Title
Brake System Modeling and Control

Authors
Hedrick, J. K.
Uchanski, M.

Publication Date
2001-09-01

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Brake System Modeling and Control

J.K Hedrick, M. Uchanski
University of California, Berkeley

California PATH Research Report
UCB-ITS-PRR-2001-25

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department of Transportation, Federal Highway Administration.

Report for MOU 308

September 2001

ISSN 1055-1425

Brake System Modeling and Control

Final Report

PATH MOU 308

J.K. Hedrick
M. Uchanski

Mechanical Engineering Department
University of California at Berkeley
Berkeley, CA 94720

Abstract

MOU 308 provides solutions to two common brake control problems: variable brake torque gain and brake rotor-induced brake torque oscillations. The adaptive control solution for the variable brake torque gain problem is shown to work experimentally, and the algorithm to eliminate brake torque oscillations is demonstrated in simulation. In addition, documentation is provided for differential braking hardware, wheel speed measurement hardware, and a strain-based brake torque sensor that were constructed to test these algorithms.

Keywords: Adaptive Control, Adaptive Feedforward Cancellation, Antilock Braking Unit, Brake Control, Differential Braking, Torque Sensor, Wheel Speed Measurement

Executive Summary

MOU 308 provides solutions to two common brake control problems: variable brake torque gain and brake rotor-induced brake torque oscillations. The adaptive control solution for the variable brake torque gain problem is shown to work experimentally, and the algorithm to eliminate brake torque oscillations is demonstrated in simulation. In addition, documentation is provided for the unique hardware that needed to be constructed to test these algorithms.

To compensate for variable brake pressure to brake torque gain, two different Lyapunov-based adaptation laws—one "smooth" and the other "nonsmooth"—are designed. Experiments using a strain-based brake torque sensor show that both adaptation algorithms reduce tracking error and converge to the correct parameter values.

To eliminate brake rotor-induced vibrations, a nonlinear version of a so-called Adaptive Feedforward Cancellation algorithm is developed and simulated. Roughly, the algorithm continuously approximates the vibrations with several terms of a Fourier Series and then injects a brake pressure input to cancel them.

To test the brake control algorithms presented in this report, a fairly sophisticated vehicle test bed was developed with differential braking hardware, ABS wheel speed sensor conditioning hardware, and a strain-based brake torque sensor. Detailed documentation of this hardware is provided so that future researchers can either make use of the test bed or use the experience gained in its construction to create an improved test vehicle.

Contents

List of Figures

Chapter 1

Introduction

This research report presents findings obtained under PATH MOU 308, a project devoted to developing advanced brake control methods for automated highway systems.

Brake control at the level of precision needed for platooning is a difficult problem for two main reasons. First, the vacuum booster-based actuator that is standard on almost all passenger cars introduces many factors that complicate control into the braking system. These include pure time delays, slow dynamics, and hysteresis. Under MOU 104 and MOU 237, the effect of the booster was studied, and a new brake actuator was constructed that controlled the brake system's hydraulic pressure directly, eliminating the booster and its difficulties from the brake control loop.

The second difficulty with brake control for platooning is the highly variable gain between the brake hydraulic pressure and the brake torque. For example, in a two-hour testing period, the brake pressure to brake torque gain often changes by 50% or more, leading to suboptimal platooning performance. Maciuca partially dealt with this problem under a previous MOU, but his solution could never be fully tested because the proper hardware did not exist. Under MOU 308, we constructed the hardware that was lacking and completely tested the algorithms to compensate for the variable brake gain. Chapter 2 derives the algorithms and describes our tests, the results of which showed remarkable agreement between theory and experimentation.

Chapter 3 addresses another, more subtle, problem in brake control—that of brake rotor vibrations. Simulation results show that nonlinear version of a method called Adaptive Feedforward Cancellation is useful in attenuating

these vibrations. Roughly, the algorithm continuously approximates the vibrations with several terms of a Fourier Series and then injects a brake pressure input to cancel them.

In the process of constructing the hardware needed to get the results of Chapter 2, the red Lincoln Towncar evolved into a fairly sophisticated brake control test bed which is unique to the PATH fleet. To make this test bed accessible to other researchers, Chapter 4 provides detailed documentation of the hardware we added to this vehicle. In addition to "standard" PATH measurements and actuators, the red Towncar can measure brake torque with a strain-based sensor, actuate its brakes separately, and measure its four wheel speeds separately. These new capabilities (as well as more standard capabilities) are accessible through a Matlab interface which automatically generates the vehicle's control code.

Finally, Chapter 5 summarizes the results of work on MOU 308 and provides recommendations for other researchers.

Chapter 2

Experimental evaluation of adaptive brake control laws

2.1 Introduction

"Brake fade"—the situation where elevated temperatures at the brake rotor decrease braking gain—is a phenomenon which most drivers have experienced. In an automated highway, or short-headway automated cruise control, even a moderate change in the brake system gain could lead to unacceptable tracking errors. Even if tracking errors are acceptable, it is often desirable to know the brake system gain in order to estimate other parameters. For example, [7] introduces a novel road friction force observer which requires an estimate of the brake gain to give quantitatively accurate results.

Here, we develop and experimentally demonstrate two versions of an adaptive sliding mode brake controller which compensates for an unknown brake gain. One of the algorithms, originally presented by Maciuca [15] uses the concept of nonsmooth Lyapunov functions [21] to derive an adaptation law which may have more desirable convergence properties than standard adaptation laws.

The remainder of this chapter is organized as follows: Section 2.2 introduces the vehicle model and derives a sliding surface controller. Section 2.3 then introduces two adaptive algorithms—one "smooth," and one "non-smooth"—to compensate for the unknown brake pressure to brake torque gain. The smooth algorithm arises out of a standard quadratic

Lyapunov function, and the nonsmooth algorithm is derived from a Lyapunov function with a discontinuous derivative at the origin. In Section 2.4 a test vehicle with a strain based brake torque sensor is used to experimentally verify the stability and parameter convergence of the adaptive controller. Finally, Section 2.5 offers conclusions and directions for future work.

2.2 Vehicle Model and Control Algorithm

In this section, we present a simplified longitudinal vehicle model and derive a brake controller for velocity profile tracking.

2.2.1 Vehicle Model

An $F = ma$ force balance for the vehicle results in the equation

$$F_{x_f} + F_{x_r} - F_d = ma_x \tag{2.1}$$

with

F_{x_f}	road force on the front wheels
F_{x_r}	road force on the rear wheels
F_d	drag forces due to wind and grade
m	vehicle mass including wheels
a_x	vehicle longitudinal acceleration

Moment balances for the front (not connected to engine) and rear (connected to engine) wheels give

$$J_{w_f}\dot{\omega}_{w_f} = -rF_{x_f} - M_f - T_{b_f} \tag{2.2}$$

and

$$J_r\dot{\omega}_r = T_e - rF_{x_r} - M_r - T_{b_r} \tag{2.3}$$

with

$\dot{\omega}_{w_f}$	angular acceleration of front wheels
$\dot{\omega}_r$	angular acceleration of rear wheels
J_{w_f}	moment of inertia of front wheels
J_r	J of rear wheels, differential, engine, gears
M_f	front wheel rolling resistance moment
M_r	rear wheel rolling resistance moment
T_{b_f}	front brake rotor torque
T_{b_r}	rear brake rotor torque
T_e	torque converter output torque as seen at wheel
F_{x_f}	road force on the front wheels, as above
F_{x_r}	road force on the rear wheels, as above
r	wheel radius

The moment balance for the rear wheel assumes that the entire drive train can be treated as one lumped inertia. Although this is not true in general, especially under high torque conditions, it is a reasonable approximation for generating a relatively simple control-oriented model.

If we assume wheel slip to be negligible, then the kinematic rolling condition gives that $\dot{\omega}_{w_f} = \dot{\omega}_r = a/r$, and we can solve equations 2.2 and 2.3—the torque balances at the wheel—for F_{x_f} and F_{x_f}, substitute into the $F = ma$ equation for the vehicle (equation 2.1), and group terms multiplying a to yield

$$[T_e - T_{b_f} - T_{b_r} - M_f - M_r]/r - F_d = \left(m + \frac{J_{w_f} + J_r}{r^2}\right) a \tag{2.4}$$

If we group the brake torques T_{b_f} and T_{b_r} together into one effective brake torque T_b, lump the rolling resistance moments into an effective rolling resistance moment M, and substitute $J_{w_r} + \frac{J_e}{R_g^2}$ (R_g is the ratio of the wheel rotational speed to the engine rotational speed, J_e is the moment of inertia of the engine, and J_{w_r} is the moment of inertia of the rear wheels) for the lumped wheel/drivetrain/engine inertia J_r, we get

$$\frac{T_e - T_b - M}{r} - F_d = \left(m + \frac{J_{w_r} + \frac{J_e}{R_g^2} + J_{w_f}}{r^2}\right) a \tag{2.5}$$

which simplifies to

$$T_e - T_b - M - rF_d = \frac{1}{rR_g^2}\left(J_e + R_g^2(mr^2 + J_{w_r} + J_{w_f})\right) a \tag{2.6}$$

Defining $\beta = \frac{1}{rR_g^2}\left(J_e + R_g^2(mr^2 + J_{w_r} + J_{w_f})\right)$ gives

$$T_e - T_b - M - rF_d = \beta a \tag{2.7}$$

as our final equation of motion for the vehicle. To simplify notation in the following sections, we define $T_{ext} = T_e - M - rF_d$ to be the sum of the engine, rolling resistance, and wind drag terms to get

$$T_{ext} - T_b = \beta a \tag{2.8}$$

Equation 2.8 serves as the basis of our longitudinal control and parameter adaptation algorithms in the following sections.

2.2.2 Brake Controller

Using the vehicle model developed above, we now design a sliding surface velocity tracking brake controller which we will augment in the next section with adaptation. (For details on sliding control, see [22]). An engine controller and switching logic between throttle and brakes have been previously designed (see, for example, [10]), but are omitted here for brevity.

Unfortunately, the control input for the brakes is not the T_b appearing in the vehicle's equation of motion, $T_{ext} - T_b = \beta a$. Instead, the input is the brake pressure at the master cylinder which, neglecting brake cylinder "pushout pressure" and hydraulic dynamics and assuming low slip, is linearly related to the brake torque:

$$T_b = K_b u \tag{2.9}$$

Until brake pressures are large enough to cause wheel lock-up, the linearity assumption is quite good. However, the gain, K_b, can change by more than 50% under normal driving conditions. A change in the gain due to heat is commonly called "brake fade" and is a noticeable, even dangerous, problem on long downhill sections of road. Water and brake pad wear can also significantly affect the brake gain. For now, we proceed to design the control law assuming that the gain, K_b, is known and then introduce adaptation in the next section to compensate for its excursions from the nominal value.

The surface, S, is defined to be the velocity error, so that when the controller reaches $S = 0$ we have that the actual velocity v equals the desired desired velocity v_{des}:

$$S := v - v_{des} \tag{2.10}$$

To assure that the surface will reach zero, the control, u, is chosen so that the time derivative of a Lyapunov function $V := \frac{1}{2}S^2$ along the closed-loop state trajectories is negative, ie

$$S\dot{S} \leq 0 \tag{2.11}$$

Often, u is chosen with a switching term to assure that $S = 0$ is reached, but to avoid the chattering problems often associated with the switching term, we use an asymptotic approximation to sliding mode and choose u so that for $\lambda > 0$, the derivative of the surface S is $\dot{S} = -\lambda S$, giving $\dot{V} = -\lambda S^2$. To find the control, we substitute the vehicle dynamics of equation 2.8 into the desired surface dynamics $\dot{S} = -\lambda S$:

$$\dot{S} = \dot{v} - \dot{v}_{des} = \frac{T_{ext} - K_b u}{\beta} - \dot{v}_{des} = -\lambda S \tag{2.12}$$

Solving for u gives the rule for choosing brake pressure:

$$u = \frac{1}{K_b}(T_{ext} + \beta(\lambda S - \dot{v}_{des})) \tag{2.13}$$

2.3 Adaptation Algorithms

A myriad of factors—among them temperature, pad material, wear, and moisture—can profoundly affect the gain K_b used in the control law of equation 2.13, leading to unacceptable tracking errors. Thus, two schemes to adapt on the value of K_b were designed. The "smooth" formulation follows a standard design process, and is used as a benchmark. The "nonsmooth" formulation uses the theory of nonsmooth Lyapunov functions and is of interest because it may offer better parameter convergence.

2.3.1 Smooth K_b Adaptation

Adding dynamics to the estimated brake torque gain, $\hat{K}_b$, makes the system two states instead of one, so we construct a Lyapunov function, V_1, out of the velocity error, S, and the parameter error, $\tilde{K}_b := K_b - \hat{K}_b$, both of which we would like to make stable :

$$V_1 = \frac{1}{2}S^2 + \frac{\gamma}{2}\tilde{K}_b^2 \tag{2.14}$$

Taking the time derivative and substituting the definition of S gives

$$\dot{V}_1 = S(\dot{v} - \dot{v}_{des}) - \gamma \tilde{K}_b \dot{\tilde{K}}_b \tag{2.15}$$

Employing equation 2.8—the vehicle equation of motion—and using the control law of equation 2.13 with our current estimate of the torque gain, $\hat{K}_b$, replacing K_b yields

$$\dot{V}_1 = S\left(\frac{1}{\beta}(T_{ext} - \frac{K_b}{\hat{K}_b}T_{ext} - \frac{K_b}{\hat{K}_b}\beta(\lambda S - \dot{v}_{des})) - \dot{v}_{des}\right) \tag{2.16}$$

$$-\gamma \tilde{K}_b \dot{\tilde{K}}_b$$

Using the identity $\frac{K_b}{\hat{K}_b} = \frac{\tilde{K}_b}{\hat{K}_b} + 1$ and assuming that K_b varies slowly so that $\dot{\tilde{K}}_b = -\dot{\hat{K}}_b$ gives

$$\dot{V}_1 = -\lambda S^2 - \frac{S\tilde{K}_b}{\hat{K}_b}\left(\frac{1}{\beta}T_{ext} + \lambda S - \dot{v}_{des}\right) - \gamma \tilde{K}_b \dot{\hat{K}}_b \tag{2.17}$$

so that choosing

$$\dot{\hat{K}}_b = -\frac{S}{\gamma \hat{K}_b}\left(\frac{1}{\beta}T_{ext} + \lambda S - \dot{v}_{des}\right) \tag{2.18}$$

gives $\dot{V}_1 = -\lambda S^2$, implying that V_1 does not grow. This, combined with the fact that V_1 is radially unbounded, allows us to conclude that S and $\tilde{K}_b$ are bounded. Since the system equations are driven by $v_{des}(t)$, they are non-autonomous, so we cannot use invariant set theorems to argue that $S \to 0$. Instead, we use a Barbalat's Lemma-based argument. We first calculate $\ddot{V}_1 = 2\lambda^2 S^2$. Since S is bounded $\ddot{V}_1$ is also bounded, implying that $\dot{V}$ is uniformly continuous. Since V is lower bounded by zero and negative semidefinite and $\dot{V}$ is uniformly continuous, Barbalat's lemma can be used to conclude that $\dot{V} \to 0$ so $S \to 0$. Convergence of $\tilde{K}_b$ is achieved if the system is persistently excited.

2.3.2 Non-Smooth K_b Adaptation

A nonsmooth adaptation algorithm for $\hat{K}_b$ was also derived using a similar procedure, but starting with a nonsmooth Lyapunov function:

$$V_2 := Ssgn(S) + \frac{\gamma}{2}\tilde{K}_b^2 \tag{2.19}$$

The same procedure as above holds everywhere except at $S = 0$ where the $sgn(\cdot)$ function is not defined, giving the adaptation law

$$\dot{\hat{K}}_b = -\frac{sgn(S)}{\hat{K}_b}\left(\frac{1}{\beta}T_{ext} + \lambda S - \dot{v}_d\right) \tag{2.20}$$

Stability analysis is complicated by the slope discontinuity in the Lyapunov function at zero and proceeds according to an extension to Lyapunov theory which treats the derivative of the Lyapunov function at this point as the convex closure of the derivatives on either side. More detailed treatment of these concepts can be found in Maciuca [15] and Shevitz and Paden [21].

2.4 Experiments on Vehicle

The controller and adaptive algorithms developed in the previous section were tested on a 1990 Lincoln Towncar belonging to the California PATH automated highway program. The test vehicle is equipped with a data acquisition and control computer, strain-based brake torque sensor, throttle actuator, brake pressure actuator, wheel speed sensors, manifold pressure sensors, and differential braking capabilities.

In addition to monitoring velocity tracking error, we used the torque sensor to check that the adaptive controller's estimate of the brake torque given by $\hat{K}_b u$ corresponded to the actual torque, indicating that the parameter $\hat{K}_b$ had converged to the correct value.

Unfortunately, the strain-based torque sensor was only installed on the vehicle's *left front wheel*, while the adaptive algorithms developed in the previous section estimate the gain between the brake system pressure and the *total* brake torque on the vehicle—from all four brakes. To assure that the total brake torque on the vehicle equaled the left front wheel torque, braking to the other three wheels was disabled by over-riding the vehicle's ABS controller. The non-symmetrical braking caused a small yaw moment on the vehicle, but decelerations were kept low enough for it to have no noticeable effect.

The same trapezoidal profile for the desired velocity, v_{des}, was used in each test: The vehicle started at 6 m/s and maintained this speed until $t = 5s$, at which point it accelerated for $7.5s$ at $0.8m/s^2$ to reach $12m/s$. It

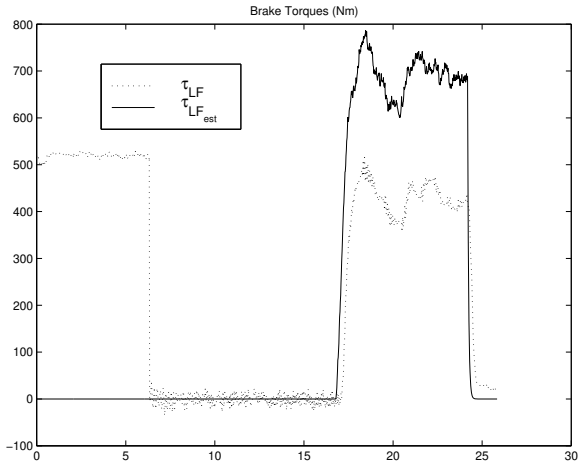

Figure 2.1: Measured and estimated brake torque (Nm) vs. time (s) for baseline test with no adaptation. The initial high value of the measured brake torque is an artifact of the offset removal scheme used for the sensor.

then maintained this speed for $5s$ before decelerating at $0.8m/s^2$ for $7.5s$ to return to $6m/s$. Braking occurs only between $t = 17.5$ and $t = 25$.

