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

Fault Detection and Identification with Application to Advanced Vehicle Control Systems: Final Report

Randal K. Douglas, Jason L. Speyer,
D. Lewis Mingori, Robert H. Chen,
Durga P. Malladi, Walter H. Chung

California PATH Research Report
UCB-ITS-PRR-96-25

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

September 1996

ISSN 1055-1425

Fault Detection and Identification With Application to Advanced Vehicle Control Systems

Award No. 65H998, M.O.U. 126

Randal K. Douglas, Durga P. Malladi, Robert H. Chen, Walter H. Chung, Thanh M. Nguyen, Jason L. Speyer and D. Lewis Mingori

Mechanical and Aerospace Engineering Department
University of California, Los Angeles
Los Angeles, California 90095

June 20, 1996

Fault Detection and Identification With Application to Advanced Vehicle Control Systems

Award No. 65H998, M.O.U. 126

Randal K. Douglas, Durga P. Malladi, Robert H. Chen, Walter H. Chung, Thanh M. Nguyen, Jason L. Speyer and D. Lewis Mingori

Mechanical and Aerospace Engineering Department
University of California, Los Angeles
Los Angeles, California 90095

June 20, 1996

Abstract

A preliminary design of a health monitoring system for automated vehicles is developed and tests in a high-fidelity nonlinear simulation are very encouraging. A new detailed nonlinear vehicle simulation which extends the current simulation is documented and will be used as a future testbed for evaluating the performance of the health monitoring system. A health monitoring system has been constructed for the lateral and longitudinal modes that monitors twelve sensors and three actuators. The approach is to fuse data from dissimilar instruments using modeled dynamic relationships and fault detection and identification filters. The filters are constructed so that the residual process has directional characteristics associated with the presence of a fault, that is, static patterns. Sensor noise, process disturbances, system parameter variations, unmodeled dynamics and nonlinearities can distort these static patterns. Two candidate residual processing schemes are developed and tested. A Bayesian neural network is trained to announce a fault and the probability of fault occurrence by recognizing fault patterns embedded in the residual. A new multiple hypothesis Shiryayev probability ratio test is also developed. Finally, development of a

time-varying fault detection filter, applicable to maneuvering vehicles with time-varying dynamics, is described.

Keywords. Automated Highway Systems, Automatic Vehicle Monitoring, Fault Detection and Fault Tolerant Control, Neural Networks, Reliability, Sensors, Vehicle Monitoring.

Fault Detection and Identification With Application to Advanced Vehicle Control Systems

Award No. 65H998, M.O.U. 126

Executive Summary

A preliminary design of a health monitoring system for automated vehicles is developed and tests in a high-fidelity nonlinear simulation are very encouraging. A new detailed nonlinear vehicle simulation which extends the current simulation is documented and will be used as a future testbed for evaluating the performance of the health monitoring system. A health monitoring system has been constructed for the lateral and longitudinal modes that monitors twelve sensors and three actuators. The approach is to fuse data from dissimilar instruments using modeled dynamic relationships and fault detection and identification filters. The filters are constructed so that the residual process has directional characteristics associated with the presence of a fault, that is, static patterns. Sensor noise, process disturbances, system parameter variations, unmodeled dynamics and nonlinearities can distort these static patterns. Two candidate residual processing schemes are developed and tested. A Bayesian neural network is trained to announce a fault and the probability of fault occurrence by recognizing fault patterns embedded in the residual. A new multiple hypothesis Shiryayev probability ratio test is also developed. Finally, development of a time-varying fault detection filter, applicable to maneuvering vehicles with time-varying dynamics, is described.

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

A PROPOSED TRANSPORTATION SYSTEM with vehicles traveling at high speed, in close formation and under automatic control demands a high degree of system reliability. This requires a health monitoring and maintenance system capable of detecting a fault as it occurs, identifying the faulty component and determining a course of action that restores safe operation of the system. This report is concerned with vehicle fault detection and identification and describes a vehicle health monitoring system approach based on analytic redundancy.

Analytic redundancy methods for fault detection and isolation use a modeled dynamic relationship between system inputs and measured system outputs to form a residual process. Nominally, the residual process is nonzero only when a fault has occurred and is zero at other times. For an observable system, this simple definition is met by the innovations process of any stable linear observer. A detection filter is a linear observer with the gain constructed so that when a fault occurs, the residual responds in a known and fixed direction. Thus, when a nonzero residual is detected, a fault can be announced and identified.

A complication arises when there are many possible faults because a fault detection filter can only be designed to detect a limited number of faults. This is related to the order of the vehicle dynamics. When more faults need to be identified, several fault detection filters have to be used with each filter designed to detect and identify some but not all possible faults. The vehicle fault detection system described in this report has four fault detection filters. This raises two difficult design issues. First, some and probably all faults will not be included in the design of one or more fault detection filters. When such a fault occurs, the residual of all filters will respond, even the residuals of the filters that do not have the fault included in their design. If a fault is not included in a fault detection filter design, the directional characteristics of the residual will be undefined and the fault cannot be properly identified. The challenge is to build a mechanism that recognizes when a fault detection filter is responding to a fault for which it has not been designed and then to exclude the residual of all such filters from the fault identification process. If it can be assumed that only one fault occurs at a time, then the residual processor can exclude the residual of any fault detection filters that point to two or more faults.

A second design issue is how the faults should be grouped and identification delegated among the fault detection filters. Several approaches are taken in the design described in this report. In one, the functional form of a given sensor is restricted. In particular, it is assumed that the sensor fault is a bias of unknown magnitude. The assumption allows this sensor and a certain actuator to be isolated by a single fault detection filter. This point is significant because the conventional approach to fault detection filter design would not allow a single filter to isolate these two faults and would require this task to be passed on to the residual processing module.

A second fault grouping design consideration is a newly apparent tradeoff between filter parameter robustness, as determined by the eigenvector conditioning, and fault input observability. Using an eigenstructure assignment algorithm, a design objective is to place well-conditioned eigenvectors. However, it has been found recently that for some fault groups, a fault might have only one highly observable direction. This means that while a fault might be large and dynamically active, the residual is small most of the time. The residual would always be large if all associated eigenvectors were placed close to being collinear with the most observable direction. Hence a tradeoff exists. An objective in assigning faults to fault groups is to minimize the impact of this tradeoff.

