and operate without a trailer. The full functionality of a RAS system will not be modelled in detail. Instead, the system will be represented in the simulations by directly applying steering angles to the rear wheels. However, certain important aspects of the RAS system will be included in the model. In particular, the steering angles commanded to the rear wheels must comply with the specifications of a generic RAS system. RAS systems often have several physical limitations; not all of them will be considered in this thesis. The most significant limitations are the maximum steering angle and the angle rate. Other limitations will be ignored, as the system will be used at high speeds and during braking, even though they are not intended for such conditions.

1.6 Thesis Outline

The thesis is structured as follows. Chapter 1 introduces the background and the problem addressed in this thesis, including relevant previous research, the overall purpose of the work, the research questions that guide the thesis, and its limitations. Chapter 2 provides the theoretical background necessary to understand the methods and analyses presented later in the thesis. Chapter 3 describes the methodology and

workflow used in the thesis. The results obtained using these methods are presented and discussed in Chapter 4, while Chapter 5 concludes the thesis and addresses the research questions formulated in Chapter 1.

2
Theory

This chapter presents the theoretical foundation for the work carried out in this thesis. It begins with a general description of a rear-axle steering system, outlining its fundamental functionality, properties, and current implementations and uses for low-speed manoeuvrability. Following this, a brief background on some aspects of vehicle dynamics is presented, and other subsystems that this thesis will come into contact with are introduced and explained, such as ABS.

2.1 Rear Axle Steering Systems

Rear axle steering systems are primarily used to improve manoeuvrability, reduce turning radius and decrease tyre wear. These systems generally consist of two primary components: a power unit and a hydraulic cylinder unit. The power unit integrates an electric motor, hydraulic pump, and an electronic control unit (ECU). The cylinder unit includes a hydraulic piston mechanism, valve system, and a linear position sensor for feedback control.

At low vehicle speeds, the rear wheels are steered in counterphase with the front wheels. This means that the rear axle steers in the opposite direction from the front axle. This counter-phase behaviour is consistent with Ackermann steering geometry principles, where the wheel angles are arranged to minimise tyre slip during low-speed cornering. The steering angle of the front axle is transmitted via the vehicle controller area network (CAN) bus to the RAS control unit. Based on this input, the control unit calculates the desired rear axle steering angle and actuates the electric motor, which drives the hydraulic pump. The generated flow extends or retracts the cylinder piston, thereby steering the rear wheels through a mechanical system. The integrated linear sensor provides closed-loop feedback to ensure accurate positioning of the rear axle.

With the rear steering angle gradually reduced and eventually deactivated as vehicle speed increases, RAS systems are not currently designed or used to actively influence stability during high-speed driving or asymmetric braking, such as split-μ.

2.2 Vehicle Dynamics

This section aims at describing the mathematical and vehicle dynamics theory that lays the foundation for the modelling part of this thesis. Concepts and methods

used are described in order to follow the derivation of the reference model described in Chapter 3.

2.2.1 Vehicle Coordinate System

The coordinate system used for the truck follows the ISO 8855 convention for vehicle-fixed coordinate systems and rotations [10]. It is defined as a right-handed Cartesian coordinate system attached to the vehicle body. Figure 2.1 presents the vehicle-fixed coordinate system.

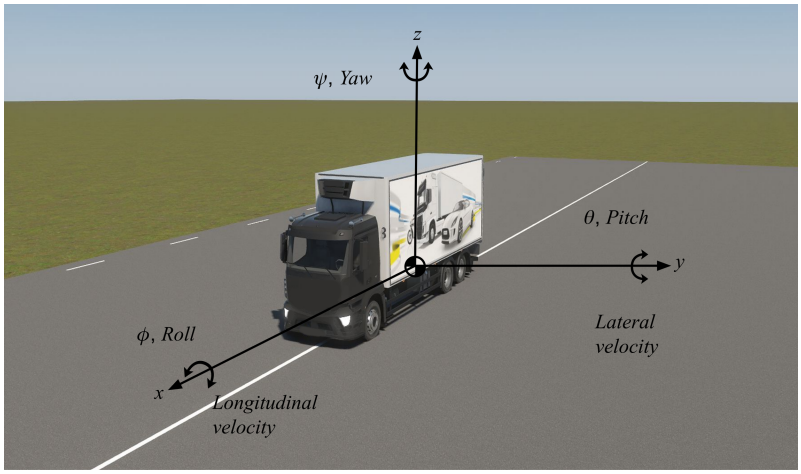

Figure 2.1: Vehicle coordinate system.

Translational motion along the x-axis represents longitudinal velocity, while motion along the y-axis represents lateral velocity. Rotational motion is described using Euler angles, where roll, ϕ, is defined as rotation about the x-axis, pitch, θ, as rotation about the y-axis and yaw, ψ, as rotation about the z-axis.

2.2.2 Reference Model

A reference model is a mathematical representation of the desired system behaviour, used to define how the system should respond under certain operating conditions. It provides a target trajectory or signal that the controller aims to follow.

In the context of vehicle dynamics, a reference model is essential because the control objective is not only to maintain stability but also to achieve a desired motion. For lateral dynamics, this is typically expressed in terms of a desired yaw rate or vehicle response to steering inputs. The reference model, therefore, captures what is considered stable, predictable, and physically reasonable vehicle behaviour. In this thesis, the reference model is used to generate a desired yaw rate, which serves as a benchmark for the controller. By comparing the measured yaw rate to this reference, a tracking error can be defined and minimised through control action. This allows the controller to both stabilise the vehicle and enforce a desired dynamic response.

Additionally, the use of a reference model provides a consistent framework for evaluation. While different control strategies may use different formulations, they can all be evaluated based on their ability to track the same reference behaviour. This ensures a logical comparison between the controllers. One common way to construct a reference model in vehicle dynamics is by using the bicycle-/one-track model.

2.2.3 One-track Model

The one-track model is a mathematical representation of a vehicle in which the left and right wheels on each axle are combined into a single equivalent wheel, resulting in one wheel per axle. This simplification reduces complexity while retaining key behaviours such as yaw rate, lateral motion, and side-slip angle. The one-track model is industry standard and is generally sufficient to capture the essential vehicle dynamics. The linear one-track models assume small steering angles, linear tyre models, low lateral velocities and allow for the use of linear ordinary differential equations (ODEs) to describe vehicle motion. Due to their linear structure, linear control methods can be applied. Non-linear vehicle models exist and better represent a vehicle's dynamics, but they are, in turn, more complex.

2.2.4 Tire Dynamics

Tire modelling plays a crucial role in vehicle dynamics, as the tyres are the primary interface through which forces and moments are transmitted between the vehicle and the road. Accurate tyre models are therefore essential for describing longitudinal, lateral, and yaw dynamics of the vehicle. The generated tyre forces depend primarily on the tyre slip ratio and slip angle.

The slip angle is defined as the angle between the direction of the wheel's heading and the actual direction of travel of the tyre contact patch. Slip angles mainly give rise to lateral tyre forces, which are responsible for vehicle cornering behaviour. The slip ratio, on the other hand, describes the relative difference between the wheel's circumferential speed and the longitudinal vehicle speed. Slip ratio governs the generation of longitudinal tyre forces during acceleration and braking.

For small slip angles and slip ratios, the tyre operates in the so-called linear region. In this region, the relationship between lateral force and slip angle can be approximated as linear. The proportionality constant between the lateral tyre force and slip angle is known as the cornering stiffness. Cornering stiffness is a key parameter in vehicle dynamics, as it characterises how responsive the tyre is to steering inputs and directly influences vehicle handling and stability. When the tyre operates beyond the linear region, non-linear effects such as force saturation and reduced responsiveness occur due to limitations in the tyre. However, for many vehicle models and stability analyses, assuming linear tyre behaviour is a reasonable approximation under normal driving conditions.

2.2.5 ABS

An anti-lock braking system (ABS) is a fundamental active safety system designed to prevent wheel lock-up during braking events. ABS works by maintaining wheel slip within an optimal range, typically close to the peak of the force generated by the tyre. ABS ensures that both longitudinal stability and steering capability are preserved during high deceleration manoeuvres. The system operates by continuously monitoring wheel rotational speeds and comparing them to an estimated vehicle reference speed. When excessive wheel slip is detected, the braking pressure is modulated through rapid actuation of the brake valves. This modulation follows a cycle of increasing and decreasing brake pressure, which allows the wheel to regain traction before reapplying braking force. In heavy-duty vehicles, ABS plays a particularly important role due to large vehicle mass, varying load conditions, and complex pneumatic braking systems. These factors increase the risk of wheel lock-up and instability, especially under low-friction or split-μ road conditions. This system is of importance in the scope of this thesis, as it is implemented in all modern trucks and therefore cannot be ignored or considered a limitation.

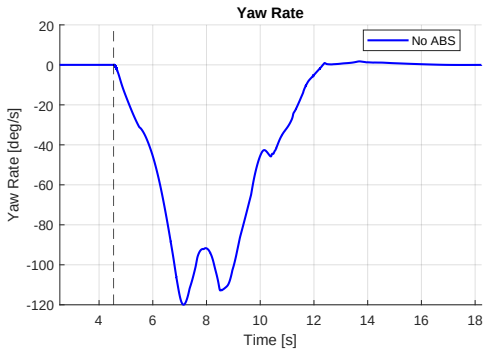

Figure 2.2: Split-μ braking without ABS.

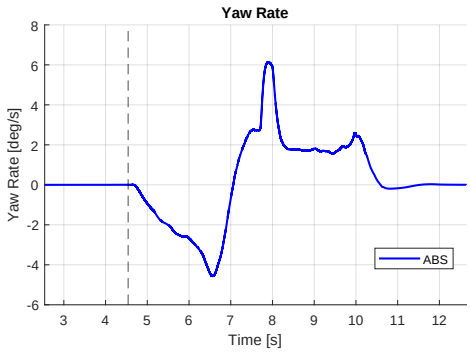

Figure 2.3: Split-μ braking with ABS.

The effects of ABS can be observed in Figures 2.2 and 2.3. The results were obtained from a split-μ braking scenario at a speed of 80 km/h, where the road friction differs between the left and right wheel paths ($\mu = 0.2$ and $\mu = 1.0$). The dashed line indicates when the split-μ braking begins.

Without ABS, Figure 2.2, the yaw rate increases significantly, indicating a strong yaw moment caused by the asymmetric braking forces. This results in highly unstable vehicle behaviour. In contrast, Figure 2.3 shows the yaw rate for the same scenario but with ABS activated. A substantial reduction in yaw rate can be observed, demonstrating the effectiveness of ABS in stabilising the vehicle under split-μ conditions.

This observation forms the starting point of this thesis and provides a baseline for comparisons. While ABS improves stability by regulating longitudinal slip, residual yaw motion remains. The goal of this thesis is therefore to further enhance vehicle stability by reducing yaw rate as well as other stability metrics through the devel-

opment of control strategies utilising rear wheel steering. This setup with ABS only will hereafter be referred to as the baseline and will be represented by the colour black in all plots.

2.2.6 Axle-ackermann Controller

As mentioned in the introduction, current implementations of rear axle steering systems are primarily designed to improve low-speed manoeuvrability. This behaviour can be achieved through axle-Ackermann steering geometry. While this reduces tyre wear and improves low-speed manoeuvrability, it is of interest for this thesis to test and evaluate its performance during a split-μ braking scenario.

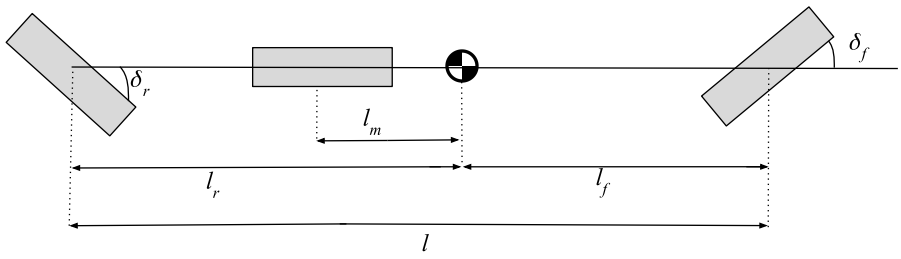

Figure 2.4: Three axle truck with steerable rear axle.

The rear axle steering angle using axle-Ackermann steering is determined as a function of the front steering angle. The rear wheels are steered in counter-phase relative to the front wheels, meaning that the rear axle steers in the opposite direction. This behaviour is consistent with the axle-Ackermann steering principle, in which the wheel angles are arranged to minimise tyre slip during low-speed cornering.

Based on this concept, the rear steering angle can be computed using a geometric relationship derived from the vehicle model in Figure 2.4:

$$\delta_{\mathrm{r}} = -\arctan\left(\tan(\delta_{\mathrm{f}}) \cdot \frac{l_r - l_m}{l_f + l_m}\right), \tag{2.1}$$

where:

- δ_{f} is the front steering angle,
- l_f is the distance from the CoG to the front axle,
- l_m is the distance from the CoG to the middle axle,
- l_r is the distance from the CoG to the rear axle.

The negative sign ensures counter-phase steering between the front and rear axles. This formulation provides a purely kinematic feed-forward control law that does not account for more complex vehicle dynamics such as tyre forces.

The results obtained from testing the axle-Ackermann controller, which computes

the rear steering angle according to equation 2.1 and with ABS active, are presented in Figure 2.5, where the resulting yaw rate is shown together with the baseline. The test is conducted under the same conditions as the scenario comparing no ABS to ABS active, Figures 2.2 and 2.3, a split-μ braking scenario at 80 km/h with friction coefficients of 0.2 and 1.0. The dashed line indicates when the split-μ braking begins.

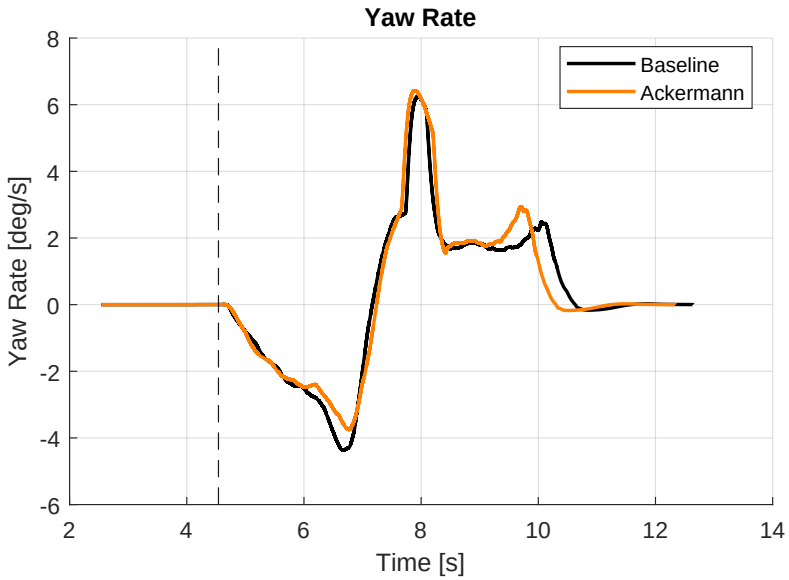

Figure 2.5: Split-μ with axle-Ackermann controller.

Compared with the baseline, a slight reduction in the initial yaw rate is observed. The initial yaw rate for the baseline after braking reaches slightly more than 4 deg/s, while the Ackerman controller manages to reduce the yaw rate down to around 3 deg/s in Figure 2.5. However, the overall improvement is not substantial, and a significant yaw rate remains. This indicates that the implemented Ackerman-based control strategy provides only marginal stabilisation under these conditions. These results highlight the need for a more advanced control approach, as the current implementation does not adequately address vehicle stability in demanding scenarios such as split-μ braking at higher speeds. Although the axle-Ackermann-based controller is not intended for such operating conditions, this evaluation provides useful insight into its limitations when applied outside its design range.

2.3 Tools

This section presents the primary tools used in this thesis. An in-depth technical description of each tool will not be provided; instead, a brief introduction and an overview of how they interact will be provided. Detailed explanations of how specific components were applied and used are provided later in Chapter 3.

2.3.1 MATLAB/Simulink

MATLAB and Simulink were used for controller development, implementation, and analysis. MATLAB was used for offline computations, including model generation, linearization, gain scheduling, and analysis of system properties such as controllability and stability, as well as post-processing of data to generate plots. Simulink provides a graphical environment for modelling dynamic systems and designing control algorithms, and it was used to architect the controller logic and manage the real-time data and signal flow. Its block-diagram structure allows for intuitive visualisation of the entire control loop. This thesis was conducted using MATLAB version R2021b.

2.3.2 IPG Truckmaker

To develop and validate the proposed control strategies, a high-fidelity vehicle simulation environment is required. IPG TruckMaker was used as the primary vehicle dynamics simulation platform. TruckMaker is a physics-based simulation software specifically designed for heavy-duty vehicles such as trucks and buses. It supports complex vehicle configurations, including multi-axle systems, suspensions, and advanced braking systems, making it well-suited for studying the dynamics of commercial vehicles. It allows accurate simulation of non-linear effects that are particularly important during critical manoeuvres such as split-μ braking. In addition to the vehicle model, TruckMaker provides modules for driver behaviour and environment modelling. The built-in driver model, IPGDriver, is used to follow predefined steering and speed profiles, while the road model allows specification of varying friction conditions. TruckMaker was selected due to its ability to accurately capture heavy-vehicle dynamics, particularly lateral behaviour and tire-road interactions, and its integration with MATLAB/Simulink. This makes it suitable for evaluating rear-wheel steering control strategies under realistic operating conditions.

Truckmaker allows for the creation of custom road profiles, including variations in geometry, curvature, and surface friction. This makes it possible to design and test specific scenarios such as split-μ conditions, where different friction levels exist on the left and right wheel paths. In addition, the built-in driver model can be modified to better represent human driving behaviour. Parameters such as reaction time and steering aggressiveness can be adjusted, enabling more realistic interaction between the driver and the vehicle control system. Finally, the vehicle model can be configured to match specific heavy-duty configurations, such as standard 6×2 truck layouts. This allows the simulation environment to closely replicate real-world vehicle properties, including mass distribution, axle configuration and suspension characteristics. This thesis was conducted using IPG TruckMaker version 13.1.3.

2.3.3 Co-simulation

To evaluate the developed control strategies, the co-simulation framework between MATLAB/Simulink and IPG TruckMaker was used. In this setup, both environments run simultaneously and exchange signals in a closed-loop configuration. As

mentioned above, Simulink was used for the implementation of the control algorithms, but it also enables clear visualisation of signals and internal variables, which simplifies debugging and analysis of the controller's behaviour. TruckMaker provides a predefined Simulink interface, where the vehicle model is integrated as a subsystem. Within this framework, the control algorithms are implemented in dedicated blocks, while vehicle states such as velocity, yaw rate, and steering inputs are accessed through interface blocks provided in the TruckMaker library. During simulation, TruckMaker computes the vehicle dynamics and makes the relevant signals available to Simulink. The controller processes these signals and computes the corresponding control inputs, which are then fed back to TruckMaker. This establishes a closed-loop interaction between the high-fidelity vehicle model and the control system. The use of co-simulation allows the advantages of both environments to be utilised. Simulink offers flexibility in controller design and signal monitoring, while TruckMaker provides a detailed representation of vehicle dynamics. This combination enables efficient development and evaluation of control strategies close to realistic operating conditions.

3
Methods

This chapter describes the methodology used in this thesis. It includes the evaluation framework, test scenarios and performance metrics, the development of the simulation environment, the implementation of the ABS, the construction of the reference model, and the design of the different controllers.

3.1 Tests and Evaluation

To evaluate the performance of the different controllers, they were compared with a baseline. As mentioned in 2.2.5, the baseline consisted of a truck with a non-steerable rear axle but including ABS. The different braking scenarios are meant to represent common driving scenarios and are defined below.

- Straight road with split-μ braking, $v_x = 30, 50, 80$ [km/h]
- ISO Lane change during split-μ braking, $v_x = 30, 50, 80$ [km/h]
- Curved road with split-μ braking, $v_x = 30, 50, 80$ [km/h]
- Scenario with split-μ braking, delay and noise, $v_x = 30, 50, 80$ [km/h]

The delay represents the time from when the controller calculates the requested angle until it reaches the rear axle steering system. The noise represents sensor noise, influencing yaw rate measurements and consists of band-limited white noise.

In each scenario, different split-μ cases were tested. The pairs of coefficients of friction for split-μ that were tested are $[(1.0, 0.8), (1.0, 0.5), (1.0, 0.2)]$, representing increasing asymmetry between the left and right wheel paths. These cases were chosen to evaluate controller performance under progressively more severe disturbance conditions, in which larger friction differences generate stronger yaw moments and greater stability challenges.

The performance of the controllers on the different braking scenarios will be evaluated based on the following metrics:

- Yaw rate [deg/s]
- Yaw angle [deg]
- Breaking distance [m]
- Lateral deviation [m]
- Steering wheel input [deg]
- Rear wheel angle [deg]

These metrics were selected to capture different aspects of vehicle stability and control performance. The yaw rate reflects the vehicle's rotational behaviour and is a direct indicator of directional stability. The braking distance is included to assess longitudinal performance and ensure that stability improvements do not come at the cost of increased braking distance. Lateral deviation, along with yaw angle, measures the vehicle's ability to maintain its intended path and steering wheel input represents the driver's effort and provides insight into how much assistance the controllers offer. A reduction in required steering input indicates that the controller supports the driver in maintaining stability.

Additionally, the magnitude of the control input will be taken into account when evaluating the controllers. The demanded rear steering angle is considered a measure of control effort. If one controller achieves only marginally better performance but requires significantly larger control inputs, it may be less desirable. Excessive rear wheel angles can be limited by actuator constraints, energy consumption, and mechanical wear. Therefore, a controller that achieves comparable performance with smaller control inputs may be preferable, especially if the objective is to maintain moderate rear steering angles. This provides additional insight when comparing controllers, allowing for discussion of advantages and disadvantages. The control input is therefore taken into account when ranking the controllers from a practical implementation perspective.

Further analysis of the requirements for the rear-axle steering system to achieve even greater stability and braking improvements will be conducted. Simulation will reveal when the truck is limited by the RAS or the control strategy. Simulation with a better performance than the current rear axle system could provide insights into what performance can be gained by such a system.

During the development of the controllers, it must also be acknowledged that the solution must be general enough to function and operate in the same way, independent of the truck configuration it is used in. It should not require any major modifications or tailoring to function on a new vehicle. For example, the controllers should not need to know the tyre compound to properly perform. In other words, the controllers should be based on measurable, inherent properties of the vehicle. An example of such a system is ABS. This system does not change its functionality for a heavier, longer, or wider truck. It maintains the same algorithm and bases its control decisions on the same type of data regardless of what vehicle it is applied to.

