Control Design for Differential Lock Synchronization in Heavy-Duty Trucks

Master's thesis in Systems, Control and Mechatronics

Hampus Johansson
Lukas Karlhager

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

Master's thesis 2026

Control Design for Differential Lock Synchronization in Heavy-Duty Trucks

Hampus Johansson
Lukas Karlhager

Department of Electrical Engineering
Division of Systems and Control
Chalmers University of Technology
Gothenburg, Sweden 2026

Control Design for Differential Lock Synchronization in Heavy-Duty Trucks
HAMPUS JOHANSSON
LUKAS KARLHAGER

Supervisors: Basar Özkan and Esteban Gelso, Volvo Group Trucks Technology
Examiner: Jonas Fredriksson, Department of Electrical Engineering

Master's Thesis 2026
Department of Electrical Engineering
Division of Systems and Control
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: A heavy-duty truck driving off-road, the kind of low-traction environment where differential locks are used.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Control Design for Differential Lock Synchronization in Heavy-Duty Trucks
HAMPUS JOHANSSON
LUKAS KARLHAGER
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Heavy-duty trucks operating in low-traction environments rely on differential locks to maintain traction when a wheel spins out. These locks are commonly implemented with dog clutches, which require the connected shafts to be speed-matched before they can engage. Following a spin-out, achieving this match can force the driver to slow down or stop, wasting vehicle momentum and creating a safety risk on slopes. This thesis develops and compares active control strategies that synchronize the differential shafts after a wheel spin-out, enabling faster and safer dog clutch engagement. Individual wheel brakes and engine torque are used as actuators. A driveline model is derived for both the open and locked inter-axle differential configurations. A tire force estimator based on a Kalman filter provides feedforward disturbance cancellation, and a state transformation resolves an observability problem that arises when the inter-axle differential is locked. Three model-based controllers are designed and evaluated: a Linear-Quadratic Regulator (LQR), a Model Predictive Controller (MPC), and a Sliding Mode Controller (SMC). They are compared in simulation across split-friction and gravel road scenarios, using performance metrics for synchronization time, velocity loss, driver disturbance, and control effort, with tuning parameters swept to map the trade-offs between objectives. No significant trade-off is found between synchronization time and the remaining metrics: faster synchronization consistently coincides with lower velocity loss and does not worsen driver disturbance or control effort. A control strategy that follows the principles of the SMC is found to be best suited to the problem's disturbance-heavy nature. Active engine torque control reduces velocity loss when traction allows, while on low-traction surfaces it must instead be limited to avoid excessive brake demand.

Keywords: differential lock, dog clutch, synchronization, traction control, heavy-duty truck, optimal control, tire force estimation, linear-quadratic regulator, model predictive control, sliding mode control

Acknowledgements

We would like to thank our supervisors Basar Özkan and Esteban Gelso at Volvo Group Trucks Technology for their great support throughout the entire project, as well as everyone else at the company who have been involved and helped us along the way.

We also thank our examiner Jonas Fredriksson for many interesting discussions and helpful insights throughout the project.

Finally, we would like to thank our families and friends for their continuous support and encouragement throughout our studies. Hampus would like to extend a special thank you to his girlfriend, Mathilda, for her patience, understanding, and support throughout his entire education.

Hampus Johansson & Lukas Karlhager
Gothenburg, June 2026

List of Acronyms

The following acronyms are used throughout this thesis.

ABS Anti-lock Braking System.
ASR Anti-slip Regulation.

CARE Continuous-time Algebraic Riccati Equation.

ECU Electronic Control Unit.

IA Inter-Axle.
IW Inter-Wheel.

LQR Linear Quadratic Regulator.
LTI Linear Time-Invariant.

MPC Model Predictive Control.

RMS Root Mean Square.

SDG Sustainable Development Goal.
SMC Sliding Mode Control.

Nomenclature

The following notation is used throughout this thesis.

Vehicle Parameters

m	Total vehicle mass
J_e	Combined engine and transmission inertia
J_e^{eqv}	Equivalent engine inertia seen from the wheels
J_w	Wheel shaft inertia, assumed equal for all driven wheels
N_{IA}	Inter-axle differential gear ratio
N_{IW}	Inter-wheel differential gear ratio
r	Effective wheel radius

Vehicle States

v	Longitudinal vehicle velocity
$\dot{\psi}$	Vehicle yaw rate
α_{slope}	Road inclination angle
κ	Longitudinal slip ratio
ω_i	Angular velocity of wheel shaft i
ω_e	Engine shaft angular velocity
F_i	Longitudinal tire force at wheel i

Transformed Representation

$\omega_{1234}^{\mathrm{avg}}$	Average speed of the four rear wheels
ω_{12-34}^{Δ}	Front-to-rear speed difference
ω_{12}^{Δ}	Front-axle speed difference
ω_{34}^{Δ}	Rear-axle speed difference

F_{1234}^{avg}	Average tire force of the four rear wheels
F_{12-34}^{Δ}	Front-to-rear tire force difference
F_{12}^{Δ}	Front-axle tire force difference
F_{34}^{Δ}	Rear-axle tire force difference

Actuators and Torques

T_e	Engine torque
T_{bi}	Brake torque at wheel i
T_i	Driveline torque delivered to wheel i

Dog Clutch Engagement

t_{engage}	Minimum time the speed difference must remain within the engagement window before locking
$\omega_{\text{lb}}^{\Delta}$	Lower bound of the engagement window
$\omega_{\text{ub}}^{\Delta}$	Upper bound of the engagement window

Performance Metrics

t^{Δ}	Synchronization time
v^{Δ}	Velocity loss
j^{RMS}	Root mean square jerk, used as the driver disturbance metric
u_b^{tot}	Total brake torque impulse
u_e^{tot}	Total engine torque impulse

Estimator and Controller Parameters

$\omega_{\text{controller}}^{\Delta}$	Wheel speed difference threshold for controller activation
Q_{KF}, R_{KF}	Kalman filter process and measurement noise covariances
Q, R	State and control weighting matrices used by the controllers
K_{SMC}	SMC reaching-law gain matrix
ϕ	SMC boundary-layer thickness

Contents

List of Figures

List of Tables

1 Introduction

Transportation is essential to modern society, and heavy-duty trucks are a crucial part of this system, operating across diverse conditions, from paved roads to low-friction surfaces such as gravel, snow, and mud. To ensure good vehicle handling and allow for smooth turning, vehicles use differentials, which enable driven wheels on the same axle to rotate at different speeds. A standard open differential sends equal torque to all wheels. In low-traction conditions, this can cause a wheel spin-out, in which the torque applied to a wheel exceeds the available friction, causing the wheel to spin freely. When this happens, the wheel with higher friction slows down and produces less force, limiting the total traction to what the wheel with the least friction can provide.

To avoid traction loss, differential locks are used. A differential lock is a mechanical system that forces the wheels to rotate at the same speed, preventing wheel spin and improving traction. However, locking the differential worsens handling, especially when turning, so the lock should not be engaged at all times. In heavy-duty vehicles, differential locks are often implemented with a dog clutch, a mechanical device with interlocking teeth that physically connects the axles, due to the high torque demands in the driveline. For safe and reliable engagement of a dog clutch, the speed difference between the two shafts cannot be too great. If the differential lock is to be engaged after a wheel has spun out, the shaft speed difference needs to be reduced before the clutch can be engaged. This forces the driver to slow down before using the lock, thereby braking the vehicle's momentum and potentially causing a standstill. On slopes, this is especially problematic, as starting from a stop can be difficult and may require the driver to reverse and try the climb again, posing a safety risk.

This thesis investigates and develops control strategies for actively synchronizing differential shafts in heavy-duty trucks to enable faster and safer engagement of dog clutches following wheel spin in low-traction conditions. The proposed control system utilizes individual wheel brakes and engine torque as actuators, with the objective of maintaining vehicle momentum while minimizing driver intervention during the synchronization process.

1.1 Thesis Aim

The thesis aims to demonstrate a control approach that improves differential lock engagement performance compared to existing methods. A key requirement is that the controller must be suitable for implementation on a truck's Electronic Control Unit (ECU), imposing constraints on computational complexity. Furthermore, the proposed solution should be capable of synchronizing multiple differentials in vehicles equipped with multiple driven axles.

The thesis will address the following research questions:

1. What control strategy is most suitable for active synchronization of differential shafts, when evaluated in terms of synchronization time, vehicle velocity loss, driver disturbance, and control effort?
2. What trade-offs exist between these performance objectives, and which metrics can be improved simultaneously and which cannot?
3. What role does engine torque control play during the synchronization process, and how does limiting or removing it as a control input influence overall performance?

1.2 Limitations

This thesis is subject to the following limitations:

- Mechanical hardware effects, such as dog clutch wear, are not investigated.
- Interactions with other traction control systems are not considered. The controller is assumed to act independently.
- Low-level control of the lock device actuator is not addressed. It is assumed that a separate system handles actuation upon receiving a lock request signal.
- The conditions under which differential locks should be engaged, and which locks to select, are not analyzed in detail. The controller assumes all differentials should be locked following a spin-out.
- The analysis is limited to heavy-duty trucks with four driven rear wheels.

1.3 Related Work

This section presents an overview of prior work relevant to this thesis, covering the automation of differential lock engagement, traction control systems, and the synchronization of dog clutches in other applications.

1.3.1 Automatic Differential Lock Engagement

Existing studies on automating the engagement of differential locks in heavy vehicles cover both the conditions that should trigger locking and the practical challenges of achieving engagement.

Regarding when and in what order to engage the locks, it has been proposed in [1] that inter-axle differentials should be locked before inter-wheel differentials, balancing the traction gains against the negative effects on steerability. Trigger conditions studied in the literature include wheel-slip thresholds and driver-selected driving modes, with safety interlocks preventing engagement when the shaft-speed difference exceeds a defined limit to avoid powertrain damage [1], [2].

A practical challenge in differential locking is the tooth-on-tooth problem, in which near-zero shaft-speed differences cause the dog clutch faces to jam rather than engage [3]. One approach to resolving this is to use the service brakes to actively increase the shaft speed difference beyond the near-zero region, reducing the maximum lock time from 80 s to around 2 s in vehicle tests [4].

A common limitation across these works is that they rely on passively waiting for wheel speeds to naturally reach a suitable state, rather than actively controlling. This gap motivates the use of active control strategies for wheel-speed synchronization.

1.3.2 Traction Control Systems

Traction control systems represent a well-studied area of automotive control, where a common approach is to regulate the wheel slip ratio, a normalized measure of how much a wheel slides relative to the vehicle velocity, toward a reference value that maximizes tire force in the desired direction. Anti-lock Braking System (ABS) applies this principle to braking, using individual wheel brakes to prevent excessive slip during deceleration [5], while Anti-slip Regulation (ASR) applies it to acceleration, limiting driven wheel slip to maximize longitudinal traction [6]. Some works extend these concepts by incorporating additional objectives alongside slip tracking, such as minimizing the yaw moment induced by asymmetric braking forces on surfaces with different left- and right-side grip levels [7].

A phenomenon commonly neglected in the existing literature is the torque coupling that arises when shafts are connected through differential gears. Most proposed control strategies in the literature treat the driven wheels as independent dynamic systems and do not account for the effects of different brakes on different wheels.

1.3.3 Dog Clutch Synchronization in Transmission Systems

Dog clutches are already used in automotive transmissions, and numerous studies have examined their synchronization [3], [8], [9]. The findings are relevant to differential locking, since the underlying engagement mechanics are the same.

A central challenge identified across the literature is that successful engagement requires the shaft speed difference to be within a specific range at the moment of engagement. Near-zero speed differences risk tooth-on-tooth contact while excessive speed differences cause torsional vibrations, noise, and mechanical wear [3]. It has also been argued that full mechanical synchromesh is insufficient for modern heavy-duty transmissions and that control-based synchronization is necessary to meet comfort and durability requirements. To formalize this acceptable speed difference range, concepts such as engagement probability and shiftability conditions have been introduced, both of which describe the likelihood of successful engagement as a function of the initial speed mismatch [3], [8].

Various control approaches have been explored to achieve reliable engagement. These include trajectory-based methods using optimal control theory to avoid sleeve-to-gear impacts [9], as well as Linear Quadratic Regulators (LQRs) designed to track a reference engagement trajectory [8]. This shows that impact-free engagement is achievable through active control of speed and position.

1.4 Societal and Environmental Impact

Although the project focuses on a specific technical challenge, it has broader societal implications. Improved traction control has applications in both civilian and military vehicles. Civilian uses include rescue operations, forestry, mining, and other industries that require access to difficult terrain, contributing to improved safety, reduced downtime, and critical infrastructure work.

Improved off-road capability may also enhance military mobility, raising ethical concerns about potential use in armed conflict. This dual-use nature is important to acknowledge. The project itself focuses on developing a technical solution to improve safety and efficiency, and the ethical responsibility ultimately lies in how the technology is applied.

Environmental considerations must also be acknowledged. The controller developed in this thesis relies on active brake control, which generates particle emissions from brakes, tires, and clutches [10]. This is worth keeping in mind, particularly if the system operates the brakes more frequently or aggressively than a human driver would. At the same time, better traction control may reduce overall emissions by improving operational efficiency and reducing the need for repeated attempts in difficult terrain. The broader environmental impact of off-road industries such as mining and forestry must also be recognized, even if they remain necessary for societal development.

1.4.1 Sustainability Contributions

Improved off-road vehicle performance can support the United Nations' Sustainable Development Goals (SDGs) [11], by enabling the construction and maintenance of infrastructure in remote areas. Enhanced traction control helps vehicles access locations needed for projects related to:

- Clean Water and Sanitation (SDG 6) by enabling transport of equipment and materials for water supply and sanitation systems in rural regions.
- Affordable and Clean Energy (SDG 7) by supporting the construction of energy infrastructure, including renewable energy installations such as remote wind farms.
- Industry, Innovation, and Infrastructure (SDG 9) by contributing to infrastructure development in hard-to-reach environments.

While heavy-duty off-road operations have environmental impacts, the technology developed in this project can indirectly support long-term sustainable development by improving access to essential infrastructure projects.

1.5 Thesis Outline

Mathematical models of the vehicle and driveline dynamics are derived in Chapter 2 and serve as the foundation for the three model-based controllers developed in Chapter 3. The controllers are evaluated in a truck simulation environment provided by the company, described in Chapter 4, where two road scenarios and a set of performance metrics are defined. Multiple tunings are tested for each controller to reduce the influence of tuning on the comparisons, and the results are presented and discussed in Chapter 5. The thesis concludes with a summary of findings and suggestions for future work in Chapter 6.

2

System Modeling and Analysis

This chapter develops the mathematical models that form the foundation for controller design, covering vehicle and tire dynamics, dog clutch engagement conditions, driveline modeling, and state-space derivation.

2.1 Introduction to Trucks and Vehicle Dynamics

To better understand the content presented in later chapters, this section introduces fundamental concepts in vehicle and wheel dynamics, as well as truck-specific terminology.

2.1.1 Vehicle Coordinate System

The three-dimensional motion of a vehicle is described using three axes: longitudinal, lateral, and vertical. The rotations around these axes are referred to as roll, pitch, and yaw, respectively. The axes are fixed to the vehicle body as shown in Figure 2.1, defined according to ISO 8855 [12]. Rotation directions follow the right-hand rule.

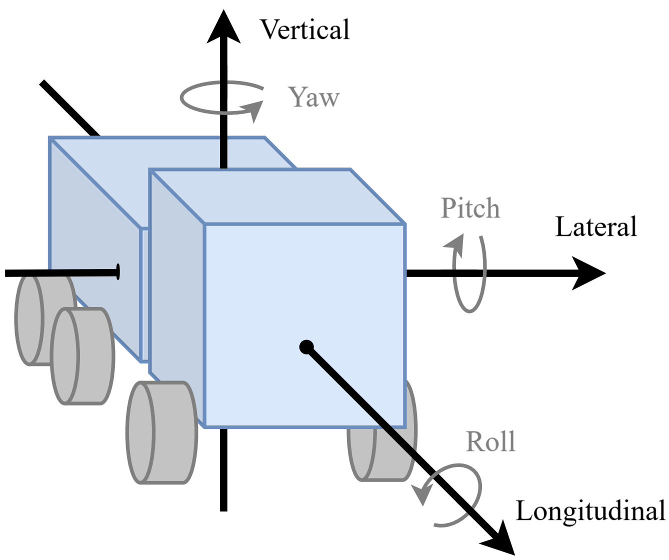

Figure 2.1: Vehicle coordinate system showing the longitudinal, lateral, and vertical axes.

2.1.2 Truck Nomenclature

Unlike passenger cars, trucks come in a wide variety of axle and wheel configurations. Key distinguishing features include:

- Number of axles.
- Number of steerable axles.
- Number of driven axles.
- Use of twin tires, i.e., two wheels mounted side by side on a single axle end.
- Liftable axles that can be raised or lowered depending on load.

A standard notation $A \times B$ denotes a truck with A total wheels and B driven wheels (twin tires counted as one). This thesis focuses on a 6×4 configuration, having four driven rear wheels with twin tires and no steering, and two undriven front wheels with steering and no twin tires. Rear axles are typically driven because the heaviest loads are distributed there, enabling greater traction.

Multiple driven axles require a hierarchy of differentials. Differentials connecting wheels on the same axle are called Inter-Wheel (IW) differentials, while those connecting multiple driven axles are called Inter-Axle (IA) differentials. The gearbox linking the front and rear driven axle groups is commonly referred to as the transfer case. For off-road applications, in-wheel transmissions may be added between the wheels and the IW differentials to provide further gear reduction.

The layout of these components for a 6×4 truck is shown in Figure 2.2, along with the wheel and differential numbering convention used throughout this thesis.

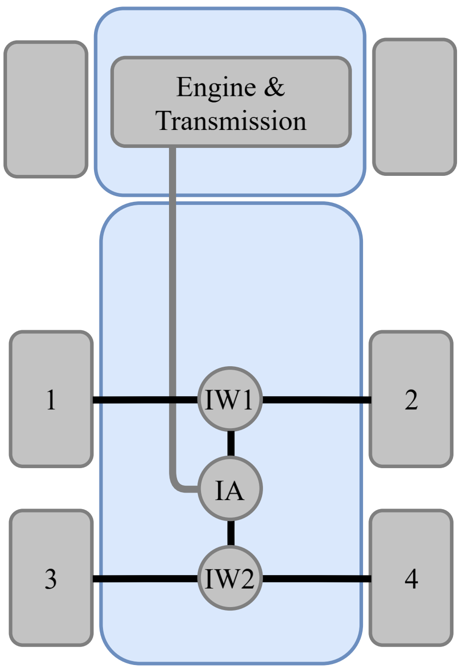

Figure 2.2: Wheel numbering and location of IA and IW differentials in a 6×4 truck.

2.1.3 Vehicle Dynamics

Traction control primarily concerns the longitudinal dynamics of the vehicle, since the tire forces acting in the longitudinal direction are responsible for vehicle acceleration and deceleration [13]. Furthermore, longitudinal tire forces are the primary controllable forces via actuation.

Figure 2.3 illustrates a simplified representation of the longitudinal dynamics of a 6×4 truck. The resistive forces, consisting of aerodynamic drag and rolling resistance, are lumped into a single term denoted by F_r. Applying Newton's second law in the longitudinal direction yields

$$m\dot{v} = F_1 + F_2 + F_3 + F_4 - mg\sin(\alpha_{\text{slope}}) - F_r, \tag{2.1}$$

where m denotes the total vehicle mass, g the gravitational acceleration, $\dot{v}$ the vehicle longitudinal acceleration, $F_i,\ i \in \{1,2,3,4\}$ the longitudinal tire forces generated by the driven wheels, and α_{slope} the road inclination angle.

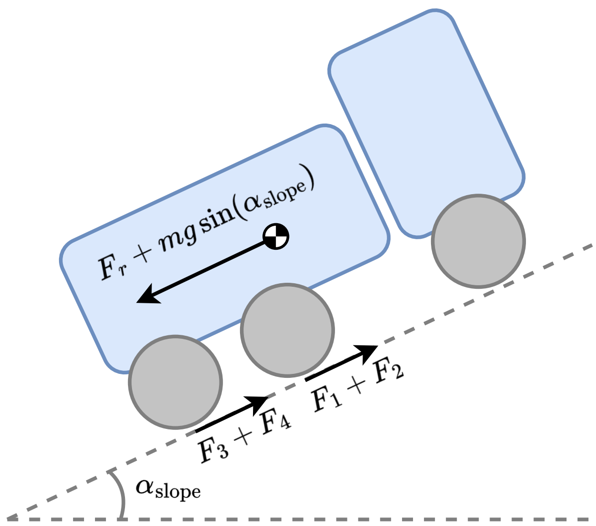

(a) Side view.

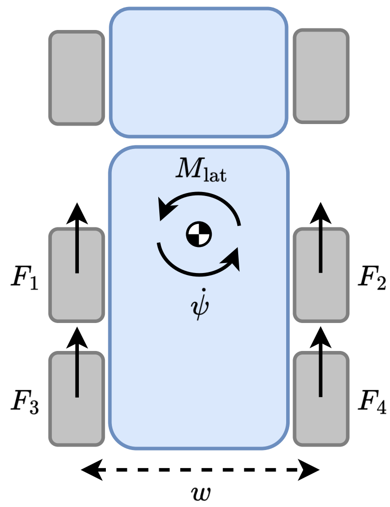

(b) Top view.

