19940020901

N94- 25383

SPACE STATION FLEXIBLE DYNAMICS UNDER PLUME IMPINGEMENT

Final Report

NASA/ASEE Summer Faculty Fellowship Program - 1993

Johnson Space Center

Prepared By: Trevor Williams, Ph.D.

Academic Rank: Associate Professor

University & Department: University of Cincinnati
Department of Aerospace Engineering
Cincinnati, Ohio 45221

NASA/JSC

Directorate: Engineering

Division: Guidance, Navigation & Aeronautics

Branch: Control & Guidance Systems

JSC Colleague: John Sunkel, Ph.D.

Date Submitted: September 14, 1993

Contract Number: NGT-44-001-800

ABSTRACT

Assembly of the Space Station requires numerous construction flights by the Space Shuttle. A particularly challenging problem is that of control of each intermediate station configuration when the Shuttle Orbiter is approaching it to deliver the next component. The necessary braking maneuvers cause Orbiter thruster plumes to impinge on the station, especially its solar arrays. This in turn causes both overall attitude errors and excitation of flexible-body vibration modes. These plume loads are predicted to lead to CMG saturation during the approach of the Orbiter to the SC-5 station configuration, necessitating the use of the station RCS jets for desaturation. They are also expected to lead to significant excitation of solar array vibrations.

It is therefore of great practical importance to investigate the effects of plume loads on the flexible dynamics of station configuration SC-5 as accurately as possible. However, this system possesses a great many flexible modes (89 below 5 rad/s), making analysis time-consuming and complicated. Model reduction techniques can be used to overcome this problem, reducing the system model to one which retains only the significant dynamics, i.e. those which are strongly excited by the control inputs or plume disturbance forces and which strongly couple with the measured outputs. The particular technique to be used in this study is the *subsystem balancing* approach which was previously developed by the present investigator. This method is very efficient computationally. Furthermore, it gives accurate results even for the difficult case where the structure has many closed-spaced natural frequencies, when standard modal truncation can give misleading results. Station configuration SC-5 is a good example of such a structure.

INTRODUCTION

Assembly of the Space Station requires numerous construction flights by the Space Shuttle, regardless of the details of the station configuration that is finally selected. A particularly challenging problem is that of control of each intermediate station configuration when the Shuttle Orbiter is approaching it to deliver the next component. The necessary braking maneuvers cause Orbiter thruster plumes to impinge on the station, especially its solar arrays. This in turn causes both overall attitude errors and excitation of flexible-body vibration modes. These plume loads are predicted to lead to CMG saturation during the approach of the Orbiter to the SC-5 Space Station Freedom (SSF) configuration, necessitating the use of the station RCS jets for desaturation. They are also expected to lead to significant excitation of solar array vibrations.

It is therefore of great practical importance to investigate the effects of plume loads on the flexible dynamics of station configuration SC-5 as accurately as possible. However, this system possesses a great many flexible modes (89 below 5 rad/s), making analysis time-consuming and complicated. Model reduction techniques can be used to overcome this problem, reducing the system model to one which retains only the significant dynamics, i.e. those which are strongly excited by the control inputs or plume disturbance forces and which strongly couple with the measured outputs. A considerable amount of work has been carried out on the topic of model reduction in the past. Well-known methods include *modal truncation* [1], based either on the natural frequencies of the structure or its modal costs, and *balancing* [2] of the entire structure and then truncation to retain a dominant model for it. An advantage of the balancing approach is that it typically yields a more accurate reduced-order model than does simple modal truncation. This is particularly true when the structure possesses clustered natural frequencies, as is often the case for realistic flexible space structures. However, the disadvantages of balancing are its high computational cost, possible numerical sensitivity problems resulting from the large matrices being operated on, and the difficulty involved in providing a physical interpretation for the resulting balanced "modes".

The purpose of the work reported here is to investigate the performance of the alternative *subsystem balancing* technique when used to study the plume impingement problem on SC-5. This method, introduced in [3][4], further developed in [5], and then applied to a simplified SSF model in [12], retains the desirable properties of standard balancing while overcoming the three difficulties listed above. This is achieved by first decomposing the structural model into subsystems of highly correlated modes, based on the *modal correlation coefficients* derived in [4] from the controllability and observability Grammian matrices [6] of the structure. Each subsystem is approximately independent of all others, so balancing each separately and concatenating the dominant reduced-order models obtained yields roughly the same result as balancing the entire structure directly. The computational cost reduction produced by this subsystem technique is considerable: an operation count reduction by a factor of roughly r^2 if the system decomposes into r equal subsystems. The numerical accuracy of the resulting reduced-order model is also improved considerably, as the matrices being operated on are of reduced dimension; this avoids the numerical conditioning problems noted in [8][9] for standard balancing. Furthermore, the modes of the reduced model do now permit a clear physical interpretation. This is a consequence of the fact that each correlated subsystem must necessarily only include modes with close natural frequencies. The balanced modes of each subsystem are therefore, to first order, linear combinations of repeated-frequency modes, and so can themselves be taken as an equally valid set of physical modes. Balancing the entire structure, on the other hand, combines modes of widely differing frequencies, making interpretation difficult.