3.2 Simulation Enviroment

The following section describes the setup of the simulation environment used in this thesis. The environment consists of different test road scenarios with varying surface friction and geometries, as well as the test vehicle and its parameters, characteristics and the properties of the driver used.

3.2.1 ABS Implementation

ABS functionality was required, but none of the available IPG truck models included an implemented ABS. However, several IPG car models provided a complete ABS implementation. Therefore, the ABS system from a car model was modified and integrated into the truck model to ensure proper functionality under truck-specific dynamics and parameters. An overview of the ABS system can be viewed in Fig. 3.1.

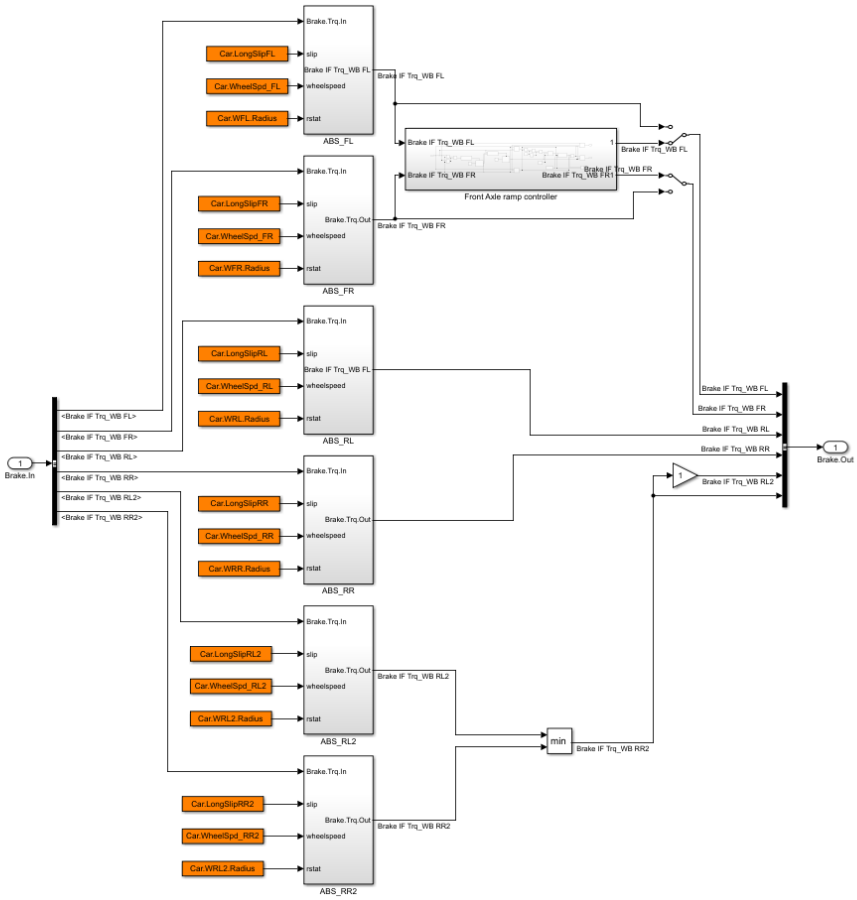

Figure 3.1: Overview of the ABS system.

The six subsystems represent the ABS control logic for the six wheels. The ABS algorithm implemented is based on the trigger criterion S_K, which combines wheel deceleration control and slip control using weighting factors. The control criterion is defined as

$$S_K = K_z \, z_R + K_\lambda \, \lambda_{\text{rel}}, \tag{3.1}$$

where z_R denotes the wheel deceleration and λ_{rel} the relative wheel slip, and K_z and K_λ are weighting factors that balance two different physical control objectives. This control strategy follows the formulation presented in [7]. Additionally, a delay in the logic was added to represent pneumatic and controller delays that a real ABS system would have.

The front axle has a ramp controller. Its purpose is to help the driver by introducing the split braking force on the front axle slowly. This controller works by restricting the brake torque, or in real-world implementations, the brake pressure, to be equal to the minimum requested by the left and right sides. It then ramps up the brake torque to its requested value over a set time. The difference in brake torque between left and right is initially zero and ramps up over a set time, and ends with individually controlled left and right tyres.

In contrast to the front axle ramp controller, the rear axle is controlled using a simpler strategy, a minimum selection function. This approach selects the smaller of the two rear brake torque commands and applies it to both rear wheels at all times.

3.2.2 Driver

The driver's properties are presented here and can be viewed in Tab. 3.1. The driver model and its properties were modified to represent an average truck driver, including their ability to act and react.

Parameter	**Value**	**Unit**
Max steering wheel angle	630	deg
Max steering wheel angle speed	500	deg/s
Max steering wheel acceleration	3000	deg/s^2
Max steering wheel torque	20	Nm
Max steering wheel torque velocity	1000	Nm/s
Lateral reaction time	0.5	s
Longitudinal reaction time	0	s

Table 3.1: Driver properties and constraints.

3.2.3 Vehicle

The truck used in this thesis is a generic 6x2 configuration. The vehicle can be viewed in Figure 3.2 and its main properties can be viewed in Table 3.2.

Figure 3.2: Simulation Vehicle.

Property	**Value**	**Unit**
m (Total mass)	25200	kg
Total weight	247111.2	N
Total roll inertia	23012.185	kgm^2
Total pitch inertia	88151.138	kgm^2
Total yaw inertia	88132.073	kgm^2
l (Total length)	7.42	m
CoG height (z)	1.435	m
CoG position (x)	5.447	m
CoG position (y)	0.000	m
l_f (Front axle to CoG)	3.843	m
l_m (Middle axle to CoG)	1.987	m
l_r (Rear axle to CoG)	3.577	m

Table 3.2: Vehicle properties with corresponding values and units.

3.2.4 Roads

Four different road scenarios were created: a straight road, an ISO lane-change manoeuvrer, a curved road with low friction on the inside and a curved road with low friction on the outside. Each road scenario was created using three different low-friction coefficients, $\mu = 0.2$, $\mu = 0.5$, and $\mu = 0.8$, representing low-, medium-, and high-friction conditions, respectively. In addition, each scenario was simulated at three vehicle speeds: 30 km/h, 50 km/h, and 80 km/h. This results in a total of 36 individual simulation tests for each controller. The test roads can be viewed in the Figures 3.3 to 3.6 below.

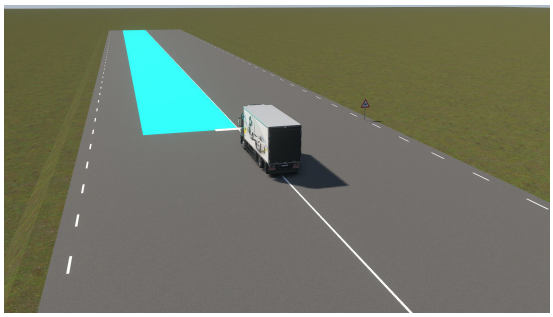

Figure 3.3: Straight road with split-μ friction.

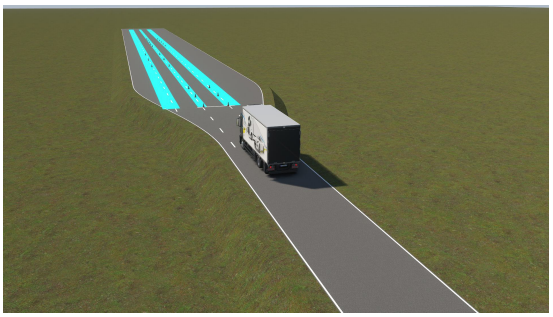

Figure 3.4: ISO lane change with winter road friction conditions.

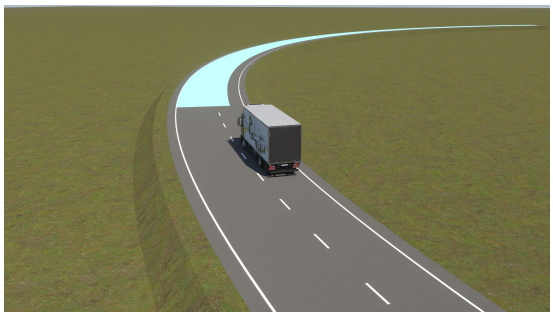

Figure 3.5: Curved road with low friction on the outside ($\varnothing 400m$).

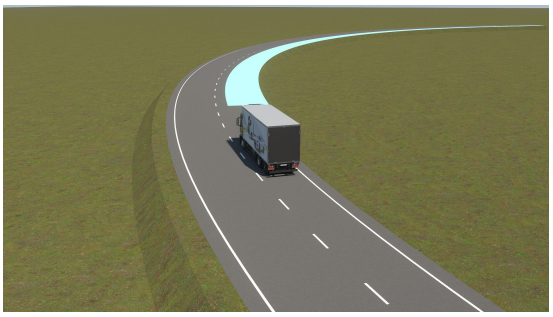

Figure 3.6: Curved road with low friction on the inside ($\varnothing 400m$).

3.3 Reference Model

The truck model being simulated in Truckmaker is of high fidelity and includes the many non-linearities of a real truck. The model includes dynamics such as body, tyres, suspension, power-train, aerodynamics, steering, brakes, and more. TruckMaker thus provides a simulation environment capable of representing real-world vehicle behaviour to a high degree, making it a suitable tool for developing vehicle dynamics-related applications. While this level of detail is advantageous, the development of controllers often benefits from a more simplified system representation to more effectively apply common control theory concepts and techniques. Therefore, a simpler reference model was required.

In order to acquire a reference model, a one-track model representation of the truck was needed. The one-track model includes yaw rate as a state, which is important for this thesis, since split-μ braking gives rise to a yaw moment that affects the dynamics of the truck.

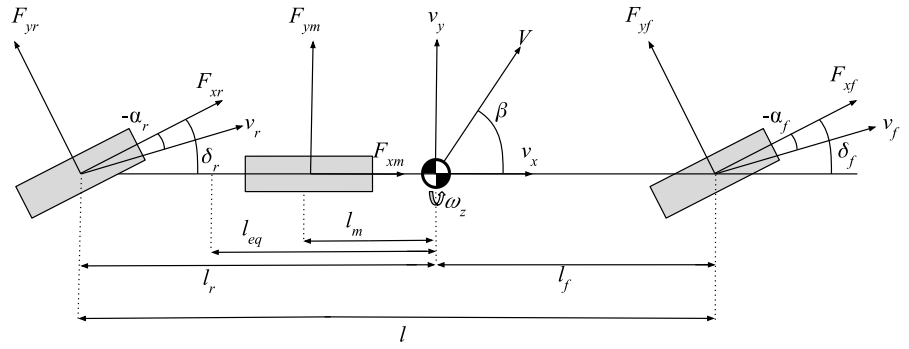

Figure 3.7: Three axle one-track model.

A three-axle one-track model with rear-axle steering was used and can be viewed in Figure 3.7. The lateral and yaw dynamics are given by the following equations of motion:

$$\sum F_y : m(\dot{v} + v_x\omega_z) = F_{yf} + F_{ym} + F_{yr}, \tag{3.2}$$

$$\sum M_z : I_z\dot{\omega}_z = F_{yf}l_f - F_{ym}l_m - F_{yr}l_r, \tag{3.3}$$

where f, m, r denote the front, middle and rear axles, respectively. The lateral force can be expressed in the following way:

$$F_{y,i} = C_i arctan(\alpha_i), \tag{3.4}$$

where C_i is the cornering stiffness, $i \in \{f, m, r\}$ and α_i is the slip angles which are defined as:

$$\tan(\delta_f - (-\alpha_f)) = \frac{v_y + l_f\omega_z}{v_x}, \tag{3.5}$$

$$\tan(\delta_r - (-\alpha_r)) = \frac{v_y - l_r\omega_z}{v_x}. \tag{3.6}$$

While the middle, non-steerable, axles slip angle is defined as:

$$\alpha_m = \arctan\left(\frac{v_y + l_m\omega_z}{v_x}\right). \tag{3.7}$$

Solving for α_f and α_r gives:

$$\alpha_f = -\delta_f + \arctan\left(\frac{v_y + l_f\omega_z}{v_x}\right), \tag{3.8}$$

$$\alpha_r = -\delta_r + \arctan\left(\frac{v_y - l_r\omega_z}{v_x}\right). \tag{3.9}$$

Assuming small slip angles ($|\alpha_i| \ll 1$) and using $\arctan(x) \approx x$, the linearized slip angles are:

$$\alpha_f \approx -\delta_f + \frac{v_y + l_f\omega_z}{v_x}, \tag{3.10}$$

$$\alpha_m \approx \frac{v_y + l_m\omega_z}{v_x}, \tag{3.11}$$

$$\alpha_r \approx -\delta_r + \frac{v_y - l_r\omega_z}{v_x}. \tag{3.12}$$

The linearized lateral forces can be expressed as:

$$F_{yf} = C_f\alpha_f = C_f\left(-\delta_f + \frac{v_y + l_f\omega_z}{v_x}\right), \tag{3.13}$$

$$F_{ym} = C_m\alpha_m = C_m\left(\frac{v_y + l_m\omega_z}{v_x}\right), \tag{3.14}$$

$$F_{yr} = C_r\alpha_r = C_r\left(-\delta_r + \frac{v_y - l_r\omega_z}{v_x}\right). \tag{3.15}$$

Substituting the linearized lateral forces into the equations of motion 3.2 and 3.3 we get the following:

$$m(\dot{v}_y + v_x\omega_z) = C_f\left(-\delta_f + \frac{v_y + l_f\omega_z}{v_x}\right) + C_m\frac{v_y + l_m\omega_z}{v_x} + C_r\left(-\delta_r + \frac{v_y - l_r\omega_z}{v_x}\right) \Rightarrow \tag{3.16}$$

$$\Rightarrow \dot{v}_y = -v_x\omega_z - \frac{C_f + C_m + C_r}{mv_x}v_y - \frac{C_f l_f + C_m l_m - C_r l_r}{mv_x}\omega_z + \frac{C_f}{m}\delta_f + \frac{C_r}{m}\delta_r \tag{3.17}$$

$$I_z\dot{\omega}_z = l_f C_f\left(-\delta_f + \frac{v_y + l_f\omega_z}{v_x}\right) + l_m C_m\frac{v_y + l_m\omega_z}{v_x} - l_r C_r\left(-\delta_r + \frac{v_y - l_r\omega_z}{v_x}\right) \Rightarrow \tag{3.18}$$

$$\Rightarrow \dot{\omega}_z = \frac{-C_f l_f + C_m l_m + C_r l_r}{I_z v_x}v_y + \frac{-C_f l_f^2 - C_m l_m^2 - C_r l_r^2}{I_z v_x}\omega_z + \frac{C_f l_f}{I_z}\delta_f - \frac{C_r l_r}{I_z}\delta_r \tag{3.19}$$

Using the equations above, the three-axle one-track model can be represented in the following state-space,

$$\dot{x} = Ax + Bu, \tag{3.20}$$

$$y = Cx + Du, \tag{3.21}$$

$$x = \begin{bmatrix} v_y \\ \omega_z \end{bmatrix}, u = \begin{bmatrix} \delta_f \\ \delta_r \end{bmatrix}, \tag{3.22}$$

$$\begin{bmatrix} \dot{v_y} \\ \dot{w_z} \end{bmatrix} = \begin{bmatrix} \frac{-C_f - C_m - C_r}{mv_x} & \frac{-C_f l_f + C_m l_m + C_r l_r}{mv_x} - v_x \\ \frac{-C_f l_f + C_m l_m + C_r l_r}{I_z v_x} & \frac{-C_f l_f^2 - C_m l_m^2 - C_r l_r^2}{I_z v_x} \end{bmatrix} \begin{bmatrix} v_y \\ \omega_z \end{bmatrix} + \begin{bmatrix} C_f/m & C_r/m \\ C_f l_f/I_z & -C_r l_r/I_z \end{bmatrix} \begin{bmatrix} \delta_f \\ \delta_r \end{bmatrix}, \tag{3.23}$$

$$y = \begin{bmatrix} 0 & 1 \end{bmatrix} x + 0 * u. \tag{3.24}$$

Given the parameters listed in Table 3.2, the one-track model can be constructed. However, the cornering stiffness of each axle, C_f, C_m, C_r [N/rad], along with the longitudinal velocity v_x [m/s], is not directly available. These parameters, therefore, had to be obtained separately or estimated.

3.3.1 Cornering Stiffness

Two approaches were used to obtain the cornering stiffness.

- Slope of linear region around origin in F_y/sideslip plot.
- Grey-box estimation using simulated data.

The F_y/sideslip plot was obtained through a manoeuvrer consisting of a ramp-steer at different velocities. To ensure accurate estimation, pure sideslip conditions are desired; therefore, longitudinal tyre forces must be minimised. This was done by initialising the simulation at a speed and then letting the truck free roll without any power-train, brakes, or aerodynamic forces slowing the truck down. Then each wheel's lateral forces can be plotted as a function of its sideslip. At each axle, the cornering stiffness is the sum of the average of the left and right wheels. Figure 3.8 shows lateral forces on each tyre with respect to each tyre's sideslip.

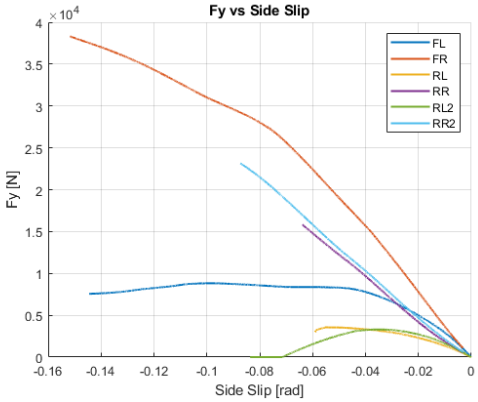

Figure 3.8: F_y with respect to wheel sideslip angle during ramp steer left.

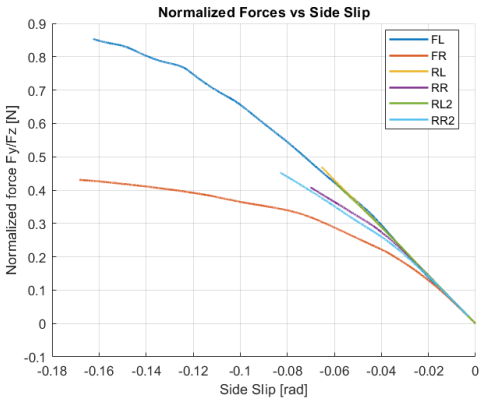

Figure 3.9: F_y/F_z with respect to wheel sideslip during ramp steer left.

The cornering stiffness for each wheel is extracted from the slope in the linear region close to zero slip according to:

$$C_i = \frac{\partial F i_y}{\partial \alpha_i}, \quad |\alpha| < 1°. \tag{3.25}$$

The total axle cornering stiffness is the sum of cornering stiffness from each wheel on the corresponding axle, averaged from tests at different speeds. Table 3.3 shows the test data and the final axle corner stiffnesses.

Cornering Stiffness	**80 km/h**	**100 km/h**	**120 km/h**
FL	350k	332k	345k
FR	378k	411k	393k
RL	126k	132k	132k
RR	177k	188k	204k
RL2	193k	190k	179k
RR2	250k	256k	258k
Average			
C_f	364k	371k	369k
C_m	152k	160k	168k
C_r	222k	223k	219k
Overall Average			
C_f		368k [N/rad]	
C_m		160k [N/rad]	
C_r		221k [N/rad]	
Normalized Cornering stiffness			
$C_{normalized}$		10.26 [N/rad$*F_z$]	

Table 3.3: Cornering stiffness estimated from F_y vs. sideslip angle plots.

The other approach used was grey-box estimation. By setting the cornering stiffnesses to be an estimated parameter, the objective becomes to find an optimal estimated cornering stiffness in order to reduce the error between the state-space output and the yaw-rate captured by simulation. This will converge the one-track model or state-space closer to the real dynamics of the truck model at a given speed.

$$\begin{bmatrix} \dot{v_y} \\ \dot{w_z} \end{bmatrix} = \begin{bmatrix} \frac{-\hat{C_f}-\hat{C_m}-\hat{C_r}}{mv_x} & \frac{-\hat{C_f}lf+\hat{C_m}lm+\hat{C_r}lr}{mv_x} - vx \\ \frac{-\hat{C_f}lf+\hat{C_m}lm+\hat{C_r}lr}{I_z v_x} & \frac{-\hat{C_f}lf^2-\hat{C_m}lm^2-\hat{C_r}lr^2}{I_z v_x} \end{bmatrix} \begin{bmatrix} v_y \\ \omega_z \end{bmatrix} + \begin{bmatrix} \frac{\hat{C_f}}{m} \\ \frac{\hat{C_f}lf}{I_z} \end{bmatrix} \delta_f \tag{3.26}$$

A chirp steer test was used at different vehicle speeds to gather data for the grey-box estimation. The chirp steer signal ranged from 0.05Hz to 2Hz, thereby exciting the truck across a frequency range representative of driver-induced steering inputs. Figure 3.10 shows the front wheel angle and the corresponding yaw rate obtained from a chirp steer test conducted at 80 km/h. The collected data, together with the one-track model, is then fed into the greyest() solver in MATLAB's System Identification toolbox. The solver then outputs the cornering stiffness that reproduces the closest result to the simulated data and minimises the error between the one-track model and the real truck. Figure 3.11 shows the estimated cornering stiffness for different speeds along with a fit percentage indicating how well the one-track model represents the high-fidelity truck. Analysing the plot it can be seen that the front axle cornering stiffness does not deviate much and could be approximated as linear or perhaps as a static value. However, the cornering stiffness for the two rear-axles changes drastically with speed, even changing sign, this reveals that non-linearities are present and play a part in describing the forces and that a static value or a linear curve can not represent the system. The fit percentage is above 95% for all speeds, which indicates that the estimated cornering stiffnesses accurately represent the truck.

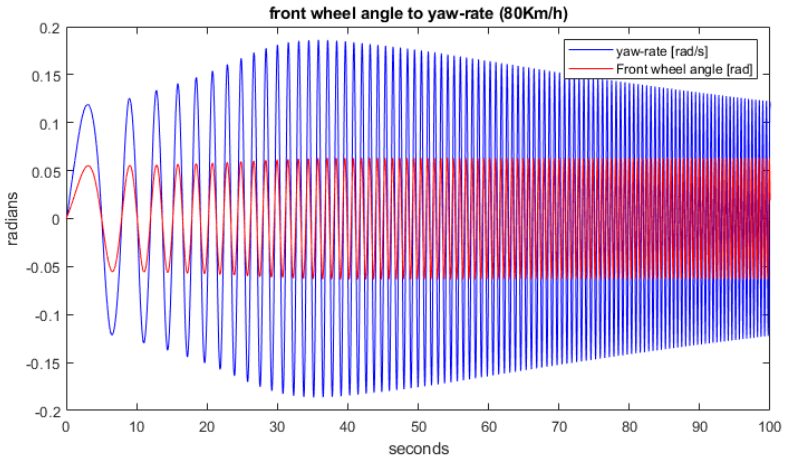

Figure 3.10: Front wheel angle to yaw rate.