Figure 2.3: Simplified longitudinal dynamics of a 6×4 truck configuration.

Differences in longitudinal tire forces between the left and right sides of the vehicle may also generate a yaw moment, thereby influencing the lateral stability of the vehicle. Although lateral and vertical dynamics are not the primary focus of this thesis, the effect of longitudinal tire force differences on the yaw behavior of the vehicle is of relevance and will therefore be briefly considered.

A simplified expression for the yaw dynamics can be written as

$$J_z\dot{\psi} = \frac{w}{2}\left(-F_1 + F_2 - F_3 + F_4\right) + M_{\text{lat}}, \tag{2.2}$$

where $\dot{\psi}$ denotes the yaw rate, w the track width, J_z the vehicle's yaw moment of inertia, and M_{lat} represents additional yaw moment contributions arising from other lateral effects, such as vehicle sideslip dynamics and steering inputs.

2.1.4 Tire Dynamics

The longitudinal tire force depends primarily on the normal load and the slip ratio κ, defined as the normalized speed difference between the wheel surface and the road:

$$\kappa = \frac{r\omega - v}{|v|}, \tag{2.3}$$

where v is the longitudinal vehicle velocity, ω the wheel angular velocity, and r the effective wheel radius (treated as constant). A schematic of the relevant forces and velocities is shown in Figure 2.4.

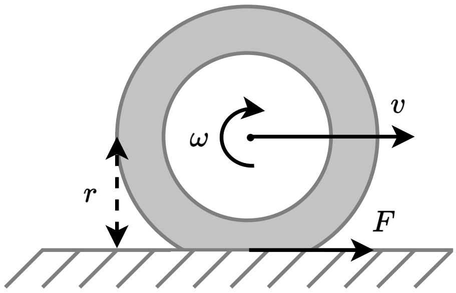

Figure 2.4: Simple wheel model showing longitudinal tire force and velocity.

The relationship between slip and tire force is nonlinear and depends on a variety of factors, including road surface, tire pressure, and weather conditions. As shown in Figure 2.5, the force increases approximately linearly with slip up to a peak value, beyond which it decreases [13]. Operating above this peak is therefore undesirable, as it reduces traction.

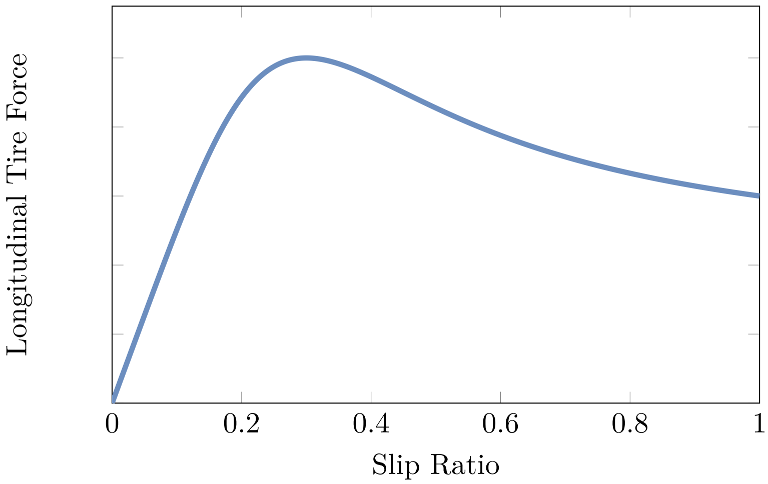

Figure 2.5: Typical longitudinal tire force vs. slip ratio characteristic.

Detailed mathematical models such as the Magic Formula [14] can capture this nonlinear behavior. Unfortunately, it requires numerous surface- and tire-specific parameters that are highly volatile and difficult to validate, making them poorly suited for real-time control applications. Linear tire models are significantly simpler to implement in real-time controllers. However, they are only valid within the approximately linear operating region and cannot capture the reduction in tire force that occurs at high wheel slip.

2.2 Dog Clutch Modeling

A dog clutch is a mechanical device that locks two shafts together using interlocking teeth, providing a rigid connection with very high torque transfer capability. Compared to a friction disc clutch, which transfers torque through surface friction without rigid locking, dog clutches are more compact and have lower energy losses [15]. These properties make them well-suited for differential locking in heavy-duty trucks, where high driveline torques are common. A conceptual illustration of a typical dog clutch used in these applications is shown in Figure 2.6.

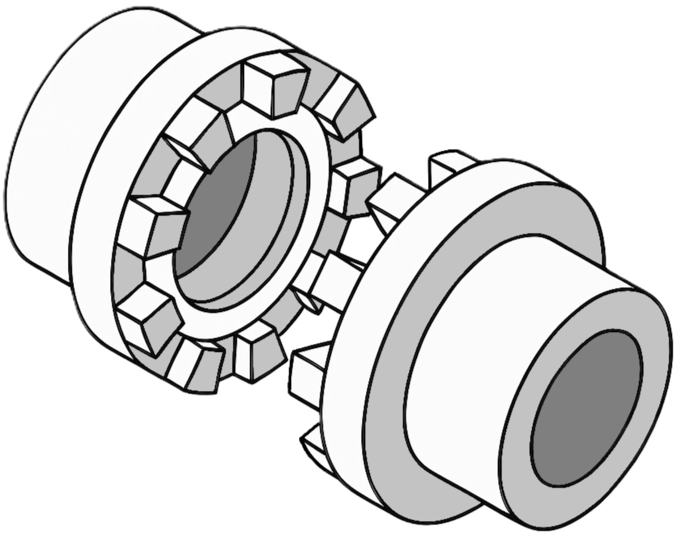

Figure 2.6: Conceptual illustration of a typical dog clutch used in truck applications.

Dog clutches must be synchronized before engagement. As mentioned in Section 1.3.3, synchromesh mechanisms are insufficient for heavy-duty applications, meaning that synchronization must be handled externally by the driver or an active control system before engagement is attempted.

2.2.1 Engagement Conditions

For reliable dog clutch engagement, the shaft speed difference at the moment of engagement must fall within a specific range, see Figure 2.7. If the speed difference is too large, the impact between the teeth causes mechanical wear and torsional vibrations in the driveline [16], [17]. The exact upper limit of this window of engagement depends on factors such as tooth geometry, chamfer angles, actuator force, and material properties [15], [17].

If the speed difference is too small, a tooth-on-tooth situation can occur, where the teeth of the two sides make face contact and jam against each other rather

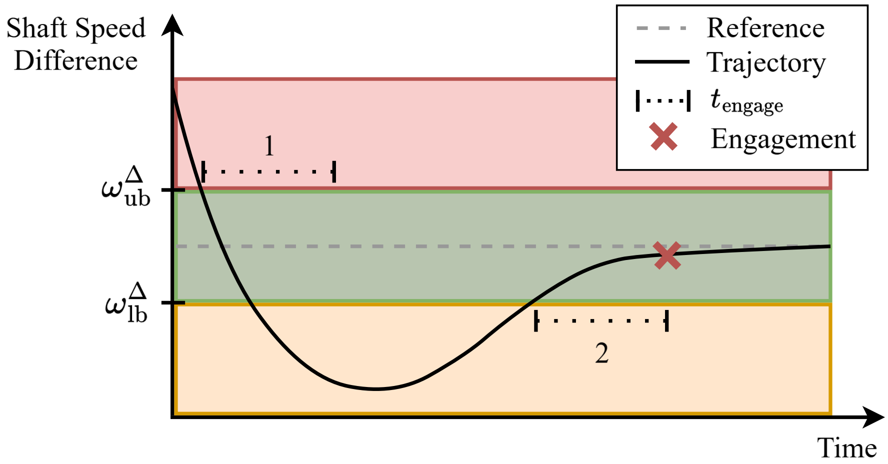

Figure 2.7: Speed difference trajectory over time, showing two instances of the speed difference entering the engagement window, shown in green. In the first instance, the difference leaves the window before t_{engage} is reached, and engagement does not occur. In the second, the condition is held, and engagement is triggered.

than sliding into the slots [3]. The lower bound on the speed difference is therefore non-zero.

Successful engagement thus requires the shaft speed difference to satisfy

$$\omega^{\Delta}_{\text{lb}} < \left|\omega^{\Delta}\right| < \omega^{\Delta}_{\text{ub}}, \tag{2.4}$$

where $\omega^{\Delta}_{\text{lb}}$ and $\omega^{\Delta}_{\text{ub}}$ are the lower and upper bounds of the engagement window. These bounds are specific to the hardware used and are treated as known parameters in this thesis. They are used to set the speed-difference references in the controller.

2.2.2 Simulation Model

Accurate modeling of dog clutch engagement is complex and highly hardware-dependent. Capturing the engagement dynamics in full requires modeling tooth geometry, tooth face friction, rebound effects, and high-frequency transients, which is computationally expensive [15]. In the simulation environment used in this thesis, the dog clutch is modeled as a friction clutch with very high torque-transfer capability.

2.3 Differentials and Driveline Modeling

As previously described, a differential is a mechanical device that connects three rotating shafts: one input shaft and two output shafts. It allows the two output shafts to rotate at different speeds while still transferring power from the input, which is necessary during cornering, where the inner and outer wheels travel different path lengths.

Differentials can operate in either an open or a locked state. In the open state, torque is split equally between the two output shafts while their speeds are free to differ. In the locked state, the output shafts are constrained to rotate at the same speed, and the torque is distributed according to the resistance on each shaft.

The following sections derive the governing equations for both states, assuming ideal differentials with no internal losses. The derivation starts with a single differential and then extends to the full double-axle configuration used in this thesis.

2.3.1 Differential Dynamics and Kinematics

Consider a single, isolated, and ideal differential, see Figure 2.8. The rotational dynamics of the connected shafts are

$$J_{\text{in}}\dot{\omega}_{\text{in}} = T_{\text{in}} - T_{\text{R}}, \tag{2.5a}$$
$$J_{\text{out},1}\dot{\omega}_{\text{out},1} = T_{\text{out},1}, \tag{2.5b}$$
$$J_{\text{out},2}\dot{\omega}_{\text{out},2} = T_{\text{out},2}, \tag{2.5c}$$

where J_{in}, $J_{\text{out},1}$, and $J_{\text{out},2}$ are the rotational inertias of the input and output shafts, ω_{in}, $\omega_{\text{out},1}$, and $\omega_{\text{out},2}$ are their angular velocities, T_{in} is the applied input torque, T_{R} is the reaction torque from the differential, and $T_{\text{out},1}$, $T_{\text{out},2}$ are the torques delivered to the output shafts. The dynamics hold for both open and locked configurations; only the kinematic constraints differ between the two.

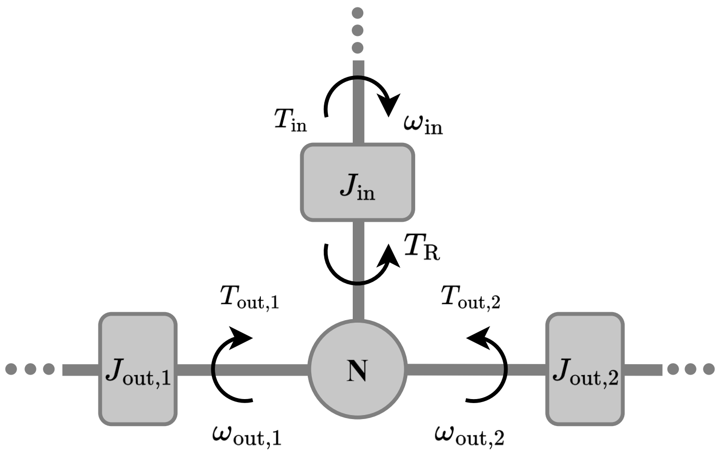

Figure 2.8: Illustration of a single, isolated differential with definitions of shaft speeds, torques, and inertias.

In the open configuration, the average output speed is related to the input speed through the gear ratio N, and the input torque is split equally between the two outputs [18],

$$\omega_{\text{in}} = \frac{\omega_{\text{out},1} + \omega_{\text{out},2}}{2}N, \qquad T_{\text{R}}\frac{N}{2} = T_{\text{out},1} = T_{\text{out},2}. \tag{2.6}$$

When locked, both output shafts are forced to rotate at the same speed, and the equal torque split is no longer enforced,

$$\omega_{\text{in}} = \omega_{\text{out},1} N = \omega_{\text{out},2} N, \qquad T_{\text{R}} N = T_{\text{out},1} + T_{\text{out},2}, \tag{2.7}$$

meaning torque is distributed according to the external load on each shaft.

2.3.2 Double-Axis Differential Configuration

In a double-axis differential configuration, see Figure 2.9, the rotational dynamics of the system are

$$\begin{aligned} J_e \dot{\omega}_e &= T_e - T_{e,\text{IA}}, &\text{(2.8a)}\\ J_f \dot{\omega}_f &= T_{f,\text{IA}} - T_{f,\text{IW}}, &\text{(2.8b)}\\ J_r \dot{\omega}_r &= T_{r,\text{IA}} - T_{r,\text{IW}}, &\text{(2.8c)}\\ J_w \dot{\omega}_i &= T_i - T_{bi} - F_i r, &\text{(2.8d)} \end{aligned}$$

where $i \in \{1, 2\}$ indexes the front axle output shafts and $i \in \{3, 4\}$ the rear, ω_e and T_e are the angular velocity and torque delivered from the engine and transmission, J_e is the combined rotational inertia of the engine and transmission, J_f and J_r are the rotational inertias of the intermediate shafts, ω_f and ω_r are their angular velocities, J_w and ω_i are the rotational inertia and angular velocity of each wheel shaft, T_i is the torque delivered to each wheel shaft from the driveline, $T_{e,\text{IA}}$ is the reaction torque from the IA differential on the engine shaft, $T_{f,\text{IA}}$ and $T_{r,\text{IA}}$ are the torques delivered by the IA differential to the front and rear intermediate shafts, $T_{f,\text{IW}}$ and $T_{r,\text{IW}}$ are the reaction torques from the front and rear IW differentials, T_{bi} and F_i are the brake torque and longitudinal tire force at each output shaft, and N_{IA} and N_{IW} are the IA and IW gear ratios respectively.

When all differentials are open, the constraints follow directly from (2.6) applied at each differential level,

$$\begin{aligned} \omega_e &= \frac{\omega_f + \omega_r}{2} N_{\text{IA}}, \qquad T_{e,\text{IA}} \frac{N_{\text{IA}}}{2} = T_{f,\text{IA}} = T_{r,\text{IA}}, &\text{(2.9a)}\\ \omega_f &= \frac{\omega_1 + \omega_2}{2} N_{\text{IW}}, \qquad T_{f,\text{IW}} \frac{N_{\text{IW}}}{2} = T_1 = T_2, &\text{(2.9b)}\\ \omega_r &= \frac{\omega_3 + \omega_4}{2} N_{\text{IW}}, \qquad T_{r,\text{IW}} \frac{N_{\text{IW}}}{2} = T_3 = T_4. &\text{(2.9c)} \end{aligned}$$

When the IA differential is locked, its constraints change according to (2.7) while the IW constraints remain unchanged,

$$\frac{\omega_e}{N_{\text{IA}}} = \omega_f = \omega_r, \qquad T_{e,\text{IA}} N_{\text{IA}} = T_{f,\text{IA}} + T_{r,\text{IA}}. \tag{2.10}$$

Locking the IA differential forces the front and rear intermediate shafts to rotate at the same speed, allowing the torque split between the two axles to vary freely with the load on each.

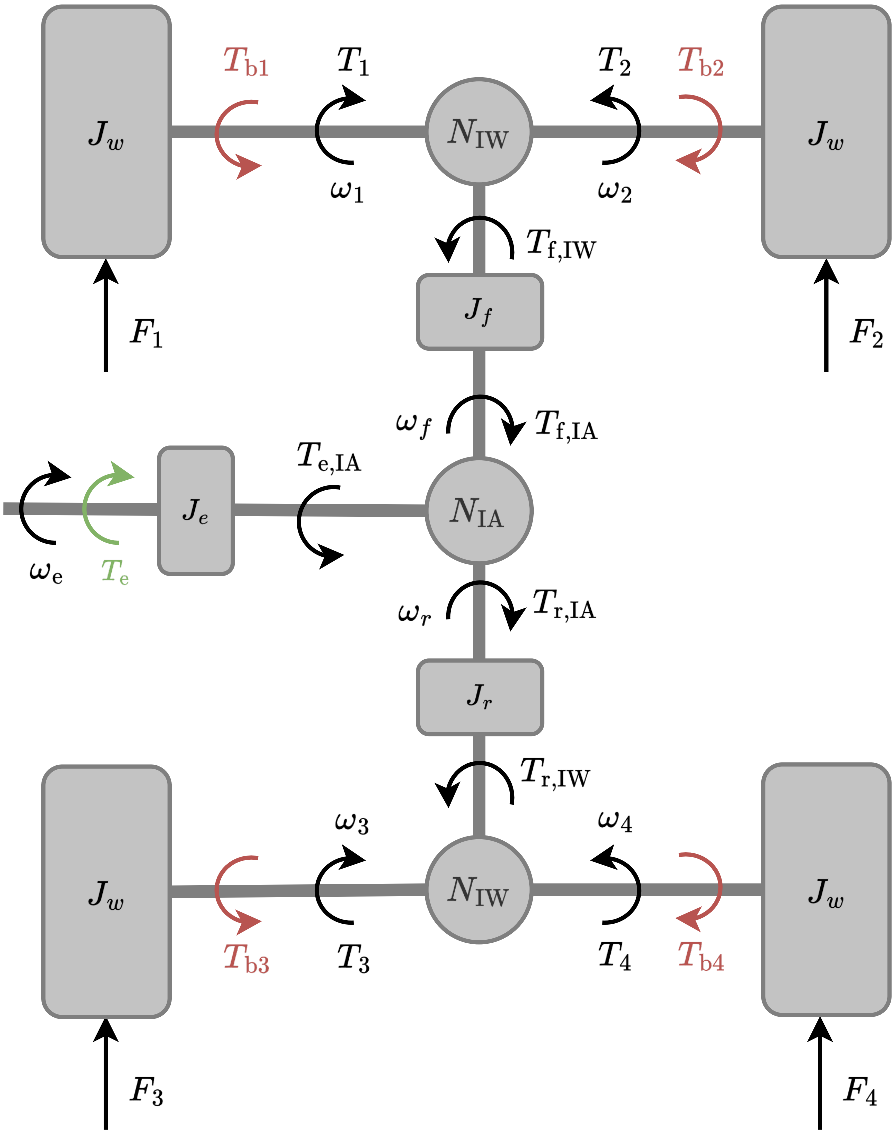

Figure 2.9: Illustration of a double-axis differential configuration with definitions of shaft speeds, torques, inertias, and tire forces.

2.4 State-Space Derivation

The driveline model is put into state-space form for use in the controller. Two models are derived: one with the IA differential open and one with it locked, corresponding to the two configurations encountered during the synchronization sequence.

2.4.1 IA Differential Open

With all differentials open, substituting the constraints (2.9) into the dynamics (2.8), and assuming negligable shaft inertias $J_f = J_r = 0$, yields the state-space representation

$$\dot{\boldsymbol{\omega}} = B_O\,\mathbf{u} + G_O\,\mathbf{v}, \tag{2.11a}$$

$$\boldsymbol{\omega} = \begin{bmatrix}\omega_1\\ \omega_2\\ \omega_3\\ \omega_4\end{bmatrix}, \quad \mathbf{u} = \begin{bmatrix}T_e\\ T_{b1}\\ T_{b2}\\ T_{b3}\\ T_{b4}\end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix}F_1\\ F_2\\ F_3\\ F_4\end{bmatrix}, \tag{2.11b}$$

$$\alpha_O = \frac{N_{\mathrm{IA}}\,N_{\mathrm{IW}}}{J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 4J_w}, \tag{2.11c}$$

$$\beta_O = \frac{J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2}{4J_w\left(J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 4J_w\right)}, \tag{2.11d}$$

$$\gamma_O = -\frac{3J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 16J_w}{4J_w\left(J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 4J_w\right)}, \tag{2.11e}$$

$$B_O = \begin{bmatrix}\alpha_O & \gamma_O & \beta_O & \beta_O & \beta_O\\ \alpha_O & \beta_O & \gamma_O & \beta_O & \beta_O\\ \alpha_O & \beta_O & \beta_O & \gamma_O & \beta_O\\ \alpha_O & \beta_O & \beta_O & \beta_O & \gamma_O\end{bmatrix}, \tag{2.11f}$$

$$G_O = \begin{bmatrix}\gamma_O & \beta_O & \beta_O & \beta_O\\ \beta_O & \gamma_O & \beta_O & \beta_O\\ \beta_O & \beta_O & \gamma_O & \beta_O\\ \beta_O & \beta_O & \beta_O & \gamma_O\end{bmatrix} r, \tag{2.11g}$$

where G_O equals the brake columns of B_O scaled by wheel radius r, meaning tire forces enter through the same channel as the brakes. This structure will be used later in the controller development.