The graphs and tabulated results given in this report demonstrate the improvements achievable by using the subsystem balancing technique for model reduction. In fact, the original 94-mode model for the SC-5 SSF configuration was reducible to a 10-mode model at the cost of a modeling error of less than 2%. This then allows analysis and simulation of the effects of plume loading to be carried out much more efficiently than if the entire model were used.

THEORETICAL BACKGROUND

Consider an n-mode model for the structural dynamics of a modally damped, non-gyroscopic, non-circulatory FSS with m actuators and p sensors, not necessarily collocated. This model can be written in modal form [1] as

$$\begin{aligned} &\ddot{\boldsymbol{\eta}} + diag(2\zeta_i\omega_i)\dot{\boldsymbol{\eta}} + diag(\omega_i^2)\boldsymbol{\eta} = \hat{B}\mathbf{u}, \\ &\mathbf{y} = \hat{C}_r\dot{\boldsymbol{\eta}} + \hat{C}_d\boldsymbol{\eta}, \end{aligned} \tag{1}$$

where $\boldsymbol{\eta}$ is the vector of modal coordinates, $\mathbf{u}$ that of applied actuator inputs and $\mathbf{y}$ that of sensor outputs, and ω_i and ζ_i are the natural frequency and damping ratio of the i[th] mode, respectively. For the typical FSS [7], the $\{\zeta_i\}$ are quite low (e.g. 0.5 %), and the $\{\omega_i\}$ occur in clusters of repeated, or nearly repeated, frequencies as a result of structural symmetry.

Defining the state vector $\mathbf{x} = (\dot{\eta}_1, \omega_1\eta_1, \cdots, \dot{\eta}_n, \omega_n\eta_n)^T$ for this structure yields the state space representation $\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}, \mathbf{y} = C\mathbf{x}$, where $A = blkdiag(A)_i$, $B = (B_1^T, \cdots, B_n^T)^T$ and $C = (C_1, \cdots, C_n)$, with

$$A_i = \begin{pmatrix} -2\zeta_i\omega_i & -\omega_i \\ \omega_i & 0 \end{pmatrix}, B_i = \begin{pmatrix} \mathbf{b}_i \\ 0 \end{pmatrix} \text{and } C_i = \begin{pmatrix} \mathbf{c}_{ri} & \mathbf{c}_{di} / \omega_i \end{pmatrix}; \tag{2}$$

$\mathbf{b}_i$ is the i[th] row of $\hat{B}$, and $\mathbf{c}_{ri}$ and $\mathbf{c}_{di}$ are the i[th] columns of $\hat{C}_r$ and $\hat{C}_d$, respectively.

The problem studied here is that of obtaining a reduced-order model

$$\begin{aligned} &\dot{\mathbf{x}}_r = A_r\mathbf{x}_r + B_r\mathbf{u}, \\ &\mathbf{y}_r = C_r\mathbf{x}_r \end{aligned} \tag{3}$$

for this structure for which the normalized output error

$$\delta^2 = \int \|\mathbf{y}(t) - \mathbf{y}_r(t)\|_2^2 dt \Big/ \int \|\mathbf{y}(t)\|_2^2 dt \tag{4}$$

is acceptably small. Of course, the size of δ will depend on the order, n_r, chosen for the reduced model. A good model reduction procedure should ideally provide information allowing an intelligent choice for n_r to be made so as to achieve a specified δ value.

Two techniques for model reduction that have been extensively studied are those of modal truncation and internal balancing. The new method implemented in this report, subsystem balancing, can be regarded as a generalization of the two established techniques; it includes both as special cases, but permits a wide range of other solutions also. It is this additional freedom that allows subsystem balancing to produce superior results to those obtainable by means of the traditional methods.

Model reduction by subsystem balancing proceeds by first dividing the given structure into subsystems of highly correlated modes. The dimensions of these subsystems are determined by the choice of a *correlation coefficient threshold* value by the analyst. This quantity is used as follows. Two modes which have nearly equal natural frequencies and non-orthogonal mode shapes will be approximately in resonance, and so can strongly excite each other. An input which excites the first mode will thus typically also indirectly excite the second. It is therefore important to treat the two modes as a single unit, or subsystem, when considering whether or not to retain them in a reduced-order model. Such a pair of close modes can be shown to have a high correlation coefficient; the basis of subsystem balancing is therefore to define subsystems made up of all modes which have correlation coefficients greater than the defined threshold value. It can be seen that selecting a threshold of 0 implies that all modes are deemed to be correlated, i.e. there will be a single "subsystem" which is in fact the entire system. This is equivalent to the case of standard balancing. Conversely, choosing a threshold of 1 leads to all modes being deemed to be uncorrelated, and so dealt with in isolation: this amounts to the modal truncation technique.

Each of the subsystems obtained for the chosen threshold value is then balanced independently, and a reduced-order model for it generated by deleting all balanced states corresponding to *Hankel singular values* [2] below some specified threshold. Note that the singular value weighting described in [10] could be applied, if desired, without changing the argument in any way. Similarly, frequency weighting of the Hankel singular values can easily be incorporated to deal with input signals which have a known frequency spectrum. This is actually done in the present application, where the inputs are steps (representing thruster firings) rather than the impulses classically considered in model reduction problems. The resulting reduced-order subsystem models so obtained are then combined to yield a dominant, approximately balanced, reduced-order model for the full system.