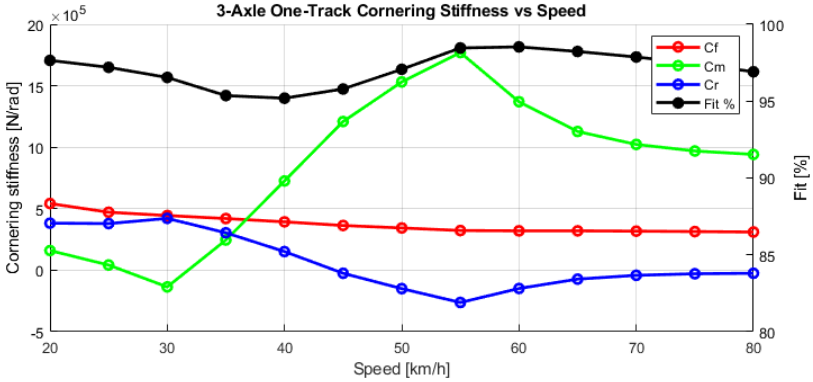

Figure 3.11: Estimated cornering stiffness with respect to speed in a 3-axle one-track model.

Due to that, the rear-axles showed signs of non-linearities, a two-axle one-track model with equivalent wheelbase was tested to see if it could represent the truck better. The following state-space describes the dynamics of the two-axle one-track model of the truck.

$$\dot{x} = Ax + Bu \tag{3.27}$$

$$y = Cx + Du \tag{3.28}$$

$$x = \begin{bmatrix} v_y \\ \omega_z \end{bmatrix} \quad u = \delta_f \tag{3.29}$$

$$\begin{bmatrix} \dot{v_y} \\ 2\dot{w_z} \end{bmatrix} = \begin{bmatrix} \frac{-C_f - C_{eq}}{mv_x} & \frac{-C_f l_f + C_{eq} l_{eq}}{mv_x} - vx \\ \frac{-C_f l_f + C_{eq} l_{eq}}{I_z v_x} & \frac{-C_f l_f^2 - C_{eq} l_{eq}^2}{I_z v_x} \end{bmatrix} \begin{bmatrix} v_y \\ \omega_z \end{bmatrix} + \begin{bmatrix} \frac{C_f}{m} \\ \frac{C_f l_f}{I_z} \end{bmatrix} \delta_f \tag{3.30}$$

$$y = \begin{bmatrix} 0 & 1 \end{bmatrix} x + 0 * u \tag{3.31}$$

Figure 3.12 shows the estimated cornering stiffnesses with respect to speed for the two axle one-track model with equivalent wheelbase. It can be observed that the cornering stiffnesses can be described linearly, but only for speeds under 45 km/h.

Below 45 km/h it fails to capture good estimations and the fit percentage drops. To address this, a linear regression was made on the data for speeds above 45 km/h and then extended backwards so the cornering stiffness could be described as a linear function over the entire speed range.

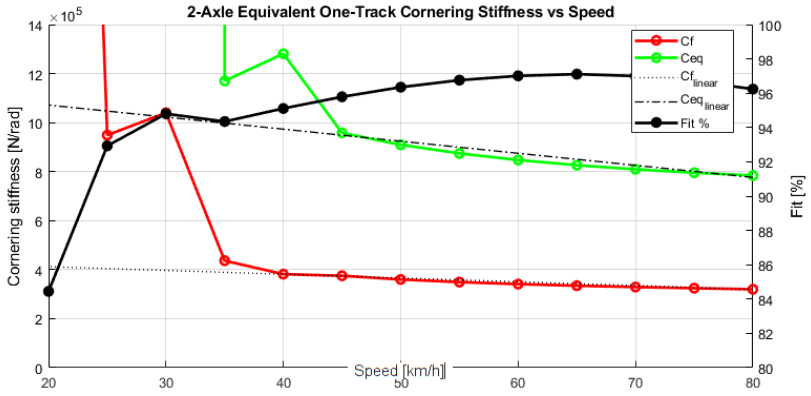

Figure 3.12: Estimated cornering stiffness with respect to speed in a 2-axle one-track model.

3.3.2 LPV Model

The cornering stiffness for the different speeds has now been estimated. However, since the cornering stiffnesses change with speed and the speed itself also changes during test runs/manoeuvres, a nominal linear time-invariant (LTI) model is not sufficient. It is possible to model the cornering stiffness and speed as uncertainties with weightings in the LTI system, but since the cornering stiffness and speed have been directly measured, the system can instead be modelled using Linear Parametric Varying (LPV) models. The LPV model can be mathematically described as follows:

$$\dot{x}(t) = A(\rho)x(t) + B(\rho)u(t), \tag{3.32}$$

$$y(t) = C(\rho)x(t) + D(\rho)u(t). \tag{3.33}$$

Where ρ is a parametric vector which is available in real-time and depends on the speed. The scheduling parameter was chosen as $\rho = v_x$. The front and rear cornering stiffnesses are not constant but vary with the longitudinal velocity.

$$C_f = C_f(v_x), \qquad C_m = C_m(v_x), \qquad C_r = C_r(v_x). \tag{3.34}$$

To construct the LPV, the non-linear dependence on v_x is approximated using a set of frozen LTI state-space models. A total of $N = 16$ operating points were used, one for each of the speeds tested in the data gathering.

$$v_x \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60, 65, 70, 75, 80\} \tag{3.35}$$

For each operating point, the corresponding cornering stiffnesses and speed is used to obtain a LTI model:

$$\dot{x} = A(v_{x,i})x + B(v_{x,i})u, \qquad i = 1, \ldots, 16. \tag{3.36}$$

The LPV system is obtained by scheduling between the individual frozen models as a function of the instantaneous velocity. Interpolation between all of the local LTI systems can be used to ensure continuous dependence on the scheduling parameter v_x. The system matrices can be expressed as a convex combination in the following way:

$$A(\rho) = \sum_{i=1}^{16} w_i(\rho)\, A_i, \tag{3.37}$$

$$B(\rho) = \sum_{i=1}^{16} w_i(\rho)\, B_i. \tag{3.38}$$

The resulting LPV one-track model with speed-dependent cornering stiffness can therefore be expressed as

$$\dot{x}(t) = \left(\sum_{i=1}^{16} w_i(v_x)A_i\right) x(t) + \left(\sum_{i=1}^{16} w_i(v_x)B_i\right) u(t). \tag{3.39}$$

A total of six LPV models were created:

- **LPV 3-axle F steering:** Speed-dependent cornering stiffness from estimation of a front wheel-steered truck.
- **LPV 2-axle equivalent wheelbase:** Speed-dependent cornering stiffness from estimation of a front wheel-steered truck.
- **LPV 3-axle Fy/sideslip:** Static cornering stiffness from obtained from plots.
- **LPV 3-axle with F&R steering:** Speed-dependent cornering stiffness from estimation of a front and rear wheel-steered truck.
- **LPV 3-axle with rear input gain:** LPV 3-axle F steering with a speed-dependent estimated gain on the rear wheel angle input.
- **LPV 3-axle with R steering:** Speed-dependent cornering stiffness from estimation of a rear wheel-only steered truck.

Figure 3.13, 3.14, 3.15, and 3.16 show the simulation yaw rate from the high fidelity TruckMaker model compared to the LPV models during a test-run on a virtual race-track. Figure 3.17 shows the speed and front wheel angle during the test run. Figure 3.18 shows the speed and front and rear wheel angles during the test run. Analysing Figure 3.13 shows that the LPV model, using speed-dependent cornering stiffness, provides an accurate representation of the high-fidelity model. The **LPV 3-axle F steering** and **LPV 2-axle equivalent wheelbase** models perform equally, while the static cornering stiffness LPV did not perform as well. For this thesis, **LPV 3-axle F steering** and **LPV 2-axle eq wheelbase** showed accurate enough performance to be used as reference generators for the control algorithm development. Figure 3.19 shows the step response with a step steer input with only front wheel steering, while Figure 3.20 shows the step response with a step steer input using both

front and rear wheel steering. Analysing the plots reveals that the **LPV 3-axle F steering** is more accurate when rear axle steering is not included. However, when both front and rear wheel steering is used, then **LPV 3-axle with F&R steering** is more accurate than **LPV 3-axle F steering** but still shows worse performance compared to only using front wheel steering.

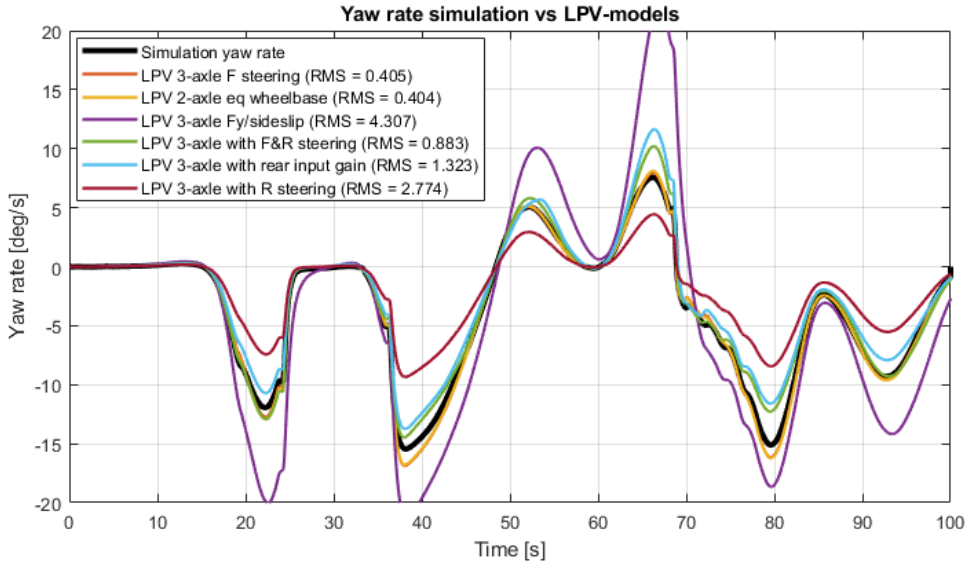

Figure 3.13: LPV models compared to simulation data from racetrack test run with only front wheel steering.

While root-mean-squared (RMS) is a good metric to evaluate model accuracy, it does not provide information about where the model is accurate or inaccurate. Therefore, a 3-dimensional mapping of the steady-state yaw rate as a function of speed and front-wheel angle was created. By simulating the truck at each operating point, one can then compare that to what the LPV produces. Figure 3.21 shows the yaw rate produced by the truck at different speeds and front wheel angles. Figure 3.22 shows the yaw rate produced by the **LPV 3-axle F steering**. Taking the error between the two, which can be viewed in Figure 3.23, the performance can be easily analysed. The region opposite to high speed and high front wheel angle is very accurate, which is indicated by the flat surface in the plot. At high speeds and high front angles, the linear tyre representation no longer holds, and the LPV overestimates the yaw rate. Additionally, the small-angle assumption becomes increasingly inaccurate under these conditions. In this region, the truck's tyres are saturated, and the slip is very high. This is acceptable, however, since such aggressive manoeuvres are unlikely to be induced by a driver. For example, a front wheel angle of 8 degrees at a speed of 80km/h could lead to a rollover or loss of control of the vehicle.

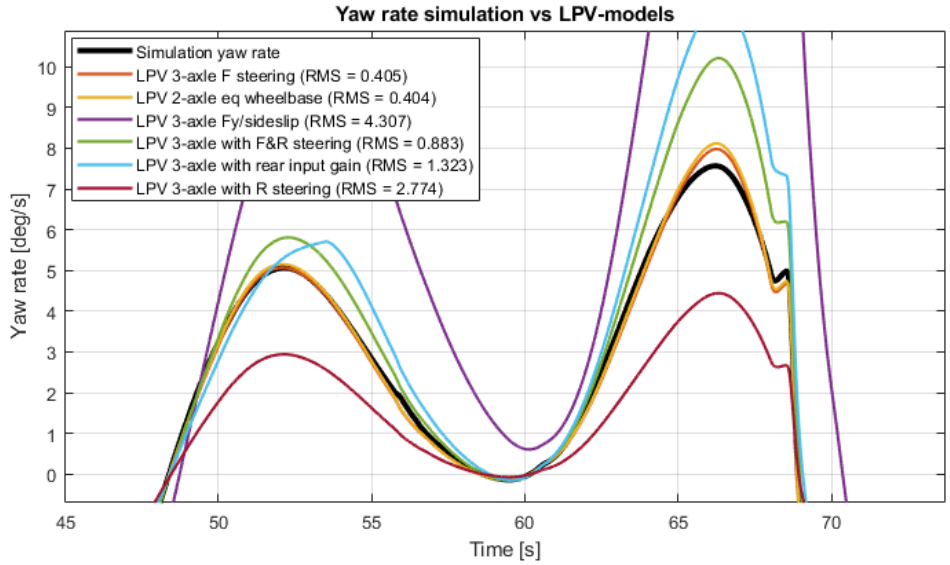

Figure 3.14: LPV models compared to simulation data from racetrack test run with only front wheel steering.

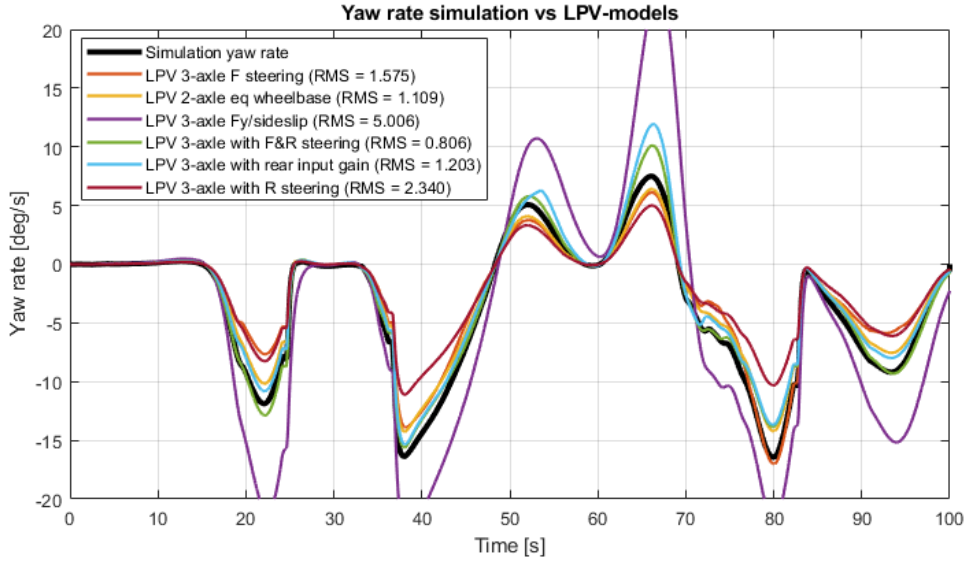

Figure 3.15: LPV models compared to simulation data from racetrack test run with front and rear wheel steering.

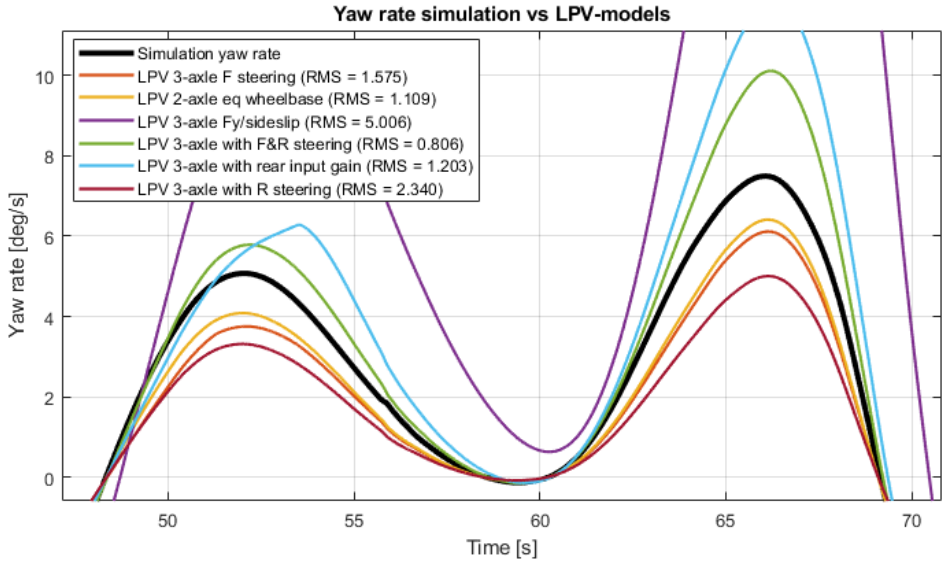

Figure 3.16: LPV models compared to simulation data from racetrack test run with front and rear wheel steering.

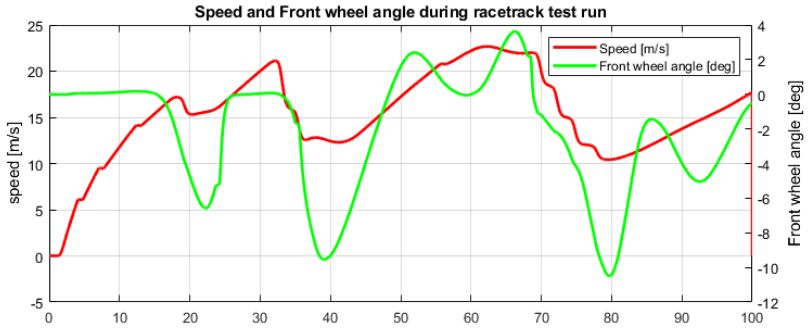

Figure 3.17: Speed and front wheel angle during racetrack test run.

Figure 3.18: Speed, front and rear wheel angle during racetrack test run.

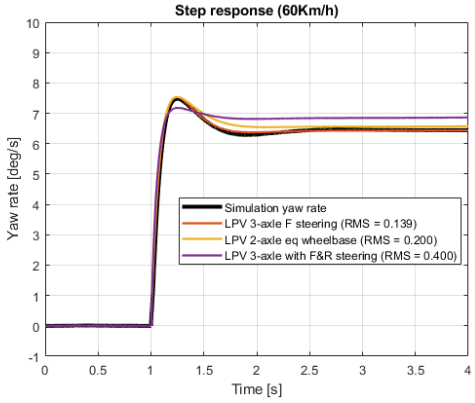

Figure 3.19: Step response from steering wheel to yaw rate with only front wheel steering.

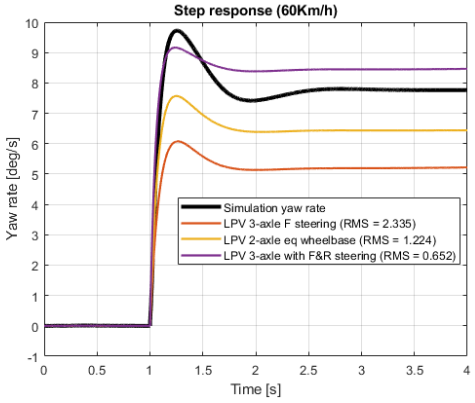

Figure 3.20: Step response from steering wheel to yaw rate with front and rear wheel steering.

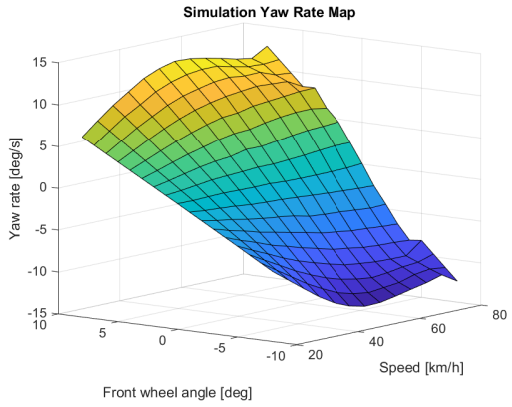

Figure 3.21: Steady state yaw rate from simulation.

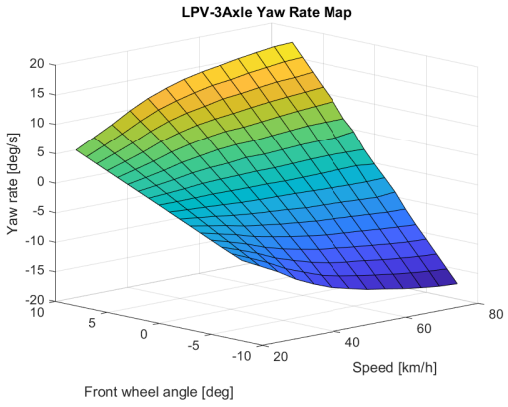

Figure 3.22: Steady state yaw rate from LPV 3-axle.

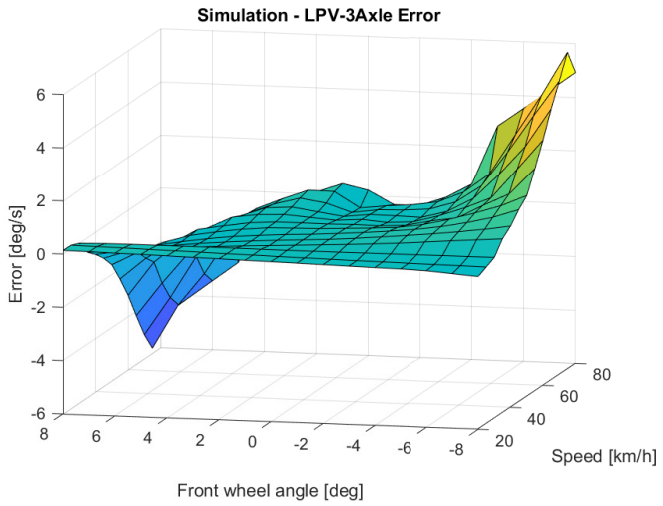

Figure 3.23: Error between predicted yaw rate from LPV and actual yaw rate from simulation.

3.4 Control

The following subsections describe the controllers developed and analysed. The general control objective is to track the reference yaw rate produced by the LPV reference generator. It could be seen as model reference control, where the controller tries to drive the truck to behave like the LPV system. Since the LPV represents how the truck behaves in normal conditions with no split-μ present, the controller will try to mimic this behaviour even during split-μ braking, hence making the truck predictable for the driver and making sure the driver's intentions are fulfilled.

