THE CONTROL OF DYNAMICALLY INTERACTING SYSTEMS

by

JAMES EDWARD COLGATE

S.B. in Physics, Massachusetts Institute of Technology
(1983)

S.M. in Mechanical Engineering, Massachusetts Institute of Technology
(1986)

SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN MECHANICAL ENGINEERING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
August 1988

Signature of Author

Department of Mechanical Engineering
August 3, 1988

Certified by

Neville Hogan
Thesis Supervisor

Accepted by

Professor Ain A. Sonin
Chairman, Departmental Committee on Graduate Studies

THE CONTROL OF DYNAMICALLY INTERACTING SYSTEMS

by

JAMES EDWARD COLGATE

Submitted to the Department of Mechanical Engineering
on August 3, 1988 in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy in
Mechanical Engineering

ABSTRACT

This thesis addresses the analysis and design of controllers for dynamically interacting systems. The example of a robotic manipulator which must interact with an arbitrary passive environment is considered in detail.

An approach to the design of *interaction controllers* is presented in terms of a heirarchical set of design specifications and is contrasted with an approach to servo design. It is shown that the emphasis on interaction creates a need for *coupled stability* and *interactive behavior* specifications.

Methods are developed for the analysis of coupled stability and interactive behavior. A necessary and sufficient condition for the stability of a linear, time-invariant plant coupled to an arbitrary passive environment is derived. An alternative test for coupled stability requiring the construction of two root loci, one representing interaction with springs, and one representing interaction with masses, is also developed.

Experiments performed to examine the utility of these methods are described. A variety of controllers were implemented on a two-link manipulator and a device for measuring the endpoint impedance of the manipulator was constructed. Impedance measurements of the closed-loop systems were made, and interaction with springs and masses was examined. These experiments indicate that a measurement of the impedance is an effective predictor of interactive behavior. In addition, these experiments demonstrate that systems with identical servo responses can exhibit significantly different interactive behaviors, and that contact instability can occur in the absence of force feedback.

The concepts of a *passive physical equivalent* and an *uncontrollable element* are introduced. These concepts are used to analyze the contact instability phenomenon associated with force feedback, and to make recommendations for improved force control.

Finally, approaches to the design of interaction controllers are presented and analyzed. The relative merits of various design methods, including "servo masking" and the "target model referenced controller", are discussed.

Thesis Supervisor: Neville Hogan
Title: Associate Professor of Mechanical Engineering
Department of Mechanical Engineering

MEMBERS OF THE COMMITTEE

Dr. Neville Hogan (Chairman)
Associate Professor
Department of Mechanical Engineering

Dr. J. Karl Hedrick
Professor
Department of Mechanical Engineering

Dr. Warren P. Seering
Associate Professor
Department of Mechanical Engineering

Dr. John L. Wyatt
Associate Professor
Department of Electrical Engineering and Computer Science

Acknowledgements

I would like to thank Neville Hogan for introducing me to this field of research, and for doing well those things an advisor should do—providing essential guidance and support. However, I am grateful to Neville for much more—for allowing me the latitude to pursue my own, sometimes off-beat notions; for sparking lively discussion one-on-one, in group meetings, or over beer; and most of all, for managing to transcend the more petty concerns (like thesis-writing) that threaten to control the life of an academic, thereby providing frequent reminder of what was so attractive about this field in the first place.

I am also grateful to the other members of my committee, Karl Hedrick, Warren Seering, and John Wyatt. As individuals, they offered encouragement and sage consultation. As a group, they trounced my pre-conceived notions of the ineffectiveness of committees.

Among my fellow students, Bill Murray stands out. In five years as an officemate, Bill was always a source of help, humor, and discussion. He also set a standard of quality in all aspects of research from which I and the entire Newman Lab have benefitted.

I would like to thank Dov Adelstein for his irreverence, which helped me keep things in perspective, and for the bull sessions with Bill and others that were usually just too entertaining. Dov also wrote and gave me access to the data analysis software used in this thesis.

Thanks to Ian, wherever he is, for his work on the two-link manipulator; to Andre, for the rousing discussions of force control; to Crispin and Ernie, for affording me a glimpse of thought on other, higher planes; to Cary, Joey B., Dave (the story-teller), and John W., who have found life beyond MIT; to Eric and Dawei, for mezzanine banter; to Keita, for nursing the microVAX and keeping me in business with Matrix-x, and to J.B., for nursing the laser printer and seeing to the purchase of Matrix-x. And thanks to the rest of the members of the Newman Lab for collectively making it a fine place to be.

I am grateful to Emilio Bizzi for allowing me to use the equipment in his lab, and to Sandro Mussa-Ivaldi, for helping me first to learn, and then to maintain the two-link manipulator.

Thanks to Bob Samuel for the good times in the machine shop, and to Deborah Faust for helping me out of many a tight spot.

Thanks to Cliff and Greg, roommates of long-standing, and all-around pals who helped me realize that there is more to life than the lab. Thanks to Marissa, who may yet convince me that there is more to life than engineering.

Finally, it would be fruitless to try to find a sentence or two that would properly acknowledge the love and support of my parents (fortunately, we're family, so they won't mind if I don't).

This research was performed in the Eric P. and Evelyn E. Newman Laboratory for Biomechanics and Human Rehabilitation at M.I.T. Partial support was supplied by a National Science Foundation Fellowship, NSF grants number ECS-8307461 and EET-8613104, NIDRR grant number G000830074, and the Whitaker Foundation.

To Mom and Dad

Contents

List of Figures

List of Tables

Nomenclature

General Notation

a	scalar
$\mathbf{a}$	vector
A	matrix or scalar, as indicated by context
A_c	matrix corresponding to closed loop system
A_s	matrix corresponding to servo system
A_t	matrix corresponding to target system
$\mathbf{a}_0$	reference value

Acronyms

LHP	left half s-plane
LTI	linear time-invariant
PD	positive definite
PSD	positive semi-definite
RHP	right half s-plane
TMRC	target model referenced controller

Mathematical Symbols

a^*	complex conjugate of a
$\mathbf{a}'$, A'	transpose of $\mathbf{a}$ or A
$\mathbf{a}^H$, A^H	complex conjugate transpose (Hermitian) of $\mathbf{a}$ or A.
$\|a\|$, $\|\mathbf{a}\|$	Euclidean norm ($\sqrt{a^*a}$, $\sqrt{\mathbf{a}^H\mathbf{a}}$)
$\langle \mathbf{a}, \mathbf{b} \rangle$	inner product
$\sigma(A)$	singular value of A
$Re\{a\}$, $\Re a$	real part of a
$Im\{a\}$, $\Im a$	imaginary part of a
j	$\sqrt{-1}$

Bond Graph Symbols

0	common effort junction
1	common flow junction
S_e	effort source
S_f	flow source
I	admittance causality storage (e.g., inertia, inductor)
C	impedance causality storage (e.g., spring, capacitor)
R	dissipator (e.g., damper, resistor)
TF	transformer
GY	gyrator

Miscellaneous

A	state matrix
b, B	viscous damping coefficient
B	input matrix
C	output matrix
$C(s)$	closed loop transfer function
D	feedforward matrix, determinant term
$\mathbf{e}$	effort, error
E	energy
$\mathbf{f}$	flow
F	force
G	force feedback gain
$G(s)$	plant transfer function
$G_{xx}(\omega)$, $G_{xy}(\omega)$, $G_{xF_x}(\omega)$, $G_{xF_y}(\omega)$	auto- and cross-spectral densities
$H_x(\omega)$, $H_y(\omega)$	transfer functions from x position to x and y forces
I	inertia matrix
J	jacobian matrix
k, K	stiffness
L	environmental input matrix
m, M	mass
P	power, PD matrix in energy function
$\mathbf{r}$	reference
s	Laplace variable ($\sigma + j\omega$)
S	scattering matrix
$S(\omega)$	transfer function from x position to y position
t	time
T	time delay

u, U	control effort
v	velocity
x, X, y	positions
$Y(s)$	admittance
$Z(s)$	impedance
α	amplification factor
β	two-link manipulator damping coefficient
γ	coherence
ε	error
ω	imaginary part of s, joint angular velocity
σ	real part of s
τ	time constant
θ	joint angle
Θ	state vector of two-link manipulator

Chapter 1

Introduction

1.1 Statement of Purpose

This thesis addresses the control of dynamically interacting systems. Interest in this problem arises in the study of manipulation. One of the major difficulties currently facing those in the field of robotics is controlling the dynamic interaction of a robot with its workpiece. This problem has engendered several approaches to robot control, such as "force control" [89] and "impedance control" [36], which have met with varying degrees of success. However, despite these approaches, basic understanding of the interaction of controlled systems—such as manipulators—with dynamic environments remains weak.

The primary goal of this thesis is to develop tools for analysis of the interactive behavior of controlled systems. These tools are intended to be sufficiently general to apply to any linear control system. Their utility will be experimentally examined.

The second goal is to better, through application of these tools, the understanding of those approaches to robot control which address interaction. Several existing approaches, as well as several proposed here, will be examined.

The strategy for realizing these two goals involves investigation of control systems' *driving point impedances* (see Section 1.5). A final goal is to demonstrate that the focus

on impedance rather than transfer functions, which lie at the heart of servo control theory, is the correct approach to the interaction problem. The distinction between interaction and servo problems will be illustrated with examples located throughout the text.

1.2 Background

The approach taken to the realization of the goals above is the application of an eclectic set of concepts and theoretical tools. The purpose of this section is to provide a brief introduction to the influential background topics, with emphasis on their relevance to the thesis goals.

1.2.1 Physical Systems Modeling

Physical systems modeling techniques, such as bond graphing [65], provide a means of analyzing the behavior of complex dynamic systems based upon an understanding of one, the dynamic behavior of the system components, and two, the energetic interaction (connection) of these components. This thesis is a systems approach to the study of manipulation (see Section 1.3.1), where the system of interest is composed of the manipulator coupled to some dynamic environment. The study of this system can be addressed with the techniques of physical systems modeling; however, there is a complication, which is that one of the components is feedback-controlled.

Because feedback is the use of measures of system states to modulate a source, one reasonable approach is to develop a model of the complete system (manipulator plus environment), and subsequently analyze the effect of adding feedback. This, however, is not the approach taken here. Primarily because a manipulator may be expected to interact with a diverse set of environments (e.g., tools), it would be convenient to treat

it as a component of the overall system[1].

Such an approach is not without its perils. Much of the power of bond graphing, for instance, lies in its parsimony. Only nine symbols (0, 1, S_e, S_f, GY, TF, I, C, R) are typically used in bond graph models, yet these symbols are sufficient to describe the behavior of nearly all physical systems. The choice not to represent explicitly the presence of feedback will result in system components which cannot necessarily be described by these symbols alone. Thus, the increased complexity of such an approach is a concern.

Nevertheless, this route has been taken because it is felt that both a *modular* treatment of a controlled manipulator, and access to the insights and theoretical results (e.g., Tellegen's theorem) of the study of physical systems, would prove quite powerful.

1.2.2 Physical Equivalence

A related concept which is also of great influence is Hogan's postulate of "physical equivalence" [36]:

> "Any controlled system will consist of 'hardware' components (e.g., sensors, actuators, and structures) combined with controlling 'software' (e.g., neural networks, brains, or computers). A unified approach to both the design of the controller and the physical hardware can be developed by postulating that, taken together, the hardware and software is still a physical system in the same sense that the hardware alone is."

The potential value of this postulate is twofold: first, in the design of controllers for physical systems, one need not consider any controllers which will result in behavior

[1]The use of feedback controlled components in the design of complex systems is actually quite common—consider the ubiquitous op-amp. Devices such as the op-amp, however, make use of high input impedance and low output impedance connections to limit energetic interaction (and don't always succeed); mechanical devices generally exhibit more significant energy exchange.

that cannot be described as that of an equivalent physical system; second, it should be possible to use physical systems modeling techniques to describe controlled systems. Thus, this postulate suggests that the approach described in the previous section is not altogether unreasonable.

The coupled stability analysis in Chapter 3, and, to a greater extent, the concept of a "passive physical equivalent" introduced in Chapter 6, attempt to exploit this concept. Throughout this thesis, feedback—whether it is introduced for servo control or for interaction control—is treated as an operation which has the effect of changing the physical behavior of the plant.

1.2.3 Control Theory and Network Analysis

The mathematical techniques of control theory and (electrical) network analysis and synthesis are used throughout this thesis. The proofs in Chapter 3 are based on Nyquist stability theory and root locus. The mathematical statements of passivity presented in Chapter 2 may all be found in the network analysis literature, and network synthesis concepts are the basis of the passive physical equivalent analysis in Chapter 6.

Although network analysis may be considered a subset of physical systems theory, it is focused to a much greater extent on the analysis of models rather than the modeling of systems. Because of this focus, many powerful mathematical results are available, some of which play a large role in this thesis.

1.3 Context

1.3.1 Manipulation

The context of this thesis is the study and design of manipulators. The study of manipulation, both human and machine, provides a strong incentive for understanding the control of interacting systems. As Hogan has often reasoned [37]:

> "By any reasonable definition, manipulation of an object implies mechanical interaction with that object."

The human arm, for instance, is a manipulator with capabilities that include well-behaved interaction with an extremely diverse set of environments. Its repertoire includes the abilities to employ eating utensils, to perform complex dynamic tasks such as catching or throwing a ball, and to perform kinematically constrained tasks such as sanding a surface or opening a drawer.

This tiny selection of tasks in fact encompasses a far broader range than that achievable by a mechanical manipulator controlled with state of the art servo (i.e., command following) techniques. Most successful industrial applications of robots have been in non-contact applications such as welding or spray-painting, or in simple pick-and-place operations. Yet, the bulk of industrial operations to which a robot might be applied require significant work transfer. These operations include parts assembly and various material removal processes. In the future, cooperation with other robots may also be included. In such a case, the conventional position servo strategy of making the robot much stiffer than the environment (to 'reject' environmental disturbances) is bound to encounter difficulties—the environment may be just another, perhaps identical, robot.

In the next section, an approach to manipulator control which explicitly considers interaction with the environment is described.

1.3.2 Impedance Control

The concept of active control of a manipulator's interactive behavior is formally treated as an aspect of "impedance control" [34,35,36]. Impedance control is an approach to manipulation which is based on the assertion that it is not sufficient to control some vector of port variables such as positions, velocities, or forces, but that it is also necessary to control the dynamic relations among the port variables, such as impedances and admittances. In particular, impedance control makes the reasonable assumption that most "environments" a manipulator will interact with are admittances (i.e., kinematic constraints or mass-like objects), and that to be causally consistent, the port behavior of the manipulator should be that of an impedance. To illustrate, a single-axis manipulator connected to an environment with admittance causality may be represented by the bond graph in Figure 1.1. The job of the impedance controller is to modulate both v_0, the "virtual trajectory" of the manipulator [36], and Z, the manipulator impedance. The reader should note that the use of a bond graph to represent an impedance controlled manipulator is an application of the postulate of physical equivalence.

Hogan has shown that a large class of nonlinear, multi-axis manipulators may be represented by the structure in Figure 1.1. He has also provided a few illustrations of impedance controller designs (see Section 4.1). The benefits of these designs are well documented [39,33]. The design methods offered do not, however, apply to a broad range of manipulators. Indeed, impedance control is intended to provide an approach to understanding manipulation rather than a prescription for designing controllers.

The design and implementation of impedance controllers is, however, an active area of research. Some of the current topics of interest are planning impedances [5,26,27], designing impedance controllers in the face of modeling errors and/or complex models [22,30,44,45,55], and applications to force control [40,92]. The latter two topics will be briefly introduced in the following two sections, and then reviewed in more depth in

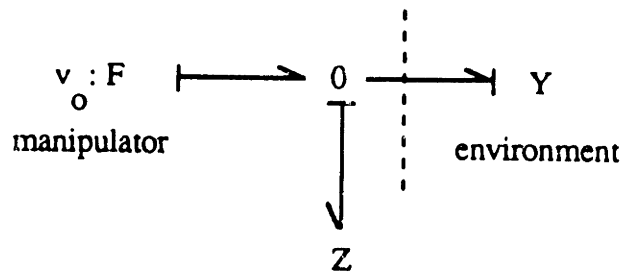

Figure 1.1: Bond graph representation of an impedance controlled manipulator.

Chapters 4, 7, and 8, as they are important background for this thesis.

The work presented in this thesis may be classed as impedance control as the focus is on interaction with a dynamic environment. This work, however, departs philosophically from impedance control in two ways: manipulators are modeled as admittances, not impedances, and the basis of the theoretical developments is a concern for the stability rather than the causal consistency of a manipulator which interacts with its environment.

The reason for these departures is the difficulty associated with modeling the coupling of dynamic systems, particularly when both systems would, if modeled independently, be assigned the same causality (see, for instance, [37]). Most manipulators and many environments are probably best modeled as admittances, thus the difficulty. Two readily apparent approaches to this problem are one, to model the manipulator as an impedance, and attempt to design controllers which make it behave as an impedance, as there is nothing that can be done about the environment; and two, to avoid the issue by considering the stability of the coupled system which, at least in the linear case, does not require consideration of causality. Other methods, such as endowing the junction between the two systems with appropriate dynamic behavior, may also be used, but tend to lead to stiff system models, which are difficult to analyze. The second approach lends itself most neatly to the application of network analysis, and is therefore used

here. Because this thesis does not focus on the manipulator impedance—as opposed to the admittance—the controller designs are frequently referred to as "interaction controllers".

1.3.3 Departures from the Ideal

Despite the desirable properties that an impedance controller would have if it could be designed, it is not always easy or even possible to implement a particular impedance. A major part of the reason for this is that real hardware exhibits various non-ideal behaviors, such as actuator, transmission, and sensor dynamics. The implementations suggested by Hogan have all been for robots with rigid links and without any of these additional dynamic effects.

Recently, efforts have been made to examine the effects of various non-idealities on the performance of impedance controllers. Kazerooni [44] has developed an approach to the design of impedance controllers which treats a linearized manipulator model with first-order actuator dynamics; Spong [74] has developed an impedance controller which is intended to perform well in the face of joint flexibilities; Fasse [22] has examined the robustness of impedance controllers to various modeling errors. More recently, Kazerooni [45] has developed an approach to the design of impedance controllers which will work with almost completely unstructured models.

This thesis takes a step back from these sundry design efforts in an attempt to establish quantitative design criteria which any of these implementations must meet if it is to be successful. Thus, Chapters 3–7 focus on the development and application of analytical techniques; these techniques are restricted to linear systems, but not to particular manipulator models or controller designs. Chapter 8 returns to consider the design of interaction controllers in light of what is learned from the analysis.

1.3.4 Force Control

Force control of robotic manipulators is analyzed in detail (see Chapter 7) as it is a research topic that is currently receiving tremendous attention in the literature. Force control is a type of interaction control as it presumes that the manipulator is attempting to exert a controlled force on some object (generally, a stiff surface as in, for instance, a deburring application).

Force control is also interesting in that it demonstrates how demanding interactive tasks can be. As described in Chapter 7, many implementations of force control (using force feedback) result in a violent chattering behavior upon contact with the workpiece; implementations which do not exhibit this behavior generally suffer from low bandwidth.

An analysis of force control, which treats the problem as one of interaction control, is presented. New insights into the difficulty with force control are generated, and novel suggestions for improved control are offered.

1.4 An Approach to the Control of Interaction

In this section an approach to the design of interaction controllers is presented. Design specifications are described and contrasted with those of a servo, and the structure of the closed loop interaction controller is presented in a fashion suited to the application of network analysis.

1.4.1 A Heirarchy of Design Specifications

To begin, consider a servo controller, which is meant to be synonymous with a "command following" or "tracking" controller. Specifications for a servo controller may be

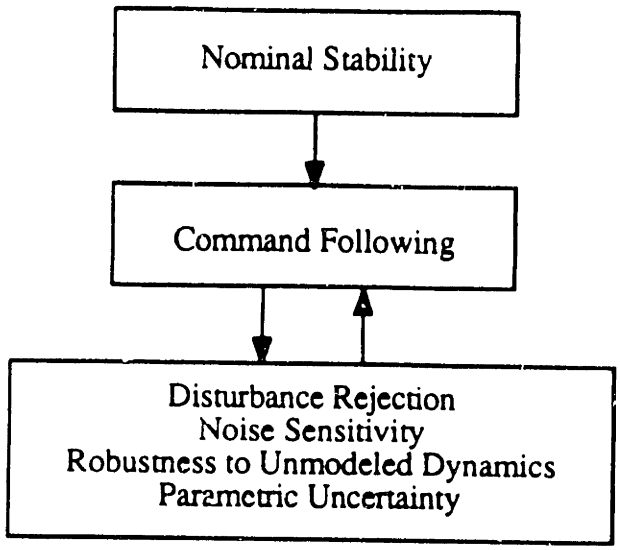

Figure 1.2: Design specifications for a servo controller.

stated in virtually any manner the designer chooses, but generally include some subset of those shown in Figure 1.2. The heirarchy of specifications is also generally as shown. Nominal stability is, of course, an absolute requirement, while tradeoffs are required for the other specifications. Classical control focuses on nominal stability, command following, and disturbance rejection. Many of the techniques of modern control provide guarantees of nominal stability, while command following and disturbance rejection are weighed against various measures of robustness. Much of the current effort in control theory is focused on robustness and sensitivity issues.

The design specifications for an interaction controller are somewhat different. The heirarchical structure is shown in Figure 1.3. Stability of the isolated system remains paramount; however, as the manipulator is expected to interact extensively with its environment, coupled stability is nearly as important. Coupled stability may, however, be traded against performance; a manipulator need not be capable of interacting stably with every conceivable environment, just those which it is likely to encounter. Interactive behavior has been added to the performance specifications as stability is clearly an

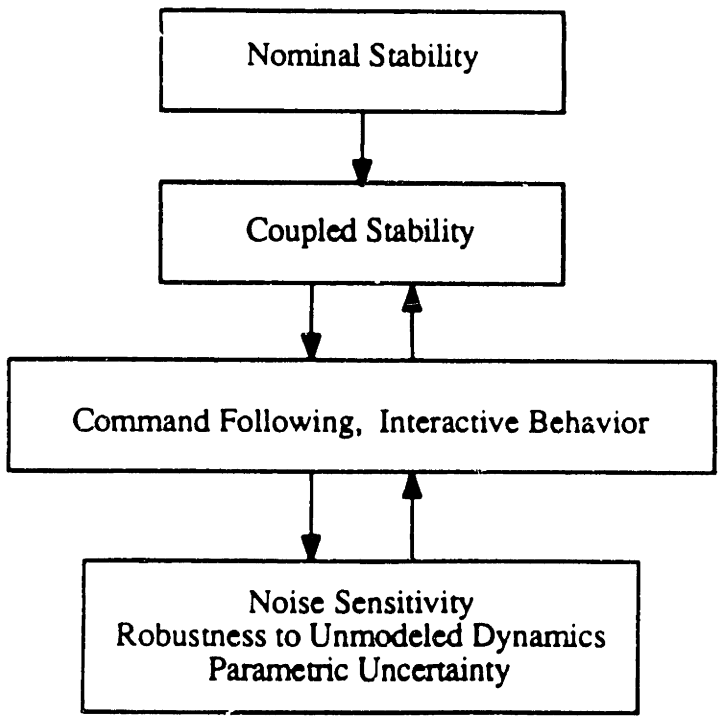

Figure 1.3: Design specifications for an interaction controller.

insufficient requirement; a reasonable behavior is also needed. The other requirements are essentially as before.

Although methods for the design of interaction controllers can be expected to develop more rapidly than methods for the design of servo controllers have evolved, they are currently in an infant state. Recently, Kazerooni has begun to apply various techniques of modern control to the design of impedance controllers, but many of his results, which will be reviewed herein, suffer poor behavior at the level of coupled stability or interactive behavior. In Chapters 8 and 9, the relation of interaction control to recent activity in control theory is considered further. The bulk of the analysis, however, will have a classical control flavor, and will concentrate primarily on coupled stability, and secondarily interactive behavior. The intent is to develop analytical tools which provide the same sort of insight to the design of interaction controllers that stand-by

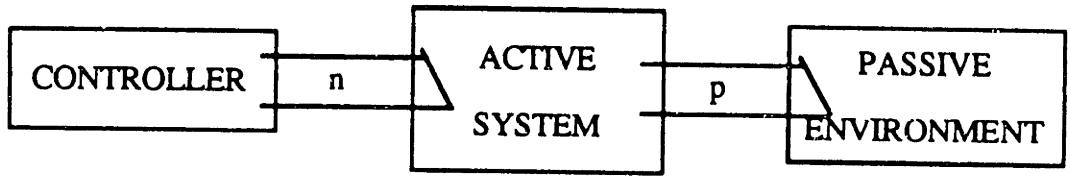

Figure 1.4: Structure of manipulator/environment system.

techniques such as root locus and Bode plots do for servo controllers.

The sets of design specifications should indicate that, at least in conception, interaction control differs from servo control. Several examples located throughout the thesis serve to bolster the argument that this is more than a matter of philosophy, that different designs and behavior result.

1.4.2 Structure of Interaction Controllers

The structure of the complete system to be studied in this thesis is shown in Figure 1.4. It consists of an active system (the manipulator) with m states, n control ports, and p interaction ports (generally, $m \geq n \geq p$), connected at the appropriate ports to a controller and a passive environment (see Section 3.1 for an explanation of the restriction to passive environments).

It is assumed that the control can be divided into a state-dependent part and a state-independent part, that is, that the control vector $\mathbf{u}$ may be expressed as $\mathbf{u} = \mathbf{h} + \mathbf{r}$, where $\mathbf{h}$ is the output of a compensator whose inputs are the states of the active system, and $\mathbf{r}$ represents the state-dependent control, which itself may be the output of a dynamic system. The complete system may now be reconfigured as shown in Figure 1.5.

As an aside, it should be noted that many of the interaction controller designs that will be presented assume that a measure of the inputs to the manipulator from the

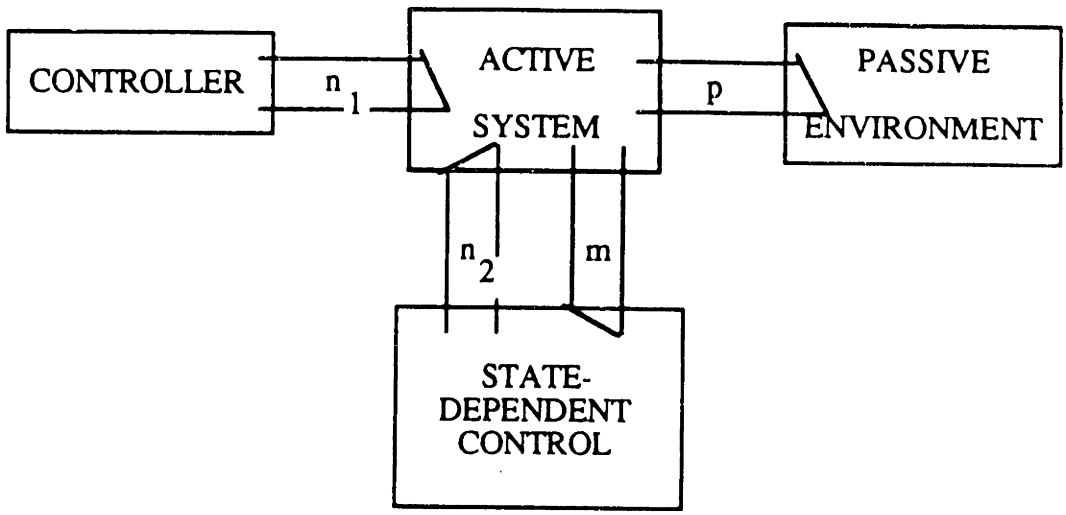

Figure 1.5: Structure of the manipulator/environment system; control divided into state-dependent and state-independent parts.

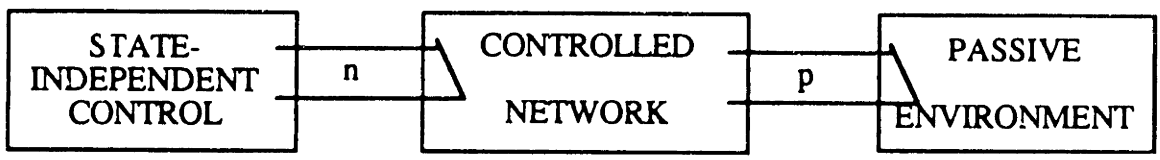

Figure 1.6: Network representation of manipulator/environment system.

environment is available. In point of fact, such a measure is a measure of sensor state, and the sensor is part of the active system, so that the structure in Figure 1.5 applies.

Finally, a change of viewpoint will be made. The combination of the active system and state-dependent control will be viewed as a controlled network, as shown in Figure 1.6. The value of this viewpoint is that it indicates that the wealth of literature and experience in the area of network analysis and synthesis may be applied to the problem of controlling interaction. This will be quite important in Chapters 3, 6, and 7.

1.5 Terminology

The terminology in this thesis reflects a rather mixed heritage. The disciplines of classical control, physical systems modeling, and network analysis all provide important background, but also contribute distinct dialects to common subject areas. This section, therefore, is intended to be as much thesaurus as glossary.

All of the examples in this thesis treat the behavior of physical systems. The physical system of interest, in the abscence of control, is referred to as a **system**, **network**, or **plant**. It is generally the case that there is control, possibly via feedback, over the behavior of the system.

A system communicates with its **environment** at **ports**. The environment is simply another physical system, but generally not one over which there is control. A port is a physical location where two systems exchange energy. It is characterized by two variables, such as force and velocity, whose product is the power flow *into* the system through that port. For an electrical system, any two **terminals** constitute a port.

A **1-port** is a system with one interaction port. A **two-terminal network** is an example of a 1-port. An **n-port** is a system with n interaction ports.

The two power variables associated with any port are generally termed **effort** and **flow** variables. Forces, torques, and voltages, for instance, are types of efforts. Translational and angular velocities and currents are types of flows. The term **force** will generally refer to a force-torque vector. The term **velocity** will generally refer to a translational velocity-angular velocity vector.

The dynamic relationship between power variables at a port is, in general, called an **impedance**. More specifically, the impedance maps a flow input to an effort output, while the **admittance** maps an effort input to a flow output. It should be clear from context whether "impedance" is being used in a general or specific sense. For n-ports,

impedances and admittances are generalized to map vectors of inputs to vectors of outputs. The terms **driving point impedance** and **driving point admittance** are equivalent to "impedance" and "admittance".

Although the terms "impedance" and "admittance" are commonly restricted to linear systems, they may be applied to nonlinear systems as well. They are causal dynamic operators which map an input time function $\mathbf{u}(t)$ onto an output time function $\mathbf{y}(t)$ such that the present value of the output $\mathbf{y}(t)$ may depend on the entire past history of the input $\mathbf{u}(t-\tau)$ for $0 < \tau < \infty$, and such that $\mathbf{u}'(t)\mathbf{y}(t)$ is the instantaneous power flow into the system [65].

A **transfer function** is distinct from an impedance in that it relates arbitrary inputs and outputs found at arbitrary locations. Transfer functions generally relate controlled outputs to reference inputs. A **servo** controller attempts to make such a transfer function equal to unity within a specified bandwidth. A good servo controller also attempts to 'reject' disturbances to the closed loop system.

The term **interactive behavior** is used loosely throughout the thesis until Chapter 8, where it is defined in terms of a frequency domain specification. In general, however, interactive behavior simply refers to the system's response to environmental disturbances.

The term "physical equivalence" was defined in Section 1.2.2. A physical system which has the same behavior as a controlled system, at least as viewed at some port, is termed a **realization**. The passive physical equivalents described in Chapter 6 are realizations.

1.6 Summary of Remaining Chapters

Chapter 2 serves as a primer in *passivity*. The intuitive notion and mathematical statements of passivity play an important role throughout the thesis. In Chapter 3, a necessary and sufficient condition for the stability of a linear control system coupled to an arbitrary passive environment is derived. This analysis lays the theoretical groundwork for the remainder of the thesis. Graphical tests for "coupled stability", based on the Nyquist and root locus constructions, are presented. Several examples of the application of the coupled stability criterion are described in Chapter 4. Some of the examples apply to a two-link manipulator. In Chapter 5, experiments performed with a two-link manipulator are described. The experiments demonstrate the utility of the analytical tools developed in Chapter 3, and corroborate the analyses presented in Chapter 4.

In Chapter 6, "passive physical equivalents" and "uncontrollable elements" are introduced. These are tools for the analysis of a control system's interactive behavior based upon the generation of an equivalent structured model. In Chapter 7, these tools are used to analyze the behavior of various force feedback controlled robot models. Suggestions for improved force control are also made. In Chapter 8, attention is focused on methods of implementing desired interactive behavior. Three methods are described, and examples are analyzed. Some conclusions are drawn with regard to the effectiveness of these methods, and the importance of the inherent dynamics of the plant is discussed. In Chapter 9, the results of the thesis are summarized. The relevance of these results to areas other than manipulation is discussed, and directions for future research are suggested.

Chapter 2

Passivity

The concept of a passive physical system will play an important role in the sequel. The purposes of this chapter are to provide the reader insight into passivity, and to derive the mathematical statements of passivity that will prove useful in subsequent analyses.

2.1 Thermodynamic Concepts

Passivity is closely related to the thermodynamic concept of *availability*, due to Keenan [15,48]. Availability may be understood as follows.

Consider (following [15]) a system immersed in an atmosphere as shown in Figure 2.1. We are interested in determining the amount of useful work that can be extracted from this system as it undergoes a change in state (from state 1 to state 2). The first law of thermodynamics requires the following:

$$(Q_{1-2})_{rev} - (W_{1-2})_{rev} = E_2 - E_1 \tag{2.1}$$

where $(Q_{1-2})_{rev}$ is the heat delivered to the system by a reversible process from state 1 to state 2; $(W_{1-2})_{rev}$ is the work done by the system; and E is the energy of the system. The second law of thermodynamics may be invoked to rewrite this equation

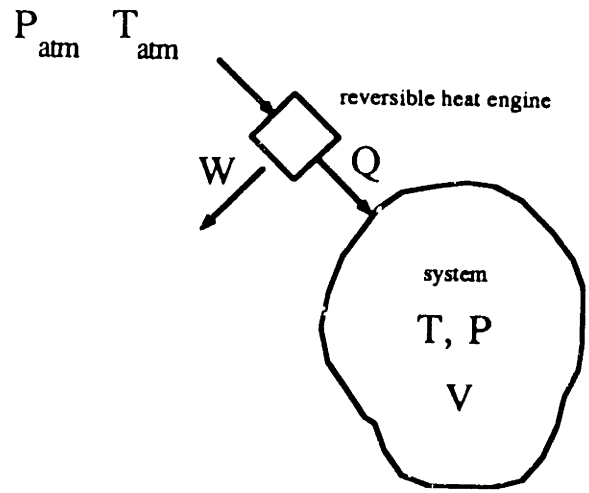

Figure 2.1: Thermodynamic system interacting with an atmosphere; the reversible engine will extract the maximum useful work from this system for any change in state. Adapted from [48].

for a general, possibly irreversible, process:

$$W_{1-2} \leq \int_1^2 T\,dS - (E_2 - E_1) \tag{2.2}$$

where T is the temperature and S is the entropy of the system. Not all of W_{1-2} will be useful work, however. Some portion of the total work transfer will be required to displace the atmosphere:

$$(W_{1-2})_{useful} \leq \int_1^2 T\,dS - \int_1^2 P_{atm}\,dV - (E_2 - E_1) \tag{2.3}$$

where V is the volume of the system. The integrals on the right side of this equation will be maximized for a reversible process which, because the atmosphere is at constant temperature and pressure, will be isothermal and isobaric. Thus:

$$(W_{1-2})_{useful} \leq -[(E + P_{atm}V - T_{atm}S)_2 - (E + P_{atm}V - T_{atm}S)_1] \tag{2.4}$$

We now define the property $\Phi = E + P_{atm}V - T_{atm}S$, and write:

$$(W_{1-2})_{useful} \leq -(\Phi_2 - \Phi_1) \tag{2.5}$$

The quantity Φ has a minimum value, corresponding to the system state $P = P_{atm}$ and $T = T_{atm}$. Thus, Keenan has defined the availability, Λ, as:

$$\Lambda = \Phi - \Phi_{min} \tag{2.6}$$

Equation 2.5 can be rewritten in terms of the availability as follows:

$$(W_{1-2})_{useful} \leq \Lambda_1 - \Lambda_2 \tag{2.7}$$

Now that we have an expression for the amount of useful work that can be extracted from a system, we are prepared to consider passivity. However, because the first and second laws of thermodynamics guarantee the existence of such a Φ_{min}, thermodynamics texts generally do not forge a link between availability and passivity. In order to do this, it is necessary to postulate the existence of systems for which Φ has no lower bound. For such systems, the availability (Λ) would be infinite. Because these systems could supply infinite useful work, they would clearly not be passive in any intuitive sense. Conversely, those systems for which Φ is bounded from below (Φ_{min} exists) could supply only finite useful work, and should be termed "passive". This is, in fact, the concept of passivity which has received the most widespread acceptance in the engineering literature [90,93].

But clearly something is amiss. Any physical system which can be described as shown in Figure 2.1, and which has no lower bound on Φ, is nothing less than a perpetual motion machine of the second kind! Is it true then, that the laws of thermodynamics prohibit the existence of active systems? The answer, which is surely obvious when considered on a universal scale, is yes. Thus, we have arrived at our first statement of passivity, which follows from the postulate that all physical systems must obey the laws of thermodynamics:

Passivity 1 (Thermodynamic). All physical systems are passive.

Of course, this statement poses a practical problem because engineers have access to a plethora of devices, such as electric motors and amplifiers, which are commonly referred to as active. The problem is that, when speaking of active (as opposed to passive) systems, engineers typically ignore interaction with certain other systems, not shown in Figure 2.1. For instance, it is inadequate from a thermodynamic point of view to model a motor as interacting with only its load and surroundings; yet, this is an approach which is frequently and justifiably taken in the process of modeling mechanical devices. The interaction with the electrical power source should also be included, and even this source should not be considered infinite if the laws of thermodynamics are to be satisfied in full.

Such a modeling effort would, however, while expanding the realm of validity, do little to improve the utility of the result. The point is that engineers are generally forced to generate models based upon limited[1] sets of observations, and to describe these models with mathematical equations which have limited bounds of validity. As a consequence, models often predict that certain systems are capable of performing infinite work on their environments, and extrapolation based on limited observations often bears this out. Within this context, the distinction between active and passive systems becomes useful. The next section explores the origin of active models.

2.2 Active Models

This section is intended to bridge the gap between the thermodynamic concepts of the last section, which do not allow active systems at all, and the concepts of the next section, which make a mathematical distinction between passivity and activity. We will make this bridge by examining the specific mechanisms with which mathematical mod-

[1] For instance, engineers do not have infinite time for observation.

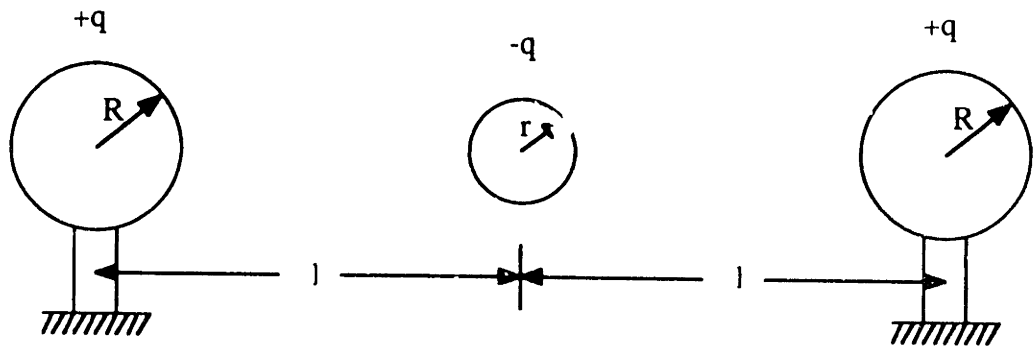

Figure 2.2: Passive system with an active linearized model

els can predict active behavior. For the purposes of illustration, these mechanisms will be divided into three categories, although these categories are not rigorously distinct.

The first and probably most ubiquitous mechanism is simply the use of sources in modeling. A battery, modeled as a voltage source, is an example. Wyatt, et al. [93] have pointed out that this sort of active model is due to multiple time scales. The time scale of interest for the electric circuit of which the battery is a part is generally too short for the battery to discharge significantly, thus the battery is modeled as not discharging at all. This type of extrapolation tends to be easily justified, because the model matches the relevant observation. Sources will be standard modeling elements for the rest of this document.

The second mechanism is the use of models with limited bounds of validity (though not necessarily limited in time), as mentioned in the previous section. To elaborate, consider a charged $(-q)$ particle located midway between two fixed spheres, each of charge $+q$ (Figure 2.2).

The force on this particle is zero, but it is clearly in a state of (locally) unstable equilibrium. If we were to linearize the restoring force as a function of displacement about the center position, the result would appear to be the constitutive equation of a negative spring:

$$\delta F = -\left(\frac{4kq^2}{l^3}\right)\delta x$$

where $k = 8.9879 \times 10^9\, Nm^2C^{-2}$. The potential energy of this linearized model (relative

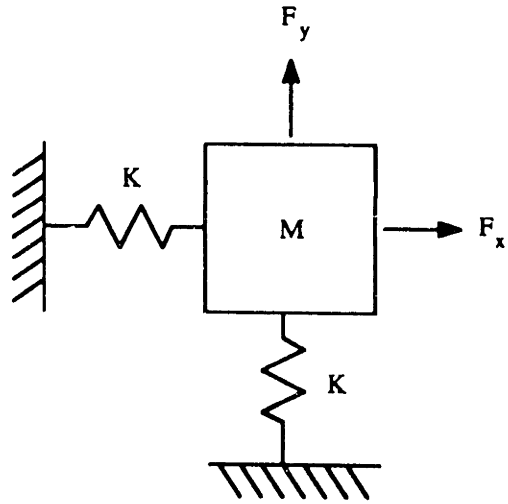

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \ddot{x} \\ \ddot{y} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\frac{K}{M} & 0 & 0 & 0 \\ 0 & -\frac{K}{M} & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ \dot{x} \\ \dot{y} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \frac{1}{M} & 0 \\ 0 & \frac{1}{M} \end{bmatrix} \begin{bmatrix} F_x \\ F_y \end{bmatrix}$$

Figure 2.3: Two-dimensional mass-spring system

to the energy as $\delta x \to \infty$) is infinite for any finite δx, suggesting that the system can supply infinite energy. However, analysis of the nonlinear model based upon Coulomb's Law, together with the assumptions that the spheres behave as point charges and cannot penetrate one another, leads to the following maximum energy which can be extracted from the system:

$$E = kq^2 \left(\frac{1}{R+r} + \frac{1}{2l-R-r} - \frac{2}{l} \right)$$

For $r > 0$ and $R > 0$, this solution is certainly not infinite. Of course, not even the nonlinear model is completely valid, but the point is that active behavior can be a consequence of a particular model's limitations.

The final mechanism, which presumes the first (sources), is somewhat more subtle. It arises because mathematical representations, such as state equations, of physical systems are quite capable of disguising energy sources. This can occur when a particular source is state-dependent; i.e., when a system is feedback controlled.

Consider, for instance, the system in Figure 2.3. If we were to implement the

following feedback law,

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} 0 & -K \\ K & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

the result would be a system described by the following state equations:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \ddot{x} \\ \ddot{y} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\frac{K}{M} & -\frac{K}{M} & 0 & 0 \\ \frac{K}{M} & -\frac{K}{M} & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ \dot{x} \\ \dot{y} \end{bmatrix}$$

Although no energy source is evident in these equations, the vector field relating steady state forces to displacements has non-zero curl. This suggests that infinite energy can be extracted from this closed-loop system, simply by allowing the mass to move continually in a closed, counterclockwise path (while turning a crank, for instance).

Note that feedback does not necessarily result in active behavior. We could just as well have chosen the following feedback law,

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} 0 & -K \\ -K & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

which would have resulted in closed-loop behavior indistinguishable from that of an appropriately designed collection of springs and masses.

The result, however, is that state-dependent sources are an important mechanism for generating active behavior, but *the presence of these sources is not necessarily evident on the basis of a state description alone.* Nevertheless, state equations rather than physical models provide the starting point for many of the theoretical developments found in the network analysis and control systems literature, so that an analysis of passivity based upon a state equation formulation proves quite useful.

2.3 General State Space Concepts of Passivity

In this section the results of Willems [90] and Wyatt, et al. [93] concerning passivity will be informally presented. This presentation is intended to provide the reader with

an introduction to the state-space theory of passivity and to pave the way for the linear analyses of the next section without fostering the naive belief that linear concepts of passivity have global application. It is not the object of this presentation to provide a rigorous statement of passivity. Rigor, as well as a number of curious examples that justify the need for rigor, can be found in [93].

The concepts presented in [90,93][2] apply to finite-dimensional, time-invariant systems described by state equations:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{g}(\mathbf{x},\mathbf{u}) \\ \mathbf{y} &= \mathbf{h}(\mathbf{x},\mathbf{u}) \end{aligned} \tag{2.8}$$

The system described by these equations is taken to have n interaction ports. A power input function, $p(\mathbf{x},\mathbf{u})$, which is the sum of the products of generalized effort and flow at each of the ports, is defined:

$$p(\mathbf{x},\mathbf{u}) = \langle \mathbf{e}(\mathbf{x},\mathbf{u}), \mathbf{f}(\mathbf{x},\mathbf{u}) \rangle \tag{2.9}$$

Before making a definition of passivity, several assumptions are necessary:

1. $\mathbf{g}$, $\mathbf{h}$, $\mathbf{e}$, and $\mathbf{f}$ are continuous functions.

2. For every initial condition $\mathbf{x}_0$ and input $\mathbf{u}(t)$, there exists a unique solution $\mathbf{x}(t)$ to the differential equation $\dot{\mathbf{x}} = \mathbf{g}(\mathbf{x},\mathbf{u})$.

3. The set of inputs is translation invariant and closed under concatenation; i.e., if $\mathbf{u}(t)$ is an admissable input, so is $\mathbf{u}(t+\tau)$, and if $\mathbf{u}_1(t)$, $t_1 \leq t \leq t_2$, and $\mathbf{u}_2(t)$, $t_2 \leq t \leq t_3$, are admissable inputs, so is $\mathbf{u}_1$ followed by $\mathbf{u}_2$.

4. $p(\mathbf{x},\mathbf{u})$ is locally L^1, i.e.:

$$\int_{t_1}^{t_2} \| p(\mathbf{x},\mathbf{u}) \| \, dt < +\infty \text{ for every choice of } t_1 \geq 0, t_2 \geq 0.$$

[2]The analyses of passivity found in Willems and Wyatt, et al. are quite similar—identical for our purposes. The notation used here will be that of Wyatt, et al

Given these assumptions, the *available energy* (Wyatt, et al.) or *available storage* (Willems) can be defined and related to passivity. The available energy is defined as follows:

$$E_A(\mathbf{x}) \triangleq \sup_{\substack{\mathbf{x}\rightarrow \\ T\geq 0}} \left\{ -\int_0^T p(\mathbf{x},\mathbf{u})\, dt \right\} \tag{2.10}$$

where the notation $\mathbf{x} \rightarrow$ indicates that the supremum is taken over all state trajectories starting in the state $\mathbf{x}$ at $t = 0$, and where the supremum is taken over all admissable inputs. Note that $E_A(\mathbf{x})$ is a non-negative function since it is the supremum over a set of numbers which include 0 ($T = 0$). The intuitive relation of the available energy to availability should be fairly evident—both represent the maximum work that can be extracted from a particular system. Specifically, the available energy for a particular initial state $\mathbf{x}$ represents the maximum energy (all of which may be considered to be useful work) which can be extracted from the system when its initial state is $\mathbf{x}$.

At this point a statement of passivity (formally, a theorem) is straightforward:

Passivity 2a (State-Space). A dynamic system is passive if, for each initial state $\mathbf{x}$, $E_A(\mathbf{x}) < +\infty$. Otherwise, it is active.

This is clearly a rather weak definition of passivity, particularly in that E_A must be calculated for every initial state $\mathbf{x}$. We can alleviate this problem somewhat by considering only those systems which are *completely controllable*. For a completely controllable system there always exists a control $\mathbf{u}(t)$, $0 \leq t \leq T < \infty$, which will take a system from each state $\mathbf{x}_0$ to any other state $\mathbf{x}_1$. Further, because of assumption (4), this transfer will consume or produce only finite energy, and, as a consequence, we need examine only one initial state:

Passivity 2b. A system with a completely controllable state representation is passive iff $E_A(\mathbf{x}_0) < +\infty$, where $\mathbf{x}_0$ is any initial state.

Finally, Wyatt et al. have strengthened the definition further for those systems

which have states of zero stored energy, called *relaxed states.* A state $\mathbf{x}$ is said to be relaxed if $E_A(\mathbf{x}) = 0$.

Passivity 2c. If a system with a completely controllable state representation has a relaxed state $\mathbf{x}^*$, then the system is passive.

An alternative statement of passivity is in terms of *internal energy functions* or *storage functions.* $E_I(\mathbf{x})$ is an internal energy function of a system if:

$$E_I(\mathbf{x}(t_2)) - E_I(\mathbf{x}(t_1)) \leq \int_{t_1}^{t_2} p(\mathbf{x}(t), \mathbf{u}(t))\, dt \tag{2.11}$$

for all inputs $\mathbf{u}(t)$, and all $0 \leq t_1 \leq t_2$. An internal energy function is a function of the state which is bounded from below and increases along trajectories in the state space more slowly than the rate at which energy is delivered to the ports. Note that, if a system is passive, E_A is an internal energy function.

The following theorem relates internal energy to passivity:

Passivity 3 (State-Space). Given a system with a state representation subject to the same assumptions used earlier, the system is passive iff there exists an internal energy function $E_I(\mathbf{x})$ defined over the state space.

If the internal energy function is sufficiently smooth, then a differential version of equation 2.11 exists:

$$\dot{E}_I(\mathbf{x}, \mathbf{u}) \leq p(\mathbf{x}, \mathbf{u}) \tag{2.12}$$

The use of an internal energy function can be attractive because of its potential application as a Lyapunov function. It is important to realize, however, that an internal energy function need not qualify as a Lyapunov function, as the latter is subject to more stringent requirements on shape [86].

This completes our brief tour of nonlinear systems passivity. The next section will treat the more specific class of linear, time-invariant, finite-dimensional systems.

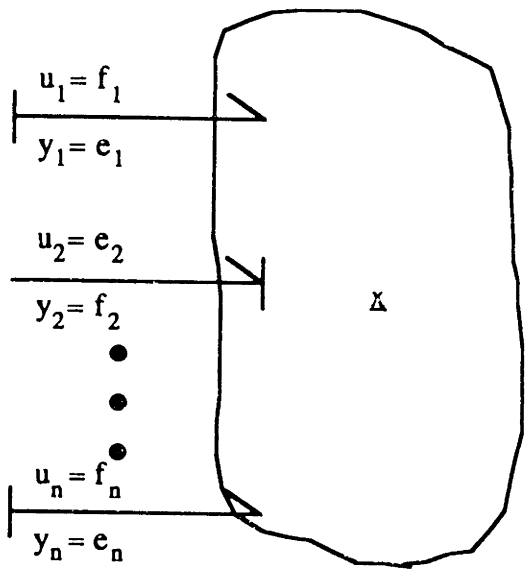

Figure 2.4: n-Port

2.4 Passivity of Linear Systems

The purpose of this section is to derive quantitative statements of the passivity of linear, time-invariant, finite-dimensional systems described as follows:

$$\begin{aligned} \dot{\mathbf{x}} &= A_{k\times k}\mathbf{x} + L_{k\times n}\mathbf{u} \\ \mathbf{y} &= C_{n\times k}\mathbf{x} + D_{n\times n}\mathbf{u} \end{aligned} \tag{2.13}$$

where $\mathbf{u}$ and $\mathbf{y}$ are a hybrid pair [4] of power duals, i.e., generalized efforts and flows associated with the n interaction ports (Figure 2.4). The developments in this section are fairly standard, similar to those found in several texts on network analysis. In particular, this presentation draws heavily upon the analyses of Kuh and Rohrer [53], Anderson and Vongpanitlerd [4], and the seminal work of Brune [13].

We will assume that the system above is a *minimal realization* [4,91], which is to say that it is both controllable and observable. The assumption of controllability was motivated in the previous section; with it, the initial state is not a consideration. The additional assumption of observability allows us to tie together the external (in-

put/output) and internal (state-based) analyses of passivity. In particular, it enables us to disregard those pathological cases in which the Lyapunov energy—defined to be a positive definite quadratic function of state[3]—grows without bound due to an unstable mode which also happens to be unobservable.

Certain assumptions must also be made regarding the inputs:

1. The set of inputs is closed under concatenation (see page 43).

2. The set of inputs is translation invariant (see page 43).

3. The set of inputs is $L^2(-\infty, t]$, i.e.:

$$\int_{-\infty}^{t} \| \mathbf{u}(t) \|^2 \, dt < +\infty$$

The third assumption is essentially a boundedness criterion on the input; the utility of each of these assumptions will become evident in the sequel. It should be noted that none of the assumptions is restrictive for systems of engineering interest.

The statement of passivity in terms of internal energy (equation 2.11) can now be applied to linear n-ports.

A linear n-port is passive iff:

$$\int_{t_1}^{t_2} \mathbf{u}'(\tau)\mathbf{y}(\tau) \, d\tau \geq E(t_2) - E(t_1) \tag{2.14}$$

for all admissable $n \times 1$ $\mathbf{u}(t)$ and $\mathbf{y}(t)$ and all t_1 and $t_2 > t_1$. $E(t)$ is now taken to be the Lyapunov energy, which, for a minimal realization, is a suitable internal energy function. The Lyapunov energy is bounded below by zero. The derivative form of

[3]The Lyapunov energy is defined as follows: $E_{Lyapunov} = \mathbf{x}'P\mathbf{x}$, where P is a positive definite matrix. It is used here rather than the internal energy to emphasize the dependence on all states of the system. The internal energy functions defined in the previous section can ignore an unstable, unobservable mode, Lyapunov energy cannot.

equation 2.14 is also a suitable statement of passivity:

$$\mathbf{u}'(t)\mathbf{y}(t) \geq \frac{dE}{dt} \tag{2.15}$$

Equations 2.14 and 2.15, despite their intuitive appeal, can be somewhat unweildy because they require measurement of the port variables $\mathbf{u}$ and $\mathbf{y}$ as well as the energy. A criterion which does not require the measurement of energy can be found as follows. Because the linear system is controllable, there exists some control which will take the system from its state at t_2 to the state $\mathbf{x} = 0$ at $t_2 < t_3 < +\infty$. Because the inputs are closed under concatenation, the following equation holds:

$$\int_{t_1}^{t_3} \mathbf{u}'(\tau)\mathbf{y}(\tau)\, d\tau \geq -E(t_1) \tag{2.16}$$

This equation essentially states that a passive system can deliver no more energy than that amount which is initially stored in it. Next, we make the assumption that the Lyapunov energy approaches zero (i.e., the system is at rest) as $t_1 \to -\infty$, thus:

Passivity 4 (Linear n-Port). A linear n-port is passive iff:

$$\int_{-\infty}^{t} \mathbf{u}'(\tau)\mathbf{y}(\tau)\, d\tau \geq 0 \tag{2.17}$$

This is the statement of passivity which will be used for most of what follows, which is the development of practical passivity criteria in terms of both input/output and state-space descriptions. The input/output criteria will be considered first, but for historical reasons only (they predate the state-space criteria by approximately 30 years); they are in no way the more fundamental of the two.

2.4.1 Input/Output Passivity Criteria

The various input/output passivity criteria will be generated in the frequency domain. When working in the frequency domain, it can be useful to express physical quantities

of interest as the real parts of complex quantities. Thus, we will consider power duals $\mathbf{u}$ and $\mathbf{y}$ of the following form:

$$\begin{aligned} \mathbf{u}(t) &= \mathbf{u}_0 e^{\sigma t}\cos(\omega t) = Re\{\mathbf{u}_0 e^{st}\} \\ \mathbf{y}(t) &= \mathbf{y}_0 e^{\sigma t}\cos(\omega t + \phi) = Re\{\mathbf{y}_0 e^{st+j\phi}\} \end{aligned} \tag{2.18}$$

where $s = \sigma + j\omega$. Here, the the vectors $\mathbf{u}_0$ and $\mathbf{y}_0$ are real, and the assumption is made (temporarily) that the components of $\mathbf{y}$ all share the same phase. The power delivered to this n-port at any time t is:

$$\mathbf{u}'(t)\mathbf{y}(t) = \mathbf{u}_0'\mathbf{y}_0 e^{2\sigma t}\cos(\omega t)\cos(\omega t + \phi) \tag{2.19}$$

With the aid of the following trigonometric identity, which identifies the non-periodic part of equation 2.19, the net energy delivered at any time t may be evaluated:

$$\cos(\omega t)\cos(\omega t + \phi) = \frac{1}{2}\cos(\phi) + \frac{1}{2}\cos(2\omega t)\cos(\phi) - \frac{1}{2}\sin(2\omega t)\sin(\phi)$$

Thus,

$$\int_{-\infty}^{t} \mathbf{u}'(\tau)\mathbf{y}(\tau)\,d\tau = \frac{1}{4\sigma}\mathbf{u}_0'\mathbf{y}_0 e^{2\sigma t}\cos(\phi) \tag{2.20}$$

for $\sigma \geq 0$.

Now consider an alternative definition of passivity for linear systems:

A linear n-port is passive iff:

$$Re\left\{\int_{-\infty}^{t} \mathbf{u}^H(\tau)\mathbf{y}(\tau)\,d\tau\right\} \geq 0 \tag{2.21}$$

for all admissable $\mathbf{u}(t)$ and $\mathbf{y}(t)$ and all t. In this case, $\mathbf{u}$ and $\mathbf{y}$ may be complex quantities, not just the real parts. Given $\mathbf{u}$ and $\mathbf{y}$ in the form of equations 2.19, definition 2.21 reduces to:

$$Re\left\{\int_{-\infty}^{t} \mathbf{u}_0' e^{s^*\tau}\mathbf{y}_0 e^{s\tau+j\phi}\,d\tau\right\} = \frac{1}{2\sigma}\mathbf{u}_0'\mathbf{y}_0 e^{2\sigma t}\cos(\phi) \geq 0 \tag{2.22}$$

where s^* is the complex conjugate of s, and $\sigma \geq 0$.

Comparison with equation 2.20, reveals that the expression in equation 2.22 is simply twice the net energy delivered to the system. This result is readily generalized to the case in which $\mathbf{u}_0$ and $\mathbf{y}_0$ are complex vectors, and to the case in which each component of the vectors $\mathbf{u}$ and $\mathbf{y}$ may have a distinct phase. The consequence is that equation 2.21 is an acceptable statement of passivity, providing results identical to equation 2.17, and further, offering simpler computation.

The strength of a complex-variable approach is that it lets us examine passivity in terms of the properties of the driving point impedance matrix[4] relating the Laplace transforms of $\mathbf{u}$ and $\mathbf{y}$:

$$\begin{aligned} \mathbf{y}(s) &= Z(s)\mathbf{u}(s) \\ \text{where } Z(s) &= D + C(sI - A)^{-1}L \end{aligned} \tag{2.23}$$

It is now necessary to consider carefully the relation of the class of inputs $\mathbf{u}(t)$ to the inputs $\mathbf{u}(s)$. The assumptions that $\mathbf{u}(t)$ is translation invariant and $L^2(-\infty, t]$ are now important. Because $\mathbf{u}(t)$ is translation invariant it is clear that the following equation holds:

$$\int_{-\infty}^{t} \| \mathbf{u}(\tau - t) \|^2 \, d\tau < +\infty$$

Introducing the change of variables $\tau^* = \tau - t$:

$$\int_{-\infty}^{0} \| \mathbf{u}(\tau^*) \|^2 \, d\tau^* < +\infty$$

Therefore the class of inputs $\mathbf{u}(t)$ is also $L^2(-\infty, 0]$. This is useful because the Fourier transform is a Hilbert space isomorphism from $L^2(-\infty, \infty)$ to L_2 [62]. In particular, $L^2(-\infty, 0]$ is mapped onto $H_2^{\perp}$, the set of vectors whose components are strictly proper and have no poles in $Re\{s\} < 0$. The partial fraction expansion of a component of such

[4] In general, hybrid impedance/admittance matrix.

a vector would result in:

$$\mathbf{u}(s) = \sum_i \frac{(\text{residue})_i}{s - s_i}, \text{ where } s_i = \sigma_i + j\omega_i$$

Because $\mathbf{u}(s)$ is in $H_2^\perp$, $\sigma_i \geq 0$. The inverse transform of each term in the sum would be of the form:

$$u(t) = u_i e^{s_i t}$$

where u_i is, in general, complex, and $\sigma_i \geq 0$. All inputs can now be formed by linear combinations of inputs of this form. Thus, we will investigate passivity for inputs and outputs of the following form:

$$\begin{aligned} \mathbf{u}(t) &= \mathbf{u}_0 e^{s_0 t}, \quad s_0 = \sigma_0 + j\omega_0, \ \sigma_0 \geq 0 \\ \mathbf{y}(t) &= Z(s_0)\mathbf{u}_0 e^{s_0 t} \end{aligned} \tag{2.24}$$

Passivity requires the following:

$$Re\left\{\int_{-\infty}^{t} \mathbf{u}_0^H Z(s_0)\mathbf{u}_0 e^{2\sigma_0 \tau}\, d\tau\right\} \geq 0 \tag{2.25}$$

for all $\mathbf{u}_0$, all t, and all $\sigma_0 \geq 0$. Carrying out the integral leads to:

$$\frac{1}{2\sigma_0} Re\{\mathbf{u}_0^H Z(s_0)\mathbf{u}_0\} e^{2\sigma_0 t} \geq 0 \tag{2.26}$$

and to:

$$Re\{\mathbf{u}_0^H Z(s_0)\mathbf{u}_0\} \geq 0 \tag{2.27}$$

for all admissable $\mathbf{u}_0$ and all $\sigma_0 \geq 0$. Equation 2.27 may be rewritten as follows:

$$\begin{aligned} \frac{1}{2}\{[\mathbf{u}_0^H Z(s_0)\mathbf{u}_0] + [\mathbf{u}_0^H Z(s_0)\mathbf{u}_0]^H\} &\geq 0 \\ \mathbf{u}_0^H [Z(s_0) + Z^H(s_0)]\mathbf{u}_0 &\geq 0 \end{aligned} \tag{2.28}$$

Equation 2.28 will be satisfied so long as the matrix $Z + Z^H$ is non-negative definite Hermitian for all $\sigma_0 \geq 0$. This is the definition of a *positive matrix*. Because we will be concerned with physical systems only, it will always be the case that Z is a *real matrix*, i.e., that $Z(s)$ is real for real s. Thus, the input/output passivity criterion is:

Passivity 5a (Linear n-Port). A linear time-invariant n-port is passive iff $Z + Z^H$ is a *positive real matrix.*

In order to explore this criterion further, consider now the case of 1-port systems, i.e., systems for which u and y are scalars. In this case, the passivity condition reduces to the following:

Passivity 6a (Linear 1-Port). A linear time-invariant 1-port is passive iff:

$$Re\{Z(s)\} \geq 0 \text{ for } \sigma \geq 0 \tag{2.29}$$

A function $Z(s)$ which is a real, rational function of s and which satisfies this condition, is called a *positive real* function. This criterion, although rather simple, can be computationally complex because it requires evaluation of $Re\{Z(s)\}$ at each point in the right half s-plane. It can be readily shown [13], however, that the following three conditions are equivalent to equation 2.29.

Passivity 6b (Linear 1-Port). A linear time-invariant 1-port is passive iff:

1. $Z(s)$ has no poles in the right half plane.
2. Any imaginary poles of $Z(s)$ are simple, and have positive real residues.
3. $Re\{Z(j\omega)\} \geq 0$.

These conditions are essentially the result of an application of the priciple of maximum modulus [85] to a contour which encloses the region over which $\sigma \geq 0$, i.e., the right half s-plane. In order to apply this principle, however, the contour cannot enclose any poles; this is the reason for condition 1. Right half plane poles (and zeros) are, in fact, prohibited by equation 2.29, because $Re\{Z(s)\}$ must take on negative values in the neighborhood of *any* singularity. Imaginary poles are acceptable so long as the negative real part occurs to the left of the pole only [13]; this is the reason for condition

2. The contour, however, must be indented to avoid any imaginary poles, thus it is just the familiar Nyquist contour. As a result, it is possible to combine conditions 2 and 3, and to arrive at the following statement of passivity:

Passivity 6c (Linear 1-Port). A linear time-invariant 1-port is passive iff:

1. $Z(s)$ has no right half plane poles.
2. $Z(s)$ has a Nyquist plot which lies wholly within the closed right half plane.

An interesting consequence of passivity which is illustrated by this statement is that the phase of $Z(s)$ must lie between $+90°$ and $-90°$. The phase of $Z(s)$ is, of course, the difference between the phase of the output waveform $(y(s))$ and the phase of the input waveform $(u(s))$. Consider the case in which both are pure sinusoids. If the phase is $0°$, the average value of the product $u(t)y(t)$ is positive, so that the system consumes energy. If the phase is $\pm 90°$, the average is zero, and the system is called lossless[5]. If the phase is $< -90°$ or $> +90°$, the average value is negative, and the system must produce energy.

One may wonder if conditions similar to those above can be found in the n-port case. In fact, they can. Consider again the n-port passivity criterion: $Z(s) + Z^H(s)$ is non-negative definite Hermitian for all $\sigma \geq 0$. As in the 1-port case, a set of equivalent conditions can be found [4]:

Passivity 5b (Linear n-Ports). A linear time-invariant n-port is passive iff:

1. $Z(s) + Z^H(s)$ has no poles in the right half plane.
2. Any imaginary poles are simple and have positive real residue matrices.
3. $Z(j\omega) + Z^H(j\omega)$ is non-negative definite Hermitian.

[5]for further discussion of losslessness, see [4].

Given these conditions, $Z + Z^H$ need be tested only for right half plane poles and along the imaginary axis, following an indented contour around any imaginary poles. Because a non-negative definite Hermitian matrix must have zero or positive eigenvalues, one test would be to plot these eigenvalues as functions of ω.

A computationally less expensive test, however, would be to plot the determinants of the n upper left submatrices of $Z + Z^H$ as functions of ω. Another consequence of the non-negative definite Hermitian property is that these determinants must all be zero or positive [77]. Finally, imaginary poles may be treated by applying the same test to their residue matrices, rather than by evaluating these determinants along an indention.

To complete this tour through input/output passivity criteria, we will return to the single interaction port case and point out several corollaries of the passivity criterion which provide some useful insights into the nature of positive real functions. These are:

Corollary 1. The relative degree of $Z(s)$ may be only -1, 0, or 1 (as implied by the previously mentioned phase angle condition).

Corollary 2. $1/Z(s)$ is also positive real; thus, like the poles, the zeros of $Z(s)$ must lie in the closed left half plane.

Corollary 3. The zeros of $Z(s)$ on the $j\omega$ axis (including those at 0 and ∞) must be simple with positive differential quotients; i.e., for a zero at $j\omega_i$:

$$\lim_{s \to j\omega_i} \left(\frac{Z(s)}{s - j\omega_i} \right) \geq 0 \tag{2.30}$$

Corollary 1 may be interpreted to mean that, at high frequency, any positive real driving point impedance must appear to be the driving point impedance of a mass, a spring, or a viscous damper (to use mechanical elements). Corollary 2 indicates that a driving point impedance may not exhibit nonminimum phase behavior. Corollary 3 is a useful tool to have in a visual inspection of a driving point impedance function;

PID controllers, for instance, often lead to closed-loop driving point impedances which violate this condition.

2.4.2 Network Synthesis

The purpose of this section is to introduce the idea of *network synthesis*, which will be a useful preliminary to the state-space criteria of the next section, as well as an essential notion in Chapter 6. The various passivity criteria which are being presented in this chapter are found almost exclusively in the electrical engineering literature; in particular, in the network analysis and synthesis literature. Various analytical tools, such as tests for positive realness, are typically intended to serve as preliminaries to the real problem of interest, which is: given an impedance which is a prescribed function of frequency, find an electrical network which has that behavior. This is the network synthesis problem.

One of the seminal results in this field concerns the synthesis of passive networks [13]:

> If a driving point impedance function is positive real, then at least one network composed of passive, linear elements (I, C, R, TF, GY) exists which exhibits that impedance.

Of course, this result can be generalized to positive real driving point impedance matrices. This theorem will not be proven here; however, it is motivated in Chapter 6 with a discussion of network synthesis procedures.

This result has certain consequences of immediate benefit. One of these is that any system with a positive real driving point impedance, even if it is feedback-controlled, has an equivalent representation, as seen at the ports, which is composed of purely passive, linear elements. This concept lies at the heart of the developments in Chapter 6.

The second consequence stems from the first. Because a passive, linear system always corresponds to some *realization* composed of passive, linear elements, the power flowing into the ports may be accounted for exactly as:

$$Power_{in} = \frac{d}{dt} Energy_{stored} + \frac{d}{dt} Energy_{dissipated}$$

where energy is stored in I elements (e.g. masses, inductors) and C elements (e.g. springs, capacitors), and dissipated through R elements (e.g. viscous dampers, ohmic resistors). The utility of this statement will soon be evident.

2.4.3 State-Space Passivity Criteria

The state-space passivity criteria will be developed in the time domain in terms of the state and output equations 2.13. To develop the necessary conditions for passivity, we will assume that, as explained in the previous section, any passive linear system can be represented by some collection of linear elements. A power balance can, therefore, be written:

$$\mathbf{u}'(t)\mathbf{y}(t) = \frac{d}{dt}(E_{stored}) + \frac{d}{dt}(E_{dissipated}) \tag{2.31}$$

The reader might be concerned that this is not a legitimate starting point as it apparently requires some *a priori* understanding of the relationship between passivity and system structure. This is, in fact, true, but is a rather fine point that should not detract from an intuitively appealing derivation. A derivation similar to the one presented here, as well as several more rigorous derivations of state-space passivity criteria, can be found in [4].

For a linear system of minimal dimension, E_{stored} and $\frac{d}{dt}E_{dissipated}$ may be represented as follows:

$$E_{stored} = \frac{1}{2}\mathbf{x}'P\mathbf{x} \tag{2.32}$$

$$\frac{d}{dt}E_{dissipated} = (\Phi\mathbf{x} + \Gamma\mathbf{u})'R(\Phi\mathbf{x} + \Gamma\mathbf{u}) \tag{2.33}$$

where P is a positive definite (PD) matrix, and R is a positive definite diagonal matrix.

The equation for E_{stored} should seem reasonable because the energy in a linear storage element (a spring or a mass) may be represented as a quadratic function of state; however, it can be motivated further by an argument based on bond graph modeling. The crux of the argument is that a bond graph model of a minimal realization of order k will contain exactly k independent storages [6]. We may reflect each of these storages through an appropriately chosen transformer so that each of the reflected storages has unity magnitude ($I = 1$ or $C = 1$), and select as state variables the flows on the inertias (f_i^*) and the efforts on the capacitances (e_i^*):

$$\mathbf{x}^* = \begin{bmatrix} \mathbf{f}^* \\ \mathbf{e}^* \end{bmatrix}$$

Because the storages in the system are all independent, the total energy is given by the sum of the energies in the individual storages:

$$E = \frac{1}{2}\mathbf{x}'^*\mathbf{x}^*$$

Now we appeal to the fact that any two state representations of a linear system can be related by a similarity transform. Thus, a nonsingular matrix T exists which satisfies:

$$T\mathbf{x} = \mathbf{x}^*$$

We can use this transform to express the stored energy in terms of the original states:

$$E = \frac{1}{2}\mathbf{x}'T'T\mathbf{x}$$

If we define $P = T'T$, P is clearly a positive definite matrix, and equation 2.32 is restored.

The equation for $\frac{d}{dt}E_{dissipated}$ is simply the sum of the power lost to each of the dissipative elements in the system. The flow into each dissipator is a linear function of the states and the inputs:

$$f_i = \phi\mathbf{x} + \gamma\mathbf{u}$$

and,

$$\frac{d}{dt}E^i_{dissipated} = r_i f_i^2$$

where r_i is the viscous coefficient (resistance) of the i^{th} element. If $\mathbf{f}_r$ is a column vector of flows into the system's dissipators, then:

$$\mathbf{f}_r = \Phi\mathbf{x} + \Gamma\mathbf{u}$$

and,

$$\frac{d}{dt}E_{dissipated} = \mathbf{f}_r' R \mathbf{f}_r$$

where R is a diagonal matrix of the r_i. Thus, equation 2.33 is restored.

The power balance may now be written as follows:

$$\mathbf{u}'\mathbf{y} = \frac{1}{2}\frac{d}{dt}\mathbf{x}'P\mathbf{x} + (\Phi\mathbf{x} + \Gamma\mathbf{u})'R(\Phi\mathbf{x} + \Gamma\mathbf{u}) \tag{2.34}$$

Taking the derivatives and using the state and output equations 2.13:

$$\begin{aligned}
\mathbf{u}'C\mathbf{x} + \mathbf{u}'D\mathbf{u} &= \frac{1}{2}\mathbf{x}'A'P\mathbf{x} + \frac{1}{2}\mathbf{u}'L'P\mathbf{x} + \frac{1}{2}\mathbf{x}'PA\mathbf{x} + \frac{1}{2}\mathbf{x}'PL\mathbf{u} \\
&\quad + \mathbf{x}'\Phi'R\Phi\mathbf{x} + \mathbf{u}'\Gamma'R\Phi\mathbf{x} + \mathbf{x}'\Phi'R\Gamma\mathbf{u} + \mathbf{u}'\Gamma'R\Gamma\mathbf{u} \\
\mathbf{u}'C\mathbf{x} + \mathbf{u}'D\mathbf{u} &= \mathbf{x}(\frac{1}{2}PA + \frac{1}{2}A'P + \Phi'R\Phi)\mathbf{x} + \mathbf{u}'(L'P + 2\Gamma'R\Phi)\mathbf{x} \\
&\quad + \mathbf{u}'(\Gamma'R\Gamma)\mathbf{u}
\end{aligned} \tag{2.35}$$

Because this equation holds for all $\mathbf{x}$ and $\mathbf{u}$, the following equations must be satisfied:

$$\begin{aligned}
0 &= \frac{1}{2}PA + \frac{1}{2}A'P + \Phi'R\Phi \\
C &= L'P + 2\Gamma'R\Phi \\
D &= \Gamma'R\Gamma
\end{aligned} \tag{2.36}$$

Now defining $M = \Phi'R^{1/2}$ and $W = R^{1/2}\Gamma$, these equations reduce to the following:

$$\begin{aligned}
PA + A'P &= -2MM' \\
PL &= C' - 2MW \\
D &= W'W
\end{aligned} \tag{2.37}$$

Very rarely in mechanical systems does one encounter a situation in which $D \neq 0$. Therefore, the version of these passivity criteria which is most frequently encountered is:

$$\begin{aligned} PA + A'P &= -Q \\ PL &= C' \end{aligned} \tag{2.38}$$

where Q is postive semidefinite (PSD).

The sufficiency of equations 2.37 in establishing passivity can be demonstrated as follows. For any matrices M and W:

$$(M'\mathbf{x} + W\mathbf{u})'(M'\mathbf{x} + W\mathbf{u}) \geq 0 \tag{2.39}$$

$$\mathbf{x}'MM'\mathbf{x} + 2\mathbf{u}'W'M\mathbf{x} + \mathbf{u}'W'W\mathbf{u} \geq 0 \tag{2.40}$$

With the aid of equations 2.37, equation 2.40 may be rewritten as follows:

$$-\frac{1}{2}\mathbf{x}'(PA + A'P)\mathbf{x} + \mathbf{u}'(C - L'P)\mathbf{x} + \mathbf{u}'D\mathbf{u} \geq 0 \tag{2.41}$$

$$\mathbf{u}'C\mathbf{x} + \mathbf{u}'D\mathbf{u} \geq \frac{1}{2}\mathbf{x}'PA\mathbf{x} + \frac{1}{2}\mathbf{x}'A'P\mathbf{x} + \frac{1}{2}\mathbf{u}'L'P\mathbf{x} + \frac{1}{2}\mathbf{x}'PL\mathbf{u} \tag{2.42}$$

$$\mathbf{u}'\mathbf{y} \geq \frac{1}{2}\frac{d}{dt}(\mathbf{x}'P\mathbf{x}) \tag{2.43}$$

Because $\frac{1}{2}\mathbf{x}'P\mathbf{x}$ can serve as an internal energy function (it need not necessarily represent the energy in the storage elements), this establishes passivity by equation 2.15. In summary:

Passivity 7a (Linear n-Ports). A linear, time-invariant n-port described by equations 2.13 is passive iff the following equations are satisfied for some PD P and some M and W:

$$\begin{aligned} PA + A'P &= -2MM' \\ PL &= C' - 2MW \\ D &= W'W \end{aligned}$$

Passivity 7b (Linear n-Ports). A linear, time-invariant n-port described by equations 2.13 with $D = 0$ is passive iff the following equations are satisfied for some PD P and for some PSD Q:

$$\begin{aligned} PA + A'P &= -Q \\ PL &= C' \end{aligned}$$

The fact that, for a 1-port, the equations 2.37 are satisfied if and only if a positive real function $Z(s) = D + C(sI - A)^{-1}L$ exists, is known as the *Kalman-Yacubovitch Lemma*, and has received widespread application in the study of absolute stability (see, for instance, Narendra and Taylor [61]). The generalization to n-ports is known as the *Positive Real Lemma* [4].

2.5 An Alternative Criterion for Positive Real Functions

The purpose of this section is to present an alternative set of necessary and sufficient conditions for positive realness, and, equivalently, an alternative criterion for the passivity of a 1-port. Although this criterion was independently developed for applications relevant to this thesis, it may be derived in a very simple fashion from one proposed by Talbot [79,80]. A physically motivated explanation of this criterion, as well as an application to the stability of coupled dynamic systems, is presented in Section 3.3.

A surprisingly large number of alternative tests for positive real functions have been reported in the literature. The more familiar of these were presented in the last section; one good source of others is Van Valkenburg [85]. Among the others are the Sturm test, which involves the construction of a Routh-like array, and a test due to Talbot which employs a conformal map to transform what is essentially a criterion on phase

to a criterion on magnitude (Anderson performs a similar transformation with positive real matrices [4]).

More relevant to the current development, however, are tests due to Brockett [12] and Talbot [79,80]. Brockett's test is the following:

> If $Z(s) = N(s)/D(s)$, where $N(s)$ and $D(s)$ are polynomials with real coefficients and without common factors, $Z(s)$ is positive real iff $N^2(s) + kD^2(s)$ is Hurwitz for all $0 < k < \infty$.

where Hurwitz means that the polynomial has no right half plane (RHP) zeros and any imaginary zeros are simple. The assumption that the polynomials have no common factors is equivalent to previous assumptions of minimality. Part of the value of this condition is that it is equivalent to the requirement that all branches of the root locus of $Z^2(s)$ lie in the left half plane (LHP). Thus, it provides a simple, graphical test for positive realness.

Approximately one and a half years after Brockett published the theorem above, Talbot published the following similar theorem:

> If $Z(s) = N(s)/D(s)$, where $N(s)$ and $D(s)$ are polynomials with real coefficients and without common factors, $Z(s)$ is positive real iff $k_1N(s) + jk_2D(s)$ has no RHP zeros for all real constants k_1 and k_2, not both zero.

This theorem also affords a graphical test, moreover, one that relates directly to the driving point impedance of interest, $Z(s)$. This test is simply that the locus of points where $\arg(Z(s)) = \pm\pi/2$ must lie in the LHP. This test has the disadvantage, however, that it does not employ so ubiquitous an analytical tool as root locus.

A minor variant of Talbot's condition, however, both applies directly to transfer functions of interest, (see Section 3.3), and takes advantage of root locus techniques. The following theorem leads to this test:

Passivity 8 (Linear 1-Ports). If $Z(s) = N(s)/D(s)$, where $N(s)$ and $D(s)$ are polynomials with real coefficients and without common factors, $Z(s)$ is positive real iff the polynomials $r_1(s) = sD(s) + kN(s)$ and $r_2(s) = D(s) + ksN(s)$ have no RHP zeros for all $0 \leq k < \infty$.

The equivalent root locus test is that all branches of the loci of $sZ(s)$ and $Z(s)/s$ must lie in the closed LHP. This theorem is evidently little more than a variant on the one proposed by Talbot; it is given special attention here because of the previously mentioned physical motivation. The proof of this theorem is as follows:

The method of proof is to demonstrate that the criterion above is equivalent to the three conditions which constitute Passivity 6b (page 52). This may be done as follows:

- For $k = 0$, $r_2(s) = D(s)$, thus $D(s)$ has no RHP zeros. This establishes conditon 1 of Passivity 6b. Furthermore, as $k \to \infty$, $r_1(s) \to kN(s)$, thus $N(s)$ has no RHP zeros.

- For all $\sigma > 0$, almost all $\sigma = 0$, all real ω, and all $k \geq 0$:

$$\begin{aligned} \frac{N(s)}{D(s)} &\neq -\frac{s}{k}, \quad s = \sigma + j\omega \\ \frac{N(s)}{D(s)} &\neq -\frac{1}{ks} \end{aligned}$$

  These two equations have the following consequences:

$$Re\{Z(\sigma \pm j\omega)\} > 0 \quad \text{for all } \sigma > 0 \tag{2.44}$$

$$Re\{Z(\sigma \pm j\omega)\} \neq 0 \quad \text{for almost all } \sigma = 0 \tag{2.45}$$

  Consider now the locus of points, shown in Figure 2.5, defined by $\varepsilon + j\omega$, $0 \leq \varepsilon \ll 1$, $-\infty < \omega < \infty$, except for a finite number of right half plane indentations of radius $0 \leq \delta \ll 1$ around the imaginary roots of $D(s)$.

  Equations 2.44 and 2.45 require that $Re\{Z(j\omega)\} \geq 0$ in the limit as $\varepsilon \to 0$. This establishes condition 3 of Passivity 6b.

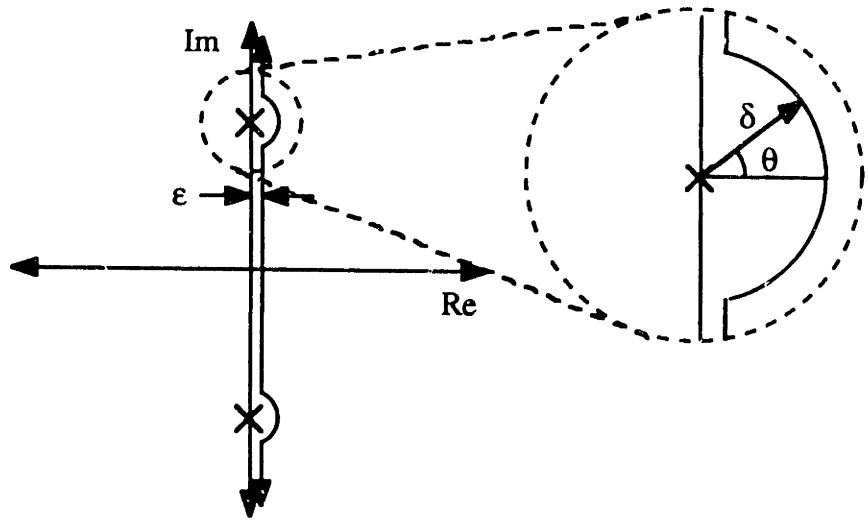

Figure 2.5: Locus of points for evaluation of $Re\{Z(\sigma + j\omega)\}$.

- Along the indentations, $Re\{Z(s)\} > 0$. The behavior of $Z(s)$ near a singularity, however, can be determined in terms of the parameter θ as follows:

$$Re\{Z(s)\} = \frac{Residue}{\delta} \cos(-n\theta) \tag{2.46}$$

where n is the multiplicity of the singularity. This expression will be greater than zero for all $\sigma > 0$ ($-90° < \theta < 90°$) only if $n = 1$ and $Residue > 0$. This establishes condition 2 of Passivity 6b, and completes the proof.

Chapter 3

Coupled Stability

The purpose of this chapter is to address the issue of the stability of a feedback-controlled system which is *interacting with a dynamic environment.* As indicated in Section 1.4.1, the stability of a controlled system coupled to its environment is the most important consideration created by our interest in the design of interaction controllers. As such, it seems natural to achieve a solid understanding of coupled stability before attempting to resolve other design issues. The next section serves to review the previous art, as well as to develop the rationale for the approach taken here.

The term "coupled stability" will be used extensively in this chapter. Although it will eventually be defined more carefully, the reader should, at this point, take it to mean quite literally the stability of the system composed of the manipulator and environment, mechanically coupled together.

3.1 Approaches to Coupled Stability

It should come as no great surprise to the reader that the bulk of the robotics literature does not address interaction with the environment. The influence of control theory is evident in that, if an environment is considered at all, it is typically treated as a source of disturbances rather than a dynamic system in its own right. Recently, however, an

increasing number of researchers have taken note of the importance of environment dynamics in overall system behavior (where the system is taken to be the robot and its environment). For instance, even if the environment is assumed to be a rigid body, it may still be useful from the point of view of control (particularly computed torque methods) to know the inertial parameters of that rigid body. Thus, several approaches to load determination, assuming rigid body loads, have been developed [9,51].

Although inertial load identification presumes stability and is intended to address performance, coupling to dynamic environments has been known to endanger robot stability. The most compelling example of current interest is the instability of force-feedback controlled robots upon contact with rigid surfaces. "Contact instability" will be analyzed in detail in Chapter 7; for now, it is simply noted that the environment is a stiff surface, which many have modeled as a stiff spring.

Most of the efforts which treat the environment as a rigid body or a linear spring have command following (trajectory or force control) rather than interactive behavior as their primary goal, and, as such, make no effort to generate a comprehensive analysis of coupled stability. Oddly enough, even those who have been concerned with the design of impedance controllers, i.e., controllers for manipulators which are intended to interact with a fairly broad class of environments, have been slow to address the coupled stability behavior of closed loop designs. Until recently, greater emphasis was placed on performance aspects, such as the "closeness" of the closed loop behavior to that of a target model. Such an approach, however, is contrary to the set of design specifications proposed in Section 1.4.2, and may well result in undesirable behavior upon implementation. This sort of pitfall is illustrated by several of the examples in Chapters 4 and 8.

The impedance control literature reveals the following four levels of concern for coupled stability: the coupled stability properties of the:

- target dynamics.
- closed loop system, assuming that the robot is composed of rigid links with torque sources at the joints.
- closed loop system, assuming that additional dynamic effects (e.g., actuator bandwidth limitations or transmission dynamics) are important, but are included in the manipulator model.
- closed loop system, while accounting for modeling uncertainties, such as unmodeled dynamics or parametric errors.

The robust coupled stability property of a given set of target dynamics is an important justification for pursuing impedance control. For instance, Hogan has analyzed a manipulator which exhibits the following behavior [33]:

$$\mathbf{F} = \mathbf{K}(\mathbf{X} - \mathbf{X}_0) + \mathbf{B}(\mathbf{V})$$

where $\mathbf{X}$ and $\mathbf{V}$ are the manipulator position and velocity vectors, $\mathbf{X}_0$ is a vector of desired positions, $\mathbf{F}$ is a force imposed by the environment, $\mathbf{K}(\cdot)$ is the force-displacement relation, and $\mathbf{B}(\cdot)$ is the force-velocity relation. He has shown that, so long as $\mathbf{K}(\cdot)$ is the gradient of some potential function, and $\mathbf{B}(\cdot)$ is chosen so that this manipulator is capable of stably positioning an arbitrarily small mass, then the manipulator can interact stably with any dissipative environment describable in Hamiltonian terms. Extensions to this analysis can be found in Fasse [22]. Kazerooni [44] makes a similar argument for the case in which both the manipulator and the environment can be described as 2-n order linear systems (e.g., n-link robots, linearized) with positive definite inertia, damping, and stiffness matrices.

Apparently this sort of behavior is quite desirable, but can it be realized with feedback control and an actual robot? If the robot has rigid links and actuators which

act as pure torque sources, the answer is yes. Hogan has presented a controller which will, under these conditions, implement the target dynamics above (with an added inertia term), and even compensate for joint friction and gravity [35]. Takegaki and Arimoto also present a rigid-link, ideal-actuator robot controller which exhibits a strong coupled stability property [78].

In practice, however, these controllers do not necessarily succeed in implementing the robust coupled stability of the target dynamics (see, for instance, [92]). The reason is, of course, that real robots do not have such simple behavior as these methods assume. Actuator and transmission dynamics, joint and link flexibilities, and computational delays are but a few of the non-idealities. It turns out, however, that to account for these non-idealities, even if they are well characterized, can be quite difficult. For instance, in [44], Kazerooni performs a linearized analysis of an n-d.o.f. system with actuator bandwidth limitations. Each actuator is assumed to have a first-order rolloff. He is careful to point out that the closed loop system will be guaranteed to have the coupled stability property of the target dynamics only if the two have identical behavior at *all* frequencies, but that this cannot happen. The coupled stability properties of his designs are, in fact, rather poor—a detailed example is presented in Section 4.1.3.

In more recent work [45], Kazerooni has developed a method for guaranteeing the stability of closed loop designs of arbitrary complexity coupled to environments which are characterized by the magnitude of their input/output mapping (e.g., for linear systems, the singular values). This method is developed quite naturally for linear systems using the techniques of multivariable analysis; he has also shown that it can be extended via the Small Gain Theorem to apply to a broad range of nonlinear systems. This technique will be analyzed in some detail in Chapter 8; for the present purpose, it is noted that the important benefit of Kazerooni's method is that it provides an analysis for a very general set of systems and even (although he doesn't treat this) provides a framework suitable for treating modeling uncertainties; however, it has the

drawback that the magnitude (e.g., stiffness) of the environment must be bounded. This is inconvenient in, for instance, the case of a rigid surface.

Fasse [22] has developed a method for guaranteeing the stability of closed-loop systems coupled to passive environments (which may be nonlinear). The basis of his method is the selection of an appropriate Lyapunov function for the controlled system. An important result of this work is that the Lyapunov function, V, need not represent an actual energy so long as the following relation holds:

$$\dot{V} = -D + \mathbf{u}'\mathbf{y}$$

where D is a positive semi-definite function of state, and $\mathbf{u}'\mathbf{y}$ is the power delivered to the ports of the system. In the terminology of the last chapter, V is an internal energy function. Fasse shows several examples of linear systems for which such Lyapunov functions can be found by restricting the feedback gains in appropriate ways.

If such a Lyapunov function can be found, it constitutes a sufficient, but not necessary condition for the stability of the closed loop system coupled to arbitrary passive environments. In the next section, a necessary and sufficient condition will be found; thus, Fasse's result is a subset of the one found here.

Fasse also addresses the robustness of coupled stability to various types of modeling errors. In these analyses, he treats the standard implementations proposed by Hogan [35,92], but subject to non-idealities such as sensor gain errors, actuator dynamics, base dynamics, friction, and gravity. His is the only systematic examination of coupled stability robustness available to date.

This thesis will not add to his results, but will focus on the criteria which must be satisfied by any physical system, controlled or not, if it is to retain stability when interacting with arbitrary passive environments. However, so that the analysis is manageable, only linear plants will be considered.

The restriction to passive environments merits some further explanation. There are,

in fact, several reasons for focusing on passive environments. The first of these is that many environments of interest are passive. Rigid bodies, rigid surfaces, springs, and any collection of these are passive. But many other systems, less easily modeled and categorized, are also passive. For instance, a crucible of molten iron is passive. A second reason is that passive systems are mathematically well-defined and understood; this was the subject of the last chapter. Furthermore, it can be argued that the restriction of the environment to be a passive system is a much more natural choice than the restriction of the environment to be a system of certain input/output magnitude bounds. Certainly it is convenient for a manipulator to be able to interact with both inertial loads and rigid surfaces. Finally, the end justifies the means. The coupled stability analysis of the next section yields significant insight into the general topic of coupled dynamic systems, so that a number of extensions to the basic analysis are possible. For instance, in Section 3.5 it is shown that the stability criterion also applies to a broad class of active systems.

Before moving on to the coupled stability criterion, the term "coupled stability" will be more carefully defined. For this document, the definition given by Fasse [22] will be used:

Coupled Stability. A system is said to have the coupled stability property if:

1. The system is stable when isolated.
2. The system remains stable when coupled to any passive environment which is also stable when isolated.

The reader may be concerned about the common situation in which the manipulator and environment are in contact, but are not mechanically coupled. The relation to "contact stability" will be discussed in Section 3.6.

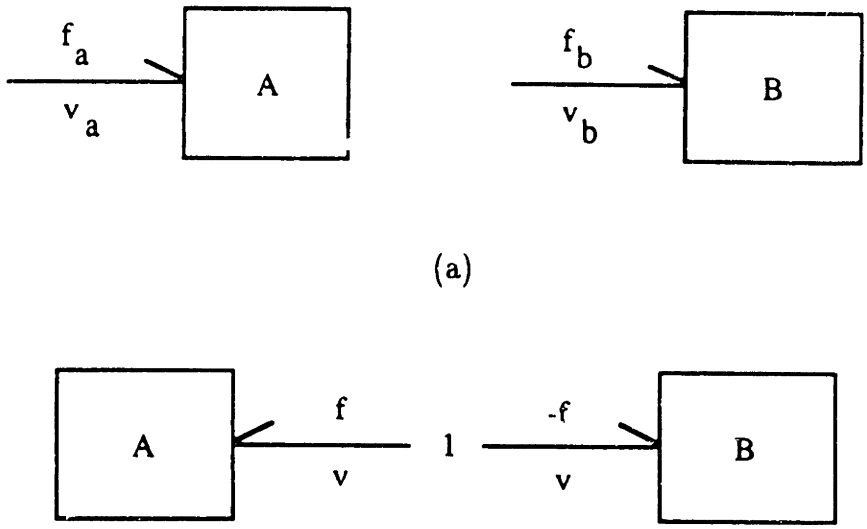

Figure 3.1: (a) Isolated systems, each with an available interaction port. (b) Coupled system. An alternative would have been to couple these systems with a 0 junction, in which case $f_a = f_b$ and $v_a = -v_b$; however, the stability result would be unaffected. See [68] for an explanation of 0 and 1 junctions.

3.2 1-Port Interaction

The purpose of this section is to derive a set of necessary and sufficient conditions for the stability of a linear, time-invariant system coupled at a single interaction port to an arbitrary linear, passive environment. The system in question (which may also be referred to as the plant) may be feedback controlled or not; the environment is assumed to be passive according to Passivity 6c of Section 2.4.1.

To commence, consider two linear, time-invariant systems, each with an open interaction port, as indicated in Figure 3.1(a). We are interested in the stability of the system that results when these are coupled at their interaction ports, i.e. the stability of the system shown in Figure 3.1(b).

Figure 3.1(b) can be rewritten in the block diagram form shown in Figure 3.2.

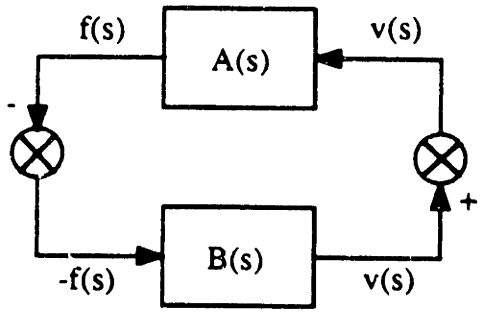

Figure 3.2: Block diagram representation of coupled system.

Implicit in this representation is the linear nature of A and B, i.e., both $A(s)$ and $B(s)$ may be expressed as the ratic of two rational polynomials in s. Furthermore, because each of these transfer functions relates power variables (f and v) measured at a given interaction port, they are *driving point impedances.*

The dynamics of the coupled system may be expressed as follows:

$$0 = \begin{bmatrix} -1 & B(s) \\ -A(s) & -1 \end{bmatrix} \begin{bmatrix} v(s) \\ f(s) \end{bmatrix} \tag{3.1}$$

The roots of the characteristic equation, $1 + A(s)B(s) = 0$, must fall within the closed LHP if a stable solution is to exist.

A Nyquist procedure can be used to study the pole locations. The basis of the Nyquist procedure is the use of the Principle of the Argument and of a well chosen mapping between complex planes. The Principle of the Argument states:

> The number of clockwise encirclements of the origin by a map through $G(s)$ of a clockwise contour in the s plane equals the number of zeros of $G(s)$ within the contour minus the number of poles of $G(s)$ within the contour.

A counterclockwise encirclement counts as minus one clockwise encirclements.

In classical control theory, the transfer function of interest is generally $1 + G(s)$ rather than $G(s)$. This is because, assuming unity feedback and $G(s) = N(s)/D(s)$ as

an open loop transfer function, $1 + G(s)$ is the quotient of the closed loop characteristic polynomial ($N(s) + D(s)$) and the open loop characteristic polynomial ($D(s)$). Rather than map through $1 + G(s)$, however, it is convenient to map through $G(s)$ and to investigate encirclements of the -1 point.

The Nyquist contour, which is shown in Figure 3.3, includes the entire right half plane, but excludes any poles on the imaginary axis. An example of a mapping into the $G(s)$ plane is also shown in this figure. This mapping is necessarily symmetric about the real axis as $G(s)$ is the ratio of two polynomials in s, $N(s)/D(s)$, the real part of which is an even function of ω on the $j\omega$ axis, and the imaginary part of which is an odd function. The behavior very near a pole on the $j\omega$ axis is determined by this pole (or poles, if they are not simple) and its residue; thus, an arc of 180° in the s plane, taken at an infinitesimal radius around an imaginary pole, and starting at $\pm 90°$, must result in an arc in the $G(s)$ plane that is some multiple of 180°, and is symmetric about the real axis. Similarly, the behavior along the arc of radius R, as $R \to \infty$, is determined by αs^n, where α is a constant, and n is the relative order of $G(s)$; thus, the mapping of this arc must also be symmetric about the real axis of the $G(s)$ plane.

Once the mapping through $G(s)$ has been generated, it is a simple matter to interpret the Principle of the Argument as follows:

> The number of clockwise encirclements of the -1 point by a mapping of the Nyquist contour through $G(s)$ equals the number of unstable closed loop poles (i.e., RHP roots of $N(s) + D(s)$) minus the number of unstable open loop poles (i.e., RHP roots of $D(s)$).

Figure 3.2 indicates that $A(s)B(s)$ plays the role of the open loop transfer function; thus, we should replace $G(s)$ with $A(s)B(s)$ in the statement above. Some further assumptions about $A(s)$ and $B(s)$ are now necessary. First $A(s)$ is identified with the plant, and $B(s)$ with the environment.

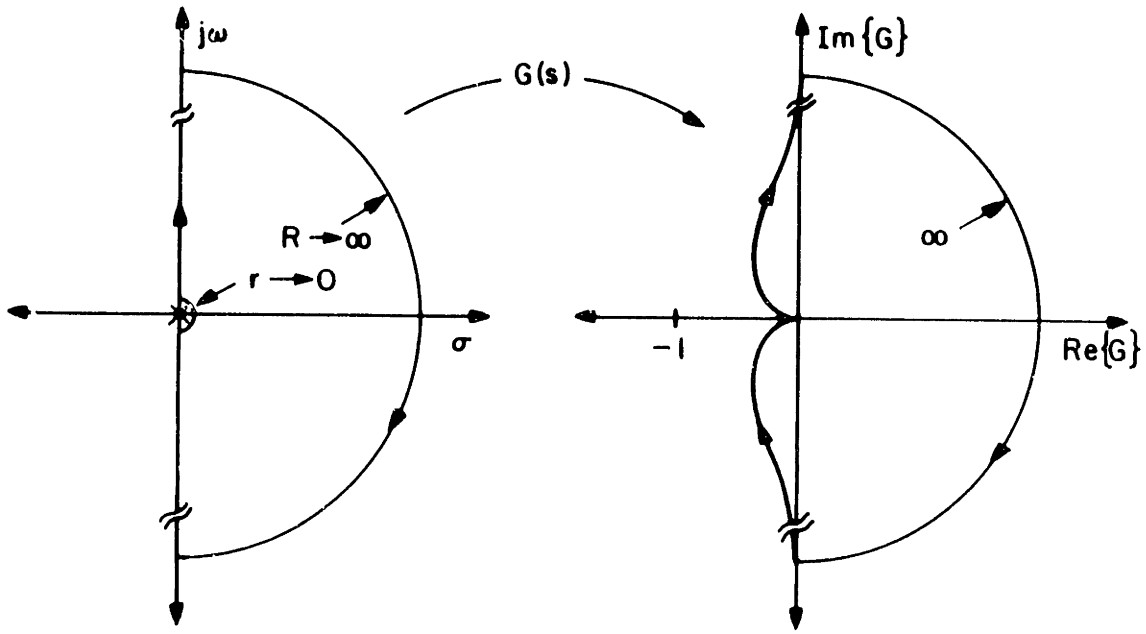

Figure 3.3: An example mapping of the Nyquist contour.

It will also be assumed that the input and output of $A(s)$ have been chosen appropriately so that a state space realization exits; i.e., so that F_a, G_a, H_a, and J_a exist which satisfy:

$$\begin{aligned} \dot{\mathbf{x}} &= F_a\mathbf{x} + G_a v \\ f &= H_a\mathbf{x} + J_a v \end{aligned} \tag{3.2}$$

and:

$$A(s) = J_a + H_a(sI - F_a)^{-1}G_a \tag{3.3}$$

This is equivalent to the requirement that $A(\infty) < \infty$. It is further assumed that the eigenvalues of F_a lie in the closed left half plane. These assumptions, taken together, guarantee that the poles of $A(s)$ lie in the closed left half plane, and are indicative of a stable system.

The environment, which was previously assumed to be passive, must satisfy the following criteria (from Passivity 6c, Section 2.4.1):

- $B(s)$ has no RHP poles.

- $B(s)$ has a Nyquist plot which lies wholly within the closed right half plane.

Otherwise stated (see Section 2.4.1), there is no restriction on the magnitude of $B(s)$, but the phase of $B(s)$ must lie between $-90°$ and $+90°$.

Recall that we are interested in guaranteeing the stability of the plant when coupled to *any* linear, passive environment. If $B(s)$ is viewed as a compensator which affects the magnitude and phase of $A(s)$, then the class of all such compensators will include those capable of altering the phase by $\pm 90°$, and changing the magnitude by a factor of 0 to ∞. Therefore, it is clear that, in order to interact with *all* passive environments, $A(s)$ must have a phase margin of $\pm 90°$, an upward gain margin of ∞, and a downward gain margin of 0.

The effect of these restricitions is to limit the mappings of the Nyquist contour through $A(s)$ to lie in the closed right half plane. For any path that enters the left half plane in a clockwise fashion, no more than 90° additional lead or lag will be necessary to cause the path to cross the negative real axis, and, once crossing the axis, some gain between 0 and ∞ will cause the path to encircle the -1 point. An example of this is illustrated in Figure 3.4(a); the Nyquist plots shown in Figures 3.4(b) and (c) are clearly inadmissable for the same reason. The Nyquist plots shown in Figures 3.4(d) and (e), however, enter the left half plane in a counterclockwise fashion, and will not be reshaped by any acceptable $B(s)$ to yield a clockwise encirclement of the origin. It may be shown, however (see Appendix A), that any Nyquist plot which contains a counterclockwise loop is indicative of a RHP pole, violating one of our previous assumptions.

Thus, the mapping through $A(s)$, like that through $B(s)$, is restricted to lie wholly within the closed right half plane. This, taken together with the stability of $A(s)$ in isolation, is a necessary and sufficient condition to ensure coupled stability. It is also,

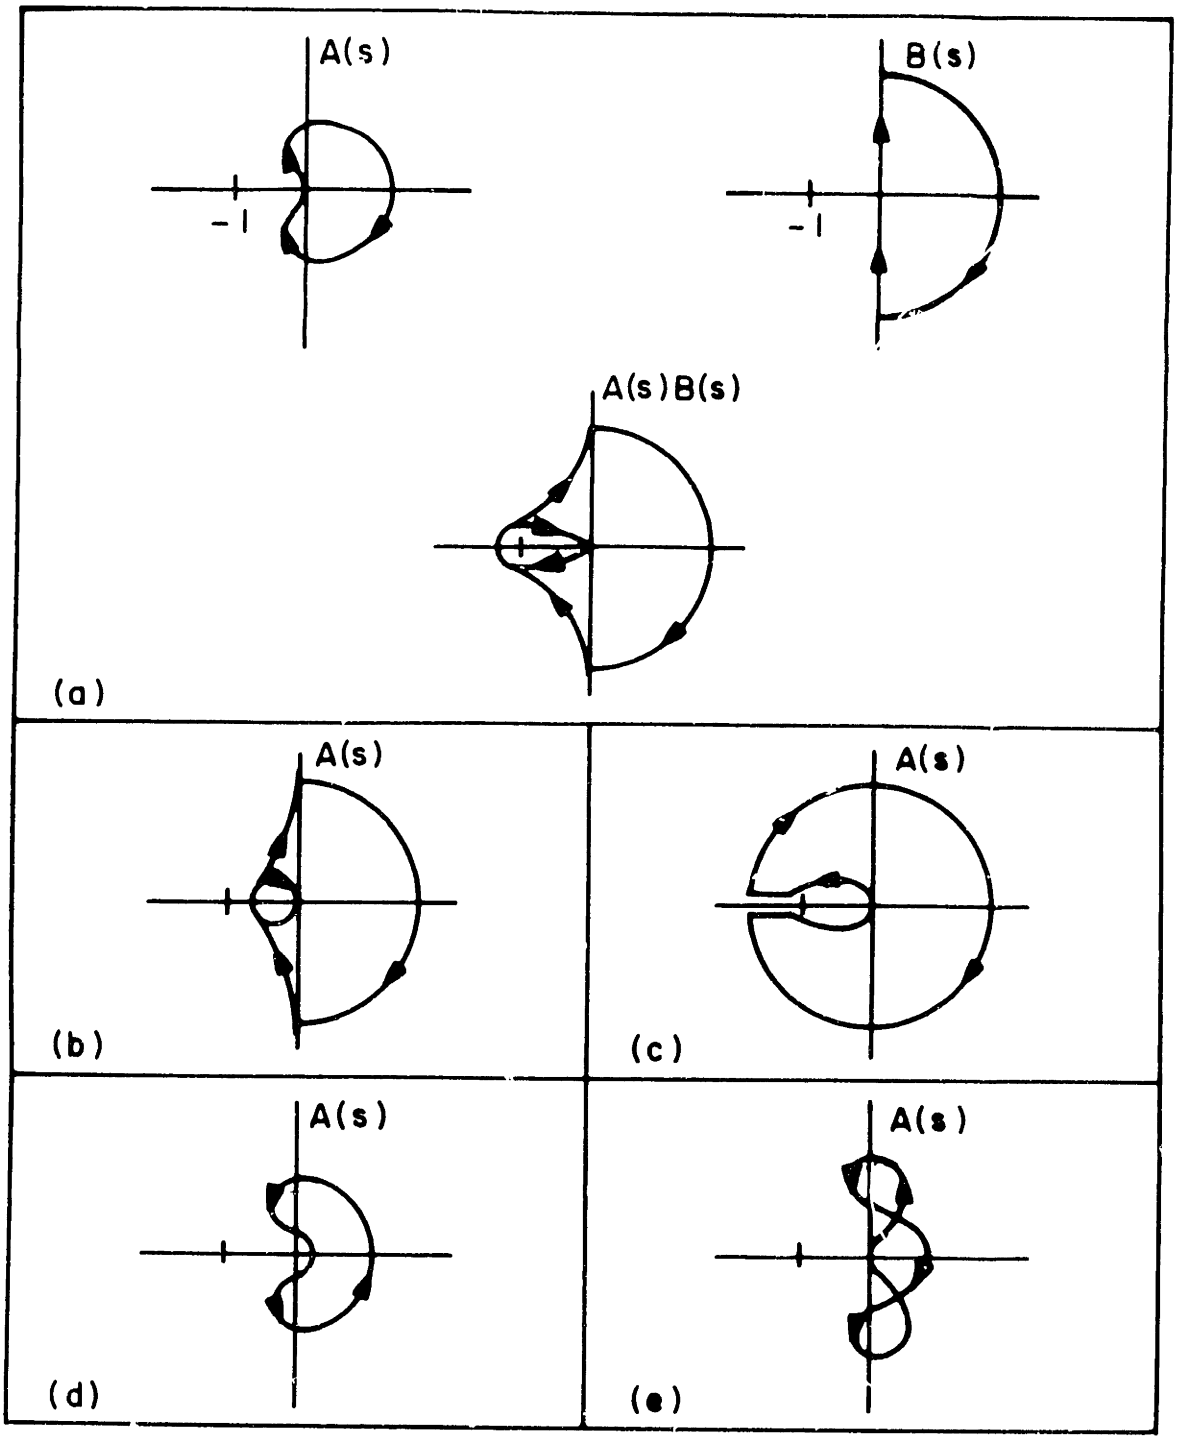

Figure 3.4: (a) Nyquist plots of a plant and a passive environment which, when coupled, lead to instability. (b),(c) Nyquist plots which fail to meet the coupled stability criterion. (d),(e) Nyquist plots which apparently cannot lead to coupled instability, but are indicative of RHP plant poles.

according to Passivity 6c, necessary and sufficient for $A(s)$ to represent the driving point impedance of a passive system. Note that this does not prohibit the plant, A, from being feedback-controlled; examples of feedback-controlled systems which obey this criterion are presented in the next chapter. The result of this analysis may be summarized as follows:

Coupled Stability 1 (1-Port). A necessary and sufficient condition to ensure the stability of a LTI, stable plant coupled at a single port to a LTI, stable, passive environment, is that the driving point impedance of the plant be positive real; or, equivalently, that the plant be passive.

3.3 n-Port Interaction

The purpose of this section is to derive a necessary and sufficient condition for the stability of a linear, time-invariant system coupled at n interaction ports to a linear, passive environment. Both the plant and environment will be assumed to be strictly proper impedances or admittances. This may seem restrictive, but one can always take a physical system, active or not, and, if it is to be described as an impedance, add an arbitrarily stiff spring to the interaction port, thereby making it strictly proper. The same can be done by adding an arbitrarily small mass to the interaction port of an admittance. It can also be argued that Nature takes care of this problem for us—that any physical system must be causal, and therefore, if linear, strictly proper. Thus, the restriction to strictly proper systems does not weaken the proof.

Consider two linear, time-invariant n-ports, A and B:

$$\begin{aligned} \dot{\mathbf{x}}_a &= A_a \mathbf{x}_a + L_a \mathbf{e}_a \\ \mathbf{f}_a &= C_a \mathbf{x}_a \end{aligned}$$

and,

$$\begin{aligned} \dot{\mathbf{x}}_b &= A_b\mathbf{x}_b + L_b\mathbf{f}_b \\ \mathbf{e}_b &= C_b\mathbf{x}_b \end{aligned}$$

where e and f represent power duals along the interaction ports of the two systems, and $\mathbf{e}_a$, $\mathbf{e}_b$, $\mathbf{f}_a$, and $\mathbf{f}_b$ are all $n \times 1$ vectors. A_a is a $k \times k$ matrix, and A_b is an $m \times m$ matrix; k and m are not necessarily equal. Now consider the following coupling of A and B:

$$\begin{aligned} \mathbf{e}_a &= -\mathbf{e}_b \\ \mathbf{f}_a &= \mathbf{f}_b \end{aligned}$$

This may be thought of as coupling A and B with a 1 junction.

The dynamics of the coupled system are represented by the following state equations:

$$\begin{bmatrix} \dot{\mathbf{x}}_a \\ \dot{\mathbf{x}}_b \end{bmatrix} = \begin{bmatrix} A_a & L_aC_b \\ -L_bC_a & A_b \end{bmatrix} \begin{bmatrix} \mathbf{x}_a \\ \mathbf{x}_b \end{bmatrix} \tag{3.4}$$

which will, for notational convenience, sometimes be referred to as:

$$\dot{\mathbf{x}}_c = A_c\mathbf{x}_c \tag{3.5}$$

Necessary and sufficient conditions for the stability of the coupled system are the existence of PD P and PSD Q that satisfy:

$$P_cA_c + A_c'P_c = -Q_c \tag{3.6}$$

which is simply the Lyapunov stability equation. For the remainder of the proof, P_c will be expressed in the following partitioned form:

$$P_c = \begin{bmatrix} P_1 & R \\ R' & P_2 \end{bmatrix} \tag{3.7}$$

where P_1 is a $k \times k$ matrix, P_2 is an $m \times m$ matrix, and R is a $k \times m$ matrix.

Some further assumptions about A and B are now necessary. The first is that A is stable in isolation. A PD matrix P_a and PSD matrix Q_a exist which satisfy the Lyapunov stability equation:

$$P_a A_a + A_a' P_a = -Q_a \tag{3.8}$$

The second assumption is that system B represents a passive environment, in which case Passivity 7b applies:

$$P_b A_b + A_b' P_b = -Q_b \tag{3.9}$$

$$P_b L_b = C_b' \tag{3.10}$$

With these preliminaries completed, the proof will proceed in two parts, sufficiency and necessity.

Sufficiency Proof

Suppose that $P_1 = P_a$, $P_2 = P_b$, and $R = 0$, then:

$$\begin{aligned} P_c A_c + A_c' P_c = & \begin{bmatrix} P_a & 0 \\ 0 & P_b \end{bmatrix} \begin{bmatrix} A_a & L_a C_b \\ -L_b C_a & A_b \end{bmatrix} \\ & + \begin{bmatrix} A_a' & -C_a' L_b' \\ C_b' L_a' & A_b' \end{bmatrix} \begin{bmatrix} P_a & 0 \\ 0 & P_b \end{bmatrix} \end{aligned} \tag{3.11}$$

After multiplying, adding, and using equations 3.8 and 3.9:

$$P_c A_c + A_c' P_c = \begin{bmatrix} -Q_a & P_a L_a C_b - C_a' L_b' P_b \\ C_b' L_a' P_a - P_b L_b C_a & -Q_b \end{bmatrix} \tag{3.12}$$

Equation 3.10 may be used to simplify this matrix:

$$P_c A_c + A_c' P_c = \begin{bmatrix} -Q_a & (P_a L_a - C_a') C_b \\ C_b' (L_a' P_a - C_a) & -Q_b \end{bmatrix} \tag{3.13}$$

The condition $P_a L_a = C_a'$, which, by Passivity 7b, is sufficient to guarantee that system A is positive real, is clearly sufficient to establish the stability of the coupled system:

$$P_c A_c + A_c' P_c = -Q_c = \begin{bmatrix} -Q_a & 0 \\ 0 & -Q_b \end{bmatrix} \tag{3.14}$$

Necessity Proof

Having demonstrated sufficiency, we now intend to show that $P_a L_a = C_a'$ is a necessary condition for the existence of positive definite P_c and positive semidefinite Q_c. To do this, we begin by expressing $P_c A_c + A_c' P_c$ in terms of the partitioned P_c (equation 3.7):

$$P_c A_c + A_c' P_c = -\begin{bmatrix} RL_bC_a + C_a'L_b'R' - (P_1A_a + A_a'P_1) & C_a'L_b'P_2 - P_1L_aC_b - (A_a'R + RA_b) \\ -C_b'L_a'P_1 + P_2L_bC_a - (R'A_a + A_b'R') & -C_b'L_a'R - R'L_aC_b - (P_2A_b + A_b'P_2) \end{bmatrix} \tag{3.15}$$

Because all of the submatrices along the main diagonal of a PSD matrix must be PSD, the diagonal terms in equation 3.15 must be PSD. These terms are sums, and, in general, the individual components of the sum need not be PSD. This brings us to a critical point in the proof. We are interested in guaranteeing the stability of system A when it may be coupled to *any* passive environment. Because a passive system must satisfy Passivity 7b, all we need do to construct the state equations of a passive system is select A_b to satisfy equation 3.9, and then select *any* L_b and find the associated C_b with equation 3.10 (or select C_b and find L_b). The point is that either L_b or C_b may be considered to be completely arbitrary.

This brings us back to the submatrices which must be PSD. Although, in general, the components in the sum need not be PSD, if L_b is taken to be completely arbitrary, then the term $RL_bC_a' + C_a'L_b'R'$ found in the upper diagonal term of equation 3.15 must be PSD independent of L_b (otherwise, a simple scaling of any given L_b would be sufficient to force the entire term to be non-PSD). This will occur only if R has a particular form:

$$R = C_a'L_b'J_1 \tag{3.16}$$

where J_1 is a PSD matrix. Then:

$$C_a' L_b' R' + R L_b C_a = 2 C_a' L_b' J_1 L_b C_a \tag{3.17}$$

Thus, although L_b is completely arbitrary, the left side of equation 3.17 is always PSD. Similar arguments may be made for the lower diagonal term in equation 3.15. The term $-C_b' L_a' R - R' L_a C_b$, which may be rewritten using equation 3.10 as $-P_b L_b L_a' R - R' L_a L_b' P_b$, must be PSD. This will occur only if R has the following form:

$$R = -J_2 L_a L_b' P_b \tag{3.18}$$

where J_2 is PSD. Then:

$$-P_b L_b L_a' R - R' L_a L_b' P_b = 2 P_b L_b L_a' J_2 L_a L_b' P_b \tag{3.19}$$

Equations 3.16 and 3.18 may be equated to solve for J_1 and J_2:

$$C_a' L_b' J_1 = -J_2 L_a L_b' P_b \tag{3.20}$$

A trivial solution to this equation is $J_1 = 0$ and $J_2 = 0$, however, it is necessary to determine if there are any others. Pre-multiplying by L_a', and post-multiplying by L_b:

$$(L_a' C_a')(L_b' J_1 L_b) = -(L_a' J_2 L_a)(L_b' P_b L_b) \tag{3.21}$$

Because L_b is arbitrary, it is clear that a nontrivial solution to equation 3.21 can be found in general only if L_b can be removed from the equation. This will occur only if we select $J_1 = \Lambda P_b$, where Λ is some PSD diagonal matrix. This done, equation 3.21 reduces to:

$$L_a' C_a' = -\Lambda^{-1}(L_a' J_2 L_a) \tag{3.22}$$

Equation 3.22 will be satisfied only if the matrix $L_a' C_a'$ has negative or zero eigenvalues. However, we already know from the sufficiency proof that one solution is:

$$L_a' C_a' = L_a' P_a L_a$$

This solution clearly requires that $L_a'C_a'$ have positive or zero eigenvalues. The only reconciliation, other than the trivial solution of $L_a = 0$ and $C_a = 0$, is that $J_1 = 0$ and $J_2 = 0$. R, therefore, must also be a zero matrix.

Now equation 3.15 may be written in a simplified form:

$$P_c A_c + A_c' P_c = -\begin{bmatrix} -(P_1 A_a + A_a' P_1) & C_a' L_b' P_2 - P_1 L_a C_b \\ P_2 L_b C_a - C_b' L_a' P_1 & -(P_2 A_b + A_b' P_2) \end{bmatrix} \tag{3.23}$$

Because the right hand side of this equation must be the negative of a PSD matrix whose submatrices along the main diagonal are PSD, the following equations must hold:

$$P_1 A_a + A_a' P_1 = -Q_1 = -V_1' V_1 \tag{3.24}$$

$$P_2 A_b + A_b' P_2 = -Q_2 = -V_2' V_2 \tag{3.25}$$

where V_1 and V_2 are unrestricted except in dimension, and Q_1 and Q_2 are PSD. Furthermore, one of the following two equations must hold:

$$\begin{bmatrix} Q_1 & P_1 L_a C_b - C_a' L_b' P_2 \\ C_b' L_a' P_1 - P_2 L_b C_a & Q_2 \end{bmatrix} = \begin{bmatrix} V_1' \\ V_2' \end{bmatrix} \begin{bmatrix} V_1 & V_2 \end{bmatrix} \tag{3.26}$$

or,

$$\begin{bmatrix} Q_1 & P_1 L_a C_b - C_a' L_b' P_2 \\ C_b' L_a' P_1 - P_2 L_b C_a & Q_2 \end{bmatrix} = \begin{bmatrix} V_1' & 0 \\ 0 & V_2' \end{bmatrix} \begin{bmatrix} V_1 & 0 \\ 0 & V_2 \end{bmatrix} \tag{3.27}$$

Equation 3.26 requires the following:

$$P_1 L_a L_b' P_b - C_a' L_b' P_2 = V_1' V_2$$

As with equation 3.21, L_b must be removed from this equation if it is to have a general solution. This would require $P_2 = \Lambda P_b$, where Λ is a PD diagonal matrix, and $V_2 = \Gamma L_b' P_2$, where Γ is some matrix of appropriate dimension. Substituting this expression into equation 3.25 produces the following result, which clearly has no solution for arbitrary L_b:

$$P_2 A_b + A_b' P_2 = -P_2 L_b \Gamma' \Gamma L_b' P_2$$

The second option requires the following:

$$P_1 L_a L_b' P_b - C_a' L_b' P_2 = 0$$

Again, $P_2 = \Lambda P_b$ is required to remove L_b from this equation. This done, the following equation must hold:

$$P_1 L_a - \Lambda C_a' = 0 \tag{3.28}$$

If we define $P_a = \Lambda^{-1} P_1$, and $Q_a = \Lambda^{-1} Q_1$, then we may write the necessary conditions for coupled stability as follows:

$$\begin{aligned} P_a A_a + A_a' P_a &= -Q_a \\ P_a L_a &= C_a' \end{aligned} \tag{3.29}$$

These equations constitute the necessary and sufficient conditions for system A to be passive, according to Passivity 7b. In summary, the following result has been reached:

Coupled Stability 2 (n-Port). A necessary and sufficient condition to ensure the stability of a LTI, stable plant coupled at n ports to a LTI, stable, passive environment, is that the plant be passive according to Passivity 7b.

3.4 The "Worst" Environment

The purpose of this section is to describe an alternative method of analyzing coupled stability in the case of 1-port interaction. It will be demonstrated that, if a manipulator is stable when coupled to a particular, but restricted, class of passive environments, the "worst" environments, then it is stable when coupled to all passive environments; and furthermore, that this restricted class of environments is easily parametrized so that coupled stability can be analyzed by the generation of two root loci.

If we return briefly to the intuitive treatment of passive environments as a class of compensators which can add up to $\pm 90°$ phase lag and any magnitude scaling to the open loop transfer function (the plant's driving point impedance), then it would seem reasonable that the most destabilizing environments would be those adding the full $\pm 90°$ phase lag and some appropriate magnitude scaling. It is lossless environments [4] which have $\pm 90°$ phase lag. Lossless environments are composed purely of storage elements such as springs or masses, with no dissipation. In fact, the simplest lossless environments are simply masses and springs which have driving point impedances and admittances of the form a/s or as.

The intuitive approach, then, would suggest that if a plant is stable when coupled to any mass or any spring, it is stable when coupled to any passive environment. In fact, this result can be formally established with the use of Passivity 8 from Section 2.5. Passivity 8 requires that all branches of the root loci of $sZ(s)$ and $Z(s)/s$ lie in the closed LHP if $Z(s)$ is to represent a passive system. Because Coupled Stability 1 (Section 3.2) requires that the driving point impedance, $A(s)$, of the plant be that of a passive system, it is evident that $sA(s)$ and $A(s)/s$ must also have loci which lie in the closed LHP.

Now consider the case in which $A(s)$ is a driving point admittance. The coupling of $A(s)$ to a spring of stiffness k_e is shown in Figure 3.5. Evidently, the stability of the coupled system can be determined for all stiffnesses k_e simply by generating a root locus for the open loop transfer function, $A(s)/s$. Similarly, if the plant is coupled to a mass, stability can be determined for all values of the mass simply by generating a root locus of $sA(s)$. These same results hold if $A(s)$ is a driving point impedance, except that stability when coupled to springs requires a root locus of $sA(s)$, and when coupled to masses, a locus of $A(s)/s$. In summary:

Coupled Stability 3 (1-Port). A LTI plant interacting with its environment at a

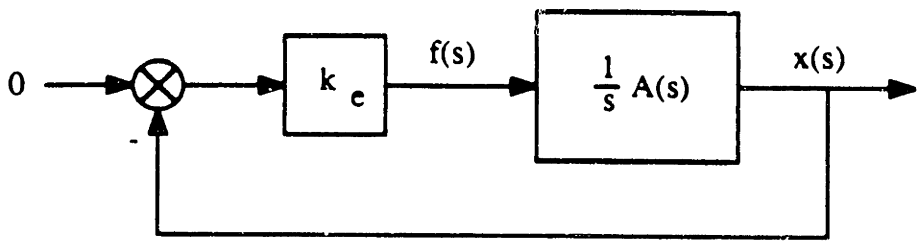

Figure 3.5: The plant A coupled to a spring.

single port will have the coupled stability property iff it is stable when coupled to all springs and all masses. Equivalently, the root loci of $A(s)/s$ and $sA(s)$, where $A(s)$ is the plant's driving point impedance, must lie completely within the closed LHP to guarantee coupled stability.

This method will be used in some of the examples of the next chapter. One important caveat: the term "worst" environment is something of a misnomer; any environment which leads to instability need not be outdone, and in most cases, if a lossless environment leads to instability, many dissipative environments will as well. Furthermore, finding those springs and masses which lead to instability says nothing about any other environments which will lead to instability.

3.5 Extensions to the Coupled Stability Analysis

The coupled stability criteria which have been derived in this chapter would not be nearly so powerful as they actually are if the plant actually had to be passive. In fact, it does not. Figure 3.6 shows how an active term may be added without jeopardizing the stability of the coupled system. Clearly, so long as u does not depend on the states of the plant or the environment, it does not affect the system stability.

Some further insight can be gained by realizing that the driving point impedance

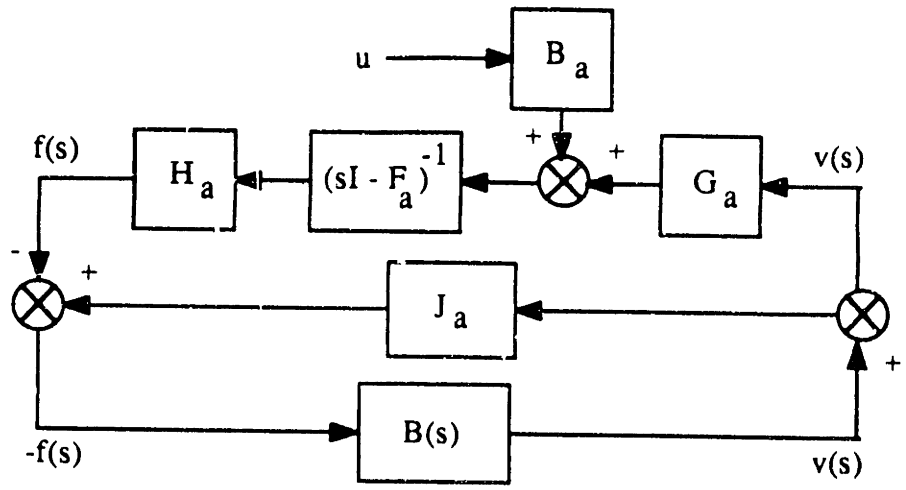

Figure 3.6: The addition of an active term to the coupled system (the plant is described by equations 3.2, with the addition of $B_a u$).

treated so far is simply a Thevenin or Norton equivalent impedance seen at that port, and that the presence of a general Thevenin or Norton equivalent source does not change the stability result[1]. In fact, the source term can itself be the output of some dynamic system, so long as the behavior of this system does not depend on the states of the plant or environment. Finally, it is evident that the environment may be active as well, so long as it has the equivalent impedance of a passive system.

Another important extension to the basic result is the inclusion of nonlinear passive environments. If the environment is nonlinear, passive, and stable, then there must exist some Lyapunov function which is also an internal energy function. Because a linear, passive plant also has a Lyapunov function which will serve as an internal energy function, then the two may simply be summed, following the arguments of Fasse [22], to demonstrate the stability of the coupled system. This provides a sufficient condition for coupled stability; necessity follows from the linear analyses, as linear systems are a subset of nonlinear. Thus:

[1] Anderson and Spong [5] have also used Thevenin and Norton equivalents to describe a robot and its environment.

Coupled Stability 4 (n-Port). A LTI n-port plant will be stable when coupled to an arbitrary passive environment iff it has the driving point impedance of a passive system.

3.6 Contact Stability

The relation of contact stability to coupled stability can now be understood in a straightforward manner. A system which can make or break contact with its environment is essentially operating between two modes—uncoupled and coupled. If it is stable in both modes, then it must be stable in general, as the contact cannot create energy. Thus, coupled stability is a sufficient condition for contact stability.

Now suppose that the plant is unstable when coupled to some environment. If it is not secured to the environment, then it must eventually lose contact. At this point the plant will return to stable behavior and begin to dissipate energy. However, assuming that its equilibrium configuration brings it once again into contact with the surface, limit cycles will begin. These limit cycles cannot die away because of the coupled instability, but they do not necessarily have to increase without bound, either. At some amplitude, the energy loss when uncoupled may balance the energy gain when coupled.

Thus, bounded-input bounded-output stability can sometimes be established for contact stability when the coupled stability condition indicates instability. This cannot necessarily be resolved by requiring asymptotic stability, either, as our statement of coupled stability allows lossless systems, which are not asymptotically stable.

One resolution, which has been used in this thesis, is simply to disallow sustained oscillations which include a change in mode. If such behavior is considered unstable, coupled instability is also a sufficient condition for contact instability.

3.7 Summary

The most important result of this chapter is a necessary and sufficient condition for coupled stability:

> A LTI n-port plant will be stable when coupled to an arbitrary passive environment iff it has the driving point impedance of a passive system.

It has also been shown that the plant need not be passive. Arbitrary state-independent source terms can be introduced without affecting the stability result. This is the reason for the structure presented in Section 1.4.2, which separates state-dependent and state-independent source terms.

An alternative criterion for 1-port coupled stability, based on the generation of two root loci, one for coupling to springs and one for coupling to masses, has also been presented. Finally, it has been shown that coupled stability and instability are sufficient conditions for contact stability and instability, so long as sustained limit cycles which include loss of contact with the environment are considered unstable.

Chapter 4

Examples of the Coupled Stability Criterion

The purpose of this chapter is to provide a number of examples that illustrate the coupled stability criterion and its application to the analysis of closed loop designs. The examples will focus on the consequences of failure to satisfy the criterion. A number of the examples are of controllers designed for a two-link manipulator. These controllers were implemented on an actual manipulator, and interactive behavior was experimentally examined. The experiments and results are described in the next chapter.

4.1 Impedance Controllers

This section examines the coupled stability properties of various impedance control implementations that have been described in the literature. Of special interest is the linearized control presented in Section 4.1.2, because an experimental investigation of its properties is described in the next chapter.

4.1.1 Feedback Linearized Impedance Control

In [35], Hogan presents an impedance controller for a manipulator with the following rigid body description:

$$I(\theta)\frac{d\omega}{dt} + \mathbf{C}(\theta,\omega) + \mathbf{V}(\omega) + \mathbf{S}(\theta) = \mathbf{U} + J'(\theta)\mathbf{F}$$

where θ and ω are the generalized joint angles and velocities, respectively; $I(\theta)$ is the configuration-dependent inertia tensor; $\mathbf{C}(\theta,\omega)$ is a vector of inertial coupling terms; $\mathbf{V}(\omega)$ includes any velocity-dependent forces such as friction; $\mathbf{S}(\theta)$ includes any static configuration-dependent forces such as gravitational forces; $\mathbf{U}$ is a vector of actuator torques; $J(\theta)$ is the configuration-dependent jacobian relating joint velocities to endpoint velocities; and $\mathbf{F}$ is a vector of endpoint forces imposed by the environment (also to be called interaction port forces).

The desired endpoint behavior of the manipulator is given by:

$$M\frac{d\mathbf{V}}{dt} - B[\mathbf{V}_0 - \mathbf{V}] - K[\mathbf{X}_0 - \mathbf{X}] = \mathbf{F}$$

where $\mathbf{X}$ and $\mathbf{V}$ represent the position and velocity of the endpoint, and $\mathbf{X}_0$ and $\mathbf{V}_0$ represent the commanded position and velocity. This behavior is clearly that of a passive, linear system for PD M, B, and K.

Simply by equating the endpoint accelerations of the target and actual models, a control (expression for $\mathbf{U}$ in terms of θ, ω, and $\mathbf{F}$) can be derived which will generate the target behavior exactly! Thus, even though this manipulator is nonlinear, it should be possible to satisfy the coupled stability criterion.

4.1.2 A Linearized Impedance Controller for a Two-Link Manipulator

The geometry of the two-link manipulator used in the experiments is shown in Figure 4.1. Nonlinear equations describing the dynamic behavior of this manipulator are

derived in [23]. They are:

$$I(\theta)\ddot{\theta} = \beta\omega + \mathbf{C}(\theta,\omega) + \mathbf{U} + J'(\theta)\mathbf{F}$$

where:

$$I(\theta) = \begin{bmatrix} J_1 + l_1^2 m_2 & l_1 h_2 m_2 \cos(\theta_2 - \theta_1) \\ l_1 h_2 m_2 \cos(\theta_2 - \theta_1) & J_2 \end{bmatrix}$$

$$\beta = \begin{bmatrix} -(B_1 + B_c) & B_c \\ B_c & -(B_2 + B_c) \end{bmatrix}$$

$$\mathbf{C}(\theta,\omega) = \begin{bmatrix} l_1 h_2 m_2 \omega_2^2 \sin(\theta_2 - \theta_1) \\ -l_1 h_2 m_2 \omega_1^2 \sin(\theta_2 - \theta_1) \end{bmatrix}$$

$$J(\theta) = \begin{bmatrix} -l_1 \sin\theta_1 & -l_2 \sin\theta_2 \\ l_1 \cos\theta_1 & l_2 \cos\theta_2 \end{bmatrix}$$

θ_1 = absolute angle of link 1 (rad)
θ_2 = absolute angle of link 2 (rad)
ω_1 = angular velocity of link 1 (rad/sec)
ω_2 = angular velocity of link 2 (rad/sec)
l_1 = length of link 1 (m)
l_2 = length of link 2 (m)
h_2 = length from pivot to c.g. of link 2 (m)
m_2 = mass of link 2 (kg)
J_1 = moment of inertia of link 1 (about joint) (kg-m^2)
J_2 = moment of inertia of link 2 (about joint) (kg-m^2)
B_1 = viscous damping, link 1 to ground (kg-m^2/s)
B_2 = viscous damping, link 2 to ground (kg-m^2/s)
B_1 = viscous damping between links (kg-m^2/s)

Linearization

Because the experiments described in the next chapter were performed only in the vicinity of certain zero-velocity operating points, these equations of motion were linearized about those operating points. The reader may question the need for this if the

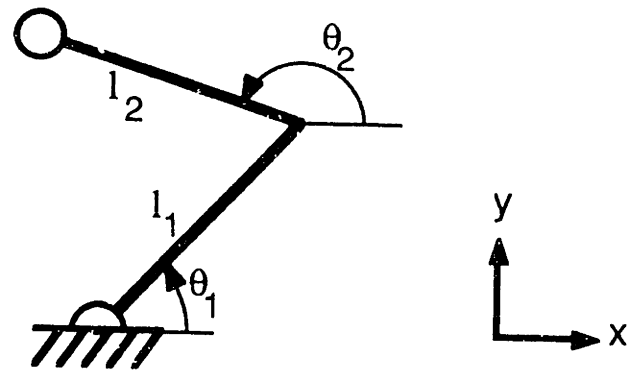

Figure 4.1: Two-link manipulator

impedance control law described in the last section could be used to linearize the end-point behavior. That control law, however, is computationally expensive, and was not judged to be necessary for the small motion experiments. The linearized manipulator description is:

$$I(\Delta\theta_0)\delta\ddot{\theta} = \beta\delta\dot{\theta} + \mathbf{U} + J'(\theta_0)\mathbf{F}$$

where $\delta\theta = \theta - \theta_0$, $\omega_0 = 0$, and $\Delta\theta_0 = \theta_{2_0} - \theta_{1_0}$. Dropping the δ, this can be written in standard state variable form as follows:

$$\begin{bmatrix} \dot{\theta} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & I^{-1}\beta \end{bmatrix} \begin{bmatrix} \theta \\ \omega \end{bmatrix} + \begin{bmatrix} 0 \\ I^{-1} \end{bmatrix} \mathbf{U} + \begin{bmatrix} 0 \\ I^{-1}J' \end{bmatrix} \mathbf{F} \tag{4.1}$$

or:

$$\dot{\Theta} = A\Theta + B\mathbf{U} + L\mathbf{F} \tag{4.2}$$

Control Derivation

For what will be called the "simple impedance controller" to be designed and implemented here, no attempt will be made to measure and use force information; only joint

angles and velocities will be used. The derivation of this control law follows that given by Hogan in [35]. The dynamics of the manipulator are ignored, and the principle of virtual work is used to relate the actuator torques to the desired endpoint force with the following result:

$$\mathbf{U} = J'\mathbf{F}$$

The desired endpoint behavior is given by:

$$\mathbf{F} = -K_x\mathbf{X} - B_x\mathbf{V}$$

where the commanded position and velocity are zero. Then, for small displacements:

$$\mathbf{U} = -J'K_xJ\theta - J'B_xJ\omega$$

It is a simple matter to provide additional compensation for the viscous losses in the manipulator:

$$\mathbf{U} = -J'K_xJ\theta - (J'B_xJ + \beta)\omega$$

or,

$$\begin{aligned} \mathbf{U} &= -G\boldsymbol{\Theta} \\ G &= \begin{bmatrix} J'K_xJ & J'B_xJ + \beta \end{bmatrix} \end{aligned}$$

The closed loop state equations are then:

$$\begin{aligned} \dot{\boldsymbol{\Theta}} &= (A - BG)\boldsymbol{\Theta} + L\mathbf{F} \\ \mathbf{X} &= \begin{bmatrix} J & 0 \end{bmatrix}\boldsymbol{\Theta} \end{aligned}$$

Figure 4.2 is a Nyquist plot of the closed loop driving point admittance[1] of this system for typical values. The closed loop system clearly appears passive.

[1] Actually, because the manipulator has two degrees of freedom, an n-port passivity criterion should be used. This will be done in the next chapter; for the present, however, only the diagonal terms will be considered. These can be analyzed with the simpler 1-port techniques, and are representative (for the examples considered here) of all interesting features.

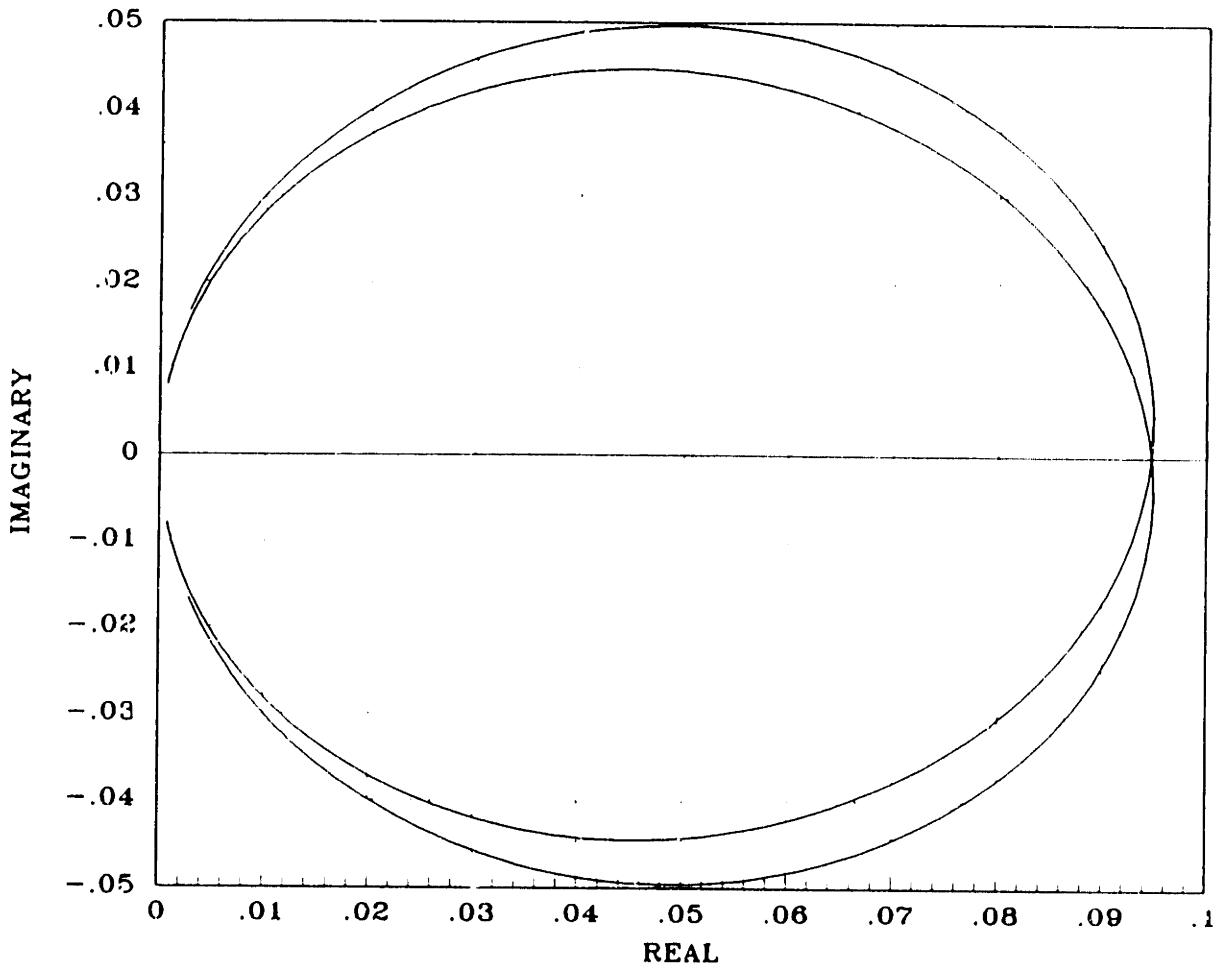

$$\theta_{1_0} = 30°, \quad \theta_{2_0} = 150°$$

$$I = \begin{bmatrix} 0.1954 & -0.0560 \\ -0.0560 & 0.0932 \end{bmatrix}, \quad \beta = \begin{bmatrix} -0.01070 & 0.01066 \\ 0.01066 & -0.01070 \end{bmatrix}$$

$$J = \begin{bmatrix} -0.183 & -0.168 \\ 0.316 & -0.291 \end{bmatrix}$$

$$K_x = \begin{bmatrix} 60 & 0 \\ 0 & 60 \end{bmatrix}, \quad B_x = \begin{bmatrix} 10 & 0 \\ 0 & 10 \end{bmatrix}$$

Figure 4.2: Simple impedance controller. Nyquist plot of $Y_{xx}(s)$. Both axes have units of sec/kg.

4.1.3 The Addition of First-Order Decoupled Actuator Dynamics

Kazerooni [44] developed a methodology for the design of impedance controllers for plants of the following form:

$$\begin{bmatrix} \Delta\dot{\Theta} \\ \Delta\ddot{\Theta} \\ \Delta\dot{\mathbf{T}} \end{bmatrix} = \begin{bmatrix} 0 & I & 0 \\ -M^{-1}(\Theta_0)GR(\Theta_0) & C & M^{-1}(\Theta_0)T_s \\ 0 & 0 & A_a \end{bmatrix} \begin{bmatrix} \Delta\Theta \\ \Delta\dot{\Theta} \\ \Delta\mathbf{T} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ B_a \end{bmatrix} \Delta\mathbf{U} + \begin{bmatrix} 0 \\ M^{-1}(\Theta_0)J_c' \end{bmatrix} \Delta\mathbf{D}$$

Or, in short:

$$\Delta\dot{\mathbf{X}} = A\Delta\mathbf{X} + B\Delta\mathbf{U} + L\Delta\mathbf{D}$$

where all of the submatrices of A are $n \times n$ (n d.o.f.); Θ_0 ($\dot{\Theta}_0 = 0$) is the operating point of the system, about which the linearization occurs; M is the inertia matrix of the system; GR is a gravity term; T_s is a nonsingular square matrix which represents the effect of the actuator torques, $\Delta\mathbf{T}$, on the coordinates, $\Delta\Theta$; A_a is a diagonal matrix of actuator bandwidths; $B_a = -A_a$; $\Delta\mathbf{U}$ is the input to the actuators; $\Delta\mathbf{D}$ is the force applied by the environment; and J_c is a jacobian matrix relating the coordinates of the interaction ports to $\Delta\Theta$.

This form is quite similar to that shown in the last section, except that gravity is included, joint damping is omitted, and, most importantly, the actuators no longer behave as pure sources. Instead, each actuator has a first order rolloff described by one of the diagonal terms in the matrix A_a. The actuator dynamics remain decoupled, however, in the sense that there are no back effects from the dynamics of the hardware.

Kazerooni's method involves matching the behavior of this system to that of the following set of target dynamics, over some frequency range $0 < \omega < \omega_0$:

$$\begin{bmatrix} \Delta\dot{\Theta} \\ \Delta\ddot{\Theta} \end{bmatrix} = \begin{bmatrix} 0 & I \\ -J_c^{-1}J^{-1}KJ_c & -J_c^{-1}J^{-1}CJ_c \end{bmatrix} \begin{bmatrix} \Delta\Theta \\ \Delta\dot{\Theta} \end{bmatrix} + \begin{bmatrix} 0 \\ J_c^{-1}J^{-1} \end{bmatrix} \Delta\mathbf{D}$$

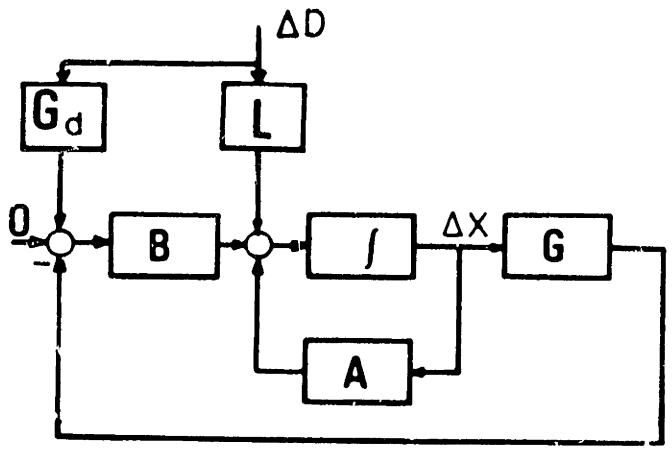

Figure 4.3: Closed loop control system, taken from [44]. G is a state feedback matrix, and G_d is an input feedback matrix.

where J is an inertia matrix, C is a damping matrix, and K is a stiffness matrix, all expressed in interaction port coordinates. The matrices J, C, and K are chosen by the designer to represent the desired dynamic response.

The details of this technique will not be reviewed here, as they are rather involved, but they include the use of an eigenvector assignment technique to design a state feedback matrix, as well as a novel technique for designing an input feedback matrix. The general form of the controller is shown in Figure 4.3.

Kazerooni provides a number of quantitative examples in [44]. One of these (example 1, case 1) applies to a two-link manipulator like that described in the previous section. In this example, the target dynamics are chosen to be soft in the x-direction at the endpoint, and about 20 times stiffer in the y-direction, as might be used to track a surface parallel to the y-axis.

The results may be summarized with Bode plots of the driving point admittance (Figure 4.4). The design specifically calls for matching the target behavior out to a bandwidth of 1 Hz (6.28 rad/sec), which is clearly achieved.

There is a problem, however. The phase of the impedance clearly dips below $-90°$,

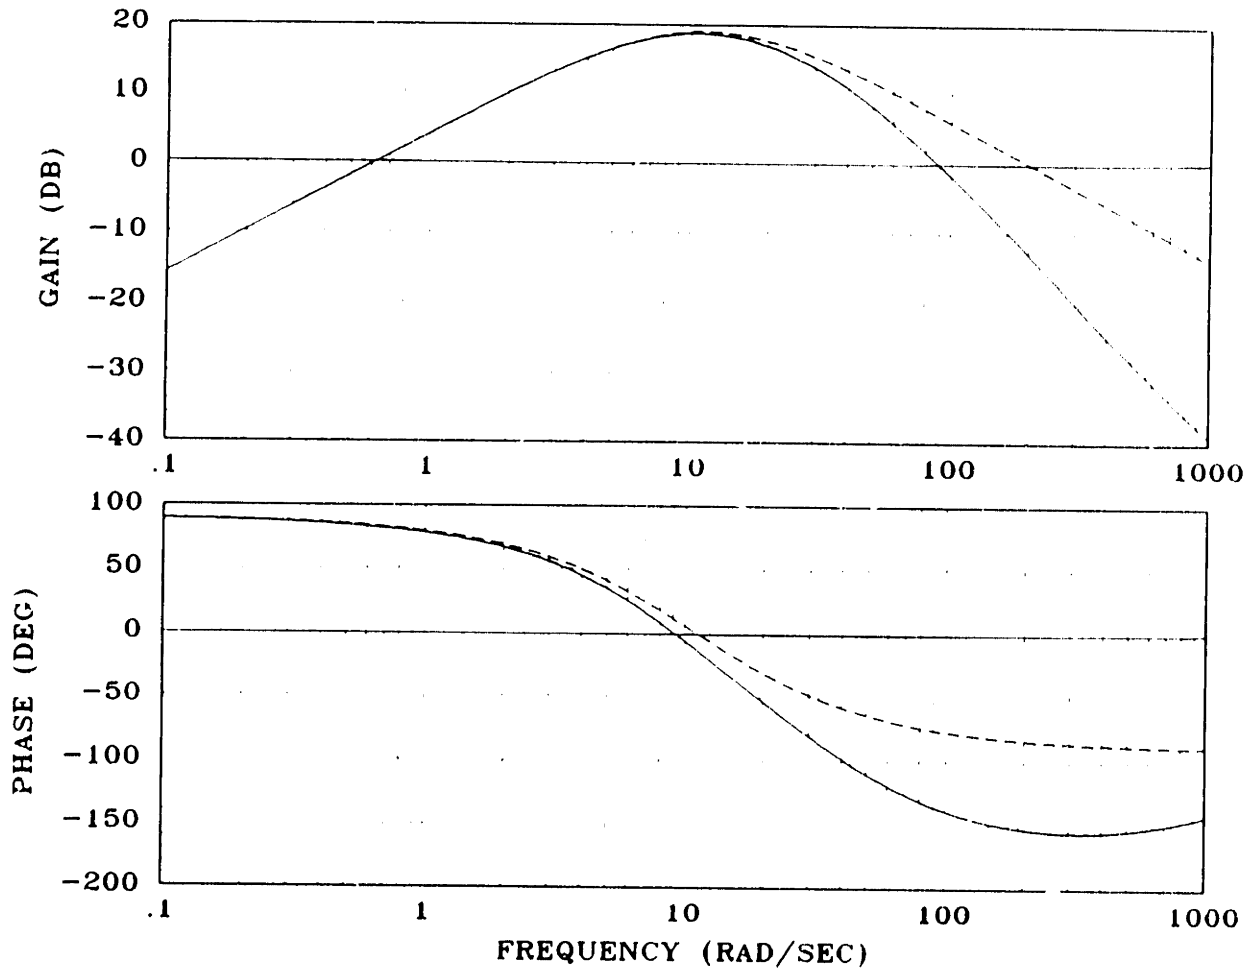

Figure 4.4: Bode plots of closed loop driving point admittance (solid lines) and target model admittance (dashed lines). Example taken from example 1, case 1, [44]; upper left components of admittance matrices.

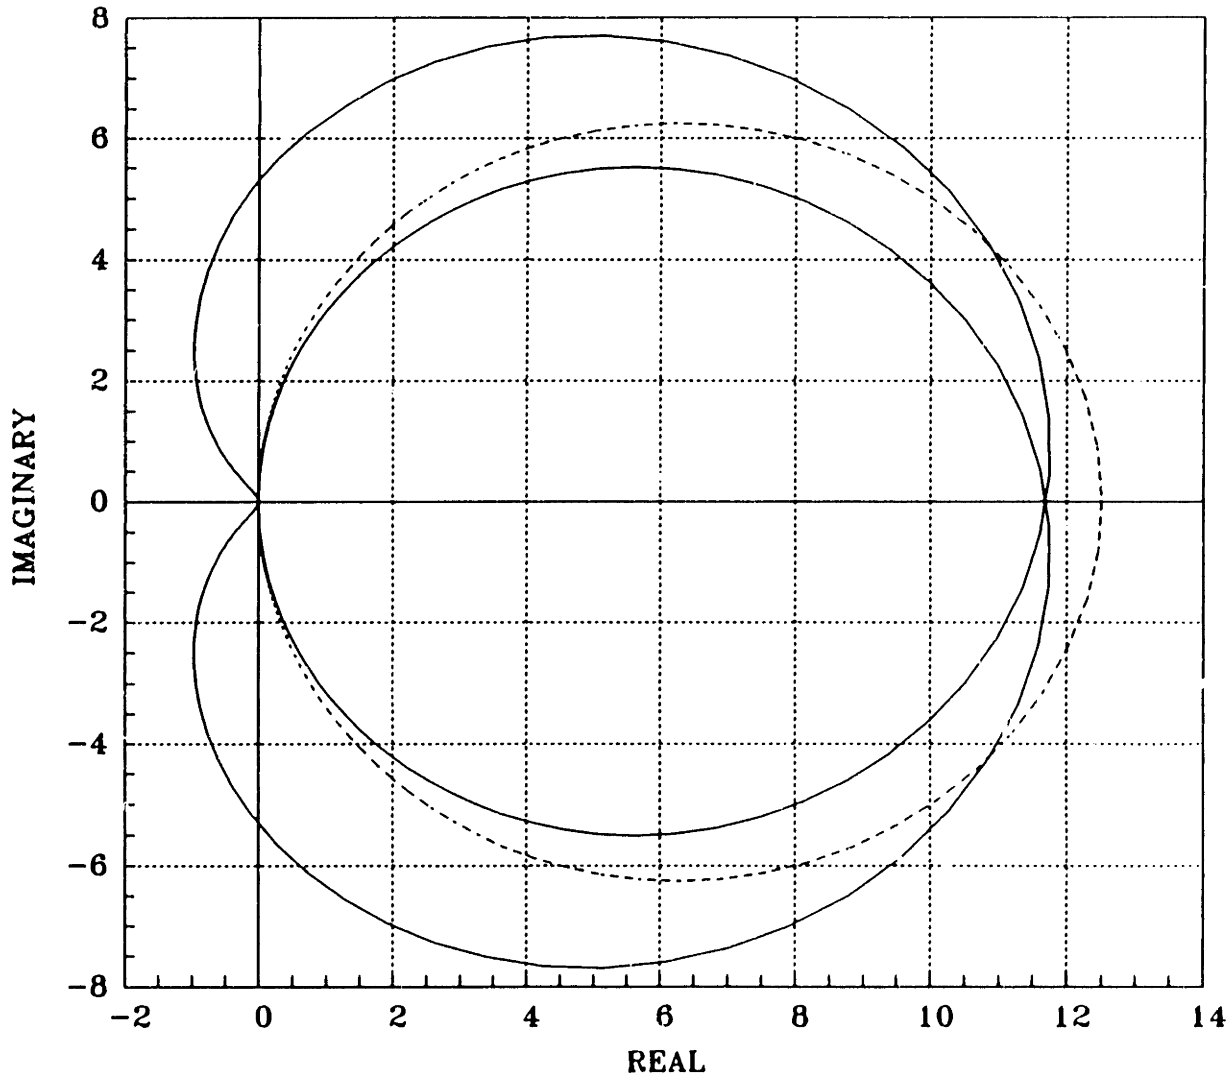

Figure 4.5: Nyquist plots of closed loop driving point admittance (solid lines) and target model admittance (dashed lines). Example taken from example 1, case 1, [44]; upper left components of admittance matrices.

and this appears in the Nyquist plot (Figure 4.5) as a loop in the left half plane. The consequence (as a worst environment test would reveal) is that this system is unstable upon interaction with stiff environments. This might not be so damaging, except that the goal of the procedure was to develop a controller suitable for contact applications.

Thus, even when explicitly included in a design procedure, first-order actuator dynamics complicate matters severely. In Chapter 8 it will be shown that the greatest source of difficulty is not the actuator dynamics alone, but rather the combination of input feedback and actuator dynamics.

4.2 Force Feedback

This section presents a simple analysis of a force-feedback controlled manipulator in an effort to explain the limitations of its interactive behavior. The control of force exerted on a rigid surface is certainly the example of manipulator interaction which has received the greatest attention in the literature. Part of the reason for this attention (besides the potential value of a force-controlled robot in a manufacturing setting) is the notorious "contact instability" associated with force-feedback. Contact instability is the phenomenon of violent chatter which can occur when a force-feedback controlled robot is brought into contact with a rigid surface. A detailed account of this phenomenon, along with a review of the relevant literature, is presented in Chapter 7.

For the present, an examination of the coupled stability properties of two simple models should elucidate one source of trouble. Consider, to begin, the manipulator model shown in Figure 4.6—a rigid body model reduced to its barest essence. If the force control law $u = G(F_0 - F)$ is chosen for this manipulator, then the closed loop driving point admittance is:

$$Y(s) = \frac{1 + G}{Ms}$$

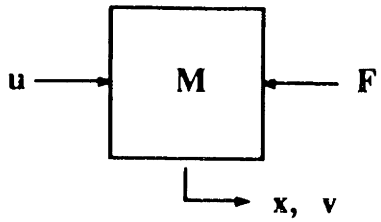

Figure 4.6: Rigid body manipulator. u is the actuator force; F is a force imposed by the environment.

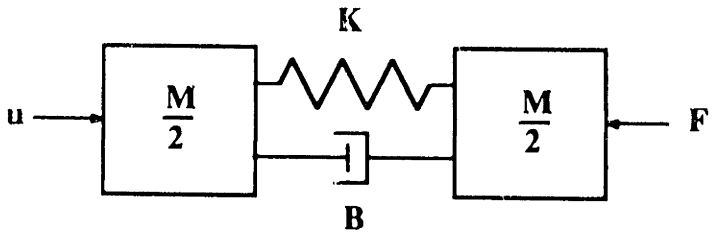

Figure 4.7: Manipulator with non-colocated actuation and sensing.

This admittance will be positive real so long as $G \geq -1$. This is encouraging, because servo theory would suggest that the route to high-bandwidth force control would be to make G as large as possible.

Now consider another simple manipulator model, shown in Figure 4.7. An essential feature of this model is that the actuator (u) and sensor (F) are non-colocated [25,20]. Non-colocation is given quantitative meaning in Appendix D; however, it can be understood as a "dynamic" separation of two physical locations—dynamic in the sense that power flowing between the locations must pass through energy storing elements. The reason for introducing non-colocation is described in Chapter 7; it is an essential factor in predicting contact instability.

The closed-loop admittance of this manipulator is:

$$Y(s) = \frac{s^2 + 2B((1+G)/M)s + 2K((1+G)/M)}{\frac{1}{2}Ms^3 + 2Bs^2 + 2Ks}$$

Passivity 6b can be used to check the positive realness of this admittance. Because $M > 0$, $B > 0$, and $K > 0$, $Y(s)$ has no poles in the RHP. It does, however, have

a pole at the origin. Although this pole is simple, its residue must be checked:

$$\lim_{s \to 0} sY(s) = \frac{1+G}{M} \geq 0$$

Given that this is the residue associated with the rigid body mode, it should not be too surprising that it leads to the same result as the rigid body example above, namely, that $G \geq -1$.

In addition, however, it is necessary that $Re\{Y(j\omega)\} \geq 0$. It is straightforward to show that:

$$Re\{Y(j\omega)\} = \frac{(1-G)B\omega^4}{M^2\omega^6 + 4(B^2 - MK)\omega^4 + 4K^2\omega^2}$$

The denominator of this expression is the sum of squares, so that the only new condition is $G \leq 1$.

The interesting result is that, no matter what the exact values of M, B, and K are, so long as they are greater than zero, the coupled stability condition is $-1 \leq G \leq 1$.

Nyquist plots of this system are shown in Figure 4.8 for the following parameter values: $M = 2$, $B = 2$, and $K = 10{,}000$. The loops correspond to behavior near the structural resonance; the remainder of each Nyquist plot lies on the imaginary axis. The top plot is for $G = 0$, so that the system *is* passive, and as passivity requires, it lies completely in the closed right half plane. The lower plot is for the seemingly moderate force gain of $G = 2$, but the difference is rather dramatic. The loops now lie completely within the closed left half plane, and coupled stability is clearly in jeopardy.

The root locus test can be used to learn more about the types of environments which cause problems; this is shown in Figure 4.9. Evidently, it is springs which cause the difficulty. The worst instability occurs when the stiffness of the environment is comparable to the structural stiffness of the manipulator, however, any stiffer environment will also cause instability. To verify this, the centroid of the root locus can be found. The centroid is the point on the real axis through which the two asymptotes pass. Because the pole and zero locations in this example may be found analytically,

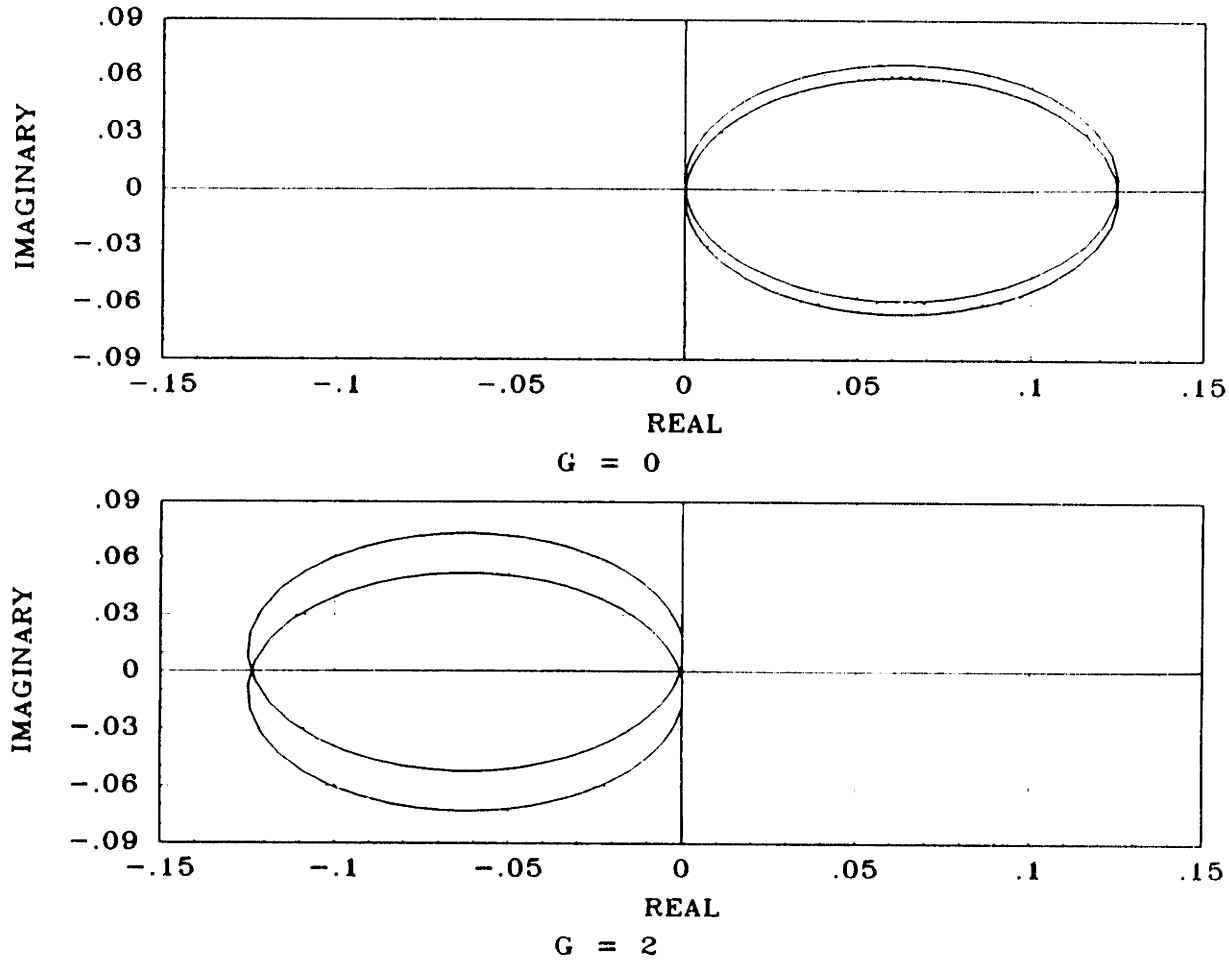

Figure 4.8: Nyquist plots of a force controlled manipulator's admittance. Open loop ($G = 0$), top; closed loop ($G = 2$), bottom.

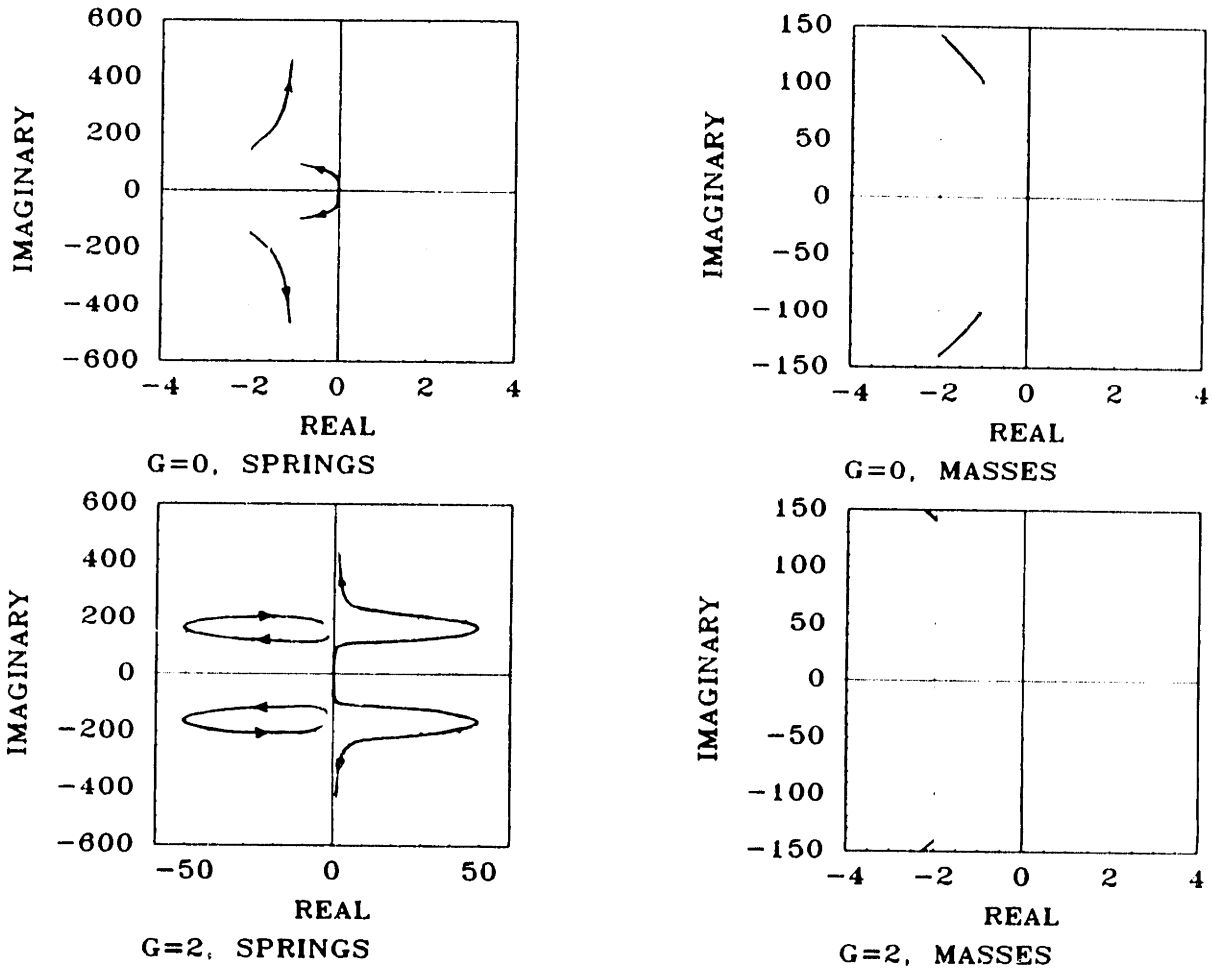

Figure 4.9: Worst environment root loci for a force controlled manipulator.

it is straightforward to calculate the centroid location. The result is:

$$\sigma_{centroid} = (G-1)\frac{B}{M}$$

The conclusion is that, for $G > 1$, stiff enough environments must lead to instability.

A final note: springs are not the only environments which can lead to instability. A simple scaling of the Nyquist plot for $G = 2$, as would occur if the system were coupled to a large enough viscous damper, would be sufficient to create instability.

4.3 PID Control

This section examines the PID control of second order systems. A simple SISO control example is considered in some detail, and a digital implementation is developed for the two-link manipulator.

4.3.1 SISO Example

Consider the PID control of a second order system described by the following state equations ($M > 0$; $B, K \geq 0$):

$$\begin{bmatrix} \dot{x} \\ \dot{v} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -K/M & -B/M \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} + \begin{bmatrix} 0 \\ 1/M \end{bmatrix} u + \begin{bmatrix} 0 \\ 1/M \end{bmatrix} F$$

A block diagram representation of the closed loop system is shown in Figure 4.10, and the driving point admittance of this system is:

$$Y(s) = \frac{s^2}{Ms^3 + (a+B)s^2 + (b+K)s + c}$$

According to Corollary 3 of Section 2.4.1, the zeros of a positive real function must be simple, a condition which is violated by the double zero at the origin. A necessary condition for $Y(s)$ to be positive real is $c = 0$, in other words, the removal of all integral action. This is not very surprising in terms of conventional notions of passivity. The resulting admittance of the PD controlled system is:

$$Y(s) = \frac{s}{Ms^2 + (a+B)s + (b+K)}$$

It is relatively easy to show that the PD controller results in a positive real $Y(s)$ so long as $a \geq -B$ and $b \geq -K$. These are unrestrictive conditions, as it is nearly always the case that a and b are greater than zero (corresponding to negative feedback and a minimum phase compensator zero). It should also be noted that the simple impedance

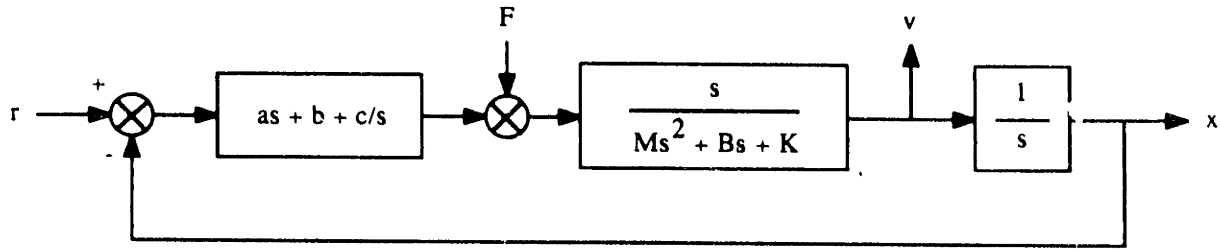

Figure 4.10: PID control of a second order system.

controller introduced previously is essentially a multivariable PD controller. Its robust coupled stability property should, therefore, be expected.

Returning to PID control, it is reasonable to ask what type of environment will lead to instability. Again, the root locus test may be used. Before proceeding with this, however, it would be useful to specify the compensator parameters more carefully.

Most techniques for selecting the three gains (a, b, and c) of a PID controller are based upon tradeoffs between certain performance specifications, such as time to peak, maximum overshoot, steady state error, and stability margin [17]. If the plant is second order, however, the gains may be selected to invert the plant. If we select $a = M/\tau$, $b = B/\tau$, and $c = K/\tau$, then the zeros of the compensator exactly cancel the poles of the plant, leaving the well behaved loop transfer function $1/\tau s$. τ is a free parameter which can be used to set the closed loop bandwidth. Given this choice of parameters, the driving point admittance is:

$$Y(s) = \frac{s^2}{(s + 1/\tau)(Ms^2 + Bs + K)}$$

The coupled stability properties of this admittance can be examined by generating the root loci for spring environments and mass environments. Examples of these are shown in Figure 4.11. Apparently, it is large masses that lead to difficulty. The minimum mass which will cause instability may be found analytically to be:

$$(M_e)_{min} = \frac{BM}{K\tau} + \frac{B^2}{K} + B\tau$$

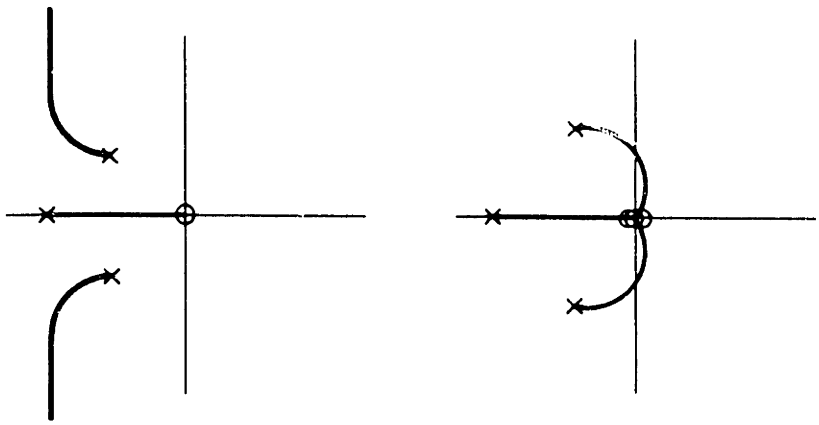

Figure 4.11: Second order system with a PID compensator: interaction with springs (left) and masses (right).

$$= \frac{B}{K\tau}(M + B\tau + K\tau^2)$$

The dependence upon plant parameters indicates that increasing M and B increases the M_e necessary to create instability, whereas increasing K decreases the necessary M_e; the dependence on τ is not monotonic. In any case, a stiff, underdamped system will be driven unstable by a small mass. This may also be seen on the root locus; the closer the complex poles move toward the imaginary axis, the smaller the mass that is needed to make the locus cross the axis.

This analysis, of course, has its limitations, the most severe of which is that a pure derivative term (as) cannot be implemented. In practice, the derivative action must be rolled off with at least one additional pole ($as/(Ts+1)$). This, of course, complicates the selection of τ; if one chooses T too small and τ too large, the closed loop behavior may be unacceptably oscillatory.

4.3.2 Two-Link Manipulator Implementation

For the set of experiments involving PID control, the two-link manipulator was used as a test-bed for second order systems. To do this, the nominal elbow angle of the manipulator was set at 90° to decouple the links, and proportional position and velocity feedback loops were closed around each joint. The position and velocity feedback gains (K and B) could be used to set the pole locations for each link independently. The goal of the PID control would then be to cancel the pair of poles for each joint, and to insert instead the simple loop transfer function $1/\tau s$.

This may appear to be a rather peculiar approach to manipulator control, adding dynamics with one feedback loop, then cancelling these dynamics with another. In fact it is odd, but it must be understood that the emphasis here is not on manipulator control, but rather on the use of the manipulator plus control to represent a generalized second-order system. In particular, this has been done to demonstrate that two control systems with identical loop transfer functions may have significantly different driving point impedances. Toward this end, three different PID controllers have been implemented to compensate for three different "plants".

For each link, the control is of the form:

$$\begin{aligned} U &= U_1 + U_2 \\ U_1 &= -K\theta - B\omega \\ U_2 &= \frac{1}{\tau}(B + \frac{K}{s} + \frac{Mas}{s+a})e \end{aligned}$$

where $e = \theta_0 - \theta$. U_1 is the control which places the two poles. U_2 is the PID control; a sets the bandwidth of the derivative compensation. Of course, these controllers are digitally implemented, and therefore must be converted to the z-domain:

$$\begin{aligned} U_{1k} &= -K\theta_k - B\omega_k \\ U_{2k} &= \frac{1}{\tau}(B + \frac{KT}{z-1} + Ma\frac{z-1}{z-\delta})e_k \end{aligned}$$

where k is the discrete time index, T is the sampling period, and $\delta = e^{-aT}$. For ease of implementation, these equations are converted to state space form:

$$\begin{bmatrix} \mu_{k+1} \\ \nu_{k+1} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\delta & 1+\delta \end{bmatrix} \begin{bmatrix} \mu_k \\ \nu_k \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} \theta_k \\ \omega_k \end{bmatrix}$$

$$U_k = \begin{bmatrix} \beta_1 & \beta_2 \end{bmatrix} \begin{bmatrix} \mu_k \\ \nu_k \end{bmatrix} + \begin{bmatrix} -(\beta_3 + K) & -B \end{bmatrix} \begin{bmatrix} \theta_k \\ \omega_k \end{bmatrix}$$

where:

$$\beta_1 = (Ma(1-\delta) - KT\delta)/\tau$$

$$\beta_2 = (KT - Ma(1-\delta))/\tau$$

$$\beta_3 = (B + Ma)/\tau$$

and where μ and ν are the states of the compensator. If these equations are treated as those of link 1, and similar equations for link 2 are augmented, the result may be written in the following compact notation:

$$\chi_{k+1} = \Phi_c \chi_k + \Gamma_c \Theta_k$$

$$U_k = C_c \chi_k + D_c \Theta_k$$

where:

$$\chi' = \begin{bmatrix} \mu_1 & \mu_2 & \nu_1 & \nu_2 \end{bmatrix}$$

$$\Phi_c = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\delta_1 & 0 & 1+\delta_1 & 0 \\ 0 & -\delta_2 & 0 & 1+\delta_2 \end{bmatrix}, \quad \Gamma_c = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{bmatrix}$$

$$C_c = \begin{bmatrix} \beta_{1_1} & 0 & \beta_{2_1} & 0 \\ 0 & \beta_{1_2} & 0 & \beta_{2_2} \end{bmatrix}, \quad D_c = \begin{bmatrix} -(\beta_{3_1} + K_1) & 0 & -B_1 & 0 \\ 0 & -(\beta_{3_2} + K_2) & 0 & -B_2 \end{bmatrix}$$

This is the form in which the compensator is implemented.

At this point it would be convenient to examine the closed loop Nyquist plot of the driving point admittance. There is a problem, however, in that the compensator is

discrete time, and the plant is continuous time. The plant model may be discretized, but it must be understood that the driving point impedance is fundamentally a continuous time function, regardless of the control implementation.

This suggests that a continuous time approach is appropriate. However, while simulation is possible, an exact analytical solution for the driving point impedance of a continuous time plant with a discrete time controller does not exist. Of course, a continuous time approximation to the discrete time controller can be used, but it would be useful to investigate the effects of the finite sampling period.

One approach is to discretize the plant assuming a zero-order hold, which is correct at the control input, but incorrect at the environmental input. This approach leads to the following closed loop state equations:

$$\begin{bmatrix} \Theta_{k+1} \\ \chi_{k+1} \end{bmatrix} = \begin{bmatrix} \Phi + \Gamma_B D_c & \Gamma_B C_c \\ \Gamma_c & \Phi_c \end{bmatrix} \begin{bmatrix} \Theta_k \\ \chi_k \end{bmatrix} + \begin{bmatrix} \Gamma_L \\ 0 \end{bmatrix} \mathbf{F}_k$$

$$\mathbf{V}_k = \begin{bmatrix} (0 \quad J) \quad 0 \end{bmatrix} \begin{bmatrix} \Theta_k \\ \chi_k \end{bmatrix}$$

where $\mathbf{V}$ is the cartesian endpoint velocity of the manipulator, and:

$$\Phi = e^{AT}$$

$$\Gamma_B = \int_0^T e^{A\sigma}\, d\sigma\, B$$

$$\Gamma_L = \int_0^T e^{A\sigma}\, d\sigma\, L$$

where A, B, and L are the matrices defined in equation 4.2. Defining $\Sigma' = [\ \Theta' \quad \chi'\ ]$, the state equations can be written in a compact form:

$$\Sigma_{k+1} = \Phi_{cl}\Sigma_k + \Gamma_{cl}\mathbf{F}_k$$

$$\mathbf{V}_k = C_{cl}\Sigma_k$$

The discrete time admittance is now defined as:

$$Y^*(z) = C_{cl}(zI - \Phi_{cl})^{-1}\Gamma_{cl}$$

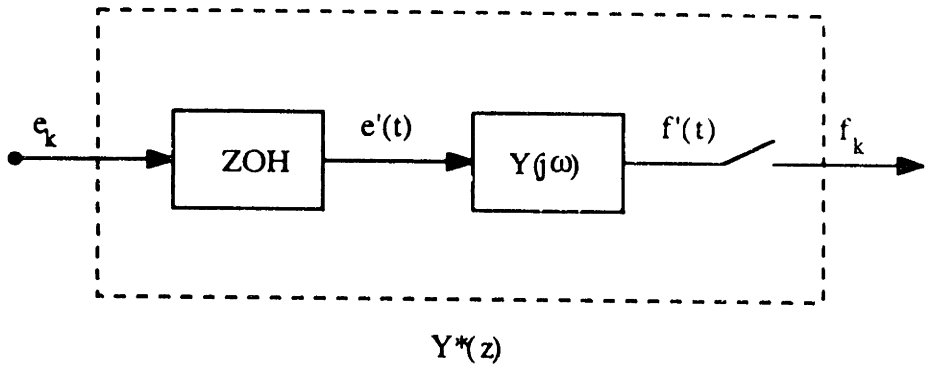

Figure 4.12: Discrete representation of a driving point admittance.

The frequency response plots of discrete time systems generated by commercial software packages involve computation of the gain and phase of $Y^*(e^{j\omega T})$ for ω in the range $0 < \omega < \omega_N$, where ω_N is the Nyquist frequency. If a Bode or Nyquist plot is generated in such a fashion, the question becomes, how does the result relate to the continuous time admittance?

The procedure is to assume that $Y(s)$ exists, as indicated in Figure 4.12, imbedded between a zero-order hold and a sampler, and that the collection is equivalent to $Y^*(z)$. Thus, to extract the frequency response of $Y(s)$ from that of $Y^*(z)$, we need only subtract the effects of the zero-order hold and the sampler. The hold has the following transfer function:

$$H_{ZOH}(j\omega) = T\left[\frac{\sin\left(\frac{\omega T}{2}\right)}{\left(\frac{\omega T}{2}\right)}\right] e^{-j\left(\frac{\omega T}{2}\right)}$$

The sampler has a transfer function of $1/T$. It is straightforward to find the magnitude and the phase of the product of these transfer functions at any frequency ω, and to subtract these from the overall frequency response to arrive at that of $Y(s)$.

A drawback to this method, however, is that $Y^*(z)$ is meaningful only at frequencies less than the Nyquist frequency (π/T rad/sec), so that the high frequency behavior of the impedance is unavailable. For the experiments described in the next chapter, all interesting phenomena occur at low enough frequencies that this is not an issue. This,

therefore, is the analytical approach that has been taken.

At frequencies well below the Nyquist frequency, the combination of zero-order hold and sampler behaves simply as a delay of $T/2$, so that the magnitude scaling is unity, and the phase lag is $\omega T/2$ radians. For the examples presented in this thesis, the Nyquist frequency is 50 Hz and the maximum frequency of interest is 8 Hz, so that it is a simple matter to correct the product of the digital analysis by subtracting the phase lag $\omega T/2$.

This analysis was performed for three different controllers, each designed to compensate for a different second-order plant. The relevant parameters for the three designs are:

Case 1 (subscripts refer to link number):

$$\begin{aligned} \tau_1 &= \tau_2 = 0.2 \text{ sec} \\ a_1 &= a_2 = 75.0 \text{ rad/sec} \\ M_1 &= 0.1869 \text{ kg-m}^2 \\ M_2 &= 0.0860 \text{ kg-m}^2 \\ K_1 &= 25.0 \text{ kg-m}^2/\text{s}^2 \\ K_2 &= 11.5 \text{ kg-m}^2/\text{s}^2 \\ B_1 &= 2.00 \text{ kg-m}^2/\text{s} \\ B_2 &= 0.92 \text{ kg-m}^2/\text{s} \end{aligned}$$

Case 2 (same as Case 1 except as noted):

$$\begin{aligned} B_1 &= 1.00 \text{ kg-m}^2/\text{s} \\ B_2 &= 0.46 \text{ kg-m}^2/\text{s} \end{aligned}$$

Case 3 (same as Case 1 except as noted):

$$\begin{aligned} B_1 &= 0.50 \text{ kg-m}^2/\text{s} \\ B_2 &= 0.23 \text{ kg-m}^2/\text{s} \end{aligned}$$

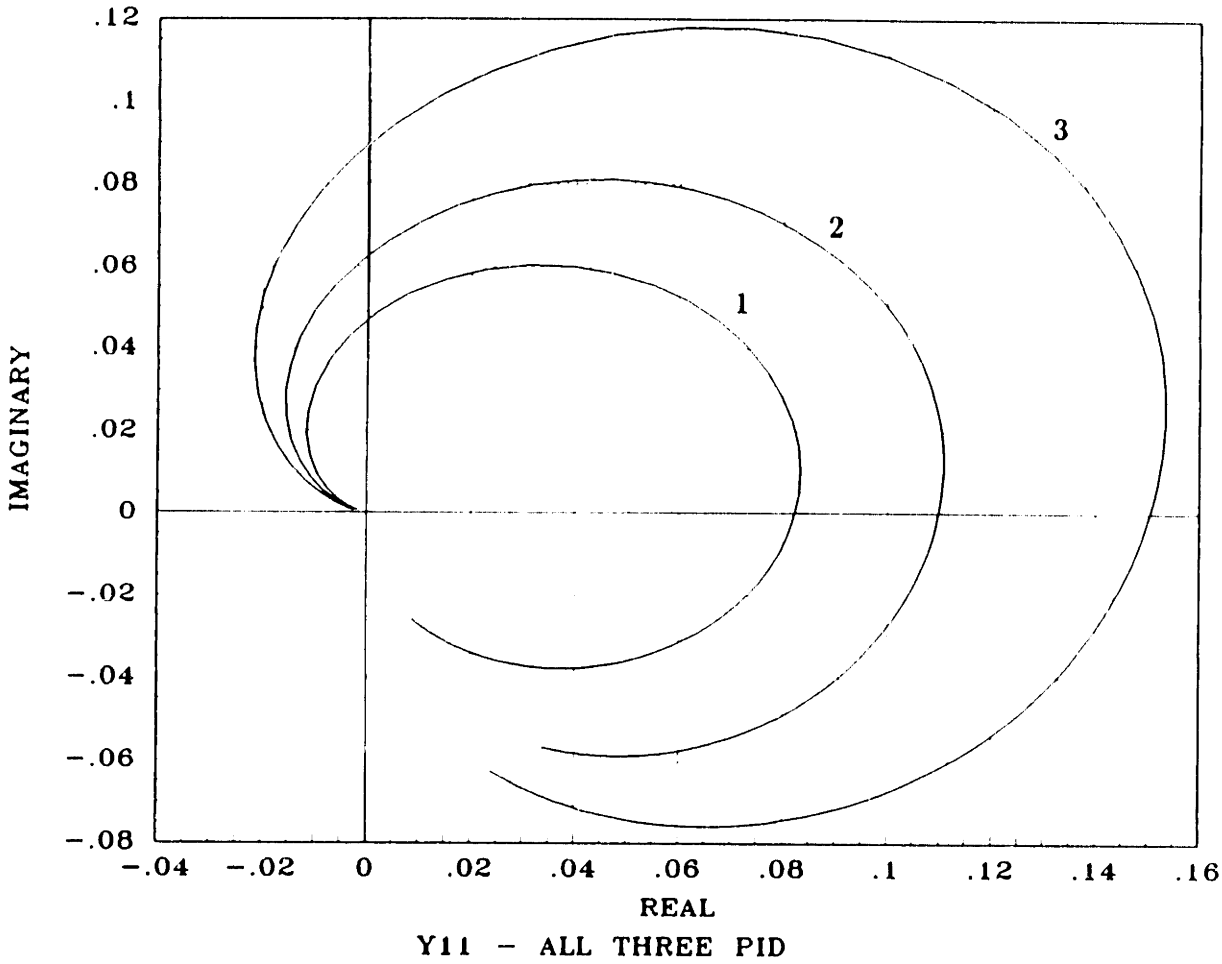

Figure 4.13: Driving point admittance, $Y_{xx}(s)$. All three PID controlled systems. Both axes have units of sec/kg.

The expectation is that these three designs will exhibit similar servo behavior, but due to the different selections of B_1 and B_2, different coupled stability properties. This is supported by the Nyquist plots of $Y_{xx}(s)$ and $Y_{yy}(s)$ shown in Figures 4.13 and 4.14. Although all three designs violate the coupled stability criterion, the masses which will lead to instability will be substantially different in the different cases.

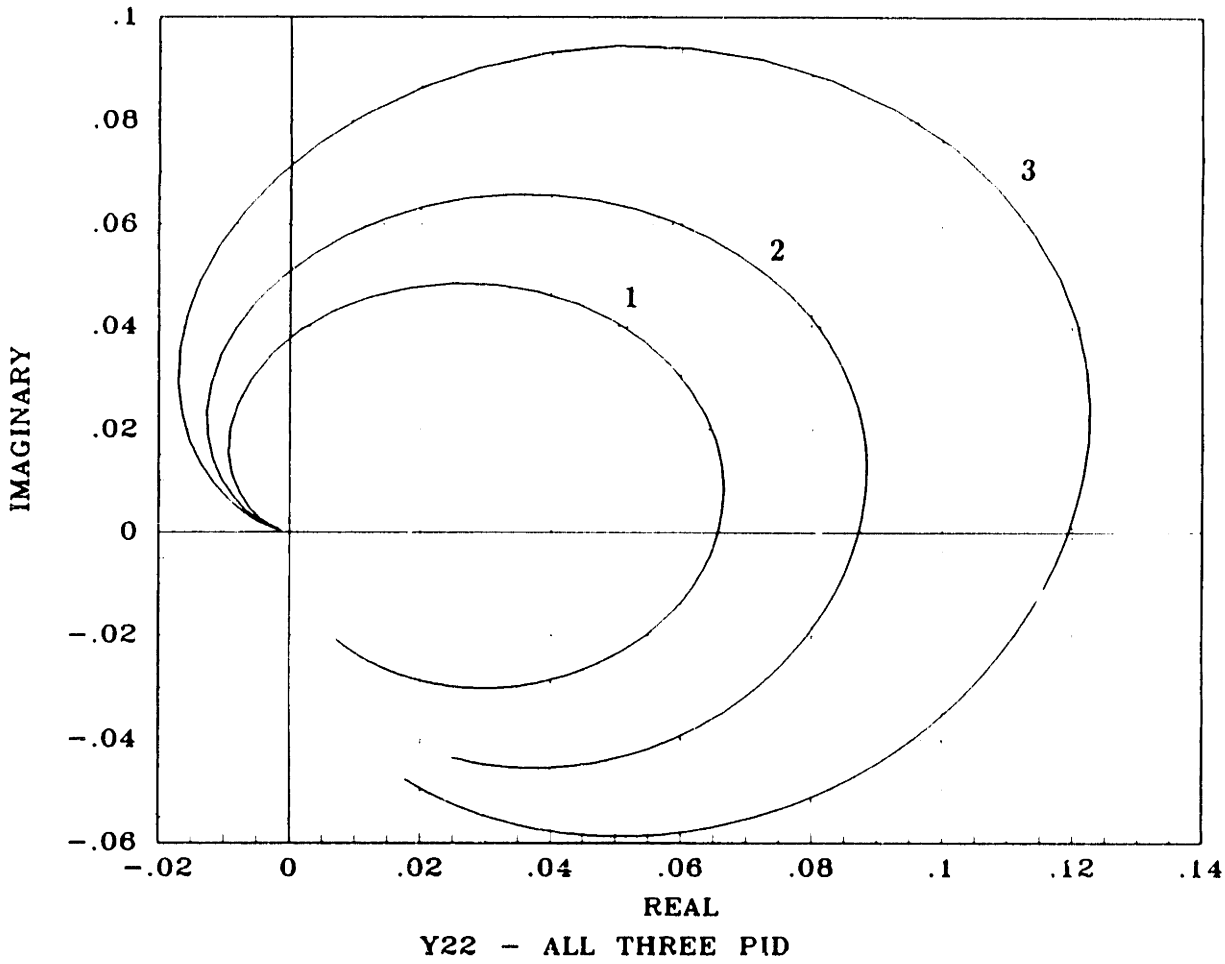

Figure 4.14: Driving point admittance, $Y_{vv}(s)$. All three PID controlled systems. Both axes have units of sec/kg.

4.4 LQG/LTR Control

This section describes an LQG/LTR controller which is implemented on the two-link manipulator. The LQG/LTR method is essentially a multivariable servo design technique; it is described in greater detail in Appendix B. For the purposes of this section, it is important only to understand the following points:

- This implementation of LQG/LTR includes integral action in each control channel.

- The primary goal of the procedure is to "shape" the frequency domain behavior of the singular values of the loop transfer function. In particular, this implementation attempts to "match" the magnitudes of all the singular values, and to make these have a constant slope of -20 dB/decade in the frequency domain, at least for some desired bandwidth. Furthermore, this is done by generating a plant inversion. In other words, this is just a multivariable generalization of the PID strategy presented in the last section.

- To emphasize the multivariable nature of the design, the nominal manipulator configuration is chosen not to be decoupled. The link angles are: $\theta_1 = 30°$, $\theta_2 = 150°$.

- Because the formal LQG/LTR design procedure described in Appendix B is ill-conditioned for a double-integrator plant (which the manipulator very nearly is), position and velocity feedback are used to move the plant poles into the LHP.

The LQG/LTR compensator for this problem is eighth order. The closed loop singular values, assuming no errors in the plant model, are shown in Figure 4.15. The LQG/LTR design procedure is continuous time, but it typically results in high order compensators which require digital implementation, so singular values corresponding

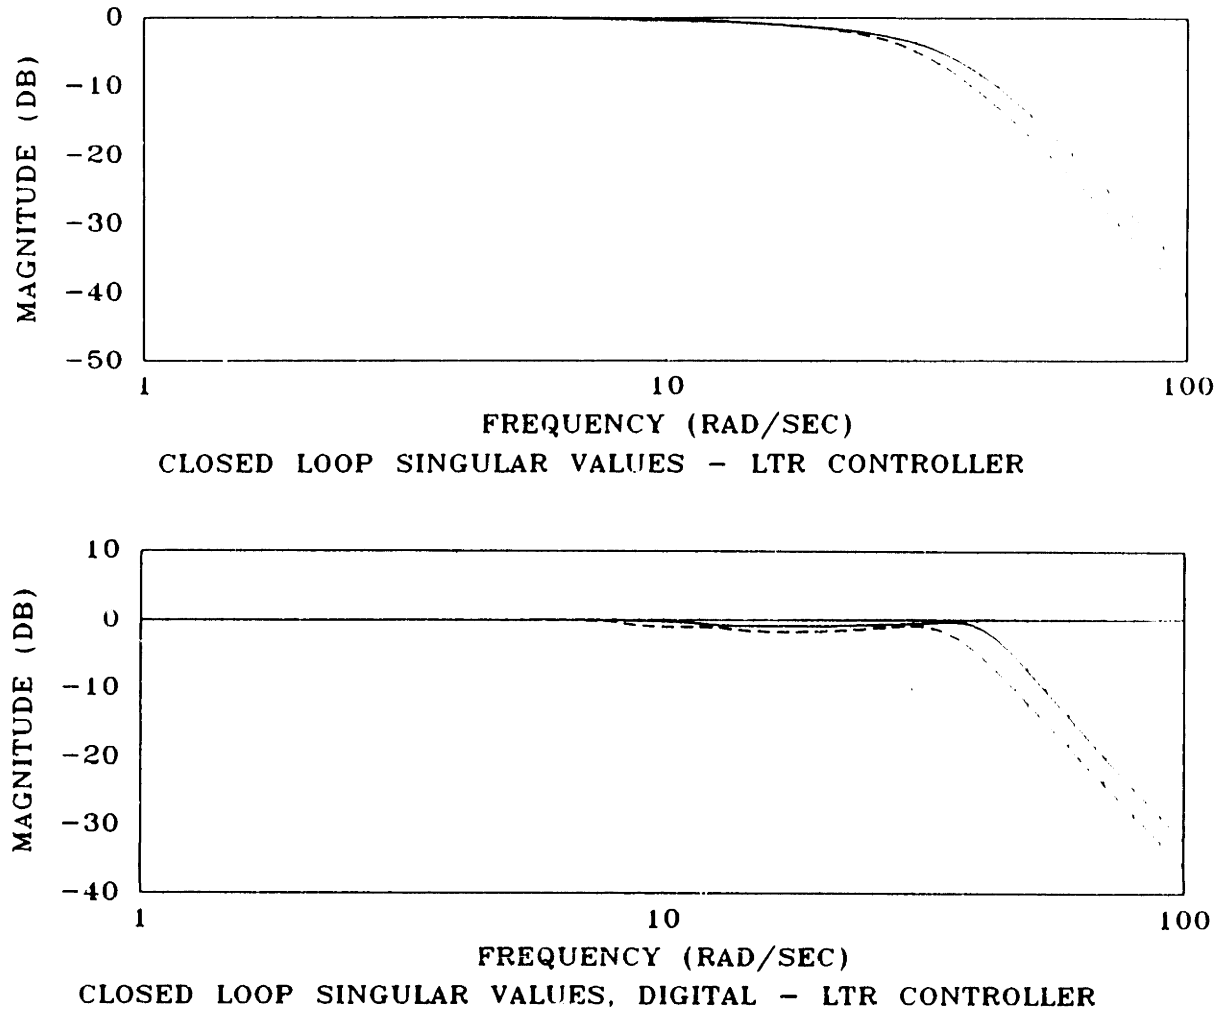

Figure 4.15: Closed loop transfer function singular values, LQG/LTR control. Top: continuous time compensator. Bottom: discrete time compensator.

to both cases are shown. Evidently, the discrete time implementation degrades performance slightly, but both plots have the classic shape of a good command-following control system.

A Nyquist plot of the driving point admittance (discrete time compensator) is shown in Figure 4.16, and Bode plots are shown in Figure 4.17. It is evident that the coupled stability property is not achieved. The large left half plane loop in the Nyquist plot occurs at low frequency and is due to integral action; it suggests stability problems upon coupling to masses. The smaller loop occurs for a range of higher frequencies (5.6

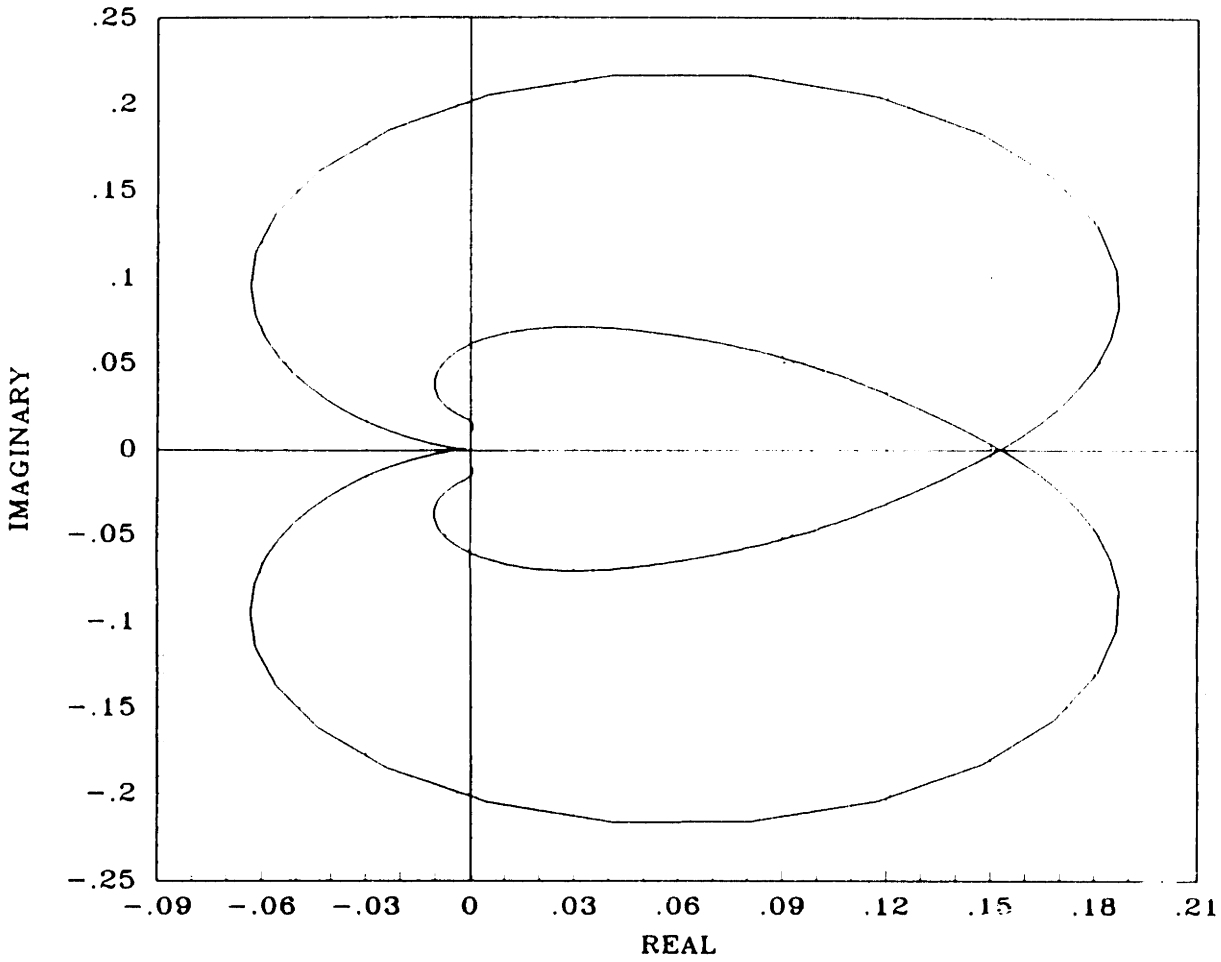

Figure 4.16: Nyquist plot of closed loop admittance ($Y_{vv}(s)$). LQG/LTR control. Both axes have units of sec/kg.

Hz - 9.5 Hz), and suggests stability problems upon coupling to springs with stiffnesses within a certain range.

This controller is intended to have good servo behavior, not necessarily good interactive behavior. The point of introducing it is to show that the two goals are quite distinct. An added benefit of this implementation is that it is also an example of a controller *without* force feedback that can exhibit instability upon contact with a stiff surface. It is important to understand that contact instability is a consequence of the driving point impedance, *not* necessarily of force feedback.

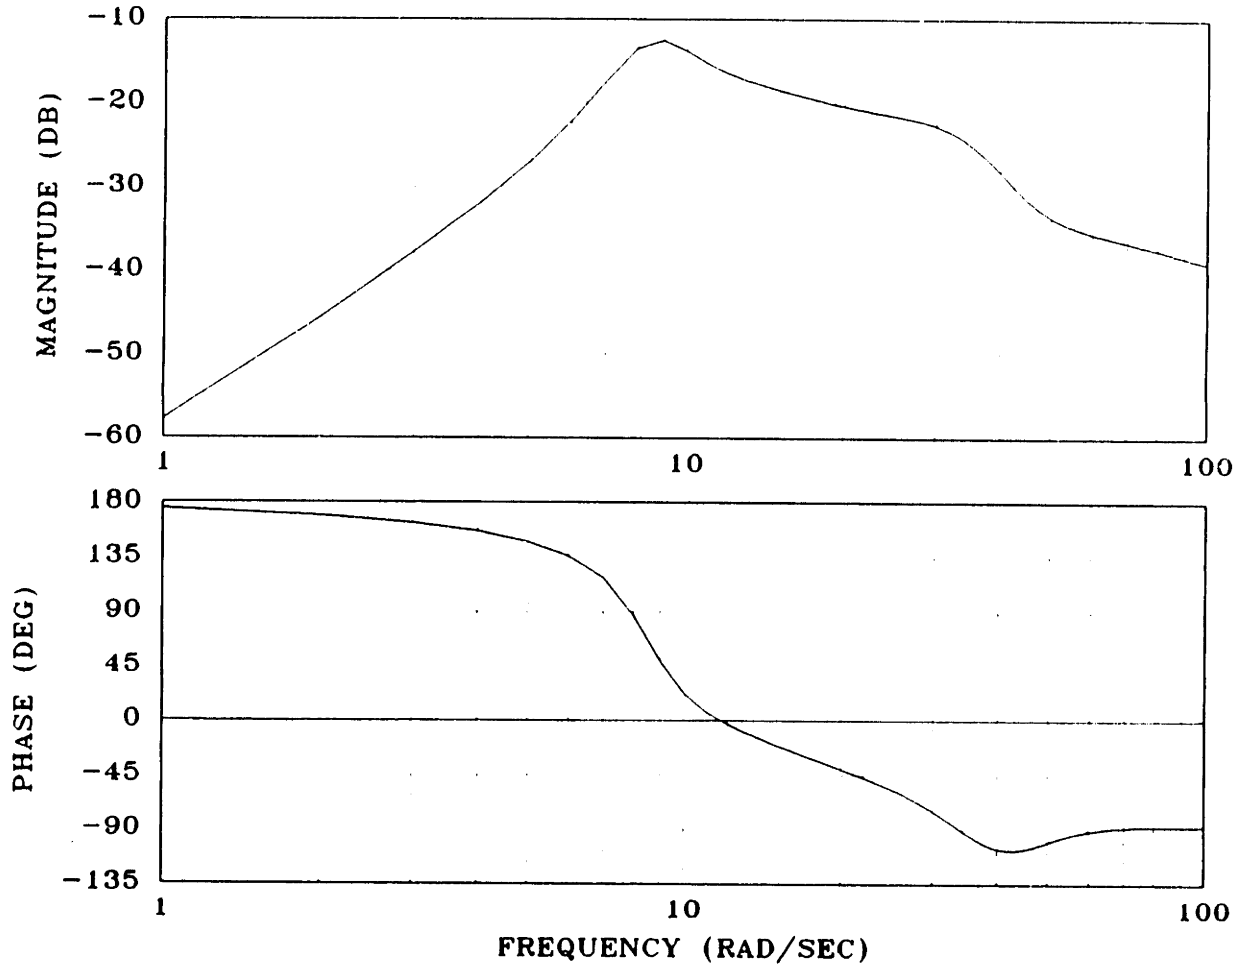

Figure 4.17: Bode plots of closed loop admittance ($Y_{vv}(s)$). LQG/LTR control.

Chapter 5

Experimental Investigation of the Interactive Behavior of a Two-Link Manipulator

5.1 Overview

This chapter describes an experimental investigation of the coupled stability property and its consequences. Two types of experiments are described—the first is a direct measurement of the driving point admittance of a feedback controlled two-link manipulator, and the second is an investigation of the behavior of this manipulator when coupled to spring and mass environments.

One of the primary goals of these experiments is to provide an experimental verification of the relationship between the driving point impedance and coupled stability. The experiments, however, provide insight into a number of other important issues.

For instance, if the properties of the driving point impedance are shown to be valuable indicators of interactive behavior, then it is important to understand those issues relevant to measuring the impedance. One issue is that the phase of the impedance is the critical property for predicting coupled stability, yet phase is notoriously difficult to measure. This is, in fact, the consideration which drives the somewhat elaborate

data acquisition and analysis procedures described in Section 5.2.

Furthermore, the "worst environment" criterion (Section 3.4) indicates that, at least in the case of 1-port linear systems, springs and masses are quite important in their own right. If one's primary concern is that of coupled stability rather than shaping the driving point impedance, it may be appropriate to focus on interaction with springs and masses rather than on the measurement of impedance. The interaction tests have the promise of simple implementation, as well as a closer relationship to practical applications. The interaction tests also provide an opportunity to examine behavior not covered by the linear analysis, such as the effects of joint friction.

Finally, these experiments are intended to indicate the value of a unified approach to interaction based upon the properties of the driving point impedance. Toward this end, it is shown (Section 5.4.4) that a manipulator which does not use force feedback may nevertheless exhibit the contact instability which has come to be associated with force feedback. The point is that contact stability is not fundamentally a force feedback phenomenon, but *is* fundamentally an impedance phenomenon, and may always be analyzed as such.

5.2 Measurement of the Driving Point Impedance

5.2.1 Approach

The goal of these experiments is to measure the driving point admittance at the end-point of the manipulator. Recall that driving point impedance (or admittance) is the transfer function between input and output at a given port. As such, to measure impedance, it necessary to disturb the system at this port, and to measure the input and output. Note that the disturbance need not be created by a pure source, or even

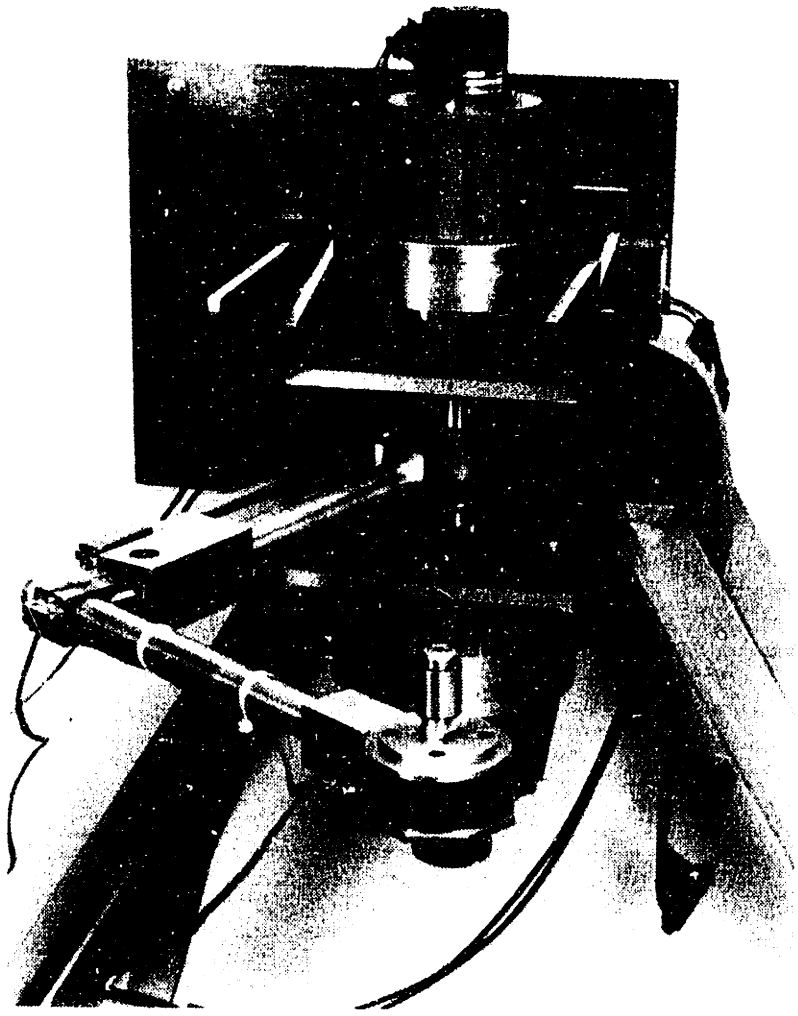

Figure 5.1: Two-link manipulator.

a dynamic system with a very high or a very low impedance. It is necessary only that the inputs and outputs at the port of interest be measured.

Typical disturbances for frequency domain analysis would be either white noise or sinusoids. This apparatus lends itself to the latter because the sinusoidal input can be created simply by rotating the endpoint of the manipulator in a constant speed circle. In this way, both the x and y axes are sinusoidally perturbed. The sinusoids, however, are not independent. The x and y positions are 90° out of phase. To extract the full admittance matrix, it is necessary to take data for both clockwise and counterclockwise rotation at a given frequency (see Section 5.2.4).

The input to the manipulator is assumed to be force, and the output position. Note that a true admittance produces a velocity output, so we will be speaking loosely when the measured quantity is referred to as an admittance. The selection of position rather than velocity as an output hinges simply on the fact that position can, in this implementation, be measured much more accurately. The important consequence of this measure is that the phase of the "admittance" must now lie between 0° and $-180°$ to guarantee passivity.

Despite the sinusoidal nature of the data, random data analysis techniques— the computation of auto- and cross-spectral power densities from an FFT—are employed. Besides being readily available and applicable, these techniques obviate the need for filtering the data, and automatically separate the information of interest from that generated by other frequency components, such as harmonics.

To summarize, the driving point admittance is measured as follows: At a given frequency, forces and positions are measured for both clockwise and counterclockwise rotations. The components of the 2×2 admittance matrix at that frequency are calculated from these data, and the process is repeated at each frequency of interest. Further details are supplied in the next four sections.

5.2.2 Apparatus

The major components of the apparatus for these experiments are the two-link manipulator, a PDP 11/73 computer for control as well as data acquisition, and the "spinner". The first two components are described in detail in [23], however, some relevant features will be reviewed here.

The two-link manipulator (Figure 5.1) is direct drive. The inner link is connected directly to a motor shaft; the outer link is driven through a parallelogram linkage. The two motors are d.c. torque motors with low inductance pancake armatures. The motors are powered by transconductance amplifiers which behave as voltage-controlled current sources, thus, the amplifier-motor combination is accurately modeled as voltage-controlled torque source driving the armature inertia.

The manipulator is equipped with joint angle, joint angular velocity, and endpoint force sensors. Optical encoders with a resolution of 8000 counts/revolution (0.045°) measure the angles. Tachometers measure the angular velocities. The tachometers are filtered with four cascaded first order, 50 Hz filters, for an overall cutoff of 21 Hz. The force transducer (Lord F/T, L01052) mounts on the end of the manipulator. It can measure three axes of force and three axes of torque, although only the two forces in the plane of the linkage are collected. The force resolution is 1/5 ounce; the maximum force that can be measured is 15 lbs. The force and torque data is delivered to the 11/73 via a parallel port at a rate of 104 Hz.

The 11/73 is responsible for both implementing the various control laws, and collecting the data. All controllers are implemented with a 0.01 second update rate; the data acquisition is triggered by interrupts from the force transducer. All software is written in C.

The spinner (Figure 5.2) was designed and built for these experiments. It consists of a d.c. torque motor with the shaft oriented vertically (perpendicular to the plane of the

manipulator), mounted on a concrete pedestal. The pedestal, which weighs about 300 lbs, rests on an aluminum base plate which has three jack-style legs. The legs are used to elevate and align the motor until it can be coupled to the end of the manipulator.

A 6 inch diameter plate with mounting holes at various radii is fixed to the motor shaft. A bearing block connects the force transducer at the end of the manipulator to any of these mounting locations on the plate. Thus, when the motor is powered, the end of the manipulator rotates in a circle of the desired radius.

A tachometer is mounted on the motor shaft and is used in a velocity control loop. The motor has a pancake armature with a large number of poles so that smooth, low velocity operation is possible.

The hardware determines the bandwidth of the admittance measurements. At low enough frequency, the integral action of the PID and LQG/LTR controllers is sufficient to overcome the spinner, and at high enough frequency the force transducer will saturate and the forces on the spinner bearings will exceed manufacturer's specifications. Each of these effects, however is a function of the radius. The bandwidth can be extended at both the lower and upper limits by reducing the radius. The resolution of the encoders set a lower limit to the radius, which is approximately 0.25 inches. Given this limit, the usable bandwidth is $0.2 \text{ Hz} < f < 8.0 \text{ Hz}$.

The controllers are designed to provide their most interesting behavior in this range. However, in order to do this, the full performance capabilities of the manipulator cannot be exploited. Thus, the impedance controller is more compliant than might be desired, and the PID and LQG/LTR controllers exhibit lower closed loop bandwidth than necessary.

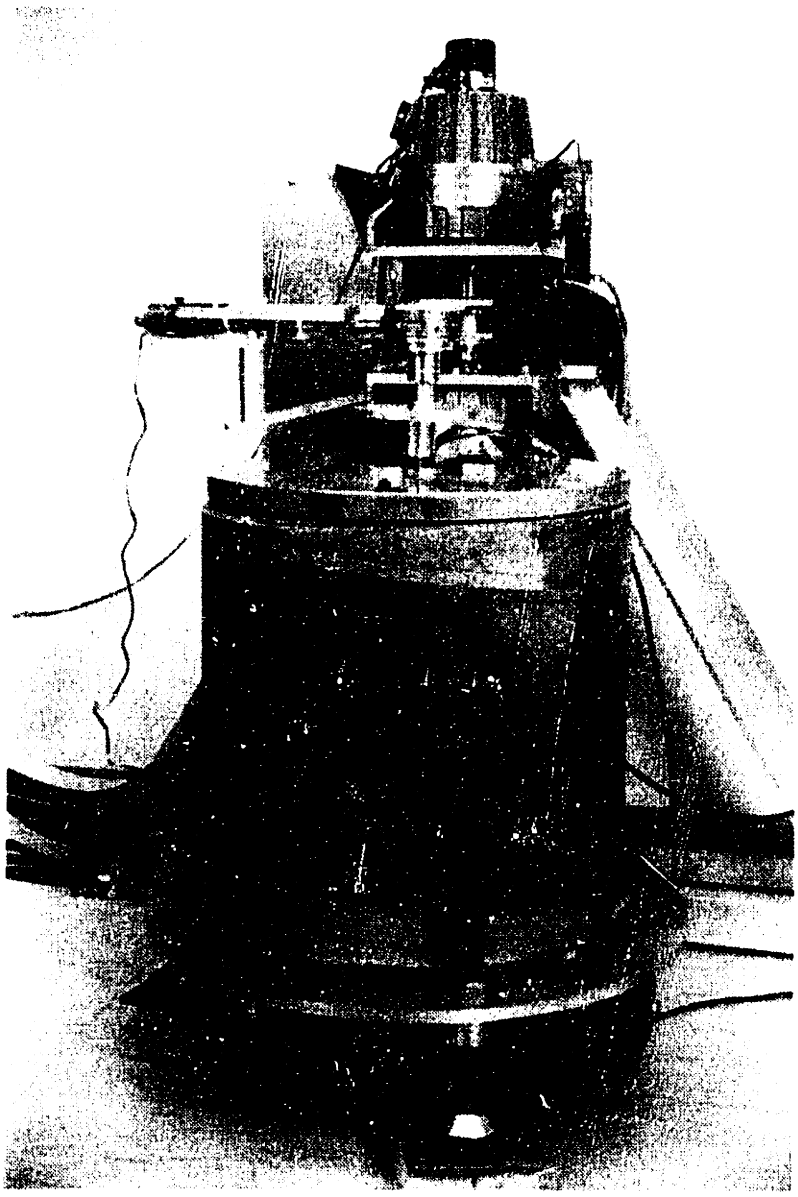

Figure 5.2: The spinner.

5.2.3 Data Acquisition

The raw data collected in a given experiment consist of the x and y forces in the sensor frame, as well as the absolute joint angles. The x and y endpoint positions, rather than the joint angles, are the port variables of interest. However, the manipulator is modeled assuming that the links are rigid (see Section 4.1), so that joint angles and endpoint positions are related through a purely kinematic transformation. This modeling assumption appears to be quite valid for the frequency range of interest.

Data are collected at some multiple ($n_\omega \geq 2$) of 9.6 msec, which is the update rate of the force transducer. The multiple n_ω is chosen to make the sampling rate approximately proportional to the frequency of rotation. The intent is to keep the number of samples/cycle at about 16. The emphasis of the data collection is on record length (1024 points) rather than samples/cycle, due to the behavior of the FFT algorithm: the random error decreases with increasing record length, while the frequency resolution is inversely proportional to the sampling rate.

Figures 5.3 and 5.4(a) explain the data collection scheme. Interrupts generated by the force transducer occur at 104 Hz. The interrupt service routine responds to the interrupt by immediately reading in the forces and angles. The angles are then put into storage for one cycle. This has the effect of adding a delay to the angle measurement comparable to the computational delay of 9.6 ms for the forces. The routine then checks if the appropriate number of cycles (n_ω) have passed since the last storage. n_ω is calculated according to an estimate of the rotational speed, ω, which is generated before data-taking begins. Once n_ω cycles have passed, the data is stored and the entire cycle begins again.

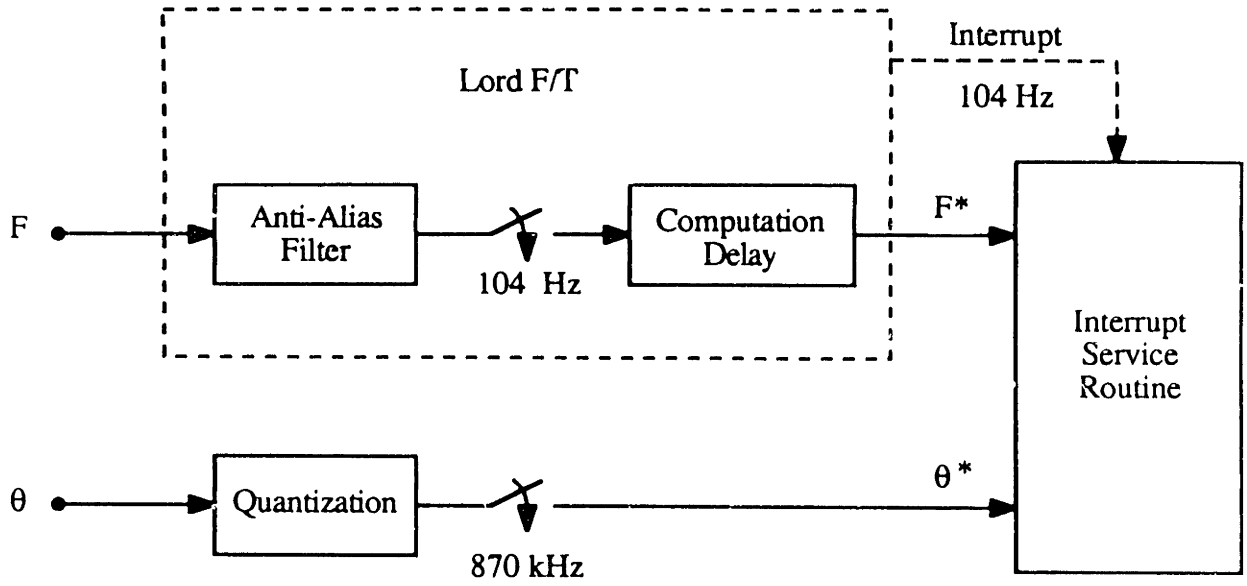

Figure 5.3: Data acquisition scheme

5.2.4 Data Analysis

The analysis of a single set of raw data is flowcharted in Figure 5.4(b). The endpoint positions are determined from the joint angles and the link lengths. The origin of the cartesian frame is taken to be center of rotation, which is also the nominal configuration about which linearization occurs. The endpoint forces are rotated from the sensor frame to the lab frame; the angle of rotation is the instantaneous angle of the outer link, θ_2. Examples of the data at this stage of processing are shown in Figures 5.5 and 5.6.

The mean and linear trends are then removed from the data in preparation for the FFT algorithm. The FFT algorithm divides each record into four contiguous sections of 256 samples each, as well as an overlapping group of three contiguous sections shifted half a section length (128 samples) relative to the others. Each section is shaped with a Hanning window to reduce side-lobe leakage. The purpose of the overlapping sections is to restore some of the power lost to the Hanning windows.

From the FFT results, any of the combinations of auto- and cross-spectral densities

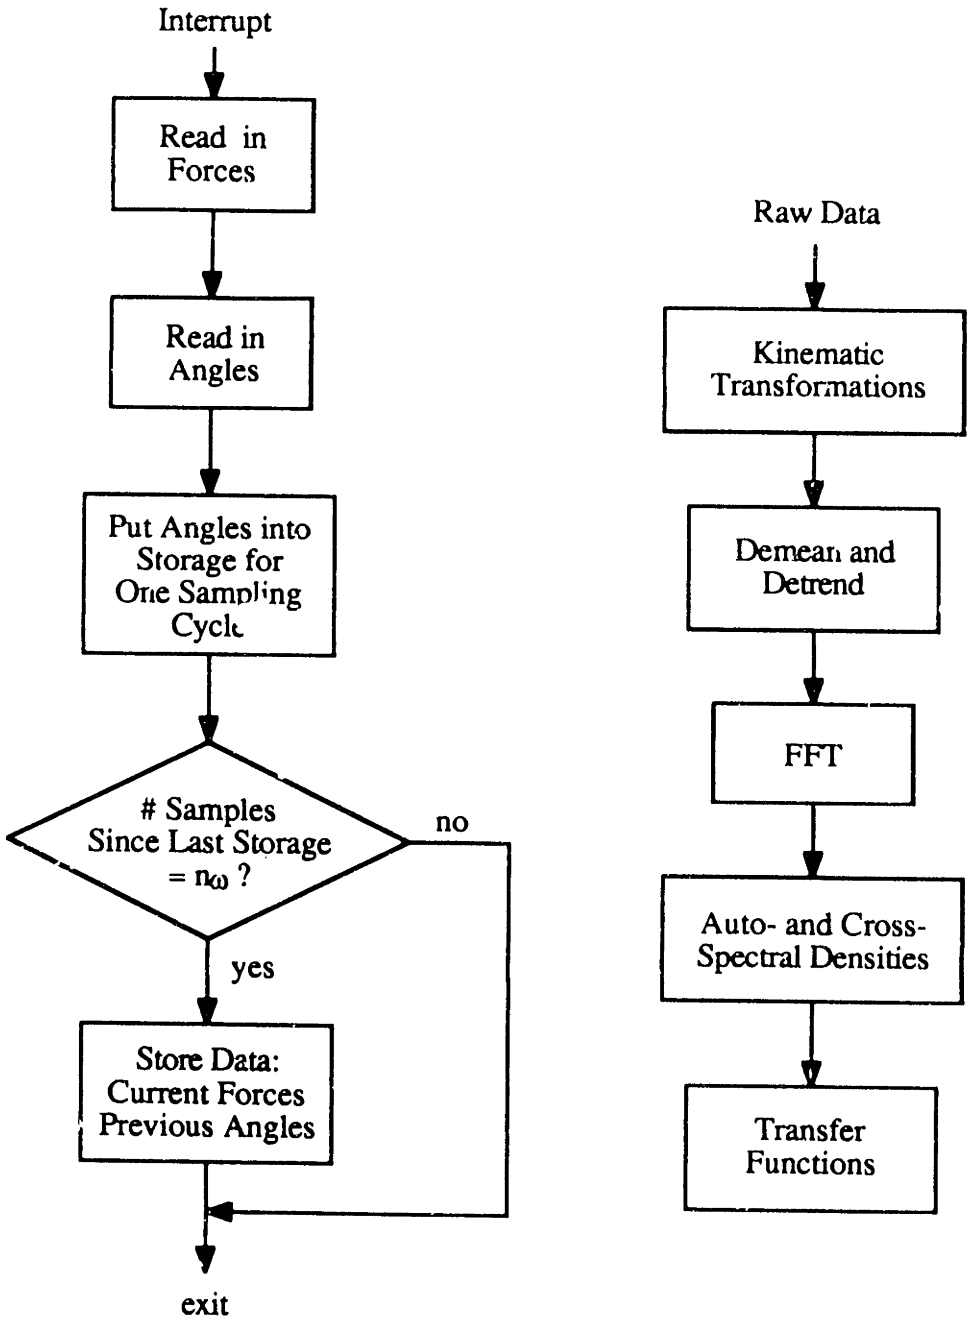

Figure 5.4: (a) Interrupt service routine. (b) Analysis of a single set of raw data.

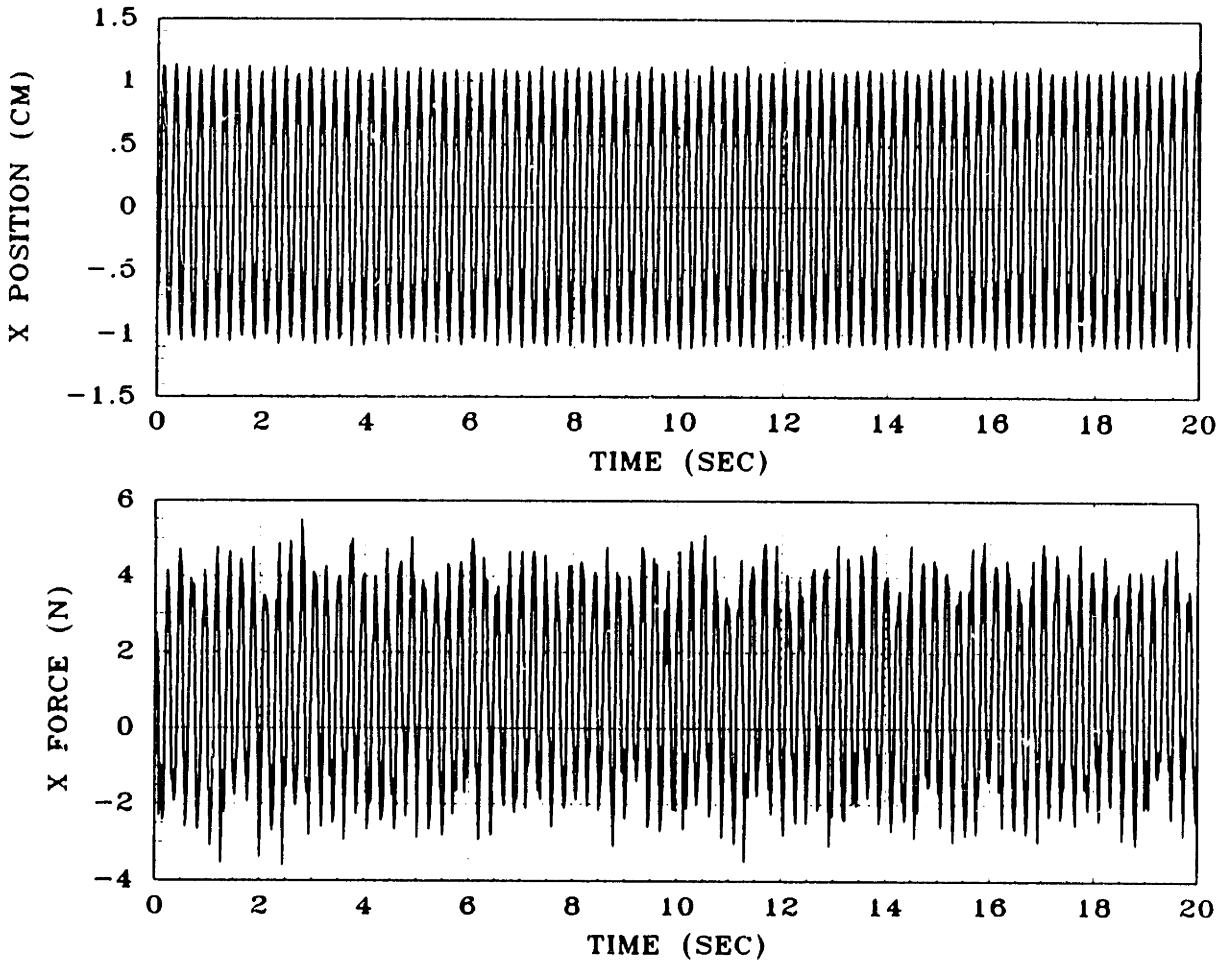

Figure 5.5: An example of the position and force data (PID control, case 2). At the frequency of 4.90 Hz, the coherence between these two records is 0.9996.

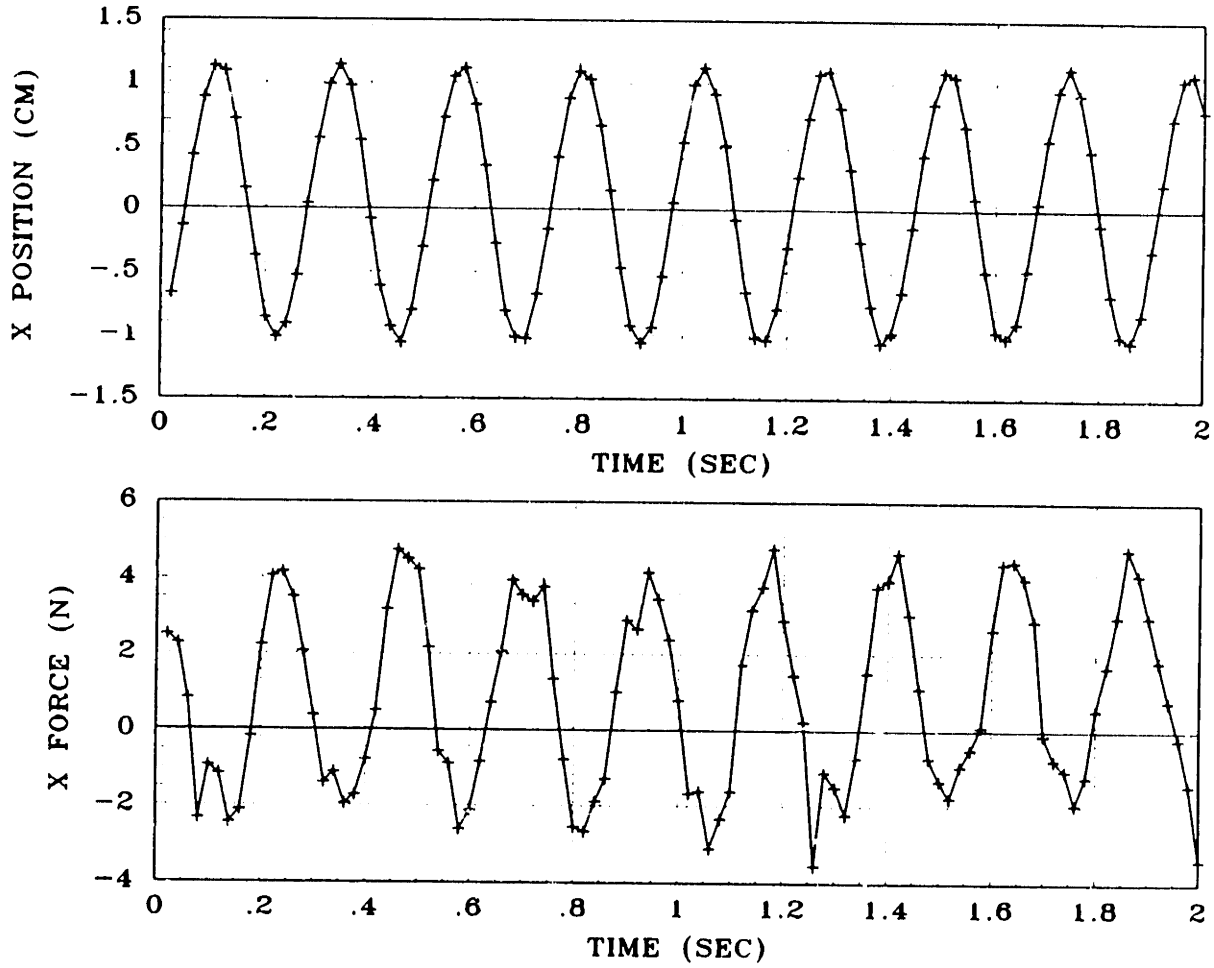

Figure 5.6: An expanded view of the position and force data (PID control, case 2).

may be found. In particular, the following four are evaluated: $G_{xx}(\omega)$, $G_{xy}(\omega)$, $G_{xF_x}(\omega)$, and $G_{xF_y}(\omega)$. At this point, attention is focused on the one frequency of interest. Because the data are largely sinusoidal, these spectra are all fairly sharply peaked at the frequency of the sinusoid. Furthermore, the coherences are generally quite high at these peaks. Typical values are greater than 0.9995. The data were spot-checked for low coherence values, but none were found below 0.9990.

The high coherences justify the calculation of linear transfer functions. The ultimate product of each single set of data is the following collection of transfer functions, each calculated at the one frequency of interest:

$$S(\omega) = \frac{G_{xy}(\omega)}{G_{xx}(\omega)} \tag{5.1}$$

$$H_x(\omega) = \frac{G_{xF_x}(\omega)}{G_{xx}(\omega)} \tag{5.2}$$

$$H_y(\omega) = \frac{G_{xF_y}(\omega)}{G_{xx}(\omega)} \tag{5.3}$$

These transfer functions all take x position as the input, although we are interested in calculating an admittance which takes force as an input. The result, therefore, will eventually need to be inverted. The reason for the circuitous route is simply that the position records contain much less noise than the force records; treating position as the input propagates less error to the result.

The immediate goal, therefore, is to find the impedance matrix Z:

$$\mathbf{F} = Z\mathbf{x}$$

or, in detail:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} Z_{11}(\omega) & Z_{12}(\omega) \\ Z_{21}(\omega) & Z_{22}(\omega) \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \tag{5.4}$$

Because each component of this impedance matrix is a complex quantity, there are, at any frequency, eight unknowns. The following equations can be derived from the

calculated transfer functions:

$$H_x(\omega) = Z_{11}(\omega) + Z_{12}(\omega)S(\omega) \tag{5.5}$$

$$H_y(\omega) = Z_{21}(\omega) + Z_{22}(\omega)S(\omega) \tag{5.6}$$

As these quantities are complex, each single set of data yields four equations—four short of what is needed to solve for the impedance matrix. Two options present themselves. The first is to make do with the four equations. The passivity criterion (Passivity 5a) is actually that $Z + Z^H$ be positive real, which requires that the following equations hold for all $\Re\omega \geq 0$:

$$\begin{aligned} \Re Z_{11} &\geq 0 \\ (\Re Z_{11})(\Re Z_{22}) - \frac{1}{4}(\Re Z_{12} + \Re Z_{21})^2 - \frac{1}{4}(\Im Z_{12} - \Im Z_{21})^2 &\geq 0 \end{aligned}$$

Thus, the only unknowns which necessarily need be found are $\Re Z_{11}$, $\Re Z_{22}$, $(\Re Z_{12} + \Re Z_{21})$, and $(\Im Z_{12} - \Im Z_{21})$. Unfortunately, it is not at all clear how to extract these particular unknowns from equations 5.5 and 5.6. A superior option, however, is simply to find four more equations and to solve for all components of Z. This may be accomplished by taking another set of data at the same frequency, but with the opposite direction of rotation, i.e., counterclockwise instead of clockwise. The result is the following set of equations:

$$H_x^*(\omega) = Z_{11}(\omega) + Z_{12}(\omega)S^*(\omega) \tag{5.7}$$

$$H_y^*(\omega) = Z_{21}(\omega) + Z_{22}(\omega)S^*(\omega) \tag{5.8}$$

Note that, in an ideal world, clockwise rotation would produce $S(\omega) = j$, and counterclockwise rotation would produce $S^*(\omega) = -j$. The solution would then be trivial. Errors in kinematic calibration, however, tend to perturb $S(\omega)$ and $S^*(\omega)$ slightly, and as a consequence, the full set of equations must be solved. The solution

can be broken into two sets of equations in four unknowns. The solution for Z_{11} and Z_{12} is as follows:

$$\begin{bmatrix} \Re Z_{11} \\ \Im Z_{11} \\ \Re Z_{12} \\ \Im Z_{12} \end{bmatrix} = \begin{bmatrix} 1 & 0 & \Re S & -\Im S \\ 0 & 1 & \Im S & \Re S \\ 1 & 0 & \Re S^* & -\Im S^* \\ 0 & 1 & \Im S^* & \Re S^* \end{bmatrix}^{-1} \begin{bmatrix} \Re H_x \\ \Im H_x \\ \Re H_x^* \\ \Im H_x^* \end{bmatrix} \tag{5.9}$$

A similar solution exists for Z_{21} and Z_{22}. Finally, Y is found by inverting Z.

5.2.5 Error Analysis

There are two sources of random error in the results. The first is simply noise; therefore it shows up in each data set. The second is the mis-matching of frequencies for clockwise and counterclockwise sets.

The first source turns out to be the more benign because of the high coherences. Formulas for the normalized random error ε (variance over the mean) can be found in [10]. They are:

$$\varepsilon_{magnitude} = \varepsilon_{phase} = \frac{\sqrt{1-\gamma^2(\omega)}}{\|\gamma(\omega)\|\sqrt{2n}} \tag{5.10}$$

where $\gamma(\omega)$ is the coherence calculated at ω, and n is the number of independently-analyzed sections of data in a set (in this case, $n = 7$). For coherences above $\gamma = 0.999$, the normalized random error is less than 1%. Phase accuracy is, of course, the chief concern. These numbers suggest a phase error on the order of 1°, which is not terribly significant.

The second source of error can be more serious, depending on the frequency resolution and the local rate of change of $Y(j\omega)$. Using an elementary tangent plane approximation, the following approximate expressions are found for the errors in magnitude and phase:

$$\delta\|Y_{11}\| = \frac{\delta\omega}{\|Y_{11}\|}\left[\Re Y_{11}\left(\frac{d\Re Y_{11}}{d\omega}+\frac{d\Im Y_{12}}{d\omega}\right)+\Im Y_{11}\left(\frac{d\Im Y_{11}}{d\omega}-\frac{d\Re Y_{12}}{d\omega}\right)\right] \tag{5.11}$$

$$\delta\angle Y_{11} = \frac{\delta\omega}{2\|Y_{11}\|^2}\left[-\Im Y_{11}\left(\frac{d\Re Y_{11}}{d\omega}+\frac{d\Im Y_{12}}{d\omega}\right)+\Re Y_{11}\left(\frac{d\Im Y_{11}}{d\omega}-\frac{d\Re Y_{12}}{d\omega}\right)\right] \tag{5.12}$$

where $\delta\omega$ is the difference in the frequencies of the clockwise and counterclockwise data sets. Similar expressions may be found for the other components of Y. These formulas are most easily applied in the case of passive data (Section 5.4.1), when the phase is approximately constant at 180°. In this case, the formulas reduce to the following:

$$\delta\|Y_{11}\| = 17.372\frac{\delta\omega}{\omega}\ \text{dB} \tag{5.13}$$

$$\delta\angle Y_{11} = \frac{\|Y_{12}\|}{\|Y_{11}\|}\frac{\delta\omega}{\omega}\ \text{rad} \tag{5.14}$$

$$\delta\angle Y_{12} = \frac{\|Y_{11}\|}{\|Y_{12}\|}\frac{\delta\omega}{\omega}\ \text{rad} \tag{5.15}$$

The formula for $\delta\angle Y_{12}$ has been included because its relationship to $\delta\angle Y_{11}$ is now quite clear. In particular, it is generally the case that $\|Y_{12}\|$ is quite a bit smaller than $\|Y_{11}\|$, so that the errors in the phase of the off-diagonal terms tend to be quite a bit higher than the errors in the diagonal terms.

Of course, this error is not truly random—very poorly matched data points are easily identified. As a consequence, this source of error generally adds less than ±3° uncertainty to the diagonal terms, and less than ±8° of uncertainty to the off-diagonal terms.

In addition to the random errors, there are sources of systematic error in this experiment. The most critical of these is the time delay between the measurement of force and the measurement of position. At the highest frequency, 8 Hz, pure delay accounts for 2.9°/msec of phase lag. A delay as small as 1 msec is extremely difficult to achieve, particularly given that the force transducer makes its measurements over an 800 μsec period, then spends 9.6 msec transforming these measurements from the sensor frame to an appropriate cartesian frame. Furthermore, the data is pre-filtered before sampling, producing additional lag. The prefilter is part of the transducer,

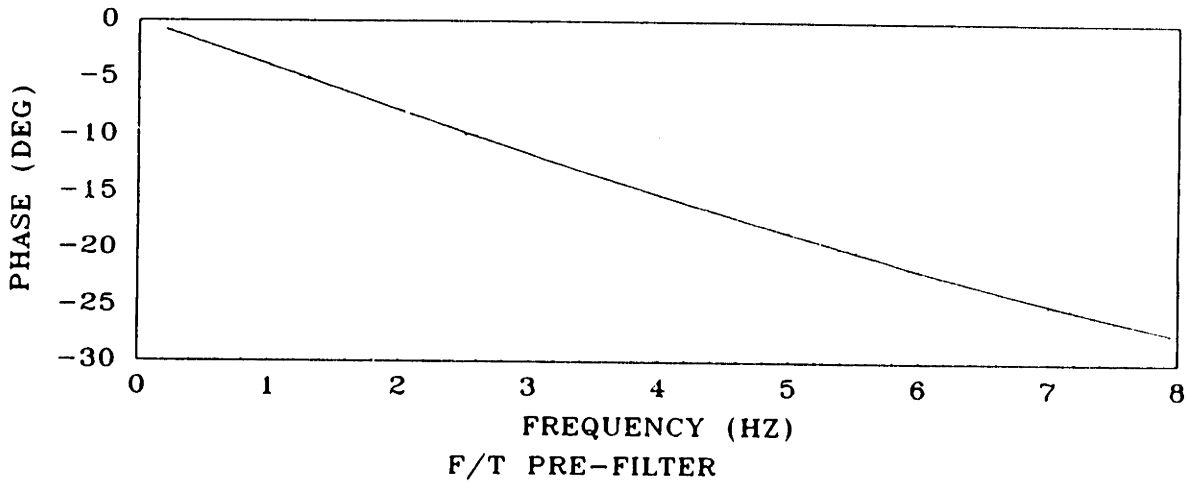

Figure 5.7: Phase of the force transducer pre-filter, assuming nominal values. Note that the phase is very nearly proportional to frequency, as would be the case for a pure delay.

supplied by the manufacturer. It has the following nominal form, although the exact coefficient values are suspect:

$$G_{filter}(s) = \frac{152s + 3.44 \times 10^4}{s^2 + 531s + 3.44 \times 10^4}$$

It is fortunate, however, that the phase of this filter varies approximately linearly over the frequency range of interest (see Figure 5.7). Although the slope of this relationship is uncertain, the filter behaves very much like a pure delay, and is therefore lumped together with the computational delay.

The point of the intentional delay in the storage of angles (Figure 5.4(a)) is to compensate for all or part of those delays in the force measure. But how to know how much? The approach taken here has been to gather data when all manipulator control is turned off—when it is truly a passive system. In this case, we can feel confident that the phase lag (force to position) should be close to 180° at high frequencies. Certainly the phase lag should not be greater than 180°, as this would violate passivity. The use of this information is described further in Section 5.4.1.

5.3 Interaction Tests

By contrast with the impedance measurement experiments, the interaction experiments are quite straightforward. These experiments are not meant to be nearly so exacting, but rather are meant to demonstrate that the consequences of meeting or failing to meet the coupled stability criterion are quite real. However, it is also felt that experiments such as these are an important first step toward understanding many of the more difficult issues that affect any real problem. Two such issues that are addressed by these experiments are the effect of friction (rather than viscous damping) at the joints, and the differences and similarities between contact and coupled stability.

5.3.1 Mass Interaction

Figure 5.8 shows the manipulator interacting with a mass. So that the manipulator need not be burdened with weight, the mass is hung from the ceiling. For the small motions in these experiments, the gravitational restoring force is not noticeable. In fact, due to friction at the joints of the manipulator, there is no single equilibrium point. The mass is also free to spin relative to the end of the manipulator.

The protocol is extremely simple. The manipulator is given a 4 cm diameter, 2 sec/cycle reference trajectory. Data (positions) are stored for 10 sec, in which time the manipulator should negotiate approximately 5 cycles, at least if it is stable. This simple procedure is repeated for all controllers and varying levels of added mass.

The added masses have an uncertainty of approximately 0.1 kg because the portion of the force transducer between the strain elements and the environment must be treated as part of the environment, but this mass is not exactly known.

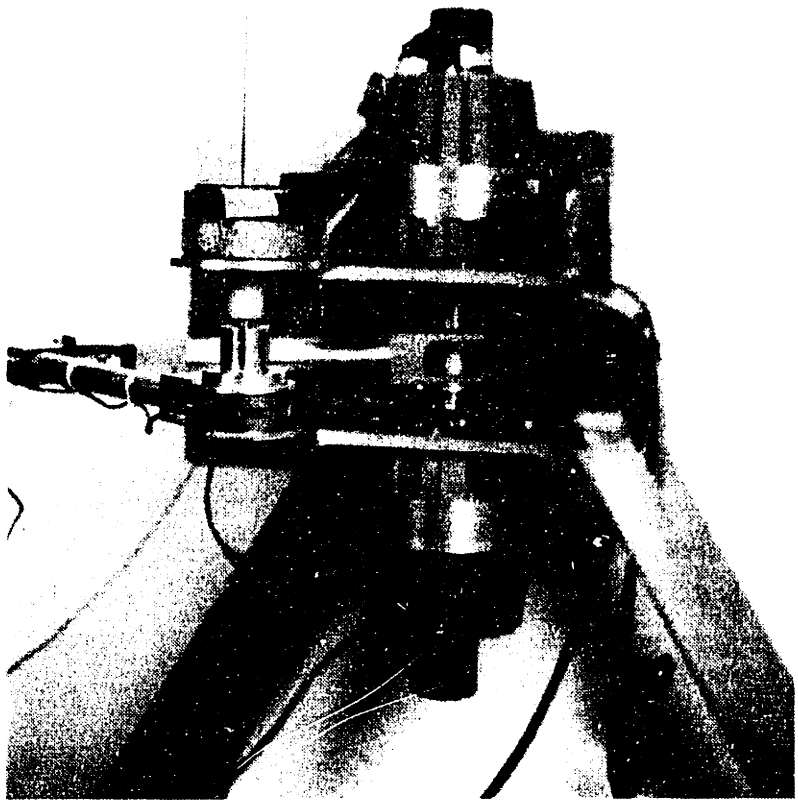

Figure 5.8: The manipulator coupled to a mass. The mass-holder is hung from the ceiling.

5.3.2 Spring Interaction

Figure 5.9 shows the manipulator interacting with a spring. The spring is a simply supported beam whose stiffness depends on the cantilevered length. The spring was calibrated by hanging weights from it and measuring the deflection with a micrometer. Nevertheless, the stiffness varies as the cube of the length of the cantilever and during any given experiment the the distance of the coupling location from the support is sufficiently uncertain to make this only a crude quantitative experiment.

For the coupled stability experiments, the manipulator holds the beam between a pair of pinch rollers. The manipulator is given a ramp input which reaches a plateau at 0.5 cm. The slope of the ramp is variable, but is held at 10 cm/sec for most of the experiments.

For the contact stability experiments, only a single roller is mounted on the end of the manipulator. It is given the same circular reference trajectory as in the mass interaction experiments, however, the trajectory intersects the spring.

5.4 Results

5.4.1 No Control

As described in Section 5.2.5, the strictly passive (no control) data provides an important benchmark. There is sufficient confidence in the predicted behavior of the passive system that it can be used to correct, to some extent, for the systematic errors. It will be shown that it can also be used to correct certain modeling errors.

Figure 5.10 shows the phase data for each component of $Y(j\omega)$. Error bars corresponding to frequency mis-matching are shown, and, as suggested in Section 5.2.5, the errors are much smaller for the diagonal terms. The data, as shown, is corrected for

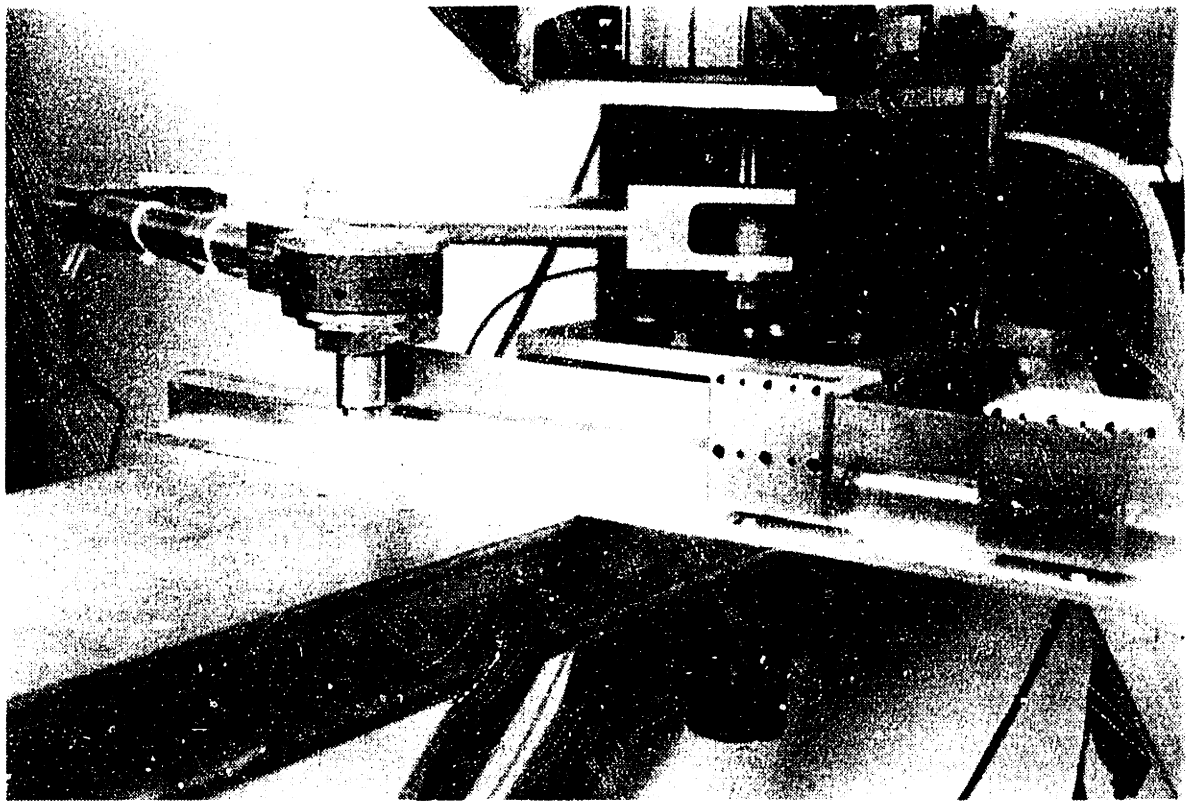

Figure 5.9: The manipulator coupled to a spring (cantilever beam).

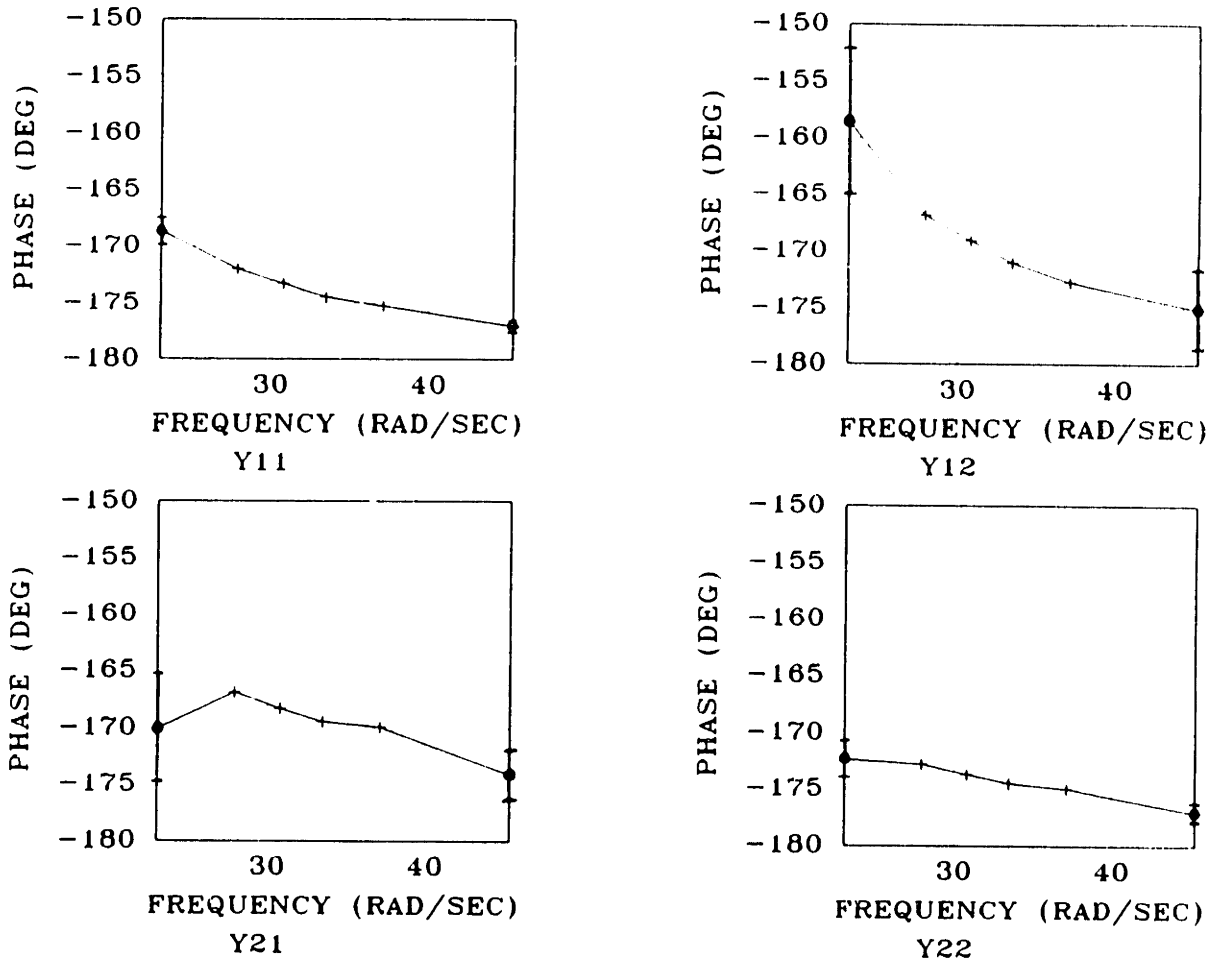

Figure 5.10: Phase of $Y(j\omega)$: no control data.

the pure delay. The total correction is the one sample period storage delay shown in Figure 5.4(a). One would expect the phase plot to tend to zero slope at high enough frequency. If this such data were available, it would be possible to figure and correct for the delay nearly exactly. Unfortunately, the data continue to show some slope at the highest measurement frequency so that the systematic error may still amount to about 5° at this frequency. Nevertheless, this is more than adequate accuracy for the analysis of the closed loop systems.

If the manipulator were purely inertial, one would expect the phase to be $-180°$ exactly. The downward slope suggests the presence of frictional losses. This is not

too surprising, but the magnitude of these losses is considerably greater than that which could be accounted for by the losses in the model (Section 4.1.2). There are two probable explanations for this. The first is that the damping in the model is simply too low. This is likely, because the conditions under which it was measured were not the conditions of normal usage [23]. The second is that the manipulator is loaded out-of-plane when it is coupled to the spinner. Even a very small out-of-plane deflection can amplify the frictional effects substantially.

Figure 5.11 shows the magnitude data, along with the magnitude estimates of the model. The difference is substantial. The primary reason for this is that the model includes the entire mass of the force transducer in the manipulator inertia matrix, whereas only that part of the transducer on the manipulator side of the strain elements should be included. However, the difference is actually larger than can be accounted for by this, suggesting that the estimates of link inertias are in error. In any case, the manipulator model is readily corrected to match the data. The results are shown in Figure 5.12.

This sort of correction has been made before each of the subsequent experiments with closed loop controllers. The driving point impedance predictions have updated to reflect the improved inertial parameters, but the controller designs have not been altered.

5.4.2 Impedance Control

Figures 5.13 and 5.14 are the magnitude and phase plots of $Y(j\omega)$ for the simple impedance controller. The diagonal terms match the model predictions fairly well, but the off-diagonal terms do not match as well, particularly at low frequency. This is actually not very surprising, as the impedance controller is attempting to implement

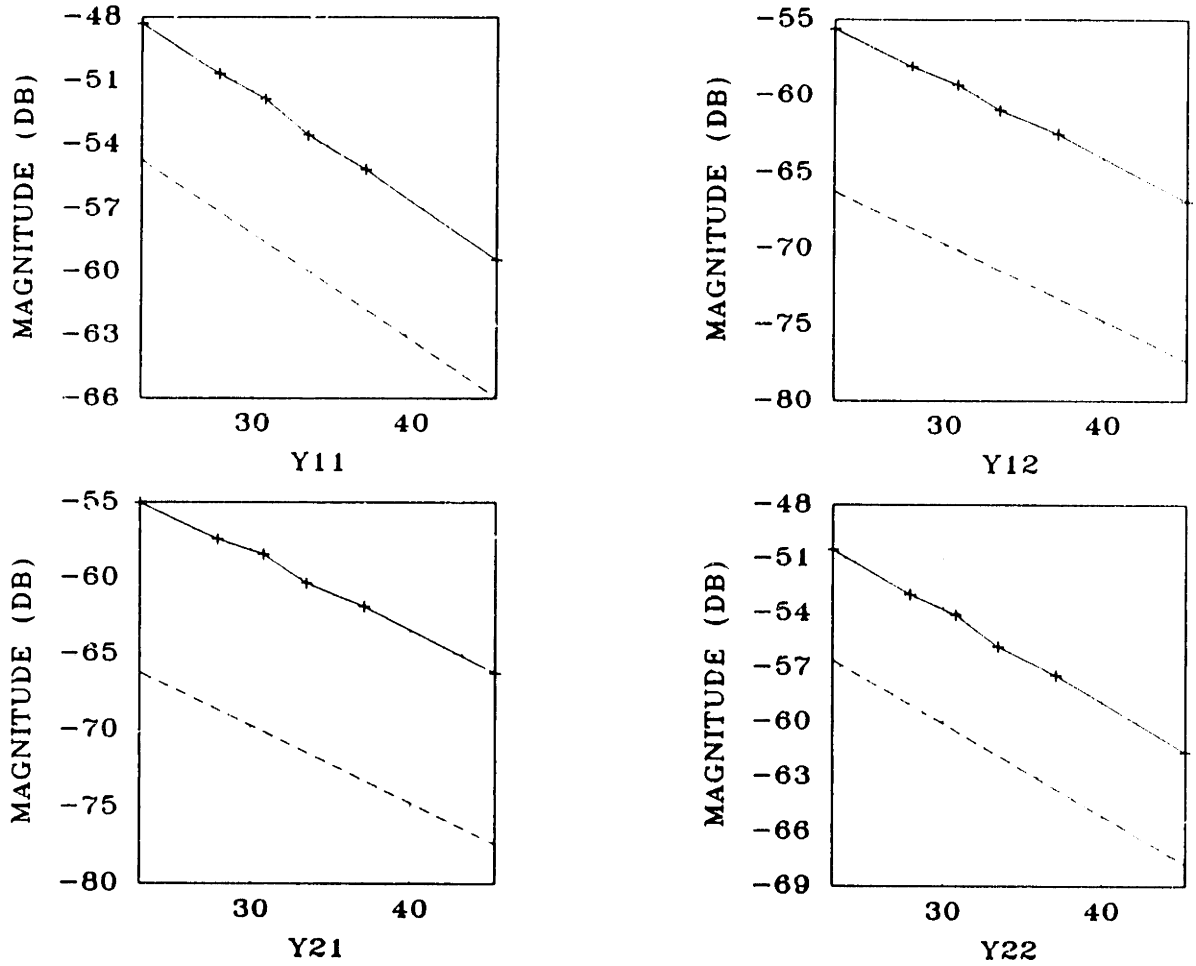

Figure 5.11: Magnitude of $Y(j\omega)$: no control data and original model. Horizontal axis has units of rad/sec.

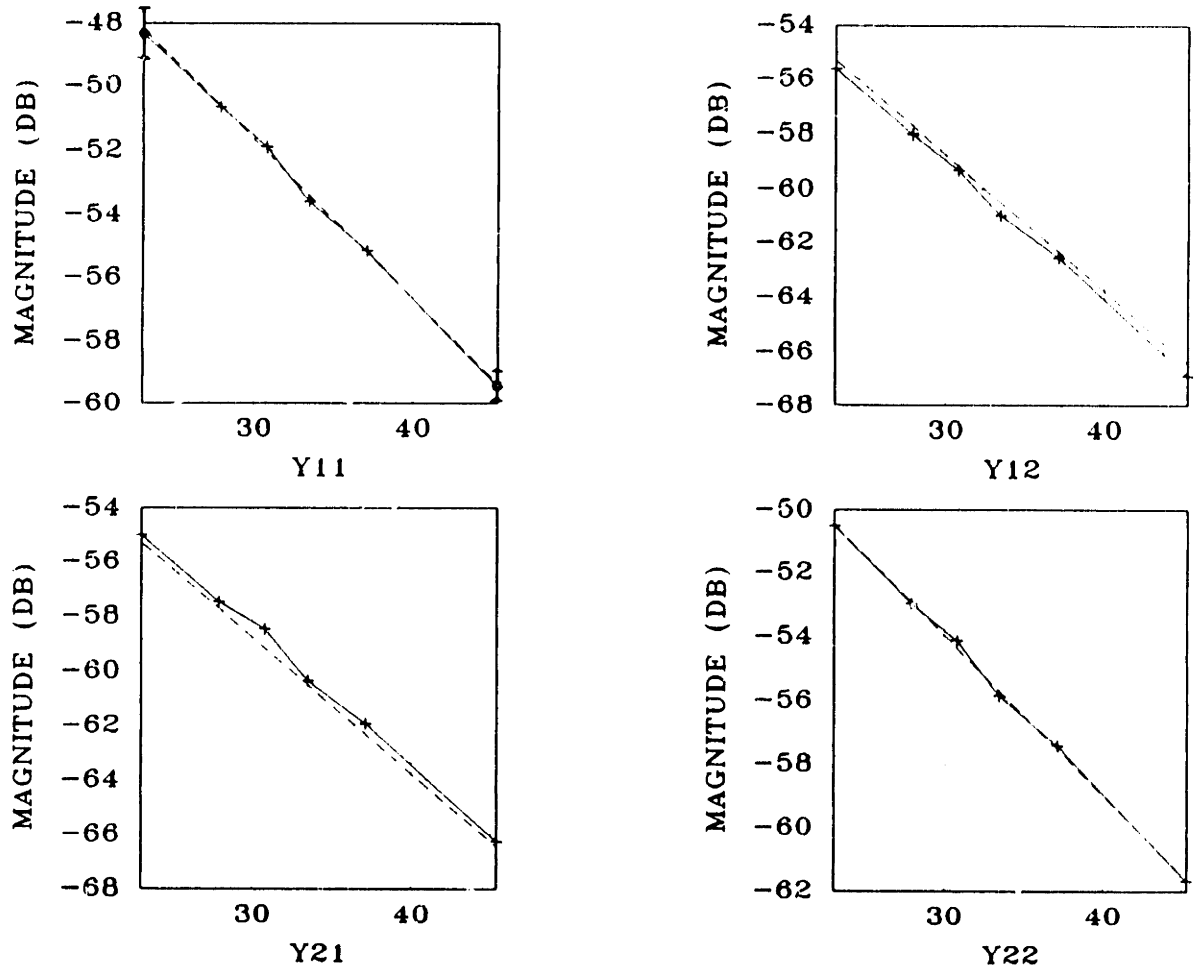

Figure 5.12: Magnitude of $Y(j\omega)$: no control data and improved model. Horizontal axis has units of rad/sec.

the following endpoint behavior (see Section 4.1.2):

$$\mathbf{F} = -K_x\mathbf{X} - B_x\mathbf{V}$$

But K_x and B_x are diagonal matrices so that the off-diagonal terms of the admittance are zero. In practice, due to the inherent damping and inertia of the manipulator, these terms are not zero, but are very small. Furthermore, if the estimated friction at the joints is in error, as is surely the case (see the previous section), then the predicted off-diagonal terms can be expected to be quite far afield. The other controllers that have been implemented all have high impedances at low frequency, and do not suffer from this problem.

Figure 5.15 is a Nyquist plot of the true driving point admittance with velocity as output. The velocity is calculated as the position times $j\omega$. The two large loops are the diagonal terms of Y. The smaller loop is the so-called "determinant term", D, which arises in the multivariable test of positive realness (Section 2.4.1):

$$\begin{aligned}
\Re D &= (\Re Y_{11})(\Re Y_{22}) - \frac{1}{4}(\Re Y_{12} + \Re Y_{21})^2 - \frac{1}{4}(\Im Y_{12} - \Im Y_{21})^2 \\
\Im D &= (\Im Y_{11})(\Im Y_{22}) - \frac{1}{4}(\Re Y_{12} - \Re Y_{21})^2 - \frac{1}{4}(\Im Y_{12} + \Im Y_{21})^2
\end{aligned}$$

In Figure 5.15, as well as in the other Nyquist plots in this chapter, D has been scaled up by a factor of 5 so that it can be easily viewed alongside the diagonal terms. A necessary and sufficient condition for passivity (Passivity 5b and following) is that $\Re D \geq 0$ for all ω and that either $\Re Y_{11} \geq 0$, or $\Re Y_{22} \geq 0$ for all ω. The Nyquist plot, therefore, is redundant as far as passivity is concerned, but the extra information can be useful in terms of predicting specific stability problems.

This particular example, however, would not be expected to have stability problems, as the Nyquist plot lies completely in the right half plane. Figure 5.16 shows samples of the interaction with springs and masses. As a position servo, this controller would not

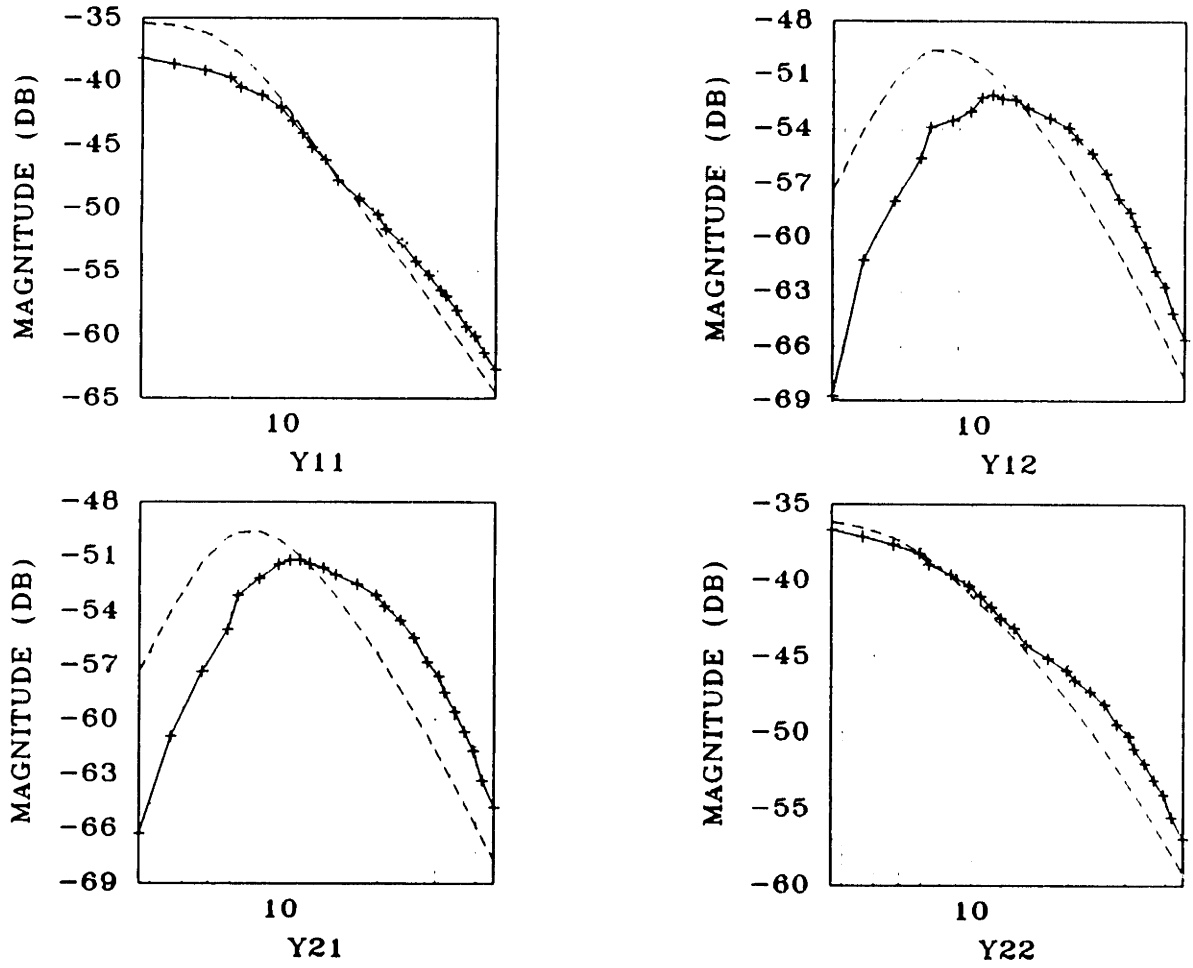

Figure 5.13: Magnitude of $Y(j\omega)$: simple impedance controller, data and model. Horizontal axis has units of rad/sec.

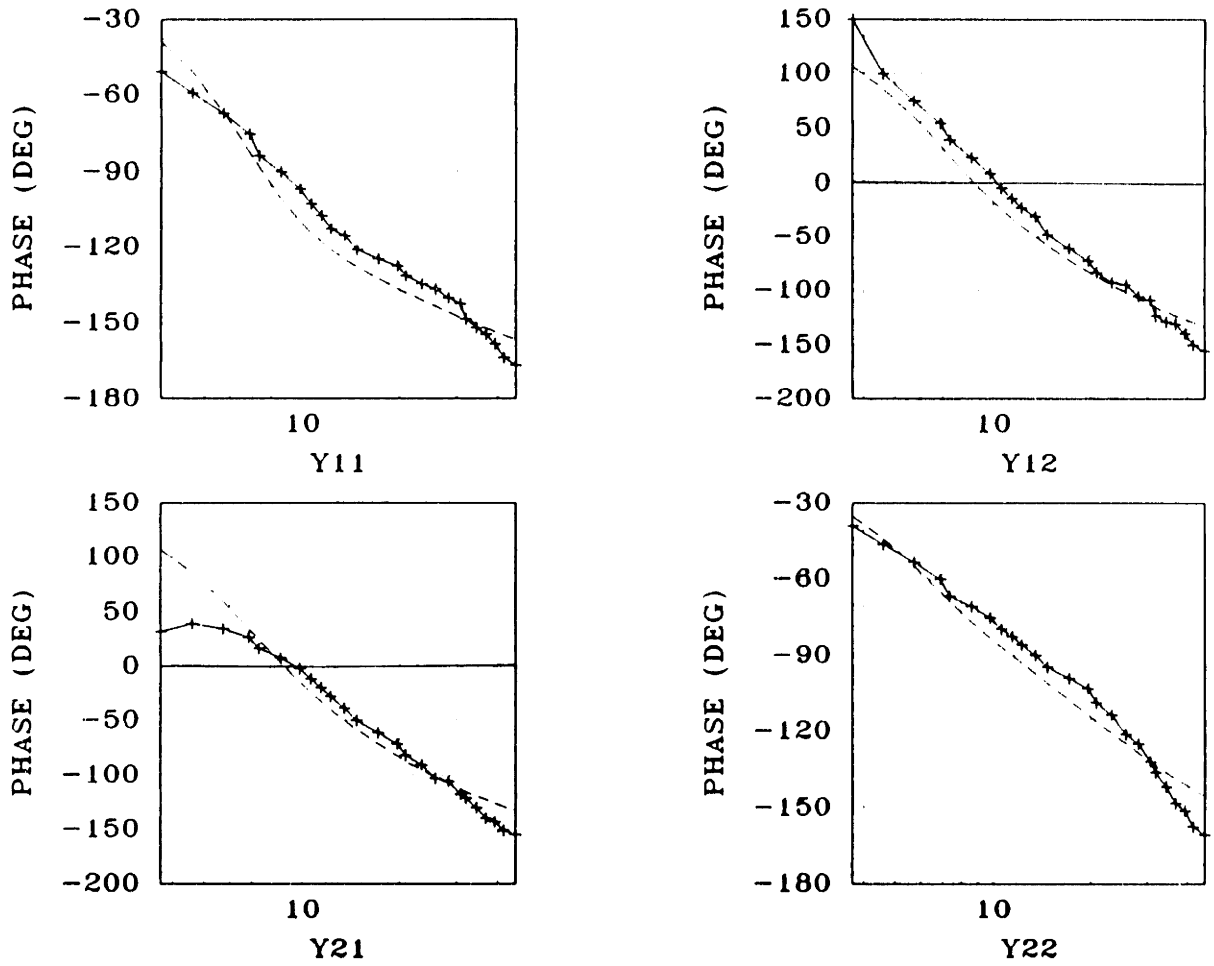

Figure 5.14: Phase of $Y(j\omega)$: simple impedance controller, data and model. Horizontal axis has units of rad/sec.

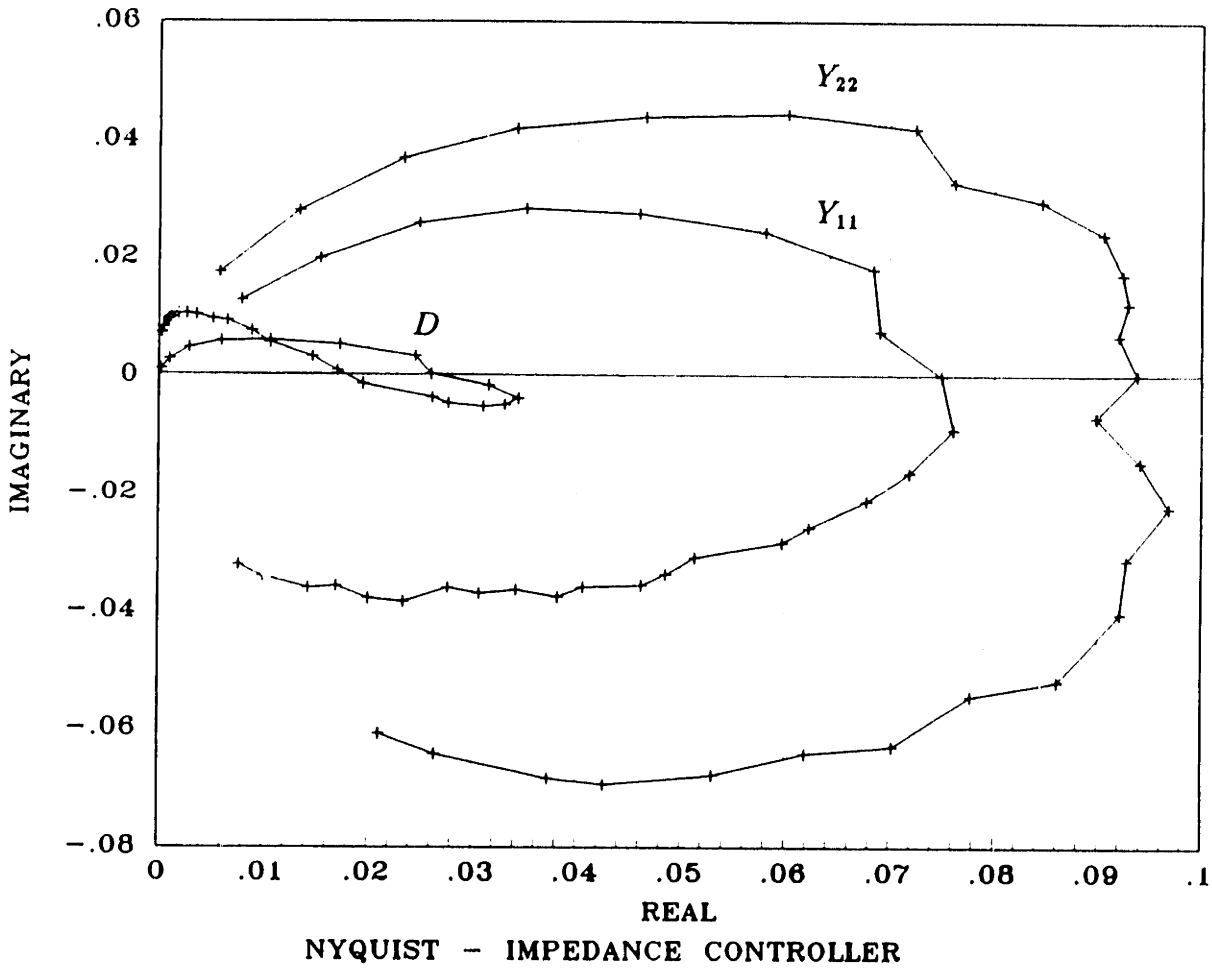

Figure 5.15: Nyquist plots of Y_{11}, Y_{22}, and D: simple impedance controller. Both axes have units of sec/kg.

fare too well[1]—it does not come close to either the 0.5 cm deflection of the spring or to tracking the circular path. However, in every case investigated, the simple impedance controller was stable, as predicted.

Contact stability was also investigated; the results are shown in Figure 5.17. As might be expected, the manipulator always retains stability upon contact with the surface, regardless of the stiffness of the surface. The force traces, however, are quite noisy, suggesting that the contact excites high frequency, underdamped dynamics.

5.4.3 PID Control

The three PID controllers described in Section 4.3.2 were investigated. Bode and Nyquist plots of the three closed loop systems are shown in Figures 5.18–5.26. In all three cases, the measured behavior is reasonably faithful to the predicted behavior. The Nyquist plots also clearly indicate the left half plane loop which was predicted in Section 4.3.2.

Of the greatest interest are Figures 5.27 and 5.28 which superimpose the Nyquist plots of the three cases. These plots may be compared to those in Figures 4.13 and 4.14. It should be abundantly clear, on the basis of these plots, that the three cases will exhibit different coupled stability properties. But Figure 5.29 shows that the uncoupled ramp responses of the three cases are essentially identical. Thus, systems with identical servo responses can exhibit different interactive behaviors.

Recall (Section 4.3.1) that the model predicts that some finite M_0 exists such that, for all $M_e > M_0$, the system composed of the manipulator coupled to M_e is unstable. Because no data exists at d.c., the data cannot predict the absence of an upper bound on M_e. The data does, however, predict some minimum mass M_0 which will lead to

[1] It should be understood that this type of impedance controller is not an inherently poor servo, but the experiments call for a rather low bandwidth implementation.

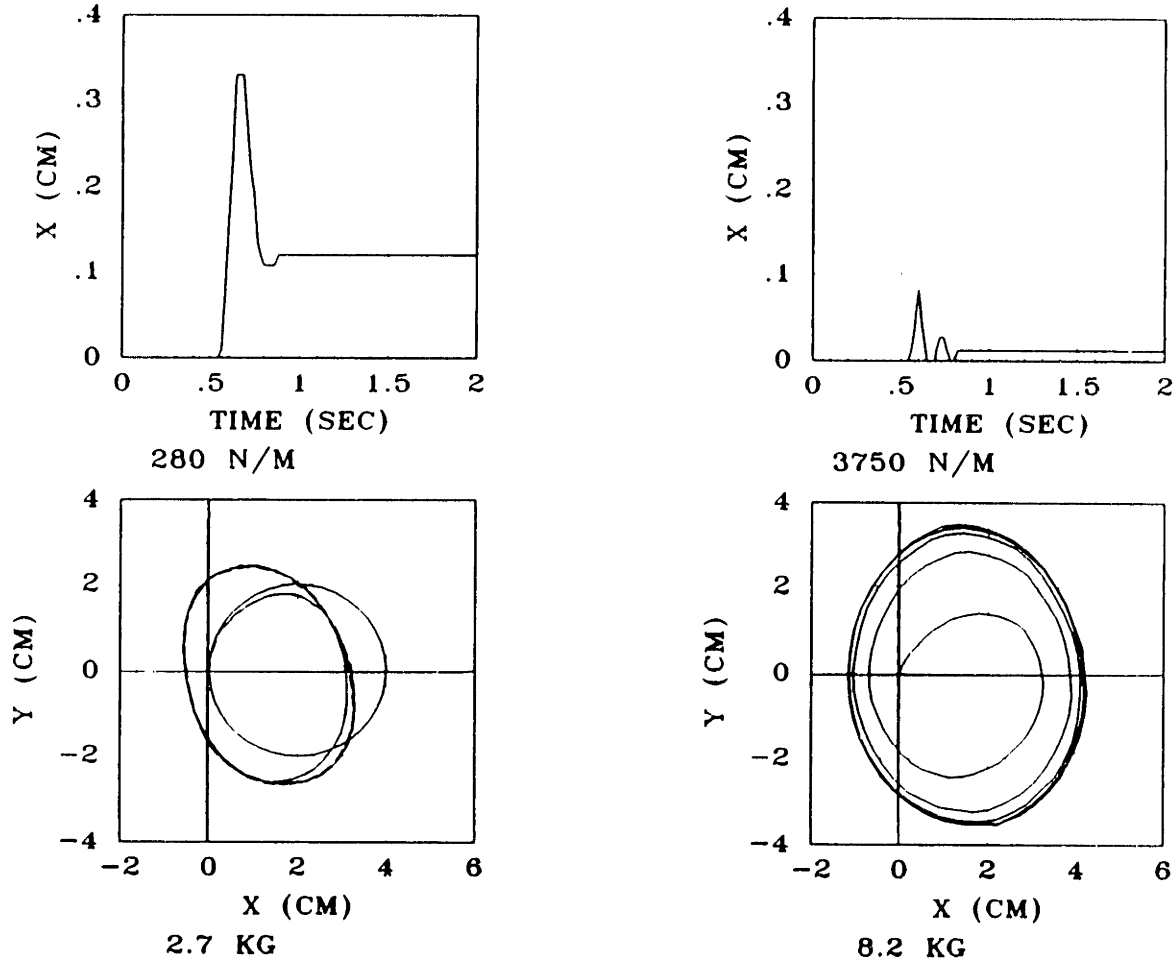

Figure 5.16: Examples of interactive behavior: simple impedance controller. Top: ramp resonses when coupled to a spring (reference trajectory is a 5 cm high ramp). Bottom: circle tracing when coupled to a mass (circular reference trajectory is shown in left plot).

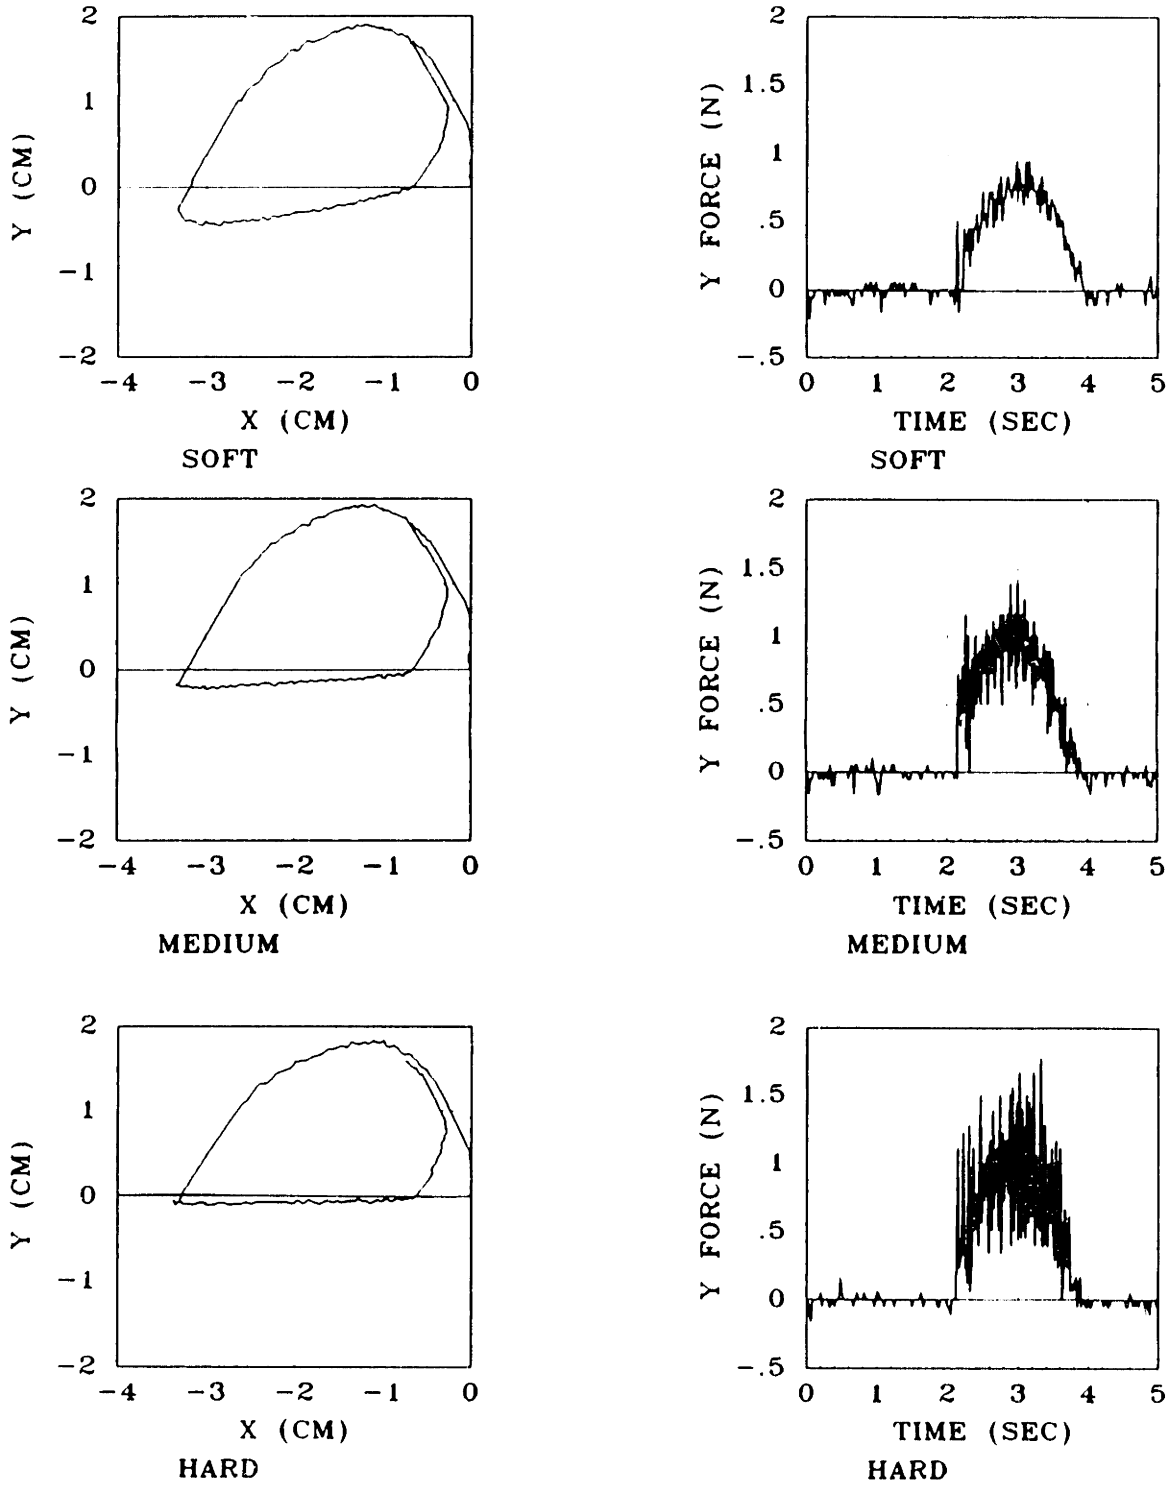

Figure 5.17: Position and force behavior during a contact task: simple impedance controller. Stiffness of soft environment is ~200–400 N/m; medium is ~1300–3000 N/m; hard is ~4000–12000 N/m.

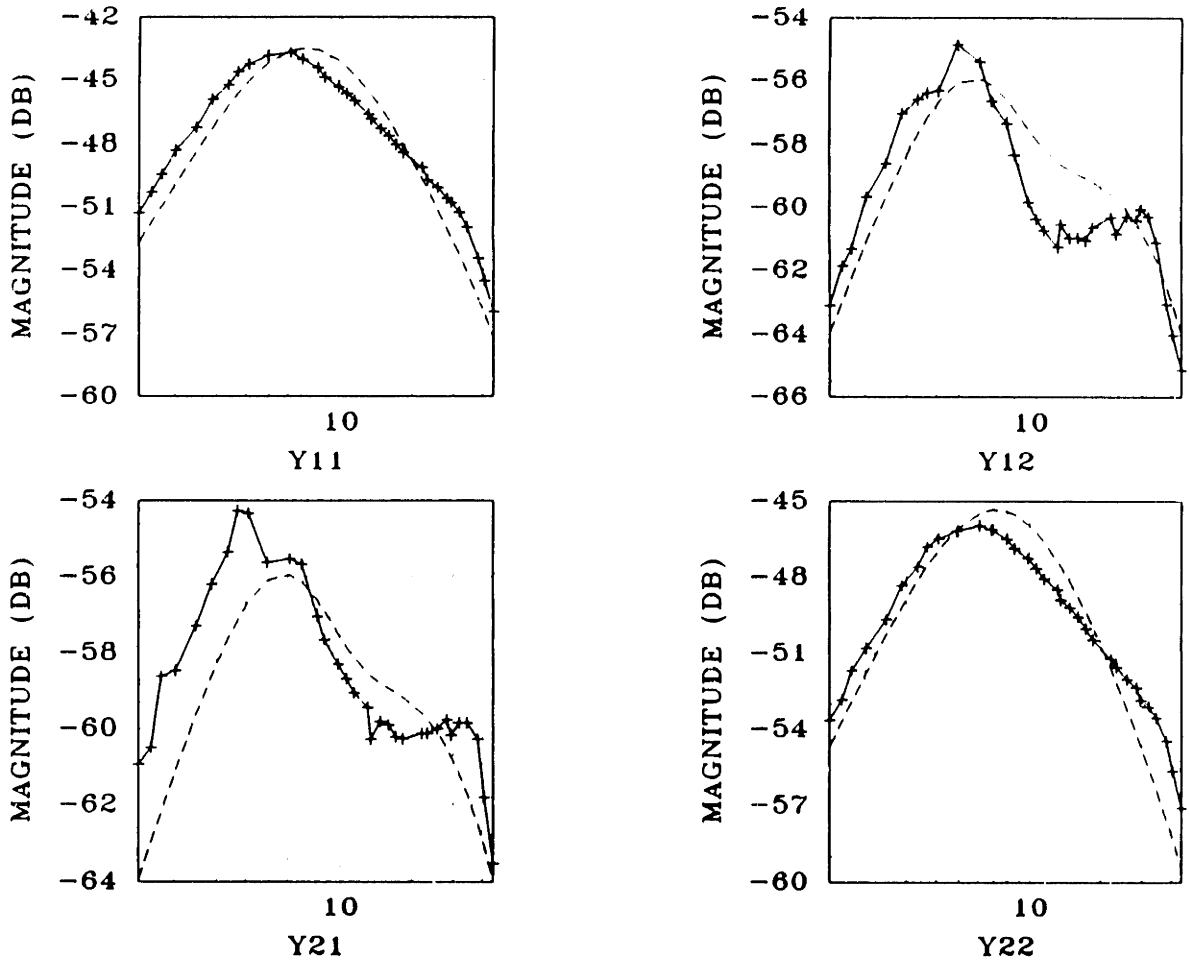

Figure 5.18: Magnitude of $Y(j\omega)$: PID controller, Case 1 ($B_1 = 2.0, \quad B_2 = 0.92$), data and model. Horizontal axis has units of rad/sec.

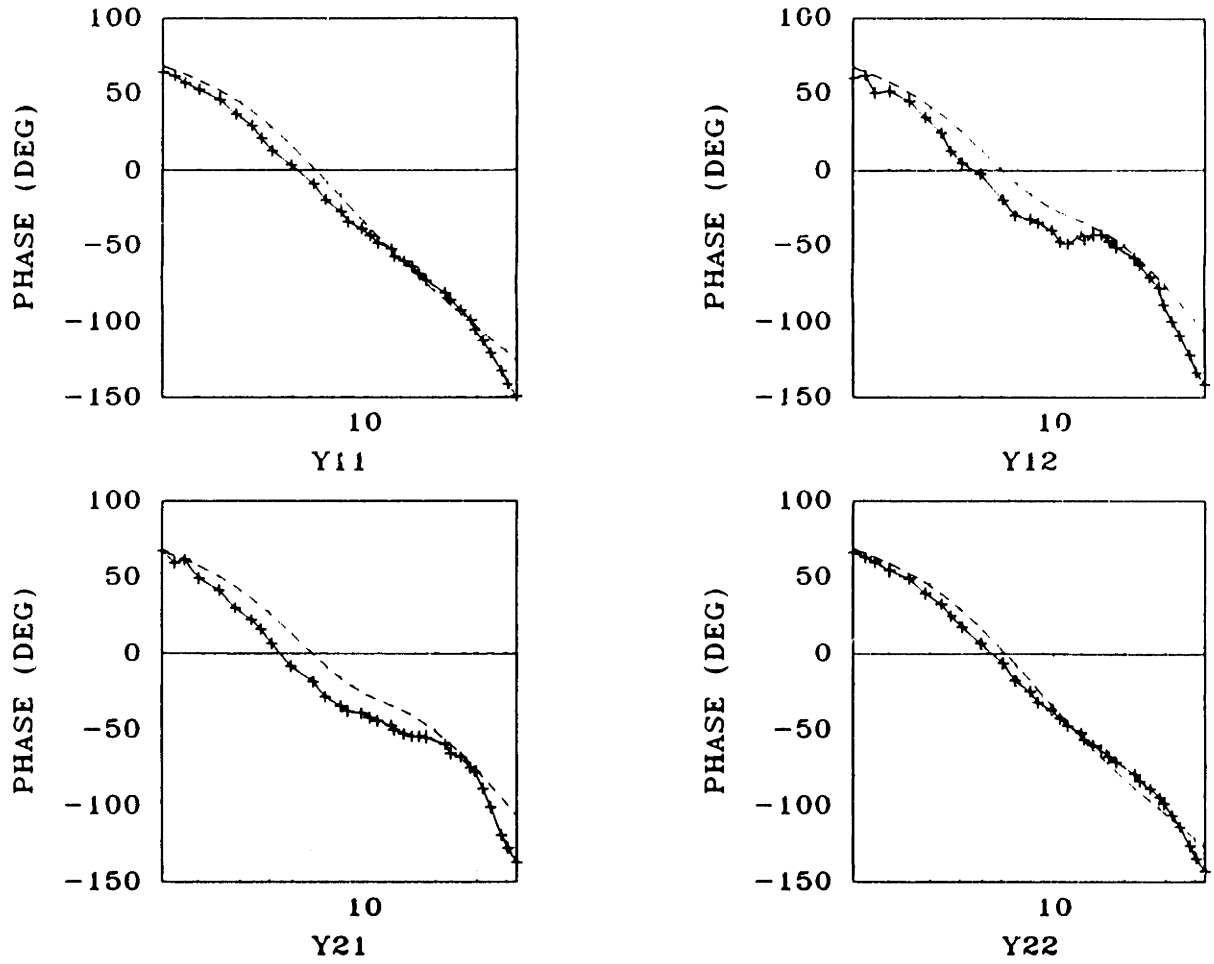

Figure 5.19: Phase of $Y(j\omega)$: PID controller, Case 1 ($B_1 = 2.0, \quad B_2 = 0.92$), data and model. Horizontal axis has units of rad/sec.

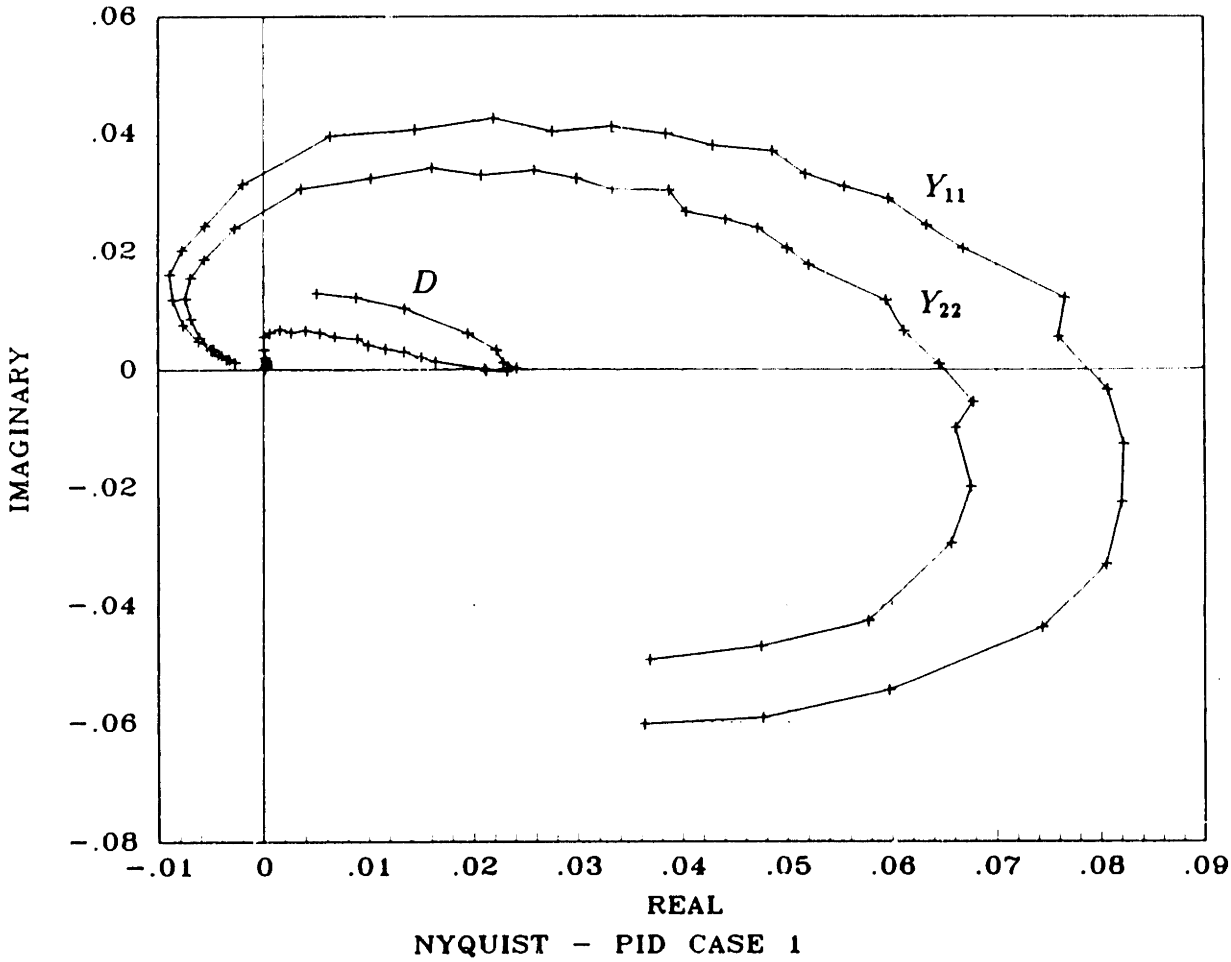

Figure 5.20: Nyquist plots of Y_{11}, Y_{22}, and D: PID controller, Case 1 ($B_1 = 2.0$, $B_2 = 0.92$). Both axes have units of sec/kg.

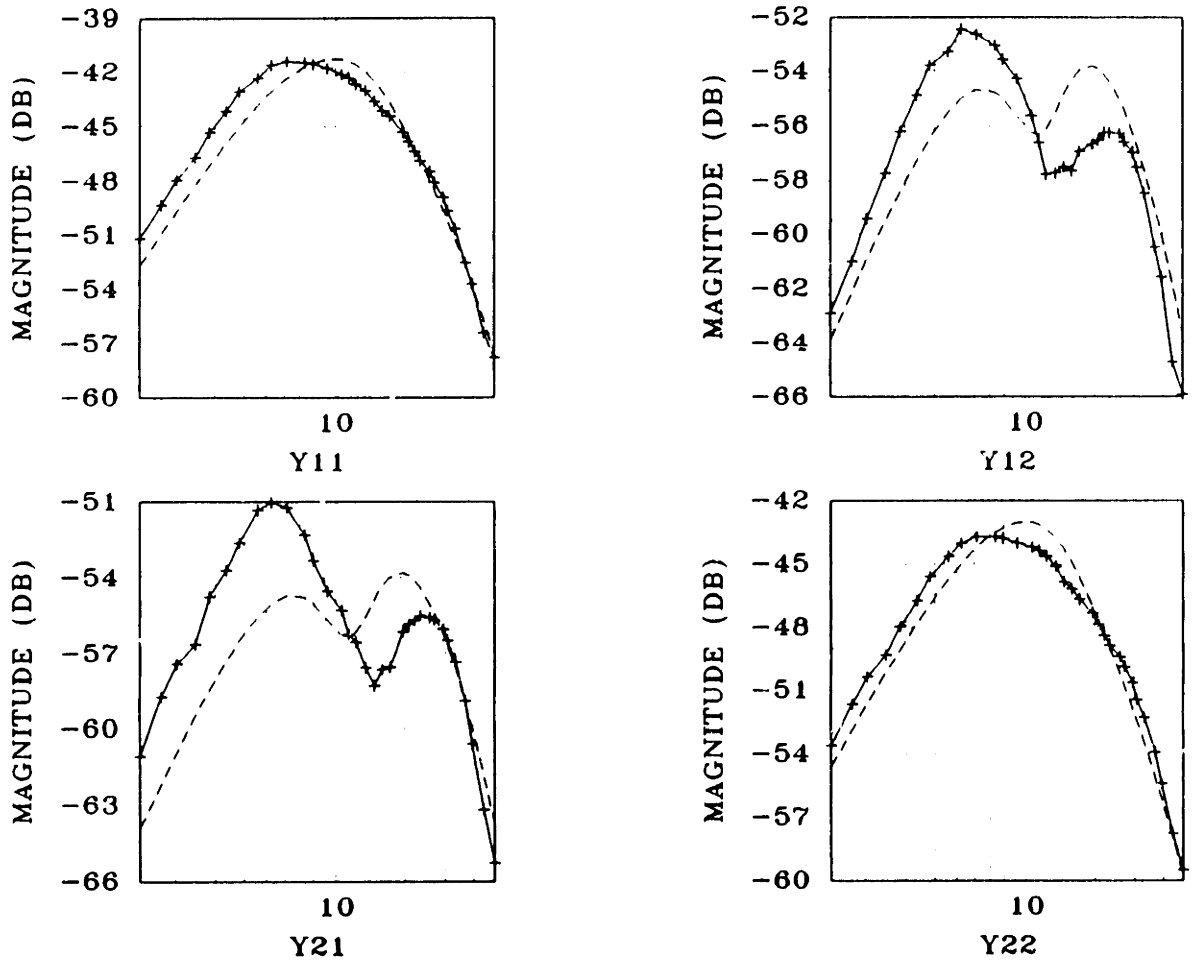

Figure 5.21: Magnitude of $Y(j\omega)$: PID controller, Case 2 ($B_1 = 1.0, \quad B_2 = 0.46$), data and model. Horizontal axis has units of rad/sec.

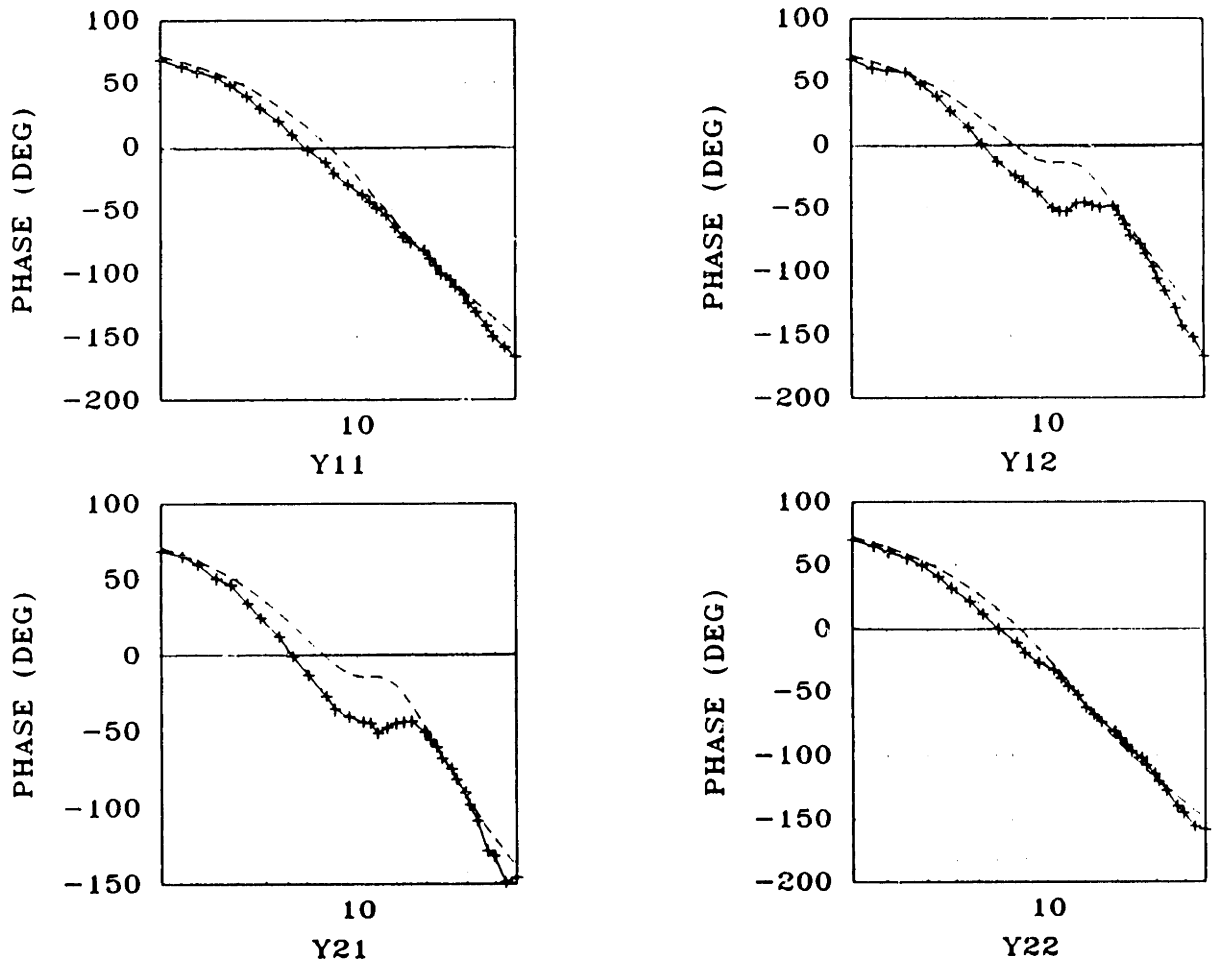

Figure 5.22: Phase of $Y(j\omega)$: PID controller, Case 2 ($B_1 = 1.0, \quad B_2 = 0.46$), data and model. Horizontal axis has units of rad/sec.

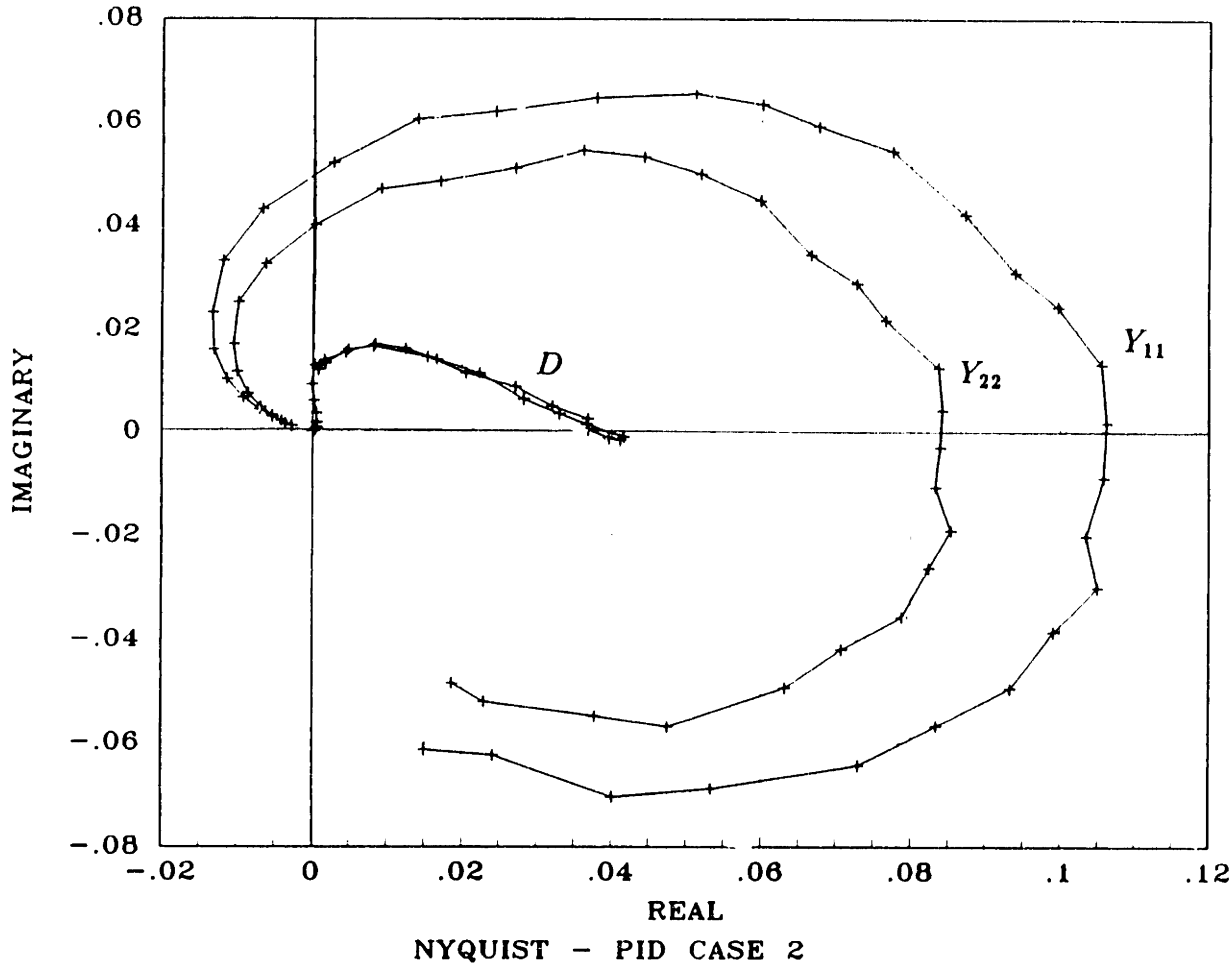

Figure 5.23: Nyquist plots of Y_{11}, Y_{22}, and D: PID controller, Case 2 ($B_1 = 1.0$, $B_2 = 0.46$). Both axes have units of sec/kg.

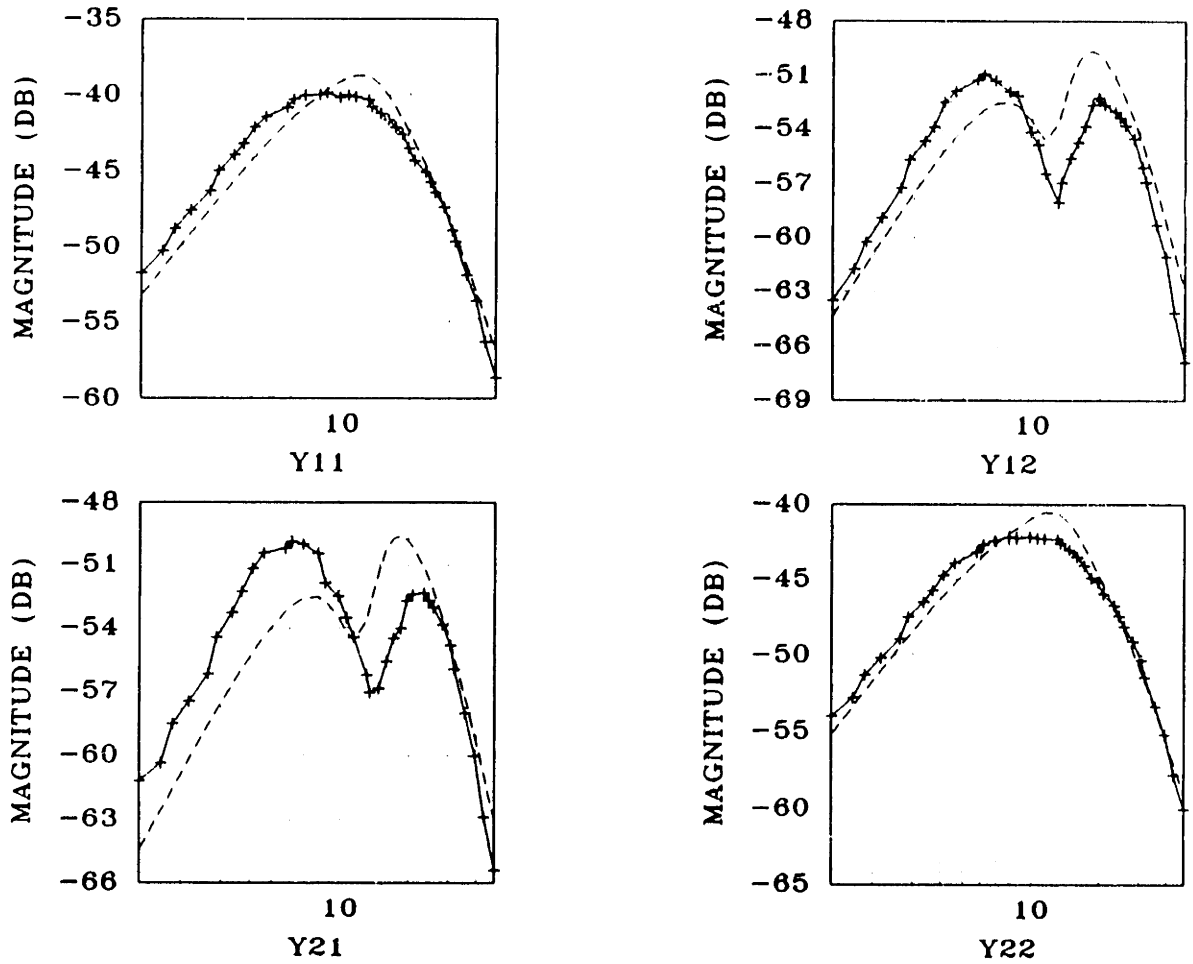

Figure 5.24: Magnitude of $Y(j\omega)$: PID controller, Case 3 ($B_1 = 0.5, \quad B_2 = 0.23$), data and model. Horizontal axis has units of rad/sec.

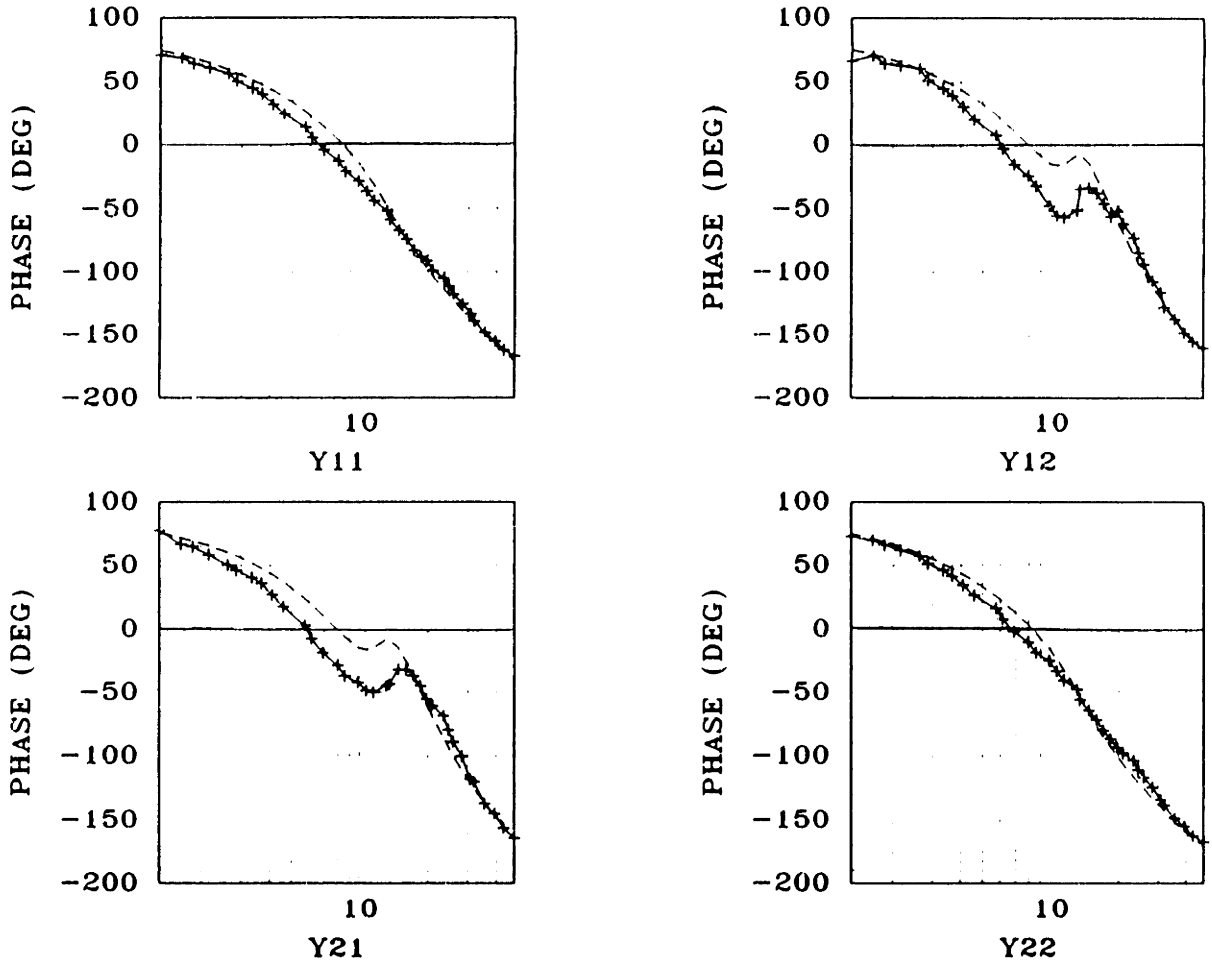

Figure 5.25: Phase of $Y(j\omega)$: PID controller, Case 3 ($B_1 = 0.5, \quad B_2 = 0.23$), data and model. Horizontal axis has units of rad/sec.

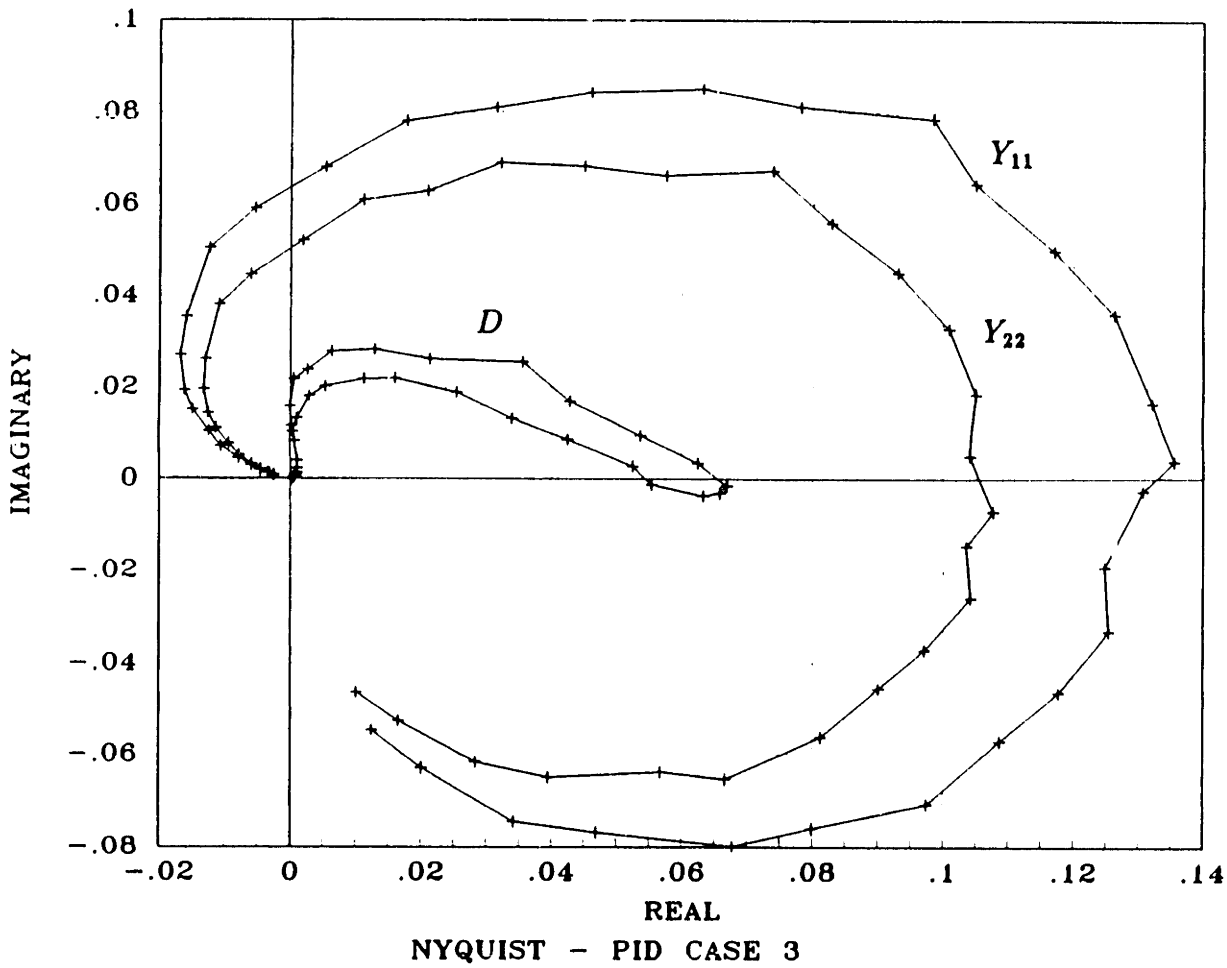

Figure 5.26: Nyquist plots of Y_{11}, Y_{22}, and D: PID controller, Case 3 ($B_1 = 0.5$, $B_2 = 0.23$). Both axes have units of sec/kg.

coupled instability. These predictions, along with the results of the mass interaction tests, are listed in Table 5.1.

Examples of the mass interaction test results are shown in Figures 5.30–5.32. The data clearly indicate instability associated with each controller as the mass increases. It was found that some motion was necessary to initiate the instability; hence, the circle test. There are two reasons for this. The first is that the position feedback is from encoders which require some finite motion, greater than the noise in the system, to produce a change in output. The second is that the presence of static friction makes what would be an unstable equilibrium point locally stable. Some motion that takes the manipulator beyond the region of local stability is therefore necessary. Because of these factors, it is difficult to accurately assess the minimum added mass which will lead to instability. Crude estimates, however, are based upon the results of the circle experiments, and are listed in Table 5.1.

One of the interesting results is that, except for Case 1, which is unclear, the masses which cause instability are less than those predicted by measurement of either Y_{11} or Y_{22}. This is not unreasonable as there is no particular reason to believe that, as mass is added, the initial instability will occur in either of these directions (x or y). In fact, the initial instability is most easily located in joint coordinates—it corresponds to an

Table 5.1: Minimum masses which cause instability, predicted from experimental Nyquist plot and observed: PID control.

	x-Direction (Predicted) (kg)	y-Direction (Predicted) (kg)	Circle Tracking (Observed) (kg)
Case 1	5.83 ($\pm$ .05)	6.86 ($+$ .05)	6.0 ($\pm$ .40)
Case 2	3.33 ($\pm$ .05)	4.17 ($\pm$ .05)	3.0 ($\pm$ .20)
Case 3	2.44 ($\pm$ .05)	3.03 ($\pm$ .05)	2.1 ($\pm$ .20)

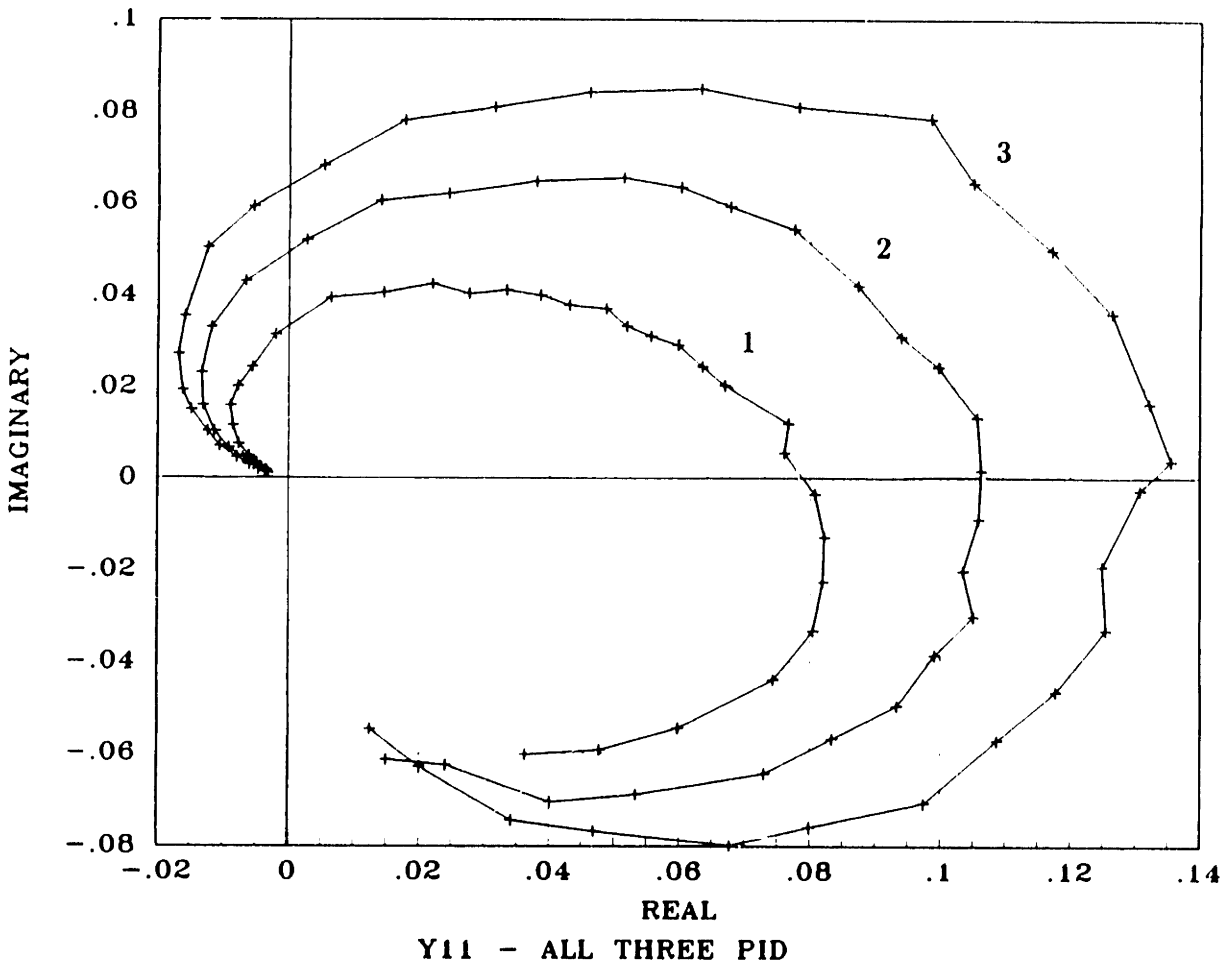

Figure 5.27: Nyquist plots of Y_{11}: PID controllers, Cases 1, 2, and 3. Both axes have units of sec/kg.

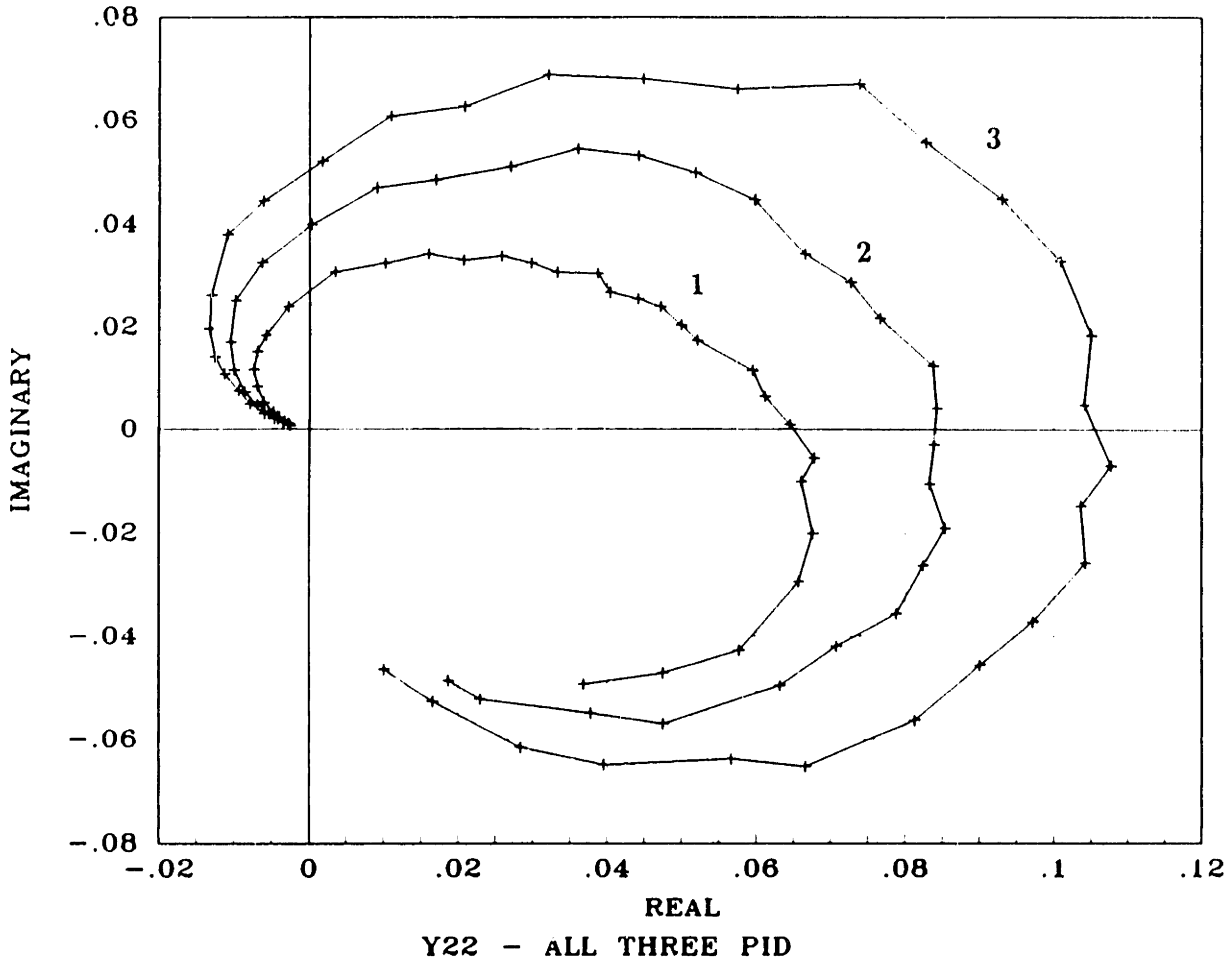

Figure 5.28: Nyquist plots of Y_{22}: PID controllers, Cases 1, 2, and 3. Both axes have units of sec/kg.

oscillation of link 2 about the elbow joint (recall that, for these experiments, the links are decoupled by selecting a nominal elbow angle of 90°). The directional quality of the instability is apparent in the data.

This result brings up an interesting point. In Section 2.4.1, two routes to checking the positive realness of an impedance Z were suggested. The first was to plot the eigenvalues of $Z + Z^H$, and the second was to plot the upper left determinants of $Z + Z^H$. The second was chosen because it is easier to compute. However, we can now see why the first method might be preferable in some instances, despite the computational expense. The method can be implemented by plotting the eigenvalues of $Z(j\omega)$, which must have positive real parts to guarantee passivity. The eigenvalues are the frequency-dependent "principle axes" of $Z(j\omega)$, and the plot may be treated just like a Nyquist plot. Therefore, the minimum and maximum masses and stiffness which lead to instability should be directly available.

Because the spring interaction tests are for x and y motions, and because the eigenvalue test does not produce a substantial difference in any of the examples considered here, we will continue to use the determinant test.

5.4.4 LQG/LTR Control

Because the LQG/LTR controller employs a model-based compensator, it is considerably more sensitive to modeling errors than the impedance or PID controllers. Thus, it is not surprising that, when corrected for inertial parameters, the Nyquist plot (Figure 5.33) has a shape which is significantly different from that predicted in the last chapter (Figure 4.16). Despite the change in shape, however, the two excursions into the left half plane still appear.

The frequency response of the LQG/LTR controller is shown in Figures 5.34 and 5.35; the Nyquist plot is shown in Figure 5.36. Comparing these plots, a major feature

Table 5.2: Minimum masses which cause instability, predicted from experimental Nyquist plot and observed: LQG/LTR control.

x-Direction (Predicted) (kg)	y-Direction (Predicted) (kg)	Circle Tracking (Observed) (kg)
2.53 ($\pm$ .06)	1.44 ($\pm$ .06)	1.4 ($\pm$ .10)

is the resonant peak which occurs at about 5.25 Hz, and shows up on the Nyquist plot as a fourth quadrant distension. The resonant peak is the consequence of inexact pole-zero cancellation.

The model makes the following predictions: there exist finite M_0, K_1, and K_2 such that, for all $M_e > M_0$ and $K_1 < K_e < K_2$, the coupled system (composed of the manipulator coupled to either the mass M_e or the spring K_e) is unstable. As described in the last section, the data cannot predict the absence of an upper bound on M_e. It also happens that the data are not collected at a high enough frequency to predict the existence of K_2. The data do, however, predict values for M_0 and K_1.

Results of the interaction tests are shown in Figures 5.37–5.39. The predictions and measurement of M_0 are summarized in Table 5.2. In the previous section it was noted that the observed instability occurred at a smaller added mass than the instability predicted by either Y_{11} or Y_{22}. This was explained by a "principle axes" argument, and it was also argued that the natural selection of principle axes was dictated by the joint geometry. In this case, however, the observed instability is very close to that predicted by Y_{22}. A priciple axis approach—diagonalizing $Y(j\omega)$—should still work, but the results can no longer be viewed naturally in terms of the joint geometry because the LQG/LTR controller is designed to match the singular values in endpoint coordinates, and, therefore, to remove the dependence upon the geometry.

The data for interaction with springs clearly support the predictions of the model—

Table 5.3: Minimum stiffnesses which cause instability, predicted from experimental Nyquist plot and observed: LQG/LTR control.

	Predicted (N/m)	Observed (N/m)
x-Direction	114	700 ($\pm$ 100)
y-Direction	70.5	420 ($\pm$ 20)

that K_1 and K_2 exist which bound the class of destabilizing springs. The predicted and measured values of K_1 for x and y motions are listed in Table 5.3. The disparity between the predictions and observations is far more than can be accounted for with an error analysis. The source of the difficulty is that the predictions assume that the environment is a pure spring. In fact, this is not true, as the mass between the strain elements and the spring must be considered a part of the environment. This mass includes part of the force transducer, as well as that of the pinch roller assembly. Accounting for this mass, the impedance of the environment is:

$$Z_e(s) = \frac{M_e s^2 + K_e}{s}$$

The stiffness K_1 which leads to instability may be found as follows:

$$K_1 = K_e = \frac{\omega}{\|Y_{ii}(\omega)\|} + M_e\omega^2$$

where ω is the frequency at which the Nyquist plot of Y_{ii} crosses the imaginary axis. Consider the behavior along the y-axis; $\omega = 37.7$ rad/sec, and $\|Y_{yy}\| = 0.535$ sec/kg, so that:

$$K_e = 70.5 + 1421 M_e \text{ N/m}$$

where M_e is measured in kg. Given that the mass contributed by the force transducer alone is expected to be over 0.1 kg, the effect of this additional term is significant.

M_e, however, is not known very accurately. But, in the y-direction we are fortunate to have a case of spring interaction which exhibits marginal stability, and therefore

provides a fairly accurate estimate of K_e (420 N/m). The approach that has been taken, therefore, is to use this value of K_e to solve for M_e. The result is $M_e = 0.25$ kg, which is quite plausible.

Figure 5.40 is an indication of the important role that static friction can play in quasistatic tasks. It was previously indicated that some motion is necessary to excite any instabilities. However, it turns out that, at least in the case of spring interaction, the motion may need to exceed some finite speed to excite the instabilities. Figure 5.40 shows two ramp responses, each for the same system of manipulator coupled to a spring. One of the ramps has a slope of 10 cm/sec, the other 0.5 cm/sec. The faster one excites the instability, but the slower one does not *although it successfully executes the move.*

In addition to the coupled stability tests, contact stability was investigated as described in Section 5.3.2. Representative data are shown in Figure 5.41. The essential result is unchanged. Interaction with very compliant or very stiff environments is stable, whereas there is a range of values for K_e in which the interaction is violently unstable, reminiscent of force feedback induced contact instability.

It is interesting that a correlation between contact stability and coupled stability exists in both the impedance control and LQG/LTR control implementations. Although this result is only qualitative, it lends credence to the result of Section 3.6, that coupled stability and instability are sufficient to predict contact stability and instability. This is a useful result, because it allow engineers to perform a rather straightforward coupled stability analysis in order to understand something about the more complicated contact behavior.

5.5 Summary

The major points presented in this chapter may be summarized as follows:

- A measure of the driving point admittance is an effective tool for predicting coupled stability properties.
- It is possible to design feedback systems which are stable when coupled to a broad class of passive environments. The simple impedance controller is an example.
- Integral control of joint angles leads to instability upon coupling to sufficiently large masses.
- Systems which exhibit identical servo response need not have identical coupled stability properties. This is illustrated by the three cases of PID control.
- Manipulator control schemes which do not employ force feedback may nevertheless exhibit contact instability. The LQG/LTR controller is an example.
- The experiments support the theoretical result of Section 3.6, that coupled stability is a necessary and sufficient condition for contact stability.

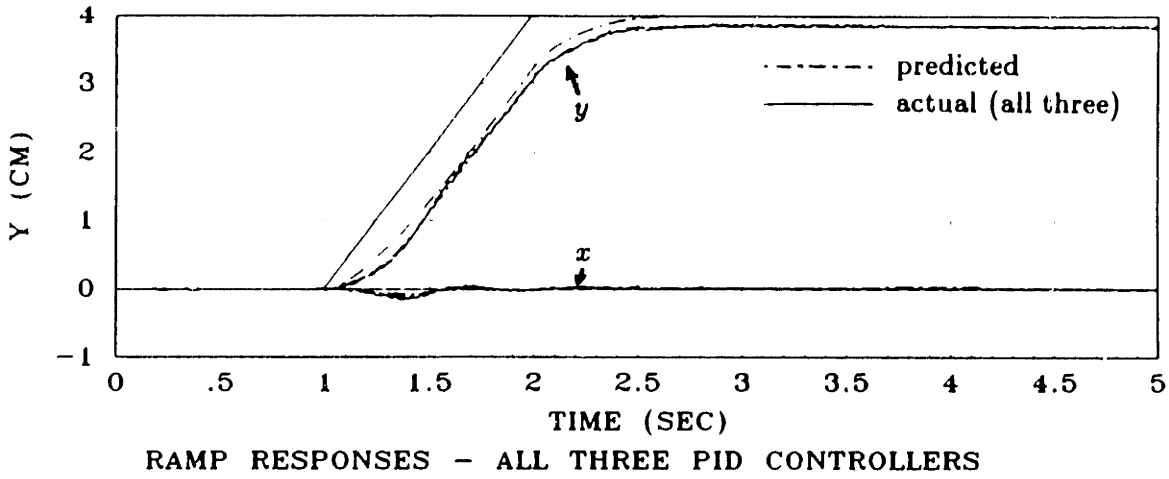

Figure 5.29: Analytically predicted and actual ramp responses of the three PID controllers. The responses of the three controllers lie directly on top of one another. Differences from the predicted response are due to friction (delay) and kinematic errors caused by the use of a fixed jacobian matrix for the entire move (apparent steady state error).

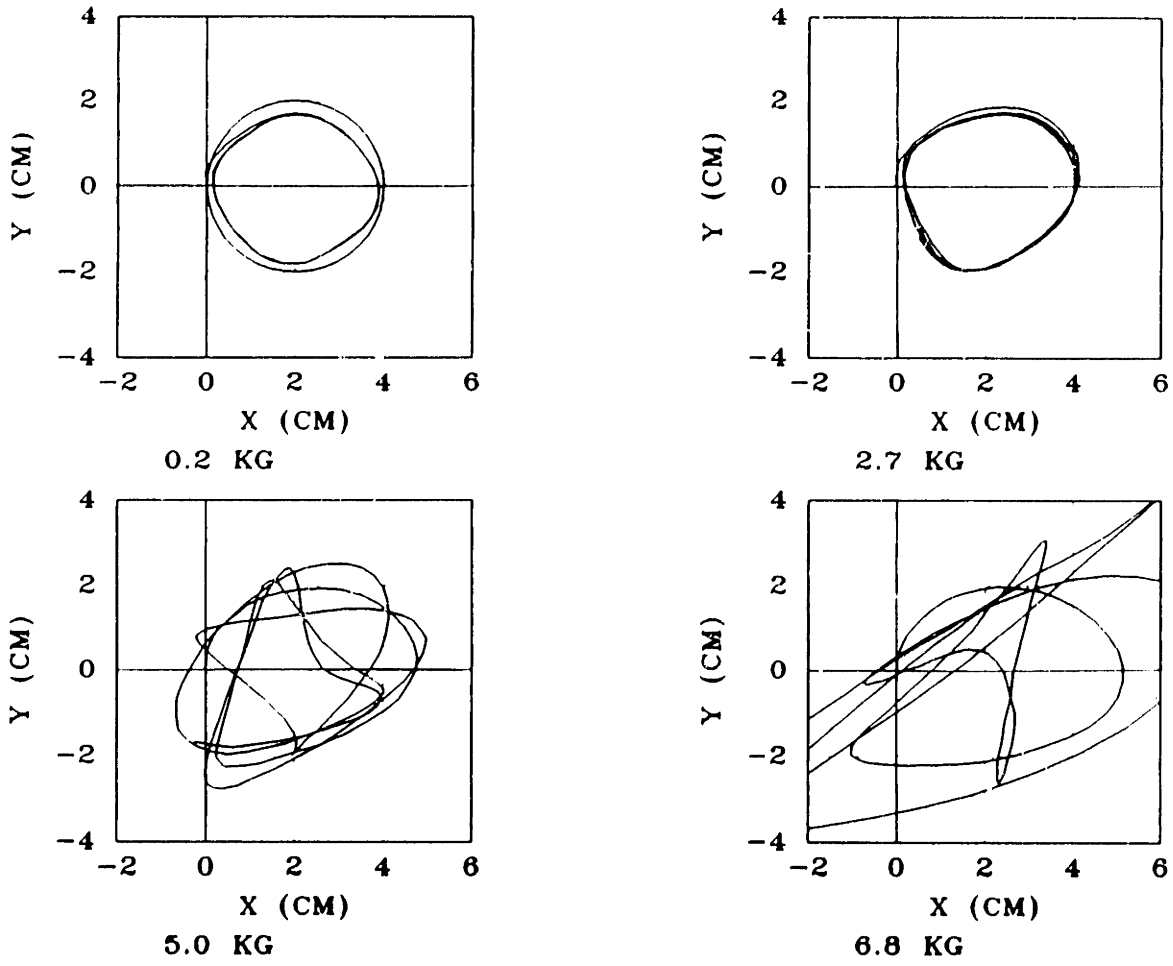

Figure 5.30: Circle tracking performance when coupled to masses: PID control, Case 1: $B_1 = 2.0, \quad B_2 = 0.92$.

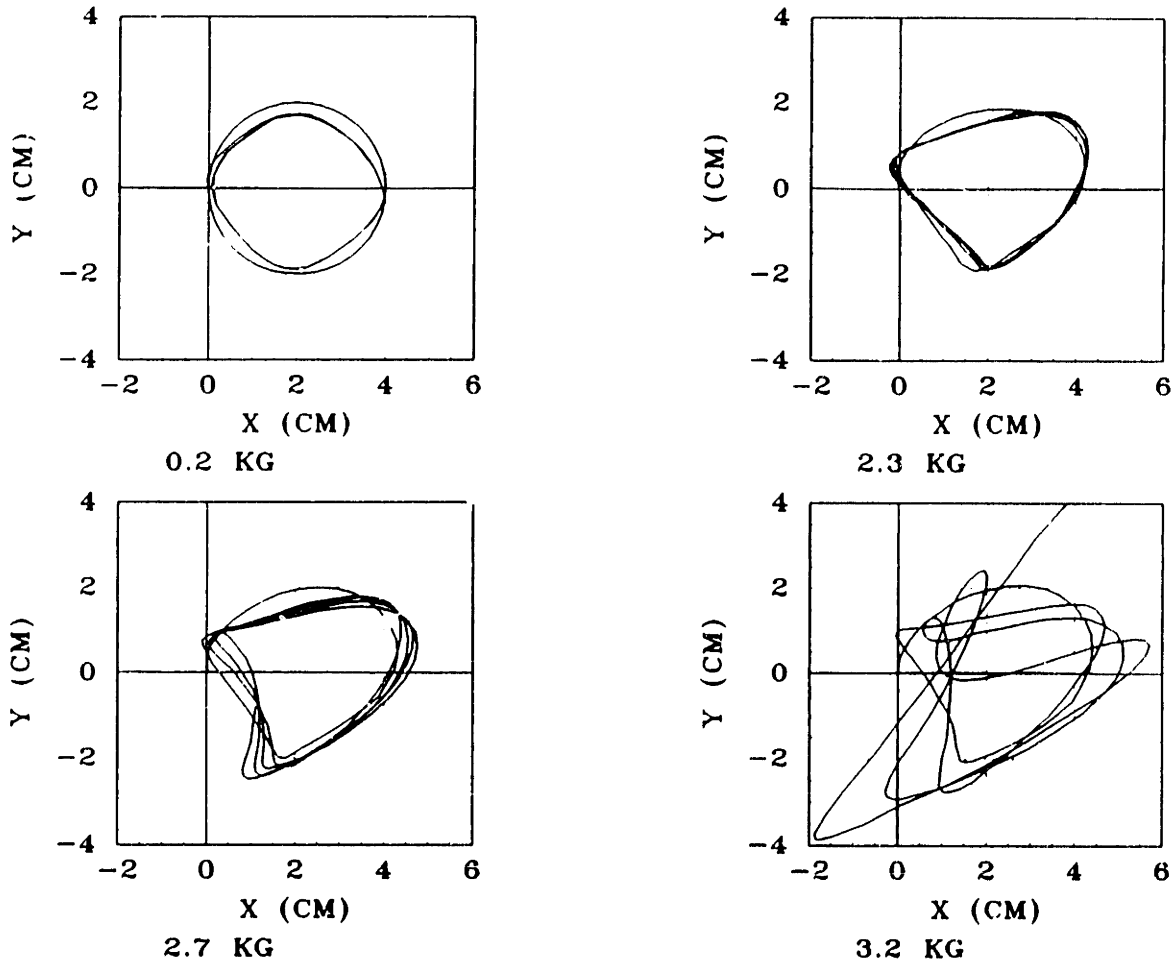

Figure 5.31: Circle tracking performance when coupled to masses: PID control, Case 2: $B_1 = 1.0, \quad B_2 = 0.46$.

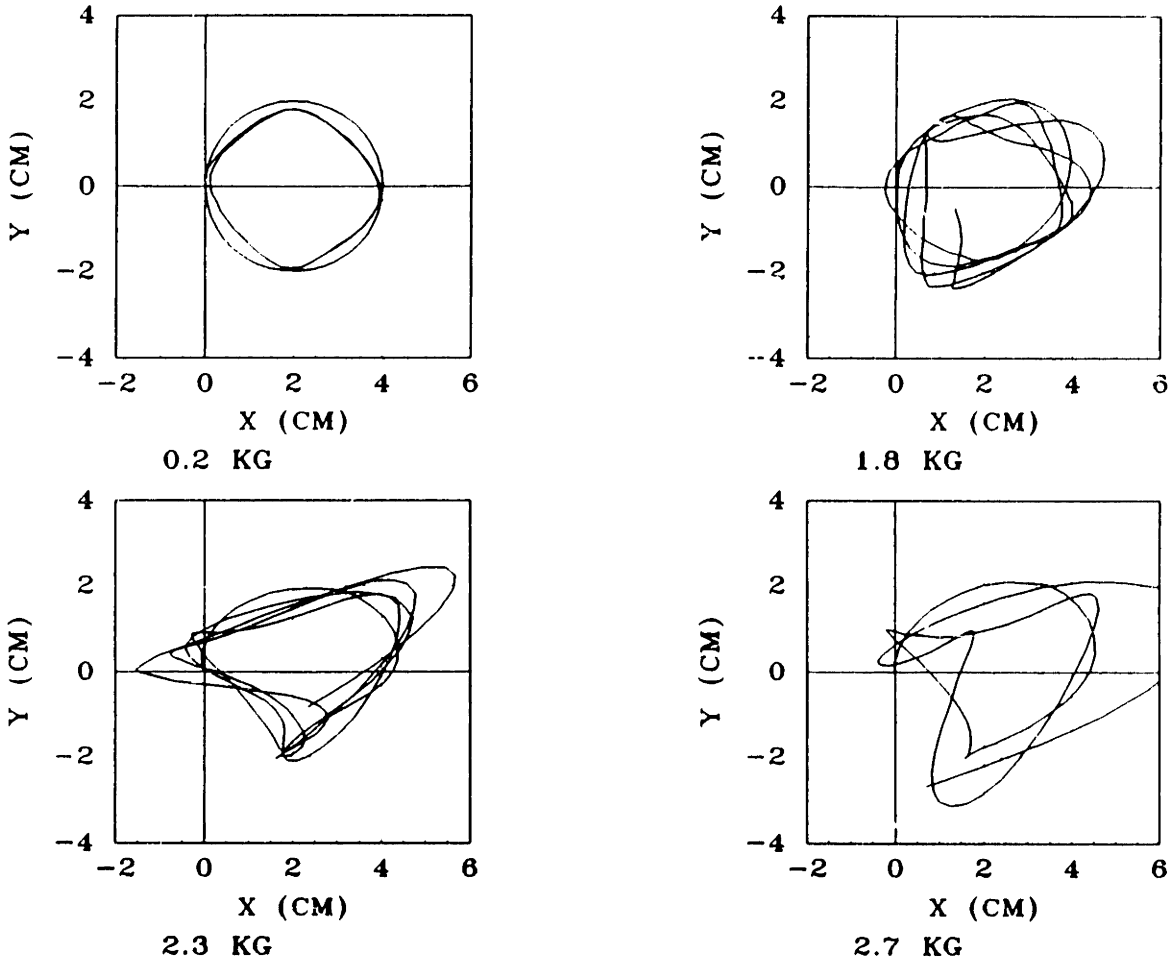

Figure 5.32: Circle tracking performance when coupled to masses: PID control, Case 3: $B_1 = 0.5, \quad B_2 = 0.23$.

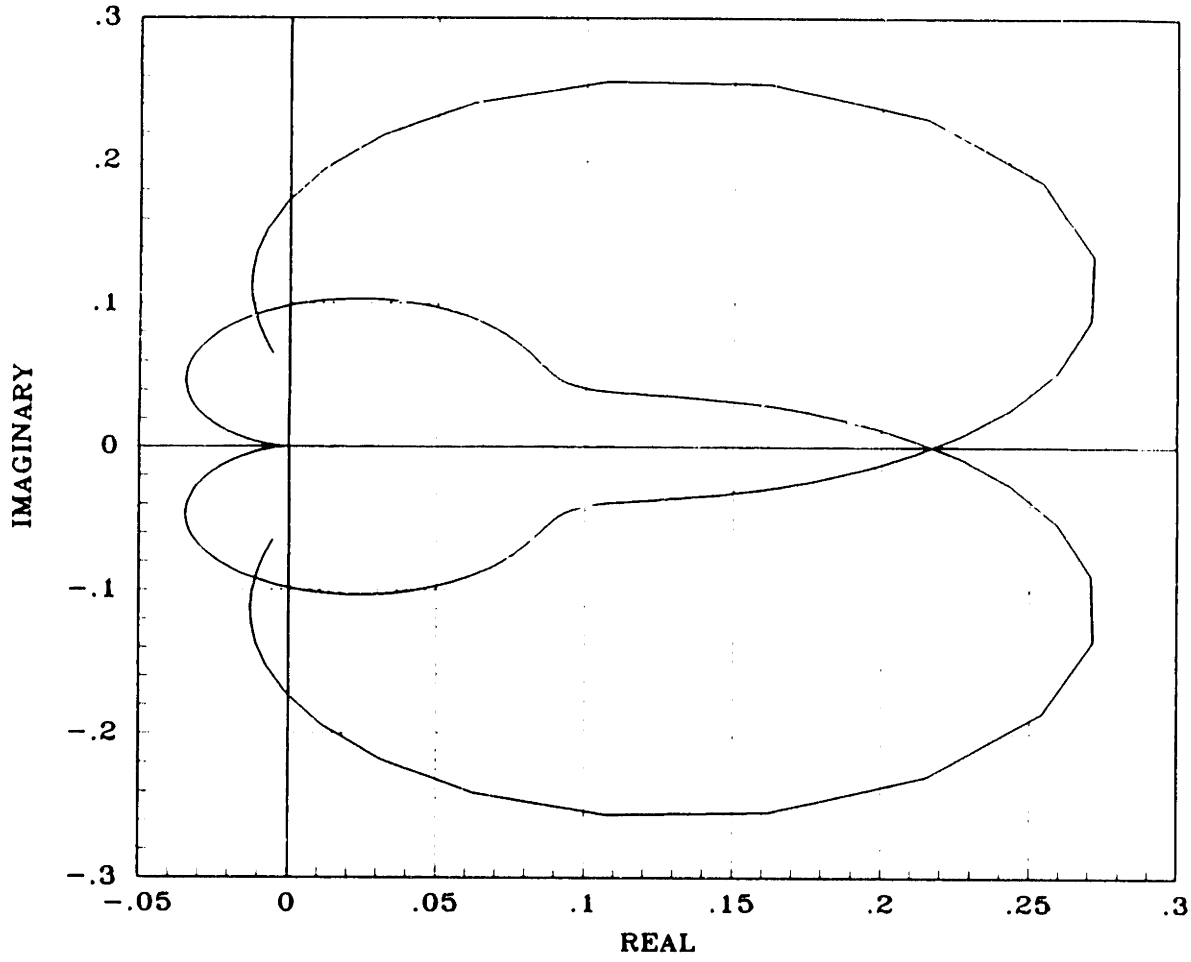

Figure 5.33: Nyquist plot of $Y_{22}(s)$, corrected inertial parameters: LQG/LTR controller. Both axes have units of sec/kg.

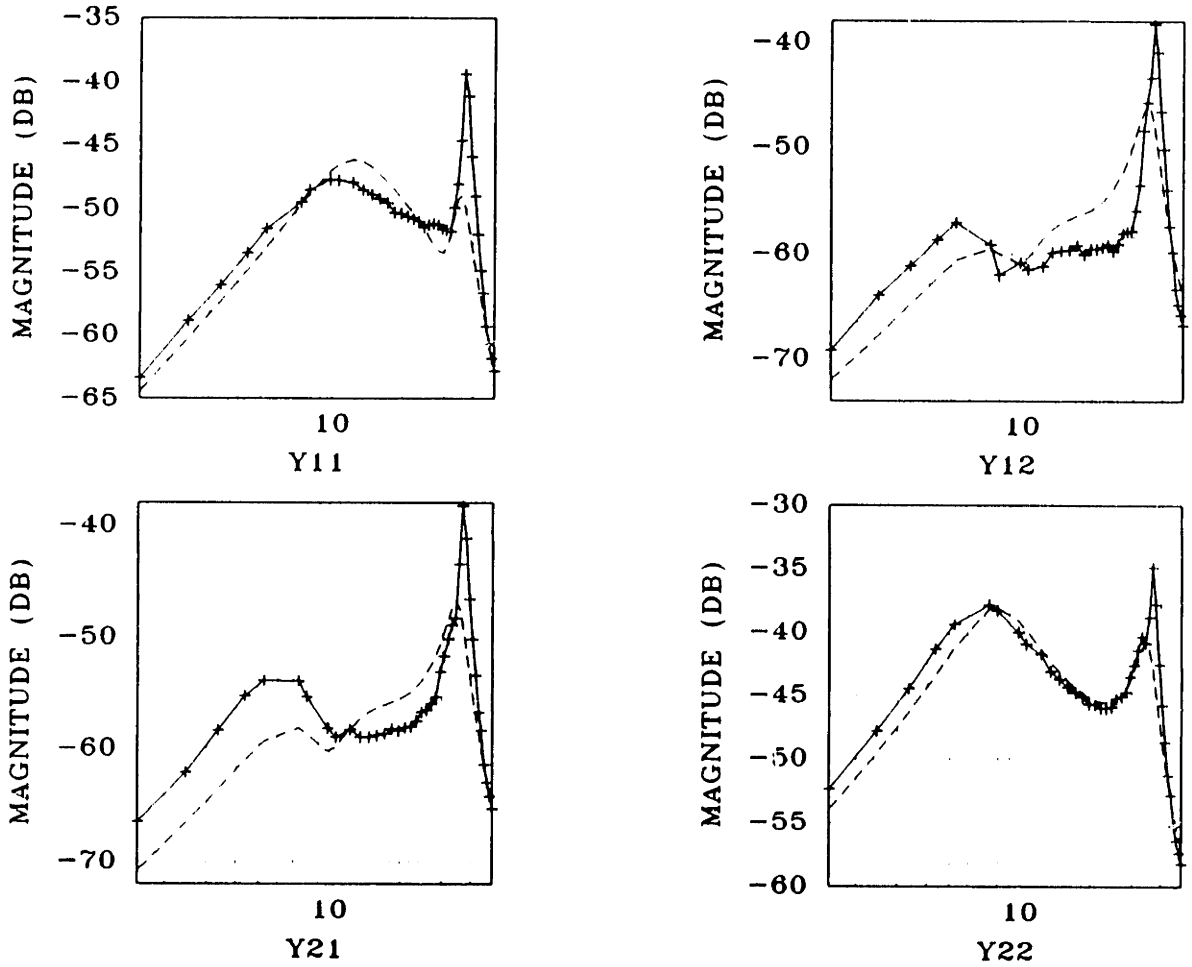

Figure 5.34: Magnitude of $Y(j\omega)$: LQG/LTR controller, data and model. Horizontal axis has units of rad/sec.

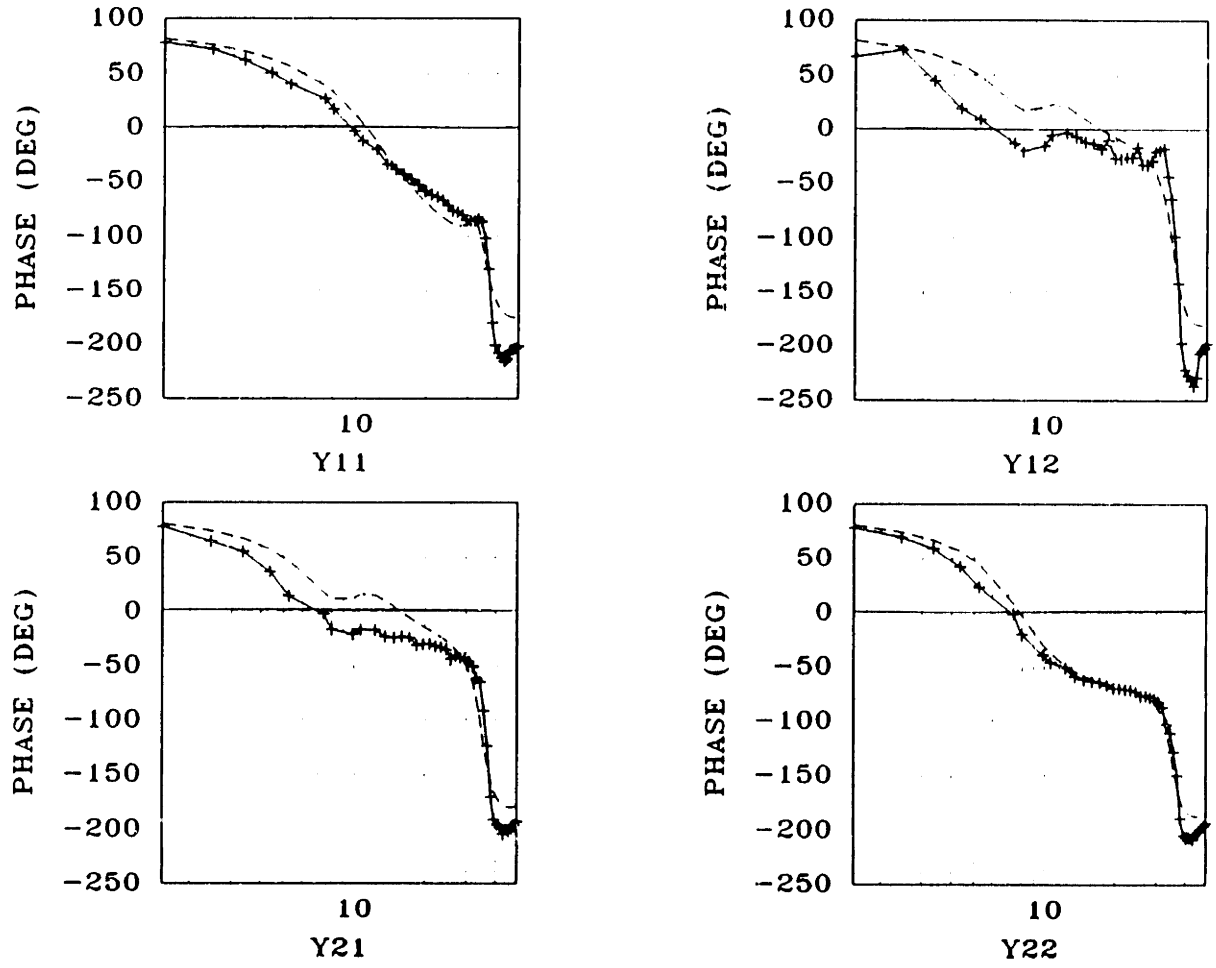

Figure 5.35: Phase of $Y(j\omega)$: LQG/LTR controller, data and model. Horizontal axis has the units of rad/sec.

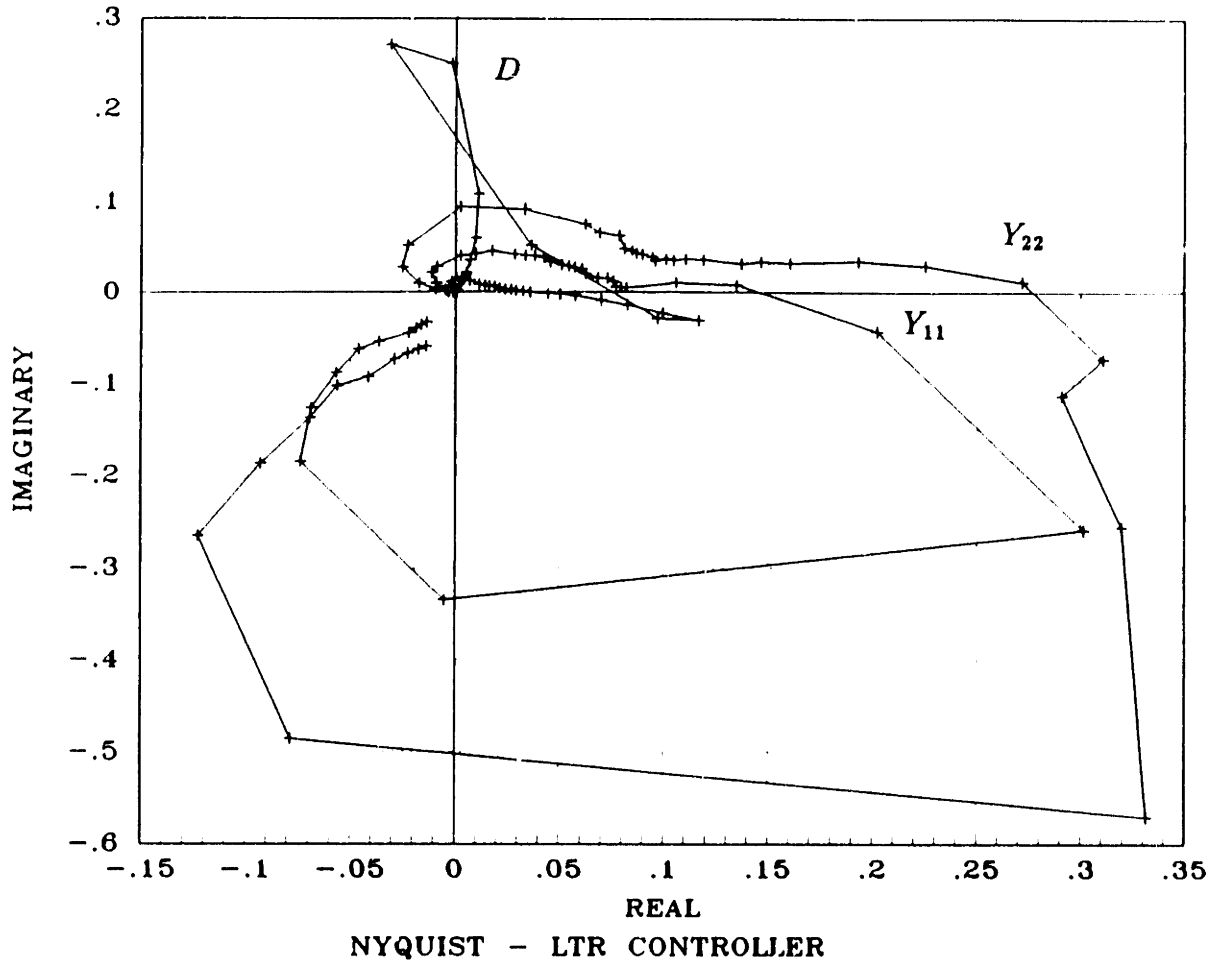

Figure 5.36: Nyquist plots of Y_{11}, Y_{22}, and D: LQG/LTR controller. Both axes have units of sec/kg.

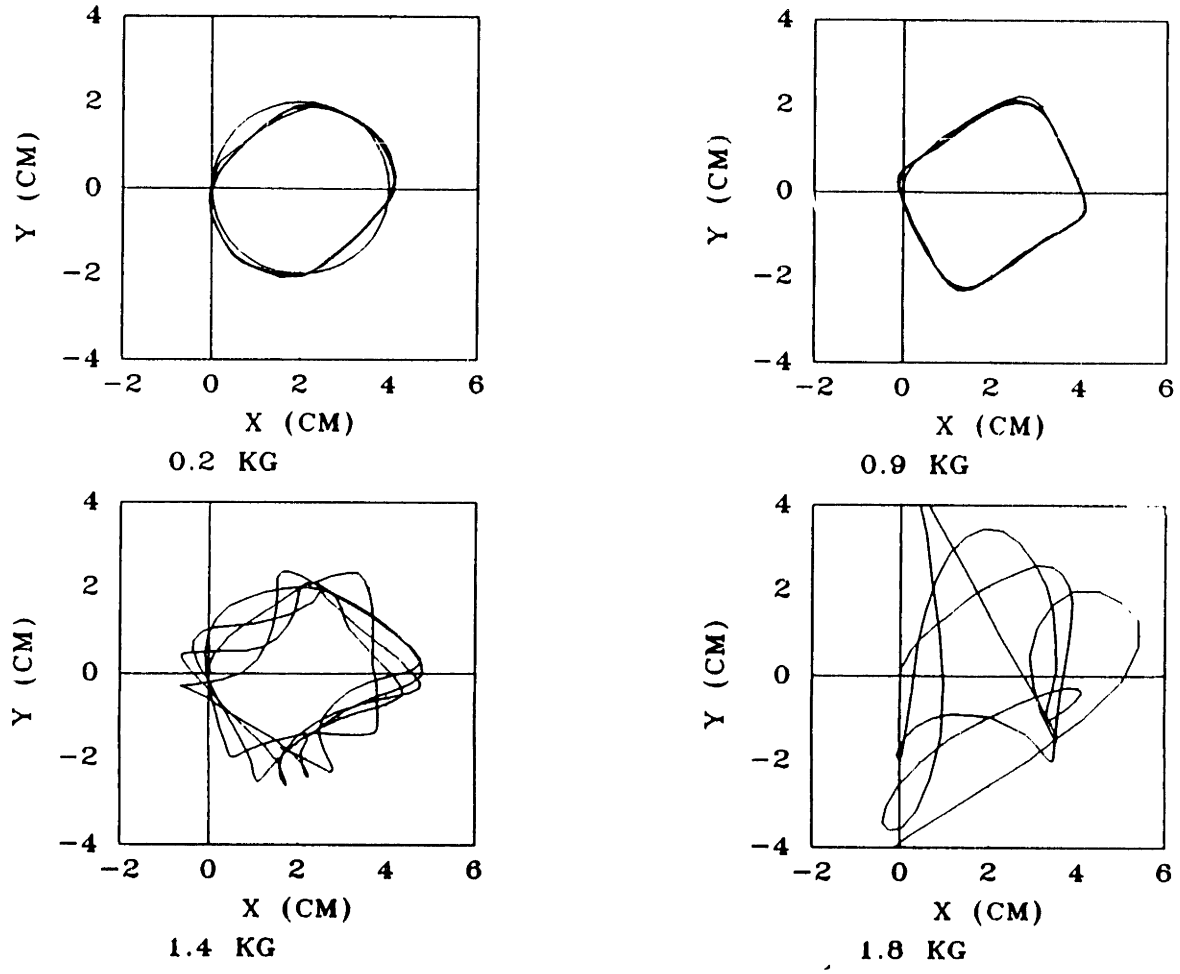

Figure 5.37: Circle tracking performance when coupled tc masses: LQG/LTR control.

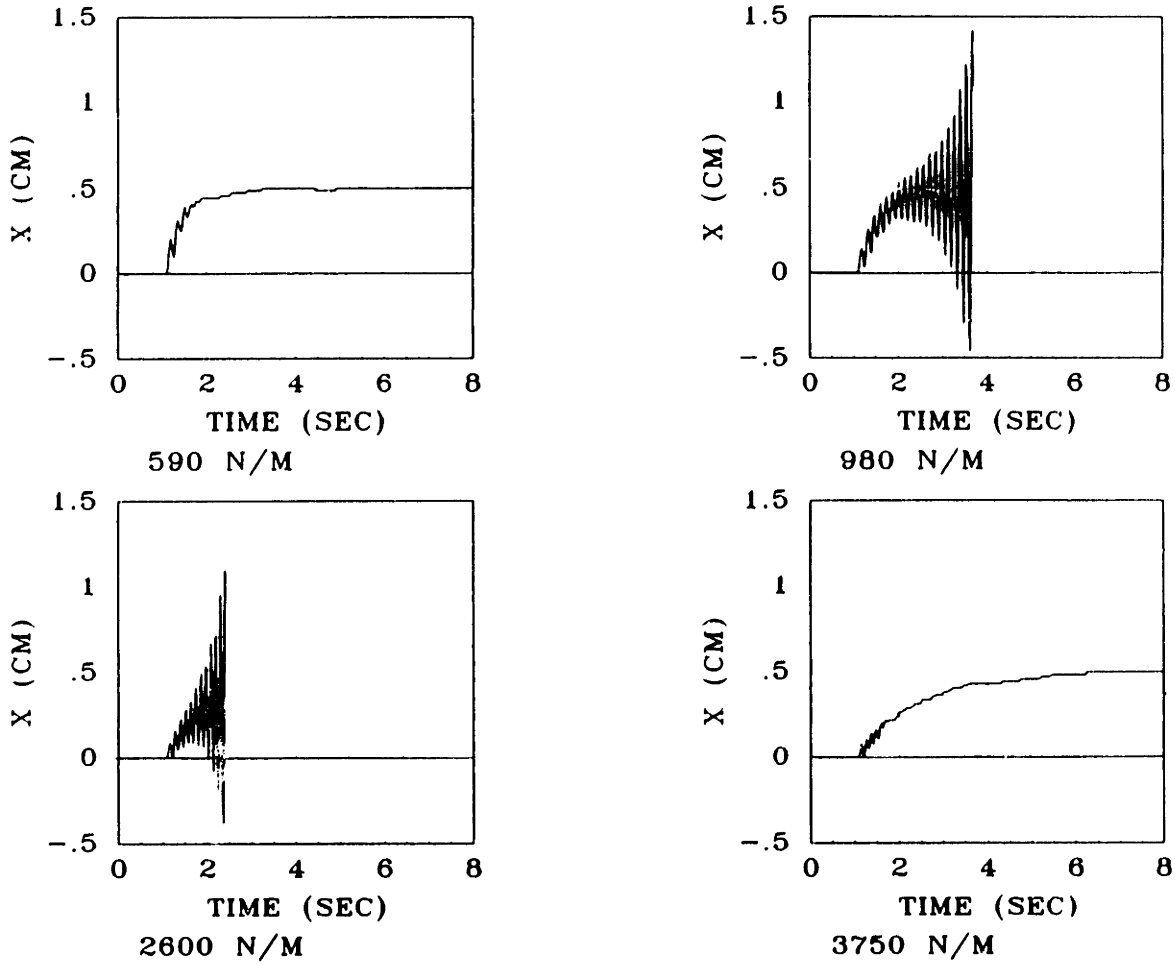

Figure 5.38: Ramp response when coupled to springs (x-direction): LQG/LTR control.

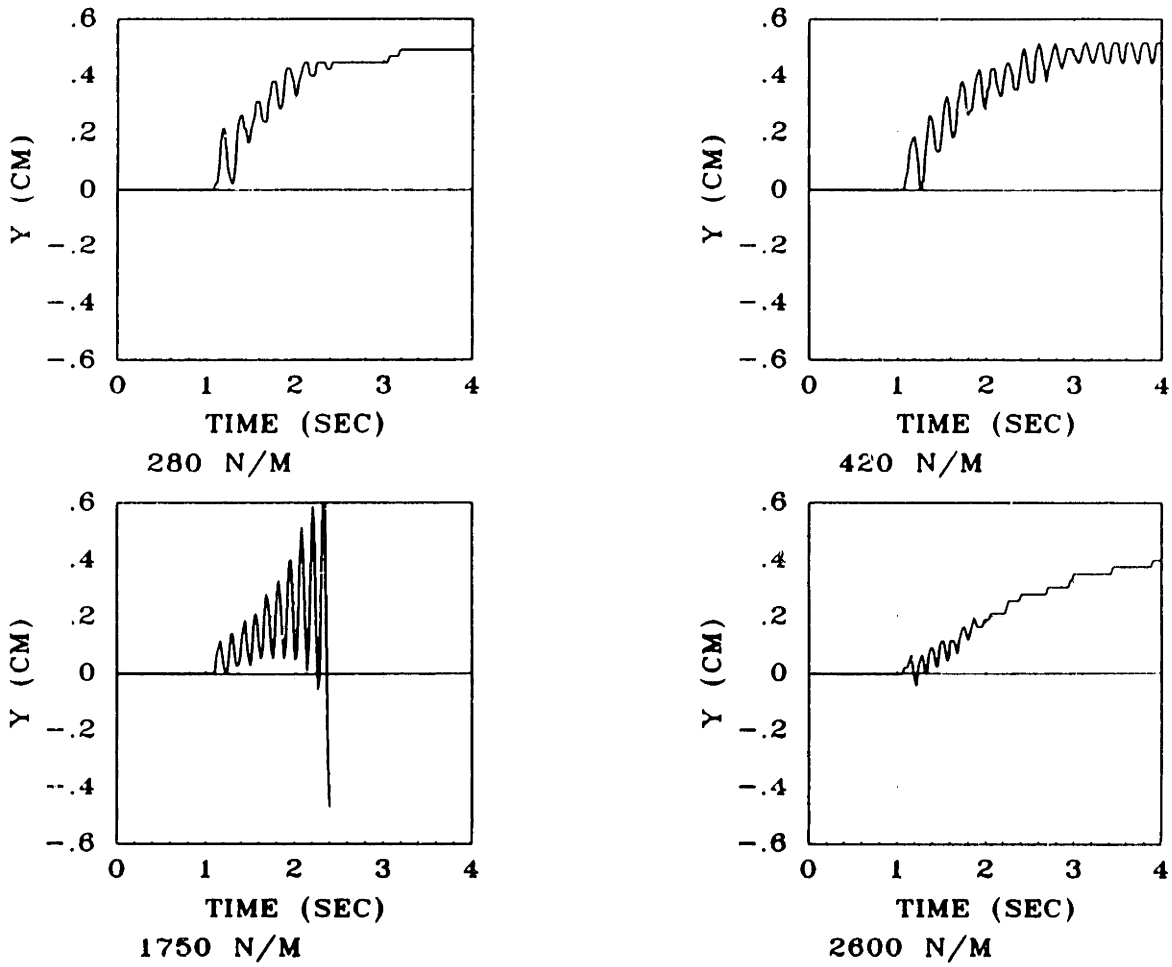

Figure 5.39: Ramp response when coupled to springs (y-direction): LQG/LTR control.

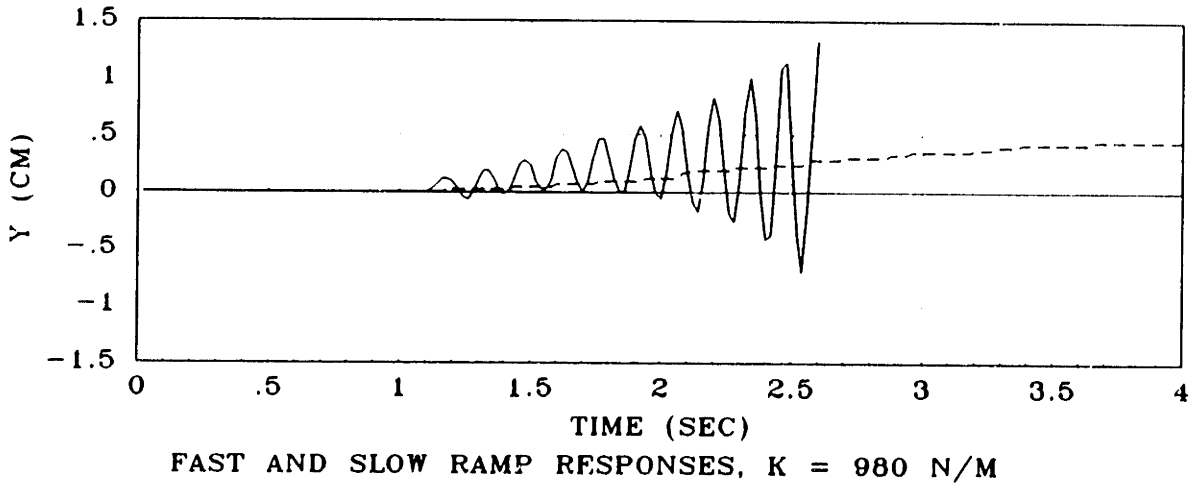

Figure 5.40: The effect of friction and speed on the stability of a ramp response: LQG/LTR control. Stable response—0.5 cm/sec ramp input; unstable response—10 cm/sec ramp input.

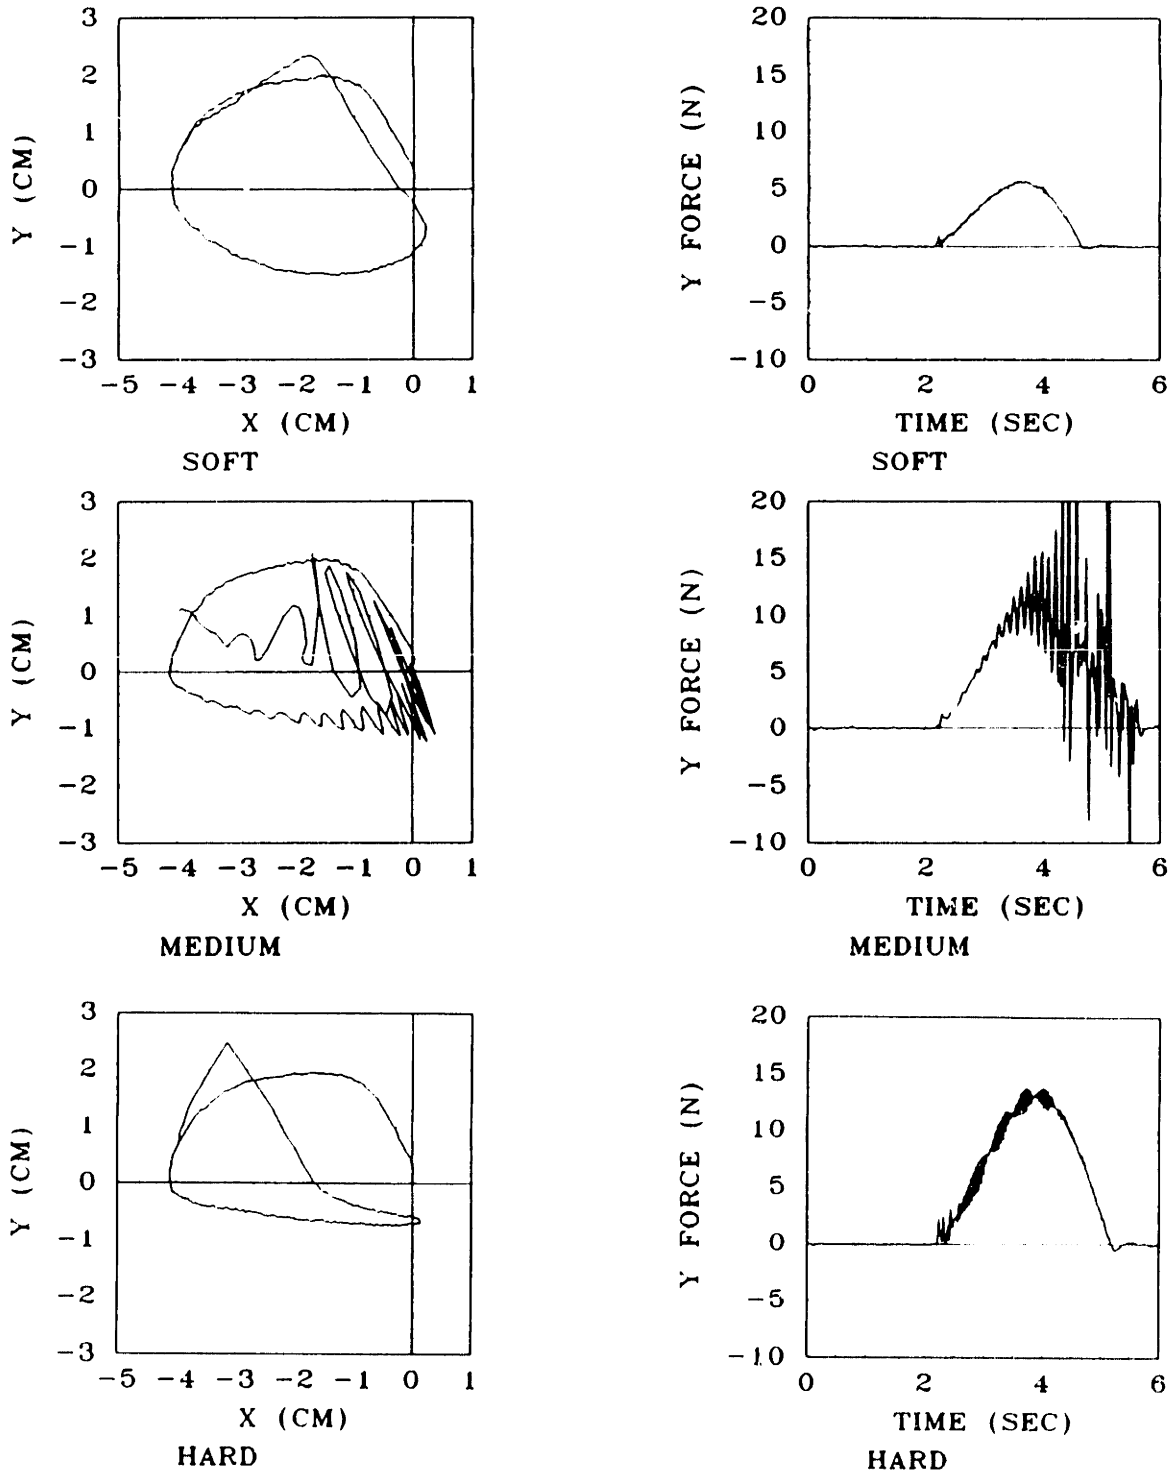

Figure 5.41: Endpoint position and (y-force) during the execution of a contact task: LQG/LTR control. Stiffness of soft environment ~300–400 N/m; medium ~1000-1800 N/m; hard ~2200–10000 N/m.

Chapter 6

An Application of Network Synthesis to the Analysis of Interactive Behavior

In this chapter the concepts of a *passive physical equivalent* and an *uncontrollable element* are introduced. These concepts are intended to exploit physical systems modeling techniques to lend insight into the interactive behavior of controlled systems. In the next chapter passive physical equivalents are used to analyze the behavior of a force feedback controlled manipulator. In both this chapter and the next, attention will be restricted to 1-port systems.

6.1 Passive Physical Equivalents

The previous chapters have established that the properties of the driving point impedance dictate interactive behavior. For instance the driving point impedance contains the only information necessary to predict the coupled stability properties of a linear manipulator. The impedance, however, is limited in one important respect: it can be an unwieldy tool for design.

A comparison with SISO servo design will illustrate this limitation. Consider the

Nyquist plot of an open loop transfer function $K(s)G(s)$, where $K(s)$ represents the compensator, and $G(s)$ the plant. The phase and gain margins which are taken from the Nyquist plot can be used directly for design. For instance, because the phase of the compensator adds to that of the plant, a change in the phase margin can be implemented simply by changing the phase of the compensator by an equal amount. The effect of changes in the plant dynamics on overall system behavior are equally apparent.

The Nyquist plot of the driving point impedance, however, despite its importance as an analytical tool, does not let effects due to the compensator and plant be separated in such a simple fashion. For instance, consider a manipulator with an open loop admittance $Y(s)$ relating the input effort, e, to the output flow, f, and an open loop transfer function $G(s)$ relating the control input, u, to f. If the control law $u = -K(s)f$ is implemented, then it is readily shown that the closed loop admittance is:

$$Y_c(s) = (I + G(s)K(s))^{-1}Y(s)$$

The role of $K(s)$ is now clouded because it does not show up as a simple multiplier. Whereas the concepts of gain and phase margin can be used to restrict the class of compensators which will lead to closed loop stability, there is no apparent method to restrict the class of compensators which lead to a positive real driving point impedance. The problem is only exacerbated if the sensed outputs are not the same as the interaction port outputs, or if input feedback is used.

What is needed is a way to simplify the interpretation of a closed loop driving point impedance. A successful method would clarify the contributions of the compensator and the plant to the overall behavior, perhaps going so far as to suggest changes to the design of either that would improve the behavior. The proposed method, the generation of a *passive physical equivalent*, has, in theory, this as well as other advantages to recommend it.

As indicated in Section 2.4.2, if a driving point impedance is positive real, then some passive system (called a *realization*) exists which has that impedance at an interaction port. If the driving point impedance is that of some controlled system, such a passive realization is termed a passive physical equivalent. In other words, although the system is feedback controlled, and therefore active at some sufficiently detailed level of description, there exists a passive system[1] which has the identical port behavior. Passive physical equivalents may be found from driving point impedances using the network synthesis techniques described in the next section.

It should be evident that passive physical equivalents are a form of Hogan's "physical equivalence" described in Section 1.2.2. The two stated goals of the physical equivalence approach are one, to bring to bear on the behavior of a controlled system those powerful techniques developed for the analysis of physical systems, and two, to restrict in some useful fashion the class of controllers that need be considered. Passive physical equivalents will to a certain extent, given that stable interaction with arbitrary passive environments is a primary consideration, satisfy both goals.

For instance, consider Tellegen's theorem, which is one of the more powerful tools of physical systems theory. One consequence of Tellegen's theorem is that, if a system is composed of passive components, it is itself passive [4]. Although earlier we assumed that the driving point impedance was positive real, we evidently need not have—if we can show that each element of the passive physical equivalent is passive (i.e. positive viscous coefficients, stiffnesses, and masses), then we are certain that the overall system is passive[2]. This is an interesting result, because it is potentially much easier to determine whether a given element value, say a stiffness, is positive than it is to check the positive realness of a high-order impedance function. Furthermore, the dependence of an element value on control gains and plant parameters can be exploited either to

[1] Generally, multiple passive physical equivalents exist. See the next section.

[2] The converse (an active element implying an active overall system) is not necessarily true, however. See Section 6.3.

bound gains or to direct tradeoffs between plant design and compensator design.

Finally, it is felt that the most important reason for using passive physical equivalents is to reap the insight that they can provide into the workings of control systems. For instance, it may be seen that the effect of a particular controller is to increase some stiffness or to decrease a mass. These are effects for which great insight exists—based, of course, on physical systems modeling. This capability and the others described above are used used in the analysis of force feedback control presented in the next chapter.

Before introducing the network synthesis techniques, a simple example should help illustrate the passive physical equivalent concept. Consider the system shown in Figure 6.1(a), which has the following driving point admittance:

$$Y(s) = \frac{s}{Ms^2 + K}$$

If the control law $u = -kx - bv$ is implemented, the closed loop driving point admittance is:

$$Y_c(s) = \frac{s}{Ms^2 + bs + (k + K)}$$

Because this controlier is simply generating a force proportional to displacement and a force proportional to velocity, it is doing the equivalent of connecting the mass to ground through a spring of stiffness k and a damper of viscous coefficient b. The two springs (k and K) are in mechanical parallel, so their stiffnesses add. The passive physical equivalent is shown in Figure 6.1(b). It is readily shown that this system has the driving point admittance $Y_c(s)$. For the equivalent to be passive, it is sufficient (according to Tellegen's theorem) that $b \geq 0$ and $k \geq K$; it is fairly easy to convince oneself that these conditions are also necessary for passivity.

This example points out several basic, but interesting facts:

- Control does not necessarily affect all elements of the plant model. In this case, the endpoint mass is unchanged.

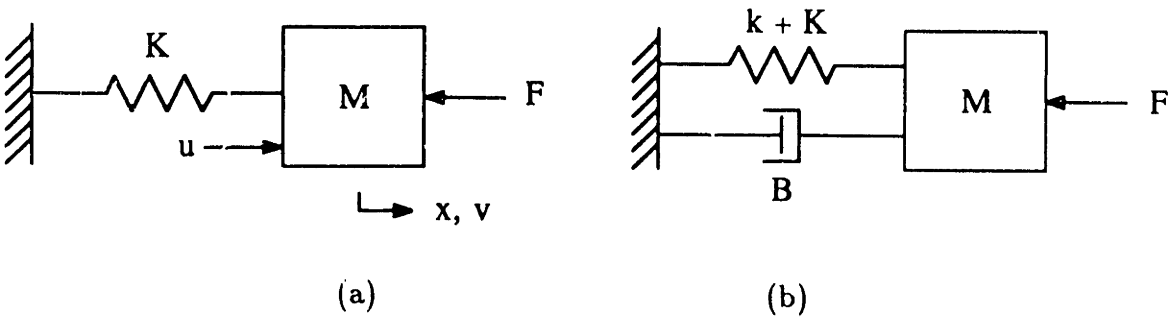

Figure 6.1: An example of a passive physical equivalent. (a) The open loop system. u is the control force, and F is the force imposed by the environment. x and v are position and velocity states, measured relative to ground. (b) The passive physical equivalent, given the control law $u = -kx - bv$ and $b \geq 0$ and $k \geq -K$.

- Control can introduce elements, such as the damper.
- Control can affect the magnitude of elements, such as the spring.
- The coupled stability criterion does not always require that the control be dissipative. In this case, negative stiffness can be added. If inherent damping existed, negative damping could also be added.

In the next section, the network synthesis procedures which allow the generation of more complicated passive physical equivalents will be introduced.

6.2 Network Synthesis Procedures

In the examples of the last section, the result was correctly guessed by correlating the effect of the control to the behavior of physical elements. In general, this is both difficult and unnecessary. Several techniques can be found in the electrical engineering literature for the synthesis of two-terminal (1-port) networks from driving point impedance functions. There are several excellent texts which describe these techniques, including those by Van Valkenburg [85] and Weinberg [87].

The oldest and most widely used procedures apply to electric networks which contain two types of elements only, i.e., LC, RC, or RL networks. We will consider LC networks, as the techniques are easily generalized to the others. The first two methods are due to Foster, and date to 1924. Both are based on a partial fraction expansion of the driving point impedance. The first uses the impedance form, and the second the admittance form. The partial fraction expansions, along with circuit diagrams and bond graph realizations, are shown in Figures 6.2 and 6.3.

The Foster synthesis techniques, and, for that matter all the techniques to be considered here, proceed in cycles. The first cycle consists of extracting a pole at $s = \infty$, if one exists. This is done by finding the residue, α_1, of $Y(s)/s$ as $s \to \infty$. The second cycle consists of extracting a pole at $s = 0$, if it exists; this leads to the term α_2/s. All subsequent Foster cycles consist of identifying sets of imaginary poles, and extracting them by finding their residues. The result has one of the two forms shown in Figures 6.2 and 6.3.

The second two methods for the synthesis of LC networks are due to Cauer, and date to 1927. These are: "continued fraction expansion about infinity" (Figure 6.4), and "continued fraction expansion about zero" (Figure 6.5).

Continued fraction expansion about infinity consists of the following steps: $Z(s)$ is written in acausal form so that there is a pole at $s = \infty$. This pole is removed by long division, leaving $Z_1(s)$. $Z_1(s)$ is inverted, yielding $Y_1(s)$ which has a pole at $s = \infty$. Each subsequent cycle consists of an inversion and the removal of a pole at $s = \infty$. As indicated in Figure 6.4, the result of this process is a chain of inductors and capacitors, with the inductors in series, and the capacitors in parallel.

Continued fraction expansion about zero is similar, except that each cycle consists of an inversion and the removal of a pole at $s = 0$. The resulting network has the same structure, but with the capacitors in series and the inductors in parallel.

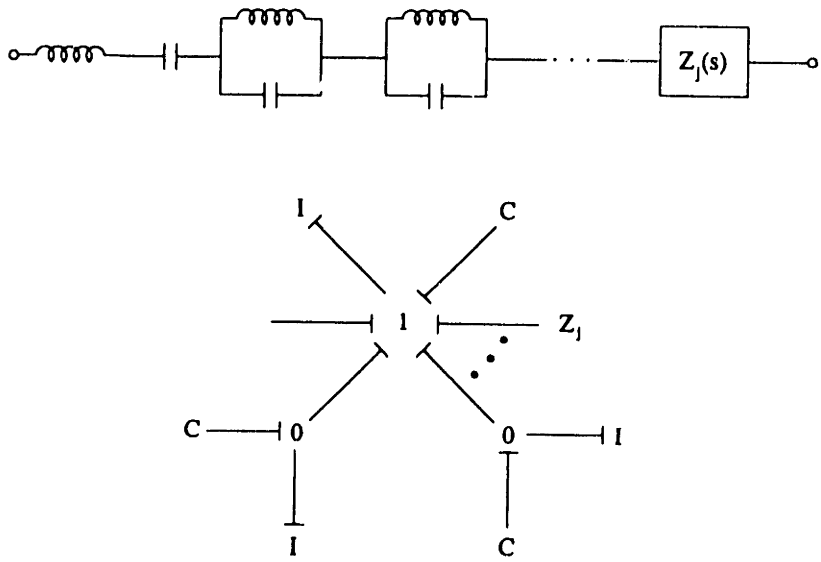

$$Z(s) = \alpha_1 s + \frac{\alpha_2}{s} + \sum_{i=3}^{i=j-1} \frac{2\alpha_i s}{s^2 + \omega_i^2} + Z_j(s)$$

Figure 6.2: First Foster form—partial fraction expansion in the impedance form. The components of this expansion are all connected in series.

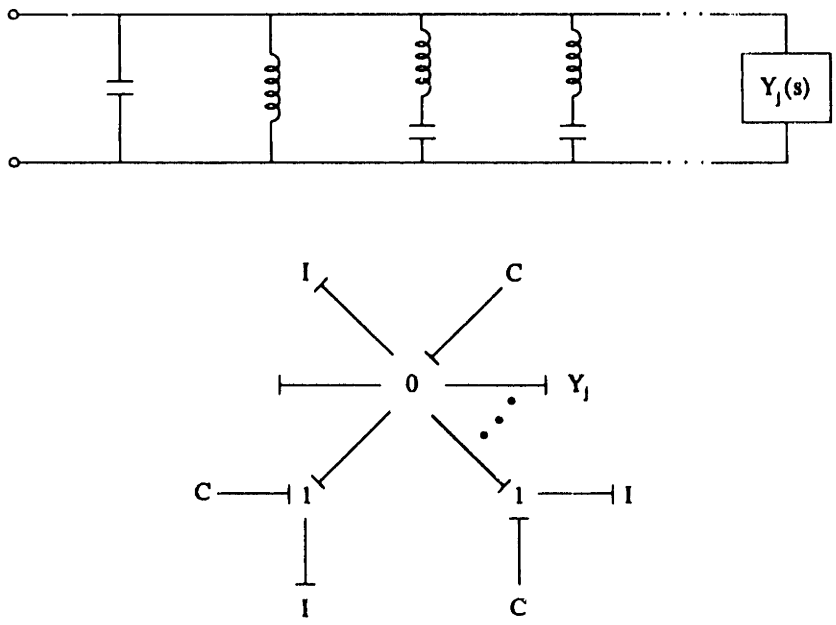

$$Y(s) = \alpha_1 s + \frac{\alpha_2}{s} + \sum_{i=3}^{i=j-1} \frac{2\alpha_i s}{s^2 + \omega_i^2} + Y_j(s)$$

Figure 6.3: Second Foster form—partial fraction expansion in the admittance form. The components of this expansion are all connected in parallel.

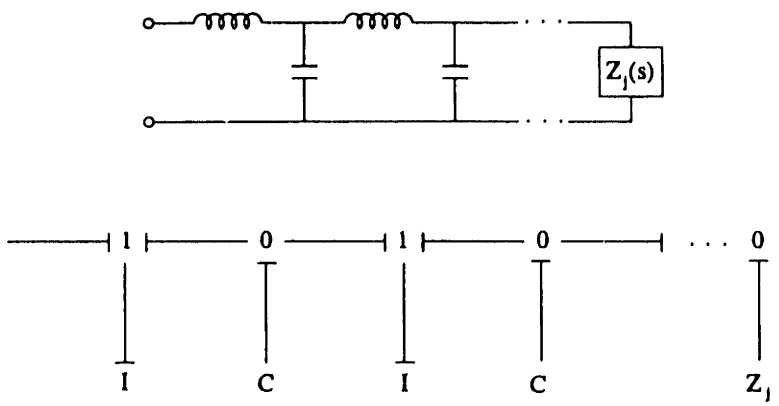

$$Z(s) = L_1 s + \cfrac{1}{C_2 s + \cfrac{1}{L_3 s + \cfrac{1}{C_4 s + \ddots}}}$$

Figure 6.4: First Cauer form—continued fraction expansion about infinity. When applied to the impedance of a positive real LC network, this synthesis technique results in a chain of series inductors and parallel capacitors.

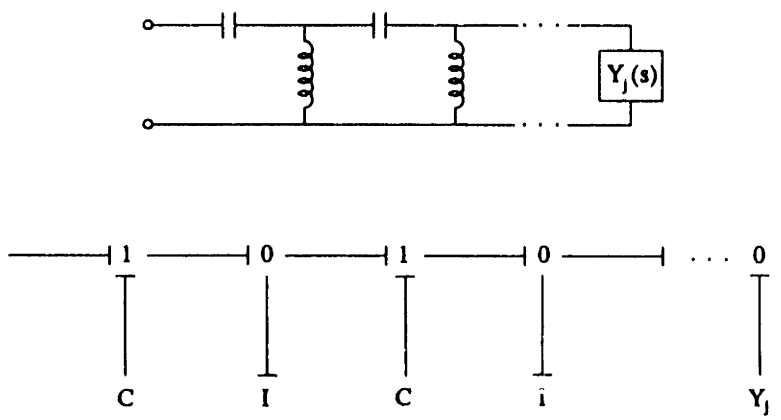

$$Z(s) = \cfrac{1}{C_1 s} + \cfrac{1}{\cfrac{1}{L_2 s} + \cfrac{1}{\cfrac{1}{C_3 s} + \cfrac{1}{\cfrac{1}{L_4 s} + \ddots}}}$$

Figure 6.5: Second Cauer form—continued fraction expansion about zero. When applied to the impedance of a positive real LC network, this synthesis technique results in a chain of series capacitors and parallel inductors.

The synthesis of RC and RL networks is essentially identical. Either the Foster or the Cauer techniques can be applied to these networks with little difficulty. One can even find a change of variables which will convert an RC network to an LC network or vice versa, or an RL to an LC or vice versa [87]. One difference, however, is worth noting. Whereas the real part of the driving point impedance of an LC network is zero for all $s = j\omega$, it is greater than zero for an RC or RL network. In these cases, however, the real part is a monotonic function of ω which has its minimum at either 0 or ∞. This minimum, if not zero, is generally removed as a first step, corresponding to the removal of some series or parallel resistance.

An example of the generation of a passive physical equivalent using Cauer synthesis is presented in Appendix C.

The Foster and Cauer syntheses owe much of their simplicity to the fact that, for networks composed of two kinds of elements only, after any cycle, the remaining impedance is guaranteed to be positive real [85]. For RLC circuits, however, this need not be the case. At some point in the synthesis process, it is possible that the remaining impedance will be a *minimum function*, i.e. an impedance whose real part is not a monotonic function of ω, but has a minimum value at (at least one) finite frequency, ω_1. If an appropriate resistor is removed so that this minimum value becomes zero, it can be shown that no poles or zeros remain on the imaginary or real axes, and that none of the previous techniques can be successfully applied for another cycle without a remainder function which is not positive real. As a result, a new type of cycle is needed.

Such a cycle was provided by Brune in 1931 [13]. The Brune cycle entails the removal of the circuit element shown in Figure 6.6(a). This element consists of an ideal[3] transformer (unity coefficient of coupling: $M/\sqrt{L_p L_s} = 1$) connected to a capacitor.

[3] "Ideal" in the electrical engineering sense of no flux leakage rather than the bond graph sense of no energy storage.

Brune showed that when such an element, with L_p, L_s, M, and C properly chosen, is removed from a minimum function, the result is positive real, and both numerator and denominator are reduced in order by two.

A bond graph of this circuit element is shown in Figure 6.6(b). As the bond graph indicates, energy is stored but not dissipated in an ideal transformer. Energy is stored in a magnetic field (in the soft iron loop connecting the coils) with the consitutive relation $\dot{\mathcal{F}} = \mathcal{R}\dot{\Phi}$ (see figure legend) which is considered to be that of a capacitor. Energy is transformed[4] from the magnetic domain to the electrical when changes in the magnetic flux, Φ, induce electromotive force in the coils. Of course, the reverse process also occurs, i.e., current in the coils induces a magnetomotive force.

Given this understanding, it is fairly straightforward to find a mechanical equivalent. Gyrators, however, do not exist in the translational mechanical realm. Transformers, therefore, are used instead, as shown in Figure 6.7(a) Ignoring for a moment the storage, we note that two gyrators connected in series constitute a transformer, as do two transformers connected in series. Both descriptions, therefore, have the basic behavior of a transformer.

The mechanical equivalent shown in Figure 6.7(b) is somewhat awkward. A simpler equivalent can be found by generating its dual, which is shown in Figure 6.8. In this case, the ideal transformer is simply a massless lever; the storage in the transformer is the rotational spring at its pivot. The capacitor in the Brune element becomes a mass upon which the lever assembly rests.

Although the dual is a somewhat simpler equivalent, it is unfortunate to resort to ideal transformers at all, as they surely add a level of complexity to the interpretation of a passive physical equivalent. Ideal transformers are a source of concern for those interested in network synthesis also, primarily for reasons of parameter sensitivity which are discussed in the next section. After Brune published his method in 1931,

[4]Given v and $\mathcal{F}$ as efforts, the relations $v = N\dot{\Phi}$ and $\mathcal{F} = Ni$ are actually those of a gyrator.

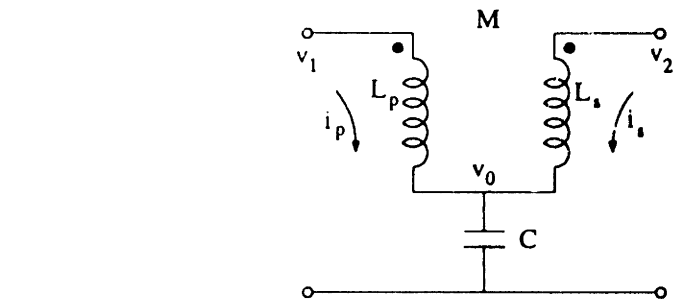

(a)

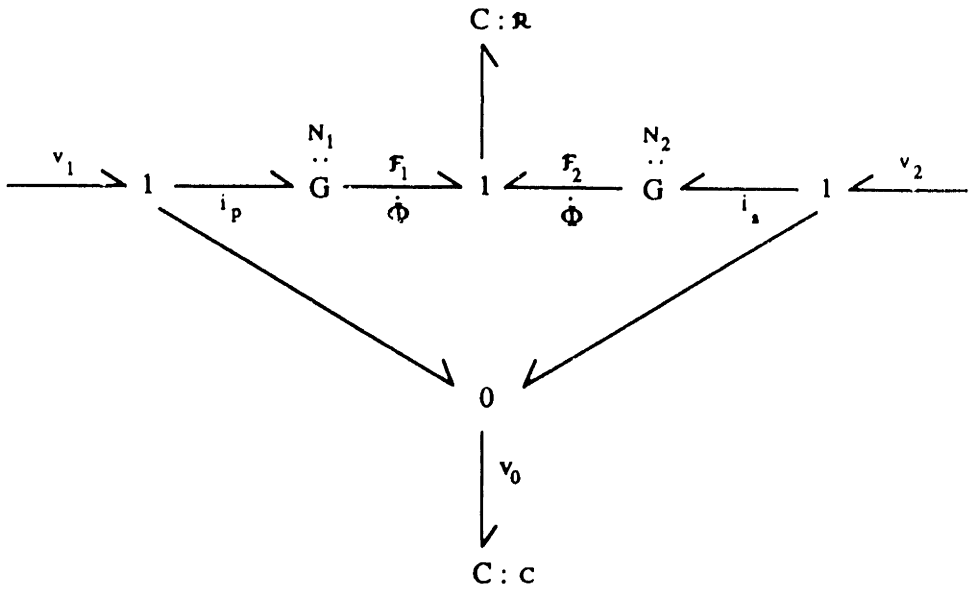

(b)

$$
\begin{aligned}
v_1 - v_0 &= L_p \frac{di_p}{dt} + M \frac{di_s}{dt} \\
v_2 - v_0 &= M \frac{di_p}{dt} + L_s \frac{di_s}{dt} \\
i_p + i_s &= C \frac{dv_0}{dt}
\end{aligned}
$$

Figure 6.6: The Brune element. (a) Circuit diagram of the circuit element extracted by one Brune cycle. (b) Bond graph equivalent to the circuit diagram. v is voltage, i is current, N is number of turns, $\dot{\Phi}$ is the rate of change of magnetic flux, $\mathcal{F}$ is magnetomotive force, and $\mathcal{R}$ is reluctance. The constitutive equation of the magnetic storage is $\dot{\mathcal{F}} = \mathcal{R}\dot{\Phi}$, and of the gyrators are $v = N\dot{\Phi}$ and $\mathcal{F} = Ni$. The bond graph model leads to the same equations as the circuit model provided we identify the inductances as follows: $L_p \equiv N_1^2/\mathcal{R}$, $L_s \equiv N_2^2/\mathcal{R}$, and $M \equiv N_1 N_2/\mathcal{R}$.

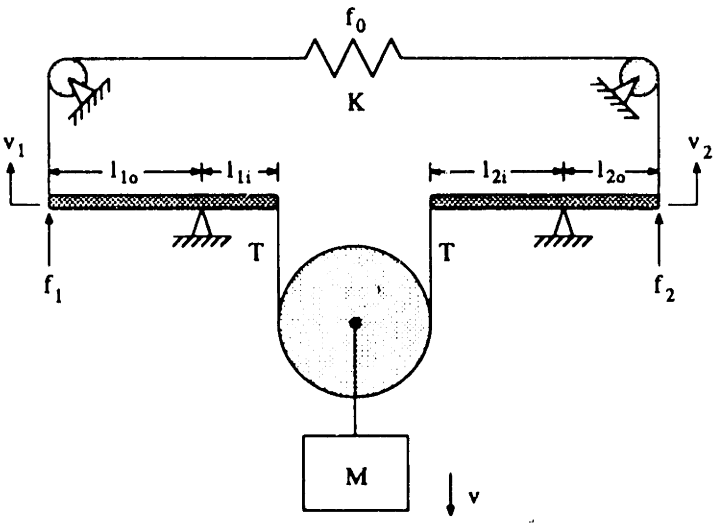

(a)

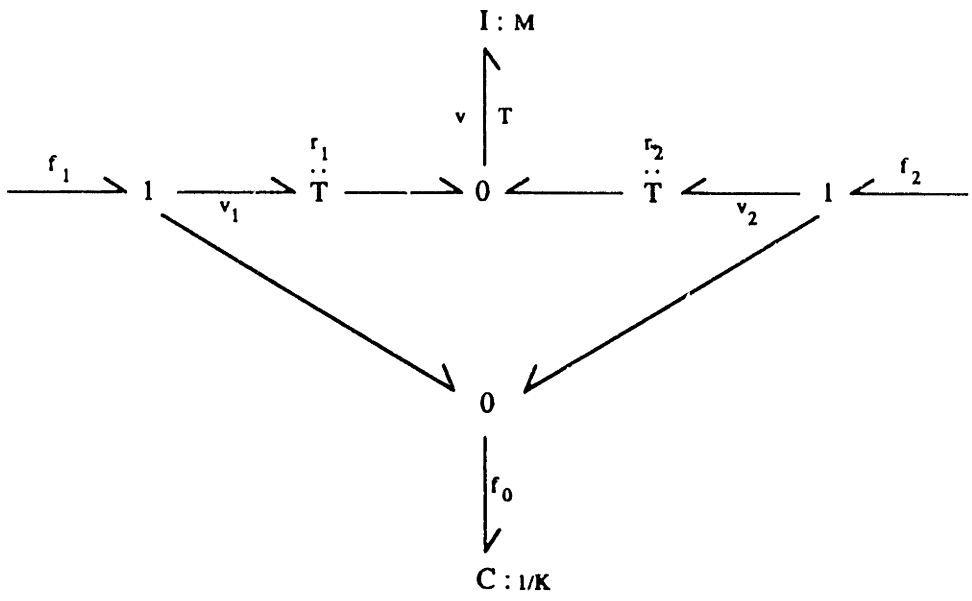

(b)

$$
\begin{aligned}
f_1 - f_0 &= \frac{Mr_1^2}{4}\frac{dv_1}{dt} + \frac{Mr_1r_2}{4}\frac{dv_2}{dt} \\
f_2 - f_0 &= \frac{Mr_1r_2}{4}\frac{dv_1}{dt} + \frac{Mr_2^2}{4}\frac{dv_2}{dt} \\
v_1 + v_2 &= \frac{1}{K}\frac{df_0}{dt}
\end{aligned}
$$

Figure 6.7: Mechanical equivalent to Brune element. (a) Bond graph model. (b) Mechanical equivalent. Definitions and assumptions: $r = l_i/l_o$; no gravity; cables are inextensible and never go slack; levers and pulleys are massless; f_0 increases in compression.

a tremendous amount of effort was put into the search for "transformerless" synthesis methods.

This search bore no fruit until 1949 when Bott and Duffin [11] showed that such a method did indeed exist. This method was an important milestone in the history of network synthesis. It is explained in detail in [85,87]. We will not review it here, however, because it results in a realization with many more elements than that generated by the Brune synthesis. The Brune realization is minimal, meaning that no realization can be found with fewer elements, whereas the Bott-Duffin realization is non-minimal, meaning that it results in a higher-order impedance which relies upon pole-zero cancellations to generate the correct driving point behavior. Thus, the Bott-Duffin synthesis has two serious drawbacks for our purposes: it generates a more complex realization measured in terms of the number of elements, and it relies upon exact pole-zero cancellation, which is a phenomenon we associate with control systems, usually not with physical systems. It is therefore inconsistent with one of our primary reasons for generating a passive physical equivalent, to generate insight through the application of physical systems modeling techniques.

This completes our brief tour of network synthesis procedures. In summary, given a positive real driving point impedance, a passive physical equivalent can always be found, but there is no guarantee that it will be simple to interpret. As for electrical network synthesis, the Bott-Duffin and all subsequent synthesis techniques turned out to exhibit extreme parameter sensitivity and to be of little value in practice. In the 1980's, analytical network synthesis procedures are rarely ever taught or used. The next section describes how parameter sensitivity and various other difficulties affect passive physical equivalents.

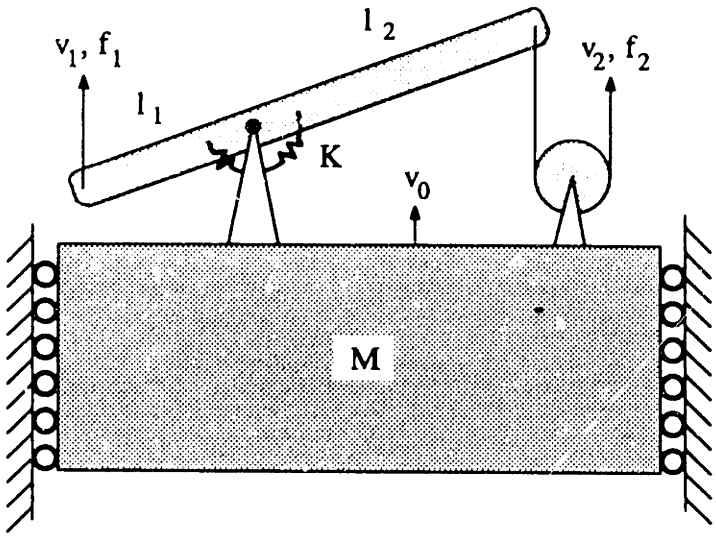

(a)

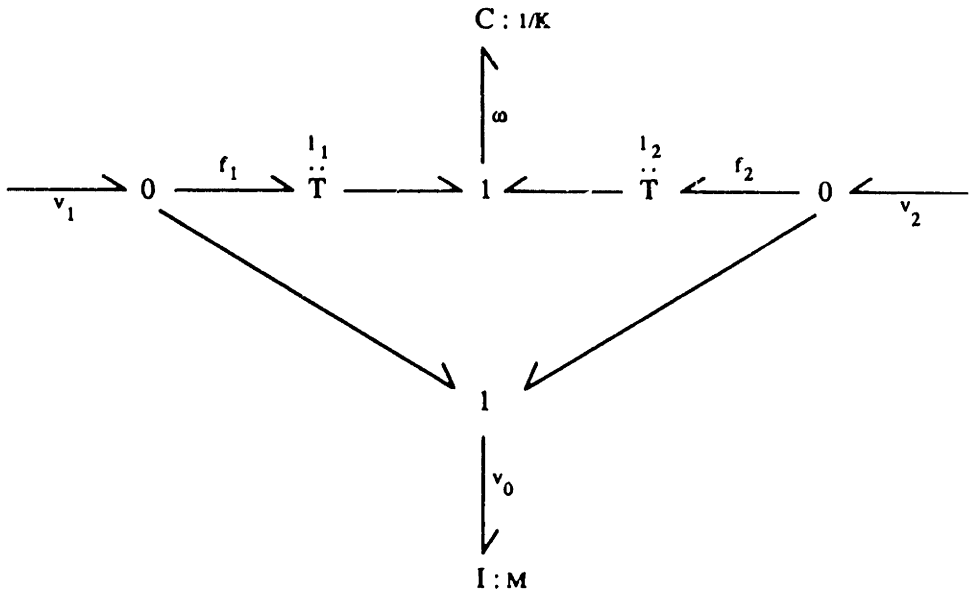

(b)

$$
\begin{aligned}
v_1 - v_0 &= \frac{l_1^2}{K}\frac{df_1}{dt} + \frac{l_1 l_2}{K}\frac{df_2}{dt} \\
f_2 - f_0 &= \frac{l_1 l_2}{K}\frac{df_1}{dt} + \frac{l_2^2}{K}\frac{df_2}{dt} \\
f_1 + f_2 &= M\frac{dv_0}{dt}
\end{aligned}
$$

Figure 6.8: Dual of mechanical equivalent. (a) Bond graph model. (b) Dual. Definitions and assumptions: no gravity; cable is inextensible and never goes slack; lever and pulley are massless.

6.3 Difficulties

After the previous section, it should be evident that there are several potential difficulties with the generation and interpretation of passive physical equivalents. The various difficulties as they are currently perceived are discussed in this section. They are:

- Complexity introduced by ideal transformers.
- Parameter sensitivity of realizations.
- Non-uniqueness of realizations.
- Structural stability of realizations.
- Complexity of an analytical approach.

Upon close inspection, the complexity introduced by the ideal transformers that appear in the Brune synthesis is not really a problem at all. Despite the complicated junction structure created by the lever arrangement in, for instance, Figure 6.7, the passivity of the element is determined only by the signs of the stiffness, K, and the mass, M, exactly as it is for any other second order lossless mechanical element. The junction structure is simply a price that must be paid to understand the behavior of a sufficiently complicated control system.

It is more important to realize that the parameter sensitivity of the Brune (and other) syntheses is not a difficulty, either. The problem with electrical network synthesis is that the realizations are meant to be built. However, if the coefficient of coupling of the transformer is not exactly unity (or the pole-zero cancellations of the Bott-Duffin synthesis are not exact), the circuits may end up exhibiting behavior which differs substantially from that of the desired impedance. But the passive physical equivalent is not meant to be built—it has already been implemented by the control system. We need be no more concerned with an ideal transformer than with a massless spring.

And, although we do not recommend the use of the Bott-Duffin technique, any pole-zero cancellation would be exact by definition, as the order of the closed loop admittance is fixed.

The biggest problem with passive physical equivalents is probably one of the most useful features for the network synthesist. This is the non-uniqueness of the realizations. Consider, for instance, LC networks. Four different synthesis techniques are available, and there is nothing that prevents the mixed use of these techniques. Thus, for a network of high enough order, the number of valid but distinct realizations may be quite large.

This is not necessarily a problem if a single realization can be selected for analysis. Of course, there is the problem of selecting the most appropriate realization, but one can generally use a model of the open loop system as a guide by attempting to select the realization with the structure closest to this model. This is the approach that is taken in Chapter 7. It tends to work well for low-order compensators and controllers (e.g. impedance controllers) which are intended to exhibit physically-interpretable behavior, but not as well for high-order compensators which are designed to meet certain performance criteria.

A more serious problem, however, is that, although there is, by definition, a one to one mapping from the driving point impedance of each realization to that of the closed loop system, there is not necessarily a one to one mapping from the element by element passivity of a given realization to the passivity of the closed loop system.

For instance, suppose that the controller includes some gain factor, G, and that we know from an analysis of the driving point impedance that the positive real condition is satisfied so long as $a \leq G \leq c$. Then it does not follow that any realization will contain all positive element values for the same range of G. What does follow is that some set of realizations must exist which collectively contain instances of all positive element

values for all $a \leq G \leq c$[5]. For instance, one realization may be valid for $a \leq G \leq b$, but at $G = b$ the viscous coefficient of some damper goes to 0, and is negative for $G > b$. Meanwhile, another realization, which had a negative damper in some other location for $G < b$, has this damper become positive for $b \leq G \leq c$. An interesting consequence is that changing feedback gains can change the *structure* of the passive physical equivalent.

Another important consequence is that negative element values do not necessarily indicate a violation of the positive real condition. For this, it is necessary that *every possible* minimal realization have at least one negative element. By "minimal" we mean the equivalent of "reduced" in the bond graph sense, so that no excess storages or dissipators exist. If not reduced, it is possibie that a negative damper, for instance, could be added to a larger positive damper, thus removing the negative element value. To summarize:

- If all element values of a realization are positive, the driving point impedance is positive real.
- If a minimal realization contains a negative element value, the driving point impedance violates the positive real condition only if all other minimal realizations derived from that impedance also contain negative elements.

The fourth area of concern is the "structural stability" of the realizations. We have already indicated that, as a gain changes, the structure of the realization may change. Presumably, the same thing could happen due to a change in a plant parameter. If the nominal value of the plant parameter was close to its critical value, modeling uncertainties would introduce uncertainty with regard to the structure of the passive physical equivalent. Of course, so long as the modeling uncertainties introduce only

[5]Of course, *no* realization can exhibit all positive element values outside this range.

small errors to the driving point impedance, the two structures would necessarily exhibit similar port behavior.

A reasonable approach to treating this problem would be to develop software for the generation of passive physical equivalents. A software package similar to those which have been developed for the analysis of electrical networks and physical systems would graphically display the parameter and structural changes of the passive physical equivalent resulting from changes in plant or compensator parameters. The software might also address the uniqueness problem by sorting among the realizations to find those with the "best" structures (based, perhaps, on similarity to the plant structure).

Software is also the logical approach to the fifth area of concern, the complexity of an analytical generation of passive physical equivalents. All passive physical equivalents in this thesis were determined symbolically in terms of plant and compensator parameters. Even though only simple systems were treated, a great deal of time and effort was involved. In conclusion, it is felt that the future of passive physical equivalents as a technique for the analysis of interactive behavior depends in large part on the development of appropriate software. Certainly the Ricatti equation would never have developed such fame without the help of digital computers.

6.4 Uncontrollable Elements

Although the generation of a passive physical equivalent may in some cases be prohibitively difficult, a complete realization is not always necessary to learn much that is useful. For instance, it may be valuable to know, given particular actuator and sensor locations, which elements of the plant model *cannot* be changed by any compensator; i.e., which elements must appear, unaltered, in any passive physical equivalent, given any (causal) compensator. Such "uncontrollable elements" are the subject of this

section.

It can be shown that, given a plant model sufficiently detailed to include actuator and/or transmission dynamics separating the control and interaction ports, that at least one uncontrollable element must exist. The reasoning is as follows.

In Appendix D it is shown that, given any causal compensator and non-colocated control and interaction ports, the relative order of the driving point impedance will be unchanged and furthermore, that:

$$\lim_{s\to\infty} Z(s) = \lim_{s\to\infty} Z_c(s)$$

where $Z_c(s)$ is the closed loop driving point impedance. If $Z(s)$ is strictly proper, this implies that the zero at $s = \infty$ is unaffected by compensation. This in turn implies that any of the possible realizations contains some element which is unaffected by compensation. This is because all of the synthesis techniques result in a term of the form a/s, where $a = \lim_{s\to\infty} sZ(s)$, if a zero exists at $s = \infty$.

The preservation of the limit as $s \to \infty$ may be understood physically to mean that the high frequency behavior, which is dictated by the element located at the interaction port[6], cannot be altered by a controller which is necessarily of limited bandwidth. Therefore, that element located at the interaction port is preserved in any passive physical equivalent. Such an element is termed "uncontrollable" because no causal compensation scheme can alter its value.

Barring colocated control and interaction ports and barring actuators and transmissions of infinite bandwidth, *every* control system will contain an uncontrollable element. However, the element found at the interaction port is not necessarily the only one which is uncontrollable. Given a particular configuration of actuators and sensors, there may be others. Examples of multiple uncontrollable elements will play an important role in the analyses of the next chapter.

[6]If the plant is an admittance, this element must be a mass, if the plant is an impedance, it must be a spring.

The concept of an uncontrollable element illustrates a very important idea, which is that there is at least one aspect of the behavior of a physical system which no controller, no matter how sophisticated, can alter. Furthermore, this aspect is not one of small consequence, rather it is a critical determinant of the behavior which the control system represents to the outside world.

Chapter 7

Force Feedback as an Example of Interaction Control

In this chapter, passive physical equivalents are used to analyze the coupled stability properties of a force feedback controlled manipulator. The analysis demonstrates that there are certain physical limitations to the effectiveness of any force feedback scheme. Suggestions for improved force control are also made.

7.1 Background

The last several years have witnessed explosive growth in the study of robot force control. The development of successful strategies and implementations for force control is seen as a crucial step in enabling robots to perform tasks, such as deburring, drilling, and parts assembly, which require significant interaction with the environment. The implementation of high-bandwidth, high-accuracy force control, however, has proven to be quite difficult, primarily due to stability problems that occur upon contact with a rigid surface.

A full scale, practical implementation of force control, like other types of robot control, calls for a heirarchical architecture. The breakdown of this heirarchy is a choice of the designer, but it is clearly useful to separate the higher level supervisory

and planning designs, which must address the execution of a task, from the lower level control design, which must address the stability and performance issues. Much of the force control literature has been divided roughly along these lines.

Paul and Shimano [64], Mason [59], and Raibert and Craig [66] were responsible for much of the early work in planning. In particular, the hybrid position/force control proposed by Raibert and Craig and based upon the theoretical framework of Mason has evolved into a fairly active field of research [5,63,71,82,95]. The basic idea behind hybrid control is that the physical constraints of the task should dictate those axes along which force is controlled and those axes along which position is controlled. For instance, in a surface grinding task, force should be controlled normal to the surface, and position tangential to the surface.

Early hybrid schemes implemented force control based upon a kinematic rather than a dynamic description of the robot. More recently, Khatib [49,50] proposed an approach to hybrid control based upon the operational space formulation which accounts for the rigid body dynamics of a robot. Although it is evident on the basis of efforts such as this that the gap between those concerned with higher level issues and those concerned with lower level issues is beginning to close, it remains the case that most of the papers in the area of hybrid control have not addressed the lower level issues of stability or bandwidth in any detail.

But stability and high bandwidth are important issues in force control. Whitney [88] was the first to provide a stability analysis of a force controlled manipulator. He modeled a manipulator as a velocity-input integrator, and assumed that proportional position, velocity, and force feedback were implemented in discrete time. He modeled the environment as a spring and derived the following stability result:

$$0 < TGK_e < 1$$

where T is the sampling rate, G is the force feedback gain, and K_e is the combined

stiffness of the sensor and the environment. For fixed T, this indicates that a tradeoff exists between G and K_e. In other words, high bandwidth force control requires a compliant sensor or environment.

In the years since Whitney published this result, the essential stability tradeoff between force feedback gain and stiffness of the environment[1] has been substantiated many times [20,21,32,45,67,92]. It is now known that the problem transcends the digital implementation, that even for analog control the tradeoff exists [92]. In a recent review [63], Paul indicated that the contact problem with a rigid manipulator, rigid sensor, and rigid environment is still unsolved.

Various explanations for contact instability have been offered. Kazerooni [45] and An [3] have both shown that the tradeoff may be attributed to unmodeled dynamics. Eppinger and Seering [19,20,21] have carried out extensive stability analyses based on a series of increasingly sophisticated single-axis models. They have shown that stability problems arise due to the *non-colocation*[2] of sensors and actuators. There are two important consequences of their observations: first, that it is specifically those dynamic elements (modeled or not) separating the actuator and sensor which are most seriously implicated in contact instability; second, that, even if the robot is exquisitely modeled, high gain force control is a challenging proposition.

Eppinger and Seering also performed simulations which indicated that the force discontinuity that occurs upon contact with a surface may decrease the range of acceptable gains (which disagrees with the result of Section 3.6). However, not all nonlinearities necessarily degrade force control performance; Townsend [84] has shown that coulomb friction can actually have a stabilizing effect.

Many techniques have been tested in efforts to improve force control fidelity. A number of these techniques involve the use of dynamic compensation. Interpreting the

[1]The stiffness of the sensor is generally included as part of that of the environment.

[2]The importance of non-colocation in the stability of feedback systems was first noted by Gevarter [25] in the context of flexible space vehicles.

literature in this area is complicated by the variety of manipulator models that are used. Although a rigid body description generally serves as the base model, it is frequently assumed that the force control loop is closed *around* some higher-bandwidth inner loop, such as a position [45,67] or velocity [29,76,88] controller. Others, however, have analyzed force feedback in the absence of such inner loops [3,20,21,96]. A review of the different control architectures in which force feedback appears is presented by Whitney [89]. De Schutter [16] has contrasted the use of position, velocity, and acceleration[3] controlled plants. His analysis, however, focuses on performance criteria such as steady state error and accuracy rather than contact stability.

Kazerooni [45] has shown, for (inner loop) position controlled manipulators, that the gain of any force feedback compensator is inversely proportional to the stiffness of the position servo and to the stiffness of the environment, if stability is to be guaranteed by the multivariable Nyquist criterion. If the position controlled manipulator is non-backdriveable, as is generally the case, and the environment is a rigid surface, the conclusion is that very low force feedback gains should be sufficient to create contact instability. Roberts, et al. [67] experimentally confirmed the tradeoff between servo/environment stiffness and the gain of an integral force controller. They found that low stiffness was necessary to achieve reasonable bandwidth.

Various dynamic compensators—PD, PI, and PID—have been implemented on velocity controlled manipulators with similar results. For instance, Haefner, et al. [29] implemented a PID force controller for a deburring task. Although they achieved stable control, they found that the bandwidth was too low to be of practical value.

A similar tradeoff exists in the absence of an inner loop control, and in this case the relative merits of various sorts of dynamic compensation have been more thoroughly investigated. Eppinger and Seering suggest that compensators which add lead (e.g.,

[3] In the case of a rigid body manipulator, acceleration control corresponds to the absence of an inner loop control.

PD) should allow larger gains, whereas those which add lag (e.g., PI or slow filters) should have the opposite effect. However, both An [3] and Youcef-Toumi [96] have achieved promising results with a first-order lag compensator. This inconsistency can perhaps be resolved by noting that these implementations require that the compensator rolloff occur at very low frequency, which Eppinger and Seering do not necessarily assume.

A number of other approaches to improved force control exist. For instance, Hogan [35,36] has shown that impedance control can be used to control the force exerted on an environment without the need for force feedback. This approach is to implement an endpoint impedance which is a dynamic relationship between motions input by the environment and forces output by the robot. Certain classes of impedances have been shown to guarantee contact stability [38]. Hogan and Wlassich [40,92] have also investigated impedance control implementations which make use of force feedback. They have shown that force feedback acts primarily to modulate the inertia of the robot as seen by the environment, and that high gain force feedback does the equivalent of attempting to mask this inertia. One interesting result is that *positive* force feedback (which acts to increase the inertia) was shown to have a strong stabilizing effect on contact. The utility of this result in the context of high bandwidth force control has not yet been demonstrated.

Other approaches to improving force control involve robot design modifications. A number of investigators have suggested the use of passive compliance between the robot and environment [3,16,18,64,88]. One example of passive compliance is the use of an RCC [18], which is well-known as a solution to the peg-in-hole insertion problem. De Schutter [16] has made the case that passive compliance is essential to reduce impact loads and therefore to increase task execution speeds. Of course, passive compliance is not satisfactory in all applications due to the positioning inaccuracies which may be created. Roberts, et al. [67] show that it is possible to compensate for this effect, but

their implementation exhibits fairly low bandwidth.

Part of the value of force feedback from a wrist sensor is that, in theory, undesirable effects such as motor cogging and joint friction may be masked. If these effects can be eliminated through design, however, open loop force control becomes a conceivable alternative. Youcef-Toumi [96] has made the case that, due to the inherently simple dynamics of a direct-drive arm, improved force control may be achieved, possibly in an open loop fashion. Townsend [83] has designed a cable-driven manipulator which is also intended to exhibit particularly good open loop force control.

Another design alternative is the use of a macro/micro architecture. Sharon, et al. [69] have argued that the use of a high bandwidth micro-robot mounted on the end of a standard macro-robot should enable the implementation of higher bandwidth force controllers. Tilley, et al. [82], have presented encouraging data with a similar system. Their manipulator is composed of a very fast wrist subsystem mounted on a very flexible arm. An integral force controller at the wrist was shown to exhibit reasonable stability and bandwidth.

In summary, high bandwidth control of the force exerted on a rigid environment remains elusive. None of the sundry compensation schemes which have been introduced have achieved tremendous success, and many current efforts are addressing better dynamic modeling of robots as well as better design of robots for force control.

7.2 Analysis of Lumped Models

In this section the technique of passive physical equivalents derived in the last chapter is used to analyze the coupled stability properties of several simple robot models. The reader may find it convenient to refer to Section 4.2, as the models treated in that section are the same as the first two considered here. Derivation of the passive physical

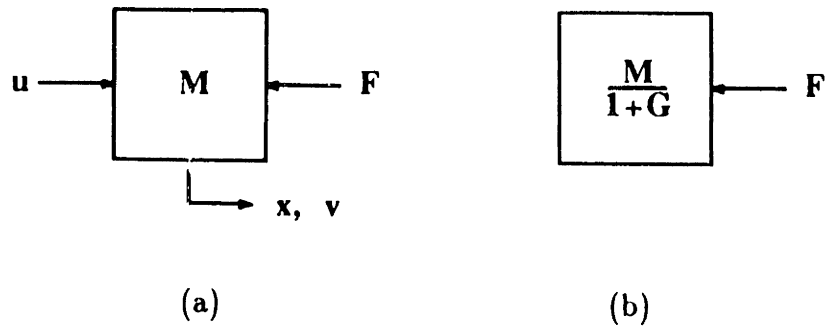

(a) (b)

Figure 7.1: (a) Rigid body robot model. (b) Passive physical equivalent for $u = -GF$.

equivalents is not presented in this chapter; an example is provided in Appendix C. It is left as an excercise for the reader to verify that the passive physical equivalents which are presented predict the limits of passivity exactly as the driving point impedance does (see Section 6.3).

The first robot model to be considered is shown in Figure 7.1(a). It is a mass M acted on by an actuator force u and an environmental force F. This model is extremely simple, yet it captures the essential behavior of a robot (in particular, a direct drive robot), at least at low frequency. If $u = -GF$ is selected as the control law, the passive physical equivalent is a mass $M/(1+G)$ acted on by the environmental force F (Figure 7.1(b)).

Figure 7.1(b) is extremely valuable insofar as it represents the relationship between force feedback and robot inertia. As described in detail in [40,82], negative force feedback acts to reduce the robot's inertia; conversely, *positive* force feedback acts to increase the robot's inertia. Note that $G \geq -1$, corresponding to $M < \infty$, is necessary for the closed loop system to be passive.

In general, a connection can be made between force servoing and admittance amplification. An ideal force control robot would act as an effort source[4], exhibiting infinite admittance. In actual implementations, the goal of a force controller should be to in-

[4]By analogy, an ideal tracking robot would act as a flow source.

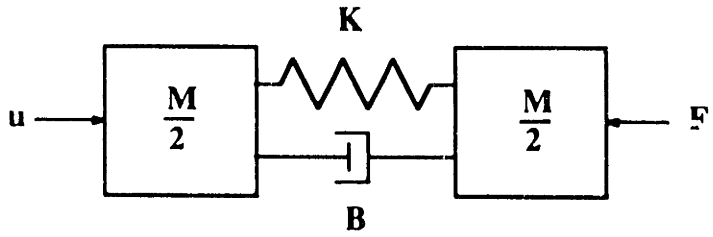

Figure 7.2: Manipulator with non-colocated actuation and control.

crease the admittance as much as possible over the desired bandwidth without risk of coupled instability. Because most robots are characterized primarily by inertial dynamics, increasing the admittance consists primarily of decreasing the inertia. However, other factors, such as joint friction and stiffness, also contribute to the admittance, and should, in general, be reduced as well. Reducing joint friction is discussed in Section 7.5.

The rigid body model, however, does not predict the contact instability that has been observed so frequently. This is not very surprising, as the actuator and sensor are colocated. Therefore, the non-colocated model shown in Figure 7.2 is considered next. In this model, K and B generally represent structural dynamics, although they may represent other effects, such as actuator or transmission dynamics. Note that the overall mass of this manipulator is the same as that of the first model.

A passive physical equivalent for this manipulator, given the same control law, $u = -GF$, is shown in Figure 7.3. Note that the endpoint mass, that acted on by the environment, is unchanged—it is an uncontrollable element. Despite the fact that this mass does not change, one would expect this model to predict the same low-frequency behavior as the rigid body model; thus, the total mass of the realization should be $M/(1+G)$. An examination of Figure 7.3 shows that this is the case. A consequence of this low frequency behavior is that the the mass near the actuator exhibits a very strong dependence upon G.

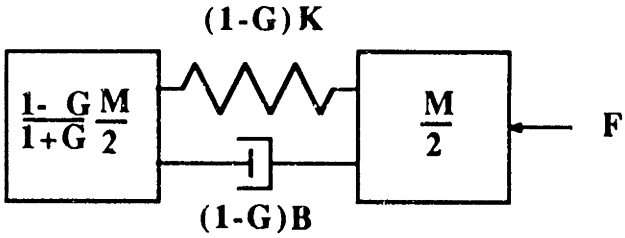

Figure 7.3: Passive physical equivalent of the two-mass model.

Clearly, the force feedback gain is restricted to lie between -1 and 1 for this to be a passive realization. Large negative force feedback gains ($G > 1$) would result in a negative spring, mass, and damper. Oddly enough, even negative values such as these do not suggest that the isolated manipulator is unstable. In fact, it is not, and the reasoning is quite simple: because force feedback is not state feedback, the poles of the closed loop system, and therefore of the realization, must equal the poles of the open loop system, but the open loop system is stable.

The negative system elements do, however, indicate trouble when coupled to some passive environments. For a detailed analysis of the coupled stability properties of this system, the reader can refer to Section 4.2. The result is that stiff environments are among the most destabilizing, a finding which is consistent with observation. Notice also that the coupled stability result is independent of the values of M, B, and K, so long as they are positive. To be sure, the worst environment test in Section 4.2 indicates that the environment must have a stiffness on the order of the structural stiffness of the robot to create the instability. This does not weaken the result significantly, because many potential workpieces (e.g., an engine block) are at least as stiff as commercially available robots.

To summarize, the behavior of this system may be attributed to one, the uncontrollability of the endpoint mass, and two, the preservation of the rigid body low frequency

behavior. One may ask, however, if, in light of the apparent importance of the structural dynamics, this intuitive explanation and the quantitative limits on G will hold for more complicated robot models.

One approach to this question is simply to examine higher-order models. Passive physical equivalents for three-mass and four-mass uniform[5] models are shown in Figure 7.4. These realizations support the qualitative and quantitative behavior shown by the two-mass model. They further indicate that springs, masses, and dampers are all affected by the same scaling factors near the ends of the structure. The four-mass model also shows symmetry virtually identical to that of the two-mass model. The logical culmination of this line of investigation is a distributed beam model, which is treated in the next section.

7.3 Analysis of a Distributed Model

In this section, the uniform distributed model of a robot shown in Figure 7.5 is analyzed. The compression/tension beam has uniform density ρ, modulus of elasticity E, area A, and length L. As before, the control law is $u = -GF$. This is a better robot model than any of the lumped parameter versions previously analyzed, but is not as realistic as a cantilever beam. It will be seen, however, that the primary value of the distributed model is in elucidating the role of transmission delays—toward such an end, the compression/tension beam is sufficient. Further, it is somewhat easier to analyze than a cantilever beam, and is the natural extension to the lumped parameter models of the last section.

The network synthesis procedures become decreasingly tractable as the order of the system increases; using them to analyze an infinite dimensional system would be asking

[5] Uniform means that all springs have the same value, all masses do, and all dampers do. Transmission lines are frequently modeled as uniform dynamic chains.

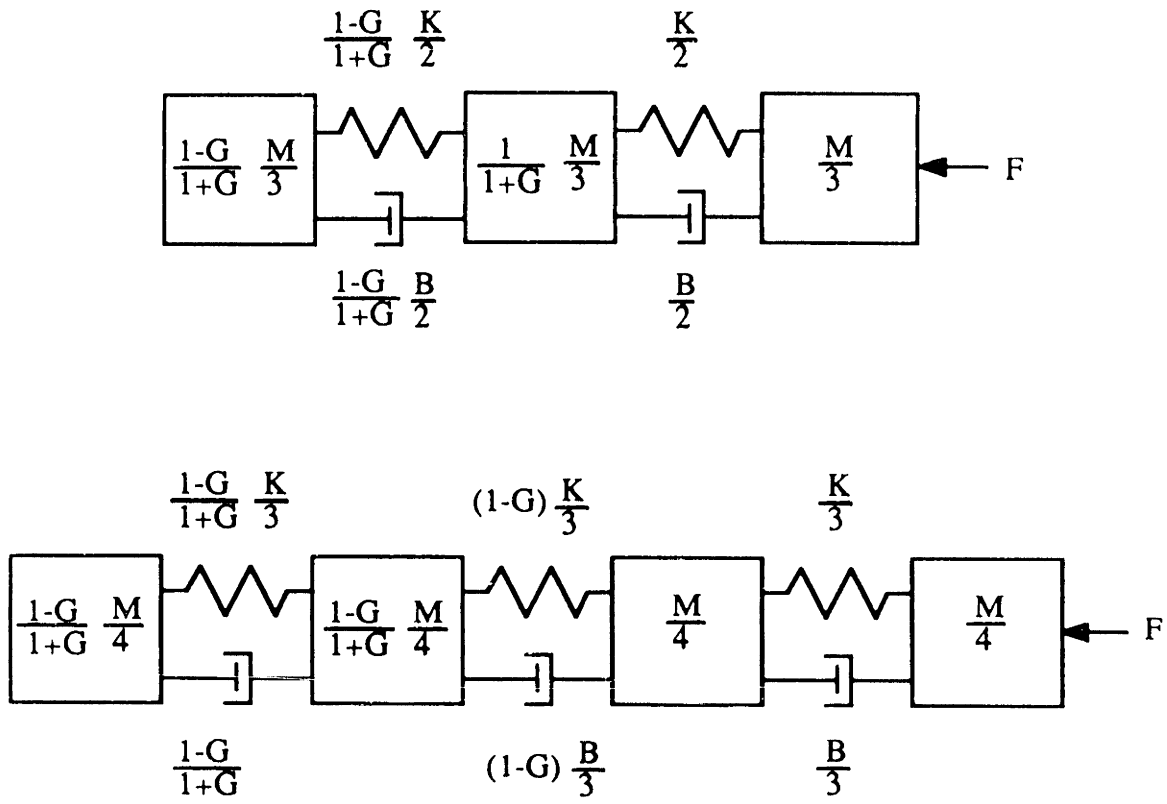

Figure 7.4: Passive physical equivalents of three-mass and four-mass uniform beam models.

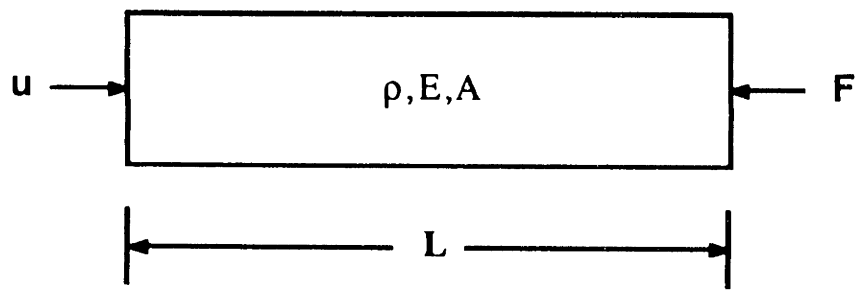

Figure 7.5: Uniform distributed model of a robot.

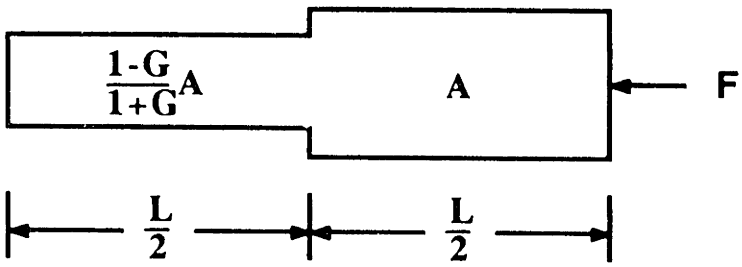

Figure 7.6: Passive physical equivalent of the uniform beam model.

too much. However, the insight that was created by the analyses of the last section may be used to guess a solution, which may then be checked by assuming a sinusoidal environmental input, $F = f\sin(\omega t)$, and showing that the endpoint behaviors of the controlled system and the realization match for all f and ω. The realization is shown in Figure 7.6.

Note that, whereas the element values of the lumped parameter models were scaled by feedback, the *area* of this model is scaled. In fact, this is done simply for the purposes of illustration, as both the mass per unit length and stiffness per unit length of a distributed beam are proportional to cross-sectional area. We could just as well have kept the area constant and have indicated that ρ and E were scaled by $(1-G)/(1+G)$.

This realization has the expected low frequency behavior—that the total mass is scaled by a factor of $1/(1+G)$. In addition, this realization has what should be the expected high frequency result—that the half of the beam close to the endpoint is unaffected by feedback. Why should this be the expected result? The answer follows from a consideration of transmission delays.

If the transmission time over one length of the uniform beam (Figure 7.5) is T, and if the sensing and actuation are instantaneous (as has been assumed throughout), then there will be no effect at the endpoint of the beam due to the control until time T has elapsed since the first effect due to the environment (assumed to occur at $t = 0$).

If the realization is to be consistent with this, which it must be to have the correct high frequency behavior, then there can be no change in the response to an endpoint disturbance until time T has passed. This requires that the realization be uniform over half its length, and that at its midpoint, a finite change occur. In this way, there will be a reflection from the midpoint, and a response will be felt at the endpoint after a delay of $T/2 + T/2 = T$.

This argument based on time delays is clearly independent of the particular control law that is used, so long as the control law is causal. Thus we have a fundamental result—the best that can be done without violating the positive real condition, given a uniform robot model, is to mask half the mass. In fact, this is rather optimistic, as typical digital controllers have lengthy computational delays of a few milliseconds, and analog controllers have finite delays as well.

7.4 Dynamic Compensation

In the last section it was demonstrated that, for the uniform robot, no controller could do better than mask half the mass without violating the positive real condition. But we may ask if it is possible to violate this condition in such a way (e.g., only at low frequency) as to cause no difficulty with stiff environments.

A reasonable place to start would be with the non-colocated two-mass model (Figure 7.2), and a first-order lag compensator:

$$u = -\frac{G}{\tau s + 1} F$$

Unfortunately, there is no "good" realization, meaning that there is no single realization which makes the closed loop structure and the dependence upon G and τ apparent. It should be understood that this is because we are seeking a *symbolic* realization. If plant parameters, a range of values for G and τ, and appropriate software were available,

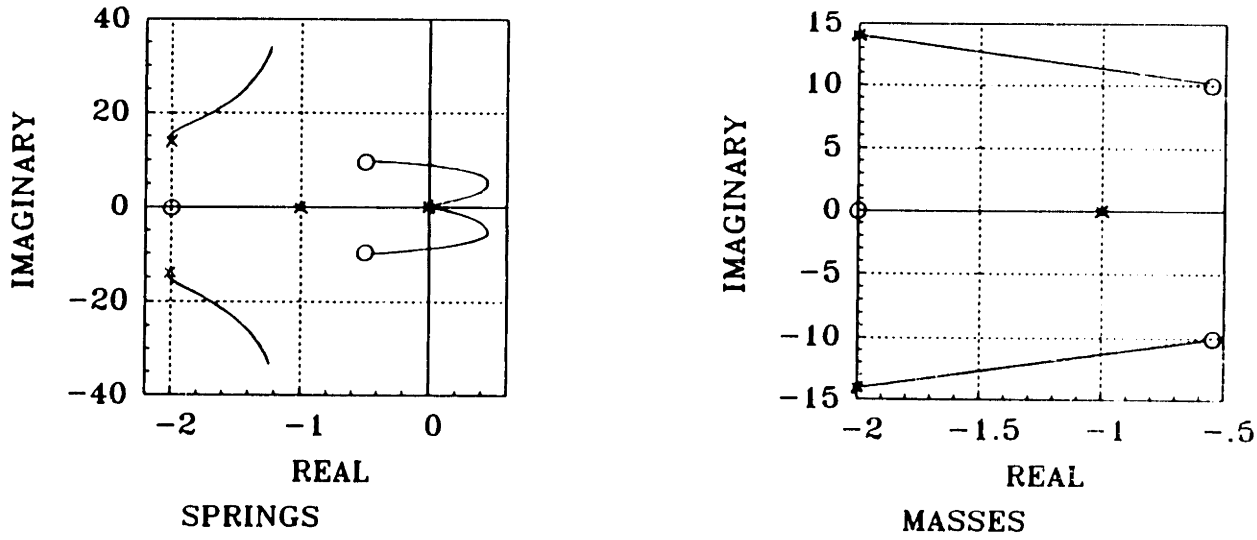

Figure 7.7: Worst environment root loci for first-order lag force feedback compensator and two mass plant. Parameters are $M = 1$, $B = 2$, $K = 100$, and $G = 1$.

passive physical equivalents would continue to be used. For the current purposes, however, a straightforward analysis of the driving point admittance will suffice. The closed loop admittance is:

$$Y_c(s) = \frac{\tau s^3 + (1 + B\tau/M)s^2 + ((1+G)B/M + K\tau/M)s + (1+G)K/M}{M\tau s^4 + (M + 2B\tau)s^3 + (2B + 2K\tau)s^2 + 2Ks}$$

The positive real condition requires the following in addition to $G \geq -1$:

$$(B\tau^2)\mu^3 + (GK\tau + (1-G)B - 2GB^2\tau/M)\mu^2 + (-2GK^2\tau/M)\mu \geq 0$$
$$\text{for} \quad 0 \leq \mu < \infty$$

where $\mu = \omega^2$. An analysis of $Y_c(s)$ and the passivity condition leads to the following results:

- At low enough frequency, the term $(-2GK^2\tau/M)\mu$ dominates, so that, for any $G > 0$, $\tau > 0$, the driving point admittance is not positive real. The consequence is that the closed loop system will always exhibit instability upon coupling to soft environments, as illustrated by the worst environment root loci in Figure 7.7.

- The centroid of the root locus for coupling to springs (see Section 4.2) always lies in the left half plane ($\sigma_{centroid} = -B/2M$), exactly where it lies in the case of

no control, regardless of the values of τ and G. In other words, this closed loop system is always stable when coupled to very stiff environments, so long as there are no nonminimum phase zeros. This disagrees with the results of Eppinger and Seering [21].

- There are gains, G, for which there are no nonminimum phase zeros. This is a possible explanation for the success that An [3] and Youcef-Toumi [96] reported with this compensator.

- The following relations, which can be found by applying a Routh-Hurwitz test to the numerator of $Y_c(s)$, guarantee minimum phase zeros[6]:

$$\begin{aligned} \tau &> 0, \quad G \geq -1 \\ G &\leq \frac{BM + 2B^2\tau + 2BK\tau^2}{(KM - 2B^2)\tau - BM} \quad \text{for} \quad \tau > \frac{BM}{KM - 2B^2} \\ G &\geq \frac{BM + 2B^2\tau + 2BK\tau^2}{(KM - 2B^2)\tau - BM} \quad \text{for} \quad \tau < \frac{BM}{KM - 2B^2} \end{aligned}$$

The final pair of conditions is not very simple to interpret. However, if we make the simplifying assumptions that $2B^2 \ll KM$ (low structural damping) and $1/\tau \ll \sqrt{2K/M}$ (control bandwidth well below first structural resonance), which would be reasonable for a typical robot, then the conditions above reduce to:

$$\tau > 0, \quad -1 \leq G \leq 2B\tau/M$$

The result is that increasing the compensator bandwidth $(1/\tau)$ decreases the allowable steady state gain. At low frequency, the driving point impedance is at best $Y(s) = 2M/(1 + 2B\tau/M)$. Thus it is possible that, for very slow control (large τ), some mass reduction beyond the limit of $M/2$ is possible. The addition of structural damping should also help.

[6]Section 8.4 provides an explanation for the generation of nonminimum phase zeros by input feedback compensators which add lag.

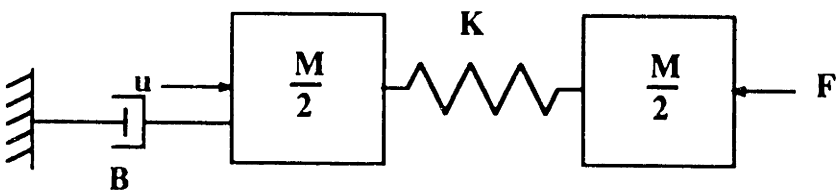

Figure 7.8: Two-mass model for the analysis of integral force control.

In conclusion, this form of compensation shows some promise for stable control of the force exerted on a rigid surface, but only when high gains are not involved.

Integral control ($u = -(G_I/s)F$) is another form of compensation which has received some attention in the past, particularly because of the promise of zero steady state error. Of course, since integral action should mask *all* inertia at low frequency, it is not reasonable to expect a passive physical equivalent. In fact, it is fairly easy to show that integral control applied to any of the models treated so far will lead to instability when coupled to environments of any stiffness.

In order for integral control to satisfy the positive real condition at all, the robot model must include some damping to ground. One such model is shown in Figure 7.8. Structural damping is left out of this model for simplicity's sake. Again, there is no good realization; nevertheless, the bottom line is passivity, and analytic limits on the values of G_I which will satisfy the positive real condition may be found. The limits are:

$$0 \leq G_I \leq \frac{B}{M}$$

At low frequencies, the driving point admittance is $Y(s) = G_I/Bs$; thus, at best, the apparent mass of the robot is M. In other words, integral control results in no mass reduction at all without jeopardizing coupled stability.

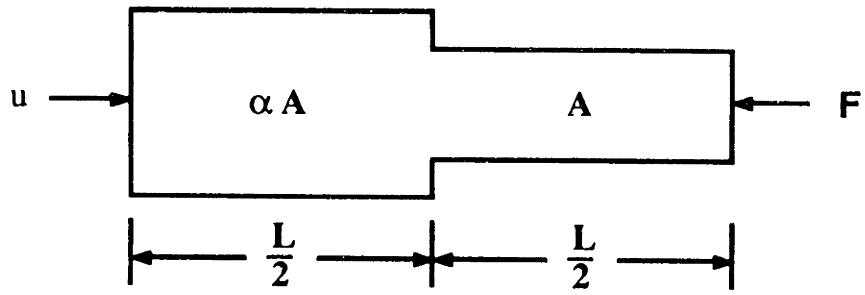

Figure 7.9: Non-uniform distributed beam model.

integral control, does, however, generate high gains at low frequencies—potential benefits of high gains which are independent of mass reduction are discussed in the next section. Note that, although the high gain is the result of integral action, the gain G_I may also be increased by increasing the viscous damping to ground, B.

7.5 Non-Uniform Manipulators

With a fundamental limit on what control can achieve in terms of masking a manipulator's inertia, an alternative approach to the force control problem is redesign of the manipulator. In this section we consider one type of redesign which fits neatly into the context of the models considered so far.

The non-uniform model shown in Figure 7.9 is considered. The area has been increased by some factor $\alpha > 1$ over the half close to the actuator. The reason for doing this is to offset the very strong dependence on the force feedback gain G of that half. Given such a manipulator and the control law $u = -GF$, the realization is shown in Figure 7.10.

The limits on G are now $-1 \leq G \leq \alpha$. There are two ways to interpret this result in terms of robot design. One is simply that it pays to put the bulk of the mass near the

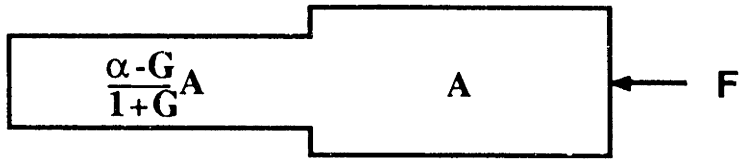

Figure 7.10: Passive physical equivalent to non-uniform beam model.

actuator, leaving the distal part of the robot quite light[7]. The second interpretation is that it would be useful to add weight near the actuator, even if nothing is done to the weight near the endpoint. No more mass could be masked than before, but a higher gain G could be used. One benefit of higher gain is that the steady state error of the force servo could be reduced. Note that if a set-point is included in the control law, $u = G(F_0 - F)$, then the steady state force seen by the environment is $GF_0/(1+G)$. If G is limited to 1, the steady state error is 50%, but if G is limited to α and α is large, the steady state error may be small.

An additional benefit of the higher gain is that it helps to decrease the effect of joint friction. Although joint friction has been left out of the models so far for simplicity's sake and because it does not seem to be a deterent to coupled stability (see the next section), it is a significant impediment to accurate low frequency force control. This is because some (load-dependent and, in general, unknown) amount of torque is needed to overcome joint friction, so that a virtual work relation between joint torques and endpoint forces is insufficient to accurately predict endpoint forces.

To investigate the tradeoff between α and the apparent joint friction, consider the model in Figure 7.11. Here, a viscous damper (b) plays the role of friction. Given the control law $u = -GF$, the low frequency admittance is $(1+G)/b$. Thus, increasing G decreases the apparent damping. Although the addition of this damper complicates

[7]Note that the robot end-effector is typically beyond the force sensor (usually a wrist sensor), and therefore does not add to the weight near the endpoint; i.e., this strategy does not prohibit end-effectors.

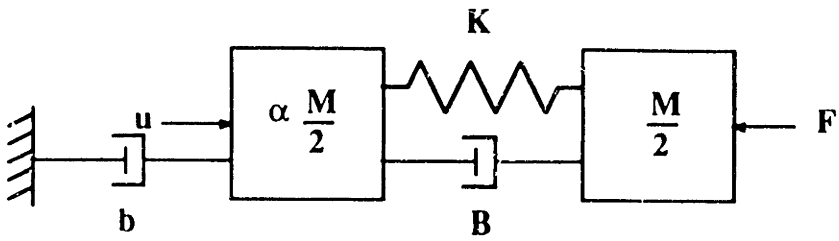

Figure 7.11: Two-mass model with damping to ground to illustrate the effect of joint friction.

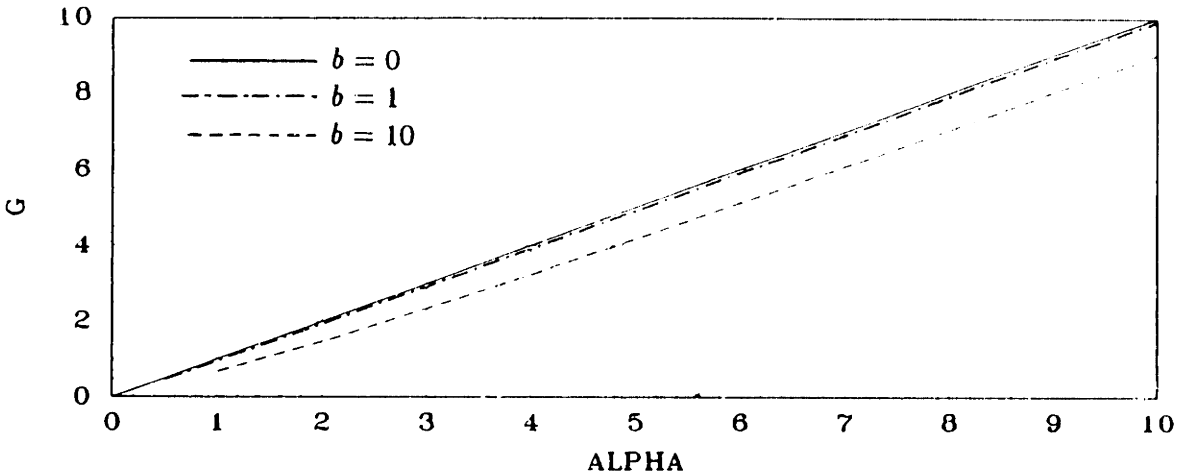

Figure 7.12: Maximum values of force feedback gain G plotted as a function of mass amplification factor α (see Figure 7.11). Parameter values are: $M = 2$, $B = 2$, $K = 10{,}000$. Values of b are indicated on plot.

the analysis somewhat, it turns out that the upper limit on G is approximately equal to α. Some example results of maximum G versus α are plotted in Figure 7.12.

In conclusion, the act of adding mass near the actuator can improve the steady state response by allowing larger gains, which have the effects of making the set-point $GF_0/(1+G)$ approach F_0, and of reducing the apparent joint friction. It should also be noted that the integral control described in the last section is theoretically capable of entirely eliminating viscous joint friction, but not while simultaneously reducing the inertia and providing a guarantee of coupled stability.

7.6 The Effect of a Computational Delay

The preceeding analyses have shown that a fundamental limit exists on the amount of inertia reduction which can be achieved with force feedback without risk of coupled instability. The models used, however, have been idealized in a number of important respects. One of these is that the effect of a computational time delay has not been considered. But, given the crucial role of inherent delays in the limitations of force feedback, it should be not be surprising that computational delays will also play an important role in real implementations.

To investigate this role, consider the two-mass model shown in Figure 7.2 and the control law $u(t) = -GF(t-\tau)$, where τ is the computational delay. In the frequency domain, $u(s) = -GFe^{-\tau s}$. The driving point admittance is now of infinite order, mixing transcendental functions and polynomials, and is rather difficult to analyze short of generating a Nyquist plot for a given set of parameter values. Fortunately, there is little need for a complete symbolic analysis. It has already been established that force feedback has a profound effect on the closed loop admittance at frequencies close to the structural resonance, and that this effect is due to an inherent storage delay. An added computational delay may exacerbate the effect, but should not change its nature. It has also been shown that high frequency behavior is dictated by the uncontrollable element no matter what causal compensation scheme is used. On the basis of these observations, we should expect that any interesting effects due to a computational delay should occur at frequencies below the structural resonance. In fact, interesting effects do occur, as a low frequency approximation to $Y_c(s)$ reveals:

$$\lim_{j\omega\to 0} Y_c(j\omega) = -\frac{G\sin(\omega\tau)}{M\omega} + j\frac{1+G\cos(\omega\tau)}{M\omega}$$

For $G = 0$, $\mathrm{Re}\{Y_c(j\omega)\} = 0$, which is due to the fact that this manipulator behaves as a mass at low frequency, and the Nyquist plot of a mass's admittance lies directly on the imaginary axis. However, for *any* $G > 0$, the Nyquist plot is shifted into the left half

plane (assuming that τ is small enough that $\sin(\omega\tau) > 0$), and the coupled stability condition is violated. For $G < 0$, on the other hand, the Nyquist plot is shifted into the right half plane, indicating apparent energy dissipation and behavior more robust to interaction than with no feedback at all! This result is consistent with the findings of Wlassich and Hogan [92,40].

The result that $G > 0$ cannot be used at all is, however, inconsistent with observation. This result can be reconciled by realizing that, although the effect of a delay is to create an apparent negative damping, some amount of energy dissipation is inherent to all real robots. For small enough force feedback gains, the inherent dissipation may offset this apparent negative damping. It is theoretically possible to supplement the inherent losses with velocity feedback[8]; however, when interacting with a stiff surface, velocities may be very low, straining the resolution of a velocity sensor.

In conclusion, some amount of low frequency energy dissipation[9] is a useful feature, as is a minimal time delay in the control implementation.

7.7 Summary: Improving Force Control

The main result of this chapter has been the demonstration of a fundamental limitation of force feedback controllers: given any manipulator with non-colocated actuation and force sensing, there is a limit to the amount of inertia reduction which can be achieved without risk of coupled instability. But the various analyses have also suggested a number of potential routes to improved force control. The object of this section is to summarize these routes.

[8] Computational delays do not have the same effect on velocity feedback that they do on force feedback. The Nyquist plot of a damper lies on the real axis, and a small angular shift (e.g., low frequency phase lag due to a computational delay) will not affect the coupled stability properties.

[9] Townsend and Salisbury [84] also find that joint friction can increase the region of stability of a force feedback controlled robot, but point out various other deleterious effects of the friction.

Before attempting to improve force control, however, it is important to decide what must be improved. For instance, this chapter has focused on mass reduction, and to a lesser extent, upon general admittance amplification. This is despite the fact that control system performance is typically rated in terms of bandwidth and disturbance rejection. Such specifications have not been used here because they make no sense in an interaction control problem unless a specific model of the environment is chosen so that force can be measured as an output. The properties of the admittance (e.g., total inertia), however, are well-defined with or without an environment.

Nevertheless, there is an intuitive mapping from mass reduction and admittance amplification to bandwidth and disturbance rejection. For instance, it is clear that (in the absence of static friction or grounded springs) the steady state force applied to the manipulator equals the steady state force applied to the environment; however, as the frequency of the actuator force increases, an increasing amount of effort is used to accelerate the manipulator rather than apply force to the environment. The bandwidth, therefore, scales inversely with the manipulator inertia.

Disturbances may arise as irregularities in the environment, or as undesirable features of the manipulator dynamics, such as joint friction. The admittance of the manipulator *is* its response to environmental disturbances—thus disturbance rejection correlates to admittance amplification. The masking of joint friction and other plant non-idealities also falls under the umbrella of admittance amplification.

If mass reduction is the primary design goal, it may be more fruitful to consider *design of the hardware* rather than design of the controller, because of the fundamental limit described in Sections 7.2 and 7.3. To increase the bandwidth, it is necessary to decrease physically the mass near the endpoint. One viable way of doing this is to use a macro/mini design—a small robot mounted on the end of a larger robot—and to close the feedback loop around the mini, thus reducing the mass to be masked [70,82].

If admittance amplification or steady state error is the primary goal, it is useful to

achieve high gain. For instance, integral control, while it does not result in any mass reduction, ensures that the manipulator behave as a perfect force source at d.c., so long as some amount of viscous damping is present at the joints. High gain and low steady state error can also be achieved without sacrificing mass reduction by adding mass near the actuator, as described in Section 7.5. We make special note of the fact that this is an approach to improved force control which would likely have gone unnoticed without the use of passive physical equivalents. Futhermore, it is an approach which requires redesign of the plant rather than redesign of the controller, a result which is quite natural in the context of passive physical equivalents.

The ability to achieve high gains may also be limited by computational delays. This problem may be addressed by using analog control (which can be digitally supervised, as described in [·]), or by using velocity feedback, as described in Section 7.6. It should be noted that the use of velocity feedback (usually in a fast inner loop) is quite common in actual implementations (see Section 7.1.). It may seem, however, that the goals of reducing joint friction and introducing viscous damping are somewhat at odds. But joint friction is a poorly characterized, nonlinear effect which can cause steady state error, whereas proportional velocity feedback is well characterized and causes no steady state errors. Thus it can make sense to replace the former with the latter.

Finally, of course, one need not use force feedback to implement force control. A conventional treatment of joint friction as well as other nonlinearities, such as motor cogging, is to implement a torque control loop around each joint. Although this eliminates structural dynamics as a source of non-colocation, it does not eliminate actuator or transmission dynamics. A close inspection of the literature would probably reveal that joint torque feedback has been most successful in non-contact tasks.

Another approach is to use impedance control to establish a known dynamic relation (impedance) between the position of the end of the robot and the force that it exerts on the environment [40]. Current implementations of impedance control, however, do

not ensure that the desired impedance is implemented, so that effects such as joint friction are still liable to cause significant steady state error. A final approach is to design robots that do a better job of controlling force in an open loop fashion [83,96], therefore avoiding control problems altogether.

Chapter 8

Techniques for the Design of Interaction Controllers

The last five chapters focused on the *analysis* of interactive behavior; this chapter focuses on the *implementation* of a desired interactive behavior. The term "implementation" is used to refer to the process of designing a specific controller for a specific plant, rather than the process of executing a specific design in hardware and software.

The design problem will be formulated as follows:

Interaction Control. Given a finite-dimensional, LTI plant model, design a controller to implement the following:

- $\|Z_c(j\omega) - Z_t(j\omega)\| < \epsilon \quad \text{for} \quad \omega < \omega_b$
- $Z_c(s)$ positive real

where $Z_c(s)$ is the closed loop impedance, $Z_t(s)$ is the *target impedance*, ϵ is an arbitrarily small number, and ω_b is the desired bandwidth over which Z_c and Z_t are "matched". The norm is a suitably chosen matrix norm, and the dimension of Z_t must equal that of Z_c.

A shortcoming of this formulation is that, although it addresses only interaction, manipulators must also move. Therefore, whenever possible, a *target model* [44] will

be chosen rather than a target impedance. This model will consist of a passive physical system with state-independent control which has the desired command following and interactive behavior. In this case, the formulation includes a term for command following:

Interaction and Command Following Control. Given a finite-dimensional, LTI plant model, and a positive real target model, design a controller to implement the following:

- $\|Z_c(j\omega) - Z_t(j\omega)\| < \epsilon_z \quad \text{for} \quad \omega < \omega_b$
- $\|C_c(j\omega) - C_t(j\omega)\| < \epsilon_c \quad \text{for} \quad \omega < \omega_b$
- $Z_c(s)$ positive real

where $C_c(s)$ is the closed loop transfer function, and $C_t(s)$ is the transfer function of the target model.

An alternative to the target model is to design a feedback controller to implement the desired interactive behavior, and a feedforward compensator to shape the command following transfer function. The feedforward compensator will not affect the closed loop impedance, but can be used to "invert" the closed loop transfer function to provide good command following. The method, however, is subject to the same sensitivity problems as computed torque methods [7], and will not be pursued here.

Another question that will not be treated here is that of selecting either $Z_t(s)$ or the target model. Hogan addresses the "impedance selection" problem at some length in [34]. He suggests several approaches (for linear as well as nonlinear systems) based upon optimization of particular objective functions, such as the power transmitted to the environment. Although these methods are not restricted to $2n$ order systems[1], these are the focus of his examples.

[1] A "$2n$ order system" is an n degree of freedom system, second order in each d.o.f., e.g., a serial rigid-link robot. Although redundant robots may have $n > 6$, a target model would generally have $n \leq 6$.

Kazerooni also addresses the selection of target dynamics. He focuses exclusively on $2n$ order systems, making the case that the three generalized parameters, M, B, and K, provide a designer control over the performance specifications of static stiffness, bandwidth, and stability.

Furthermore, the bulk of the references to impedance control that are found in the literature speak in terms of a generalized impedance as characterized by inertia, stiffness, and damping. However, this is not so (the distributed beam considered in the last chapter had an infinite order impedance), nor is is clear that target impedances should be restricted to $2n$ order systems. For instance, if there are multiple uncontrollable elements in the plant model (as in the three-mass and four-mass manipulator models of the last chapter), then it may make sense to include these in the plant model. Also, Sharon, et al. [69] make the postulate that:

> "Any physical system can be made to approximate the dynamic behavior of any other physical system of similar kinematic structure, provided that actuators and sensors can be placed at any point in the system."

The authors suggest that, although this is probably not feasible in practice, the appropriate selection of actuator and sensor locations as well as the appropriate selection of a target system make an approximate solution realizable. This philosophy promotes target models that may well be other than $2n$ order.

The following policy will be used in this chapter:

- When possible, the only restriction placed on $Z_t(s)$ will be that it is positive real.
- Examples, however, will be performed with second order target dynamics to provide more convenient comparison with available techniques.

8.1 Review of Available Techniques

The purpose of this section is to review the impedance control design techniques that have appeared in the literature. None of these techniques has addressed the problem of implementing a closed loop impedance such that $\|Z_c(j\omega) - Z_t(j\omega)\| < \epsilon$ for $\omega < \omega_b$, given an arbitrary LTI, controllable and observable plant, and an arbitrary positive real driving point impedance. The primary reasons for this are one, the perception that second order target dynamics are sufficient for most or all needs, and two, the prevalence of rigid body/ideal actuator robot models[2] in the literature. As described in the previous section, however, $2n$ order target dynamics are not necessarily always desirable. Furthermore, because the positive real condition requires evaluation of $Z(j\omega)$ at all ω, rigid body models are not necessarily sufficient—the force feedback induced contact instability described in the last chapter should be ample evidence of this shortcoming.

Methods in the literature may be classified according to the type of robot models they treat, ideal or non-ideal. Those who have studied the ideal robot problem include Hogan, Wlassich, Anderson and Spong, Goldenberg, and Hamilton. The impedance control implementation due to Hogan [35], which assumed an n-link ideal robot model and a $2n$ ($n \leq 6$) order linear target impedance, was described in Section 4.1.1. A similar implementation is investigated by Wlassich [92]. This controller differs from the one described by Hogan in that it does not require that the jacobian relating changes in joint angles to changes in endpoint positions be inverted. It generates a weighted pseudoinverse of the jacobian instead.

Anderson and Spong [5] describe a "hybrid impedance controller" which combines hybrid position/force control and impedance control. This technique makes use of the causality concept to select those axes along which the manipulator should behave as an admittance and those axes along which it should behave as an impedance. They

[2]henceforth referred to as "ideal robot models".

use an ideal robot model and second order target dynamics, although their formulation permits the introduction of higher order target dynamics in a rather straightforward manner.

Goldenberg [26,27] describes an approach to simultaneous implementation of force and impedance control. He also uses an ideal robot model and shows examples of second order target dynamics, but formulates the problem in such a way that higher order target dynamics could be used. It should be noted that the approaches of both Anderson and Spong and Goldenberg assume that an inverse of the robot's impedance (or admittance) can be implemented. While there is an assumption that cancellation occurs at low frequencies only it remains a dangerous approach in terms of coupled stability, for reasons which will be described in Section 8.4.

Finally, Hamilton [30] has also developed an approach to impedance control based on an ideal robot model and second order target dynamics. His method is essentially an extension (without actuator dynamics) of one proposed by Kazerooni [44], adapted to ensure global stability of the manipulator.

The impedance control method described by Kazerooni in his doctoral thesis [44] was the first to treat a non-ideal robot model. This method includes the effects of first-order actuator dynamics. It is shown in Section 4.1.3, however, that this approach fails to meet the coupled stability criterion, and is of little practical value.

More recently, Spong [74] has introduced a method for the implementation of impedance control on a flexible joint robot, i.e. a robot with rigid links and ideal actuators, but flexibility concentrated in the joints and transmission elements coupling the links to the actuators. The inclusion of joint flexibility rather than link flexibility appears to be a reasonable modeling assumption for many commercial robots. Spong selects second-order actuator dynamics. His method uses a singular perturbation [52] model of the robot—if the joint stiffness is k, the parameter $\mu = 1/k$ is defined, and it is shown that, as $\mu \to 0$, the rigid system is recovered. For $\mu = 0$, a *rigid manifold*

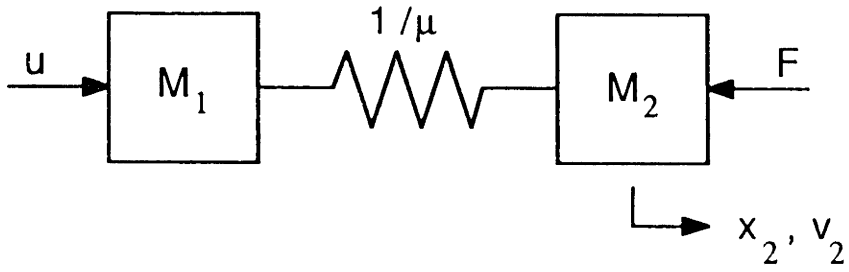

Figure 8.1: Single-axis, linear, flexible joint robot model.

is defined which relates the forces at the joints to the position and velocity states of the rigid system, as well as the control torques and forces imposed by the environment. Finally, the control is divided into slow and fast parts. The slow part implements the impedance control on the rigid manifold, and the fast part ensures that the joint forces remain on the manifold. A fast control is not needed if the manifold happens to be sufficiently attractive (which, however, is unlikely if the application is sufficiently demanding).

The slow control can be illustrated in terms of the single-axis linear model in Figure 8.1. The slow control for this model is:

$$u_{slow} = -Kx_2 - Bv_2 - \frac{\mu M_1 M_2}{M_1 + M_2}(K\dot{v}_2 + B\ddot{v}_2) + \frac{\mu M_1^2}{M_1 + M_2}\ddot{F}$$

where K and B are the desired stiffness and damping. An examination of the impedance for typical plant parameters shows that this control law does indeed recover the rigid body behavior and satisfy the coupled stability criterion. A potential cause for concern, however, is the need for measurements of jerk ($\ddot{v}_2$) and the second derivative of force. These are extremely difficult measures to make in practice due to noise amplification. Furthermore, a *causal* implementation of this controller will probably preclude the recovery of the rigid system at higher frequencies, so that the higher-order derivatives may be superfluous. It will probably take an actual hardware implementation to understand these issues fully, however, the method shows considerable promise.

By far the most general approach to implementing impedance control is the recent method due to Kazerooni [45]. Kazerooni's robot model is two nonlinear mappings—

one (G) from a vector of input trajectories to the robot position vector, and another (S) from a vector of environmentally imposed forces to the position—each subject to L_p boundedness. The environment is also treated as a nonlinear mapping (E) subject to L_p boundedness, which maps position to force. The reason for such an approach is to develop a methodology general enough to apply to all existing industrial manipulators, including those with significant nonlinearities, actuator and transmission dynamics, and any other difficult-to-model features. In particular, the hope is to implement impedance control on the many robots which have been designed as good positioning systems.

This approach may be readily understood by considering the linear case, which is also treated in detail by Kazerooni. The mappings G, S, and E become transfer function matrices, and their L_p norms are replaced with singular value norms. Because G and S represent a good positioning robot, the following conditions hold:

$$\sigma_{max}(G - I) < \epsilon_g \quad \text{for all } \omega < \omega_b \tag{8.1}$$

$$\sigma_{max}(S) < \epsilon_s \quad \text{for all } \omega < \omega_b \tag{8.2}$$

where σ_{max} is the maximum singular value[3], I is the identity matrix, ϵ_g and ϵ_s are arbitrarily small numbers, and ω_b is the bandwidth. Only the first of these two conditions (command following) is necessary in a world free of disturbances, the second supplements command following with disturbance rejection. The impedance controller proposed by Kazerooni is shown in Figure 8.2. The transfer function matrix H is a compensator to be designed. However, H also embodies the desired impedance, at least at low frequency when $G \approx I$ and $S \approx 0$.

Kazerooni provides a stability analysis of the *coupled system* based on the multivariable Nyquist theorem. He derives the following stability condition:

$$\sigma_{max}(GHE) \leq \sigma_{min}(SE + I) \quad \text{for all } \omega \tag{8.3}$$

[3] $\sigma_{max} = \max\{\|G\mathbf{x}\|/\|\mathbf{x}\|\}$, $\mathbf{x} \neq 0$, where $\|\mathbf{x}\| = \sqrt{\mathbf{x}'\mathbf{x}}$.

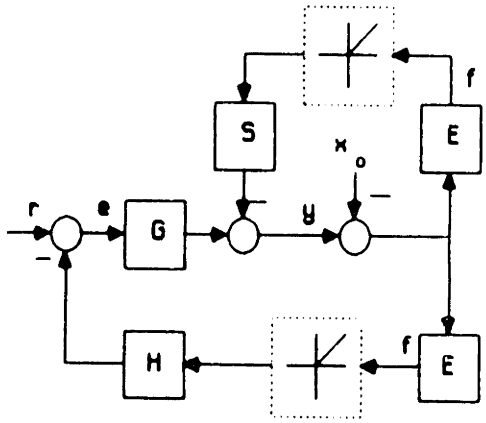

Figure 8.2: Structure of Kazerooni's impedance controller. The nonlinear blocks represent the fact that negative forces will not occur because the manipulator will leave the surface. Taken from [45].

This result does not require that G and S describe a good positioning system, and any compensator which satisfies this condition is acceptable. A method for selecting a particular H or relating it to a set of target dynamics is not discussed. We will return to the discussion of this method in Section 8.4.

8.2 A Naive Approach

The purpose of this section is to demonstrate that, given a particular plant model and control law, the positive real condition can be explicitly enforced to ensure coupled stability. However, it will also be shown that enforcing coupled stability in no way guarantees that the target impedance is recovered.

The example system that will be used for the analyses in this section as well as the next two is shown in Figure 8.3. This system is a linear representation of a manipulator. The spring k_2, the damper b_2, and the mass M_2 represent, roughly, the "hardware" dynamics, i.e. factors such as configuration-dependent forces (e.g., gravity), friction at the joints, and link inertias. The damper b_1 and the mass M_1, on which the control E acts, represent the actuator dynamics, and the spring k_2 represents the transmission

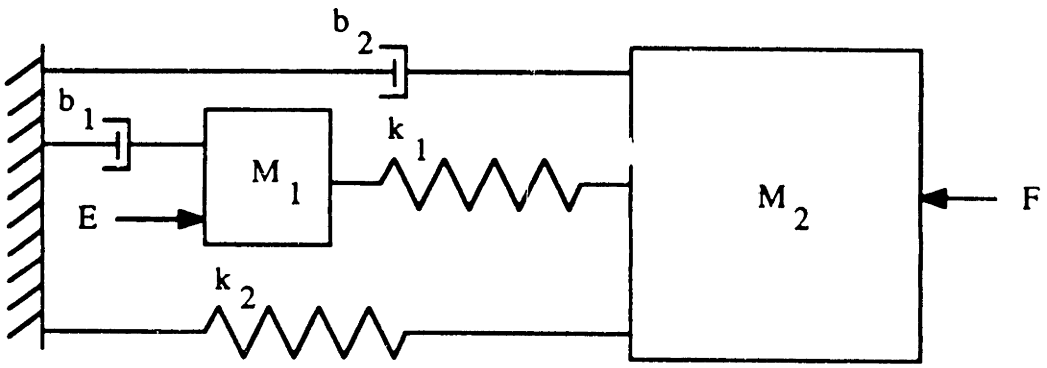

Figure 8.3: Linear manipulator with actuator and transmission dynamics.

dynamics. F represents forces imposed by the environment. This model, while crude, will prove to be reprentative of some of the difficulties of controlling the interactive behavior of a manipulator which has significant actuator and transmission dynamics.

State equations of this manipulator may be written as follows:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{v}_1 \\ \dot{v}_2 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -k_1/M_1 & k_1/M_1 & -b_1/M_1 & 0 \\ k_1/M_2 & -(k_1+k_2)/M_2 & 0 & -b_2/M_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ v_1 \\ v_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1/M_1 \\ 0 \end{bmatrix} E + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1/M_2 \end{bmatrix} F \tag{8.4}$$

where x_1 and v_1 refer to the position and velocity of mass M_1, and x_2 and v_2 refer to M_2.

The following control law, motivated by the simple impedance controller discussed in Section 4.1.1, will be considered:

$$E = -K(x_2 - x_0) - B(v_2 - v_0) \tag{8.5}$$

Were it not for the actuator and transmission dynamics (i.e., if E acted directly on M_2), the effects of this control law would be to adjust the stiffness and damping of the

manipulator, while preserving the coupled stability condition for all positive K and B, and to introduce x_0 and v_0 which would act as the "virtual trajectory" [36]. However, as will be demonstrated, these are not the actual effects.

The naive approach that will be taken in this section is to explicitly determine constraints on the values of the feedback gains K and B that will ensure the satisfaction of closed loop stability and closed loop coupled stability criteria. Closed loop stability constraints can be derived from the Routh-Hurwitz array, and coupled stability contraints can be derived from the positive real condition. The results, expressed as explicit and inequality constraints on the values of K and B are:

Closed loop stability:

1.

$$-b_2 - b_1\frac{k_1+k_2}{k_1} \leq B \leq \frac{b_1^2 b_2 + k_1 b_1 M_2}{k_1 M_1} + \frac{b_1 b_2^2 + b_2 M_1(k_1+k_2)}{k_1 M_2}$$

2.

$$K \geq -k_2$$

3.

$$K \leq \frac{\alpha_1(B)\alpha_2(B)}{\alpha_3} - k_2$$

$$\begin{aligned}
\text{where: } \alpha_1(B) &= b_1(k_1+k_2) + k_1(b_2+B) \\
\alpha_2(B) &= b_1^2 b_2 M_2 + b_1 b_2^2 M_1 + b_1 k_1 M_2^2 \\
&\quad + b_2 M_1^2(k_1+k_2) - k_1 M_1 M_2 B \\
\alpha_3 &= k_1(b_1 M_2 + b_2 M_1)^2
\end{aligned}$$

Coupled stability:

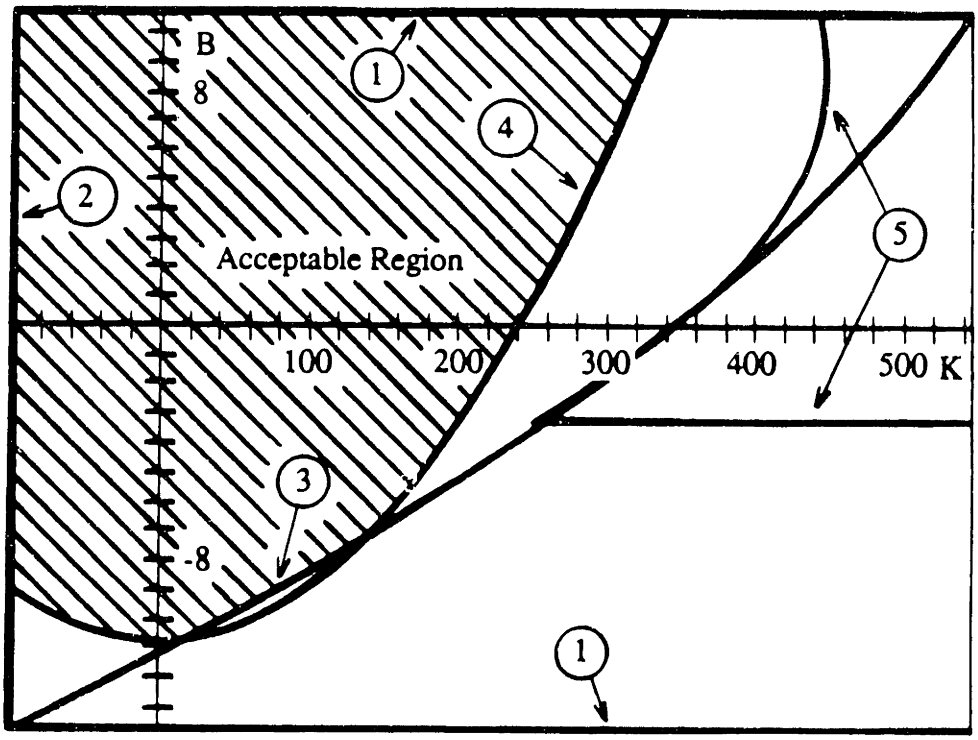

Figure 8.4: Feedback gains which yield closed loop and coupled stability. Numbers refer to inequality constraints in text. Note that constraint 5 does not come into play.

4.

$$k_1^2(b_1 + b_2 + B) - b_1 K^2 \geq 0$$

5.

$$\text{while} \quad \beta(B) = b_1^2 b_2 - b_2 k_1 M_1 \quad k_1 M_1(b_2 + B) \leq 0,$$

$$\text{then:} \quad \beta^2(B) - 4b_2 M_1^2(k_1^2(b_1 + b_2 + B) - b_1 k_1 K) \leq 0$$

Figure 8.4 is a graphical representation of these inequalities given the following parameter values:

b_1	=	4.8	b_2	=	6.0
k_1	=	160.0	k_2	=	100.0
M_1	=	0.1	M_2	=	1.0

As a concrete example of the design process, consider the target model shown in Figure 8.5 which embodies the desired behavior of the manipulator. The parameters

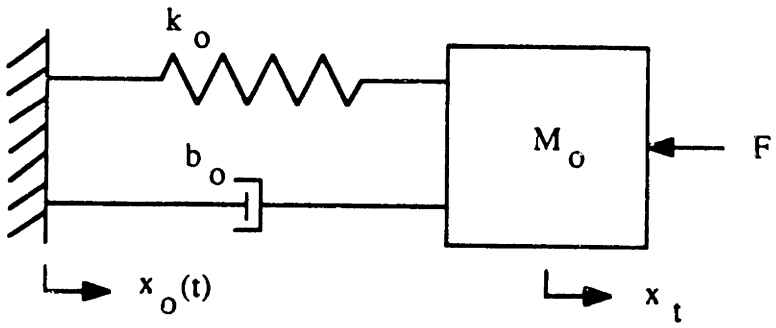

Figure 8.5: Target model.

chosen for this target model are:

$$\begin{aligned} b_0 &= 4.24 \\ k_0 &= 9.0 \\ M_0 &= 1.0 \end{aligned}$$

These parameters correspond to a considerably more compliant system than the uncontrolled manipulator—a reminder that it is not always desirable in the case of manipulation to generate the very high stiffnesses that a servo typically would.

One technique for choosing K and B is to ignore the actuator and transmission dynamics of the manipulator, in which case (since $M_2 = M_0$) the very simple conditions $k_2 + K = k_0$ and $b_2 + B = b_0$ appear to generate the target system exactly. The results are $K = -91$ and $B = -1.86$[4]. These values fall within the acceptable region of the K–B plane shown in Figure 8.4, guaranteeing closed loop and coupled stability.

As suggested in the introduction to this chapter, performance may be evaluated in the frequency domain by comparing magnitude plots of the closed loop and target systems. Magnitude and phase plots of $C_c(s) = x_2(s)/x_0(s)$ and $C_t(s) = x_t(s)/x_0(s)$ are shown in Figure 8.6. Although no attempt will be made to provide a quantitative assessment, it is clear that the command following of the manipulator is not very good.

The magnitude and phase plots of $Y_c(s) = v_2(s)/F(s)$ and $Y_t(s) = v_t(s)/F(s)$ are

[4]Note that this approach encourages a significant departure from servo philosophy: we are using positive feedback to add negative stiffness and negative damping.

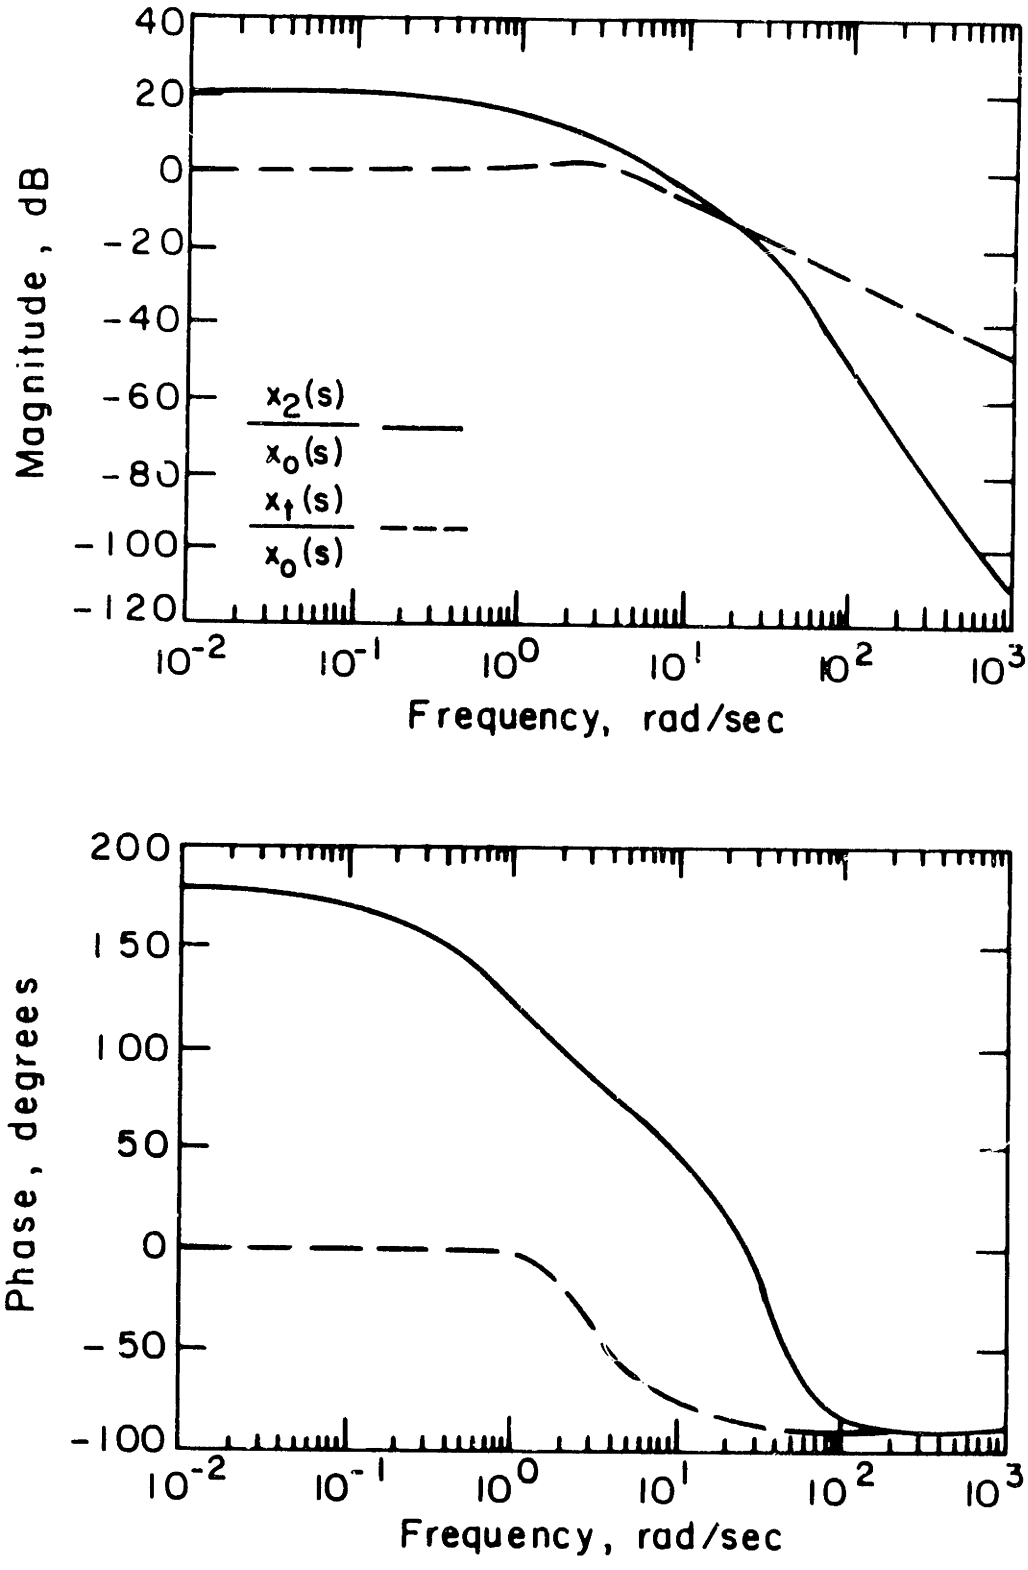

Figure 8.6: Command following performance: naive controller.

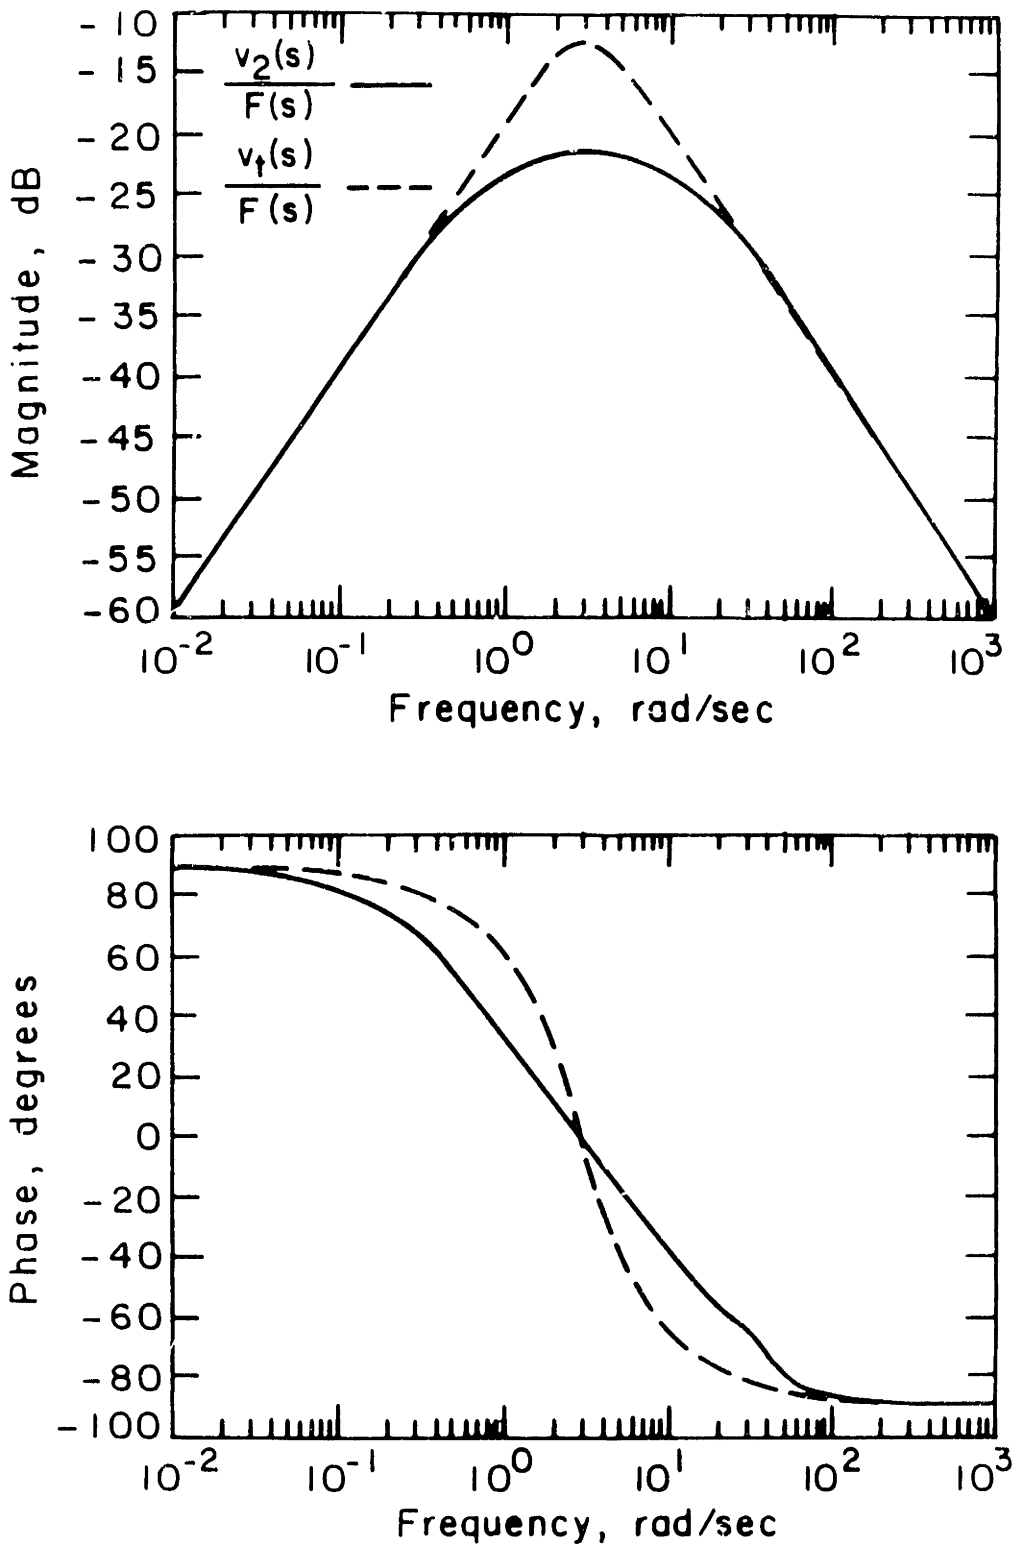

Figure 8.7: Interactive behavior: naive controller.

shown in Figure 8.7. The interactive behavior is promising at low and high frequencies, but shows a significant departure from that of the target system in the middle range.

The point of this analysis has been that ensuring coupled stability is possible given a particular control law, but that this does not necessarily constitute an effective design procedure. In the next section, a technique which attempts to satisfy both stability and performance specifications will be considered.

8.3 Servo-Masking

The purpose of this section is to demonstrate one systematic approach to the design of an interaction controller which addresses both stability and performance requirements.

We will begin with a foray into nonlinear controller design. Assume that we have a nonlinear manipulator which may be described by equations in the same canonical form as that used for the plant model in the last section:

$$\begin{aligned} \dot{x}_1 &= v_1 \\ \dot{x}_2 &= v_2 \\ \dot{v}_1 &= f_1(x_1, v_1, x_2, v_2, E) \\ \dot{v}_2 &= f_2(x_1, v_1, x_2, v_2, F) \end{aligned} \tag{8.6}$$

As before, x_1 and v_1 correspond roughly to actuator and transmission states, and x_2 and v_2 correspond roughly to hardware (linkage) states. E is the control input, and F is an environmentally imposed force. We will also define a nonlinear target model which contains no actuator or transmission dynamics and is, therefore, of lower order than the manipulator model:

$$\begin{aligned} \dot{x}_t &= v_t \\ \dot{v}_t &= f_t(x_t, v_t, x_0, v_0, F) \end{aligned} \tag{8.7}$$

x_0 and v_0 represent the virtual trajectory, and F is the environmental force. One way of deriving a control law is to equate the motion of the manipulator to that of the target model:

$$x_2 = x_t \tag{8.8}$$

$$v_2 = v_t \tag{8.9}$$

$$f_2(x_1, v_1, x_2, v_2, F) = f_t(x_2, v_2, x_0, v_0, F) \tag{8.10}$$

Equations 8.9 and 8.10 represent successive differentiations of equation 8.8, the "constraint relation", which defines the desired behavior. At this point, there are various paths from which to choose. For instance, equation 8.10 may be differentiated again, as often as necessary, to expose the control E. Presumably, it would then be possible to solve for E to yield a control law.

There are problems with this approach, however. One is simply that it may not be possible to solve explicitly for E. Another problem is that implementing a particular control law does not guarantee that the equation which was differentiated to obtain it is implemented—there is an arbitrary constant of integration that occurs in the reverse process. This typically results in behavior like that caused by a pole at the origin—slow drift. For instance, if we succeed in implementing $v_2 = v_t$, then $x_2 = x_t$ is not guaranteed; the values are likely to drift apart. Of course, multiple differentiations exacerbate the problem.

As an aside, sliding mode control [73] is a technique which ensures that the constant of integration is (on average) zero. However, it works only when the desired behavior and the control law are separated by a single differentiation; the author knows of no technique that may be used when multiple differentiations are needed to derive the control law.

The problem of constants of integration is a difficult one, and we will not attempt to deal with it here. However, rather than make matters worse by differentiating equation

8.10 again, this equation will be taken as a prescription for a servo design. That is, we will require that f_t is of such a form that we may separate the virtual trajectory terms and write:

$$f_2(x_1, v_1, x_2, v_2, F) - f_{to}(x_2, v_2, F) = f_{ti}(x_0, v_0) \tag{8.11}$$

The left side of this equation may be taken to define the output of a servo, and the right side the input, and we may use any available servo design technique to implement it. This technique will be called "servo-masking," as the effect of the servo is to mask undesirable dynamic effects in the plant model. Notice that the technique requires, in general, the use of full state feedback as well as the measurement of interaction port variables, such as the force applied to the manipulator.

This technique may be applied to the linear manipulator and target model described in the last section. The result is the following equation:

$$\begin{aligned}\left(\frac{k_t}{M_t}\right) x_0 + \left(\frac{b_t}{M_t}\right) v_0 \;=\; & \left(\frac{k_1}{M_2}\right) x_1 + \left(\frac{k_t}{M_t} - \frac{k_1 + k_2}{M_2}\right) x_2 \\ & + \left(\frac{b_t}{M_t} - \frac{b_2}{M_2}\right) v_2 + \left(\frac{1}{M_2} - \frac{1}{M_t}\right) F \end{aligned} \tag{8.12}$$

For the particular parameter values chosen in the last section, equation 8.12 reduces to:

$$9x_0 + 4.24v_0 = 160x_1 - 251x_2 - 1.76v_2 \tag{8.13}$$

Almost certainly the simplest method of implementing this equation is to use classical control techniques to design a proportional controller, as shown in Figure 8.8. An important advantage of this technique is that it enables the use of a single parameter, G, to describe the allowable spaces for closed loop and coupled stability. The results are:

$$\text{Closed loop stability:} \quad G \geq -0.848$$

$$\text{Coupled stability:} \quad G \geq -0.737$$

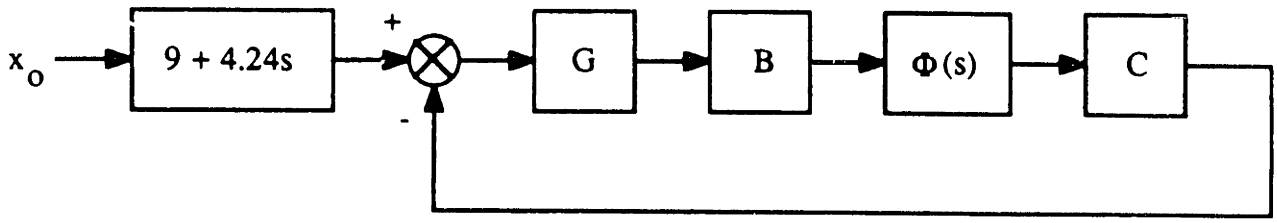

$$C = [160\ -251\ 0\ -1.76]$$

Figure 8.8: Block diagram of "servo-masking" interaction controller.

There is no upper bound on G which means that, for this problem, we may come arbitrarily close to matching both the command following and interactive behavior of the target model. Of course, actuator limitations provide practical limits on the value of G.

Figures 8.9 and 8.10 show the manipulator performance in terms of Bode plot comparisons with the target model when $G = 100$. In this case, both the command following and the interactive behavior are excellent. As might be expected, rolloff and accompanying phase lag limit the command following bandwidth. There is, however, no bandwidth limitation on the interactive behavior. This is primarily due to the selection of a target model sufficiently similar to the manipulator model. Both models include a unity mass at the interaction port (this mass is an uncontrollable element of the manipulator), and this mass is principally responsible for the driving point behavior at high frequency.

In summary:

- The servo masking approach demonstrates that it is possible to achieve the target impedance. The target impedance, if chosen properly, can be achieved at all frequencies.

- Servo masking requires, in general, full state and input feedback, which is not usually feasible.

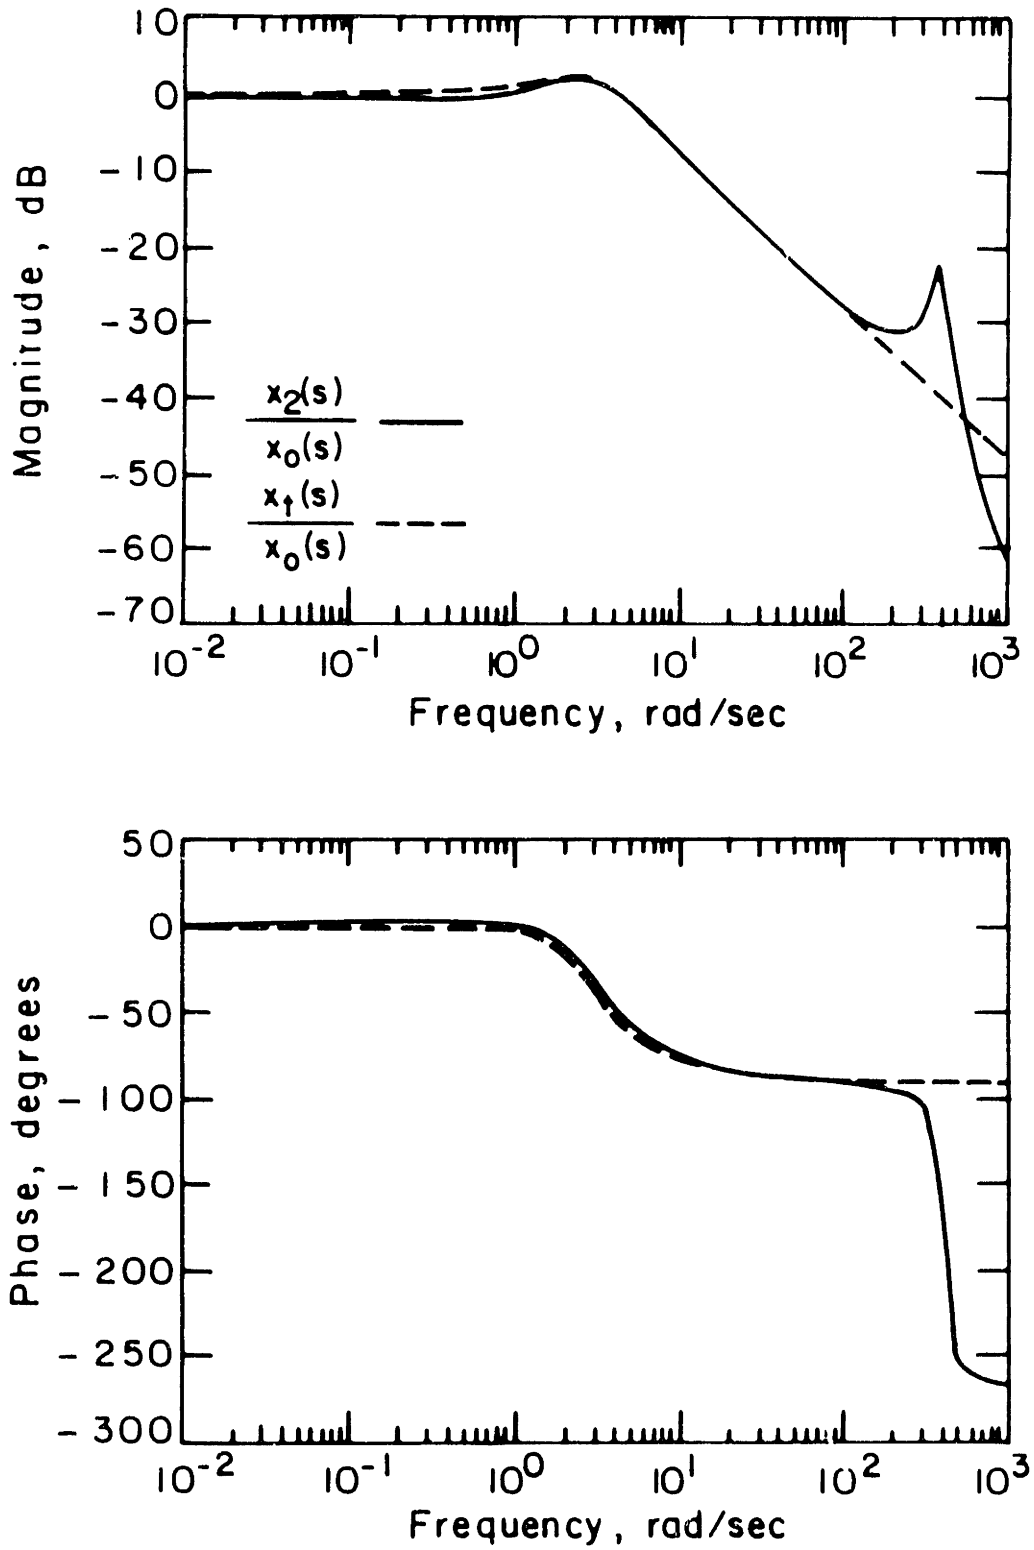

Figure 8.9: Command following performance: servo-masking controller.

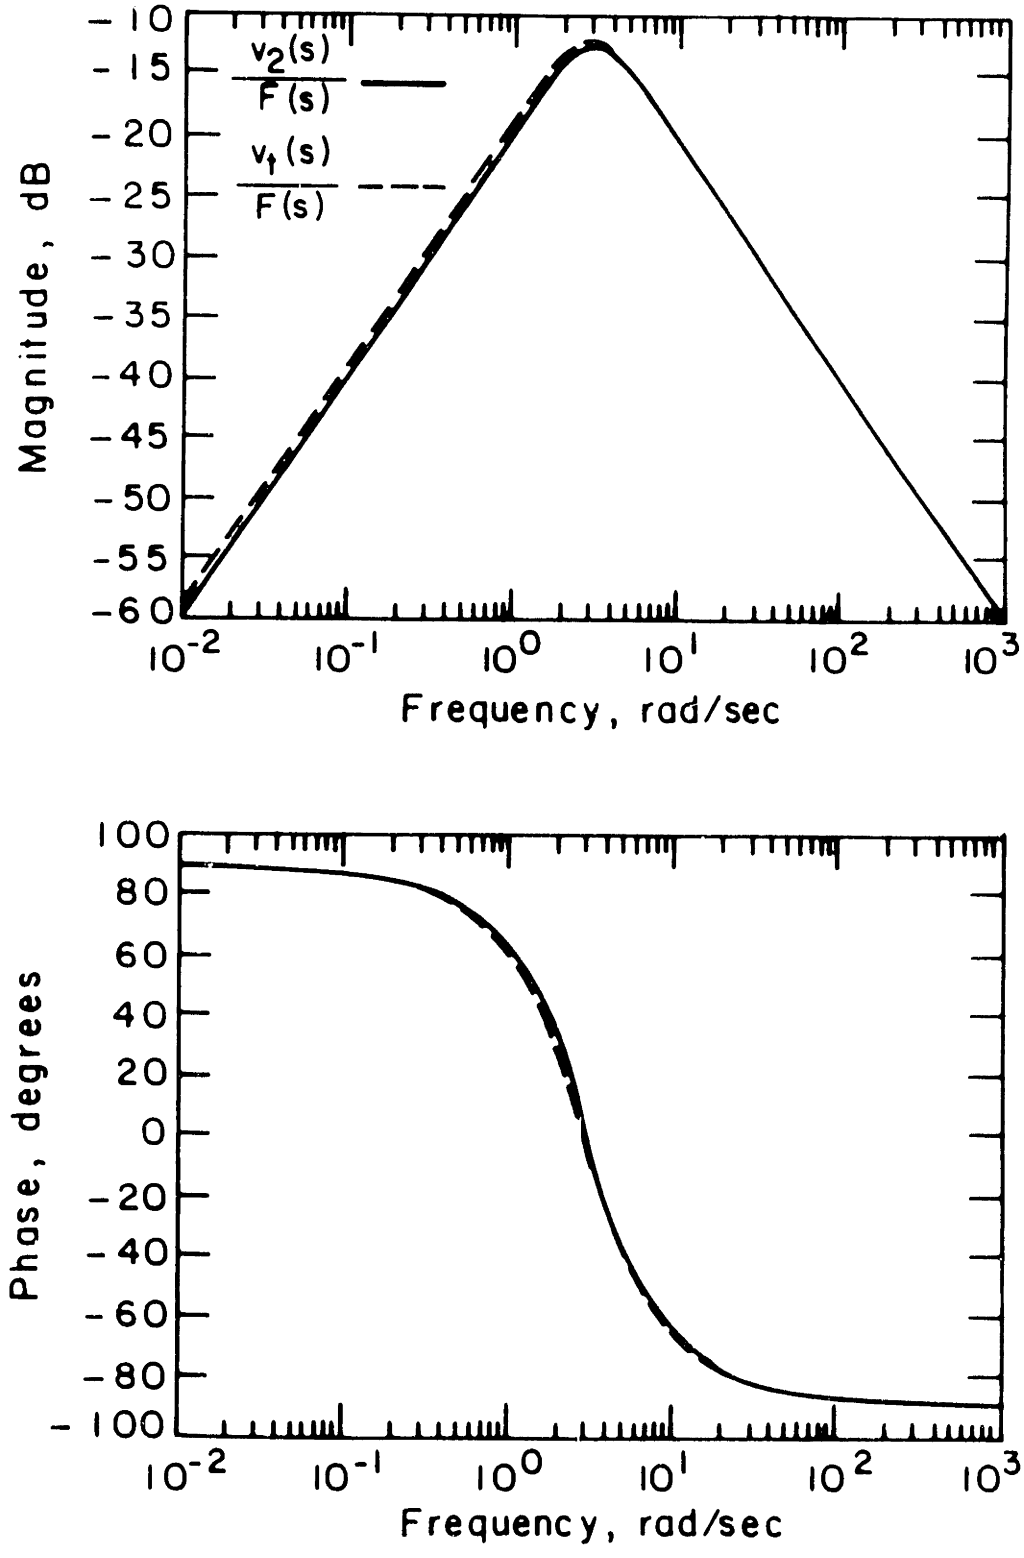

Figure 8.10: Interactive behavior: servo-masking controller.

- Servo masking is a systematic procedure. It does not, however, guarantee that the closed loop behavior will converge upon the desired result. The success of the method will degrade if the structure of the target model differs substantially from that of the plant model, if uncontrollable elements of the plant model are not represented in the target model, or if successive differentiations of the constaint relation are needed.

8.4 The Target Model Referenced Controller

This section describes an interaction controller implementation which employs a servo design loop and a pre-compensator composed of a target model. Although this method was independently conceived, it is identical in structure to Kazerooni's second method described in Section 8.1. The analysis presented here, however, is quite different from that presented by Kazerooni. Comparisons will be offered.

A general block diagram of the target model referenced controller (TMRC) is shown in Figure 8.11. The plant and target model as illustrated have admittance causality, however, the procedure is identical for impedance causality. This structure is stimulated by a very simple philosophy: a servo controller with good command following and disturbance rejection properties, designed with whatever method is deemed best, is implemented to control the interaction port output ($\mathbf{f}$). A target model with the desired interactive and command following behavior is also designed. The target model takes a reference command and a measure of the interaction port effort ($\mathbf{e}$) as inputs. It outputs the desired interaction port flow ($\mathbf{f}_d$), which is an input to the servo. Thus, over some limited bandwidth, the command following and interactive behavior of the closed loop system should match those of the target model.

Note that this structure applies to linear or nonlinear systems (Kazerooni performed

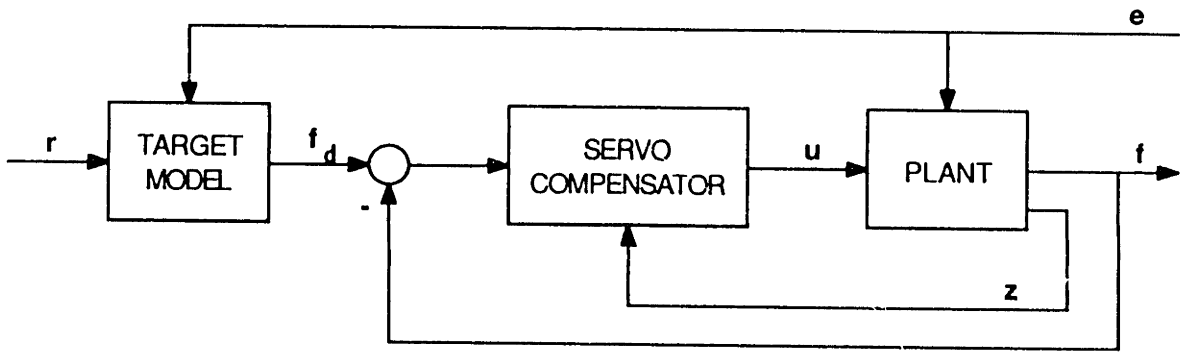

Figure 8.11: Block diagram of the target model referenced controller. **r** is the reference input to the target model, **e** is the environmental input to both the plant and target model, **f** is the interaction port output of the plant, $\mathbf{f}_d$ is the output of the target model, **u** is the control input to the plant, and **z** is a vector of measured outputs which can be fed back to improve servo performance.

a nonlinear stability analysis of this type of controller). For linear systems, it is easily seen that the pre-compensator does not affect the closed loop stability so long as it is stable itself. The closed loop coupled stability, however, requires a closer look.

To facilitate the analysis of coupled stability, several transfer functions are defined. $Y(s)$ is the admittance of the plant; $Y_t(s)$ is the admittance of the target model; $Y_s(s)$ is the admittance of the servo controller in the absence of a pre-compensator; $C_s(s)$ is the closed loop transfer function of the servo—it maps $f_d(s)$ to $f(s)$. The admittance of the complete system is:

$$Y_c(s) = C_s(s)Y_t(s) + Y_s(s) \tag{8.14}$$

Clearly, for a good servo:

$$C_s(s) \approx I \quad \text{for} \quad \omega < \omega_b \tag{8.15}$$

$$Y_s(s) \approx 0 \quad \text{for} \quad \omega < \omega_b \tag{8.16}$$

These equations may be compared to equations 8.1 in Section 8.1. Thus:

$$Y_c(s) \approx IY_t(s) + 0 = Y_t(s) \quad \text{for} \quad \omega < \omega_b \tag{8.17}$$

This is, of course, the desired result.

As an example of the application of this method, a TMRC will be designed for the plant and target model introduced in Section 8.2. The servo is designed with the LQR method. Although the LQR method does not currently receive widespread use, it remains an effective technique given an accurate model and access to all system states, both true in this academic exercise. In particular, the "cheap control" LQR problem [54] is solved, i.e., the following objective function is chosen:

$$J = \frac{1}{2}\int_0^\infty (\mathbf{x}'C'C\mathbf{x} + \rho\mathbf{u}'\mathbf{u})\,dt$$

where C is a matrix which weights selected outputs, and ρ is a control weighting which is adjusted until desirable behavior is achieved. C is chosen to be [1 10 0 0] for this problem, i.e., the highest penalties are placed on position deviations of the endpoint mass. Because the servo is designed to control endpoint position rather than velocity it must take a reference position as the input. This simply requires that the chosen output of the target model be position, but it does not require any change in the behavior of the target model. All closed loop admittance plots are of the true admittance, i.e., force to velocity.

In the first case considered, $\rho = 10^{-6}$ is selected as the control weighting. The results are shown in Figures 8.12 and 8.14(a). Figure 8.12 displays the Bode plots of $Y_t(s)$ and $Y_c(s)$, as well as the Bode plots of $C_s(s)Y_t(s)$ and $Y_s(s)$. The magnitude plot shows that the closed loop admittance matches the target admittance at all frequencies except for a small deviation in the vicinity of $\omega = 40\,\text{rad/sec}$. It is further apparent that, for frequencies below $\sim 40\,\text{rad/sec}$, the main contributor to $Y_c(s)$ is $C_s(s)Y_t(s)$ (the two plots are indistinguishable), and that, for frequencies above $\sim 40\,\text{rad/sec}$, the main contributor to $Y_c(s)$ is $Y_s(s)$. These results are all as expected.

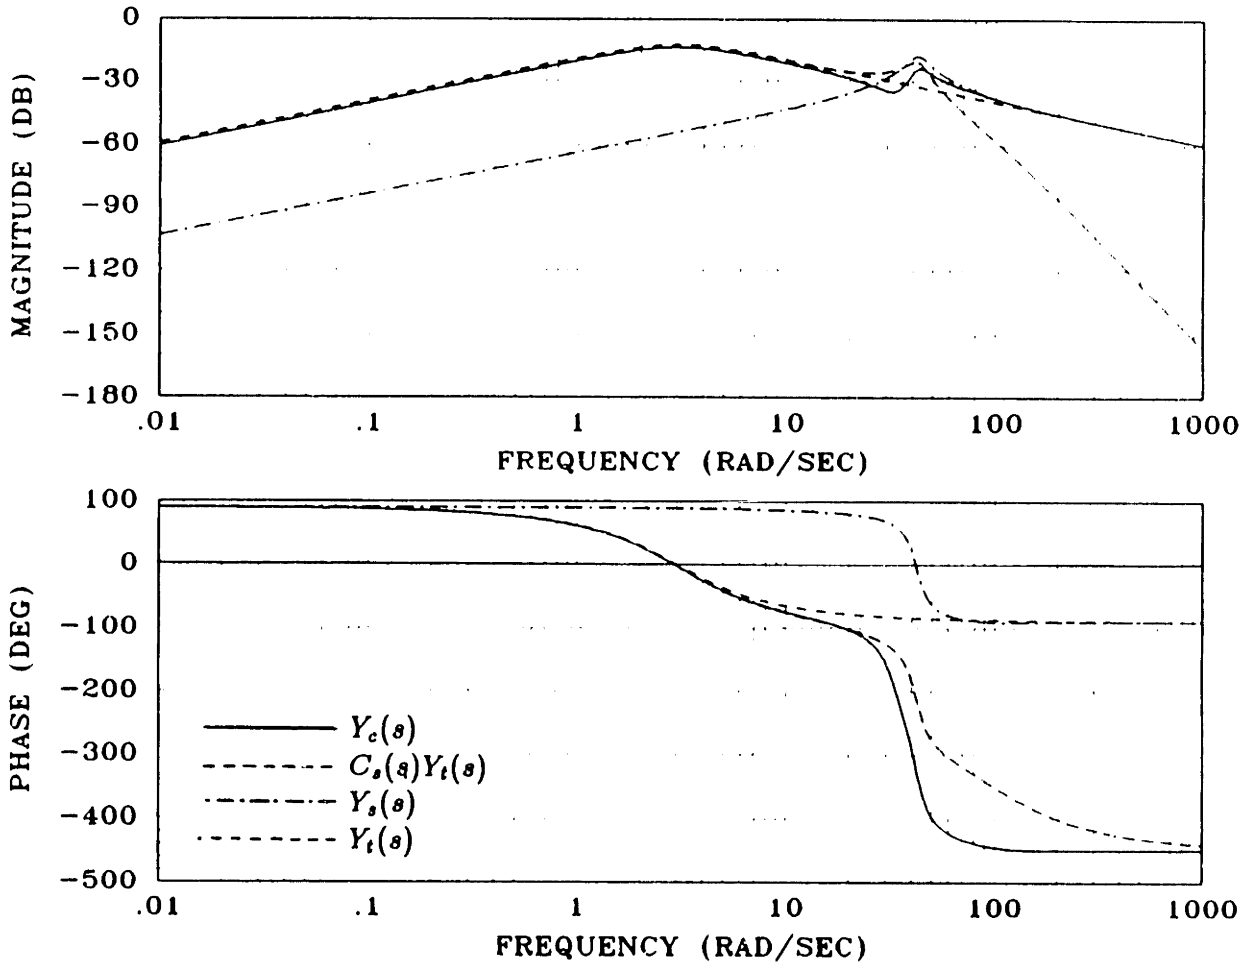

Figure 8.12: Bode plots of TMRC behavior, $\rho = 10^{-6}$.

The phase, however, behaves in an unexpected fashion. Although the phase of $Y_c(s)$ tracks that of $Y_t(s)$ at low frequencies, at high frequencies it roughly follows that of $C_s(s)Y_t(s)$, which eventually exhibits a phase lag of 450°. The result is that $Y_c(s)$ is clearly not positive real. Certain comments are in order. The first is that the *expected* behavior at high frequencies, due to rolloff in $C_s(s)$, is $Y_c(s) \rightarrow Y_s(s)$. This is desirable behavior only if $\mathrm{Re}\{Y_s(j\omega)\} \geq 0$ for large ω. Although this is true in this example, we know from the analysis of Section 4.4 that it need not be.

The second comment is that an apparent inconsistency exists. It is shown in Appendix D that a broad class of causal compensators, which includes the one considered here, cannot change the relative order of $Y(s)$, the admittance of the plant, and that the following condition must hold:

$$\lim_{s\rightarrow\infty} sY_c(s) = \lim_{s\rightarrow\infty} sY(s) = \alpha > 0$$

But the Bode and Nyquist plots clearly indicate that the high frequency behavior of $Y_c(s)$, which is characterized by a phase lag of 450°, is quite different from that of $Y(s)$, which must have a phase lag of only 90°. The only resolution to this conflict is that some of the zeros of $Y_c(s)$ must be *nonminimum phase*, therefore adding phase lag rather than lead to the high frequency result. A closer look at $Y_c(s)$ reveals that it does, indeed, contain two nonminimum phase zeros.

Nonminimum phase zeros are an insidious violation of the positive real condition. If one thinks of coupled stability in terms of the root locus test (Section 3.4), it is clear that some branches of either $sY_c(s)$ or $Y_c(s)/s$ must eventually converge on the nonminimum phase zeros. Therefore, a large enough stiffness or a large enough mass will lead to instability. The same is true of a large enough viscous damping. This is clearly behavior to be avoided.

Figure 8.15 provides an explanation of the nonminimum phase behavior. The Nyquist plots of $C_s(s)Y_t(s)$ and $Y_s(s)$ are shown. It is clear that, even if $Y_s(s)$ is posi-

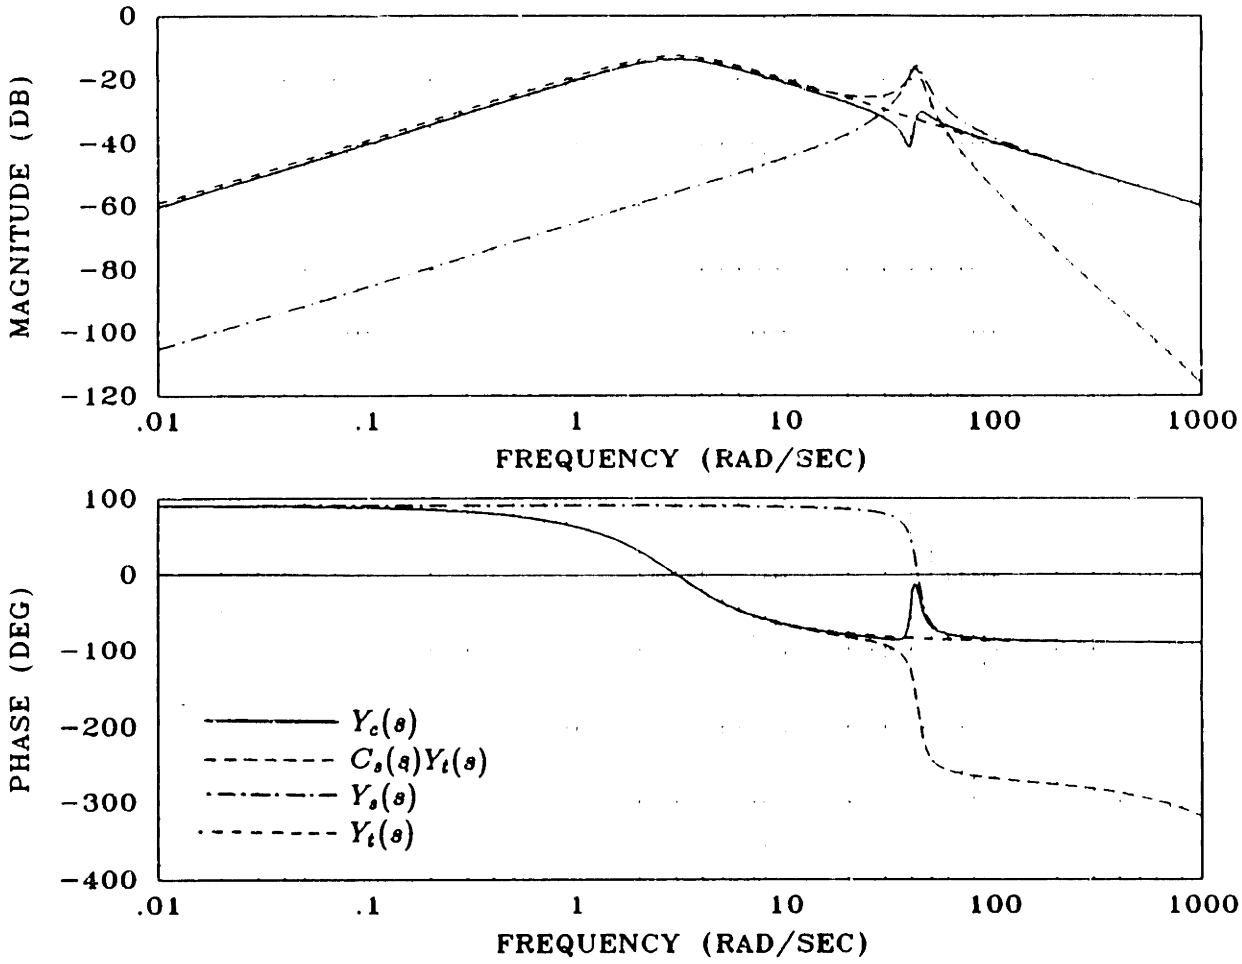

Figure 8.13: Bode plots of TMRC behavior, $\rho = 10^{-11}$.

tive real, $C_s(s)Y_t(s)$ *cannot* be. Because $C_s(s)$ is a transfer function, not an impedance, the phase lag of $C_s(s)Y_t(s)$ must be greater than 90° at high enough frequency. The consequence is, intuitively speaking, that at high enough frequency this term instructs the system to move in the 'wrong' direction in response to environmental disturbances. Furthermore, if the magnitude of this term is large enough relative to that of $Y_s(s)$, it is possible that $Y_c(s)$, which is the complex sum of the two terms, will also have a large phase lag at high frequency, which requires nonminimum phase zeros as described above.

The second case (Figures 8.13 and 8.14(b)), which is the same except for a control

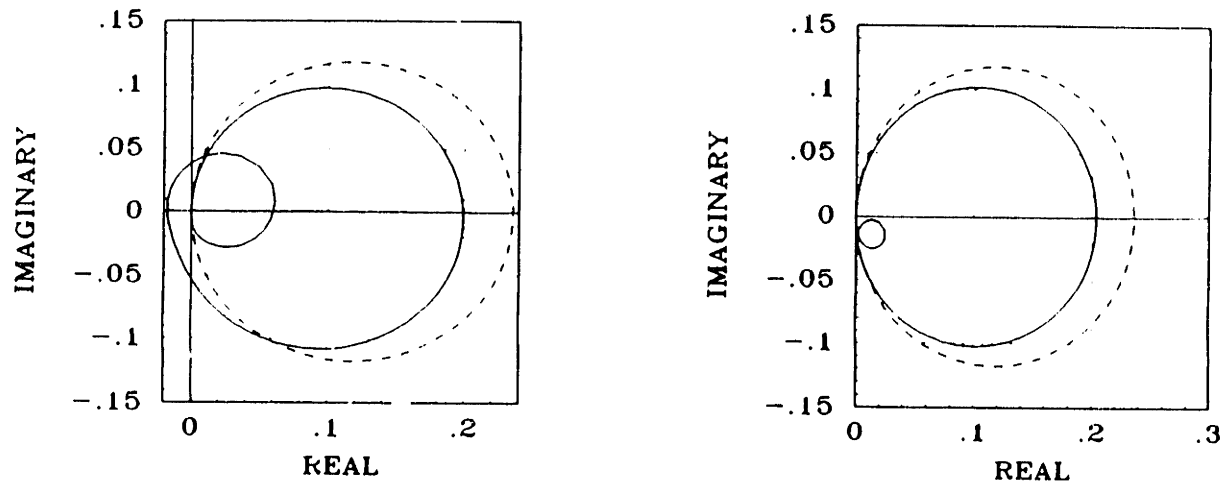

(a) (b)

Figure 8.14: Nyquist plots of target (dashed) and closed loop (solid) admittances. (a) $\rho = 10^{-6}$. (b) $\rho = 10^{-11}$.

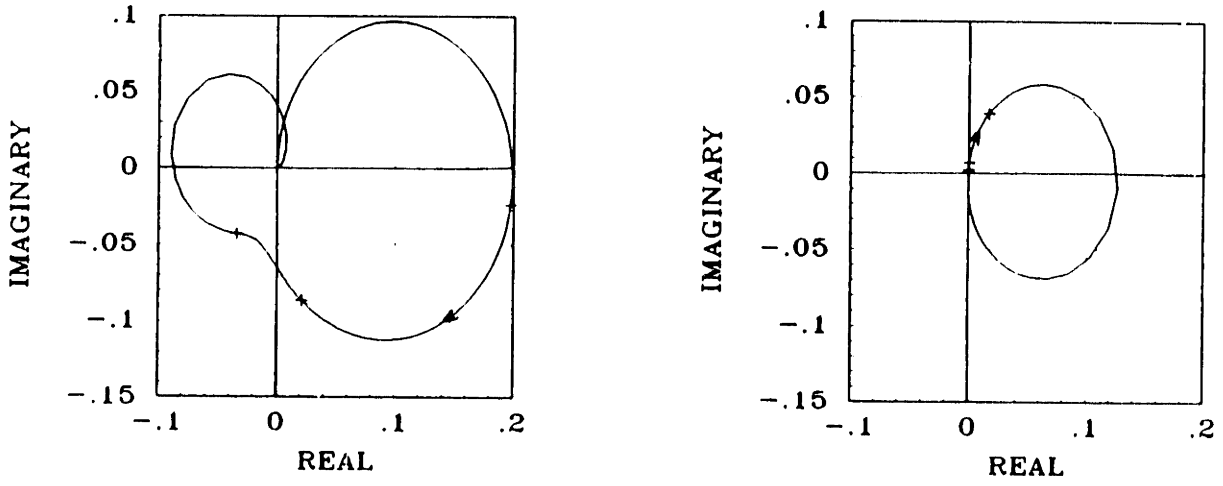

(a) (b)

Figure 8.15: (a) Nyquist plot of $C_s(s)Y_t(s)$, $\rho = 10^{-6}$. (b) Nyquist plot of $Y_s(s)$, $\rho = 10^{-6}$. These two Nyquist plots can be vectorially summed, pointwise in frequency, to produce $Y_c(s)$. Three different frequencies have been label on the plots. For the third, the sum is a vector which points into the left half plane.

weighting of $\rho = 10^{-11}$, demonstrates that nonminimum phase zeros are not, however, an essential by-product of the technique. The effect of this reduction in ρ is an increased (by ~ 5 rad/sec) bandwidth for the servo. The Bode plots of $C_s(s)Y_t(s)$ and $Y_s(s)$ are also more sharply peaked at the frequency $\omega = 42$ rad/sec, and their phase changes are more rapid. Apparently these changes conspire to form a positive real $Y_c(s)$. There are, however, several problems with this result:

- It is not systematic. Although a change in the parameter ρ reshapes both $C_s(s)Y_t(s)$ and $Y_s(s)$, it is not clear exactly how to predict whether the sum will be positive real or not, or even if ρ should be increased or decreased to improve matters. Nor is it clear whether another approach to servo design will help.

- The gains of the servo controller are unrealistic. For $\rho = 10^{-11}$, the state feedback gain matrix is $G = 10^6[0.3162\ 3.1622\ 0.0002\ 0.0025]$. This can be compared to the more successful servo masking design of the last section, for which $G = 10^4[-1.6\ 2.51\ 0\ 0.0176]$. Even if sufficient actuator effort were available, it is unlikely that such a high-gain controller could be stably implemented in the face of unmodeled dynamics.

- There is every reason to expect that computational delays and unmodeled dynamics will add more phase lag to $C_s(s)$, so that, in practice, even the design for $\rho = 10^{-11}$ might be expected to be non-positive real.

Although the outlook is not very promising, there are other approaches to the problem at hand. One of these is to introduce a feedforward compensator between the target model output and the servo input. The feedforward compensator could be used to change $C_s(s)$ without affecting $Y_s(s)$; thus, their bandwidths could be independently set. The problem now becomes one of effectively exploiting this capability. Reducing the bandwidth of $C_s(s)$ is certainly not desirable in terms of performance, and the extra

lag added by the feedforward compensator might only exacerbate the problem. But increasing the bandwidth of $C_s(s)$ is even worse—at high frequencies, when the phase of $C_s(s)Y_t(s)$ is $-180°$, the magnitude of $Re\{Y_s(j\omega)\}$ will certainly be too small to pull the complex sum into the right half plane. Numerical analysis bears out the folly of the feedforward compensator.

Another approach is to take advantage of the freedom offered by the target model. If the magnitude of $Y_t(s)$ is reduced, then the magnitude of $C_s(s)Y_t(s)$ is reduced and the sum $C_s(s)Y_t(s) + Y_s(s)$ is more likely to be positive real. Figures 8.16 and 8.17 show that this approach can indeed work. But these figures also indicate another interesting consequence of reducing $\|Y_t(s)\|$, that the low frequency condition of $\|C_s(s)Y_t(s)\| \gg \|Y_s(s)\|$ no longer holds. At low frequency, $Y_c(s)$ and $Y_s(s)$ have the same shape and differ in magnitude by about 10 dB. If this is the only effect, one might question whether the TMRC is even worth the trouble—a slight change in the servo design (e.g., lowering a position gain) might have just as profound a result. Furthermore, as the Nyquist plot indicates, the recovery of the target admittance is extremely poor.

We now return briefly to consider Kazerooni's stability criterion (equation 8.3), with the following notation changes: $G \rightarrow C_s$, $S \rightarrow Y_s$, and $H \rightarrow Y_t$:

$$\sigma_{max}(C_s Y_t E) \leq \sigma_{min}(Y_s E + I) \text{ for all } \omega \tag{8.18}$$

E now represents the impedance of the environment. Assuming now that E represents a passive environment, the following equivalent condition can be found:

$$\sigma_{max}(C_s Y_t) \leq \sigma_{min}|(Y_s E + I)E^{-1}| \tag{8.19}$$

For a sufficiently stiff environment (which is generally the case of interest in force control problems, and is also the example provided by Kazerooni), $\|E\|$ is large, and the condition 8.19 reduces to:

$$\sigma_{max}(C_s Y_t) \leq \sigma_{min}(Y_s) \tag{8.20}$$

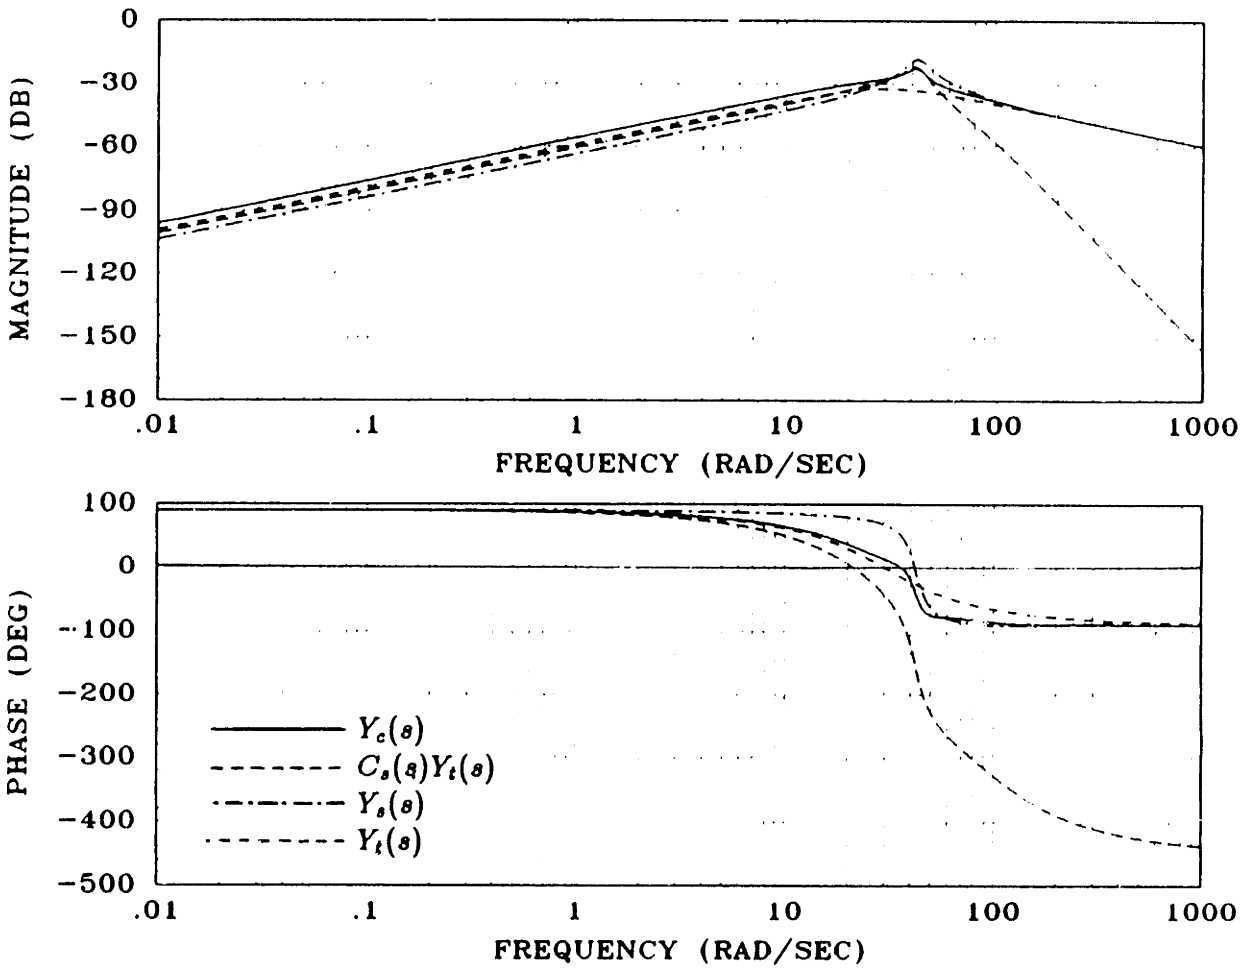

Figure 8.16: Bode plots of TMRC behavior, $\rho = 10^{-6}$, and target model characterized by $M_0 = 1$, $b_0 = 42.426$, and $k_0 = 900$. This target model has a reduced magnitude at low frequencies.

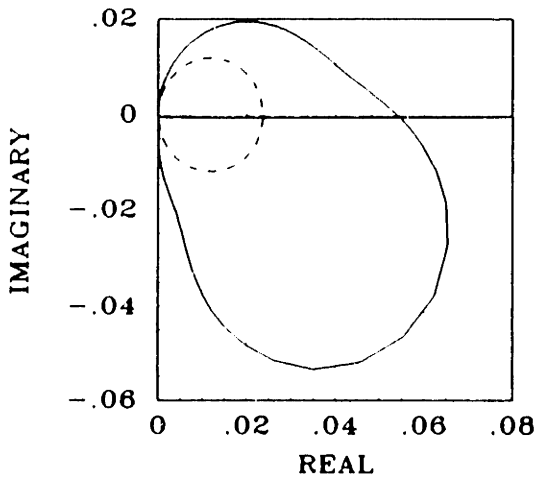

Figure 8.17: Nyquist plots of target (dashed) and closed loop (solid) admittances. Target model characterized by $M_0 = 1$, $b_0 = 42.426$, and $k_0 = 900$.

In the example above we showed what happens if $\|C_s Y_t\| \approx \|Y_s\|$. The result: coupled stability may be guaranteed, but the TMRC does not have a significant effect on the closed loop behavior. With Kazerooni's criterion, it will have even less effect.

In light of this shortcoming, we might ask what would happen if, instead of simply scaling $Y_t(s)$, its output was low-pass filtered. Would this not at least preserve the low frequency behavior of the TMRC? The answer is yes, but this is essentially the same as the earlier idea of adding a feedforward compensator to reduce the bandwidth of $C_s(s)$. The added phase lag virtually guarantees that $Y_c(s)$ will violate the coupled stability condition.

In summary, the TMRC is a critically flawed approach to implementing a desired impedance. The reader might wonder why so much effort was expended to describe a technique that has so little chance of working. The reason is that the TMRC behaves in essentially the same fashion as several approaches to compliant motion which are currently receiving much attention in the literature. Kazerooni's approach is one of these. In a recent paper, Lawrence [55] referred to this structure as a common approach to impedance control. Apparently it is considered common because it is a natural

approach given a position controlled robot, such as a PUMA. Unfortunately, the very concept of implementing a compliant controller on a robot designed to be a positioner (i.e., containing non-backdriveable reductions at the joints) appears to be ill-conceived. The approach described by Goldenberg (and to a lesser extent, Anderson and Spong; see Section 8.1), which involves inverting the robot impedance and then replacing it with some desired impedance, also behaves like the TMRC. Even the force feedback controllers described in the last chapter have a structure similar to the TMRC except for the abscence of a servo controller. Nevertheless, the inherent time delay plays the same debilitating role for the force feedback controller as the closed loop transfer function does for the TMRC. Also, the correlation between dynamic compensation and nonminimum phase zeros (Section 7.4) should now be clear.

We should be careful not to implicate all input feedback schemes in this assessment. The use of positive force feedback described by Hogan and Wlassich [40,92] is an example of input feedback which leads to a more robust interactive behavior than that of the open loop system (see the last chapter). The effect of the positive force feedback, however, is to *enhance* the inherent behavior of the manipulator rather than mask and replace it. The next section contains further discussion of the role of inherent dynamics. It should also be noted that even for positive force feedback, the magnitude of the gain must be less than unity (given a uniform manipulator); for larger gains, nonminimum phase zeros appear.

The point that we hope to have made by this lengthy analysis is that the concept of *replacing* the inherent driving point impedance with a target impedance through the action of state and input feedback is incompatible with robust coupled stability of the closed loop system. In the next chapter, alternative approaches will be suggested.

8.5 The Role of Inherent Dynamics

After the analyses of the last section, the following question naturally arises: why don't other types of impedance controllers, such as the simple impedance controller described in Section 4.1.1, also have difficulty with nonminimum phase zeros, or at least some violation of the positive real condition? After all, these controllers are also subject to computational delays and actuator dynamics which should, at some frequency, add sufficient lag to change the effective sign of the stiffness and damping terms.

To understand this issue, consider the example of a rigid body manipulator like that shown in Figure 7.1(a) with the control law $u(t) = -K_f x(t-T)$. It is straightforward to derive the following result:

$$Re\{Y_c(j\omega)\} \geq 0 \quad \Rightarrow \quad -K_f \frac{\sin(\omega T)}{\omega} \geq 0$$

The coupled stability condition is clearly violated for certain values of ω, including very small values. But suppose the plant included some inherent viscous damping (b) connecting the mass to ground. Then the following condition could be derived:

$$b - K_f \frac{\sin(\omega T)}{\omega} \geq 0$$

Now for $b \geq K_f T$, the closed loop system is positive real.

This simple example points out the importance of inherent dynamics, especially of inherent dissipation. But the existence of some amount of joint friction is not enough to explain the observed robustness to interaction of simple impedance controllers. This is illustrated by another example. Suppose that the damped manipulator is controlled with the following feedback law: $u(t) = -K_f x(t-T) - B_f \dot{x}(t-T)$. Then the coupled stability condition is the following:

$$b + B_f \cos(\omega T) - K_f \frac{\sin(\omega T)}{\omega} \geq 0$$

The following two conditions are now sufficient to make $Y_c(s)$ positive real: $b + B_f \geq K_f T$, and $b \geq B_f$. The first condition poses no difficulty, particularly for small T. However, the second condition, which comes into play when $\cos(\omega T) \approx -1$, is often violated in practice. Because discrete time implementation requires that $1/T$ be large compared to the servo bandwidth ($\sqrt{K_f/M}$), the violation will occur at frequencies well above the closed loop bandwidth.

How, then, do we explain the robustness? The solution lies in the *non-colocation* of actuator and interaction ports. To illustrate this point, consider the two-mass non-colocated robot model shown in Figure 7.2. Suppose that the control law $u(t) = -K_f x(t-T) - B_f \dot{x}(t-T)$ is implemented. The closed loop admittance is now a rather complicated expression, but at high frequencies the positive real condition has the following simple form:

$$\lim_{\omega \to \infty} Re\{Y_c(j\omega)\} = \frac{4B}{M\omega^2}$$

where $M/2$ is the endpoint mass and B is the structural damping. Apparently neither of the feedback gains affect the impedance at high frequency, a result which is consistent with the derivation in Appendix D. Although it is difficult to prove positive realness without selecting parameter values, this result suggests that—due to non-colocation—at those frequencies where the delay might be expected to have detrimental effects, the environment sees almost exclusively the inherent dynamics of the part of the manipulator close to the interaction port.

Consider again the question posed at the beginning of this section, which was why the coupled stability properties of the TMRC and simple impedance controller differ so. The following two explanations are offered:

- In the TMRC, the problem is tied to $C_s(s)$, and therefore to the servo bandwidth. At those frequencies that the term $C_s(s)Y_t(s)$ is driving the system in the wrong direction, its magnitude may still be large compared to that of $Y_s(s)$. In the

simple impedance controller, the problem is tied to computational delays and actuator dynamics which generally become important at 5–10 times the closed loop bandwidth. At these frequencies, the effect of the controller is already quite small compared to that of the inherent dynamics.

- In the TMRC, the servo is designed to make $\|Y_s(s)\|$ small so that it can be replaced. In the simple impedance controller, the feedback is used to augment $Y(s)$, and $\|Y_s(s)\|$ is not intended to be small.

The heart of the matter is that the TMRC and similar methods attempt to *suppress* the inherent dynamics, while the simple impedance controller and similar methods *exploit* the inherent dynamics. This is yet another indication that, insofar as interaction is concerned, mechanical design and controller design should be made to complement one another.

8.6 Summary

The major results of this chapter may be summarized as follows:

- It is possible to explicitly enforce the positive real condition by appropriately restricting the gains of any closed loop design. This, however, does not constitute an effective procedure for implementing a desired interactive behavior.

- Servo masking is an approach to implementing desired interactive behavior and command following behavior. This method, based upon the use of a servo to implement a desired constraint relation, proved very successful in the example considered. There is, however, no formal method for predicting the success of servo masking in general.

- The TMRC, which is very similar to a number of impedance control and force control schemes that have been proposed elsewhere, tends to have limited success because it attempts to mask the manipulator's natural response and implement interactive behavior through a term $(C_e(s)Y_t(s))$ which *cannot* be positive real.

- Inherent dynamics play an important role in allowing certain types of interaction controllers to meet the coupled stability criterion. Because these inherent dynamics do not need to be modeled, this suggests an interesting perspective on unmodeled dynamics as friend rather than foe.

Chapter 9

Conclusions and Discussion

9.1 Conclusions

The major contributions of this thesis may be summarized as follows:

- The derivation of a necessary and sufficient condition for the stability of a feedback controlled system coupled to an arbitrary passive environment.
- The derivation of an alternative criterion for the stability of a 1-port system coupled to an arbitrary passive environment based on the generation of two root loci—one for coupling to springs, and one for coupling to masses.
- An analytical and experimental confirmation that the driving point impedance is an effective predictor of interactive behavior. Results include:
  - The implementation of a simple impedance controller which satisfies the coupled stability condition.
  - A demonstration that integral control of joint angles leads to instability upon coupling to sufficiently large masses.
  - A demonstration that control schemes which do not employ force feedback can nevertheless exhibit contact instability.

- An experimental verification that coupled instability is sufficient to predict contact instability.

- The introduction of passive physical equivalents, which allow the interactive behavior of a closed loop, LTI 1-port to be examined with the same graphical techniques that are used to model the plant.

- The introduction of uncontrollable elements, which demonstrate in a quantitative and graphical manner fundamental limits on the performance of any interaction controller.

- An analysis of force feedback contact instability in terms of passive physical equivalents and uncontrollable elements. The analysis indicates that there is a fundamental limit to the effectiveness of *any* force feedback controller, and suggests a novel approach to improving force control—increasing the inertia close to the actuator to allow increased gain.

- The introduction of "servo-masking" and "target model referenced controller" design strategies, as well as a thorough analysis of the latter. The analysis demonstrates that the TMRC, as well as several similar techniques proposed by other authors, are ineffective because, in order to have a significant effect on the interactive behavior, they are likely to produce nonminimum phase zeros in the driving point impedance.

The coupled stability criterion formed the centerpiece of this thesis. In most examples it was treated as an absolute criterion, but in a few cases it was knowingly violated because the class of passive environments could be restricted (see, for example, Section 7.4). This suggests two interpretations of the criterion: one, as a robustness condition which should always be satisfied to guarantee coupled stability, and two, as an analytical tool for identifying those environments which will endanger coupled stability.

Which of these is preferable will depend on the application, but it should be evident that, when interaction is involved, there should always be some level of concern for the driving point impedance.

One of the goals of this thesis was to demonstrate the distinction between servo control and interaction control. It should now be evident on the basis of several examples that this distinction exists and that it should influence the way that manipulator controllers are designed. For instance, the three PID controllers described in Chapters 4 and 5 exhibit identical servo response, but substantially different interactive behavior. The LQG/LTR controller exhibits good command following, but poor coupled stability properties. Traditional servo approaches to force control (large gain proportional control or integral control) also cause contact instability.

It is further evident that interaction controllers must be designed as such from the outset. The band-aid approach of closing an input feedback loop around a servo controller, as embodied in the TMRC, is prone to failure. The servo-masking technique is more successful for the examples considered because it does not attempt to mask and replace the entire impedance, it simply attempts to mask those portions contributed by the actuator and transmission dynamics (which disappear at high frequency anyway).

Servo control and interaction control, however, are by no means mutually exclusive. A high bandwidth servo, for instance, can satisfy the coupled stability criterion, or an interaction controller can exhibit good servo properties (consider the simple impedance controller with very large K). Furthermore, a feedforward compensator can be used to tune servo behavior without affecting interactive behavior, as described in the introduction to Chapter 8.

Traditional approaches to servo control and many approaches to interaction control share the property that they readily destroy the behavior which they attempt to implement. Feedback controllers, for instance, perform best with large gains, but are virtually always unstable for large enough gains. Similarly, the use of input feedback,

which has become prevalent only since the advent of interaction problems, is a primary contributor to coupled instability. This paradoxical behavior, which is severe, is not necessarily unavoidable. High-bandwidth, stable servos can certainly be designed, but no successful, general approach to interaction control has yet been identified. One potential approach is described in Section 9.3.

9.2 Related Areas of Research

This thesis has focused on the interaction of a robot manipulator with a passive environment; however, there are a number of other potential applications of the concepts that have been developed. Several of these are also in the area of manipulation.

For instance, a growing number of researchers are beginning to investigate the problem of two or more robots cooperating on some task [2,31,81]. Much of the interest has been with regard to kinematic issues, although Kazerooni has performed a stability analysis of two cooperating robots using the Small Gain Theorem [46]. His results guarantee stability in the following three cases:

- One robot behaves as an effort source, the other as a flow source (no dynamics).
- One robot behaves as an effort source, the other as a bandwidth-limited position servo with admittance causality.
- One robot behaves as a flow source, the other as a bandwidth-limited force servo with impedance causality.

One weakness of this result is that stability is guaranteed only if at least one robot is free of dynamic effects; another is that the two robots cannot have similar behavior. An obvious alternative which addresses both of these problems is to ensure that both robots have positive real driving point impedances. In fact, an arbitrary number of

robots, all of which may be identical, may be stably coupled so long as they exhibit passive port behaviors.

Another natural application of these concepts is to the design of powered upper-extremity prostheses. Because upper-extremity prostheses, as partial replacements of human arms, are intended to interact with extremely diverse sets of dynamic environments, the issue of coupled stability is of crucial importance. However, coupled stability problems have not yet surfaced in the field of prosthetics. There are doubtlessly a number of factors which contribute to this apparent robustness, possibly including the soft mechanical coupling of the device to the wearer as well as the high-level supervision that the wearer provides. Another factor, however, is that conventional prostheses employ designs which would not be expected to exhibit coupled instability.

Cable-operated prostheses *are* passive devices, and state-of-the-art cybernetic prostheses such as the Boston Arm [57] and the Utah Arm [43] employ proportional velocity controllers, which are typically robust to interaction. This is because the impedance of a velocity controlled system is much like that of a viscous damper. Furthermore, these devices use large transmission ratios (the Utah Arm has a 323:1 reduction) which act to amplify inherent energy dissipation.

Of course, these devices do not restore full functionality to the lost limb, nor does such a device appear in the foreseeable future [57]. There are, however, several paths toward improved rehabilitation with a prosthetic limb. One that has been proposed is to provide the amputee the ability to tune the impedance of the limb in much the same way that intact individuals can via the coactivation of antagonist muscles [1,41].

To investigate the potential of such an approach, Abul-Haj [1] and Mansfield [58] have developed "prosthesis emulators". These laboratory devices have inherently low impedances which are meant to be augmented with the use of feedback controllers. Therefore, although these devices may provide the capability of more natural control, they also increase the risk of coupled stability problems. To date they have been used

only in constrained tasks, e.g. crank turning with the prosthesis securely attached to the crank, and have not exhibited coupled stability problems. One of the challenges of transferring the results of this research to a commercial device will be ensuring the coupled stability property.

The human arm is itself a manipulator, one which has provided much inspiration for this research. Voluntary control, reflexes, and inherent dynamics of the musculoskeletal system all contribute to the behavior of the arm, although the recipe for control is, as yet, poorly understood. The properties of the driving point impedance provide one important means of improving this understanding.

For instance, Hogan, et al. [42] showed that when the arm is held in a horizontal plane, measurements of the 2×2 stiffness tensor at the hand show very nearly zero curl component. This is, in fact, evidence that in steady state the arm behaves as a passive system composed of springs. Furthermore, the authors show that non-zero curl can be due only to intersegmental feedback (e.g., the motion of muscles acting about the elbow joint affecting the activation of muscles acting about the shoulder joint; see Section 2.2). The measurement of zero curl is a surprising result because it means that the intersegmental feedback paths, which do occur physiologically, must either be inactive, or their effects must be carefully balanced to avoid curl.

In either case, this is one data point which suggests that the central nervous system ensures that the closed loop impedance of the arm is that of a passive system. Of course, in order to verify this conjecture, measurements of the endpoint impedance at frequencies other than d.c. are necessary. Possible approaches are direct measurement of the endpoint impedance with a device such as the one described in [14], or an inferred measure of passivity through interaction with springs and masses. Either approach must carefully consider the ability of the human to adapt to the environmental disturbance. The role of adaptation is considered further in the next section.

The design and implementation of "simulated environments" is another potential

application area for these concepts. One approach to creating a simulated environment is the combined use of impedance-controllable hardware to provide a haptic image, and computer graphics to provide a visual image of the environment. Potential applications are the design of teaching aids for arthroscopic surgery, and the design of robot submarines which will explore the ocean floor while providing a remote operator a simulated copy. The task of the hardware in such devices would be to generate the correct passive (or possibly active) physical equivalent.

The concept of coupled stability, however, need not be restricted to manipulators. For instance, Wyatt and Standley [94] have applied a very similar analysis to the design of VLSI chips which are meant to emulate the behavior of lateral inhibition neural networks. Lateral inhibition neural networks are found in biological systems, such as the vertebrate retina. In such networks, the output of each neuron inhibits the response of its neighbors. In the retina, this type of connectivity acts as an automatic gain control and as a means for enhancing intensity edges.

In one VLSI implementation of a lateral inhibition neural network, each neuron is represented by a photoreceptor and two cascaded MOS transconductance amplifiers. Each of these elements is connected to its neighbors through a resistive array. The network may be extremely large, effectively precluding dynamic analysis. Yet, it is evident that the stability of such a network is in danger if the output impedances of the active "neurons" are inappropriately shaped in the Nyquist plane. The approach taken by Wyatt and Standley is essentially to ensure that the output impedance of each neuron is positive real. However, they have gone further to show that in cases where the network contains no inductance, a less conservative condition, termed "θ-positive" can be applied. They further show that certain classes of nonlinear resistors and capacitors can be included in the analysis if the linear portions of the circuit satisfy a Popov criterion.

It has also been suggested that related concepts may be applied to the design of

axial compressors [72]. An axial compressor, which may be thought of as an effort source coupled to a nonlinear resistance, typically exhibits a non-monotonic relation between velocity input and pressure output [28]. Because this characteristic is non-monotonic, the compressor appears to be active for certain operating conditions and may, depending on the dynamics of the system to which it is connected, cause severe limit cycles. For this to be the case, the nonlinear resistance must be negative at certain operating points. The suggestion is that feedback may be used to reshape the characteristic to a monotonic form, thereby forcing the nonlinear resistance to always be positive and the compressor to always be dissipative. Global stability would then be guaranteed so long as the remainder of the system was composed of passive components (as is typically the case).

9.3 Future Directions

In this section, a number of suggestions for future research are discussed. These suggestions are divided into those which address the general theory of dynamically interacting systems, and those which address application to manipulators in particular.

General Theory

One of the more obvious needs in the future is suggested by the final example of the previous section—extension of the coupled stability criterion to nonlinear plants. It is straightforward to prove that the passivity of a nonlinear system is *sufficient* to guarantee stability when coupled to an arbitrary nonlinear, passive environment, provided that all internal energy functions also serve as Lyapunov functions, which is usually the case. Fasse [22] has provided such a proof. A proof of necessity, however, is expected to be more difficult.

Even if necessary and sufficient conditions can be established, finding an internal energy function is bound to be just as difficult as finding a Lyapunov function. Several alternative approaches to this situation come to mind. One is to use feedback control to construct an internal energy function in much the same fashion as sliding mode techniques [73] are used to construct Lyapunov functions. Another approach, which would use the Small Gain Theorem [86] to establish input/output stability, is suggested below in the discussion of scattering variable methods.

A third approach would start with a structured physical model rather than a set of state equations. In this approach, the nonlinear elements of the model would be either linearized or simply replaced with linear elements. A linear feedback loop would be closed and a passive physical equivalent found. Based on the structure and the element values of the equivalent, an educated guess of the passive realization corresponding to the nonlinear system with linear feedback would be made. The behavior of the nonlinear candidate realization would be compared to that of the closed loop system, and iteratively updated as necessary until the behaviors matched. The result would be a nonlinear passive physical equivalent. Of course, there is no guarantee that the structure of the linear passive physical equivalent would be that of the nonlinear passive physical equivalent, or that a suitable method of iteration could be found.

Linear systems also provide a number of interesting research topics, particularly with regard to the design problem discussed in Chapter 8. One approach to the design problem which seems logical in light of the current feverish activity in linear control theory is to recast the interaction control problem in a form suitable for formal control system synthesis [24,75]. The standard block diagram describing the generalized control system is shown in Figure 9.1. A block diagram depicting a generalized interaction controller is shown in Figure 9.2. The interaction controller has provision for both output and interaction port input feedback.

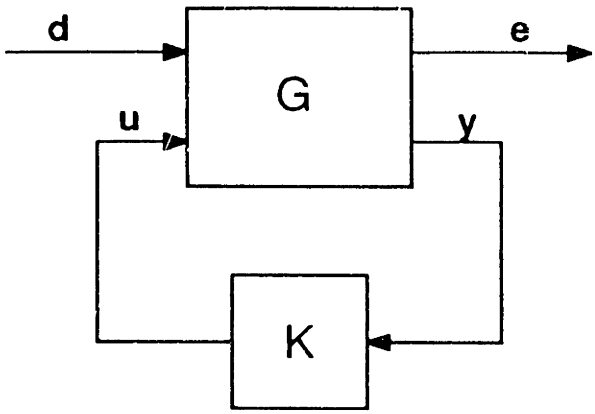

$$\begin{bmatrix} \mathbf{e} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix} \begin{bmatrix} \mathbf{d} \\ \mathbf{u} \end{bmatrix}$$
$$\mathbf{u} = K\mathbf{y}$$

Figure 9.1: The "standard block diagram" for formal control system synthesis, adapted from [24]. **d** is a vector of exogenous inputs, **u** is the control signal, **e** is an error output, and **y** is the measured output.

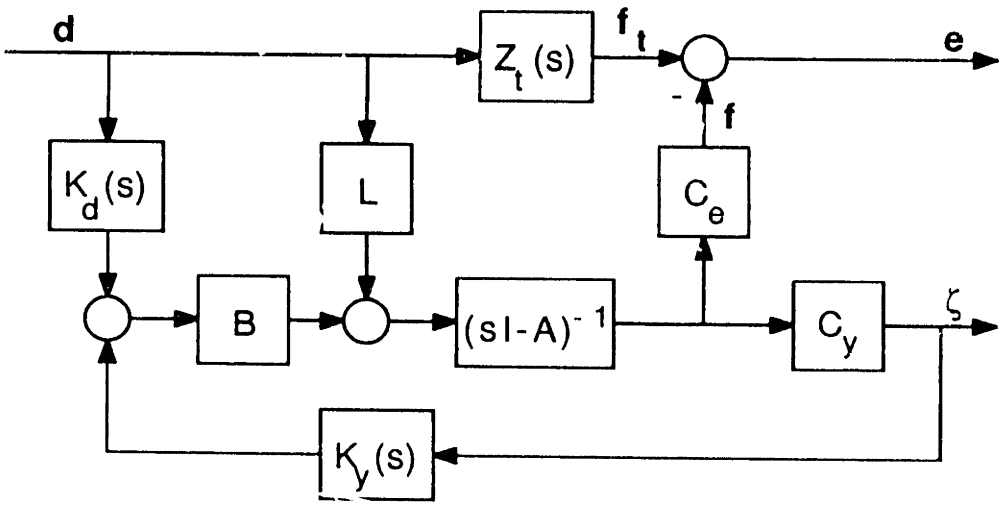

$$Z(s) = C_c(sI - A)^{-1}L$$
$$G_c(s) = C_c(sI - A)^{-1}B$$
$$G_y(s) = C_y(sI - A)^{-1}B$$
$$S(s) = C_y(sI - A)^{-1}L$$

Figure 9.2: Block diagram of a generalized interaction controller. Note that the target impedance, $Z_t(s)$, does not affect either the closed loop stability or the coupled stability.

The generalized interaction controller may be expressed in the standard form by making the following definitions:

$$\begin{aligned} \mathbf{y} &= \begin{bmatrix} \mathbf{d} \\ \varsigma \end{bmatrix} \\ G_{11} &= -G_e \\ G_{12} &= Z_t - Z \\ G_{21} &= \begin{bmatrix} I \\ \mathcal{S} \end{bmatrix} \\ G_{22} &= \begin{bmatrix} 0 \\ G_v \end{bmatrix} \\ K &= \begin{bmatrix} K_d & K_v \end{bmatrix} \end{aligned}$$

Now, by restricting in some useful manner the class of inputs $\mathbf{d}$ and the class of allowable errors $\mathbf{e}$, a control system synthesis procedure can be performed. For instance, if both $\mathbf{d}$ and $\mathbf{e}$ are required to be L_2, then an H_∞ optimization can be used to design the controller which will minimize the L_2 norm of $\mathbf{e}$ given worst-case inputs, $\mathbf{d}$ [75]. Whatever the procedure selected, we would like to add the constraint that the closed loop driving point impedance be positive real. The closed loop impedance, $Z_c(s)$, is given by:

$$Z_c(s) = G_e(I - K_v G_v)^{-1}(K_d + K_v S) + Z$$

The positive real constraint poses a difficulty in that it is essentially a phase condition, whereas all of the techniques of formal control system synthesis operate on suitably chosen norms. A way of converting the phase condition to a magnitude condition is discussed below.

There are, however, other potential difficulties with this approach. For instance, the TMRC described in Chapter 8 fits the general form in Figure 9.2, but is demonstrably ineffective as an interaction controller. This raises the suspicion that the method cannot, in general, produce a small error without sacrificing coupled stability. Another problem is that the method requires a bound on the interaction port input. This was

not required by any of the examples in this thesis, nor is it clear how such a bound should be selected in general, particularly if the environment is a passive system. A final concern is that the method does nothing to help understand the effect of plant design, which we know plays a critical role in the interactive behavior of closed loop systems.

In summary, this approach has the promise of making powerful synthesis techniques available, but it has several potential liabilities as well. At this stage, further investigation appears to be justified.

As mentioned above, it is possible to convert the positive real criterion to a magnitude condition. This may be accomplished by transforming from an impedance description to a scattering description [4]:

$$S(s) = [Z(s) + I]^{-1}[Z(s) - I]$$

The scattering matrix, $S(s)$, is a mapping from power flow into the interaction ports to power flow out of the ports. It is therefore fairly evident that passivity can be defined as follows [4].

Passivity 9 (Scattering Variable). A LTI n-port is passive iff $S(s)$ has no poles in the closed RHP, and:

$$\sigma_{max}\{S(j\omega)\} \leq 1 \quad \text{for all } \omega$$

In other words, passivity requires that the power out be less than the power in. Passivity has now been stated in a form which is identical to that of standard robustness specifications [56]. Another advantage of this form is that imaginary poles do not need special consideration—they are not allowed.

The use of scattering variables has other potential benefits, as well. For instance, scattering variables allow a symmetric representation of the coupling of two systems, because each system takes power as an input and produces power as an output. This is

a convenient form for modeling the interaction of identical systems, such as two robots. Another potential benefit is that scattering variables apply to nonlinear systems. This suggests that the Small Gain Theorem can be used to identify a class of nonlinear passive systems, and even to prove coupled stability in the nonlinear case.

A drawback to the scattering variable formalism is that the concepts of power in and power out are fictitious for lumped parameter models. Scattering variables arise naturally in the analysis of distributed systems where left-traveling and right-traveling waves carry the power. Impedances and admittances, which relate the measurable port variables force and velocity, are more natural descriptions of lumped parameter systems. In any case, the scattering variable approach is a realistic reformulation of the concepts presented in this thesis.

Another area of research is understanding the roles that real-time adaptation and learning[1] might play. Although it can be argued that the coupled stability criterion must be satisfied if stable interaction with arbitrary passive environments is to be guaranteed, this does not rule out adaptation and subsequent violation of the criterion *once coupled* to a particular environment. Whether or not they result in non-positive real behavior, real-time adaptation and learning are potentially valuable features of an interaction controller.

We make note of the fact that the human arm can stably interact with a tremendously diverse set of environments, suggesting that the arm has the port behavior of a passive system; yet, the arm is also able to execute highly skilled and learned movements, such as swinging a golf club or a tennis racket. In other words, robust interaction and high performance are by no means mutually exclusive goals, but adaptation and learning may play important roles.

[1] Real-time adaptation refers to the adjustment of compensator parameters on-line, whereas learning is an adjustment—generally of a feedforward signal—which occurs off-line between subsequent trials of some repetitive task.

Finally, the discussions of control system synthesis and adaptation may be coalesced to suggest a new perspective on interaction. Previous approaches have focused on achieving a desired servo behavior while ignoring robust coupled stability (e.g., force control), or on implementing a desired impedance while attempting to redefine performance measures in terms of a target impedance or target model (e.g, impedance control). Although we have attempted to demonstrate in this thesis that servo control and interaction control are distinct, we have also shown that they are not mutually exclusive. The new approach, therefore, would be to achieve both; for instance, to focus on servoing some output, but with a positive real constraint on the impedance, at least before coupling. Adaptation and learning after coupling would be used to improve the servo performance without jeopardizing stability.

Application to Manipulation

It is evident that any of the ideas discussed above is applicable to manipulation in some form. The object here is to discuss potential research activities that are particular to manipulation. For instance, it was suggested in Chapter 7 that certain robot design strategies should improve force control. Given the current prominence of force control research in the field of robotics, these design strategies seem worthy of implementation and investigation.

The coupled stability concept also provides an additional experimental approach to understanding human manipulation, as described in the previous section. For instance, understanding the constraints imposed by stable interaction should help distinguish the various roles of voluntary control, reflexes, and inherent muscle and limb dynamics.

In particular, it seems likely that the success of the human arm is due not only to the immense computational capabilities of the brain, including the abilities to learn and adapt, but also to the intelligent use of feedback and, perhaps most importantly, to the excellent design of the "hardware"—muscles, joints, and limbs. The lesson that

seems appropriate for those in the field of robotics is not to emulate the design and control of the human arm, but to understand that a successful approach to robotics will require the same sort of synergistic use of design, control, and computation that evolution has provided in the physiological system.

For instance, it was shown in Chapter 8 that the use of feedback to reduce the impedance of a manipulator which is designed to be non-backdriveable can result in disaster. On the other hand, the use of feedback to add stiffness and damping to a backdriveable manipulator produces a much more robust result (in part due to the presence of higher-order dynamics). This approach, however, is energetically expensive, especially when one considers that the end result has the behavior of a passive system. The real value of using feedback rather than actual masses, springs, and dampers, is twofold. First, feedback can be used to mask undesirable effects, such as friction and other nonlinearities; second, feedback allows the impedance to be tuned, simply by changing gains. The human arm, however, addresses these issues through design as well as control. Important design features include very low friction joints and actuators (muscles) whose impedance can be tuned open loop. The use of tuneable impedance actuators is energy efficient[2], alleviates the compuational burden on the controller, and is extremely robust to interaction. It is felt that one of the most important research activities for the future will be the construction of robotic manipulators which fully exploit the capabilities of both inherent dynamics and feedback control.

[2]Because muscles require metabolic energy to stiffen, there are instances in which the use of feedback is an important mechanism for reducing energy expenditure. The maintenance of upright forearm posture is an example [60].

Appendix A

Counterclockwise Loops in Nyquist Plots

In this appendix, it is shown that a Nyquist plot which contains a counterclockwise loop (such as those in Figures 3.4(d) and (e)) represents a transfer function with at least one RHP pole. Transfer functions of the following form will be considered:

$$A(s) = \frac{\alpha_0 s^m + \alpha_1 s^{m-1} + \cdots + \alpha_m}{s^n + \beta_1 s^{n-1} + \cdots + \beta_n}$$

where $m \leq n$. N will refer to the number of *clockwise* encirclements of the origin by the Nyquist plot of $A(s)$, P to the number of RHP poles of $A(s)$, and Z to the number of RHP zeros of $A(s)$. There are two cases:

Case 1. This is the trivial case in which the counterclockwise loop encircles the origin; therefore the Principle of the Argument requires:

$$Z - P = N = -1$$

Clearly, as $Z \geq 0$, it is necessary that $P \geq 1$.

Case 2. This is the general case in which the counterclockwise loop does not encircle the origin; therefore:

$$Z - P = 0$$

In this case, we find some point $a+bj$ which is encircled by the loop, and consider a mapping of the Nyquist contour through $A(s) - (a + bj)$. This mapping has a shape identical to that of $A(s)$, but is shifted so that the loop now encircles the origin. Furthermore, it is evident that $A(s)$ and $A(s) - (a + bj)$ share the same poles:

$$A(s) - (a + bj) = \frac{(\theta_0 + \phi_0 j)s^n + (\theta_1 + \phi_1 j)s^{n-1} + \cdots + (\theta_n + \phi_n j)}{s^n + \beta_1 s^{n-1} + \cdots + \beta_n}$$

For the shifted transfer function, the Nyquist plot reveals:

$$Z' - P = -1$$

where Z' is the number of RHP zeros of $A(s) - (a + bj)$. Thus:

$$P = Z' + 1 \geq 1$$

This completes the proof.

Appendix B

The LQG/LTR Design Procedure

This appendix provides a brief overview of the Linear Quadratic Gaussian / Loop Transfer Recovery design procedure, with emphasis on the design of a servo for the two-link manipulator (Section 4.4). The results presented in this appendix are based upon the description of LQG/LTR found in Athans [8].

The LQG/LTR procedure is a servo design technique which essentially provides a prescription for a compensator to invert a multi-input, multi-output plant. The compensator is model-based, as shown in Figure B.1. The expression for the compensator, $K(s)$, is:

$$K(s) = G[(sI - A)^{-1} + BG + HC]^{-1}H \tag{B.1}$$

In this development, we will not be concerned with the effects of modeling errors, so the matrices A, B, and C in the compensator are assumed to be identical to their counterparts in the plant. A discussion of the robustness of LQG/LTR compensators to modeling errors is provided in [56].

The "control gain matrix", G, and "filter gain matrix", H, in Figure B.1 are the choice of the designer, and their appropriate selection lies at the heart of the technique. The loop transfer recovery result determines the selection of G. This result is as follows: if H is chosen such that the eigenvalues of $A - HC$ lie in the LHP, and if a family of control gain matrices parametrized by ρ (termed G_ρ) can be found such that

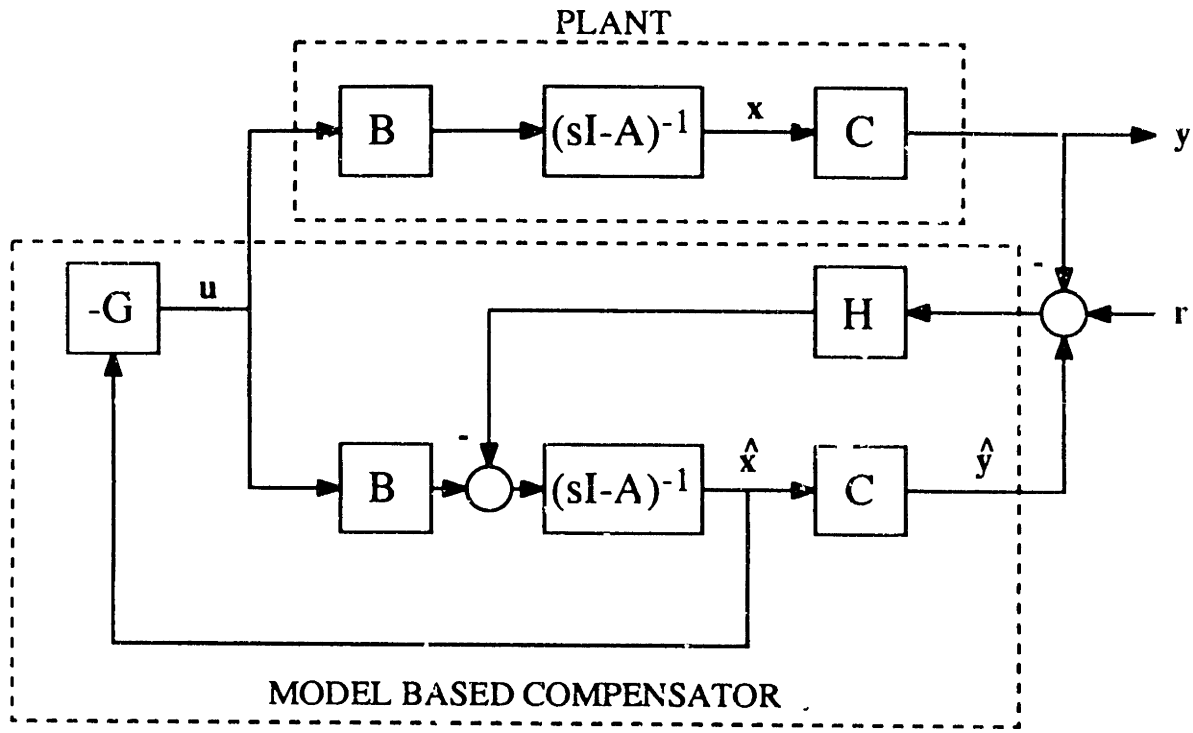

Figure B.1: Structure of the model-based compensator. The matrix G generates a control based upon the state estimate, $\hat{\mathbf{x}}$; the matrix H ensures that the error, $\mathbf{x} - \hat{\mathbf{x}}$, does not diverge.

the eigenvalues of $A - BG_\rho$ lie in the LHP for all ρ, and:

$$\lim_{\rho \to 0} \sqrt{\rho} G_\rho \to WC; \quad W'W = I \tag{B.2}$$

then, for each value of s, as $\rho \to 0$:

$$T(s) = C(sI - A)^{-1} B K_\rho(s) \to C(sI - A)^{-1} H \tag{B.3}$$

In other words, the behavior of the control system approaches that of the system shown in Figure B.2. The extent to which this "recovery" is effected is a function of ρ. Furthermore, such a G_ρ can be found by solving the Control Algebraic Riccati Equation (CARE) for minimum phase, stabilizable $C(sI - A)^{-1}B$:

$$0 = -\Phi_\rho A - A'\Phi_\rho - C'C + \frac{1}{\rho}\Phi_\rho B B' \Phi_\rho \tag{B.4}$$

$$G_\rho = \frac{1}{\rho} B' \Phi_\rho \tag{B.5}$$

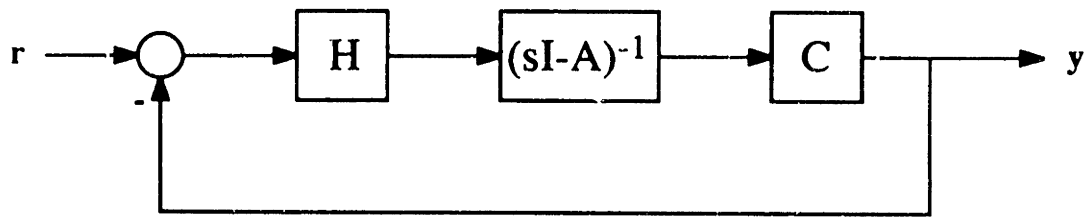

Figure B.2: The "target loop". This is the system which is "recovered" by the loop transfer recovery procedure. It therefore represents the targeted behavior.

It can be shown that G_ρ defined in this manner satisfies the conditon B.2 [8]. An alternative, but less frequently used method of selecting G_ρ is described by Kazerooni [47].

With the LTR result established, the problem becomes one of selecting H to generate a desirable shape for the loop transfer function $C(sI - A)^{-1}H$. One systematic method of selecting H is to use the Filter Algebraic Riccati Equation (FARE) in the following manner:

$$0 = A\Sigma + \Sigma A' + LL' - \frac{1}{\mu}\Sigma C'C\Sigma \tag{B.6}$$

$$H = \frac{1}{\mu}\Sigma C' \tag{B.7}$$

The matrix L and the scalar μ are now design parameters. L may be used to "shape" the singular values (as functions of frequency) of $C(sI - A)^{-1}H$ (see below), and μ may be chosen to set the bandwidth.

This completes our brief overview of the procedure, which will now be illustrated with the design of a servo for the two-link manipulator. A block diagram of the servo to be designed is shown in Figure B.3.

The manipulator dynamics are represented by the state matrices A and B (defined in equations 4.1 and 4.2) and the output matrix C, which simply selects the velocity

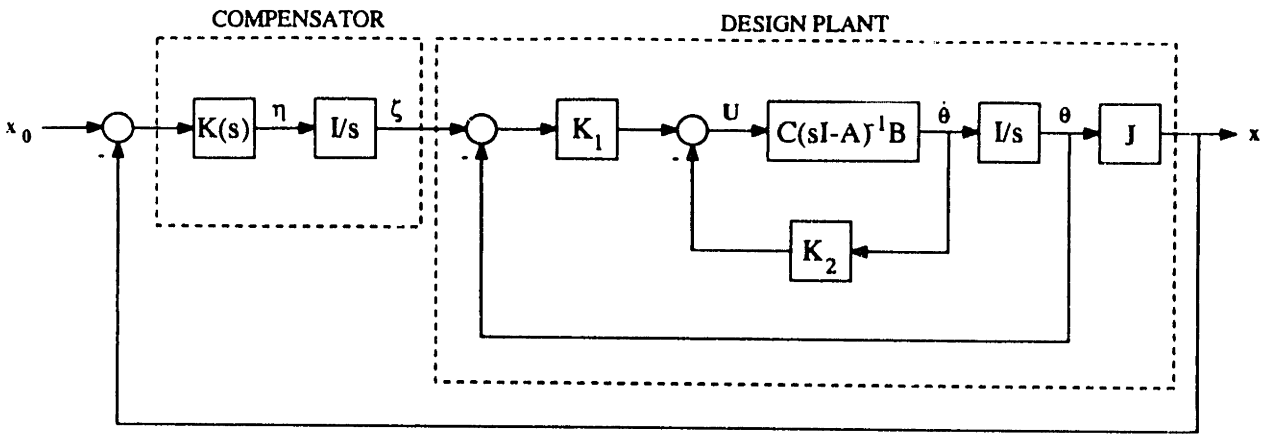

Figure B.3: Block diagram of LQG/LTR controller for the two-link manipulator. Note that this implementation contains integrators in the forward loop.

states, $\dot{\theta}$, as outputs. The position states, θ, can then be found by integrating the velocity, and the endpoint positions can be found from the joint positions through the jacobian matrix, J (recall that this model represents the manipulator behavior only for small excursions about some equilibrium point). The "inner loop" feedback represented by K_1 and K_2 is intended to condition the manipulator dynamics for the LQG/LTR procedure. All four eigenvalues of A lie very close to the origin of the s-plane, a condition which makes the standard approach to selecting H fail. To avoid this difficulty, the eigenvalues are moved into the LHP with K_1 and K_2. The "design plant" (including the inner loop feedback) has the following state and output equations:

$$\begin{bmatrix} \dot{\theta} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -I^{-1}K_1 & I^{-1}(\beta - K_2) \end{bmatrix} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} + \begin{bmatrix} 0 \\ I^{-1}K_1 \end{bmatrix} \zeta$$

$$\mathbf{x} = \begin{bmatrix} J & 0 \end{bmatrix} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix}$$

This implementation includes an integrator in each forward loop control channel, so that zero steady-state error is achieved. The design plant is augmented with these integrators, and it is the "augmented plant" which is used in the design of the LQG/LTR

compensator. The augmented plant has the following state and output equations:

$$\begin{bmatrix} \dot{\varsigma} \\ \dot{\theta} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ I^{-1}K_1 & -I^{-1}K_1 & I^{-1}(\beta - K_2) \end{bmatrix} \begin{bmatrix} \varsigma \\ \theta \\ \dot{\theta} \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \eta$$

$$\mathbf{x} = \begin{bmatrix} 0 & J & 0 \end{bmatrix} \begin{bmatrix} \varsigma \\ \theta \\ \dot{\theta} \end{bmatrix}$$

or, in a condensed form:

$$\dot{\xi} = A_a \xi + B_a \eta$$

$$\mathbf{x} = C_a \xi$$

The matrix H is chosen to shape the singular values of the target loop transfer function, $C_a(sI - A_a)^{-1}H$. It can be shown [8] that, if the matrix L in the FARE is chosen as follows:

$$L = \begin{bmatrix} -[C(A - B[\ K_1 \quad K_2\ ])^{-1}B]^{-1} \\ C'[CC']^{-1} \end{bmatrix}$$

the singular values of $C_a(sI - A_a)^{-1}H$ will be "matched" (all have the same magnitude) at low and high frequencies, and that at low and high frequencies, these singular values will drop off at 20 dB/decade as a function of frequency. The targeted loop transfer function, then, is simply a multivariable integrator. The magnitude of this integrator, and therefore the open loop crossover frequency and closed loop bandwidth, may be adjusted with the parameter μ.

Once the target loop has been designed, the matrix G is found by solving the CARE (equation B.5). The parameter ρ is adjusted (decreased) until suitable recovery of the target loop is obtained. The selection of both μ and ρ must be moderated by attention to the necessary control effort—excessively high bandwidth or close recovery will require excessively large control effort. Step responses of the two-link manipulator were simulated to ensure that the controls were within the capabilities of the hardware.

The complete controller consists of the LQG/LTR compensator, $K(s)$, plus the integrators. Together, these constitute an eighth-order compensator for a fourth order

plant. Such a high-order compensator calls for a digital implementation. Therefore, a discrete time compensator was derived from the contiuous time compensator based upon the first-order approximation, $\mathbf{x}_{k+1} = (1 + AT)\mathbf{x}_k + BT\mathbf{U}_k$, where T is the sample period, and k is the discrete time index. The eighth-order discrete time compensator, as implemented, is described by the following state and output equations:

$$\begin{bmatrix} \zeta_{k+1} \\ \mathbf{z}_{k+1} \end{bmatrix} = \begin{bmatrix} 1_{2\times 2} & -GT \\ 0_{6\times 2} & 1_{6\times 6} + (A_a - B_a G - HC_a)T \end{bmatrix} \begin{bmatrix} \zeta_k \\ \mathbf{z}_k \end{bmatrix} + \begin{bmatrix} 0_{2\times 2} & 0_{2\times 2} & 0_{2\times 2} \\ -HT & HJT & 0_{6\times 2} \end{bmatrix} \begin{bmatrix} \mathbf{r}_k \\ \theta_k \\ \dot{\theta}_k \end{bmatrix}$$

$$\mathbf{U}_k = \begin{bmatrix} K_1 & 0_{2\times 6} \end{bmatrix} \begin{bmatrix} \zeta_k \\ \mathbf{z}_k \end{bmatrix} + \begin{bmatrix} 0_{2\times 2} & -K_1 & -K_2 \end{bmatrix} \begin{bmatrix} \mathbf{r}_k \\ \theta_k \\ \dot{\theta}_k \end{bmatrix}$$

The performance of this controller is described in Sections 4.4 and 5.4.4.

Appendix C

An Example of Passive Physical Equivalent Synthesis

In this appendix, a passive physical equivalent is synthesized. The system of interest is shown in Figure 7.2; the control law is $u = -GF$. The closed loop driving point admittance is:

$$Y(s) = \frac{v(s)}{F(s)} = \frac{s^2 + (1+G)B\frac{2}{M}s + (1+G)K\frac{2}{M}}{\frac{M}{2}s^3 + 2Bs^2 + 2Ks}$$

It is straightforward to generate an equivalent with the Cauer synthesis techniques. The first step is a continued fraction expansion about infinity. The admittance is written in an anticausal form (i.e., as an impedance), and the pole at infinity is removed:

$$\begin{aligned} Z(s) = Y^{-1}(s) &= \frac{M}{2}s + \frac{(1-G)Bs^2 + (1-G)Ks}{s^2 + (1+G)B\frac{2}{M}s + (1+G)K\frac{2}{M}} \\ &= \frac{M}{2}s + Z_1(s) \end{aligned}$$

This result may be interpreted to mean that a mass, $M/2$, has a common velocity with the rest of the system, represented by Z_1, and that the sum of the forces on the mass and Z_1 is equal to the force imposed by the environment A bond graph representation of this result is shown in Figure C.1(a).

The next step is continued fraction expansion about zero. Z_1 is inverted, and the

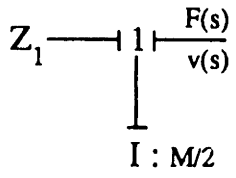

(a)

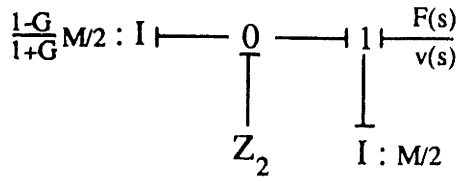

(b)

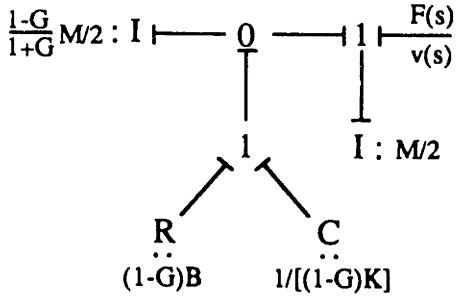

(c)

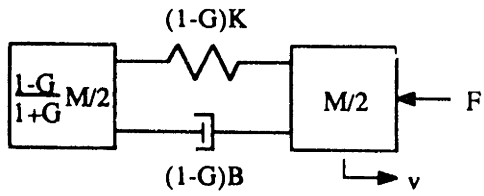

(d)

Figure C.1: Synthesis of a passive physical equivalent. (a) Step 1. (b) Step 2. (c) Bond graph model of complete realization. (d) Mechanical element model of complete realization.

pole at zero is removed:

$$\begin{aligned} Y_1(s) = Z_1^{-1}(s) &= \frac{1+G}{1-G}\frac{2}{Ms} + \frac{s}{(1-G)Bs + (1-G)K} \\ &= \frac{1+G}{1-G}\frac{2}{Ms} + Y_2(s) \end{aligned}$$

The bond graph model, updated to reflect this cycle, is shown in Figure C.1(b).

As a final step, Y_2 can be written in impedance form, and a pole at zero can be removed to reveal the contributions of a spring and a damper:

$$Z_2(s) = Y_2^{-1}(s) = (1-G)B + \frac{(1-G)K}{s}$$

A bond graph of the complete realization is shown in Figure C.1(c), and a mechanical element drawing is shown in Figure C.1(d).

Appendix D

Preservation of Relative Order

The purpose of this appendix is to show that state and input feedback do not change the relative order of a system's driving point impedance, or the behavior of this impedance at high frequency. The system under consideration is described as follows:

$$\begin{aligned} \dot{\mathbf{x}} &= A_{n\times n}\mathbf{x} + B\mathbf{u} + L\mathbf{e} \\ \mathbf{f} &= C\mathbf{x} \end{aligned}$$

where $\mathbf{u}$ is the control and $\mathbf{e}$ and $\mathbf{f}$ are power duals. It is also assumed that the control ports and interaction ports are non-colocated; the reason for this assumption will become clear in the sequel.

To begin, consider the zeros of the open loop ($\mathbf{u} = 0$) driving point admittance (or impedance). The zeros of the admittance function relating the i^{th} input (e_i) to the j^{th} output (f_j) are the roots of:

$$Y_{ij} = C_j(sI - A)^{-1}L_i = 0$$

where C_j is the j^{th} row of C, and L_i is the i^{th} column of L. Consider now $(sI - A)^{-1}$, which may be written in terms of the elements of A:

$$(sI - A)^{-1} = \begin{bmatrix} s - a_{11} & -a_{12} & \cdots & -a_{1n} \\ -a_{21} & s - a_{22} & & \\ \vdots & & & \\ -a_{n1} & & & s - a_{nn} \end{bmatrix}^{-1}$$

$$= \frac{1}{\det(sI-A)} \begin{bmatrix} M(s-a_{11}) & -M(-a_{21}) & \cdots & (-1)^{n+1}M(-a_{n1}) \\ -M(-a_{12}) & M(s-a_{22}) & & \\ \vdots & & & \\ (-1)^{n+1}M(-a_{1n}) & & & M(s-a_{nn}) \end{bmatrix}$$

where $M(\xi)$ is the determinant of the minor matrix found by eliminating from A the row and column in which ξ resides. If $\xi = s - a_{ii}$, the minor matrix has only one term containing an s eliminated (namely, $s - a_{ii}$), and M is a polynomial in s of order $n-1$. If $\xi = -a_{ij}$ $(i \neq j)$, two terms containing s are eliminated, and M is a polynomial in s of order $n-2$. Thus, we may write:

$$(sI-A)^{-1} = \frac{1}{\det^{(n)}} \begin{bmatrix} M_{11}^{(n-1)} & -M_{21}^{(n-2)} & \cdots & (-1)^{n+1}M_{n1}^{(n-2)} \\ -M_{12}^{(n-2)} & M_{22}^{(n-1)} & & \\ \vdots & & & \\ (-1)^{n+1}M_{1n}^{(n-2)} & & & M_{nn}^{(n-1)} \end{bmatrix}$$

where the superscripts in parentheses refer to the order of the polynomial. We also note that each of the terms $M_{ii}^{(n-1)}$ is of the form $s^{n-1} + m_{n-2}s^{n-2} + \cdots + m_0$, where the leading coefficient (m_{n-1}) is unity.

Our goal is to determine the number of zeros of $Y_{ij}(s)$ and, therefore, the relative order of $Y_{ij}(s)$. In particular, for reasons which will soon be evident, we wish to show that each admittance $Y_{ii}(s)$ contains $n-1$ zeros. To do this, it is necessary to show that $Y_{ii}(s)$ is of the form $y_{n-1}s^{n-1} + y_{n-2}s^{n-2} + \cdots + y_0$, where $y_{n-1} \neq 0$

To simplify the task of demonstrating that $y_{n-1} \neq 0$, we note that the off-diagonal terms of the matrix above all have order $n-2$, and cannot affect y_{n-1}. Thus, we may write:

$$\begin{aligned} \text{L.C.}\{\text{numerator } Y_{ij}(s)\} &= \text{L.C.}\left\{ C_j \begin{bmatrix} s^{n-1}+\cdots & 0 & \cdots & 0 \\ 0 & s^{n-1}+\cdots & & \\ \vdots & & & \\ 0 & & & s^{n-1}+\cdots \end{bmatrix} L_i \right\} \\ &= C_j L_i \end{aligned}$$

where L.C.$\{\cdot\}$ refers to the leading coefficient of the argument. It is now evident that the matrix CL is simply the matrix of leading coefficients of the numerators of $Y(s)$.

The state space passivity criterion (equation 2.38) provides some important information about this matrix:

$$PL = C'$$
$$CL = CP^{-1}C'$$

Because P is a PD matrix, it is clear that CL is a PD or PSD matrix. Furthermore, although C may be singular, no row of C will contain all zeros. It follows from an application of Rayleigh's principle [77] that all diagonal elements of CL will be positive. Thus, each diagonal element of $Y(s)$ will have n poles, $n-1$ zeros, and therefore a relative order of 1.

Now consider the zeros of the *transfer function* $G(s) = C(sI - A)^{-1}B$. If the control and interaction ports were colocated, we would expect to reach the same conclusions that we did for $Y(s)$. However, we have assumed that these ports are non-colocated, and the consequence is that the relative order of each element of $G(s)$ must be greater than 1. Consider, for instance, the example in Figure D.1. As indicated by the transfer function, the relative order is 2. The higher relative order makes intuitive sense, as standard sources of non-colocation, e.g. actuator rolloff or transmission dynamics, should add significant attenuation and lag at high frequencies. We might even expect infinite high frequency lag and infinite relative order in real situations. In conclusion, the matrix CB, representing the coefficients of the numerator s^{n-1} terms, should be identically zero:

$$CB = 0$$

We are now prepared to consider the effect of a feedback controller on the relative order of the elements of $Y(s)$. Figure D.2 is a block diagram of the controller. $K(s)$ is a matrix of proper transfer functions which represents the combined effects of sensor and compensator dynamics. A state-space representation of $K(s)$ may be written as

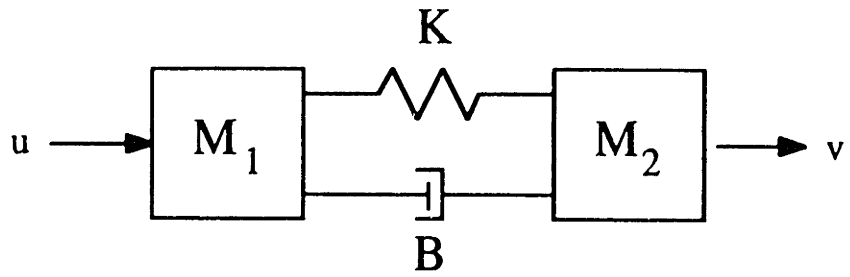

$$G(s) = \frac{v(s)}{u(s)} = \frac{Bs + K}{M_1 M_2 s^3 + (M_1 + M_2)Bs^2 + (M_1 + M_2)Ks}$$

Figure D.1: A system with non-colocated control and interaction ports. The relative order of $G(s)$ is 2.

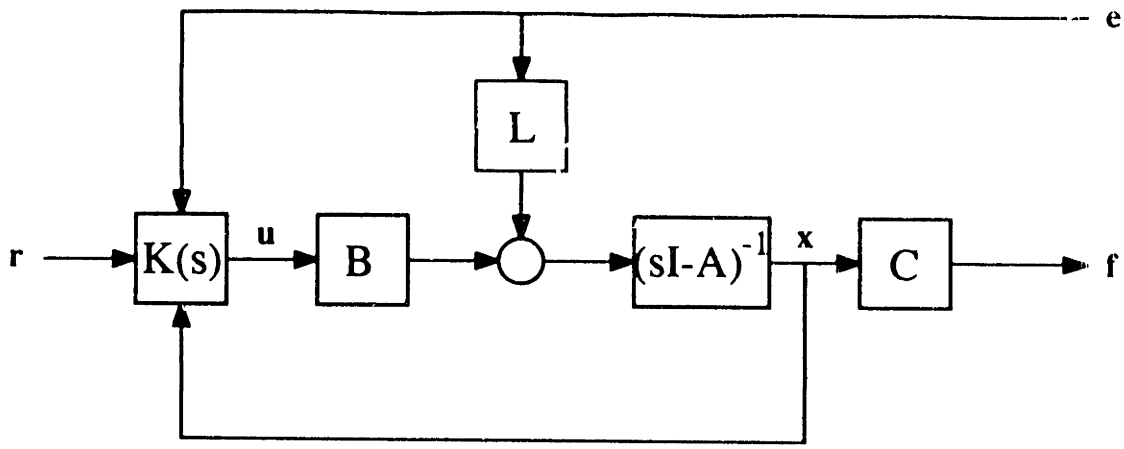

Figure D.2: Block diagram of a general input and state feedback controller with dynamic compensation. The system states are **x**, the control inputs are **u**, the exogenous inputs are **r**, the interaction port inputs are **e**, and the interaction port outputs are **f**.

follows:

$$\begin{aligned}\dot{\mathbf{z}} &= A_K\mathbf{z} + B_{Kx}\mathbf{x} + B_{Ke}\mathbf{e} + B_{Kr}\mathbf{r}, \quad A_K \text{ is } m \times m \\ \mathbf{u} &= C_K\mathbf{z} + D_{Kx}\mathbf{x} + D_{Ke}\mathbf{e} + D_{Kr}\mathbf{r}\end{aligned}$$

where the elements of $\mathbf{z}$ are the sensor and compensator states. The exogenous input, $\mathbf{r}$, will not affect the admittance, and is therefore set to zero. The closed loop state equations are:

$$\begin{aligned}\begin{bmatrix}\dot{\mathbf{x}} \\ \dot{\mathbf{z}}\end{bmatrix} &= \begin{bmatrix}A + BD_{Kx} & BC_K \\ B_K & A_K\end{bmatrix}\begin{bmatrix}\mathbf{x} \\ \mathbf{z}\end{bmatrix} + \begin{bmatrix}L + BD_{Ke} \\ B_{Ke}\end{bmatrix}\mathbf{e} \\ \mathbf{f} &= \begin{bmatrix}C & 0\end{bmatrix}\begin{bmatrix}\mathbf{x} \\ \mathbf{z}\end{bmatrix}\end{aligned}$$

Following the same reasoning as before, the matrix of leading coefficients of the closed loop driving point admittance, $Y_c(s)$, is:

$$\text{L.C.}\{\text{numerator}\, Y_c(s)\} = \begin{bmatrix}C & 0\end{bmatrix}\begin{bmatrix}L + BD_{Ke} \\ B_{Ke}\end{bmatrix}$$

where the leading coefficient now multiplies s^{n+m-1}, and the closed loop system is of order $n + m$. Thus:

$$\text{L.C.}\{\text{numerator}\, Y_c(s)\} = CL + CBD_{Ke}$$

But we know that $CB = 0$, so that the leading coefficients are all the same as they were before the control was implemented:

$$\text{L.C.}\{\text{numerator}\, Y_c(s)\} = CL$$

Thus the relative order of $Y(s)$ is unchanged by state or input feedback and, in particular, the relative order of the diagonal terms of $Y(s)$ remains exactly 1.

Because the leading coefficients of the numerator of $Y(s)$ are unchanged by feedback, and the leading coefficients of the denominator are all unity by convention, we have also shown that:

$$\lim_{s\to\infty} Y(s) = \lim_{s\to\infty} Y_c(s)$$

In other words, the high frequency behavior of the closed loop system is identical to that of the open loop system.

Bibliography

[1] C.J. Abul-Haj. *Elbow-Prosthesis Emulation: A Technique for the Quantitative Assessment of an Assistive Device.* PhD thesis, M.I.T. Department of Mechanical Engineering, June 1987.

[2] T.E. Alberts and D.I. Soloway. Force Control of a Multi-Arm Robot System. In *Proceedings of the IEEE Internation Conference on Robotics and Automation*, pages 1490–1496, 1988.

[3] C.H. An. *Trajectory and Force Control of a Direct Drive Arm.* PhD thesis, M.I.T. Department of Electrical Engineering and Computer Science, August 1986.

[4] B.D.O. Anderson and S. Vongpanitlerd. *Network Analysis and Synthesis: A Modern Systems Theory Approach.* Prentice-Hall, Englewood Cliffs, New Jersey, 1973.

[5] R.J. Anderson and M.W. Spong. Hybrid Impedance Control of Robotic Manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1073–1080, 1987.

[6] A.N. Andry, Jr. and R.C. Rosenberg. On the Dimension of State Space for Physical Systems. In *Proceedings of Joint Automation Control Conference*, San Francisco, CA, August 1980.

[7] H. Asada and J.-J.E. Slotine. *Robot Analysis and Control.* John Wiley and Sons, New York, 1986.

[8] M. Athans. *Lectures in Multivariable Control Systems, I and II.* MIT Department of Electrical Engineering and Computer Science, Cambridge, MA, 1986. Lecture notes.

[9] C.G. Atkeson, C.M. An, and J.M. Hollerbach. Estimation of Inertial Parameters of Manipulator Loads and Links. In *International Symposium of Robotics Research*, pages 221–228, 1985.

[10] J.S. Bendat and A.G. Piersol. *Random Data: Analysis and Measurement Procedures.* John Wiley and Sons, New York, 2 edition, 1986.

[11] R. Bott and R.J. Duffin. Impedance Synthesis without Use of Transformers. *J. Appl. Phys.*, 20:816, August 1949.

[12] R.W. Brockett. A Theorem on Positive Real Functions. *IEEE Trans. Circuit Theory (Correspondence)*, CT-11:301–302, June 1964.

[13] O. Brune. Synthesis of a Finite Two-Terminal Network whose Driving-Point Impedance is a Prescribed Function of Frequency. *J. Math. Phy., MIT*, 10:131–236, August 1931.

[14] J.E. Colgate. *The Design of a Dynamics Measuring Device.* Master's thesis, M.I.T. Department of Mechanical Engineering, January 1986.

[15] E.G. Cravalho and J.L. Smith, Jr. *Engineering Thermodynamics.* Pitman Publishing, Marshfield, Massachusetts, 1981.

[16] J. De Schutter. A Study of Active Compliant Motion Control Methods for Rigid Manipulators Based on a Generic Scheme. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1060–1065, 1987.

[17] R.C. Dorf. *Modern Control Systems.* Addison-Wesley, Reading, Massachusetts, 1980.

[18] S. Drake. *Using Compliance Instead of Sensory Feedback for High Speed Robot Assembly.* Technical Report CSDL-81-833, C. S. Draper Laboratory, 1981.

[19] S.D. Eppinger and W.P. Seering. Modeling Robot Flexibility for Endpoint Force Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 165–170, 1988.

[20] S.D. Eppinger and W.P. Seering. On Dynamic Models of Robot Force Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 29–34, 1986.

[21] S.D. Eppinger and W.P. Seering. Understanding Bandwidth Limitations in Robot Force Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 904–909, 1987.

[22] E.D. Fasse. *Stability Robustness of Impedance Controlled Manipulators Coupled to Passive Environments.* Master's thesis, M.I.T. Department of Mechanical Engineering, June 1987.

[23] I.C. Fayé. *An Impedance Controlled Manipulandum for Human Movement Studies.* Master's thesis, M.I.T. Department of Mechanical Engineering, June 1986.

[24] B.A. Francis. *A Course in H_∞ Control Theory. Lecture Notes in Control and Information Sciences*, Springer-Verlag, Berlin, 1987.

[25] W.B. Gevarter. Basic Relations for Control of Flexible Vehicles. *AIAA Journal*, 8(4):666–672, April 1970.

[26] A.A. Goldenberg. Force and Impedance Control of Robot Manipulators. In R. Shoureshi, K. Youcef-Youmi, and H. Kazerooni, editors, *Modeling and Control of Robotic Manipulators and Manufacturing Processes*, pages 161–168, ASME, 1987.

[27] A.A. Goldenberg. Implementation of Force and Impedance Control in Robot Manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1626–1632, 1988.

[28] E.M. Greitzer. Surge and Rotating Stall in Axial Flow Compressors: Part I: Theoretical Compression System Model. *ASME J. of Engineering for Power*, 98:190–198, April 1976.

[29] K.B. Haefner, P.K. Houpt, T.E. Baker, and M.E. Dausch. Real Time Robotic Position/Force Control for Deburring. In F.W. Paul and K. Youcef-Youmi, editors, *Robotics: Theory and Applications*, ASME, New York, 1986.

[30] W.E. Hamilton, Jr. Globally Stable Compliant Motion Control for Robotic Assembly. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1179–1184, 1988.

[31] S. Hayati. Hybrid Position/Force Control of Multi-Arm Cooperating Robots. In *Proceedings of the IEEE Internation Conference on Robotics and Automation*, pages 82–89, 1986.

[32] G. Hirzinger. Direct Digital Robot Control Using a Force-Torque Sensor. In *Proceedings of the IFAC Symposium on Real Time Digital Control Applications*, pages 243–255, 1983.

[33] N. Hogan. Control Strategies for Complex Movements Derived from Physical Systems Theory. In H. Haken, editor, *Complex Systems—Operational Approaches in Neurobiology, Physics, and Computers*, pages 156–168, Springer-Verlag, New York, May 1985.

[34] N. Hogan. Impedance Control: An Approach to Manipulation: Part III—Applications. *J. Dynamic Systems, Measurement, and Control*, 107:17–24, March 1985.

[35] N. Hogan. Impedance Control: An Approach to Manipulation: Part II—Implementation. *J. Dynamic Systems, Measurement, and Control*, 107:8–16, March 1985.

[36] N. Hogan. Impedance Control: An Approach to Manipulation: Part I—Theory. *J. Dynamic Systems, Measurement, and Control*, 107:1–7, March 1985.

[37] N. Hogan. Modularity and Causality in Physical System Modelling. *J. Dynamic Systems, Measurement, and Control*, 109:384–391. December 1987.

[38] N. Hogan. On the Stability of Manipulators Performing Contact Tasks. *IEEE Journal of Robotics and Automation*, 1988. in press.

[39] N. Hogan. Some Computational Problems Simplified by Impedance Control. In *Proceedings of ASME Computers in Engineering*, page , 1984.

[40] N. Hogan. Stable Execution of Contact Tasks Using Impedance Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1987.

[41] N. Hogan. Tuning Muscle Stiffness can Simplify Control of Natural Movement. In V.C. Mow, editor, *1980 Advances in Bioengineering*, pages 279–282, ASME, New York, 1980.

[42] N. Hogan, E. Bizzi, F.A. Mussa-Ivaldi, and T. Flash. Controlling Multijoint Motor Behavior. In K.B. Pandolf, editor, *Exercise and Sport Sciences Reviews*, pages 153–190, MacMillan Publishing Company, 1987.

[43] S.C. Jacobsen, D.F. Knutti, R.T. Johnson, and H.H. Sears. Development of the Utah Artificial Arm. *IEEE Trans. on Biomedical Engineering*, BME-29(4):249–269, April 1982.

[44] H. Kazerooni. *A Robust Design Method for Impedance Control of Constrained Dynamic Systems.* PhD thesis, M.I.T. Department of Mechanical Engineering, February 1985.

[45] H. Kazerooni. Robust, Non-Linear Impedance Control for Robot Manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 741–750, 1987.

[46] H. Kazerooni. Stability Analysis of Two Cooperating Robot Manipulators. In R. Shoureshi, K. Youcef-Toumi, and H. Kazerooni, editors, *Modeling and Control of Robotic Manipulators and Manufacturing Processes*, pages 203–212, ASME, New York, 1987.

[47] H. Kazerooni and P.K. Houpt. On the Loop Transfer Recovery. *International J. of Control*, 43(3), March 1986.

[48] J.H. Keenan. *Thermodynamics.* John Wiley and Sons, New York, 1957.

[49] O. Khatib. A Unified Approach for Motion and Force Control of Robot Manipulators: The Operational Space Formulation. *IEEE Journal of Robotics and Automation*, RA-3(1):43–53, February 1987.

[50] O. Khatib and J. Burdick. Motion and Force Control of Robot Manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1381–1386, 1986.

[51] P.K. Khosla and T. Kanade. Parameter Identification of Robot Dynamics. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1754–1760, 1985.

[52] P.V. Kokotovic. Applications of Singular Perturbation Techniques to Control Problems. *SIAM Review*, 26(4):396–410, 1984.

[53] E.S. Kuh and R.A. Rohrer. *Theory of Linear Active Networks.* Holden-Day, San Francisco, 1967.

[54] H. Kwakernaak and R. Sivan. *Linear Optimal Control Systems.* Wiley-Interscience, New York, 1972.

[55] D.A. Lawrence. Impedance Control Stability Properties in Common Implementations. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1185–1190, 1988.

[56] N.A. Lehtomaki. *Practical Robustness Measures in Multivariable Control.* PhD thesis, MIT Department of Electrical Engineering and Computer Science, 1981.

[57] R.W. Mann. Cybernetic Limb Prostheses. *Annals of Biomedical Engineering*, 9(1):1–43, 1981.

[58] J.M. Mansfield. *The Design of a Lightweight Elbow Prosthesis Emulator.* Master's thesis, M.I.T. Department of Mechanical Engineering, May 1988.

[59] M.T. Mason. Compliance and Force Control for Computer Controlled Manipulators. *IEEE Trans. on Systems, Man, and Cybernetics*, SMC-11(6):418–432, June 1981.

[60] W.R. Murray. *Essential Factors in Modeling Modulation of Impedance About the Human Elbow.* PhD thesis, M.I.T. Department of Mechanical Engineering, June 1988.

[61] K.S. Narendra and J.H. Taylor. *Frequency Domain Criteria for Absolute Stability.* Academic Press, New York, 1973.

[62] R.E.A.C. Paley and N. Wiener. *Fourier Transforms in the Complex Domain.* American Mathematical Society, New York, 1934.

[63] R.P. Paul. Problems and Research Issues Associated with the Hybrid Control of Force and Displacement. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1966–1971, 1987.

[64] R.P. Paul and B. Shimano. Compliance and Control. In *Proceedings of the Joint Automatic Control Conference*, 1976.

[65] H.M. Paynter. *Analysis and Design of Engineering Systems.* M.I.T. Press, Cambridge, Mass, 1961.

[66] M. Raibert and J. Craig. Hybrid Position/Force Control of Manipulators. *J. Dynamic Systems, Measurement, and Control*, 102:126–133, June 1981.

[67] R.K. Roberts, R.P. Paul, and B.M. Hillberry. The Effect of Wrist Force Sensor Stiffness on the Control of Robot Manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 269–274, 1985.

[68] R.C. Rosenberg and D.C. Karnopp. *Introduction to Physical Systems Dynamics.* McGraw-Hill, New York, 1983.

[69] A. Sharon, N. Hogan, and D. Hardt. High Bandwidth Force Regulation and Inertia Reduction Using a Macro/Micro Manipulator System. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1988.

[70] A. Sharon, N. Hogan, and D. Hardt. More Analysis and Experimentation on a Macro/Micro Manipulator System. In R. Shoureshi, K. Youcef-Youmi, and H. Kazerooni, editors, *Modeling and Control of Robotic Manipulators and Manufacturing Processes*, pages 417–422, ASME, 1987.

[71] K.G. Shin and C. Lee. Compliant Control of Robotic Manipulators with Resolved Acceleration. In *Proceedings of the Twenty-Fourth Conference on Decision and Control*, pages 350–357, 1985.

[72] J. Simon. Personal communication.

[73] J.-J.E. Slotine and S.S. Sastry. Tracking Control of Nonlinear Systems Using Sliding Surfaces, with Applications to Robot Manipulators. *Int. J. Control*, 38(2):465–492, 1983.

[74] M.W. Spong. Force Feedback Control of Flexible Joint Manipulators. In R. Shoureshi, K. Youcef-Toumi, and H. Kazerooni, editors, *Modeling and Control of Robotic Manipulators and Manufacturing Processes*, pages 177–184, ASME, 1987.

[75] G. Stein. *Lectures in Multivariable Control Systems, II.* MIT Department of Electrical Engineering and Computer Science, Cambridge, MA, May 1988. Lecture notes.

[76] T.M. Stepien, L.M. Sweet, M.C. Good, and M. Tomizuka. Control of Tool/Workpiece Contact Force with Application to Robotic Deburring. *IEEE Journal of Robotics and Automation*, RA-3(1):7–18, February 1987.

[77] G. Strang. *Linear Algebra and its Applications.* Academic Press, New York, 1980.

[78] M. Takegaki and S. Arimoto. A New Fedback Method for Dynamic Control of Manipulators. *J. Dynamic Systems, Measurement, and Control*, 103:119–125, June 1981.

[79] A. Talbot. Some Theorems on Positive Functions. *IEEE Trans. Circuit Theory (Correspondence)*, CT-12:607–608, December 1965.

[80] A. Talbot. Transformation of Positive Functions by Linear Operators. *IEEE Trans. Circuit Theory*, CT-19(5):460–465, September 1972.

[81] T.J. Tarn, A.K. Bejczy, and X. Yun. Coordinated Control of Two Robot Arms. In *Proceedings of the IEEE Internation Conference on Robotics and Automation*, pages 1193–1202, 1986.

[82] S.W. Tilley, R.H. Cannon, Jr., and R. Kraft. End Point Force Control of a Very Flexible Manipulator with a Fast End Effector. In F.W. Paul and K. Youcef-Youmi, editors, *Robotics: Theory and Applications*, ASME, New York, 1986.

[83] W.T. Townsend. *The Effect of Transmission Design on the Performance of Force Controlled Manipulators.* PhD thesis, M.I.T. Department of Mechanical Engineering, May 1988.

[84] W.T. Townsend and J.K. Salisbury. The Effect of Coulomb Friction and Stiction on Force Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 883–889, 1987.

[85] M.E. Van Valkenburg. *Introduction to Modern Network Synthesis.* John Wiley and Sons, New York, 1960.

[86] M. Vidyasagar. *Nonlinear Systems Analysis.* Prentice-Hall, Englewood Cliffs, New Jersey, 1978.

[87] L. Weinberg. *Network Analysis and Synthesis.* McGraw-Hill, New York, 1962.

[88] D.E. Whitney. Force Feedback Control of Manipulator Fine Motions. In *Proceedings of the Joint Automatic Control Conference*, 1976.

[89] D.E. Whitney. Historical Perspective and State of the Art in Robot Force Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 262–268, 1985.

[90] J.C. Willems. Dissipative Dynamical Systems. Part 1: General Theory. *Arch. Rational Mech. Anal.*, 45(5):321–351, 1972.

[91] J.C. Willems. Dissipative Dynamical Systems. Part 2: Linear Systems with Quadratic Supply Rates. *Arch. Rational Mech. Anal.*, 45(5):352–393, 1972.

[92] J.J. Wlassich. *Nonlinear Force Feedback Impedance Control.* Master's thesis, M.I.T. Department of Mechanical Engineering, February 1986.

[93] J.L. Wyatt, Jr., L.O. Chua, J.W. Gannett, I.C. Göknar, and D.N. Green. Energy Concepts in the State-Space Theory of Nonlinear n-Ports: Part I—Passivity. *IEEE Trans. Circuits Syst.*, CAS-28(1):48–61, Jan 1981.

[94] J.L. Wyatt, Jr. and D.L. Standley. Circuit Design Criteria for Stable Lateral Inhibition Neural Networks. In *Proceedings of IEEE International Symposium on Circuits and Systems*, pages 997–1000, Espoo, Finland, June 1988.

[95] T. Yabuta, A.J. Chona, and G. Beni. On the Asymptotic Stability of the Hybrid Position/Force Control Scheme for Robot Manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 338–343, 1988.

[96] K. Youcef-Toumi and D. Li. Force Control of Direct Drive Manipulators for Surface Following. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 1987.