Optimal Control Allocation on Over-Actuated Vehicles

Maximizing vehicle performance and safety through advanced vehicle control strategies

E.H. van den Berg

Master of Science Thesis

TU Delft Delft University of Technology

Delft Center for Systems and Control

Optimal Control Allocation on Over-Actuated Vehicles

Maximizing vehicle performance and safety through advanced vehicle control strategies

Master of Science Thesis

For the degree of Master of Science in Systems and Control at Delft University of Technology

E.H. van den Berg

August 10, 2016

Faculty of Mechanical, Maritime and Materials Engineering (3mE) · Delft University of Technology

Abstract

Modern passenger vehicles are equipped with an increasing number of actuators that may be used to actively control the lateral and longitudinal dynamics of the vehicle. During limit-handling situations, proper coordination of all the available actuators by the vehicle stability control (VSC) can lead to improved overall control authority, which in turn may lead to improved handling performance and decreased intrusiveness to the driver. The difficulty, however, is coordinating the available actuators, given that the addition of actuators typically leads to the vehicle becoming over-actuated.

The state-of-the-art method of solving the over-actuation is optimal control allocation, often referred to as Global Chassis Control (GCC). A drawback of most formulations of GCC proposed in literature is the prohibitive computational burden associated with the optimization. To obtain a real-time feasible GCC algorithm, the hybrid steepest descent optimization method is applied to the control allocation problem in this work. The optimization problem is set up to minimize tracking errors of virtual control inputs whilst minimizing actuator effort. A high level controller is designed that produces these virtual control input targets, which are a yaw moment used to stabilize the vehicle, and longitudinal force, to represent the acceleration intent of the driver. Using online linearization of a nonlinear vehicle model the yaw moment and yaw moment effectiveness of the available actuators is estimated and used to update the optimization problem. Furthermore, improvements are suggested to the hybrid steepest descent method to accelerate convergence and reduce or eliminate chatter at constraint boundaries by dynamically scaling the constraints.

The optimal control allocator is shown, using a validated simulation model, to produce improved allocation performance when compared to a simpler control allocation method that is similar to the industry-standard. Furthermore, the proposed algorithm was converted to C-code and implemented on one of the on-board ECUs of a Tesla Model S, demonstrating real-time feasibility. Experimental results for an aggressive single lane-change using this implementation show the algorithm provides good performance compared to an industry standard brake-based stability control system.

Table of Contents

Acknowledgements

First of all, I would like to thank my supervisors Prof.dr.ir.J. Hellendoorn and Dr.ir. M. Corno for their guidance, critical comments, ideas and support throughout the entire process that culminated in this thesis.

I'd also like to thank my colleagues at Tesla for their flexibility, patience and support, allowing me to combine both my research work as well as my day-to-day job.

Furthermore, I want to extend my appreciation to all my friends for their encouragements and help where possible in completing this thesis.

Finally, I want to thank my parents for their continued support and encouragements throughout the entire process, from the subtle and not-so-subtle hurry-ups they have given me along the way, to their thorough and critical reviews of my work.

Delft, University of Technology E.H. van den Berg
August 10, 2016

Chapter 1

Introduction

Over the past century the automobile has become an integral part of modern society, enabling affordable mobility to the masses. The numbers speak for themselves: the US alone was home to over 240 million registered passenger vehicles in 2012, which traveled a combined 4.7 trillion kilometers. This unprecedented level of mobility is however not without its cost: over 30 thousand people died in car accidents in the US in 2012 which, for some perspective comes down to a fatality in the US alone every 16 minutes. Vehicle related fatalities are the leading cause of death in the ages between 11 and 27 in the US [1].

Events such as high speed swerving or sudden changes in surfaces friction can make it hard for the untrained driver to keep control of the vehicle and are often the cause of a crash. A significant number of these types of accidents can be reduced in severity or even completely prevented [2] by making it easier for the driver to maintain control over the vehicle in such events. Active safety systems, systems intended to help the driver in these kinds of limit-handling situations, first started appearing back in 1980's and have progressively become more commonplace in commercial vehicles to the point that most are now fitted as standard. An overview of the various driver aids introduced in the last three decades is shown in Figure 1-1.

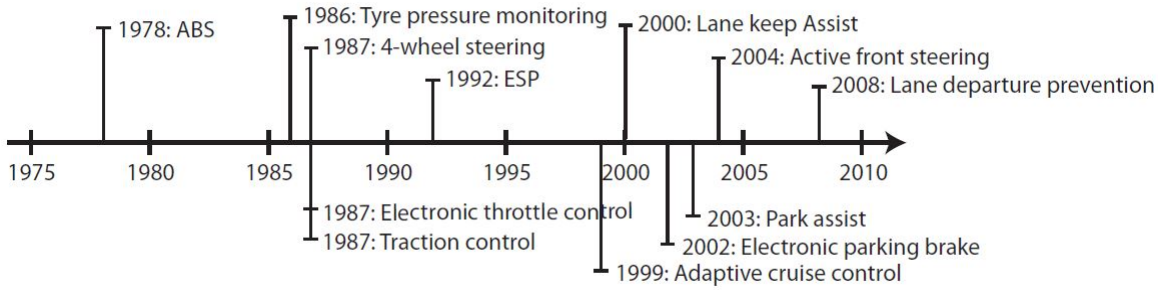

Figure 1-1: Introduction of various driver aids [3]

This thesis will focus on one of these active safety systems, being Vehicle Stability Control

(VSC, also referred to as ESP or ESC). This system is there to aid the driver with the lateral control of the vehicle. Despite the fact that control systems like this have been around for over two decades, continual improvements have been proposed due to the emergence of new actuators and increasing computational power, allowing for more advanced control strategies to be used.

1-1 Vehicle limit handling

Safety systems such as VSC should engage when there is a significant chance that the driver is losing control or is close to losing control. Loss of control, like spinning out, is typically the result of the front or rear tires reaching their limit of grip, which subsequently leads to non-linear and often unstable vehicle behavior in combination with incorrect or delayed corrective driver actions. This nonlinear and unstable vehicle behavior is inherent to nearly all vehicles, mostly due to the force producing characteristics and limitations of the tires. Pneumatic tires, found on nearly every commercial passenger car behave predictably and almost linearly when producing small forces generated during normal operation, but have inherently nonlinear and highly variable characteristics when required to generate large forces. These characteristics depend on, for example, the type of tires (e.g. summer/winter), the state of the tires (e.g. new/worn), the quality and friction of the surface it is on (e.g. high grip tarmac/ice) etc.

To understand how the tire forces interact with the main chassis dynamics of the vehicle some definitions of vehicle chassis states need to be set. Figure 1-2 shows a heavily simplified vehicle model:

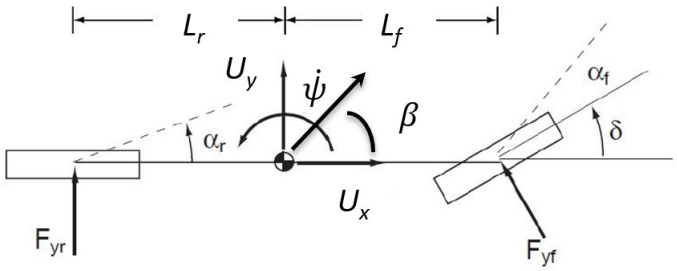

Figure 1-2: Vehicle dynamics coordinate frame from [4]

In Figure 1-2 a vehicle is shown in a typical coordinate system. This coordinate system and corresponding sign convention will be used as standard throughout the rest of this thesis. The origin of the coordinate system is in the vehicle center of gravity (CG). The positive longitudinal direction is defined in the direction of U_x, the positive lateral direction as defined by U_y. In plan view, a positive yaw rate $\dot{\psi}$ is a rotation of the vehicle body around its center of gravity in counterclockwise direction. The angle between the longitudinal and lateral velocity at the CG is defined as sideslip angle β.

Sideslip angle β and yaw rate $\dot{\psi}$ are key variables when it comes to stability control design for vehicles [5, 6]. Sideslip, yaw rate and steering wheel angle determine for a large part the

amount of lateral slip on each tire. If these slips reach appreciable values it typically means the tire is saturated or close to saturation. When the front tires reach this state before the rear tires the car is said to understeer. This is a stable state of the car, but the fact that the front tire can't produce additional grip means the driver has lost directional control authority of the vehicle, which can be dangerous and disconcerting to the driver. Even more dangerous is when the rear tires reach their limits first, resulting in the vehicle building sideslip rapidly and without fast corrective measures, will be said to 'spin out'. It is the goal of a stability control system to regulate vehicle states such as the yaw rate and sideslip in some way and keep both extremes from occurring by keeping these vehicle states within desired boundaries to help the driver maintain control over the vehicle.

1-2 Vehicle stability control: an overview

Vehicle stability control systems, from here on abbreviated as VSC systems, attempt to regulate these vehicle states to ensure these stay within certain acceptable boundaries. The way this is achieved in most passenger vehicles, in a very general sense, is by applying a yaw moment to the car, either to directly or indirectly regulate sideslip and yaw rate. On current VSC systems this yaw moment is generated by modulating the individual brakes of the car to generate a differential longitudinal force across the front and/or rear axle [7, 2]. This differential force produces a yaw moment around the CG of the vehicle. A very basic example is shown in Figure 1-3. However, using solely brakes might not always be the best

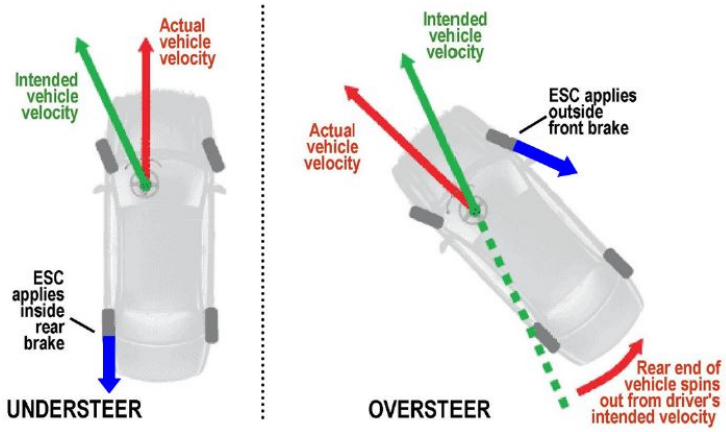

Figure 1-3: Example of under- and oversteer correction using differential braking

way of achieving this yaw moment. Brakes typically have limited control bandwidth and accuracy, can be noisy and brake usage means that energy is being wasted to perform the regulation, because the brakes slow the vehicle down when the driver might not be intending to slow down. This makes pure brake based stability control intrusive to the driver [8] and relatively inefficient which means it is only used when strictly necessary. With the rise in popularity of the electric car (Tesla Roadster/Model S/Model 3, BMW i3/i8), the VSC system has access to high bandwidth, accurate and efficient control of wheel or axle torque using electric motors. Furthermore, axle torques may be actively controlled using electronic differentials [9, 10] in 4WD vehicles and systems like active front steering or active

rear steering. These are becoming commonplace features on high-end commercially available vehicles. However, the various actuators that may be placed on a car don't all offer the same level of control authority and intrusiveness. Electric motors may have limited power and torque outputs, differentials can only bias certain percentages of torque etc. Given that most actuators have a specific limitation and operating area where they are effective, control strategies tailored to a specific actuator layout have been presented in literature. However, if the best possible use is to be made of the available actuation power and bandwidth, the actuation effort needs to be coordinated appropriately. Research into Global Chassis Control (GCC) has attempted to come up with a generalized approach of handling multiple actuators. However, the generalization is difficult given the wide variety of actuators that may be used [3].

1-2-1 Actuation

Given that actuator effectiveness and limitations play a defining role in VSC design, this topic will be studied in further depth. A list of some of the common actuators used in literature on VSC is presented below:

1. **Differential braking**: the first among the actuation systems that found its widespread use in the commercial car was the ABS system [7], allowing for individual control of brake pressure of a wheel. Although initially designed to prevent tire slip during heavy braking, it was quickly realized that this system could be exploited for other purposes such as generating yaw moments by coordinating wheel slips across an axle, referred to as differential braking in this thesis. This actuation method is available on nearly all modern vehicles

2. **Torque vectoring differentials**: these allow for controllable left/right motor torque distribution [11, 12] and on 4WD vehicles the possibility of front-to-rear torque distribution [10]

3. **Active steering systems**: both front wheel and rear wheel steering have also been well covered in literature [13, 14, 3] and are becoming increasingly popular on high-end luxury vehicles

4. **Adaptive roll moment distribution**: this indirectly impacts the lateral forces distribution between the front and rear axle. This can be done for example, using magneto rheological dampers, an often found feature on high performance sports cars [13]

5. **Electric motors acting on a single wheel**: found in either a 4WD or 2WD drive configuration [15, 16, 17, 3]

Electric vehicles have become increasingly popular and a large amount of academic research has been dedicated to exploring its possibilities for yaw moment control. Most often this research is focused on vehicles with electric motors acting on a single wheel, given that this actuator layout gives the greatest amount of yaw moment authority. The relevance of that research to the current breed of electric vehicles is however limited as the number of commercially available cars with independently driven wheels on an axle is extremely limited (e.g. Honda NSX, Porsche 918). However, there is an increasing number of commercially

available hybrids such as the BMW i3 and BMW i8 or full electric vehicles such as Tesla's Model S and X with completely independently driven axles. Such vehicles could benefit from using torque split strategies for yaw moment control, possibly in combination with differential braking and other actuators, if present.

1-2-2 Control allocation problem

When a vehicle is equipped with a number of actuators that may be independently controlled and that can be used to achieve the same effect, the system is said to be over-actuated [18, 19]. A practical example that this thesis will focus on is a vehicle that is equipped with differential braking capabilities and a torque split system. In order to exert a yaw moment on the vehicle the algorithm that arbitrates this yaw moment, the control allocator, has to decide if it either wants to exert this yaw moment using solely torque split, solely differential braking, or a combination of the two. In order to make such a decision additional information in the form of some decision variable or constraints are required. These decision variables or constraints may be simple, or they may be cast in the form of an optimization problem, where the decision variable is typically derived from some cost function which describes trade-offs between certain actuator properties such as saturation limits, desired equilibrium positions, actuation effort, tracking performance etc. In its most general sense the control allocation problem has the following form:

$$\dot{x} = f(t, x, u) \tag{1-1}$$

$$u = h(t, x, v) \tag{1-2}$$

where $x \in \mathbb{R}^n$ the states, $u \in \mathbb{R}^m$ the actuator inputs, $v \in \mathbb{R}^k$ the virtual control input and $k < m$. The actuator inputs u exert an effective control input v, where the mapping between u and v is determined by $h(t, x, v)$. The virtual control input v is typically generated by a control system, such as VSC. The necessity for $h(t, x, v)$ arises as, for over-actuated systems, this mapping from virtual control input v to actual control input u is non-unique and hence a control allocation algorithm is needed to provide this mapping.

In most production cars the control allocation is performed by applying a rule-based algorithm [20] which, although robust and relatively simple, typically doesn't make optimal use of the combined capabilities of the actuators due to a lack of coordination and optimization. Furthermore, it might be rather cumbersome to tune in order for such a control allocation algorithm to have acceptable performance. With further increases in complexity of the controls and additional actuators, such allocation schemes will become intractable. Therefore, more integrated schemes are required to solve this control allocation problem.

1-3 Control allocation in literature

The control allocation problem was first tackled in the aerospace industry, given that most aircraft are over-actuated by design. This is typically done to provide some level of redundancy in case of an actuator or actuator set failure. This also meant methods were required to distribute control effort accordingly over the available actuators. Some of the lessons learned

in the aerospace field can be applied to control allocation for road going vehicles given that the problem, in essence, is very similar. A short overview of methods for the aerospace and automotive fields is presented below:

1-3-1 Static control allocation

When constraints are added between actuators this decreases the number of degrees of freedom which can help solve the indeterminacy of the problem. Often used methods are daisy chaining of actuators or, as is common in aerospace, ganging of actuators [19, 21].

Ganging of actuators is achieved by adding ratio-metric constraints to the control effort of different actuators and in that way grouping them together, eliminating degrees of freedom. Ganging of actuators, although common in the aerospace industry (controlling the ailerons, for example) seems to be an uncommon approach for automotive control allocation in GCC due to a lack of publications on this particular topic.

More common is daisy chaining or sequential control, where, as the name implies, actuators are used in sequence [19, 21]. One actuator is used first until it saturates, after which the next actuator becomes operational. Such a daisy chaining approach can be found in [14] where this control allocation method is applied in an automotive setting. Daisy chaining performance is compared to a system that uses no actuator coordination for a vehicle with active suspension and differential braking. Much improved performance is found when using daisy chaining.

Both methods, although effective in solving the indeterminacy of the problem, might not be optimal. 'Optimal' in this context can refer to: actuator effort, tracking performance, tire potential utilization or other secondary optimization goals such as intrusiveness to the driver.

1-3-2 Dynamic control allocation

There are alternatives for solving the control allocation problem that have been proposed in literature which dynamically select the optimal actuator configuration. The optimality implies that there are some optimization criteria which allows an allocation algorithm to find the 'best' control input u that satisfies the virtual control request v. These optimization criteria can take different forms, such as error minimization between virtual control and actual control output in the face of actuator constraints, actuator effort minimization, actuator rate minimization etc. In [18, 22] a number of these methods are discussed, where [21] gives a good overview of the different optimization methods that are used in aerospace to solve control allocation problems and provides comparisons on optimality and computational burden using a few examples between the different algorithms. The control allocation algorithms discussed in [21] all use a form of non-predictive allocation. However, most works on GCC apply a form of predictive control allocation. Justification and comparison to non-predictive methods often lack in the research, but some works, e.g. [18, 23, 24] show improved performance compared to the usage of simpler control methods.

1-3-3 Optimization schemes

In most publications the dynamic control allocators show good performance and give the control designer a powerful way to determine how the allocation is performed. There is however a large drawback with dynamic control allocators: some form of optimization is required in order to find the optimal control output. The computational burden associated with these optimizations, especially in a predictive setting, can be prohibitively large. The additional performance gained with such control allocation methods over much simpler and computationally lighter allocation schemes is in practice often not enough to justify the extra costs that have to be incurred to allow for sufficient computation power. However, with decreasing costs in commercially available processors and advances in computational efficiency of optimization schemes, dynamic control allocation can be considered as a serious candidate over static control allocation methods.

The main optimization methods proposed in literature are:

1. **Repeated optimization**, where at every time control cycle the optimization for the dynamic control allocation is performed. The methods proposed in e.g. [21, 23, 24, 25] use this approach. This method is generally computationally demanding, as the full optimization needs to carried out at every time-step.
2. **Precomputed laws**, used in the field of Explicit Model Predictive Control [26, 27, 28] where the optimization is solved offline. When the right formulation is used for the optimization the entire space of optimal points can be mapped, and the results stored in a set of look-up tables. These look-up tables can reach impractical sizes if the size of the optimization problem grows in terms of variables or constraints.
3. **Update laws**, such as used in [3], where the optimization variables are treated as states of a dynamical system. The optimal solution is not directly calculated, but the variables move towards the optimal point over time.

1-4 Scope of thesis

Given the lack of publications on yaw moment control using torque split strategies between front and rear axles and the increasing number of vehicles on the road with these capabilities, research in this actuation method alone would already be worthwhile. However, given that nearly all vehicles nowadays are equipped with differential braking capabilities, the scope of this thesis was expanded to include this as well. This means a control allocation problem needs to be solved. Dynamic control allocation is state of the art in this field and one such method will be worked out in further detail in this thesis. Testing and tuning of the dynamic control allocation algorithm will be done in simulation and performance will be compared to a simpler allocation strategy. Additionally the dynamic control allocator will be tested in an actual vehicle to validate the performance. The test vehicle used for research is a Tesla Model S P90DL, which has two independently electrically driven axles and four independently controllable hydraulic brakes.

1-4-1 Control allocator design goals

As a guideline for the design of the dynamic control allocator, two main design goals were identified based on the literature research and conversations with people in the industry. The control allocator should:

1. **Achieve a high level of performance** by fully exploiting control authority of the available actuators whilst minimizing the intrusiveness to the driver and actuator usage
2. **Ease of implementation** using a problem formulation and optimization technique that is computationally simple, transparent and easy to tune

As optimization method, an update law approach will be used. The selected method is hybrid steepest descent optimization proposed in [3]. The approach suggested in this work will be modified to better fit the design goals and actuator architecture investigated in this thesis.

1-4-2 Thesis outline

An outline of the covered topics per chapter is shown below:

Chapter 2: Vehicle and Tire Models

Chapter 2 will cover the vehicle model used in this thesis and the validation of this model. The simulation environment that this vehicle model is part of will be used to evaluate the performance of the control allocator.

Chapter 3: Problem Definition

Chapter 3 defines the tasks of the different modules in the VSC system. It discusses the basic control strategies used for the high and low level controller, as well as determining the interfaces of those modules with the control allocator.

Chapter 4: Hybrid Steepest Descent Optimization

This chapter will focus on the formulation of the hybrid steepest descent optimization problem. It covers a simple example to explain the basics of hybrid steepest descent optimization and provides the full formulation of the optimization problem that can be applied to control allocation on the test vehicle. Furthermore, this chapter will present a theoretical analysis of actuator yaw moment control authority, focusing on an actuator layout such as on the test vehicle, to gain an understanding of the control authority of the various actuators present on the test vehicle. Finally, this chapter will also cover the design of the simple control allocator, which will serve as a benchmark for the hybrid steepest descent optimization control allocator.

Chapter 5: Control Allocator Evaluation

This chapter will cover the evaluation of the dynamic control allocator, comparing its performance to that of the benchmark control allocator on a set of maneuvers. The dynamic control allocator is incrementally given access to extra actuators and the performance improvements are studied. Furthermore, two example cases to provide an indication of fault tolerance and robustness are also treated in this chapter.

Chapter 6: Experiments

Finally, a high level controller and hybrid steepest descent optimization control allocator will be implemented on the ECU of the test vehicle, and one of the maneuvers will be performed to validate the simulation results.

Chapter 2

Modeling Environment

The investigation of control allocation using hybrid steepest descent optimization was first performed in a simulation environment. The tire and vehicle model used in this simulation environment will be covered in this chapter, as well as the validation of this vehicle model.

2-1 Tire models

Tires play a central role in vehicle handling and tire modeling has been an ongoing field of study and research. As tires are the main way for the vehicle to exert forces on the environment, having a model that accurately reflects the way a tire generates longitudinal and lateral force is critical for obtaining a model with good accuracy.

A widely adopted method of tire modeling is through a physical approach of the tire force generating characteristics. This method assumes that the circumference of the tire consists of brushes. With the assumptions of a certain pressure distribution in the contact patch it is possible to analytically calculate the tire force as function of lateral and longitudinal slip and normal load, for example, using Fiala's theory [29]. There is a large number of variants to be found for these brush-type tire models with some of the more common ones being the Dugoff tire model [30], also referred to as the HSRI model. This tire model is used for control design and analysis in [31, 10, 32]. A series of dissertations from Stanford [4, 5] use Fromm's brush model to good effect, a model which is similar to Dugoff's model in basic construction. Both models are computationally quite simple, capable of combined lateral and longitudinal force calculations and capture the most important characteristics of the tire. The choice for either seems arbitrary and is hardly ever motivated in the papers. The Dugoff model was selected as tire model for this thesis due to its simplicity and widespread use.

The models discussed above are static tire models, which assume that the force generated is instantaneous. Its widely known however that this is not the case for pneumatic tires and there are multiple methods that may be used to model tire dynamics [33, 34]. Some of the more complicated tire models might have tire dynamics as an inherent part of the model, but

as is the case with the static models mentioned above, some dynamics of to the way the slip or force is generated may be added to approximate dynamic behavior.

2-1-1 Tire models: definitions

All of the tire models mentioned use a formulation where the force is generated as a function of the slip of the tire, both lateral and longitudinal. These slips feed into the tire model where the force producing characteristics of a tire are typically of the following form:

$$F_i = f(\alpha, \lambda, F_z), \quad i = \{x, y\} \tag{2-1}$$

Where F_i is broken down as a force in the x and the y direction of the tire.

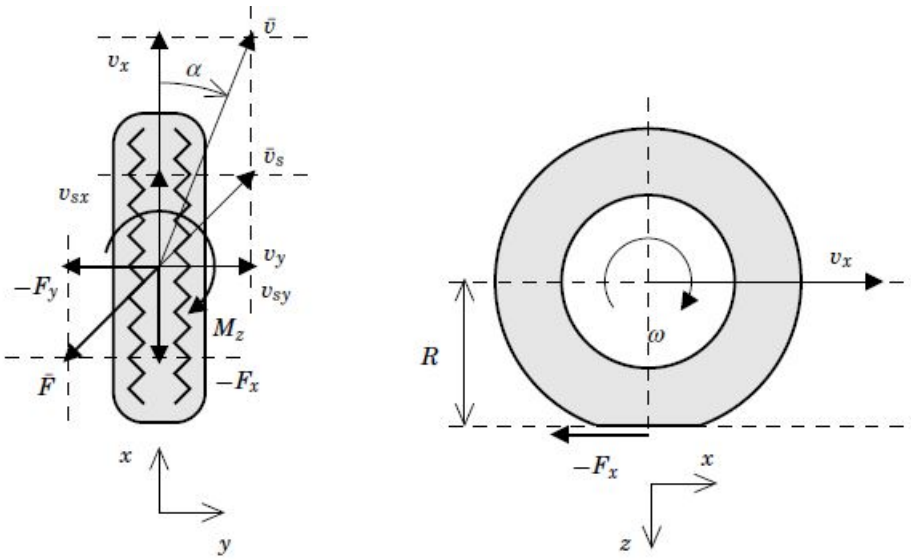

Figure 2-1: Tire kinematics definitions [33]

The following definitions of slip angle α of the tire (lateral slip) and the longitudinal slip λ are used:

$$\lambda = \frac{\omega r_w - v_x}{v_x} \tag{2-2}$$

$$\alpha = \tan^{-1}\left(\frac{v_y}{v_x}\right) \tag{2-3}$$

With r_w the effective rolling radius of the tire, ω the rotational speed of the wheel and v_x the velocity of the hub. The physical interpretation of the slips can be seen in Figure 2-1.

These equations may be expanded to include the earlier mentioned tire dynamics. A very simple and effective method, proposed in [33], is by approximating the slip dynamics as a

first-order transfer function of the form:

$$\frac{L_{relax}}{v_x}\dot{\lambda} = \tilde{\lambda} - \lambda \tag{2-4}$$

$$\frac{L_{relax}}{v_x}\dot{\alpha} = \tilde{\alpha} - \alpha \tag{2-5}$$

Where L_{relax} is the tire relaxation length. This is usually a fixed value in the order of 0.1 to 0.2 meter and describes the distance the tread of the tire needs to travel in order to build up a certain amount of slip. λ and α are the values that feed into a tire model and $\tilde{\lambda}$, $\tilde{\alpha}$ the actual instantaneous slips.

2-1-2 The Dugoff tire model

The Dugoff tire model will be used for the vehicle model in this thesis and is worked out in detail below. It describes both the lateral and longitudinal force generated by the tire as function of the tire slips α, λ, the peak friction coefficient of the tire with the surface μ_s, the normal force on the tire F_z, the velocity dependency factor ϵU_x and C_x and C_y and the tire's longitudinal and lateral slip stiffness at zero slip respectively.

$$F_x = \frac{C_x\lambda}{1-\lambda} \cdot f(\theta) \tag{2-6}$$

$$F_y = \frac{C_y \tan(\alpha)}{1-\lambda} \cdot f(\theta) \tag{2-7}$$

$$f(\theta) = \begin{cases} \theta(2-\theta) & \theta \leq 1 \\ 1 & \theta > 1 \end{cases}$$

Where θ defined as:

$$\theta = \frac{\mu_s F_z((1-\epsilon U_x\sqrt{\lambda^2 + \tan^2(\alpha)})(1-\lambda)}{2\sqrt{C_x^2\lambda^2 + C_y^2\tan^2(\alpha)}} \tag{2-8}$$

An example of combined-slip tire curves this produces is shown below. The parameters are chosen based on a typical pneumatic automotive tire. Note that the behavior is nearly linear at low slip angles in Figure 2-2 and that at the extremes the tire force producing capability is limited to what is often referred to as the 'friction ellipse' or 'friction circle' as can be seen in Figure 2-3 and is roughly described by:

$$F_{max}^2 = (\mu_s F_z)^2 = F_x^2 + F_y^2 \tag{2-9}$$

This ellipse or circle results because the peak force producing capacity F_{max} of a typical pneumatic tire is close to constant, no matter which way the slip speed vector v_s is pointing. An example of the circular approximation of the peak force producing capabilities of the tire is shown in Figure 2-3.

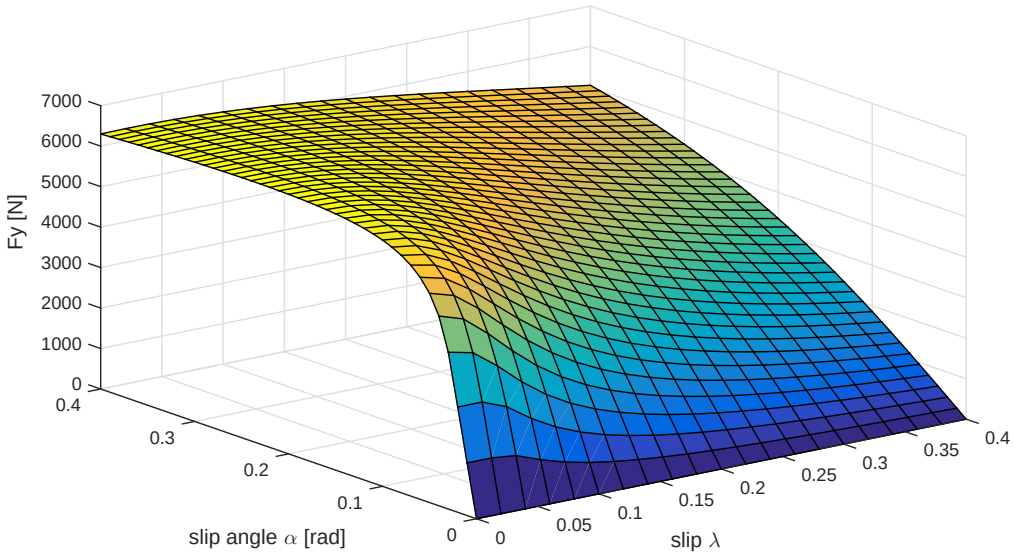

Figure 2-2: Dugoff tire model for lateral force as function of slip and slip angle for a typical pneumatic tire at a normal load of 6 [kN] and velocity of 15 [m/s]. The velocity dependency factor, taken from [31], has a relatively small effect on this particular tire, mostly scaling down grip as velocity increases. The difference between 15 and 45 [m/s], the speed range of interest, is an approximately 10% overall decrease in grip

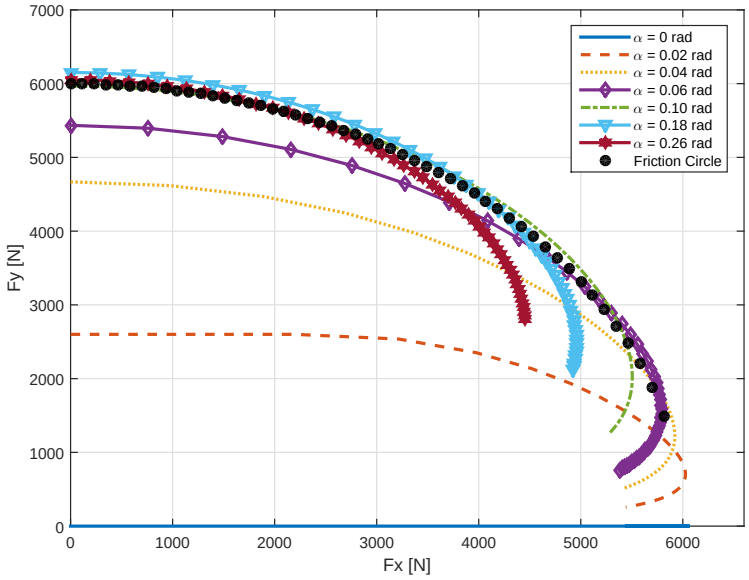

Figure 2-3: Quarter of the friction ellipse formed by the same Dugoff model in combined slip, showing the lateral force produced by the tire for a number of different slip angles as a function of the longitudinal force on the tire, shown on the x-axis. Normal load on the tire is again 6 kN

2-1-3 The linear tire model

An approximation often made is that the force producing characteristics of a tire are linear near zero slip ratio or slip angle and this is a reasonably accurate assumption up to approximately halfway the tire's force producing limit. This assumption greatly simplifies the equations and allows for a completely linear vehicle model to be derived. Such a model will be used and analyzed in Section 3-2-1.