2.4.2 IA Differential Locked

With the IA differential locked, applying the same procedure to (2.10) yields

$$\dot{\boldsymbol{\omega}} = B_L\,\mathbf{u} + G_L\,\mathbf{v}, \tag{2.12a}$$

$$\alpha_L = \frac{N_{\mathrm{IA}}\,N_{\mathrm{IW}}}{J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 4J_w}, \tag{2.12b}$$

$$\beta_L = \frac{J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 2J_w}{2J_w(J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 4J_w)}, \tag{2.12c}$$

$$\gamma_L = -\frac{J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 6J_w}{2J_w(J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 4J_w)}, \tag{2.12d}$$

$$\delta = -\frac{1}{J_e N_{\mathrm{IA}}^2 N_{\mathrm{IW}}^2 + 4J_w}, \tag{2.12e}$$

$$B_L = \begin{bmatrix} \alpha_L & \gamma_L & \beta_L & \delta & \delta \\ \alpha_L & \beta_L & \gamma_L & \delta & \delta \\ \alpha_L & \delta & \delta & \gamma_L & \beta_L \\ \alpha_L & \delta & \delta & \beta_L & \gamma_L \end{bmatrix}, \tag{2.12f}$$

$$G_L = \begin{bmatrix} \gamma_L & \beta_L & \delta & \delta \\ \beta_L & \gamma_L & \delta & \delta \\ \delta & \delta & \gamma_L & \beta_L \\ \delta & \delta & \beta_L & \gamma_L \end{bmatrix} r. \tag{2.12g}$$

Compared to (2.11), a fourth scalar δ has appeared. This reflects the fact that braking on one wheel now affects the wheels on the opposite axle differently than it affects the opposite wheel on the same axle. Since the locked IA differential constrains the average speeds of both axles to be equal, a braking torque on one axle also reduces the speeds on the other axle through this constraint.

2.4.3 Transformed State-Space Representation

The state-space representations, (2.11) and (2.12), use individual shaft speeds as states. This parametrization has two drawbacks. First, control tuning is less direct because differential locking depends on shaft-speed differences rather than on absolute speeds. Second, as will be shown in Section 3.2, the states become unobservable when the IA differential is locked. Both issues can be resolved by the linear state transformation

$$\tilde{\boldsymbol{\omega}} = \begin{bmatrix} \omega_{1234}^{\mathrm{avg}} \\ \omega_{12-34}^{\Delta} \\ \omega_{12}^{\Delta} \\ \omega_{34}^{\Delta} \end{bmatrix} = \underbrace{\begin{bmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} & -\frac{1}{2} \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}}_{T} \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \\ \omega_4 \end{bmatrix}, \tag{2.13}$$

where the states represent the overall average shaft speed, the front-to-rear average speed difference, and the left-to-right speed differences on the front and rear axles, respectively. Applying the transformation T to the system matrices yields the transformed state space representation

$$\dot{\boldsymbol{\omega}} = \tilde{B}\,\mathbf{u} + \tilde{G}\,\mathbf{v}, \qquad \tilde{B} = TB, \quad \tilde{G} = TG. \tag{2.14}$$

For the all-open configuration, the transformed matrices are

$$\tilde{B}_{\text{open}} = \begin{bmatrix} \frac{N_{\text{IA}}N_{\text{IW}}}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} \\ 0 & -\frac{1}{2J_w} & -\frac{1}{2J_w} & \frac{1}{2J_w} & \frac{1}{2J_w} \\ 0 & -\frac{1}{J_w} & \frac{1}{J_w} & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{J_w} & \frac{1}{J_w} \end{bmatrix}, \tag{2.15a}$$

$$\tilde{G}_{\text{open}} = \begin{bmatrix} -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} \\ -\frac{1}{2J_w} & -\frac{1}{2J_w} & \frac{1}{2J_w} & \frac{1}{2J_w} \\ -\frac{1}{J_w} & \frac{1}{J_w} & 0 & 0 \\ 0 & 0 & -\frac{1}{J_w} & \frac{1}{J_w} \end{bmatrix}. \tag{2.15b}$$

where $J_e^{\text{eqv}} = J_e N_{\text{IA}}^2 N_{\text{IW}}^2$ is the equivalent engine inertia as seen from the wheels, with the gear ratios squared reflecting the inertia scaling through the driveline. When the IA differential is locked, the corresponding state ω_{12-34}^{Δ} is forced to zero and can simply be removed from the state vector. The locked IA matrices are therefore obtained by omitting the second row,

$$\tilde{B}_{\text{IA-locked}} = \begin{bmatrix} \frac{N_{\text{IA}}N_{\text{IW}}}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} \\ 0 & -\frac{1}{J_w} & \frac{1}{J_w} & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{J_w} & \frac{1}{J_w} \end{bmatrix}, \tag{2.16a}$$

$$\tilde{G}_{\text{IA-locked}} = \begin{bmatrix} -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} & -\frac{1}{J_e^{\text{eqv}}+4J_w} \\ -\frac{1}{J_w} & \frac{1}{J_w} & 0 & 0 \\ 0 & 0 & -\frac{1}{J_w} & \frac{1}{J_w} \end{bmatrix}. \tag{2.16b}$$

2.5 Coupling Analysis

Due to the off-diagonal elements of the braking columns in the input matrices B_O and B_L, braking torque applied at one wheel affects not only the corresponding wheel speed, but also the speeds of the other wheels. This effect, referred to as coupling, is analyzed in this section for the two different differential configurations.

2.5.1 IA Differential Open

The coupling strength is defined as the braking effect on the remaining wheel speeds relative to the braking effect on the directly actuated wheel. From (2.11), the coupling ratio is

$$\frac{\beta_O}{\gamma_O} = -\frac{J_e^{\text{eqv}}}{3J_e^{\text{eqv}} + 16J_w}, \tag{2.17}$$

which depends only on the ratio

$$k = \frac{J_e^{\text{eqv}}}{J_w}. \tag{2.18}$$

Figure 2.10a shows how coupling varies with k. As $k \to 0$, the coupling approaches zero. This represents the case where the equivalent engine inertia is small compared to the wheel inertia, such as when the clutch is disengaged. As k grows, it goes towards a maximum value of $\frac{1}{3}$.

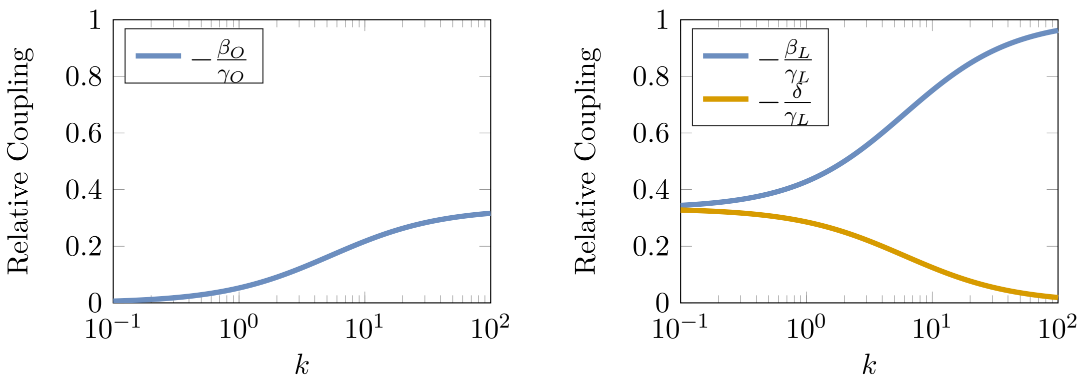

(a) All-open differential configuration. **(b)** Locked IA differential configuration.

Figure 2.10: Coupling between shaft speeds as a function of $k = J_e^{\text{eqv}}/J_w$.

2.5.2 IA Differential Locked

Locking the IA differential changes the coupling structure. Braking at a front wheel now affects rear wheel speeds differently than braking at the opposite front wheel, creating asymmetric coupling reflected by the coefficient δ in (2.12).

The coupling ratios become

$$\frac{\beta_L}{\gamma_L} = -\frac{J_e^{\text{eqv}} + 2J_w}{J_e^{\text{eqv}} + 6J_w}, \qquad \frac{\delta}{\gamma_L} = \frac{2J_w}{J_e^{\text{eqv}} + 6J_w}. \tag{2.19}$$

Figure 2.10b shows both ratios remain substantially different from zero across all values of k. This contrasts sharply with the all-open case, where coupling vanishes when k is small. Locking the IA differential therefore creates strong coupling between wheel speeds regardless of operating conditions.

2.5.3 Control Implications

The coupling analysis explains why decentralized wheel control is common in the literature. For typical vehicles, k is small, and the all-open configuration exhibits negligible coupling in this range. Decentralized control assumptions become reasonable under these conditions.

The locked IA configuration, however, shows strong coupling regardless of k. Since the controllers developed in this thesis operate in this configuration, a centralized control structure that explicitly accounts for coupling between wheel dynamics is required.

3

Controller Development

Developing a control system for differential synchronization presents several challenges. Synchronization is almost exclusively needed under poor traction conditions, such as on ice, where no tire has adequate grip, or on uneven gravel roads, where a tire can lose ground contact almost instantaneously. These operating conditions are inherently disturbance-heavy, placing strong demands on disturbance rejection. At the same time, the controller must achieve synchronization within short time windows while minimizing impacts on driver comfort and vehicle stability. The controller must also be computationally lightweight enough for deployment on an ECU.

An additional challenge arises from the system's over-actuated nature. When both the engine and individual wheel brakes are available as actuators, multiple actuator combinations can produce the same net effect on the system. For instance, an increase in engine torque can be entirely negated by an equal and opposite braking torque, leaving the system state unchanged. This redundancy constitutes a control allocation problem that can complicate the design of model-based and optimal controllers.

This chapter presents the development of three distinct control strategies: an LQR, followed by Model Predictive Control (MPC), and finally Sliding Mode Control (SMC). Each controller strategy is first presented in theory, followed by a problem-specific design description and implementation.

3.1 Shared Controller Design Principles

The operating challenges outlined above guide the design of all three controllers proposed in this chapter. This section describes the principles and structural choices shared across all three formulations.

3.1.1 Controller Engagement Logic

A detailed analysis of exactly when and under what conditions the differential locks should be engaged is outside the scope of this thesis. However, a simple engagement logic is needed to evaluate the controller. The controller activates when a wheel spin-out is detected, defined as any wheel speed difference among the driven wheels exceeding a threshold $\omega^{\Delta}_{\text{controller}}$. This condition is evaluated across three speed differences: the left-right difference on each of the two driven axles, and the difference between the average speeds of the two axles. When the controller activates, its goal is to synchronize and lock all three differentials.

Once activated, the IA differential is synchronized and locked first, followed by both IW differentials. This order is motivated by the fact that locking the IA differential gives the largest traction gains relative to the handling losses it introduces [1]. A lock request is sent to the dog clutch actuator when the relevant shaft speed difference falls under the safe upper limit $\omega_{\mathrm{ub}}^{\Delta}$, as described in Section 2.2.

3.1.2 Model-Based Control Architecture

All three controllers are model-based and use a centralized state-space representation of the wheel speeds, derived in Section 2.4.3. This gives each controller explicit knowledge of the coupling between wheel speeds through the driveline, rather than treating these couplings as unexpected disturbances.

Since the locking sequence changes the driveline configuration, two models are needed. The first is used while the IA differential is open, and the second is used once it is locked, with ω_{12-34}^{Δ} removed since it is forced to zero by the lock. The state matrices are modified accordingly, as described in Section 2.4.3.

3.1.3 Tire Force Disturbance Cancellation

Tire force-induced torques represent the dominant external disturbance acting on the driveline. Since these torques enter the system through the same channels as the control inputs, they can be directly compensated for through feedforward action. In practice, the tire forces are not directly measured and must instead be estimated, for example, using a dedicated tire force estimator. The estimates are then used to compute opposing torques, which are applied through the engine and brakes to cancel the disturbance.

Ideally, perfect cancellation would reduce the closed-loop dynamics to the wheel and driveline inertias alone, leaving a simpler system for the feedback controller to regulate. By directly compensating for the estimated disturbance, there is less need for aggressive feedback or integral gains to achieve a good response speed of the controller. In practice, the disturbance estimates are never perfect due to modeling inaccuracies and measurement uncertainties. Consequently, the feedforward term does not eliminate disturbances entirely but rather reduces their influence, allowing the feedback controller to handle the remaining errors more effectively.

3.1.4 Reference Setting

Each controller state has a corresponding reference. The speed difference references $\omega_{\mathrm{ref}}^{\Delta}$ and the average speed reference $\omega_{\mathrm{ref}}^{\mathrm{avg}}$ are set differently and serve different purposes. The speed difference references are set to a value within the safe engagement window of the dog clutch, as described in Section 2.2. A speed difference that is too large risks mechanical damage on engagement, while one that is too small risks a tooth-on-tooth situation. The reference is therefore set to a value between these limits, as illustrated in Figure 3.1.

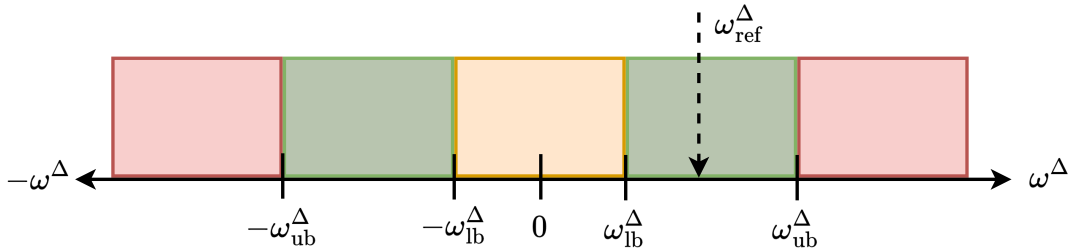

Figure 3.1: Illustration of dog clutch engagement capability as a function of shaft speed difference. Green indicates feasible engagement, red indicates infeasible engagement, and yellow indicates risk of a tooth-on-tooth situation.

Even though the IA differential is synchronized first, references are set for all speed difference states simultaneously. The priority between states can be controlled through controller tuning by assigning a higher state-deviation cost to the state corresponding to the differential currently being synchronized.

The reference for a given speed difference state can also be set to target a specific side, since both green regions in Figure 3.1 are valid. For example, if the driver is turning right, the references for the IW states could be set such that the left-side wheels spin faster than the right-side wheels, supporting the intended yaw behavior. Similarly, the IW references could be set opposite in the case that the driver drives straight in order to minimize the induced yaw moment.

The average speed reference $\omega^{\text{avg}}_{\text{ref}}$ is set to maintain a target longitudinal slip ratio κ_{ref}, maximizing forward tire force during synchronization and thereby minimizing vehicle speed loss. This is achieved by setting

$$\omega^{\text{avg}}_{\text{ref}} = \frac{(1+\kappa_{\text{ref}})v}{r}, \tag{3.1}$$

where κ_{ref} is chosen near the peak of the tire force curve, around 0.3 as shown in Figure 2.5.

3.1.5 Actuator Constraints

All actuators are subject to physical torque limits. The engine cannot reliably produce negative torque and is therefore constrained to non-negative torque requests. Its upper limit is set either by the physical engine limit or, for safety reasons, by the torque currently requested by the driver. The wheel brakes are strictly dissipative, so brake torque is constrained between zero and the physical brake capacity.

Respecting these constraints is critical. Commanding torques outside the physically realizable range breaks the assumptions of the underlying model and can compromise the validity of the control solution when deployed on a real system.

3.2 Tire Force Estimator

Accurate knowledge of longitudinal tire forces is important for the controller, as they are the primary disturbances. Rather than relying on a tire model, the forces are treated as unknown states and estimated alongside the wheel speeds using a Kalman filter [19]. This approach is preferred for the same reasons discussed in Section 2.1.4: linear tire models are invalid at the high slip ratios of interest and nonlinear models are computationally expensive and require tire- and surface-specific parameters that are difficult to obtain and validate in practice.

3.2.1 Kalman Filter

Consider a Linear Time-Invariant (LTI) system

$$\dot{x} = Ax + Bu + Dw, \tag{3.2}$$

$$y = Cx + v, \tag{3.3}$$

where $x \in \mathbb{R}^n$ is the state vector, $u \in \mathbb{R}^m$ the input, w the process disturbance, and $y \in \mathbb{R}^p$ the measurement. The process noise w and measurement noise v are modeled as zero-mean white noise processes with covariances $Q \succeq 0$ and $R \succ 0$, respectively.

The Kalman-Bucy filter provides the optimal continuous-time estimate $\hat{x}$ of the state through

$$\dot{\hat{x}} = A\hat{x} + Bu + L\Big(y - C\hat{x}\Big), \tag{3.4}$$

where $L = PC^\top R^{-1}$ is the filter gain and P is the estimation error covariance [20]. P evolves according to the Riccati differential equation

$$\dot{P} = AP + PA^\top + DQD^\top - PC^\top R^{-1}CP. \tag{3.5}$$

If $(A,\, DQ^{1/2})$ is stabilizable and (A, C) is detectable, P converges to a steady-state value P_∞ satisfying the Continuous-time Algebraic Riccati Equation (CARE)

$$AP_\infty + P_\infty A^\top + DQD^\top - P_\infty C^\top R^{-1}CP_\infty = 0, \tag{3.6}$$

yielding a constant gain $L_\infty = P_\infty C^\top R^{-1}$ and exponentially stable error dynamics.

3.2.2 Estimator Implementation

The tire force estimator is implemented to support both IA and IW differential synchronization. Since the IA differential is locked during the IW synchronization process, the estimator must use different system models depending on the current lock status.

3.2.2.1 Observability Problem with a Locked IA Differential

Estimating tire forces from wheel-speed measurements requires that the estimator system be observable. This section shows that when the IA differential is locked, a state estimator based on individual wheel speeds, together with random-walk tire-force states, becomes unobservable, and identifies the specific tire-force combination that cannot be reconstructed.

As discussed previously, the operating conditions considered in this thesis motivate modeling the tire forces as a random walk

$$\dot{\mathbf{F}} = \begin{bmatrix}\dot{F}_1 & \dot{F}_2 & \dot{F}_3 & \dot{F}_4\end{bmatrix}^\top = 0_{4\times1}. \tag{3.7}$$

Assuming the individual wheel speeds as states,

$$\boldsymbol{\omega} = \begin{bmatrix}\omega_1 & \omega_2 & \omega_3 & \omega_4\end{bmatrix}^\top, \tag{3.8}$$

together with the individual tire force states, yields the estimator state vector

$$\mathbf{z} = \begin{bmatrix}\boldsymbol{\omega}^\top & \mathbf{F}^\top\end{bmatrix}^\top \tag{3.9}$$

Assuming that the wheel speeds are measurable while the tire forces are not, yields the output matrix

$$C = \begin{bmatrix}I_4 & 0\end{bmatrix}, \tag{3.10}$$

where I is the identity matrix. Combining the driveline dynamics presented in Section 2.4 with the random walk tire force model gives the augmented estimator dynamics

$$\dot{\mathbf{z}} = \underbrace{\begin{bmatrix}0 & G\\ 0 & 0\end{bmatrix}}_{A_{\text{aug}}}\mathbf{z} + \begin{bmatrix}B\\ 0\end{bmatrix}\mathbf{u}, \qquad \mathbf{y} = C\mathbf{z}. \tag{3.11}$$

The corresponding observability matrix is given by

$$\mathcal{O} = \begin{bmatrix}C\\ CA_{\text{aug}}\\ CA_{\text{aug}}^2\\ \vdots\\ CA_{\text{aug}}^7\end{bmatrix}. \tag{3.12}$$

Since

$$CA_{\text{aug}}^i = 0 \quad \text{for } i \in \{2, 3, \ldots, 7\} \tag{3.13}$$

all higher-order terms vanish, reducing the observability matrix to

$$\mathcal{O} = \begin{bmatrix}C\\ CA_{\text{aug}}\end{bmatrix}. = \begin{bmatrix}I_4 & 0\\ 0 & G\end{bmatrix}. \tag{3.14}$$

Full observability of the augmented system requires

$$\operatorname{rank}(\mathcal{O}) = 8, \tag{3.15}$$

which in turn requires
$$\operatorname{rank}(G) = 4. \tag{3.16}$$

However, when the IA differential is locked, the matrix G attains rank 3, implying that one state combination becomes unobservable. By analyzing the null space of the observability matrix, it follows that vectors in the direction of

$$\begin{bmatrix} 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} & -\frac{1}{2} \end{bmatrix}^\top \tag{3.17}$$

span the null space. Consequently, the estimator cannot uniquely distinguish how the tire forces are distributed between the front and rear axles. In particular, the quantity
$$\frac{(F_1 + F_2) - (F_3 + F_4)}{2} \tag{3.18}$$
is unobservable from the available wheel speed measurements.

3.2.2.2 Transformed State Formulation

By using the transformation introduced in Section 2.4.3, the system can be described using,

$$F_{1234}^{\text{avg}} = \frac{F_1 + F_2 + F_3 + F_4}{4}, \tag{3.19a}$$
$$F_{12-34}^{\Delta} = \frac{(F_1 + F_2) - (F_3 + F_4)}{2}, \tag{3.19b}$$
$$F_{12}^{\Delta} = F_1 - F_2, \tag{3.19c}$$
$$F_{34}^{\Delta} = F_3 - F_4, \tag{3.19d}$$

and the unobservable state combination F_{12-34}^{Δ} can be removed explicitly from the model when the IA is locked.