The overall control structure is as follows. Initially, sensor readings are retrieved using the Truckmaker-specific Read CM Dict blocks in Simulink. These signals are then provided to the LPV model, which generates a reference signal. This reference, together with other sensor measurements such as yaw rate, is then supplied to the different controllers. Each controller computes a rear wheel angle (RWA), which serves as the control input. The control signal is then saturated and rate-limited according to Volvo confidential parameters. The resulting RWA is then applied to the rear wheels via the Write CM Dict blocks.

The trigger criteria for the controllers are that the brake pedal is pressed more than 50% and that the vehicle speed is non-zero. Thus, the controllers are not continuously active. This is intentional as small yaw rate errors are always present and continuous activation could lead to unnecessary control actions, potentially compromising safety.

3.4.1 PID

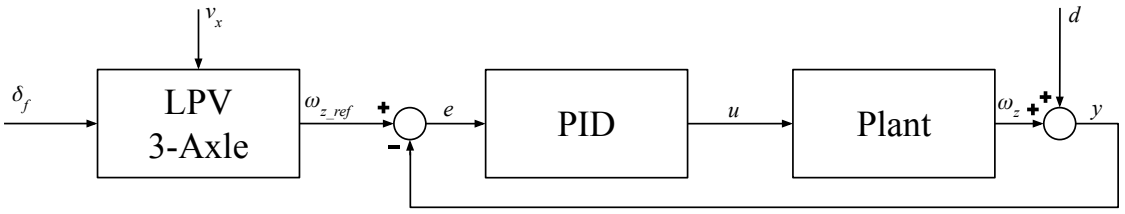

Figure 3.24: PID block diagram.

PID control is a classic control method and is very versatile since it can be easily adapted to systems where little or no information is known about the plant dynamics. In the case of this thesis control objective, a PID could be a valuable solution since it is often difficult to accurately derive the equations for the highly uncertain truck dynamics. The PID state-feedback control law is defined as:

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau) + d\tau + K_d \frac{de(t)}{dt}, \tag{3.40}$$

where $e(t)$ is the error signal, which is equal to the difference between the reference yaw rate and the measured plant yaw rate, $e(t) = ref_{\text{yawrate}} - y_{\text{yawrate}}$. Since the rear

axle has constraints regarding absolute max angle and rate limit, there is a need to include these constraints in the controller to ensure it still works properly even when close to the limits of the rear axle steering system. This was done by output saturation and integral anti-wind-up. This ensures that the PID won't continue to accumulate the error when the output is already at its limit. This gives a faster response once the system returns to the valid range. The saturation and integral wind-up are computed in the following way:

$$u_{\text{sat}}(t) = \text{clip}(u(t),\, u_{\min},\, u_{\max}) \tag{3.41}$$

$$\dot{I}(t) = \begin{cases} e(t), & u_{\text{sat}}(t) = u(t) \\ e(t), & u_{\text{sat}}(t) = u_{\max} \ \wedge \ e(t) < 0 \\ e(t), & u_{\text{sat}}(t) = u_{\min} \ \wedge \ e(t) > 0 \\ 0, & \text{otherwise} \end{cases} \tag{3.42}$$

where $u_{\min}, u_{\max}$ are the smallest and largest absolute angles that the rear axle system can produce. The reference yaw rate for the PID controller is provided from the **LPV 3-axle F steering** reference generator. The gains were determined by manual tuning during step inputs at different speeds. This tuning led to the following gains:

$$K_p = 10, \quad K_i = 5, \quad K_d = 0.5 \tag{3.43}$$

3.4.2 Gain-Scheduled LQR

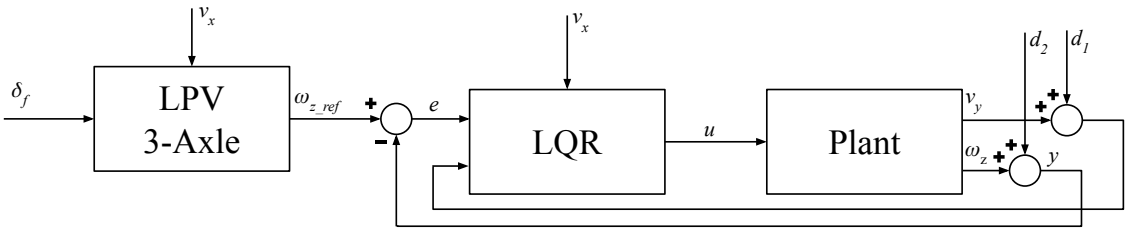

Figure 3.25: LQR block diagram.

To achieve stable and optimal control of the vehicle dynamics, a Linear Quadratic Regulator (LQR) can be used. LQR is an optimal state-feedback controller that determines the control input by minimising a cost function. For this thesis, the LQR framework provides a way to compute feedback gains for rear axle steering, enabling stabilisation of yaw dynamics during split-μ braking scenarios. A gain-scheduled LQR was developed using one of the LPV representations of the truck, specifically using the **LPV 3-axle with F&R steering** model. The LPV allows the controller to adapt to changes in longitudinal velocity, which is an important parameter influencing lateral vehicle behaviour during split-μ braking. The vehicle

model was expressed in state-space form, where the system matrices vary as a function of the longitudinal velocity v_x. A discrete set of operating points was defined over a velocity range of 20–80 km/h:

$$v_x \in [20, 25, 30, \ldots, 80] \text{ [km/h]}. \tag{3.44}$$

This range was selected to reflect the relevant speeds for the braking scenarios considered and was used as scheduling points for the LPV model. For each velocity, an LQR controller was synthesised based on the corresponding linearised state-space model:

$$\dot{x} = A(v_x)x + B(v_x)u. \tag{3.45}$$

The control objective is formulated through a quadratic cost function:

$$J = \int_0^\infty \left(x^T(t)Qx(t) + u^T(t)Ru(t) \right) dt, \tag{3.46}$$

where Q and R are weighting matrices. The matrix Q penalises deviations in the system states, while R penalises control effort. These matrices serve as tuning parameters that define the trade-off between performance and actuator usage.

By minimising the cost function, the optimal control law is obtained as a linear state-feedback:

$$u(t) = -Kx(t). \tag{3.47}$$

The gain matrix K is computed from:

$$K = R^{-1}B^T P, \tag{3.48}$$

where P is the solution to the continuous-time algebraic Riccati equation:

$$A^T P + PA - PBR^{-1}B^T P + Q = 0, \tag{3.49}$$

where the system matrices are parameter-dependent:

$$\begin{aligned} A &= A(v_x^{(i)}), \\ B &= B(v_x^{(i)}). \end{aligned} \tag{3.50}$$

The solution P is symmetric and positive semidefinite [6], and guarantees that the closed-loop system is stable provided that the system is controllable. Controllability was evaluated for each operating point by analysing the rank of the controllability matrix:

$$\mathcal{C} = \begin{bmatrix} B & AB & A^2B & \ldots & A^{n-1}B \end{bmatrix}, \tag{3.51}$$

where n is the number of states. The system is considered controllable if the controllability matrix has full rank, which is verified by checking the following:

$$\text{rank}(\mathcal{C}) = n. \tag{3.52}$$

Full rank was confirmed across all velocity points, ensuring that all system modes can be influenced by the rear steering input. Furthermore, the stability of the resulting

closed-loop system was verified by computing the eigenvalues of the closed-loop system matrix:

$$A_{cl} = A - BK. \tag{3.53}$$

The system is asymptotically stable if:

$$\mathbf{R}\{\lambda_i(A_{cl})\} < 0 \quad \forall i, \tag{3.54}$$

where $\lambda_i(A_{cl})$ denotes the eigenvalues of A_{cl}. All eigenvalues were found to lie in the left-half complex plane, confirming asymptotic stability of the closed-loop system for all scheduled velocities. This confirms that the gain-scheduled LQR controller provides a stabilising feedback law over the entire operating range considered.
The LQR gain matrix $K(v_x)$ was obtained by minimizing the quadratic cost function 3.46 where the weighting matrices were selected as:

$$Q = \begin{bmatrix} 1 & 0 \\ 0 & 100 \end{bmatrix}, \tag{3.55}$$

$$R = 5. \tag{3.56}$$

This places strong emphasis on yaw rate regulation while limiting excessive rear steering effort. This resulted in a set of gain matrices stored as:

$$K_{\text{LPV}} = \{K(v_x^{(i)})\}_{i=1}^{N}, N = \text{length}(v_x). \tag{3.57}$$

The controller implementation uses a one-dimensional lookup table of precomputed LQR gains, K_{LPV}, indexed by longitudinal velocity. The active gain is selected using a nearest-grid-point rule based on the current vehicle speed. The control law is defined and applied as:

$$u = -K(v_x)x_e, \tag{3.58}$$

where the state error is defined as:

$$x_e = \begin{bmatrix} v_y \\ r - r_{\text{ref}} \end{bmatrix}. \tag{3.59}$$

Here, v_y is the lateral velocity and r is the yaw rate, while r_{ref} denotes the reference yaw rate. Using this formulation, lateral velocity is regulated towards zero, while yaw rate is tracked with respect to a reference signal. This design choice was motivated by the fact that the LPV model provides reliable yaw rate tracking, whereas lateral velocity, v_y, tracking performance was less accurate. Therefore, v_y is directly included in the feedback to promote lateral stability rather than reference tracking. The gain matrices were precomputed offline for each velocity point using the MATLAB `lqr()` function.

3.4.3 $\mathcal{H}_\infty$

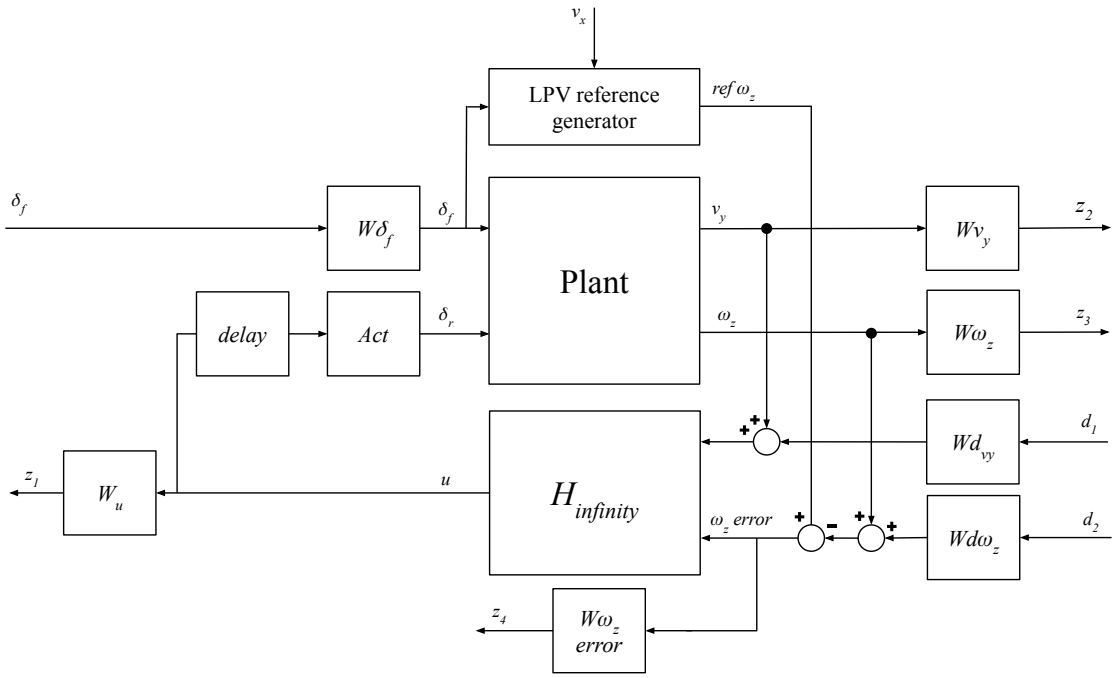

Figure 3.26: $\mathcal{H}_\infty$ block diagram.

The dynamics of a truck can change drastically due to tyre wear, weight, speed, etc. Therefore, a robust control method was implemented so that these changes can be modelled as uncertainties in the system. This allows for a more accurate plant model and will work under more circumstances with guaranteed stability and robust performance. $\mathcal{H}_\infty$ control, where with weightings one can shape the controller in the frequency domain to meet the requirements set for the controller, was used. $\mathcal{H}_\infty$ synthesis returns a controller, K, that stabilizes the plant, P, and achieves a performance level γ which is also the $\mathcal{H}_\infty$-norm of the closed loop system. The $\mathcal{H}_\infty$-norm is defined as:

$$||G(s)||_\infty = ||G(iw)||_\infty = max_w \bar{\sigma}(G(iw)). \tag{3.60}$$

It can be interpreted as the maximum amplification, which is the peak induced 2-norm [8]. In the time domain, this is formulated as:

$$||g(t)||_\infty = max_{u(t)\neq 0} \frac{||y(t)||_2}{||u(t)||_2} = max_{||u(t)||_2=1} ||y(t)||_2, \tag{3.61}$$

which can be viewed as the worst-case ratio between input and output signal [8]. The $\mathcal{H}_2$-norm is defined as:

$$||G(s)||_2 = ||G(iw)||_2 = \sqrt{\frac{1}{2\pi}\int_{-\infty}^{\infty} tr(G(-iw)^T tr(G(iw))dw}, \tag{3.62}$$

Where $tr(G(-iw_0)^T tr(G(iw)) = \sum_{i,j} |g_{i,j}(iw_0)|^2$. Or in the time domain as:

$$||g(t)||_2 = \sqrt{\int_0^\infty tr(g(-\tau)^T tr(g(\tau))d\tau}, \tag{3.63}$$

where $tr(g(-\tau))^T tr(g(\tau)) = \sum_{i,j} |g_{i,j}(\tau)|^2$.

Synthesising the controller was done in the following steps. The first step was to design the open loop connections, which is the block diagram without the $\mathcal{H}_\infty$ block and weights, presented in Figure 3.26. The weights were then added to shape tracking performance and reject disturbances. The Wd_{vy} and $Wd\omega_z$ weights capture measurement noise of the plant output vy and ω_z.

$$Wd_{vy}, Wd\omega_z = \frac{0.01}{s}. \tag{3.64}$$

The rear wheel steering actuator dynamics are added in the *Act* weight and allow the controller to take into account the limits of the actuator.

$$Act = \frac{12.25}{s^2 + 7s + 12.25}. \tag{3.65}$$

The delay from the requested rear wheel angle u until it reaches the rear axle steering actuator is modelled as a first-order system.

$$delay = \frac{1}{0.005s + 1}. \tag{3.66}$$

The weight for tracking error and $W\omega_z$ *error* is a low-pass filter that should provide good steady state tracking while rejecting high frequency yaw rate error.

$$W\omega_z = \frac{20}{s + 1}. \tag{3.67}$$

Both $W\omega_z$ and Wv_y emphasise suppression at low frequency where the vehicle handling region lies, while rejecting high frequency where noise and unmodeled dynamics exist. v_y is less penalised since that state is not equally important. $W\omega_z$ is used in conjunction with $W\omega$ *error* since there is a balance between low yaw rate and tracking the reference.

$$Wv_y = \frac{1}{s + 1}, \quad W\omega_z = \frac{6}{s + 1}. \tag{3.68}$$

Control effort is limited by W_u and limits the rear axle steering input by a static gain and a dynamic term that reduces chattering and oscillations.

$$W_u = 1 + \frac{s}{50}. \tag{3.69}$$

A second-order Butterworth low-pass filter is used for W_{ref}, whose purpose is to reject any input that is not induced by the driver. This could be a sudden pothole that makes the front wheel turn and, in return, generates a small yaw rate reference to track. This weighting tries to mitigate anything that a driver can't induce or cause:

$$W_{ref} = \frac{3948}{s^2 + 88.86 + 3948}. \tag{3.70}$$

The reference yaw rate is calculated using the LPV reference generator seen in Figure 3.26. The LPV reference generator provides the dynamics of the truck without rear axle steering active. The complete system can be viewed as a model reference robust controller where the control objective is to make the system reject disturbance and act like a truck without rear axle steering, so the driver should not feel any difference during changing circumstances, and the truck behaves the same.

Instead of synthesising a controller at every operating point of the LPV plant, similar to the LQR controller, a single operating point at 80 km/h was chosen and instead included modelled uncertainties in the system that covered the necessary region. Weight, inertia, and speed were modelled as uncertainties or uncertain parameters, with 50% allowed deviation and cornering stiffness was modelled with 20% uncertainty. This was done to cover an unloaded and loaded truck, as well as different tyres or worn tyres. Since lower speeds provide greater stability and less controllability, 80 km/h was chosen as the operating point, as it represents a typical truck speed where stability is most critical.

3.4.4 Adaptive MPC

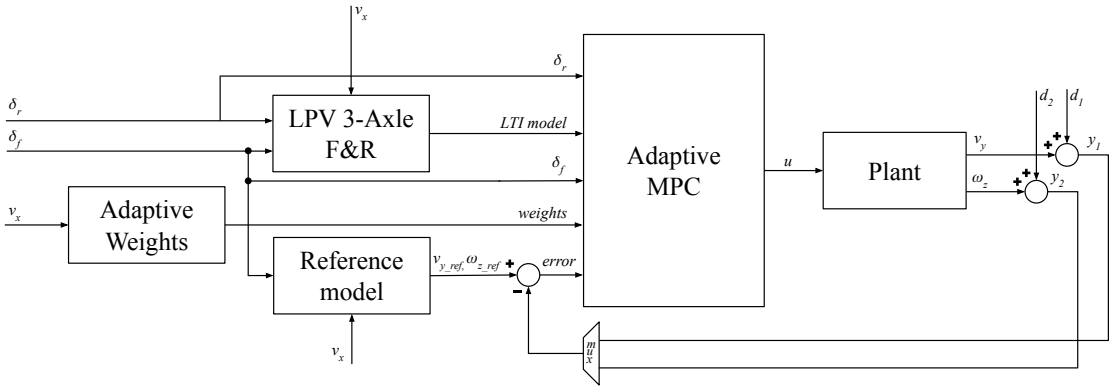

Figure 3.27: MPC block diagram.

To further improve control performance and handle constraints and disturbances, a Model Predictive Control (MPC) strategy was also developed. In contrast to LQR, MPC allows for the inclusion of system constraints and the prediction of future system behaviour over a finite horizon. MPC is based on predicting future system behaviour using a linear time-invariant (LTI) model. However, for systems with strong non-linearities or significant parameter variations, the accuracy of a fixed LTI model can degrade, which may result in reduced control performance. To address this, an adaptive MPC was implemented. In adaptive MPC, the internal prediction model is updated online to reflect the current operating conditions. This allows the controller to maintain prediction accuracy across a wide operating range. In this thesis, adaptation is introduced both through the velocity-dependent LPV prediction model and through scheduling of the MPC weighting matrices. Both the model parameters and the weighting matrices are scheduled as functions of the longitudinal velocity v_x. This results in a controller that adapts both its internal

model and its control objectives based on the current operating conditions.

The truck dynamics are represented using an LPV model where the system matrices depend on the longitudinal velocity v_x. The discrete-time state-space model used by the MPC controller is given by:

$$x(k+1) = A(v_x)x(k) + B(v_x)u(k) + B_d(v_x)d(k), \tag{3.71}$$

$$y(k) = C(v_x)x(k) \tag{3.72}$$

where:

- $x = [v_y \;\; r]^T$ is the state vector,
- u is the rear steering angle or manipulated variable (MV),
- d is the front steering angle provided by the driver, which is treated as a measured disturbance (MD).

The front steering input, δ_f, was explicitly modelled as a measured disturbance, allowing the controller to compensate for driver actions. This allows the MPC to anticipate and respond to the driver's steering inputs, improving disturbance rejection by including them in the prediction model.

The previous RWA, δ_r, is also provided to the controller as an external manipulated variable. This corresponds to the control input applied the previous control interval. Including this signal as an external manipulated variable improves state estimation, as the controller does not have full authority over the plant, the truck, since a driver is also present in the loop. The actual applied control input also differs from the controller output due to the application of saturation and rate limiting.

At each sampling instant, the MPC controller solves a finite-horizon optimisation problem:

$$J = \sum_{k=0}^{N_p} \left(y_k^T Q_y y_k \right) + \sum_{k=0}^{N_c} \left(u_k^T R_u u_k + \Delta u_k^T R_{\Delta u} \Delta u_k \right). \tag{3.73}$$

where:

- Q_y penalises output deviations,
- R_u penalises steering magnitude,
- $R_{\Delta u}$ penalizes steering rate.

The cost function used in MPC can be interpreted as a finite-horizon, discrete-time extension of the LQR objective, presented in equation 3.46. While both methods are based on quadratic optimisation, there are several important differences. In LQR, the objective is defined over an infinite time horizon and is formulated in terms of the system states. In contrast, MPC optimises over a finite prediction horizon and is typically formulated in terms of system outputs, which can be directly related to the physical performance variables such as yaw rate. Additionally, LQR does not explicitly account for system constraints, whereas MPC can incorporate constraints on inputs, input rates, and outputs directly in the optimisation problem. MPC also

explicitly predicts future system behaviour over the prediction horizon, while LQR provides a static state-feedback law without prediction. Despite these differences, both approaches share the same underlying principle of minimising a quadratic cost function to achieve optimal control performance.

The output and input weights were scheduled as functions of longitudinal velocity using predefined breakpoints and computed via piecewise interpolation. This enables the controller to adapt its behaviour depending on the operating conditions in the following way:

- Increased yaw rate weighting at higher speeds improves directional stability,
- Reduced weighting at lower speeds allows greater control flexibility,
- Input and input-rate weights regulate steering aggressiveness as a function of vehicle speed.

This scheduling strategy ensures that the controller behaviour remains consistent and well-tuned across the full operating range. The prediction and control horizons, as well as the adaptive weights, were determined through manual tuning to achieve the lowest possible yaw rate during split-μ braking.

The MPC controller requires estimates of the system states for prediction and control. For this thesis, state estimation is handled internally by MATLAB's MPC toolbox using a Kalman filter. Since the prediction model is updated at each sampling instant because of the LPV, the state estimator operates as a linear time-varying Kalman filter (LTVKF). The estimator gains are automatically updated based on the current system matrices, ensuring consistency between the prediction model and the estimated states. The estimated state vector is defined as:

$$x = \begin{bmatrix} v_y \\ r \end{bmatrix}, \tag{3.74}$$

where v_y is the lateral velocity and r is the yaw rate. The LTVKF is described by the recursive equations:

$$x_{k+1|k} = A_k x_{k|k} + B_k u_k, \tag{3.75}$$

$$x_{k|k} = x_{k|k-1} + L_k \left(y_k - C_k x_{k|k-1} \right), \tag{3.76}$$

where L_k denotes the time-varying Kalman gain, and the system matrices A_k and C_k depend on the longitudinal velocity v_x through the LPV model. The Kalman gain L_k is computed from the error covariance matrix:

$$L_k = P_{k|k-1} C_k^T \left(C_k P_{k|k-1} C_k^T + R \right)^{-1}. \tag{3.77}$$

And the covariance matrix is updated according to:

$$P_{k+1|k} = A_k P_{k|k-1} A_k^T + Q - A_k P_{k|k-1} C_k^T \left(C_k P_{k|k-1} C_k^T + R \right)^{-1} C_k P_{k|k-1} A_k^T. \tag{3.78}$$

where Q and R are the process and measurement noise covariance matrices. The Kalman filter is not explicitly designed or tuned, but is automatically generated and

updated by MATLAB's MPC toolbox.

