

ARPA/S30M2/93/01
MRC/AST/FSR/11.0

**ADVANCED SUBMARINE TECHNOLOGY - PROJECT M
CONTROL THEORY REPORT**

M.A. Swinbanks
S. Daley



GEC-Marconi Research Centre

1993
Interim Technical Report

ARPA Order Number: 7813
Contract Number: N00014-93-C-0063
Effective date of contract: 04 January 1993
Expiration date of contract: 04 March 1995
Principal Investigator: Dr F A Johnson (tel: +44 245 473331)

Advanced Research Projects Agency
Arlington, Va.
Office of Naval Research
Department of the Navy
Arlington, Va.

This document has been approved for public release and sale; its distribution is unlimited.

© 1993 GEC-MARCONI LTD The copyright in this published document is the property of GEC-Marconi Limited. The information and data contained herein is provided without responsibility and GEC-Marconi Limited disclaims all liability arising from its use.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency of the U.S. Government.

Limited rights shall be effective until 31st January 1997, thereafter the limited rights shall expire and the Government shall have unlimited rights in the technical data.

The restrictions governing the use and disclosure of technical data marked with this legend are set forth in the definition of "limited rights" in paragraph (a)(15) of the clause at 252.227-7013 of the contract listed above.

19941207 003

ARPA/S30M2/93/01
MRC/AST/FSR/11.0

ADVANCED SUBMARINE TECHNOLOGY - PROJECT M CONTROL THEORY REPORT

Prepared by:
M.A. Swinbanks (Programme Consultant)
S. Daley

GEC-Marconi Research Centre
GEC-Marconi Limited
Great Baddow, Chelmsford
Essex, U.K., CM2 8HN

September 1993

ARPA Order Number: 7813
Contract Number: N00014-93-C-0063
Effective date of contract: 04 January 1993
Expiration date of contract: 04 March 1995
Principal Investigator: Dr F A Johnson (tel: +44 245 473331)

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Avail	Unannounced
Avail	Unannounced
A-1	

Prepared for
Advanced Research Projects Agency
Arlington, Va.
Office of Naval Research
Department of the Navy
Arlington, Va.

© 1993 GEC-MARCONI LTD The copyright in this published document is the property of GEC-Marconi Limited. The information and data contained herein is provided without responsibility and GEC-Marconi Limited disclaims all liability arising from its use.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency of the U.S. Government.

Limited rights shall be effective until 31st January 1997, thereafter the limited rights shall expire and the Government shall have unlimited rights in the technical data.

The restrictions governing the use and disclosure of technical data marked with this legend are set forth in the definition of "limited rights" in paragraph (a)(15) of the clause at 252.227-7013 of the contract listed above.

Abstract

This interim report presents the broad band control theory aspects of a programme whose objective is to control the dynamics of a large, lightweight marine machinery raft supported on an array of electromagnetic actuators. The report shows how it is possible to achieve, simultaneously, the normally conflicting requirements of accurate alignment of the raft under sea-way or manoeuvring motions, of the elimination of mount resonances, of stiffening long wavelength bending resonances without incurring a radiation hazard, of decoupling shorter wavelength bending resonances and of greatly reducing any residual force transmission through the mounts at higher frequencies. Simulations of the control theory applied to a simple model raft, show that very large performance improvements are possible. These findings go some way to giving the assurance that large, lightweight machinery rafts could be used in future installations without having to make any compromises in overall performance. This, in turn points to a new way for reducing overall costs.

Contents

Abstract	ii
Summary	v
1 Introduction	1
1.1 General Perspective	1
1.2 Specific Design Objectives	3
1.3 Sensors and Performance Criteria	3
2 General Algorithm Principles	5
2.1 Physical Principles	5
2.2 General Algorithm Characteristics	7
2.3 Low Frequency Control Algorithm	9
2.3.1 Modal Gaps	12
2.3.2 Non-Reciprocal Criterion	12
2.3.3 Flexural Mode Accelerations	13
2.3.4 Spatial Control Force Distribution	13
2.3.5 Decoupling of Rigid Body and Flexural Modes of Motion	14
2.4 High Frequency Control Algorithm	15
3 Low Frequency Controller	18
3.1 System Model	18
3.2 Rigid Body Controller	19
3.3 Flexing Mode Controller	21
3.4 Summary of Low Frequency Control Algorithm	22
4 High Frequency Control Algorithm	24
4.1 Basic Feedforward Adaptive Algorithm	24
4.2 Correction for Downstream Propagation	25
4.3 Deconvolution of Downstream Interaction Paths	25
4.4 Prevention of Interaction with Low-Frequency Algorithm	27
4.5 The Overall High Frequency Algorithm	29
5 Numerical Calculations and Simulations of Performance	31
5.1 Theoretical Basis for Raft Dynamical Model	31
5.2 Off-Line Frequency Response Calculations	32
5.3 Simulations of Real-Time Low-Frequency Response	41
5.4 Simulation of High Frequency Algorithm	46
6 Conclusions	50
7 References	52
Distribution	53
REPORT DOCUMENTATION PAGE	54

Figures

Figure 1.1: Typical Marine Raft Installation	1
Figure 1.2: Experimental Raft Structure	3
Figure 2.1: ZBD Control	7
Figure 2.2: Low Frequency Algorithm	8
Figure 2.3: High Frequency Algorithm	8
Figure 2.4: Overall Algorithm	9
Figure 2.5: Model Reference State-Space	10
Figure 2.6: Modified Observation Process	11
Figure 2.7: Detail of Observation Process	12
Figure 2.8: Front and Rear End Processes	14
Figure 4.1: Basic FeedForward Algorithm	24
Figure 4.2: Filtered-X Adaptive Algorithm	25
Figure 4.3: Multi-Degree-of-Freedom Adaptive Algorithm	26
Figure 4.4: Adaptive Algorithm with Low Freq. Constraint	28
Figure 4.5: Adaptive Algorithm with Simplified LF Constraint	29
Figure 4.6: Overall High Frequency Algorithm	30
Figure 5.1: Experimental Raft Structure	31
Figure 5.2: Mean Square Acceleration - Soft Model (Off-line)	33
Figure 5.3: Mean Square Acceleration - Med. Model (Off-line)	34
Figure 5.4: Mean Square Acceleration - Stiff Model (Off-line)	34
Figure 5.5: Wavenumber Distribution - Soft Model (Off-line)	36
Figure 5.6: Wavenumber Distribution - Med. Model (Off-line)	37
Figure 5.7: Wavenumber Distribution - Stiff Model (Off-line)	38
Figure 5.8: Typical Wavenumber Cost Function	39
Figure 5.9: Radiated Noise Performance - Soft Model	40
Figure 5.10: Radiated Noise Performance - Medium Model	41
Figure 5.11: Response to Seaway Disturbance	42
Figure 5.12: Mean Sq. Acceleration - Soft Model (Real-time)	43
Figure 5.13: Mean Sq. Acceleration -Medium Model (Real-time)	44
Figure 5.14: Mean Sq. Acceleration - Stiff Model (Real-time)	44
Figure 5.15: Wavenumber Distribution - Soft Model (Real-time)	45
Figure 5.16: Wavenumber Distribution - Med. Model (Real-time)	45
Figure 5.17: HF Algorithm - Soft Raft	47
Figure 5.18: HF Algorithm - Medium Raft	47
Figure 5.19: HF Algorithm - Stiff Raft	48

Summary

Objectives

The overall objective of the research described in this report has been to investigate the dynamics of broad-band active control applied to large, lightweight marine machinery rafts with distributed active mounts. The performance of large, passively mounted marine machinery rafts is limited by the presence of low-frequency internal structural resonances, which can compromise the isolation effectiveness of the installation. In order to overcome these effects, such rafts have to be of stiff and massive construction. The use of active techniques can provide selective dynamic stiffening and damping, which should permit much more flexible, lighter structures to be constructed. Moreover, by appropriate choice of the control force distributions, very substantially improved isolation can be achieved while still providing the necessary stiffening and damping actions.

Such active mountings can also be used to provide accurate raft alignment under seaway and manoeuvring disturbance, and to suppress fundamental raft-mounting resonances. At higher frequencies, local control strategies can cancel the residual vibration transmission through each mount and thus provide enhanced isolation performance.

Thus the capability of an actively controlled distributed raft extends over a wide frequency range, from low-frequency alignment, through suppression of mounting and internal raft resonances, through to decoupling of higher frequency resonances and local high frequency vibration suppression.

The purpose of this report has therefore been to establish the performance achievable by a broad-band, distributed active mounting system over the frequency range from DC to 500Hz, which covers the entire set of effects described.

Specifically, broad-band active control of a distributed mounting system should simultaneously achieve the following objectives:

- (1) Greatly improved machinery alignment under seaway or manoeuvring motions
- (2) The elimination of conventional mounting resonances
- (3) The stiffening of the raft at low frequencies in order to eliminate internal structural resonances, while simultaneously minimising acoustic radiation.
- (4) The decoupling of higher frequency structural resonances, so that the mounting system does not respond to their effects.
- (5) Greatly improved local vibration isolation at high frequencies.

If the above objectives can be achieved, this provides the opportunity to construct large, lightweight platforms which can provide very high levels of isolation for any arbitrary items of machinery mounted on them. This in turn should enable more relaxed vibration specifications to be applied to those machinery items, with subsequent reduction in their cost.

Critical Questions

Five critical questions were considered:

- (1) What are the principal dynamical features associated with the response of large machinery rafts.
- (2) What control strategies are required to deal with these separate features.
- (3) Can an overall low-frequency feedback control strategy be formulated to suppress the structural motions while simultaneously providing high levels of vibration isolation.
- (4) Can a local high-frequency feedforward control strategy suppress the broad-band vibration transmission through individual mounts, while taking into account interactive effects associated with multiple transmission paths.
- (5) Can the separate low-frequency and high-frequency control algorithms be integrated into an overall control algorithm, with smooth transition between the control regimes, and without interaction between the separate algorithms.

Approach

The approach adopted has been to examine the behaviour of a flexible rafted structure consisting of coupled longitudinal and transverse beams, to define the overall system dynamics. A low-frequency state-space control strategy has then been formulated, to describe the modal response of this structure. This control strategy is a global feedback strategy, which senses the entire structural motions and feeds back appropriate control forces to drive the motions in the required manner.

The optimisation of this control strategy involves the application of appropriate cost criteria, to represent the trade-offs associated with the specific design requirements, namely to suppress the unwanted structural motions while maximising the acoustic isolation of the installation.

These cost criteria must include not only the frequency regimes over which the various constraints apply, but also the precise way in which the control forces are spatially distributed. This latter aspect is very important, since it is the spatial distribution of the control forces which defines the efficiency with which these forces couple into the receiving hull structure and hence radiate into the sea.

A central feature of the control strategy is that all the performance requirements are condensed into a specific set of spatial/temporal filters which are applied to the control and observation signals. The outputs of these filters are used to define a new set of control and observation parameters, which then enable a standard optimisation procedure to be applied to determine the overall control strategy. In this way, the designer's performance requirements can be fully accounted for within the specification of these space-time filters, while the subsequent determination of the optimal control strategy then becomes a routine, fully-defined procedure which is common to all requirements.

The high-frequency control algorithm has been formulated as a local feedforward control algorithm, which senses the local motion of the raft and defines appropriate output control forces to minimise transmission into the receiving base-structure. Account is taken of the finite transmission time between introduction of the control actions and observation of the base response. Moreover, the effects of interaction from multiple transmission paths are also taken into account, thus enabling the individual transmission paths to be separately optimized.

The two distinct algorithms are then combined into a single overall algorithm; in order to achieve this, a method of fully decoupling the actions of the low-frequency and high frequency algorithms has been defined.

In order to assess the performance, off-line calculations of the expected frequency responses of the low-frequency control algorithm were carried out. The overall system model together with its associated controller were then simulated as a time-stepped process, enabling the stability and performance of the algorithm to be confirmed. The high frequency algorithm was then introduced into this same simulation, so that the ability of this algorithm to suppress local vibration transmission paths could be determined.

These simulations enabled the entire range of design requirements to be reproduced. The performance of a conventional 4Hz passively mounted raft was simulated as a baseline, so that the achievable improvements in performance associated with the active system could then be quantified.

Results

Simulations performed with the low-frequency algorithm have shown that it is possible to satisfy the apparently conflicting requirements of actively aligning and stiffening the raft structure, while maintaining very high levels of isolation. In this respect, the objective becomes to constrain the raft to move primarily as a rigid body, so that the entire mounted-mass then becomes a very efficient absorber of the momentum associated with unwanted vibration excitation forces. This whole mass can then move accurately to absorb unwanted momentum with high precision.

It has been shown that the use of a dual sensing capability, namely observation of the gap between the raft and the hull, together with separate sensing of the raft acceleration response, enables cross-processing to be performed which discriminates hull-seaway induced motions from raft vibration. This enables separate, non-reciprocal responses to be obtained for these different types of disturbance. As a result large control forces can be applied to enable the raft to track seaway motions very accurately, yet it still remains very soft down to very low frequency in response to vibration inputs.

In the context of controlling the very low-frequency flexural modes of the structure, it has been shown that the associated low-frequency spatial control force distributions can be chosen to have preferentially high wave-number content. This means that when coupling into the hull structure, the array of control forces is largely self-cancelling and represents a very poor source of radiation. At the same time, these control forces can still be chosen to maintain adequate leverage over the resonant modal response of the raft.

As the frequency is increased, it is inappropriate to continue to constrain the motions of the raft; the correct strategy becomes to filter out and decouple the higher frequency resonances. The modal control strategy enables this to be accomplished, so that the control system does not respond at all to these motions, and as a consequence they no longer represent a transmission hazard.

The high frequency algorithm has been shown to be very effective in cancelling out unwanted local leakage-path transmission (associated with electromagnetic eddy current or other forms of residual mount transmission) It has been shown that without such control, the high frequency performance of the installation can be significantly degraded. Introduction of the high frequency control strategy enables these effects to be overcome so that very good vibration isolation can be maintained at each individual mount.

The low-frequency and high-frequency algorithms have been simulated together in parallel, and it has been shown that they can co-exist without interaction. Thus the overall system correctly makes the transition between low-frequency global feedback control, and high frequency local feedforward control.

The overall performance which can be achieved with such a broad-band active control system is well in excess of that offered by an equivalent passive installation. The alignment of the raft relative to the hull can be made extremely precise, so that relative raft-hull displacements are constrained to 1mm or less, in contrast to the 10mm motions associated with passive mounts. Low frequency mounting resonances can be completely suppressed, and 20dB reductions achieved in the acceleration response due to low-frequency raft structural resonances.

At the same time, the corresponding vibration isolation is substantially better than a passive installation, with theoretical performance improvements in excess of 40dB.

At the higher frequencies, the raft is allowed to "run free", so that high frequency resonances do not transmit through the control system. Thus, in theory, there should be no vibration transmission at all associated with these modes. In practice, local leakage paths through the mounts will compromise this isolation, but the high frequency adaptive algorithm can subsequently provide 30-40dB suppression of such transmission paths.

Conclusions

It has been shown that the use of a distributed array of active mounts can enable large, flexible lightweight machinery rafts to be constructed which offer substantial performance improvements over equivalent stiff and massive, passively mounted installations. This improvement is achieved by accurately controlling the motions of the structure, so that it moves correctly to absorb the unwanted momentum associated with vibration forces, while unwanted resonant motions are constrained and suppressed.

By using dual detection strategies, namely raft-hull gap sensors, and raft-acceleration sensors, it is possible to achieve separate dynamical response to hull-seaway and raft-vibration disturbances. Thus the raft installation can track very accurately the hull-seaway and manoeuvring disturbances, yet still remain very soft down to very low frequencies in response to vibration force inputs. At the same time, all resonances associated with the fundamental raft-mounting dynamics can be completely suppressed.

The correct use of distributed spatial control force arrays can enable large-scale, low-frequency flexural motions of the raft structure to be damped out, yet vibration transmission associated with these distributed control forces is largely self-cancelling and minimised as a result.

At the higher frequencies, the raft resonances are allowed to "run-free" and are fully decoupled from the control system. So these resonances cannot (in theory) transmit at all into the receiving structure. In practice, local residual mount transmission paths can cause re-coupling of these modes, but a local high frequency adaptive feedforward algorithm can then be brought into play, which will completely suppress such unwanted transmission.

The associated control strategies are broad-band controllers, so that all these performance improvements are achievable regardless of the precise nature of the excitation signals. So the performance is as equally applicable to random or transient disturbances as for steady, discrete frequency tones.

The overall control strategies which have been proposed define a comprehensive method of isolating large, flexible, lightweight structures over a broad bandwidth which fully encompasses both large-lengthscale, low- frequency structural motions and localised high frequency vibration. The various conflicting design requirements, which the passive designer has hitherto been forced to trade-off against each other can now be separated out and achieved independently. Such requirements can be explicitly defined as appropriate cost functions within a control optimisation strategy, so that the design task then becomes associated with the specification of reconfigurable software, rather than of compromise mechanical installation.

This page is intentionally left blank.

1 Introduction

1.1 General Perspective

An important feature of many marine installations is that several items of machinery are coupled together onto a raft structure, and provided with a common mount isolation system. The advantage of such an approach is that the task of vibration isolation can be greatly simplified, since the alignment between individual machinery items on the raft is more easily maintained, while the need for flexible couplings and pipework is much reduced.

A further advantage of this approach is that the raft structure can generally provide a more rigid, massive assembly than the local mounting points associated with individual machinery casings, thus in principle offering a more dynamically satisfactory isolation system. An example of a large raft structure is sketched in Figure 1.1, showing a diesel electric propulsion system and associated auxiliaries mounted on a common platform (this specific example relates to a fisheries research ship installation).

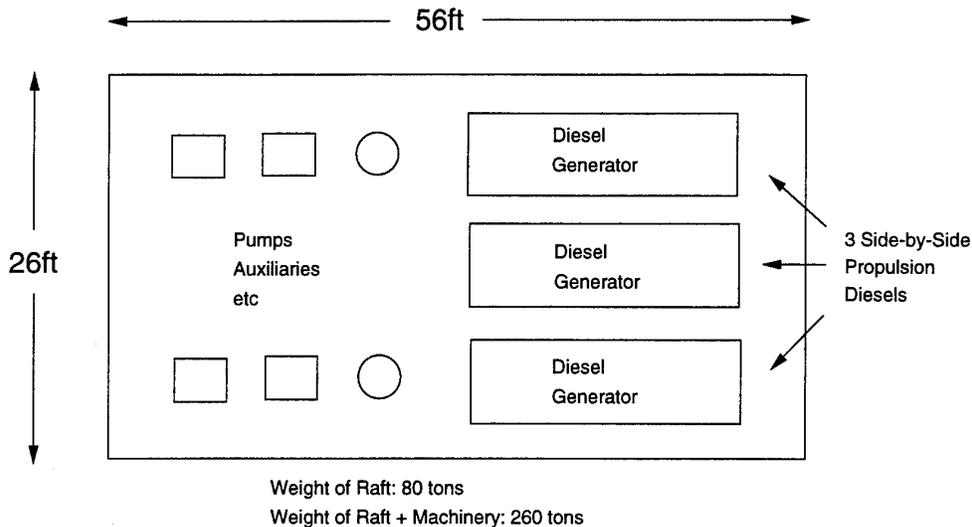


Figure 1.1: Typical Marine Raft Installation

Unfortunately, as the size of such rafted structures increases, the overall isolation performance can become increasingly compromised by the effects of structural flexure; large extended structures are inevitably prone to resonating in their fundamental bending and flexing modes, and such effects can substantially offset the benefits associated with mounting a large number of items on a common, isolated bedplate. In the example shown in Figure 1.1, it was necessary to provide very substantial steel beams to stiffen adequately the basic raft assembly, with the result that the total raft weight represented 1/3 the weight of the entire machinery installation.

Modern development of active control techniques, however, offers the opportunity to overcome this constraint by providing appropriately controlled input signals to damp out and control many of the low frequency flexural motions. Moreover, it then becomes possible to ensure that the residual raft motions more closely approximate to an ideal freely suspended rigid body, thus considerably enhancing the efficiency of the isolation. In this way, the possibility opens up of designing large extended raft structures using lighter structural components than hitherto, which can yield corresponding reductions in the overall weight of the isolated assembly while providing extremely high isolation performance.

The overall purpose of the analyses described in this report is therefore to investigate the dynamics and operation of actively controlled, distributed machinery isolation systems. This work represents the theoretical component of a practical research program leading to the construction of an experimental rig to demonstrate the principles which have been identified. This demonstrator rig consists of a lightweight, resonant raft assembly supported on electromagnetic levitation mounts. The motion of the raft structure will be detected by appropriate gap and acceleration transducers, and the distributed active mount system will be controlled in response to these inputs.

The objective is to demonstrate that such isolation systems, correctly designed, can achieve substantial performance improvements over conventional passive raft installations. Such improvements are sought in two separate respects. The first is entirely local to each active mount, which should be capable of providing a much higher degree of softness and isolation than conventional passive elements, while possessing much higher stiffness at very low frequencies in order to satisfy the constraint of accurate low frequency alignment. The second relates to the co-ordinated operation of distributed mount arrays; the correct design of a distributed active isolation system opens up the possibility of exploiting spatial impedance mismatch to enhance isolation performance. Specifically, this involves transforming low-wavenumber disturbances into higher wavenumber distributions, which couple less efficiently into the receiving structure. Correct use of such techniques can provide the opportunity to deploy large, lightweight flexible raft structures, which are actively stiffened to minimise the effects of large-scale internal structural resonances, while providing a degree of vibration isolation well in excess of conventional, massive, stiff rafts.

The overall specification for the design of a high-performance distributed isolation system involves several requirements to be addressed in parallel. These requirements can be set out as follows:

- (1) Accurate Low Frequency Alignment (<1 Hz): the ability to maintain constant position of the machinery raft relative to the hull under seaway and manoeuvring disturbances.
- (2) Active Stiffening of Internal Raft Resonances with simultaneous Low Frequency Isolation (1Hz-50Hz). The raft must provide a stiff platform for mounted machinery, while at the same time providing a high degree of isolation from the hull. This requires the actively controlled mounts to be controlled in a collective fashion, taking into account the spatial distribution associated with the modal deformations, and the corresponding wavenumber distributions associated with the control force array.
- (3) High Frequency Local Mount Isolation (> 50Hz). At the higher frequencies of operation, mount vibration transmission becomes more localised, and individual mount paths can be regarded as separate, independent sources of excitation. The associated control strategies seek to minimise the vibration transmission through the individual mounts, irrespective of the behaviour of other elements in the active array.

The purpose of this report is to address all three requirements, and to define an overall control strategy which correctly combines these features, with smooth transition from one regime to the next. Thus the overall bandwidth of the control system (DC-500Hz) must correctly encompass the transition from large-scale global dynamics to small-scale local detail. Moreover, the design objective is to achieve these performance requirements for fully broad-band, random inputs, so that the isolation achieved will be largely independent of the precise nature of the disturbance; the raft assembly can then isolate both transient and steady- state disturbances with equal facility.

1.2 Specific Design Objectives

The experimental demonstration rig will consist of two parallel flexible beams, coupled together by eight stiff transverse members. (Figure 1.2) The assembly will be supported by an array of 16 triaxial active electromagnetic mounts at each intersection point, giving control forces on demand in each of the vertical, fore-aft and transverse directions. The objective of the theoretical analyses is to identify and define a software control specification which will enable the isolation of this structure to be demonstrated, so as to meet the following requirements.

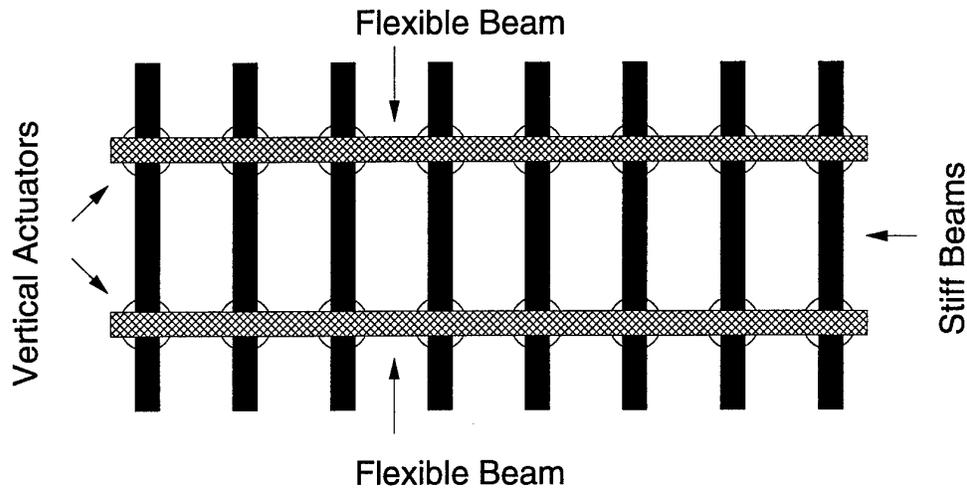


Figure 1.2: Experimental Raft Structure

- (1) To provide substantial low-frequency stiffness of the overall assembly when subjected to simulated seaway motions (< 1Hz).
- (2) To suppress the fundamental mount resonances (1Hz-10Hz).
- (3) To suppress the fundamental raft structural resonances using force distributions which simultaneously minimise low-wavenumber coupling into the hull structure.
- (4) To isolate higher frequency raft structural resonances
- (5) To maximise the isolation of localised high frequency vibrations

1.3 Sensors and Performance Criteria

The raft instrumentation will consist of the following components:

- (1) Gap Transducers (to define low frequency alignment)
- (2) Accelerometers (to define the raft dynamical response)
- (3) Flux transducers (to control and linearize individual magnet operation, while providing direct monitoring of the output transmitted forces).

The objective is to minimise the response of the gap and accelerometer transducers (1) & (2) by application of appropriate control forces (3). At the same time, however, such control forces must be constrained to minimise the coupling and transmission of vibration to the hull structure.

The performance requirement is to achieve a consistent improvement in both precision of alignment and low frequency isolation compared to conventional passive mounting. In this respect, a starting-datum can be identified as a 4Hz-resilient mounted passive system: the corresponding seaway response and force transmissibility characteristics are well-known. The objective is to seek an order-of-magnitude performance improvement in respect of both low frequency alignment and vibration isolation, relative to the 4Hz-passive datum. The precisely achievable improvement is the subject of the theoretical research and demonstrator program; the contractual requirement is for a consistent 10dB improvement in every respect, but the theoretical simulations have indicated that considerably greater improvement should be achievable.

2 General Algorithm Principles

2.1 Physical Principles

Conventional passive isolation of machinery on rafted platforms relies on the existence of an impedance mismatch between the raft and the hull structure. This is achieved by the use of compliant elements, such as rubber mounts, for which the natural impedance reduces as a function of frequency. At very low-frequencies such mounts provide a degree of restraint against seaway motion (high impedance relative to the mounted structure), while at higher frequencies they are very much softer (low impedance) and thus provide isolation. This natural roll-off in impedance means that there must inevitably be a transition regime where the mount and structural impedances are equal; in this regime, mount-structural resonances occur.

The behaviour of idealised passive mass-spring isolators is well-known; in particular, it is useful to note that the static deflection of the mount can be calculated directly from the specified mount resonant frequency regardless of the size of the supported mass, and is given by the expression:

$$d = g_o / (2\pi f) \qquad g_o = 9.81 \text{ m/sec}^2$$

where d is the deflection (mm), and f is the mount resonant frequency. Thus, for any 4Hz mount (typical of modern marine installations), the vertical deflection under load = 16mm. For a vessel undergoing a slow 45 degree roll, one can immediately expect athwartship displacements of +/- 10mm simply as a result of the change in direction of the static load vector.

In contrast, actively controlled magnetic mounts provide the opportunity for using integrating feedback to reduce the deflections under seaway motion by a full order of magnitude, while simultaneously achieving a much greater high frequency impedance mismatch. This can be achieved by introducing the characteristics of a "frequency dependent spring" ie by ensuring that the force/deflection characteristic of the magnetic mount rolls-off as a function of frequency much faster than would be achievable by purely passive means.

It must be noted that a control strategy which generates its output force in response to measured mount deflection will respond with the same dynamics whether those deflections are induced by seaway disturbances or vibration. Moreover one will still encounter a transition regime where the impedance of the isolating element matches the impedance of the adjacent surrounding structure, giving rise to mount/structural resonances. Care must then be taken to ensure that the resultant interactions are of a stable nature; the process of "rolling off" the active spring stiffness introduces phase lags which tend to destabilize these interactions. However, more sophisticated control strategies can provide improved response, by using dual detection systems (gap, accelerometer) to discriminate between external seaway induced motions and internal vibration induced motions, thus allowing separate isolation characteristics to be obtained for each separate mode of transmission (ie non-reciprocal behaviour). In this way, the installation can be made much softer in response to vibration inputs, while still accurately tracking seaway motions. In addition, the unwanted mount resonant effects can be suppressed completely, in a stable manner, without compromising either seaway restraint or isolation performance.

The above comments apply to idealised installations, where the mounted assembly is assumed to represent a compact, localised mass. However, under circumstances where the installation is an extended, distributed raft (eg rafted installations in surface vessels up to 20 metres by 10 metres have been constructed), there is a very important additional effect to take into account. The natural flexural response of the raft and the hull/water structure defines an extended interface where wavenumber

matching plays an important role in the vibration transmission process. If the natural flexural wavenumbers of the raft and hull/water interface are closely related, then vibration coupling can be significant, whereas if the flexural wavenumbers are substantially mismatched, this coupling is correspondingly reduced.

This particular effect (ie wavenumber mismatch) is well-established as a mechanism which minimises response at extended interfaces, and is exploited (for example) in the design of distributed conformal hydrophone arrays, to decouple the low wavenumber acoustic reception from the high wavenumber disturbances associated with flow noise. Similar principles are also used in the design of the supporting structures for such arrays, to minimise the interference from high-wavenumber hull vibration (self-noise).

The use of an actively controlled array to support a large distributed raft assembly provides the opportunity to implement control force arrays which exploit this same feature, leading to improved isolation while at the same time permitting high-wavenumber control force distributions to suppress the internal resonances of the structure.

Thus the design of an actively controlled, extended distributed raft can permit three separate effects to be taken into account:

- (1) The introduction of controlled impedance mismatch between the raft and hull structure
- (2) The use of dual detection systems (gap/accelerometers) to discriminate between externally induced and internally induced motions, and to suppress fundamental mount/structural resonant interactions.
- (3) The use of controlled spatially distributed force arrays to stiffen the internal resonant modes of the raft, while maximising the wavenumber mismatch between this control force distribution and the characteristic hull/water wavelengths, thus still further reducing the vibration transmission.

The comprehensive design of an actively controlled distributed machinery raft requires the above aspects to be addressed in parallel. i.e. one seeks to control simultaneously the response according to frequency, according to direction of transmission, and according to spatial distribution. This overall combination of techniques can provide a very substantial enhancement in performance compared to passive methods of isolation.

However, a first requirement is to provide a basic method of stable support for the electromagnetically suspended raft; this will enable test signals to be injected into the structure, and thus detailed assessment of its dynamics to be carried out. Based on such testing, the necessary control algorithms can then be introduced to exploit the more sophisticated isolation features referred to above.

This particular requirement can be met in a straightforward and reliable fashion if each magnetic mount is configured to simulate a conventional highly damped spring. (Figure 2.1) It is well-known that passive structures supported on damped spring suspensions are extremely stable, so one can guarantee from the outset that this method of magnetic suspension will also be very stable, regardless of the precise dynamics of the supported structure.

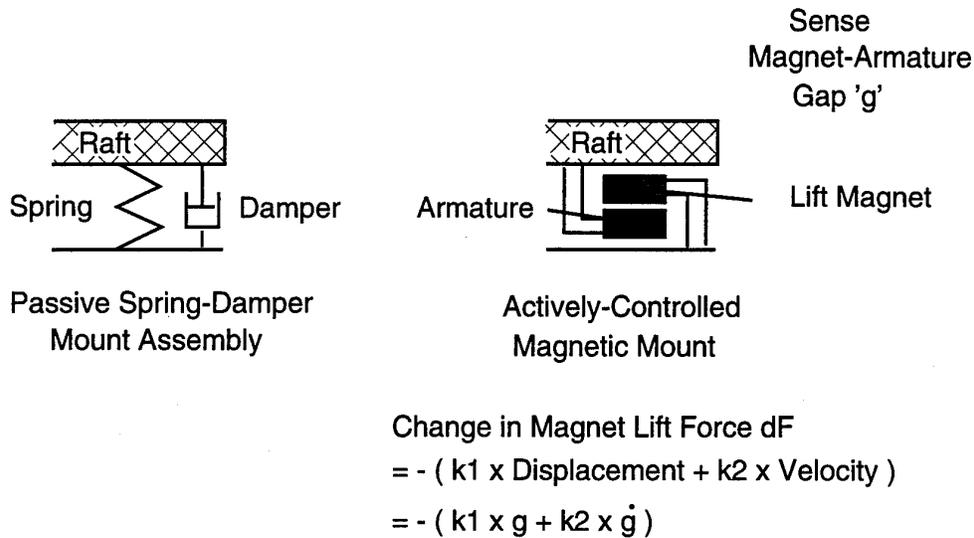


Figure 2.1: ZBD Control

This particular method of control is referred to as the "Zebedee (ZBD) Algorithm"; it has been found in practice to be an extremely useful and reliable approach.

Having thus achieved a basic, stable electromagnetic support, it is then appropriate to configure more sophisticated algorithms to seek out and exploit the features referred to in (1)-(3) above. The general features of such algorithms will now be described in the following section.

2.2 General Algorithm Characteristics

The physical features which have been described in 2.1 require three separate algorithm structures to be considered. The first of these has already been outlined, namely the use of a basic support algorithm which simulates a damped, passive mount (the ZBD algorithm) to enable the raft to be stably supported and located while more detailed testing and system identification is performed.

The control algorithms which are subsequently implemented take two separate forms, namely a low-frequency feedback algorithm, and a high frequency feed-forward algorithm. These are illustrated in Figures 2.2 and 2.3 respectively. Essentially, at low frequencies, the input to the control system is provided by gap and accelerometer signals, and the control system acts directly to suppress fluctuations in these components (a closed-loop feedback system). On the other hand, at higher frequencies the phase lags involved in feeding-back signals become too large for stable operation and it becomes necessary to implement a feed-forward open-loop control system.

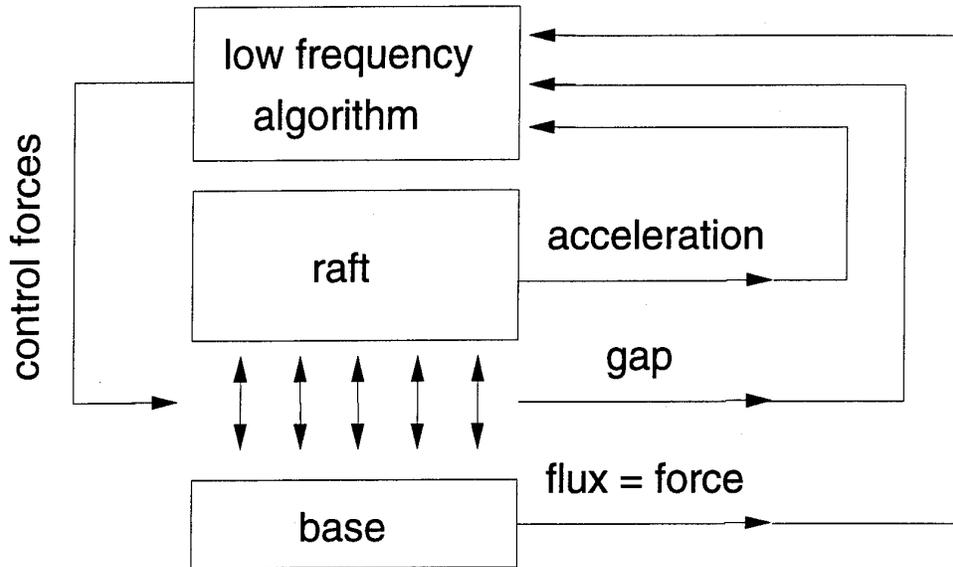


Figure 2.2: Low Frequency Algorithm

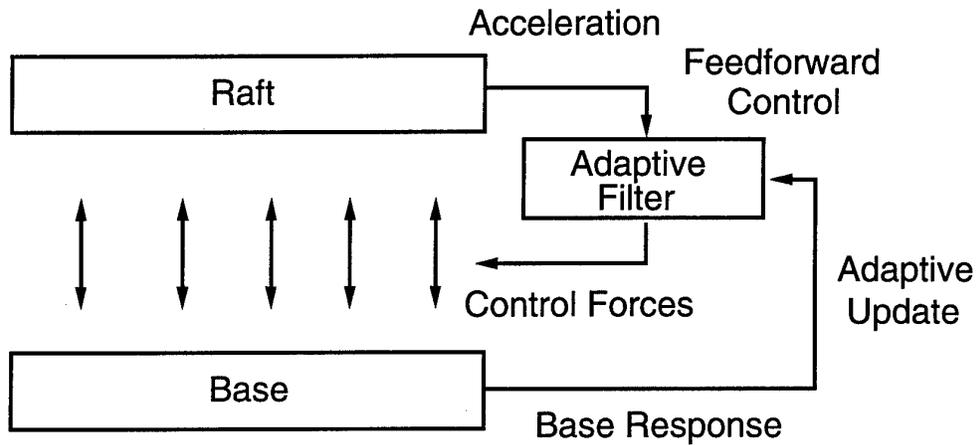


Figure 2.3: High Frequency Algorithm

In the latter case, the input to the control system is still provided by the accelerometers and the output is still represented by the control forces, but now these control actions take place at a very low gain so that the control forces have negligible feedback influence on the raft acceleration levels. Rather, the objective becomes to define the control system transfer functions so as to minimise the net response of the receiving hull-structure. This is achieved by performing an adaptive update of the transfer functions as a result of observation of the receiving structure response.

The overall control system is then constructed by implementing these two control strategies in parallel, to give the final structure sketched in Figure 2.4. This essentially consists of an inner high-gain loop representing the low-frequency feedback algorithm, and an outer low-gain path which represents the high frequency feedforward process. An important requirement will then be to ensure that these two processes can successfully coexist, without the action of one compromising the action of the other.

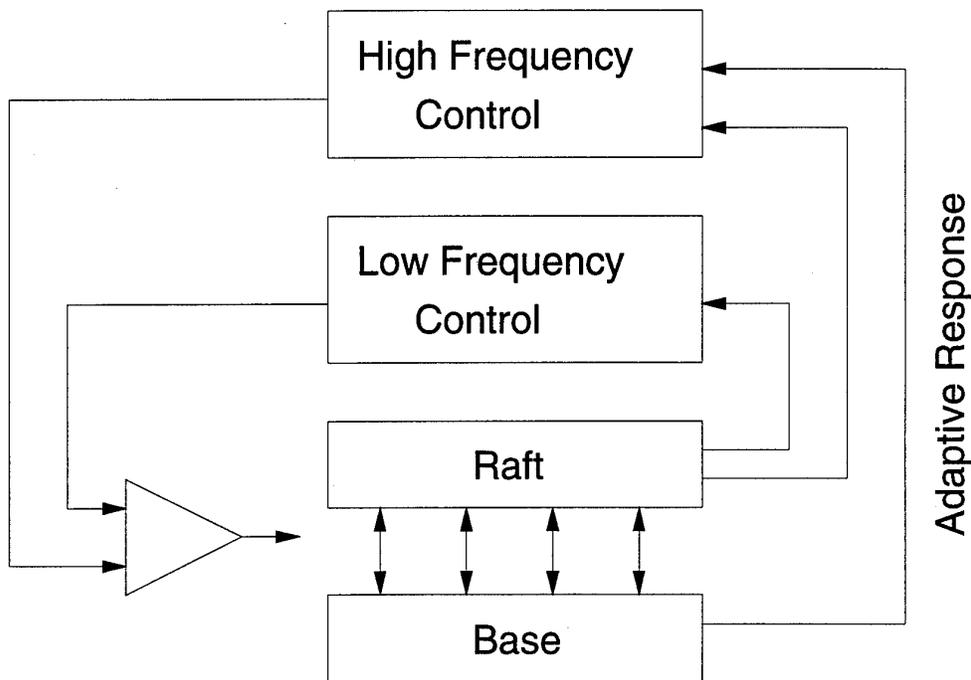


Figure 2.4: Overall Algorithm

2.3 Low Frequency Control Algorithm

As indicated above, the low frequency control algorithm is a feedback controller, which detects the raft gap and acceleration signals and outputs the necessary control forces to suppress unwanted low frequency disturbances in these signals. Since the objective is also to isolate the raft structure from the hull, these control forces must reduce rapidly in amplitude with increasing frequency, so that the overall controller exhibits a markedly low-pass characteristic.

An important requirement is to consider also the spatial distribution associated with these control forces, in addition to their frequency dependence, so that in the region where they exert significant restraint on the raft motions, they are arranged to provide the least possible coupling interaction with the hull/water radiation characteristics.

This algorithm is specified using optimal state-space feedback control techniques. A feature of this approach is that provides a "model reference" controller (Figure 2.5), whereby the operation of the physical system is simulated in parallel by a real-time state-space model. At any instant, this model possesses a comprehensive set of estimates of the instantaneous state-variables of the overall system. The control output is derived directly from these estimated states, while the tracking of the model is maintained by a Kalman Filter update to the state-equations, based on feeding-back the difference between the predicted outputs of the model and the actual observed system outputs.

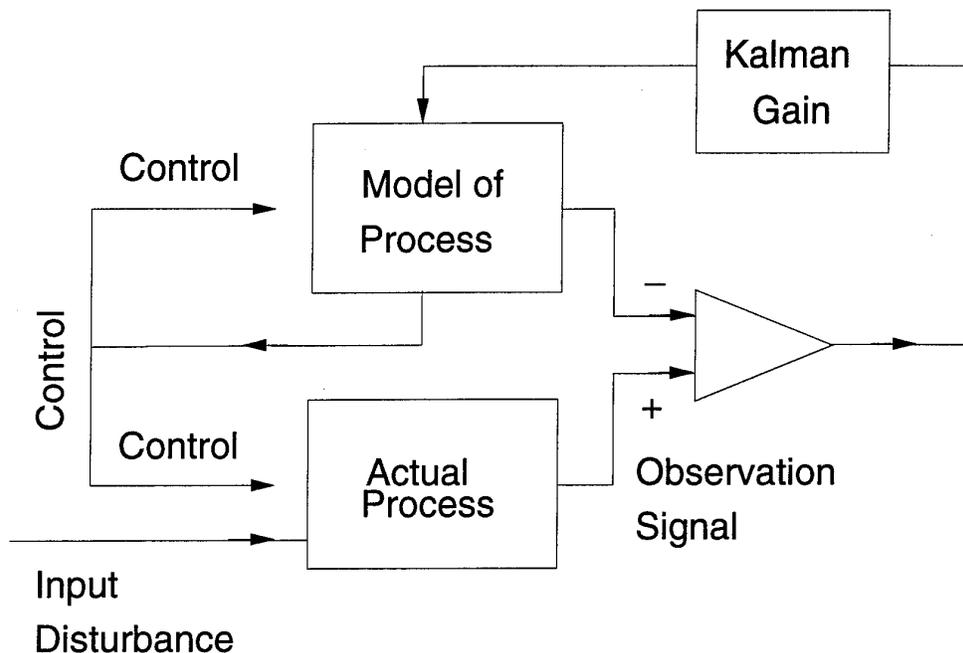


Figure 2.5: Model Reference State-Space

There are several advantages associated with this approach, as follows. First, once the system model has been accurately established, the technique permits rigorous, stable optimal feedback control to be applied, since the control can be derived directly from each and every (estimated) state of the system. Secondly, the associated feedback of the Kalman filtered observations can be specified according to optimal criteria for the rejection of unwanted noise in the observations. Finally, the existence of a parallel system model enables a continuous comparison to be made between the behaviour of the model and the behaviour of the actual system, thus providing a basis for future adaptive modification of the model, together with capability for rapid detection of transducer faults in the real system.

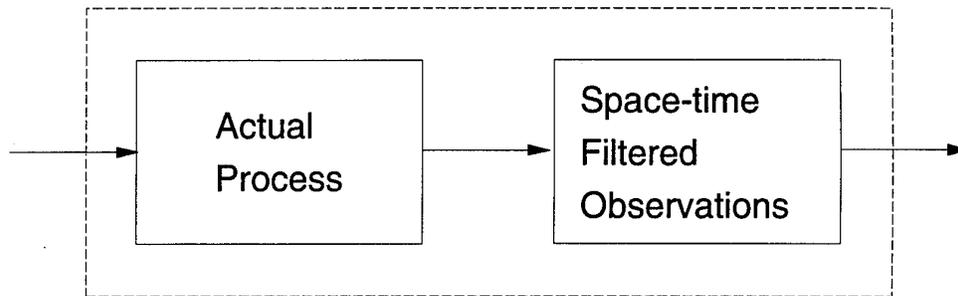
In order to pursue such an approach, several requirements must be satisfied. First, it is important to have an accurate model of the real system, expressed as a parallel set of differential state equations. In practice this is most easily achieved by decomposing the motions of the structure into a finite set of modes, each described by a conventional second-order equation of motion. The corresponding states

can then be defined as the amplitudes and velocities of each mode respectively. Since the system in this case is a finite structure possessing conservative boundary conditions (it is assumed that the only external connection and support is provided through the active mounts), such a unique modal decomposition undoubtedly exists, and can be established by a variety of alternative testing procedures.

Having defined a suitable model for the raft platform itself, the subsequent control task would be greatly simplified if state-space control could be applied as a set of fully independent modal control actions. Unfortunately, the rigorous solution requires additional optimality criteria to be satisfied both with respect to the control and observation dynamics. The specific constraints imposed by these optimality conditions lead inevitably to a degree of cross-coupling between the modes. Fortunately, however, it is found that the rigid-body motions of the assembly can be decoupled from the flexural motions, so that the problem can proceed with separate optimality conditions for these separate components of motion.

Turning now to the optimality conditions, it is necessary to summarise the designer's requirements in an appropriate set of quadratic cost functions. In this respect, it is convenient to impose an additional space/time filtering action on the actual system observation and control variables, so that the output of this filtering action completely summarises the various space/time weighted performance criteria which the designer wishes to optimise. Then, by specifically incorporating this filtering action into the overall definition of the system model, the quadratic cost functions for the modified output signals take a very much simpler form, and the subsequent solution of the problem becomes relatively straightforward.

This approach is illustrated in Figure 2.6; the control and observation parameters are passed into a filtering process which generates a new set of control and observation components. This transformed set of components are then incorporated into the overall system model, thus defining a modified model which includes these filtering processes as an integral part of its system state equations.



**Translates Actual Observations Into
Noise Control Requirements**

Figure 2.6: Modified Observation Process

The advantage of this approach is that the optimality condition for the modified parameters now takes a particularly simple form, namely to minimise the mean square sum of the filtered observation and control signals. The detailed trade-offs associated with the designer's various performance requirements are now contained within the central definition of the space/time filtering process, which in turn becomes an integral component of the system model. The designer can therefore concentrate

his effort on specifying this space/time filtering process to satisfy his particular objectives, without undue concern for the subsequent solution procedure. Once the filtering processes have been accurately defined, the derivation of stable, optimal feedback gains becomes a clear-cut calculation, using well-established techniques.

A brief summary will now be given of the typical space/time filtering processes which are required, with reference to Figure 2.7. The objective is that the outputs of the filtering process should represent the significant contributions to the cost function.

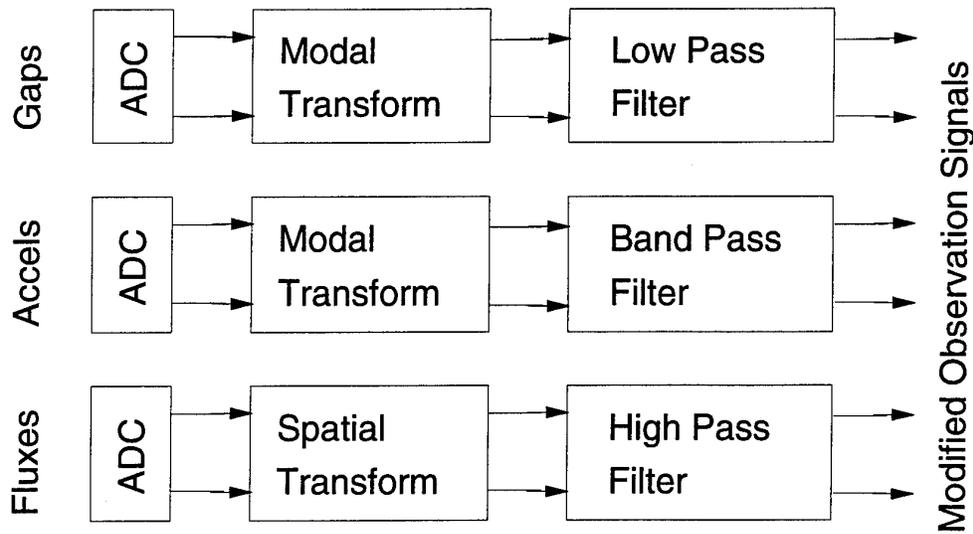


Figure 2.7: Detail of Observation Process

2.3.1 Modal Gaps

The objective is to ensure that at very low frequencies the actual gaps are accurately aligned, but this requirement is relaxed as the frequency increases and the raft is then allowed to vibrate freely. Thus the corresponding modal gap outputs are low-pass filtered with a cut-off around 1Hz, so that these output components assume significant values only at very low frequencies.

2.3.2 Non-Reciprocal Criterion

The requirement is to ensure that those rigid-body accelerations which are necessary to follow hull seaway motions are accurately driven by the control forces, while rigid-body accelerations which are induced by vibration excitation are ignored, and excite no corresponding force action. This requirement can be met by imposing a combined constraint on the rigid-body modal gap and acceleration components, namely

$$\text{Minimise } \{ F + M (\ddot{g} - a) \}^2$$

where the F is the control force, g is the rigid-body modal measurement gap, and a is the rigid-body raft modal acceleration component.

Under circumstances where the motions are entirely due to vibration, the gap derivatives g and raft acceleration components a are equal; the cost function then seeks to drive F to zero. However, when the motions are due to seaway demands, the function $(\ddot{g} - a)$ effectively measures the hull acceleration. The overall cost function is then at minimum when the control force is of precisely the size necessary to drive the raft with an acceleration matching that of the hull.

2.3.3 Flexural Mode Accelerations

The requirement is to ensure that low frequency internal resonances of the raft structure are suppressed, so the corresponding acceleration components are band-pass filtered. Each filter is centred on the resonant frequency of its respective mode, so that its contribution to the cost function is greatest at its resonant frequency. For higher order modes, modal accelerations are allowed to run free, and the corresponding band-pass filter gains are set to zero. So the components of acceleration associated with the higher order modes play no part in the control response.

2.3.4 Spatial Control Force Distribution

An important requirement is that the control actions which are brought into play to suppress the low frequency resonant modes should possess spatial distributions which minimise the efficiency of transmission into the supporting hull structure. A regularly distributed control force array can be decomposed into spatial spectral components, simply by imposing a discrete spatial Fourier Transform operation. The corresponding transform components are referred to as spatial wavenumber components, and it is well-known that it is the long wavelength, low-wavenumber components which transmit most efficiently into the hull receiving structure and the sea. In contrast, the troughs and crests of the shorter wavelength high wavenumber components tend to mutually interfere, giving rise to a spatial cancellation effect. As a result, they transmit much less efficiently into the receiving structure.

Thus when seeking to control the internal resonances of the raft structure, it is preferable to do so wherever possible with short wavelength, rather than long wavelength control force distributions. This in turn can be quantified by imposing suitable constraints on the choice of wavenumber component; the low wavenumber force components carry a high penalty, while the high wavenumber force components carry a much lower penalty.

The control force array can therefore be transformed using an orthogonal spatial transform, to define an equivalent set of controls in the spatial transform domain. The cost penalty is then applied most severely to the low wavenumber components. This constraint takes the form of a set of conceptual high-pass filters cutting on at progressively higher frequencies. Thus, for the lowest wavenumber component, the high pass filter comes into play at a very low frequency, whereas for the higher wavenumber components the cut-on becomes effective at correspondingly higher frequencies. In general, these cut-on parameters should be chosen to mirror directly the transmission and radiation cut-on characteristics of the corresponding wavenumber patterns.

The above discussions relate to the cost penalty which should be applied to the observation of the control action, namely the control fluxes. However, in practice it is equally possible to impose the equivalent cost constraint directly on the demand for control action.

But in the latter context, it becomes easier to choose the central mathematical set of control parameters as those variables which would correspond to the outputs of this high-pass filtering process; the real control signals can then be derived by low-pass filtering and inverse spatial transforming these central control parameters. The necessary low-pass filters become naturally

defined as the complementary inverse of the corresponding high-pass filters. This feature is illustrated by transforming Figure 2.7 to the equivalent representation of Figure 2.8. It can be seen that the model of the complete system now incorporates both front-end and back-end filtering and transformation processes.

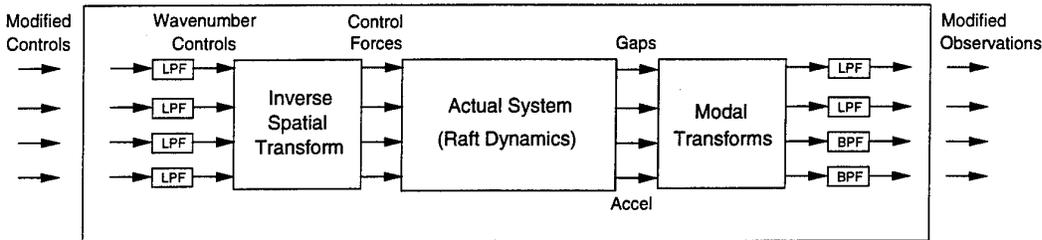


Figure 2.8: Front and Rear End Processes

The advantage of this latter approach is that all the filters which are actually implemented become low-pass components, and so are much less sensitive to high-frequency noise. The net cost-function can then be applied as an equal weighting against the central control components; the subsequent low-pass filtering to derive the final control outputs will automatically suppress the low-wavenumbers and bias the outputs in favour of the high-wavenumber components.

2.3.5 Decoupling of Rigid Body and Flexural Modes of Motion

It has previously been remarked that for simplicity one would wish to control each mode of motion of the structure independently, but the imposition of the cost constraints leads to cross-coupling between the different modes of motion. This cross-coupling arises as a consequence of the dual orthogonal structure of the control/optimisation problem. The modal decomposition of the structure defines one complete set of orthogonal eigenvectors which can describe fully any force distribution. But the spatial decomposition of the control force array provides a second complete set of orthogonal eigenvectors which can also fully describe any force distribution. These two orthogonal sets are not the same; indeed it is the very fact that they represent two different orthogonal sets spanning the same vector space of force distributions that provides the opportunity to choose higher wavenumber distributions which can still exert leverage over low-order modal deformations.

This is most easily expressed mathematically as follows.

There exist two complete sets of orthogonal eigenvectors, ϕ_n and ψ_n , such that any force vector can be expressed as a sum of either set:

$$\text{Structural Modal Decomposition: } f = \sum_{n=1}^N a_n \phi_n$$

$$\text{Spatial Wavenumber Decomposition: } f = \sum_{n=1}^N b_n \psi_n$$

Each individual set of eigenvectors is mutually orthogonal, so that

$$\int_{\text{Array}} \phi_n \phi_m = 0, \quad n \neq m \quad \text{and} \quad \int_{\text{Array}} \psi_n \psi_m = 0, \quad n \neq m$$

But the two separate sets are not the same, with the result that cross-products between the two sets do not vanish

i.e. In general, $\int_{\text{Array}} \phi_n \psi_m \neq 0$ for arbitrary n, m

It is this very feature which enables alternative choices of control action to be made. For example if one wishes to control a specific mode ϕ_s , the fact that

$$\int_{\text{Array}} \phi_s \psi_L \neq 0 \quad \text{and} \quad \int_{\text{Array}} \phi_s \psi_H \neq 0$$

means that one can influence the mode either with a low-wavenumber component ψ_L or a higher wavenumber component ψ_H .

The purpose of the optimisation is to determine the specific choice of control action which can adequately leverage a given mode (belonging to the first set of vectors) while minimising the spatial cost function; the latter function is defined on the second set of vectors.

Although these two orthogonal sets of vectors are not the same, they can be arranged to share one set of eigenvectors in common, namely the rigid body modes. These components of motion are always orthogonal to the higher bending modes of the structure, while momentum considerations ensure that the rigid body modes will not be excited by any force distribution which has zero net linear or angular momentum. So neither a higher flexural mode of the structure, nor a zero-momentum wavenumber mode of the spatial force distribution can influence the rigid body modes.

As a result, the optimality conditions for the rigid body motions of the uniform raft can be fully decoupled from both the flexural structural motions and the zero-momentum wavenumber components. Thus, the control problem can be separated-out into two distinct problems, namely the optimal control of the rigid body modes, and the optimal control of the flexing modes.

While this separation of the rigid body modes from the flexural modes does not in any way modify the general philosophy of the control problem as previously set out and described, it does make for a useful simplification in the present numerical calculations. Instead of handling a single large matrix optimisation, the practical optimisation can be performed as two separate, non-interacting smaller calculations. This has been found to substantially reduce the sheer size of the numerical arrays which are required for each specific calculation, and also to improve the speed of convergence of the associated numerical solutions.

Having described the general features and philosophy of the low-frequency control algorithm, the precise mathematical calculations are set out in Section 3. Meanwhile, a corresponding overview of the high frequency algorithm will now be given.

2.4 High Frequency Control Algorithm

The high frequency control algorithm operates according to an entirely different philosophy from the low frequency algorithm. Whereas the low frequency algorithm makes use of substantial feedback gain to suppress unwanted disturbances, the high frequency algorithm operates in a feedforward mode with gains which are generally of order unity or less. Moreover, whereas the low frequency algorithm

operates in a global sense, ie forming each control output from a linear combination of all the sensor inputs, the high frequency algorithm operates entirely local to each individual mount, sensing the immediate input acceleration to the mount, and driving the control electromagnets to minimise transmission into the base structure.

Since this control process operates at low gain in an essentially open loop fashion, it is necessary to use a secondary observation process to monitor the overall performance and adjust the control parameters to achieve the desired response.

At the same time, it is necessary to ensure that operation of the high frequency algorithm does not interfere with and compromise the low frequency algorithm, which it is assumed has already been optimized over its working frequency range.

A feature of such adaptive feedforward processes is that the processing workload is generally very much more intensive than the corresponding demand for a feedback controller. This is because all the details of the transmission process must be reproduced accurately within the feedforward path; the requirement for precision in this respect is much tighter than for a conventional feedback process. A typical requirement is to reproduce features of detail with 2Hz resolution using a 2kHz sampling rate (500Hz bandwidth). This necessitates a 1000-point adaptive FIR filter, with a corresponding workload of 4MFlops for the baseline filtering requirement, and 8Mflops if the adaptive update is to be maintained every sample.

One can summarise the following requirements which the algorithm must satisfy.

- (1) It must seek to reproduce unwanted transmission into the base structure (eg arising from eddy current and local air compression effects) and regenerate this in antiphase, thus cancelling these components of transmission.
- (2) It must possess an appropriate method of performance monitoring and optimisation, to maintain the net transmission process at the least possible level.
- (3) Although individual mounts are controlled on a local basis, the subsequent base structural response is due to the combined effects of several transmission paths. An appropriate deconvolution process must be implemented to ensure that the minimisation of each transmission path takes place without cross-interaction.
- (4) The optimisation process must take place over the higher frequency regime, without compromising the existing optimized performance associated with the low frequency algorithm.

The proposed algorithm seeks to address all of the issues (1)-(4) above. It must be noted, however, that in certain respects the interaction with the low frequency algorithm is already minimal, since the latter is a global algorithm while the high frequency algorithm is entirely local in its operation. The net coherence between any control signal which has been derived purely locally, and the low frequency control signal which is derived from a distributed global set of inputs will generally tend to be very poor if there are multiple uncorrelated sources present on the raft. Thus there will be little opportunity for the local control signal to influence and modify the response induced by the global controller.

The advantage of this feature is that the interaction between the global low frequency and local high frequency algorithms will be correspondingly reduced, but the corresponding disadvantage is that the high frequency algorithm will be unable to "make up" for any poor design in the roll-off associated with the low frequency algorithm. Thus the rate of roll-off associated with the global algorithm defines a lower bound on the overall performance, and the purpose of the high frequency algorithm becomes to ensure that this lower bound is indeed achieved, and is not compromised by local leakage paths associated with the mount transmission.

The high frequency algorithm is implemented, in the first instance, by using an appropriate modification of the standard "filtered-x" algorithm. The central feature of this algorithm is that the transmission process from the controller output to the observation points is simulated in parallel (the filtered-x component). Any residual error at the downstream observation locations can then be correlated with the response which would have been expected at these locations, given the known signals which have been applied to the actuators.

Any residual correlation is then used as a correction to the actuator drive filters, in order to drive the correlation to a minimum. An important modification to this algorithm in the present context is that it includes the cross-correlation of the responses at all the downstream observation points, thus enabling the cross-interactions between the different actuators to be taken into account.

As already remarked, it is essential that the algorithm should only modify the high frequency components of the actuator drive signals, and should not seek to modify the existing optimized low-frequency behaviour. In this respect, it is not sufficient simply to high-pass filter the information supplied to the algorithm, since this leaves the algorithm free to assume arbitrary values at low frequency. A positive constraint must be applied to the low-frequency behaviour.

This constraint is achieved by simulating the known effects of the low frequency algorithm, and subtracting these off from the error correlation process. This ensures that the low-frequency operation is left intact, and that non-zero error correlations relate only to the high frequency behaviour. More importantly, since the information supplied at the input to the algorithm is now all-pass, any attempt to modify the low-frequency behaviour will definitely register as an "error", so the constraint is imposed in a positive fashion, rather than simply by default.

The precise operation of this algorithm is set out in detail in section 4.

3 Low Frequency Controller

3.1 System Model

At start-up, the raft is levitated using a robust control algorithm designed to mimic as closely as possible the properties of a passive spring-damper device. The design and performance characteristics of this local point controller have been documented elsewhere [1]. Once a stable lift has been achieved, a disturbance can be added to each actuator drive in turn and the acceleration and/or gap signals from all locations captured. Standard techniques ([2], for example) can then be used to process this data in order to provide a truncated modal model of the form

$$\ddot{\underline{\phi}} + \Lambda \dot{\underline{\phi}} + \Omega \underline{\phi} = V^T \underline{u} + V^T \underline{\Gamma}$$

$$\underline{y} = V \underline{\phi} + \underline{\tilde{y}} \quad \dots(1)$$

where $\underline{\phi}$ is the q dimensional vector of modal displacements, \underline{u} is the r dimensional vector of actuator forces, $\underline{\Gamma}$ is the r dimensional vector of vibration forces, \underline{y} is the r dimensional vector of measured gaps and $\underline{\tilde{y}}$ is the r dimensional vector of seaway induced gap changes. The matrices Λ and Ω are diagonal with i th elements $2\zeta_i \omega_i$ and ω_i^2 respectively. ω_i and ζ_i are the natural frequency and damping of the i th mode. The i th column of the matrix V is the mode shape vector corresponding to the i th mode. If it is assumed that the hull is much stiffer than the raft structure then the seaway motion will only couple in to the rigid body modes of the raft. As a result the seaway input can be defined in terms of a vector of rigid body modal displacements (heave, pitch, roll etc.), \underline{y}_r , and $\underline{\tilde{y}}_r$ in equation 1 can be expressed as

$$\underline{\tilde{y}}_r = V_r \underline{y}_r \quad \dots(2)$$

V_r results from the partitioning of V into its rigid body (denoted by the subscript r) and flexing (denoted by the subscript f) components,

$$V = [V_r \quad V_f] \quad \dots(3)$$

On the assumption that the columns of V form an orthonormal set ($V^T V = I$), then

$$\underline{y}_m = \begin{bmatrix} \underline{y}_r \\ \underline{y}_f \end{bmatrix} \stackrel{\Delta}{=} V^T \underline{y} = \underline{\phi} + \begin{bmatrix} I_6 \\ 0 \end{bmatrix} \underline{y}_r \quad \dots(4)$$

or

$$\underline{y}_r = \underline{\phi}_r + \underline{y}_r \quad \dots(5)$$

and

$$\underline{y}_f = \underline{\phi}_f \quad \dots(6)$$

where

$$\underline{\phi} = \begin{bmatrix} \underline{\phi}_r \\ \underline{\phi}_f \end{bmatrix} \quad \dots(7)$$

The assumption of orthonormal V is made for convenience and any transform that separates the rigid body and flexing modes can be used.

If

$$\underline{u}_m = \begin{bmatrix} \underline{u}_r \\ \underline{u}_f \end{bmatrix} \stackrel{\Delta}{=} V^T \underline{u} \quad \dots(8)$$

then the differential equation in (1) can be rewritten as

$$\ddot{\underline{\phi}}_r = \underline{u}_r + V_r^T \underline{\Gamma} \quad \dots(9)$$

and

$$\ddot{\underline{\phi}}_f + \Lambda_f \dot{\underline{\phi}}_f + \Omega_f \underline{\phi}_f = \underline{u}_f + V_f^T \underline{\Gamma} \quad \dots(10)$$

Since the flexing mode dynamics (equation 10) are not directly influenced by the seaway input, \underline{y}_s , then the controller design is most efficiently dealt with by constructing independent controllers for the rigid body and flexing modes.

3.2 Rigid Body Controller

As discussed above, the controller is to be designed to minimise a set of filtered observations. For the rigid body controller, each mode is considered separately using the model

$$\begin{aligned} \ddot{\phi}_{r,i} &= u_{r,i} + \eta_{r,i} \\ y_{r,i} &= \phi_{r,i} + y_{s,i} \end{aligned} \quad \dots(11)$$

where $\eta_{r,i}$ represents the component of the vibration force in the i th mode, $i = 1 \rightarrow 6$.

For the rigid body motions the improvement over a typical passive performance is obtained by using the dual sensor capability to provide non-reciprocal behaviour. This can be achieved by seeking to minimize

$$\ddot{g}_s \stackrel{\Delta}{=} \ddot{y}_{r,i} - \ddot{\phi}_{r,i} + u_{r,i} \quad \dots(12)$$

which ensures that the controller responds to hull motions induced by seaway and manoeuvring but rejects raft motions induced by vibration. This feature can be demonstrated by considering the limiting case in which $\ddot{g}_{s,i} = 0$. The control signal is then given by

$$u_{r,i} = \ddot{\phi}_{r,i} - \ddot{y}_{r,i} \quad \dots(13)$$

For the vibration free case, combination of equations 11 and 13 yields

$$\ddot{\phi}_{r,i} = -\ddot{y}_{s,i} \quad \dots(14)$$

It can be seen from this expression that the raft acceleration is identical to the hull and the gap is therefore maintained. In the case of a stationary hull

$$\ddot{y}_{r,i} = \ddot{\phi}_{r,i} \quad \dots(15)$$

therefore from equation (13)

$$u_{r,i} = 0 \quad \dots(16)$$

and so the force applied to the hull is totally decoupled from the vibration force.

In practice, this feature is only sought in a particular frequency band (centred on 4Hz) and therefore the controller is aimed at minimising a signal, g_n , which is the result of band-pass filtering \ddot{g}_s . This operation can be represented in state space form as

$$\begin{aligned}\dot{\underline{x}}_n &= A_n \underline{x}_n + \underline{b}_n \ddot{g}_s \\ g_n &= \underline{c}_n^T \underline{x}_n + d_n \ddot{g}_s\end{aligned}\quad \dots(17)$$

Other signals which should be minimised are the low pass filtered modal gap, g_i , to ensure that the mount is stiff at very low frequencies regardless of the excitation source, and the band pass filtered modal acceleration, a_i , to provide a means for suppressing raft acceleration in specific frequency bands. These signals can be represented as the outputs of the following state space filters

$$\begin{aligned}\dot{\underline{x}}_g &= A_g \underline{x}_g + \underline{b}_g y_{r,i} \\ g_i &= \underline{c}_g^T \underline{x}_g + d_g y_{r,i}\end{aligned}\quad \dots(18)$$

and

$$\begin{aligned}\dot{\underline{x}}_l &= A_l \underline{x}_l + \underline{b}_l \ddot{\phi}_{r,i} \\ a_i &= \underline{c}_l^T \underline{x}_l + d_l \ddot{\phi}_{r,i}\end{aligned}\quad \dots(19)$$

In addition to the above requirements, it is important that the control force is *rolled out* beyond the seaway frequencies in order to minimise its action at higher frequencies. This can be achieved by defining a notional control signal, u_i , which is included in the overall cost function. The actual control signal, $u_{r,i}$, is calculated as the output of the following low pass filtering operation

$$\begin{aligned}\dot{\underline{x}}_u &= A_u \underline{x}_u + \underline{b}_u u_i \\ u_{r,i} &= \underline{c}_u^T \underline{x}_u + d_u u_i\end{aligned}\quad \dots(20)$$

To ensure that control forces are only applied where required, it is necessary to add shaping filters for the vibration forces and seaway disturbance that approximate their spectral characteristics. To this end, the seaway acceleration, $\ddot{y}_{s,i}$ is assumed to be formed from low pass filtering a white noise source,

$$\begin{aligned}\dot{\underline{x}}_s &= A_s \underline{x}_s + \underline{b}_s \varepsilon_s \\ \ddot{y}_{s,i} &= \underline{c}_s^T \underline{x}_s + d_s \varepsilon_s\end{aligned}\quad \dots(21)$$

and the vibration signal $\eta_{r,i}$ is formed from high pass filtering a white noise source

$$\begin{aligned}\dot{\underline{x}}_v &= A_v \underline{x}_v + \underline{b}_v \varepsilon_v \\ \eta_{r,i} &= \underline{c}_v^T \underline{x}_v + d_v \varepsilon_v\end{aligned}\quad \dots(22)$$

In order to solve the optimisation problem, an overall state space description can be obtained by combining equations 11-22 and a standard LQG design obtained. A suitable overall description is

$$\begin{aligned}\dot{\underline{x}}_o &= A_o \underline{x}_o + \underline{b}_o u_i + G_o \underline{E} \\ \underline{y}_o &= C_o^T \underline{x}_o + \underline{d}_o u_i + J_o \underline{E}\end{aligned}\quad \dots(23)$$

where

$$\underline{x}_o^T = [y_{r,i} \quad \dot{y}_{r,i} \quad \underline{x}_g^T \quad \underline{x}_l^T \quad \underline{x}_n^T \quad \underline{x}_u^T \quad \underline{x}_f^T \quad \underline{x}_v^T]$$

$$\underline{E}^T = [\underline{\varepsilon}_v \quad \underline{\varepsilon}_s] \quad , \quad \underline{y}_o^T = [g_l \quad a_l \quad g_n]$$

With this description, the LQG controller

$$u_l = -\underline{k}_{x_o}^T \underline{\hat{x}}_o$$

$$\dot{\underline{\hat{x}}}_o = A_o \underline{\hat{x}}_o + \underline{b}_o u_f + K_{f_o} (\underline{y}_o - C_o \underline{\hat{x}}_o - \underline{d}_o u_f) \quad \dots(24)$$

that minimizes the cost function

$$J = E \left\{ \frac{1}{2\pi} \int_{-\infty}^{\infty} \underline{y}_o^*(j\omega) Q_c \underline{y}_o(j\omega) + u_l(j\omega)^* r_c u_l(j\omega) d\omega \right\} \quad \dots(25)$$

where Q_c and r_c are a weighting matrix and scalar respectively, is readily constructed [3]. The controller gain vector, \underline{k}_{x_o} is calculated according to

$$\underline{k}_{x_o} = (r_c + \underline{d}_o^T Q_c \underline{d}_o)^{-1} (C_o^T Q_c \underline{d}_o + P_c \underline{b}_o) \quad \dots(26)$$

where P_c is the solution of the stationary Riccati equation

$$P_c A_o + A_o^T P_c - (P_c \underline{b}_o + C_o^T Q_c \underline{d}_o) (r_c + \underline{d}_o^T Q_c \underline{d}_o)^{-1} (\underline{b}_o^T P_c + \underline{d}_o^T Q_c C_c) + C_o^T Q_c C_o = 0 \quad \dots(27)$$

Similarly, the Kalman gain, K_{f_o} is calculated according to

$$K_{f_o} = (R_f + J_o Q_f J_o^T)^{-1} (G_o Q_f J_o^T + P_f C_o^T) \quad \dots(28)$$

where P_f is the solution of the stationary Riccati equation

$$P_f A_o^T + A_o P_f - (P_f C_o^T + G_o Q_f J_o^T) (R_f + J_o Q_f J_o^T)^{-1} (C_o P_f + J_o Q_f G_o^T) + G_o Q_f G_o^T = 0 \quad \dots(29)$$

Q_f is the covariance matrix of the disturbance vector \underline{E} and R_f is the covariance matrix of the notional noise on the measurement \underline{y}_o .

3.3 Flexing Mode Controller

The flexing mode controller can be developed in a similar fashion to the rigid body controller, however, individual modes cannot be controlled in isolation. This is because the performance benefits are to be achieved by using higher wavenumber force distributions to control low wavenumber modes and this feature could not be employed using single mode control. The design therefore makes use of the flexing mode model represented by equation 10.

The major difference from the approach used for the rigid body controller is that the spatial wavenumber filter discussed in section 2 is to be included in the LQG cost function. The spatial filtering process can be represented by

$$\underline{\tilde{p}} = N \underline{u} \quad \dots(30)$$

where the projection matrix N provides the wavenumber components. Since the control force vector can be expressed as

$$\underline{u} = V \underline{u}_m \quad \dots(31)$$

equation 30 can be rewritten as

$$\tilde{\underline{p}} = NV\underline{u}_m = N(V_r\underline{u}_r + V_f\underline{u}_f) = \tilde{\underline{p}}_f + \tilde{\underline{p}}_r \quad \dots(32)$$

The wavenumber contribution from the rigid body control $\tilde{\underline{p}}_r$ is minimized using the non-reciprocal feature detailed in the previous section and the flexing mode design is aimed at finding a control distribution \underline{u}_f which minimizes $\tilde{\underline{p}}_f$. As each wavenumber radiates in a distinct frequency range, in practice a controller which minimizes a set of signals, \underline{p}_f , formed by high pass filtering $\tilde{\underline{p}}_f$ is sought.

By noting that

$$\underline{u}_f = V_f^T N^{-1} \tilde{\underline{p}}_f \quad \dots(33)$$

a similar operation to that defined in equation 20 is used by formulating \underline{u}_f as the output of the low pass filter

$$\begin{aligned} \dot{\underline{x}}_p &= A_p \underline{x}_p + B_p \underline{p}_f \\ \underline{u}_f &= C_p \underline{x}_p + D_p \underline{p}_f \end{aligned} \quad \dots(34)$$

Using similar filters on gap, acceleration and vibration signals as used in equations 18, 19 and 22 an overall description of the form

$$\begin{aligned} \dot{\underline{x}}_o &= A_o \underline{x}_o + B_o \underline{p}_f + G_o \underline{\epsilon}_v \\ \underline{y}_o &= C_o^T \underline{x}_o + D_o \underline{p}_f + J_o \underline{\epsilon}_v \end{aligned} \quad \dots(35)$$

can be obtained where the following definitions now hold

$$\underline{x}_o^T \triangleq [\phi_f^T \quad \dot{\phi}_f^T \quad \underline{x}_g^T \quad \underline{x}_l^T \quad \underline{x}_p^T \quad \underline{x}_v^T]$$

$$\underline{y}_o^T \triangleq [\underline{g}_l \quad \underline{a}_l]$$

With this description, the LQG controller

$$\begin{aligned} \underline{p}_f &= -\underline{K}_{x_o} \hat{\underline{x}}_o \\ \dot{\hat{\underline{x}}}_o &= A_o \hat{\underline{x}}_o + B_o \underline{p}_f + K_{f_o} (\underline{y}_o - C_o \hat{\underline{x}}_o - D_o \underline{p}_f) \end{aligned} \quad \dots(36)$$

that minimizes the cost function

$$J = E \left\{ \frac{1}{2\pi} \int_{-\infty}^{\infty} \underline{y}_o^*(j\omega) Q_c \underline{y}_o(j\omega) + \underline{p}_f(j\omega)^* R_c \underline{p}_f(j\omega) d\omega \right\} \quad \dots(37)$$

can be calculated in an identical fashion to that defined for the rigid body control (equations 26-29).

3.4 Summary of Low Frequency Control Algorithm

The low frequency control algorithm can be summarised by the following steps

- (i) Obtain gap and acceleration measurement vectors, \underline{y} and \underline{a} respectively.
- (ii) Form modal gap and acceleration vectors according to

$$\begin{bmatrix} \underline{y}_r \\ \underline{y}_f \end{bmatrix} = V^T \underline{y}$$

$$\begin{bmatrix} \ddot{\underline{\Phi}}_r \\ \ddot{\underline{\Phi}}_f \end{bmatrix} = V^T \underline{a}$$

and non-reciprocal signal \underline{g} , for each rigid body mode (equation 12).

- (iii) Form filtered measurement vector, \underline{y}_o , for each rigid body mode and the flexing body modes in overall description (equations 23 and 35) using filters described by equations 17-19.
- (iv) Calculate notional control signals for rigid body modes and flexing modes using equations 24 and 36 respectively.
- (v) Form modal control signals, $\underline{u}_{r,i}$ and \underline{u}_f according to equations 20 and 34.
- (vi) Apply local control vector

$$\underline{u} = V \begin{bmatrix} \underline{u}_r \\ \underline{u}_f \end{bmatrix}$$

In order to calculate the controller parameters the modal parameters in the raft dynamic model (equation 1) and the spatial filter N (equation 30) must be provided. These are defined by the physics of a particular installation and must adequately describe its behaviour.

Other parameters required are the characteristics of the filters defined by equations 17-22 (for both rigid body and flexing modes) and the cost function weighting matrices. These are tuning parameters that can be selected by the design engineer to achieve a particular performance.

4 High Frequency Control Algorithm

As outlined in 2.3, the high frequency control algorithm takes a very different form from the low-frequency algorithm, since it is implemented as a low-gain feedforward process, rather than high-gain feedback control. Since the feedforward gains are small compared to unity, the possibility of feedback instability should not arise, but the precision of operation must be maintained using an adaptive update operation, which optimizes the control filters to minimise vibration transmission to the hull structure.

This high frequency algorithm can be regarded as an "outer loop" which is added external to the existing low-frequency algorithm. An important requirement is that it should not interfere with or compromise the operation of the existing low frequency controller.

The description of the algorithm will now be given using a sequence of block diagrams, enabling the various requirements to be added-in progressively. The specific implementation of each block can be achieved in a number of different ways; in the present case it is proposed that these blocks should generally take the form of simple FIR (finite-impulse- response) filters, with time-domain adaptive update where appropriate.

4.1 Basic Feedforward Adaptive Algorithm

The basic operation associated with an LMS adaptive algorithm is shown in Figure 4.1 The objective is to control an FIR filter so as to reproduce in antiphase a given process, by modifying the filter tap weights (coefficients) in response to the observed performance error. The simplest form of adaptive update algorithm is the Widrow algorithm, which updates the array of tap-weights based on the instantaneous cross-correlation between the filter input data array $X(t)$ and the observed error $\epsilon(t)$.

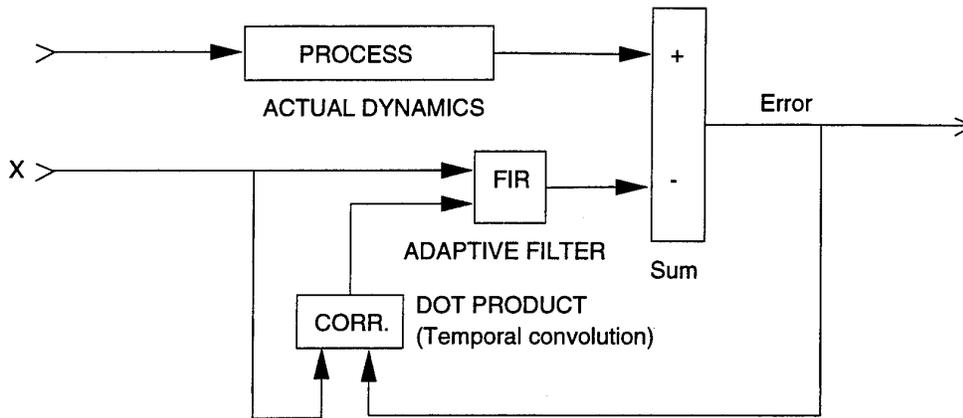


Figure 4.1: Basic FeedForward Algorithm

The objective is:

$$\text{Minimise } \overline{(y - F \cdot X)^2} \text{ by choice of filter coefficients } F$$

and noting $\nabla_F (y - F \cdot X)^2 = -2(y - F \cdot X) X$

the instantaneous gradient is $-2\epsilon X$ where $\epsilon = (y - F \cdot X)$

which can be interpreted as the cross-correlation between the filter input data X and the observed error ϵ .

The optimal performance is achieved when all components of the cross-correlation are driven to zero, so the adaptive update seeks to modify each tap by an amount proportionate to any residual non-zero correlation.

$$\text{i.e. } \delta F = 2\mu\epsilon X$$

4.2 Correction for Downstream Propagation

In practice, observation cannot take place immediately downstream of the adaptive filter, but rather after the combined signals have propagated through a further downstream process. Thus the x -input to the filter is no longer accurately time-aligned with the observed error, and steps must be taken to realign this data. This can be achieved using the "filtered- x " algorithm, which passes the x -input data through a simulation of the downstream process, prior to performing the correlation update. This restores the alignment, and ensures that the residual error is accurately correlated with the expected response at the observation point. (Figure 4.2)

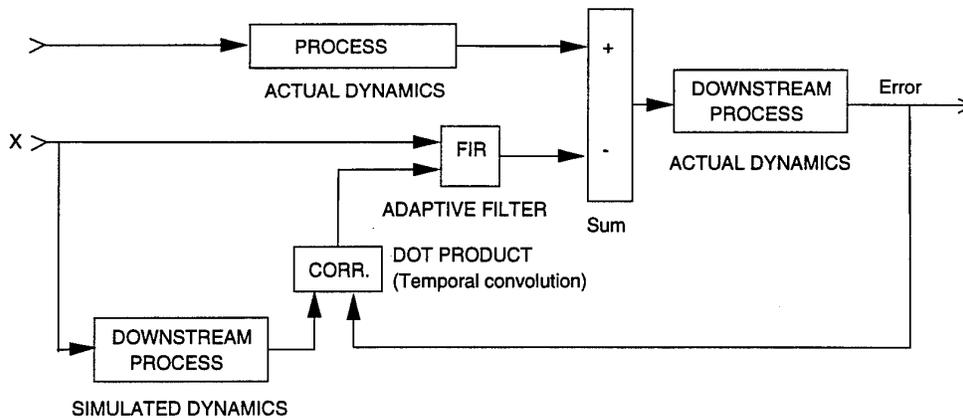


Figure 4.2: Filtered-X Adaptive Algorithm

This additional simulation requirement may itself be implemented using a separate adaptive filter; this involves injecting test signals into the downstream process, and observing the associated response at the observation point. An adaptive filter can be used to characterise this transmission directly. Once it has adequately converged, it can be "frozen" and inserted as the downstream simulation element of Figure 4.2.

4.3 Deconvolution of Downstream Interaction Paths

The algorithm described above enables the effects of temporal misalignment arising from downstream propagation to be compensated. But in the present context one must also take account of the fact that the error responses are induced by a multiplicity of separate mount-transmission paths. Thus in practice, one seeks to update a parallel array of adaptive filters, each simulating a separate mount-transmission path, by observing a corresponding array of error responses.

Under such circumstances, one must implement the correlation process with respect to the spatial characteristics of the observed response, as well as its temporal behaviour. Effectively, in order to minimise the response due to any specific mount path, one seeks out those components of response which simultaneously exhibit the correctly matching spatial and temporal pattern.

This can be achieved using the multi-degree-of-freedom filtered-x algorithm as shown in Figure 4.3. Each input to each adaptive filter is passed through a simulated downstream process, in order to generate an array of outputs corresponding to the responses which would be expected at the various observation points.

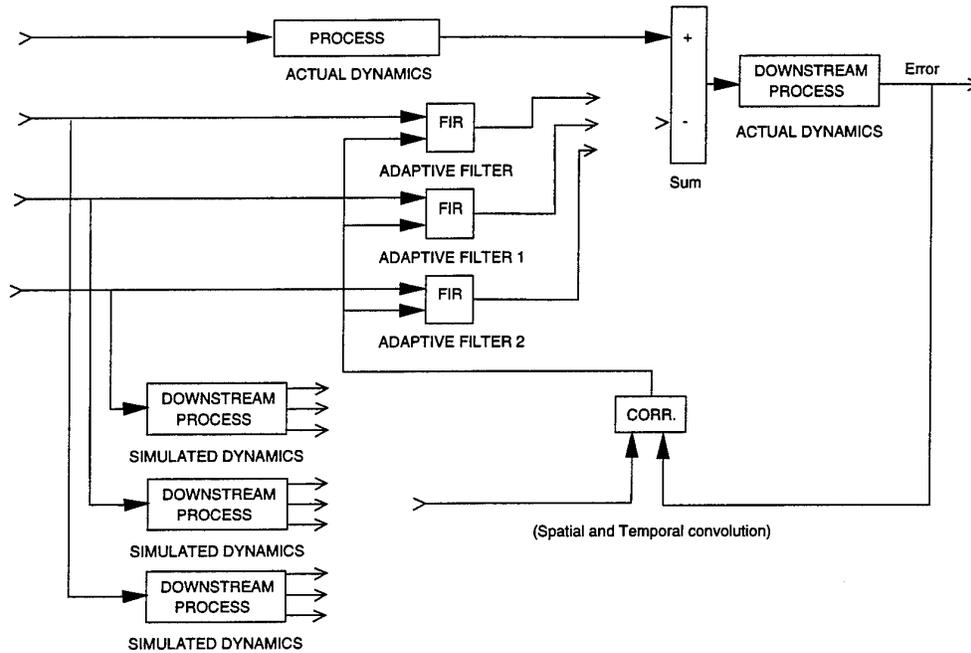


Figure 4.3: Multi-Degree-of-Freedom Adaptive Algorithm

Each individual array of simulated responses is then correlated with the full array of observation error-responses. This defines a set of scalar products of space-time correlations, which are then used to define the coefficient updates for each corresponding adaptive filter.

It can be seen that the computational workload associated with this adaptive update can very rapidly exceed the workload within the primary filtering process, since the simulation of the downstream array dynamics involves a multiplicity of cross-paths. For example, for a 48-input, 48-output system, there are 48 feedforward paths, but depending on the severity of the cross-interactions, the downstream filtered-x process can involve simulating up to 48 x 48 transfer functions.

It should be noted, however, that various simplifications can be applied to this downstream process. In particular, if the updates are not applied on a sample-by-sample basis, but rather in sequential rotation, the workload can be substantially reduced (at the expense of longer convergence times). Moreover, many of the downstream processes share common dynamical features (e.g. resonant frequencies), so it can be more economic to implement these processes using IIR (infinite-impulse-response) filters, with common poles.

But the most effective technique is undoubtedly to try to ensure that each downstream observation only responds to a local subset of inputs, thus enabling the 48 x 48 matrix array of transfer functions to be condensed to a banded or near-diagonal array.

In the present context, it is proposed to achieve this by using the magnet flux response as a measure of vibration transmission. Clearly, the flux induced within any particular mount is not sensitive to the action of other mounts, whereas the vibration response at the base of each mount will be influenced by the combined actions of all the mounts.

For this reason, each mount is instrumented with passive flux-sensing coils, intended to provide clean, low-interaction measures of the residual transmitted forces.

Thus, it is intended that the downstream process should be simplified to yield a diagonal or near-diagonal form with only local interactions. However, the practical extent of such interactions can only be determined by experiment. Nevertheless, the algorithm structure presented in Figure 4.3 can in principle handle any such level of interaction.

4.4 Prevention of Interaction with Low-Frequency Algorithm

Since the low frequency algorithm has been carefully designed to achieve optimal performance while satisfying a number of constraints, one does not wish the high frequency algorithm to introduce subsequent low frequency perturbations which compromise this careful optimisation.

In this respect, it is not considered sufficient simply to high-pass filter the data which is supplied to the adaptive algorithm, since this will mean that there is no information at all to constrain its low-frequency characteristics; the latter will then be established at random according to the specific circumstances. Indeed the algorithm will almost certainly seek to equalise the effects of any high-pass filters, and will seek to boost its low-frequency gain in order to do so. So it is important that relevant low-frequency information is still made available to the algorithm, in such a manner that if it departs from its constrained characteristic, errors are registered and positive corrections are applied.

This requirement can be achieved using the circuit of Figure 4.4. It can be seen that this is a basic filtered-x algorithm, but the output of the adaptive filter is split into complementary high-pass and low-pass components (the sum of these complementary filters is unity). While the high-pass component is then injected into the actual system, the low-pass component is passed through a parallel simulation of the downstream process, and then subsequently subtracted-off from the error-sensor.

Providing the simulation dynamics is an accurate representation of the actual downstream process, the overall dynamics of this circuit are identical to those of a standard filtered-x adaptive algorithm. The fact that the output is split into two complementary high-pass and low-pass components does not affect any of the data supplied to the algorithm.

The adaptive filter will then correctly seek to reproduce both the low-frequency and high-frequency characteristics of the desired process. But only the high frequency component, with controlled high-pass cut-on, is introduced into the real cancellation process. The low-pass components are kept out of the real system, but they are re-introduced into the error path prior to the update.

So the error-update is an all-pass process, which seeks to maintain accurate alignment over the full frequency range, but only the constrained high frequency components are introduced into the actual dynamics of the real system.

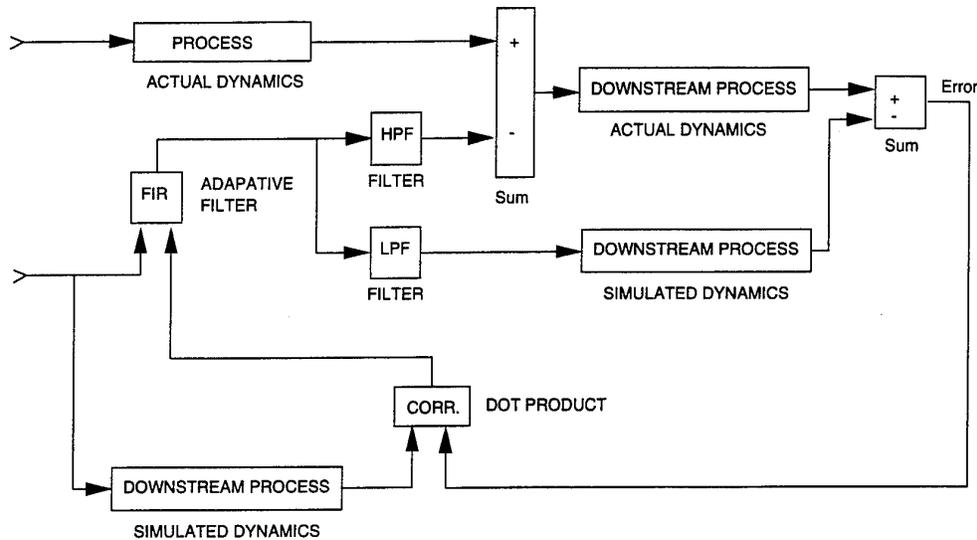


Figure 4.4: Adaptive Algorithm with Low Frequency Constraint

Having thus identified an algorithm which will achieve the desired objective, it is appropriate to try to simplify it.

In Figure 4.5, this is achieved by noting that the parallel low-pass components could equivalently be derived from the true signals emerging from the actual process, rather than from the output of an adaptive filter which is seeking to reproduce this process. This relieves the adaptive filter of the need to converge so accurately at low frequencies, since the requisite low-frequency component of output data is now automatically available to full precision.

The requirement now becomes that the adaptive filter should accurately converge at high frequencies; moreover it will no longer seek to compensate for, or "equalise", the effects of its own output high-pass filter, since the overall process which it must now reproduce is also a matching high-pass process. Yet, if any significant low frequency error is generated, this will propagate through to the error observation point, where it will be taken into account and corrected in the correlation update process.

In its present form, this algorithm requires a matching high-pass filter to be placed in the filtered-x portion of the circuit. But this filter roll-off could be flattened-out or removed completely, provided a compensating low-frequency weighting correction is applied to the error- output path as indicated. The use of such a filter would recover the all-pass characteristics of the adaptive update, and would serve to tighten-up still further the low-frequency precision of operation.

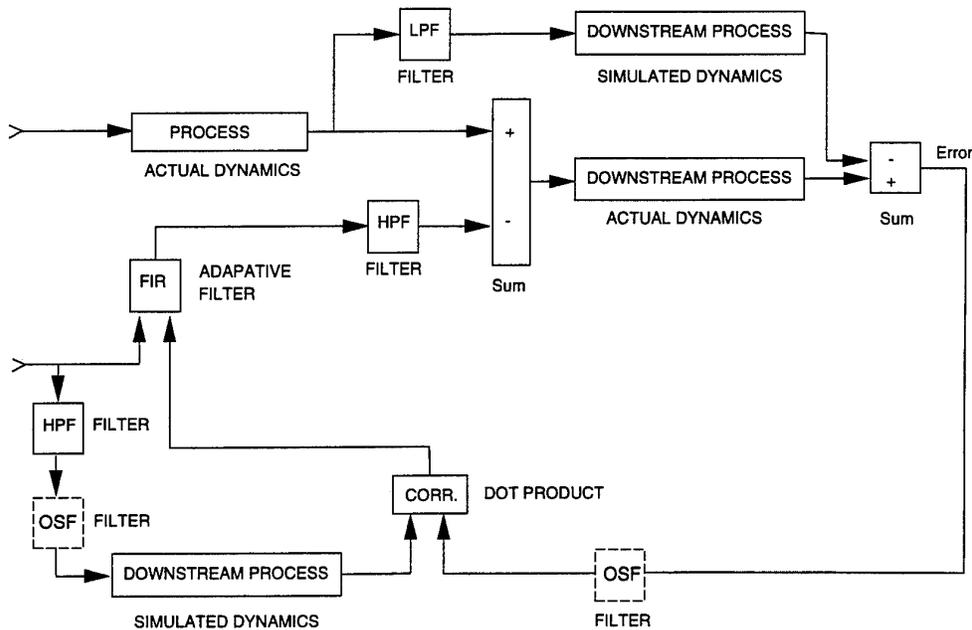


Figure 4.5: Adaptive Algorithm with Simplified Low Frequency Constraint

4.5 The Overall High Frequency Algorithm

Having discussed all the separate requirements for the overall algorithm in sections 4.1 to 4.4, it is appropriate to combine these together to form the overall algorithm for the system. This algorithm is set out in Figure 4.6.

It can be seen that the central portion of the circuit contains the primary low-frequency feedback algorithm, which has been described in Section 3. For the high frequency algorithm, the acceleration response of the raft is detected, and used as input to the adaptive FIR filters which simulate the mount transmission characteristics. The outputs of the FIR elements are high-pass filtered and fed-forward for addition into the active mount drive signals.

The adaptive update is performed by applying a matching high-pass filter to the acceleration response, and feeding it into a simulation of the dynamics of the base supporting structure. This simulation calculates the separate response at each of the observation points which would be generated by the individual accelerometer signals, and correlates this information with the actual error signals observed in practice. This correlation defines the adaptive update for the FIR filters.

The low-frequency constraint is applied by passing the actual low-frequency control signals through a separate simulation of the base dynamics, and subtracting the simulated response signals from the actual observed response of the base.

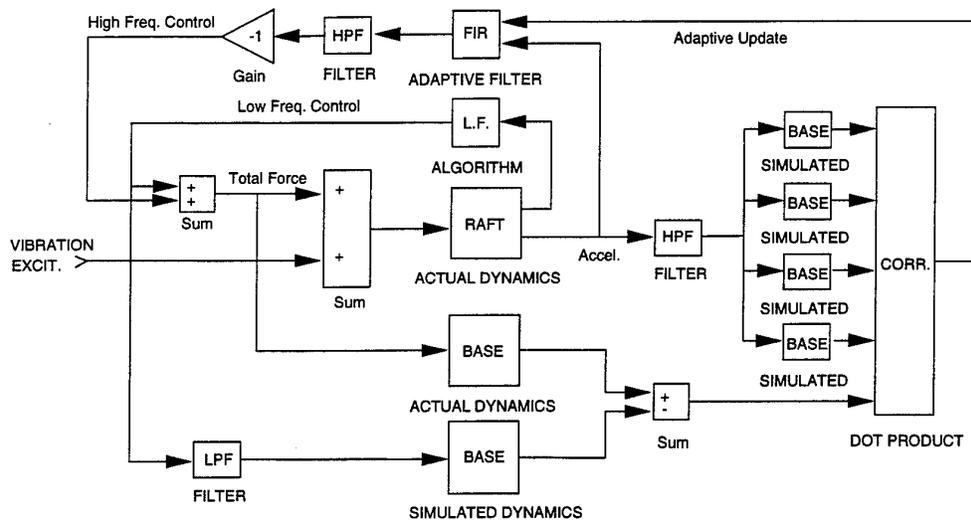


Figure 4.6: Overall High Frequency Algorithm

In this form the algorithm contains all the features which have previously been identified separately. Specifically, the downstream response is simulated by a set of filtered-x paths, the correlation is performed taking into account spatial cross-interactions in the response, and the algorithm is constrained so that it does not seek to modify the existing low-frequency operation of the system.

It should be noted that specification of the separate components in the algorithm is largely defined by the dynamical characteristics of the actual installation, so there is little scope for ambiguity in this respect. The principal variable which can be chosen independently is the exact specification of the complementary high-pass/low-pass filters. The characteristics of these filters will define the precise cut-on point of the algorithm, while the precision with which these filters approximate to unity in their pass-band will define the limiting precision of the algorithm.

5 Numerical Calculations and Simulations of Performance

A series of calculations and simulations have been undertaken to establish the theoretical performance of the algorithms which have been described in Sections 3 and 4. These calculations have taken two distinct forms. First, in the case of the low frequency algorithm, the frequency response associated with a particular choice of control parameters was calculated, having assumed a dynamic system with given frequency response characteristics.

This same system was then simulated as a time-stepped model, in which the control algorithm ran exactly as it would in practice. This enabled the stability of the algorithm to be established. Finally, the output time series from this simulation was analysed to determine the frequency response characteristics which were actually being achieved; these could then be compared with the previous off-line calculations.

In the case of the high frequency model, it was inappropriate to perform off-line calculations, since the adaptive behaviour is an evolutionary process. All these calculations were therefore performed as time-stepped simulations, with the low-frequency algorithm running in parallel as it would in practice.

Once again, the output time traces were then analysed to determine the frequency responses before and after adaption. In this way, the stability and improvement in performance associated with the adaptive algorithm could be established.

5.1 Theoretical Basis for Raft Dynamical Model

A first requirement in the simulation was to establish a suitable dynamical model for the raft. This was defined so as to represent the actual raft structure of Figure 1.2, reproduced as Figure 5.1, but only the vertical components of flexure were simulated. It was assumed that the horizontal flexure would be dynamically similar, but (in theory) such motions should be fully decoupled from the vertical motions.

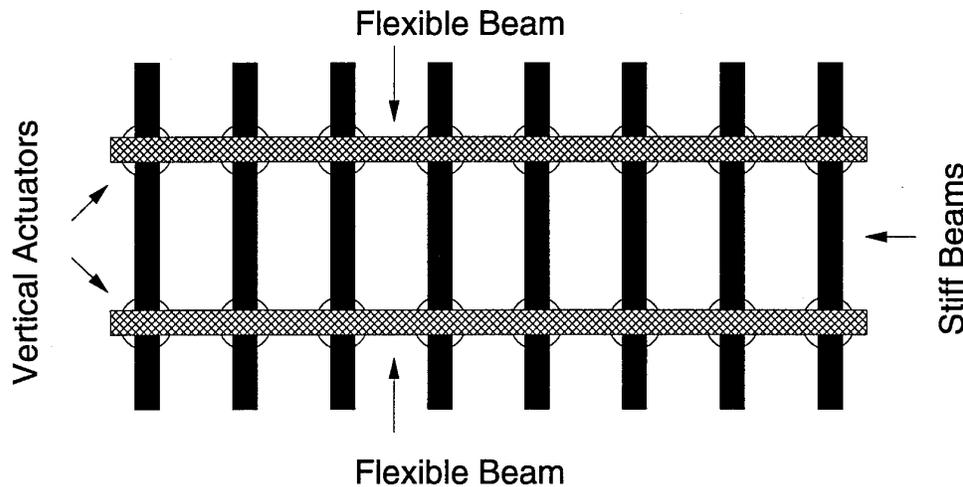


Figure 5.1: Experimental Raft Structure

The start point for the simulation was the partial differential equation for a simple beam with free-free boundary conditions. The exact analytic solutions for such a beam driven by a point force at arbitrary location were used to define the frequency response characteristics for each of the longitudinal and transverse beams of Figure 5.1

These beams were then considered to be coupled together with pin-joints, exerting vertical forces only, and the complete set of transfer functions for the coupled structure were then calculated.

In order to progress the control calculation and subsequent real-time simulations, it was then necessary to define an equivalent state-space formulation of this model. In this respect, it was convenient to make use of the same identification procedures as might be used in practice. Using the transfer function data base, the response of the model to simulated time series was calculated, and this data was used to extract a parallel set of modal responses, each defining a second-order modal equation of motion. These latter equations were then cast into conventional first-order state-space form, with associated modal matrices defined from the corresponding mode shapes.

The accuracy of this state-space description was then confirmed by comparing the transfer function response characteristics of the state-space model, with the corresponding transfer functions derived from the continuous beam representation.

Three different raft models were simulated in this way, and are designated respectively a soft raft, a medium raft, and a stiff raft. The difference between these models lay in the fundamental resonant frequencies of the longitudinal beams. The "medium" raft was chosen to be representative of a typical very large machinery raft, having its primary longitudinal resonant frequencies over the frequency range 6Hz-40Hz. The "stiff" raft represented a much stiffer structure, with corresponding resonances over the range 15Hz-200Hz, while the "soft" raft was a structure with very low frequency longitudinal resonant frequencies spanning the range 3Hz-20Hz.

These state-space descriptions of the raft models were then used to calculate the optimal feedback and Kalman filter gains according to the procedures set out in Section 3. The same equations were also used to define the discrete-time evolution operators for time-stepped simulation of the overall control algorithms.

The results of the calculations will now be presented, starting with the off-line frequency response characteristics, followed by the time-stepped simulations of the low frequency algorithm, and finally the complete low-frequency/high-frequency algorithms.

In order to define a reference datum for the calculations, the response of an equivalent structure mounted on conventional 4Hz passive mounts was also calculated. In this case, the control force was defined as being directly proportional to the local mount displacement and velocity. The displacement-force gain was chosen to represent a passive mount with stiffness corresponding to 4Hz (rigid body) resonant frequency, while the velocity-force gain was chosen to give a damping factor equivalent to a typical rubber mount (approx 0.05).

5.2 Off-Line Frequency Response Calculations

The first set of results show the acceleration response to vibration of the three different raft models using optimal low-frequency control, compared to the corresponding response of the same structures on passive mounts. The measured parameter is the frequency distribution of the mean square acceleration of the raft structure, in response to vibration excitation from a random point force.

Figure 5.2 shows the response of the soft raft, Figure 5.3 shows the response of the medium raft, and Figure 5.4 shows the response of the stiff raft. It is immediately apparent that in the case of the actively controlled soft raft, the first 3 low frequency raft resonances are completely suppressed, and the mean acceleration response approximates nearly to a straight line. This means that the raft is behaving almost as a rigid body (the latter would yield a constant acceleration mass-line), with almost no evidence of low-frequency resonances. In contrast, it can be seen that the response of the passively mounted structure is extremely resonant, with amplification of up to 20dB over the actively controlled structure. This corresponds to motions which are a factor of 10 greater in amplitude.

The 4th resonance is not suppressed, although it is still at a comparatively low frequency (20Hz), but this is primarily due to the fact that there are only 8 mounts supporting the structure. It is believed that with an increased distribution of mounts this resonance could also be controlled.

Similar results are obtained for the medium raft (Figure 5.3), but in the case of the stiff raft (Figure 5.4), the control system can only partially suppress the first resonance, and thereafter the response is very similar to the passive raft. The explanation is quite straightforward; suppression of resonances in the stiff raft requires much larger forces to be exerted at higher frequencies, and the active mounts cannot supply such forces without compromising the isolation of the structure from the hull. Thus it is better to "let these resonances go". On the other hand, for the soft raft at very low frequencies, the active mounts can indeed supply the requisite control forces while still maintaining overall isolation of the structure.

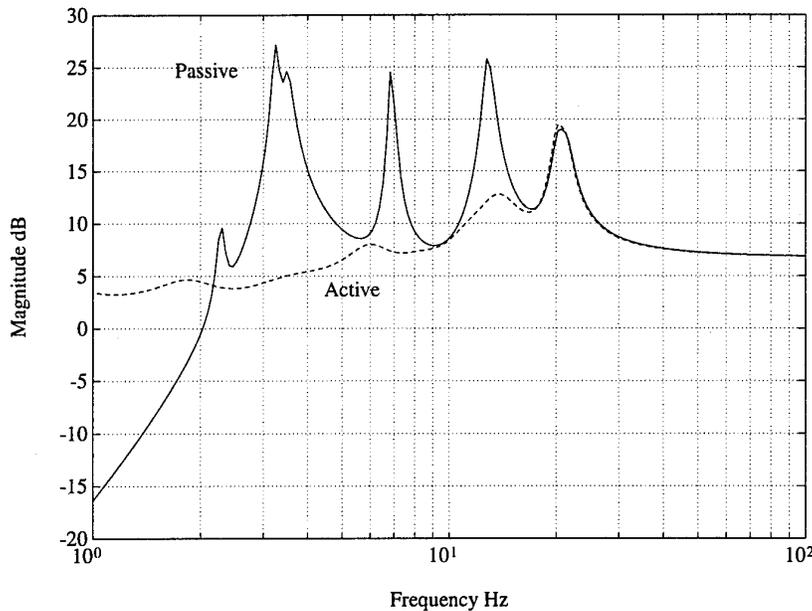


Figure 5.2: Mean Square Acceleration - Soft Model (Off-line)

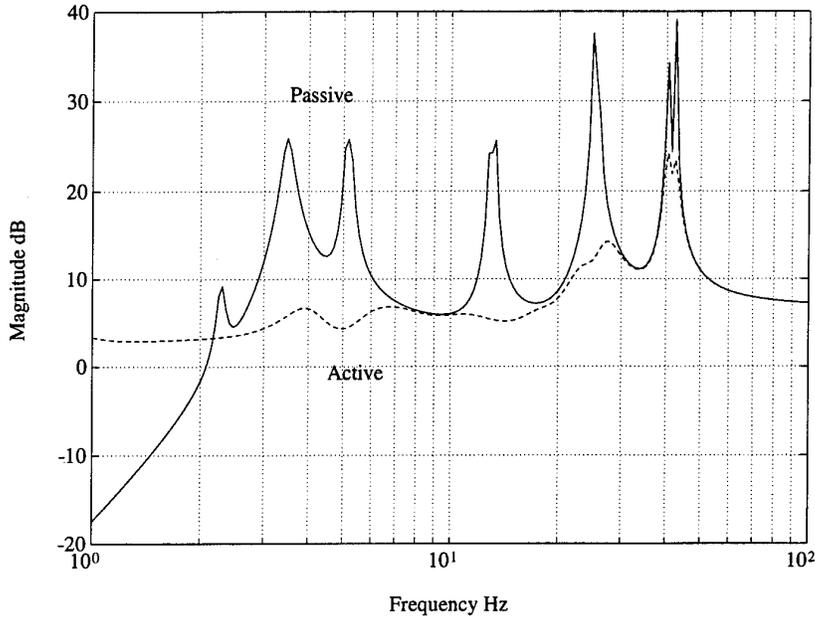


Figure 5.3: Mean Square Acceleration - Medium Model (Off-line)

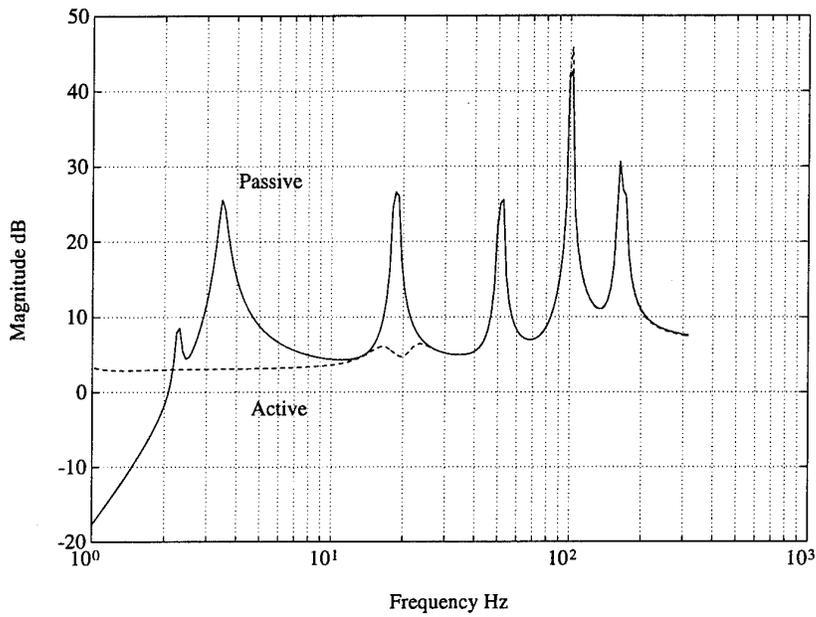
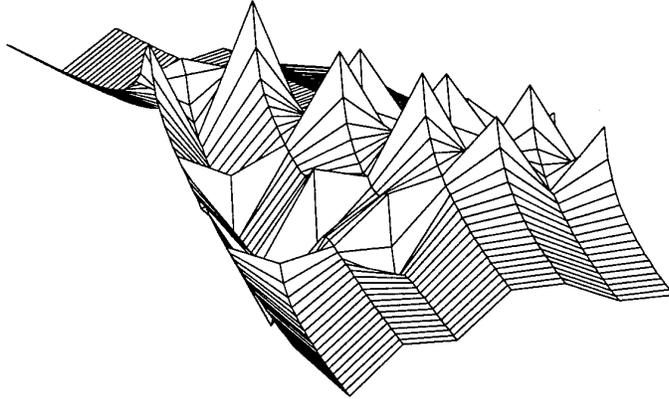


Figure 5.4: Mean Square Acceleration - Stiff Model (Off-line)

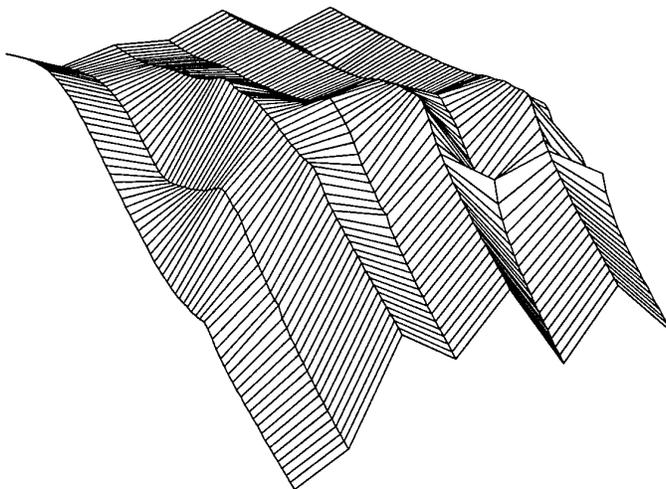
It is clear from these results that the ability to suppress low frequency resonances is related to the frequency range of the resonances; for frequencies up to 30Hz-40Hz, very effective suppression is possible, but at higher frequencies it is better not to even attempt to suppress the structural motions. Thus this active control strategy is most effective on flexible structures which have a number of low frequency resonances, but is less effective on stiffer structures. However, it is precisely this feature that one wishes to exploit in the case of large, extended machinery raft installations, where achieving adequate passive stiffness is extremely difficult.

Corresponding to the calculations of Figures 5.2 to 5.4, the control force distributions were examined, and their spatial characteristics determined. This involved calculating the frequency response of the outputs of the spatial wavenumber filters. These wavenumber filters effectively measure the extent to which the control force distribution will couple into the hull structure; at low frequencies, low wavenumber outputs will couple well, but the higher wavenumbers tend to self-cancel and couple poorly as a result. However, as the frequency increases, this self-cancellation effect reduces, so that at the higher frequencies all wavenumbers would be expected to couple with equal efficiency.

The results of these calculations are most easily presented as 2- dimensional wavenumber-frequency plots, in Figures 5.5, 5.6 and 5.7. The interpretation of these figures is that the frequency response of each wave-filter output is plotted side-by-side, with the lowest wavenumbers to the left, and the highest wavenumbers to the right.

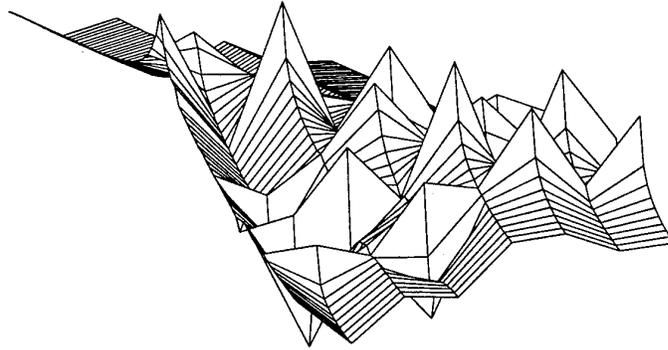


Passive

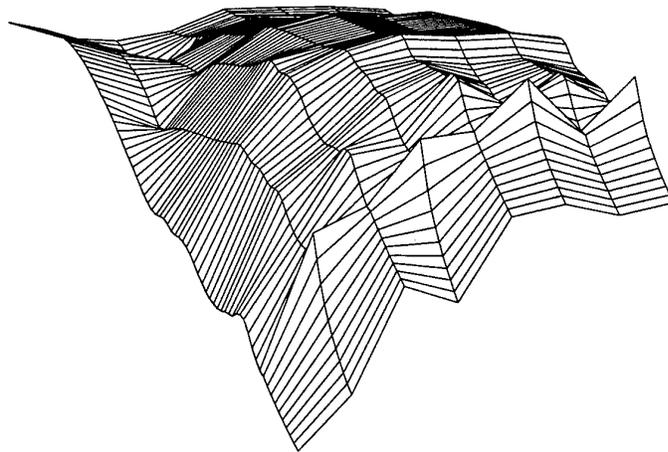


LQG

Figure 5.5: Wavenumber Distribution - Soft Model (Off-line)

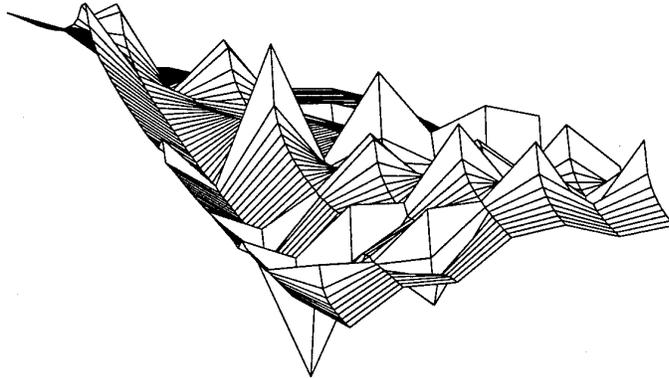


Passive

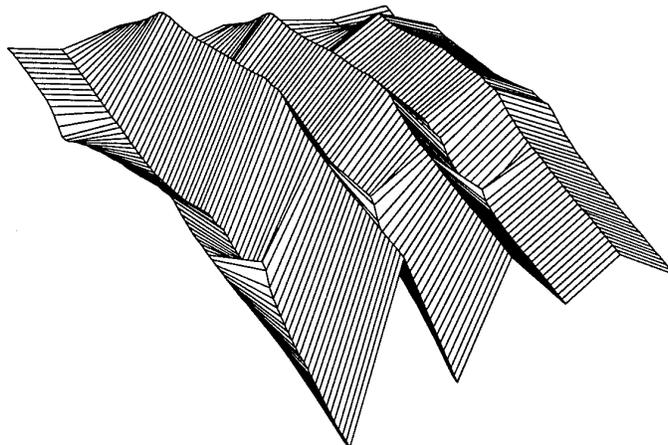


LQG

Figure 5.6: Wavenumber Distribution - Medium Model (Off-line)



Passive



LQG

Figure 5.7: Wavenumber Distribution - Stiff Model (Off-line)

In each case, the distributions for the actively controlled raft are compared with the distributions for the equivalent passively mounted structure. It is immediately apparent that for the soft and medium rafts, the response is much smoother than in the passive case; the passive rafts have very "peaky" distributions, each peak corresponding to the forces transmitted by the resonances.