2.4.1 Nonadaptive Controller

The first test demonstrates tracking performance of the controller developed in Section 2.2 without parameter adaptation. The actual value of K_b is $0.39\frac{Nm}{kPa}$, but the value which the controller uses is 50% too high at $0.58\frac{Nm}{kPa}$. The surface is the velocity error: $S := v - v_{des}$.

The upper left plot of Figure 2.2 shows the velocity tracking maneuver. Figure 2.1 shows that, due to the erroneous brake torque gain, the brake torque value which the controller uses is significantly higher than the actual torque. This results in a velocity tracking error of approximately $0.2\frac{m}{s}$, as shown in the upper right of Figure 2.2.

2.4.2 Smooth Adaptation

The second test demonstrates tracking performance with smooth parameter adaptation. Like before, the actual value of K_b is $0.39\frac{Nm}{kPa}$, but the value which the controller uses is 50% too high at $0.58\frac{Nm}{kPa}$. The surface in the sliding controller is just the velocity error: $S_1 := v - v_{des}$.

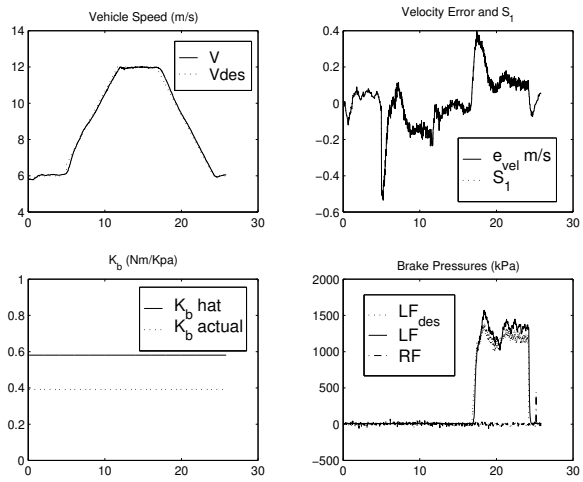

Figure 2.2: Vehicle speed (m/s), velocity error (m/s), K_b (Nm/kPa), and brake pressures (kPa), vs. time (s) for baseline test with no adaptation.

Figure 2.3 shows that the estimated torque converges towards the measured value, and the upper right plot of Figure 2.4 shows that the velocity error converges to zero and the parameter $\hat{K}_b$ converges. When the initial estimate of the brake torque gain differed by a large amount (a factor of three or four) from the actual brake torque gain, the parameter estimate, $\hat{K}_b$, sometimes would overshoot before converging to the correct value, causing uncomfortable oscillations in the brake pressure.

2.4.3 Non-Smooth Adaptation

The final test demonstrates tracking performance with nonsmooth parameter adaptation. As above, the actual value of K_b is $0.39\frac{Nm}{kPa}$, but the value which the controller uses is 50% too high at $0.58\frac{Nm}{kPa}$.

Figure 2.5 shows that the estimated torque converges towards the measured value, and Figure 2.6 shows that the velocity error converges to zero and the parameter $\hat{K}_b$ converges. When the initial parameter error was large, the uncomfortable oscillations which were present for the smooth parameter adaptation law did not occur.

Interesting behavior occurs once the parameter converges to the correct value. Due to modeling uncertainties, the parameter "chatters" around the correct value at a very low frequency. As a result, the surface also shows very

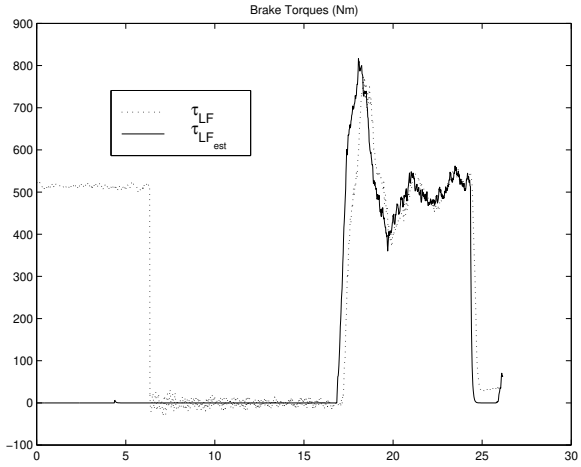

Figure 2.3: Measured and estimated torques (Nm) vs. time (s) using smooth adaptive control law to adjust value of $\hat{K}_b$ and surface definition $S_1 := v - v_{des}$.

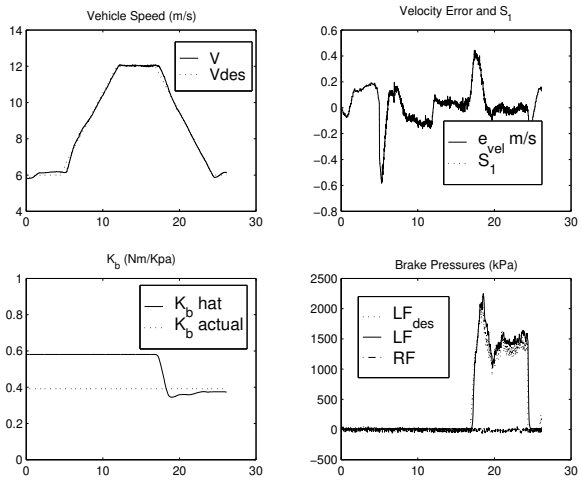

Figure 2.4: Vehicle speed (m/s), velocity error (m/s), K_b (Nm/kPa), and brake pressures (kPa), vs. time (s) for smooth $\hat{K}_b$ adaptation with $S_1 := v - v_{des}$.

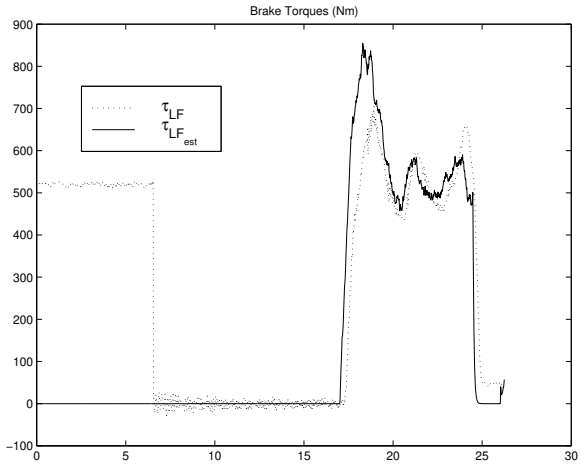

Figure 2.5: Measured and estimated torques (Nm) vs. time (s) using switching adaptive control law to adjust value of $\hat{K}_b$ and surface definition $S_1 := v - v_{des}$.

low frequency chattering. The chattering was invisible to passengers and, in cases with many parameters, could prove to be beneficial by providing excitation to force convergence.

2.5 Conclusion

The non-adaptive sliding mode brake controller performed well but showed a tracking error when there were parameter mismatches. Although this could be remedied with a higher surface gain, an integral term in the surface, or a switching term, it would likely be at the expense of increased control effort.

Both adaptive algorithms reduced velocity tracking error and had their parameters converge to the correct value. The smooth adaptation scheme converged as expected, and the nonsmooth adaptation algorithm converged in a linear matter, chattering at a low frequency once it reached the correct parameter value. For this application, the nonsmooth parameter adaptation law gave better results because it had less of a tendency to over/undershoot the correct parameter value when the initial parameter error was large.

An interesting future direction for work would be to compare the convergence properties of the nonsmooth adaptation algorithm with those of the smooth adaptation algorithm in cases where persistence of excitation

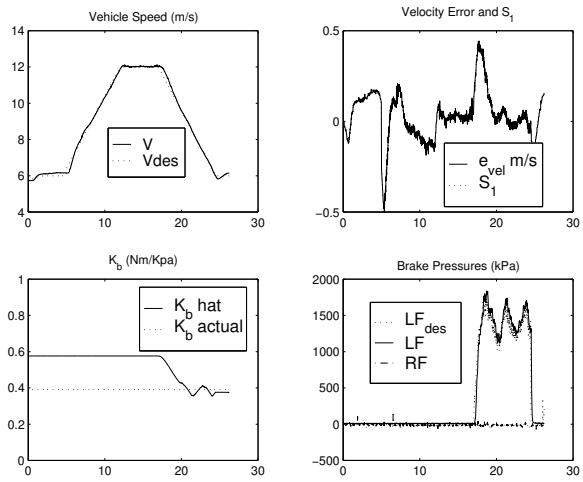

Figure 2.6: Vehicle speed (m/s), velocity error (m/s), K_b (Nm/kPa), and brake pressures (kPa), vs. time (s) for test with switching $\hat{K}_b$ adaptation with $S_1 := v - v_{des}$.

is difficult to achieve.

Chapter 3

Rejection of Periodic Disturbances Via Adaptive Control

3.1 Introduction

Eliminating the noise, vibration, degradation of steering control, and premature wheel lock up caused by rotor induced vibrations in a brake systems has been a goal of brake hardware designers for many years. These vibrations originate from two sources: 1) Thickness variations due to manufacturing limitations and 2) Rotor warpage due to uneven heating and cooling during braking. A comprehensive study quantifying the causes, effects, and magnitudes of brake vibration can be found in [23].

Until recently, better hardware design was the primary way to deal with this problem. However, with the increasing popularity of systems like ABS, TCS, VDC, and CW/CA which can automatically actuate the brakes with a high bandwidth, it may be possible to eliminate rotor-induced brake vibrations through electronic control.

This chapter demonstrates the feasibility of a tracking brake controller which eliminates brake vibrations using a technique called Adaptive Feedforward Cancellation (AFC). Essentially, the system consists of a stable tracking feedback controller combined with a scheme that estimates the periodic disturbance and then feeds forward a signal to cancel it.

Section 3.2 gives a background on the literature of AFC and explains the

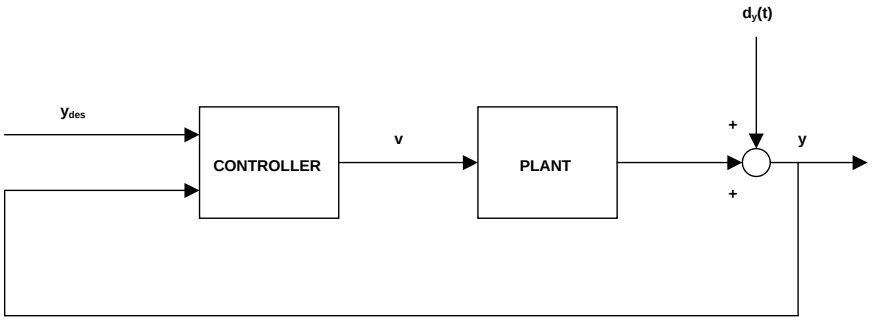

Figure 3.1: Control system with disturbance entering at the output.

technique in detail. Since most studies in the literature involve linear plants and a brake system is a nonlinear plant, Section 3.3 extends AFC results to include nonlinear plants that are affine in the input by proving the asymptotic stability of n^{th} order systems under sliding mode/AFC control. Section 3.4 then illustrates the difficulties encountered in proving the stability of AFC schemes for non-affine nonlinear systems but shows through simulation that the scheme nevertheless works very well. Section 3.5 then successfully applies the technique to the system of interest, an automobile braking system. Finally, Section 3.6 offers conclusions and suggestions for future work.

3.2 Adaptive Feedforward Cancellation

Figure 3.1 shows a typical feedback control system with a disturbance entering at the output. If the disturbance is periodic and has a known fundamental frequency, an Adaptive Feedforward Cancellation scheme might be appropriate for cancelling it. Examples of applications where the frequency of a disturbance is known include the control of disk drives, CD players, motors, cutting operations, and automobile brakes.

The idea behind the AFC scheme is to map the output disturbance, $d_y(t)$, through the plant to a fictitious (or real) disturbance in the input $d_u(t)$ that is causing it and then add a term to the controller output, v, which cancels this disturbance, as shown in Figure 3.2. For linear plants, a sinusoidal disturbance at the output maps to a sinusoidal disturbance at the input with only the magnitude being altered. Nonlinear plants are more complex because extra frequencies may be generated in the mapping from ouput to input, but it is still reasonable to expect that the fundamental frequency of

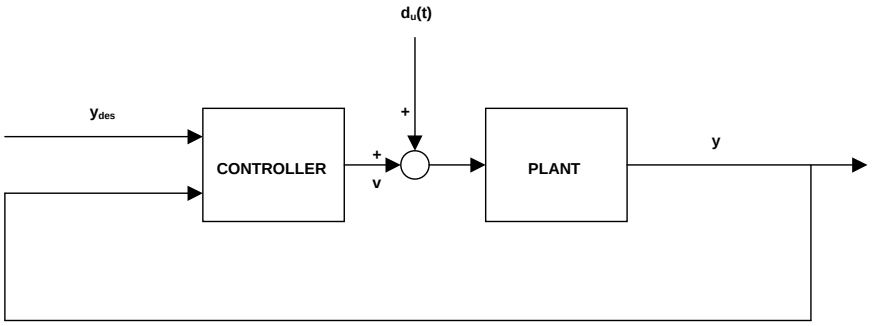

Figure 3.2: Control system with disturbance moved to plant input.

the disturbance will retain most of the energy of the disturbance.

There has been a fair amount of work done successfully using this technique on linear systems. In [6], Chen and Paden apply the technique to reduce torque ripple in a stepper motor (linear plant). In addition to successful experimental results, they show that the system is robust to disturbance frequencies beyond those that the system was designed to cancel. Bodson *et. al.* [3] examine the cancellation of higher order harmonics in linear systems in greater detail and show that AFC is in some sense equivalent to the Internal Model Principle. In [18], further experimental results are presented.

To date, though, there seems to have been little interest in applying this technique to nonlinear plants. In this chapter the technique is applied to both plants that are affine and are not affine in the input and it is successful in both cases. Unfortunately, a stability proof was only derived for the affine case as the non-affine case proved to be very challenging to handle (for the time being anyways). However, as the encouraging simulation results of Section 3.5 show, AFC may still be of use for controlling brake rotor vibration.

3.3 Affine Nonlinear System

This section extends the AFC method to affine nonlinear plants and verifies the results via simulation. First we derive a stable adaptation law for this class of plants which guarantees the boundedness of the parameter errors and the convergence of the tracking error to zero and then we demonstrate persistency of excitation, leading to the conclusion that the parameter errors will also converge to zero and that the system is asymptotically stable.

3.3.1 Stable Adaptation Law

Consider the affine nonlinear plant in canonical form:

$$\begin{aligned} \dot{x_1} &= x_2 \\ \dot{x_2} &= x_3 \\ &\vdots \\ \dot{x_n} &= f(\mathbf{x}) + g(\mathbf{x})u \end{aligned}$$

where the "control" u is the sum of a periodic disturbance $d(t)$ and the input to the plant from the controller, v (ie $u = v + d(t)$) and where the goal is for the state $\mathbf{x}$ to track some desired trajectory $\mathbf{x}_d$ given by a reference model. We first consider the problem of deriving a control law which will give asymptotic tracking when the external disturbance is zero. Since the plant that originally motivated this work is a hydraulic system (which by its "stiff" nature does not respond very well to discontinuous control laws), the focus here will be on using a smooth version of sliding mode control:

Define $\tilde{x}_1 = x_1 - x_{1d}$ and

$$S = \left(\frac{d}{dt} + \lambda\right)^{n-1} \tag{3.1}$$

At $S = 0$, the tracking is perfect. To force the system towards the surface, we define the desired surface dynamics as

$$\dot{S} = -\lambda S \tag{3.2}$$

so that the system will approach the surface exponentially. Differentiating S gives

$$\dot{S} = \tilde{x}_1^n + \ldots + \lambda^{n-1}\dot{\tilde{x}}_1 = \dot{x}_n - \dot{x}_{nd} + CE(\tilde{x}_1) \tag{3.3}$$

Substituting this in the expression for the desired surface dynamics (equation 3.2) and replacing $\dot{x}_n$ with the expression from the plant dynamics (equation 3.1) gives

$$\dot{S} = f(\mathbf{x}) + g(\mathbf{x})u - \dot{x}_n d + CE(\tilde{x}_1) = -\lambda S \tag{3.4}$$

Finally, solving for the desired input, u^* yields

$$u^* = \frac{\dot{x}_{nd} - \lambda S - CE(\tilde{x}_1) - f(\mathbf{x})}{g(\mathbf{x})} \tag{3.5}$$

In the absence of the periodic disturbance, this control input would give us the desired closed loop tracking performance. However, since there is a periodic disturbance and the control law attempts to cancel it through a feedforward term as outlined in Section 3.2, we choose our controller's output, v, to be the control input above minus our best estimate of the disturbance, $\hat{d}(t)$, ie

$$v = u^* - \hat{d}(t) = u^* - \sum_{i=1}^{N} (\hat{a}_i sin i w_f t + \hat{b}_i cos i w_f t) \tag{3.6}$$

where N is the number of frequencies we choose to model, and w_f is the fundamental frequency of the disturbance. The input to the plant, u, is then

$$u = v + d(t) = u^* - \tilde{d}(t) = u^* - \sum_{i=1}^{N} (\tilde{a}_i sin i w_f t + \tilde{b}_i cos i w_f t) \tag{3.7}$$

where $\tilde{a}_i = \hat{a}_i - a_i$, $\tilde{b}_i = \hat{b}_i - b_i$, and $\tilde{d}(t) = \hat{d}(t) - d(t)$. Substituting this expression for u into equation 3.4, the expression for $\dot{S}$ yields the surface dynamics in the presence of a disturbance:

$$\dot{S} = -\lambda S - g(\mathbf{x})\tilde{d}(t) = -\lambda S - \sum_{i=1}^{N} g(\mathbf{x})(\tilde{a}_i sin i w_f t + \tilde{b}_i cos i w_f t) \tag{3.8}$$

With this equation for the perturbed closed-loop feedforward system's surface dynamics in hand, it is straightforward to derive a stable Lyapunov-based adaptation law. To start out, define a positive definite Lyapunov function of the sliding surface and the parameter errors:

$$V(S, \tilde{a}_i, \tilde{b}_i) = \frac{S^2}{2} + \sum_{i=1}^{N} \left(\frac{\gamma_{a_i} \tilde{a}_i^2}{2} + \frac{\gamma_{b_i} \tilde{b}_i^2}{2} \right) \tag{3.9}$$

Taking its time derivative, using the fact that $\dot{a}_i = \dot{b}_i = 0$, and substituting equation 3.8 for $\dot{S}$ gives

$$\dot{V} = -\lambda S^2 - \sum_{i=1}^{N} g(\mathbf{x}) S \left(\tilde{a}_i sin iw_f t + \tilde{b}_i cos iw_f t \right) + \sum_{i=1}^{N} \left(\gamma_{a_i} \tilde{a}_i \dot{\hat{a}}_i + \gamma_{b_i} \tilde{b}_i \dot{\hat{b}}_i \right) = -\lambda S^2 \tag{3.10}$$

if $\dot{\hat{a}}_i$ and $\dot{\hat{b}}_i$ are chosen as

$$\begin{aligned} \dot{\hat{a}}_i &= \frac{S g(\mathbf{x})}{\gamma_{a_i}} sin iw_f t \\ \dot{\hat{b}}_i &= \frac{S g(\mathbf{x})}{\gamma_{b_i}} cos iw_f t \end{aligned}$$

The fact that $\dot{V}$ is negative semidefinite implies that all of the parameter errors and the value of the sliding surface, S, are bounded. Furthermore, calculation of $\ddot{V}$ reveals that it is bounded, implying the uniform continuity of $\dot{V}$. This gives the welcome result that $\dot{V} \to 0$ as $t \to \infty$, meaning that $S \to 0$ as $t \to \infty$. To make conclusions about the parameter errors, other than the fact that they are bounded, we need to make the persistence of excitation arguments presented in the next section.

