A Comparison of Controller Designs for an Active, Electromagnetic, Offroad Vehicle Suspension System Traveling at High Speed

By:

K.T. Schuetze
J.H. Beno
W.F. Weldon
S.V. Sreenivasan

1998 SAE International Congress and Exposition, February 23 - 26, 1998, Cobo Center, Detroit, Michigan

PR - 257

Center for Electromechanics
The University of Texas at Austin
PRC, Mail Code R7000
Austin, TX 78712
(512) 471-4496

October 10, 1997

A Comparison of Controller Designs for an Active, Electromagnetic, Off-Road Vehicle Suspension System Traveling at High Speed

K.T. Schuetze, J.H. Beno, W.F. Weldon, and S.V. Sreenivasan

Center for Electromechanics
The University of Texas at Austin

ABSTRACT

This paper discusses controller development for an active, off-road vehicle suspension system. A brief review of electronic filters and their characteristics is used to provide insight on the difficulties of designing a control algorithm for negotiating hilly and rough terrain at higher speeds. Two controller designs are presented. One was designed by pole placement and causes the suspension response to approximate a Type 1 Chebychev filter. The other was designed using constrained optimization. A comparison and discussion of simulation results leads to the conclusion that the suspension should be adaptively or predictively controlled for arbitrary terrain and velocity conditions.

INTRODUCTION

The goal of the research to be discussed in this paper was to establish system dynamics for an active suspension that will allow the vehicle sprung mass to follow larger low frequency terrain components (i.e., hills) while still having adequate attenuation of smaller high frequency terrain components (i.e., bumps) at high speed. This separation of low and high frequency terrain components is related to signal processing from the standpoint that the low frequency information must be extracted from a more complex signal. A brief discussion of signal processing techniques will be conducted to shed light on the difficulties associated with separating the low and high frequency components. Next, two controller designs will be presented with simulation results. The first design will be constructed using pole placement to approximate the dynamics of a filter. The second design will be determined by solving a constrained optimization problem. Conclusions about the appropriate dynamic characteristics of the suspension will be based on a comparison of the simulation results for both controller designs.

This research was conducted as part of The University of Texas at Austin Center for Electromechanics (UT-CEM) Electromechanical Suspension (EMS) program, which uses high bandwidth electromechanical actuators to develop positive and negative forces, as required at each wheel for high speed travel over rough off-road terrain.[1,2] Currently, UT-CEM has a four-corner test rig in operation to test advanced EMS algorithms.[3] The test rig was designed to represent a military off-road high mobility multi-purpose wheeled vehicle (HMMWV) at one third scale (except the subscale wheel travel is set at 0.127 m, corresponding to a higher performance off-road vehicle with 0.387 m of travel). Consequently, to capitalize on the availability of this test rig, the research described here was conducted at the one third HMMWV scale, which will enable future testing of the final algorithms that emerge from this ongoing research.

BASIC SYSTEM MODEL, TERRAIN MODEL AND PERFORMANCE GOALS

A basic system model, representing one quarter of a vehicle with an EMS system, is shown in figure 1. The suspension model is composed of a linear bidirectional force actuator (f), a linear passive spring (k) to support the vehicle static weight, inherent damping (b), a tire model (k_t and m_w), and sprung mass (m). The analysis shown from this point forward assumes a linear actuator and the 0.3 scale test rig mechanical properties. The state equations for the system model are shown below.

$$m\ddot{y}_b + b\dot{y}_b + ky_b = f + b\dot{y}_w + ky_w$$
$$m_w\ddot{y}_w + b\dot{y}_w + (k + k_t)y_w = k_t yg - f + b\dot{y}_b + ky_b$$

It should be noted that the dynamics of the actual test rig are nonlinear. However, extensive system identification has been performed so that the forces associated with these parameters can be approximated as linear to within a few percent. This fact justifies the use of a linear model for analysis and development of the control algorithm. The system is fully controllable through the actuator force (f).

The selection of a controller form is dependent upon which measurements are available or readily calculated. For these reasons, the relative position between the sprung mass and wheel are used instead of absolute position. With proper accounting for kinematics, a tachometer, and an encoder or

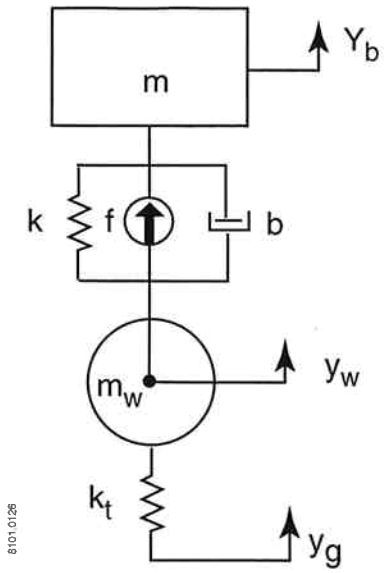

Figure 1. Basic system model and definition of terms

resolver could be attached to the motor shaft of the actuator to obtain relative velocity and displacement. This equipment is readily available and relatively inexpensive. Sprung mass absolute vertical velocity is calculated from accelerometer output, which is filtered and integrated to obtain an absolute velocity. The integration is only valid within a certain frequency band because the filtering process removes dc and very low frequency signal components. Obtaining measured values for absolute vertical position, however, has proven difficult, therefore this quantity is not used.

TERRAIN MODEL -- Terrain displacements can be subdivided into two classes: bumps and hills. Bumps describe terrain displacements less than the vehicle wheel travel and hills describe terrain displacements larger than the allowable wheel travel. In the ideal case, hills would be followed perfectly by the sprung mass and the bumps would be attenuated completely. Generally speaking, the frequency content of the terrain disturbance is dependent on vehicle speed and type of terrain. The combined terrain displacements cause random, non-stationary motions of the wheel within the approximate range of 0 to 7 Hz. Beyond this range, the disturbances are usually small and are beyond the bandwidth of the wheel/tire assembly [4]. The largest hills will generally have frequencies well below 0.1 Hz. However, the main band of concern falls between 0.1 to 1 Hz which is within the range of possible natural frequencies for the sprung mass.

Figure 2 shows the terrain that is used in all of the analysis examples that follow. The hill that is superimposed on a series of high frequency bumps is 0.33 Hz in frequency and 0.127 m peak to peak in amplitude. This hill height violates the original definition for a hill stated earlier because it is within the allowable wheel travel limits. However, the hill combined with the bumps does exceed suspension travel limits. The desired forward speed is 17.9 m/s. Since the simulated vehicle is only 0.3 scale, its speed must be scaled upward to reflect kinetic energy possessed by the full scale vehicle traveling at a slower speed. The value of the scale factor (SF) is:

$$SF = \frac{1.0}{\sqrt{0.3}}$$

For the terrain shown, a scaled speed of 32.6 m/s was used to simulate a full scale vehicle traversing the same terrain at 17.9 m/s.

PERFORMANCE GOALS -- Performance goals of the EMS system are to:

1. Operate without any prior knowledge of the upcoming terrain. The large bandwidth of the control system-actuator assembly allows the suspension to react very quickly to any disturbances.

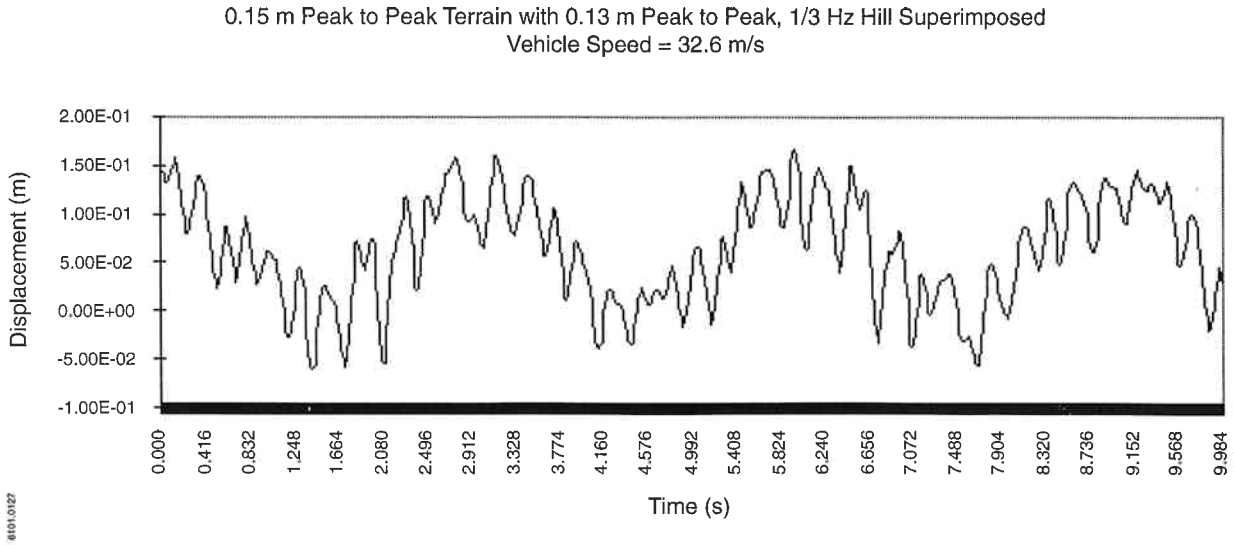

Figure 2. Example terrain

2. Limit sprung mass peak accelerations within a band of ±0.3 g and rms accelerations within a band of ±0.05 g. This requirement was selected to ensure passenger comfort. In the case of a combat vehicle, limiting accelerations below this level would likely allow soldiers to use firing control systems while maneuvering at high speed.

3. Maintain wheel travel within design limits (i.e., to avoid bump stops). This requirement is necessary for the stability and handling of the vehicle. In addition, exceeding wheel travel limits may result in damage to the vehicle or injury to its crew. For this case, ±0.1 m is the wheel travel limit for the 0.3 scale test rig.

SUSPENSION SYSTEMS VIEWED AS A SIGNAL PROCESSING PROBLEM

Vehicle suspension systems are analogous in function to electronic filters. They separate desired information from undesired information. Also, electronic filters can have an ideal magnitude response because they can have approximately unit gain on frequency components in the pass band and around the natural frequency. In light of these facts, a brief review of signal processing techniques provides insight on how to appropriately design the new control algorithm and shows the difficulties of distinguishing bumps from hills without access to future information (i.e, no terrain "look-ahead").

Magnitude and phase characteristics are key factors in the design of an electronic filter for this application. Four specifications are usually required to design the magnitude response of a basic filter: pass band gain, pass band cutoff frequency, stop band start frequency, and stop band gain. The separation in frequency between hill and bump components is often small. Therefore, the magnitude drop off between the pass band and stop band must be sharp to achieve the necessary attenuation of components in close proximity to the pass band cutoff frequency. Although it is often ignored in filter design specifications, the phase response of the filter is also an important consideration. It is well known that the steady state response of a linear system to a sinusoidal input has the same frequency as the input, but has a scaled magnitude and lagging phase. The phase lag delays the response relative to the input. For digital filters, the time delay in samples is equal to the negative of the phase shift divided by frequency. Time delay imparted to large, low frequency signal components must be minimized because it results in excessive wheel travel. In the next section, four different filtering techniques [5-9] will be used in an attempt to isolate the 0.33 Hz component from the example terrain shown in figure 2. The sampling frequency is 20 Hz.

