A Regression-based Approach for Simulating Feedforward Active Noise Control, With Application to Fluid-Structure Interaction Problems

by

Christopher E. Ruckman

Dissertation submitted to the faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Department of Mechanical Engineering

APPROVED:

Dr. Chris R. Fuller, Chairman

Dr. Ricardo A. Burdisso

Dr. Yuan-Ning Liu

Dr. Harry H. Robertshaw

Dr. Craig A. Rogers

22 January, 1994

Blacksburg, Virginia

A Regression-based Approach for Simulating
Feedforward Active Noise Control,
With Application to Fluid-Structure Interaction Problems

by

Christopher E. Ruckman
Dr. Chris R. Fuller, Chairman
Department of Mechanical Engineering

(ABSTRACT)

This dissertation presents a set of general numerical tools for simulating feedforward active noise control in the frequency domain. Feedforward control is numerically similar to linear least-squares regression, and can take advantage of various numerical techniques developed in the statistics literature for use with regression. Therefore, an important theme of this work is to look at the control problem from a statistical point of view, and explore the analogies between feedforward control and basic statistical principles of regression.

Motivating the numerical approach is the need to simulate active noise control for systems whose dynamics must be modeled numerically because analytical solutions do not exist, e.g., fluid-structure interaction problems. Plant dynamics for examples in the present work are modeled using a finite-element / boundary-element computer program, and the associated numerical methods are general enough for us with many types of problems. The derivation is presented in the context of active structural-acoustic control (ASAC), in which sound radiating from a vibrating structure is controlled by applying time-harmonic vibrational inputs directly on the structure.

First, a feedforward control simulation is developed for a submerged spherical shell using both analytical and numerical techniques; the numerical formulation is found by discretizing the integrations used in the analytical approach. ASAC is shown to be effective for controlling radiation from the spherical shell. For a point-force disturbance at low frequencies, a single control input can reduce the radiated power by up to 20 dB (ignoring the possibility of measurement noise). A more general numerical methodology is then developed based on weighted least-squares regression in the complex domain. It is shown that basic regression diagnostics, which are used in the statistics literature to describe the quality and reliability of a regression, can be used to model the effects of error sensor measurement noise to produce a more realistic simulation. Numerical results are presented for a finite-length, fluid-loaded cylindrical shell with clamped, rigid end closures. It is shown that when the controller reduces the radiated power by less than 2 dB, the control simulation is usually invalid for statistical reasons. Also developed are confidence intervals for the individual control input magnitudes, and prediction intervals which help evaluate the sensitivity to measurement noise for the regression as a whole.

Collinearity, a type of numerical ill-conditioning that can corrupt regression results, is demonstrated to occur in an example feedforward control simulation. The effects of collinearity are discussed, and a basic diagnostic is developed to detect and analyze collinearity. Subset selection, a numerical procedure for improving regressions, is shown to correspond to optimizing actuator locations for best control system performance. Exhaustive-search subset selection is used to optimize actuator locations for a sample structure. Finally, a convenient method is given for investigating alternate controller formulations, and examples of several alternate controllers are given including a wavenumber-domain controller. Numerical results for a cylindrical shell give insight to the mechanisms used by the control system, and a new visualization technique is

used to relate farfield pressure distributions to surface velocity distributions using wavenumber analysis.

Acknowledgments

Any doctoral degree is a credit not only to the recipient but to those around him. I was fortunate to have support from many more people than I can list here. First and foremost I wish to thank Dr. Chris Fuller, who proposed the original line of research and took a very active role in guiding my progress and developing the final product. Dr. Fuller helped me in ways too numerous to recount, and was a inspiration to me both professionally and personally. The members of my Committee made many helpful suggestions along the way, especially during the oral examination. Dr. Y.-N. Liu was particularly helpful. I would also like to thank: Felix Rosenthal for discussions regarding singular-value decomposition; Gordon Everstine and Tony Quezon for help with NASHUA; R. Vasudevan for discussions about dispersion curves and acoustic radiation; my parents for supporting my decision to drop everything and go back to school; Dawn Williams, Jeff Vipperman, and Jim Collier for a thousand small (and not-so-small) favors. Dr. Paul Shang and Dr. Ted Farabee provided essential guidance and moral support at difficult moments.

I would like to acknowledge the financial support of Carderock Division, Naval Surface Warfare Center. I received a one-year Extended-Term Training grant and generous support from the IR/IED program under Bruce Douglas. The Center also provided me with computing facilities that proved invaluable. Without their support this work would not have been possible. I would also like to acknowledge the financial support of the Office of Naval Research (Grant no. ONR-N00014-88-K-0721) during the first two years of this project. Kam Ng and Geoff Main of ONR were instrumental in helping me decide what material to include in the dissertation.

Through the wonders of USENET and electronic mail, I had fascinating conversations with statisticians around the world. Of particular help were Alan Miller (CSIRO, Australia), who provided source code from his book and reviewed an early manuscript, and Ken White (University of British Columbia) who also reviewed some of my work. Both deserve my thanks.

Finally, to my wife, Diane Wilson, I owe the deepest thanks of all. I have heard it said that doctoral students stand a higher-than-average risk for marital problems, but I think this grueling ordeal actually brought us closer together. Without Diane's support, especially during the final few months of research and thesis-writing, my sanity and health would surely have suffered. I say to her now, in writing: there will be paybacks!

Table of Contents

List of Illustrations

List of Tables

CHAPTER 1: INTRODUCTION

1.1. Background and motivation

The goal of this research is to develop a general numerical frequency-domain approach for simulating the performance of feedforward active noise control systems. With active control gaining popularity as a practical solution for low- to mid-frequency noise problems, researchers are considering systems that are more complex and ambitious than those now in use. Feedforward control has proven effective when the disturbance is periodic and a reference signal is available. By working in the frequency domain, we can in some cases derive convenient simulations to predict how well the control system will perform. However, such simulations usually require a closed-form solution for the system dynamic response. Some problems of practical interest, including most structural-acoustic interaction problems, do not posses closed-form solutions: plant dynamics must be found with numerical methods such as finite-element methods (FEM) or boundary-element methods (BEM). The present work formulates and validates algorithms for simulating active control when plant dynamics are modeled numerically.

Another important theme of this work is to illustrate how numerical techniques developed in the statistics literature may be applied to active control problems. If we examine feedforward active control in the frequency domain, the process of solving for control inputs is numerically similar to multiple linear least-squares regression. However, the analogy extends beyond simply solving for the control inputs. Experience has shown that analytical predictions are nearly always optimistic, i.e., physical implementations do not perform as well as predicted. One cause of the disparity is error sensor measurement noise, which is usually ignored by simulations. A second reason is that

analytical simulations usually presume to know the system response as a continuous function of space, whereas a real control system knows the response only at a finite number of discrete error sensor locations. Both problems are analogous to problems encountered in performing regressions, and their effects are conveniently addressed by adapting regression diagnostic techniques to the active control problem. To the author's knowledge, the statistical techniques used in this thesis have not been applied in the context of active noise control.

While most of the ideas discussed in this thesis apply to any feedforward control system, for purposes of illustration we focus on active structural-acoustic control (ASAC). ASAC involves controlling the acoustic response of a fluid-structure system by applying time-varying force inputs directly on the structure. Recent theoretical and experimental studies use ASAC to control radiation from 1-dimensional and 2-dimensional structures including beams, plates, submerged plates; similar techniques are used to control noise transmitted through flat plates or into infinite cylindrical shells [1-12]. The concept is similar to active vibration control (AVC) since the actuators are applied directly on a flexible structure, but the goal of reducing acoustic response often differs from the goal of reducing a purely structural response. ASAC also differs from active noise cancellation (ANC), since ASAC actuates the structure itself rather than actuating the acoustic medium with loudspeakers. For structure-borne noise, ASAC can often produce widespread farfield reductions with fewer actuators as compared to AVC and ANC. Also, because vibrational actuators tend to be more compact than loudspeakers, ASAC can be used in certain situations where ANC is impractical.

One reason for developing control simulations based on numerical approaches is the need to apply ASAC to submerged, thin-walled shells and other fluid-loaded structures. Submerged shells are of significant practical interest, but they also represent a fundamental departure from existing

work. Because no analytical expressions exist for the dynamic response of a general, 3-dimensional, fluid-loaded shell, analytical simulations of active control cannot be used. In the present work the fluid-structure dynamic response is found from transfer functions obtained numerically. The method is therefore quite general: we can develop an ASAC simulation for any system that can be modelled numerically. Two similar studies exist. Mollo and Bernhard [13] use the indirect boundary-element method to formulate active noise cancellation systems for both interior and exterior noise fields. A related technique is proposed by Song, Koopman, and Fahnline [14], who use the principle of superposition to solve for source strengths to cancel radiation from an arbitrary structure. These works do not model fluid-structure interaction, and they model the structure only as a velocity distribution rather than modeling the structure implicitly as with a FEM model. Another recent work is Nelson and Elliott [15], which presents both analytical and numerical techniques for frequency-domain feedforward problems. None of these works address modelling the effects of error sensor measurement noise, nor do they discuss the other statistical aspects of the feedforward control problem discussed in this thesis.

A numerical simulation capability would help the development of *concurrent passive/active design* techniques. Most ASAC systems are *sequential* designs: first one designs and builds a structure from strength requirements, weight requirements, and/or passive acoustic considerations, and then one retrofits an active control system if needed. In a concurrent design approach, both structural properties and active control features would be considered simultaneously during design. Concurrent design would likely produce a structure more amenable to active control than one designed sequentially. Concurrent design approaches are beyond the scope of the present work, but basing the control simulation on a finite-element structural model would enhance a concurrent design approach because of the close linkage between structural and active control models.

Since the physical locations of actuators can dramatically influence system performance, a concurrent design approach would benefit from a general method for optimizing actuator locations. For geometrically simple systems the best actuator locations may be obvious; for more complicated systems, perhaps not. Also, the performance bottleneck for some control systems involves not algorithms or hardware, but weak coupling between the actuators and system response (poor controllability). Both these points are reasons to seek a formal, structured procedure for optimizing actuator locations. However, methods commonly in use are not computationally efficient, and are therefore limited to simplified problems. An optimization approach better suited to discretized structural models would provide computational efficiency gains, allowing the study of larger and more practical problems. One such approach is discussed further in Section 1.2.4, and a detailed derivation with numerical results is presented in Chapter 5.

The problem of simulating ASAC for fluid-loaded shells can be cast into the larger class of structural-acoustic systems that do not posses analytical solutions. Therefore, the overall research objectives may be set out as follows:

1. Apply known analytical methods to develop an analytical ASAC simulation for a fluid-loaded spherical shell, and then discretize the analytical approach to obtain a numerical approach suitable for use with spherical or non-spherical geometries. Validate the numerical method by comparing results with the analytical method for both the uncontrolled and controlled cases.
2. Conduct a survey of elementary statistical techniques that have direct bearing on the ASAC problem. Develop and study statistical aspects of ASAC performance, emphasizing methods for modeling the effects of error sensor measurement noise and assessing model reliability.
3. Develop a general method for optimizing actuator locations for best control system performance. The method should take advantage of computational opportunities that exist when the problem is discretized and numerical methods such as FEM are used to model the structural-acoustic response.

4. Give numerical examples for several alternate controller formulations in the spatial domain and in the wavenumber domain, using a non-spherical structural geometry; use the results to shed light on mechanisms of ASAC for fluid-loaded structures.

1.2. Overview

The entire investigation is conducted in the frequency domain rather than the time domain. The disturbance signal is time-harmonic with time dependence $e^{i\omega t}$, and thus only one frequency component and is stationary in time. We assume steady-state operation such that all transient responses have been damped out, and we express the dynamic response as a complex-valued spatial distribution in the frequency domain. This simplification provides insight to the more general problem of multi-frequency or broad-band disturbances, which can always be decomposed into their frequency-domain components. For simplicity of notation, however, all time-varying quantities are shown without the explicit time dependence $e^{i\omega t}$.

Figure 1 shows a block diagram of a general single-input, single-output (SISO) feedforward active control system. A *primary* or *disturbance source* attached to the structure injects energy (the *disturbance input*) into the structure. The actuators used to influence system response are termed *secondary sources* or *control inputs*. The system contains both a *detection sensor*, which measures the disturbance input, and an *error sensor*, which measures the system output. The *disturbance input S* propagates through the *primary path A* to form the *disturbance response* P_p. The disturbance input is also measured by a detection sensor, which possibly adds measurement noise ε_d. The signal from the detection sensor is operated on by the controller to produce the control input b, which propagates through the secondary path X to form the control response P_c. The total response $P_p + P_c$ is measured by the error sensor, which possibly adds measurement noise ε; the error sensor output is used in determining the controller gain.

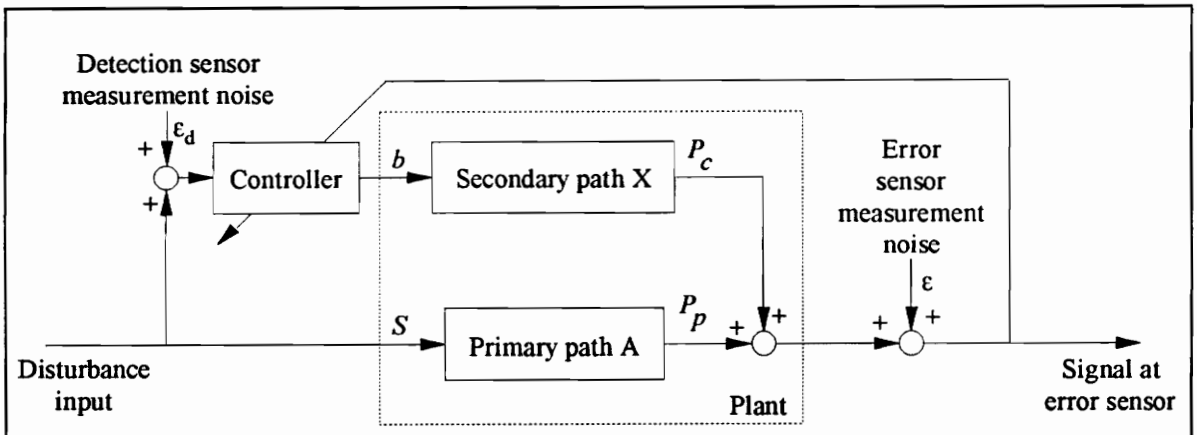

Figure 1: Block diagram of general SISO feedforward control system with measurement noise occurring at detection sensor and error sensor.

Measurement noise in the error sensor is analogous to observation noise in regression, and therefore various statistical methods are available for investigating the effects of measurement noise in the error sensor signal. We shall not treat the case of measurement noise in the detection sensor signal, since that type of noise has no direct analog in the regression problem. Both single-input and multi-input cases may be treated, with the multi-input cases being similar to Fig. 1.

1.2.1. Validating the numerical method: the spherical shell

The thin-walled spherical shell is often used to build confidence in numerical procedures for structural-acoustic interaction problems because it is one of the few separable geometries for which closed-form analytical solutions are available. The first step in developing a numerical approach for ASAC simulations is to develop an analytical method for simulating ASAC for a fluid-loaded spherical shell. The next step is to develop a numerical method by discretizing the analytical method, and compare results between the numerical and analytical methods. These comparisons and their development are the subject of Chapter 2.

The analytical methods of Nelson et. al. [16-20] are used to obtain the analytical solution, whose objective is to reduce farfield radiation from a fluid-loaded spherical shell excited by a point-force disturbance. The total radiated power is formulated from a linear sum of the primary (disturbance) response plus the secondary (control) response. The simulation consists of solving for a set of complex-valued actuator strengths (control forces) that minimize a radiation-related cost function. We develop the simulation using both analytical (closed-form) and discretized (numerical) approaches, and then compare the analytical results to the numerical results. Treatment of error sensor measurement noise, however, is reserved for Chapter 3.

1.2.2. A more general method: using least-squares regression

A more general method than the method of Chapter 2 involves posing the feedforward control problem as a linear least-squares regression of complex-valued variables. This type of analysis was first proposed by Snyder [21] and by Snyder and Hansen [22]. Forming the problem as a regression provides an appealingly logical framework for discussing and analyzing system performance. Furthermore, because least-squares regression is a mature topic supported by a large body of literature and algorithms from the statistical sciences, we can adapt well-established regression methods to the feedforward control problem. The most useful example is regression diagnostics, which allow a way to model the effects of error sensor measurement noise. For example, the "F-test" measures whether the control simulation is "statistically significant", and prediction intervals indicate whether the results are sensitive to measurement noise. (Similar statistical concepts are used in applying system identification procedures both in the frequency domain and in the time domain; see, for example, [23] or [24]). More importantly, regression diagnostics can be used to generate confidence bounds describing how well the controller cancels the offending sound field. Chapter 3 examines feedforward control in light of applicable statistics

concepts, and discusses some basic diagnostic procedures. Example results are presented for a clamped, finite-length, axisymmetric, cylindrical shell with fluid loading.

It should be noted that all the techniques developed in this work are equally applicable to 2-dimensional and 3-dimensional problems. The structures treated in Chapter 3 are axisymmetric and thus effectively 2-dimensional. Extending the techniques to 3-dimensional problems is straightforward, but produces much more stringent requirements for computing time and output postprocessing. A simple 3-dimensional case is treated in Chapter 2 by superposing the axisymmetric solutions for the spherical shell to obtain a radiated field with a plane of symmetry. Similar techniques may be used with the procedures outlined in Chapters 3-6.

1.2.3. Collinearity and the need for collinearity diagnostics

It is well-known in the statistics community that regressions are prone to a type of numerical ill-conditioning known as "collinearity," but it is less well-known in the active control community that frequency-domain simulations of active control can suffer from collinearity as well. Such ill-conditioning is mentioned by Rosenthal [25], who notes that redundancies between the sensors in an active noise filtering system can produce numerical problems limiting the number of sensors that can be used. The issue is also discussed by Nelson and Elliott [15], who use iterative gradient descent solution methods to circumvent the collinearity problem. The purpose of Chapter 4 is to explain how collinearity enters an active control system and how it affects feedforward control simulations. Also discussed are a set of diagnostic techniques analogous collinearity diagnostics used in the statistics literature. A collinearity diagnostic is presented as a step-by-step outline, with sample results given for a clamped cylindrical shell with fluid-loading. When the sole purpose of the diagnostic is to detect rather than analyze collinearity, the collinearity diagnostic consists of simply examining the condition number of a special scaled design matrix.

1.2.4. Optimizing actuator locations for best possible system performance

Chapter 5 proposes a novel method for optimizing the locations of multiple actuators in a feedforward ASAC problem. The method is computationally efficient because it takes advantage of numerical properties already present in the simulation, and therefore it can be used in situations for which other methods are unsuitable. Geometrically complex structures require multiple actuators whose optimal placement may be difficult to anticipate solely from knowledge of the system dynamics. Optimizing actuator placement by manual trial and error can be costly and time-consuming, especially when many possible configurations exist. The method given in Chapter 5 allows for comprehensive studies using large numbers of actuators. The combinatorial nature of the method is ideally suited for practical design problems, in which the control actuators can often be placed only in certain discrete locations rather than arbitrarily on the structure. Rather than providing a single "optimum" actuator configuration, the method used here provides access to a large number of near-optimal configurations.

1.2.5. Method for investigating alternate controller formulations

In formulating the control problem as a regression, we can conveniently model alternative controller types by using the same analysis with different types of sensors and, therefore, different cost functions. This approach can be useful for performing numerical experiments to explore or optimize a system configuration. By taking weighted sums of various structural or acoustical quantities, we can derive quadratic, frequency-dependent "response functions" that measure different types of vibrational or acoustic energies. Designating one of the available response functions to be the "cost function" minimized by the controller, we can obtain different controller formulations depending on which response function is minimized. The quantity from which the cost function is computed would determine the type of error sensor used in the control system.

For example, we can derive a response function based on farfield pressure such that, with an appropriate weighting function, the response function is proportional to the total radiated power. If we use this radiated power response function as our cost function, we are effectively using many farfield pressure sensors as error sensors and minimizing the radiated power. However, we could just as easily derive a response function that measures radiation into some chosen portion of the farfield rather than the entire farfield, and use that as the cost function. We could use surface vibration data to derive a kinetic energy response function, or transform the vibration data into the wavenumber domain to derive a wavenumber-domain response function. Chapter 6 explores these suggestions and presents numerical examples for the clamped, cylindrical shell with fluid loading. The approach involves forming various response functions and examining how they change depending on which functions is used as the cost function in the controller. The results help illustrate the mechanisms used to control radiation, and could be used to develop guidelines for implementing an active control system.

1.3. Numerical modeling of fully-coupled dynamic response

Numerical solutions for structural-acoustic dynamic responses are found using the computer program NASHUA [26]. NASHUA calculates low frequency, farfield acoustic pressure field radiated or scattered by an arbitrary, submerged, three-dimensional or axisymmetric elastic structure with internal or external time-harmonic loading. The finite-element program NASTRAN is used to compute structural mass, damping, and stiffness matrices, which are then used with a discretized Helmholtz surface integral equation to obtain the fully-coupled response of the fluid-structure system. In the present research, all structures are analyzed using the axisymmetric capability of NASHUA to simplify analysis and reduce computing costs.

There are two significant benefits to using an approach based on a FEM/BEM model of the structure/fluid system. Firstly, altering the structural geometry involves minimal effort because the boundary element formulation does not require discretization of the ambient fluid. Secondly, the FEM/BEM approach can treat nearly any structure amenable to FEM modeling, making it an excellent companion for theoretical and experimental studies of submerged shells. Both these features work to support the development of concurrent passive/active design techniques as discussed above.

The following sections describe the geometries, coordinate systems, and NASHUA models for the two structures described in this thesis: a spherical shell and a cylindrical shell. Both models use axisymmetric cone elements that include both bending and membrane stresses. The structures are undamped steel shells submerged in water. The shell material has Young's modulus of 1.85×10^{11} Pa, density of 7670. kg/m^3, and Poisson's ratio of 0.3. The fluid medium surrounding the structure has a density of 1000. kg/m^3 and an acoustic velocity of 1460. m/s. All farfield quantities are normalized to a distance of 1.0m from the structure, i.e., the pressures are calculated using farfield equations but a radius of 1.0m is assumed for purposes of comparison with other calculations.

1.3.1. Numerical modeling of spherical shell

The first structure to be examined is a thin-walled spherical shell, shown in Fig. 2. The shell is subject to a concentrated point-force disturbance (not shown) acting normal to the shell surface. The global z-axis is defined as the axis along which the disturbance force acts. Farfield locations are defined in global spherical coordinates (R,θ,ϕ), where R is the radial distance, θ is the latitudinal angle measured from the global z-axis, and ϕ is the longitudinal angle measured from the x-z plane. The structure is completely submerged in a dense fluid, with material properties

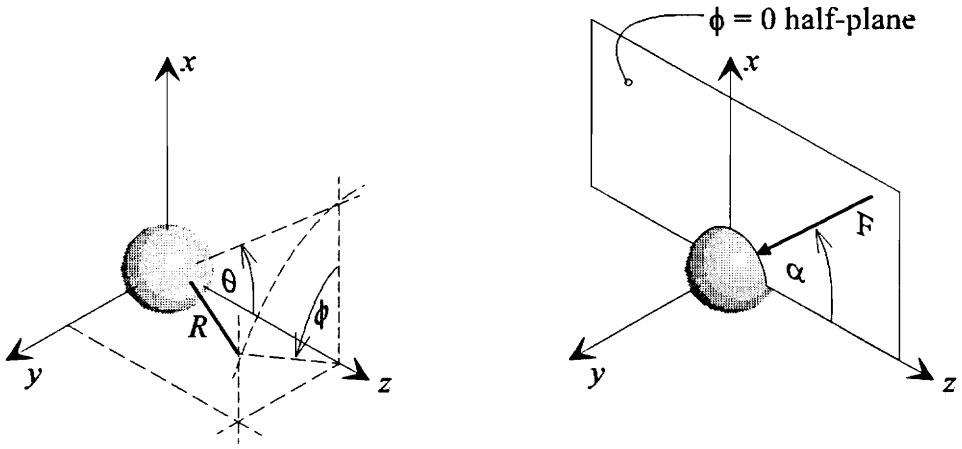

a) farfield locations

b) control force locations

Figure 2: Coordinate systems for defining a) locations in the farfield, and b) locations on the shell. Control forces are restricted to lie within the ϕ=0 half-plane, so that force locations are defined by a single angle α.

chosen for a steel shell in water. The ratio of shell thickness to radius is 0.01, and the shell radius is a=1m. We examine frequencies up to $k_o a$=1.7, where k_o is the acoustic wavenumber; this frequency range includes the resonance frequencies of several low-order modes.

The finite-element portion of the NASHUA spherical shell model contains 129 grid circles. Since NASTRAN does not allow grid circles with zero radius, the poles at θ=0 and θ=π are actually "pinholes" or grid circles of radius 0.001a. To approximate a point force, we apply a ring-force on one of the small-radius grid circles. This approximation casts some doubt on the validity of the model at high frequencies. However, even with this approximation the agreement between analytical and numerical results is sufficient for the problem at hand, as is shown in Chapter 2.

1.3.2. Numerical modeling of cylindrical shell

The second structure to be examined is a finite-length, fluid-loaded cylindrical shell with rigid, flat end closures as shown in Fig. 3. The shell length L is ten times the shell radius a, and the ratio of wall thickness to shell radius is 0.005. The global z-axis is coincident with the shell symmetry axis, with the right-hand end closure intersecting the positive z-axis. Farfield locations are described by the same global spherical coordinates (R,θ,ϕ) used for the spherical shell. Surface locations are described by a normalized distance coordinate $\delta=(z+L/2)/L$, where $\delta=0.0$ at the left-hand end closure and $\delta=1.0$ at the right-hand end closure. (The farfield location $\theta=0$ corresponds to the positive z-axis.) We examine frequencies as high as $k_0a=0.95$.

The cylindrical portion of the NASHUA finite-element model contains 50 grid circles. The end closures, being rigid, contain only 4 grid circles each. The grid circles at which the cylinder intersects the end closures are constrained both in translation and in rotation, so that the cylindrical section is clamped. Three types of forcing functions are examined (see Fig. 3b). The first is an axisymmetric ring-force, that is, a concentrated line force applied normal to the shell surface. The second type is an axisymmetric ring-moment about the circumferential direction. The third type is an axisymmetric distributed force whose amplitude varies as $\cos(3\pi\delta)$, not shown in Fig. 3b.

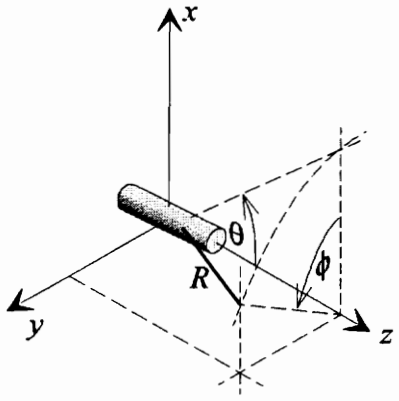

a) farfield locations

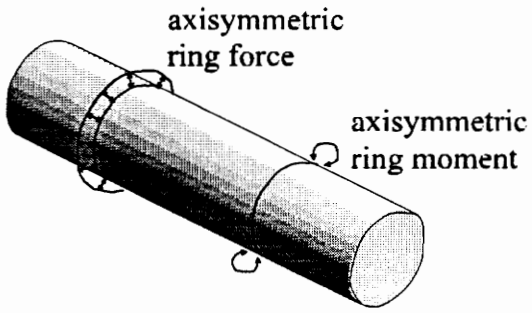

b) typical control forces

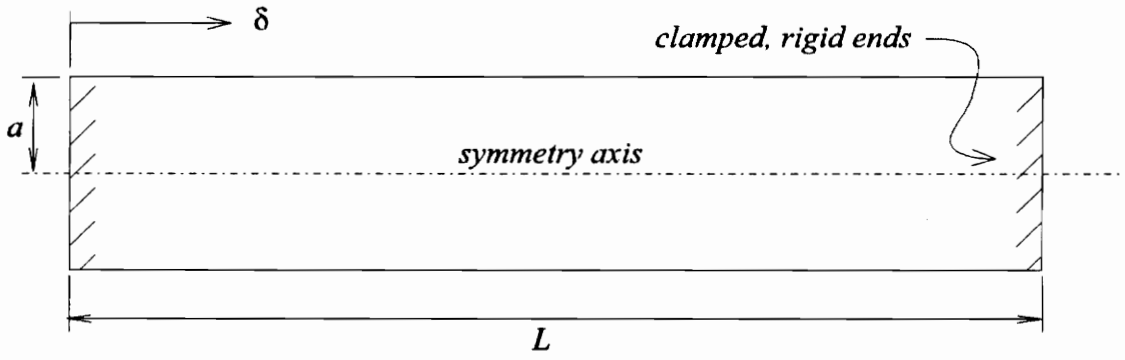

c) surface locations

Figure 3: Geometry of clamped cylindrical shell. Figure 3a: Farfield coordinate system for cylindrical shell. Figure 3b: Cylindrical shell has length L=1.0m and radius a=0.1m, with clamped, rigid end closures; shell material is steel, ambient fluid is water. Figure 3c: Typical control forces; actuators are concentrated line loads applied either as ring-forces in the radial direction or ring-moments about the circumferential direction.

CHAPTER 2: ANALYTICAL AND NUMERICAL STUDIES OF SPHERICAL SHELL WITH FLUID LOADING

In this chapter we present the analytical and numerical models of the thin-walled, fluid-loaded spherical shell, and derive the feedforward control approach used to find control force magnitudes and phases. We also present results to validate the numerical approach by comparing with analytical results. Some of the results are also presented in [12]. For this chapter we shall not treat the effects of measurement noise at the error sensors, discussion of which is postponed until Chapter 3.

2.1. Theoretical background

2.1.1. Analytical modeling of spherical shell

To simplify the problem, we consider a restricted set of forcing functions for control forces. The forces are concentrated loads (point forces) normal to the shell surface, and are applied only in the $\phi=0$ plane as shown in Fig. 2b. Force locations are specified by a single angle α. The disturbance is always a single point-force applied at $\alpha=0$. The total farfield response is therefore either axisymmetric (if a single control force is applied at $\theta=\pi$), or symmetric about the $\phi=0$ plane (if control forces are applied off the z-axis.)

Multi-force responses can be found by combining the axisymmetric responses due to each of the individual forces using superposition, even if their combined response is not axisymmetric. Given the restrictions on force locations described above, the coordinate transformation used to

find radiation and shell vibration due to any individual force is a simple rotation whose effective rotation angle can be found from the following relation.

$$\cos\theta_{eff} = \left\{\cos^2\theta + \sin^2\theta\sin^2\phi\right\}^{1/2}\cos\left\{\alpha + \tan^{-1}\left(\frac{\sin\theta\sin\phi}{\cos\theta}\right)\right\} \tag{1}$$

The analytical model for a thin-walled, fluid-loaded, spherical shell with a normal, persistent, point-force is due to Junger and Feit [27]. The model includes membrane stresses but ignores bending stresses, and is therefore limited to low frequencies. Omitting the e^{iwt} time dependence, the farfield pressure due to a force b_i applied on along the z-axis is

$$P(R,\theta,b_i) = \sum_{n=0}^{\infty} b_i \Gamma_n(R) P_n(\cos\theta) \tag{2}$$

where $\mathbf{P}_n$ is the Legendre polynomial of order n, and

$$\Gamma_n(R) = \frac{2n+1}{4\pi a^2}\left(\frac{\rho c}{k_a R}\right)\frac{\left(-\sqrt{-1}\right)^n}{(Z_n + z_n)h'(k_a a)}. \tag{3}$$

The symbols ρ and c represent the fluid density and acoustic velocity, h' is the derivative of the spherical Hankel function, and Z_n and z_n are the modal structural and acoustic impedances as defined in [27] .

2.1.2. Feedforward control approach

The control approach taken here is a single-frequency, steady-state, feedforward approach based on linear quadratic optimal control theory [6,16,17,20]. Consider a spherical shell subjected to a a total of N_d coherent, persistent, time-harmonic disturbance forces. The *disturbance response* P_p, also known as the *primary response*, is defined as the farfield pressure due to the disturbance forces. In the absence of measurement noise at the error sensors, the

disturbance response corresponds to the *open-loop* signal measured by the error sensor in Fig. 1. To make the development more clear we stipulate that the disturbance response must be axisymmetric, $P_p = P_p(\theta)$. In Section 2.1.4 we describe how to use the same techniques for non-axisymmetric problems. Suppose the i^{th} disturbance force has magnitude s_i, and would cause farfield pressure distribution $A_i(\theta)$ if acting alone. Assuming the force magnitudes are small enough that the response is linear, the disturbance response can be written as a linear sum of the individual disturbance responses.

$$P_p(\theta) = \sum_{i=1}^{N_d} A_i(\theta) s_i \,. \tag{4}$$

Suppose we now apply a total of N control forces to control the disturbance response, where the i^{th} control force has magnitude b_i. The *control response* $\hat{P}$, also known as the *secondary response*, is defined as the farfield pressure due to only the control forces. The control response corresponds to the controller output after propagating through the error path in Fig. 1. Now we stipulate, again for convenience of formulation, that both the control response and the disturbance response must be symmetric about the same axis so that their sum is axisymmetric. Thus the control response $\hat{P}(\theta)$ is a linear sum of the individual control responses $X_i(\theta)$.

$$\hat{P}(\theta) = \sum_{i=1}^{N} X_i(\theta) b_i \,. \tag{5}$$

The total response is the sum of the disturbance response plus the control response; this residual corresponds to the closed-loop signal being measured by the error sensor in Fig. 1.

Now let us define a quadratic *cost function* χ^2 as a weighted integral of squared magnitude of the total response over the farfield domain,

$$\chi^2 = \int w(\theta) \left| P_p(\theta) + \hat{P}(\theta) \right|^2 d\theta, \tag{6}$$

where $w(x)$ is a weighting function that can be chosen to give physical significance to χ^2 if desired. The feedforward control problem can be restated as an unconstrained quadratic minimization, i.e. we seek a set of control forces b_i to minimize χ^2. Recalling that the total radiated power is

$$\Pi = \int \frac{\pi R^2 \sin\theta}{\rho c} |p(\theta)|^2 \, d\theta, \tag{7}$$

we can obtain a cost function equal to the total radiated power by setting

$$w(\theta) = \frac{\pi R^2}{\rho c} \sin\theta. \tag{8}$$

To minimize χ^2 we form its derivatives with respect to each of the control forces, and require that all the derivatives vanish. This produces an N x N system of complex-valued, simultaneous equations, the so-called *normal equations*, which have as their solutions the complex control force magnitudes b_i. We can generate the coefficients of the normal equations in either of two ways, depending on whether $P_c(\theta)$ and $\hat{P}(\theta)$ are given as continuous functions of θ or as discretized functions defined only at discrete values of θ. If we have continuous expressions for the variables, we can integrate Eq. (6) in closed form. On the other hand, if the variables are defined only at discrete locations, we can generate discrete (analytical or numerical) transfer functions and integrate Eq. (6) numerically. These two methods are outlined below.

Note that each of the variables A_i and X_i from Eqs. (4-5) is merely a transfer function between a unit force somewhere on the structure and the farfield pressure (which may be either a continuous pressure distribution or a set of pressures at discrete locations). In the present context these transfer functions are found using either the FEM/BEM solution or an analytical solution. However, for purposes of simulation they could be obtained from any source such as experimental data or other numerical solutions. Note that although Eq. (6) contains a line integral for a one-

dimensional (axi-symmetric) system, it could easily be extended to an integral over an area or through a volume.

2.1.3. Obtaining the normal equations from continuous variables

Since we have an analytical solution available, we can formulate the problem in terms of continuous functions of farfield location θ. This allows us to integrate Eq. (6) in closed form, evaluate the coefficients of the normal equations, and solve for the required control forces. Suppose we write the control forces and disturbance forces in vector form so that

$$B = \{b_1 \quad b_2 \quad \cdots \quad b_N\}^{\mathrm{T}} \quad \text{and} \quad S = \{s_1 \quad s_2 \quad \cdots \quad s_N\}^{\mathrm{T}}. \tag{9}$$

We now write the disturbance response as $P_p(\theta)=AS$ and the control response as $\hat{P}(\theta)=XB$, where A and X are row vectors containing the continuous variables $A_i(\theta)$ and $X_i(\theta)$ from Eqs. (4) and (5), respectively. After some algebra, the cost function from Eq. (6) becomes

$$\chi^2 = \int B^T X^T X^* B^* w(\theta) d\theta - 2\,\mathrm{Re}\left\{\int S^T X^T A^* B^* w(\theta) d\theta\right\} + \int S^T A^T A^* S^* w(\theta) d\theta. \tag{10}$$

Using the orthogonality of the Legendre polynomials in Eq. (2), we can perform the integrations above in closed form and reduce the cost function to

$$\chi^2 = B^T C B^* - 2\,\mathrm{Re}\{S^T D B^*\} + S^T E S^*, \tag{11}$$

where a typical element of the matrix C is

$$C_{ij} = \frac{\pi R^2}{\rho c} \sum_{n=0}^{\infty} \Gamma_n(R)\Gamma_n^*(R) \frac{2}{2n+1} P_n(\alpha_i) P_n(\alpha_j) \tag{12}$$

and elements of D and E are similar. By differentiating the cost function with respect to each of the control forces f_i and setting all the differentials to zero, we can show that the vector of control forces B_{opt} which minimizes the cost function is the solution to the normal equation

$$CB_{opt} = S^T D. \tag{13}$$

In the next section we develop similar equations using a numerical formulation.

2.1.4. Obtaining the normal equations from discrete variables

On the other hand, suppose the $B_i(\theta)$ and $X_i(\theta)$ are transfer functions defined only at discrete farfield angles locations θ_m, m=1,2,..M_q. Any numerical solution produces a discretized farfield pressure distribution, including the analytical solution when computed numerically using a computer program. By summing the contributions made to the cost function at each farfield location θ_m, we essentially perform a numerical integration of Eq. ([7]) to give an estimated radiated power

$$\hat{\Pi} = \frac{\pi R^2}{\rho c} \sum_{i=1}^{M} w_i \left| p(\theta_i) \right|^2 . \tag{14}$$

The constant coefficients w_i, discussed below, account for the area-weighting and sine term in the integration. Equation (14) allows us to evaluate the coefficients of the normal equations and solve for the required control forces. For the numerically computed cost function to closely approximate the radiated power, we must satisfy two conditions. First, there must be enough farfield locations to adequately characterize the response. Second, the weighting function $w(\theta)$ must be discretized to account for the contributions of each location. This is accomplished by defining a diagonal matrix W in which each entry on the diagonal represents a different farfield location. Using the notation of the previous section, the C, D, and E matrices become

$$\begin{aligned} C &= X^T W X^* \\ D &= A^T W X^* \\ E &= A^T W A^* \end{aligned} \tag{15}$$

and the normal equations are essentially identical to Eq. (13).

If the total response is axisymmetric, then the variables contain acoustic pressures at farfield locations θ_m, m=1,2, ... M_q (at constant ϕ). The row vectors A and B become matrices of column dimension M_q. The entries in the diagonal weighting matrix W are weighting coefficients w_i proportional to the areas associated with each of the farfield locations. If the farfield locations are evenly spaced at intervals of $\Delta\theta$, then

$$w_i = \begin{cases} \dfrac{\pi R^2}{\rho c}[1-\cos\Delta\theta/2], & i=1 \text{ or } i=M_\theta \\ \dfrac{\pi R^2}{\rho c}[\cos(\theta_i-\Delta\theta/2)-\cos(\theta_i+\Delta\theta/2)], & 1<i<M_\theta \end{cases} \tag{16}$$

If the total response exhibits a plane of symmetry, then the variables must be partitioned to contain farfield pressures at θ_m, m=1,2,...M_q for *each* of the longitudinal locations ϕ_i, i=1,2,...M_f. Thus the row vectors X and A become matrices whose column dimension is the product of M_q and M_f. The entries in the weighting matrix W must be similarly partitioned, and scaled proportional to the areas associated with each of the farfield locations. If the farfield locations are evenly spaced at intervals of $\Delta\theta$ in the latitudinal direction and $\Delta\phi$ in the longitudinal direction, then the entries in W are the same as in Eq. ( ) but divided by the number of intervals in the ϕ-direction.

2.2. Results and discussion

2.2.1. Analytical and numerical response to disturbance force

Figure 4 compares the uncontrolled radiated power spectra for the analytical model vs. the numerical model. The response contains sharp peaks at $k_o a$=1.14, 1.44, 1.64, and 1.8, the nature of which are detailed below. The two methods agree well in terms of magnitudes and overall trends, as could be expected for this simple structure. At frequencies above $k_o a$>1.7, the peaks in the numerical model are shifted to slightly lower frequencies. Reasons for this shift include both modeling simplifications in analytical solution and discretization approximations in the numerical

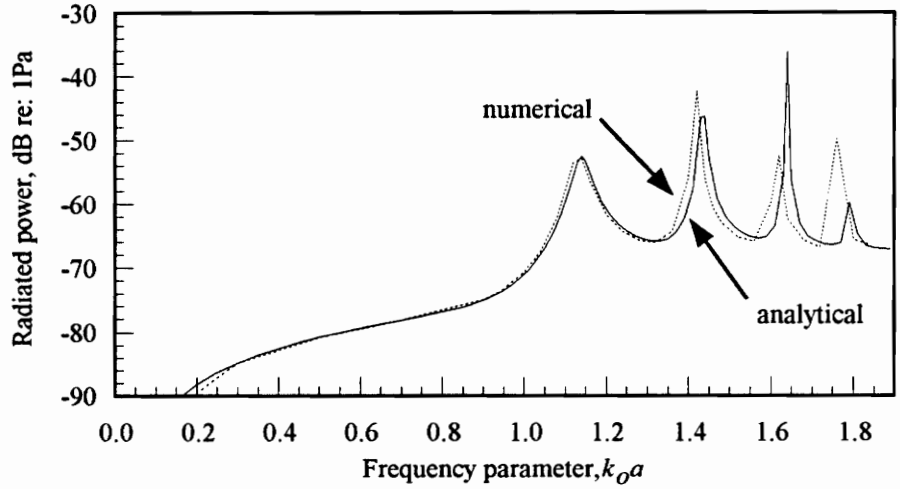

Figure 4: Comparison of radiated power computed with analytical and numerical (NASHUA) solutions. Frequency regime $0<k_o a<1.7$ contains three peaks; agreement is good between analytical and numerical results. Some frequency shifting occurs at higher frequencies.

model. For the purposes of this study, the region below $k_o a=1.7$ contains sufficient features that higher frequencies need not be treated here.

One difference between the analytical and numerical solutions appears in Fig. 5, which shows the distribution of surface velocity at a frequency of $k_o a=1.14$ due to a point-force disturbance

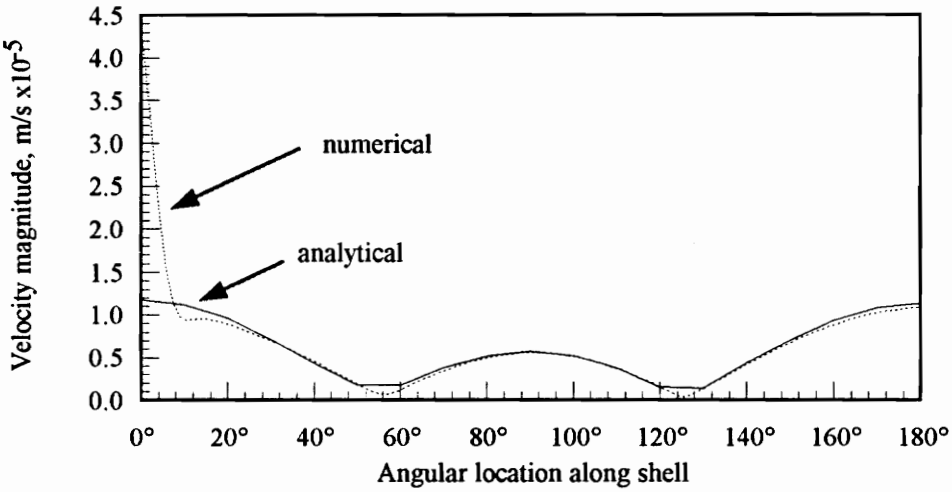

Figure 5: Sample surface velocity distributions from analytical and numerical solutions ($k_o a=1.14$). Large local deformations near drivepoint occur in numerical solution but not analytical solution.

without control. The two solutions agree well except near $\theta=0°$, the location at which the force is applied. The numerical solution, which accounts for bending as well as membrane stresses, displays large local deformations near the drive-point location; the analytical solution, which truncates high-order terms and does not account for bending stresses, does not exhibit local deformations. Such short-wavelength disturbances do not propagate efficiently to the farfield; they can be ignored without consequence if a radiation-based cost function is used. However, when analytical and numerical velocity distributions are compared to check solution accuracy, we must expect such differences to occur near drivepoints.

The solution accuracy also depends on the farfield discretization interval $\Delta\theta$ used to compute radiated power, and we must verify that our choice of $\Delta\theta$ is sufficiently small to characterize the farfield behavior. (Recall that farfield locations are separated by an angle $2\Delta\theta$.) Since the analytical solution can use either continuous or discrete variables, we can compare the radiated power found using continuous variables to that found using discrete variables with different values of $\Delta\theta$. A small value of $\Delta\theta$ should provide the same radiated power found using continuous variables, and the agreement should deteriorate as $\Delta\theta$ increases. Figs. 6a and 6b compare the radiated power spectra for continuous variables to that found using farfield intervals of $\Delta\theta=2.5°$, 5°, and 15°. Figure 6a shows the frequency range $0<k_0a<1.9$, and Fig. 6b shows an enlarged view of the frequency range $1.6<k_0a<1.9$. With $\Delta\theta=2.5°$ and 5°, the agreement with the continuous solution is excellent at all frequencies. Even with $\Delta\theta=15°$ the solution is adequate through $k_0a=1.7$. Results in this thesis use a farfield interval of $\Delta\theta=5°$.

Figure 7 shows the shapes of the first few vibrational modes, denoted by the number n of nodal circles present. The mode shapes shown are proportional to the surface velocity distributions

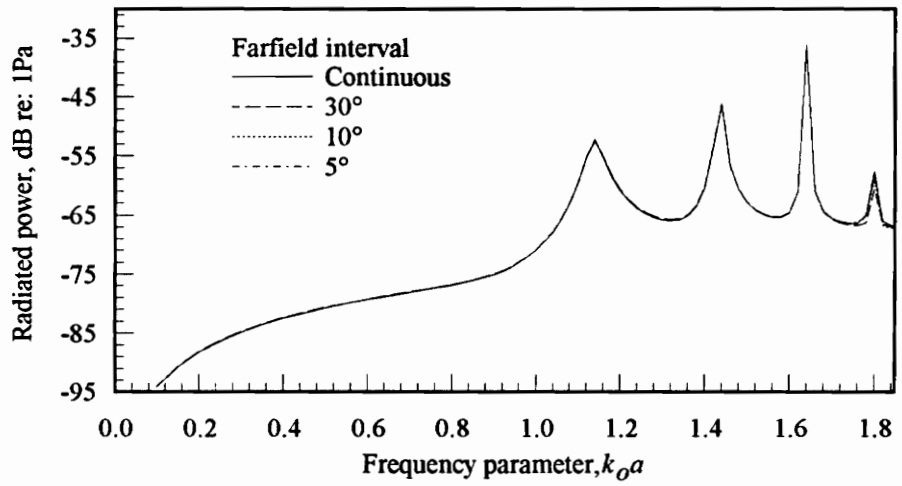

Figure 6a

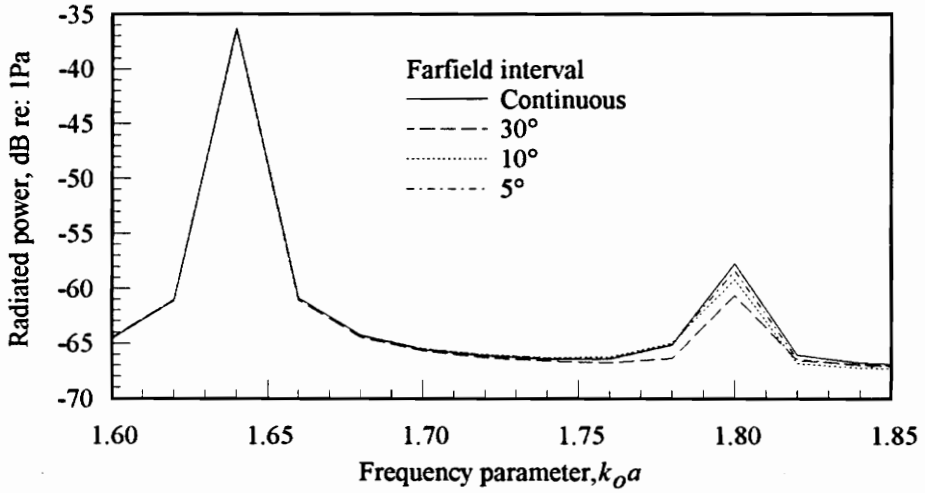

Figure 6b

Figure 6: Effect of farfield discretization on accuracy of radiated power calculations. Solid line represents continuous variables; other lines represent discretized variables with various mesh sizes. Figure 6a shows entire frequency range; Fig. 6b enlarges upper frequencies to show detail. At these low frequencies, 10° farfield interval is sufficient to capture dynamic response.

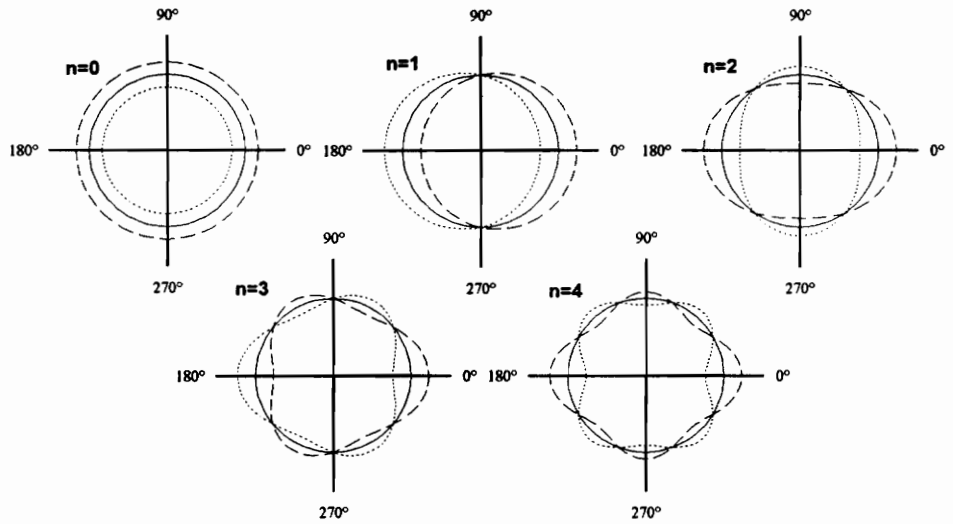

Figure 7: Mode shapes of first few vibrational modes of the spherical shell, denoted by the number n of nodal circles present.

obtained using only the n^{th} term of Eq. (2). The n=0 mode is a breathing mode or uniform radial expansion-contraction. The n=1 mode is a rigid-body translation, the n=2 involves stretching the sphere along the symmetry axis, and so forth. Figure 8 shows the contributions of the individual modes to the radiated power for a single point-force excitation. Since the modes are orthogonal and are not coupled by the presence of the fluid, there are no cross-terms in the radiated power.

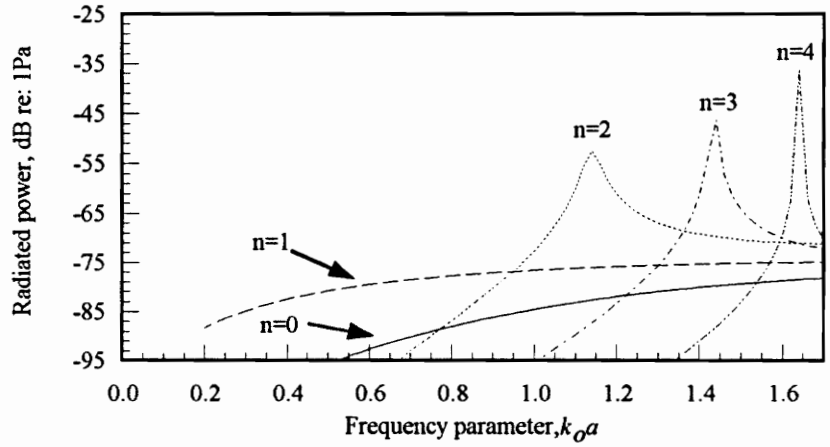

Figure 8: Modal contributions to radiated power spectrum. The n=0 mode does not cause a separate peak in the response because of high radiation damping.

Note that because of high radiation damping the n=0 mode does not display a separate peak in this frequency range, although it contributes to the radiation at all frequencies. The n=1 mode dominates the radiated power at very low frequencies, which can be considered the "resonant frequency" of the rigid-body mode. The n=2, 3 and 4 modes each dominate the response at the frequencies k_0a=1.14, 1.44, and 1.64, respectively. These modal contributions are used below to help explain controller performance in terms of modal spillover.

2.2.2. ASAC using a single point force

To see whether a single control force can produce significant *global* attenuation, i.e., attenuation throughout the entire radiated field rather than in a localized area, we can examine how control affects the total radiated power. Figure 9 shows the radiated power spectra due to a unit-magnitude disturbance at α=0°, before and after adding a single control force at α=180° (numerical solution). The control force reduces the radiated power by up to 20 dB at resonance frequencies and very low frequencies; between resonances the reductions are smaller. At several frequencies there is no attenuation, a phenomenon that is explored later in this section.

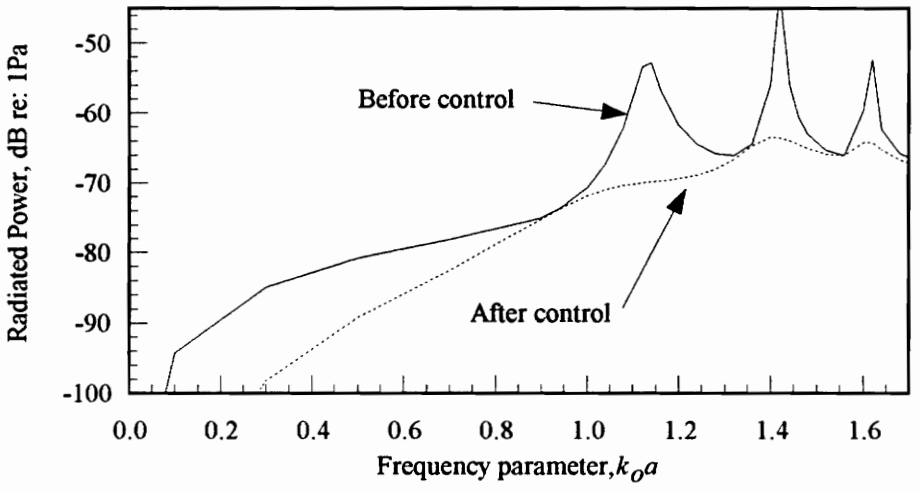

Figure 9: Radiated power spectrum before and after applying a single control force at α=180°. Reductions of up to 20 dB are seen at resonance frequencies and at very low frequencies. Between resonances, little or no attenuation is seen.

Figure 10a compares the analytical and numerical solutions for the residual (closed-loop) radiated power. The agreement between the two solutions is excellent, with discrepancies of less than 0.5 dB at frequencies below k_0a=1.4. At higher frequencies the levels agree quite well, but exhibit a slight frequency shift consistent with the frequency shift seen in Fig. 4. Figure 10b

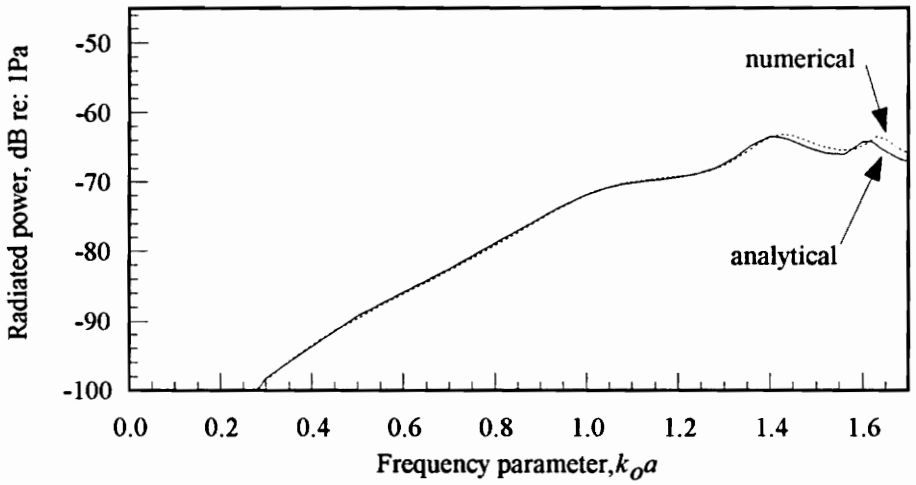

Figure 10a

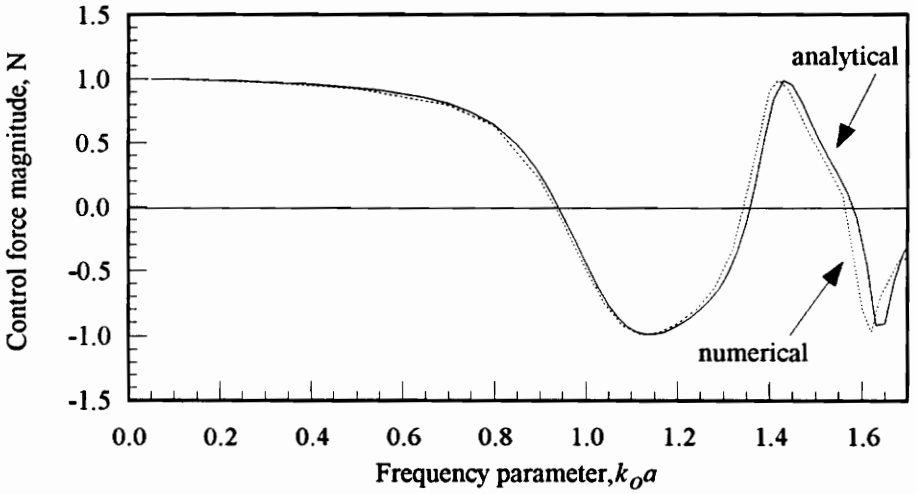

Figure 10b

Figure 10: Comparison of numerical and analytical solutions for 10a) radiated power residual vs. frequency, and 10b) control force magnitude. Levels agree well between numerical and analytical solutions; slight frequency shift is observed at higher frequencies. For disturbance force of unit magnitude, control force varies between +1 and -1 with extrema occurring at resonance frequencies.

compares the analytical and numerical solutions for the control force magnitude as a function of frequency. The magnitude varies between +1 and -1 for a unit disturbance force; the largest control forces are required at resonance frequencies. Between resonances the control force drops to zero at frequencies corresponding to those at which there is no reduction in radiated power. The numerical and analytical results agree very closely in levels, although they exhibit the same frequency shift seen in Fig. 4 and Fig. 10a.

We can better understand system behavior by examining which modes contribute to the radiated power, a convenience possible here but not in the general case of a numerical model. Because of symmetry, the control response and disturbance response contain the same modes except for phase changes. The control and disturbance forces are exactly 180° apart, so the phase differences between modes in the control and disturbance responses are either 0° or 180°. In other words, the control inputs are either in-phase or out-of-phase with the modes in the disturbance response. Each disturbance mode is either eliminated or doubled in strength; cancelling the dominant mode increases contributions from other modes via "modal spillover."

Figure 11 shows the modal contributions to radiated power before and after adding the control force for k_0a=1.14, 0.94 and 1.32. In Fig. 11a the disturbance response is clearly dominated by the n=2 mode. The control force eliminates the contributions of the n=2 mode and, by coincidence, the n=0 mode. The response is reduced globally both in the farfield and on the structure, an behavior sometimes called "modal suppression" because the dominant radiating mode is eliminated [8]. Note that contributions from the n=1 and n=3 modes *increase* by 6 dB. The controller focuses on eliminating the dominant mode; the net effect is a reduction in radiated power, even though some energy "spills over" into less efficient modes that are in-phase with the control force.

In Fig. 11b, the frequency has been chosen so that the n=1 and n=2 modes make equal contributions to the radiated power. Since only one control force is present, it cannot eliminate the radiation from one of the modes without spilling energy into the other mode. Any non-zero control force only increases the total radiated power, so the optimum solution is to set the control force to zero. This appears in Fig. 9 as a point at which no power reduction takes place, and in Fig. 10b as a point at which the control force vanishes.

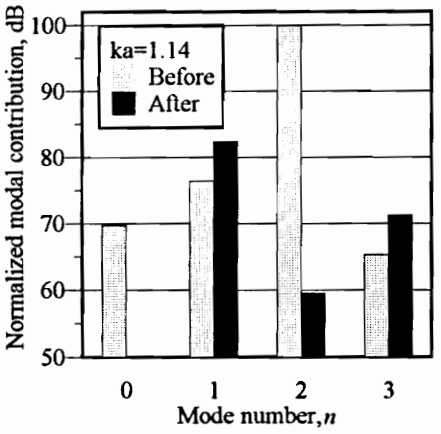

Figure 11a

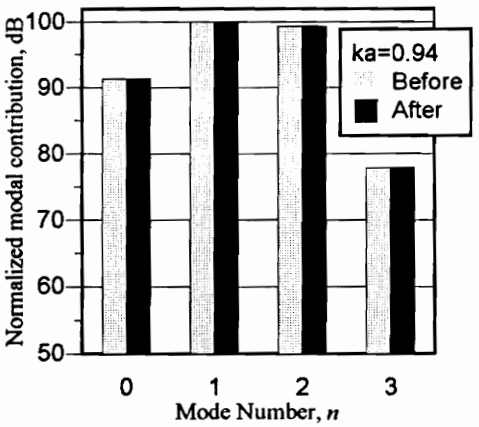

Figure 11b

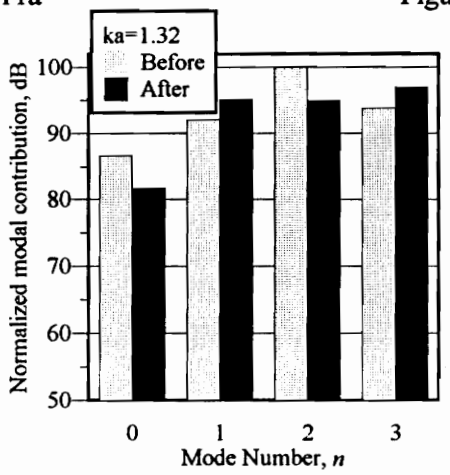

Figure 11c

Figure 11: Modal components of radiated power before and after applying control for k_0a=1.14 (Fig. 11a), k_0a=0.94 (Fig. 11b), and k_0a=1.32 (Fig. 11c). Figure 11a exhibits modal suppression; the n=0 and n=2 modes are reduced by 40 dB or more, but the n=1 and n=3 modes are boosted 6 dB by modal spillover. In Fig. 11b, no reduction is possible with one control force. Figure 11c exhibits modal restructuring; energy is transferred from n=2 mode to n=3 mode.

In Fig. 11c both the n=2 and n=3 modes contribute significantly to radiation, but the largest contribution comes from the n=2 mode. Applying a control force reduces the contribution from the n=2 mode, but not beyond the balancing point at which the control energy "spilled over" into uncontrolled modes would begin to increase the radiated power. The net decrease in radiated power is 3 dB. Interestingly, although the radiated power is reduced, the vibration levels on the shell actually increase slightly. Figure 12 shows the numerically computed surface velocity distribution before and after applying the control force. As in Fig. 5, the solution exhibits relatively large local deformations near the drive-points at 0° and 180°. Elsewhere the surface velocity is actually higher after the control is applied even though the radiated power has decreased. This type of behavior is often termed "modal restructuring" because energy from modes that radiate efficiently is redirected into less efficient modes rather than being suppressed altogether [8].

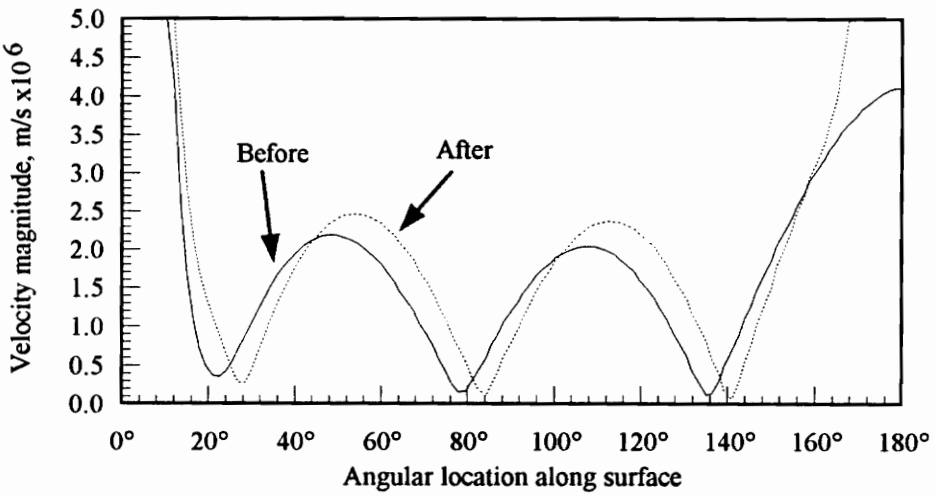

Figure 12: Surface velocity magnitude before and after applying control force, $k_o a$=1.32. Modal restructuring transfers energy from n=2 mode to n=3 mode to decrease radiated power, but vibration levels actually increase.

Numerical accuracy is reflected in part by the agreement between analytical and numerical methods in predicting the *residual response* (closed-loop response). The residual pressures are (hopefully) much smaller than the open-loop pressures, and should therefore provide a more stringent test of numerical accuracy. Figures 13a and 13b show the residual farfield pressure directivity and surface velocity distribution at k_0a=1.14, while Fig's 14a and 14b show similar plots for k_0a=1.32. (The frequencies correspond to those used for Fig's 11a and 11c, respectively.) In each case the numerical solution reproduces the analytical results to a satisfactory degree, with reasonable agreement on the general character of the residual if not on absolute levels. As seen in Fig. 5, the localized drive point deformations that appear in the numerically computed surface velocities do not appear in the analytical solutions. The agreement for farfield pressure residuals is closer because the local surface deformations to not propagate efficiently to the farfield. Overall, the numerical approach appears sufficiently accurate for use in simulating active noise control, at least at the low frequencies examined here.

2.2.3. ASAC using multiple forces

With multiple control forces we face the potential of numerical ill-conditioning known as collinearity, a topic discussed further in Chapter 4. Collinearity generally occurs when too many actuators are used for the frequency being considered, or the actuator positions are poorly chosen. For the present discussion, it will suffice to say that we can circumvent the collinearity problem when necessary by using singular-value decomposition (SVD) to solve the normal equations (see Section 4.1.3). If no collinearity is present, SVD yields the same solution as other solution techniques; in the presence of collinearity, it yields an approximate solution that minimizes the effects of collinearity. The SVD technique is used for all the results discussed below.

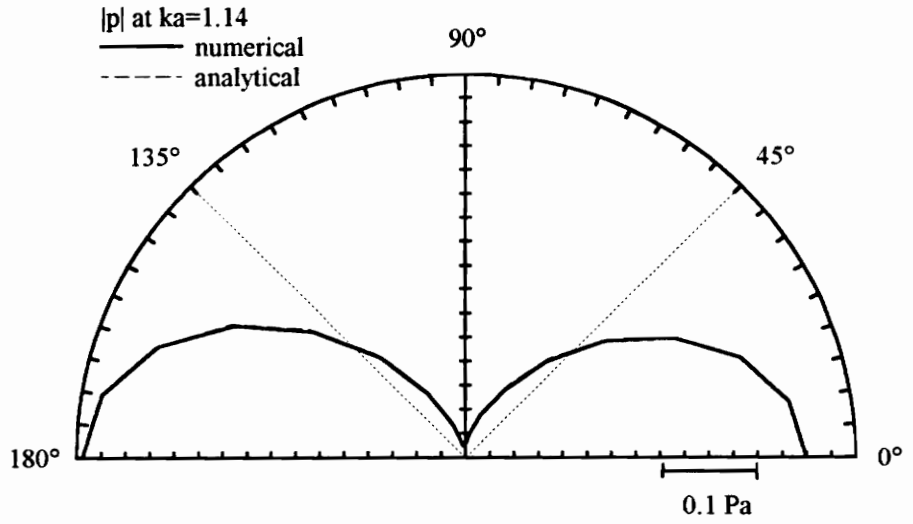

Figure 13a

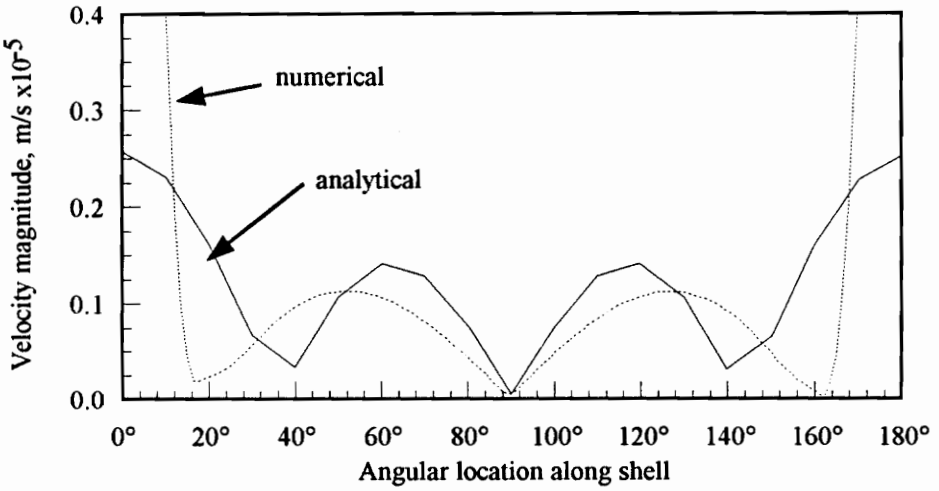

Figure 13b

Figure 13: Comparison of residual response found with analytical vs. numerical methods at $k_o a$=1.14 (first resonance frequency). Solid line represents analytical results; the broken line represents numerical results. Figure 13a shows farfield pressure directivity; Fig. 13b shows surface velocity distribution. Farfield solutions agree extremely well. Local deformations near drive-points in numerical solution do not propagate to farfield, and thus do not affect control behavior.

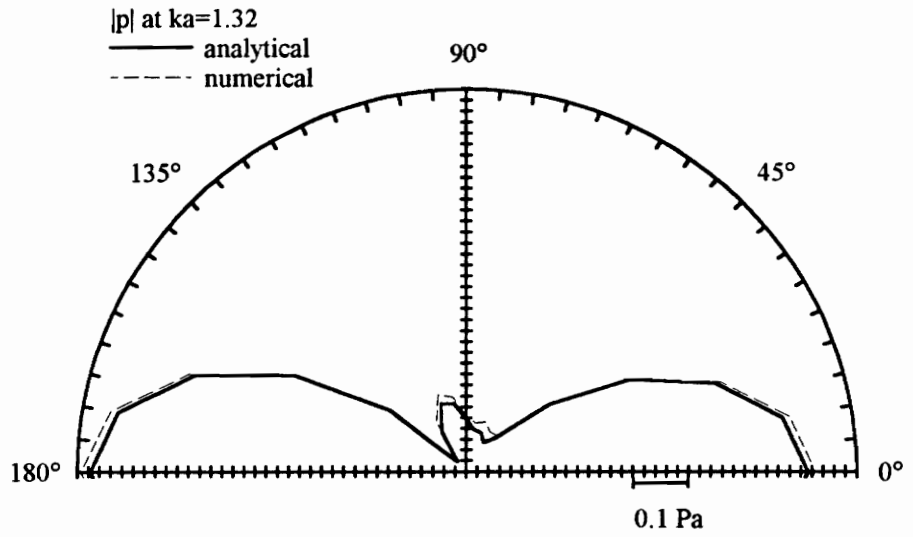

Figure 14a

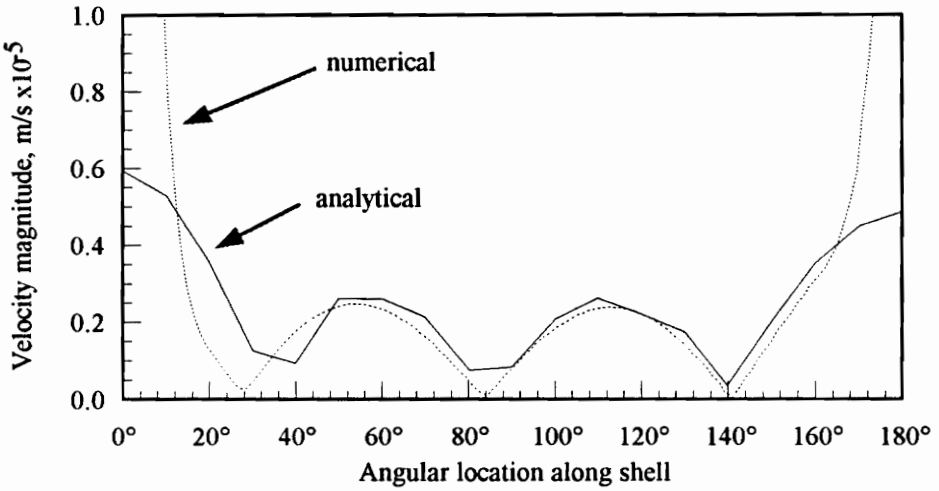

Figure 14b

Figure 14: Comparison of residual response found with analytical vs. numerical methods at $k_o a$=1.32. Solid line represents analytical results; the broken line represents numerical results. Figure 14a shows farfield pressure directivity; Fig. 14b shows surface velocity distribution. Farfield solutions agree extremely well. Local deformations near drive-points in numerical solution do not propagate to farfield, and thus do not affect control behavior.

Figure 15 which shows numerical predictions of the system performance with multiple actuators. The solid line represents the disturbance response; the other lines represent the controlled response with one, two, four, and seven control forces. Each configuration has one control force at α=180°. The configurations with two, four, and seven forces also have control forces evenly spaced at intervals of 90°, 35°, and 25° respectively. Controller performance increases markedly as more control forces are added. This is especially true at off-resonance frequencies where multiple modes contribute to the cost function and thus multiple forces are required to control the radiation. Four control forces can achieve roughly 10 dB attenuation at k_0a=0.94, where a single actuator achieves no attenuation. Increasing the number of control forces to seven gives attenuations of 6 dB or more throughout the entire frequency range.

It should be noted here that no attempt has been made to optimize the locations of the control forces. Similar reductions could possibly be obtained with a smaller number of optimally placed

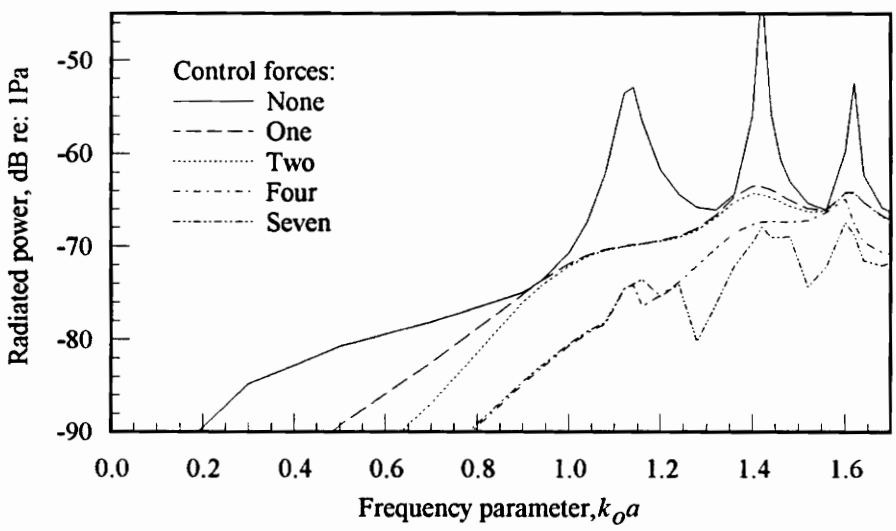

Figure 15: Radiated power spectrum with multiple control forces. Solid line represents response with no control. Other lines represent response with one, two, four, and seven control forces. Increasing number of control forces improves performance, especially between resonances where multiple modes contribute to the response.

control forces. The number of control forces needed could be as low as one control force per contributing mode, if control force locations were optimized so that each force could control a separate mode without spillover into other modes.

2.3. Summary

This chapter presents both an analytical simulation of feedforward control for a spherical shell, and its numerical equivalent. The analytical and numerical models for the spherical shell agree well both in response characterizations and in solutions to the active control problem. Numerical and analytical solutions for the uncontrolled response display excellent agreement below k_0a=1.4, with some slight frequency-shifting in the region $1.4<k_0a<1.7$. Predictions of controller performance, including control force magnitudes and reductions of radiated power, agree well but are subject to the same frequency shift. Short-wavelength structural responses that appear in the numerical solution but not the analytical solution do not affect the farfield behavior, and therefore do not affect the controller performance. Numerical predictions of residual responses agree with analytical predictions to the same extent that uncontrolled predictions agree. In general, the results give confidence that the methods described could be used to examine other, more complicated structures for which no analytical models exist.

For the spherical shell, ASAC is an efficient method for controlling radiated noise at low frequencies ($0<k_0a<1.7$). At resonance frequencies, radiation due to a point-force disturbance can be reduced by up to 20 dB using only one actuator. The mechanism for reducing radiated power on-resonance is modal suppression, in which the dominant response mode in terms of sound radiation is completely suppressed. Between resonances, multiple actuators are required to obtain large reductions. The main mechanism for reducing radiated power off-resonance is modal restructuring, in which modal energy is transferred to less-efficient modes and can actually

increase vibration levels while rediation is suppressed. Increasing the number of control forces improves controller performance. With seven control forces, attenuations of at least 6 dB can be obtained over the entire frequency range $0<k_0a<1.7$.

The results of this section leave unanswered some important questions. We have shown that the numerical and analytical approaches give the same results when we ignore measurement noise in the error sensors. However, in reality each error sensor measurement contains random measurement noise. Furthermore, we have said nothing about the statistical implications of using a finite number of sensors. These topics are addressed in Chapter 3.

CHAPTER 3: USING REGRESSION TO SIMULATE FEEDFORWARD CONTROL

The primary goal of this chapter is to present least-squares multiple linear regression as another way to formulate the numerical approach derived in Chapter 2 for simulating feedforward active control. In doing so, we shall show that regression diagnostics commonly used in statistical regressions can also be applied to the active control problem. We shall discuss why and how regression diagnostics should be used, paying special attention to the physical significance of the quantities involved. Finally, we shall examine some numerical results showing how regression diagnostics can be used in simulating ASAC problems.

The need for regression diagnostics in a numerical simulation of ASAC stems from two facts. Firstly, error sensor outputs in a real control system are contaminated with random measurement noise whose properties we know only in a statistical sense. Secondly, we do not know the error as a continuous function in space; instead, we sample the error at a finite number of discrete error sensor locations, allowing the possibility of undersampling. Because of measurement noise and/or undersampling, real control systems generally do not attain the level of performance predicted by numerical or analytical simulations. To improve simulations and check their reliability, we can use regression diagnostics to characterize the effects of measurement noise and undersampling. Only if the system being modeled has continuously distributed error sensors that provide noise-free measurements (a situation not at all common in practical applications) can we be assured that the simulation will accurately predict the control system response. Otherwise, we must employ some statistical approach such as regression diagnostics to guage how reliable the results are.

This implicit modeling of measurement noise differentiates the statistical methods outlined below from previous works.

By way of introduction, we shall now summarize as briefly as possible the relation between regression and active control. Multiple linear least-squares regression is a mathematical technique used to describe or model a statistical relation between two or more variables. The numerical procedures used to solve a regression problem can also be used to solve for the control inputs in a feedforward, frequency-domain active control problem, provided the techniques are modified to operate on complex-valued variables. Similarly, regression diagnostics can be used to model the effects of measurement noise and/or undersampling. In keeping with tutorial nature of the discussion, this chapter discusses only the most basic regression diagnostics: the F-test, the t-test, confidence intervals, and prediction intervals.

The link between regression and feedforward control was first noted by Snyder and Hansen [21, 22]. Their work describes how to simulate a feedforward control problem using commercial regression software. Since commercial regression packages generally operate only on real-valued data, Snyder and Hansen partitioned their problem into separate real and imaginary parts. In the present work we use a more general formulation that expresses the regresson in the complex domain, removing the need to separate the real and imaginary parts. Formulating the problem with complex variables also allows the use of regression diagnostics, which were not discussed by Snyder and Hansen because the diagnostics are difficult to use when the real and imaginary parts of the data are treated separately. Much of the theoretical development for least-squares regression of complex variables is taken from Miller [28, 29].

Various methods for assessing the quality of a model fit are used in the system identification literature, and there is some overlap with the present work [23, 24]. Despite the overlap,

regression diagnostics are discussed in detail here to provide a unified development, and to emphasize that regression diagnostics are among the many statistical techniques that can potentially be applied to the feedforward active control problem.

We first review the fundamentals of multiple linear regression in Section 3.1, and then examine how a feedforward active control problem can be expressed as a complex-valued multiple linear regression. We examine how regression diagnostics can be interpreted for active control problems, paying particular attention to the relation between statistical quantities and physics of the active control problem. Finally, we discuss the relation between the weighting function and the physical significance of the cost function. Included in Section 3.2 are sample calculations for the cylindrical shell example problem of Fig. 3.

3.1. Regression and feedforward control

Much of the following discussion may be found in textbooks and articles on linear algebra or least-squares techniques; see, for example, [30-32]. In particular, a comprehensive discussion of weighted least-squares regression may be found in [33]. The discussion is included here to introduce terminology which may not be familiar to a general acoustics-oriented or control-oriented audience, and to help illustrate the link between feedforward control and regression. In addition to the discussion here, the reader is urged to review basic statistics in order to ease understanding of the following theory and modeling.

The reader is also reminded here that the following analysis implicitly assumes, as do most statistical analyses, the presence of random measurement noise. In this sense it is unlike most previous works, which take a deterministic view of active control systems.

3.1.1. Fundamentals of multiple linear regression

In regression one seeks a model $\tilde{Y}(\beta_i, X_i)$, the *true model*, to approximate or predict a set of *measured data* Y, also called the *dependent variable*. The true model is a function of β_i (the *true regression coefficients*) and a collection of *predictor variables* X_i (also called *independent variables*). The measured data and predictor variables contain a total of M observations, j=1,2,...M, and are written in vector form as $Y=\{y_1...y_M\}^T$ and $X_i=\{X_{1i}...X_{Mi}\}^T$. For example, a model of human body weight would have the body weights of a group of people as the measured data Y, while the predictor variables X_i would be other measurable characteristics such as height, age, etc.

Suppose now that each observed value y_i contains a small amount of random measurement noise ε_i. Note the use of the term "noise" rather than "error", in keeping with the usage of Nelson and Elliott [15] and Fig. 1. In statistical applications, measurement noise is referred to as "measurement error" or "observation error", whereas in active control literature the term "error" refers to the residual signal that remains after the controller has attempted to cancel the offending signal. The present text uses the terms "error" and "residual" somewhat interchangeably, but the term "noise" is reserved only for the random component of the measured data Y.

In addition to assuming that each y_i contains measurement noise ε_i, let us assume that the true model gives a good approximation of the data:

$$y_i = \tilde{y}_i + \varepsilon_i, \tag{17}$$

where the expected value of ε_i is E[ε_i]=0, the expected value of y_i is E[y_i]=$\tilde{y}_i$. Because we are only *assuming* the validity of Eq. (17), it is important to return to this assumption later and check whether the data cause us to reject the assumption. One such check, the F-test, is discussed in the next section.

To use the regression diagnostics described in this Chapter, we must make certain assumptions about the type of measurement noise we expect to encounter. These assumptions are *not* needed to perform the regression; we can always use regression to find a set of regression coefficients *regardless* of the form of the measurement error. However, sometimes we wish to go beyond merely computing the regression coefficients, estimated model, and residuals. To do this, we must make assumptions about the measurement error, and the assumptions are presented here to introduce nomenclature used later.

The assumptions we must make concern $V(y_i)$, the variance of the i^{th} data measurement. (A brief review of variances and other descriptive statistics is given in Section 3.2.1.) A general form for the variance is

$$V(y_i) = \sigma_Y^2 g^2(\tilde{y}_i, w_i). \tag{18}$$

where σ_Y^2 is a constant measurement noise variance coefficient and $g^2(\tilde{y}_i, w_i)$ is a measurement noise variance function. By including $\tilde{y}_i$ in the variance function $g^2(\tilde{y}_i, w_i)$, we allow the possibility that measurement noise may depend on the magnitude of the quantity being measured. The weighting coefficients w_i allows us to "weight" some observations more heavily than others, for example placing more emphasis on those observations known to be most reliable. Let us assume that the ε_i are normally distributed and pairwise uncorrelated, that is, ε_i and ε_j are uncorrelated when $i \neq j$. (When the ε_i are *not* normally distributed, we must question whether a least-squares approach should be discarded in favor of a more robust approach.) Next we make a simplifying assumption to be scrutinized later: we assume that $V(y_i)$ is independent of $\tilde{y}_i$. This can be written

$$V(y_i) = \frac{\sigma_Y^2}{w_i} \quad \text{or, equivalently,} \quad V(Y) = \sigma_Y^2 W^{-1} \tag{19}$$

where we have defined a *weighting matrix* $W = \text{diag}(w_1, w_2, \ldots w_N)$.

Deciding on a value for σ_Y^2 is an important step. Ideally, the value of σ_Y^2 should come from some *a priori* knowledge of the physical measurement system to be used. For example, we could estimate σ_Y^2 by repeatedly measuring a calibration signal and examining how much variation exists in the measurements. In a preliminary investigation, however, we may not know σ_Y^2. In the absence of a measured value of σ_Y^2, we can use the data to estimate σ_Y^2. In the present work, we use the *estimated residual variance* $s_{Y-\hat{Y}}^2$ (see Eq. 47) as an estimate of σ_Y^2.

Let us now assume a linear regression model of form

$$\tilde{Y} = \sum_{i=1}^{N} \beta_i X_i . \tag{20}$$

Because we do not know the true model $\tilde{Y}$, we cannot directly find the true regression coefficients β_i. But we can use the information contained in the variables X_i to compute an estimated model $\hat{Y}$ with estimated regression coefficients b_i:

$$\hat{Y} = \sum_{i=1}^{N} b_i X_i \tag{21}$$

Our objectives are twofold. First, we must find a set of predictor variables X_i such that the true model $\tilde{Y}$ gives a good approximation of the measured data Y. Second, we must solve for a set of estimated regression coefficients b_i such that the estimated model $\hat{Y}$ closely approximates the true model $\tilde{Y}$. The *F*-test and *t*-test, developed in Section 3.2, allow us to measure how well we have achieved these objectives.

To simplify the notation, we write the estimated regression coefficients as a vector of the form $B=\{b_1 \;\; b_2 \;\; \ldots \;\; b_M\}^T$, define a *design matrix* $X=\{X_1 \;\; X_2 \;\; \ldots \;\; X_M\}$, and express the estimated model as

$$\hat{Y} = XB. \tag{22}$$

Note that the design matrix X generally has more rows than columns, i.e., there are more observations than there are variables.

If our weighting coefficients w_i are independent of i, we would use a standard least-squares regression, seeking regression coefficients that minimize an *unweighted cost function*

$$\chi^2 = \sum_{i=1}^{M} |y_i - \hat{y}_i|^2, \tag{23}$$

where $y_i - \hat{y}_i$ is the *residual*. By requiring the derivatives $\partial\chi^2/\partial b_i$ to vanish, we would obtain the *unweighted normal equation* whose solution is the vector of estimated regression coefficients B.

$$X^T XB = X^T Y. \tag{24}$$

Of interest here, however, is the case of nonuniform w_i. To include weighting we must transform of the original regression into a weighted regression. The transformation consists of replacing Y with $W^{1/2}Y$, and replacing X with $W^{1/2}X$. The *weighted cost function* becomes

$$\chi^2 = \sum_{i=1}^{M} w_i |y_i - \hat{y}_i|^2, \tag{25}$$

and the *weighted normal equation* is

$$X^T WXB = X^T WY. \tag{26}$$

Rather than defining new variables to represent $W^{1/2}Y$ and $W^{1/2}X$, the presence and use of the weighting coefficients will be shown explicitly where needed.

3.1.2. Formulating feedforward active control as a complex-valued regression

First, let us define the *primary response*, P_p, to be the vector of error sensor outputs due to the disturbance input. For example, if the error sensors are hydrophones, the primary response is a

vector containing the complex pressures at each hydrophone due to the disturbance input. Similarly, let us define a *transfer function X* as the vector of error sensor outputs that would result from a single actuator of unit strength acting alone. Finally, let us define the *secondary response*, $\hat{P}$, as the vector of error sensor outputs that would result from all the actuators acting together but without the disturbance. (Thus $\hat{P}$ depends on all the transfer functions as well as all the actuator strengths.)

To formulate the problem as a multiple linear regression, we must have a linear system as defined by three conditions. The response due to any given actuator acting alone must equal its transfer function multiplied by its complex strength. The secondary response must be a linear sum of all the actuator transfer functions multiplied by their respective complex strengths. Finally, the total or residual response must equal the sum of the primary response plus the secondary response.

Now we can list one by one the elements of the feedforward control problem and their counterparts in the regression problem. The transfer functions are analogous to the predictor variables. The control forces are analogous to the estimated regression coefficients b_i. The secondary response $\hat{P}$ is analogous to the estimated model $\hat{Y}$. Lastly, $P = -P_p$ (the *inverse* of the primary response) is analogous to the dependent variable Y. By choosing control forces such that $\hat{P}$ approximates P, we cause the secondary response to "cancel" the primary response, so that the residual response $P + \hat{P} = -P_p + \hat{P}$ is approximately zero. See Table 1 for a list of active control terms and the analogous terms used in statistical regressions.

Recalling that our variables are complex-valued, we rewrite the cost function of Eq. (25) as

$$\chi^2 = \sum_{i=1}^{M} w_i (p_i + \hat{p}_i)^* (p_i + \hat{p}_i). \tag{27}$$

After some algebra [32], we can show that the weighted normal equation takes the form

$$X^H W X B = X^H W P, \tag{28}$$

where the superscript H indicates the Hermitian conjugate. Similarly, we must modify the way in which we characterize the error sensor measurement noise. If $\tilde{P}$ is the true model approximated by $\hat{P}$, we now assume that

$$p_i = \tilde{p}_i + \varepsilon, \quad E(p_i) = \tilde{p}_i, \quad \text{and} \quad V(|\varepsilon_i|) = \sigma_P^2. \tag{29}$$

Note that Eq. (29) only characterizes the variance of the noise magnitude $|\varepsilon|$, not the phase. Modeling phase errors explicitly would be desirable, but the method for doing so is not clear from existing statistical regression techniques. See Section 3.1.4 for further discussion.

3.1.3. Physical significance of the cost function

In many cases, the cost function being minimized by the controller can be an approximation of some physical quantity. The physical significance of the cost function depends on how we define

Table 1: Comparison of statistics terminology to active control terminology.

Statistical regression terminology	Active control terminology
dependent variable; data	inverse of primary response; inverse of open-loop response
model; estimated model	secondary response; control response
independent variable; predictor variable	transfer function
regression coefficient	control force; control input
residual	error; closed-loop response; residual response
observation or measurement	error sensor
observation or measurement error	error sensor measurement noise

two quantities: the predictor variables (i.e. the sensors) and the weighting coefficients. In the present example the sensors are farfield pressure sensors and the predictor variables are vectors of pressure per unit force of length M, where M is the number of sensors. The N predictor variables are nothing more than transfer functions between unit forces on the structure and the resulting farfield pressures $P_j(\theta_i)$, i=1, 2, ... M, j=1, 2, ... N. Consider the following expressions for the radiated power [27]:

$$\Pi = \oint_S \frac{\pi r^2}{\rho c}\left|P+\hat{P}\right|^2 dS \quad \cong \quad \sum_{i=1}^{M} \frac{\pi r^2 a_i}{\rho c}\left|p_i+\hat{p}_i\right|^2, \tag{30}$$

where S' is a spherical surface of farfield radius r with the origin at its center, and a_i is the area on the surface S' associated with the i^{th} farfield location. If $a_{\max}$ is the largest value of a_i, we can rewrite the estimated radiated power as

$$\hat{\Pi} = \frac{\pi r^2 a_{\max}}{\rho c} \sum_{i=1}^{M} \frac{a_i}{a_{\max}}\left|p_i+\hat{p}_i\right|^2 \tag{31}$$

By comparison with Eq. (25), the cost function will be proportional to the radiated power if we specify weighting coefficients of the following form:

$$w_i = \frac{a_i}{a_{\max}} \tag{32}$$

If there are enough error sensors to completely characterize the radiated field, then the cost function χ^2 is proportional to the radiated power Π; minimizing one minimizes the other. To obtain the radiated power from the cost function, we must use

$$\hat{\Pi} = \frac{\pi r^2 a_{\max}}{\rho c}\chi^2 \tag{33}$$

For the remainder of this text, all references to the total radiated power actually refer to the estimated radiated power in Eq. (33).

Note that we have not put forth any restrictions regarding the physical configuration of the control system. We have chosen an axisymmetric example problem for simplicity of notation, since farfield locations may be specified by only one angle rather than two. However, the method is not limited to axisymmetric cases. The transfer functions can contain pressures from any number of sensors in any physical locations desired. Assuming the number of sensors is sufficient, and the sensor locations and weighting coefficients are appropriate, the cost function can approximate the radiated power.

Another important point is that *any* combination of actuators and sensors can be used. Physical interpretation of the cost function, however, will only be possible with certain combinations of actuators, sensors, and weighting coefficients. Also, sensors need not be distributed throughout the entire radiated field; they could be concentrated in one area to reduce radiation in a certain direction. The reduction of radiated power is a specific case of a more general cost function. Several examples of different cost functions are given in Chapter 6, along with example numerical results.

3.1.4. Assumptions regarding variance of measurement noise

In Section 3.1.1 we made some simplifying assumptions about the variance of the measurement noise in a general regression problem; see Eq. (19). In this Section we examine Eq. (19) in the context of the ASAC problem to determine whether it represents the type of sensor measurement noise we would expect to find in a real control system. Recall that the regression itself does not depend on the distribution of measurement noise, which plays a part only in the regression diagnostics. As discussed below, Eq. (19) may not be completely appropriate for some types of sensors likely to be used in active noise control systems.

Suppose our error sensors will be hydrophones. Suppose we repeatedly measure a calibration signal P_{true} and find that the measurements, when expressed in decibels, exhibit Gaussian measurement noise with a sample standard deviation of, say, $sp = 0.1$ dB. For real hydrophones, this form of approximation may be reasonable. But what if the measurement noise is Gaussian only when expressed in decibels, while we have defined the regression in linear units? Consider Fig. 16, which shows a hypothetical set of hydrophone measurements of a "known" pressure P_{true} over a range of pressure levels. The measurement noise is the difference in magnitudes between P_{true} and the measured value. When the measurement noise is expressed decibels, as in the left plot, the noise variance appears to be constant with respect to the magnitude of P_{true}. But when the same data are plotted in pascals instead of decibels, it is obvious that the noise is much larger at high pressures than at low pressures. In other words, the measurement noise has a non-constant variance. To describe the variance properly we must use Eq. (18) rather than using the simplified Eq. (19).

Because the primary response varies in space, our supposed hydrophone error sensors will measure different pressure magnitudes depending on their locations. A sensor measuring a large

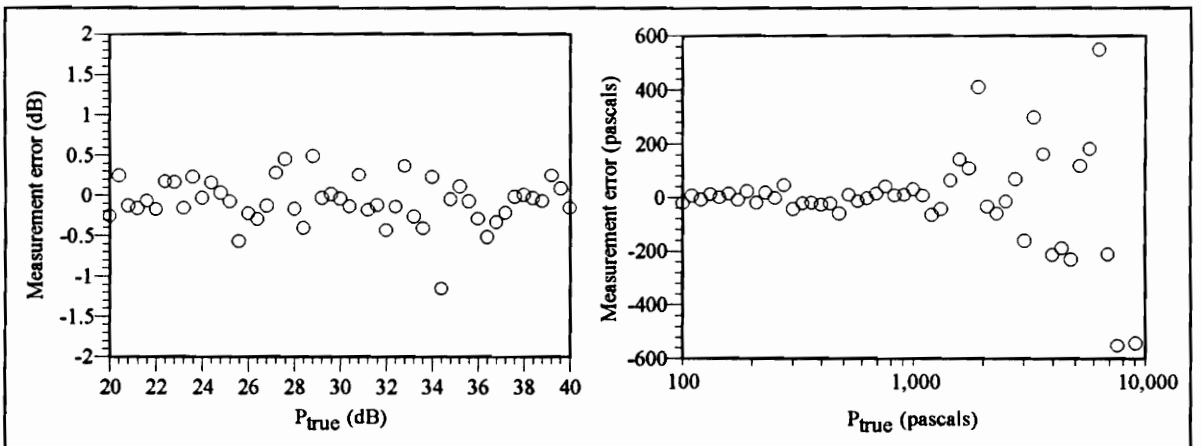

Figure 16: Example of measurement noise with non-constant variance. When expressed in decibels, left, the noise variance appears constant with respect to signal amplitude. When the same data are converted to linear units, right, the variance is clearly a function of signal magnitude.

pressure will thus have a larger noise component (in linear units of pressure) than a sensor measuring a small pressure. This highlights an important point: the assumptions used to derive Eq. (19) might not provide an accurate representation of the noise present with real sensors. Perhaps the variance should instead be proportional to the signal magnitude as in Eq. (18). For simplicity in the remainder of this thesis, *we shall assume Gaussian measurement noise with a constant variance as in Eq.* (19). The point of the above discussion is that one must carefully evaluate the mathematical form of measurement noise based on the specific sensors being used, and then formulate the regression diagnostics appropriately.

More sophisticated treatments of variance are available, but they are beyond the scope of the present analysis. Most involve performing a transformation and then using of some form of generalized least squares; see [33]. For example, it may be necessary to convert the transfer functions and error sensor measurments to decibels before using them in a regression. In many cases it will likely be seen that the use of standard constant-variance techniques is well-justified, and that more complicated techniques would be of limited benefit.

3.1.5 The squared multiple correlation coefficient R^2

One convenient statistic, the squared multiple correlation coefficient R^2, is commonly used to measure the quality of a regression but also has an interesting relation to a quantity used in predicting the effectiveness of active control systems. Ross [34] showed that for a SISO control system, the controller reduces the open-loop response by an amout expressed in decibels as

$$\Delta_{dB} = -10\log\left(1-\gamma^2\right) \tag{34}$$

where γ^2 is the *ordinary coherence function* between the signals of the detection sensor and the error sensor or, in the case of a MIMO system, γ^2 is the *multiple coherence function* (see the discussion of coherence functions in [15].) The ordinary coherence function is easily measured in

practice, and provides a best-case estimate of the effectiveness of a SISO active control system. The squared multiple correlation coefficient $\mathbf{R}^2$ (also known as the squared coefficient of determination) provides an analogous relation in a multiple linear regression. In both cases we are essentially quantifying the fit between the disturbance response and the control response. The main difference is that γ^2 is defined in the time domain, and $\mathbf{R}^2$ is defined in the frequency domain.

For a regression that does not include a constant term as one of the predictor variables, $\mathbf{R}^2$ may be written

$$\mathbf{R}^2 = \frac{\sum_{j=1}^{M} |\hat{y}_j|^2}{\sum_{j=1}^{M} |y_j|^2} = 1 - \frac{\sum_{j=1}^{M} |y_j - \hat{y}_j|^2}{\sum_{j=1}^{M} |y_j|^2} \tag{35}$$

as described in, for example, Appendix II of [23]. A value of $\mathbf{R}^2$ near unity indicates that the estimated model closely fits the measured data. Since M is the number of sensors in the control system, y is the open-loop response, and $\hat{y}$-y is the residual or closed-loop response, we see that

$$\frac{\sum_{j=1}^{M} |y_j - \hat{y}_j|^2}{\sum_{j=1}^{M} |y_j|^2} \tag{36}$$

is exactly the ratio we seek to quantify the controller performance. Thus it is straightforward to show that $\mathbf{R}^2$ and γ^2 have the same form, and the controller reduces the open-loop response by an amount expressed in decibels as

$$\Delta_{dB} = -10 \log\left(1 - \mathbf{R}^2\right) \tag{37}$$

The reader should bear in mind that because $\mathbf{R}^2$ is based on the measured data, which includes measurement noise, it is a best-case measure of performance. Measurement noise will cause an actual control system to perform worse than would be indicated by $\mathbf{R}^2$.

3.2. The role of regression diagnostics

The previous sections have described how regression can be used to predict the performance of an ASAC system, and discussed some basic statistical concepts that must be considered. Measurement noise in the error sensors has been discussed, but until now the measurement noise has not been explicitly considered in the calculations. If the predicted attenuation and the residual field are all that is required, there is no need to explicitly model the measurement noise. However, to improve system design we may need a more detailed analysis including modeling the impact of measurement noise on the controller performance. For this we must assume a value for the variance of the measurement noise, and use one or more regression diagnostics based on that value.

This section discusses the most basic regression diagnostics (the *F*-test, the *t*-test, confidence intervals, and prediction intervals) by defining them and relating them to the feedforward active control problem. To introduce notation, we first define elementary descriptive statistics in a manner that allows for either real-valued distributions or complex-valued distributions.

3.2.1. Descriptive statistics: the mean, variance, and standard deviation

For a given data sequence y_i , i=1, 2, ... M, we can describe the basic features of the data using two familiar statistics: the *mean*, which provides a measure of central tendency, and the *variance*, which provides a measure of dispersion about the mean value. The *true mean* μ_Y is the value that satisfies

$$\sum_{i=1}^{M}(y_i - \mu_Y) = 0, \tag{38}$$

where $(y_i - \mu_Y)$ is called the *deviation* from the true mean. The variance $V(Y)$ is the sum of the squared deviations divided by the number of values:

$$V(Y) = \frac{1}{M}\sum_{i=1}^{M}\left|y_i - \mu_y\right|^2. \tag{39}$$

Note that we have defined the variance using the squared magnitude of the deviation rather than the squared deviation. This distinction will allow us to treat either real or complex-valued distributions in the same fashion; the variance measures the dispersion of the magnitudes, not of the real and imaginary parts.

To find the true mean we must measure all possible values of y. In practice we can only measure a representative sample of the possible values of y, so we must estimate the mean and variance based on the sample. The *sample mean* $\overline{Y}$ is

$$\overline{Y} = \frac{1}{M}\sum_{i=1}^{M} y_i. \tag{40}$$

To estimate the variance of y we use the familiar *standard deviation* s_Y, where

$$s_Y^2 = \frac{1}{M-1}\sum_{i=1}^{M}\left|y_i - \overline{y}\right|^2 \tag{41}$$

Many other descriptive statistics are available, but the sample mean and standard deviation are reviewed here because they are used in developing the regression diagnostics of the following sections.

3.2.2. Testing the regression for lack of fit: The *F*-test

The first test we can perform is the so-called "*F*-test", which measures our confidence in the regression as a whole. The *t*-test and other regression diagnostic to follow are only valid if the estimated model gives an accurate representation of the true model; however, we have no way of verifying this directly because we do not know the true model. Fortunately, we need not prove that the estimated model *equals* the true model to make practical use of the regression diagnostics. If we can show by statistical inference that the variance of the estimated model is not drastically different from the variance of the true model, then we can proceed with a reasonable amout of confidence. The *F*-test allows us to test whether this condition is true.

We begin by assuming that the estimated model $\hat{P}$ is very close to the true model $\tilde{P}$; call this assumption $H_0: \hat{P} \approx \tilde{P}$. One way of verifying $H_0: \hat{P} \approx \tilde{P}$ is to show that the variance of $\hat{P}$ approximately equals the variance of $\tilde{P}$:

$$V(\tilde{P} - \overline{P}) \approx V(\hat{P} - \overline{P}). \tag{42}$$

where $\overline{P}$ is the estimated mean of the disturbance response. Since we cannot know $\tilde{P}$ directly, we make use of a relation from regression theory [31] which states that

$$V(\tilde{P} - \overline{P}) = V(P - \hat{P}) + V(\hat{P} - \overline{P}) \tag{43}$$

Loosely speaking, Eq. (43) states that the variance of the true model about $\overline{P}$ equals the residual variance plus the variance due to the regression. We can rearrange Eq. (43) as

$$V(\tilde{P} - \overline{P}) = V(\hat{P} - \overline{P})\left\{1 + \frac{1}{F}\right\} \quad \text{where} \quad F = \frac{V(\hat{P} - \overline{P})}{V(P - \hat{P})} \tag{44}$$

and F is the so-called *variance ratio*. Therefore, to satisfy Eq. (42) as closely as possible we must have $F \gg 1$, in which case we do not reject $H_0: \hat{P} \approx \tilde{P}$ and thus the *F*-test is not violated.

To determine whether F is large enough, we compare F to a predetermined value. Let $Q(F, \nu_1, \nu_2)$ be F-distribution probability function with ν_1 degrees of freedom in the numerator and ν_2 degrees of freedom in the numerator. Then, given X, Y, M, and N for the specific problem at hand, $Q(F, \nu_1, \nu_2)$ is the probability that F could be as large as it is *by chance alone* rather than because of the regression. When $Q(F, \nu_1, \nu_2)$ exceeds some specified risk level, we must reject $H_0: \hat{P} \approx \tilde{P}$. In the present context with N actuators and M sensors, we have

$$\nu_1 = N \quad \text{and} \quad \nu_2 = M - N - 1. \tag{45}$$

We can obtain values for $Q(F, \nu_1, \nu_2)$ from tables included in most statistics texts, or we may easily generate them on a computer. We do not know the variances required to compute F exactly, but we can approximate F by computing the ratio of standard deviations

$$F = \frac{s^2_{\hat{P}-\overline{P}}}{s^2_{P-\hat{P}}} \tag{46}$$

where

$$s^2_{\hat{P}-\overline{P}} = \frac{1}{N}\sum_{i=1}^{M} w_i \left|\hat{p}_i - \overline{p}\right|^2 \quad \text{and} \quad s^2_{P-\hat{P}} = \frac{1}{M-N-1}\sum_{i=1}^{M} w_i \left|p - \hat{p}_i\right|^2 . \tag{47}$$

If F is large enough, that is, if $Q(F, \nu_1, \nu_2)$ does not exceed our specified acceptable risk level, then we need not reject $H_0: \hat{P} \approx \tilde{P}$ and we can proceed with testing the individual regression coefficients as described in the next section. If F is too small, then it is too likely that we could have chosen regression coefficients at random and still achieved the same fit to the data.

When a regression fails the F-test, it generally means this: given X, Y, M, and N, there is too much variance (scatter) in the measured data Y for us to confidently say that we have modeled the data. In terms of an ASAC problem, it means that one or more of the following conditions exists:

The data and predictor variables are undersampled, i.e., there are not enough sensors to confidently say that the data has been modeled. The solution is to add more sensors; this will provide better estimates of the variances, and may increase F unless other problems are present.

The dependent variable contains components that are not present in the predictor variables, i.e., there are one or more significant modes in the disturbance response that are not coupled into by the actuators. The solution is to add more actuators or use different actuators that couple into different modes.

There is too much measurement noise in the measured data. In other words, even if the actuators couple into the same modes that appear in the disturbance response, severe measurement noise can have the same effect as modes that are not coupled into by the actuators. This seems rather extreme and probably would not occur in real circumstances, since usually great pains will have been taken to improve signal-to-noise ratios.

Generally speaking, a regression will only fail the F-test when the estimated model is very different from the shape of the data. When the F-test fails in an ASAC problem, very little cancellation takes place. Any cancellation that does occur is purely coincidental and occurs only in localized areas rather than globally throughout the domain. (After all, if the shape of the estimated model was close to that of the data, and cancellation occurred everywhere, then the F-test would not have failed in the first place.) The relation between F-test failure and control system performance is explored in Section 3.3.

3.2.3. The t-test and confidence intervals for individual regression coefficients

The t-test, which examines the individual regression coefficients corresponding to the control inputs in the active control problem, allows us to detect whether any actuators are completely

uncoupled from the cost function. If a given actuator has no beneficial effect on the regression, then its regression coefficient vanishes and we could obtain a more reliable regression by reformulating the problem without that actuator. However, determining whether a given coefficient is zero can be difficult because problem scaling may yield arbitrarily small coefficients. To decide whether a coefficient should be removed, we can define a statistical "confidence interval" for the magnitude of that coefficient, and then check whether the confidence interval contains zero.

Because we have assumed Gaussian measurement noise, we can calculate variances and confidence intervals for each of the regression coefficients. The variance of the magnitude of the i^{th} coefficient is

$$\sigma_{b_i} = \sqrt{e_i} \tag{48}$$

where e_i is the i^{th} diagonal of the *variance-covariance matrix*

$$V(b) = \sigma_p^2 \left(X^H W X \right)^{-1} \tag{49}$$

and σ_p is the variance of the error sensor measurement noise. Since the true error variance σ_p is generally unknown, we can use the estimated error variance s_p in its place to obtain an estimated variance s_{b_i}.

If an estimated variance is available for a specific type of error sensor likely to be used by the control system, that variance may be used in performing a simulation. If no such estimate is available, the residual variance $s_{\hat{P}-\overline{P}}$ may be used as an estimate of the error variance. Because the present work emphasizes numerical methods rather than results for specific types of hardware, *all numerical results in this thesis use error variances that are estimated from the data.* The reader is reminded that a specific value of the error variance is required *only* for the *t*-test,

confidence intervals, and prediction intervals; it is not used by the actual regression, which does not consider measurement noise directly.

Next we decide on a level of significance α_t, say 90%, and compute a quantity known as *Student's t*. Student's t is the value of the Student's probability distribution for the chosen level of significance α_t and ν degrees of freedom, where ν=*M*-*N*-1. We can then state that *given the data*, the probability of $|b_i - \beta_i| < ts_{b_i}$ is α_t. In other words, there is a 90% probability that the magnitude of the error between the true control input and the control input computed by the regression is less than ts_{b_i}. This gives us circular confidence region on b_i in the real-imaginary plane (see Fig. 17) such that

$$|b_i| - ts_{b_i} \le |\beta_i| \le |b_i| + ts_{b_i} \tag{50}$$

The largest magnitude inside the confidence region, which represents the largest likely control input magnitude, is $|b_i| + t\, s_{b_i}$.

As noted in [24], we could also consider the real and imaginary parts of the problem separately to produce an elliptical confidence region. This would essentially preserve phase information in

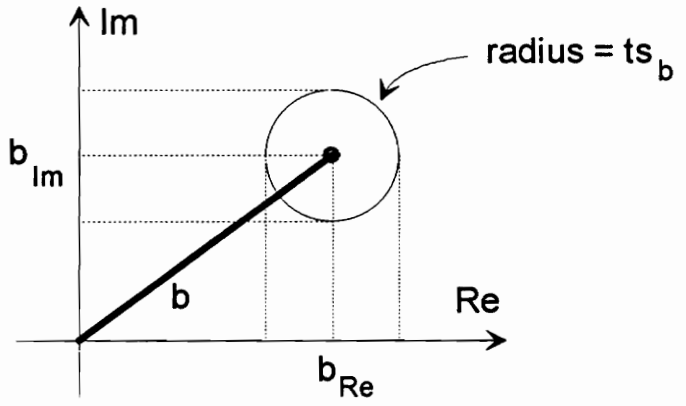

Figure 17: Circular confidence limit on a regression parameter, where Re{b}=b_{Re} and Im{b}=b_{Im}. The confidence limit specifies a circular region of radius ts_b centered at (b_{Re}, b_{Im}) such that the true parameter lies within the circle.

the confidence region rather than only magnitude, but for the present work we shall consider only circular confidence regions.

We can also perform a related statistical test, the so-called "*t*-test". The *t*-test measures whether a given regression coefficient is *significantly* different from zero. Essentially, if the range $\{|b_i| - ts_{b_i}, |b_i| + ts_{b_i}\}$ contains zero at the specified level of significance, then we cannot conclude that b_i is significantly different from zero. Said another way, if a coefficient fails the "*t*-test", we have no statistically significant indication that the coefficient has nonzero magnitude. Note that increasing the number of sensors for a given number of actuators has the effect of decreasing *t*, and thus producing tighter confidence intervals and a more reliable model.

In terms of the ASAC problem, a control input will only be set equal to zero if the actuator is completely uncoupled from the disturbance response. Therefore, any actuator that fails the *t*-test should be removed from the problem.

3.2.4. Prediction intervals

Taking the results of the previous section one step further, we can condense the confidence intervals for all the actuators into a single quantity: the prediction interval, which provides a confidence interval on the estimated model $\hat{P}$. The regression provides an estimated residual $P + \hat{P}$ corresponding to the pressure when the control actuators are turned on. The estimated model $\hat{P}$ is based on error sensor outputs that contain measurement noise, and its reliability is therefore related to the accuracy of the error sensors. By applying the appropriate regression diagnostics, we can estimate confidence intervals on $\hat{P}$ and provide a single quantity that reflects the sensitivity to measurement noise for the regression as a whole.

Note that the prediction interval only reveals how close the estimated model is to the true model. Therefore, the prediction interval only examines sensitivity to measurement noise and, specifically, it does not allow us to compute confidence intervals on the residual or the coefficient of determination $\mathbf{R}^2$, which measures overall controller performance. Because the measurement noise levels at all sensor locations are uncorrelated from one another, we cannot directly calculate their cumulative effect on $\mathbf{R}^2$. However, confidence intervals for $\mathbf{R}^2$ may be estimated using bootstrapping with Monte Carlo methods [35,36,39]. The idea behind bootstrapping is that if we know the probability distribution of the measurement noise, and hence the control inputs and estimated model, we can find the probability distribution of the residual cost function. The technique uses Monte Carlo methods to simulate a statistical universe of "bootstrap" datasets, recalculate the regression for each dataset, and observe how little or how much the residual cost function varies. The technique is widely used, but is beyond the scope of the present analysis.

At a given sensor location, the estimated model is a linear sum

$$\hat{P}_j = \sum_{i=1}^{n} b_i X_{ij} \tag{51}$$

where the control inputs b_i have variances as expressed as in the previous section, and the transfer functions X_i are considered constant and completely accurate for the present purposes. If we define the vector $X_{i,1:N}$ as the i^{th} row of the design matrix X,

$$X_{i,1:N} = \begin{Bmatrix} X_{i,1} & X_{i,2} & \cdots & X_{i,N} \end{Bmatrix} \tag{52}$$

then it can be shown that the variance of the estimated model at the location of the j^{th} sensor is

$$V\left(\hat{P}_j\right) = \sigma_P^2 X_{i,1:N}\left(X^H W X\right)^{-1} X_{i,1:N}^H \tag{53}$$

This is also equal to the j^{th} diagonal of the variance-covariance matrix

$$V(\hat{P}) = \sigma_P^2 X\left(X^H W X\right)^{-1} X^H \tag{54}$$

As in the previous section, we do not know the error variance σ_p unless we have decided upon using a specific type of error sensor with known characteristics. In the present work, we use the estimated residual variance $s_{\hat{P}-\bar{P}}$ as a convenient estimate of the error variance σ_p. Then, using Student's t for a chosen level of significance, the prediction interval for the estimated model at the jth sensor location is defined by the standard error

$$\pm t\sigma_{\bar{P}+P}\sqrt{V(\hat{P}_j)} \tag{55}$$

As with the confidence intervals of the previous section, the prediction interval describes only the magnitude of the estimated model, not the phase.

The prediction interval summarizes the sensitivity to measurement noise for the entire regression, whereas the confidence intervals address individual actuators. Prediction intervals can be used to determine whether a given combination of structure, actuators, and sensors is particularly sensitive to noise and, if so, what portions of the radiated field are most likely to be degraded by noise. As with confidence intervals, increasing the number of sensors for a given number of actuators has the effect of decreasing the width of the prediction interval.

3.3. Results and discussion

Using the analysis of Sections 3.1 and 3.2, we can perform a basic regression analysis for the example clamped cylinder problem of Fig. 3. The cylinder is disturbed by a unit-magnitude axisymmetric ring-force b_d placed near the right-hand end closure at δ=0.96. The two actuators (M=2) are axisymmetric ring-forces b_1 and b_2 placed at δ=0.25 and δ=0.50. The 25 error sensors (N=25) are pressure sensors located at 7.5° intervals from 0° to 180° in the farfield. The

objective of the control will be to minimize the radiated power cost function Π, where Π is the weighted sum of squared pressure magnitudes at the 25 error sensors. (A separate set of calculations, not included here, confirmed that 25 sensors are sufficient to estimate the total radiated power for the frequencies considered.)

Figure 18 shows the radiated power cost function Π vs. nondimensional frequency $k_o a$ for the frequency range $0.2<k_o a<1.0$. The solid line represents the open-loop case. The broken line represents the closed-loop case, and the difference between the two is the reduction in the radiated power cost function. The actuator magnitudes and phases have been optimized separately at each frequency in the region, since each regression uses data for a single frequency. The actuators reduce the radiated power cost function by as much as 14 dB depending on the frequency. Detailed analysis of the structural and acoustic responses before and after applying control is postponed until Chapter 6. Here, where the intent is to illustrate the use of regression diagnostics, we restrict our attention to three specific frequencies as described in the following paragraphs.

The first frequency we shall examine is $k_o a=0.40$. The regression predicts that the radiated power cost function for this frequency will be reduced by 5 dB, and the control inputs are

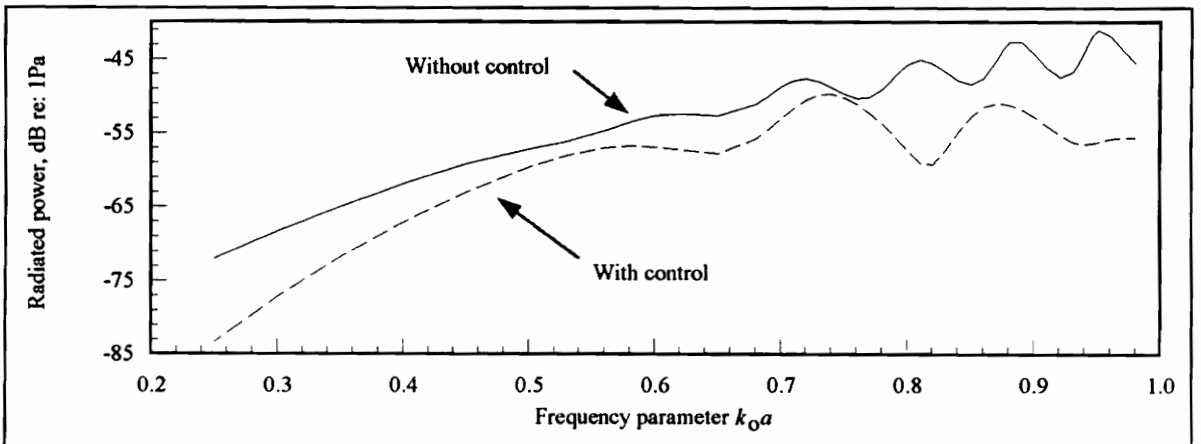

Figure 18: Radiated power vs. frequency for clamped cylindrical shell. Radiated power without control actuators shown as solid line; with control actuators shown as broken line

b_1= 1.06 N and b_2= -1.08 N. These results may are obtained directly from the regression, and do not make use of any assumptions about measurement noise. Now, we invoke the assumptions described in Section 3.1 in order to employ the F-test (to determine whether the regression is valid), the t-test (to determine whether any control forces do not contribute to the fit), and the prediction interval (to examine sensitivity to measurement noise). For the F-test, we compute $Q(F, \nu_1, \nu_2)$=1.6x10^{-6}, which means: *given the data*, there is only a 0.00016% probability that our estimated model is significantly different from the true model. We therefore feel confident enough about the regression to continue with further tests.

To use the t-test and other diagnostics, we must estimate the error sensor noise variance. As in previous sections, we shall use the residual variance rather than some measured value for a specific sensor. Computing confidence intervals at the 90% significance level, we see that the control inputs at this frequency are b_1= 1.06±0.27 N and b_2= -1.08±0.28 N. Since neither confidence interval contains zero, we can say with 90% confidence that, given the data and our estimate of the error variance, both forces are significantly different from zero. Figure 19 shows

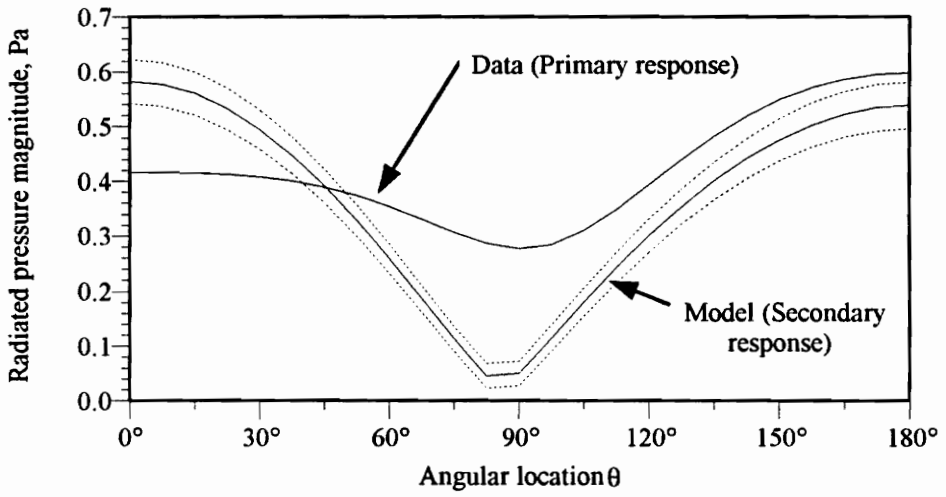

Figure 19: Prediction intervals for radiated pressure estimated model, k_0a=0.4. Estimated model pressure magnitude is shown with 90% prediction based on estimated error variance.

the magnitude of the estimated model as a function of farfield location, with prediction intervals at the 90% significance level; also shown are the data Y. The average width of the prediction interval is roughly 0.06 Pa, while the signal levels are on the order of 0.6 Pa.

Consider a second single-frequency case with k_0a=0.50. The regression indicates that the control inputs b_1 = 0.61 N and b_2 = 0.01 N reduce the radiated power cost function by 2.4 dB. We find that $Q(F, \nu_1, \nu_2)$=0.0041, e.g., there is less than one-half percent probability that this solution arose by chance, so we shall consider the regression statistically significant. However, one of the control forces fails the t-test: we find that the confidence interval b_2 = 0.01±0.31 N contains zero, and thus we cannot conclude that $|b_2|$ is significantly different from zero. Figure 20 shows the prediction interval. At this frequency, the average width of the prediction interval is 0.08 Pa, slightly larger than for the frequency discussed above. To improve the reliability of the regression, we could reformulate it without b_2. (Again, the reader is reminded that the error variance is estimated from the data rather than from knowledge of the sensors, and the results should be considered only as an example of how regression diagnostics may be used.)

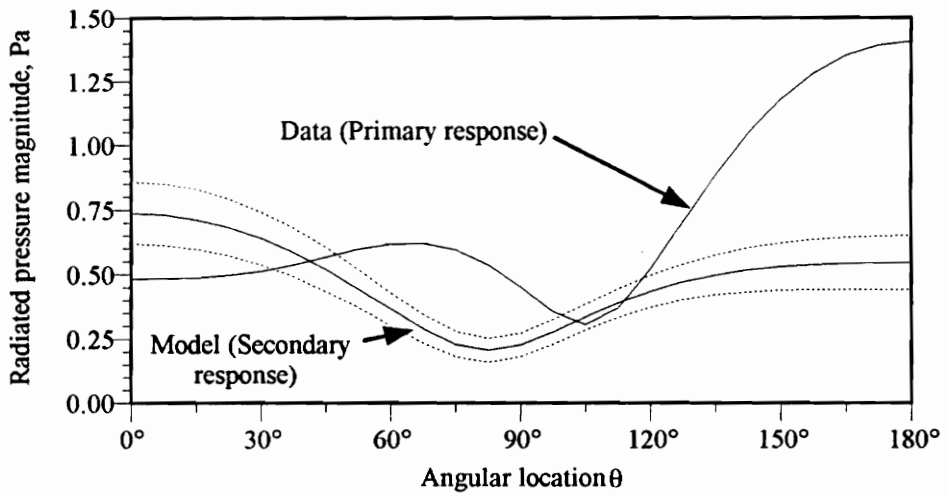

Figure 20: Prediction intervals for k_0a=0.5. Estimated model pressure magnitude is shown with 90% prediction interval based on estimated error variance.

Now consider a third single-frequency case with k_0a=0.74. At this frequency the regression predicts a reduction in the radiated power cost function of less than 1 dB. Performing the F-test, we find that $Q(F, \nu_1, \nu_2)$=0.114. In other words, there is an 11% chance that the estimated model is significantly different from the data. The regression at this frequency fails the F-test by most reasonable standards, and therefore we cannot make use of regression diagnostics such as confidence intervals and prediction intervals. An examination of the data and the estimated model, shown in Fig. 21, confirms that there is little relation between the data and the model. Significant cancellation occurs by chance in a localized area near θ=90°, but in general the shapes of the two curves differ significantly. .

In Fig. 22 we explore the relationship between the F-test results and the reduction of the radiated power cost function, which is related to the squared coefficient of determination $\mathbf{R}^2$. The figure is a scatter plot in which each data point represents a regression at a different frequency in the range $0.2<k_0a<1.0$. The quantity $Q(F, \nu_1, \nu_2)$ is shown on the vertical axis, where $Q(F, \nu_1, \nu_2)$ is the probability that the regression fails the F-test. The upper and lower plots

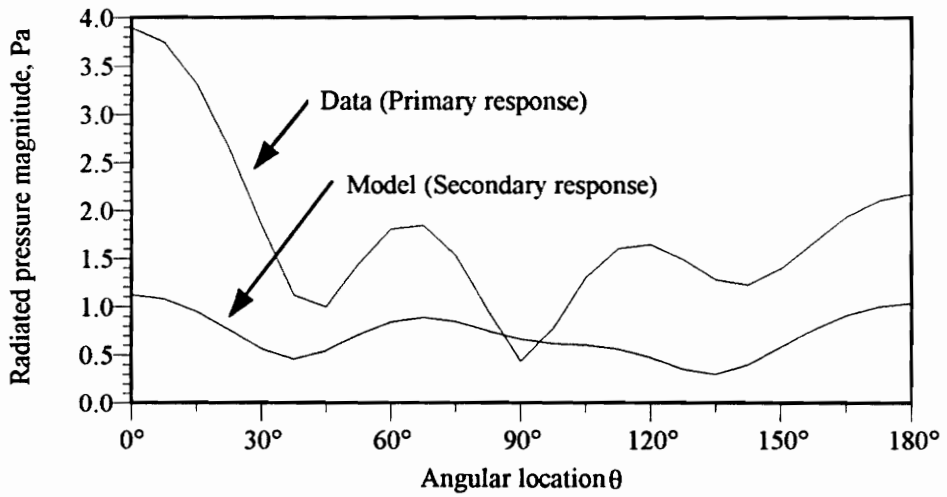

Figure 21: Data and estimated model for k_0a=0.74. Prediction intervals are not computed at this frequency because the regression does not pass the F-test.

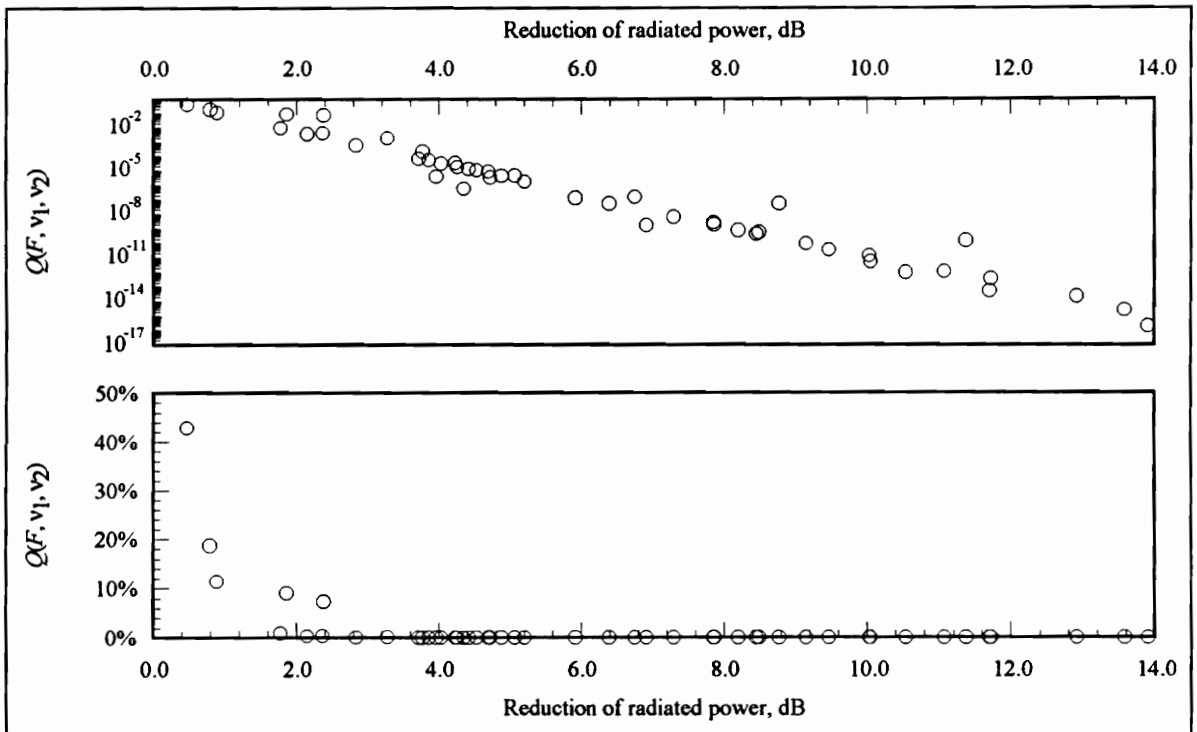

Figure 22: *F*-test probability vs. reduction of radiated power cost function. Each data point represents a different frequency in the range $0.2<k_0a<1.0$. When $Q(F, \nu_1, \nu_2)$ is greater than a specified acceptable risk level, say, α_f=1%, the regression is not statistically significant.

display the same data with different scales on the vertical axes: loglinear on the upper plot to show all the data, and linear on the lower plot to more clearly display values between 1% and 50%. Wherever the radiated power cost function is reduced by 2 dB or more, $Q(F, \nu_1, \nu_2)$ is less than 1% and we have no reason to reject the model. Only when the radiated power cost function is reduced by less than 2 dB does $Q(F, \nu_1, \nu_2)$ rise to appreciable levels. This result is consistent with other force configurations investigated but not shown here. However, the reader should bear in mind that the results could differ greatly for different combinations of X, Y, M, and N.

Thus it appears that failing the *F*-test is *always* accompanied by a low coefficient of determination R, meaning that any attenuation that does occur is purely coincidental and likely

occurs only in localized regions rather than globally throughout the control domain. Likewise, a small value of R appears to always signal that the regression has failed the F-test. One might be tempted to wonder why both measures are necessary. The answer is that while R (and thus the cost function attenuation in decibels) is necessary to provide a familiar measure of controller performance, only the F-test provides a means to decide whether the model is valid or not. The 2 dB criterion discussed above is specific to the X, Y, M, and N used in this example. To obtain a similar result for a different example, the F-test would be required in addition to R. Note, however, that one can always calculate R, even for a SISO system. On the other hand, F is only defined for $M > 1$, that is, more than one sensor.

Figure 23 shows the magnitudes and magnitude confidence intervals of the two control forces as a function of frequency. The solid curves represent the control force magnitudes, with the upper and lower 90% confidence limits displayed as broken curves. The upper and lower plots represent b_1 and b_2, respectively. There are several frequency ranges in which one of the control forces "fails the t-test", meaning that one of the confidence intervals contains zero. For example, the confidence interval for b_2 contains zero near k_0a=0.50, k_0a=0.69, and k_0a=0.91. When a confidence interval contains zero, the associated control force does not contribute significantly to the regression at that frequency and could probably be removed without significantly affecting the results.

3.4. Summary

Multiple least-squares regression provides a numerical approach for simulating feedforward active control in the frequency domain. Solving the regression for a given frequency provides the complex control inputs and the amount by which the cost function is reduced. The squared multiple correlation coefficient $\mathbf{R}^2$ provides the amount of attenuation possible in the absence of

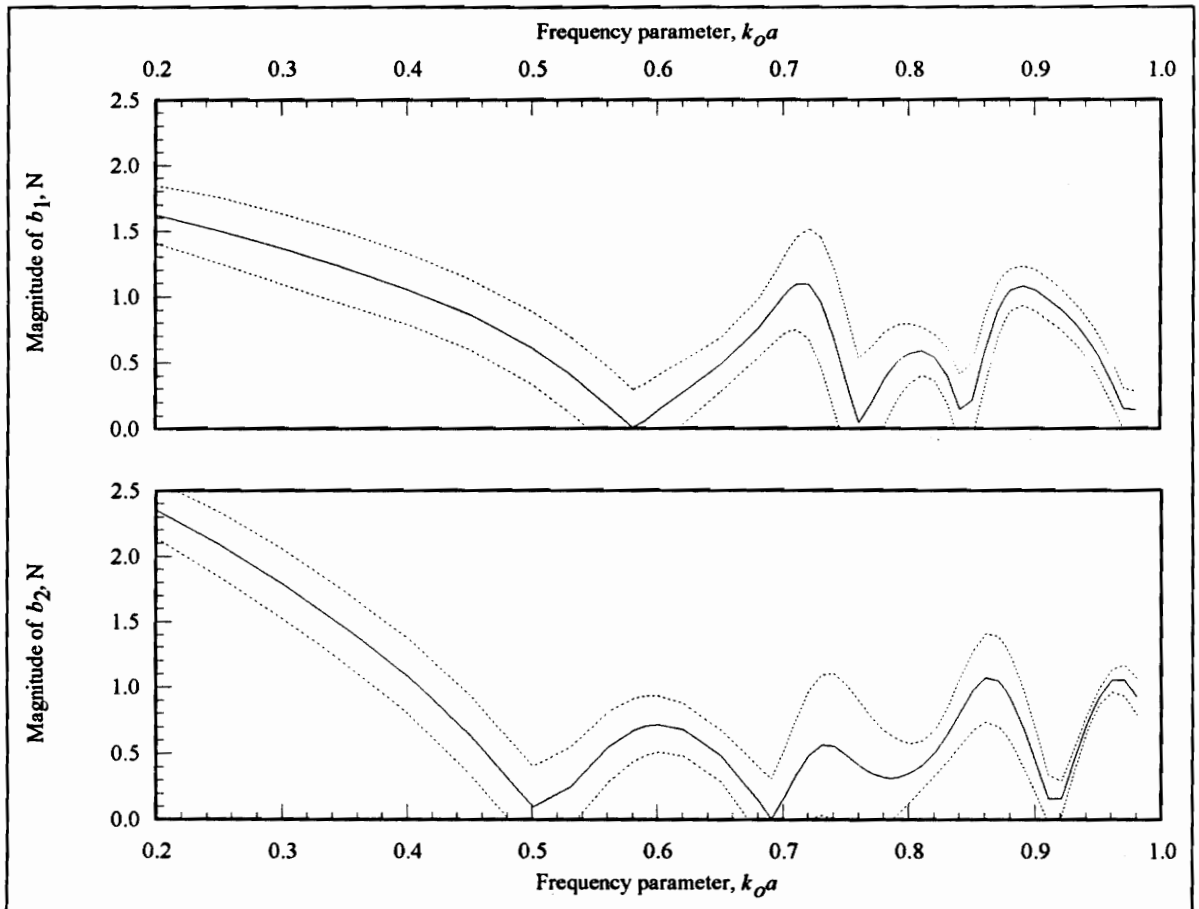

Figure 23: Magnitude of control forces vs. frequency, with 90% confidence intervals on control force magnitude.

measurement noise. In some cases, particularly where the effects of error sensor measurement noise are a concern, these quantities alone do not adequately characterize the regression. To help model the effects of measurement noise, one may consider the F-test, the t-test, confidence intervals, and prediction intervals.

The F-test provides a statistical measure of the integrity of one our basic assumptions: that the model for which we have estimated coefficients is close to the "true model" about which the data are distributed. When the F-test fails, as it does in the example problem whenever the cost

function is reduced by less than 2 dB, it means that the estimated model is significantly different from the true model, and the other regression diagnostics may not be used for that particular frequency. (Note that the reverse is not true: merely passing the F-test does not "prove" our assumption.) The t-test provides a statistical measure of the reliability of the estimated control input magnitudes. When the confidence interval for a given control input contains zero, we cannot conclude that the control input magnitude is significantly different from zero; the width of the confidence interval for an actuator is a measure of its sensitivity to measurement noise. Prediction intervals combine the information from all the actuators to help describe the sensitivity to noise for the regression as a whole.

Before using regression diagnostics, it is important to carefully examine the statistical properties of the measurement noise present in the proposed error sensors. Strictly speaking, the noise must be normally distributed for the F-test, t-test, confidence intervals, and prediction intervals to be valid. Furthermore, the noise should have constant variance, i.e., the variance should be independent of the magnitude of the signal being measured. For some types of sensors, measurement noise is neither normally distributed nor independent of signal magnitude. Although a constant-variance measurement noise is a convenient first approximation, a more sophisticated analysis may be warranted in some cases.

CHAPTER 4: COLLINEARITY AND COLLINEARITY DIAGNOSTICS

Chapter 3 describes how feedforward control in the frequency domain is similar to multiple linear least-squares regression in terms of the numerical procedures involved. Given the similarities between regression and feedforward control, it is logical to assume that numerical difficulties commonly experienced by regressions might also appear in feedforward control simulations. Such is the case with collinearity, the topic of this chapter. The methods developed in previous chapters use transfer functions between control inputs and system dynamic responses. Consider the transfer functions as vectors in M-dimensional space: collinearity is a type of numerical ill-conditioning that occurs when two transfer functions are parallel or "co-linear" with each other. But it is possible for a control system to exhibit collinearity even if two transfer functions are not exactly parallel. In general, collinearity occurs whenever two *or more* transfer functions are not mutually orthogonal to one another. Since it is unlikely that the transfer functions in a real dynamic system will all be mutually orthogonal, collinearity is almost always present to some degree. In fact, control simulations are perhaps more prone to collinearity than other regressions. To simulate a feedforward controller without checking for collinearity is to risk producing meaningless results. In this chapter we define collinearity and give a physical explanation for its equivalent in the feedforward control problem. Also given is a basic diagnostic procedure for detecting and analyzing collinearity.

Two statistical references that provide a comprehensive discussion of collinearity diagnostics for real-valued regressions are Belsley, Kuh, and Welsch [37] and by Belsley [38]. All of the concepts discussed in this chapter are based on the techniques in [38]. The extensions to complex-valued regressions are straightforward, and are based on the work of Miller [29].

4.1. Collinearity in feedforward control

4.1.1. How collinearity enters a feedforward control problem

Suppose we apply two control forces at exactly the same location on the structure. Obviously their transfer functions are identical to each other, the normal equations are singular, and we have no unique solution for the control force magnitudes. Now suppose the two control forces are close together, but not quite collocated. Their transfer functions are similar, but not identical; the normal equation are ill-conditioned, but not quite singular. In fact, the possibility for such near-dependencies to cause ill-conditioning exists whenever the transfer functions are not all mutually orthogonal.

Suppose we plunge forward and solve the normal equation; what sort of solution can we expect? Often the adjacent (and partially redundant) control forces have opposite phases and very large magnitudes. In some circumstances this is purely a numerical artifact, the result of ill-conditioning of the normal equation. In other circumstances such results could be interpreted as a force couple approximating a concentrated moment, and there might be no indication of whether the solution was genuine or spurious. This disturbing ambiguity is the motivation for investigating collinearity diagnostics.

Control forces need not be located next to each other to produce collinearity. If the modal density is low and the structure is very lightly damped, there may be several locations at which applied forces would produce nearly the same response. Control forces applied at two or more such locations would produce collinearity even though they were distant from each other. Another source of collinearity arises when one control force can be replaced by a linear combination of other control forces in the system.

4.1.2. An example showing the effects of collinearity

A simple example shows how collinearity can affect even a relatively simple simulation. Consider the clamped cylinder (see Fig. 3, p.14) with a disturbance force b_d acting at a frequency of $k_o a$=0.95. Let us choose three control actuators b_1, b_2, and b_3, whose locations we shall discuss after performing a brief numerical experiment. Solving the regression, we predict that the control inputs reduce the radiated power cost function by 23 dB. Since the F-test and t-test results (see Chapter 3) appear reasonable, we accept as significant the following complex control inputs:

$$b_1 = (-18.1,-0.8) \quad b_2 = (-21.3,-2.2) \quad b_3 = (131.9,11.4) \tag{56}$$

where, for example, b_1=(-18.1,-0.8) means that Re$\{b_1\}$=-18.1 and Im$\{b_1\}$=-0.8.

The purpose of our numerical experiment is to see how measurement noise in the error sensor data $Y=\{y_1, y_2, \ldots y_M\}$ affects the regression results. The method is to simulate the presence of measurement noise by artificially injecting various amounts of random noise into Y, recomputing the regression, and examining the control input magnitudes. The results are summarized in Table 2. The first pair of columns, labeled "No simulated noise," contains the original regression results given in Eq. (56). Next, each error sensor output y_i is perturbed with 0.5% simulated measurement noise, that is, a noise component with random phase and random magnitude up to 0.5% of the maximum value in Y, and the regression is recomputed using the perturbed Y. As shown in the second pair of columns in Table 2, 0.5% simulated measurement noise did not significantly change the predicted control inputs. Next we increase the simulated measurement noise to 5% and recompute the regression. As shown in the third pair of columns in Table 2, the control input magnitudes are changed by less than 2%. In fact, we can increase the simulated measurement noise to 25% or more before the control inputs change substantially. Thus the regression appears relatively insensitive to measurement noise.

Table 2: Effect of simulated measurement noise without collinearity.

	No simulated noise		0.5% simulated noise		5% simulated noise	
	b_i	$\Delta\|b_i\|$	b_i	$\Delta\|b_i\|$	b_i	$\Delta\|b_i\|$
b_1	(-18.1,-0.8)	--	(-18.0,-0.8)	<0.5%	(-17.7,-0.8)	2%
b_2	(-21.3,-2.2)	--	(-21.3,-2.2)	<0.5%	(-21.0,-2.2)	1%
b_3	(131.9,11.4)	--	(131.9,11.4)	<0.5%	(131.8,11.3)	1%

Table 3: Effect of simulated measurement noise with collinearity.

	No simulated noise		0.5% simulated noise		5% simulated noise	
	b_i	$\Delta\|b_i\|$	b_i	$\Delta\|b_i\|$	b_i	$\Delta\|b_i\|$
b_1	(-627.7,-37.0)	--	(-604.5,-49.0)	5%	(-195.1,-78.9)	67%
b_2	(613.5,36.1)	--	(589.7,48.5)	2%	(164.9,78.2)	70%
b_3	(117.3,9.6)	--	(118.1,9.4)	1%	(128.2,9.6)	9%
b_4	(24.1,1.4)	--	(23.2,1.9)	2%	(7.0,3.1)	68%

Now consider another example in which we add a fourth control input b_4 to the configuration. (For the sake of illustration we choose a particularly bad location for b_4.) For the new configuration with no simulated measurement noise, we predict that the control inputs reduce the radiated power by 29 dB, an improvement of 6 dB over the previous case. Again the F-test and t-test results appear reasonable, so we accept as significant the control inputs shown in the first pair of columns in Table 3. As in the previous example, we add simulated measurement noise to Y and recompute the regression. But now, even with an error of only 0.5%, we see changes of up to 5% in the predicted control inputs. If we increase the simulated measurement noise to 5%, the control inputs are all but unrecognizable, although somehow the third input seems less drastically affected. Still, the F-test and t-test show no signs of trouble, and the subroutine used to solve the normal equation gives no warning or error messages. Why is the solution so sensitive to

measurement noise? The reason, of course, is collinearity between the particular control forces we chose.

If we now examine the locations of the control inputs, shown in Fig. 20, we can clearly see why we encounter numerical difficulties. The first two inputs, b_1 and b_2, are axisymmetric ring-forces at δ=0.17 and δ=0.31, respectively, while b_3 is a distributed axisymmetric ring-force. But the control input used in the second example, b_4, is an axisymmetric ring-moment applied at δ=0.25, exactly between b_1 and b_2. The two ring-forces, when combined with opposite phases, produce a force couple that has nearly the same effect as the ring-moment. The ring-moment is redundant, and in this simple example can easily be seen as the cause of the numerical problems. The collinearity diagnostics presented in the next section provide a general method that could be used to detect and analyze collinearities in more complicated systems where relations between actuators might not be so clear.

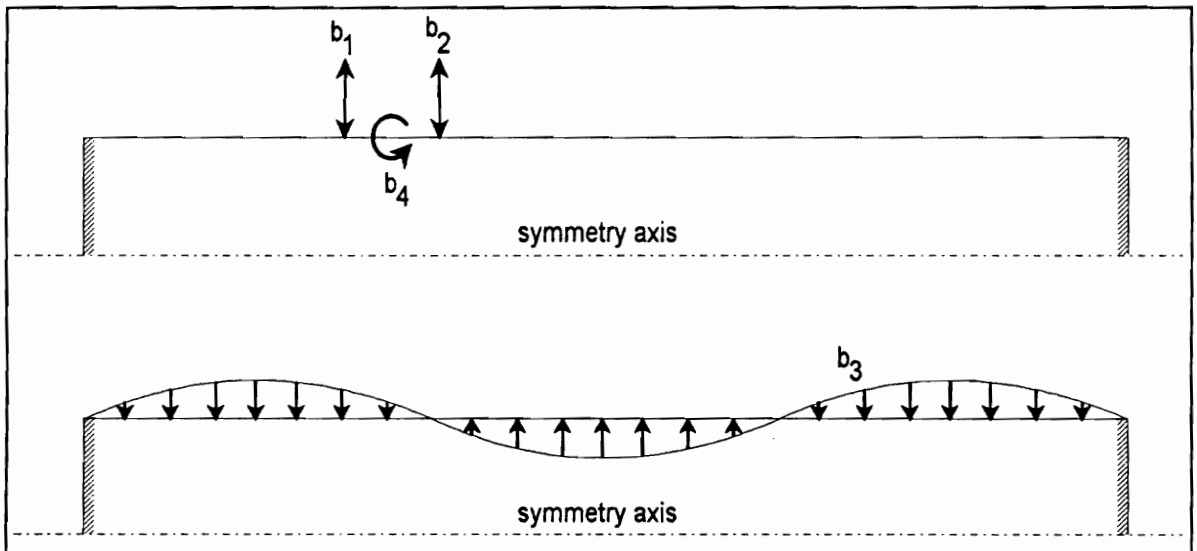

Figure 20: Example problem for collinearity diagnostics. The first two actuators are axisymmetric ring-forces b_1 and b_2 applied at δ=0.25 and δ=0.31. The third actuator is a axisymmetric distributed force b_4=cos(3$\pi\delta$). The fourth actuator is an axisymmetric ring moment b_3 applied at δ=0.31. (The distributed force is also used as the disturbance input for subset selection in Chapter 5.)

4.1.3. The role of approximate solution techniques

Generally some redundancy between the variables is unavoidable, and must be expected and planned for in one of several ways. One method, discussed in the next section, uses collinearity diagnostics to analyze sources of collinearity and suggest strategies for removing collinearity from the problem. Another method is to simply circumvent the effects of collinearity by employing an approximate solution method to solve the normal equations. The former method appears to be the more appropriate: merely solving the normal equations does not eliminate the collinearity problem, and does nothing to illuminate its causes and possible solutions. However, if the situation at hand requires no more than workable solutions to the normal equations, the latter method may be used. Two examples exist in the literature, and are described below.

Rosenthal [25] finds approximate solutions to the normal equations by using singular-value decomposition (SVD). (This method was used to find multi-force results in Chapter 2 of this thesis.) The SVD approach, described in [39], provides an approximate solution that minimizes the length of the solution vector. This has an interesting interpretation in the context of ASAC, since the length of the solution vector is a measure of the amount of control effort expended. The SVD method also provides convenient access to the covariances. A drawback of using SVD is that one must manually decide on a threshold beyond which singular values are discarded, and the threshold is somewhat problem-dependent. Nevertheless, the SVD provides a reliable way to solve the normal equations in the presence of collinearity, and the singular vectors may have interesting physical interpretations as well.

Nelson and Elliott [15] discuss the use of iterative gradient descent techniques as an approximate solution technique. Iterative techniques are able to find a reliable and stable solution to the normal equations even in the presence of collinearity. They are computationally efficient

for practical applications, in which the problem parameters may be slowly varying in time. Also, they do not require the user to select a threshold value as in SVD, although the user must determine when to discontinue the iterations.

4.2. A collinearity diagnostic

Rather than trying to present a comprehensive review of collinearity diagnostics and their possible applications in studying feedforward control, this section presents one diagnostic method with some basic example results. For more detail the reader is referred to Belsley [38].

4.2.1. Background discussion

Recall the normal equation for a feedforward control system with M sensors and N actuators,

$$\left(X^H W X\right) B = X^H W Y. \tag{57}$$

As the reader might expect, many collinearity diagnostics involve the condition number of $X^H W X$, the matrix that forms left-hand side of the normal equation. Clearly a large condition number announces the presence of ill-conditioning, but simply examining the condition number leaves unanswered three questions. First, how can we detect whether more than one near-dependency is present? Second, how large a condition number can we tolerate and still have confidence in the regression results? And third, how can we determine which variables (or, in the context of feedforward control, which control inputs) are involved in near-dependencies? We will briefly discuss these three questions before presenting the diagnostic method.

The first question, regarding the number of near-dependencies present, is answered by examining not just the condition number of $X^H W X$ but the eigenvalues of $X^H W X$. We accomplish this by examining the singular-value decomposition of $W^{1/2} X$, since the singular values of $W^{1/2} X$

are the positive square roots of the eigenvalues of $X^H W X$. A zero singular value indicates a perfect linear dependency in $W^{1/2}X$; a "small" singular value indicates a near-dependency that may cause conditioning problems. In fact, the number of "small" singular values indicates the number of near-dependencies present in the problem.

The second question is: how do we make quantitative comparisons to determine whether the condition number is "large", or equivalently, whether a given singular value is "small"? The answer is that we must scale the columns of $W^{1/2}X$ before computing the singular values,. Column scaling , described in the next section, permits such quantitative comparisons.

To answer the third question of determining which variables are involved in near-dependencies, we turn to a result from the least-squares regression. One of the effects of collinearity is variance inflation: when two or more actuators are involved in a near-dependency, their variances become large. We can express the variances in terms of the singular values of $W^{1/2}X$. When we find two or more actuators whose variances are primarily associated with a "small" singular value, then those actuators are involved in a near-dependency.

We now describe the diagnostic method in somewhat of a cookbook fashion. The intent is not to present a detailed derivation, but rather to outline the method and show a simple numerical example related to feedforward control. We modify the technique slightly to account for complex-valued variables and weighted least squares, which are not included in the discussion of [38].

4.2.2. Step 1: Determine $W^{1/2}X$

Each column of $W^{1/2}X$ represents a transfer function between a given control input and the error sensor array. These may be obtained via analytical approximations, numerical calculations, or any other means available. For N actuators and M sensors, X has the form

$$X = \{X_1 \quad X_2 \quad \cdots \quad X_N\} \tag{58}$$

where X_j is a complex vector of length M containing the transfer function for the j^{th} actuator, and W is the weighting matrix. In our example the columns of X are transfer functions between the four control inputs and the farfield pressure at 25 sensor locations, thus N=4 and M=25. We obtain the transfer functions by numerical modeling of the structure using a finite-element/boundary-element approach.

4.2.3. Step 2: Apply column-equilibration to $W^{1/2}X$

The purpose of column-equilibration is to produce a set of singular values that do not depend on the units of the problem. We scale each column in $W^{1/2}X$ to have unit vector length after scaling. This produces a new matrix Z, the scaled weighted design matrix:

$$Z = \Omega^{-1/2} W^{1/2} X, \tag{59}$$

where $\Omega^{-1/2}$ is a diagonal matrix whose diagonals are $\Omega_j^{-1/2} = \left\{X_j^H W X_j\right\}^{-1/2}$. The relation between Z and the original problem may be seen from

$$X^H W X = Z^H \Omega Z \tag{60}$$

4.2.4. Step 3: Obtain scaled condition indexes and variance-decomposition proportions

First we generate the singular-value decomposition of Z,

$$Z = UDV^H, \tag{61}$$

where U is column-orthogonal, V is column-orthogonal and row-orthogonal, and D is non-negative and diagonal. (Note the distinction between the matrix V and the variance V). The elements of D are the singular values μ_k, k=1,2...N, which are real and non-negative since Z is Hermitian. The *scaled condition indexes* are then defined as

$$\widetilde{\eta}_k = \frac{\mu_{\max}}{\mu_k}, \quad k = 1,2...N. \tag{62}$$

If we express the elements of the matrix V as v_{kj}, we can write the variances of the regression coefficients in the following useful form.

$$V(b_k) = \sigma_P^2 \sum_{j=1}^{N} \frac{v_{kj}^2}{\mu_j^2} \tag{63}$$

The utility of Eq. (63) lies in the fact that each term in the summation is associated with only one scaled condition index $\widetilde{\eta}_k$. By noting which of the "large" scaled condition indexes affect which variances, we can pick out near-dependencies among the columns of $W^{1/2}X$. To this end, let us define *variance-decomposition proportions*,

$$\pi_{jk} \equiv \frac{\gamma_{kj}}{\gamma_k}, \quad k = 1,...N., \tag{64}$$

where

$$\gamma_{kj} \equiv \frac{v_{kj}^2}{\mu_j^2} \quad \text{and} \quad \gamma_k \equiv \sum_{j=1}^{N} \gamma_{kj}, \quad k = 1,...N. \tag{65}$$

The variance-decomposition proportions are interpreted as follows: π_{jk} is the portion of the variance for the k^{th} regression coefficient that is associated with the j^{th} scaled condition index.

Table 4 contains the scaled condition indexes and variance-decomposition proportions for the example problem of Fig. 20. The second column contains the scaled condition indexes ranked by

Table 4: Variance decomposition proportions for example problem.

j	$\widetilde{\eta}_j$	π_{j1}	π_{j2}	π_{j3}	π_{j4}
1	1.0	.000	.000	.000	.000
2	1.1	.000	.000	.000	.000
3	6.3	.000	.000	.871	.000
4	11502.1	.999	.999	.129	.999

size, with the largest shown in the last row. The remaining four columns contain the variance-decomposition proportions for each of the control forces b_1, b_2, b_3, and b_4. The next sections explain how to interpret Table 4.

4.2.5. Step 4: Determine the number of near-dependencies.

The method for counting the number of near-dependencies involves choosing a threshold $\tilde{\eta}^*$ such that when $\tilde{\eta}_{\max} < \tilde{\eta}^*$, the columns of $W^{1/2}X$ may be said to be free of collinearity. The value suggested by Belsley [38] is

$$\tilde{\eta}^* = 30, \tag{66}$$

although slightly larger or smaller values could also be argued for. This is the value used for all numerical results in this dissertation. Having chosen a threshold value for $\tilde{\eta}^*$, we now have the most basic collinearity diagnostic: whenever $\tilde{\eta}_{\max} < \tilde{\eta}^*$, at least one near-dependency is present and the regression should not be used without further examination of its collinearity properties. Thus, $\tilde{\eta}_{\max}$ should be computed as a routine part of every regression, and displayed along with the *F*-test and *t*-test results described in Chapter 3.

If more than one scaled condition index exceeds the threshold $\tilde{\eta}^*$, then multiple near-dependencies are present. The number of near-dependencies is determined according to the so-called "progression of 10/30": the first near-dependency corresponds to the first $\tilde{\eta}$ >30; the second near-dependency corresponds to $\tilde{\eta}$ >100, the third to $\tilde{\eta}$>300, the fourth to $\tilde{\eta}$>1000, and so forth. For example, the sequence of scaled condition indexes (1, 9, 33, 35) indicates only one near-dependency, while the sequence (1, 9, 33, 105, 108) indicates two near-dependencies.

Referring to the first column of Table 4, it is evident that only one scaled condition index exceeds $\tilde{\eta}^*$, so we conclude that only one near-dependency is present in the example.

4.2.6. Step 5: Determine actuator involvement

To determine which actuators are affected by a near-dependency, we examine the variance-decomposition proportions π_{jk}. Having associated a particular $\tilde{\eta}$ with a near-dependency, we examine the π_{jk} in that row of the table. If two or more actuators have π_{jk} in that row larger than some threshold value π^*, then those actuators are involved in the near-dependency. The value of π^* suggested by Belsley [38] and used in our example is

$$\pi^* = 0.5. \tag{67}$$

In Table 4 we examine the last row since that row is associated with a near-dependency. Noting that the π_{jk} for b_1, b_2, and b_3 are all well above π^*, we can clearly see that these actuators are involved in the near-dependency. (Recall that b_1 and b_2 are the ring-forces and b_4 is the ring-moment.) The third actuator, with π_{jk}=0.129, is not involved in the near-dependency.

The rules for determining variate involvement when multiple near-dependencies are present are somewhat more complicated and need not be discussed here to illustrate the basic technique. Again, the interested reader is invited to examine Belsley [38].

4.2.7. Step 6: Determine auxiliary regressions

Once we know which actuators are involved in a near-dependency, we can regress the individual transfer functions against each other to determine the relationships between them. For example, we have determined that b_1, b_2, and b_4 are involved in a near-dependency. If we set up a new regression using b_4 as the disturbance input with b_1 and b_2 as control inputs, we can see that the combination of b_1 and b_2, if given equal magnitudes and opposite signs, reproduces b_4 almost exactly (as we suspected.) Thus the control input b_4 could be removed to obtain a more stable regression result.

4.2.8. Step 7: Determine unaffected actuators

When an actuator associates most of its variance with "small" values of $\tilde{\eta}$, we can consider that actuator to be uninvolved with any near-dependency. To be more precise, when the total proportion of variance associated with "small" $\tilde{\eta}^*$ is less than π^*, that actuator is uninvolved with any near-dependency. In the example problem, the variance of b_3 is associated almost entirely with the first three "small" $\tilde{\eta}^*$, and therefore we conclude that it is not involved in any near-dependency. This is supported by the fact that b_3 was the least affected by the simulated measurement noise introduced in the example discussed earlier.

4.3. Summary

The collinearity problems common in many multiple linear regressions can also occur when regression is used for a numerical simulation of feedforward active control in the frequency domain. Without proper collinearity diagnostics, numerical problems can go undetected and produce results that are very sensitive to measurement noise in the error sensors. Collinearity diagnostics can detect and analyze numerical ill-conditioning. Collinearity diagnostics for complex-valued regressions follow directly from their real-valued counterparts, which are developed in considerable detail in the statistics literature.

The primary collinearity diagnostic is the scaled condition number. Scaling allows for a quantitative comparison between the scaled condition number and a predetermined threshold value to determine whether or not collinearity is present. Even when no other collinearity diagnostics are performed, the analyst may compute the scaled condition number for each regression and compare it to the threshold value in a somewhat mechanical fashion. If the scaled condition number exceeds the threshold, the regression should be discarded because of collinearity.

More detailed diagnostics are available for determining the number of near-dependencies and which control forces are involved. These diagnostics are less mechanical, requiring more interaction on the part of the analyst; however, they may provide insights when investigating complicated control systems.

CHAPTER 5: OPTIMIZING ACTUATOR LOCATIONS USING SUBSET SELECTION

In most multi-actuator control systems, the physical locations of the actuators can profoundly affect the performance of the control system. Properly located actuators can potentially give dramatically better performance than actuators whose locations are chosen at random [40]. For simple systems the optimal actuator locations can be fairly easy to predict based on physical knowledge of the system. But for more complicated systems the optimal actuator locations may be far less obvious, especially when many actuators are used. The purpose of this chapter is to present a formal, general method for optimizing the actuator locations for a given disturbance.

Optimization schemes can be classified as using either continuous or discrete. In a continuous problem, the domain of possible solutions is a continuous function in parameter space, while in discrete problems only certain discrete points in the parameter space are admissible solutions. The technique presented in this Chapter is a form of discrete optimization.

Continuous problems are generally approached using numerical optimization in which the actuator locations are included as optimization parameters [40,41]. For example, [40] optimizes the locations of piezoelectric actuators used to control acoustic radiation from a rectangular plate, achieving dramatically better performance than with randomly located actuators. Such studies are instructive, but are limited to relatively simple problems with closed-form solutions because of the high computational expense of repeated response evaluations. Also, minimization routines can become trapped by local minima and fail to find global minima.

Discrete problems require different optimization techniques than continuous problems; such techniques are often referred to as "integer programming" (see, for example, [42]). Discrete optimization has been applied to truss structures or other systems for which the permissible actuator locations are naturally discrete, but to the author's knowledge has not been used for optimizing transducer arrangements in ASAC systems. In this Chapter, we analyze a continuous structure that has been discretized by finite-element modelling. Therefore, discrete optimization techniques are a natural way to examine the problem. The method has practical appeal for applications in which only certain locations are acceptable for mounting actuators; for example, an aircraft fuselage might require actuators mounted on structural reinforcements.

The optimization technique used in the present work is based on the concept of *subset selection*. The method begins with the analogy between feedforward control and complex-domain multiple linear regression discussed in Chapter 4. The main benefit of subset selection over other methods is its extreme computational efficiency, which allows more comprehensive studies with larger numbers of actuators. Since the method is based on discretized transfer functions, it is quite general and is ideal for simulations involving numerical models or experimental data.

To help explain subset selection, we optimize the locations of ring-force actuators on the clamped cylindrical shell of Fig. 3 (page 14). The set of candidate actuator locations, shown in Fig. 25, includes 25 axisymmetric ring-forces evenly spaced along the length of the cylinder. They are numbered sequentially with actuator 1 at δ=0.02, actuator 13 at δ=0.50, and actuator 25 at δ=0.98. The reader should note that the use of subset selection to optimize the actuator locations is completely separate from any statistical modeling of error sensor measurement noise. The regression diagnostics of Chapter 3 may be used to model some of the effects of measurement noise, but in this chapter we assume that no measurement noise is present.

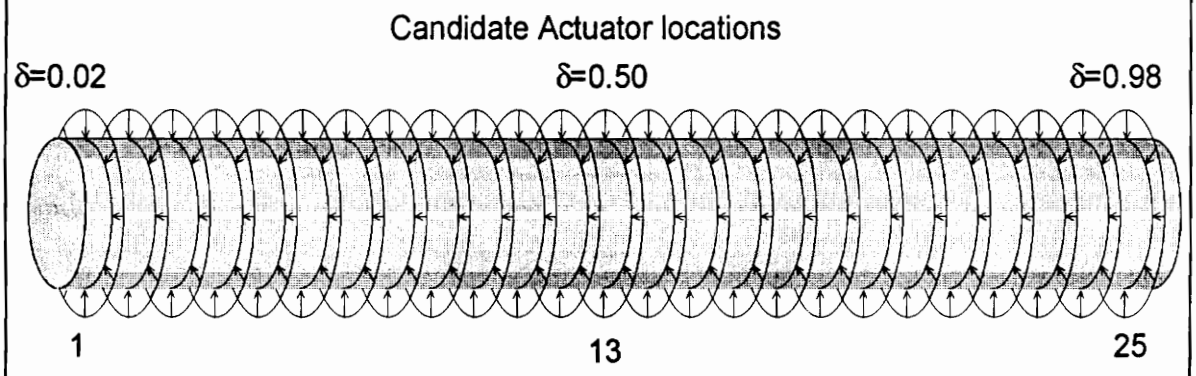

Figure 25: Locations of candidate actuators. All 25 candidate control forces are axisymmetric ring-forces.

5.1. Theoretical background

5.1.1. Using subset selection to optimize actuator locations

In the statistical sciences, one often has access to many more predictor variables than are necessary (or prudent) to include in a model. One might be tempted to include all available variables in hopes of devising the most complete model, but usually this is difficult for practical reasons. Some variables might be too expensive to measure; others might duplicate information already contained in other variables, or might not contribute any new information to the fit. For these reasons and others, statisticians have developed a large body of literature describing the concept of *subset selection* or *variable reduction* (see [43] for a detailed discussion and extensive bibliography; see also [44-46]). The goal is to pick out, from among a large set of candidate variables, some "best subset" of variables that provides an adequate fit to the data.

Similarly, in the feedforward control problem, there are usually many possible locations for the actuators. Cost and complexity preclude covering the entire structure with actuators at every

possible location. (Also, using too many actuators would produce massive collinearity as discussed in Chapter 4.) Since it is usually neither necessary nor practical to use a large number of actuators, subset selection lets us choose an optimal subset of actuator locations from among all the candidates, in effect optimizing their locations for the given disturbance and frequency. The method requires computation of one transfer function per candidate actuator location. Therefore, a numerical approach based on finite elements is ideally suited: when we compute transfer functions for one actuator location, we can at no additional cost obtain transfer functions for all the other actuator locations.

The literature describes a variety of approaches to subset selection, each with different strengths and levels of complexity; the most straightforward method is *exhaustive search subset selection*. In an exhaustive search, we compute a separate regression for *every possible combination* of actuator locations. For each regression, we compute the value of the residual cost function χ^2, and keep track of which subsets produce the smallest χ^2 values. The subset that produces the lowest χ^2 is the optimal subset; other subsets that produce near-optimal values of χ^2 may be useful also. An exhaustive search is guaranteed to find the global minimum χ^2 rather than becoming trapped in a local minimum.

Because exhaustive-search techniques can require long computation times, the present work uses a "reduced-order" form of exhaustive-search subset selection. The number of possible subsets of one or more variables out of N is $(2^N - 1)$, which roughly doubles each time N increases by one. A true exhaustive search with 20 or more variables is not practical with current facilities. However, in the present context we are not interested in large subsets, that is, subsets that include many actuators. It is more logical to only investigate subsets of, say, n or fewer actuators out of the possible N candidates, where n may be specified for the problem at hand. Suppose we wish to

choose from among 25 candidate actuator locations to optimize a 6-channel control system (N=25, n=6). It would be inefficient or impossible to examine all subsets of *all sizes*, including subsets of 24 actuators or 25 actuators, when the control system can use no more than six actuators. Limiting the search to subsets of six or fewer actuators reduces the computation time by many orders of magnitude. Reduced-order exhaustive search permits the use of many more candidate actuator locations, and is used exclusively in the present work.

Other types of subset selection include forward selection, backward elimination, sequential replacement, and various "stepwise" methods. These require less computing effort than exhaustive searches, but bring into question whether the solutions found are global minima. For purposes of demonstration, only exhaustive-search subset selection is used in this thesis.

Before describing the mechanics used to perform subset selection on complex-valued regressions, a few words of caution are in order. As with most forms of numerical optimization, one should not treat subset selection as a "black box" technique that will produce the single best configuration of actuators. The analyst should examine a large number of near-optimal subsets over a range of frequencies looking for actuator locations that appear again and again. This point cannot be overstated; one must vigilantly guard against putting too much confidence in a numerical technique without a good understanding of the physics involved.

5.1.2. Exhaustive search subset selection

Exhaustive-search subset selection performs regressions for many different subsets of actuators to find which subset produces the smallest residual cost function χ^2. Using a reduced-order exhaustive search cuts computation times by orders of magnitude as described below, but the nature of subset selection suggests two other time-saving measures. Firstly, computing the

regression coefficients requires extra computational effort beyond that required to compute χ^2. Therefore, we should not waste effort solving for the regression coefficients because they are not needed to evaluate the worth of a particular subset. Secondly, we should choose a solution technique that allows us to remove or add variables from the regression without regenerating the entire computation. Both goals are accomplished by using a *QR* factorization as described below.

Gentlemen [48, 49] gives an efficient computational method for *QR* factorization of real-valued systems of normal equations. After performing the factorization and including the effects of weighting coefficients, the normal equations may be written

$$\overline{R}B = \overline{\theta} \tag{68}$$

where $\overline{R}$ is an N by N upper triangular matrix and $\overline{\theta}$ is a vector of length N. If the actual values of the regression coefficients are required, solving for B is trivial by back-substitution.

For subset selection, the utility of Eq. (68) is twofold. First, noting that each row of Eq. (68) represents a different actuator, it can be shown that the i^{th} element of $\overline{\theta}$ is proportional to the value of χ^2 that would result from including only actuators 1, 2, ... i. Secondly, the rows may be re-ordered with minimal effort by taking advantage of the matrix structure already present in the equations. By successively re-ordering the equations and examining the diagonals of $\overline{\theta}$, we can rapidly examine the necessary subsets to determine which subsets produce the most effective controller. Subsets are examined in order of increasing subset size: subsets of one actuator are examined first, then two actuators, three actuators, and so forth. For a reduced-order exhaustive search, the process is terminated after subsets of a given size have been examined.

A general set of routines for weighted least-squares regression is presented by Miller [50]. However, these routines are for real-valued variables only. The routines can easily be modified to

operate on complex-valued variables; the modifications are described in Appendix B. For further details on exhaustive search using *QR* decomposition, the reader is referred to [43].

All regression and subset selection computations for this paper were performed on a desktop computer, and none required more than a few minutes of computing time. (The transfer functions were computed separately, before performing subset selection.) The high computational efficiency of the *QR* decomposition allows large systems to be investigated economically.

5.1.3. The effect of collinearity on subset selection

Collinearity degrades the accuracy of the regression, and if severe can completely destroy the results. A common result is *bias error*, that is, an erroneously small χ^2 which is actually a numerical artifact. Thus we see the impact of collinearity on subset selection: if we record the subsets with the smallest χ^2 values, some of them may be contaminated by collinearity. We must test the near-optimal subsets to see which, if any, exhibit collinearity and should be discarded.

Detecting collinearity is not difficult. We simply compute the largest scaled condition index (see Section 4.1.4) and compare it to the scaled condition index threshold π^* (Section 4.1.5). If the threshold is exceeded, we discard the regression as a numerical artifact.

As described in Section 4.1.3, we could still solve the regression with collinearity present by using an approximate method such as singular-value decomposition (SVD). However, SVD and other approaches are not suitable for use with subset selection because of the computing time involved. Recall that the exhaustive search is based on repeatedly re-ordering the equations and solving for χ^2. Singular value decomposition would require the algorithm to completely regenerate the equations from scratch each time the variables are rearranged, rather than taking advantage of previous computations.

5.2. *Results and discussion*

To help explain subset selection, we optimize the locations of ring-force actuators on the clamped cylindrical shell of Fig. 3 (page 14). The disturbance is the distributed axial force shown in Fig. 20 (page 73), and the disturbance frequency is $k_o a$=0.59. The set of candidate actuator locations, shown in Fig. 25, includes 25 axisymmetric ring-forces evenly spaced along the length of the cylinder. They are numbered sequentially with actuator 1 at δ=0.02, actuator 13 at δ=0.50, and actuator 25 at δ=0.98. There are 25 error sensors (M=25) measuring farfield pressure, and the cost function being minimized is the radiated power cost function. The reader is reminded that this chapter, unlike Chapter 3, assumes there is no error sensor measurement noise.

Figure 26 shows the distribution of normal surface velocity due to the distributed disturbance force. The horizontal axis shows the location coordinate δ on the shell, where the end closures have been omitted and $0.0 \leq \delta \leq 1.0$. The vertical axis shows the magnitude of the normal surface velocity. The surface velocity has local maxima at approximately δ=0.1, δ=0.5, and δ=0.9. Judging solely from this plot, an intuitive guess for the optimal actuator locations would be to locate them near these "anti-nodes" or local surface velocity maxima, since the actuators would be ineffective if located near the nodes of the vibration pattern where the mobility is low.

Following the procedure outlined in previous sections, we must first form and evaluate all

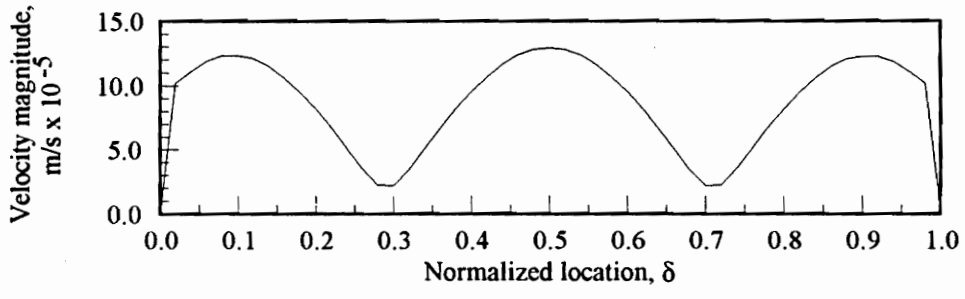

Figure 26: Operating shape at $k_o a$=0.59 with disturbance force used in subset selection calculations.

subsets of n or fewer actuators, and then check the resulting near-optimal subsets for collinearity. In this case the number of candidates is N=25, and the maximum subset size was chosen as n=5. Using the QR decomposition method and the techniques described in Appendix B, we evaluate the residual χ^2 for all subsets of five or fewer actuators out of the possible 25 candidates, recording the optimal subset and a number of near-optimal subsets. For subsets of one actuator, only 25 different subsets are possible (one per candidate actuator location). Subsets of one actuator cannot display collinearity because collinearity requires at least two actuators. Among subsets of two actuators, the 75 best subsets were tested for collinearity. No collinearity was found, i.e., in each case the scaled condition number was below the threshold. Likewise, among subsets of three actuators and four actuators, no collinearity was found among the 75 best subsets.

Among subsets of five actuators, every one of the best 75 subsets was contaminated by collinearity and was therefore discarded. With five actuators on such a geometrically simple structure, it becomes difficult to find actuator locations that do not produce collinearity. Subsets of five actuators could possibly exist without collinearity, as could subsets of six or more actuators, but for the discussion below only subsets of four or fewer actuators are discussed.

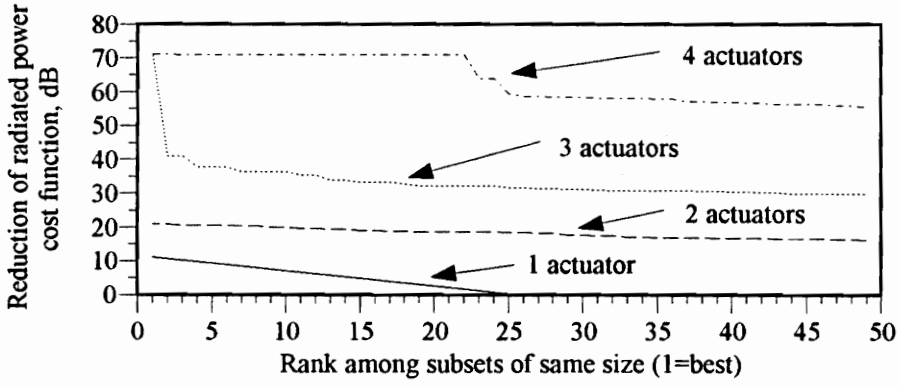

Figure 27: Reductions in radiated power cost function with best actuator subsets of 1, 2, 3, and 4 actuators.

5.2.1. Residual radiated power levels with optimized actuator locations

Figure 27 shows the extent to which radiated power cost function is reduced by optimal or near-optimal subsets of one, two, three, or four actuators. Each curve represents a different subset size, with the solid curve representing subsets of a single actuator, the dashed curve representing subsets of two actuators, and so forth. The horizontal axis shows the rank among subsets, where the subsets are ranked by the amount of reduction they produce in the radiated power cost function. For a given subset size, the "optimal" subset has rank one and produces the largest reduction of radiated power cost function. With one actuator, the reduction ranges from 12 dB to less than 1 dB depending on which actuator is used. With multiple actuators reductions of over 70 dB are possible (assuming perfect sensing and actuation). In general, each additional actuator brings a substantial improvement on the order of 10 dB or more.

The risk involved in examining only the "optimal" subset, rather than examining a number of near-optimal subsets, is clearly demonstrated by the subsets of three actuators. The best subset (rank 1) appears to reduce the radiated power by 70 dB as compared to 40 dB for the second-best subset (rank 2). At first glance one might be enthusiastic about the prospect of a 70 dB reduction, but the large difference between the best subset and the second-best subset means that the performance is extremely sensitive to the actuator locations and/or measurement noise. If the actuators are not placed in *precisely* the correct locations, or if sensor measurement error is present, the performance will likely be closer to that of the second-best subset. (This issue is revisited in the next section when the actuator locations for the optimal subsets are discussed.)

In general, for a given subset size many near-optimal subsets are available that perform almost as well as the "optimal" subset. The designer therefore has some freedom in choosing actuator locations. This would be important in designing for broadband control; one would seek subsets

that are near-optimal at several or many frequencies, rather than subsets that are exactly optimal for one specific frequency.

5.2.2. Optimal locations for actuators

Figure 28 shows the actuator locations for the 15 best subsets of one, two, three, or four actuators at k_0a=0.59 with the uncontrolled cylinder response shown in Fig. 26. The horizontal axis shows where each actuator is located on the cylinder, with δ=0 being the left end closure and δ=1 being the right end closure. Each row on the chart represents a different subset, with actuator locations for that row appearing as squares in the appropriate locations. The reduction in radiated power cost function for each subset is shown in decibels at the left side. The best subset is shown in the top row, second-best in the second row, etc.

For the case of a single actuator, Fig. 28a, the optimal location corresponds to what a designer might anticipate from prior knowledge of the system dynamics. Judging from the distribution of surface velocity shown in Fig. 26, one might guess that the most effective location for a single control force would be near one of the "anti-nodes" at δ=0.1, 0.5, and 0.9. But optimizing for two or more actuators leads to configurations that are less intuitive. In [40], a single optimally located actuator generally retained its location when a second actuator was added and both locations were optimized. In the present example, there appears to be little relation between actuator locations for different subset sizes. With two actuators (Fig. 28b) the algorithm requests actuators not at the anti-nodes, but rather in pairs on either side of a nodal circle. In almost all cases, the paired forces have opposite signs so that they resemble line moments. With three actuators (Fig. 28c), the optimal actuator placement again calls for actuators near the anti-nodes.

Among subsets of three actuators, note that in comparing between the best subset and the second-best subset, the seemingly minor difference in actuator placement produces a 30 dB difference in the achievable performance. The controller performance can thus be very sensitive to actuator locations. Note also that the best subsets of four actuators (Fig. 28d) are all nearly identical to the best subset of three actuators, i.e., three actuators are placed at δ=0.08, 0.50, and 0.92 with a fourth actuator placed elsewhere on the structure. The fourth actuator has almost no

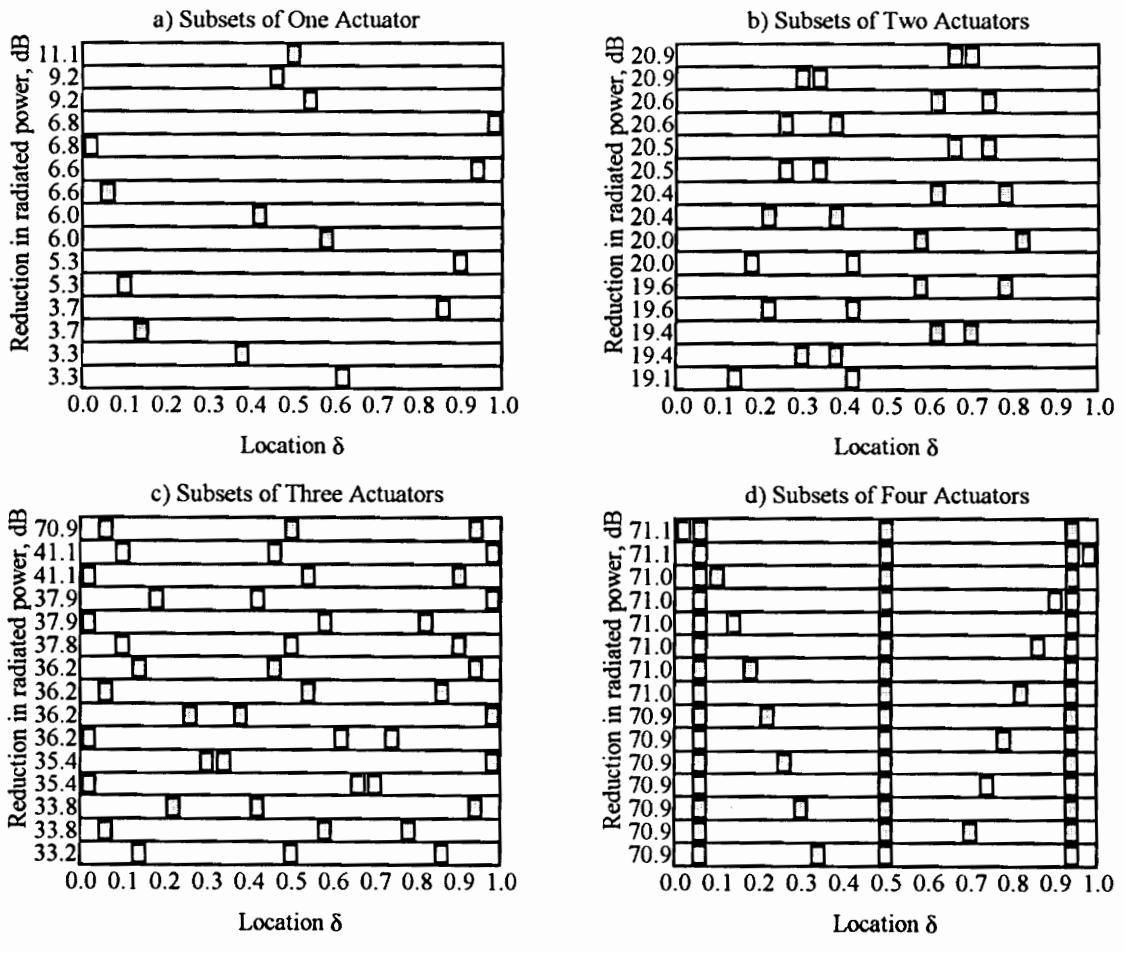

Figure 28: Best 15 actuator subsets at k_0a=0.59. Figures 28a, 28b, 28c and 28d represent subsets of one, two, three, and four actuators, respectively. Each row represents a different subset; radiated power for that subset is shown at left side of each plot. Actuator locations are shown at bottom of each plot.

influence on the system performance, and can therefore be located almost anywhere without affecting the results. However, as soon as one of the other three actuators is positioned away from the locations δ=0.08, 0.50, and 0.92, there is a significant decrease in system performance (see Fig. 27).

Less important than the specific results of this particular example is the fact that even this relatively simple structure provides results that are not immediately obvious. For more realistic and complicated systems, it may be difficult or impossible to guess the best actuator locations. The more complicated the physical system is, the more important it is to have formal, general methods for optimizing the system configuration.

5.3. Summary

This chapter discusses a numerical method for optimizing actuator locations in a feedforward active control system to provide the best system performance. The method involves using exhaustive search subset selection to choose the most effective subset of actuators from a set of candidates. Computational efficiency and a logical, general approach make this method ideal for analyzing complex systems with large numbers of actuators. Although presented here for an axisymmetric structure, the method can be extended to more general situations including three-dimensional structures because it is based purely on transfer functions.

The method is illustrated by optimizing a simple active structural-acoustic control system consisting of a finite-length, fluid-loaded cylindrical shell with flat, rigid, clamped end closures. The locations of control forces on the cylinder are optimized to reduce the total radiated power caused by a harmonic, distributed force. It is shown that while the optimal location of a single actuator can be predicted intuitively from knowledge of system dynamics, multi-force

configurations can be difficult to predict from intuition alone. It is risky to examine only one "optimal" actuator configuration because the system performance can be sensitive to actuator locations. Also, there may be many near-optimal actuator configurations that perform nearly as well as the "optimal" configuration.

Subset selection provides an efficient, versatile way to optimize actuator locations, but one must test the results for collinearity (numerical ill-conditioning). In the example problem investigated here, all subsets of five or more actuators were discarded because of collinearity. Judging from the small differences between subsets of three actuators and four actuators, the addition of a fifth actuator leaves the problem ill-conditioned for most combinations of actuators (though possibly not all combinations.)

CHAPTER 6: ALTERNATIVE COST FUNCTIONS

This chapter investigates the idea that the physical significance of the cost function depends on the type of sensor and on the weighting function. Numerical experiments, particularly in the form of varying the cost function, can help a designer better understand the relevant system dynamics, and ultimately lead to a more efficient control system design. Examining the problem in the context of its structural-acoustic behavior also helps illustrate how to use the regression diagnostics developed in Chapter 3 for modeling the effects of measurement noise. To this end, we define below the concept of a *response function*, i.e., a weighted sum of squares of the form of Eq. (25) (page 43) that measures some aspect of the dynamic response. A variety of different types of response functions may be considered, since both the quantity being measured and the weighting function may be chosen as needed. The designer can then experiment with an alternate controller formulation simply by using a different response function as the cost function. The underlying regression analysis does not depend on the nature of the cost function, although the quantity used in the cost function determines the type of error sensor that would be used.

In this chapter we present some alternative controller formulations for the clamped cylindrical shell. Measurement noise is modeled implicitly using the techniques of Chapter 3. Again, rather than assuming measurement noise levels for a particular type of sensor, the error sensor measurement noise variance is approximated from the data (see Section 3.1). We investigate three types of response functions, corresponding to three types of sensors: farfield radiated pressure sensors, surface vibration sensors, and wavenumber sensors. The pressure sensors would correspond to farfield microphones or hydrophones for the external radiation problem. The vibration sensors could be surface-mounted accelerometers or, equivalently, velocity sensors. (We

shall in this section use surface velocities, which happen to be the surface quantity most easily obtained from the computer program used to compute the transfer functions.) Structural wavenumber quantities are found by transforming the spatial-domain surface velocity data into the wavenumber domain using a Fourier transform. For a description of practical ramifications of wavenumber sensing in real time, see [52] and [51].

To illustrate the use of alternative controller formulations, we examine how the dynamic response (as characterized by various response functions as well as other data) changes depending on which response function is used as the cost function in the controller. Conceptually, this would correspond to having *observation sensors* that provide information about the structural-acoustic response, as shown in Fig. 29. The data measured by the observation sensors are separate from the data that are available to the controller, and are not used in computing the controller gains, but are used only to provide diagnostic information. The concept of separate *observation* sensors and *error* sensors is convenient for discussion purposes, and will be used throughout this Chapter.

Alternative cost functions are also examined in [12], which uses spatial-domain and wavenumber-domain cost functions for a spherical shell. Other studies include [2], [3], and [4].

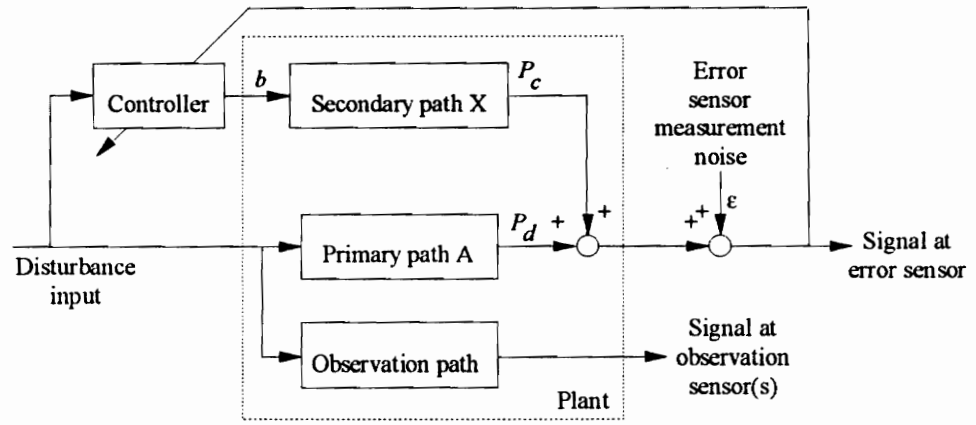

Figure 29: Block diagram showing SISO feedforward control system with observation sensors (not used by controller) and no measurement noise at detection sensor.

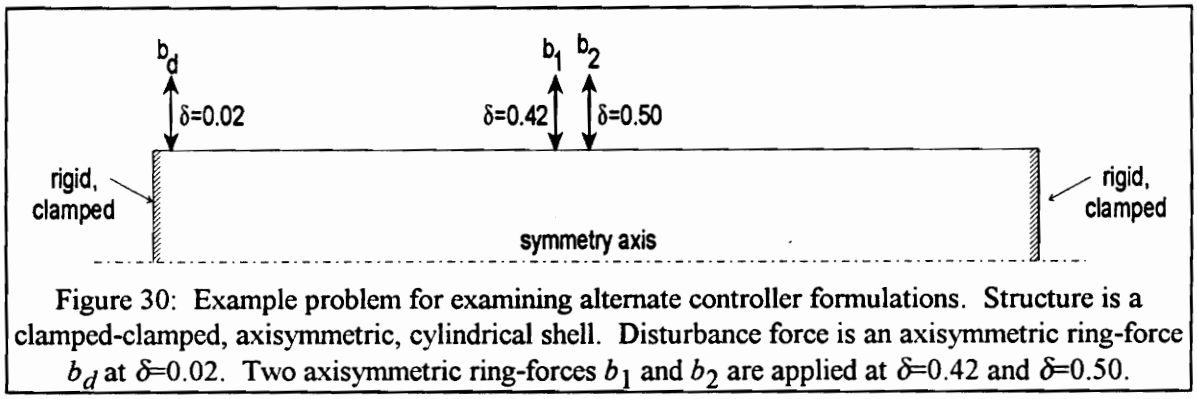

Figure 30: Example problem for examining alternate controller formulations. Structure is a clamped-clamped, axisymmetric, cylindrical shell. Disturbance force is an axisymmetric ring-force b_d at δ=0.02. Two axisymmetric ring-forces b_1 and b_2 are applied at δ=0.42 and δ=0.50.

For the discussion below we consider the cylindrical shell with disturbance and actuator inputs as shown in Fig. 30. The disturbance is an axisymmetric ring-force b_d acting normal to the shell surface at δ=0.02. The actuators are two axisymmetric ring-forces b_1 and b_2 acting normal to the shell surface at δ=0.42 and δ=0.50, respectively. The actuator locations are not necessarily optimal for this disturbance and frequency, but were chosen to illustrate certain properties of the alternative cost functions considered here. The sensors are assumed to be perfect, that is, we assume there is no error sensor measurement noise.

6.1. Response function definitions

The following equation presents the regression problem of Eq. (21) and Eq. (25) in a more general form that aids in discussing alternate controller formulations:

$$\min_{b_j} \chi^2 = \sum_{i=1}^{M} w_i(\boldsymbol{r}_i)\left|y_i(\boldsymbol{r}_i)+\hat{y}_i(\boldsymbol{r}_i)\right|^2 \quad \text{where} \quad \hat{y}_i(\boldsymbol{r}_i)=\sum_{j=1}^{N} b_j X_{ij}(\boldsymbol{r}_i). \tag{69}$$

Here we explicitly show the dependence on the error sensor location $\boldsymbol{r}_i$ in the cost function χ^2. In general $\boldsymbol{r}_i$ is a vector. However, for axisymmetric problems $\boldsymbol{r}_i$ specializes to either an angular coordinate θ, a surface location δ, or an axial wavenumber k_z. Recall that $X_{ij}(\boldsymbol{r}_i)$ is a transfer function between a unit input to actuator j and the response at $\boldsymbol{r}_i$. By combining different response

data and weighting coefficients, we can perform numerical experiments using a variety of different controller formulations. Some combinations produce response functions that cannot be readily interpreted as physically significant quantities. However, examining several different response functions can help one better understand the physics of the problem for a particular cost function. With each type of response data, i.e. each sensor type, we will consider various forms of weighting functions as described below.

All response functions used in this thesis have the same form as the cost function in Eq. (69) above. The response data come from observation sensors in the case of a response function, or error sensors in the case of the cost function. In some instances the "observation sensors" and "error sensors" are identical, and are simply referred to as "sensors", although the distinction is emphasized where needed.

6.1.1. Farfield pressure sensors

When the sensors are farfield pressure sensors such as hydrophones, a proper weighting function can produce a response function $\hat{\Pi}$ that is proportional to the total radiated power as in Chapter 2 and Chapter 3. Of course, in order to reasonably characterize the farfield over a range of frequencies, the sensors must be sufficient in number and distributed throughout the farfield. The number of sensors required depends on the radiation characteristics of the specific structure under consideration, but the number increases as frequency increases and the farfield becomes more spatially complicated. If necessary, the required number of farfield sensors can be found via a convergence study as described in Fig. 6 (page 24). For the clamped cylindrical shell considered here, a sensor spacing of $\Delta\theta=7.5°$ is sufficient in the frequency range of interest, thus the number of sensors is $M=25$.

The weighting function coefficients in this case represent the relative areas associated with each of the sensors. If the structural geometry and sensor configuration are axisymmetric, and the sensors are placed at a uniform distance from the structure, then the sensor location $\boldsymbol{r}_i$ may be replaced by the angular coordinate θ. With sensors uniformly spaced at intervals of Δθ between θ=0° and θ=180°, the desired weighting function is

$$w_i = \frac{a_i}{a_{\max}} \quad \text{where} \quad a_i = \begin{cases} 1-\cos\Delta\theta/2, & i=1 \text{ or } i=M \\ \cos(\theta_i-\Delta\theta/2)-\cos(\theta_i+\Delta\theta\setminus 2), & 1<i<M \end{cases} . \tag{70}$$

and $a_{\max}$ is the maximum value of a_i, the area associated with the i^{th} sensor. (One could select sensor spacing for uniform relative areas rather uniform spacing, in which case w_j=1.) Since the response function with this sensor/weighting combination is proportional to the radiated power, it is referred to as the *radiated power response function* and denoted by the symbol $\hat{\Pi}$.

By using a different weighting function, we can restrict our attention to a conical sector of the radiated field rather than the entire farfield as described above. (Results for this type of approach are not given in the present work, but the approach is included for completeness.) For example, to minimize the acoustic power radiated into a cone θ≤30° we could specify

$$w_i = \begin{cases} a_i/a_{\max} & \theta_i \le 30° \\ 0 & \theta_i > 30° \end{cases} . \tag{71}$$

By including zeros in the weighting function, we obviously perform some unnecessary computations that could be avoided if we formed a new regression without the sensors corresponding to the zeros. However, in the context of performing a series of numerical experiments to improve our understanding of underlying physics, we may find it more convenient to simply change the weighting function. Furthermore, we can reduce the impact of the unnecessary calculations on computation time by designing the associated computer programs to take advantage of sparsity (as in Appendix B.)

Other types of weighting functions are also possible. For example, we could use zeros for all weighting coefficients except one, thereby using only a single farfield sensor in the cost function. We could weight some regions of the farfield more heavily based on practical requirements of the particular application. Many variations are possible.

6.1.2. Surface velocity sensors

We can also investigate the use of structure-based sensors such as accelerometers, which may have advantages in applications for which farfield sensors are impractical. For a numerical simulation, we can use the complex surface velocity to simulate a structurally mounted vibration sensor. Suppose we know the radial surface velocity u at a number of locations $r_j=\delta_j$ over the vibrating portion of the structure. For the cylindrical shell example, the δ_j are the grid circle locations of the axisymmetric finite-element model used to predict the dynamic response. If the structural mass per unit distance along the shell is $m(\delta)$, we can write $\mathbf{T}$, the portion of the total kinetic energy associated with radial motion, as

$$\mathbf{T} = \frac{1}{2}\int_0^1 m(\delta)\left|u(\delta)\right|^2 d\delta . \tag{72}$$

Given a FEM model of the structure, the radial surface velocity u is not a continuous function of space, but instead a discrete sequence of radial velocities at the M sensor locations. (For the cylindrical shell used in the present work, M=50.) By summing the squared radial velocity magnitudes [53] we can form a *kinetic energy response function* $\hat{\mathbf{T}}$ which approximates the radial kinetic energy:

$$\hat{\mathbf{T}} = \frac{1}{2}\sum_{i=1}^{M} u_i m_i u_i^{*} , \tag{73}$$

where m_i is the structural mass associated with location r_i. This expression is subject to the same sensor distribution considerations as the radiated pressure cost function from the previous section, namely, the sensors must be sufficient in number and adequately distributed to reasonably characterize the vibration patterns at all frequencies of interest.

Comparing Eq. (73) with weighted cost function defined in Eq. (25), we see that each weighting coefficient w_i must be proportional to m_i. For the clamped cylinder under consideration, the mass distribution and sensor spacing are uniform over the entire vibrating surface, so

$$w_i = \frac{m_i}{m_{\max}} \quad \text{where} \quad m_i = \begin{cases} 1/2, & i = 1 \text{ or } i = M \\ 1, & 1 < i < M \end{cases} . \tag{74}$$

Noting that the surface velocity is zero for i=1 and i=M because of the clamped boundary conditions, we can simply set the weighting coefficients to unity to obtain a response function that is proportional to the estimated total kinetic energy.

Had we used acceleration transfer functions rather than velocity transfer functions, the response function would still be proportional to the estimated total kinetic energy. Because the system is in harmonic steady-state vibration, the acceleration differs from the velocity by a factor of $\omega\sqrt{-1}$, where ω is the circular frequency. Thus the response function $\hat{\mathbf{T}}$ with acceleration sensors would differ from the estimated total kinetic energy by a factor of $-\omega^2$.

As with the previous section, other response functions may be examined by including some zeros among the weighting coefficients. For example, rather than characterizing the radial kinetic energy over the entire structure, we could restrict our attention to a portion of the structure by using a weighting function similar to Eq. (71).

6.1.3. Wavenumber sensors

The concept of a wavenumber-domain controller is attractive for two reasons. In some situations the use of farfield error sensors is inconvenient or impossible, which suggests the use of structurally mounted sensors. However, it is neither necessary nor desirable to control the entire vibrational field on the structure, since some components of vibration do not propagate energy to the farfield. A wavenumber-domain controller could provide both an effective use of structure-borne sensors and a way to expend control effort only on propagating vibrations [5,12,51,52]. Generally speaking, only supersonic components of vibration propagate to the farfield [27,56]. If we could restrict our active control efforts to only the supersonic portion of the vibration field, we could effectively apply a high-pass filter (with cut-off wavenumber k_o) to the closed-loop structural wavelength and thus reduce acoustic radiation.

Assuming we have an array of surface radial velocity sensors as in the previous section, we can perform a wavenumber transform on these data to simulate wavenumber sensors. Maillard and Fuller [51,52] discuss a practical method for using accelerometers to obtain wavenumber sensor information in both the time domain and the frequency domain. (Their frequency-domain methods are similar to the method used here, although the measured quantity is acceleration rather than velocity.) For a three-dimensional cylindrical shell, the wavenumber transform is found by taking the spatial Fourier transform of the normal surface velocity $\dot{r}(z,\theta)$:

$$\tilde{U}(k_z,k_\theta)=\int_0^{2\pi}\int_{-\infty}^{\infty} u(z,\theta)e^{-k_z z\sqrt{-1}}e^{-k_\theta \theta\sqrt{-1}}dzd\theta. \tag{75}$$

For the axisymmetric cylindrical shell considered here, we consider only the zero harmonic in θ. Given surface velocity data for a region of the structure $0\le\delta\le L$, we obtain the wavenumber transform $\tilde{\mathbf{U}}(k_{zi})$ at wavenumber k_{zi} as

$$\tilde{\mathbf{U}}(k_{zi}) = \int_0^L \boldsymbol{u}(\delta) e^{k_z \delta \sqrt{-1}} d\delta . \tag{76}$$

Since we know the velocity only at the sensor locations, we must compute the wavenumber transfer function as a discrete Fourier transform. A standard FFT can be used if the sensors are evenly spaced along the structure. FFT cannot be used when the sensors are not evenly spaced, as may be the case when finite-element data are used to represent sensors. The present work uses a direct-integration approach as in [54]. The velocity data could also be interpolated into an evenly spaced grid allowing a standard FFT to be used, and other methods are also available [39].

We can compute wavenumber transforms for a spectrum of wavenumbers $-k_{\max} \le k_j \le k_{\max}$, where $k_{\max}$ is determined based on the spacing between velocity sensors. If the M sensors are evenly spaced between δ=0 and δ=1, and are separated by a distance $\Delta\delta$, then the maximum wavenumber for which the Fourier transform can be computed is

$$k_{\max} = \frac{\pi M}{2} \tag{77}$$

which corresponds to the largest wavenumber that may be computed without spatial aliasing. In reality the structural wavenumber spectrum will be band limited by the physical properties of the shell, which prohibits it from exhibiting extremely short-wavelength responses [55]. For a finite-element model of the structure, we must use a grid-circle spacing $\Delta\delta$ small enough that $-k_{\max} \le k_j \le k_{\max}$ completely characterizes the velocity distribution, i.e., we must avoid aliasing at the wavenumbers of interest. Once we compute the wavenumber transfer function at M wavenumbers, we can write the wavenumber response function for weighting w_i as

$$\sum_{i=1}^{M} \tilde{\mathbf{U}}(k_i) w_i \tilde{\mathbf{U}}(k_i)^* . \tag{78}$$

We can use the entire computed wavenumber spectrum $-k_{\max} \le k_j \le k_{\max}$ and set the weighting coefficients to unity to form the *wavenumber response function* Λ:

$$\Lambda = \sum_{i=1}^{M} \tilde{\mathbf{U}}(k_i) w_i \tilde{\mathbf{U}}(k_i)^* \quad \text{where} \quad w_i = 1 \text{ for } i = 1,\ldots M. \tag{79}$$

If the wavenumber spectrum contains the *complete* bandwidth present in the velocity data, then minimizing Λ would actually minimize the entire vibration pattern and would therefore have the same effect as minimizing the kinetic energy response function $\hat{\mathbf{T}}$.

Another possible controller formulation would involve controlling only the supersonic portion of the wavenumber spectrum, as outlined above. To characterize only supersonic wavenumbers, we would define a *supersonic wavenumber response function* Λ_{sup} as

$$\Lambda_{sup} = \sum_{i=1}^{M} \tilde{\mathbf{U}}(k_i) w_i \tilde{\mathbf{U}}(k_i)^* \quad \text{where} \quad w_i = \begin{cases} 1, & |k_i| \le k_o \\ 0, & |k_i| > k_o \end{cases}. \tag{80}$$

Since there is likely to be a strong relationship between the computed supersonic wavenumber spectrum and the farfield radiation, minimizing Λ_{sup} should have roughly the same effect as minimizing the estimated radiated power.

A third interesting wavenumber-domain controller would be to examine one specific structural wavenumber component. Consider an infinitely long cylindrical shell in steady-state vibration with a velocity distribution whose wavenumber spectrum contains only one nonzero component, say, $-k_z$, corresponding to a positive-traveling structural wave. As outlined in Fahy [56] or Junger and Feit [27], the farfield would consist of radiation towards a single direction, i.e., a conical wave front propagating at a coincidence angle θ_Z, where

$$\cos\theta_z = -k_z / k_o. \tag{81}$$

We could characterize the farfield radiation at the coincidence angle by examining the structural wavenumber component k_z. If we consider instead a velocity distribution containing a range of wavenumber components, we can still characterize farfield radiation at a desired angle by defining

a *narrow-band wavenumber response function* to include only the associated wavenumber component:

$$\Lambda_{k_z} = \sum_{i=1}^{M} \tilde{\mathbf{U}}(k_i) w_i \tilde{\mathbf{U}}(k_i)^* \quad \text{where} \quad w_i = \begin{cases} 1, & k_i = k_z \\ 0, & \text{elsewhere} \end{cases} \tag{82}$$

6.2. Numerical results

Having defined in the preceding sections a variety of response functions that can be used in alternative controller formulations, we next examine the clamped cylindrical shell and see how its closed-loop response changes depending on which response function is minimized by the controller. When the radiated power response function is used, the number of farfield pressure error sensors is M=25. When the radial kinetic energy response function or the wavenumber response function is used, the number of surface velocity error sensors is M=50.

The reader is reminded that error sensor measurement noise is modeled using the same technique as in Chapter 3. Rather than assuming a specific value for the variance of the error sensor measurement noise, we use the data to estimate the variance. See Chapter 3.

6.2.1. Minimizing radiated power response function at a single frequency

Figure 31 shows a typical set of dynamic response data for the system configuration of Fig. 30. The frequency is $k_o a$=0.80, and the function being minimized is the radiated power response function $\hat{\Pi}$. The regression predicts that the radiated power cost function is reduced by 10.5 dB, and neither the F-test nor the scaled condition number indicate problems with the regression. The optimal control forces for this frequency are b_1=(-0.027, -0.014) and b_2=(-0.856, -0.005), but the t-test indicates that b_1 is not statistically significant and that only b_2 contributes to the regression at this frequency. (Performing a separate regression without b_1 confirms this, as the residual

response is not affected by the absence of b_1.) Figure 31 contains three plots which are discussed in the following paragraphs to present a more complete picture of the controller performance. All three plots show the response both without the controller (solid lines) and with the controller (broken lines). Figure 31a shows the pressure directivity, i.e., farfield pressure magnitude is shown on the vertical axis vs. farfield angle θ on the horizontal axis. Figure 31b shows the distribution of normal surface velocity, with velocity magnitude on the vertical axis and location δ along the cylinder wall on the horizontal axis. Figure 31c shows the magnitude of the wavenumber transform used to compute the wavenumber cost function vs. the nondimensional axial wavenumber $k_z a$, with the sonic lines $-k_o a$ and $+k_o a$ shown as vertical dotted lines.

To better interpret the structural response, we must consider the relation between the disturbance frequency and the *ring frequency* ω_r, where the non-dimensional frequency Ω is

$$\Omega = \frac{\omega a}{c_l} = \frac{\omega}{\omega_r}, \text{ and } \quad c_l = \sqrt{\frac{E}{\rho(1-\nu^2)}}. \tag{83}$$

The ring frequency is the frequency for which a circumferential wave has wavelength equal to the shell circumference. For the clamped cylinder of Fig. 30, which is a steel shell submerged in water, the disturbance frequency range $k_o a<1.0$ corresponds to $\Omega<0.03$ and is thus well below the ring frequency.

The predominant mechanism for radiation in the frequency range of interest is the coupling of radial and axial displacements on the shell. The shell curvature causes coupling between axial displacement and radial displacement in a manner similar to Poisson's effect. Any excitation that is not purely torsional will excite both axial (compressional) motion and radial (flexural) motion on the shell surface.

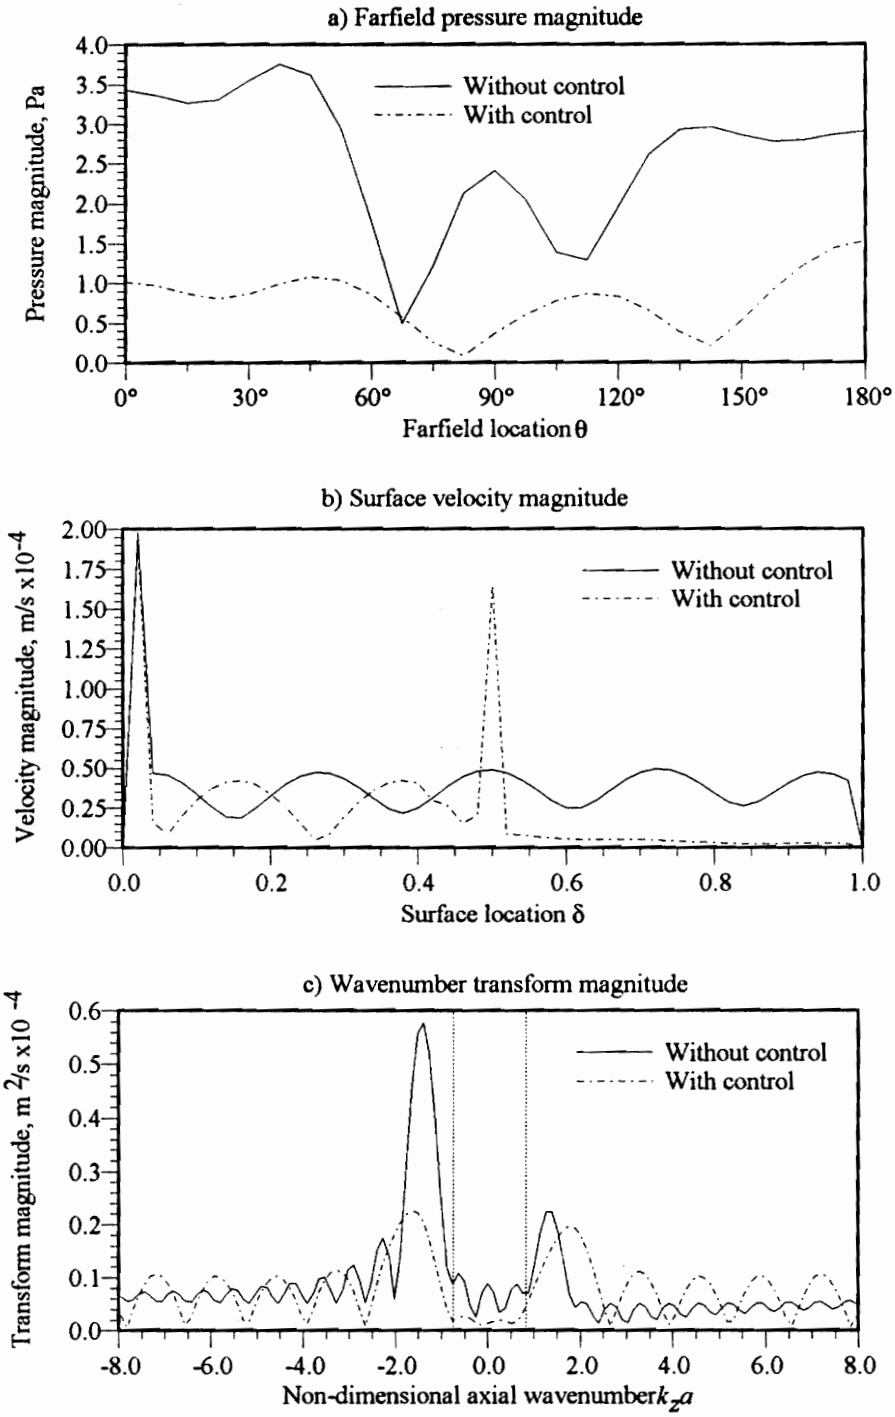

Figure 31: Residual farfield pressure, surface velocity, and wavenumber spectrum at $k_o a$=0.80 when forces are optimized to reduce radiated power response function. Disturbance and control forces are as shown in shown in Fig. 30. Solid curves represent response without control, broken curves represent response with control. Shown are farfield pressure magnitude (Fig. 31a), surface velocity magnitude (Fig. 31b) and wavenumber transform magnitude (Fig. 31c).

Figure 32 shows a typical dispersion curve for n=0 motion of an unstiffened, infinite cylindrical shell with fluid loading [58]. The roots of the shell equation of motion are plotted as a function of axial structural wavenumber $k_z a$ on the horizontal axis and frequency Ω on the vertical axis. The sonic line is shown for reference as a dashed line. The compressional wave is non-dispersive and supersonic, and therefore appears as a straight line to the left of the sonic line. The flexural wave, shown as the solid curve, is dispersive and subsonic. At the low frequencies considered here, the line representing the compressional wave should be straight and supersonic, and the curve representing the flexural wave should be slightly subsonic and curving away from the sonic line. The flexural wave, however, should be excited much more strongly than the compressional wave by the radially-oriented disturbance and control forces used in the example problem. Thus we should expect that, although the wavenumber spectrum will contain both compressional and flexural waves, the flexural waves will likely be of much larger magnitude.

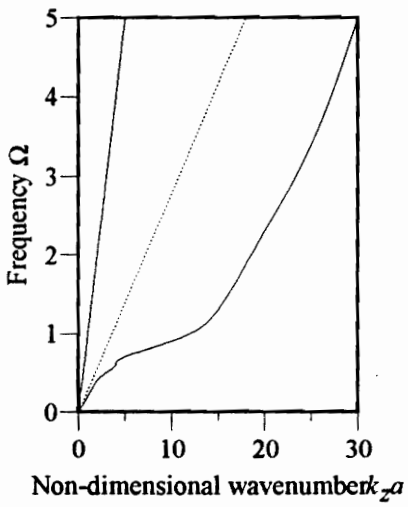

Figure 32: Typical dispersion curve for n=0 modes of an infinite, thin-walled, fluid-loaded, cylindrical shell (from Scott [58]). Solutions to shell equation of motion are plotted vs. non-dimensional structural wavenumber $k_z a$ and non-dimensional frequency Ω. Dashed line is sonic line, solid line is compressional wave, solid curve is flexural wave.

Let us first examine the primary response, that is, the response due to the disturbance only without the controller (Fig's 31a - c, solid lines). Looking at the surface velocity distribution (Fig. 31b), we see a large velocity magnitude at the disturbance drivepoint location δ=0.02. Away from the disturbance drivepoint we see a roughly sinusoidal variation of wavelength λ=0.44L, with five half-wavelengths (m=5) along the length of the cylinder. Notice also that the average velocity is slightly higher on the right half of the cylinder than on the left half, indicating an m=2 component of vibration that adds to the vibration on the left half and subtracts from the vibration on the right half. Though small, this m=2 component is important and appears in the discussion of residual response later in this section.

Turning next to the wavenumber spectrum (Fig. 31c), we see large spectral components at $k_z a$=-1.39 and $k_z a$=+1.39, which correspond to right-traveling and left-traveling flexural waves of the m=5 component. Since the acoustic wavenumber is $k_o a$=$\pm$0.8, the flexural wave is slightly subsonic. However, because the vibration appears over a cylinder of finite length rather than an infinitely long cylinder, the m=5 contribution to the wavenumber spectrum appears as a wide peak rather than as a delta function of infinitesimal width. This finite-length effect is similar to the effect of using a rectangular window in computing a Fourier transform: peaks that would have infinitesimal width are distorted by the finite extent of the window and become peaks of finite width. The m=5 peak becomes wide enough that it overlaps the edge of the supersonic region, and contributes to farfield radiation even though the center of the peak is subsonic.

Note that the wavenumber spectrum is generally higher for negative wavenumbers (right-traveling waves) than for positive wavenumbers (left-traveling waves). The disturbance drivepoint is near the left end closure. Flexural waves propagate from left to right, losing energy to radiation, and are then reflected by the right end closure to travel from right to left at reduced

magnitude. The loss of energy to acoustic radiation is reason for the disparity between the spectral magnitudes of right- and left-traveling waves.

We also note the presence of three smaller peaks within the supersonic region at $k_z a$=-0.63, 0.0, and 0.63, as well as numerous peaks in the regions $k_z a$<-2.0 and $k_z a$>2.0. These components are present because of the interaction of flexural waves with the boundary conditions presented by the clamped end closures. The interactions between the subsonic bending waves and the end closure discontinuities scatter energy from the bending wave into both subsonic and supersonic nearfield components, and it is the supersonic components that provide most of the acoustic radiation present. (Some high-wavenumber components may also be due to aliasing of the wavenumber transform, which could occur because of the high velocity gradient near the drivepoint. Aliasing could be reduced by using a finer mesh of grid-points in the finite-element structural model and the Fourier transform.) Since the relation between structural wavelength and structural wavenumber is $\lambda_z=2\pi/k_z$, we can easily show that the $k_z a$=+/- 0.63 components in the wavenumber spectrum correspond to the m=2 components in the vibration, and $k_z a$=0.0 corresponds to m=1. Recalling Eq. (81), these three peaks should correspond to peaks in the farfield pressure at angles of 38°, 0°, and 142°. Looking at the pressure directivity (Fig. 31a), we see peaks in the radiated pressure at exactly the angles predicted by Eq. (81). In addition, we see large components at angles of 0° and 180°, which correspond to the portions of the positive-traveling and negative-traveling flexural waves that radiate because of finite-length effects.

Now consider the residual response, that is, the closed-loop response with the controller acting to minimize the radiated power cost function (Figs. 31a - c, broken lines.) In Fig. 31b we see at δ=0.50 the large drivepoint velocity of control input b_2. (Recall that control input b_1 has zero magnitude at this particular frequency.) The presence of the control input does not significantly

alter the response at the disturbance drivepoint, δ=0.02. On the left half of the cylinder, between the disturbance drivepoint and the control drivepoint, the velocity peaks do not experience a drastic reduction in magnitude. To the right of the control drivepoint, however, the controller reduces the surface velocity by an order of magnitude or more. Thus the flexural wave is allowed to remain on at least part of the structure, while the control effort appears to be expended on removing the longer-wavelength components. One might expect an associated decrease in the total kinetic energy, but the decreased velocity on the right half of the cylinder is partially offset by the increased velocity at the control drivepoint. The total kinetic energy drops by only 1.1 dB.

The residual wavenumber spectrum (Fig. 31c) shows a substantial restructuring of the wavenumber spectrum. The wavenumber transform is smaller in the supersonic region, but somewhat larger in the subsonic region indicating spillover of energy into shorter wavelengths. The peaks corresponding to the flexural waves are diminished in magnitude, but are also moved to slightly higher wavenumbers. The residual pressure directivity (Fig. 31a) shows both a large decrease in pressure magnitude and a shifting of the radiation peaks to different angles. The two residual pressure peaks at θ=45° and θ=115° correspond to the peaks in the wavenumber spectrum at $k_z a$=-0.57 and $k_z a$=0.34, respectively.

6.2.2. Examining primary response over a range of frequencies

Controller performance for a given disturbance and actuator configuration varies with frequency, especially when only one or two actuators are used. For example, a controller with two actuators will perform poorly at a frequency for which both actuators are situated on nodes of the vibration pattern. Therefore, it is important to examine the primary and residual responses for a range of frequencies. In this section we examine the primary response, and in following sections we examine the residual responses associated with different controller formulations. Because of

the large volume of data associated with presenting the results for each frequency separately, we shall instead use the more compact format provided by color-scaled image plots as described below. Figure 33 characterizes the primary response in terms of the farfield pressure, surface velocity, and velocity autospectrum. The figure contains four color-scaled plots in which the vertical axis represents the nondimensional frequency $k_o a$. Each of the four plots is discussed in the following paragraphs.

Figure 33a shows the farfield pressure directivity as a function of frequency. The color scale is proportional to farfield pressure magnitude, with black/blue representing the lowest pressures and yellow/white representing the highest pressures. The horizontal axis represents angular location in the farfield, given as the angle θ. If one takes a horizontal "slice" at a frequency of $k_o a$=0.80, for example, one obtains the single-frequency disturbance response shown in Fig. 31a. In general the radiation levels increase as frequency increases, seen in the progressively lighter colors at the top of the plot. At most frequencies the farfield pressure exhibits one or more peaks between 0° and 180°, and at almost all frequencies there is substantial radiation at 0° and 180°. The plot is not symmetric about θ=90° because the disturbance force is placed near one end of the structure.

Figure 33b shows the radial surface velocity distribution as a function of frequency. The color scale is proportional to surface velocity magnitude and the horizontal axis represents location on the cylinder. A horizontal slice from this plot at $k_o a$=0.8 corresponds to the single-frequency disturbance response in Fig. 31b. The most obvious feature of Fig. 33b is that the velocity magnitude at the drivepoint location (δ=.02) is much larger than at other locations on the shell. The difference between the radial velocity magnitude at the drivepoint and away from the drivepoint is much larger at lower frequencies; at frequencies above $k_o a$=0.8, the difference is

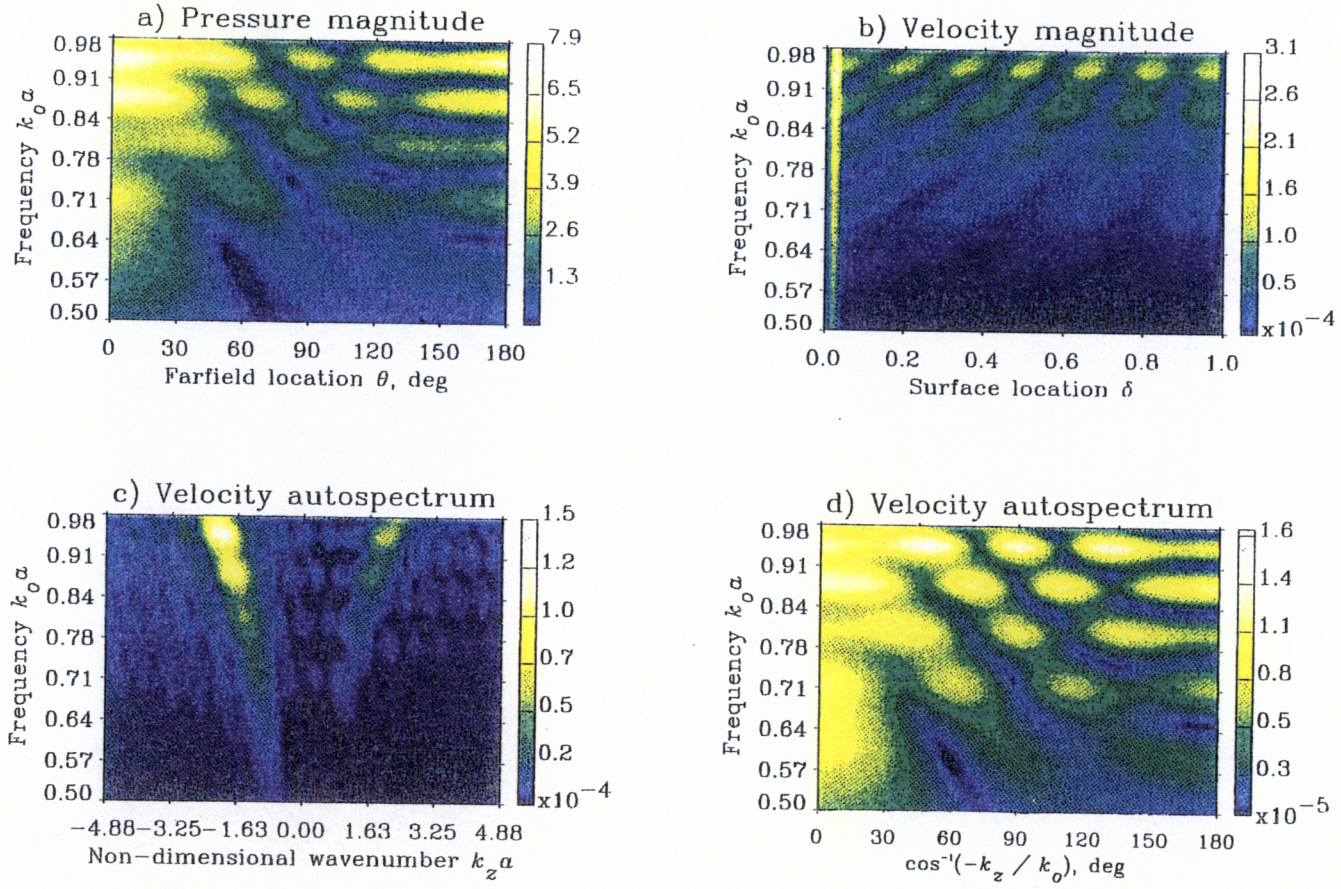

Figure 33: Primary response due to disturbance shown in Fig. 30. Shown are farfield pressure magnitude vs. frequency (Fig. 33a), surface velocity magnitude vs. frequency (Fig. 33b), and velocity autospectrum vs. frequency (Fig. 33c and 33d). Color corresponds to magnitude of quantity being plotted; see color scale at right of each plot.

roughly a factor of two or three. Away from the drivepoint, the velocity displays a smooth sinusoidal variation, with the number of half-wavelengths m ranging up to roughly seven at k_0a=0.95. At certain frequencies, e.g., k_0a=0.88 and k_0a=0.94, the velocity away from the drivepoint becomes relatively larger than at surrounding frequencies in what resembles a structural resonance. These are the frequencies for which the flexural wavelength coincides with a multiple of the cylinder length L.

Figure 33c shows the wavenumber spectrum as a function of frequency. The color scale is proportional to the velocity autospectrum (wavenumber transform), and the horizontal axis represents the axial wavenumber k_z nondimensionalized by the cylinder radius a. The flexural wave is clearly visible as a pair of curves that form the most easily visible feature of the plot. The compressional wave is not visible because its magnitude is too small compared to other components, i.e., it is not excited by the radially-directed excitation. Also visible are the low-magnitude supersonic wavenumber components caused by interaction between the subsonic flexural wave and the end closures, as well as higher-order wavenumber components. Note that the supersonic region is a V-shaped region extending between $k_za=\pm 0.5$ at the bottom of the plot and $k_za=\pm 0.98$ at the top of the plot.

Figure 33d presents the same wavenumber information as Fig. 33c, but the information is displayed using a novel format that proves useful in relating surface velocities to farfield pressures. Rather than normalizing the structural wavenumber to the cylinder radius as in Fig. 33c, the structural wavenumber is normalized to the acoustic wavenumber k_0 and plotted between -1 and +1. This corresponds to the V-shaped supersonic region inside the sonic line of Fig. 33c, except that it is expanded to fill a rectangular region. There are two benefits to this type of plot. First, the supersonic contributions are not overshadowed by the subsonic contributions, which are

of much larger magnitude but do not radiate to the farfield. Second, the relation between supersonic wavenumbers and farfield pressures becomes immediately obvious by comparing Fig. 33d to Fig. 33a. For every peak in the wavenumber spectrum there is a corresponding peak in the farfield pressure as predicted by Eq. 81. The relation is difficult to see in 33c because the corresponding plot regions have different shapes, but it is quite clear in 33d.

Figure 34(a-c) show the magnitudes of three response functions as frequency varies between $k_o a$=0.5 and $k_o a$=1.0: the radiated power response function $\hat{\Pi}$, the kinetic energy response function $\hat{\mathbf{T}}$, and the supersonic wavenumber response function Λ_{sup}. Only the primary response is shown, that is, the response without control. Referring to Fig. 34a, we see that the radiated power response function increases with increasing frequency, and displays resonance peaks at $k_o a$=0.59, 0.71, 0.80, 0.88, and 0.95. (We show below that these are the resonances of the m=3,4,5,6, and 7 modes, respectively.) The supersonic wavenumber response function (Fig. 34c) is almost identical to the radiated power response function except for a shift of scale, which is not surprising given the good agreement between the supersonic wavenumber spectrum and the farfield pressure observed in Fig. 33.

The kinetic energy response function (Fig. 34b) also displays resonance peaks, but they are less pronounced than the peaks in the radiated power response function or the supersonic wavenumber response function. This is because the kinetic energy response function includes the velocity at the disturbance drivepoint. Throughout the frequency range, and especially at low frequencies, the velocity at the disturbance drivepoint is much larger than elsewhere. The disturbance drivepoint velocity thus dominates the kinetic energy response function so that fluctuations due to structural resonances appear less pronounced than the fluctuations seen in other response functions. The

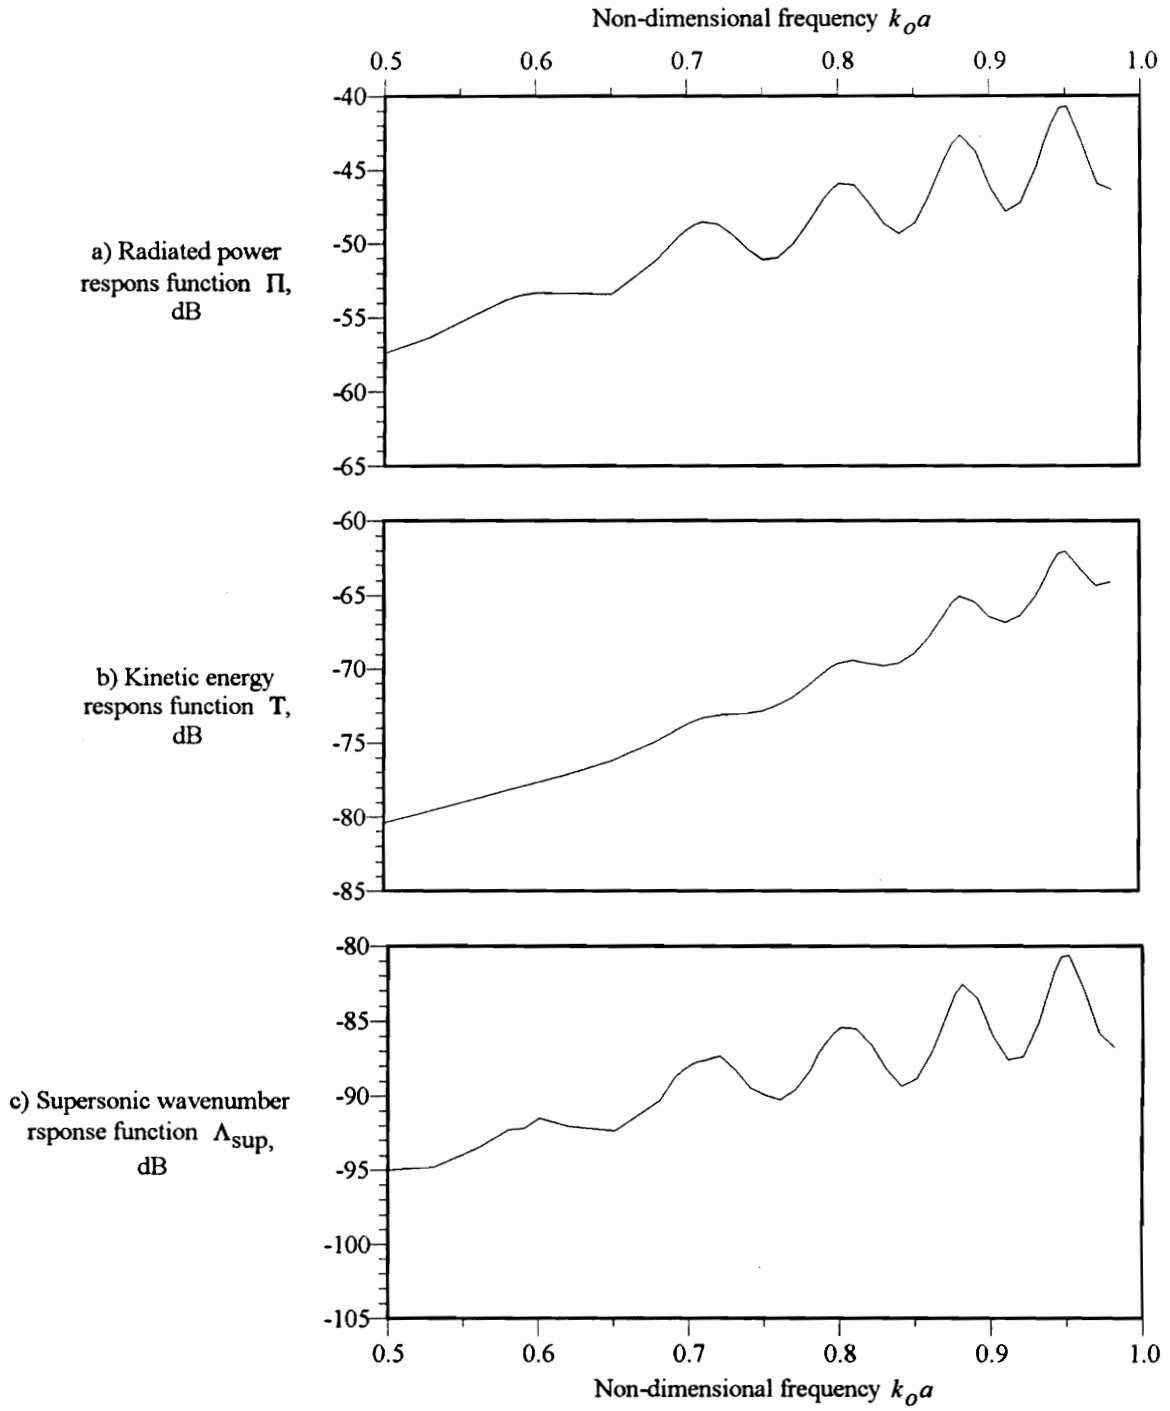

Figure 34: Response functions due to disturbance force shown in Fig. 30. Shown are radiated power response function vs. frequency (Fig. 34a), kinetic energy response function vs. frequency (Fig. 34b), and supersonic wavenumber response function vs. frequency (Fig. 34c).

inclusion of the disturbance and control drivepoints in the kinetic energy response function also causes controller problems, as discussed later.

The next sections use the same presentation formats used in Fig. 34 and Fig. 33 to examine the residual responses that remain after applying several different types of controller formulations.

6.2.3. Controller performance and residual response using radiated power cost function

In this section we examine the effect of minimizing the radiated power response function $\hat{\Pi}$ over a range of frequencies; the results are described by Figures 35 and 36. As in the previous section, each frequency represents a different regression with a different set of controller gains and a different effect on the response. For each regression the F-test and t-test results were examined along with the scaled condition number; the results of these diagnostic tests were satisfactory except where noted in the discussion, and are therefore not included below.

Figure 35(a-c) show the magnitudes of the three response functions ($\hat{\Pi}$, $\hat{\mathbf{T}}$, and Λ_{sup}, respectively) as frequency varies between k_0a=0.5 and k_0a=1.0. The solid curves represent the value of the function without control (primary response), and the broken curves indicate the value of the function with control (residual response). In each case the quantity being minimized is the radiated power response function $\hat{\Pi}$. The two control forces are able to reduce the radiated power response function by 4-12 dB over the entire frequency range, with the largest reductions appearing at higher frequencies. Not surprisingly, the residual supersonic wavenumber response function (Fig. 35c) is nearly identical to the residual radiated power response function; thus the wavenumber sensors could serve as an effective replacement for the farfield pressure sensors (assuming they could resolve the entire supersonic portion of the wavenumber spectrum.)

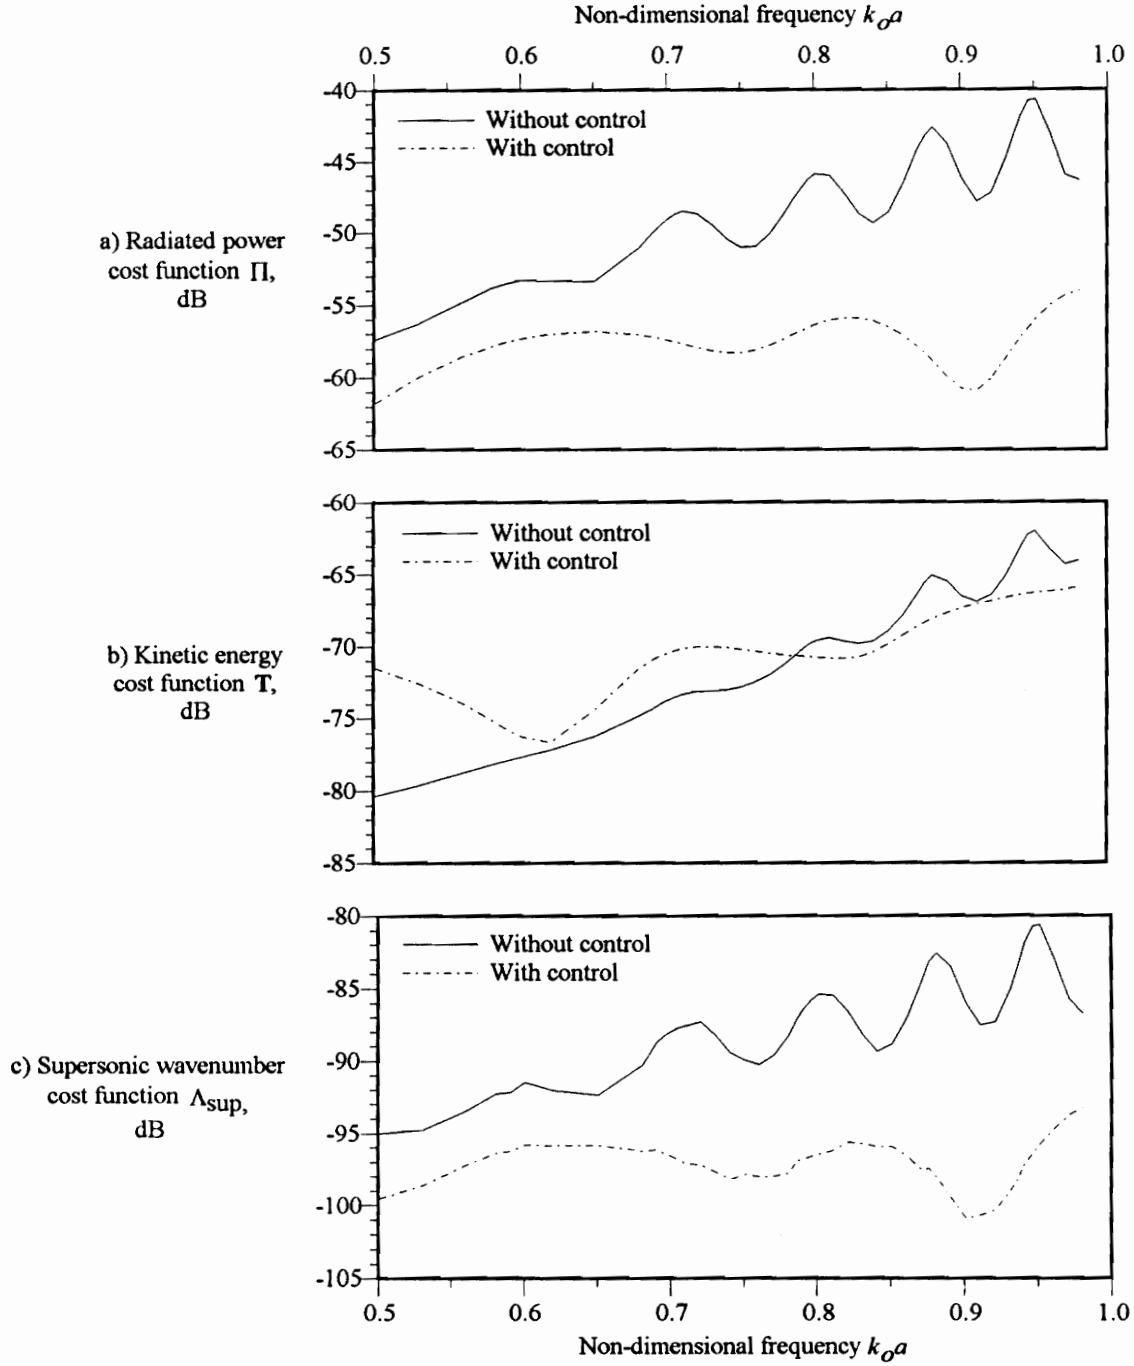

Figure 35: Residual response functions when control minimizes the radiated power response function. For each frequency, a separate regression is performed. Solid curve represents response function without control, broken line represents response function with control. Shown are radiated power response function $\hat{\Pi}$ (Fig. 35a), kinetic energy response function $\hat{\mathbf{T}}$ (Fig. 35b) and supersonic wavenumber response function Λ_{SUP} (Fig. 35c).

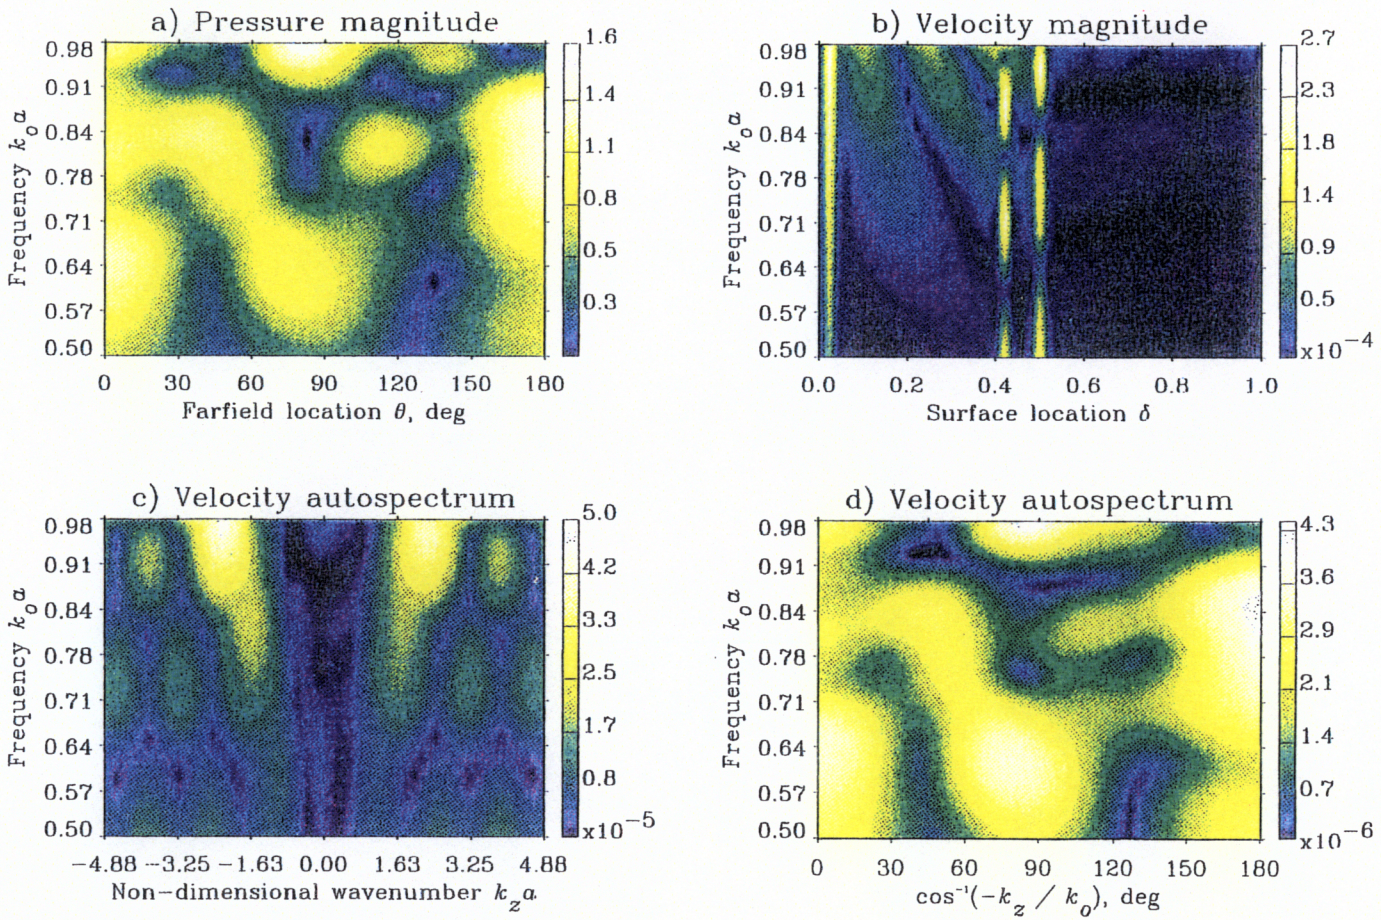

Figure 36: Response after controller minimizes radiated power response function. Shown are farfield pressure magnitude vs. frequency (Fig. 36a), surface velocity magnitude vs. frequency (Fig. 36b), and velocity autospectrum vs. frequency (Fig. 36c). Color corresponds to magnitude of quantity being plotted; see color scale at right of each plot.

The reduction in estimated kinetic energy (Fig. 35b) is rarely more than 2 dB, and at most frequencies the estimated kinetic energy actually increases even though the estimated radiated power decreases. The kinetic energy response function includes the velocities at all points, including the control drivepoints. When the control inputs are large, the control drivepoint velocities are large and their contributions to the kinetic energy response function are large. Even though the overall vibration levels may decrease over most of the structure at a given frequency, the kinetic energy response function can increase if the control drivepoint velocities are large.

Figure 36 shows the residual response as a function of frequency for the case where we minimize the radiated power response function. Comparing these plots with Fig. 33 helps illustrate the mechanisms used by the controller to reduce the cost function. The results discussed in the previous section for k_0a=0.80 (see Fig. 33) are typical of the entire frequency range, so only a brief discussion is included here. Figure 36a shows that not only is the pressure magnitude reduced by a factor of roughly four on average, but the locations of pressure magnitude peaks have also changed substantially. In other words, the controlled structure essentially has new eigenvalues as discussed by Burdisso and Fuller in [57]. In Fig. 36b, the velocity to the left of the control forces is not reduced in magnitude, but the velocity to the right of the control forces is reduced almost to zero. The control forces effectively prevent the right-traveling structural waves from propagating to the right half of the structure by creating a reactive discontinuity which reflects waves. In Fig. 36c, the spectral magnitudes of supersonic components of vibration decrease while the spectral magnitudes of most subsonic components actually increase slightly. Finally, the supersonic wavenumber spectrum shown in Fig. 36d bears a substantial resemblance to the farfield pressure distribution in Fig. 36a.

6.2.4. Controller performance and residual response using kinetic energy cost function

In this section we examine results similar to those of the previous section, except that the quantity being minimized by the controller is the kinetic energy response function $\hat{\mathbf{T}}$ rather than the radiated power response function $\hat{\Pi}$. As above, the response functions $\hat{\Pi}$, $\hat{\mathbf{T}}$, and Λ_{sup} before and after control are shown in Fig. 37 and the residual pressures, velocities, and wavenumber spectra are shown in Fig. 38.

As shown by Fig. 37, the kinetic energy response function turns out to be a poor choice for a controller cost function. The kinetic energy includes contributions from the surface velocities at all structural points, including the control drivepoints. Unfortunately, this has the effect of imposing a penalty on large controller gains. Large control inputs produce large drivepoint surface velocities which increase the kinetic energy response function, regardless of any associated reductions in the radiated field. At low frequencies, where the control drivepoint velocities are roughly two orders of magnitude larger than the velocity away from the drivepoints, applying even marginally effective control inputs raises the kinetic energy response function because of the control drivepoint contributions. As shown in Fig. 37b, $\hat{\mathbf{T}}$ is reduced by less than one decibel for frequencies below $k_o a$=0.75. (Note that these solutions do not pass the F-test, and are therefore suspect for statistical reasons. See Chapter 3.) Above $k_o a$=0.75, the control drivepoint velocities are only larger than the velocities at surrounding points by a factor of two or three rather than by orders of magnitude, so the effective penalty for applying a control input is less severe. Consequently, the controller is able to reduce $\hat{\mathbf{T}}$ by up to 5 dB, and thus reduce the radiated power response function by up to 12 dB.

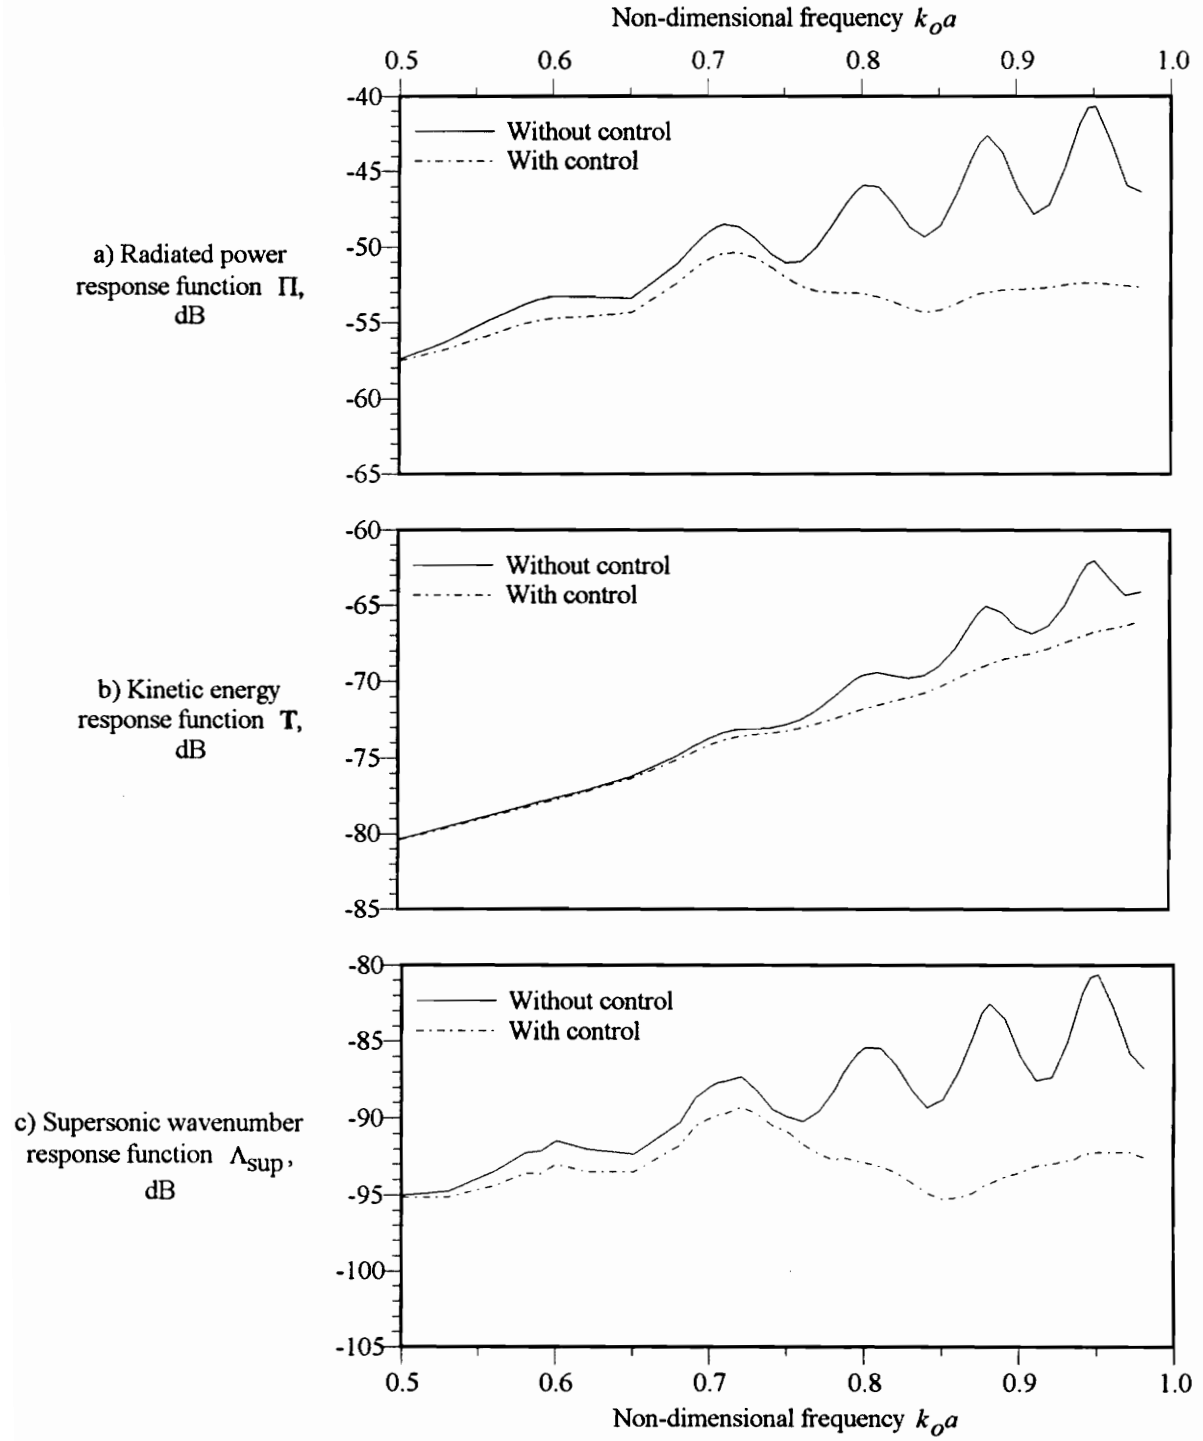

Figure 37: Response functions when controller minimizes kinetic energy response function. For each frequency, a separate regression is performed. Solid curve represents cost function without control, broken line represents cost function with control. Shown are radiated power response function $\hat{\Pi}$ (Fig. 37a), kinetic energy response function $\hat{\mathbf{T}}$ (Fig. 37b) and supersonic wavenumber response function Λ_{sup} (Fig. 37c).

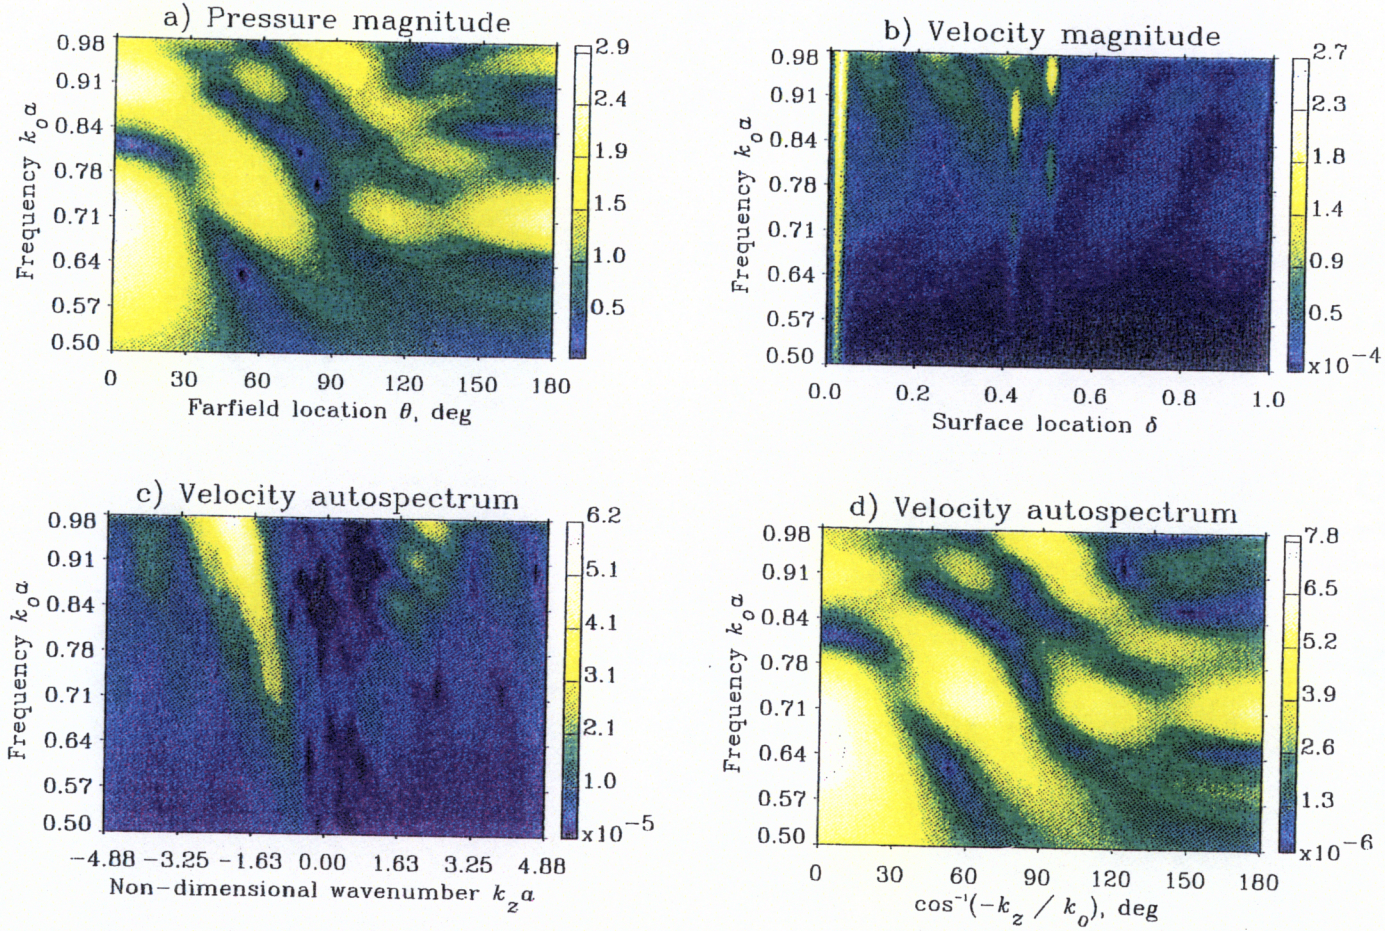

Figure 38: Response after controller minimizes kinetic energy response function. Shown are farfield pressure magnitude vs. frequency (Fig. 38a), surface velocity magnitude vs. frequency (Fig. 38b), and velocity autospectrum vs. frequency (Fig. 38c). Color corresponds to magnitude of quantity being plotted; see color scale at right of each plot.

Looking at the residual velocity magnitude for the kinetic energy controller (Fig. 38b), we see that the control inputs are indeed quite small at low frequencies since the drivepoint velocities at low frequencies are much smaller than at higher frequencies. Above $k_o a$=0.75, the controller behaves in a manner similar to the radiated power controller: waves in the left half of the cylinder are prevented from traveling to the right half. However, the isolation of the right half of the cylinder is not as complete as in previous cases, because this would likely cause higher velocities on the left half of the cylinder and thus increase the kinetic energy response function.

6.2.5. Controller performance and residual response using supersonic wavenumber cost function

As above we examine the effect of the controller on the response functions and on the residual response, but in this section the controller minimizes the supersonic wavenumber response function Λ_{sup}. The response functions $\hat{\Pi}$, $\hat{\mathbf{T}}$, and Λ_{sup} for the primary and residual responses are shown in Fig. 39 and the residual pressures, velocities, and wavenumber spectra are shown in Fig. 40. Given the degree of similarity between the farfield pressure distribution and the supersonic wavenumber distribution, one would expect that the control performance when minimizing Λ_{sup} would be nearly identical to the response when the radiated power response function is minimized; Fig. 39 confirms this expectation. Even though the controller uses only velocity sensors instead of farfield pressure sensors, the performance is very similar to that obtained using farfield sensors directly. Comparing Fig. 40 to Fig. 36 reveals that the residual responses are also nearly identical between the two different controller formulations. The close agreement confirms that the supersonic wavenumber response function encompasses all radiating portions of the response, making it a suitable choice for a cost function.

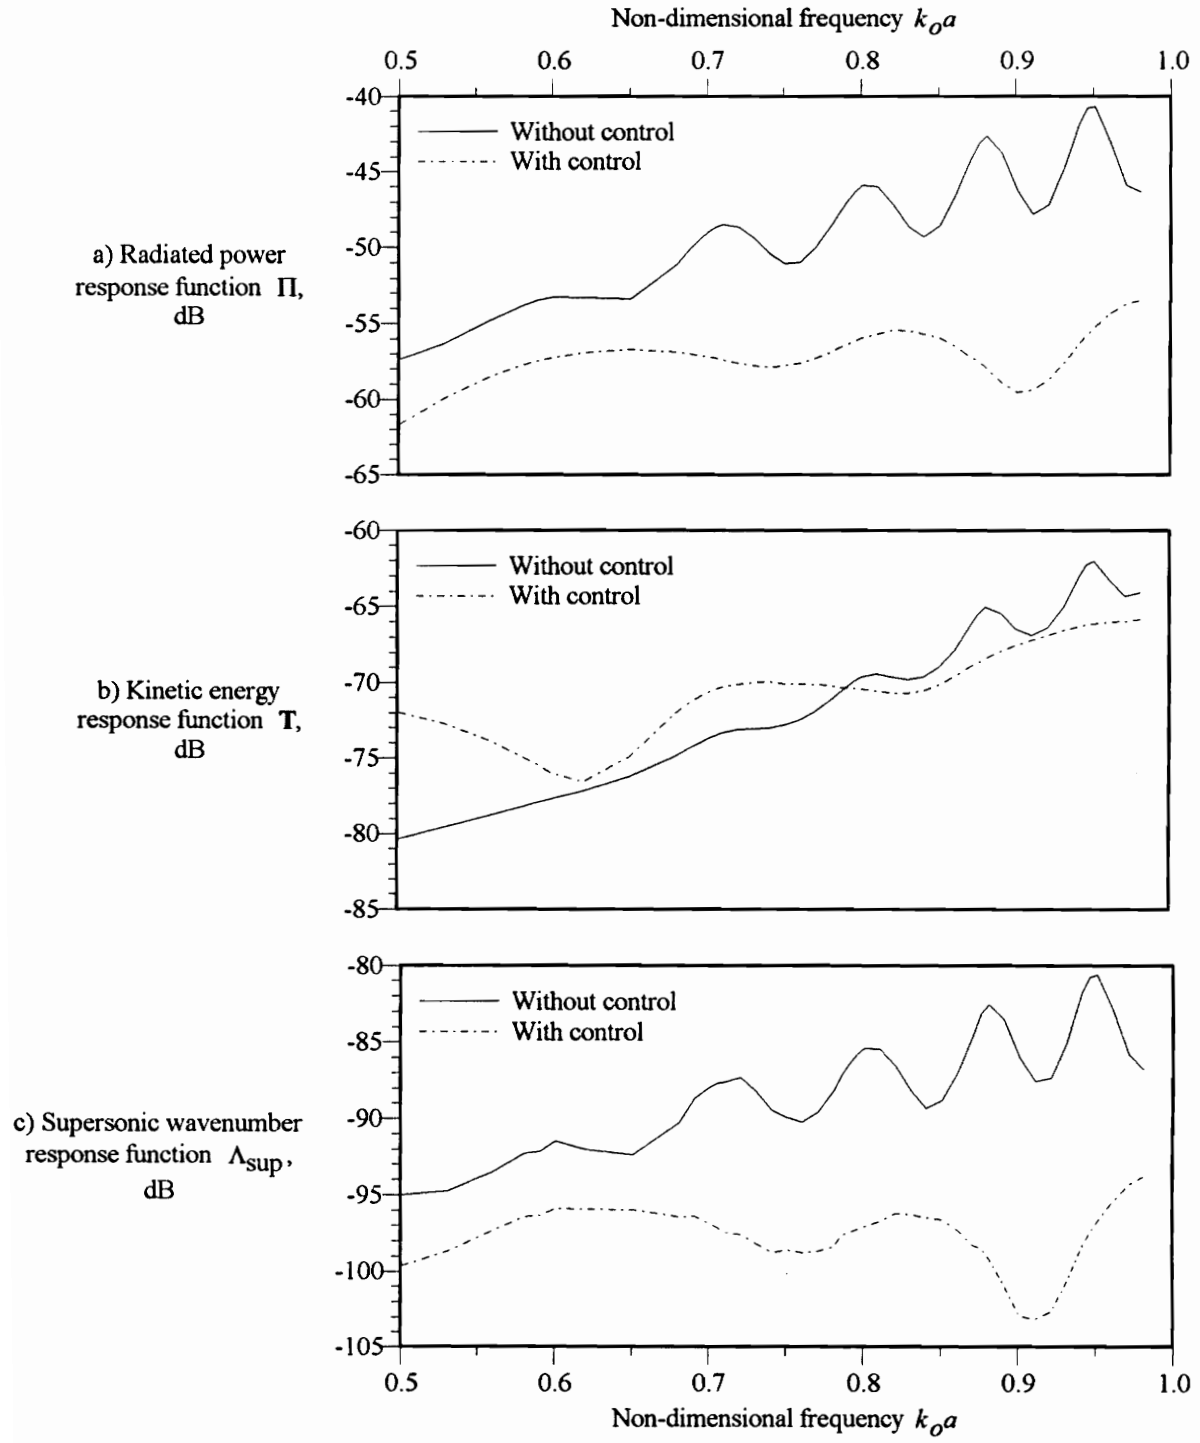

Figure 39: Response functions after minimizing supersonic wavenumber response function. Solid curve represents response function without control, broken curve represents response function with control. Shown are radiated power response function $\hat{\Pi}$ (Fig. 39a), kinetic energy response function $\hat{T}$ (Fig. 39b) and supersonic wavenumber response function Λ_{sup} (Fig. 39c).

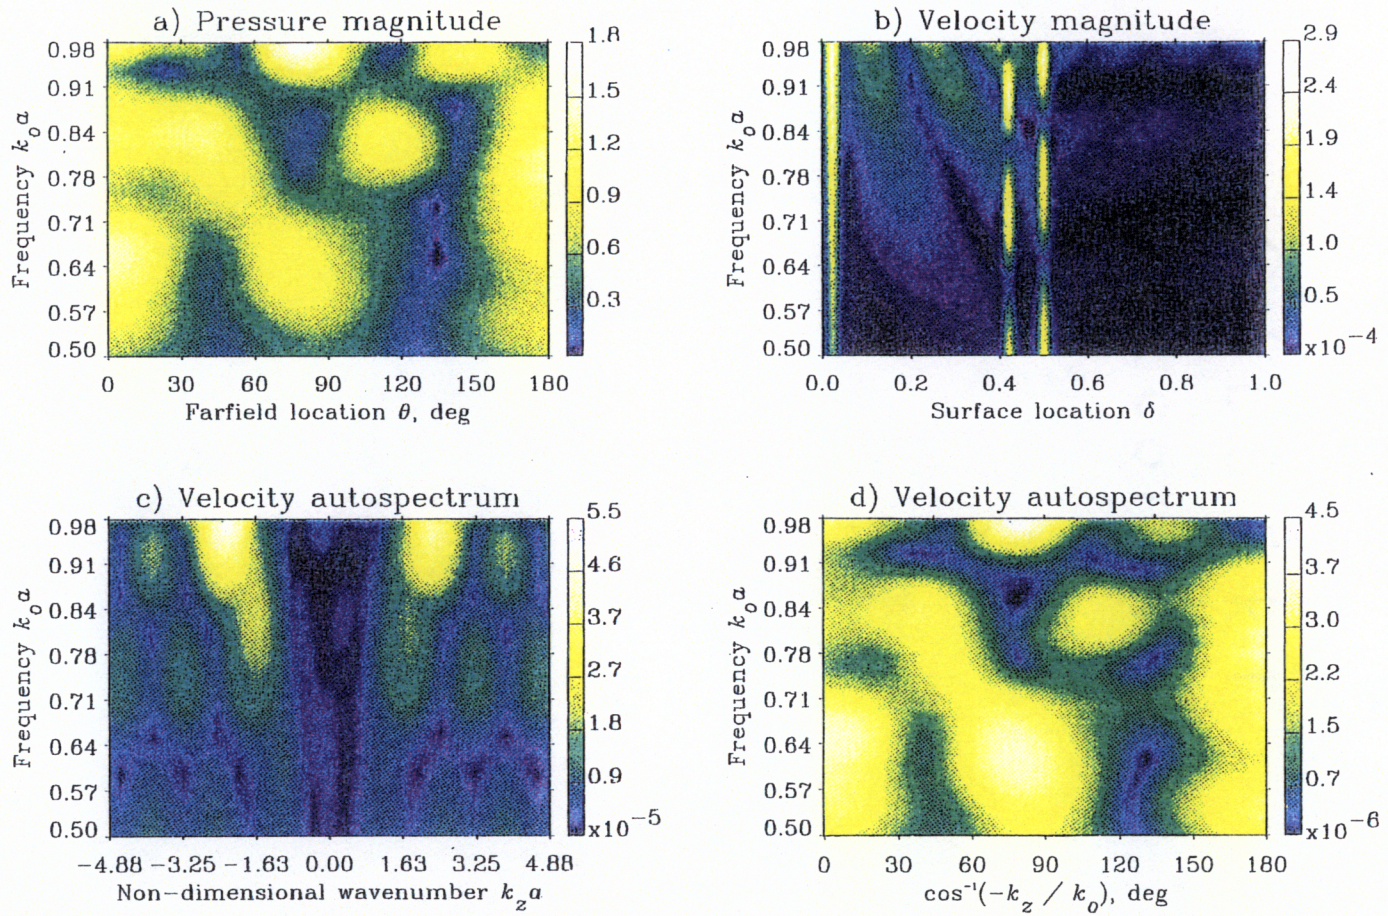

Figure 40: Response after controller minimizes supersonic wavenumber response function. Shown are farfield pressure magnitude vs. frequency (Fig. 40a), surface velocity magnitude vs. frequency (Fig. 40b), and velocity autospectrum vs. frequency (Fig. 40c and 40d). Color corresponds to magnitude of quantity being plotted; see color scale at right of each plot.

6.2.6. Controller performance and residual response using single-wavenumber cost function

For a final example, we define a narrow-band wavenumber controller that includes *only a single wavenumber component*. In other words, we use the weighting function of Eq. (82) centered at a structural wavenumber k_z. However, since we are examining a range of frequencies, let us define the weighting function such that the center wavenumber k_z is proportional to frequency:

$$w_i = \begin{cases} 1, & k_i = 0.5k_o \\ 0, & \text{elsewhere.} \end{cases} \tag{84}$$

Since there is only one non-zero component in the cost function, we can have only one control input. (If the number of control input exceeds the number of sensors, the system is underdetermined and has no unique solution.) For this section only, we use a single control input b_1 to minimize the cost function rather than using both control inputs b_1 and b_2 as in previous sections.

Figure 41 shows the residual response after applying the narrow-band wavenumber controller described above. The effect of the controller is most evident in Fig. 41d, where the wavenumber transform at k_z=$0.5k_o a$ is reduced almost to zero and appears as a vertical line in wavenumber space. This corresponds via Eq. (81) (page 106) to a reduction in the farfield pressure at 120°, which can be seen in Fig. 41a. Again, the closed-loop system has different eigenvalues as discussed in [57]. Thus it appears possible to selectively cancel radiation in desired directions by using a wavenumber controller that minimizes only certain components of the wavenumber transform.

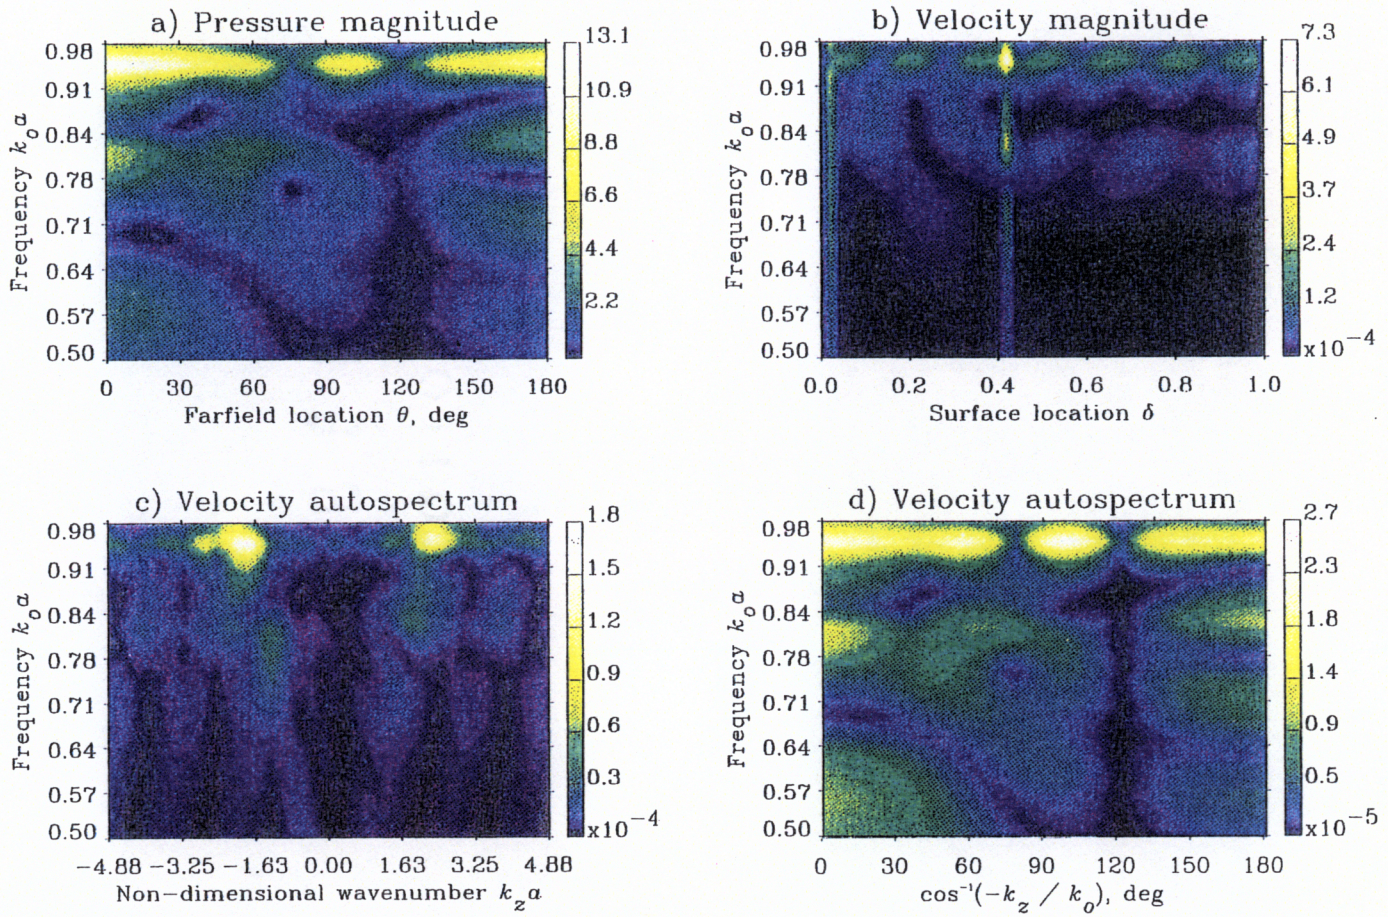

Figure 41: Response after controller minimizes wavenumber component at k_z=0.5$k_o a$. Shown are farfield pressure magnitude vs. frequency (Fig. 41a), surface velocity magnitude vs. frequency (Fig. 41b), and velocity autospectrum vs. frequency (Fig. 41c and 41d). Color corresponds to magnitude of quantity being plotted; see color scale at right of each plot.

6.3. Summary

When regression is used to simulate feedforward active control in the frequency domain, convenient methods are available for investigating different cost functions. By simply changing the type of transfer function used and the form of the weighting function, it is possible to compute numerous different "response functions" that help illustrate the system response. By choosing one response function for the controller cost function and examining the effect on other response functions, the analyst can perform numerical experiments to help understand the physical mechanisms of control, and help optimize the control system configuration.

For the clamped, cylindrical shell investigated in this chapter, ASAC appears to be an effective means of controlling acoustic radiation due to axisymmetric vibrations at low frequencies. With two control inputs, the radiated power controller reduces the estimated radiated power by up to 16 dB. The supersonic wavenumber controller also reduces the estimated radiated power, but does so without requiring the use of farfield sensors. The kinetic energy controller, which also does not require farfield sensors, performs poorly because the control drivepoints make disproportionately large contributions to the cost function, effectively penalizing the use of control inputs at most frequencies. To be more effective, the cost function would need to exclude the regions surrounding the drivepoints.

The novel data format used to present the supersonic wavenumber spectrum in, for example, Fig. 33 provides an easily understandable way to link wavenumber information with farfield information. By plotting the wavenumber transform against a structural wavenumber that has been normalized by the acoustic wavenumber, we can immediately verify whether the computed wavenumber spectrum displays sufficient relation with the farfield pressure field.

CHAPTER 7: SUMMARY AND RECOMMENDATIONS

7.1. Summary

This work establishes a general set of regression-based numerical procedures for investigating feedforward active structural-acoustic control (ASAC) in the frequency domain. Not only do the procedures accurately reproduce the results of analytical methods, but they also provide a more general paradigm in which to consider the control problem. The language and concepts of regression can improve understanding of how feedforward control systems work for a given structure, and provide insights into the physical mechanisms of active noise control. In fact, many concepts from statistical science can be applied more or less directly to the feedforward control problem. The present analysis includes only the most fundamental regression diagnostics, and will hopefully set the stage for adapting other statistics methods. Some possible adaptations are discussed in Section 7.2.

Although most of the numerical methods refined in this thesis are illustrated using axisymmetric examples, they are easily extendible for 3-dimensional problems. An example of how to approach a 3-dimensional problem is given in Chapter 2, which discusses a spherical shell with multiple control forces. Analyzing detailed results other than the radiated power would be more difficult for a 3-dimensional problem than for a 2-dimensional problem because of the amount of data involved. Data such as farfield pressure directivities that are plotted using line plots in an axisymmetric problem would become surface plots in a 3-dimensional problem; data surfaces would become data volumes, and so forth. However, none of the methods described here place any restrictions on the geometry of the structure.

7.1.1. Spherical shell

Chapter 2 develops frequency-domain analytical and numerical methods for investigating feedforward ASAC of a fluid-loaded spherical shell. The analytical method is derived using methods similar to those of Elliot and Nelson [15], and the numerical method is derived by discretizing the integrations of the analytical method. The controller minimizes the estimated radiated power at a single frequency. A FEM/BEM approach is used to calculate structural-acoustic dynamic responses for the numerical method. To check the validity of the numerical method, results are compared between the numerical and analytical methods.

The numerical and analytical solutions, considering the fundamentally different processes by which they are derived, show satisfactory agreement both with and without control. They also demonstrate that ASAC is an efficient method for controlling radiated noise at low frequencies ($0<<1.7$). At resonance frequencies, radiation due to a point-force disturbance can be reduced by up to 20 dB using only one actuator. Similar levels of attenuation are achieved using multiple control forces.

7.1.2. Regression and regression diagnostics

A numerical simulation of feedforward ASAC in the frequency domain can be represented as a multiple linear least-squares regression; the main difference from regressions presented in the statistics literature is that the variables are complex-valued rather than real. Regression provides the magnitudes and phases of the control forces, and can be used in the general case of an axisymmetric or 3-dimensional structure modeled using numerical methods. After making suitable assumptions about the error sensor measurement noise, we can use regression diagnostics to model the effects of measurement noise and of using a finite number of sensors. Thus, in

addition to detecting potential problems with the regresson itself, regression diagnostics can help provide more realistic simulations.

The F-test measures the statistical significance of the regression as a whole. When the F-test fails, the shape of the control response is significantly different from the shape of the disturbance response. Any acoustic cancellation that takes place is local rather than global, and is due more to chance than to a statistically significant model. Failing the F-test is always accompanied by poor controller performance; in the example results for the clamped cylinder, the F-test fails whenever the reduction in the cost function is smaller than 2 dB.

The t-test and associated confidence intervals provide a statistical measure of the probable magnitude errors in our estimated contol input estimates. When the confidence interval for a given control input contains zero, the associated actuator does not contribute significantly to the regression and should be removed to make the regression more reliable.

Similarly, we can find prediction intervals on the secondary response to help measure how sensitive the regression is to error sensor measurement noise. Whereas the confidence intervals address individual control inputs, the prediction intervals address the regression as a whole. Prediction intervals do not provide confidence limits for the reduction in radiated power cost function, but could be used with bootstrapping techniques to provide such estimates.

In developing the regression diagnostics described above, it is important to examine the statistical properties of measurement noise. The present analysis assumes normally distributed measurement noise with constant variance, and the measurement noise at each error sensor is independent from the measurement noise at other sensors. Often with real sensors the error variance is proportional to the signal being measured rather than being constant. Assuming

constant-variance measurement noise is a convenient approximation which may be used to avoid the complications associated with more sophisticated analyses. For a numerical simulation the analyst must estimate a value for the variance. If characteristics are known *a priori* for a specific type of sensor that will be used, they may be used to estimate the variance. Otherwise, the variance may be estimated form the data as is done in the present work.

7.1.3. Collinearity and collinearity diagnostics

The collinearity problems common in many statistical regressions can occur when regression is used for a numerical simulation of feedforward active control in the frequency domain. Without proper collinearity diagnostics, numerical problems can go undetected and produce results that are excessively sensitive to error sensor measurement noise. Collinearity diagnostics for complex-valued regressions follow directly from their real-valued counterparts, which are developed in considerable detail in the statistics literature.

The primary collinearity diagnostic is the scaled condition number. Scaling allows for a quantitative comparison between the scaled condition number and a predetermined threshold value to determine whether or not collinearity is present. Even when no other collinearity diagnostics are performed, the analyst may compute the scaled condition number for each regression and compare it to the threshold value in a somewhat mechanical fashion.

More detailed diagnostics are available for determining the number of near-dependencies and which control forces are involved. These diagnostics are less mechanical, requiring more interaction on the part of the analyst, but may provide interesting insights when simulating complicated structures.

7.1.4. Optimizing actuator locations

Subset selection provides an efficient, versatile, general way to optimize actuator locations for the best possible system performance. The method is computationally efficient for subsets of fewer than, say, ten actuators. It requires the generation of a large number of transfer functions (one per candidate actuator location per frequency). However, generating transfer functions for many different loading cases does not significantly increase computational expenses if a finite-element-based method is used for generating transfer functions.

For complex geometries and/or MIMO control system it is important to have some sort of formal procedure to optimize actuator locations. Optimal locations are difficult to predict *a priori*, even for the relatively simple clamped cylinder investigated here. For example, a logical assumption for optimal actuator locations on the clamped cylinder would be to place actuators at the anti-nodes. However, a pair of actuators placed near a node proved more effective at reducing radiation. For more complex structures, guessing the optimal actuator locations would be even more difficult. Furthermore, it is important to examine not just one optimal subset, but rather a number of near-optimal subsets. In some circumstances the system performance can be extremely sensitive to the actuator locations, a condition that becomes obvious when examining a number of near-optimal configurations.

In performing subset selection, the analyst must be aware of possible collinearity problems and take steps to correct for them. In exhaustive-search subset selection, this means the analyst must check for collinearity in each near-optimal subset being investigated. This is possible using the collinearity diagnostics of Chapter 4. Any near-optimal subset that is contaminated by collinearity must be discarded.

7.1.5. Cost functions and weighting functions

When we use regression to perform a numerical simulation of feedforward control, a convenient method is available for investigating alternate controller formulations. Various response functions can be formed from different combinations of transfer functions and weighting functions. Performing numerical experiments that choose different response functions for the controller cost function can be instructive in determining the physical mechanisms used by the controller. For example, we can use area-weighting with farfield pressure sensors to produce a response function proportional to the radiated power. We can use mass-weighting with surface velocity sensors to produce a response function proportional to the total kinetic energy. We can transform the surface velocities to produce a wavenumber-domain response function. Any of these response functions may be used as a cost function by the controller.

7.2. Recommendations for further research

Several opportunities for further research concern the analogy between feedforward control and multiple regression. The statistical methods described in the present work, namely regression diagnostics, collinearity diagnostics and subset selection, constitute only the most basic statistical considerations. Many other opportunities exist to apply statistical methods to the active control problem, although most will require slight modification to operate on complex variables rather than real variables. Some recommendations are listed here:

1. A more detailed study of the role of error sensor measurement noise should be performed. The noise variance could be estimated *a priori* for a specific type of sensor, rather than being estimated from the residual variance as in the present work. Also, regression transformations could be used to provide a more realistic model of noise variance than that

described in Section 3.1. For example, the noise variance could be made proportional to the magnitude of the signal being measured by performing a loglinear transformation.

2. Robust estimation techniques could be used to estimate regression parameters in the case that measurement error is known to be significantly non-Gaussian.

3. Bootstrapping with Monte Carlo simulation, or other statistical methods, could be used to obtain confidence bounds on the residual cost function. The present analysis computes confidence limits only on the magnitudes of the control inputs and the estimated model, not on the amount of reduction in the cost function.

4. It may be possible to optimize sensor locations using regression diagnostics originally intended for detecting outliers and influential data. Standard methods are available for detecting whether some observations in a regression have disproportionately high influence on the regression results. In the feedforward control problem, this corresponds to some error sensors being more influential than others. Detecting influential sensors would aid the designer both in understanding the system dynamics and in reducing the number of sensors required.

5. A more detailed study of collinearity diagnostics and their application in system design would be appropriate. This thesis only discusses the most elementary concepts of collinearity diagnostics. Much more sophisticated approaches are possible, and may prove useful in analyzing system configurations. In particular, a more detailed study should be made of what threshold values are appropriate for the case of feedforward control.

6. The effects of measurement noise on subset selection should be studied to determine whether it is possible to compute confidence intervals for optimal actuator locations. The

present analysis ignores measurement noise, although clearly some of the optimal subsets are very sensitive to small perturbations in the measured data.

7. The influence of measurement noise in the transfer functions X_i should be studied. The present analysis assumes that X_i is known with a high degree of accuracy, since measurement noise in X_i can be ignored when it is small in comparison with the measurement noise in Y.

8. Non-linear multiple regression should be investigated for situations in which the actuators do not respond linearly. The present analysis assumes linear actuators, e.g., the response doubles when the actuator signal is doubled.

9. Other subset selection techniques should be investigated for optimizing actuator location. Methods such as stepwise regression, backward elimination, and others would use far less computing time than exhaustive search. However, the effects of collinearity on such methods would require special caution. Furthermore, other methods can become trapped in local optima instead of always finding global optima.

A study should be made of the effects of structural damping on frequency-domain active control simulations. One likely effect would be that actuator locations far removed from the disturbance force location would be less effective than actuator locations near the disturbance force location; this would effect results of actuator optimization studies.

Another interesting topic would be to investigate the use of *constrained* least-squares. For example, we could pose the constrained least squares problem as

$$\min_{B} \sum_{j=1}^{M} w_j \left| y_j + \hat{y}_j \right| \quad \text{subject to} \quad \mathrm{GB} = \mathrm{H} \tag{85}$$

where G and H contain constraint information. One potential use for constrained least-squares would be to implement a penalty for large control force magnitudes; this would prevent small improvements in the residual cost function at the price of large increases in control effort. Another potential use for constrained least-squares would be to implement magnitude limits for some or all of the control forces.

Another possible application of constrained least-squares would be optimizing actuator locations for multi-frequency disturbance excitation. For example, suppose the normal equations for two frequencies f_1 and f_2 were

$$\begin{aligned} &\text{at frequency } f_1, \quad Y_1 = X_1 B_1 \\ &\text{at frequency } f_2, \quad Y_2 = X_2 B_2 \end{aligned} \tag{86}$$

We could then write both regressions as the single augmented normal equation

$$\begin{Bmatrix} Y_1 \\ Y_2 \end{Bmatrix} = \begin{bmatrix} X_1 & 0 \\ 0 & X_2 \end{bmatrix} \begin{Bmatrix} B_1 \\ B_2 \end{Bmatrix}, \tag{87}$$

and use subset selection to optimize the actuator locations. However, we would need to ensure that each subset contains the same number of actuators from X_1 and X_2, rather than just selection columns at random from the entire augmented matrix. This could be accomplished by performing subset selection with equality constraints.

A more general topic for development would be to incorporate the regression and subset selection methods described here into a coherent, generalized suite of computer programs for concurrent passive/active design. Close integration with existing structural analysis codes would be an important part of such an approach, as would integration with visual data analysis software.

Finally, and perhaps most importantly, a detailed comparison of experimental measurements and numerical simulations for a specific structure would provide further verification of the

method, particularly the modeling of sensor measurement noise. Numerical techniques could be used for an in-depth study of the control mechanisms applicable to that specific structure. The flexibility of the numerical methods based on finite-element models would compliment the experimental results to provide a reliable database for use in a concurrent design effort.

APPENDIX A: LIST OF SYMBOLS

Superscripts

*	complex conjugation
H	Hermitian transpose
T	transpose

Symbols

a	shell radius
A_i	transfer function of i^{th} disturbance force
a_i	area associated with the i^{th} sensor
$a_{\max}$	maximum value of a_i
α	force location on spherical shell, given as latitudinal angle measured from global z-axis
α_t	level of significance used in t-test
B	vector of complex control forces; in a regression, vector of estimated regression coefficients
b_i	complex strength of i^{th} control force; in a regression, estimated regression coefficients
β	true regression coefficients
c	acoustic velocity in fluid
χ^2	cost function to be minimized
δ	normalized location on surface of cylinder ($\delta = z+L/2$)
D	diagonal matrix of singular values

$\Delta\phi$	farfield interval in ϕ-direction (pressure sensors are separated by longitudinal angle $\Delta\phi$)
$\Delta\theta$	farfield interval in θ-direction (pressure sensors are separated by latitudinal angle $\Delta\theta$)
$E(x)$	expected value of x
ε_i	measurement noise of i^{th} member of Y
F	variance ratio used in performing the F-test
ϕ	longitudinal angle measured from x-z plane
$g^2(\tilde{y}_i, w_i)$	variance function
$\Gamma_n(R)$	See Eq. (3).
h'	derivative of spherical Hankel function
H	hat matrix
$\tilde{\eta}_k$	the k^{th} scaled condition index
$\tilde{\eta}^*$	threshold value for scaled condition indexes
k_z	axial structural wavenumber
k_o	acoustic wavenumber
L	length of cylindrical shell
Λ_{sup}	supersonic wavenumber response function
m_i	relative mass associated with the sensor location $\boldsymbol{r}_i$
M	number of sensors
M_f	number of sensors in ϕ-direction
M_q	number of sensors in θ-direction
μ_k	the k^{th} singular value of $\boldsymbol{Z}$
μ_Y	true mean of Y

n mode number for spherical shell; equals number of nodal circles present in surface velocity

N number of control forces

N_d number of disturbance forces

P farfield acoustic pressure

$\hat{P}$ control response (secondary response), or estimated model in a regression

$\tilde{P}$ true model in a regression

P_p disturbance response (primary response)

P_n Legendre polynomial of order n

Π total radiated power

$\hat{\Pi}$ radiated power response function

π^* threshold value for variance-decomposition proportions

π_{kj} variance-decomposition proportion for k^{th} regression coefficient and j^{th} scaled condition index

θ latitudinal angle measured from global z-axis

Θ right-hand side of normal equation after using QR decomposition of X

$Q(F, \nu_1, \nu_2)$ F-distribution probability function with degrees of freedom ν_1 in the numerator and ν_2 in the numerator

$\boldsymbol{Q}$ left-hand side of a QR decomposition

$\boldsymbol{r}_i$ vector denoting location of i^{th} sensor

R farfield radius for computing pressure

$\boldsymbol{R}$ right-hand side of a QR decomposition

$\mathbf{R}^2$ squared multiple correlation coefficient (squared coefficient of determination)

R_{xy} correlation coefficient between x and y

ρ	mass density of fluid
s_i	complex strength of i^{th} disturbance force
s_Y	standard deviation of Y
S	vector of complex disturbance forces
σ_Y	constant portion of variance of Y
t	Student's probability distribution for level of significance α_t with ν degrees of freedom
T	transformation matrix used in generating a QR decomposition
$\boldsymbol{T}$	total kinetic energy of a structure in steady-state harmonic motion
$\hat{\mathbf{T}}$	total kinetic energy response function
u	radial component of surface velocity
$\tilde{\mathbf{U}}(k_{zi})$	wavenumber transform at wavenumber k_{zi}
U	left-hand side of singular-value decomposition of design matrix
v_{kj}	elements of the matrix V
V	variance operator
V	right-hand side of singular-value decomposition of design matrix
w_i	weighting coefficient for the i^{th} sensor
$w(x)$	weighting function
W	diagonal matrix containing the weighting coefficients w_i
$W^{1/2}$	diagonal matrix whose elements are the square roots of corresponding elements of W
ω	circular frequency
$\boldsymbol{\Omega}^{-1/2}$	scaling matrix used to scale the weighted design matrix in collinearity diagnostics
Ω	dimensionless frequency ω/ω_r
X	matrix of transfer functions; in a regression, the design matrix

X_i	transfer function of i^{th} control force; in a regression, independent or predictor variable
Y	measured or observed data in a regression
$\tilde{Y}(\beta_i, X_i)$	true model in a regression
$\hat{Y}$	estimated regression model
$\overline{Y}$	sample mean of Y
z	longitudinal coordinate of cylindrical shell
z_n	modal acoustic impedance
$\boldsymbol{Z}$	weighted, scaled design matrix used in collinearity diagnostics
Z_n	modal structural impedance

APPENDIX B: LINEAR LEAST-SQUARES REGRESSION USING COMPLEX VARIABLES

B.1. Description of routines

In [50], A.J. Miller presents a set of FORTRAN routines for performing linear, weighted, least-squares regression on real-valued data. The routines are based on [48] and [49], which give procedures for linear least squares problems using a *QR* decomposition. The method allows for changing the order of variables and calculating regression coefficients for subsets of variables. Also included in [50] are routines for calculating correlations and partial correlations, covariances, and various diagnostic quantities described in Belsley, Kuh, and Welsch [37].

This appendix describes modifications to [50] by which it is possible to perform linear least-squares regression on complex-valued data. The theory behind complex-valued regression is fairly well-established, having been developed primarily for the frequency-domain analysis of time series. Relevant works include [28, 29, 59, 60, 61].

The modifications required to adapt [50] to the case of complex variables are trivial, amounting to little more than a few complex conjugations sprinkled among the program statements. For brevity and to avoid duplication with published material, this appendix focuses only on the modifications rather than attempting to present a completely self-contained description of the algorithm. For a more complete picture, the interested reader is invited to follow the development in [48] and [50] while reading this section.

The most obvious modifications to [50] are that complex variables must be declared as complex rather than real. Rather than showing each such declaration separately, the modified code segments given in this article assume that complex variables have been properly declared. Another obvious modification required in some routines is that intrinsic functions such as ABS must be modified to accommodate complex variables; these changes will be obvious to the programmer, and need not be included here. In one of the routines, a new variable is added and is noted in the accompanying discussion. Otherwise, no changes are required to portions of [50] not specifically mentioned.

For clarity in code segments all complex variables are shown in boldface type. Table 5 lists the variables in each routine which must be declared as complex. Program statements that require modifications other than those above are denoted by the comment "!! MODIFIED !!" near the right margin.

Table 5: List of variables which must be declared complex.

Routine:	Variables to be declared complex:
CLEAR	RBAR THETAB
COV	RBAR RINV COVMAT SUM
HDIAG	XROW RBAR WK SUM
INV	RBAR RINV SUM
INCLUD	XROW YELEM RBAR THETAB Y XI WXY SBAR XK XSTAR
REGCF	RBAR THETAB BETA
SING	RBAR THETAB
SS	THETAB
TOLSET	RBAR
PCORR	RBAR THETAB CORMAT YCORR
COR	RBAR THETAB CORMAT YCORR SUM
VMOVE	RBAR THETAB SBAR X Y XSTAR
REORDR	RBAR THETAB

B.2. Basic procedure: QR decomposition of a complex matrix

Consider an n x 1 vector of observations y and an n x p design matrix X which define a linear least squares problem to compute a p x 1 vector of regression coefficients β, where β minimizes the sum of the squares of the elements of the n x 1 residual vector r defined by

$$r = y - X\beta. \tag{88}$$

If all of the above quantities are real-valued, solves the well-known normal equations

$$X^T X\beta = X^T y. \tag{89}$$

To solve Eq. (89), [50] uses the QR decomposition

$$X = \mathbf{QR}, \tag{90}$$

where $\boldsymbol{Q}$ is an n x p matrix whose columns are orthonormal and $\boldsymbol{R}$ is p x p upper triangular. The normal equations then become

$$\mathbf{R}\beta = \mathbf{Q}^T y = \Theta, \tag{91}$$

where $\boldsymbol{R}$ and Θ are obtained by applying a sequence of Givens transformations. Each such transformation takes the form

$$T_{ij} = \begin{bmatrix} I_{i-1} & & & & (0) \\ & \mathbf{c} & \cdots & \mathbf{s} & \\ & \vdots & I_{j-i} & \vdots & \\ & -\mathbf{s} & \cdots & \mathbf{c} & \\ (0) & & & & I_{p-i} \end{bmatrix} \tag{92}$$

where I_n is an n x n identity matrix. The quantities c and s are chosen such that T_{ij} is orthogonal, and multiplying X by T_{ij} obliterates the (i,j) element of X.

Consider now a similar system in which all the quantities are complex-valued. The normal equations now become

$$X^{*T} X\beta = X^{*T} y. \tag{93}$$

After decomposing X, the normal equations can be written

$$\mathbf{R}\beta = \mathbf{Q}^{*T} y = \Theta. \tag{94}$$

To obtain the QR decomposition of the complex matrix X, we must have a Givens transformation that is unitary, rather than merely orthogonal. As in [62], the following matrix satisfies our requirements.

$$T_{ij} = \begin{bmatrix} I_{i-1} & & & & (0) \\ & \mathbf{c} & \cdots & \mathbf{s} & \\ & \vdots & I_{j-i} & \vdots & \\ & -\mathbf{s}^* & \cdots & \mathbf{c}^* & \\ (0) & & & & I_{p-i} \end{bmatrix} \tag{95}$$

Propagating these modifications through [48], we can easily show that only two formulas must be changed to account for complex variables. In Eq. (9) of [48], the real variable $d' = d + \delta x_i^2$ becomes $d' = d + \delta x_i x_i^*$, and thus remains a real quantity. Also in Eq. (9) of [48], the real variable $\bar{s} = \delta x_i / d'$ becomes $\bar{s} = \delta x_i^* / d'$, and thus becomes a complex quantity. (Note that the variable δ used here is different from the axial coordinate δ used elsewhere in this thesis; the notation is used only in this paragraph for consistency with [48], and hopefully will not cause confusion with other sections of the text.)

B.3. Modifications to INCLUD and VMOVE

Turning now to the actual algorithm from [50], two of the routines contain variations of code for implementing the method of [48], and must be modified in a similar manner: INCLUD creates

the original decomposition and VMOVE restores it after rearranging the order of the rows. In routine INCLUD, we modify only two lines of code as shown below. (Recall that variables which must be declared as complex are shown in boldface.) The first modification involves calculating the intermediate product WXI as

```
      WXI = W * CONJG(XI)                                    !! MODIFIED!!
```

rather than the original statement which contains WXI = W * XI. Since WXI is used to compute both DPI and SBAR, no further modifications are needed to actually form the *QR* decomposition.

Only one other modification is needed in INCLUD. The variable SSERR is incremented by an amount proportional to the squared magnitude of Y rather than the square of Y, and must be computed as shown below.

```
C     Y * SQRT(W) is now equal to Brown & Durbin's recursive residual.
      SSERR = SSERR + W * Y * CONJG(Y)                       !! MODIFIED!!
```

In routine VMOVE, the modifications are more numerous but still quite trivial. In the segment of code for dealing with "special case" of D2 equal to zero, we define a new variable XSTAR equal to the complex conjugate of X. The variables D(M), RBAR(M1), and DTHETAB(M) are calculated using XSTAR instead of X.

```
      ELSE IF (D2 .EQ. ZERO) THEN
         XSTAR = CONJG(X)                                    !! MODIFIED!!
         D(M) = D1 * X*XSTAR                                 !! MODIFIED!!
         RBAR(M1) = ONE / XSTAR                              !! MODIFIED!!
         DO 20 COL = M+2, NP
            M1 = M1 + 1
            RBAR(M1) = RBAR(M1) / XSTAR                      !! MODIFIED!!
   20    CONTINUE
         THETAB(M) = THETAB(M) / XSTAR                       !! MODIFIED!!
         GO TO 40
      END IF
```

In the segment of code for dealing with the "regular case" of D1, D2, and X all nonzero, we find similar modifications. In this case, the variable XSTAR is multiplied by XSTAR to save multiplications in computing D1NEW and SBAR.

```
XSTAR = D1 * CONJG(X)                                    !! MODIFIED!!
D1NEW = D2 + X*XSTAR                                     !! MODIFIED!!
CBAR = D2 / D1NEW
SBAR = XSTAR / D1NEW                                     !! MODIFIED!!
```

Finally, residual sums of squares are calculated as the product of THETAB(MP1) and its complex conjugate.

```
RSS(M) = RSS(MP1) + D(MP1) * THETAB(MP1)*CONJG(THETAB(MP1))!! MODIFIED!!
```

B.4. Modifications to SING and SS

Routines SING and SS each require only one modification. As in VMOVE, the residual sums of squares are calculated by multiplying THETAB(COL)by its complex conjugate.

```
SSERR = SSERR + D(COL) * THETAB(COL)*CONJG(THETAB(COL))     !! MODIFIED!!
```

B.5. Modifications to COV and INV

Although routine INV requires no modifications other than declaring complex variables, routine COV does require minor modifications. The covariance matrix takes the form

$$\text{cov} = \sigma^2 \left[X^{*T} X \right]^{-1}, \tag{96}$$

where σ is the variance. Substituting $X = QR$, we can obtain

$$\text{cov} = \sigma^2 \mathbf{R}^{-1} \left[\mathbf{R}^{-1} \right]^{*T}, \tag{97}$$

where R^{-1} is obtained by calling INV. In the code fragment below, note that the diagonal terms of the array COVMAT are purely real.

```
      POS = 1
      START = 1
      DO 40 ROW = 1, NREQ
         POS2 = START
         DO 30 COL = ROW, NREQ
            POS1 = START + COL - ROW
            IF (ROW .EQ. COL) THEN
               SUM = ONE / D(COL)
            ELSE
               SUM = CONJG(RINV(POS1-1)) / D(COL)                    !! MODIFIED!!
            END IF
            DO 20 K = COL+1, NREQ
               SUM = SUM + RINV(POS1) * CONJG(RINV(POS2)) / D(K)   !! MODIFIED!!
               POS1 = POS1 + 1
               POS2 = POS2 + 1
20          CONTINUE
            COVMAT(POS) = SUM * VAR
            IF (ROW .EQ. COL) STERR(ROW) = SQRT(DBLE(COVMAT(POS)))!! MODIFIED!!
            POS = POS + 1
30       CONTINUE
         START = START + NREQ - ROW
40    CONTINUE
```

B.6. Modifications to PCORR and COR

Routine PCORR calls routine COR to compute partial correlations. While PCORR requires no modifications other than declaration statements, slight modifications are required for COR. The complex correlation coefficient between two complex variables x and y may be written

$$R_{xy} = \left\{ \sum_i x_i^* x_i \sum_i y_i^* y_i \right\}^{-1/2} \sum_i x_i^* y_i \,. \tag{98}$$

The correlation coefficient given above exhibits the required property that $R_{xy} = 1$ when there is perfect correlation between x and y, and $R_{xy} = -1$ for perfect negative correlation, and $R_{xy} = 0$ for no correlation. However, it does little to show the relation between the real and imaginary parts of x and y. There may be other definitions which are more useful. For example, it may be more illuminating to manually separate x and y into their real and imaginary parts and compute

correlations separately, e.g. one statistic for the correlation relating Re{x} to Re{y}, and a separate statistic relating Im{x} to Im{y}.

To compute the complex correlation coefficient shown in Eq. (16), we first calculate the lengths of the two vectors by summing their squared magnitudes. We accomplish this using

```
10      SUMY = SUMY + D(ROW) * THETAB(ROW)*CONJG(THETAB(ROW))        !! MODIFIED!!
```

and using the following set of statements.

```
          DO 20 ROW = 1, MIN(COL1-1, NP)
             SUM = SUM + D(ROW) * RBAR(POS1)*CONJG(RBAR(POS1))      !! MODIFIED!!
             POS1 = POS1 + NP - ROW - 1
20        CONTINUE
          WORK(COL1) = DSQRT(DBLE(SUM))                             !! MODIFIED!!
```

Although the length being evaluated above is a real quantity, there is no need to define a separate real accumulator variable for the summation. The complex quantity **SUM** may be used without consequence. The partial correlations are then found using

```
          SUM = SUM + D(ROW) * CONJG(RBAR(POS1)) * THETAB(ROW)      !! MODIFIED!!
```

to compute the Y-correlations, and

```
          SUM = SUM + D(ROW) * CONJG(RBAR(POS1)) * RBAR(POS2)!! MODIFIED!!
```

to compute the covariances.

B.7. Modifications to HDIAG

Routine HDIAG computes quantities useful in diagnostic procedures as outlined in [37]. The hat matrix H, used to find confidence limits on the estimated model as in Chapter 3, may be found as shown below.

$$H = X(X^{*T}X)^{-1}X^{*T} = X\mathbf{R}^{-1}(X\mathbf{R}^{-1})^{*T} \tag{99}$$

Aside from changing the workspace array WK and the summation variable SUM to complex variables, we make only one modification to HDIAG. To calculate HII we must sum the squared magnitudes of the columns of XR^{-1}, as shown below.

```
      HII = HII + SUM*CONJG(SUM) / D(COL)                    !! MODIFIED!!
```

REFERENCES

[1] C.R. Fuller, "Apparatus and method for global noise reduction," US Patent No. 4,715,59 (1987).

[2] B.-T. Wang, "Active control of sound transmission/radiation from elastic plates using multiple piezoelectric actuators," Unpublished doctoral dissertation, Department of Mechanical Engineering, Virginia Polytechnic Institute and State University (June, 1991).

[3] R.L. Clark and C.R. Fuller, "Control of sound radiation with adaptive structures," Journal of Intelligent Material Systems and Structures, **2**, pp.431-45 (1991).

[4] R.L. Clark and C.R. Fuller, "Active structural acoustic control with adaptive structures including wavenumber considerations," Proceedings of the Conference on Recent Advances in Active Control of Sound and Vibration (April 1991).

[5] C.R. Fuller and R.A. Burdisso, "A Wavenumber Domain Approach to the Active Control of Structure Borne Sound," Journal of Sound and Vibration, **148**(2), pp.355-360 (1991).

[6] C.R. Fuller and J.D. Jones, "Experiments on reduction of propeller-induced interior noise by active control of cylinder vibration," Journal of Sound and Vibration, **112**(2), pp.389-395 (1987).

[7] C. Guigou and C.R. Fuller, "Active control of edge sound radiation from a semi-infinite elastic beam with a clamped edge: Influence of bending near-field waves," Journal of the Acoustical Society of America, **93**(5), pp.2716-25 (1992).

[8] C.R. Fuller, C.H. Hansen and S.D. Snyder, "Active control of sound radiation from vibrating rectangular panel by sound sources and vibration inputs: An experimental comparison," Journal of Sound and Vibration, **145**(2), pp.195-215 (1991).

[9] H.C. Lester and C.R. Fuller, "Active control of propeller induced noise fields inside a flexible cylinder," AIAA Journal, **28**(8), pp.1374-1380 (1989).

[10] C.R. Fuller, "Active control of sound transmission/radiation from elastic plates by vibration inputs: I. Analysis," Journal of Sound and Vibration, **136**(1), pp.1-15 (1990).

[11] Y. Gu and C.R. Fuller, "Active control of sound radiation due to subsonic wave scattering from discontinuities on fluid-loaded plates. I: Far-field pressure," Journal of the Acoustical Society of America, **90**(4-1), pp.2020-2026 (1991).

[12] C.E. Ruckman and C.R. Fuller, "Numerical simulation of active structural-acoustic control for a fluid-loaded spherical shell," Proceedings of the Second International Congress on Recent Developments in Air- and Structure-borne Sound and Vibration, Auburn University, **1**, pp.377-385 (March 4-6, 1992).

[13] C.G. Mollo and R.J. Bernhard, "Generalized method of predicting optimal performance of active noise controllers," AIAA Journal **27**(11), pp.1473-78 (1989). Note: author's name misspelled as "Molo".

[14] L. Song, Gary H. Koopmann and J.B. Fahnline, "Active control of the acoustic radiation of a vibrating structure using a superposition formulation," Journal of the Acoustical Society of America, **89**(6), pp.2786-2792 (1991).

[15] P.A. Nelson and S.J. Elliott, *The Active Control of Sound,* Academic Press, London (1992).

[16] P.A. Nelson, A.R.D. Curtis, S.J. Elliott and A.J. Bullmore, "The minimum power output of free field point sources and the active control of sound," Journal of Sound and Vibration, **116**(3), pp.397-414 (1986).

[17] P.A. Nelson, A.R.D. Curtis, S.J. Elliott and A.J. Bullmore, "The active minimization of harmonic enclosed sound fields, Part I: Theory," Journal of Sound and Vibration, **117**(1), pp.1-13 (1987).

[18] A.J. Bullmore, P.A. Nelson, A.R.D. Curtis and S.J. Elliott, "The active minimization of harmonic enclosed sound fields, Part II: A computer simulation," Journal of Sound and Vibration, **117**(1), pp.15-33 (1987).

[19] S.J. Elliott, A.R.D. Curtis, A.J. Bullmore and P.A. Nelson, "The active minimization of harmonic enclosed sound fields, Part III: Experimental Verification," Journal of Sound and Vibration, **117**(1), pp.35-58 (1987).

[20] P.A. Nelson, A.R.D. Curtis and S.J. Elliott, "Quadratic optimization problems in the active control of free and enclosed sound fields," Proceedings of the Institute of Acoustics **7**, pp.45-53 (1985).

[21] S.D. Snyder, "A fundamental study of active noise control system design," Unpublished dissertation, Department of Mechanical Engineering, University of Adelaide (1991).

[22] S.D. Snyder and C.H. Hansen, "Using Multiple Regression to Optimize Active Noise Control System Design," Journal of Sound and Vibration, **148**(3), pp.537-542 (1990).

[23] L. Ljung, *System Identification: Theory for the User,* Prentice-Hall, Inc., Englewood Cliffs, New Jersey (1987).

[24] B. Wahlberg and L. Ljung, "Hard Frequency-Domain Model Error Bounds from Least-Squares Like Identification Techniques.," IEEE transactions on automatic control, **37**(7), pp.900-912 (1992).

[25] F. Rosenthal, "A new technique for the active cancellation of wide-band noise using multiple sensors," Proceedings of the Second International Congress on Recent Developments in Air- and Structure-borne Sound and Vibration, **1**, pp.337-344 (March, 1992).

[26] G.C. Everstine and A.J. Quezon, "User's guide to the coupled NASTRAN/Helmholtz equation capability (NASHUA) for acoustic radiation and scattering," CDNSWC-SD-92-17, Carderock Division, NSWC, Bethesda, MD, USA (1992).

[27] M.C. Junger and D. Feit, *Sound, Structures, And Their Interaction, Second Edition,* The MIT Press, Cambridge (1986).

[28] K.S. Miller, "Complex Gaussian processes," SIAM Review, **11**(4), pp.544-567 (1969).

[29] K.S. Miller, "Complex linear least squares," SIAM Review, **15**(4), pp.707-726 (1975).

[30] A.M. Neville and J.B. Kennedy, *Basic Statistical Methods for Engineers and Scientists,* International Textbook Company, Scranton, Pennsylvania (1964).

[31] N.R. Draper and H. Smith, *Applied Regression Analysis,* John Wiley & Sons, Inc., New York (1966).

[32] C.L. Lawson and R.J. Hanson, *Solving Least Squares Problems,* Prentice-Hall, New Jersey (1974).

[33] R.J. Carroll and D. Ruppert, *Transformation and Weighting in Regression,* Chapman and Hall, New York (1988).

[34] C.F. Ross, "Active control of sound," Unpublished doctoral dissertation, University of Cambridge, England (1980).

[35] P. Diaconis and B. Efron, "Computer intensive methods in statistics," Scientific American, pp.116 (May, 1983).

[36] B. Efron, "Computer intensive methods in statistical regression," SIAM Review, **30**(3), pp.421 (1988).

[37] D.A. Belsley, E. Kuh and R.E. Welsch, *Regression diagnostics: Identifying influential data and sources of collinearity,* Wiley, New York (1980).

[38] D.A. Belsley, *Conditioning diagnostics: collinearity and weak data in regression,* John Wiley & Sons, Inc., New York (1991).

[39] W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, *Numerical Recipes: The Art of Scientific Computing, 2nd Ed.,* Cambridge University Press, Cambridge (1989).

[40] B.-T. Wang, R.A. Burdisso and C.R. Fuller, "Optimal placement of piezoelectric actuators for active control of sound radiation from elastic plates," Proceedings of NOISE-CON 91, pp.267-274 (1991).

[41] R.E. Skelton and M.L. DeLorenzo, "Selection of noisy actuators and sensors in linear stochastic systems," Journal of Large Scale Systems, Theory, and Applications, **4**, pp.109-136 (1983).

[42] R.T. Haftka and H.M. Adelman, "Selection of actuator locations for static shape control of large space structures by heuristic integer programming," Computers & Structures, **20**(1-3), pp.575-582 (1985).

[43] A.J. Miller, *Subset Selection in Regression,* Chapman and Hall, New York (1990).

[44] A.J. Miller, "Selection of subsets of regression variables," Journal of the Royal Statistical Society, **147**(3), pp.389-425 (1984).

[45] G.M. Furnival and R.W. Wilson Jr., "Regressions by leaps and bounds," Technometrics, **16**(4), pp.499-511 (1974).

[46] M.A. Efroymson, "Multiple Regression Analysis," Mathematical Methods for Digital Computers, pp.191-203 (1960).

[47] D.M. Smith and J.M. Bremner, "All possible subset regressions using the QR decomposition," Computational Statistics & Data Analysis, **7**, pp.217-235 (1989).

[48] W.M. Gentleman, "Least squares computations by Givens transformations without square roots," Journal of the Institute of Mathematics Applications, **12**, pp.329-336 (1973).

[49] W.M. Gentleman, "Algorithm AS75: Basic procedures for large, sparse, or weighted linear least squares problems," Applied Statistics, **23**(3), pp.448-454 (1974).

[50] A.J. Miller, "Algorithm AS274: Least squares routines to supplement those of Gentleman," Applied Statistics, **38**, pp.458-478 (1991).

[51] J.P. Maillard and C.R. Fuller, "Advanced time domain sensing for structural acoustic systems, II. Radiation control of a simply-supported beam," Journal of the Acoustical Society of America (to be submitted 1993).

[52] J.P. Maillard and C.R. Fuller, "Advanced time domain sensing for structural acoustic systems, I.," Journal of the Acoustical Society of America (to be submitted 1993).

[53] L. Meirovitch, *Analytical Methods in Vibrations,* Macmillan Publishing Co., Inc., New York (1967).

[54] C.E. Ruckman and E.A. Unver, "A Fourier transform for unevenly spaced, complex, periodic sequences," Signatures Directorate Technical Memorandum, Carderock Division, NSWC (1987).

[55] E.G. Williams and H.D. Dardy, "Generalized nearfield acoustical holography for cylindrical geometry:Theory and experiment," Journal of the Acoustical Society of America, **81**(2), p.389 (1987).

[56] F. Fahy, *Sound and Structural Vibration: Radiation, Transmission, and Response,* Academic Press, London (1985).

[57] R.A. Burdisso and C.R. Fuller, "Dynamic behavior of structural-acoustic systems in feedforward control of sound radiation," Journal of the Acoustical Society of America, **92**(1), pp.277-286 (1992).

[58] J.F.M. Scott, "The free modes of propagation of an infinite fluid-loaded thin cylindrical shell," Journal of Sound and Vibration, **125**(2), pp.241-280 (1988).

[59] J.S. Bendat and A.G. Piersol, *Random Data,* John Wiley & Sons, New York (1986).

[60] S.L. Marple Jr., *Digital Spectral Analysis,* Prentice-Hall Inc., Englewood Cliffs, NJ (1987).

[61] G.E.P. Box and G.M. Jenkins, *Time Series Analysis,* Prentice-Hall Inc., Englewood Cliffs, NJ (1976).

[62] J.H. Wilkinson, *The Algebraic Eigenvalue Problem,* Clarendon Press, Oxford (1965).

VITA

Christopher E. Ruckman was born February 19, 1962, while astronaut John Glenn was preparing to orbit the Earth in Freedom 7. Mr. Ruckman lived in Prescott, Wisconsin, and Syracuse, New York before moving to Alexandria, Virginia in 1978. In 1980 he graduated from Hayfield Secondary School in Alexandria, where he was co-captain of the varsity wrestling squad. He attended Virginia Polytechnic Institute and State University (VPI&SU) in Blacksburg, Virginia, where he received a B.S. in Aerospace and Ocean Engineering (1984) and an M.S. in Aerospace Engineering (1986). Mr. Ruckman worked as an acoustical engineer for three years at Carderock Division, Naval Surface Warfare Center (a.k.a. David Taylor Model Basin, formerly David Taylor Research Center, formerly David W. Taylor Naval Ship Research and Development Center, formerly David Taylor Model Basin) before returning to Blacksburg in 1989 to pursue his doctorate. Upon completion of his coursework in Blacksburg he returned to NSWC where he works in the Signatures Directorate.

Mr. Ruckman presently lives in Germantown, Maryland, with his wife, Diane Wilson, and their dog, Rheba Mae Razzbutin. He enjoys music, computers, restoring antique automobiles, woodworking, juggling, hiking, and fishing.