USER INTERFACE

This section describes the user interface to the model reduction package which was developed as part of this contract. This software consists of a library of Matlab m-functions, with mrmain3 calling all the other functions internally. The package is installed on the Sun SparcStation 2 *carmel* in the Interactive Analysis Laboratory in Building 16, and has also been provided in a Macintosh version to personnel at JSC and Dynacs Engineering Co. The information that follows details the user interface for mrmain3. This function, together with the purpose-built second level functions that it calls, have extensive in-line documentation, facilitating future use and/or modification. Listings of all these functions are provided in the report [12].

Input arguments

om: The natural frequencies (rad/s) of the structure, input as either a row or column vector. Any rigid-body modes must precede the flexible modes and be represented by hard zero frequencies.

phia: The influence matrix, in mass-normalized coordinates, corresponding to the specified actuator locations. If the structure has n modes and m actuators, *phia* will be an (n x m) matrix.

phis: Similar to *phia*, but for sensor stations or positions of outputs of interest (e.g. solar array tips).

Prompts to user

The following prompts are generated by mrmain3 when running. They allow each run to be customized as desired by the user.

Output the time taken for each step?: The time required for each matrix decomposition, etc., is output to the screen if requested. This allows the progress of the model reduction procedure to be monitored, as well as giving an indication of which steps are the most computationally intensive.

Vectorize? (Faster, but requires more storage): In Matlab, `for` loops are typically an order of magnitude slower to execute than the equivalent "vectorized" operation. For instance, `s=0; for i=1:n, s=s+x(i); end;` runs considerably slower than does `s=x*ones(n,1)`. If vectorization is requested, computation of the system Grammian matrices and correlation coefficients is put into the form of vector-matrix operations rather than loops; this is indeed considerably faster, but requires some additional temporary storage arrays.

Structural damping ratio, % (default is 0.5%): The specified damping ratio is applied uniformly to all flexible modes of the full structural model.

Print frequencies in Hz?: The mean frequency of each subsystem can be output in either rad/s or Hz, as desired.

Desired controllability threshold?: This threshold value is used to determine which modes are correlated in a controllability sense. The system is then broken down into disjoint sets of modes (subsystems), where modes with a controllability correlation coefficient greater than the specified threshold are deemed to be correlated. Taking a threshold value of 0 implies that all modes are considered correlated, i.e. the method reduces to standard balancing. Conversely, a threshold value of 1 implies that no modes are taken together: this is modal truncation. Intermediate values allow the dimensions of the resulting subsystems to be specified to a large extent; reducing the threshold reduces the number of subsystems, so increasing their dimension. The program allows the user to try various different values for this threshold until he is satisfied with the resulting set of subsystems. At this point, a negative value is input to cause the program to proceed to the next step.

Desired overall threshold?: This threshold is used in a similar fashion to the controllability threshold, but both controllability and observability are now taken into account. This yields the final subsystem distribution output by the program (in *modemap*) and used to obtain the reduced-order model.

Compare step responses?: If requested, the step responses of the full and reduced-order models are computed, plotted, and the relative differences (i.e. reduced-order model error) output for each input-output channel.

Compare frequency responses?: If requested, the Bode amplitude plots of the full and reduced-order models are similarly computed and plotted, and the relative differences (i.e. reduced-order model error) output for each input-output channel.

Desired truncation measure?: Two types of measure can be used to define the number of modes retained in the reduced-order model. If a positive integer is input, this is taken to be precisely the desired reduced model order. On the other hand, if a real number in the interval $[0, 1)$ is input, this is taken to be the desired relative error in the reduced-order model step response, and the model order required to achieve this is computed. (Note that this second option is only an approximation, and should be treated as such.) The program loops over this step, allowing the user to try a sequence of various different reduced model orders to determine which one is most satisfactory for his purposes. This is quite a rapid operation when using mrmain3, as nearly all computations need only be done once. Essentially the only step that must be repeated when trying a new model order is that of calculating the step and/or Bode plots of the reduced-order model.

Output arguments

am, bm, cm: The reduced-order state-space model matrices $\{A_r, B_r, C_r\}$ obtained.

modemap: This matrix specifies which physical modes are grouped into which subsystems in the decomposition based on overall correlation coefficients. In particular, the i^{th} column of *modemap* lists the modes making up the i^{th} of these subsystems.

DISCUSSION OF RESULTS

Results will now be provided which illustrate the behavior of the subsystem balancing technique when applied to a structural model [11] of the SC-5 configuration of Space Station Freedom. This structure possesses light damping (estimated to be 0.5% of critical), and a large number of closely-spaced vibration modes (94 flexible modes in the model considered, of which 89 are below 5 rad/s). The model has 12 inputs: 6 Reaction Control System (RCS) thrusters, and 6 disturbance inputs which represent the lumped forces and torques about all three axes that result from Shuttle Orbiter plume impingement during approach. Note that the Control Moment Gyroscopes (CMGs) that would also be used on SC-5 do not need to be considered as inputs here: this is because CMGs incorporate torque shaping techniques so as to not significantly excite flexible vibration modes. The measured outputs are the 3 angular rates sensed by the rate gyros on the station. (The movements at other positions of interest, for instance the solar array tips, could also be considered if desired; the method remains exactly the same.)