3.4.5 Brake Pressure Feed Forward

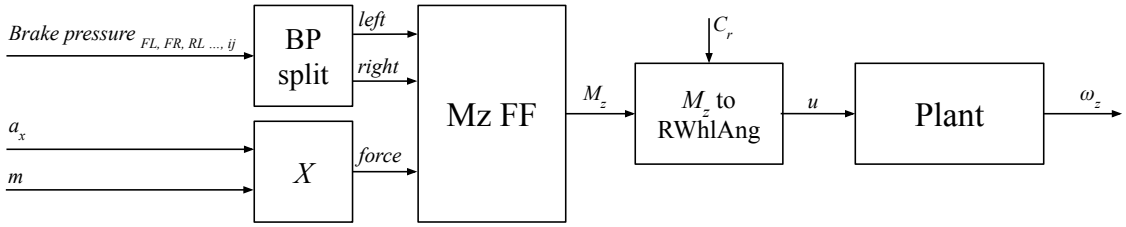

Figure 3.28: Brake pressure feed forward block diagram.

In contrast to the other controllers, the brake pressure controller is a feed-forward controller compared to the other controllers, which are feedback controllers. Unlike feedback controllers, which rely on yaw rate measurements and therefore react only after the vehicle dynamics have responded, the feed-forward controller, in a sense, bypasses the yaw dynamics and sensor delays. This results in a significantly faster response during braking, where early compensation of asymmetric brake forces is critical for yaw stability. Feedback control is still required, but the feed-forward approach improves performance by reacting to the disturbance at its source instantaneously.

The brake pressure feed-forward controller exploits the left–right brake force imbalance to generate a corrective yaw moment during split-μ braking scenarios. This corrective yaw moment is calculated by first summing the brake pressures after ABS modulation at each wheel in the following manner:

$$BP_{\text{tot}} = \sum_{i \in \mathcal{L}} BP_{\text{Lwhl},i} + \sum_{i \in \mathcal{R}} BP_{\text{Rwhl},i}, \tag{3.79}$$

where $\mathcal{L}$ and $\mathcal{R}$ denote the left and right side sets of tires. The total longitudinal force acting on the vehicle is calculated from the vehicle weight and the inertial measurement unit (IMU) longitudinal acceleration measurements.

$$F_x = ma_x, \tag{3.80}$$

$$K_{fx} = \frac{F_x}{BP_{\text{tot}}}. \tag{3.81}$$

K_{fx} thus represent a linear relation between brake-pressure and longitudinal force. Together with the brake pressure difference between the left and right sides, a yaw moment acting on the vehicle is calculated:

$$\Delta BP_{\text{whl}} = \sum_{i \in \mathcal{L}} BP_{\text{Lwhl},i} - \sum_{i \in \mathcal{R}} BP_{\text{Rwhl},i}, \tag{3.82}$$

$$M_z = \begin{cases} \frac{t_w}{2} K_{fx} \Delta BP_{\text{whl}}, & \text{if } F_x < 0 \ \wedge \ T_{\text{whl}} > T_{\text{min}} \\ 0, & \text{otherwise} \end{cases} \tag{3.83}$$

where T_{min} sets the limit for when the controller should activate, T_{whl} is the actual wheel torque and t_w is the track-width. During runtime, these equations are computed at every time step and provide a real-time estimation of yaw moment caused by differential braking. The yaw moment is then converted to a rear wheel angle or control input u in the following way:

$$u = \frac{M_z}{l_r * C_{r(vx)}} * K_p, \tag{3.84}$$

where K_p is a tunable gain added to tune the aggressiveness of the controller.

3.5 Stability

In the context of vehicle dynamics, stability refers to the ability of the vehicle to maintain a controlled and predictable motion following a disturbance, such as asymmetric braking or steering inputs. For this thesis, stability is evaluated primarily in terms of lateral behaviour, with a particular focus on the yaw rate. A stable response is characterised by bounded system states and convergence towards a desired operating condition without excessive oscillations or divergence. In contrast, unstable behaviour is indicated by growing oscillations, sustained deviations, or loss of control. There are many ways to interpret stability, depending on the goal, conditions, and architecture of the control system. In this case, where a driver is part of the loop together with an LPV system, classical stability analysis methods such as gain and phase margins may not be applicable or as valuable as other methods, and can come with inherent limitations.

The nominal plant, together with the controller, can be assessed by analysing the system's closed-loop poles. For continuous time systems, if the poles lie in the left half-plane of the complex plane, the system is internally stable. However, a limitation of this approach is that it provides no information about how driver input is propagated, nor any robustness guarantees. The most important and safety-critical requirement for this thesis control objective is whether the controller introduces unintended yaw rate beyond the driver's intention. In order to evaluate this, another approach to analyse stability is needed.

Gain and phase margins can be used to evaluate robustness to model uncertainties or unmodelled dynamics. It provides how much margin the feedback loop has to gain or phase perturbations. While this is valuable information, it does not apply to reference-driven behaviour, and it can not distinguish between whether the driver causes the system to go into an unstable state or if it is the controller that amplifies the yaw rate without the driver's intention.

The $||T_{ref\_yawrate \to yawrate}||_\infty$ induced norm of the closed loop from reference yaw rate to actual yaw rate is another way to assess stability. This measures the worst-case amplification of the driver-induced bounded reference yaw rate to the actual vehicle yaw rate. This indicates that the above-mentioned stability assessments cannot be the driver's intention. If the norm is finite, then any bounded input from the driver yields a bounded yaw rate response. This can be applied to all controllers and checked for each operating point given by the LPV. This is important because the system can enter an unstable state. If the driver wants to spin the truck, it can and should be allowed to; what is safety-critical is that the controller does not introduce anything unintended by the driver. The stability metric in terms of the infinity norm can be interpreted in the following way:

$\|\|T\|\|_\infty < 1$	Stable.
$\|\|T\|\|_\infty > 1$	Could be stable, depends on magnitude.
$\|\|T\|\|_\infty \approx 1$	Stable.
$\|\|T\|\|_\infty \to \infty$	Unstable, induced instability by the controller.

Table 3.4: Infinity norm evaluation.

For the MPC, driver-intent stability was evaluated using the linear equivalent MPC representation, with the reference yaw rate directly applied to the controller reference input. This avoids artificial error injection and ensures a physically consistent input–output stability assessment.

4

Results and Discussion

This chapter presents the results obtained from the simulations conducted in this thesis. The developed controllers are evaluated and compared using a set of predefined scenarios designed to assess vehicle stability under split-μ braking. The performance of each controller is analysed based on the metrics formulated in Chapter 3 . The stability characteristics of the controllers are also examined to provide insight into their behaviour. The results are presented through plots and quantitative comparisons, and are discussed within this chapter, forming the basis for the conclusions presented in the following chapter.

4.1 Performance Plots

This section presents the results for the controllers across the test scenarios defined in Chapter 3. The test scenarios consist of split-μ braking on different test tracks. The different test tracks are: a straight road, a curved road with low friction on the inside, a curved road with low friction on the outside, and a lane-change manoeuvre. Additionally, tests using a more aggressive ABS were conducted to assess its effect on braking distance, as well as tests using both feedback and feed-forward control and one including simulated sensor noise and actuator delay. All tests were performed using the friction coefficients and speeds also described in Chapter 3; this chapter only includes the results from using a split-μ surface of (0.2,1) and a speed of 80 km/h. This is because during these conditions, the most interesting and telling results were captured. Using the other friction coefficients and speeds, the results showed similar behaviour as the ones presented below, but not as extreme in terms of magnitude. These were therefore omitted from the report.

The dashed line in the plots shows where the braking and split-μ surface begins. It is also important to note that the first 2–4 seconds of the manoeuvre are the most critical and are therefore the primary focus of the analysis. For the steering wheel angle, the interval between the dashed lines is the region of interest. Beyond this interval, the vehicle speed is very low, and the driver, or driver model, primarily aims to realign the vehicle after braking. To achieve this at low speed, the driver applies disproportionately large steering wheel angles, which are not representative of the split-μ braking scenario. All controllers are compared against each other as well as to the baseline and are used together with the ABS system presented in section 3.2.1. As mentioned in 2.2.5, the baseline consists of no RAS control but with ABS active.

4.1.1 Straight Road

Below are the results from split-μ braking at a $\mu = (0.2, 1)$ surface at 80 km/h on a straight road. This test corresponds to the road presented in Figure 3.3.

Disclaimer: There is no need to analyse Figure 4.1, 4.2, 4.3, 4.4, 4.5, and 4.6 beyond 9 seconds because the truck is almost at a standstill and the driver model produces extreme steering wheel angles to compensate for the lateral deviation caused by the scenario and steer the truck back to its intended route on the road.

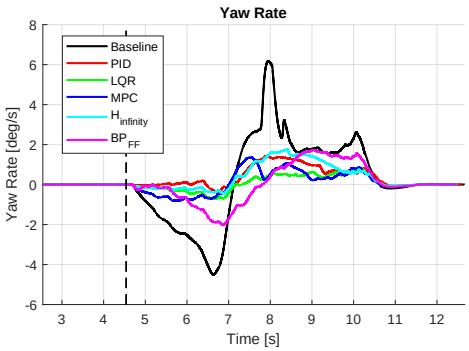

Figure 4.1: Yaw rate on straight road (0.2,1) 80 km/h.

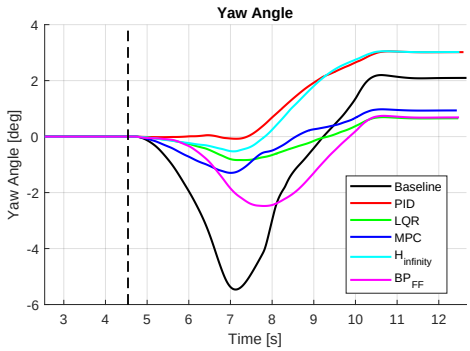

Figure 4.2: Yaw angle on straight road (0.2,1) 80 km/h.

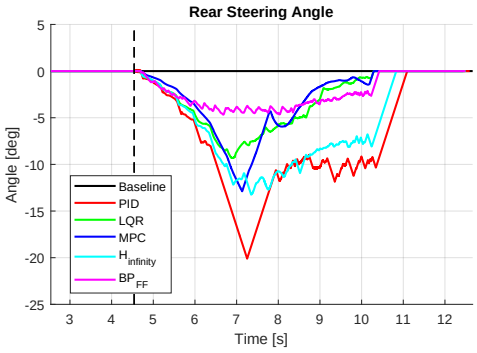

Figure 4.3: Rear steering angle on straight road (0.2,1) 80 km/h.

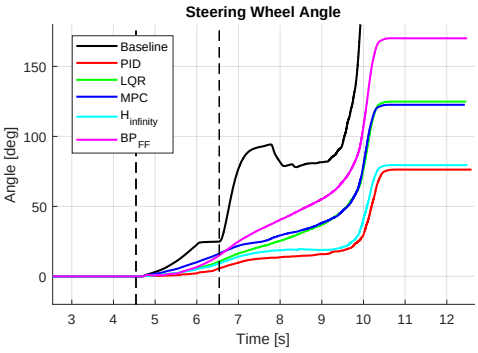

Figure 4.4: Steering wheel angle on straight road (0.2,1) 80 km/h.

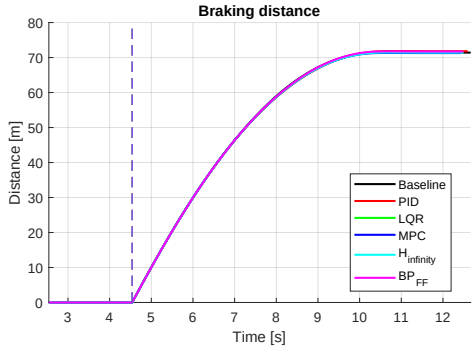

Figure 4.5: Braking distance on straight road (0.2,1) 80 km/h.

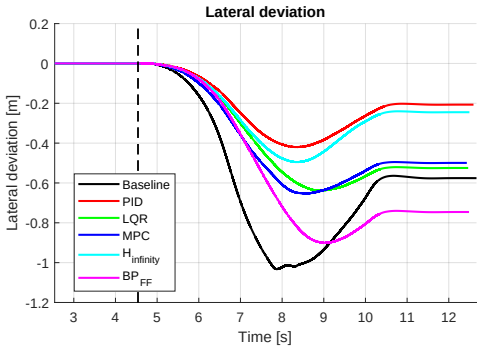

Figure 4.6: Lateral deviation on straight road (0.2,1) 80 km/h.

The braking distances and RMS for the different controllers can be viewed more clearly in the table below.

Setup	Brake Distance [m]
Baseline	71.461
PID	71.855
LQR	71.593
MPC	71.324
$\mathcal{H}_\infty$	71.460
BP_{FF}	71.764

Table 4.1: Brake distance for each setup on straight road (0.2,1) at 80 km/h.

Setup	RMS Yaw Rate Error [deg/s]
Baseline	2.3294
PID	0.7403
LQR	0.4426
MPC	0.6314
$\mathcal{H}_\infty$	0.9112
BP_{FF}	1.1170

Table 4.2: RMS of tracking error between reference and actual yaw rate on straight road (0.2,1) at 80 km/h.

Setup	RMS yaw rate tracking error [deg/s]	RMS yaw rate [deg/s]
Baseline	4.6689	2.3294
PID	0.6339	0.7403
LQR	1.6101	0.4426
MPC	1.7464	0.6313
$\mathcal{H}_\infty$	0.8250	0.9112
BP_{FF}	not applicable	1.1170

Table 4.3: RMS of tracking error between reference and actual yaw rate and RMS of yaw rate on straight road (0.2,1) 80 km/h.

Overall, significant improvements are observed in nearly all performance metrics except braking distance. None of the controllers achieves a substantial reduction in braking distance compared to the baseline, and some increase it. On the other hand, what is important is that even though no large reduction in braking distance can be observed, the controllers do not increase it by a significant amount. This is an important result, as increased vehicle stability achieved through RAS control could otherwise come at the cost of longer braking distances. This outcome would

introduce a trade-off between stability and stopping performance, requiring prioritisation between the two.

The PID controller could be regarded as the best overall performer in this specific test in terms of stability and driver effort reduction, but at the expense of larger control inputs. The LQR controller achieves performance comparable to the PID controller, while yielding a slightly shorter braking distance and using a smaller RWA during the manoeuvre. Less aggressive control inputs but still yielding comparable performance could motivate that the LQR is better than the PID in this sense. The worst performance is observed for the feed-forward controller. The feed-forward controller, however, is primarily intended to be used in combination with a feedback controller. The feed-forward controller reacts very quickly and is effective at cancelling the initial yaw rate. Therefore, combining the feed-forward controller with a high-performing feedback controller could yield additional performance improvements. These improvements are expected to be most noticeable at the start of the manoeuvre, as the feed-forward controller's effectiveness diminishes over time. Since the PID controller already achieves a low yaw rate early in the manoeuvrer, the added benefit of feed forward may be marginal. However, the feed-forward controller could provide more noticeable improvements when combined with the LQR controller by assisting it in reducing the initial yaw rate toward zero and decreasing the driver effort even further.

For all controllers, both yaw rate and yaw angle are reduced, which is beneficial from a vehicle stability perspective. However, a more notable result is the influence of the controllers on driver effort, represented by the steering wheel angle. Compared to the baseline case, all controllers significantly reduce the required steering effort. Within the region of interest, indicated by the dashed lines, the PID controller is the most effective at maintaining a low steering wheel angle, with a peak value of approximately 10 degrees. This result shows a large reduction in driver workload during the most critical phase of the manoeuvre. A noticeable improvement in lateral deviation can also be observed, with the maximum deviation reduced from approximately 1.0 meters in the baseline case to about 0.4 meters for the best-performing controller, which is the PID controller in this metric.

The MPC does not decrease the initial yaw rate as well, and overall shows a slower, more sluggish performance. This is expected, and MPC might not be optimal for this use case since it is better at predicting further into the future, using a larger prediction horizon. It slightly decreases braking distance but does not reduce the steering wheel angle as much as LQR and PID. The RWA the MPC uses is similar to the PID in shape but not in magnitude. The PID is faster and more aggressive, while the MPC is slower and not as aggressive; this could explain the difference in the performance in the other metrics. This could also be a tuning issue, but generally, MPC might not be suited for these types of short-duration, high-dynamics manoeuvres unless a longer prediction horizon or more aggressive state weighting is used.

The $\mathcal{H}_\infty$ controller performs somewhere in-between the PID and the LQR. It uses a higher and more aggressive RWA, which may be why it performs better than the LQR. The $\mathcal{H}_\infty$ controller is not the best in any metric, but performs well in all of them. This is the strength of robust control, and the results show that $\mathcal{H}_\infty$ is not optimal but prioritises robustness, leading to overall reliable performance.

4.1.2 Curved Road

Below are the results from split-μ braking on the two different curved roads, both with a radius of 200 meters, are presented. The difference between these tests is where the low friction is located.

4.1.2.1 Low Friction on Inside

Below are the results from split-μ braking at a $\mu = (0.2, 1)$ surface at 80 km/h on a curved road with low friction on the inside of the curve. This test corresponds to the road presented in Figure 3.6.

Disclaimer: There is no need to analyse Figure 4.7, 4.8, 4.9, 4.10, 4.11, and 4.12 beyond 19 seconds because the truck is almost at standstill and the driver model produces extreme steering wheel angles to compensate for the lateral deviation caused by the scenario and steer the truck back to it's intended route on the road.

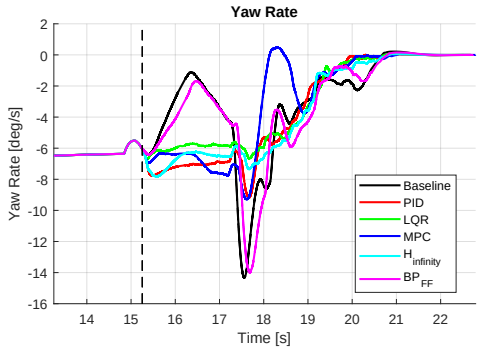

Figure 4.7: Yaw rate on curved road with low friction on inside (0.2,1) 80 km/h.

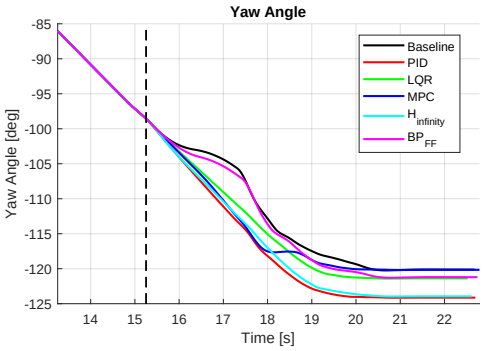

Figure 4.8: Yaw angle on curved road with low friction on inside (0.2,1) 80 km/h.

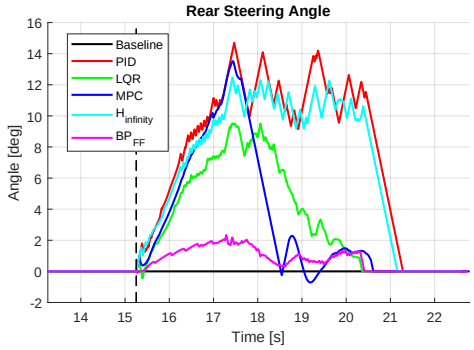

Figure 4.9: Rear steering angle on curved road with low friction on inside (0.2,1) 80 km/h.

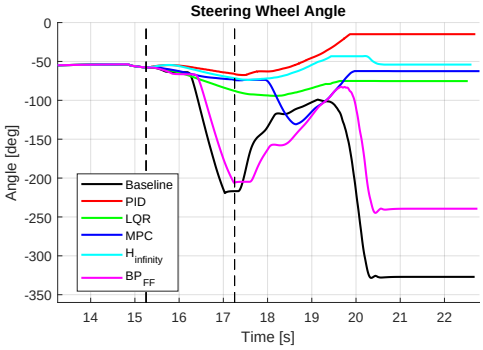

Figure 4.10: Steering wheel angle on curved road with low friction on inside (0.2,1) 80 km/h.

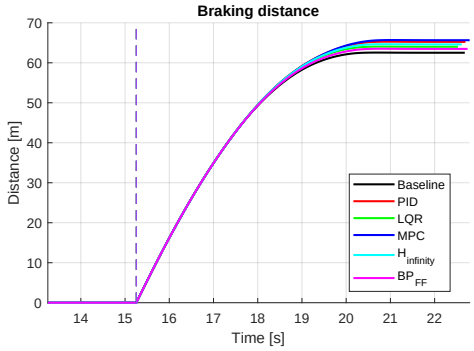

Figure 4.11: Braking distance on curved road with low friction on inside (0.2,1) 80 km/h.

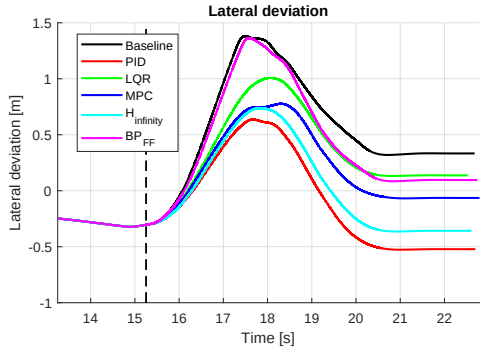

Figure 4.12: Lateral deviation on curved road with low friction on inside (0.2,1) 80 km/h.

The braking distances and RMS for the different controllers can be viewed more clearly in the table below.

Setup	Brake Distance [m]
Baseline	62.554
PID	65.272
LQR	64.049
MPC	65.689
$\mathcal{H}_\infty$	64.645
BP_{FF}	63.515

Table 4.4: Brake distance for each setup on curved road with low friction on inside (0.2,1) 80 km/h.

Setup	RMS Yaw Rate Error [deg/s]
Baseline	8.5869
PID	0.5343
LQR	2.4513
MPC	3.3640
$\mathcal{H}_\infty$	0.6921
BP_{FF}	8.5826

Table 4.5: RMS of tracking error between reference and actual yaw rate on curved road with low friction on inside (0.2,1) 80 km/h.