INFINITE IMPULSE RESPONSE (IIR) FILTERS

The IIR filter class includes Butterworth, Chebychev, Bessel and Elliptic filters. In digital form, their output is a weighted sum of input and past output values.

$$y(n) = -\sum_{j=1}^{N} a(j)y(n-j) + \sum_{k=0}^{M} b(k)x(n-k)$$

Theoretically, the filter response continues for infinite time because of dependence on its own output. Several advantages and disadvantage of this type of filter are:

- Advantages
  - Sharp magnitude drop off
  - Easily designed from analog prototypes via transformation
  - Small number of computations allows for efficient real-time implementation
- Disadvantages
  - Nonlinear phase response results in nonlinear time delay and signal distortion
  - Output may be noisy or unstable due to finite precision numerical effects

To investigate these filters, the terrain show in figure 2 was filtered first with a second order Butterworth filter and then with a fourth order Butterworth filter, both with pass band cutoff frequencies of 0.5 Hz. These filters can be quickly designed using the Signal Processing Toolbox in MATLAB or any one of a number of commercial signal processing software packages. The frequency responses and outputs for both filters are shown in figures 3 to 6. The output of the fourth order filter is noticeably smoother than that of the second order. Phase delay occurs in both signals, but is more drastic in the higher order filter. Figure 7 shows the filter time delay as a nonlinear function of frequency. This nonlinearity is most obvious in a region near the cutoff frequency and becomes stronger with increasing filter order. The second order filter has approximately 0.45 s of delay on frequencies within the pass band while the fourth order filter has delays of about 0.85 s in the pass band.

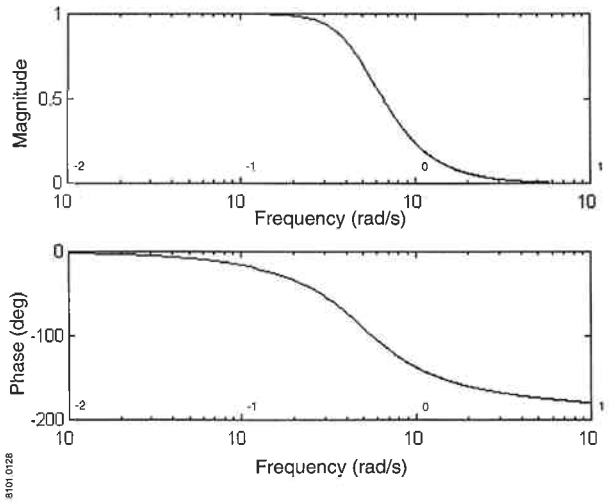

Figure 3. Frequency response of a 2nd order Butterworth filter

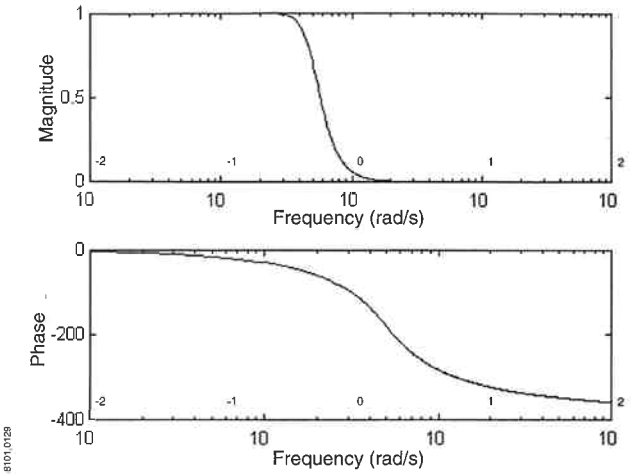

Figure 4. Frequency response of a 4th order Butterworth filter

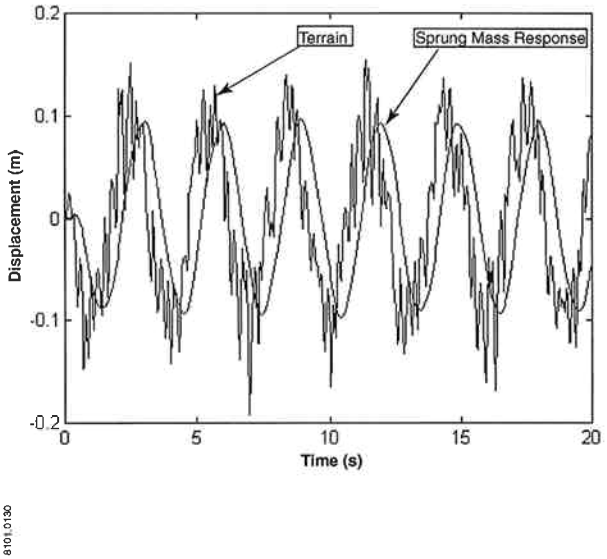

Figure 5. Response of 2nd order Butterworth filter to example terrain

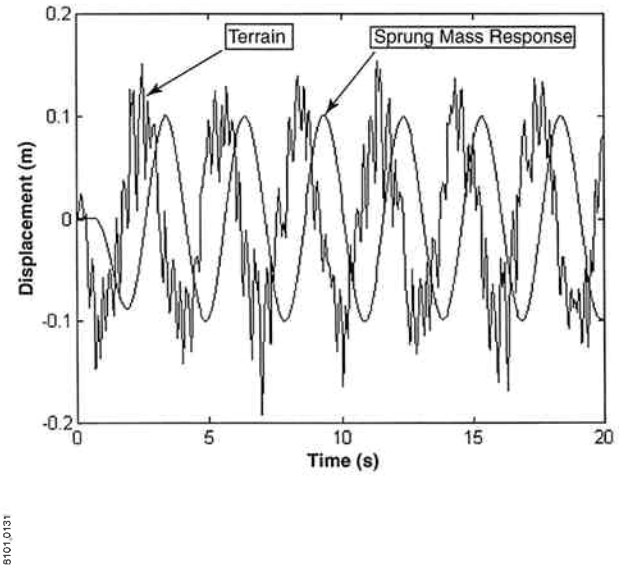

Figure 6. Response of 4th order Butterworth filter to example terrain

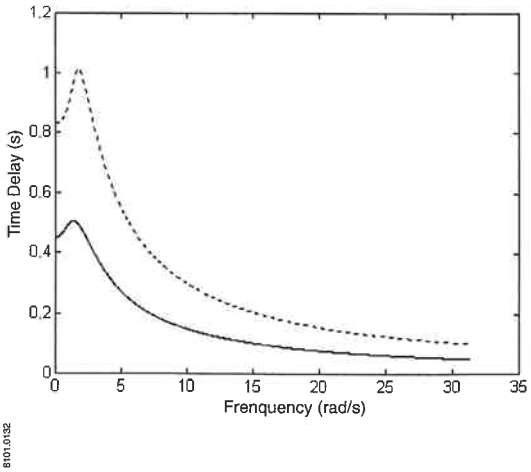

Figure 7. Plot of delay vs. frequency of 2nd (solid) and 4th order (dashed filters)

FINITE IMPULSE RESPONSE (FIR) FILTERS

FIR filters are a class of filters that can only be implemented using digital electronics. Their output (y) is a weighted sum of the current and past input sample values (x).

$$y(n) = \sum_{m=0}^{N-1} b(m)x(n - m)$$

When the input ceases, the filter output will end at a time equal to the filter length in samples. The advantages and disadvantages of this filter type are:

- Advantages
  - Has constant time delay at all frequencies in the pass band
  - Always stable
- Disadvantages
  - Can require large amounts memory
  - Require more computations than IIR filters and may require DSP hardware

Figures 8 to 11 illustrate the frequency response, output and time delay versus frequency plots for a 24 coefficient FIR filter. The filter was designed to approximate the magnitude response of the 2nd order Butterworth filter using the REMEZ Exchange Algorithm [5] that is packaged in MATLAB's Signal Analysis Toolbox. The time delay imposed by an N coefficient FIR filter is always N/2 samples for an even number of coefficients and (N-1)/2 samples for an odd number of coefficients. For the sample rate given earlier, the time delay for this filter is 12*(1/20) = 0.6s. This is higher than the delay created by the second order Butterworth filter. Generally speaking, an IIR filter imparts less delay than the FIR filter for the same level of attenuation. However, in cases where signal quality is an issue FIR filters are preferred because they have constant time delay at all frequencies in the passband.

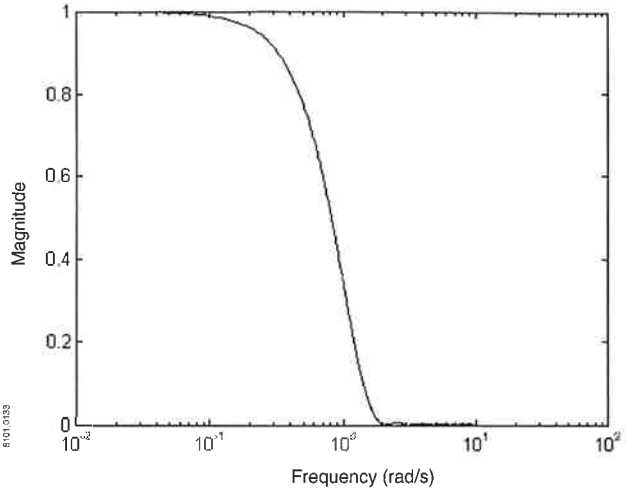

Figure 8. Magnitude response of 24 tap FIR filter

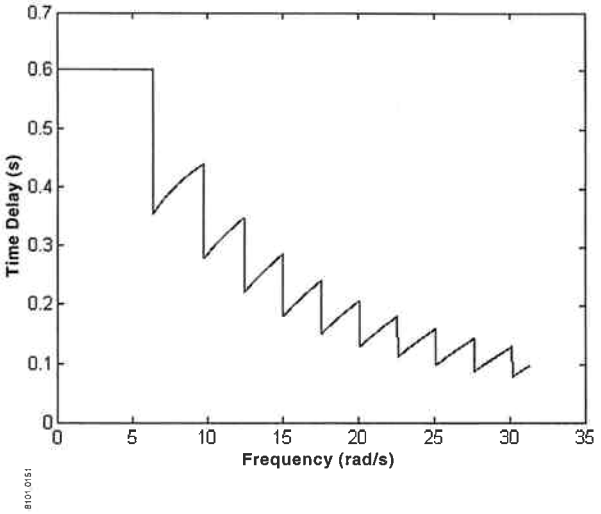

Figure 11. Plot of delay vs. frequency for a 24 tap FIR filter

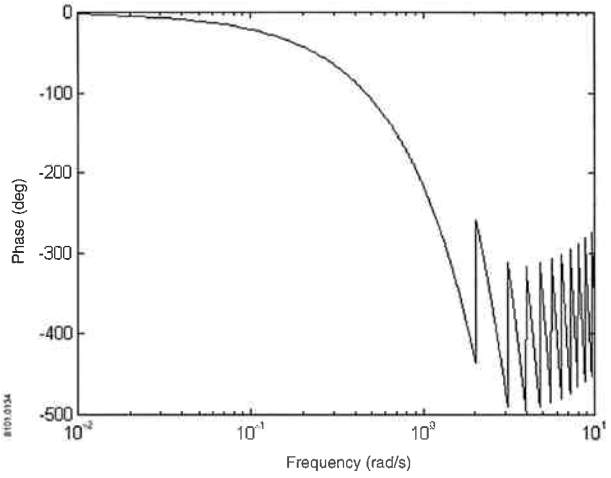

Figure 9. Phase response of 24 tap FIR filter

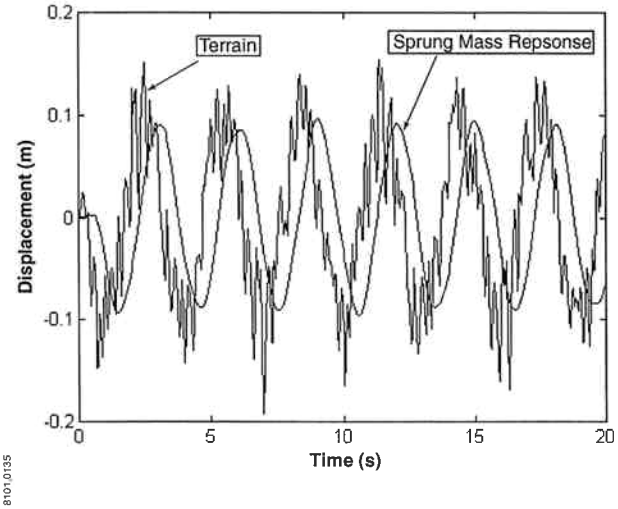

Figure 10. Response of 24 tap FIR filter to the example terrain

DISCRETE FOURIER TRANSFORM (DFT) TRUNCATION

This filter method begins by transforming a sequence of data into the frequency domain via DFT. The DFT of an N point time sequence of samples yields the frequency spectrum of components present in the sequence. It is computed as [5,6]:

$$C(k) = \sum_{k=0}^{N-1} \sum_{n=0}^{N-1} x(n)\, e^{-j2\pi nk/N}$$

The time length of the data sequence determines the frequency resolution in the frequency domain. Filtering is performed by multiplying unwanted frequency components, or bins, by zero and then inverse transforming the remaining low frequency content back into the time domain. The inverse transform is:

$$y(m) = \sum_{m=0}^{N-1} \sum_{k=0}^{N-1} C_f(k)\, e^{j2\pi km/N}$$

where C_f represents the truncated DFT of the N point time sequence. The advantages and disadvantages of the DFT approach are:

- Advantages
  - Always stable
  - Can be implemented in real time using special DSP hardware
- Disadvantages
  - Requires large amounts of memory
  - Spectral leakage occurs with random signals
  - Frequency components lower than the resolution of the window are averaged

Figures 12 and 13 show the process of filtering the sample terrain by this method. The results shown in figure 14 are somewhat deceiving. This particular case was done off-line for

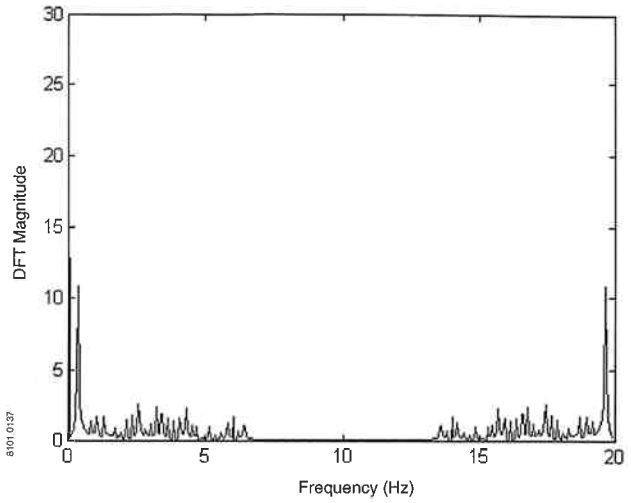

Figure 12. DFT of 20 s of the example signal

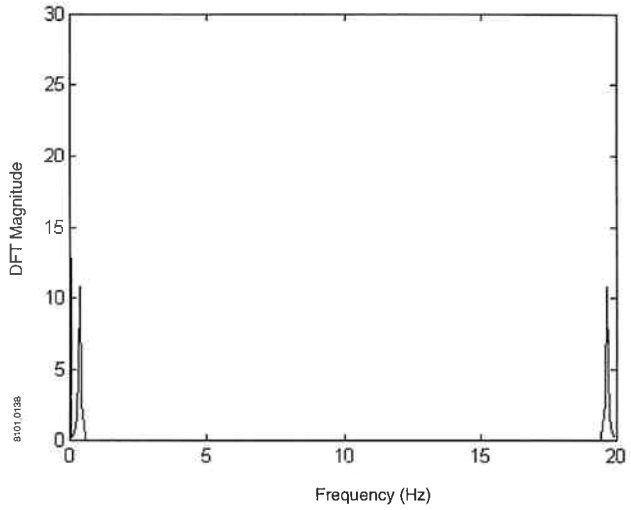

Figure 13. DFT Truncated at 0.5 Hz

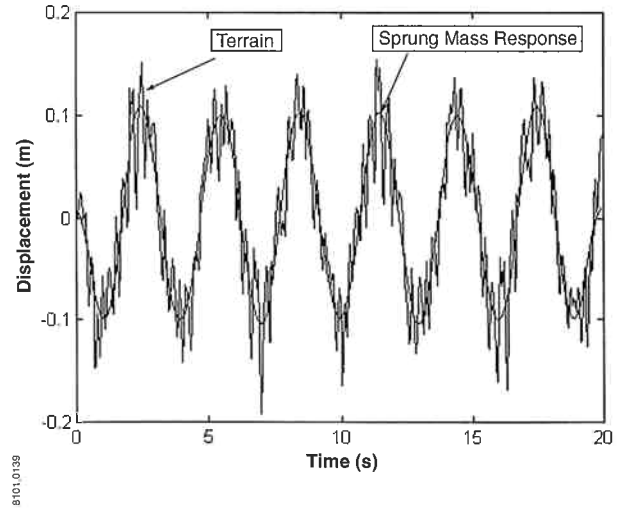

Figure 14. Example terrain and filter output using a stationary window

a situation in which the signal is not evolving in time. Only under these conditions will the filtered data look like this. To implement the DFT filter in real time requires an array of data that must be updated and filtered at each time step. The most current value in the output array is selected as the filter output for that particular time step. The sliding window DFT filter output is shown in figure 15. In this case, the window length is 200 samples, or 10 s and is truncated at 0.5 Hz. It should also be noted that the magnitude of the error in the first 10 s of output can be attributed to the fact that the time sequence is not yet "full" and that there are still several zero valued samples in the sequence. The sliding window filter output suffers from a significant time delay. Since the data in the time sequence changes at each sample instant, the average or dc value of the data changes with time. This time varying average results in the phase delay of the filtered signal. Both filtered signals show some distortion due to spectral leakage. Spectral leakage, a smearing of a primary frequency component over two or more frequency bins, occurs when there are partial cycles of sinusoidal components captured in the time sequence. For a stochastic signal, it is impossible to capture an exact number of cycles of data in the window. Consequently, there will always be spectral leakage of the primary components into surrounding frequency bins.

ZERO PHASE FILTERS

Zero phase filtering [8,9] is performed by passing a sequence of data through the standard digital IIR or FIR filter, reversing its order, and then filtering the reversed sequence again. The resulting output from the filter has twice the attenuation of one filtering process and no delay. Figure 16 illustrates a magnitude response curve for a zero phase filter constructed from two second order Butterworth filters. However, zero phase filtering is a non-causal technique that is usually performed off line because it requires sequences of data that are 3 to 4 time constants, of the filter being used, long. The advantages and disadvantages of this technique are listed below.

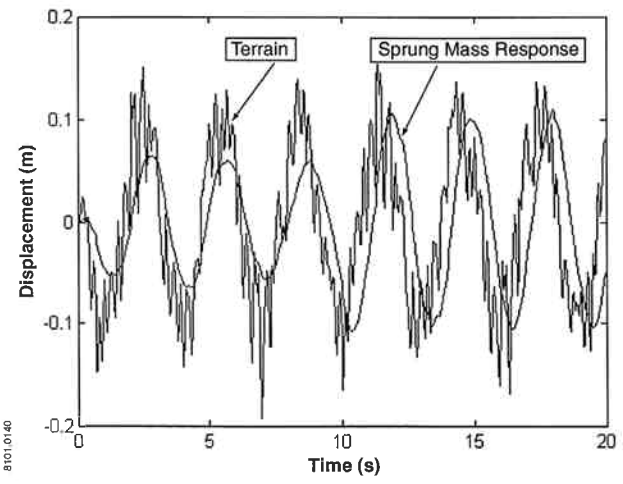

Figure 15. Example terrain and filter output using a sliding window

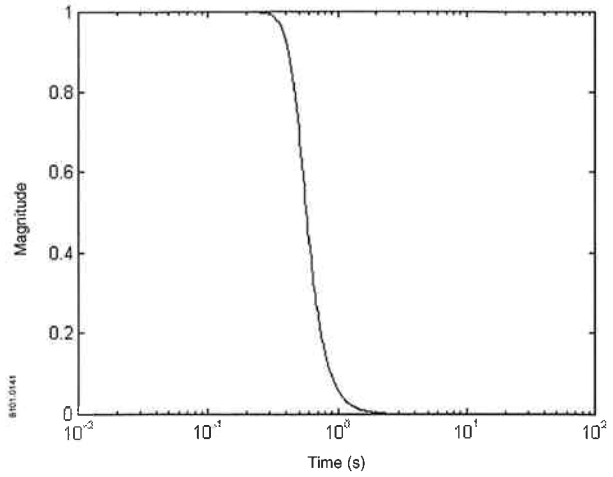

Figure 16. Magnitude response of a zero phase filter using a 2nd order Butterworth filter

- Advantages
  - No phase delay
  - Twice the attenuation of the filter used
  - Real time implementation is possible for high frequency applications [8]
- Disadvantages
  - Non-causal
  - Large number of computations required

Figure 17 depicts the results that are obtained using the second order Butterworth filter as a zero phase filter. The results are as smooth as those obtained from the fourth order Butterworth filter except that there is no time delay. A real time implementation of a zero phase filter has computational time delay equal to the length of three data sequences being processed [9]. If each data sequence must be three to four time constants long, then a second order Butterworth filter with a pass band cutoff frequency of 0.5 Hz would have a delay of 3 to 4 s.

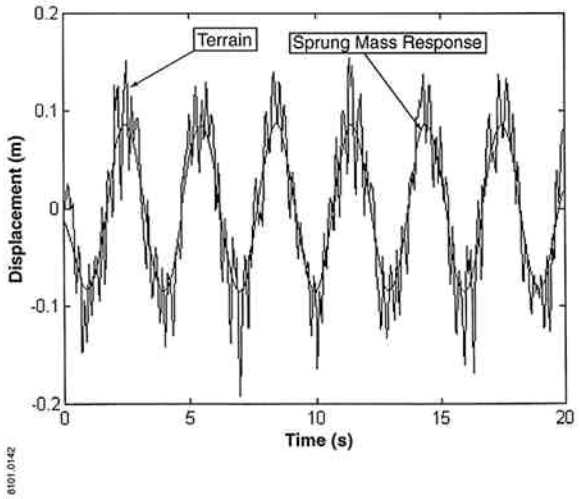

Figure 17. Example terrain and response of a zero phase filter

The goal of these filtering efforts has been to isolate a 0.33 Hz frequency signal under low signal to noise ratio conditions. In all cases, significant delays resulted from the filtering process. If an attempt is made to distinguish a sinusoid from a ramp or other type of function, it would take approximately 1/4 cycle of the sinusoid to make an accurate determination [10]. Therefore, significant delays should be expected when filtering at frequencies in the 0.1 to 1 Hz band. A comparison of controller designs in the next section will lead to an appropriate course of action to minimize the effects of the delay.

CONTROLLER DESIGNS

Previously, a model was presented for the suspension system shown in figure 1. The model is:

$$m\ddot{y}_b + b\dot{y}_b + ky_b = f + b\dot{y}_w + ky_w$$

$$m_w\ddot{y}_w + b\dot{y}_w + (k + k_t)y_w = k_t yg - f + b\dot{y}_b + ky_b$$

The first equation represents the sprung mass dynamics while the second equation represents the wheel dynamics. Since the sprung mass dynamics are of the most interest, only the first equation will be considered in the controller design. The two equations are linked by relative motion and it is not possible to change the dynamics of one without affecting the dynamics of the other. If only one of the equations is used in the design of the controller it is necessary to check the stability of the other equation at the end of the process. Simulation results for the controller designs presented in this section showed that the wheel dynamics are stable for both controllers.

Early work on the UT-CEM EMS system began with the assumption that if the low frequency components of the relative displacement or velocity could be isolated, then the system could be controlled in a manner to drive those components to zero. This idea is equivalent to placing the filter in the control loop. Unfortunately, the filter time constant is on the same order as the time constant of the suspension system. This means that the controller action will be sluggish, making it difficult to control the system in a stable manner while meeting all performance objectives. Consequently, efforts were focused on designing a controller that would force the system to perform analogous to an appropriate filter design.

For control purposes, relative displacement measurements are necessary because of the difficulty in obtaining accurate values of the absolute positions. For this reason, a controller form was chosen that avoids the use of absolute positions. The chosen control design consists of three constant gains applied to relative displacement, relative velocity and absolute velocity. This combination of control terms allows complete specification of the second order sprung mass dynamics. Replacing the force term (f) with the appropriate control terms yields:

$$m\ddot{y}_b + b\dot{y}_b + ky_b = -k_1\Delta - k_2\dot{\Delta} - k_3\dot{y}_b + b\dot{y}_w + ky_w$$

where

$$\Delta = y_b - y_w$$
$$\dot{\Delta} = \dot{y}_b - \dot{y}_w$$

Expanding the relative coordinate measurements into absolute coordinates, the equation becomes:

$$m\ddot{y}_b + b\dot{y}_b + ky_b = -k_1y_b + k_1y_w - k_2\dot{y}_b + k_2\dot{y}_w - k_3\dot{y}_b + b\dot{y}_w + ky_w$$

With some further simplification, the final equation form is:

$$\ddot{y}_b + \frac{1}{m}[(b + k_2 + k_3)\dot{y}_b + (k + k_1)y_b] = \frac{1}{m}[(k_1 + k)y_w + (k_2 + b)\dot{y}_w]$$

Let

$$c_1 = \frac{(k_1 + k)}{m}$$
$$c_1c_2 = \frac{(k_2 + b)}{m}$$
$$c_3 = \frac{k_3}{m}$$

Then the final equation form becomes:

$$\ddot{y}_b + (c_1c_2 + c_3)\dot{y}_b + c_1y_b = c_1y_w + c_1c_2\dot{y}_w$$

It is also helpful to have the transfer function of this differential equation in standard form. Using Laplace transforms, the transfer function is:

$$\frac{Y_b(s)}{Y_w(s)} = \frac{c_1(c_2s + 1)}{s^2 + (c_1c_2 + c_3)s + c_1}$$

Because the system is analog, the only filter type that can be approximated by this system is an IIR filter. The best filter design that can be approximated by the system is the Type 1 Chebychev filter. This filter has the following transfer function form:

$$\frac{Y_{out}(s)}{Y_{in}(s)} = \frac{b_0}{s^2 + a_1s + a_0}$$

where $b_0 \approx a_0$ for a magnitude ripple less than 0.1 dB in the pass band. The sprung mass transfer function is first order in the numerator which means that the system output is dependent on wheel velocity as well as wheel position. To approximate the Type 1 Chebychev filter, the effects of this velocity dependence can be minimized by making c_2 sufficiently small. Making c_2 small is equivalent to minimizing the relative damping in the system. The remaining coefficients are easily specified by equating the denominator coefficients of the system and filter transfer functions.

To demonstrate this controller design, gains c_1, c_2 and c_3 were computed to force the suspension model to behave approximately like a Type 1 Chebychev filter with a pass band cutoff frequency of 0.33 Hz and with a pass band ripple of no more than 0.1 dB. The magnitude and phase plots for the closed loop system are shown in figures 18 and 19. These plots show that the block dynamics approximate the filter dynamics very closely within the band of interest.

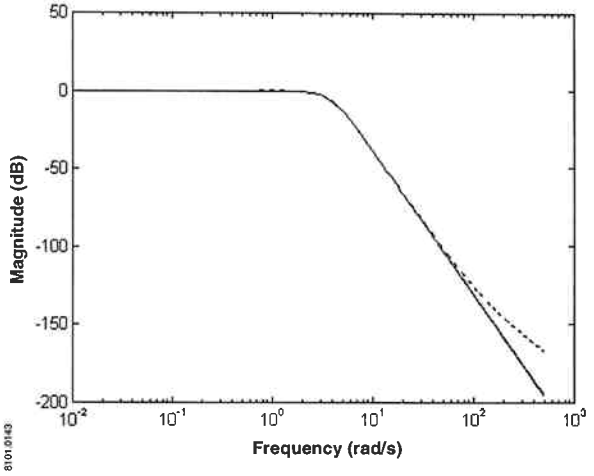

Figure 18. Magnitude response of Type 1 Chebychev filter (solid) and suspension system (dashed)

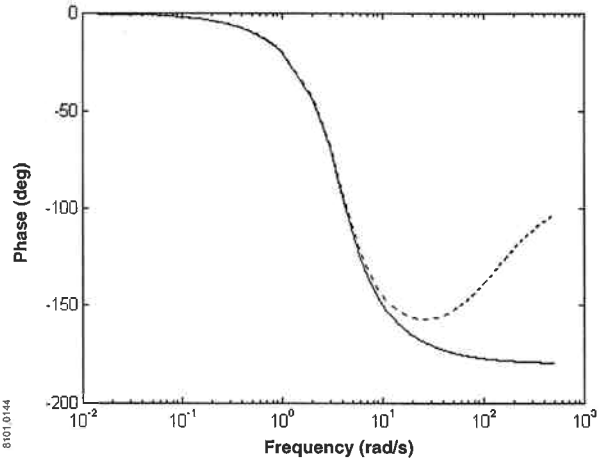

Figure 19. Phase response of Type 1 Chebychev filter (solid) and suspension system (dashed)

Figures 20 and 21 illustrate the simulated performance of the suspension system traversing the terrain shown in figure 2. The sprung mass peak accelerations and rms average acceleration, 0.021 m/s^2, are well within desired levels. On the other hand, The relative displacements exceed allowable limits by as much as 40%. There is significant phase shift in the block trajectory relative to the actual hill trajectory which manifests itself as an obvious low frequency component in the relative

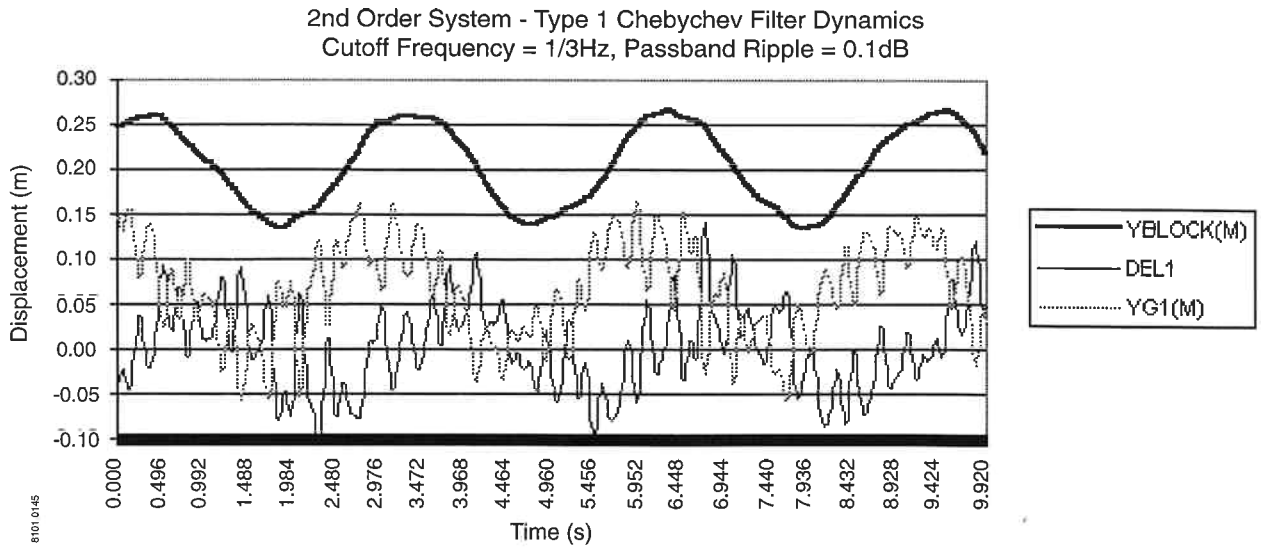

Figure 20. Sprung mass response and relative displacement over simulated terrain, y_g

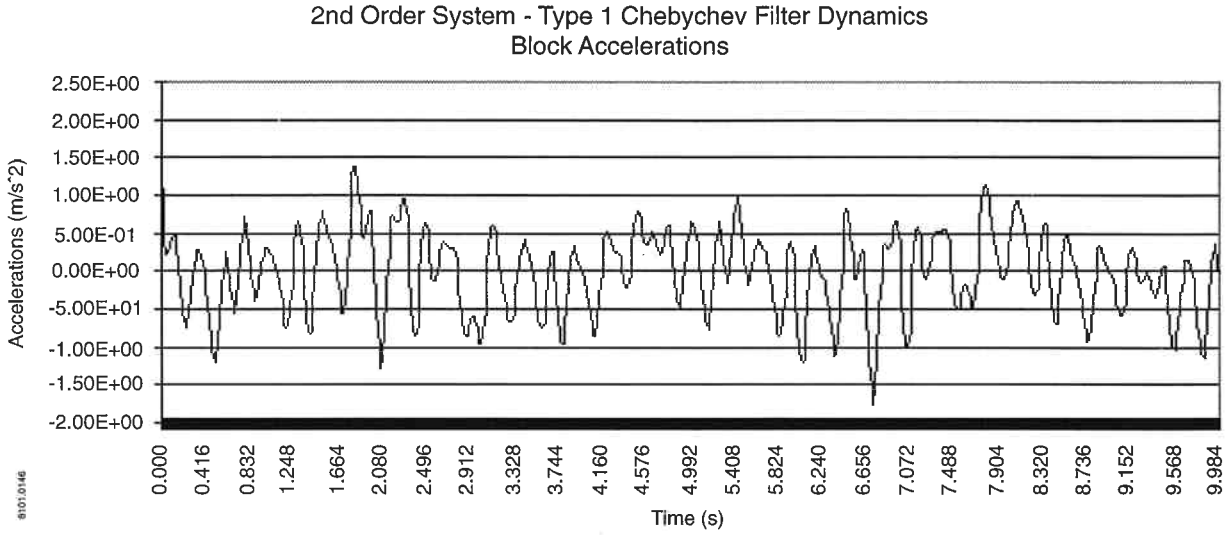

Figure 21. Sprung mass accelerations over simulated terrain, y_g

displacement. The 0.127 m peak to peak terrain without hills or bias has an rms average displacement of 0.0012 m over the 10 s interval. In contrast, the rms average relative displacement for the suspension with the filter dynamics is 0.0019 m over the 10 s interval. The results of this simulation show that it should be possible to decrease the relative displacements by allowing acceleration levels on the chassis to increase.

The final objective was to determine what values of the controller gains c_1,c_2 and c_3 would allow the sprung mass to follow the hill as closely as possible yet still be able to adequately attenuate high frequency bumps. Optimization of the control parameters was carried out using the downhill simplex optimization routine, AMOEBA [11], and an objective function based on hill tracking error with additional penalties for exceeding wheel travel and sprung mass acceleration limits (±0.1 m and ±0.3 g respectively). Several solutions were computed using different initial conditions to ensure that a global minimum had been obtained. Figures 22 and 23 show the frequency response of the optimized system. The optimum suspension system for this terrain has a natural frequency of 0.72 Hz and a damping ratio of 0.27 as opposed to the filter based suspension which had a natural frequency of 0.61 Hz and a damping ratio of 0.65. The net result is that the time constant of the optimized suspension is twice that of the suspension with filter based dynamics.

This suspension model was simulated for the same terrain as the filter suspension. The results are shown in figures 24 and 25. When compared to the simulation for the previous design, the ride is not as smooth. However, wheel travel limits are exceeded only one time by 10%. Also, the rms value of wheel travel has dropped from 0.0019 m, in the previous design, to 0.0015 m. This statistic, along with visual inspection, shows that the optimized system does a much better job of following the hills. By contrast, the peak values of acceler-