Linearizing around zero slip gives the approximate linear relations between force and slips:

$$C_x = \left.\frac{\partial F_x}{\partial \lambda}\right|_{\lambda=0,\alpha=0} \qquad C_y = \left.\frac{\partial F_y}{\partial \alpha}\right|_{\lambda=0,\alpha=0} \tag{2-10}$$

$$F_x = C_x \lambda \qquad F_y = C_y \alpha \tag{2-11}$$

2-1-4 Sigmoid tire model

A tire model with a complexity between that of the Dugoff and the linear tire model was also developed to calculate the lateral force of the tires, under the assumption the longitudinal force is know. This model is based around sigmoid functions that approximate the slip characteristics of the typical automotive tire. Combined slip behavior was approximated using the friction circle as described by Eq. (2-9).

The sigmoid function chosen for this is the hyperbolic tangent function, tanh (). The function roughly resembles a tire curve, as can be seen in Figure 2-4, is continuous unlike the Dugoff model and has a straightforward derivative that contains its original value. This is a property that could be exploited if it is considered for fast real-time applications as it reduces the number of required calculations to find the derivative:

$$\frac{d}{dx}\tanh Cx = C(1 - \tanh^2 Cx)$$

The sigmoid tire model with circular approximation of combined slip behavior is:

$$F_{y,ij} = \sqrt{F_{max,ij}^2 - F_{x,ij}^2}\tanh(C_{y,i}\alpha_{ij}), \quad i = \{f, r\}, \quad j = \{l, r\} \tag{2-12}$$

where $C_{y,i}$ the cornering stiffness of the front and rear tires respectively. $F_{max,ij}$ is defined as:

$$F_{max,ij} = \mu_s F_{z,ij}, \quad i = \{f, r\}, \quad j = \{l, r\}$$

where $F_{z,ij}$ the normal force on the tire and $F_{x,ij}$ the longitudinal force exerted by the tire. The approximation is quite accurate up to moderate levels of slip angle, even in combined slip. It is only for larger slip angles that deviations between the two models reach appreciable levels as can be seen in Figure 2-4.

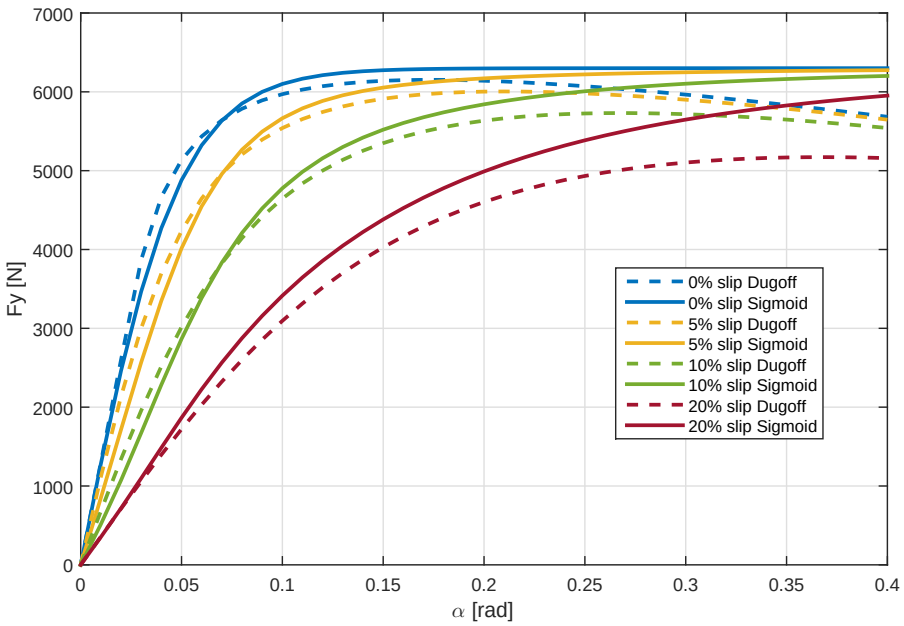

Figure 2-4: Comparing the Dugoff and sigmoid tire models for varying levels of longitudinal slip. The normal load on the tire is 6 kN

2-2 Vehicle equations of motion

The most common approach in literature on vehicle dynamics [29, 35] is to reduce the equations of motion of the vehicle to that of planar motion of a single mass with inertia around one axis resulting in a model with three degrees of freedom; one rotational and two planar degrees of freedom. These models are typically quite accurate, capturing the most important dynamics in lateral and longitudinal direction yet simple enough to apply to traditional control analysis methods to determine stability and other basic characteristics such as input-output behavior. Common ways of describing the vehicle planar motion is using a two-track model [4, 5, 36, 37, 23, 16], allowing for effects such as body roll (which might add another state) and load transfer to be included. Such a model was also selected for the simulation model of this thesis, and proved to provide sufficient fidelity as is shown in Section 2-3.

2-2-1 Two-track kinematic vehicle model

A schematic picture of a front steered four wheel vehicle, commonly known as two-track model, is shown in Figure 2-5.

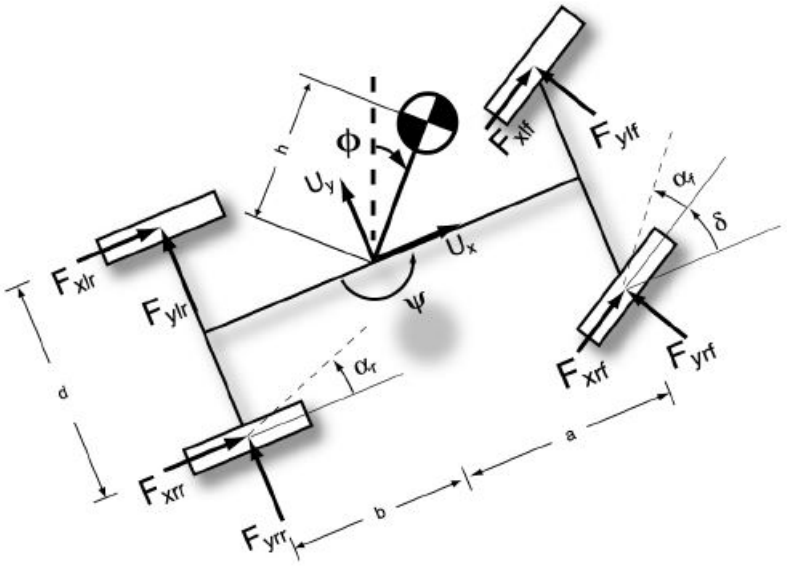

Figure 2-5: Front steered two-track model [5]

The assumptions made are that the wheels are rigidly attached to the body and the body rolls around a fixed axis defined by front and rear roll centers $h_{rc,f}$ and $h_{rc,r}$ (not shown) and δ_f is the driver's steering input to the front wheels. The equations of motion for this system can be described as follows:

$$\dot{U_x} = \frac{1}{M_{car}}\Big((F_{x,fl} + F_{x,fr})\cos(\delta_f) + (F_{x,rl} + F_{x,rr})\cos(\delta_r) - ... \\ (F_{y,fl} + F_{y,fr})\sin(\delta_f) + F_{x,disturbance}\Big) + U_y\dot{\psi} \tag{2-13}$$

$$\dot{U_y} = \frac{1}{M_{car}}\Big((F_{x,fl} + F_{x,fr})\sin(\delta_f) + (F_{y,fl} + F_{y,fr})\cos(\delta_f) + ... \\ (F_{y,rl} + F_{y,rr}) + F_{y,disturbance}\Big) - U_x\dot{\psi} \tag{2-14}$$

$$\ddot{\psi} = \frac{1}{I_{zz}}\Big(a(F_{x,fl} + F_{x,fr})\sin(\delta_f) + a(F_{y,fl} + F_{y,fr})\cos(\delta_f) + d/2(F_{x,rr} - F_{x,rl}) + ... \\ d/2(F_{x,fr} - F_{x,fl})\cos(\delta_f) + d/2(F_{y,fl} - F_{y,fr})\sin(\delta_f) - b(F_{y,rl} + F_{y,rr})\Big) \tag{2-15}$$

Where $F_{x,fr}$ denotes the force in x-direction of the *front right* tire and $F_{y,rl}$ the force in y-direction of the *rear left* tire, U_x and U_y the longitudinal and lateral velocities at the CG respectively and $\dot{\psi}$ the yaw rate. $F_{x,disturbance}$ accounts for rolling friction and aerodynamic drag components and $F_{y,disturbance}$ can be used to represent lateral disturbances such as cross-winds. The forces are calculated using the Dugoff tire model in this thesis. This tire model requires that both the lateral slip α, the longitudinal slip λ and the normal force are known.

The slip angles α_{ij} on the tires may readily be calculated using rigid body motion of the vehicle:

$$\alpha_{fl} = \delta_f - \frac{U_y + \dot{\psi}a}{U_x - \dot{\psi}d}, \qquad \alpha_{fr} = \delta_f - \frac{U_y + \dot{\psi}a}{U_x + \dot{\psi}d} \tag{2-16}$$

$$\alpha_{rl} = -\frac{U_y - \dot{\psi}a}{U_x - \dot{\psi}d}, \qquad \alpha_{rr} = -\frac{U_y - \dot{\psi}a}{U_x + \dot{\psi}d} \tag{2-17}$$

2-2-2 Wheel and powertrain dynamics

To find the longitudinal slips, λ additional state information is needed about the relative velocity of the wheel with regard to the ground. Therefore, it requires equations that describe wheel dynamics. Wheel dynamics may be approximated as pure rotational dynamics around a single axis. The resultant torque on the wheel is a function of the drive or brake torque applied to the hub and the reaction force of the tire with the ground. These equations take the following form:

$$\dot{\omega}_{ij} = \frac{1}{I_{wheel}}(T_{hub,ij} - F_{x,ij}r_w), \quad i = \{f, r\}, \quad j = \{l, r\}$$

Where i is either the front or rear axle and j either the left or right side of the vehicle, $T_{hub,ij}$ the respective torque applied to that wheel and I_{wheel} the wheel inertia. This torque maybe be generated by a number of actuators, but typically this would be either a torque from the brakes or from the connected half-shafts or hub motors. Torque production may be assumed to be instantaneous, or driveline and actuator dynamics can be added to account for effects such as motor inertia, halfshaft stiffness, brake pressure build-up etc.

As the test vehicle is equipped with both front and rear electric motors, driving the wheels through an open differential and four brakes, one on each wheel, the hub torques can be calculated as:

$$T_{hub,ij} = 0.5\tilde{T}_{m,i} + \tilde{T}_{br,ij}, \quad i = \{f, r\}, \quad j = \{l, r\} \tag{2-18}$$

With $\tilde{T}_{m,f}, \tilde{T}_{m,r}$ the dynamic equations describing the dynamics of the front and rear motor torques at the hubs and $\tilde{T}_{br,ij}$ a function that describes brake torque dynamics at the respective wheel. Rate limits and first order filters will be applied to both the brakes and motor torques to approximate their respective dynamics:

$$\begin{aligned} \tau_i\dot{\tilde{T}}_i &= T_{i,cmd} - \tilde{T}_i \\ \dot{T}_{i,min} \leq \dot{\tilde{T}}_i &\leq \dot{T}_{i,max} \end{aligned} \tag{2-19}$$

where τ_i the time constant of the system, $\dot{T}_{i,min}, \dot{T}_{i,max}$ the highest achievable downward and upward rates of the actuators and $T_{i,cmd}$ the commanded actuator output.

2-2-3 Vertical tire forces

Nonlinear tire models such as Dugoff's need the normal force F_z on each tire to calculate the respective lateral and longitudinal forces generated by the tire. To determine the normal

load distribution over the four corners of the car, some assumptions need to be made as this is a statically indeterminate system. Assuming rigid suspension, there are three equilibrium equations (roll moment, pitch moment, forces in vertical direction) and four wheels to react these forces. This requires that some assumptions are made about how the normal forces are distributed as a function of the stiffness of the suspension/chassis, better known as tire lateral load transfer distribution (TLLTD). In order to describe how these forces develop dynamically, either an eighth state, body roll, may be added, or the assumption can be made that the normal forces develop instantaneously [23, 29, 31]. Using the assumption of instantaneous forces allows one to directly calculate the normal loads. The TLLTD equations can take a few forms depending how many further assumptions are made. The equations for determining wheel loads from [29] will be used:

$$F_{z,fl} = M_{car}g\frac{b}{2(a+b)} - (U_x\dot{\psi} + \dot{U}_y)K_{roll,front} - (\dot{U}_x - U_y\dot{\psi})K_{accel} \tag{2-20}$$

$$F_{z,fr} = M_{car}g\frac{b}{2(a+b)} + (U_x\dot{\psi} + \dot{U}_y)K_{roll,front} - (\dot{U}_x - U_y\dot{\psi})K_{accel} \tag{2-21}$$

$$F_{z,rl} = M_{car}g\frac{b}{2(a+b)} - (U_x\dot{\psi} + \dot{U}_y)K_{roll,rear} + (\dot{U}_x - U_y\dot{\psi})K_{accel} \tag{2-22}$$

$$F_{z,rr} = M_{car}g\frac{b}{2(a+b)} + (U_x\dot{\psi} + \dot{U}_y)K_{roll,rear} + (\dot{U}_x - U_y\dot{\psi})K_{accel} \tag{2-23}$$

$$K_{accel} = M_{car}\frac{h}{2(a+b)} \tag{2-24}$$

$$K_{roll,front} = \frac{M_{car}}{a}\left(\frac{K_{\phi,front}h_s}{K_{\phi,front} + K_{\phi,rear} - M_{car}h_s}\right) + \frac{b \cdot h_{rc,f}}{(a+b)} \tag{2-25}$$

$$K_{roll,rear} = \frac{M_{car}}{b}\left(\frac{K_{\phi,rear}h_s}{K_{\phi,front} + K_{\phi,rear} - M_{car}h_s}\right) + \frac{a \cdot h_{rc,r}}{(a+b)} \tag{2-26}$$

where $K_{\phi,front}$ and $K_{\phi,rear}$ the total roll stiffness of each respective axle, h_s the vertical distance of the CG from the roll axis of the vehicle, h the height of the CG from the ground plane, g is the gravitational constant and M_{car} the mass of the vehicle.

A number of other effects such as suspension compliance and roll steer might have effects on these equations [29]. To reduce complexity most of these effects will be left out given that their effect is relatively minor for the vehicle that is modeled. The only compliance that was included, given its significance, was the front steering system compliance, which was modeled as:

$$\delta_f = \tilde{\delta}_f - (F_{y,fl} + F_{y,fr})C_{st}$$

where C_{st} the steering system compliance, $\tilde{\delta}_f$ the measured steering input and δ_f the steering input that is seen by the front wheels.

2-3 Model validation

To ensure the simulation model developed in this chapter provides an accurate representation of the dynamics of the test vehicle, a set of validation maneuvers were performed and compared

to the predicted model outputs. All data was collected during a Tesla-organized test trip to CRTC proving grounds in Alaska. This test site has both reasonably sized tarmac and snow vehicle dynamics areas (VDA), where testing took place. The overall goal of this test session was to collect data to validate the lateral dynamics of the vehicle. Because the vehicle model uses Tesla proprietary parameters the presented data is normalized.

2-3-1 Maneuvers and measurements

The maneuvers performed were limited by the size of the vehicle dynamics areas, but most of the relevant dynamics were captured despite the size restrictions. The inputs were human generated, with targeted lateral accelerations between 70 and 80 % of the peak tire capabilities to prevent potentially unstable behavior, but far enough to enter into the stable nonlinear region of the tire curve. Three different maneuvers were used for validation:

1. Step steers, constant speed, tarmac VDA with a μ of approximately 1
2. Short ramp steers, constant speed, tarmac VDA with a μ of approximately 1
3. Sine wave steering input, constant speed, snow VDA with a μ of approximately 0.4

Additional maneuvers such as step steers were also performed on snow. However, it was found that most of the data captured on the snow VDA proved unsuitable for analysis, as the maneuvers would very quickly reach or exceed the tire limits leading to inconsistent results. The only useable data set that was obtained was a sine wave maneuver. Additional data was also captured on the tarmac VDA, but will be left out for brevity as the above data sets show sufficiently clearly how well the model matches the actual vehicle behavior. Validation of the following vehicle states was done:

1. yaw rate ψ, measured by the on board ESP module
2. lateral acceleration A_y, measured by the on board ESP module
3. longitudinal acceleration A_x, measured by the on board ESP module
4. longitudinal velocity U_x, measured by an externally mounted dual antenna GPS system
5. sideslip β, measured by an externally mounted dual antenna GPS system

2-3-2 Validation results

The model was fed with the same inputs (steering wheel angle, brakes torques, motor torques) and initial conditions as the real vehicle and outputs were compared to the measured values. Furthermore, the model parameters that were used for the vehicle model were sourced from Tesla, and converted or simplified where necessary to fit within the format required by the vehicle model that has been developed in this chapter.

Given that this model is relatively complicated and nonlinear, using parameter identification methods to fine-tune parameters was not performed, as the large number of tunable

parameters could very quickly lead to over-fitting. Therefore it was decided to judge model fidelity mostly by visual inspection and RMS errors for each state that was validated. A comparison between model output and measured data for the three different maneuvers is shown in Figure 2-6, Figure 2-8 and Figure 2-10 and the corresponding inputs in Figure 2-7, Figure 2-9 and Figure 2-11. Furthermore, the RMS errors are shown in the table below:

	A_x RMS error	A_y RMS error	β RMS error	ψ RMS error
Tarmac, step steer	0.0400	0.0774	0.0478	0.0309
Tarmac, ramp steer	0.0213	0.0525	0.0727	0.0187
Snow, sine wave	0.0439	0.1462	0.2109	0.0573

Table 2-1: RMS errors in normalized units for the three maneuvers shown. Note the relatively large errors on snow, indicating that the model does not perform particularly well on this surface

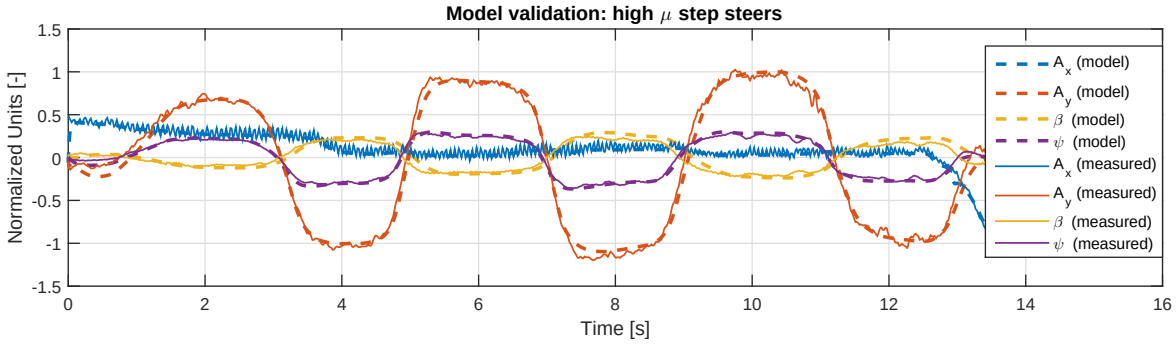

Figure 2-6: Normalized data for the tarmac VDA steps steers, $U_x \approx 34$ [m/s]

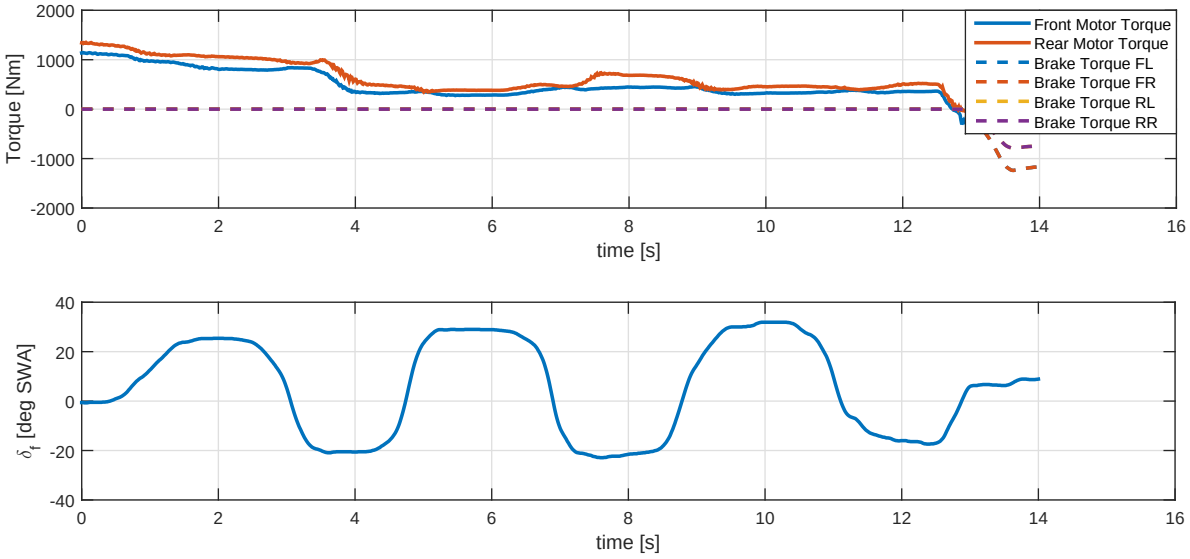

Figure 2-7: Inputs for tarmac VDA steps steers. All torques shown are at the wheel hub

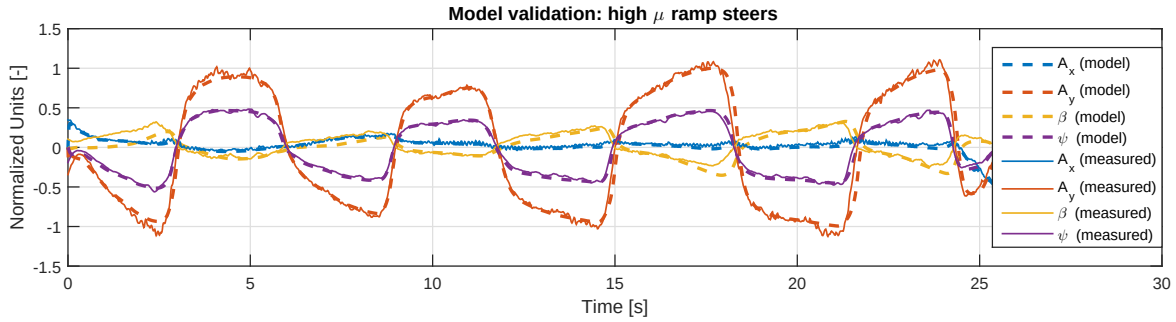

Figure 2-8: Normalized data for the tarmac VDA ramp steers, $U_x \approx 21$ [m/s]

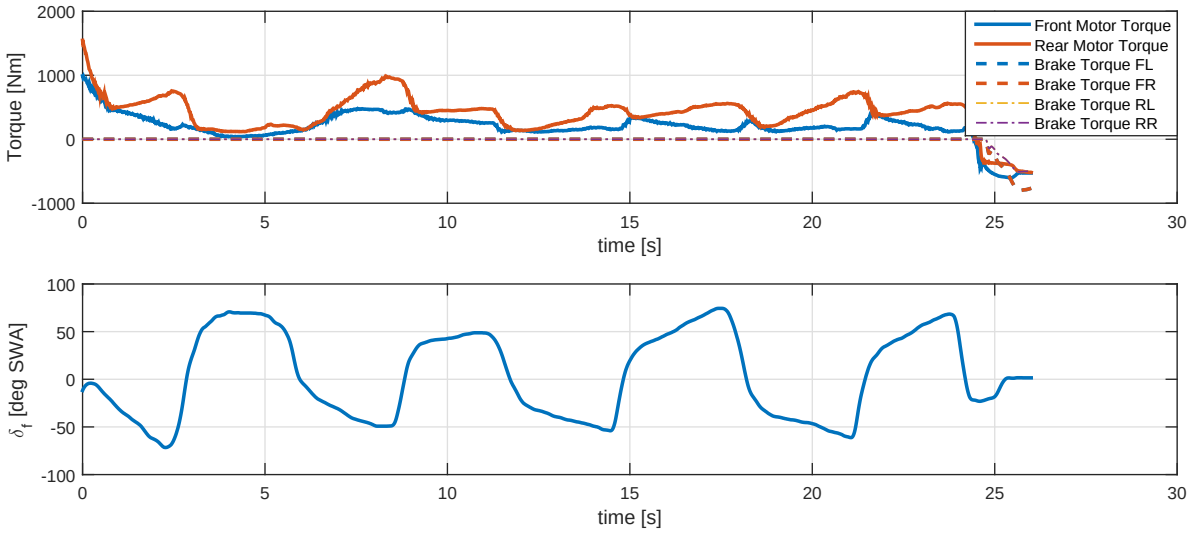

Figure 2-9: Inputs for tarmac VDA ramp steers. All torques shown are at the wheel hub

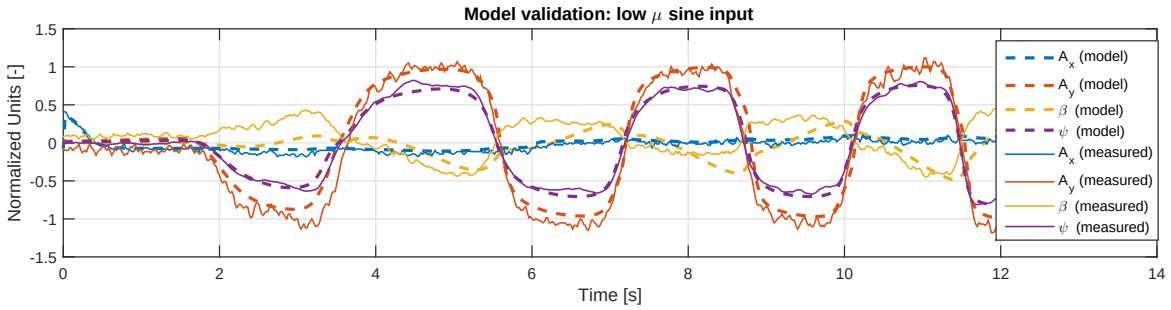

Figure 2-10: Normalized data for the snow VDA sine wave steering inputs, $U_x \approx 15$ [m/s]

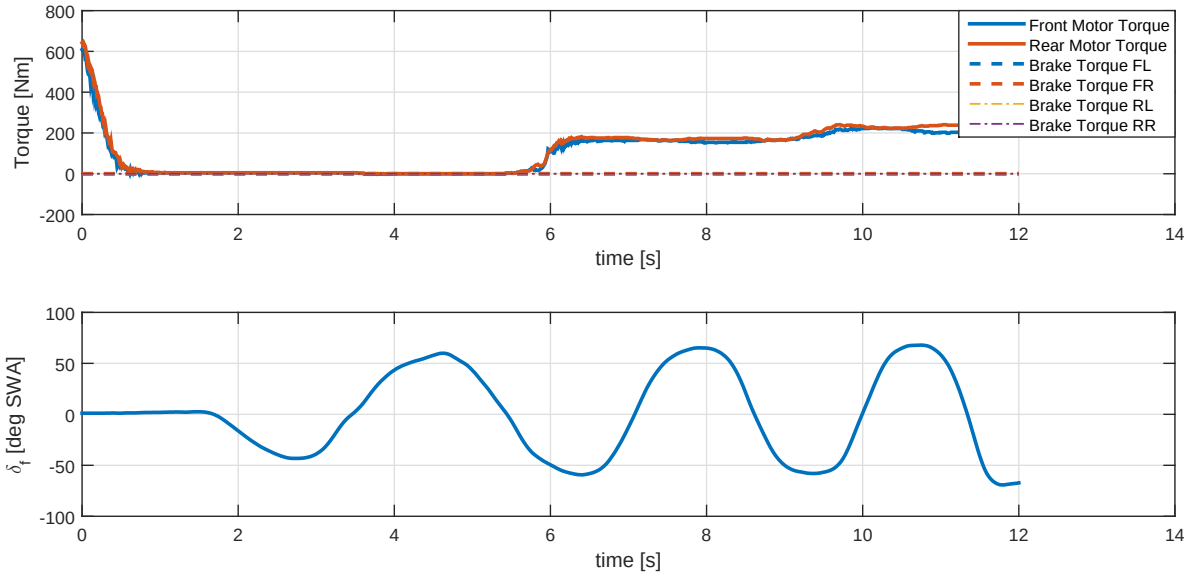

Figure 2-11: Inputs for snow VDA sine steering inputs. All torques shown are at the wheel hub

2-3-3 Conclusions

In general, the model correlates well with the measurements, especially in steady state or quasi-steady state. The correlation for the data captured on the tarmac VDA is especially good, indicated by the small RMS errors for all measured states. This was found to be the case across a wider spectrum of tests besides the ones shown, indicating that the model is a useful tool for predicting lateral vehicle behavior on tarmac. Closer inspection of the data does reveal a small amount phase lag in the sideslip dynamics on the tarmac data sets, visible in e.g. Figure 2-6. However, the effect of this on overall vehicle response is minor, hence no additional investigation was done to improve this.

The errors in sideslip dynamics in the single snow data set are much more pronounced. This does not come as a surprise given the difficulty of modeling tire behavior on low μ deformable surfaces. It is expected that major modifications will be required to the tire models to achieve higher accuracy on this type of surface. Given that this thesis is focused on control system design and not on vehicle or tire model validation, the lower accuracy on low μ surfaces is accepted as it is.

Chapter 3

Vehicle Stability Control Architecture

Despite the fact that this thesis focuses on the control allocation part of the VSC system, the full system needs to be designed to evaluate the performance of the control allocation, as it is an integral part of the VSC system. This chapter will define the goals and interfaces for the separate modules that make up the VSC system.

3-1 Control architecture

There is no consistent naming convention or control architecture for VSC systems that is used in literature. The reason for this is the wide variety of high level control objectives and available actuators that are used. In [38] an investigation was performed into the various control architectures used in literature and industry. Using this information, an architecture was selected for this thesis that is flexible and future proof, allowing for easy expansion due to its modularity with each module level having distinct functionality. The architecture is shown in Figure 3-1.

