

Modeling Robot Dynamic Performance for Endpoint Force Control

by
Steven D. Eppinger

BSME Massachusetts Institute of Technology
(February 1983)

MSME Massachusetts Institute of Technology
(June 1984)

Submitted to the
Department of Mechanical Engineering
in partial fulfillment of the
requirements for the degree of

Doctor of Science

at the

Massachusetts Institute of Technology
September 1988

Copyright ©1988 Massachusetts Institute of Technology

Signature of Author _____
Department of Mechanical Engineering
September 1, 1988

Certified by _____
Professor Warren P. Seering
Thesis Supervisor

Accepted by _____
Ain A. Sonin
Chairman, Departmental Graduate Committee

Archives

MASS. INST. TECH.

MAR 16 '88

Acknowledgments

This thesis describes research conducted at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Office of Naval Research University Research Initiative Program under Office of Naval Research contract N00014-86-K-0685, and in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-85-K-0124.

Support for this research project was also provided in part by generous grants from TRW. Jack Searle, arranged for my initial funding through the TRW Fasteners Division. Arden Bement coordinated the continuation of this fellowship through the TRW Foundation. I greatly appreciate their confidence in my ability and their tireless efforts on my behalf. It has been a sincere pleasure to work with TRW.

I feel privileged to have been able to work in the MIT Artificial Intelligence Laboratory, perhaps the finest research environment in the world. Patrick Winston and Tomas Lozano-Perez have made a tremendous effort to bring together an expanse of state-of-the-art computers, a variety of robot hardware, a collection of very bright people, and even a machine shop. This combination has made my doctoral research a truly remarkable educational experience.

The members of Warren Seering's research group have shared with me the joys and struggles of graduate student life at MIT. Each of the following has contributed directly to my research: Mike Benjamin, Mike Caine, Andy Christian, Alfonso Garcia-Reynoso, Steve Gordon, Peter Meckl, Ken Pasch, Jose Rivero, Neil Singer, Karl Ulrich, Erik Vaaler, and Al Ward. I am proud to call this group my closest peers.

There are many other graduate students at MIT who also shaped this work through enlightening conversations on robotics, dynamics, control, computation, and design. Some of these include: Chae An, Chris Atkeson, Tom Bailey, Dave Brock, Ed Colgate, Anita Flynn, Jon Hirschtick, Jeff Kochling, Sundar Narasimhan, Bill Townsend, and John Wlassich. I appreciate their input and look forward to continued association with them.

My Doctoral Thesis Committee, consisting of Karl Hedrick, Don Pieper, Marc Raibert, Ken Salisbury, and Warren Seering, has been extremely helpful over the course of this research. Their direction of my efforts helped this project produce useful results, and a readable document.

By far the greatest influence on my research has been through the guidance and encouragement of my advisor and friend, Warren Seering. Our frequent and lengthy discussions ranged in topic from the behavior of machine systems to the behavior of student systems. The quality of the results in this document and their acceptance in the technical community are largely due to Warren's close attention to my research progress and its presentation. I will model my interactions with students after Warren's fine example.

The support of my parents, Robert and Paula Eppinger, and my brothers, Jeffrey and Bruce over the recent years has been invaluable. Mom and Dad have offered words of wisdom and confidence. Jeff's parallel effort in doctoral research has given us frequent opportunities to step back for a broader view of the process. Family brings me great happiness and I hope we continue to grow together for years to come.

Finally, Julie Laukkanen, my wonderful wife, has completed my life with all the necessary diversions. Her countless contributions to our success and happiness include loving affection, financial support, personal ambition, and most spectacular of all, our daughter Lauren.

Table of Contents

Abstract	2
Acknowledgments	3
Table of Contents	5
List of Figures	8
List of Models	11
Thesis Summary	12
1. Introduction	14
1.0.1 A Note on Terminology	15
1.1 Force Control Research Background	15
1.1.1 Passive Compliance	16
1.1.2 Force Monitoring	16
1.1.3 Active Compliant Control	17
1.1.4 Assembly Strategy and Planning	22
1.1.5 Trials Using Force Control	23
1.1.6 Robot Hardware Design	24
1.1.7 Sensor Development	25
1.1.8 Modeling and Stability Analysis	27
1.2 Motivation	28
1.2.1 Some Observations	28
1.2.2 Causes of Force Control Instability	29
1.2.3 Problem Statement	30
1.3 Overview of This Thesis	31
2. Modeling Machine Performance	32
2.1 Introduction	32
2.2 Rigid-Body Dynamics	34
2.2.1 The Simplest Robot Model	35
2.2.2 Model Behavior	35
2.2.3 Interpretation	37
2.3 Actuator Performance	37
2.3.1 Perfect and Real Actuators	37
2.3.2 Actuator/Load Coupling	38
2.3.3 A Simple Actuator Example	39
2.3.4 Actuator-Coupled Rigid-Body Model Performance	39
2.4 Modeling Flexibility	42
2.4.1 Types of Machine Flexibility	42
2.4.2 Transmission and Link Flexibility	42
2.4.3 Base Dynamics	46
2.4.4 Better Bending Models	48
2.5 Task Dynamics	55
2.5.1 The Dynamic Workpiece Model	55
2.5.2 Adding Robot Dynamics	56
2.5.3 Adding Actuator Dynamics	56

2.6 Grip Compliance	62
2.7 Nonlinearities	66
2.7.1 Discontinuity at the Workpiece Contact	67
3. Measuring Robot Performance	70
3.0.1 Deciding What to Measure	70
3.1 Experimental Hardware	71
3.1.1 Frequency Response Tests	71
3.1.2 Transient Response Tests	73
3.1.3 The Test Workpiece	74
3.2 Frequency Response Tests	75
3.2.1 Actuator Transfer Functions	76
3.2.2 Robot System Transfer Functions	76
3.2.3 Comments on the Measured Frequency Response	79
3.3 Time Response Tests	82
3.3.1 Position-Control Response	82
3.3.2 Force Control Transient Response	82
4. Modeling the MIT Precision Assembly Robot	87
4.1 Actuator Dynamics	87
4.1.1 Amplifier and DC Motor Model	88
4.1.2 Amplifier and Motor Parameters	89
4.1.3 No-Load Motor Performance	90
4.2 Polynomial Resonance Model	92
4.3 Lumped-Mass Models	95
4.3.1 Coupled Rigid-Body Model Performance	95
4.3.2 Modeling Everything	98
4.3.3 Reduced-Order Models	98
4.3.4 Arm Flexibility	99
4.3.5 Base Dynamics	100
4.3.6 Workpiece Dynamics	101
4.3.7 Grip Compliance	103
4.4 Finding the Lumped-Mass Model Parameters	107
4.4.1 Rigid-Body Robot Parameters	107
4.4.2 Flexible Robot Parameters	107
4.4.3 Base Model Parameters	108
4.4.4 Sensor Model Parameters	108
4.4.5 Grip Compliance	108
4.4.6 Workpiece Model Parameters	108
4.5 Nonlinear Model Simulations	109
4.5.1 Position Control Step Response	109
4.5.2 Force Control Transient Response	111
5. Improving Performance	112
5.1 Grip Compliance	112
5.1.1 Model Predictions	112
5.1.2 Experimental Data	113
5.1.3 Interpretation	113
5.2 Workpiece Damping	115
5.2.1 Workpiece Damping Tests	115

6. Conclusion	122
6.1 Contributions of This Research	122
6.1.1 Modeling Machine Performance	122
6.1.2 Understanding Robot Bandwidth Limitations	123
6.1.3 Force Control Performance Improvements	124
6.1.4 Insights for Machine Design	124
6.2 Recommendations for Future Research	126
6.2.1 Robot Control	126
6.2.2 Machine Design	127
6.2.3 System Dynamics	128
6.2.4 Modeling of Machine Systems	128
References	130

List of Figures

Figure 1.1:	Position Control and Force Control Tasks	14
Figure 1.2:	Position Control	18
Figure 1.3:	Explicit Force Control	19
Figure 1.4:	Hybrid Position/Force Control	20
Figure 1.5:	Stiffness Control	21
Figure 1.6:	Damping Control	21
Figure 1.7:	Impedance Control	22
Figure 2.1:	Development of One-Axis Robot Models	33
Figure 2.2:	Single-Axis Explicit Force Control	34
Figure 2.3:	Rigid-Body Robot Model with Rigid Workpiece	35
Figure 2.4:	Bode Plot for the Rigid-Body Robot Model of Figure 2.3	36
Figure 2.5:	Root-Locus Plot for the Rigid-Body Robot Model of Figure 2.3	36
Figure 2.6:	Actuator Decoupling Model	38
Figure 2.7:	Simple Actuator Model	39
Figure 2.8:	Simple Actuator Model Coupled to the Rigid-Body Robot Model	40
Figure 2.9:	Bode Plot for the Rigid-Body Coupled Model of Figure 2.8	41
Figure 2.10:	Root-Locus Plot for the Rigid-Body Coupled Model of Figure 2.8	41
Figure 2.11:	Robot Model Including Transmission or Link Flexibility	43
Figure 2.12:	Bode Plot for the Flexible Robot Model of Figure 2.11	44
Figure 2.13:	Root Locus Plot for the Flexible Robot Model of Figure 2.11	44
Figure 2.14:	Bode Phase Plot Showing Phase Shift Caused by Noncollocation	45
Figure 2.15:	Mode Shape of the 12 Hz Resonance	47
Figure 2.16:	Two Robot Models Including Base Flexibility	48
Figure 2.17:	Bode Plot for the First Base Flexibility Model of Figure 2.16a	49
Figure 2.18:	Root Locus Plots for the First Base Flexibility Model of Figure 2.16a	49
Figure 2.19:	Bode Plot for the Second Base Flexibility Model of Figure 2.16b	50
Figure 2.20:	Root Locus Plots for the Second Base Flexibility Model of Figure 2.16b	50
Figure 2.21:	Distributed-Parameter Beam Robot Model	51
Figure 2.22:	Bode Plots for the Distributed-Parameter Beam Model of Figure 2.21	52
Figure 2.23:	Pole-Zero Arrays for the Distributed-Parameter Beam Model of Figure 2.21	53
Figure 2.24:	Lumped-Parameter Beam Model	54
Figure 2.25:	Pole-Zero Arrays for the Lumped-Parameter Beam Model of Figure 2.24	54
Figure 2.26:	Robot Model Including Workpiece Dynamics	55
Figure 2.27:	Bode Plots for the Workpiece Model of Figure 2.26	57
Figure 2.28:	Root Locus Plot for the Workpiece Model of Figure 2.26	57
Figure 2.29:	Robot Model Including Workpiece Dynamics and Robot Flexibility	58
Figure 2.30:	Root Locus Plot for the Workpiece Model with Robot Flexibility of Figure 2.29	59
Figure 2.31:	Root Locus Plot for the Workpiece Model Showing a Possibly Stabilizing Case	59
Figure 2.32:	Robot Model Including Workpiece and Actuator Dynamics	60
Figure 2.33:	Root Locus Plot for the Coupled Workpiece Model of Figure 2.32	61
Figure 2.34:	Root Locus Plot for the Workpiece Model Showing Another Stabilizing Case	61
Figure 2.35:	Robot Model Including Grip Compliance	62
Figure 2.36:	Bode Plots for the Uncoupled Grip Model of Figure 2.35	64
Figure 2.37:	Root Locus Plot for the Uncoupled Grip Model of Figure 2.35	64

Figure 2.38:	Bode Plots for the Coupled Grip Model of Figure 2.35	65
Figure 2.39:	Root Locus Plot for the Coupled Grip Model of Figure 2.35	65
Figure 2.40:	Rigid-Body Robot Model with Discontinuity	69
Figure 2.41:	Rigid-Body Discontinuous Model Response	69
Figure 2.42:	Flexible Robot Model with Discontinuity	69
Figure 2.43:	Flexible Discontinuous Model Response	69
Figure 3.1:	The MIT Precision Assembly Robot	72
Figure 3.2:	Robot Frequency Response Test Setup	72
Figure 3.3:	Structure of the Robot/Actuator System	73
Figure 3.4:	Robot Control Hardware and Experimental Setup	74
Figure 3.5:	The Test Workpiece Mounted to the Robot Frame	75
Figure 3.6:	Preamplifier Velocity Command Transfer Function	77
Figure 3.7:	Preamplifier Tachometer Feedback Transfer Function	77
Figure 3.8:	Amplifier/Motor Current Transfer Function	78
Figure 3.9:	Amplifier/Motor Velocity Transfer Function	78
Figure 3.10:	Measured Robot Colocated Acceleration Transfer Function	80
Figure 3.11:	Measured Robot Noncolocated Acceleration Transfer Function	80
Figure 3.12:	Measured Robot Noncolocated Force Transfer Function	81
Figure 3.13:	Robot Position-Control Step Response Data	83
Figure 3.14:	Robot Force-Control Step Response Data	84
Figure 3.15:	Robot Contact Stability Depends on Force Gain and Force Setpoint	85
Figure 3.16:	Summary of Robot Contact Stability for Various Gains	86
Figure 4.1:	Robot Actuator Model	88
Figure 4.2:	Simple Saturation	89
Figure 4.3:	PWM Amplifier Two-Gain Saturation	90
Figure 4.4:	Motor Model Parameters	91
Figure 4.5:	No-Load Motor Model with Rotor Inertia and Damping Only	91
Figure 4.6:	Bode Plot for the No-Load Coupled Actuator Model of Figure 4.5	92
Figure 4.7:	Polynomial Resonance Model Load	93
Figure 4.8:	Bode Plot for the Polynomial Resonance Model of Figure 4.7	94
Figure 4.9:	Pole-Zero Array of the Polynomial Resonance Model of Figure 4.7	95
Figure 4.10:	Rigid-Body Robot Model with Rigid Workpiece	96
Figure 4.11:	Rigid-Body Robot Model Load Coupled to the Motor Model	96
Figure 4.12:	Bode Plot for the Rigid-Body Coupled Model of Figure 4.11	97
Figure 4.13:	Root-Locus Plot for the Rigid-Body Coupled Model of Figure 4.11	97
Figure 4.14:	A Robot Model Combining All of the Previous Lumped-Mass Model Effects	98
Figure 4.15:	Robot Model Including Arm Flexibility	99
Figure 4.16:	Root-Locus Plot for the Coupled Model with Arm Flexibility of Figure 4.15	100
Figure 4.17:	Robot Model Including Base Flexibility	101
Figure 4.18:	Root-Locus Plot for the Coupled Model with Base Flexibility of Figure 4.17	101
Figure 4.19:	Robot Model Including Workpiece Dynamics	102
Figure 4.20:	Root-Locus Plot for the Coupled Model with Workpiece Dynamics of Fig. 4.19	102
Figure 4.21:	Root-Locus Plot for the Coupled Robot Model with Stiffer Workpiece Dynamics	104
Figure 4.22:	Root-Locus Plot for the Coupled Robot Model with Softer Workpiece Dynamics	104
Figure 4.23:	Robot Model Including Grip Compliance	105
Figure 4.24:	Root-Locus Plot for the Coupled Robot Model with Stiff Grip	106
Figure 4.25:	Root-Locus Plot for the Coupled Robot Model with Soft Grip	106
Figure 4.26:	Robot Model Parameters	109
Figure 4.27:	Nonlinear Robot Position-Control Step Response Simulation	110
Figure 4.28:	Nonlinear Coupled Flexible Robot Model Force-Control Simulation	111
Figure 5.1:	Force Data Showing Instability with the Rigid Grip	114
Figure 5.2:	Force Data Showing Instability with the Hard Grip Pad	114
Figure 5.3:	Force Data Showing Improved Stability with the Soft Grip Pad	114

Figure 5.4: Root Locus for the Robot Model with the Original Workpiece Damping	116
Figure 5.5: Root Locus for the Robot Model with Increased Workpiece Damping	116
Figure 5.6: Workpiece Damping Trials	118
Figure 5.7: Design of the Workpiece Damper	119
Figure 5.8: Robot Force-Control Step Response Data for the Original Workpiece	120
Figure 5.9: Robot Force-Control Step Response Data for the Damped Workpiece	120
Figure 5.10: Contact Stability Is Increased with Workpiece Damping	121
Figure 6.1: Three Dynamic Problems in Machine Control	123
Figure 6.2: Types of Robot Flexibility	125
Figure 6.3: Nonlinear Two-Link Rigid Robot Model with Nonminimum Phase Zeros	128

List of Models

1. Rigid-Body Robot Model with Rigid Workpiece	35
2. Actuator Decoupling Model	38
3. Simple Actuator Model	39
4. Simple Actuator Model Coupled to the Rigid-Body Robot Model	40
5. Robot Model Including Transmission or Link Flexibility	43
6. Two Robot Models Including Base Flexibility	48
7. Distributed-Parameter Beam Robot Model	51
8. Lumped-Parameter Beam Model	54
9. Robot Model Including Workpiece Dynamics	55
10. Robot Model Including Workpiece Dynamics and Robot Flexibility	58
11. Robot Model Including Workpiece and Actuator Dynamics	60
12. Robot Model Including Grip Compliance	62
13. Rigid-Body Robot Model with Discontinuity	69
14. Flexible Robot Model with Discontinuity	69
15. Robot Actuator Model	88
16. No-Load Motor Model with Rotor Inertia and Damping Only	91
17. Polynomial Resonance Model Load	93
18. Rigid-Body Robot Model with Rigid Workpiece	96
19. Rigid-Body Robot Model Load Coupled to the Motor Model	96
20. A Robot Model Combining All of the Previous Lumped-Mass Model Effects	98
21. Robot Model Including Arm Flexibility	99
22. Coupled Robot Model with Arm Flexibility	100
23. Robot Model Including Base Flexibility	101
24. Coupled Robot Model with Base Flexibility	101
25. Robot Model Including Workpiece Dynamics	102
26. Coupled Robot Model with Workpiece Dynamics	102
27. Robot Model Including Grip Compliance	105
28. Coupled Robot Model with Grip Flexibility	106
29. Nonlinear Models of the MIT Precision Assembly Robot	109
30. Nonlinear Two-Link Rigid Robot Model with Nonminimum Phase Zeros	128

Thesis Summary

The goal of this research project is to discover and explain the factors which can cause machines to be unstable under closed-loop sensory feedback control. We develop insights for machine designers which enable them to **build systems capable of performing with a desired bandwidth**. The example we consider is **robot endpoint force control**, also known as manipulator compliant motion control. Force control research has been quite popular recently, yet the dynamics involved have not been adequately explained in the literature. The design issues addressed in this thesis include: actuator characteristics, structural dynamics, sensor placement, and process interaction. While we consider these issues as they apply specifically to robotics, they also apply directly to machine design in general.

Heavy emphasis is placed upon learning from **mathematical models of machine performance**. A series of models is developed in an effort to explain observed system behavior. Specifically, robots can exhibit violent unstable response under endpoint force control. The model developed here show why this behavior occurs and predicts under what conditions stable interactions can be made. The models development is guided by analysis of **laboratory measurements of machine performance**. The experimental research serves to verify that the models do indeed depict real robot systems.

To improve robot force control performance, we rely on the dynamic models to evaluate proposed changes to the system. Two schemes are analyzed and implemented on the MIT Precision Assembly Robot for verification. These performance improvements are adding **grip compliance** and adding **workpiece damping**. Grip compliance can simply be in the form of a soft pad between the robot endpoint and the environment's contact surface. Low grip stiffness essentially decouples the robot from the workpiece. This allows the robot to achieve higher bandwidth although it is less responsive to the interactions, requiring higher force control gains and larger motions to accomplish its programmed task. Workpiece damping, added in parallel to the workpiece stiffness, does not change the loop gain of the system. To the extent that the dominant poles become better damped, higher control gains can be used and higher bandwidth is achieved.

Machine performance limitations are discussed in terms of rigid-body rolloff, dynamically colocated modes, and dynamically noncolocated modes. The limited **rigid-body performance** results from the coupling of an imperfect actuator to a rigid-body load. **Dynamically colocated modes** result from flexibility of structural elements supporting the actuator or a dynamic

environment. In these modes, the actuator and feedback sensor move together (with little phase shift) throughout the frequency range of interest. **Dynamically noncolocated modes** result from flexible elements between the actuator and sensor. In these modes, the flexibility allows the actuator and sensor to move opposite in phase (from their rigid-body motion). Colocated modes can cause instability in machine systems, however this effect can be alleviated by the addition of suitable damping. **Noncolocated modes present a fundamental performance limitation**, adding 180° of phase shift to the open-loop transfer function.

Machine designers should use a dynamic performance model as part of the system development effort. Rigid-body response can be predicted before actuators and transmissions are specified. Estimates of the frequencies of structural resonances need to be made and incorporated into the model. Low-frequency colocated modes should be well damped, and noncolocated modes should be raised above the phase crossover of the rigid-body model. Use of the proper dynamic model is essential. **A good robot model should demonstrate the difference between colocated control and noncolocated control.** This thesis shows how to develop and learn from models of machine performance.

Chapter One

Introduction

Most robot applications are limited to operations where the manipulator has little or no interaction with its environment. Control systems have been developed to guide the manipulator through a programmed trajectory consisting of a sequence of joint positions and velocities. Position and velocity sensors located at the robot actuators or joints are used for closed-loop feedback. We call this scheme **robot position control**, and this is the conventional method of robot servo control. Position control is useful in tasks where the robot is not constrained by objects in the workspace. Examples of successful industrial applications of position control include spray painting and pick-and-place tasks.

However, in many automated manufacturing processes, the robot is required to interact with its environment in a more controlled manner, but that environment is only imprecisely known. In mechanical assembly, for example, the manipulator must mate parts whose position, orientation, size, and shape are somewhat uncertain. **Robot force control** is the term used to describe control schemes where measurements of the interaction forces are used to alter the commands given to the joint servos. Many different force control schemes have been developed by researchers in the robot control field. The more common ones will be reviewed here.

Figure 1.1 illustrates two distinctly different types of robot tasks. Manipulation tasks involving little or no interaction with the environment are well-suited to robot position control. Tasks with significant interaction require some form of robot force control.




Figure 1.1: Position Control and Force Control Tasks

1.0.1 A Note on Terminology

In robotics, the term *force* often refers to the generalized force/torque vector. Similarly, the term *position* refers to the generalized position/orientation vector. We will use the term *compliant control* to mean the ability of a robot to react in response to endpoint forces. Strictly speaking, *compliance* is the inverse of mechanical *stiffness*. However, in common usage, *compliant control* refers to all of the schemes developed where joint servo commands depend upon both commanded trajectories and measured forces of interaction. Finally, the terms *force control* and *compliant control* are often used interchangeably. Generally speaking, the robot's *environment* is that part of the manipulator's immediate surroundings with which the endpoint makes contact. An *obstacle* is part of the environment that is not intended to be contacted. A *workpiece* is the specific object in the environment with which the robot does attempt to interact.

An element is said to be *rigid* if it has no *compliance* or *flexibility*, and the object is called *stiff* if it has high *stiffness*, which is low *compliance*. A robot *axis* is generally one *link* and one *joint* of the manipulator. Feedback *sensors* are usually located at the *actuators*, or at the *joints*, or at the *endpoint*. An actuator and joint are possibly separated by a *transmission*, while the joint and endpoint are separated by the link itself. The *tip* is the same as the *endpoint*, where a *gripper* may be located.

1.1 Force Control Research Background

Research in robot compliant motion control has been in these areas:

- **Passive compliance** schemes for immediate industrial application.
- **Force monitoring** and branching upon force thresholds.
- **Active compliant control** strategies implementing six-axis force control.
- **Assembly strategy and planning** systems which program robots for compliant motion.
- **Trials using force control** to accomplish simulated industrial tasks.
- **Robot hardware design** to create new systems using force control.
- **Sensor development** to devise new ways to measure robot interaction.
- **Modeling and stability analysis** to understand how these systems behave.

This thesis presents research which fits best in the final category, robot modeling and stability analysis, however it applies more broadly to the design and control of machines in general. This chapter introduces robot force control, briefly reviews some of the work in this area, and motivates the research to be presented in the rest of the thesis.

1.1.1 Passive Compliance

Robot force control schemes describe ways to achieve compliant motion. The simplest form of force control uses conventional robot position control and some form of passive compliance. The more complex procedures use active compliance, the distinction between active and passive compliance being whether or not forces of interaction are used for feedback.

In industrial applications, where robustness and speed are of great concern, compliance is currently achieved passively, rather than through active control. **Passive compliance** strategies have emerged which allow position-controlled manipulators to perform mechanical assembly in the face of limited uncertainty. These schemes are not strictly force control in that they do not program the servos to actively respond to the forces at the endpoint. They include:

- robot joint and structural compliance
- the RCC compliant robot wrist

Joint compliance is achieved through "softening" the position control loops by lowering the loop gains. **Structural compliance** is the inherent flexibility of the robot links themselves. The combination of low servo stiffness and link flexibility allows the endpoint to respond in a "natural" way to the forces applied to the end effector, often resulting in successful task completion. Unfortunately, it is frequently the case that this "natural" compliant behavior is unable to successfully perform the task required.

The most successful passive strategy involves an end-effector tool which allows a position-controlled manipulator to perform very accurate mechanical assembly tasks. The **remote center compliance** [96] (RCC) is a passive device configured to automatically execute a successful peg-in-hole strategy. The RCC mounts to the robot wrist and creates an effective compliance center at a fixed distance from the wrist. Forces acting at the compliance center result in pure translations of that point, while torques about the center of compliance result in pure rotations. With the center of compliance at the tip of the "peg", the RCC strategy allows workpieces to be assembled with little or no risk of jamming or wedging, even with workpiece clearances well below the robot positioning capability. Commercially-available RCC wrists [12] work well for parts that can be gripped so that the tip is very near the compliance center. Otherwise, the RCC strategy fails.

1.1.2 Force Monitoring

Early force control schemes use relatively simple sensors and threshold monitoring to determine when contact geometry changes have been achieved or when the gripper has grasped or dropped a part. These schemes usually consist of serially-executed statements in a robot command language, and the comparisons made are basically scalar in nature. Typical examples are the **guarded move** ("MOVE IN $-Z$ UNTIL F_Z EXCEEDS F_{ZMAX} .") and **logic branching** ("IF $T_Y > 0$ THEN

CONTINUE ELSE GOTO PICKUP.") The limitation of these schemes is their binary control logic. They can generally detect process changes only after they have occurred, rather than while they are taking place.

1.1.3 Active Compliant Control

Much attention has been given to the development of **compliant motion control** algorithms where contact force information is used continuously in the arm control. Multi-axis force/torque signals are used to alter the trajectory commands given to some or all of the joint servo controllers. These strategies provide the greatest flexibility in terms of the variety of tasks that can be programmed. Compliant control schemes have been developed which allow the robot to react to sensed forces in many different ways. However, they have demonstrated serious stability problems giving rise to bandwidth limitations and system performance that is generally unsatisfactory for automated industrial applications. Nevertheless, compliant motion control remains an active research topic and will someday become practical for a broad range of tasks.

Active force control research began in the 1950s when Goertz added force reflection to his master/slave manipulators [35]. Force feedback was also used by Mann in the control of an artificial elbow [76]. Whitney traces the development of force control algorithms and applications, also providing numerous references [98]. Early robot active compliance strategies were suggested by Inoue [46] and Paul and Shimano [65].