The key property is that the difference-state dynamics do not depend on the unknown IA torque split. Subtracting the dynamic equations for wheels 1 and 2 gives

$$\begin{aligned} J_w\Big(\dot{\omega}_1 - \dot{\omega}_2\Big) &= (T_1 - T_{b1} - F_1 r) - (T_2 - T_{b2} - F_2 r) \\ &= -(T_{b1} - T_{b2}) - (F_1 - F_2)r, \end{aligned} \tag{3.20}$$

where T_1 and T_2 cancel since the IW differential remains open. The dynamics of ω_{12}^{Δ} are therefore well-defined regardless of the IA lock state, with an analogous cancellation holding for ω_{34}^{Δ}.

Instead of using a tire model, all tire force states are modeled as random walks with time derivatives
$$\dot{F}_{1234}^{\text{avg}} = \dot{F}_{12-34}^{\Delta} = \dot{F}_{12}^{\Delta} = \dot{F}_{34}^{\Delta} = 0. \tag{3.21}$$

The uncertainty in their evolution is absorbed into the process noise covariance Q. Their estimates are corrected implicitly through the measured wheel speed dynamics, since the effect of each force component on the measurable states is known from the driveline model.

3.2.2.3 All-Open Differential Configuration

For the all-open configuration the full set of transformed wheel speed states is used. The estimator state vector is

$$x_{\text{open}} = \begin{bmatrix}\omega_{1234}^{\text{avg}} & \omega_{12-34}^{\Delta} & \omega_{12}^{\Delta} & \omega_{34}^{\Delta} & v & F_{1234}^{\text{avg}} & F_{12-34}^{\Delta} & F_{12}^{\Delta} & F_{34}^{\Delta}\end{bmatrix}^{\top}, \tag{3.22}$$

where v is the vehicle longitudinal velocity. The control input vector is

$$u = \begin{bmatrix}T_e & T_{b1} & T_{b2} & T_{b3} & T_{b4} & \alpha_{\text{slope}}\end{bmatrix}^{\top}, \tag{3.23}$$

where the road gradient α_{slope} is included as a known input. Including v introduces the dynamic equation

$$m\dot{v} = 4F_{1234}^{\text{avg}} - mg\sin(\alpha_{\text{slope}}), \tag{3.24}$$

where resistive forces are neglected, linking the total tire force to a directly measurable quantity and thereby improving the force estimates. For small road inclinations, $\sin(\alpha_{\text{slope}}) \approx \alpha_{\text{slope}}$, which is applied here to keep the system linear. The state-space matrices are derived from the driveline dynamics of Section 2.4.3 together with the longitudinal vehicle dynamics, creating

$$A = \frac{\partial \dot{x}}{\partial x}, \quad B = \frac{\partial \dot{x}}{\partial u}. \tag{3.25}$$

The individual wheel speeds and vehicle velocity are measured. Expressing the individual wheel speeds in terms of the transformed states yields the measurement vector

$$\underbrace{\begin{bmatrix}\omega_1\\ \omega_2\\ \omega_3\\ \omega_4\\ v\end{bmatrix}}_{y} = \underbrace{\begin{bmatrix}1 & \frac{1}{2} & \frac{1}{2} & 0 & 0\\ 1 & \frac{1}{2} & -\frac{1}{2} & 0 & 0\\ 1 & -\frac{1}{2} & 0 & \frac{1}{2} & 0\\ 1 & -\frac{1}{2} & 0 & -\frac{1}{2} & 0\\ 0 & 0 & 0 & 0 & 1\end{bmatrix}}_{C} \underbrace{\begin{bmatrix}\omega_{1234}^{\text{avg}}\\ \omega_{12-34}^{\Delta}\\ \omega_{12}^{\Delta}\\ \omega_{34}^{\Delta}\\ v\end{bmatrix}}_{x}. \tag{3.26}$$

A steady-state filter gain L is computed offline from the CARE, as described in Section 3.2.1, using the derived matrices and the noise tuning parameters Q_{KF} and R_{KF}. The estimator and its integration with the controller in the all-open differential configuration is illustrated in Figure 3.2.

3.2.2.4 Locked IA Differential Configuration

When the IA differential is locked, $\omega_{12-34}^{\Delta} = 0$ and that state is removed from the estimator together with the corresponding tire force state F_{12-34}^{Δ}. The reduced state vector becomes

$$x_{\text{locked}} = \begin{bmatrix}\omega_{1234}^{\text{avg}} & \omega_{12}^{\Delta} & \omega_{34}^{\Delta} & v & F_{1234}^{\text{avg}} & F_{12}^{\Delta} & F_{34}^{\Delta}\end{bmatrix}^{\top}. \tag{3.27}$$

The inputs, disturbances, state-space matrix derivation, and filter gain computation follow the same procedure as for the all-open case. With $\omega_{12-34}^{\Delta} = 0$, the wheel speed expressions in (3.26) simplify and the measurement matrix C is updated accordingly.

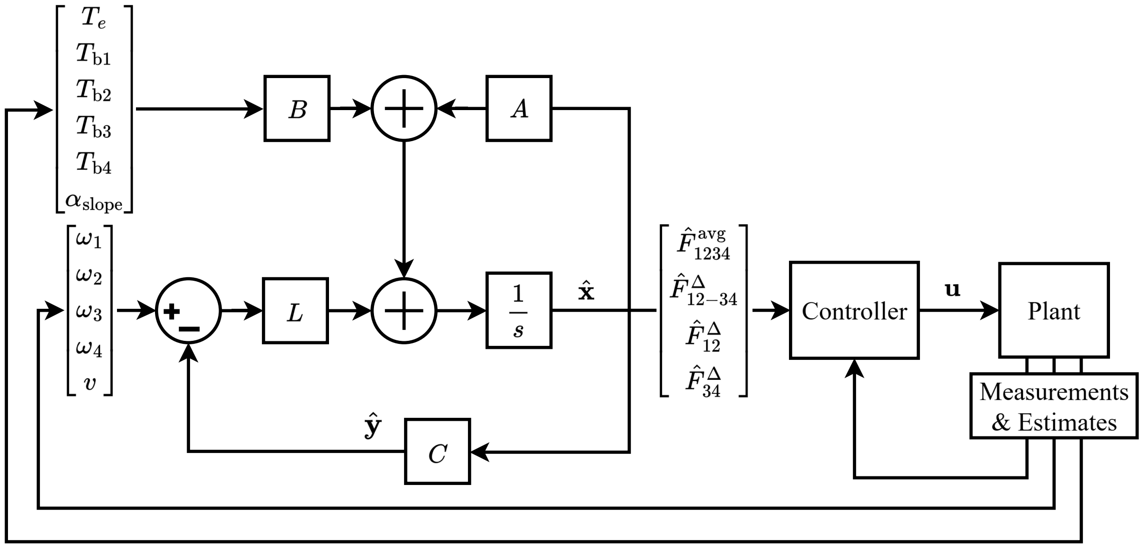

Figure 3.2: Block diagram of the tire force estimator operating together with a controller.

3.2.2.5 Configuration Switching

Separate state-space models and steady-state filter gains are derived for each differential configuration and switched between according to the current lock status. The transformed state representation makes this transition straightforward: When the IA differential opens, the state vector is expanded by appending ω_{12-34}^{Δ} and F_{12-34}^{Δ}, both initialized to zero, while all remaining states and covariances carry over directly from the locked-configuration estimator. This warm initialization avoids full re-convergence and reduces the time until the estimates become reliable.

3.3 Linear Quadratic Regulator

An LQR is an optimal state-feedback control method designed for linear dynamical systems. The main idea behind LQR is to determine a control law that minimizes a quadratic cost function, thereby balancing state-tracking performance and control effort by solving an optimization problem analytically.

One of the primary advantages of LQR is its systematic tuning procedure, in which the control behavior can be adjusted via weighting matrices that determine the relative importance of state deviations and control effort. Furthermore, the resulting controller is computationally efficient since the optimal feedback gain can be computed offline, making LQR attractive for real-time applications with limited computational resources.

Considering the generic LTI system

$$\dot{x} = Ax + Bu \tag{3.28}$$

with state vector $x \in \mathbb{R}^n$ and control vector $u \in \mathbb{R}^m$. The infinite-horizon LQR problem is stated as

$$\min_{u} \int_{0}^{\infty} \left(x^{\top} Q x + u^{\top} R u\right) dt, \qquad Q \succeq 0,\ R \succ 0. \tag{3.29}$$

If (A, B) is stabilizable and $(A, Q^{1/2})$ detectable, the optimal control that minimizes 3.29 is linear state feedback

$$u = -Kx, \qquad K = R^{-1} B^{\top} P, \tag{3.30}$$

where $P \succeq 0$ is the stabilizing solution

$$A^{\top} P + PA - PBR^{-1} B^{\top} P + Q = 0 \tag{3.31}$$

of the CARE [20].

The LQR is globally optimal for the given linear dynamics and quadratic cost (3.29). The closed-loop dynamics $\dot{x} = (A - BK)x$ are asymptotically stable when P is the stabilizing CARE solution.

While LQR offers several advantages, including computational efficiency, straightforward implementation, and systematic tuning via the weighting matrices, these properties are accompanied by inherent limitations. Most notably, standard LQR does not explicitly handle constraints on control inputs, system states, or their rates of change. If actuator saturation or operational constraints become active, the resulting optimal control policy is generally no longer linear and can therefore not be represented by the standard linear state-feedback formulation. Consequently, unlike optimization-based controllers such as MPC, standard LQR does not naturally redistribute or reprioritize control objectives when limitations are reached, since the feedback gain matrix K remains fixed regardless of the operating conditions.

3.3.1 LQR with Integral Action

Steady-state offsets arising from constant disturbances or reference signals can be eliminated by augmenting the LQR formulation with integral action on the output error. Consider the regulated output

$$y = Cx, \tag{3.32}$$

where $r \in \mathbb{R}^p$ denotes the reference signal. An integral state $\xi \in \mathbb{R}^p$ is introduced, governed by

$$\dot{\xi} = r - y = r - Cx. \tag{3.33}$$

Augmenting the original state vector with the integral state, $x_e = \begin{bmatrix} x^{\top} & \xi^{\top} \end{bmatrix}^{\top}$, yields the extended system dynamics

$$\dot{x}_e = \underbrace{\begin{bmatrix} A & 0 \\ -C & 0 \end{bmatrix}}_{A_e} x_e + \underbrace{\begin{bmatrix} B \\ 0 \end{bmatrix}}_{B_e} u + \underbrace{\begin{bmatrix} 0 \\ I \end{bmatrix}}_{E_e} r. \tag{3.34}$$

The standard LQR problem is then posed on the augmented system (A_e, B_e) with the block-diagonal weighting matrix

$$Q_e = \begin{bmatrix} Q & 0 \\ 0 & Q_\xi \end{bmatrix}, \qquad Q_\xi \succeq 0, \tag{3.35}$$

where Q_ξ penalizes the integrated output error. The resulting optimal state-feedback law takes the form

$$u = -K_{\mathrm{e}}x = -[K_{\mathrm{FB}}, \mathrm{K_I}]x \tag{3.36}$$

where the gain partitions K_{FB} and K_I handle the proportional and integral parts respectively.

3.3.2 LQR Implementation

A block diagram of the implemented LQR can be seen in Figure 3.3. The controller is formulated using the transformed wheel-speed states

$$\tilde{\boldsymbol{\omega}} = \begin{bmatrix} \omega_{1234}^{\mathrm{avg}} & \omega_{12-34}^{\Delta} & \omega_{12}^{\Delta} & \omega_{34}^{\Delta} \end{bmatrix}^\top. \tag{3.37}$$

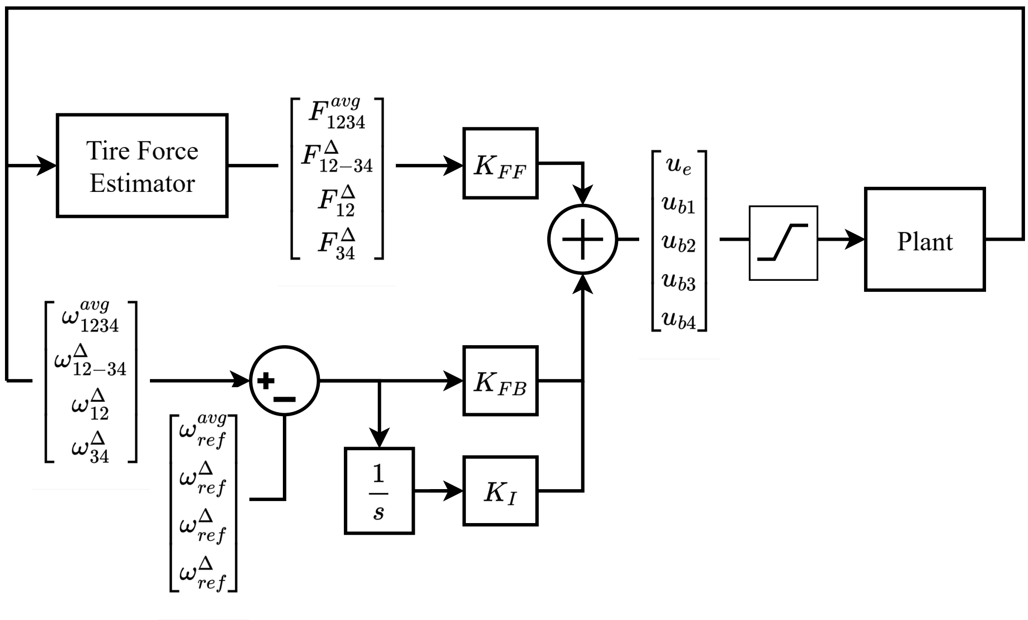

Figure 3.3: Block diagram of the implemented LQR.

The corresponding reference vector is defined as

$$\tilde{\boldsymbol{\omega}}_{\mathrm{ref}} = \begin{bmatrix} \omega_{\mathrm{ref}}^{\mathrm{avg}} & \omega_{\mathrm{ref}}^{\Delta} & \omega_{\mathrm{ref}}^{\Delta} & \omega_{\mathrm{ref}}^{\Delta} \end{bmatrix}^\top, \tag{3.38}$$

where the average wheel speed reference $\omega_{\mathrm{ref}}^{\mathrm{avg}}$ and the wheel speed-difference reference $\omega_{\mathrm{ref}}^{\Delta}$ were selected in accordance with Section 3.1.4.

The estimated tire force disturbance is defined as

$$\tilde{\mathbf{v}} = \begin{bmatrix} F_{1234}^{\mathrm{avg}} & F_{12-34}^{\Delta} & F_{12}^{\Delta} & F_{34}^{\Delta} \end{bmatrix}^\top, \tag{3.39}$$

where each component represents transformed tire forces, as described in Section 3.2.

The continuous-time system dynamics are expressed as

$$\dot{\tilde{\boldsymbol{\omega}}} = \tilde{A}\tilde{\boldsymbol{\omega}} + \tilde{B}\mathbf{u} + \tilde{G}\mathbf{v}, \tag{3.40}$$

where $\tilde{A}$ is the null matrix for the considered system. The control input vector is defined as

$$\mathbf{u} = \begin{bmatrix} u_e & u_{b1} & u_{b2} & u_{b3} & u_{b4} \end{bmatrix}^\top, \tag{3.41}$$

where u_e denotes the engine torque and u_{bi} the braking torque applied at wheel i.

To eliminate steady-state tracking errors, integral action was incorporated into the controller design. The system was therefore augmented with integral states corresponding to the accumulated tracking error

$$\dot{\tilde{\boldsymbol{\omega}}}_I = \tilde{\boldsymbol{\omega}} - \tilde{\boldsymbol{\omega}}_{\text{ref}}, \tag{3.42}$$

where $\tilde{\boldsymbol{\omega}}_I$ is the integral error states. The resulting augmented state vector becomes

$$\tilde{\boldsymbol{\omega}}_e = \begin{bmatrix} \tilde{\boldsymbol{\omega}} \\ \tilde{\boldsymbol{\omega}}_I \end{bmatrix}. \tag{3.43}$$

Using this augmented formulation, the system dynamics can be written as

$$\dot{\tilde{\boldsymbol{\omega}}}_e = A_e\tilde{\boldsymbol{\omega}}_e + B_e\mathbf{u}, \tag{3.44}$$

where the augmented system matrices are given by

$$A_e = \begin{bmatrix} \tilde{A} & 0 \\ -I & 0 \end{bmatrix}, \tag{3.45a}$$

$$B_e = \begin{bmatrix} \tilde{B} \\ 0 \end{bmatrix}. \tag{3.45b}$$

The LQR cost function becomes

$$J = \int_0^\infty \left(\tilde{\boldsymbol{\omega}}_e^{\mathrm{T}} Q_{\mathrm{LQR}} \tilde{\boldsymbol{\omega}}_e + \mathbf{u}^{\mathrm{T}} R_{\mathrm{LQR}} \mathbf{u} \right) dt, \tag{3.46}$$

where the weighting matrix Q_{LQR} penalizes both wheel speed tracking errors and accumulated integral errors, while R_{LQR} penalizes the control effort.

The optimal feedback gain matrix was obtained by solving the CARE using MATLAB's `icare` solver, after which the optimal state-feedback gain matrix was obtained with

$$K = [K_{\mathrm{FB}}, K_{\mathrm{I}}] = R_{\mathrm{LQR}}^{-1} B_e^{\mathrm{T}} P. \tag{3.47}$$

In addition to the feedback, a feedforward term was introduced to compensate for the effects of tire forces. The feedforward gain matrix was computed as

$$K_{\mathrm{FF}} = \tilde{B}^\dagger \tilde{G}, \tag{3.48}$$

where $(\cdot)^\dagger$ denotes the pseudo-inverse. The resulting control law was implemented as

$$\mathbf{u} = -K\tilde{\boldsymbol{\omega}}_e + K_{\mathrm{FF}}\tilde{\mathbf{v}}. \tag{3.49}$$

Before being sent as actuator torque requests, the control inputs were saturated according to

$$\boldsymbol{u}_{\min} \leq \boldsymbol{u} \leq \boldsymbol{u}_{\max} \tag{3.50}$$

3.4 Model Predictive Control

MPC is an optimization-based control strategy that utilizes a dynamic model of the system to predict its future behavior over a finite prediction horizon [21]. At each sampling instant, MPC computes the control inputs by solving an optimization problem that minimizes a predefined cost function while simultaneously satisfying system and actuator constraints. Only the first control input of the optimized sequence is applied to the system, after which the optimization is repeated at the next sampling instant using updated state measurements.

Consider a generic discrete-time system described by

$$x_{k+1} = A_d\mathbf{x}_k + B_d\mathbf{u}_k, \tag{3.51}$$

where $\mathbf{x}_k \in \mathbb{R}^n$ denotes the system state vector and $\mathbf{u}_k \in \mathbb{R}^m$ represents the control input vector. The matrices A_d and B_d describe the discrete-time system dynamics.

Using the system model, MPC predicts the future state evolution over a prediction horizon of length N. The control objective is typically formulated through a quadratic cost function of the form

$$J = \sum_{i=0}^{N-1} \left(\|\mathbf{x}_{k+i} - \mathbf{x}_{\mathrm{ref}}\|_Q^2 + \|\mathbf{u}_{k+i}\|_R^2 \right), \qquad Q \succeq 0,\ R \succ 0. \tag{3.52}$$

where x_{ref} is a reference trajectory, while Q and R are weighting matrices penalizing tracking error and control effort, respectively. The weighting matrices determine the trade-off between tracking performance and actuator usage.

One of the primary advantages of MPC compared to the previously considered LQR formulation is its ability to explicitly incorporate system constraints into the controller design. Constraints on actuator limits, actuator rate changes, and state variables can be included directly in the optimization problem, ensuring that the computed control actions remain physically feasible and safe.

Similar to LQR, the weighting matrices allow different system states and control objectives to be prioritized relative to one another. Retaining such a model-based formulation is desirable for the application, since the wheel torques simultaneously influence both the average wheel speeds and the relative wheel speed differences between wheels and axles, resulting in coupled multivariable dynamics.

Justified by these properties, the MPC was selected for further investigation.

3.4.1 MPC Design

Although MPC is fundamentally based on predicting the future system evolution over a finite prediction horizon, the predictive capability itself was of limited importance for the control problem considered in this work. The reference signals and disturbance estimates were assumed constant over the prediction horizon, and the system model contains no internal state propagation (continuous state matrix $\tilde{A} = 0$). Consequently, extending the prediction horizon provided little additional control performance while significantly increasing the computational complexity of the optimization problem, which is undesirable for real-time implementation.

Motivated by these observations, together with the computational limitations associated with real-time implementation, the prediction horizon was selected as

$$N_{\mathrm{MPC}} = 1. \tag{3.53}$$

With this choice, the MPC formulation effectively reduces to a single-step constrained optimization problem, where the primary benefits stem from the weighting of competing objectives and the explicit handling of actuator constraints, rather than from long-term trajectory prediction.