3.3.2 Persistence of Excitation

This section demonstrates that persistence of excitation of the system is maintained independent of the reference trajectory, allowing us to conclude that the parameter errors converge to zero. To simplify notation, we can re-write the adaptation laws of equation 3.11 as

$$\dot{\hat{D}} = S \cdot g(\mathbf{x}) \Gamma w \tag{3.11}$$

with

$$\begin{aligned} \hat{D} &= \left[\hat{a}_1, \hat{b}_1 \ldots \hat{a}_N, \hat{b}_N \right]^T \\ w &= \left[sin(w_f t), cos(w_f t) \ldots sin(N w_f t), cos(N w_f t) \right]^T \\ \Gamma &= diag(\gamma_{a_1}, \gamma_{b_1} \ldots \gamma_{a_N}, \gamma_{b_N}) \end{aligned}$$

To show persistence of excitation, we need to show that there exist some positive scalar constants α and T such that for all $t > 0$

$$\int_t^{t+T} w w^T d\tau \geq \alpha I \tag{3.12}$$

where I is the appropriately sized identity matrix.

Constructing the integral, we get

$$\int_t^{t+T} \begin{bmatrix} sin^2(w_f t) & & \dots & sin(w_f t)cos(Nw_f t) \\ & cos^2(w_f t) & & \\ \vdots & & sin^2(2w_f t) & \\ & & & \ddots \\ sin(w_f t)cos(Nw_f t) & & \dots & cos^2(Nw_f t) \end{bmatrix} d\tau \tag{3.13}$$

Selecting $T = \frac{2\pi}{w_f}$ (so that we integrate over one period) and exploiting the orthogonality of sines and cosines at differing frequencies simplifies this integral to

$$\begin{bmatrix} \frac{\pi}{w_f} & 0 & \dots & 0 \\ 0 & \frac{\pi}{w_f} & & \\ \vdots & & \ddots & \\ 0 & \dots & 0 & \frac{\pi}{w_f} \end{bmatrix} \geq \alpha I \tag{3.14}$$

for $\alpha = \frac{\pi}{w_f}$ so persistence of excitation is established. Note that persistence of excitation is satisfied for *any* reference trajectory and as long as the fundamental frequency of the sinusoids is nonzero. Therefore, the parameter errors will always converge to zero and the control law/parameter adjustment mechanism derived in the previous section is shown to be asymptotically stable in the presence of an unknown disturbance. As the simulations in the next section will show, this result comes with one caveat: the regressor, w, must contain at least as many components as the disturbance. If the regressor lacks sufficient richness to fully model the disturbance, the system still appears to be stable, but the errors and the surface only appear to converge to a small neighborhood around zero and the Lyapunov based stability proof in the form it is presented above is lost. To circumvent this problem for a linear plant with SPR transfer function, Chen and Paden [6] prove the exponential stability of the adaptive system in the absence of a disturbance and then treat the disturbance aas a robustness problem, using a small gain theorem-like argument.

3.3.3 Simulation Results

The simulations presented in this section were performed on the plant

$$\dot{x} = -x + (1 + x^2)u \tag{3.15}$$

and the control goal was to track a trajectory given by a "reference" system. A "smooth" sliding control with the adaptation equations given above were used. Adaptation for only two parameters is shown here, but simulations were also done with three parameters and gave similar results. For completeness, the control law used and parameter update laws were:

$$\begin{aligned} v &= \frac{\dot{x}_d - \lambda S + x}{x^2 + 1} - \hat{a}_1 sin(w_f t) - \hat{b}_1 cos(w_f t) \\ \dot{\hat{a}}_1 &= \frac{(1 + x^2) \cdot S}{\gamma_{a_1}} sin(w_f t) \\ \dot{\hat{b}}_1 &= \frac{(1 + x^2) \cdot S}{\gamma_{b_1}} cos(w_f t) \end{aligned}$$

Figure 3.3 shows the tracking ability of the closed loop, sliding mode controlled system without the adaptive feedforward component of the control law. Thus, the Fourier coefficients $\hat{a}_1$ and $\hat{b}_1$, which are shown in the lower left plot, do not adjust. The upper left plot shows the desired and actual values of x. The upper right plot shows the controller output, v, and the input to the plant, u, (includes the external disturbance, ie $u = v + d(t)$). Although the sliding mode controller rejects the disturbance somewhat, there is clearly room to improve.

Figure 3.4 is the same as Figure 3.3, except that the adaptation is turned on. The tracking is very good (upper left graph) and the parameters converge to the correct values very quickly. In addition, the value of the time derivative of the Lyapunov function in equation 3.9 is plotted and is never greater than zero, as predicted. Notice the interesting u that is required to create a sinusoidal x. The parameter convergence is independent of the desired trajectory.

Figure 3.5 is the same as the previous figure, except that an unmodeled higher harmonic has been added to the disturbance. The most noticeable effect of the extra harmonic is that it causes the parameters representing the disturbance at the fundamental frequency to vary with time. For linear

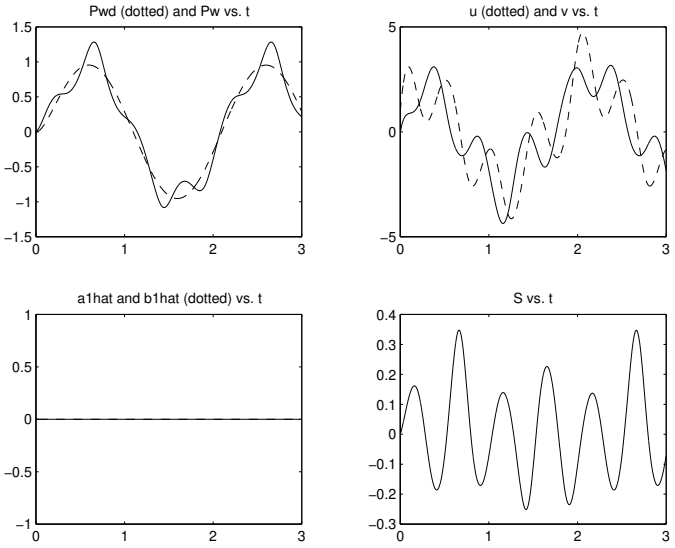

Figure 3.3: Affine plant with sliding mode control, no AFC, and a periodic disturbance with only one frequency. Upper left graph is actual state x and desired state x_d. (These will soon become our actual wheel pressure, P_w and the desired wheel pressure, P_{wd} when we apply AFC to the brake problem). $\lambda = 10, a_1 = 2, b_1 = 1$.

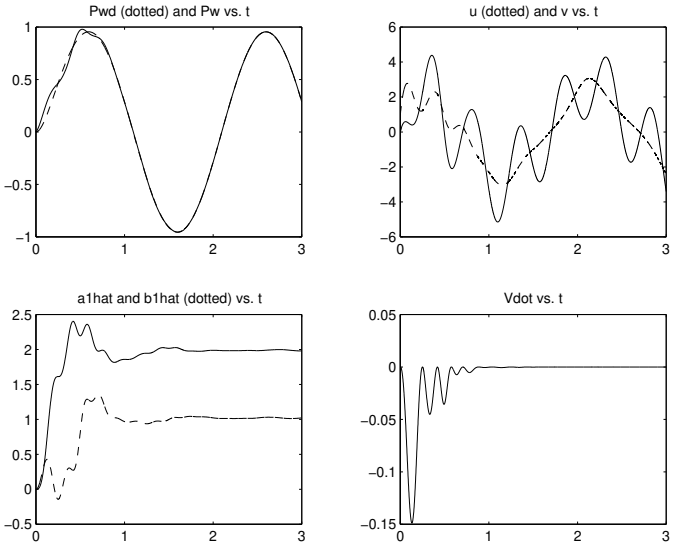

Figure 3.4: Affine plant with AFC and sliding mode control. $\lambda = 10, a_1 = 2, b_1 = 1, \gamma_{a_1} = \gamma b_1 = 0.01$

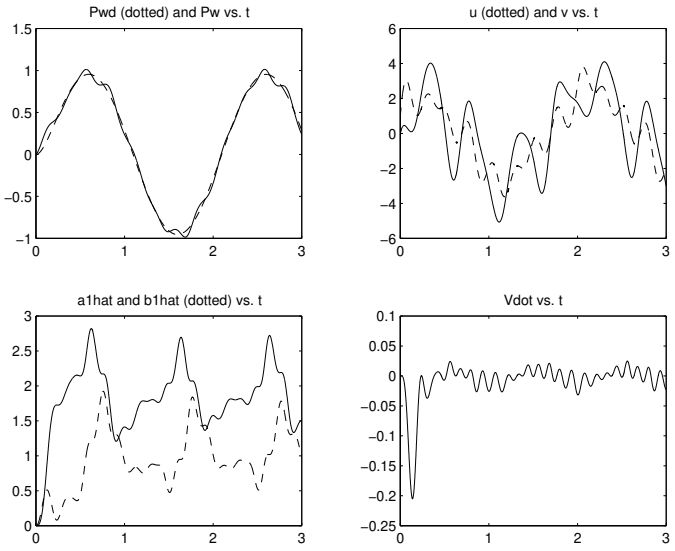

Figure 3.5: Affine plant with AFC and sliding mode control but with an unmodeled frequency. $\lambda = 10, a_1 = 2, b_1 = 1, a_2 = 0.5\gamma_{a_1} = \gamma b_1 = 0.01$

systems, this phenomenon has been studied [3] and was explained as the modulation of the higher order harmonic with the regressor. In some cases this time variation of the parameters actually *reduced* the magnitude of the higher harmonic compared to the case where the parameters were "frozen" to their correct values. In this case, the time variation did not noticeably help or hurt the response compared to the case where the parameters were frozen. Notice in this case that the derivative of the Lyapunov function (equation 3.9) is greater than zero. While this does not imply pending instability, it does imply that this Lyapunov function is not particularly useful for showing robustness to unmodeled harmonics.

3.4 Non-affine Nonlinear System

Since the automotive braking system that inspired this work is a nonaffine nonlinear system, it is natural to try to extend the results from the previous section to nonaffine plants. Towards this end, this section is divided into two parts: Section 3.4.1 attempts to use a derivation similar to that of Section 3.3 to arrive at a stable adaptation law, demonstrating the difficulties presented by a nonaffine system. It then discusses some of the approaches

tried (all unsuccessfully) to derive a stability guarantee for AFC used on nonaffine nonlinear systems. Section 3.4.2 then presents results of a series of simulations used to explore different adaptation laws on a very simple plant. Although none of these laws is proven stable rigorously, the fact that they all work very well indicates that nonlinear nonaffine systems are quite receptive to the AFC approach to periodic encouraging further theoretical exploration in the area.

3.4.1 Difficulties of Non-Affine System

Consider the non-affine nonlinear system in canonical form given by

$$\begin{aligned} \dot{x_1} &= x_2 \\ \dot{x_2} &= x_3 \\ &\vdots \\ \dot{x_n} &= f(\mathbf{x}, u) \end{aligned}$$

where, as before, the "control" u is the sum of a periodic disturbance $d(t)$ and the input to the plant from the controller, v (ie $u = v + d(t)$) and where the goal is for the state $\mathbf{x}$ to track some desired trajectory $\mathbf{x}_d$ given by a reference model.

As before, we consider the problem of controlling the "undisturbed" system first and then consider disturbances. As before, a smooth formulation of sliding mode control is used, the idea being that switching should be kept to a minimum, especially in systems like the stiff hydraulic one that inspired this work. The derivation of the control law is given only tersery treatment here because it is a standard application of sliding mode control and is very similar to Section 3.3.1 above:

Define

$$S = \left(\frac{d}{dt} + \lambda\right)^{n-1} \tag{3.16}$$

The desired surface dynamics are

$$\dot{S} = -\lambda S \tag{3.17}$$

Differentiating S, substituting this result in the desired surface dynamics equation, utilizing the plant dynamics, and solving for the input u^* gives

$$u^* = f^{-1}(\dot{x}_{nd} - \lambda S - CE(\tilde{x}_1)) \tag{3.18}$$

where f^{-1} is the inverse of $f(\mathbf{x}, u)$ given that $\mathbf{x}$ is the current state.

Following the procedure of Section 3.3.1 above, we then choose the controller output, v, as $v = u^* - \hat{d}(t)$ and note that this makes the input to the plant, u, equal to $u^* - \tilde{d}(t)$ where $\tilde{d}(t)$ is the same as in Section 3.3.1. Next, we substitute this u into the equation for $\dot{S}$ to get a surface dynamics equation for the system with a disturbance. This yields

$$\dot{S} = f(\mathbf{x}, u^* - \tilde{d}(t)) - \dot{x}_n d + CE(\tilde{x}_1) \tag{3.19}$$

Since the control was multiplied by a function of the states, $g(\mathbf{x})$ in the affine case, u^* canceled out the $-\dot{x}_n d + CE(\tilde{x}_1$ terms and gave the desired surface dynamics, $\dot{S} = -\lambda S$ while the $\tilde{d}(t)$ term acted as a disturbance. The left hand side of Figure 3.6 illustrates how u^* and $\tilde{d}(t)$ affect $\dot{S}$ for a given value of the state $\mathbf{x}$ and a given desired trajectory $\mathbf{x_d}$. The important element is that the gain between $\tilde{d}(t)$ and $\dot{S}$ is independent of the size of the disturbance. This makes it possible to "cancel" the effect of the disturbance on the derivative of the Lyapunov function, as in equation 3.10 above. As the right hand side of Figure 3.6 shows, for a given value of the state $\mathbf{x}$ and a given desired trajectory $\mathbf{x_d}$, the gain between the disturbance and $\dot{S}$ is dependent on the value of the disturbance itself, but since the size of the disturbance is unknown, the gain is also unknown.

3.4.2 Simulation Testing of Adaptation Laws

Rather than experiment with adaptation laws on the "stiff" automotive braking system which has a tendency to be numerically difficult, experiments were ran on the simple plant

$$\dot{x} = -x + sgn(u)u^2, \quad y = x \tag{3.20}$$

using a variety of adaptation laws and a smooth sliding mode controller. Although none of them were proven stable, all were motivated by some physical reasoning. The two most successful adaptation laws are presented here:

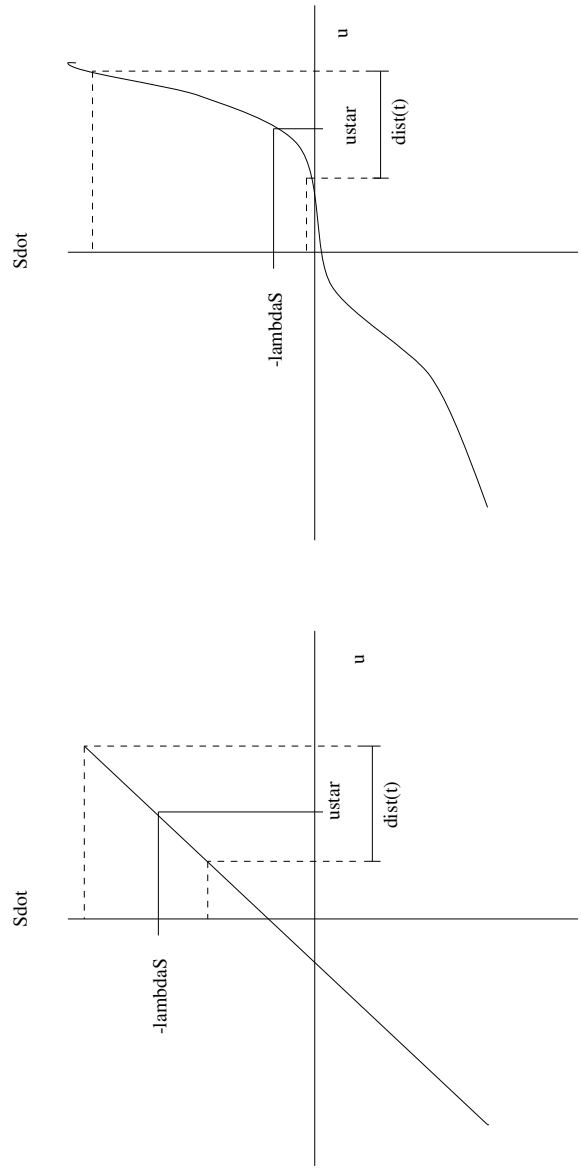

Figure 3.6: Mappings between u and $\dot{S}$ for fixed $\mathbf{x}$ and $\mathbf{x_d}$ for affine (left) and nonaffine(right) systems.

