Design and Adaptive Control of a Lab-based, Tire-coupled, Quarter-car

Suspension Test Rig for the Accurate Re-creation of Vehicle Response

by

Justin D. Langdon

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

Master of Science

In

Mechanical Engineering

_____

Dr. Steve Southward, Chairman

_____

Dr. John Ferris

_____

Dr. Corina Sandu

January 31, 2007
Danville, VA

Keywords: quarter-car, test rig, adaptive inverse control, system identification, suspension, vehicle response replication

Design and Adaptive Control of a Lab-based, Tire-coupled, Quarter-car Suspension Test Rig for the Accurate Re-creation of Vehicle Response

by

Justin Langdon

Dr. Steve Southward, Chairman

Abstract

The purpose of this study has two parts directed toward a common goal. First, a state-of-the-art quarter-car test platform has been designed and constructed to offer increased testing flexibility at a reasonable cost not found commercially. With this new test rig completed, the second objective is a proof-of-concept evaluation of a well known adaptive control algorithm applied to this new quarter-car test rig for the purpose of replicating the dynamic suspension response, such as a response that was recorded during a road test. A successful application of this control algorithm on the quarter-car rig is the necessary first step toward its application on an 8-post test rig for a direct comparison to current practices.

Before developing a new test rig, the current state-of-the-art in quarter-car rigs was first evaluated as well as indoor vehicle testing in general. Based on these findings, a list of desired functional requirements was defined for this new design to achieve. The new test rig was built and evaluated to determine how these goals were met and what the next steps would be to improve the rig. The study then focused on evaluating control policies used for reproducing dynamic responses on vehicle road simulators such as 4-post and 7-post shaker rigs. A least-mean squares (LMS) adaptive algorithm is introduced and applied first in software using a linear two-mass quarter-car model, and then to the actual hardware-in-the-loop quarter-car rig.

The results of the study show that the resulting quarter-car test rig design is quite flexible in its ability to test a multitude of suspension designs and also its ability to accommodate new hardware in the future such as a body loaders. The study confirms that this particular implementation of the LMS algorithm is a viable option for replicating test vehicle response on an indoor quarter-car test rig. Thus, a future study to compare the use of this algorithm to the current industry standard batch processing method is possible.

Acknowledgments

I really want to thank my girlfriend Stephanie for putting up with so much crap. You supported me in ways that I would never dream of asking anybody. You handled my traveling and stress way better than I ever could.

Most of all I want to thank God and my parents, Bonnie and Dave. Without you, none of my achievements in both school and life would be possible. I will never be able to repay you for the doors you have been able to open up for me. Without you guys I would probably be looking up car parts behind a counter in town. I don't think a person could ask for better parents. I want to thank all of my family for their support in my school and work endeavors.

Contents

List of Nomenclature

$\mathbf{X}_k$	discrete input vector at time k
x_k	input vector element at time k
T_s	sample time in seconds
z^{-n}	n-sample delay
w_{kn}	n^{th} sample weight at time k
y_k	output at time k
$\mathbf{W}_k$	weight vector (FIR model) at time k
d_k	desired response signal at time k
ε_k	error at time k
ξ	expected value function
$\mathbf{R}$	input auto-correlation matrix
$\mathbf{B}$	cross-correlation between desired and input signals
$\mathbf{W}^*$	optimal weight vector
μ	small gain constant or step size
∇	gradient
$\mathbf{H}$	linear transfer function matrix
$\mathbf{x}$	state space state vector
$\mathbf{y}$	state space output
$\mathbf{u}$	state space input
$\mathbf{A},\mathbf{B},\mathbf{C},\mathbf{D}$	continuous-time state space matrices
$\mathbf{A}_k,\mathbf{B}_k,\mathbf{C}_k,\mathbf{D}_k$	discrete-time state space matrices
g_k	unknown plant output
n_k	plant noise
$\mathbf{P}$	unknown plant
$\mathbf{R}_k$	filtered-x desired signal

$\hat{\mathbf{P}}$	identified plant (FIR model)
m_s	sprung mass
m_u	unsprung mass
k_s	suspension stiffness
c_s	suspension damping coefficient
k_u	tire stiffness
ζ	damping ratio
z	sprung mass degree-of-freedom coordinate
y	unsprung mass degree-of-freedom coordinate
x_{in}	road input to tire patch
T	kinetic energy function
V	potential energy function
D	damping pseudo-energy function
q_i	i^{th} degree-of-freedom coordinate

List of Figures

List of Tables

1. INTRODUCTION

This chapter provides motivation for the research presented in this thesis by describing some of the difficulties inherent in vehicle development. Next, the objectives and approach to achieve the goals of this study are explained. The chapter ends with an outline of the thesis.

1.1. Motivation

For as long as products have been under development, engineers have struggled with the trade-off between research and development time and quality or performance of the product under development. This is especially true in the automotive and motorsports industries. In the automotive industry, it was said that the necessity to continuously improve the efficiency of research and development was due to things such as changing market demands [1]. In the motorsports industry the same is true, although the 'market' in this case is dictated by the performance of one's competition. Other constraints in motorsports are restrictive rules that govern the specifics of a racecar's design or that limit actual track testing that a team may perform in a given season.

A major trend in both industries is to utilize more indoor lab-based test equipment. It was noted that it is near 'compulsory' for a F1 race team to have a 7-post shaker at its disposal [2]. Testing in a laboratory environment allows for greater control of the experiment, less time required which allows for more experimentation, and, in many cases, much lower costs and liabilities [3]. Specifically, it was noted that a laboratory test generally requires one-fourth of the time as a road test and is less expensive per test. It reduces the need to have support personnel such as safety crews during a track test. Lab testing also reduces the liability as a test driver could be injured during a road test. In many cases, the driver can be eliminated during a laboratory test. One primary mission of the Virginia Institute for Performance Engineering and Research (VIPER) Lab, the research group from which this study stems, is to develop and evaluate the technology that has the potential to advance the state-of-the-art in 8-post testing. The goals of this study seek to support this mission.

1.2. Objectives

The primary goal for this study is to develop and evaluate technology that can improve the state-of-the-art in indoor vehicle simulation testing on a quarter-car vehicle. Because this is a start-up lab, an additional, necessary objective is to develop a state-of-the-art quarter-car test bed for this research.

Currently, the desired functions of a quarter-car rig include the ability to mount several different designs of actual car suspensions, have the ability to perform a wide range of tests, which include body loads on the sprung mass, and still have the ability to expand for future developments. There are several quarter-car test rigs that are currently available. Many of those rigs do not have the functionality desired for the current and future research projects. The commercially available rigs are quite expensive and yet most still do not offer the amount of flexibility needed. In order to meet the combination of functional requirements necessitated by VIPER, a commercially developed rig from a company such as MTS or ServoTest would require close development with the researchers. This would prove to be an even more expensive proposition. Thus, the goal of this research was to develop the new rig, in house, in order to achieve the desired functionality at a reasonable expense.

With a competed test rig the other goal is to perform a study on the application of a control system to the rig for reproducing test vehicle response. This is not a new concept. In fact, much of the literature suggests that several methods have been attempted, and commercial solutions are available. However, most of the literature is vague in terms of the technical details. Based on personal communications with engineers in both the racing industry and the test equipment industry there is a need to make indoor testing more efficient, by reducing time on the rig. Most manufacturers of indoor test equipment, like a 7-post or now an 8-post, supply software that allows for the replication of vehicle response. Very little information is available in the public domain that discusses the details of how these software tools function.

After a review of the literature, an application of a different control algorithm, not currently in use, is merited. However, to simplify the proof-of-concept evaluation of its application to indoor vehicle testing, this control scheme has been applied to a quarter-car rig as opposed to the more complex test beds. Because of this simplification, a true

comparison of the performance of this algorithm to existing software, on the more complex test rigs, is not yet possible. However, value can still be gained since the more complex systems implement an extension of the same principles. Thus with this concept of application proved, a future study will be to apply these principles to an 8-post rig for a direct comparison of performance between the proposed method and those that are currently in use.

In summary, a new quarter-car test rig has been designed and built. In an attempt to prove the concept, a well known control scheme, not currently used for this application, is applied to the problem of replicating vehicle response signals. Finally, the last goal is to be the first to publish results, such as convergence rates, of such a study.

1.3. Approach

To achieve these goals the following approach is taken. Current quarter-car test rigs and the state of indoor testing in general are first evaluated. The approach is to develop a quarter-car test rig that addresses many of the shortcomings found with the state-of-the-art while trying to minimize expense. In the future, the design will allow for expanded functionality as defined in this paper.

Having constructed a quarter-car test rig, the next step is to evaluate an adaptive inverse control (AIC) algorithm as it applies to vehicle dynamics testing. Specifically, a least-mean squares (LMS) algorithm is utilized to reproduce vehicle acceleration response from a previous test. This algorithm is first simulated purely in software to evaluate its effectiveness. Finally, a real-time implementation of the algorithm is applied to the quarter-car test bed for hardware-in-the-loop testing to validate the control scheme. Actual acceleration response will be reproduced on the sprung and unsprung mass of the quarter-car rig.

1.4. Outline

The following is a brief outline of the chapters to come. Chapter two provides the background for this study which includes a literature review of current quarter-car test

rigs. Also reviewed are the current control strategies used on more complicated rigs such as 4-post, 7-post and 8-post test rigs. Chapter three details the development of the new quarter-car test rig completed for this study. The functional requirements are stated and achieved. Chapter four introduces the reader to the control algorithm concepts utilized in this study. The building blocks for the adaptive algorithm are explained in detail. Chapter five applies these concepts in simulation with the use of a simplified linear quarter-car model. The results of this simple study are then discussed. Similarly, in chapter six, the software simulated quarter-car model is now replaced with actual hardware. The same tests are then run with hardware-in-the-loop. The results of this study are provided and compared to those of the simulation work. Finally, the thesis ends with the conclusions and recommendations provided in chapter seven.

2. LITERATURE REVIEW

In this chapter, a background of the current state-of-the-art in vehicle testing rigs and the controls they utilize is reviewed. The chapter begins with a survey of the current test rig technology and some of the issues or deficiencies found with them. This will lay the groundwork for defining the new requirements of the quarter-car rig design presented in this thesis. The second half of the chapter reviews the current control algorithms in use on indoor vehicle shaker rigs. The differences between these current algorithms and the one used in this study are highlighted. These differences will become more obvious when the control details are discussed in Chapter 4.

2.1. Vehicle Testing Rigs

It is clear that the main purposes of a shaker rig regardless of the number of posts are to evaluate noise, vibration and harshness (NVH), perform durability tests, and/or improve vehicle performance [4, 5, 6]. These goals vary slightly depending on the nature of the industry in which they are applicable. The automotive manufacturing segment would primarily be interested in NVH and durability but on some occasions may want to improve the handling of their vehicle without spending countless hours on a proving ground. The racing industry is slightly different. They are not as likely to be interested in NVH, however durability and particularly performance metrics, such as handling and suspension tuning, are critical. For these tests, knowledge of the road input is extremely important, especially if the desire is to simulate the surface of the track and characterize vehicle response. If complete knowledge of the road profile were known, this would undoubtedly improve the efficiency of 7-post software. However, this information is seldom actually known, particularly in the motorsports industry. Even if the road profile is known precisely, it tells the test engineer nothing about other dynamics the vehicle endures such as inertial and aerodynamic loading. This information is often calculated based on maps, lookup tables, or vehicle models running in software such as ADAMS. These types of software run in conjunction with the simulator. [5, 6]

This section discusses the current state-of-the-art in vehicle testing rigs. Particularly, a survey of quarter-car test rig technology is presented which details the

need for increased functionality of the new quarter-car test rig. Finally, the section closes with the proposed new functional requirements of said test rig.

2.1.1. Complex Shakers

Among the most complex test equipment are the 4-post, 7-post and 8-post shakers and kinematic and compliance rigs. A typical 4-post rig is comprised of 4 servo actuators. If the test rig is tire-coupled, each actuator (post) supports the vehicle under each tire.   If spindle-coupled, each spindle on the vehicle is mounted directly to the actuator. Thus, the test rig can input various signals into the vehicle and responses may be measured. The 7-post works in a similar fashion with the addition of three extra actuators (4 extra actuators if it is an 8-post) attached between ground and the sprung mass of the vehicle. These offer increased capability in the form of simulating vehicle response from inputs such as braking, acceleration and cornering as well as aerodynamic loading.

Figure 2-1 Image of ServoTest 7-post Test Rig (reproduced with permission)

Figure 2-1 represents a ServoTest tire-coupled 7-post rig with a Formula 1 racecar. These complex test rigs offer an immense amount of capability, however they

are very expensive to build and maintain. They also present other difficulties. These rigs are very sophisticated muli-input/multi-output (MIMO) systems which require a high degree of control knowledge and understanding to use properly. Often, the complex nature of these multivariable problems requires multi-step iteration to obtain a suitable drive file for commanding each of the actuators. Once converged data is extracted from tests run on these systems, it is often very difficult to interpret and correlate to the real-world counterpart. Some reasons for these issues with more complex test rigs are the lack of literature and other available documentation [7, 8]. To the authors knowledge only a handful of papers that discuss multi-post test to any detail exist [2, 4, 9, 10]. It is likely that the lack of available information is partially due to race teams and automotive companies trying to protect their competitive advantage.

2.1.2. Current Quarter-car Rigs

As an answer to the high complexity and expense of these systems, simpler test beds such as the quarter-car test rig are used. A rig such as this reduces the complexity greatly by only focusing on one corner or quarter of the vehicle. These may be considered one-post or two-post systems. Often, these systems can be viewed as a single-input/single-output (SISO). This greatly reduces computational time and complexity and often a closed form solution may be reached. This allows for much better understanding of both the problem and results.

There are some inherent deficiencies with many of the existing test rigs. Often, the suspension components of the vehicle are simplified or removed altogether. To the author's knowledge, the literature does not discuss a test rig that can incorporate multiple suspension designs. This may be desired to reduce the expense and to add flexibility. Several companies such as MTS and ServoTest are capable of developing such rigs but the expense and time required for such designs would be high. Many times, the tire dynamics are omitted altogether and represented by a simple linear spring. Removing the non-linear geometry does make the analysis simpler. This can allow for reduced computations and easier implementation into simulation with the quarter-car as hardware-in-the-loop. While this does tend to simplify the understanding of the problem, it

introduces a new issue. Good correlation of the results from a quarter-car test with real data from the represented vehicle is difficult to obtain. This is often due to the choice of sprung and unsprung masses being represented as linear moving masses. To constrain this linear motion, guide bearings are used which are often insufficient. If the bearings have a highly non-linear friction, this can create problems in the dynamics of the represented response. Often, the suspension and tire are represented by equivalent springs which likely do not have the same overall non-linear characteristics, or may not have the same linear characteristics over the range in which the test is operating. These shortcomings, along with eliminated or simplified suspension geometry, can lead to a gross error in the representation of the actual vehicle. Figure 2-2 is an example of one such quarter-car test rig. In this system, the suspension compliance of a vehicle is wholly represented by a set of air springs. The tire is represented by elastomeric mounts. The sprung mass is represented by a sliding carriage carrying lead weights. The bearings are roller wheels and Teflon bushings running in grooves on extruded aluminum uprights [11]. The rigidity of the load frame is also a potential issue due to the vast amount of bolted joints and aluminum extrusion construction.
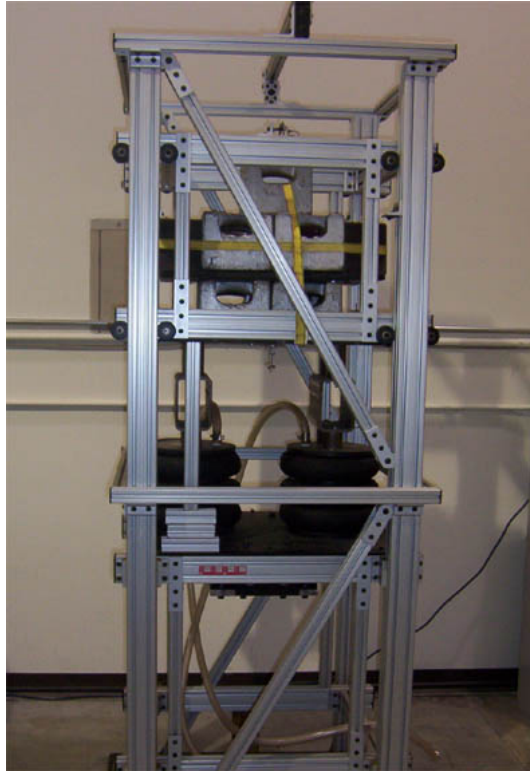
Figure 2-2 Simplified Quarter-car Test Rig (VT AVDL)

Other limitations of conventional quarter-car rigs include the inability to introduce dynamics such as lateral forces and weight transfer of the vehicle in events such as cornering, braking, acceleration, and aerodynamic loading. It was asserted that in a cornering event the outside front wheel may be carrying as much as the entire load of the front half of the vehicle while the inner wheel may be carrying a negligible load [12]. Thus, the vertical loads that a particular vehicle must carry will change dramatically during such events. Also, during such cornering events the vehicle may roll causing geometry orientation of the suspension to change. This will, in turn, alter the kinematics and dynamics of the suspension. Rigs that do not incorporate actual or equivalent suspension geometry do not accommodate such dynamics at all. Some rigs, primarily those used in durability testing, may incorporate some of the actual vehicle components. This often limits the use of the testing to component tests only. One such platform provided by Labtech International is shown in Figure 2-3. In this example the tire is supported by a rotating drum with a bump excitation. The suspension is simplified such that the engineer can concentrate on the analysis of the spring and damper performance.

Also, the rotating drum only allows for simple cyclic road inputs to the suspension. Thus a simulated non-periodic road profile excitation on this rig would be impossible.

Permission could not be obtained as the company would not respond to request. Virginia Tech and the authors could not reproduce without permissions.
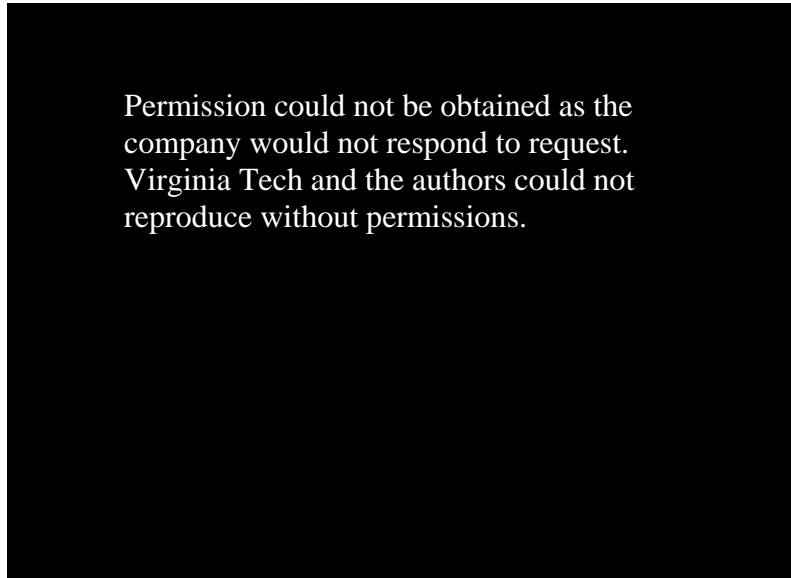
Figure 2-3 Quarter-car Rig for Component Testing ([13])

A couple of commercially developed rigs that offer some of same capability were found in the literature search. MTS Systems Corporation, in cooperation with dSPACE, has developed a proof-of-concept mechatronics development and validation (MDV) bench. The proof-of-concept combines MTS Modeled and Mechanical Simulation Technology with dSPACE hardware-in-the-loop simulation to permit real-time, integrated physical and electronic development validation platform. The mechatronics test system provides an accurate, repeatable and highly efficient means of performing software and hardware development and validation tasks currently executed with prototypes on the proving grounds. This quarter-car rig has much of the flexibility seen in the rig developed for this study. A picture of this rig is found in Figure 2-4. This quarter-car rig has a sprung mass loader based on a force feedback servo-hydraulic actuation system. It also incorporates an actual vehicle suspension. This test bed is only a prototype and is not commercially available.
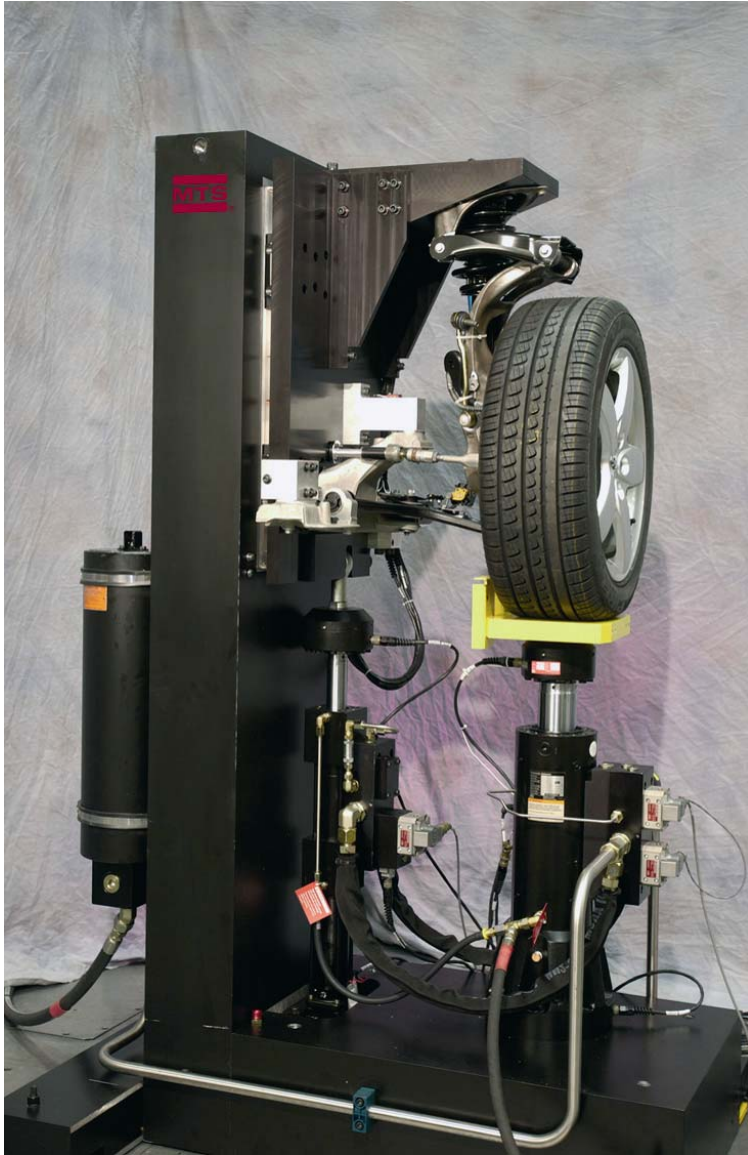
Figure 2-4 MTS Mechatronics Development and Validation Bench (reproduced with permission)

A short video of another quarter-car rig developed by ServoTest was found. This example is captured in Figure 2-5. Here a MacPhereson strut type suspension from a World Rally Championship Car is being tested. The details of this rig and origin of the video were not disclosed however.

Figure 2-5 ServoTest Quarter-car Rig (reproduced with permission)

2.1.3. Functional Requirements

After reviewing the operational functions offered by the current state-of-the-art, the following requirements were proposed for a new quarter-car test rig:

- Design for a variety of different, actual suspension hardware to be mounted and tested thus including the kinematics and dynamics of the vehicle's suspension geometry
- Design for a large range of vehicle corner weights (racecar to HMMWV)
- Design for sprung mass external forces such as aerodynamic loading and/or weight transfer
- Design in flexibility to add future functionality such as vehicle roll or rotating and/or steering the tire

These functional requirements are made such that a new state-of-the-art test rig would be as flexible as possible, allow for more accurate and realistic representation of the test vehicle, and achieve these goals as inexpensively as possible.

2.2. Control of Rig Response

A literature review was performed on the control of vehicle shaker rigs. All servo-hydraulic actuation systems have displacement feedback control loops at the innermost level. These loops accept a displacement input reference command which the servo-hydraulic system tracks within the designed PID inner loop bandwidth. Here, interest is in placing an additional feedback control loop around this inner loop. Particularly, the application of adaptive inverse control was searched. The usual function of a control system on such test rigs is to identify a drive file such that the response of the specimen recorded during a road test is replicated on the rig. Several methods for performing this task on quarter-car and complex shaker rigs exist, however none utilize the particular algorithm implemented in this study. In this study, a proof of concept is being evaluated to find another viable method for controlling these shakers. Once proven feasible, the algorithm used in this study could be applied to a direct comparative test to rate the performance against the existing art for 7- or 8-post test rigs.

2.2.1. Industry Standard

The current industry standard control method is a batch processing algorithm [7, 14, 15]. The various software packages using this method do so in distinct iterations as follows. To start, data is collected during the road test of a vehicle. Next, the vehicle or equivalent specimen is installed on the shaker rig. The specimen is then excited by independent shaped white noise signals running to each actuator simultaneously. The response from vehicle-mounted sensors is collected and a linear model is estimated to match the multi-input/multi-output relationship of the inputs to the response. This model is usually based on a frequency response function. This model must be of high quality because it is used in each subsequent step.

An inverse MIMO transfer function of this model is then calculated and convolved with the data recorded from the road test to create a road or drive file. The resulting drive file is a collection of signals, one for each actuator, that are the same length as the original recorded data. This is the first iteration. The resulting drive file is then scaled to protect the equipment and vehicle and then played through the 7-post

actuators. The response of the vehicle is recorded once again and compared to the desired response. The error between these is then convolved with the inverse transfer function. The resulting drive file correction is again scaled and added to the first iterated drive file. This constitutes the second iteration. This process is repeated until the error between the response of the specimen and the response during the road test is acceptable.

It can be deduced, that one drawback of this current method is that the iteration is an off-line, 'batch' process that can take a large amount of time. It is likely that this large convergence time is measured in hours, as it is indicated that many days can be spent testing on a 7-post during one session [5, 6, 7, 8, 9, 15].

2.2.2. On-line Adaptive Algorithm

Another algorithm is a form of on-line adaptive filter [16]. This algorithm appears in a patent along with another algorithm listed as prior art. The block diagram for the prior art is show in Figure 2-6.
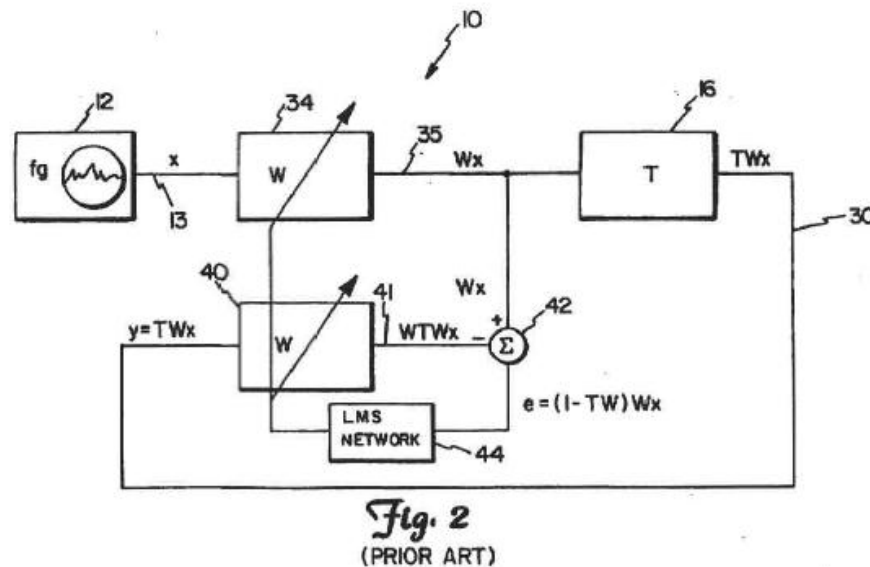
Figure 2-6 Prior Art Algorithm from U.S. Patent No. 5,394,071

Though this algorithm is much closer to the one used in this study it still has some very distinct differences. In this figure, the plant, designated by 'T', is the unknown specimen whose response is being controlled. The desired signal, 'x', is first filtered by

an inverse model, 'W', of the unknown plant. This model is the adaptive filter that is continuously changing as dictated by the LMS gradient descent algorithm. In this case, the desired signal is sent through the inverse adaptive model. The output of this model, 'Wx', is then fed to the unknown plant, in this case a vehicle suspension. The output of the unknown plant, 'y', which equals 'TWx', is then fed into a copy of the adaptive filter. The outputs of each adaptive filter are then compared to create an error. The error is:

$$e = (1 - \mathbf{TW})\mathbf{W}x \qquad (1)$$

This error is then fed into the LMS algorithm which then updates both adaptive filters identically. This method is different from that presented in this study for several reasons. First, the data used to update the filters via LMS is different. Second the algorithm here uses two identical adaptive filters where as that presented in this research only uses one adaptive control filter and one identified model. Finally, the error calculated in this control scheme is compared from the outputs of two adaptive filters rather than comparing the response of the unknown plant to the desired response directly.

The second algorithm, presented in this patent, is shown in Figure 2-7. In this algorithm a different approach is taken. A derived or corrected drive file indicated by '102' is fed directly to the unknown plant. The output of this plant is then fed to an inverse plant identifier, '36', which contains the LMS algorithm. At the same time, this corrected drive file is also fed into this same plant identifier. Meanwhile the output of the unknown plant shown as '52' is then compared to the actual desired response, '12', of the specimen to calculate an error denoted by '106'. This error is fed into an adaptive filter, shown by '34', whose coefficients are adapted by the inverse plant identifier. The output of this filter is then added to the corrected drive file to create a new corrected drive file at '112'. This iteration of the drive file is then collected and stored at '54'. Once the entire corrected drive file is played through once it is then overwritten by the newly stored drive file completing an iteration of the control loop. This process is then repeated until the corrected drive file converges to create a response of the plant that closely matches the desired response.

15
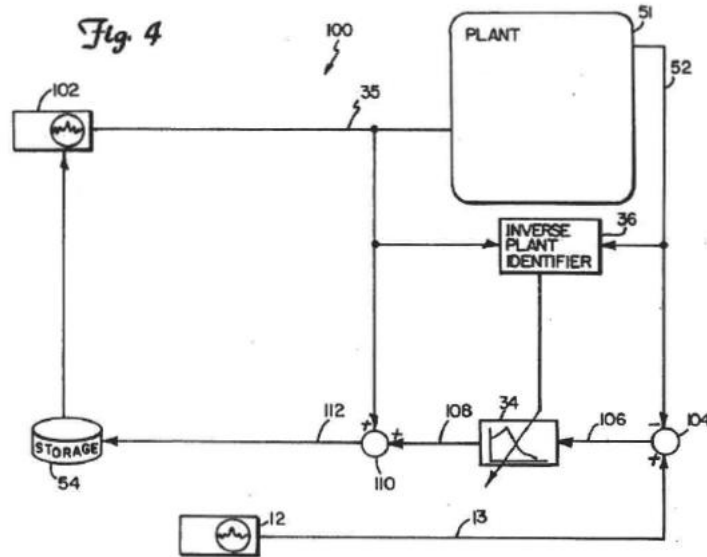
Figure 2-7 Algorithm Presented in U.S. Patent 5,394,071

This method is also very different from the algorithm applied in this research. One primary difference is that even though the process is continuously online, the changes to the drive signal occur in 'batches' where as the drive signal in this study changes each sample. Also, the error computed in this algorithm is fed directly in to the adaptive filter rather than to the algorithm that adapts the filter. Finally, another difference that is apparent is that some base-line drive file must first be computed off-line to begin the process.

2.2.3. Summary of Literature

To summarize the literature, there are many methods for computing a drive file that will cause the test specimen to replicate signals. Acquiring the specific technical details of how these algorithms work is very difficult. The methods appear to have many small discrepancies among one another. It is quite clear that the methods presented in this literature review are different than that which is used in this study. These distinctions will be become more apparent as the details of the algorithm presented in this study are discussed.

3. HARDWARE DEVEOLPMENT

Chapter 3 discusses the engineering approach to the design and construction of the quarter-car test rig used in this research. It begins with a general discussion of the quarter-car rig's design. Following, are detailed discussions of how each major component was developed and/or specified. Included is discussion of the implementation of the first suspension tested with the new rig. The chapter closes by summarizing the functionality of the new quarter-car rig and presents some future developments being planned.

3.1. General Description

The quarter-car test rig is designed to represent one corner of a test vehicle. The design of the rig is such that the rig will go through several configuration phases to achieve the full functionality desired. For this study the rig is in Phase 1. The number of phases is not actually a set number, but progress as time and funding dictate. There will be a brief general discussion of these configuration developments at the end of the chapter. Figure 3-1 is a schematic of this new design.
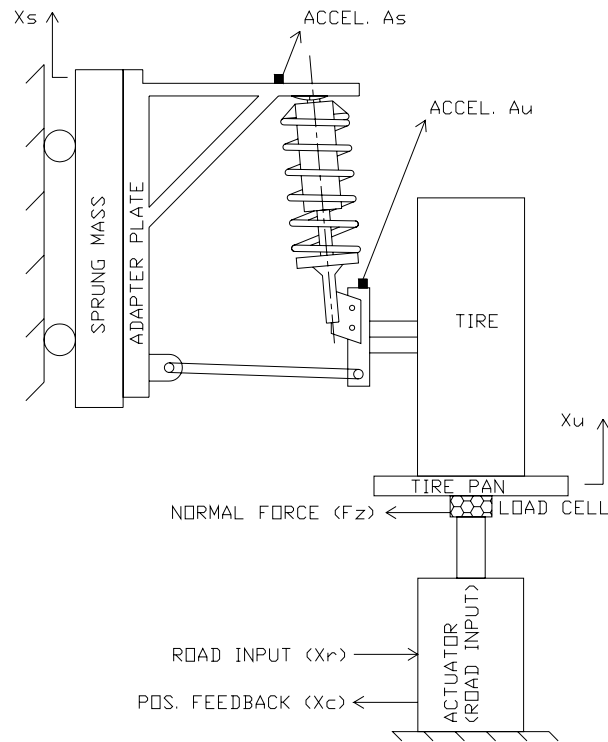
Figure 3-1 Schematic of Quarter-car Test Rig

This representation shows the primary components of the rig. Those components are the sprung mass, adapter plate, vehicle suspension, tire, tire pan, actuator and ground or load frame. The sprung mass and suspension adapter plate are constrained with bearings to move in the vertical direction. The actual suspension of the vehicle is attached to the sprung mass via the adapter plate and fixtures. The actuator is fixed to the ground and supports this represented vehicle via the tire resting on the tire pan. A displacement command signal is input to the servo-hydraulic actuator which then excites the system through the tire contact so that the suspension response matches the measured signal. Since the sprung mass is free to move in the vertical direction only, the vertical dynamic response of the specimen may be tested.

To ensure that the requirements of the new rig were fulfilled, special attention was paid to the specifications and design of the key components. The result of this design was a state-of-the-art quarter-car rig that is able to fulfill many of the functional requirements listed, and yet flexible enough to allow for implementation of hardware that

18

will allow for fulfillment of the remaining and added requirements. Figure 3-2 is picture of the solid model of the new design. This model was created in Solidworks 2006.
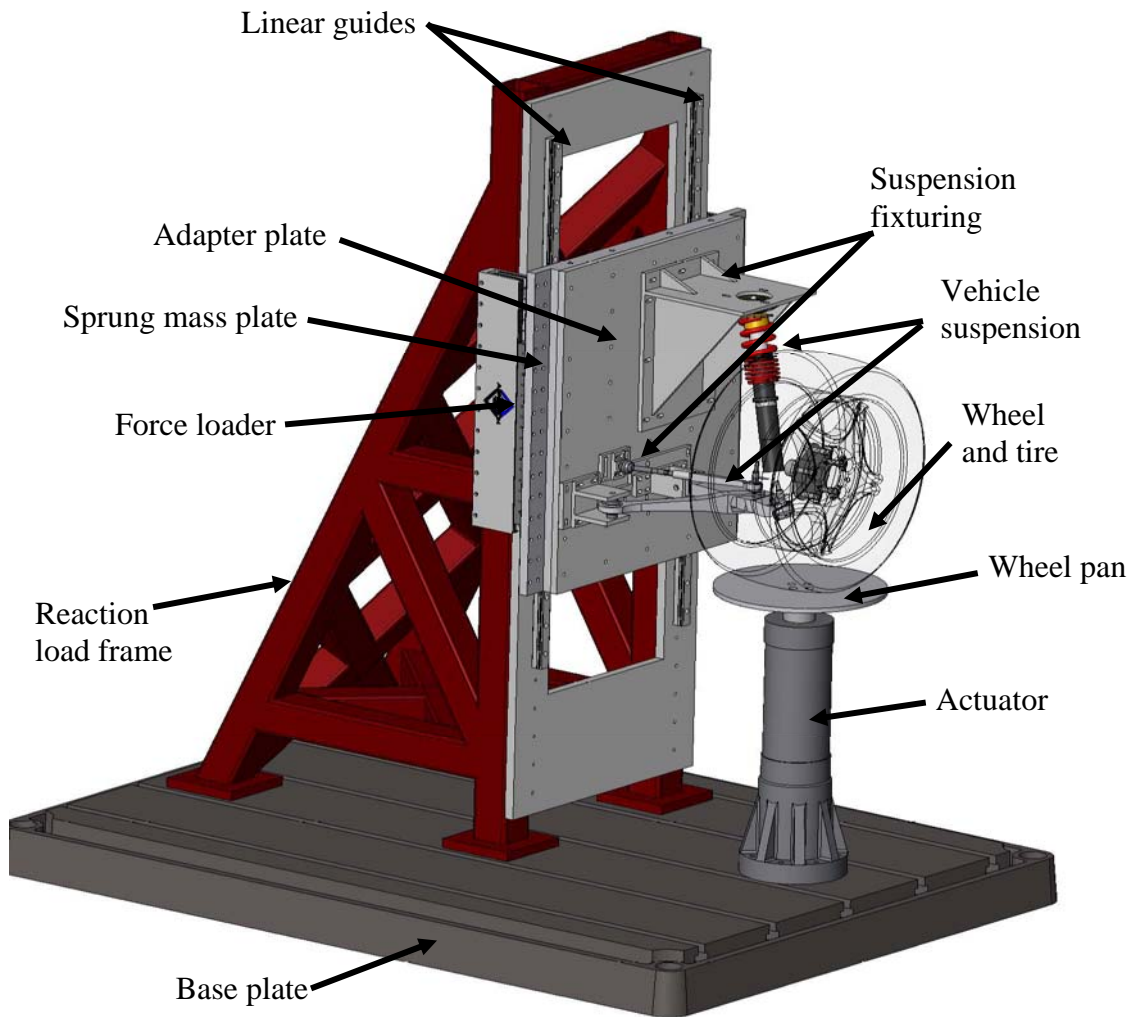
Figure 3-2 Solid Model of Quarter-car Rig in Design Phase

From the figure the base plate and reaction frame bolted together are the support structure for the entire rig. The linear guides are actually six parts. There are two guide rails each with two bearing carriers that move in one degree of freedom. The linear guide rails are bolted to the load frame vertically. The sprung mass is constrained to the rail

using the four carriers. There are two force loaders, which are known as aeroloaders in the 7-post testing vernacular. There is one loader on each side of the sprung mass. One loader can be seen in the figure. The force loaders, which will be implemented in a later phase of construction, are two piece electro-magnetic linear motors. One piece is bolted to the load frame and one is bolted the sprung mass. The suspension is bolted to the sprung mass via the adapter plate and suspension fixturing. With the tire attached, the representative corner of a car is supported by the actuator via the wheel pan. The sections that follow detail the design of many of these major components of the rig.

3.2. Base Plate

In any laboratory test an engineer wishes to have as much control of the test specimen's environment as possible. Particularly, in vibrations and control tests any fixturing used must be designed such that the dynamics of the fixture do not affect the dynamics of the test specimen. Some studies used a welded steel structure and base plate which weighed in excess of 12 tons [17, 18]. To this end the base plate needed to be heavy and rigid. The base plate was specified and sourced by BayCast Technologies [19].

3.2.1. Plate Design

The specified base plate is an 84 in x 60 in x 7 in purpose build, cast iron plate. The plate is hollow with a stiffening rib structure on the under-side. It weighs approximately 2500 lbs. The base plate is anchored floor via four 1 in anchors, one for each corner of the base plate. The anchoring system works as follows. The 1 in all-thread anchors were secured to the 8 in thick concrete floor with epoxy. With the epoxy cured, the base plate was lowered down over the anchors. Fixed nuts in the base plate thread the plate down these anchors such that it was actually floating over the floor. The full weight of the plate was supported temporarily by the anchors. This allowed the plate to be leveled to within 0.1° via BayCast's proprietary Level-Tite system. Once the plate was leveled, jam nuts fixed the position of the plate. With the plate fixed in place it was then filled with a non-shrinking grout. The grout filled the underside completely and

20

allows the plate to distribute loads over its entire footprint on to the floor. The fully anchored plate is actually an extension of the concrete floor. The working face of the plate is machined flat to within 0.005 in and has cast in 7/8 in T-slots.

3.2.2. Functionality

The anchoring system and structural design of the plate makes it extremely heavy and rigid. This gives the plate much higher natural modes than the test specimen which helps to reduce error in tests. The extremely flat surface ensures that the forces being introduced into the suspension are well defined. Misalignment of the actuator to the motion of the sprung mass would introduce undesired forces such as lateral force when only vertical response is of interest. These lateral forces can cause extra non-linearities between the input-to-output relationships of the system.

The load frame and actuator are secured to the base plate via the T-slots in Figure 3-3 and the appropriate hardware shown in Figure 3-4. The hardware uses a T-nut that is designed to fit in the T-slots in the base plate. The stud threads into the T-nut. The block has stepped teeth and the clamp strap has the same teeth as well. The clamp strap has a slot in the middle for the stud to go through. One end of the strap is resting on the object to be held in place and the other is resting on the step block with the teeth of each being meshed. The top nut is then turned down to apply a clamping force on the clamped object.

Figure 3-3 T-slots in Base Plate
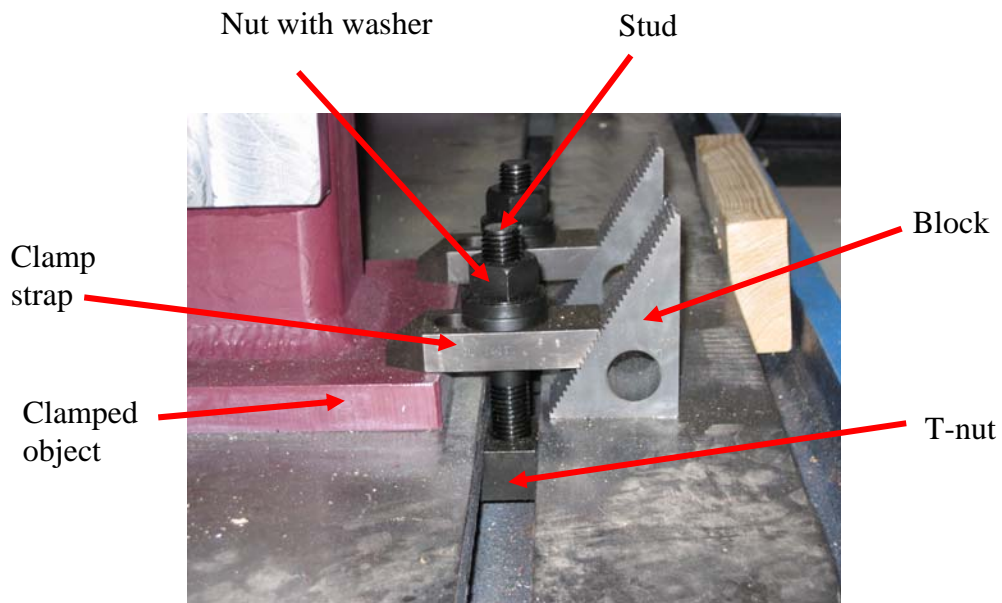
Figure 3-4 Clamping Hardware

The combination of the machined surface and T-slot system allows for a vast amount of adjustability and also for precise positioning of the actuator under the tire. This adjustability allows the rig to accommodate a wide variety of suspension component sizes like long control arms of up to 22 inches, as well as wide racing tires. In summary

the base plate's functionality lies in it rigid design to reduce unwanted noise in test signals and its ability to accommodate a vast variety of suspension configurations.

3.3. Reaction Load Frame

For the same reasons as the base plate design the reaction load frame also needed to be extremely strong. The reaction frame is the tall triangulated steel structure that is clamped to the base plate. Together, the reaction frame and base plate work to create an excellent rigid ground for the vibration experiment to work on. The reaction frame's additional functionality lies in its ability to be placed at various positions on the base plate allowing for a large range of suspension designs and sizes.

3.3.1. Frame Design

The first step in the frame design was deciding on the material and construction of the frame. In several cases an extruded aluminum frame was assembled using bolts, nuts and brackets to connect the pieces. First-hand accounts of this type of structure indicate that this type of joining can lead to some unwanted flexibility in the structure. The goal is to reduce excitation of the sensors from outside sources including the rig's structure. Thus, a welded low-carbon steel space frame design was chosen.

The overall dimensions of the frame are 84 in tall with a 28 in x 42 in foot print. The rig is made from 4 in x 4 in steel square tubing with a 3/8 in wall thickness. The tubes create a triangulated frame seen in Figure 3-2. A 1 in thick steel plate is welded to the front of the space frame as a location for mounting the bearing rails. Completed, the frame weighs approximately 1500 lbs.

3.3.2. Frame Analysis

To optimize the rigidity of the frame, a finite-element model of the rig was created in ANSYS using beam elements. This model was simplified in that the steel plate on the front was not modeled due to computational restrictions with this particular license of the software. This was not a problem as it merely added a safety factor to the analysis.

If the welded frame was strong enough and had a high enough first mode, then welding a steel plate to the front would only strengthen the design. The plate has a window cut into it as well so the center section is lightened to increase the natural modes of this constrained membrane.

Figure 3-5 is an image from ANSYS showing the finite-element beam model created. Two analyses were performed to aid in the design of the frame. First, deflection of the frame in the out-of-plane direction was simulated with a lateral suspension load. Though this rig is not designed to test lateral forces at this point, the suspension components will resolve normal forces at the contact patch into some combination of lateral forces on the rig. Thus, the lateral force used for the deflection analysis was well beyond that which the suspension could create in pure jounce. Minimizing the deflection of the frame was necessary to reduce friction created by the bearing rails which mount directly to the frame. Just to ensure the strength of the load frame allowed for testing flexibility, a steady-state lateral force of a 6000 lb vehicle in an extreme cornering event was simulated to be acting on the load frame. The result of this simulation showed that the deflection of the frame was less than 0.010 in, which was significantly below the allowable misalignment of 375μm/500mm (rail length), specified by the bearing manufacturer. This ensured proper low friction bearing operation.
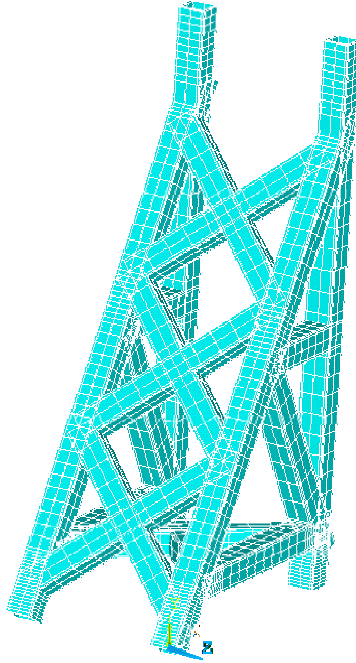
Figure 3-5 Finite-element Beam Model of Reaction Frame

The modal natural frequencies of the structure were investigated with finite-element analysis as well. Using the same beam model as before, the feet of the structure were constrained in all directions and a free vibration modal analysis was performed in ANSYS. The results of this analysis showed that the first modal frequency of concern was nearly 200 Hz. This is comparable to the results of another study which yielded a 500 Hz first resonant frequency for the test specimen's superstructure [20]. These resonant frequencies are much higher than those of interest in vehicle dynamics. In many tests the frequency band of interest is only 1 to 25 Hz [2].

3.3.3. Functionality

The results of the analyses of the frame indicate that it is very rigid compared the rest of the experiment. This minimizes sensor noise from excitation of the load frame. To make the frame fully functional, an 8 in x 8 in x 1 in steel pad was welded to the bottom of each corner. These pads are used to clamp the fixture to the base plate. Once fully welded, the entire load frame was machined as a unit. The feet were all milled to within 0.003 in. flatness of one another. This ensures that the rig is very stable and does

not flex when clamped to the base plate. Next, the front 1 in plate was milled to within 0.0003 in flatness. This was within the allowable misalignment of the bearing rails per the manufacturer. Also, the front face was milled to within 0.005 in perpendicularity of the bottom feet. This ensured that the forces entering the suspension were well defined by reducing the misalignment of the actuator with the motion of the sprung mass.

The design of the load frame as a single unit allows it to be placed all over the base plate for various configurations. The large cutout in the front face allows for protrusions from the rear of the sprung mass such as added mass plates or motors for an active geometry type of suspension.

3.4. Linear Guides

The linear guides are LH35 series linear bearings manufactured by NSK. This particular bearing design is used due to its proven performance in machine tool and automation industries. A larger, custom variation of this particular design can be found in many Toyoda vertical high-speed milling machines [21]. These bearings are a high accuracy, high load and low friction design.

3.4.1. Sizing

The bearings were sized with assistance from NSK Corporation engineers and product documentation [22]. NSK has a standard sizing practice that is based on life, load, and load offset among others. In this study the driving load was calculated from a simple two-mass quarter-car model detailed later. The corner weight chosen was 900 lbs. The offset distance, or the distance at which the actuator applies a force to the tire relative to the bearing, was chosen to be 20 in. These somewhat extreme values were chosen to ensure a safety factor in the design and to ensure a high load capacity for future use of the rig.

3.4.2. Design and Functionality

The design of LH35 bearings is a two groove, gothic-arch guide rail made from hardened and ground steel. The carrier, which runs along the rail, has a recirculating ball bearing design. Figure 3-6 shows a cross section of the rail and cutout of the carrier. The ball bearings are of angular contact design, which yields the high load capacity and low friction design. These bearings are an interchangeable design between rail and carrier making them a cost effective solution.
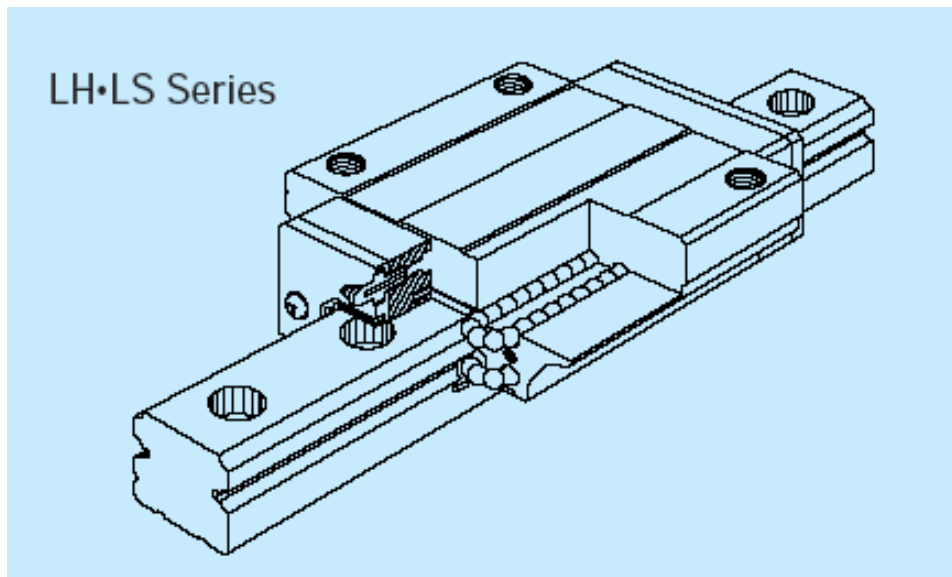
Figure 3-6 Section of LH Series Bearing

The guide rails are 1500 mm long. This was chosen to allow the sprung mass to have a high range of motion allowing for various applications. For this application the specification for the rail to carrier clearance was given such that there is a slight amount of preload in the bearings. This was done to eliminate rattle space between the carrier and rail which could cause parasitic vibrations during a test. The preload can not be too great however, as friction is a function of the preload. Based on NSK documentation, the dynamic friction force for the four carriers used in this design should total less than 10 lbf in the velocity range of interest. The design, of these bearings, also gives them a low variation between static and dynamic friction. This is useful because it makes controlling a force loader a much easier task if force feedback were used.

3.5. Moving Mass

The moving mass is the part of the rig that replicates the vehicle's sprung mass, which includes the body and chassis. The most difficult functional requirement of this new rig was addressed with a sprung mass modular design. One design goal was to be able to attach various suspension designs to the same test rig without major modifications. To this end the moving mass was made into a modular two piece design. A sprung mass plate was designed to be the permanently installed moving plate. The sprung mass is bolted to the linear bearing carriers which are constrained by the rails. An adapter plate along with the appropriate fixtures is designed to be the interface between the sprung mass and the vehicle suspension.

3.5.1. Sprung Mass Plate

The primary goals of the sprung mass plate were to be lightweight, rigid and functionally flexible. Aluminum 6061-T6 was chosen for the material. The plate is 2 in thick and has dimensions of 33 in x 36 in. The long dimension is in the direction of travel. The area defined by these dimensions allow for a large working area for various suspension designs. To reduce weight the bearing side of the plate was pocketed heavily, leaving a ribbed support structure. The plate has an array of threaded thru-holes which are used to fasten the adapter plate to the sprung mass. Not all holes are used at the same time as some suspensions may block certain holes depending on their locations. The finished plate weighs approximately 150 lb. Though, this is extremely light for the corner weight of a vehicle, the mass increases with the addition of the adapter plate, fixturing, and fasteners. Extra weights have been made that can bolt to the back of the sprung mass to adjust the mass according to the corner weight of the vehicle in increments of 30 lb. As much as 270 lb can be added to the back side alone.

Again, ANSYS was utilized as a design tool for the sprung mass. The plate was loaded as a quasi-simply-supported beam to find the maximum deflection and stress in the plate. For this study it was assumed that the rig could be used for lateral loading

some point in the future. Therefore, a large vehicle weighing 6000 lbs was assumed to be in a tight steady state turn. The assumption is that a single wheel might be supporting as much as half the weight of the vehicle. Thus, a lateral force of 3000 lbs was chose for the deflection load case. This load was applied to a small pad in the middle of the plate and the plate was constrained in all degrees of freedom as the bearing carriers.

Figure 3-7 shows the out of plane deflection and stress distribution in the plate. The simulation revealed that the plate would have no more than 0.008 in of deflection out of plane. This corresponded to a von-Mises stress of 5000 psi [23]. This was substantially less than the yield strength of 40000 psi. Based on maximum-distortion-energy theory yielding will not occur [24]. The design of the plate is more than sufficient in strength and rigidity. The rigidity helps extend the life of the bearings and reduces friction as well.

(a)   **(b)**
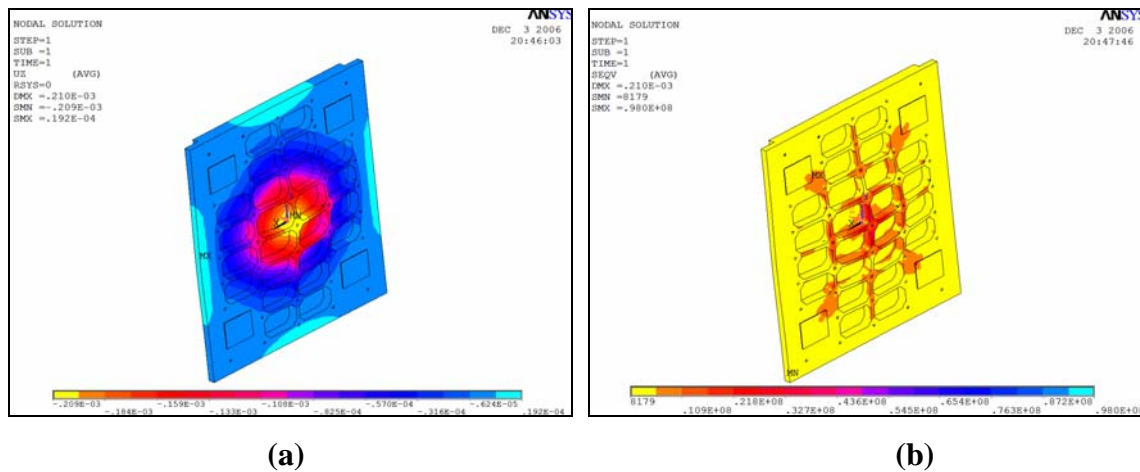
Figure 3-7 ANSYS Sprung Mass Results (a) Out-of-plane Deflection (b) von-Mises Stress

3.5.2. Adapter Plate and Fixturing

The latter component of the modular moving mass design is the adapter plate. This plate is the interface between the suspension mounts and the sprung mass plate. The design of this plate is specific to the suspension design. The idea is to have one adapter plate per suspension tested on the rig. For this study the adapter plate is made from ½ in thin cold-rolled steel plate. This plate is designed in conjunction with the suspension

mounts to ensure that clearance and fastener issues do not arise. The addition of this plate fastened to the sprung mass helps to further stiffen the entire moving mass. More detail of the design of this plate is included later in the discussion of the first suspension application.

3.6.   First Application

The first suspension implemented is the left-front suspension from a 2004 Porsche 996 Grand American Cup GS Class racecar, similar to that seen in Figure 3-8. The car has a minimum weight of 3000 lb [25]. Using scale pads the left-front corner weight was found to be 630 lb. The suspension and mass was setup on the quarter-car rig to have the same corner weight, springs, damper and suspension geometry as that of the actual racecar [26].

Figure 3-8 Porsche 996 Grand Am Cup GS Racecar

3.6.1. 996 Suspension

The Porsche 996 suspension, shown in Figure 3-9 is a typical variation of the MacPherson strut type independent suspension. The primary components are the upright, two piece lower control arm and strut comprised of a coil-over spring and damper. The suspension requires two mounts for the lower control arm. A bracket is necessary for

connecting the strut mount and a fully adjustable inner tie-rod mount is needed to constrain the steering motion of the wheel.
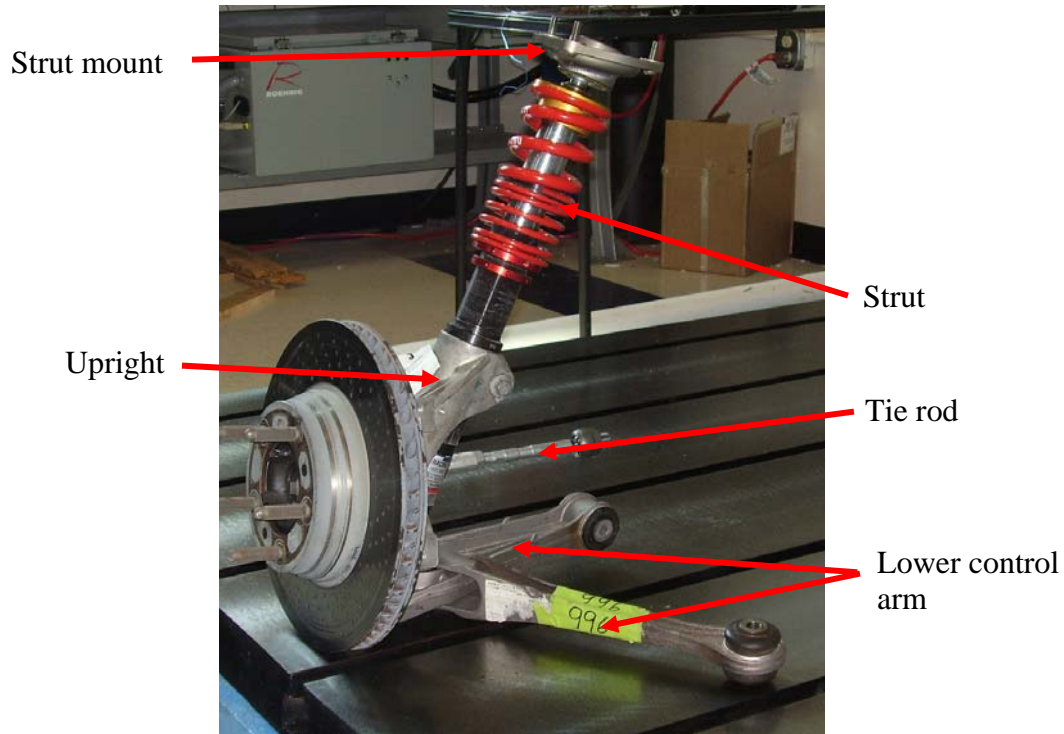
Figure 3-9 LF Porsche 996 MacPherson Strut Type Suspension

First, a solid model, of each suspension component, was created such that the mounting points were located as accurately as possible. Detailed blue prints for these components and chassis mounting points were not found in the public domain making this process somewhat difficult. The lower control arm, tie-rod, strut and strut mount were measured directly with a scale and calipers due to their relatively simple geometry. The upright was taken to a local tool and die shop and all of the mounting points were measured with a coordinate measuring machine (CMM). Data from the CMM was imported directly into Solidworks to create the mating surfaces of the upright. The mounting locations modeled on the upright were the strut bore, tie-rod taper bore, lower control arm taper bore and center of the hub to locate the wheel.

3.6.2. Fixture Design

The required mounting brackets and fixtures need to offer at least the same level of adjustability as found on the car. Provisions for adjusting camber, caster and toe, based on the required race setup, were considered during the design of these components. Figure 3-10 shows the final design of all fixturing for the 996 suspension.
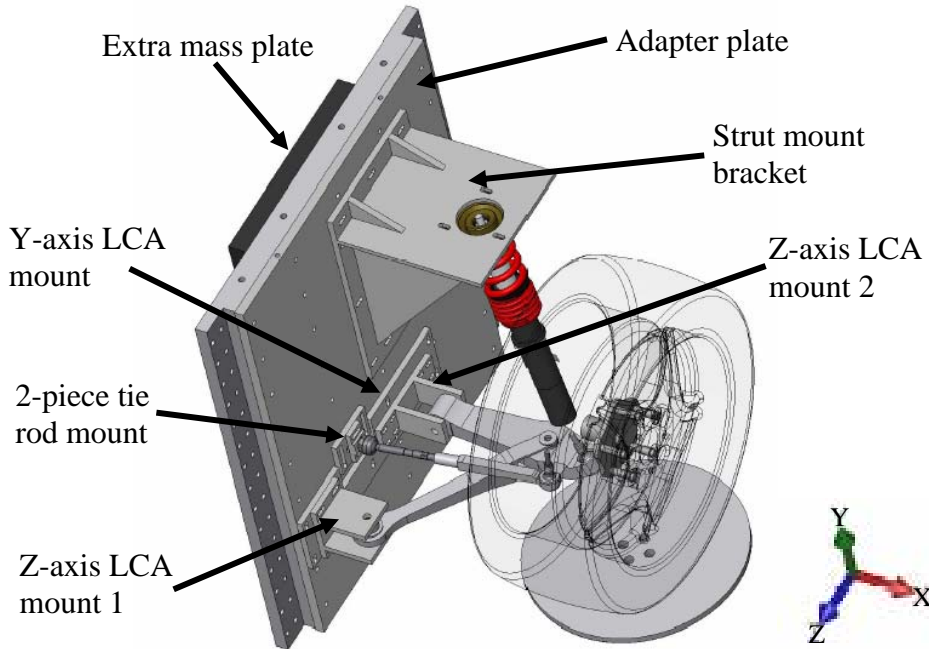
Figure 3-10 Detailed View of Modeled Suspension and Mounts

The figure shows the lower control arm (LCA) mount is a three piece design. Because the LCA is actually a two piece design, there needed to be separate brackets for each part of the arm to allow for adjustment in the Z-direction, when referring to the axis coordinate system in the figure. The two ends of the LCA bolt to mounts 1 and 2 respectively. The back faces, of these mounts, have slotted holes for adjustment in the Z-direction independent of one another. The first two brackets bolt to a third plate that has holes slotted for adjustment in the Y-direction. This plate mounts to the adapter plate. A detail of the adjustment motions of this three-piece design is shown in Figure 3-11.
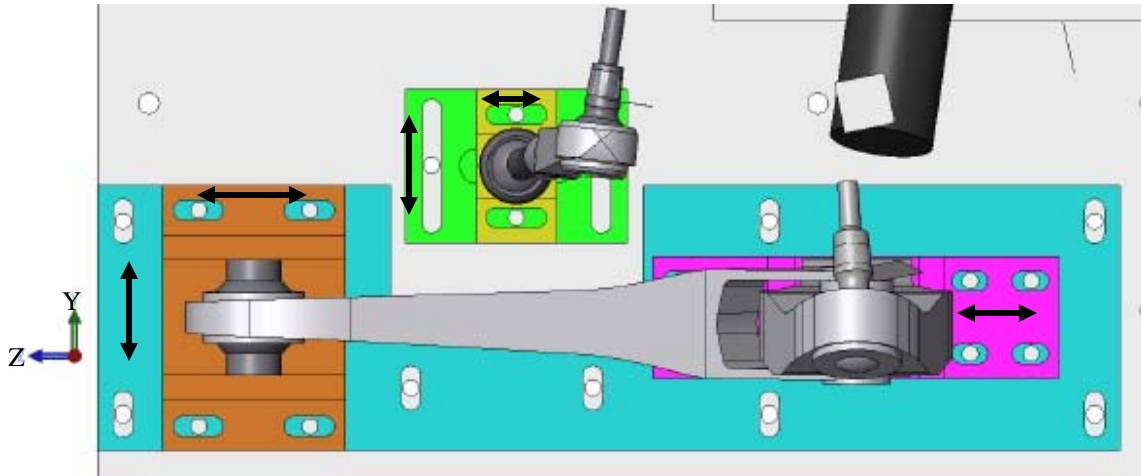
Figure 3-11 Detailed View of LCA and Tie-rod Brackets Showing Adjustment Directions

A similar design for the inner tie-rod mount allows for two axes of adjustment of the inner tie-rod pivot. The tie rod threads into the Z-axis mount, which bolts to the Y-axis mount, which bolts to the adapter plate. This two-piece bracket design is shown in the previous figure as well. The strut mount bracket seen in Figure 3-10 offers two axes of adjustment. This bracket mounts to the adapter plate and has slotted holes allowing for Z-axis adjustment. The strut mount seen in Figure 3-9 has three studs that go through slotted holes in the bracket. These holes are slotted in the X-axis direction. The strut mount connects the strut to the strut bracket which then fixes the assembly to the adapter plate.

With all of the brackets designed and modeled, everything was assembled in Solidworks along with the suspension. Once assembled, the components were adjusted until the actual 996 suspension geometry was replicated. At this point the appropriate holes were positioned in the adapter plate such that the plate could bolt to all of the brackets and the sprung mass plate in order to complete the adaptation.

3.6.3. Adjustments

With the manufactured parts in hand and the quarter-car rig assembled, some adjustments of the suspension were necessary. The adjustments are meant to match the

suspension set up of the represented 996. Synergy Racing was the source for the front end alignment figures. To protect their setup information, ranges were provided rather than actual specific numbers. Table 1 is a list of the geometry setup numbers.

Table 1 Setup range for 996

Setting	Range	Setup on Rig
Camber	0° to -4°	-0.5°
Caster	5° to 10°	8°
Toe	1/8" out to 1/8" in	0°
Corner weight	560 lb to 640 lb	625 lb
LCA Ride-height angle	1° to -1°	-0.5°

Based on the figures in Table 1 the suspension is set up to meet to attempt to meet the median of these ranges. The camber is adjusted by sliding the location of the strut mount on the strut mount bracket in the X-direction and by adjusting the Y-position of the LCA inner mounts. Care is take here as changing the LCA mount position also has some affect on the angle of the LCA at ride height. The caster was set by adjusting the Z-positions of the LCA mounts and also of the strut mount bracket. Finally, the toe is set by adjusting the position of the inner tie-rod mount and the length of the tie-rod.

3.7.   Actuation

Road input to the suspension is supplied via a tire coupled servo-hydraulic system manufactured by MTS Systems Corporation. The servo hydraulic system is comprised of an actuator with position feedback, pump, manifold, servo valves and PID controller. The hydraulic actuation system was specified with help from MTS engineers.   The functional requirements of this system were to have a high load capacity for a wide range of vehicles and a high actuation bandwidth to allow for a wide range of tests.

The actuator is an MTS Model 248.03 linear hydraulic actuator capable of 5.5 Kip force. It has a dynamic stroke of approximately ±3 in. The actuator, shown in Figure 3-12, has hydrostatic bearings to give it high side load capability should some lateral force become generated in the tire contact patch. The piston was custom made to be light weight and strong to provide a bandwidth of over 150 Hz at low amplitude. The actuator

34

has an inline linear-variable displacement transducer (LVDT) and a delta-pressure cell for position or force feedback to the controller respectively. A wheel pan was custom made from 6061 aluminum to keep the working mass low. The first natural frequency of the wheel pan was found to be over 200 Hz using finite-element analysis. This ensures that unwanted excitation does not appear in a test.

Figure 3-12 MTS 248.03 Hydraulic Actuator

The flow of hydraulic oil is regulated by two Moog Model 252.25G-01 4-Port servovalves. These valves seen in Figure 3-12 are two stage valves giving them a bandwidth of nearly 300 Hz. Each valve is rated at 15 gpm flow rate. These valves work together to increase the response of the actuator and yield as much bandwidth as possible.

Hydraulic power is supplied by an MTS Model 505.03 SilentFlo hydraulic power supply. The power unit shown in Figure 3-13 runs on 480 VAC 3-phase power and is water cooled. It operates at 3000 psi and can supply up to 30 gpm.

Figure 3-13 MTS SilentFlo Hydraulic Power Supply

Hydraulic power is regulated to the actuator by an MTS 293.11 Hydraulic Service Manifold. It also supplies the oil pressure for the hydrostatic bearings in the actuator. The manifold seen in Figure 3-14 was specified to have an over-sized accumulator on the pressure side to aid in bandwidth requirements of the system.

Figure 3-14 MTS Hydraulic Service Manifold

Finally, the controller shown in Figure 3-15 provides the real-time closed-loop control for the hydraulic system. The controller is an MTS Model 493.02 FlexTest SE Controller. This stand-alone controller can close the loop with force or position feedback. It accepts one analog input for external command signal generation and it also has an internal function generator. The controller also has provisions for four digital inputs for triggers and interlocks, four digital outputs and three analog outputs for monitoring signals or closing an outer loop.

Figure 3-15 MTS FlexTest SE Controller

3.8. Maintenance

For the purpose of maintaining the quarter-car test rig and supporting apparatus a bulleted list of maintenance items is found in Appendix A.

3.9. Summary and Future Developments

The following section provides a brief summary of the functionality of the newly designed quarter car test rig and how it has fulfilled the functional requirements laid forth. Next, a brief discussion of future developments for the test bed is provided.

3.9.1. Summary of Functional Requirements Fulfilled

Due to the modular design of this new quart-car test rig it is able to accommodate a multitude of different vehicle suspension designs. A photograph of the fully assembled test rig is shown in Figure 3-16 with the Porsche 996 suspension installed. The design of the sprung mass to receive an adapter plate rather than a single specific suspension keeps the rig from being purpose built for one suspension design and corresponding car mass. The configurability allotted by the base plate design and the single-piece design of the reaction load frame allow for a large window of placement of the tire relative to the sprung mass. The moving mass has provisions for adding mass which allows for vehicles of various weights to be replicated. Finally, provisions for additional functionality are designed into the rig making it useful for future studies. A set of specifications for the rig are included in Appendix B.
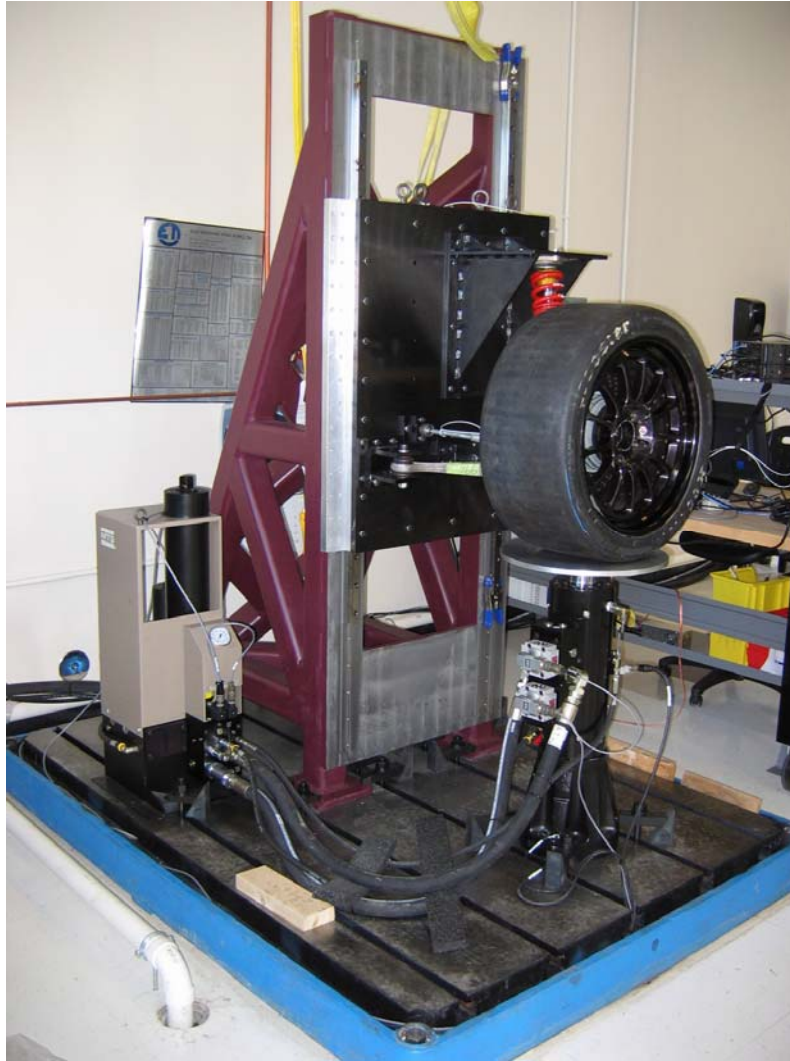
Figure 3-16 Completed Quarter-car Test Rig with Suspension Installed

3.9.2. Future Enhancements

Future additions in functionality will occur in steps. The first of which is to add force loading capability on the sprung mass. The results of this research effort can only be improved by replicating inertial and aerodynamic loading on the sprung mass. The sprung mass has been designed to accommodate Aerotech electro-magnetic motors, one on each side. These will be somewhat innovative as most loaders in 7-post systems are electro-hydraulic with force feedback. The design is such that the magnets of these motors will bolt directly to the load frame. Flanges on the sprung mass have already

been designed such that the armatures will bolt directly to the mass. This combination will apply the desired inertial and aerodynamic load required.

Another future project is to demonstrate the flexibility of the test bed by implementing the suspension of a completely different vehicle type. Plans are in place to replicate the suspension of an AM General HMMWV. This vehicle is substantially different in that it has a corner weight that is nearly 2000 lbs more than the Porsche and experiences much larger amplitude vibrations in the sprung mass. Concurrently, there is also discussion of installing a NASCAR Nextel Cup suspension on the rig.   These racecars have a short-long arm suspension which is a completely different design compared to both the Porsche and HMMWV. At any given time, the quarter-car rig can be reconfigured from testing one type of suspension to another in a matter of a few hours.

Another possible function of the rig is to lock out the motion of the sprung mass and perform compliance or durability tests. Additional functionality is to install a dynamically controllable roll degree of freedom on the sprung mass. This additional degree of freedom would allow the suspension geometry to be changed by simulating a low frequency body roll as would occur in a cornering event. The body roll changes the suspension geometry and replicates a non-linearity that most other single-post rigs do not account for. Many of the findings that may result from these potential studies could be applied to more complex testing rigs.

4. CONTROL THEORY

The control of this test setup is based on two well known principles. These are the adaptive linear combiner (ALC) and the least-mean squares algorithm (LMS). Chapter 4 introduces these two concepts in detail. The chapter then discusses the application of LMS and the ALC to system identification. The chapter ends with a similar discussion of application to adaptive control. This will familiarize the reader with the control concepts used in this study.

4.1. The Adaptive Linear Combiner

The ALC is an adaptive finite-impulse response (FIR) filter that is a fundamental building block in adaptive signal processing. This time-varying filter is shown in Figure 4-1. In this instance the ALC is represented for a single input, $\mathbf{X}_K$.
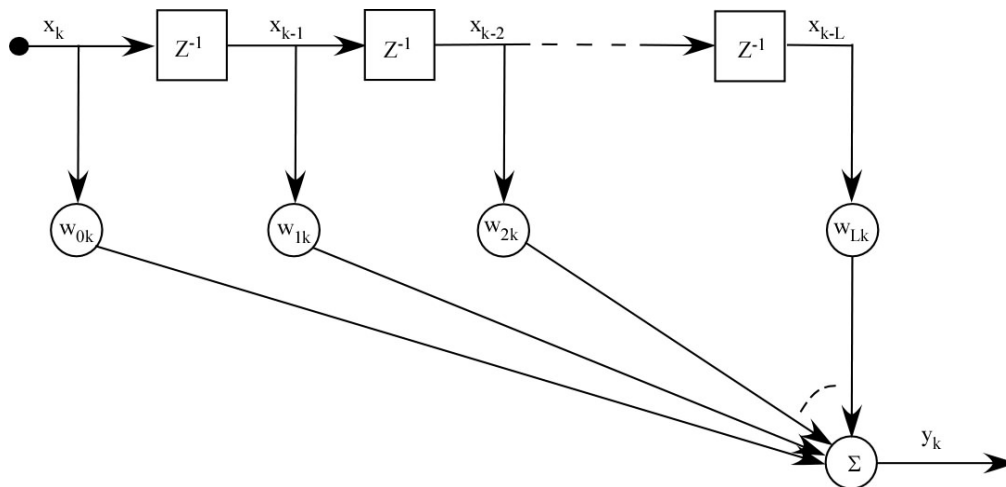
Figure 4-1 The Single Input Adaptive Linear Combiner

For the single input case, the ALC functions as follows. The input, $\mathbf{X}_K$, is sampled by a tapped-delay line, which creates a sequence of delayed values from the same source sampled every T_s seconds. The z^{-1} blocks in the figure represent a unit or sample delay. Thus, the input is successively sampled and multiplied by a gain, w_{nk}, which multiplies the n^{th} input sample at time k. In other words, an input can be thought of as a vector with

the first element being the most recent sample, x_k, followed by the previous sequence of data, $x_{k-1}, x_{k-2} \ldots$, going back in time. The weights are also contained in a vector that when dotted with the input vector yields the output at that particular moment in time. This can be expressed as:

$$y_k = \sum_{l=0}^{L} w_{lk} x_{k-l} \tag{2}$$

This can also be represented in vector form in (3)

$$y_k = \mathbf{W}_k^T \mathbf{X}_k \tag{3}$$

If the weights were constants this would be a linear system. However, the adaptive part of the filter is the time-varying weight vector. The weights are time-varying because they are continuously changed to meet some performance criteria. Normally, the goal of changing the weights is to reduce the error between the output of the FIR filter and some known signal. This desired signal is added to the ALC shown in Figure 4-2.
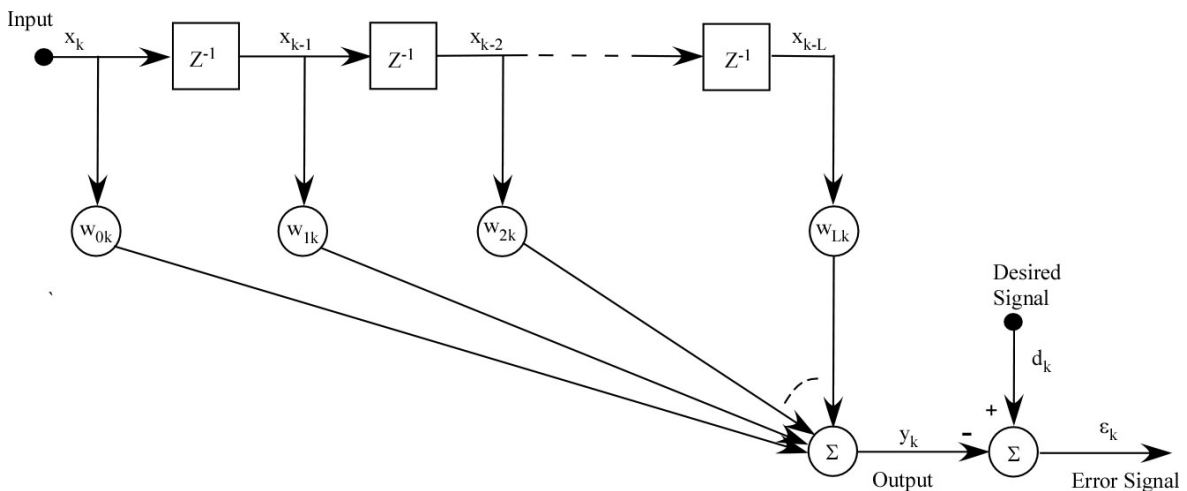
Figure 4-2 The Adaptive Linear Combiner Compared to a Desired Signal

42

The output of the filter, y_k, is normally subtracted from this desired signal, denoted by d_k, to produce an error at time, k. This error is denoted by ε_k. Thus the error signal is given in (4) when a substitution is made from (3):

$$\varepsilon_k = d_k - y_k = d_k - \mathbf{W}_k^T \mathbf{X}_k \qquad (4)$$

At this point, we construct a quadratic cost function which is the instantaneous squared error.   We can study the optimal solutions and the dynamics of weight trajectories by assuming that d_k and $\mathbf{X}_k$ are statistically stationary. The details of this analysis can be found in the literature [27, 28]. Taking the expected value of equation (4), we have:

$$\xi = E[\varepsilon_k^2] = E[d_k^2] + \mathbf{W}^T \mathbf{R} \mathbf{W} - 2\mathbf{B}^T \mathbf{W} \qquad (5)$$

Equation (5), represents the mean-square error function. Here, $\mathbf{R}$ is the input auto-correlation matrix and $\mathbf{B}$ is the cross correlation between the desired response and the input signal. It is clear that this equation is a quadratic function of the weights, $\mathbf{W}$. Thus the mean-square error cost function is 'bowl' shaped. Naturally, this bowl shape has a bottom or minimum value. In order to determine this minimum, the partial derivative (gradient) of the cost function is taken with respect to the weight vectors. The result of differentiation is set equal to zero to find the minimum error value. This minimum is found when the weights are at their optimal values, $\mathbf{W}^*$. In practice, most adaptive algorithms utilize noisy instantaneous gradients to slowly converge on the optimal value as opposed to single-step deadbeat control algorithms. In this study the LMS gradient descent adaptation algorithm is utilized to optimize the weights of the adaptive linear combiner.

4.2. Least-Mean Squares

The LMS algorithm is a steepest-descent method for finding the minimum value of a function. That is, it estimates the gradient of a function and then travels in the opposite direction of the positive slope. The algorithm is extremely useful for applications that use the ALC because computations are relatively simple which make it very useful for on-line applications. Many identification or optimization processes that are more complicated require some form of off-line adaptation.

The LMS algorithm begins with the definition of the error signal given by (4). Typically, equation (5) is used to define the error function. Instead, the squared error, ε_k^2, is used as the estimation of the error function rather than ξ_k. Taking the square of (4) yields the following:

$$\varepsilon_k^2 = d_k^2 + \mathbf{X}_k^T \mathbf{W}_k \mathbf{W}_k^T \mathbf{X}_k - 2d_k \mathbf{W}_k^T \mathbf{X}_k \tag{6}$$

Next, the gradient is estimated by taking the partial derivative of the squared error with respect to the weight vector. In vector form this appears as:

$$\nabla = \frac{\partial \varepsilon_k^2}{\partial \mathbf{W}_k} = \begin{bmatrix} \dfrac{\partial \varepsilon_k^2}{\partial w_{1k}} \\ \vdots \\ \dfrac{\partial \varepsilon_k^2}{\partial w_{Lk}} \end{bmatrix} = 2\varepsilon_k \begin{bmatrix} \dfrac{\partial \varepsilon_k}{\partial w_{1k}} \\ \vdots \\ \dfrac{\partial \varepsilon_k}{\partial w_{Lk}} \end{bmatrix} = -2\varepsilon_k \mathbf{X}_k \tag{7}$$

The second partial derivative shown comes from the derivative of the error function in (4). The negative value indicates that the algorithm is descending the performance surface. The resulting gradient in (7) is attenuated with a small gain constant or step size, μ. Finally, the attenuated gradient is then added to the weight vector to yield the weight vector at the next time step. This is shown in (8):

$$\mathbf{W}_{k+1} = \mathbf{W}_k - \mu\nabla = \mathbf{W}_k + 2\mu\varepsilon_k \mathbf{X}_k \tag{8}$$

The algorithm is simplistic because a process such as averaging is not necessary for good convergence. The gradient is taken at each time step, which can be noisy at times. This is not a problem as the LMS method will still generally move towards the steepest descent on the performance surface. Normally, the '2' in (8) is dropped leaving only the gain constant, μ. Also called the step size, this constant is a tuning variable that is used to control the rate of convergence and stability of the adaptation. If the step size is set too high overshoot can occur causing an unstable system and excessive error. In a physical sense, this could lead to loss of control of the inputs to the specimen. The other tuning knob in LMS is the length of the weight vector. This generally depends on the physical system being optimized but can also be used to aid in the quality of the error's convergence towards a minimum value.

4.3. System Identification

The concept of system identification or system ID is relatively straight forward. The process begins with an unknown system. This can be a physical plant such as a vehicle or quarter-car suspension test rig. It can also simply be a software FIR filter or model. Typically, being an unknown system means the user does not know the input to output relationship or transfer function of that system. An example is represented as a linear system such as:

$$\mathbf{y} = \mathbf{Hx} \tag{9}$$

Here, the input, $\mathbf{x}$, might be known and the output, $\mathbf{y}$, might be a measurable quantity. However, the matrix, $\mathbf{H}$, which represents the transfer function of the system is unknown. This means that knowledge of only the input will not allow one to calculate the output and visa-versa. The system ID algorithm used in this study solves this problem using the ALC algorithm adapted using the LMS algorithm described in the previous sections. Figure 4-3 is a block diagram of a typical system ID scheme. In this diagram the LMS

algorithm is included in the adaptive filter block. The LMS algorithm uses the same input signal for adaptation as that which is fed into the adaptive filter.
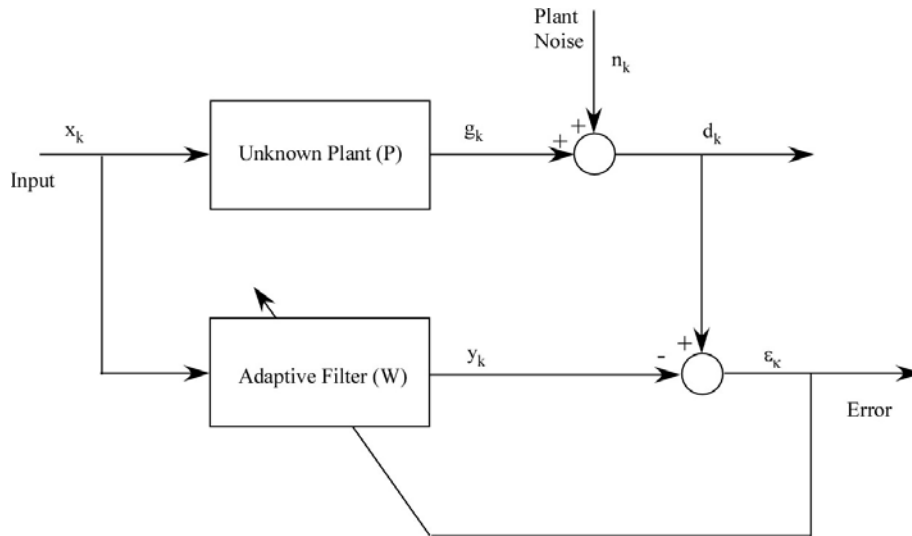
Figure 4-3 Typical System Identification Block Diagram

In this block diagram it is desired to learn the input to output relationship of the unknown plant, **P**. The adaptive filter, **W**, is a form of the adaptive linear combiner introduced in the first section of this chapter. Referring back to Figure 4-2 the desired signal d_k is now actually a combination of the unknown plant output and noise. The noise might be measurement noise for example. The noise does not have any effect on the convergence of the adaptive filter if it is not correlated to the plant input. The input to the adaptive linear combiner is the same in both the system ID problem and the ALC discussed earlier. In the system ID problem, the same input that is fed to the adaptive filter is also fed into the unknown plant. This input signal is usually a filtered white noise with a high spectral content in the frequency range of interest. The output of the adaptive filter is subtracted from the output of the unknown plant to form an error. This calculated error looks identical to that defined in equation (4). The calculated error is then used to adjust the weights of the adaptive linear combiner. This process may occur with several different methods. However, the LMS algorithm is selected for this study. The LMS algorithm uses the approximated error gradient to incrementally adjust the values of the weights. Again, the step size and the size of the weight vector are the adjustments.

These 'knobs' control the convergence rate, stability and quality of the identified model. With the weights of the adaptive filter converged to values that minimize the error, the adaptive filter will replicate the input to output relationship of the unknown plant. Thus the adaptive filter has identified the plant. The difference, between this and other plant models, is that this is not a parametric model with coefficients that have physical meaning. It requires virtually no knowledge of the plant. In the scope of this project a rough physical model of the quarter-car suspension and test rig is not needed to correctly identify the system. The model, created by this method, is purely empirical and is, in fact, an FIR filter that replicates the input to output relationship of the physical system.

The adaptive linear combiner used to perform this task works best identifying a linear system or a non-linear system operating in some linear range. The ALC does do a relatively good job of identifying a linear approximation of a non-linear system. The algorithm is extremely robust and lends itself to various forms of implementation in a controls problem. The identification process may be continuously updated to track non-linear changes in a system. For this research the system ID algorithm presented will be used in the subsequent steps to assist in controlling the quarter-car rig to replicate a desired response.

4.4.   Adaptive Control

The final processor used in this study is an adaptive controller. This controller is based on the same building blocks as the system identifier. The only real differences are the signals fed to the adaptive filter and LMS algorithms. Specifically, the controller used in this study is an inverse adaptive controller using the filtered-X LMS algorithm.

4.4.1. General Description

Figure 4-4 shows the basic block diagram for filtered-x LMS adaptive inverse control. In this diagram, the LMS algorithm is separated from the adaptive block to explicitly show the inputs to the algorithm. One channel of the desired response of the plant is fed into an adaptive filter, **W**. The output of the adaptive filter is fed into the

47

plant, **P**. The output of the plant is the measured response. The goal is to match this to the desired response. Meanwhile, the same desired response is delayed through a delay element. This delayed, desired signal is then subtracted from the output of the plant to give the error. This error is then fed into the LMS algorithm which adapts the weight coefficients of the adaptive filter. The other input to the LMS algorithm is the filtered-X signal. This signal is created from concurrently sending the desired signal through a plant model. The adaptive filter is adjusted until it creates an input to the plant that causes the plant to respond as desired. This minimizes the error between the plant response and the desired response. This process is discussed in detail in the following section.
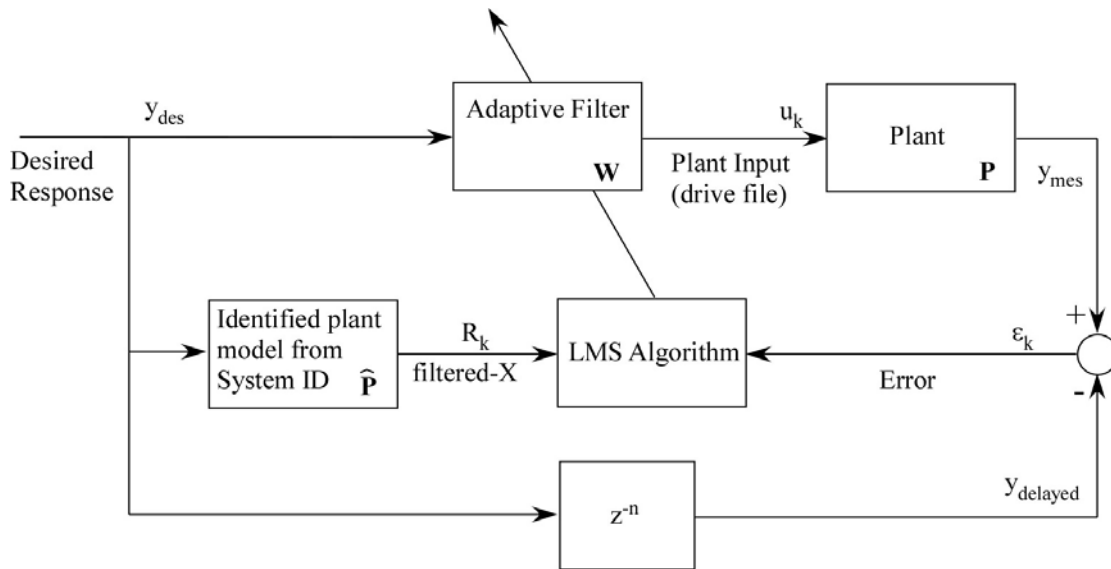
Figure 4-4 Filter-X LMS Adaptive Inverse Control Block Diagram

In this study the input to the plant, or drive file, needed to create this desired response is unknown but it is known that the input is correlated to the desired response. In order to create this drive file, the desired response is fed to the adaptive filter as it is adapting. While it converges, the drive file signal converges as well. This optimal input drives the plant to produce the desired response. Upon convergence, the adaptive filter is actually an inverse of the plant over the frequency range covered by the desired response spectrum. In an ideal system, the converged filter convolved with the plant will give a

perfect unity, or feed through, input to output relationship over the frequency range of interest.

4.4.2. Filtered-X LMS

The adaptation of the inverse adaptive filter occurs in much of the same way as the adaptive filter used in system identification. Unlike system ID, the adaptive filter is now matching the inverse dynamics of the plant. The filter weights are adapted using a virtually identical LMS algorithm. The adaptation is based on the error between the desired and measured system responses and also the filtered-X signal. The filtered-X signal is created by playing the desired response through an FIR filter model of the plant. In this study the model is identified using the system ID technique described in the last section. A nice feature of this method is that the identified model does not need to be extremely accurate for the controller to work well [27]. This speaks to the robustness of the algorithm.

The filtered-X version of LMS works exactly the same as in system ID with the exception of the input vector used in the gradient term. In system ID, the input to the adaptive filter was the same input going into the plant being identified. In the case of filtered-X, the input signal is first filtered through the identified plant model. Referring to equation (8) the input vector, $\mathbf{X}_k$, is now replaced with the filtered, desired response, $\mathbf{R}_k$ in Figure 4-4. The gradient is still estimated in the same manner. Only the input vector in the gradient term is changed. Therefore, the weights are adjusted at each time step by:

$$\mathbf{W}_{k+1} = \mathbf{W}_k + \mu\varepsilon_k\mathbf{R}_k \tag{10}$$

where,

$$\mathbf{R}_k = \hat{\mathbf{P}}\mathbf{y}_{des} \tag{11}$$

49

In control, the step size still dictates the rate of convergence and determines stability. The length of the weight vector can also still be tuned to affect convergence and quality of the inverse model. Additionally, the amount of delay of the desired response signal may be tuned to help produce a more accurate response.

5. SIMULATION AND RESULTS

The following chapter presents the first step taken to apply the principles discussed in the previous chapter. Before proving the worth of applying these control theories to the quarter-car rig, they were first tested purely in software. This was to help understand the process before risking damage to the hardware due to some unforeseen instabilities. The chapter begins by discussing the linear quarter-car model used to represent the unknown plant in software. System ID is then applied and the results of this identification are discussed. Finally, a desired response is produced and the inverse control algorithm is applied to replicate this response on the quarter-car model. The chapter closes with the results of the control study.

5.1. Quarter-car Model

To prove the concept purely in simulation, a plant was required. Though the plant is completely known for this simulation study, it is treated as an unknown system by the software for the purposes of identification and control. It is also useful to know the plant because this allows a direct measure of how well the FIR filters identify and control the system. For this study a simple linear two-mass quarter-car model was chosen. This was an obvious choice because it has the same types of input and outputs that the real quarter-car test rig has. In this case one velocity input and two acceleration outputs. The use of a quarter-car model as the plant was not required to demonstrate the concept. The quarter-car model was chosen simply to allow for a chance at being able to compare data from the test rig later in the study. Its use also aids in making physical sense of the results given.

5.1.1. Mathematical Model

The quarter-car model is the usual two degree-of-freedom vibration model of a single corner of a vehicle. Figure 5-1 is a diagram of the quarter-car model. This model is a two mass model which only concentrates on the vertical motion of the vehicle on one corner. The model contains a sprung mass and unsprung mass denoted by m_s and m_u

respectively. The coordinate associated with the motion of the sprung mass is called z and the coordinate associated with the unsprung mass is y. The suspension is modeled with a simple linear spring, k_s, and damper, c_s. The tire is modeled as a linear spring denoted by k_u. The road input to the tire is modeled as a velocity input called x_{in}.
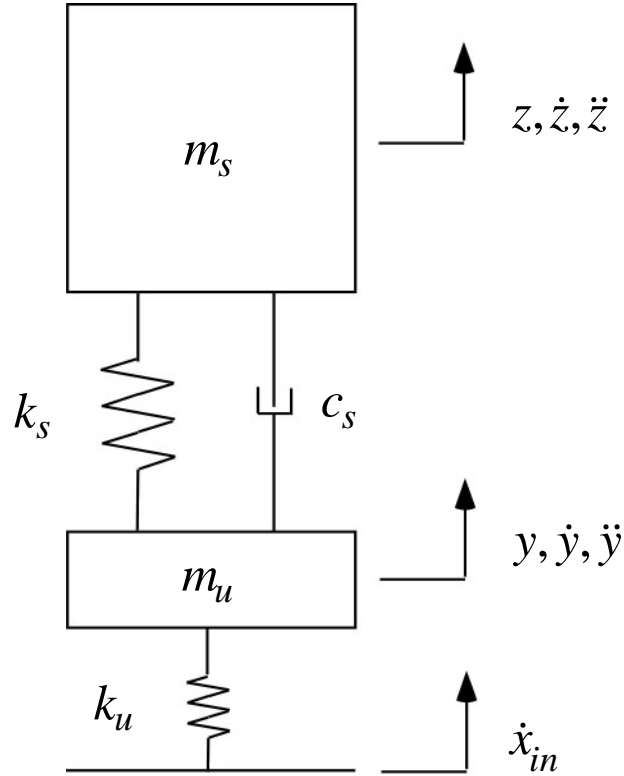
Figure 5-1 Two Degree-of-freedom Quarter-car Model

The schematic defined above was modeled using Lagrangian dynamics. In modeling, energy functions are defined for kinetic, T, and potential energies, V, as well as a pseudo-energy damping function, D. These functions are defined for this model as:

$$T = \frac{1}{2}m_s\dot{z}^2 + \frac{1}{2}m_u\dot{y}^2 \tag{12}$$

$$V = \frac{1}{2}k_s\left(z - y\right)^2 + \frac{1}{2}k_u\left(y - x_{in}\right)^2 \tag{13}$$

$$D = \frac{1}{2}c_s\left(\dot{z} - \dot{y}\right)^2 \tag{14}$$

52

With these functions defined the equations of motion can be computed for each coordinate based on an expanded form of the Lagrange's equation of motion. This equation is defined for each coordinate, q_i as follows:

$$\frac{d\left(\dfrac{\partial T}{\partial \dot{q}_i}\right)}{dt} - \frac{\partial T}{\partial q_i} + \frac{\partial D}{\partial \dot{q}_i} + \frac{\partial V}{\partial q_i} = 0 \tag{15}$$

The equation is set to zero because of a lack of forces or other loads from external sources. Because there are two degrees of freedom, there will are two coordinates and hence, two equations of motions. The model is a linear time-invariant system and can be put into a state space form for ease of computation in MATLAB. To put the equations in state form a vector of states must be defined. The state vector defined is

$$\mathbf{X} = \begin{Bmatrix} y \\ \dot{y} \\ z \\ \dot{z} \\ x_{in} \end{Bmatrix} \tag{16}$$

These states include the positions and velocities of the two masses and the displacement of the input. The input, $\mathbf{u} = \begin{bmatrix} \dot{x}_{in} \end{bmatrix}$, for this system is the velocity of the road. The standard continuous-time, state space is then:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \\ \mathbf{y} &= \mathbf{C}\mathbf{x} + \mathbf{D}\mathbf{u} \end{aligned} \tag{17}$$

The state matrices **A**, **B**, **C**, and **D** are defined by the two equations of motion and are found in Appendix C.

5.1.2. Parameter Values

For the purpose of having a ballpark comparison of experimental results such as frequency response to those in this simulation, the parameters were chosen reasonably close to the real Porsche suspension on the test rig. A high fidelity parameter identification was not performed for this study. The following table is a list of those parameters used in the model.

Table 2 Quarter-car Model Parameters

Parameters	Value	Unit
m_u	540	lbf
m_s	90	lbf
k_s	1480	lbf/in
k_u	5*k_s	lbf/in
zeta	0.4	

The values selected are meant to be a fairly accurate representation of the suspension on the quarter-car rig. The masses were measured directly as the components were in place at the time of this study. The suspension spring rate was measured directly. The tire spring rate was approximated based on an assumption that a tire spring rate was on the same order of magnitude as a race car suspension spring rate. It was not possible to directly measure the damping rate. Instead a damping ratio, ζ, was selected to be around 0.4. The damping coefficient was then calculated from this assumption based on the simple calculation:

$$c_s = 2\zeta \sqrt{k_s m_s} \tag{18}$$

These values were then used to fill in the elements of the state variables to complete the continuous-time, s-domain, state space model. The parameters were put into a MATLAB parameter file for ease of making changes. The state space variables were set up in a MATLAB script as well. With the variables filled the state space was created using the 'ss.m' MATLAB function.

5.1.3. Discretization

Because the programs are simulated in Simulink and later implemented in a dSPACE rapid-prototyping environment, the state space system was discretized. The Simulink simulation environment will automatically discretize equations of motion; however, if continuous elements are used in the code, there is little control over how the software discretizes them. For this reason, all of the code was discretized up front to have full control on how this process takes place. Thus, the state space system was discretized using a bilinear Tustin transform. Before discretizing the system, a sample rate and corresponding sample period was chosen. The frequency was chosen to be 1000 Hz which corresponds to a sample time of $T_s = 0.001$ s.

In MATLAB the function used for discretizing the state space system is called 'c2dm.m'. This function was run using the Tustin transformation. This type of transformation transforms a filter from the continuous-time s-domain to the discrete-time z-domain such that:

$$s = \frac{2(z-1)}{T_s(z+1)} \tag{19}$$

Here T_s is still the sample time in seconds. The 'c2dm.m' function also allows one to input the state space matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$. The resulting discrete state space system is:

$$\begin{aligned}
\mathbf{x}_{k+1} &= \mathbf{A}_d\mathbf{x}_k + \mathbf{B}_d\mathbf{u}_k \\
\mathbf{y}_k &= \mathbf{C}_d\mathbf{x}_k + \mathbf{D}_d\mathbf{u}_k
\end{aligned} \tag{20}$$

5.1.4. Frequency Response

For later comparison to the identified numerical model, the frequency response of the quarter-car model was examined. Figure 5-2 is a plot of the frequency response of the continuous time state space model. The frequency response shows the two expected resonance frequencies. The sprung mass appears to have a resonance near 5.3 Hz and the

unsprung mass has a resonance near 26 Hz. The phase remains relatively unchanged at low frequencies approaching DC. This would be expected as such low frequency excitation does not tend to excite the dynamics of the system. This low frequency phase is at 90 degrees because the transfer function is between a velocity input and acceleration outputs.

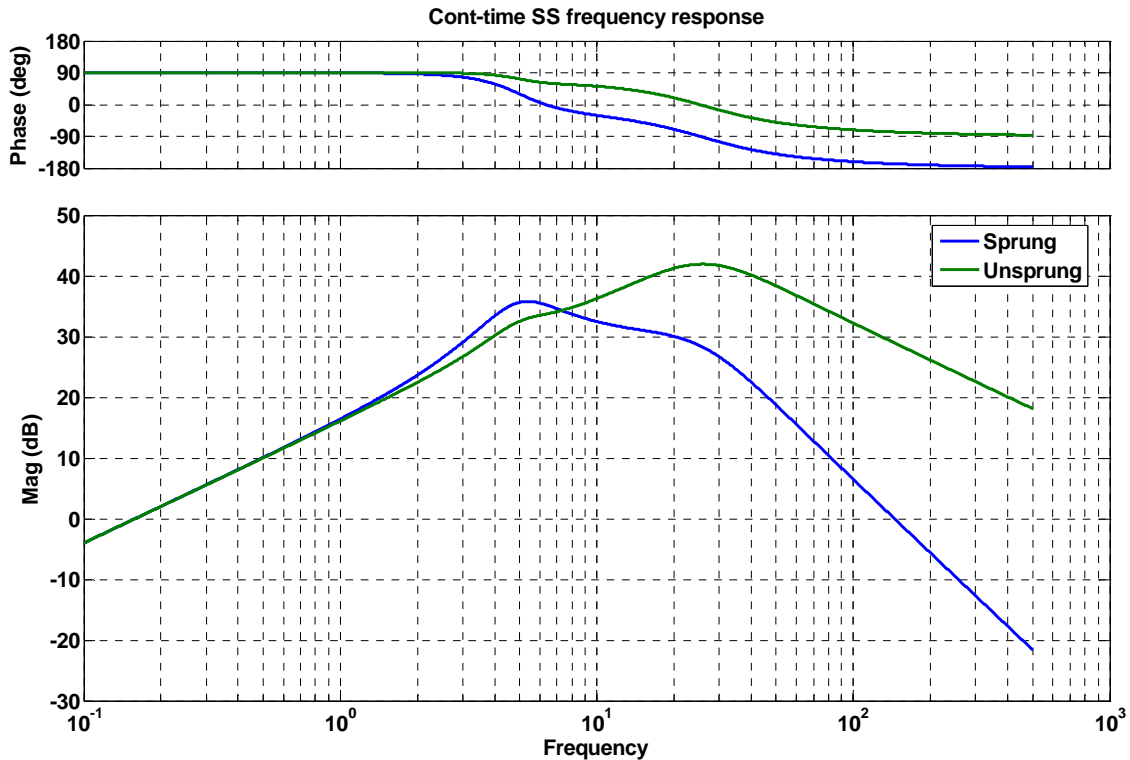**Figure 5-2** Frequency Response of the Analytical Quarter-car State-space Model

5.2. System Identification Study

With a quarter-car model created the next step was to test system ID in an analytical simulation. First, a Simulink model was created to implement the system ID technique based on the algorithm introduced in Chapter 4. Next, an excitation signal was defined as an input for the model and the simulation is run.

5.2.1. Simulink Model

Since there are actually one input and two outputs, there are really two transfer functions to be identified in this quarter-car model. Thus, each of the transfer functions must be identified simultaneously. Therefore, two system ID blocks must be used to identify these two distinct transfer functions. As mentioned, all of the software development was performed in discrete-time. This was to control how Simulink discretizes the model and also to make the transition into real-time control much more seamless. A model of the system ID algorithm is found in Figure 5-3.
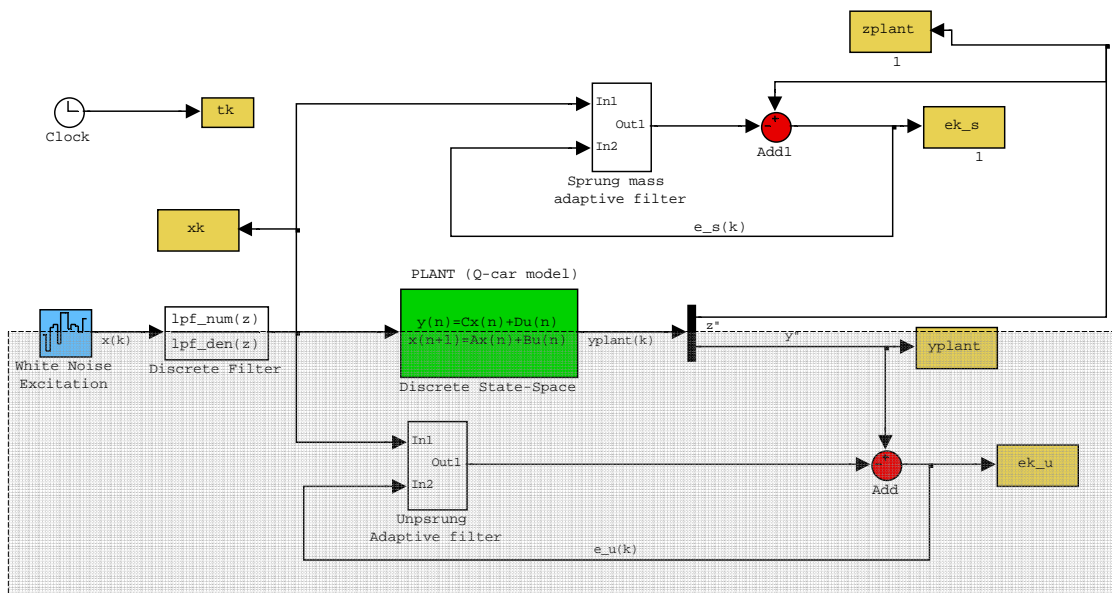

Figure 5-3 Simulink Block Diagram of Basic ID Scheme

In this diagram two adaptive filter blocks are introduced to empirically identify the dynamics of the quarter-car state space system. This is the same process described by Figure 4-3. Because each of these adaptive filters function the same way, only the operation of one loop will be discussed in detail. In this example the focus is on the adaptation of the unsprung mass numerical model. The loop for this algorithm is located in the shaded box.

The input to the in this case is a filtered white noise excitation. This signal is fed to the adaptive filter and also to the discrete state space model of the quarter-car. The

state space system outputs an acceleration signal for the unsprung mass. At the same time the adaptive filter block is adapting weights of the ALC using the LMS algorithm introduced earlier. The output of the filter is also an acceleration signal.

To understand the mechanics of implementing the algorithm in a Simulink model the adaptive filter blocks labeled 'Sprung mass adaptive filter' and 'Unsprung adaptive filter' are expanded from Figure 5-3. This detailed representation is in Figure 5-4. Here the same unsprung mass adaptive loop is still highlighted by the shaded box. Following the input to the filter more closely, the input signal goes to a tapped delay block. This block samples the input signal and files it into a vector format such that the first entry is the most recent signal and then the one before and so on. This is block performs the exact function of the tapped delay line discussed in section 4.1. This vector of tap-delayed inputs, X(k) is then split. The signal is then dotted with the vector of weight coefficients labeled Wu(k). This dot product is the same mathematical process as in equation (3). The result of which is the output acceleration of the adaptive filter.
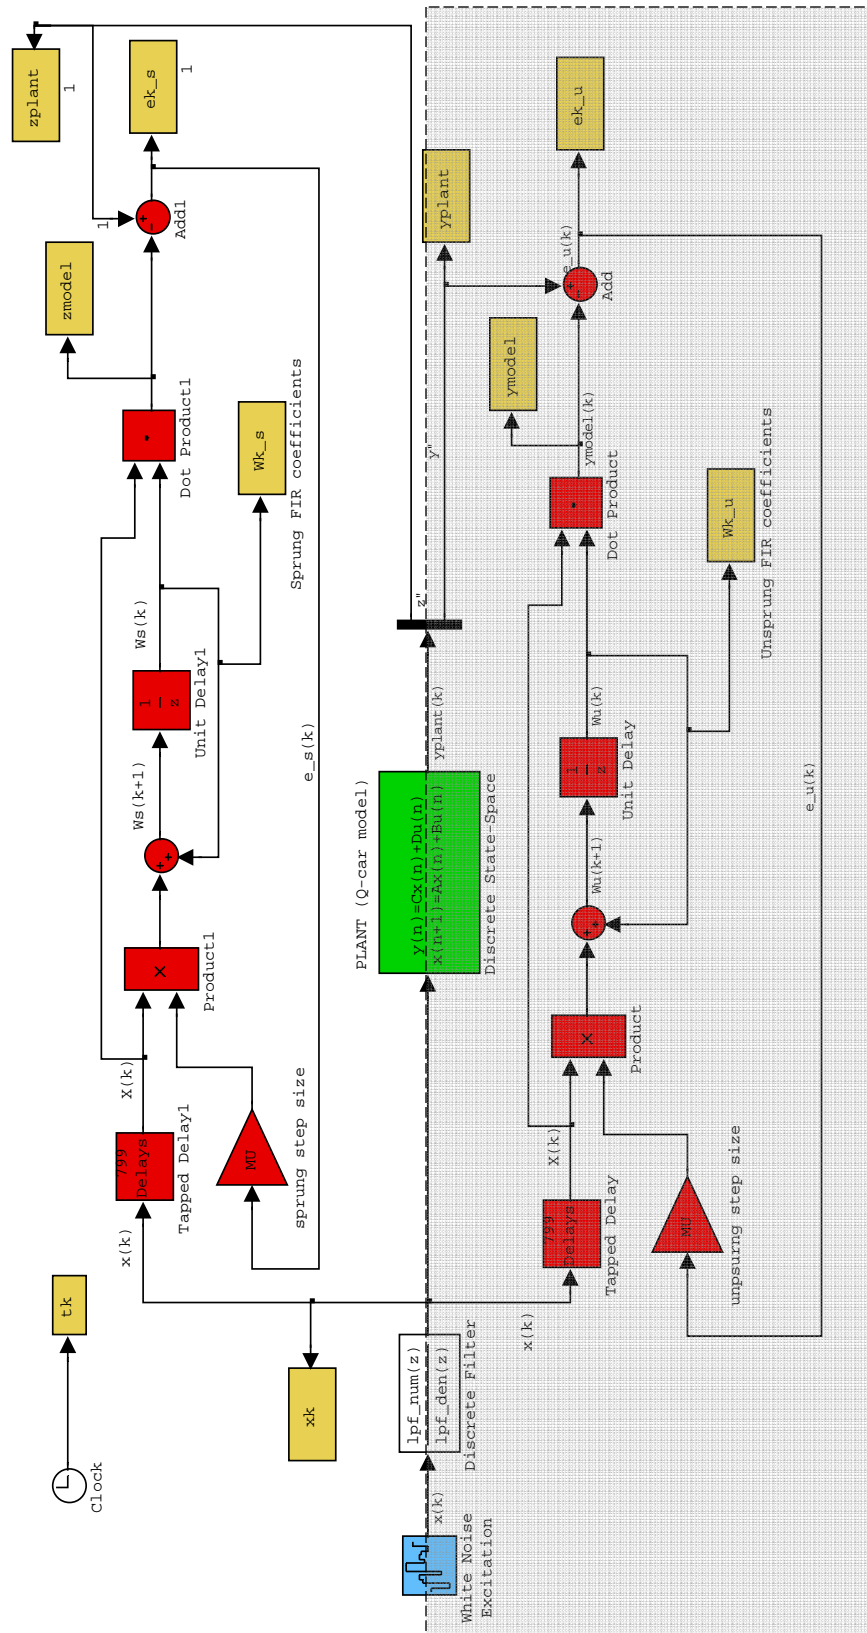
Figure 5-4 Simulink Block Diagram Model of Simulation ID Algorithm

59

The error created from subtracting the output signals is then fed back with a gain or step size multiplying it. This step size is then multiplied by the same tap-delayed input signal, X(k) and then added to the previous set of weights Wu(k). This addition provides the next corrected set of weights ,Wu(k+1) at the following time step. Using the appropriate math for a block diagram and referring to Figure 5-4, the following equation may be written:

$$Wu(k+1)=Wu(k)+MU*e\_u(k)*X(k) \tag{21}$$

This equation is basically the same as (8) in section 4.2. The identical process takes place for the sprung mass loop as well. The simulation is allowed to run until both of the error signals are minimized. All of the yellow 'To Workspace' blocks such as 'xk' or 'tk' were used to store the data to the workspace in MATLAB for post processing.

5.2.2. Excitation Signal Shaping

Care must be taken when choosing the filter for shaping the input excitation. The idea is to use a band limited white noise excitation. However, this signal should be limited such that it is allowed to excite all of the dynamics of the system that are of interest in modeling. The frequency response study of the linear quarter-car model indicated that the system had two resonances which did not exceed 30 Hz. Therefore, exciting the system, at frequencies much higher than this frequency, does not really help the model converge. Also, because the sample rate was set at 1000 Hz, care was taken not have high power signal near the Nyquist frequency of 500 Hz which would cause aliasing problems in the data.

The solution was to use a fourth-order, low-pass Butterworth filter with a 60 Hz break frequency. Coefficients for the numerator and denominator of the filter were created with the 'butter.m' function in MATLAB. A discrete filter block was then used in the Simulink model to actually filter the white noise before it was fed into the rest of the model. The two blocks are cut from Figure 5-4 and shown below in Figure 5-5.
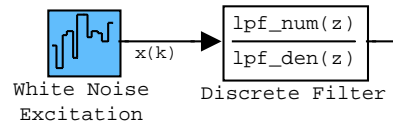
60

Figure 5-5 Model Detail of Filtering the White Noise Excitation

5.2.3. Numerical Results

Several tests were run to achieve the best error convergence possible. The 'tuning knobs' noted in the previous chapter were used to control the convergence and quality of the model. For system ID the same model size and step sizes were used for both identification paths. The best results occurred with a step size of 0.0002 and a model size of 150 coefficients.
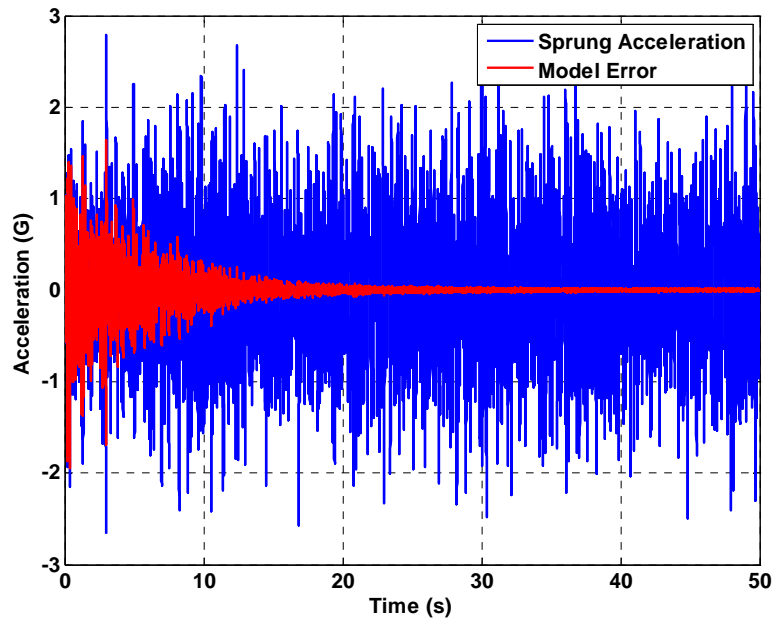

Figure 5-6 Error Convergence of Sprung Mass Numerical Model Identification
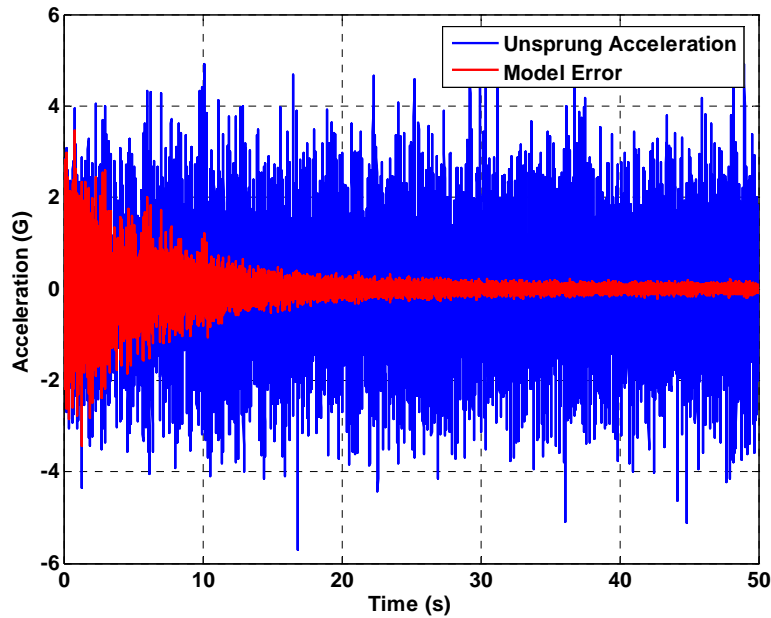
Figure 5-7 Error Convergence of Unsprung Mass Numerical Model Identification

Figure 5-6 and Figure 5-7 are plots showing how the error converges with time. The error is placed over top of the mass acceleration response for both the unsprung and sprung masses. The error convergence happens quite quickly. For the sprung mass identified model the error is minimized in about thirty seconds. It takes up to about fifty seconds for the error to converge to a comparable level for the unsprung mass path. The root-mean square (RMS) of a converged section of the error signal for the unsprung mass was compared to the RMS of the unsprung response signal. This comparison showed that the RMS value of the error converged to within 2.3% of the RMS of the unsprung acceleration signal. Likewise, the RMS value of the error signal for the sprung mass converged to less than 1% of the RMS value of the sprung mass response.

Another way to interpret error convergence is by looking at the accelerometer signal power versus the error signal power. The signal in either case is calculated by squaring the signal and the filtering it with a low-pass filter. Again a Butterworth filter was utilized to create this low-pass filter. The 2-pole filter had a break frequency of 0.3 Hz. The use of a low-pass filter removes most of the dynamics from the signal leaving only a relatively clean DC power signal. The plots in Figure 5-8 and Figure 5-9 show the power levels of the signals as the error signal is close to converging on its minimum. The

power levels are displayed on a decibel scale. In general, a 10 dB difference between the desired signal and the error signal is considered acceptable and 20 dB or more of a difference is considered excellent. These plots show that the power level of the sprung mass error signal decreases to nearly 60 dB lower than the accelerometer signal. Likewise, the error of the unsprung mass error decreases to nearly 30 dB lower than the unsprung accelerometer signal. Both of these converged power level drops are considered fantastic, which is to be expected for a simulation without external noise added.
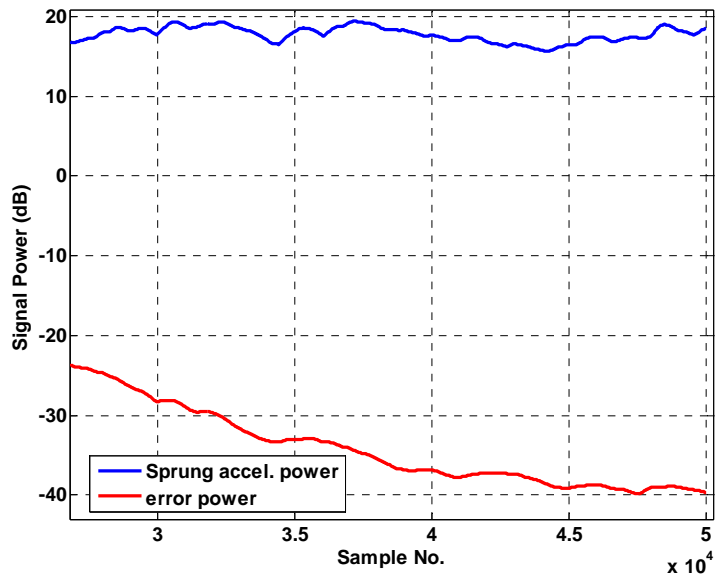

Figure 5-8 Sprung Accelerometer and Error Signal Powers
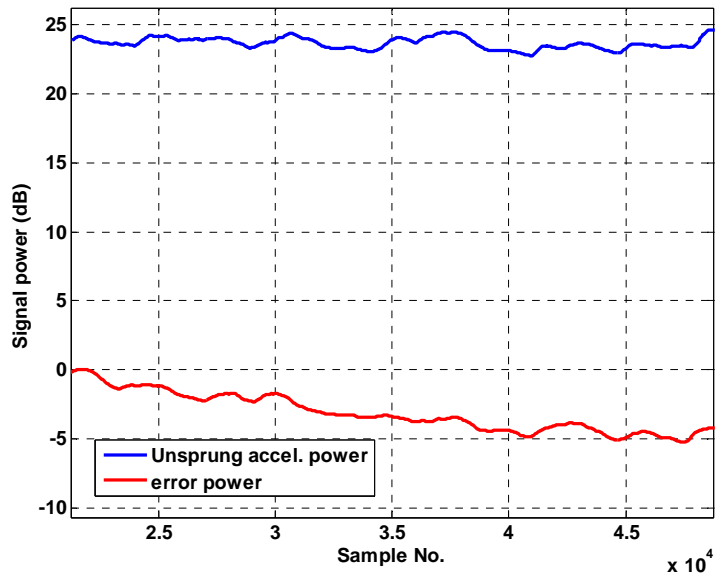
Figure 5-9 Unsprung Accelerometer and Error Signal Powers

Finally, the converged weights of the FIR models were studied. The weights of the FIR filter actually represent an impulse response of the adaptive model when plotted versus time. In this study, the weights of the adaptive filters actually match the impulse response of the unknown system they are modeling very well, further showing how well the filter adapts to look like the identified plant. This is demonstrated in Figure 5-10. It was expected that the adaptive filters would do a good job of identifying the quarter-car model extremely well since the quarter-car model was a linear system.
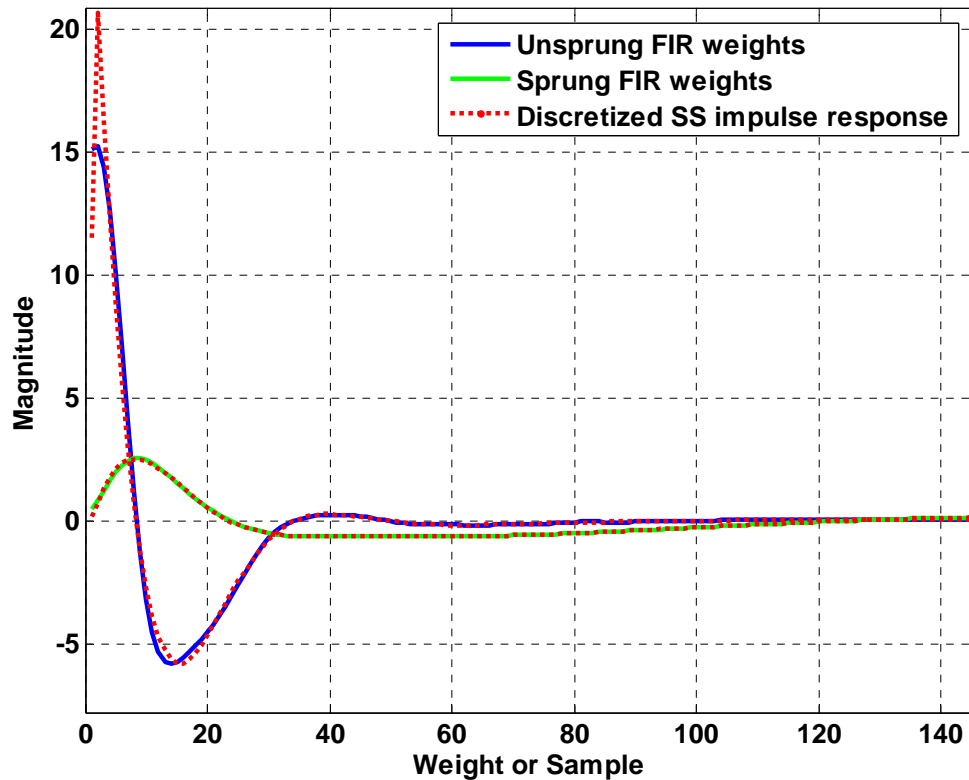
Figure 5-10 Impulse Response Compared to Converged ID Weights

A frequency response of the adaptive FIR model was computed and compared to the frequency response of the original linear state space model. The results are plotted in Figure 5-11. The figure shows that the frequency response of the adaptive filter is extremely accurate in replicating the magnitude and phase response of the quarter-car plant. One place that the model shows to have a slight amount of trouble is in the extremely low frequency region of the unsprung mass model.
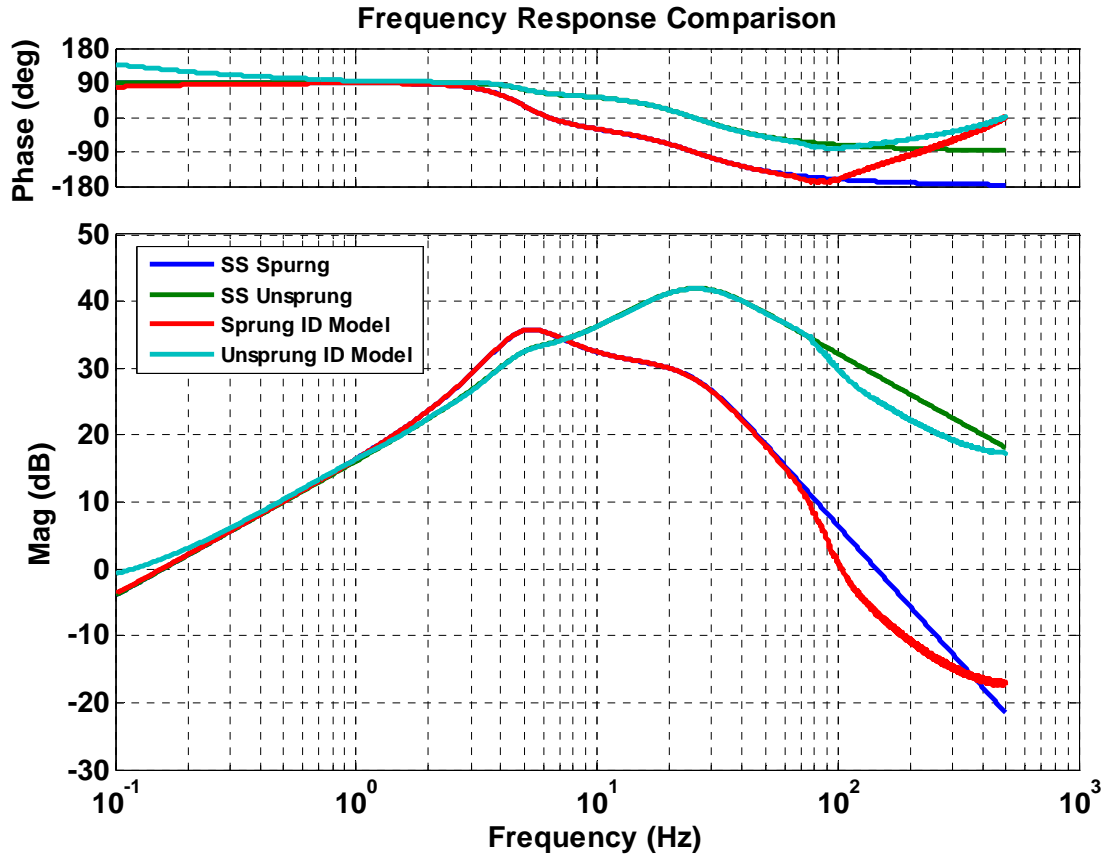
Figure 5-11 Comparison of Model and Adaptive Filter Frequency Responses

Referencing the phase portion of the figure the FIR model for the unsprung path does not match the state space quarter-car model below 1 Hz. This is because accurate phase measurements cannot be made when the transfer function magnitude is so small. In any case, the FIR adaptive filter does still replicate the dynamics of the quarter-car analytical model very well.

5.3. Adaptive Control Study

The second half of the simulation study was to control the quarter-car model such that it replicated a pre-made acceleration signal. Again a Simulink model was created to implement the simulation. The adaptive inverse control algorithm also introduced in Chapter 4 was implemented to accomplish this task. The identified FIR filters from the previous section were used in the filtered-X LMS algorithm.

5.3.1. Simulink Model

The Simulink model was designed to take the basic shape of the model which would eventually be run in dSPACE. This was to help the transition into real-time control easier. Figure 5-12 shows a block diagram illustrating the basic setup of the control software. The control model was broken up into four distinct blocks: analog-to-digital converter (ADC), digital-to-analog converter (DAC), system ID and control.
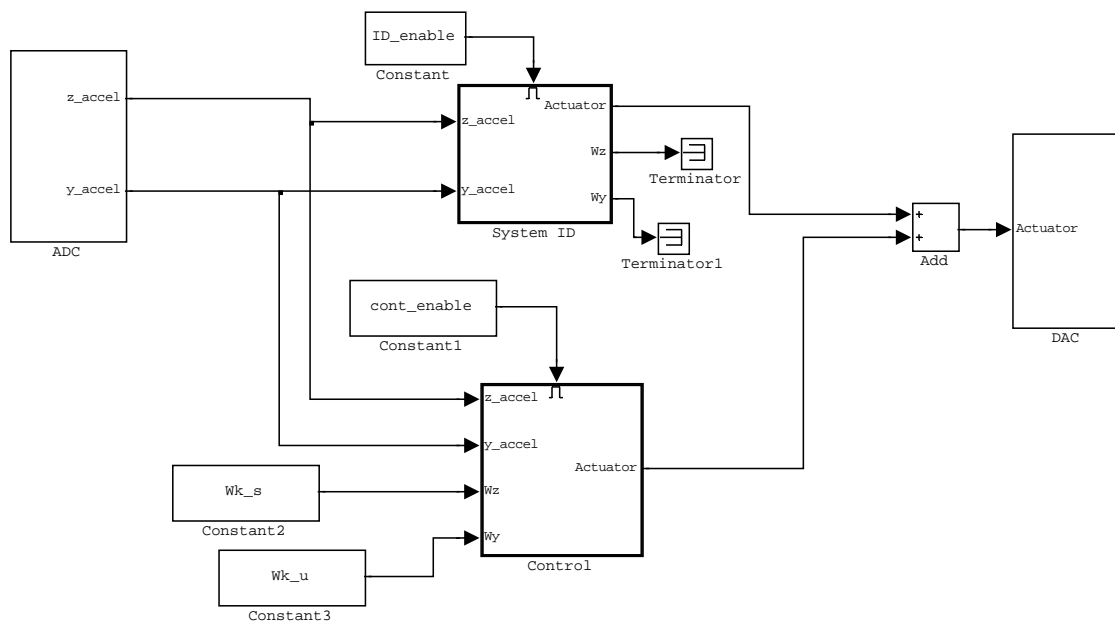
Figure 5-12 Basic Construction of Simulink Control Block Diagram

Of course, in a pure simulation study there are no analog to digital or digital to analog converters. Inside these blocks there is basically just a wire to feed the signals through. These wires will later be replaced with the appropriate converters such that dSPACE can route the signals properly. Thus, actuator signals that are fed into the DAC block in simulation are simply fed through directly to the ADC block where the model of the plant resides. The system ID and control blocks are enabled subsystems. The 'ID_enable' and 'cont_enable' blocks are constants with either a 1 or 0 value. These constant blocks are used to turn the respective subsystems on or off. Thus, when the simulation is running, the quarter-car model is first identified using the system ID

subsystem. Once the model is identified the system ID is switched off and the controller is turned on to allow for replication of signals on the model. The identified FIR models are stored in the two blocks labeled as 'Wk_s' and 'Wk_u' for the sprung and unsprung masses respectively.   When required, these weights can be fed into the controller subsystem.

The quarter-car state space model has been placed inside the ADC block. In a real-time simulation the plant would obviously be hardware and the ADC and DAC blocks would be the means to get signals from the software environment to real signals in the hardware environment.   This software design mimics this by simply keeping everything in a software environment. Detail of the DAC block is shown in Figure 5-13 (a) and the ADC block is shown in Figure 5-13 (b).
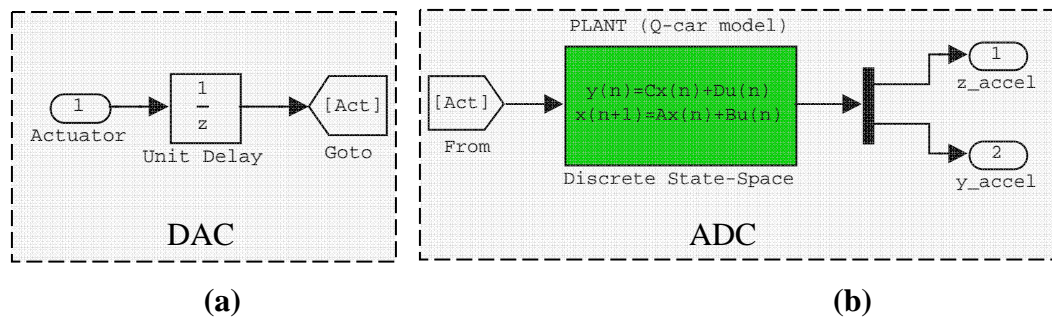

Figure 5-13 Detail of (a) DAC Block and (b) ADC Block

Following the figures from left to right, the feed-through nature of these blocks is now apparent. An actuator signal is fed into the DAC block. There is a one sample delay on the actuator signal to break an algebraic loop in the software. This is a simulation anomaly and is not required in the real-time implementation. The actuator signal is then fed directly from the DAC block to the ADC block via 'Goto' and 'From' blocks. These blocks are Simulink's way of transmitting signals without using wires. The actuator signal is then fed to the discretized state-space quarter-car model developed at the beginning of the chapter. The outputs are then separated or 'demuxed' and output to the ID and controller.

Inside the ID subsystem are the exact same inner workings discussed in the previous section. The only exception is the relocation of the state space model such that

the controller is able to have access to it. The details of the controller are illustrated in Figure 5-14. The goal of this simulation is to replicate a pre-defined acceleration signal on the previously identified quarter-car model. When viewing Figure 5-14 in a landscape format the block diagram will be discussed starting with the desired signals on the far left.

For this problem there are actually two signals for which replication is desired. This is obvious because there are two moving masses in the quarter-car model. Unfortunately, there is only one input to the system. Therefore, there can only be one controller. Only one of the desired signals can actually be input to the controller as the reference signal. The controller is an inverse of the plant being controlled. Since the plant actually has two transfer functions associated with it only one path can be inverted. The unsprung mass signal was chosen as the reference due to its higher frequency content and what would likely be a more interesting result in the racing industry. To begin the discussion the focus will be on shaded portion of the diagram.
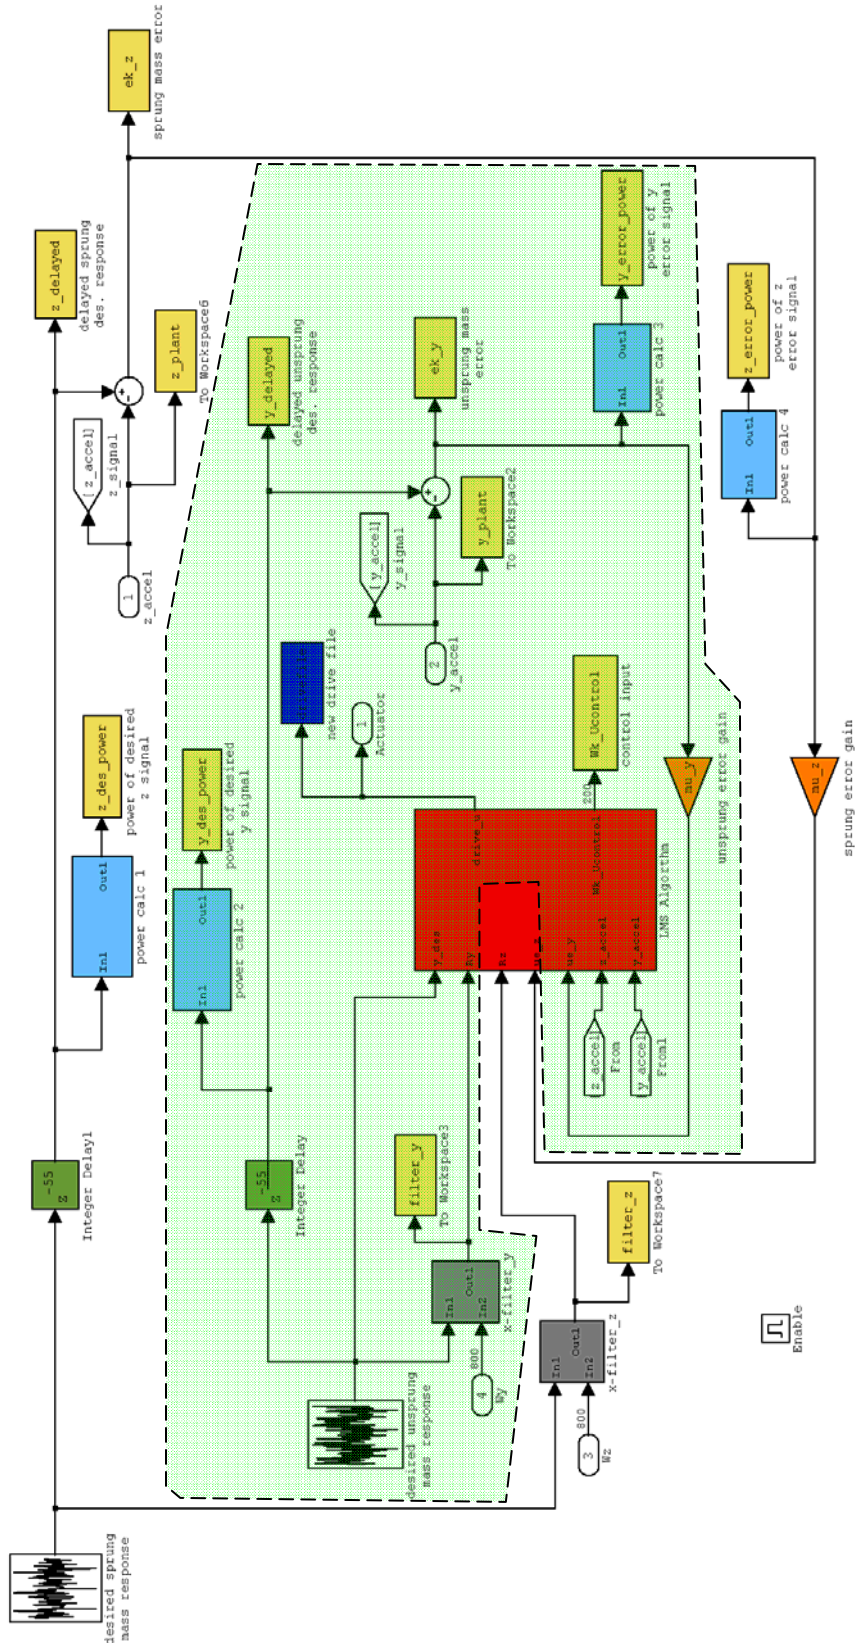
Figure 5-14 Detailed Simulink Diagram of Control Block

70

The desired unsprung mass acceleration signal is fed into the adaptive filter which is being adapted with the usual LMS algorithm. The output of the filter is a drive profile, called 'drive_u' in the block diagram. This signal then sent out of the control subsystem to the DAC block and thus fed to the state space model. The response of the model is then fed back into the control subsystem at the input blocks labeled 'y_accel'. The desired signal is also fed into a delay block called 'Integer Delay'. This block delays the desired signal before the signal is compared to the output of the quarter-car model. The error is produced by subtracting the actual response with the delayed desired signal. A programmable gain or step size, 'mu_y' is then placed on the error signal which is then fed into the adaptive algorithm. The other input to the adaptive algorithm is the filtered-X signal. This signal is the output from the 'x-filter_y' block. This block has two inputs: the desired unsprung mass signal and the FIR filter model identified in the previous step. This vector of weights is brought into the subsystem at the 'Wy' block in Figure 5-14. Inside the 'x-filter_y' block the desired signal is run through a tapped delay to create a vector the same size as the identified model. This desired response vector is then dotted to the vector of FIR weights to create a single r_k value. This value is later fed through a different tapped delay to create the vector of filter-X coefficients, $\mathbf{R}_k$, in equation (10).
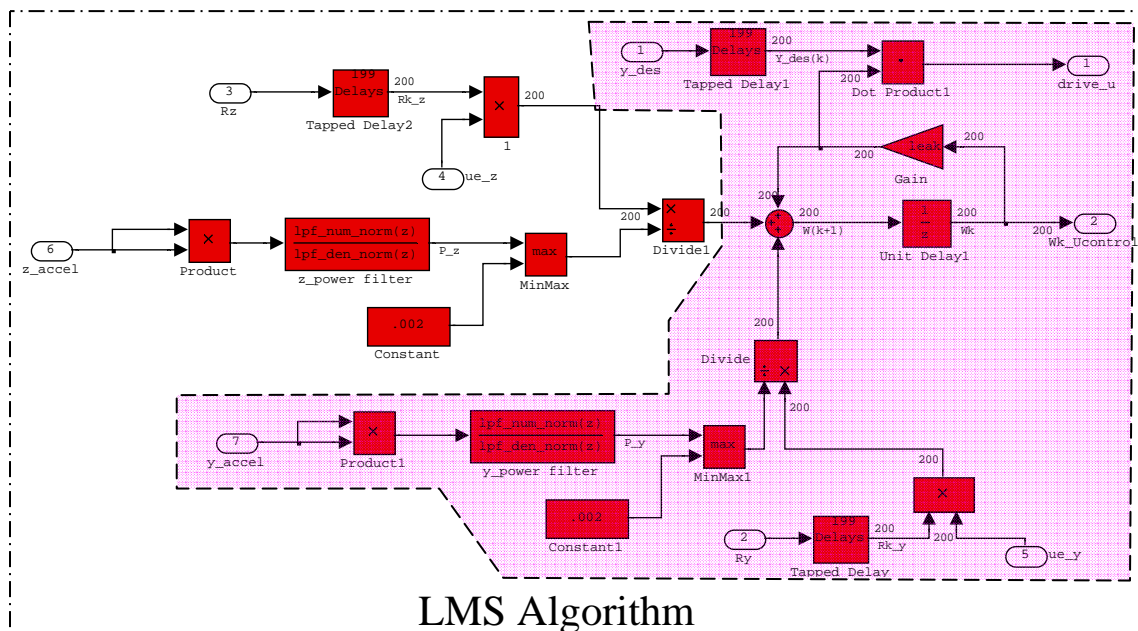

Figure 5-15 LMS Algorithm Detail

71

The 'LMS Algorithm' subsystem from Figure 5-14 is expanded in Figure 5-15. Again, only the loop containing the unsprung mass, shown in the shaded area, is the current focus. In this figure the r_k signal comes in towards the bottom of the figure at block 'Ry'. The signal is tapped-delayed and then multiplied by the gained error, 'ue_y'. To make tuning the step size easier, an option to attenuate the error signal with the signal power of the desired signal is included. This power signal is calculated and fed into a division block, thus attenuating the signal. For this discussion the assumption is that the attenuation is unity. Thus the error signal is not actually affected by signal power.

The weights of adaptive inverse control work on the same ALC principals as before. The weight at step k+1 is defined by adding the previous set of weights to the product of the filtered-X signal, step size and error signal. This is represented by equation (22), referencing Figure 5-15.

$$W(k+1)=W(k)+u*e\_y*Ry \qquad (22)$$

Again this equation is similar to that of (10).

The filter is then sent through a unit delay and then dotted with the tapped delayed desired signal. The product of this dot product is the output of the adaptive inverse control filter. This output signal is also considered the new drive signal for the quarter-car model. This process continues to adapt the FIR filter until the error between the desired response and actual response is minimized.

Closer examination of Figure 5-14 and Figure 5-15 reveal that there is an additional error loop. There is not an additional set of control filter weights but this is not a problem. The aforementioned control filter may also be adapted using the error between the sprung mass desired response and actual response from the quarter-car model. Referring back to Figure 5-14, the desired sprung mass response data is introduced into the control subsystem the same way but this time it is only used to produce an error signal and not fed into the controller. This signal is routed into the adaptive LMS algorithm the same way as the unsprung error signal. This signal is run through the identified input-to-sprung-mass output model at the 'x-filter_z' block. The

resultant 'Rz' is fed into the adaptive algorithm and the same tapped-delay and multiplication process takes place to help adapt the FIR controller weights. The resulting equation for the weight adaptation, shown below, is a slight modification of (10):

$$\mathbf{W}_{k+1} = \mathbf{W}_k + \mu_{yk}\varepsilon_{yk}\mathbf{R}_{yk} + \mu_{zk}\varepsilon_{zk}\mathbf{R}_{zk} \tag{23}$$

Thus the adaptive inverse controller for a single path may be adapted by the error signals of several paths in the same dynamic quarter-car model. This is a very common modification to filtered-x LMS gradient descent adaptation algorithms.

5.3.2. Desired Response Generation

To run this controller simulation some 'experimentally collected' data was required. Since this portion of the study was performed purely in simulation, the response data was also created using a MATLAB program. To produce the response data, a filtered white noise was input to the original continuous-time quarter-car simulation. This white noise was shaped in much the same way as the excitation noise used for system identification. The filter chosen was a four-pole Butterworth filter with a 50 Hz break frequency. The break frequency was chose to be slightly lower than the identification white-noise filter break frequency. It is good practice to use higher identification bandwidth than what the specimen normally sees in a test. This is to ensure that all of the dynamics are properly captured by the ID model.

A vector of white noise was created using the 'randn.m' function in MATLAB. The low-pass filter was created and the white noise was filtered through it using the 'filter.m' function. Finally, the filtered white noise input was simulated through the continuous-time state space model using the 'lsim.m' command. The resulting accelerations were stored for later use. A sample of the desired acceleration signals are plotted in Figure 5-16. The acceleration magnitudes of the desired signals are somewhat small. This is not an issue in practice because the system is linear.
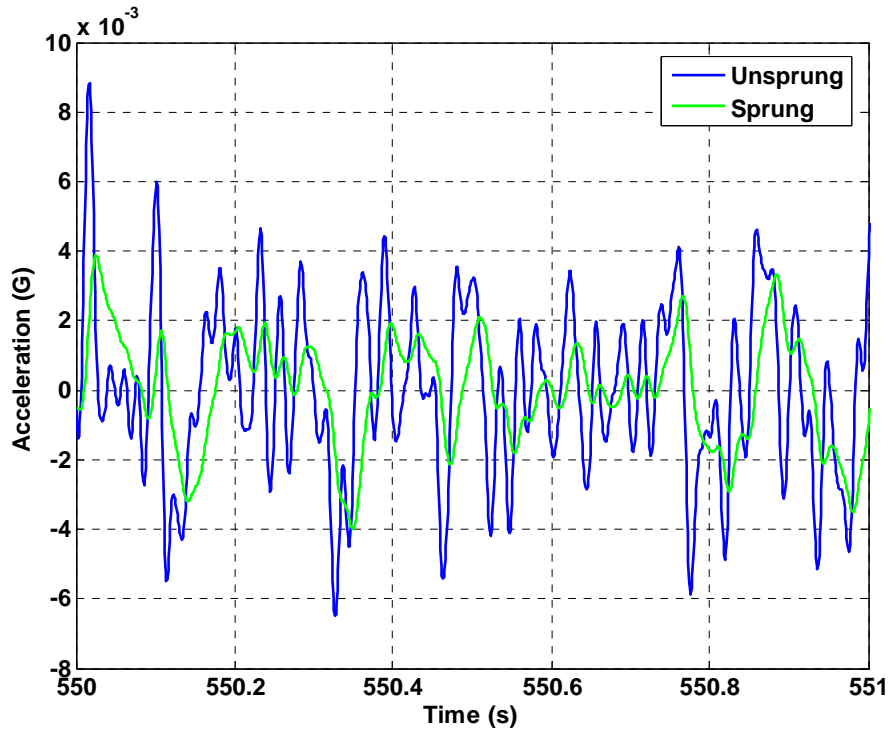
73

Figure 5-16 Sample Desired Acceleration Response Generated with State-space
Quarter-car Model

5.3.3. Numerical Results

The results of this numerical study indicate that adaptive inverse controller was able to produce an input, to the quarter-car model, that reproduced the desired accelerometer signals very well. The usual tuning parameters were adjusted to improve the convergence and results. For the final configuration the number of inverse adaptive control filter weights was set to 350. The step size for the unsprung mass was set to 1.3e-8 and twice that for the sprung mass loop. The z-delay on the desired signal was set to 115 samples.
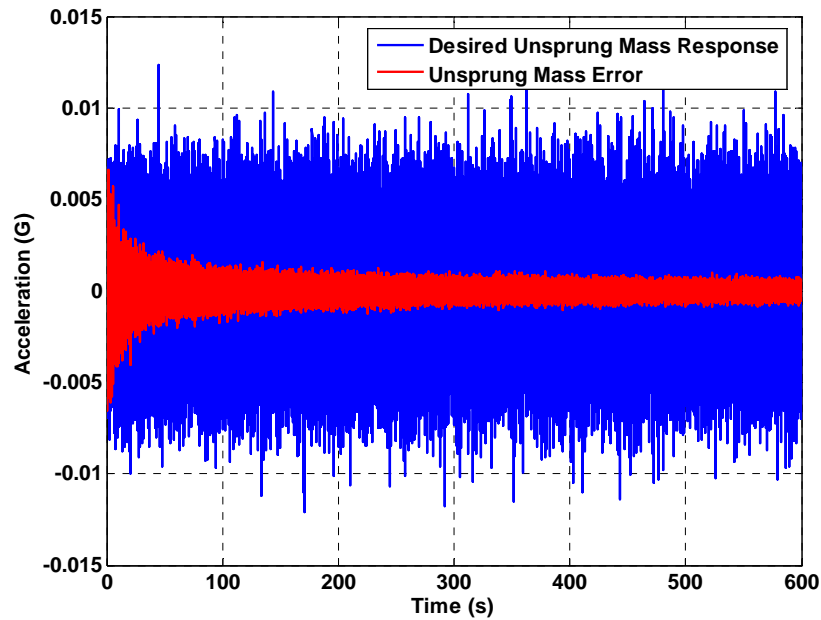
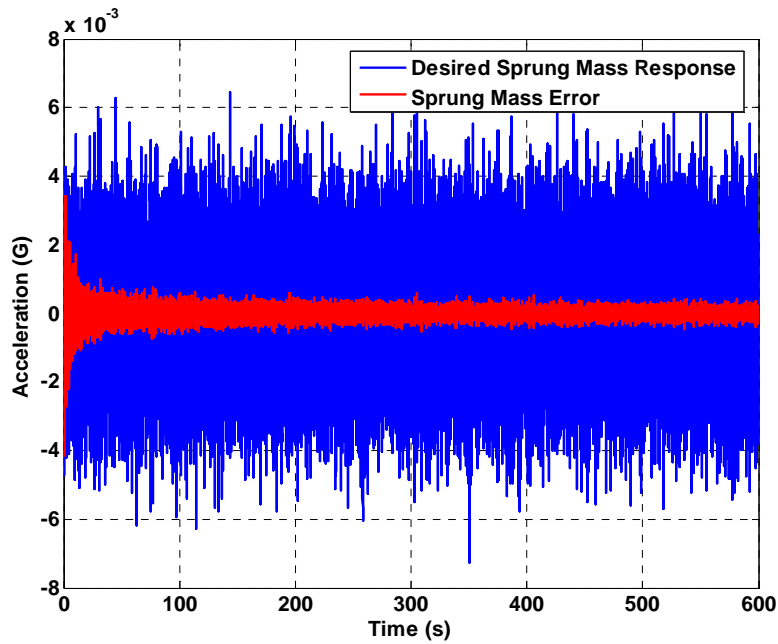**Figure 5-17** Unsprung Mass Desired Response and Error Convergence

Figure 5-18 Sprung Mass Desired Response and Error Convergence

The plots in Figure 5-17 and Figure 5-18 show that the error did not converge as well as in the system ID simulation. It did occur very rapidly however. The RMS value

of the unsprung mass response error converged to within 8.4% of the RMS of desired signal. Similarly, the RMS of the sprung mass response error was 7.5% of the RMS of the desired sprung mass response. Figure 5-19, found below, is a half-second window showing a comparison between the desired and actual responses. The controller appears to make the signal shapes and bandwidth match very well.
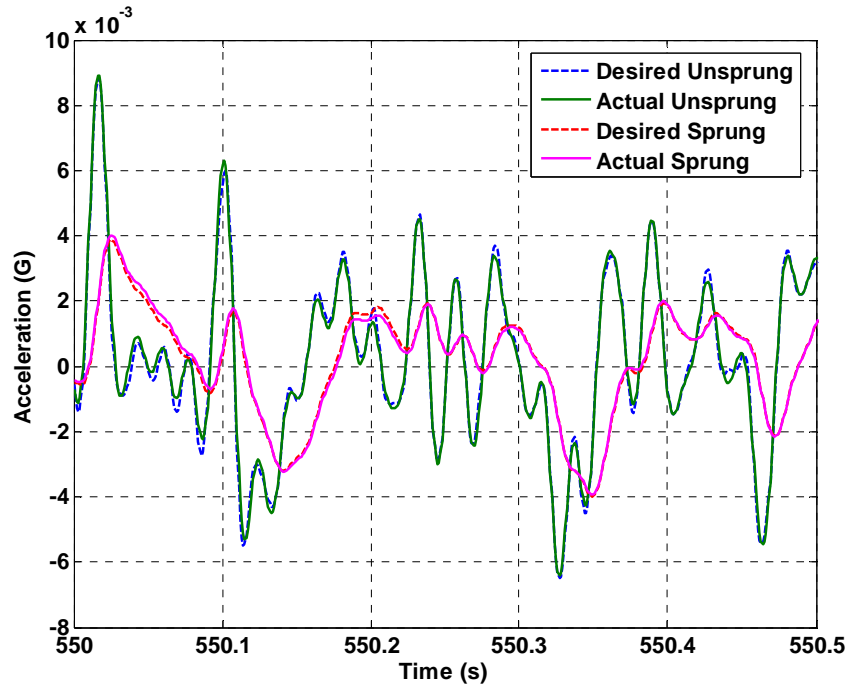

Figure 5-19 Sample of Desired Compared to Actual Response

Similar to the system ID the signal powers of the error and desired responses were investigated. Figure 5-20 and Figure 5-21 show the signal powers during convergence. Again, the criterion for good convergence in vibrations is a 10 dB reduction from the source signal to the error signal. The error signal power of the unsprung mass is approximately 22 dB lower than that of the desired signal. Likewise, the error signal of the sprung mass is about 23 dB down from the desired sprung mass signal.
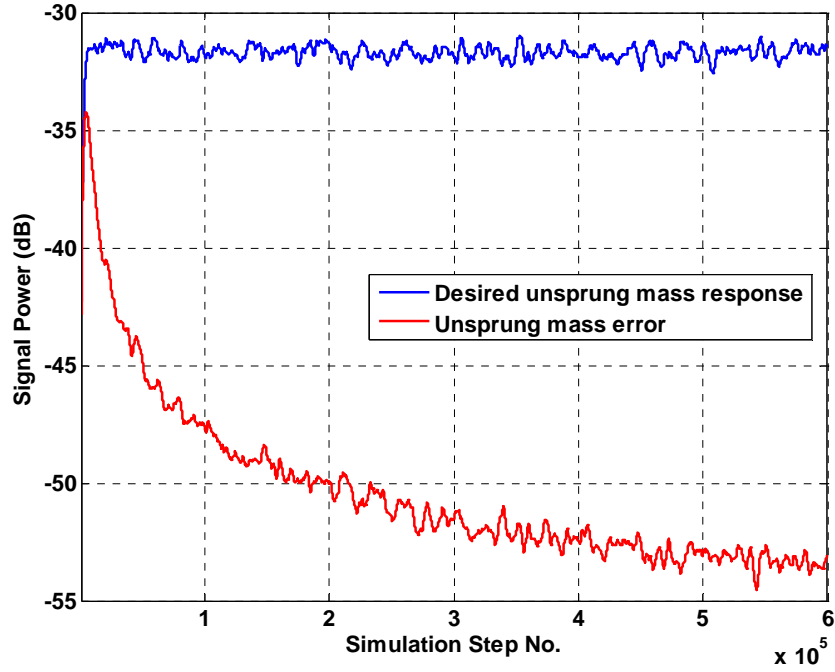
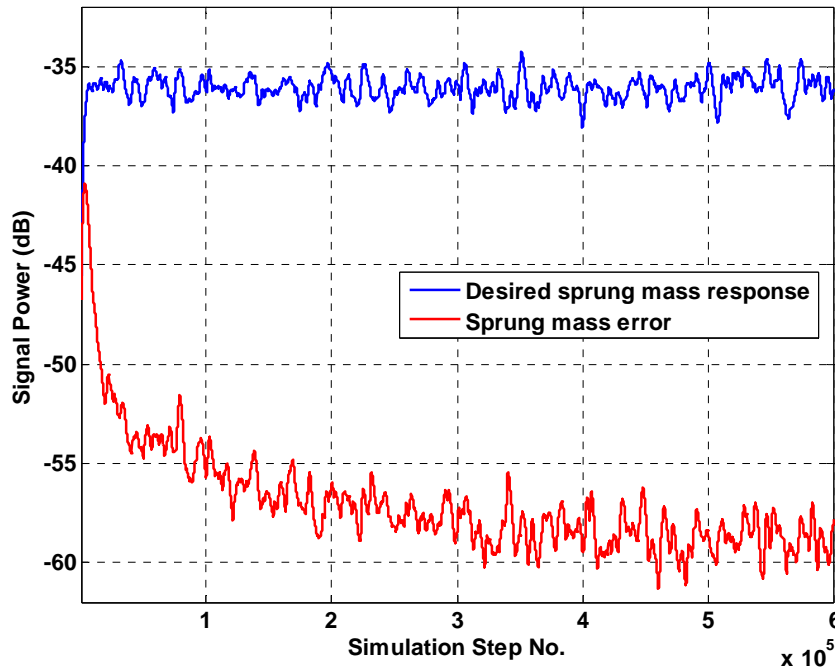**Figure 5-20** Signal and Error Power for Unsprung Mass

Figure 5-21 Signal and Error Power for Sprung Mass

The control weighs are plotted in Figure 5-22. The weights represent the impulse response of the inverse of the quarter-car transfer function. This response can be moved

forwards and backwards by adjusting the amount of delay in the desired signal before comparing it to the actual signal. A goal is to move the response such there are near-zero weights on both the leading and trailing tails of the larger response.
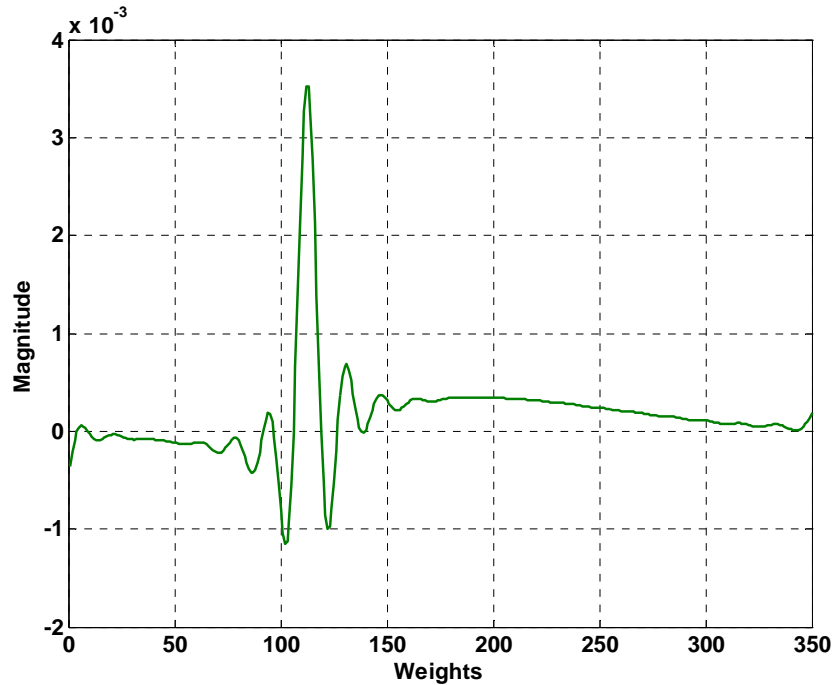
Figure 5-22 Converged Control Weights

Another way to gauge the quality of the inverse model is to convolve the FIR controller weights with the original, unknown plant. To do this, the frequency responses of the controller weights and the original plant were computed. These two frequency responses were multiplied together. Also, the frequency response of a 115 sample delay was calculated and then divided into the resulting frequency response. Both of these calculations were done on an element by element basis at each frequency. Division by the delay response function was performed in order to see more clearly the phase response. The resulting frequency response, of the convolved transfer functions, is shown in Figure 5-23.
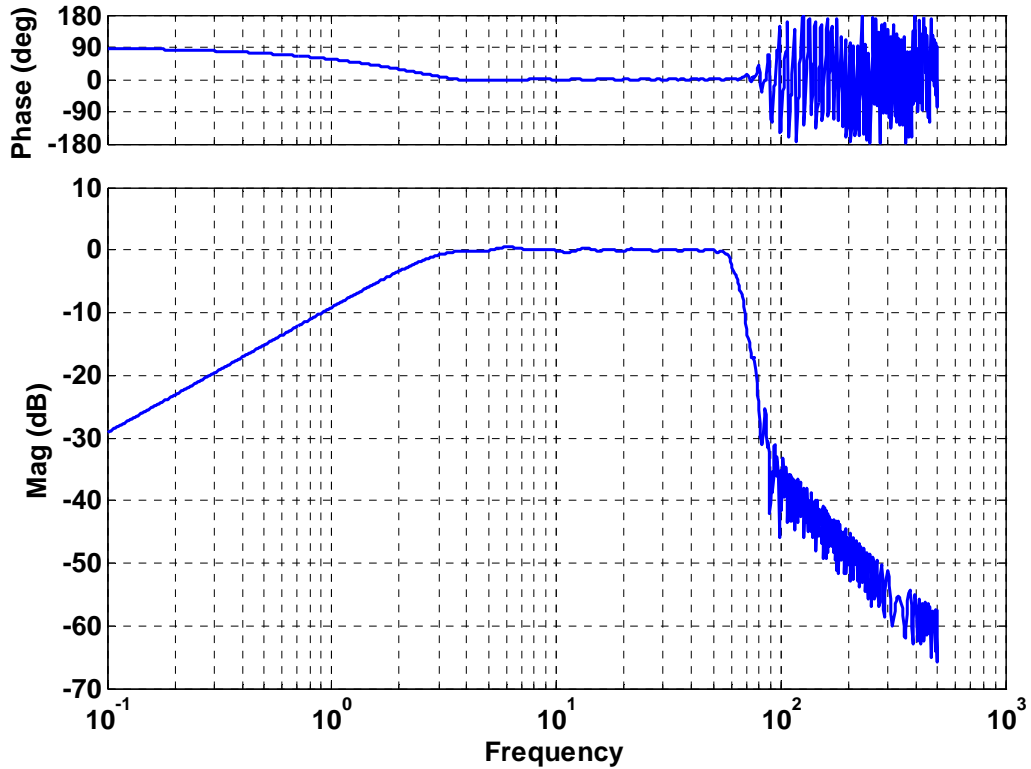
Figure 5-23 Frequency Response Function of the Convolved Controller and Plant

This plot clearly shows how well the adaptive inverse controller works. As expected the convolved controller and plant act as a feed through device only over the frequencies of interest. This is indicated by the 0 dB magnitude of the frequency response between 3 – 60 Hz. It is also noted that there is a zero phase shift in this frequency range as well. Below 3 Hz the controller appeared to have some trouble identifying the inverse of the plant. This is likely due to lack of power in the low frequency signals.

6. EXERIMENTAL PROCEDURES AND RESULTS

With proof of concept in simulation, the final step is implementation to the quarter-car test rig. This chapter follows much of the same structure as Chapter 5. However, this is where everything comes together with hardware-in-the-loop. The chapter begins with a brief description of the test setup including sensors and data acquisition. Next, the system identification setup and tests are discussed along with the results of the study. Finally, the control algorithm is implemented on the hardware. The chapter closes with a discussion of the results.

6.1. Test Setup

This section details the equipment used for the physical testing. The instrumentation and sensors will be covered along with the basic methods of performing a test.

6.1.1. Sensors

The sensors used to measure the mass accelerations are PCB model number 333B40 accelerometers. These sensors use a powered piezoelectric shear crystal and have a frequency range of 0.5-30000 Hz. These accelerometers are rated at ± 10 G and output approximately 500mV/G. The sensors were positioned so that they were exposed to as little transverse vibration as possible. These vibrations could cause a small non-linearity error in the measurement. Installation was a fairly easy task for installation on the sprung mass. The top of the mass was drilled and tapped to use a threaded stud to hold the accelerometer down. This face was perpendicular to the motion of the mass making it an ideal location for this measurement. The installed sensor is pictured in Figure 6-1.

Figure 6-1 Accelerometer Installed on Sprung Mass

The installation on the unsprung mass was a bit more difficult. The wheel and rotor were removed to expose the suspension upright. The top of the upright just over the wheel bearing was filed flat such that the surface was perpendicular to the motion of the suspension. This surface was then drilled and tapped and a threaded stud was installed to secure the other accelerometer to the upright. This was an ideal location because it allows the accelerometer to have the most sensitivity to the motion of the suspension while being as isolated as possible from the motion of the sprung mass. Figure 6-2 shows the installed sensor on the assembled suspension. Care was taken insure that the signal cables were protected from being pinched, kinked, vibrated, or rubbed by sharp or rough edges. The only other sensors used in this test setup were the LVDT and delta-P transducers used in feedback control of the servo-hydraulic actuation system. Access to these signals was possible using analog outputs on the servo controller.
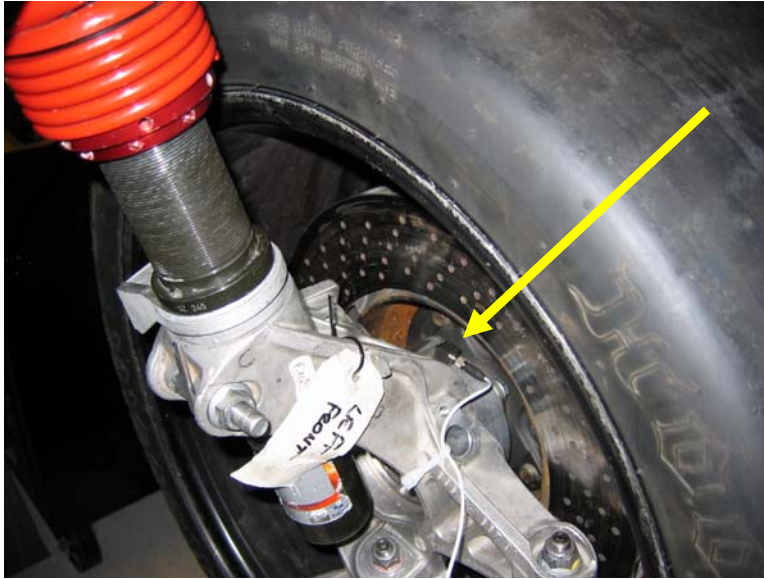
Figure 6-2 Accelerometer Installed on the Suspension Upright

6.1.2. Instrumentation

Several lab instruments were implemented during the experiment. The primary pieces of equipment were a data acquisition box, real-time control box, oscilloscope, lab computer, and servo-hydraulic controller. The data acquisition box is an IOtech WaveBook/516E. This is an expandable data acquisition unit. For this study an IOtech model number WBK18 8-channel high speed signal conditioning module was added. This expansion unit is capable of powering the accelerometers and has programmable signal conditioning such as low-pass filtering and AC coupling filtering built in. The main unit has an Ethernet connection for direct communication with the lab PC. An analog break-out box was built to allow the acquisition system to output the analog accelerometer signals to the control box. The data acquisition system is shown in Figure 6-3 with two of these expansion units.

Figure 6-3 IOtech Data Acquisition System

The real-time control software was run on dSPACE AutoBox see in Figure 6-4. This unit is a high speed real-time control prototyping box. The nice feature of this unit is the ability to communicate directly with MATLAB and Simulink on the lab PC. All of the software was written as Simulink models and then uploaded directly to the AutoBox via a special network adapter. The AutoBox has 16 analog inputs linked to high speed analog to digital converters and it also has 6 analog outputs that can output up to ± 10 V.

Figure 6-4 dSPACE AutoBox Control Prototyping Box

While running a test, all of the real-time software is running on board dSPACE. A PC user interface was designed using dSPACE Control Desk software. This software

allowed control of the experiment and collection of data with a series of windowed layouts. With this interface, the hydraulics could be controlled and the system ID and control algorithms could be run and monitored. Figure 6-5 is a general layout of the experiment.
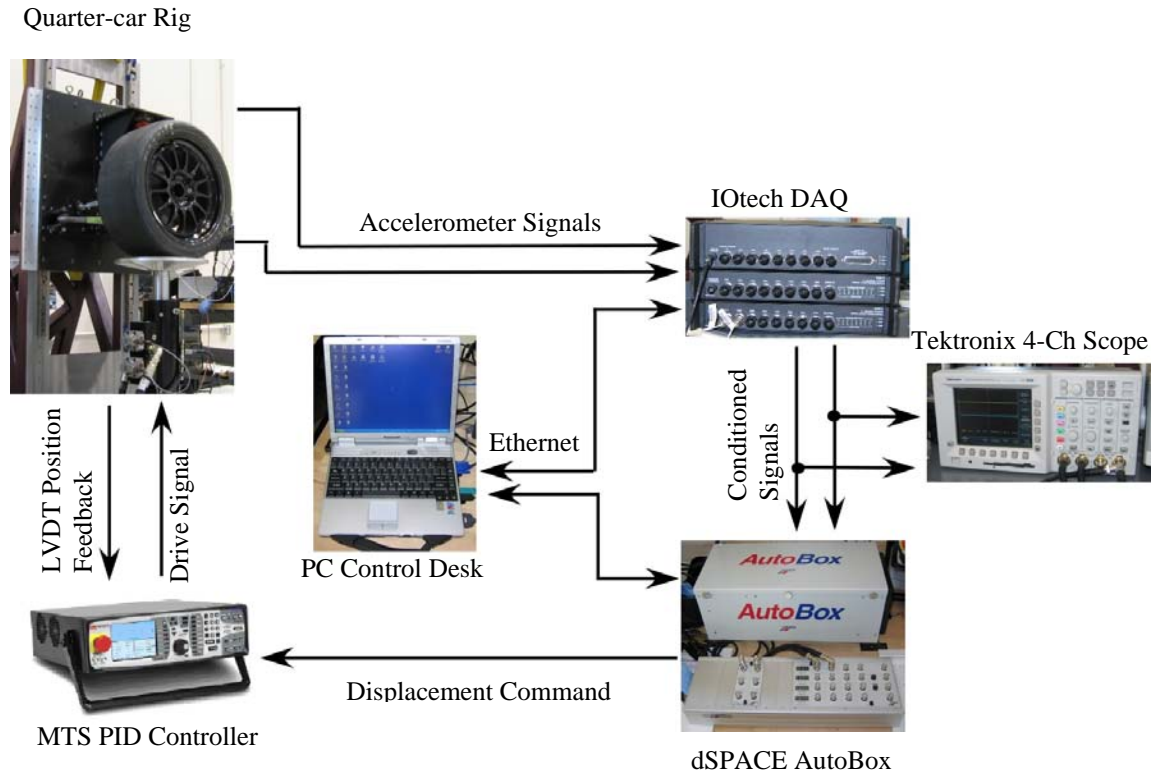

Figure 6-5 Instrumentation Layout

Generally, in an experiment the MTS controller was set to accept and follow an external displacement command. This PID controller's purpose was to insure that displacement command signals were followed by the actuator. The external command comes directly from dSPACE based on the software running at the time. If system ID is enabled, the excitation comes from an attenuated, filtered white noise signal generator. If adaptive inverse control is enabled, then the command signal for the actuator comes from the output of the adaptive control filter.

The IOtech system anti-aliases and low-pass filters the accelerometer signals before sending them to dSPACE. The IOtech is configured for the accelerometers with PC based software called EZ Analyst. The digitized, filtered signals may also be viewed on

the PC with this software. The oscilloscope was used to monitor the analog signals moving into dSPACE. It was found that the signals picked up a lot of noise from the unshielded break-out box. To correct this dSPACE was run at a high sample rate and the accelerometer signals were low-pass filtered in the digital domain. With the sample rate high enough this would effectively perform anti-aliasing with a conventional multi-rate approach.

6.2. Basic dSPACE Code

The Simulink model that was compiled for running in dSPACE was a slightly modified version of the code used in Chapter 5. The basic format of the code is similar to that introduced in Figure 5-12 in section 5.3. Only slight modifications were made on this primary structure to aid in some signal routing as there are small discrepancies when moving from the simulation environment into dSPACE. Figure 6-6 is block diagram of this modified code. This high level code in the program is basically the same setup as in Figure 5-12. In this instance the ID weights have been routed directly into the control algorithm rather than being stored for later use. Also there is a 'Goto' block labeled 'ID_state' which controls a switch inside the DAC block. This switch controls whether the system ID or control algorithm controls the actuator.
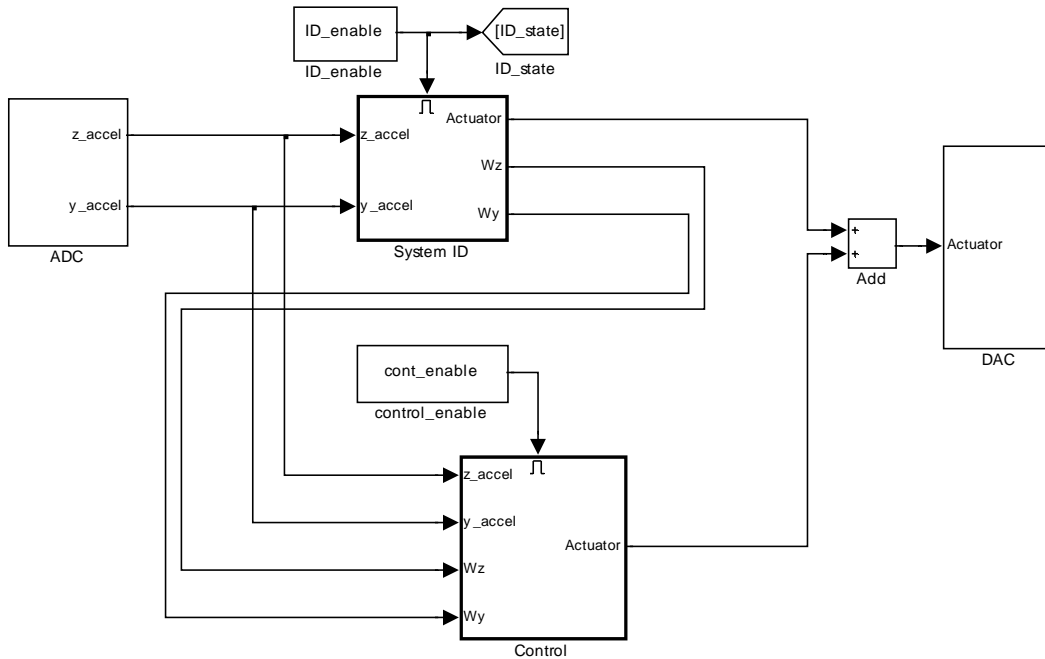
Figure 6-6 High Level Simulink Model Modified for dSPACE

The major difference between this code and that used in simulation is the replacement of the state-space quarter-car model with DAC and ADC dSPACE blocks and the physical quarter-car hardware. A detailed of view of the DAC and ADC blocks are represented with along with the connection to the hardware in Figure 6-7.
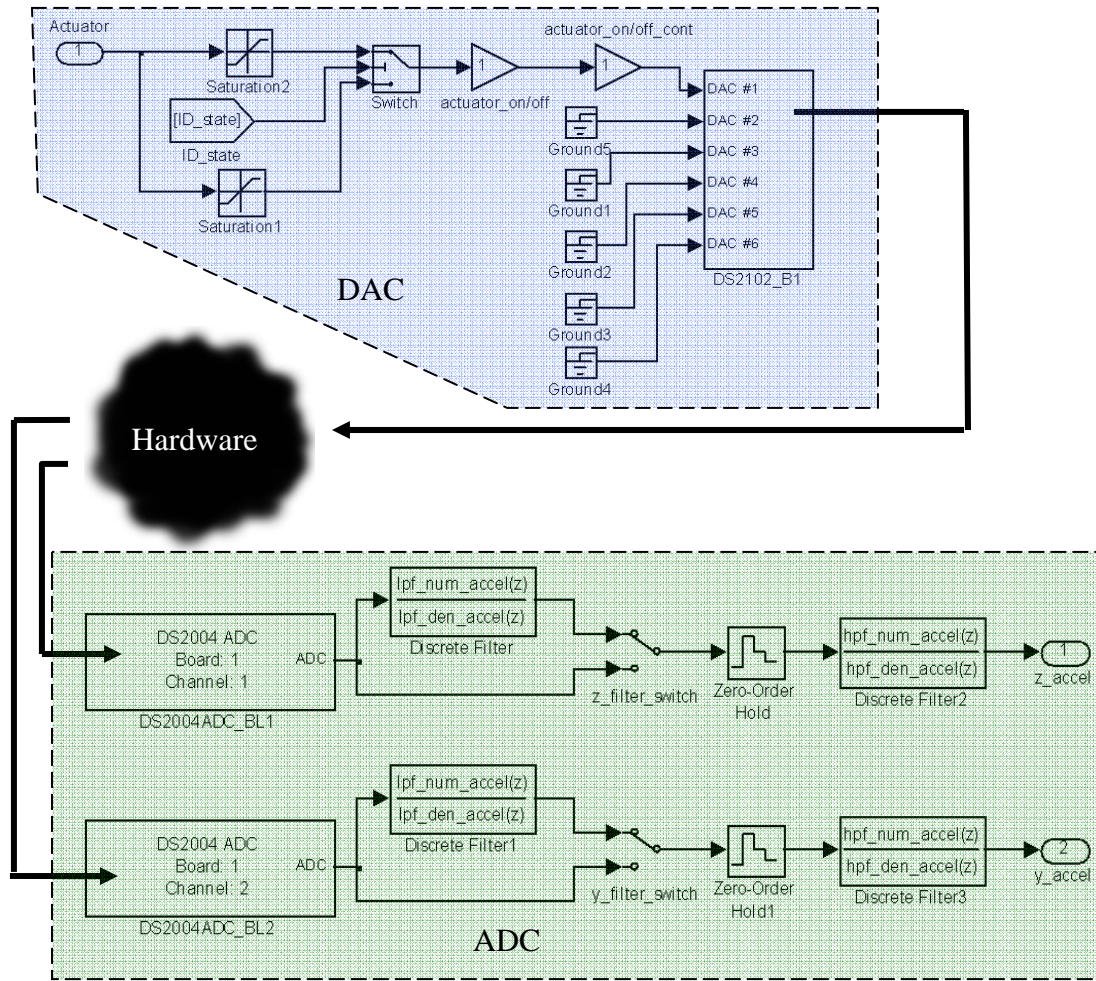
Figure 6-7 Detail of DAC and ADC and Connection to External Hardware

In this image the dSPACE specific blocks are the 'DS2101_B1', 'DS2004ADC_BL1', and 'DS2004ADC_BL2'. The DS2101 block is the DAC configuration block. Once compiled, this block tells dSPACE how to configure its analog outputs. Each analog input to dSPACE has its own DS2004 block. These blocks configure the ADC and tell the software where to route the signal to.

The DAC block contains the switch described earlier to control what part of the software is outputting a drive signal to the actuator. Also, to protect the hardware the saturation blocks were installed to limit the magnitude of the signal output. The ADC block contains two filters and a zero-order-hold. The ADC and low-pass filters run at a higher rate than the rest of the software. This helps filter line noise picked up between the WaveBook and dSPACE. The zero-order-hold block down-samples the data to the

appropriate 1000 Hz frequency. The filters on the right side of the ADC are high-pass filters. These work like AC coupling filters to remove any DC component in the signal picked up by the wiring or equipment.

6.3. System Identification

This section details the process of system identification in the hardware environment. The test performed was virtually the same as in simulation but now the quarter-car state-space model was replaced with the real hardware.

6.3.1. dSPACE Model

The Simulink model compiled for dSPACE was virtually the same as the code for the simulation environment seen in Figure 5-4. Figure 6-8 is the system ID enabled subsystem in Figure 6-6 expanded. Again two identical blocks were used for each of the input to output paths. All of the output blocks seen in Figure 5-4 are no longer necessary as the code is running on the stand-alone control box. Instead the data can be collected using capture functions in the dSPACE Control Desk Software. The product and low-pass filter blocks which are terminated are to compute the power level of the accelerometer and error signals. These power levels along with the regular signals can be viewed in and saved with Control Desk.
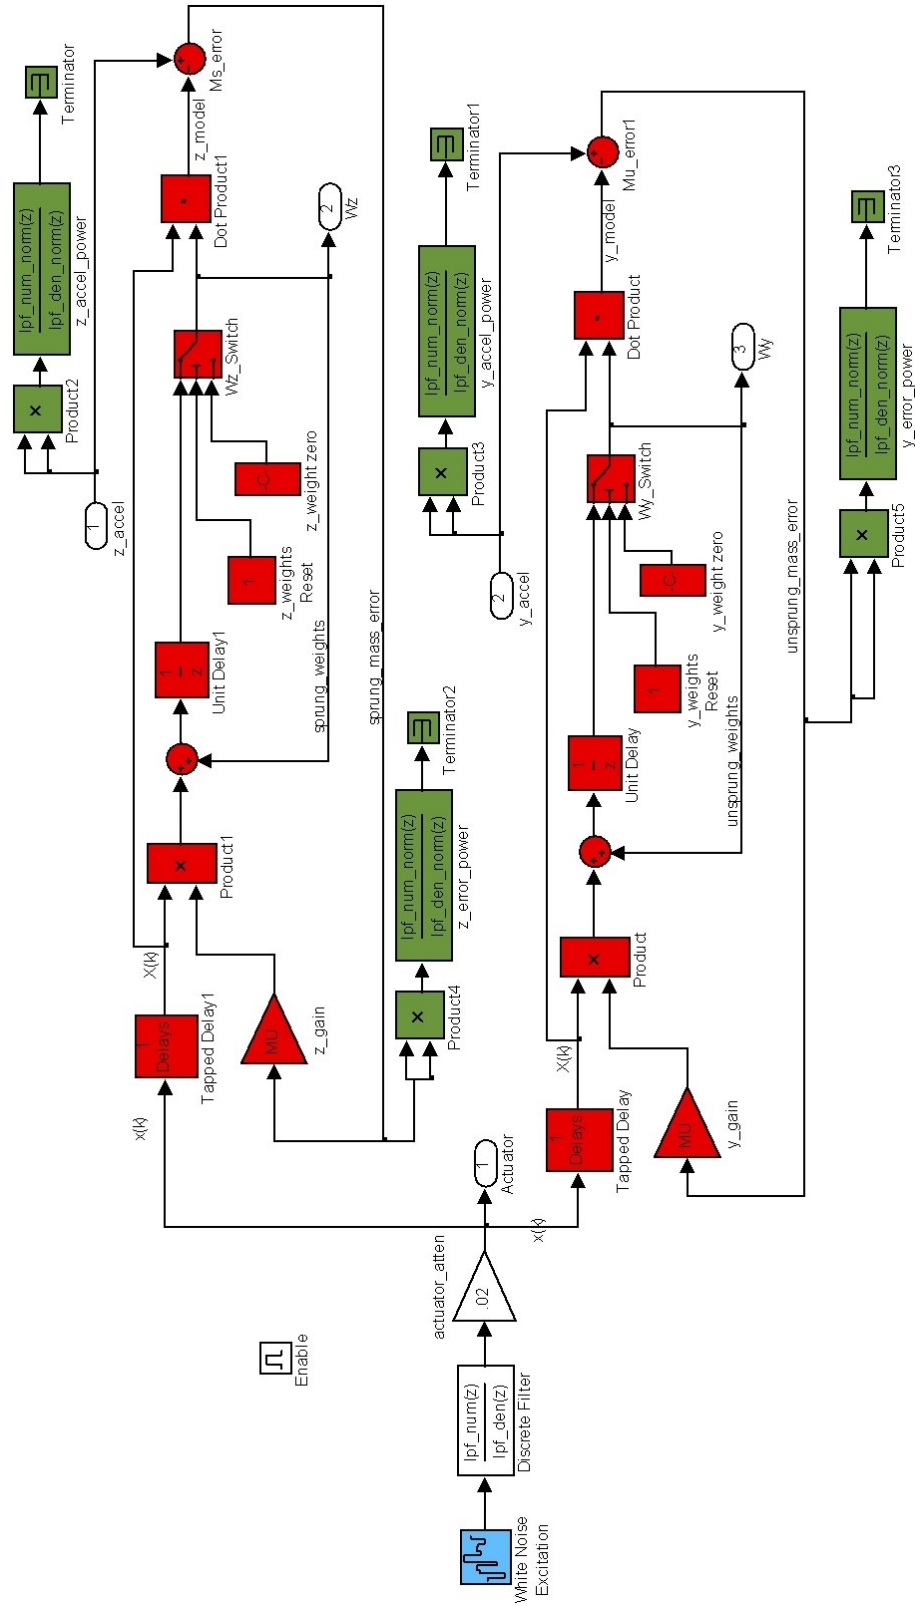
Figure 6-8 Detailed Simulink Model for ID

The LMS algorithm works to adapt the weights of the FIR filters the same way as in simulation. The step sizes 'z_gain' and 'y_gain' in this case can be changed form the front panel during simulation in Control Desk. The output block labeled 'Actuator' is where the signal is sent out to the DAC block. The addition of weight switches 'Wz' and 'Wy' allows for the ability to reset the weights while the real-time processor is running. This is to perform multiple tests without reloading the software or in the event that the weights become unstable during a test. Otherwise, the algorithm runs the same way to adapt the weights per equation (8).

For the system ID test, a front panel in Control Desk was created to have control of the tests. An images of two of the layouts from Control Desk are found in Figure 6-9 and Figure 6-10.
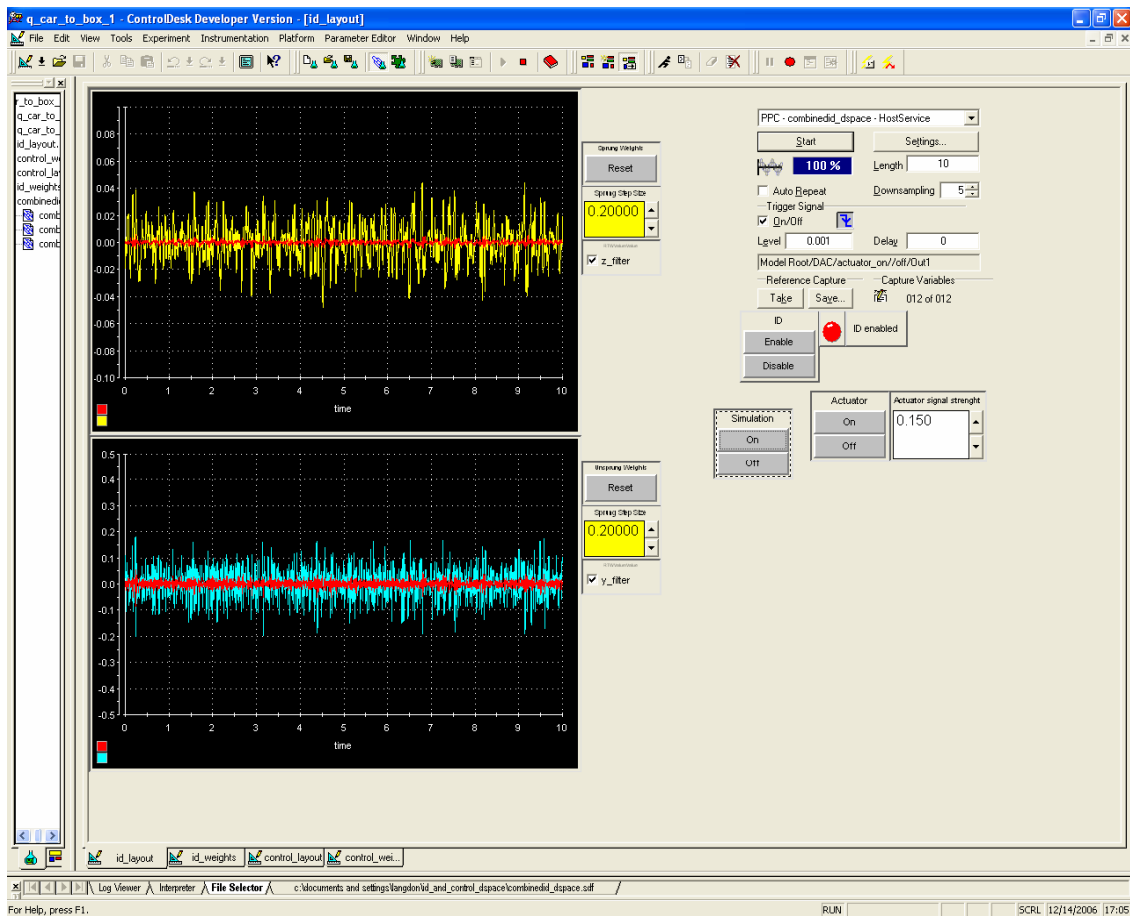

Figure 6-9 Control Desk Layout for System ID Accelerometer/Error Plots and Control Panel
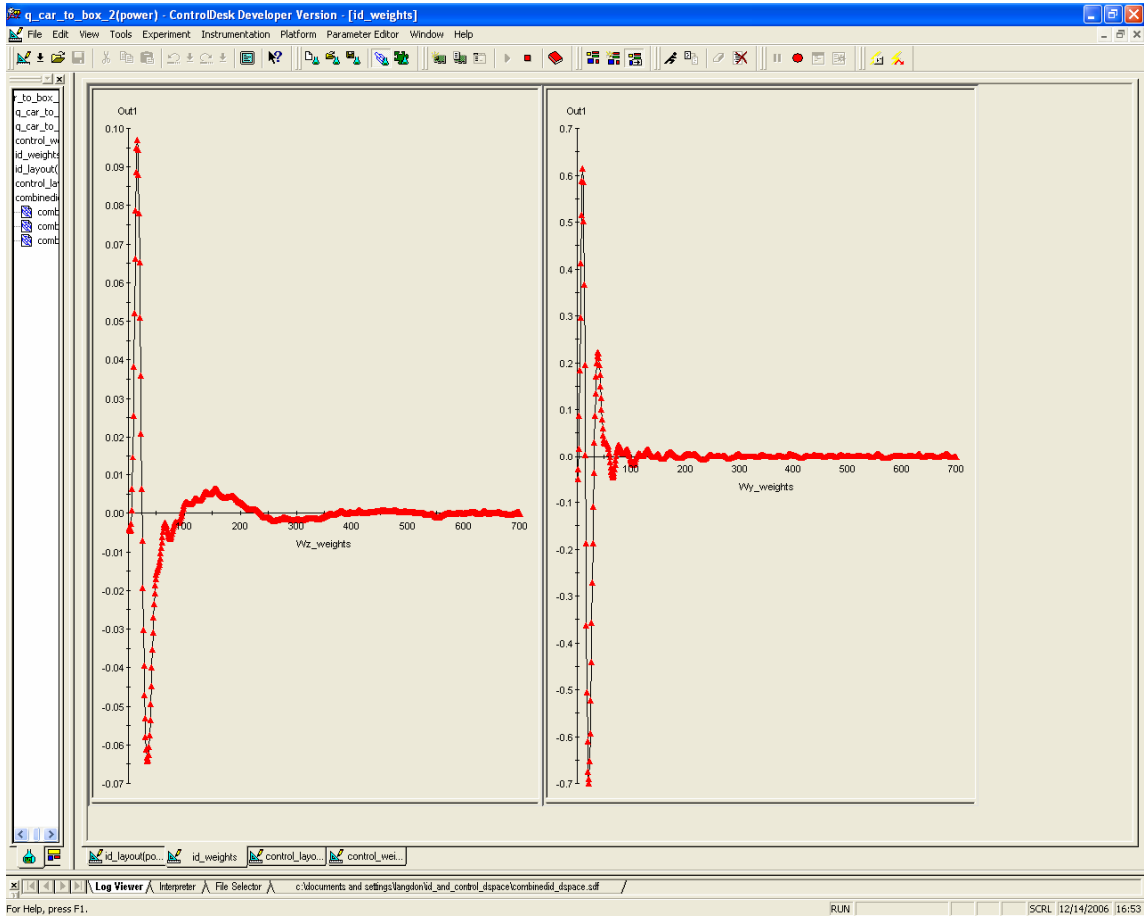
Figure 6-10 Control Desk Layout for System ID FIR Weight Vector Plots

These panels allowed for ease in knowing when the system ID weights converged. The weights continuously updated as they were adjusted. The control panel has the ability to adjust the step size and the amplitude of the excitation signal. It was also possible to toggle the low-pass filters on the accelerometer signals on and off as well as capture all of the signals. Also included, were buttons for enabling and disabling the ID subsystem as well as turning the actuator on and off. If the ID subsystem were disabled the actuator output was programmed to reset to zero and the FIR weights were programmed to hold their current values at the time. Thus, once the identification routine was completed output to the actuator form the ID subsystem would not interfere with that from the adaptive controller. Also, the converged weights when held could then be used in the adaptive inverse control subsystem.

91

6.3.2. Excitation Signal Shaping

Considerations similar to that in the simulation had to be made for the shaping of the excitation of the input signal for system identification. In addition to the desire to excite the system in the frequency range of interest care also must be taken not to exceed the limits of the actuator. For this test a band limited white noise excitation was filtered with two cascaded low-pass filters. A single resulting filter was constructed as the convolution of two Butterworth filters. One was a one-pole filter with a 4 Hz break frequency and the other was a four-pole with an 85 Hz break frequency. The frequency response of this shaping filter is shown in Figure 6-11. Clearly the filter does not pass much signal past 100 Hz. The magnitude crosses -20 dB at 40 Hz.
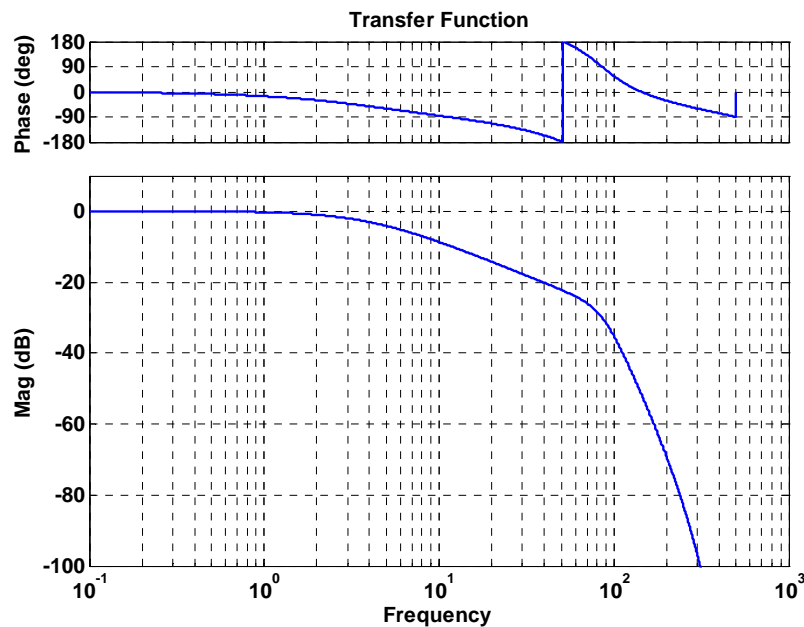
Figure 6-11 Frequency Response of White Noise Excitation Low-Pass Filter

6.3.3. Experimental Results

Several tests were run while trying to optimize the balance between fast convergence and stability. It was found that tuning the step sizes of both FIR models while running was an excellent way to achieve this. The step sizes were set relatively high during the early part of convergence. Later on the step size was reduced to help fine

tune the weights. This offered a great balance between quick convergence and nice results. The sizes of the FIR filters were also altered and tested several times to get the best results. Ultimately, the model size was set at 700 coefficients. The step size for each adaptive filter started at around 0.2 and then was fine tuned down to around 0.05.

Figure 6-12 and Figure 6-13 contain plots of the quarter-car accelerations and the error associated with how well the FIR models identified the plant.
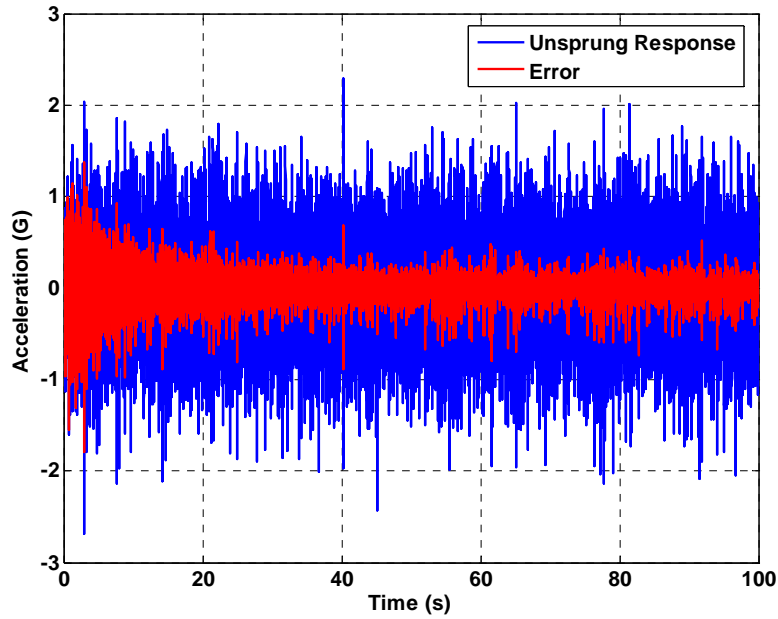

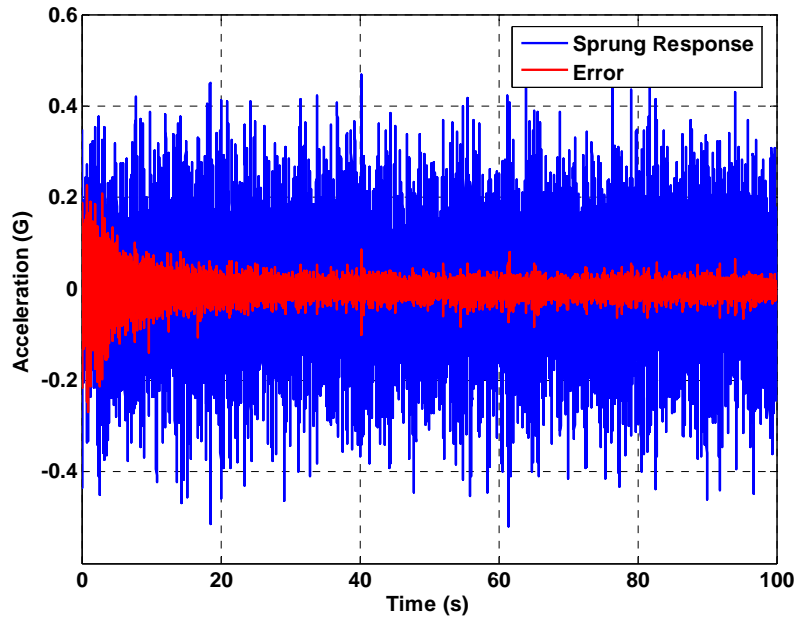**Figure 6-12** Unsprung Mass Acceleration Response and Model Error

Figure 6-13 Sprung Mass Acceleration Response and Model Error

The adaptive filters converged on their minimum very quickly just as they did in the simulation. As expected, they do not do as good of a job identifying the quarter-car plant as in simulation. The last 10,000 points of each signal were used to calculate an RMS value of each signal after convergence. These RMS values were used to calculate an error based on how small the model error was compared to the acceleration signals. The results are compared to the simulation results shown in Table 3 below.

Table 3 ID Error Comparison

Model Metric	Unsprung TF	Sprung TF
Simulation Error	2.30%	<1%
Experimental Error	16.90%	9.90%
dB Reduction (Sim)	~30 dB	~60 dB
dB Reduction (Exp)	~15 dB	~20 dB

The values in the table indicate that the FIR filters are much better at modeling the dynamics of the simulated model than they are at modeling the real physical system. This is not surprising because the simulated model was an ideal linear system. The

94

quarter-car rig has a lot of non-linearity associated with the suspension geometry, bearings, bushing and the like as well as additive measurement noise.

Figure 6-14 and Figure 6-15 are plots of the power levels of the mass accelerations and the modeling error. Much like in the simulation analysis, examining these power plots further show the quality of the model. The power levels of the signals were computed by squaring the respective signals and then low-filtering the squared signal with a fourth-order 0.1 Hz filter. These product and filter blocks are also seen in Figure 6-8. Similar to the simulation the decrease in power from the acceleration to the error signals is very good. Here the unsprung mass model error signal converges to around 15 dB lower than the accelerometer signal power. The sprung mass error signal has a slightly larger decrease in power of 20 dB. Both indicate that the model is doing a fairly accurate job of modeling the dynamics of the quarter-car rig.
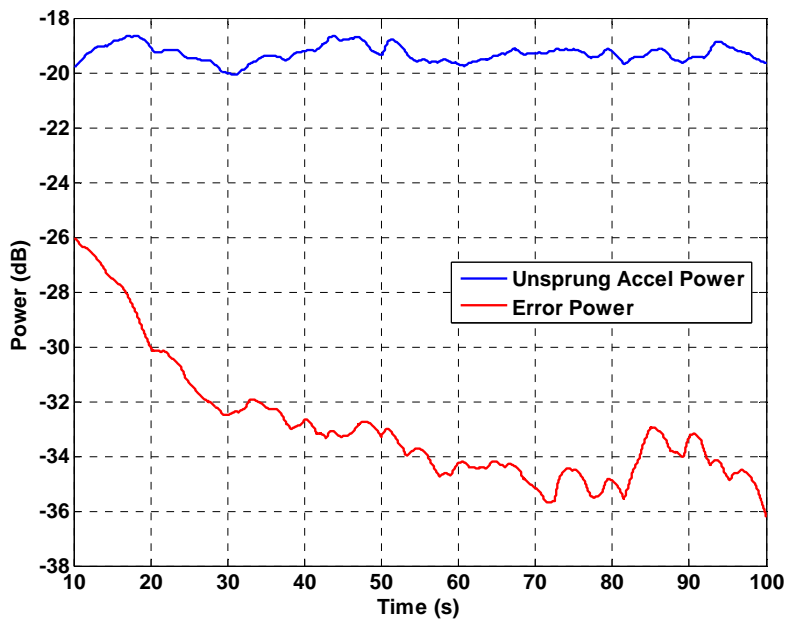

Figure 6-14 Unsprung Acceleration and Error Power
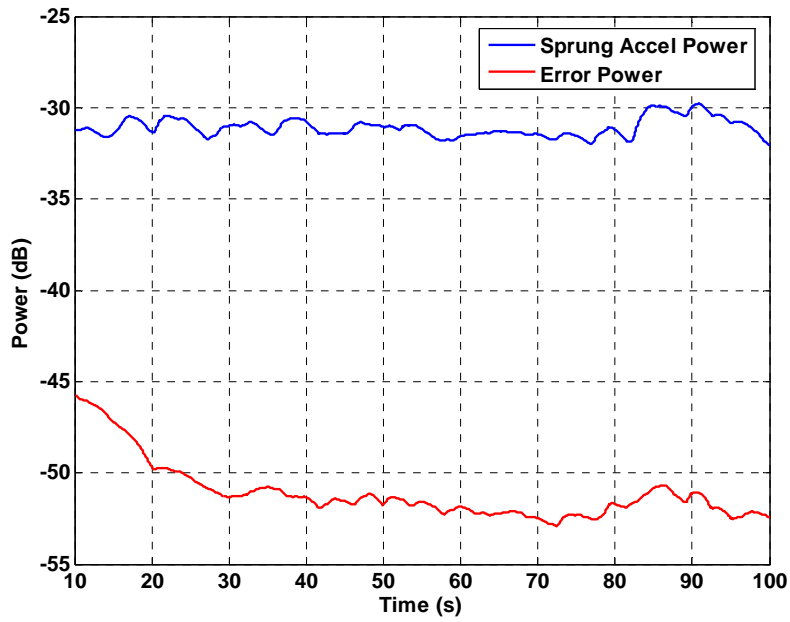
95

Figure 6-15 Sprung Acceleration and Power

Other interesting results include plots of the FIR model weights. These can be found in Figure 6-16. The weights take the shape of an impulse response of the modeled system.

Figure 6-16 Experimental FIR Model Weights
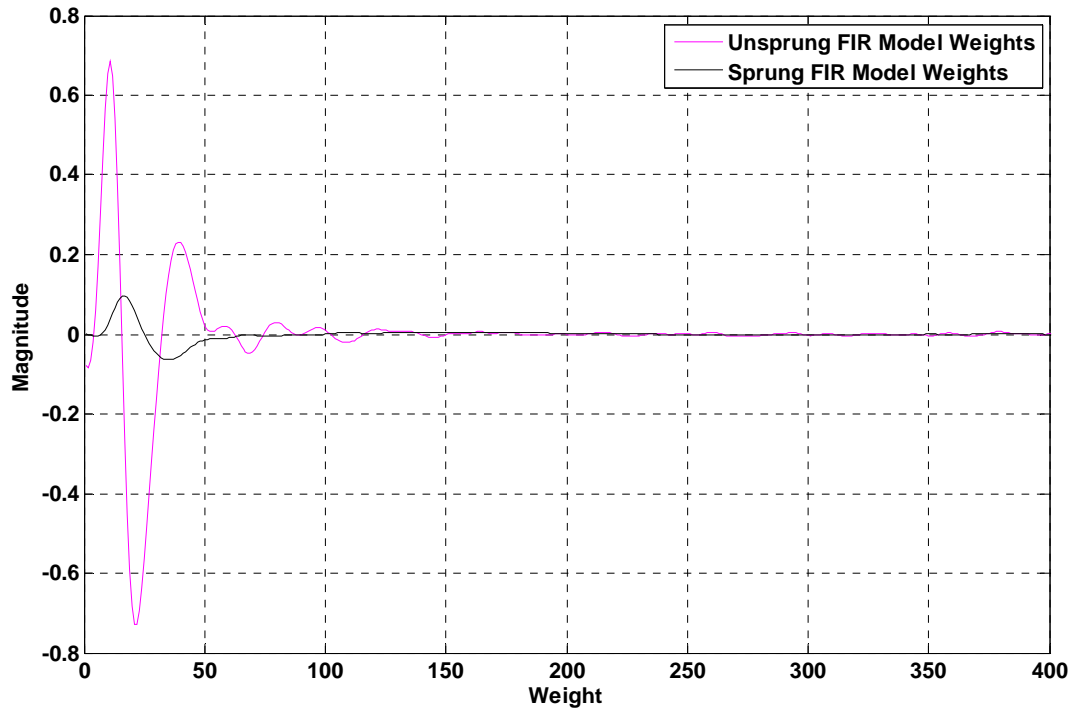
It is clear from the plot that there are at least two distinct dynamics which the FIR filters successfully identified. It is also clear that the FIR model of the unsprung mass transfer function has a much higher amplitude and frequency content than that of the sprung mass transfer function. Finally, the frequency response of the identified models is investigated in Figure 6-17.
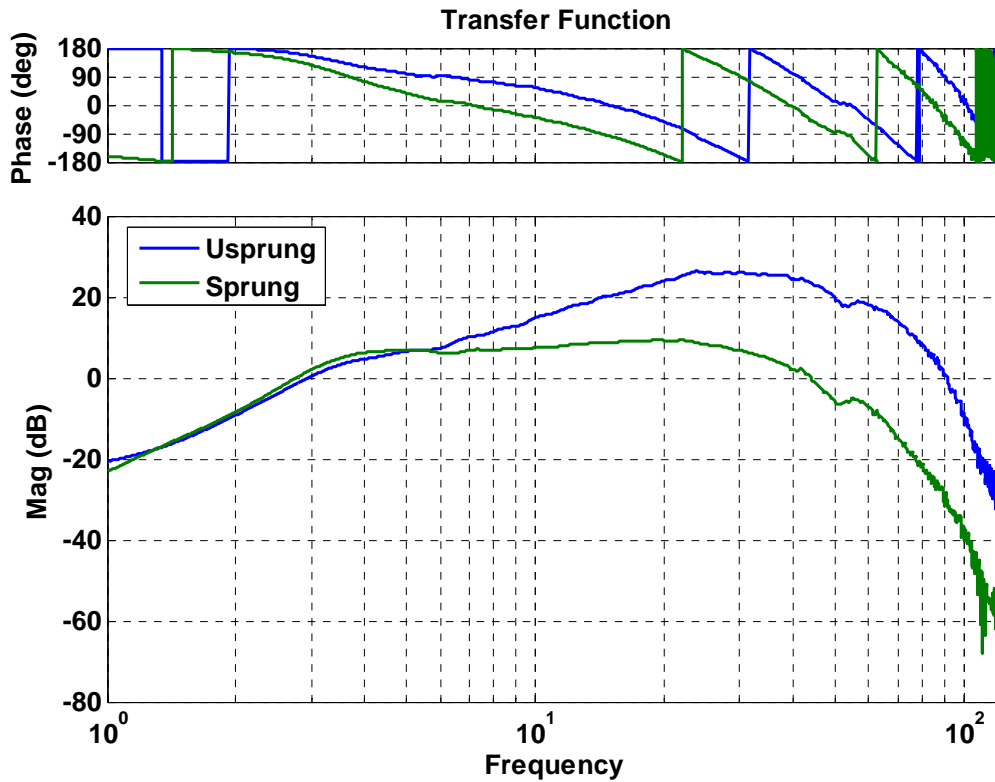
Figure 6-17 Frequency Response of the FIR Identified Model

The frequency response of the FIR model has two resonances, one near 5.2 Hz and one near 29 Hz. This closely agrees with the simulation model. This agreement with simulation is fairly reasonable to expect because the masses and spring rates used were well known when used in the state-space model. It is curious that the sprung mass is sensitive to a wide band of frequencies between the two resonance peaks. This is likely due to non-linear properties of the bearings, suspension bushing and damper and high damping in the system. The resonance peak at 29 Hz is more defined for the unsprung mass. There also appears to be a small spike near 55 Hz. This could be an un-modeled resonance within the suspension such as a second mode of the tire or wheel.

6.4. Adaptive Control

In this chapter the outer control loop is closed on the quarter-car rig to recreate acceleration data on the rig. This section discusses the software developed to close this loop in the real-time analysis. A discussion follows detailing the generation of desired response for the adaptive control tests. Finally, the section closes with a discussion of the experiment and results.

6.4.1. dSPACE Model

The Simulink model compiled for the closed loop control was very similar to that used in the simulation. Figure 6-18 is the expanded 'Control' block from Figure 6-6 detailing the control algorithm used. Like the simulation program, the 'to workspace' blocks have been removed as all signals were access directly by Control Desk. The addition of the product and discrete filter blocks were used to calculate the power of the signals for later viewing. These blocks were set up the same way as those in the ID software. The desired signals are loaded into dSPACE as a set when the code is compiled. Again the usual tuning parameters were accessible. The step sizes could be altered while the real-time test was running to fine tune the convergence rate and maintain stability. The number of FIR filter weights and desired response delay had be set up front and recompiled each time. Like the control algorithm in simulation only one of the desired responses were fed into the inverse controller. The other desired response was used to help adapt the filter weights only.
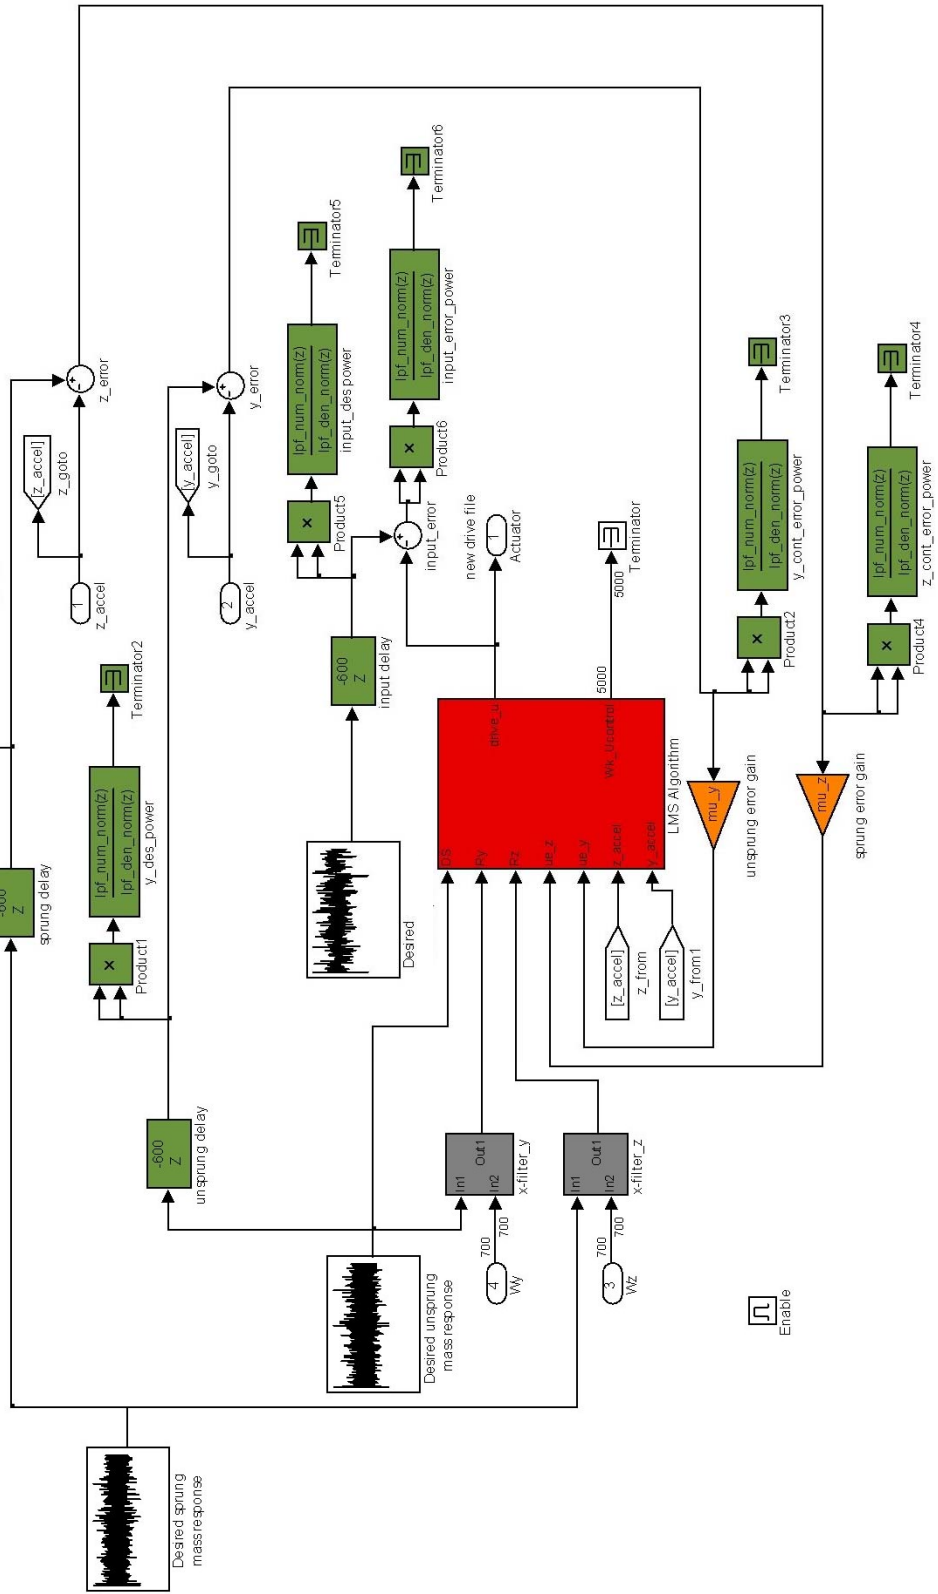
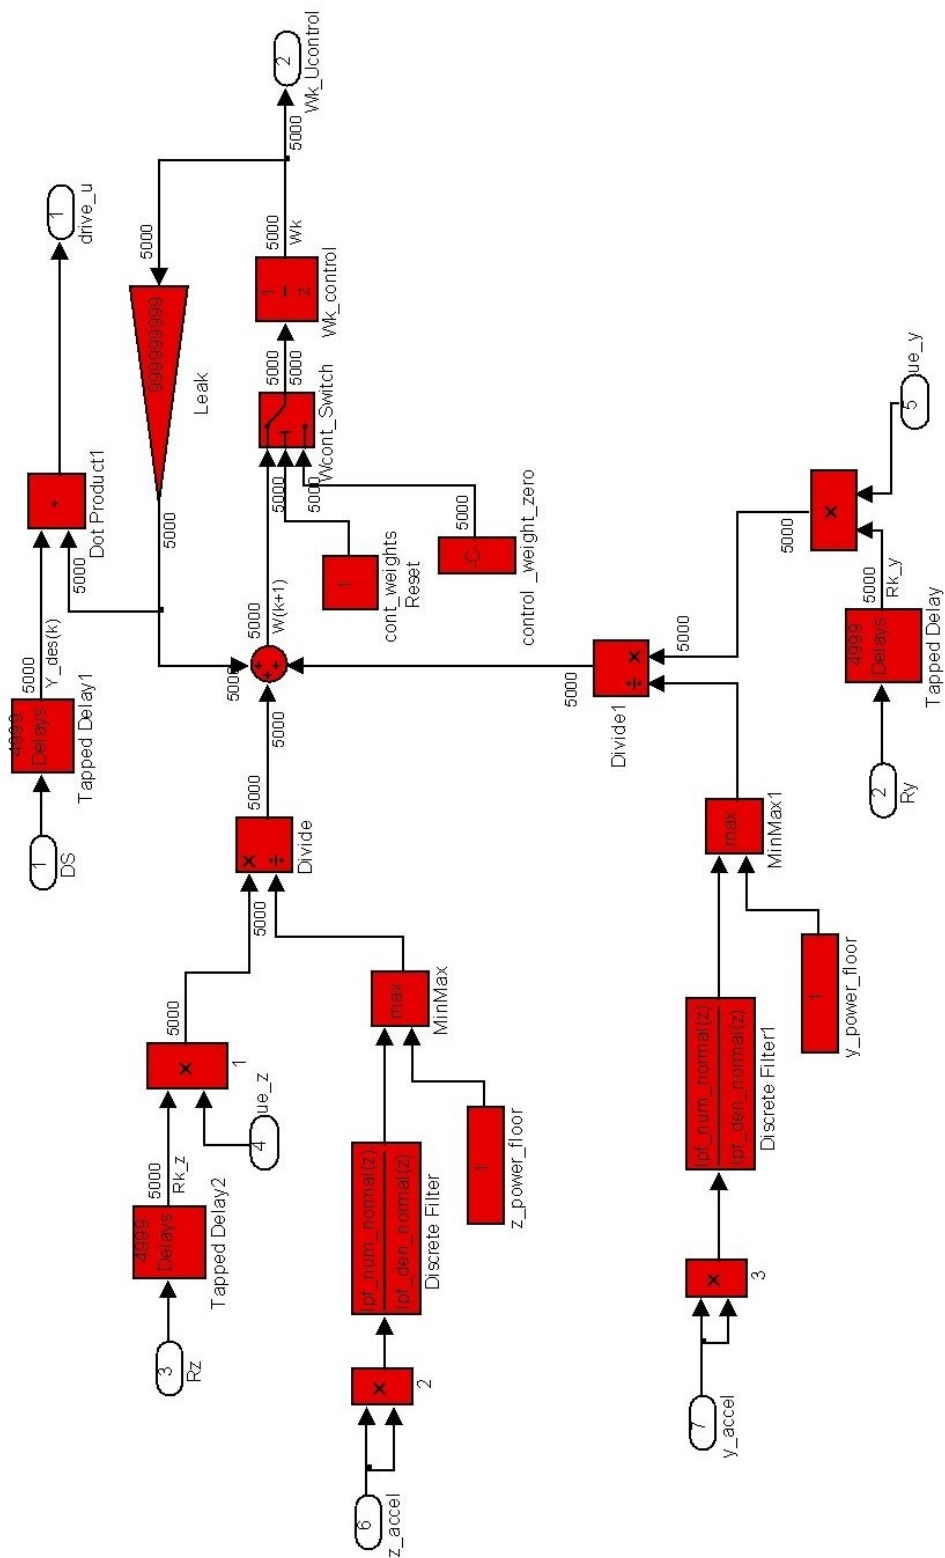**Figure 6-18** Real-time Control Simulink Block Diagram

Figure 6-19 Details of the Real-time Filtered-X LMS Algorithm

The 'LMS Algorithm' block in Figure 6-18 is further expanded in Figure 6-19. This code has virtually the same functionality as that used in simulation. The only alteration is the addition of a switch at the output of the summing junction for the weight vector. This switch is included to be able to reset the weights to zeros should the system become unstable or run several test using the same identified plant models.

Control Desk panels were developed to command the control software. Screenshots of the two layouts are shown in Figure 6-20 and Figure 6-21. This software was used to enable and disable the control algorithm and tune the step size for each error path. This software could also reset the control filter to zero coefficients, change the leak in the algorithm, and graphically monitor the convergence of the error. A second layout was used to watch the weights move and help determine when they have converged.
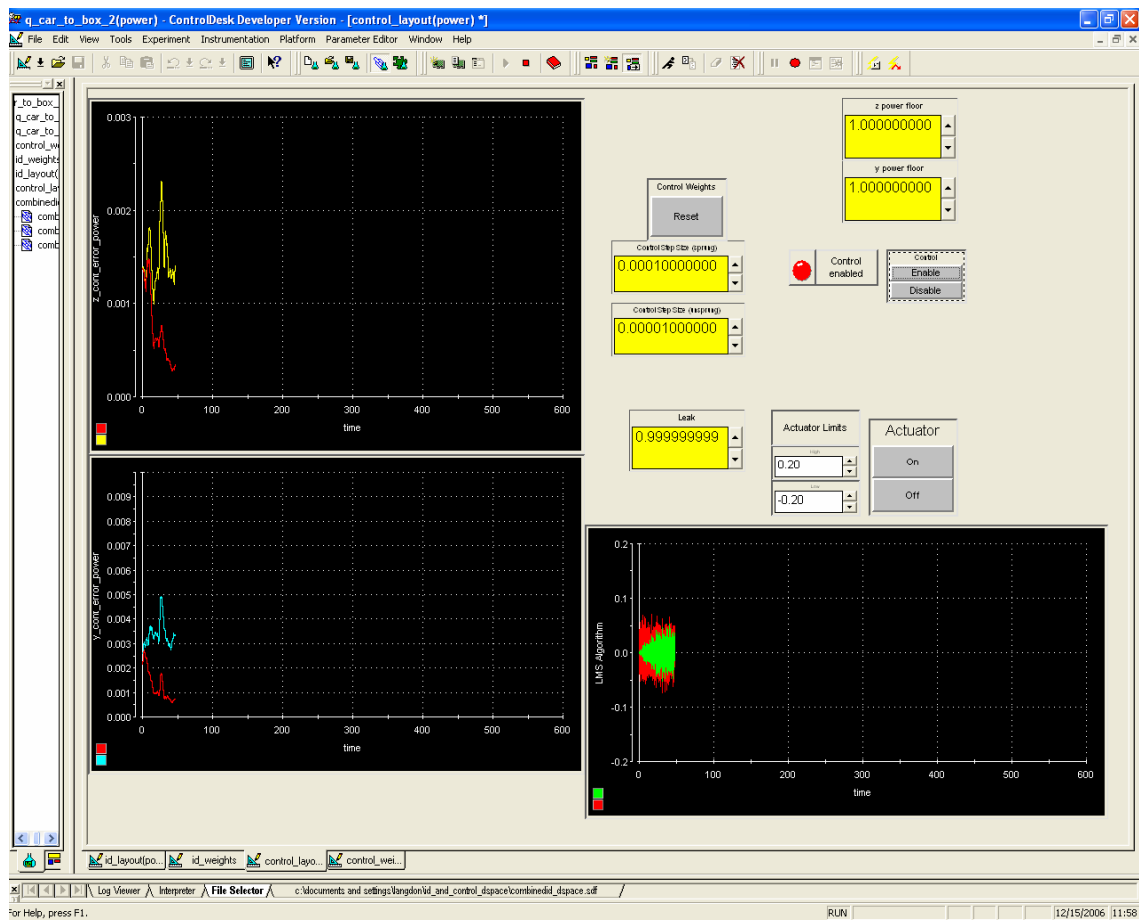

Figure 6-20 Screenshot of Control Desk Real-time Control Interface
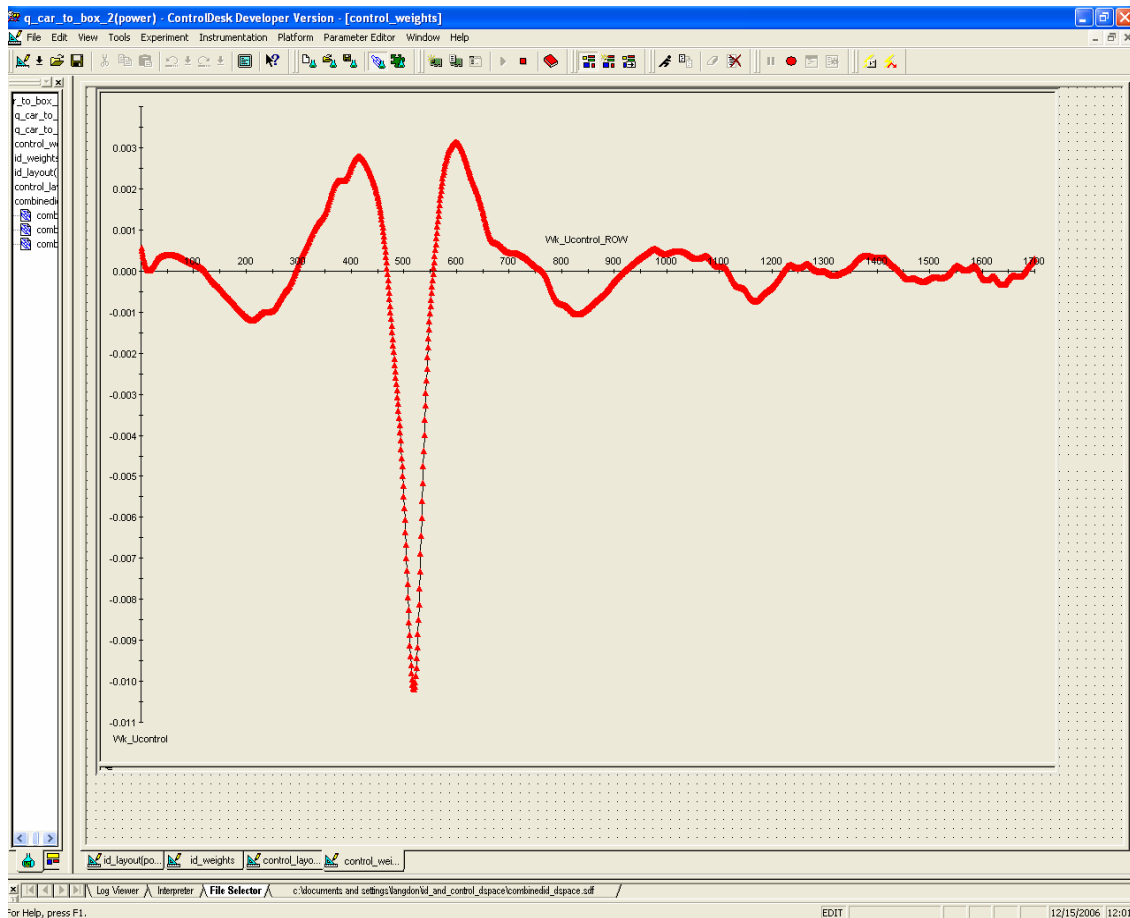
Figure 6-21 Screenshot of Control Desk Real-time Control Weights

6.4.2. Desired Response Generation

For the control experiment, a desired response was first needed for reproduction. The first step in proving this concept was to collect data off of the quarter-car rig itself. Obviously, if data from the rig could not be reproduced then the problem would be even more difficult to bring track data in for replication using this algorithm. To collect data a filtered white noise excitation was input to the rig. This noise was filtered with a discrete filter convolved from two low-pass filters. The two filters were a single pole, 4 Hz filter and a four-pole 50 Hz filter. The filtered noise was played into the rig and the accelerometer signals were recorded directly with the EZ Analyst software running the IOtech system at 1000 Hz. The signals were anti-aliased by the data acquisition system. The same signals were filtered and recorded two different times. One data set was

103

created by filtering the data with a 15 Hz low pass filter. The second data set was created by low-pass filtering the accelerometer data with a 30 Hz filter. Two control experiments were performed to replicate each set of data to see how the excitation bandwidth would affect the quality of convergence on a solution. The desired response data size was recorded for 60 s.

6.4.3. Experimental Results – 15 Hz Data

Reproduction of each data set was attempted several times to learn the best settings for the adaptive filter coefficients and step sizes. Each time the number of filter weights or delay size was changed the software had to be recompiled and uploaded to dSPACE. The controller step sizes could be changed during the test via Control Desk. For all tests the unsprung mass response was fed directly to the controller as the desired response. Thus, the adaptive inverse controller was ideally trying to represent the inverse of the path from the tire input to the unsprung mass acceleration.

The first set of tests was performed to recreate the 15 Hz filtered data. After several tests, the controller size was set to 1700 weights and the sample delay was set to 550 samples. For the low frequency data replication the step size for the sprung mass error was set to 0.0001 and the step size for the unsprung mass was set to 0.00001. Figure 6-22 and Figure 6-23 show the converging error overlaid on the desire response for both masses.
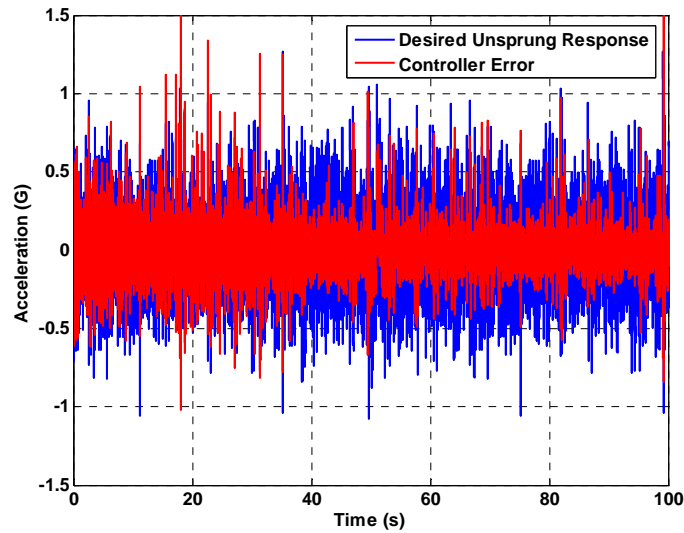
Figure 6-22 15 Hz Filtered Unsprung Mass Desired Response Controller Error
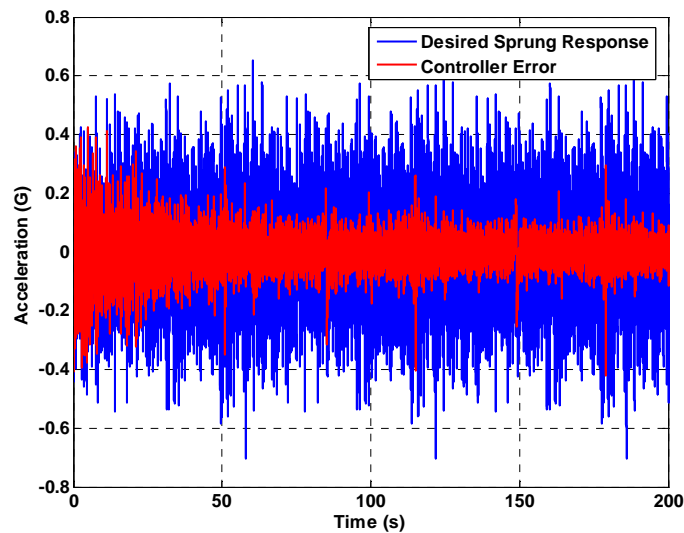
Figure 6-23 15 Hz Filtered Sprung Mass Desired Response Controller Error

It is noted that the error convergence of the unsprung mass error does not appear to be very good, particularly at the time scale of the plot. Further examination of the error plot in Figure 6-22 shows why. Figure 6-24 is a one second detailed view of a converged part of the unsprung mass error plot. It is clear in this plot that the high magnitude peaks have been captured the best. These higher magnitude accelerations are associated with lower frequencies. Because this data set was collected with a 15 Hz filter most of the dynamics

105

of the unsprung mass were not really captured. Thus the controller has a much harder time replicating any dynamics at the higher frequency. Most of the error in this plot appears to be associated with higher frequency excitation.
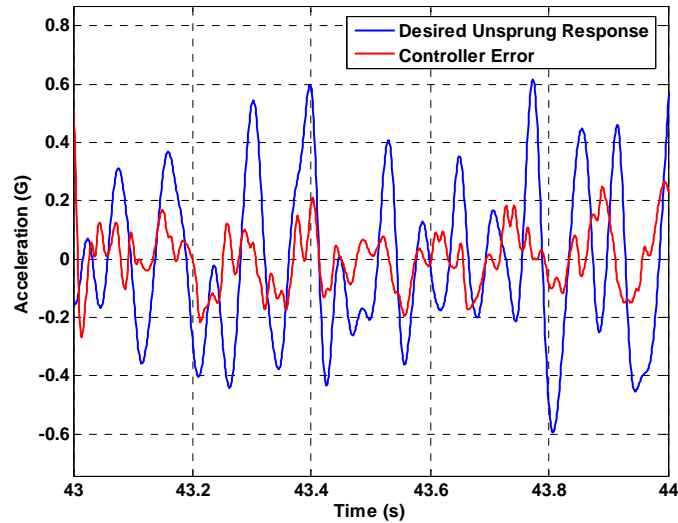
Figure 6-24 Detail of Unsprung Response Controller Error

Perhaps the better way to look at the error convergence is based on power levels of the signals. Figure 6-25 and Figure 6-26 are the power levels of the desire signals and their respective error signals. These plots tell more about how well the controller is actually replicating the signal. The power level drop of the unsprung mass error signal from the desired response is nearly 10 dB. Considering that the large amplitude, low frequency response seemed to be replicated with more accuracy than the higher frequency, low amplitude signals, indicates that there is a lot of signal power associated with the higher amplitude response.
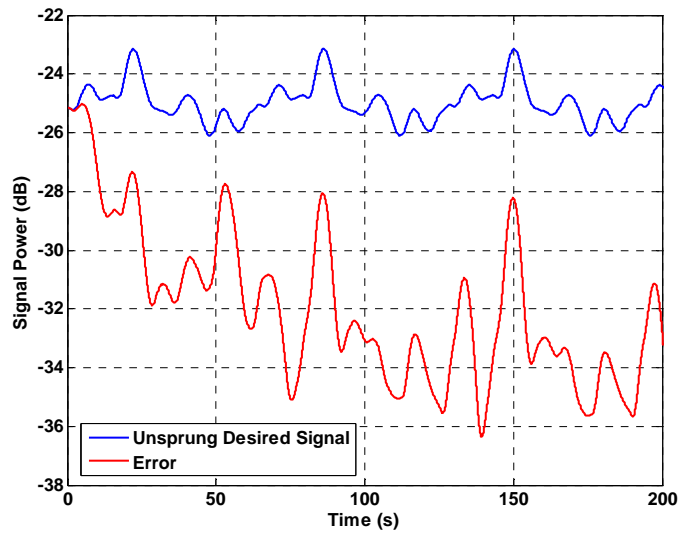
Figure 6-25 15 Hz Filtered Unsprung Desired Response and Error Signal Powers

Figure 6-26 15 Hz Filtered Sprung Desired Response and Error Signal Powers

The RMS values of the error and desired signals were computed over a converged section of data. An error was calculated from these RMS values in a similar fashion as the identifi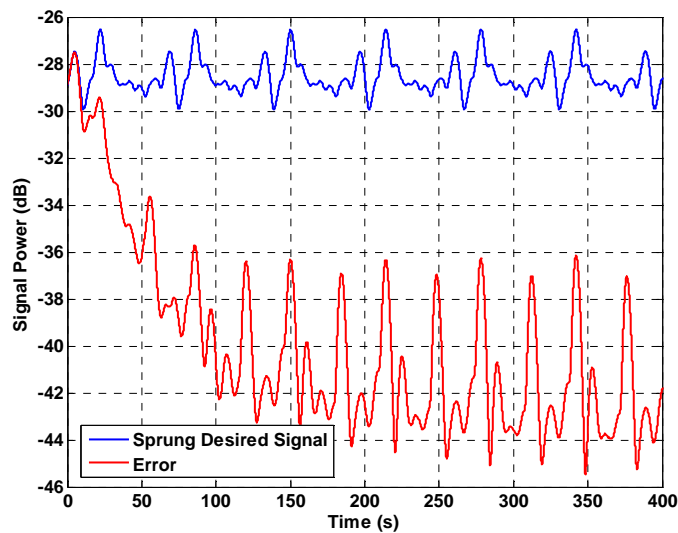cation error. This error along with the power level reduction is compared to the results from the simulation analysis in Table 4.

Table 4 15Hz Replication Controller Error Comparison

Model Metric	Unsprung TF	Sprung TF
Simulation Error	8.40%	7.50%
Experimental Error	36.20%	23.40%
dB Reduction (Sim)	~22 dB	~23 dB
dB Reduction (Exp)	~9 dB	~13 dB

Figure 6-27 is a plot of the converged weights of the inverse adaptive controller for the 15 Hz reproduction data. These weights represent the inverse impulse response of the quarter-car test rig. It is clear in the plot that there is a lot of low frequency content. This is expected as the controller was only trying to replicate the inverse of a plant that was creating 15 Hz and lower responses.
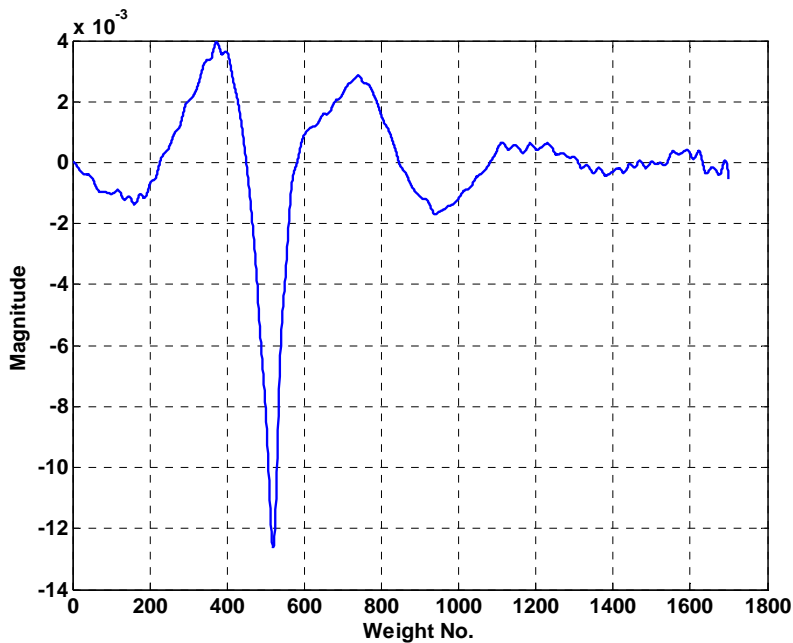

Figure 6-27 FIR Adaptive Control Filter Weight Coefficients for 15 Hz Data

Because the unsprung mass acceleration was the desired response being input to the controller, this plot represents the inverse of the transfer function between the road input and the unsprung mass acceleration.

To better understand how well the filter performed this task, the FIR controller was convolved with the identified model of the unsprung mass transfer function. A frequency response of this new convolved filter is plotted in Figure 6-28. Again the frequency response of the 550 sample delay was divided into the response of the convolved filters. This removes the phase warping caused by the delay.
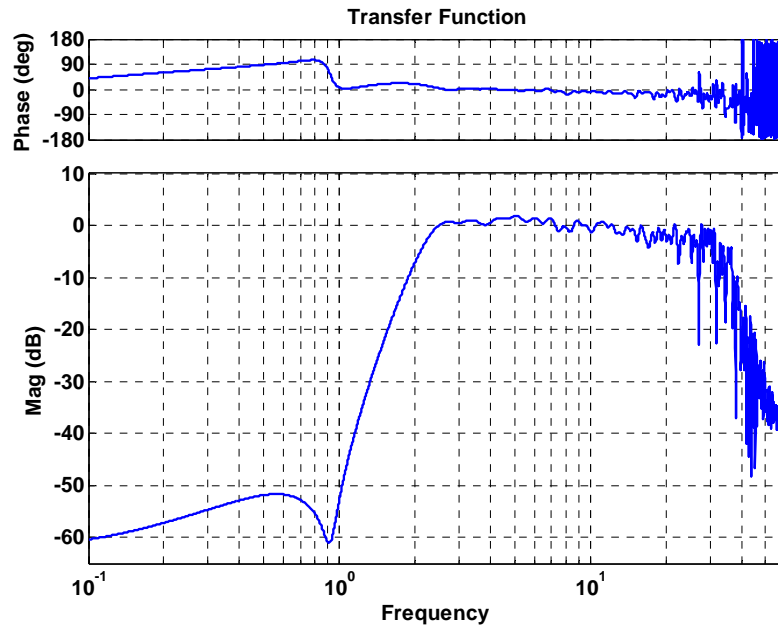
Figure 6-28 Frequency Response of FIR Controller Convolved with Plant for 15 Hz Unsprung Mass Response Data

The frequency response shows that the controller does a good job as an inverse of the plant. The magnitude is 0 dB or feed through between 2.5 Hz and starts to roll off after 15 Hz and the phase is basically zero in this range as well. Over this range, the controller does a great job of inverting the dynamics of the quarter car rig. Finally, Figure 6-29 shows a 1 s sample of the actual response of both the sprung and unsprung mass accelerations overlaid on their respective desired responses. This plot shows the accuracy of the inverse controller on a detailed response scale.
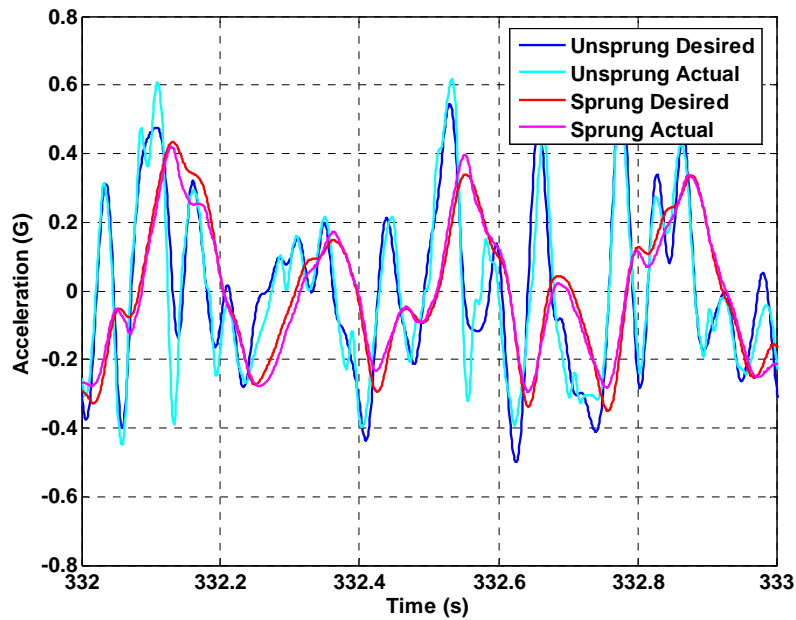
Figure 6-29 Detail of Actual Response Over Desired Response

6.4.4. Experimental Results – 30 Hz Data

The same tests were performed to replicate acceleration data that was filtered with a 30 Hz roll-off filter. Ideally, the extra frequency content would allow the controller to do a much better job at replicating the unsprung mass response, which again was the input to the filter. The number of filter weights and sample delay size was set to be the same as the 15 Hz replication. However, because of the higher signal power levels due to the higher frequency content, the step sizes needed to be lowered to 0.00001 for the sprung mass error and 0.000001 for the unsprung mass error. Figure 6-30 and Figure 6-31 show the error and desired responses for each path. These plots are somewhat surprising. The error still does not appear to converge well for the unsprung mass, although it does converge very fast. Of course the scale does much to distort the look of the error. The fact that it still has a fair amount of error may be attributed to non-linearity or other dynamics not in the frequency range of interest. The sprung mass error converges relatively slowly but do seem to do a nice job once converged.

110
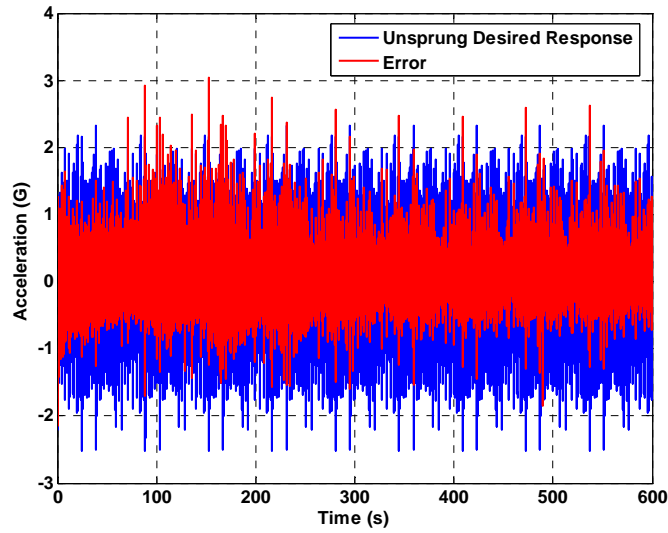
Figure 6-30 30 Hz Filtered Unsprung Mass Desired Response Controller Error
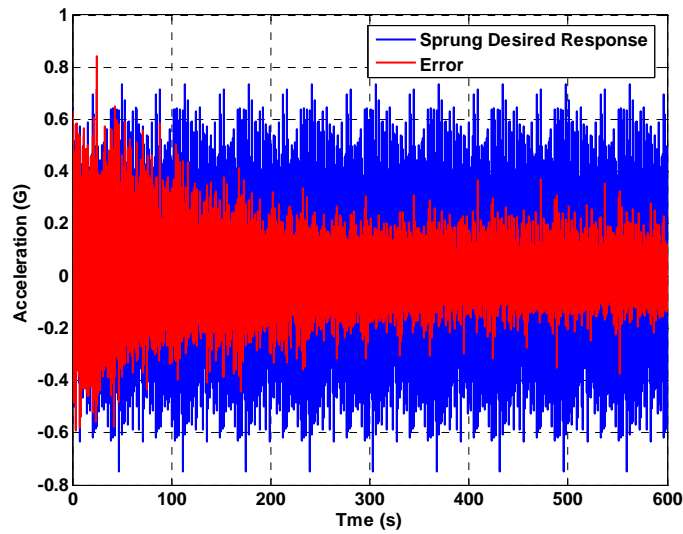
Figure 6-31 30 Hz Filtered Sprung Mass Desired Response Controller Error

Again, to get a better understanding of the unsprung mass response a zoomed-in look at the desired response and error is found in Figure 6-32. Here it is clear that the controller is still capturing the majority of the high power signals.
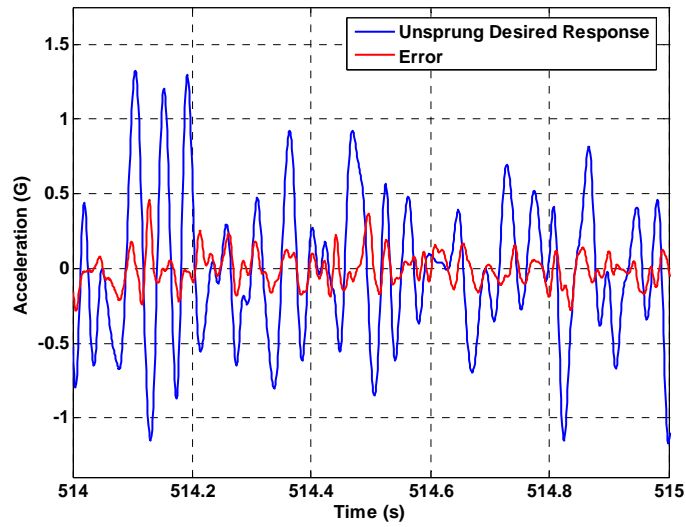
Figure 6-32 Detail of 30 Hz Unsprung Mass Response and Error

Furthermore, the quality of the error convergence was again gauged by the signal power levels are plotted in Figure 6-33 and Figure 6-34.
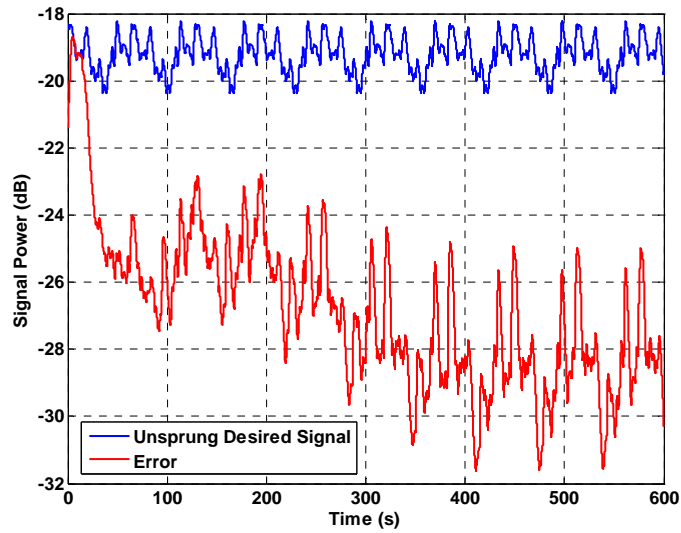
Figure 6-33 30 Hz Filtered Unsprung Desired Response and Error Signal Powers
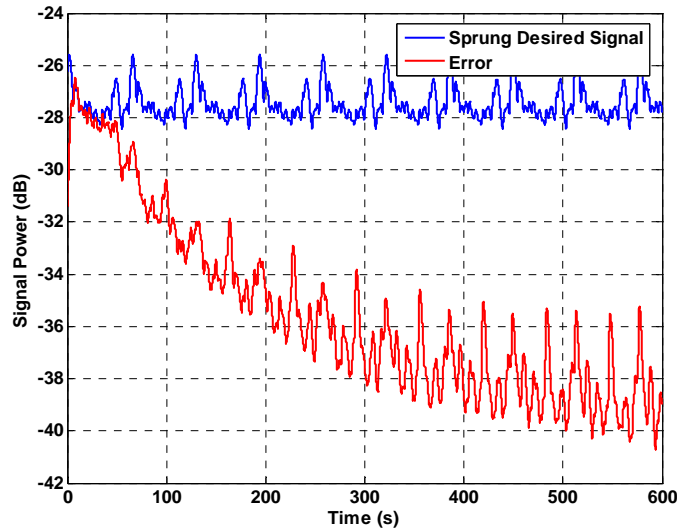
Figure 6-34 30 Hz Filtered Sprung Desired Response and Error Signal Powers

Table 5 is assembled below comparing these error and power level results to the control simulation results. Again, the RMS values of the response and error signals were used to form the error metric for convergence and the decibel levels are directly from the above plots.

Table 5 30Hz Replication Controller Error Comparison

Model Metric	Unsprung TF	Sprung TF
Simulation Error	8.40%	7.50%
Experimental Error	34.40%	31.10%
dB Reduction (Sim)	~22 dB	~23 dB
dB Reduction (Exp)	~10 dB	~12 dB

The error of the controller in replicating the unsprung mass dropped slightly compared to the 15 Hz data however it ability to replicate the inverse of the sprung mass path was degraded. The power level drops are still considered very acceptable.

The inverse controller weight coefficients are plotted in Figure 6-35. These represent the inverse impulse response dynamics of the unsprung mass to road input of the quarter-car. These weights do not appear to have much more high frequency content

113

than the controller weights from the 15 Hz experiment. This may help explain why the error for the unsprung mass path is not improved that much.
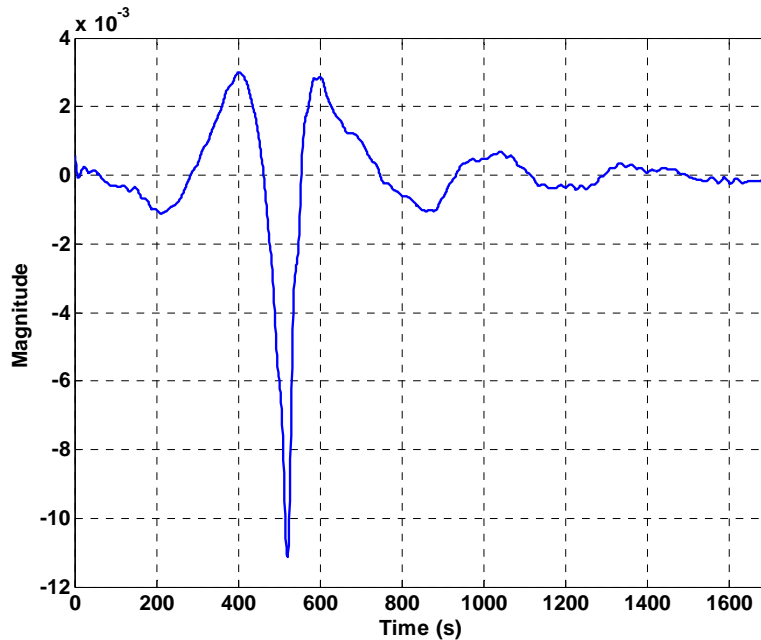
Figure 6-35 30 Hz FIR Controller Weight Coefficients

These controller weights were convolved with the unsprung mass identified model. The frequency response of these convolved weights is found in Figure 6-36. This plot explains why the error convergence and power level drop is not as good as was expected. The inverse FIR model does not create a feed through transfer function over enough bandwidth. The convolved filters only have a unity transfer between 2 ~ 8 Hz. This is only enough to do a fair job for the filter to invert the sprung mass dynamics. However, the filter appears to just not invert the higher frequency dynamics of the unsprung mass. The fact that the error stated is still in the mid 30 % range, considering how poorly the convolved transfer function looks, shows the robustness of these inverse adaptive filters.
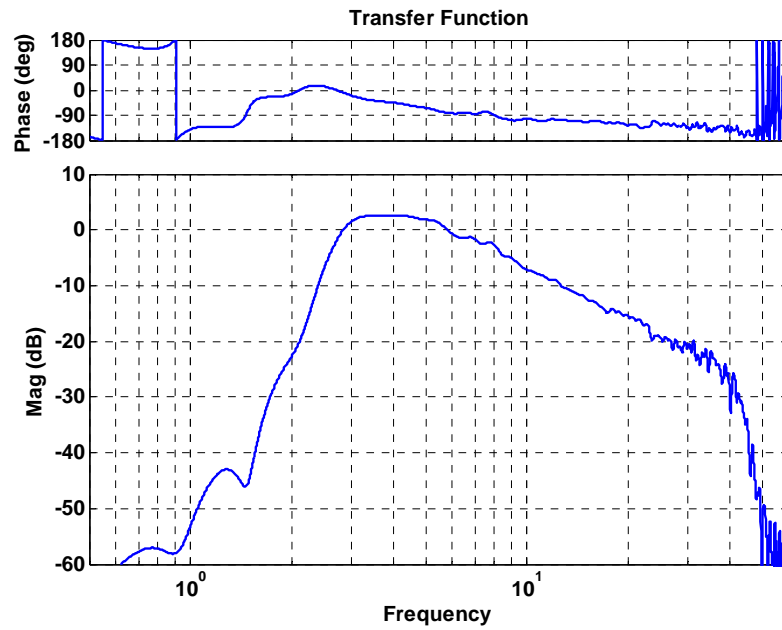
Figure 6-36 Frequency Response of FIR Controller Convolved with Plant for 30 Hz

Unsprung Mass Response Data

7. CONCLUSIONS AND RECOMMENDATIONS

The concept of this new quarter-car test bed proved to be a very flexible and robust design for the purposes of this and future experiments. Future plans set forth for the rig also prove that its design will lend its use for many studies to come with varying purposes. The rig's stability provided a great deal of confidence in the results obtained during the tests in this study. The design goals for the rig in the scope of this study were achieved in their entirety and future phase developments of the rig look to be performed in a very seamless manner.

The control algorithm implemented proved to be a different adaptation from current practices. Though the RMS errors quoted in this study appear to be slightly higher in some tests, the proof of concept of applying these control methods for replication of data has been completed successfully. A careful explanation of the concepts of system identification and adaptive inverse control performed with the use of an adaptive linear combiner and least-mean squares algorithm was presented. These concepts were successfully demonstrated using a purely software simulation applied to a linear quarter-car model. Once the process of tuning the adaptive filters was practiced in simulation and the concept proved in software the algorithms were then implemented on the new developed quarter-car hardware using a real-time prototyping control system. The concept of on-line adaptive inverse control presented in this study was successfully demonstrated on hardware-in the loop. The system ID technique introduced works extremely well considering the linear model is modeling a non-linear system. The controller does appear to have some room for improvement.

There are some recommendations that extend from this study which might be considered before taking future steps. For the quarter-car platform these recommendations are simply to continue the development of the rig set forth in this thesis. The first step is to add the electromagnetic force loaders to the sprung mass to simulate body and aerodynamic loads. Another step is to demonstrate the flexibility of the test rig to mount a different design of suspension such as a Nextel Cup Stock Car or HMMWV suspension. Another recommendation is to develop a roll degree-of-freedom for the sprung mass which could actively replicate the roll dynamics of the body of a vehicle with a variable roll-center.

In terms of the control algorithm and testing performed in this study, there are several recommendations which arise. First, it would be ideal to perform several more tests with the quarter-car rig to learn more about tuning the adaptive algorithms. There were certain tests performed along the way where the control filter weights had a totally different appearance which included some higher frequency response. To the authors knowledge these test could not be repeated on demand. It seems that there are some unknown subtleties that affect the results and convergence of the adaptive controller. In light of some of the more surprising results noted in this study further tests may also add insight into how to make the test more repeatable. Also, more tests would be useful to better understand how the frequency content of the desired response and also the shaping of the system ID excitation affect the performance of the inverse controller. Further alterations of the adaptive linear combiner might be necessary to help capture some of the non-linear dynamics of the hardware. This may help the adaptive controller do a better job of replicating the true dynamics of the system.

Furthermore, a similar set of tests should be performed on the quarter-car while feeding the sprung mass desired response to the adaptive filter. During this study the unsprung acceleration was exclusively used as a reference signal. Even so, the sprung mass signal was replicated better than the unsprung mass signal. It would be interesting to see if using the sprung mass signal as reference would further improve the error convergence of the sprung mass replication. Finally, the next step in a future study would be to attempt to implement this algorithm on a more complex system such as an 8-post shaker to yield a fair comparison of this method to existing art.

REFERENCES

1. Bigliant U., Piccolo R., Vipiana C., "On Road Test vs. Bench Simulation Test: A Way to Reduce Development Time and Increase Product Reliability," *SAE Technical Paper Series, No. 905207*, Warrendale, PA, 1990.

2. Kelly J., Kowalczyk H., Oral H. A., "Track Simulation and Vehicle Characterization with 7 Post Testing," *SAE Technical Paper Series, No. 2002-01-3307,* Warrendale, PA, 2002.

3. Mianzo L., Fricke D., Chabaan R., "Road Profile Control Methods for Laboratory Vehicle Road Simulators," *Proceedings of the 1998 IEEE AUTOTESTCON,* Salt Lake City, UT, p. 222-228, 1998.

4. Vetturi D., Magalini A., "Road Profile Excitation on a Vehicle Measurements and Indoor Testing Using a Four-post Rig," Dipartimento di Ingengeria Meccanica – Universita degli Studi di Brescia, 2002.

5. Castor A., "Tune Up At Jaguar," Article, *Design News,* Reed Business Information, Waltham, MA, May 5, 2003.

6. Wright P., "Movers and Shakers," Article, *GrandPrix.com,* The Motorsport Company, October 19, 2000.

7. Personal communication on 7-post test road-profile generation methods with Dr. Christoph Leser, *MTS Systems Corporation,* December 5, 2006.

8. Solderling S., Sharp M., Leser C., "Servo Controller Compensation Methods Selection of the Correct Technique for Test Applications," *SAE Technical Paper Series, No. 1999-01-3000,* Warrendale, PA, 1999.

9. Kowalczyk, H., "Damper Tuning with the Use of a Seven Post Shaker Rig," *SAE Technical Paper Series, No. 2002-01-0804,* Warrendale, PA, 2002.

10. Vaes D., Souverijns W., De Cuyper J., Swevers J., Sas P., "Optimal Decoupling for Improved Multivariable controller Design, Applied on an Automotive Vibration Test Rig," *Proceedings of the American Control Conference, v 1*, p. 785-790, 2003.

11. Goncalves F., "Dynamic Analysis of Semi-Active Control Techniques for Vehicle Applications," Master's Thesis, Virginia Tech, Blacksburg, VA, August 2001.

12. Milliken W. F., Milliken D. L., *Race Car Vehicle Dynamics,* Society of Automotive Engineers Inc., Warrendale, PA, 1995.

13. PT. Labtech Penta Internationa, http://labtech.org/, website, Batam, Indonesia, 2006.

14. FCS Control Systems B. V., *FasTEST Manager User Manual*, Issue 1.2.4, Oude Meer, Netherlands, 2004.

15. MTS Systems Corporation, *RPC Pro,* Brochure, Eden Prairie, MN, 2003.

16. Thoen, B. K., MTS Systems Corporation, Control Network with On-line Iteration and Adaptive Filter*,* US Patent No. 5,394,071, 1995.

17. Lauwerys C., Swevers J., Sas P., "Robust Linear Control of an Active Suspension on a Quarter Car Test-Rig," *Control Engineering Practice 13,* p. 577-586, 2005.

18. Vandersmissen B., Kennes P., Maes M., Reybrouck K., "Skyhook Control and Performance Evaluation of an Active Suspension System on a Quarter Car Test Rig," *Proceedings of the 2004 International Conference on Noise and Vibration Engineering, ISMA,* p. 103-114, 2004.

19. Bay Cast Technologies, http://baycasttech.com/, website, Bay City, MI, 2006.

20. Gobbi M., Giorgetta F., Guarneri P., Rocca G., Mastinu G., "Experimental Study and Numerical Modelling of the Dynamic Behavior of Tyre/Suspension while Running Over an Obstacle," *Proceedings of IMECHE2006, 2006 ASME International Mechanical Engineering Congress and Exposition, No. IMECHE2006-14804,* Chicago, IL, 2006.

21. Toyoda Machine Works Ltd., *PV4-IIA High-speed Vertical Milling Machine User's Manual and Technical Drawings,* Okazaki City, Japan, 1996.

22. NSK Ltd., *Technical Journal – Linear Rolling Guides*, Ann Arbor, MI.

23. Ansys Inc. *User's Manual,* ANSYS 9, Canonsburg, PA, 2004.

24. Marshek K. M., Juvinall R. C., *Fundamentals of Machine Component Design 3rd Edition,* John Wiley & Sons, Hoboken, NJ, 2003.

25. Synergy Racing, http://www.synergyracing.com/, website, Alton, VA, 2006.

26. Personal communication on Porsche 996 suspension setup and Grand American Racing regulations with Greg Jones and Richard Binzer, *Synergy Racing*, 2006.

27. Widrow B., Stearns S. D., *Adaptive Signal Processing*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1985.

28. Widrow B., Walach E., *Adaptive Inverse Control*, Prentice Hall P T R, Upper Saddle River, NJ, 1996.

29. dSPACE digital signal processing and control engineering GmbH, *User's Manual,* Control Desk, 2005.

30. Vaes D., Swevers J., Sas P., "Experimental Multivariable Tracking Control on an Automotive Vibration Test Rig," *ISMA Proceedings of the 2004 International Conference on Noise and Vibration Engineering,* ISMA, p. 311-323, 2004.

31. Oral H. A., "A Tool for Control Algorithm Development of an Active Vehicle Suspension Using Hardware-in-the-Loop," *SAE Technical Paper Series, No. 2002-01-1597,* Warrendale, PA, 2002.

32. Fricke D. M., Hansen M. D., Chabaan R. C., Ford Motor Company, Effective Road Profile Control Method for a Spindle-coupled Road Simulator, US Patent No. 5,610,330, 1997.

33. Kino H., Iwai M., Tamura M., "A Study of Road Load Severity Prediction in Market for Power Spectrum Density," *SAE Technical Paper Series, No. 2003-01-2867,* Warrendale, PA, 2003.

34. Moran T., Sullivan M., Menmuir D., Mahoney J., "Replication of Heavy Truck Dynamic Wheel Loads Using a Road Simulator," *Road Transport Technology – 4,* University of Michigan Transportation Research Institute, Ann Arbor, MI, p. 309-315, 1995.

35. MTS Systems Corporation, Eden Prairie, MN, 2007 (image provided).

36. Servotest Systems Ltd., Slough, Berkshire, England, 2007.

APPENDIX

Appendix A - Rig Maintenance

There are few simple maintenance procedures to ensure good operation of the quarter-car test bed. The following is a short bulleted list of items that require periodic attention.

- NSK Linear guides require Alvania (AS2) grease approximately once a year.

- Keep the guide rails lightly coated with grease or other rust inhibitor during long periods of down time to prevent rust from developing.

- Periodically check the torque of the T-bolt clamps. They can become loose during some vibration tests.

- Monitor the status lights on the hydraulic power unit. The unit has over-temperature and filter clogged lights. If the unit requires filter replacement one can be ordered from the phone number and part number listed directly on the filter housings inside the pump enclosure. MTS documentation contains the changing intervals otherwise.

- The service manifold also has a filter with a mechanical indicator. Periodically check the indicator while a test is running to ensure that the pressure drop across the filter is not too high. If it is replace the filter.

- Inspect service manifold and pump accumulator pressures semiannually or as system performance dictates. Proper accumulator pressure for each may be found on the accumulators or in MTS documentation.

- Inspect rubber hydraulic lines twice per year for rubs or cracks. Fix/protect as necessary.

- The clean the bare steel parts of the load frame and the black-oxide coated parts apply a light coating of WD-40 or equivalent with a paper towel or clean cloth. This will help prevent rust from forming.

- If disassembly of sprung mass and bearings is required please refer to NSK documentation first. Improper disassembly or reassembly misalignment can cause bearing failure.

Appendix B - Rig Specifications

The following section is a short list of physical specification for the quarter-car rig.

- Overall mass of grounding fixtures (base plate and reaction fixture) is approximately 4500 lbs.

- Sprung mass has a rail spacing of 22 in and currently has a max travel of ±13 in. this travel allows for 1.5 in of rail at either end to ensure there is no dislodging of the bearing carriers.

- There is currently a 33x36 in window on which to mount a suspension to the sprung mass plate.

- The bare aluminum sprung mass plate and carriers weigh approximately 170 lbs. This is the absolute minimum sprung mass weight allows by this rig.

- The maximum sprung mass weight is limited by the hydraulic actuator.

- The servo-hydraulic actuator has an amplitude of ±3 in and a maximum static load of 5.5 kip.

- Each bearing carrier has a basic dynamic (carrier moving) load rating of 13,800 lbf in the direction normal to the motion of the carrier. Each carrier can handle an actuator induced moment of 1120 ft-lb (static loading).

- Linear rail parallelism is 0.003 in over the entire length of the guide rails.

Appendix C – Linear State Space Matrices

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ \left(-\dfrac{k_s}{m_u} - \dfrac{k_u}{m_u}\right) & -\dfrac{c_s}{m_u} & \dfrac{k_s}{m_u} & \dfrac{c_s}{m_u} & \dfrac{k_u}{m_u} \\ 0 & 0 & 0 & 1 & 0 \\ \dfrac{k_s}{m_s} & \dfrac{c_s}{m_s} & -\dfrac{k_s}{m_s} & -\dfrac{c_s}{m_s} & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathbf{B} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

$$\mathbf{C} = \begin{bmatrix} \dfrac{k_s}{m_s} & \dfrac{c_s}{m_s} & -\dfrac{k_s}{m_s} & -\dfrac{c_s}{m_s} & 0 \\ \left(-\dfrac{k_s}{m_u} - \dfrac{k_u}{m_u}\right) & -\dfrac{c_s}{m_u} & \dfrac{k_s}{m_u} & \dfrac{c_s}{m_u} & \dfrac{k_u}{m_u} \end{bmatrix}$$

$$\mathbf{D} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

VITA

Justin Langdon was born in Bristol, Pa on June 2, 1981. He grew up in Morristown, TN graduating from Morristown-Hamblen High School West in 1999. One strong influence on his career path becoming engineering related, was Gene Quarles, an excellent math teacher and great person at West High. From there Justin went to Tennessee Technological University to pursue a degree in Mechanical Engineering. While at TTU he developed a strong interest in the automobile and decided to pursue a career in the automotive industry with hope of one day racing professionally. Part of this desire stemmed from working on the Formula SAE team at TTU. Many summers and off-time from school was spent working for an automotive supplier in Morristown. Justin, a KA, graduated from Tech in 2004.

After finishing a Bachelor's degree Justin decided that he just couldn't get enough and came to Virginia Tech to work on a Master's Degree in Mechanical Engineering. He spent the first year taking classes in Blacksburg and helped the VT FSAE team compete in 2005. He ended up in Danville, VA and teamed up with Dr. Steve Southward to work on this research project under the Virginia Institute for Performance Engineering and Research (VIPER) project based at the Institute for Advanced Learning and Research (IALR) in Danville, VA. Justin finally finished the project in late 2006 and received the Master's degree in 2007. At this point the pursuit of a racing career finally came to fruition. After a short stint as a mechanic in the Nextel Cup series with Front Row Motorsports, he finally scored a full time gig as an engineer at Hall of Fame Racing in Charlotte, NC. Justin decided to take a chance and go race. Perhaps he'll make it back to VT one day to continue on in pursuit of a PhD. This would make Steve, John, Jimmie and Wilbur very happy… probably Justin too.