The state, reference, and disturbance vectors are defined analogously to those used in the LQR formulation. To enable real-time implementation of the MPC, the continuous-time systems, presented in Section 2.4.3 are discretized using a zero-order hold assumption over the controller sampling time T_s. The equivalent discrete-time system is obtained as

$$\tilde{\boldsymbol{\omega}}_{k+1} = \tilde{A}_d \tilde{\boldsymbol{\omega}}_k + \tilde{B}_d \boldsymbol{u}_k + \tilde{G}_d \tilde{\boldsymbol{v}}_k \tag{3.54}$$

with the discrete-time matrices being defined by

$$\tilde{A}_d = e^{\tilde{A}T_s}, \tag{3.55a}$$

$$\tilde{B}_d = \int_0^{T_s} e^{\tilde{A}\tau} \tilde{B}\, d\tau, \tag{3.55b}$$

$$\tilde{G}_d = \int_0^{T_s} e^{\tilde{A}\tau} \tilde{G}\, d\tau. \tag{3.55c}$$

For the one-step horizon considered in this thesis, the quadratic cost function is formulated as

$$J = (\tilde{\boldsymbol{\omega}}_{k+1} - \tilde{\boldsymbol{\omega}}_{\mathrm{ref}})^{\mathrm{T}} Q_{\mathrm{MPC}} (\tilde{\boldsymbol{\omega}}_{k+1} - \tilde{\boldsymbol{\omega}}_{\mathrm{ref}}) + \mathbf{u}_k^{\mathrm{T}} R_{\mathrm{MPC}} \mathbf{u}_k, \tag{3.56}$$

where Q_{MPC} is the state tracking weight matrix and R_{MPC} is the control effort weight matrix. Substituting the predicted state, using (3.54), into the cost function yields

$$\begin{aligned} J = {} & (\tilde{A}\tilde{\boldsymbol{\omega}}_k + \tilde{B}\mathbf{u}_k + \tilde{G}\tilde{\mathbf{v}}_k - \tilde{\boldsymbol{\omega}}_{\mathrm{ref}})^{\mathrm{T}} Q_{\mathrm{MPC}} (\tilde{A}\tilde{\boldsymbol{\omega}}_k + \tilde{B}\mathbf{u}_k + \tilde{G}\tilde{\mathbf{v}}_k - \tilde{\boldsymbol{\omega}}_{\mathrm{ref}}) \\ & + \mathbf{u}_k^{\mathrm{T}} R_{\mathrm{MPC}} \mathbf{u}_k. \end{aligned} \tag{3.57}$$

3.4.2 MPC Implementation

A block diagram of the implemented MPC can be seen in Figure 3.4. The optimization problem can be written in the quadratic programming form required by the `quadprog` solver in MATLAB,

$$\min_{\mathbf{u}_k} \quad \frac{1}{2}\mathbf{u}_k^{\mathrm{T}} H \mathbf{u}_k + f^{\mathrm{T}}\mathbf{u}_k, \tag{3.58}$$

where the Hessian matrix H and gradient vector f are given by

$$H = \tilde{B}^{\mathrm{T}} Q_{\mathrm{MPC}} \tilde{B} + R_{\mathrm{MPC}} \tag{3.59a}$$

$$f = \tilde{B}^{\mathrm{T}} Q_{\mathrm{MPC}} (\tilde{A}\tilde{\boldsymbol{\omega}}_k + \tilde{G}\mathbf{v}_k - \tilde{\boldsymbol{\omega}}_{\mathrm{ref}}) \tag{3.59b}$$

and the optimization variables correspond to the actuator torque commands

$$\mathbf{u}_k = \begin{bmatrix} u_e & u_{b1} & u_{b2} & u_{b3} & u_{b4} \end{bmatrix}^{\mathrm{T}}. \tag{3.60}$$

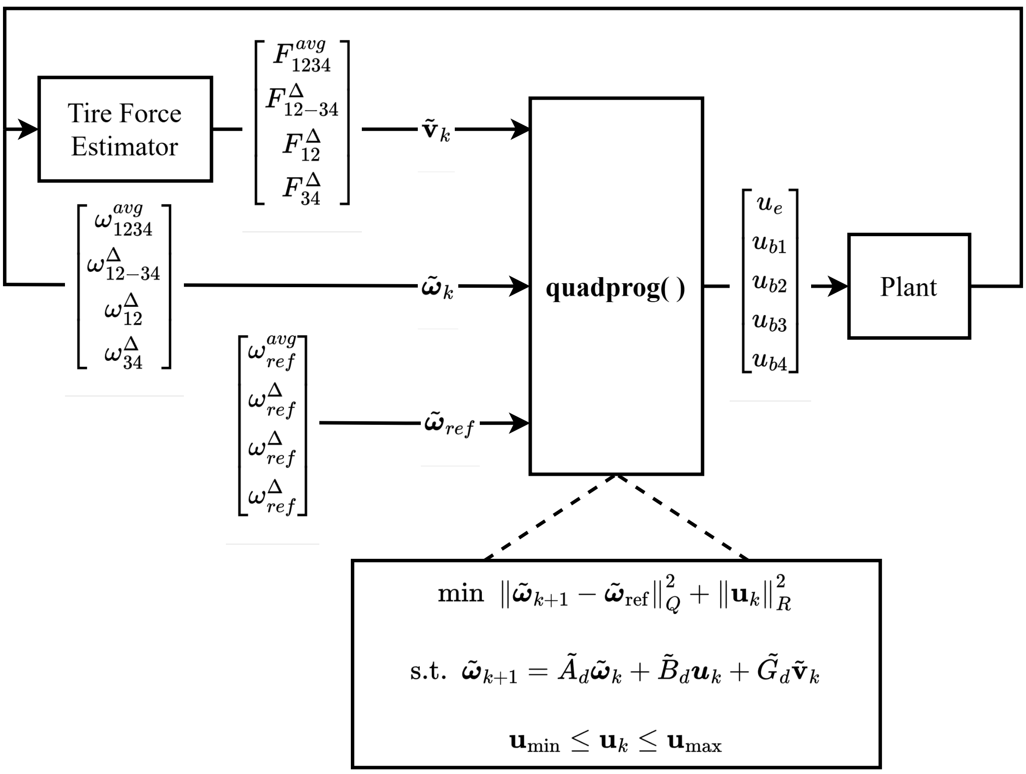

Figure 3.4: Block diagram of the implemented MPC.

Finally, the optimization problem is solved subject to actuator constraints

$$\mathbf{u}_{\min} \leq \mathbf{u}_k \leq \mathbf{u}_{\max}, \tag{3.61}$$

ensuring that the computed torque commands remain physically feasible.

3.5 Sliding Mode Control

SMC is a nonlinear control method known for its robustness against model uncertainties and external disturbances. The main idea behind SMC is to force the system trajectories onto a predefined manifold, referred to as the sliding surface, and to then maintain the motion on this surface. Once the system reaches the sliding surface, the closed-loop dynamics become insensitive to certain parameter variations and disturbances.

Consider a general nonlinear system described by

$$\dot{x} = f(x) + g(x)u, \tag{3.62}$$

where $x \in \mathbb{R}^n$ denotes the system states, u is the control input, and $f(\cdot)$ and $g(\cdot)$ represent the system dynamics.

A sliding surface is introduced as

$$s(x) = 0, \tag{3.63}$$

where s is typically designed such that the desired closed-loop dynamics are obtained when the system evolves on the surface. The control objective is therefore often divided into two phases: The reaching phase, during which the trajectories are driven toward the sliding surface, and the sliding phase, where the trajectories remain on to the surface.

To ensure convergence toward the sliding surface, the reaching condition

$$s \cdot \dot{s} < 0 \tag{3.64}$$

is commonly imposed. This condition guarantees that the distance to the sliding surface decreases over time. A common SMC control law can be expressed as

$$u = u_{\text{eq}} - K \operatorname{sign}(s), \tag{3.65}$$

where u_{eq} is the equivalent control responsible for maintaining motion on the sliding surface, $K > 0$ is a design parameter, and $\operatorname{sign}(s)$ is the sign function introducing the discontinuous switching action characteristic of SMC.

A common issue with SMC is that the discontinuous switching may lead to high-frequency oscillations in the control signal, commonly referred to as chattering. In practical implementations, this issue is often mitigated by replacing the sign function with a saturation function

$$\operatorname{sat}(s) = \begin{cases} 1, & s > \phi, \\ \dfrac{s}{\phi}, & |s| \leq \phi, \\ -1, & s < -\phi \end{cases} \tag{3.66}$$

where $\phi > 0$ is the boundary-layer thickness, yielding the control law

$$u = u_{\text{eq}} - K \operatorname{sat}(s). \tag{3.67}$$

SMC has been previously applied in automotive control applications such as ABS [5] and ASR [6]. A key reason is that the dominant disturbances in these problems, tire force induced torques, enter the wheel dynamics through the same channel as the applied brake or engine torque. Such disturbances are therefore matched with respect to the control input, a condition under which SMC can compensate for them effectively.

The control problem considered in this work exhibits similar structural properties. Motivated by these similarities and the robustness characteristics of SMC, the method was selected for further investigation.

3.5.1 SMC Design

The sliding surface was defined based on the wheel speed states introduced previously. The objective was to regulate both the average wheel speed and the relative wheel speed differences using a common surface formulation. The resulting sliding surface was chosen as

$$\mathbf{s} = \begin{bmatrix} \omega_{1234}^{\text{avg}} - \omega_{\text{ref}}^{\text{avg}} \\ \omega_{12-34}^{\Delta} - \omega_{\text{ref}}^{\Delta} \\ \omega_{12}^{\Delta} - \omega_{\text{ref}}^{\Delta} \\ \omega_{34}^{\Delta} - \omega_{\text{ref}}^{\Delta} \end{bmatrix}, \tag{3.68}$$

where the first component represents the tracking error of the average wheel speed, while the remaining components represent the wheel speed difference errors with respect to their corresponding references.

As described in the previous section, the SMC design requires the time derivative of the sliding surface, denoted $\dot{\mathbf{s}}$. Since the sliding variables are constructed from both the transformed wheel speed states and their corresponding references, the surface derivative is obtained by directly differentiating the surface definition with respect to time. This yields

$$\dot{\mathbf{s}} = \begin{bmatrix} \dot{\omega}_{1234}^{\text{avg}} - \dot{\omega}_{\text{ref}}^{\text{avg}} \\ \dot{\omega}_{12-34}^{\Delta} - \dot{\omega}_{\text{ref}}^{\Delta} \\ \dot{\omega}_{12}^{\Delta} - \dot{\omega}_{\text{ref}}^{\Delta} \\ \dot{\omega}_{34}^{\Delta} - \dot{\omega}_{\text{ref}}^{\Delta} \end{bmatrix}. \tag{3.69}$$

The wheel speed derivatives are obtained directly from the transformed state-space formulation presented in Section 2.4.3. This relation couples the sliding surface dynamics to the system model according to

$$\dot{\mathbf{s}} = \tilde{B}\mathbf{u} + \tilde{G}\tilde{\mathbf{v}} - \begin{bmatrix} \dot{\omega}_{\text{ref}}^{\text{avg}} \\ \dot{\omega}_{\text{ref}}^{\Delta} \\ \dot{\omega}_{\text{ref}}^{\Delta} \\ \dot{\omega}_{\text{ref}}^{\Delta} \end{bmatrix}. \tag{3.70}$$

The average wheel speed reference was defined from the longitudinal slip relationship as

$$\omega_{\text{ref}}^{\text{avg}} = \frac{(1+\kappa_{\text{ref}})v}{r}, \tag{3.71}$$

where κ_{ref} denotes the reference slip ratio, v the vehicle longitudinal velocity, and r the effective wheel radius. Assuming a constant reference slip ratio yields the reference derivative

$$\dot{\omega}_{\text{ref}}^{\text{avg}} = \frac{(1+\kappa_{\text{ref}})\dot{v}}{r}, \tag{3.72}$$

where $\dot{v}$ is the longitudinal vehicle acceleration. Furthermore, the remaining reference terms were selected as constant values, resulting in

$$\dot{\omega}_{\text{ref}}^{\Delta} = 0. \tag{3.73}$$

The sliding surface dynamics can therefore be written as

$$\dot{\mathbf{s}} = \tilde{B}\mathbf{u} + \tilde{G}\tilde{\mathbf{v}} - \begin{bmatrix} \frac{(1+\kappa_{\text{ref}})\dot{v}}{r} \\ 0 \\ 0 \\ 0 \end{bmatrix}. \tag{3.74}$$

Having obtained the sliding surface dynamics, the control input can now be formulated such that the system trajectories are driven toward the sliding surface, i.e., $\mathbf{s} = 0$. As discussed previously, convergence toward the surface is guaranteed if the reaching condition (3.64) is satisfied. A straightforward way of enforcing this condition is to impose the sliding dynamics

$$\dot{\mathbf{s}} = -K_{\text{SMC}}\,\text{sat}(\mathbf{s}), \tag{3.75}$$

where K_{SMC} is a positive diagonal gain matrix and $\text{sat}(\cdot)$ denotes the element-wise saturation function. This formulation ensures that each sliding variable is driven toward zero with a direction determined by the sign of the corresponding surface error.

3.5.2 SMC Implementation

A block diagram of the implemented SMC can be seen in Figure 3.5. In the proposed implementation, the desired sliding dynamics in (3.75) were not enforced analytically through an explicit control law. Instead, a constrained optimization problem is formulated, solved at each sampling instant using quadratic programming via MATLAB's `quadprog` solver. The optimization objective was defined such that the resulting surface dynamics followed the desired reaching law as closely as possible in a weighted least-squares sense.

By rewriting the sliding dynamics (3.75) into a residual that should be minimized

$$\mathbf{e} = \dot{\mathbf{s}} + K_{\text{SMC}}\,\text{sat}(\mathbf{s}) = \tilde{B}\mathbf{u} + \tilde{G}\tilde{\mathbf{v}} - \begin{bmatrix} \frac{(1+\kappa_{\text{ref}})\dot{v}}{r} \\ 0 \\ 0 \\ 0 \end{bmatrix} + K_{\text{SMC}}\,\text{sat}(\mathbf{s}). \tag{3.76}$$

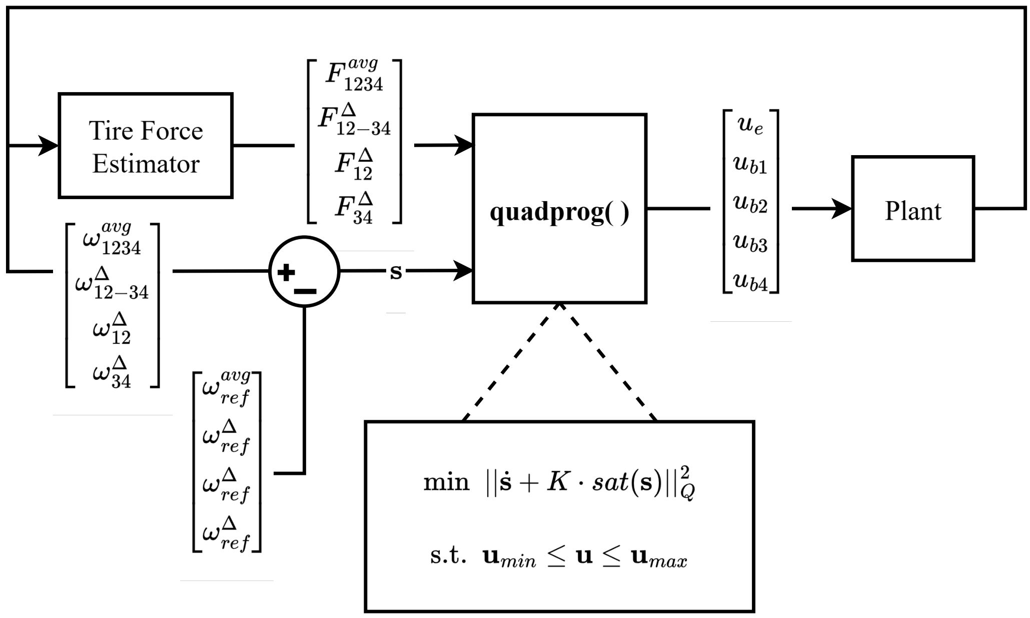

Figure 3.5: Block diagram of the implemented SMC.

with control input vector

$$\mathbf{u} = \begin{bmatrix} u_e & u_{b1} & u_{b2} & u_{b3} & u_{b4} \end{bmatrix}^{\mathrm{T}}. \tag{3.77}$$

The optimization problem can be formulated as

$$\min_{\mathbf{u}} \quad \mathbf{e}^{\mathrm{T}} Q_{\mathrm{SMC}} \mathbf{e}, \tag{3.78}$$

where $Q_{\mathrm{SMC}} \succeq 0$ is a weighting matrix used to prioritize the different sliding surface components.

To ensure physically reachable torque requests, inequality constraints

$$\mathbf{u}_{\min} \leq \mathbf{u} \leq \mathbf{u}_{\max}, \tag{3.79}$$

were imposed on all control inputs.

It should be noted that the weighting matrix Q_{SMC} only influences the solution when the actuator constraints are active. When one or more inputs reach their physical or safety-related limits, the optimization problem becomes over-constrained and the desired sliding dynamics can no longer be fulfilled perfectly for all sliding variables simultaneously. In these situations, Q_{SMC} determines how the remaining control authority is distributed among the sliding surface components. For instance, assigning a larger weight to $Q_{\mathrm{SMC}}(2,2)$ prioritizes synchronization of the IA speed difference when the control inputs become constrained.

This formulation also enables additional operational constraints to be incorporated seamlessly into the controller. For example, the upper bound on the propulsion torque can be limited by the driver's requested torque demand, ensuring that the controller never commands more propulsion torque than requested while still maintaining stable and feasible control behavior. Furthermore, actuator rate limits could also be incorporated within the same optimization framework if required.

4

Simulation Setup

This chapter describes the simulation environment and test setup used to evaluate the controllers, covering the simulation software, test scenarios, performance metrics, and tuning parameters.

4.1 Simulation Software

The simulation tests were conducted using the ASM vehicle dynamics model software, developed by dSPACE. A 6×4 truck model developed by Volvo was used within this environment. The model captures longitudinal, lateral, and vertical vehicle dynamics, and includes engine, suspension, and brake models as well as a tire model. The driveline includes two IW differentials and one IA differential, all of which are lockable. The differential lock clutch is modeled as a friction clutch with high torque transfer capability rather than as a dog clutch, as discussed in Section 2.2.

4.1.1 Test Automation

To enable testing of a large number of controller tunings across multiple scenarios, an in-house developed Python-based scripting framework was used. While the simulation still ran in real time, the framework automated starting and stopping runs, resetting all simulation variables between tests, and saving data for post-processing. This ensured that all tests were run under identical conditions and made the volume of testing required practically feasible.

4.2 Scenario Modeling

All tests are conducted in simulation on straight roads. Since the simulations involve lateral dynamics, a built-in driver model is used to keep the vehicle on a straight path, serving as a simple feedback controller for the vehicle's lateral position.

The gear is set manually for each scenario rather than relying on the automatic transmission, preventing gear shifts mid-test and ensuring sufficient engine torque response throughout. The gear is selected based on the road incline and vehicle speed of each scenario.

The vehicle is initialized on a high-friction surface before reaching a low-friction road section, ensuring all vehicle states are stable at the start of each test.

4.2.1 Road Conditions

The road conditions are among the most significant external factors affecting both the difficulty and realism of the simulation scenarios. Consequently, representative road models are essential for obtaining meaningful evaluations of controllers. Two different road models were used in the simulations presented in this thesis.

The first road model consists of a split-friction surface with a longitudinal slope of 10 %, where the left side of the road has a significantly lower friction coefficient than the right side, as illustrated in Figure 4.1a. Such conditions are representative of scenarios that could be encountered during winter driving, where one side of the vehicle may travel on ice or packed snow while the other side remains on a higher-friction surface.

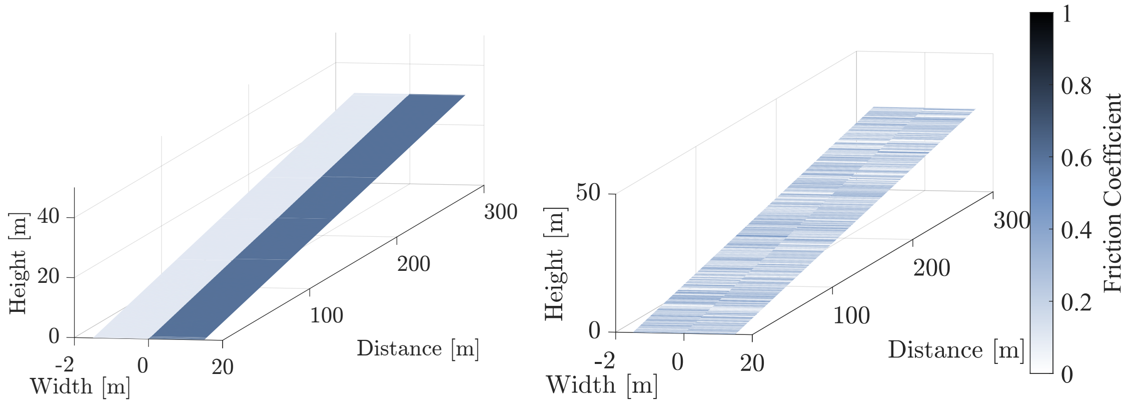

(a) Split-friction road. **(b)** Gravel road.

Figure 4.1: Road surface friction models used in the simulations.