There are many possible control implementations which can achieve compliant motion. This section will overview some of the more popular schemes. They all allow the arm to respond to contact forces by altering the programmed trajectory. In the block diagrams of the control schemes presented here, we use the following nomenclature:

θ	measured joint positions
$\dot{\theta}$	measured joint velocities
τ_c	commanded joint actuator torques
x	computed endpoint positions
x_d	desired endpoint positions (input trajectory)
\dot{x}_d	desired endpoint velocities (input trajectory)
f_s	measured endpoint contact forces
f_d	desired endpoint contact forces
J	arm Jacobian matrix ($\dot{x}=J\dot{\theta}$)
J^T	transpose Jacobian matrix ($\tau=J^T f_s$)
Λ	arm forward kinematics matrix ($x=\Lambda\theta$)
K_{fj}	joint torque gain matrix (unitless)
K_v	joint velocity gain matrix (torque/velocity)
K_{fs}	cartesian stiffness matrix (force/position)
K_{fd}	cartesian damping matrix (force/velocity)
S	hybrid control cartesian selection matrix
I	identity matrix

The input to the control system is the desired trajectory, possibly including desired positions, velocities, and forces, \mathbf{x}_d , $\dot{\mathbf{x}}_d$, and \mathbf{f}_d . The outputs of the controller are in the form of joint torque commands, τ_c . The robot arm accepts the joint torque commands and interacts with the environment. The arm/environment plant outputs are the sensor signals, θ , $\dot{\theta}$, and \mathbf{f} . Note that the position and velocity sensors give joint information, since these sensors are located at the joints (or at the actuators). The force sensor output, on the other hand, gives endpoint force information in tool coordinates. Transformation of the joint position and velocity vectors to cartesian coordinates is done by the arm forward kinematics Λ and Jacobian \mathbf{J} , respectively. Texts by Craig [23] and Paul [66] discuss robot kinematics.

In the control implementations presented here in block diagram form, we have carefully avoided using the Jacobian inverse, \mathbf{J}^{-1} , for its computational burden and stability problems at kinematic singularities. For reference, we show a position-control servo loop first, and then discuss five schemes which implement active compliant control:

1. position control (for reference)
2. explicit force control
3. hybrid position/force control
4. stiffness control
5. damping control
6. impedance control




Figure 1.2: Position Control

Position control, shown in Figure 1.2, is the conventional method of robot servo control. In this implementation, the measured joint position is converted to cartesian (tool) coordinates and compared to the desired position trajectory. The gain matrix \mathbf{K}_{fs} sets the endpoint servo stiffness in tool coordinates, while \mathbf{K}_{fj} sets the joint servo stiffnesses. Note that these stiffnesses give the passive joint compliance mentioned above, not to be confused with active stiffness control. Joint damping is added by \mathbf{K}_v . This controller implements the control equation

$$\tau_c = \mathbf{K}_{fj} \mathbf{J}^T \mathbf{K}_{fs} (\mathbf{x}_d - \Lambda \theta) - \mathbf{K}_v \dot{\theta}$$

which is analogous to a simple proportional servo with velocity damping that embodies the control law:

$$f_a = k_p(x_d - x) - k_d \dot{x}$$




Figure 1.3: Explicit Force Control

Explicit force control is the most straightforward compliant control scheme. Essentially, this is a classical endpoint force "servo" with actuator velocity feedback for damping. Figure 1.3 shows this scheme implemented as a simple proportional control loop, implementing the following control equation:

$$\tau_c = \mathbf{K}_{fj} \mathbf{J}^T (\mathbf{f}_d - \mathbf{f}_s) - \mathbf{K}_v \dot{\boldsymbol{\theta}}$$

which is the multi-axis equivalent to the simple proportional force control law with velocity damping:

$$f_a = k_p(f_d - f_s) - k_d \dot{x}$$

The reference input is the desired force "trajectory" vector, \mathbf{f}_d . The "error" vector represents the force errors in tool coordinates. The (unitless) gain matrix \mathbf{K}_{fj} is usually chosen to be diagonal, and it is sometimes replaced by $(\mathbf{K}_{fj} + s\mathbf{K}_{fjd})$ or $(\mathbf{K}_{fj} + \frac{1}{s}\mathbf{K}_{fji})$ to give PD or PI force control. The inner loop is simply joint velocity feedback for damping. Note that the joint position signals are not used, and we have therefore no control over absolute position. It is for this reason that this scheme is of little practical use, however its simplicity makes it popular as a research tool.

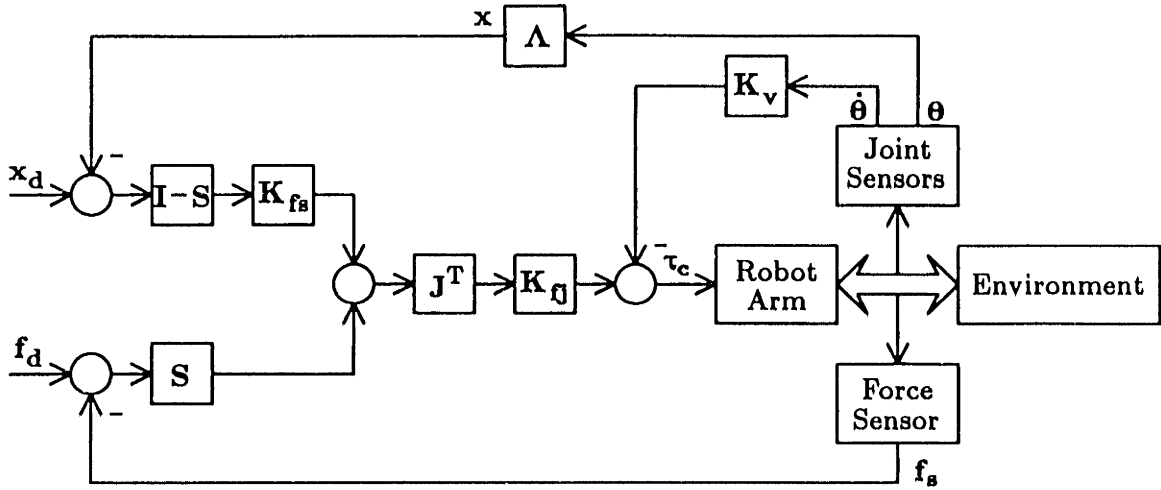


Figure 1.4: Hybrid Position/Force Control

Hybrid position/force control [71], shown in Figure 1.4, combines conventional position control and explicit force control. The natural constraints imposed by the environment dictate that only force control, and not position control, can be achieved in the contact directions. Conversely, only position control can be achieved in the directions of no contact [58]. In hybrid control, we can choose the force-control directions by setting the appropriate elements in a diagonal selection matrix S to unity, leaving zeros in the other diagonal positions to specify position-control directions. In this manner, the hybrid control scheme implements both force control and position control simultaneously, in orthogonal directions along the tool coordinate axes. For more complex tasks, which may not be aligned with the tool coordinate frame, we can choose a more general form of S and then calculate an orthogonal matrix S^{ORTHO} (in the tool coordinates) to use in place of S [58]. The hybrid controller shown implements the following control equation:

$$\tau_c = K_{fj} J^T [S(f_d - f_s) + K_{fs}(I - S)(x_d - \Lambda\theta)] - K_v \dot{\theta}$$

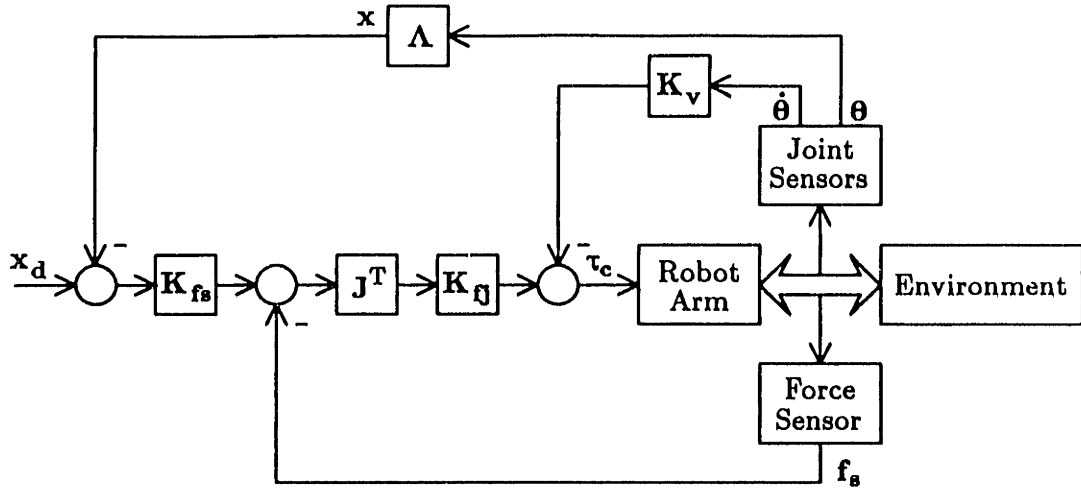


Figure 1.5: Stiffness Control

Stiffness control [77] implements a six-axis active "spring" in tool coordinates. The sensed forces are converted to offsets for the commanded position trajectory. The stiffness matrix K_{fs} in Figure 1.5 therefore sets the effective spring rates in cartesian tool coordinates. The various implementations of stiffness control include converting the sensed forces into offsets for an inner joint position control loop, or converting the position errors into a desired force setpoint. Here we have taken the latter approach, implementing this control equation:

$$\tau_c = K_{fj} J^T [K_{fs}(x_d - \Lambda\theta) - f_s] - K_v \dot{\theta}$$

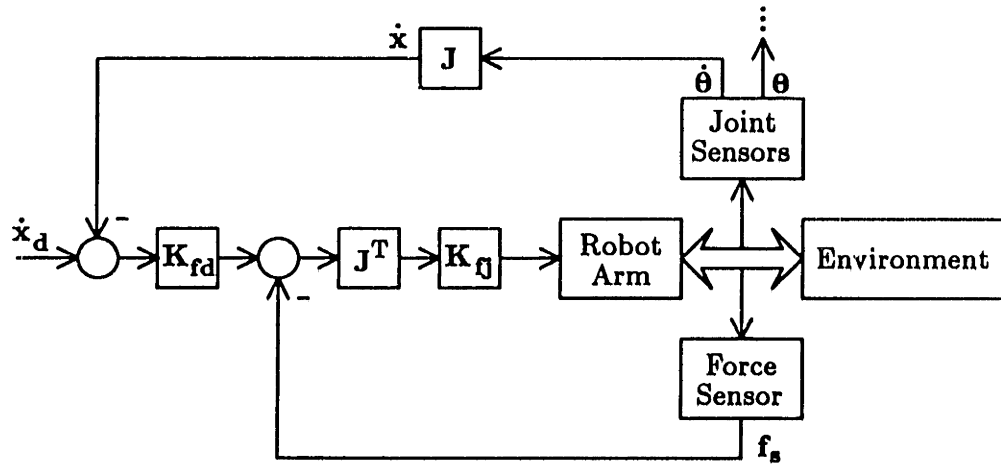


Figure 1.6: Damping Control

Damping control [97] implements a six-axis active "damper" in tool coordinates. The sensed forces are converted to offsets from the commanded velocity trajectory, or equivalently, the velocity error is converted to a force setpoint.

$$\tau_c = K_{fj} J^T [K_{fd}(\dot{x}_d - J\dot{\theta}) - f_s]$$

Figure 1.6 shows this form of damping control. The damping gain matrix K_{fd} sets the effective damping rates in cartesian tool coordinates. Choosing K_{fd} diagonal, this scheme becomes a generalized damper. This force control implementation has a behavior which is popular for many planning algorithms.

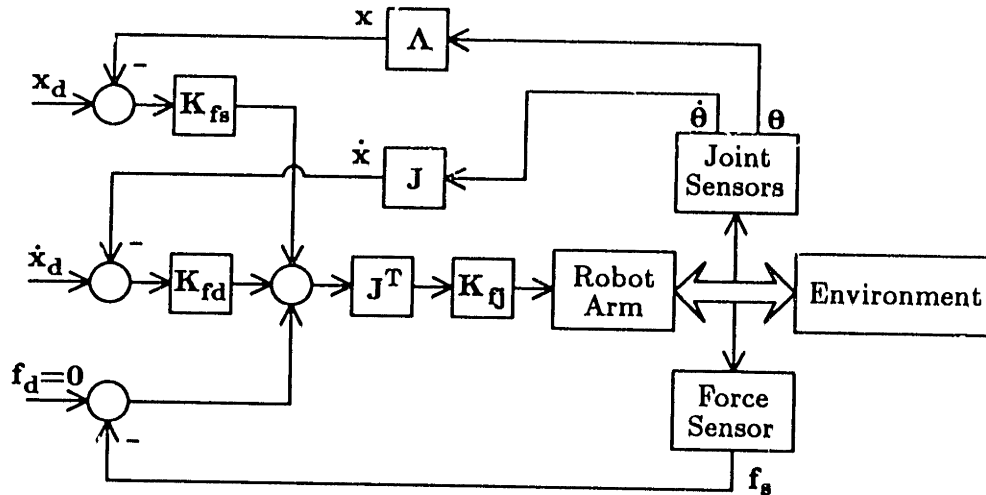


Figure 1.7: Impedance Control

Impedance control [43] is an approach to compliant manipulation based upon physical systems theory. A common implementation of impedance control achieves compliant motion by combining stiffness control and damping control. Errors in position, velocity, and force are used to determine the joint torque commands.

$$\tau_c = K_{fj} J^T [K_{fs}(x_d - \Lambda\theta) + K_{fd}(\dot{x}_d - J\dot{\theta}) + (f_d - f_s)]$$

Figure 1.7 shows this form of this scheme where all of the errors are formed in tool coordinates, and the force setpoint is zero. The three gain matrices K_{fs} , K_{fd} , and K_{fj} set the effective stiffness, damping, and inertia of the manipulator. These map neatly into stiffness, stability, and bandwidth requirements [48].

1.1.4 Assembly Strategy and Planning

Mechanical assembly and other force control applications are tasks where the robot actions are partially constrained by the environment due to contact with one or more surfaces. If we consider a manipulation task as a sequence of subtasks defined by the various contact situations which are planned, we can analyze each subtask in terms of the **natural constraints** imposed upon the robot by the contact geometry and the **artificial constraints** necessary to execute the subtask. Mason [58] has formalized a framework for analyzing the natural constraints and choosing artificial constraints appropriate for compliant motion strategies.

An assembly strategy includes a set of artificial constraints and/or compliant-control gain

matrices as well as methods for recognizing changes in contact geometry which should trigger transitions to subsequent subtasks. Programming assembly strategies for complex assembly tasks is quite difficult, and automatic planning systems are being developed [31, 54, 55].

Many researchers have studied variations of the peg-in-hole assembly task. Whitney [99] analyzed the mechanics of this insertion process and describes how the RCC strategy works [96]. Caine [19] analyzed the square peg-in-hole problem in two and three dimensions. He shows how to find a force and torque to apply which guarantee successful assembly for the given geometry and friction. Gustavson [40] derived the mechanics of three-dimensional chamfered assembly, and defined desirable chamfer characteristics. Rivero and Seering [72] empirically studied compliant peg-in-hole insertion. They show typical quasi-static insertion-force time histories, and give guidelines for setting robot wrist stiffnesses to reduce insertion forces.

1.1.5 Trials Using Force Control

Force control applications to date have included edge following, assembly, and grinding. Edge following is perhaps the simplest force control task, a sub-goal of which is to make initial stable contact. Some authors have shown experimental results from which we can infer bandwidth and other performance specifications.

Maples and Becker [57] suggest that the inner control loop be a fast, joint-based, position-control servo. The outer, endpoint force-control loop can be slower. They implemented cartesian stiffness control with 500 Hz digital servo rates and achieved 140 ms rise time in Z-axis force steps. The Z-axis is in the stiffest direction of the AdeptOne arm.

Khatib and Burdick [51] describe an operational space compliant control formulation. Their implementation on a PUMA 560 robot showed 20 ms rise times (very fast). They used an impact transition control stage (high joint velocity damping) to absorb the energy of impact. The authors credit the high performance achieved to "nearly-perfect" dynamic decoupling by the operational space formulation and a good six degree-of-freedom PUMA model with very accurate identification of the arm's dynamic parameters [5]. Starr [87] implemented accommodation control [97] on a PUMA 560 robot for edge following, achieving 1.7 Hz force control bandwidth.

Wlassich [101] designed a nonlinear impedance control algorithm for a two-link robot, using position, velocity, and force feedback. He achieved stable control of the contact with a stiff surface without switching control schemes upon contact.

Many authors have applied force control to deburring castings or grinding weld beads. Plank and Hirzinger [69] are among the early applications of force feedback for grinding. They used a six-axis force sensor [80] on an ASEA robot to improve deburring performance. Asada and Goldfine

[8] analyzed the grinding process and found that using a wrist with high stiffness normal and low stiffness tangent to the work surface reduced robot vibrations during grinding, improving accuracy and surface finish. Kazerooni, Bausch, and Kramer [47] implemented the high and low impedances necessary in the normal and tangential directions, respectively, using Kazerooni's impedance control design methodology [49]. They used a servo-positioning table to move the workpiece, and they designed a compliant end effector to hold the tool. Stepien [88] achieved stable force control for deburring in spite of a flexible drive transmission (harmonic drive) on the ASEA robot.

1.1.6 Robot Hardware Design

Force control strategies can be applied to all types of robot arms. Tests have been performed on serial-link arms such as the PUMA robot [19, 51, 87]. Many other types of robots have been used for force control development, including direct-drive arms, SCARA-type robots, cartesian machines, flexible arms, and dextrous hands.

Youcef-Toumi and Asada [103] showed how to design a direct-drive arm with decoupled and configuration-invariant inertia tensors. They used graphite composites to raise the lowest natural frequency up to 68 Hz and to achieve 5 g acceleration [6].

Cartesian robot arms such as the IBM 7565 have been popular for assembly research [7, 37]. The MIT Precision Assembly Robot (MITPAR) [14, 93] has the advantage that the axes are decoupled and the arm is very stiff, with its first bending mode at 33 Hz.

The SCARA-type arms are inherently stiffest in the vertical direction, where Nevins and Whitney have shown that most assembly operations take place [61]. Maples and Becker's force control experiments were done in the vertical direction with the AdeptOne robot [57].

Research on flexible robot arms has gained popularity in recent years [16, 42, 67], and has drawn upon efforts in control of other flexible structures [41, 52]. Cannon's work is aimed at understanding how to control flexible systems [20, 21, 44]. All systems are inherently flexible, and sometimes we wish to control them with bandwidth above the first mode of vibration. The growing interest in control of flexible links should not be misinterpreted as a suggestion that robot arms should be deliberately made flexible. On the contrary, we should study flexibility to learn how to control it, because even stiff arms have flexible modes.

Reducing vibrations in flexible structures can be very difficult. Much work has been done in vibration damping [10, 18, 24, 70]. Two successful vibration control tricks are: first, not to put in energy at the resonances [60]; and second, to put in energy as fast as you like, but then remove it in the very same way, so as to cancel all but the rigid-body effects [86].

New arms are being developed which may perform significantly better in force control.

Seering [82] argues that we should not design anthropomorphic devices for factory automation. If we are trying to put robots where people are insufficient, then robots should be thought of as machines instead of as automatons. Townsend [91] analyzed transmission dynamics and he shows why robot drive system reducers should be placed near the joints rather than at the actuators to minimize compliant effects of the transmission. Pasch and Seering [64] give guidelines for choosing actuators and transmissions for servo-controlled machinery. They point out that maximizing acceleration does not necessarily minimize move time.

Whole-arm manipulation [79] (WAM) is an extension of force control in which the robot system is designed to interact stably with the environment using all of its available link surfaces instead of just at the endpoint. Townsend [78, 91] implemented this concept in the design of a four-degree-of-freedom WAM arm. The rubber-covered links are long and slender to allow them to reach between and around objects in the workspace and to enhance the size objects that can be grasped between consecutive links. To improve the bandwidth of forces transmitted bilaterally through the link structure, the transmission ratios are limited so that the joints are backdrivable in both the inertia and friction senses. By using accurate joint-torque sensors rather than a force/torque wrist sensor, they hope to, with appropriately designed kinematics, compute the magnitude, line of action, and contact location of forces along extended regions of the robots links.

Active end effectors have gained some popularity recently. Sharon and Hardt [83, 84] developed a small, high-performance robot to mount at the tip of a larger, slower robot arm. The *macrolmicro manipulator* system achieves the small manipulator's bandwidth for tip motions within the range of micro-manipulator motion (0.1 inch). Kazerooni [50] designed a two degree-of-freedom mini-manipulator end-effector with active compliance (Active RCC). He implemented closed-loop impedance control and achieved 15 Hz bandwidth within the 0.3 inch square workspace. Tilley [90] built a fast wrist for Cannon's flexible (1Hz) arm [21]. The wrist regulates the contact force long before vibrations in the arm itself have settled out.

1.1.7 Sensor Development

Force sensors generally use strain gauges and analog instrumentation to provide voltage signals proportional to the forces and torques acting upon the robot arm [12, 30, 53, 80]. Many other transducer technologies have also been used in the design of robot force-sensing devices, including inductive displacement transducers [68], magnetoelastic materials [95], and photodiodes [25].

The transducers are mounted to a compliant element within the robot structure. Since the transducers are actually measuring position changes in the compliant element, the sensitivity of the sensor is inversely proportional to its stiffness. Since both high stiffness, and high sensitivity are desirable, this tradeoff must be carefully considered when choosing a force sensor.

There are four locations in robot systems where sensors for force feedback are commonly placed:

1. at the actuators
2. at the joints
3. in the wrist
4. on the gripper fingers

Actuator effort is often sensed for internal feedback, and this signal can also be used to compute joint torques and/or endpoint forces. For example, current in the motor windings may be sensed and then motor and joint torque can be computed with the knowledge of the motor torque constant and the transmission ratio. With accurate modeling and high backdrivability, it is even possible to calculate the position of the point of contact on the arm [78].

A joint torque sensor is usually placed at the output of the joint power transmission [56]. This allows a loop to be closed around the joint power train to compensate for disturbances, cogging, and joint friction [9]. While it is often possible to create high-bandwidth motor current or joint torque servo loops, the joint torque information may not accurately represent the endpoint interaction, which we desire to control.

The wrist-mounted force sensor is by far the most popular in current use, and there are several commercially-available models [12, 53]. The sensor is mounted to the endpoint of the arm, and the end-effector tooling is then attached to the force sensor. Wrist sensors have been built which provide sensed information about only one component or up to all six components of the applied force/torque vector. The more sophisticated sensors perform digital transformation of the many strain gauge outputs into the force and torque components in the tool coordinate system.

Force sensing gripper fingers have been developed which give up to six components of force information for each fingertip. While these can be successfully used in the compliant motion control schemes described below [51], force sensing fingers are often incorporated into the grippers on commercial robots and are then used in the simpler logic branching schemes. Some advanced fingertip sensors have been recently developed which give both the applied force/torque vector, and the contact position/orientation vector [17].

There are many other types of sensors that may enhance a robot's ability to perform useful tasks. These include endpoint relative position sensors [11, 38], tactile array sensors [85], sonar range sensors [32], and many vision algorithms [2, 39, 45].

1.1.8 Modeling and Stability Analysis

Many authors have made efforts to model robot behavior, mostly these have been rigid-body models. A few pieces of research stand out, though.

An and Hollerbach discuss some dynamic and kinematic stability issues of robot force control. The dynamic analysis [3] shows how dealing with stiff environments is difficult with unmodeled dynamics. Some proposed solutions include compliant coverings and low-pass filtering. The example of high-frequency unmodeled dynamics given is arm flexibility in a single-link arm. The kinematic stability analysis [4] shows that Raibert and Craig's hybrid control is stable only for some arm configurations, while Salisbury's stiffness control method is universally stable. The difference is due to use of the Jacobian inverse in this particular hybrid control implementation instead of the Jacobian transpose in the stiffness control law.

Colgate and Hogan [22] present *interaction control* as a new controller design specification in which "coupled stability" is added to the usual nominal stability and performance specifications. A model of the environment dynamics is unnecessary. They claim that robots should be stable when coupled with any passive environment. The most destabilizing passive environment is always a pure mass-spring system, providing $\pm 90^\circ$ of rotation to the robot's driving-point impedance transfer function plotted in the Nyquist plane. Therefore, we should restrict the driving-point impedance of the robot to be positive real ($\pm 90^\circ$ phase). McClamroch and Wang [59] also derived a model of constrained manipulator dynamics and global stability conditions.

Good, Sweet, and Strobel [36] showed that transmission flexibility (harmonic drive) gives rise to the low-frequency resonant modes in the ASEA robot, while link and joint flexibility give rise to higher modes. They developed a one-axis model with three flexible elements (four degrees of freedom) and some nonlinearities. This is perhaps the most sophisticated robot model found in the literature. They conclude that due to drive system compliance and motor nonlinearities, one cannot rely on actuator sensors to give accurate joint information. Therefore, endpoint sensing is needed to improve performance.

Roberts [75] studied the effect of wrist sensor stiffness on active force control performance. He suggests using a soft sensor and compensating for static positioning error actively. He developed a good general model, with actuator, rigid arm, and sensor dynamics, including the environment as a disturbance force. He used integral force control to eliminate steady-state force errors, and a compensator for the position errors that become large due to the soft sensor. The controller, which includes saturation nonlinearities and filtering of the force feedback signal, runs at about 70 Hz.

Townsend and Salisbury [92] studied the effect of two types of friction on force control with integral feedback. They show that Coulomb friction can extend the system stability bounds, while stiction can cause limit cycling.

Whitney [98] discusses sampling rate requirements and the difficulty of contacting stiff environments. Yabuta and Chona [102] analyze force-control stability using a one-axis rigid-body robot model attached to a linear environment stiffness. They consider the effects of static friction (piecewise, either sticking or slipping), environment stiffness, and digital sampling rate on the stability of force control implemented with underlying position, velocity, and torque servos.

The earlier publications of this research project are mentioned here, although most of this work is presented in detail in the remainder of this thesis. Eppinger and Seering [26, 29] first presented a series of dynamic models of robot force control, which included some effects of workpiece dynamics, sensor compliance, and robot structural or transmission stiffness. These models show that the dynamics of the robot first mode may contribute to force-controlled instabilities. Later the effects of the discontinuity at the workpiece contact were included [81]. The actuator bandwidth limitations, filtering, and classical controllers were also investigated [27]. Finally, we discussed some of the various ways to model robot flexibility for force control [28]. We have stressed the importance of demonstrating (with the correct model) the difference between joint position control and endpoint force control, and presented the idea of individual modes being *dynamically colocated* in a noncolocated transfer function.

1.2 Motivation

Over the past 15 years, researchers have conceived and implemented force control algorithms on many robot arms. The resulting robot force control performance has been notoriously slow and/or unstable. For this reason, robots operating under force control have not yet played a significant role in the manufacturing environment. Despite some success in the laboratory, force controlled robots do not have adequate bandwidth for industrial applications.

1.2.1 Some Observations

Examples of observed poor performance are difficult to find in the published literature. Nevertheless, all of the laboratory experiments published had their performance limitations. Industrial robots typically have their lowest frequency modes of vibration in the 5 to 20 Hz range. Some of the resonances observed in the PUMA 560 robot in the "arm-stretched" configuration are: 6.5, 11, 15, and 25 Hz. These low-frequency modes are essentially "unmodeled dynamics" for the PUMA rigid-body motion control system.

Robot dynamics couple with the task dynamics in a very complex way. When the robot makes contact with the environment, the impact excites the resonant modes of the coupled robot/environment system. When the robot loses contact, the separate robot and workpiece systems now respond differently. These impact forces on contact are often blamed for "causing" unstable

force control behavior. However, even when the contact is maintained, these systems do still exhibit instabilities under some conditions.

