Active Vibration Isolation Using an Induced Strain Actuator with Application to Automotive Seat Suspensions

by

Mark G. Malowicki

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

Donald J. Leo, Chair
Mehdi Ahmadian
Doug Nelson

June 2000

Blacksburg, Virginia

Active Vibration Isolation Using an Induced Strain Actuator with Application to Automotive Seat Suspensions

Mark G. Malowicki, M.S.
Virginia Polytechnic Institute and State University, 2000
Advisor: Donald J. Leo

Abstract

The characteristics of an automotive passenger seat in response to vibrational excitations are examined and an active vibration isolation system incorporating smart materials is designed, built, and tested. Human sensitivity to vibration is discussed. Characteristics of road roughness are discussed and used to implement a representative test input to a passenger seat system. MATLAB is used to model the car seat and vehicle system with four degrees of freedom to determine actuator requirements. Selection and implementation of a low–profile, prestressed piezoceramic device into an active seat suspension system is described, and experimental results of the actuator assembly performance are presented. Vibration isolation is realized in an experimental setup representing one quarter of a seat and passenger's total mass, using one actuator assembly (representing one corner of the seat suspension).

For an input power spectrum representative of a passenger vehicle environment, the smart material actuator assembly, as applied to a quarter seat experimental setup, is proven to be capable of isolating vibration with an isolation frequency of 2Hz and no resonant peak, versus 6Hz and a resonant peak of 2g/g for an actual passenger seat tested.

Acknowledgments

This research was made possible by a generous grant from the Daimler-Chrysler corporation, under the direction of Stephen J. Buckley. Experimental materials were also donated by Johnson Controls. Thanks to Nagi Naganathan and the seat project team at the Univeristy of Toledo for their collaboration.

Many thanks to Don Leo, my advisor and mentor, for his tireless support and guidance. Throughout my time at Virginia Tech, he has been a constant source of technical expertise and enthusiastic encouragement. None of my work would have been possible without his help. I would also like to thank the other members of my committee, Dr. Ahmadian and Dr. Nelson, for their input.

Thanks also to all of my fellow graduate students at the Center for Intelligent Material Systems and Structures, as well as all of my friends here at Virginia Tech, for their friendship and comraderie.

Thanks to Matt Groening, who has inspired me to constantly strive to challenge conventional thinking and maintain an appropriate view of reality.

Thanks to my parents and family for their eternal love and support.

Mark G. Malowicki

Virginia Polytechnic Institute and State University

June 2000

Contents

List of Tables

List of Figures

Chapter 1

Motivation

We live in a society of ever increasing mobility. The ease of transportation which we enjoy makes the road to our friends and relatives' houses ever shorter and continues to make the world a smaller and more accessible place.

Not only do our forms of transportaion become faster and faster as our world evolves, but hopefully, they become more comfortable and enjoyable as well. After all, tools such as planes, trains, and automobiles must ultimately meet their earthly creators and end users with a friendly and convenient interface. Just as an uncomfortable pair of boots can ruin our favorite hike, an uncomfortable car can ruin our drive across town. Thankfully, as we evolve the systems that transport us, we wisely incorporate features to enhance our comfort. In the automobile, we even have two distinct systems to guard us against the harsh road inputs that impact the tires' footprints — a primary suspension (between the wheels and the vehicle body) which filters out most of the road roughness, and a nice comfy car seat that further isolates the passenger from the vibrations that manage to make it through the primary suspension.

The automobile has been around for some time, and so we've ended up with some pretty comfortable car seats. Our automotive design goals, however, continue to evolve just as everything else does. The massive and comfortable seats of yesterday must become more streamlined and lighter in the face of today's more energy conscious society. After all, a lighter vehicle is easier to accelerate, allowing the use of smaller

and more fuel efficient powerplants. This poses a trade–off — how much of a seat can we take away to save weight while still maintaining the level of comfort we desire? Today's car seats achieve their dynamic vibration isolation characteristics mainly through the use of a large piece of compliant foam. Is there a better way?

Perhaps the current method of vibration isolation, a passive block of foam, can be replaced with an active vibration isolating mini-suspension for the seat. If this active system can be made light enough, not only is mass reduced, but many additional advantages of an active system are realized as well. There is now control over the excitations that reach the passenger, and the seat's transmissibility can be shaped as desired (within limits). Passive isolation systems, while offering simplicity, do have their disadvantages. The isolation achieved at higher frequencies requires that the system have a low frequency resonance. These systems behave like low pass filters, with design tradeoffs such as the size of the resonant peak versus how much isolation is achieved at higher frequencies, and how low the resonant frequency can be made without requiring too much displacement. (It should be noted that at this point, the cost is not considered to be a factor in the analysis, as the current focus is to determine the feasibility of applying smart materials to the active seat suspension concept.)

Furthermore, the design of an automotive suspension is always a tradeoff between ride comfort and handling performance. A soft suspension with low damping results in good passenger comfort, while a stiff suspension with more damping is desired for a responsive, well-handling vehicle. By incorporating an active seat suspension, the seat can do more of the work in terms of achieving passenger comfort, and thereby shift the primary suspension setup more towards the handling end of the spectrum. Instead of a primary suspension that must be a compromise between ride and handling, the primary suspension can have an increased focus on handling when combined with an active seat's improvements in comfort.

An active system has other implications outside of passenger comfort, as well. The potential benefits are limited only by the force and stroke of the actuator, the software capabilities, and the imagination. The seat suspension could conceivably

show application to impact mitigation, by isolating the passenger from the damaging accelerations that result from vehicular accidents. The required force and stroke levels for this specific case are not analyzed in this thesis, since the focus is on passenger comfort.

Is an active vibration isolation mount feasible for use in a lightweight automotive passenger seat?

Chapter 2

Literature Review and Background

This chapter provides a general overview of concepts relevant to active seat suspension analysis and design. Engineering fundamentals as well as current research topics are discussed as appropriate. Material from the literature review are included within each section alongside the related subject matter.

2.1 Vibration Isolation

A common problem in the study of vibrations is that of "base excitation," in which a dynamic system is excited by motion of the supporting structure to which it is attached. Examples of this would include an automobile suspension (with the varying road surface as the source of excitation), an expensive and fragile payload aboard a rocketship which is violently thrusting itself away from the earth's surface, or any similar system. Figure 2.1 shows a simple model of the situation, with a mass M connected to a moving base via a spring with rate k and a damper with rate c.

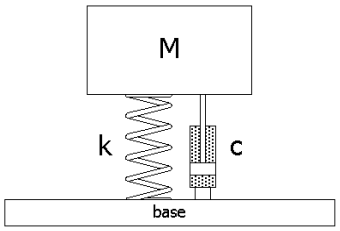

Figure 2.1: Model of Base Excitation Problem

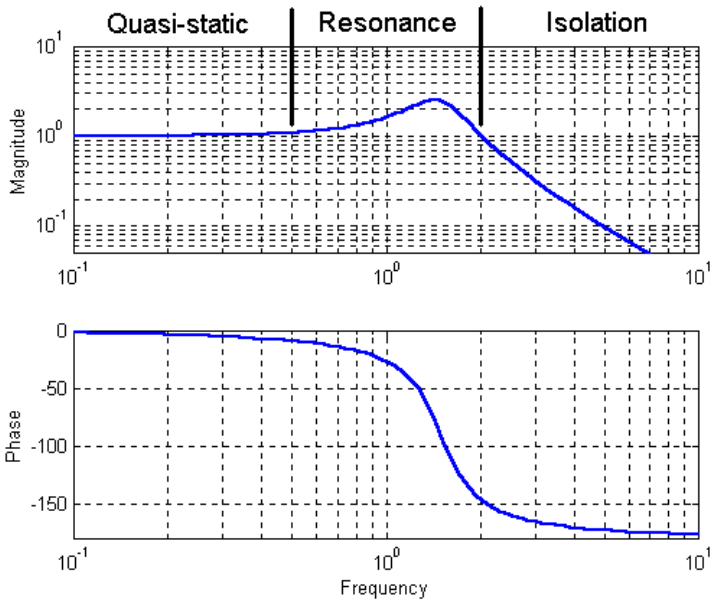

Figure 2.2: Bode Plot of Typical 2nd Order Mechanical System

The base may have some characteristic motion which is undesirable or even damaging to the mass or payload. Hence, it is desirable to isolate the mass from the excitation of the base by a careful selection of an appropriate attachment scheme with a specific spring rate and damping ratio. [Inman, 1996]

In a seat suspension system, the floor of the vehicle to which the seat is attached (the base) is constantly moving due to road surface irregularities, and the mass of the seated passenger is connected to this base by an effective spring rate and damping ratio (which is characteristic of the seat foam or seat suspension system).

It is interesting to note that mechanical systems such as in Figure 2.1 generally behave as low pass filters. At very low frequencies, the mass will move exactly as the base moves, with the same motion and displacement as the base. However, at higher input frequencies, the inertia of the mass will begin to dominate the response, and the result will be a smaller displacement of the mass per given displacement of the base. This follows from the fact that the force required to maintain a certain displacement increases as frequency increases, since the acceleration of the mass is increasing with increasing frequency. A larger amount of energy would be required for the same displacement at a higher frequency. Figure 2.2 shows a frequency response function

for such a system, with three recognizable regions: low frequency in–phase motion, amplification near the resonant frequency of the system, and attenuation at higher frequencies. The magnitude ratio is shown in the top graph, with relative phase shown in the bottom graph. Note that the 2nd order system shown in the graph has a resonant frequency at 1.5, with attenuation of the input signal after an isolation frequency of 2. For an input signal at a frequency higher than 2, the output will be a signal with a smaller amplitude than the input. Hence the origin of the term "transmissibility" — the frequency response function gives an indication of how much of the input signal is transmitted through the system of interest.

If the base excitation has a known frequency content, it is easy to customize the isolation system for the particular application. If there is one known excitation frequency coming from the base, or perhaps only a few known excitation frequencies together, the system can be designed such that the isolation frequency is well below the lowest excitation frequency. That is, all excitation frequencies should fall within the region of attenuation, if possible. The location of the isolation frequency is controlled by varying the spring rate, which shifts the resonant peak to the left or right in Figure 2.2.

It is important to note that the tradeoff for lowering the isolation frequency is a resulting increase in required travel. As the spring is softened, it becomes easier to move the mass relative to the base, and a given input force will cause a larger displacement of the mass. This becomes an important design consideration in the real world, since systems are generally designed to occupy as little space as possible, and too much travel eventually results in a collision with some other structure or system.

Figure 2.3 demonstrates the change in resonant frequency with changing spring rate. The lower natural frequencies have a lower spring rate, or softer springs, while higher natural frequencies accompany higher spring rates. Note that in this figure, mass is *not* normalized, causing the height of the resonant peak to vary with varying spring rate.

Similarly, the damping ratio can be fine–tuned in order to control the height

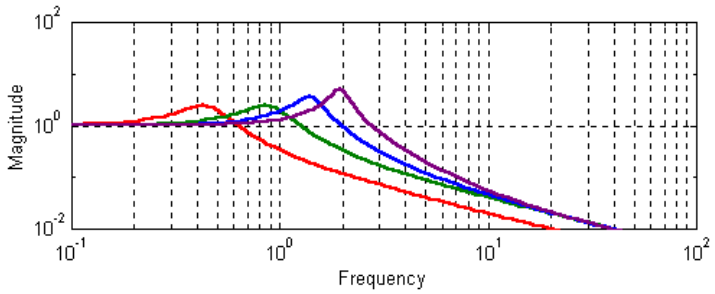

Figure 2.3: Effect of Varying Spring Rate

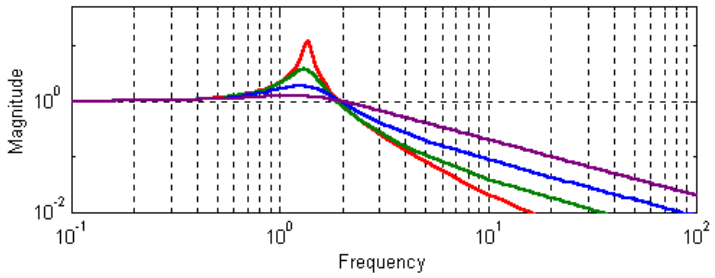

Figure 2.4: Effect of Varying Damping Ratio

of the resonant peak. This is yet another engineering design tradeoff, in that with a *decrease* in the height of the resonant peak, there is a corresponding *increase* in the high frequency transmissibility. Figure 2.4 demonstrates the situation, with the low damping case having a sharp peak (with a high transmissibility) and the high damping case having little to no peak (and decreased isolation). A small resonant peak is desirable, so as not to drive the system out of control in the presence of an input at the natural frequency. High frequency isolation is the real goal, and some of this isolation is sacrificed with increased damping. With a little engineering work, an acceptable medium can usually be found.

2.2 Road Roughness

If all roads were perfectly flat, cars wouldn't even need suspensions in the first place since no vertical excitation forces would be impacting the vehicle. There would be no vibration from which to isolate the passenger. In the imperfect world that is

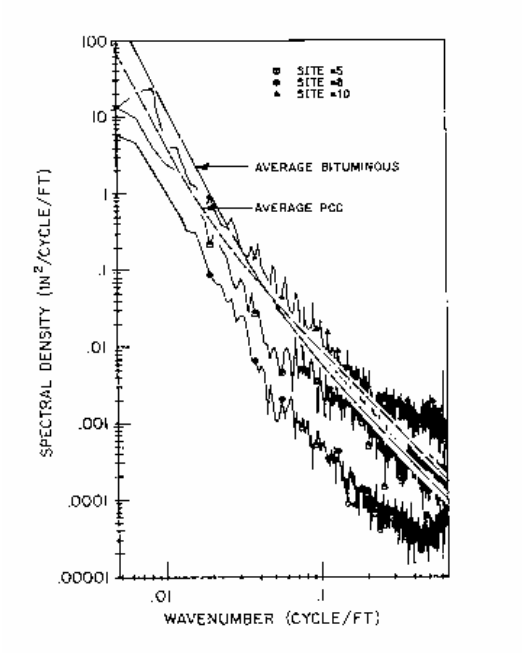

Figure 2.5: PSD of Typical Road Surfaces [Gillespie, 1985]

reality, roads are not nice and smooth due to construction limitations and also due to road surface deterioration over time. This road roughness is the ultimate input to the mechanical system under consideration, and so an understanding of the car seat suspension system must start here.

In the beginning, it is necessary to know what the actual elevation profile of the road surface looks like — all of the nooks and crannies, potholes, and even hills. An actual road can be measured to determine the change in vertical height versus distance along some length of the road. This data can then be manipulated into forms that are a little more meaningful in the context of automotive suspensions. [Gillespie, 1985]

Road roughness is a random signal, but even random signals have a characteristic signature that can be expressed by decomposing the elevation profile into a frequency–based expression of its content. The road surface profile can be broken down by the Fourier Transform process into many superimposed sine waves, each with a specific amplitude and phase. Plotting the magnitude of each component at each frequency yields a Power Spectral Density (PSD) graph. Typical roads are rep-

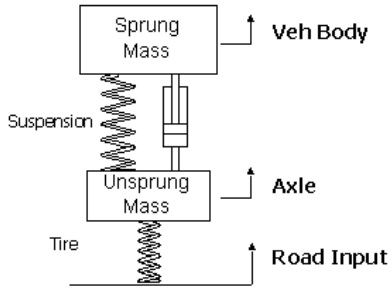

Figure 2.6: Quarter–Car Vehicle Model

resented in the data displayed in Figure 2.5 [Gillespie, 1992]. The x-axis is labeled in cycles per foot, representing spatial frequency rather than temporal frequency. Once a speed of travel is known, the data can be converted into cycles per second.

The PSD amplitude falls with increasing frequency, which is typical of all roads due to the fact that low wavenumber (high wavelength) surface disturbances have larger height variations than do high wavenumber (low wavelength) disturbances. For example, traveling over a hill is a low wavenumber event with a very high amplitude, while traveling over a speedbump is a higher wavenumber, lower amplitude event.

Figure 2.5 shows various types of paved surfaces, with the Portland Cement Concrete (PCC) and bituminous asphalt datasets labeled respectively. Basically, a rougher road has a higher amplitude level, and all road types demonstrate the characteristic falloff with increasing frequency.

This road profile data becomes more meaningful when converted into the temporal frequency domain based upon an assumed speed of travel. Also, since acceleration is more commonly used in the measure of ride quality, the displacement (road surface elevation variation) can be expressed as an acceleration by simply differentiating twice. The first graph in Figure 2.7 shows the result for a vehicle traveling at 50mph. Since the differentiation involves multiplying by the frequency, the acceleration spectrum now shows a characteristic *increase* in amplitude with increasing frequency.

To finally reach the level of the vibration experienced at the vehicle level, the

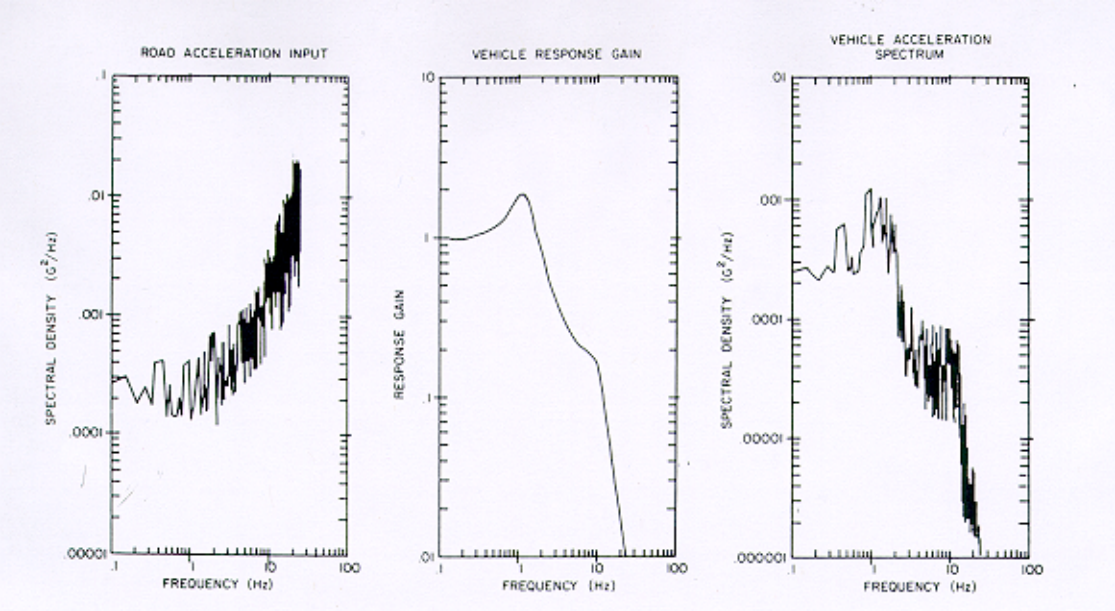

Figure 2.7: PSD After Vehicle Suspension [Gillespie, 1992]

isolation that is provided by the vehicle's primary suspension must be taken into consideration. Taking a simple quarter car model, the road roughness is transmitted through the suspension as illustrated in Figure 2.7 [Gillespie, 1992]. The middle graph is the frequency response function of a typical automotive suspension, with the resonance of the vehicle mass occurring at 1Hz, and the wheel hop mode at 10Hz. Such a response is the result of a simple quarter car model, as shown in Figure 2.6. Taking the road acceleration input and multiplying it through the suspension response, we finally arrive at the acceleration spectrum at the vehicle level.

From the third graph in Figure 2.7, the vehicle suspension has an isolation frequency on the order of 2–3Hz, and indeed does a good job of filtering out the high frequency end of the road roughness input. At the level of the vehicle floorboard, only lower frequency disturbances will now be experienced. This is where the automotive seat enters the picture — the vibration that has managed to survive the primary suspension's filtering efforts is now left as the input to the base of the automotive seat. The seat is the final level of isolation between road roughness and the passenger. It is worthy to note that the input signals at this point fall into a low frequency range, with excitations higher than 20-30Hz sufficiently filtered out by the vehicle's primary

suspension.