The results demonstrate substantial improvements in lateral stability during braking in a corner with low friction on the inside wheels. In contrast to the straight-line scenario, this case is characterised by a transition toward understeer, requiring the controller to actively restore yaw response rather than counteract oversteer.

Figure 4.7 shows significant improvements in yaw rate, especially using LQR, $\mathcal{H}_\infty$, and PID, where the understeer is effectively removed. The baseline and BP_{FF} understeer, which can be seen in Figure 4.7, since the absolute yaw rate decreases after braking is initiated. This indicates that the truck loses cornering capability and understeers away from the corner. Using LQR, $\mathcal{H}_\infty$, and PID the yaw rate before braking is sustained and smoothly transitions to zero once the truck slows down to a stop. The yaw angle in Figure 4.8 also confirms the behaviour where an ideal curve is linear initially with a smooth transition to a horizontal line. LQR, $\mathcal{H}_\infty$, and PID achieve this while the baseline and BP_{FF} fail to follow the corner and stop turning. The sudden increase in yaw rate and yaw angle at 17.5 seconds for the baseline and BP_{FF} is not to be considered a valid run since the truck leaves the low-friction surface and instantly regains more control authority because of the sudden increase in traction caused by the high-friction side. This is to be considered a failure since the truck did not manage to stay in its own lane. Understeering out into the other lane imposes great danger and increases the risk of a frontal crash with vehicles coming from the opposite direction. That is why we also see a reduction in braking distance for baseline and BP_{FF} seen in Table 4.4. The truck simply leaves the lane and stops faster due to the higher friction surface.

Figure 4.10 shows the steering wheel angles the driver induced during this scenario. It is evident that using RAS controls the effort required of the driver in order to stabilise the vehicle. The baseline and BP_{FF} required the driver to induce 150 degrees of steering correction relative to the initial angle during the first 2 seconds of the braking scenario, while LQR, $\mathcal{H}_\infty$, and PID almost completely remove the driver's need to intervene at all. This allows the driver to maintain an almost constant steering wheel angle throughout the entire scenario. This indicates that the truck maintains its intended trajectory with minimal driver input. The control objective is to realise the driver's intended trajectory, which is successfully achieved in this scenario.

The MPC controller, on the other hand, behaves somewhat differently. At approximately 18 seconds, it produces a large change in yaw rate and steering wheel angle. This behaviour is likely caused by the adaptive weighting, where penalties on state deviations and control effort change at lower vehicle speeds. Nevertheless, this does not introduce stability issues, as the vehicle speed at this point, near the end of the manoeuvre, is relatively low. All controllers use relatively large and aggressive control inputs, as indicated by the frequent activation of rate limiting. As in 4.1.1, among the feedback controllers, the LQR controller is the most conservative in terms of control effort, with the exception of the feed-forward controller. Despite using a smaller RWA, the LQR controller still remains competitive with the other controllers in terms of performance.

4.1.2.2 Low Friction on Outside

Below, the results from split-μ braking at a $\mu = (0.2, 1)$ surface at 80 km/h on a curved road with low friction on the outside of the curve are presented. This test

corresponds to the road presented in Figure 3.5.

Disclaimer: There is no need to analyse Figure 4.13, 4.14, 4.9, 4.16, 4.17, and 4.18 beyond 20 seconds because the truck is almost at a standstill, and the driver model produces extreme steering wheel angles to compensate for the lateral deviation caused by the scenario and steer the truck back to its intended route on the road.

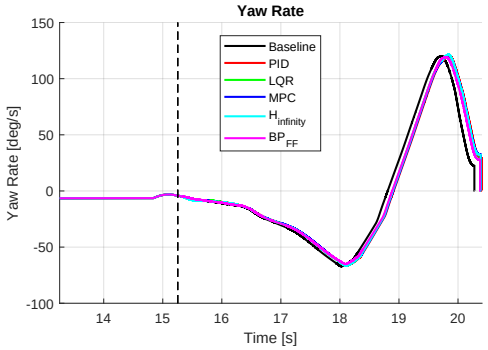

Figure 4.13: Yaw rate on curved road with low friction on outside (0.2,1) 80 km/h.

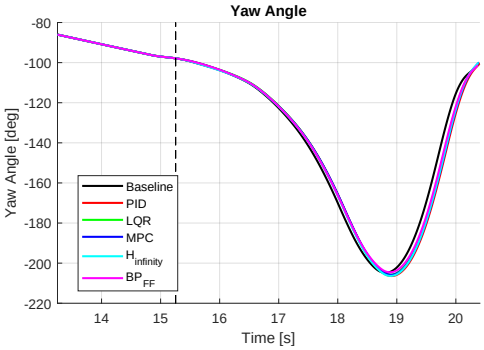

Figure 4.14: Yaw angle on curved road with low friction on outside (0.2,1) 80 km/h.

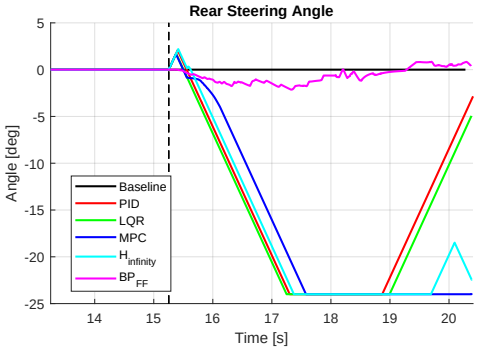

Figure 4.15: Rear steering angle on curved road with low friction on outside (0.2,1) 80 km/h.

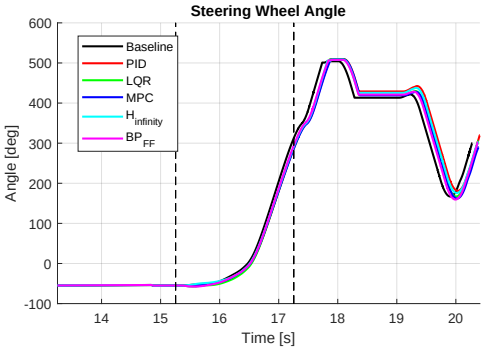

Figure 4.16: Steering wheel angle on curved road with low friction on outside (0.2,1) 80 km/h.

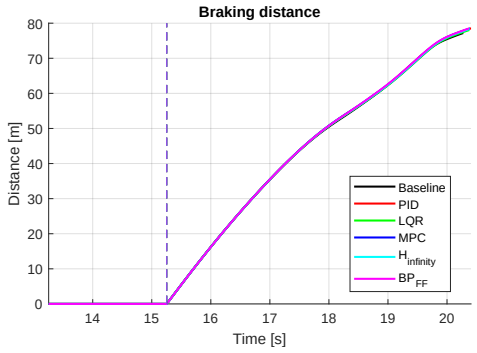

Figure 4.17: Braking distance on curved road with low friction on outside (0.2,1) 80 km/h.

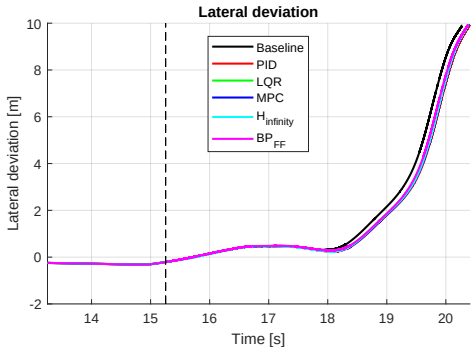

Figure 4.18: Lateral deviation on curved road with low friction on outside (0.2,1) 80 km/h.

The braking distances and RMS for the different controllers can be viewed more clearly in the table below.

Setup	Brake Distance [m]
Baseline	77.204
PID	78.513
LQR	78.048
MPC	78.468
$\mathcal{H}_\infty$	78.336
BP_{FF}	78.495

Table 4.6: Brake distance for each setup on curved road with low friction on outside (0.2,1) 80 km/h.

Setup	RMS Yaw Rate Error [deg/s]
Baseline	58.2115
PID	59.7709
LQR	58.8592
MPC	59.6546
$\mathcal{H}_\infty$	57.9914
BP_{FF}	61.4057

Table 4.7: RMS of tracking error between reference and actual yaw rate on curved road with low friction on outside (0.2,1) 80 km/h.

Figures 4.13, 4.14, 4.15, 4.16, 4.17, and 4.18 show the baseline and the different controllers during split-μ braking in a corner with low friction on the outside. The results show that at this speed and with a corner radius of 200 meters, the truck enters an unstable state. More precisely, the truck oversteers and spins out. This can be seen by the extreme yaw rate and yaw angles in the plot. No controller manages to stabilise the truck in this scenario, nor does it significantly worsen the stability. The control output, and consequently the rear wheel angle, saturates for the feedback control strategies, indicating that the controller demands more from the rear axle steering system than it can deliver.

The reason this combination of speed, corner radius, and friction coefficient is likely unrealistic for the rear axle steering system to stabilise is the load transfer that occurs during braking and cornering. Figure 4.19 shows the vertical load on each tyre during the scenario, but with a constant friction coefficient of 1 on all tyres. The figure reveals that the load transfers to the left side and towards the front when turning right and braking, which is expected. This reduces the controllability of the RAS system since it is provided with less load that the tyre can use to generate forces. Low friction is especially bad on the outside tyre in a corner since it has

the most load on it, which could be used to create forces. But the outside tyre is instead limited by the low friction, reducing the tyre's effectiveness to recover the vehicle in this scenario. This scenario requires stabilising the truck using the axle

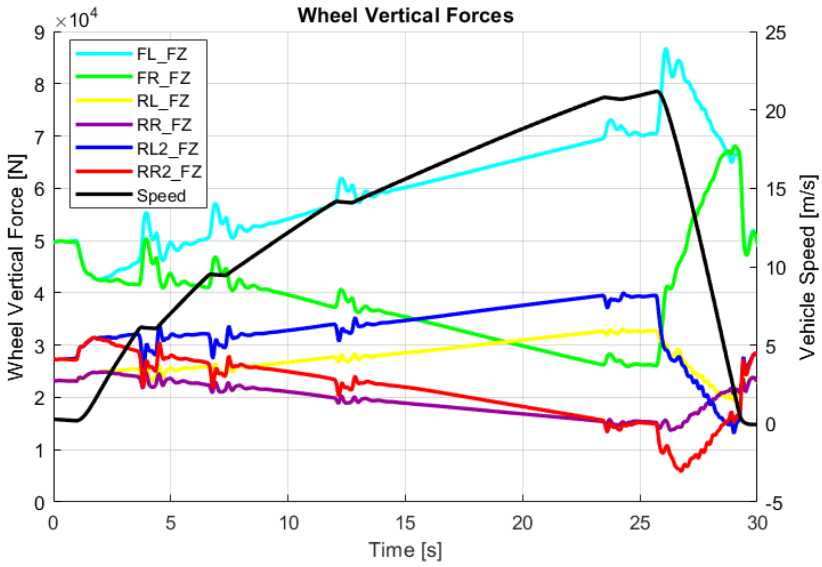

Figure 4.19: Wheel vertical forces during braking in a corner.

with the lowest vertical load, while the most heavily loaded tyre on the rear axle operates on a low-friction surface. This is not feasible given the scenario's speed, corner radius and friction coefficient. The jitter in Figure 4.19 is the driver model manually shifting the truck to reach the requested speed. The maximum steering angle of the rear axle steering system is not the limiting factor in this scenario, as counter-steering the rear tyres to their limits, combined with the yaw dynamics, results in a tyre side-slip so aggressive that the tyres lose their ability to generate forces in the intended direction and instead just provide drag. To extreme rear wheel steering angles relative to the body side-slip results in tyre slip-angles beyond the optimal range, and the tyre is heavily saturated, which decreases performance.

Figures 4.20, 4.21, 4.22, and 4.23 show the yaw rate and steering wheel angle during the same scenario but using the split-μ pairs (0.5,1) and (0.8,1) correspondingly. Under these friction conditions, the truck does not spin out and completes the scenario. The truck passes all curved roads with low friction on the outside tests except at 80 km/h with $\mu = (0.2, 1)$. Analysing the performance reveals a slight improvement when rear axle steering is used, but its effectiveness decreases when the difference in friction goes to zero. At 80 km/h and with $\mu = (0.5, 1)$, a noticeable improvement can be observed. The LQR controller shows the best performance, exhibiting a smoother yaw rate transition from cornering to standstill. Additionally, the peak steering wheel effort is reduced by approximately 40 degrees within 2 seconds after the brake begins. All tests at speeds 50 and 30 km/h with $\mu = (0.2, 0.5, 0.8)$ were stable and showed decreasing difference between the baseline and the RAS controllers when speed decreases and the low friction coefficient increases.

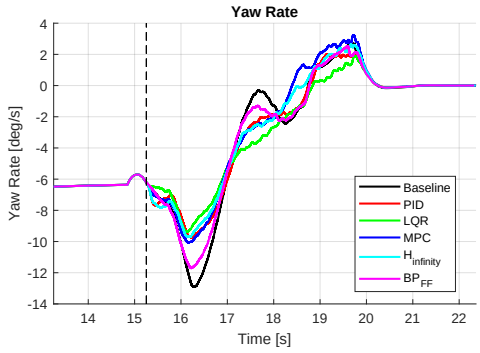

Figure 4.20: Yaw rate on curved road with low friction on outside (0.5,1) 80 km/h.

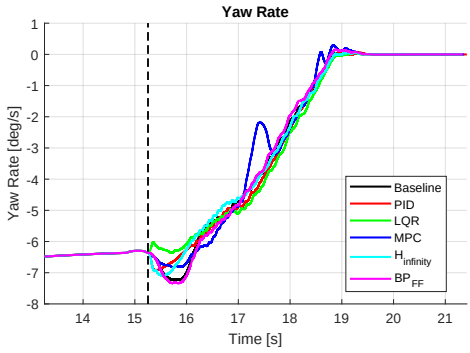

Figure 4.21: Yaw rate on curved road with low friction on outside (0.8,1) 80 km/h.

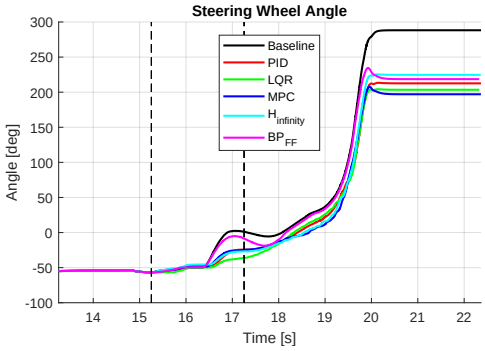

Figure 4.22: Steering wheel angle on curved road with low friction on outside (0.5,1) 80 km/h.

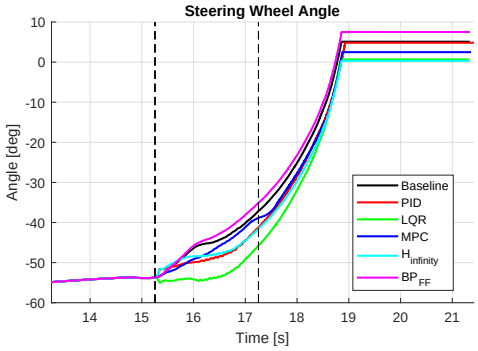

Figure 4.23: Steering wheel angle on curved road with low friction on outside (0.8,1) 80 km/h.

4.1.3 ISO Lane Change

Below, the results from split-μ braking at a $\mu = (0.2, 1)$ surface at 80 km/h during a lane change are presented. This test corresponds to the road presented in Figure 3.4.

Disclaimer: There is no need to analyse Figure 4.24, 4.25, 4.26, 4.27, 4.28, and 4.29 beyond 6 seconds because the truck is almost at standstill and the driver model produces extreme steering wheel angles to compensate for the lateral deviation caused by the scenario and steer the truck back to it's intended route on the road.

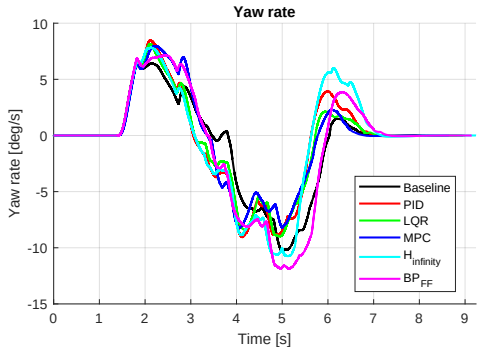

Figure 4.24: Yaw rate during lane change (0.2,1) 80 km/h.

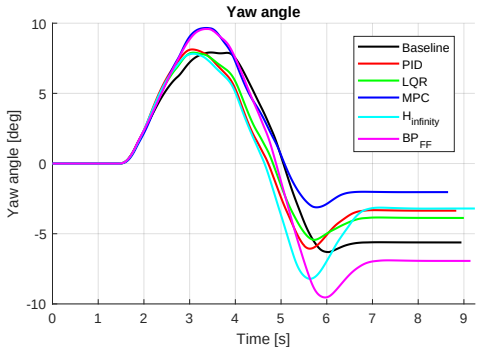

Figure 4.25: Yaw angle during lane change (0.2,1) 80 km/h.

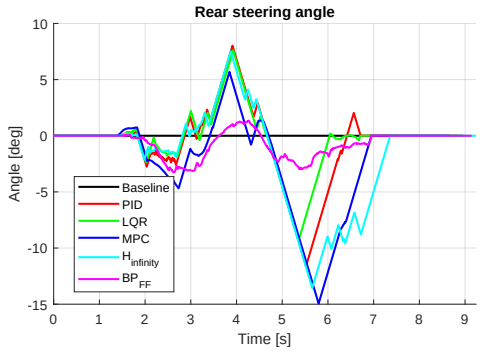

Figure 4.26: Rear steering angle during lane change (0.2,1) 80 km/h.

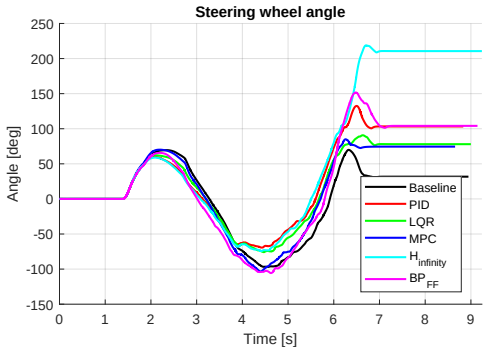

Figure 4.27: Steering wheel angle during lane change (0.2,1) 80 km/h.

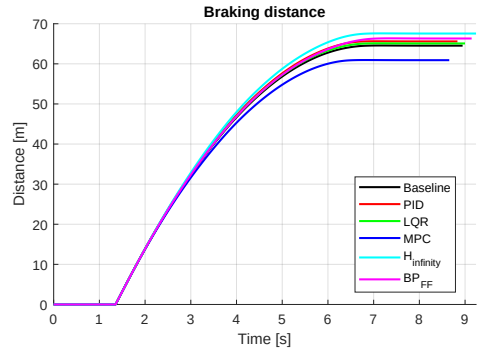

Figure 4.28: Braking distance during lane change (0.2,1) 80 km/h.

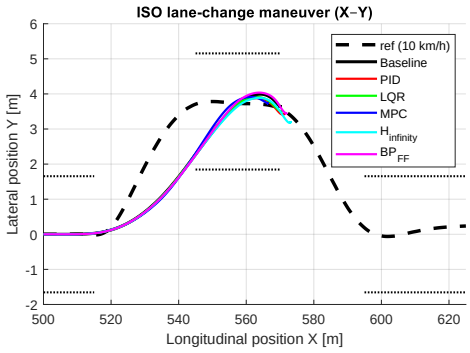

Figure 4.29: Lateral deviation during lane change (0.2,1) 80 km/h.

The braking distances and RMS for the different controllers can be viewed more clearly in the table below.

Setup	Brake Distance [m]
Baseline	64.575
PID	65.630
LQR	65.126
MPC	60.951
$\mathcal{H}_\infty$	67.596
BP_{FF}	66.360

Table 4.8: Brake distance for each setup during lane change (0.2,1) 80 km/h.

Setup	RMS Yaw Rate Error [deg/s]
Baseline	3.0393
PID	1.2201
LQR	1.3951
MPC	1.5319
$\mathcal{H}_\infty$	1.3032
BP_{FF}	2.4926

Table 4.9: RMS of tracking error between reference and actual yaw rate during lane change (0.2,1) 80 km/h.

The ISO lane change scenario differs from the other scenarios because the driver must also steer into the other lane while braking. This simulates a scenario where the driver realises that just braking won't be enough to prevent a collision, and the object in front has to be avoided by also steering away from it. In addition, winter road conditions are present, with a dry, high-friction surface resulting from many vehicles clearing a path in the lane, and snow and ice are present in the middle of the road and on the sides, where vehicles don't usually travel when going down the lane. This is seen in Figure 3.4, where the truck travels only on high-friction surfaces when positioned in the middle of the lane. A lane change during this scenario leads to the truck entering and exiting the low-friction surface multiple times, which could lead to unpredictable and unstable behaviour.

Analysing the Figures 4.24, 4.25, 4.26, 4.27, 4.28, and 4.29 reveal that the controllers barely provide any benefit during this scenario. LQR and PID slightly reduce driver effort, with the driver having to steer approximately 30-40 deg less than the baseline at 4.5 seconds. However, no controller manages to make the truck pass the ISO lane change. Figure 4.29 shows the trajectory of the truck and the ISO lane change boundaries. It might seem that the truck passes the test, but the position of the truck is recorded in the middle, so half the width has to be applied, and then the truck runs over the cones it should avoid, which was visually confirmed in the simulations.

The reason the controllers do not show any significant improvement during this scenario could be attributed to the controller design and driver model. For the controllers, the objective is to enhance the driver's will to follow the driver's intended reference yaw rate. The controllers do not act predictively based on upcoming road conditions, as this would require additional sensing technologies such as cameras, LiDAR, or similar systems. Using or applying those types of systems would contradict the objective of achieving a straightforward implementation using only signals and information already available in existing systems. The rear steering angle induced by the controller corresponds to the control action required to follow the expected vehicle behaviour defined by the LPV reference generator. Since the yaw rate input from the driver is already close to this expected behaviour, the required control intervention is limited and not so significant.

Together with the low-friction surface during cornering, this leads to minimal intervention of the truck's yaw rate. The limited need for controller intervention in this scenario can be attributed to the driver model. The driver model setup in TruckMaker is configured to represent a generic truck driver with respect to steering wheel velocity, reaction time and longitudinal and lateral acceleration limits. As a result, the driver balances the driver model parameters while completing the lane change and therefore does not induce an aggressive enough manoeuvre to destabilise the truck or create conditions where the controllers can significantly contribute. The controllers have the potential to improve performance; this requires the driver to initiate a manoeuvre that would normally be impossible without rear-axle steering, and subsequently trust the RAS controller to fulfil the intended vehicle response. In this scenario, the low-friction surface does not introduce a sufficiently large deviation from the vehicle's expected high-friction behaviour, thereby limiting the potential impact of the controller. The controllers are better at rejecting disturbances, in other words, deviation in yaw rate when the driver's intent is steady, for example, in a straight driving scenario or a constant turning scenario.