Two documented examples of force control instability are on the PUMA 560 arm and on the MIT Precision Assembly Robot. Caine [19] used the PUMA to test his precomputed force trajectory for three-dimensional peg insertion. He set up the controller to maintain the desired forces and torques which would guarantee successful assembly. The arm quickly went unstable under force control, and the controller gains had to be turned down so low that the assembly process took one hour under force control with the parts in stable contact. (The assembly should take only one second, considering only the motions required and the robot's acceleration capabilities.)

The MIT Precision Assembly Robot is a very different configuration machine, yet it also displays the same phenomenon. Benjamin [13] implemented closed-loop force control with a simple force sensor at the tip of the robot arm. Stable contact could be maintained with some hand-held objects, but not with others. It seemed that the robot could make stable contact with "soft" workpieces such as a foam block or a bare hand. However, "hard" surfaces, such as a metal block, it appeared could not be contacted without violent limit cycles. This response was recorded on videotape and is included in Whitney's compilation of force control demonstrations [100].

Some authors mention limit cycle behavior as one of the performance problems of force control, and yet no solid explanation of this response is offered in the literature. These observations served as motivation for this research project to understand the dynamics of robot force control.

1.2.2 Causes of Force Control Instability

As in classical servo control, we can generally raise force control gains to increase effective compliance and damping until the larger control bandwidth allows higher-order dynamics to be excited. Since in practice one generally only models the rigid-body dynamics of a robot arm for the purposes of controller design, there remain many forms of unmodeled higher-order dynamics. Some sources of unpredicted instability may be:

- digital sampling rate
- environment stiffness
- actuator dynamics
- sensor filtering
- arm flexibility
- contact discontinuity
- transmission backlash
- friction
- control saturation

Historically, digital sampling rate has been thought to cause force control instabilities.

However, today's computers are fast enough to sample the dynamics involved. Since many researchers have noticed that contacting stiff environments is difficult, and so the workpiece has often been accused of causing the observed force control instabilities. Arm flexibility between the joint actuators and the force sensor is indeed a fundamental performance limitation [27], since this is a form of noncolocated control [34], which will be discussed in detail in the next chapter. Control schemes where the actuators and sensors are physically located on different parts of a flexible structure are termed **noncolocated control** [34]. Generally speaking, however, unmodeled effects (nonlinearities or higher-order dynamics) are believed to cause instabilities in force-controlled systems.

1.2.3 Problem Statement

This research attempts to understand what factors can cause machines to be unstable under closed-loop sensory feedback control. Linear control theory tells us that if open-loop phase drops below 180° while at the same frequency the open-loop amplitude remains above unity gain, then the closed-loop system will be unstable. Machine designers however, do not think in terms of amplitude and phase. They work with elements like joints, links, motors, electronics, and sensors. The control engineer can tell the designer what bandwidth is achievable for the machine he has built, but this research is aimed at telling the designer how to build systems capable of performing with the desired bandwidth. The design issues addressed in this thesis include: actuator characteristics, structural dynamics, sensor placement, and environment interaction. While we consider these issues as they apply specifically to robotics, they also apply directly to machine design in general.

The example we consider is robot force control. For the purpose of this research, the details of the controller implementation are not important. In fact, much of the analysis is done with detailed single-axis robot models and simple linear controllers. The goal of the modeling work is to develop insight. Dynamic systems can be very difficult to understand. So very simple models are studied first, and then higher-order effects are added, generally one at a time. Some laboratory studies are included in this research; they serve to assure that the mathematical models developed do realistically represent actual machine systems.

The problem statement under which this thesis is written is the following:

Study the behavior of robot systems under endpoint force control in order to understand why their performance is limited, and to learn how we can design better-performing robots in particular and machine systems in general.

1.3 Overview of This Thesis

This thesis presents analytical modeling, laboratory experimentation, and discussions of issues in machine and robot performance.

- Chapter Two presents the dynamic models. It begins with a simple rigid-body robot model, then actuator dynamics, various types of robot flexibility, task dynamics, and nonlinearities are considered.
- Chapter Three discusses measuring robot performance in the laboratory. Both frequency response and time response data are shown.
- Chapter Four develops detailed models of the MIT Precision Assembly Robot and compares the model predictions to the actual robot performance.
- Chapter Five is devoted to performance improvements. The factors contributing to bandwidth limitations are explored. Then ideas are presented for improving performance and design guidelines are developed for high-performance machinery. The chapter concludes with an example which is verified in the laboratory.
- Chapter Six reviews the insights gained through this research and makes some recommendations for future work.

Chapter Two

Modeling Machine Performance

This chapter discusses how to model machine systems in order to predict their performance and to understand their behavior. A series of dynamic performance models is developed, attempting to depict the behavior of robot systems. We begin with relatively simple, rigid-body models, and then add higher-order effects one at a time.

2.1 Introduction

Rigid-body models representing six degrees of freedom are quite common in robotics. To control robot motions in cartesian coordinates, we need to derive both the **forward and inverse kinematics** of the manipulator. In fact, to issue commands that compensate for gravity, centrifugal, and coriolis terms, we also need to compute the **rigid-body dynamics** of the arm. The analytical methods used are rather straightforward and are discussed in many books on robot mechanics [23, 66].

On the other hand, including **higher-order arm dynamics** in a six-axis robot model is extremely difficult. One noteworthy effort to do so resulted in a 25th-order model of the NASA Space Shuttle Remote Manipulator System. The state variables include seven movable joints, three orbiter rigid-body degrees of freedom, eight bending modes, two torsional modes, and five additional states needed to represent freeplay in the joints and grip. Development of this model [1] took about 10 man-years of effort, and the simulation includes almost 14,000 lines of FORTRAN code.

This chapter shows how to develop and learn from higher-order models of a single robot degree of freedom. We will not use multi-axis models, although they are essential for gross motion planning and position control, where large motions are executed. The one-axis modeling techniques that we will use can also be applied to multi-axis analysis. However, it is much more insightful to consider the higher-order effects for one robot direction (or one degree of freedom) at a time. Figure 2.1 shows this strategy. After linearizing the complex multi-axis dynamics, we can decouple the response into six rigid-body robot models, one for each cartesian direction at the endpoint. Once we understand the rigid-body behavior, we can then add various higher-order effects.

In force control, the robot executes fine motions intending to control its constrained interaction

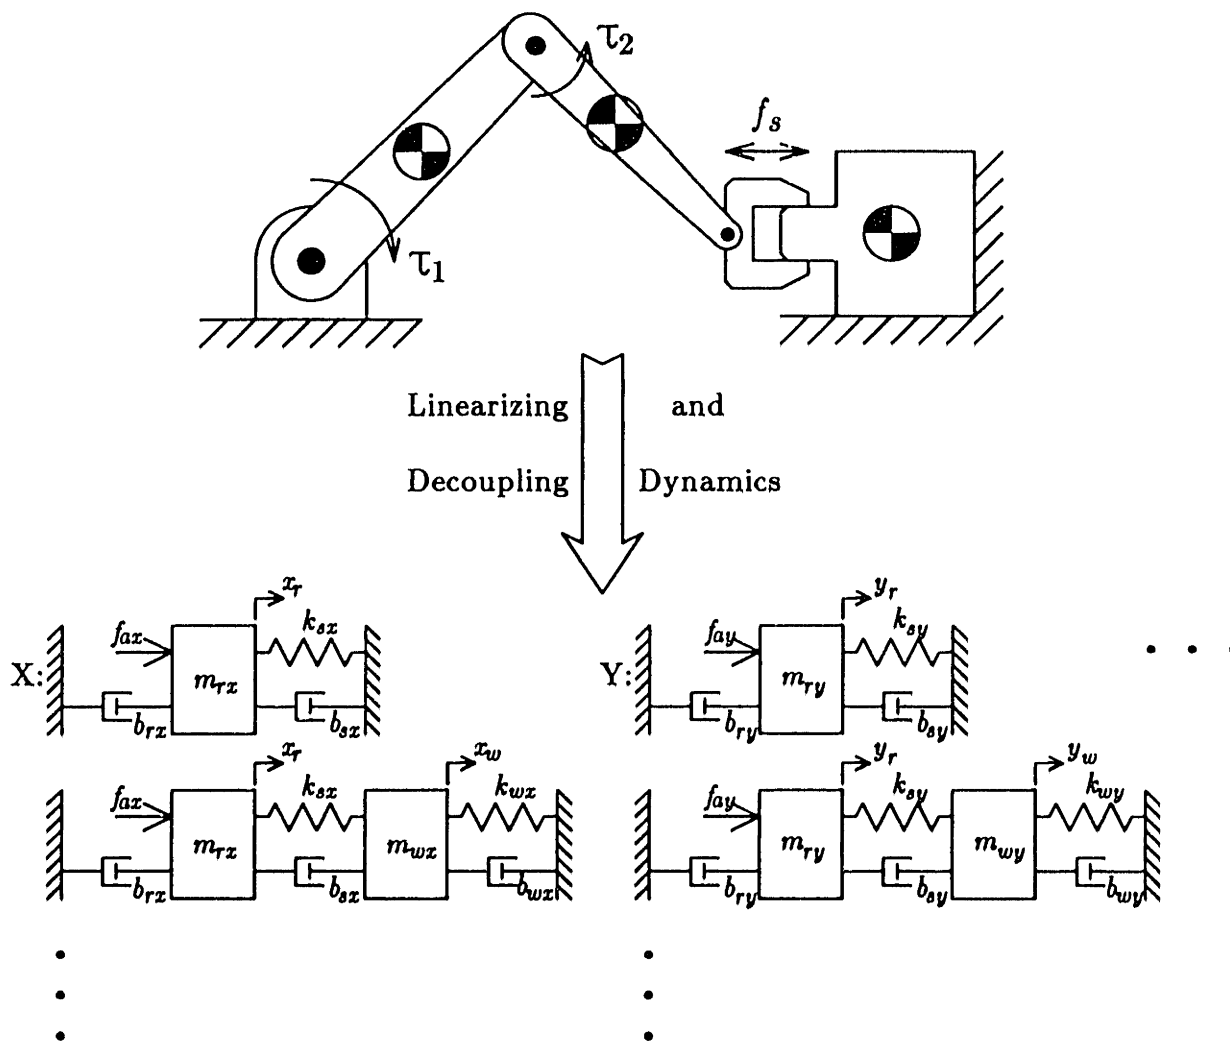


Figure 2.1: Development of One-Axis Robot Models

with the environment. We will learn a great deal by considering the linear, higher-order behavior of a single axis. These models, however, are generally only valid for small displacements, not for large motions. For force control, this is appropriate.

We will begin by discussing a simple one-mass model, representing the rigid-body behavior of one axis. We then consider the actuator dynamics and couple the actuator model to the rigid robot model. Next various forms of flexibility are added to the rigid-body model, including transmission, link, and base compliance. Then the task and grip dynamics are modeled and coupled to the robot models, and finally we add some nonlinear effects.

Most of the models developed in this chapter will show the robot in contact with either a rigid or a dynamic workpiece. This emphasis on constrained dynamics is meant only to help us think in terms of the coupled behavior of the robot in force control. Nevertheless, the very same robot models apply equally well to unconstrained robot motion. So these models can certainly be used to evaluate position-control performance, simply by eliminating the sensor and workpiece.

We will analyze the model's closed-loop performance using the explicit force control implementation (described in Section 1.1.3). We make this seemingly arbitrary choice for two reasons. First, this is the simplest form of force control. Second, and more importantly, all of the other compliant-control implementations use joint position and/or velocity feedback in addition to endpoint force feedback. The force control loop can be seen as (and is often implemented as) an outer loop closed around inner joint servo-control loops. The robot system to model is then the effective impedance of the inner closed-loop plant. From this viewpoint, the importance of understanding the stability of the explicit force control loop is clear. The results presented here also give insight into the behavior of the other compliant control schemes. For a single robot axis, explicit force control becomes quite simple, as shown in Figure 2.2.




Figure 2.2: Single-Axis Explicit Force Control

Most of the models presented are known as **lumped-parameter dynamic models**. This name correctly implies that we will "lump" all of the mass, stiffness, and damping into single-parameter linear elements, like point masses, ideal springs, and perfect viscous dampers. In the real system however, these physical properties may be nonlinear and are distributed throughout the machine. The lumped-parameter approximation is very often a good one. It is ideally suited to our purposes, since our primary modeling goal is to develop useful insight as to how systems behave.

We will make extensive use of two linear system analysis tools, namely the bode plot and the root-locus plot. These are explained in detail in texts on control system design [63]. In this chapter, we will perform the analysis without explicitly giving numerical parameters for the models. The resulting plot shapes are generated by using typical robot parameters and therefore intend to represent typical robot and machine systems, not any specific robot. In Chapter Four, we will develop models specific to the MIT Precision Assembly Robot.

2.2 Rigid-Body Dynamics

This section presents the rigid-body robot model, on which all of the models of higher-order dynamics are based.

2.2.1 The Simplest Robot Model

Let us consider one robot degree of freedom (a single axis, link, or joint). Figure 2.3 shows the effective inertia (total moving mass) of the axis as m_r , and the effective viscous damping to ground is b_r . For force control, we model the force sensor as a linear spring, with stiffness k_s and damping b_s , between the robot and the workpiece. For now, we model the workpiece as rigid, or ground. The actuator effort, f_a , is the system input, and the outputs are the robot position, x_r , and the force across the sensor, $f_s = k_s x_r$.




Figure 2.3: Rigid-Body Robot Model with Rigid Workpiece

2.2.2 Model Behavior

The open-loop transfer function for this system from input force at the actuator to output force across the sensor is

$$\frac{F_s(s)}{F_a(s)} = \frac{k_s}{m_r s^2 + (b_r + b_s)s + k_s}$$

This transfer function gives the system frequency response characteristics shown in the open-loop bode plot, Figure 2.4. The bode plot shows no surprises for this simple system. There is unity gain at zero frequency, and both the amplitude and phase roll off above the natural frequency of $\sqrt{k_s/m_r}$.

To investigate the closed-loop behavior of this system, we need to formulate a control law. A simple force controller implementing explicit force control is of the form

$$f_a = k_f [f_d - f_s]$$

This very simple force control law states that the actuator effort shall be the *force gain* times the *force error*. The force gain, k_f , is unitless, and the force error is the difference between the desired force setpoint, f_d , and the force measured across the sensor, $f_s = k_s x_r$.

Now we can draw a root-locus plot for this system as the force feedback gain k_f varies. Figure 2.5 shows the positions in the s -plane of the closed-loop roots as the gain is varied. The roots start out, for zero gain, at the open-loop poles, and as the gain is increased, the two roots move straight up and down, and the closed-loop system remains stable for all positive gains.




Figure 2.4: Bode Plot for the Rigid-Body Robot Model of Figure 2.3




Figure 2.5: Root-Locus Plot for the Rigid-Body Robot Model of Figure 2.3

2.2.3 Interpretation

An interpretation of this root-locus plot is that force feedback changes the effective stiffness of the system, and this can be seen in the closed-loop characteristic equation.

$$m_r s^2 + (b_r + b_s)s + k_s(1+k_f) = 0$$

That is, for this rigid-body robot model, endpoint force feedback has the same effect as joint position feedback. This is only true to the extent that this simple model characterizes a real robot system. Specifically, if the endpoint force sensor and joint position sensor actually feed back the very same information, then it is adequate to represent the robot axis as a rigid body. If, on the other hand, the robot axis has any flexibility between the actuator and the sensor, then this model is limited in its ability to represent the robot behavior. Since all robots have some flexibility, we will discuss flexible robot models in detail in Section 2.4.

Before considering higher-order dynamics, we should take a step back and note how this model behaves, since we will refer to this rigid-body model as a baseline for comparing the response of more complex systems. This model shows explicit force control to be stable for all gains, which has been repeatedly observed in the laboratory to be not the case. In practice, force control is stable only for low gains. We must dismiss this model as inadequate to represent robot systems for the purpose of controller design. Just which higher-order effects need to be modeled, is not yet clear. We will begin with actuator dynamics.

2.3 Actuator Performance

In this section, we consider the response of the robot actuators. We construct an actuator model and couple that to the robot model. A typical bandwidth-limited actuator is used as an example. In Chapter Four, a more-detailed model is developed.

2.3.1 Perfect and Real Actuators

In the previous section, we did not explicitly include an actuator model. Implicitly however, we assumed that the actuator response is perfect. In reality, of course, the actuator response is limited. All actuators have both bandwidth limitations and effort limitations. In addition, when the actuator is coupled to a load, the load characteristics can affect the actuator performance. The nonlinear saturation effects are discussed in Section 2.7, while here we restrict our analysis to linear models.

First we will formulate a coupled actuator/load model which allows us to derive equations for the decoupled systems separately, and then couple them mathematically. As an example, we will consider an actuator that has a first-order bandwidth limitation, which is typical of many machine systems. We then couple this simple actuator model with the rigid-body robot model load of Section 2.2.

2.3.2 Actuator/Load Coupling

The actuator and its load are dynamically coupled, since the load dynamics physically "load" the actuator, limiting its performance. Some examples of this effect are load inertia, load damping, and back EMF. It is generally necessary to derive the equations governing the coupled system with full knowledge of the structure of both the actuator and the load models. However, if we first derive the algebra of the dynamic coupling, it will then be possible to derive the transfer functions of the actuator and load separately. Of course, to analyze the coupled system, we do still need to consider the fully coupled dynamics, but this can be done numerically, and not symbolically.

Consider the generic actuator and load shown in the block diagram of Figure 2.6. The actuator dynamics are represented by two transfer functions, G_{af} and G_{av} . The outputs of the two actuator transfer functions sum to give the total actuator output force. Their inputs are the actuator command, v_{cmd} (usually volts), and the actuator velocity, v_{act} , which is specified by the load. The load dynamics are similarly represented by two transfer functions, G_{lf} and G_{lv} , specifying the load output force (if this is the variable of interest) and velocity, respectively. The input to both load transfer functions is the actuator force, f_a .

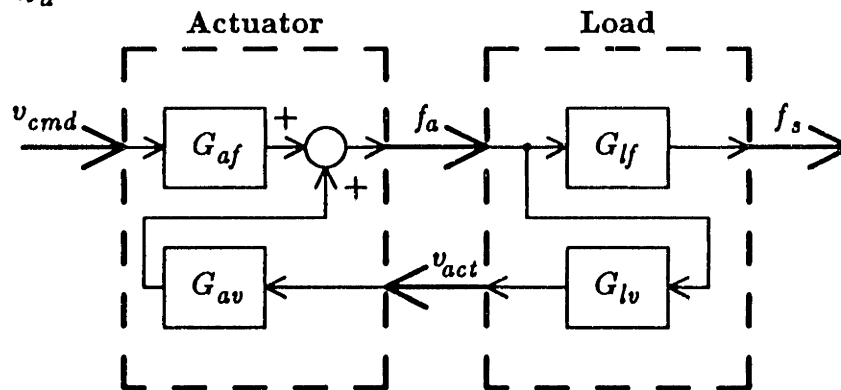


Figure 2.6: Actuator Decoupling Model

The equation which dynamically couples the general actuator to the general load is this:

$$\frac{F_s(s)}{V_{cmd}(s)} = \frac{G_{lf}G_{af}}{1 - G_{av}G_{lv}}$$

This actuator model allows us to derive the actuator and load dynamic equations entirely separately, only combining them numerically for analysis.

2.3.3 A Simple Actuator Example

The actuator model to be presented here represents a generic bandwidth-limited actuator. A sketch of the simple actuator model is shown in Figure 2.7. The input to this model is the actuator velocity command.




Figure 2.7: Simple Actuator Model

The two actuator transfer functions required to couple the simple actuator model to a robot model are:

$$G_{af}(s) = \frac{F_a(s)}{V_{cmd}(s)} = \frac{k_{act}a_{act}}{s + a_{act}}$$

and

$$G_{av}(s) = \frac{F_a(s)}{V_{act}(s)} = \frac{-k_{act}a_{act}}{s + a_{act}}.$$

2.3.4 Actuator-Coupled Rigid-Body Model Performance

The rigid-body robot model of Figure 2.3 depicts a rigid robot in contact with a rigid environment. We found this simplified model to be stable for all force feedback gains. Now we will couple the rigid robot model to the simple actuator model and evaluate its coupled performance, as shown in Figure 2.8. The transfer functions needed to couple the models are:

$$G_{lf}(s) = \frac{F_s(s)}{F_a(s)} = \frac{k_s}{m_r s^2 + (b_r + b_s)s + k_s}$$

and

$$G_{lv}(s) = \frac{V_{act}(s)}{F_a(s)} = \frac{s}{m_r s^2 + (b_r + b_s)s + k_s}$$

The explicit force control equation for the actuator-coupled system is:

$$v_{cmd} = k_f [f_d - f_s]$$

The coupled system's open-loop bode plot is shown in Figure 2.9, and the root-locus plot for varying

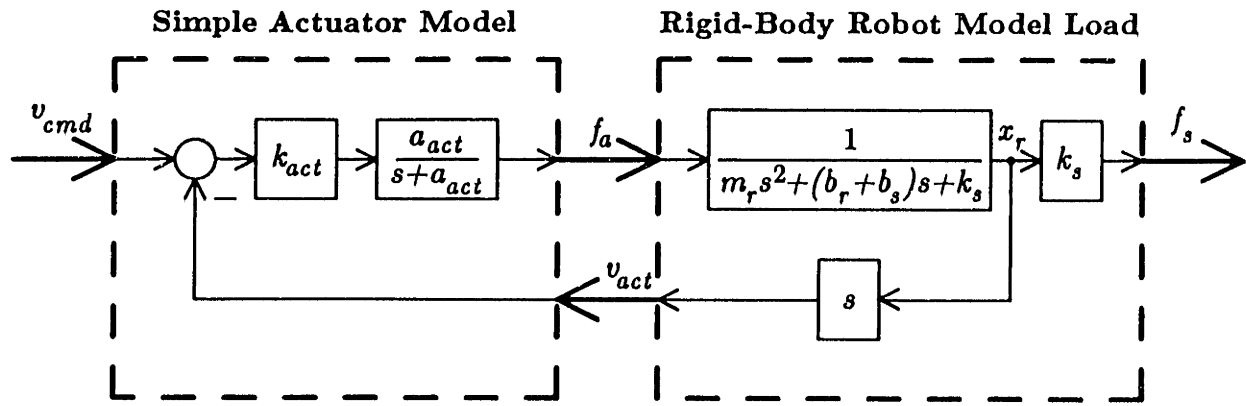


Figure 2.8: Simple Actuator Model Coupled to the Rigid-Body Robot Model

k_f is shown in Figure 2.10. The actuator model has dramatically changed the rigid-body robot model response, and made the closed-loop system now **unstable** for high gains. As expected, this less-than-perfect actuator has a great effect on the stability of the model. We can no longer increase the speed of the robot response without limit by turning up the force gain. This system is limited in bandwidth by the rolloff of the coupled actuator/rigid-robot system.




Figure 2.9: Bode Plot for the Rigid-Body Coupled Model of Figure 2.8




Figure 2.10: Root-Locus Plot for the Rigid-Body Coupled Model of Figure 2.8

2.4 Modeling Flexibility

This section discusses machine flexibility. There are many kinds of flexible elements and just as many ways to model them. Most importantly however, some compliant elements affect machine performance quite dramatically, while other forms of flexibility have more subtle effects.

2.4.1 Types of Machine Flexibility

It was mentioned in Section 2.2 that a rigid-body model is severely limited in its ability to represent real machine systems, which always include some forms of flexibility. Here are a few examples of flexible elements in robot systems:

- transmission compliance:
 - gear tooth bending
 - elasticity of tension elements (cables, belts, and chains)
 - bearing mount compliance
 - lead or ball screw torsion
- structural compliance:
 - link bending or torsion
 - mounting base compliance
 - bending of axis guides or ways
- wrist and contact compliance:
 - force sensor compliance
 - wrist joints servo stiffness
 - gripper flexibility
 - grasp compliance
 - workpiece compliance

Note that some of these flexible elements are physically located between the actuator and sensor, while others are not. This distinction will be important when we evaluate their effects on dynamic performance. It will be very instructive to compare the performance of the various flexible robot models to the rigid-body robot model of Section 2.2. Referring to the above list, the force sensor compliance is already included in our rigid-body robot model, and we will discuss workpiece compliance in the section on task dynamics, Section 2.5. The other forms of compliance will be covered in this section.

2.4.2 Transmission and Link Flexibility

The transmission and link dynamically couple the actuator to the endpoint. In the models discussed so far, we have assumed that the actuator velocity is the same as the endpoint velocity. In fact, with any flexibility in either the transmission or the link, there will be frequencies above which the actuator and endpoint velocities will not only be different, but even completely out of phase.

Since both transmission compliance and link or axis compliance can give rise to this same dynamic effect, we say that the distinction between transmission flexibility and link or axis flexibility is not particularly important.

We can include the compliance of the transmission or link by adding another lumped mass to the rigid-body robot model of Section 2.2, and placing a spring between the two masses, as shown in Figure 2.11. The robot itself is now represented by the two masses, m_{r1} and m_{r2} , and the flexibility between them, given by k_{rf} and b_{rf} . For now, we will use the simple sensor and workpiece model, as before, and here we will analyze the model without coupling to an actuator model.




Figure 2.11: Robot Model Including Transmission or Link Flexibility

If the stiffness of the transmission is low, as in a harmonic drive [36] or with flexible tendons [44], then the robot compliance, k_{rf} would represent this transmission flexibility. If the link itself is flexible [3, 21], then the robot compliance should represent the link flexibility. Determining the parameters for the flexible robot model is not so straightforward as for the rigid-body robot model. One method for finding the necessary values is given in Section 4.4.

The transfer function required to analyze this model is the endpoint (sensed) force transfer function.

$$\frac{F_s(s)}{F_a(s)} = \frac{k_s X_{r2}(s)}{F_a(s)} = \frac{k_s(b_{rf}s + k_{rf})}{[m_{r1}s^2 + (b_r + b_{rf})s + k_{rf}][m_{r2}s^2 + (b_{rf} + b_s)s + k_{rf} + k_s] - [b_{rf}s + k_{rf}]^2}$$

The endpoint force transfer function can be analyzed alone to show the performance of this model with a perfect actuator. The open-loop bode plot, Figure 2.12, and the root-locus plot, Figure 2.13, demonstrate the effect of adding the robot flexibility to the rigid-body model.

The rigid-body robot model bode plot, Figure 2.4, in the absence of actuator dynamics, never dropped past 180° in phase. The flexible robot model bode plot, Figure 2.12, does. This allows the closed-loop system to be unstable for some gains. The root locus plot, Figure 2.13, shows that for high gains, the system has closed-loop roots in the right-half of the s -plane.

We refer to this effect as **noncolocation**, first discussed by Gevarter [34] in the context of




Figure 2.12: Bode Plot for the Flexible Robot Model of Figure 2.11




Figure 2.13: Root Locus Plot for the Flexible Robot Model of Figure 2.11

controlling flexible vehicles. Gevarter showed that if an actuator and sensor are physically located at different points on a flexible structure, then there will be unstable modes in the closed-loop system. He also showed that collocation does not guarantee stable closed-loop control. (Even if an actuator and sensor are physically colocated, there can still be modes in which the actuator output and the sensor signal are out of phase, while for the rigid-body mode, they are in phase. It depends upon what is being measured and where the actuator/sensor pair is placed on the structure.) Gevarter presents examples of noncolocated control systems, including a rail train autopilot with the locomotive at the head of the train and the speedometer in the center.

In our example, the addition of a robot vibrational mode has allowed the actuator and the endpoint to now respond opposite in phase to the control inputs at certain frequencies. The bode plots of the two open-loop position transfer functions demonstrate this. Figure 2.14 shows the phase response of the mass that moves with the actuator $\frac{X_{r1}(s)}{F_a(s)}$ and of the mass that moves with the endpoint $\frac{X_{r2}(s)}{F_a(s)}$. Above the frequency of the robot resonance, the actuator and endpoint are responding considerably out of phase. This means that even though the force signal may be giving accurate information about the endpoint interaction, this results in the controller sending exactly the wrong signals to the actuators.

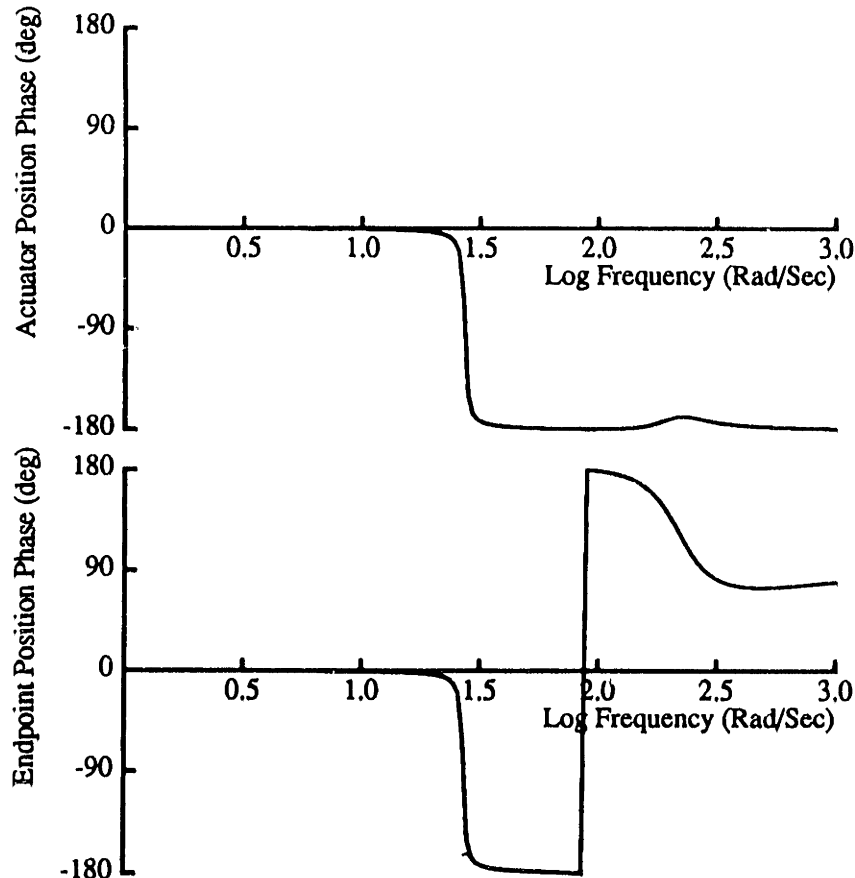


Figure 2.14: Bode Phase Plot Showing Phase Shift Caused by Noncollocation

Robot flexibility is certainly one possible cause of robot force control instability. The robot-flexibility model explains how joint position control can achieve higher bandwidth than endpoint force control. This insight was lacking in the actuator model, which displayed an instability that would be present even for joint position control. Nevertheless, it remains to be seen how other types of flexibility affect machine performance. In the next section, we consider the flexibility of the base structure which supports the manipulator.

2.4.3 Base Dynamics

A robot can be quite a massive machine tool. The base on which it is mounted is never perfectly rigid. In fact, the base compliance can even give rise to the lowest-frequency mode of vibration observed in the robot structure. The simple explanation for this mode is that before the robot was bolted to the floor, while in free space, it had six floating rigid-body modes. After attachment of the robot base to the floor, the rigid-body modes now become vibrational modes. It is quite possible that the effective combined stiffness of the robot base structure and the floor flexibility gives rise to low-frequency modes in which the entire mass of the arm (actuators, links, and sensor) moves against the effective base stiffness. In fact the workpiece too may be moving with the arm and not the base. This is the case with the MIT Precision Assembly Robot, and the lowest-frequency mode is the "base mode", at approximately 12 Hz. Garcia-Reynoso [33] measured the mode shape to be as shown in Figure 2.15.

We will consider two robot models in which we include base dynamics. In the first model, shown in Figure 2.16a, the robot has a compliant base, and the workpiece is still mounted to ground. In the second model, Figure 2.16b, the base also supports the workpiece.

In both base models, m_b represents the moving mass of the base, while k_b and b_b give its stiffness and damping. The actuator effort is applied to both the base and the robot. For now, we use the simple rigid-body robot model of Section 2.2 to represent the arm and sensor. In both of these base models the workpiece is given no dynamics yet, however in the second base model, the dynamic base is also the workpiece support. This configuration best describes the setup of many cartesian robots, including the MIT Precision Assembly Robot and the IBM 7565 arms where the workpiece is generally mounted to the robot base plate.

The equations of motion for the base models are quite simple to derive, as before. However, in formulating the transfer functions, we have a bit more algebra to do. The actuator velocity is given by this equation:

$$V_{act}(s) = s[X_r(s) - X_b(s)]$$

The transfer functions needed for the first base model are:

$$\frac{F_s(s)}{F_a(s)} = \frac{k_s X_r(s)}{F_a(s)} = \frac{k_s(m_b s^2 + b_b s + k_b)}{[m_b s^2 + (b_b + b_r)s + k_b][m_r s^2 + (b_r + b_s)s + k_s] - b_r^2}$$

The transfer functions needed for the second base model are:




Figure 2.15: Mode Shape of the 12 Hz Resonance

$$\frac{F_s(s)}{F_a(s)} = \frac{k_s X_r(s)}{F_a(s)} = \frac{k_s[(m_b+m_r)s^2+b_b s+k_b]}{[m_b s^2+(b_b+b_r+b_s)s+k_b+k_s][m_r s^2+(b_r+b_s)s+k_s]-[(b_r+b_s)s+k_s]^2}$$

The bode and root locus plots for the base dynamics models are shown in Figure 2.17 through 2.20. These plots demonstrate very clearly how this type of structural flexibility differs from the link flexibility considered in the previous section. The bode phase plots show a sharp lag in phase followed quickly by phase lead. This behavior is explained by the root locus plots which show that a pair of poles is very near a pair of zeros all of which are contributed by the base dynamics. As the base damping, b_b , is increased, the poles and zeros move away from the imaginary axis and have less effect on the response. The base dynamics effects become more important when other lags are introduced by either actuator dynamics or arm flexibility. In Section 4.3.5, we will discuss an application of the base models where this is the case.

In contrast with the link and transmission compliance, the base compliance is less detrimental to the bandwidth of the closed-loop system. The phase dip contributed by the base dynamics would not cause instability if the rigid-body system had sufficient phase margin to keep the total phase above -180° . Also, if the base mode is sufficiently damped, then the pair of poles will lie very close




Figure 2.16: Two Robot Models Including Base Flexibility

to the pair of zeros, and this phase dip is hardly even noticeable. Finally, we call the base mode **dynamically colocated**, since the base resonance does not cause the actuator and sensor to vibrate opposite in phase.

A comment is in order about the difference between the two base models. When the base dynamics are very underdamped and have a resonance near the crossover of the rigid-body model, this mode can contribute significantly to the response of the system. In general, however, when the base dynamics are significant, the configuration of the base B model appears to be more stable for force control. This is because the endpoint force comes from the sensor which is contacting the dynamic base, instead of the ground. The base motion accomplishes some work on behalf of the endpoint force control task.

2.4.4 Better Bending Models

So far, we have used lumped-parameter models used to represent the robot as a flexible system. These models do adequately model some flexible systems, however, other types of compliance may require **bending models**. The models in this section are motivated by the question: "*How well can a serial lumped-mass model represent a beam in bending?*" The model which we have used to represent the transmission flexibility in Section 2.4.2 looks like it would also model the first mode of




Figure 2.17: Bode Plot for the First Base Flexibility Model of Figure 2.16a




Figure 2.18: Root Locus Plots for the First Base Flexibility Model of Figure 2.16a




Figure 2.19: Bode Plot for the Second Base Flexibility Model of Figure 2.16b




Figure 2.20: Root Locus Plots for the Second Base Flexibility Model of Figure 2.16b

a beam in axial (tension/compression) vibration. It just does not appear to model bending with the same fidelity. In search of models to use for real machine systems, we investigate two models of beams in bending: a **distributed-parameter beam model** and a **lumped-parameter beam model**.

In particular, we would like to be able to derive both colocated and noncolocated transfer functions for the beams. Perhaps these models will more closely depict the behavior of the robot structure, whose modes look as if they are comprised of beams in bending.

Figure 2.21 shows the distributed-parameter beam robot model. We consider the cantilever beam shown with a sliding constraint at one end and free at the other end. At the sliding end, we apply the actuator force and measure the actuator motion. At the free end, we measure the tip displacement. For now, we will not include the contact interactions and instead leave the endpoint free.




Figure 2.21: Distributed-Parameter Beam Robot Model

Derivation of equations of motion for such a beam is fairly simple, however, we desire transfer functions. Vaughan [94] has developed a method for the derivation of such transfer functions, although he does not present this particular set of end conditions. Using his procedure, however, we can derive the following relations representing the colocated (actuator) and noncolocated (tip) transfer functions for the sliding-free beam shown.

$$\frac{V_{act}(s)}{F_a(s)} = \frac{-1}{2} \sqrt{2/p} \frac{(S^2+C^2)^2+2S^2+2C^2+1}{(S^2+C^2)^2-4CS-1}$$

$$\frac{V_{tip}(s)}{F_a(s)} = -2 \sqrt{2/p} \frac{C(S^2+C^2+1)}{(S^2+C^2)^2-4CS-1}$$

$$\begin{aligned} \text{where } S &= e^{-\sqrt{T}p} \sin \sqrt{T}p \\ C &= e^{-\sqrt{T}p} \cos \sqrt{T}p \\ p &= \frac{s}{a} \\ a^2 &= \frac{EI}{\rho A} \\ T &= \frac{l^2}{2} \end{aligned}$$

and $\rho, E, I, A,$ and l are the usual beam parameters.

These transfer functions are non-polynomial, and have an infinity of poles and zeros. Nevertheless, we can plot their frequency response in the normal manner, by substituting $s = j\omega$. Figure 2.22 shows these frequency response plots, using sample beam parameters.

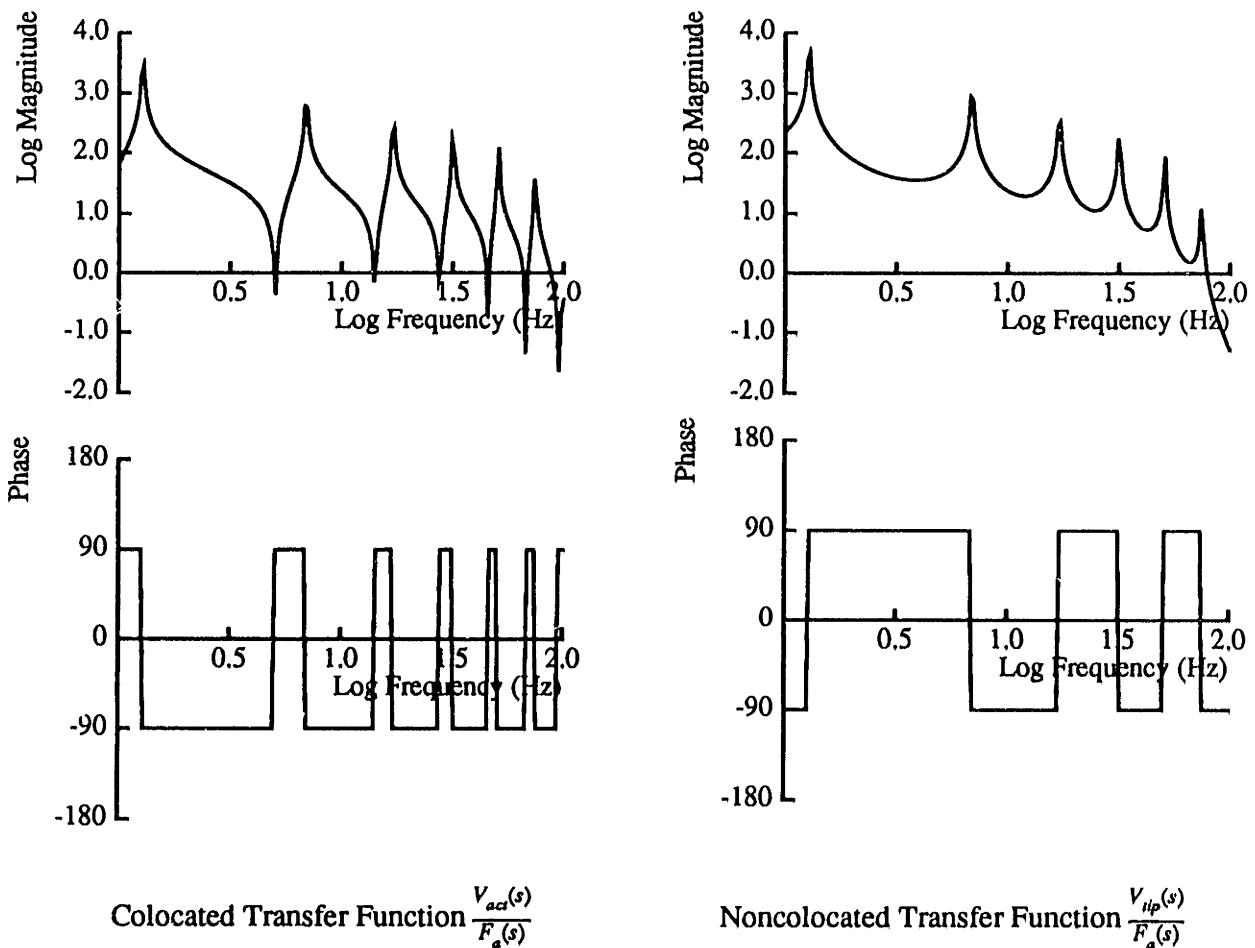


Figure 2.22: Bode Plots for the Distributed-Parameter Beam Model of Figure 2.21

Now, to investigate the suitability of this type of beam model to robot data, we consider the poles and zeros of these transfer functions. While we cannot explicitly solve for the poles and zeros as the roots of the numerators and denominators of the transfer functions, we can still find the poles and zeros as the points in the s -plane at which the transfer functions become infinite or zero,

respectively. In fact, by doing so, we can construct polynomial approximations to these beam transfer functions which are valid up to any frequency we choose. Shown in Figure 2.23 are the pole-zero arrays for the colocated and noncolocated transfer functions of the beam. These polynomial approximations match the true beam transfer functions perfectly through the fourth modes. The colocated beam transfer function has poles and zeros in alternating pairs along the imaginary axis. The noncolocated beam transfer function has pairs of poles along the imaginary axis (in the same locations, since the two transfer functions have the same characteristic equation), but the zeros are spread out symmetrically along the positive and negative real axes! The real zeros appear in these "pairs" so as to contribute no net phase shift. (Without these zeros, the magnitude response would drop off much faster.)




Figure 2.23: Pole-Zero Arrays for the Distributed-Parameter Beam Model of Figure 2.21

The nonminimum phase zeros which appear in the noncolocated transfer function of the distributed-parameter beam robot model are interesting. They were also reported by Cannon and Schmitz to exist in the transfer functions of their very flexible beam structure [21]. The nonminimum-phase zeros are curious because they do not appear in any of the lumped-mass models derived above. Is this because the distributed-parameter beam model uses distributed elements instead of lumped? Or is modeling the bending instead of the tension/compression the significant difference which gives rise to the nonminimum phase zeros?

We hypothesize that the right-half plane zeros come from the bending effects and not the distributed nature of the model. However, to find out for sure, we derive transfer functions for a lumped-parameter bending-beam model shown in Figure 2.24. This model is a limited form of the lumped-parameter Timoshenko beam. Here the left-end mass slides, while the right-end mass is free to rotate and translate, as in the distributed-parameter model discussed above.




Figure 2.24: Lumped-Parameter Beam Model

The equations of motion for this model can be derived and we can solve for the transfer functions from the input force to various outputs. Figure 2.25 shows the pole-zero arrays for two such transfer functions. The collocated transfer function, from the force input to the actuator output, v_{act} , is again minimum phase. The noncollocated transfer function, from the force input to the tip output, v_{tip} , is again nonminimum phase. In fact, this noncollocated transfer function also has a "pair" of real zeros at equal frequencies on the positive and negative real axes.




Figure 2.25: Pole-Zero Arrays for the Lumped-Parameter Beam Model of Figure 2.24

The appearance of the nonminimum phase zeros in the noncollocated transfer functions for both of these bending-beam models is particularly interesting. It leads us to wonder if the nonminimum phase noncollocated transfer functions result in increased tendency for closed-loop roots to cross into the right-half plane. For now, we will close the discussion of the bending models until we see if the experimental data justify their further development. To summarize, in the bending models, the collocated transfer functions showed poles and zeros alternating in order along the imaginary axis.

The noncolocated transfer functions had the same poles, but the zeros were arranged symmetrically about the origin on the positive and negative real axes.

2.5 Task Dynamics

Some researchers that have implemented closed-loop endpoint force control have found that stiff environments are more difficult to make contact with. Others have had trouble with more massive workpieces. The role of environment stiffness (and more generally workpiece dynamics) is to be investigated in this section. So far, we have modeled the environment as a rigid "wall", thought by some to be the worst-case task. Colgate [22] has shown that for some robot systems, a pure mass might even be the most destabilizing environment. In this section, we will model the task as a second-order lumped-mass system, and investigate how these dynamics affect the stability of robot force control.

2.5.1 The Dynamic Workpiece Model

Figure 2.26 shows the rigid-body robot model augmented with a dynamic workpiece. The workpiece has mass m_w , stiffness k_w , and damping b_w . The robot has only its rigid-body dynamics, and no vibrational modes. The force sensor, with stiffness k_s , couples the robot to the workpiece.

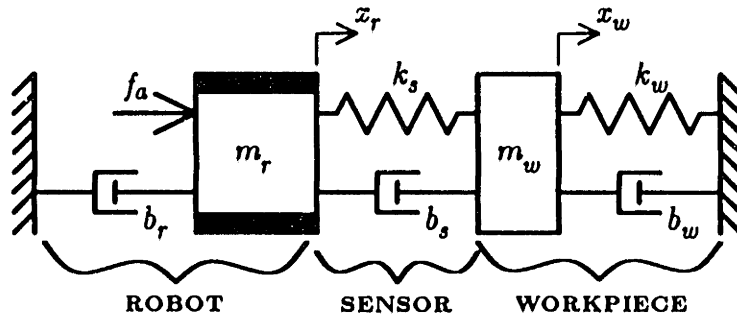


Figure 2.26: Robot Model Including Workpiece Dynamics

The transfer function from input actuator force to output sensed force is:

$$\frac{F_s(s)}{F_a(s)} = \frac{k_s(m_w s^2 + b_w s + k_w)}{[m_r s^2 + (b_r + b_s)s + k_s][m_w s^2 + (b_s + b_w)s + k_s + k_w] - (b_s s + k_s)^2}$$

The bode and root locus plots are shown in Figure 2.27 and 2.28. The second-order workpiece dynamics contribute a pair of poles and a pair of zeros. This gives a phase dip to the bode plot, just like the base dynamics did. Adding more damping to the workpiece moves the pair of poles away from the imaginary axis and toward the pair of zeros, so they almost cancel. Increasing the workpiece stiffness makes the workpiece poles faster, moving them away from the origin in the vertical direction. (Of course, strictly speaking, all of the roots are eigenvalues of the whole system. We find

it instructive, nevertheless to speak in looser terms, by naming the specific poles by the prominent deflecting features of their mode shapes. In this case, the "workpiece poles".)

In terms of colocation, the workpiece dynamics add a **dynamically colocated mode**. That is, just like the base dynamics, the workpiece dynamics do not contribute to the actuator-to-endpoint out-of-phase response. In fact, for the rigid-body robot model shown, the workpiece dynamics have very little effect on stability. In the absence of other dynamics, the model shows that robot force control is stable with all workpieces. Next we will consider what happens when we also include in the model robot flexibility or actuator dynamics.

2.5.2 Adding Robot Dynamics

Figure 2.29 shows the robot now represented by two lumped masses, as in the robot flexibility model of Figure 2.11, and the workpiece has the dynamics described above. In Figure 2.30, the root locus plot shows, as expected, that the workpiece adds a pair of poles and a pair of zeros to the robot flexibility model that we analyzed earlier. As before, the system is unstable for high gains, and the workpiece dynamics actually have little effect unless their frequency is near the crossover of the system without the workpiece. In this special case, illustrated in Figure 2.31 the workpiece can actually help to stabilize the system, since its zeros can attract the roots which would be crossing the imaginary axis, and the system becomes unstable only at a higher frequency, increasing the achievable bandwidth.

2.5.3 Adding Actuator Dynamics

Figure 2.32 shows the rigid-body robot with workpiece model of Figure 2.26 coupled to the simple actuator model of Section 2.3. Recall that the actuator model, when coupled with the rigid-body robot model did indeed show a stability limitation. Now we add the workpiece and Figure 2.33 shows that the workpiece dynamics have again contributed the pair of poles and pair of zeros that we have come to expect from colocated modes. These added dynamics generally have little effect on the system bandwidth. The special case, as before, is when the workpiece dynamics have nearly the same frequency as the crossover of the system without the workpiece mode. Figure 2.34 shows that the workpiece can again help to stabilize the system, when its zeros attract the roots which would have been crossing the imaginary axis, and the system then goes unstable at a slightly higher frequency, increasing the overall bandwidth.

The important conclusion about colocated dynamics is that they have little effect on the overall response of the system, however they do effect the response over a small range of frequency. This phenomenon is quite similar to the contribution of a second-order lead filter. If the characteristics are chosen carefully, the effect can be somewhat helpful.




Figure 2.27: Bode Plots for the Workpiece Model of Figure 2.26




Figure 2.28: Root Locus Plot for the Workpiece Model of Figure 2.26




Figure 2.29: Robot Model Including Workpiece Dynamics and Robot Flexibility




Figure 2.30: Root Locus Plot for the Workpiece Model with Robot Flexibility of Figure 2.29




Figure 2.31: Root Locus Plot for the Workpiece Model Showing a Possibly Stabilizing Case

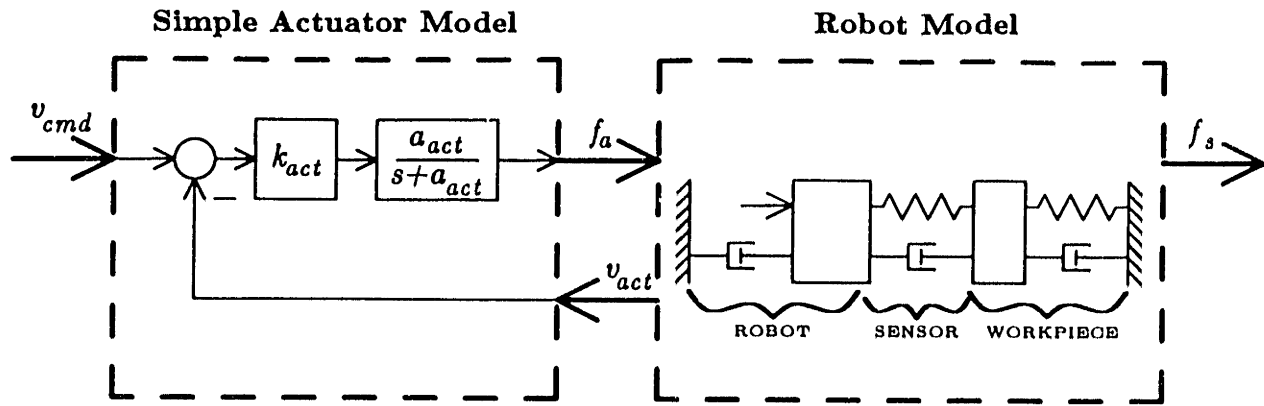


Figure 2.32: Robot Model Including Workpiece and Actuator Dynamics




Figure 2.33: Root Locus Plot for the Coupled Workpiece Model of Figure 2.32




Figure 2.34: Root Locus Plot for the Workpiece Model Showing Another Stabilizing Case

2.6 Grip Compliance

Some researchers have claimed that contact with stiff tasks can be stabilized by using soft gripper pads between the robot endpoint and the workpiece. Many researchers in fact do use compliant contact surfaces either intentionally to help stability, or without knowing that this has any effect on the robot performance. This section will model how this works and shows how this trick can be used to increase system bandwidth.

Figure 2.35 shows a model which includes a rigid robot and dynamic workpiece, as in Figure 2.26. We have now added the grip compliance to the model, which allows us to also include the sensor dynamics. The grip stiffness k_g and damping b_g may be relatively low as compared to the other springs and dampers in the model. The force sensor, which has stiffness k_s , now also has mass m_s , which was previously lumped into the workpiece mass. We do not include the robot flexibility this time. However, we would like to consider a case where without the grip compliance the system does go unstable, and so we will couple this grip model with the simple actuator dynamics. We know that the actuator model combined with the rigid robot and workpiece model, does indeed have a stability limit.




Figure 2.35: Robot Model Including Grip Compliance

First, we will look at the effect of the grip on the uncoupled (perfect-actuator) robot model. The analysis shown in Figures 2.36 and 2.37 is for the model uncoupled from the actuator. Here we still have colocated control so the system is stable even for high gains, and the grip dynamics have had little effect.

The analysis shown in Figures 2.38 and 2.39 is for the model coupled with the simple actuator model of Section 2.3. The lag introduced by the actuator poles give the system severe stability problems. Now the low-gain stability depends upon all the system parameters, particularly the stiffness and damping of the workpiece, grip, and sensor. At high gain, the coupled system will always be unstable. Since it is difficult to draw conclusions without specific model parameters, we will defer most of this discussion to Chapters Four and Five, where we model a specific case more carefully.

Some general observations can be made, however, about the effect of grip compliance. High grip stiffness has relatively little effect on the system performance. Specifically, modeling a stiff grip adds a colocated mode just like the base and workpiece dynamics do. This pair of poles and pair of zeros close together are not particularly interesting since we know they have fairly little effect. Low grip stiffness has a very interesting effect, however. The added pair of poles and zeros are not close together. The very compliant grip allows the robot and the workpiece to act more like their decoupled systems. In Chapter Five, we will further investigate how grip compliance affects force control performance.




Figure 2.36: Bode Plots for the Uncoupled Grip Model of Figure 2.35




Figure 2.37: Root Locus Plot for the Uncoupled Grip Model of Figure 2.35




Figure 2.38: Bode Plots for the Coupled Grip Model of Figure 2.35




Figure 2.39: Root Locus Plot for the Coupled Grip Model of Figure 2.35

2.7 Nonlinearities

All of the dynamic phenomena that we have attempted to model are actually nonlinear effects. We have already learned a great deal about the performance of machines from linear models. Nonlinear models, however are quite necessary for evaluating the behavior in more detail.

There are many possible sources of nonlinearity in machines. Below some of the more important robot nonlinearities are listed.

- saturation
 - sensor output
 - controller command
 - analog signal conditioning
 - actuator effort
- digitization
 - sensor digital conversion
 - controller servo rate
 - output analog conversion
- discontinuities
 - backlash
 - workpiece contact
- friction
 - actuator mechanism
 - transmission elements
 - joint bearing surfaces
 - coulomb friction
 - stiction
 - nonlinear viscous friction
 - drag
- structural elements
 - stiffening and softening springs
 - limit stops
 - load-dependent compliance
- geometry
 - gravity changes with position
 - acceleration coupling
 - harmonics of modeled modes

This section will not cover all of these nonlinear effects. Instead, we will recommend when to use a nonlinear model for design, and present a nonlinear example. The example we consider

investigates the discontinuity at the workpiece contact. Linear models are unable to represent the very nonlinear force control applications where the robot is only sometimes in contact with the environment, and at other times not in contact.