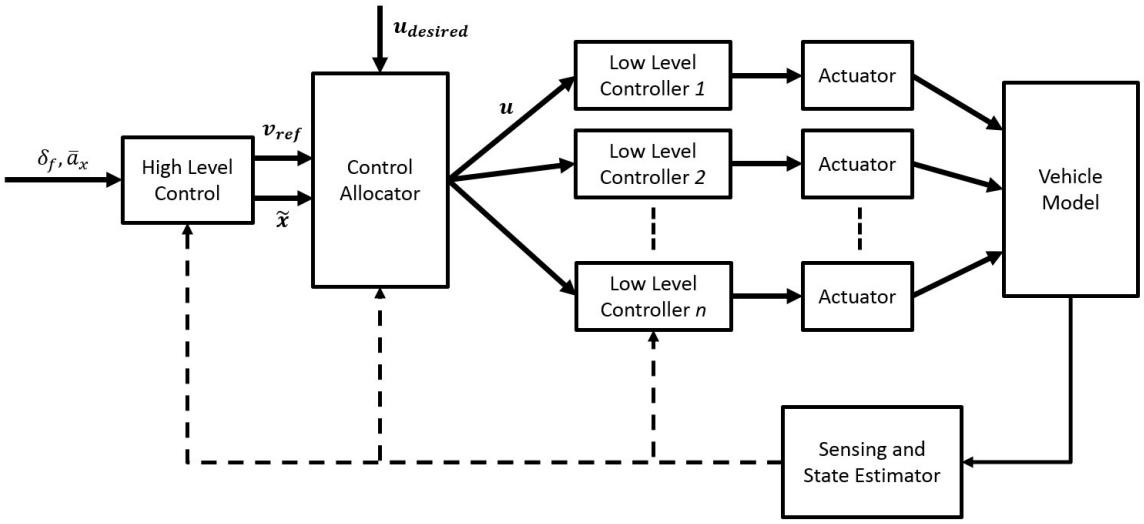

Figure 3-1: Control Architecture

Where $\mathbf{v}_{ref}$ a virtual control input request from the high level controller, $\tilde{\mathbf{x}}$ additional state information from the vehicle state estimator, $\mathbf{u}$ the actual actuator targets, e.g. torques or steering angles, and $\mathbf{u}_{desired}$ the desired steady state positions of the actuators.

Each module has a well-defined task within the VSC system:

1. **Reference generation and high level controls**. This is a module that interprets and processes driver commands and converts that to vehicle level control objectives, e.g. a reference yaw rate, sideslip bounds. A high level controller then attempts to track these references by requesting, for example, a yaw moment. Controllers for this high level control task may range from the extremely simple, such as PID controllers [25, 16, 36], to the highly complex, such as (nonlinear) MPC schemes [23, 39].

2. **Control allocation**. The control allocator that is responsible for distributing control effort over the available actuators. When multiple actuators are available for control, which is very often the case in passenger cars, the control allocator has to provide a way to allocate the control effort appropriately. This may be done using some ruled based algorithm such as actuator ganging or daisy chaining, or what this work will focus on, dynamic allocation of the control effort.

3. **Low level controls**. The control allocator may act directly upon the actuators, but usually it produces target values for the low level controllers which attempt to meet these control targets and/or adds robustness by limiting actuator targets when the control allocator requests more control effort than the system is physically capable of supporting. Example of this is traction control which limits wheel or axle torques when the tires have reached their maximum longitudinal force capability.

3-2 Reference generation and high level controls

On the passenger vehicle, inputs to the VSC system are either human generated, usually through the steering wheel angle δ_f and generalized acceleration $\bar{a_x}$ derived from throttle position or brake pedal pressure/force. Similar inputs may be requested by an on-board autonomous system. These inputs are then interpreted by the high level controller and internally translated to reference vehicle states.

3-2-1 Reference generation

An often used approach is assuming that the driver desires a linear yaw response from the vehicle [23, 40]. Some means of translating steering angle input to yaw rate is required. This may be done in a number of ways such as steady state yaw rate gains derived from a bicycle model or using a full linear bicycle model of the form:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \tag{3-1}$$

where $\mathbf{x} = \begin{bmatrix} U_y \\ \dot{\psi} \end{bmatrix}$ and $\mathbf{u} = \delta_f$. The matrix $\mathbf{A}$ and $\mathbf{B}$:

$$\mathbf{A} = \begin{bmatrix} \frac{-C_f - C_r}{mU_x} & \frac{-aC_f + bC_r}{mU_x} - U_x \\ \frac{-aC_f + bC_r}{I_{zz}U_x} & \frac{-a^2C_f - b^2C_r}{I_{zz}U_x} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \frac{C_f}{mU_x} \\ \frac{aC_f}{I_{zz}} \end{bmatrix} \tag{3-2}$$

where C_f and C_r are double the respective tire cornering stiffness and δ_f the steering wheel angle. The outputs of this model are a reference yaw rate and sideslip:

$$\mathbf{x}_{ref} = \begin{bmatrix} U_{y,ref} \\ \dot{\psi}_{ref} \end{bmatrix}$$

Given that high quality measurements of the lateral velocity or sideslip are usually lacking in commercially available vehicles, the reference lateral velocity $U_{y,ref}$ is ignored and only $\dot{\psi}_{ref}$ is used for feedback control.

3-2-2 Controllability study

In order to control the vehicle it can be shown that regulating just the yaw rate using a yaw moment is sufficient to stabilize the vehicle, even if the vehicle itself is in an unstable state. The formulation below is set in a somewhat contrived state, where the virtual control input is the yaw moment *acceleration*. This will allow for the investigation of controllability of the system in the face of absolute actuator limits and actuator rate limits:

$$\begin{bmatrix} \dot{\beta} \\ \dot{r} \\ \dot{M}_z \\ \ddot{M}_z \end{bmatrix} = \begin{bmatrix} \frac{-C_f - C_r}{mV_x} & \frac{-L_fC_f + L_rC_r}{mV_x^2} - 1 & 0 & 0 \\ \frac{-L_fC_f + L_rC_r}{I_{zz}} & \frac{-L_f^2C_f - L_r^2C_r}{I_{zz}V_x} & \frac{1}{I_{zz}} & 0 \\ 0 & 0 & 0 & \rho_1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \beta \\ r \\ M_z \\ \dot{M}_z \end{bmatrix} + \begin{bmatrix} \frac{C_f}{mV_x} \\ \frac{L_fC_f}{I_{zz}} \\ 0 \\ 0 \end{bmatrix} \delta_f + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \rho_2 \end{bmatrix} \ddot{v} \tag{3-3}$$

with $\rho_1, \rho_2 \in \mathbb{R}$ the gains of the input $\ddot{v}$.

Controllability of the system is maintained when the Kalman matrix K is full rank:

$$K = (B \quad AB \quad A^2B \quad A^3B) \tag{3-4}$$

For this system, with yaw moment acceleration $\rho_2 \cdot \ddot{v}$ as input:

$$K = \begin{bmatrix} 0 & 0 & 0 & \left(\frac{-L_fC_f + L_rC_r}{mV_x^2} - 1\right) \cdot \frac{1}{I_{zz}} \cdot \rho_1 \cdot \rho_2 \\ 0 & 0 & \frac{1}{I_{zz}} \cdot \rho_1 \cdot \rho_2 & \frac{-L_f^2C_f - L_r^2C_r}{I_{zz}V_x} \cdot \frac{1}{I_{zz}} \cdot \rho_1 \cdot \rho_2 \\ 0 & \rho_1 \cdot \rho_2 & 0 & 0 \\ \rho_2 & 0 & 0 & 0 \end{bmatrix} \tag{3-5}$$

If K is full rank, the system is controllable. K will be full rank except for the following four cases:

1. $C_f = 0, C_r = 0$. This can occur when the tires are fully saturated, i.e. when the vehicle is in a severe sideways slide. Therefore, the high level controls should try to keep the vehicle away from such an extreme state.
2. $\rho_1 = 0$ or $\rho_2 = 0$. This value goes to zero when the controller saturates. Saturation occurs when there is no further increase or decrease in yaw moment possible or the ability to increase or decrease the yaw moment is limited. This situation will lead to a complete loss of controllability, which is obviously undesirable and potentially dangerous if the system is in an unstable state.
3. $L_fC_f - L_rC_r = mV_x^2$. This condition only occurs for a single, low velocity for the vehicle under consideration, hence it will not be taken into further consideration in the rest of this thesis.

From the analysis it can be concluded that as long as there is sufficient yaw moment authority and yaw moment bandwidth, the VSC will be able to control the system, and *it can drive the system states to wherever the driver might want them to be.* It is interesting to note this, because that means that the main performance limitation for lateral vehicle handling is in the ability of the system to generate yaw moments. As long as the control allocator does not saturate, the control designer can basically make the vehicle do whatever is desired. The yaw moment that may be generated is of course physically limited by both the maximum

capabilities of the tires and the maximum capabilities of the actuators. A good control allocator should therefore be able to handle both and find the right actuator efforts to maximally exploit the tire capabilities whilst working around actuation limits.

3-2-3 Yaw rate controller

Given the yaw rate reference, a controller can be designed that takes in the system states, and attempts to regulate these using the virtual control input v, a target yaw moment. Many different control strategies have been applied in literature to this high level control task. However, as the design of the high level control is not the scope of this thesis, a simple PID controller was selected for this task. It is easy to tune, achieves acceptable levels of performance and it behaves predictably, making it very suitable for investigating the control allocation part of the system:

$$\begin{aligned} \dot{\psi}_{ref} &= \text{sat}(\dot{\psi}_{ref}, -\dot{\psi}_{max}, \dot{\psi}_{max}) \\ \dot{\psi}_{error} &= \dot{\psi}_{ref} - \dot{\psi} \\ M_{z,ref} &= K_{p,yaw}\dot{\psi}_{error} + K_{i,yaw}\int_0^t \dot{\psi}_{error}dt + K_{d,yaw}\ddot{\psi}_{error} \end{aligned} \tag{3-6}$$

where $K_{p,yaw}$, $K_{i,yaw}$ and $K_{d,yaw}$ the proportional, integral and derivative gains respectively, $M_{z,ref}$ the yaw moment request to the control allocator and sat() the saturation function that saturates $\dot{\psi}_{ref}$ between lower and upper bound. These lower and upper bounds were imposed to prevent the driver from requesting non-achievable yaw rates from the vehicle. Given that this VSC system is mostly targeted for passenger vehicles where safety comes first and performance second, some sensible bounds on what the driver is allowed to request in terms of references is imposed. The limits are based on the maximum tire capabilities. How maximum tire capabilities relate to maximally achievable steady state yaw rate can be approximated as [36]:

$$\dot{\psi}_{max} = \mu_s \cdot g/U_x$$

where μ_s the estimate of the surface friction coefficient, g the gravitational constant and U_x the longitudinal velocity. ψ_{max} is then used to provide upper and lower bounds to driver requested yaw rates.

3-3 Control allocator

Remembering Eq. (1-1) and Figure 3-1, the main input to the control allocator is the virtual control input $\mathbf{v}_{ref}$, generated by the high level controller. It is the control allocator's primary goal to ensure this virtual control input is actually exerted on the system by properly distributing the actual control effort $\mathbf{u}$.

Two virtual control inputs were selected for this work, one focusing on the lateral control of the vehicle and one on the longitudinal control of the vehicle. Lateral control is provided by the high level controller through the yaw moment $M_{z,ref}$. However, $M_{z,ref}$ cannot be applied to the vehicle directly. The mapping from $\mathbf{u}$ to $M_{z,ref}$ is non-unique as will be shown in Section 4-3-2, which gives rise to a control allocation problem. Therefore, $M_{z,ref}$ was

selected as one of the virtual control inputs of the control allocator. Furthermore, to limit intrusiveness to the driver, deviations from the desired longitudinal acceleration $F_{x,ref}$ are also penalized. This forces the control allocator to attempt to decouple the effects of lateral control through yaw moment $M_{z,ref}$ on the longitudinal control. Finally, as certain actuators will be preferred over other actuators in achieving either task, penalties on the individual actuator efforts are applied. This leads to the following input and output definitions for the control allocator and the optimization problem formulation:

$$\begin{aligned} \min_{u} \quad & \mathbf{R} \parallel \mathbf{v}(u,x,t) - \mathbf{v}_{ref}(t) \parallel^{n_1} + \mathbf{Q} \parallel (\mathbf{u} - \mathbf{u}_{desired}(t)) \parallel^{n_2} \\ \text{subj. to} \quad & \mathbf{u} - \mathbf{u}_{lim}^{max}(t) \leq 0 \\ & -\mathbf{u} - \mathbf{u}_{lim}^{min}(t) \leq 0 \\ & \dot{\mathbf{u}} - \dot{\mathbf{u}}_{lim}^{max}(t) \leq 0 \\ & -\dot{\mathbf{u}} - \dot{\mathbf{u}}_{lim}^{min}(t) \leq 0 \\ & \{n_1, n_2\} \in 1, 2, \ldots \end{aligned} \tag{3-7}$$

where n_i the norm (usually 1 or 2), $\mathbf{v}(u,x,t) \in \mathbb{R}^m, \mathbf{u} \in \mathbb{R}^n$, $\mathbf{R} \in \mathbb{R}^{m \times m}$ and $\mathbf{Q} \in \mathbb{R}^{n \times n}$ diagonal weighting matrices to penalize tracking errors and actuator usage respectively. Furthermore, $\mathbf{v}_{ref}(t), \mathbf{u}_{desired}(t), \mathbf{u}_{lim}^{max}(t), \mathbf{u}_{lim}^{min}(t), \dot{\mathbf{u}}_{lim}^{max}(t), \dot{\mathbf{u}}_{lim}^{min}(t)$ are all time-varying values, which means both the unconstrained and the constrained optimum will be time-varying as well. The result of the optimization is a set of actuator targets $\mathbf{u}$ which serve as the low level control inputs.

The optimization inputs, outputs and constraints are defined as:

$$\begin{aligned} \mathbf{v}_{ref} &= \begin{bmatrix} M_{z,ref} & F_{x,ref} \end{bmatrix}^T, \quad \tilde{\mathbf{x}} = \begin{bmatrix} \delta_f & \mu & v_x & v_y & r & a_x & a_y \end{bmatrix}^T \\ \mathbf{u} &= \begin{bmatrix} T_{m,f} & T_{m,r} & T_{br,fl} & T_{br,fr} & T_{br,rl} & T_{br,rr} \end{bmatrix}^T \\ \mathbf{u}_{desired} &= \begin{bmatrix} T_{m,f}^{desired} & T_{m,r}^{desired} & T_{br,fl}^{desired} & T_{br,fr}^{desired} & T_{br,rl}^{desired} & T_{br,rr}^{desired} \end{bmatrix}^T \\ \mathbf{u}_{lim}^{min/max} &= \begin{bmatrix} T_{m,f}^{min/max} & T_{m,r}^{min/max} & T_{br,fl}^{min/max} & T_{br,fr}^{min/max} & T_{br,rl}^{min/max} & T_{br,rr}^{min/max} \end{bmatrix}^T \\ \dot{\mathbf{u}}_{lim}^{min/max} &= \begin{bmatrix} \dot{T}_{m,f}^{min/max} & \dot{T}_{m,r}^{min/max} & \dot{T}_{br,fl}^{min/max} & \dot{T}_{br,fr}^{min/max} & \dot{T}_{br,rl}^{min/max} & \dot{T}_{br,rr}^{min/max} \end{bmatrix}^T \end{aligned} \tag{3-8}$$

where the virtual control inputs are: $M_{z,ref}$, the requested yaw moment and $F_{x,ref}$, the desired longitudinal force. The outputs of the control allocator are: $T_{m,f/r}$, the front and rear motor torques and $T_{br,ii}$, the brake torques for the brakes on all four corners of the vehicle. The desired steady state position of those actuators is given by $\mathbf{u}_{desired}$. $\mathbf{u}_{desired}$ is included in the optimization problem to prevent penalties from being applied to driver generated commands. For example, when the driver attempts to accelerate using a certain predetermined $T_{m,f}$ and $T_{m,r}$ these torques should be used without penalty within the optimization as these inputs represent the driver intent. It is *deviations* from the driver intent that should be penalized, not the actual actuator commands $\mathbf{u}$.

Constraints are imposed as well, both absolute and rate limits. $\dot{T}_{m,f/r}^{min/max}$ and $\dot{T}_{br,ii}^{min/max}$ are the maximum allowable upward and downward rates of the actuators at a certain time

instance and $T_{m,f/r}^{min/max}$ and $T_{br,ii}^{min/max}$ are the maximum and minimal allowable actuator efforts. Furthermore, the assumption is made that no forward prediction is performed in the control allocator. The proposed architecture is flexible enough to allow incorporation of forward prediction, but this was left out to narrow the scope of the thesis. The detailed design of the control allocator will be elaborated upon in Chapter 4.

3-4 Low level control

For each of the available actuators a low level controller was designed and implemented. The goal of these low level controllers is to regulate wheel slip in case the control allocator is requesting more torque than the tire is capable of handling longitudinally due to internal model or estimation errors, both for drive torques and braking torques. The control allocator that was designed has internal estimates of the tire capabilities, and expects a certain amount of wheel slip to correspond to the torque it is applying to that wheel. However, in the case of model or estimation errors, this feedforward approach can't be relied on and might lead to excessive wheel flares or wheel locking, given that the control allocator is not set up provide closed loop control of wheel slip. Therefore these low level controllers were implemented to provide closed loop control of wheel slip to add robustness to control allocator performance. Given that the desire is to give maximal control authority to the control allocator, these low level controls were set up to be minimally intrusive which was achieved by setting high activation thresholds. This ensures they only act when strictly necessary.

For the motors a simple traction control scheme was developed based on a PI regulator that engages when one of the wheels on that axle exceeds a certain slip threshold λ_{lim}:

$$\begin{aligned} \lambda_{error} &= 0 \quad \text{for } \|\lambda\| < \lambda_{lim} \\ \lambda_{error} &= \lambda - \text{sign}(\lambda) \cdot \lambda_{lim} \quad \text{for } \|\lambda\| \geq \lambda_{lim} \\ T_{m,f/r}^{min/max} &= K_{p,TC}\lambda_{error} + K_{i,TC}\int_0^t \lambda_{error} dt \end{aligned}$$

An example of traction control performance is shown in Figure 3-2. A similar approach was used for the brake low level controls where some simple heuristics were applied that reduce brake torque when the wheel enters into deep slip to prevent the wheels from fully locking up.

All state information required to calculated slip for each tire was assumed to be known without noise or estimation errors.

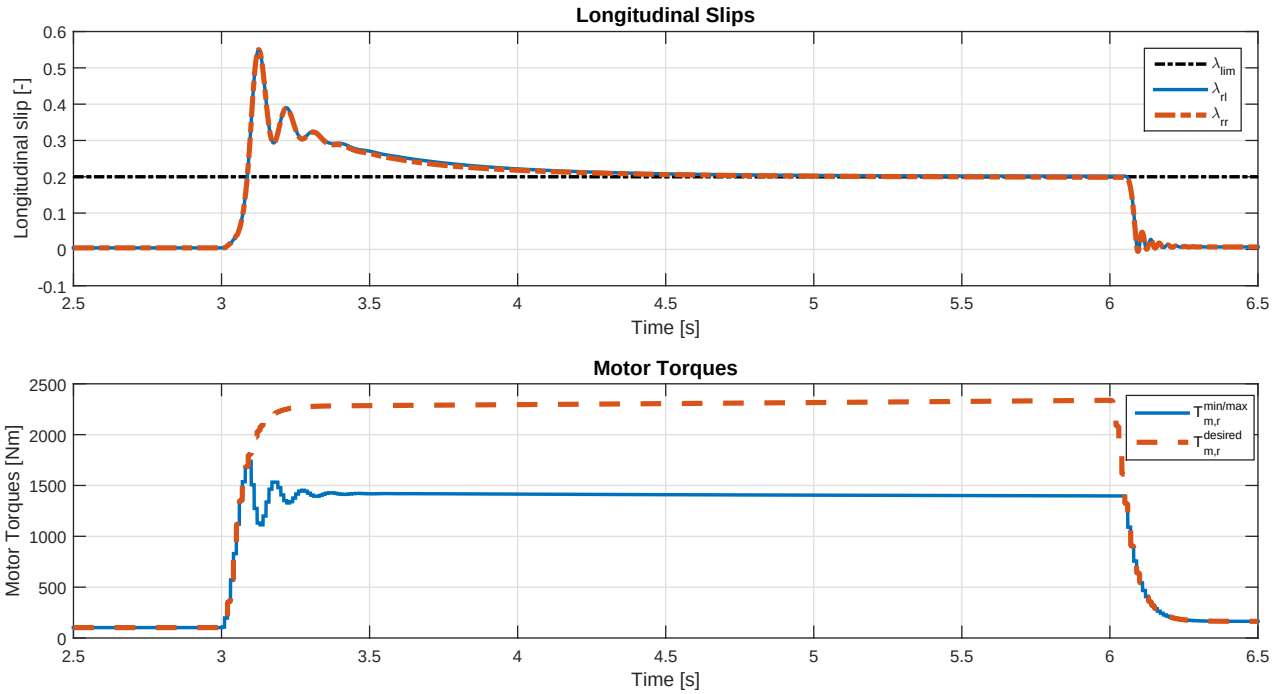

Figure 3-2: Traction control performance. Straight rolling launch on snow, starting at 6 [m/s], with rear motor torque only. Slip limit λ_{lim} was set at 20% longitudinal slip

3-5 Conclusion

The VSC system was successfully broken into separate modules, each with distinct functionality. Both systems that are not the direct subject of study, the high and low level controllers, were kept simple to ensure predictable performance. The control allocation algorithm was posed as a general optimization problem that should work for any actuator layout and gives the control designer full control over trade-off's between actuator usage and the coupling between yaw moment tracking and longitudinal control. The proposed formulation is not unique in any sense as there has been plenty of work that has used similar formulations, e.g. [3, 38]. However, the true difficulty here is not the formulation, but finding an optimization technique that is real-time feasible to solve the proposed optimization problem. One such technique, hybrid steepest descent optimization is where the following chapters will focus on and its application to this particular formulation of the control allocation problem.

Chapter 4

Dynamic Control Allocation using Hybrid Steepest Descent Optimization

This chapter covers the design of the control allocator. The rationale behind choosing the hybrid steepest descent optimization method will be presented, after which the theoretical background of the method will be covered. Two modifications to the optimization method are proposed that enable accelerated convergence and reduced chatter. It is followed by an illustrative example, demonstrating how the optimization method works and how the modifications improve performance. The second part of this chapter will explain how the control allocation problem formulated in Chapter 3 may be converted to make it suitable for hybrid steepest descent optimization. The chapter closes with the design of a simple static benchmark control allocator that will be used to provide reference performance in Chapter 5.

4-1 Hybrid steepest descent optimization

The solution to the control allocation problem using optimization in passenger vehicles is not trivial, as indicated by the large variety of control architectures and methodologies that have been proposed over the years. The problem is challenging for a number of reasons:

1. The vehicle and the actuators acting on it usually show nonlinear behavior during limit-handling, mostly caused by the nonlinear force producing characteristics of tires. Therefore, some method of handling these non-linearities is required, either by including nonlinear models or sacrificing performance to make the system robust to these nonlinearities.

2. The presence of time-varying constraints, imposed by either the actuators (saturation limits or rate limits) or tire limits, both of which can rapidly change, appear or disappear. This limits the number of optimization methods that may be used as not all optimization methods are capable of handling time-varying constraints.

3. The computational power available on current generation ECUs is limited, requiring computationally efficient optimization schemes.

These challenges provide the requirements that drive the selection for a suitable optimization method for this problem. Methods that make use of precomputed laws are excluded by 2, as these methods cannot handle time-varying constraints. Furthermore, methods that apply repeated optimization tend to be hard to solve efficiently when constraints are present, as some form iteration is required to find the optimum. However, as was also stated [3], it is not strictly necessary for the optimization method to find the exact optimum at every time instance, as long as it is close to the optimal point and the solution moves towards the optimal point sufficiently quickly over time. This suggests using an update law approach to the optimization problem which has the potential of being computationally efficient as it does not require iterations at every time step to find the optimum. This makes the hybrid steepest descent method a suitable candidate for this problem. The method makes use of update laws and it efficiently handles (time-varying) constraints, fulfilling both requirement 2. and 3. The only problem that remains is 1, the non-linearity of the problem. The solution to this problem will be covered in detail in the latter part of this chapter.

4-1-1 Hybrid steepest descent optimization: background

The hybrid steepest descent method assumes a convex cost function with convex constraints:

$$\begin{aligned} \min_{u} \quad & q(u(t)) \\ \text{subj. to} \quad & g(u(t)) \leq 0 \end{aligned} \tag{4-1}$$

with $u(t) \in \mathbb{R}^n, q : \mathbb{R}^n \to \mathbb{R}$ and $g_i(u(t)) : \mathbb{R}^n \to \mathbb{R}^m$ for $i = 1 \dots m$ and $q(u(t)), g(u(t))$ differentiable convex functions and where the optimal point is denoted as q^*.

This optimization problem may be solved by setting it up as a hybrid dynamical system of the form:

$$\dot{u}(t) = f(u(t)) \tag{4-2}$$

which has the following properties in continuous time:

1. for some $u(\bar{t})$ not in the feasible set at time $\bar{t}$, the trajectory $u(t)$ returns to the feasible set, for example $\exists t_f > \bar{t}$ s.t. $g(x(t_f)) \leq 0$
2. the trajectory $u(t)$ remains in the feasible set as soon as $u(t_f)$ is in this set
3. when $u(t)$ in the feasible set, the trajectory of $u(t)$ decreases the cost function $q(u(t))$ at all time until it reaches the optimal point q^*

Whilst in the feasible set, an efficient way of reducing the cost function is through the use of the gradient descent method. Different directions are also possible, such as Newton's direction, but these fall outside the scope of this thesis. The novelty of the HSDO is based on the way the constraints are handled. In earlier works, the gradient projection method was used to

handle the constraints. However, the authors of [3] note that this method is computationally demanding and therefore unattractive for the kind of fast-real time applications these optimization techniques are intended for. What is proposed in [3] is using some form of barrier generated by constraint violations. The barrier is implemented by forcing $u(t)$ to move back towards the feasible set using the gradient of the violated constraint(s). This ensures that there are descent directions outside the feasible set, which is a very useful property when dealing with time-varying constraints as will be later demonstrated. If the trajectory is on the boundary of the feasible set, it will alternately be pushed by the gradient of $q(u(t))$ and the active constraint $g_i(u(t))$. This will make $u(t)$ naturally slide along the constraint in the direction of the optimal point q^*.

The hybrid feedback law this results in:

$$f(u) = \begin{cases} -\nabla q(u) & \text{if } g_j(u) \leq 0 \forall j \\ -\sum_{i \in C(u)} \nabla g_j(u) & \text{if } j : g_j(u) > 0 \end{cases}$$

with $C(u) = \{l : g_l(u) \geq 0\}$ This may alternatively be written as:

$$\dot{u}(t) = -\gamma_0 \nabla q(u) - \sum_{i=1}^{m} \gamma_i(u) \nabla g_i(u) \tag{4-3}$$

for $\gamma_j(u) \geq 0, \quad j \in \{0, \ldots, m\}$ where:

- $\gamma_0 = 1 \quad \text{if } g_j(u) \leq 0 \quad \forall j \text{ or } \gamma_0 = 0 \quad \text{if } j : g_j(u) > 0$
- $\gamma_j = 1 \quad j \in C(u)$ and 0 otherwise

Then finally the trajectory of $u(t)$ is obtained by integrating Eq. (4-3), either in continuous time:

$$u(t) = \int_0^t \dot{u}(t) dt$$

or in discrete time, using forward Euler integration for example:

$$u(n+1) = \dot{u}(n) \cdot t_s + u(n)$$

where t_s the sample time.

4-1-2 Discrete time implementation

Given that the system will implemented in discrete time, the behavior of this algorithm will be covered in some extra detail. The discretization has some potentially negative effects on the smoothness of the outputs of the optimization algorithm, as the discrete time implementation means it will move with discrete steps towards the optimum over time, where the step size towards the optimum or away from the infeasible region in case of constraint violation is determined by $\dot{u}(n) \cdot t_s$. Due to the hybrid approach, the discretization means the trajectory of $u(n)$ is expected to show some chatter when moving along constraints, where the gradient of

$\dot{u}(n)$ will be alternating between $\nabla q(u)$ and $\nabla g_j(u)$. Given that the outputs of the optimization are expected to be actuator efforts for the application under consideration in this thesis, keeping the outputs $u(n)$ relatively smooth will be important to prevent unnecessary actuator movement. However, step sizes should also not be too small, otherwise the trajectory of $u(n)$ might converge on the optimum too slowly, also leading to deteriorated performance. As the step size in $\dot{u}(n) \cdot t_s$ is set by the magnitude of the gradients $\|\nabla q(u)\|$, $\|\nabla g_j(u)\|$ and sample time t_s, it will be important to understand how these should be chosen to to prevent unacceptable levels of chatter whilst still ensuring that the system converges towards the optimum sufficiently quickly. An example of a discrete time hybrid steepest optimization problem will be covered in Section 4-2 to illustrate this behavior and the effect of these trade-off's.

4-1-3 Modifications to $\nabla q(u)$, $\nabla g_j(u)$ for improved convergence and reduced chatter

For the subsequent discussion the sample time t_s is assumed to be a fixed parameter determined by the firmware architecture the optimization will be running in. This reduces the scope to finding a good trade-off between $\nabla q(u)$ and $\nabla g_j(u)$ for fast convergence and limited chatter. The hybrid feedback law is modified with extra parameters α, β_j for the following analysis:

$$f(u) = \begin{cases} -\alpha \nabla q(u) & \text{if } g_j(u) \leq 0 \forall j \\ -\sum_{i \in C(u)} \beta_j \nabla g_j(u) & \text{if } j : g_j(u) > 0 \end{cases} \tag{4-4}$$

where $\alpha, \beta_j > 0$ scalar values that may be used to scale the magnitude of the respective gradients.

To ensure fast convergence of $u(n)$ it will be desirable to use large values of α and β_j, where especially β_j is important. This is intuitive, because when $u(n)$ is in the infeasible region, $u(n)$ is not descending towards the optimum q^*. Therefore, if $u(n)$ spends a large number of time steps $\bar{n}$ in the infeasible region, convergence speed towards q^* is expected to suffer. This situation is expected to happen if during a sliding mode $\|\alpha \nabla q(u)\| \gg \|\beta_j \nabla g_j(u)\|$. In this case a single step in the direction $\alpha \nabla q(u)$ will push $u(n)$ relatively far into the infeasible region, requiring a large number of steps in the direction of $\nabla g_j(u)$ to return $u(n)$ to the feasible region again. This highlights the importance of ensuring that both the magnitude of $\alpha \nabla q(u)$ and $\beta_j \nabla g_j(u)$ are approximately the same to prevent this kind of behavior. To ensure that $\|\alpha \nabla q(u)\| \approx \|\beta_j \nabla g_j(u)\|$ a modification is proposed to the algorithm that should enforce this by scaling β_j appropriately:

$$\beta_j = \|\nabla g_j(u)\|^{-1} \|g_j(u)\|$$

Where $\|g_j(u)\|$ denotes the magnitude of the constraint violation. This method basically seeks to solve the constraint violation of $g_j(u)$ in a single step, projecting it back to the border of the constraint using its own gradient.

However, this modification does not provide any indication of how large α should be chosen. Intuitively α should still be as large as possible for fast convergence, but this should not go at the cost of unacceptable levels of chatter when sliding along a constraint. This trade-off can be made less critical however if an extra step is allowed within the optimization. The extra step

basically enforces the optimization algorithm to only output $\dot{u}(n)$ if $g_j(u) \leq 0 \forall j$. This means that the optimization will always start with a feasible step determined by $\dot{u}(n) = -\alpha \nabla q(u) t_s$. If this step results in infeasibility, such as is expected close to constraints, extra iterations are performed during the same time step n that force $u(n)$ back to the feasible region, with the direction of the extra iteration(s) determined by $\sum_{i \in C(u)} \beta_j \nabla g_j(u)$. When β_j is properly scaled the extra iteration steps required to achieve this are expected to be minimal, in the order of 1 or 2 for relatively simple problems which should mean minimal increase of computation time.

It must be noted that these modifications void the mathematical proofs of convergence and asymptotic stability presented for the hybrid steepest descent method in [3]. However, given that fundamentally nothing has changed to the way the algorithm works it is expected that it is possible, with some extra work, to show that the modifications don't impact convergence and asymptotic stability. This however will be left for future work.

4-2 Example of hybrid steepest descent optimization

To illustrate the effect of the proposed algorithm enhancements and the general behavior of the hybrid steepest descent optimization method a simple example is investigated first.

4-2-1 Cost function and constraints

The example cost function is taken to be a two-dimensional quadratic cost function. The constraints are chosen similar to the formulation in Eq. (4-22):

$$\begin{aligned} q(x_1, x_2) &= \alpha(x_1(n)^2 + x_2(n)^2) \\ \text{subj. to} \quad g_1(x_1): &\quad \beta_1(x_1(n) - 5) \leq 0 \\ g_2(x_2): &\quad \beta_2(-x_2(n) + 3) \leq 0 \end{aligned} \tag{4-5}$$

where α and β_j are scaling factors that determine the size of the descent steps.

Translating the equations above to the form of Eq. (4-3) and converting to discrete time format:

$$\nabla q = \alpha \begin{bmatrix} 2x_1(n) \\ 2x_2(n) \end{bmatrix}, \quad \nabla g_1 = \beta_1 \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \nabla g_2 = \beta_2 \begin{bmatrix} 0 \\ -1 \end{bmatrix}$$

$$\dot{x}(n) = -\gamma_0 \nabla q - \gamma_1 \nabla g_1 - \gamma_2 \nabla g_2 \tag{4-6}$$

$$x(n+1) = x(n) + \dot{x}(n) t_s \tag{4-7}$$

with t_s the sample time, which was chosen as 1 sec for this example for simplicity.

4-2-2 Example case

As mentioned in Section 4-1-3, the values for α and β_j determine how fast the trajectory of $x(n)$ moves towards the optimum and how high the level of chatter will be around constraints.

An example illustrating the effects of this trade-off is shown in Figure 4-1 followed by a discussion. The initial point in this example is chosen at $[10, 10]$, which is well into the infeasible region. The trajectory first has to move back to the feasible region. After reaching the feasible region it slides towards the optimal point using the steepest descent direction. At some point it hits constraint g_2 after which the trajectory enters a sliding mode, alternately moving towards the optimum and moving back towards the feasible region, asymptotically approaching the optimum at $[0, 3]$.

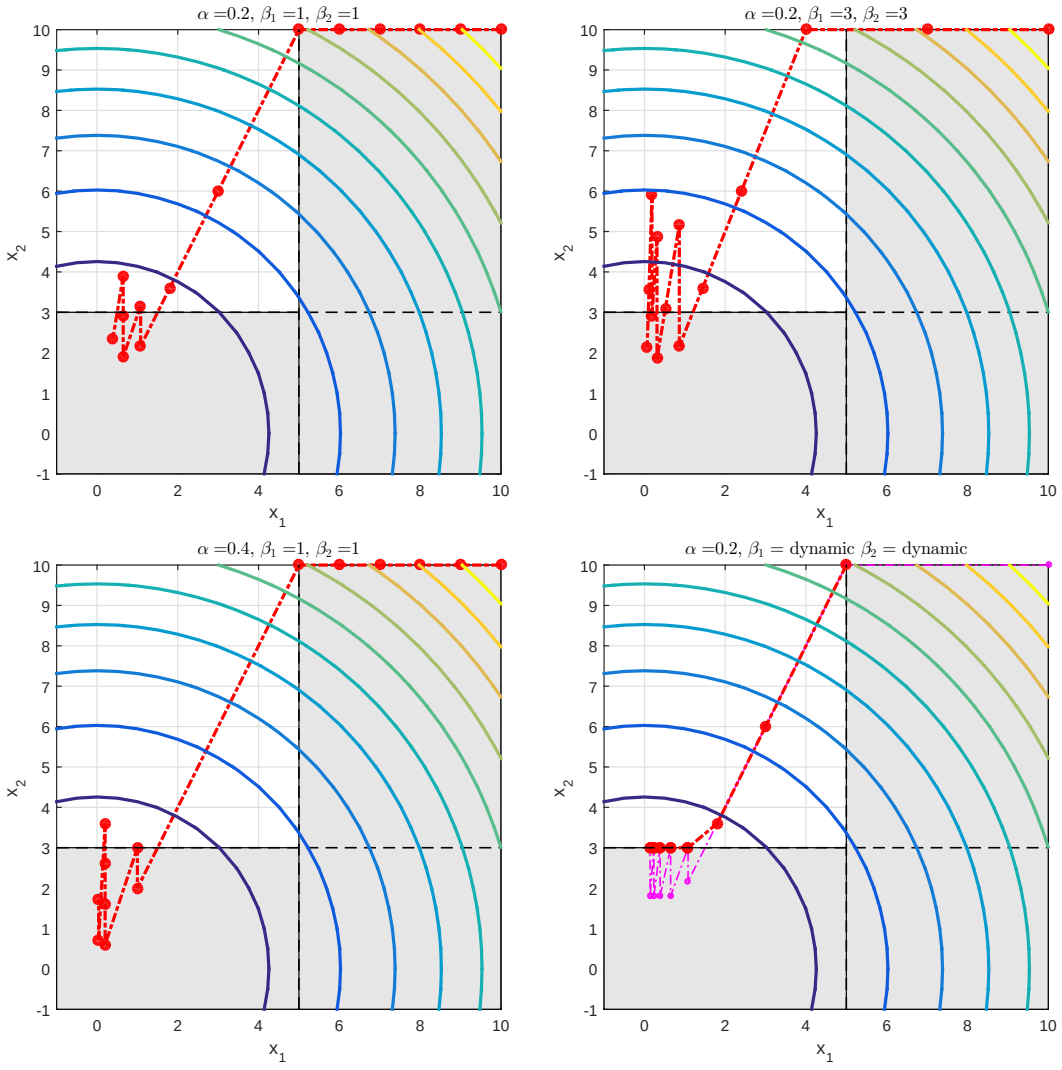

Figure 4-1: Four examples demonstrating the effect of choices for α and β. The initial point is chosen at $[10, 10]$ and the constrained optimal point is at $[0, 3]$. The contours of the cost function are shown, as well as the constraints in black and the infeasible region they define is grayed out. The steps the trajectory of x traverses towards the constrained optimum is shown with red dots. The method with dynamically scaled β_j also applied the extra iteration steps to force $x(n)$ to be feasible at every step. The extra steps that were required to achieve this are shown with small purple dots

The impact for the choice of larger or smaller values of α and β is visually obvious. However, to get a better sense for the relative performance two metrics were compared:

1. the number of steps required to get within a radius of 1 from the constrained optimal point. If there is an extra iteration every time-step, such as used by the dynamic scaling method, these iterations are *not* counted as extra steps
2. the largest deviation around the optimal point after it has entered the radius of 1 from the optimal point

These metrics are chosen somewhat arbitrarily, but they clearly demonstrate the effect between the different values chosen for α and β:

1. $\alpha = 0.2$, $\beta = 1$ leads to a long time spent in the infeasible region, resulting in a large step count to reach the optimal point (≈ 10 steps total) but the amount of chatter is relatively small (≈ 1)
2. $\alpha = 0.2$, $\beta = 3$ gets the trajectory out of the infeasible region much more quickly (≈ 5 steps total), but the high value of β leads to a large amount of chatter (≈ 3)
3. $\alpha = 0.4$, $\beta = 1$ means the trajectory still spends a long time in the infeasible region, but converges rapidly when in the feasible region (≈ 9 steps total) but at the cost of relatively large amount of chatter (≈ 2)
4. $\alpha = 0.2$, $\beta =$ *dynamic* uses both enhancements covered in Section 4-1-3, moving the trajectory out of the infeasible region instantly and then converges towards the optimum at a normal speed (≈ 3 steps), whilst the amount of chatter is 0 due to the extra iteration steps that force feasibility of $x(n)$

The approach with dynamically scaled β_j and extra iteration step clearly show improved performance over the other methods that use the unmodified formulation. Chatter is completely removed for this example and convergence is by far the fastest.

4-3 Hybrid steepest descent optimization applied to the control allocation problem

The proposed formulation for the control allocation problem is re-iterated below:

$$\begin{aligned} \min_{u} \quad & \mathbf{R} \parallel \mathbf{v}(u,x,t) - \mathbf{v}_{ref}(t) \parallel^{n_1} + \mathbf{Q} \parallel (\mathbf{u} - \mathbf{u}_{desired}(t)) \parallel^{n_2} \\ \text{subj. to} \quad & \mathbf{u} - \mathbf{u}_{lim}^{max}(t) \leq 0 \\ & -\mathbf{u} - \mathbf{u}_{lim}^{min}(t) \leq 0 \\ & \dot{\mathbf{u}} - \dot{\mathbf{u}}_{lim}^{max}(t) \leq 0 \\ & -\dot{\mathbf{u}} - \dot{\mathbf{u}}_{lim}^{min}(t) \leq 0 \\ & \{n_1, n_2\} \in 1, 2, \dots \end{aligned} \tag{4-8}$$

The formulation above can be made suitable for HSDO, provided that the right norms are selected for both parts of the cost function. Although the cost function depends convexly on $\mathbf{v}(u,x,t)$, the relation between the virtual control inputs (yaw moment M_z and longitudinal force F_x) and the real actuator inputs $\mathbf{u}$ is non-trivially modeled. This will need to be investigated more closely to obtain a computationally simple formulation for the optimization problem.

4-3-1 F_x as function of $\mathbf{u}$

Determining the total longitudinal force produced by the actuators is straightforward. Using the equation for longitudinal dynamics in Eq. (2-13) and expressing it in terms of the actuator efforts $\mathbf{u}$:

$$F_x = 1/r_w \begin{bmatrix} \cos(\delta_f) & 1 & \cos(\delta_f) & \cos(\delta_f) & 1 & 1 \end{bmatrix} \mathbf{u}$$

Or using small angle approximations $\cos(\delta_f) \approx 1$:

$$F_x = 1/r_w \sum_{i=1}^{n} u_i, \quad \mathbf{u} = \begin{bmatrix} u_1 & u_2 & \dots & u_i & \dots & u_n \end{bmatrix}^T \tag{4-9}$$

This formulation makes F_x a simple linear function of $\mathbf{u}$, which means this can be used without modification in the problem formulation.

4-3-2 M_z as function of $\mathbf{u}$

Finding the yaw moment M_z as function of $\mathbf{u}$ is unfortunately not as straightforward as finding F_x as function of $\mathbf{u}$. The yaw moment exerted by the actuators on the vehicle does not appear explicitly in the equations of Eq. (2-13). The yaw moment balance of Eq. (4-11) does provide information about the combined effect of all inputs, both controlled ($\mathbf{u}$) and uncontrolled (such as steering wheel angle δ_f) on the system but a means to separate the contributions of uncontrolled and controlled inputs will be required. The problem is that the yaw moment contribution of the uncontrolled inputs is, in part, a function of the controlled inputs and vice versa. One could attempt to solve the equations explicitly to separate the two, but this is expected to get very involved or might even be impossible. Therefore, the yaw moment exerted by the controlled inputs $\mathbf{u}$ will be inferred through a different approach which will be worked out below and is similar to the method used in [31] and [15]. The vehicle model used for this is the double-track model, where the longitudinal dynamics are removed for simplification, given that their impact on the overall results is relatively minor. Re-iterating the equations from Section 2-2-1, without longitudinal dynamics:

$$\begin{aligned} \dot{U}_y = \frac{1}{M_{car}} \Big(& (F_{x,fl} + F_{x,fr}) \sin(\delta_f) + (F_{y,fl} + F_{y,fr}) \cos(\delta_f) + \dots \\ & (F_{y,rl} + F_{y,rr}) + F_{y,disturbance} \Big) - U_x \dot{\psi} \end{aligned} \tag{4-10}$$

$$\begin{aligned} \ddot{\psi} = \frac{1}{I_{zz}} \Big(& a(F_{x,fl} + F_{x,fr}) \sin(\delta_f) + a(F_{y,fl} + F_{y,fr}) \cos(\delta_f) + d/2(F_{x,rr} - F_{x,rl}) + \dots \\ & d/2(F_{x,fr} - F_{x,fl}) \cos(\delta_f) + d/2(F_{y,fl} - F_{y,fr}) \sin(\delta_f) - b(F_{y,rl} + F_{y,rr}) \Big) \end{aligned} \tag{4-11}$$

And the longitudinal forces are assumed to be instantaneous:

$$F_{x,fl} = 1/r_w(T_{m,f}/2 - T_{br,fl}) \tag{4-12}$$
$$F_{x,fr} = 1/r_w(T_{m,f}/2 - T_{br,fr}) \tag{4-13}$$
$$F_{x,rl} = 1/r_w(T_{m,r}/2 - T_{br,rl}) \tag{4-14}$$
$$F_{x,rr} = 1/r_w(T_{m,r}/2 - T_{br,rr}) \tag{4-15}$$

Estimating yaw moment induced by controlled inputs u

The yaw moment exerted by the actuators will be defined as the difference between the yaw acceleration of the vehicle with all the actuators at their current effort $\mathbf{u}$ for a given vehicle state $\mathbf{x}_0 = \{\delta_f, U_x, U_y, \dot{\psi}, \mu_s\}$, and all the actuators set to 0 for the same $\mathbf{x}_0 = \{\delta_f, U_x, U_y, \dot{\psi}, \mu_s\}$. The latter state is referred to as the *un-actuated* state of the vehicle, with the forces corresponding to this state denoted by $F_i^{\mathbf{u}=0}$. The difference between the yaw acceleration in the actuated state and un-actuated state, $\Delta\ddot{\psi}$, when multiplied with the vehicle's yaw inertia I_{zz} is then defined as the yaw moment M_z exerted by the actuators on the system:

$$\begin{aligned} I_{zz}\Delta\ddot{\psi} = M_z = \Big(& a(F_{x,fl} - F_{x,fl}^{\mathbf{u}=0} + F_{x,fr} - F_{x,fr}^{\mathbf{u}=0})\sin(\delta_f) + ... \\ & a(F_{y,fl} - F_{y,fl}^{\mathbf{u}=0} + F_{y,fr} - F_{y,fr}^{\mathbf{u}=0})\cos(\delta_f) + ... \\ & d/2(F_{x,rr} - F_{x,rr}^{\mathbf{u}=0} - F_{x,rl} + F_{x,rl}^{\mathbf{u}=0}) + ... \\ & d/2(F_{x,fr} - F_{x,fr}^{\mathbf{u}=0} - F_{x,fl} + F_{x,fl}^{\mathbf{u}=0})\cos(\delta_f) + ... \\ & d/2(F_{y,fl} - F_{y,fl}^{\mathbf{u}=0} - F_{y,fr} + F_{y,fr}^{\mathbf{u}=0})\sin(\delta_f) - ... \\ & b(F_{y,rl} - F_{y,rl}^{\mathbf{u}=0} + F_{y,rr} - F_{y,rr}^{\mathbf{u}=0}) \Big) \end{aligned} \tag{4-16}$$

Or written alternatively, with $\Delta F_i = F_i - F_i^{\mathbf{u}=0}$:

$$\begin{aligned} I_{zz}\Delta\ddot{\psi} = M_z = \Big(& a(\Delta F_{x,fl} + \Delta F_{x,fr})\sin(\delta_f) + ... \\ & a(\Delta F_{y,fl} + \Delta F_{y,fr})\cos(\delta_f) + ... \\ & d/2(\Delta F_{x,rr} - \Delta F_{x,rl}) + ... \\ & d/2(\Delta F_{x,fr} - \Delta F_{x,fl})\cos(\delta_f) + ... \\ & d/2(\Delta F_{y,fl} - \Delta F_{y,fr})\sin(\delta_f) - ... \\ & b(\Delta F_{y,rl} + \Delta F_{y,rr}) \Big) \end{aligned} \tag{4-17}$$

This method allows us to explicitly express the amount of yaw moment M_z that is generated as a function of the change in forces ΔF_i. When the forces are mapped to actuator efforts an expression is obtained that gives yaw moment as a function of the actuators, where the mapping between force and actuator effort is described by:

$$\Delta F_i = B_e(\mathbf{x}_0, \mathbf{u})$$

where $B_e(\mathbf{x}_0, \mathbf{u})$ a nonlinear set of functions that maps $\mathbf{u}$ to ΔF_i for a given $\mathbf{x_0}$. For this system, with motor and brake torques the part of $B_e(\mathbf{x}_0, \mathbf{u})$ is easily found for the longitudinal forces $\Delta F_{x,ij}$, as it is given directly by Eq. (4-12). For the lateral forces $\Delta F_{y,ij}$ these calculations becomes more involved. An example of this is worked out in appendix A.

Yaw moment authority examples

To illustrate the proposed method with an example, an investigation was done to get an idea of the yaw moment authority of the individual actuators on a vehicle. For this example case, the states of the vehicle are chosen as such that the vehicle is in steady state when un-actuated. For a given actuator input $\mathbf{u}$, steering angle δ_f and velocity U_x the system has a unique steady state point when using tires without a peak, such as the sigmoid tire model used in this example. Next, every individual actuator is swept through its entire operation range $\{T_{min} \dots T_{max}\}$ whilst keeping all other actuators at 0 and the resultant yaw moment M_z is calculated for a given steering angle δ_f and velocity U_x. The unreachable space, the places where one of the tires is fully saturated, is left blank. Furthermore, to clarify how the differential braking torques $T_{br,f}$ and $T_{br,r}$ are defined in the equations below:

$$
\begin{aligned}
T_{br,fl} &= |T_{br,f}| && \text{for } T_{br,f} > 0, T_{br,fl} && = 0 \text{ otherwise} \\
T_{br,fr} &= |T_{br,f}| && \text{for } T_{br,f} < 0, T_{br,fr} && = 0 \text{ otherwise} \\
T_{br,rl} &= |T_{br,r}| && \text{for } T_{br,r} > 0, T_{br,rl} && = 0 \text{ otherwise} \\
T_{br,rr} &= |T_{br,r}| && \text{for } T_{br,r} < 0, T_{br,rr} && = 0 \text{ otherwise}
\end{aligned}
$$

The yaw moment analysis was performed for a surface with a friction coefficient of 1 and friction coefficient of 0.4, which resemble dry tarmac and snow respectively. The steering angle is swept from 0 to the tire saturation point. Results are shown for only one speed for each surface, but this illustrates sufficiently clearly how the yaw moment authority per actuator changes as function of surface friction coefficient and lateral acceleration.

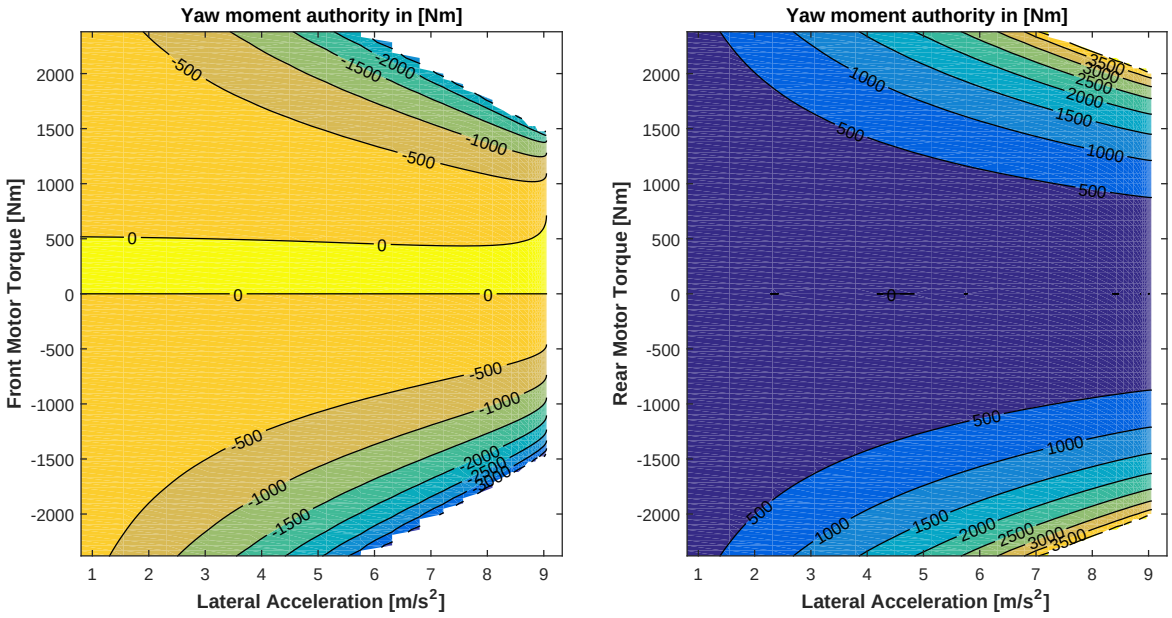

Figure 4-2: Yaw moment effectiveness in [Nm] at a μ of 1, $U_x = 25$ [m/s] for motor torques. Negative yaw moments are stabilizing yaw moments, as would be used during oversteer events and vice versa, positive yaw moments would be used to combat understeer. Note that using motor torques for yaw moment control on this surface is relatively ineffective when compared to using the differential braking (see figure below)

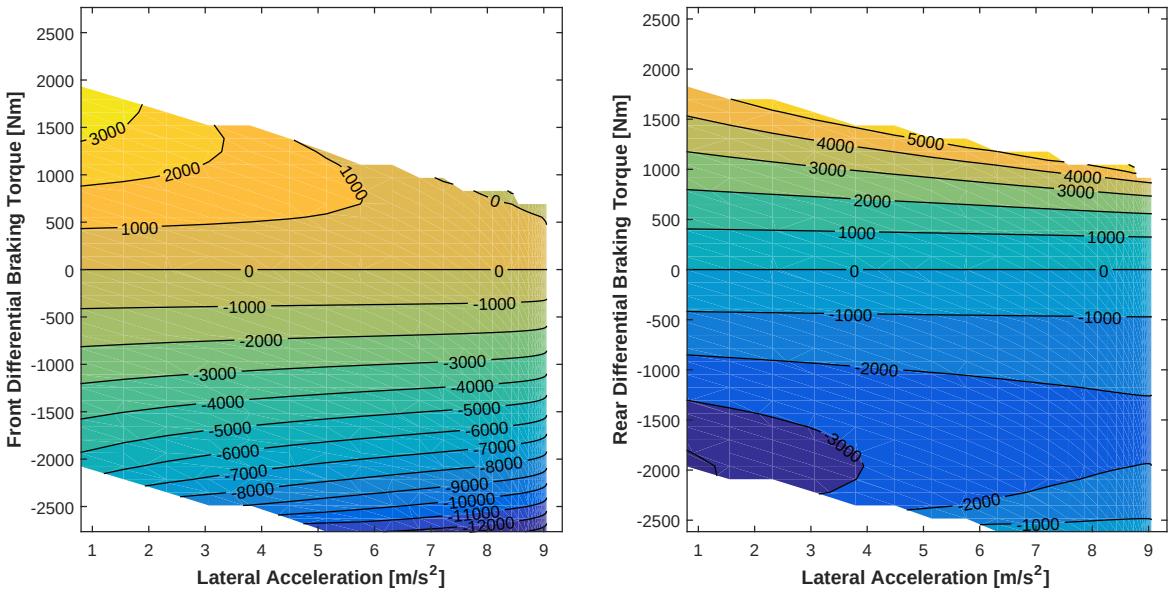

Figure 4-3: Yaw moment effectiveness in [Nm] at a μ of 1, $U_x = 25$ [m/s] for differential braking

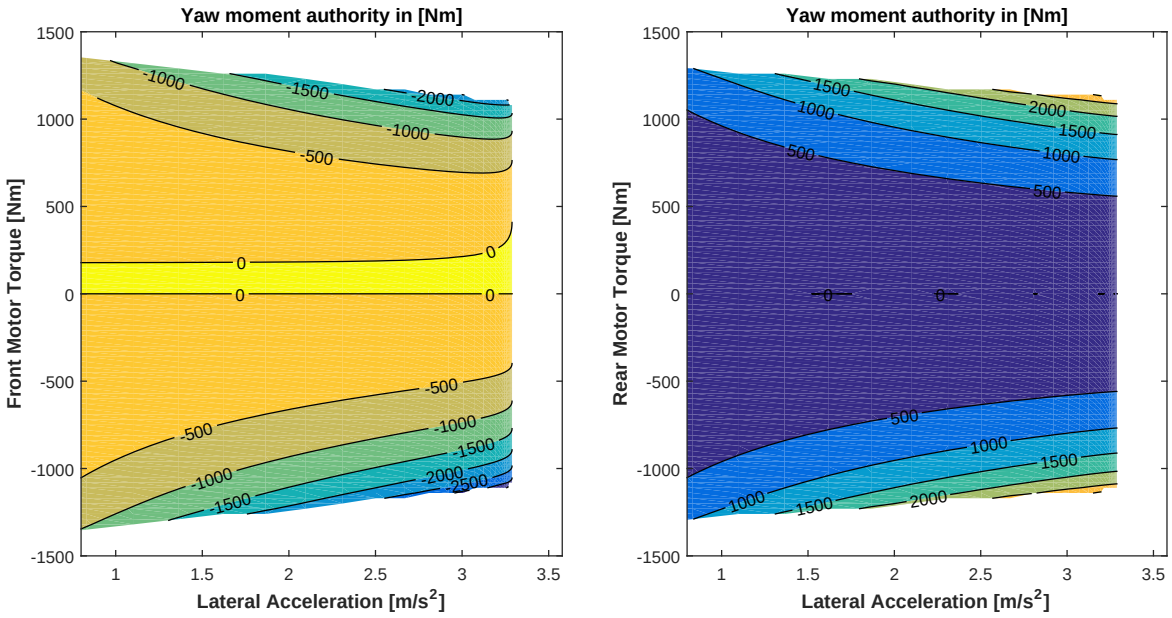

Figure 4-4: Yaw moment effectiveness in [Nm] at a μ of 0.4, $U_x = 15$ [m/s] for motor torques. On this low μ surface the motor torques have become significantly more effective compared to differential braking (see figure below)

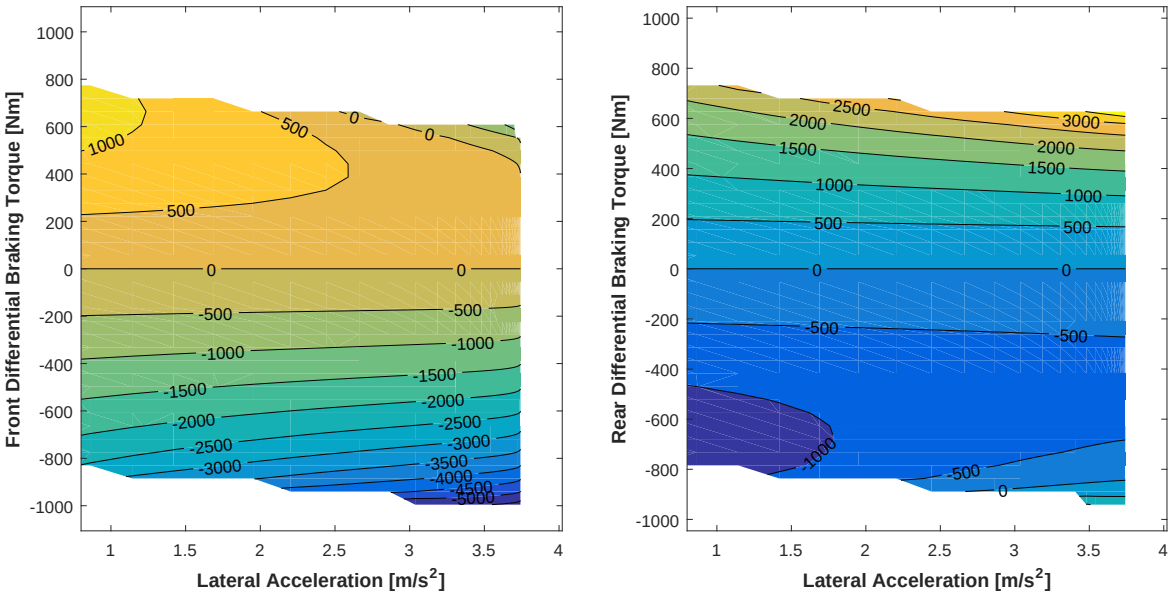

Figure 4-5: Yaw moment effectiveness in [Nm] at a μ of 0.4, $U_x = 15$ [m/s] for differential braking

Yaw moment authority examples: discussion

It is clear, even from this very limited set of plots, that the absolute effectiveness of an actuator can vary significantly as a function of the vehicle states. The same can be said of the relative effectiveness between the various actuators. Differential braking can be seen to be far more effective than motor torques for example. To provide some background as to why this is so, a short discussion of the way yaw moments are generated is presented below.

Differential braking generates a yaw moment directly through the differential longitudinal force applied across an axle, mostly determined by $d/2(F_{x,fr} - F_{x,fl})\cos(\delta_f)$ for the front brakes and $d/2(F_{x,rr} - F_{x,fr})$ for the rear brakes. This way of generating yaw moments will be referred to as a *direct* yaw moment, an effect that stays almost constant over the range of lateral accelerations because it is only dependent on track width (and for a small part, the steering angle δ_f). Besides the direct yaw moment, the differential braking force also produces an *indirect* yaw moment, which is defined as the resultant yaw moment through a reduction in lateral force due to the application of a longitudinal force. Remembering Figure 2-2 and Figure 2-3, it can be seen that the application of a longitudinal force (or equivalently, longitudinal slip) will lead to a reduction in lateral force. This reduction in lateral force will in turn produce a yaw moment. The indirect method is only capable of *reducing* lateral force, hence this is why for both positive and negative front motor torques it produces a negative (stabilizing) yaw moment, and conversely, only a positive yaw moment for the rear motor torques. Furthermore, it is important to note that this indirect effect is insignificant at low lateral accelerations, where the lateral forces are low, but as soon as the tires approach their peak lateral force capability this indirect effect becomes significant.

Because the simulated vehicle is equipped with open differentials, which distribute longitudinal force equally between left and right wheels on an axle, the motor torques do not benefit from the direct yaw moment effect and only rely on indirect yaw moments, which explains why their effectiveness only reaches appreciable values at high acceleration levels and is, in general, lower than differential braking on both surfaces.

In summary, it is clear that the yaw moment authority varies significantly between actuators, surface conditions and the state of the car. In order to achieve the design goal of 'high performance', which includes maximally exploiting actuator authority, it therefore seems necessary to include an accurate estimate of the effectiveness of each actuator in the control allocation problem.

Yaw moment effectiveness applied to the optimization problem

The equation for yaw moment M_z as shown in Eq. (4-17) will be used to estimate the yaw moment exerted on the vehicle by the actuators. However, this equation does not provide an explicit expression of M_z as function of $\mathbf{u}$, which means this formulation of M_z cannot be directly applied to HSDO. To obtain an explicit expression of M_z as function of $\mathbf{u}$ that is suitable for HSDO an online linearization of Eq. (4-17) is performed at every time step. The linearization approach is attractive as it is computationally simple and leads to a very natural problem formulation, as will be shown in the following sections.

The partial derivatives of Eq. (4-16) with respect to the control inputs $\mathbf{u}$ around the point $\mathbf{x}_0, \mathbf{u}_0$ can be analytically expressed when a simple tire model is used, such as the sigmoid tire model. The partial derivatives of $M_{z,i}$:

$$E_i(\mathbf{x}_0, \mathbf{u}_0) = \frac{\partial M_{z,i}(\mathbf{x}_0, \mathbf{u}_0)}{\partial u_i}, \quad i = \{1, \ldots, n\} \tag{4-18}$$

This produces a vector with actuator effectiveness values, which can change over time as a function of $\mathbf{x}$ and $\mathbf{u}$. Using these effectiveness values estimating the linearized yaw moment, $\bar{M}_{z,i}$ per actuator i is approximated as:

$$\bar{M}_{z,i}(x, u) = E_i(x_0, u_0) \cdot (u_i(n) - u_{i,0}) + M_z(\mathbf{x}_0, \mathbf{u}_0) \tag{4-19}$$

Therefore, the total yaw moment for a vector of control inputs $\mathbf{u}$ becomes:

$$\bar{M}_z(x, u) = \bar{E}^T \cdot (\mathbf{u} - \mathbf{u}_0) + M_z(\mathbf{x}_0, \mathbf{u}_0) \tag{4-20}$$

Where $\bar{E} \in \mathbb{R}^n$ a vector containing the effectiveness of each actuator. An example is worked out in detail in appendix A

It must be noted that by using this linearization approach, it can't be guaranteed that if an optimal point q^* is reached that this point is a global optimum. The local linear approximation might yield a local optimum, but given that the optimization is based on a nonlinear, non-convex function for the yaw moment it can't be guaranteed that this point is also globally optimal.

4-3-3 Choice of norms n_R and n_Q

The most commonly selected norm for optimization problems is the ℓ_2-norm. This norm produces a convex optimization problem when its argument is linear or affine, and due to its quadratic nature produces smooth outputs. The ℓ_2-norm is therefore selected as the norm that penalizes the errors in virtual control effort. However, to penalize the actuator usage, using the ℓ_2-norm might not be the most suitable when looking at the design goals. As was noted in the introduction, the brakes tend to be an intrusive and energy-inefficient way of intervening and it is therefore preferred to delay the usage of brakes as much as possible and favor usage of other actuators, such as the front and rear electric motors for the vehicle under consideration. However, if a point is reached where the brakes need to be used to meet the virtual control requests, they might as well be used with relatively little penalty.

The modified ℓ_2-norm

Converting these desires into a cost function that reflects this would mean that:

1. at low actuator efforts, the cost of actuation, q_Q, needs to be relatively high for the brakes, or any undesirable actuator for that matter.
2. whenever the tracking error of the virtual control command becomes relatively large the undesired actuators need to engage and their usage should not be heavily penalized.

For 1 an ℓ_2-norm with a high cost at low actuator efforts for the undesirable actuators can be used. This however directly contradicts 2 and therefore a modification to the ℓ_2-norm is proposed. Above a certain actuation effort $u_{Q,lim}$ the cost function for the actuators is changed from growing quadratically to growing linearly:

$$q_Q = \begin{cases} \mathbf{u}^T \mathbf{Q}^{\ell_2} \mathbf{u} & \text{if } \mathbf{u} \leq u_{Q,lim} \\ 2u_{Q,lim}^T \mathbf{Q}^{\ell_2} |\mathbf{u}| - u_{Q,lim}^T \mathbf{Q}^{\ell_2} u_{Q,lim} & \text{if } \mathbf{u} > u_{Q,lim} \end{cases} \tag{4-21}$$

where, $u_{Q,lim} \in \mathbb{R}^n$ a row vector containing the individual actuation efforts where it switches and $\mathbf{Q}^{\ell_2} \in \mathbb{R}^{n \times n}$ the traditional, diagonal cost matrix with weights:

$$\mathbf{Q}^{\ell_2} = \begin{bmatrix} Q_1^{\ell_2} & 0 & \dots & 0 \\ 0 & Q_2^{\ell_2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & Q_n^{\ell_2} \end{bmatrix}$$

This modified ℓ_2-norm for q_Q is still convex and continuous which means it is suitable for use within the hybrid steepest descent framework.

Illustrative example of modified ℓ_2-norm

To show how this modified norm is different from a 'regular' ℓ_2-norm a simple scalar example will be worked out. The control allocation is heavily simplified, with a single actuator u and virtual control input v:

$$q = \begin{cases} (u-v)^2 + Q \cdot u^2 & \text{if } u \leq u_{Q,lim} \\ (u-v)^2 + 2u_{Q,lim} Q|u| - u_{Q,lim}^2 Q & \text{if } u > u_{Q,lim} \end{cases}$$

Calculating the derivative for q and setting it to 0 allows us to find the optimal value u as a function of v and $u_{Q,lim}$:

$$\frac{dq}{du} = \begin{cases} (Q+1)u - v & \text{if } u \leq u_{Q,lim} \\ u - v + u_{Q,lim} Q & \text{if } u > u_{Q,lim} \end{cases}$$

when restricting the analysis to $u \geq 0$.

Setting these equations to 0 and solving for u allows us to find the optimal value of u. However, the interest is not in the value of u but in the relative error between input u and output v, which will be defined as $e = \frac{u}{v}$:

$$e = \begin{cases} \dfrac{1}{1+Q} & \text{if } u \leq u_{Q,lim} \\ 1 - \dfrac{u_{Q,lim} Q}{v} & \text{if } u > u_{Q,lim} \end{cases}$$

These equations show that using only the ℓ_2-norm is inflexible, given that it introduces a fixed error $\dfrac{1}{1+Q}$ regardless of the magnitude of v. When using the modified norm, above a

certain $u = u_{Q,lim}$ the relative tracking error will start decreasing in an inversely proportional way to v: $1 - \frac{u_{Q,lim}Q}{v}$, meaning that at low virtual control inputs the focus can be more on minimizing the actuator effort (if desired) whereas if the virtual control input grows relatively large, the focus of the optimization shifts to just tracking the virtual control input with a low relative error. This is precisely what this modified norm intended to achieve.

4-3-4 Total cost function

With the norms selected and definitions for the virtual control inputs as a function of $\mathbf{u}$, a convex formulation is obtained for the cost function q with tracking error cost q_R and actuation cost q_Q:

$$q = q_R^{M_z} + q_R^{F_x} + q_Q \text{ where } \begin{cases} q_R^{M_z} & = R_{1,1}\Big(\bar{E}^T \cdot (\mathbf{u} - \mathbf{u}_0) + M_z(\mathbf{x}_0, \mathbf{u}_0) - M_{z,ref}\Big)^2 \\ q_R^{F_x} & = R_{2,2}\Big(1/r_w \sum_{i=1}^{n} u_i - F_{x,ref}\Big)^2 \\ q_Q & = \begin{cases} (\mathbf{u} - \mathbf{u}_{desired})\mathbf{Q^{\ell_2}}(\mathbf{u} - \mathbf{u}_{desired}) & \text{if } q_Q \leq q_{Q,lim} \\ 2\sqrt{\mathbf{Q^{\ell_2}} q_{Q,lim}}|\mathbf{u} - \mathbf{u}_{desired}| - q_{Q,lim} & \text{if } q_Q > q_{Q,lim} \end{cases} \end{cases} \tag{4-22}$$

The next step is to compute the gradients ∇q and ∇g as a function of $\mathbf{u}$, again separating tracking error and actuation:

$$\nabla q = \nabla q_R^{M_z} + \nabla q_R^{F_x} + \nabla q_Q \text{ where } \begin{cases} \nabla q_R^{M_z} & = 2R_{1,1}\bar{E}\Big(\bar{E}^T \cdot (\mathbf{u} - \mathbf{u}_0) + M_z(\mathbf{x}_0, \mathbf{u}_0) - M_{z,ref}\Big) \\ \nabla q_R^{F_x} & = 2R_{2,2}\Big(1/r_w \sum_{i=1}^{n} u_i - F_{x,ref}\Big) \\ \nabla q_Q & = \begin{cases} 2\mathbf{Q^{\ell_2}}(\mathbf{u} - \mathbf{u}_{desired}) & \text{if } q_Q \leq q_{Q,lim} \\ 2\sqrt{\mathbf{Q^{\ell_2}} q_{Q,lim}}\text{sgn}(\mathbf{u} - \mathbf{u}_{desired}) & \text{if } q_Q > q_{Q,lim} \end{cases} \end{cases} \tag{4-23}$$

and the constraints:

$$\nabla g = \begin{cases} \nabla g_i = \beta_i \mathbf{e}_i & \text{for } \mathbf{u} - \mathbf{u}_{lim}^{max} > 0 \\ \nabla g_i = -\beta_i \mathbf{e}_i & \text{for } -\mathbf{u} - \mathbf{u}_{lim}^{min} > 0 \end{cases} \tag{4-24}$$

where $i = \{1, 2, \ldots, n\}$, $\mathbf{e}_i$ the unit vector in the i-th direction and β_i the scaling factor to determine how fast $\mathbf{u}$ is forced back towards the feasible region. The dynamic scaling method as explained in Section 4-1-3 will be used to determine the values of β_i.

If the linearization is performed at every time step n the formulation for $\nabla q_R^{M_z}$ can be further simplified because in that case $\mathbf{u_0} = \mathbf{u}$. This results in:

$$\nabla q_R^{M_z}|^{\mathbf{u_0}=\mathbf{u}} = 2R_{1,1}\bar{E}\Big(M_z(\mathbf{x}_0, \mathbf{u}_0) - M_{z,ref}\Big) \tag{4-25}$$

Given the gradients, Eq. (4-3) is then used to provide the rate of change of the actuator efforts $\dot{\mathbf{u}}$:

$$\dot{\mathbf{u}} = -\gamma_0 \nabla q - \sum_{i=1}^{m} \gamma_i(u) \nabla g_i(u), \quad i = \{1, 2, \ldots, 12\} \tag{4-26}$$

Integrating this provides the outputs of the control allocator: the actuator targets for the low level controllers $\mathbf{u}$. Because the implementation is done in discrete time, the integration is performed using the the forward-Euler approach:

$$\mathbf{u}(n) = \dot{\mathbf{u}} \cdot t_s + \mathbf{u}(n-1) \tag{4-27}$$

Note that the actuator rate limits have not yet been applied. This may be done by saturating Eq. (4-26) between upper and lower rate limits:

$$\dot{\mathbf{u}} = \text{sat}(\dot{\mathbf{u}}, \mathbf{u}_{lim}^{min}, \mathbf{u}_{lim}^{max}) \tag{4-28}$$

This method is also proposed in [3]. The modification means the trajectory is no longer following the steepest descent direction, but it will still be descending towards the optimal point q^*.

4-4 Alternate perspective on hybrid steepest descent optimization method

The presented optimization problem in the previous section, although relatively simple for an optimization problem, might still be difficult to understand as to how it works. This section will provide an alternative perspective on the problem that is more intuitive. By using the transfer function formalism, the optimization problem is translated from a dynamic system where the dynamics of this system are determined by ∇q or ∇g, to a closed loop controller with gain scheduling. The transfer function approach, despite requiring some simplifications, provides good insight into how the optimization works and how the optimization parameters, such as the optimization weight, influence allocation performance. The following section will cover the simplifications and steps that were taken to convert the system to a transfer function form, focusing on the dynamics of the yaw moment request $M_{z,ref}$ for this particular example. However, the analysis presented below can be extended to the other optimization objectives as well.

4-4-1 Problem formulation

The goal of this analysis is to set the problem up to obtain a function of the form that will allow investigation of input-output behavior of the optimization algorithm, where the input is $M_{z,ref}$ and the output M_z:

$$M_z = f(M_{z,ref}) \rightarrow \mathcal{L}\{f(M_{z,ref})\} = \frac{M_z(s)}{M_{z,ref}(s)}$$

When none of the constraints in Eq. (4-26) are active the actuator dynamics are described by:

$$\dot{\mathbf{u}} = -\nabla q_R^{M_z} - \nabla q_R^{F_x} - \nabla q_Q$$

Lumping parts of $\nabla q_R^{M_z}$ together and renaming them:

$$\nabla q_R^{M_z} = \overbrace{2R_{1,1}\bar{E}}^{K}\left(\bar{E}^T \cdot (\overbrace{\mathbf{u}-\mathbf{u}_0}^{\tilde{u}}) + \overbrace{M_z(\mathbf{x}_0,\mathbf{u}_0) - M_{z,ref}}^{-\tilde{v}}\right)$$

The next step requires a large simplification that assumes $\tilde{u}$ to be a linear function of $\nabla q_R^{F_x} + \nabla q_Q$. $\nabla q_R^{F_x} + \nabla q_Q$ in the suggested formulation is in part a linear function of $\tilde{u}$. However, it also contains inter-dependencies between $\tilde{u}$ as well as constant terms and terms that depend on other virtual inputs. This makes translation to transfer function form more complicated if not impossible, depending on the exact problem formulation. These non-linear parts will therefore be left out, reducing the accuracy of this approach. However, the results will still prove insightful despite this step. The approximation:

$$\nabla q_R^{F_x} + \nabla q_Q \approx G\tilde{u}$$

where $G \in \mathbb{R}^{n\times n}$ a diagonal matrix where its entries on the diagonal approximate the effects that other virtual control inputs and actuator penalties have on $\dot{\mathbf{u}}$

4-4-2 Converting optimization problem to transfer functions

With all components of the equations now explicit, linear functions of $\tilde{u}$ the system may be transformed to a transfer function. $\mathbf{u}_0$ is treated as a linearization constant, hence its derivative is zero:

$$\begin{aligned}\dot{\tilde{u}} &= \dot{\mathbf{u}} - \dot{\mathbf{u}}_0, \quad \text{where } -\dot{\mathbf{u}}_0 = 0\\ \dot{\tilde{u}} &= -K(\bar{E}^T\tilde{u} - \tilde{v}) - G\tilde{u}\end{aligned}$$

Rearranging:

$$K\tilde{v} = \dot{\tilde{u}} + (K\bar{E}^T + G)\tilde{u}$$

Applying the $\mathcal{L}$-transform yields a vector $A(s)$ with transfer functions:

$$\begin{aligned}a_i(s) &= \frac{\tilde{u}_i(s)}{\tilde{v}(s)} = \frac{K_i}{K_i\bar{E}_i + G_{i,i} + s} = \frac{K_i \cdot (K_i\bar{E}_i + G_{i,i})^{-1}}{(K_i\bar{E}_i + G_{i,i})^{-1}s + 1}\\ A(s) &= \begin{bmatrix}a_1(s) & a_2(s) & \dots & a_n(s)\end{bmatrix}^T\\ i &= \{1, 2, \dots, n\}\end{aligned}$$

Realizing that the change of yaw moment caused by the actuators around $\mathbf{x}_0, \mathbf{u}_0$ was approximated as $\tilde{M}_z = \bar{E}^T \cdot \tilde{u}$, the sum of the vector of transfer functions may be used to approximate the transfer function from $\tilde{v}$ to M_z, denoted by $\tilde{M}_z(s)$:

$$\begin{aligned}\tilde{M}_z(s) &= \bar{E}^T\frac{\tilde{u}(s)}{\tilde{v}(s)} = \bar{E}^T A(s)\\ \tilde{M}_z(s) &= \bar{E}_1\tilde{a}_1(s)\tilde{v}(s) + \bar{E}_2\tilde{a}_2(s)\tilde{v}(s) + \dots + \bar{E}_n\tilde{a}_n(s)\tilde{v}(s)\end{aligned}$$

The resultant transfer function of $\tilde{M}_z(s)$ is precisely what this analysis set out to define.

Several useful things can be learned from this transfer function $\tilde{M}_z(s)$:

1. $\tilde{M}_z(s)$ consists of the sum of first-order transfer functions, where the bandwidth of each transfer function is determined by $(K_i\bar{E}_i + G_{i,i})^{-1}$
2. The DC gains of each transfer function are determined by $\bar{E}_i K_i \cdot (K_i\bar{E}_i + G_{i,i})^{-1}$

Although these equation cannot be used directly for control analysis due to the large simplifications, they can be used to provide an indication of how large the optimization weights need to be chosen to achieve sufficient bandwidth of the control allocator. K and G both contain optimization weights so the magnitude of the weights will directly impact bandwidth. Obviously the bandwidth should be chosen as high as possible, however this should not go at the cost of too much chatter around constraints or the optimum. The equation for $\tilde{M}_z(s)$ can be used to make an informed trade-off between bandwidth and chatter.

Furthermore, the effect of the actuator effort weights, which in this approach are contained in G can be seen to result in a non-unity DC gain of the system if $G \neq 0$. If a large G is chosen it is clear that the DC gains will tend towards 0 and steady state tracking performance of the the yaw moment request will be relatively bad. Conversely if a small G is chosen it will make the DC gain tend towards 1, resulting in good steady state tracking of the reference. Therefore the DC gain can be used to get an understanding of how the trade-off between the gains in G and K should be made to get the desired behavior of the control allocator.

4-5 Benchmark control allocator

The benchmark control allocator serves the purpose of providing lower performance bounds against which the performance of the HSDO allocation algorithm can be measured. A very simple allocator will be designed and implemented to provide this. The goal is that the HSDO allocator should at the very least perform better than the simple control allocator, and preferably demonstrate superior performance.

4-5-1 Design of the simple control allocator

The simple control allocator will be made to resemble the performance of current industry standard VSC, which only makes use of the brakes. The focus of this simple control allocator will be on producing the yaw moment $M_{z,ref}$ requested by the high level controller. This decision was made to keep this control allocator simple and targeted. As soon as other high level objectives are added (minimizing deviations from $F_{x,ref}$ or minimizing actuator effort) some heuristics or control dead zones would need to be introduced, which can quickly lead to increasing levels of complexity and unpredictable behavior that is very sensitive to tuning. This is not desired for this benchmark controller, hence it is set to track only the most important objective.

Given that there are four brakes on the car, two on either side, a way needs to be devised to split the yaw moment request between front and rear axle and left and right side. For the front-to-rear split, the decision was made to use actuator ganging, which in its simplest form

is a fixed split value σ_{split} that determines the amount of yaw moment produced by the front and rear brakes:

$$\begin{aligned} M_{z,f} &= \sigma_{split} M_z \\ M_{z,r} &= (1 - \sigma_{split}) M_z \end{aligned} \tag{4-29}$$

When investigating Eq. (2-13) it can be seen that the yaw moment on the vehicle is mainly determined by the differential of longitudinal forces across an axle, the direct yaw moment. Given that brakes can only provide negative forces (whilst traveling in forward direction) finding the amount of brake force required to meet the per axle yaw moment is approximated as:

$$\begin{aligned} T_{br,fl} &= r_w d/2 |M_{z,f}| && \text{for } M_{z,f} > 0, F_{br,fl} = 0 \text{ otherwise} \\ T_{br,fr} &= r_w d/2 |M_{z,f}| && \text{for } M_{z,f} < 0, F_{br,fr} = 0 \text{ otherwise} \\ T_{br,rl} &= r_w d/2 |M_{z,r}| && \text{for } M_{z,r} > 0, F_{br,rl} = 0 \text{ otherwise} \\ T_{br,rr} &= r_w d/2 |M_{z,r}| && \text{for } M_{z,r} < 0, F_{br,rr} = 0 \text{ otherwise} \end{aligned}$$

The value of σ_{split} was chosen to be 0.65 after some tuning and was found to produce good performance across the various test cases.

The motor torques are left unmodified. The original brake based stability control designed by Bosch [7] is able to modify engine torque, but it was only allowed to modify in the direction of 0 torque and is focused on combustion engine powertrains. In order to avoid further complexity, it is therefore decided to leave motor torques as they are.

4-6 Conclusion

The formulation presented in this chapter for the control allocation problem, made suitable for hybrid steepest descent optimization, seems to fulfill the requirement of 'ease of implementation'. The modifications proposed to the hybrid steepest descent optimization method, dynamically scaling the constraint gradients and iterating to ensure feasibility provide improvements in terms of convergence speed of the algorithm and smoothness of the outputs. Despite the modifications the method is still computationally simple making it transparent and its behavior easy to understand. The linearization approach to estimate the yaw moment control effectiveness per actuator elegantly removes a lot of the difficult nonlinearities inherent to vehicle models and is an important step in simplifying the formulation. Furthermore, the number of tuning variables of this formulation is limited to the optimization weights which should make the tuning straight-forward and intuitive. The modified ℓ_2-norm proposed in this chapter allows the optimization algorithm to naturally shift focus from one optimization objective to the other. This lets the control designer make trade-off's between prioritizing actuator effort reduction and virtual control input tracking performance as a function of the virtual control inputs without the need of additional heuristics or fuzzy logic.

Chapter 5

Control Allocator Evaluation: Simulation

To evaluate the performance of the proposed algorithm, it was implemented in the simulation model and compared to the benchmark control allocator for a number of different test cases. First of all, to assess if the advanced allocator improves performance over the simple control allocator, the advanced control allocator is allowed to use only the brakes for control. Secondly, the advanced control allocator is also allowed to modify motor torques and performance is again assessed against the advanced control allocator using only the brakes to see how the control allocator utilizes the full control authority of the entire system. Finally, to highlight some other features, two special cases are shown of the advanced control allocator. The first case is where the brake system is faulted and the control allocator has to stabilize the vehicle with only motor torques. This was done to show-case the ability of the advanced control allocator to deal with actuator failures. The second case was performed to assess the robustness of the advanced control allocator to parameter mismatch. One of the most likely sources of parameter mismatch, an incorrect estimate of surface friction coefficient, was applied to the advanced control allocator and performance is once again assessed.

A note about the presented data in this chapter; all results shown have been scaled with base values. This was done to protect Tesla's intellectual property. The base values scale torques, forces and accelerations by a fixed amount to bring them into approximately the same order of magnitude. The base values were kept constant across all tests. To give the reader a frame of reference for the orders of magnitude, all torques are scaled by approximately 5000 [Nm], accelerations by approximately 10 [m/s^2] and angular values by approximately 10 [deg]

5-1 Maneuver and performance metric definition

The proposed hybrid steepest descent optimization control allocator will be compared to the simple control allocator on a set of maneuvers. The rationale behind the maneuver selection

will be presented in this section, as well as the performance metrics that were chosen to compare the performance between the various control allocators.

5-1-1 Maneuvers

The choice was made to perform maneuvers with open loop driver inputs where the open loop inputs are the steering wheel angle δ_f and the desired acceleration. This is subsequently translated to motor torque $T_{m,f}^{desired}$ and $T_{m,r}^{desired}$.

Alternatively, closed-loop driver inputs could have been selected where the driver is modeled as a closed loop controller that attempts to track a trajectory and desired speed. However, this would introduce extra variability in the simulation results which makes direct comparisons more difficult, hence this method was not used.

Two different types of maneuvers were selected, each performed on two different surfaces. The selected maneuvers are:

1. Single lane change. This is a standard vehicle dynamics test performed at constant speed with aggressive steering inputs, resulting in fast transients that will usually destabilize an uncontrolled vehicle. A number of standards exist for this maneuver, e.g. [41]. The inputs were chosen similar to these standards, but adapted to make sure it would lead to sufficiently unstable behavior in the speed range of interest. This particular test is used mostly to highlight the ability of the control allocator to produce a large amount of yaw moment very quickly whilst dealing with rate limits and tire traction limits.
2. The highway on-ramp. The maneuver starts by bringing the vehicle into a steady state turn at approximately 70% of its tire limits. After reaching steady state, the driver aggressively commands longitudinal acceleration whilst attempting to maintain the same level of lateral acceleration. All the motor torque is initially applied on the rear axle which would lead to aggressive sideslip growth in the uncontrolled state. This maneuver tests the capabilities of the control allocation in combined limit handling, facing potentially conflicting objectives as there is a request for aggressive longitudinal acceleration as well as a need for corrective yaw moment to reduce the on-power oversteer induced by the rear motor torque.

Both maneuvers are performed on surfaces resembling dry tarmac with a μ_s of around 1 and on snow with a μ_s of around 0.4. The uncontrolled maneuvers with the corresponding driver inputs for high μ_s are shown in Figure 5-1 and Figure 5-5.

5-1-2 Performance metrics

To get an indication of the performance of the various control allocators in the maneuvers defined above, a number of performance metrics are defined, some of which will provide an indication of how the performance of the control allocator impacts system level performance, and some that indicate how well the allocator itself performs. It will also provide an indication of how the trade off is made between the three main optimization objectives, these being yaw

moment tracking, longitudinal force tracking and actuator effort reduction. The selected system level metric is:

1. RMS yaw rate error ψ_{error}

and the selected control allocator performance metrics are:

1. RMS tracking error between produced yaw moment and requested yaw moment $M_{z,error}$
2. RMS longitudinal force error $F_{x,error}$
3. Energy dissipated by the brakes over the maneuver as a measure of actuation effort. Energy consumption is not calculated using wheel speeds, but approximated by the vehicle speed to account for locked wheels: $E_{brake,ij} = \int_0^t U_x/r_w T_{br,i,} dt$. During extreme maneuvers the brakes are capable of driving the wheels into deep slip, which results in a relatively low rotation velocity of the wheel and hence low energy consumption by the brakes. However, considerably more energy is lost during these events in the form of tire slip. To account for this the vehicle speed is used to better approximate the combined energy loss through tire slip and heat dissipation of the brakes.

Motor energy consumption is not reported, given that electric motors are actuators with high energy conversion efficiency. Furthermore, for this particular vehicle the slip losses due to the motor torques don't need to be taken into account as the slips are kept relatively small due to the way the traction control was tuned.

5-1-3 Control allocator settings

For the simulation, the tuning parameters and saturation limits need to be set for the advanced control allocator. This section will cover all the tuning parameters that were used for simulation.

Saturation limits

The settings for absolute saturation limits $\mathbf{u}_{lim}$ and rate limits $\dot{\mathbf{u}}_{lim}$ for the control allocator were kept constant across all the test cases. The limits approximate the actual capabilities of the actuators used in the simulation model. It must be noted that the brakes do not have a lower saturation limit on the simulated vehicle, hence the control allocators saturation limit was set as to not engage during normal operation. Dynamic saturation limits can be also imposed on actuators. These dynamic limits $T_i^{max,dyn}$ and $T_i^{min,dyn}$ can represent whatever limit is desired. For these simulations the only dynamic limits imposed are the limits produced by the low level controllers, denoted by $T_{m,i}^{TC}$ for motor traction control and $T_{br,ij}^{BTC}$ for brake traction control. All the limits shown are scaled by the same base values as the rest of the presented data.

$$\dot{\mathbf{u}}_{lim} = \begin{bmatrix} \dot{T}_{m,f}^{min/max} \\ \dot{T}_{m,r}^{min/max} \\ \dot{T}_{br,fl}^{min/max} \\ \dot{T}_{br,fr}^{min/max} \\ \dot{T}_{br,rl}^{min/max} \\ \dot{T}_{br,rr}^{min/max} \end{bmatrix} = \begin{bmatrix} \pm 3.4 \\ \pm 3.4 \\ \pm 1.4 \\ \pm 1.4 \\ \pm 1.4 \\ \pm 1.4 \end{bmatrix}, \mathbf{u}_{lim}^{max} = \begin{bmatrix} T_{m,f}^{max} \\ T_{m,r}^{max} \\ T_{br,fl}^{max} \\ T_{br,fr}^{max} \\ T_{br,rl}^{max} \\ T_{br,rr}^{max} \end{bmatrix} = \begin{bmatrix} \min(0.5, T_{m,f}^{max,dyn}) \\ \min(1, T_{m,r}^{max,dyn}) \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \mathbf{u}_{lim}^{min} = \begin{bmatrix} T_{m,f}^{min} \\ T_{m,r}^{min} \\ T_{br,fl}^{min} \\ T_{br,fr}^{min} \\ T_{br,rl}^{min} \\ T_{br,rr}^{min} \end{bmatrix} = \begin{bmatrix} \max(-0.25, T_{m,f}^{max,dyn}) \\ \max(-0.5, T_{m,r}^{max,dyn}) \\ \max(-2, T_{br,fl}^{max,dyn}) \\ \max(-2, T_{br,fr}^{max,dyn}) \\ \max(-2, T_{br,rl}^{max,dyn}) \\ \max(-2, T_{br,rr}^{max,dyn}) \end{bmatrix} \tag{5-1}$$

Optimization weights

One set of optimization weights will be used throughout the analysis. To keep it comparable to the simple control allocator, it was designed to focus on yaw moment tracking, with lower priority for longitudinal acceleration tracking and reduced actuator usage. Given that using the front and rear electric motors for achieving the virtual control inputs is preferred over using the brakes, a lower weight is put on motor usage to favor using these actuators.

After some tuning the following set of parameters was found to provide good performance across the whole range of tests:

$$\mathbf{R} = \begin{bmatrix} 0.2 & 0 \\ 0 & 0.02 \end{bmatrix}, \quad \mathbf{Q}^{\boldsymbol{\ell_2}} = \begin{bmatrix} 0.3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.3 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad u_{Q,lim} = \begin{bmatrix} 0.08 \\ 0.08 \\ 0.04 \\ 0.04 \\ 0.04 \\ 0.04 \end{bmatrix} \tag{5-2}$$

Yaw moment estimation model and sample time

No further modifications were made to the advanced control allocator; the implementation was done as described in Chapter 4. The two-track model using sigmoid tires, as shown in appendix A was used to estimate the exerted yaw moment M_z as well as to calculate the yaw moment effectiveness.

Furthermore, the sample time for both allocators was set to 10 [ms]. This provided good performance in simulation and the computational load associated with running the advanced control allocator on an actual ECU was found to be manageable, as will be shown in Chapter 6.

5-1-4 Low level controllers and high level controller

Both the low level controllers and high level controller were experimentally tuned to achieve good and robust performance across the various test cases. The tuning for both was kept the same across all the test cases for both the simple and advanced control allocator. Furthermore, the low level controllers and high level controller run at the same sample rate as the control allocators, which is 10 [ms]. Given the frequency content of the chassis dynamics of the typical car, the high level controller could have been designed to run at a lower sample rate to reduce computational load, but it was kept the same as the other modules for simplicity.

5-2 Performance evaluation in simulation: lane change maneuver

In this section the performance of the various control allocators will be assessed and compared with each other for the single lane change. First, the simple control allocator is compared against the advanced control allocator, where the advanced control allocator is allowed to only modify brake torques. This allows a direct assessment of how much the optimization improves the allocation performance over the simple brake control allocator. Secondly, the advanced control allocator will also be given access to motor torques. Performance is compared to the version using only the brakes to see how the optimization handles the addition of extra actuators. The following pages show the time domain plots of the uncontrolled case (Figure 5-1), the simple control allocator (SCA, Figure 5-2), the brake only HSDO control allocator (B.O. HSDO, Figure 5-3) and the full HSDO control allocator with motor torques (Figure 5-4). For brevity only time domain plots are shown for the tarmac case. The performance metrics will be covered for the single lane change on snow, given that in terms of trends, behavior was found to be similar to the tarmac maneuvers.

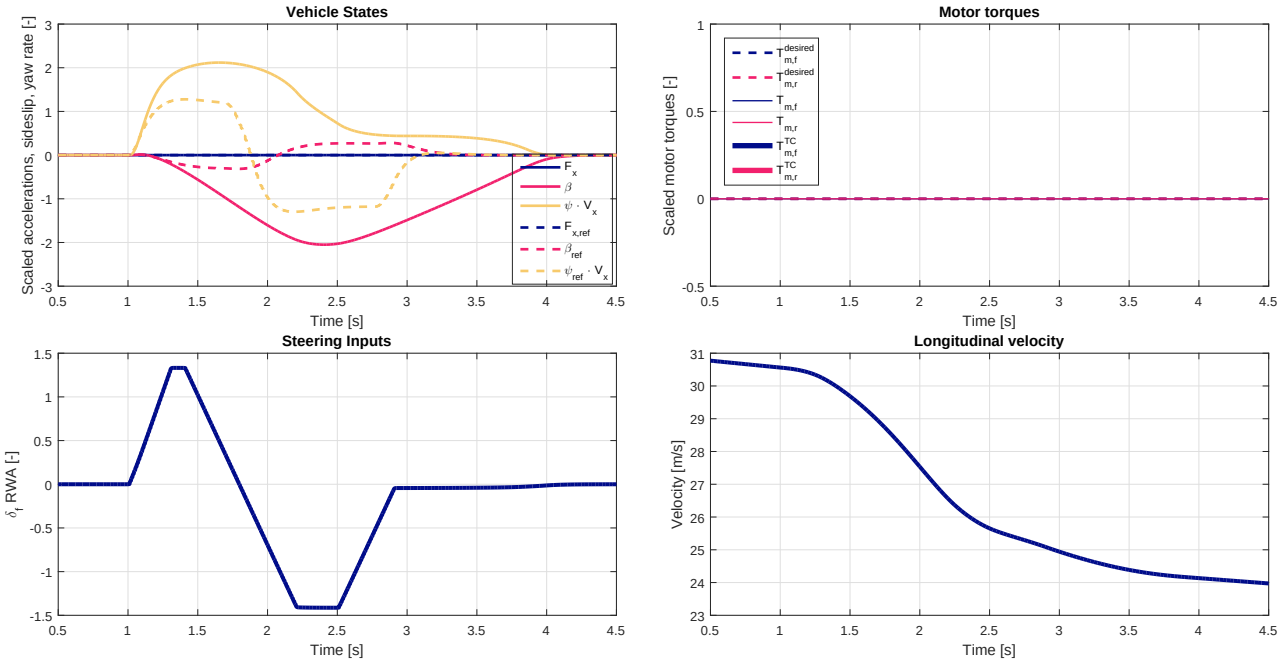

Figure 5-1: Uncontrolled single lane change on tarmac

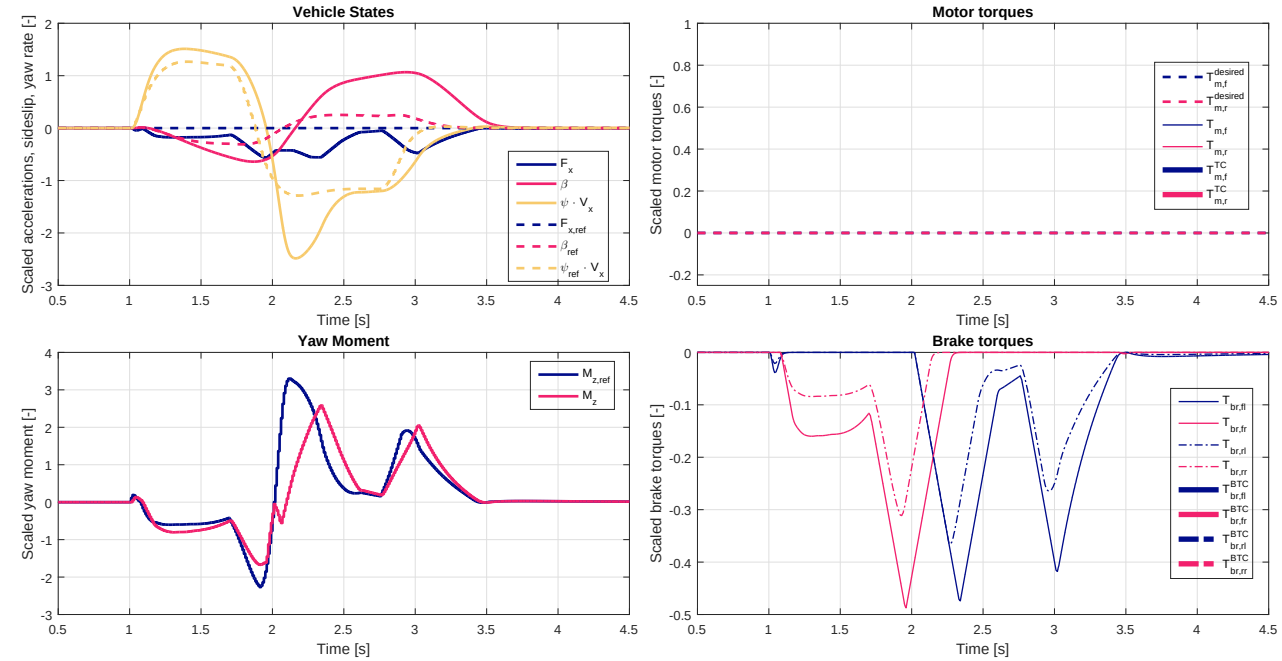

Figure 5-2: Simple control allocator, single lane change on tarmac

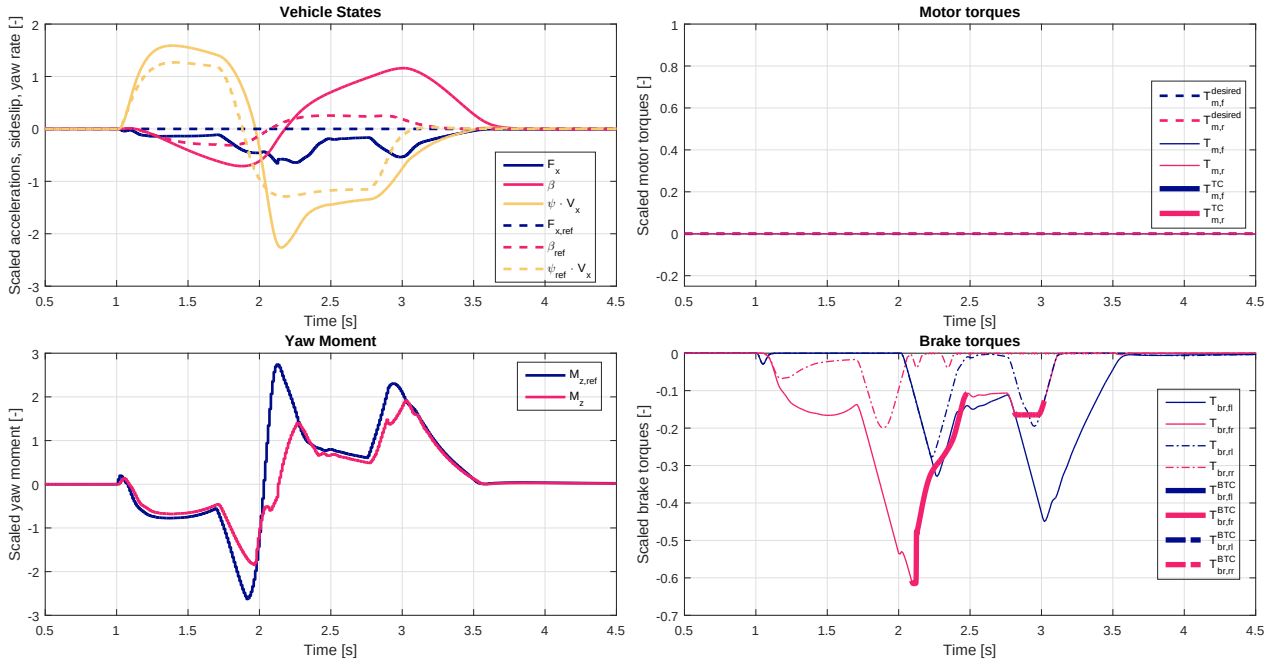

Figure 5-3: HSDO control allocator using only brakes, single lane change on tarmac

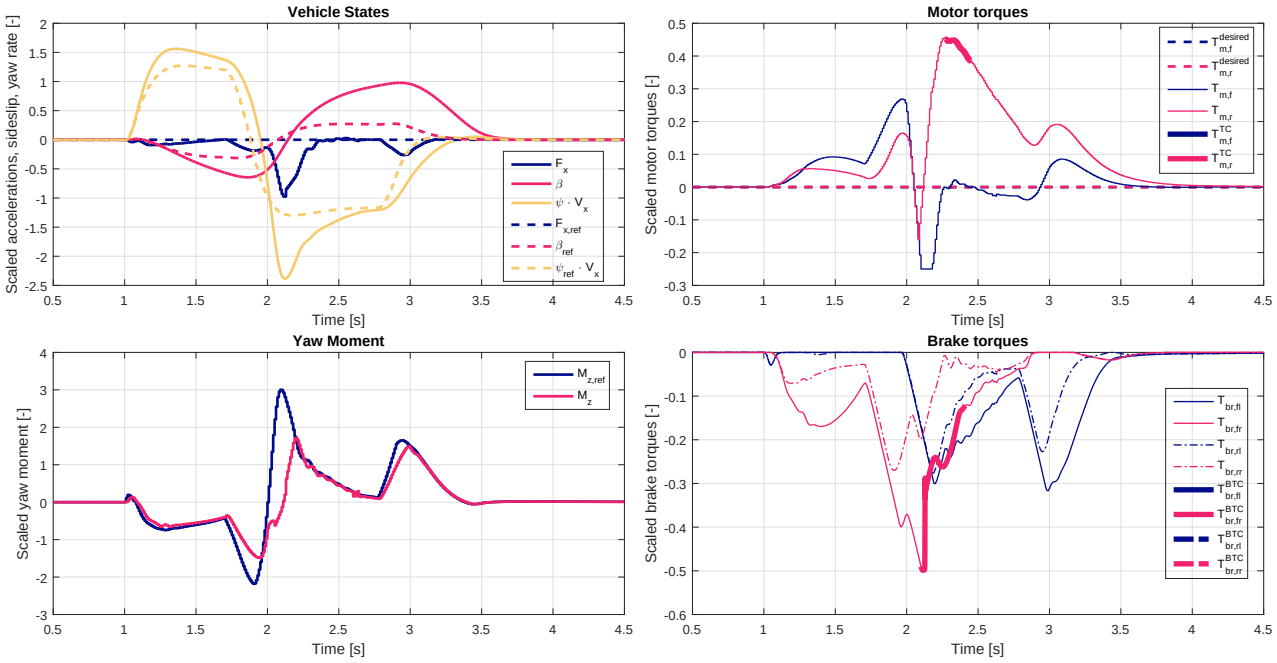

Figure 5-4: Full advanced control allocator, single lane change on tarmac

5-2-1 Discussion of results: single lane change

A progression of improvements is visible across the various tests. It can be seen that the uncontrolled single lane-change is dominated by unstable behavior, indicated by the large sideslip growth (β) and the vehicle completely failing to follow the driver's requested yaw rate $\dot{\psi}_{ref}$. All the subsequent controlled cases show much improved performance over the uncontrolled case. This indicates that the VSC control architecture that was designed is capable of stabilizing the vehicle, regardless of the type of control allocator that is used. From a control perspective, the maneuver is completely dominated by three large spikes in stabilizing yaw moment $M_{z,ref}$. The ability for the control allocators to track this large and rapidly changing request is what differentiates the performance of the control allocators.

5-2-2 Comparing SCA and B.O. HSDO control allocator

When comparing Figure 5-2 and Figure 5-3 it can be seen that the HSDO control allocator seems to better track the yaw moment, resulting in slightly better yaw rate control. This observation is echoed by the performance metrics:

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
HSDO	30.98	18.87	0.2523	0.1827	-66.19	0.3837
SCA	30.98	19.22	0.2591	0.1722	-62.85	0.4505
% Improved	-	-	2.62	-6.1	-5.31	14.8

Table 5-1: Performance comparison between HSDO control allocator using only brakes and SCA, single lane change on tarmac

When assessing the results in more detail it can be seen why the advanced control allocator is better at tracking the yaw moment. As mentioned before, this lane-change maneuver is dominated by three large spikes of stabilizing yaw moment request from the high level controller. It can be seen that during every fast yaw moment transient, the advanced allocator uses both the front and rear brake on one side of the vehicle to bring the yaw moment up as quickly as possible, maximally exploiting the total available actuation bandwidth. As soon as the error is reduced the optimization switches to using the most efficient actuator as much as possible, which are the front brakes for a stabilizing yaw moment. Despite the fact that the optimization attempts to use the most efficient actuator as much as possible, it can still be seen that the energy consumption is slightly higher. That is mostly due to the large yaw moment spike at the 2.3 second mark, where the yaw moment error is so large that actuator usage is hardly penalized, which is by design (Section 4-3-3). During this spike the advanced control allocator makes a non-intuitive decision that upon further reflection makes sense. At the two second mark the yaw moment request changes from a large negative to a large positive value. This would intuitively result in the inside brakes reducing torque as quickly as possible, and the outside brakes increasing as fast as possible. This is also seen to happen with the simple control allocator. The advanced control allocator does indeed increase both outside wheels at a maximum rate, however, the inside front wheel (front right) is not decreased as expected, but actually *increased* to drive that wheel into deep longitudinal slip to prevent it from producing

any lateral force. This should help produce a stabilizing yaw moment. After the wheel is driven into deep slip it can be seen that the brake traction control starts reducing torque to prevent the wheel from fully locking up. It does something similar at the three second mark, deciding to drive the inside front wheel into deep slip to produce extra stabilizing yaw moment.

5-2-3 Comparing B.O. and full HSDO control allocator

When comparing Figure 5-3 and Figure 5-4 it must be noted that the yaw moment tracking performance seems very similar, but the tracking of the longitudinal acceleration is considerably improved:

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
Full HSDO	30.98	22.41	0.2233	0.1274	-64.04	0.3819
B.O. HSDO	30.98	18.87	0.2523	0.1827	-66.19	0.3837
% Improved	-	-	11.5	30.3	3.25	0.469

Table 5-2: Performance comparison between full HSDO control allocator and the brake-only variant, single lane change on tarmac

Given the limited control authority of motor torques on tarmac, it is not surprising that the improvement in yaw moment tracking is insignificant. The control allocator can be seen to be using the motor torques mainly to reduce the deviations in longitudinal acceleration F_x, which is where a significant improvement is noted. It is only during the largest spike in yaw moment where the yaw moment tracking error is large, around the 2.3 second mark, that the control allocator starts using the front motor fully to aid the brakes with stabilizing the car. It does this by driving the front motor with maximum negative torque in an attempt to provide an indirect, stabilizing yaw moment. As soon as the brakes are capable of handling the yaw moment request, the motor torques return to being used for longitudinal acceleration control. This is as expected, given that the brakes are far more efficient at producing the yaw moments on tarmac compared to the front and rear motors.

5-2-4 Performance metrics of single lane change on snow

Both test cases were also performed on snow, but will not be discussed as elaborately. The lane change on snow shows a more significant improvement compared to the one on tarmac, in this case mostly because the HSDO control allocator makes efficient use of both the indirect and direct way of generating yaw moments. The simple control allocator, by design, only relies on the direct approach and has no notion of the indirect yaw moment it is generating. The control allocator performance metrics all show significant improvements. Interesting to note is that the yaw rate tracking performance has decreased somewhat, despite the better yaw moment tracking. This happens because the simple control allocator does not take into account the additional indirect yaw moment it is generating. Therefore, it actually overproduces yaw moment at times which leads to a higher effective gain for the high level controller which, in turn, leads to tighter yaw rate control.

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
HSDO	16.73	12.99	0.05062	0.03742	-9.089	0.1062
SCA	16.73	12.78	0.04691	0.04854	-10.86	0.1687
% Improved	-	-	-7.91	22.9	16.3	37

Table 5-3: Performance comparison between HSDO control allocator using only brakes and SCA, single lane change on snow

When the brakes are added to the system further improvements are noted. Given that the motors have relatively more control authority on snow than they do on tarmac, it seems consistent that the addition of the motor torques leads to overall improved yaw moment tracking. Once again, considerably improved longitudinal force tracking is noted, where the full HSDO control allocator uses the motor torques in a similar way to the lane change on tarmac, shown in Figure 5-4.

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
Full HSDO	16.73	13.7	0.04848	0.02157	-10.58	0.07996
B.O. HSDO	16.73	12.99	0.05062	0.03742	-9.089	0.1062
% Improved	-	-	4.23	42.4	-16.4	24.7

Table 5-4: Performance comparison between full HSDO control allocator and brake-only variant, single lane change on snow

5-3 Performance evaluation in simulation: highway on-ramp maneuver

Similar to the previous section, the performance of the various control allocators will be assessed and compared to each other, starting again by comparing the simple control allocator against the advanced control allocator, where the advanced control allocator is only allowed to modify brake torques. Secondly, the advanced control allocator will be given access to motor torques as well and performance is compared to using only the brakes to see how the optimization handles the addition of extra actuators. The following pages show the time domain plots of the uncontrolled case (Figure 5-5), the simple control allocator (SCA, Figure 5-6), the brake only HSDO control allocator (B.O. HSDO, Figure 5-7) and the full HSDO control allocator with motor torques (Figure 5-8). Only the performance metrics will be covered for the on-ramp maneuver on snow, given that in terms of trends, behavior was found to be similar to the tarmac maneuvers.

5-3-1 Discussion of results: highway on-ramp maneuver

Despite the completely different nature of this maneuver, a progression of improvements is again seen across the various tests. Similar to the lane-change, it can be seen that the uncontrolled single lane-change is dominated by unstable behavior, this time induced by the application of a large amount of rear motor torque. The rear motor torque request $T_{m,r}^{desired}$ is bounded significantly by traction control, denoted by $T_{m,r}^{TC}$. Despite the traction control intervention the torque application on the rear axle still causes the vehicle to oversteer, indicated by a large persistent error in yaw rate $\dot{\psi}$ and resultant sideslip growth. All the subsequent controlled cases show considerably improved performance over the uncontrolled case, especially for the full HSDO control allocator. From a control perspective, unlike the lane-change, this maneuver is not clearly dominated by one objective. The yaw moment request $M_{z,ref}$ is relatively low and the longitudinal acceleration request $F_{x,ref}$ relatively high, meaning the control allocator will need to find a good balance between the two.

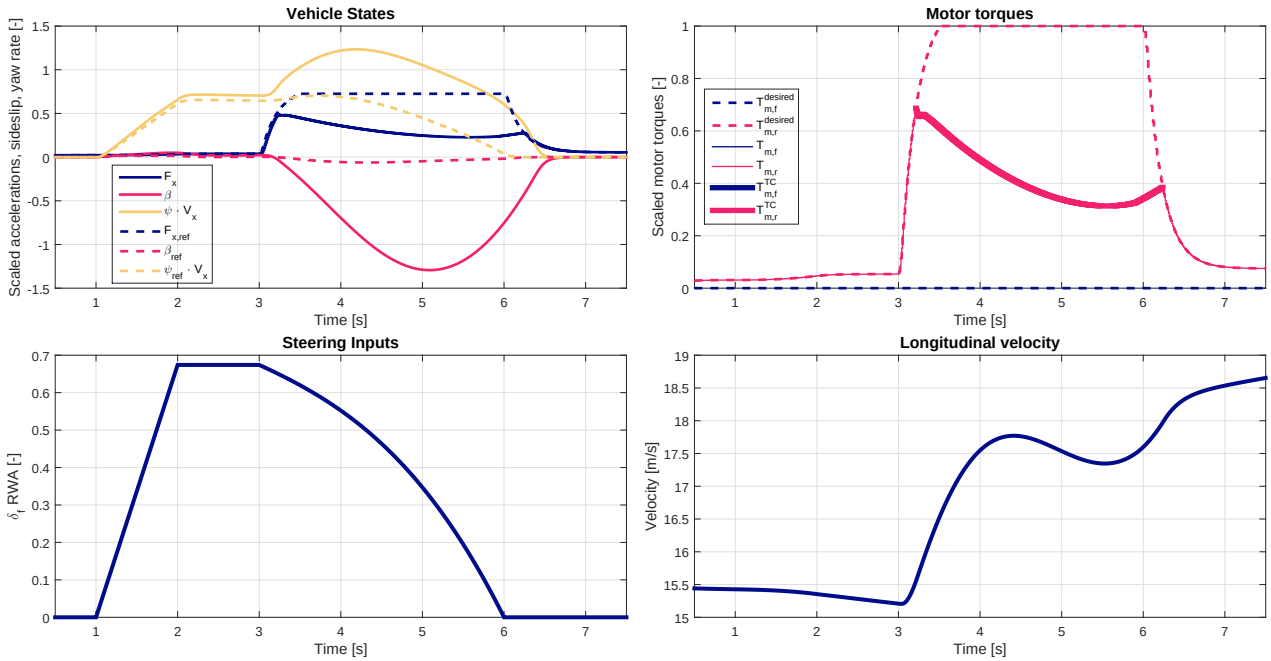

Figure 5-5: Uncontrolled highway on-ramp maneuver on tarmac

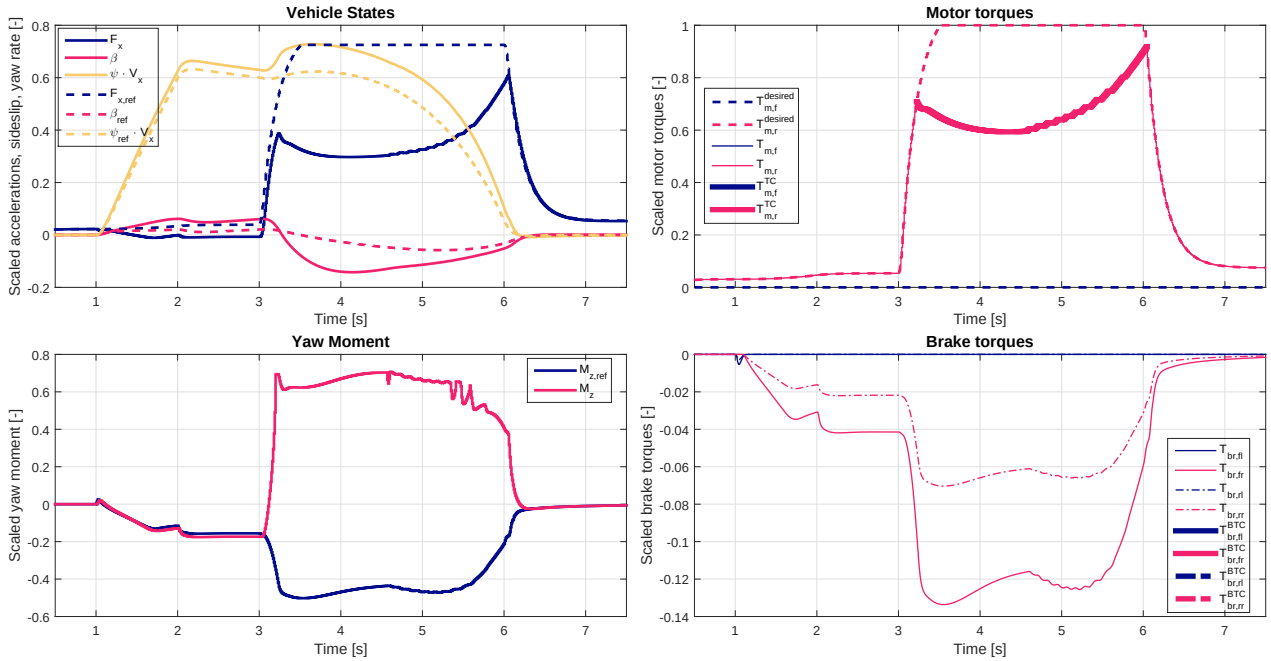

Figure 5-6: Simple control allocator, on ramp maneuver on tarmac

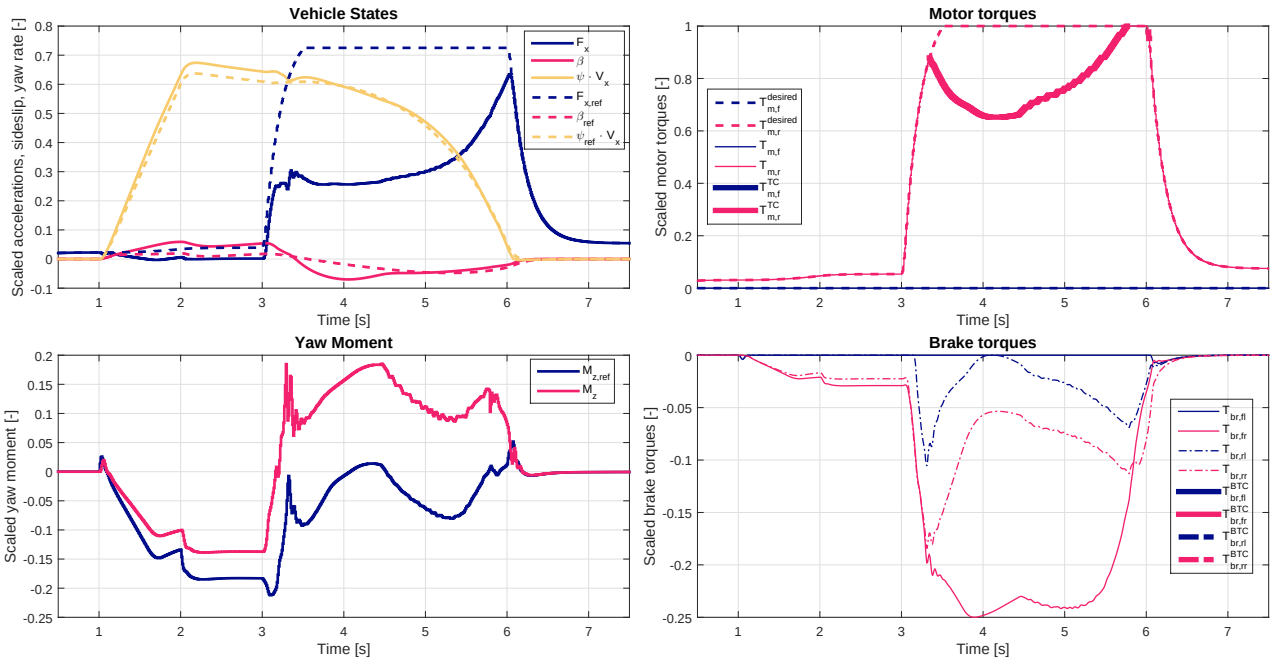

Figure 5-7: HSDO allocator using only brakes, on ramp maneuver on tarmac

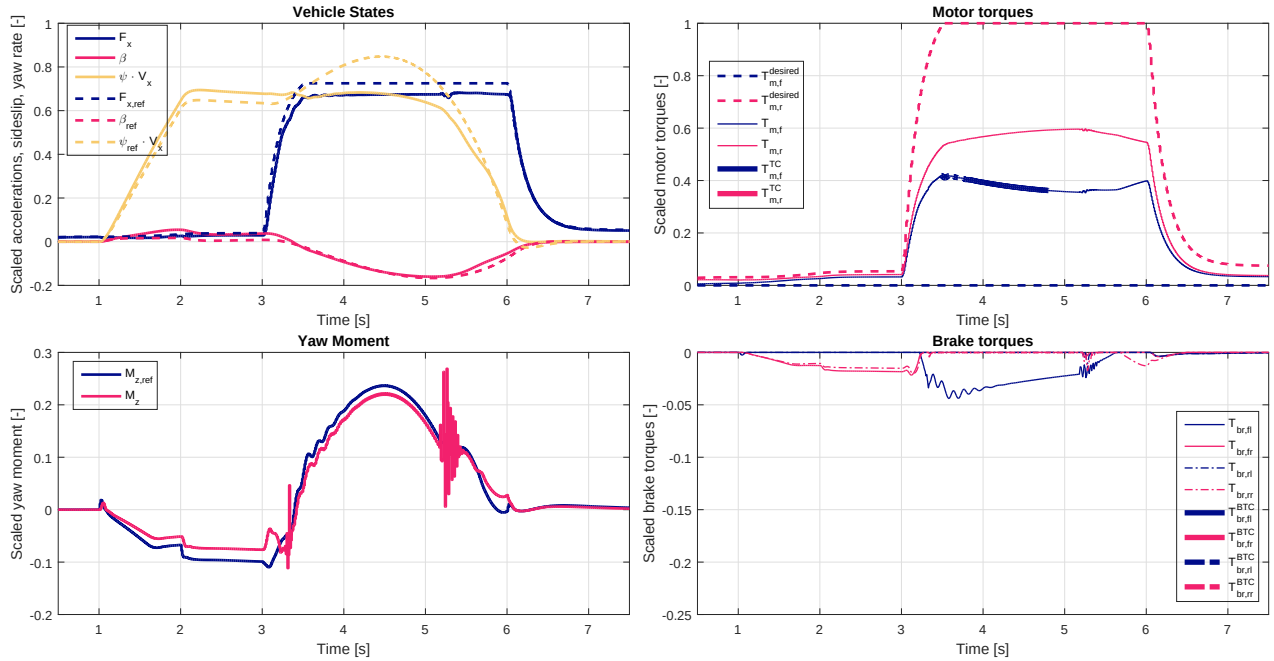

Figure 5-8: Full HSDO control allocator, on ramp maneuver on tarmac

5-3-2 Comparing SCA and B.O. HSDO control allocator

When comparing Figure 5-6 and Figure 5-7 it can be seen that the HSDO control allocator seems better in tracking the yaw moment, resulting in slightly better yaw rate control. However, it does so at the cost of brake usage. This observation is echoed by the performance metrics:

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
HSDO	15.49	22.34	0.01824	0.2494	-51.52	0.1062
SCA	15.49	23.02	0.06249	0.2303	-31.33	0.6673
% Improved	-	-	70.8	-8.29	-64.4	84.1

Table 5-5: Performance comparison between HSDO control allocator using only brakes and SCA, highway on-ramp on tarmac

In the initial phase of the maneuver, the yaw moments are relatively small and it can be observed that the advanced allocator does not track the yaw moment request very well (Figure 5-7). It does this because the yaw moment is small and it focuses instead on reducing brake usage which leads to a fixed offset in the yaw moment (Section 4-3-3). It is during the phase where rear motor torque is applied that the biggest advantage is gained by the HSDO allocator. The HSDO allocator tracks the yaw moment request considerably better because it has an understanding of the destabilizing yaw moment being generated by the rear motor torque due to its internal yaw moment estimate. This means it can compensate for the destabilizing rear motor torque by applying extra stabilizing brake torques, whereas the simple allocator does not take that information into account and tries to directly apply the commanded yaw moment. However, even for the HSDO control allocator the error is still relatively large visually, which is caused by the fact that in, an absolute sense, the yaw moment error is not very large compared to the longitudinal acceleration error. This means the optimization focuses actuator efforts between these two objectives and tries to find a good compromise determined by the relative difference between the optimization weights.

The way in which the HSDO allocator achieves this compromise is interesting. Firstly, it understands that overall force needs to be reduced on the rear axle to stabilize the car. It does this by applying both inside *and* outside rear brakes. Furthermore, it brakes the outside wheel more than the inside wheel to provide an additional stabilizing yaw moment, turning the rear axle into a very inefficient torque vectoring differential. This is also why the energy consumption is relatively high. In reality the motor torque would probably be reduced as well to prevent this inefficient mode of operation, but, in this contrived scenario, this seems the best solution to achieve both good tracking of the yaw moment and the longitudinal force targets. Furthermore, the optimization uses the outside front brake most heavily to meet the stabilizing yaw moment request, because the outside front brake is the most efficient actuator for producing stabilizing yaw moments in situations like this.

5-3-3 Comparing B.O. and full HSDO control allocator

When comparing Figure 5-7 and Figure 5-8 a considerable improvement can be seen, both in yaw moment as well as longitudinal acceleration tracking which is reflected by the performance metrics:

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
Full HSDO	15.49	31.09	0.05318	0.03675	-10.23	0.03121
B.O. HSDO	15.49	22.34	0.01824	0.2494	-51.52	0.1062
% Improved	-	-	-192	85.3	80.1	70.6

Table 5-6: Performance comparison between full HSDO control allocator and the one using only brakes, highway on-ramp on tarmac

Intuitively it is clear that the optimal torque distribution out of a sharp corner should be close to a 50/50 torque split between front and rear motors, not 100 % rear. During the initial phase of the maneuver, the motor torques are still low and the allocator is forced to use some light braking to track the yaw moment request, but as soon as the acceleration request goes up, so do the torques and it finds a way to balance the torques front to rear to track both the longitudinal acceleration request as well as the yaw moment request and does so using very little braking.

Some interesting things should be noted about these results:

1. At the 3.5 second mark a small damped oscillation can be seen both in the yaw moment request and front left brake torque. After further investigation this oscillation was found to be mostly caused by an interaction between the high level controller and the limited bandwidth of the brakes which was set at around 5 [Hz]. The low bandwidth leads to reduced system phase margin which can lead to oscillatory behavior. In situations where limited actuator bandwidth leads to deteriorated performance, it is anticipated that adding prediction to the problem formulation will help by allowing the allocation algorithm to account for actuator dynamics.

2. Some chatter can be observed at the 5.2 second mark in the yaw moment request and brake torques. This is because the optimization is chattering around the optimal point. The magnitude of the optimization weights was chosen to maximize the bandwidth of the control allocator without causing chatter, but it was tuned to be right on the edge and in this case, it briefly goes over the edge. It does not however impact performance negatively, but if smooth outputs are desired, it is advisable to lower the overall magnitude of the optimization weights. If performance bandwidth is found to be unacceptable then it will mean sample time needs to be decreased.

3. Yaw rate tracking did not improve despite better yaw moment tracking. This is simply a shortcoming of the simple design of the high level controller which relies mostly on proportional action. This can easily be improved by adding more aggressive integral action or more advanced high level control strategies, if so desired.

5-3-4 Performance metrics of highway on-ramp on snow

Trends for both comparisons on snow are very similar, the brake-only HSDO control allocator showing much improved yaw moment tracking at the cost of brake usage:

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
HSDO	8.367	10.83	0.006091	0.1813	-12.89	0.08425
SCA	8.367	10.79	0.01294	0.1859	-6.255	0.1588
% Improved	-	-	52.9	2.47	-106	46.9

Table 5-7: Performance comparison between HSDO control allocator using only brakes and SCA, highway on-ramp on snow

Furthermore, the same big performance gain is seen when the HSDO control allocator is given access to motor torques. The control allocation manages to improve tracking of yaw moment and longitudinal acceleration whilst hardly using the brakes:

	V_{init}	V_{end}	RMS $\dot{\psi}_{error}$	RMS $F_{x,error}$	Brake Energy	RMS $M_{z,error}$
Full HSDO	8.367	15.34	0.003422	0.07723	-1.819	0.01919
B.O. HSDO	8.367	10.83	0.006091	0.1813	-12.89	0.08425
% Improved	-	-	43.8	57.4	85.9	77.2

Table 5-8: Performance comparison between full HSDO control allocator and the one using only brakes, highway on-ramp on snow

5-4 Fault tolerance and robustness

The performance improvements using the advanced control allocator are clear for standard operation. This section will cover some non-standard operation to investigate fault-tolerance and robustness to model parameter mismatch of the advanced control allocator. The investigated fault tolerance case is a full brake failure, giving the allocator authority over only the motor torques to stabilize the vehicle. The robustness case focuses on the allocator using a bad surface friction coefficient in its yaw moment effectiveness calculations.

5-4-1 Fault tolerance: brake system failure

In this example case the entire brake system is faulted, meaning that the control allocator cannot command any brake torques. The control allocator is not informed of the brake system failure so the brakes are still part of the optimization. As the online linearization and yaw moment estimation use actual brake torque measurements it does at least estimate the exerted yaw moment and actuator effectiveness accurately. The fact that it uses measured motor and brake torques for these calculations is very important. If it had used its internal actuator commands for this, it would not be aware at all of this brake failure and performance would most likely suffer.

The resultant performance of the advanced control allocator is shown in Figure 5-9. Even without the brakes the system manages to stabilize the vehicle, although the performance is obviously not as good as that of the one that also has brakes available. This is shown as comparison in Figure 5-10. The HSDO control allocator makes heavy use of the front motor torque, pushing the front motor to maximum negative torque every time a large stabilizing yaw moment is required and at the same time keeping rear torque relatively low to give the rear axle maximum lateral grip. During the brief moments the vehicle understeers the control allocator increases the rear motor torque to reduce rear lateral grip and reduces front motor torque towards 0 to increase front axle lateral grip. This behavior is all consistent with the findings in Section 4-3-2.

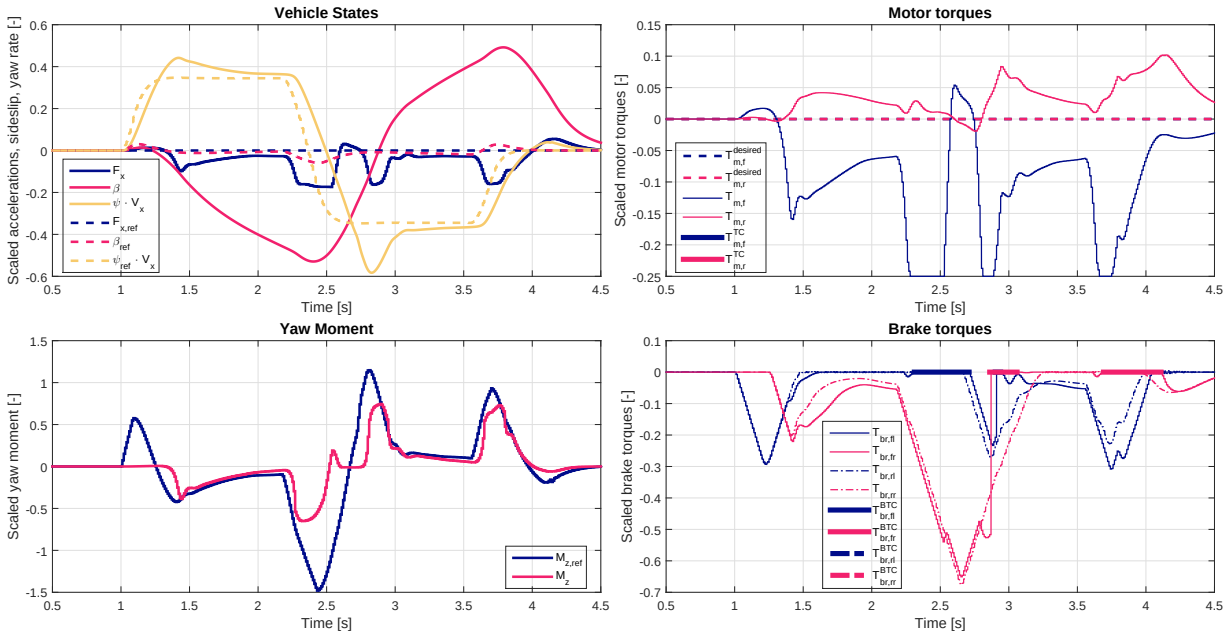

Figure 5-9: Full advanced control allocator, lane-change on snow with faulted brake system. Commanded brake torques are shown, actual brake torques are at 0

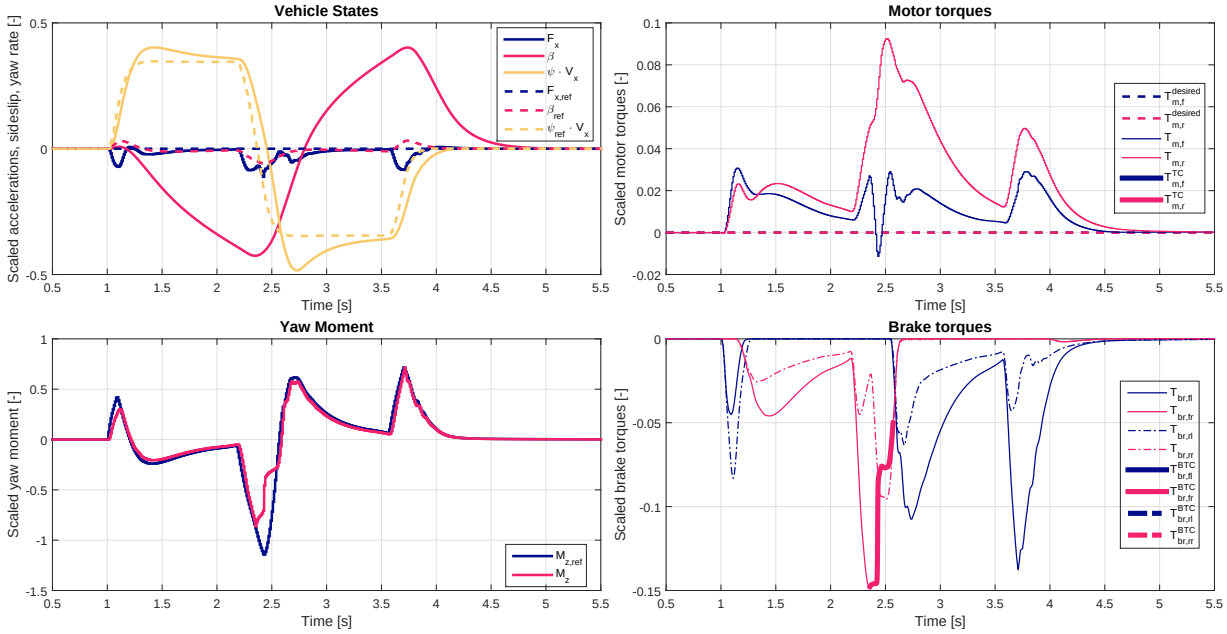

Figure 5-10: Full advanced control allocator, lane-change on snow

5-4-2 Robustness: μ estimation error

Because the surface friction coefficient μ_s can't be measured directly, vehicles need to be equipped with some form of surface friction estimation. This estimation can be challenging and vehicle dynamics control systems should be robust to errors in this estimate. Therefore, to investigate robustness of the advanced control allocator to such estimation errors, it was subjected to an extreme test where it performs the lane-change on snow, whilst assuming it is still on tarmac. This is an over-estimation of the friction coefficient of a factor of 2.5. The bad estimate will mean it estimates the yaw moment it is producing incorrectly as well as getting the yaw moment effectiveness estimates incorrect. The performance is slightly reduced compared to the nominal system (as shown in Figure 5-10), but given the extremity of the estimation error the loss of performance seems wholly acceptable.

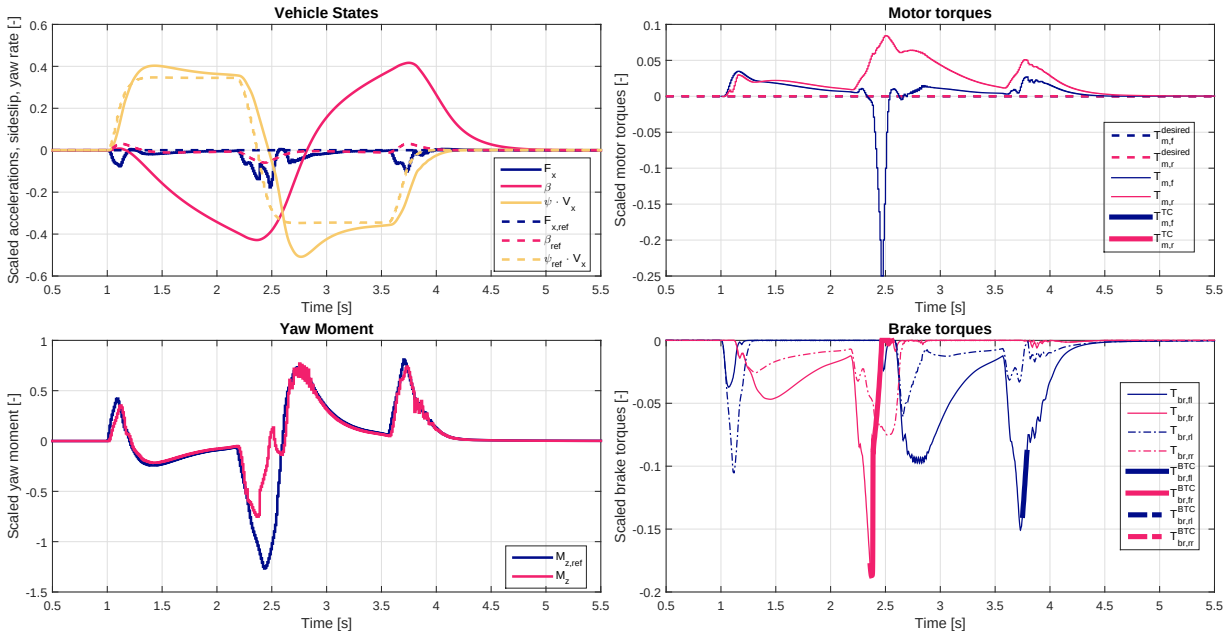

Figure 5-11: Full advanced control allocator, lane-change on snow with incorrect μ_s estimate

This performance cannot however be fully attributed to the control allocator, but also for a large part to the high level controller being robust to this. Therefore, for parameter mismatch scenario's, both the high level controls as well as the control allocator need to be investigated as one system.

5-5 Conclusion

The overall results show significantly improved performance when using the HSDO control allocator, especially if it is given access to all the available actuators. Even in its most basic form, when compared to the simple control allocator, the HSDO control allocator showed improved allocation performance, mostly due to its ability to handle actuator constraints efficiently as well as having the internal yaw moment estimate which allows it to compensate

for disturbance yaw moments such as those introduced by motor torques. The addition of the motors to the allocation problem proved straightforward and, without any retuning, provided a significant performance improvement. This indicates that the problem formulation is chosen well, allowing for easy addition or removal of actuators from the system. Remembering that the design goals are to design an algorithm capable of providing high performance whilst being easy to use and intuitive, the simulation results suggest that this has been achieved with the proposed HSDO control allocator. Finally, two examples were shown to demonstrate robustness to estimation errors and actuator failures. In both cases performance was found to be good despite the handicaps that were imposed.

Chapter 6

Experimental Validation

In the previous chapters the hybrid steepest descent optimization control allocator was worked out in detail and shown, using simulations, to be a promising method of performing dynamic control allocation on an over-actuated vehicle. To strengthen these conclusions, experimental validation of the model with the VSC system developed in this thesis was performed using a Tesla Model S. The performance of the motor-only hybrid steepest descent control allocator to the performance of the Bosch VSC that is currently used on the test vehicle was also compared. This chapter will cover the experimental setup, the maneuvers performed and a discussion of the results.

6-1 Experimental setup

As mentioned before, the test vehicle used for the experiments is a Tesla Model S P90DL. The vehicle is used in an unmodified state, except for the special code which was running on one of the on board ECUs. Unfortunately, the available vehicle did not allow brake torques to be commanded, meaning that solely motor torques were available for actuation. This meant that only a stripped down version of the hybrid steepest descent optimization allocator could be tested. However, as the results of Section 5-4-1 showed, using only motor torques can be sufficient to stabilize the vehicle and it will be interesting to see if these conclusions are echoed by the experimental results.

6-1-1 Test equipment and measurements

The vehicle is equipped with a VSC system from Bosch which provides a number of measurements of vehicle states. Furthermore, the ECU on which the code is executed includes some additional vehicle state estimates. The only missing vehicle state is lateral velocity U_y. This requires external equipment to be fitted to the vehicle to be measured. The equipment was not available at the time of testing. However, the available measurements are more than sufficient to provide an indication of the predictive capabilities of the model. The lack of lateral

velocity measurement did require some modifications to the yaw moment and yaw moment effectiveness calculations, but the modifications do not cause a significant loss of accuracy. The modifications are shown in appendix A along with the original calculation method that does utilize lateral velocity.

An overview of the measurements that were taken:

1. yaw rate $\dot{\psi}$, measured by the VSC module
2. lateral acceleration A_y, measured by the VSC module
3. longitudinal acceleration A_x, measured by the VSC module
4. longitudinal velocity U_x, internal estimate
5. front and rear motor torques $T_{m,f}, T_{m,r}$, internal estimate
6. steering wheel angle δ_f, internal measurement

No post-processing of the data was required as the data proved sufficiently smooth and drift-free.

6-1-2 Test site and maneuvers

Testing was performed at Valkenburg Naval Airport in the Netherlands. The site has a reasonably sized tarmac VDA, but it was deemed insufficiently large to perform the on-ramp maneuvers safely. Therefore, tests were restricted to the single lane-change maneuver. The maneuver was performed manually during testing, which introduced small variances in initial velocity and steering inputs from test-to-test. Therefore, every test was conducted at least twice to ensure some level of repeatability. The speed at which the maneuver was performed was reduced from 31 [m/s] to 23 [m/s] for safety reasons, but this was still sufficiently fast to make the vehicle spin out approximately 180^o after exiting the lane-change with all the controls deactivated. A data trace of the uncontrolled lane change is shown in Figure 6-4.

6-1-3 Code implementation on the Tesla ECU

No external prototyping ECU was required for these tests. The high level controller and hybrid steepest descent control allocation algorithm were converted, by hand, to C-code with only minor modifications and integrated within the torque management algorithms aboard one of the Tesla ECU's. The conversion process to C-code proved straightforward given the relative simplicity of the HSDO control allocator.

Both high level controller and dynamic control allocator are executed inside a 10 [ms] loop, just as was used for the simulations. To prevent any unintended interactions with the Bosch VSC system or Tesla's traction control, these systems were completely disabled to give the high level controller and HSDO control allocator full authority.

6-1-4 Modifications to simulation environment

To compare the measurements with the simulated results, some modifications were made to the simulation environment. The C-code implementation of the HDSO control allocator running on the Tesla ECU is slightly different to the implementation in the simulation environment used to generate the results show in Chapter 5. These were due to safety modifications that were made to safe-guard against inadvertent torque modifications from the control allocator during normal driving. To eliminate any errors arising from different algorithmic implementation, the C-code running on the Tesla ECU was adapted to be executed inside a software-in-the-loop (SIL) environment. This was connected to the vehicle model used in this thesis to allow it to perform closed loop control model validation.

To get an idea of how the system was going to perform, this SIL environment was used before testing to perform some re-tuning. Given that the system only has the front and rear motors for yaw moment regulation which are expected to have limited control authority on tarmac, more emphasis was put on tracking $M_{z,ref}$ rather than $F_{x,ref}$ by reducing the weight $R_{2,2}$.

6-2 Test results

The measurements collected at Valkenburg were processed and compared with the outputs of the SIL simulation environment to assess the fidelity of the modeling environment. The inputs to the simulation used for validation are the measured steering wheel angle, initial vehicle states and the driver requested motor torques. Time domain results for one of the data sets is shown in Figure 6-1, showing both the measured data and the simulation outputs. Furthermore, the data collected from the lane change with the VSC system designed in this thesis was compared with that of the performance of the Bosch VSC which is currently used on the vehicle. The Bosch system is a brake-based stability control system which is expected to be well tuned and should therefore provide a good indication of the current industry standard.

6-2-1 Model validation and discussion of results

The single lane change test case is dominated, just as it was in the simulation results of Chapter 5, by a number of large spikes in stabilizing yaw moment. Recalling the results from Figure 4-2 and Figure 5-9 the system is expected to rely heavily on negative front motor torques to stabilize the vehicle. Given the limited yaw moment authority of the motor torques on tarmac, the control allocator is expected to struggle in meeting the yaw moment request from the high level controller.

When investigating the time domain plots in Figure 6-1 it is clear that the system responds as expected. The maneuver is entered with the driver requesting practically 0 torque, but as soon as the high level controller requests stabilizing yaw moments, the front motor torque is aggressively reduced to its maximum negative output by the control allocator. This is precisely as expected. As soon as the yaw moment requests reduce again the motor torque is rapidly brought back to the driver requested torque to minimize the deviation in acceleration

request. Furthermore it can be observed that the yaw moment request is not tracked particularly well, as was expected, because the yaw moment authority of motor torques is relatively low on tarmac.

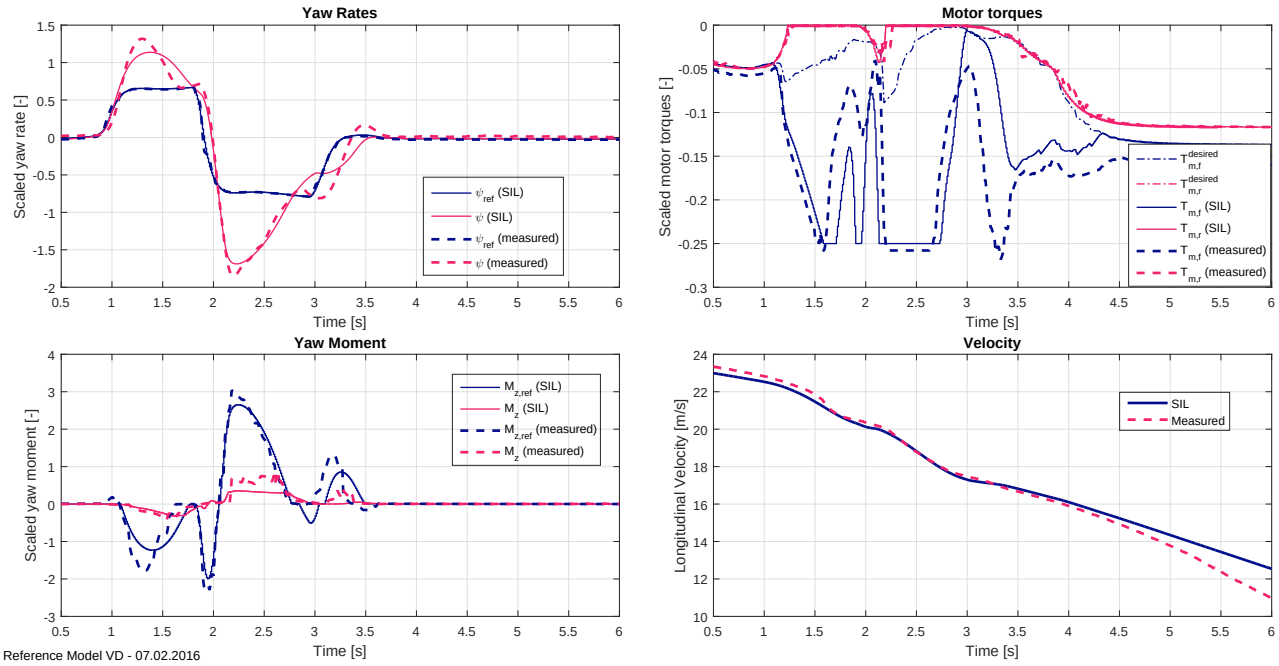

Figure 6-1: Comparison between measured and simulation data. Data is shown for the HSDO control allocator stabilizing the Model S during an aggressive lane change with only the motor torques (same as Figure 6-2). Simulation results of the SIL simulation environment of the same maneuver with the same inputs are also shown

It can be observed that the SIL simulation environment shows good agreement with the measured data, indicating that the model that was developed has good fidelity even when performing closed loop control during limit handling. Some deviations in state trajectories and control inputs can be seen during the maneuver which was to be expected, given the extremity of the maneuver. The tires are both laterally and longitudinally heavily saturated during various stages in the maneuver and operate deeply within their nonlinear region, an area where tire models typically tend be less accurate.

6-2-2 Comparison to Bosch VSC system

To get an impression how good the performance of the system developed in this thesis is, it is compared to the Bosch VSC system. The comparison is shown below in Figure 6-2:

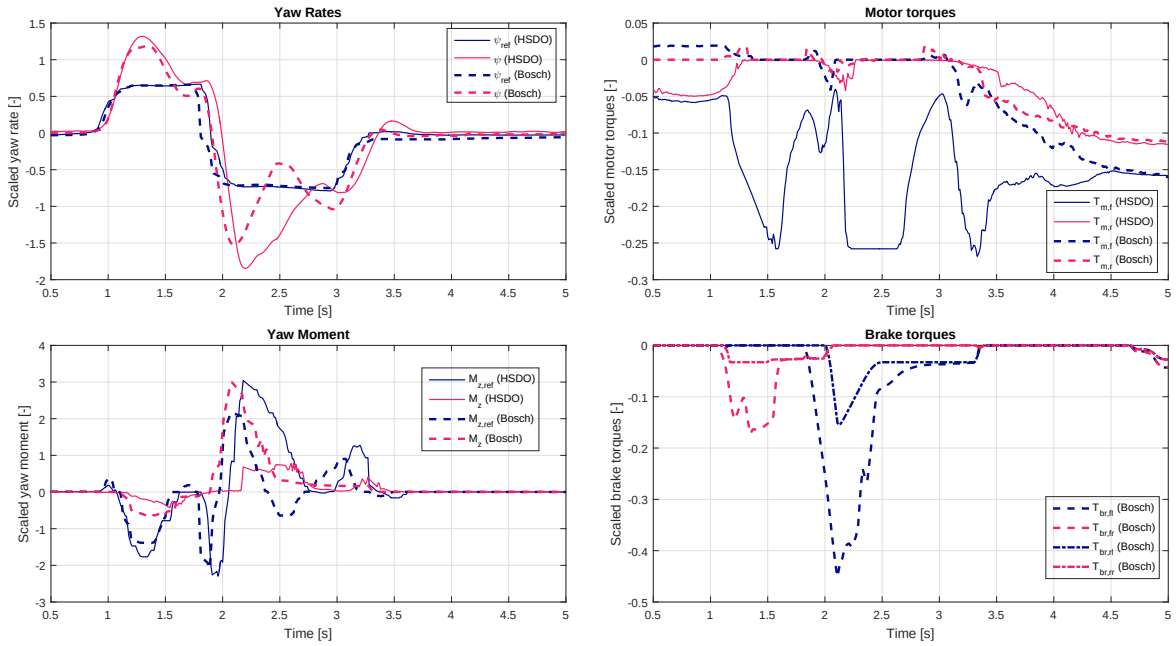

Figure 6-2: Comparison of Bosch VSC and the VSC system developed in this thesis with HSDO motor-only control allocator stabilizing the Model S during an aggressive lane change. The only data traces that were actually collected from the Bosch VSC system are the yaw rate, brake and motor torques. All the other values such as reference yaw rate $\dot{\psi}_{ref}$ and yaw moments M_z were inferred by running the VSC system developed in this thesis in the background with its outputs disabled, making the assumption that the Bosch system operates in a similar way to this system

And the corresponding steering inputs and velocities during both maneuvers in Figure 6-3:

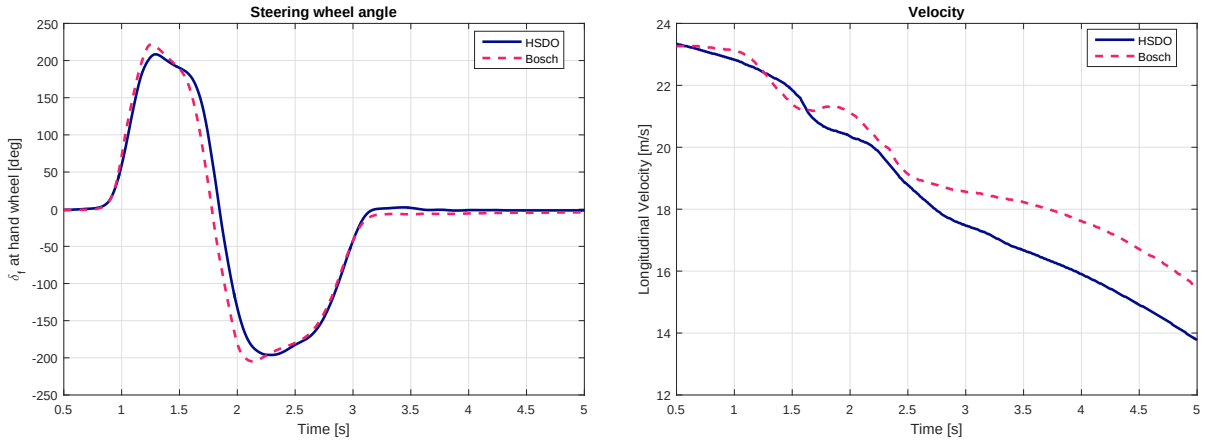

Figure 6-3: Inputs for Bosch VSC and the VSC system developed in this thesis with HSDO motor-only control allocator stabilizing the Model S during an aggressive lane change

Furthermore, to provide a frame of reference for what happens when all the controls are turned off, the uncontrolled lane change is shown in Figure 6-4:

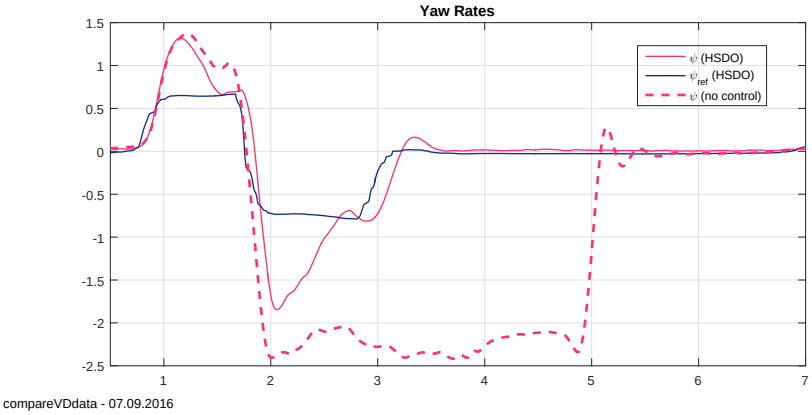

Figure 6-4: Comparison of yaw rate during an uncontrolled lane change and the HSDO motor-only control allocator. For the uncontrolled case, the vehicle starts spinning out at the 2-second mark and makes a full 180^0 rotation in the three following seconds

It can be seen that the VSC system developed in this thesis performs well compared to the Bosch VSC system currently on the car, given how limited the yaw moment authority of pure motor torque based yaw moment regulation is supposed to be. The overshoots are slightly larger in magnitude because of the motor's limited yaw moment authority, but in general behavior is not significantly different. This suggests there is room for improvement in the current implementation of the Bosch VSC system. It was shown using simulations that the HSDO control allocator effectively makes use of additional actuators when they are added to the system. It is expected that overall system performance will be improved when the HSDO control allocator gains access to brake torques. This however will be left for future work.

6-2-3 Observations during testing

During testing it was observed that the run-to-run variability was relatively large despite the initial conditions and inputs being very similar. This variability is mostly attributed to the fact that the maneuver is performed in a way that briefly puts the vehicle in an unstable state, so even minor differences in initial conditions or inputs can lead to large deviations in state trajectories and controller response, even within a short time-frame. This hypothesis was validated by looking for similarly variable behavior in simulation. This was indeed found to be the case. As illustration, another validation data set is shown in the appendix in Figure B-2. This has very similar inputs as Figure 6-1, but shows considerably different behavior during the latter part of the maneuver. The fact that both measurements and the SIL model show consistent behavior support the hypothesis that the variability is caused by the small variance in initial conditions, inputs and the nature of this test and not by inconsistent vehicle or algorithm behavior or other external circumstances. The downside of the variability means that the effect of different tuning parameters cannot be effectively analyzed without a significant number of repetitions per tuning set to get statistical significance. Due to time constraints testing was therefore left to one set of tuning parameters.

6-3 Conclusions

Since the experiments were limited to a single maneuver, it is difficult to draw any general conclusions. However, for this particular case it can be said that the developed simulation environment has sufficient fidelity to provide useful predictions of the combined performance of the high level controls and the HSDO control allocator. Furthermore, the implementation of the HSDO control allocator on the ECU proved straightforward and provided the predicted level of performance, suggesting that the HSDO control allocator can be considered a realistic option even on current generation vehicles to perform the control allocation. Finally, it was shown that the performance of the current implementation of the motor-only VSC system developed in this thesis is not far from from the performance of the brake-based Bosch VSC system. It is therefore, expected that if the HSDO control allocator also has access to the brake torques, overall system performance will be significantly improved.

Chapter 7

Conclusions and Recommendations

This chapter will cover the conclusions from each separate chapter and and provide an overall conclusion that reflects on whether the design goals set out in the introduction have been reached with the hybrid steepest descent optimization control allocator. Recommendation for further research into this methodology are also presented.

Vehicle model
The nonlinear two track model with HSRI tire models and instantaneous load transfer developed in this thesis proved to provide useful and accurate predictions of vehicle performance, both in the uncontrolled case as well as the closed loop experiments performed in Chapter 6.

Vehicle stability control architecture
Due to the lack of a common definition for control system architectures for vehicle stability control, an architecture suitable for the hybrid steepest descent control allocator needed to be designed. To allow for a detailed study of control allocation, the vehicle stability controller was divided into several control levels, leading to a modular control structure with clear functional separation between the levels. A high level controller interprets driver commands and determines which virtual control inputs are to be exerted on the vehicle. These virtual control inputs are sent to the control allocation module which attempts to best meet these requests by dynamically allocating control effort across the available actuators. The dynamic allocation comes courtesy of an optimization problem that is solved in the control allocation module. Here the optimization problem is formulated as an optimization problem with longitudinal acceleration and yaw moment request as virtual control inputs and additionally minimizing actuator effort. The actuator targets produced by the optimization in the control allocation module are fed to low level controllers that attempt to track these actuator targets and provide limiting action where necessary. Both the high and low level controllers were kept simple in this thesis to ensure consistent performance.