Table 4.8 show the braking distances during braking under the ISO lane change manoeuvrer. The different controllers and the baseline are quite similar, with the exception of the MPC, which reduces the braking distance slightly. However, this should not be interpreted as an improvement in braking distance attributable to the MPC. Rather, it is a consequence of the control actions taken by the MPC, which coincidentally caused the truck to remain on the high-friction surface slightly longer during the test compared to the other controllers. Additionally, it should be emphasised that none of the controllers incorporates any mechanism to identify or actively seek regions of higher friction to decrease the braking distance or gain stability.

4.1.4 Straight Road with Delay and Noise

Below are the results from split-μ braking at a $\mu = (0.2, 1)$ surface at 80 km/h on a straight with delay and noise. The results show the impact of delay and measurement noise in the system. The delay represents the time between the requested angle, calculated by the controllers, until it reaches the rear axle steering system. The noise represents sensor noise, influencing yaw rate measurements and consists of band-limited white noise. The noise is artificially added in simulations. The noise can be viewed and analysed in Figure 4.30. This test corresponds to the road presented in Figure 3.3.

Disclaimer: There is no need to analyse Figure 4.31, 4.32, 4.33, 4.34, 4.35, and 4.36 beyond 9 seconds because the truck is almost at a standstill and the driver model produces extreme steering wheel angles to compensate for the lateral deviation caused by the scenario and steer the truck back to its intended route on the road.

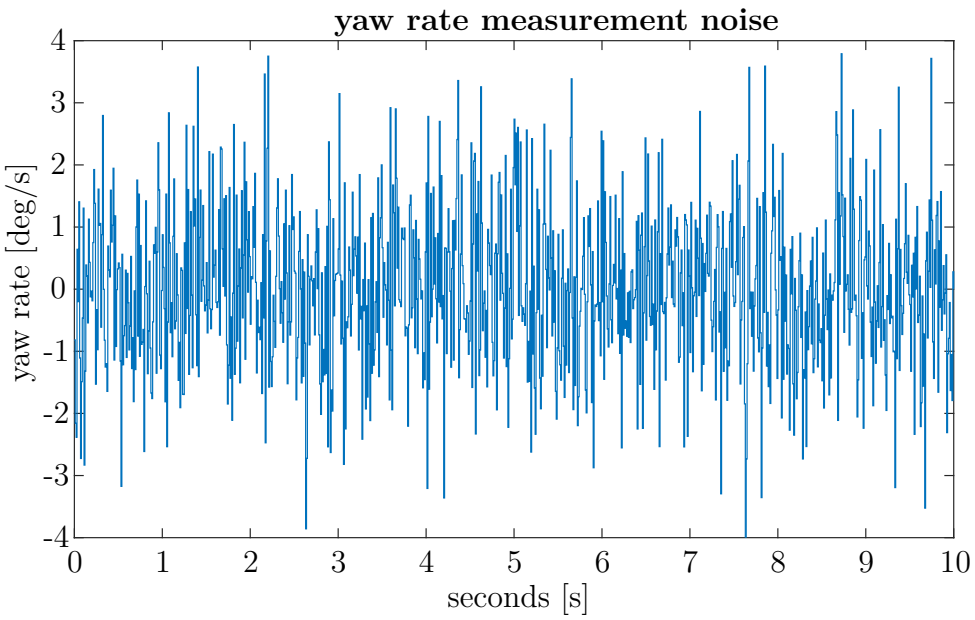

Figure 4.30: Simulated sensor noise.

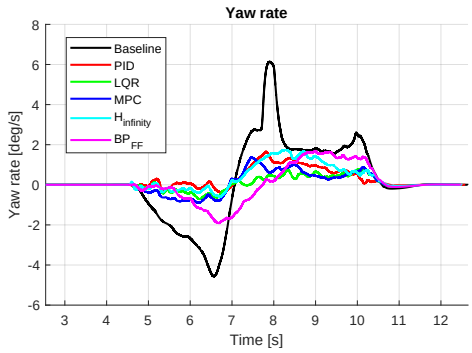

Figure 4.31: Yaw rate on straight road with delay and noise (0.2,1) 80 km/h.

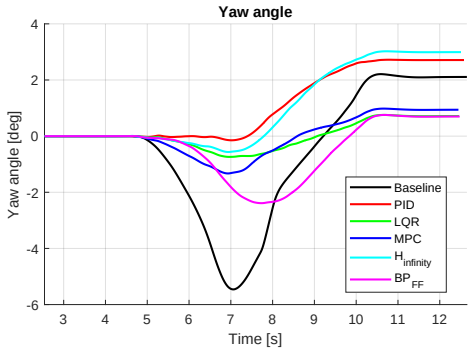

Figure 4.32: Yaw angle on straight road with delay and noise (0.2,1) 80 km/h.

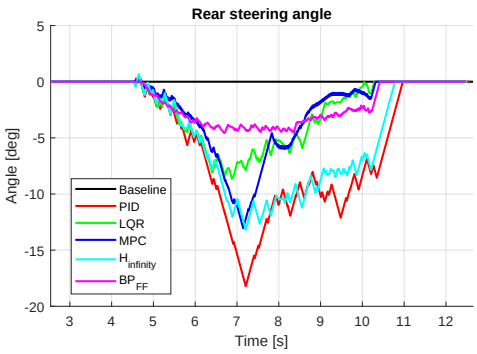

Figure 4.33: Rear steering angle on straight road with delay and noise (0.2,1) 80 km/h.

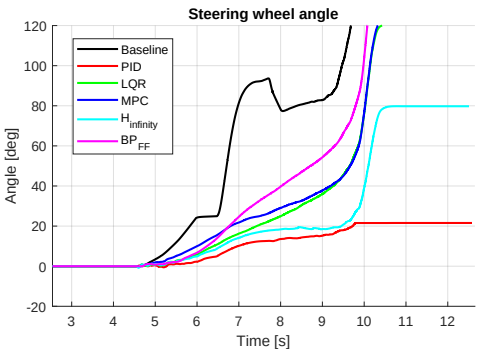

Figure 4.34: Steering wheel angle on straight road with delay and noise (0.2,1) 80 km/h.

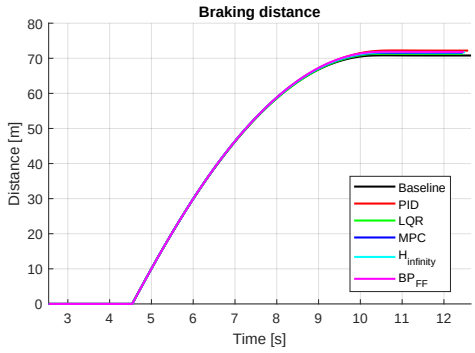

Figure 4.35: Braking distance on straight road with delay and noise (0.2,1) 80 km/h.

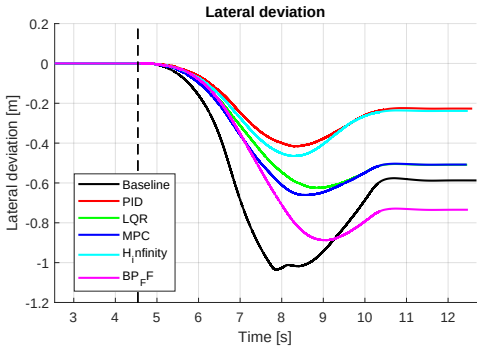

Figure 4.36: Lateral deviation on straight road with delay and noise (0.2,1) 80 km/h.

The braking distances and the RMS for the different controllers can be viewed more clearly in the table below.

Setup	Brake Distance [m]
Baseline	70.830
PID	72.229
LQR	71.310
MPC	71.535
$\mathcal{H}_\infty$	71.724
BP_{FF}	71.862

Table 4.10: Brake distance for each setup on straight road with delay and noise (0.2,1) 80 km/h.

Setup	RMS yaw rate tracking error [deg/s]	RMS yaw rate [deg/s]
Baseline	4.5550	2.3096
PID	1.3418	1.3451
LQR	1.5399	0.5022
MPC	1.7927	0.6168
$\mathcal{H}_\infty$	0.8756	1.0758
BP_{FF}	not applicable	0.9998

Table 4.11: RMS of tracking error between reference and actual yaw rate and RMS of yaw rate on straight road with delay and noise (0.2,1) 80 km/h.

The effects of noise and delay are most notable in the RWA signal and the yaw rate response. The RWA is visibly noisier, particularly for the PID and $\mathcal{H}_\infty$ controllers. Simulating using delay and noise the controllers still perform very similarly to the case without noise and delay. This is a positive result, as the introduction of noise and delay is intended to represent real-world conditions. The fact that the controllers continue to perform well and still demonstrate a clear improvement over the baseline controller indicates that they are inherently robust in a sense. The RWA behaviour

of the controllers differs slightly from the previous test, both in magnitude and in the rate of change. In this scenario, the controllers reach the rate limits more frequently, indicating that the RWA must be adjusted more rapidly to achieve the same performance as in the previous test in 4.1.1. The noise and delay primarily affect the actuator, RAS, rather than system stability. Frequent rate saturation could increase actuator wear and energy consumption and have implications for actuator wear and efficiency in practical implementations.

4.2 Braking Performance

Rear axle steering does not directly generate longitudinal braking force, which is expected. This is supported by the results, where the braking distance remains approximately unchanged across the evaluated scenarios. However, the results also show that the introduction of rear axle steering does not significantly increase the braking distance. One key insight is that a more aggressive ABS strategy can be applied. The ABS was made more aggressive by decreasing the ramp-up time between the allowed braking torques on the front axle. The time was decreased from 2 to 0 seconds, leading to continuous individual ABS operation on the front axle at all times. The controller used for this evaluation was the LQR controller, as it demonstrated stable and consistent performance across the braking scenarios.

4.2.1 Aggressive ABS on Straight Road

Below, the results from split-μ braking at a $\mu = (0.2, 1)$ surface at 80 km/h on a straight road using default and aggressive ABS are presented.

Disclaimer: There is no need to analyse Figure 4.37, 4.38, 4.39, 4.40, 4.41, and 4.42 beyond 9 seconds because the truck is almost at a standstill and the driver model produces extreme steering wheel angles to compensate for the lateral deviation caused by the scenario and steer the truck back to its intended route on the road.

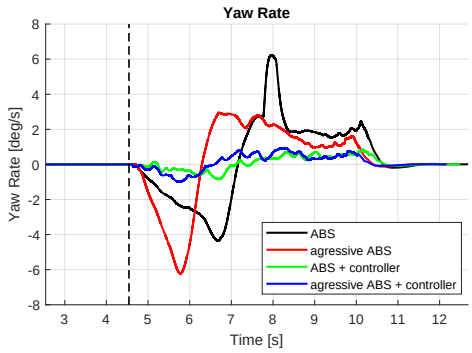

Figure 4.37: Yaw rate using different ABS settings and controller on straight road (0.2,1) 80 km/h.

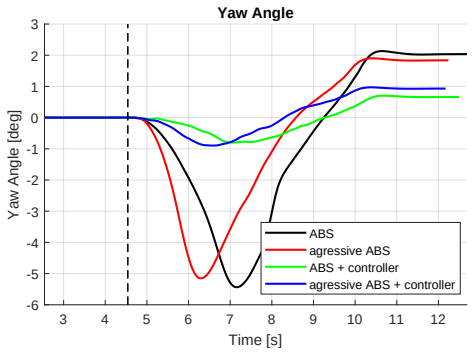

Figure 4.38: Yaw angle using different ABS settings and controller straight road (0.2,1) 80 km/h.

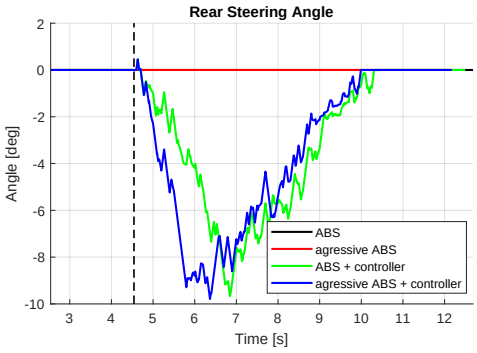

Figure 4.39: Rear steering angle different ABS settings and controller on straight road (0.2,1) 80 km/h.

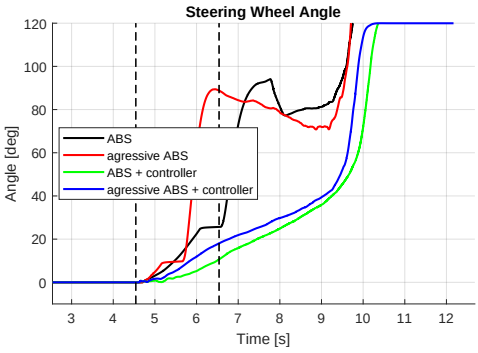

Figure 4.40: Steering wheel angle, different ABS settings, and controller on a straight road (0.2,1) 80 km/h.

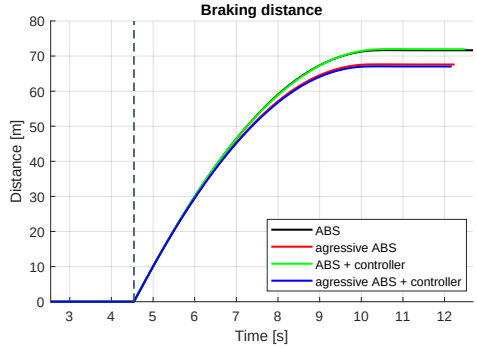

Figure 4.41: Braking distance using different ABS settings and controller on straight road (0.2,1) 80 km/h.

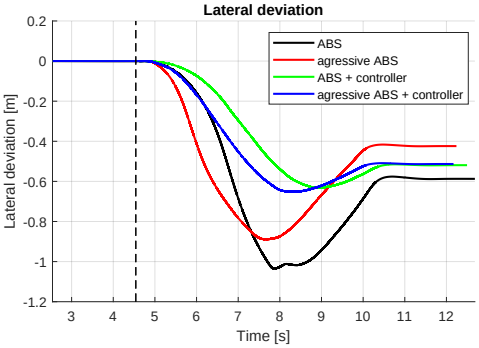

Figure 4.42: Lateral deviation using different ABS settings and controller on straight road (0.2,1) 80 km/h.

Figures 4.37, 4.38, 4.39, 4.40, 4.41, and 4.42 shows the difference between using only ABS and aggressive ABS together with the LQR controller. It is evident from the yaw rate and yaw angle that RAS control enhances both stability and brake distance, without introducing a trade-off between them. In both evaluated configurations, the LQR-based RAS controller outperforms the baseline case without RAS, for both the standard ABS and the more aggressive ABS. The results indicate that an aggressive ABS strategy, which may otherwise harm yaw stability, can be combined with RAS control to achieve superior overall performance compared to using the default ABS alone. Thus, improvements are obtained in both stability and braking distance. Looking at Figure 4.40 also reveals that the driver effort is reduced in both scenarios using RAS control. The slope also reveals that the behaviour of the truck is more predictable using RAS control compared to only using ABS and aggressive ABS. Since the tests were conducted at 80 km/h and stability increases as speed decreases, it can be concluded that the combination of RAS control and aggressive ABS will generate improvements for speeds from 80 km/h down to 0 km/h.

4.3 Straight Road with Feedback and Feed Forward

An additional test was conducted using a combination of feedback and feed-forward control, since in practical control applications, feedback and feed-forward controllers are typically used together. However, in the previous sections, the controllers were assessed individually in order to clearly isolate their respective contributions and performance traits. Based on these earlier assessments, the LQR controller was identified as the most suitable feedback controller due to its balance between stability and control effort when compared to the PID controller.

The feed-forward component used is the BP_{FF} controller, which was previously evaluated as a stand-alone controller in all test scenarios. While the feed-forward controller on its own demonstrated poor overall performance, it showed a fast response and was particularly effective at counteracting the initial yaw moment. For this reason, it is of interest to investigate whether the feed-forward controller can enhance the performance of a high-performing feedback controller like the LQR.

The objective of this test is to evaluate if the fast disturbance cancellation characteristics of the feed-forward controller can complement the LQR feedback controller by further reducing the initial yaw rate response and steering effort, while maintaining overall vehicle stability. The performance of the combined LQR + BP_{FF} controller is compared against both the baseline and the LQR controller operating without feed-forward control. Below, the results from split-μ braking at a (0.2,1) surface at 80 km/h on a straight road are presented. This test corresponds to the road presented in 3.3.

Disclaimer: There is no need to analyse Figure 4.43, 4.44, 4.45, 4.46, 4.47, and 4.48 beyond 9 seconds because the truck is almost at a standstill and the driver model produces extreme steering wheel angles to compensate for the lateral deviation caused by the scenario and steer the truck back to its intended route on the road.

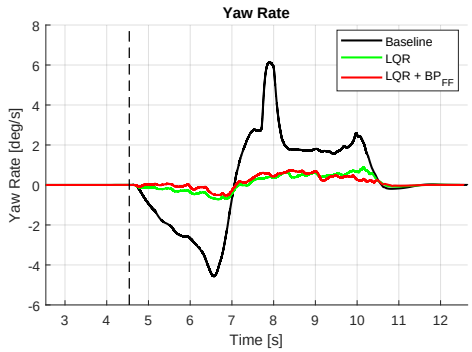

Figure 4.43: Yaw rate on straight road (0.2,1) 80 km/h (FB+FF).

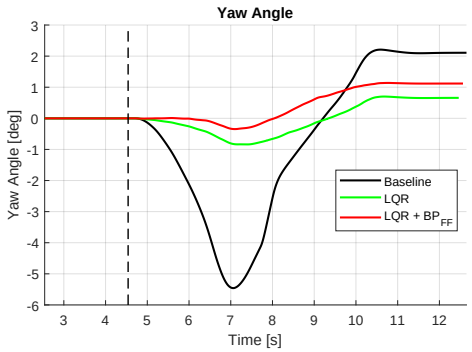

Figure 4.44: Yaw angle on straight road (0.2,1) 80 km/h (FB+FF).

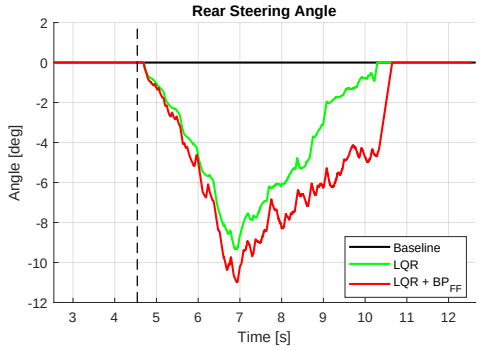

Figure 4.45: Rear steering angle on straight road (0.2,1) 80 km/h (FB+FF).

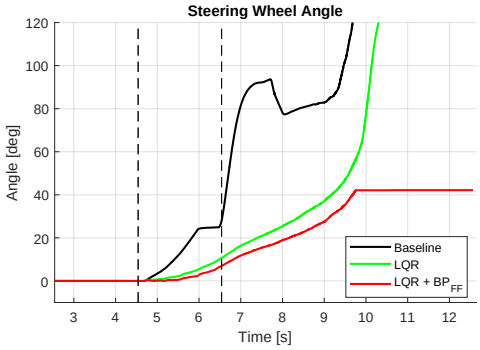

Figure 4.46: Steering wheel angle on straight road (0.2,1) 80 km/h (FB+FF).

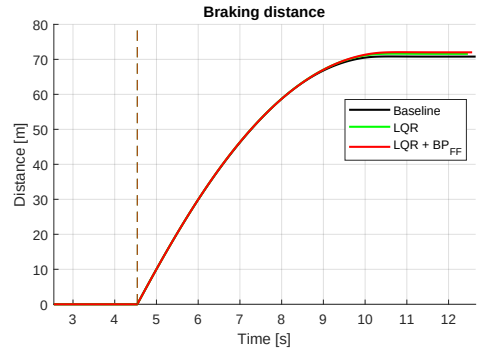

Figure 4.47: Braking distance on straight road (0.2,1) 80 km/h (FB+FF).

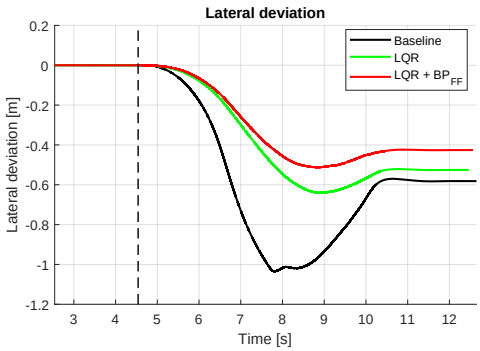

Figure 4.48: Lateral deviation on straight road (0.2,1) 80 km/h (FB+FF).

The braking distances for the different controllers and RMS can be viewed more clearly in the tables below.

Controller	Brake Distance [m]
Baseline	70.830
LQR	71.618
LQR + BP_{FF}	72.054

Table 4.12: Brake distance for each controller on straight road (0.2,1) at 80 km/h (FB+FF).

Controller	RMS yaw rate [deg/s]
Baseline	2.3294
LQR	0.4426
LQR + BP_{FF}	0.3556

Table 4.13: RMS for each controller on straight road (0.2,1) at 80 km/h (FB+FF).

Compared to both the baseline and the LQR-only configuration, the combined LQR + BP_{FF} controller shows a reduced initial yaw rate peak immediately following the braking on the split-μ surface. This indicates that the feed-forward action effectively counteracts the asymmetric braking forces before the feedback controller's response fully develops. This reduction in yaw rate directly translates to a smaller yaw angle during the first seconds after braking is initiated. A reduction in steering wheel angle can also be observed compared to the baseline and LQR-only case, showing that the addition of feed-forward control reduces the corrective steering effort required by the driver.

Looking at the RWA, a bit more aggressive behaviour is observed with LQR + BP_{FF} compared to LQR alone. It uses a slightly higher RWA consistently, but overall follows the same shape, oscillates more, and changes at a higher rate. However, this shows that stability improvements are achieved without introducing overly aggressive rear axle steering inputs. Additionally, lateral deviation is also reduced. The inclusion of feed-forward control does, however, slightly increase the braking distance compared to both the baseline and LQR-only configurations. This increase is, however, marginal and remains within an acceptable range.

A notable improvement can also be observed in the RMS yaw rate error, where the combined LQR + BP_{FF} controller achieves the lowest value among all tested configurations. Overall, the results confirm that combining feedback and feed-forward control results in improved yaw stability and reduced driver effort compared to using feedback alone.