Besides representing a practically relevant operating condition, the split-friction road mainly serves as a controlled test scenario for the controllers. Due to the abrupt but constant friction difference, the scenario behaves similarly to a step response, allowing the controller behavior and synchronization performance to be analyzed in a systematic and repeatable manner. The driver aims to enter the split-friction section at 40 km/h.

The second road model was designed to emulate highly varying traction conditions typically encountered on gravel roads. The same longitudinal road slope was used as in the split-friction scenario. Of particular importance for the control problem are the rapid and localized reductions in available friction caused by loose gravel patches, as well as transients in normal load caused by small bumps, potholes, or uneven terrain. However, directly modeling detailed road height variations was found to introduce simulation instability. Instead, a custom friction map was constructed to reproduce the primary traction variations associated with gravel driving. The longitudinal friction profile $\mu(d)$ was generated as a smoothed normal distribution

$$\mu_{\text{raw}}(d) \sim \mathcal{N}(\mu, \sigma^2), \tag{4.1a}$$

$$\mu_{\text{smooth}}(d) = \texttt{movmean}(\mu_{\text{raw}}(d),\, n_{\text{avg}})\,, \tag{4.1b}$$

$$\mu(d) = \max(\mu_{\text{smooth}}(d),\, 0)\,, \tag{4.1c}$$

$$d \in \{0,\, \Delta d,\, 2\Delta d,\, \ldots,\, L\}, \tag{4.1d}$$

where $\mu = 0.20$ is the mean friction coefficient and $\sigma = 0.30$ its standard deviation, $n_{\text{avg}} = 10$ the moving-average window length used to smooth the profile, $\Delta d = 0.1\,\text{m}$ is the sampling distance and $L = 300\,\text{m}$ is the total road length. The resulting gravel road is illustrated in Figure 4.1b.

Compared to the split-friction scenario, the gravel road model introduces significantly more disturbances and continuously varying traction conditions, thereby providing a more realistic and demanding test environment for evaluating controller robustness and realistic synchronization performance. The driver aims to enter the gravel section at 15 km/h.

4.3 Performance Metrics

To evaluate and compare the controllers, a set of performance metrics are defined. Each metric is evaluated over the synchronization window, spanning from the time the controller is activated, t_{start}, to the time at which all differentials have successfully locked, t_{end}.

4.3.1 Synchronization Time

The synchronization time

$$t^{\Delta} = t_{\text{end}} - t_{\text{start}}, \tag{4.2}$$

is the duration of the synchronization. It is one of the most direct measures of controller performance, as a shorter synchronization time means that the differential locks can be engaged more quickly.

4.3.2 Velocity Loss

The velocity loss v^{Δ} is the change in longitudinal vehicle velocity over the synchronization window,

$$v^{\Delta} = v(t_{\text{start}}) - v(t_{\text{end}}). \tag{4.3}$$

A positive value indicates a loss of speed. Minimizing v^{Δ} is important for reducing the risk of the truck getting stuck, and is a measure of how well the controller's slip regulation works.

4.3.3 Driver Disturbance

Driver comfort during synchronization is assessed using Root Mean Square (RMS) jerk, defined as

$$j^{\mathrm{RMS}} = \sqrt{\frac{1}{t^{\Delta}} \int_{t_{\mathrm{start}}}^{t_{\mathrm{end}}} \dddot{v}^2 \, dt}, \tag{4.4}$$

where $\dddot{v}$ is the time derivative of the longitudinal acceleration, i.e., the jerk. It captures the smoothness of the vehicle's response during synchronization and is a common metric for evaluating driver comfort [22].

4.3.4 Control Effort

The control effort metrics are a measure of the total energy consumed by the actuators over the synchronization window. The total brake impulse is

$$u_b^{\mathrm{tot}} = \sum_{i=1}^{4} \int_{t_{\mathrm{start}}}^{t_{\mathrm{end}}} |T_{bi}| \, dt, \tag{4.5}$$

and the total engine impulse is

$$u_e^{\mathrm{tot}} = \int_{t_{\mathrm{start}}}^{t_{\mathrm{end}}} |T_e| \, dt. \tag{4.6}$$

Lower brake usage corresponds to reduced particle emissions and lower maintenance costs, while lower engine usage corresponds to reduced fuel consumption.

4.4 Tuning Parameters

This section presents the tuning parameters for the tire force estimator and the three controllers. For each controller, certain parameters are fixed while others are swept across a range of values. Sweeping parameters reduces the influence of any single tuning on the comparison and allows the trade-offs between competing performance objectives to be mapped out. The total number of parameter configurations grows as

$$N = n^p, \tag{4.7}$$

where n is the number of values per parameter and p is the number of swept parameters. To keep the number of simulations manageable, $p = 3$ and $n = 5$ are used for all three controllers, yielding $5^3 = 125$ different tunings for each controller.

4.4.1 Tire Force Estimator Tuning

Both Q_{KF} and R_{KF} are diagonal, meaning no cross-correlations between states or measurements are assumed. The tuning values are given in Table 4.1. These values apply to both the open and locked IA configurations. In the locked case, the rows corresponding to ω_{12-34}^{Δ} and F_{12-34}^{Δ} are not used, since those states are removed when the IA differential is locked.

Table 4.1: Tire force estimator tuning parameters.

(a) Q_{KF}.

State	Variance
$\omega_{1234}^{\text{avg}}$	10^0
ω_{12-34}^{Δ}	10^0
ω_{12}^{Δ}	10^0
ω_{34}^{Δ}	10^{-1}
v	10^{-1}
F_{1234}^{avg}	10^6
F_{12-34}^{Δ}	10^8
F_{12}^{Δ}	10^8
F_{34}^{Δ}	10^8

(b) R_{KF}.

State	Variance
ω_1	10^{-1}
ω_2	10^{-1}
ω_3	10^{-1}
ω_4	10^{-1}
v	10^{-3}

4.4.2 LQR Tuning

Both Q_{LQR} and R_{LQR} are diagonal. The control penalty matrix R_{LQR} is fixed. The state penalty matrix Q_{LQR} contains both swept and fixed entries: the state deviation weights q_1, q_2, and q_3 are swept logarithmically, meaning the values are evenly spaced in log-scale between the minimum and maximum values, while the integral action weights q_4, q_5, and q_6 are fixed. All parameters are given in Table 4.2.

The weights are defined as follows: q_1 penalizes reference deviation of $\omega_{1234}^{\text{avg}}$, q_2 the primary differential state, and q_3 the secondary differential states. During IA synchronization, q_2 is mapped to ω_{12-34}^{Δ} and q_3 to ω_{12}^{Δ} and ω_{34}^{Δ}. Once the IA locks, q_2 is remapped to ω_{12}^{Δ} and ω_{34}^{Δ}. The integral weights q_4, q_5, and q_6 follow the same mapping as q_1, q_2, and q_3 respectively.

Table 4.2: LQR tuning parameters.

(a) Q_{LQR} (swept).

Weight	Min	Max
q_1	10^4	10^8
q_2	10^4	10^8
q_3	10^4	10^8

(b) Q_{LQR} (fixed).

Weight	Value
q_4	10^4
q_5	10^7
q_6	10^4

(c) R_{LQR} (fixed).

Weight	Value
u_e	1
u_{b1}	1
u_{b2}	1
u_{b3}	1
u_{b4}	1

4.4.3 MPC Tuning

The MPC shares the same cost matrix structure as the LQR, but without integral action. R_{MPC} uses the same values as R_{LQR}, and Q_{MPC} is swept over the same logarithmic grid with the same weight mapping as q_1, q_2, and q_3 in Q_{LQR}. The prediction horizon is fixed at $N_{\text{MPC}} = 1$, for the reasons described in Section 3.4.1.

4.4.4 SMC Tuning

The sliding-mode gain matrix K_{SMC} is diagonal. Three parameters are swept: k_1 penalizing reference deviation of $\omega_{1234}^{\text{avg}}$, k_2 the primary differential state, and k_3 the secondary differential states. They are linearly swept, meaning the values are evenly spaced between the minimum and maximum. The weighting matrix Q_{SMC} and boundary-layer thickness $\phi = 0.25$ are fixed. The gains k_1, k_2, k_3, and the weights in Q_{SMC} follow the same state mapping as described for the LQR. All parameters are given in Table 4.3.

Table 4.3: SMC tuning parameters.

(a) K_{SMC} (swept).

Weight	Min	Max
k_1	5	50
k_2	5	50
k_3	5	50

(b) Q_{SMC} (fixed).

Weight	Value
q_1	10^0
q_2	10^3
q_3	10^0

5 Simulation Results and Analysis

This chapter presents the simulation results for the three controllers developed in Chapter 3 to assess which control strategy best meets the objectives outlined in Section 1.1. All tests are conducted in the ASM simulation environment described in Section 4.1.

The chapter begins by evaluating the tire force estimator. Each controller is then evaluated individually using a representative tuning, meaning a tuning that produces typical and well-functioning behavior without being optimized for any single performance metric. The purpose of these sections is to illustrate how each controller operates and how the control signals differ between them, rather than to make a direct comparison. The controllers are then compared against each other across both road scenarios, using results from multiple tunings to map out the trade-offs among performance metrics. The chapter concludes with a specific investigation into the effect of engine torque control on performance.

Note that, for confidentiality reasons, the numerical values on all axes throughout this chapter have been scaled and therefore do not correspond to the true values. The qualitative behavior, trends, and relative comparisons between controllers remain fully representative of the underlying results.

5.1 Tire Force Estimator

The tire force estimator developed in Section 3.2 is evaluated over a full synchronization sequence on the split-friction uphill road, covering both the open and locked IA differential configurations. The estimated tire forces are compared against the true values taken directly from the simulation model. The results are shown in Figure 5.1, where the forces are plotted in the transformed representation. The three vertical dashed lines mark key events in the synchronization sequence: the first marks the start of IA synchronization, the second marks the IA lock and the start of IW synchronization, and the third marks when the last of the two IW differentials has locked and synchronization ends.

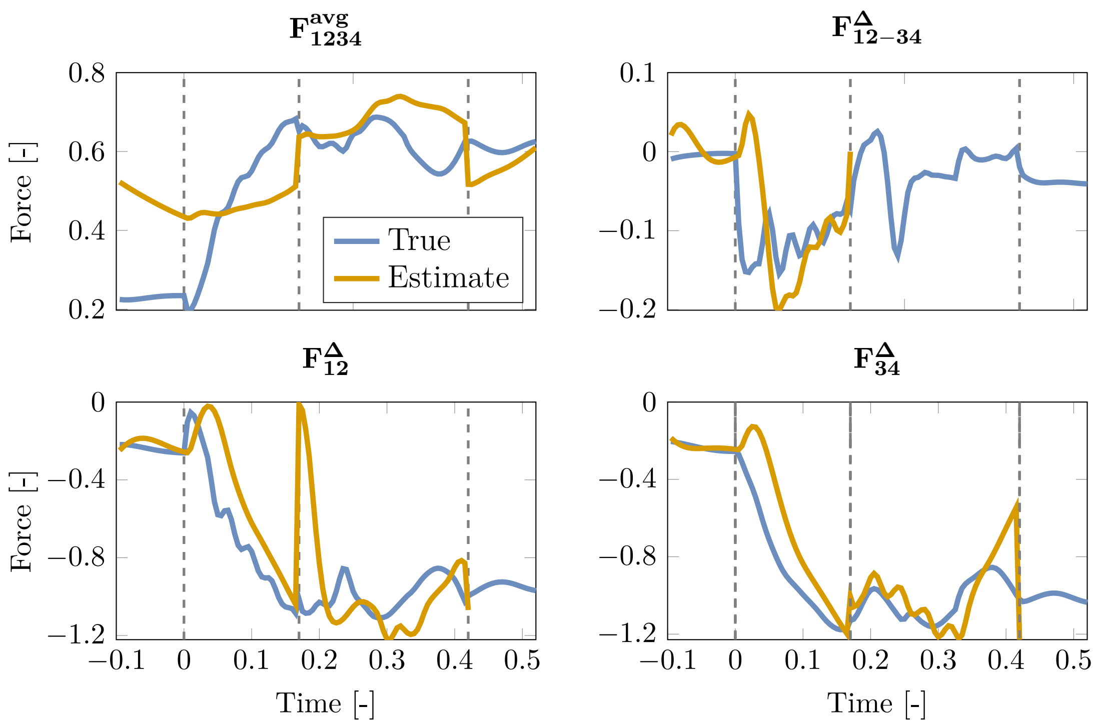

Figure 5.1: Estimated and true tire forces in the transformed representation over a full synchronization sequence on the split-friction road. Vertical dashed lines mark the start of IA synchronization, the IA lock, and the final IW lock.

5.1.1 Analysis

The estimates track the true tire forces reasonably well throughout the sequence. The largest deviations occur during rapid transients, where the random walk model has limited ability to track fast changes. A notable jump in the estimates is visible at the second dashed line, when the IA differential locks. This is caused by the clutch's locking force, which is not modeled in the estimator and therefore appears as a sudden, unexplained change in tire force estimates. The estimator recovers from the transient, and the estimates resume tracking the true values. The jump could potentially be mitigated by temporarily freezing or limiting the estimator update during the locking event. Overall, the estimator performance is considered sufficient for the purposes of this thesis. Further improvements could be achieved with a more sophisticated approach, for example, by incorporating lateral dynamics and nonlinear estimation methods such as an extended Kalman filter, but this is outside the scope of this work.

5.2 LQR Behavior on Split-Friction Road

The first controller evaluated is the LQR. This section presents its behavior on the split-friction road, using a representative tuning, to illustrate the controller's general characteristics and to introduce concepts referenced in subsequent sections.

The resulting state trajectories, brake torque requests, engine torque request, and longitudinal vehicle velocity are shown in Figure 5.2. The green lines mark the acceptable engagement interval defined by $\omega_{\text{lb}}^{\Delta}$ and $\omega_{\text{ub}}^{\Delta}$, the red line indicates the wheel speed corresponding to $0\,\%$ wheel slip, and the three vertical dashed lines marks the key events in the synchronization sequence.

5.2.1 Analysis

The LQR locks all differentials with near-zero velocity loss. Only the left brakes are active during synchronization, as expected, since the left wheels have spun up on the low-friction surface, and any braking on the right side would indicate an overshoot of the reference. The front left brake applies slightly more torque than the rear left during the initial peak in order to reduce the front axle average speed and drive the IA state to its reference.

The speed-difference states approach their references asymptotically, with the largest change occurring early in the synchronization and then tapering off as the error decreases. The engine torque request is held at a high level throughout which keeps the average wheel speeds close to the $20\,\%$ slip reference and the vehicle moving with minimal speed loss.

The brake torque requests remain relatively constant throughout the synchronization due to the integral action. Testing showed that without integral action, the LQR struggled to reach the speed-difference references, likely due to the constraints not being handled properly. When brake torque requests for the right-side wheels are computed but saturated to zero, a persistent low-frequency error is introduced, which the integral action effectively compensates for in this scenario.

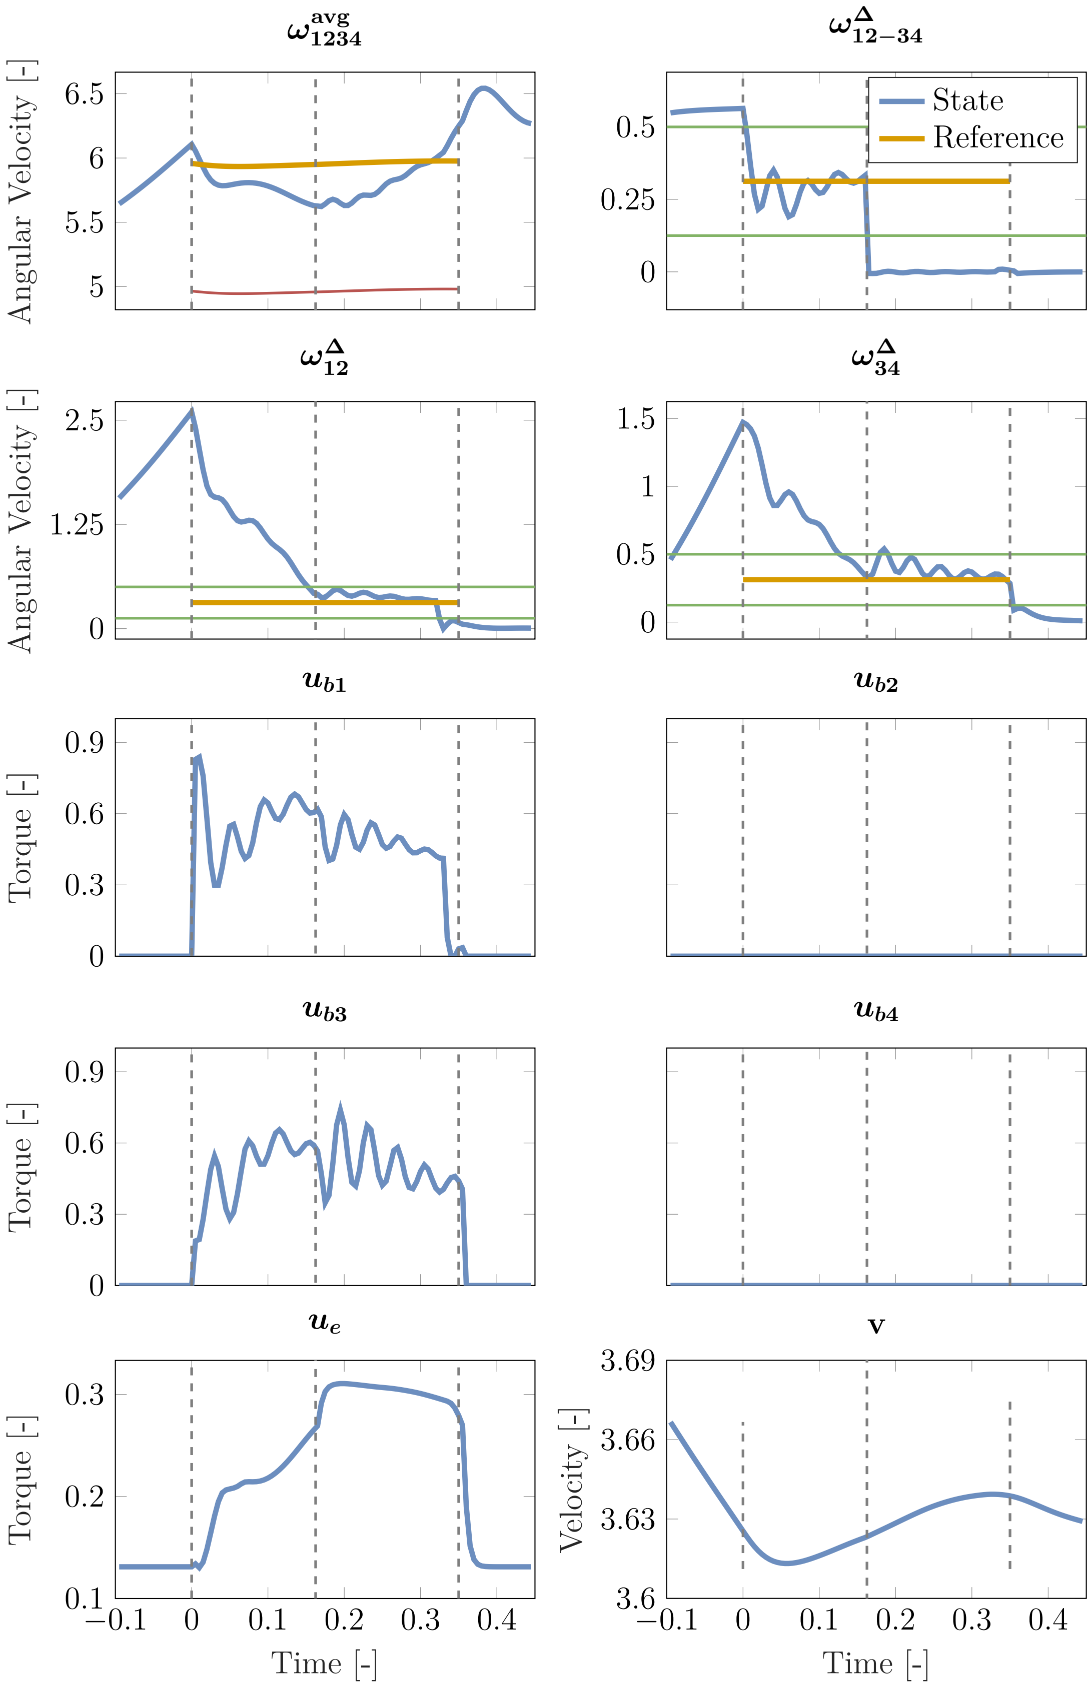

Figure 5.2: State trajectories, brake torques, engine torque, and longitudinal vehicle velocity during synchronization using the LQR on the split-friction road. Green lines mark the engagement interval, and the red line marks 0 % wheel slip.

5.3 MPC Behavior on Split-Friction Road

The MPC was evaluated on the same split-friction road scenario, also using a representative tuning. The resulting state trajectories, brake torque requests, engine torque request, and longitudinal vehicle velocity are shown in Figure 5.3.

5.3.1 Analysis

The MPC achieves synchronization and the differential states enter the acceptable engagement region smoothly and track the reference without overshoot. Both the brake and engine torque requests are significantly lower than those of the LQR, and they are more clearly proportional to the state error because there is no integral action. Only the left brakes are active, consistent with the LQR.