2.3 Passenger Comfort

The natural mode of human transportation is travel by foot. Any other form, whether it be a bicycle or an airplane, is a form that presents us to an unnatural environment with unique and characteristic vibrational excitations. While the human body is generally tolerant and adaptable, some vibrational environments can have serious adverse effects on our comfort and well being. Certain vibration inputs to the human body can even cause back disorders, decreased eye-hand coordination, and vision impairment [Berger and Gilmore, 1993]. Is there a way to quantify what type and level of vibration inputs are unpleasant or harmful? This question has been asked and answered by a variety of researchers.

An automobile driving along a real, imperfect road surface at varying speeds is a complex system that is difficult to characterize, especially when it is recognized that excitations can also arise from onboard sources such as the engine, drivetrain, wheel imbalances, and so on. This input acceleration is a complex combination of many inputs at various frequencies. Also, while there may be no absolute standard of human comfort due to the fact that no two human beings are exactly alike (in fact, the array of body weights, sizes, and shapes seems limitless), the research results have proven that we can define a zone or area in which people generally experience discomfort. Above this region the vibration is certain to be intolerable, and below this region the vibration level is acceptable. In order to quantify this region, we can begin by simplifying our considerations to a human being subjected to simple sinusoidal motion in one direction. By taking data points at specific frequencies, we can define a curve which defines the tolerance level versus frequency.

The subject of human tolerance to vibration appeared in SAE publications as early as the 1950s, when the work of several researchers was included in the Ride and Vibration Data Manual [SAE, 1965]. An excellent summary is given in Gillespie's text, and is shown in Figure 2.8.

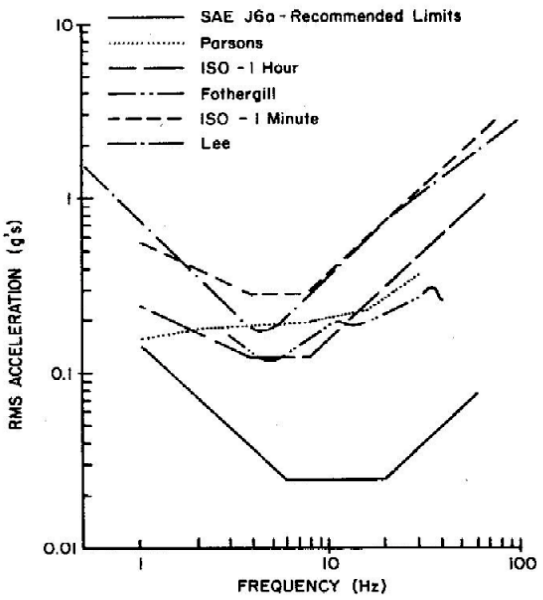

Figure 2.8: Tolerance to Vertical Vibration [Gillespie,1985]

Each line in the figure represents a line of constant comfort above which a test subject will complain of excessive vibration. While the absolute level of vibration at which a person will complain varies from study to study, the characteristic shape is seen to be similar. Human beings are especially sensitive to vibration in the range of 4-10Hz, where each curve dips to its minimum. This is often attributed to resonances in the abdominal cavity, but this frequency has also been shown to excite resonant for-aft motion of the human head in response to vertical excitation [Paddan and Griffin, 1988]. The skull also exhibits a vertical resonance around 10-14Hz, and some comfort curves show a slight dip in this region. Above and below this 4-14Hz range, sensitivity decreases and humans are tolerant of higher acceleration levels.

The variance seen between researchers is due to many factors. Differing seating positions have been used, along with presence or absence of backrests. Most studies use rigid, flat seating surfaces while some researchers have used compliant seating. No standard definition of "discomfort" exists either, with some researchers using nine point scales and some using a reference vibrational level which is repeated for each data point. Furthermore, some data suggests that the curves vary from subject to subject and that differences can be correlated to variables such as subject size [Griffin,

Whitham, and Parsons, 1982].

Research done by M. J. Griffin exhaustively analyzed comfort contour curves in all three translational directions, as well as rotational excitations about all three axes. The discomfort levels he reports in the vertical direction agree well with those of other researchers. When examining human response to each of the specific directional inputs, vibration in the vertical direction was shown to be the most upsetting to test subjects. Vertical seat excitations caused discomfort at lower vibration levels for a greater frequency range than other directional inputs. Griffin's experiments also involved vibrational inputs through the backrest and footrest, with vertical motion of the seat still seen as causing the most discomfort [Griffin, Parsons, and Whitham, 1982].

It is interesting to tie this together with the characteristics of road inputs and with the response behavior of automotive primary suspensions. Humans are most tolerant of high frequency vibration, yet the magnitude of road inputs falls off with increasing frequency. Also, vehicle suspensions do a good job of filtering out high frequency inputs, but allow lower frequencies to pass through to the body of the vehicle. By the time we end up at the interface of the vehicle and our seated human at the surface of the seat, high frequency vibration is of lesser concern. In designing a seat or a seat suspension, it is really the low frequency content of the signal that will tend to bite us, perhaps in the ischial tuberocities.

2.4 Smart Materials

> The concept of a "smart" structure or material is a total anathema to some and a technical and intellectual stimulus to the majority. The concept is all about materials and structures that can react to the world within which they operate and thereby enhance their functionality or survivability.
>
> *–Brian Culshaw [Culshaw, 1996]*

So what is a smart material, anyway? Have steel I–beams evolved to the point of growing brains and plotting to overtake the human race? Thankfully not. The

term "smart" here is used in a broad sense, implying some form of ability to take in and process information, make a decision based upon that information, and choose whether or not to act upon the outside environment. It's a lot to ask for from a sheet of aluminum. Some might argue that a material all by itself can indeed be smart, such as a living plant whose individual cells react to sunlight and appropriately bend a stem or leaf in the sunlight's direction, or such as the bimetallic strip in your thermostat that turns the furnace on or off based on ambient temperature. A more realistic concept might be to extend the consideration of intelligence to a system or structure.

On the level of a system or structure, it is much easier to envision a system such as a bridge with some sort of strain sensing elements incorporated into the structure itself, whose signals are gathered at a central processor for analysis, and can then cause an alarm signal to be sent out when damage is detected. Even better would be a bridge that could then repair itself. Such a system would indeed be deserving of a label of "smart". The vital components of this scheme would be a material that could both sense and act upon its environment.

Certain classes of materials do lend themselves well to use as both sensors and actuators, and some can even be tuned to act based on certain conditions, such as the bimetallic strip in home thermostats. These materials have come to be labeled "smart materials" for these reasons. Typically, it is this class of materials that make our smart systems and structures possible and/or feasible, so perhaps it's reasonable that some intelligence is given to them by association. Included in the list of everyone's favorite smart materials are shape memory alloys, electrorheological and magnetorheological fluids, and piezoelectric materials. These materials are considered here for application to an active vibration isolation mount based on their actuator characteristics.

Shape memory alloys (SMAs) are alloys which, when plastically deformed at one temperature, will return to their undeformed state after being heated above a certain material–specific transformation temperature [Gandhi and Thompson, 1992]. The most common SMA is Nitinol, which is a nickel titanium alloy, but many copper alloy systems and some plastics also exhibit the shape memory effect. The precise

composition of the alloy determines the critical temperature at which the material will transform between a martensitic and austenitic phase, with common transition temperatures being in a range of -150° to 150°C. Shape recovery strains on the order of 6% can be achieved, and the materials are typically relatively heavy but reasonably strong. The energy storage capacity of SMAs is very large. [Culshaw, 1996], [Gandhi and Thompson, 1992]

The thermal drive requirement of shape memory alloys is typically a limiting factor. In order to achieve good response times, a high power heating device is required. Furthermore, in order to minimize power consumption, the SMA would have to be thermally insulated from it's surroundings and possibly even from the structure it is actuating. This thermal insulation implies long thermal time constants, which limits the response time and may require forced cooling to be used. The typical frequency response of shape memory alloys, when used as an actuator, is 0-5Hz. [Culshaw, 1996], [Gandhi and Thompson, 1992]

Electrorheological fluids and **magnetorheological fluids** exhibit changes in rheological properties (such as viscosity or elasticity) in the presence of electrical and magnetic fields, respectively. The fluids incorporate a suspension of particles which are uniformly distributed throughout. With the application of an electric or magnetic field, the particles stick together and create fibrils, modifying the viscous properties of the fluid considerably. Response time is typically less than one millisecond. Power requirements are similar for the two, but MR fluids require low voltage and low current, while ER fluids require high voltage with low current. ER and MR fluids are well suited to systems in which variable damping is desired, and have found commercial application in the dampers used in heavy truck seat suspensions [Song, 1999]. The properties of the fluid depend on a uniform distribution of particles, with some fluids exhibiting the potential for sedimentation.

Piezoelectric materials are materials which both generate an electrical charge in response to mechanical deformations as well as exhibiting mechanical deformation in response to an applied electric field. This phenomenon occurs naturally in quartz and other crystalline materials, but usually to a small degree. Manmade piezoelectrics

such as $BaTiO_3$ and Lead Zirconate Titanate (PZT) show improved properties. PZT's are ceramics and are therefore brittle, but they lend themselves well to being manufactured in a vast array of shapes and sizes. A common configuration is to stack many layers of thin sheets in order to realize any substantial displacement. PZT actuators typically have plenty of force, but small displacements.

A summary of these smart materials is given in table 2.1.

Table 2.1: Comparison Of Smart Material Actuators [Gandhi and Thompson, 1992], [Culshaw, 1996]

Actuator	Max Strain	Max Stress	Response
Shape Memory Alloys	5-6%	50 GPa	0-5 Hz
Magnetorheological Fluid	n/a	~500 Pa shear	0-12,000 Hz
Piezoelectrics	0.20%	50-100 GPa	1-20,000 Hz

Chapter 3

Modeling

3.1 4DOF Model

Design of a feasible actuator requires an up–front understanding of the required forces and strokes involved. Given the list of actuator materials to choose from, each material's potential can be evaluated based upon an analysis of what level of actuator performance is required for an effective seat suspension.

Consider the quarter car vehicle suspension model previously shown in Figure 2.6. Road inputs enter at the base, travel through a tire to the axle, and then through the vehicle's primary suspension to the vehicle level. Once we have reached the vehicle level and a significant portion of the road roughness has been filtered out, we have reached the input to our seating system. Just as a seat is bolted to the floorboard of the vehicle, a seat model can now be added onto the top of the quarter car model.

The resulting four degree of freedom model is shown in Figure 3.1. Two additional degrees of freedom are added in order to consider the presence of both foam cushioning and an active suspension between the passenger and the vehicle. The foam is modeled with a spring rate K_f and damping rate C_f, while the actuator is represented by a spring rate K_a in parallel with the force generating actuator.

The model allows a good degree of flexibility in the analysis. Various vehicle applications can be considered by modifying the primary suspension spring rate and

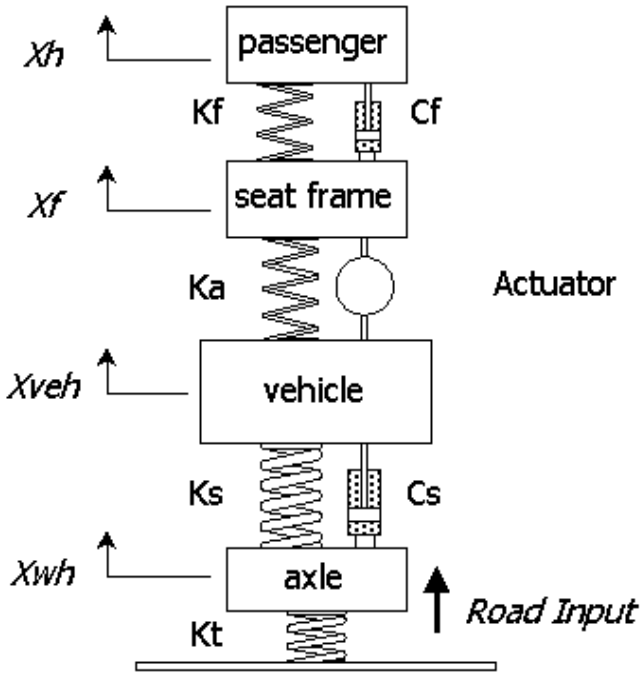

Figure 3.1: Four Degree Of Freedom Model

damping ratio. Seat isolation can be evaluated for the case of a foam seat only and a rigid connection in place of the actuator (with K_a set very large to imitate infinite stiffness) or a fully active seat with no foam (K_f set to a very large value). Even when removing foam to save weight in the vehicle, some foam will be required beneath the seated passenger in order to comfortably conform to the contours of the passenger's posterior. A seat suspension can easily take over the task of serving as the spring and damper that isolate vibration, but some foam will still exist on the seat for the purpose of properly interfacing the seat and the human. Thus, the lightweight seat application under consideration in this thesis maintains some amount of foam even in the active case.

The pitch and roll modes of the vehicle are not included in the model. Not only is the vertical input the dominant excitation arising from a paved road (in most cases), but in addition, human discomfort is known to be most sensitive to vibration in the vertical direction. Roll excitations do arise from road surfaces, such as localized potholes (which are high wavenumber events), but due to the requirement that the general elevation of both wheel tracks remain consistent, roll excitation has a power

spectral density that falls off with *decreasing* wavenumber [Gillespie, 1985], and thus is largely filtered out by the suspension. Pitch may become critically important if the road contains a specific forcing frequency, such as would be present in a fixed slab length concrete road, which excites a pitch resonance of the vehicle. Such a special case is not considered here. A vertical only model serves well to determine approximate actuator stroke and force requirements.

The equations of motion for a position vector x_1 are given by

$$x_1 = \begin{bmatrix} x_h \\ x_f \\ x_{veh} \\ x_{wh} \end{bmatrix}$$

$$\begin{bmatrix} M_h & 0 & 0 & 0 \\ 0 & M_f & 0 & 0 \\ 0 & 0 & M_{veh} & 0 \\ 0 & 0 & 0 & M_{wh} \end{bmatrix} \ddot{x_1} + \begin{bmatrix} c_f & -c_f & 0 & 0 \\ -c_f & c_f & 0 & 0 \\ 0 & 0 & c_s & -c_s \\ 0 & 0 & -c_s & (c_s + c_t) \end{bmatrix} \dot{x_1} + \begin{bmatrix} k_f & -k_f & 0 & 0 \\ -k_f & (k_f + k_a) & -k_a & 0 \\ 0 & -k_a & k_s + k_a & -k_s \\ 0 & 0 & -k_s & k_s + k_t) \end{bmatrix} x_1 = \begin{bmatrix} 0 \\ -F_a \\ F_a \\ F_{road} \end{bmatrix} \tag{3.1}$$

where vertical positions are given by x_h for the human passenger, x_f for the seat frame, x_{veh} for the vehicle (sprung mass), and x_{wh} for the wheels and axle (unsprung mass). The equations can then be written in the state space form according to

$$\begin{aligned} \dot{x} &= Fx + Gu \\ y &= Hx + Ju \end{aligned} \tag{3.2}$$

with variables defined as

x	state vector
y	output vector
u	input or control vector
F	system matrix
G	input matrix
H	output matrix
J	feedforward matrix

The input vector u contains both a defined road input signal, and an actuator input

force as determined by the feedback system. The state vector x is now defined as

$$x = \begin{bmatrix} x_h \\ x_f \\ x_{veh} \\ x_{wh} \\ \dot{x}_h \\ \dot{x}_f \\ \dot{x}_{veh} \\ \dot{x}_{wh} \end{bmatrix} \tag{3.3}$$

The state space matrices for the system, according to equation 3.2, become:

$$F = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ \frac{-k_f}{M_h} & \frac{k_f}{M_h} & 0 & 0 & \frac{-c_f}{M_h} & \frac{c_f}{M_h} & 0 & 0 \\ \frac{-k_f}{M_f} & \frac{-(k_f+k_a)}{M_f} & \frac{-k_a}{M_f} & 0 & \frac{c_f}{M_f} & \frac{-c_f}{M_f} & 0 & 0 \\ 0 & \frac{k_a}{M_{veh}} & \frac{-(k_s+k_a)}{M_{veh}} & \frac{k_s}{M_{veh}} & 0 & 0 & \frac{-c_s}{M_{veh}} & \frac{c_s}{M_{veh}} \\ 0 & 0 & \frac{k_s}{M_{wh}} & \frac{-(k_s+k_t}{M_{wh}} & 0 & 0 & \frac{c_s}{M_{wh}} & \frac{-(c_s+c_t)}{M_{wh}} \end{bmatrix}$$

$$G = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & \frac{-1}{M_f} \\ 0 & \frac{1}{M_{veh}} \\ \frac{k_t}{M_{wh}} & 0 \end{bmatrix} \quad H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & -1 & 0 \end{bmatrix} \quad J = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \tag{3.4}$$

For these particular H and J matrices, the values in the output vector y will be:

$$y = \begin{bmatrix} \text{passenger displacement} \\ \text{vehicle displacement} \\ \text{actuator force} \\ \text{passenger displacement relative to vehicle} \\ \text{passenger velocity relative to vehicle} \end{bmatrix}$$

For the input matrix G, the first column represents the road input and the second

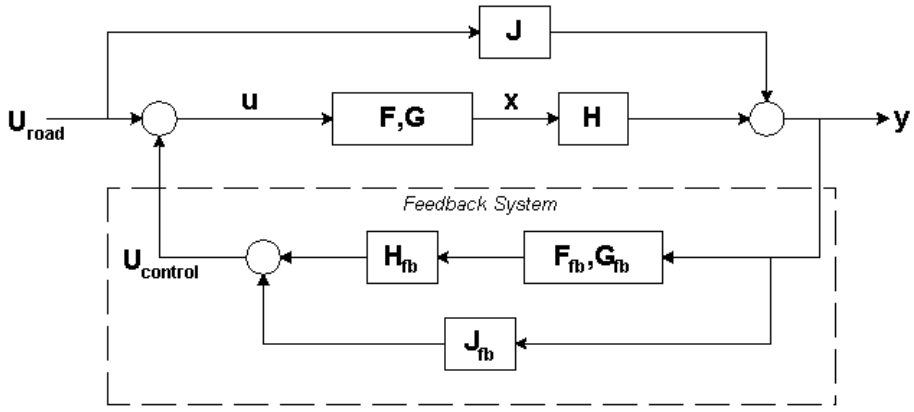

Figure 3.2: State Space Block Diagram

column represents the actuator input force. The J matrix has been chosen to output the actuator force in the third row of y. The H matrix as shown here outputs the position difference and velocity difference (of the attachment points of the actuator) in the fourth and fifth rows of y, respectively.

With the complete model of the system in the free response condition, the next step is to close the loop and apply some form of feedback in order to alter the system response as desired. The feedback system can also be modeled in state space form. Figure 3.2 shows a block diagram representation of the complete closed loop system. The complete set of feedback matrices allows the model to incorporate representative actuator dynamics in addition to the system dynamics. To simply apply feedback gains based upon the output of the model, gains are placed in the J_{fb} matrix, with the other feedback matrices set to zero. While actuator dynamics are ignored, this provides a sufficient analysis of required actuator performance to size the seat suspension actuator.

3.2 Matlab Results

The model was coded into MATLAB for analysis. The complete code is given in Appendix A. MATLAB version 5.3 was used. Parameters are as shown in Table 3.1, and are easily modified to cater to differing vehicle configurations. Known mass values

Table 3.1: 4DOF Matlab Parameters

Vehicle Suspension Resonance	1.5	Hz
Wheel Hop Resonance	12	Hz
Vehicle Mass	1200	kg
Unsprung Mass	150	kg
Human Mass	65	kg
Seat Mass	18	kg
Primary Suspension Stiffness	106,500	N/m
Tire Stiffness	852,700	N/m
Seat Cushion Stiffness	50,000	N/m
Primary Suspension Damping	18,000	Ns/m
Tire Damping	10	Ns/m
Seat Cushion Damping	600	Ns/m

are used for the seat, and the remaining mass values are approximated for a typical small to mid–size sedan type automobile. Spring and damper rates are chosen so as to have a typical primary suspension resonance at 1.5Hz and a wheel hop mode at 12Hz. Seat cushion damping is based upon experimental measurement of an actual passenger seat.

Three cases were examined: 1) the open loop, free response of the baseline seat having full foam thickness, 2) the open loop free response of a seat with half the foam removed, and 3) the closed loop response of a seat which has half the of the cushion removed and an active seat suspension added in. Case 2 demonstrates the degradation in vibration isolation that results from removing half of the foam out of the seat, and case 3 represents the attempt at forcing case 2 back to at least the level seen in case 1. The goal of the active system is to recover the performance lost in removing foam and so case 1 is taken to be the desired level of isolation. Of course, improvements beyond this point are welcome in that they prove the actuator to be more than capable of the task at hand.

Feedback gains are applied to two outputs of the system — the position difference between the passenger and the vehicle, and the velocity difference between the passenger and the vehicle. Individual gains for each case are chosen separately based on a root locus diagram of a single valued feedback matrix, then these gains are

employed as a starting point for the dual valued feedback case, and scaled back until an appropriate response is achieved with good stability. Even with this admittedly crude method of gain selection, a more than sufficient level of feedback control is achieved.

A sample time history of MATLAB output is given in Figure 3.3, which shows the vertical displacement of the passenger (x_h) for each of the three cases. It is obvious from the figure that removing foam from the seat causes larger excursions in x_h, which will be experienced by the passenger as a less comfortable ride. The closed loop case demonstrates a marked improvement in the figure, more than recovering the performance lost with the removal of foam from the seat.

The isolation achieved with the active seat is better examined in the frequency domain. Figure 3.4 shows the frequency response function for all three cases. With the decrease in foam thickness, the performance loss is seen as a higher peak at both the seat suspension resonance at 4Hz and the wheel hop resonance at 12Hz. This is due to the loss of damping with the decrease in foam thickness. Note that the resonant frequencies remain constant, since in the half foam case the spring in parallel with the actuator is assumed to recover the loss in spring rate with the removal of foam. That is, overall spring rate of the seat foam and seat suspension system is held constant, and so the stiffer rate of the half foam thickness case is balanced by an appropriate value of K_a. The spring shown by K_a in Figure 3.1 is required to offload the actuator by supporting the static weight of the seat and passenger. The actuator then needs only to apply force for the dynamic movement of the mass, and not to support the mass in a static condition. This greatly decreases actuator force requirements.

With the actuator incorporated in feedback control of the system, a performance improvement is obvious in Figure 3.4. The isolation frequency of the system has been decreased from 5Hz to 2.5Hz, and the transmissibility at both resonant frequencies has been decreased.

Now that the J_{fb} matrix has been determined for a sufficient level of performance, we need only examine what actuator forces and strokes occur in the MATLAB model. Figure 3.5 shows a time history of both actuator force, in Newtons, and actua-

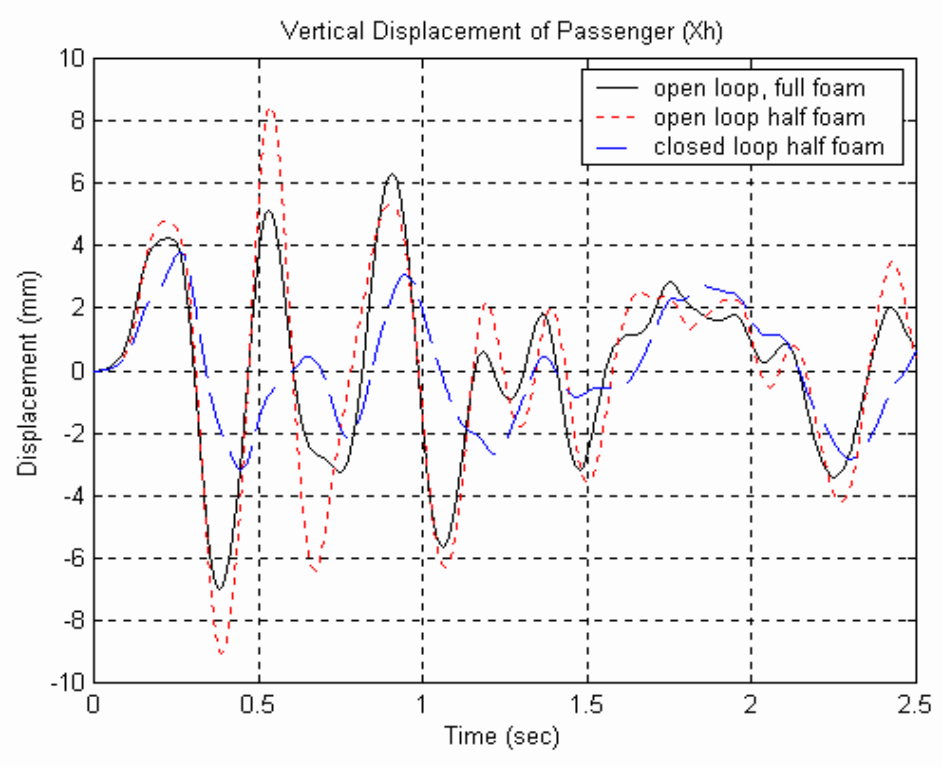

Figure 3.3: Matlab Output – Time History

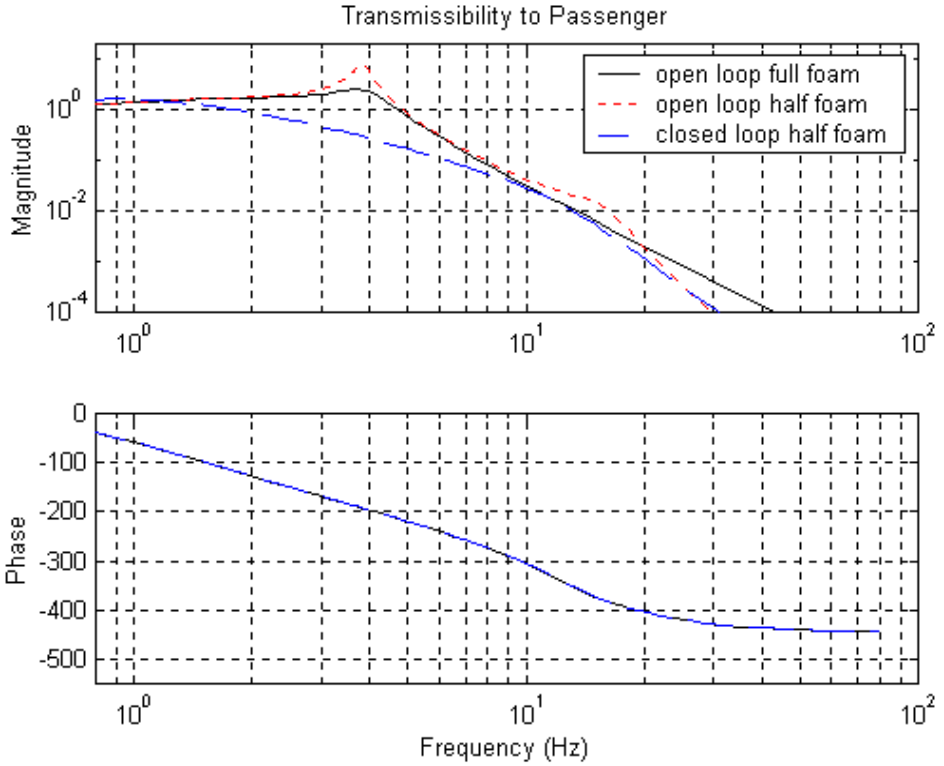

Figure 3.4: Matlab Output – Frequency Response

tor stroke, in millimeters. Actuator force can be seen to reach peaks of approximately ± 400 Newtons, with actuator stroke reaching ± 4 millimeters. RMS actuator force is 288N, and RMS actuator stroke is 2.9mm. Results are summarized in Table 3.2.

The power consumption of the actuator can also be estimated at this point from the model. Ideal power is calcuated as 4 Watts average.

This is now the starting point in choosing what type of actuator our system should employ. Knowing the force and stroke requirements narrows down the field considerably, as will be discussed further in the next chapter.

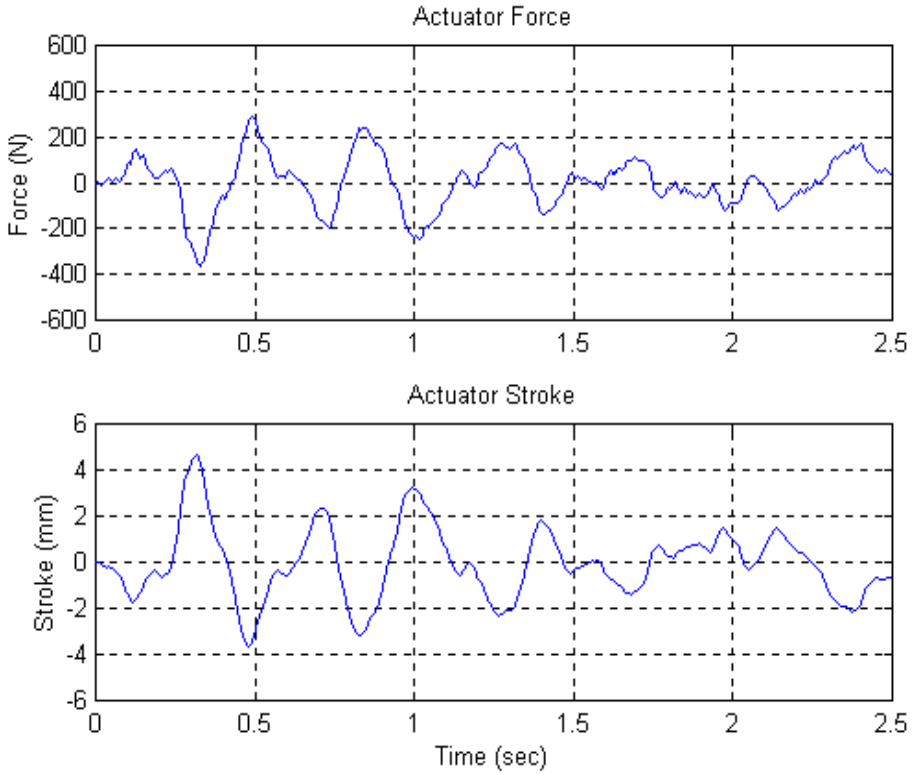

Figure 3.5: Matlab Output – Force And Stroke Levels

Table 3.2: Matlab Results For Baseline Seat And Active Suspension Seat

	Foam Seat	Active Suspension
Peak Amplification	2	1
Isolation Frequency	6 Hz	2 Hz
RMS Passenger Displacement	4.07 mm	2.12 mm
Peak Actuator Force	--	±400 N
RMS Actuator Force	--	288 N
Peak Actuator Stroke	--	±4.0 mm
RMS Actuator Stroke	--	2.9 mm

Chapter 4

Seat Suspension

4.1 Design Constraints

The challenge that rises before us is to recover the performance of a baseline car seat, when a large portion of the foam is removed, by use of a smart material actuator. Several design constraints are implied in this statement alone.

Perhaps the most salient design goal is the issue of mass reduction. The reason that foam is being removed in the first place is to reduce mass, and so the actuator must be lightweight as well as meeting force and stroke requirements. Specifically, our actuator should be lighter than the foam that it replaces. This is a tall order indeed! For the passenger seat donated by Johnson Controls, the mass breakdown is as shown in Figure 4.1. The foam weighs in at 2.66 kilograms or 13% of the total mass of 21kg. (That's 5.8 pounds of foam for the metric system inept among us.) If half of the foam is removed and replaced with an actuator, the actuator mass should be less than 1.33kg (2.9lb) in order to maintain mass reduction in the overall system. A lightweight actuator is a must.

Another consideration is the possible range in variation of passenger mass. The system needs to be able to cope with a small, lightweight person as well as with a larger, higher mass person. This limits the choice of spring rate for the spring placed in parallel with the actuator, k_a. With too soft of a spring rate there will be too much deflection difference between a light and a heavy passenger, changing the static

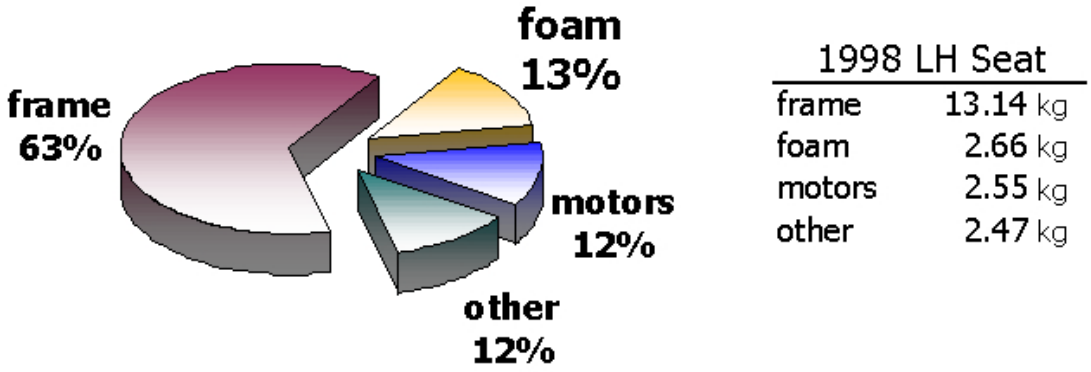

Figure 4.1: Seat Mass Breakdown

operating point of the actuator. Ideally, with the spring and actuator in parallel, the applied passenger weight would result in zero load on the actuator at static deflection.

There is also the question of whether the seat suspension should attempt to isolate vibration in all three axes and also in all three rotational directions. It would be possible to independently apply actuator forces at each corner of the seat so as to realize multidirectional control, but control algorhythm complexity and actuator duty cycle requirements would increase. Furthermore, we have seen evidence implying that isolation in the vertical direction is the main concern. Therefore, we can choose to control seat motion in the vertical direction only, and constrain the seat in the lateral, fore–aft, roll, pitch, and yaw directions.

Additionally, there isn't an enormous amount of room to work with beneath a car seat. An actuator must either be slim enough that it doesn't raise seat height above an acceptable level, or in the case of a long piston type actuator or damper, must be oriented near parallel to the vehicle floorboard and then actuate or dampen the seat through a set of linkages.

4.2 Actuator Design

4.2.1 Material Considerations

Now that the expectations of the actuator are defined, the various classes of smart materials can be re–examined to determine if any options exist that meet the force and stroke requirements, while at the same time being very lightweight. The field quickly dwindles in the face of the system requirements.

Shape memory alloys immediately show several limitations. The need for heating and/or cooling the material (in order to alter material properties by crossing the transition temperature) necessitates a high level of complexity in the design of the system. The very idea of having a heating element underneath a car seat conjures up images of undesirable thermal events (such as fire) in the vehicle interior. The thermal insulation required for safety and for power consumption issues also limits the response time to very low frequencies, on the order of 0–5Hz. Shape memory alloys, in their current state, are most likely not the best option.

Magnetorheological (MR) fluids show promise in their application to variable rate dampers. The classic piston and valve damper is easily modified to accept the addition of a magnetic field at the valve, allowing control of damping rate by modifying localized fluid properties. The resulting system will be a semi–active suspension, since it has both a spring and damping element, one of which has time varying (controlled) properties. Damper stroke is limited only by the spatial envelope in which the damper must be placed, and high forces are easily achievable. Magnetorheological dampers are already in use in the primary suspension of race cars, and also in the seat suspensions of tractor–trailers [Song, 1999]. Force levels for the damper made by Lord Corporation for heavy truck applications reach levels as high as 2000N. [Carlson, 1999]

Disadvantages still exist. A fluid damper is typically heavy, violating our mass reduction requirements. Also, typical piston type fluid dampers would be a tight squeeze under a car seat, most likely acting through some form of linkage which allows the damper to be laid horizontal. The small displacements involved in an

automotive application may not involve enough relative displacement for effective control by variable damping, either. The commercial viability of MR fluid dampers in large truck seat suspensions is possible due to the increased design flexibility allowed by the application. Truck seats sit high enough to allow plenty of space for a seat suspension. The amount of mass in commercial heavy trucks is so large that the additional weight of a fluid damper becomes much less significant than it would be in a passenger vehicle. Furthermore, a tractor–trailer operator spends his or her entire day in the driver's seat, whereas a passenger vehicle's expected use is shorter commute–length trips. Hence, comfort is an even higher priority in a heavy truck than in a passenger vehicle, allowing more room for sacrifice in the areas of added weight and added complexity.

A recent development which combats some of these disadvantages is a sponge type MR fluid device. Instead of employing a volume of fluid which is valved between two chambers, a sponge damper contains the fluid in an absorbent matrix. In this way, a small volume of fluid is held at the sliding interface, and a magnetic field is applied across this active region. An MR sponge device uses a much smaller quantity of fluid, implying a substantial weight savings over a fluid filled damper. This also implies a drastic cost savings due to the much smaller volume of expensive MR fluid required. The seals within a piston type damper are no longer needed, which eliminates the need for small tolerances and opens up great freedom in terms of the geometrical design of the damper. Just as with a fluid MR damper, stroke is limited only to geometric constraints. High force levels are possible if enough sponge area is present, with an approximate damping force of 5 Newtons per cm^2 of sponge area. Unfortunately, only a semi-active suspension is possible with an MR sponge damper (rather than a fully active suspension which would be capable of generating force). Just as with an MR fluid damper, an automotive application involves small displacements, which do not lend themselves well to variable damping control. [Carlson, 1999]

Piezoceramics also show significant limitations. In order to realize significant displacements at the required level, a stack of many thin wafers must be used. Even with a stack arrangement, stroke remains low (tenths of a millimeter). Force is plen-

tiful (10,000 Newtons is easily feasible), so one can envision simply levering the stack assembly in order to increase displacement at lower force levels. This necessitates a complex and most likely heavy arrangement of lever arms and linkages, in addition to a heavy actuator.

The piezoceramic story takes a turn for the better with the innovative arrangement of a prestressed layer of PZT sandwiched between two curved plates. Such devices are made by Face–International Corporation, and known as THUNDER© actuators. The name derives from the attempted acronym "***Th***in Layer ***Un***imorph Ferroelectric ***D***river and S***e***nso***r***." The actuator is constructed of a bottom layer of stainless steel, a middle layer of PZT, and a top layer of aluminum, all sandwiched together. The layers are assembled at an elevated temperature and bonded together. When later allowed to cool, the difference in coefficients of thermal expansion between the materials causes the sandwich to take on a curved shape and prestresses the bonded assembly. This configuration results in surprisingly large displacements for a PZT actuator. The advertised displacement of the devices ranges as high as 7 millimeters for a model TH 7–R device used in a cantilevered configuration. Blocked force for this same model is 133 Newtons. This particular actuator weighs in at 18 grams, which is perhaps the lightest among the actuators considered here. Although force is somewhat on the low side based on the MATLAB model, THUNDERs have the advantage of being stackable, one atop another, to increase force level. Physical dimensions of the TH 7–RA are 96.5mm length, 71.1mm width, and 0.635mm thickness, with an arch height of 9mm (3.8in by 2.8in by 0.025in, 0.354in height). A sample device is pictured in Figure 4.2.

4.2.2 Material Selection

For the options considered here, two of the smart materials stand out as plausible candidates. The biggest design consideration is weight, and so the magnetorheological sponge damper and THUNDER actuator are the most appealing candidates. The main limitation of the MR damper is that it is not really an actuator, but merely a variable damping rate device. The PZT actuator can generate whatever force level

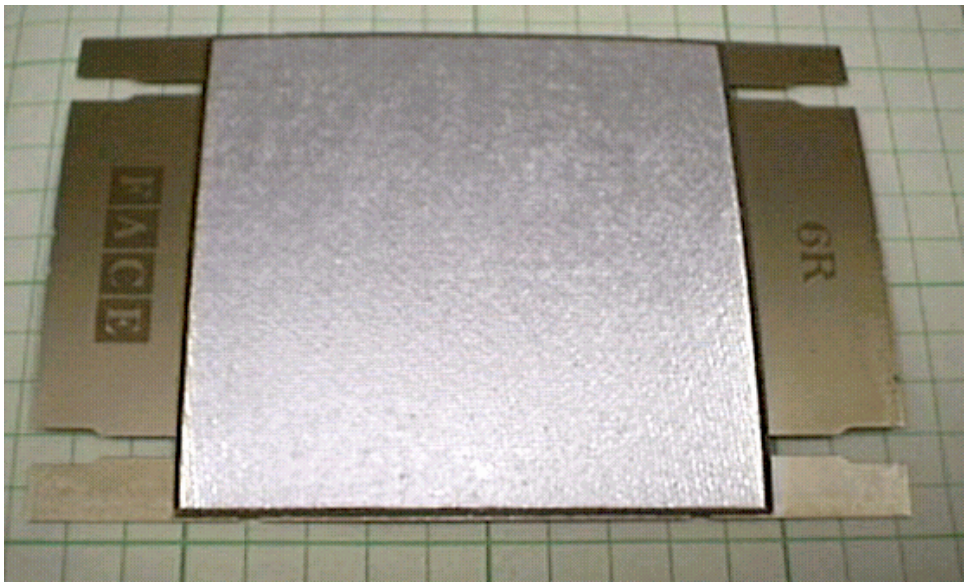

Figure 4.2: Thunder Actuator Model TH 7–RA (3.8" x 2.8")

asked for (within its own limitations) at any time; the sponge damper develops a force opposing velocity and so is dependent upon the instantaneous relative velocity of the system. Furthermore, an MR sponge device for this application would require a large sponge area, which might be difficult to fit beneath a seat. The small, flat shape of the sandwiched PZT device is very appropriate for the geometric constraints beneath a car seat. For all of these reasons, the THUNDER actuator is selected as the actuator of choice for the active seat suspension system.

4.2.3 Piezoceramic Actuator Implementation

While the prestressed piezoceramic devices exhibit many enticing features and specifications, the actual implementation into a mechanical system is far from simple. Displacement is developed by a change in radius of curvature, with displacement measured at the center of the arc. This vertical motion also carries with it a corresponding horizontal movement of the endpoints. We have three distinct load–bearing points — one at the center and one at each end. In the horizontal direction, one of these must be constrained to keep the system in place, but the other two *must* be allowed to slide freely in order to accommodate the length change that accom-

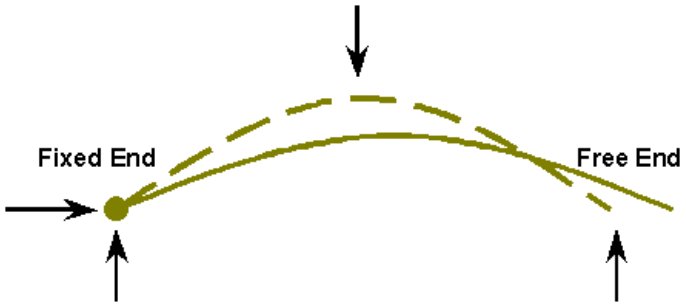

Figure 4.3: Actuator Endpoint Movement Under Deflection

panies change in curvature. At the same time, all three attachment points must be constrained in the vertical direction, since this is the direction in which we wish to generate force in our system. Figure 4.3 demonstrates the situation, with the arrows representing the free body reaction forces required in the mounting scheme. Ideally, the horizontal motion at the non-fixed points should have zero friction; in the real world some sliding friction will exist and detract from the amount of force we achieve in the vertical direction. Furthermore, the fixed end must allow rotational movement due to the small angular change that accompanies deflection.

In our application, two actuators will be stacked on top of one another in order to increase force level. A thin Teflon sheet placed between the two devices serves to electrically isolate the oppositely charged surfaces, and also to permit the small amount of sliding that will occur. The midpoint of the actuator will be attached to the seat frame, while the endpoints will be attached to the floorboard, with a static weight supporting spring in parallel with the actuator assembly.

In order to fix one end in both the vertical and horizontal directions while still allowing rotation, the devices are attached to an aluminum bar which is held in place by two spherical bearings. Two Igus pillow blocks, model JAM-10, are used due to their light weight, low profile, electrical isolation, and low cost. The pillow blocks can be seen in the actuator as pictured in Figure 4.4, which shows the actuator components and the assembled actuator.

The opposite end is also attached to the vehicle floorboard, with constraint in

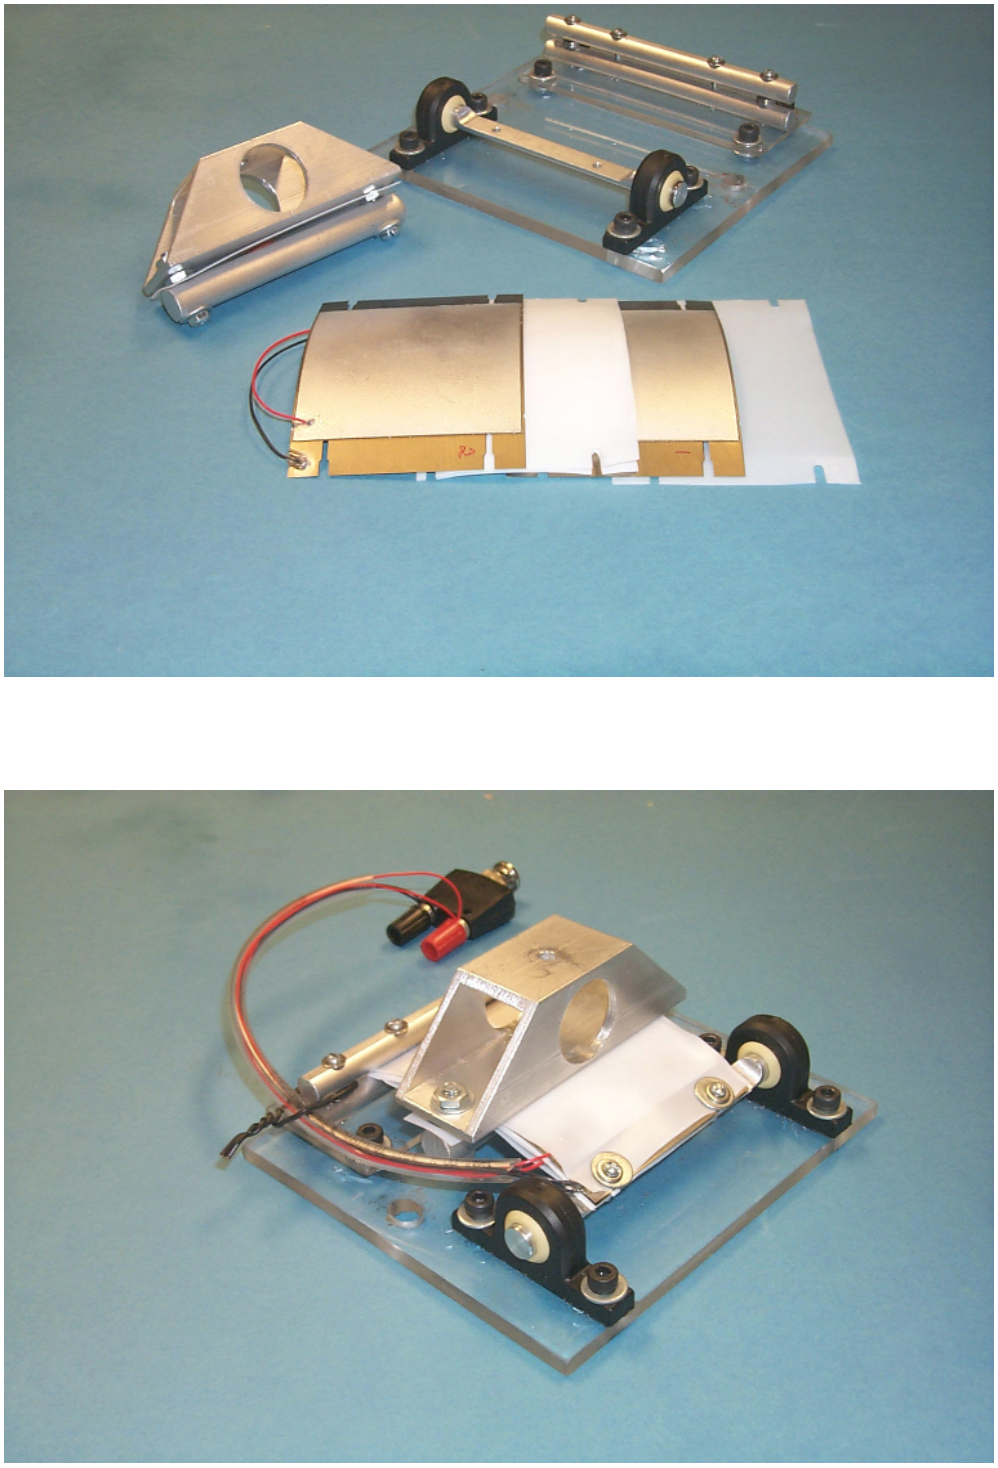

Figure 4.4: Actuator Assembly

Table 4.1: Actuator Mass Breakdown

component	quantity	grams each
Thunder Device	2	18
Pillow Block	2	10
Pillow Block Crossbar	1	18
Fixed Crossbars	2	22
Trapezoid	1	114

Complete Assembly: 488 grams

the vertical direction only. The THUNDER actuator stack is held in place between two round aluminum rods. Clearance is such that sliding is allowed in the horizontal direction, while clearance in the vertical direction is near zero. Thin Teflon layers are again employed in order to allow sliding of the ends of the piezoceramic devices with minimal friction. The two ends of the assembly are fixed to an electrically isolating lexan base.

A similar arrangement is used at the point of attachment to the seat frame, where we again desire constraint only in the vertical direction. A square aluminum tube, cut to a trapezoidal shape, is used in order to allow ease of assembly and disassembly during testing, with another round aluminum bar underneath so as not to interfere with the curvature of the actuator. The Teflon sheets needed to allow sliding at the ends are continued across the top and bottom surfaces of the actuators, easing horizontal movement of the trapezoid and providing electrical isolation.

In the as–tested configuration, the complete actuator assembly weighs in at a total of 488 grams. A breakdown of selected components of the assembly is given in Table 4.1. Revisiting the goal of having an actuator that is lighter than foam, we can compare this 488 gram mass to the total foam mass of 2,660 grams. Our two THUNDER actuator assembly is applied to a quarter seat system (one quarter of the total seat and passenger mass), meaning that in order to actuate the whole seat, we would require a total of eight PZT devices to achieve the necessary force levels. Two options exist: use the as–tested assembly at each corner of the seat, or stack two more THUNDERs onto the assembly and use one assembly at the front and one

at the rear of the seat. The former results in a total actuator mass of 1,952 grams, while the latter results in 1,048 grams (2 times the assembly plus two additional 18 gram PZT devices). Even with a pair of static weight springs in parallel with each actuator, at 50grams per spring, the latter active suspension system tips the scales at 1,248 grams, or 47% of the total seat foam mass of 2,660 grams. Of course, this does not include the additional mounting hardware that may be required in an actual in vehicle implementation, nor does it include the added weight of a controller or power amplifier.

The total cost of the prototype actuator is $340. This includes material and machining, with the two THUNDER devices costing $97 each. Of course, the cost of such a device would be significantly lower in a mass–produced form. The author chooses not to compare the prototype cost to the cost of the foam cushioning that is typically used in an automotive seat, for obvious reasons.

4.3 Control Design

Efficient use of the whatever actuator forces and strokes are generated requires an appropriate control strategy. Consider again a system driven by base excitation. The displacement response of the suspended mass excited by a force input is shown in the frequency domain in Figure 2.2. The resonant peak is a major problem in that when the system sees an input at or near the resonant frequency, it will actually *amplify* this signal. The goal is to isolate vibration, and so any region of amplification should be eliminated.

To control the system, a gain or set of gains based on measured signals is used. The seat suspension system will most likely have two directly measured variables: acceleration of the vehicle floorboard, and acceleration of the seat itself. From these, the velocity and displacement of each location can easily be obtained by integration.

Figure 4.5 again shows the base excitation (force to displacement) transfer function, but now also shows the additional cases of force to velocity and force to acceleration transmissibility. The shape of the velocity transmissibility curve is par-

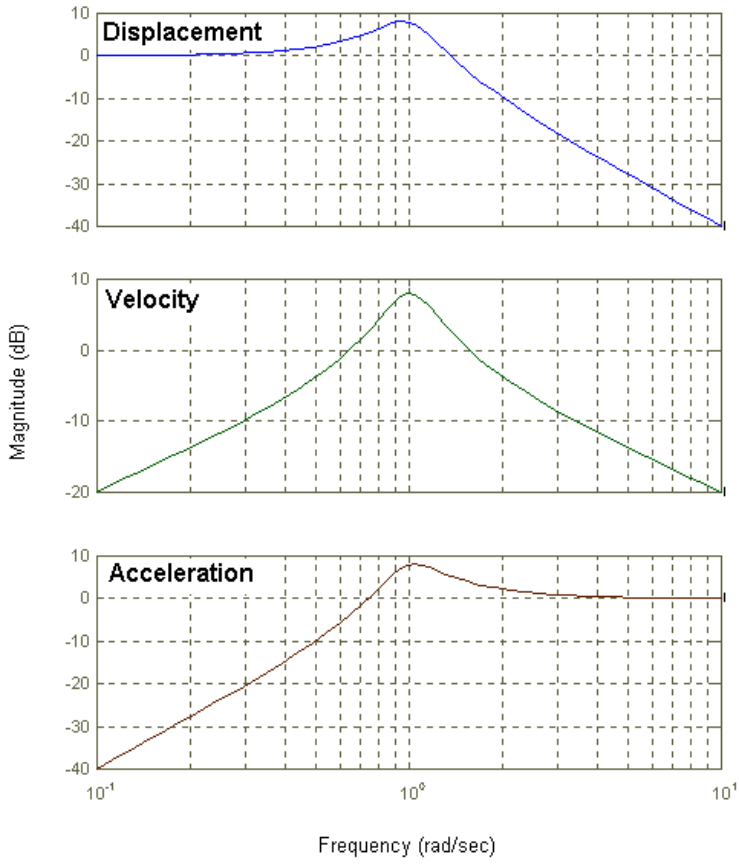

Figure 4.5: Force To Displacement, Velocity, And Acceleration

ticularly interesting in that it has a magnitude greater than one only at frequencies surrounding the resonant peak. Hence, a gain based upon the velocity of the mass will be of particular interest in the control strategy.

Chapter 5

Experimental Analysis

5.1 Baseline Seat Testing

In order to understand the real world application of our actuator, we must have an accurate data set for the transmissibility of the car seat in question. A test stand was constructed to excite the seat in the vertical direction, allowing measurement of the passive seat's characteristics. This test stand was also later used for actuator testing and control algorhythm testing.

5.1.1 Test Stand Construction

The stand is constructed using extruded aluminum components which allow ease of assembly and are extremely modular. The general configuration of the test stand is a four post structure with a vertical motion table actuated by an electrodynamic shaker. Linear bearings at each post constrain motion to the vertical direction. A drawing of the stand is shown in Figure 5.1 (dimensions in centimeters).

An MB Dynamics PM–50 shaker is used to drive the system, and is capable of 222 Newtons (50 pounds) of force and 12.7mm (0.5in) of displacement. The shaker is strong enough to provide the dynamic acceleration of the seat and passenger mass, but is not strong enough to do so while also supporting the static weight of the seat and passenger (approximately 120kg). To offload the shaker, low rate springs are used to

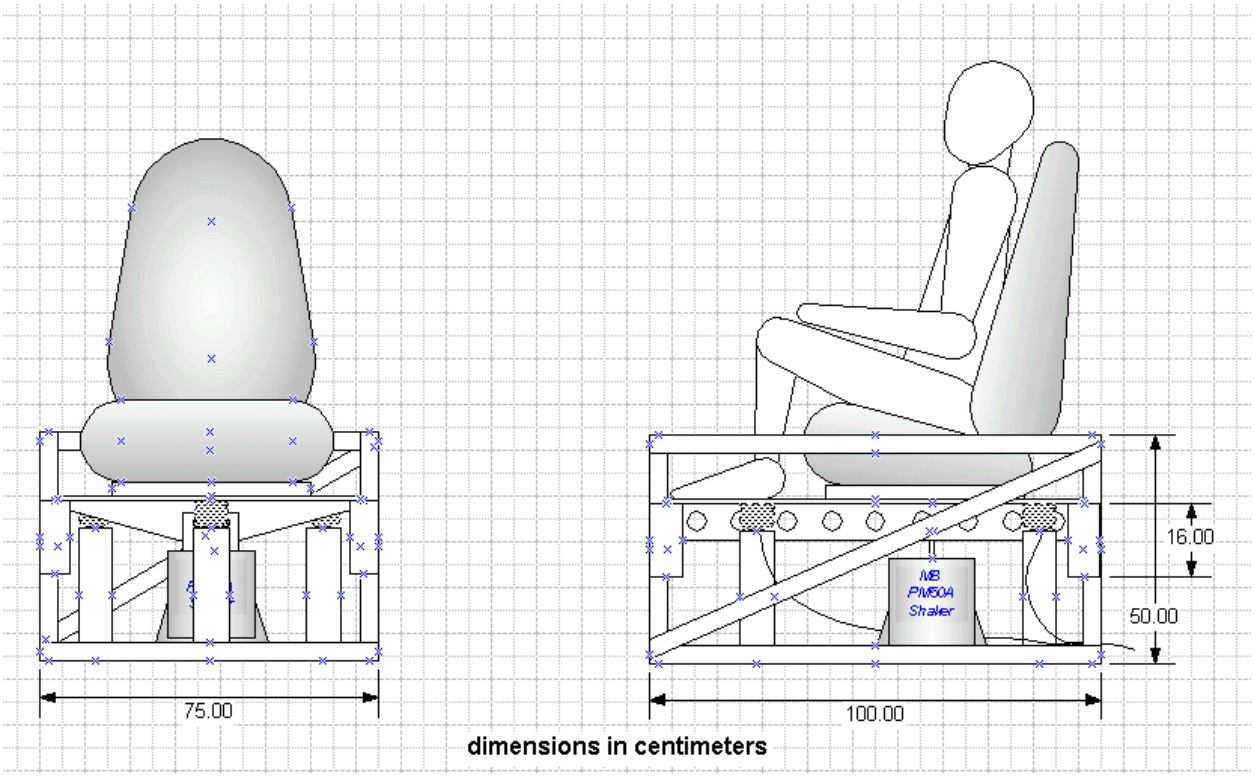

Figure 5.1: Test Stand Drawing

support the static weight. Four extension coil springs with a 2.4 Newtons/millimeter rate are used, and so the test stand differs slightly from Figure 5.1, which shows an air spring design.

Several pictures of the test stand are shown in Figures 5.2 and 5.3. The luxurious leather trimmed seat donated by Johnson Controls is seen attached to the table, and two bags of concrete are placed atop the seat to simulate the distributed mass of a human passenger. Figure 5.3 shows a closeup view of the PM–50 shaker which is used to excite the table.

5.1.2 Seat Transmissibility Data

In order to determine the transfer function of the passive passenger seat, the test stand is used in conjunction with a Tektronix model 2630 Fourier Analyzer. The analyzer has the capability of generating a bandwidth limited random output signal, which feeds into the amplifier for the electrodynamic shaker. For vertical transmissi-

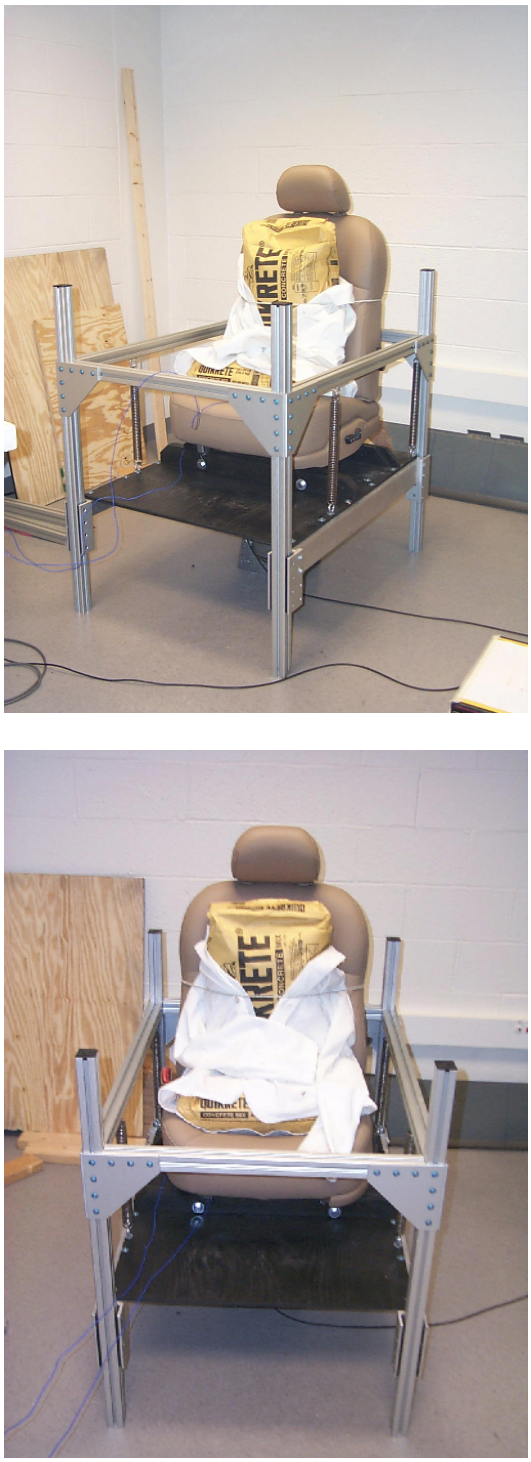

Figure 5.2: Test Stand Pictures

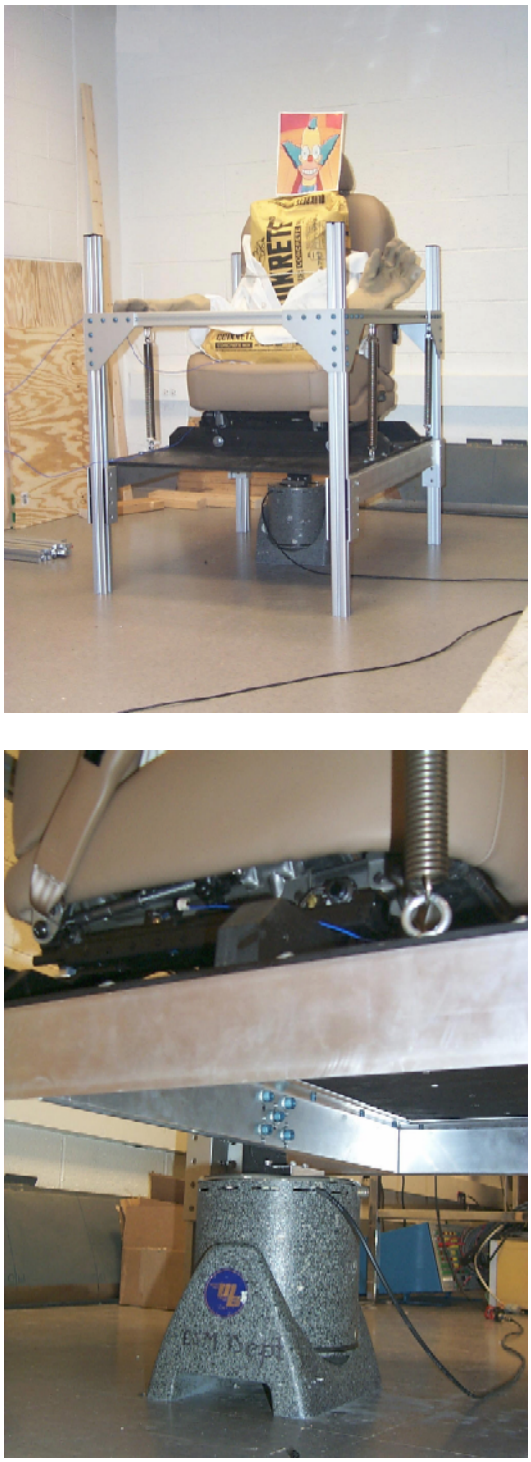

Figure 5.3: More Test Stand Pictures

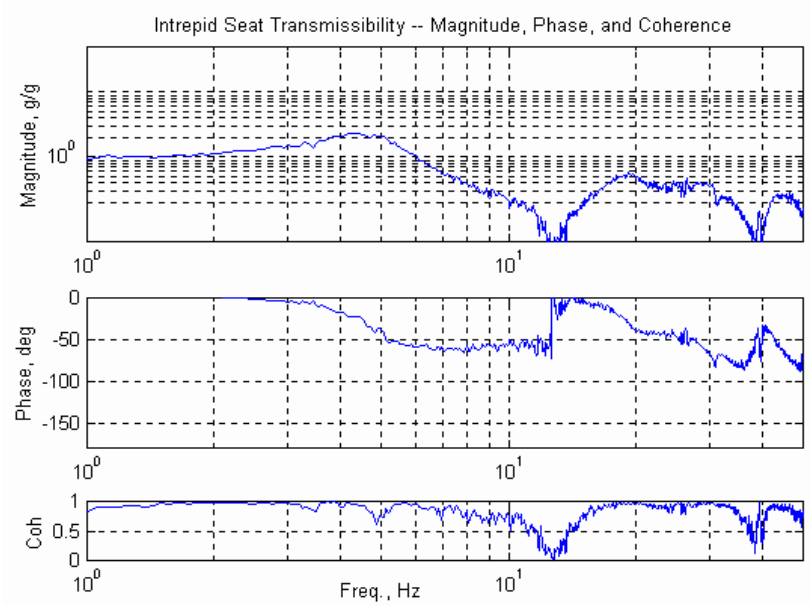

Figure 5.4: Vertical Transmissibility

bility, an accelerometer placed on the table directly above the shaker attachment point measures the input acceleration, and an accelerometer placed at the seat/passenger interface, again directly above the attachment point, measures the output acceleration after the seat has done its best to isolate the input vibration. Frequency range is limited to a maximum of 50Hz. The analyzer is used to process the signal into a frequency domain transmissibility dataset.

The measured vertical transmissibility is shown in Figure 5.4. For this dataset, an actual human dummy (grad student) was used for the passenger mass. This insures proper mass value and mass distribution. The resonant peak is seen to occur at 4Hz, with an isolation frequency (frequency at which magnitude ratio drops below one) of 6Hz. The resonant peak has relatively low damping, which can be calculated from the data as ζ=0.2 using the half power method ($2\zeta = \frac{\omega_2 - \omega_1}{\omega_{res}}$, where ω_{res} is the resonant frequency and ω_2 and ω_1 are the frequencies on either side of resonance at which the amplitude has fallen to 3dB below its maximum value).

Figure 5.4 demonstrates that there is definitely some room for improvement in seat characteristics. The underdamped response, while resulting in improved high frequency isolation, causes a defined resonant peak with a magnitude ratio greater than

two. Hence, at an input frequency of 4Hz, the acceleration level at the seat/passenger interface will be *twice* the level of the input acceleration. Recalling the human sensitivity data in Figure 2.8, we see that the resonant peak occurs in the very region where humans are *least* tolerant of vibration. This is necessary given the fact that the primary suspension typically has a resonant frequency in the 1-2Hz region. The seat must be designed so as to avoid lining up with other system resonances, which implies a resonant frequency slightly above the vehicle mass resonance. An active system, since its response properties can be varied with frequency, would ideally eliminate any magnitude ratios greater than one while still retaining good high frequency isolation characteristics. An active seat system could be designed to maintain a highly damped natural frequency at 4Hz while at the same time having low damping at higher frequencies. Recall the MATLAB output of Figure 3.4, which demonstrated these very characteristics.

We can also measure the degree to which the seatback allows vibrational inputs to be transmitted to the passenger. Figure 5.5 shows three datasets measuring the transmissibility of the seatback and seatback foam. Accelerometer locations are as shown in Figure 5.6. For the first dataset, labeled "to back frame" (#1), the accelerometer measuring input acceleration is placed on the shaker table, directly above the shaker attachment point, measuring acceleration in the vertical direction. The output accelerometer is placed on the internal frame of the seatback, measuring acceleration perpendicular to the seatback. The next dataset, labeled "to back cushion" (#2), holds input accelerometer location constant and moves the output measurement to the interface of the seatback and human passenger. The third dataset, "through cushion" (#3) holds the output accelerometer at the seatback/passenger interface, with the input accelerometer location on the internal seatback frame.

With these three datasets, we gain an understanding of how vertical acceleration of the vehicle floorboard is transmitted to the back of a seated passenger. From #1, we see that the seatback is not perfectly stiff in its attachment to the rest of the seat, and exhibits resonant frequencies at 20 and 40Hz. Dataset #3 measures the transmissibility solely due to the seatback foam. Combining the two yields the overall

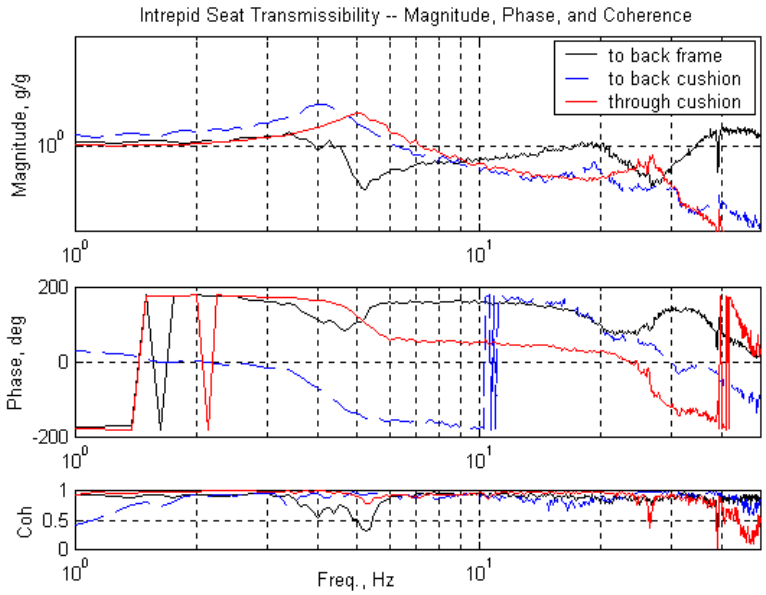

Figure 5.5: Transmissibility Through Seat Back

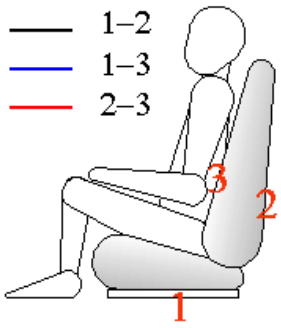

Figure 5.6: Accelerometer Locations

situation in #2, which is the total effective transmissibility from a vertical floorboard input to the back of a seated passenger. Although the foam by itself exhibits resonance at 5Hz, the overall resonant frequency is 4Hz with an isolation frequency is 6Hz. The response of the seatback is strikingly similar to the response of the seat bottom. Thus, the seatback would also benefit from an active system which could selectively filter out vertical inputs in the 4–10Hz frequency range, where humans are most sensitive to vibration.

5.2 Actuator Testing

The next step is to characterize the performance of the THUNDER actuator assembly. The real world implementation of the THUNDER device into a usable assembly undoubtedly causes some small deviation from its completely unconstrained performance.

5.2.1 Free Displacement

Free displacement is easily measured by applying a voltage signal to the actuator, in the absence of any external load, and measuring displacement. A sinusoidal input voltage is applied at frequencies of 1Hz, 5Hz, and 10Hz. Peak to peak voltage is varied in 50V increments from zero to 600V. The input voltage is limited to 600V peak to peak in this case since the THUNDER model TH–7R is limited to +600V and -300V. A Polytech OFV 303 laser vibrometer is used to measure displacement. Test setup is as shown in Figure 5.7, with the first picture showing the actuator and laser vibrometer setup, and the second picture showing a closeup of the actuator while being measured. Results are shown in Figure 5.8. The maximum stroke achieved is 2.7mm, as seen in the figure. The small degree of scatter in the data implies the absence of any frequency dependent trends.

In our quest for more displacement, we realize that a 600V peak to peak sinusoidal signal ($\pm$300V) falls short of the piezoceramic's capability of a positive input voltage of 600V, due to a limit on the negative side of -300V. More displacement

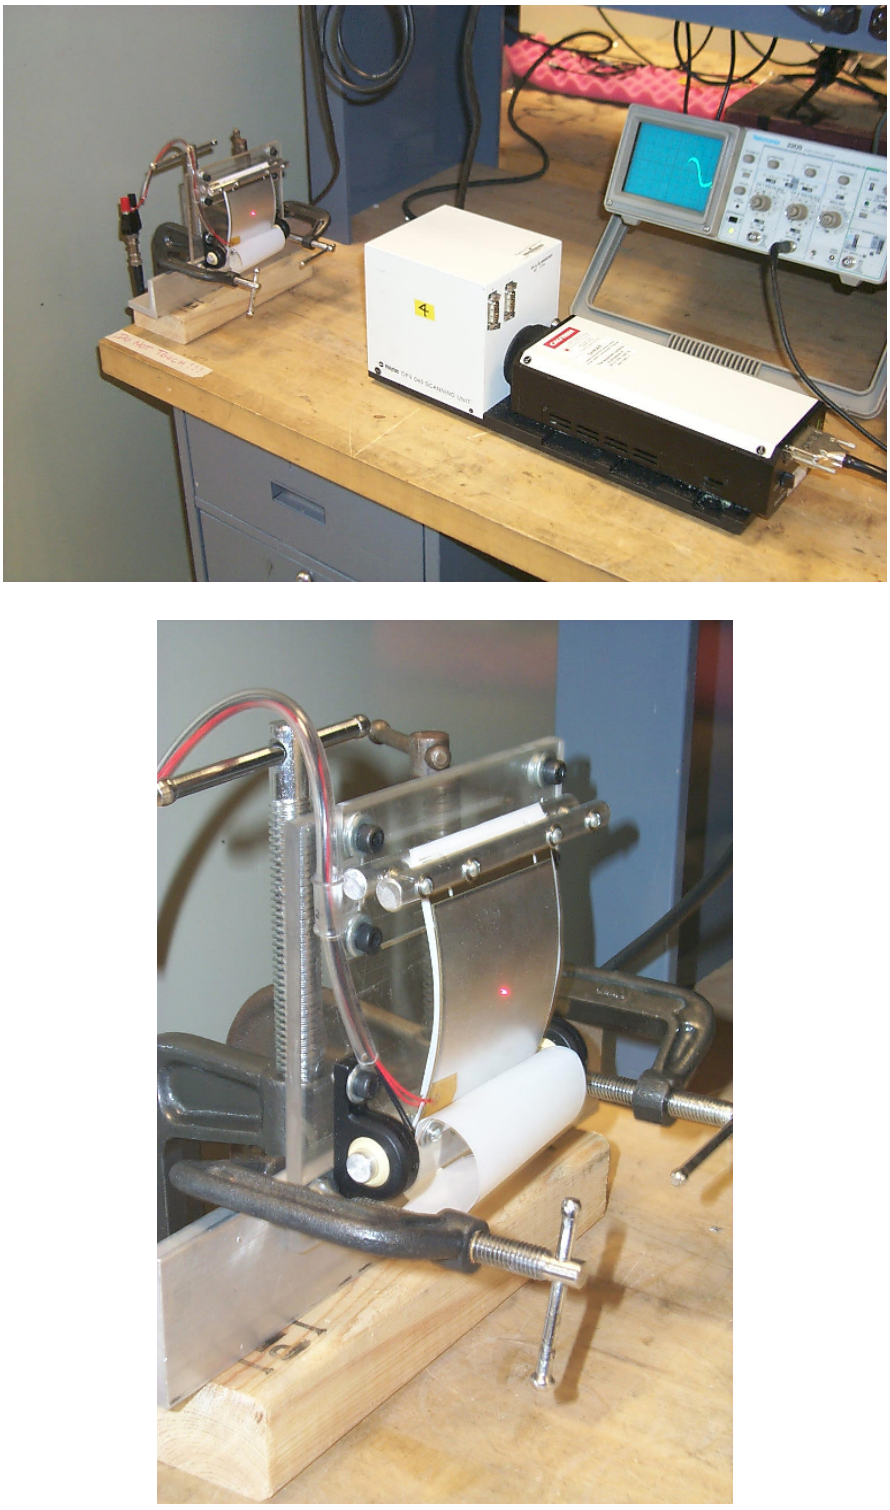

Figure 5.7: Experimental Setup For Free Displacement Measurement

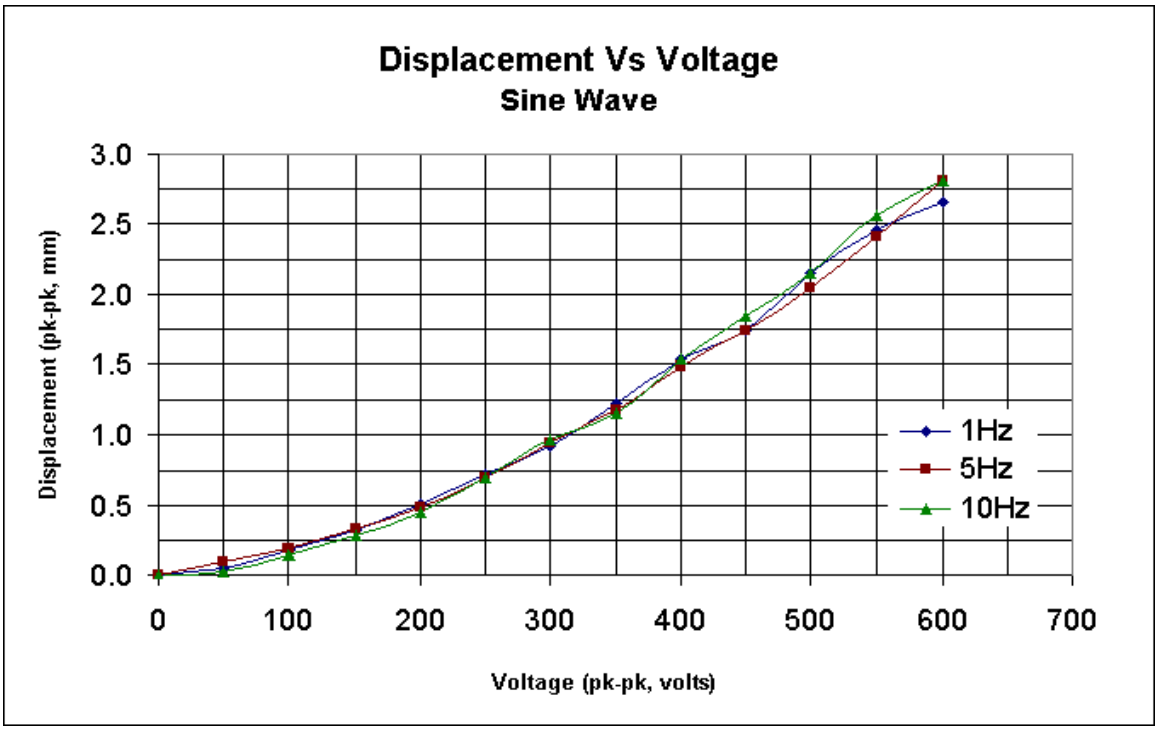

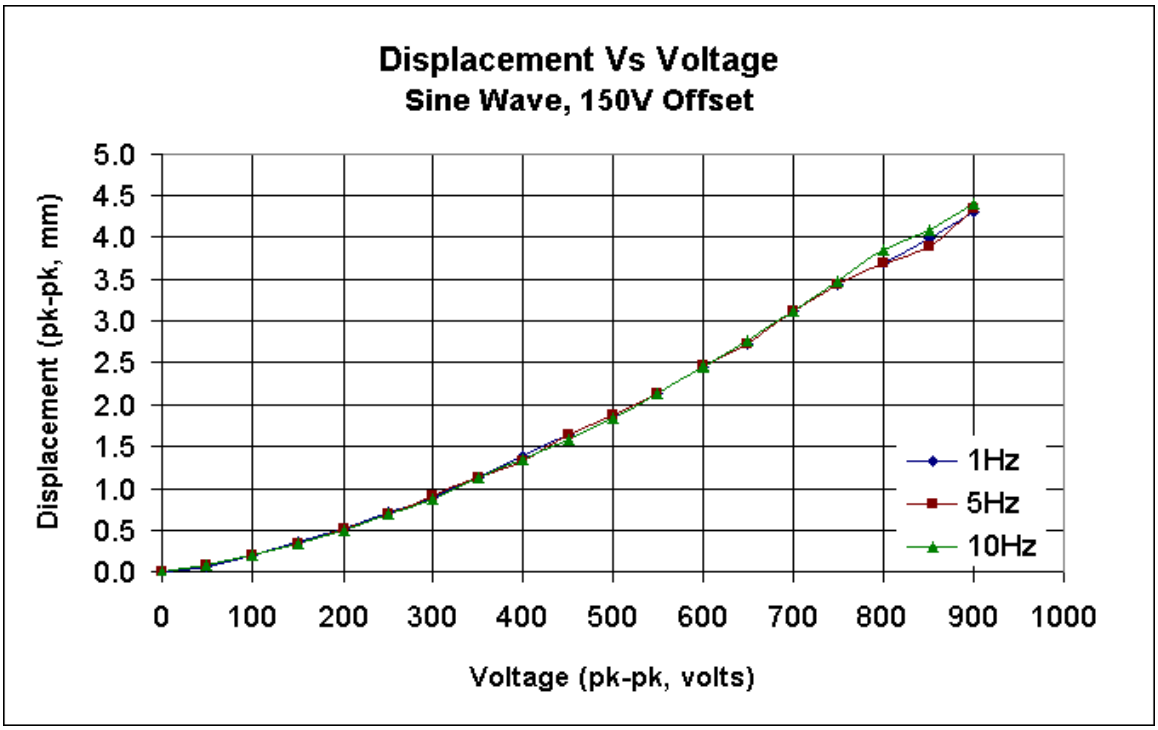

Figure 5.8: Actuator Displacement Vs. Voltage, No Load

can be achieved with application of a waveform that varies between -300V and +600V. Free displacement is measured again, this time with a +150V DC offset added to the sinusoidal signal, allowing a maximum of 900V peak to peak. Results are shown in the second graph of Figure 5.8. Maximum stroke is now increased to 4.4mm peak to peak. Both datasets show a reasonably linear response beyond 300V, with a gain of 0.00625mm/V.

5.2.2 Blocked Force

At the other extreme of the operating envelope is the case of zero displacement with maximum actuator force, or blocked force. The actuator is constrained to prevent any displacement, energized, and blocked force is measured. Experimental setup is as shown in Figure 5.9, with a steel plate used to constrain actuator motion. At maximum actuator force, the deflection of the plate is calculated to be less than one thousandth of an inch. A quartz force transducer is affixed between the actuator and the plate

A sinusoidal input voltage signal is again applied. The two cases of 1) zero DC voltage offset and 2) +150V DC offset voltage are again examined. Peak to peak input voltage is varied from 0 to 600V or 900V, in increments of 50V. Data is taken for sinusoids of 1Hz, 5Hz, and 10Hz. The datasets are presented in Figure 5.10. For the zero offset case, a maximum of 200 Newtons peak to peak is realized. The addition of the voltage offset results in greater force generation with the increased voltage level, and shows a maximum of 320N peak to peak. Again, response is linear beyond the 300V level, with a gain value of 0.41N/V.

5.2.3 Force Versus Displacement

We can now combine the blocked force and free displacement data and characterize the actuator's performance in a force versus displacement realm, which is shown in Figure 5.11. A family of curves is shown for varying voltage levels. The blocked force and free displacement data are shown as the endpoints which fall on

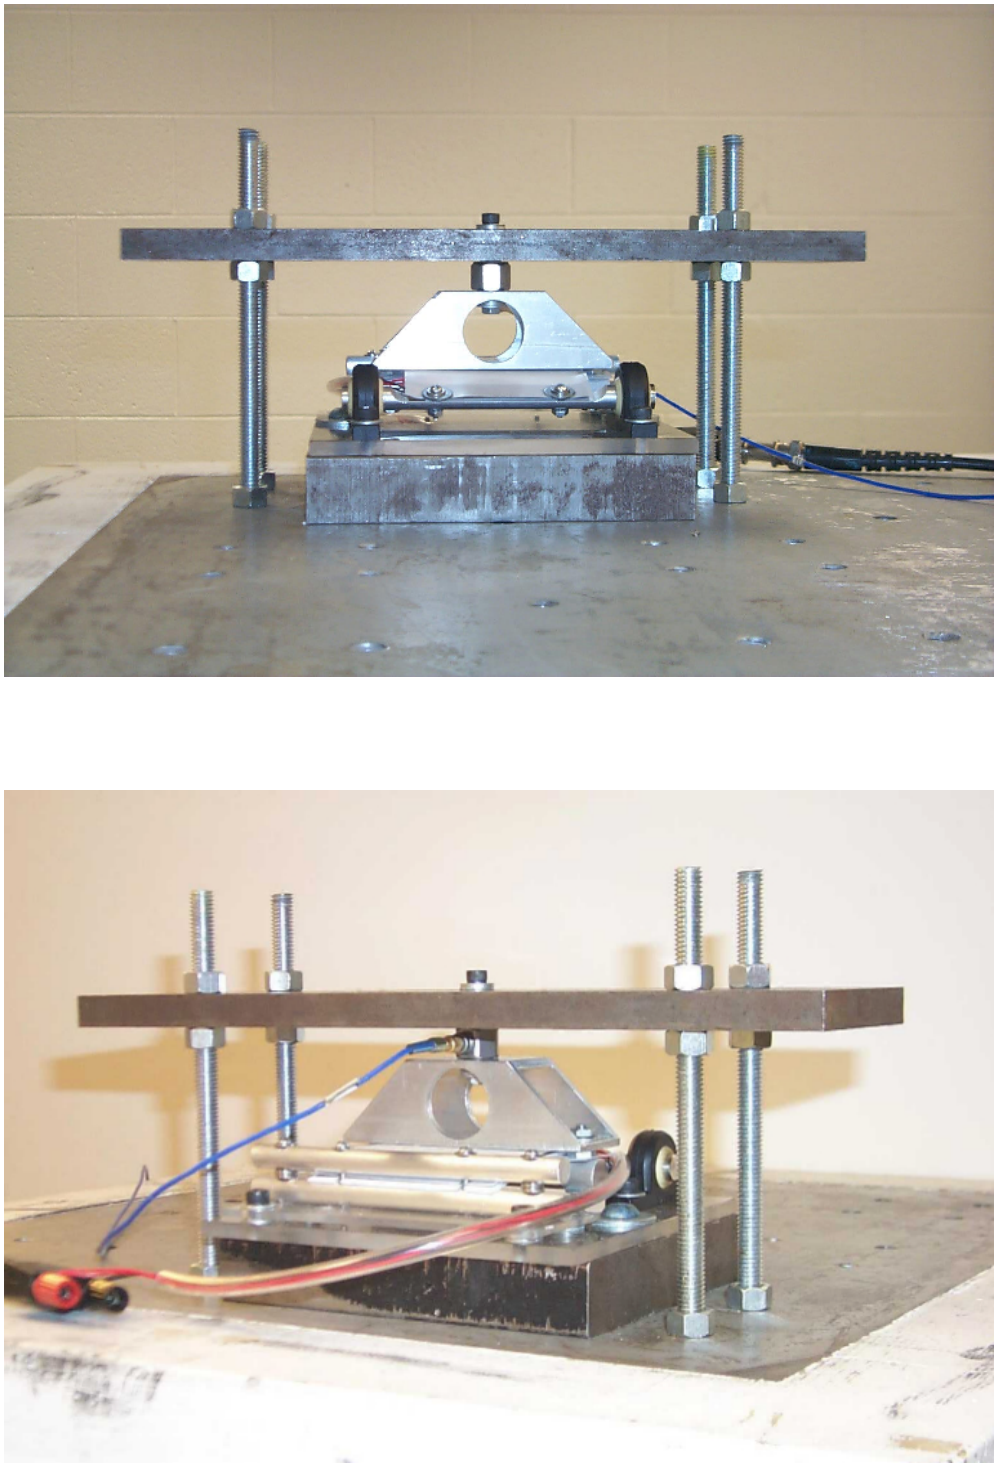

Figure 5.9: Experimental Setup For Blocked Force Measurement

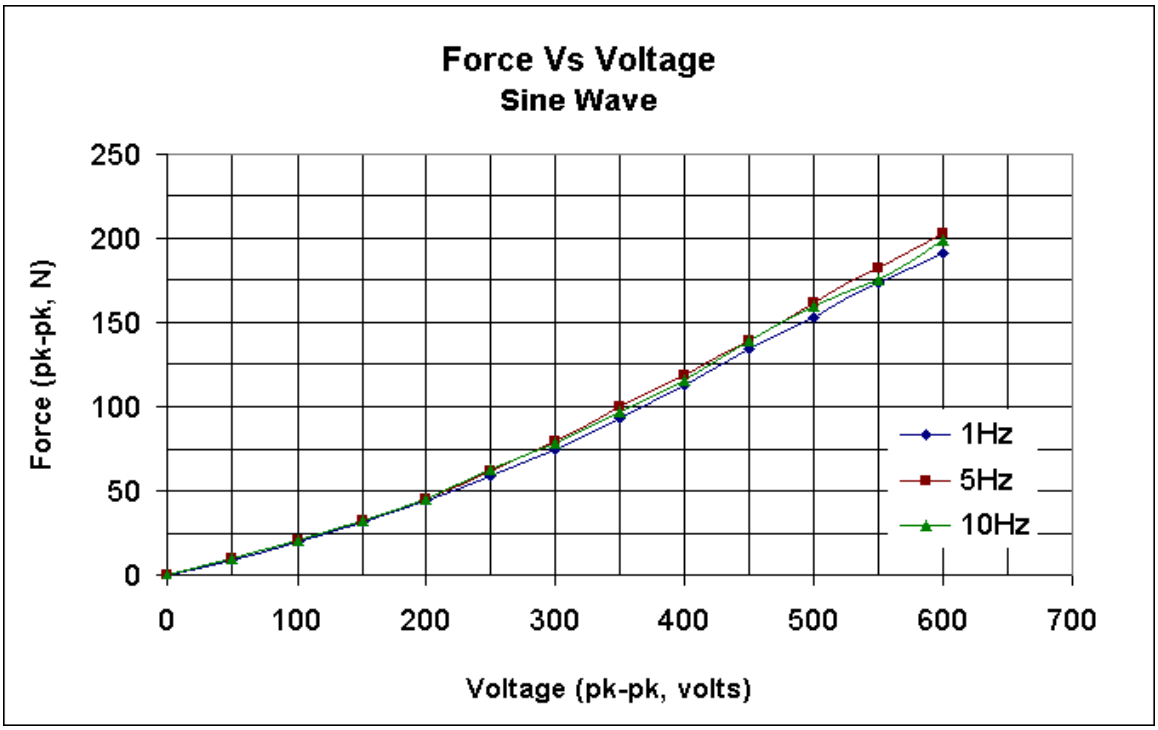

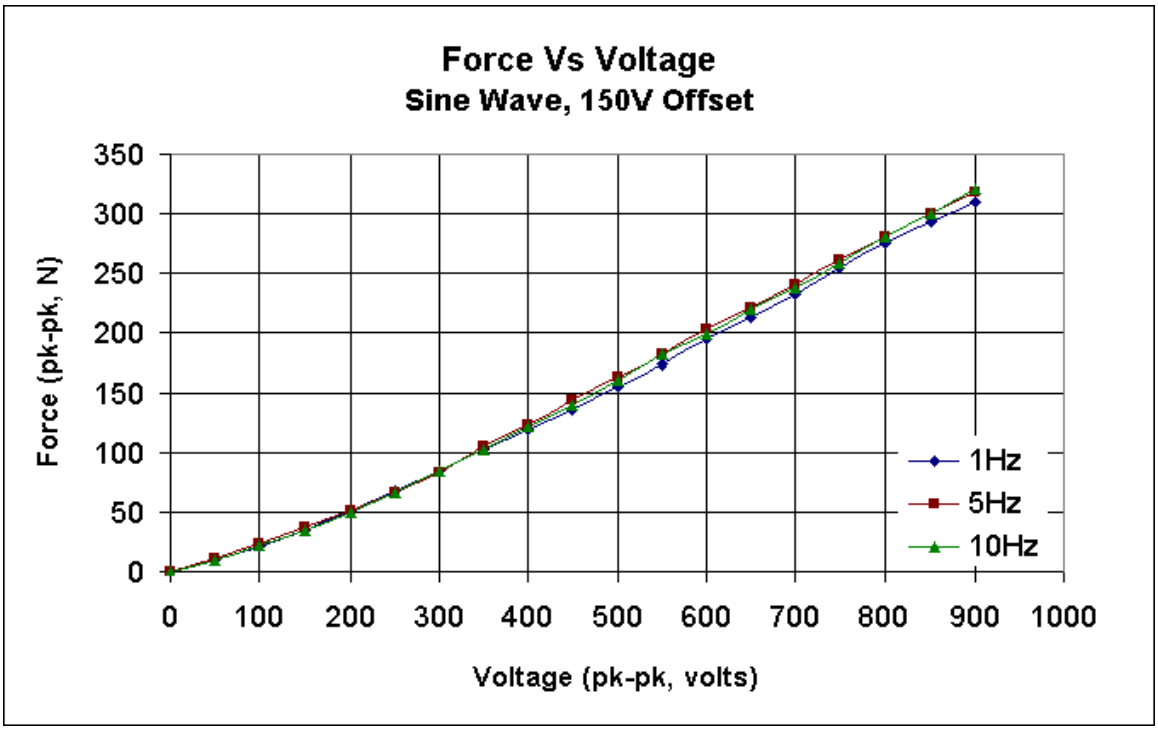

Figure 5.10: Blocked Force Vs. Voltage

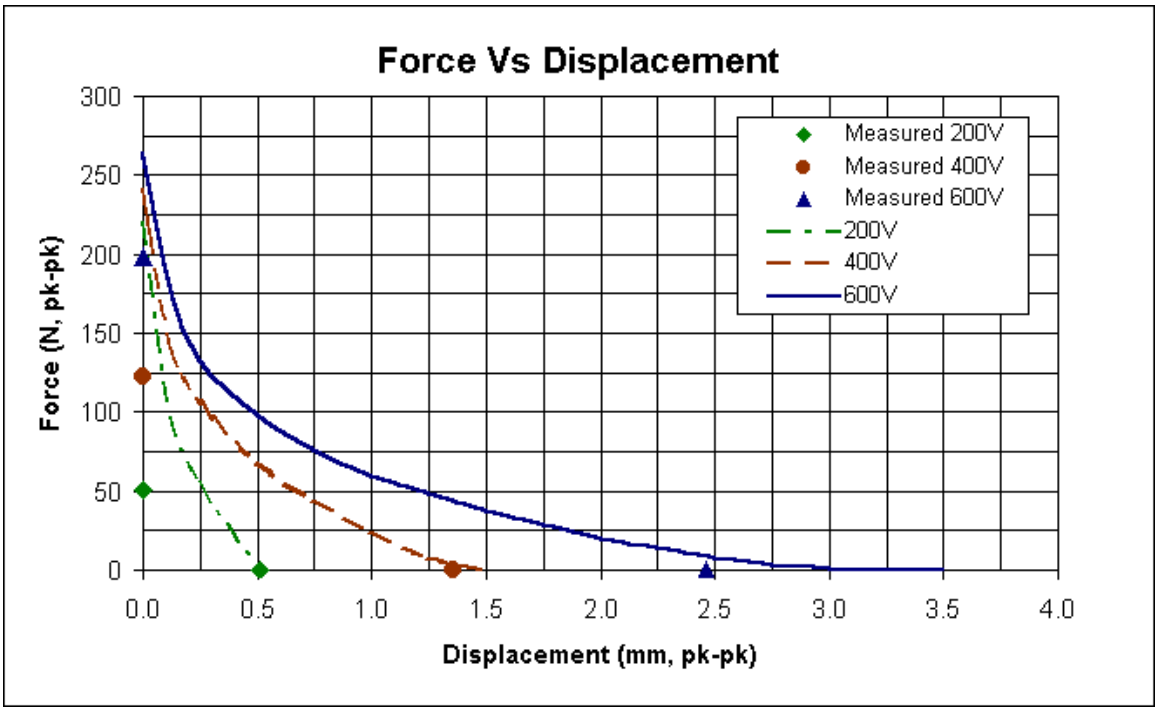

Figure 5.11: Force Vs. Displacement As A Function Of Voltage

either axis, while the complete datasets provided for the THUNDER actuators are plotted as continuous functions. The complete datasets provided for a single bare PZT device are doubled to match our setup of two devices stacked on top of each other. The force levels achieved with our configuration are noticeably lower than the data provided for a bare actuator, due to the inefficiencies imposed by our mounting system. The sliding and rotational friction at each mounting point, as well as the sliding friction between the two devices, subtracts from the total output force of the actuator assembly. It is interesting to note that even for a constant voltage, when displacement is varied, the force does not vary with a constant slope, as would be expected.

5.3 Dead Mass Setup

Before tackling a complete setup of an entire car seat with four actuator assemblies (one at each corner), the simplified case of a single actuator assembly controlling the vertical motion of a representative amount of dead mass can be examined. The

test stand is again used to measure the seat transmissibility, and the shaker is now used to force a simulated road input into the system. In this way, the table is vibrating just like a vehicle floorboard would when traveling down the road, with the actuator located between that floorboard and a mass which represents the seat plus a human passenger. Since only one actuator assembly is being tested, one quarter of the total mass is used. This amounts to 24.4kg, which is one fourth of the 21kg seat mass plus 77kg (170lb) of human mass.

The setup is shown is Figure 5.12. The orange colored steel mass is constrained to vertical motion by means of the linear bearings sliding along the aluminum posts on either side. The actuator assembly is directly below the dead mass, and is situated on an adjustable height platform. The two static weight supporting springs in parallel with the actuator are readily visible in the picture, as are the two accelerometers (one on top of the mass, one on the table surface directly below the actuator). The two accelerometers are used to determine the transmissibility through this prototype seat suspension system. The calculated natural frequency of the 24.4kg dead mass with two 12.3N/mm static weight springs is 5.05Hz. This does not account for the added spring rate of the THUNDER devices themselves, which will raise the resonant frequency slightly.

To produce a shaker table noise signal that mimics the floor of a vehicle driving down the road, dSpace is used in symphony with MATLAB's simulink environment. The dSpace module allows the input and output of voltage signals to and from the simulink environment, and so is ideal for not only producing the shaker table signal, but also for later efforts to send an appropriate control signal to the actuator.

For now, the focus is on making the shaker table shake like the floor of a vehicle. The complete simulink model is shown in Figure 5.13. The MATLAB code which initializes the simulink model parameters is given in Appendix B. The upper portion of the model shows a bandwidth–limited noise signal which is fed through two transfer functions. The first transfer function converts the noise from a broad spectrum (constant Power Spectral Density (PSD) level) noise to a signal that is representative of a real road surface. Gillespie's road profile analysis shown in Figure 2.5 provides

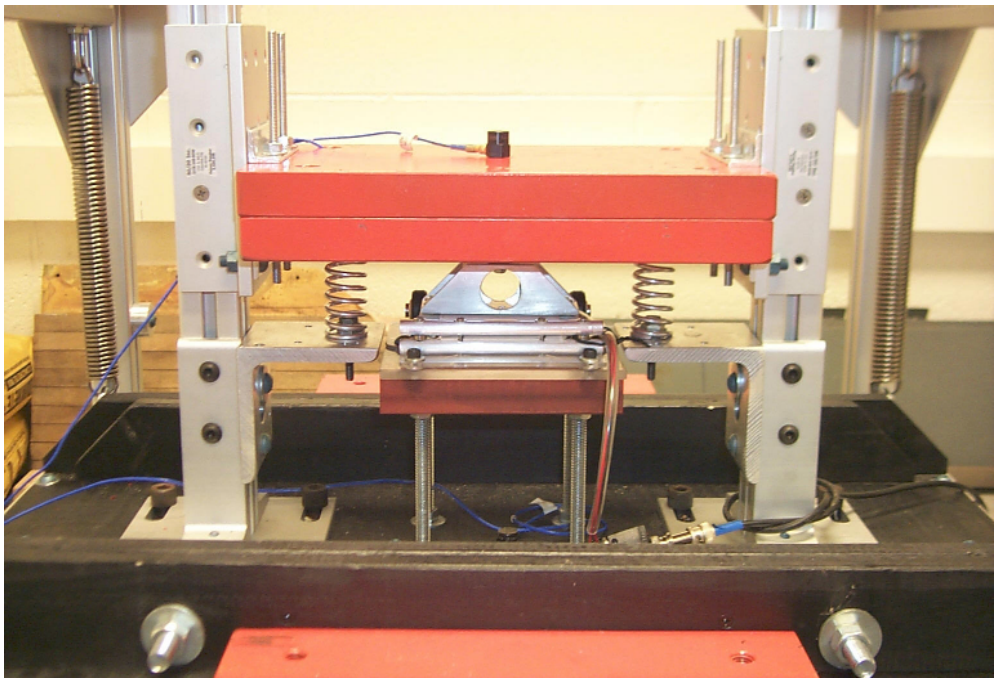

Figure 5.12: Dead Mass Setup on Test Stand

actual road data to emulate.

After achieving a signal similar to a road surface, that signal must also pass through the transfer function of a vehicle primary suspension to reach the level of the floorboard. A second transfer function, representing the primary suspension in state space form, is used to convert the road noise to a vehicle floorboard noise. The final signal is fed into the dSpace digital to analog converter, and the actual output voltage signal is fed into the amplifier driving the electrodynamic shaker.

The noise signals, as well as the measured acceleration levels, are shown in Figure 5.14. The first graph shows the constant level PSD (broadband) noise that is fed through the two transfer functions, and the second (upper righthand) graph shows the spectrum of the signal sent out to the shaker. The lower two graphs show actual measured data from the accelerometer on the shaker table surface and from the accelerometer on top of the mass.

With the shaker table shaking just like a car's floor would when driven down a typical road, the next step is to characterize the prototype seat suspension system. The transmissibility in the frequency domain is a good place to start, with an eval-

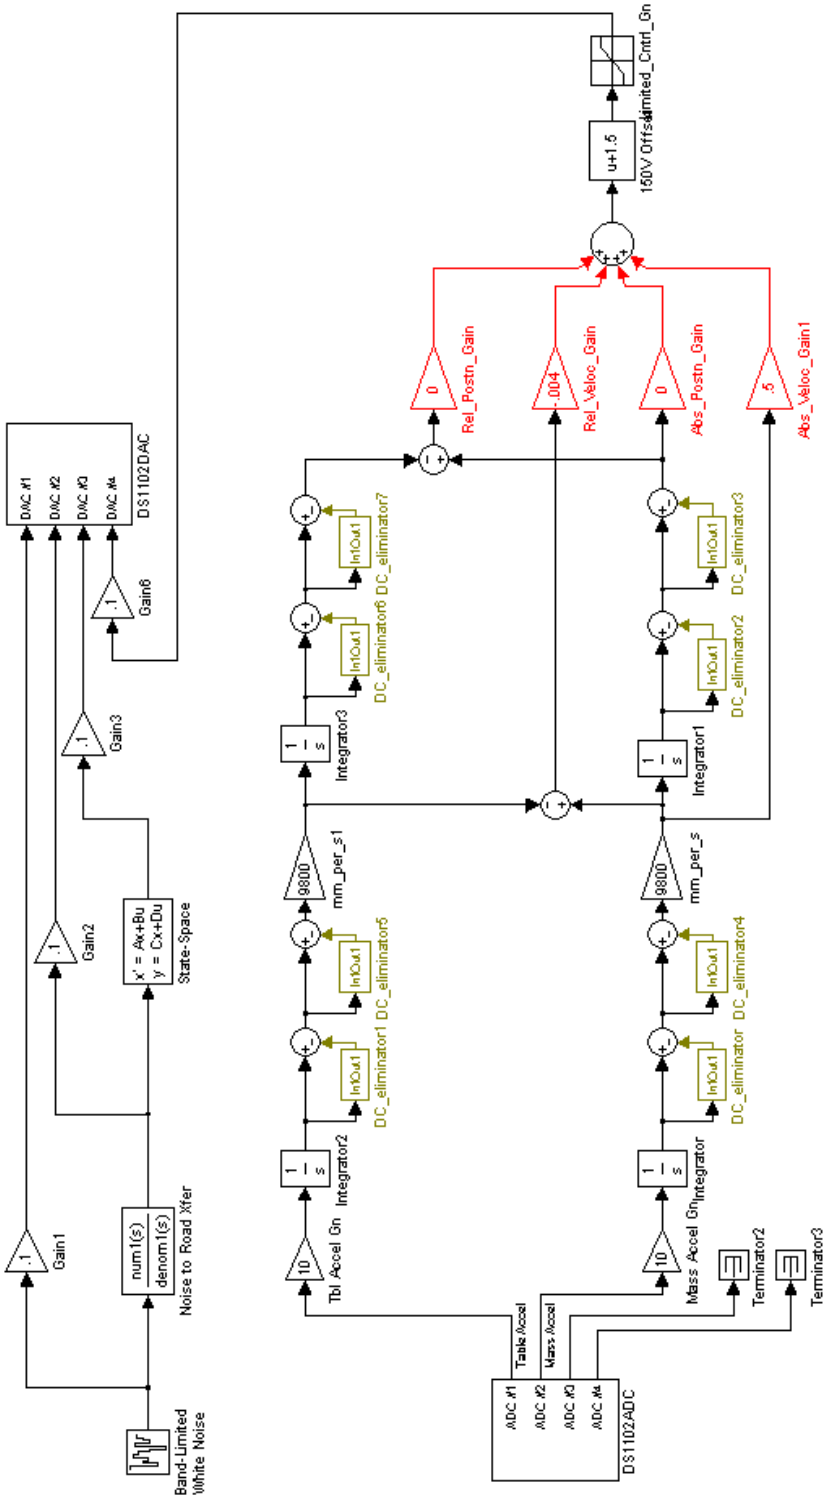

Figure 5.13: Simulink/dSpace Model

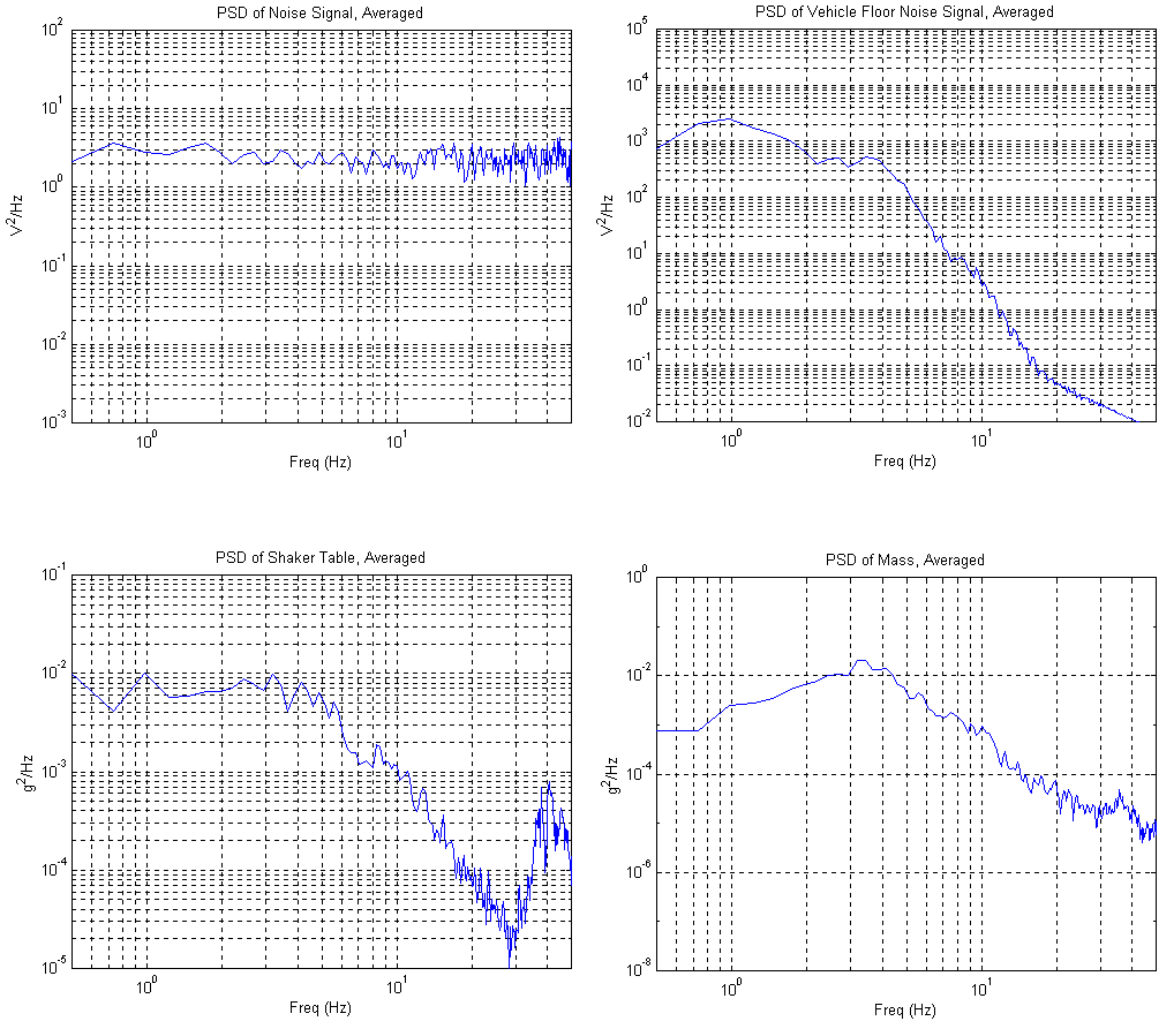

Figure 5.14: PSD Of Noise Signals, Measured Table Acceleration, and Measured Mass Acceleration With Representative Road Input

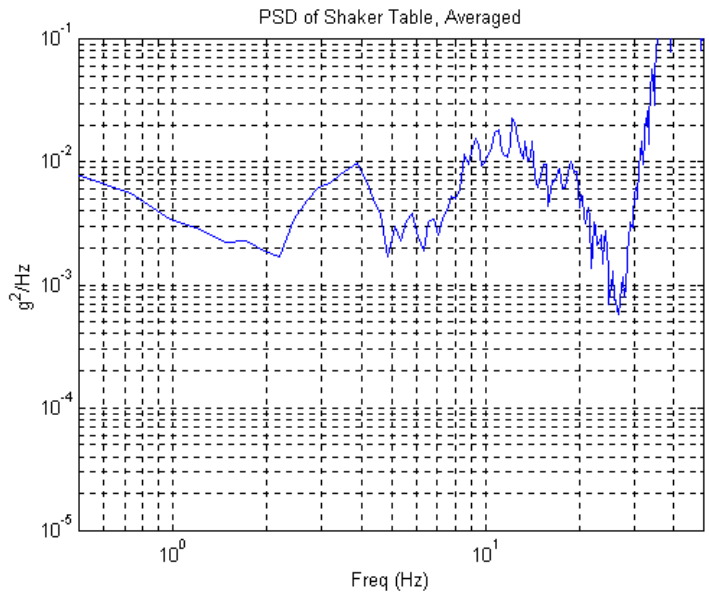

Figure 5.15: PSD of Shaker Table With Broadband Noise Signal

uation of the frequency response function using the shaker table acceleration as the input and the mass acceleration as the output.

Problems arise. After all of the efforts to achieve a representative input signal, the resulting input signal contains only lower frequency content. Upon taking the frequency response function (FRF), the coherence in the higher end of the frequency range is very low. The filtering has done such a good job screening out all of the higher frequency content of the noise signal that there is nothing left in the frequency range above 10Hz to excite the system. When the input signal is so small, any noise in the accelerometers, wires, or instrumentation becomes large in comparison, and results in low coherence. To remedy this, a more broadband input signal is used to excite the shaker table during the frequency domain analysis. The PSD of the broadband noise, as measured by the table mounted accelerometer, is shown in Figure 5.15. The shaker table now has an adequate input signal throughout the analysis frequency range. The dips in the 6Hz and 25Hz range will be shown to be attributable to the energy being absorbed by resonances of the table and dead mass system.

With the broadband input signal, a suitable FRF of the system is obtained, as shown in Figure 5.16. The resonance of the system is at 6Hz, with an isolation

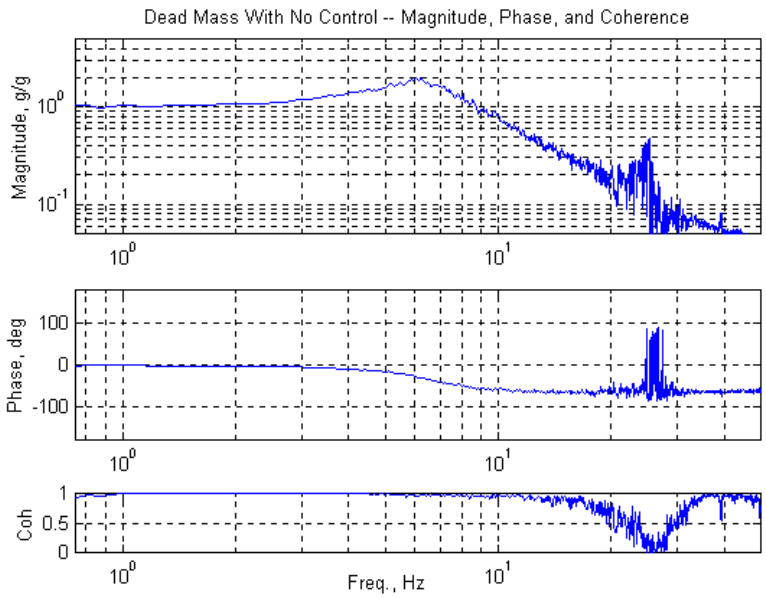

Figure 5.16: Response of Dead Mass With No Control

frequency of 9Hz. The poor coherence at 25Hz is due to a resonance of the surface of the shaker table itself. This was determined by moving the output accelerometer to various locations, and noting that the table surface itself resonates at 25Hz.

5.4 Control Experiment

Recalling the MATLAB modeling, four gains are under consideration for controlling the system: relative position of the mass, relative velocity of the mass (both with respect to the shaker table), absolute position of the mass, and absolute velocity of the mass. Only acceleration is being measured directly (by being fed into the dSpace analog to digital converter), but the signal can easily be integrated to obtain velocity and then integrated again to yield position. While integration tends to clean unwanted high frequency noise from the signal, any DC offset of the acceleration signal will produce unwanted drift in the velocity measurement. The simulink model in Figure 5.13 contains a DC offset filter immediately following each integrator to solve this problem.

Figure 5.17 shows control using only the absolute velocity gain, with a broad-

band input signal. The actuator does a more than adequate job of isolating vibration from the suspended mass. The gain is increased up to a level showing significant clipping at the maximum allowable actuator voltage of +600V and -300V. Note that the simulink model is set up to clip the control voltage to these limits. The resonant peak is eliminated by even modest gains, and as gain is increased further, isolation frequencies in the 2–3Hz range are easily achievable.

The reduction in RMS acceleration produced by the control used in Figure 5.17 is given in Table 5.1. The table also shows the reduction in RMS acceleration achieved by the MATLAB simulation for the modeled foam seat verses the modeled active seat, which produced an 80% reduction in acceleration level. The dead mass experimental setup did not include a test of actual seat foam, and so a direct comparison cannot be made for the passive cases, but the active cases can be compared directly. The input acceleration level was normalized such that both the modeled and the experimental cases experienced the same level of acceleration input. A more than adequate level of isolation is achieved. The message to be emphasized is that a more than sufficient level of vibration isolation is indeed possible by replacing seat foam with the lightweight actuator in question.

Recalling Figure 5.4, the baseline seat has a resonant peak at 4Hz with an isolation frequency at 6Hz. At 4Hz, the seat has a magnitude ratio of two, meaning any signal near 4Hz that manages to pass through the primary suspension will actually be amplified by the baseline seat. This is one of the advantages of active suspension systems – Figure 5.17 clearly shows that it is easy to actively control the system to eliminate *any* amplification.

A velocity dependent gain is similar to merely increasing damping of the system. However, the beauty of active control is now evident in the fact that we can increase damping only in an area surrounding the resonant peak, while leaving the rest of the frequency response function unchanged.

A time history of the dead mass setup being driven with a representative road input signal is shown in Figure 5.18. Both the open loop and closed loop responses are shown. This case shows the level of improvement for an absolute velocity gain of 0.5.

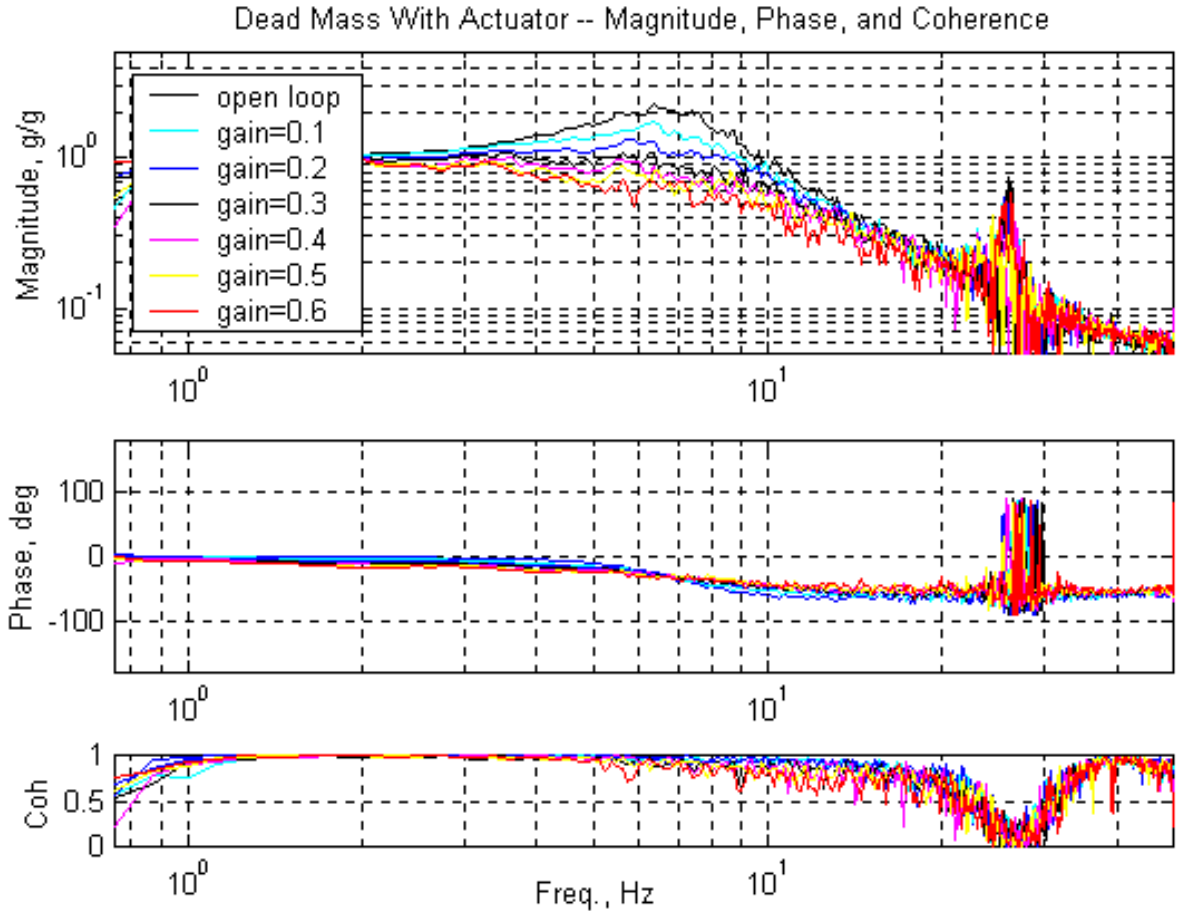

Figure 5.17: Successful Active Vibration Isolation

Table 5.1: Reduction in RMS Acceleration vs. Model

Acceleration Levels (g_{rms})

simulation		experimental	
passive foam	active	control off	active
0.078	0.015	0.025	0.017

Identical random inputs are obtained for the two cases by using the exact same seed value for the noise generator in simulink. The results are not as drastic as predicted by the MATLAB model, since greater effort was made to obtain a representative road signal in the dead mass experiment than was made in the simulation. The active system shows greatest improvement over the baseline seat in the 4–10Hz region, while the frequency content of the representative road signal begins to fall off above 3–4Hz. The representative road input does not incorporate distinct road events such as potholes or railroad tracks which might better emphasize the improvement of the active system.

Efforts to use the three remaining control gains in the simulink model did not show any improvement over the use of the absolute velocity gain alone. Since a desired level of isolation was achieved with the velocity gain in the frequency domain, a full evaluation of all control gains was not completed. Use of this single gain has the distinct advantages of reducing complexity in terms of data collection and processing, since only one acceleration signal is required.

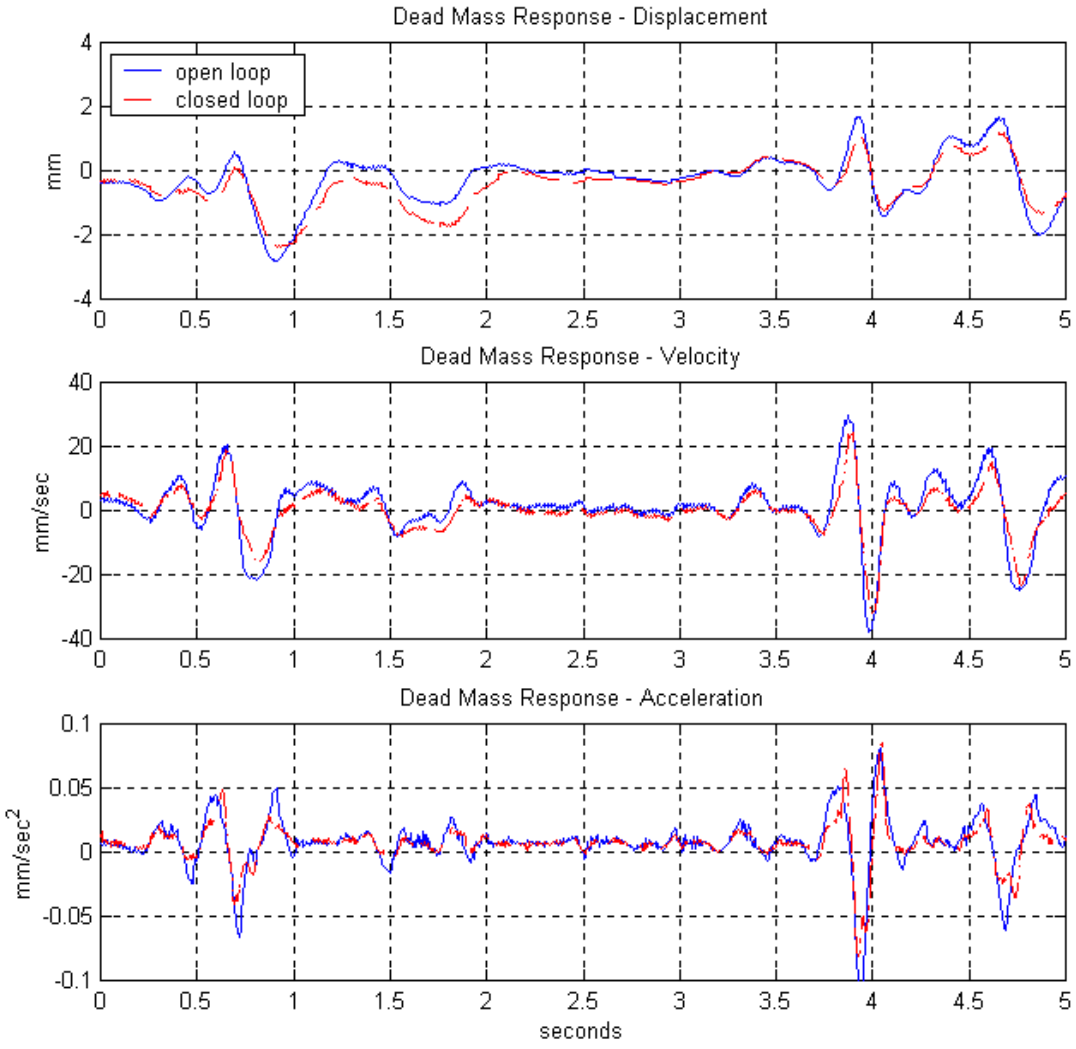

Figure 5.18: Time History of Dead Mass With Representative Road Input

Chapter 6

Conclusions

6.1 Degree Of Isolation Realized

The prototype actuator is successfully applied to achieve significant vibration isolation of an experimental setup representing one–quarter of a vehicle seat and passenger. The active system eliminates any amplification in the frequency response function, whereas the baseline seat tested has a magnitude ratio of two at its resonant frequency of 4Hz. Isolation frequency is reduced to 2-3Hz, while the baseline seat does not begin isolation until 6Hz.

The passenger comfort data suggests that a car seat with a resonant peak in the 4-10Hz region is directionally *incorrect*, since this is the frequency range in which human beings are most sensitive to vertical vibration. By eliminating any amplification in this range, the potential for a passenger to experience discomfort is reduced.

The prototype active suspension has a total mass of 588 grams. The total amount of foam in the baseline seat is 2,660 grams, meaning that the active system indeed shows a possibility of replacing seat foam with potential mass reduction. For a complete active suspension system using the prototype assembly, with four THUNDER devices stacked in each of two actuators, total mass would be 1,248 grams (not including the added mass of a controller or power amplifier).

Although the simulink model allows control using four different control gains,

the gain based solely on absolute velocity of the mass yielded more than sufficient results. A control algorhythm using only the velocity of the mass eliminates the need for an additional measurement at the vehicle floorboard level, and so reduces cost and complexity.

6.2 Recommendations For Future Work

The obvious next step is to apply the actuator assembly to a complete automotive seat and passenger system. Two actuator assemblies with four THUNDER devices each would have sufficient force and stroke, according to our quarter–seat model. An actuator at both the front and rear of the seat would also allow control of any pitching of the seat, should it be desired in the future.

Power consumption of the system was examined only in passing, and determined to be 4 Watts average in the model. The piezoceramic actuator would be well–suited for a regenerative system in which motion of the seat could also produce energy to be fed back into the system.

A consideration of crucial importance in the automotive realm is the lifetime of the device in a real world application, which was not investigated in this thesis. Fatigue testing of the actuator would answer the question of whether it is suitable for long term exposure to an automotive environment. Cost is the other real–world constraint which was conveniently ignored in this thesis. While the cost of the prototype actuator is less than astronomical, commercial viability requires questioning how much an actual, mass–produced seat suspension would truly cost.

An interesting sidenote is that an actuator under the seat not only allows vibration isolation, but would also allow the input of vibration to produce a relaxing massage. The thin PZT devices used here seem ideal for embedding a small back massager in a lightweight seat...

Bibliography

E. Berger, and B. J. Gilmore "Seat Dynamic Parameters for Ride Quality", *Seat System Comfort and Safety*, SP–963, Society of Automotive Engineers, Inc., 1993

J. D. Carlson, "Low–Cost MR Fluid Sponge Devices", 7th ERMR Conference, paper AI1, 1999

J. D. Carlson, D. M. Cantanzarite, and K. A. St. Clair, "Commercial Magneto–Rheological Fluid Devices", dave_carlson@crd.lord.com

B. Culshaw, *Smart Structures And Materials*, Artech House Inc., 1996

M. V. Gandhi and B. S. Thompson, *Smart Materials And Structures*, Chapman And Hall, 1992

T. D. Gillespie, *Heavy Truck Ride*, SP–607, Society of Automotive Engineers, Inc., 1985

T. D. Gillespie, *Fundamentals of Vehicle Dynamics*, Society of Automotive Engineers, Inc., 1992

M. Griffin, and K. Parsons, E. Whitham, "Vibration And Comfort, IV Application of Experimental Results", *Ergonomics* Vol. 25 No. 8, pp. 721–739, 1982

M. Griffin, E. Whitham, and K. Parsons, "Vibration And Comfort, I. Translational Seat Vibration", *Ergonomics* Vol 25 No. 7, pp. 603–630, 1982

D. J. Inman, *Engineering Vibration*, Prentice Hall, Inc., 1996

G. Paddan, and M. Griffin, "The Transmission of Translational Seat Vibration To The Head, I. Vertical Seat Vibration", *Biomechanics* 21(3), pp.191–197, 1988

SAE, "Ride And Vibration Data Manual", SAE J6a, Society of Automotive Engineers, Inc., 1965

X. Song, *Design Of Adaptive Vibration Control Systems With Application Of Magneto–Rheological Dampers*, PhD. Dissertation, Virginia Polytechnic Institute and State University, Dec. 1999

Appendix A

MATLAB Code For 4DOF Seat Model

```
% Mark Malowicki     vertmodel2.m
%
% model the vertical motion of a seated passenger
% model includes seat foam, primary suspension, and tire stiffness
% state space form
% 4 DOF


clear all; close all;

% Define constants, SI UNITS
% note values are for complete vehicle (include all four corners)
% desired natural frequencies used to determine spring rates

wnveh=1.5;        %primary suspension nat'l freq, Hz
wnwh=12;          %wheel hop nat'l freq, Hz
wnh=4;

Mh=65;            %mass of human passenger, kg
Mveh=1200;        %sprung vehicle mass, kg
Mwh=150;          %unsprung (wheel) mass (total all 4 corners), kg
Mf=18;            %foam cushion mass + seat mass, kg
Mf2=0.6;          %foam cushion mass, kg, half foam thickness

Kf=50000;         %seat cushion (foam) stiffness, N/m
Kf2=100000;       %seat cushion (foam) stiffness, N/m, Half foam (actuator case)

Ka=99999999;      %actuator stiffness, no actuator case (solid connection)
Ka2=100000;       %actuator stiffness, actuator case
```

```
Ks=Mveh*(wnveh*2*pi)^2
Kt=Mwh*(wnwh*2*pi)^2 -Ks
%Kf=Mh*(wnh*2*pi)^2

Cf=600;           %seat cushion (foam) damping, Ns/m
Cf2=900;          %seat cushion (foam) damping, Ns/m, Half foam (actuator case)
Cs=18000;         %primary suspension damping, Ns/m
Ca=1000;          %actuator damping, no actuator case
Ca2=10;           %actuator damping, actuator case
Ct=10;            %tire damping, Ns/m


% Define state space system

%           F = state matrix        G = input matrix
%           H = output matrix       J = feedforward matrix
%
%           state vector = [Xh; Xf; Xveh; Xwh; Vh; Vf; Vveh; Vwh]
%
%           input vector = [uroad; uactuator]

F = [zeros(4) eye(4);
   -Kf/Mh Kf/Mh 0 0                -Cf/Mh Cf/Mh 0 0;
   Kf/Mf -(Kf+Ka)/Mf Ka/Mf 0       Cf/Mf -(Cf+Ca)/Mf Ca/Mf 0;
   0 Ka/Mveh -(Ks+Ka)/Mveh Ks/Mveh 0 Ca/Mveh -(Cs+Ca)/Mveh Cs/Mveh;
   0 0 Ks/Mwh -(Ks+Kt)/Mwh         0 0 Cs/Mwh -(Cs+Ct)/Mwh];

% modified F matrix for case of actuator and half foam thickness
F2 = [zeros(4) eye(4);
   -Kf2/Mh Kf2/Mh 0 0                     -Cf2/Mh Cf2/Mh 0 0;
   Kf2/Mf2 -(Kf2+Ka2)/Mf2 Ka2/Mf2 0       Cf2/Mf2 -(Cf2+Ca2)/Mf2 Ca2/Mf2 0;
   0 Ka2/Mveh -(Ks+Ka2)/Mveh Ks/Mveh      0 Ca2/Mveh -(Cs+Ca2)/Mveh Cs/Mveh;
   0 0 Ks/Mwh -(Ks+Kt)/Mwh                0 0 Cs/Mwh -(Cs+Ct)/Mwh];

G = [0 0;
   0 0;
   0 0;
   0 0;
   0 -1/Mf;
   0 1/Mveh;
   Kt/Mwh 0];

G2 = [0 0;
   0 0;
   0 0;
   0 0;
   0 0;
   0 -1/Mf2;
   0 1/Mveh;
   Kt/Mwh 0];

% define a 'root locus' G matrix, for case of acuator input only (no road input)
Grl=G2(:,2);
```

```
H=[1 0 0 0 0 0 0 0;          % vertical position of passenger, meters
   0 1 -1 0 0 0 0 0;         % actuator stroke, meters
   0 0 1 0 0 0 0 0;          % vertical position of vehicle mass, meters
   0 0 0 0 0 0 0 0;          % acuator force output, Newtons
   0 1 -1 0 0 0 0 0;         % last two outputs used for control, meters
   0 0 0 0 0 1 -1 0];        % last two outputs used for control, meters

J =[0 0;
   0 0;
   0 0;
   0 1;                      % acuator force output
   0 0;
   0 0];

sys1=ss(F,G,H,J);            % baseline system - no actuator & full seat foam
sys2=ss(F2,G2,H,J);          % 'improved' system - actuator & half seat foam
damp(sys1)                   % output nat'l freqs & damp ratios
damp(sys2)

% check root locus to find gains one at a time, individually
% (one gain for position difference Xh-Xveh, one for velocity difference Vh-Vveh
Hrl=H(5,:); Jrl=[0];
sysrl=ss(F2,Grl,Hrl,Jrl);
figure(1); rlocus(sysrl); grid on; axis([-50 30 -40 40]);
kgain(1)=rlocfind(sysrl);
kgain(2)=0;
%kgain(1)=0;

% do root locus for 2nd gain
Hrl=H(6,:); Jrl=[0];
sysrl=ss(F2,Grl,Hrl,Jrl);
figure(2); rlocus(sysrl); %axis([-100 60 -80 80]);
grid on; kgain(2)=-rlocfind(sysrl);
%kgain(2)=0;

kgain

% define feedback loop

Ffb=zeros(8);                % actuator dynamics not included
Gfb=zeros(8,2);
Hfb=zeros(1,8);
Jfb=[kgain];                 % feedback gains

sysfb=ss(Ffb,Gfb,Hfb,Jfb);

sys2cl=feedback(sys2,sysfb,[2],[5 6]);

% build an input vector
nsteps=2500;                 % # of steps
tstep=.001;                  % timestep, seconds

% pothole, 5cm (or depth) deep, 50cm long
```

```
%depth=-.05;                  % pothole depth in m
%for n=1:(.5/tstep)
%   u(:,n)=[0; 0];
%end;
%for n=(.5/tstep + 1):(.51/tstep)   % .26m pothole length = .01s at 60mph
%   u(:,n)=[depth; 0];
%end;
%for n=(.51/tstep + 1):nsteps
%   u(:,n)=[0; 0];
%end;
%for n=1:nsteps
%   time(n)=n*tstep-tstep;
%end;

% random input
for n=1:nsteps
   u(:,n)=[(rand/10)-.05;0]*1;      % u(1,:) is (-5cm < u < +5cm) * scaling factor
   time(n)=n*tstep-tstep;
end;

figure; grid on; lsim(sys2cl,u,time);
title('Closed Loop System Response to Road Input');
[y,t]=lsim(sys2cl,u,time);          % closed loop response
[yol,tol]=lsim(sys1,u,time);        % open loop response
[y2ol,t2ol]=lsim(sys2,u,time);      % open loop response


figure;
plot(tol,1000*yol(:,1),'b',t,1000*y(:,1),'r--',t2ol,1000*y2ol(:,1),'g:');
grid on; title('Vertical Displacement of Passenger (Xh)');
xlabel('Time (sec)'); ylabel('Displacement (mm)');
legend('open loop, full foam','closed loop half foam','open loop half foam');
axis([0 2.5 -10 10]);

figure(5);
plot(tol,1000*y(:,2));grid on; title('Actuator Stroke');
xlabel('Time (sec)'); ylabel('Stroke (mm)');
axis([0 2.5 -10 10]);

% just look at bode plots of response at xh...
%figure(3); grid on; bode(sys1cl,{1,500});
[magcl,phcl,wvectcl]=bode(sys2cl,{5,500});
[magol,phol,wvectol]=bode(sys1,{5,500});

figure(6);
for m=1:length(wvectcl);
   magclp(m)=magcl(1,1,m);
   phclp(m)=phcl(1,1,m);
end;
for m=1:length(wvectol);
   magolp(m)=magol(1,1,m);
   pholp(m)=phol(1,1,m);
end;
```

```
subplot(2,1,1); loglog(wvectol,magolp); grid on; axis([5 500 .001 20]);
title('Trans to Passenger');
hold on; loglog(wvectcl,magclp,'r--');
legend('open loop','closed loop'); ylabel('Magnitude');

subplot(2,1,2); semilogx(wvectcl,phclp); grid on; axis([5 500 -360 0]);
hold on; semilogx(wvectcl,phclp,'r--');
ylabel('Phase'); xlabel('Frequency (rad/sec)');
```

Appendix B

MATLAB Code To Initialize Simulink Model

```
% initialize.m
% initialize parameters for roadinc.mdl

Ts=.01;
Tfilt=.02         % Filter timestep
Tfilt2=.15


% Random noise --> representative road acceleration transfer function
wc=5.5;           % cutoff freqency, rad/sec
num1=[wc^2 0 0];  % multiplied by s^2 to get acceleration
denom1=[1 2*wc wc];
sys1=tf(num1, denom1);
[num1,denom1]=tfdata(sys1,'v');


% road input --> veh floorboard transfer function, state space

% desired natural frequencies used to determine spring rates
wnveh=1;          %primary suspension nat'l freq, Hz
wnwh=11;          %wheel hop nat'l freq, Hz
wnh=4;

Mh=77;            %mass of human passenger, kg
Mveh=1400;        %sprung vehicle mass, kg
Mwh=150;          %unsprung (wheel) mass (total all 4 corners), kg

Ks=Mveh*(wnveh*2*pi)^2;
Kt=Mwh*(wnwh*2*pi)^2 -Ks;
Kf=Mh*(wnh*2*pi)^2;

Cf=900;           %seat cushion (foam) damping, Ns/m
Cs=6000;          %primary suspension damping, Ns/m
Ct=10;            %tire damping, Ns/m
```

```
% Define state space system
%           F = state matrix        G = input matrix
%           H = output matrix       J = feedforward matrix
%           state vector = [Xh Xveh Xwh Vh Vveh Vwh]

F = [zeros(3) eye(3);
   -Kf/Mh Kf/Mh 0 -Cf/Mh Cf/Mh 0;
   Kf/Mveh -(Kf+Ks)/Mveh Ks/Mveh Cf/Mveh -(Cs+Cf)/Mveh Cs/Mveh;
   0 Ks/Mwh -(Ks+Kt)/Mwh 0 Cs/Mwh -(Cs+Ct)/Mwh];

G = [0; 0; 0; 0; 0; Kt/Mwh];

H = [1 0 0 0 0 0];

J = [0];

sys=ss(F,G,H,J);          % transfer function road input to floorboard, continuous
```

Vita

It was a cold day in September of 1971 when Mark G. Malowicki first entered the world at a hospital in Utica, New York. Born to the loving parents Edward and Irene Malowicki, he would spend his formative years in the rural splendor near a town named North Western in upstate New York. He earned his high school diploma from Adirondack Senior High School in Boonville, New York, in 1989. He then attended Cornell University in Ithaca, New York, where he graduated with a Bachelor of Science degree in Mechanical and Aerospace Engineering in 1994. He was then plunged into the heartless yet financially rewarding world of gainful employment in the fall of 1994 when he began work as a Durability Test Engineer with General Motors Powertrain Group in Ypsilanti, Michigan. After a few years spent breaking every automatic transmission part you've ever heard of (and even more that you may never hear of), Mark then became a Controls Development Engineer at GM's Proving Grounds in Milford, Michigan. Sensing an ever–present need to expand his horizons in a quest to meet the ever–widening universe, in August of 1998 he enrolled at Virginia Polytechnic Institute and State University, in Blacksburg, Virginia, to pursue a Master of Science degree in Mechanical Engineering, concentrating in vibrations and controls. If you are reading this Vita now, it is safe to assume that Mark earned his Masters degree in the summer of 2000, and returned to General Motors, accepting a position in automatic transmission calibration.

This thesis was typeset with LaTeX 2$_\varepsilon$ by the author. LaTeX 2$_\varepsilon$ is an extension of LaTeX. LaTeX is a collection of macros for TeX. TeX is a trademark of the American Mathematical Society. The macros used in formatting this thesis were written by Greg Walker, Department of Mechanical Engineering, Virginia Tech.