4.4 Stability Assessment

This section presents and analyses the results of the stability assessments done on the controllers, where stability is calculated as described in Chapter 3 in section 3.5. The BP_{FF} controller is not included since the aim is to assess the driver-intent stability, in other words, the reference yaw rate to actual yaw rate.

Figures 4.49, 4.50, 4.51, and 4.52 show the stability in the form of the induced norm of the closed loop from reference yaw rate to actual yaw rate. Table 3.4 presents the assessment guidelines for the different degrees of stability. Additionally, the frequency response of the closed-loop systems with the controllers across different speeds is shown in Figures 4.53, 4.54, 4.55, and 4.56. The colour map is in decibels (dB) and illustrates the stability as a function of frequency (rad/s) and speed (km/h). Blue denotes a region of high attenuation through the system, green denotes good stability and tracking performance and red indicates unstable behaviour.

The closed-loop induced norm of the PID controller, shown in Figure 4.49, increases with speed, starting at approximately 1 at low speeds and reaching about 4.3 at 80 km/h. This behaviour indicates good stability and tracking at low speeds but deteriorates as speed increases. At 80 km/h, the worst case amplification of the driver's intent or disturbance is ≈ 4.3, which suggests poor robustness and a large tendency to induce unstable behaviour.

Figure 4.53 helps identify the region of worst-case behaviour. The induced norms location in frequency is given by the maximum peak at each speed and is approximately equal to 70 rad/s or 11 Hz. However, the possible input the driver can induce is below this region and is approximately 0.62 to 19 rad/s or 0.1 to 3 Hz. This indicates that the region of high unwanted gain is outside the bandwidth of the driver. It will, however, still be excited by any disturbances that enter the system

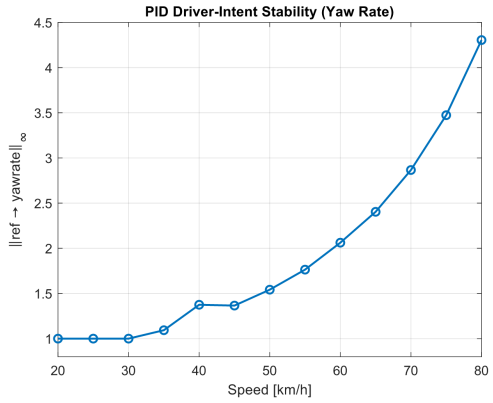

Figure 4.49: PID closed loop induced norm stability.

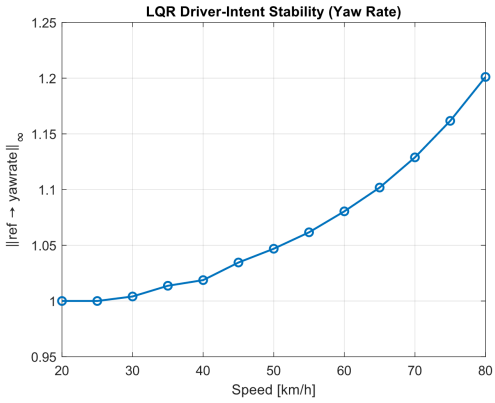

Figure 4.50: LQR closed loop induced norm stability.

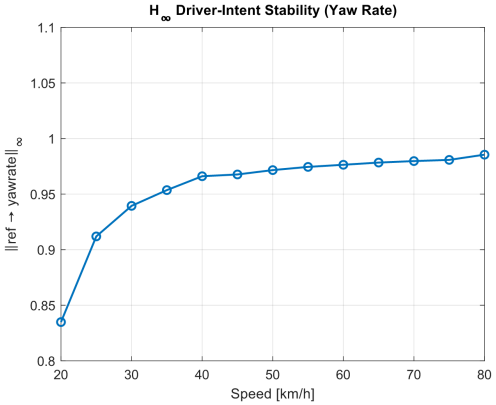

Figure 4.51: $\mathcal{H}_\infty$ closed loop induced norm stability.

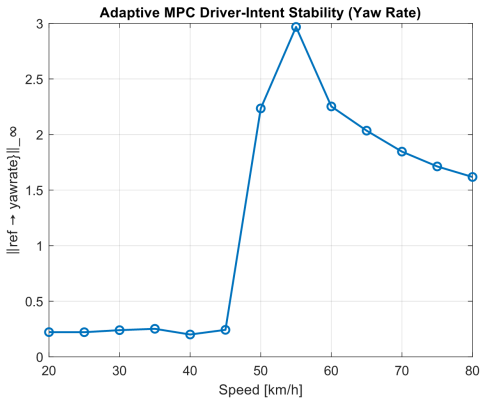

Figure 4.52: MPC closed loop induced norm stability.

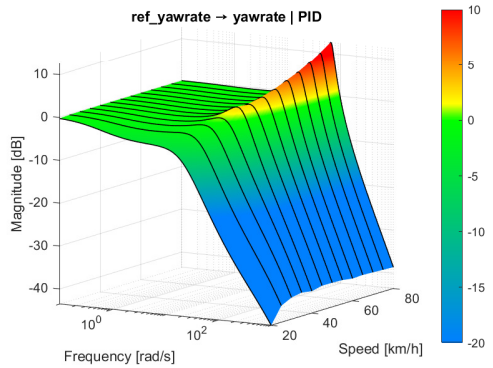

Figure 4.53: PID closed loop response.

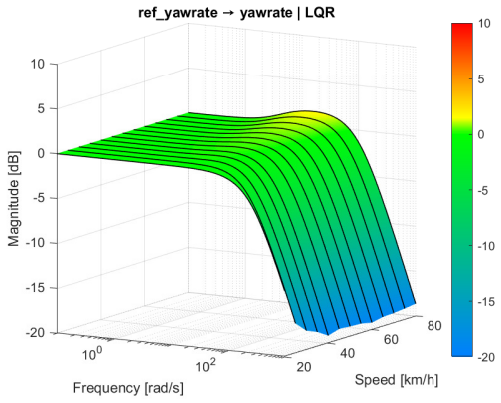

Figure 4.54: LQR closed loop response.

in this region. High-speed tracking performance is still possible, but the system is increasingly sensitive, which can compromise stability and robustness.

In contrast to the PID, the LQR controller shows consistent, stable, and robust

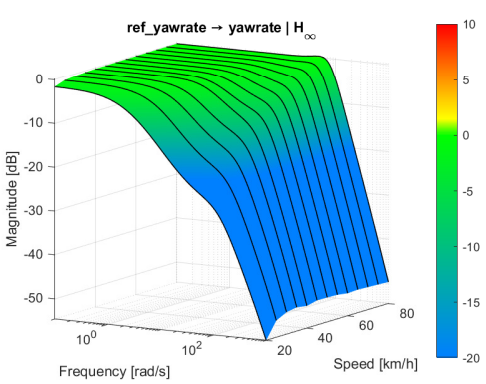

Figure 4.55: $\mathcal{H}_\infty$ closed loop response.

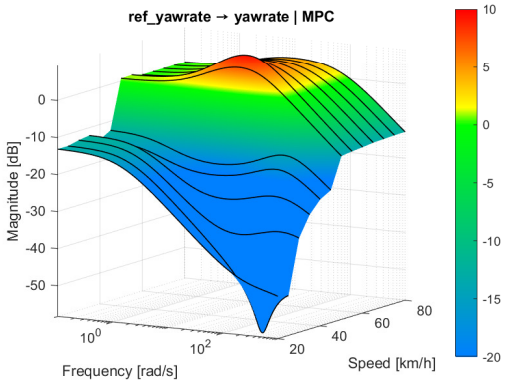

Figure 4.56: MPC closed loop response.

performance across the full speed range and can be seen in Figure 4.50. The growth across speed is not as large as the PID and only reaches a magnitude of 1.2 compared. The smooth increase suggests predictable performance degradation rather than instability as speed increases. The induced norm close to 1 implies that the system does not amplify the driver's intention or disturbance at any frequency, which is confirmed by Figure 4.54 where the entire region is mostly green and thereby stable.

The H_∞ controller is explicitly designed to handle worst-case input-to-output gain and model uncertainty. By minimising the $||G(s)||_\infty$, the controller guarantees robust stability with a performance level γ over the given model uncertainty. The Figure 4.51 shows that the induced norm is below unity across all speeds and can therefore be considered internally stable and well damped. Since $||T||_\infty < 1$, the vehicle response is bounded, and no disturbances are amplified through the system. At low speeds, the driver's intention is conservative and slightly attenuated, but as speed increases, tracking becomes more accurate, and the vehicle's responsiveness increases. The smooth convergence to ≈ 1 reveals good scheduling of the controller.

The Adaptive MPC demonstrates non-uniform stability and robust performance across the speed range. Figure 4.52 and 4.56 show this non-uniform behaviour. At low speeds, the MPC is very stable with high attenuation but consequently lacks perfect tracking performance. This can be regarded as acceptable because the system is limited towards the safe side. At approximately 50 km/h, a sudden increase in the induced norm is observed. This is highly undesirable and introduces a sudden, abrupt change that the driver might feel when traversing through that speed region, along with amplification of the driver's intention and disturbance. A probable cause of this is the difficulty of tuning the adaptive MPC, more precisely, the adaptive weights on yaw rate and lateral velocity that are scheduled on speed. In conjunction with the adaptive model in the MPC, this results in this non-uniform behaviour that was not foreseen as tuning of the controllers was done solely by taking the performance metrics from Chapter 3 subsection 3.1 into account during the split-μ braking scenarios.

The BP_{FF} is not a closed-loop controller and does not depend on driver input. Instead, it only reacts by trying to counteract the torque produced by asymmetric braking forces. To ensure stability of this controller, the output gain was tuned so that the counteracting yaw torque produced by the rear axle steering system is slightly less than that of what the asymmetric braking forces produce. Consequently, the gain is always less than 1 to ensure stability. This design choice reduces performance slightly but increases stability.

5
Conclusion

This chapter provides a reflection on the results and discussion presented in the previous chapter and addresses the research questions formulated in Chapter 1. The findings are discussed in relation to the overall purpose of the thesis. The chapter concludes with recommendations for future work and potential extensions of this research.

5.1 Summary of Work

In summary, the purpose and research questions formulated in sections 1.2 and 1.3, respectively, were fulfilled and answered. This thesis investigated the potential of RAS to improve the stability and controllability of heavy-duty vehicles during split-μ braking. The primary objective was to evaluate whether RAS, when controlled using different control strategies and constrained by realistic actuator limitations, can enhance yaw stability and controllability. Through simulations using a high-fidelity truck model, this thesis demonstrates that RAS can provide substantial stability benefits in split-μ braking situations.

An accurate LPV reference model, derived from one-track vehicle dynamics and validated against the high-fidelity model, was developed in MATLAB and Simulink and integrated with Truckmaker through a co-simulation framework. This was used to generate the desired yaw behaviour under nominal conditions. From this LPV model, several rear axle steering control strategies were developed and evaluated, including PID control, gain-scheduled LQR, adaptive MPC and robust H_∞ control as well as a brake-pressure-based feedforward controller. The controllers were evaluated on different split-μ braking scenarios with varying friction pairs and speeds, including straight and curved roads, lane-change manoeuvres, and scenarios with realistic sensor noise and actuator delay.

Overall, the results demonstrate that rear-axle steering significantly improves vehicle yaw behaviour across a wide range of split-μ braking scenarios by reducing peak yaw rate during the initial braking phase, decreasing lateral deviation, and reducing the corrective steering effort required of the driver. Although braking distance is largely unaffected by rear-axle steering alone, improved yaw stability enables a more aggressive ABS braking strategy without sacrificing stability or controllability.

5.2 Main Findings

The conclusion of this thesis is that rear axle steering substantially increases stability during split-μ braking, particularly during the initial phase immediately following brake application. The early yaw rate suppression is critical for preventing vehicle instability. An increase in stability during the initial phase of braking is more significant than in the later phases since the driver has limits when it comes to reaction time and making the correct steering actions. Since the controllers prove that they are capable of combating the asymmetric forces acting on the vehicle arising from braking at split-μ surfaces, the driver is not required to adjust their intended path by moving the steering wheel to the same extent as in the case without rear axle steering control. This can be attributed to the controller's objective of enforcing the expected vehicle behaviour regardless of road surface or operating conditions. The conclusion of this approach is that the transition into split-μ braking is less pronounced compared to the case without the RAS control, and driver effort is reduced across all scenarios.

Rear-axle steering control alone does not meaningfully increase or decrease braking distance. However, more importantly, when combined with aggressive ABS strategies, rear axle steering enables simultaneous improvements in braking performance and stability, removing the trade-off between stopping distance and yaw stability. A more aggressive ABS strategy cannot be used in isolation, as it leads to increased yaw instability and significantly higher driver steering effort. However, when combined with RAS control, aggressive ABS can be applied without increasing driver workload while maintaining vehicle stability and controllability. These two systems work well together and should be designed with each other in mind.

The actuator's steering rate limitation is the primary performance constraint, rather than the maximum steering angle. The results show that stability improvements can be achieved by using rear-axle steering angles within the actuator's capabilities. However, the controllers frequently reached rate saturation, indicating that faster rear axle steering dynamics would enhance effectiveness. The conclusion is that rear axle steering systems, intended for stability applications, benefit more from faster rear axle actuation rather than increased maximum steering angles.

Another finding concerned the necessity of an explicit split-μ detection system or algorithm capable of detecting a split-μ braking situation and, upon its triggering, activating the controller. Split-μ detection was initially considered central to this thesis, which motivated its inclusion as a research question. However, early simulation results revealed that effective yaw stabilisation can be achieved without explicit detection of split-μ conditions. The primary purpose of a split-μ detector would be to tell the controller when to start controlling or intervene. However, if the controllers regulate the error between the expected behaviour of the truck and the actual behaviour, the detection becomes obsolete. Instead, any deviation from the truck's expected behaviour is the controller's task to cancel. It does not matter where it comes from, and it could, for example, be wind, tyre wear, weight, etc.

When no friction asymmetry is present, no yaw moment arises, and the yaw rate error remains negligible, which will result in no intervention from the controller. In contrast, split-μ conditions generate a yaw moment that produces a measurable yaw rate error, which directly triggers the control action. This works because the main control strategy is feedback control, whereas the brake-pressure feedforward only works during split-μ braking. Also, detection of the split-μ surface takes time and will add a small delay before the controller can intervene. Since the initial part is the most critical, that time is not worth scarifying just to say that the vehicle is on a split-μ surface. Since the controller treats any deviation as a disturbance, the need to explicitly know when a split-μ surface is present is removed. Although split-μ detection remains an interesting and relevant problem, it was ultimately excluded from the scope of this thesis because the simulation results did not support its use. This insight revealed itself during the course of the work and was not known or considered in the beginning.

A finding that was particularly interesting concerns the LPV model. The LPV reference model proved to be an accurate way of representing the dynamics of the high-fidelity truck, but it could also be utilised to serve another purpose. The LPV reference model is needed to synthesise the LQR, MPC, and $\mathcal{H}_\infty$ stability controllers, but as a yaw rate reference generator, it could be modified to shape the preferred behaviour of the truck across different speeds, steering wheel angles, and scenarios, all using one controller. Figure 3.22 shows the LPV system used for reference generation of a yaw rate that the feedback controllers aim to track. This three-dimensional mapping of yaw rate with respect to speed and front-wheel angle can be modified to align with a preferred behavioural profile, rather than simply matching or replicating the real truck dynamics. As an example, the benefits of axle-Ackermann steering, which is commonly used at low speeds in trucks equipped with rear axle steering systems, can be preserved using the controllers developed in this thesis, without requiring any modifications to the controllers themselves to include Ackermann steering dynamics. This is possible since the axle-Ackerman behaviour can be mapped or described in the 3D LPV map by just increasing the absolute magnitude of the squares around low speeds. This will result in a steeper angle at low speeds in the 3D LPV map, which corresponds to the desired yaw rate, the reference, becoming higher than the real yaw rate produced by the truck with rear axle steering angle set to zero at low speed. This will cause the controller to countersteer the rear axle in order to produce a higher yaw rate to eliminate the error between the plant and the modified reference. Through this, low manoeuvrability is increased using the same architecture as the stability controllers, and no switching between modes and or controllers is necessary. Instead of defining different modes for the rear axle steering system based on specific scenarios, one single controller could be utilised together with a modified reference generator that represents the desired behaviour continuously across speeds and steering inputs.

Another use of the rear axle steering is for high-speed stability, where the goal could be to reduce yaw acceleration or reduce yaw rate during lane change. This could also be done using the stability controllers developed in this thesis in the same way

as for low manoeuvrability by just modifying the 3D LPV map at high speeds. For example, flattening the slope of the yaw rate at high speed by reducing the absolute magnitugde of the squares by some factor will result in a reference or desired yaw rate that is lower than the real measured yaw rate that the truck produces. This yaw rate error will cause the controller to steer the rear axle in the same direction as the front wheels in order to reduce the error and drive the system towards the desired yaw rate. This will reduce the yaw rate of the truck and could reduce the risk of unstable behaviour at high speeds. Effectively, what is done is that a safety limit is implemented by reducing the driver's ability to produce extreme yaw rates at high speeds. Setting the desired yaw rate lower than the plant dynamics of the truck at high speed also introduces the ability to crab-walk. With a lower reference yaw rate at high speeds, the stability controller will steer the rear axle in the same direction as the front wheels in order to reduce the yaw rate. Since the only way to keep a low yaw rate, as demanded by the reference, and still move laterally is to crab-walk, this results in a crab-walking behaviour.

The benefit of incorporating low-speed manoeuvrability, high-speed stability, crab-walk, etc., into the stability controllers, by modifying the LPV model developed in this thesis, is that the stability and robust performance guarantees of the control design are preserved, compared to defining the rear axle steering angle directly as a fixed relation to the front wheel angle. Also, any disturbance or uncertainties affecting the system, such as wind, tyre wear or weight, can be mitigated by the stability controller. High-speed stability as a fixed relation between front and rear wheel angles will produce different results based on these factors, while the stability controller always tries to drive the system to the desired behaviour.

In summary, the main finding regarding the LPV approach is that control design can be shifted from acting directly at the rear axle to shaping the reference instead. By tuning the reference model to create desired errors across different speeds and steering inputs, multiple aspects of vehicle behaviour can be captured. By using rear axle steering, low-speed manoeuvrability, crab-walking, and split-μ braking, stability can all be achieved using a single controller that simply fulfils the behaviour dictated by the reference model.

5.3 Future Work

One of the major conclusions drawn from the results of this thesis was that the dominant factor limiting performance is the steering rate rather than the maximum steering angle. Therefore, investigating the effect of a higher rear-axle steering rate on vehicle stability during split-μ braking could be an interesting direction for future work.

Trucks with trailers were outside the scope of this thesis and were therefore not included due to the additional complexity they would introduce in terms of dynamics, modelling, and control. However, future work could explore the use of rear axle steering for both the truck and the trailers' rear wheels to study its potential in

mitigating or preventing jackknifing during split-μ braking scenarios. The controller could possibly use articulation angle as a tracking reference for the trailer's rear axle steering controller.

Based on the findings, future work could also investigate a yaw rate reference generator that encapsulates multiple behaviours, such as low-speed manoeuvrability, high-speed, crab-walk, etc., while preserving system stability. It should probably be investigated how much the reference generator can deviate from the plant dynamics until it is no longer feasible or stable to command the controller to drive the system to that state. This could be investigated in both the time and frequency domains, since frequency-domain analyses might reveal inputs that the controller should enhance or reject.

Multiple controllers were employed throughout this project, providing valuable insights into different control strategies. Since multiple controllers were implemented, there was limited time to tune each. Proper tuning of the weighting matrices in LQR, $\mathcal{H}_\infty$, and MPC is essential for good performance and can substantially improve a controller's performance. Due to limited time, the authors believe there is room for improvement in most of the controllers through further tuning of the weights. Future work could therefore focus on one of the controllers developed in this thesis and further optimise its weighting parameters to achieve its maximum performance.

Improving the model fidelity is another area that could enhance controller performance and provide a closer representation of reality. In the current implementation, the truck dynamics under split-μ conditions are not explicitly modelled and are instead treated as disturbances. Future work could therefore be to model the truck under split-μ conditions. This could yield valuable insights but requires real-time knowledge of the road surface ahead, which is often difficult to obtain and requires sensors that are not standard on most vehicles.

Finally, future work could also include experimental validation by implementing a selection of the developed controllers on a real truck and evaluating their performance during a real split-μ braking scenario on a test track.

Bibliography

[1] National Highway Traffic Safety Administration, *Large Trucks: 2023 Data*, DOT HS 813 717, U.S. Department of Transportation, Apr. 2025. [Online]. Available: https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813717

[2] Federal Motor Carrier Safety Administration, *The Large Truck Crash Causation Study – Analysis Brief*, FMCSA-RRA-07-017, U.S. Department of Transportation, Jul. 2007. [Online]. Available: https://www.fmcsa.dot.gov/safety/research-and-analysis/large-truck-crash-causation-study-analysis-brief

[3] S. Kharrazi, M. Lidberg, P. Lingman, J.-I. Svensson, and N. Dela, "The effectiveness of rear axle steering on the yaw stability and responsiveness of a heavy truck," *Vehicle System Dynamics*, vol. 46, no. S1, pp. 365–372, 2008, doi: 10.1080/00423110801958568.

[4] Yangyan Gao, Timothy Gordon, Leon Henderson, Leo Laine, Shammi Rahman, and Anders Ivarsson "Control of Heavy Truck Braking on Split-Mu Surfaces Including Steering Response Adaptation"

[5] Sebastian Haglund and Henrik Johansson. "Steering Control During μ-split Braking for an Autonomous Heavy Road Vehicle," Master's Thesis, Department of Electrical Engineering, Linköping University, Linköping, Sweden, 2020. LiTH-ISY-EX–20/5304–SE.

[6] Torkel Glad and Lennart Ljung. *Control Theory: Multivariable and Nonlinear Methods*. Studentlitteratur, 2000. ISBN: 0-7484-0878-9.

[7] M. Burckhardt, *Fahrwerktechnik: Radschlupfregelsysteme*, Vogel Verlag, 1993.

[8] B. Kulcsar, *Lecture notes for Robust and Nonlinear control*, Chalmers univerity of technology, 2025.

[9] Trafikverket, *Tillsammans för Nollvisionen*. Available at: `https://bransch.trafikverket.se/for-dig-i-branschen/samarbete-med-branschen/Samarbeten-for-trafiksakerhet/tillsammans-for-nollvisionen/`, accessed April 27, 2026.

[10] International Organization for Standardization, *Road vehicles – Vehicle dynamics and road-holding ability – Vocabulary*, ISO 8855:2011, 2011. [Online]. Available: https://www.iso.org/standard/51180.html

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY