AN ELECTRO-HYDRAULIC TRACTION CONTROL SYSTEM FOR HEAVY DUTY OFF-ROAD VEHICLES: FORMULATION AND IMPLEMENTATION

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Addison B. Alexander

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

December 2018

Purdue University

West Lafayette, Indiana

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF DISSERTATION APPROVAL

Dr. Andrea Vacca, Chair

School of Mechanical Engineering

Dr. Gregory Shaver

School of Mechanical Engineering

Dr. John Lumkes

School of Agricultural and Biological Engineering

Dr. José Garcia Bravo

School of Agricultural and Biological Engineering

Approved by:

Dr. Jay P. Gore

Head of the School Graduate Program

To Stephanie, an island of tranquility in the tempestuous sea of my life.

ACKNOWLEDGMENTS

This work would not have been possible without countless hours of help from my mentor, Prof. Andrea Vacca. I am grateful for his tireless assistance and constant example of professional excellence. Second, I would be remiss if I did not thank Prof. Monika Ivantysynova, a bright light in the fluid power community, whose legacy we at the Maha Fluid Power Research Center strive to advance. I would also like to thank my family and friends for their love and support these long years. Without them, I would never have made it through. Finally, I thank God, "who is able to do immeasurably more than all [I] ask or imagine." To Him alone be the glory.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SYMBOLS

a	additive perturbation gain for extremum-seeking optimization
a_x	longitudinal vehicle acceleration [m/s^2]
A_i	area of component i [m^2]
B	tire longitudinal relaxation length [m]
B_x	stiffness factor for use in Magic Formula tire model
b_i	fraction of engine torque distributed to wheel i
C	system controller transfer function
C_x	shape factor for use in Magic Formula tire model
c_i	damping rate of component i [N·s/m or N·m·s/rad]
D_x	peak value for use in Magic Formula tire model
d	vehicle track width [m]
d_i	component distance i [m]
E_x	curvature factor for use in Magic Formula tire model
e	system error signal
F_i	force related to parameter i [N]
F_{Ni}	normal force at component i [N]
F_p	resistive force generated by a work pile [N]
F_{roll}	rolling resistance [N]
F_{resist}	other resistive forces (e.g. drag) [N]
F_{xi}	longitudinal force at wheel i [N]
F_z	vertical force [N]
G_i	closed-loop system transfer function from input i
g	acceleration of gravity (9.81 m/s^2)
g_i	system coordinate transformation i
J	objective function for optimization
K_i	gain of system component i
k_BT	braketorque ratio [N·m/bar]
k_i	spring rate of component i [N/m or N·m/rad]

I_i	moment of inertia of component i [kg·m^2]
L	approximation of function Hessian matrix for quasi-Newton method
l	vehicle wheelbase length [m]
l_i	length of split wheelbase section i [m]
M_y	moment about vehicle lateral axis [N·m]
m	mass [kg]
N_{diff}	gear ratio at vehicle differential
n	total number of system states
P	open-loop system plan transfer function
p	high-pass filter pole for extremum-seeking optimization [s^{-1}]
p_i	hydraulic pressure at component i [bar]
q	relative degree of system (feedback linearization)
R_i	steering radius i [m]
r	system reference trajectory
r_B	average braking torque lever arm [m]
r_e	effective radius of wheel [m]
s^k	optimization search direction at iteration k
T	total time in segment [s]
T_i	torque at component i [N·m]
t_i	length of time or point in time relating to parameter i [s]
u	system input
$v_{s,i}$	slip velocity at wheel i [m/s]
v	system virtual input
v_x	longitudinal vehicle velocity [m/s]
w_i	weighting value for system error (feedback linearization)
x	system state vector (x_i being state i)
x_i	longitudinal position of component i [m]
y	system output signal
z	transformed system state vector
z_i	vertical position of component i [m]

α_i	rotational acceleration of component i [rad/s^2]
β	optimization variable
Γ	adaptation law matrix (adaptive control)
γ	vehicle steering angle [rad]
γ^k	optimization step length at iteration k
δ_i	multi-objective function weight for objective i
ζ	system damping ratio
η	system internal states (feedback linearization)
$\eta_{tr,i}$	traction effectiveness at wheel i
θ	vector of uncertain system parameters (sliding mode control)
θ_i	rotational position of component i [rad]
κ_i	slip ratio at wheel i
λ_{ij}	absolute angle measure of the ray from point i to point j
μ_{roll}	rolling resistance coefficient
μ_{xi}	longitudinal friction force coefficient at wheel i
ξ	system external states (feedback linearization)
σ	boundary layer thickness for smoothed sliding mode control
ϕ	supplement of steering angle [rad]
χ	angle used in steering radius calculation [rad]
ψ	system time delay [s]
ω	perturbation frequency for ES optimization [rad/s]
ω_i	rotational velocity of component i [rad/s]
ω_n	system natural frequency [rad/s]

ABBREVIATIONS

ABS	anti-lock braking system
CAN	controller area network
ECU	engine control unit
EH	electro-hydraulic
ES	extremum seeking
GPS	global positioning system
LTI	linear time-invariant
MSE	mean squared error
PID	proportional-integral-derivative (controller)
SISO	single input-single output
SMC	sliding mode control(ler)
TC(S)	traction control (system)

ABSTRACT

Alexander, Addison B. Ph.D., Purdue University, December 2018. An Electro-Hydraulic Traction Control System for Heavy Duty Off-Road Vehicles: Formulation and Implementation. Major Professor: Andrea Vacca, School of Mechanical Engineering.

Traction control (TC) systems have become quite common in on-road passenger vehicles in recent years. However, for vehicles in other applications, they are not as widely available. This work presents a methodology for the proper design and implementation of a traction control system for heavy duty off-road machines, using a wheel loader as a reference vehicle.

A simulation model was developed, using standard vehicle dynamics constructs, including equations of motion and a description of the distribution of weight between the axles for different operating conditions. This model contains considerations for resistive forces acting on the machine implement, such as that generated by a work pile. The simulation also incorporates a detailed representation of the slip-friction characteristics between the vehicle tires and the road surface. One objective of this research was to model this interaction accurately, because the system traction behavior is dependent on it. Therefore, a series of tests was run using a state estimator to generate data on the slip-friction relationship at various ground conditions, and the results were incorporated into the simulation model. The dynamics of the machine braking system pressure were also modeled to give a more accurate description of the system response. The result is a mathematical model capable of accurately reproducing the behavior of the real-world system.

One of the primary goals of this work was the description of the traction control strategy itself, which should work as effectively and efficiently as possible. Several different aspects of the system were taken into consideration in generating this control structure. First, a relatively simple controller based on a PID control law was created. This controller was updated to account for peculiarities of the traction control system, as well as aspects like time delay. From there, more advanced controllers were created to address certain aspects of the system in greater detail. First, a self-tuning controller based on real-time optimization strategies was developed, to allow the controller to quickly adapt to changes in ground condition. Then, different nonlinear controllers were synthesized which were designed to address the

theoretical behavior of the system. All of these controllers were simulated using the system model and then some were run in experiments to show their potential for improving system performance. To improve system efficiency, the machine drivetrain itself was also examined to develop a more efficient control algorithm. By designing a more efficient methodology, traction control configurations which had previously seen increases in fuel consumption of 16% were now able to actually reduce fuel usage by 2.6%.

Another main goal of this work was the development of a prototype system capable of implementing the formulated control strategies. The reference machine was modified so that the brakes could be controlled electronically and independently for implementation of the TC system. The vehicle was instrumented using a wide array of sensors, and estimation methodologies for accurately determining vehicle speed and implement forces were designed. The velocity estimator designed in this work is more accurate and more reliable than an industry standard sensor, which is important for traction control implementation. The implement force estimate was also quite accurate, achieving payload estimate errors of less than 2.5%, comparable to commercially-available measurement systems. This setup allowed for tests to be accurately compared, to assess the traction control performance.

With the objective of performing experiments on the traction control system, many tests were run to assess its capabilities in various situations. These tests included experiments for characterizing the vehicle behavior so that the simulation model could be updated to accurately reflect the physical machine performance. Another task for the experimental work was the generation of useful metrics for quantifying traction control performance. Laboratory experiments which were very controlled and repeatable were also run for generating data to improve the system model and for comparing traction control performance results side-by-side. The test metrics proposed for these experiments provided for accurate, repeatable comparisons of pushing force, tire wear, and brake consumption. For each of these tests, the traction control system saw an increase in pushing force of at least 10% when compared with the stock machine, with certain operating conditions showing increases as high as 60%. Furthermore, every test case showed a decrease in wheel slip of at least 45% (up to 73% for some cases), which translates into increased tire longevity.

Other tests were conducted in the field, designed to mimic the real-world operating conditions of the wheel loader. Various performance comparisons were made for different configurations in which traction control could provide potential benefits. These included parameters for comparing overall vehicle performance in a typical truck loading cycle, such as tire wear, fuel consumption, and material moved per load. Initial results for this testing showed a positive result in terms of wheel slip reduction, but other performance parameters such as fuel consumption were negatively impacted. Therefore, the control structure was reexamined extensively and new methods were added to improve those results. The final control implementation saw a 12% reduction in tire slip, while also reducing fuel consumption by 2.6% compared to the stock system. These results show significant potential for traction control as a technology for maximizing the performance output of construction machines.

1. INTRODUCTION

The development of well-designed vehicle systems always maintains in tension three distinct goals: safety, productivity, and cost. Often these objectives tend to work against each other, with solutions that improve one factor negatively impacting the other two. However, innovation in vehicle systems occurs when systems are developed which can simultaneously benefit all three of these metrics.

Traction control (TC) systems are an example of such a solution, which can have benefits in terms of all three metrics. For this reason, it has gained prominence over the past few decades in the field of on-road automobiles. Many modern passenger vehicles incorporate TC systems, which are touted as state-of-the-art safety features. More recently, academic and industry research groups in other fields have begun investigating the feasibility of incorporating traction control into their existing systems. This is particularly true of off-road vehicles such as construction equipment (Fig. 1.1), where safety, productivity, and cost are emphasized to a heightened degree, due to the breakneck operating pace and the immense loads and forces which are involved in their day-to-day operations in many industries.

1.1 Wheel Slip in Off-Road Machines

The importance of a properly-developed TC system for off-road equipment is difficult to overstate. There are many factors that such a mechanism would improve. From a safety perspective, the ability to maintain a proper traction with the ground surface is crucial. If the tires (or tracks, etc. for other machines) are slipping significantly against the ground, the operator is not able to control the system precisely, increasing the probability that the machine will collide with other nearby objects. From a cost perspective, significant wheel slip poses a problem in terms of wear, which is increased as the tire abrades more against the ground surface. Furthermore, traction control can be beneficial when driving in conditions like steep inclines. On top of the safety benefits of simply increasing the system controllability, the TC system can expand the range of ground conditions in which the machine can reliably operate, thereby also potentially increasing system productivity.

This possibility to improve safety, productivity, and cost makes traction control a particularly appealing technology for construction equipment manufacturers. In fact, some manufacturers have already implemented traction control systems on existing product lines [1, 2].

Figure 1.1. Examples of off-road machines.

Off-road construction machines represent a unique case for traction control system design, due to several peculiarities. First, construction machines tend to operate at relatively low velocities, often at or near zero. This causes issues in terms of simulation modeling, as traditional vehicle dynamics equations are often designed for vehicles in motion. Second, construction machines have unusual dynamics, with specific working conditions and external forces which can vary widely from moment to moment. On top of this, they are frequently used to transport materials, meaning the mass of the machine itself can change significantly and rapidly. This makes them significantly different than highway vehicles, which typically have very few external loads (mostly resistant forces like drag and rolling resistance) that are relatively well-behaved and simple to model. Finally, construction machines must drive on a very wide range of ground conditions. Whereas passenger automobiles stay mainly on asphalt and concrete (with infrequent exceptions), off-road equipment is routinely driven on

surfaces like mud, snow, gravel, and other diverse materials. This means that any traction control strategy for these machines must be robust enough to satisfy performance criteria under a wide array of operating conditions. All of these contribute to the fact that designing a TC system for construction machines is not as simple as mimicking the current systems used in standard automobiles [3].

1.2 Reference Machine

The machine used as a reference vehicle in this work is a wheel loader. Wheel loaders are large front-loading tractors used primarily for digging and transporting material at work sites. They can also be used for pushing and pulling in applications such as demolition, ground grading, and snow removal. In all of these operating modes, large forces and suboptimal ground conditions create the potential for significant wheel slip. Because of the prevalence of these machines and their propensity for wheel slip, they represent one of the most significant cases for the implementation of traction control. The incorporation of a TC system will aid the operation of the wheel loader, potentially cutting down costs (such as tires and fuel), while increasing productivity (in terms of controllability and tractive force). Despite the relatively narrow focus on wheel loaders for this research, the work presented here is

Figure 1.2. Wheel loader used as a reference vehicle in this research.

applicable to a wide array of off-road systems. The overall simulation process and controls methods are generally valid for many different machines.

A Case 721F 14-metric ton wheel loader (Fig. 1.2) was provided to the research team for use as the development prototype. It is equipped with a standard-configuration boom and bucket as the front implement, which is actuated by a high-pressure hydraulic circuit. Power is transferred from the combustion engine to the wheels through use of a torque converter, a locked transfer case to transmit torque to the driveshaft, and open differentials at each axle.

The stock braking system on the vehicle is a relatively simple hydraulic circuit wherein the pressure at all four brakes is controlled by the position of a valve mechanically linked to the brake pedal position. In order to create a prototype capable of implementing traction control, the stock braking system was modified to allow for electronic control of the braking pressure in parallel with the existing circuit. This electro-hydraulic braking system enables the implementation of the TC methods described in this report.

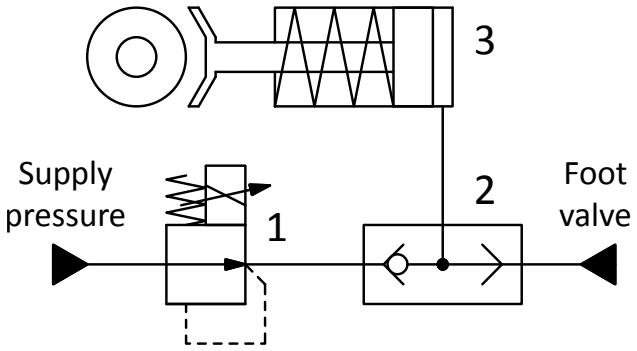

Figure 1.3. Simplified hydraulic schematic of the electro-hydraulic braking setup installed at each wheel.

Figure 1.3 shows a simplified hydraulic schematic of the modified system. The details of the implemented system are protected by confidentiality, but a general description is included here. The machine operator is able to electronically command a braking pressure through the pressure reducing valve added into the braking circuit (labeled "1" in Fig. 1.3). Due to the shuttle valve (labeled "2"), the braking cylinder ("3") ultimately sees whichever commanded pressure is higher (between the foot valve and the pressure reducing valve). Therefore, the electronic command is only able to *increase* the braking pressure at the braking cylinder

above the pressure commanded by the foot valve. This means that the modified circuit does not decrease the safety of the original system by allowing the braking pressure to be lowered.

1.3 Research Goals

The goal of this work is to develop and describe a sound methodology for the design of traction control systems for heavy duty off-road equipment. To achieve this, the following objectives were defined:

- Realize a prototype hydraulic system architecture that allows for safe, independent actuation of the individual wheel brakes.
- Determine the behavior of the tire-ground surface interaction, with the purpose of identifying the best conditions for tire slip.
- Develop and simulate controllers that allow for optimal tire slip, including a controller capable of tuning itself to that best condition.
- Establish a method for enacting efficient wheel slip control by taking advantage of the vehicle drivetrain architecture, with considerations for engine torque control.
- Define suitable metrics for assessing the performance of such a traction control system.
- Design and perform experiments to demonstrate the effectiveness of the proposed controller implemented on the reference machine, accounting for operator variability.

To that end, one of the original contributions of this work is the creation of a vehicle dynamics simulation model. The simulation model is based on standard vehicle dynamics modeling methods, with appropriate simplifications and assumptions. Model parameters were chosen based on experimental data to ensure that the model matched the machine behavior as closely as possible. The simulation makes it possible to understand the vehicle system behavior under a wide range of operating conditions without having to conduct tests for each case.

Using the simulation model, a control system is developed using common controls methodologies, with some adjustments for effective real-world implementation. The idea is to create

a controller capable of actuating the brakes independently at each wheel in order to keep the friction between ground and tire as high as possible. Different controller designs and their performance for several types of terrains are also investigated. Along with the proposed controller, other advanced topics are discussed, such as a novel design that implements a self-tuning algorithm which enables the vehicle to work at the optimal tire slip, which maximizes the propulsion force of the machine in any ground condition.

The controls methods were then tested using the simulation model to understand their potential for increasing system performance. A wide array of simulations were run to assess several different facets of the machine behavior. Based on these simulation results, the control systems were modified to increase the performance of the system and see what the potential limits and pitfalls of each strategy are.

Finally, the system was implemented on a physical system in order to assess the tangible benefits of traction control in realistic conditions. The reference machine available to the research team was modified to create a prototype capable of braking the wheels electronically, and a computer code and control system were installed in order to test the controllers which were previously assessed in simulation. Based on performance feedback, the control methods were again altered to optimize their performance. The real-world performance of the machine with traction control implemented is then compared to that of the standard vehicle system, and the results are discussed.

In order to demonstrate how the research met these goals, this dissertation has the following structure. Chapter 2 lays out the current state of the art in vehicle dynamics modeling and traction control design, in order to put the present work in context and lay the foundation for the rest of the document. Chapter 3 describes the simulation model which was built to characterize the machine behavior, including braking hydraulics and a method for determining implement forces using only position and pressure measurements. Chapter 4 outlines the generation of the Augmented PID control, which was the primary TC strategy used in experimental investigations for this vehicle. Chapters 5 and 6 expand on this by introducing an online setpoint optimization procedure, as well as more advanced nonlinear control methods which are capable of yielding better results than the Augmented PID, although they increase the system complexity significantly. Chapter 7 describes the

experimental setup needed to conduct the tests in the subsequent chapters. Chapters 8 and 9 discuss the experiments themselves and important results, starting with more controlled laboratory testing and proceeding on to the more realistic field tests. Finally, Chapters 10 and 11 conclude the discussion and provide suggestions of possibilities for expanding on this work in the future.

2. STATE OF THE ART AND BACKGROUND

This chapter examines the past research conducted into traction control systems. Though the present investigation has some novel aspects, it is built on a foundation laid by several previous works from various fields. This includes contributions from automotive design, vehicle dynamics, and control systems, among others.

2.1 Traction Control Systems (TCS)

Reif [4] discusses in detail the current framework around which modern automobile traction control systems are designed. Of course, the basic philosophy behind these systems is to prevent excessive slip of the wheels in an attempt to maximize the propulsion force of the automobile. They can be considered in many ways to be the mirror image of the anti-lock braking system (ABS) [5]. Often these TC strategies involve examining the individual wheel speeds to see if one of them exceeds a certain limit, and then applying brakes at the appropriate wheels. Early TC systems used electrical architectures like comparative circuits to generate control signals for the appropriate brakes [6]. As on-board computers became more prevalent, their processing capabilities were leveraged to monitor and control wheel traction.

There are many existing methods for determining the proper wheel slip reduction strategy. For two-wheel drive vehicles, the imposed wheel speed limit is often set using the velocity of the non-driven wheels, as they are generally assumed not to be slipping against the driving surface [7]. Four-wheel drive vehicles (such as the wheel loader under investigation in this work), provide more of a challenge, as all four wheels can be slipping against the ground simultaneously. Therefore, TC systems created for these vehicles must incorporate a creative method for limiting wheel slip. One such method involves monitoring engine speed and requested torque along with wheel speeds, incorporating a model of the engine dynamics [8].

Similarly, several different strategies also exist for decreasing wheel slip once it has been identified. As alluded to above, the most common construction is the electro-hydraulic actuation of individual brakes to slow down slipping wheels [9]. Another solution involves the use of what is called a *limited-slip differential*, wherein the standard open differential is

locked so that the two wheels on the axle in question no longer spin completely independently of each other. In this way, the wheel which is in a higher friction condition will work to keep the other from slipping [10]. Furthermore, the proliferation of electric cars has enabled certain operations which were previously impractical, such as traction control through direct torque control of the motors driving the wheels [11, 12].

It should also be noted that TC systems are already present to some degree in commercially available construction machines [2, 13, 14]. This includes investigations into machines which use continuous tracks for propulsion [15]. However, the research work behind these systems is typically rudimentary or not widely available.

2.2 Vehicle Dynamics

Many different methods have been developed for modeling the dynamic behavior of vehicles [16]. The overall vehicle dynamics model incorporated into this work is based on models like those in Gillespie [17], Jazar [18], and Rajamani [19]. These vehicle dynamics models tend to follow the same pattern.

In general, vehicle dynamics modeling starts by describing the equations of motion of various vehicle components (typically the machine chassis and wheels). This is done through the use of Newton's Second Law of Motion. The system is broken into the *longitudinal* and *lateral* directions, which represent the forward-reverse direction of travel and the side-to-side transverse direction, respectively. Vertical forces and movements (i.e. normal to the driving surface) are also taken into consideration.

The equations of motion for the vehicle system are connected through reactive forces and torques which must also be described. Often, the longitudinal motion of the chassis is simply a function of the sum of forces arising from other components. On the other hand, the wheels in many vehicle systems are linked mechanically through one or more components in the transmission. The nature of this connection can vary from machine to machine, but many of the core components (differentials, gearboxes, torque converters) are common to most transmission setups. It is important to model these components properly, as they have a strong effect on the motion of the wheels.

The specific dynamics model employed can have a wide range of complexity depending on the application. Expanding the directions of motion and operating modes of the vehicle, among other aspects, can greatly increase the number of equations and states required to model the system, so it is vital to understand the purposes of a given simulation to avoid generating a model which is too simplistic for the application or too complicated to be implemented efficiently. For example, some models include detailed descriptions of dynamic handling performance of the vehicle [20, 21], while others examine the transmission system more closely [22]. More recently, quite a lot of attention has been given to modeling the driver himself through various means [23]. Regarding construction machines, some work has also been done with respect to developing the dynamics of heavy machines which operate on a wide range of surfaces [24, 25]. This work builds on some of those previous discussions.

2.3 Tire Fundamentals

For considering the traction at the tire-road interface, a more specific model is needed. As the purpose of the TCS is to prevent wheel macro-slip, it is important to study the effect of forces and torques applied to the wheel. Tire models allow the forces transmitted from the road to the vehicle to be studied, along with the interaction between the two elements.

There are several different methods of approaching the concept of tire modeling [26]. One of the most common approaches, and the one used in this work, is to develop a mathematical model of tire-road interactions which relies on replicating behavior patterns without necessarily building a conceptual construct purely from physics. Despite the prevalence of these models, however, other physics-based models do exist for more intensive applications. These include models based on finite-element analysis [27], models derived from materials science (which can include phenomena like membrane vibration) [28], and others [29].

2.3.1 Tire Slip

Wheels which are driven or braked do not maintain a pure rolling motion against the driving surface. In fact, any wheel which is transmitting a force onto a surface experiences some amount of *tire slip*, or simply *slip*. Tire slip arises due to the phenomena involved in

generating the wheel forces, many of which have to do with the material deformation of the tire [30–33]. A driven wheel will turn faster than its rate of travel; a braked wheel will turn slower than its rate of travel (Fig. 2.1). When this slip is controlled within the normal bounds of the wheel operation, it is known as *micro-slip*, while excessive tire slip which causes the system to lose traction force is called *macro-slip*. Typically, wheel slip which is visible to the naked eye falls into the category of macro-slip.

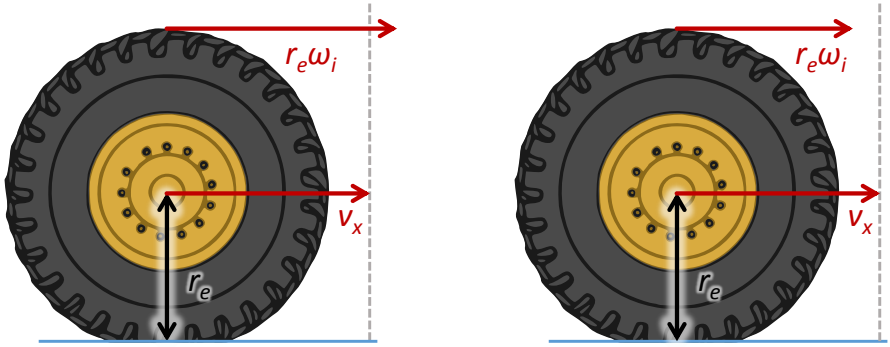

(a) (b)

Figure 2.1. Effect of torque applied to a rolling wheel for (a) a driven wheel ($r_e\omega_i > v_x$) and (b) a braked wheel ($r_e\omega_i < v_x$).

A common way to quantify tire slip is through use of the so-called *slip ratio*, represented by κ:

$$\kappa = \frac{r_e\omega}{v_x} - 1 \ , \ v_x \neq 0 \ , \tag{2.1}$$

in which ω represents the angular velocity of the wheel, r_e is its effective radius, and v_x is the longitudinal linear velocity of the vehicle. This linear velocity is compared to the theoretical linear wheel velocity $r_e\omega$. The slip rate with the formulation in (2.1) is defined on the interval $[-1, \infty)$ when ω and v_x have the same sign (i.e. the same direction). When $\kappa = 0$, there is no wheel slip (pure rolling), while a slip of $\kappa = -1$ represents a locked wheel sliding across the travel surface.

The expression in (2.1) does not include considerations of wheel force dynamics, which could cause the model to be inaccurate in situations where the wheel forces change rapidly. It

also encounters stability problems near zero velocity, because of the v_x term in the fraction denominator. In cases where low-speed calculations are necessary (a relatively common condition when dealing with construction equipment), it is better to define the slip ratio in another way to avoid this numerical instability. Bernard's formulation [34] defines the slip in a first-order differential equation rather than an algebraic equation.

$$\dot{\kappa} + \frac{|v_x|}{B}\kappa = \frac{|v_x| - r_e\omega\,\mathrm{sgn}(v_x)}{B} \ , \tag{2.2}$$

where B is the relaxation length of the tire. The *relaxation length* can be defined as the wheel arc length of the tire required to build up adequate friction to move the vehicle. In other words, it is the distance required for a wheel to reach the steady-state value of the friction force. It is needed as a "time constant" for the wheel to calm down the otherwise erratic behavior of the slip ratio, particularly at low velocity.

2.3.2 Friction Force and the Magic Formula

The prevalent tire friction models used in literature are those of slip-friction maps, which are defined as one-to-one relationship between the longitudinal slip κ_i at a given wheel and the resulting friction force F_{xi}. Examples of such slip-friction models can be found in [35–38]. Perhaps the most well known and widely used model to relate slip ratio and generated longitudinal force is Pacejka's model, also known as the *Magic Formula tire model* [39], which is the model used in this work. This formulation is a semi-empirical model, which means that it is not based on physics phenomena, but rather a mathematical formula representative of the shape of most slip-friction relationships. The parameters within this mathematical formula are adjusted to match experimental data. Therefore, for a given tire-road interface, the modeled force generated should be reasonably close to the real-world force for a certain tire slip.

For a given vehicle system, the longitudinal force at tire i, also known as *tractive force* or *pushing force*, is described by the relation:

$$F_{xi} = \mu_{xi}(\kappa_i)F_{Ni} \ , \tag{2.3}$$

where μ_{xi} is the coefficient of friction in the longitudinal direction and F_{Ni} is the normal force pressing the wheel down onto the road surface.

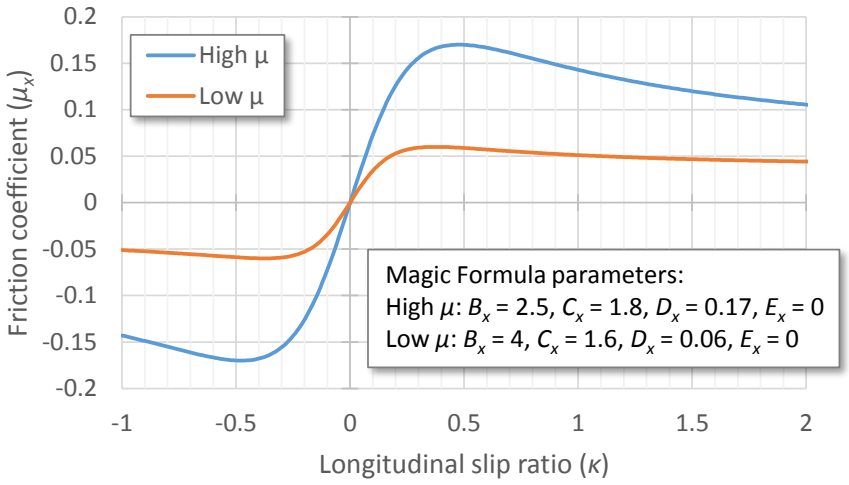

Figure 2.2. Examples of slip-friction curves.

The Pacejka model has the form:

$$\mu_x(\kappa) = D_x \sin\left(C_x \tan^{-1}\left[B_x\kappa - E_x(B_x\kappa - \tan^{-1}\left[B_x\kappa\right])\right]\right) , \tag{2.4}$$

with the following descriptions for the model parameters:

- B_x: Stiffness factor
- C_x: Shape factor
- D_x: Peak value
- E_x: Curvature factor

These parameters depend on the tire characteristics (such as material, tread type, tread depth, inflation pressure, and temperature), on the road conditions (e.g. dry asphalt, loose gravel, ice), and on the vehicle operation (velocity, loading, etc.). Typical variations of the tire-road friction profiles for different road surface conditions are represented in Fig. 2.2. As part of this work, data for slip-friction profiles for the wheel loader on different ground

conditions was taken to generate a more accurate tire force model (see Section 8.1). Dynamic tire behavior may be covered by this model as long as the frequency of the wheel motion remains well below the first natural frequencies of the belt with respect to the rim (typically less than around 15 Hz) [40].

Additionally, the slip-friction relationship (Fig. 2.2) has certain important features. The maximum value of the Magic Formula plot is the point where the longitudinal force is maximized (i.e. where it is possible to accelerate most quickly or push the strongest). Beyond that point, there is a region wherein the vehicle is in a visible slip condition (macro-slip), meaning that the wheel in question is wasting some amount of energy.

3. SYSTEM MODELING

Having discussed in general the background behind the work under consideration, it is now necessary to build more specifically the theory composing the modeling of the dynamic systems involved. This is crucial to all of the subsequent activities of this work, as the simulation model is used to help foster new ideas about controlling the traction more effectively and to evaluate proposed control structures before implementation on the real-world machine. Therefore, this chapter is devoted to the different elements involved in constructing the dynamic model of the reference machine.

The system model is built starting with the simplest components, and it grows in complexity from there. From a relatively simple setup including only some applications of Newton's Second Law of Motion, the model is expanded to include interactions among the wheels. A simplified braking dynamics model is also drawn up. These submodels together form the overall system plant, which was used extensively in both simulation and controller development activities.

3.1 Tire Modeling

In order to model the traction force as a function of wheel slip, Pacejka's Magic Formula (2.4) is used. The wheel slip itself is modeled in a manner based on the work done by Bernard [34]. The slip ratio dynamic equation formulated by Bernard (2.2) requires the implementation of a sign-switching algorithm to function correctly. In that model, positive slip yields forces acting always against the motion v_x of the vehicle, whereas negative slip generates forces acting in the same direction as the vehicle velocity. This requires that the longitudinal forces be multiplied by the sign of v_x and that when v_x switches signs during a spinout event or a stopping maneuver, the sign of the tire slip for implementation in the ongoing differential equation must be changed.

For the model implemented in this work, a formulation for the slip ratio has been used which does not match the exact formulation of Bernard but which is derived from that work. This definition is as follows:

$$\dot{\kappa}_i = \frac{1}{B}(-v_x + r_e\omega_i - |v_x|\,\kappa_i)\ , \tag{3.1}$$

wherein it is not necessary to utilize the sign-switching algorithm, as described above. This simplifies the implementation considerably, with positive slip always yielding a force in the forward direction of travel and negative slip always yielding a force in the reverse direction.

3.2 Quarter-Car Vehicle Model

The basic starting point for modeling vehicle dynamics is the quarter-car model, which is very common in literature because it allows the complicated system to be simplified heavily. This model consists of a mass m supported by a single wheel with radius r_e (Fig. 3.1).

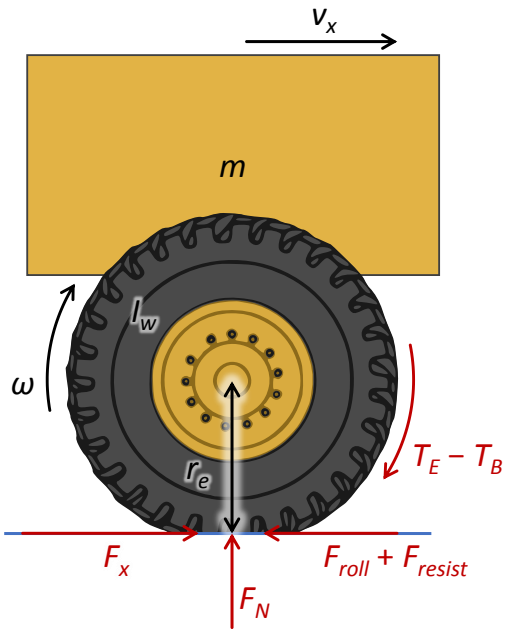

Figure 3.1. A quarter-car vehicle dynamics model.

Applying a torque to the wheel, the contact between tire and the ground creates a longitudinal force F_x that generates a moment which opposes the turning motion of the wheel. The sum of moments on the wheel causes an angular velocity ω_i. The various forces on the mass act to generate a velocity of the entire system v_x. The equations that describe the motion of the quarter car are:

$$m\dot{v}_x = F_x - F_{roll} - F_{resist} \tag{3.2}$$

$$I_w\dot{\omega}_i = T_E - T_B - r_e F_x \tag{3.3}$$

in which:

- F_x: friction force [N]
- F_{roll}: rolling resistance [N]
- F_{resist}: other resistive forces [N]
- I_w: wheel moment of inertia [kg·m^2]
- T_E: engine torque [N·m]
- T_B: braking torque [N·m]

Equations (3.2) and (3.3) come from Newton's Second Law. The first one represents the linear motion along the x-axis, where the force lost through the rolling resistance of the tire and other resistive forces (i.e. drag force, slope resistance) are subtracted from the longitudinal force F_x, which is attempting to propel the vehicle forward. The second equation, instead, evaluates the angular acceleration generated from the torques applied to the wheel: the torque coming from the engine, which is opposed by the braking torque, and the reaction moment created by the friction force at the tire-ground interface.

For the wheel loader system, F_{resist} has been neglected. This is due to the fact that most resistive forces are much lower in magnitude than the longitudinal force of the machine. For instance, air resistance is negligible at the low velocities seen by the wheel loader. When considering rolling resistance F_{roll}, there are many different methods which can be used [18,

41]. These models can become quite elaborate, but their added complexity generally comes with only a marginal improvement in accuracy. It is quite common in many applications to simply use a constant value of rolling resistance or a constant friction coefficient. This work uses a constant rolling resistance coefficient μ_{roll}, as shown in (3.4).

$$F_{roll} = \mu_{roll} F_N \operatorname{sgn}(v_x) \tag{3.4}$$

Values for rolling resistance coefficient commonly range between 0.008 and 0.3 for various tire constructions and ground surfaces [18].

3.3 Longitudinal Dynamics Model

In order to have a better understanding of the longitudinal weight distribution in a moving vehicle, a two-axle model for the reference machine was investigated, based on the work in Jazar [18]. This study is particularly important because there is a direct relationship between the normal force acting on the wheel F_N and the force generated along the longitudinal direction F_x, as seen in (2.3). In an accelerating vehicle, the normal force distribution at each wheel is changing over time, and this must be accounted for in calculating the tractive force at the wheels. For construction machines, this is further complicated by the inclusion of outside forces, such as the resistive force of a work pile.

3.3.1 Stationary Vehicle

When a vehicle is parked on a flat road (Fig. 3.2), the normal forces under the front and rear wheels (F_{N1} and F_{N2}) are:

$$F_{N1} = \frac{1}{2} mg \frac{l_2}{l} \tag{3.5}$$

$$F_{N2} = \frac{1}{2} mg \frac{l_1}{l} \ , \tag{3.6}$$

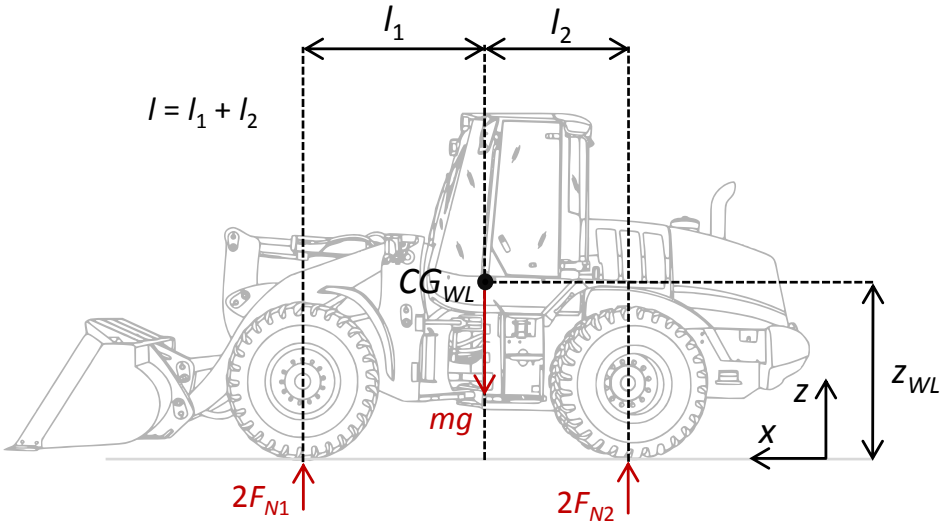

Figure 3.2. Stationary wheel loader on a level surface.

in which l_1 and l_2 are the distances from the wheel loader center of gravity CG_{WL} to the front and rear axles, respectively, and g is the acceleration of gravity. The distance l instead is the length of the entire wheelbase ($l = l_1 + l_2$).

These two equations come from a force equilibrium along the z-axis and a moment equilibrium about CG_{WL}:

$$\sum F_z = 0 \tag{3.7}$$

$$\sum M_y = 0 \tag{3.8}$$

Then, applying the static equilibrium equations:

$$2F_{N1} + 2F_{N2} - mg = 0 \tag{3.9}$$

$$-2F_{N1}l_1 + 2F_{N2}l_2 = 0\ . \tag{3.10}$$

Solving these two equations for F_{N1} and F_{N2} produces Equations (3.5) and (3.6).

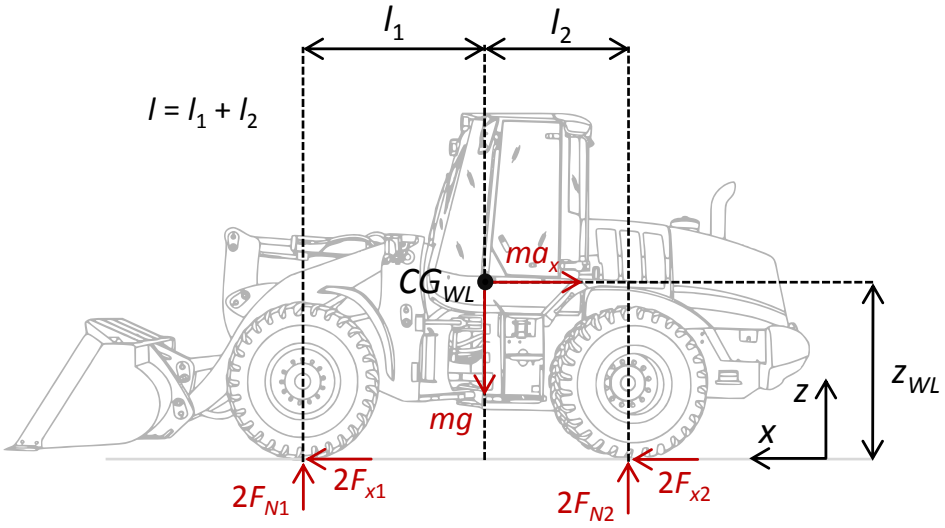

Figure 3.3. Wheel loader accelerating on a level surface.

3.3.2 Accelerating Vehicle

It is now possible to consider the case in which the vehicle is accelerating on a level road as shown in Fig. 3.3, using the same approach as before and applying the equilibrium equations along the y- and z-directions plus the linear dynamics equation in the x-direction:

$$\sum F_x = ma_x \tag{3.11}$$

$$\sum F_z = 0 \tag{3.12}$$

$$\sum M_y = 0 \tag{3.13}$$

Expanding the three equations:

$$2F_{x1} + 2F_{x2} = ma_x \tag{3.14}$$

$$2F_{N1} + 2F_{N2} - mg = 0 \tag{3.15}$$

$$-2F_{N1}l_1 + 2F_{N2}l_2 - 2z_{WL}(F_{x1} + F_{x2}) = 0\ , \tag{3.16}$$

there are three equations for four unknowns, but it is possible to eliminate $(F_{x1} + F_{x2})$ between (3.14) and (3.16) and then solve for F_{N1} and F_{N2}.

$$F_{N1} = \frac{1}{2}mg\frac{l_2}{l} - \frac{1}{2}ma_x\frac{z_{WL}}{l} \tag{3.17}$$

$$F_{N2} = \frac{1}{2}mg\frac{l_1}{l} - \frac{1}{2}ma_x\frac{z_{WL}}{l} \tag{3.18}$$

The normal forces are composed of two different components, a static part and a dynamic part. The static component is the weight distribution for a stationary vehicle, and it depends on the horizontal position of the center of gravity CG_{WL}. The dynamic part, instead, indicates how the normal forces are modified by the dynamics of the vehicle, and it depends on the linear acceleration of the vehicle a_x and the vertical position of CG_{WL} (labeled here as z_{WL}). During an acceleration phase, the force of inertia generates a moment about CG_{WL} that causes the rear axle load to increase and the front axle load to decrease.

3.3.3 Vehicle Pushing Against a Work Pile

For the purposes of this project, it is important to also study the case in which the off-road vehicle pushes against a material pile, to see how this affects the normal forces acting on the wheels. To this end, an additional pair of forces and a moment were applied to the wheel loader bucket (see Fig. 3.4). These represent the horizontal and vertical components of the pile force (F_{xp} and F_{zp}, respectively) acting on the bucket, as well as a correcting moment M_p. The forces and moment are assumed to be acting on the center of gravity of the wheel loader bucket CG_B. As any set of forces applied to a body can be represented as a force-moment pair acting on a single point [42], this formulation is capable of supporting any arbitrary pile force model. The development of the pile force model itself is outside the scope of this work, and simulations here instead focus on fairly simplistic resistive force generation models, like the one described in Section 3.7.

A complete kinematic model of the wheel loader implement system was created, allowing the simulation to accurately track the location of the bucket center of gravity (x_B, z_B), which changes with respect to the wheel loader center of gravity (x_{WL}, z_{WL}). This kinematic model is described in more detail in Section 3.8, as it is also used in force estimation during

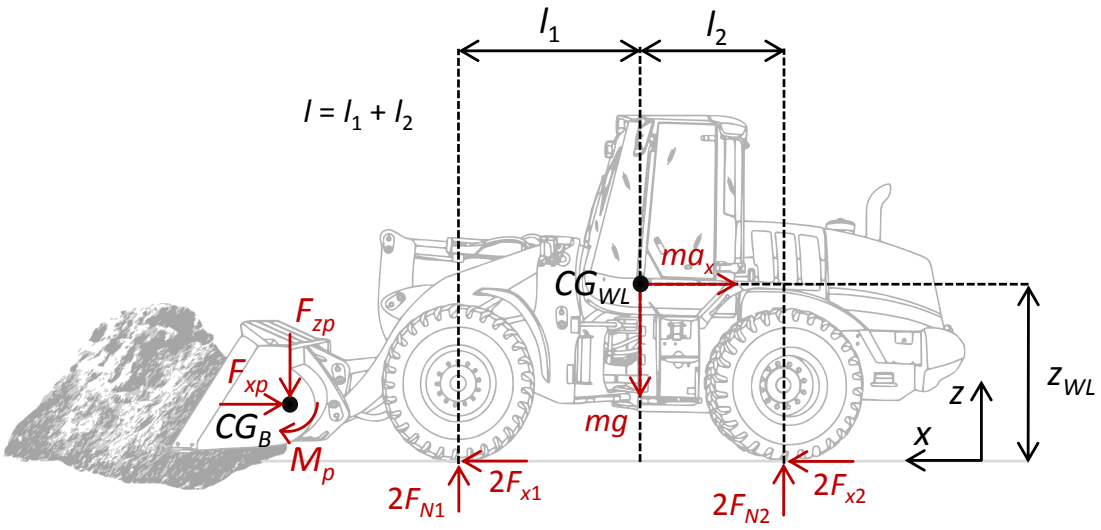

Figure 3.4. Wheel loader pushing against a work pile.

laboratory testing. The simulation model uses the length of the boom lift and bucket tilt cylinders as an input, applying standard trigonometric procedures to define the location of pertinent points on the implement. This kinematic model was based on geometric drawings of the machine components, and it was validated through experimental testing.

As before, expanding the three equilibrium equations and solving for the normal forces F_{N1} and F_{N2}, the following is found:

$$F_{N1} = \frac{1}{2l}\left(mgl_2 - ma_x z_{WL} - F_{xp} z_B + F_{zp}\left(x_B - x_{WL} + l_2\right) - M_p\right) \tag{3.19}$$

$$F_{N2} = \frac{1}{2l}\left(mgl_1 + ma_x z_{WL} + F_{xp} z_B - F_{zp}\left(x_B - x_{WL} - l_1\right) + M_p\right) \ , \tag{3.20}$$

in which there are three additional terms derived from the forces and moment created by the pile resistance force. These equations also take into account the planar location of the centers of gravity of the wheel loader CG_{WL} and the bucket CG_B. As the longitudinal resistive force F_{xp} and correction moment M_p generated by the pile increase, so does the normal force applied to the rear axle. On the other hand, as the vertical pile force F_{zp} (which can be thought of as a combination of resistive force and added material weight) increases, the share of the overall normal force sustained by the wheel loader shifts toward

the front axle. This makes sense intuitively, and it is consistent with experimental results generated in the course of this research.

3.4 Vehicle Transmission Model

In order to study how the dynamics of one wheel are connected to all the others, it is important to extend the model to include at least part of the vehicle drivetrain [43]. The term *drivetrain* describes the main components that deliver power from the engine to the road surface. In a four-wheel drive off-road vehicle, it includes the torque converter, transmission gearbox, driveshafts, differentials, and the final drive gears. The main purposes of the drivetrain are to decouple the engine rotation from the vehicle's wheels, distribute the torque to the wheels, and vary the direction and magnitude of the torque transmitted by the engine to the wheels.

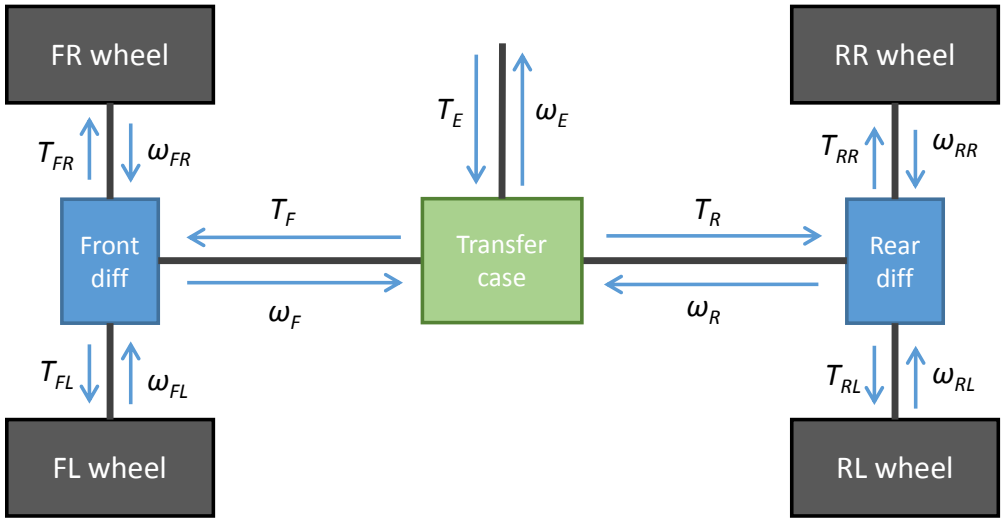

Figure 3.5. Flow of physical signals through the wheel loader transmission system.

Off-road vehicles can vary widely in terms of specific drivetrain components. The reference vehicle considered for this work (Fig. 3.5) has rear and front differentials, which connect the wheels on each axle together, and a single central transfer case, which links the front driveshaft to the rear driveshaft and both driveshafts to the engine via the torque converter.

The next two sections explain the approaches for modeling these two component blocks: the open differential and the locked transfer case.

3.4.1 Differential

The differential is needed in vehicles to allow the outer drive wheel to rotate faster than the inner wheel during turns. This is necessary to avoid either that the internal wheel rotates too quickly or that the outer wheel drags during turning maneuvers (due to the fact that the wheel on the outside has to roll farther than the inside wheel in the same amount of time). The resulting average speed between the two half axles connected to the driving wheels is equal to the angular velocity of the drive shaft. For a constant driveshaft speed, an increase in the speed of one wheel is balanced by a decrease in the speed of the other. There are different kinds of differentials, but this work considers the open differential, which is used on the reference vehicle. It is the most common type of differential in automotive systems.

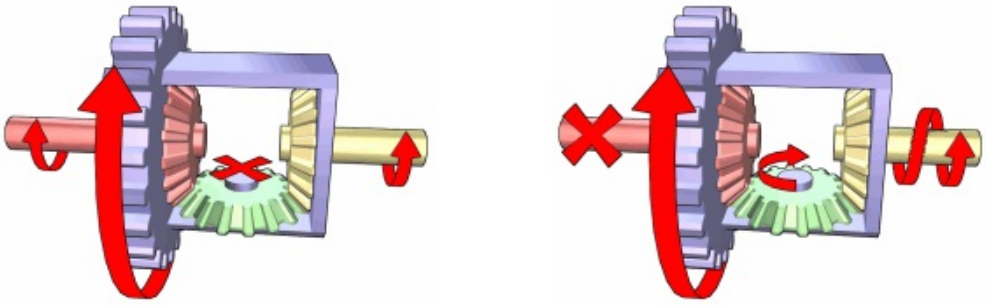

(a) (b)

Figure 3.6. Automotive differential operation for (a) the same resistance at both wheels and (b) more resistance at one wheel.

The differential has three functions:

- To deliver the engine torque evenly to both wheels on an axle
- To act as a gear reduction or increase for the vehicle drivetrain
- To allow the two wheels on an axle to rotate at different speeds

A differential model consists of one input, the driveshaft, and two outputs, which are the two driven wheels. The rotations of the driven wheels are determined by their equations of motion discussed above. The equations characterizing the differential's speed and torque relationships are:

Velocity-

$$\omega_{DS} = \frac{1}{2} N_{diff}(\omega_R + \omega_L) \tag{3.21}$$

Torque-

$$T_R = T_L = \frac{1}{2} N_{diff} T_{DS} \ , \tag{3.22}$$

where ω_{DS}, ω_R, and ω_L are the rotational velocities of the driveshaft, right half axle, and left half axle, respectively; T_{DS}, T_R, and T_L are the torques at the driveshaft, right half axle, and left half axle, respectively; and N_{diff} is the differential gear ratio.

3.4.2 Transfer Case

The transfer case's role is to distribute torque from the engine to the driveshaft. The transfer case of the reference vehicle is a *locked* transfer case, which means that the front and rear drive shafts do not behave as two different bodies, but instead they are connected mechanically. Because of the locked transfer case, the rotational motion of the front and rear driveshaft sections is identical, while the torque distribution to each axle can vary. This is different from the central differential found in many four-wheel drive passenger automobiles, which splits the input torque half to the front driveshaft and half to the rear.

Therefore the transfer case model must allow for the following:

- Front and rear driveshaft speeds should be equal (or roughly equal)
- Torque distribution should vary based on load at each wheel

Taking advantage of the work done by Tinker [44], the model of the locked transfer case is seen as a rotational spring-mass-damper system. The following is the governing equation for the rotational acceleration of the transfer case α_{TC}:

$$\alpha_{TC} = \frac{1}{I_{TC}}((T_{TC}) - (\theta_{TC} - \theta_F)k_F - (\omega_{TC} - \omega_F)c_F - (\theta_{TC} - \theta_R)k_R - (\omega_{TC} - \omega_R)c_R) \ , \tag{3.23}$$

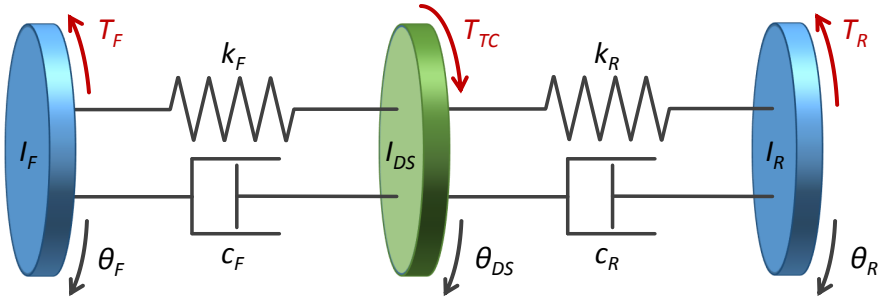

Figure 3.7. Transfer case and driveshaft model representation.

where θ_{TC}, θ_F, and θ_R are the rotational position of the transfer case, front drive shaft, and rear drive shaft; ω_{TC}, ω_F, and ω_R are their rotational velocities; I_{TC} is the transfer case equivalent inertia; k_F and k_R are the spring constants of the front and rear drive shafts; and c_F and c_R are the damping constants of the front and rear drive shafts, respectively.

This equation can be split to find the torque at each axle:

$$T_F = (\theta_{TC} - \theta_F)k_F - (\omega_{TC} - \omega_F)c_F \tag{3.24}$$

$$T_R = (\theta_{TC} - \theta_R)k_R - (\omega_{TC} - \omega_R)c_R \ . \tag{3.25}$$

Through a series of simplifications to take away the strong dependence on k_i and c_i and assuming equal k and c for each driveshaft section, the following torque distribution is reached:

$$T_F = \frac{1}{2}T_{TC} - \frac{k}{2}(\theta_F - \theta_R) - \frac{c}{2}(\omega_F - \omega_R) \tag{3.26}$$

$$T_R = \frac{1}{2}T_{TC} + \frac{k}{2}(\theta_F - \theta_R) + \frac{c}{2}(\omega_F - \omega_R) \ , \tag{3.27}$$

where T_F and T_R are the torques at the front and rear driveshaft sections, respectively.

By adjusting the values of parameters k and c so that they are sufficiently high, the above model allows the torque from the transfer case to be distributed unequally to the front and

rear axles, while keeping the rotational positions (and velocities) essentially equal at all three components.

3.5 Braking System Model

It has already been noted in this work that the reference wheel loader's braking system was modified to allow for electro-hydraulic independent control of the brakes (see schematic in Section 1.2). The basic layout of the braking system model is shown in Fig. 3.8. Using the command generated by the TC system controller (based on sensor data), the model incorporates the dynamics of the braking valve itself to generate a braking pressure which is then converted into braking torque.

Figure 3.8. Simulation model diagram for vehicle braking system.

This research is not especially focused on describing the behavior of the individual braking hydraulic components in detail. Instead, it is more important to understand how the hydraulic components affect the performance of the system at large. Therefore, the hydraulic braking system model contained herein functions simply as a way of defining the braking torque dynamics. There are many different ways to model braking systems, with some applications requiring increased complexity [3]. For the purposes of this study, it is assumed that the braking pressure p_B and torque T_B are proportionally related.

$$\begin{aligned} T_B &= \begin{cases} r_B A_B (p_B - p_{B,low}) \ , & p_B \geq p_{B,low} \\ 0 \ , & p_B < p_{B,low} \end{cases} \\ &= \begin{cases} k_{BT} (p_B - p_{B,low}) \ , & p_B \geq p_{B,low} \\ 0 \ , & p_B < p_{B,low} \end{cases} , \end{aligned} \tag{3.28}$$

where A_B is the area of the braking cylinder piston in contact with the brake disc, r_B is the distance from the center of the wheel to the average point of application of the braking cylinder force (i.e. the average braking torque lever arm), $k_{BT} = r_B A_B$ is the so-called *braketorque ratio*, and $p_{B,low}$ is the minimum pressure at which the brake calipers make contact with the disks. This relationship (shown in Fig. 3.9) was verified experimentally.

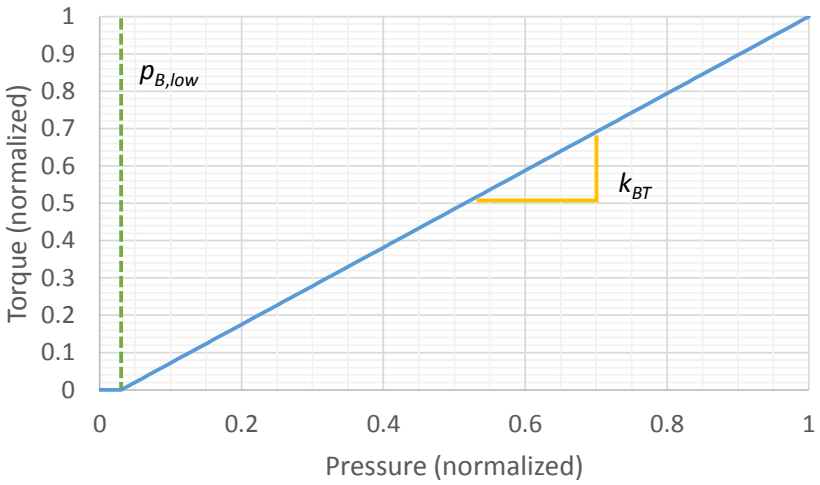

Figure 3.9. Braking torque vs. brake pressure for a single wheel.

Equation (3.28) shows that, assuming a constant cylinder area and lever arm, the braking torque of the system is dependent solely on the braking pressure. Therefore, it is important to model the braking pressure response dynamics with respect to a given command from the EH system controller. The primary component controlling the braking pressure from the TC system standpoint is the pressure reducing valve (see Section 7.1 for more detail). This valve controls the pressure in the braking line directly, proportional to a command from the controller. Due to this, it was decided that modeling the response of this component was sufficient to describe the braking pressure dynamics. An investigation into the response characteristics of the pressure reducing valves installed on the reference wheel loader was conducted, and the result was a second order dynamic model from input command to braking pressure (output).

$$P_B(s) = \frac{\omega_n^2 p_{B,max} U(s)}{s^2 + 2\zeta\omega_n s + \omega_n^2} e^{-\psi s} \tag{3.29}$$

In (3.29), P_B is the Laplace transform of the braking pressure, ω_n and ζ are the natural frequency and damping ratio, respectively, of the second order system, ψ is the time delay of the system, $p_{B,max}$ is the maximum braking pressure, and U is the Laplace transform of the input signal into the system (assumed to be in the range 0 to 100%). For most simulations the time delay ψ of the braking system is assumed to be zero to simplify the controller formulations (the other parameters were adjusted to slow the simulated system down accordingly). However, there were also some investigations done into the effects of and possible solutions for a system incorporating nonzero time delay (see Section 4.6).

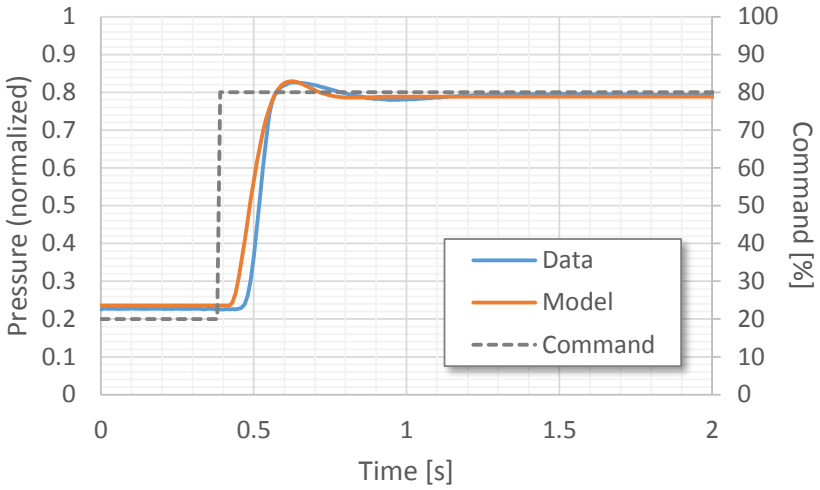

Figure 3.10. Model matching for braking system pressure response.

The braking pressure model for the pressure reducing valve is based on the response of the system as measured from the real-world valves. Indeed, the various parameters defining the model (natural frequency, damping ratio, and time delay) were all determined by recording step responses of the system at various levels and adjusting the parameters until the modeled step response matched the real-world response as closely as possible (Fig. 3.10). Therefore, the braking pressure model is built on a second-order differential equation model which approximates the response of the pressure reducing valve in the EH braking system. For the purposes of this research, this behavior is sufficient to represent the effect of the braking system on the torque response.

3.6 Complete Vehicle System Model

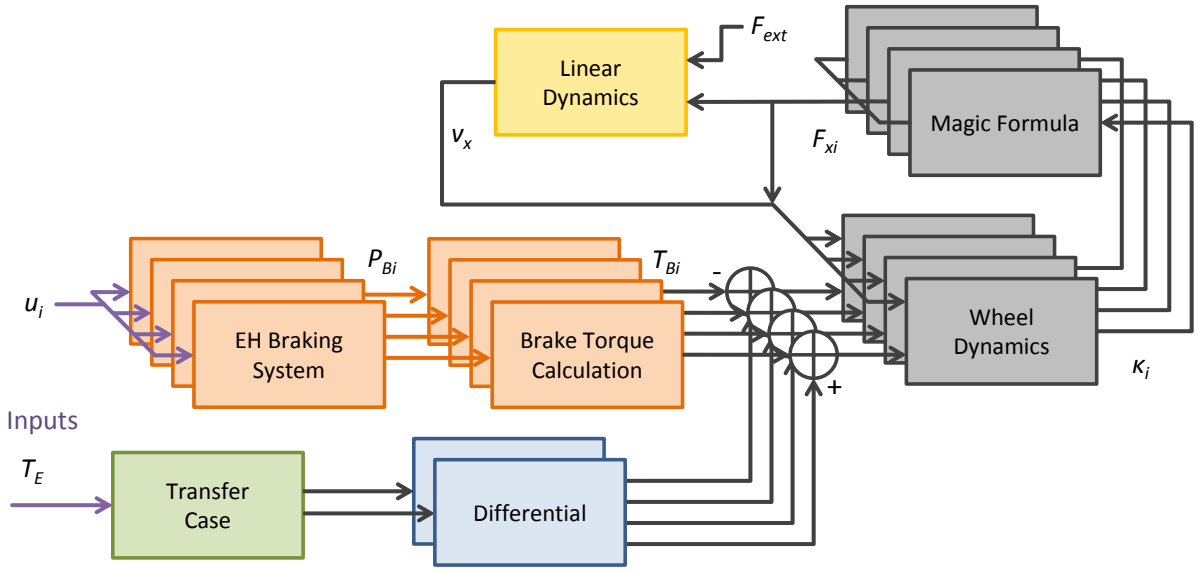

Figure 3.11. Block diagram of full vehicle dynamics model.

Having constructed all of the various components of the vehicle dynamics model, it is possible to simulate the complete vehicle for assessing traction control methods. A diagram showing the connections among the various submodels is shown in Fig. 3.11.

The primary inputs to the system are the engine torque T_E and the commands to the brakes u_i. The brake commands can come either from the operator pressing the brake pedal or from the TC system utilizing the electro-hydraulic capacity of the modified system. The engine torque is distributed to each of the wheels through the transfer case and differentials, then those distributed torque values are summed with the braking torque. The resultant torque is used in the equations of motion for the wheels. Using the wheel slip at each wheel κ_i, the longitudinal force F_{xi} is calculated. From that and other forces, the linear motion of the simulated vehicle is updated. All of these various submodels work together to generate a suitable estimate of the wheel loader behavior under various conditions and inputs.

3.7 Simulation Description and Results

Throughout the course of this work, it was necessary to simulate the various control structures and other ideas to get some idea of how the system would behave in experiments. Therefore, a standard test simulation was developed using MATLAB Simulink for the purpose of comparing different system configurations side by side. This simulation setup is designed to reflect the real-world laboratory testing conditions described in Section 8.2 of this report, and it is pictured in Fig. 3.12.

The simulated wheel loader begins from a certain position with an initial velocity v_{x0} (typically a very low value). At the beginning of the simulation ($t = 0$), the engine torque into the transfer case is set to a certain value T_E. The input torque depends on the purpose of the simulation and the operating conditions. Once the simulated vehicle has traveled a certain distance x_{tire}, it comes into contact with a resistive force which arrests its forward motion. As the wheel loader continues to push against this resistance, its wheels begin to slip against the ground. Simulating this condition allows for implementation of various control strategies, and it allows access to many parameters which would be difficult or impossible to measure on a real-world setup.

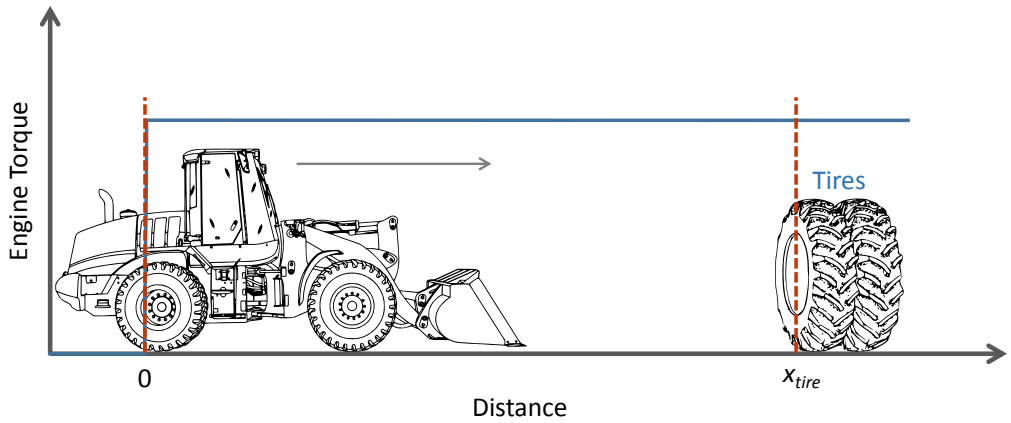

Figure 3.12. Test setup for model simulations.

The resistive force could be modeled any number of ways. Much work has previously been done into modeling the forces generated in digging soil and other materials [45, 46].

However, for the investigative work done here, it was decided that a more useful simulation scheme would be one that better approximates the test setup run at the laboratory (see Chapter 8). Therefore, the resistive force included in the model was instead designed to represent the tires from the lab tests (as shown in Fig. 3.12). The resistive force of the tires was modeled as a spring-damper system which works in only one direction. That is,

$$F_{tire} = \begin{cases} k_{tire}x_p + c_{tire}\dot{x}_p \ , & x_p \geq 0 \text{ and } \dot{x}_p \geq 0 \\ 0 \ , & \text{else} \end{cases} \ , \tag{3.30}$$

where $x_p = x - x_{tire}$ is the distance the wheel loader has pushed into the tires, and k_{tire} and c_{tire} are the spring and damping rates of the tires, respectively.

The results of a system simulation without traction control are shown in Fig. 3.13. In this simulation, the wheel loader has an initial velocity of 2 km/h with zero wheel slip (i.e. the vehicle is coasting with no engine torque or brake torque applied). All four wheels experience the same ground condition (slip-friction characteristic), so their motion is identical. At time $t = 0$, a significant engine torque is applied to the system, and this torque remains constant throughout the simulation. The wheel loader accelerates briefly, with the wheels remaining in good traction with the road surface. At the time marked with the red dashed line, the wheel loader impacts the tires and begins pushing against them. The resultant resistive force causes the machine to slow to a halt (Fig. 3.13(a)), while all four wheels begin to slip significantly against the ground (Fig. 3.13(b)). Because the linear velocity of the system is very low and the wheels' velocity so relatively high, the slip ratios at each wheel achieve values which are quite large (Fig. 3.13(c)). The result of this uncontrolled increase in wheel slip is the reduction in pushing force achieved by the wheel loader (Fig. 3.13(d)). This reduction in force is related to the shape of the slip-friction curve of the model, and it represents the potential of the TC system to increase the machine pushing force when implemented correctly.

At this point, it should be noted that this model does not include additional resistances on the wheels, so they will continue to increase in velocity beyond the real-world system limit if the engine torque remains constant. As the outcome of this work is to control the

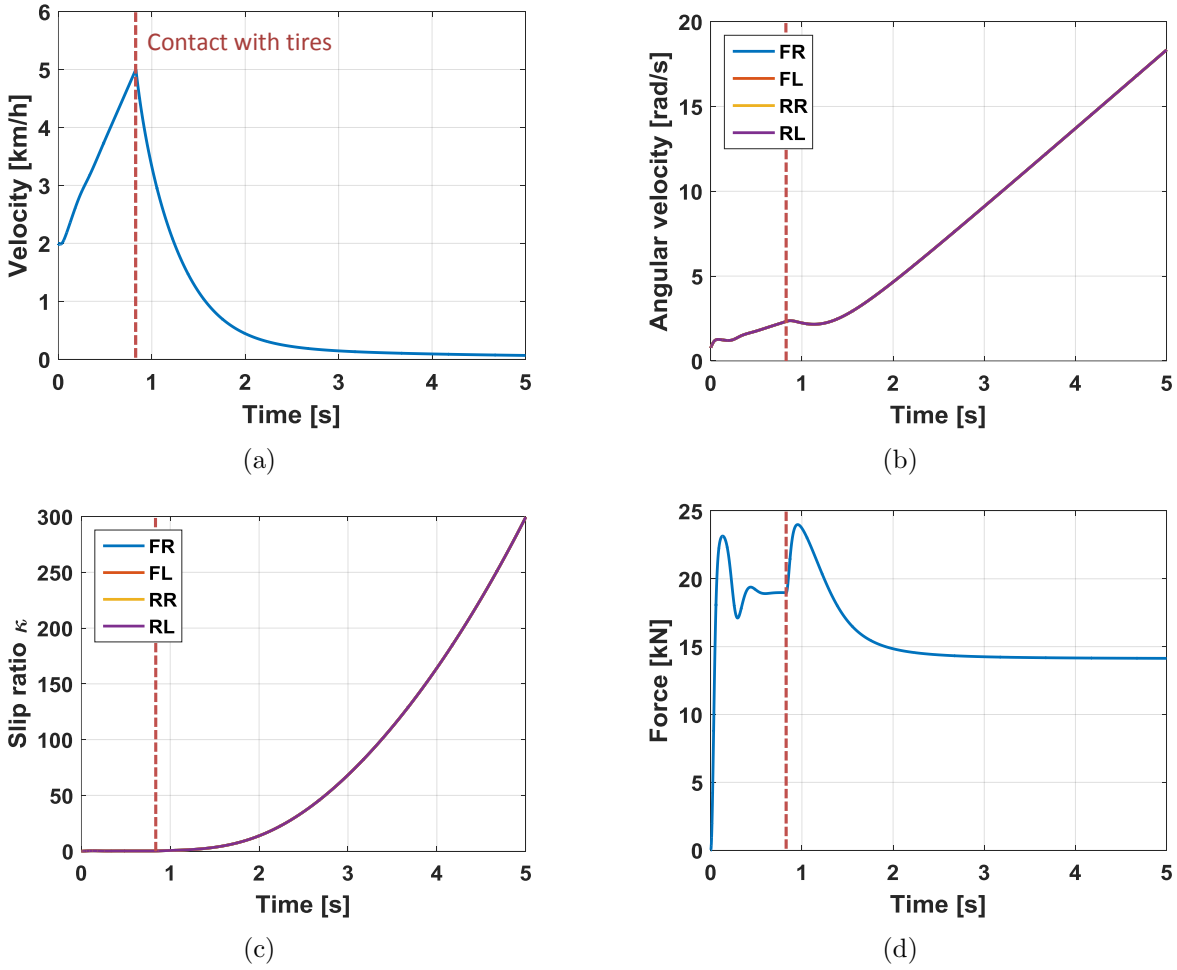

Figure 3.13. Simulation results for system with no traction control implemented, showing (a) vehicle linear velocity, (b) wheel rotational velocities, (c) wheel slip ratios, and (d) total longitudinal force.

wheel speeds, this is not a significant issue. However for other systems, it may be necessary to include considerations such as wheel resistive torque or engine/drivetrain dynamics which can limit the maximum wheel speeds.

This simulation setup is the same that is used throughout this work for comparing the performance of the different control strategies. It generates results which can be compared to the experimental laboratory tests, and it gives a very repeatable, controlled test condition with quickly changing dynamics that can test the ability of the TC system to effectively and accurately control the wheel slip. By using this simulation, multiple different control

schemes can be efficiently tested and compared without all the time required to implement and test every controller on the real-world machine.

3.8 Implement Force Estimation

One of the most important metrics for the performance of a traction control system is pushing force. There are many benefits that come from introducing a traction control system onto a vehicle, including safety and controllability. However, when considering performance, aspects such as acceleration (for on-road vehicles) and work cycle productivity (for construction machines) are of primary concern. Both of these are largely functions of the amount of force generated at the wheels.

For vehicles which are simply driving on the ground surface (termed "roading" in many heavy-duty system applications), it is possible to generate an estimate of the wheel force using a method such as that proposed in Section 8.1 of this work. However, that methodology assumes little to no force is generated when there is no vehicle motion. This assumption is obviously incorrect in situations such as loading cycles, wherein a wheel loader may produce extremely high wheel forces while traveling at little to no velocity.

Therefore, a method is needed which can estimate the force produced by the wheel loader when it is pushing against a barrier or a material pile. Fortunately, the system architecture allows for the development of this force estimate, by taking advantage of the properties of the hydraulic actuators used to manipulate the wheel loader implement. As long as neither hydraulic cylinder is being pressed against its endstop (that is, neither fully extended nor fully retracted), then the hydraulic pressure differences across those cylinders are directly proportional to the forces they are generating. Therefore, by measuring the implement position and hydraulic cylinder pressures, an estimate of the forces acting on the bucket can be obtained. Similar systems have already been developed for estimating only the payload in the machine bucket [47, 48].

In order to calculate the reactive forces across the implement system, the component positions must first be known. The basic system geometry to be solved is shown in Fig. 3.14 (though it should be noted that this drawing is for explanation purposes only and does not

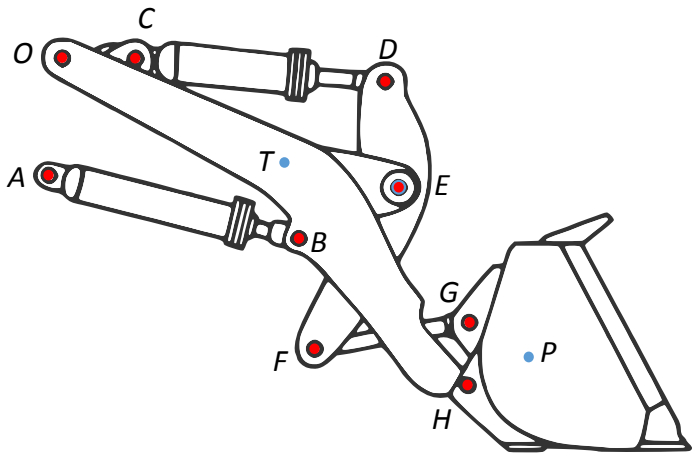

Figure 3.14. Point locations used for implement force calculations.

match the actual geometry of the reference machine). There are six distinct components involved in the system. Two of them are hydraulic cylinders: the boom lift cylinder AB, which actually represents two cylinders acting symmetrically on the real-world system, and the bucket tilt cylinder CD. The boom itself contains pin O, which connects it to the wheel loader chassis, point E, connecting it to the bellcrank, and point H, which connects it to the bucket. The boom center of gravity is located at point T. Link FG transfers the tilting motion from the bellcrank to the bucket. The bucket center of gravity is located at point P, while for this work the bellcrank center of gravity is considered collocated with point E (the distance between the two points is negligible on the reference machine).

As described in Section 7.2.1, the prototype machine has two position sensors: one for boom angle and one for bellcrank angle. The boom angle sensor measures the angle that ray $\overrightarrow{OH}$ makes with respect to the horizontal (that is, the angle is positive if H is above O), while the bellcrank sensor measures angle $\angle OED$. Once these two angles have been obtained, it is relatively straightforward to ascertain the locations of all points in the implement system by constructing constituent triangles and using trigonometric identities and the Law of Cosines.

Point O is used as the origin for this analysis, with positive vertical direction z defined as pointing the opposite direction of the gravitational force and positive longitudinal direction x normal to z and pointing toward the front of the machine. Furthermore, points O, A, and

C are all on the wheel loader chassis, so the locations of A and C with respect to O are fixed (assuming the machine stays on level ground). Because the boom angle sensor gives the angle of ray $\overrightarrow{OH}$ and the length of segment $\overline{OH}$ is known, determining the location of H is very straightforward. And as the boom is a rigid body, points B, E, and T can be found easily through their fixed relationship between points O and H.

The next component whose position can be solved is the bellcrank. The bellcrank angle sensor gives angle $\angle OED$, and the length of segment $\overline{DE}$ is known; therefore, point D can be defined from the known location of E. Again, triangle $\triangle DEF$ has a fixed shape, so point F can be found knowing the locations of D and E. As the lengths of $\overline{FG}$ and $\overline{GH}$ are fixed, and the positions of F and H are known, triangle $\triangle FGH$ can now be solved to find the location of G. Finally, triangle $\triangle GHT$ on the rigid-body bucket is also fixed in shape, so the location of point T can also be solved. This defines the entire implement geometry for all possible positions.

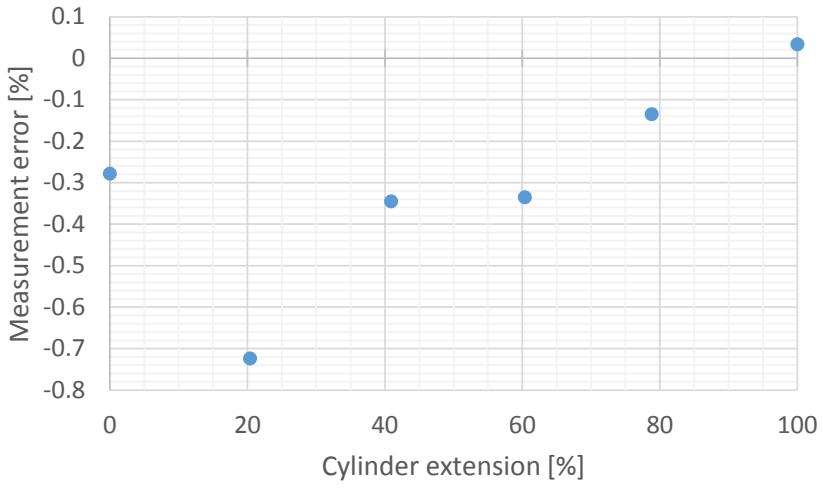

Figure 3.15. Estimation error for bucket tilt cylinder total length.

The geometric model of the system was validated via comparison with real-world measurements of angles and point locations. As shown in Fig. 3.15, the values matched quite well, with estimated position errors of only a few millimeters (less than 1% error). It should also be noted that the reference machine contains some physical endstops which prevent the hydraulic cylinders from fully extending or retracting in certain positions. These addi-

tional endstops were added into the simulation model to prevent it from achieving geometric configurations which are not feasible for the real-world machine.

Having calculated the component positions, a force balance model for the system can now be constructed. As they are relatively light-weight compared to the system forces, and because they have no other external forces acting on them, cylinders AB and CD and link FG are treated as two-force members, meaning that they only transmit equal and opposite forces at each pin connection, which act axially through those points [49]. That leaves only the boom, bellcrank, and bucket as components which require free-body force analysis.

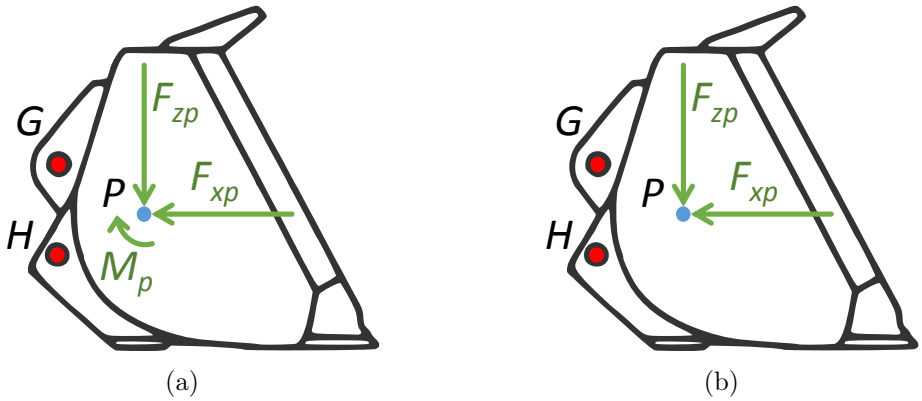

Figure 3.16. Estimated forces acting on wheel loader bucket (a) with a correction moment and (b) without a correction moment.

The force balance in question is, roughly speaking, the mirror image of the analysis conducted in Section 3.3.3. In that section, the machine response was calculated from known forces at the bucket. Inversely, in the current discussion, the forces acting on the bucket are being estimated from the machine reactions. Ideally, the outputs of this model would be the inputs of the dynamic simulation model, with a horizontal resistive force F_{xp}, a vertical resistive load F_{zp}, and a correction moment M_p (as shown in Fig. 3.16(a)). This would allow for complete definition of an arbitrary set of forces acting on the implement. Using only the measurements of hydraulic cylinder pressures, this system is ill-defined and not solvable. However, by removing the correction moment and simply solving for F_{xp} and F_{zp} as in Fig. 3.16(b), the system of equations can now be solved. This simplification means that the new

system has no ability to correct for placement of the forces. It imposes the assumption that all forces are acting on the bucket center of gravity P. This assumption is typically not perfectly valid; however, in many cases, the lines of action of forces on the bucket run close enough to P that the error is small.

What results by applying the unknown forces in Fig. 3.16(b) to the system in Fig. 3.14 is a coupled free-body system for which all unknowns can be calculated. For this work, this system is treated as a static force balance, as the motions involved are quite slow. Again, there are some instances where this assumption is unreasonable. However, for most operating conditions (and especially in the tests in which this force estimate were used), the implement is stationary or nearly stationary during data generation.

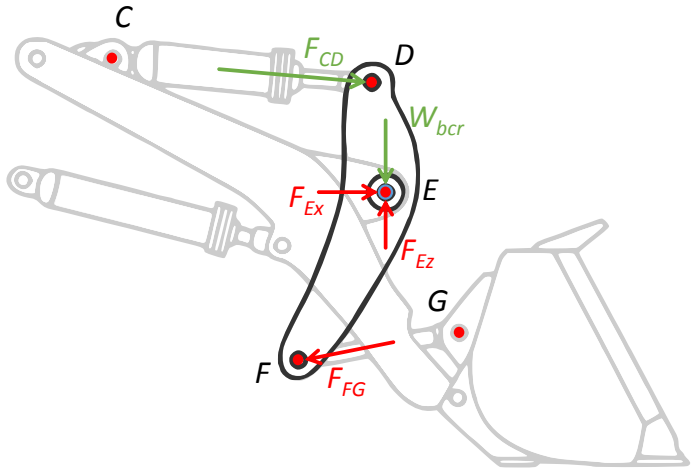

Figure 3.17. Free-body diagram for bellcrank force balance.

Starting from the known hydraulic cylinder forces at points B and D, the coupled force system can be solved in almost every configuration using matrix inversion. (Certain specific implement positions do result in a singular matrix.) The first step to determining the other forces in the implement system is to solve the force balance on the bellcrank (Fig. 3.17). In this discussion, known forces are represented in green while unknowns are represented in red. Starting with the known cylinder force F_{CD}, which is calculated from the differential pressure across cylinder CD and acts along its axis, and the known weight of the bellcrank W_{bcr}, the reaction force components at the boom pin F_E (F_{Ex} and F_{Ez}) and the force at

the bucket link F_{FG} (assumed to be acting along segment $\overline{FG}$) can be calculated using the static force and moment balance equations for the bellcrank.

$$\sum F_{xi} = F_{CD}\cos(\lambda_{CD}) + F_{Ex} - F_{FG}\cos(\lambda_{FG}) = 0 \tag{3.31}$$

$$\sum F_{zi} = F_{CD}\sin(\lambda_{CD}) + F_{Ez} - F_{FG}\sin(\lambda_{FG}) - W_{bcr} = 0 \tag{3.32}$$

$$\begin{aligned}\sum M_{yE} =& F_{CD}\left[\sin(\lambda_{CD})(x_D - x_E) - \cos(\lambda_{CD})(z_D - z_E)\right] \\ &- F_{FG}\left[\sin(\lambda_{FG})(x_F - x_E) - \cos(\lambda_{FG})(z_F - z_E)\right] = 0\end{aligned} \tag{3.33}$$

In Eqs. (3.31) through (3.33), λ_{ij} represents the angle created by the ray drawn from point i to point j (for instance, λ_{CD} is the angle that ray $\overrightarrow{CD}$ makes with respect to the x-axis). Also, x_i and z_i are the horizontal and vertical locations of point i. The moment balance for the bellcrank, Eq. (3.33), is taken about point E, simplifying it by eliminating force F_E and the bellcrank weight W_{bcr}, which is assume to be acting through that point. It should be noted that in this notation, it is assumed that F_{FG} acts in the opposite direction of the positive sign convention. Therefore, its signs are reversed from the other terms.

This construction forms a system of linear equations, which can be represented as the matrix equation:

$$Kv = b \ , \tag{3.34}$$

where v is the vector of unknown values K is the matrix of equation coefficients, and b is the vector of known results for each system. For the bellcrank force balance, this forms a 3×3 system, with the following components.

$$\begin{aligned} v_{bcr} &= \begin{bmatrix} F_{Ex} \\ F_{Ez} \\ F_{FG} \end{bmatrix}, \ b_{bcr} = \begin{bmatrix} -F_{CD}\cos(\lambda_{CD}) \\ -F_{CD}\sin(\lambda_{CD}) + W_{bcr} \\ -F_{CD}\left[\sin(\lambda_{CD})(x_D - x_E) - \cos(\lambda_{CD})(z_D - z_E)\right] \end{bmatrix} \\ K_{bcr} &= \begin{bmatrix} 1 & 0 & -\cos(\lambda_{FG}) \\ 0 & 1 & -\sin(\lambda_{FG}) \\ 0 & 0 & -\left[\sin(\lambda_{FG})(x_F - x_E) - \cos(\lambda_{FG})(z_F - z_E)\right] \end{bmatrix} \end{aligned} \tag{3.35}$$

As long as K is nonsingular (i.e. $\tan(\lambda_{FG}) \neq \frac{z_F - z_E}{x_F - x_E} = \tan(\lambda_{EF})$), this system can be solved by matrix inversion. That is,

$$v = K^{-1} b \ . \tag{3.36}$$

This operation results in the forces F_{Ex}, F_{Ez}, and F_{FG}, which can now be used for solving the rest of the system.

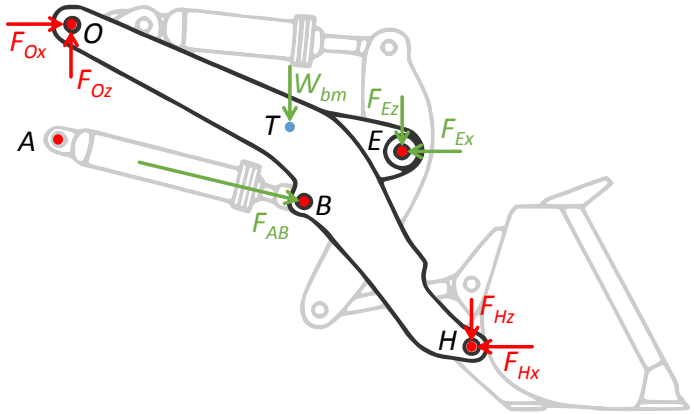

Figure 3.18. Free-body diagram for boom force balance.

The same solution method can now be applied for the boom. Its free-body diagram is constructed as in Fig. 3.18. It can be seen that, whereas F_{Ex} and F_{Ez} were treated as unknown values for the bellcrank force balance, they are now considered known. The same force and moment balances are done again, with the moment balance being taken about point O.

$$\sum F_{xi} = F_{AB}\cos(\lambda_{AB}) + F_{Ox} - F_{Ex} - F_{Hx} = 0 \tag{3.37}$$

$$\sum F_{zi} = F_{AB}\sin(\lambda_{AB}) + F_{Oz} - F_{Ez} - F_{Hz} - W_{bm} = 0 \tag{3.38}$$

$$\begin{aligned} \sum M_{yO} =& F_{AB}\left[\sin(\lambda_{AB})(x_B - x_O) - \cos(\lambda_{AB})(z_B - z_O)\right] \\ &+ F_{Ex}(z_E - z_O) - F_{Ez}(x_E - x_O) + F_{Hx}(z_H - z_O) \\ &- F_{Hz}(x_H - x_O) - W_{bm}(x_T - x_O) = 0 \end{aligned} \tag{3.39}$$

The variable conventions used for Eqs. (3.37) through (3.39) are the same as in Eqs. (3.31) through (3.33). Before constructing the matrix equations as with the bellcrank force balance, it should be noted that this system, which has only three equations, contains four unknown values: F_{Ox}, F_{Oz}, F_{Hx}, and F_{Hz}. This poses a problem, as the system is therefore not solvable.

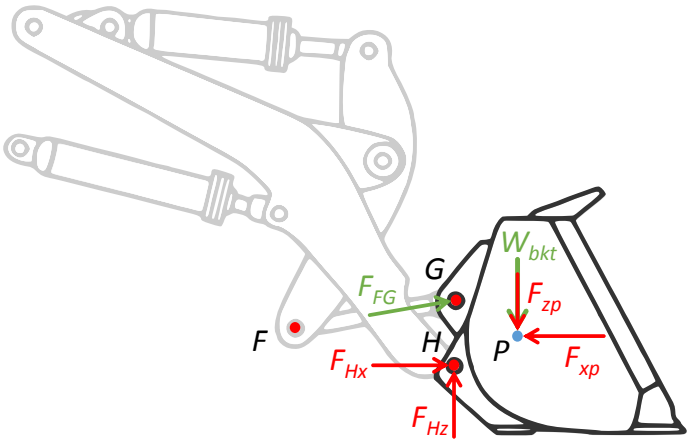

Figure 3.19. Free-body diagram for bucket force balance.

Instead, the bucket force balance must also be taken into consideration. Its free-body diagram can be seen in Fig. 3.19. Again, the link force F_{FG}, which was an unknown in the bellcrank balance, has become a known value here. Constructing the balance equations with the moment balance done about point H gives the following result.

$$\sum F_{xi} = F_{FG}\cos(\lambda_{FG}) + F_{Hx} - F_{xp} = 0 \tag{3.40}$$

$$\sum F_{zi} = F_{FG}\sin(\lambda_{FG}) + F_{Hz} - F_{zp} - W_{bkt} = 0 \tag{3.41}$$

$$\begin{aligned}\sum M_{yH} =& F_{FG}\left[\sin(\lambda_{FG})(x_G - x_H) - \cos(\lambda_{FG})(z_G - z_H)\right] \\ &+ F_{xp}(z_P - z_H) - F_{zp}(x_P - x_H) - W_{bkt}(x_P - x_H) = 0\end{aligned} \tag{3.42}$$

This system also has four unknowns (F_{Hx}, F_{Hz}, F_{xp}, and F_{zp}), rendering it unsolvable. However, noting that two of these unknowns (F_{Hx} and F_{Hz}) are common to both the boom and the bucket subsystems, it can be seen that these two systems contain six equations

for six unknowns. Therefore, the two systems can be combined into a single 6×6 matrix equation, with the following values:

$$v_{bb} = \begin{bmatrix} F_{Ox} \\ F_{Oz} \\ F_{Hx} \\ F_{Hz} \\ F_{xp} \\ F_{zp} \end{bmatrix}, \; b_{bb} = \begin{bmatrix} F_{Ex} - F_{AB}\cos(\lambda_{AB}) \\ F_{Ez} - F_{AB}\sin(\lambda_{AB}) + W_{bm} \\ f_1 \\ -F_{FG}\cos(\lambda_{FG}) \\ -F_{FG}\sin(\lambda_{FG}) + W_{bkt} \\ f_2 \end{bmatrix}$$

$$K_{bb} = \begin{bmatrix} 1 & 0 & -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & -1 & 0 & 0 \\ 0 & 0 & (z_H - z_O) & -(x_H - x_O) & 0 & 0 \\ 0 & 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & (z_P - z_H) & -(x_P - x_H) \end{bmatrix}, \tag{3.43}$$

where:

$$\begin{aligned} f_1 = &- F_{AB}\left[\sin(\lambda_{AB})(x_B - x_O) - \cos(\lambda_{AB})(z_B - z_O)\right] \\ &- F_{Ex}(z_E - z_O) + F_{Ez}(x_E - x_O) + W_{bm}(x_T - x_O) \end{aligned} \tag{3.44}$$

$$\begin{aligned} f_2 = &- F_{FG}\left[\sin(\lambda_{FG})(x_G - x_H) - \cos(\lambda_{FG})(z_G - z_H)\right] \\ &+ W_{bkt}(x_P - x_H)\ . \end{aligned} \tag{3.45}$$

With a nonsingular K_{bb}, this system can be solved system using (3.36).

The system solution contains an estimate of both horizontal and vertical forces being applied to the wheel loader implement. This is a significant result, as pushing force (horizontal) and payload weight (vertical) are important performance metrics for these machines in everyday operation. The research team did not have a good method for validating the horizontal pushing force results. However, it was possible to verify the values for payload weight. Therefore, a multi-point comparison was run to asses the capability of the estimation model to determine payload weight. To give true value of payload weight, a truck scale

was used to measure the wheel loader weight with an empty bucket and again with three payloads of different size. By comparing the force estimate from this section with the true value, a percentage error was calculated. Furthermore, the same payload was weighed using the on-board payload weighing system mentioned in Section 7.2.1. The results of these tests are shown in Table 3.1, with the heading "State of the Art" referring to the existing on-board payload weighing system and "Force Estimate" referring to the methodology set forth in this section.

Table 3.1. Comparison of payload estimate to state of the art.

Load (% of full capacity)	**State of the Art (% error)**	**Force Estimate (% error)**
30%	1.34%	2.19%
65%	−1.04%	1.70%
100%	−1.20%	1.31%

The results were very positive, with the force estimate developed here achieving estimates within 2–3% of the payload scale measurements, which is comparable to the state-of-the-art system. In fact, the existing system has a few limitations in its operation. The most significant drawback is that the payload measurement only occurs when the boom is raised to a certain height. This means that if the edge of the loading truck or container is not already high enough off the ground, the operator must pay careful attention to raise the bucket to a sufficient height before dumping or risk losing payload data. Instead, the force estimate from the current work can be taken with a static implement position, and is not dependent on a particular work cycle. Furthermore, the state-of-the-art system requires data about engine speed, which is not needed for the force estimate developed here.

Not only does the vertical force estimate match the expected value quite well, the estimate of horizontal force for these tests was relatively quite low. This means that the model is capable of distinguishing properly between horizontal and vertical forces. The same trend exists in the lab tests, wherein the vertical force remains near zero while the pushing force estimate mimics the observed behavior of the machine. Therefore, the force estimation

methodology is giving reasonable results for both vertical and horizontal forces. Once again, the assumption that the applied forces are acting through the bucket center of gravity is not viable for all operating conditions. But for the test cases examined in this work, this assumption provides good results. The horizontal pushing force estimate was used extensively for data generation in the lab tests described in Sections 8.2 through 8.4.

4. CONTROL CONSIDERATIONS

This chapter focuses on the proper development of a control system designed to eliminate excessive tire slip in vehicle systems. The discussion begins with a relatively simplistic controller built using a proportional-integral-derivative (PID) construct. Some logic considerations are added to this basic structure to generate a more robust Augmented PID controller, which is capable of functioning well on a real-world machine. After that, some further control theories are described in a more theoretical sense.

At this juncture, it is worth noting that the Augmented PID controller is the control structure implemented on the reference machine during the tests outlined in Chapters 8 and 9. The more advanced controllers were developed with the understanding that their implementation on a real-world machine was beyond the scope of the present project. Some of these more advanced ideas are present in this chapter; others are expounded upon in Chapters 5 and 6.

4.1 A Simple Traction Control Methodology

Developing a strategy for a basic control algorithm for the wheel loader traction is not a particularly difficult task, as the system itself has a straight-forward relationship between inputs and outputs. The basic structure of the traction control system is shown in Fig. 4.1. Based on the difference between some setpoint value $r(t)$ (typically the optimal slip ratio κ^* for this work) and the current value of that parameter $y(t)$, an error signal $e(t)$ is given to a controller which generates the command signal vector $u(t)$ for the system. In this case, the system plant is represented as two separate blocks, the braking system, which outputs a braking torque T_{Bi} for each wheel, and the vehicle dynamics block, which determines the vehicle motion based on the braking torque and the engine torque (which is treated as an uncontrolled input to the system). The dynamic expressions which describe these system models are developed in Chapter 3 of this work.

As shown in Fig. 4.1, the control methodology being used to improve vehicle traction is a simple PID controller. This controller was selected for use in this project because it is

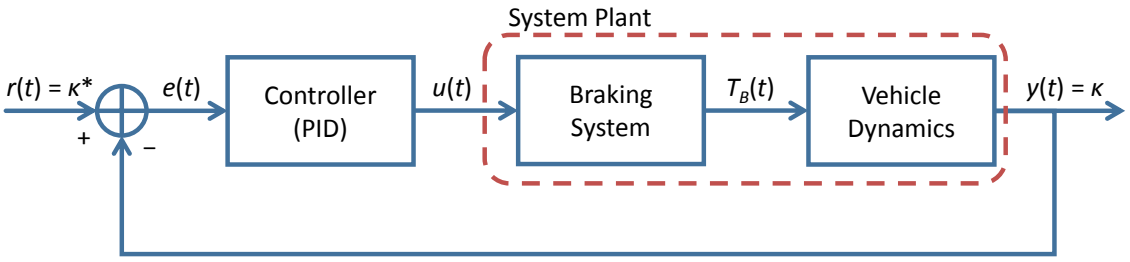

Figure 4.1. Basic system control structure incorporating a PID controller.

relatively simple and well-understood, and because it is applicable to a wide range of system types. As shown in Fig. 4.2, the PID controller works by using the error signal, its derivative, and its integral, scaling those values, and summing the result together. The resultant control law is as follows:

$$u(t) = K_P e(t) + K_I \int e(t)dt + K_D \frac{de(t)}{dt} \ . \tag{4.1}$$

The proportional component of the controller acts to drive the current value of the error toward zero, while the derivative component aids in improving the dynamic response of the controller, and the integral contribution works to drive the steady-state system error toward zero [50]. By setting the proportional, integral, and derivative gains (K_P, K_I, and K_D, respectively) correctly, the system response can be tuned to a better performance. For most systems, the closed-loop response with a PID controller is not the best possible system response, but it is a good starting point which is relatively reliable.

Care should be taken when setting the values of the controller gains, as performance and stability can both be significantly impacted by these parameters. If they are set too low, the controller will have trouble tracking the reference trajectory. However, if they are too high, many systems can become unstable. This restriction is further complicated if the system plant contains elements such as nonlinearities, time delays, and so on (which the vehicle system at hand does contain). Therefore, much time was spent adjusting and fine-tuning these controller parameters both in simulation and in real-world testing to ensure the traction control was as effective as possible.

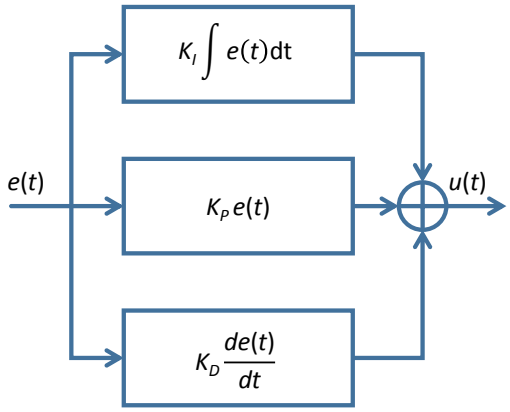

Figure 4.2. PID control operating principle.

4.1.1 Simulation Results

Once the simple PID controller had been formulated, it was implemented in simulation. Figure 4.3 shows the results of this simulation. The simulation setup is the same as that described in Section 3.7. The vehicle accelerates toward the tires and impacts them just before time $t = 1$ second. As the resistive force slows the wheel loader to a stop, all four wheels begin to slip (Fig. 4.3(a)). Due to the fact that all four wheels are on the same ground condition (i.e. they have the same slip-friction curve), their behavior is identical.

As the wheels experience macro-slip, the pushing force of the system begins to decrease (Fig. 4.3(b)). The simple PID controller recognizes the slip has increased beyond the optimal value κ^*, and it begins braking the wheels to slow them down. The behavior of the PID is such that at first it brakes the system too harshly (overshoot), causing it to actually experience a negative force momentarily (though the operator likely would not notice this much, as the system is barely moving). Around two seconds after the initial impact, however, the wheel slip has converged roughly to the setpoint, and the force has been increased close to its maximum value.

This basic simulation shows already the potential of the traction control system. With TC implemented, wheel slip is greatly reduced and tractive force is maximized. Of course, the performance of the controller (overshoot, convergence time, etc.) can be improved, but this result is still a significant improvement over the case with no traction control. Different

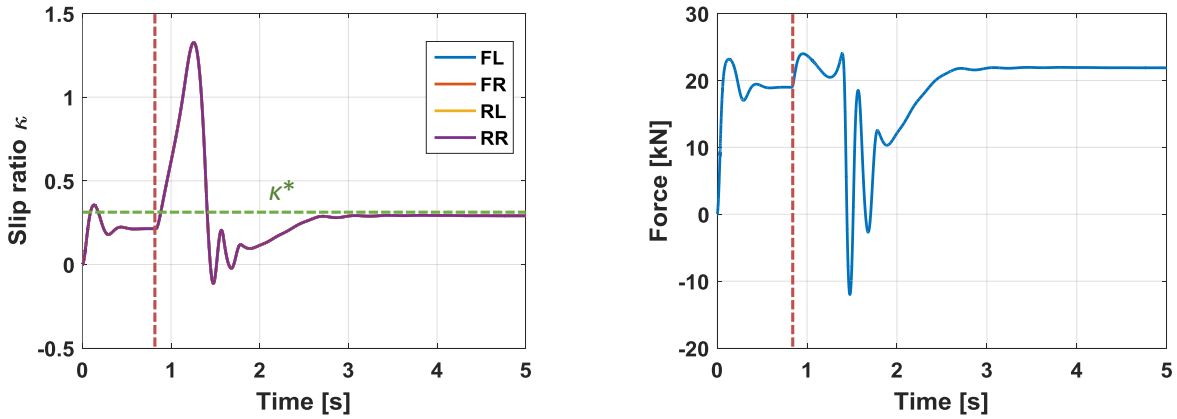

(a) (b)

Figure 4.3. Simulation results for traction control using a simple PID controller, showing (a) wheel slip ratio and (b) total longitudinal force.

gain parameters for the PID controller will allow for different system responses, but there are always tradeoffs. For instance, decreasing the gains will decrease the overshoot of the system, but it will require more time to settle to the final value. All of this should be taken into account when designing a TC system using this simple PID structure.

4.2 The Augmented PID Structure

One of the more difficult aspects of designing a traction control system is the fact that it is a single-sided control system. That is to say, the actuators being used to control wheel slip (the brakes) are only capable of slowing down the wheels. They cannot add more torque into the system to increase wheel speeds. Furthermore, the error signal for this system has the unique property that it is bounded from below. That is to say, for a driven system, the slip of the system is always positive or zero. Therefore, subtracting a setpoint, the resulting error signal can theoretically have any positive value, but it is limited to only a small range of negative values. This means that there is a limit to the rate at which the integral component of the PID control signal can be decreased, while there is no limit to the rate at which it can increase. What results is a phenomenon called *integral windup*, wherein the behavior of

the integral term of the PID controller is impacted by these physical limitations on the error signal.

Most controller algorithms are developed under the assumption that the controlled actuator is has a continuous influence on the system and that the actuators are double-sided (i.e. they can influence the system by both increasing and decreasing certain parameters). The traction control system, on the other hand does not satisfy these assumptions. Therefore, care should be taken in designing controllers for this system, as the single-sided nature of the brakes can cause serious issues if it is not given careful consideration.

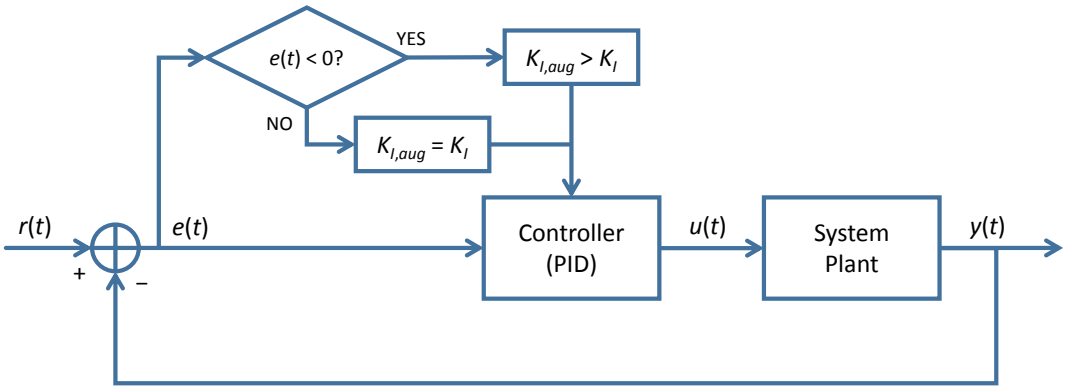

Figure 4.4. Augmented PID structure.

To this end, a logic structure was added to the simple PID structure used in early versions of this work (Fig. 4.4). Instead of using a constant value for the integral gain of the PID controller, it was given two different values, determined by the sign of the error signal. Typically, the controller integral gain is scaled by some factor when the slip error signal is negative. While this is not a perfect solution (the falling rate of the signal still has an upper bound), it improved the controller performance by allowing the control signal to fall more rapidly when braking is unnecessary and thus alleviating the effects of integral windup.

4.2.1 Simulation Results

In order to assess the benefits of using the Augmented PID controller, it was also implemented in simulation. The results of one such simulation is shown in Fig. 4.5. This

simulation contains all of the events shown in the previous simulations: as the wheel loader impacts the tires, the wheels slip and traction force is lost, but it is regained once the TC system activates.

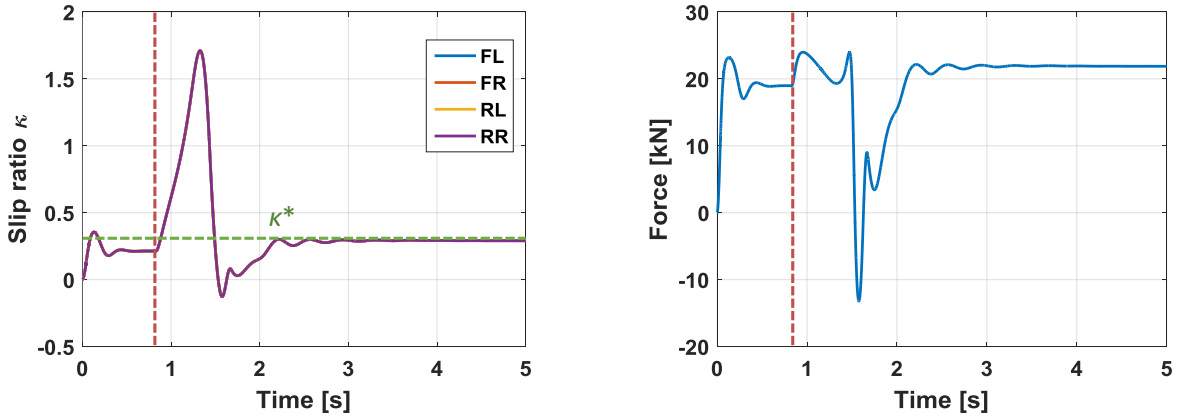

(a) (b)

Figure 4.5. Simulation results for traction control using the Augmented PID controller, showing (a) wheel slip ratio and (b) total longitudinal force.

In this case, the advantage of the Augmented PID is clear. With the additional integral gain during moments of negative slip error, the system response is significantly faster. The slip and force signals converge to their final values roughly half a second (25%) faster than in the case of the simple PID controller. Furthermore, the oscillating overshoot has been reduced significantly, and the controlled system response is much smoother. All of these together show a much improved performance for TC via Augmented PID compared to the standard PID case.

4.3 Effects of Ground Conditions and Braking on Wheel Resistive Torque

There are several different factors which contribute to the philosophy behind traction control system architectures. Of course, the basic idea is to prevent macro-slip of the tires and maximize the pushing force of the machine. However, there are multiple aspects to consider in moving toward these goals. Based on the vehicle's drivetrain, braking system,

steering capability, and other factors, the optimal traction control configuration can vary from machine to machine.

For the wheel loader used as a reference machine in this work, the transmission system comprises a differential gear at each axle and a locked transfer case connecting the driveshaft ultimately to the engine (see Section 3.4 for more information). This has some important implications for the system as opposed to a setup which includes a central differential. The locked transfer case imposes the restriction that the front and rear axles must have the same average velocity, while it allows the torque distributed to each axle to vary. In order to take advantage of this, it is important to consider possibilities with respect to engine management and torque control.

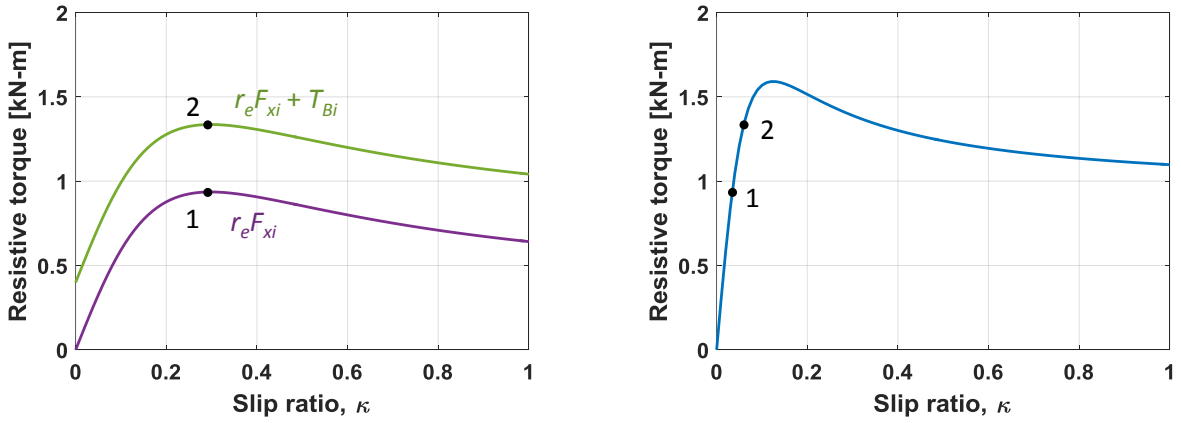

Figure 4.6. (a) Resistive torque seen by the differential from a given wheel and (b) effect on max resistive torque (i.e. pushing force) at other wheel on same axle.

Due to the behavior of the differentials at each axle, traction control via braking can be used to increase the potential tractive force of the system above that allowed by simply reducing the engine torque to the system. This can be explained in part by the process shown in Fig. 4.6. It is known that the differential gears at each axle split the torque equally to each wheel (Section 3.4.1). Therefore, the torque to each wheel on a given axle is limited by whichever wheel demands less torque. In a driven wheel, the torque demand is dependent on the pushing force generated by that wheel F_{xi}. This force creates a *resistive torque* ($r_e F_{xi}$)

when viewed from the differential gear. As the torque must be split evenly between wheels, the torque into the other wheel on the same axle is therefore limited by the resistive torque at the first wheel. For this discussion, it is assumed that the TC system is capable of controlling wheel slip so that the longitudinal force achieves its peak value. Under this assumption, for the unbraked wheel (purple line) shown in Fig. 4.6(a), the opposite wheel on the same axle will see a maximum input torque only as high as the maximum resistive torque value of the first wheel (point 1).

By incorporating a proper traction control strategy, however, the tractive force of the second wheel can actually be increased. If the brakes on the first wheel were to be pressurized, the resistive torque seen by the differential gear will increase, due to the added braking torque T_{Bi}. This increased resistive torque value ($r_e F_{xi} + T_{Bi}$) at one side of the differential causes the torque input to the other wheel to be increased as well, assuming more input torque is available from the driveshaft. This allows the maximum torque (and hence the maximum possible tractive force) at the wheel with the higher friction condition to be increased (from point 1 to point 2 in Fig. 4.6(b)).

The primary takeaway of this is that, in situations where wheels on the same axle are not both on the same ground condition, it is actually possible to increase the net traction force of a machine by a proper traction control using individually brakes wheels.

4.4 Torque Control

The discussion thus far has neglected what is happening at the other axle. However, it is important to control slip on both axles simultaneously, and as the axles are mechanically linked, the conditions at one axle will affect the other axle.

4.4.1 Governing Theory

Consider the example shown in Fig. 4.7. Say that the front right wheel and both rear wheels are on ground which is slick relative to the front left wheel (Fig. 4.7(a)). For the system at hand, the locked transfer case imposes the restriction that the front and rear axles

must have the same average velocity. Therefore, if the input torque is high enough to cause wheels on both axles to slip, further considerations should be made.

For this example, it is possible to imagine an input torque to the system which is large enough to cause the front right wheel and both rear wheels to slip significantly (point 1 in Fig. 4.7). With a TC system based on independent wheel braking alone, the only option is to brake all three wheels to bring them out of macro-slip. However, in this case, both wheels on one axle are being braked. At this point, the braking on that axle is not providing a benefit in terms of pushing force; it is simply increasing the resistive torque at the axle. Therefore, at least some amount of torque being distributed to that axle is now being lost as heat in the brakes. While this setup allows for traction force to be optimized at all four wheels, the overall system is not operating at the most efficient condition.

Therefore, it makes sense to decrease the total input torque into the system enough to bring these three wheels out of macro-slip (i.e. moving from point 1 to point 2). At this point, braking is no longer needed at any wheel to maintain traction with the ground. However, there is still some potential to be gained from the system in terms of traction force. By increasing the input torque marginally and braking the front right wheel, the front axle behaves exactly as the system considered in Section 4.3 above. This brings the system to the maximum possible tractive force condition at point 3.

It is therefore very likely that the most efficient type of traction control system for vehicles with a similar mechanical drivetrain will also incorporate both an independent braking strategy and an engine control strategy, which limits the engine torque alongside taking action to brake wheels which are prone to slipping due to a low-friction ground condition. In the ideal structure, a single wheel slipping on either axle would be controlled by actuating the brake at that wheel (up to one wheel on each axle). When both wheels on one axle begin to slip (implying that at least one wheel on the other axle is slipping), however, the engine control should take effect and decrease the input torque to the system until the only one wheel per axle needs to be braked. In theory, this should maximize the tractive force for a given ground condition, while ensuring that the engine torque is being used as efficiently as possible.

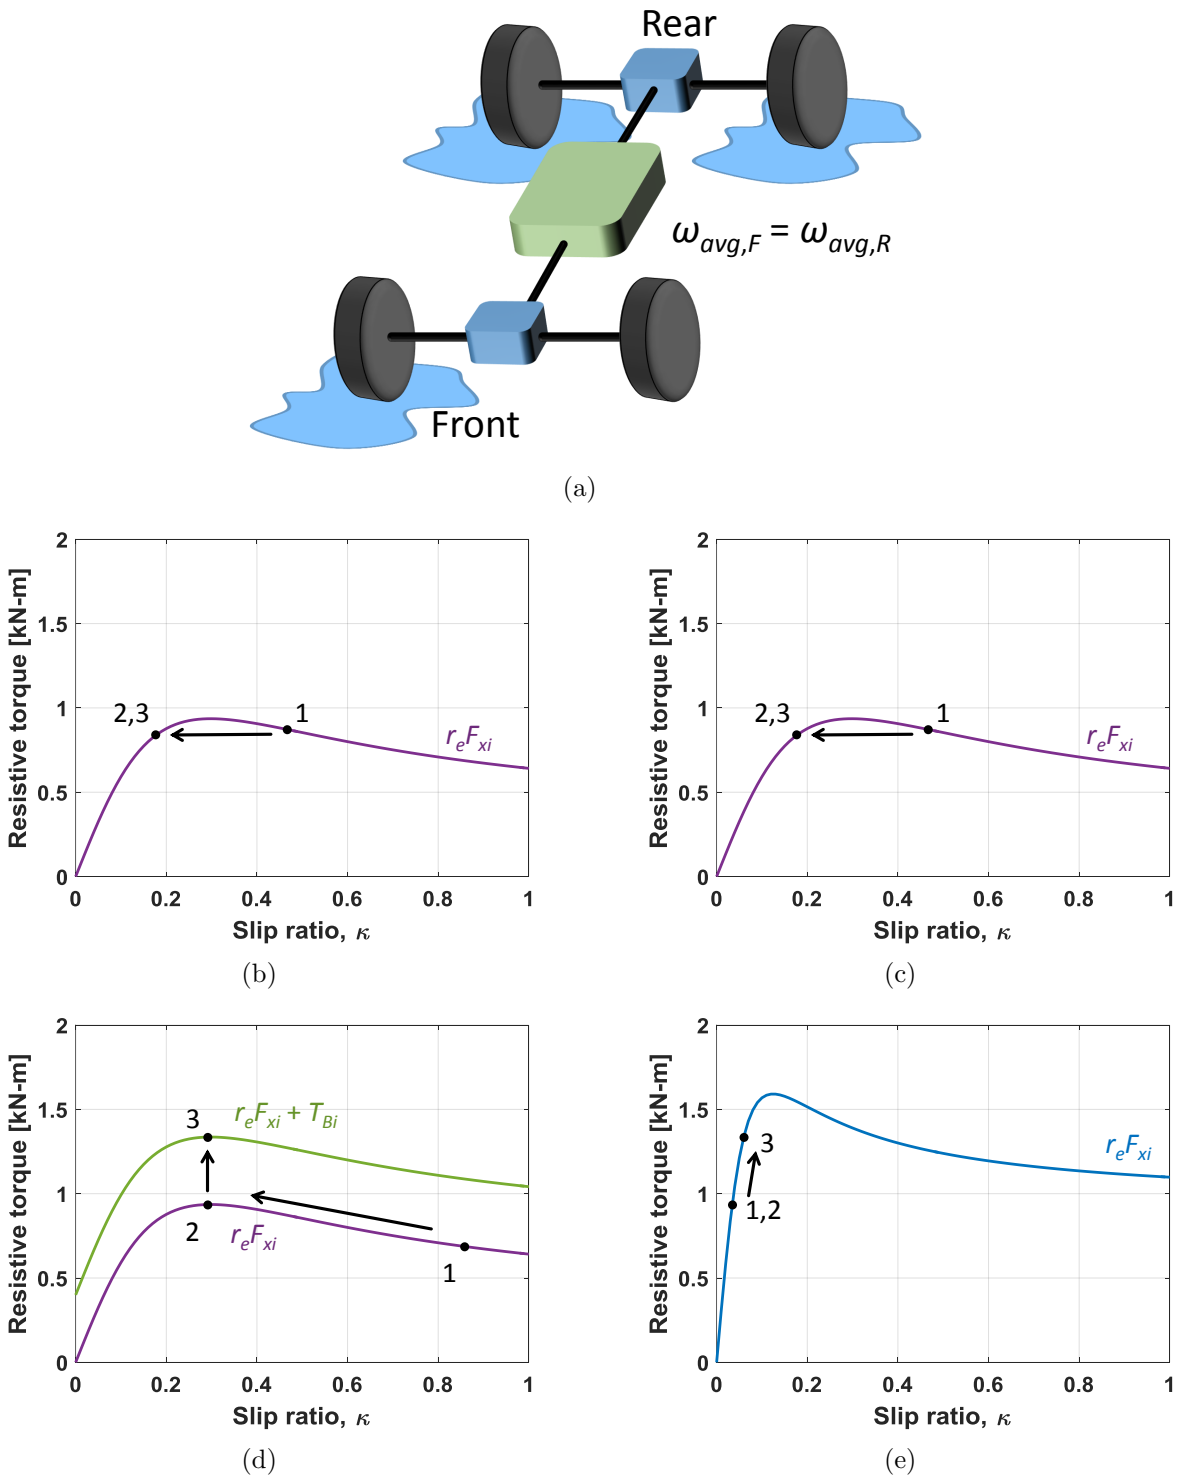

Figure 4.7. Example using combined torque control and braking for traction control, showing (a) the machine setup and the resistive torques at (b) rear right wheel, (c) rear left wheel, (d) front right wheel, and (e) front left wheel.

4.4.2 Simulation Results

This sort of torque-limiting traction control structure was successfully constructed in simulation. One very simple representative case which was simulated is a situation where both front wheels are on a relatively slick surface compared to the rear wheels (Fig. 4.8). In this case, when the simulated system reaches a resistive force and begins to push against it (at around the 1-second mark on the plot), both front wheels have a tendency to slip. Due to the constraints of the vehicle transmission, this causes all four wheels to slip. With the standard TC design (only using brakes to decrease wheel slip), the previously optimal braking strategy would be to actuate both front brakes, causing some engine torque to be lost to heat.

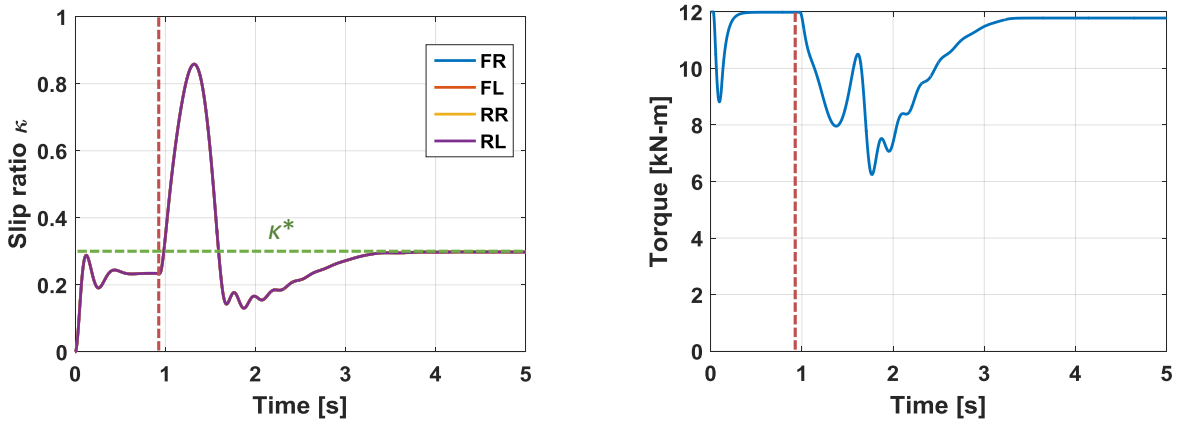

(a) (b)

Figure 4.8. Simulation results for traction control via torque reduction, showing (a) wheel slip ratio and (b) engine torque.

Instead, with an engine torque reduction strategy implemented, the brakes do not activate at all. As the wheel slip begins to increase beyond the optimal value of $\kappa^* = 0.3$ (in this case, the setpoint of the PID controller), the engine torque begins to decrease, eventually slowing the wheels down and preventing excessive slip. Within three seconds, the wheel slip has converged, with the engine torque achieving a lower steady-state value than its initial value (12 kN·m). The result is well-behaved and avoids the negative aspects of same-axle braking.

As the preliminary simulation results showed significant potential for TC via torque reduction, it was eventually incorporated into the prototype TC strategy. More information about the implementation can be found in Section 7.4. The final version of the TC system used in the Productivity Tests highlighted in Section 9.4 included this torque reduction approach.

4.5 Steering Compensation

The traction control system must be able to fully function given any situation in which tire slip can occur. Also, it must be able to assist the operator in any driving maneuvers that the operator may take. Therefore, any TC system on a four-wheel machine such as a wheel loader needs to account for possible steering. This section describes the steering compensation for a vehicle which incorporates articulated steering (such as the reference wheel loader).

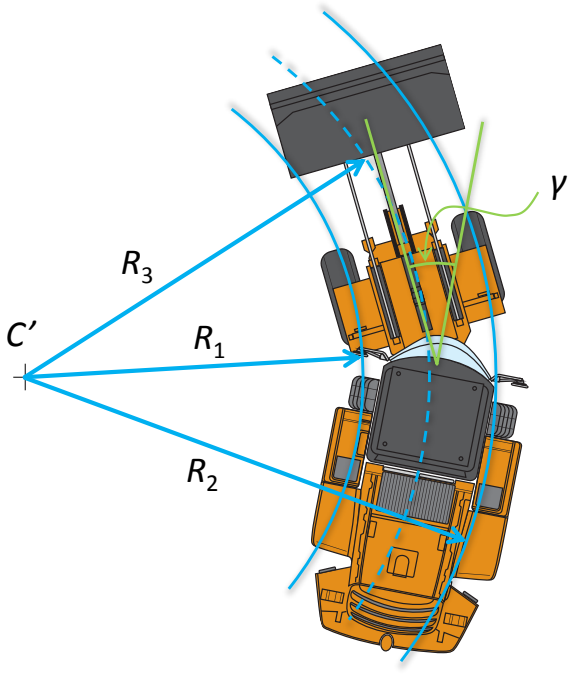

Figure 4.9. Wheel loader steered at angle γ, showing pertinent radii.

To account for steering angle of the wheel loader and its effects on the velocity of the inner and outer radius wheel speeds, a kinematic analysis of the instant center of rotation was implemented. This analysis is needed to determine the varying values of the steering radii. As seen in Fig. 4.9, when the wheel loader steers, an instantaneous center of rotation C' is created. From this center, 3 radii are also formed with respect to different points on the wheel loader. Radius R_1 is formed by the distance from C' to the inner wheel loader tires. Similarly, R_2 is the distance from C' to the outer wheel loader tires. Radius R_3 is the distance from C' to the center of the wheel loader.

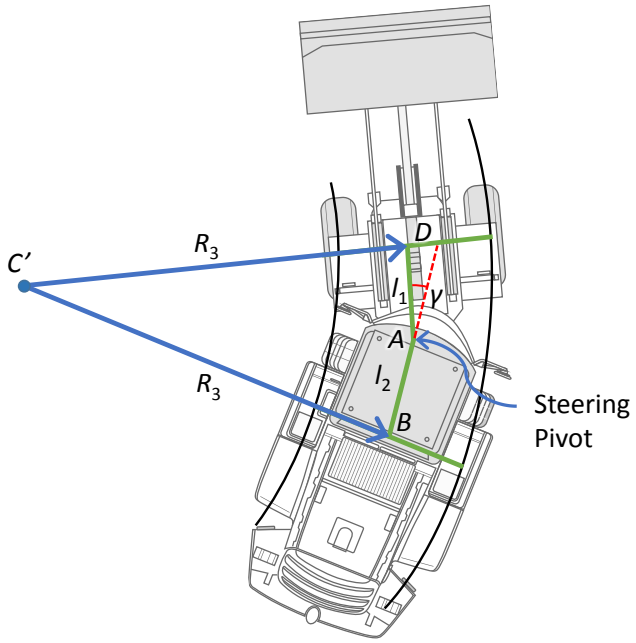

Figure 4.10. Kinematic representation of steered wheel loader.

In the figure, the wheel loader is being steered at an angle γ. The goal of the analysis is to find the relationship that links the radius R_3 with angle γ.

$$R_3 = f(\gamma) \tag{4.2}$$

The track width (the distance between the outer and inner radii) is given by:

$$d = R_2 - R_1 \ . \tag{4.3}$$

Another known value that is useful for this analysis is the wheelbase l. This value is the distance between the two axles. Since the pivot point of the steering lies exactly in the middle of the two axles, whenever the wheel loader steers, distance l is split by half. The two resulting lengths l_1 and l_2, which have the same value, are used to create the quadrilateral $ABC'D$ shown in Fig. 4.10. By tracing perpendicular lines with respect to l_1 and l_2 towards the direction of steering, the radius R_3 is obtained with respect to the instantaneous center of rotation C'.

From this, a simplified geometric diagram of the previous figure is generated. It takes the quadrilateral with vertices $ABC'D$ and sides R_3, l_1, and l_2. Two additional angles are introduced, ϕ and χ. Figure 4.11 displays the simplified diagram.

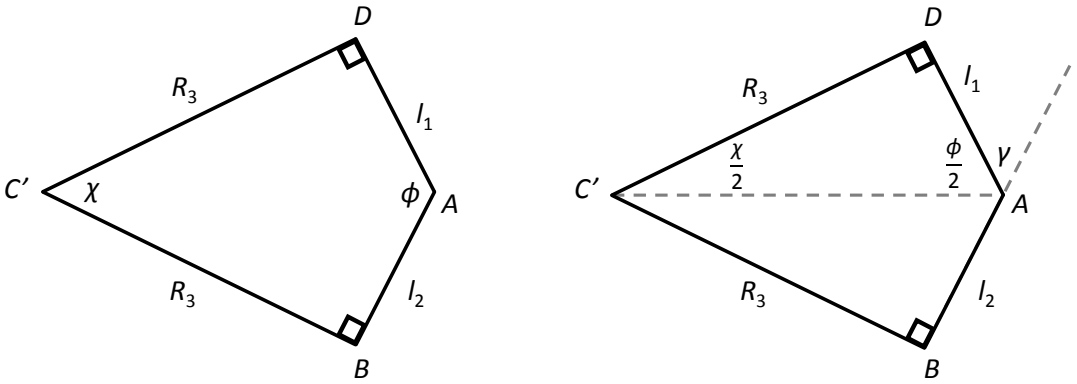

Figure 4.11. Geometric diagram of steered system.

Since $l_1 = l_2$, the quadrilateral proposed is symmetrical. A bisecting line is made between vertices C' and A, which yields two triangles, $AC'D$ and ABC'. Either of these triangles can be used to solve for the value of R_3. For this explanation, $AC'D$ is used. Based on Fig. 4.11(b), it can be seen that:

$$\phi = \pi - \gamma \ . \tag{4.4}$$

The relation of the angles of the triangle is given by:

$$\pi = \frac{1}{2}(\pi + \phi + \chi) \ . \tag{4.5}$$

By substituting the value of ϕ from (4.4), the following is obtained:

$$\frac{\chi}{2} = \frac{\gamma}{2} \ . \tag{4.6}$$

R_3, l_1, and angle γ are related through the tangent function.

$$\tan\frac{\gamma}{2} = \tan\frac{\chi}{2} = \frac{l_1}{R_3} \tag{4.7}$$

The final step is to solve for R_3:

$$R_3 = \frac{l_1}{\tan\frac{\gamma}{2}} = \frac{l}{2\tan\frac{\gamma}{2}} \ . \tag{4.8}$$

From the value of R_3, the values of R_1 and R_2 are simple to find.

$$R_1 = R_3 - \frac{d}{2} \tag{4.9}$$

$$R_2 = R_3 + \frac{d}{2} \tag{4.10}$$

Knowing these three radii, the steering compensation coefficient can be calculated and, in theory, can compensate for the steering of the machine in order for the controller to function correctly while steering. Of course, it is necessary to include logic to determine which side of the machine (right or left) is the outer radius and which is the inner radius. This determination is relatively straightforward using the sign of the steering angle measurement. When implemented correctly, the steering angle correction compliments the vehicle velocity estimation algorithm.

4.6 Time Delay Considerations

Time delays can be a major source of problems in control systems. When a delay exists between a certain command given by a controller and the actual execution of that command (e.g. rising current in a valve driver), it is possible for the added phase lag to create problems with the controlled system's ability to correctly follow a reference trajectory. Indeed, it is

even possible for a system with a stable open-loop plant and otherwise stable closed-loop dynamics to become unstable when incorporated into a control structure with a significantly long time delay. Therefore, it is important to determine, especially for a real-world system, whether or not the time delay between command and response is large enough to cause such issues.

For the system at hand, time delay tends to arise primarily from the electronic components. The computer acquisition module, valve drivers, and even sensors all typically incorporate at least a small amount of delay, the cumulative effects of which could noticeably harm TC system performance. Therefore, an investigation was conducted into potential methods for counteracting the effect of time delay. In the end, a method known as the *Smith predictor* was implemented to accomplish that. The following section describes this control structure.

4.6.1 The Smith Predictor

The Smith predictor attempts to quantify the effect of time delay on a dynamic system and correct the controller output accordingly [51]. A standard configuration of a system containing a Smith predictor is shown in Fig. 4.12. $P(s)$ is the transfer function of the open-loop system plant, which consists of a stable rational transfer function $P_0(s)$ and a time delay of ψ seconds. That is,

$$P(s) = P_0(s)e^{-\psi s} \ . \tag{4.11}$$

$\hat{P}_0(s)$ and $\hat{P}(s)$ are model representations of $P_0(s)$ and $P(s)$, respectively. The objective of the Smith predictor is to modify the error signal e using the difference between the predicted system response with and without the time delay, in order to compensate for the effects of the delay. The modified error e' is then used to generate the appropriate control signal using the original controller C_0.

The area outlined in red in Fig. 4.12 represents the Smith predictor, which consists of a primary controller $C_0(s)$ and a minor feedback loop, which contains the model of the process

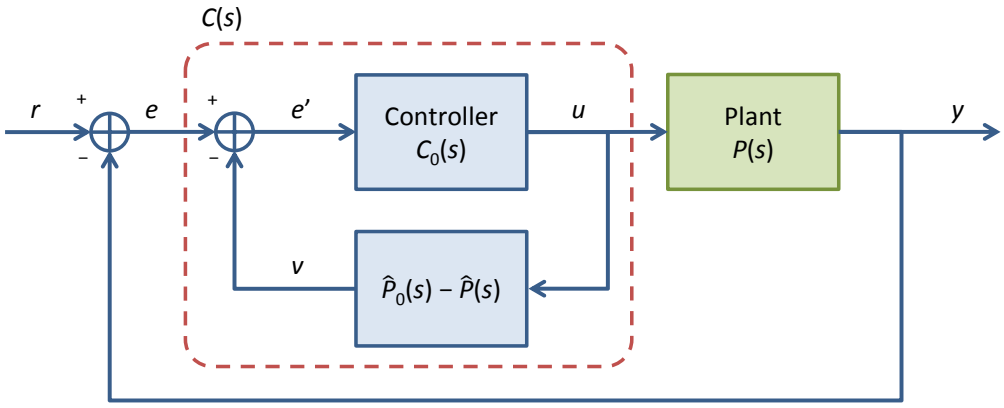

Figure 4.12. Control system incorporating a Smith predictor.

with and without the time delay. The overall transfer function of the time delay compensator is given by:

$$C(s) = \frac{C_0(s)}{1 + C_0(s)(\hat{P}_0(s) - \hat{P}(s))} \ . \tag{4.12}$$

The signal $v(t)$ contains a prediction of $y(t)$ occurring ψ seconds into the future. It is observed that $e' = rP_0u$, where $e = rPu$. Therefore, the *adjusted* error $e'(t)$, which is fed into the primary controller, brings that part of the error that is caused by the primary controller. This reduces the overcorrection related to the conventional controllers.

Assuming perfect model matching (that is, $\hat{P}(s) = P(s)$), the transfer function of the closed loop system in Fig. 4.12 from the reference trajectory to the output is:

$$G_r(s) = \frac{y(s)}{r(s)} = \frac{C_0P}{1 + C_0P_0} \ . \tag{4.13}$$

It can be noted that the time delay has been removed from the denominator of (4.13) thanks to the inclusion of the predictor. In fact, the denominator of (4.13) is the same as a feedback system with the same system plant and controller, but with no time delay (and of course, also removing the Smith predictor feedback loop). What this means is that, assuming perfect model matching, the Smith predictor removes all effects of the time delay from the feedback system. The only remaining impact of the time delay is now a

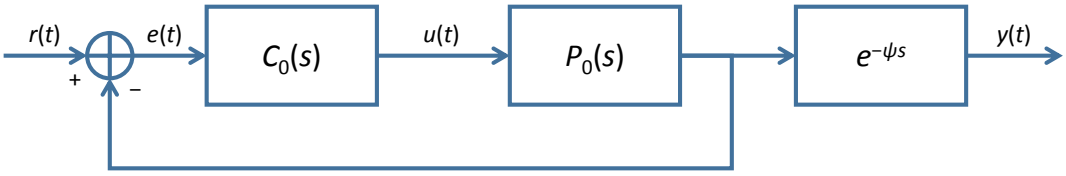

Figure 4.13. Equivalent closed-loop system assuming Smith predictor with perfect model matching.

delayed output signal. Essentially, a time-delayed system incorporating a Smith predictor with perfect model matching is equivalent to the system shown in Fig. 4.13.

4.6.2 Simulation Results

Figure 4.14 shows the results of two simulations, one with a Smith predictor to account for time delay and one without a predictor. The simulated braking system delay time was 50 ms, which is not an unreasonable number for real-world electronic systems. The differences in these two simulations are quite clear. For the system without the predictor (Fig. 4.14(a)) the controller overshoots the slip ratio setpoint, causing the slip ratio to oscillate around the setpoint value for more than 3.5 seconds before settling at the correct value. When incorporating the Smith predictor (Fig. 4.14(b)), on the other hand, the slip response has a much better behavior. The peak slip value increases very slightly (due to the fact that the Smith predictor decreases the severity of the controller somewhat), but it is quickly corrected and does not suffer from the same sustained oscillations as the uncorrected system. The settling time improves a little bit, as well. The main payoff, however, is the added system stability and the relatively smooth controller operation bringing the slip ratio down to the desired value.

This simulation shows the potential value of the Smith predictor, should system time delay ever become a significant issue for the real-world system. In simulation, increasing the braking system time delay from 50 to 100 ms (again, not unreasonable for the reference machine) caused the controlled system without a Smith predictor to become unstable, its oscillations growing higher and higher as time progressed. This behavior has not been

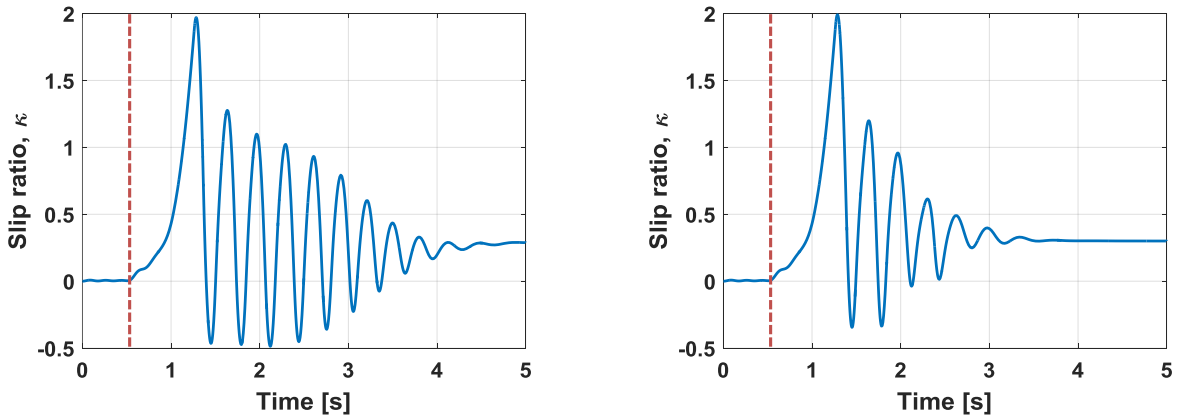

(a) (b)

Figure 4.14. Slip ratio responses for dynamic simulations of a system incorporating a time delay (a) without a Smith predictor and (b) with a Smith predictor.

seen on the actual machine (likely due to dissipative factors which are not included in the simulation). But it is conceivable that if different braking valves are tested which have a longer time delay, some problems like this could arise. In such a case, the Smith predictor should be revisited to assess its suitability at that time.

5. REAL-TIME CONTROLLER OPTIMIZATION

Chapters 5 and 6 of this work describe controls methods which go beyond the basic control strategy implemented in the traction control system. In general, a standard PID controller with the modifications discussed in Section 4.2 is adequate for wheel slip reduction, and it has been implemented on the real-world system successfully. Nevertheless, these discussions are included to give an idea of other frameworks which may be incorporated into the system to potentially improve results further.

This section specifically focuses on an optimization strategy which, in conjunction with a controller like those presented in this work, should be able to adapt the controller parameters to their best values. By incorporating a feedback in terms of an objective function to be maximized or minimized, the optimizer is able to modify the controller setpoint in real time, creating an adaptive controller which can update to maximize tractive force even in conditions where the tire-ground interface characteristics are changing over time.

5.1 The Extremum-Seeking Algorithm

One optimization strategy selected for this particular work is the extremum-seeking (ES) algorithm. It is a non-model-based approach which uses sinusoidal perturbations to ascertain the gradient of an objective function in real time. The parameter or parameters to be optimized are then adjusted according to the objective function gradient to either minimize or maximize it. The general continuous-time implementation of this algorithm is shown below (as formulated by Ariyur [52]). This approach has previously been successfully implemented on other systems by the author's research team [53, 54].

Essentially, the ES algorithm begins with an estimate of the optimal parameter value $\hat{\beta}$. By applying an additive sinusoidal perturbation to this signal, the actual parameter value β which is used in the system in question causes a response in the objective function $J(\beta)$. By passing the resulting signal through a high-pass filter (labeled "HPF" in Fig. 5.1), demodulating it by using a multiplicative perturbation of the same frequency, and using

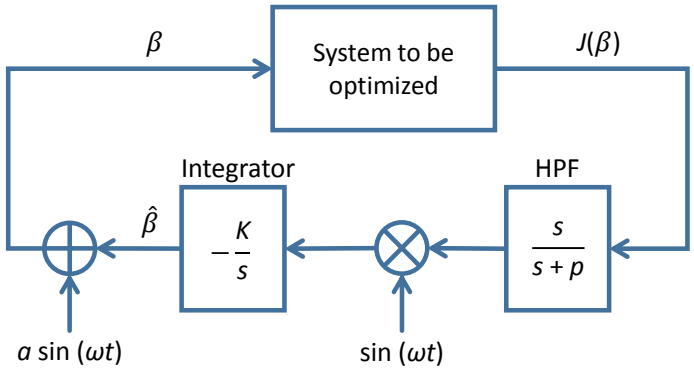

Figure 5.1. Basic extremum-seeking optimization scheme.

an integrator to remove high-frequency components, the estimate of the optimal parameter value $\hat{\beta}$ is driven toward the actual optimal value for the system β^*.

If the objective function of the system $J(\beta)$ has a well-defined singular optimum J^* at β^*, the ES algorithm has been shown to converge to that value, given proper setting of the ES parameters. Perhaps the most difficult aspect of implementing the ES algorithm is this: it introduces four new terms which must themselves be set before using the optimization strategy: ω, the frequency of the perturbations, a, the additive perturbation gain, p, the high-pass filter pole, and K, the integrator gain. These parameters do add a degree of difficulty to implementing the optimization, but once they are tuned sufficiently, the algorithm can work quite well.

It should also be noted that this particular system plant is quite well-suited to the ES optimization technique. In general, the slip-friction characteristic of tire-road surface interfaces looks something like the plot in Fig. 2.2. These plots are continuous, and most have a single, easily definable maximum at some wheel slip $\kappa^* > 0$. Therefore, given proper settings for the ES algorithm, the system should typically be able to find the best setpoint for the controller relatively efficiently.

5.2 Extremum-Seeking PID Control Structure

There are multiple important parameters used by a PID controller, including the gains K_P, K_I, and K_D for the proportional, integral, and derivative components of the control law, respectively. These can have a strong impact on the system if set incorrectly, even causing an otherwise stable system to become unstable in certain circumstances. However, once set correctly within reason, small changes in other traction control parameters, specifically the tire slip setpoint, should not cause the system to become unstable. Therefore, because they represent multiple parameters to be optimized simultaneously for each wheel, the controller gains were not chosen as the optimization parameters for the ES algorithm.

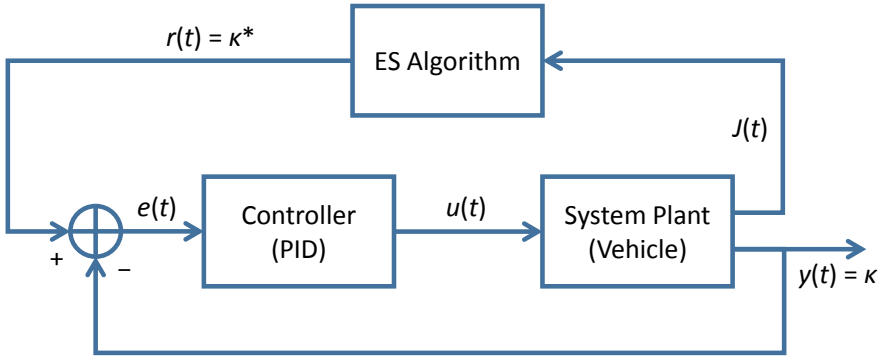

Figure 5.2. Control system block diagram for setpoint optimization.

Instead, the parameter to be optimized is the tire slip setpoint for each wheel (Fig. 5.2). By selecting the slip setpoint instead of the controller gains, at most one value per wheel needs to be optimized, instead of potentially three or more for each wheel. This greatly simplifies the optimization process. Furthermore, the tire slip setpoint is controlled directly by the operating condition for the vehicle, so it should be a fitting parameter to use for optimizing the controller.

In order to do this, the ES algorithm must be inserted into the system. What results is a system in which the tracking setpoint $r(t)$ is modified in real time. This reference trajectory is the equivalent of β, the output of the ES algorithm in Fig. 5.1. As optimization methods require a quantifiable objective function, the ES algorithm in Fig. 5.2 must receive feedback

information $J(t)$ from the system plant. The signal chosen for this feedback can have a drastic effect on the capability of the ES algorithm to optimize the controller performance.

In general, $J(t)$ needs to be chosen as a signal which quantifies the system performance in some meaningful way. For traction control systems, the primary desired output of the system is tractive force at the wheels. In standard on-road vehicle applications, the force can be approximated using the linear acceleration and wheel velocities, etc., as the vehicle is only propelling itself and is not attempting to interact with an external body. For construction machines like the wheel loader, however, the range of operating conditions and uses is significantly higher. In fact, the typical driving operation for this machine takes place at relatively low speeds, and its weight is enough that loss of traction when driving is usually not an issue.

For implementation in a real-world machine, a representative objective function needs to be constructed from available sensor data. This may not always be simple, but when specifically considering wheel loaders, a good estimate of the wheel force can be generated using implement boom and bucket cylinder pressures and angles (see Section 3.8).

On the other hand, simulations allow for much more direct access to the system states. Therefore, the objective function in simulation can simply be the total pushing force $F_{x,tot}$ of the machine. That is,

$$J(t) = F_{x,tot} = \sum_i F_{xi} \ . \tag{5.1}$$

This objective function may be overly simplistic for use in the real world. Therefore, the prototype implementation of this methodology includes a term for *traction efficiency*, similar to the work done by Osinenko [55], which can be included as an optimization tradeoff. Other systems may want to take into account the amount of wheel slip when creating the feedback objective function. However, for the sake of this simulation study, the purpose of the optimizer is simply to maximize the total pushing force of the machine.

The ES algorithm updates parameters based on the assumption that all changes in the parameters are reflected in the objective function. However, if the braking control signal is negative, the wheel speeds are not affected at all. That is to say, brakes are only capable of slowing down the wheels; they cannot increase wheel speed. However, the standard ES

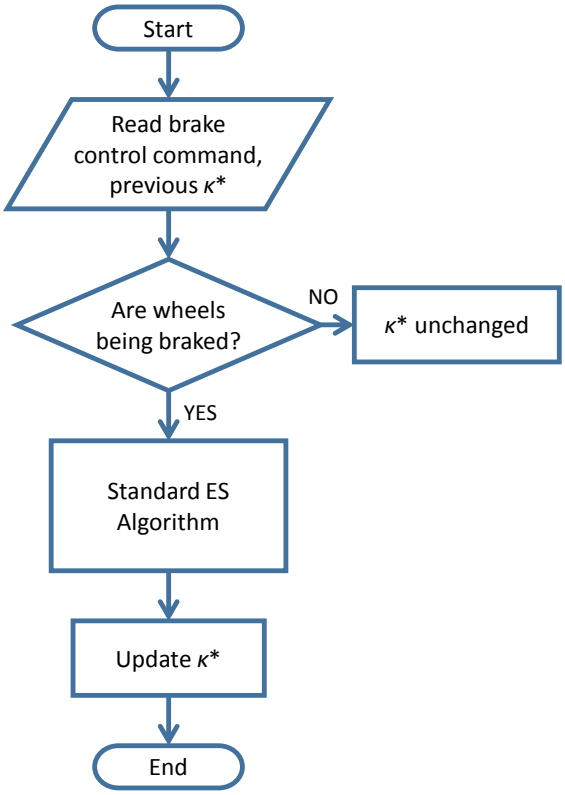

Figure 5.3. Augmented extremum-seeking algorithm.

algorithm does not take this into account. Therefore, during instances where the control signal is negative, the optimization code receives incorrect information about its effect on the system. This can cause the algorithm to converge to an incorrect point. Therefore, some logic has been added to the ES algorithm so that it does not modify the guess of κ^* unless the wheels are being braked (Fig. 5.3). The scheme incorporating this update logic is referred to in this work as the *augmented extremum-seeking algorithm*.

5.3 Simulation Results

Now that the system has been properly constructed and the optimization algorithm set in place, it is necessary to determine how well the system can perform. The simulation results shown here are based on the same setup described in Section 3.7.

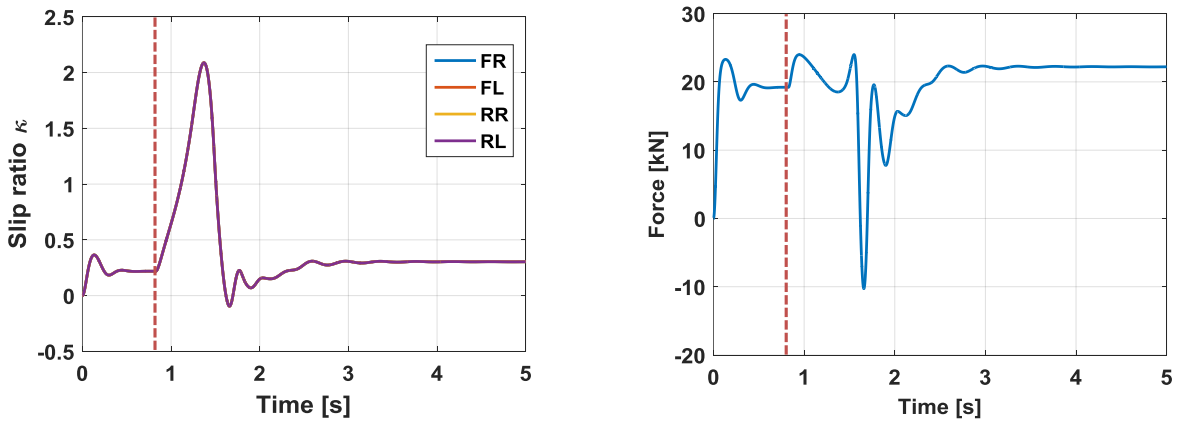

(a) (b)

Figure 5.4. Simulation results for extremum-seeking traction control in high-friction condition, showing (a) wheel slip ratio and (b) total longitudinal force.

For the first simulation (Fig. 5.4), the vehicle is on a ground condition with relatively high friction for all four wheels (the high-μ example from Fig. 2.2). The wheel loader impacts the tires, and its wheels begin to slip against the ground. The response of the TC system with extremum-seeking is quite consistent with that of the previously shown PID system (Section 4.2), even though the system did not start with a correct value of wheel slip setpoint. This indicates that the ES algorithm is allowing the system to quickly find an acceptable value of slip setpoint.

Figure 5.5 shows the result of the setpoint optimization for this simulation using the augmented extremum-seeking algorithm, as well as a comparison to the same simulation using an unmodified ES algorithm (i.e. exactly as shown in Fig. 5.1 with no added logic). The tire-road interface has a known maximum friction coefficient at a slip ratio of around $\kappa^* \approx 0.50$. This was found by analytically determining the location of the maximum friction coefficient for the slip friction curve.

It can be seen that, even when the system is not being braked, the unmodified ES algorithm still attempts to use the force information to update the setpoint. As a result, it ends up overshooting the actual optimum value and, as the force continues to change, it could potentially wander farther from the correct point. If, on the other hand, the optimization is

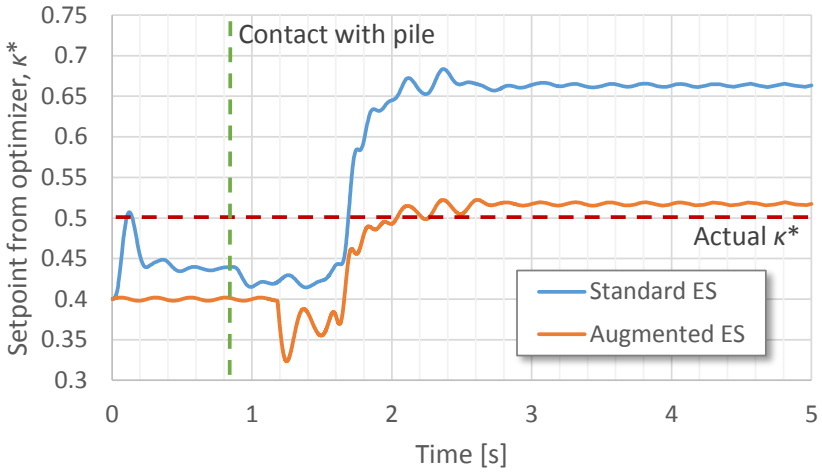

Figure 5.5. Performance comparison of standard ES algorithm with updated version (high-friction).

updated only when the brakes are being actuated (as in the augmented ES algorithm), the optimization avoids the effect of bad objective function information. Therefore, instead of the haphazard development of the standard ES algorithm, the updated ES algorithm pauses when no wheels are being braked and consistently approaches the correct setpoint as time progresses.

A second simulation was then run using the same system in a different ground condition. This time, all four wheels were simulated as being on slick ground (the low-μ example from Fig. 2.2). In this condition, the traction control system must work much harder to keep the wheels from spinning. Furthermore, the maximum pushing force is found at a different value of slip ratio ($\kappa^* \approx 0.36$). This means that, whereas the ES algorithm needed to increase the setpoint in the previous simulation, with this condition it must decrease the setpoint somewhat. The results of this simulation are shown in 5.6.

In this simulation, the controller setpoint was again initialized at a non-optimal value. In fact, it started at the same point as the high-friction simulation ($\kappa_0 = 0.40$). In this case, the ES algorithm decreased the setpoint during the simulation, so that it achieved a better value. Furthermore, the traction control system was able to contain the wheel slip to an appropriate value such that the tractive force was maximized (again, after some initial

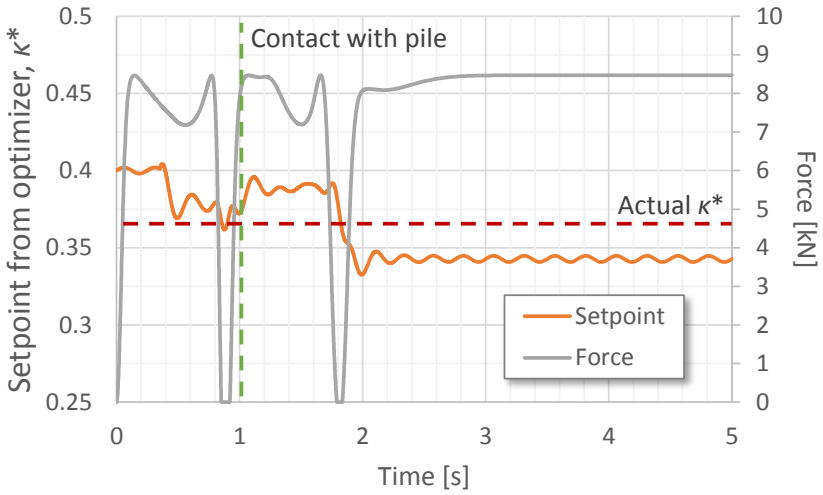

Figure 5.6. Simulated system performance (low friction).

overshoot where the tractive force was briefly decreased). These simulations indicate that the system developed in this work is capable of determining and reaching an optimal slip value given the proper feedback signal for the optimization objective function.

5.4 Adaptation for Prototype Implementation

To apply this methodology to a real-world system, certain considerations had to be made. The first is the identification of a better objective function. Tractive force is often the value which is most intuitive as an objective. However, the slip-friction relationship has some characteristics that do not always lend themselves well to efficient optimization. The second consideration deals with the selection of the optimization methodology used to alter the controller setpoint. The extremum-seeking algorithm is an excellent method for real-time optimization in simulated systems. It has many benefits in that it is a true "real-time" approach: that is, it is continuously updating the decision variable based on the real-time value of the objective function. However, the ES algorithm struggles in some ways when it is applied to real-world systems.

This is due to several factors. First, the ES algorithm can be very sensitive to small changes in operating condition. ES settings which converge quite well to one operating

condition can have trouble when the condition is changed. Furthermore, as the ES algorithm works with current system output values, it can be very susceptible to effects such as noise, which can be quite noticeable in mobile machine applications. Therefore, the online controller setpoint optimization scheme was modified to allow it to better cope with the difficulties of running on a prototype machine. This work is outlined in more detail in [56].

5.4.1 Tractive Energy Considerations

Thus far in the discussion, optimization of the traction control system has been centered on maximizing the pushing force of the system. This is an intuitive choice to make, as pushing force is most often a proxy for other values which are of interest for the particular application (vehicle acceleration, material mass excavated from a pile, etc.). However, there are some aspects of pushing force which can cause difficulties when attempting to optimize the system setpoint.

The main issue at hand stems from the fact that the slip-friction relationship for a given tire and ground condition may not have a well-defined maximum value. As discussed in Section 2.3.2, the slip-friction relationship depends on many factors. Even in the example shown in Fig. 2.2, the location of the maximum friction value for the low-friction condition is not very well defined (hence the slight inaccuracy resulting in the ES optimization shown in Fig. 5.6). It is entirely possible that physical system effects such as noise could cause a significant amount of error when determining the slip value to define the maximum pushing force. Furthermore, there is no guarantee that all systems will have a bounded maximum value at all. From the experimental data shown in Fig. 8.4, the data taken in the ground condition "Snow on grass" increases monotonically, never reaching a maximum value. Therefore, any optimization function attempting only to maximize pushing force would increase the setpoint to the maximum allowable boundary. Based on the work done in this research, this monotonic behavior is not uncommon for ground conditions which are easily deformed, such as soil.

It is obvious that allowing the machine wheels to slip uncontrollably is not the optimal solution for this problem. To this end, it is important to ensure that the optimization

function used in this work is capable of finding a suitable reasonable value for optimal slip ratio. This implies that the objective function of the optimization needs to be adapted from that in Eq. (5.1). Wheel friction coefficient is still an important factor, but a new term is needed which captures the performance of the system. For off-road vehicles such as the wheel loader considered in this work, a suitable parameter for capturing these effects is the *traction effectiveness* (also known as *traction efficiency*). The goal of the traction effectiveness metric is to determine the amount of power transferred into each wheel that is converted into useful pushing force [25].

The traction effectiveness $\eta_{tr,i}$ at wheel i is typically modeled as follows (see [55]).

$$\eta_{tr,i}(\kappa_i) = \frac{\mu_{xi} - \mu_{roll}}{\mu_{xi}}(1 - \kappa_i) \tag{5.2}$$

Other considerations have been made for certain applications and operating points, such as the work done by Schreiber [57, 58]. However, for this research, definition (5.2) is sufficient. The traction effectiveness takes into consideration both the friction force at the wheel μ_{xi} and the rolling resistance μ_{roll}, so having reasonable models for those terms is important in determining traction effectiveness.

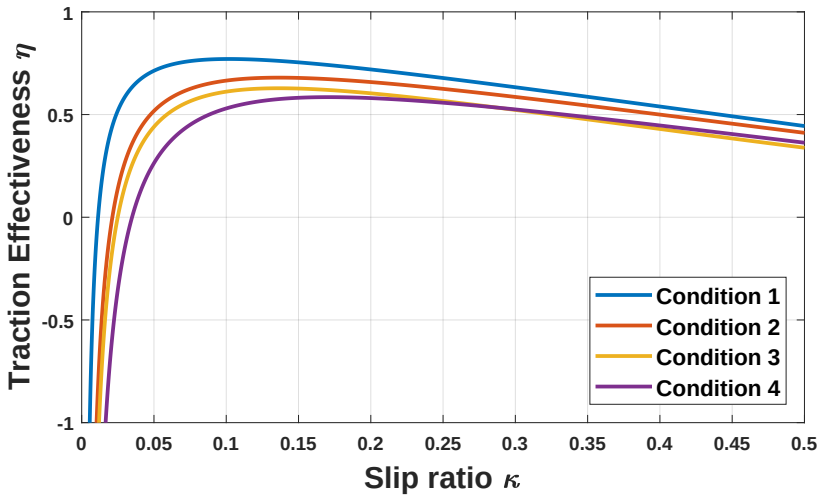

Figure 5.7. Examples of traction effectiveness curves.

Figure 5.7 shows four examples of traction effectiveness curves from various ground conditions. Two aspects of this plot should be noted. First, as shown in Eq. (5.2), traction effectiveness is a function of the slip ratio. Therefore, for any given ground condition it can be optimized over the slip ratio without danger of other factors affecting performance. Second, for all four ground conditions shown in the figure, there is a well-defined maximum at some reasonable slip ratio value. This implies that combining the traction efficiency with the friction coefficient should lead to an objective function which allows the optimization scheme to efficiently converge to an acceptable value of slip setpoint.

5.4.2 Redefinition of Optimization Problem

The optimization problem itself was also reexamined to assess whether there was a more effective setup which would allow for fast, consistent results in experimentation. To this end, a new problem definition was formulated. Most commonly available optimization methods do not handle real-time date well, so some amount of basic processing needed to be done to construct an optimization problem which is easily solvable by one of those methods.

Instead of examining the data one point at a time, the optimization scheme needs to act on aggregate data, so that a proper objective function can be constructed. Therefore, data needs to be stored in some memory structure (see the discussion in Section 5.4.4), and then the aggregate data is used for the optimization.

In fact, the final optimization implementation consists of two minimization problems. The first actually behaves more as a system identification method.

$$\underset{B_x, C_x, D_x}{\text{minimize}} \quad J_0(x) = \text{MSE}\left(\hat{\mu}_x(B_x, C_x, D_x)\right) \tag{5.3}$$

In this equation, the decision variables which are being optimized are actually the Magic Formula parameters B_x, C_x, and D_x, as in Eq. (2.4) (E_x is assumed to be zero, which forces the function to be odd). For a given set of averaged wheel slip values κ and corresponding friction coefficient values μ_x, this function finds the Magic Formula parameter values which

minimize the mean squared error between the data and the Magic Formula model. Mean squared error is defined as shown in (5.4).

$$\text{MSE}\left(\hat{\mu}_x(B_x, C_x, D_x)\right) = \frac{1}{n}\sum_{j=1}^{n}\left(\mu_{x,j} - \hat{\mu}_{x,j}(\kappa_j, B_x, C_x, D_x)\right)^2 \tag{5.4}$$

Essentially, for all n data points, the estimated (model) friction coefficient value $\hat{\mu}_{x,j}$ is subtracted from the actual data value $\mu_{x,j}$ and the result is squared. These values are summed and divided by n. By minimizing this function, the first minimization results in values B_x, C_x, and D_x which minimize the model error. That is to say, this function automatically fits the Magic Formula to the aggregate data.

Having constructed a model fit to the data, it is much simpler to construct a second optimization problem which will result in consistent behavior for determining the best wheel slip setpoint. The objective function for the new optimization is a combination of the friction coefficient as formulated in Eq. (2.4) and the traction effectiveness from Eq. (5.2).

$$J(\kappa) = -\delta_1\hat{\mu}_x(\kappa) - \delta_2\eta_{tr}(\kappa) \tag{5.5}$$

In this equation, δ_1 and δ_2 are weighting coefficients which control the relative contribution of the friction coefficient and the traction effectiveness to the overall multi-objective optimization. The larger the value of δ_1 compared to δ_2, the more strongly the optimization will favor maximizing friction coefficient over traction effectiveness. It is common practice (including in this work) to impose the restriction $\delta_1 + \delta_2 = 1$; however it is not necessary to do so. Finally, both terms in Eq. (5.5) are negative, as standard optimization practice is to represent the task as a minimization problem [59].

Having defined this new multi-objective function, the new optimization problem for automatic parameter tuning is as follows.

$$\begin{aligned} \underset{\kappa}{\text{minimize}} \quad & J(\kappa) \\ \text{subject to} \quad & \kappa > 0 \end{aligned} \tag{5.6}$$

Given a suitable optimization technique, this problem formulation should be able to set κ^* to a value of slip ratio which provides a good compromise between pushing force and tractive efficiency. The constraint $\kappa > 0$ exists in order to keep the optimization algorithm from encountering the numerical singularity in the traction effectiveness term at $\kappa = 0$. However, as it is known that the optimal slip setpoint should always be positive, this restriction also serves to keep the algorithm from attempting nonsensical values.

5.4.3 Alternative Optimization Methods

The ES algorithm represents one paradigm for real-time optimization, wherein a time-based signal is conditioned in such a way that the output of the signal operation is itself driven to the optimal value as a response to the changing objective function. Many optimization techniques, however, do not behave in the same manner. Generally, an optimization works on a function which is already mapped and has no concept of values changing with time. It is for this reason that the optimization scheme developed in Section 5.4.2 is constructed to handle a set of data instead of a single point.

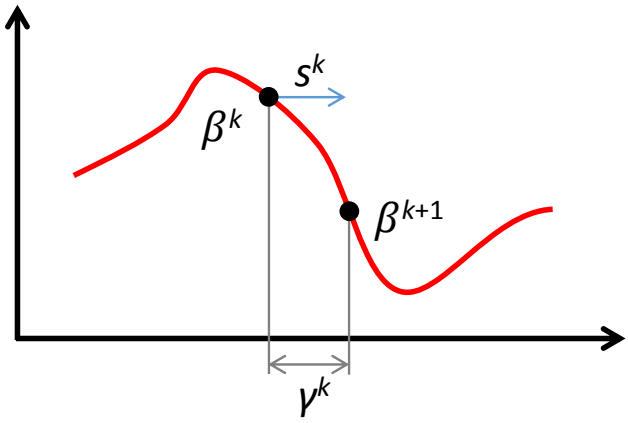

Figure 5.8. Standard optimization procedure.

The typical procedure for finding the location of a function's minimum value is shown in Fig. 5.8. Starting from a location β^k at iteration k of the optimization with an objective function value $J(\beta^k)$, a search direction s^k is first chosen based on some criteria. Then a

step length γ^k is determined such that the value of J is reduced after the iteration (that is, $J(\beta^{k+1}) < J(\beta^k)$). In equation form, this is represented as:

$$\beta^{k+1} = \beta^k + \gamma^k s^k \ . \tag{5.7}$$

What remains is to select a suitable optimization technique for determining s^k and γ^k. There are several factors which go into determining which optimization technique is best suited for a given problem. The type of problem itself plays a large role. The differentiability of the objective function is an important aspect, as well as desired computation time. In the end, the method must be based on decision variable type, objective function type, computational power, desired computing speed, and desired precision, among others.

As the objective functions at hand are relatively well-behaved and do not contain any discontinuities or other irregularities, there are many suitable optimization techniques. For this work, the quasi-Newton method with a Broyden-Fletcher-Goldfarb-Shanno (BFGS) update rule was selected. This is a first-order approach which uses the gradient of the objective function $\nabla J(\beta^k)$ and L, an approximation of the objective function's Hessian matrix, to generate the search direction s^k.

$$s^k = -L^k \nabla J(\beta^k) \tag{5.8}$$

The step length γ^k is then found using a simple line search minimization.

The quasi-Newton approach is a first-order optimization method (i.e. only utilizing the first derivative of the objective function) which can approach the performance of a second-order method [60]. This is useful, because it reduces the amount of computation effort needed to determine second derivatives of all the objective functions. Because of its reliability and speed, the quasi-Newton approach with BFGS update rule was used for both minimizations (system identification and setpoint optimization) within the parameter update algorithm for this work.

5.4.4 Data Buffering

As has already been mentioned, the new setpoint tuning technique requires a whole set of data in order to run. Therefore, it is necessary to construct a method for how this data set will be stored. To create an aggregate data set from a time-varying signal, a structure called a *data buffer* is typically used. The data buffer takes the real-time data and stores values until they are needed.

The standard buffering strategy is simply to store the most recent n data points, and delete all older data. Then, when an operation is performed on the data set, it has only the most recent values. This is sufficient for many applications; however, there are some weaknesses to this approach. The primary drawback is that this buffering approach assumes that all data collected within an arbitrary time window is equally useful. For the current system, that is not necessarily the case.

As shown in Fig. 2.2, it is very common for the Magic Formula model of the slip-friction relationship to have two separate regions: a low-slip regime in which the relationship increases quickly and somewhat linearly, and a high-slip regime wherein the data flattens out and either continues rising or begins to fall with a much shallower slope. Both of these regimes contain important information about the slip-friction relationship which cannot be ignored without losing significant model fidelity.

These considerations led to the development of a so-called *enhanced buffer*, comprising two *selective buffers* running in parallel (Fig. 5.9). Based on the value of the slip (high or low) at any given point in time, the corresponding data is allowed to enter one of the two buffers. When performing the optimization, the data from both selective buffers is used. What this means is that the optimization scheme captures a significant amount of data from both sections of the curve. The selective buffer for the high-slip regime is allowed more data points, as more variability exists in this regime.

Once the optimization is run, the high-slip regime buffer is emptied, while the low-slip buffer never empties, but maintains the most recent values. As more data is generated, both buffers are updated to include new slip values. When the high-slip buffer is filled, the optimization runs using both selective buffers to capture as much of the curve as possible.

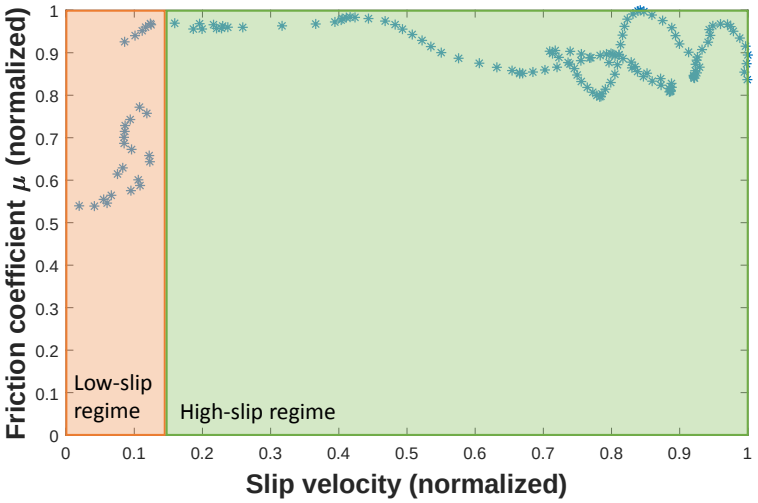

Figure 5.9. Low- and high-slip regimes for enhanced data buffer, shown with experimental data.

Using both buffers in this way has the downside of potentially using older data for the low-slip regime, if the system has remained in the high-slip regime for an extended length of time. However, if the optimization were allowed to run with data from only a small part of one regime, the resulting system identification phase could result in a Magic Formula model which is far from matching the real-world conditions outside the zone of the measured slip. Therefore, including what could be relatively old data is still preferable to simply using the most recent values. The enhanced buffer is an intelligent data-handling structure which is capable of overcoming the limitations of inconsistent data generation in the system.

5.4.5 Simulation Results

To show the potential of this optimization scheme for updating with changing ground condition, a simulation was run similar to previous simulations, with the vehicle accelerating toward a barrier until it makes contact and the wheels begin to slip against the ground. However, in this case, at time $t = 4$ seconds, the ground condition is changed to one which is significantly different. It should be noted that, as the simulation does not suffer from the same noise problems as the experiments, a standard data buffer was used for this simulation.

Therefore, the slip ratio setpoint updates at regular intervals. This is not the case for the enhanced data buffer described in Section 5.4.4.

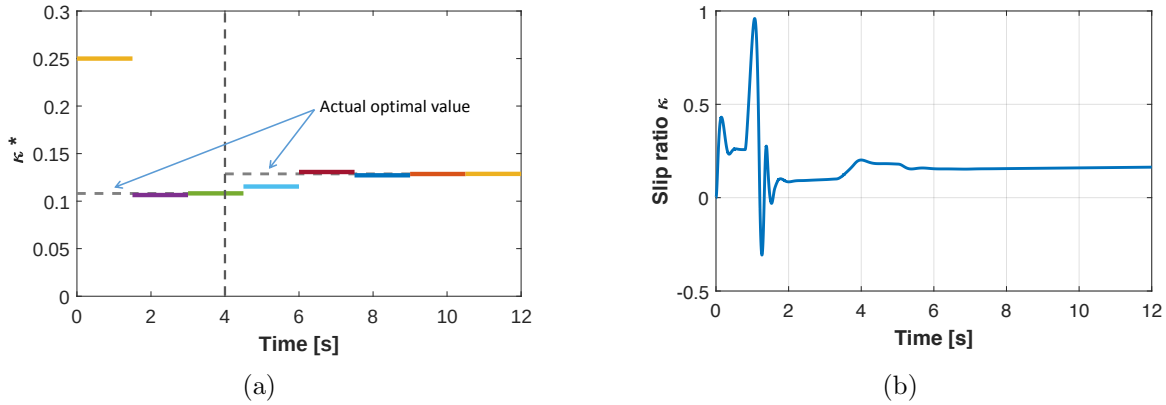

(a) (b)

Figure 5.10. Simulation results for optimal slip setpoint using quasi-Newton method with buffered data, showing (a) slip ratio setpoint and (b) actual controlled slip ratio.

Figure 5.10 shows the results of this simulation. It can be seen from Fig. 5.10(a) that the initial guess of slip setpoint $\kappa_0^* = 0.25$ is not very close to the actual optimal value for the system. However, once the buffer fills up and the optimization algorithm runs, the new setpoint (purple) is very close to the correct value. Again, when the ground condition changes after 4 seconds, there is a delay due to waiting on the data buffer to fill. The first setpoint after the ground condition change (light blue) is between the two optimal setpoints, because it contained data from both ground conditions. Nevertheless, at the very next iteration (red), the system has again converged to essentially the correct value. Furthermore, as shown in Fig. 5.10(b), the traction controller is maintaining the correct slip values (after brief transient stages).

The two objective functions used for this multi-objective optimization are shown in Fig. 5.11. In this particular simulation, the weights of the two objective functions from Eq. (5.5) are $\delta_1 = 0.2$ and $\delta_2 = 0.8$. This means that the traction effectiveness was much more heavily emphasized than the friction coefficient. This is easily seen from Fig. 5.11(a). At the initial guessed setpoint of $\kappa_0^* = 0.25$, the traction effectiveness only achieves a value of 0.68. However, after the optimization routine and subsequent action of the traction control system,

that value rises to 0.77 (an increase of 14%. After the ground condition changes to the lower-friction condition, the optimization maintains the traction effectiveness at 0.75. However, this emphasis on traction effectiveness comes with a loss in terms of friction coefficient (Fig. 5.11(b)). At the outset, the friction coefficient is around 0.15. However, once the first optimization takes place, the traction controller alters the behavior of the system such that friction coefficient decreases to 0.072 for the initial ground condition and 0.062 for the lower-friction condition.

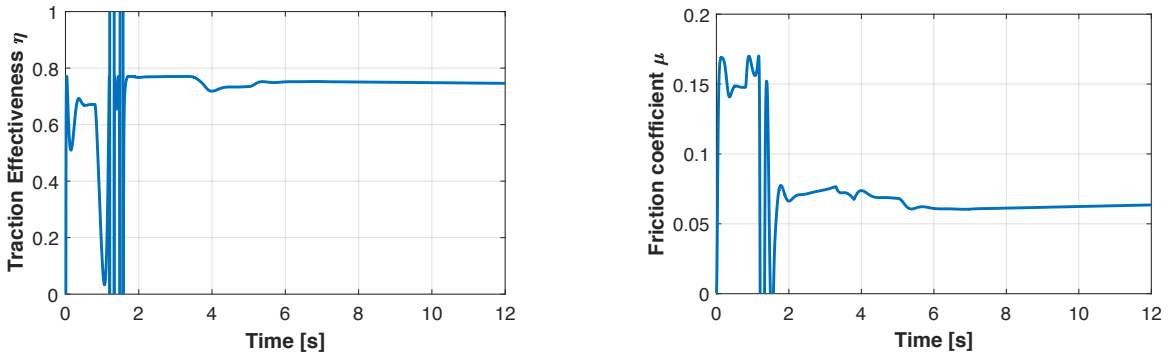

(a) (b)

Figure 5.11. Simulation results for (a) traction effectiveness and (b) friction coefficient, using quasi-Newton method with buffered data.

This result is heavily dependent on the weighting coefficients of Eq. (5.5). At first glance, it would appear that the optimization is not working properly, as the resultant friction coefficient has been so significantly decreased. However, examining the actual multi-objective function shows that the optimization is indeed performing well. Had the initial setpoint been kept (i.e. $\kappa^* = 0.25$ for the whole simulation), the objective function $J(\kappa^*)$ would have had a value of -0.57 for the first ground condition and -0.52 for the second. The optimization, on the other hand, achieves objective function values of $J(\kappa^*) = -0.63$ and $J(\kappa^*) = -0.55$ for the same conditions, improvements of 11% and 6%, respectively. Therefore, the objective function is certainly being minimized.

In this case, friction coefficient (i.e. traction force) and traction effectiveness are competing objectives. That is to say, when the system is optimized, an increase in one objective

tends to lead to a decrease in the other. Therefore, it is possible to represent this tradeoff in terms of the amount lost in one variable for an incremental increase in the other. By varying the values of δ_1 and δ_2, the Pareto front in Fig. 5.12 was constructed. This figure shows the tradeoff between friction coefficient and traction effectiveness for two different ground conditions. It is clear that these two objectives do indeed compete. Because of this, the balance between them (dictated by δ_1 and δ_2) must be determined by the system designer. It is likely that this balance will be application-dependent, with some users preferring to maximize traction force at the expense of efficiency, while others choose a more even approach. There is no one correct balance to be struck.

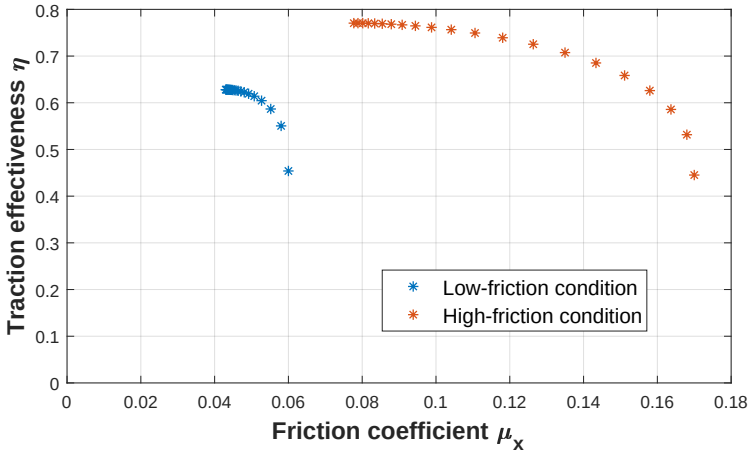

Figure 5.12. Pareto front showing tradeoff between friction coefficient and traction effectiveness at different ground conditions.

Having shown the capability of the optimization scheme in simulation, it was then implemented on the real-world prototype system. The resulting proof-of-concept tests are documented in Section 8.3.

6. NONLINEAR CONTROLLERS

The majority of the activities done in this work of the project have focused on preparing the wheel loader system for experimental activities and creating a control scheme which would be able to work on the final production system, which has limitations in terms of available sensors and computing power. However, a separate parallel investigation was conducted into alternate control methodologies which may have some benefit for the system beyond that of the Augmented PID structure used throughout the project.

Therefore, four different controller designs were realized in the MATLAB Simulink environment. The first is a feedback linearization controller, which seeks to cancel out the effect of the system nonlinearities in order to produce a system which is capable of being conditioned in terms similar to a linear system. The second and third, both sliding mode controllers, are capable of controlling the nonlinear system even in the presence of parametric uncertainties. Finally, an adaptive controller was developed which is capable not just of overcoming the effect of parametric uncertainties, but also of correcting the real-time parameter estimates used in the model. While the Augmented PID was used for most of the real-world testing, a sliding mode controller was also implemented on the prototype and its performance compared with the Augmented PID. The sliding mode control shows a noticeable improvement over the Augmented PID control in laboratory testing (see data and discussion in Section 8.4).

6.1 Feedback Linearization

In order to assess the system performance, both with and without the controller, it is necessary to create a system model out of the physical equations generated in Chapter 3. These equations lend themselves well to a state-space formulation, so such a system has been

arranged here. Defining states x_1, x_2, and x_3 as v_x, ω_i, and κ_i, respectively, and taking the inputs u_1 and u_2 to be T_E and T_B, respectively, the model for a single wheel becomes:

$$\begin{aligned}
\dot{x}_1 =& \frac{1}{m}(-\mu_{roll} F_N \operatorname{sgn}(x_1) + \mu_x(x_3) F_N) = f_1(x) \\
\dot{x}_2 =& \frac{1}{I_w}(-r_e \mu_x(x_3) F_N + u_1 - u_2) = f_2(x) \\
\dot{x}_3 =& \frac{1}{B}(-x_1 + r_e x_2 - |x_1|\, x_3) = f_3(x) \ .
\end{aligned} \tag{6.1}$$

This is a relatively simplified version of the system dynamics. Also, for the purposes of this discussion, it will be assumed that both the engine torque and braking torque can be controlled simultaneously. While this may not be the case for all systems, it is certainly a reasonable assumption, and it significantly simplifies the analysis for this paper by making the system a single-input system. Thus, u_1 and u_2 are combined into a single input u, that is:

$$\begin{aligned}
\dot{x}_1 =& \frac{1}{m}(-\mu_{roll} F_N \operatorname{sgn}(x_1) + \mu_x(x_3) F_N) = f_1(x) \\
\dot{x}_2 =& \frac{1}{I_w}(-r_e \mu_x(x_3) F_N + u) = f_2(x) \\
\dot{x}_3 =& \frac{1}{B}(-x_1 + r_e x_2 - |x_1|\, x_3) = f_3(x) \ .
\end{aligned} \tag{6.2}$$

The idea behind feedback linearization control is to use a series of nonlinear coordinate transformations to create a new "equivalent linear time-invariant (LTI) system" which is linear from some virtual input v to the system output y. To generate this controller, first examine the system:

$$\begin{aligned}
\dot{x} =& f(x, u) \\
y =& h(x) \ ,
\end{aligned} \tag{6.3}$$

where f and h are nonlinear functions. It is desired to use a nonlinear coordinate transformation to generate a new state vector z and coordinate transformation g_2 such that:

$$z = g_2(x) \quad \text{and} \quad x = g_2^{-1}(z) \ . \tag{6.4}$$

Also needed is a feedback linearization control law,

$$u = g_1(x, v) \ , \tag{6.5}$$

generated such that the transformed system is LTI from input v to output y. The overall schematic of the feedback linearization controller is shown in Fig. 6.1. Since the equivalent system is LTI, standard linear methods can be used to generate the control signal v.

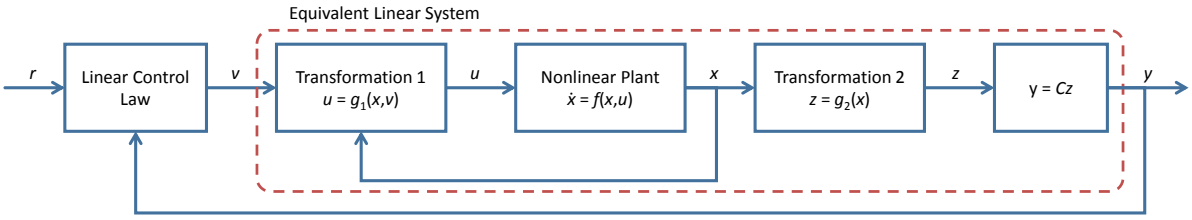

Figure 6.1. Feedback linearization control schematic.

The transformation g_2, the control transformation law g_1, and the LTI control law which generates v have still not been chosen. The process for doing this begins by putting the system in the normal single input-single output (SISO) form:

$$\begin{aligned} \dot{x} &= f(x) + g(x)u \\ y &= h(x) \ . \end{aligned} \tag{6.6}$$

For the present system, the functions f and g can be determined from the state-space equations in Eq. (6.2). For the nonlinear controllers in simulation, the slip ratio κ_i of the system was chosen as the process variable. Therefore, the system output was chosen as follows.

$$y = \kappa_i = x_3 = h(x)$$

In order to determine the appropriate coordinate transformation for the system, it is necessary to find the relative degree q of the system. The relative degree is the number of

time derivatives of y (that is, dh/dt) it takes to have a form explicitly containing the input u. Therefore, the first derivative of y is taken.

$$\dot{y} = \frac{\partial h}{\partial x}\dot{x} = \dot{x}_3 = \frac{1}{B}(-x_1 + r_e x_2 - |x_1|\, x_3)$$

This does not contain u explicitly, so another derivative must be taken. In this case, the process makes use of the Lie derivative L_f:

$$L_f h(x) = \frac{\partial h(x)}{\partial t} f(x) \ . \tag{6.7}$$

This operator is used to compute the second derivative of y.

$$\ddot{y} = \frac{\partial}{\partial x}\left(L_f h(x)\right)\dot{x} = \frac{\partial}{\partial x}(\dot{y})\dot{x} \tag{6.8}$$

For the present system in vector form, this comes out to be:

$$\ddot{y} = \frac{1}{B}\begin{bmatrix} -1 - x_3\,\mathrm{sgn}(x_1) & r_e & -|x_1| \end{bmatrix}\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix}$$

Expanding for the state derivatives:

$$\begin{aligned}
\ddot{y} =& \frac{1}{B}\left[\left(\frac{-1 - x_3\,\mathrm{sgn}(x_1)}{m}\right)\left(-\mu_{roll}F_N\,\mathrm{sgn}(x_1) + \mu_x(x_3)F_N\right)\right. \\
&\left. + \frac{r_e}{I_w}\left(-r_e\mu_x(x_3)F_N + u\right) - \frac{|x_1|}{B}\left(-x_1 + r_e x_2 - |x_1|\, x_3\right)\right] \\
=& \frac{F_N}{mB}\left(\mu_{roll}\,\mathrm{sgn}(x_1) - \mu_x(x_3)\right) + \frac{F_N}{mB}x_3\left(\mu_{roll} - \mu_x(x_3)\,\mathrm{sgn}(x_1)\right) \\
&- \frac{r_e^2 F_N}{I_w B}\mu_x(x_3) + \frac{1}{B^2}\left(x_1\,|x_1| - r_e x_2\,|x_1| + x_1^2 x_3\right) + \frac{r_e}{I_w B}u \ .
\end{aligned}$$

This derivative explicitly contains the input u. Therefore, the relative degree of the system is $q = 2$. This means that in the coordinate transformation will have two external state variables. These state variables are defined from the selected output and its derivative.

$$\xi_1 = y = x_3 = h(x)$$
$$\xi_2 = \dot{y} = L_f h(x) = \frac{1}{B}(-x_1 + r_e x_2 - |x_1|\, x_3)$$
$$\xi = \Psi(x) = \begin{bmatrix} h(x) \\ L_f h(x) \end{bmatrix}$$

The state vector ξ only has $q = 2$ states. Therefore, in order to characterize the full system (with $n = 3$ states), an additional $n - q = 1$ internal state is needed. This state must be independent of the other two states. In this case, either remaining original state (x_1 or x_2) can be chosen. Since the dynamics for x_1 do not contain u, it was chosen as the internal state (the absence of u simplifies the process later on).

$$\eta = \sigma(x) = x_1$$

Therefore, a new set of state variables is defined:

$$z = \begin{bmatrix} \xi \\ \eta \end{bmatrix} = \begin{bmatrix} h(x) \\ L_f h(x) \\ \sigma(x) \end{bmatrix} = g_2(x) = T(x)\ , \tag{6.9}$$

which characterizes the system completely if the Jacobian of T is nonsingular. For the case of the present system, this is true. Thus, the system is well defined by the new states.

Along with the new transformed states comes a new set of nonlinear state-space dynamics.

$$\dot{\xi}_1 = \xi_2$$
$$\dot{\xi}_2 = \frac{F_N}{mB}\left(\mu_{roll}\operatorname{sgn}(x_1) - \mu_x(x_3)\right) + \frac{F_N}{mB}x_3\left(\mu_{roll} - \mu_x(x_3)\operatorname{sgn}(x_1)\right) - \frac{r_e^2 F_N}{I_w B}\mu_x(x_3)$$
$$+\frac{1}{B^2}\left(x_1|x_1| - r_e x_2|x_1| + x_1^2 x_3\right) + \frac{r_e}{I_w B}u = \alpha(x) + \beta(x)u$$
$$\dot{\eta} = \dot{x}_1 = \frac{F_N}{m}(-\mu_{roll}\operatorname{sgn}(\eta) + \mu_x(\xi_1)) = \gamma(\xi,\eta)$$

This leads to the following nonlinear functions α, β, and γ.

$$\alpha(x) = \frac{F_N}{mB}\left(\mu_{roll}\operatorname{sgn}(x_1) - \mu_x(x_3)\right) + \frac{F_N}{mB}x_3\left(\mu_{roll} - \mu_x(x_3)\operatorname{sgn}(x_1)\right)$$
$$-\frac{r_e^2 F_N}{I_w B}\mu_x(x_3) + \frac{1}{B^2}\left(x_1|x_1| - r_e x_2|x_1| + x_1^2 x_3\right)$$
$$\beta(x) = \frac{r_e}{I_w B}$$
$$\gamma(\xi,\eta) = \frac{F_N}{m}(-\mu_{roll}\operatorname{sgn}(\eta) + \mu_x(\xi_1))$$

It is relatively straightforward to cancel out the nonlinear dynamics in ξ_2. From this, the input-output linearization control law is synthesized.

$$u = \frac{1}{\beta(x)}(-\alpha(x) + v) = g_1(x,v) \tag{6.10}$$

What remains is to generate a control law for v which is capable of controlling the linear equivalent system in an acceptable manner. One such suitable controller has the form:

$$v = \ddot{y}_d - w_2(\xi_1 - y_d) - w_1(\xi_2 - \dot{y}_d) = \ddot{y}_d - w_2 e - w_1\dot{e} \tag{6.11}$$

where w_1 and w_2 are gains chosen to make the controller stable. This control law focuses on reducing the error in the output y and its derivative. Furthermore, if w_1 and w_2 are chosen such that $s^2 + w_1 s + w_2$ is Hurwitz, then this controller is stable, that is:

$$e(t) \to 0 \quad \text{as} \quad t \to \infty\ .$$

For this work, w_1 and w_2 were taken to be 20 and 100, respectively (which gives a Hurwitz controller).

With the application of this control law, the dynamics of the system have taken the form:

$$\dot{\xi} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \xi + \begin{bmatrix} 0 \\ 1 \end{bmatrix} v \tag{6.12}$$
$$\dot{\eta} = \gamma(\xi, \eta)$$

And with v selected as in Equation (6.11), the external state dynamics are stable. The internal state η still has some dynamics of interest. Even though the external states are stable, it is possible that the internal state could become unstable, resulting in an unstable system. Therefore, it must be verified that internal dynamics are stable. In order to do this, it is convenient to start with the zero dynamics:

$$\dot{\eta} = \gamma(0, \eta)$$
$$\gamma(0, \eta) = \frac{F_N}{m}(-\mu_{roll}\,\mathrm{sgn}(\eta) + \mu_x(0)) = -\frac{F_N}{m}\mu_{roll}\,\mathrm{sgn}(\eta)$$

This is stable for all η (with the exception of $\eta = 0$, where the sgn function is ill-defined). Therefore, internal dynamics are stable, and the overall controlled system should be stable, as well.

The feedback linearization approach has some major advantages. First, the simplicity of the resulting equivalent LTI system allows for the control signal v to be designed like any linear feedback controller. This means that the controller performs well and is relatively simple to change. However, there are also some underlying assumptions which could cause trouble in implementation. Perhaps the most important assumption is that the nonlinear dynamics of the system are perfectly captured in the state-space formulation. This is almost certainly not true for most systems, and with nonlinear systems, such model uncertainties can cause significant error.

6.1.1 Simulation Results

The simulations run for the nonlinear controllers were different from the setup described in Section 3.7 for the other controllers. Instead of a full simulation involving a wheel loader pushing against a resistive force, the nonlinear controllers were simply tested to see how well they could track a wheel slip trajectory. This trajectory included a brief positive-slip acceleration phase followed by a negative-slip deceleration phase (Fig. 6.2(a)). This shows the capability of the controller to perform under dynamic conditions.

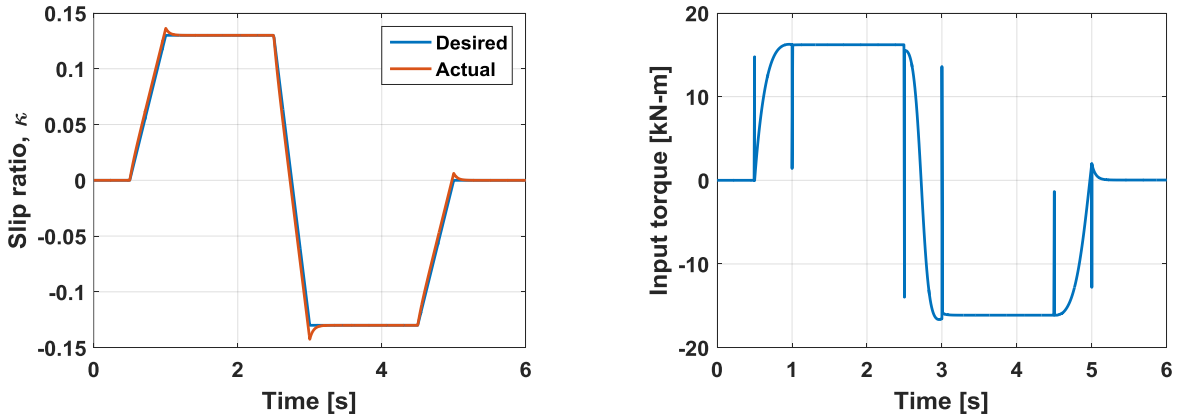

Figure 6.2. Simulation results for Feedback Linearization controller, showing (a) wheel slip ratio and (b) commanded input torque.

Figure 6.2(a) shows that the reference trajectory is followed quite closely by the system, with a maximum error value over the entire simulation of around 0.01 (less than 10%). However, Fig. 6.2(b) shows the primary drawback of the feedback linearization controller. In this case, the commanded torque values necessary to achieve such a performance have sharp spikes which are unachievable by the real-world system. Despite this, the feedback linearization controller has a reasonably good torque command profile, and it tracks the reference trajectory well for a system where the model parameters are accurately known.

6.2 Sliding Mode Control

Sliding mode control (SMC) is a control formulation designed to deal with bounded uncertainties in the model parameters. As it happens, this consideration is particularly important for the application at hand, since some parameters could change throughout the course of operation and others are very difficult to measure to an appropriate degree of accuracy. Due to this, sliding mode control has already been successfully applied to similar systems [61, 62].

Taking into account the state-space model (6.1) of the system, and considering parameter uncertainties (represented here as values θ_i), the system model becomes:

$$\begin{aligned}\dot{x}_1 =& \frac{F_N}{\theta_2}(-\theta_1 \operatorname{sgn}(x_1) + \mu_x(x_3)) \\ \dot{x}_2 =& \frac{1}{I_w}(-r_e F_N \mu_x(x_3) + u) \\ \dot{x}_3 =& \theta_3(-x_1 + r_e x_2 - |x_1|\, x_3),\end{aligned} \tag{6.13}$$

where $\theta_1 = \mu_{roll}$, $\theta_2 = m$, $\theta_3 = 1/B$. Whereas the dynamic radius and moment of inertia of the wheel (r_e and I_w, respectively) are relatively straightforward to determine, the coefficient of resistance (μ_{roll}) and relaxation length of the tire (B) are much more difficult, and it is possible that they could change depending on certain conditions. Also, as the wheel loader is a load-handling machine, its mass changes with changing load. Therefore, it is convenient to represent the mass m of the machine as an uncertain parameter.

One caveat of this design is that the while the parameters are uncertain, they must be bounded. Therefore, each parameter needs limits on its values. Fortunately for this system, there are some acceptable ranges these values could take. Reasonable limits for the parameters are [0.01, 0.09], [3200, 4500], [6.25, 100] for θ_1, θ_2, and θ_3, respectively.

As sliding mode control needs a structure similar to the feedback linearization form, it was decided to utilize the same coordinate transformation as above.

$$\dot{\xi}_1 = \xi_2$$

$$\dot{\xi}_2 = \frac{F_N\theta_3}{\theta_2}\left(\theta_1 \operatorname{sgn}(x_1) - \mu_x(x_3)\right) + \frac{F_N\theta_3}{\theta_2}x_3\left(\theta_1 - \mu_x(x_3)\operatorname{sgn}(x_1)\right) - \frac{r_e^2 F_N}{I_w}\theta_3\mu_x(x_3)$$
$$+\theta_3^2\left(x_1|x_1| - r_e x_2|x_1| + x_1^2 x_3\right) + \frac{r_e}{I_w}\theta_3 u = \varphi^T(x)\hat{\theta} + k\theta_3 u$$

$$\dot{\eta} = \dot{x}_1 = \frac{F_N}{\theta_2}(-\theta_1 \operatorname{sgn}(x_1) + \mu_x(x_3))$$

As above, the internal states are guaranteed to be stable, as long as θ_1 and θ_3 are greater than zero. Therefore, the controller can be designed as an input to cancel out the dynamics of ξ_2.

The sliding mode controller is designed such that the system states reach a desired surface, which is related to the state errors. For this system, the target surface and its dynamics are defined as follows:

$$s = \tilde{\xi}_2 + w\tilde{\xi}_1 \tag{6.14}$$

$$\rightarrow \dot{s} = \dot{\tilde{\xi}}_2 + w\dot{\tilde{\xi}}_1 = \varphi^T(x)\theta + k\theta_3 u - \ddot{y}_d + w\tilde{\xi}_2 \ . \tag{6.15}$$

The constant w in these equations is simply a weighting constant. A value of $w = 100$ was used for this work. Of course, the idea of the controller is to reduce the state errors to zero. Therefore, the sliding mode control structure is defined as:

$$u = -\frac{1}{k\hat{\theta}_3}(u_a + u_s) \ , \tag{6.16}$$

where u_a and u_s are the model compensation and sliding mode components of the control signal, respectively.

As its name suggests, the model compensation control signal seeks to cancel out the nonlinearities present in the error dynamics of the system. The actual parameter values θ are not known, so an estimate of those parameters $\hat{\theta}$ is used instead.

$$u_a = -\varphi^T(x)\hat{\theta} + \ddot{y}_d - w\tilde{\xi}_2 \tag{6.17}$$

$$\tilde{\xi}_2 = \xi_2 - \dot{x}_{3d} = \dot{x}_3 - \dot{x}_{3d} \tag{6.18}$$

The values of $\hat{\theta}$ are typically chosen as the average of the lower and upper limits to minimize the potential error in their values. This value selection method was done for this project, with the exception of θ_3, which was chosen as 11.765 (the reciprocal of the average of the limits on the relaxation length). It is possible, however, that one might use anther value within the range for the parameter estimate, for instance if the probability distribution of the parameter value is not centered on the range.

Having applied the model compensation control to the system, the resulting error dynamics come out to be:

$$\rightarrow \dot{s} = -\varphi^T(x)\tilde{\theta} + u_s \ , \tag{6.19}$$

where $\tilde{\theta}$ is the error in the parameter estimates. From here, the sliding mode control law us must be synthesized in order to ensure that the error dynamics are stable (i.e. that the error converges to zero).

6.2.1 Ideal Sliding Mode Control Law

The first candidate for this control law is known as the *ideal sliding mode controller*, and its structure is as follows:

$$u_s = -(h(x,t) + h_0)\,\mathrm{sgn}(s) \ , \tag{6.20}$$

where

$$h(x,t) \geq \left| -\varphi^T(x)\tilde{\theta} \right| \ . \tag{6.21}$$

The function $h(x,t)$ is chosen such that it is always greater than the parameter estimation error of the system. For this project, it was taken as a constant value greater than or equal

to the maximum effect of the parameter estimation error (this value was calculated to be 362.8702). The error dynamics $\dot{s}$ are thus negative-definite, and it is seen that the addition of constant h_0 provides a margin on this negative-definite guarantee. Noting that the following inequality holds globally for the error dynamics:

$$s\dot{s} \leq -h_0 \left| s \right| \ , \tag{6.22}$$

it is seen by Lyapunov analysis that the state error reaches zero in a finite time and stays at zero. For the simulation results shown below, h_0 was taken to be 100.

The ideal sliding mode control law suffers from controller chattering (i.e. rapid switching between two levels) due to the sgn function discontinuity in the formulation of u_s. This is easily seen in the simulated control signal for this controller. It is likely that this control action is impossible for the system to achieve. In fact, for the current system, the braking hydraulic dynamics are much slower than this desired action.

6.2.2 Smoothed Sliding Mode Control Law

Therefore, another controller is proposed which removes the discontinuity in the ideal controller law. This structure is called the *smoothed sliding mode controller*, and it uses a saturation block to approximate the action of the sgn function.

$$u_s = -(h(x,t) + h_0) \, \text{sat} \left(\frac{s}{\sigma} \right) \tag{6.23}$$

In this formulation, σ is known as the *boundary layer thickness*, and it controls how rapidly the output of the saturation block changes from -1 to 1 with a varying input. For this project, the values of $h(x,t)$ and h_0 were taken to be the same for the smooth SMC as for the ideal SMC. Of course, the guarantee of finite-time trajectory tracking is lessened by this approximation. However, the tradeoff in terms of control signal shape is significant, as shown in the simulation results below.

The sliding mode controller is an excellent way to handle parameter uncertainties, but the problem still persists that SMC requires some knowledge of the uncertain parameters

to create the reasonable parameter ranges *a priori*. Therefore, it could be useful to utilize a controller which is capable of updating the parameter estimates in real time. This is the justification for the Adaptive Controller examined in Section 6.3.

6.2.3 Simulation Results

The simulations run for the sliding mode control were the same as those used to demonstrate the feedback linearization controller performance. Table 6.1 shows that in the simulation, the estimated values of the three uncertain parameters were different from their actual values (but still within the bounds used in creating the controller). These values were used for simulating both versions of the sliding mode controller (ideal and smoothed). Therefore, this simulation also examines the controller's capability to deal with uncertainties in the system, something of which feedback linearization is incapable.

Table 6.1. Comparison of actual and estimated parameter values used in Sliding Mode controller simulations.

Parameter	**Actual Value θ_i**	**Estimated Value $\hat{\theta}_i$**
Resistance coefficient, μ_{roll}	0.03	0.05
Vehicle mass, m	4000 kg	3850 kg
Tire relaxation length, B	0.05 m	0.085 m

Figure 6.3 shows the result of the sliding mode controller simulation for the ideal control law. This controller performs quite well, in terms of meeting the desired state trajectories. It can be seen in Fig. 6.3(a) that there is in fact some steady-state error at both levels. This matches with theory, as sliding mode does not guarantee zero steady-state error, only bounded error for a bounded input. The major drawback of this control structure is seen in the control input torque plot (Fig. 6.3(b)). It contains copious amounts of chattering, which a real system would not be able to realize. Furthermore, the large signals (around 350 kN·m) are likely outside of the operable range for the system. Therefore, this control structure is likely not very practical for real-world systems.

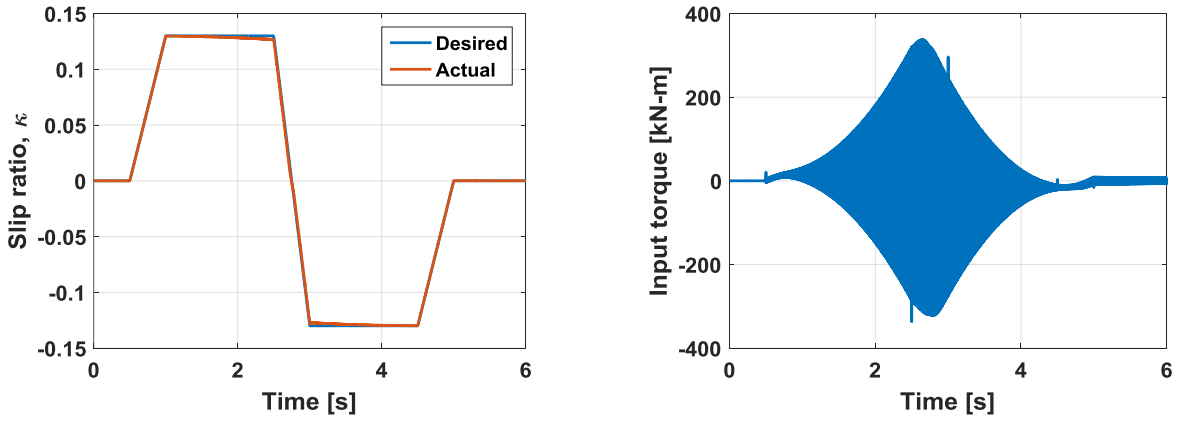

(a) (b)

Figure 6.3. Simulation results for Ideal Sliding Mode controller, showing (a) wheel slip ratio and (b) commanded input torque.

The smoothed sliding mode control, on the other hand, is much better behaved. The severe chattering has been basically eliminated, and the command torque has returned to much more reasonable values (Fig. 6.4(b)), which are comparable to the results from the feedback linearization controller (which had zero parameter error). The slip ratio tracking (Fig. 6.4(a)) is roughly as good as the ideal sliding mode case, although it does sacrifice some precision to achieve the better-conditioned control signal (so equivalent performance to the ideal case is not guaranteed).

All in all, the sliding mode control is a suitable solution in this case for a system which has uncertain model parameters. The trajectory tracking performance is comparable to that of the feedback linearization control, and with the smoothed version of the control, it achieves similar command actions. All of this is achieved with three parameter estimates which are significantly different from the actual values used in the simulation. Because of these distinct advantages, the smoothed sliding mode control was actually implemented on the real-world prototype system, and the results are discussed in Section 8.4 of this work.

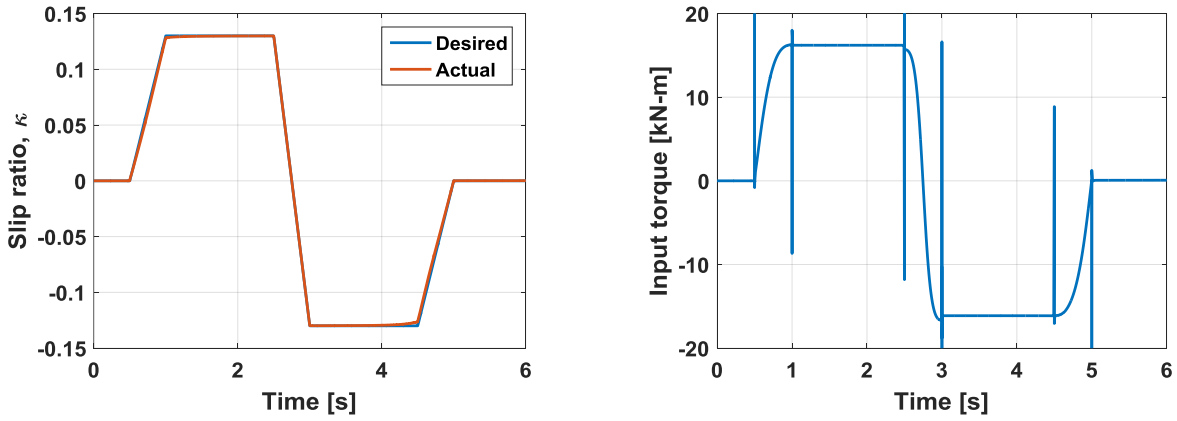

(a) (b)

Figure 6.4. Simulation results for Smoothed Sliding Mode controller, showing (a) wheel slip ratio and (b) commanded input torque.

6.3 Adaptive Control

The final control structure analyzed in this work is an adaptive control law. This control methodology, like sliding mode control, is specially designed to work with systems that have uncertainties in their model parameters. It has also been shown as suitable for use in traction control applications [63,64]. In this case, the control structure is actually designed to change the controller model parameters in real time. Given enough time, with a proper system structure, an adaptive controller will modify its model compensation control to correctly cancel out the nonlinearities of the system dynamics. A typical linear feedback control law is then used to stabilize and control the system. The general structure of such a controller is shown in Fig. 6.5.

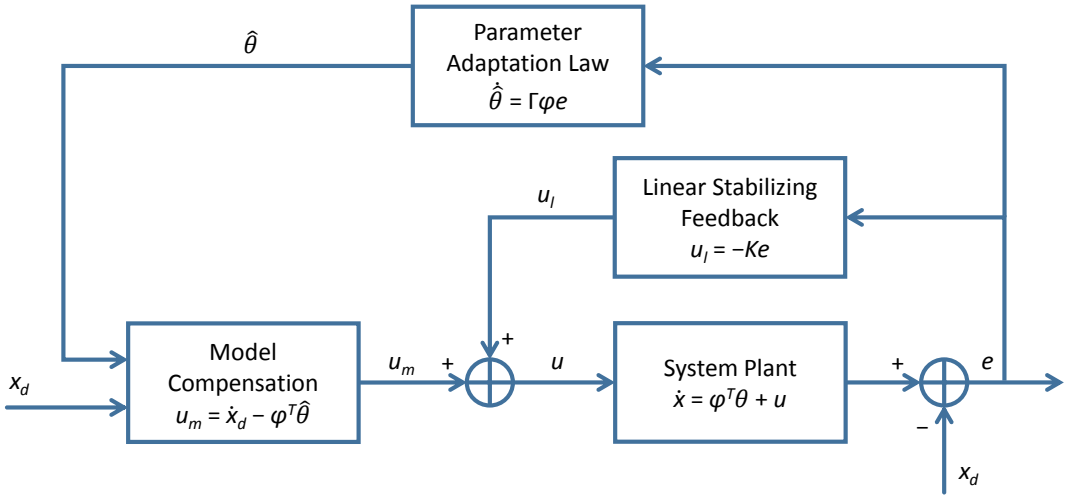

Figure 6.5. Adaptive control schematic for a system with parameter uncertainties.

Much like for the SMC, it was decided that the simplest way to go about designing the controller for this system was to utilize the transformed coordinates of the feedback linearization design.

$$\dot{\xi}_1 = \xi_2$$

$$\dot{\xi}_2 = \frac{F_N\theta_3}{\theta_2}\left(\theta_1\,\mathrm{sgn}(x_1) - \mu_x(x_3)\right) + \frac{F_N\theta_3}{\theta_2}x_3\left(\theta_1 - \mu_x(x_3)\,\mathrm{sgn}(x_1)\right) - \frac{r_e^2F_N}{I_w}\theta_3\mu_x(x_3)$$

$$+\theta_3^2\left(x_1\,|x_1| - r_e x_2\,|x_1| + x_1^2x_3\right) + \frac{r_e}{I_w}\theta_3 u = \varphi^T(x)\hat{\theta} + k\theta_3 u$$

$$\dot{\eta} = \dot{x}_1 = \frac{F_N}{\theta_2}(-\theta_1\,\mathrm{sgn}(x_1) + \mu_x(x_3))$$

From this modified system, it is possible to generate φ and θ for the purposes of constructing the parameter adaptation law.

$$\varphi = \begin{bmatrix} (\mathrm{sgn}(x_1) + x_3) \\ -\mu_x(x_3)(1 + x_3\,\mathrm{sgn}(x_1)) \\ -\frac{r_e^2F_N}{I_w}\mu_x(x_3) \\ x_1|x_1| - r_e x_2\,|x_1| + x_1^2x_3 \end{bmatrix}$$

$$\theta = \begin{bmatrix} \frac{\theta_1 \theta_3}{\theta_2} \\ \frac{\theta_3}{\theta_2} \\ \theta_3 \\ \theta_3^2 \end{bmatrix}$$

Choosing Γ as an appropriate positive-definite matrix, the parameter adaptation law can be synthesized as:

$$\dot{\hat{\theta}} = \Gamma \varphi e \ , \tag{6.24}$$

where e is the system error. This adaptation law controls the development of the parameter estimates in such a way that they converge to the actual parameter values. Essentially, $\hat{\theta}$ becomes a new state vector for the system which must also be considered when assessing system stability.

If the parameter estimates converge to their correct values, then the model compensation control will adequately cancel the nonlinearities in the model dynamics, meaning that the system can be controlled with a simple linear feedback law. That is to say, the control input u to the system is taken as:

$$u = u_l + u_m \ , \tag{6.25}$$

where u_l is the linear stabilizing feedback law, in this case:

$$u_l = -Ke \ , \tag{6.26}$$

and u_m is the model compensation control law seen in the previous sections.

$$u_m = -\varphi^T(x)\hat{\theta} + \ddot{y}_d - w\tilde{\xi}_2$$

$$\tilde{\xi}_2 = \xi_2 - \dot{x}_{3d} = \dot{x}_3 - \dot{x}_{3d}$$

From this control law, the dynamics of the system error e can be calculated as:

$$\rightarrow \dot{e} = -\varphi^T(x)\tilde{\theta} - Ke \ . \tag{6.27}$$

In order to assess the stability of the control law, the error dynamics are analyzed using the Lyapunov function:

$$V_a = \frac{1}{2}e^2 + \frac{1}{2}\tilde{\theta}^T\Gamma^{-1}\tilde{\theta} \ , \tag{6.28}$$

which has derivative:

$$\dot{V}_a = e(-\tilde{\theta}^T\varphi(x) - Ke) + \tilde{\theta}^T\Gamma^{-1}\dot{\hat{\theta}} = -Ke^2 \leq 0 \ . \tag{6.29}$$

As this Lyapunov derivative is negative semi-definite, all signals in the system are bounded. It can be shown that for this system, e asymptotically converges to zero. This is an important consideration for the adaptation law to be considered valid.

The adaptive control law is indeed very useful for a system like the one presented in this paper. However, it does have some limitations as well. Perhaps the most important limitation is that it assumes there are no uncertain nonlinearities in the system (i.e. $\Delta = 0$). If this is not the case, asymptotic tracking is not guaranteed. In such systems, it is necessary to use a form of (probably nonlinear) robust control. If needed, robust control can be combined with the adaptive control seen here to produce an adaptive robust control.

6.3.1 Simulation Results

As with the sliding mode controller simulations above, the Adaptive controller was simulated starting with an initial guess for the uncertain parameters which was incorrect. However, the adaptive control law allows the system to update these values in real-time, so they change throughout the simulation to improve system performance. Table 6.2 shows the initial values $\hat{\theta}_{0,i}$ used for the simulation.

The first condition which was simulated for the adaptive controller was the same as the feedback linearization and sliding mode controllers (Fig. 6.6). The primary purpose of this test is to assess the dynamic tracking capabilities of the controller. There are several conclusions which can be drawn from these plots. First, the adaptive control has a relatively poor transient response (Fig. 6.6(a)). This is due to the fact that there is some error in the parameter estimates throughout the simulation. These poor parameter estimates keep

Table 6.2. Comparison of actual and starting parameter estimate values used in Adaptive controller simulations.

Parameter	**Actual Value θ_i**	**Initial Estimate Value $\hat{\theta}_{0,i}$**
Resistance coefficient, μ_{roll}	0.03	0.05
Vehicle mass, m	4000 kg	3850 kg
Tire relaxation length, B	0.05 m	0.085 m

the model compensation control signal from doing its job properly. With a better tuned adaptation law matrix Γ and perhaps a different linear feedback gain K, this system would perform better. Nevertheless, zero steady-state error is indeed achieved in this case. This result is guaranteed by the adaptive controller. The input torque command for the system is well in line with the responses seen in the feedback linearization and smoothed sliding mode control cases. Therefore, the system responds reasonably well, even starting from parameter estimate values which are incorrect.

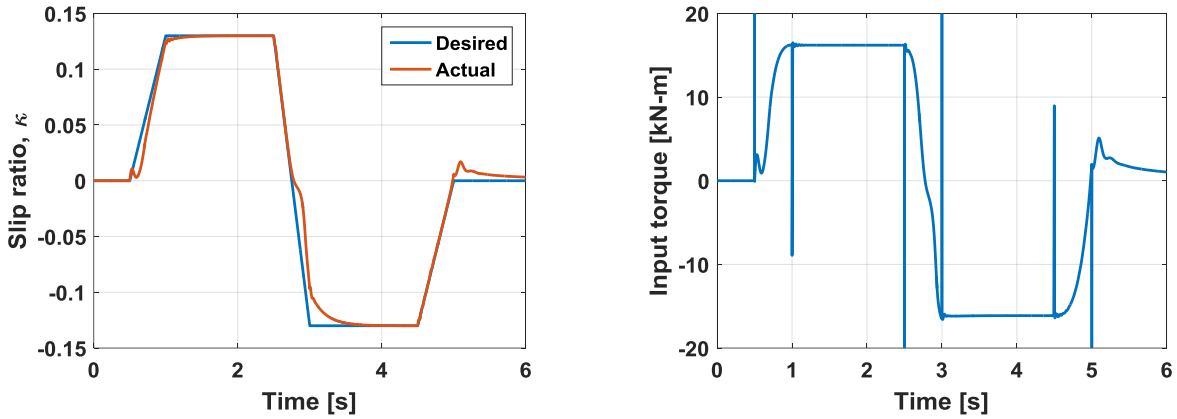

(a) (b)

Figure 6.6. Simulation results for Adaptive controller, showing (a) wheel slip ratio and (b) commanded input torque.

One thing which should be noted with respect to the adaptive control in this first simulation is that, while the parameter estimate values do change, they do not necessarily converge

to the correct values. There is no general guarantee of the adaptive control parameters converging to the correct value. Instead, in order for the parameters to converge properly, the system input must persistently excite the system [65]. Therefore, another simulation was run with a sinusoidal reference trajectory (Fig. 6.7), instead of the sudden, abrupt reference changes seen in the previous simulation.

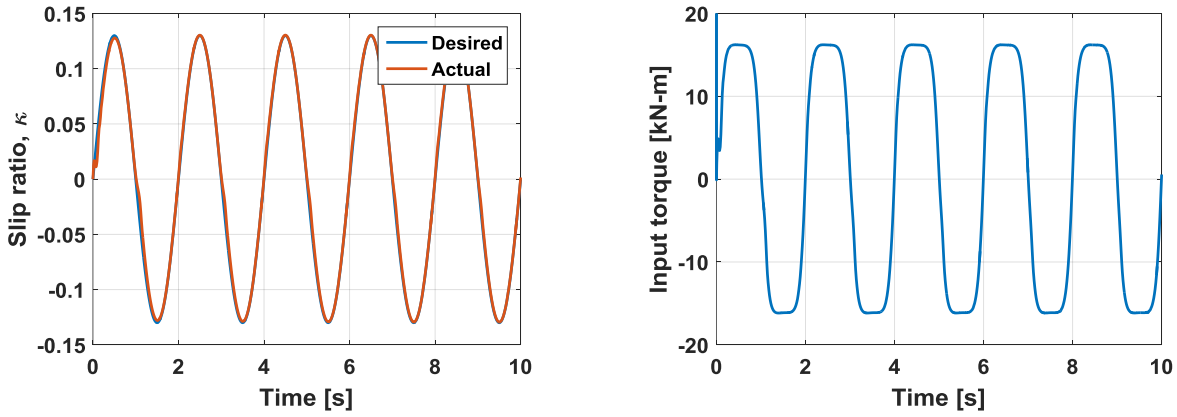

(a) (b)

Figure 6.7. Simulation results for Adaptive controller with sinusoidal reference trajectory, showing (a) wheel slip ratio and (b) commanded input torque.

For this simulation, the results are consistent with what was seen previously. There is some small error at the beginning of the simulation, but it follows the sinusoidal reference slip quite well. Furthermore, the slip ratio error is decreased over time, as the parameter estimates achieve better values. The commanded input torque is also quite reasonable.

Figure 6.8 shows the parameter estimate values over the course of the sinusoidal input simulation. In this case, only one of the two estimates converges to the proper value (the vehicle mass in Fig. 6.8(b)). The coefficient of rolling resistance (Fig. 6.8(a)) finds the correct value initially, but then diverges to a different point. Finally, the estimate of the tire relaxation length in Fig. 6.8(c) does not change much. It begins converging in the correct direction, but it does not approach the correct value during this simulation.

There are a few different reasons why only one parameter estimate was able to converge correctly. First, it is possible that with more time, the parameter estimates will indeed find

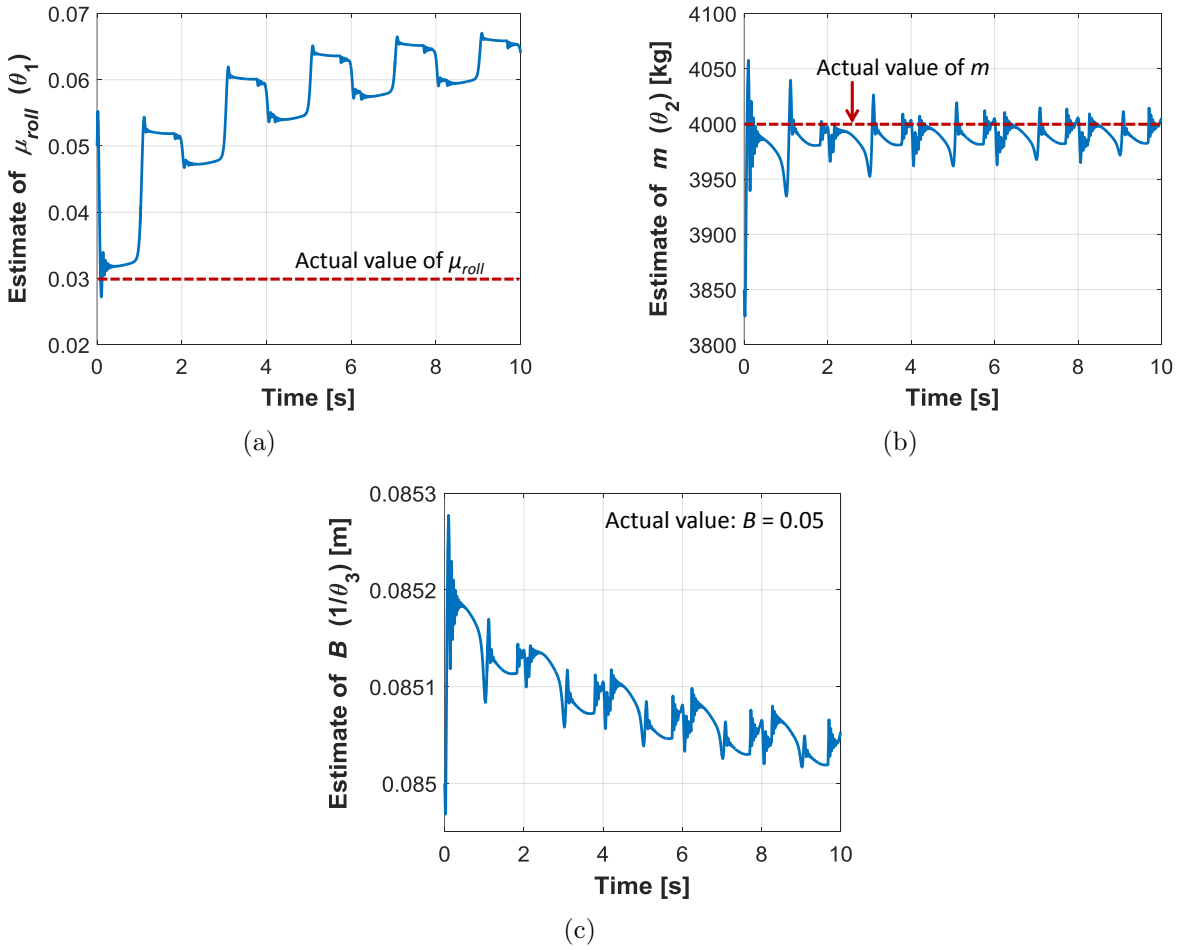

Figure 6.8. Simulation results of Adaptive controller parameter estimate convergence for (a) coefficient of rolling resistance, (b) vehicle mass, and (c) tire relaxation length.

their correct values. However, it is more likely that the adaptation law matrix Γ has been chosen in such a way that the parameter estimate convergence is poor. With the proper value for Γ and a persistent excitation, better estimate convergence should be achievable.

7. EXPERIMENTAL SETUP

Having developed the various appropriate control schemes for the system, the next step of the process was to implement the controller on the prototype machine to test its feasibility for real-world applications. Chapters 8 and 9 describe the different categories of tests which were run over the course of this research. These tests comprised a wide array of investigations which were intended to measure many different values. Therefore, the reference machine was modified in several ways to create the new functions which were needed to run the tests, as well as adding instrumentation for acquiring data.

This chapter outlines the modifications made to the reference wheel loader to prepare it for the tests described in later chapters. All of the machine updates were done at Purdue University in collaboration with an industry sponsor, starting from an essentially unmodified stock vehicle.

7.1 Electro-Hydraulic Braking Circuit

One of the primary modifications to be made to the reference machine was altering the braking system so that it could be controlled electronically. As mentioned in Section 1.2, the modifications allow for increasing the braking pressure above that commanded by the foot valve, but they do not allow for decreasing it. What this means is that, while this system is capable of implementing a traction control system with independent braking at each wheel, it is unable to work as an anti-lock braking system (ABS). To allow the system to function as an ABS, it must be modified to be more like the system presented in [4], which can both raise and lower the pressure at each brake. However, ABS functionality is not within the scope of this work. As the details of the EH braking system are protected by a confidentiality agreement, only a general description of the system could be included here.

A computer system, running LabVIEW development software, was connected to valve drivers which operated pressure reducing valves at each wheel (Fig. 7.1). Each pressure reducing valve is connected to one inlet of a shuttle valve, and the other inlet is connected directly to the pressure set by the braking foot valve. The outlet of the shuttle valve is

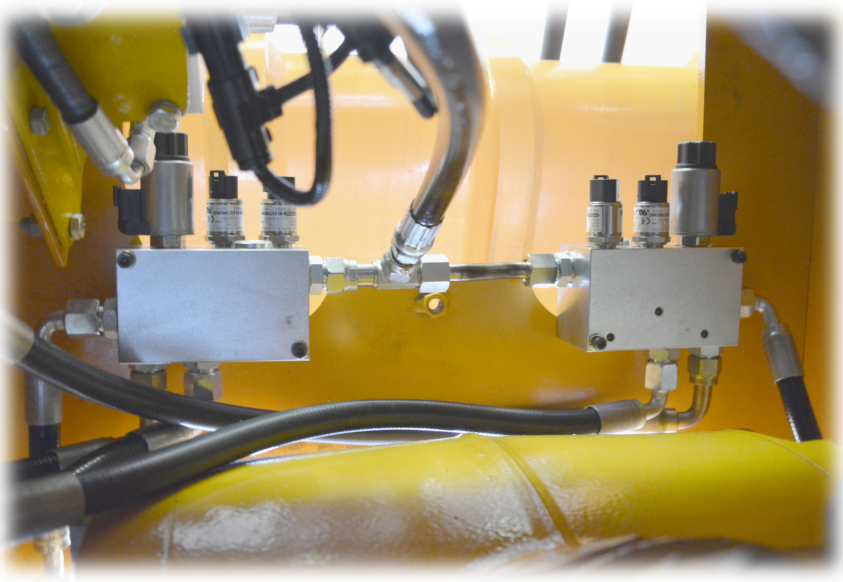

Figure 7.1. Electro-hydraulic braking valves installed on reference wheel loader.

connected to the brake cylinder itself, meaning that the brake cylinder pressure is equal to whichever pressure is higher (between the pressure-reducing valve and the foot valve). The supply to the pressure reducing valves at each wheel comes from the same pressure source as the braking foot valve, which can be shut off from the EH braking circuit by use of an on-off valve. In the case that the supply pressure is not connected to the EH circuit, the vehicle's braking behavior is essentially the same as it would be without the system modifications.

7.2 Instrumentation

This section enumerates the various sensors installed on the wheel loader system. First, the sensors used in the development stage of the process are described. Each sensor has a crucial function in some part of the research activity; however, the system is designed such that not all of these sensors are required for the final system implementation. It was determined early on that the number of additional sensors needed for the system (i.e. on top of the various sensors already installed on the stock machine) should be minimized.

Therefore, the final TC system setup includes only wheel speed sensors and a steering angle sensor.

7.2.1 Development Setup

Table 7.1. List of sensors for development testing.

Sensor	Qty.	Description
Ground speed	1	Doppler effect microwave sensor, frequency output
Wheel speed	4	Inductive proximity sensors with steel polewheels, frequency output
Accelerometer	1	3-axis accelerometer with 2-axis gyroscope, analog voltage output
Brake supply pressure	2	250 bar measuring range, analog voltage output
Braking pressure	8	100 bar measuring range, analog voltage output
Boom and bucket work pressure	4	400 bar measuring range, analog voltage output
Boom and bellcrank angle	2	180° measuring range, analog voltage output
Steering angle	1	180° measuring range, analog voltage output
Brake temperature	4	T-type thermocouples (−270 to 400°C measuring range)

Table 7.1 lists all of the sensors used on the development prototype along with brief descriptions of their function. The tire slip velocity used as the TC system controller process variable is calculated by subtracting the linear velocity of the wheel loader from the theoretical linear velocity of each of the wheels. Therefore, the wheel velocities and wheel loader speed need to be known. In order to obtain a measurement of vehicle speed, a Doppler effect ground speed sensor was installed on the vehicle (Fig. 7.2). However, this particular sensor has some known issues erroneous measurements under certain operating conditions. Furthermore, the cost of this sensor was problematic for including it in the final system

implementation. Because of this, a velocity estimator based on wheel speeds (see Section 7.3) was instead used for wheel slip calculations, with the ground speed sensor used only to verify the estimated speed values.

Figure 7.2. Ground speed sensor mounted on reference machine.

For the correct calculation of tire slip velocities at each wheel, a set of accurate high-resolution wheel speed sensors was needed. The prototype solution for determining wheel speed was to weld steel encoding wheels on the inside edge of the wheel rims (Fig. 7.3). These encoding wheels had protruding teeth, which generated a digital frequency signal when rotated past an inductive proximity sensor. These worked quite well, and were capable of accurately measuring the wheel speeds.

There are some important aspects to consider when designing a frequency-based digital sensor. The primary method of measuring frequency is to mark falling or rising edges (or both) and measure the time between each consecutive edge. The resulting time interval gives the average frequency of the signal in that time span. Due to this principle of operation, the following factors are important. First, the encoding wheel must generate enough pulses to give a meaningful signal at the lowest desired wheel speed. Frequency-based sensors do not have a resolution in the standard sense, but they do have a *spatial resolution*, which for this application can be thought of as the number of "pulses per revolution" for the signal

encoding wheel. As a zero-speed signal would simply cause the system never to update, the general solution is to set a time interval Δt_{zero} after which the frequency will be assumed zero if no new pulses have occurred. Therefore, the minimum spatial resolution n_{pulse} for the polewheel can be computed as follows

$$n_{pulse} \geq \frac{2\pi}{\Delta t_{zero}\omega_{min}} \ . \tag{7.1}$$

where ω_{min} is the desired minimum rotational velocity (in rad/s) for the sensor to read without reading zero intermittently. That is to say, the encoder must generate enough pulses per revolution to allow for an acceptably small zero-speed wait time at minimum measurable speed.

Conversely, the system which is reading the frequency signal is necessarily limited in its maximum sampling rate f_s. Therefore, another constraint is generated when considering the maximum wheel velocity. Generally, this means the maximum signal content of the measurement should be less than the Nyquist frequency of the acquisition system ($f_{Nyq} = 0.5f_s$) to avoid aliasing. This consideration places a limit on the maximum value of the polewheel spatial frequency.

$$n_{pulse} < \frac{2\pi f_{Nyq}}{\omega_{max}} \ , \tag{7.2}$$

where ω_{max} is the maximum desired speed of interest for the wheel. This represents the maximum possible value of the spatial resolution. However, it is best practice to have the sampling frequency at least six to eight times the maximum signal frequency, which would correspond to scaling the maximum resolution down by a factor of at least three to four.

The inductive speed sensors were adequate for many tests; however, for others, the test conditions allowed for implementation of a system which had a better signal behavior at very low velocities. Therefore, in these tests the inductive wheel speed sensors were switched to optical sensors capable of achieving significantly higher spatial resolution (Fig. 7.4). The encoding wheels for the optical sensors (a series of black and white stripes) were printed on magnet-backed vinyl, and they were attached to the steel encoding wheels magnetically.

This optical speed measurement is ideal for laboratory applications, where there are no external factors such as flying debris or excessive dust that might affect the measurement

Figure 7.3. Inductive wheel speed sensor with encoding wheel mounted on vehicle wheel rim.

of the sensor. However, in field tests, the optical sensor sees severe noise problems due to dust and vibrations. In the final production vehicle implementation, due to the fact that the machine is designed for heavy duty work, another type of sensor is needed which is more robust. Ideally, the wheel speed sensors will be located inside the wheel axle housings, like they are in on-road automotive wheel bearings.

One of the first experiments run with the reference wheel loader was a series of measurements to characterize the slip-friction characteristics of the vehicle wheels, as mentioned in Section 2.3.2. These tests were done using a method suggested by Rajamani [66], wherein the sensors required for the tire force estimation at each wheel are vehicle and wheel speed sensors, along with longitudinal acceleration data. Therefore, an accelerometer was mounted at a central location on the vehicle, at a height close to that of the wheel axles. This accelerometer was not used for controls purposes, and it does not appear in the final system instrumentation.

A very important part of the project was generating a suitable comparison between tests where traction control was implemented and tests without control. To do this, a test barrier

Figure 7.4. Optical wheel speed sensor with encoding wheel mounted on vehicle wheel rim.

was made and several tests were conducted (see Chapter 8 for more information). Each test was conducted by approaching the tires on the barrier and pushing against them. The angles of the boom and bucket are required for experimental testing purposes in order to ensure similar testing conditions (that is, keeping the orientation of the wheel loader implement constant among tests). To this end, angle position sensors were installed on the machine to measure the boom position (Fig. 7.5(a)) and bellcrank position (Fig. 7.5(b)). These signals were also used, along with pressure measurements at the hydraulic cylinders that control boom lift and bucket tilt, in calculating an estimate of the machine pushing force to compare different system performances (see Section 3.8).

Even a relatively small steering maneuver can cause the wheels on the outside of the steering arc to have a noticeably higher velocity than the wheels on the inside. Without any compensation for steering angle, the controller will read that the outside wheels are slipping and apply the brakes accordingly. This can cause the system to seize up and block the outside tires unnecessarily, keeping the wheel loader from advancing. To overcome this issue, steering compensation was added to the slip estimation algorithm. Based on the

(a) (b)

Figure 7.5. Angle sensors installed at (a) boom and (b) bellcrank pins.

steering angle and machine geometry, the system estimates the inner and outer radii of the wheel loader travel path. Using this information, the center-line ground velocity is scaled appropriately to compare the wheel speeds. In order to determine the steering angle, another angular sensor had to be mounted on the machine. The sensor used is of the same type as the boom and bellcrank angle sensors. More on the steering angle compensation is explained in Section 4.5.

In addition to the velocity, wheel speed, and angular position sensors already discussed, several pressure sensors were installed to monitor hydraulic pressures in specific points on the system. There are two braking supply control sensors mounted, as there are two separate hydraulic pressure supply subsystems for the electro-hydraulic braking circuit. The implementation of these sensors is not necessary on the final iteration, as their function was simply to ensure that the braking systems were receiving the correct supply pressure. There are also a total of eight braking pressure sensors installed (two per braking module). These sensors were needed to verify that the pressure commanded corresponded with the command signal. They are mounted such that the pressure controlled by the pressure reducing valves and the pressure at the braking disks are measured independently (in case the disks see a higher pressure from the foot valve than from the electronic braking valves). The braking pressure sensors were also used to calibrate the command voltage to the valve drivers for controlling the pressure at each of the brakes. Unless some sort of feedback structure is

desired for controlling the braking pressure, these sensors will also not be required on the final implementation of this system.

Figure 7.6. Wheel hub thermocouple mounting.

Additional pressure sensors were installed on the wheel loader to measure the hydraulic pressure at either side of the main boom lift cylinders and the bucket tilt cylinder. The primary purpose of these sensors is to help generate the estimate of the pushing force of the wheel loader during the laboratory tests (see Chapter 8). Details of this pushing force estimate can be found in Section 3.8. Again, these pressure sensors are used only for monitoring and testing purposes, and will not be needed in the final system implementation.

Brake wear is an important factor which has not been mentioned heavily in this work. Traction control can decrease tire wear, but if the brakes are actuated too often or too harshly, the resulting wear to the brake discs themselves could be severe. It is not feasible to measure real-time brake wear for a system like the reference wheel loader. Instead, to gain some estimate of whether the brakes are being stressed excessively by the increase of braking action, it was decided to monitor the temperature inside the braking assembly. A T-type thermocouple was mounted inside the hub assembly of each wheel through the maintenance port (Fig. 7.6).

The thermocouple is installed so that its floating distance from the rotating brake disks d_{TC} is only 2 mm (Fig. 7.7). Therefore, the sensor does not directly measure the temperature of the braking discs. Instead, the sensor reads the ambient temperature inside the wheel hub assembly. The main purpose of the mounting of this sensor is to monitor the temperature

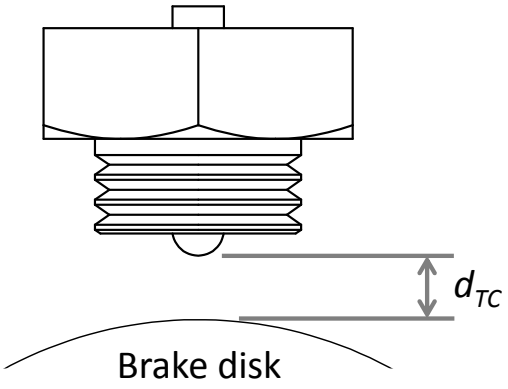

Figure 7.7. Diagram of thermocouple and brake disk.

in the hub in search for any drastic temperature change whenever the TC system is active. If a more thorough investigation of brake temperature is desired in the future, a different sensor setup will be required. However, in terms of the planned final implementation of the traction control system into a production machine, the braking temperature sensors will not be needed.

In addition to the sensors already described, the LabVIEW program which maintained the control and data acquisition code was connected to the wheel loader CAN bus in order to acquire other signals. These included machine velocity (redundant with the ground speed sensor and velocity estimate based on wheel speeds), ambient temperature/humidity, fuel rate, throttle position, engine speed, and system torques. The transmission gear and direction of travel were also monitored and used by the control system for resetting and disabling the controller at appropriate times. Figure 7.8 shows all of the instrumentation for the reference vehicle system, along with the approximate location of each sensor.

For the purposes of the productivity field tests conducted with the wheel loader (see Chapter 9), two other sensing systems were installed on the machine. The first is an on-board scale, which provides an estimate of the mass of material in the wheel loader bucket each time it is emptied. It does this by monitoring the boom cylinder pressures, boom position, and engine speed, among other things. The other component which was added to the wheel loader for these tests is a position and velocity measurement system. This system uses a GPS sensor, among others, to determine wheel loader longitude, latitude, altitude,

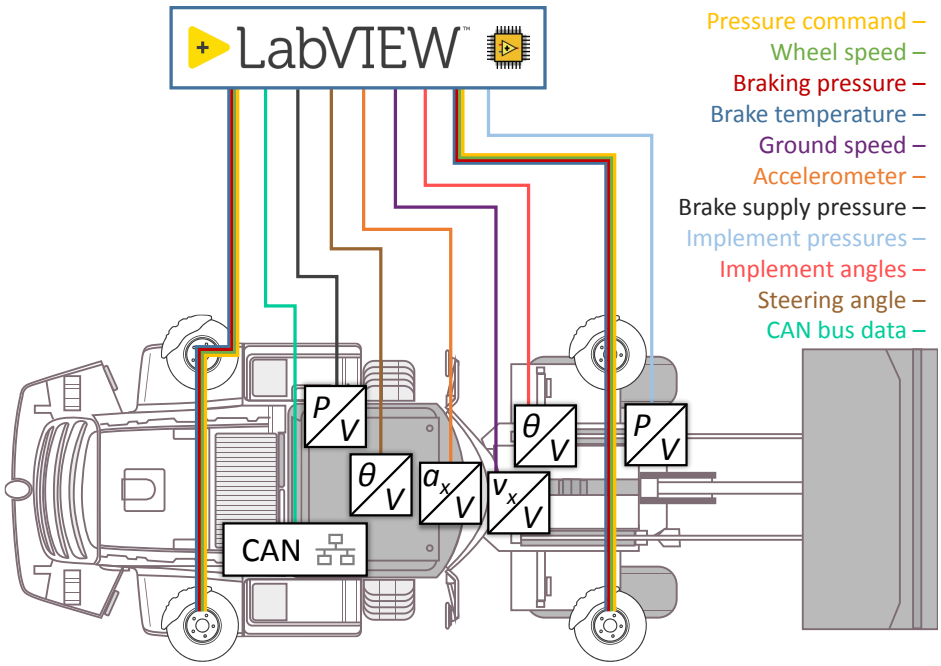

Figure 7.8. Sensor diagram for reference wheel loader.

and speeds, as well as a host of other parameters, most of which were not recorded for the tests. Both the load-weigh system and the GPS system were commercially available packages provided by the company sponsor, so they are not described in great detail here. As with most of the other sensors described in this chapter, these systems will not be necessary in the final machine TC implementation. They were simply used to generate more experimental data for the field tests, so that they could be compared in more detail.

7.2.2 Sensors Required for Control Functionality

In this section, a list of the sensors needed for operation of the traction control system is given. Table 7.2 lists all the important sensors that must be implemented, along with important specifications. A quick comparison between this and Table 7.1 shows the drastic reduction of sensors required for TC functionality. As mentioned above, one of the objectives of this work was to minimize the number of additional sensors needed for the system so

that sensor cost does not become a factor prohibiting the TC system from being put into production.

Table 7.2. Required sensor specifications for machine controller implementation.

Wheel Speed Sensors		
Parameter	**Value**	**Notes**
Quantity	4	Wheel speeds should be controlled independently
Sensor type	Frequency	Hall effect or other inductive sensor with metal polewheel interior to the wheel hub housing
Measuring range	0.1 – 30 km/h	Minimum achievable speed measurement is based on wait time before setting speed to 0, maximum is based on sampling frequency, both are affected by number of pulses per revolution (spatial resolution)
Encoder pulses/rev	> 160	Calculated from minimum desired speed (0.1 km/h); this could potentially be done with 80 "stripes," measuring both rising and falling edges
Steering Angle Sensor		
Parameter	**Value**	**Notes**
Quantity	1	Could be done with angular position sensor or by measuring extension of steering cylinders
Measuring range	-40 – 40°	Based on maximum steering angle of the wheel loader
Error	$< 2^\circ$	Small amount of error will not cause drastic problems, but significant error (or drift) could create severe issues

The choice of sensor installation is relatively straightforward. All four wheel velocities must be measured, as the system must be able to identify if any one wheel is slipping excessively against the ground. Due to the inclusion of a ground velocity estimate (see Section 7.3), there is no need for a linear velocity sensor. It is also assumed that the braking pressures can be controlled to a reasonable degree without additional pressure feedback. The boom and bellcrank angle sensors, boom and bucket cylinder pressure sensors, and accelerometer were all only used for experimental data generation, and are not required for controller implementation. Therefore, the only other sensor required for the TC system to function properly is the steering angle sensor.

7.3 Vehicle Velocity Estimation

The controller is designed to act whenever the wheel slip surpasses the setpoint of permissible wheel slip velocity established in the controller. Thus, a precise calculation of slip velocity is needed, which requires an accurate knowledge of the longitudinal vehicle speed.

At the beginning of this project, the velocity of the wheel loader was obtained with the Doppler effect ground speed sensor, which is discussed in Section 7.2.1. However, as stated in that section, the ground speed sensor has some known issues with accuracy, especially at low velocities. There is also some lag of the measurement value, which can cause issues for the controller. Furthermore, the ground speed sensor is quite expensive. These factors motivated the development of a different way of acquiring the velocity of the wheel loader. The solution to this problem is a logic-based velocity estimator based on the wheel speeds.

The algorithm uses the wheel speeds ω_i and the effective radius of the tire r_e to calculate the theoretical linear velocity of each of the tires. For the majority of the time, the vehicle is assumed to be in a condition wherein at least one wheel is maintaining proper traction against the ground surface (i.e. avoiding the region of macro-slip). It has been shown through experimental results that this is a reasonable assumption for almost all conditions where no traction control is active, or where the traction control is limited to fewer than three wheels being braked at once. Under these assumed conditions, the slowest linear velocity of the four tires is closely representative of vehicle velocity.

The algorithm (Fig. 7.9) starts by reading the wheel velocities and braking signals from the controller. If the estimator sees that fewer than three wheels are being braked, then the vehicle velocity estimate will be equal to the slowest longitudinal velocity read by the estimator. (This is based on the above assumption that at least one of the four tires is not slipping significantly and travels at the vehicle velocity.) Therefore:

$$v_{x,est} = \min(R_{st,i} r_e \omega_i) \ , \tag{7.3}$$

where $R_{st,i}$ is a correction factor for wheel i based on the steering angle.

Whenever three or more tires are being braked, there is an increased likelihood that all four tires could be slipping heavily against the ground. Therefore, the algorithm must be

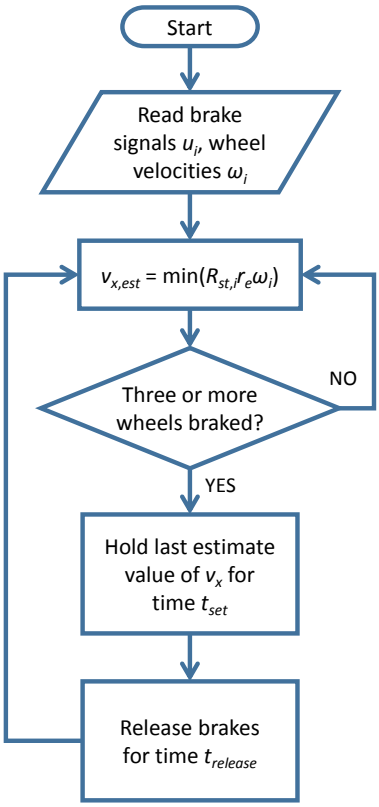

Figure 7.9. Flowchart of logic-based vehicle velocity estimation algorithm.

able to maintain a reasonable estimate of linear velocity even in that condition. This is done by storing the previous estimate $v_{x,est}$ and trusting this value for a given time t_{set} when three or more brakes are being actuated. Once the time for trusting this value has elapsed, the algorithm commands the system to reduce the number of wheels braked to two for a given time $t_{release}$. Typically, the algorithm is constructed so that both wheels on one of the two axles will be unbraked during $t_{release}$ (which axle to release will depend on the application). During the release command, it is expected that at least one of the two unbraked wheels will slow down to the velocity of the vehicle.

This logic-based algorithm allows for an accurate vehicle velocity estimation (and hence a good control signal). This is true even when the system is being steered severely or more than two wheels are being braked. Of course, in the case that the torque limitation concept

from Section 4.4 has been implemented, the TC system will never brake more than two wheels, rendering this consideration irrelevant.

7.4 Throttle Reduction

Section 4.4 of this work presents a methodology for improving the system performance through a combination of active braking and throttle reduction. This scheme takes advantage of the vehicle powertrain architecture to keep the system in the optimal slip condition while mitigating the loss of power due to braking extra wheels. After the results of the field tests in Section 9.2 showed increased fuel consumption with the standard braking scheme, the torque control method was implemented on the prototype system.

There are multiple ways in which this could be accomplished. Perhaps the most straightforward would be to modify the CAN messages being sent to the vehicle engine control unit (ECU). The ECU contains the actual software which controls the engine speed and torque output. However, the researchers did not have adequate access to modify the CAN messages sent to the ECU, so a different solution had to be found.

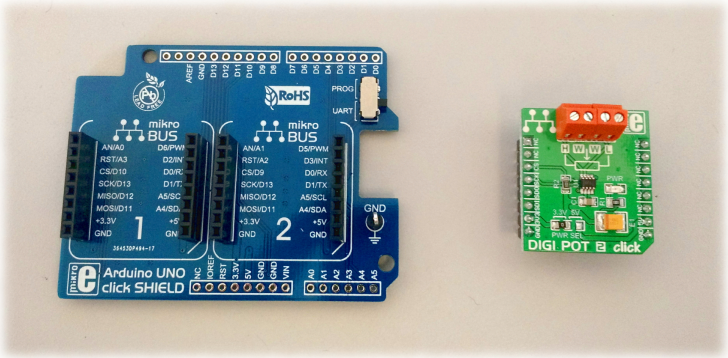

Figure 7.10. Digital potentiometer used for overriding user throttle commands.

The final implemented solution for reducing the engine torque into the system was to override the analog voltage signal sent directly from the operator throttle pedal. By using a digital potentiometer (Fig. 7.10) as a voltage divider installed downstream of the throttle pedal, it was possible to reduce the operator command as seen by the ECU. This, in turn,

reduced the torque output of the engine. The digital potentiometer was controlled via Arduino, which itself received commands from LabVIEW by way of the National Instruments hardware.

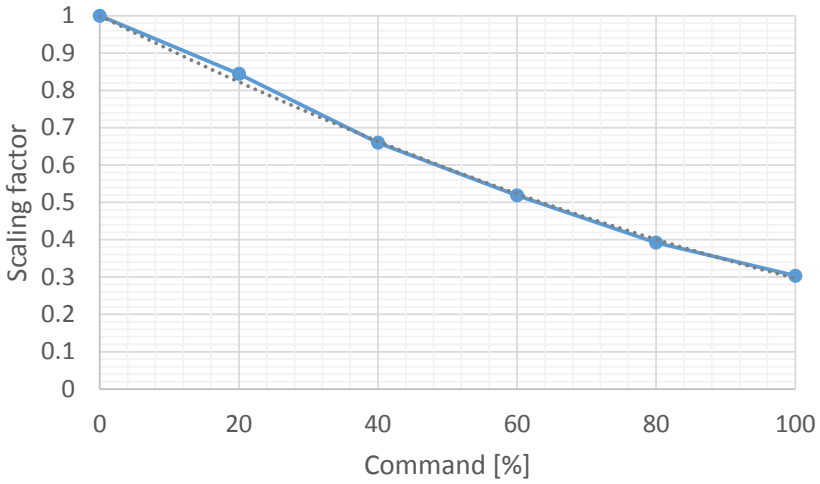

Figure 7.11. Experimental voltage divider signal reduction response.

In practice, this setup is capable of reducing the throttle command seen by the ECU, but only to a certain degree. Figure 7.11 shows the voltage scaling response of the voltage divider in practice. With a 100% reduction command, the divider is only capable of reducing the throttle command to 30.4% of its initial value. However, it is extremely rare that the engine torque should need to be lowered by more than 70%. Therefore, this response is more than sufficient for this work. To generate a proper command, the data was fit to a best-fit line (also shown in Fig. 7.11). As the response is not completely linear, a second-order polynomial best fit line was used, which achieved an R^2 value of 0.998. The inverse of this best-fit line's function was then used to generate the command to the voltage divider based on the desired throttle reduction.

Once the torque control electrical architecture was finalized, it was necessary to determine the proper algorithm for generating the torque reduction signal. This methodology is outline in Fig. 7.12. At first, the controller generates braking signals as if all four wheels are capable of being braked. The algorithm examines the four braking signals, and makes a decision of whether or not to reduce the user throttle command based on whether there are more

than two wheels (i.e. more than one wheel per axle) which need to be braked concurrently (see Section 4.4 for further discussion on this logic). If so, the algorithm does two things. First, it generates a throttle reduction command u_{th} based on the third-highest braking command of the four wheels (shown here as u_{B3}). This command is scaled by some factor r_{th}, which is determined in the course of designing the traction control system. Second, the algorithm reduces all four braking signal commands by an amount equal to u_{B3} (negative braking commands are simply saturated to zero). This is done because the resulting torque reduction from the engine affects all four wheels. Therefore, the amount of braking torque applied to all wheels should be reduced.

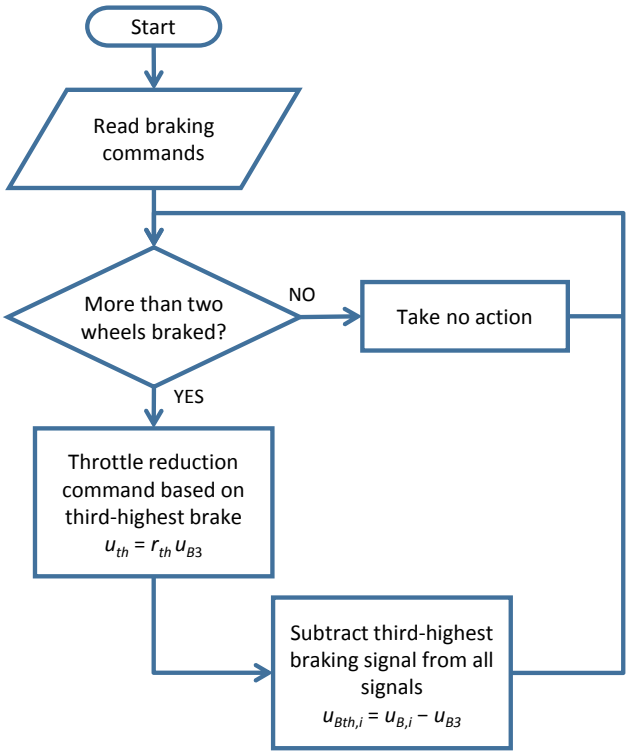

Figure 7.12. Algorithm for torque control implementation.

The resulting system was implemented for the final system productivity tests (Section 9.4) as one of many considerations geared toward reducing fuel consumption with traction control active. It successfully enacts the strategy outlined in Section 4.4 on the prototype machine, reducing overall braking action without a noticeable loss in traction control performance.

However, the throttle reduction architecture presented here is not perfect. The throttle pedal output can only be modified to a certain degree, and the altered electrical setup caused occasional ECU errors to be thrown (as the vehicle on-board diagnostics system is designed to identify when it has been tampered with). For future development work, it would be better if this algorithm were implemented directly on the control logic of the ECU. This would allow for a better response, both in terms of throttle reduction range and stability and in terms of response time.

7.5 Operator-in-the-Loop Control

Different ground conditions yield different friction coefficients and in result different slip-friction curves (see Section 2.3.2). This directly affects the mechanisms involved in producing wheel slip. It is intuitive that wheels will tend to slip more easily on wet conditions than in dry conditions. Originally, the controller was tuned to have the best performance when pushing against the test barrier at the laboratory. This controller tuning was designed for a fast and aggressive response to slip, in order to perform well on the very repeatable test conditions of the lab tests (see Chapter 8 of this work).

When the field tests began (Chapter 9), it quickly became apparent that the controller parameters used in lab tests were not well suited to the field test conditions. Based on feedback from expert operators after the first field tests, it was decided that the controller should incorporate different aggressiveness settings for adjusting the severity of the tore slip correction. Weather conditions, ground surface material, and surface conditions all affect the system behavior. Therefore, multiple levels of traction control where defined, with each level aiming to satisfy certain ground conditions during the work cycle.

As seen in Table 7.3, each level of the controller assigns different constant values of PID gains (proportional gain K_P, integral gain K_I, and derivative gain K_D), output rate limit and slip velocity set point. As the level number increases so does the controller aggressiveness. For instance, levels 1 to 3 were said by the operators to be the best settings for most field test digging operations. Levels 4 and 5 were generally seen as too severe for those operating conditions, braking the system too harshly and resulting in a negative feel. In practice, the

Table 7.3. Controller level settings.

Level	**Slip Velocity Setpoint [km/h]**	**Output Rate Limit [V/s]**	K_P	K_I	K_D
0 (OFF)	5	0	0	0	0
1	1.0	60	0.5	1.0	0.0001
2	0.8	60	0.6	1.2	0.0001
3	0.6	85	0.7	1.4	0.0001
4	0.5	100	0.8	1.5	0.0001
5	0.4	150	0.9	1.75	0.0001

operator can manually choose among the five levels by pushing a button and cycling through the levels. The level chosen by the operator can be seen on the added screen installed on the cab of the machine (Fig. 7.13).

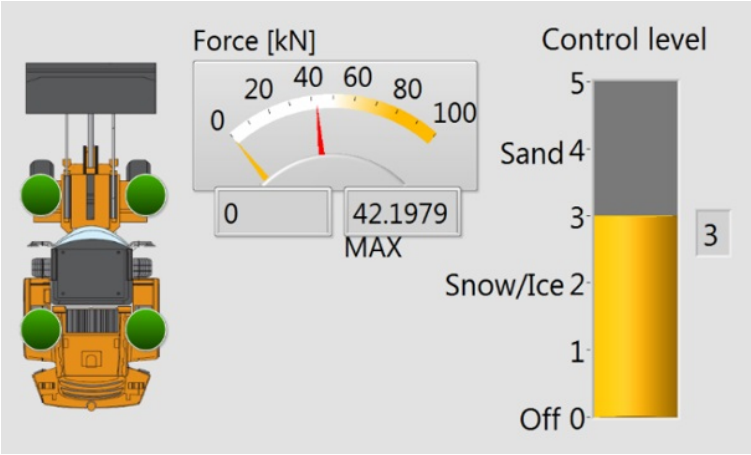

Figure 7.13. Auxiliary screen readout to WL operator.

As the operator cycles through the levels, the display changes according to the level selected. This is a relatively simple mode of operation which should encourage operators to use the traction control in ways that they may see fit. The low complexity of choosing a level should also avoid causing confusion for the machine operator. The auxiliary display is configured to show other pertinent data to the operator, such as an indication of which

wheels are currently being braked and an estimate of the current and maximum pushing force of the wheel loader in a given test.

This rudimentary severity level selection method serves much the same purpose as the online controller optimization discussed at length in Chapter 5. Ideally, the optimization would be incorporated into the machine control structure to adapt the controller to the ground condition automatically. However, the additional sensors and computing power necessary to run such an optimization online prevented that system from being used in the final production controller implementation. Allowing the operator to adjust the controller parameters himself is a necessary compromise which allows for the TC system to be adjusted to different operating conditions without the added system complexity of the online optimization.

8. LABORATORY TESTS

A wide array of different tests were conducted at Maha Laboratory (Purdue University) with various objectives. These include tests to characterize the system for improving the simulation model, proof-of-concept demonstrations, and system performance characterizations once the traction control system had been fully implemented. This chapter describes some of those tests in detail.

8.1 Slip-Friction Curve Generation

One of the most important details for the system simulation model discussed in Chapter 3 of this work is the inclusion of a detailed description of the relationship between tire wear and longitudinal wheel friction force (see Section 2.3 for details). It is important that this interaction is modeled correctly, as it is directly tied to the potential of the TC system to improve system performance.

Therefore, experiments were designed in order to capture the interaction between tire slip and longitudinal force under different operating conditions. These tests are based on the work done by Rajamani [66]. Other strategies for wheel force estimation exist, including using unscented Kalman filters [67, 68] and neural networks [69].

8.1.1 State Observer Construction

In order to assess the slip-friction characteristics of the reference machine's tires, it is necessary to know, not just the total longitudinal force generated by the wheel loader, but the individual forces at each tire. However, there is no easily implementable sensor setup which will allow for direct independent measurement of these forces. Instead, Rajamani's method is based on a state estimator which approximates the force at each wheel using measurements of wheel speed, vehicle speed, and vehicle acceleration. With the sensor setup described in Section 7.2.1, all of these can be measured.

To begin, a state-space system model is proposed with a state observer to estimate the force at each wheel. The observer structure is as follows:

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx \\ \dot{\hat{x}} &= A\hat{x} + Bu + K_o(y - C\hat{x}) \ , \end{aligned} \tag{8.1}$$

where x is the state vector for the system, u is the input vector, y is the output vector, and $\hat{x}$ is the estimate of the real state values. A, B, and C are the typical state-space system matrices which describe the dynamics of the system, and K_o is the *observer gain matrix*, which controls the convergence of the state estimate to the actual state values.

State observers are used when information about one or more states is not directly available in the measurements of a system. The goal is to drive the state estimate $\hat{x}$ to the actual value of the state x using the available system measurements (or outputs) y. Essentially, an estimate of the system output $\hat{y} = C\hat{x}$ is compared with the actual measurements y, and an input signal is added to the state estimate dynamics in order to drive the estimation error $\hat{y} - y$ to zero. If the system is observable (i.e. if its observability matrix has full rank), this will cause the state estimate $\hat{x}$ to converge to the actual state vector x over time, with a convergence rate based on the strength of K_o. Therefore, states which were not explicitly represented in system measurements are observed via the available outputs.

In order to estimate the forces at each wheel, the following state, output, and input vectors were defined:

$$x = \begin{bmatrix} v_x \\ \omega_{FL} \\ \omega_{FR} \\ \omega_{RL} \\ \omega_{RR} \\ F_{x,FL} \\ F_{x,FR} \\ F_{x,RL} \\ F_{x,RR} \\ T_E \end{bmatrix} = \begin{bmatrix} v_x \\ \omega_i \\ F_{x,i} \\ T_E \end{bmatrix} \quad y = \begin{bmatrix} \sum F_{x,i} \\ v_x \\ \omega_{FL} \\ \omega_{FR} \\ \omega_{RL} \\ \omega_{RR} \end{bmatrix} = \begin{bmatrix} \sum F_{x,i} \\ v_x \\ \omega_i \end{bmatrix} \quad u = F_{roll} \ ,$$

where v_x is the system linear velocity, ω_i and $F_{x,i}$ are the rotational velocity and longitudinal force at wheel i, respectively, T_E represents the total engine torque into the system, and F_{roll} is the rolling resistance. The state vector here contains individual values of the force at each wheel $F_{x,i}$, while the output contains only the sum of these four forces (found by using the data acquired by the accelerometer on the system). The other outputs used here are the vehicle velocity and the individual wheel velocities (each of which has a dedicated sensor on the machine). Furthermore, the state observer contains an estimate of the total torque into the system, for which no sensor is installed.

It is important to ensure that the system dynamic equations that define A, B, and C are well formulated. In the case of the present system, as described in Chapter 3, the state-space system matrices are as follows:

$$A = \begin{bmatrix} 0 & 0 & p & 0 \\ 0 & 0 & r & q \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad B = \begin{bmatrix} -\frac{1}{m} \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 0 & s & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \ . \tag{8.2}$$

These have the following submatrices:

$$p = \frac{1}{m}\begin{bmatrix}1 & 1 & 1 & 1\end{bmatrix} \quad q = \frac{1}{I_w}\begin{bmatrix}R_{FL}\\ R_{FR}\\ R_{RL}\\ R_{RR}\end{bmatrix} \quad r = -\frac{r_e}{I_w}I_{4\times 4} \quad s = \begin{bmatrix}1 & 1 & 1 & 1\end{bmatrix} .$$

In these matrices, R_i represents the percentage of engine torque distributed to wheel i. By examining matrices A and C, it can be seen that this system is in fact observable. Therefore, a proper setting of the observer gain matrix K_o will allow the state estimate to converge to the correct values. For the present work, K_o was determined using the differential matrix Riccati equation (as outlined in [51]). Essentially K_o is calculated as follows:

$$K_o = PC^T R^{-1} , \tag{8.3}$$

where P is the solution to the Riccati equation:

$$\dot{P} = AP + PA^T - PC^T R^{-1} CP + Q . \tag{8.4}$$

In the above equations, Q is a positive-semidefinite matrix and R is a positive-definite matrix, both related to the spectral densities of the observation noise, while P is the estimation error covariance matrix. With a proper setting of Q and R to adjust the intensity of the observation noise, a proper K_o is generated which gives suitable state observer dynamics.

8.1.2 Test Setup

In order to simulate a wide range of operating conditions, multiple different ground conditions were tested. Tests were conducted at normal and low tire pressures to assess how changing pressure affects the force generation characteristics of a tire. The different conditions which tested are shown in Table 8.1.

In order to have the best possible representation of a given condition, it is important to test at extreme values of slip. Therefore, the tests were conducted in alternating cycles of

Table 8.1. Test conditions for slip-friction curve tests.

Ground Condition	Tire Pressure
Dry concrete	Normal
	Low
Snow on concrete	Normal
	Low
Snow on grass	Normal
	Low

rapid acceleration and rapid deceleration, as shown in Fig. 8.1. Using the state estimator described above, the force generated at each wheel was estimated and then correlated with the corresponding slip value for that wheel.

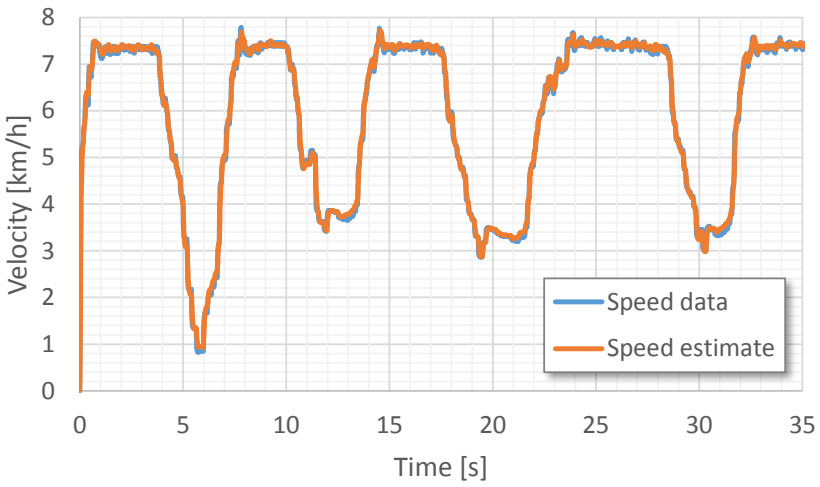

Figure 8.1. Representative test cycle for slip-friction curve generation tests.

Figure 8.1 also demonstrates the reliability of the state estimator used in the tire force estimation. For the vehicle velocity, both the measured data and the estimated values are shown. The estimate matches the data quite well, indicating that the estimator has a good matching performance for that state. Since all measured states match their estimates very well, it is likely that the unmeasured state estimates are also reliable, assuming a reasonably accurate system model [51].

8.1.3 Experimental Results

Each test from this particular method produces a series of points relating the slip ratio at a given tire to its tractive force (which is essentially the friction coefficient scaled by the normal force at that tire). Therefore, these points can be plotted to give an idea of the slip-friction characteristics for the wheel loader.

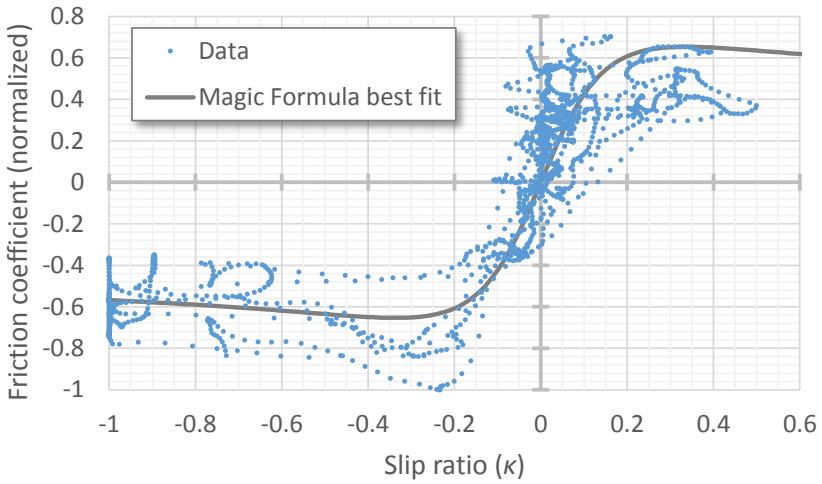

Figure 8.2. Slip-friction data generated for reference wheel loader system (snow on concrete).

Figure 8.2 shows such a plot for a particular test condition. For this case, it can be seen that it is much easier to bring the system to high values of negative slip (i.e. wheel lock-up when braking) than it is to achieve high values of positive slip (i.e. wheels spinning out against the ground). This is in part due to the fact that the wheel loader engine management is designed to avoid such spin-out scenarios.

Nevertheless, this plot still represents a relatively wide range of slip values for the system. For the purposes of this research, there is one major assumption which controls the Magic Formula modeling taken from these data points. This is the assumption that the slip-friction relationship has odd symmetry. This gives two important constraints to the model: first, it will pass through the origin; and second, the behavior with positive slip mirrors that with negative slip. Due to these restrictions, the entire Magic Formula model can be generated from the data shown above.

In order to properly fit the modeled curve to the data (as in Fig. 8.2), the Magic Formula parameters (B_x through E_x) were adjusted manually to reduce the mean square error for the model with respect to the measured data. This method was sufficient for the sake of this research; however, should a more rigorous curve-fitting be required, there are other strategies for tuning the Magic Formula parameters. For example, Rajamani [66] suggests a recursive least-squares algorithm. A similar method incorporating mean squared error is used in the online optimization formulation presented in Section 5.4.2 of this work.

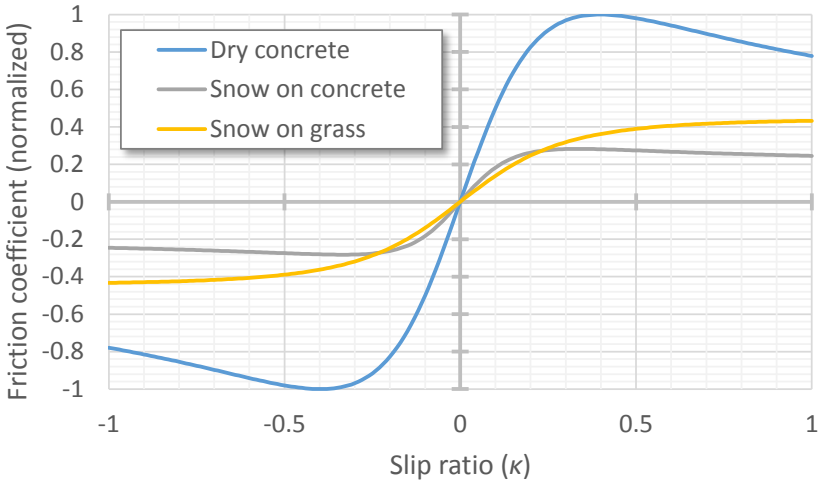

Figure 8.3. Slip-friction model ground condition comparison.

The data for different road and tire conditions was taken, and curves were fit to each data set. Figure 8.3 shows a comparison of the resultant curves. This figure has some important results for the research here. First, the trend of road conditions is quite clear and also reasonable to what would be expected for this system. Dry concrete, the best case for friction force generation between the ground and the wheels, has by far the highest friction coefficient for a given slip ratio. The other cases exhibit a similar behavior, with snow on concrete having the lowest overall friction coefficient. Snow on grass was the most variable case, but it is more or less comparable to snow on concrete.

Furthermore, the condition of the tires also has an impact on friction coefficient, with tires having a lower inflation pressure achieving the higher friction coefficient (Fig. 8.4). This

makes sense intuitively, as tires with lower inflation pressure deform more readily and thus have a larger contact surface area with which to generate friction force against the ground.

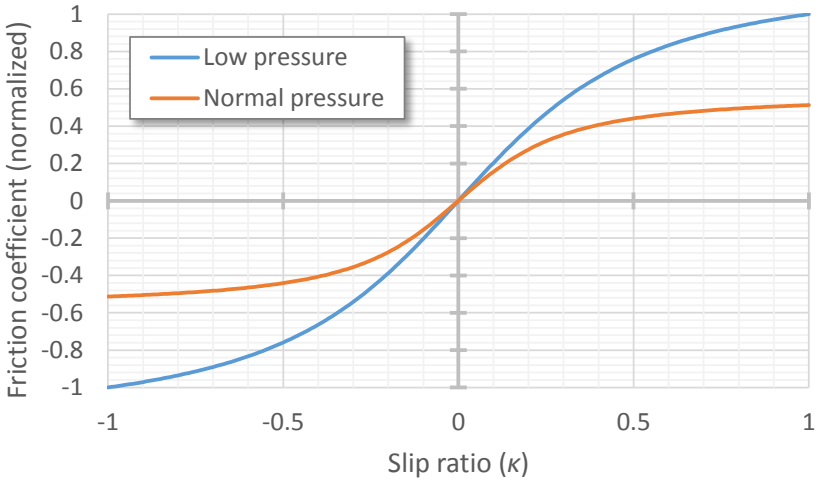

Figure 8.4. Slip-friction model tire pressure comparison.

It should be noted that these slip-friction curves are not intended to always match the actual system performance exactly. There are a great many different potential ground conditions which cannot all be tested in this way, and even within the same condition there can be a high degree of variability in the resultant friction force. This data was taken primarily to give the investigators a good idea of reasonable values for the friction coefficient of each wheel. Once it was obtained, these slip-friction relationships were incorporated into the system model described in Chapter 3 to allow the model to better approximate the real-world behavior of the vehicle.

8.2 Traction Control Performance Comparison

Once the modeling and simulation phase was completed for the TC system, it needed to be experimentally tested. This testing was accomplished in two phases. First, a series of very controlled tests was conducted to assess the performance increase of the system with traction control implemented. Those tests are the subject of the current section. Other

tests were run at a proving ground to mimic more closely the real-life operation of the wheel loader. Those tests are discussed in Chapter 9.

Figure 8.5. Test barrier used for experimental testing at Maha Laboratory, showing the tires and two of the four steel plates.

The laboratory experiments were designed to create a repeatable test condition in which the tires slipped reliably for comparing control strategies side-by-side. To accomplish this, a barrier was designed to provide a resistive force for the vehicle to push against (Fig. 8.5). This test barrier contains three large tires against which the wheel loader would push, using a steel plate attached attached across the front of the bucket (Fig. 8.6). The barrier was supported from behind by large trusses made of steel I-beams, which kept the structure from being destroyed by the immense forces involved. Finally, steel plates were anchored into the concrete driving surface, located so that the tires would be on top of the plates when the wheel loader is pushing against the barrier. These plates provide a repeatable driving surface different from the concrete surface. It also allows for removing individual plates for generating mixed ground conditions at the wheels.

Figure 8.6. Plate mounted on wheel loader bucket for pushing against tires.

8.2.1 Test Plan

A braking test plan was developed in order to test the performance of the traction control system. The list of tests conducted is shown in Table 8.2. The tests are grouped based on the number of steel plates used under the tires. The "Four plates" tests included plates under all four wheels, the "Asymmetric" tests had plates under the front left and rear right tires, and the "One plate" tests had only a single steel plate under the front left tire. As each configuration was run both with and without traction control, this made six test conditions total. These tests gave a range of different operating conditions which caused the wheels to slip in different ways with varying resultant forces.

Three repetitions were done of each test condition (for a total of eighteen repetitions). The order of the tests was determined by the number of plates needed, so that plates would not have to be removed and reset. Therefore, all "Four plates" tests were grouped together, as were all "Asymmetric" tests, and so on. Within these groupings, however, the repetitions were run in a random order to minimize other effects influencing the data.

Each test was conducted by approaching the tires in front of the barrier and pushing against them with the accelerator at full throttle (i.e. fully depressed) and the machine in first gear, with automatic gear shifting disabled. The accelerator remained at full throttle

Table 8.2. Test matrix for laboratory tests.

A. Baseline Tests (No Control)	
1 – 3	Four plates (under all wheels)
4 – 6	Asymmetric (two plates at opposite corners)
7 – 9	One plate (under front left tire)
B. Augmented PID Tests	
1 – 3	Four plates (under all wheels)
4 – 6	Asymmetric (two plates at opposite corners)
7 – 9	One plate (under front left tire)

for at least five seconds after making contact with the tires. Once the system converged to a steady-state slip condition (i.e. all four wheels turning at a constant velocity or not at all), the operator released the accelerator. All of the recorded data was stored for post-processing.

8.2.2 Objective Function Definitions

In order to assess the performance of the different controllers, objective functions were defined which quantified various aspects of each test. It was decided to use two separate function sets, one which took into account only the first five seconds of data after the wheel loader comes into contact with the tires and one which took only the first three seconds. These lengths of time were chosen based on representative loading cycles for the wheel loader, where the operator typically pushes into the work pile for three to four seconds before reversing. For the sake of the control tests, five seconds was also chosen since pushing against tires is a slightly different condition and the operator is more likely to approach them more carefully than he would a work pile. It also gives the data more time to develop, to contrast the results with the relatively transient three-second window data.

The first performance parameter which needed to be quantified is pushing force for the machine. While pushing force cannot easily be measured on the system, a force estimate was generated using the pressures in the implement boom hydraulic cylinder and the geometry of the system (see Section 3.8 for more information). A typical “no control” test is shown in

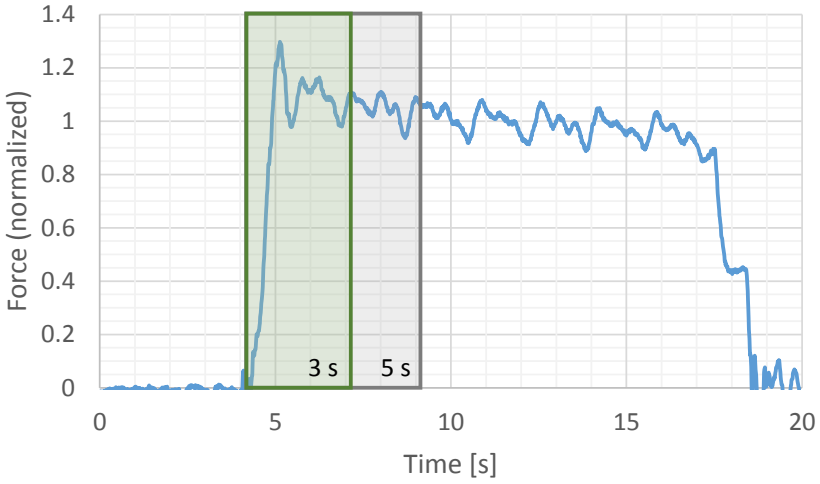

Figure 8.7. Typical lab test force curve, with objective function time windows shown.

Fig. 8.7, with boxes highlighting the first three and five seconds of the test after the wheel loader begins pushing against the tires. There are two data values of concern when it comes to examining force. The first is the average force over the time span, which is an integral term attempting to roughly approximate the amount of material the wheel loader could be able to excavate from a material pile during the time window. The second objective function value of note is the final force at the end of the time window. This gives an indication of the controller performance, as it is seeking to maximize the instantaneous force at the wheels by controlling wheel slip. The force signal is filtered using a low-pass filter to remove effects from measurement noise when computing the objective functions.

The definitions for the average force and final force objective functions are as follows:

$$F_{avg,i} = \frac{1}{T} \int_{t_0}^{t_1} F dt \tag{8.5}$$

$$F_{final,i} = F(t_1) \ , \tag{8.6}$$

with the following definitions.

- t_0 and t_1 are the start and end times of the test data, respectively

- T is the total test time $(t_1 - t_0)$, which will be either three or five seconds for defined test cycle
- F is the total pushing force of the wheel loader (function of time t)

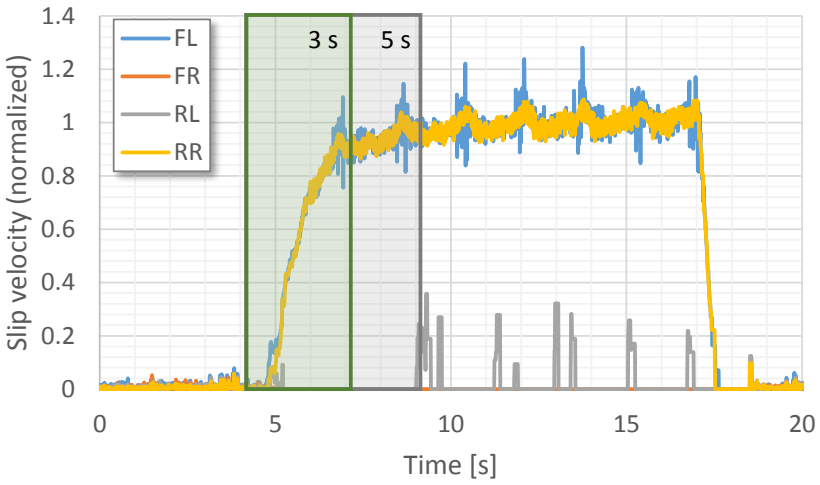

Figure 8.8. Typical lab test wheel slip curve, with objective function time windows shown.

The next objective functions to be defined were related to the amount of tire wear. For this study, tire consumption is primarily represented by slip velocity, a description of the relative velocity between the driving surface and the contact patch of the tire (Fig. 8.8). This quantity was found using the rotational velocities of each tire (shown above) and the linear velocity of the vehicle (which is close to zero through most of the test). In order to give a picture of the overall tire wear, the cumulative slip distance (integral of the slip velocity over the time interval) was used. Another data point which may be relevant is the peak slip velocity seen for each case, as this gives an idea of the performance of the controller.

The objective function values for cumulative slip distance and peak slip velocity are defined as:

$$x_{slip} = \sum_{k} \int_{t_0}^{t_1} v_{s,k} dt \tag{8.7}$$

$$v_{s,max,i} = \sup\{\boldsymbol{v_s}(t) : t_0 \leq t \leq t_1\} \tag{8.8}$$

$$\boldsymbol{v_s} = \begin{bmatrix} v_{s,FL} \\ v_{s,FR} \\ v_{s,RL} \\ v_{s,RR} \end{bmatrix} \tag{8.9}$$

$$v_{s,k} = \omega_k r_e - v_x \ . \tag{8.10}$$

- $v_{s,k}$ is the slip velocity at wheel k (k can be FL, FR, RL, or RR)
- The sup function returns the maximum value of any element of vector $\boldsymbol{v_s}$ over the whole time span from t_0 to t_1
- ω_k is the rotational velocity of wheel k
- r_e is the dynamic radius of the tire
- v_x is the linear velocity of the wheel loader

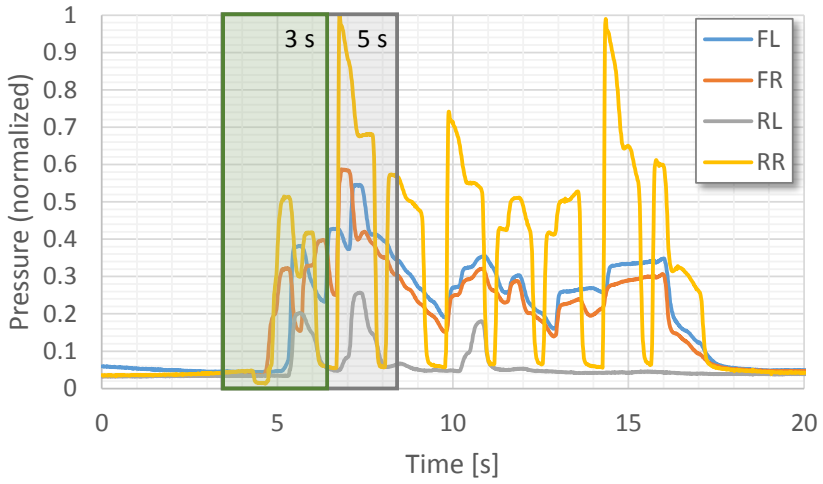

Figure 8.9. Typical lab test braking pressure curve, with objective function time windows shown.

The final parameter which needed to be quantified is the brake consumption when using the controlled system. As an analog for brake consumption, the braking power is estimated

using braking pressure (see Fig. 8.9) and wheel speed. The braking power objective function is:

$$P_B = \sum_k \int_{t_0}^{t_1} p_{B,k}\omega_k dt \ . \tag{8.11}$$

- $p_{B,k}$ is the braking pressure at wheel k

In order to ensure consistency between tests, external tire temperature was also measured. This helped to keep differences in tire performance based on temperature from impacting test results. Furthermore, CAN data from the machine itself was collected to provide even more data for each test.

8.2.3 Test Results Matrix

The data acquired from the proposed tests above was used to populate the test result matrix shown in Table 8.3. The matrix contains objective functions whose values are used to compare the system performance with and without the traction control active.

Using these expressions above, the objective functions $\hat{J}_i$ can be found by averaging the three values for each test condition (i.e. from each of the three repetitions). The normalized objective functions for the cases with traction control implemented are compared to the baseline (no control) case. That is, when the objective function is denoted with "% ↑" (percent increase):

$$\bar{J}_{Bi} = \frac{\hat{J}_{Bi}}{\hat{J}_{Ai}} - 1 \ , \tag{8.12}$$

and when the objective function is denoted with "% ↓" (percent decrease):

$$\bar{J}_{Bi} = 1 - \frac{\hat{J}_{Bi}}{\hat{J}_{Ai}} \ . \tag{8.13}$$

A total of eighteen repetitions were taken in the laboratory testing. The results from these tests were processed to populate the test matrix. To ensure similar conditions among the tests, the temperature of the tires and plates was recorded. Once a test was completed, time was given for the tires and plates to cool so that they had the same starting temperature

Table 8.3. Test objective function matrix for laboratory tests.

Controller	**Tests**	**Avg. Force**	**Final Force**	**Slip Distance**	**Peak Slip Velocity**	**Brake Wear**
A (No control)	1 – 3	$\hat{F}_{avg,A1-3}$ [N]	$\hat{F}_{final,A1-3}$ [N]	$\hat{x}_{slip,A1-3}$ [m]	$\hat{v}_{s,max,A1-3}$ [m/s]	0
	4 – 6	$\hat{F}_{avg,A4-6}$ [N]	$\hat{F}_{final,A4-6}$ [N]	$\hat{x}_{slip,A4-6}$ [m]	$\hat{v}_{s,max,A4-6}$ [m/s]	0
	7 – 9	$\hat{F}_{avg,A7-9}$ [N]	$\hat{F}_{final,A7-9}$ [N]	$\hat{x}_{slip,A7-9}$ [m]	$\hat{v}_{s,max,A7-9}$ [m/s]	0
B Augmented (PID)	1 – 3	$\bar{F}_{avg,B1-3}$ [% ↑]	$\bar{F}_{final,B1-3}$ [% ↑]	$\bar{x}_{slip,B1-3}$ [% ↓]	$\bar{v}_{s,max,B1-3}$ [% ↓]	$\hat{P}_{b,B1-3}$ [Pa]
	4 – 6	$\bar{F}_{avg,B4-6}$ [% ↑]	$\bar{F}_{final,B4-6}$ [% ↑]	$\bar{x}_{slip,B4-6}$ [% ↓]	$\bar{v}_{s,max,B4-6}$ [% ↓]	$\hat{P}_{b,B4-6}$ [Pa]
	7 – 9	$\bar{F}_{avg,B7-9}$ [% ↑]	$\bar{F}_{final,B7-9}$ [% ↑]	$\bar{x}_{slip,B7-9}$ [% ↓]	$\bar{v}_{s,max,B7-9}$ [% ↓]	$\hat{P}_{b,B7-9}$ [Pa]

in all test repetitions. Weather was also taken into consideration, as ambient temperature and humidity can have a significant impact on tire behavior.

The test results matrix (Table 8.4) was split into two parts, one for each time window discussed in Section 8.2.2. Both time segments were assessed starting the moment the wheel loader made contact with the tires on the test barrier.

Table 8.4 shows that the traction control system is capable of improving all four performance objective functions of interest for this investigation. In all cases, the pushing force was increased and the tire slip decreased. In general, the best results come from the "Four plates" tests (repetitions 1–3). This makes sense, because the steel plates are more slick than the concrete, meaning that there is more room for improvement when there are more tires on plates. The difference in the three- and five-second time windows can also be seen. As the control action is still changing and has not yet fully converged in the three-second time window, the test results are generally better when using the five-second window for evaluating objective functions.

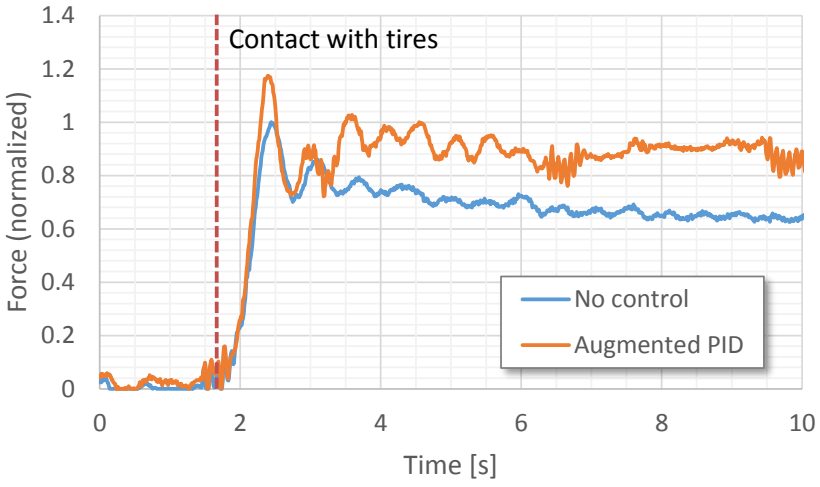

Figure 8.10. Comparison of pushing force with and without traction control (4 plates test).

Figure 8.10 shows a comparison of two different lab tests run with and without traction control implemented. It can clearly be seen from this plot that the traction control is able to quickly and effectively increase the overall pushing force of the machine. Less than two

Table 8.4. Test results matrix for laboratory tests.

3 Seconds

Controller	Tests	Avg. Force	Final Force	Slip Distance	Peak Slip Velocity	Brake Wear
B Augmented (PID)	1 – 3	14.7 % ↑	19.9 % ↑	50.1 % ↓	36.9 % ↓	0.583*
	4 – 6	10.0 % ↑	11.2 % ↑	60.7 % ↓	53.4 % ↓	0.427*
	7 – 9	12.2 % ↑	24.2 % ↑	46.7 % ↓	57.4 % ↓	0.279*

5 Seconds

Controller	Tests	Avg. Force	Final Force	Slip Distance	Peak Slip Velocity	Brake Wear
B Augmented (PID)	1 – 3	18.6 % ↑	31.8 % ↑	69.0 % ↓	41.9 % ↓	1.000*
	4 – 6	11.6 % ↑	12.2 % ↑	70.5 % ↓	43.6 % ↓	0.746*
	7 – 9	14.3 % ↑	18.6 % ↑	61.2 % ↓	46.7 % ↓	0.532*

* *Values normalized*

seconds after impacting the tires, the uncontrolled system has lost about 20% of its peak pushing force, whereas the system with TC implemented is able to maintain forces close to the uncontrolled system's peak value (after a brief time wherein the traction control has not yet converged). It is clear from this example data why both pushing force-related objective functions had significant increased for the four plates tests.

All in all, the results for the lab experiments were quite positive, with pushing force increases consistently above 10% and decreases in slip typically at or above 45%. For worse testing conditions where tire slip was easy to generate, the improvements were even more significant, with pushing force seeing increases of up to 60% and tire slip being reduced by as much as 73%. This indicates that there is quite a lot to be gained from the traction control system in terms of productivity and cost effectiveness.

8.3 Real-Time Controller Optimization Tests

The online control setpoint optimization scheme set forth in Section 5.4 of this work was added to the LabVIEW data acquisition code, and additional tests were run. These tests were not meant to be a comprehensive assessment of the optimization performance; instead, they serve more as a proof of concept.

The basic test was the same as the other laboratory tests, with the wheel loader pushing against the tires on the barrier. Once the wheel loader impacted the tires, the throttle was modulated, causing the wheels to slip significantly. There was no traction control active during these tests, as the objective was to read a large range of wheel slip values. Figure 8.11 shows the normalized wheel slip for one such test. The large slip values allowed the enhanced data buffer to fill in a reasonable amount of time, causing the slip setpoint optimization to run multiple times during a single test.

These tests were run on multiple different ground conditions to assess the ability of the system to converge to different values. The results for two such tests are shown in Fig. 8.12. It should be noted that these results are for an average slip setpoint, as it is assumed all four wheels are on the same ground condition. This assumption can be removed in the future, should it be desired to have independent setpoints for the four wheels. In these two tests, as

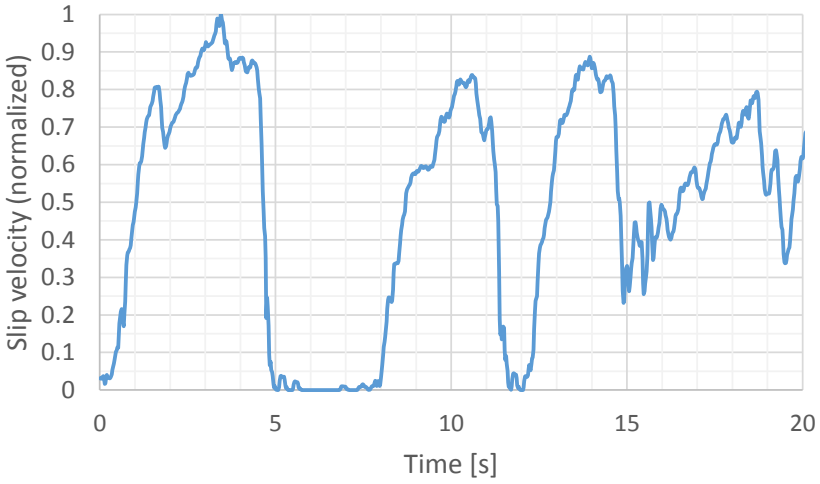

Figure 8.11. Wheel slip velocity for typical optimization test.

with the simulation shown in Section 5.4.5, it takes some time for the optimizations to begin converging. This is because the system had data stored from previous ground conditions. However, in both cases, after the second optimization (which only uses data from the ground condition in question), the results of the optimizations are very consistent, varying by at most $\pm 10\%$ from the average value. Furthermore, it is clear that the optimization is achieving different optimal setpoint results for the two ground conditions, with the optimal soil setpoint averaging about 17% higher than the setpoint for the steel plates. This is consistent with the previous experimental data, which showed that the maximum friction coefficients for soil conditions tended to be higher than that of hard surfaces like steel and concrete.

The data indicates that the optimization scheme developed in Section 5.4 is suitable for implementation on a real-world system. This is a significant result for systems like the wheel loader considered in this work, which operate on ever-changing ground conditions. The ability to automatically adapt to different surfaces would be a boon for any such traction control methodology.

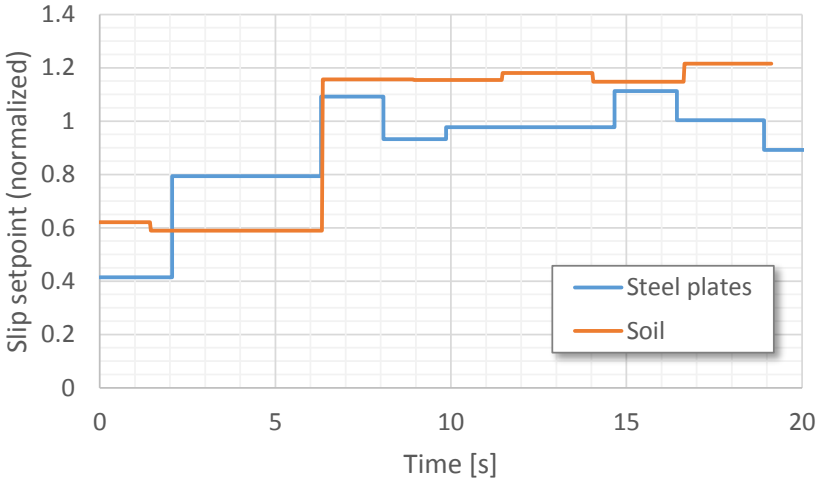

Figure 8.12. Optimized slip velocity setpoints for two different ground conditions.

8.4 Nonlinear Controller Performance Tests

The majority of the research in this project centered on developing a system capable of being implemented on a current-production construction machine. Therefore, most of the more exotic control systems were overlooked in favor of the relatively simple Augmented PID control which has been discussed at length in this work (see Chapter 4). This simpler approach makes more sense for existing machines, as it minimizes the number of added sensors and the additional computational power needed to implement the controller. However, Chapter 6 presents the effort made in developing a suitable nonlinear controller for the system. These controllers represent some of the best theoretical approaches for designing a traction controller.

The controller selected for implementation was the sliding mode control, presented in detail in Section 6.2. This controller was selected because it is capable of handling variations and uncertainties in model parameters, which is important for control systems in the real world. However, the added complexity of the adaptive control would not provide a significant increase in terms of performance, and it was deemed unnecessary for this work. The sliding mode controller used for system implementation was slightly different from that developed in Section 6.2, but the basic components are the same.

Much like the work in Section 8.3, these tests were conducted primarily as a proof of concept, and they are not meant to represent the final performance of the nonlinear control system. The tests were essentially identical to those described in Section 8.2, with the wheel loader approaching and then impacting the tires against the barrier and then the operator pressing the throttle to cause the wheels to slip significantly against the ground. All tests in this investigation were conducted under the "Four plates" setup, with all four tires driven on steel plates.

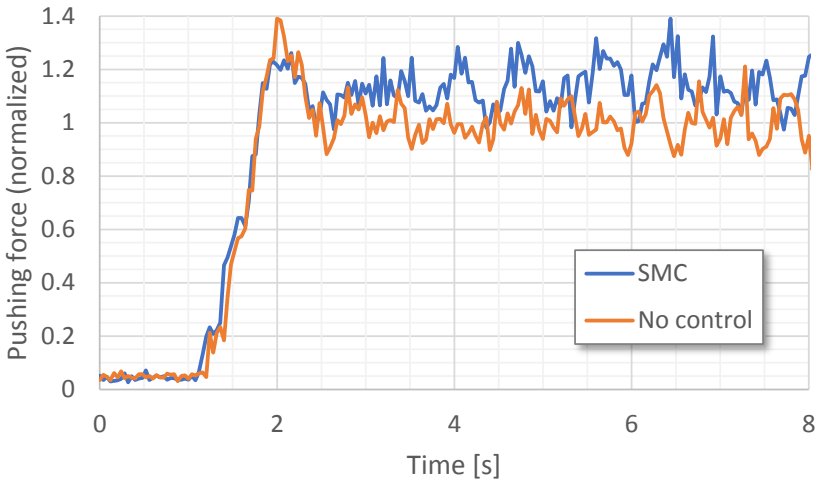

Figure 8.13. Experimental pushing force results with sliding model control implemented.

The results of two representative tests are shown in Fig. 8.13. This plot is normalized with respect to the average pushing force of the system with no traction control implemented. For both tests, the wheel loader impacts the tires at around time $t = 1$ second/ After that, there is a very sharp peak in pushing force, which decreases as the wheels begin to slip. For the case without traction control, the pushing force diminishes significantly, with the steady-state pushing force during tire slip being quite a bit below the peak force value. On the other hand, the sliding mode controller quickly identifies the wheels that are slipping and acts to slow them down. The result is that the system with sliding mode control never sees the same drop-off in force as the uncontrolled system, and its average force value in steady state is approximately 12% greater than the case without traction control.

In order to assess the overall performance of the system, repetitions of these comparison tests were run. Furthermore, test repetitions were also run for the Augmented PID controller, so that the sliding mode control performance could be compared with the current system implementation. Figure 8.14 shows the results of that performance comparison. The objective function values being compared here were taken in the three-second time window as described in Section 8.2.2.

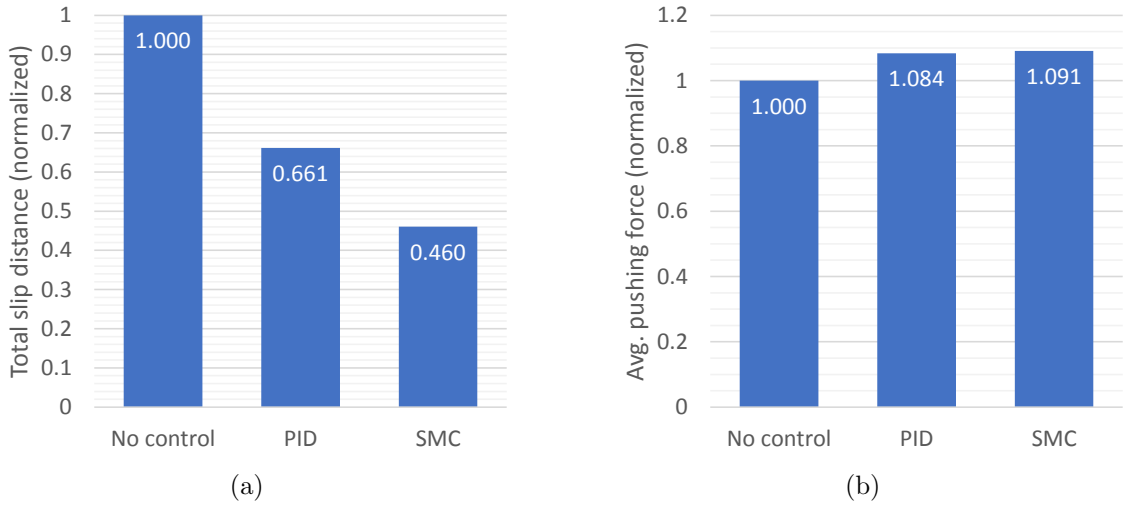

Figure 8.14. Comparison of sliding mode control (SMC) with baseline (No control) and Augmented PID controller performance, showing (a) wheel slip distance and (b) average pushing force.

From this figure, it is clear that the sliding mode control has some advantages even over the Augmented PID. The most obvious improvement comes in the form of wheel slip distance. The Augmented PID shows a 34% reduction in slip compared to the case with no traction control. However, the sliding mode control reduces wheel slip by 54%, a very significant result. Pushing force is not as strongly affected, however, with Augmented PID and sliding mode controllers increasing the average pushing force by 8% and 9%, respectively.

There are a few primary takeaways from this result. First, the sliding mode control, when properly constructed, acts much more quickly than the Augmented PID. This is to be expected, because even after the introduction of the added logic to the controller, the Augmented PID control is still susceptible to effects like integrator wind-up, which slow down

its response time. However, this advantage of the sliding mode control essentially disappears when the objective functions are analyzed over five seconds. Therefore, the Augmented PID is capable of performing roughly as well as the sliding mode control, given enough time. Also, it should be noted that the sliding mode control has not been tuned nearly as extensively as the Augmented PID. It is entirely possible that some excess braking is occurring with the sliding mode control, which would give it a faster response while negatively impacting other performance metrics, such as fuel consumption. However, it is also likely that even better results could be achieved with further investigation into the proper sliding mode control settings.

9. FIELD TESTS

Having conducted many tests in a controlled environment at Maha Laboratory, further testing was needed in conditions which were more realistic compared to the real-world operation of a wheel loader. Therefore, a series of experiments was designed which incorporated common wheel loader operating modes. These tests examined such operating cycles as digging, grading, and driving up an incline.

Furthermore, a standard Y-cycle test was conducted in order to assess the impact of the traction control system on the performance of the machine. This test strategy was developed in conjunction with an industry sponsor, and it incorporates many of the elements of industry standard productivity tests used to assess machine efficiency. The result of these tests is data on fuel consumption, loading efficiency, and tire slip, among other important parameters.

9.1 Test Descriptions

The first field tests were run at a testing ground facility at Purdue University campus in West Lafayette, Indiana. There were different kinds of tests, each designed to assess a different aspect of the machine performance (Table 9.1).

Table 9.1. List of field tests with descriptions.

Test	**Purpose**
Articulated	Assess the performance of steering angle compensation in the system
Climbing	Determine if traction control provides more stability or control when scaling a steep incline
Grading	Test the traction control performance when pushing through material (instead of digging into it)
Productivity	Examine the impact of traction control on a standard material loading cycle (Y-cycle)

The first three categories of tests were more subjective in nature. The first test was called the *Articulated Test*. Its purpose was to determine if the controller modifications

Figure 9.1. Wheel loader pushing a material pile during a Grading Test.

which allowed for steering angle compensation were adequate. By steering at severe angles throughout a digging cycle, the operator put the traction control system in many conditions where it might tend to lock up the wheels due to the speed difference when steering, as described in Section 4.5. If the system did not lock up due to the steering, then the steering compensation likely was acting correctly. The second series of tests conducted on the wheel loader was the *Climbing Test*. In this test, the machine was driven up a steep incline to examine the effect of the controller on the behavior of the system. A proper traction control system can improve system stability in such operations and potentially allow the wheel loader to climb higher. The third category of tests was the *Grading Test*, in which the wheel loader was used to level the ground surface by pushing excess material away (Fig. 9.1). These tests were conducted by first making several small piles of material. Next, the wheel loader was driven to push through these small piles in an attempt to create a flat surface. The resistance caused by pushing through the small piles can easily cause the wheels to slip, resulting in significant wheel ruts and uneven ground surface. Ideally, the traction control will help in avoiding these unwanted phenomena.

The test scenario was the *Productivity Test*. There was much more time spent on the Productivity Tests than on any of the other three tests. This is due to the fact that they are designed to mimic real-world operating conditions and are based on standard produc-

Figure 9.2. Productivity Test setup, showing wheel loader at the digging pile with dump pile in the foreground (white pole and orange traffic cone indicate dumping height and location).

tivity tests conducted for these systems by the company sponsor. Many repetitions of the Productivity Tests were made to ensure that the results were well understood.

The setup for the Productivity Test is relatively straightforward (Fig. 9.2). A short Y-cycle for loading material is outlined with a material digging pile and a dump pile located approximately 10 m from each other. The operator is tasked with transporting as many loads as possible within a predefined time frame (for this work, typically either five or ten minutes). A marker was used to show the height at which the bucket should release the material, simulating the height of a dump truck. With this defined test cycle, an on-board scale was also used to give an estimate of the mass of material in the bucket for each load.

At the end of the allotted time, the number of loads, mass of material moved, and amount of fuel consumed were all noted from the on-board scale system and the wheel loader console. A scale reading of the empty bucket was also taken before and after each test in order to ensure that no significant change in the load-weigh system had occurred. The data acquisition system was recording data through the whole test, including vehicle and

wheel velocities, fuel rate, throttle command, GPS position, controller command, and much more.

Figure 9.3 shows GPS data from a single test to give an idea of the motion of the machine during Productivity Tests. This motion is typically broken into four phases: Phase I, wherein the machine approaches the material pile to fill the bucket; Phase II, reversing from the digging pile; Phase III, when the wheel loader approaches the dumping pile to empty the bucket; and Phase IV, reversing from the dumping pile. These phases can be identified in the data using the acquired transmission gear data.

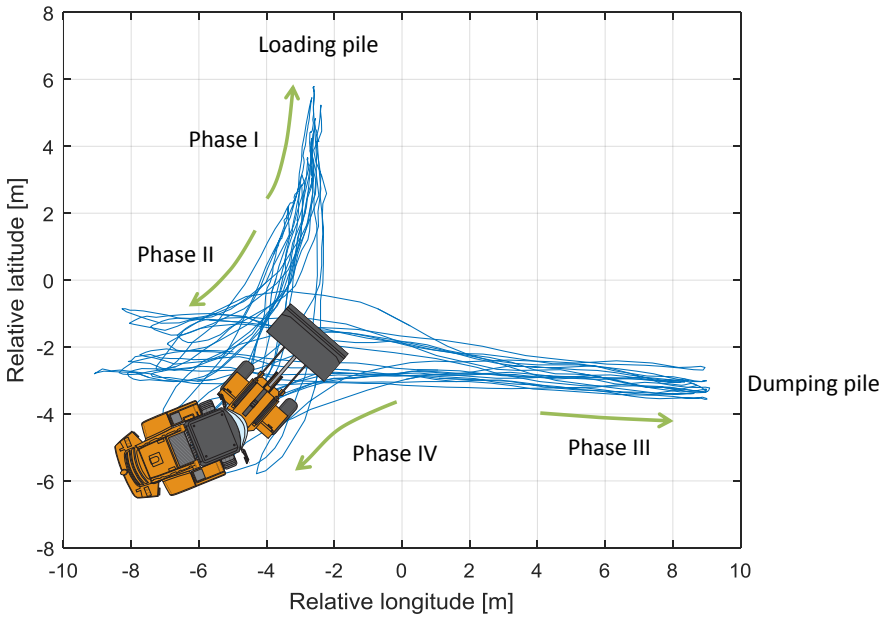

Figure 9.3. Typical wheel loader travel path for Productivity Test.

As the name implies, the objective of the Productivity Tests was to ascertain the effect of the traction control system on the wheel loader's productivity. Both the ability to quickly move material and the ability to minimize the fuel consumed by the system are important for companies operating these machines. If the traction control system impacts one of these in too negative a fashion, the reduction in tire wear and machine stability may not be worth the price in productivity.

9.2 Performance Feedback and Results

Each of the tests conducted targets specific issues for the system. For the more subjective tests (Articulated Test, Climbing Test, and Grading Test), the system performance feedback was very positive. During these tests, the system behaved exactly as expected, with the controller actuating the brakes to keep the wheels from slipping excessively against the ground. There was a significant amount of data taken for these tests as well, but in the end they served more as a verification that the system was functioning as intended.

Instead, much more focus was given to the Productivity Tests. These tests were designed to assess measurable changes in system performance with the controller active. Originally the Productivity Tests were run using material piles of both gravel and mud, but the mud proved too difficult to generate consistent bucket loads, so in general the results cited in this report were generated by digging in the gravel pile.

The field tests represent a much more difficult operating condition for the system than the relatively sterile conditions under which the lab tests were run. The field tests were run on dirt and gravel surfaces which were quite rough and resulted in serious machine vibrations throughout the tests. Dust, mud, rain, and debris from the material piles also had the potential to interfere with some of the sensors crucial to the proper operation of the TC system. Furthermore, while the test condition was relatively well defined, the operator had full control of the machine and could potentially introduce unforeseen driving behaviors. Due to these and other factors, some minor system modifications were made before the field tests were run.

One such change was the re-installation of the inductive wheel speed sensors to replace the optical sensors. These sensors do not have as high a spatial resolution as the optical sensors, but they provide a much cleaner signal that is not affected as much by dirt and vibration in the system. The computer code was also upgraded to read both the rising and falling edges of the wheel speed signal in order to improve the wheel speed update frequency. Another update to the system was the tuning of a so-called “standby control,” which kept a small amount of pressure in the braking lines to improve the dynamic response of the braking system. The standby pressure for lab tests did have a significant impact on system

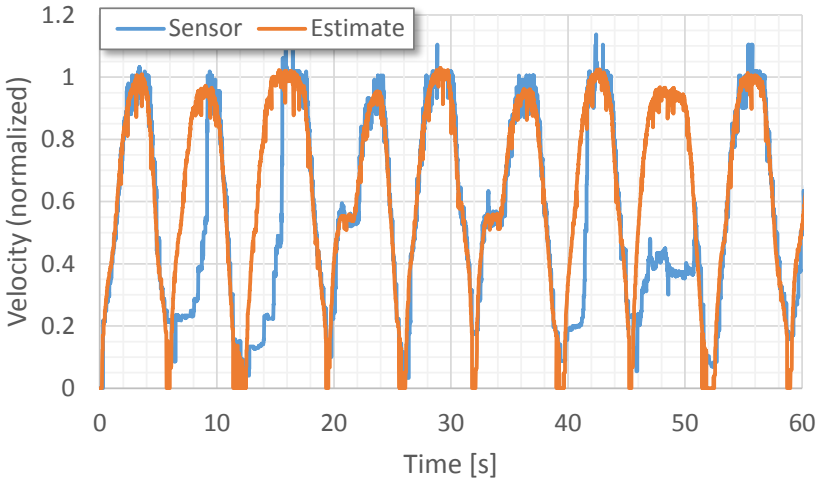

Figure 9.4. Documented signal delay issues for ground speed sensor.

behavior. But for the field tests, where the driving style is much more variable and the system starts and stops frequently, even a small amount of braking could have significant impacts on system performance. Therefore, the standby control was tuned rigorously to ensure that the standby pressure did not exceed the contact pressure of the brake calipers.

It was also decided that for the field tests, the velocity estimation algorithm would be used instead of the linear velocity sensor installed on the system. Upon close investigation, the ground speed sensor showed some instances where, when shifting the machine quickly from forward to reverse or vice-versa, the velocity reading would lag behind the actual value. These signal errors were quite severe in some cases (Fig. 9.4). The velocity estimation algorithm (outlined in Section 7.3), on the other hand, was quite reliable and allowed the controller to behave much better.

Results from the Productivity Tests are shown in Fig. 9.5. For each test, a five-minute window was analyzed wherein many different data signals were acquired to assess the system performance. For the loading cycle performed in these tests, the ground was not particularly slick, and the operator was able to keep the wheels from slipping most of the time by simply controlling the throttle position. However, when the wheel loader was digging into the material pile, there was some amount of significant wheel slip, which the controller was able to prevent with the TC system active. Because of this, the traction control system was able

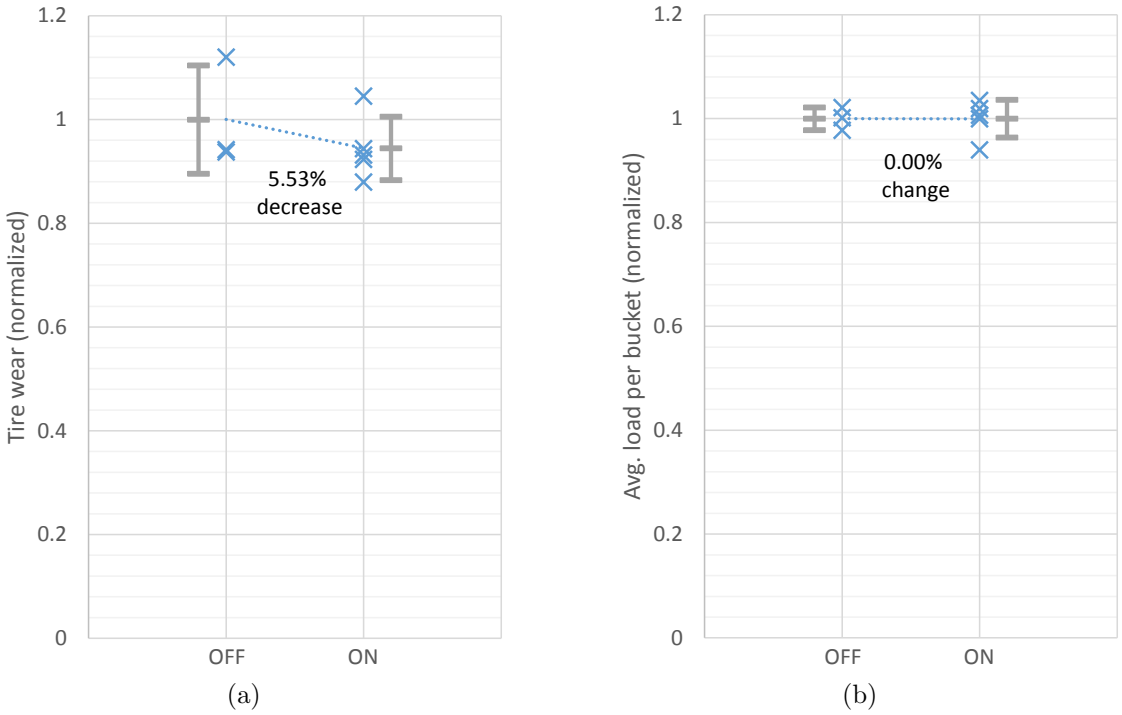

Figure 9.5. Comparison of Productivity Test system performance metrics with traction control OFF and ON, showing (a) total tire wear and (b) material moved per bucket load.

to reduce tire slip overall by 5.53% (Fig. 9.5(a)). Due to the fact that the conditions were not very slick, however, there was no appreciable change in the operator's ability to fill the bucket with material (Fig. 9.5(b)). In fact, the operator was able to over-fill the bucket even without the TC system turned on, so there was not much room for improvement in this area.

As the implemented TC system does not have the capability to decrease engine torque to the wheels, the only method available is to brake wheels which are slipping (see the discussion in Section 4.3). This is adequate for reducing wheel slip and increasing pushing force, but this improvement is paid for by an increase in fuel consumption. Therefore, the system fuel consumption was also monitored and compared for the system with TC on and off (Fig. 9.6). For the Productivity Tests run here, the system saw an average fuel increase of 16.10%. This is a significant increase in fuel consumption, and one which represents a potentially serious barrier to the current system being implemented in a production machine. Therefore, the

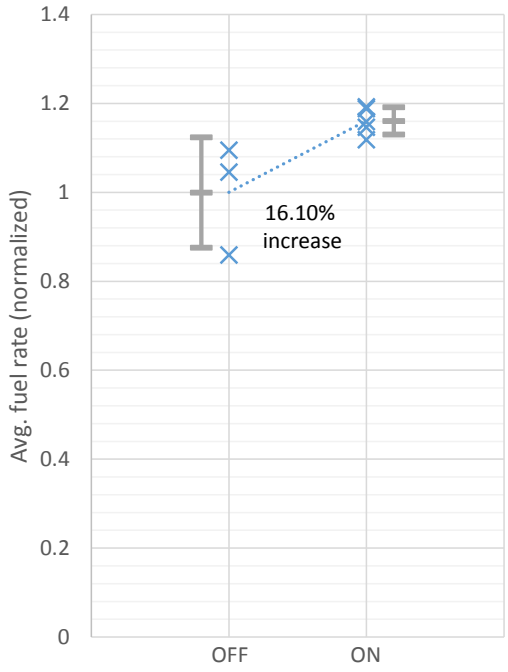

Figure 9.6. Comparison of Productivity Test system fuel consumption with traction control OFF and ON.

Productivity Tests were analyzed further to determine the primary causes contributing to this change.

One investigation which was conducted was an examination of the driving style of the operator with the traction control system on versus when the system was off. In order to do this, an analysis of the machine throttle position during each test was made. From Fig. 9.7, it can be seen that with the TC system on, the operator spends much more time with the throttle at a very high position (90–100%), which would necessarily correlate to a higher fuel consumption. It is clear that there is a significant difference in driving style for these tests. This trend holds even for a comparison of individual tests; that is, tests with a higher throttle distribution consistently have higher fuel consumption values.

Despite this discovery, the reason for this driving style disparity is not so easily understood. It is possible that a small amount of excess braking may have caused the operator to subconsciously increase his throttle position, though that is unlikely to explain most of the throttle increase. Instead, it was noted by those observing the test that the operator seemed

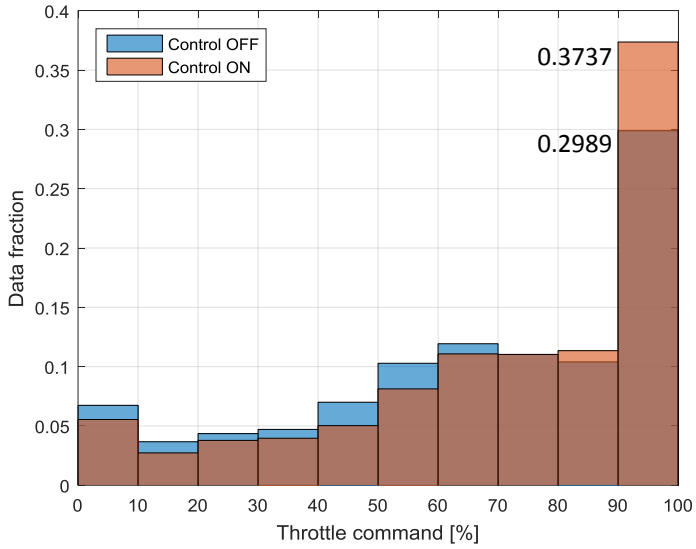

Figure 9.7. Throttle command distribution comparison for Productivity Tests.

less cautious with the system when the TC system was on. The fact that the system was able to control the tire slip automatically allowed the operator to drive at a very high throttle position without worrying about causing the system to slip. What results is a significantly higher throttle curve with much time spent at full throttle (100% command). Nevertheless, the fuel consumption increase is troubling, and further investigation of the causes should be conducted if the TC system is to be implemented on a production vehicle.

9.3 System Improvements

Whatever the cause of the increased fuel consumption, it was clear that in order for the traction control system to make sense for a production machine, this result would have to be improved. Therefore, several investigations were made into different aspects of the system where fuel economy could be improved.

The primary system to be considered was engine torque limitation. This concept is developed in detail in Section 4.4 of this work, and its implementation on the prototype system is described in Section 7.4. Several iterations of tests were run with different parameters to ensure that the the traction control performance of the system (i.e. its ability to keep tires

from slipping against the ground) was not reduced too greatly. In the end, when the torque reduction control parameter values were set correctly, the system performance with torque reduction and only braking a maximum of two wheels was essentially equivalent to that of the system allowed to brake all four wheels. Furthermore, the reduction in engine torque represents a reduction in fuel demand (on top of the obvious reduction in brake wear), so this methodology truly does represent a more efficient method for reducing wheel slip.

Electro-hydraulic control of the brakes also affords some other system-level benefits. Among these is the ability to automatically brake the system to a stop during reversing maneuvers. It is quite common in high-volume applications for operators to drive construction machines recklessly. This is evident in the so-called "shuttle shift" maneuver, wherein the operator shifts the machine from forward to reverse (or vice versa) while the it is still moving (i.e. not braking the system to a stop or actuating the clutch before the reversing command is given). In many applications with mechanical transmissions, this maneuver results in the system being brought to a halt using the engine as a brake. What results is added complexity in the engine controller and an increase in fuel consumption during the reversing cycle. In operations like wheel loader loading cycles, this fuel increase can become quite significant, as the machine direction is changed very frequently. By using the electro-hydraulic brakes, however, it is possible to command the braking pressure to assist the machine in slowing down. In addition, it is possible to override the operator direction commands and force the vehicle into neutral while the brakes slow it down, and then shift back into gear once it has come to a halt. This strategy, which also shows a marked improvement in fuel economy for reversing cycles, was implemented on the prototype machine, as well.

After conducting many proof-of-concept tests, wherein the parameters for all of the added control considerations were tuned to proper values, preliminary data indicated that the adjusted system was capable of giving a performance similar to the stock system, with the added benefits of traction control. Therefore, more performance tests were ordered to assess the overall potential of the final control system.

9.4 Final Performance Tests

The final set of Productivity Tests had a very similar setup to the those described in Section 9.1. However, these tests were much more rigorous, as they were considered the final proof-of-concept tests for the traction control system. A series of eighty-six repetitions was run, testing eighteen different configuration combinations of the traction control system. Two different expert operators were used to account for variations in driving style. In order to ensure consistent behavior of the material in the work pile, backfill loaders were used to replenish material that had already been dumped (see Fig. 9.8). Each test lasted thirty minutes, and all available data was stored for analysis. Operator subjective feedback was also recorded for many tests, to assess how the perceived performance was changing.

Figure 9.8. Test setup for final Productivity Tests, showing test wheel loader (left) and backfill loader (top) at the digging pile with dumping area to the right.

Some of the final results for this series of tests are shown in Fig. 9.9. There are two different metrics shown here. First, Fig. 9.9(a) shows the average fuel consumption over the tests. This comparison is quite significant, as it shows that the changes implemented on the system (mentioned in Section 9.3) were able to eliminate the additional fuel consumption seen in the previous Productivity Tests with the traction control system active. In fact, this data shows a small decrease in fuel consumption (around 2.6%) with the traction control

on. There were never any specific claims made about the traction control system impact on fuel consumption. However, it appears that the engine torque control and added benefits of the electro-hydraulic braking system have resulted in a system which is somewhat more fuel efficient in a loading cycle operating mode. The impact on other operations may not reflect exactly the same trend; however, this is generally viewed as the most common application for wheel loaders.

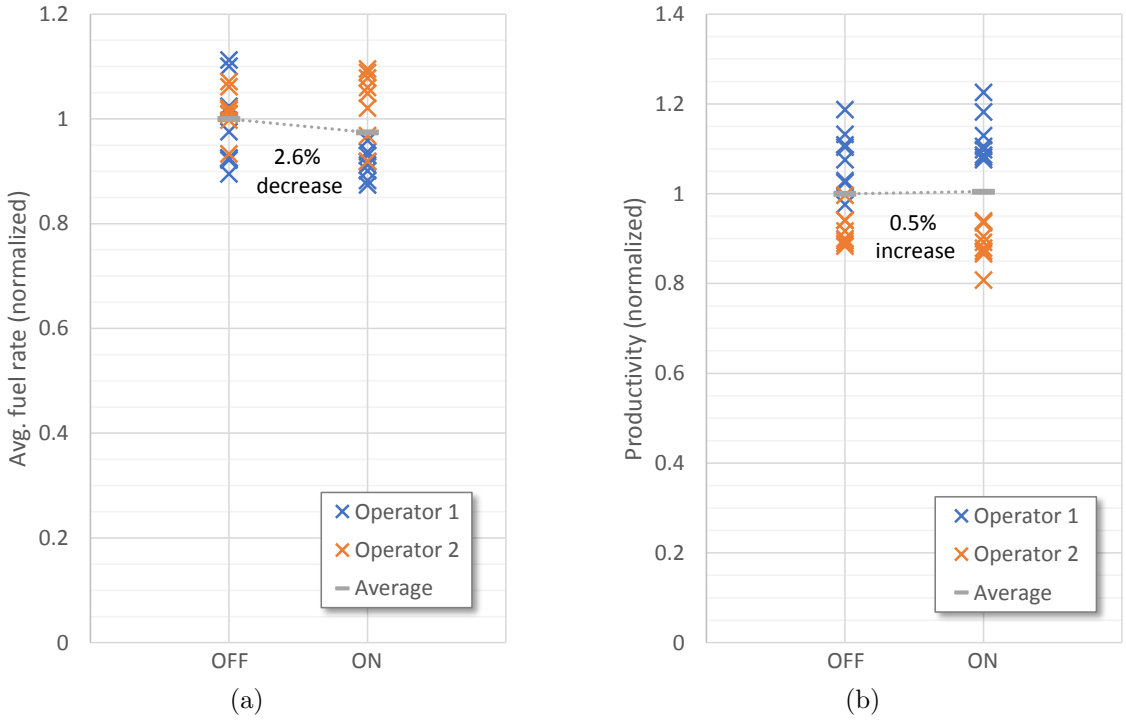

Figure 9.9. Comparison of final Productivity Test system performance metrics with traction control OFF and ON, showing (a) fuel consumption and (b) material moved per unit fuel.

With that in mind, Fig. 9.9(b) shows the overall impact of the TC system on the overall productivity of the wheel loader (represented here as the mass of material moved per unit fuel consumed). In this case, the productivity with TC active is essentially the same as that with TC inactive (showing an average increase of just 0.5%, well within the margin of error). As the fuel consumption was decreased with TC on, this means that, on average, less material was moved during the tests with TC active. The material moved per load

was essentially unchanged; therefore, this impact on material per test represents a slightly slower cycle time with TC on. This could indicate a small amount of excess braking in the system, or a change in operator behavior for some other reason. On top of this, Fig. 9.9(b) reveals another curious phenomenon with respect to differences in operator. Whereas Operator 1 showed a clear increase in overall productivity with TC active, it appears that the traction control actually caused a decrease in Operator 2's productivity. It is unclear why the operators would respond differently to the TC system, but this does represent a significant result. This would indicate that there is an interaction between operator and system type. There is likely no single solution which will work best for every single operator, as their preferences for machine behavior are too subjective.

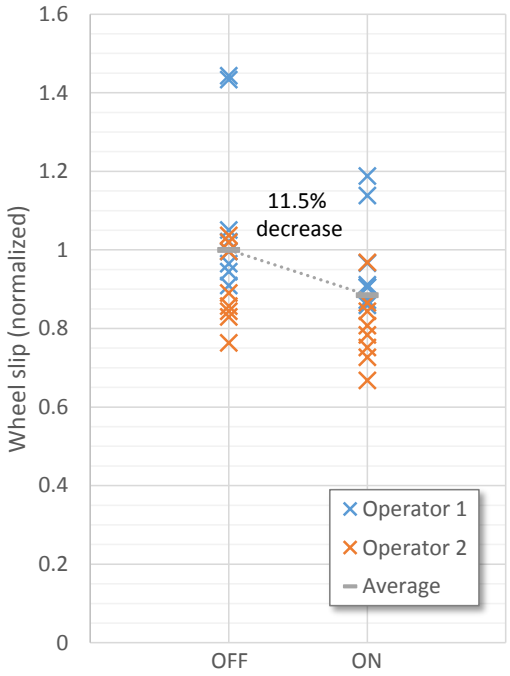

Figure 9.10. Comparison of final Productivity Test system wheel slip with traction control OFF and ON.

While these results are positive, they are pointless if the traction control does not achieve its primary goal: reduction of wheel slip. To that end, Fig. 9.10 shows the overall tire slip from these Productivity Tests. There is a wide range of values, as each category (TC OFF and TC ON) had multiple different configurations within it. Nevertheless, the data

shows an overall 11.5% decrease in wheel slip over the entire range of tests. This is a very significant result, because, along with the data in Fig. 9.9, it indicates that the traction control system is indeed capable of significantly reducing wheel slip for this construction machine without negatively impacting other performance metrics. Furthermore, operator perception was generally positive when using TC. Of course, there are other factors, such as brake consumption, which have not been heavily considered in this work. Therefore, other considerations would likely need to be made before this system is implemented on a production vehicle.

9.5 Recommendations

Based on the results of these field tests, certain considerations should be made when designing a traction control system for production machines. First, it is clear that steering angle compensation crucial to the correct operation of a TC system. It should not be neglected in any final implementation. Also, wheel speed sensor signals should be well-conditioned and accurate. Noise and other deviations from the correct value can cause major issues with the EH braking system locking up the wheels of the machine or not braking properly once the wheels start slipping. On top of this, it is important to make sure the controller logic is robust enough to handle many different situations. If necessary, a PID-type structure should incorporate some sort of reset logic for the integral term, to avoid issues occurring from integrator wind-up.

Finally, the TC system should be thought of, not as a replacement for the operators own feel and perception of the machine, but as an assisting mechanism which works alongside the operator in reducing tire slip. For conditions in which the tires are not very prone to slipping, an experienced operator will be able to consistently keep the wheels from slipping when digging by adjusting the throttle in real time based on feel. Traction control should only be necessary in slippery conditions, or when the operator's intuition is not capable of maintaining traction without assistance. For inexperienced operators, the decrease in wheel slip with TC active may be more noticeable, and therefore improvements can be expected in terms of tire wear and machine controllability. However, this improvement may require some

price in brake consumption. For this reason, the traction control system is not a panacea that can improve every situation, but it is still likely useful under the correct circumstances.

10. CONCLUSIONS

This work outlines the activities involved in the development and implementation of an electro-hydraulic traction control system designed for heavy duty off-road vehicles. Using a wheel loader as a reference machine, the TC system was developed through constructing a simulation model, defining a proper control structure, and conducting a series of rigorous experimental tests.

A basic vehicle dynamics simulation model was developed, accounting for both the linear motion of the vehicle and the rotational motion of the wheels. Some aspects of the vehicle drivetrain were modeled to allow the system to represent the interactions among the four wheels. Also, the shifting normal force distribution between the front and rear axles was defined for the case the reference wheel loader accelerating and the case where it is pushing against an arbitrary resistive force, which could represent a work pile or other feature which impedes the machine motion. The model was further improved by the incorporation of a wheel force model which defines a relationship between wheel slip and tire-ground friction for generating force. The slip-friction model was created using experimental data and varied with the ground type simulated.

Next, an appropriate controller for the system was designed, taking into account the peculiarities and difficulties of traction control. A relatively simple controller using a PID structure with incorporated logic was developed for implementation on the real-world machine. The resulting Augmented PID control incorporated considerations for reducing the effects of integrator wind-up, among other factors. The vehicle drivetrain architecture was then considered in order to determine the most efficient TC strategy for this particular reference machine. It was determined that the most logical approach should incorporate both independent wheel braking and an engine torque reduction strategy to allow the TC system to maximize pushing effort without excessive braking. A method for accounting for the vehicle steering angle was also discussed. Steering correction is a necessity for any correctly functioning TC system, so it was incorporated in the prototype controller implementations. The torque limitation strategy was also incorporated in final versions of the controller to

improve system efficiency. Considerations were also made for mitigating the negative impact of system time delay, but these were never needed for the prototype controller.

Other controller schemes were investigated from a more theoretical standpoint. These include controllers which run real-time optimization strategies in concurrently with the controller to find the optimal operating point for the system. Multiple nonlinear controller were also described and simulated which are capable of improving the performance of the system. These controllers are significantly more complex and require more sensors than the PID structure, so they were not intended for final implementation on a production machine. However, the self-tuning controller and a nonlinear (sliding mode) controller were tested on the prototype system, which has a wider array of sensors and more computing power than a final production vehicle. In both cases, the laboratory testing reflected the simulated performance of these systems, showing potential for increasing the output of future TC designs.

Extensive work was done to assess the real-world TC system performance in an attempt to demonstrate its feasibility for application in the real world. The reference wheel loader's braking system was modified to allow for independent electro-hydraulic control of the brakes, and the prototype vehicle itself was instrumented with a wide array of sensors to generate adequate data for capturing system performance. On top of this, two different estimation schemes were designed to improve system reliability and data acquisition capability. The first was a vehicle velocity estimate, designed as a logic-based structure which leveraged the wheel speeds to calculate an accurate estimate of linear vehicle velocity. The resulting estimate was more accurate and better conditioned than a commercially-available Doppler effect sensor. An implement force estimator was also developed based on a kinematic analysis and a force balance analysis of the wheel loader implement system. Component position measurements and the operating pressure readings of the hydraulic cylinders for boom lift and bucket tilt were used to generate an estimate of the horizontal and vertical forces acting on the wheel loader bucket. This method was able to achieve payload weight estimates within 2.5% of their actual values, which is comparable to state-of-the-art load weighing systems without the same drawbacks.

Once the updates had been made, several different tests were run. The first tests run were designed to characterize the slip-friction behavior of the tire-ground surface interface for different operating conditions. The data generated in these tests guided the construction of the vehicle model to allow the resulting simulations to be as accurate as possible. Next were a series of tests conducted at Maha Laboratory, in which the wheel loader pushed against a barrier, causing the tires to slip. These tests gave an indication of machine pushing force and controller behavior under very controlled circumstances. The performance comparison showed an increase in average pushing force of at least 10% in every case, with the potential for much higher increases (up to 60%) at certain operating conditions. This increase in force also came with a decrease in tire slip of more than 45% in every case (with some cases as high as 73%). Of course, some price was paid in brake wear to achieve these results.

Field tests were also run which were more in line with the real-world operation of the wheel loader. There were several different types of experiments designed to assess different aspects of the TC system performance. Perhaps the most important of these tests were the Productivity Tests, which analyzed the impact of traction control on the entire system's performance. These tests showed that the system had some clear advantages in specific applications, and wheel slip was reduced in every case. However, fuel consumption represented a critical issue which had to be addressed in order for the TC system to be viable for real-world implementation. Therefore, more work was done to update the controller, including incorporating an engine torque reduction strategy. The updated system showed a nearly 12% reduction in tire slip and even managed to reduce fuel consumption by 2.6% with respect to the stock system.

It is clear that the incorporation of a traction control system into the reference machine driveline can have significant benefits to its overall operation. Data from various simulations and tests has consistently shown a reduction in tire wear and an increase in pushing force with a proper control strategy. However, the specifics of this control system are crucial, as a poorly designed controller can have significant negative impacts on other performance parameters. Therefore, it is imperative that the future of traction control system design for heavy-duty equipment be done by starting from an examination of the system dynamics, then building a suitable control structure that works alongside the existing system architecture to provide

the most efficient slip reduction scheme possible. Vehicle behavior and limitations, typical operating conditions, and operator perception must all be considered in order to see the best results for any given machine. As long as this paradigm is maintained, traction control should have a place among the assistive technologies required for optimal performance of construction equipment.

10.1 Novel Contributions

There were several objectives laid out for this work, each of which forms a novel contribution to the field. The first goal was the development of a prototype system capable of actuating brakes independent of each other, for implementation of the traction control system. This was accomplished through modification of a wheel loader provided to the research group by an industry sponsor. The resulting prototype incorporated accurate control of the braking pressure at all four wheels independently, along with other features like engine torque limitation. This allowed the various controllers developed to be enacted on a physical system to test their real-world capabilities.

Another goal was the characterization of the tire-ground surface interaction for various conditions, as the performance of the traction control system hinges on this relationship. By using a method involving a state estimator running during repetitive acceleration cycles, slip-friction data was generated for several different operating points, including different ground conditions and tire pressure levels. This information was used to generate an accurate simulation model, which successfully mimics the motion of the prototype vehicle. This model was a crucial component to the research, as all controllers were first tested in simulation to assess their capabilities before ever being applied to the real-world machine.

With regard to the controller itself, it was desired to create and simulate control structures which maintained optimal tire slip in the machine, along with a self-tuning controller. Multiple different algorithms were formulated using various aspects of control theory, from simplistic PID structures to complex adaptive methods. Optimization methods like the extremum-seeking algorithm and the quasi-Newton approach were also incorporated into the controller to determine the best slip ratio setpoint for a given ground condition during

operation. The controllers were able to significantly reduce tire slip in both simulation and experiments, while also increasing the tractive force of the vehicle. Furthermore, the online optimization was able to quickly and consistently react to shifting ground conditions and increase the output of the traction control system in real time.

With the control structure in mind, another goal for the research was a higher-level algorithm capable of controlling wheel slip in the most efficient manner. As the wheels are mechanically linked to each other, the vehicle driveline architecture was considered to determine what the best control action should be for each possible slip condition. In the end, it was decided that the best approach should combine independent wheel braking with engine torque limitation. This investigation was incorporated into the traction control system on the prototype, and as a result, tests which had seen an increase of 16% in fuel consumption with TC active now actually saw a 2.6% decrease under the same conditions.

The final two goals center on defining a metric for numerically comparing experimental results for traction control and then drawing up and executing the tests to generate that data. For laboratory tests, five objective functions were formulated, two for assessing vehicle pushing force, two for assessing tire slip, and one for quantifying brake wear. The tests themselves were conducted using a barrier designed and constructed specifically for this work, with provisions made for consistently altering ground condition. The objective functions provided repeatable, intuitive comparisons among test cases, showing that traction control is capable of increasing pushing force by 10% to 60%, while decreasing tire wear by between 45% and 73%. Further tests were also designed for more realistic productivity cycles run in the field. For these tests, the metrics quantified tire slip, fuel consumption, and material moved per bucket load. Again, these parameters were sensible and consistent, and the TC system reduced tire slip in these tests by 12% with a 2.6% decrease in fuel consumption and essentially no impact on bucket load. All tests were run with multiple operators and many repetitions to mitigate the impact of operator variability.

11. FUTURE WORK

The research presented in this work is an extensive investigation into traction control systems applied to heavy-duty off-road machinery. The focus was on describing the philosophy of proper TC design for these types of vehicles, from fundamental modeling approaches to more practical application-specific considerations. However, there remain many areas for expanding the scope of this work. Some of these topics are expansions on concepts already explored in this work, while others represent altogether new ideas which leverage the knowledge gained from the present study to push the boundaries of the field.

The first topic of interest is an investigation into the impact of traction control on the longevity of the friction brakes. The only consideration made for assessing brake wear in this model is a measurement of brake temperature. It was thought that excessive braking would cause the temperatures in the brake housing to increase at a more rapid rate than normal. While this is true, it is hardly a direct measurement, and the model relating true brake wear to temperature rate of change is not well understood. It would be worthwhile to conduct a study on how overall brake life is affected by the introduction of a TC system. However, it would take a very long time to generate all the comparison data needed to assess this impact. Another approach would be to generate a simulation model based on instantaneous brake wear. Understanding brake consumption is a quite complex modeling problem, as there are a great many factors which contribute to this phenomenon. Brakes abrade differently depending on the type of braking system, the materials involved, rates of rotation, application pressure, operating temperatures, and more. There exists a body of literature about brake wear (see for example [70]), which include discussions on metallurgical properties and other design features. All of this goes well beyond the scope of the current work, but these are important considerations for the end user of the construction machine, as increased brake consumption ultimately translates into increased operating cost.

As has been stated throughout this work, the heavy emphasis on traction control design has largely ignored the development of another similar system, the anti-lock braking system (ABS). ABS is in many ways the mirror image of traction control, releasing the brakes when wheels begin to lock up, instead of actuating the brakes when the wheels begin to spin

out. Much of the modeling and controls work which went into designing the TC system here can be directly applied to anti-lock braking. For construction equipment, the speeds of travel tend to be low, with very little danger of wheel lockup during braking. For this reason, ABS is not a priority for these machines in the same way it is for on-road passenger vehicles. Furthermore, ABS requires additional fail-safe considerations, as overriding the vehicle brakes to decreasing braking action is inherently less safe than increasing that action. Therefore, great care needs to be taken to ensure that the machine is never placed in an unsafe operating condition by the action of the controller. Finally, the electro-hydraulic braking circuit discussed in this work is not capable of enacting an ABS on the prototype machine. Therefore, this architecture needs to be redesigned in order for ABS to be tested alongside traction control. Should this work be undertaken, however, it is likely that ABS will produce some positive results in terms of safety or productivity for heavy-duty machines. For examples of recent advances in system design for ABS, see [71–73].

Now that a good estimate of work pile resistance forces has been achieved, that information can be leveraged for even more advanced machine behavior optimization. By examining the forces generated under specific digging configurations, an optimal excavation trajectory can be plotted for the system. This optimization can be based on any number of objective functions: for instance, minimizing the amount of work needed to achieve a fully-loaded bucket or minimizing the maximum engine torque required throughout the digging maneuver. Along the lines of the present work, a digging trajectory could even theoretically be determined which maximizes the tractive effort at the wheels. By updating the prototype data acquisition code accordingly, it should be possible to implement a real-time optimization for excavation trajectory on the prototype machine. This work should take into account previous research into excavation trajectory planning for earthmoving machines, such as that described by Singh [74]. Some of these investigations have already been applied to wheel loaders [75], while others focused on excavators [76] and rope shovels [77]. The incorporation of such an online optimization could lead to increased machine productivity, and this work is especially relevant as an ever-increasing number of industrial processes are being automated.

Finally, the vehicle system architecture itself should be considered. The scope of this work was limited to retrofitting an electro-hydraulic braking circuit onto a stock vehicle, without

modifying the basic drivetrain of the machine. However, the increasing prevalence of alternative drivetrain designs such as distributed electric motors, hydrostatic and hydraulic hybrid transmissions, and others suggest that those systems may have advantages over the conventional mechanical drivetrain considered in this work. In fact, architectures which allow for independent control of the torques at each wheel do not require analysis like that in Section 4.3, as the torque at each wheel can be optimized separately. Furthermore, vehicles incorporating hybrid drivetrain technologies commonly allow for energy recovery when decelerating. This is typically more fuel efficient then using wheel braking for traction control, as friction brakes simply convert wheel rotational energy into heat instead of capturing that energy for use later. By comparing the control capabilities of the other systems with the existing powertrain, their advantages could be examined more thoroughly. Currently, these alternative drivetrain technologies are uncommon on heavy-duty machines as large as the reference wheel loader used in this work, but it is quite possible that they will become more widespread in the near future, especially if they show significant benefits when combined with advanced technologies like traction control.

There is still much room for innovation in the field of heavy-duty off-road equipment, and the demands of consumers and production trends will continue to push these technologies forward. The machines of today are more intelligent, more efficient, more reliable, and safer than their predecessors of only a few years ago, and the systems of tomorrow will be even better still. As long as there continues to be a market for expansion and improvement of industries and communities, construction machines will be pushed to their limits. Correspondingly, the technology incorporated in these systems must rise to meet the challenge.

REFERENCES

[1] J. Hosseini, M. A. Cabo, N. J. Rytter, and A. G. Verheyen. Method and apparatus for controlling differentially driven wheel-slip for an articulated machine, 5,535,124. July 1996.

[2] J. Sjögren. Anti spinning device, a vehicle comprising the device and a method for reducing slip during advancing of a vehicle, US 8226177 B2. July 2012.

[3] S. M. Savaresi and M. Tanelli. *Active Braking Control Systems Design for Vehicles*. Advances in Industrial Control. Springer London, London, 2010. DOI: 10.1007/978-1-84996-350-3.

[4] K. Reif, editor. *Brakes, Brake Control and Driver Assistance Systems*. Springer Fachmedien Wiesbaden, Wiesbaden, 2014.

[5] M. H. Cardon, G. B. Hickner, and R. W. Rothfusz. Development and Evaluation of Anti-Lock Brake Systems. SAE Technical Paper 760348, February 1976.

[6] W. C. Eddy. Traction control apparatus, US3981545 A. September 1976.

[7] Robert Bosch GmbH, editor. *Bosch Automotive Electrics and Automotive Electronics*. Bosch Professional Automotive Information. Springer Fachmedien Wiesbaden, Wiesbaden, 5th edition, 2014. DOI: 10.1007/978-3-658-01784-2.

[8] P. A. Fawkes and S. M. Dunning. Traction control system, US6755488 B2. June 2004.

[9] S. K. Mohan and R. C. Williams. A Survey of 4WD Traction Control Systems and Strategies. SAE Technical Paper 952644. SAE International, November 1995.

[10] D. Piyabongkarn, J. Lew, J. Grogg, and R. Kyle. Stability-Enhanced Traction and Yaw Control using Electronic Limited Slip Differential. SAE Technical Paper 2006-01-1016, Detroit, MI, USA, April 2006. SAE International.

[11] Y. Lin and S. Anwar. A Traction Enhanced On-Demand All Wheel Drive Control System for a Hybrid Electric Vehicle. SAE Technical Paper 2007-01-0299, Detroit, MI, USA, April 2007. SAE International.

[12] S. Kuntanapreeda. Traction Control of Electric Vehicles Using Sliding-Mode Controller with Tractive Force Observer. *International Journal of Vehicular Technology*, 2014:1–9, 2014.

[13] B. J. Holt, J. E. Jensen, S. Marathe, and S. A. Marks. Electronic traction control system, US 6631320 B1. October 2003.

[14] K. Uematsu, K. Hatake, and A. Nomura. Vehicle speed estimator and traction control device, US 8538635 B2. September 2013.

[15] Z. Fan, Y. Koren, and D. Wehe. Traction Control and Modeling of Tracked Vehicles. SAE Technical Paper 942375, São Paulo, Brazil, November 1994. SAE International.

[16] S. Yang, Y. Lu, and S. Li. An overview on vehicle dynamics. *International Journal of Dynamics and Control*, 1(4):385–395, December 2013.

[17] T. D. Gillespie. *Fundamentals of Vehicle Dynamics*. Society of Automotive Engineers, Warrendale, PA, 1992.

[18] R. N. Jazar. *Vehicle Dynamics: Theory and Application*. Springer, New York, 2nd edition, 2014.

[19] R. Rajamani. *Vehicle Dynamics and Control*. Mechanical Engineering Series. Springer, New York, NY, 2nd edition, 2012.

[20] M. Abe and W. Manning. *Vehicle Handling Dynamics: Theory and Application*. Butterworth-Heinemann, Amsterdam, 1st edition, 2009.

[21] M. Guiggiani. *The Science of Vehicle Dynamics*. Springer Netherlands, Dordrecht, 2014. DOI: 10.1007/978-94-017-8533-4.

[22] F. Jamzadeh, T.-M. Hsieh, and K. Struthers. Dynamic Simulation Modeling For Heavy Duty Automatic Transmission Control Development. SAE Technical Paper 922441, Toledo, OH, USA, November 1992. SAE International.

[23] M. Tanelli, M. Corno, and S. M. Savaresi, editors. *Modelling, Simulation and Control of Two-Wheeled Vehicles*. John Wiley & Sons, Ltd, Chichester, UK, February 2014.

[24] J. Y. Wong. *Theory of Ground Vehicles*. Wiley, Hoboken, N.J, 4th edition, 2008.

[25] A. F. Andreev, V. I. Kabanau, and V. V. Vantsevich. *Driveline Systems of Ground Vehicles: Theory and Design*. Ground Vehicle Engineering Series. CRC Press, Boca Raton, 2010.

[26] D. Schramm, M. Hiller, and R. Bardini. *Vehicle Dynamics: Modeling and Simulation*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.

[27] M. Gipser. FTire: a physically based application-oriented tyre model for use with detailed MBS and finite-element suspension models. *Vehicle System Dynamics*, 43(Sup1):76–91, January 2005.

[28] H. Willumeit and F. Böhm. Wheel Vibrations and Transient Tire Forces. *Vehicle System Dynamics*, 24(6-7):525–550, July 1995.

[29] H. Dugoff, P. S. Fancher, and L. Segel. An Analysis of Tire Traction Properties and Their Influence on Vehicle Dynamic Performance. SAE Technical Paper 700377, February 1970.

[30] H. B. Pacejka. In-Plane and Out-of-Plane Dynamics of Pneumatic Tyres. *Vehicle System Dynamics*, 10(4-5):221–251, September 1981.

[31] G. Mastinu, S. Gaiazzi, F. Montanaro, and D. Pirola. A Semi-Analytical Tyre Model for Steady- and Transient-State Simulations. *Vehicle System Dynamics*, 27(Sup1):2–21, January 1997.

[32] J. Svendenius. *Tire Models for Use in Braking Applications*. Licentiate thesis, Department of Automatic Control, Lund Institute of Technology (LTH), 2003.

[33] M. Gipser. FTire - the tire simulation model for all applications related to vehicle dynamics. *Vehicle System Dynamics*, 45(Sup1):139–151, January 2007.

[34] J. E. Bernard and C. L. Clover. Tire Modeling for Low-Speed and High-Speed Calculations. In *New Developments in Vehicle Dynamics, Simulation, and Suspension Systems*, pages 85–94, Detroit, MI, USA, February 1995. Society of Automotive Engineers.

[35] P. Lugner, H. Pacejka, and M. Plöchl. Recent advances in tyre models and testing procedures. *Vehicle System Dynamics*, 43(6-7):413–426, June 2005.

[36] A. J. C. Schmeitz and W. D. Versteden. Structure and Parameterization of MF-Swift, a Magic Formula-based Rigid Ring Tire Model. *Tire Science and Technology*, 37(3):142–164, September 2009.

[37] F. Braghin and E. Sabbioni. A Dynamic Tire Model for ABS Maneuver Simulations. *Tire Science and Technology*, 38(2):137–154, June 2010.

[38] R. van der Steen, I. Lopez, H. Nijmeijer, A. Schmeitz, and B. de Bruijn. Experimental and Numerical Study of Friction and Braking Characteristics of Rolling Tires. *Tire Science and Technology*, 39(2):62–78, June 2011.

[39] H. B. Pacejka and E. Bakker. The Magic Formula Tyre Model. *Vehicle System Dynamics*, 21(S1):1–18, January 1992.

[40] H. B. Pacejka and I. J. M. Besselink. Magic Formula Tyre Model with Transient Properties. *Vehicle System Dynamics*, 27(S1):234–249, January 1997.

[41] H. B. Pacejka and I. Besselink. *Tire and Vehicle Dynamics*. Elsevier/Butterworth-Heinemann, Amsterdam, 3rd edition, 2012.

[42] I. Emri and A. Voloshin. *Statics: Learning from Engineering Examples*. Springer New York, New York, NY, 2016.

[43] H. W. Müller, W. G. Mannhardt, and J. H. Glover. *Epicyclic Drive Trains: Analysis, Synthesis, and Applications*. Wayne State University Press, Detroit, 1982.

[44] M. M. Tinker. *Wheel Loader Powertrain Modeling for Real-Time Vehicle Dynamic Simulation*. Master's thesis, The University of Iowa, Iowa City, July 2006.

[45] E. McKyes. *Soil Cutting and Tillage*. Number 7 in Developments in Agricultural Engineering. Elsevier, Amsterdam, 1985.

[46] M. D. Worley and V. LaSaponara. Development of a Simplified Load-Cycle Model for Wheel Loader Design. In *Proceedings of IMECE2006*, volume 2006, pages 641–654, Chicago, IL, USA, 2006. ASME.

[47] H. Kang, W. Jung, and C. Lee. Modeling and Measurement of Payload Mass of the Wheel Loader in the Dynamic State based on Experimental Parameter Identification. SAE Technical Paper 2016-01-0469, April 2016.

[48] A. Shatters. Method and system for estimating payload weight with tilt position compensation, US 9464403 B2. October 2016.

[49] W. J. Spencer. *Fundamental Structural Analysis*. Springer New York, New York, NY, 1988.

[50] N. S. Nise. *Control Systems Engineering*. Wiley, Hoboken, NJ, 6th edition, 2011.

[51] W. S. Levine, editor. *The Control Handbook*. Electrical Engineering Handbook. CRC Press, Boca Raton, Fla., 2nd edition, 2011.

[52] K. B. Ariyur and M. Krstić. *Real-Time Optimization by Extremum-Seeking Control*. John Wiley & Sons, Inc., Hoboken, NJ, 2003.

[53] D. Cristofori, A. Vacca, and K. Ariyur. A Novel Pressure-Feedback Based Adaptive Control Method to Damp Instabilities in Hydraulic Machines. *SAE International Journal of Commercial Vehicles*, 5(2):586–596, September 2012.

[54] G. F. Ritelli and A. Vacca. Energy Saving Potentials of a Novel Electro-Hydraulic Method to Reduce Oscillations in Fluid Power Machines: The Case of a Hydraulic Crane. *SAE International Journal of Commercial Vehicles*, 6(2):269–280, September 2013.

[55] P. V. Osinenko, M. Geissler, and T. Herlitzius. A method of optimal traction control for farm tractors with feedback of drive torque. *Biosystems Engineering*, 129:20–33, January 2015.

[56] A. Alexander, A. Sciancalepore, and A. Vacca. Online Controller Setpoint Optimization for Traction Control Systems Applied to Construction Machinery. In *Proceedings of the 2018 Bath/ASME Symposium on Fluid Power and Motion Control*, Bath, UK, September 2018. ASME.

[57] M. Schreiber and H. Kutzbach. Comparison of different zero-slip definitions and a proposal to standardize tire traction performance. *Journal of Terramechanics*, 44(1):75–79, January 2007.

[58] M. Schreiber and H. Kutzbach. Influence of soil and tire parameters on traction. *Research in Agricultural Engineering*, 54(No. 2):43–49, June 2008.

[59] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.

[60] K. Lange. *Optimization*, volume 95 of *Springer Texts in Statistics*. Springer New York, New York, NY, 2013.

[61] M. Schinkel and K. Hunt. Anti-lock braking control using a sliding mode like approach. pages 2386–2391 vol.3. IEEE, 2002.

[62] J. Li, Z. Song, Z. Shuai, L. Xu, and M. Ouyang. Wheel Slip Control Using Sliding-Mode Technique and Maximum Transmissible Torque Estimation. *Journal of Dynamic Systems, Measurement, and Control*, 137(11), August 2015.

[63] H. Lee and M. Tomizuka. Adaptive Vehicle Traction Force Control for Intelligent Vehicle Highway Systems (IVHSs). *IEEE Transactions on Industrial Electronics*, 50(1):37–47, February 2003.

[64] T. Nakakuki, T. Shen, and K. Tamura. Adaptive control approach to uncertain longitudinal tire slip in traction control of vehicles. *Asian Journal of Control*, 10(1):67–73, January 2008.

[65] I. D. Landau, R. Lozano, M. M'Saad, and A. Karimi. *Adaptive Control*. Communications and Control Engineering. Springer London, London, 2011. DOI: 10.1007/978-0-85729-664-1.

[66] R. Rajamani, G. Phanomchoeng, D. Piyabongkarn, and J. Y. Lew. Algorithms for Real-Time Estimation of Individual Wheel Tire-Road Friction Coefficients. *IEEE/ASME Transactions on Mechatronics*, 17(6):1183–1195, December 2012.

[67] S. Antonov, A. Fehn, and A. Kugi. Unscented Kalman filter for vehicle state estimation. *Vehicle System Dynamics*, 49(9):1497–1520, September 2011.

[68] H. Hamann, J. K. Hedrick, S. Rhode, and F. Gauterin. Tire force estimation for a passenger vehicle with the Unscented Kalman Filter. In *Intelligent Vehicles Symposium Proceedings*, pages 814–819, Dearborn, MI, USA, June 2014. IEEE.

[69] J. Matuško, I. Petrović, and N. Perić. Neural network based tire/road friction force estimation. *Engineering Applications of Artificial Intelligence*, 21(3):442–456, April 2008.

[70] J. Kukutschová, V. Roubíček, K. Malachová, Z. Pavlíčková, R. Holuša, J. Kubačková, V. Mička, D. MacCrimmon, and P. Filip. Wear mechanism in automotive brake materials, wear debris and its potential environmental impact. *Wear*, 267(5-8):807–817, June 2009.

[71] M. Mirzaei and H. Mirzaeinejad. Optimal design of a non-linear controller for anti-lock braking system. *Transportation Research Part C: Emerging Technologies*, 24:19–35, October 2012.

[72] X. Zhang, Y. Xu, M. Pan, and F. Ren. A vehicle ABS adaptive sliding-mode control algorithm based on the vehicle velocity estimation and tyre/road friction coefficient estimations. *Vehicle System Dynamics*, 52(4):475–503, April 2014.

[73] J. J. Castillo, J. A. Cabrera, A. J. Guerra, and A. Simon. A Novel Electrohydraulic Brake System With TireRoad Friction Estimation and Continuous Brake Pressure Control. *IEEE Transactions on Industrial Electronics*, 63(3):1863–1875, March 2016.

[74] S. Singh and H. Cannon. Multi-resolution planning for earthmoving. In *Proceedings. 1998 IEEE International Conference on Robotics and Automation*, volume 1, pages 121–126, Leuven, Belgium, 1998. IEEE.

[75] P. J. A. Lever. An automated digging control for a wheel loader. *Robotica*, 19(05), September 2001.

[76] Y. B. Kim, J. Ha, H. Kang, P. Y. Kim, J. Park, and F. Park. Dynamically optimal trajectories for earthmoving excavators. *Automation in Construction*, 35:568–578, November 2013.

[77] M. Dunbabin and P. Corke. Autonomous excavation using a rope shovel. *Journal of Field Robotics*, 23(6-7):379–394, June 2006.

VITA

Addison Alexander was born April 18, 1991 in Elizabethtown, Kentucky, USA. After graduating high school in 2009, he attended the University of Kentucky, achieving a Bachelor of Science in Mechanical Engineering in May 2013. In August of that year, he started his doctoral studies as a direct-Ph.D. student of Mechanical Engineering at Purdue University. He received his Masters of Science in Mechanical Engineering from Purdue University in August 2016. Since then, he has continued his doctoral studies, the result of which is the work laid out in this document.

PUBLICATIONS

[1] R. Bianchi, A. Alexander, and A. Vacca. Active Vibration Damping for Construction Machines Based on Frequency Identification. SAE Technical Paper 2016-01-8121, September 2016.

[2] A. Alexander and A. Vacca. Real-Time Parameter Setpoint Optimization for Electro-Hydraulic Traction Control Systems. In *Proceedings of 15th Scandinavian International Conference on Fluid Power*, pages 104–114, Linköping, Sweden, June 2017. Linköping University Electronic Press.

[3] A. Alexander and A. Vacca. Longitudinal vehicle dynamics model for construction machines with experimental validation. *International Journal of Automotive and Mechanical Engineering*, 14(4):4616–4633, December 2017.

[4] A. Alexander, A. Vacca, and D. Cristofori. Active Vibration Damping in Hydraulic Construction Machinery. *Procedia Engineering*, 176:514–528, 2017.

[5] A. Alexander, A. Sciancalepore, and A. Vacca. Online Controller Setpoint Optimization for Traction Control Systems Applied to Construction Machinery. In *Proceedings of the 2018 Bath/ASME Symposium on Fluid Power and Motion Control*, Bath, UK, September 2018. ASME.

[6] A. Alexander and A. Vacca. Traction Control Development for Heavy-Duty Off-Road Vehicles Using Sliding Mode Control. *Control Engineering Practice*. Currently under review.

[7] D. Colombara, A. Alexander, and A. Vacca. Analysis of Dynamics and Force Estimation for a Wheel Loader Implement. Under preparation.