The linear models are quite adequate in general, since force control usually involves fairly little motion and it is the linear dynamic terms that generally dominate the robot behavior for fine motions. For large robot motions (which we have not been considering here), multi-axis kinematic models are certainly required, and the geometry can give rise to some nonlinear terms.

For controller design, saturation should be included in the simulations. Saturation nonlinearities are very easy to include, and can cause very significant limitations in robot performance. Sweet and Good [89] found that motor torque limits were the most significant nonlinearities in their system. We include saturations in the simulation model of the MIT Precision Assembly Robot system in the next section.

Discretization of the analog signals can be a nasty problem; however, if it is done with 12-bit conversions and properly set scales, this effect usually is not important. Simulating this nonlinearity is not advisable, since the numerical integrator will too often come across the small discontinuities. A fixed-timestep integrator will accumulate large errors, and a variable-timestep integrator will proceed very slowly, since it will do too much backtracking to find the discontinuities.

Digital sampling is indeed an important nonlinearity. Authors have analyzed the effect of low servo rates on force control [98, 102], and not surprisingly they have found that you cannot sample very slowly and retain stability. A safe rule of thumb is to set the sample rate five to ten times faster than the closed-loop bandwidth desired. A z-plane analysis is required to determine if slower sampling can be tolerated.

Friction has been analyzed in some detail [92, 102]. Townsend and Salisbury showed that pure coulomb friction can actually help to stabilize force control. It increases the stable region so that higher force gains can sometimes be used. Stiction, on the other hand, where the static friction is higher than the dynamic friction, can be destabilizing, causing sustained limit cycles.

2.7.1 Discontinuity at the Workpiece Contact

The discontinuity at the task interface (workpiece contact) is quite a severe nonlinearity. When the robot loses contact with the environment, its dynamics are changed and the force feedback goes away. The collision which occurs when the contact is again made is thought to cause an impact which excites the arm, causing the system to go unstable. The analysis presented here shows that the discontinuity is not the cause of the observed instability, but rather it is the higher-order dynamics instead.

Figure 2.40 shows a nonlinear version of the rigid-body robot model. In this model, the robot and workpiece can separate when the contact force becomes zero or would be negative. When the linear control law, $f_a = k_f [f_d - f_s]$, is implemented as above, nonlinear simulations have shown time responses resembling Figure 2.41. With this rigid-body robot model, no limit cycle response is displayed, even for very high gains, large initial displacements, and large impact velocities.

Figure 2.42 shows a discontinuous version of the flexible robot model of Figure 2.11. In this model, the robot and workpiece can also separate when the contact force becomes zero. When the linear control law is implemented as above, nonlinear simulations have shown time responses resembling Figure 2.43.

The nonlinear system exhibits limit cycles for some values of gain. For **all** gains higher than the critical gain predicted by the linear analysis, the system will limit cycle. For **some** gains lower than the critical gain predicted by the linear analysis, the system will also limit cycle. For **other** low gains, however, the system is always stable. The limit cycle performance depends upon all the system parameters, as well as the initial conditions, and of course the feedback gain.

The conclusion about this particular nonlinearity is that the linear model is useful for calculating an upper bound on the allowable force feedback gain. However, since Townsend [92] showed that for coulomb friction, the opposite is true, that the linear model predicts a lower bound on the stable gain limit, we have to be somewhat cautious. It suffices to say that it is very difficult to draw general conclusions about nonlinear system behavior, and therefore simulations are required to make reasonable performance predictions when severe nonlinearities are involved.




Figure 2.40: Rigid-Body robot Model with Discontinuity




Figure 2.41: Rigid-Body Discontinuous Model Response




Figure 2.42: Flexible Robot Model with Discontinuity




Figure 2.43: Flexible Discontinuous Model Response

Chapter Three

Measuring Robot Performance

Modeling machine performance is very important for developing intuitive understanding of how systems behave. In addition, we must continually test this intuition and our models by comparing them to actual machine performance. This chapter describes an effort to measure robot response in the laboratory with three goals in mind. First, we require measurements of the parameters unknown in the robot models. Second, we need real data against which to judge the validity of the mathematical models. Third, we wish to observe the real system behavior so that if our models are wrong, we can create new models with the right effects included. (In fact, a great deal of qualitative laboratory data has already been used in the creation of the models, but here we intend to collect quantitative data.)

3.0.1 Deciding What to Measure

Most of the measurements made are frequency response data. Just why this is the case, and when to take other types of data is not obvious. All of the characteristics of a linear model are visible in the frequency domain, including DC gain and phase, resonant and antiresonant frequencies, phase shifts, damping ratios, and even nonminimum phase. Time response data is used as a benchmark for the nonlinear model performance. We can check the various saturation levels by monitoring the appropriate signals.

There are various performance indices which apply to robotics and, of these, closed-loop bandwidth is one which will be considered often in this and in the later chapters. The open-loop phase crossover point determines the bound on achievable closed-loop bandwidth, so particular attention is paid to this specification. These frequency-domain tests will be made by instrumenting both the actuator and the robot itself, measuring inputs and outputs, and finally calculating transfer functions.

The time-domain tests include both position-control and force-control transient response to step inputs, impacts with a workpiece under closed-loop force control, and closed-loop task-disturbance rejection. We pay particular attention to characterizing the nonlinearities with the step response tests, and we look for global stability in the impact and disturbance tests.

While tests were run on more than one robot axis, we will describe and show data for the X axis only (Figure 3.1), since the methods and results are typical.

3.1 Experimental Hardware

The MIT Precision Assembly Robot (MITPAR) was designed and built in the MIT Artificial Intelligence Laboratory by Vaaler and others in the robotics group under Professor Seering [93]. The design specifications included that the robot perform an order of magnitude better than commercially-available robot arms. A gantry-type arm, shown in Figure 3.1 was built with three cartesian axes, labeled X, Y, and Z in the sketch. At the tip of the arm, where most manipulators have a three-axis wrist, this arm has a single, continuous-rotation degree of freedom, labeled Θ in the sketch. Finally, the base holds a two-axis rotary table with the last two degrees of freedom, labeled Φ and Ψ . This unique configuration is ideally-suited to mechanical assembly, where most of the work is done from the Z-direction [61]. Generally, the rotary table is not even required, but can be used to mount the workpiece for access from five of its six sides.

The force sensor is a two-axis transducer which uses four strain gages and analog instrumentation (two Wheatstone half bridge circuits and two AD521 precision instrumentation amplifiers) to provide decoupled voltages corresponding to the X- and Y-direction interaction forces. The compliant element is a .375 inch square aluminum beam 1.5 inch long. A glue joint was added to the sensor as a precaution to limit damage caused by unstable robot response.

3.1.1 Frequency Response Tests

The test setup for the frequency response tests is shown in Figure 3.2. The hardware used includes: a Hewlett-Packard 5423A Structural Dynamics Analyzer; the robot's Automatix amplifiers and Aerotech motors; the robot structure; and a Bruel & Kjaer piezoelectric accelerometer and charge amplifier. The input given to the system is band-limited white noise generated by the analyzer. Various outputs are measured in the tests, and the analyzer compares the frequency components of the input and output signals, then calculates a transfer function for plotting. In this manner, we can measure transfer functions from any input voltage to any other output voltage. With the appropriate transducers, we can then effectively measure input or output forces, accelerations, positions, and more.

The structure of the robot actuator system is shown in Figure 3.3. The detailed motor model, developed in Section 4.1.1 requires parameters for the preamplifier, amplifier, motor, and tachometer. To characterize the Automatix amplifiers and Aerotech motors, we measure the following transfer functions:

- velocity command to current command
- tachometer feedback to current command
- current command to current monitor with motor shaft locked
- current command to current monitor with motor shaft free
- velocity command to tachometer voltage with motor shaft free




Figure 3.1: The MIT Precision Assembly Robot




Figure 3.2: Robot Frequency Response Test Setup

Measuring all of these transfer functions allows us to develop small sections of the proposed motor model at a time. For example, the velocity command to current command transfer function has a DC

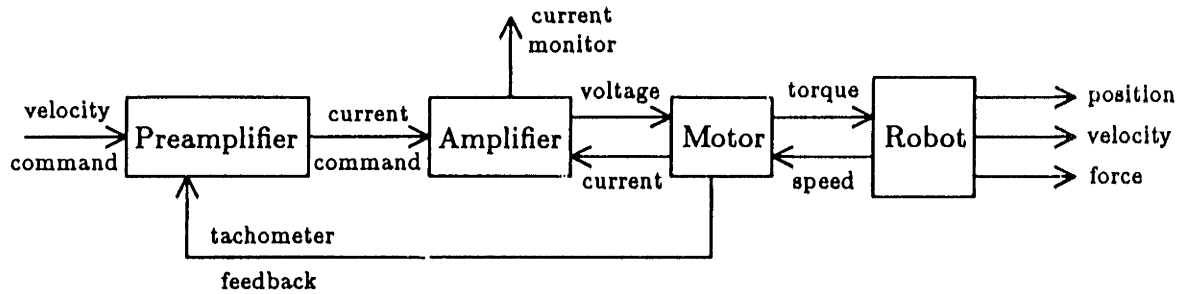


Figure 3.3: Structure of the Robot/Actuator System

gain which is equal to the preamplifier feedforward gain. Its flat amplitude and zero phase confirms that there is no significant filtering in the feedforward section of the preamplifier. Similarly, the measurement using the tachometer feedback input gives the values of the preamplifier feedback input gain and filter frequency. The remaining three tests are used to confirm the motor parameters and the structure of the proposed motor model. In the test using the current command as input, the noise signal was actually given to the velocity command, with the current command monitored as input to the measured transfer function.

To understand the behavior of the robot arm itself, we use as input the velocity command voltage at the amplifier. We measure the following three outputs:

- robot acceleration measured at the motor (colocated)
- robot acceleration measured at the endpoint (noncolocated)
- sensed contact force measured at the tip (noncolocated)

The acceleration measurements are made for the robot not in contact with the environment, and the force transfer function is measured while the robot is in contact with the dynamic workpiece. The colocated and noncolocated transfer functions are needed to help explain the difference in the achievable bandwidths of colocated joint position control and noncolocated endpoint force control. The endpoint force transfer function shows the open-loop plant for the force control system, including base, actuator, robot, and task dynamics.

3.1.2 Transient Response Tests

We measure transient response quite differently, using the robot controller to monitor the necessary signals. The controller includes for each axis an Intel 8031 microprocessor, two AD7574 8-bit A/D converters, one AD7542 12-bit D/A converter, and an encoder tracking circuit. The microprocessors run the axis servo loops at 1500 Hz, and they have up to 12K of memory each, so they can store a modest amount of data. They are supervised by a Digital PDP-11/23 minicomputer, which uses FORTH to send the servo loops, gains, and motion commands down to the microprocessors. The data are transferred from the microprocessors, through the PDP-11, to a Digital VAX 11/750 for storage. The data processing is done in LISP on a Symbolics 3600 machine. Figure

3.4 shows this setup where sensed force and acceleration are being measured for the closed-loop force control tests. In the position-control tests, we take the motor current monitor signal from inside the amplifier and use the two A/D converters to monitor tachometer voltage and motor current.




Figure 3.4: Robot Control Hardware and Experimental Setup

3.1.3 The Test Workpiece

For the force control tests, we require a workpiece to make contact with. In designing this workpiece, we make the following functional requirements:

- The workpiece needs to be stiff enough to be considered rigid by most robot modeling standards, since force control researchers have noted difficulty making stable contact with very stiff environments.
- The workpiece should be quite massive, large enough to bring its lowest natural frequency down within the range of modeling interest, say about 40 Hz.
- The workpiece requires a flat surface against which the robot can make contact while leaving most of the workspace available for other tests and also room to start from a distance to reach top speed before making contact.

The workpiece design is sketched in Figure 3.5. It is a cast iron angle jig, 1 inch thick, with milled surfaces. Its support is a steel beam, .5 inch thick by 1.5 inches high, bolted at both ends to the robot frame. The support effective stiffness (including 10% twist and 90% bending) is measured to be 412,700 N/m. The workpiece itself has an effective mass (including some of the beam mass) of 6.6 kg giving its lowest measured natural frequency of 39.8 Hz.




Figure 3.5: The Test Workpiece Mounted to the Robot Frame

3.2 Frequency Response Tests

The goal of the frequency response testing is to measure some transfer functions to compare with the models, and to determine experimentally the system bandwidth. First the actuator response is measured to characterize the amplifier and the motor. Then the robot is tested in both colocated and noncolocated configurations. We attempt to predict from these tests the closed-loop system bandwidth that can be achieved.

3.2.1 Actuator Transfer Functions

The motor and amplifier are tested extensively to determine the parameters for the motor model described in Section 4.1.1 and to verify that the structure of the model is correct.

The preamplifier is tested in two parts. First, the velocity command to current command transfer function is measured. Figure 3.6 shows the preamp response to be flat up to 1600 Hz. The DC gain is 4.1 volts/volt. Second, the tachometer feedback to current command transfer function is measured. Figure 3.7 shows this preamp response to be filtered with a first-order cutoff of 150 Hz. The DC gain is -3.9 volts/volt. (This is an inverting input.)

The power amplifier stage is more difficult to test. It is a pulse-width modulated (PWM) power amplifier with current feedback, so it looks like a transconductance amplifier (current amplifier). The current error essentially determines the duty cycle of the voltage output to the motor. The motor current is sensed and the measured current signal is subtracted from the current command to yield the current error. Since the motor and power amplifier are coupled in this manner, we must test the two together. We cannot test the power amplifier separately, since it requires the motor inductance for stable operation.

We can however, test the motor without its inertia, damping, and back-EMF. We do this by locking the motor shaft so there is no rotation whatsoever. Then we test the power amplifier and motor electrical dynamics only. The noise from the analyzer cannot be given directly to the current command, so it is applied to the velocity command, and then the current command is measured as the input to the transfer function, while the current feedback signal is measured as the transfer function output. Figure 3.8 shows the transfer function to have flat DC response with roughly unity gain at low frequency, and first-order rolloff at about 2000 Hz. This good current-command-following ability implies high gain in the inner current loop of the power amplifier.

Now, to include the motor mechanical dynamics without the robot dynamics, we simply unlock the motor shaft but do not couple it to the robot transmission. The transfer function of interest measures the tachometer voltage as output and this compared to the velocity command as input gives the amplifier/motor unloaded bandwidth. Figure 3.9 shows this transfer function to have a DC gain of about 0.85 and a rolloff (which may be first order) at about 46 Hz.

3.2.2 Robot System Transfer Functions

When we attach the motor to the robot, we can test the system as a whole. We would rather, for the purpose of model verification, test the robot separately. However, with the motor used as the system input, its dynamics are included. Actually, it might be possible to excite the robot with another input, such as an impact from a force hammer on the motor shaft. This technique is used extensively in modal analysis, and was used by Garcia-Reynoso to take data to build his dynamic




Figure 3.6: Preamplifier Velocity Command Transfer Function




Figure 3.7: Preamplifier Tachometer Feedback Transfer Function




Figure 3.8: Amplifier/Motor Current Transfer Function




Figure 3.9: Amplifier/Motor Velocity Transfer Function

model of the robot structure [33]. For our tests, we require a great deal of energy to be put into the system, and so we use the motors as actuators. We will send noise in to the preamplifier at the velocity command input and measure three outputs

- motor acceleration (colocated transfer function)
- tip acceleration (noncolocated transfer function)
- contact force (noncolocated transfer function)

The acceleration measurements are made with the robot unconstrained, that is, not in contact with the workpiece. These two transfer functions measure the colocated response and the noncolocated response. The system is the same for both tests, only the position of the accelerometer is changed. In the motor acceleration measurement, the accelerometer was placed on a bracket attached to the ball screw at the motor coupling. The data are scaled to units of the linear (translational) acceleration of the axis. The endpoint acceleration measurement was made with the accelerometer placed at the tip of the robot, near the force sensor. These transfer functions, shown in Figures 3.10 and 3.11 represent the open-loop arm response for the robot in free space, unattached from the workpiece. The robot models aim to match this dynamic data.

The contact force transfer function is measured with the endpoint in contact with the workpiece. Actually, the endpoint must be clamped to the workpiece to maintain contact since the noise signal has no offset to give a bias force. The force measurement is the normal output of the force sensor strain gage amplifier circuit. The force transfer function measurement is shown in Figure 3.12. This measurement represents the open-loop transfer function for the attached robot/workpiece system model to match.

3.2.3 Comments on the Measured Frequency Response

It is interesting to note a few features of the robot system frequency response. The open-loop colocated and noncolocated transfer functions differ greatly. Both measurements were made on identical systems; only the accelerometer placement differed. Therefore, these transfer functions have the same characteristic equation, and different numerator dynamics. The noncolocated transfer function shows system resonances at 12, 38, 44, 67, and 77 Hz. The colocated transfer function does not show these resonant frequencies as clearly.

More importantly, the two acceleration transfer functions display the open-loop phase crossover frequencies for colocated (joint) and noncolocated (endpoint) position control. We cannot infer the force control bandwidth from these measurements because they had the robot unattached from the environment. Note that the measurements are acceleration transfer functions. The position transfer functions would be shifted by 180° in phase. So we should measure phase margin for these plots from 0° , instead of -180° . The colocated transfer function shows that the phase drops below 0° at 50 Hz, while the noncolocated transfer function dips below 0° from 11 to 13 Hz then crosses




Figure 3.10: Measured Robot Collocated Acceleration Transfer Function




Figure 3.11: Measured Robot Noncollocated Acceleration Transfer Function




Figure 3.12: Measured Robot Noncolocated Force Transfer Function

finally at 28 Hz. We can see why the system with colocated feedback would be able to achieve greater closed-loop bandwidth. The noncolocated system shows two particular problems for endpoint position control. At 12 Hz, the system loses its phase margin due to the phase dip from the colocated base mode. Even in the absence of the 12 Hz base mode, at 28 Hz the system would drop below 0° anyway due to the first noncolocated mode, which is primarily the bending of the X-axis carriage and the Z-axis boom. Colocated control is simply more stable. Even though the colocated system has all of the same poles, their phase effects are largely cancelled, and the colocated system has less phase lag.

The force transfer function, which also measures a noncolocated situation, but now with the robot attached to the workpiece, is equally interesting. It also shows the system resonances much less than the unattached noncolocated acceleration transfer function. This time, we cannot attribute this to changed numerator dynamics. The poles have actually moved because the robot is now coupled to the environment. Since the workpiece is quite stiff, it is reasonable to assume that some of the lower resonances have moved up in frequency. The base mode is not likely to have changed much though.

We hope to be able to match the dynamics of this transfer function with a detailed robot model capable of predicting closed-loop performance and useful for evaluating proposed changes to the robot design.

3.3 Time Response Tests

This section discusses the time-domain measurements required to compare with the nonlinear simulations of the robot behavior. Two types of time response data are taken: position control step response, and force control contact response. The position control data are mainly used to verify the nonlinearities in the actuator model, while in the force control data, we are concerned primarily with the higher-order dynamics and the workpiece contact discontinuity.

3.3.1 Position-Control Response

PD joint position control is implemented in the microprocessors by subtracting the measured encoder position from the desired position setpoint, multiplying by the proportional gain, and adding the product of the derivative gain and the joint velocity (calculated by differences). In these tests, the microprocessor has been programmed to save the position, velocity, current, and output velocity command values in memory. Figure 3.13 shows the step response data for a typical move of 36,000 encoder counts (9 inches).

The encoder position shows some overshoot, which can be tuned out with the proper choice of gains. The velocity data show ramping up to constant speed. The current and microprocessor output curves show saturations which must be included in the actuator model. When the setpoint becomes close, the controller tells the system to slow, reverse, and then finally stop. These data display clearly only the rigid-body response, so they will be useful only to confirm the actuator/rigid-body model, and the force data will be necessary to confirm the higher-order models.

3.3.2 Force Control Transient Response

The force control loop is also implemented in the microprocessor, running at 1500 Hz. The sensed force (digitized, 8 bits only) is subtracted from the desired force setpoint, and this force error is multiplied by the force gain to give the velocity command output to the amplifier. The microprocessor is also programmed to save the sensed force, and the encoder position data in memory for later analysis.

The set of data presented in Figure 3.14 shows the transient response for the robot making initial contact with the workpiece. The robot starts some distance away and reaches constant velocity. This part of the data is not shown. The data taking begins just before contact is made, and shown here are both the sensed force and encoder position data. A series of tests were run, varying the force setpoint and the force gain. Figure 3.14 shows a sampling of this data. The top two plots show the encoder position data and the sensed force data for $k_f=5$ with $f_d=30$, which gives stable response. The bottom two plots show the encoder position data and the sensed force data for $k_f=10$ with $f_d=15$, which gives unstable response. Figure 3.15 summarizes the stability conditions.

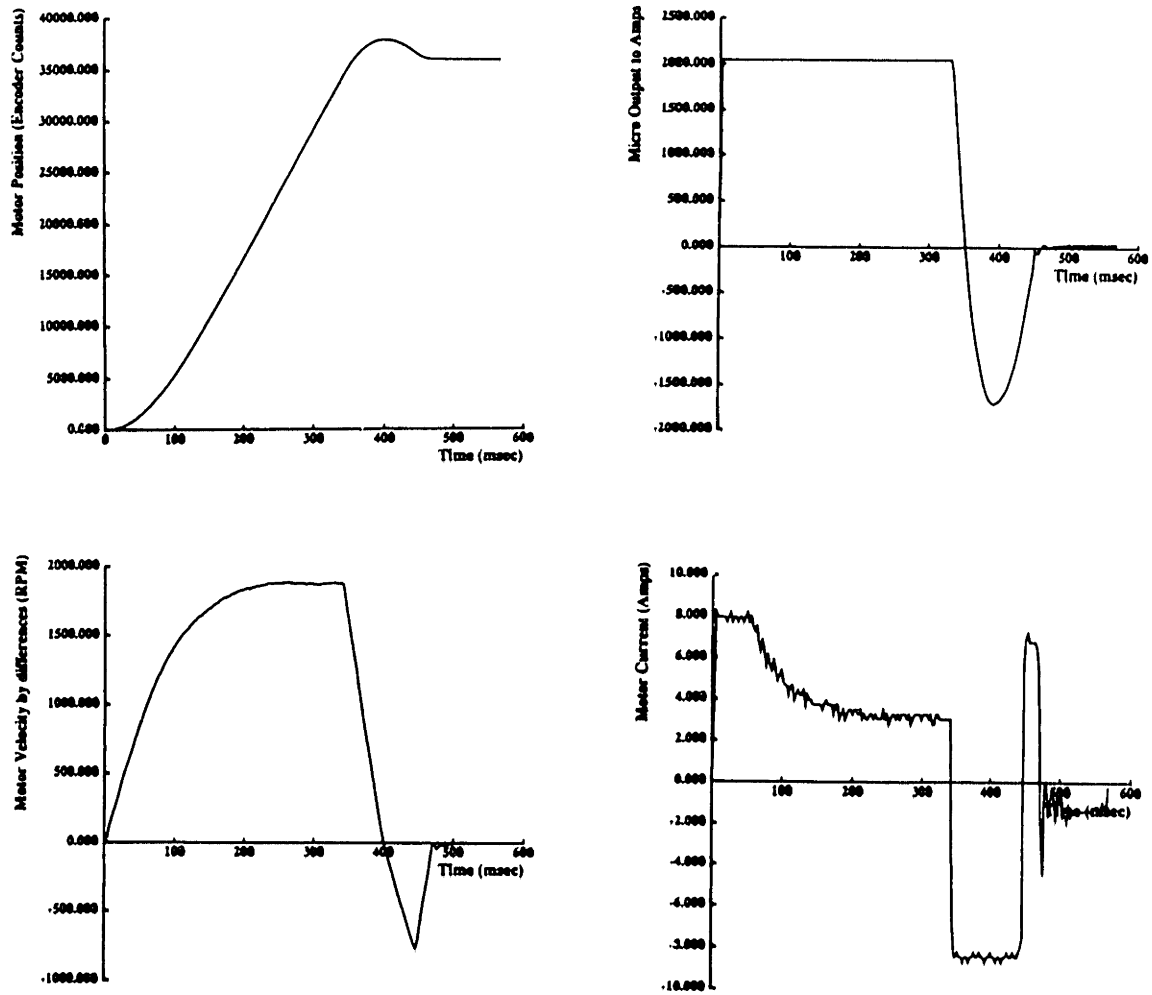


Figure 3.13: Robot Position-Control Step Response Data




Figure 3.14: Robot Force-Control Step Response Data

		Force Setpoint				
		5.0	10.0	15.0	20.0	30.0
Force Gain	5.0	stable	stable	stable	stable	stable
	10.0	stable	stable	<i>unstable*</i>	<i>unstable</i>	<i>unstable</i>
	15.0	<i>unstable</i>	<i>unstable</i>	<i>unstable</i>	<i>unstable</i>	<i>unstable</i>

Figure 3.15: Robot Contact Stability Depends on Force Gain and Force Setpoint

In reviewing this data, we pay particular attention to the stability of the closed-loop system. For low gain, the system is quite stable, and for high gain the system limit cycles. The stability also depends on the force setpoint. We would like to develop a robot model which is able to predict this performance, particularly the stability limit on the force feedback gain.

For any particular force setpoint, there is a limit on the force gain that can be used (while maintaining stability) in the above contact test. There is also a different (lower) limit on the force gain that maintains stability when the stable contact is **perturbed**. Perhaps this is because when the contact is first made, the impact velocity is fairly low, being limited by a function of the force gain, force setpoint, back EMF, and other factors. While in stable contact, a perturbation of the workpiece could involve a much greater velocity than the initial contact condition. For whatever reason, the stability limit is observed to be considerably lower when disturbances in the form of workpiece perturbations are present. In an effort to quantify the stability limit under both conditions, numerous tests were run, with the behavior summarized in the chart of Figure 3.16. The workpiece perturbation in these tests was in the form of a hand sharply striking the workpiece after the robot had made stable contact.

In Figure 3.15 the one unstable case marked with an asterisk* was able to be stabilized using a soft pad on the contacting surface. This grip compliance effect will be discussed in detail in Chapter Five.

Force Gain	Initial Contact	Workpiece Disturbance
4.0	stable	stable
5.0	stable	stable
6.0	stable <i>unstable</i>
7.0	stable	<i>unstable</i>
8.0	stable	<i>unstable</i>
9.0	stable	<i>unstable</i>
10.0 <i>unstable</i>	<i>unstable</i>
11.0	<i>unstable</i>	<i>unstable</i>
12.0	<i>unstable</i>	<i>unstable</i>

Figure 3.16: Summary of Robot Contact Stability for Various Gains

Chapter Four

Modeling the MIT Precision Assembly Robot

The models that we analyzed in Chapter Two did not attempt to rigorously match the behavior of any particular robot. They were developed in an effort to understand the fundamental behavior of machine systems and robots in general. Now we will apply these tools to model the performance of a specific robot system.

The MIT Precision Assembly Robot closely resembles a computer-controlled machine tool, and so it is considerably stiffer and faster than most manipulators. Garcia-Reynoso [33] measured the modal characteristics of the arm to confirm his very detailed model of its structural dynamics. He measured the first vibratory mode to be 52 Hz in free space. After bolting the robot to the floor, the lowest resonance became 12 Hz. A sketch of the mode shape for this "base mode" is shown in Figure 2.15.