The three numerical parameters which the user must select when running mrmain3 are: the controllability coefficient threshold value; the overall (controllability and observability combined) coefficient threshold; and the order of the final reduced-order model. It should be noted that, as previously pointed out, the program allows the user to try a sequence of values for each of these quantities until satisfactory results are obtained. The effects of different choices for these three parameters will now be examined.

Figure 1 illustrates the role of the controllability threshold coefficient ρ_{ct} in determining the subsystem decomposition. The two solid graphs give the maximum and minimum subsystem dimensions that result for values of ρ_{ct} ranging between 0 (standard balancing) and 1 (modal truncation). The dashed graph then shows the number of subsystems obtained for each threshold value. It can be seen that the system decomposition does indeed change, as expected, from that of standard balancing (the only subsystem is the entire model, with 94 modes) to that of modal truncation (there are 94 single-mode subsystems) as the threshold increases from 0 to 1. It can be noted that the evolution of subsystem dimensions is somewhat discontinuous: for instance, large changes occur for thresholds between 0.05 and 0.15, whereas there are hardly any differences between 0.55 and 0.70. A consequence of this is that it is not always possible to find a threshold value which will yield a particular maximum subsystem order. However, it *is* always possible to obtain a good working value which gives a totally acceptable subsystem partition. From practical experience, it has been observed that a good choice of threshold value is generally one which gives the maximal subsystem dimension approximately equal to the number of subsystems. The reason for this is that it provides a good balance between having subsystems which are too large (and so susceptible to numerical problems) and ones which are too small (and so neglect appreciable cross-mode interaction). For the system studied here, such a choice can be seen to be $\rho_{ct} = 0.2$, giving 27 individual subsystems (some of which consist of single modes), and a maximum subsystem dimension of 17.

Figures 2-4 similarly illustrate the role of the overall correlation coefficient threshold ρ_{ot}. These figures can be seen to be broadly similar to Figure 1. The main difference is that the choice made for ρ_{ct} affects the initial, controllability-based, subsystem decomposition, which is then used to compute the final subsystem decomposition based on ρ_{ot}. Figures 2-4 demonstrate this by letting ρ_{ot} run through all possible values between 0 and 1 for ρ_{ct} equal to 0.1, 0.2 or 0.3, respectively. The differences between the three plots can clearly be seen: the main distinction is that the low controllability threshold value of 0.1 tends to give rise to larger final subsystems. This makes sense, as a lower threshold leads to more modes being considered sufficiently correlated that they must be grouped in the same subsystem. A rule-of-thumb for overall threshold that has generally been found to work well is simply to take $\rho_{ot} = \rho_{ct}$.

Table 1 then quantifies the effect of different choices for the reduced model order n_r. In all cases, the controllability and overall thresholds are taken to be equal, $\rho_{ot} = \rho_{ct} = \rho_t$, and this value is allowed to range over values between 0 (balancing) and 1 (modal truncation). The numbers that are tabulated are the step response relative error norms, as defined by (4), for the resulting reduced-order models. It is significant to note that the new subsystem balancing technique gives better results than those of standard balancing or modal truncation for all the reduced orders considered. For instance, reducing the original 84-mode system to a 10-mode model by subsystem balancing with $\rho_t = 0.2$ gives rise to a relative error of less than 2%. By contrast, the error obtained for a model of the same dimension by means of balancing is over twice as large, and that produced by modal truncation is nearly three times as great.

TABLE 1.- STEP RESPONSE RELATIVE ERRORS

n_r	$\rho_t = 0$ Balanc	$\rho_t = .1$	$\rho_t = .2$	$\rho_t = .3$	$\rho_t = .4$	$\rho_t = .5$	$\rho_t = .6$	$\rho_t = .7$	$\rho_t = 1$ Modal
10	0.0416	0.0235	**0.0196**	**0.0196**	0.0218	0.0198	0.0233	0.0233	0.0556
11	0.0370	0.0226	0.0169	0.0169	0.0191	**0.0164**	0.0203	0.0203	0.0408
12	0.0188	0.0230	0.0161	0.0161	0.0138	**0.0117**	0.0203	0.0203	0.0179
13	0.0199	0.0207	0.0125	0.0125	0.0136	0.0117	**0.0082**	**0.0082**	0.0153
14	0.0096	0.0215	0.0129	0.0129	0.0145	**0.0057**	0.0058	0.0058	0.0080
15	0.0070	0.0180	0.0072	0.0108	0.0094	0.0071	**0.0050**	**0.0050**	0.0058
16	0.0073	0.0081	0.0067	0.0067	0.0091	0.0072	**0.0050**	0.0051	0.0058
17	0.0070	0.0081	0.0064	0.0063	0.0087	**0.0041**	**0.0041**	**0.0041**	0.0049
18	0.0049	0.0046	0.0038	0.0037	0.0042	**0.0035**	0.0036	0.0036	0.0046
19	0.0052	0.0050	0.0036	0.0036	0.0040	**0.0032**	0.0033	0.0033	0.0044
20	0.0044	0.0050	0.0042	0.0042	0.0041	0.0043	**0.0040**	**0.0040**	0.0052
21	0.0044	0.0052	0.0045	0.0045	0.0043	0.0045	**0.0036**	0.0037	0.0060
22	0.0043	0.0041	0.0035	0.0035	**0.0030**	0.0035	0.0040	0.0040	0.0079
23	0.0034	0.0034	0.0033	0.0033	**0.0029**	0.0032	0.0031	0.0031	0.0059
24	0.0038	0.0033	0.0026	0.0026	**0.0024**	0.0025	0.0032	0.0031	0.0039
25	0.0026	0.0020	0.0020	0.0020	0.0018	0.0018	**0.0016**	**0.0016**	0.0029