A notable limitation is the conservative engine torque request. More aggressive engine torque tunings were found to cause instability, which restricts the tunable parameter space and results in greater velocity loss during synchronization compared to the LQR. Since the MPC handles actuator limits directly through its optimization, control saturation does not occur, and integral action is not needed. Without integral action, the brake torque requests are more clearly proportional to the state error: when the speed difference error is small, the brake actuation is low. This means the controller has no mechanism to maintain a constant braking level to counteract the increase in wheel-speed difference caused by increased engine torque. In principle, this could be addressed by including a model that captures how an increase in engine torque leads to greater wheel-speed differences, which the current driveline model does not account for. However, capturing this behavior requires a tire model, since an equal increase in engine torque on both wheels of a differential will produce different speed increases depending on the current slip conditions. This comes with the modeling difficulties discussed in Section 2.1.4.

While the LQR benefits from integral action in this well-conditioned scenario, the same integral action is expected to struggle in more demanding conditions where the error changes too rapidly for the integrator to track. The advantage of the MPC formulation is therefore expected to be more apparent in the gravel road scenario.

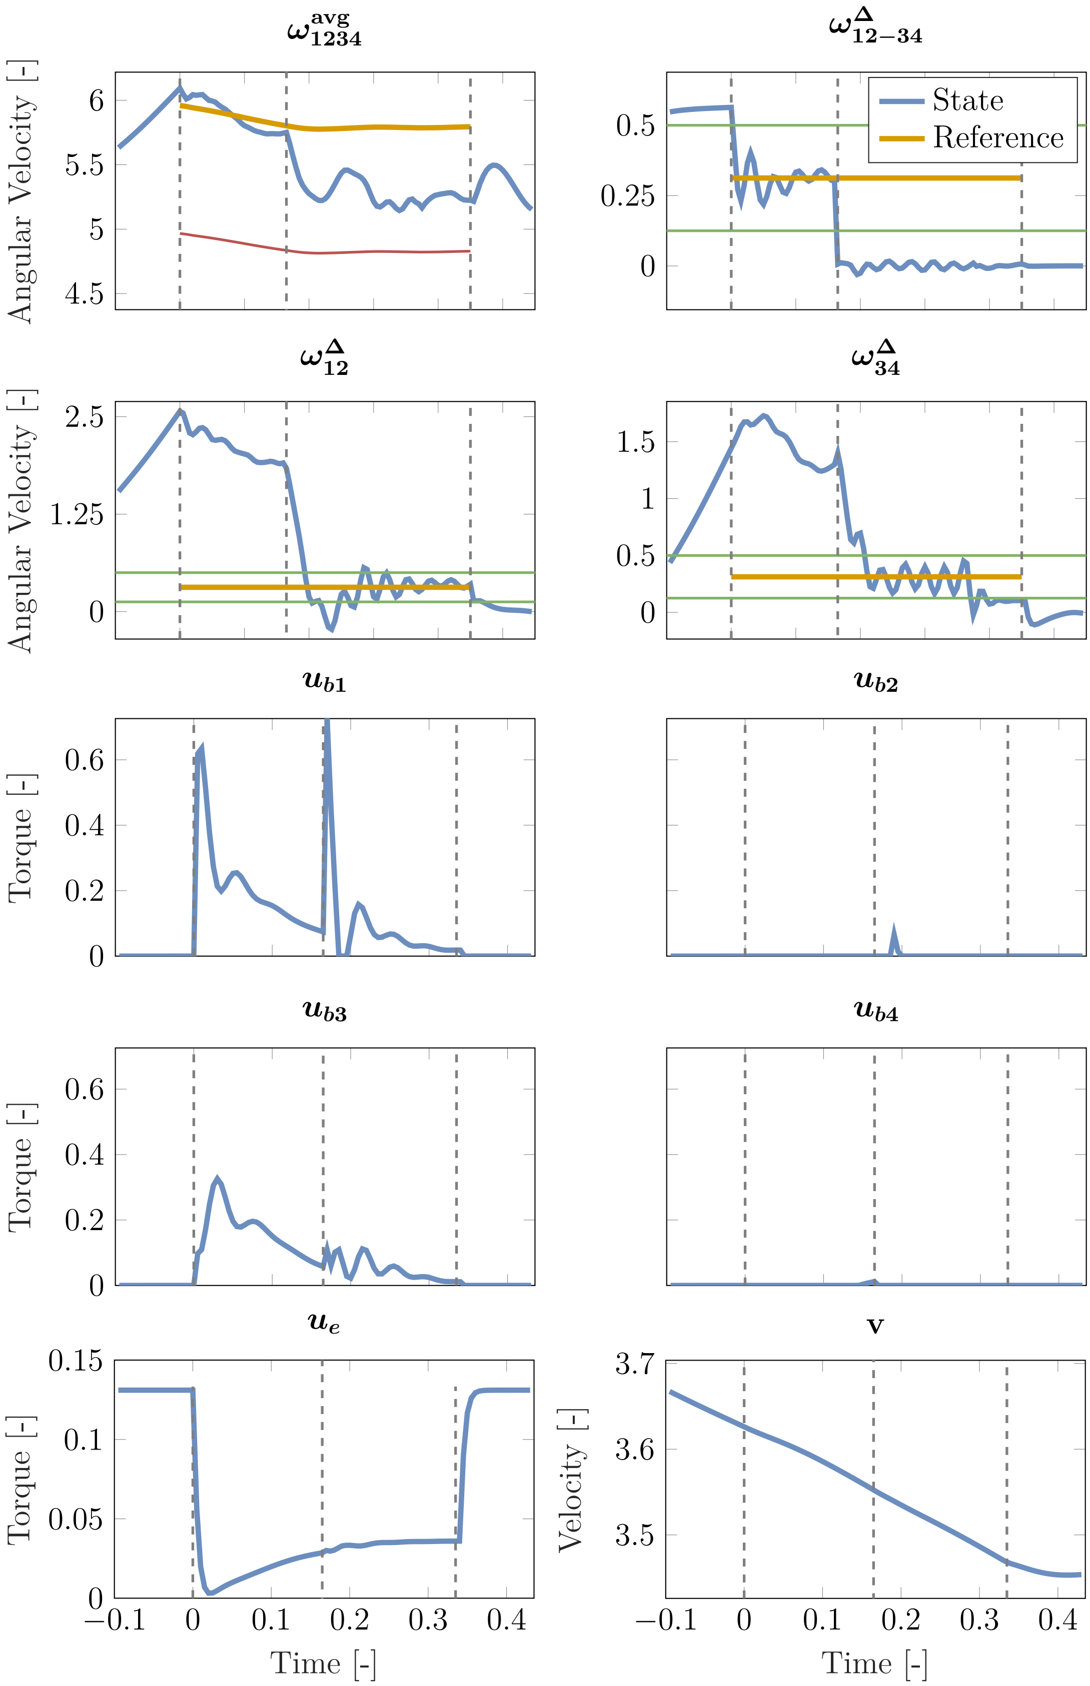

Figure 5.3: State trajectories, brake torques, engine torque, and longitudinal vehicle velocity during synchronization using the MPC on the split-friction road. Green lines mark the engagement interval, and the red line marks 0 % wheel slip.

5.4 SMC Behavior on Split-Friction Road

The SMC was evaluated on the split-friction road scenario using a representative tuning. The resulting state trajectories, brake torque requests, engine torque request, and longitudinal vehicle velocity are shown in Figure 5.4.

5.4.1 Analysis

The SMC locks all differentials with near-zero velocity loss, similar to the LQR. The engine torque request also reaches levels comparable to the LQR.

A small amount of braking is applied to the front right wheel during IA synchronization. This is because the front axle hits the low-friction surface before the rear, creating a small front-to-rear speed difference that the controller correctly compensates for.

A key behavioral difference compared to the LQR and MPC is how the speed difference states approach their references. Rather than the asymptotic approach inherent to the other two controllers, the SMC drives the states toward the reference more linearly. This is a consequence of the reaching law, where the brake actuation is independent of the magnitude of the speed difference error. The brake torque request is therefore not proportional to the state error. This allows the SMC to maintain high engine torque without the wheel-speed differences increasing, since the brakes actuate aggressively enough to counteract them regardless of the size of the state error. This is in contrast to the MPC, where the low brake actuation at small state errors limits how much engine torque can be applied.

Some chattering is visible in both the state trajectories and the brake torque requests, which is an expected consequence of the switching action inherent to SMC. The chattering remains contained within the acceptable engagement region and does not prevent successful synchronization.

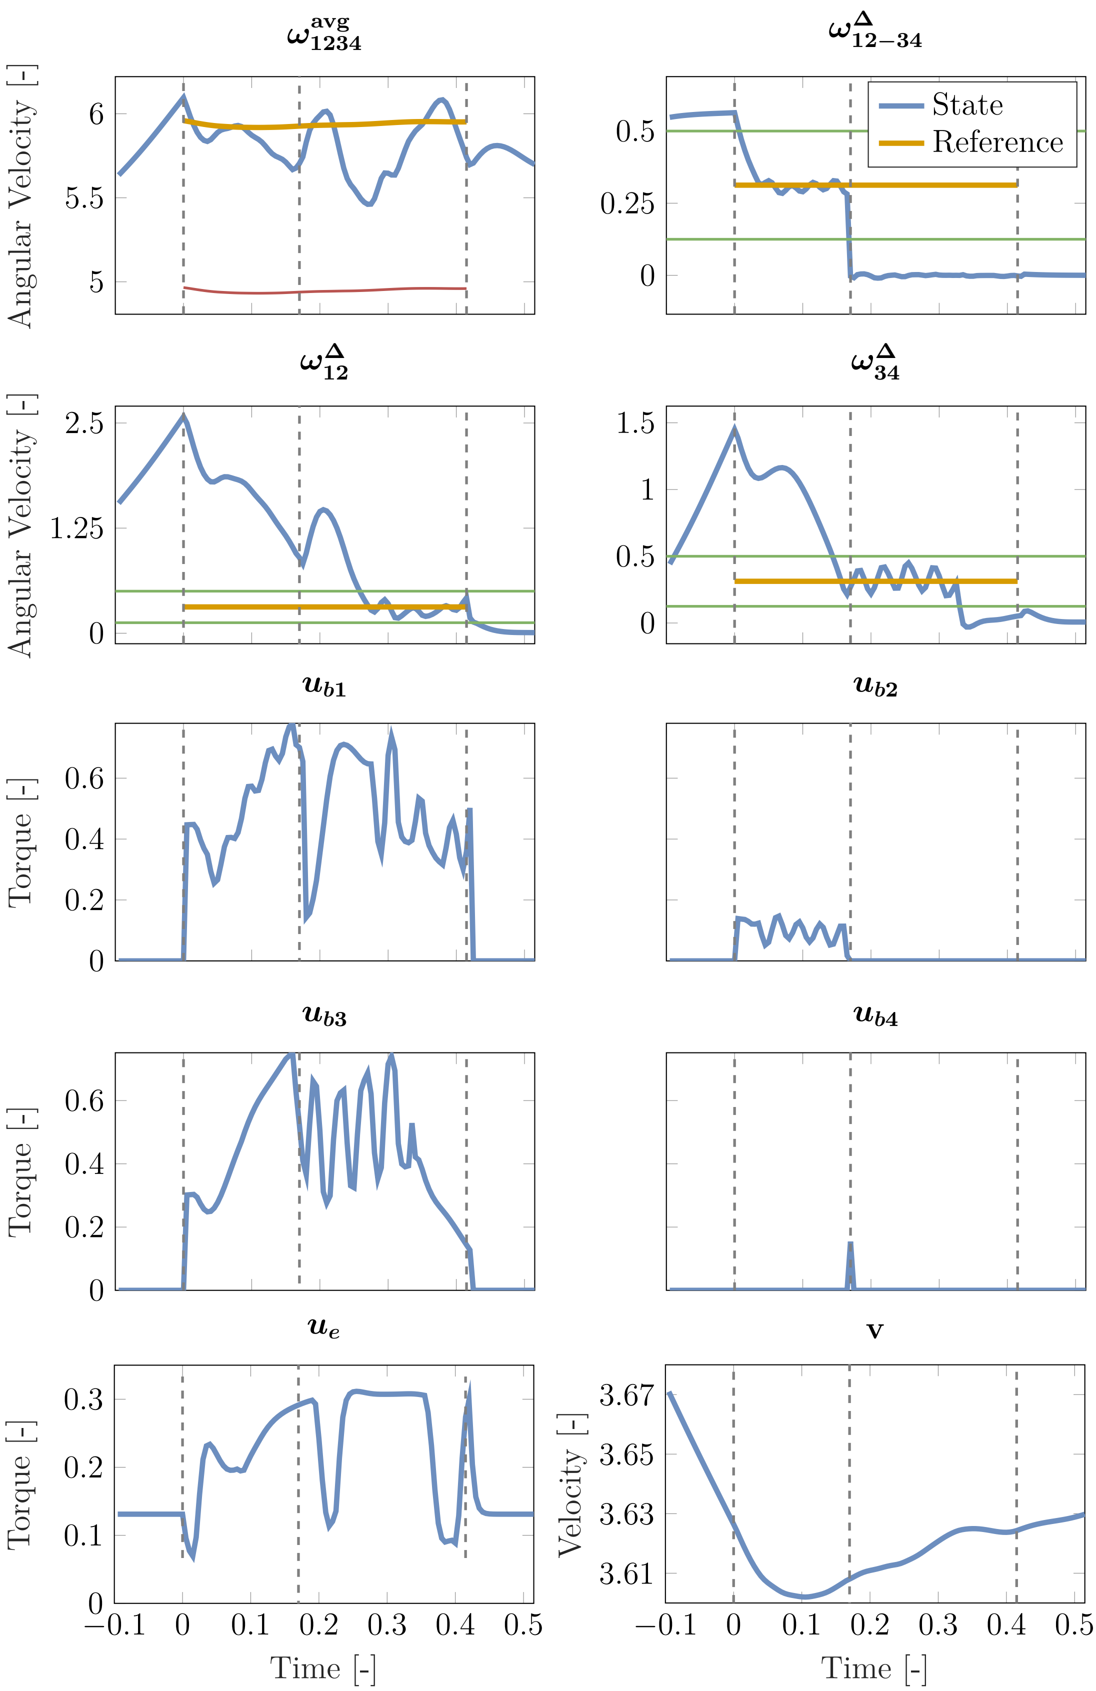

Figure 5.4: State trajectories, brake torques, engine torque, and longitudinal vehicle velocity during synchronization using the SMC on the split-friction road. Green lines mark the engagement interval, and the red line marks 0 % wheel slip.

5.5 Controller Comparison on Split-Friction Road

All three controllers were evaluated on the split-friction road with the tuning parameters swept as described in Section 4.4. The performance metrics from each run are plotted in Figure 5.5, where each dot represents one tuning. The lower-left boundary of each scatter represents the set of tunings where no further improvement in one metric can be achieved without worsening the other, forming the optimal trade-off boundary between the two metrics. Tunings far from this boundary are suboptimal and are of less interest for the comparison. Since the parameters are swept over a wide range, many tunings will naturally fall away from this boundary.

5.5.1 Analysis

On the split-friction road, the MPC achieves the shortest synchronization time but at the cost of the largest velocity loss, which is consistent with its limited engine torque usage observed in the individual controller section. The MPC also has the lowest total brake energy and lowest RMS jerk, both of which follow from the low overall actuation levels.

The LQR is the second fastest, while the SMC is the slowest of the three. The slower convergence of the SMC is related to its linear approach to the reference, as opposed to the asymptotic convergence of the LQR and MPC. Increasing the SMC gain to speed up convergence increases chattering near the reference, which can prevent the speed difference from staying within the engagement window long enough to trigger a lock. This trade-off limits how aggressively the SMC can be tuned in this scenario. One way to address this would be to change the reaching law to include a term proportional to the sliding surface, which would give faster convergence from far away while reducing chattering near the reference.

The LQR and SMC perform similarly across the remaining metrics, with higher brake and engine actuation than the MPC, consistent with observations in the individual sections. The split-friction road is a well-conditioned scenario, and the advantage of explicit constraint handling in the MPC is limited here. A more informative comparison is expected in the gravel scenario.

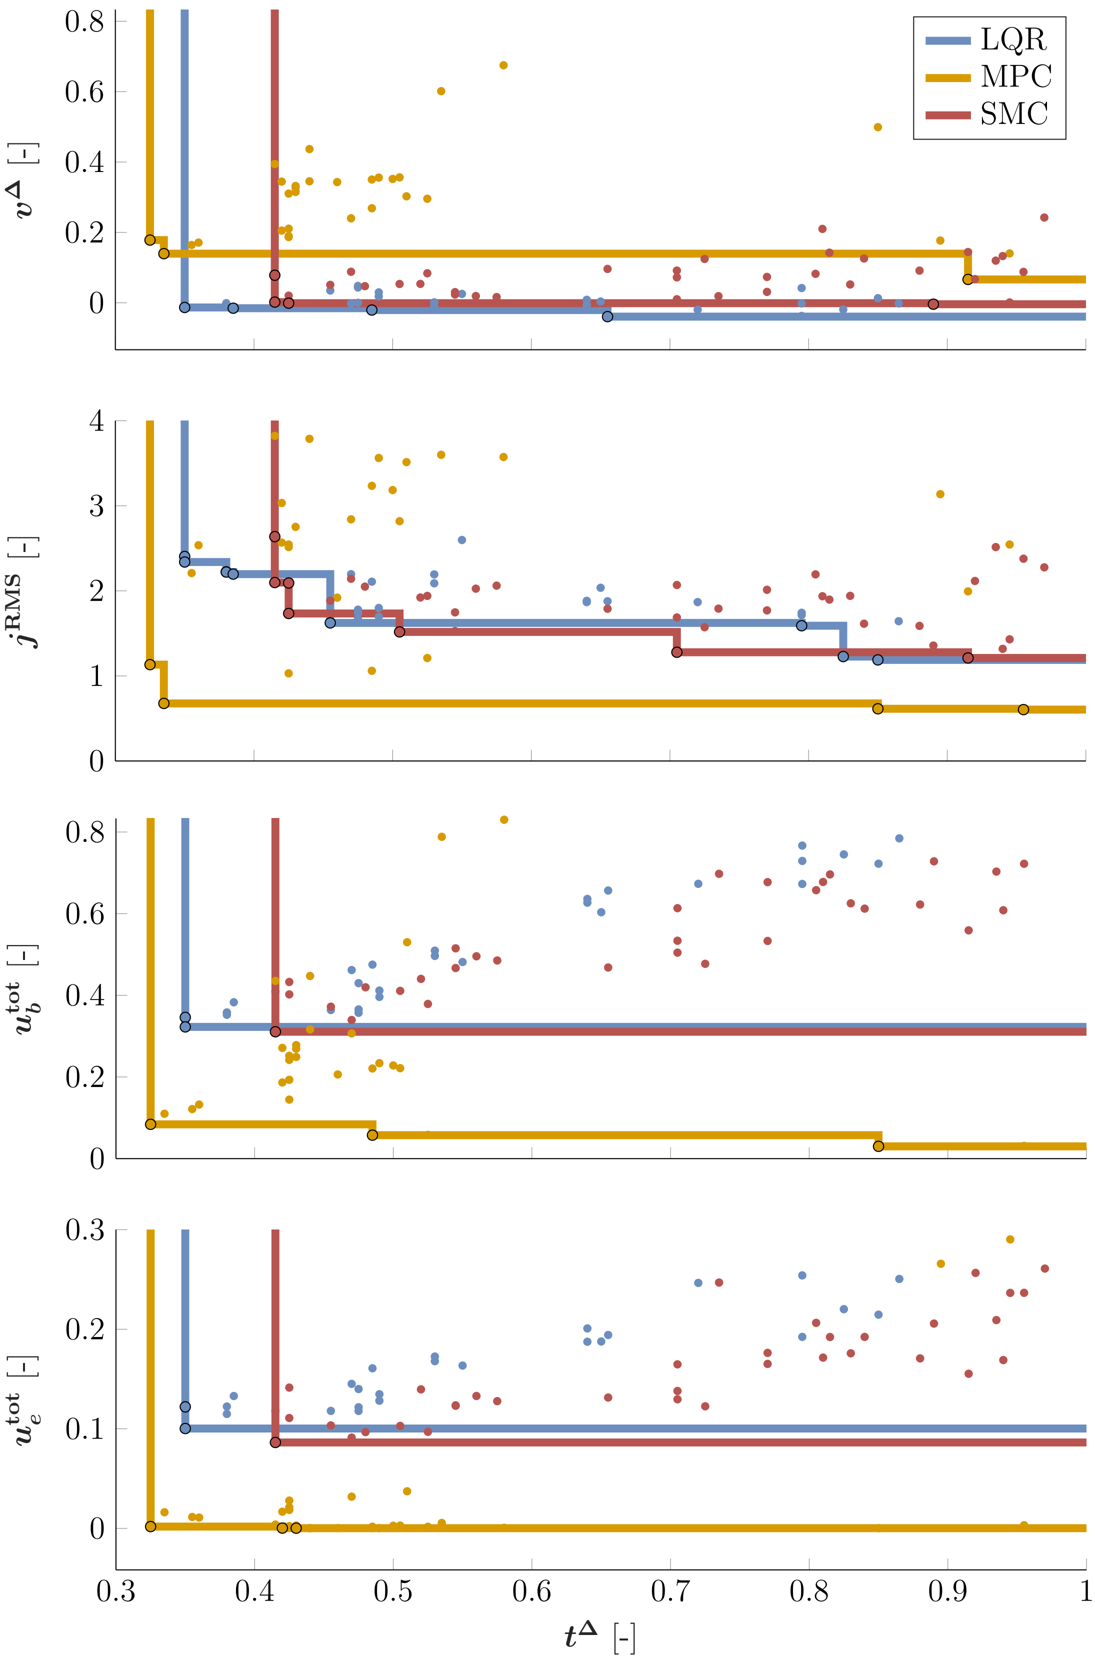

Figure 5.5: Trade-off plots comparing the three controllers on the split-friction road. Each dot represents one tuning. The lower-left boundary of each scatter represents the optimal trade-off between the two metrics shown.