Chapter Three described the experiments conducted to measure the response of a single axis of this robot arm. Our goal in this chapter is to develop a model which predicts the behavior that is observed in the real system. First we develop a detailed actuator model. We then create a fitted-polynomial model which matches the experimental data, and explain why this may be inappropriate. Then we propose a lumped-mass model, which includes all of the elements that we think might affect the robot dynamics. This model is unnecessarily complex, and we need to look at the actual system response in order to decide how to reduce the model. A series of reduced-order lumped-mass models is analyzed and we show how they compare to the actual response data.

4.1 Actuator Dynamics

The actuator model to be presented here represents a typical DC electric servomotor drive system. This configuration does closely resemble that of a machine tool, however it also describes many robot actuators. In particular, this model is constructed to match the response of the MIT Precision Assembly Robot X, Y, and Z axis motors.

4.1.1 Amplifier and DC Motor Model

In this electric-motor model we include the significant dynamics of the analog preamplifier circuit, the power amplifier stage, and the DC permanent-magnet electric motor. We couple the actuator system to the load system through a transmission.

A sketch of the robot actuator model is shown in Figure 4.1. The input to this model is the velocity-command voltage given to the preamplifier (this signal usually comes from the digital controller). The preamplifier includes analog tachometer feedback. The power amplifier is configured as a current (transconductance) amplifier, which is simply a voltage amplifier with sensed current feedback. The power amplifier is actually a pulse-width modulated (PWM) device with sufficiently high switching frequency (19.4 kHz) that it can be modeled as a linear amplifier. The current feedback comes from the motor. (Note that modeling the power amplifier simply as a current amplifier would, by causality, force us to ignore the motor dynamics.) The motor input is its terminal voltage. The motor electrical characteristics are constant inductance and resistance, a torque constant, and back EMF. The motor torque output is then proportional to its current. The transmission ratio (ball screw pitch) converts motor torque to the motor model force output. Finally, the rest of the feedback loops are closed, with the load-determined velocity also passing through the transmission to get the motor speed.

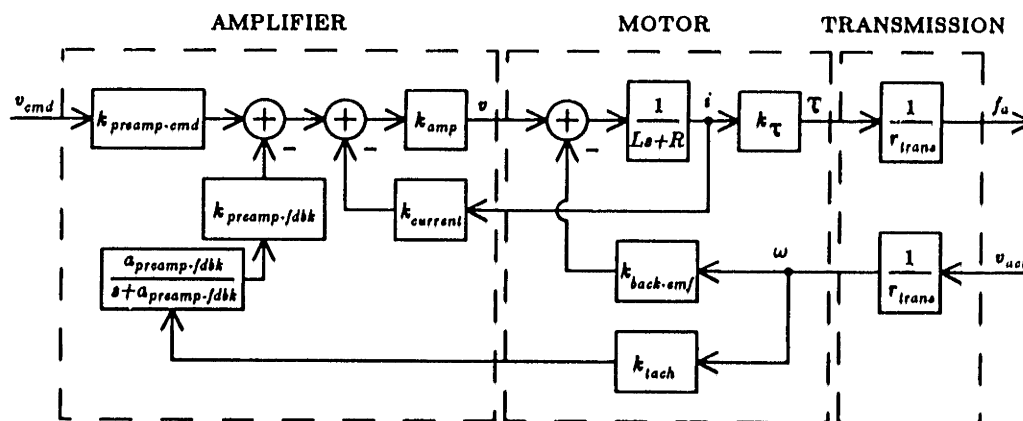


Figure 4.1: Robot Actuator Model

Many nonlinearities present in the actual system are neglected in this model. The most significant of these nonlinearities are the various saturations, which will be included in the nonlinear simulation model of Section 4.5.

The equation governing the dynamics of the motor model is:

$$F_a(s) = \frac{k_\tau k_{amp} k_{preamp-cmd}}{r_{trans}(Ls+R+k_{amp}k_{current})} V_{cmd}(s) - \frac{k_\tau k_{amp} k_{preamp-fdbk} + k_{tach} k_{back-emf} (s+a_{preamp-fdbk})}{r_{trans}(Ls+R+k_{amp}k_{current})(s+a_{preamp-fdbk})} V_{act}(s)$$

The two transfer functions required for the model coupling are:

$$G_{af}(s) = \frac{k_t k_{amp} k_{preamp-cmd}}{r_{trans}(Ls + R + k_{amp} k_{current})}$$

and

$$G_{av}(s) = \frac{k_t k_{amp} k_{preamp-fdbk} + k_{tach} k_{back-emf} (s + a_{preamp-fdbk})}{r_{trans}(Ls + R + k_{amp} k_{current})(s + a_{preamp-fdbk})}$$

4.1.2 Amplifier and Motor Parameters

In the Automatix amplifier used, the preamplifier gains and filter frequency can be set by the user. In our lab, we changed a few resistors to give fairly low preamplifier gains. These were calculated from the circuit schematic to be $k_{preamp-cmd} = k_{preamp-fdbk} = 4.0$. The measured values come from the preamplifier transfer functions of Figures 3.6 and 3.7, giving $k_{preamp-cmd} = 3.9$ and $k_{preamp-fdbk} = 4.2$. We also find $a_{preamp-fdbk} = 942 \text{ rad/sec} = 150 \text{ Hz}$ from Figure 3.7.

The amplifier parameters k_{amp} and $k_{current}$ were found by reading the Automatix documentation [62]. Our amplifiers are set up with 30 amps motor current giving 10 volts of current feedback, so $k_{current} = \frac{1}{3} \text{ volt/amp}$. This signal saturates at ± 10 volts, and this saturation, shown in Figure 4.2, is included in the nonlinear motor model of Section 4.5. The power amplifier itself is a pulse-width modulation (PWM) device with a switching frequency of 19.4 kHz. We model it as a simple voltage amplifier with its voltage characteristic as shown in Figure 4.3. This two-gain saturation results from the offset triangle-wave modulation used. For the linear model, we simply linearize about zero current error and use $k_{amp} = 15$. The nonlinear simulation model does use the saturation curve of Figure 4.3.




Figure 4.2: Simple Saturation

The motor parameters L , R , k_t , $k_{back-emf}$ and k_{tach} are published in the Aerotech motor documentation, but some of these values were found to be off by up to 30%. The motor terminal resistance plus the lead wire impedance was measured with an ohm meter. The armature inductance is very difficult to measure, so the manufacturer's value was used. The tachometer constant was




Figure 4.3: PWM Amplifier Two-Gain Saturation

fairly easy to measure with the following procedure. With the motor shaft free, a constant voltage was applied to the motor terminals, and the tachometer output voltage was recorded. The speed was measured using a Starrett revolution counter. This test was repeated at many speeds and a line was fit by linear recursion, to give its slope as the value of the tachometer constant, $k_{tach} = .039$ volt-sec/rad.

In Section 4.1.3, we will require the motor inertia and damping. The rotor inertia, J_m , is given in the motor specifications. The damping, b_m , is also given by the manufacturer, and this value was verified by the following test: With the shaft still uncoupled, the motor was again driven by a steady voltage, v_{motor} , and the tachometer voltage, v_{tach} , was again recorded. The DC value of the transfer function, $\frac{V_{tach}(0)}{V_{motor}(0)}$ is $\frac{k_t k_{tach}}{Rb_m + k_{back-emf} k_t}$, and we know all the values except b_m . The test is repeated at various speeds to give a reasonable value for the motor damping. The transmission ratio is simply the ball screw pitch $r_{trans} = 1 \text{ inch/rev} = .04867 \text{ m/rad}$.

These motor model parameters are summarized in the table of Figure 4.4.

4.1.3 No-Load Motor Performance

As an example of how to couple a load to the motor model, and to show its no-load performance, we will first analyze the motor behavior under the load of its own rotor inertia and damping. For now, the motor has nothing attached to its output shaft. This analysis corresponds directly with the unattached motor velocity transfer function measured and displayed in Figure 3.9, Figure 4.5 shows the load model of the motor's inertia and damping. The transmission ratio is needed to convert the motor output force back to torque and linear velocity to angular velocity, since the motor model also has these transmission ratios built in. This seems awkward for this simple example, but for the one-axis linear-motion models to be used later, it is convenient not to have to deal with the transmission ratio any further.

The load model transfer functions are:

parameter	value	units	description
$k_{preamp-cmd}$	3.9	unitless	preamplifier velocity command gain
$k_{preamp-fdbk}$	4.2	unitless	preamplifier tachometer feedback gain
$a_{preamp-fdbk}$	942	rad/sec	preamplifier tachometer feedback filter cutoff
k_{amp}	15	unitless	power amplifier linear gain
$k_{current}$.333	volts/amp	current feedback gain
k_{tach}	.039	volt-sec/rad	tachometer constant
k_{torque}	.50	N-m/amp	motor torque constant
$k_{back-emf}$.50	v-sec/rad	motor back-EMF constant
L	.0158	henries	motor armature inductance
R	4.5	ohms	motor and lead terminal resistance
J_m	.0013	kg-m ²	motor rotor inertia
b_m	.02	N-m-sec/rad	motor viscous damping
r_{trans}	.004043	m/rad	ball screw transmission ratio

Figure 4.4: Motor Model Parameters

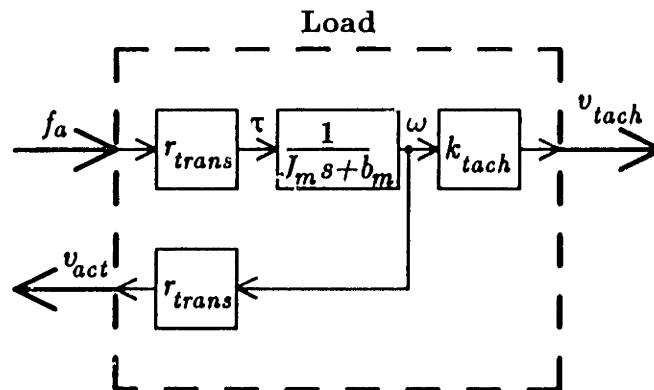


Figure 4.5: No-Load Motor Model with Rotor Inertia and Damping Only

$$\frac{V_{tach}(s)}{F_a(s)} = \frac{k_{tach}r_{trans}}{J_m s + b_m}$$

and

$$\frac{V_{act}(s)}{F_a(s)} = \frac{r_{trans}^2}{J_m s + b_m}$$

The load output chosen is the tachometer signal, which compared with the actuator input command, gives the tracking performance of the actuator model. Figure 4.6 shows the bode plot of the transfer function $\frac{V_{tach}(s)}{V_{cmd}(s)}$ using the parameters for the MIT Precision Assembly Robot actuator system,

The motor model includes two analog feedback loops: the current feedback in the power amplifier stage, and the tachometer feedback in the preamplifier stage. These have the effect of giving the motor fairly good command following ability, as shown by the bode plot. Comparing this analytical response with the measured response of Figure 3.9, we can see that the model matches well both the DC and rolloff characteristics of the data,




Figure 4.6: Bode Plot for the No-Load Coupled Actuator Model of Figure 4.5

4.2 Polynomial Resonance Model

This section develops a model which is a polynomial fit to the structural dynamics measured in the laboratory. The polynomial model is useful for predicting response only under conditions which are exactly like the experimental test. This model lacks the ability to adapt to changes in the system, such as contact with the environment. We would require polynomial models for the colocated transfer function, the noncolocated noncontact transfer function, and the noncolocated in-contact transfer function. As an example, we construct a polynomial resonance model to fit the data obtained for the noncolocated noncontact (endpoint acceleration) transfer function, shown in Figure 3.11. Since this modeling approach is not robust to system changes, and it offers little insight as to how we would design better systems, it is not recommended for design evaluation. This section can be omitted without loss of continuity.

A simple model of the structural resonances within the frequency range of interest can be obtained by constructing a polynomial transfer function to represent the dynamics in the transfer function which was obtained by test. The fitted polynomial function, when cascaded with a rigid-body motor and robot model derived, should match both the magnitude and phase of the experimental transfer function.

We will attempt to match the first five dominant modes in the noncolocated acceleration data. Moreover, we propose to do so by cascading five *simple resonance transfer functions*, each of which has the following form:

$$R_i(s) = \frac{\omega_{pi}^2}{\omega_{zi}^2} \frac{s^2 + 2\zeta_{zi}\omega_{zi}s + \omega_{zi}^2}{s^2 + 2\zeta_{pi}\omega_{pi}s + \omega_{pi}^2}$$

Each simple resonance transfer function includes a pair of poles and a pair of zeros only. The i^{th} simple resonance transfer function has poles at ω_{pi} with damping ζ_{pi} , and zeros at ω_{zi} with damping ζ_{zi} . These are each scaled to have essentially no effect on the response at $\omega < \min[\omega_{pi}, \omega_{zi}]$.

The overall transfer function that should match the data is then

$$G_p(s) = G_{rb}^a(s) \prod_{i=1}^5 R_i(s)$$

where $G_{rb}^a(s)$ is a rigid-body model transfer function of tip acceleration output to velocity command input. Figure 4.7 shows a rigid-body robot model transfer function cascaded with the polynomial resonance model, giving the acceleration output which we will compare with the experimental tip acceleration transfer function.

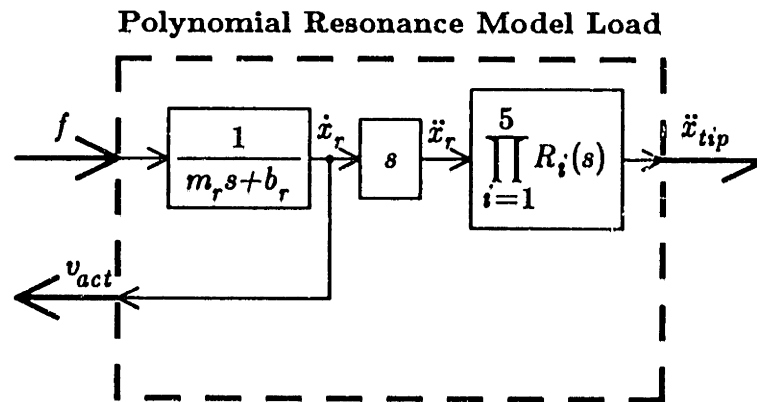


Figure 4.7: Polynomial Resonance Model Load

Construction of the polynomial model involves building up the transfer function to fit the bode amplitude to the data. For each resonance to be matched, starting with the lowest frequency:

1. Compare $|G_p(j\omega)|$ with the experimental transfer function magnitude plot.
2. Find the peak corresponding to the resonance, and set ω_{pi} to the peak frequency.
3. Find the trough corresponding to the resonance, and set ω_{zi} to the trough frequency.
(Typically, $\omega_{zi} \approx 1.1 \omega_{pi}$ for structural modes.)
4. Choose ζ_{pi} to match the height of the peak. (Typically, $\zeta_{pi} \approx .05$ for structural modes.)
5. Choose ζ_{zi} to match the depth of the trough. (Typically, $\zeta_{zi} \approx .05$ for structural modes.)

Note that this method uses only the magnitude plots of the polynomial model and the experimental data, and we construct a transfer function to match them. The phase plots should match also, and they will if the data to be matched actually represent a minimum-phase system. Figure 4.8 shows the frequency response of the polynomial resonance model, which can be compared to the experimental response data in Figure 3.11.




Figure 4.8: Bode Plot for the Polynomial Resonance Model of Figure 4.7

Since the resonance model is made up by cascading the simple resonance transfer functions with the rigid-body model, its pole-zero array includes the poles of the rigid-body model, the zero at the origin which results from choice of acceleration output, and the five pairs of poles and five pairs of zeros of the fitted-polynomial resonance model. This pole-zero array is shown in Figure 4.9.

The polynomial model response does indeed match the experimental transfer function very well (as it should, since it was constructed specifically to match the data). The close resemblance to the distributed-parameter beam result of Section 2.4.4 encourages the development of sophisticated beam models, since the beam models also showed pairs of poles and pairs of zeros alternating along the imaginary axis. Nevertheless, this model does not help our efforts to understand the nature of robot dynamic behavior. In particular, this polynomial model lacks the ability to show how and why the




Figure 4.9: Pole-Zero Array of the Polynomial Resonance Model of Figure 4.7

colocated and noncolocated robot transfer functions differ. Actually, the distributed-parameter beam model and this polynomial model have a significant discrepancy, which is that the beam model showed the noncolocated transfer function to have real zeros, yet our polynomial fit to the robot data shows that complex pairs of zeros match the noncolocated transfer function well. Moreover, this model cannot predict the different robot behavior in and out of contact with the environment. Perhaps the development of more sophisticated fitted-polynomial dynamic models would be a fruitful area for future research.

4.3 Lumped-Mass Models

Now we will couple the detailed motor/amplifier model of Section 4.1.1 to some lumped-mass models. First we investigate how the motor model affects the performance of the rigid-body robot model which was analyzed in Section 2.2 without actuator dynamics. Then we consider higher-order dynamics in an effort to find the model which best explains the behavior of the robot in the laboratory.

4.3.1 Coupled Rigid-Body Model Performance

The rigid-body robot model of Figure 4.10 depicts a rigid robot in contact with a rigid environment. We found this simplified model to be stable for all force feedback gains. In this section, we will couple the rigid robot model to the motor model, as shown in Figure 4.11, and evaluate the closed-loop performance.

The transfer functions needed to couple the models are:

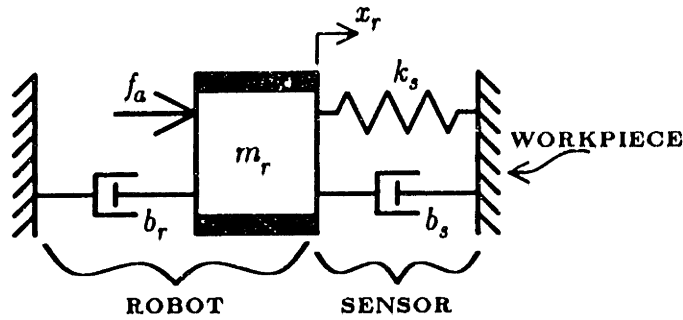


Figure 4.10: Rigid-Body Robot Model with Rigid Workpiece

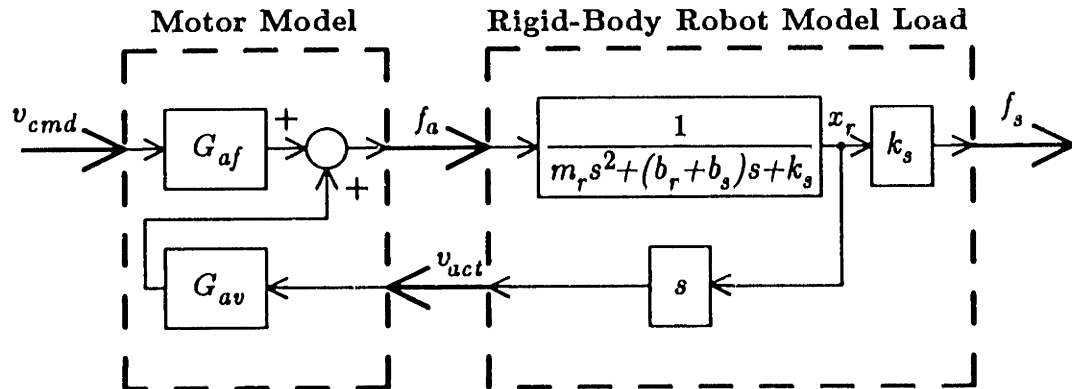


Figure 4.11: Rigid-Body Robot Model Load Coupled to the Motor Model

$$\frac{F_s(s)}{F_a(s)} = \frac{k_s}{m_r s^2 + (b_r + b_s)s + k_s}$$

and

$$\frac{V_{act}(s)}{F_a(s)} = \frac{s}{m_r s^2 + (b_r + b_s)s + k_s}$$

The explicit force control equation for the actuator-coupled system is:

$$v_{cmd} = k_f [f_d - f_s]$$

The coupled system's open-loop bode plot is shown in Figure 4.12, and the root-locus plot for varying k_f is in Figure 4.13. Now we can see how the actuator model has affected the rigid-body robot model response. As expected, this less-than-perfect actuator has a great effect on the stability of the model. (The poles contributed by the motor are beyond the scale of this plot. This scaling, chosen to use for the root-locus plots in this chapter, accentuates the structural dynamics.) For high gains, the closed-loop system is now unstable. We cannot increase the speed of the response without limit by turning up the force gain. This system is limited in bandwidth to about the same crossover seen in the measured open-loop force transfer function of Figure 3.12.

Before considering the higher-order arm dynamics, we should note that this simple rigid-body




Figure 4.12: Bode Plot for the Rigid-Body Coupled Model of Figure 4.11




Figure 4.13: Root-Locus Plot for the Rigid-Body Coupled Model of Figure 4.11

robot model, coupled with the detailed motor model, has predicted the overall behavior of the robot system. The remainder of this chapter is devoted to investigating the significant higher-order dynamics.

4.3.2 Modeling Everything

In Chapter Two, we discussed lumped-parameter models of arm flexibility, base flexibility, workpiece dynamics, and grip compliance. It is not difficult to create a lumped-mass model that combines all of these effects. Figure 4.14 shows such a model.

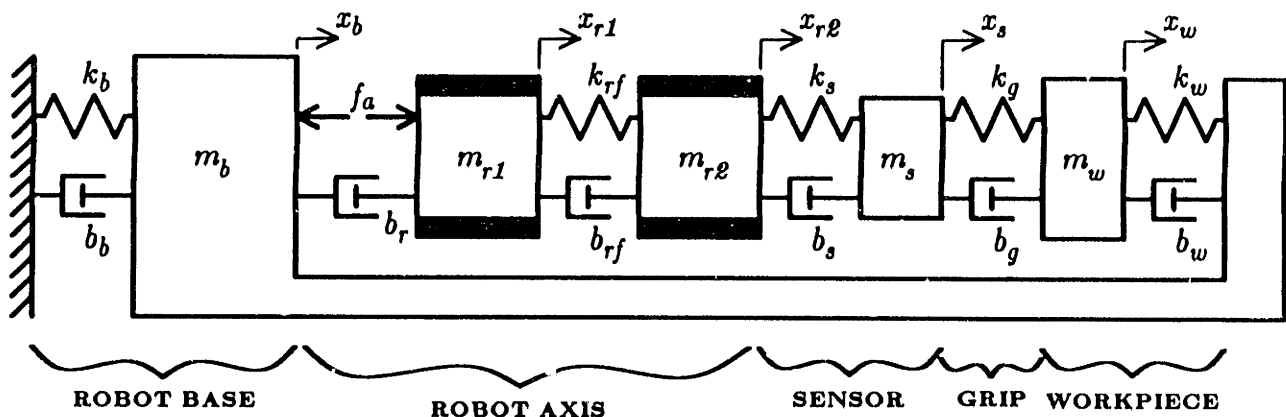


Figure 4.14: A Robot Model Combining All of the Previous Lumped-Mass Model Effects

Deriving the equations of motion would be tedious without help from MACSYMA, a symbolic manipulation package. Nevertheless, the transfer functions for this model can be derived and we can couple this model to the motor model of Section 4.1.1 and conduct the linear analysis.

The response of this very complex lumped-mass model would not surprise us, and it may include unnecessary dynamics. It will be more instructive to instead consider the separate effects of the arm flexibility, base flexibility, workpiece dynamics, and grip compliance. We analyze the effects separately for two reasons:

1. We desire to understand how each of these elements affects the rigid-body response.
2. All of the dynamics are not actually influencing the robot behavior as strongly as this model would lead us to believe.

4.3.3 Reduced-Order Models

We need to develop a procedure for reducing the complexity of this model. Since we cannot model absolutely everything, we must eliminate some effects from the model. Three schemes for deciding what to include in the model come to mind:

1. Model the low-frequency dynamics first, adding modes one at a time until the experimental data is matched well.
2. Include the lowest stiffness elements first. In order, these would be grip, sensor, workpiece, arm, and base.
3. Model everything, then exclude dynamics which do not enhance the match between the analytical and experimental responses.

The first method seems to make the most sense (at least to a control engineer who thinks in the frequency domain). However, we have observed that grip compliance may be important, and although it is the lowest stiffness in the model, it serves to decouple the sensor and workpiece, giving rise to high-frequency modes. The third method is tedious since it requires starting with models that are too complex. The second method, therefore seems to be the correct heuristic for guiding machine model development.

Nevertheless, we will use yet another method, which is more instructive, since our purpose is to learn about machine design. We will consider, as we did in Chapter Two, the higher-order robot system dynamics, one element at a time, in the following order:

- arm flexibility
- base dynamics
- workpiece dynamics
- grip compliance

We will include in the model, any effect which helps to explain the observed (measured) behavior.

4.3.4 Arm Flexibility

The robot model with arm flexibility is shown in Figure 4.15. Note that this model also represents transmission flexibility; however, in this particular example system, it is the arm flexibility (bending of the X-carriage and Z-axis) that we suspect may be significant, not the (relatively stiff) transmission. In other systems, the transmission has been found to be the most significant source of flexibility [36].

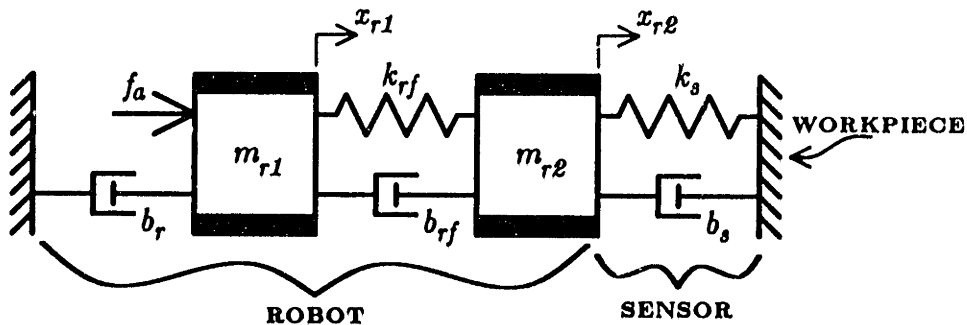


Figure 4.15: Robot Model Including Arm Flexibility

The root-locus plot for the robot model with arm flexibility is shown in Figure 4.16. This root locus shows that the arm bending mode would cause the system to have lower bandwidth than we measured. If this flexible mode were significant in the system, it would cause a large phase shift and instability at about 20 Hz, which is not what was measured.




Figure 4.16: Root-Locus Plot for the Coupled Model with Arm Flexibility of Figure 4.15

4.3.5 Base Dynamics

The appropriate robot model with base flexibility is shown in Figure 4.17. We use this form of the base model because it depicts the workpiece also supported by the robot's dynamic base. The root-locus plot for the rigid-body robot, motor, and base, is shown in Figure 4.18. This root-locus shows that the base mode at 12 Hz adds the pair of poles and pair of zeros which we have come to expect with dynamically colocated dynamics.

Since the base inertia is large compared to the moving robot axis, the frequency of the base mode does not change when contact is made. This mode appears to be damped enough to have little effect on the closed-loop force-control performance. However, we would expect that disturbances and commands with strong 12 Hz frequency components might easily excite the base dynamics. It is particularly interesting to note that the base mode appeared very pronounced in the measured non-contact endpoint acceleration transfer function yet it does not appear in the contact force transfer function data. This phenomenon is not well explained by the base model.