Figure 5 illustrates this further by plotting the step responses of both the full system and the 10-mode reduced-order model obtained for $\rho_t = 0.2$. It can be seen that the plots are nearly indistinguishable, despite the nearly order-of-magnitude reduction of system dimension. Figure 6 confirms this by plotting the step response error between the reduced model and the original system: note the different y-axis scale factors. For comparison, Figure 7 gives the error that would be obtained for a 10-mode model generated using modal truncation; again note the change in y-axis scale. The increase in this error over that obtained by subsystem balancing is clear. Finally, Figure 8 overlays the frequency response plots of the full system and the 10-mode reduced model obtained by subsystem balancing for $\rho_t = 0.2$. (These plots are weighted by a ω^{-1} term, so as to reflect the fact that a step input is being applied.) The close match between the two models is clearly visible in the frequency domain, just as it was in the corresponding time histories. It is also easy to identify which flexible modes are strongly excited by a prolonged, step-like, plume excitation force.

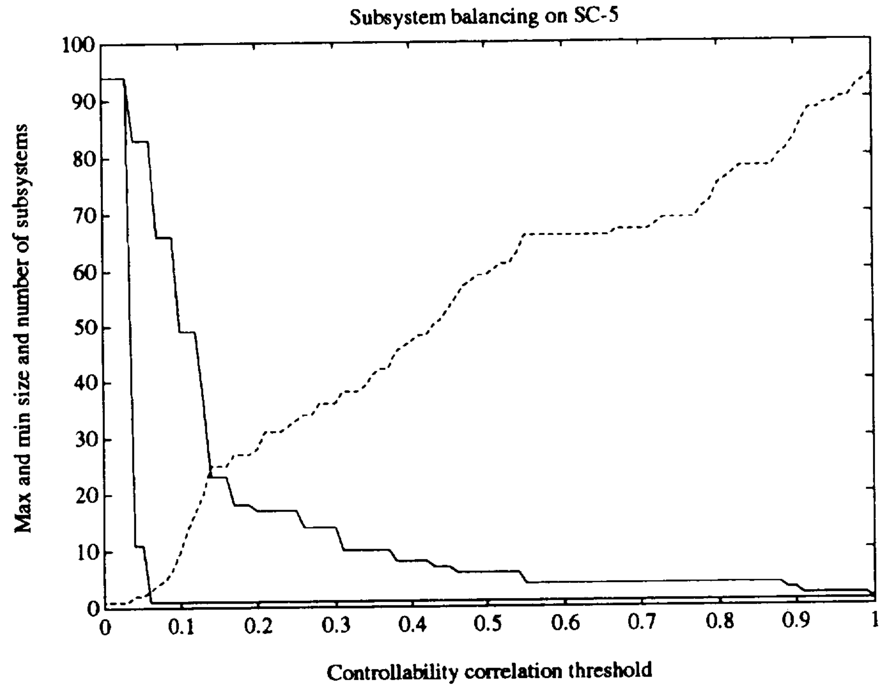

Figure 1.- Effect of varying controllability threshold level.

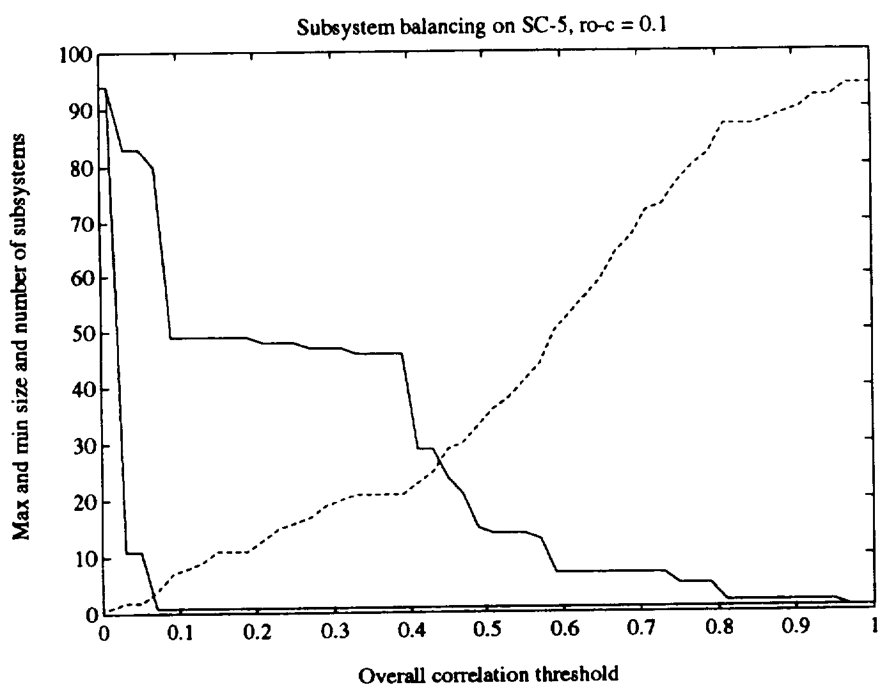

Figure 2.- Effect of varying overall threshold level for $\rho_{ct} = 0.1$.