The most important feature of these 2-D curves is the fact that the lower left triangular area of each active response is substantially flatter and lower than the equivalent regime of the passive response. This demonstrates that the low-wavenumber responses are substantially suppressed compared to the higher wavenumbers, and indeed are significantly lower than for a passively mounted raft. It is these components that are associated with efficient vibration transmission into the receiving hull structure, so it can be seen that the control algorithm does indeed preferentially suppress these components of transmission.

In the case of the stiff raft, however, these effects are not obviously present; it is immediately clear that the control force distribution is not so easily "flattened out".

In order to quantify the potential improvements implicit in Figures 5.5 and 5.6 (the soft and medium rafts), it is first necessary to define a typical cost function associated with each wavenumber distribution. Such a cost function is shown in Figure 5.8. The lowest wavenumber is assumed to couple efficiently (ie 1.0) over the full frequency range 0-100Hz, but the higher wavenumbers are assumed to have linear cut-ons, each starting at a progressively higher frequency over the range to 100Hz. This is typical of the radiation characteristics of the higher-order spatial wavenumber distributions.

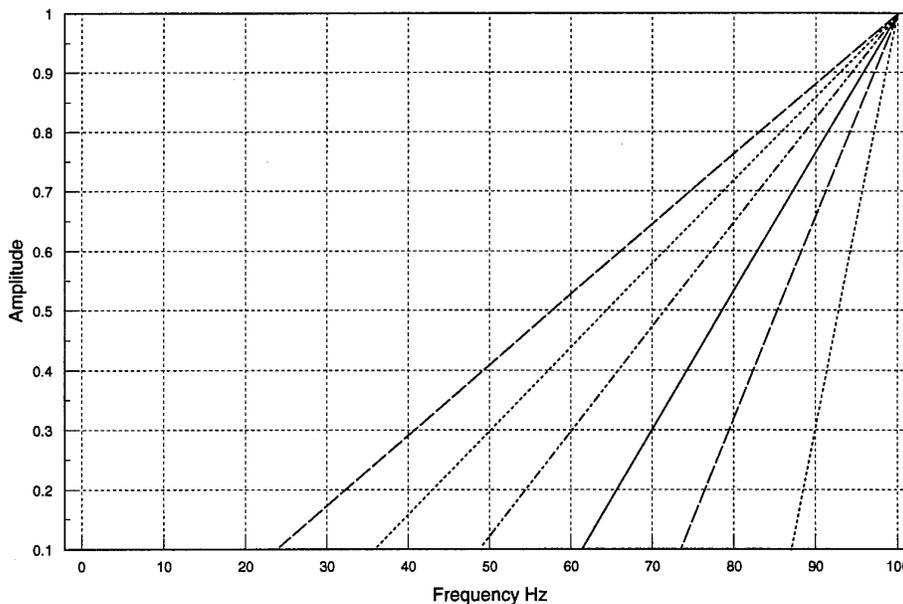


Figure 5.8: Typical Wavenumber Cost Function

One can then integrate the responses of Figures 5.5 and 5.6 with respect to this wavenumber cost-function, to obtain a measure of the performance characteristic of each distributed control array. The results are shown in Figures 5.9 and 5.10. It is apparent in both cases that the transmission efficiency is very much reduced compared to the equivalent passively mounted structure.

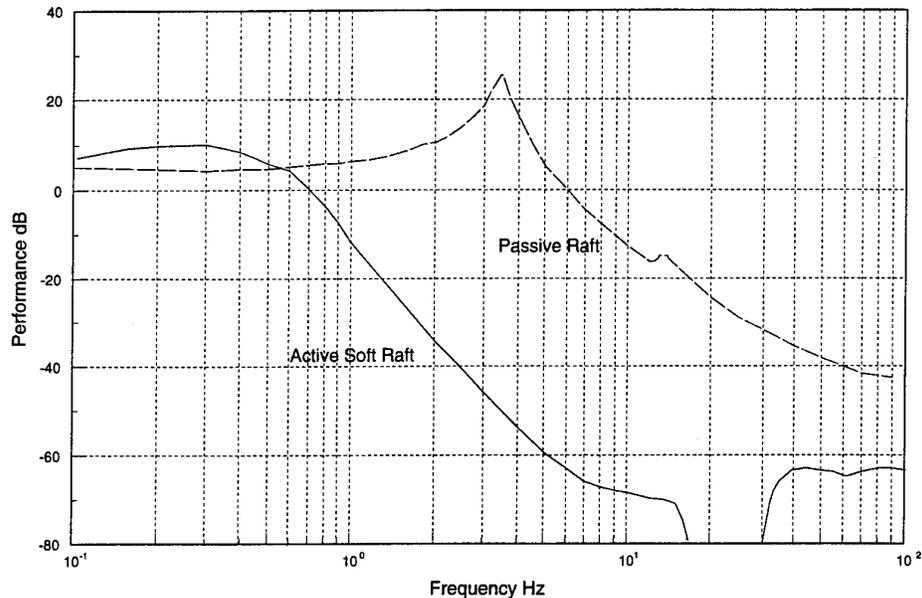


Figure 5.9: Radiated Noise Performance - Soft Model

It is particularly important to note that the performance improvements are most marked at the lower frequencies, where the redistribution of the forces into the higher wavenumber components is undoubtedly most effective. At the higher frequencies, as progressively more spatial wavenumbers cut-on, there are fewer wavenumbers which couple poorly, and the freedom of manoeuvre is restricted accordingly.

Nevertheless, the performance improvements which are predicted are very substantial indeed, amounting to 40dB-60dB over much of the frequency range. This shows that these techniques offer an opportunity for very substantial improvements over passive installations; even if these gains are only partially realised in practice, it is clear that there is the opportunity for constructing vibration isolation platforms with very much enhanced performance.

The essential feature involved is that the actively controlled raft is constrained to move almost as a rigid body, and this rigid body motion is in turn matched accurately to the vibration force inputs. Thus it becomes a very efficient platform for absorbing the momentum associated with these excitation forces.

At the same time, the force distributions which are used to stiffen the flexural motions of the raft are themselves preferentially chosen to have high wavenumber characteristics, and thus they have a very poor coupling efficiency into the receiving hull structure.

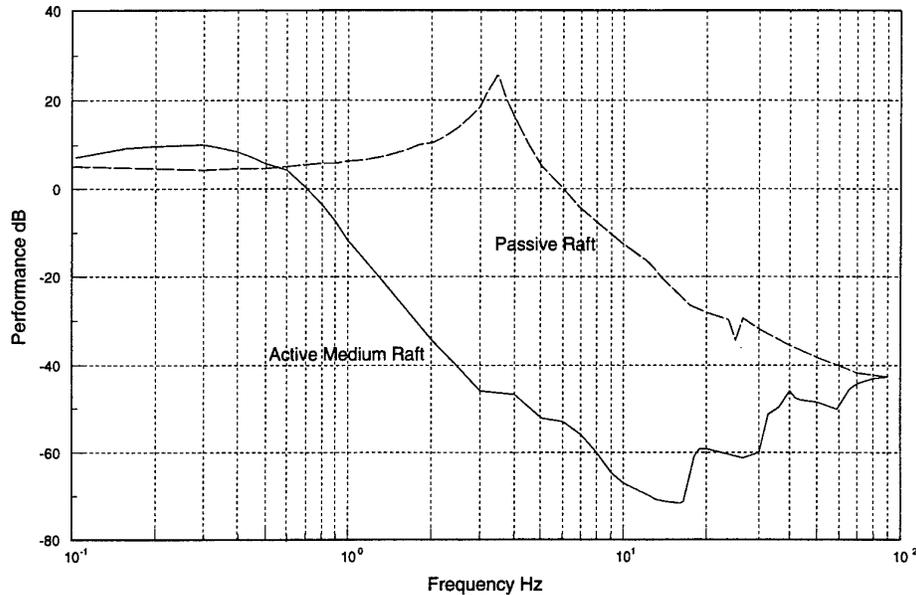


Figure 5.10: Radiated Noise Performance - Medium Model

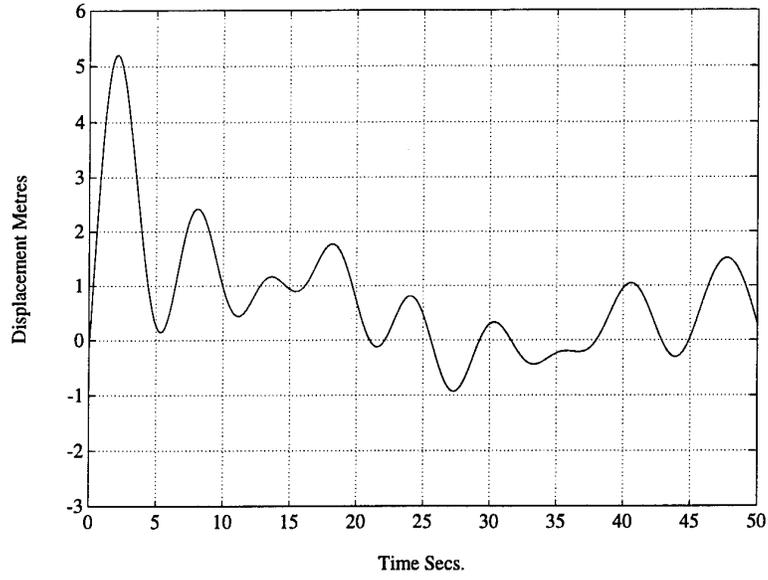
These performance improvements are not obtained at the expense of a raft which is poorly located under seaway motions; on the contrary, the results in the next section will show that its seaway tracking performance is extremely good.

5.3 Simulations of Real-Time Low-Frequency Response

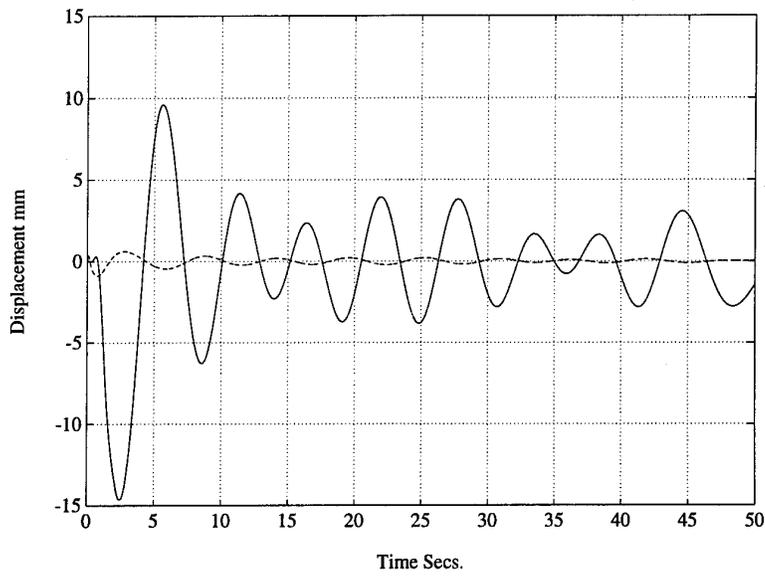
Having established the off-line frequency responses associated with the low frequency control algorithm, a time-stepped simulation of the structural response and control system operation was then performed. The purpose of this was to ensure that the control algorithms were indeed stable, while enabling their precise implementation to be simulated.

The optimal control parameter calculation had involved separating out the rigid body modes from the flexural modes, as described in Section 3, so the very low-frequency seaway response characteristics of the 3 separate raft models were identical. A typical response to random seaway disturbance is shown in Figure 5.11. This shows a hull input motion representing a peak displacement of 5 metres, and an average displacement of 1-2 metres.

The resultant motion of a passively mounted raft relative to the hull would amount to some 15mm peak, with average displacements of 3-5mm. In contrast, the actively mounted raft shows displacements which at all times are within 1mm, and generally are considerably less). Thus it is clear that the actively controlled raft can track and follow the hull motions extremely accurately.



Seaway Disturbance



Response

Figure 5.11: Response to Seaway Disturbance

For these same time-stepped models, the corresponding response of the raft acceleration to vibration forces was simulated. The resultant time series outputs were then frequency analysed, and the mean square raft acceleration calculated. The results are shown in Figures 5.12, 5.13 and 5.14. These results can be compared directly to the corresponding Figures 5.2, 5.3 and 5.4 for the off-line frequency calculations, and it can be seen that good agreement is obtained. Thus the time-stepped simulations are indeed stable, and reproduce exactly the behaviour which was predicted from the frequency response calculations.

The resonant responses in Figure 5.14 appear somewhat sharper than Figure 5.4; quite simply, the time-series transform resolution in Figure 5.14 was higher, with the result that the resonances are more precisely resolved.

Similarly, the time-series outputs of the spatial wavenumber filters were derived and frequency-transformed. The results are shown as 2- dimensional wavenumber-frequency plots in Figures 5.15 and 5.16 for the soft and medium rafts respectively. These can be compared directly with the corresponding Figures 5.5 and 5.6; once again, the time-series simulations have reproduced accurately the characteristics of the off-line calculations.

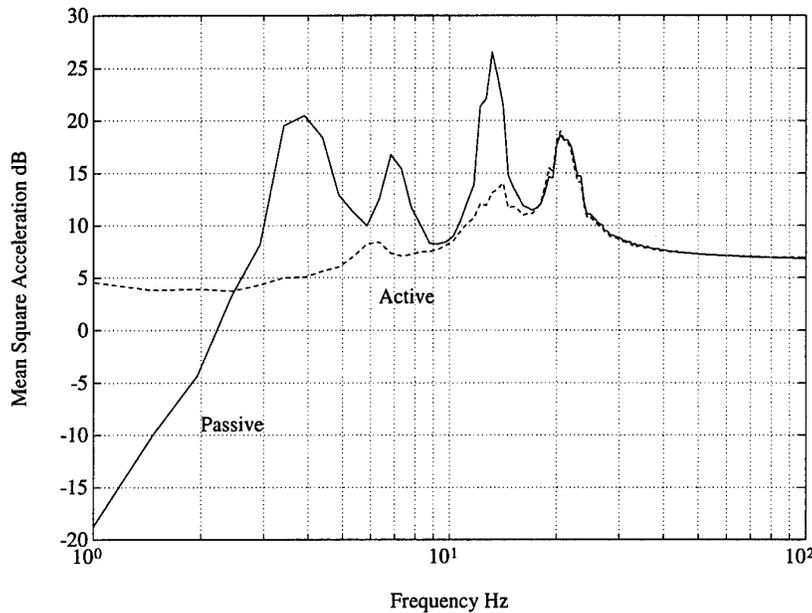


Figure 5.12: Mean Square Acceleration - Soft Model (Real-time)

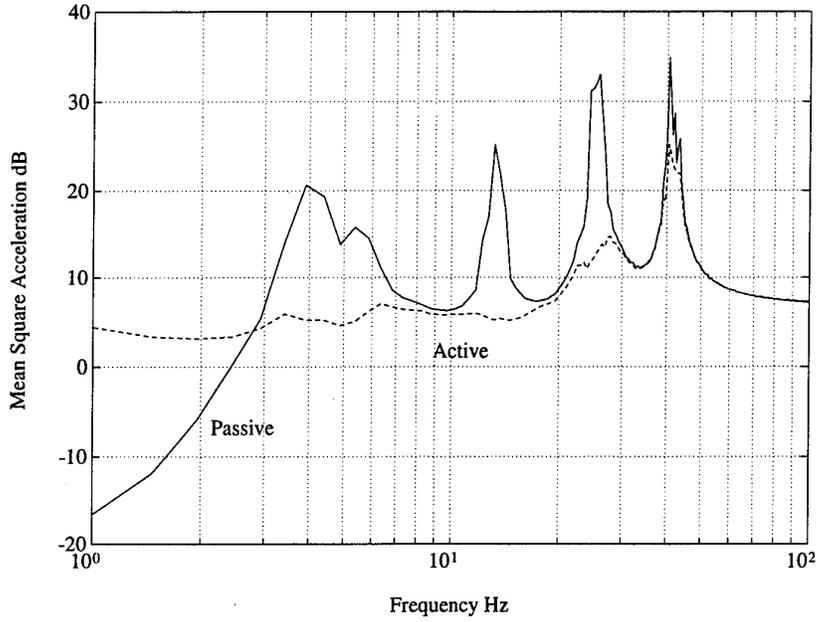


Figure 5.13: Mean Square Acceleration - Medium Model (Real-time)

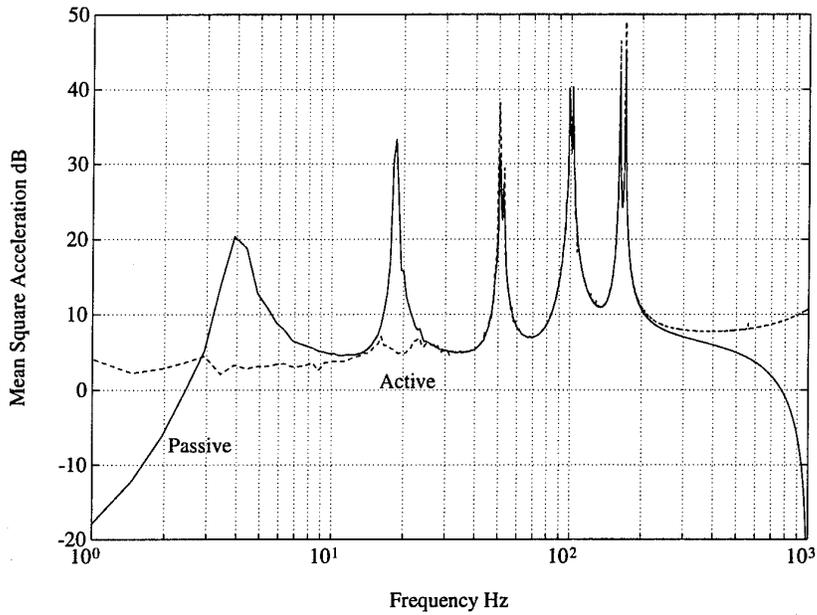


Figure 5.14: Mean Square Acceleration - Stiff Model (Real-time)