Figure 4.17: Robot Model Including Base Flexibility




Figure 4.18: Root-Locus Plot for the Coupled Model with Base Flexibility of Figure 4.17

4.3.6 Workpiece Dynamics

The robot model with workpiece dynamics is shown in Figure 4.19. The root-locus plot for the coupled motor/robot system with the workpiece dynamics is shown in Figure 4.20. Again, we have added a dynamically-located mode so we find a pair of poles and a pair of zeros added to the coupled rigid-body model root-locus plot. This time, however, our choice of workpiece has placed the (underdamped) workpiece mode at almost the exact same frequency that the rigid-body model showed its crossover. This is merely a coincidence. The frequency 40 Hz seemed like a nice place to put the workpiece mode, so it was designed with this natural frequency.




Figure 4.19: Robot Model Including Workpiece Dynamics




Figure 4.20: Root-Locus Plot for the Coupled Model with Workpiece Dynamics of Figure 4,19

The effect of these workpiece dynamics is that its pair of zeros attracts the roots which would have crossed the imaginary axis, and the workpiece poles then cross instead. Since it appears that the workpiece zeros attract the rigid-body poles, then we should try stiffening the workpiece to give the higher bandwidth. We will now see why this does not work.

Figures 4.21 and 4.22 show results from models including stiffer and softer workpieces, respectively (raising and lowering k_w and b_w by factors of four). Neither change has significantly affected the bandwidth of the system. That is, we have added some interesting dynamics, but we have not changed the open-loop crossover. The low-frequency workpiece dynamics (from the soft workpiece) may give a very underdamped mode, which is dominant and detrimental to the response even at gains well below the limit.

The conclusion here is that workpiece dynamics are important, but it is difficult to make them change the system bandwidth by much. Section 2.5 showed an example where carefully placed workpiece dynamics were able to increase the bandwidth slightly. A similar technique will be used to our advantage in the next chapter, and the key is adding damping to the workpiece as well as choosing the appropriate stiffness.

4.3.7 Grip Compliance

Finally, in an attempt to explain how grip compliance appears to help stabilize robot force control, we add grip flexibility to the robot model with workpiece dynamics. The robot model with grip compliance and sensor mass is shown in Figure 4.23. The force control root locus plot for this model coupled to the robot motor model is shown in Figure 4.24. The grip stiffness used here is high, showing the added dynamically-colocated poles and zeros that we expect. When the grip stiffness is made low, we find a very dramatic change, demonstrated by the root-locus plot of Figure 4.25. The grip zeros prevent the low-frequency instability, and the crossover then occurs at the higher sensor dynamics frequency (120 Hz).

The low grip compliance seems to stabilize our simple robot force control model. In fact, the low grip stiffness values used in this example are taken from the soft gripper pad used in the experiment of Section 3.3.2 which did indeed stabilize one unstable case tested. In the next chapter, we will discuss various interpretations of the grip compliance effect:

1. damping the robot motions
2. lowering the loop gain
3. mechanical low-pass filtering of the robot motions
4. decoupling the robot from the environment

We will also consider at what expense we can stabilize force control with low grip stiffness.




Figure 4.21: Root-Locus Plot for the Coupled Robot Model with Stiffer Workpiece Dynamics




Figure 4.22: Root-Locus Plot for the Coupled Robot Model with Softer Workpiece Dynamics




Figure 4.23: Robot Model Including Grip Compliance




Figure 4.24: Root-Locus Plot for the Coupled Robot Model with Stiff Grip




Figure 4.25: Root-Locus Plot for the Coupled Robot Model with Soft Grip

4.4 Finding the Lumped-Mass Model Parameters

The model parameters are found by considering each part of the robot system separately. We have already found the motor model parameters by considering the actuator system decoupled from the robot. In the same manner, we find the rigid-body, arm flexibility, base dynamics, sensor dynamics, grip compliance, and workpiece parameters.

4.4.1 Rigid-Body Robot Parameters

The robot effective mass m_r is the total moving mass of the robot's X axis, including the entire Z axis, its motor, and the Θ axis. These masses were measured by Garcia-Reynoso [33]. The ball screw inertia can be lumped with the motor rotor inertia. The robot effective damping b_r is small when reflected through the transmission and compared to the motor damping b_m . Benjamin [13] measured b_r to give a nominal value.

4.4.2 Flexible Robot Parameters

In the robot flexibility model, Figure 4.15, the robot compliance, k_{rf} has the units of a linear spring. The two masses, m_{r1} and m_{r2} , are translational masses. (Of course, we could have just as easily used rotational components in our model, but the parameters would still be tricky to measure and compute.) The total moving mass of the axis is already known, $m_{r1} + m_{r2} = m_r$. The effective endpoint stiffness (before, not including, the force sensor), k_{rf} was measured with the motor shaft locked to ground and the robot endpoint is in free space. Next, the natural frequency and damping ratio were measured for the arm also in this condition. (There was no rigid-body motion, since the actuator was locked.) The mode that we are trying to represent was excited and its (underdamped) time response was recorded. The period of the vibration oscillations gives the natural frequency, ω_n . The damping ratio, ζ , is calculated from the ratio of the vibration amplitudes before and after a number of cycles, and using this handy formula:

$$\zeta = \frac{\pm \ln(\text{Amplitude Ratio})}{2\pi(\text{Number of Cycles})}$$

Finally, the two masses were computed by solving the following two equations:

$$\begin{aligned} m_{r1} + m_{r2} &= m_r \\ \omega_n^2 &= \frac{k_{rf}(m_{r1} + m_{r2})}{m_{r1}m_{r2}} \end{aligned}$$

The symmetry of this problem gives two solutions which involve the interchange of m_{r1} and m_{r2} . This final conflict was resolved by inspection of the robot structure. Since the actuator inertia accounts for more than half of the total inertia, then m_{r1} must be greater than m_{r2} .

4.4.3 Base Model Parameters

The base dynamics are found in a very similar manner. Garcia-Reynoso measured the base mode frequency and damping ratio. He also gives the masses of the components which make up the X-axis base. Knowing m_b , the base natural frequency, and its damping ratio, we compute k_b and b_b .

4.4.4 Sensor Model Parameters

The sensor stiffness k_s was measured with a spring scale and a dial indicator. The sensor mass m_s was calculated from its dimensions. The natural frequency and damping ratio were measured, and then the mass calculation was confirmed and the sensor damping b_s could be calculated.

4.4.5 Grip Compliance

The stiffnesses k_g of the two grip pads used in the tests were measured using a special jig with a spring scale and a dial indicator. The natural frequency and damping ratio were then measured for each jig/pad system. The natural frequency and stiffness allowed the jig effective mass to be calculated, and then the grip damping b_g could be found by analyzing the jig/pad mass/spring system.

4.4.6 Workpiece Model Parameters

The workpiece stiffness k_w was measured with the spring scale and the dial indicator. The workpiece mass m_w was calculated from its dimensions, then confirmed by measuring the workpiece natural frequency. Since these matched well, we know that the X-direction horizontal-displacement mode (not the twist mode) is dominant. The damping ratio was measured and b_w was computed from that.

parameter	value	units	description
m_r	132.4	kg	effective moving mass
b_r	500	N-sec/m	effective damping to ground
m_b	500	kg	base mass
k_b	3,600,000	N/m	base stiffness
b_b	9500	n-sec/m	base damping
m_{r1}	114.5	kg	robot actuator lumped mass
m_{r2}	17.9	kg	robot endpoint lumped mass
k_{rf}	665,000	N/m	arm bending stiffness
b_{rf}	640	N-s/m	arm damping
m_w	6.6	kg	workpiece mass
k_w	412,700	N/m	workpiece stiffness
b_w	8.25	N-sec/m	original workpiece damping
m_s	.14	kg	sensor mass
k_s	110,100	N/m	sensor stiffness
b_s	16.5	N-s/m	sensor damping
k_g	25000	N/m	low grip stiffness
b_g	30	N-s/m	grip damping

Figure 4.26: Robot Model Parameters

4.5 Nonlinear Model Simulations

The strength of this thesis is not in the treatment of nonlinearities. Nor are the nonlinear simulations necessary to understand the fundamental behavior robot force control. Nevertheless, since it was not clear from the start that this would be the case, many nonlinearities were included in the robot models.

This section will present a few nonlinear simulations for comparison with the laboratory data, although these results are not used to confirm the validity of the flexible robot models.

4.5.1 Position Control Step Response

The position control step response was intended to verify the modeling of the nonlinearities in the motor/amplifier/rigid-body system, since the flexible modes are not noticeable in this response. The simulated position-control response is shown in Figure 4.27, and it can be compared to the experimental data of Figure 3.13. The four curves show the encoder position, motor velocity, motor current, and controller velocity-command output. These curves match the experimental data nicely and show that the amplifier and controller saturations have been included correctly.

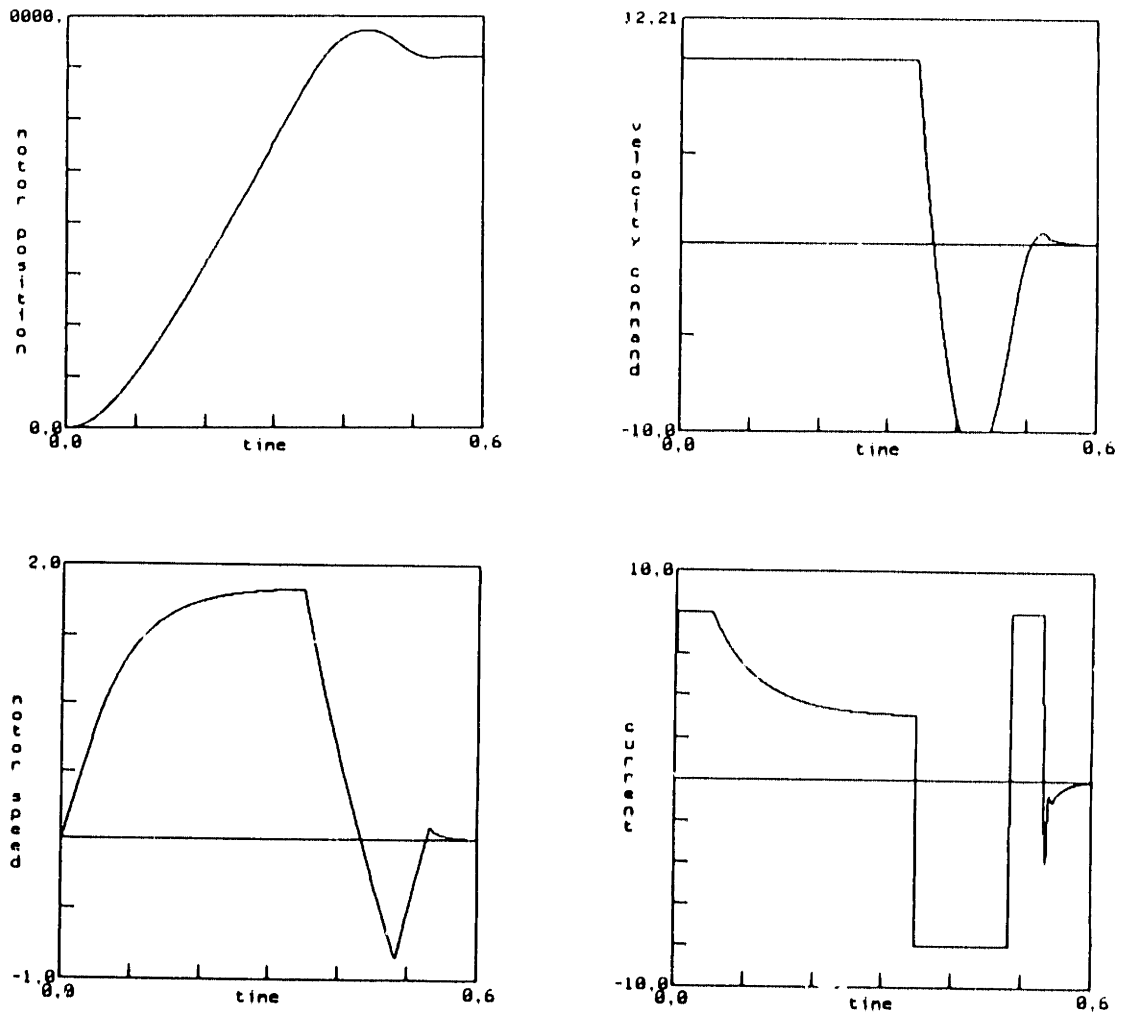


Figure 4.27: Nonlinear Robot Position-Control Step Response Simulation

4.5.2 Force Control Transient Response

A sample of the force control transient response will be shown here to display the modeling of the contact discontinuity. Figure 4.28 shows the motor position and the contact force for the robot flexibility model, with contact discontinuity, and coupled to the motor model. These plots demonstrate the contact transient response limit cycle for an unstable force control case.




Figure 4.28: Nonlinear Coupled Flexible Robot Model Force-Control Simulation

Chapter Five

Improving Performance

This chapter discusses improving the performance of robots under closed-loop force control. We emphasize use of the models developed to evaluate proposed schemes and to explain how they work. Two methods are suggested for improving force control performance. First, we will discuss adding low grip compliance, and second, we will consider adding workpiece damping.

5.1 Grip Compliance

Chapter Two introduced a model of grip compliance, which we used again in Chapter Four to model the specific case observed to stabilize an unstable contact situation in the laboratory. In this section, we will review the grip compliance model, show some actual data, and discuss various interpretations of how it works.

5.1.1 Model Predictions

The model of grip compliance, shown in Figure 2.35 includes the robot effective inertia, the sensor mass, and the workpiece mass. In Section 4.3.7, we coupled the grip model to the detailed motor model of Section 4.1.1. Using the parameters which represent the MIT Precision Assembly Robot, we found that, for a stiff grip, a high-frequency pair of poles and nearby pair of zeros was added by the sensor/grip dynamics (Figure 4.24). The robot response is basically unaffected by the stiff grip, since we are comparing it to a rigid--grip model.

When the grip stiffness is very low, to model the grip pad used in the one stabilizing laboratory experiment, we find a very dramatic effect. Figure 4.25 shows that the grip zeros move to a much lower frequency, adding enough phase lead to make the system much more stable. (In the limiting case, as k_g approaches zero, the grip zeros tend to cancel the workpiece poles, since the workpiece dynamics would no longer influence the robot transfer function response.)

5.1.2 Experimental Data

Two grip pads were used in the experiments. Figure 3.15 shows that for $k_f=10$ and $f_d=15$, the contact response is unstable with no grip pad ($k_g=\infty$). The sensed force data for this case is shown in Figure 5.1. For the hard grip pad ($k_g=120,000$ N/m), the system is still not stable, as shown in Figure 5.2. When the very soft grip pad ($k_g=25,000$ N/m) analyzed above is used, the response is indeed stable, as shown in Figure 5.3. The dominant poles, as predicted by the analysis, are at the workpiece dynamics frequency, shown in the root locus plot of Figure 4.25.

5.1.3 Interpretation

There are many explanations for the stabilizing low grip compliance effect, however, not all of them are complete and correct. We will discuss these possible explanations:

1. damping of the robot motions
2. mechanical low-pass filtering of the robot motions
3. lowering the loop gain
4. decoupling the robot from the workpiece

At first glance, it might appear as though the grip pad is actually **damping the robot motions** and thereby stabilizing the response. More careful analysis, however, reveals that while the grip damping seems high for its low stiffness, it is still much too small to affect the large robot inertia. The grip pad may be able to damp the sensor quite well, and perhaps also the workpiece, but not the robot itself.

Since the robot motions pass through the grip to interact with the workpiece, it seems as though the grip compliance may act as a **mechanical filter** to pass only the low-frequency components of the robot motions. This argument is in fact, to some extent, quite true. The grip dynamics actually add a pair of zeros and a pair of poles which look like a mechanical lead compensator. The zeros give phase lead at low frequency and the poles take it away at the higher frequency.

More dramatic, however, (and this is not shown by the root-locus plots) is that higher force control gains can now be used since the low grip stiffness **lowers the loop gain** of the open-loop system. The loop gain is a function of the sensor stiffness, workpiece stiffness, grip stiffness, and force feedback gain. When either the sensor, grip, or workpiece stiffness is lowered, the force feedback gain can be raised. The higher feedback gain alone cannot be interpreted as achieving better performance, since it is the total loop gain (and also the phase) that matters. Roberts [75] found that lowering sensor stiffness allowed higher feedback gains to be used, and he noted that the contact frequency also lowered. This effect is indeed explained by the grip model. If the only effect of adding a compliant element were to lower the loop gain, then the root locus plot would not change shape. In fact, the poles move around considerably, and the coupled system must be analyzed as a whole in order to understand all the dynamics involved.




Figure 5.1: Force Data Showing Instability with the Rigid Grip




Figure 5.2: Force Data Showing Instability with the Hard Grip Pad




Figure 5.3: Force Data Showing Improved Stability with the Soft Grip Pad

Finally, when we analyze the entire robot/grip/workpiece system, we realize that when the grip stiffness is low, it **decouples the robot from the workpiece**. With high grip stiffness, the sensor and workpiece move together. With low grip stiffness, the sensor and workpiece can move independently. This allows the robot to behave more like when it is unattached, and the workpiece can act as if it is not coupled to the robot. The robot becomes *less responsive* since the endpoint interactions now require greater arm motions and result in smaller force sensor signals. Higher force feedback gains are needed, and overall bandwidth may be improved, but by other measures of the quality of active force regulation, the performance may suffer.

At what expense do we stabilize force control with low grip compliance? We certainly lose endpoint positioning accuracy. In many compliant control tasks, however, the position is not even important. When it is important, we may be able to compensate for the position errors [74]. For larger displacements, the gripper pads may become much stiffer as the compliant elements compress. This stiffness now increases the loop gain, which destabilizes the system. Also, for larger motions, the actuator may saturate and this can limit the achievable performance.

5.2 Workpiece Damping

When the workpiece is very stiff, it couples the robot to ground. Adding damping to the workpiece can in fact damp the entire robot response. The well-damped workpiece adds a resonant mode which has poles and zeros away from the imaginary axis. The placement of these dynamics can be used to increase the system bandwidth.

Figure 5.4 shows a root-locus plot for the MIT Precision Assembly robot model with workpiece dynamics as analyzed in the previous chapter. Figure 5.5 shows the root locus plot changed when damping has been added to the workpiece (with b_w increased by a factor of 20). The workpiece zeros still attract the rigid-body poles, and the workpiece poles emit roots which cross the imaginary axis, however, in this case, the roots start out better damped and cross into the right-half plane at a higher frequency. Higher bandwidth can therefore be achieved if these model predictions are correct.

5.2.1 Workpiece Damping Tests

Adding damping to the workpiece seems quite distant from suggestions made by other researchers, so the validity of this idea is verified in the laboratory. The workpiece design was discussed in Section 3.1.3, and as constructed, the workpiece had a very underdamped vibratory mode of about 40 Hz. The most significant damping was the low internal damping of the steel supporting beam. A method for adding a great deal of damping in parallel had to be found. A series damper would change or eliminate the workpiece stiffness to ground. The parallel damper design should not change the workpiece stiffness, since we would like to test a damping parameter change, not stiffness.




Figure 5.4: Root Locus for the Robot Model with the Original Workpiece Damping





Figure 5.5: Root Locus for the Robot Model with Increased Workpiece Damping

Adding a little damping is in fact quite easy, but this workpiece needed very much. To quantify the damping methods tested, an accelerometer was placed on the workpiece, and its impact response was measured. The damping ratio can be calculated from the time record of the acceleration. The workpiece, as originally constructed, showed vibrations which looked like the time history in Figure 5.6a, with a damping ratio of $\zeta=0.002$. Two damping configurations were tested: damping the beam motion, and damping the workpiece motion. The beam motion damper involved placing various damping materials between the steel workpiece support beam and ground. This worked, and the proposed damping materials gave results looking like the response shown in Figure 5.6b, with damping ratios of about $\zeta=0.015$. The damping materials tested include: styrofoam, clay, liquids, cloth, cardboard, wood, oil, grease, and a special viscous damping fluid from Dow-Corning. (The Dow fluid is a very linear viscous damping fluid. Even very small forces result in small velocities. It has almost no "viscous stiction".) Damping the workpiece itself worked better because the larger workpiece allowed a greater surface area over which to place the viscous fluid for shear between two plates. A flat plate was mounted above the workpiece, and fixed also to the robot frame. A very thick grease (external gear lubricant) and the special Dow fluid were applied between the contacting surfaces. This workpiece damping configuration is shown in Figure 5.7. Both materials achieved excellent damping over the large area. The grease was preferred for the remainder of the tests since it does not drip off overnight. The nonlinearity of the grease damping is not severe by robot standards. The final damping ratio measured was $\zeta=0.11$, which is shown in Figure 5.6c.


The robot dynamic tests of Section 3.3.2 showed that the robot is stable under force control only for low gains. Here we will present data which shows the robot more stable in contact with the damped workpiece. Figure 5.8 shows the encoder position data and the sensed force data for $k_f=7$ with $f_d=40$, which gives unstable response for the original workpiece design. Figure 5.9 shows the response data for the same gain and setpoint giving dramatically more stable response in contact with the damped workpiece. The range of stable force control gains is now larger. For initial contact, roughly 10% higher gains can be used. For workpiece disturbances, roughly 40% higher gains can be used, as shown in Figure 5.10.

The workpiece damper has helped to stabilize the workpiece contact. With higher allowable gains, the robot is able to make initial contact at much higher velocities and can maintain stability for considerable workpiece position and velocity disturbances. Also higher gains give higher bandwidth, since in this case, we have not reduced the stiffness to ground.

a)
Initial Workpiece Response
 $\zeta = .002$



b)
Some Damping Added
 $\zeta = .015$



c)
Final Workpiece Response
 $\zeta = .11$




Figure 5.6: Workpiece Damping Trials




Figure 5.7: Design of the Workpiece Damper




Figure 5.8: Robot Force-Control Step Response Data for the Original Workpiece




Figure 5.9: Robot Force-Control Step Response Data for the Damped Workpiece

Force Gain	Initial Contact		Workpiece Disturbance	
	internal damping only	external damping also	internal damping only	external damping also
4.0	stable	stable	stable	stable
5.0	stable	stable	stable	stable
6.0	stable	stable	unstable	stable
7.0	stable	stable	unstable	stable
8.0	stable	stable	unstable	stable
9.0	stable	stable	unstable	unstable
10.0	unstable	stable	unstable	unstable
11.0	unstable	unstable	unstable	unstable
12.0	unstable	unstable	unstable	unstable

Figure 5.10: Contact Stability Is Increased with Workpiece Damping

Chapter Six

Conclusion

The goal of this thesis research was to understand the dynamic behavior of robots and to learn how to design and build better-performing machines. While some mysteries still remain, this work has come a long way toward the understanding we seek. The first part of this chapter reviews the lessons learned about the dynamic performance of machines. The latter part discusses some of the ways in which this work can be extended in future research.

6.1 Contributions of This Research

This thesis presents contributions in these areas:

1. modeling machine performance
2. understanding robot bandwidth limitations
3. force control performance improvements
4. insights for machine design

These areas are discussed in the sections which follow.

6.1.1 Modeling Machine Performance

Many dynamic models were developed, from which we were able to learn a great deal about the various effects that higher-order dynamics can have. The series of lumped-parameter robot models was used to describe arm flexibility, transmission flexibility, base dynamics, workpiece dynamics, sensor dynamics, and grip compliance. The beam models showed that bending modes give rise to quite different dynamics (including nonminimum phase zeros) than the tension/compression modes which we usually think about.

Deciding what to include in a model is not straightforward. We hypothesize that it is not necessarily the lowest-frequency modes that are the most significant, but rather it is the lowest stiffness elements that should be considered first. The most instructive (and most reliable) method seems to be to model everything, and then use only the effects that show up in the actual system response. When equations for many models are to be derived, a suitable algebraic decoupling can be used to allow the equations for parts of the models to be derived separately and then coupled mathematically. In our examples, we derived the robot and actuator models separately, coupling them only for analysis.

The strengths and weaknesses of the linear modeling technique have been illustrated with an example robot system. This thesis stresses that the astute control engineer should make open-loop frequency-response measurements of the actual transfer function involved. If remote sensing is to be used, such as an endpoint force sensor, then the noncolocated (actuator-to-sensor) transfer function is important. The colocated (actuator-to-actuator) transfer function will show less phase lag and is therefore more stable. Finally, for robot compliant control, the in-contact transfer function, which is different from the not-in-contact transfer function, must be measured.

6.1.2 Understanding Robot Bandwidth Limitations

There are many features of a machine design which determine the achievable closed-loop bandwidth for the system. Figure 6.1 shows a hypothetical open-loop bode phase plot for a machine system or a robot. This phase plot has three important features:

1. the rigid-body rolloff
2. a low-frequency colocated mode
3. a higher-frequency noncolocated mode

Each of these three dynamic forms can contribute to or entirely determine the achievable bandwidth.




Figure 6.1: Three Dynamic Problems in Machine Control

The rigid-body model coupled with the actuator dynamics generally has at least three more poles than zeros (two from the rigid axis, and one from the actuator dynamics). This situation causes the open-loop phase to drop below -180° as shown in the dashed line of Figure 6.1. So even in the absence of flexibility, there is a phase crossover, which we discovered in Section 2.3.

If there are dynamically colocated modes at frequencies below that of the rigid-body crossover, then these may further limit the bandwidth. Underdamped base or workpiece dynamics can give a pair of poles and a pair of zeros lying together near the imaginary axis, which can cause a sharp phase dip of 50° or more at that frequency. This dip may be enough to push the total phase below -180° . If the colocated mode occurs at a very low frequency, where there is greater phase margin, then the system may retain stability. If a low-frequency colocated mode is very well damped, however, the phase dip will be greatly diminished. Adding damping to the low-frequency modes of this type is an interesting design solution.

Dynamically noncolocated modes, such as those contributed by arm or transmission flexibility, come with more poles than zeros. These poles add phase lag, and could give 180° of phase shift. The frequency of the lowest dynamically noncolocated mode is a fundamental performance limitation.

6.1.3 Force Control Performance Improvements

The robot models were particularly useful in identifying ways to change the system for improved performance. The models showed that reducing workpiece, grip, or sensor stiffness serves to decouple the robot from the environment and to reduce the effect of the endpoint feedback. Higher feedback gains are then needed to compensate for the lower loop gains. In particular, we studied the effect of low grip stiffness, which has been observed to stabilize contact with stiff environments. Greater bandwidth can be achieved with higher force control gains, however performance may suffer as larger robot motions are required to execute the tasks involved.

As an example of a performance improvement that would not be predicted by other analysis techniques, the effect of workpiece damping was explored and tested in the laboratory. A well-damped workpiece places the poles further from the imaginary axis so greater bandwidth can be achieved. Various workpiece damping methods were tested to design an effective parallel workpiece damper. The robot system is considerably more stable in contact with the damped workpiece.

6.1.4 Insights for Machine Design

Throughout the detailing phase, the designer is faced with decisions like: "*How thick shall I specify this wall?*" and "*Where should I place this support?*" Machine designers basically understand the tradeoffs involving mass and stiffness. They know that adding stiffness raises the resonant frequencies, which somehow helps performance, but that this stiffness can increase the moving mass, which certainly slows the machine down. While designers do generally understand the rigid-body motion problem, they do not usually know about the effects of the various flexible modes. Designers need to understand enough about structural dynamics to "design the mode shapes". This thesis hopes to have addressed this very important concern.

We have, for the purposes of discussion, reduced the interpretation of modal analysis to two

types of resonances: dynamically colocated modes, and dynamically noncolocated modes. The distinction between them is very simple. Flexible modes in which the actuator and sensor move together (in phase) despite their noncolocated physical placement on the structure are dynamically colocated modes. Flexible modes in which the actuator and sensor move apart (out of phase) are dynamically noncolocated modes. In fact there is a third type of flexible mode, and that is one in which there is not any relative motion between the actuator and sensor (a mode in an orthogonal direction, for example), and this type of mode is unimportant for this discussion.