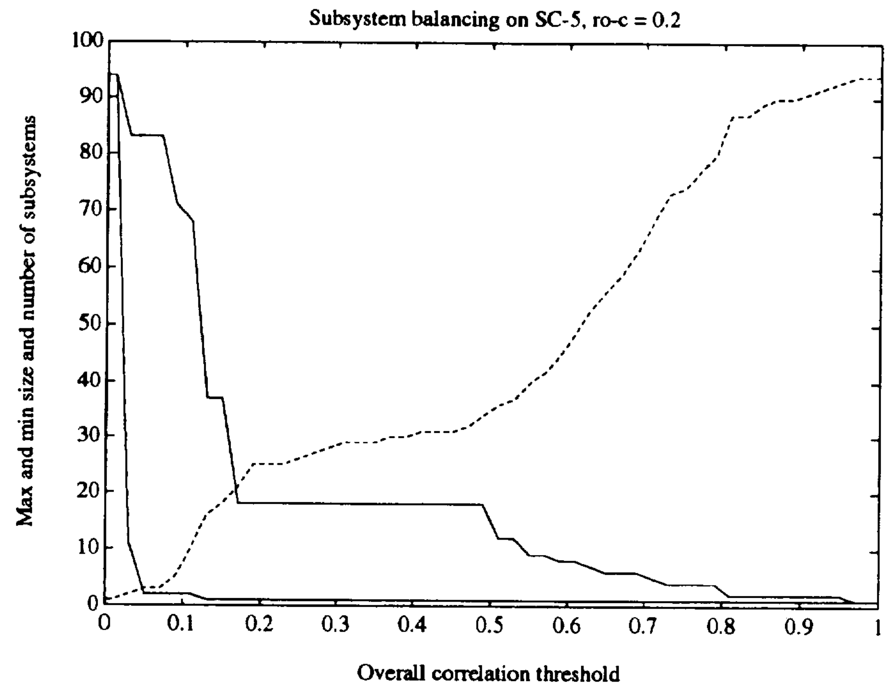

Figure 3.- Effect of varying overall threshold level for $\rho_{ct} = 0.2$.

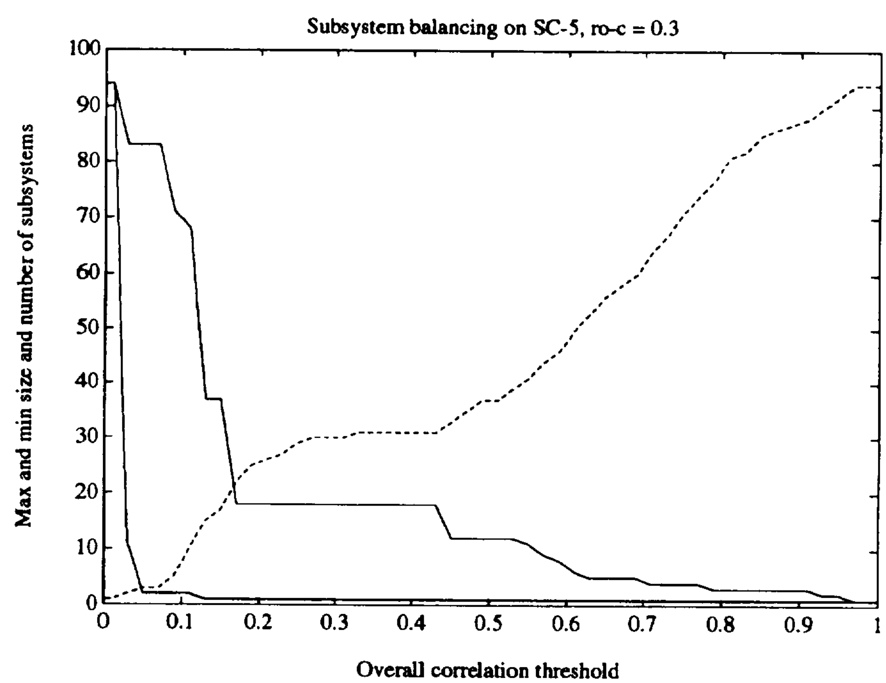

Figure 4.- Effect of varying overall threshold level for $\rho_{ct} = 0.3$.