A third fault grouping design consideration is discussed in (Douglas et al. 1995). In a fault detection system that consists of a bank of fault detection filters and a residual processor such as a neural network, fault isolation is done through the combined effort of both system elements. The fault detection filter is a carefully tuned device that uses known dynamic relationships to isolate a fault. The neural network residual processor combines the residuals from several filters and resolves any ambiguity. It is suggested that identifying a fault among a group of dynamically similar faults requires the precision of and is best delegated to the fault detection filters. Furthermore, it is suggested that the reliability of the neural network training would be improved if the fault groups associated with each of the fault detection filters are dynamically dissimilar.

In applications it is unrealistic to expect that a residual process would be nonzero only when a fault has occurred. Sensor noise, process disturbances, system parameter variations, unmodeled dynamics and nonlinearities all contribute to the magnitude of a residual. There are many methods to reduce the impact of these effects on the residual but none reduce their effect to zero. This means that some threshold detection mechanism must be built.

A simple threshold detection mechanism announces a fault when the size of a residual exceeds some prescribed value. This prescribed value could be determined from empirical studies which balance a rate of false alarm against a rate of miss alarm. A more complicated residual processor might take into account the thresholds of all other residuals as well. Reasoning that if the probability of simultaneous failures is very small, no fault is announced when more than one residual exceeds a threshold. It is more likely that the nonzero residuals are caused by noise or nonlinearities or some cause other than multiple faults.

Two residual processing systems are described in this report. In the first, a Bayesian neural network considers the residuals from all fault detection filters as constituting a pattern, a pattern which contains information about the presence or absence of a fault. Hence, residual processing is treated as a pattern recognition problem.

The objective of a neural network as a feature classifier is to associate a given feature vector with a pattern class taken from a set of pattern classes defined apriori. In an application to residual processing, the feature vector is a fault detection filter residual and the pattern classes are a partitioning of the residual space into fault directions which include the null fault. A Bayesian neural network also provides probabilities of feature classification conditioned on an observation history. A stochastic training algorithm enhances robustness by treating training sets as as sample sets providing information about the entire population.

A second approach to residual processing described in this report is a modified Shiryayev sequential probability ratio test extended to include multiple hypotheses. The algorithm, which is derived as a dynamic programming problem, detects and isolates the occurrence of a failure in a conditionally independent measurement sequence in minimum time. The test has been further extended to the detection and identification of changes with unknown parameters.

This report is organized as follows. Section 2 describes the car models used for fault detection filter design and evaluation. A nonlinear model is derived directly from one provided by the Berkeley PATH research team (Peng 1992). Low-dimensional linear models that include coupled longitudinal and lateral vehicle dynamics are used for fault detection filter design. The high fidelity nonlinear model is used for evaluation and to obtain the linear models used for design. Section 3 describes the faults to be identified by the fault detection system. Section 4 describes the design of the fault detection filters. This includes how the faults are grouped for each fault detection filter design and how the fault detection filter eigenstructure placement is done. Section 5 presents an evaluation of the performance of the fault detection filters in a nonlinear simulation.

Sections 6 and 7 describe two candidate fault detection filter residual processing systems. In Section 6 a Bayesian neural network is developed and in Section 7 a multiple hypothesis Shiryayev sequential probability ratio test is described. Both are used to process residuals from all fault detection filters to detect and identify which if any fault has occurred.

In section 8 a six degree of freedom nonlinear vehicle model is developed independently of the model used for the Berkeley simulation of Section 2. This work is done to provide a model that better accommodates a nonplanar, rough road surface, one where the road gradient is different for all four wheels. The model will be used to evaluate the robustness of the health monitoring system to road excitation. This effort is a continuation of the work reported in (Douglas et al. 1995).

In section 9 describes a new, disturbance attenuation approach to fault detection filter design. Here, a differential game is defined where one player is the state estimate and the adversaries are all the exogenous signals except for the fault to be detected. By treating faults as disturbances to be attenuated, the usual invariant subspace structure associated with fault detection filters is not present except in the limit. By treating model uncertainty as another element in the differential game, sensitivity to parameter variations can be reduced.

Section 9 also introduces the notion of a fault detection filter for time-varying systems. This is especially important in applications where a vehicle follows a maneuver such as a merge or a split. While first considered in the game theoretic filter derivation, it is expected that the Beard-Jones fault detection filter definition will be extended to time-varying systems in the same way.

Appendix A provides a theoretical review of the Beard-Jones detection filter problem. This appendix also includes some early work in extending the Beard-Jones fault detection filter definition to time-varying systems.

Appendix B provides a review of a fault detection filter left eigenvector assignment design algorithm (Douglas and Speyer 1996). The algorithm gives the user eigenvector conditioning information and provides a direct method for achieving maximally achievable eigenvector conditioning. This algorithm is used for the designs in this report.

Appendix C describes a stabilizing fault detection filter gain that bounds the $\mathcal{H}_\infty$ norm of the transfer matrix from system disturbances and sensor noise to the residual. For multi-dimensional faults, a residual direction is identified that enhances the fault signal to noise ratio while maintaining the $\mathcal{H}_\infty$ norm bound.

CHAPTER 2

Vehicle Model and Simulation Development

IN THIS SECTION, vehicle models are developed for the design and evaluation of fault detection filters. The starting point is a model obtained from the Berkeley PATH research team and derived in (Peng 1992). A version of this model coded in C also is available from Berkeley.

Two variations of the Berkeley model are considered in this section. First, modifications are made to allow for variations in road slope and road noise. The slope is restricted to a constant because of assumptions made in the original derivation of the equations of motion. After modifications, the nonlinear model has 32 states, 3 control inputs and 3 noise inputs. Second, reduced-order linearized models used for detection filter design are developed for a vehicle in a constant radius turn. Linearized models developed for a vehicle operating with zero steering angle are described in (Douglas et al. 1995).

An independent derivation of a six degree of freedom nonlinear vehicle model is also developed to be sure that we understand all the assumptions, definitions and issues which underlie the Berkeley model. This model allows for arbitrary variations in road slope and road noise. Since this model represents a significant effort that was not completed in time to be used in the fault detection filter development, it is described later in Section 8.

2.1 Modification of Berkeley's Model

Primary sources of vehicle dynamic disturbances are road roughness and variations in the road slope. First, allowing for a road roughness disturbance requires a modification to the suspension system of the Berkeley nonlinear model. A simple tire model is introduced so that high bandwidth road noise generates physically realistic suspension damping forces. Next a road noise model is desribed. Finally, the nonlinear model is modified to allow for nonzero road slope. It is important to note that because of assumptions made in the original derivation of the equations of motion, the slope is still a constant although now not necessarily zero.