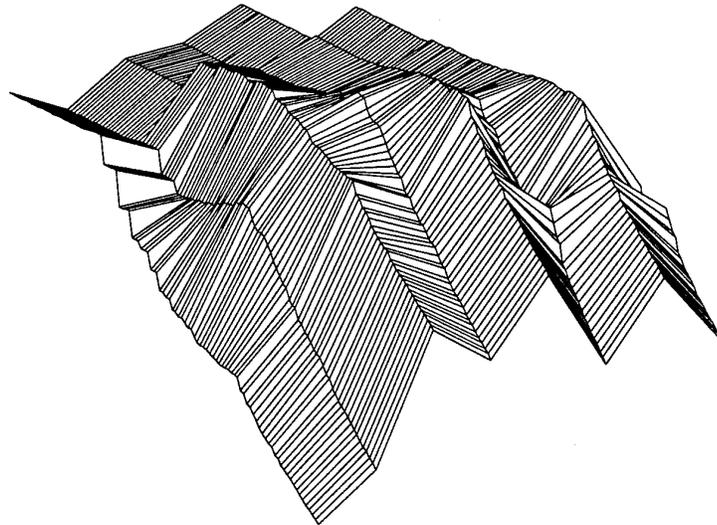


Figure 5.15: Wavenumber Distribution - Soft Model (Real-time)

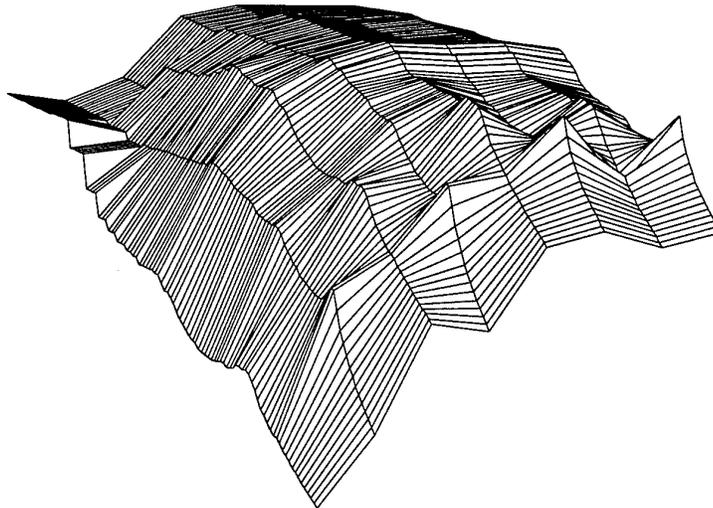


Figure 5.16: Wavenumber Distribution - Medium Model (Real-time)

Thus it is clear that the detailed operation of the low-frequency control algorithm, acting on a simulation of the raft structure, reproduces precisely the features which are required. Specifically, the raft accurately tracks hull-seaway disturbances, while low-frequency mount and structural resonances are very substantially suppressed. The associated control force distributions are biased towards the higher wavenumbers, with the result that they have very poor coupling efficiency with the hull. As a result, although the motions of the raft are accurately constrained, the isolation of the raft from the hull is very substantially improved compared to a passive installation.

5.4 Simulation of High Frequency Algorithm

Following the satisfactory implementation of the low-frequency algorithm, the high-frequency algorithm was introduced in parallel, for each of the three raft models. Thus the following simulations relate to the complete algorithm structure, with both low-frequency feedback and high-frequency feedforward components.

The coding was written to provide the choice of observing the response of either the receiving structure accelerations, or the residual forces transmitted through each mount. It was found, however, that the computational workload associated with monitoring the acceleration response, including simulation of a flexible receiving structure and every possible interaction path, was too demanding for results to be obtained on a reasonable timescale even using a very fast desktop computer. Although initial attempts were made to run the algorithm in this form, it was found that the calculation loop-time was very long, and convergence was extremely slow as a result.

But the decision had been taken, in the practical experiment, that the initial performance would be monitored by measuring residual force transmission using flux loops, rather than baseplate-accelerometers. Moreover, the assumption was made that these sensors would not be subject to cross-interaction. So it was decided to run the algorithm simulations to correspond to this condition; with this simplification, the algorithm ran very much faster and results were much more easily achieved.

The assumption was also made that each mount was subject to a component of unwanted leakage flux variation, due to the local vibration levels associated with each magnet armature. This situation is very likely to arise in practice; once the low frequency control has been rolled-off, the flux variations in each magnet air-gap will be entirely due to inductive effects associated with the local armature vibration.

Having made these assumptions, the high frequency algorithm was then run in parallel with the low-frequency algorithm, with an appropriate constraint to avoid interaction between the algorithms as described in Section 4. The results are shown in Figures 5.17, 5.18 and 5.19

In this case, the excitation inputs generating the raft vibration were taken to be a distribution of 16 random, uncorrelated forces, so that the corresponding motions of the raft structure were of a completely random nature, with poor spatial correlation at the higher frequencies.

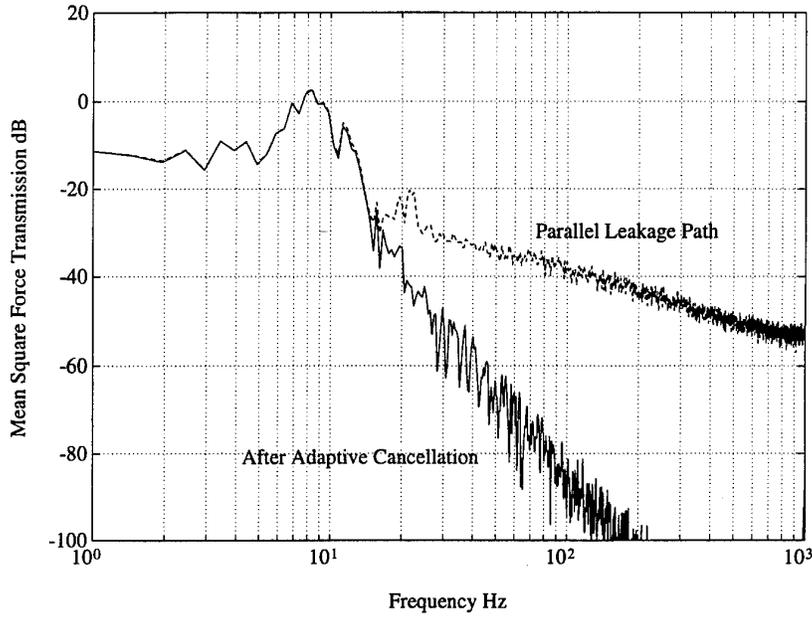


Figure 5.17: HF Algorithm - Soft Raft

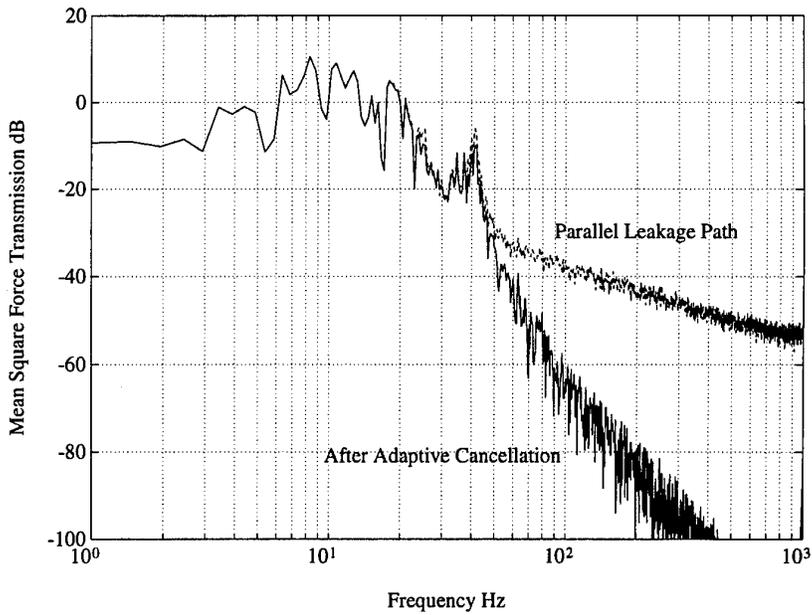


Figure 5.18: HF Algorithm - Medium Raft

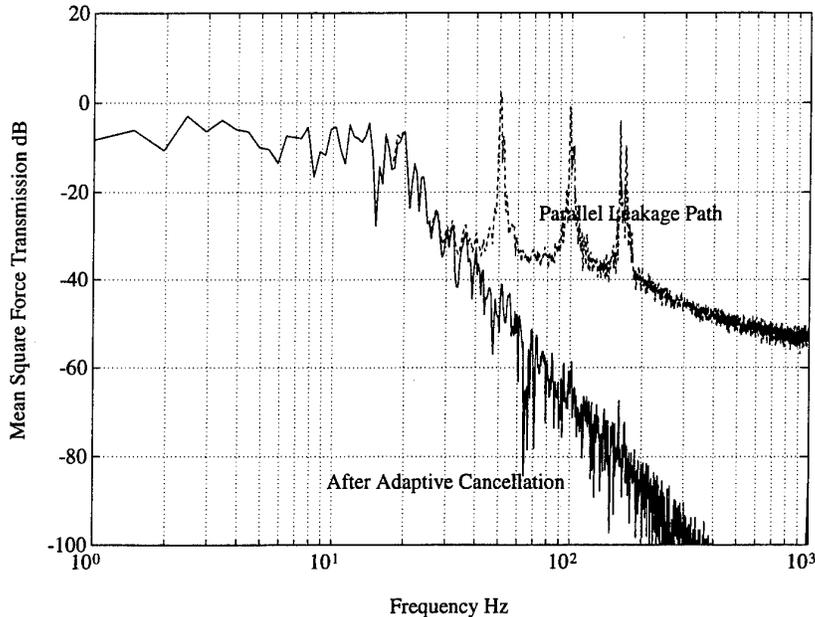


Figure 5.19: HF Algorithm - Stiff Raft

The criterion which was adopted for assessing performance was now taken to be the summed-square force transmission through all the mounts. Unlike the situation at low frequencies, where the correlation and relative phasing between different forces is a central feature defining the performance of an isolation system, at higher frequencies the forces are sufficiently separated that their transmission characteristics no longer interact when averaged over any practical bandwidth. Thus it is appropriate to regard each separate force as an independent transmission component; the appropriate measure of total transmission then becomes the total summed-square force.

Figures 5.17-5.19 show the calculated summed-square transmitted force for each of the three raft models, before and after adaptive cancellation. It can be seen immediately that in each case the adaptive process can cancel out the unwanted local components of flux transmission. The overall transmitted forces then reduce very substantially, to a lower bound which is actually defined by the residual forces associated with the low-frequency algorithm.

The adaptive algorithm cannot drive the response below this lower bound, because the low-frequency algorithm operates using input information derived from across the entire raft structure. The high-frequency adaptive algorithm does not have access to this information, and so it can only cancel those components of transmission which are associated with the local mount motions.

Nevertheless, in practice it is anticipated that it is the conditions local to each mount which will dominate individual mount transmission at the higher frequencies; the simulations show clearly that it is possible for the high frequency adaptive algorithm to cancel out these components of transmission.

An interesting feature is also observed with respect to Figure 5.19, namely the response for the stiff raft. It can be seen that the presence of leakage transmission paths can cause resonant peaks to appear and dominate the transmission process. The adaptive algorithm can knock out these resonant peaks, and cause the net transmission to reduce to that defined by the low-frequency, global controller. Since the latter specifically filters out and ignores these resonances, it follows that they are then excluded from the residual transmission process.

This example shows clearly how local mount leakage transmission can cause resonances which have been filtered out by the low-frequency algorithm to reappear. The high frequency algorithm, by cancelling these leakage paths, can ensure that these resonances no longer compromise the performance.

Thus, to summarise, it has been shown that the high frequency algorithm can successfully co-exist with the low-frequency algorithm. It serves to reinforce the actions of the low-frequency algorithm by cancelling-out unwanted transmission paths which are correlated with the local above-mount acceleration. In the absence of such adaptive cancellation, the roll-off of the low-frequency algorithm can be compromised; resonances which have been filtered-out can reappear and dominate the transmission. But the high frequency algorithm will respond to such resonances and suppress them, enabling the overall installation to achieve its full performance potential.

6 Conclusions

The dynamics of an active mounting system for large, lightweight flexible marine machinery rafts has been examined, with the objective of achieving very substantial improvements in raft isolation compared to existing passive techniques. For conventional rafts, the designer has to make trade-offs between several conflicting requirements in order to obtain a satisfactory overall compromise. Active techniques can reduce or eliminate the need for such compromise; it becomes possible to separate out these design requirements and address each aspect separately.

Specifically, the design of a machinery raft requires the installation to be stiffly mounted to resist seaway motions, yet flexibly mounted to provide good vibration isolation. As the size of the raft increases, so it can provide a useful massive bedplate for a large number of machinery items, but its isolation performance becomes compromised by the existence of low-frequency internal structural resonances. This in turn requires very stiff and massive construction if the raft is to provide a satisfactory dynamic platform for machinery mounting.

In contrast, the active mounting techniques which have been examined in this report provide the opportunity to use lightweight flexible raft structures, while still achieving levels of isolation substantially in excess of conventional passively mounted rafts. In this respect, a number of design features have emerged.

Specifically, it becomes possible to define separate response characteristics for seaway and vibration disturbances, so that the mountings of the structure can appear stiff in response to seaway motions, while being extremely soft in response to vibration excitation.

In addition, it is possible to design a distributed array of active mounts which will selectively stiffen and damp low frequency resonant modes of the raft structure, while permitting higher frequency modes to run free. Moreover, the force distributions which are used to control these low frequency modes can be optimized, so that their low-wavenumber spectral content is minimised; as a consequence they represent a much reduced vibration and sound radiation hazard.

This aspect of the active control system involves the use of global, feedback control techniques which filter out the large-scale modal characteristics of the structure, and introduce appropriate dynamic responses for each controlled mode.

As the frequency increases, so the coherent vibration of the structure becomes more localised, and vibration transmission no longer involves large-scale modal behaviour. The requirement becomes to focus down onto individual mounts, and define a low-gain, feedforward control system which seeks to minimise vibration transmission through each individual mount. This control system is adaptive, and optimizes its performance in response to observation of the forces and vibration transmitted to the receiving structure.

This adaptive high frequency controller must operate side-by-side with the optimized low-frequency feedback controller, so an important requirement is that the adaption should operate without interfering with or compromising the operation of the low-frequency controller.

An overall control structure has been defined which simultaneously achieves all the above objectives. Specifically, at low frequencies the raft is supported by a state-space feedback controller. This filters out specific modes of the structure and seeks to maximise the isolation performance, while minimising the motion of the structure with regard to both seaway disturbances and low frequency internal raft resonances.

The formulation of this aspect of the control system involves the use of optimal feedback control techniques; a control structure has been derived which enables the designer's requirements to be summarised in a specific set of space-time filters applied to both the observations and the controls. So the entire design requirement condenses down to the task of appropriately specifying these filters according to the criteria required by the designer. A straightforward optimisation procedure then enables the appropriate control gains to be fully specified.

The high frequency controller operates as an outer feedforward loop superposed on the low frequency feedback control. This outer loop provides local pointwise control, and seeks to cancel adaptively the residual components of vibration transmission through the mounts. Although the control itself is of a pointwise form, the adaptive update takes into account cross-coupling of response in the supporting hull structure.

A parallel simulation of the expected effects of the low-frequency controller enables the constraint to be applied, namely that the high frequency adaption should not compromise the existing performance of the low frequency control system.

The overall control structure provides the full bandwidth of operation, namely from low-frequency seaway motion, through mount and large-scale structural resonance suppression, to localised suppression of high frequency vibration transmission at individual mounts. This involves a corresponding transition from low-frequency global feedback control to high frequency pointwise feedforward control.

The specification of all aspects of this control system has been rigourously defined to take into account the constraints imposed by the designer, so that the system will automatically take into account these various design requirements.

The overall active control capability offers a very substantial improvement over conventional passive design.

7 References

- [1] Fitzpatrick, G. W., "Zebedee Control on 4M1 Raft", AST/ERC/7.0/387/EN, November 1991. (Project document)
- [2] Goyder, H. G. D., "Methods and application of structural modelling from measured structural frequency response data", Journal of Sound and Vibration, Vol 68, 1980, pp209-230.
- [3] Kwakernaak, H. And R. Sivan, "Linear Optimal Control Systems", John Wiley & Sons, 1972.

Distribution

1 - 2	Director, Advanced Research Projects Agency	ARPA, 3701 N Fairfax Dr. Arlington, Va. 22203-1714 DODAAD Code: HR0011
3 - 5	Scientific Officer	ONR, Code 452, 800 North Quincy St, Arlington Va. 22217-5000
6	Administrative Contracting Officer	ONR, 495 Summer St., Room 103, Boston, Ma. 02210-2109
7	Director, Naval Research Laboratory	NRL, Code 2627, Washington, DC. 20375 DODAAD Code: N00173
8 - 9	Defense Technical Information Centre	DTIC, Building 5, Cameron Station, Alexandria, Va. 22304-6145

REPORT DOCUMENTATION PAGE	1. REPORT NO. ARPA/S30M2/93/01	2.	3. Recipient's Accession No.
4. Title and Subtitle Advanced Submarine Technology - Project M Control Theory Report		5. Report Date September 1993	
7. Author(s) M A Swinbanks, S. Daley		8. Performing Organisation Rept. No. MRC/AST/FTR/11.0	
9. Performing Organization Name and Address GEC- Marconi Research Centre GEC-Marconi Limited Great Baddow, Chelmsford Essex, U.K., CM2 8HN		10. Project/Task/Work Unit No. S30M2 Phase 1	
12. Sponsoring Organization Name and Address ARPA 3701 N Fairfax Dr., Arlington Va. 22203-1714 ONR 800 N Quincy St, Arlington Va. 22217-5000		11. Contract(C) or Grant(G) No. (C) N00014-93-C-0063 (G)	
		13. Type of Report & Period Covered Interim Jan/93 - Jun/93	
15. Supplementary Notes		14.	
16. Abstract (Limit: 200 words) This interim report presents the broad band control theory aspects of a programme whose objective is to control the dynamics of a large, lightweight marine machinery raft supported on an array of electromagnetic actuators. The report shows how it is possible to achieve, simultaneously, the normally conflicting requirements of accurate alignment of the raft under sea-way or manoeuvring motions, of the elimination of mount resonances, of stiffening long wavelength bending resonances without incurring a radiation hazard, of decoupling shorter wavelength bending resonances and of greatly reducing any residual force transmission through the mounts at higher frequencies. Simulations of the control theory applied to a simple model raft, show that very large performance improvements are possible. These findings go some way to giving the assurance that large, lightweight machinery rafts could be used in future installations without having to make any compromises in overall performance. This, in turn points to a new way for reducing overall costs.			
17. Document Analysis a. Descriptors Active control, Active suspension, Active vibration isolation, Actuators, Digital signal processors, Flexible support, Machinery isolation, Magnetic levitation, Magnetic suspension, Submarines, Vibration control. b. Identifiers/Open-Ended Terms c. COSATI Field/Group			
18. Availability Statement Unlimited		19. Security Class (This Report) Unclassified	21. No. Of Pages 64
		20. Security Class (This Page) Unclassified	22. Price