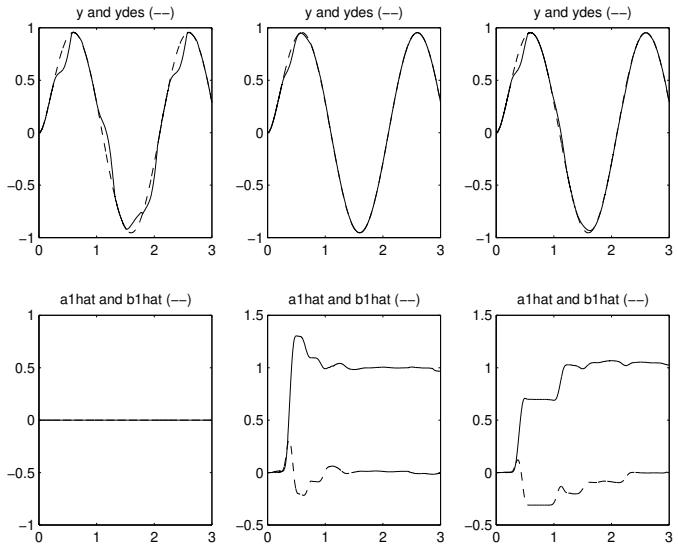

Figure 3.7: System output, y, and desired output, y_{des} vs. time, with corresponding Fourier coefficients, $\hat{a}_1$ and $\hat{b}_1$, vs. time. *Left*: Closed loop, no AFC. *Middle*: $\dot{\hat{a}}_i = \frac{S}{\gamma_{a_i}} sin(w_f t), \dot{\hat{b}}_i = \frac{S}{\gamma_{b_i}} cos(w_f t)$. *Right:* $\dot{\hat{a}}_i = \frac{S \cdot f_u}{\gamma_{a_i}} sin(w_f t) \dot{\hat{b}}_i = \frac{S \cdot f_u}{\gamma_{b_i}} cos(w_f t)$

The first column of plots in Figure 3.7 shows the response of the closed loop system without AFC. Tracking is affected significantly. The second column shows the response of a system with AFC, using the parameter adaptation laws

$$\begin{aligned}\dot{\hat{a}}_i &= \frac{S}{\gamma_{a_i}} sin(w_f t) \\ \dot{\hat{b}}_i &= \frac{S}{\gamma_{b_i}} cos(w_f t)\end{aligned}$$

These are the what the parameter laws would come out to be if the system were linear. Essentially, using these parameter laws assumes that even though the system is nonlinear, the direction we should push the parameters should be the same as for its linear equivalent. Intuitively, this assumption seems like it would be correct if the nonlinear system's f is monotonic in u.

The last column shows the response with AFC and the parameter law

$$\begin{aligned}\dot{\hat{a}}_i &= \frac{S \cdot f_u}{\gamma_{a_i}} sin(w_f t) \\ \dot{\hat{b}}_i &= \frac{S \cdot f_u}{\gamma_{b_i}} cos(w_f t)\end{aligned}$$

where f_u is the partial of f with respect to u evaluated at the value u^* and the current state $\mathbf{x}$. This adaptation law was inspired by the idea of approximating the nonlinear function that relates $\dot{S}$ to u by a linear one. Again, results are very good and parameters converge to the correct values.

3.5 AFC on Brake System

Although there were no definitive stability proofs for the AFC/sliding mode control law applied to non-affine systems, the encouraging results of the previous section motivated experimentation with the plant of interest–an automotive braking system. Section 3.5.1 outlines the plant model and "base" control law while Section 3.5.2 shows simulation results.

3.5.1 Brake Model/Controller

Figure 3.8 shows the braking system under consideration. The system output is the brake pressure at the wheel, denoted P_W, and the control input is the master cylinder pressure, P_{MC}. The disturbance is the lateral movement of the point where the distorted brake rotor touches the brake pad as it spins. Using simple fluid mechanics, this positional disturbance can be translated into a pressure disturbance. If we assume that the clearance between the rotor and the pads is small and that the system's hydraulic capacitance is linear and that the speed of the vehicle changes very little in the period of one brake rotor oscillation, then the equation governing the output is

$$\dot{P}_W = C_q C_v sgn(P_{MC} - P_W)\sqrt{|P_{MC} - P_W|} + \sum_{i=1}^{n}(a_i sin\omega_i t + b_i cos\omega_i t) \quad (3.21)$$

where C_q is a lumped flow coefficient, C_v is a hydraulic capacitance, ω_i is the rotational frequency of the wheels, and a_i and b_i are disturbance dependent parameters which must be estimated. Note that to put this system into the form we have been dealing with so far (disturbance entering at the input to the plant) we must use our imagination and pretend that the pressure disturbance which is actually caused by the rotor pushing against the pads is actually caused by some renegade element of the master cylinder pressure.

The "nominal" control law designed to drive the wheel pressure P_W to a desired wheel pressure $P_{W,des}$ is again a "smooth" exponentially convergent sliding mode controller, given by

$$P_{MC} = P_W + sgn(S)\left[\frac{\dot{P}_{W,des} + \lambda}{C_q C_v}\right]^2 \quad (3.22)$$

with the sliding surface S defined as $S := P_{W,des} - P_W$.

3.5.2 Brake Simulations

For the simulations, $C_v = 480kPa/cc$ and $C_q = 1.4cc/s\sqrt{kPa}$. Figure 3.9 shows the response of the system using the control law given by equation 3.22 but no AFC in the presence of a disturbance. The controller is partially successful in canceling the disturbance, but tracking is still severly hindered.

In Figure 3.10 a discontinuous sliding mode controller has been implemented which treats the disturbance as an uncertainty. Note that

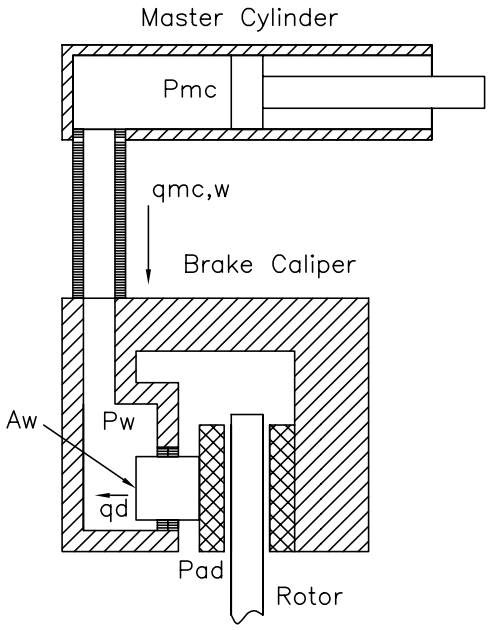

Figure 3.8: The braking system under consideration in this study.

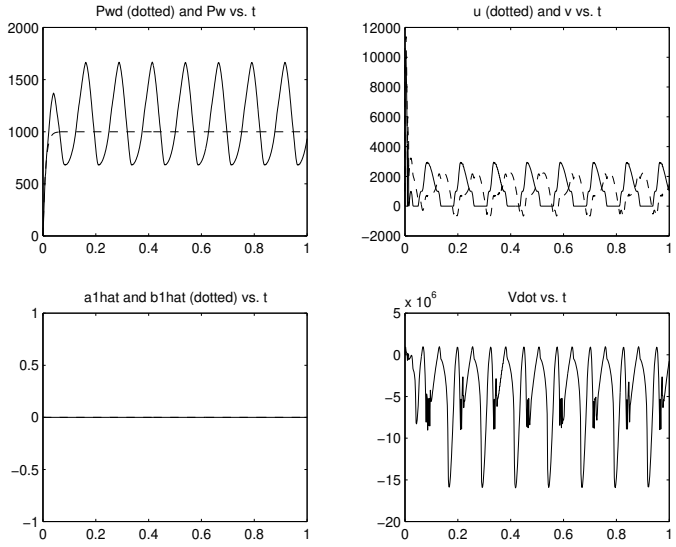

Figure 3.9: Brakes under sliding mode control with no AFC. $\lambda = 100, a_1 = 2000, b_1 = 1000$

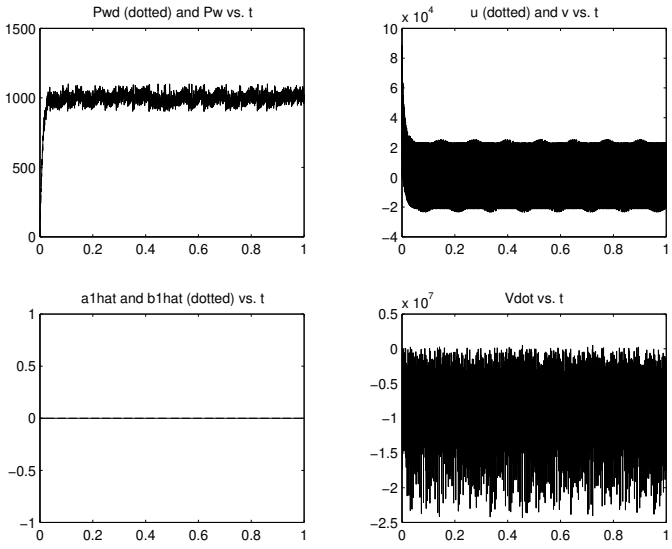

Figure 3.10: Brakes under switching sliding mode control with no AFC. $a_1 = 2000, b_1 = 1000$

although the disturbance does not make its way into the output, it is eliminated at the expense of very high frequency and high amplitude control effort.

Figure 3.11 shows the closed-loop AFC compensated system response. The parameters are adjusted according to the equations

$$\begin{aligned} \dot{\hat{a}}_1 &= \frac{S}{\gamma_{a_1}} sin(w_f t) \\ \dot{\hat{b}}_1 &= \frac{S}{\gamma_{b_1}} cos(w_f t) \end{aligned}$$

The parameters converge to the correct values and the error converges to zero very quickly. Again, the graph labeled $\dot{V}$ refers to the Lyapunov function given by equation 3.9.

Finally, Figure 3.12 shows the response of the closed loop AFC/sliding mode system to an unmodeled harmonic. Again, the response shows the "modulation" observed earlier, but the higher harmonic appears to be attenuated very little in the output.

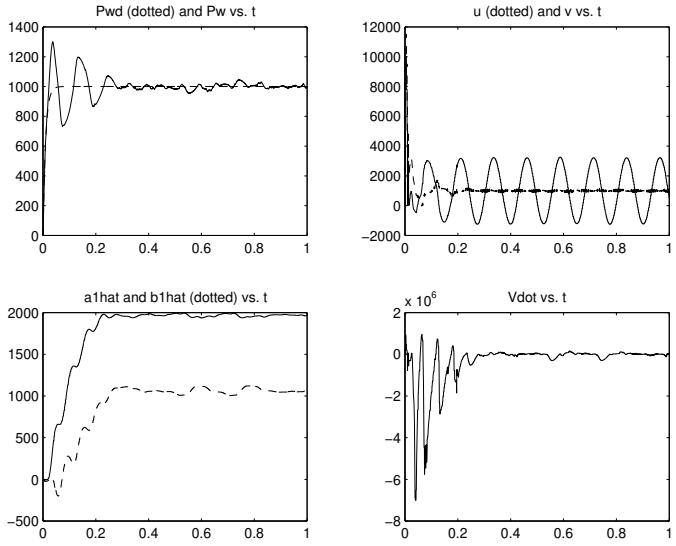

Figure 3.11: Brakes with AFC and adaptation laws $\hat{a_1} = S{\cdot}sin(w_f t)/\gamma_{a_1}, \hat{b_1} = S \cdot cos(w_f t)/\gamma_{b_1}$. $\lambda = 100, a_1 = 2000, b_1 = 1000, \gamma_{a_1} = \gamma_{b_1} = 0.01$

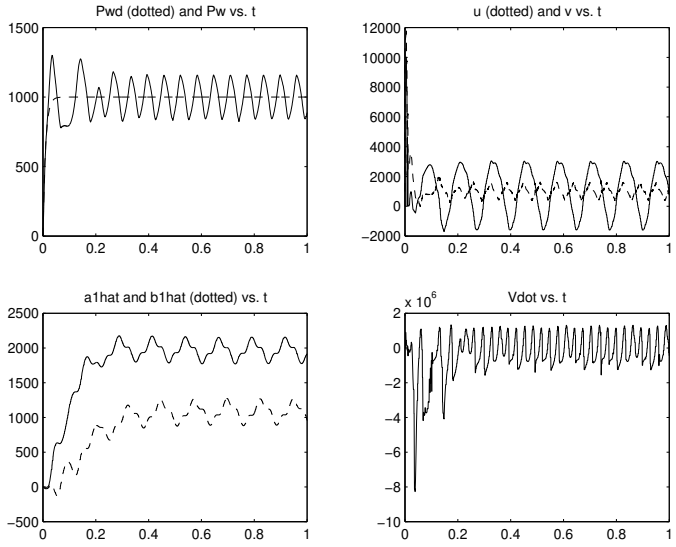

Figure 3.12: Brakes with AFC, unmodeled disturbances and same adaptation laws as in Figure 3.11. $\lambda = 100, a_1 = 2000, b_1 = 1000, a_2 = 500\gamma_{a_1} = 0.01, \gamma_{b_1} = 0.01$

3.6 Conclusions

The stability, error convergence, and persistence of excitation for an affine nonlinear system using AFC and smooth sliding mode control was established analytically and it was shown through simulation that the AFC approach seems to work for non-affine systems as well, including the brake system of interest in this study. There are several areas which need further investigation:

1. It seems extremely likely that some kind of proof of the stability of a nonaffine system which is nondecreasing (or nonincreasing) in u under AFC and sliding mode control should exist, especially if f_u is assumed to be bounded.

2. The way stability was proven for the affine system in this project (Lyapunov based approach, Lyapunov a function of state errors) does not lend itself very well to proving robustness to unmodeled disturbances. A better approach might be to show the stability of the unperturbed adaptive system and then use a small gain theorem-like argument to show robustness to disturbances and possibly modeling error.

3. The results here for brake control would be interesting to try out experimentally.

4. The attenuation of higher order harmonics than those included in the model of the disturbance by time-variation of the parameters is an interesting phenomenon and should be investigated further.

Chapter 4

Strain-based Torque Sensor, Differential Braking, and Wheel Speed Hardware

4.1 Introduction

This chapter documents the hardware produced under MOU 308. Specific developments that are addressed include a torque sensor, differential braking hardware, and wheel speed conditioning circuitry.

The torque sensor and differential braking were developed to verify the algorithms of Chapter 2, but they have proved indispensable in numerous other applications, including slip measurement and road condition estimation under MOU 388. The torque sensor provides direct measurements of braking effort that other sensors can only approximate. It accurately measures torques from 0 Nm to approximately 2500 Nm where wheels typically lock up on dry pavement. The differential braking hardware allows for brakes to be applied only at selected wheels and has mainly been used to isolate the wheel with the brake torque sensor on it so that it is the only one applying a braking force on the car. It could, however, also be used to apply yaw moments to help bring an out-of-control vehicle under control. The four wheel speed measurements were added at the same time as the differential braking hardware and quickly became the most useful new addition to the test vehicle.

Figure 4.1 illustrates the locations of all the hardware additions to the

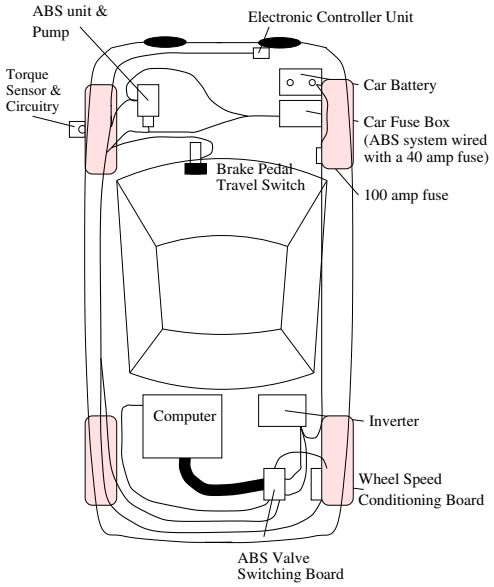

Figure 4.1: Wiring Layout of Hardware Additions to the Red Lincoln Towncar

red Lincoln Towncar produced under MOU 308. It also shows the routing of all the wires used for carrying the power and control signals.

4.2 Hardware Documentation

4.2.1 Torque Sensor

The project group constructed and installed a strain-based brake torque sensor (Figure 4.2 on the Red Lincoln Towncar in order to experimentally verify brake torque estimation algorithms.

Braking Model for Undriven Wheel

Figure 4.3 shows the forces acting on a non-driven wheel during braking. An intuitive account of how the brake slows the car is as follows:

As the wheel cylinder hydraulic pressure increases, a torque of magnitude $K_b P_w$ is generated, decreasing the wheel's angular velocity. Since the wheel's angular velocity is smaller than its translational velocity, there is "slip."

Figure 4.2: Brake torque sensor installed on the Red Towncar.

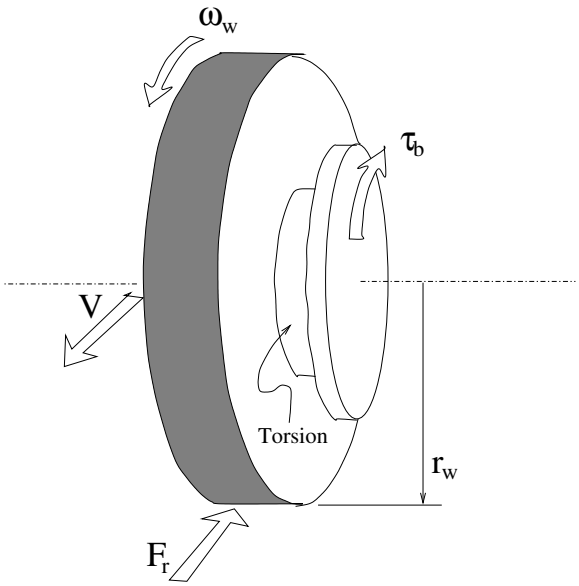

Figure 4.3: Forces acting on an undriven wheel.

"Slip," as we use it here, means longitudinal slip, often abbreviated by λ in the vehicle dynamics and control literature, and defined as

$$\lambda = \begin{cases} \frac{R_w\omega_w - v}{v} & \text{for braking} \\ \frac{R_w\omega_w - v}{R_w\omega_w} & \text{for accelerating} \end{cases}$$

As Figure 4.4 shows, the slip generates a longitudinal tire force. The longitudinal tire force acts in the direction opposite of the direction of travel, simultaneously slowing the vehicle and counteracting the brake pad torque of magnitude $K_b P_w$.

For the purposes of vehicle dynamics simulations and analysis, it is accurate to model the brake rotor, its attachment to the wheel, the wheel, and the tire as a single, rigid unit. In this case, a torque balance on the wheel/rotor unit, and a force balance on the vehicle as a whole gives the following equations of motion:

$$\begin{array}{rcl} \dot{\omega}_w & = & \frac{r_w}{J_w} F_r(v, \omega_w) - \frac{1}{J_w}\tau_b(P_w, \omega, F_r) \\ \dot{v} & = & \frac{1}{M}\sum_{i=1}^{4} F_{r_i}(v, \omega_{w_i}) \end{array} \tag{4.1}$$

where J_w is the moment of inertia of the wheel/rotor unit about the axle, F_r is the road force, r_w is the radius of the wheel, τ_b is the brake torque, and M is the mass of the vehicle plus wheels.

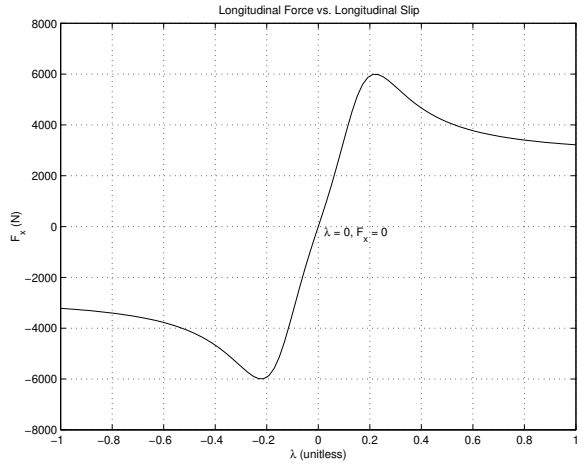

Figure 4.4: Typical slip curve showing longitudinal tire force, F_x, vs. longitudinal slip.

Equation 4.1 implies that through the longitudinal slip, the road force, F_r, is a function of the vehicle velocity, v, and the wheel's angular velocity, ω_w. Through this velocity feedback, the road force couples the equations for the vehicle and the wheels.

Also, equation 4.1 shows that the brake torque, τ_b, is not directly controlled. Instead, we control the brake pressure at the wheel cylinder, P_w, and model the generation of brake torque at the pads as a reactive "stiction" phenomenon governed by separate equations for when the wheel is moving and when it is stopped:

$$\tau_b = \begin{cases} -sgn(r_w F_r)\min(|r_w F_r|, K_b P_w) & \text{if } \omega_w = 0 \\ -sgn(\omega_w) K_b P_w & \text{if } \omega_w \neq 0 \end{cases} \tag{4.2}$$

The $\omega_w = 0$ part of equation 4.2 expresses the fact that, once the wheel is stopped, the brake pads will resist any torques attempting to start the wheel moving until the torque from the road exceeds $K_b P_w$. The second part of equation 4.2 expresses the fact that while the brake rotor is moving through the brake pads, the pads generate a torque of magnitude $K_b P_w$ in a direction to resist movement. If the brake pressure to torque gain, K_b, and wheel radius, r_w, are constant, we obtain the functional relationship employed in equation 4.1, ie $\tau_b = \tau_b(P_w, \omega_w, F_r)$. Figure 4.5 graphs the $\omega_w = 0$ part of equation 4.2.

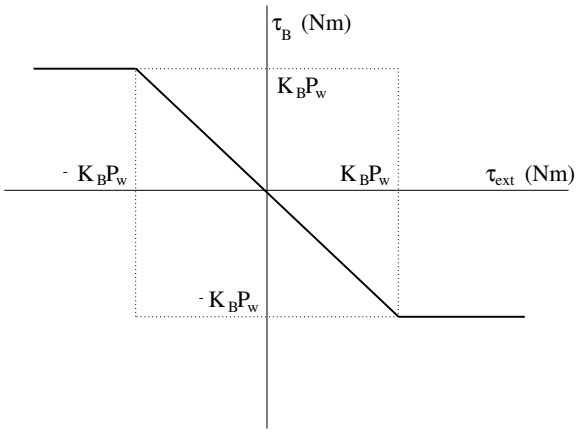

Figure 4.5: Brake torque, τ_b, vs. the sum of all other torques acting on a wheel, τ_{ext}, for a locked wheel ($\omega_w = 0$).

Strain Gage

Figure 4.3 shows that during braking, the brake rotor's short "shaft" between the wheel attachment bolts and the brake disk is in torsion. As the sensor calibration results of Section 4.2.1 show, the deflection of the shaft is proportional to the torque on it. The torque sensor uses an extremely sensitive strain gage to measure this minute deflection.

Figure 4.6 shows the full bridge strain gage configuration used to measure the deflection in the brake rotor. Because the deflections of interest are so small, half of the bridge employs extremely high gage-factor (strain sensitivity) semiconductor-based strain gages. To help offset the poor temperature stability of the semiconductor gages, the other half of the bridge uses standard metal, lower gage factor gages.

To reduce the effects of bending moments and other parasitic strains on the bridge output, R_1 and R_2 were positioned 180° away from R_3 and R_4. The strain gages were applied by Rotating Measurements Systems Incorporated. Figure 4.7 shows the brake rotor instrumented with the strain gages and on-wheel conditioning circuitry discussed in the next section.

On-Wheel Conditioning Circuitry

With approximately 9V of excitation, the signal from the strain gage bridge mounted on the brake rotor is a few tens of millivolts in amplitude. To avoid

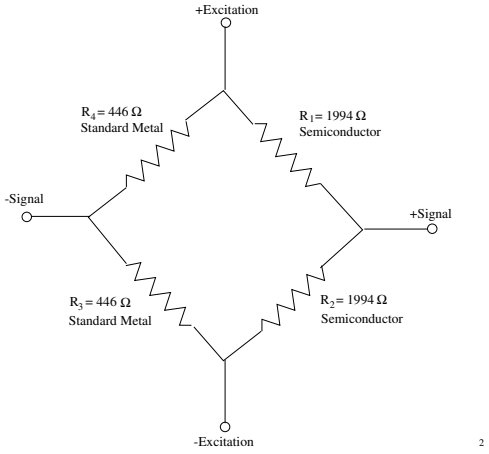

Figure 4.6: Mixed semiconductor/metal strain gage bridge used in torque sensor.

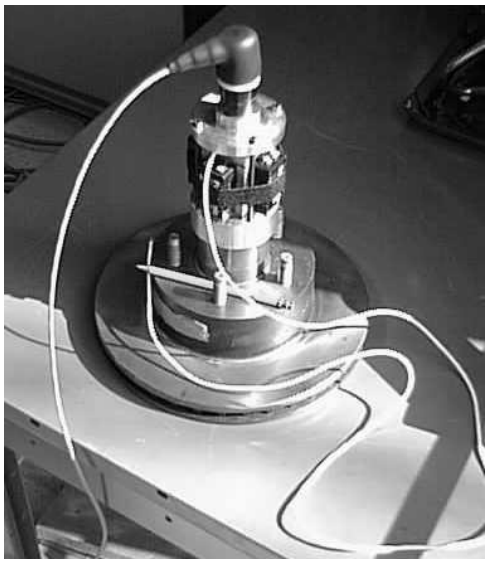

Figure 4.7: Instrumented brake rotor. Top to bottom: Slip ring, mounting bracket with signal conditioning circuitry and batteries, wheel mounting bolts, "shaft" in torsion with strain gage (light colored area below pencil), braking surface (on table).

transmitting this very small amplitude signal across a slip ring or by radio, we built on-wheel conditioning circuitry to amplify and filter the strain signal. In addition, the circuitry provides a regulated 8.73V excitation signal for the bridge.

The conditioning board rides on the wheel inside of a protective bracket assembly (see Section 4.2.1) which replaces the wheel bearing dust cap. Four standard 9V batteries riding on the outside of the bracket provide power for several weeks of testing.

Figure 4.8 shows a schematic and pinouts for the conditioning board, and Figure 4.9 shows an approximately 1:1 scale photograph of the board. Although the board was designed to accommodate two signals—one for strain and one for rotor temperature—we only implemented the strain circuitry because we found later that rotor temperature data was not necessary.

The first stage of signal conditioning is a differential amplifier which turns the differential, millivolt level strain signal into a signal ended signal with a magnitude of several volts. The INA118, a precision differential amplifier chip manufactured by Burr Brown, accomplishes this. It takes a differential input at pins 2 and 3 and outputs a single ended signal referenced to pin 5 and amplified by the gain set at pins 1 and 8 (set to 51 with a 1k resistor).

The second stage of signal conditioning is a lowpass Butterworth filter, formed with an LM324 op-amp and RC network. The design is a fixed gain VCVS type, taken from "Art of Electronics," by Horowitz and Hill, Second Edition, Page 274. We chose the breakpoint of the filter to be 100Hz, working under the assumption that we would use a 5ms data acquisition interval and that some aliasing at high frequencies would be tolerable. The gain of the amplifier/filter circuit is that of the amplifier alone (51) multiplied by the fixed gain of the Butterworth filter (1.586), giving a total of 81. A gain of 81 allows the board to condition the largest strain signals achievable—about 100mV, which occur during an ABS stop on a dry road—without exceeding rail voltages.

There are three important "rail" voltages on the board:

1. 0V from batteries: This voltage serves as the ground for the INA118 differential amplifiers, the LM78L09 8.73V voltage regulator, and the LM324 op-amps used in the Butterworth lowpass filters.

2. 18V from batteries: This voltage actually varies between 19V and about 15V, depending on the state of the batteries. This voltage is used as

the "V+" input to the INA118 differential amplifiers and LM324 op-amps. Since the actual magnitude of this voltage is a function of the battery state, we designed the circuit so that the level of this voltage is unimportant.

3. 8.73V from LM78L09 Regulator: The LM78L09 provides a stable voltage 8.73V above battery ground. This voltage is used as the reference for both the differential amplifiers (INA118) and the lowpass filters (LM324). Thus, when the strain signal at the inputs to the differential amplifier is at 0V, the amplified/filtered output of the circuit is 8.73V. The wires from the strain gage are attached so that the strains generated during braking decrease the voltage level of the output. Thus, with an overall board gain of 81 as discussed above, a 10mV strain signal gives an output of 0.81V below 8.73V, or 7.92V. We handled the strain signal in this way because it guarantees that for normal strains generated during braking, the output signals will have an 8.73V range before saturation. (The rail above 8.73V varies between 15 and 19V, depending on the battery level, so it would be hard to guarantee anything if the signal occupied the voltages above the 8.73V reference.) In addition to providing the reference for the circuit, the 8.73V signal from the regulator provides excitation to the strain gage bridge.

Figure 4.10 shows how the pinouts of the conditioning board connect to the rest of the torque sensor. The other major on-wheel electrical elements to the sensor are as follows:

- Batteries: Four nine volt batteries, connected to provide 18V.
- Main Power Switch: `ON` or `OFF`.
- Bridge Excitation Switch and LED: `BRIDGE_AND_AMPLIFIER_ON`—For normal operation (LED = Red). `AMPLIFIER_ONLY_ON`—Useful for analyzing environmental and circuit noise (LED = Green). Excitation to the bridge is cut off, but the bridge output is still amplified and filtered.
- External Socket for Shunt Resistor: Resistors inserted in this socket connect in parallel across the 1994Ω +Excitation to +Signal resistor of the strain gage bridge, allowing offset adjustment.

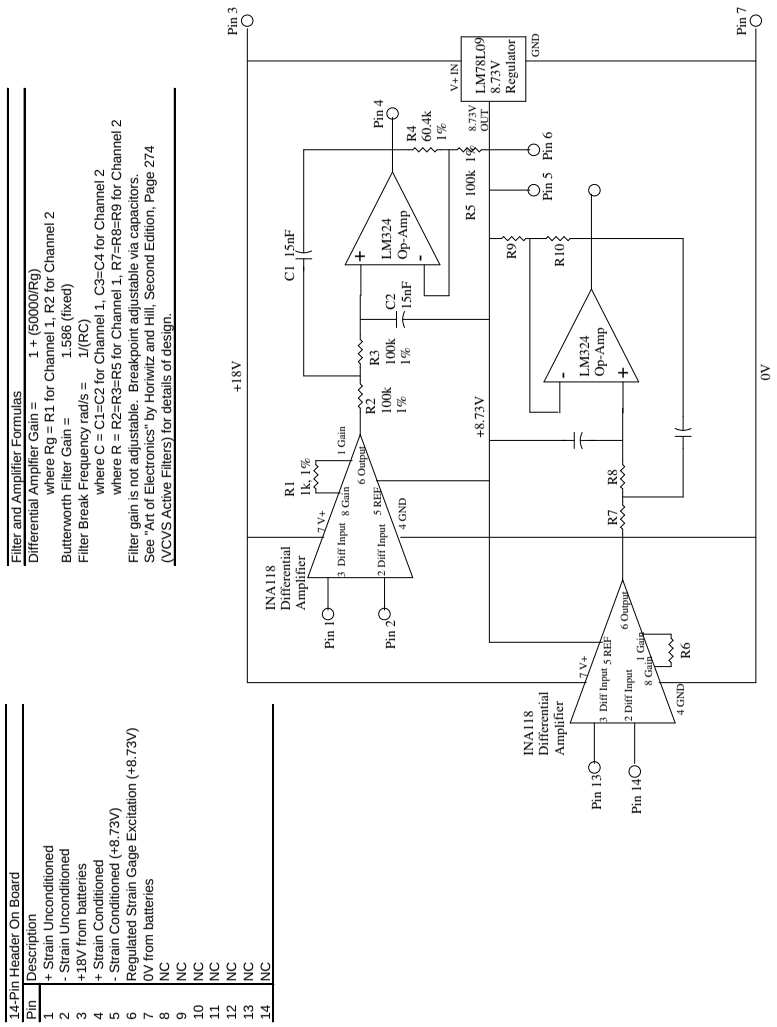

Figure 4.8: The differential amplifier and anti-aliasing filter board riding on-board the wheel conditions the noisy millivolt level strain signals from the strain gage to low-noise 0-8.73V signals suitable for transmission across slip ring. This diagram shows two differential amplifier/filter channels, but only one of them has been implemented at the time of this writing. The second channel (pin 13, pin 14, the associated INA118, R6 (gain resistor), R7, R8, R9, R10, and the two associated capacitors) was intended for a thermocouple if temperature compensation was needed, but it proved to be unnecessary.

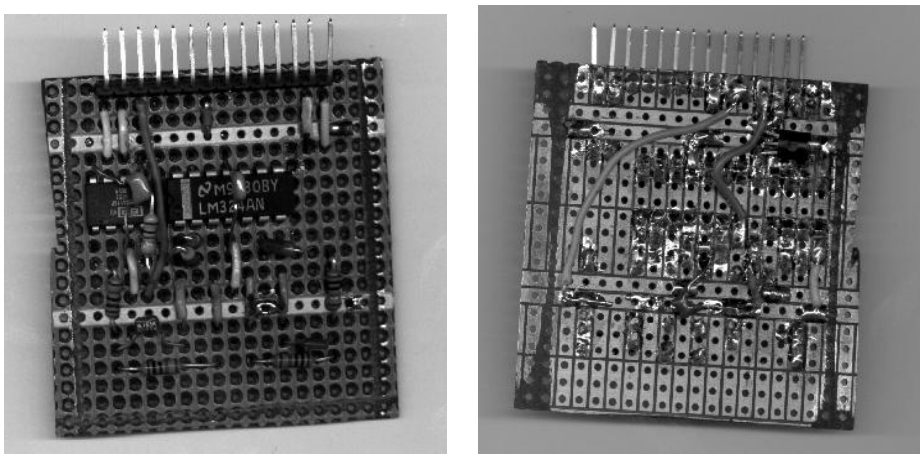

Figure 4.9: On-wheel amplifier and low-pass filter board.

- Strain Bridge: The signal from the bridge is typically 0-20mV, but can reach as much as 100mV during an ABS stop on dry pavement.
- Slip Ring: A four contact slip ring manufactured by Cavotec, Inc., Statesville, NC. The contacts of the ring are pools of mercury, eliminating brush noise.

With the exception of the strain gage bridge, all of the torque sensor elements were mounted on the mechanical bracket described in Section 4.2.1 which follows.

Mechanical Elements

A three piece aluminum bracket houses the electrical elements of the torque sensor and provides a mounting point for the mercury contact slip ring which is used to transmit the torque signal off of the spinning wheel. Figure 4.11 shows an exploded view of the bracket with slip ring (the pencil is for scale). The round bracket at far right press fits into the brake rotor's bearing opening, replacing the standard bearing dustcap. Four small screws assist in holding this bracket to the wheel. The square bracket is hollow and holds the conditioning board inside. In addition, it holds the switches, LED, and batteries. A Velcro strap assures that the batteries will not be torn from their holders as the wheel spins. The square bracket is sandwiched between the round bracket to its right and another round bracket to its left. Long

Figure 4.10: On-wheel wiring of switches, sockets, LED's, and slip ring for torque sensor.

Torque Sensor Wiring External to Conditioning Board

	12 Pin Connector Pinouts
Pin	Description
1	Unconditioned Strain +
2	Unconditioned Strain -
3	+18V From Batteries to Conditioning Board
4	Conditioned Strain + From Conditioning Board to Slip Ring
5	Conditioned Strain - From Conditioning Board to Slip Ring (+8.73V)
6	+8.73V Regulated Strain Gage Excitation from Conditioning Board
7	0V Strain Gage Ground
8	Spare
9	Spare
10	Spare
11	Spare
12	Spare

	Disassembly Instructions
1	Unscrew *double* set screws holding slip ring in place
2	Gently remove slip ring and unplug wires
3	Disassemble battery bracket by removing long allen bolts
4	Locate, disconnect small 2-pin connector
5	Locate, disconnect small 1-pin connector
6	Push large 2-pin connector out bottom of bracket and disconnect
7	Push circuit board and its 12-pin connector out bottom of square bracket
8	Remove 12-pin connector
9	Pull 12-pin connector through top of square bracket

* Shunt resistor is used to imbalance the strain gage. External socket is intended for temporary changes.
More permanent changes to the strain gage offset can be made by soldering a resistor across pins 1 and 6 of 12-pin connnector

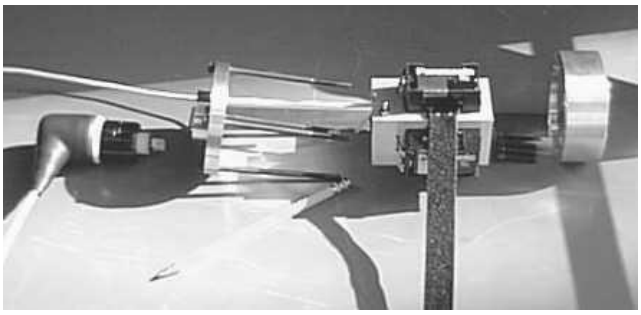

Figure 4.11: Exploded photo of three part bracket mounting batteries, signal conditioning circuitry, and slip ring to wheel.

bolts extend from the left bracket, past the square bracket, and into the right bracket to hold the sandwich tightly together. The left round bracket provides a mounting point for the slip ring (far left). Double set screws and a friction fit assure that the slip ring will not spin free of the bracket.

Figures 4.12, 4.14, and 4.13 show the mechanical details of the three brackets. The project team thanks Kevin Strauss in the Mechanical Engineering Machine shop at UC Berkeley for his excellent design advice and infinite patience in teaching the project team the ways of the machine shop.

Calibration

Figure 4.15 shows the experimental setup we used to calibrate the sensor. A hydraulic lift held the car in the air and a closed-loop brake pressure control system held the brake cylinder pressure at 3447kPa, preventing the wheel from spinning. The tire itself was used as the "lever arm" for the torque, and a hanging mass provided the force for the torque. Thin, lightweight nylon strapping was wrapped around the tire and allowed to hang free on one side. To the end of this strapping, we attached a lightweight aluminum carabiner, from which we were able to hang zero, one, two, three, or four precision 22.69kg masses. The "lever arm" length was calculated by measuring the length of webbing required to wrap around the tire and then calculating the radius.

We recorded the sensor output for each of the five possible loading situations along six locations 60° apart on the wheel. These thirty data points were collected at three strap positions on the tire—one centered, one

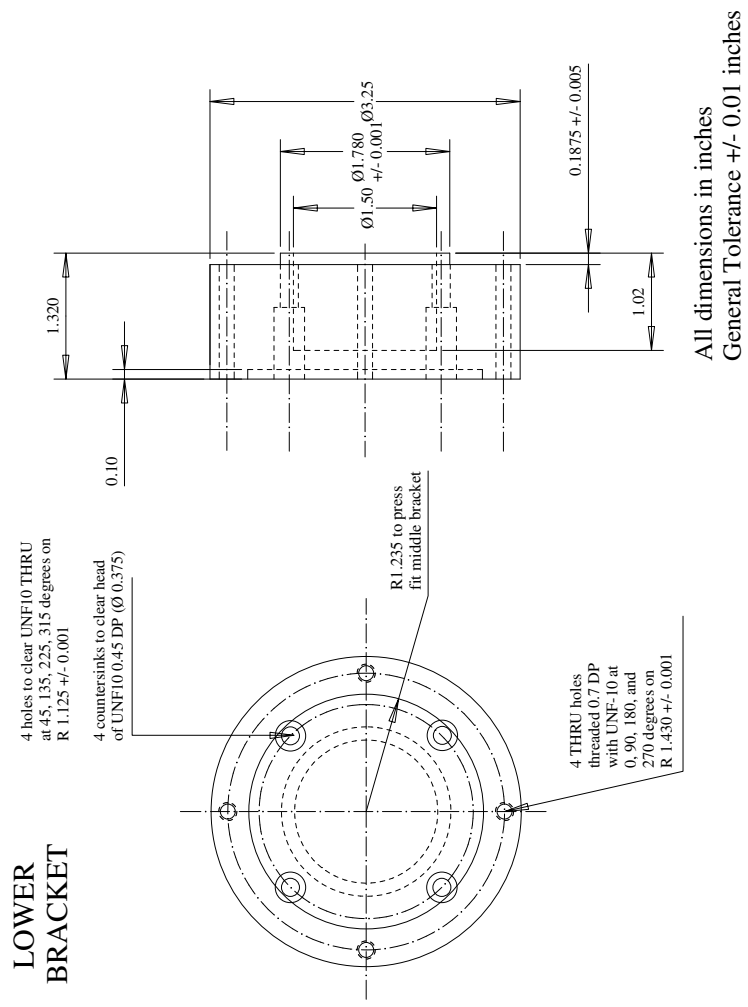

Figure 4.12: Portion of the three part torque sensor bracket which bolts to brake rotor.

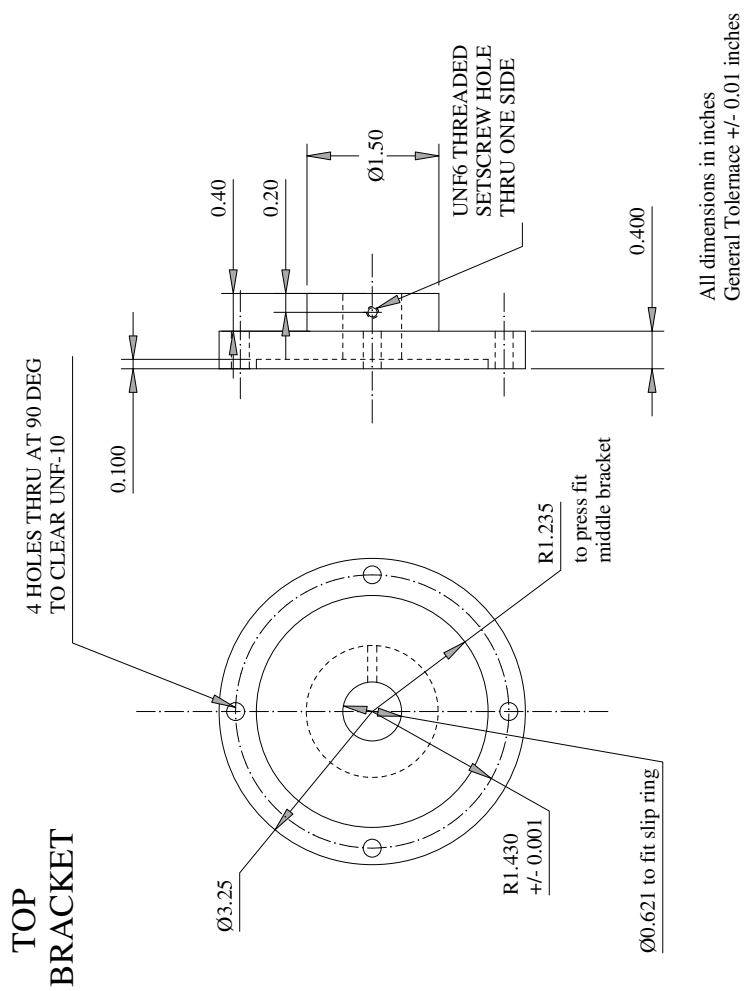

Figure 4.13: Portion of the three part torque sensor bracket housing the slip ring.

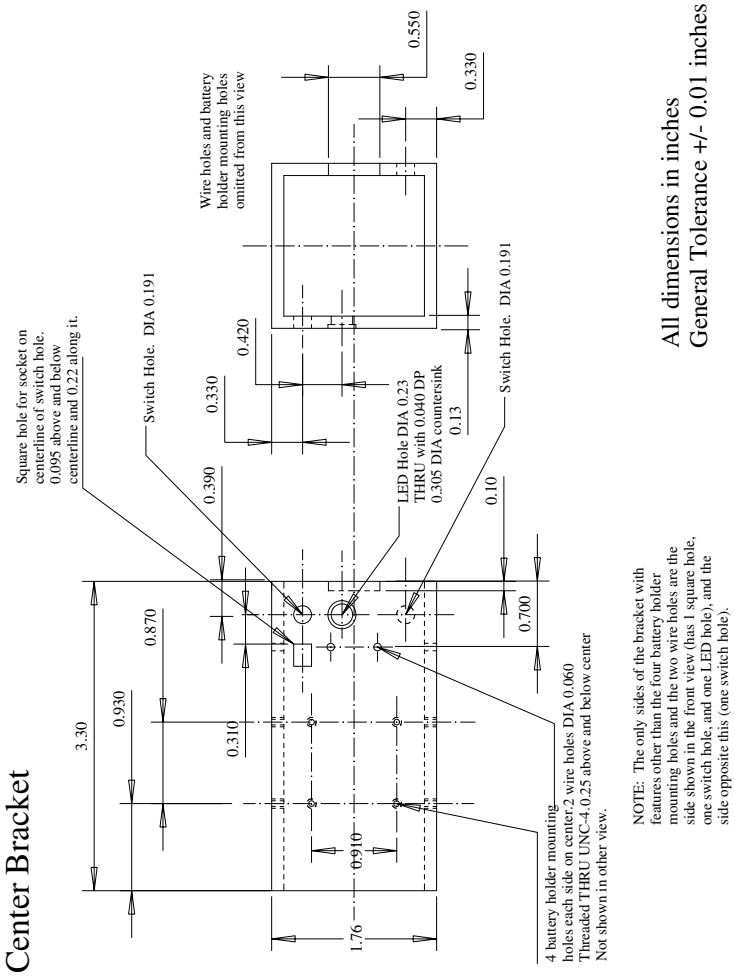

Figure 4.14: Middle portion of the three part torque sensor bracket housing the amplifier/filter board and switches.

shifted towards the outside of the tire, and one shifted towards the inside of the tire. Figure 4.16 shows the results of the calibration done with the strap at the center of the tire. The upper left graph shows the output voltage vs. angular position for five torques. If the only strains in the "shaft in torsion" part of the torque sensor (see Section 4.2.1) were due to torsion, these lines would be flat with higher voltages corresponding to higher torques. Instead, they oscillate about a mean value that increases as the torque increases, indicating that a parasitic strain that depends on angular position is also being measured.

Fortunately, the upper right portion of Figure 4.16 shows that at any given angular position, the torque-to-voltage relationship of the sensor is linear. Even more encouraging is the result shown in the lower left part of Figure 4.16 where the torque is plotted against the average of the voltages at each of the six positions. The result is almost perfectly linear.

To check the effect of parasitic strains introduced the brake pads, the lower right portion of Figure 4.16 plots sensor output voltage vs. position at several brake pressures and zero torque. We see from the large oscillations in the voltage that the brake pads squeezing on the rotor do introduce significant deformation. However, we also see that this pressure-dependent output averages out to be approximately zero over the course of one revolution of the wheel.

To summarize, the calibration shows that the torque sensor's output voltage is linear to brake torque as long as it is averaged over one revolution of the wheel. Both the oscillations in the upper left and the oscillations in the lower right of Figure 4.16 disappear when the brake torque signal is averaged over one revolution of the wheel. Thus, *we can use the lower left curve of Figure 4.16 as our master calibration curve as long as the torque sensor output voltage is averaged over one revolution of the wheel.*

Signal Conditioning

The signal conditioning software for the torque sensor signal is devoted to achieving two things: first, averaging the sensor output voltage over exactly one revolution of the wheel, and second, removing voltage offset. As we mentioned above, the signal needs to be averaged over one revolution of the wheel to average out the effects of parasitic strains in the rotor due to vertical forces and caliper squeezing. The need to remove offset was not mentioned above, but it becomes fairly obvious when one examines Figure 4.16, which

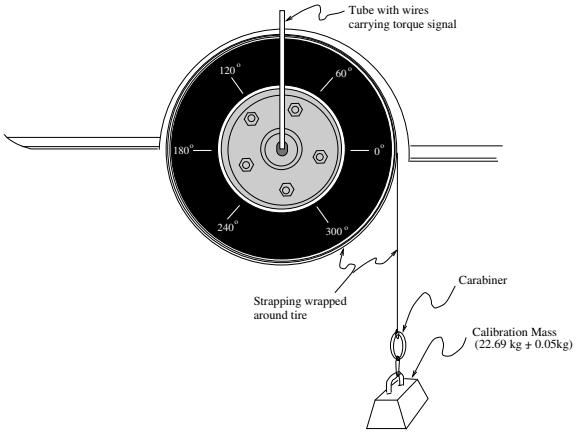

Figure 4.15: Experimental setup for calibrating torque sensor. The car was on a hydraulic lift and the brakes were applied to prevent the wheel from spinning

shows that the sensor outputs a volt or two even when there is no torque.

To average the signal over one revolution, the signal conditioning software continuously tracks at what past sampling time the wheel was at a position 360 degrees away from its current position. It maintains a buffer of the sensor output at all times between present and this "one revolution ago" time and takes the arithmetic average of the samples in the buffer to arrive at the current filtered output. Of course, this introduces lag into the filtered signal, but at high wheel speeds, the lag is only a few tens of milliseconds, and the almost complete elimination of cyclic parasitic signal that the filter provides is worth it. When the wheel slows to near-zero speed the lag becomes unacceptable (and the buffer very large) so the filtering switches to a low-pass type. However, this gives very little improvement in the torque signal and could well be eliminated.

To detect when the wheel on which the sensor is mounted has turned one revolution, a PCTIO-10 digital input/output board counts the teeth of the ABS wheel speed sensor on that wheel and the count at each sample interval is stored in a buffer in the signal conditioning software. To find the time of "one revolution ago" at each sample time, the signal conditioning algorithm notes the current ABS count and then searches the ABS count buffer going backwards into time until it finds a sample interval where the ABS count was 50 less (there are 50 ABS sensor teeth in one wheel revolution).

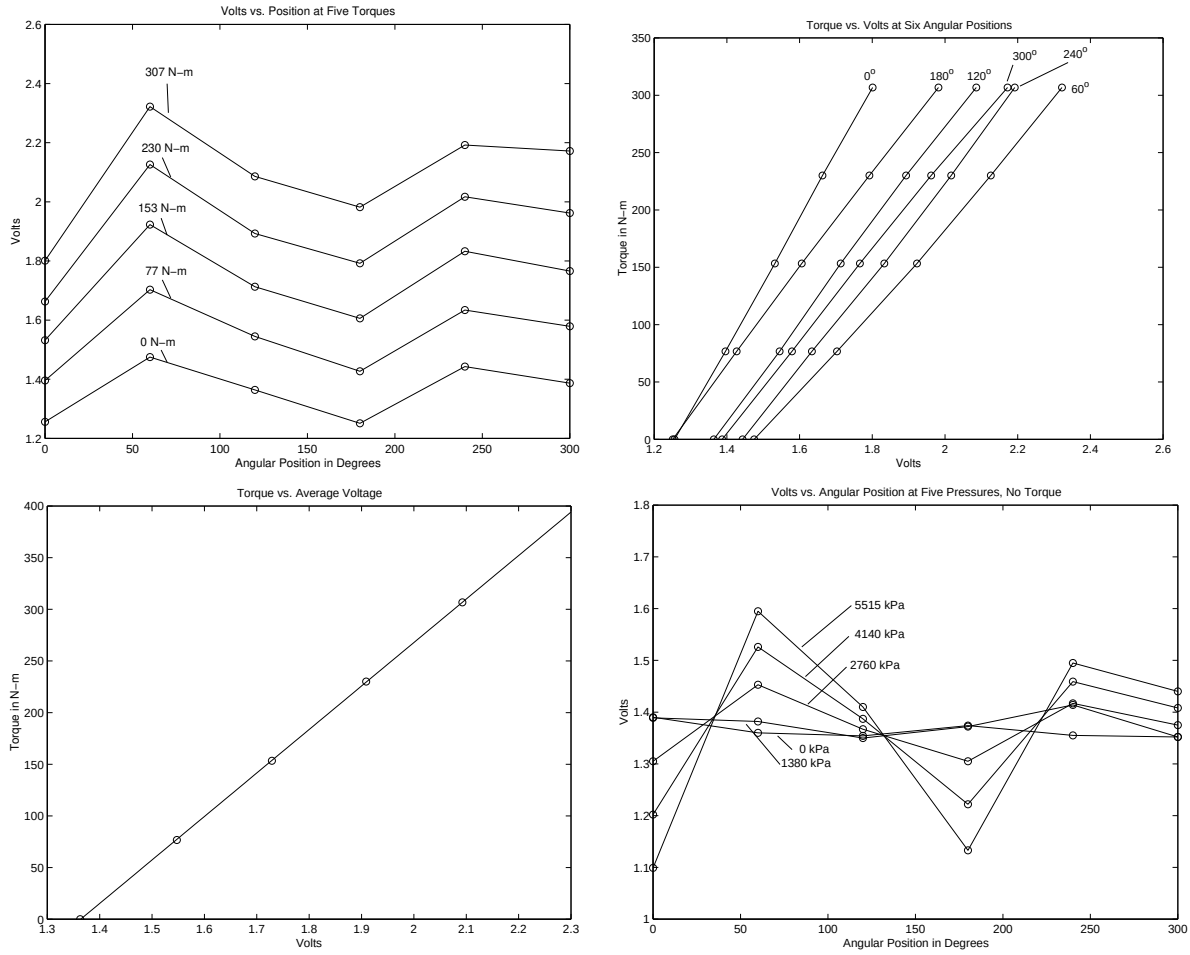

Figure 4.16: Torque sensor calibration data. All tests performed with a brake cylinder pressure of 3447 kPa *Upper Left:* Volts vs. Angular Position, showing variation in torque signal due to parasitic strains. *Upper Right:* Torque vs. Voltage at different angular positions, showing linearity of signal at any given position. *Lower Left:* Torque vs. Average Voltage over one rotation, showing linearity of sensor signal if it is averaged over one cycle. *Lower Right:* Volts vs. Position at several brake cylinder pressures, showing potential corruption that varying brake pressures introduce to torque measurement.

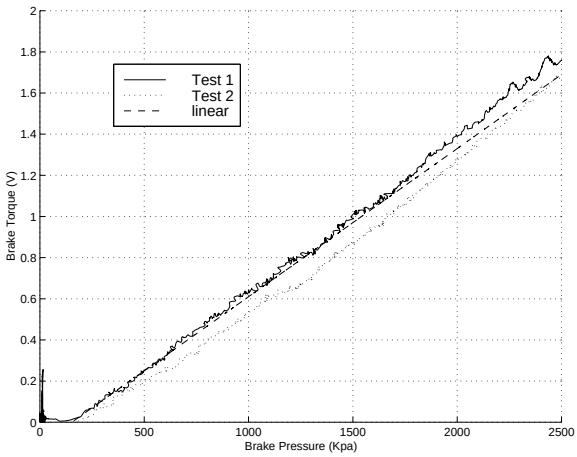

Figure 4.17: Linearity between brake pressure and brake torque sensor output during normal braking.

The offset removal algorithm uses some of the same principals as the algorithm that removes the cyclic component from the torque signal. Whenever the algorithm notices that the brake pressure is close to zero for several hundred milliseconds or more—indicating that the wheel with the torque sensor is free-rolling—the algorithm averages buffered voltage outputs of the torque sensor over a time period corresponding to several wheel revolutions. The resulting voltage offset is then subtracted from the filtered torque signal before scaling to engineering units.

Experimental Results

Figure 4.17 shows the torque sensor voltage output as a function of brake pressure for a maneuver where the brake pressure is gradually increased but does not lock the wheel. As the brake torque model mentioned above predicts, the torque is very nearly linear with brake pressure.

Figure 4.18 shows the torque sensor voltage output as a function of brake pressure for two maneuvers where the brake pressure was increased until the wheel locked. The torque-to-pressure relationship is linear until the wheel locks, at which point the torque drops drastically. The locking pressure is less on gravel than it is on pavement, coinciding with driving experience.

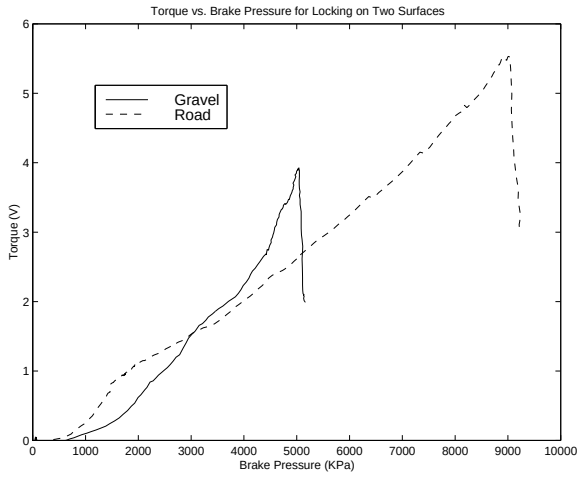

Figure 4.18: Brake torque sensor output vs. brake pressure for a braking maneuver that caused lock-up on gravel and on asphalt.

4.2.2 Wheel Speed Signal Circuitry

Conditioning the torque signal of the previous section required knowledge of wheel position, and almost any vehicle dynamic study needs measurements of the four wheel speeds. Unfortunately, the vehicles at PATH (which were instrumented with platooning in mind) were not outfitted with four individual wheel speed measurements. Therefore, we modified the Red Lincoln Towncar so that the speeds of its four wheels would be recorded by the data acquisition computer.

The output of the inductive, 50-tooth-per-revolution ABS wheel speed sensors that are standard on the car has a slowly time varying offset on the order of two volts. The passing teeth superimpose a malformed sine wave on top of this offset. Both the wave's frequency and its amplitude are proportional to the wheel speed.

To convert this rather odd analog signal to a digital signal that the data acquisition board could read, we constructed the circuit shown in Figures 4.19 and 4.20, and photographed in Figure 4.21.

Experimental results using this new wheel speed acquisition hardware are shown in the next two figures. In Figure 4.22, the car executes longitudinal maneuver and the four wheel speeds are very similar. The small differences are due to radius differences and longitudinal slip. In Figure 4.23, the wheel turns in a tight circle to the left, and the outer wheels travel faster than the

inner wheels.

4.2.3 Differential Braking Hardware

The capability of differential braking was developed on the Red Lincoln Towncar in the PATH fleet. The capability was accomplished by using the ABS hardware, pictured in Figure 4.24, that already existed on the vehicle. Using the existing computer processor that handles all of the real-time control, each of the valves of the ABS unit are controlled to actuate the brakes at each of the wheels individually. In order to functionally accomplish this, software had to be written, and circuitry had to be developed to interface between the ABS hardware and the onboard computer. The circuitry is necessary to buffer the onboard computer so that it won't be damaged from things like voltage spikes or large current draws. Furthermore, the circuit board is needed to amplify control signals from the computer to such a level that the ABS hardware can be actuated.

Figure 4.25 illustrates the layout of the whole ABS Switching Circuit Board used to achieve the capability of differential braking. In addition, photographs of the actual board are displayed in Figures 4.26 and 4.27. Within the car's onboard computer, PC-TIO-10 card #2 is dedicated to handling all of the digital I/O. This card is connected to the switching board via a ribbon cable plugged into the 50 pin connector shown in Figure 4.25. The outputs of the computer consist of control signals for the ABS unit on the car. In all there are eight signals going to each of the ABS valves, and one going to the ABS pump. These signals must each first go through a circuit described in the section below before going out to the hardware via a 12 pin connector. As well as handling outputs, this board takes in inputs as well. Signals from wheel speed sensors on three of the wheels are first transferred in through the 15 pin connector, then jumpered directly to the 50 pin connector before going onto the computer. The board also sends a signal to the onboard computer when the brake pedal is depressed past a certain limit. This information is transmitted through the 12 pin connector and alerts the onboard computer of the need to activate the ABS pump to replenish the brake fluid reservoir. Power for these circuits is supplied by the 15V power supply of the computer, as well as from the car's 12V battery. Furthermore, power for the ABS unit itself is supplied by the car's battery through a 40 amp fuse located in the vehicle's fuse box.

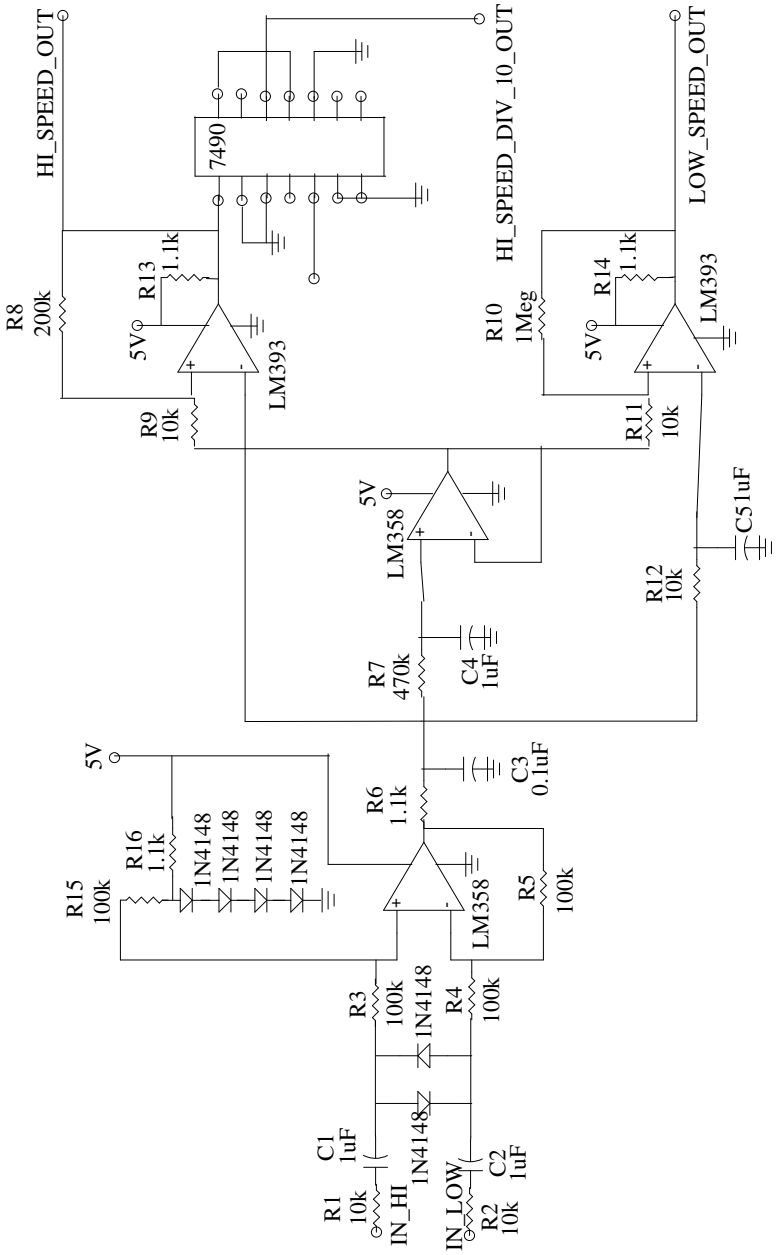

Figure 4.19: A single channel of the circuit used to convert the signal from the wheel speed sensor's magnetic pickups to digital logic suitable for the PCTIO-10. All three channels are identical, so only a single channel is shown.

NINE PIN CONNECTOR PINOUTS

Pin Number	Description
1	Channel 1 raw wheel speed input hi
2	Channel 2 raw wheel speed input hi
3	Channel 3 raw wheel speed input hi
4	Not Used
5	Not Used
6	Channel 1 raw wheel speed input low
7	Channel 2 raw wheel speed input low
8	Channel 3 raw wheel speed input low
9	Not Used

15 PIN CONNECTOR PINOUTS

Pin Number	Description
1	Channel 3 Divided by 10
2	Channel 3 High Speed
3	Channel 3 Low Speed
4	Channel 2 Divided by 10
5	Channel 2 High Speed
6	Channel 2 Low Speed
7	Channel 1 Divided by 10
8	Channel 1 Low Speed
9	Not Used
10	Not Used
11	Not Used
12	Not Used
13	5V Power
14	Channel 1 High Speed
15	Ground

Figure 4.20: Overall layout of the three channel wheel speed signal conditioning board photographed in Figure 4.21. The detailed circuit diagram for the three (identical) channels is shown in Figure 4.19.

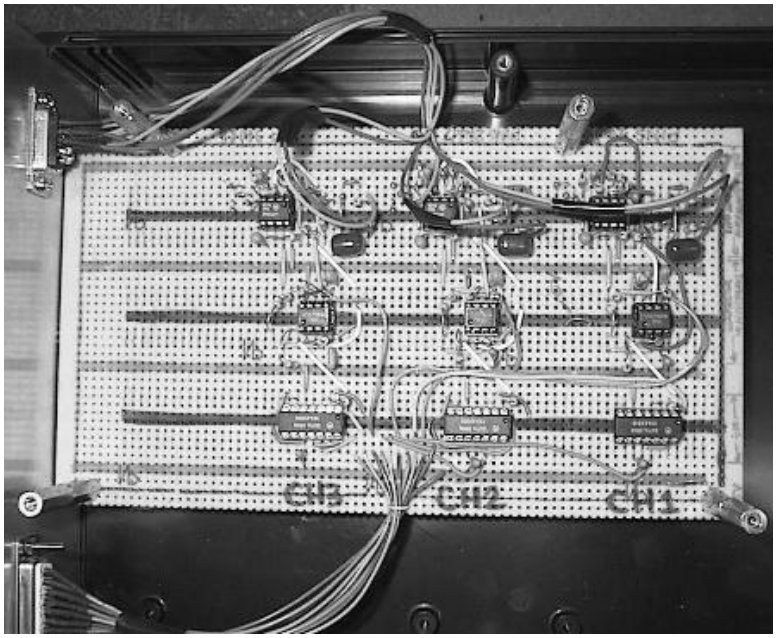

Figure 4.21: Photograph of three speed wheel speed conditioning board detailed in Figures 4.19 and 4.20.

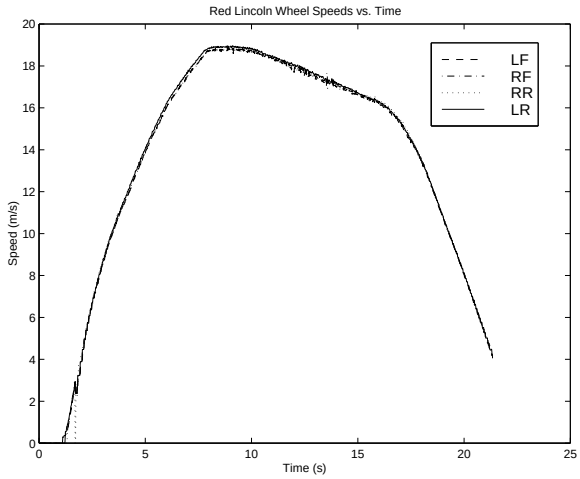

Figure 4.22: Speed measurement from each of the Red Lincoln's four wheels during a longitudinal maneuver.

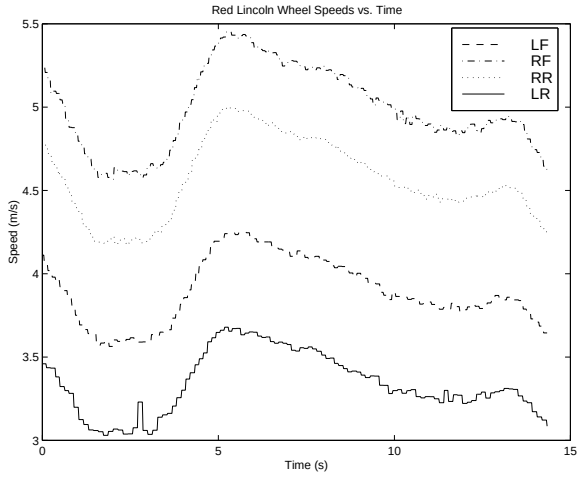

Figure 4.23: Speed measurements from each of the Red Lincoln's four wheels as the car turns circles to the left.

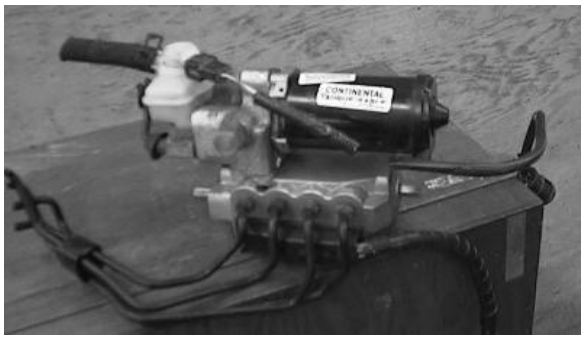

Figure 4.24: Photograph of the Red Lincoln Towncar's ABS unit

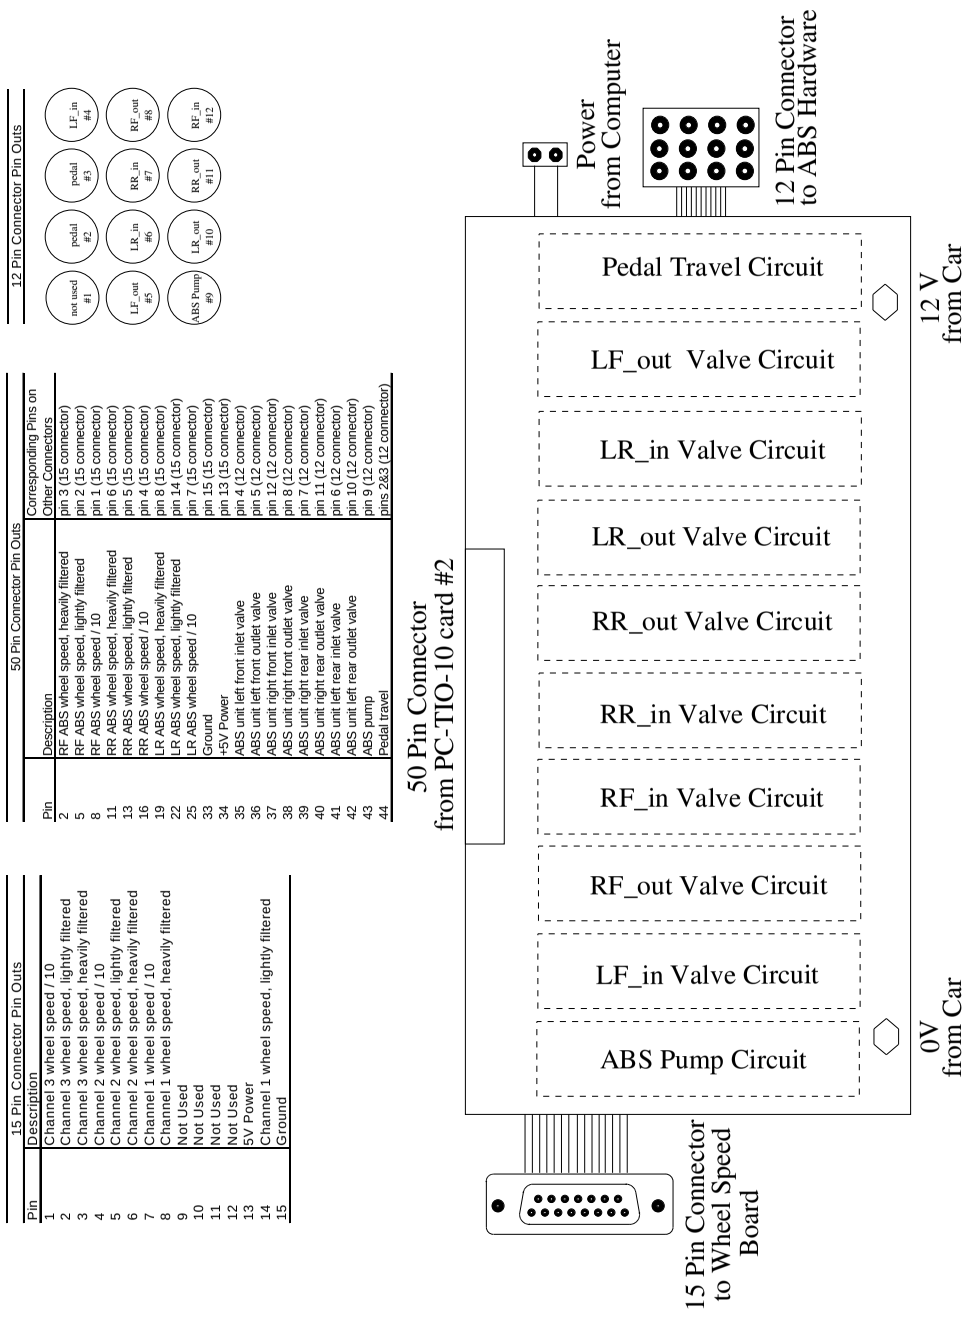

Figure 4.25: Diagram of ABS Switching Circuit Board. Figure 4.28 shows a detailed diagram of the switching circuits for the pump and the eight valves. All nine of these circuits are identical. Figure 4.29 shows details of the pedal travel circuit.

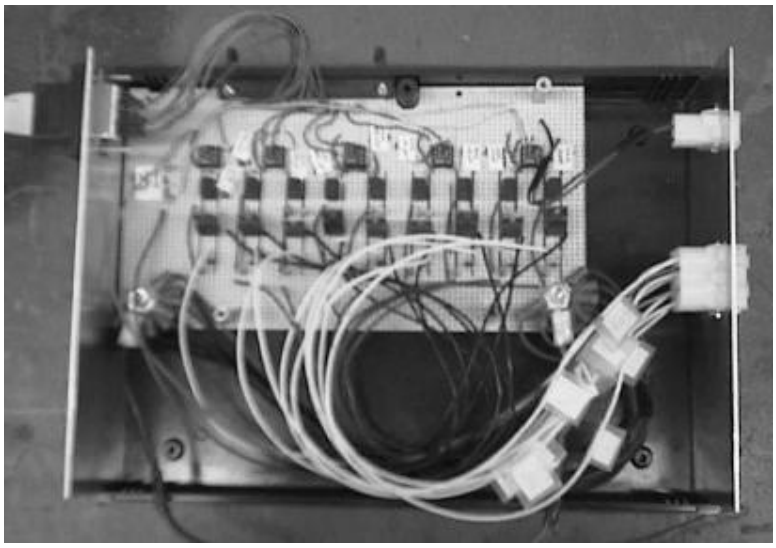

Figure 4.26: Photograph of ABS Switching Circuit Board

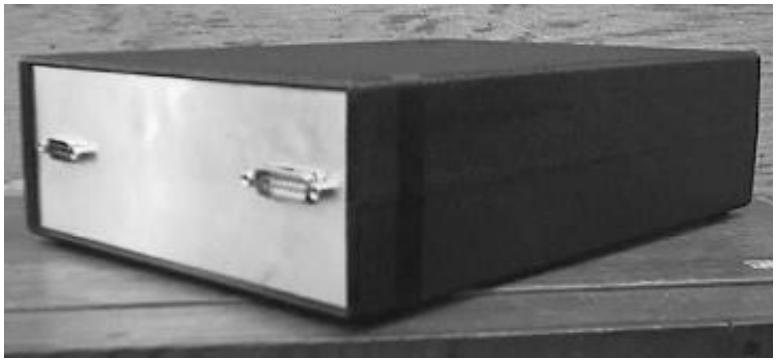

Figure 4.27: Circuit Board Packaging

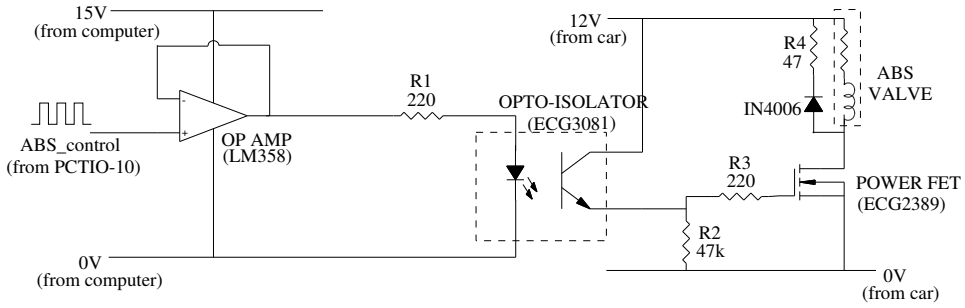

Figure 4.28: Diagram of Circuit Used for Switching the Valves and Pump of the ABS Unit

Switching Circuit

The circuit diagrammed in Figure 4.28 is duplicated nine times on the circuit board, and is used for switching each of the eight ABS valves, as well as the ABS pump. The control signal for each circuit originates from the car's onboard computer. Upon entering the circuit, the signal is passed through an operational amplifier(LM358) configured to be a voltage follower. This effectively acts like a current amplifier to increase the amperage of the control signal so that it can trigger an opto-isolator(ECG3081). The opto-isolator transmits the signal and serves to buffer the computer from the hardware. This isolation protects the onboard computer from voltage spikes and from being forced to source too much current. The signal then goes into a power FET transistor(ECG2389). The resistances R2 and R3 are chosen to limit the current going into the transistor.

The transistor is what is actually responsible for switching the valves and pump on and off. The small level signal entering the gate triggers an amplified level of current to flow from the drain to the source and consequently turn a valve or the pump on. The use of a FET transistor allowed for this switching action to take place very quickly, allowing for high bandwidth control of the valves. The transistor chosen is quite robust in the amount of power it can switch, and in the amount of reverse bias voltage it can withstand. The second characteristic is important due to the back emf from the valve that the transistor experiences. In order to help dissipate this back emf, a diode and resistor are put in parallel with the valve. The value of R4 had to be chosen through trial and error. A large value for R4 increases the rate at which the emf is dissipated, but at the expense of a larger spike. However, if

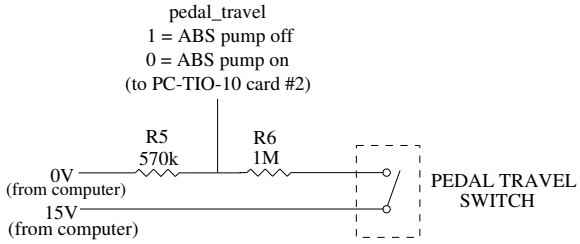

Figure 4.29: Diagram of ABS Switching Circuit Board. Lines labeled (from computer) are attached to rail voltages on the computer's motherboard.

too small a resistance is used for R4, the decay is too slow, and it decreases the bandwidth with which the valves can be controlled.

Brake Pedal Travel Circuit

Another circuit on the ABS Switching Circuit Board is shown in figure 4.29. This circuit is used to indicate when the brake pedal has been depressed past a certain point. This is accomplished using the pedal travel switch located at the brake pedal. 15V are supplied to the circuit by the onboard computer. A voltage divider is set up using resistors R5 and R6 to set the voltage being sent to the onboard computer to 5V(TTL high). This signal stays at the same level until the Pedal Travel Switch is opened, at which point the signal is pulled down to 0V(TTL low). The Pedal Travel Switch is opened when the brake pedal is depressed almost to the point of bottoming out. For this set up, a signal of 0V indicates the need for the ABS pump to turn on and replenish the brake fluid in the brake fluid reservoir. The onboard computer uses the information from this circuit to operate the ABS pump.

4.3 Instructions for Enabling Differential Braking Capability

The capability of differential braking is accomplished through interface with the Red Lincoln TownCar's stock ABS unit. This set of instructions steps through the process for enabling the real time control of the vehicle's ABS unit to take place.

1. Leave all connections in the trunk of the vehicle as they are. The switch

between normal ABS operation and operation under the control of the personal computer onboard the vehicle can be accomplished without changing the setup in the vehicle's trunk.

2. Find the ABS unit. Open the hood, the unit is located at the front left of the car, right below the air cleaner. It is distinguishable by the fact that it has a silver metal casing (the valves) with four numbered brake lines exiting it. It also has a black cylinder (the pump) mounted on top of it.

3. Next find the 19-pin connector leading from right below where the brake lines exit the unit. Disconnect the ABS unit from the car's stock control unit.

4. Find the ten individual wires wrapped with white tape exiting from the main wire bundle that run's the length of the car. Individually connect each wire with the corresponding number indicated on the male 19-pin connector.

5. Find the relay connected to the ABS pump. This relay has two round connections on it. Remove the 6-pin connector (the larger of the two).

6. Find the four individual wires wrapped with white tape exiting the main wire bundle. Individually connect each wire to the corresponding number indicated on the relay.

7. Now find the pedal travel switch connection. It is a gray two pin connector located on the right hand side of a bracket surrounding the brake pedal. Disconnect the free side of the connector.

8. Find the black and white intertwined wires with red connectors. Attach these two wires to the male connector. Order of the wires doesn't matter.

9. To switch back to standard ABS operation, reverse the steps of this procedure.

4.4 Experimental Results

Figures 4.30 and 4.31 show experiments used to verify that the differential braking hardware worked properly. The brake actuator is used to increase

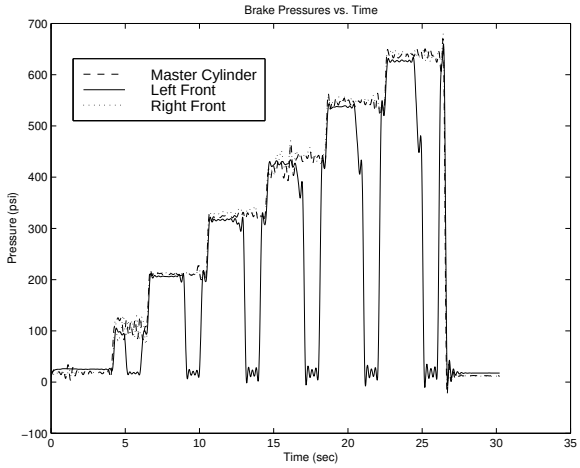

Figure 4.30: Differential Braking Hardware Demonstration: Master cylinder pressure is set to 100psi, 200psi, 300psi, 400psi, 500psi, and 600psi. For each pressure step, left front wheel is isolated by closing the inlet valve, bled by opening the outlet valve, and then re-connected to the master cylinder by closing the outlet and opening the inlet valve.

the master cylinder pressure in steps, and for each step the target wheel is disconnected from the master cylinder, bled, and then reconnected to the master cylinder.

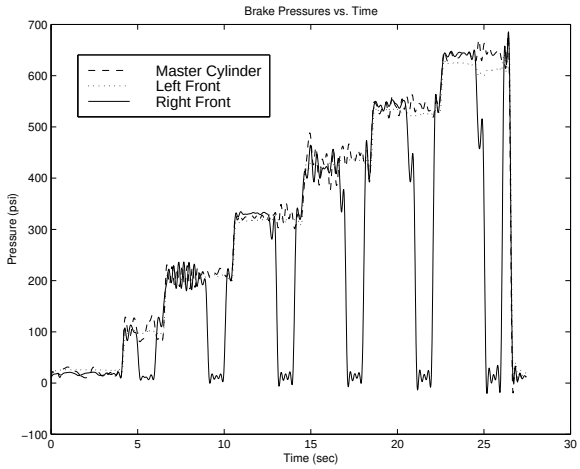

Figure 4.31: Differential Braking Hardware Demonstration: Master cylinder pressure is set to 100psi, 200psi, 300psi, 400psi, 500psi, and 600psi. For each pressure step, right front wheel is isolated by closing the inlet valve, bled by opening the outlet valve, and then re-connected to the master cylinder by closing the outlet and opening the inlet valve.

Chapter 5

Conclusion

MOU 308 offered solutions to two common brake control problems: variable brake torque gain and brake rotor-induced brake torque oscillations. It also produced a fairly advanced vehicle test bed with differential braking, torque measurement, and individual wheel speed measurements.

Two adaptive control solutions for the variable brake torque gain problem were shown to work experimentally. One of these algorithms used a "smooth" adaptive control law where the derivative of the unknown parameter was proportional to the tracking error, and the other algorithm used a "nonsmooth" adaptive control law where the derivative of the unknown parameter was proportional to the sign of the tracking error. Both laws worked well, but the nonsmooth one seemed to give less erratic results. With some "sanity checking" and other logic to prevent runaway parameters, either of these algorithms would be suitable for an implementation and would improve performance.

To eliminate brake rotor-induced vibrations, a nonlinear version of a so-called Adaptive Feedforward Cancellation algorithm was developed and simulated. Simulated results showed that the algorithm successfully reduced rotor vibrations, even when the actual vibrations were more complex than expected. Because these rotor vibrations were intermittent and did not pose much of a problem on our test vehicle, we did not test the AFC algorithm on a real vehicle. Thus, future researchers should use this algorithm with caution.

Detailed hardware documentation was provided so that future researchers can either make use of the test bed or use the experience gained in its construction to create an improved test vehicle. Although we added

numerous sensors to the Red Lincoln test bed, it would have been helpful to have more data, especially as we began to use the Red Lincoln for more fundamental vehicle dynamic studies on another project. Specifically, the yaw rate, lateral acceleration, suspension deflections, and absolute position and velocity of the car would have been useful. In addition, it would have been very helpful if our test vehicle had been street legal. This points to a need for a street-legal, highly instrumented, open-loop test vehicle for basic vehicle dynamics investigations.

Bibliography

[1] Yasushi Amano, Katsihiro Asano, Shigenobu Okada, Takahiko Mori, and Norio Iwama. Model following control of 4wd hybrid vehicles. volume 8, pages 130–135. 11th IFAC World Congress, Tallin, 1990.

[2] J.P. Aubin and A. Cellina. *Differential Inclusions.* Springer-Verlag, 1984.

[3] Marc Bodson, Alexei Sacks, and Pradeep Khosla. Harmonic generation in adaptive feedforward cancellation schemes. pages 1261–1266, Tucson, AZ, December 1992. Proceedings of the 31st Conference on Decision and Control.

[4] J. Eric Bowman and E.H. Law. A feasibility study of an automitive slip control braking system. *SAE Transactions*, (930762):1166–1192, 1993.

[5] Chieh Chen and Masayoshi Tomizuka. Steering and independent braking control for tractor-semitrailer vehicles in automated highway systems. pages 1561–1566. IEEE Conference on Descision and Control, December 1995.

[6] Degang Chen and Brad Paden. Nonlinear torque-ripple cancellation for step motors. pages 3319–3324, Honolulu, Hawaii, December 1990. IEEE Conference on Decision and Control.

[7] Sergey Drakunov, Ümit Özgüner, Peter Dix, and Behrouz Ashrafi. Abs control using optimum search via sliding modes. *IEEE Transactions on Control Systems Technology*, 1995.

[8] Leonard Evans and Peter H. Gerrish. Anitlock brakes and risk of front and rear impact in two-vehicle crashes. *Accident Analysis and Prevention*, 28:315–323, May 1996.

[9] Jonathan Frankel, Luis Alvarez, Roberto Horowitz, and Perry Y. Li. Safety oriented maneuvers for ivhs. pages 688–672. American Controls Conference, Seattle, WA, June 1995.

[10] J. Christian Gerdes. *Decoupled Design of Robust Controllers for Nonlinear Systems: As Motivated by and Applied to Coordinated Throttle and Brake Control for Automated Highways.* PhD thesis, University of California at Berkeley, 1996.

[11] Thomas D. Gillespie. *Fundamentals of Vehicle Dynamics.* Society of Automotive Engineers, Inc., Warrendale, PA, 1992.

[12] J. Karl Hedrick. Brake system modeling, control and integrated brake/throttle switching. PATH Project Proposal, 1994.

[13] Jwu-Sheng Hu and Masayoshi Tomizuka. Adaptive asymtotic tracking of repetitive signals–a frequency domain approach. pages 2621–2627, Boston, MA, 1991. American Controls Conference, Boston, MA.

[14] J.K. Hurtig, S. Yurkovich, and K.M. Passino. Torque regulation with the general motors abs vi electric brake system. pages 1210–1211. American Control Conference, Baltimore, MD, June 1994.

[15] Dragos Bogdan Maciuca. *Nonlinear Robust Control with Applications to Brake Control for Automated Highways.* PhD thesis, University of California Berkeley, 1997.

[16] Marques Monteiro and D.P. Manuel. *Differential Inclusions in Nonsmooth Mechanical Problems.* Birkhäuser, 1993.

[17] Giuseppe Puleo. Automatic brake proportioning devices. *SAE Transactions*, (700375):1181–1202, 1970.

[18] Alexei Sacks, Marc Bodson, and Pradeep Khosla. Experimental results of adaptive periodic disturbance cancellation in a high performance magnetic disk drive. pages 688–690. American Controls Conference, San Francisco, CA, June 1993.

[19] M.A. Salman. A robust servo electronic controller for brake force distribution. volume 3, pages 1963–1968. American Control Conference, 1989.

[20] Pavel A. Shavrin. Control of independent rear wheel drive vehicle. pages 4380–4385. IEEE Conference on Decision and Control, New Orleans, LA, December 1995.

[21] Daniel Shevitz and Brad Paden. Lyapunov stability theory of nonsmooth systems. In *Advances in Robust and Nonlinear Control Systems*.

[22] Jean-Jacques E. Slotine and Weiping Li. *Applied Nonlinear Control*. Prentice Hall, 1991.

[23] Walter Stringham, Peter Jank, Jerry Pfeifer, and Alex Wang. Brake roughness–disc brake torque variation, rotor distortion and vehicle response. *SAE Transactions*, (930803):1235–1247, 1993.

[24] Anton T. van Zanten, Rainer Erhardt, and Georg Pfaff. Vdc, the vehicle dynamics control system of bosch. *SAE Transactions*, (950759):1419–1436, 1995.