2.1.1 Suspension System

In the Berkeley nonlinear model, the suspension system is modeled as a spring and damper and the tire is stiff. The stiff tire causes road displacements to pass directly to the suspension system resulting in unreasonably large damping forces. Modeling the tire as a mass and linear spring allows the tire to act as a low pass filter with respect to road displacements and eliminates the unrealistic suspension damping forces. Since the mass of the tire is very small relative to the car, the tire model is simplified to a linear spring as shown in Figure 2.1. The vehicle equations are modified by adding four states, the suspension force of each wheel, which are derived as follows. The suspension force F_s acting on each wheel is given by

$$F_s = -C_1(x_2 - x_1 - x_{30})[1 + C_2(x_2 - x_1 - x_{30})^4] - D_1(\dot{x}_2 - \dot{x}_1) + mg \tag{2.1}$$

where x_{30} is the length of the suspension system when a nominal load mg is applied. The force F_t transmitted to the suspension by the tire spring is given by

$$F_t = -K_t(x_1 - r - x_{10}) \tag{2.2}$$

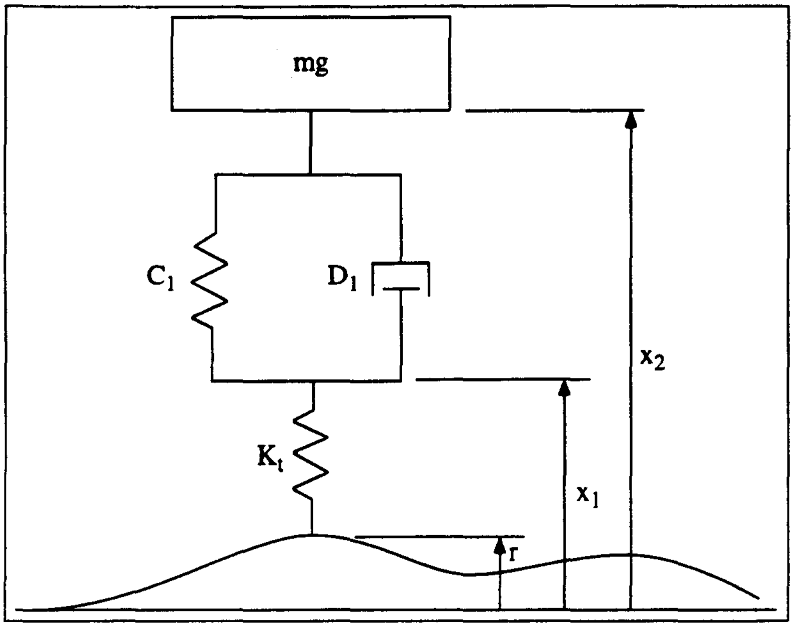

Figure 2.1: Simplified suspension and tire model.

where K_t is the tire spring stiffness and x_{10} is the nominal tire radius. Since the tire is massless, the tire spring force is equal to the suspension force.

$$F_t = F_s \tag{2.3}$$

The tire spring force F_t is eliminated by rearranging (2.2) to get

$$x_1 = r + x_{10} - \frac{F_t}{K_t} \tag{2.4a}$$

$$\dot{x}_1 = \dot{r} - \frac{\dot{F}_t}{K_t} \tag{2.4b}$$

and then combining (2.1), (2.3) and (2.4) as

$$\dot{F} = \frac{K_t}{D_1}\{-F+mg-C_1(x_2-r-x_{10}-x_{30}+\frac{F}{K_t})[1+C_2(x_2-r-x_{10}-x_{30}+\frac{F}{K_t})^4]-D_1(\dot{x}2-\dot{r})\}$$

where $F \equiv F_s$. Adding a suspension force state for each wheel to the nonlinear equations of motion brings the number of states to thirty.

2.1.2 Road Roughness

A road roughness model is derived from Robson (Robson 1980). Empirical data shows that the road displacement r can be modeled as a random process with power spectral density given by

$$P(\lambda) = \begin{cases} \frac{R_c}{\lambda_0^{2.5}}, & \lambda \leq \lambda_0 \\ \frac{R_c}{\lambda^{2.5}}, & \lambda \geq \lambda_0 \end{cases} \tag{2.5}$$

where, for a typical freeway, $\lambda_0 = 0.01$cycles/m and $R_c = 10^{-7}\mathrm{m}^2$. R_c depends on the road roughness.

The model (2.5) is difficult to apply because $P(\lambda)$ is not a rational function of λ. An obvious simplication is to take the exponent of the denominator to be 2 and write the power spectral density as

$$P(\lambda) = \frac{R_c}{\lambda^2 + {\lambda_0}^2} \tag{2.6}$$

Using (2.6), random road displacements r are given by a first-order differential equation

$$\dot{r} = -\omega_0 r + \sqrt{2\pi V R_c}\, \omega$$

where $\omega_0 = 2\pi V \lambda_0$, V is the nominal velocity of the car and ω is white noise with power spectral density one.

For simplicity, road roughness for the right and left tires is the same. Also, road roughness for the front tires is applied to the rear tires but after a time delay. The delay is given by the speed of the car and the distance between the front and rear tires. Thus, only two road roughness noise states are added to the vehicle nonlinear equations of motion. This brings the number of states to thirty two.

2.1.3 Slope

Because of assumptions made in the original derivation of the vehicle equations of motion, allowing for non-zero road slope or superelevation can only be done by rotating the gravity vector. As shown in Figure 2.2, the forces acting on a car where the road slope is γ degrees are equivalent to the forces acting on a car where the road has zero slope and the gravity vector has been rotated γ degrees. Of course, this is only true if the road slope is constant.

Because of assumptions made in the original derivation of the vehicle equations of motion, allowing for time varying road slope or superelevation requires rederivation of the vehicle equations of motion. The equations of motion are derived from first principles in Section 8.

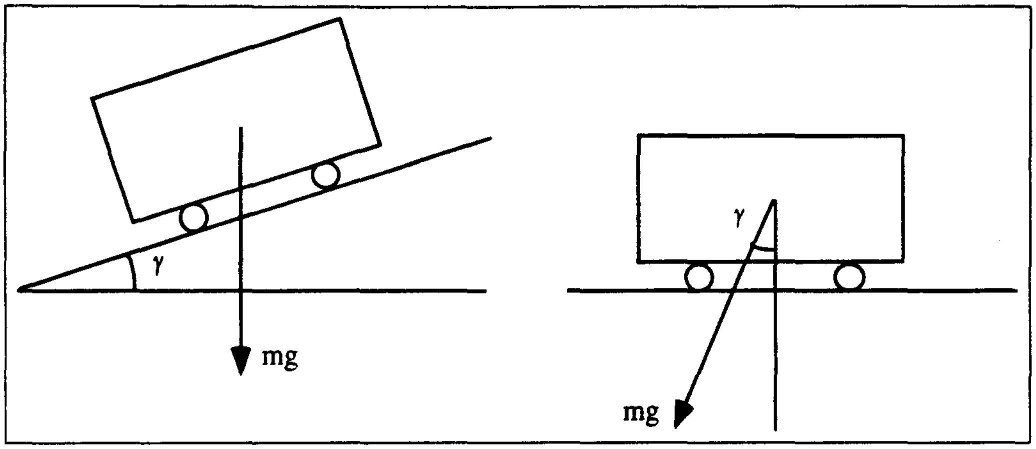

Figure 2.2: Constant non-zero road slope is simulated by rotating the gravity vector.

2.1.4 Evaluation

In this section, the modified tire, suspension system and rough road models are evaluated using the Berkeley nonlinear vehicle simulation. The car is put in a constant radius turn with a steering angle of 0.005 deg. and constant speed of $24.87\frac{\text{m}}{\text{sec}}$ which is about 56mph. In the next section, this nominal operating point is used to derive a linear model for fault detection filter design. Figures 2.3, 2.4 and 2.5 illustrate some of the more relevant vehicle states and outputs. All variables but one appear to take on reasonable and expected values.

An explanation for the large longitudinal acceleration values, which are between -0.05g and 0.05g as shown in Figure 2.4, is that the tire and suspension system is modeled as rigid along the longitudinal direction. This rigid connection allows variations in the tractive force due to the rough road to directly affect the longitudinal acceleration of the car. Since the road noise model is only used for fault detection system robustness testing, large variations in the longitudinal acceleration only imply a more conservative testing environment.

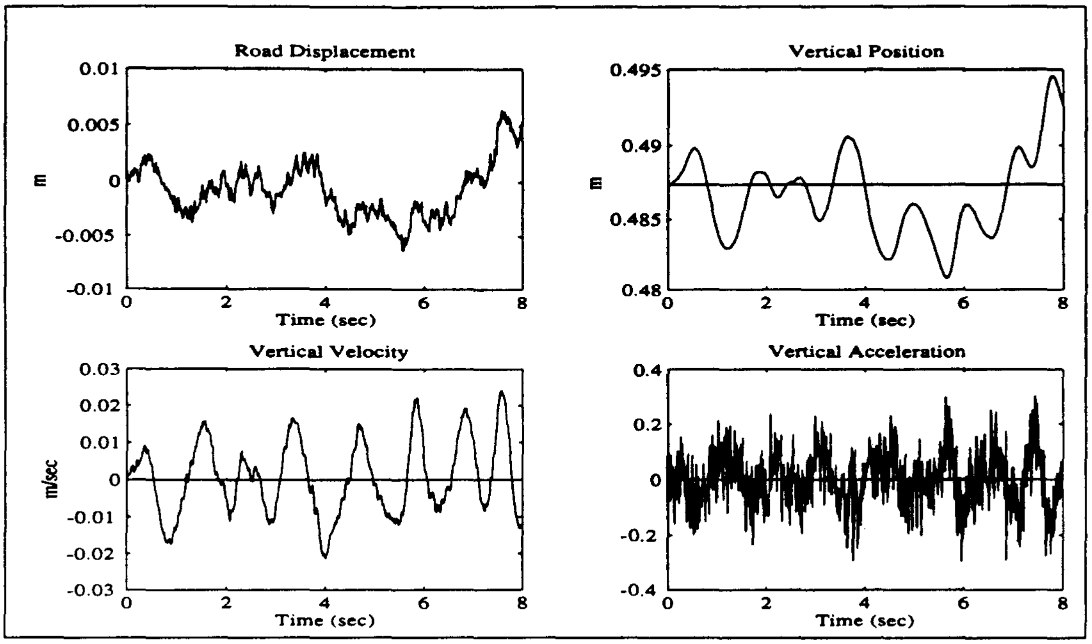

Figure 2.3: Rough road simulation.

2.2 Linear Model

In this section, a linearized model for a car making a constant radius turn is developed using the modified Berkeley nonlinear model. Linearized models developed for a vehicle operating with zero steering angle are described in (Douglas et al. 1995). Linearized models are found numerically rather than analytically. An analytical approach taking partial derivatives is impractical because the nonlinear model is too complicated. The procedure is as follows.

First, a computer run is made in which the car makes a turn at a constant speed of $24.87\frac{\text{m}}{\text{sec}} \simeq 56$mph to obtain steady state values for each state. The tire steering angle is 0.005 rad which produces about 0.1g lateral acceleration and a 638.73 meter radius turn. The nonlinear model is then linearized about this nominal operating point using the central difference method.

The nonlinear model has the form:

$$\dot{x} = f(x, u) \tag{2.7a}$$

$$y = Cx + D\dot{x} \tag{2.7b}$$

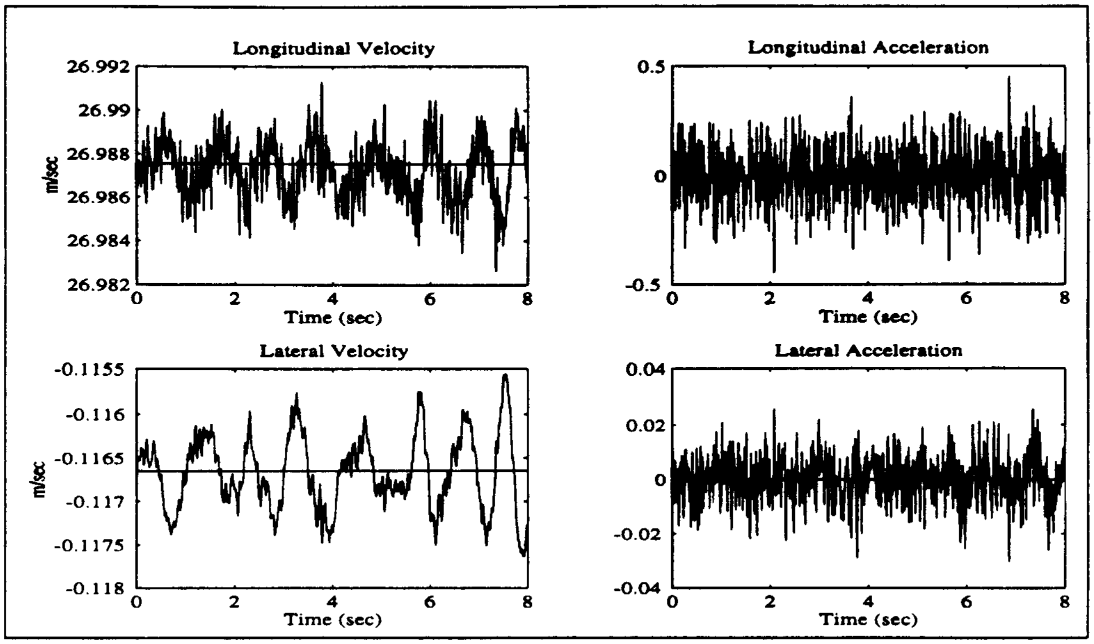

Figure 2.4: Rough road simulation.

Suppose the nominal operating point is $(x_0,\ u_0)$ where $f(x_0, u_0) = 0$. Take perturbations $\tilde{x},\ \tilde{u}$ about the nominal point, that is, let

$$\begin{aligned} x &= x_0 + \tilde{x} \\ u &= u_0 + \tilde{u} \end{aligned}$$

Also approximate $\frac{\partial f}{\partial x}$ and $\frac{\partial f}{\partial u}$ as

$$\begin{aligned} \frac{\partial f}{\partial x} &\approx \frac{\Delta f}{\Delta x} = \left.\frac{f(x+\tilde{x},u) - f(x-\tilde{x},u)}{2\tilde{x}}\right|_{x=x_0,u=u_0} \\ \frac{\partial f}{\partial u} &\approx \frac{\Delta f}{\Delta u} = \left.\frac{f(x,u+\tilde{u}) - f(x,u-\tilde{u})}{2\tilde{u}}\right|_{x=x_0,u=u_0} \end{aligned}$$

Equation (2.7a) may now be approximated as

$$\dot{x}_0 + \dot{\tilde{x}} = f(x_0, u_0) + \left.\frac{\partial f}{\partial x}\right|_{x=x_0,u=u_0} \tilde{x} + \left.\frac{\partial f}{\partial u}\right|_{x=x_0,u=u_0} \tilde{u} + \cdots$$

Truncating the higher order terms and using the approximations given above for the partial derivatives, produces the following linear equations in the perturbed state $\tilde{x}$, input $\tilde{u}$ and

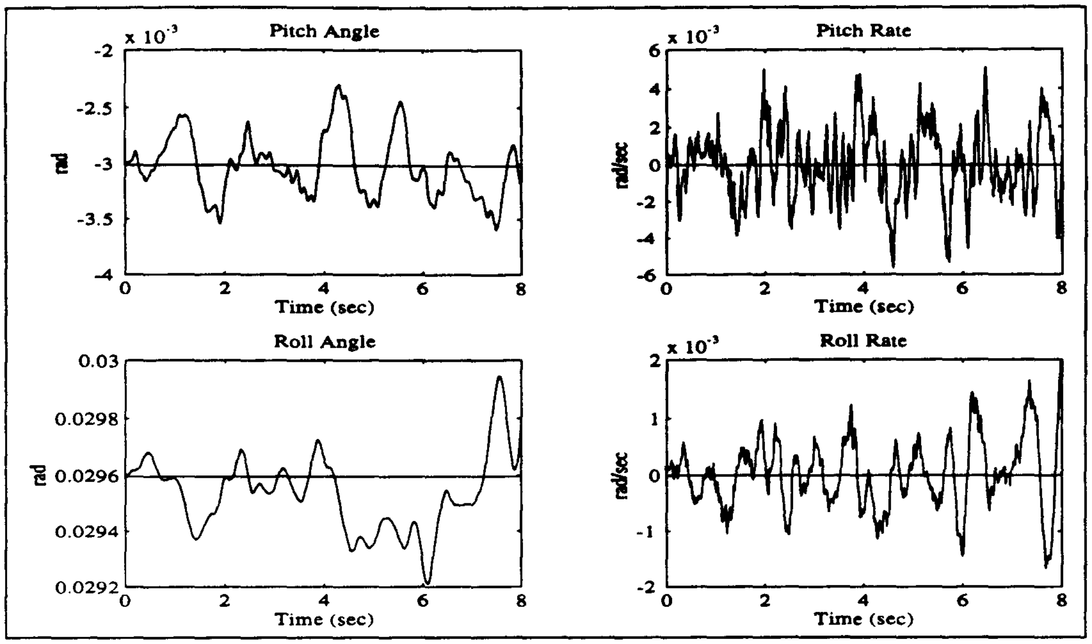

Figure 2.5: Rough road simulation.

output $\tilde{y}$

$$\dot{\tilde{x}} = A\tilde{x} + B\tilde{u} \tag{2.8a}$$

$$\begin{aligned} \tilde{y} &= C\tilde{x} + D\dot{\tilde{x}} \\ &= (C + DA)\tilde{x} + DB\tilde{u} \end{aligned} \tag{2.8b}$$

where

$$\begin{aligned}
\tilde{x} &= [m_a \quad w_e \quad v_x \quad x \quad v_y \quad y \quad v_z \quad z \quad \phi \quad \dot{\phi} \quad \theta \quad \dot{\theta} \quad \epsilon \quad \dot{\epsilon} \quad w_{fl} \quad w_{fr} \quad w_{rl} \\
&\qquad w_{rr} \quad X \quad Y \quad yr \quad \dot{yr} \quad \epsilon_{des} \quad \alpha \quad \tau_b \quad \beta \quad F_{fl} \quad F_{fr} \quad F_{rl} \quad F_{rr} \quad r_r \quad r_f]^T \\
\tilde{y} &= [m_a \quad w_e \quad \dot{v}_x \quad \dot{v}_y \quad \dot{v}_z \quad \dot{\phi} \quad \dot{\theta} \quad \dot{\epsilon} \quad w_{fl} \quad w_{fr} \quad w_{rl} \quad w_{rr}]^T \\
\tilde{u} &= [\alpha_c \quad \tau_{bc} \quad \beta_c \quad r_{rc} \quad r_{fc} \quad \gamma]^T \\
A &= \left.\left[\frac{\Delta f}{\Delta x}\right]\right|_{x=x_0, u=u_0} \\
B &= \left.\left[\frac{\Delta f}{\Delta u}\right]\right|_{x=x_0, u=u_0}
\end{aligned}$$

and where A is a 32×32 real matrix, B is a 32×6 real matrix and DB is a zero matrix.

Symbols in $\tilde{x}$, $\tilde{y}$ and $\tilde{u}$ are defined in the list of symbols.

Several sizes of perturbations must be taken to find one that gives the best approximation of the partial derivatives. If the perturbation is too small, there is a truncation error in computing the difference $f(x+\tilde{x},u)-f(x-\tilde{x},u)$. If the perturbation is too large, a roundoff error occurs in computing $f(x+\tilde{x},u)$ and $f(x-\tilde{x},u)$; also nonlinearities become important. According to our experience, $\frac{\tilde{x}}{x}$ and $\frac{\tilde{u}}{u} \approx 10^{-4}$ is a good rule for selecting the size of the perturbation when using the central differences method.

The resulting linear model is tested in a simulation to see how well it describes the nonlinear model by comparing the states of the linear and nonlinear models when various control inputs are applied. Over the speed range of $23\frac{\mathrm{m}}{\mathrm{sec}}$ to $27\frac{\mathrm{m}}{\mathrm{sec}}$, errors in the states are less than 10% except for yaw rate where the error is less than 15%.

The linear model generated as described above was intended for use in designing the fault detection filters. Since the model dimension is large with 32 states, before the model is used for design, it is simplified to the extent possible without significant loss of accuracy. The model simplification is accomplished in three steps, the first two of which result in no loss of accuracy.

First, since the present fault detection filter designs do not explicitly include either road roughness or road slope, these three noise inputs and two associated states are truncated from the model (2.8). The model (2.8) becomes

$$\begin{aligned}\dot{\tilde{x}}_r &= A_r\tilde{x}_r + B_r\tilde{u}_r\\ \tilde{y} &= C_r\tilde{x}_r\end{aligned}$$

where

$$\begin{aligned}\tilde{x}_r &= [m_a \quad w_e \quad v_x \quad x \quad v_y \quad y \quad v_z \quad z \quad \phi \quad \dot{\phi} \quad \theta \quad \dot{\theta} \quad \epsilon \quad \dot{\epsilon} \quad w_{fl} \quad w_{fr} \quad w_{rl}\\ &\qquad w_{rr} \quad X \quad Y \quad yr \quad \dot{yr} \quad \epsilon_{des} \quad \alpha \quad \tau_b \quad \beta \quad F_{fl} \quad F_{fr} \quad F_{rl} \quad F_{rr}]^T\\ \tilde{u}_r &= [\alpha_c \quad \tau_{bc} \quad \beta_c]^T\end{aligned}$$

Next, by inspection of the equations, it is possible to rearrange the sequence of states such that the linearized equations assume the following partitioned form:

$$\begin{aligned}
\dot{\tilde{x}}_r &= \begin{bmatrix} \dot{\tilde{x}}_1 \\ \dot{\tilde{x}}_2 \end{bmatrix} = \begin{bmatrix} A_1 & 0 \\ A_{21} & A_2 \end{bmatrix} \begin{bmatrix} \tilde{x}_1 \\ \tilde{x}_2 \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} \tilde{u}_r \\
\tilde{y} &= \begin{bmatrix} C_1 & 0 \end{bmatrix} \begin{bmatrix} \tilde{x}_1 \\ \tilde{x}_2 \end{bmatrix}
\end{aligned}$$

where

$$\begin{aligned}
\tilde{x}_1 &= [m_a \quad w_e \quad v_x \quad v_y \quad v_z \quad z \quad \phi \quad \dot{\phi} \quad \theta \quad \dot{\theta} \quad \dot{\epsilon} \quad w_{fl} \quad w_{fr} \quad w_{rl} \quad w_{rr} \\
&\qquad \alpha \quad \tau_b \quad \beta \quad F_{fl} \quad F_{fr} \quad F_{rl} \quad F_{rr}]^T \\
\tilde{x}_2 &= [x \quad y \quad \epsilon \quad X \quad Y \quad yr \quad \dot{yr} \quad \epsilon_{des}]^T
\end{aligned}$$

In this form, both $\tilde{x}_1$ and $\tilde{y}$ are independent of $\tilde{x}_2$. Thus $\tilde{x}_2$ can be deleted from the model without affecting the transfer function from $\tilde{u}$ to $\tilde{y}$. Based on this observation, $\tilde{x}_2$ is removed from the model, which then becomes

$$\begin{aligned}
\dot{\tilde{x}}_1 &= A_1\tilde{x}_1 + B_1\tilde{u}_r \\
\tilde{y} &= C\tilde{x}_1
\end{aligned}$$

where A_1 is an 22×22 matrix, B_1 is an 22×3 matrix and C_1 is a 12×22 matrix.

As shown in (Douglas et al. 1995), when the nominal operating point associated with the linearized system is one where the car is not making a turn, the longitudinal and lateral dynamics decouple exactly. However, the case considered here has a nonzero nominal steering angle so the longitudinal and lateral dynamics do not decouple. All 22 states are included in the linear model order reduction process explained in the next section.

2.3 Reduced-Order Model

Previous manipulation involved no approximation. For further model simplification, some approximation must occur. After the linear model is derived, the first thing one should do is check the eigenvalues. Then, two approaches are presented to get reduced-order models. The first approach is to set the derivatives of certain fast states to zero. Using this philosophy, states with large negative eigenvalues can be dropped. However, a correction should be made using the deleted states to remove the steady state error. Consider a linear system modeled as:

$$\begin{aligned}\dot{x} &= Ax + Bu \\ y &= Cx + Du\end{aligned}$$

Suppose this model is written in a partitioned form

$$\begin{aligned}\begin{bmatrix}\dot{x}_1 \\ \dot{x}_2\end{bmatrix} &= \begin{bmatrix}A_{11} & A_{12} \\ A_{21} & A_{22}\end{bmatrix}\begin{bmatrix}x_1 \\ x_2\end{bmatrix} + \begin{bmatrix}B_1 \\ B_2\end{bmatrix}u \\ y &= \begin{bmatrix}C_1 & C_2\end{bmatrix}\begin{bmatrix}x_1 \\ x_2\end{bmatrix} + Du\end{aligned}$$

where x_2 contains the *fast states.* Set the derivative of x_2 to zero and solve the resulting equations for x_2 as a function of x_1 and u. This leads to

$$x_2 = -A_{22}^{-1}A_{21}x_1 - A_{22}^{-1}B_2u$$

Substitute this result into the expressions for $\dot{x}_1$ and y to obtain the reduced order model:

$$\begin{aligned}\dot{x}_1 &= \left[A_{11} - A_{12}A_{22}^{-1}A_{21}\right]x_1 + \left[B_1 - A_{12}A_{22}^{-1}B_2\right]u \\ y &= \left[C_1 - C_2A_{22}^{-1}A_{21}\right]x_1 + \left[D - C_2A_{22}^{-1}B_2\right]u\end{aligned}$$

this model preserves the static input-output relationships.

A second approach is to use balanced realization. Balancing refers to an algorithm which finds a realization that has equal and diagonal controllability and observability grammians. The diagonal of the joint grammian can be used to reduce the order of the model. Since the diagonal elements of the grammian, the Hankel singular values $g(i)$, reflect the combined controllability and observability of each state, it is reasonable to remove those states from the model for which $g(i)$ is small. Elimination of these states retains the most important input-output characteristics of the original system. After balanced realization has been done, a truncation is used to obtain a reduced-order model. For example, if the full-order model is

$$\begin{aligned}\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix} &= \begin{bmatrix}A_{11} & A_{12}\\ A_{21} & A_{22}\end{bmatrix}\begin{bmatrix}x_1\\ x_2\end{bmatrix}+\begin{bmatrix}B_1\\ B_2\end{bmatrix}u\\ y &= \begin{bmatrix}C_1 & C_2\end{bmatrix}\begin{bmatrix}x_1\\ x_2\end{bmatrix}+Du\end{aligned}$$

then, the reduced-order model is

$$\dot{x}_1 = A_{11}x_1 + B_1u$$

This is the approach originally proposed by Moore (Moore 1981). Using this approach it is possible to calculate a bound on the error introduced by deleting states.

At the end of the previous section, Section 2.2, a linear model is developed. Its eigenvalues are -227.45, -193.79, -159.51, $-132.21 \pm 2.62i$, $-135.35 \pm 1.85i$, -138.68, $-26.16 \pm 4.47i$, $-1.99 \pm 6.63i$, $-3.10 \pm 6.07i$, $-1.31 \pm 5.60i$, -0.046, $-7.09 \pm 2.48i$, -90.91, -1.25 and -80. Observe that ten of these eigenvalues are significantly larger than the rest. Two of the *fast* eigenvalues, -90.91 and -80, happen to be associated with the actuator dynamics. These modes should be retained if the linear model is to be used to design fault detection filters for actuator faults. From this we conclude that at least eight state variables can be dropped.

In method one, by looking at the eigenvectors corresponding to the large eigenvalues, the eight fast mode states are the four wheel speeds w_{fl}, w_{fr}, w_{rl}, w_{rr} and the four suspension forces F_{fl}, F_{fr}, F_{rl}, F_{rr}. Truncating these eight states produces a fourteenth-order model. In method two, the eight states with the smallest Hankel singular values are dropped. These methods combine the states in such a way that they lose their physical significance, so explicit identification of the deleted states is not possible. Table 2.1 summarizes a comparison of the two methods for model reduction.

The eigenvalues given in Table 2.1 show that the first method produces better results because the eigenvalues are closer to the full-order model eigenvalues. The second method truncates some *slow* states which results in a large change in the eigenvalues.

[redacted]					
Method 1	$-26.46 \pm 4.32i$ $-1.44 \pm 5.50i$	$-7.53 \pm 2.94i$ $-2.82 \pm 5.60i$	-0.049 -90.91	$-2.07 \pm 6.54i$ -80	-1.25
Method 2	$-108.98 \pm 37.67i$ -79.27	-0.046 $-3.07 \pm 6.10i$	$-18.33 \pm 13.19i$ $-2.39 \pm 6.30i$	$7.10 \pm 2.49i$ $-1.31 \pm 5.62i$	

Table 2.1: Eigenvalues for the vehicle dynamics using two model reduction methods.

Another test, based on frequency response, can also be performed to see which method is best. Singular values of the multivariable input to output frequency response are plotted from frequencies of 10^{-1} to $10^{2}\frac{\text{rad}}{\text{sec}}$. The reason for choosing this frequency range is that it roughly corresponds to that of the control inputs to a car. As shown in Figure 2.6, the error of largest singular value of the reduced-order model derived from method two is slightly better than the error of method one. However, the errors of the other two singular values of method two are much worse than the errors of method one.

An interpretation of this result is that the first order reduction method tends to preserve model fidelity with respect to each input while the second method tends to preserve model fidelity for only the most important input and output pair. For the purpose of fault detection filter design, the first method is more appropriate because fault detection filters are built for each control input.

A fourteen-state model is obtained using the first model order reduction method and is used subsequently to design fault detection filters for actuators faults. The system matrices A, B, C and D are given in Appendix D. The measured outputs are

y_{m_a}	Engine manifold air mass (kg).
y_{ω_e}	Engine speed ($\frac{\text{rad}}{\text{sec}}$).
$y_{\ddot{x}}$	longitudinal acceleration ($\frac{\text{m}}{\text{sec}^2}$).
$y_{\ddot{y}}$	lateral acceleration ($\frac{\text{m}}{\text{sec}^2}$).
$y_{\ddot{z}}$	heave acceleration ($\frac{\text{m}}{\text{sec}^2}$).
$y_{\dot{\phi}}$	roll rate ($\frac{\text{rad}}{\text{sec}}$).

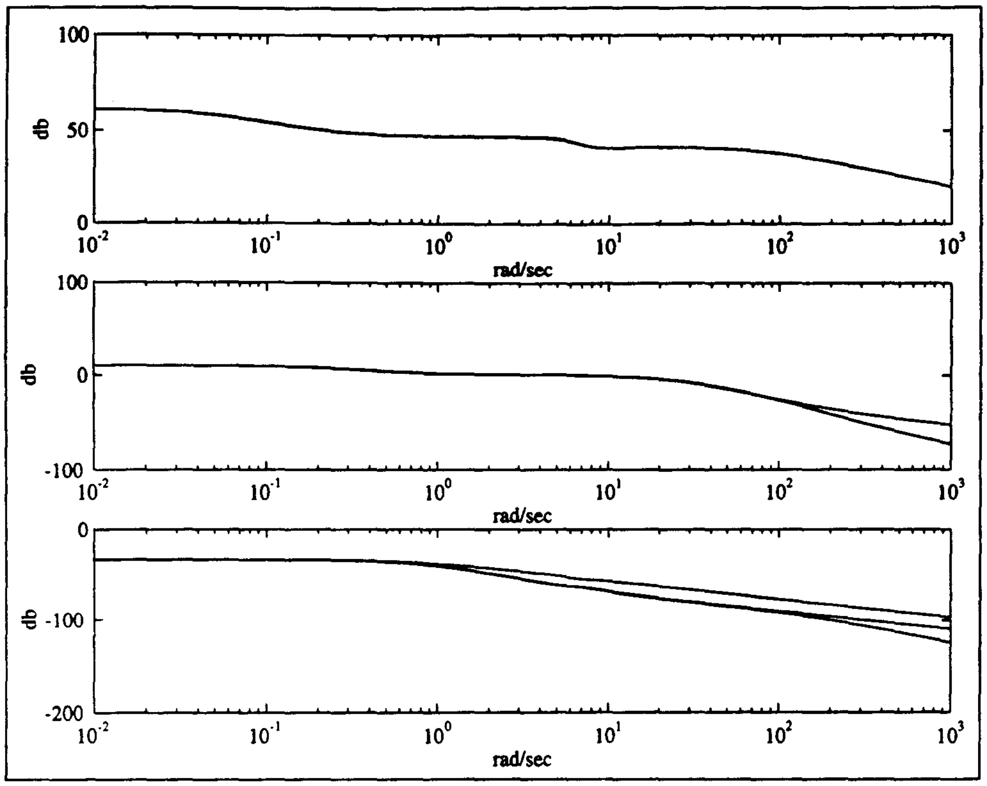

Figure 2.6: Singular value frequency response of full-order and fourteen state reduced-order models.

$y_{\dot{\theta}}$	pitch rate ($\frac{\text{rad}}{\text{sec}}$).
$y_{\dot{\epsilon}}$	yaw rate ($\frac{\text{rad}}{\text{sec}}$).
$y_{\omega_{fl}}$	front left wheel speed ($\frac{\text{rad}}{\text{sec}}$).
$y_{\omega_{fr}}$	front right wheel speed ($\frac{\text{rad}}{\text{sec}}$).
$y_{\omega_{rl}}$	rear left wheel speed ($\frac{\text{rad}}{\text{sec}}$).
$y_{\omega_{rr}}$	rear right wheel speed ($\frac{\text{rad}}{\text{sec}}$).

and the control inputs are

α	Throttle angle (deg).
τ_b	Brake torque (Nm).

β Steering angle (rad).

While a fourteen-state linear model is used for the design of actuator fault detection filters, a twelveth-order model is used to design sensor fault detection filters. Recall that two fast modes retained in the fourteen-state reduced-order model are associated with the actuator dynamics. The eigenvalue -90.91 is associated with the throttle actuator and -80 with the steering actuator. For the design of sensor fault detection filters, these modes may also be deleted. Note that since the actuator dynamics are in series with the other dynamics, the reduced-order eigenvalues do not change. Singular values of the multivariable input to output frequency response are illustrated in Figure 2.7. The model reduction error is seen to be very slightly worse than for the fourteenth-order model. The system matrices

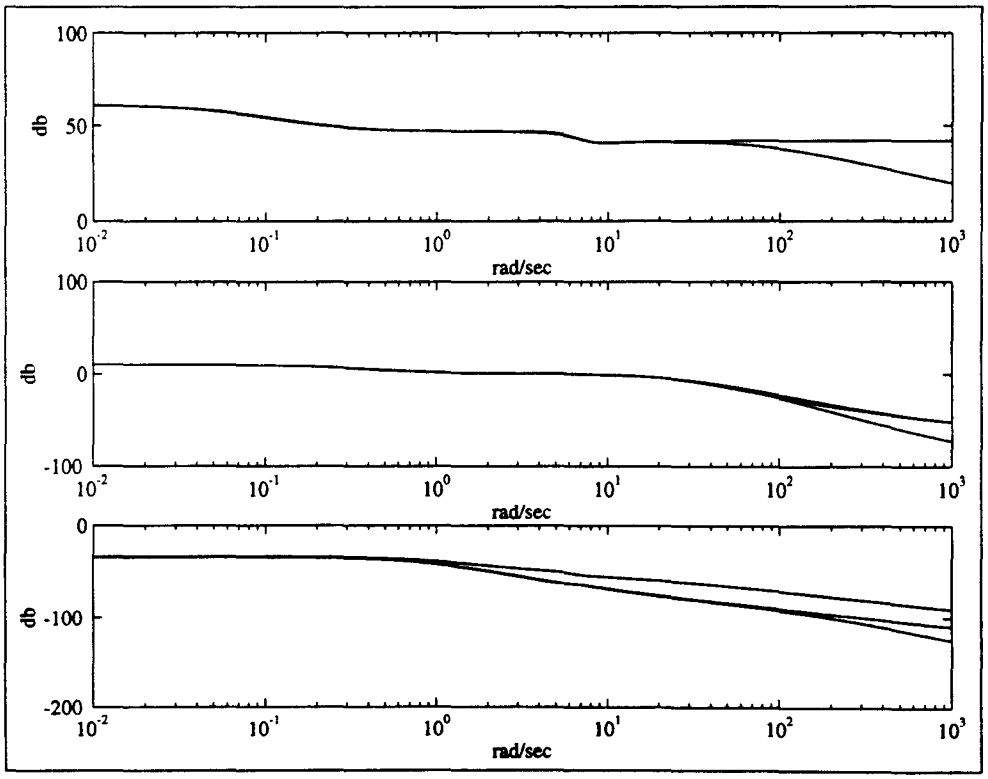

Figure 2.7: Singular value frequency response of full-order and twelve state reduced-order models.

A, B, C and D for the twelve-state reduced-order model are given in Appendix D.