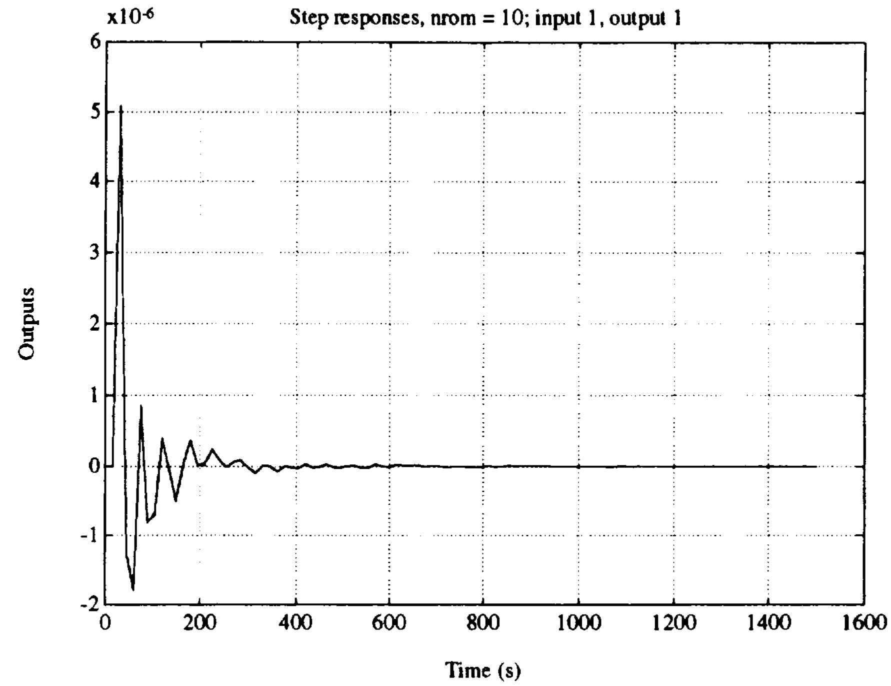

Figure 5.- Step responses of full and reduced models.

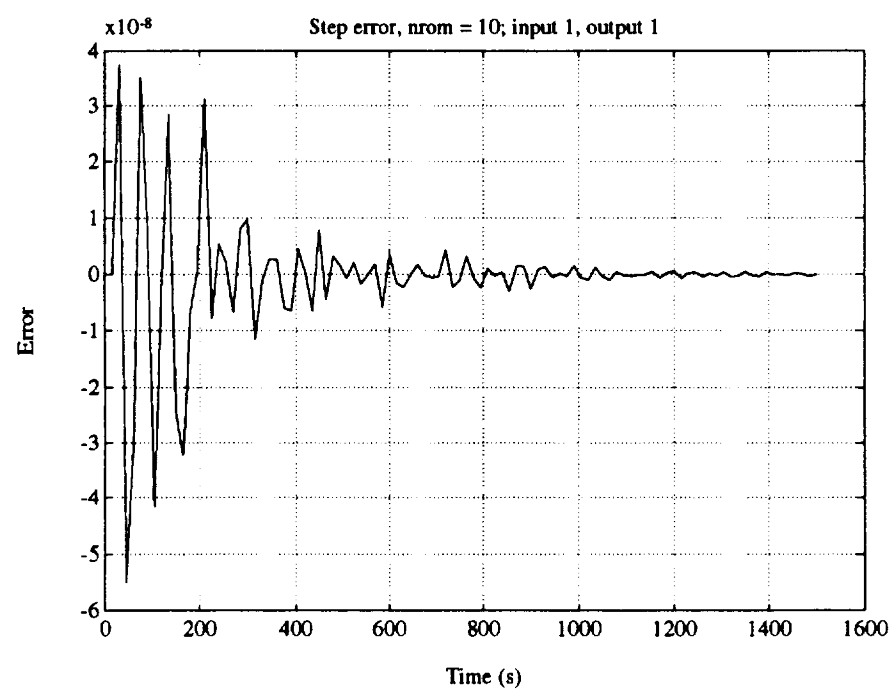

Figure 6.- Step response error of reduced model.

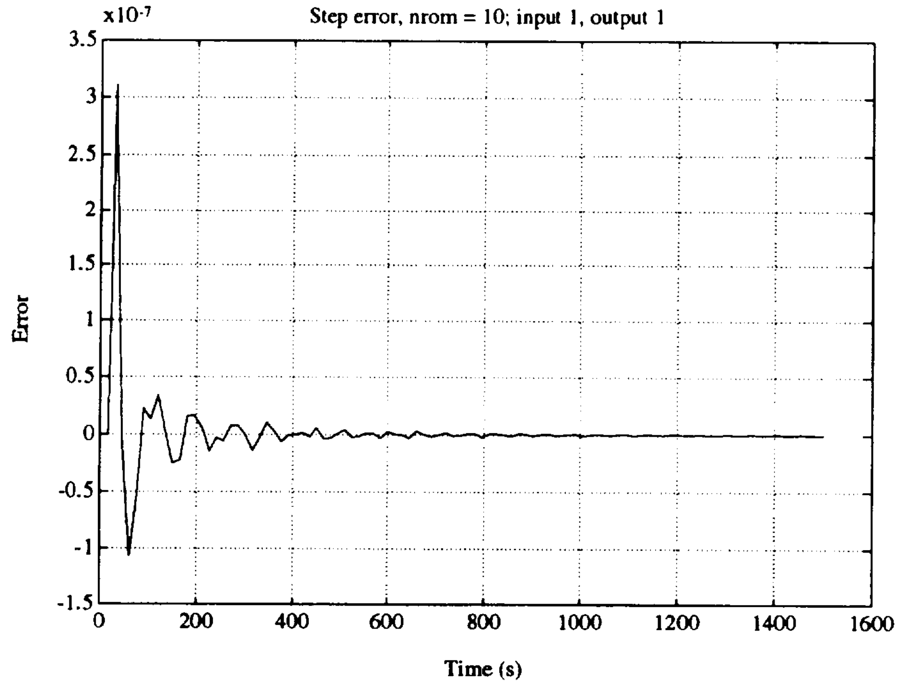

Figure 7.- Step response error of modal truncation model.

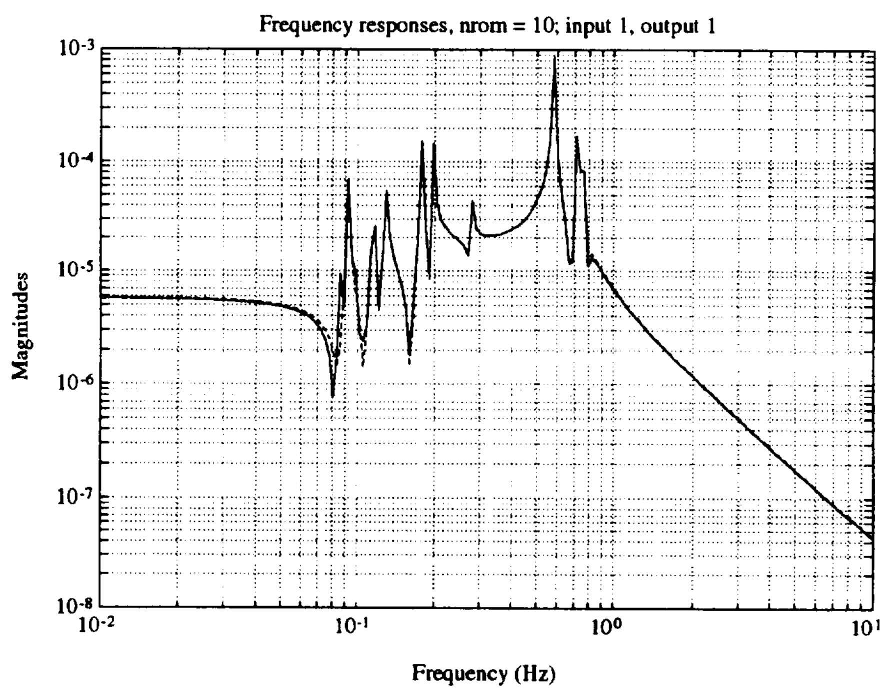

Figure 8.- Frequency responses of full and reduced models.

REFERENCES

1. R.R. Craig, *Structural Dynamics*, New York: Wiley, 1981.

2. B.C. Moore, 'Principal Component Analysis in Linear Systems: Controllability, Observability, and Model Reduction', *IEEE Transactions on Automatic Control*, Vol. 26, Feb. 1981, pp. 17-32.

3. T.W.C. Williams and W.K Gawronski, 'Model Reduction for Flexible Spacecraft with Clustered Natural Frequencies', invited paper, 3rd NASA/NSF/DoD Workshop on Aerospace Computational Control, Oxnard, CA, Aug. 1989.

4. W.K Gawronski and T.W.C. Williams, 'Model Reduction for Flexible Space Structures', *Journal of Guidance, Control, and Dynamics*, Vol. 14, Jan.-Feb. 1991, pp. 68-76.

5. T.W.C. Williams and M. Mostarshedi, 'Model Reduction Results for Flexible Space Structures', presented at 5th NASA/DoD Control-Structures Interaction Technology Conference, Lake Tahoe, Nevada, Mar. 1992.

6. T.W.C. Williams, 'Closed-Form Grammians and Model Reduction for Flexible Space Structures', *IEEE Transactions on Automatic Control*, Vol. 35, Mar. 1990, pp. 379-382.

7. M.J. Balas, 'Trends in Large Space Structure Control Theory: Fondest Hopes, Wildest Dreams', *IEEE Transactions on Automatic Control*, Vol. 27, 1982, pp. 522-535.

8. M.S. Tombs and I. Postlethwaite, 'Truncated Balanced Realization of a Stable Non-Minimal State-Space System', *International Journal of Control*, Vol. 46, 1987, pp. 1319-1330.

9. M.G. Safanov and R.Y. Chiang, 'A Schur Method for Balanced Model Reduction', *Proc. American Control Conference*, 1988, pp. 1036-1040.

10. R.E. Skelton and P. Kabamba, 'Comments on "Balanced Gains and Their Significance for L^2 Model Reduction"', *IEEE Transactions on Automatic Control*, Vol. 31, Aug. 1986, pp. 796-797.

11. H.T. Chang, 'SSF SC-4 Attitude Control Simulation', Dynacs Engineering Co. Report DYNACS-93-HM16, May 24, 1993.

12. T.W.C. Williams, 'Model Reduction for Space Station Freedom', NASA/ASEE Summer Faculty Fellowship Program Final Report, Aug. 1992.