Hybrid steepest descent optimization
The hybrid steepest descent optimization technique was successfully applied to the control allocation problem outlined in Chapter 4. Two modifications are proposed to the algorithm

that aid convergence speed and reduce chatter. The convergence speed is aided by automatically scaling the constraints for the magnitude of the constraint violation and the chatter is reduced or eliminated by performing extra iterations to ensure the outcome of the optimization at every time step is feasible. A simple example demonstrated the effectiveness of both modifications. Furthermore, a method to calculate the yaw moment exerted by actuators was proposed and converted to fit into the hybrid steepest descent optimization framework by performing on-line linearization to determine yaw moment effectiveness for each individual actuator. Additionally, a modified ℓ_2-norm was suggested to penalize the actuator efforts. The use of this norm in conjunction with a regular ℓ_2-norm for other optimization objectives allows the optimization to naturally shift focus from one optimization objective to the other without the need for heuristics or fuzzy logic.

Simulation results
The simulation results clearly showed that the HDSO control allocator provides superior performance when compared to a much simpler, static control allocator for brake torque allocation. Furthermore, it was shown that the HDSO control allocator is capable of further improving performance if actuators are added. This was shown by the addition of the front and rear motors without needing to retune the system. Finally, the system was shown to be fault tolerant to a brake failure case and robust to a large surface friction estimation error, although, as expected, in both cases deteriorated performance was noted.

Experiments
To show that the simulations provide an accurate representation of real world performance the VSC system designed in this thesis was implemented on a Tesla Model S. The high level controller and control allocation algorithm were implemented in C-code with only minor modifications and executed aboard one of the Tesla ECU's, demonstrating real-time feasibility of the algorithm. Given the limited scope of the testing, with only motor torques available for control allocation and a single maneuver that was performed, conclusions from testing are not generalizable. However, for this particular test case the model proved to provide an accurate representation of the performance of the whole system, suggesting that the developed simulation environment has the ability to predict system performance, even during limit handling.

7-1 Overall conclusion

To re-iterate, the goal was set to design a control allocation algorithm that would:

1. achieve a high level of performance: done by fully exploiting control authority of the available actuators whilst minimizing the intrusiveness to the driver and actuator effort
2. ease of implementation: done by using a problem formulation and optimization technique that is computationally simple, transparent and easy to tune

The final product of this thesis, the hybrid steepest descent optimization control allocation algorithm fulfills these design goals. Some key developments that were proposed in this thesis are:

1. Dividing the VSC system in different modules, each with distinct functionality. This allows the control allocation algorithm to perform a well defined role where performance is easy to understand and tune

2. Setting the control allocation problem as a multi-objective optimization problem. This makes it possible to intuitively trade-off between tracking performance of the virtual control inputs and actuator efforts

3. The hybrid steepest descent optimization method. This update law optimization method has proven very suitable for performing the optimization executed in the control allocator, given its inherent ability to deal with moving constraints and its computational efficiency. A modification to this algorithm was proposed that accelerates convergence by scaling constraint gradients with the magnitude of the constraint violation. Another modification was proposed that strongly reduces chatter at constraint boundaries by forcing the algorithm to ensure a feasible solution at every time-step by allowing some iterations to occur during infeasibility

4. Defining the actuator yaw moment as the difference between the yaw acceleration in actuated and un-actuated state, and subsequently linearizing the resultant equations to make them suitable for use in the hybrid steepest descent optimization framework. This method allows the control allocator to understand which way the actuators should be moved to exert a certain yaw moment, and more importantly, how efficient each actuator is at providing the requested yaw moment

The overall performance of the hybrid steepest descent control allocator proved very encouraging. It provided a very clear performance advantage of the simple control allocator, it readily accepted the addition of extra actuators and still provided good performance if some actuators were disabled. All of this was achieved without retuning the control allocator, showing that it is a very easy and predictable algorithm to work with. Furthermore, given the relative computational simplicity of the hybrid steepest descent optimization method it was straightforward to implement it on an actual ECU of an existing passenger vehicle. Finally, the performance observed in simulation was echoed by experiments performed on the test vehicle. To the author's knowledge, these experiments were the first of their kind in literature that showed, using experiments, how a vehicle may be stabilized using solely motor torques on a vehicle with a single electric motor per axle.

7-2 Recommendations

For future work, some areas that should be focused on to further improve the control allocation algorithm:

1. Provide an indication of optimality of the solution found by the hybrid steepest descent optimization scheme. Given that this optimization method moves towards the optimum q^* over time, it cannot be guaranteed the solutions found by the optimization at a given time are optimal, especially given the fact that the optimum moves as well as the constraints. It is therefore interesting to know how close the hybrid descent optimization is from the optimal point and how optimization weights and sample time influence this

behavior. The choice for sample time was found empirically in this thesis, but a study should be performed to formalize this process of finding the right sample time to achieve a certain level of optimality

2. Add prediction to the hybrid steepest descent optimization problem to deal with actuator dynamics and actuator saturation. Although the performance of the system was found to be very good, there were cases where prediction could have helped find the optimal actuator configuration more efficiently. Prediction is also expected to help if actuators are added to the system with slow dynamics

3. Additional validation of performance should be performed. The level of validation performed in this thesis has been limited, and it would be advised to extend this validation to different maneuvers on different surfaces as well as including additional actuators in the control allocator, such as the brakes. Furthermore, although the control allocator was tested for robustness against a friction estimation error and found to work satisfactorily, it is advised to perform a more comprehensive study of robustness against parameter mismatch

Appendix A

Yaw Moment Effectiveness Calculations

This appendix will work out the two methods the were used to calculated both the yaw moment and yaw moment effectiveness for the simulation model and the actual implementation in the car for the experiments of Chapter 6. The implementation between the simulation and experimental case differs in that the experimental case does not make use of a lateral velocity U_y, given that that measurement was not available during the experiments.

A-1 Yaw moment and yaw moment effectiveness used in simulation model

The section will cover the calculations performed in simulation to estimate the yaw moment and yaw moment effectiveness values for the individual actuators. Re-iterating the general yaw moment equation Eq. (4-17):

$$\begin{aligned} I_{zz}\Delta\ddot{\psi} = M_z = \Big(&a(\Delta F_{x,fl} + \Delta F_{x,fr})\sin(\delta_f) + a(\Delta F_{y,fl} + \Delta F_{y,fr})\cos(\delta_f) + ... \\ &d/2(\Delta F_{x,rr} - \Delta F_{x,rl}) + d/2(\Delta F_{x,fr} - \Delta F_{x,fl})\cos(\delta_f) + ... \\ &d/2(\Delta F_{y,fl} - \Delta F_{y,fr})\sin(\delta_f) - b(\Delta F_{y,rl} + \Delta F_{y,rr})\Big) \end{aligned} \tag{A-1}$$

This was further simplified given that $d/2(\Delta F_{y,fl} - \Delta F_{y,fr})\sin(\delta_f)$ is relatively small compared to the other terms in nearly all conditions:

$$\begin{aligned} I_{zz}\Delta\ddot{\psi} = M_z = \Big(&a(\Delta F_{x,fl} + \Delta F_{x,fr})\sin(\delta_f) + a(\Delta F_{y,fl} + \Delta F_{y,fr})\cos(\delta_f) + ... \\ &d/2(\Delta F_{x,rr} - \Delta F_{x,rl}) + d/2(\Delta F_{x,fr} - \Delta F_{x,fl})\cos(\delta_f) + b(\Delta F_{y,rl} + \Delta F_{y,rr})\Big) \end{aligned} \tag{A-2}$$

with $\Delta F_i = F_i - F_i^{\mathbf{u}=0}$.

The next step is finding: $\Delta F_i = B_e(\mathbf{x}_0, \mathbf{u})$

A-1-1 Finding $B_e(\mathbf{x}_0, \mathbf{u})$

Remembering that the un-actuated state is defined as the state where $\mathbf{u} = 0$ and that the control inputs related to forces through:

$$F_{x,fl} = 1/r_w(T_{m,f}/2 - T_{br,fl}) \tag{A-3}$$
$$F_{x,fr} = 1/r_w(T_{m,f}/2 - T_{br,fr}) \tag{A-4}$$
$$F_{x,rl} = 1/r_w(T_{m,r}/2 - T_{br,rl}) \tag{A-5}$$
$$F_{x,rr} = 1/r_w(T_{m,r}/2 - T_{br,rr}) \tag{A-6}$$

This means that $\Delta F_{x,ij}$ can be calculated as:

$$\Delta F_{x,ij} = F_{x,ij} - F_{x,ij}^{\mathbf{u}=0}, \quad \frac{d\Delta F_{x,ij}}{dF_{x,ij}} = 1 \tag{A-7}$$
$$F_{x,ij}^{\mathbf{u}=0} = 0 \tag{A-8}$$

where some additional work was performed for the linearization by calculating the derivatives w.r.t. to $F_{x,ij}$.

Calculating $\Delta F_{y,ij}$ is quite straightforward as well when using sigmoid tires such as described in Chapter 2:

$$F_{y,ij} = \sqrt{F_{max,ij}^2 - F_{x,ij}^2}\tanh(C_{y,i}\alpha_{ij}), \quad i = \{f, r\}, \quad j = \{l, r\} \tag{A-9}$$

where, important to note, α_{ij} is only a function of $\delta_f, \dot{\psi}$ and U_y (Eq. (2-16)). Furthermore, $F_{max,ij} = \mu_s F_{z,ij}$ where $F_{z,ij}$ is calculated using the equations of Eq. (2-20).

Redefining $\tanh(C_{y,i}\alpha_{ij})$ as $\mu_{y,ij}$ and calculating $\Delta F_{y,ij}$ and the derivatives:

$$\Delta F_{y,ij} = \sqrt{F_{max,ij}^2 - F_{x,ij}^2}\mu_{y,ij} - F_{max,ij}\mu_{y,ij}, \quad \frac{d\Delta F_{y,ij}}{dF_{x,ij}} = \frac{-F_{x,ij}\mu_{y,ij}}{\sqrt{F_{max,ij}^2 - F_{x,ij}^2}} \tag{A-10}$$

Given these equations, the yaw moment can now be calculated directly using Eq. (A-2), given that all actuators inputs relate explicitly to ΔF_i and all the vehicle states are assumed to be known.

A-1-2 Calculating the actuator effectiveness $\bar{E}$

The first step is calculating:

$$E_{ij} = \frac{\partial M_{z,i}(\mathbf{x}_0, \mathbf{u}_0)}{\partial F_{x,ij}}, \quad i = \{f, r\}, \quad j = \{l, r\}$$

Applying this to Eq. (A-2) gives us:

$$E_{fl}^{F_x} = a\sin(\delta_f) + a\left(\frac{-F_{x,fl}\mu_{y,fl}}{\sqrt{F_{max,fl}^2 - F_{x,fl}^2}}\right)\cos(\delta_f) - d/2\cos(\delta_f) \tag{A-11}$$

$$E_{fr}^{F_x} = a\sin(\delta_f) + a\left(\frac{-F_{x,fr}\mu_{y,fr}}{\sqrt{F_{max,fr}^2 - F_{x,fr}^2}}\right)\cos(\delta_f) + d/2\cos(\delta_f) \tag{A-12}$$

$$E_{rl}^{F_x} = -b\left(\frac{-F_{x,rl}\mu_{y,rl}}{\sqrt{F_{max,rl}^2 - F_{x,rl}^2}}\right) - d/2 \tag{A-13}$$

$$E_{rr}^{F_x} = -b\left(\frac{-F_{x,rl}\mu_{y,rl}}{\sqrt{F_{max,rl}^2 - F_{x,rl}^2}}\right) - d/2 \tag{A-14}$$

This can be reordered to express it explicitly as function of the control inputs $\mathbf{u}$:

$$\begin{aligned} E_{T_{m,f}} = 1/r_w \Bigg(& a\sin(\delta_f) + a\left(\frac{-F_{x,fl}\mu_{y,fl}}{\sqrt{F_{max,fl}^2 - F_{x,fl}^2}}\right)\cos(\delta_f)/2 + \dots \\ & a\left(\frac{-F_{x,fr}\mu_{y,fr}}{\sqrt{F_{max,fr}^2 - F_{x,fr}^2}}\right)\cos(\delta_f)/2\Bigg) \end{aligned} \tag{A-15}$$

$$E_{T_{m,r}} = 1/r_w\left(-b\left(\frac{-F_{x,rl}\mu_{y,rl}}{\sqrt{F_{max,rl}^2 - F_{x,rl}^2}}\right)/2 - b\left(\frac{-F_{x,rr}\mu_{y,rr}}{\sqrt{F_{max,rr}^2 - F_{x,rr}^2}}\right)/2\right) \tag{A-16}$$

$$E_{T_{br,fl}} = 1/r_w\left(a\sin(\delta_f) + a\left(\frac{-F_{x,fl}\mu_{y,fl}}{\sqrt{F_{max,fl}^2 - F_{x,fl}^2}}\right)\cos(\delta_f) - d/2\cos(\delta_f)\right) \tag{A-17}$$

$$E_{T_{br,fr}} = 1/r_w\left(a\sin(\delta_f) + a\left(\frac{-F_{x,fr}\mu_{y,fr}}{\sqrt{F_{max,fr}^2 - F_{x,fr}^2}}\right)\cos(\delta_f) + d/2\cos(\delta_f)\right) \tag{A-18}$$

$$E_{T_{br,rl}} = 1/r_w\left(-b\left(\frac{-F_{x,rl}\mu_{y,rl}}{\sqrt{F_{max,rl}^2 - F_{x,rl}^2}}\right) - d/2\right) \tag{A-19}$$

$$E_{T_{br,rr}} = 1/r_w\left(-b\left(\frac{-F_{x,rl}\mu_{y,rl}}{\sqrt{F_{max,rl}^2 - F_{x,rl}^2}}\right) + d/2\right) \tag{A-20}$$

A-1-3 Modifications made to yaw moment and yaw moment effectiveness for experiments

Given that a measurement or accurate estimate of the lateral velocity U_y is usually lacking, some modifications are necessary. In the derivation above, this only affects $\mu_{y,ij}$, given that this value is dependent on α_{ij} which, in turn, is a function of U_y. This means an alternative method needs to be found to calculate $\mu_{y,ij}$ which will be explained in this section.

The lateral force on the front and rear axle can be estimated from the measurements that are typically available in passenger cars, requiring the lateral acceleration A_y, yaw rate $\dot{\psi}$, steering angle δ_f and an estimate of the yaw moment resulting from longitudinal forces M_x:

$$m_{car}A_y = F_{y,f}\cos(\delta_f) + F_{y,r} \tag{A-21}$$

$$I_{zz}\ddot{\psi} = aF_{y,f}\cos(\delta_f) - bF_{y,r} + M_x \tag{A-22}$$

where M_x can be calculated using Eq. (A-2) with all the lateral force contributions set to 0. With M_x known and measurements of A_y, yaw rate $\dot{\psi}$ and steering angle δ_f, the above equations can be solved for $F_{y,f}$ and $F_{y,r}$

The next step is to assume that the slip angles left to right are the same: $\alpha_{fl} \approx \alpha_{fr}$ and $\alpha_{rl} \approx \alpha_{rr}$. This is a good approximation for the speed ranges of interest (≈ 10 [m/s] or more), where the yaw rate is relatively low. Re-iterating that the lateral forces are calculated as:

$$F_{y,ij} = \sqrt{F_{max,ij}^2 - F_{x,ij}^2}\tanh(C_{y,i}\alpha_{ij}), \quad i = \{f, r\}, \quad j = \{l, r\} \tag{A-23}$$

So this means that $\tanh(C_{y,f}\alpha_{fl}) = \tanh(C_{y,f}\alpha_{fr})$ or $\mu_{y,fl} = \mu_{y,fr} = \mu_{y,f}$. Furthermore, given that $F_{y,f} = F_{y,fl} + F_{y,fr}$ it is possible to solve for $\mu_{y,f}$:

$$\mu_{y,f} = \frac{F_{y,f}}{\sqrt{F_{max,fl}^2 - F_{x,fl}^2} + \sqrt{F_{max,fr}^2 - F_{x,fr}^2}} \tag{A-24}$$

The derivation for $\mu_{y,r}$ is identical to the one presented above and yields a similar result. These estimates of μ_y can then be used in Eq. (A-10) for the yaw moment and yaw moment effectiveness calculations.

Appendix B

Additional Experimental Plots

An additional set of plots is shown of the experimental data that was collected during testing, but for another run of the same maneuver. The inputs and entry speed to the lane-change were nearly identical to those shown for the run of Figure 6-1, repeated below to allow for easier comparison. It can be seen that behavior is very similar between the two runs, until the last part (around the three-second mark) where suddenly quite large deviations show up. At this point, although not visible in the data, the rear end of the car was right at the point of saturation, beginning to slide out. It is expected that the severity of this slide, although relatively minor in duration is what causes the relatively large differences between the two runs.

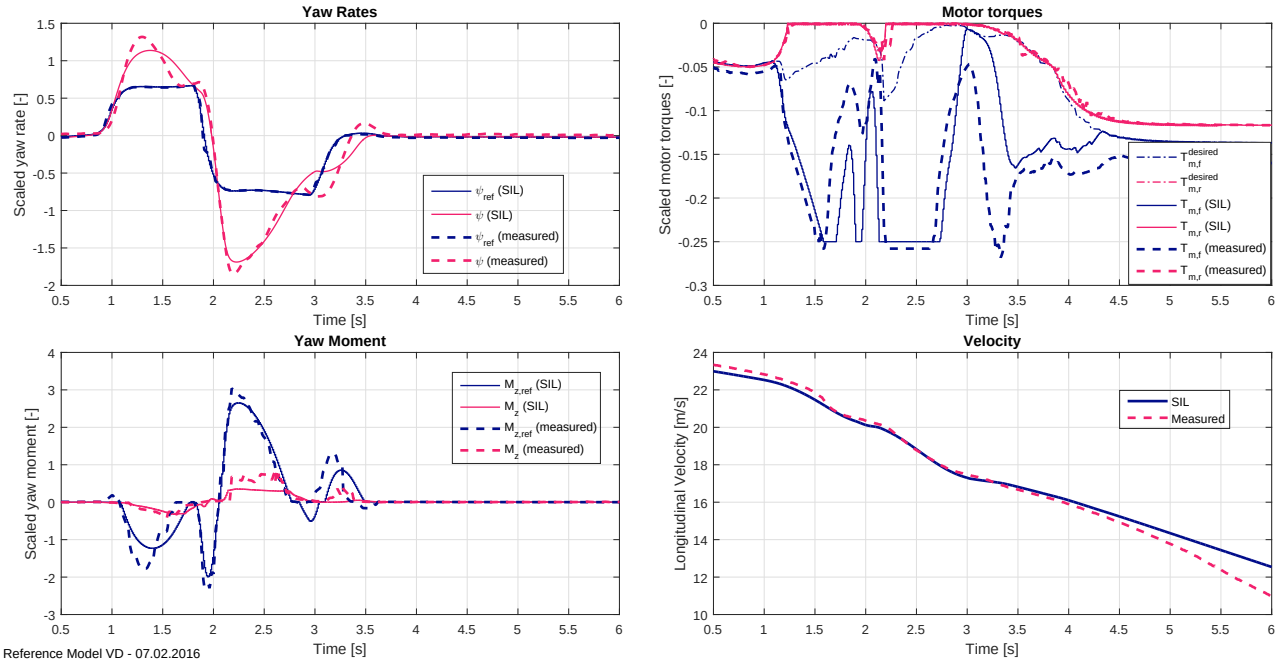

Figure B-1: Comparison between measured and simulation data, same as Figure 6-1

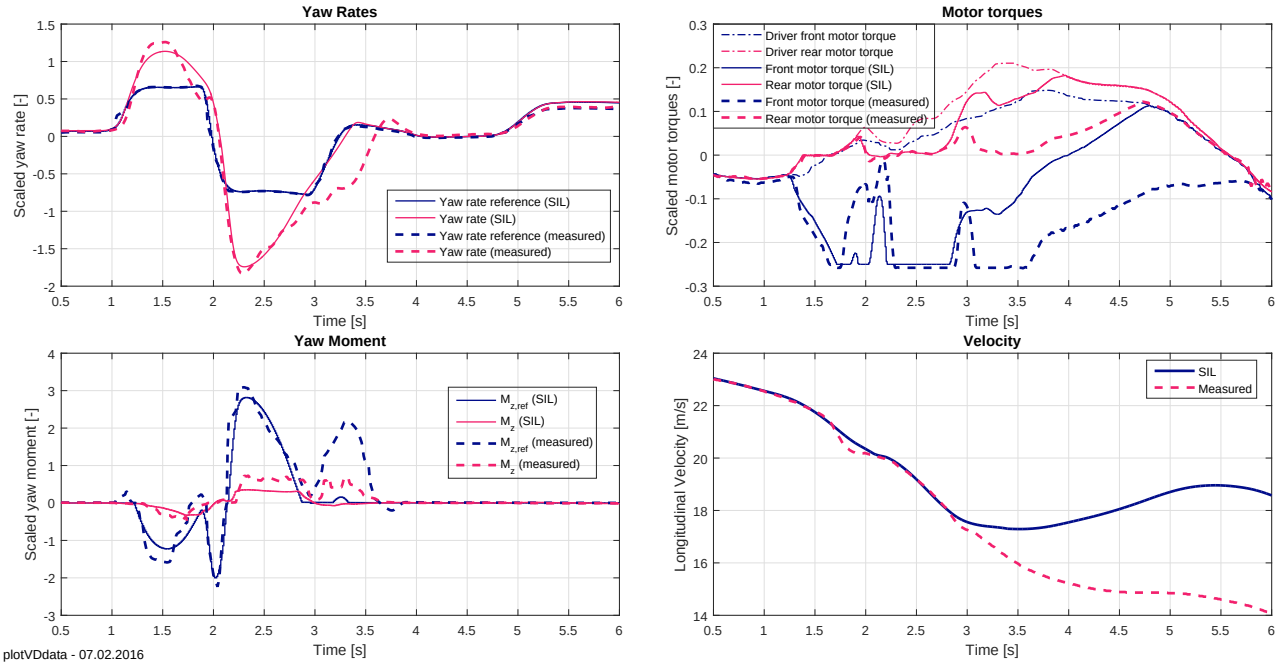

Figure B-2: Comparison between measured and simulation data of addition run

Bibliography

[1] N. C. for Statistics and Analysis, "Traffic safety facts 2012 data: Passenger vehicles," September 2014. http://www.nhtsa.gov/ [Online; posted September 2014].

[2] E. Liebemann, K. Meder, J. Schuh, and G. Nenninger, "Safety and performance enhancement: The bosch electronic stability control (esp)," *SAE Paper*, vol. 20004, pp. 21–0060, 2004.

[3] M. Gerard, *Global chassis control and braking control using tyre forces measurement*. TU Delft, Delft University of Technology, 2011.

[4] C. E. Beal, *Applications of model predictive control to vehicle dynamics for active safety and stability*. Stanford University, 2011.

[5] C. G. Bobier, *A Phase Portrait Approach to Vehicle Stabilization and Envelope Control*. PhD thesis, Stanford University, 2012.

[6] S. Inagaki, I. Kshiro, and M. Yamamoto, "Analysis on vehicle stability in critical cornering using phase-plane method," in *International Symposium on Advanced Vehicle Control (1994: Tsukuba-shi, Japan). Proceedings of the International Symposium on Advanced Vehicle Control 1994*, 1994.

[7] A. T. Van Zanten, "Evolution of electronic control systems for improving the vehicle dynamic behavior," in *Proceedings of the 6th International Symposium on Advanced Vehicle Control*, pp. 1–9, 2002.

[8] H. E. Tseng, B. Ashrafi, D. Madau, T. A. Brown, and D. Recker, "The development of vehicle stability control at ford: Focused section on mechatronics in automotive systems," *IEEE/ASME transactions on mechatronics*, vol. 4, no. 3, pp. 223–234, 1999.

[9] M. Canale, L. Fagiano, A. Ferrara, and C. Vecchio, "Vehicle yaw control via second-order sliding-mode technique," *Industrial Electronics, IEEE Transactions on*, vol. 55, pp. 3908–3916, Nov 2008.

[10] D. Piyabongkarn, J. Lew, R. Rajamani, J. Grogg, and Q. Yuan, "On the use of torque-biasing systems for electronic stability control: Limitations and possibilities," *Control Systems Technology, IEEE Transactions on*, vol. 15, pp. 581–589, May 2007.

[11] T. Yoshioka, T. Adachi, T. Butsuen, H. Okazaki, and H. Mochizuki, "Application of sliding-mode theory to direct yaw-moment control," *JSAE review*, vol. 20, no. 4, pp. 523–529, 1999.

[12] M. Canale, L. Fagiano, A. Ferrara, and C. Vecchio, "Vehicle yaw control via second-order sliding-mode technique," *Industrial Electronics, IEEE Transactions on*, vol. 55, pp. 3908–3916, Nov 2008.

[13] W. Manning and D. Crolla, "A review of yaw rate and sideslip controllers for passenger vehicles," *Transactions of the Institute of Measurement and Control*, vol. 29, no. 2, pp. 117–135, 2007.

[14] H. Smakman, *Functional integration of slip control with active suspension for improved lateral vehicle dynamics.* PhD thesis, Delft University of Technology, 2000.

[15] M. Jonasson, J. Andreasson, B. Jacobson, and A. S. Trigell, "Global force potential of over-actuated vehicles," *Vehicle System Dynamics*, vol. 48, no. 9, pp. 983–998, 2010.

[16] Y. Luo, K. Cao, and K. Li, "Coordinated control of longitudinal/lateral/vertical tire forces for distributed electric vehicles," in *American Control Conference (ACC), 2014*, pp. 3905–3910, June 2014.

[17] A. Masato, "Trends of intelligent vehicle dynamics controls and their future," *NTN TECHNICAL REVIEW*, vol. 81, 2013.

[18] C. Vermillion, J. Sun, and K. Butts, "Model predictive control allocation for overactuated systems - stability and performance," in *Decision and Control, 2007 46th IEEE Conference on*, pp. 1251–1256, Dec 2007.

[19] M. W. Oppenheimer, D. B. Doman, M. Bolender, *et al.*, "Control allocation for over-actuated systems," in *Control and Automation, 2006. MED'06. 14th Mediterranean Conference on*, pp. 1–6, IEEE, 2006.

[20] A. Shaout and K. McGirr, "Real-time systems in automotive applications: Vehicle stability control," *Electrical Engineering Research*, vol. 1, no. 4, pp. 83–95, 2013.

[21] M. Bodson, "Evaluation of optimization methods for control allocation," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 4, pp. 703–711, 2002.

[22] O. Härkegård, "Dynamic control allocation using constrained quadratic programming," *Journal of Guidance, Control, and Dynamics*, vol. 27, no. 6, pp. 1028–1034, 2004.

[23] J. H. Sill, *Predictive tire force saturation management for vehicle stability control.* PhD thesis, CLEMSON UNIVERSITY, 2012.

[24] C. Satzger, R. de Castro, and T. Bünte, "A model predictive control allocation approach to hybrid braking of electric vehicles," in *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pp. 286–292, IEEE, 2014.

[25] P. Tondel and T. A. Johansen, "Control allocation for yaw stabilization in automotive vehicles using multiparametric nonlinear programming," in *Proceedings of the American Control Conference*, vol. 1, p. 453, 2005.

[26] F. Borrelli, A. Bemporad, M. Fodor, and D. Hrovat, "An mpc/hybrid system approach to traction control," *Control Systems Technology, IEEE Transactions on*, vol. 14, pp. 541–552, May 2006.

[27] A. Bemporad, "Explicit model predictive control," 2009. [Online; accessed 10-June-2015].

[28] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.

[29] M. Abe, *Vehicle handling dynamics: theory and application*. Butterworth-Heinemann, 2015.

[30] H. Dugoff, P. S. Fancher, and L. Segel, "Tire performance characteristics affecting vehicle response to steering and braking control inputs," tech. rep., Highway Safety Research Institute, 1969.

[31] A. Hac and M. O. Bodie, "Improvements in vehicle handling through integrated control of chassis systems," *International journal of vehicle autonomous systems*, vol. 1, no. 1, pp. 83–110, 2002.

[32] X. Yang, Z. Wang, and W. Peng, "Coordinated control of afs and dyc for vehicle handling and stability based on optimal guaranteed cost theory," *Vehicle System Dynamics*, vol. 47, no. 1, pp. 57–79, 2009.

[33] J. Svendenius, *Tire Modeling and Friction Estimation*. PhD thesis, Lund University, 2007.

[34] C. De Wit and P. Tsiotras, "Dynamic tire friction models for vehicle traction control," 1999.

[35] W. F. Milliken and D. L. Milliken, *Race car vehicle dynamics*, vol. 400. Society of Automotive Engineers Warrendale, 1995.

[36] S. Chang and T. J. Gordon, "Model-based predictive control of vehicle dynamics," *International Journal of Vehicle Autonomous Systems*, vol. 5, no. 1-2, pp. 3–27, 2007.

[37] T. Pilutti, G. Ulsoy, and D. Hrovat, "Vehicle steering intervention through differential braking," *Journal of dynamic systems, measurement, and control*, vol. 120, no. 3, pp. 314–321, 1998.

[38] F. Yu, D.-F. Li, and D. Crolla, "Integrated vehicle dynamics control x2014;state-of-the art review," in *Vehicle Power and Propulsion Conference, 2008. VPPC '08. IEEE*, pp. 1–6, Sept 2008.

[39] P. Falcone, H. Eric Tseng, F. Borrelli, J. Asgari, and D. Hrovat, "Mpc-based yaw and lateral stabilisation via active front steering and braking," *Vehicle System Dynamics*, vol. 46, no. S1, pp. 611–628, 2008.

[40] W. Cho, J. Choi, C. Kim, S. Choi, and K. Yi, "Unified chassis control for the improvement of agility, maneuverability, and lateral stability," *Vehicular Technology, IEEE Transactions on*, vol. 61, pp. 1008–1020, March 2012.

[41] ISO, "Standards catalogue," 2016.

Glossary

List of Acronyms

VSC	Vehicle Stability Control
VDA	Vehicle Dynamics Area
GCC	Global Chassis Control
ESP	Electronic Stability Program
ESC	Electronic Stability Control
ECU	Electronic Control Unit
TLLTD	Tire Lateral Load Transfer Distribution
CG	Center of Gravity
HSDO	Hybrid Steepest Descent Optimization
SIL	Software in the Loop

List of Symbols

Abbreviations

α	Scaling value for cost function
α_{ij}	Slip angle of one of the four tires
$\bar{E}$	Actuator effectiveness estimates
β	Sideslip angle of vehicle
β_j	Scaling factor for constraint j
δ_f	Steering wheel angle
$\dot{\psi}$	Yaw rate of vehicle

γ_i	Switching function between constraints and cost function, as used in hybrid steepest descent method
λ_{ij}	Slip ratio of one of the four tires
$\mathbf{Q}$	Matrix with actuator weights
$\mathbf{R}$	Matrix with tracking error weights
$\mathbf{u}$	Actuator target for low level controls
$\mathbf{u}_{desired}$	Desired steady-state position of actuators
$\mathbf{v}$	Estimate of virtual control input
$\mathbf{v}_{ref}$	Reference virtual control input
μ_s	Surface friction coefficient
ω	Rotational velocity of wheel
a	Distance from center line of front axle to CG
a	Track width of vehicle
A_x	Longitudinal acceleration of vehicle
A_y	Lateral acceleration of vehicle
b	Distance from center line of rear axle to CG
C_x	Cornering stiffness of tire
C_y	Slip stiffness of tire
C_{st}	Steering system compliance
g	Gravitational constant
h_s	CG height
h_{rc}	Roll center height
K_d	Derivative gain for PID controller
K_i	Integral gain for PID controller
K_p	Proportional gain for PID controller
q	Total cost of cost function
q^*	Optimal cost
T_m	Motor torque
T_{br}	Brake torque
U_x	Longitudinal velocity of vehicle
U_y	Lateral velocity of vehicle