Figure 6.2 illustrates some of the different types of robot flexibility that we have modeled and analyzed in this thesis. Transmission and link flexibility result in noncolocated control modes, which will limit the achievable bandwidth. The base and workpiece dynamics contribute dynamically colocated modes, which do not severely affect the system performance.




Figure 6.2: Types of Robot Flexibility

Dynamically colocated modes can, if their frequency is below the phase crossover of the rigid-body system, cause an instability. However, physically adding damping to these modes can eliminate their destabilizing effects (although it is not clear how to do this).

Dynamically noncolocated modes present a fundamental performance limitation. The phase lag which accompanies a bending mode arising from flexibility between the actuator and sensor can

be 180° . Adding damping to these modes only makes the phase transition less sharp, not less severe. The designer must pay particular attention to increasing the frequency of these flexible modes. Every element between every actuator/sensor pair should be designed with high stiffness in mind. Low frequency noncolocated modes should be eliminated (raised in frequency) by stiffening the elements between the actuators and sensors. It is difficult to do this without increasing the effective moving mass of the axis, however the resonance only needs to be increased beyond the open-loop phase crossover of the rigid-body model. The frequency of the lowest mode in which the actuator and sensor move opposite in phase (from their rigid-body mode motion) represents a fundamental performance limitation. All of the other modes are less important.

The system designer should develop a dynamic performance model as part of the machine design process. A rigid-body model of the proposed system, like the one developed in Section 4.1.1, can be analyzed before the actuator and transmission specifications are finalized. The effective stiffness and inertia for each suspected flexible mode can also be estimated in the detailed design stage in order to keep tabs on the resonant frequencies of the machine. Methods for doing this are discussed in books on stress calculation and structural analysis [15, 73]. Finally, these estimates can be used as parameters for performance models like those presented in this thesis. There are many methods for developing complex structural dynamics models. However, it is generally not straightforward to derive transfer functions from these models, so the models need to be chosen carefully.

6.2 Recommendations for Future Research

This research project has indeed shown some useful results. However, there are questions which remain unanswered, and there are ideas which have surfaced that have yet to be verified as well. The recommended work falls into the following categories:

1. robot control
2. machine design
3. system dynamics
4. modeling of machine systems

6.2.1 Robot Control

There are many ways to improve robot performance. Some of these should be implemented to see if they work. Derivative force control and lead filtering have been suggested as ways to increase the achievable closed-loop force control bandwidth. Implementing the differentiation may be difficult, and if it is to be done digitally on the MIT Precision Assembly Robot system, we need to first add 12-bit analog-to-digital converters and eliminate a few noise sources.

Perhaps a new type of endpoint force sensor can be developed. Conventional sensors provide a

voltage signal proportional to the displacement only of the compliant sensor element. If a sensor could also provide a signal proportional to the velocity across the sensor, then this signal could be used for derivative force control, adding a zero to the open-loop system.

If higher-order robot dynamics are causing some problems, then perhaps it would be better if we be sure to not excite them at all. Recently, schemes have been developed which promise to reduce endpoint vibrations for open-loop moves [86]. A scheme of this sort may be useful in force control, and could, for example, be used at the force command input.

Soft environments are easier to make stable contact with, and this is for three reasons. First, the compliant workpiece, grip, or sensor moves zeros closer to the rigid-body poles. Second, the compliant element decouples the environment from the robot, allowing the robot to take on more of its unattached behavior. Third, the overall loop gain is reduced, since the robot, sensor, grip, and workpiece stiffness act in series to couple the robot to ground. Sometimes the bandwidth is no better than would be achievable with stiffer elements, and turning down the gains. The tradeoff seems to be that lower stiffness allows higher force gains to be used, which gives faster response and lower steady-state error; on the other hand, the positioning accuracy is poor and greater arm motions are required. This area should be studied further to develop guidelines for robot and task designers.

6.2.2 Machine Design

Current robots have generally been developed with position control in mind. The design rules developed here may be able to inspire a superior robot configuration. An interested designer may wish to take on the development of a new robot with force control performance in mind. There are many appropriate suggestions in the literature which can be applied [6, 64, 82, 91, 93, 103].

Adding damping to the workpiece and the other colocated modes is an interesting design issue. If a case can be found where a low-frequency colocated mode is limiting closed-loop machine performance, then damping should be added. This did not appear to be the case in the MIT Precision Assembly Robot system.

Adding workpiece damping looks very promising. This should be tried for an application using the PUMA robot. For assembly, the jig which holds the parts can be made with a large amount of damping. It is not clear how to add viscous damping in a practical way. Nevertheless, this is a solvable design problem. A general method for adding damping is to use a very viscous fluid, in a thin layer, between large shear surfaces, with as much relative motion as possible.

6.2.3 System Dynamics

Many interesting dynamics issues came up in this project, and some questions remain open and are particularly important in furthering the understanding of how to develop better-performing equipment. There are significant differences between the colocated and noncolocated transfer functions. The beam models shed some light on this subject by showing the different arrangements of the zeros. In the more complex robot system tested, we also observed much less phase lag from input to output in the colocated cases. Just how and why the zeros move around still remains a mystery.

It would also be nice to be able to predict the changing of the system dynamics when contact is made. The models here attempted to do this but they still showed some modes being more influential than they really were in the data. The base mode and the robot arm flexibility mode turned out to be less important than the models predicted. Perhaps we need more sophisticated models.

Nonminimum phase zeros can also show up in a simple two-link robot system, such as that shown in Figure 6.3, depending on the mass properties of the links and on the control structure. These right-half plane zeros can attract roots and limit the stability of the system. We should attempt to understand the linearized dynamics of multi-link systems.




Figure 6.3: Nonlinear Two-Link Rigid Robot Model with Nonminimum Phase Zeros

6.2.4 Modeling of Machine Systems

We need a better method for deciding what to include in a robot model. The lowest-frequency modes are not always the most important. The lowest stiffness elements might be always needed. An interesting project would involve modeling and testing various widely different systems to develop modeling heuristics.

The polynomial fit models were useful because they show the features that are important in the measured data. Perhaps they could be useful if separate polynomial models could be made for both the constrained and unconstrained systems as well as both the colocated and noncolocated transfer functions. Then these models could be mathematically coupled to the actuator model and used in nonlinear simulation.

Finally, some new modeling techniques could prove useful in showing robot behavior. Perhaps the series lumped-mass models can be enhanced by placing "transmissions" between each mass to set their relative effect on the overall response. This may allow the models to describe the effects of complex geometries which change the excitation of the modes.

The beam models seem to be quite promising. They are pure and simple, and they can explain the effects of bending, which is one type of higher-order dynamics that is quite common in systems. Perhaps the continuous bending beam model can be extended to include two-dimensional shapes and outputs at positions along the beam other than the endpoints. The lumped-parameter bending model may be more practical for control systems analysis, and it still shows the most important beam dynamics feature, namely the nonminimum phase zeros.

References

- [1] Abelow, A.V.
Dynamic Equation Set for a Simplified Simulation of the Space Shuttle Remote Manipulator System.
Technical Report, Charles Stark Draper Laboratory, February 1980, no. R-258.
- [2] Albus, J.S.
Brains, Behavior, and Robotics.
BYTE Books, 1981.
- [3] An, C.H. and Hollerbach, J.M.
Dynamic Stability Issues in Force Control of Manipulators.
IEEE International Conference on Robotics and Automation, April 1987, pp. 890-896.
- [4] An, C.H. and Hollerbach, J.M.
Kinematic Stability Issues in Force Control of Manipulators.
IEEE International Conference on Robotics and Automation, April 1987, pp. 897-903.
- [5] Armstrong, B., Khatib, O., and Burdick, J.
The Explicit Dynamic Model and Inertial Parameters of the PUMA 560 Arm.
IEEE International Conference on Robotics and Automation, April 1986, pp. 510-518.
- [6] Asada, H., Youcef-Toumi, K., and Ramires, R.
Design of MIT Direct-Drive Arm.
International Symposium on Design and Synthesis, Tokyo, Japan, July 1984.
- [7] Asada, H. and By, A.B.
Kinematics of Workpart Fixturing.
IEEE International Conference on Robotics and Automation, March 1985, pp. 337-345.
- [8] Asada, H. and Goldfine, N.
Process Analysis and Compliance Design for Grinding with Robots.
In Donath, M. and Leu, M. (editors), *Robotics and Manufacturing Automation.* ASME Winter Annual Meeting, November 1985, vol. PED-15, pp. 79-88.
- [9] Asada, H. and Lim, S-K.
Design of Joint Torque Sensors and Torque Feedback Control for Direct-Drive Arms,
In Donath, M. and Leu, M. (editors), *Robotics and Manufacturing Automation.* ASME Winter Annual Meeting, November 1985, vol. PED-15, pp. 277-284.
- [10] Bailey, T.L. and Hubbard, J.E., Jr.
Distributed Piezoelectric-Polymer Active Vibration Control of a Cantilever Beam,
AIAA Journal of Guidance and Control, September-October 1985, vol. 8, no. 5, pp. 605-611.
- [11] Balek, D.J. and Kelley, R.B.
Using Gripper-Mounted Infrared Proximity Sensors for Robot Feedback Control.
IEEE International Conference on Robotics and Automation, March 1985, pp. 282-287.

- [12] Barry Wright Corporation, Products for Flexible Automation. Astek Robotic End-of-Arm Products. Watertown, MA
- [13] Benjamin, M.H. *Design and Analysis of a Control System for the M.I.T. Precision Assembly Robot.* Master's thesis, Massachusetts Institute of Technology, January 1985.
- [14] Benjamin, M.H., Garcia-Reynoso, A., and Seering, W.P. *Dynamic and Vibration Modeling of a Cartesian Robot.* IEEE International Conference on Robotics and Automation, March 1985.
- [15] Blevins, R.D. *Formulas for Natural Frequency and Mode Shape.* Van Nostrand Reinhold Company, New York, 1979.
- [16] Book, W.J. Recursive lagrangian Dynamics of Flexible Manipulator Arms. *International Journal of Robotics Research*, 1984, vol. 3, no. 3, pp. 87-101.
- [17] Brock, D. and Chiu, S. Environment Perception of an Articulated Robot Hand Using Contact Sensors. In Donath, M. and Leu, M. (editors), *Robotics and Manufacturing Automation.* ASME Winter Annual Meeting, November 1985, vol. PED-15, pp. 89-96.
- [18] Burke, S.E. and Hubbard, J.E., Jr. Active Vibration Control of a Simply-Supported Beam Using a Spatially-Distributed Actuator. *IEEE Control Systems Magazine*, August 1987, pp. 25-30.
- [19] Caine, M.E. *Chamferless Assembly of Rectangular Parts in Two and Three Dimensions.* Master's thesis, Massachusetts Institute of Technology, June 1985.
- [20] Cannon, R.H. and Rosenthal, D.E. Experiments in Control of Flexible Structures with Noncolocated Sensors and Actuators. *AIAA Journal of Guidance and Control*, September-October 1984, vol. 7, no. 5, pp. 546-553.
- [21] Cannon, R.H. and Schmitz, E. Initial Experiments on the End-Point Control of a Flexible One-Link Robot. *International Journal of Robotics Research*, Fall 1984, vol. 3, no. 3, pp. 62-75.
- [22] Colgate, J.E. and Hogan, N. Robust Control of Manipulator Interactive Behavior. In Shoureshi, R., Youcef-Toumi, K. and Kazerooni, H. (editors), *Modeling and Control of Robotic Manipulators and Manufacturing Processes.* ASME, New York, December 1987.
- [23] Craig, J.J. *Introduction to Robotics: Mechanics and Control.* Addison-Wesley, Reading, Mass., 1986.
- [24] Crawley, E.F. and de Luis, J. *Experimental Verification of Distributed Piezoelectric Actuators for Use in Precision Space Structures.* AIAA Structures, Structural Dynamics, and Materials Conference, San Antonio, Texas, May 1986, pp. 116-124.

- [25] DeFazio, T.L., et al.
Feedback in Robotics for Assembly and Manufacturing.
Technical Report, Charles Stark Draper Laboratory, 1981, no. R-1450.
- [26] Eppinger, S.D. and Seering, W.P.
On Dynamic Models of Robot Force Control.
IEEE International Conference on Robotics and Automation, April 1986, pp. 29-34. (Also in *MIT Artificial Intelligence Laboratory Memo*, no. AIM-910).
- [27] Eppinger, S.D. and Seering, W.P.
Understanding Bandwidth Limitations in Robot Force Control.
IEEE International Conference on Robotics and Automation, April 1987, pp. 904-909. (Also in *MIT Artificial Intelligence Laboratory Memo*, no. AIM-948).
- [28] Eppinger, S.D. and Seering, W.P.
Modeling Robot Flexibility for Endpoint Force Control.
IEEE International Conference on Robotics and Automation, April 1988, pp. 165-170. (Also in *MIT Artificial Intelligence Laboratory Memo*, no. AIM-1046).
- [29] Eppinger, S.D. and Seering, W.P.
An Introduction to Dynamic Models for Robot Force Control,
IEEE Control Systems Magazine, vol. 7, no. 2, pp. 48-52, April 1987.
- [30] Eppinger, S.D.
Development of a Force-Feedback Sensor for Robotic Manipulators.
Bachelor's thesis, Massachusetts Institute of Technology, January 1983.
- [31] Erdmann, M.E.
On Motion Planning with Uncertainty.
Master's thesis, Massachusetts Institute of Technology, August 1984.
- [32] Frost, D.
Controlling a Robot Arm Using an Ultrasonic Range Sensor.
Bachelor's thesis, Massachusetts Institute of Technology, May 1985.
- [33] Garcia-Reynoso, A.
Structural Dynamics Model of a Cartesian Robot.
Doctoral thesis, Massachusetts Institute of Technology, October 1985.
- [34] Gevarter, W.B.
Basic Relations for Control of Flexible Vehicles.
AIAA Journal, April 1970, vol. 8, no. 4, pp. 666-672.
- [35] Goertz, R.C.
Fundamentals of General Purpose Manipulators.
Nucleonics, November 1952, vol. 10, pp. 36-42.
- [36] Good, M.C., Sweet, L.M., and Strobel, K.L.
Dynamic Models for Control System Design of Integrated Robot and Drive Systems.
ASME Journal of Dynamic Systems, Measurement and Control, March 1985, vol. 107, pp. 53-59.
- [37] Gordon, S.J. and Seering, W.P.
Programmable Tools for Flexible Assembly Systems.
ASME International Computers in Engineering Conference, August 1985.

- [38] Gordon, S.J. and Seering, W.P.
Locating Polyhedral Features from Sparse Light-Stripe Data.
IEEE International Conference on Robotics and Automation, April 1987, pp. 801-806.
- [39] Grimson, W.E.L.
From Images to Surfaces.
MIT Press, Cambridge, MA, 1981.
- [40] Gustavson, R.E.
A Theory for the Three-Dimensional Mating of Chamfered Cylindrical Parts.
ASME Journal of Mechanisms, Transmissions, and Automation in Design, vol. 107, 1985.
- [41] Hanselmann, H. and Moritz, W.
High-Bandwidth Control of the Head-Positioning Mechanism in a Winchester Disk Drive.
IEEE Control Systems Magazine, vol. 7, no. 5, pp. 15-19, October 1987.
- [42] Hastings, G.G. and Book, W.J.
Experiments in Optimal Control of a Flexible Arm.
American Control Conference, 1985, pp. 728-729.
- [43] Hogan, N.
Impedance Control of Industrial Robots.
Robotics and Computer-Integrated Manufacturing, vol. 1, no. 1, pp. 97-113, 1984.
- [44] Hollars, M.G. and Cannon, R.H.
Initial Experiments on the End-Point Control of a Two-Link Manipulator with Flexible Tendons.
ASME Winter Annual Meeting, November 1985.
- [45] Horn, B.K.P. and Ikeuchi, K.
Picking Parts out of a Bin.
Technical Report, MIT Artificial Intelligence Laboratory, October 1983, no. AIM-746.
- [46] Inoue, H.
Force Feedback in Precise Assembly Tasks.
Technical Report, MIT Artificial Intelligence Laboratory, 1974, no. AIM-308.
- [47] Kazerooni, H., Bausch, J.J., and Kramer, B.M.
Automated Deburring by Robot Manipulators.
American Control Conference, June 1986, pp. 1749-1755.
- [48] Kazerooni, H., Houpt, P.K., and Sheridan, T.B.
The Fundamental Concepts of Robust Compliant Motion for Robot Manipulators.
IEEE International Conference on Robotics and Automation, April 1986, pp. 418-427.
- [49] Kazerooni, H., Houpt, P.K., and Sheridan, T.B.
Design Method for Robust Compliant Motion for Manipulators.
American Control Conference, June 1986, pp. 1897-1906.
- [50] Kazerooni, H.
Direct-Drive Active Compliant End Effector (Active RCC).
IEEE Journal of Robotics and Automation, June 1988, vol. 4, no. 3, pp. 324-333.
- [51] Khatib, O. and Burdick, J.
Motion and Force Control of Robot Manipulators.
IEEE International Conference on Robotics and Automation, April 1986, pp. 1381-1386.

- [52] Kissel, G.J. and Hegg, D.R.
Stability Enhancement for Control of Flexible Space Structures.
IEEE Control Systems Magazine, June 1986.
- [53] Lord Industrial Automation Division, Lord Corporation.
F/T Series: Force/Torque Wrist Sensing Systems.
Cary, NC
- [54] Lozano-Perez, T., et al.
Handey: A Task-Level Robot System.
4th International Symposium on Robotics Research, Santa Cruz, California, 1987.
- [55] Lozano-Perez, T., Mason, M.T., and Taylor, R.H.
Automatic Synthesis of Fine-Motion Strategies for Robots.
International Journal of Robotics Research, Spring 1984, vol. 3, no. 1.
- [56] Luh, Y.H.S., Fisher, W.D., and Paul, R.P.C.
Joint Torque Control by a Direct Feedback for Industrial Robots.
IEEE Transactions on Automatic Control, February 1983, vol. AC-28, no. 2, pp. 153-161.
- [57] Maples, J.A. and Becker, J.J.
Experiments in Force Control of Robotic Manipulators.
IEEE International Conference on Robotics and Automation, April 1986, pp. 695-702.
- [58] Mason, M.T.
Compliance and Force Control For Computer Controlled Manipulators.
IEEE Transactions on Systems, Man, and Cybernetics, June 1981, vol. SMC-11, no. 6, pp. 418-432.
- [59] McClamroch, N.H. and Wang, D.
Feedback Stabilization and Tracking of Constrained Robots.
IEEE Transactions Automatic Control, May 1988, vol. 33, no. 5, pp. 419-426.
- [60] Meckl, P.H. and Seering, W.P.
Controlling Velocity-Limited Systems to Reduce Residual Vibrations.
IEEE International Conference on Robotics and Automation, April 1988, pp. 1690-1695.
- [61] Nevins, J.L. and Whitney, D.E.
Computer-Controlled Assembly.
Scientific American, February 1978, vol. 238, no. 2.
- [62] Nussbaum, M.
Notes on the AI Servo Drive.
Technical Report, Automatix Inc., Billerica, Mass., Nos. ER-GN-11, ER-RB-09, and EDM-RB-02, 1984.
- [63] Ogata, K.
Modern Control Engineering.
Prentice-Hall, Englewood Cliffs, NJ, 1970.
- [64] Pasch, K.A. and Seering, W.P.
On the Drive Systems for High-Performance Machines.
ASME Journal of Mechanisms, Transmissions, and Automation in Design, 1983.
- [65] Paul, R. and Shimano, B.
Compliance and Control.
Joint Automatic Control Conference, 1976, pp. 694-699.

- [66] Paul, R.P.
Robot Manipulators: Mathematics, Programming, and Control,
MIT Press, 1981.
- [67] Pfeiffer, F. and Gebler, B.
A Multi-Stage Approach to the Dynamics and Control of Elastic Robots.
IEEE International Conference on Robotics and Automation, April 1988, pp. 2-8.
- [68] Filler, G.
A Compact Six Degree-of-Freedom Sensor for Assembly Robot.
In *Proceedings of the 12th International Symposium on Industrial Robots*. 6th International
Conference on Industrial Robot Technology, IFS Publications, Bedford, England, 1982,
pp. 121-129.
- [69] Plank, G. and Hirzinger, G.
Controlling a Robot's Motion Speed by a Force-Torque-Sensor for Deburring Problems.
In Hardt, D. (editor), *Information Control Problems in Manufacturing Technology*, 4th
IFAC/IFIP Symposium, 1982, pp. 97-102.
- [70] Procopio, G.M.
*Active Damping of a Bernoulli-Euler Beam Via End-Point Impedance Control Using
Distributed Parameter Techniques*.
Master's thesis, Massachusetts Institute of Technology, September 1986.
- [71] Raibert, M.H. and Craig, J.J.
Hybrid Position/Force Control of Manipulators.
ASME Journal of Dynamic Systems, Measurement and Control, June 1981, vol. 103, no. 2.
- [72] Rivero, J.A. and Seering, W.P.
An Empirical Study of Compliant Assembly in Three Dimensions.
ASME Design Automation Conference, September 1987, pp. 307-313.
- [73] Roark, R.J. and Young, W.C.
Formulas for Stress and Strain.
McGraw Hill Book Company, New York, 1975.
- [74] Roberts, R.K.
The Compliance of End Effector Force Sensors For Robot Manipulator Control.
Doctoral thesis, Purdue University, 1984.
- [75] Roberts, R.K., Paul, R.P., and Hillberry, B.M.
The Effect of Wrist Force Sensor Stiffness on the Control of Robot Manipulators.
IEEE International Conference on Robotics and Automation, March 1985.
- [76] Rothchild, R.A. and Mann, R.W.
An EMG-Controlled Force Sensing Proportional Rate Elbow Prosthesis.
Proceedings, 1966 Symposium on Biomedical Engineering, Milwaukee, 1966.
- [77] Salisbury, J.K.
Active Stiffness Control of a Manipulator in Cartesian Coordinates.
19th IEEE Conference on Decision and Control, December 1980, vol. 1, pp. 95-100.
- [78] Salisbury, K., Townsend, W., Eberman, B., and DiPietro, D.
Preliminary Design of a Whole-Arm Manipulation System (WAMS).
IEEE International Conference on Robotics and Automation, April 1988, pp. 254-260.

- [79] Salisbury, J.K.
Whole-Arm Manipulation.
4th International Symposium on Robotics Research, Santa Cruz, California, August 1987.
- [80] Schott, J.
Tactile Sensor with Decentralized Signal Conditioning.
9th IMEKO World Congress, Berlin, 1982.
- [81] Seering, W.P. and Eppinger, S.D.
The Dynamics of Robot Force Control.
IFTOMM Seventh World Congress on the Theory of Machines and Mechanisms, September 1987.
- [82] Seering, W.P.
Who Said Robots Should Work Like People?
Technology Review, April 1985.
- [83] Sharon, A., Hogan, N., and Hardt, D.E.
High-Bandwidth Force Regulation and Inertia Reduction Using a Macro/Micro Manipulator System.
IEEE International Conference on Robotics and Automation, March 1988, pp. 126-132.
- [84] Sharon, A. and Hardt, D.E.
Enhancement of a Robot's Accuracy Using Endpoint Feedback and a Macro/Micro Manipulator System.
American Control Conference, June 1984.
- [85] Siegel, D.M., Garabieta, I. and Hollerbach, J.M.
A Capacitive Based Tactile Sensor.
SPIE Conference on Intelligent Robots and Computer Vision, Cambridge, Massachusetts, September 1985.
- [86] Singer, N.C. and Seering, W.P.
Preshaping Command Inputs to Reduce System Vibration.
Technical Report, MIT Artificial Intelligence Laboratory, April 1988, no. AIM-1027.
- [87] Starr, G.P.
Edge-Following with a PUMA 560 Manipulator Using VAL-II.
IEEE International Conference on Robotics and Automation, March 1986, pp. 379-383.
- [88] Stepien, T.M., Sweet, L.M., Good, M.C., and Tomizuka, M.
Control of Tool/Workpiece Contact Force With Application to Robotic Deburring.
IEEE International Conference on Robotics and Automation, March 1985.
- [89] Sweet, L.M. and Good, M.C.
Redefinition of the Robot Motion Control Problem.
IEEE Control Systems Magazine, August 1985, pp. 18-25.
- [90] Tilley, S.W., Cannon, R.H., and Kraft, R.
End Point Force Control of a Very Flexible Manipulator with a Fast End Effector.
ASME Winter Annual Meeting, December 1986.
- [91] Townsend, W.T.
The Effect of Transmission Design on the Performance of Force-Controlled Manipulators.
Doctoral thesis, Massachusetts Institute of Technology, April 1988. (Also in *MIT Artificial Intelligence Laboratory Technical Report*, no. AIM-1054).

- [92] Townsend, W.T. and Salisbury, J.K.
The Effect of Coulomb Friction and Stiction on Force Control.
IEEE International Conference on Robotics and Automation, April 1987.
- [93] Vaaler, E. and Seering, W.P.
Design of a Cartesian Robot.
In Donath, M. and Leu, M. (editors), *Robotics and Manufacturing Automation.* ASME
Winter Annual Meeting, November 1985, vol. PED-15, pp. 163-168.
- [94] Vaughan, D.R.
Application of Distributed Parameter Concepts to Dynamic Analysis and Control of Bending
Vibrations.
ASME Journal of Basic Engineering, June 1968, pp. 157-166.
- [95] Vranish, J.M., Mitchell, E.E., and Demoyer, R.
Magnetoelastic Force Feedback Sensors for Robots and Machine Tools.
In *Proceedings of the 12th International Symposium on Industrial Robots.* 6th International
Conference on Industrial Robot Technology, IFS Publications, Bedford, England, 1982,
pp. 131-142.
- [96] Whitney, D.E. and Nevins, J.L.
What is the Remote Center Compliance (RCC) and What Can It Do?
9th International Symposium on Industrial Robotics, Washington, DC, 1979. (Also in *Charles
Stark Draper Laboratory Report Number P-728*).
- [97] Whitney, D.E.
Force Feedback Control of Manipulator Fine Motions.
ASME Journal of Dynamic Systems, Measurement and Control, June 1977, vol. 99, pp. 91-97.
- [98] Whitney, D.E.
Historical Perspective and State of the Art in Robot Force Control.
IEEE International Conference on Robotics and Automation, March 1985.
- [99] Whitney, D.E.
Quasi-Static Assembly of Compliantly Supported Rigid Parts.
ASME Journal of Dynamic Systems, Measurement and Control, March 1982, vol. 104, pp.
65-77.
- [100] Whitney, D.E.
Touching Experiences.
Videotape presented at IEEE International Conference on Robotics and Automation, March
1985. Available from Charles Stark Draper Laboratory, Cambridge, Massachusetts.
- [101] Wlassich, J.J.
Nonlinear Force Feedback Impedance Control.
Master's thesis, Massachusetts Institute of Technology, February 1986.
- [102] Yabuta, T. and Chona, A.
Stability of Force Control Servomechanism for a Manipulator.
9th IASTED International Symposium on Robotics and Automation, Santa Barbara,
California, May 1987.

- [103] Youcef-Toumi, K. and Asada, H.
The Design of Arm Linkages with Decoupled and Configuration-Invariant Inertia Tensors:
Part I: Open Kinematic Chains with Serial Drive Mechanisms; and Part II: Actuator
Relocation and Mass Redistribution.
In Donath, M. and Leu, M. (editors), *Robotics and Manufacturing Automation*. ASME
Winter Annual Meeting, November 1985, vol. PED-15, pp. 145-152, 153-161.