5.6 Controller Comparison on Gravel Road

The same parameter sweep was repeated on the gravel road scenario. The results are shown in Figure 5.6.

5.6.1 Analysis

On the gravel road, the LQR performs noticeably worse than both the MPC and SMC in terms of synchronization time. This can be attributed to two factors. First, the LQR has no explicit constraint handling. Second, the integral action that worked well on the split-friction road struggles here: the rapidly varying friction conditions produce disturbances with higher frequency content than the integrator can track, leading to poor steady-state performance.

A notable difference from the split-friction scenario is that none of the three controllers uses significant engine torque. The uniformly low friction causes all wheels to slip, driving the average wheel speed above the reference. In this case, the controller needs to limit rather than increase engine torque to maintain the slip reference, making engine torque less of a differentiating factor between the controllers.

The MPC and SMC perform similarly overall, but differ in their brake actuation characteristics. The best-performing MPC tunings rely on peak brake torque requests roughly three times higher than the corresponding SMC tunings, which may be difficult to realize in a physical braking system. The SMC demands lower peak brake torque but higher rates of change, due to its switching behavior.

The SMC was expected to show a clearer advantage over the MPC in this high-disturbance scenario, due to its inherent robustness to disturbances. The difference is present but not as pronounced as anticipated. This may indicate that the gravel road scenario, while more demanding than the split-friction road, does not generate disturbances of sufficient magnitude and frequency to clearly differentiate the two controllers in this regard. A scenario with more severe and rapid traction variations would likely make this distinction more apparent.

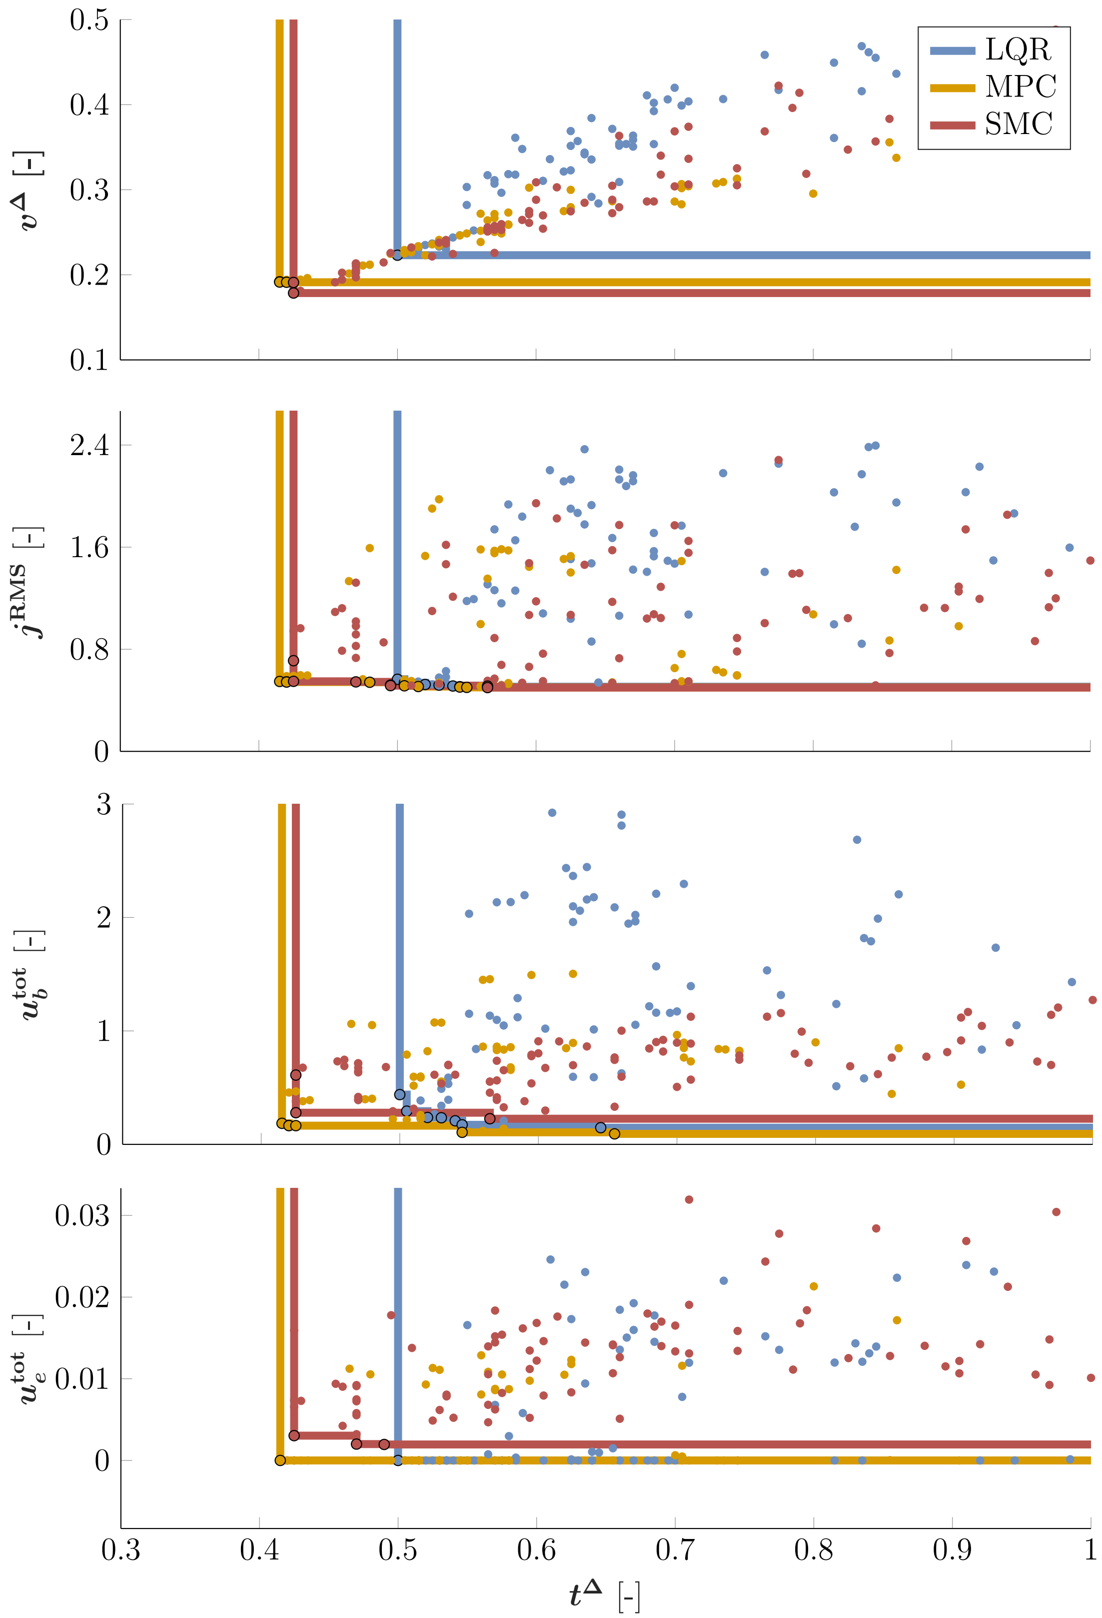

Figure 5.6: Trade-off plots comparing the three controllers on the gravel road. Each dot represents one tuning. The lower-left boundary of each scatter represents the optimal trade-off between the two metrics shown.

5.7 Effect of Engine Torque on Performance

This section investigates how the availability and control of engine torque affect synchronization performance. The investigation was conducted using the SMC on both road scenarios, testing three different engine torque modes. In the first mode, the engine torque request is set to zero during synchronization, which represents a scenario in which the accelerator pedal is released. In the second mode, the engine torque is held at the value requested by the driver just before entering the low-friction surface, which represents a scenario where the synchronization controller has no authority over the engine. The torque is kept constant rather than following a live driver model request, since the driver model would increase the torque to maintain vehicle speed during synchronization, which may not reflect realistic driver behavior. In the third mode, the controller freely adjusts the engine torque request, which is the standard mode used throughout the previous tests.

In all three modes, the controller has knowledge of the current engine torque. The results for both road scenarios are shown in Figure 5.7.

5.7.1 Analysis

On a split-friction road, engine torque clearly affects velocity loss. The zero-torque and driver cases produce similar synchronization times, but the zero-torque case results in significantly more velocity loss. The driver and SMC cases show similar total engine energy, suggesting that the controller's optimal engine torque is close to the driver's request in this scenario, consistent with the engine torque levels seen in the individual controller results.

On a gravel road, the picture is different. The controller uses negligible engine torque, similar to the zero case, consistent with the low engine torque observed in the gravel-road controller results. Velocity loss and synchronization time are largely unaffected by the engine torque mode, suggesting that engine torque plays a minor role on this surface. However, in the driver case, when engine torque is not reduced, the brake torque must increase substantially to counteract the resulting wheel-speed differences, leading to greater actuation of both the engine and the brakes without any performance benefit.

Taken together, these results highlight the value of active engine torque control. In scenarios where engine torque is beneficial, the controller can maintain it to minimize velocity loss. In scenarios where it is not, the controller reduces it to avoid unnecessary brake demand. It is also worth noting that in neither scenario does the controller request more engine torque than the driver's original request, suggesting that an implementation in which the controller only limits, rather than actively increases, engine torque would perform well and may be simpler to implement from a safety perspective.

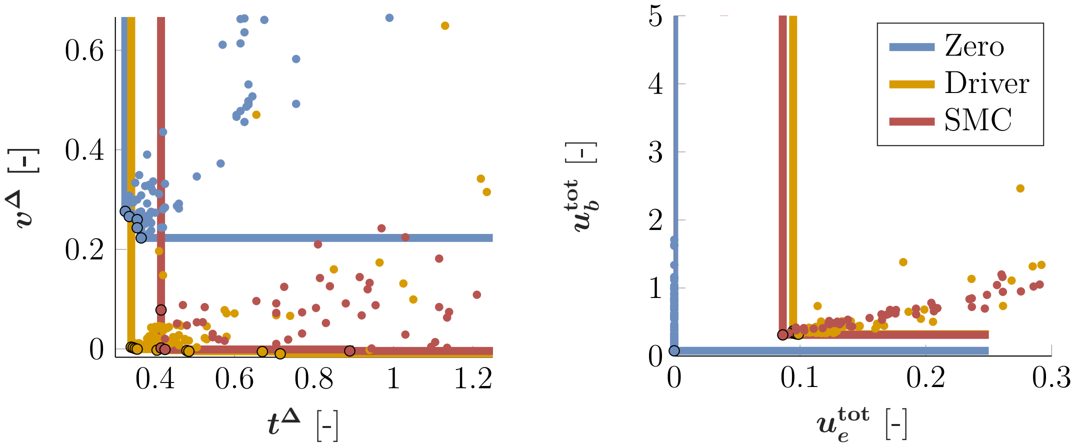

(a) Split-friction road.

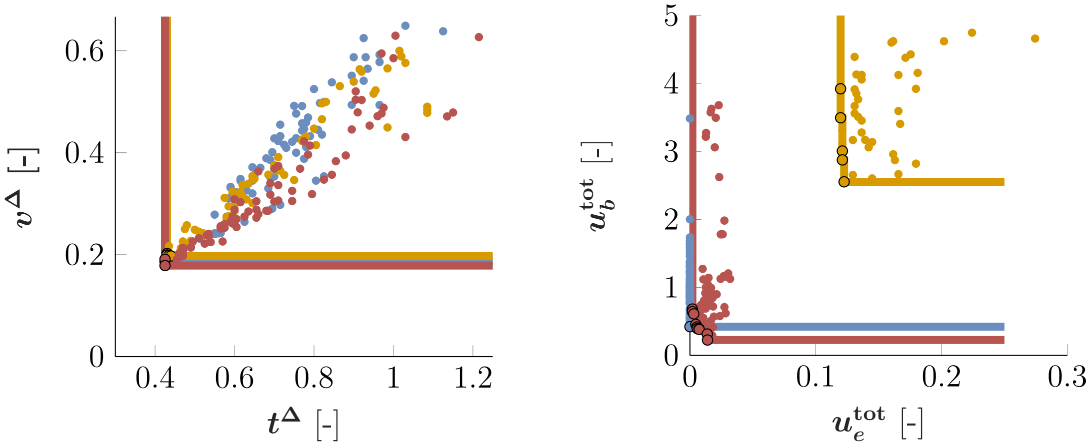

(b) Gravel road.

Figure 5.7: Trade-off plots showing the effect of engine torque mode on SMC performance for both road scenarios. Left plots show synchronization time versus velocity loss, right plots show total engine impulse versus total brake impulse.

6
Conclusion

Of the three controllers evaluated, the SMC and MPC are both suitable candidates for active differential shaft synchronization, while the LQR is considered less suitable due to its lack of explicit constraint handling and its reliance on integral action, which performs poorly under rapidly varying traction conditions. Based on the results and the nature of the control problem, a control strategy following the principles of the SMC is believed to be best suited for this application. The control problem is inherently prone to large and unpredictable tire force disturbances, and a controller whose output is not proportional to the state error maintains stronger disturbance rejection near the reference than proportional controllers such as the LQR or MPC. The boundary layer in the SMC also maps naturally onto the engagement window of the dog clutch.

No significant trade-off was found between synchronization time and the other performance metrics. Faster synchronization consistently leads to less velocity loss, and tuning for speed does not meaningfully worsen driver disturbance or total control effort. A well-tuned controller tends to perform well across all metrics simultaneously.

The problem is essentially a form of slip regulation, and actuator freedom similar to what is required in ABS and ASR systems is likely needed for the controller to perform well. Active engine torque control has a clear benefit in this regard. Maintaining high engine torque during synchronization requires correspondingly high brake actuation to counteract the induced wheel-speed differences, and if engine torque is not actively managed, this can lead to excessive brake demand without any performance gain. The controller usually does not need to increase engine torque beyond the driver's pre-synchronization request to perform well, suggesting that an implementation in which the controller limits rather than actively increases engine torque would be sufficient.

6.1 Future Work

Introducing a tire model into the controller would capture the unmodeled coupling between engine torque and wheel-speed differences, potentially enabling higher engine torque utilization without instability. Improving the tire force estimator to reduce the estimation lag inherent in the random-walk formulation would also likely improve controller performance, since more accurate and responsive force estimates directly benefit feedforward disturbance cancellation. Both extensions come with the modeling challenges associated with tire models discussed throughout this thesis.

The SMC could be augmented with an exponential reaching law to reduce synchronization time, potentially closing the gap to the synchronization times achieved by the MPC. It could also be investigated whether the control problem could be reformulated to avoid online optimization entirely, possibly through a control allocation approach in which a desired net torque per axle is first determined and then distributed to the available actuators using a less computationally demanding algorithm.

The role of the transmission clutch in the synchronization process could also be investigated, both as a means of rapidly reducing engine torque and as a way of influencing the driveline coupling during synchronization.

Finally, the evaluation could be extended to a wider range of road conditions and vehicle configurations, and the effect of actuator limitations such as rate limits, peak torque capacity, and bandwidth differences between the brake and engine actuators on controller performance could be studied more thoroughly. The results could either be incorporated directly into the controller design or assist in determining actuator authority requirements.

Bibliography

[1] M. Dabhi and K. Ramanan Vaidyanathan, "Automation and synchronization of traction assistance devices to improve traction and steerability of a construction truck," Master's Thesis, KTH Royal Institute of Technology, Stockholm, Sweden, 2017.

[2] P. Kučera, V. Píštěk, and P. Porteš, "Automatic differential lock control in a truck–straight drive," in *Proceedings of the 22nd International Conference Engineering Mechanics*, Svratka, Czech Republic, 2016, pp. 363–366.

[3] G. Bóka, L. Lovas, J. Márialigeti, and B. Trencséni, "Engagement capability of face-dog clutches on heavy duty automated mechanical transmissions with transmission brake," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 224, no. 9, pp. 1125–1139, Sep. 1, 2010, ISSN: 0954-4070, 2041-2991. DOI: 10.1243/09544070JAUTO1435.

[4] J. Hallqvist, "Faster locking differential through active brake-control," Master's Thesis, Linköpings universitet, Linköping, Sweden, 2014.

[5] T. Shim, S. Chang, and S. Lee, "Investigation of sliding-surface design on the performance of sliding mode controller in antilock braking systems," *IEEE Transactions on Vehicular Technology*, vol. 57, no. 2, pp. 747–759, Mar. 2008, ISSN: 1939-9359. DOI: 10.1109/TVT.2007.905391.

[6] H. He, J. Peng, R. Xiong, and H. Fan, "An acceleration slip regulation strategy for four-wheel drive electric vehicles based on sliding mode control," *Energies*, vol. 7, no. 6, pp. 3748–3763, Jun. 17, 2014, ISSN: 1996-1073. DOI: 10.3390/en7063748.

[7] J. Song and W. S. Che, "Comparison and evaluation of brake yaw motion controllers with an antilock brake system," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 222, no. 7, pp. 1273–1288, Jul. 1, 2008, ISSN: 0954-4070. DOI: 10.1243/09544070JAUTO839.

[8] A. Aljawabrah, "Analysis and control of the gearshift process based on a dog clutch shiftability mode," *CogSust*, vol. 3, no. 3, Sep. 30, 2024, ISSN: 2939-5240. DOI: 10.55343/CogSust.120.

[9] M. Z. Piracha, A. Grauers, E. Barrientos, H. Budacs, and J. Hellsing, "Model based control of synchronizers for reducing impacts during sleeve to gear engagement," in *SAE Technical Papers*, Detroit, Michigan, United States, Apr. 2, 2019, pp. 2019–01–1303. DOI: 10.4271/2019-01-1303.

[10] T. Grigoratos and G. Martini, "Brake wear particle emissions: A review," *Environ Sci Pollut Res*, vol. 22, no. 4, pp. 2491–2504, Feb. 2015, ISSN: 0944-1344, 1614-7499. DOI: 10.1007/s11356-014-3696-8.

[11] United Nations. "The 17 sustainable development goals," Accessed: May 18, 2026. [Online]. Available: `https://sdgs.un.org/goals`.

[12] *Road vehicles — vehicle dynamics and road-holding ability — vocabulary ISO 8855:2011*, Geneva, Switzerland, Dec. 15, 2011.

[13] B. Jacobson, *Vehicle Dynamics Compendium.* Gothenburg, Sweden: Chalmers University of Technology, 2019.

[14] E. Bakker, L. Nyborg, and H. B. Pacejka, "Tyre modelling for use in vehicle dynamics studies," in *SAE Technical Paper*, SAE International, 1987, p. 20. DOI: `10.4271/870421`.

[15] C. Duan, "Analytical study of a dog clutch in automatic transmission application," *SAE Int. J. Passeng. Cars - Mech. Syst.*, vol. 7, no. 3, pp. 1155–1162, Apr. 1, 2014, ISSN: 1946-3995, 1946-4002. DOI: `10.4271/2014-01-1775`.

[16] M. Wojtowicz, U. Wiklund, R. Elo, and S. Jacobson, "A new experimental technique for assessing impact-sliding wear resistance of materials for non-synchronous transmissions," *Wear*, vol. 586, 2026, ISSN: 00431648. DOI: `10.1016/j.wear.2025.206476`.

[17] M. Andersson and K. Goetz, "FE analysis of a dog clutch for trucks with all-wheel-drive," Master's Thesis, Linnaeus University, Växjö, Sweden, 2010.

[18] B. Källstrand, "Control allocation for vehicle motion control-maximizing traction and steering capabilities under different road conditions," Master's Thesis, Chalmers University of Technology, Gothenburg, Sweden, 2016.

[19] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, Mar. 1, 1960. DOI: `10.1115/1.3662552`.

[20] T. Glad and L. Ljung, *Control Theory: Multivariable and Nonlinear Methods*, 1st edition. London: Routledge (Taylor & Francis), Mar. 30, 2000, 482 pp., ISBN: 978-0-7484-0877-1.

[21] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model predictive control: theory, computation, and design*, 2nd edition. Madison, Wisconsin: Nob Hill Publishing, 2017, 623 pp., ISBN: 978-0-9759377-3-0.

[22] Q. Huang and H. Wang, "Fundamental study of jerk: Evaluation of shift quality and ride comfort," presented at the SAE Technical Papers, SAE International, May 2004. DOI: `10.4271/2004-01-2065`.

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY